

An Entry-Level Server with Multiple Performance Points

To address the very intense, high-volume environment of departmental and branch computing, the system design for the D-class server was made flexible enough to offer many price and performance features at its introduction and still allow new features and upgrades to be added quickly.

by **Lin A. Nease, Kirk M. Bresniker, Charles J. Zacky, Michael J. Greenside, and Alisa Sandoval**

As the computer industry continues to mature, system suppliers will continue to find more creative ways to meet growing customer expectations. The HP 9000 Series 800 D-class server, a new low-end system platform from HP, represents a radically different approach to system design than any of its predecessors (see Fig. 1).



Fig. 1. The D-class server system cabinet. The dimensions of this cabinet are 10.2 in (26 cm) wide, 23.8 in (60.4 cm) high, and 22.2 in (56.4 cm) deep.

The Series 800 D-class server comes at a time when server systems priced below U.S. \$20,000 are at a crossroads. Commodity technologies are upsizing, enterprise customers are enjoying choices of product families that offer thousands of applications, open computing and networking have blurred the distinctions between competitors' offerings, and finally, indirect marketing and integration channels can offer compelling value bundles that were once the exclusive domain of big, direct-marketing computer system suppliers. These trends have created an environment of very intense, high-volume competition for control of departmental and branch computing.

To address this environment, the system design for the D-class server had to be flexible enough to offer many price and performance features at its introduction and still allow the addition of new features and upgrades to the system quickly. The server's competitive space, being broad and heterogeneous, also demanded that the system be able to accommodate technologies originally designed for other products, including technologies from systems that cost more than the D-class server.

System Partitioning Design

In designing the D-class entry-level servers, one of the primary goals was to create a new family of servers that could be introduced with multiple performance points without any investment in new VLSI ASICs. The servers also had to be capable of supporting new advances in processors and memory and I/O subsystems with minimal system reengineering.

The server family would cover a span of performance points that had previously been covered by several classes of servers which, while they were all binary compatible with the PA-RISC architecture, had very different physical implementations. The lower-performance-point designs would be drawn from uniprocessor-only PA 7100LC-based E-class systems. The upper-performance-point designs would be drawn from the one-to-four-way multiprocessor PA 7200- based K-class systems. The physical designs of these systems varied widely in many aspects. Issues such as 5.0V versus 3.3V logic, a single system clock versus separate clocks for I/O and processors, and whether I/O devices would be located on the processor memory bus or on a separate I/O bus had to be resolved before the existing designs could be repartitioned into compatible physical and logical subsystems. Tables I and II list the key performance points for the HP 9000 E-class, K-class, and D-class servers.

Table I
Key Performance Points for HP 9000 K- and E-Class Servers

Model	Processors	Clock Speed (MHz)	Multiprocessor Configuration	Memory (M Bytes)	Cache (I/D K Bytes)	HP-PB I/O Slots	HP-HSC I/O Slots	Disk Capacity (G Bytes)
K2x0	PA 7200	120	1-to-4-Way	64 to 2048	64 to 256	4	1	3800
	PA 8000	160	1-to-4-Way	128 to 2048	256	4	1	3800
K4x0	PA 7200	120	1-to-4-Way	128 to 3840	256 to 1024	8	8	8300
	PA 8000	160 to 180	1-to-4-Way	128 to 4096	1000	5	5	8300
Ex5	PA 7100LC	48 to 96	1-Way	16 to 512	64 to 1024	4	4	144

Table II
Key Performance Points for HP 9000 D-Class Servers*

Model	Processors	Clock Speed (MHz)	Multiprocessor Configuration	Memory (M Bytes)	Cache (I/D K Bytes)	HP-HSC I/O Slots	EISA I/O Slots
D2x0	PA 7100LC	75 to 100	1-Way	32 to 512	256	4	4
	PA 7300LC	132 to 160	1-Way	32 to 1024	64 **	4	4
	PA 7200	100 to 120	1-to-2-Way	32 to 1536	256 to 1024	4	4
	PA 8000	160	1-to-2-Way	64 to 1536	512	4	4
D3x0	PA 7100LC	100	1-Way	32 to 512	256	4	7
	PA 7300LC	132 to 160	1-Way	32 to 1024	64 to 256 **	4	7
	PA 7200	100 to 120	1-to-2-Way	32 to 1536	256 to 1024	5	7
	PA 8000	160	1-to-2-Way	64 to 1536	512	5	7

* HP-PB slots = 0 and disk capacity = 5 Tbytes.

** 1M bytes of second-level cache.

Additional constraints on the design were a direct result of competitive pressures. As the presence of Industry Standard Architecture-based systems has grown in the entry-level server space, the features they offer became D-class requirements. These requirements include support for EISA (Extended Industry Standard Architecture) I/O cards and an increase in the standard warranty period to one year. Both of these requirements were new to the Series 800. Also new to the Series 800 was the desire to design a system enabled for distribution through the same type of independent distribution channels used by other server vendors. Add to these constraints the cost sensitivity of products in this price range, and we have a system that uses as many industry-standard components as possible, is extremely reliable, and is capable of being assembled by distributors, all without compromising any performance benefits of current or future PA-RISC processors.

