research

HEWLETT® PACKARD

HP V2600 HP 9000 Servers

HP 9000 V2600 Enterprise Server

Architectural Overview White Paper

This material is intended for HP internal, HP Channel Partners, and HP customers.

Table of Contents

HP 9000 Enterprise Server Product Line

The HP 9000 Enterprise Server product line running HP-UX addresses the major computing challenges that customers face today in on-line transaction processing (OLTP), enterprise resource planning (ERP), decision support, and scientific and technical applications.

The HP 9000 Enterprise Server product line includes the A-, R-, D-, K-, L-, N-, and V-Class platforms, as shown in Figure 1. Through the use of HP-UX and common components such as HP's PA-RISC processor, the HP 9000 platforms provide binary compatibility and application portability. The systems fit neatly into an existing environment and provide compute and file services for mission-critical OLTP, ERP, data warehousing, and server consolidation applications.

The HP 9000 Enterprise Server product line represents the most scalable range of computing available today. These systems feature:

- Complete up-and-down-the-line compatibility, permitting users to select the appropriate platform at the appropriate price point without concern for application availability.
- 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 consistent price/performance over the entire range of products and protecting customers' investments in hardware and software.
- Consistent programming model, presenting the same application programming environment regardless of the systems' performance levels, greatly increasing the number of available "off-the-shelf" third-party applications, and reducing porting and development costs.

This document describes the architecture, system software, and tools of the high-end HP 9000 V-Class Enterprise Server family.

Introducing the HP 9000 V2600 Enterprise Server

HP 9000 V-Class Enterprise Servers are ideally suited for commercial and technical data center applications. The latest V-Class model—the V2600—is HP's highest performing enterprise server specifically designed for the enterprise data center. Its predecessors—the V2200, introduced in November 1997; the V2250, introduced in March 1998; and the V2500, introduced in February 1999—have been installed worldwide in major Global 1000 companies, providing solutions to corporations' mission-critical commercial and technical challenges.

The V2600 greatly extends the performance and scalability of the V-Class, including:

- **High performance RISC processor (PA-8600)**
  - 552-MHz, 4-way superscalar (2.2 GFLOPS)
  - 56-way instruction reorder buffer
  - 10 functional units
  - Large on-chip cache (1.5 MB)
- **Increased Memory Subsystem Performance**
  - Double memory controller bandwidth
  - Increases bandwidth percent of peak [15.36 GB/s per cabinet, 61.44 GB/s aggregate bandwidth with HP's Scalable Computing Architecture (SCA)]
  - 8-fold increase in memory interleaving (from 32- to 256-way)
  - Higher density SDRAMS provides 32 GB per cabinet and up to 128-GB physical memory total with SCA
- **Doubled processor capacity per cabinet**
  - Configurations available from 2 to 128 CPUs
- **Doubled I/O interface performance**
  - 64-bit 33MHz (2x) industry-standard PCI
  - Eight 240-MB/s I/O channels per cabinet
  - Increased number of PCI slots (up to 112)

The V2600 runs the industry-leading HP-UX 11 operating environment. This mainframe-class, 64-bit operating system enjoys the industry's greatest support from independent software vendors—allowing you to choose from more than 15,000 applications, including native 64-bit versions of all major databases and leading ERP applications. With the 64-bit HP-UX operating system, applications that execute on the V-Class are binary compatible with future IA-64–based platforms. HP-UX 11 also turns the V-Class in an ideal Java™ platform and provides sophisticated Windows NT® integration.

The V2200, V2250, and V2500 are fully upgradable to the V2600.

Left to Right: A-Class, R-Class, D-Class, L-Class, K-Class, N-Class, and V-Class

Figure 1. The HP 9000 Enterprise Servers

V-Class System Architecture

Introduction

The V2600 server is a 2- to 128 processor system with a scalable memory and I/O subsystem. The foundation of the V-Class architecture is HP's Scalable Computing Architecture (SCA), which is based on highly scalable nodes and a proven cache-coherent, non-uniform memory access (ccNUMA) technology (see Figure 2) that is a feature of HP-UX 11. The system, composed of up to four *cabinets*, appears to the operating environment, users, I/O subsystem, and applications as a single, scalable system, while the underlying hardware-based distributed memory exhibits a high degree of scalability.

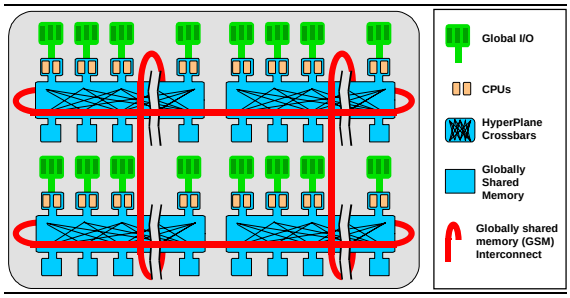

Figure 2. The V-Class Architecture

The V2600 operating environment is the industry-standard 64-bit HP-UX UNIX® operating system, release 11. This release provides a complete 64-bit environment, including a 64-bit address space, 64-bit file size and file system size, and 64-bit data types.

The V-Class system features:

- HP HyperPlane—an innovative crossbar capable of delivering 61.2 GB/s of aggregate system crossbar bandwidth
- 2 to 128 64-bit PA-8600 processors
- Up to 128 GB synchronous DRAM (SDRAM)
- Up to 7.68 aggregate GB/s I/O channel throughput
- Up to 112 industry-standard PCI I/O controllers

HP HyperPlane Memory Crossbar

The memory subsystem of each cabinet is based on HP HyperPlane, a crossbar technology that provides extremely high performance (15.36 GB/s bandwidth per crossbar) on all applications. The crossbar provides non-blocking access from CPUs and I/O channels to the memory subsystem. The use of HP HyperPlane prevents the performance drop-off associated with systems that employ a system-wide bus to handle memory and I/O traffic.

As shown in Figure 3, each V-Class cabinet contains up to 32 processors, eight memory boards, and eight I/O channels supporting up to 28 PCI I/O controllers. Each of the eight crossbar ports on the processor side connects to a single agent. Each agent supports two or four[1] PA-8600 processors and a 240 MB/s I/O channel. On the memory side of the crossbar, each port connects to a 32-way interleaved memory board.

[1]Configurations are allowed that support two processors per agent.

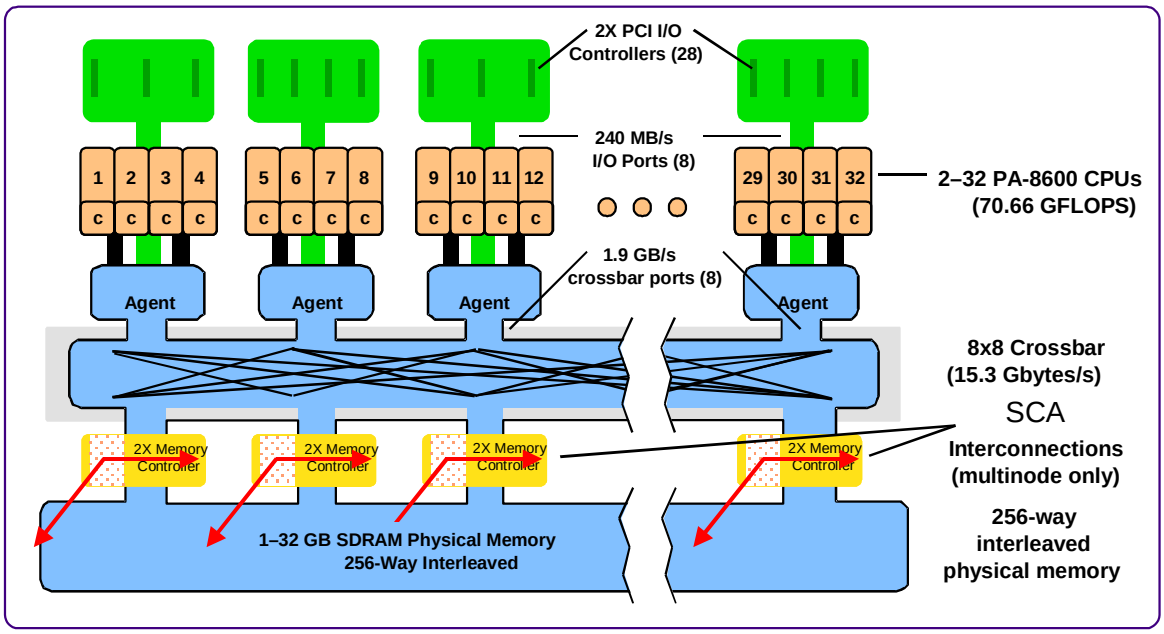

Figure 3. The V2600 HyperPlane Architecture

The crossbar operates at 120 MHz (8.33 nanosecond cycle time). The path width to the agent and memory controller chips is 64-bits wide, providing a rated bandwidth of each crossbar port of 960 MB/s (bi-directional). The crossbar is non-blocking so that all ports can be operating at full bandwidth as long as their target ports are unique. From each of the eight agents and eight memory controllers, there are two paths to the crossbar, an input and an output path. Thus, the aggregate bandwidth of each crossbar is 15.36 GB/s (120MHz * 8 B /cycle * 8 ports * 2 directions/port).

