Superdome

Hewlett-Packard Extends its High-End Computing Capabilities

An IDC White Paper

Analyst: Christopher Willard, Ph.D.

Big Iron: An Idea Whose Time Is Still Here

The history of computer technology is a real-life example of a *Matryoshka* (Russian nesting) doll, in which an initial large doll opens up to reveal a slightly smaller doll, which in turn opens up to reveal a third, smaller doll, and so on. Similarly, the computer industry started with large systems — mainframes — which led to minicomputers, which led to workstations, and so on.

Finding a smaller doll inside a larger one usually does not cause the owner to throw away the larger doll. Surprisingly, the analogy holds for the computer industry — high-end computer systems have maintained markets and advanced technologically, despite the entry of powerful midrange, low-end, and desktop computers. In 2000, computer systems sales in the $1-million-and-above price band accounted for more than $12.3 billion in revenue. If systems costing more than $500,000 are included, the total exceeds $18.4 billion. Certainly, smaller systems have expanded the overall market and captured applications from their larger predecessors, but they have not eliminated requirements for high-end systems.

This white paper first considers the major characteristics of high-end systems and the associated design goals for the latest addition to Hewlett-Packard's HP 9000 product line — the Superdome system. It then provides a top-level review of the Superdome architecture. Finally, it concludes with a brief analysis of the market opportunities and challenges for high-end systems.

Are Large Computers Just Like Small Computers, Only Bigger?

High-end computer systems represent a different class of computer capabilities. These systems can be differentiated from other classes of computers along the dimensions of scalability, capability, manageability, and architecture. This section reviews the dimensions in terms of high-end system capabilities and considers the specific design challenges these dimensions generated for HP system architects in developing Superdome. As a reference point for this discussion, we first provide a "thumbnail sketch" of the Superdome server. The second half of this paper provides a more detailed description of the architecture.

IDC
Analyze the Future

A Quick Sketch of Superdome

Superdome is a new entry in the HP 9000 line of servers. It is built from the "ground up" to provide high-end scalability ... while maintaining the compatibility of applications with other 9000 series systems and providing HP's top level of availability.

Superdome is a new entry in the HP 9000 line of servers. It is built from the "ground up" to provide high-end scalability in terms of processor performance, real memory, and I/O capability while maintaining the compatibility of applications with other 9000 series systems and providing HP's top level of availability. The system runs under HP-UX and provides a shared memory, or symmetric multiprocessing (SMP), management and programming environment. It can scale from two to 64 processors, with from 4GB to 256GB of memory, and with up to 192 PCI slots. This is the most powerful single system to come off HP's factory floor.

Scalability

Current high-end computer architectures generally represent the most capable products in a broad range of compatible scalable systems. Thus, these systems can be defined by not only how large a system can be effectively configured but also by how the architecture extends the range of system configurations and capabilities of the overall product line. Simply being big is not enough; it is also necessary for users to be able to purchase systems at different levels of performance and then grow each system's capability to match individual requirements.

Scalability is often incorrectly defined as simply the number of processors that can be wired together in a single nominal configuration. Although peak processor count defines one bound of a scalable system, one can imagine an architecture scaling boundary that must be defined in terms of how well a total system can be expanded to meet growth in application requirements. For any system, there is a point beyond which the architecture cannot be scaled without some type of breakdown of system effectiveness. This breakdown is often reflected in decreased system efficiencies (e.g., the system is grown by 20% but delivered performance increases by only 10%) or in requirements to rethink application software design in order to make full use of additional capacity.

Scaling Superdome

Scalability requirements led to several design goals for Superdome. First, of course, was that it should significantly extend the range of HP systems at the high end, allowing systems configurations with hundreds of processors with balanced memory and I/O. Second, the system should be highly modular, allowing for as much flexibility as possible in configuring processor count, memory size, and I/O capacity to meet specific user requirements. Third, Superdome was to provide a flat memory model (i.e., memory is seen by the programmer as a single, shared, contiguous space, with all data accessed in the same manner) throughout its scaling range.

Capability

Computer applications can require highly capable systems due to either the absolute amount of work that must be performed in a given time period or the application's future growth requirements.

It is not uncommon for a student leaving home for college to be able to move virtually all of his or her worldly possessions in a single carload. However, by the time we have set up households, the family car no longer works for moving (no matter how many trips one is willing to make). The job requires a moving van and expert assistance. Similarly, computer applications can require highly capable systems due to either the absolute amount of work that must be performed in a given time period or the application's future growth requirements. Examples of this type of problem include:

- **Large database problems.** Database-driven applications can require both the ability to mount a large number of disk drives and/or the ability to maintain high throughput rates to disk via large numbers of I/O channels. Examples include both decision-support applications that require searching a multi-terabyte database and online transaction processing (OLTP) applications for ecommerce.

"Christmas cannot be pushed back." eCommerce companies that are not prepared for a holiday rush simply lose sales.

- **Peak demand support.** Requirements for systems resources may fluctuate dramatically based on seasonal cycles and special events. For example, holiday buying seasons are well defined (e.g., "Christmas cannot be pushed back."). eCommerce companies that are not prepared for a holiday rush simply lose sales; pay-per-view sports events generate most of their revenue within 90 minutes before the event. Computer systems handling the transactions have to be able to handle these short-term loads. Less dramatically, within enterprises, monthly/quarterly/yearly financial reports must be completed within a fixed window for accurate legal reporting and effective corporate planning.