Feature List. The first step in the process of partitioning the system was to detail all possible features that might be desired in an entry-level server. This list was compiled by pulling features from our development partners' requirements analysis and from knowledge of our competitors' systems. Once this feature list was developed, each feature was evaluated against all of our design goals (see Fig. 2). Each feature was then ranked in terms of its relative need (must, high want, want) and technical difficulty (high, medium, low). Determining the possible feature list was the first goal of the partitioning process; the list was continually updated during the entire process.

Feature	Need	Difficulty
Front panel display	Must	Low
HP 100LX Palmtop front-panel display	Want	High
Two-line LCD display	High Want	Medium
16-character display	Must	Low
Backlit display	Want	Medium
Hexadecimal status information	Must	Low
English status information	High Want	Medium
Localized language status information	Want	High
Display suspected failure causes after faults	Must	Medium
Include power-on LED indicator	Must	Low
Include disk-access LED indicator	Want	Medium
Include blinking "Boot in progress" indicator	Want	Medium
Reset switch	High Want	Medium
Reset switch with key and lock	Want	High

Fig. 2. The feature list for the D-class server's front-panel display.

Once the initial feature list was created, a small design team consisting of a mechanical engineer, an electrical engineer, a firmware engineer, a system architect, and a system manager began analyzing the list to see how each feature would affect the physical partitioning of the system. The goal of this process was to generate a fully partitioned mock-up of the physical system. Successive passes through the feature list led to successive generations of possible designs. With each generation, the list was reevaluated to determine which features could be achieved and which features could not.

Physical Partition. After the first few generations it became clear that a few critical features would drive the overall physical partitioning. The physical dimensions would be determined by the dimensions of a few key subsystems: the disk array and removable media, the integrated I/O connectors, the I/O card cage, and the power supply (see Fig. 3). All of these components were highly leveraged from existing designs, like the hot-swap disk array module developed by HP's Disk Memory Division, and the industry-standard form-factor power supply. The first major design conflict arose when we realized that these components could not be integrated in a package short enough to fit under a standard office desk, and yet narrow enough to allow two units to be racked next to each other in a standard rack. Numerous attempts to resolve these two conflicting demands only succeeded in creating a system that would violate our cost goals or require more new invention than our schedule would allow.

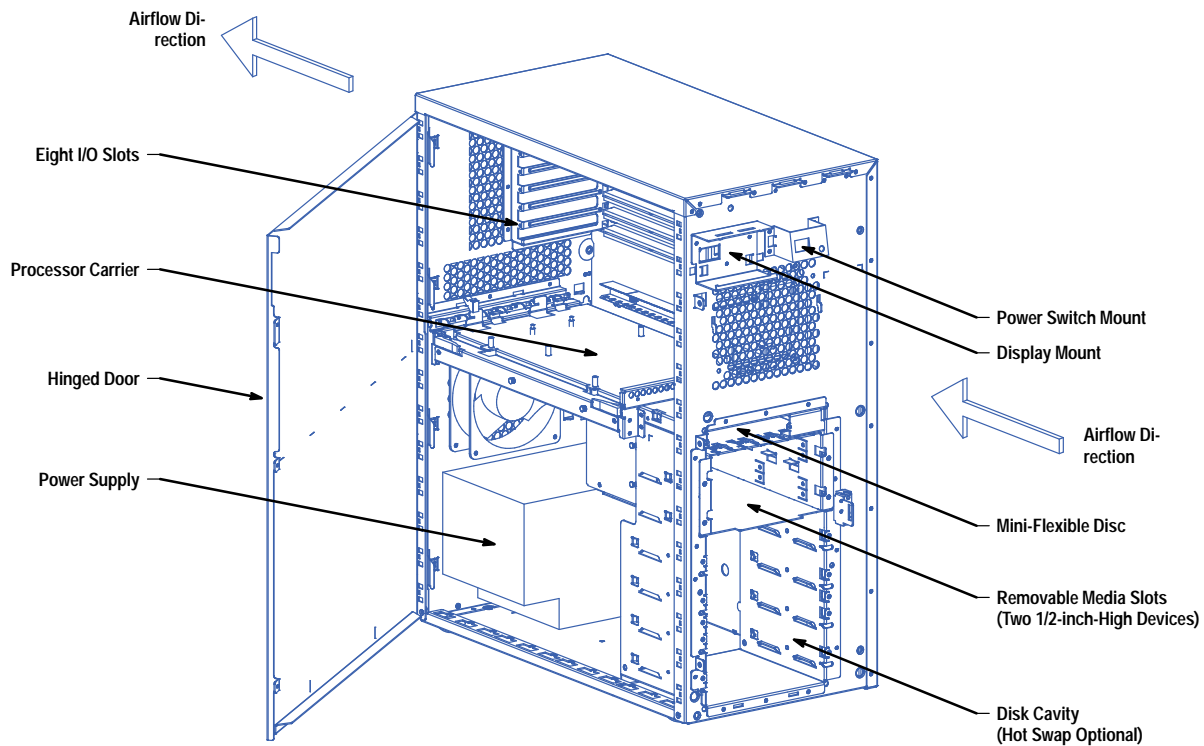


Fig. 3. The chassis for the Series 800 D-class server. The features that determined the overall size of this chassis were the disk cavity, the removable media slots, the power supply, I/O slots, and processors.