PA-8600 Processor

The V-Class is designed from the ground-up for the HP PA-8x00 processor family. These processors offer 4-way superscalar processing, yielding four operations per clock or 2.2 GFLOPS per processor at 552 MHz (see Figure 4).

The processor used in the V2600 is the PA-8600, and it has the following architectural features:

- Large on-chip cache (1 MB data, 0.5 MB instruction)
- Two 64-bit floating-point processing units each for load/store, multiply/add, divide/square root, integer, and shift/merge function (see Figure 4) for a total of ten functional units
- Dual load/store units[2]
- 56 entry instruction reorder buffer (IRB)[3]
- 10 outstanding memory requests
- Speculative execution
- Directed prefetch
- Static and dynamic branch prediction
- 2 clock maximum latency to cache (1 when pipelined)
- Dual-ported cache

The PA-8600 is binary compatible with all of the previous members of the PA-RISC family, which includes the PA-7100, PA-7200, PA-8000, PA-8200, and the PA-8500.

[2]Unique to the PA-8xxx processor family.
[3]The IRB size of 56 entries is the largest of any standard RISC processor.

PA-8600 Architecture

The PA-8600 is a true 64-bit chip with native 64-bit integer arithmetic. It supports a flat 64-bit virtual address space, although the chip exports 40 physical address bits. This corresponds to one TB (2^{40} bytes) of directly addressable memory. The chip also supports 32-bit addressing, providing backward compatibility with the PA-7100 and PA-7200 processors.

The PA-8600 is a decoupled architecture; that is, the instruction decode logic is not integrated with the functional units' pipeline logic. This architecture allows the chip to partially decode instructions well in advance of the instructions' actual execution by the functional unit(s). Decoded instructions are staged in queues within the chip. The PA-8600 can have up to 56 instructions in progress at any given time.

The processor can issue up to four instructions per clock cycle. To sustain superscalar performance to the greatest degree possible, the PA-8600 incorporates two independent floating-point functional units, two independent divide and square-root functional units, two independent 64-bit integer ALUs, two shift/merge units (although separate from the integer ALUs, only two of the possible four instruction mixes for these units can be issued per clock), and two independent load/store units. In the important case of the floating-point functional units, each is capable of issuing a multiply-add instruction on every clock.

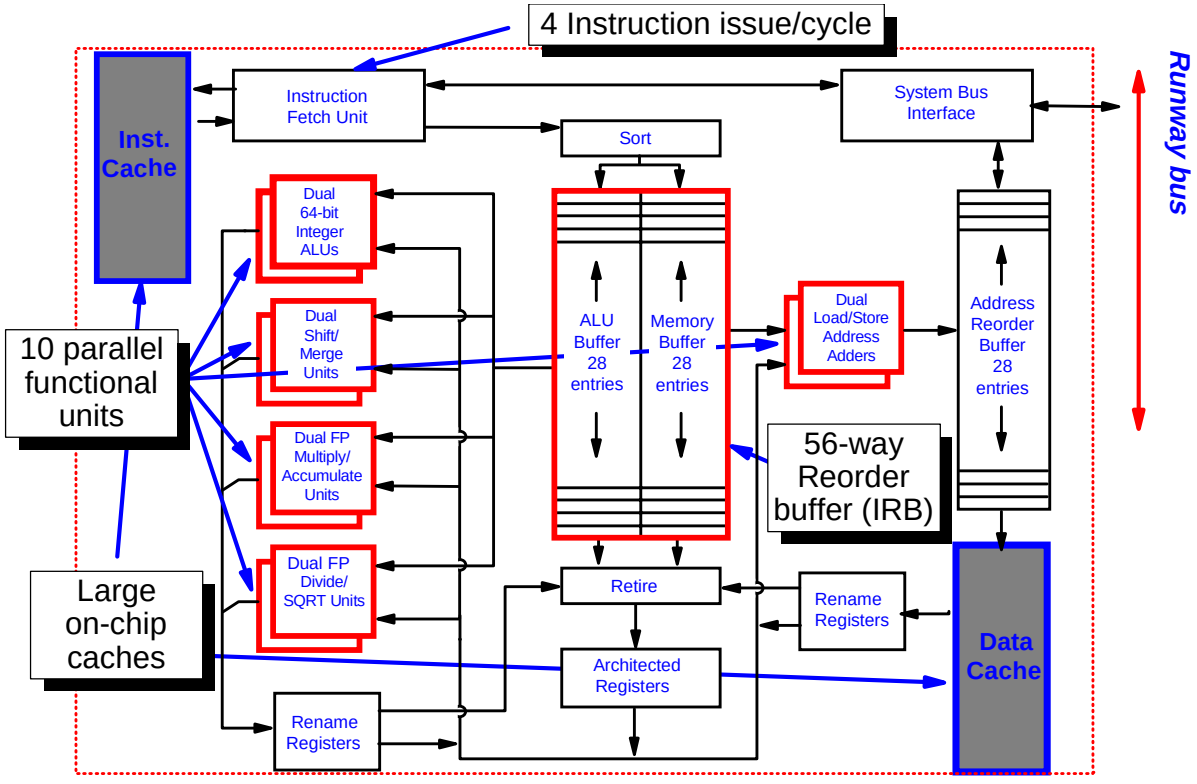

Figure 4. Overview of the Hewlett-Packard PA-8600 Processor

Thus, the peak floating-point performance of the chip is four times its clock rating. Note that the multiply/add is a compound instruction—that is, it is one instruction for which two floating-point operations are initiated. This not only increases performance, but the multiply/add combination only requires a single IRB entry and further increases the processor's efficiency. The multiply/add instruction has a 3-cycle latency, but under pipelined conditions, it delivers a result every clock.

To use the full set of functional units, the processor contains a 56-entry instruction reorder buffer (IRB), a dual-ported data cache, and the ability to fetch four instructions per clock from the large instruction cache. The processor can hold up to 56 instructions in the buffer and issue them when the requisite data and functional unit(s) are ready. The data dependencies for the instructions in the IRB buffers are known, and when the data and requisite functional units are available, the instruction is dispatched to the functional units.

PA-8600 Cache Structure

The PA-8600 provides access to more data in two cycles from its 0.5MB instruction and 1MB data caches than many systems provide in 10 or more cycles from an L2 cache. The reordering queues inherited from the PA-8000 do an excellent job extracting useful independent work from the instruction code stream. The processor's functional units can therefore keep busy, rather than stall as a conventionally pipelined processor would have to do during a two-state cache access.

One of the challenges in developing the PA-8600 is to create an on-chip cache that can fit into the allocated die area and still keep up with the reorder queues. This requires the data cache to support two simultaneous memory operations while maintaining a two-cycle access. HP accomplished this by using the same dual bank system developed for the PA-8000's off-chip data cache. With this system, the cache can be implemented with a simple single-port RAM design, conserving area. Since each lookup needs to concern itself with only half the cache, the physical dimensions traversed during a cache access are reduced, supporting a quicker cache access time.

Each 0.5MB data cache bank is implemented as four one-eighth-megabyte arrays, each providing a double-word of data plus error correction bits. Data is organized within the arrays so that either a full cache line can be addressed at one time or four ways of associativity can be addressed together. The cache line tags are held in four smaller separate and independently addressable RAM arrays. In this way, the data and tag can be accessed together, either for a data read or independently to effect a data store at the same time another store is accessing its cache line status.