Technical users regularly report that they can use computer systems with from 10 to 1,000 times more power.

- **Number crunching.** Some applications simply require high levels of consolidated computing power. Technical applications range from automotive crash testing, to structural analysis of next-generation aircraft, to seismic analysis for petroleum explorations, to chemical analysis for drug design. These applications can drive current computer systems as close to their peak performance as possible, and sites regularly report that they can use computer systems with from 10 to 1,000 times more power. Commercial users of high-end systems require raw compute power for such applications as financial analysis, econometric modeling, and tracking international exchange rates and their impact on an organization's financial status.

- **High-growth applications.** In cases such as ecommerce-based businesses and Internet service providers (ISPs), where the organization's products, services, or activities are tightly coupled to the computer system, the ability to expand the system itself can be a

In cases such as ecommerce-based businesses and Internet service providers (ISPs), where the organization's products, services, or activities are tightly coupled to the computer system, the ability to expand the system itself can be a gating factor to the overall organization's ability to expand.

gating factor to the overall organization's ability to expand. If the business is there but the computer system isn't, then the opportunity is lost and the customers go elsewhere. Large systems have the ability to expand capacity quickly and smoothly. In addition, they suffer minimal disruptions from such requirements as redesigning or redistributing the applications software or reorganizing and recabling computer-room floors.

Capability Challenges for Superdome

In addition to overall system scalability goals, capability requirements generated several other design challenges for Superdome. Such goals are requirements to support significantly larger I/O configurations, allow for the incorporation of future higher-performance processors from both HP and Intel, and provide a system interconnect that would expand performance in lockstep with the growth of the overall system.

System and Resource Management

Computer-room resource allocation options can be broken into two broad options:

- **Distributed computing.** A large number of smaller systems are each assigned more or less specific tasks.
- **Consolidated systems.** A small number of large systems are used, with each system generally performing multiple tasks.

While it is beyond the scope of this paper to review the relative merits of each approach, it is worth noting some of the reasons for considering high-end systems:

- **Resource utilization.** Systems that provide large sets of tightly coupled resources under a unified control system allow users to address a broader range of problems, bringing the appropriate resources to bear on each application. In addition, these systems allow sites to maximize the use of the resources by mixing and matching applications to make the fullest possible use of the computer.
- **Reduced duplication of data.** To the extent that applications can be broken up into small independent subtasks to be run on separate servers, there is usually some requirement for data replications, both at the system software and application software levels. This replication is not required for large systems.
- **Management cost control.** Combining the tasks run on several independent smaller systems onto a single large server (i.e., server or network consolidation) provides economies of scale in such areas as lower system staff requirements to manage one system versus several, as well as associated training and scheduling costs, reduced maintenance contract costs, and elimination of duplicate applications and system software licenses.

Management Challenges for Superdome

At first glance, the system and resource management advantages of high-end systems would appear to simply be side effects of system size and scalability. However, architectural design decisions can affect the degree to which these advantages are actually realized. Superdome's design goals in this area first center around the flat memory architecture, which allows the system to be managed by a single instance of the operating system and allows for all processors to have equal access to all system resources. Second, the company wanted to provide a highly partitionable system, which allows the computer to be divided into several independent subsystems that can each be assigned to separate tasks. Partitioning maintains a single point of control while allowing system managers to dedicate resources to applications and to easily reconfigure the systems to adjust partitions to specific changing job load requirements.

Architectural Features

With virtually any class of machine, the larger, more powerful models are designed to be structurally stronger and include additional features in order to make the additional power manageable.

With virtually any class of machine, the larger, more powerful models are designed to be structurally stronger and include additional features in order to make the additional power manageable. This is also the case for high-end computer systems, which differ from other classes of computers, based in part on architectural features. These features include:

- **Internal system interconnects.** "*Supercomputer processors are provided free of change to anyone willing to pay for the memory interconnect necessary to keep the CPUs fed.*" — *Ancient Cybernetic Proverb.* Supporting large numbers of processors, large amounts of memory, and large numbers of I/O devices requires very fast internal communications systems to move data around the system at roughly the rate that various components can create and consume it. The performance and flexibility of the interconnect distinguish high-end systems from networks of independent servers or clusters. The interconnect also determines such factors as how well the overall system scales and what types of programming models can be supported.
- **Inherent reliability.** In a case of turning necessity into a virtue, large systems are designed for high availability for several reasons: First, to account for the greater complexity of the architectures — more complexity equates to more components that can fail, which in turn leads to design features to eliminate single points of failure and to provide early warning of potential problems. Second, high availability is required to support large numbers of users and/or major organizational functions, where the failure of these systems can create major disruptions to a business. Finally, additional memory chips, disk drives, power supplies, and fans make up a smaller fraction of the total cost of a computer as the size of the system increases. Thus, the relative cost of high-availability features can decrease with system size.

Architectural Challenges for Superdome

Architectural features are almost synonymous with design requirements. The Superdome system interconnect had to be designed to support all requirements for scalability, flexibility, partitioning, performance, and so on. HP has a long history of emphasizing system availability and reliability in its computer designs. From the outset, Superdome designers were working to make this system the most reliable in HP's history.