In the end, it was determined that the desk-height requirement and cost goals were more important than the rack requirement, so the package was shortened and widened to accommodate all the critical components in a package that would fit under a standard office desk. Once this was decided, the system mock-up came together quickly, and the second goal was reached. The system partitioning shown in Fig. 3 provides several benefits necessary to achieve our goals. The standard PC-type power supply helped us to achieve new lows in cost per watt of power. The division of the box into the lower core I/O and disk array volume and the upper expansion I/O slot and processor area helped to simplify the design of the forced-air cooling system since it separates the large interior volume of the box into two more manageable regions.

Printed Circuit Assemblies. The next goal in the process was to repartition and integrate the disparate design sources* into logically and physically compatible printed circuit assemblies, while maintaining all of our design constraints on cost, expandability, and design for distribution. Again, a single crucial design decision helped to quickly partition the system: the D-class would not use any high-speed, impedance-controlled connectors. This decision was made as a direct result of the K-class development process and the success of the Series 800 G/H/I-class Model 60 and 70 systems. The K-class development process showed that although high-speed impedance-controlled connectors can add excellent flexibility and expandability to midrange systems, they require a great deal of mechanical and manufacturing infrastructure.

The processor modules for the G/H/I-class Models 60 and 70 are the uniprocessor and dual-processor versions of the same board (see Table III). The same printed circuit board is loaded with either one or two processors at the time of manufacture. To increase the number of processors in a system, the entire processor module must be replaced with a new printed circuit assembly. Other systems, like the K-class servers, allow for the incremental increase of the number of processors in the system with just the addition of new processor modules. Even though the board swap is a less desirable upgrade path than an incremental upgrade, the success of the Models 60 and 70 systems led us to believe that it was quite acceptable to our customers.

Table III
Key Performance Points for the HP 9000 Series 800
G/H/I-class Model 60 and 70 Systems*

Model	Multi-processor Configuration	Memory (M Bytes)	HP-PB I/O Slots	Disk Capacity (G Bytes)
G60	1-Way	32 to 768	4	156
G70	2-Way	32 to 768	4	156
H60	1-Way	32 to 768	8	300
H70	2-Way	32 to 768	8	300
I60	1-Way	64 to 768	12	330
I70	2-Way	64 to 768	12	330

* Processor = PA 7000, clock speed = 96 MHz, Cache per CPU (Instruction/Data) = 1M byte/1M byte, and HP-HSC slots = 0.

This decision simplified repartitioning the design sources, since it meant that the high-speed processor and memory clock domains and their data paths could remain on a single printed circuit assembly, while the moderate-speed I/O domain and its data paths could cross multiple printed circuit assemblies. It was determined that both the dual I/O bus architecture of the K-class and the single I/O bus of the E-class would be supported in the system. To do this, the connector technology used in the D-class is modular, allowing the designs to load only those portions of the connector that are supported. This lowers both the material and assembly costs. To further lower the material cost, the PA 7100LC-based processor and memory module is fabricated on a smaller printed circuit board than the PA 7200-based processor module. The difference in the size of the modules is accommodated by attaching them to a sheet-metal carrier that adapts the modules to a common set of card guides. Not only is the sheet metal cheaper than the corresponding printed circuit material would have been, but it is also stronger and easier to insert and remove.

A secondary benefit of this strategy is that it allows new investment to be made as needed. Historically, I/O subsystems and technology are much longer-lived than processor and memory technologies. The partitioning strategy we used helped to decouple the I/O subsystem from the processor and memory. As long as they remain consistent with the defined interface, processor modules are free to exploit any technology or adapt any design desirable. This also enables D-class servers to excel in meeting a new and growing requirement—design for reuse. A customer is able to upgrade through many performance points simply by changing processor modules. As some countries are investigating forcing manufacturers to accept and recycle old equipment, keeping the return stream as small as possible is highly desirable.

* The design sources were parts from E-class and K-class servers and the J-class workstations that were combined to form the D-class servers.

Once the printed circuit assembly board outlines were complete, the process of adapting the various design sources to the new partitioning was time-consuming, but relatively straightforward. Fig. 4 shows the the various design sources that were pulled together to form the PA 7200-based processor module. As portions of designs were merged, altered, and recombined, the possibility of transcription errors grew. The original designs were executed by three different labs and many different design teams. All designs were fully functional as designed, but we were extending designs as well as integrating them. In an effort to minimize the possibility of errors being introduced during the adaptation process, the schematic interconnect list was extracted and translated into a simulation model. This model was then added to the models used to verify the original designs to ensure that no new errors had been introduced.