The PA-8600 instruction cache is a 0.5MB four-way set associative pipelined cache that provides 128 bits of instruction plus pre-decode bits per cycle to the instruction fetch engine. For most workloads, a set associative cache is more likely to hit than an equivalently sized direct mapped cache. As a direct mapped cache fills with useful data, it becomes increasing unlikely that the next piece of requested data will find an unoccupied location in cache to use. Set associative caches increase hit ratios because there are multiple locations that can be used to store a given piece of data.

Speculative Execution

A powerful feature of the PA-8600 is its ability to perform speculative execution, which involves having the processor "guess" the path of execution and execute instructions in that path. If the guess is incorrect, the speculatively executed instructions are discarded. Speculative execution is aided by the sophisticated branch prediction mechanism based on a 2,048-entry branch history cache. The branch prediction determines the instruction sequence (i.e., which branch it believes will be taken) to execute and those instructions are (speculatively) executed. If the branch is mispredicted, those instructions are simply discarded since they have not yet been retired.

The Branch History Table (BHT) in the PA-8600 is a standard array of two-bit counters, but the information stored in the counters is not the direction of the branch (taken or not-taken). Rather, the counters record whether or not the branch went in the direction indicated by the static hint supplied by the compiler. If the static hint disagrees with the actual direction the branch followed, the counter is increased by one; if it agrees, the counter is decreased. Each time a branch is fetched, the BHT is consulted and, if the counter is zero or one, the static hint encoded in the instruction is followed. If the counter is two or three, the hardware predicts that the branch will go in the direction opposite the static hint.

If the compiler has done a good job of setting the static hints and most of the branches are strongly biased, most of the counters in the prediction cache will end up with counts of zero or one. For branches that are incorrectly hinted, the counts will tend to be two or three and the hardware will override the static hint provided by the compiler. With this algorithm, two branches of opposite bias which map to the same location in the BHT do not interfere with one another if each is hinted correctly.

The modified BHT enables the PA-8600 to combine the advantages of static and dynamic branch prediction methods in a single hardware structure, rather than two or three hardware arrays as some other prediction methods require.

Memory Subsystem

Each V-Class cabinet supports from one to 32 GB of SDRAM (Synchronous Dynamic Random Access Memory) physically distributed on two to eight memory boards (depending upon the packaging).

The use of SDRAMs permits the memory subsystem to operate at a higher clock frequency and at higher effective bandwidths than regular DRAMS. SDRAMs are increasingly being employed in systems at all levels (including PCs), resulting in a much higher performance memory subsystem. Additionally, the use of SDRAMs improves memory performance by preserving memory interleaving as memory is added to the system in odd increments (non-powers of 2).

Memory interleaving is in balance with the cycle time of the SDRAMs to reduce memory contention. This means that during sequential cache loads (typical of buffer copy operations found in commercial applications and most technical algorithms), the memory system will never stall due to a busy memory bank (one that is being refreshed). During a long sequential load or copy operation, the interleave factor is high enough that by the time a bank is accessed again, it will be refreshed and ready for the next operation.

Scalable Computing Architecture

With HP's Scalable Computing Architecture (SCA), the V2600 running a single copy of HP-UX can be expanded up to 128 PA-8600 processors. The SCA employs ccNUMA (cache-coherent non-uniform memory access) technology and combines the benefits of symmetric multiprocessing (SMP) and distributed memory (clustering) architectures. SMP architectures are well known, easy to program, and less costly to administer, while clustering architectures typically provide higher degrees of scalability (because of the distributed memory subsystem). As shown in Figure 5. HP's SCA delivers the programming model and appearance of an SMP system, as well as the scalability of a distributed memory subsystem.

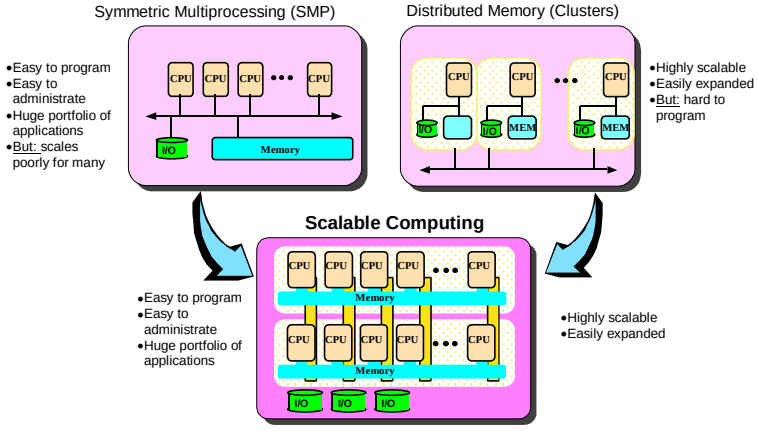

Figure 5. Scalable Computing and the V-Class

The V-Class SCA is a multilevel memory subsystem, in which each level is optimized for data sharing (see Figure 6). The first level consists of traditional SMP memory; the second level of the memory subsystem is created by tying first-level memories through a high-performance proprietary interconnect. This interconnect, called SCA HyperLink, provides multiple rings in two directions (x and y) for high bandwidth and fault resilience. The system thus becomes multiple symmetric multiprocessing cabinets connected by the SCA HyperLink, and it appears to the developer and users as a single, globally-shared-memory multiprocessor system.

The relatively large (32-way) building-block node improves overall application scalability by permitting as many as 32 processors access to up to 32 GB of local memory without off-node references. This reduces internode data traffic and greatly reduces average memory latencies.

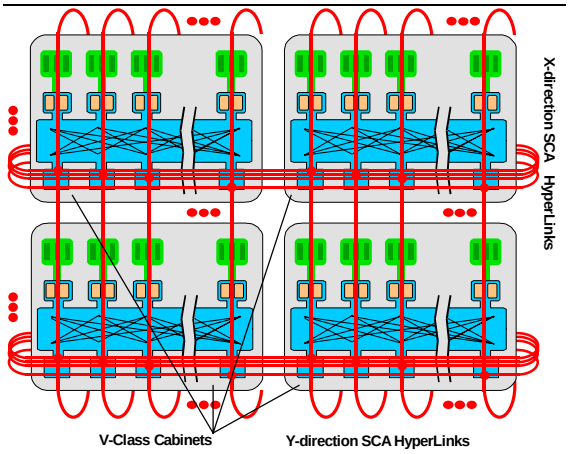

Figure 6. The V-Class Scalable Computing Architecture (SCA)

This tiered memory subsystem is optimal for several reasons:

- The low-latency shared memory of a cabinet can improve performance by effectively supporting fine-grained parallelism within applications, often by simply recompiling the program with the automatic parallelizing compilers. Cabinets may then implement coarser-grained parallelism with communication through shared memory and/or explicit message-passing mechanisms.
- The large first-level memory subsystem allows individual processors to access a very large (32 GB) local physical memory, allowing large serial and/or parallel applications to run with optimal memory reference performance and without having to make frequent off-node references.
- The distributed memory subsystem allows much greater system memory and I/O bandwidth.

SCA HyperLink

The low-latency interconnect called SCA HyperLink connects multiple V-Class cabinets within the system. This interconnect combines high bandwidth with low latency to provide system-wide coherent access to shared memory. Explicit message-passing (EMP) applications also use SCA HyperLink for extremely high-performance, inter-cabinet communication.

SCA HyperLink is implemented as a series of point-to-point unidirectional links. A single interconnect, say the x-direction, can be considered as a unidirectional ring. This ring actually consists of the eight independent data paths from each of the eight interface controllers in each cabinet. Each link between a pair of cabinets for each of the eight paths is completely independent from all others, permitting interleaved access across the rings (similar to interleaved memory requests). An important performance issue is that the SCA HyperLink controllers permit multiple outstanding requests, greatly reducing effective inter-cabinet latencies.

Owing to this split transaction protocol and the concurrency of the acknowledgment packets (which are “invisible” from a latency perspective), there is no notion of a nearest neighbor or “hops.” If the requesting cabinet and target cabinet are physically near each other in the downstream direction, then the request packet has a short path and the response packet a longer path. Likewise, if the two cabinets are far apart from each other, in the sense of directional separation on the unidirectional ring, then the request packet has a long path and the response packet’s path is short. Regardless, the combined path length of the request packet and the response packet is a full circuit around the ring.