User in the Loop Design

A final challenge for any system design is to make sure that the computer smoothly fits into real-world environments.

The systems architects' and the end users' visions for next-generation systems usually line up fairly closely, but never exactly. Thus, a final challenge for any system design is to make sure that the computer smoothly fits into real-world environments. In order to bridge this theory versus practice gap, HP instituted a partnership program with key customers to develop design requirements for Superdome. The overall goal was to provide the various Superdome development teams with a better understanding of the requirements, problems, and frustrations of end users in their day-to-day operations of computer facilities. The program specifically addressed requirements in four areas: hardware features and performance, system management solutions, support requirements, and total systems solutions packaging requirements. The program was based on visits to 25 computing facilities worldwide.

In addition to validating many Superdome design decisions, the customer feedback also provided insights that directly affected the system design. Design changes that resulted from this program include:

- Increased memory capacity
- Additional redundancy feature for high availability, such as ability to provide dual independent power sources to the system
- Instant capacity-on-demand features
- Form factor changes including two-sided access and the ability to separate fans from the main cabinet in order to easily move systems through doors lower than two meters
- Change air flow to exit out of the top of the system cabinet
- Addition of new commands in HP-UX to support customers' work practices
- Creation of the Partition Manager system management tool that allows user to configure and manage Superdome partitions. "This tool was the direct outcome of sitting next to mainframe system administrators as they were conducting their tasks."

The success of this program has led HP to implement it on an ongoing basis for Superdome enhancements and future products.

HP's Superdome: A New High-End System

This section provides a more detailed description of the Superdome architecture and considers some of the operational advantages to specific features of the system.

Meeting the Architectural Challenges

Scalable architectures are generally described in terms of the following features:

- **Building blocks.** The basic processing, memory, and I/O components from which the system is built. These components ultimately determine the amount of power that can be applied to specific problems and the range of problems that can be effectively addressed.
- **System interconnect.** The internal network that connects the building blocks. The interconnect essentially combines the various components into a single system versus a cluster of quasi-independent nodes. It determines how large a system can be effectively configured and how much control users have in configuring systems to match their site's requirements.
- **Reliability features.** The range of architectural and design decisions that determine not only the likelihood of a system failure but also the cost of that failure in terms of the length of time the system is down or running in degraded mode.
- **Operability features.** The hardware and software attributes of the system that determine how the system is managed can be logically configured and programmed.

The following section briefly reviews the architectural and design features of the Superdome server and provides our analysis on how these features can be expected to address HP's challenges/goals for the products.

Cells: Processing Building Blocks

Superdome's basic building block is the "Cell" or "Cell board." Cells can be viewed as symmetric multiprocessors (SMPs) with processors, memory, cell management functions, and an internal interconnect (i.e., the Cell controller), as shown in Figure 1:

- **Processors.** A Cell can contain from one to four processors. Currently, Superdomes use HP PA-8600 processors running at 552MHz.
- **Memory.** The CPUs share from 2GB to 16GB of memory, increasing in 2GB increments. The memory is divided into two memory banks, or the equivalent of eight Cell memory controller (DIMM) slots.

Figure 1
Cell Board Architecture

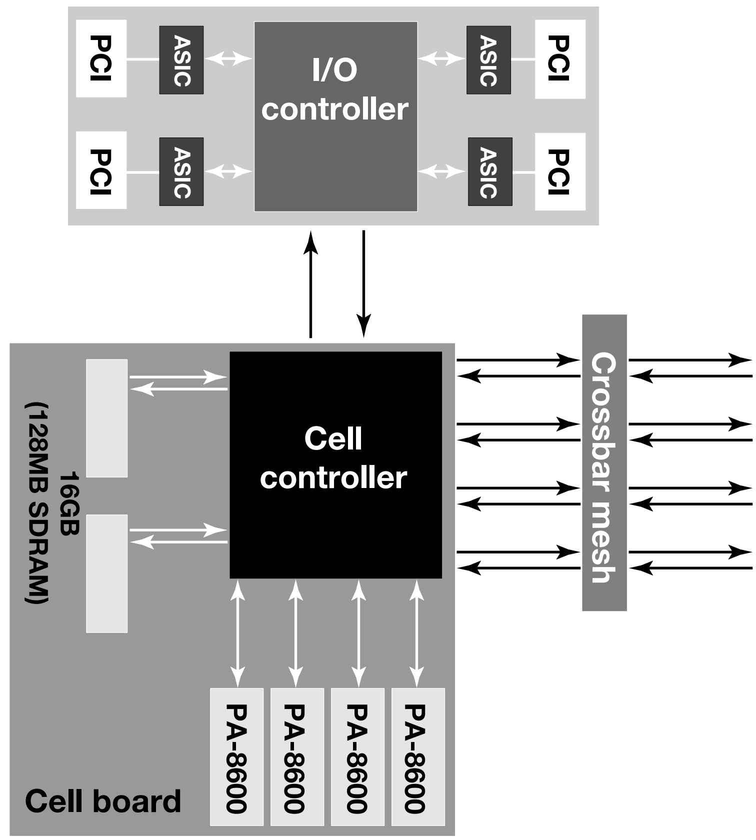

Source: Hewlett-Packard, 2000

- **Cell management.** Cells contain processor-dependent hardware (PDH), which contains the system boot firmware and resets the cell and brings it up to the point where it can be incorporated into a running system.
- **Cell controller.** The communications hub of the Cell, the controller chip manages both on-Cell and off-Cell data movement. On-Cell data operations occur between the processors and memory and between the PDH and memory. Off-Cell operations include accessing memory on other Cells and communicating with I/O subsystems. In addition, Cell controllers maintain cache coherency throughout the system.