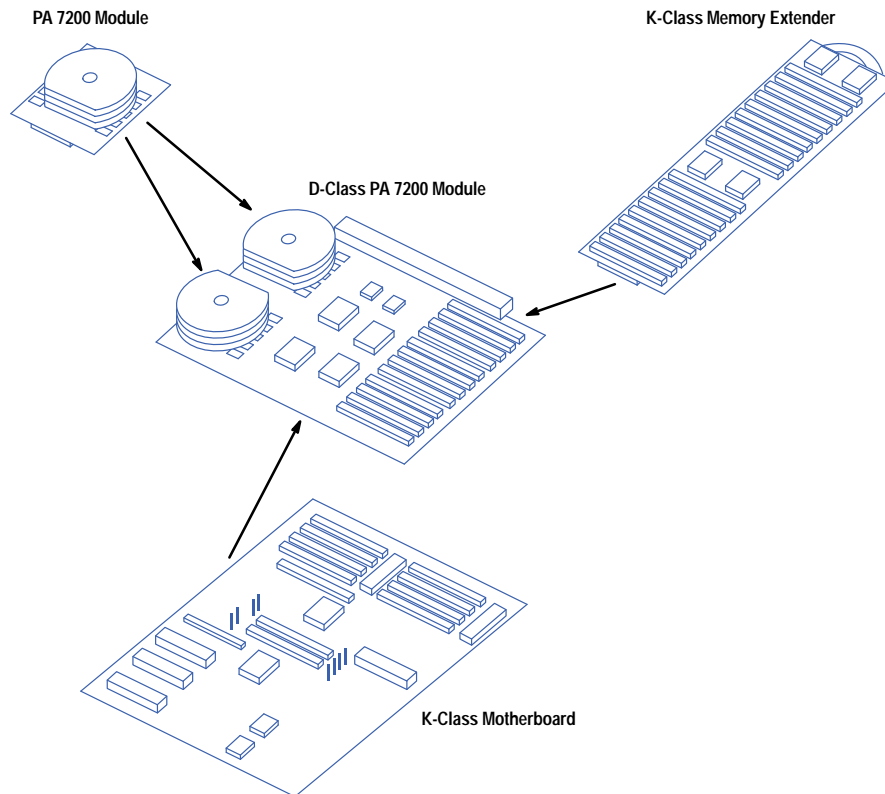


Fig. 4. The various design sources that were pulled together to form the PA 7200 D-class module.

System Partitioning and Firmware Design

Because of the partitioning scheme used for the D-class entry-level servers, the firmware design was a critical factor in achieving the overall program objective of low cost. The firmware design addressed cost issues in support of the manufacturing process, field support, and the upgrade strategy. In addition, although the underlying hardware is dramatically different depending upon which processor module is installed in the system, from the customer's perspective, the external behavior of each performance point should be the same. For the firmware design team, that meant that regardless of the underlying hardware, the entire D-class had to have the look and feel of a single product.

D-Class Subsystems. From a firmware perspective, the D-class is partitioned into two subsystems: the system board and the processor modules (see Fig. 5). The system board contains all of the I/O residing on or hanging off the HSC (high-speed system connect) bus. This includes optional I/O modules that plug into the HP-HSC slots, such as the fast-wide SCSI and graphics cards. It includes core I/O built into the system board, which provides serial interfaces, single-ended SCSI, Ethernet LAN, a parallel interface, a mouse and keyboard interface, and a flexible-disk interface. The EISA bus, which is connected at one end to the HSC bus, is also found on the system board. The Access Port/MUX card,* which contains its own HSC-to-HP-PB I/O bus converter, also plugs into an HSC slot. In addition to these I/O buses and devices, there is an EEPROM and two hardware dependent registers that hold I/O configuration information. This nonvolatile memory and the configuration registers are critical to the partitioning and upgrade strategy for the D-class server.

The core of each processor module houses the CPU, instruction and data caches, and memory subsystems. Also on the processor board is an EEPROM and two more hardware dependent registers. The PA 7100LC uses the HSC as its native bus, so its connection to the system board is relatively straightforward. However, the PA 7200 requires a bus converter between

* HP Access Port is a tool for providing remote support for HP servers.