The V-Class system provides hardware support for globally shared memory access, while a system without this feature can only emulate shared memory by moving pages from node to node under software control. The use of shared memory reduces the amount of overhead involved in parallel applications, increasing scalability and overall performance.

SCA HyperLinkcache

To minimize memory request latencies across the SCA HyperLink cache, each cabinet contains a cache of memory references that have been made over the interconnect to other cabinets. This is referred to as the SCA. Any data that has been moved from another cabinet into a processor cache on the cabinet and that is still resident in a processor cache is guaranteed to also be encached in the SCA HyperLinkcache. Consequently, HyperLinkcache directory information can be used to locate any global data that is currently encached by the cabinet. SCA HyperLinkcache is physically indexed and tagged with the global physical address.

The system guarantees cache coherence between multiple cabinets, so two or more cabinets that map the same global address will get a consistent view. This is done by maintaining a linked sharing-list that contains either all the cabinets sharing each cache line or the cabinet that exclusively owns the cache line.

The system keeps a record of which processors have encached each line in the SCA HyperLinkcache, so that interconnect coherency requests can be forwarded to the appropriate cabinets containing the processors affected.

SCA HyperLinkcache lowers the effective system-wide latencies and increases the effective bandwidth of inter-cabinet references. Application performance is increased, regardless of the programming model employed, and overall system throughput in a mix of parallel and non-parallel jobs is higher.

Enterprise Operating Environment

HP's Enterprise Operating Environment involves all aspects of meeting the business-critical needs of computing from the desktop to the data center to the e-serviceCenter—including performance and scalability, resilience (high availability), integration, security, and manageability. As shown in Figure 7, these needs are met through layered software (middleware) and HP-UX, HP's industry-standard UNIX operating system and the core of HP's Enterprise Operating Environment.

This section reviews HP's range of solutions for the enterprise for each of these needs, then describes the HP-UX in more detail.

	Performance and Scalability	Resilience and High Availability	Integration	Security	Manageability
Middleware	• Optimizing compilers	• Multi-system, High Availability MC/ServiceGuard, MC/LockManager • Enterprise Clusters • OnLine JFS	• Colliance Program for integration of HP-UX and NT • Domain Internet Platform • Enterprise Networking • Application Development, including Java	• Praesidium family of security functions: – Single Sign-On – Authorization – Authentication – B1 security	• Process Resource Manager • HP GlancePlus • HP OpenView family of Systems Management • Distributed Internet Services including DCE
HP-UX	• SMP • 64 bits • Kernel Threads • Paging • Fibre Channel • Native Java™ • Web • Networking	• Quality • Robustness • Single-system High Availability – Dynamic Processor and Memory Resilience • Fast Recovery • Journaled File System	• Web-enabled OS • Interoperability with NT, NetWare, DEC, other UNIX® systems	• C2 compliance • Cryptography APIs • Encryption, secure network connections through IPSec	• System Administration Manager • SD/UX • Ignite/UX • Distributed Print Service • Instant Ignition • DCE Client

Figure 7. Aspects of the Enterprise Operating Environment

Performance and Scalability

HP-UX is suitable for a broad range of business IT environments, from large NFS servers to very high-end computing environments. It is also ideal for a broad range of hardware platforms, from the single-user desktop workstation to very large systems (VLS) in a HyperPlex e-serviceCenter Solution clustered environment.

Key features associated with HP-UX performance include:

- Highly parallelized kernel for efficient scaling of operating system functions on multiprocessor platforms, from technical desktops all the way up to the 128-way V2600.
- The 64-bit version of HP-UX 11 allows the processing of much larger applications and data sets by addressing up to 2^{40} bytes per process of virtual memory.
- Enhanced process scheduling features for optimum performance and scalability. These features include kernel threads, gang scheduling, performance optimized page sizing (POPS), and thread affinity scheduling.

Resilience and High Availability

As companies become increasingly intolerant of computer downtime, they are turning to HP high availability technologies and our services and support programs to keep their systems and applications running. HP MC/ServiceGuard is the industry's most successful UNIX clustering solution—providing automatic failover and ensuring the availability of mission-critical applications at more than 25,000 installations worldwide. Based on MC/ServiceGuard technology, HP disaster tolerant solutions provide disaster protection and recovery for a wide range of geographically dispersed data centers.

HP offers pre-tested, pre-configured mission-critical solutions that deliver 99.95% availability—translating into downtimes totaling less than 4.3 hours per year. These high availability solutions include not only HP 9000 Enterprise Servers, but also HP MC/ServiceGuard, HP Critical System Support, high-availability storage solutions, and network and system management tools. HP extends its 99.95% uptime commitment to include the Oracle® database, Cisco network solutions, and HP's enterprise storage solutions.

The HP solution is backed by the "5nines:5minutes" program, HP's vision to deliver end-to-end solutions in the year 2000 that provide 99.999% availability—translating into only five minutes of downtime each year. HP's partnerships with Cisco and Oracle are key to delivering these true, end-to-end availability solutions.

In addition to offering highly available configurations, individual systems benefit from the resilience features of HP-UX. These features include:

- **Journaled File System:** Guarantees file system integrity through a fast recovery process in the event of a system failure
- **Dynamic Processor Resilience (DPR):** Enables taking a faulty processor offline without having to reboot the system[4]
- **Dynamic Memory Resilience (DMR):** Deallocates faulty memory to prevent data corruption or downtime in the event of failing memory[5]

Integration

The HP Colliance Program enhances HP-UX's ability to coexist and integrate with mixed environments, ensuring that IT administrators can combine HP-UX with Windows NT and Novell NetWare across their organizations.

Netscape FastTrack: Provides a complete solution for creating and managing Web sites on the Internet or an Intranet.

Oracle Web Application Server: Enhances Web servers with sophisticated features required for building and deploying real business applications on the Internet.

HP-UX Java Virtual Machine, Just-in-time compiler for Java, Java SDK: Allows developing and deploying platform-independent Java applications.

Internet and Intranet connectivity: HP-UX 11 incorporates a broad range of Internet services. These services enable any HP 9000 Enterprise Server to operate in a TCP/IP-based Internet/intranet or other corporate environment that requires business use of distributed computing technology.

[4]To be added in a future release.
[5]To be added in a future release.

Security

HP-UX 11 has a number of security features to give customers comfort that their systems are secure from external—or internal—tampering.

Compliance with DOD C2 security requirements provides server security and is required on many government-related projects.

System auditing. Improves user accountability and deters unauthorized activities.

The Praesidium Enterprise Security Framework. HP has developed the Praesidium Enterprise Security Framework to secure the Extended Enterprise and provide secure electronic commerce. Praesidium takes a comprehensive view of enterprise security that goes beyond firewalls and spans NT, UNIX, and legacy computing environments.

Manageability

The system management facilities supported by HP-UX are designed to easily manage both single-server and complex networked systems, reducing total cost-of-ownership and increasing the productivity of your IT staff.

System Administration Manager (SAM): Provides easy, graphical- or browser-based management of system resources

Software Distributor: Has robust, standards-based software and application distribution and updating

HP Ignite/UX: Enables customizable initial system configuration and easy replication across the environment

HP OpenView: Bundled in OpenView-ready CA TNG framework

HP-UX UNIX Operating System

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

Key Features

HP-UX 11:

- Includes full 64-bit functionality
- Allows the processing of large applications and data sets
- Supports 32- and 64-bit applications, ensuring investment protection
- Maximizes availability or system uptime through increased resilience
- Integrates seamlessly with the Internet for better Web enablement
- Provides binary compatibility into the next generation of processor technology (IA-64)
- Also integrates with Windows NT and Novell NetWare

HP-UX provides forward binary, in which an application developed on an earlier version of HP-UX is ensured to run smoothly on HP-UX 11. Thus, current 32-bit applications can run on either the 32- or 64-bit versions of HP-UX 11, without requiring forced recompilation.

This compatibility will carry forward into the next generation of processor technology being developed—IA-64—allowing applications currently running on PA-RISC to execute unchanged on IA-64–based platforms.