The Cell controller provides interfaces to:

1. **Memory.** There are two ports to the Cell memory controllers (DIMMs). Each port operates at 2.0GBps for peak bandwidth of 4.0GBps.

2. **Processors.** There is a dedicated port to each processor. Each port operates at 2.0GBps; thus, peak memory bandwidth on a Cell with four processors is 8.0GBps.

3. **Cell to crossbar.** There is one port to the overall system interconnect for communication to other Cells. This port operates at a peak speed of 8.0GBps.

4. **I/O subsystem.** There is one port to an I/O bus that connects the Cell to an I/O subsystem. Peak speed on this port is 2.0GBps.

5. **Processor-dependent hardware.** The PDH communicates with processors and memory via the Cell controller.

The Cell-based architectural approach for Superdome allows HP to address several of its design challenges:

A modular approach provides users with a broad range of options for specifying a system, allowing for the configuration of the system to closely match operational requirements.

- **Configuration flexibility.** Cells can be viewed as "computational building blocks" or modules. A modular approach provides users with a broad range of options for specifying a system, allowing for the configuration of the system to closely match operational requirements.

- **Scalability.** Cells allow Superdome systems to be scaled both by increasing the processor count and memory size within cells and increasing the total number of Cells in a system.

- **Investment protection.** The ability to add Cells to a system over time allows sites to expand the capability of a Superdome system without having to replace the system and sacrifice investments in existing hardware and peripherals.

I/O Module: Another Type of Building Block

A second building block in the Superdome system is the I/O module. An I/O module contains eight 2X PCI slots, four 4X PCI slots, an I/O controller chip, and a power supply. Each PCI slot has its own PCI bus. The 2X slots are 64 bits wide and run at 33MHz, thus providing about 265MBps of peak bandwidth. The 4X slots are also 64 bits wide but run at 66MHz, providing 530MBps of peak bandwidth.

I/O subsystem capacity and performance scale with the number of Cells.

One I/O module can be connected to a Cell over a 2.0GBps link. Thus, I/O subsystem capacity and performance scale with the number of Cells. Table 1 provides some examples of I/O configurations. Real estate considerations in the system cabinet may require the use of an I/O Expansion Cabinet. Cells can be configured without I/O modules attached, but I/O modules cannot be configured in the system unless attached to a cell.

The I/O subsystem extends the Cell-based modular architecture of the Superdome. Thus, I/O modules extend the configuration flexibility, scalability, availability, and investment protection advantages associated with the Cells. However, two aspects of the modular I/O approach are of particular interest:

Table 1
Sample Superdome I/O Subsystem Configurations

No. of Cells	Max. CPUs	Max. No. of I/O Modules	No. of PCI Slots	Total Bandwidth (GBps)	I/O Expansion Cabinet Required
4	16	4	48	8.0	No
8	32	8	96	16.0	Yes
16	64	16	192	32.0	Yes

Source: IDC, 2001

- **I/O capacity.** A 16-Cell (up to 64 processors) Superdome system can be configured with 192 PCI slots, with a total I/O subsystem bandwidth of 32.0GBps. This provides a significant step up in capacity for the HP 9000 product line. For comparison, a 32-processor V-class system can be configured with 28 PCI slots running at a peak bandwidth of 1.9GBps.
- **Independent I/O configuration.** With the exception that I/O modules are attached to Cells, I/O subsystem configuration is independent of any other system component. Cells can be configured either with or without I/O modules. Thus, unused I/O capacity is not automatically configured into a system in order to get required processing or memory capabilities. In addition, processor count and memory size within a Cell has no effect on how the associated I/O module is configured. Thus, at the extreme, a 16-Cell system with only 16 processors could support up to 192 PCI slots.

From Modules to Systems: Systemwide Memory Interconnect

In a nutshell, the interconnect combines the set of independent processing and I/O modules into a unified coherent system.

More than any other architectural component memory, system interconnects define the overall operational characteristics of the system. In a nutshell, the interconnect combines the set of independent processing and I/O modules into a unified coherent system. Specific attributes conferred by the interconnect include:

- **Scalability.** The memory interconnect defines the physical number of modules or nodes that can be practically configured in a system.
- **Performance.** Ultimately, computers are only as fast as their slowest components. The speed at which a system can move data between memory, processors, and I/O devices is often the determining factor in applications performance rather than processor speed.
- **Programming model.** The memory interconnect determines whether programs can be written to a flat memory model, or if a message passing approach is needed (i.e., some memory is local and accessed directly, while the majority of memory is controlled by other nodes and requires that messages be sent between nodes to request and send data). In addition, as highly scalable flat memory

systems become large, the physical distance between a memory bank that holds an item of data and the processor that needs to read or write that item can affect applications performance. The performance of the interconnect determines the point at which programmers may need to begin to address this "data locality" problem.

A major goal in designing Superdome was to create a highly scalable system that maintained the same flat memory programming model that is used on all other systems in the HP 9000 line.