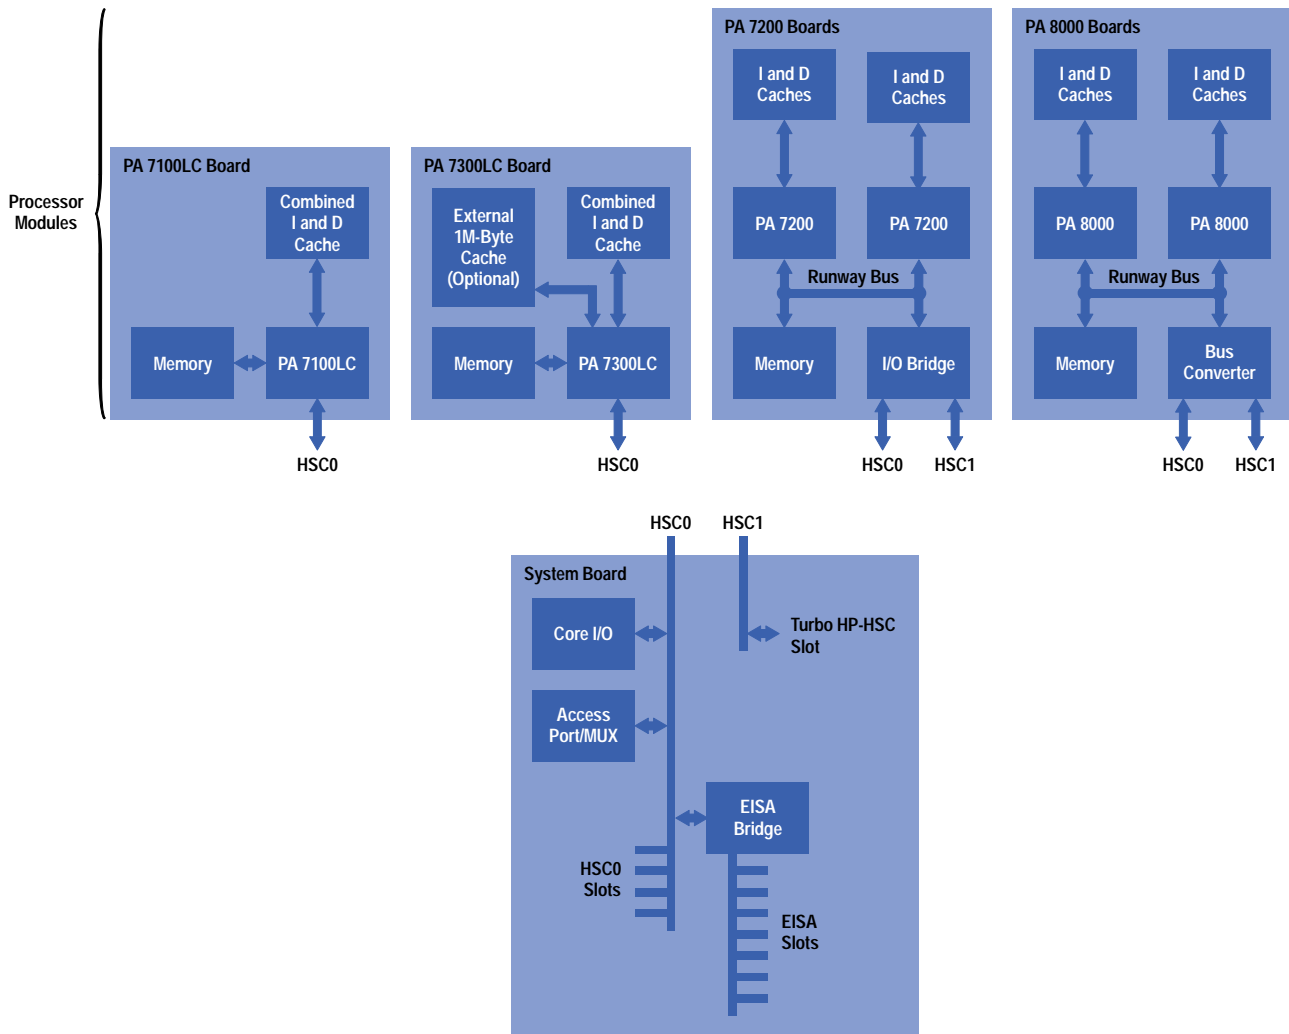


Fig. 5. A firmware perspective of the D-class subsystems.

its native bus and the HP-HSC bus. Thus, to have one system board common between the two processor modules, the PA 7200 processor board was burdened with carrying the bus converter circuitry. One other significant difference exists between the two processor modules: scratch RAM.** The inclusion of 32K bytes of static RAM on the PA 7200 module meant that system variables and a stack could be set up very early in the boot process. The lack of this scratch RAM on the PA 7100LC limited the amount of code that could be common between the two platforms.

Consistent Look and Feel. The goal of having the same look and feel for the entire product line could be met by having one common code base for all performance points, but because of the significant hardware differences mentioned above, this was not possible. The primary differences lie in the processors themselves. There is no commonality in the control and status registers of the two processors, caches are accessed differently, and the memory subsystems are too different to share code. These differences, along with other hardware incompatibilities, meant that each processor module needed to have its own separate and distinct code base. However, because the primary differences between the two PA 7100LC processor versions (75 MHz or 100 MHz) and the two PA 7200 processor modules are processor speed and cache size, all performance points that use a common CPU can be supported by one common code base.

The three areas needed to give the product line a consistent look and feel included a common feature set, similar strategies for handling and reporting errors, and a common user interface. To ensure consistency in this regard, one engineer was given responsibility for the same functional area of each platform. For example, the engineer who worked on memory code for the PA 7100LC also had responsibility for the memory code on the PA 7200 platform. Taking advantage of this synergy paid off especially well in the design and implementation of the user interface, where differences between platforms could easily lead to confusion.

System Configuration. Partitioning the EEPROMs between the processor board and the system board is a key enabler of the upgrade strategy. Since an upgrade consists of replacing the processor module, I/O configuration information must remain with the system board. The EISA configuration, graphics monitor types, and LAN MAC address are stored on the system

** The PA 7300LC and the PA 8000 also include scratch RAM.

board. Additional control information is used to check for consistency between the two EEPROMs. The firmware expects the format of the system board EEPROM to be the same, regardless of which processor module is installed. With all I/O configuration information and control variables in the same location and sharing the same set of values, processor modules can be freely swapped without changing the I/O configuration.

Dynamic configuration of the system is used to support the upgrade strategy, the manufacturing process, and field support. When a D-class server is powered on, state variables in the processor module's and system board's nonvolatile memories are tested for a value that indicates whether or not they have been initialized and configured. If they fail this test (which is always the case for initial turn-on during the manufacturing process) the system's hardware configuration is analyzed and the corresponding state and control variables are set. Much of this information is available via the hardware dependent registers located on each board. The processor frequency, system board type, and other details concerning the I/O configuration can be read from these registers.