Breadth of Applications Portfolio

To make certain that our customers can deploy optimized IT solutions with minimal cost and delays, HP has built strategic alliances with the industry's leading system integrators and developers of enterprise, Internet, and technical-computing applications. More than 4,000 software vendors have developed and tuned more than 15,000 applications for HP-UX. More mainframe-based solutions have been ported to HP systems than to any other UNIX platform.

Standards

HP-UX 11 continues HP's tradition of supporting many UNIX standards. HP-UX 11 is UNIX 95-branded, as specified by X/Open®'s Single UNIX Specification. Furthermore, the 64-bit version of HP-UX complies with the LP-64 data model standard.

Performance Features

HP-UX has many features designed specifically to enhance application performance and overall data center throughput. These features include:

- **Performance-optimized page sizing (POPS)—**the ability to vary page sizes dynamically to decrease page translation overhead.
- **Kernel threads—**support for thread-level parallelization of applications, including compiler-based or with an explicit industry-standard, POSIX-compliant applications programming interface (API) library.
- **Support for "gang scheduling"—**a mechanism for efficiently scheduling parallel applications that permits multiple threads from a single process or a set of processes to be scheduled concurrently as a group. The concurrent scheduling of threads permits low latency interactions between threads in shared memory parallel applications. The gang-scheduling feature can provide significant improvements to parallel application performance in loaded timeshare or oversubscribed[6] environments.
- **Thread affinity scheduling—**a feature that determines how long it has been since the thread last ran and, if practical, attempts to reschedule the thread back on its original processor. This helps performance if there is still data resident in the processor's caches for this process.

64-Bit Computing

Modern RISC processors have increased in performance by at least an order of magnitude in the last ten years. As a result, the amount of data that can be processed in a "short" amount of time has grown proportionally (for all types of computing) and will soon exceed the amount that can be addressed by 32 bits of address space. Consequently, operating environments (and hardware) are moving to full 64-bit environments.

Today's use of 64 bits is found in high-end database and compute-intensive scientific modeling applications. In the database arena, the major database vendors have introduced Very Large Memory (VLM) products that take explicit advantage of 64-bit architecture scalability to achieve increases in performance, especially for very large decision support and OLTP applications.

As corporations and government require increased computing power and the ability to handle larger amounts of data, the demand for 64-bit computing will grow. Other applications that will likely take advantage of 64-bit computing include:

- Internet-based commerce and large Web servers.
- Multimedia applications, including video/audio servers and 3D animation.
- Technical applications, such as ECAD/MCAD and fluid dynamics.

The phrase "64-bit computing" refers to the total system environment's ability to process 64-bit data, instructions, and addressing. A true 64-bit environment contains a 64-bit CPU with 64-bit registers and data paths, 64-bit memory addressing, 64-bit Direct Memory Access, and a 64-bit kernel as the heart of the operating system. A 64-bit environment has significantly higher scalability than its 32-bit counterpart and higher performance for many kinds of applications. Not all operating systems that claim to be 64-bit are true 64-bit environments, because they lack important elements, such as a full 64-bit kernel.

Memory addressing is one of the most important system elements that benefits from the scalability of 64-bits. While a 32-bit operating system provides flat addressing of up to 2^{32} 32-bit words, or 4 GB of memory, a 64-bit operating system provides flat addressing for 2^{64} 64-bit words, or 18 billion GB (18 exabytes) of memory[7]. As shown in Figure 8, scalability increases associated with 64 bits are dramatic.

Word Length	Mathematical Expression	Relative Scale
8-bit	$2^8 = 256$	Business Card
16-bit	$2^{16} = 65,536$	Desktop
32-bit	$2^{32} = 4.29E +09$	City Block
64-bit	$2^{64} = 1.84E +19$	Surface of the Earth

Figure 8. A Comparison of 64-Bit vs. 32-Bit Addressing

[6]Oversubscribed refers to a multiple parallel application environment in which the total thread count exceeds the number of processors in the system.

[7]Note this is the theoretical figure for 64-bit addressing. At current RAM prices, 18 exabyte (EB) would cost approximately US $288 billion.

Benefits of 64-Bit Computing

A true 64-bit environment provides a significant number of benefits, concentrated in the two major areas of scalability and performance across the operating environment. Every critical resource capacity from memory, storage, and program addressability is increased by several orders of magnitude. This enormous capacity increase translates to raw high performance when harnessed to its fullest potential.

A very large memory capacity allows a greater number of in-memory processes. The in-memory nature alone is extremely fast. Memory is accessed about 10,000 times faster than disk drives. For large applications whose virtual size exceeds the available physical memory (and therefore pages and/or swaps to disk), moving to a 64-bit operating environment with generous RAM can increase performance dramatically. A 64-bit environment can have huge system memory capacity to support the largest applications today and in the near future.

For example, the large database vendors enhance the scalability and performance of their products by exploiting 64-bit capabilities to get a higher percentage of data into memory and assure that full indexes reside in memory. By reducing swapping to disk, search and access performance significantly increases. For example, a transaction-oriented lab test at HP yielded a 34 percent reduction in I/O operations per transaction when changing from a 32-bit environment to a 64-bit environment.

The complementary increase in file sizes and file system sizes further enhances performance when moving to a 64-bit environment. Depending on the application, the use of one large file versus a large number of smaller files results in decreased system overhead. The overhead efficiency inherent in a 64-bit environment comes from the use of fewer file descriptors and provides a side benefit of easier system manageability by requiring maintenance of only one file versus several files.

The following table summarizes the sources of increased performance and scalability associated with 64-bit computing by type of application.

Example	Sources of performance and scalability gains
Large database	Reduced paging and swapping Larger memory allocation per user Many more users Large file implementations
Decision support	Reduced paging and swapping Direct addressing Large file implementations
Technical applications	Large process data space More available shared memory segments Reduced paging High-precision arithmetic

HP-UX 11 provides the 64-bit features that satisfy the requirements of IT managers and end-users, including performance and scalability, investment protection, standards conformance, and coexistence and interoperability between 64-bit and 32-bit applications running on the same platform.

Compilers

Standards compliance means easy portability and investment protection for your existing source code. HP's long-standing commitment to standards is reflected in the HP C++, HP Fortran 90, and HP C compilers.

The HP compilers support application development in 32-bit or 64-bit modes. 64-bit support means that applications can now solve dramatically larger problems, and developers can choose memory intensive algorithms to solve problems much more quickly.

HP compilers perform many powerful optimizations to assure that applications take full advantage of the V-Class platform. In addition to a wide range of traditional machine-independent optimizations, the HP compilers perform the following:

- **PA-RISC-specific optimizations—**the HP compilers strike a balance between the hardware and software that exploits the best of each technology. The HP compiler's optimizations and transformations deliver the performance required for today's applications.
- **Interprocedural optimizations—**by analyzing the application as a whole, instead of on a procedure-by-procedure basis, the optimizer uses knowledge gained in one routine to optimize others. Examples include automatic inclining and procedure cloning.

- **Profile-based optimization (PBO)—**with PBO, the compiler analyzes run-time characteristics of the application and performs optimizations that contribute to improved run-time performance.
- **Automatic parallelization—**the HP C and HP Fortran 90 compilers automatically detect loops that can be parallelized and generate code that will execute across multiple processors on the server. No source code changes are necessary, so applications remain portable and, at the same time, take maximum advantage of the multi-processor hardware for improved time to solution.

C++ is the fastest growing application development language and is quickly becoming the preferred programming language for the automotive, aerospace, defense, telecommunications, banking and other financial institutions, and manufacturing industries. HP C++ includes:

- **Standards compliance for portability and investment protection—**HP C++ supports the ANSI X3J16/ISO WG21 Draft Working Paper definition.
- **Precompiled header files for reduced compilation times and smaller object files—** a common PC feature.
- **Standard C++ Library—**offering developers full library functionality, including the Standard Template Library (STL).
- **Template instantiation and exception handling—**to provide excellent performance and industry-leading functionality.
- **Tools.h++—**a C++ foundation class library with over 120 reusable classes that provides a higher-lever interface to the Standard C++ Library.

HP DDE, HP PAK, and Blink Link

HP C++, HP C, and HP Fortran 90 each include the Distributed Debugging Environment (HP DDE), a powerful symbolic debugger; HP Performance Analysis Kit (HP PAK), a performance analysis tool; and Blink Link, a link-and-go utility, as no-cost, bundled components.