A major goal in designing Superdome was to create a highly scalable system that maintained the same flat memory programming model that is used on all other systems in the HP 9000 line. The HP approach to meet the goal is a hierarchy of memory crossbars (see Figure 2 and Table 2):

- **Cell level.** The Cell controller chip connects memory and processors within a cell and also connects the cell to the larger system interconnect. On-Cell aggregate memory bandwidth is 4.0GBps.
- **Cell to Cell.** The second level of the hierarchy is an 8-way nonblocking crossbar that connects up to four cells (up to 16 processors). This crossbar is used to implement all levels of the memory hierarchy outside of the Cell. It supports communications for both PA-RISC and IA-64-based systems. It provides 16.0GBps of bandwidth.
- **Backplane.** The third level connects two of the Cell-to-backplane level crossbars. There are two levels here: The first level is with systems with up to eight Cells and up to 32 processors, where the system provides 32.0GBps of bandwidth. The second level supports up to 16 Cells and 64 processors at 64.0GBps of bandwidth.

Superdome's memory interconnect strategy has a number of advantages:

By adding memory interconnect structure as Cells are added to a configuration, the Superdome interconnect is able to maintain a constant bandwidth per cell ratios.

- **Bandwidth growth in proportion to Cells.** By adding memory interconnect structure as Cells are added to a configuration, the Superdome interconnect is able to maintain a constant bandwidth per cell ratios (4.0GBps per Cell, assuming maximum cell configurations). Thus, adding processors and I/O capability should not throw the internal memory interconnect out of balance.
- **Latency characteristics.** In switch-based interconnect schemes, memory-access latency is a function of the switch latency (i.e., the time it takes the switch to set up a circuit) multiplied by the number of switches necessary to complete a path between the Cells exchanging data. Latency times thus depend on the number of switches, or "distance," between a processor requesting the data item and the memory bank containing the data. HP reports memory latency times in terms of an "average memory load to use" analysis for different-sized systems, assuming a uniform distribution of data across the systems. (Thus, this analysis combines both short- and long-distance memory requests with the number of requests at each distance weighted by the size of the system). Latencies reported include:

Figure 2
Interconnect Fabric: Crossbar Mesh

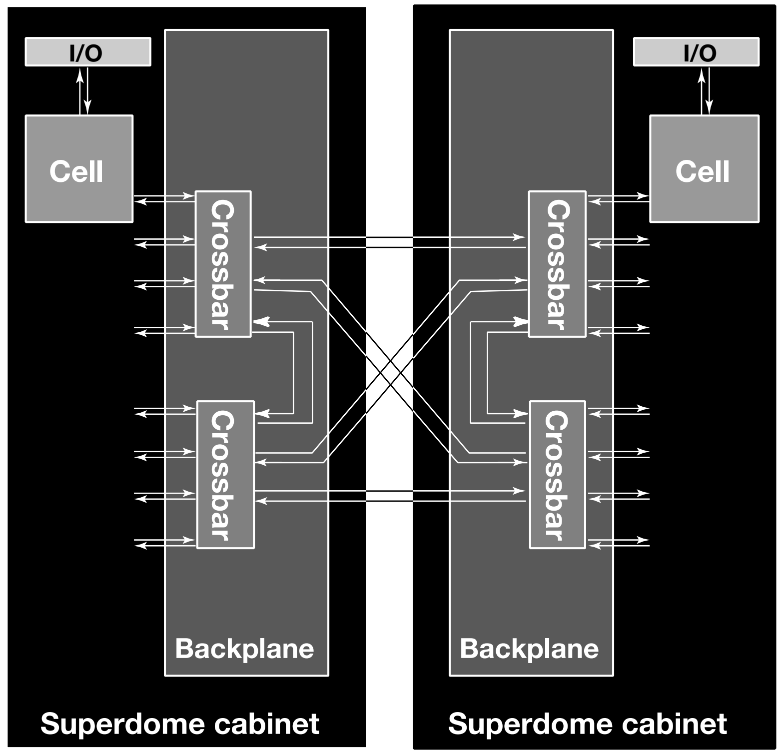

Source: Hewlett-Packard, 2000

Table 2
Superdome Memory System Hierarchy

Level	No. of Cells	Max. CPUs	Max. Memory (GB)	Aggregate Bandwidth (GBps)	Bandwidth per Cell (GBps)
Cell	1	4	16	8.0	–
Cell to cell	4	16	64	16.0	4.0
2 crossbars	8	32	128	32.0	4.0
4 crossbars	16	64	256	64.0	4.0

Source: IDC, 2001

- Average memory latency within the cell is 235ns.
- Average memory latency between cells on the same crossbar is 266ns for an eight-processor system.
- Average memory latency between cells on different crossbars ranges from 296ns for a 16-processor system to 360ns for a 64-processor system.

- **Support for a flat memory model.** The bandwidth and latency characteristics allowed HP to implement a cache coherent flat memory model on the system.

As the latest addition to and most powerful model in the HP 9000 systems series, the Superdome design represents the highest level of availability features of any HP product.