The state variables, which are set as a result of examining the hardware configuration, include the system's model identifier (e.g., HP9000/S811/D310), hardware version (HVERSION), and paths to the boot and console devices. The boot path can be either the built-in single-ended SCSI device or a hot-swappable fast-wide SCSI device. The firmware checks for the presence of a hot-swappable device which, if present, becomes the default boot path. Otherwise a single-ended SCSI device is configured as the default boot path. The actual hardware configuration is also examined to select an appropriate console path. The default console path can be either the built-in serial port, the HSC Bus Access Port/MUX card, or a graphics console. Depending upon the presence of these devices and their configurations (e.g., a graphics device must also have a keyboard attached), a console path is selected according to rules worked out in cooperation with the manufacturing and support organizations.

The same sequence of events occurs when upgrading or replacing a processor module. In this case, the system board is already initialized and only the processor module requires configuration. On every boot, information such as the model identifier is checked against the actual hardware configuration and any mismatch will invoke the appropriate configuration actions. Likewise, because some information is kept redundantly between the processor module and the system board, they can be checked for a mismatch. This redundancy means that the system board can also be replaced in the field with a minimal amount of manual reconfiguration. Because a D-class server can consist of any combination of two system boards and several different processor modules, and because further enhancements will double the number of processor modules and include two new CPUs, dynamic configuration has obviated the expense of developing external configuration tools, reduced the complexity of the manufacturing process, and simplified field repairs and upgrades.

System Packaging

Mechanical packaging is one of the key variables in maintaining a competitive edge in the server market. The challenges involved in the system package design for the HP 9000 D-class server included industrial design, manufacturability, EMI containment, thermal cooling, and acoustics, while having the design focus on low cost.

The D-class low-cost model was based on the high-volume personal computer market. However, unlike personal computers, server products must support multiple configurations with an easy upgrade path and high availability. This meant that the D-class package design had to be a highly versatile, vertical tower with the ability to be rack-mounted in a standard EIA rack (see Fig. 6). It allows for multiple processors and power supplies, and can support up to eight I/O slots for EISA and GSC cards. The design also supports up to five hot-swappable or two single-ended disk drives and two single-ended removable media devices with one IDE (Integrated Device Electronics) mini-flexible disk (Fig. 7). This diversification provides the entry-level customer with a wide range of configurations at various price/performance points.

Manufacturing and Field Engineering Support. Concurrent engineering was a key contributor to the design for assembly (DFA) and design for manufacturing (DFM) successes of the D-class server. Since we are very customer focused, we take the disassembly and repair of the unit just as seriously as the manufacturability of the product. The D-class mechanical team worked closely with key partners throughout the program to ensure the following assembly and manufacturing features:

- A single-build orientation (common assemblies)
- Multiple snap-in features
- Slotted T-15 Torx fasteners (Torx fasteners are used for HP manufacturing, and the slot is for customers and field engineers.)
- System board that slides into the chassis
- Quick access to all components
- Manufacturing line for high-volume production.

EMI. Design for EMI containment was a considerable challenge for the D-class server program. The package goal was to contain clock rates up to 200 MHz. This required a robust system design and two new designs for the EISA bulkheads and core I/O gaskets.

The system EMI design is based on a riveted sheet-metal chassis using a slot-and-tab methodology for optimum manufacturability. A cosmetic outer cover with a hinged door completes the EMI structure. EMI is contained using continuous seams and EMI gaskets with small hole patterns for airflow.

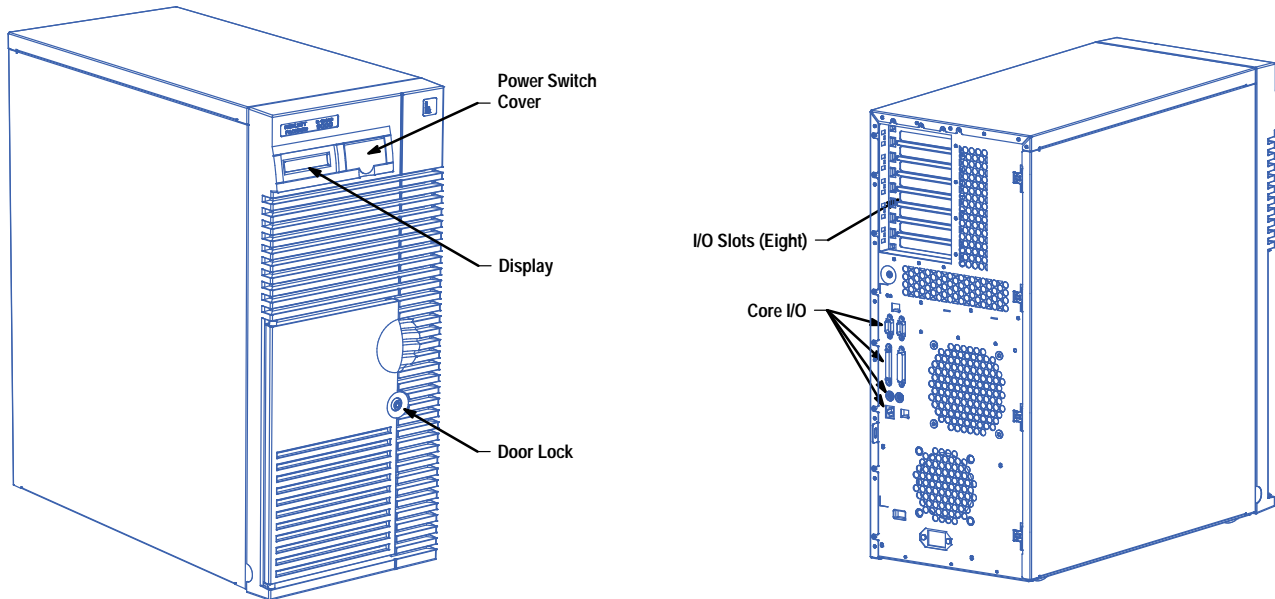


Fig. 6. A front and back view of the D-class server.