- **HP DDE—**is a symbolic debugger with an intuitive, graphical task-oriented interface. With a fully customizable debugging environment, you can control both program and process execution. It features continuously updated graphical displays that show the program state and synchronized source and PA-8000 assembly code.
- **HP PAK—**allows you to capture a rich set of performance metrics and display them in easy-to-understand graphical and tabular formats. You can analyze the relative use of system resources by each process, examine process performance at the procedure level, and analyze compute-bound procedures at the statement or instruction level. HP PAK also supports performance analysis of multithreaded programs.
- **Blink Link—**accelerates the linking cycle and provides a set of utilities to be used with incremental linking (linking only modified modules).

Development Tools

Introduction

The V-Class (and HP-UX 11) supports a sophisticated suite of program development tools, including a variety of CASE (Computer Aided Software Engineering) applications and standard compilers. A complete portfolio of performance analysis and debugging tools are especially designed for tuning, developing and porting codes in a parallel environment.

MPI

HP MPI provides application performance by harnessing the power of multiple HP processors or systems. Using this industry-standard method of implementing parallel applications, program developers can dramatically improve the performance of their software on real work. HP MPI is fully compliant with version 1.2 of the standard and supports a growing subset of the MPI-2 functionality. MPI establishes a portable and flexible approach for message passing applications. Applications developed using MPI on HP systems can easily be ported to other platforms that provide an MPI implementation.

HP MPI provides an application programming interface (API) that can be implemented on a wide range of hardware platforms, including K-Class and V-Class servers, as well as workstations. This means that MPI applications can be scaled up to execute on a single server such as the V-Class or a cluster of servers and workstations connected via high-performance networking.

HP MPI includes XMPI, a powerful monitoring, tracing, and visualization tool for MPI programs originally developed by the Ohio Supercomputing Center's LAM team. With XMPI, programmers can quickly and easily identify communications bottlenecks and ultimately speed up their application's time-to-solution.

CXperf

CXperf, an application analysis tool, helps application developers deliver the power of the V-Class server. CXperf clearly and dramatically helps improve the performance of complex real-world applications, including parallel applications.

CXperf measures wall clock timing, CPU timing, memory access counts and memory latency timing to provide programmers with a comprehensive view of how their application is performing.

To enhance understanding of the data, CXperf graphically displays the behavior of the application as it occurs during execution (see Figure 9). Using the graphical tools, the programmer can quickly identify where CPU or memory bottlenecks occur in the application, helping identify which part of an application offers the greatest potential for improved performance through optimization.

Mathematical Libraries

To meet challenging schedules and user requirements, application developers look for ways to speed creation of new applications. The MLIB Mathematics Library from HP shortens the software development cycle and reduces execution time of applications. MLIB contains robust callable subprograms, including all Basic Linear Algebra Subroutines (BLAS) Level 1, 2, and 3; LAPACK, a collection of commonly used FFTs (Fast Fourier Transforms); linear algebra solvers; and convolutions. HP chooses the best algorithm for each routine, including parallelization if appropriate, to achieve optimal performance. HP expertly codes each routine using these algorithms for maximum application performance.

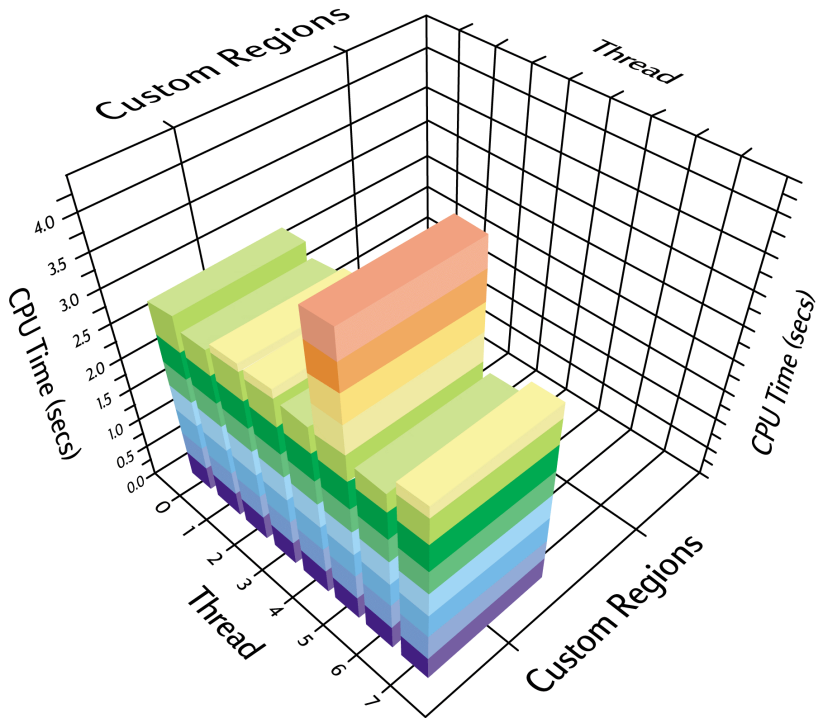

Figure 9. Sample Output from CXperf Performance Analyzer Tool

HyperPlex e-serviceCenter Solution

The HP HyperPlex e-serviceCenter Solution architecture is based upon a combination of leading-edge technologies, including HP's SCA, SMP, and parallel computing and clustering system management tools. The HyperPlex e-serviceCenter Solution provides the ability to flexibly configure central-site facilities that provide unmatched scalability and high availability.

As shown in Figure 10, the HP 9000 HyperPlex e-serviceCenter Solution provides a flexible and scalable clustering environment that can easily grow to accommodate large-scale consolidation, business-critical high availability, and performance requirements. HyperPlex solutions add value by providing this degree of scalability and flexibility while maintaining a single point of system administration and user contact across the cluster. In addition, the HyperPlex environment offers compatibility with your existing HP-UX and open systems environment.

The HyperPlex Solution also delivers investment protection. The highly adaptable HyperPlex environment allows you to add-in and scale emerging technologies such as parallel database management solutions, new platform technologies such as HP's SCA, and future servers based on the highly anticipated Intel IA-64 processor technology. In short, HP's HyperPlex clustering technology brings together the best enterprise computing products HP has to offer into a solution set that can be tailored to solve your largest computing problems.

Summary

HP-UX 11 continues HP's long-standing commitment to investment protection through binary compatibility, which allows fully linked, well-behaved HP-UX 9.x and 10.x applications to run on HP-UX 11. The 64-bit version of HP-UX 11 supports both 32- and 64-bit applications that can communicate using all standard inter-process communication methods such as IPC, pipes, and shared memory. This allows 32-bit applications to directly benefit from the performance of 64-bit databases. A wide range of applications is available on HP-UX 11, including very large memory (VLM) versions of databases from Informix, Oracle, and Sybase.