High-End Availability for a High-End System

HP sees high-availability features as crucial components of all its computer systems. As the latest addition to and most powerful model in the HP 9000 systems series, the Superdome design represents the highest level of availability features of any HP product. It is important to emphasize that availability becomes more crucial as systems take on a large proportion of an organization's computing workload and they become responsible for large, monolithic, critical applications (e.g., corporate databases). Although a complete listing of high-availability features incorporated into Superdome is beyond the scope of this paper, it is worthwhile to briefly review the top-level features:

- **N+1 redundant components.** The reduction of single points of failure is central to high availability. Superdome features include N+1 hot-swappable components and cabinet blowers (i.e., cooling fans), I/O module fans, and DC power supplies. The N+1 concept is for the system to support one more of each type of unit than is needed for operations. All units operate together, but in case of unit failure, the remaining units provide full service, keeping the system running while the failed unit is replaced. A side benefit to this level of redundancy is that one unit can be taken offline for preventative maintenance without bringing down the system.
- **Online removable and replaceable components.** These allow faulty components to be physically and logically isolated from the rest of the system and removed for maintenance or replaced without bringing down the remainder of the system. This capability requires specialized OS support to track lost memory locations and offline I/O devices, and thus it may not be available with all versions of the OS. Components can be removed and replaced include Cell boards and I/O cards. An additional advantage to this feature is that the new components can also be added without bringing down the system.
- **Input power protection.** Superdome can optionally be connected to two independent power supplies. Thus, the loss of one supply does not bring down the system.
- **Memory and data path error detection and correction.** As computer components become smaller, the risk increases that data will be modified from such uncontrollable factors as cosmic rays, or ambient Alpha particles. A flipped bit of data in the OS or an application can bring the machine down. Superdome provides error detection and correction codes on all CPU and memory data paths. In addition, I/O data transmissions are protected by data parity checks.
- **Dynamic memory resiliency.** The combination of memory chip redundancy and error correction codes on Superdome systems is

... allow the complete loss of a memory chip without causing the system to fail or result in loss of data.

sophisticated enough to allow the complete loss of a memory chip without causing the system to fail or result in loss of data.

- **Monitoring and diagnostic features.** Two major technical requirements for high availability are to provide both internal telemetry to allow the system to identify and report on actual and potential failures, and to provide the software environment to allow system management and services personnel to evaluate problems and take corrective action. Monitoring and diagnostic features include:
  - **Service management station.** This independent system connects to Superdome systems at a site and is used to run system scans. The station monitors the physical condition of the system (e.g., scans application-specific integrated circuit [ASIC] chips) and signals the operator in case of significant events. It also provides for remote access to the system.
  - **Event monitoring system.** This applications package consists of a set of system and network monitors and a management mechanism to control and monitor system activities and report on results. Examples of monitoring activities include notification of EEC errors on the main system bus and in the CPU cache, cabinet fan and power failures, and system temperature problems.
  - **High availability observatory.** This solutions package combines Superdome internal systems monitoring features with HP support engineers and HP's Mission-Critical Support center. The objective of the observatory is to allow HP to both analyze system health on an ongoing basis and respond quickly to problems that develop with the highest-level support staff necessary.

Partitioning: From One Big System of Several Small Systems and Back

Somewhat ironically, a key feature of today's high-end systems is the ability to logically divide them into several smaller systems.

Somewhat ironically, a key feature of today's high-end systems is the ability to logically divide them into several smaller systems. Partitions are used for several reasons: to assure that time-critical tasks are guaranteed systems resources and protected from interference from other applications, to allow applications to run on different versions of an OS as needed, to allow for applications development and testing without jeopardizing the operations of the overall systems, to provide high availability through multiple OS instances and failover capabilities between partitions, and simply to assure different departments that they have control over their due amount of computer resources.

The advantage of a partitioned system over multiple independent systems is that the partitions can be recombined or further subdivided on an ongoing basis.

The advantage of a partitioned system over multiple independent systems is that the partitions can be recombined or further subdivided on an ongoing basis as needed to meet regularly scheduled demand loads, adjust to changes in demand over time, and provide for one-time or emergency situations.

Superdome provides two new partitioning approaches for HP systems:

- **nPartitions.** Superdome systems include programmable hardware within the cells to define partition boundaries. This feature allows for the creation of physically isolated partitions, built up in units of complete Cells. Each partition has its own independent CPUs, memory, and I/O resources. Cells may be dynamically added or removed from a partition through the use of system management commands and without physically manipulating the hardware or affecting other partitions. nPartitions allow for multiple independent OS images to run simultaneously on different partitions (e.g., multiple versions of HP-UX, Linux, and Windows).
- **Virtual partitions.** This approach features a systems software-based partitioning capability that allows the creation of logical partitions within Superdome systems or an nPartition. Virtual partitions are built up in units of single CPUs. Processors can be dynamically migrated between virtual partitions. Virtual partitioning supports multiple independent versions of HP-UX, and individual partitions can be reconfigured and rebooted without affecting the operations of other partitions. Virtual partitions are also available on L- and N-class systems.

Staying at the High End: Processor Strategy

The continual advances in the performance of underlying component technology present a moving target for computer systems designers. To maintain both their position in the market and users' investments, system architectures must be designed to incorporate more capable components over time. The Superdome architectural strategy includes provisions for component upgrades throughout the life of the product.

HP will both provide higher-performing versions of its PA-RISC processors and incorporate Intel IA-64 processors into Superdome.