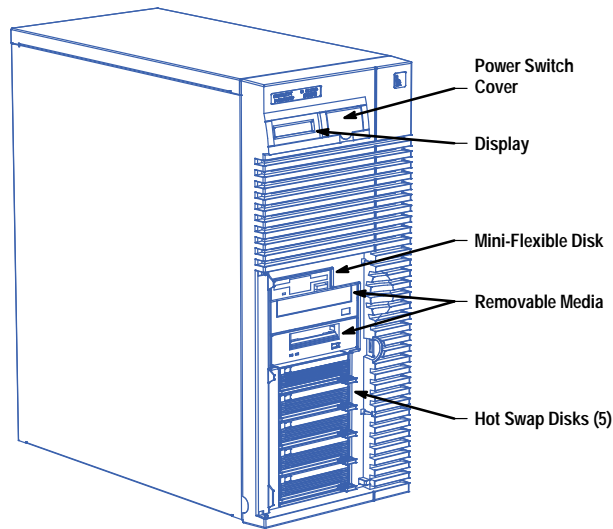


Fig. 7. The D-class cabinet with the door removed.

The EISA bulkhead gaskets required a new EMI design. The new design is a slotted pyramid that forms a lateral spring element with a low deflection force but a high contact force (see Fig. 8). The new design for the I/O gasket includes a foam core wrapped with a nickel-over-copper fabric, which provides a 360-degree contact around each connector (see Fig. 9). These new designs produced excellent results.

Thermal Design and Airflow. The thermal design for the D-class server also had some interesting challenges. The design strategy had to encompass multiple configurations and multiple processor chips and boards. Some of these options were in development, but most were future plans. A thermal analysis program, Flotherm, was used to develop the thermal solution for the system.

The Flotherm models and tests resulted in the package being separated into two main compartments. The top half, which includes the I/O and the processors (see Fig. 3), is cooled by a 12-mm tubeaxial fan. The processor chips are located side-by-side directly behind the front fan, giving an approximate air velocity to the processors of about 2.5 m/s. Heat sinks are used for processor chips that consume under 25 watts. For chips over 25 watts, a fan mounted in a spiral heat sink is used.

The bottom half of the package includes the peripheral bay and power supply and components. It is cooled using a 120-mm tubeaxial fan. However, when the hot-swap disks are in use, a separate cooling system is installed. The hot-swap bay is a

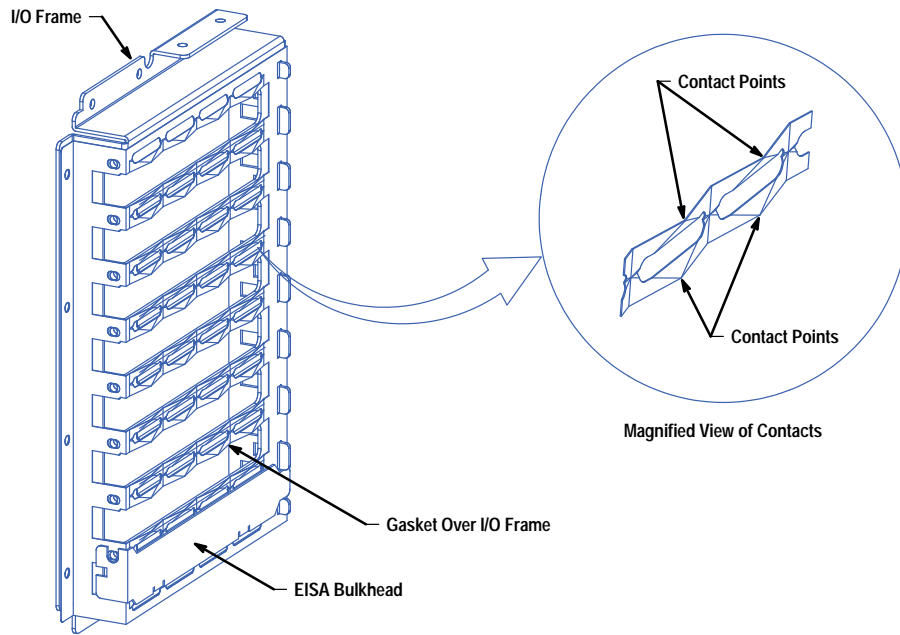


Fig. 8. The EISA bulkhead gaskets for the D-class server.