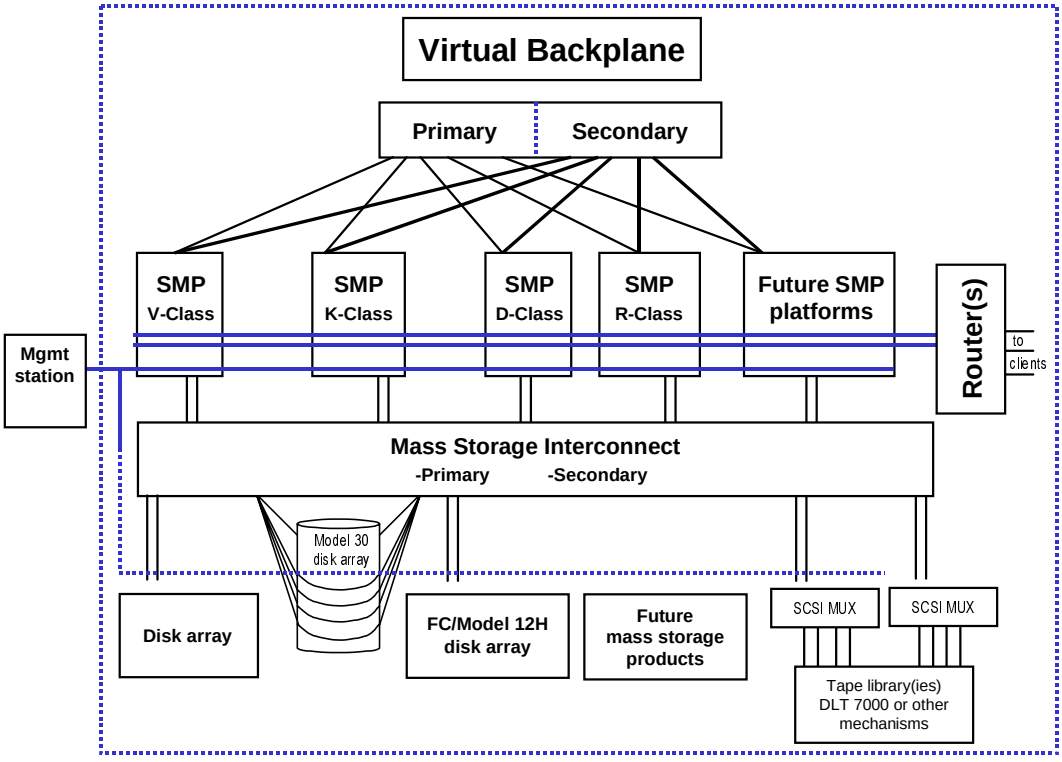

Figure 10. An Overview of the HyperPlex Architecture

I/O Subsystem

Introduction

Each V-Class system supports a highly scalable I/O subsystem based on up to 32 240-MB/s I/O channels directly connected to physical system memory through the memory subsystem (see Figure 11). Each I/O channel supports a dedicated PCI bus with three or four 2X PCI (Peripheral Component Interconnect) controllers, each capable of a peak performance of 240 MB/s[8].

Each intelligent I/O port is capable of Direct Memory Access (DMA) transfers directly to and from any physical memory units in the system. This eliminates CPU involvement in data transfers, reserving them for user work. It also streamlines data transfers for such things as large disk blocks and high-speed network connections.

In addition, high performance connectivity options, such as HP's scalable HyperFabric interconnect support direct connections to other servers and high performance networks. The I/O subsystem supports a broad range of automated tape libraries, as well as industry-leading high-performance and high-capacity tape devices.

Peripherals Subsystem

V-Class 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. The V-Class also supports Fibre Channel with arbitrated loop technology (FC-AL). Both FWD SCSI and FC are shipping in volume with V-Class to support the full range of HP on-line, near-line, and off-line storage solutions.

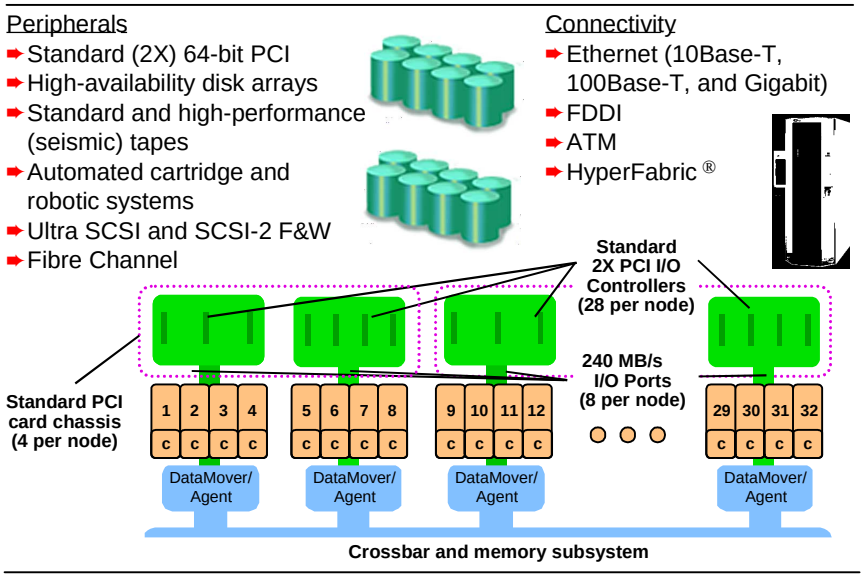

Figure 11. The V2600 Cabinet I/O Subsystem

[8] 2X PCI is 64-bits wide at 33MHz, resulting in a peak speed of 266 MB/s. This is limited, however, by the peak channel speed of 240 MB/S.

Mass Storage Options

The V-Class and K-Class processors support 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. Generally, the V-Class and N-Class systems support the same external I/O devices. A few devices are specific to either the V-Class or N-Class nodes. Some of the available mass storage devices are described below.

Inexpensive Disk Arrays—JBOD

HP offers external high-performance JBOD (just a bunch of disks) disk storage in expandable packaging for K- and V-Class servers. JBOD packaging provides a large amount of high-performance disk storage at the lowest possible price.

Each V-Class node can have up to 24 SCSI interfaces connected to the external JBOD cabinet. Disks in the external enclosure can be controlled by up to two SCSI interfaces for every ten disk drives. Each SCSI interface provides as much as 20 MB/s of I/O bandwidth.

HP Model 30/FC Disk Array/Fibre Channel

Based on an innovative architectural design that emphasizes high-availability and capitalizes on the benefits of Fibre Channel, the HP Model 30/FC disk array optimizes performance, scalability, and reliability while affording functionality to meet the storage needs of the organization.

The Model 30/FC offers Fibre Channel arbitrated loop, permitting scaling to high-capacity points and extended distances—up to 500 meters without Fibre Channel hubs or up to 3.0 km with cascaded hubs. The array supports RAID levels 1, 1/0, 3 and 5, configured within the same subsystem, and dual-active storage controllers offering automatic failover. The dual-active controllers afford dual 100 MB/s FC-AL links to a single host or multiple hosts for supporting clusters. Each array can be configured to support up to thirty 8.8-GB disks, offering 264 GB of storage with up to 64 MB of read/write cache.

HP SureStore E Disk Array

The HP SureStore E family is designed to provide the highest levels of performance, connectivity, capacity, and availability. The HP SureStore E family offers many data availability features, such as redundant power supplies, full system battery backup, non-disruptive maintenance, remote maintenance, dual-initiators, and several high-availability options, including hardware mirroring and RAID. HP SureStore E implements RAID 5 capabilities in hardware at the disk level.

HP AutoRAID

Hewlett-Packard offers a hierarchical adaptive RAID technology that fronts RAID Level 5 storage with RAID Level 1/0. Called HP AutoRAID, this technology keeps more-active data in RAID 1/0 and less-active data in RAID 5. This approach provides high availability and data security without the expense of a fully mirrored system. The AutoRAID technology uses a dynamic mapping process called log-structured RAID, a lookup table arrangement that gives HP AutoRAID a great deal of flexibility in using disk blocks and disk drives.

Implementation of HP AutoRAID gives non-technical users a way to automate RAID levels and optimize disk drive usage during online network events. AutoRAID looks at the system workload, intelligently determines the RAID level from which it can get the best read/write performance, and stores data using RAID Levels 1 (mirroring) and 5. The AutoRAID controller can convert data stored in one RAID level to another without disrupting the host’s view of the data.

An AutoRAID auto-configuration feature allows the user to add more capacity to the system without spending time formatting or configuring the added disk module. The user needs only to plug in the disk module and build another virtual disk. Another essential feature of HP AutoRAID is “active hot spare,” which uses the space that the extra disk offers and spreads it across the array rather than allowing a non-active disk to remain unused. The additional space is used to increase RAID 1/0 on the arrays.

HP Fibre Channel-to-SCSI Multiplexer

HP offers configurability and scalability through use of Fibre Channel-to-SCSI multiplexers. These devices have two Fibre Channel ports and four SCSI ports. The multiplexer can connect multiple nodes or systems to multiple disk storage devices and can be used to create large high-availability configurations.

Tape drives and tape libraries can be connected to a single system through the multiplexer. Using the Fibre Channel-to-SCSI multiplexer enables organizations to connect HP systems to tape drives and libraries that may be hundreds of feet away, which is often the case in large data center or research computing environments.

Tape Drives and Tape Libraries

HP supports the open systems line of Storage Technology, Inc., high-performance tape drives and tape libraries. Redwood D3, Timberline 9490, and DLT tape drives are supported. STK 9710 Wolfcreek with Timberline 9490 or DLT drives are also supported via a multiplexer or direct SCSI connection.