The most dramatic example of HP's plans for keeping Superdome at the top of the technology curve is the company's dual processor strategy. HP will both provide higher-performing versions of its PA-RISC processors and incorporate Intel IA-64 processors into Superdome:

- **PA-RISC.** Superdome currently ships with PA-8600 processors, with upgrades to the PA-8700, PA-8800, and PA-8900 processors planned over the next three to four years. It will be possible to mix PA-8600 and PA-8700 processors in the same system; however, the processors should be configured in separate partitions. HP will continue to invest in PA-RISC after the introduction of IA-64 to provide users as much flexibility as possible both for maintaining current applications and for migrating to IA-64.
- **Intel IA-64.** HP will begin to incorporate IA-64 processors into Superdome in early 2003. Both the IA-64 and some of the later PA-RISC processors will require new Cell boards; however, all other system components will support all processor types. HP has worked to develop a strong PA-RISC/IA-64 compatibility strategy.

Thus, current applications developed on HP-RISC systems under HP-UX, Windows 2000, or Linux will run unchanged on IA-64. In addition, application performance can be optimized by recompiling the original source code.

Conclusion: Market Challenges and Opportunities for High-End Computing

As the computer industry expanded and new classes of computer systems became available on the market, two major market dynamics came into play. First, the creation of new opportunities for the applications for computer technologies arose from the availability of new systems at different price points and with different capabilities. This in turn allowed new groups of users to take advantage of the technology and to apply it to new markets. Thus, we have the progress of computer architectures from the glass house to the desktop. Second, as all classes of systems became more powerful over time, the range of applications that each class could effectively address began to overlap. Thus, we have the development of distributed applications models.

These dynamics create both challenges and opportunities for high-end systems. The expansion of the computer market created by the entry of ever-smaller systems has the effect of pushing high-end systems into the background. This loss of visibility leads to "old technology" or "legacy system" labels for high-end systems. However, growth in system classes has also increased the total demand for data processing and computing and is in the processes of creating requirements for computing infrastructures that cross community, company, and, national boundaries. This expansion leads in part to application requirements that high-end systems are designed to fulfill.

High-end systems "have to" compete with other classes of computers, but they also "get to" and are "able to" compete with the smaller computer architectures.

Competition between classes of computers has historically been viewed in terms of a demand drain from larger systems to smaller systems as increased capabilities allow the smaller systems to "capture" applications from the larger ones. However, over the past several years, the network consolidation counter-trend has developed momentum, where high-end systems have provided the scalability and price/performance attributes to more easily manage large, diverse workloads. The point here is that high-end systems not only "have to" compete with other classes of computers but also "get to" and are "able to" compete with the smaller computer architectures.

Superdome is an excellent representative of the high-end class of computer architectures: It is a highly scalable system that is designed to provide a broad range of configurations and address the needs of highly diverse computing environments. At the same time, it is designed to maintain system balance across a broad range of configurations, be inherently reliable, and support both large, monolithic applications and sets of diverse, smaller workloads.

NORTH AMERICA

Corporate Headquarters
5 Speen Street
Framingham, MA 01701
508-872-8200

IDC Canada
36 Toronto Street, Suite 950
Toronto, Ontario
Canada M5C2C5
416-369-0033

IDC Irvine
18831 Von Karman Ave, Ste 200
Irvine, CA 92612
949-250-1960

IDC Mountain View
2131 Landings Drive
Mountain View, CA 94043
650-691-0500

IDC New Jersey
120 Wood Ave South, Suite 509
Iselin, NJ 08830
732-632-9222

IDC New York
2 Park Avenue
Suite 1505
New York, NY 10016
212-726-0900

IDC Texas
100 Congress Ave, Suite 2000
Austin, TX 78701
512-469-6333

IDC Washington
8304 Professional Hill Drive
Fairfax, VA 22031
703-280-5161

ASIA/PACIFIC

IDC Asia/Pacific (Hong Kong)
12/Floor, St. John's Building, 33 Garden Road
Central, Hong Kong
852-2530-3831

IDC Asia/Pacific (Singapore)
71 Bencoolen Street, #02-01
Singapore 189643
65-226-0330

IDC Australia
Level 4, 76 Berry Street
North Sydney
NSW 2060, Australia
61-2-9922-5300

IDC China
Room 611, Beijing Times Square,
88 West Chang'an Avenue, Beijing,
P.R. China, 100031
86-10-8391-3456

IDC (India) Limited
Cyber House
B-35, Sector 32 - Institutional
Gurgaon - 122002
Haryana, India
91-124-6381673 to 80

IDC Japan
10F The Itoyama Tower
3-7-18, Mita Minato-ku
Tokyo 108-0073, Japan
81-3-5440-3400

IDC Korea Ltd
Suite 704, Korea Trade Center
159-1, Samsung-Dong, Kangnam-Ku
Seoul, Korea 135-729
82-2-55-14380

IDC Malaysia
Suite 13-03, Level 13, Wisma KiaPeng
No. 3, Jalan Kia Peng
50450 Kuala Lumpur, Malaysia
6-03-2163 3715

IDC New Zealand
Level 7, 246 Queen Street
Auckland, New Zealand
64-9-309-8252

IDC Philippines
7F, SEDCCO 1Bldg
Rada Street Corner
Legaspi Street
Legaspi Village
Makati City, Philippines
632-894-4808

IDC Taiwan Ltd.
10F, 31
Jen-Ai Rd, Sec 4,
Taipei 106, Taiwan, R.O.C.
886-2-2731-7288

IDC Thailand
27 Soi Charoen Nakorn 14
Charoen Nakorn Road, Klongtonsai
Klongsan Bangkok 10600, Thailand
66-2-439-4591-2

IDC Vietnam
37 Ton Duc Thang Street
Unit 1606
District-1 Hochiminh City Vietnam
84-8-910-1235

LATIN AMERICA

IDC Miami
Latin America Headquarters
8200 NW 41 Street
Suite 300
Miami, FL 33126
305-267-2616

IDC Argentina
Trends Consulting
Rivadavia 413, 4th Floor, Suite 6
C1002AAC, Buenos Aires, Argentina
54-11-4343-8899

IDC Brasil
Alameda Ribeirão Preto, 130 cj 41
01331-000 São Paulo
SP Brazil
55-11-253-7869

International Data Corp. Chile
Luis Thayer Ojeda 166 Piso 12
Providencia, Santiago 9, Chile
56-2-231-0111

IDC Colombia
Carrera 40 # 103-78
Bogota, Colombia
571-533-2326

IDC Mexico
Select - IDC
Av. Nuevo Leon No. 54 Desp. 501
Col. Hipodromo, Condesa
C.P. 06100 Mexico, D.F.
52-5-256-1426

IDC Venezuela
Calle Guaicapuro
Edif. Torre Seguros Alianza
Piso 6, Ofc. 6-D, El Rosal
Caracas 1060, Venezuela
58-2-951-3270

EUROPE, MIDDLE EAST, AND AFRICA

IDC Austria
c/o Loisel, Spiel, Zach Consulting
Mayerhofgasse 6
A-1040 Vienna, Austria
43-1-50-50-900

IDC Benelux (Belgium)
29 Avenue Louis Gribaumont
B-1150 Brussels, Belgium
32-2-779-46-04

IDC Benelux (The Netherlands)
A. Fokkerweg 1
1059 CM Amsterdam
The Netherlands
31-20-669-2721

IDC Central Europe (ECE)
Male Namesti 13
Praha 1 110 00, Czech Republic
420-2-2142-3140

IDC Central Europe (Germany)
Nibelungenplatz 3, 11th Floor
60318 Frankfurt, Germany
49-69-90502-0

IDC Central Europe (Switzerland)
Niederlassung Züerich
WTC, Leutschenbachstrasse 95
CH - 8050 Züerich
Switzerland
41-1-307-1000

IDC Egypt
39 Iraq Street
Mohandesseen, Cairo, Egypt
20-2-336-7355

IDC France
Immeuble La Fayette
2, Place des Vosges, Cedex 65
92051 Paris la Defense 5, France
33-14-904-8000

IDC Hungary
Nador utca 23, 5th Floor
H-1051 Budapest, Hungary
36-1-473-2370

IDC Israel
4 Gershon Street
Tel Aviv 67017, Israel
972-3-5611660

IDC Italy
Viale Monza, 14
20127 Milano, Italy
390-2-284-571

IDC Nigeria
House 2, 'C' Close, 403 Road, 4th Avenue
New Extension, Festac Town
Lagos, Nigeria
234-1-883585

IDC Nordic (Denmark)
Jagtvej 169B
DK-2100 Copenhagen, Denmark
45-39-162222

IDC Nordic (Finland)
Jarrumiehenkatu 2
FIN-00520
Helsinki, Finland
358-9-8770-466

IDC Nordic (Sweden)
Box 1096 Kistagången 21
S-164 25 Kista, Sweden
46-8-751-0415

IDC Poland/ProMarket
Wrobla 43
02-736 Warsaw, Poland
48-22-754-0518

IDC Portugal
Av. Antonio Serpa, 36 Piso 9
1050-027 Lisbon
Portugal
351-21-796-5487

IDC Russia
c/o PX Post, RDS 186
Ulitsa Zorge 10
Moscow 125525
Russian Federation
7-501-929-9959

IDC South Africa
c/o BMI-TechKnowledge
3rd Floor, 356 Rivonia Blvd.
PO Box 4603, Rivonia, 2128
South Africa
27-11-803-6412

IDC Spain
Ochandiano, 6
Centro Empresarial El Plantio
28023 Madrid
34-91-7080007

IDC Turkey
Tevfik Erdonmez Sok. 2/1 Gul Apt.
Kat 9D; 46 Esentepe
Istanbul, Turkey
90-212-275-0995

IDC U.K.
British Standards House
389 Chiswick High Road
London W4 4AE
United Kingdom
44-20-8987-7100

IDC is the foremost global market intelligence and advisory firm helping clients gain insight into technology and ebusiness trends to develop sound business strategies. Using a combination of rigorous primary research, in-depth analysis, and client interaction, IDC forecasts worldwide markets and trends to deliver dependable service and client advice. More than 700 analysts in 43 countries provide global research with local content. IDC's customers comprise the world's leading IT suppliers, IT organizations, ebusiness companies and the financial community. Additional information can be found at www.idc.com.

IDC is a division of IDG, the world's leading IT media, research and exposition company.

Sponsored by Hewlett-Packard
00-292SYSTEM2737R
May 2001

IDC
Analyze the Future

IDC
5 Speen Street • Framingham, MA 01701
(508) 872-8200 • Fax (508) 935-4015 •
www.idc.com