Networking Subsystem

V-Class networking options are designed to meet the connectivity needs of a broad range of users. Users requiring connections to the most commonly used interfaces such as 10-Mbps Ethernet, 100-Mbps Ethernet, and FDDI can seamlessly integrate the systems into their network infrastructure.

High performance network interfaces such as 100 MB/s Fibre Channel and 155 Mbps ATM offer balanced, high-bandwidth networking options for distributed applications and global access to corporate information.

Ethernet

The V-Class systems support one or more industry-standard Ethernet connections via a single PCI controller. V-Class systems support the industry-standard 10/100Base-T, providing a low cost and high performance network implementation option.

Gigabit Ethernet

Gigabit Ethernet is an extension to the highly successful 10 Mbps and 100 Mbps IEEE 802.3 Ethernet standards. Offering a raw bandwidth of 1,000 Mbps, Gigabit Ethernet maintains full compatibility with the huge installed base of Ethernet nodes.

Gigabit Ethernet meets several essential criteria for choosing a high-speed network:

- Easy, straightforward migration to higher performance levels without disruption.
- Gigabit Ethernet follows the same form, fit and function as its 10-Mbps and 100-Mbps Ethernet precursors, allowing a straightforward, incremental migration to higher-speed networking.
- Low cost of ownership—including both purchase cost and support cost.

FDDI

The V-Class I/O system supports industry-standard FDDI (Fiber Distributed Data Interface), offering the next level of performance beyond Ethernet with a theoretical peak transfer rate of 100 Mbps. FDDI can span up to two kilometers between network nodes, support up to 500 connections on a network, and cover a total distance of 100 kilometers. On-board station management is fully integrated on the host adapter. The V-Class implementations are single-attached and dual-attached fiber optic controller using a single PCI interface.

HIPPI

HIPPI (High Performance Parallel Interface) is a dual simplex, point-to-point technology that operates at 800 Mbps (100 MB/s). HIPPI was originally developed in the late 1980s to serve the high-bandwidth needs of supercomputers and high-end workstations. It uses both copper and fiber cabling to connect systems up to 10 kilometers apart, and it operates bi-directionally offering an aggregate throughput of 1,600 Mbps. It can be used for TCP/IP traffic or operate in a “raw” based mode, called HIPPI-FP (framing protocol), to support a wide range of applications.

Serial HIPPI is the fiber-optic extension of the HIPPI standard and allows the same signals to travel over fiber-optic cabling rather than copper. A full line of network switches, gateways, and network interface cards that support fiber cabling have been developed by a number of vendors.

Fibre Channel

Fibre Channel (FC) has quickly emerged as a *de facto* industry standard for high-speed data communications. Fibre Channel provides an ideal solution for high-speed data communications in an extended environment covering long distances and remote sites (e.g., a virtual data center). HP's Fibre Channel technology supports 100 MB/s.

ATM

ATM (Asynchronous Transfer Mode) is an emerging communications technology that provides an increase in network speed and capacity compared to conventional network technologies. ATM improves the performance of today's network applications and provides a scalable roadmap for bandwidth increases required for future network growth, virtual LAN simulation, and switching circuit.

The I/O system implements a 155 Mbps OC-3 ATM interface using one PCI controller card. A future 622 Mbps enhancement is also planned. Switched virtual circuit capability providing classic IP services are fully integrated into the scalable networking infrastructure.

X.25

X.25 has been the reliable wide area asynchronous communication protocol. The V-Class supports up to 16 channels of X.25 communication pipelines.

HyperFabric Interconnect

HP's HyperPlex clustering technology is based on a high-performance switching technology called HyperFabric, which enables hosts to communicate between each other at speeds of 160 MB/s in each direction per switch port. (With cable lengths longer than five meters, the speed is limited to 80 MB/s.)

HyperFabric is based upon high-speed interface cards and fast switches using wormhole routing technology. The latency of the core switching chip is 100 nanoseconds (card-to-card latency). Current products can support up to 16 hosts with a single switch (eight hosts in a high-availability, failover environment). Specialized configurations that permit switches to be meshed together to provide arbitrarily large configurations are available.

HyperFabric provides scalability for any switch topology, because there may be as many packets traversing a switch concurrently as the switch has ports. In a 4-port switch, for example, four separate packets may traverse the switch concurrently (assuming they are destined for different nodes).

The HyperFabric interface cards adhere to the industry-standard PCI form factor. Using the scalable I/O system of the V-Class, multiple HyperFabric cards may be "stacked" to provide higher node-to-node bandwidths if required by applications. The HyperFabric product consists of support for TCP/IP and the hardware SCN feature for SAP applications running OPS.

HP-UX 11 also supports a "lowfat" protocol, which further reduces application-to-application latencies. This is important in scientific (distributed memory) applications, as well as database/OLTP applications that use a shared lock manager.

V-Class Packaging

V-Class systems are air-cooled and stand-alone (i.e., do not require a front end), and they do not require a raised floor. The system is mechanically scalable; if additional performance and capabilities are required, additional memory, processors and I/O may be added.

External peripherals like JBODs, disk arrays, and tape libraries are available as options.

V-Class systems can be stacked, providing tremendous floor-space savings often needed in data center environments. Additionally, stacked V-Class systems can be configured with HP's MC/ServiceGuard High Availability solution, providing failover capabilities and maximum system uptime. The V-Class can also act as a node within a HyperPlex cluster, providing server consolidation features and higher levels of performance and availability.

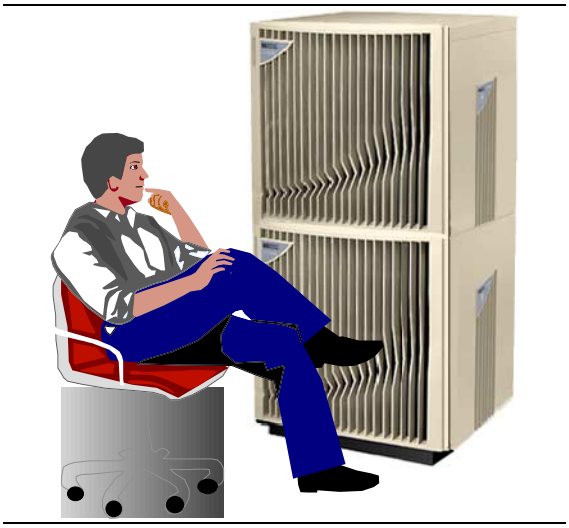

Figure 12. Two-Node (64-CPU) V-Class System

Conclusions

The e-serviceCenter requires a bulletproof infrastructure. And that infrastructure is powered by HP's industry-leading, high-performing, enterprise servers and the robust HP-UX operating environment.

Balanced and Scalable Architecture

With its large physical memory, the world's fastest RISC processor, and the world's highest performance backplane, the V2600 provides the perfect balance for situations where capacity and capability are paramount.

The V-Class is one of the highest performing servers in the market, translating to competitive advantages like faster time-to-market, broader services, and higher return on investment. It's ideal for the 24×7 operations required by today's e-services.

The V-Class is the best database server for the SAP environment, clearly demonstrating its performance over competitive products.

In technical applications, the V-Class' scalable architecture permits multiple, parallel applications to effectively use the architecture, minimizing the turn-around time for a series of simulations or design models. At the same time, applications where time-to-solution is critical may be run in a highly efficient, fully parallel environment, giving engineers and scientists opportunities to solve extremely large problems in minimal time.

Full 64-Bit Operating Environment

The V-Class' pervasive operating environment, HP-UX, supports all of HP's award-winning administration and support tools, development environments, and a huge application portfolio. The combination of these tools, the 64-bit features, and the performance of HP-UX, results in HP being the industry's leader for a sophisticated, standard UNIX operating environment.

Investment Protection for the Future

All 32- and 64-bit applications that run under HP-UX are binary-compatible with IA-64, HP and Intel's next-generation processor architecture. V-Class servers provide the highest performance today and an unsurpassed ability to upgrade to the processor technologies of the future.

For more information, contact any of our worldwide sales offices or HP Channel Partners (in the U.S. call 1 800 637 7740)

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 12 88 33 99

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 PC Products is available at www.hp.com

5968-9030E