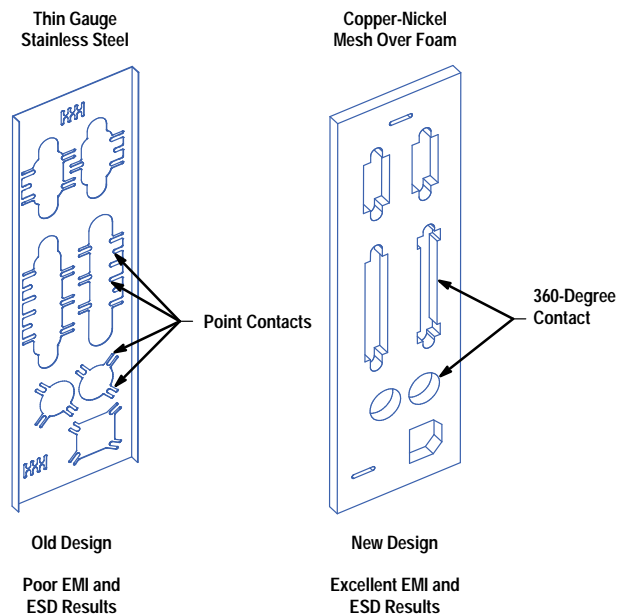


Fig. 9. A comparison between the new and old core I/O gasket designs. The new design provides a 360-degree contact around each connector.

sealed subsystem that uses a small blower to pull air through the disks. Any disk can be pulled out and the airflow to the other disks remains relatively unchanged. The power supply has its own 92-mm fan.

Acoustics. The acoustical goal for the D-class server was designed to be 5.4 bels at the low end, which was the same as for earlier server products. This package has a higher power density than previous products, more versatility, higher-speed discs, and an off-the-shelf power supply rather than a custom one. The fan in the power supply ended up being the loudest component of the system. Still, the system came in at 5.4 bels at the low end by custom tuning sheet-metal parts, baffles, and fan speeds.

High Availability and Ease of Use

Hot-Plug Internal Disks. An important feature that the D-class server has brought to the Series 800 product line is hot-pluggable internal disk drives. While commodity servers had once provided forms of internal hot-pluggable disks, these solutions were deemed too likely to cause data corruption for use in Series 800 systems. For example, the commodity

hot-plug solutions evaluated for use in Series 800 systems had the issue of “windows of vulnerability” in which sliding contacts from a swapping disk on a SCSI bus could cause data bits to actually change after parity had already been driven, causing undetected data errors. With a low probability of corruption, this approach may have made sense for a workgroup file or print server. However, Series 800 servers are expected to operate in mission-critical environments. Therefore, a more robust approach to hot-swapping disk drives had to be developed for these products.

The approach used for the D-class server was to provide a hot-swap solution using logical volume management (LVM)* and mirroring middleware facilities, and to offer a disk-drive carrier common to the standalone enclosure disk carrier being separately developed for the rest of the Series 800 product line. The common carrier approach would allow the field to learn only one solution and guarantee a higher volume of parts. In addition, solutions for the data corruption problem, the use of sequenced SCSI resets, and an automated swap script could be shared by both the enclosure team and the D-class server team.

Mission-Critical ServiceGuard and EISA I/O. HP’s MC/ServiceGuard product (portions of which were previously called Switchover/UX) has been an unofficial standard in the application-server industry for several years now. This middleware product allows Series 800 systems to operate as standby servers for one another, transferring mission-critical workloads from a failed system to its standby. This feature requires that a system and its standby host transactions on the same mass storage buses, enabling the standby system to have access to all of the primary system’s data. Multiple hosting (known as multi-initiator) on mass-storage interconnects requires significant design attention. In addition, the LANs that allow these systems to communicate with one another must offer special capabilities in the areas of error handling and reporting.

A significant challenge for the D-class design was merging the workstation, PC, and I/O infrastructures with the Series 800 infrastructure. This requires supporting MC/ServiceGuard capabilities, higher slot counts, more extensive errorhandling, and a remote support infrastructure. Before the D-class server, all Series 800 systems that supported MC/ServiceGuard were HP Precision Bus (HP-PB) based.

The D-class server’s use of EISA and HP-HSC (with the EISA form factor) required the team to implement, debug and verify MC/ServiceGuard functionality on these new I/O buses. In the process, various I/O implementations had to be modified, such as special EISA bus operating modes, SCSI adapter functionality, periodic EISA cleanups, guaranteed arbitration for core LAN, queueing transactions on HP-HSC, extra buffering of data signals, and slot configurations. These modifications made the HP-HSC and EISA I/O infrastructure more robust in a highly available departmental branch server.

In addition, HP’s Access Port product, which is used for remote support of HP servers, only existed on the HP-PB. Without HP-PB slots, the D-class server would have had to either forego remote support or develop a new card. The answer was not to develop a new card, but to leverage existing logic by supplying a buried HP-PB on the new Access Port card. To the user, the D-class server’s Access Port is an HP-HSC card. However, to the operating system and response center engineer, the D-class server’s Access Port looks like the familiar, compatible HP-PB card found on all the other server products.

Pushbutton Graceful Power Shutdown

The D-class server is the first Series 800 system to offer pushbutton graceful power shutdown. Basically, when a D-class system is up and running, the power button is equivalent to the command `reboot -h`, which causes the system to synchronize its buffer cache and gracefully shut down. This feature is most useful in a branch office or department where the server is minimally managed by local personnel. Single-user HP workstations had introduced this feature to PA-RISC systems.

Built-in Remote Management

With an emphasis on remote server management, the D-class server team decided to offer the same robust, worldwide-usable internal modem as the K-series products. This modem offers support for transfer rates well beyond those used by today’s HP response centers and is integrated with the remote support assembly in D-class systems. The product also offers a special serial port for controlling optional uninterruptible power supplies (UPS), as well as pinout definitions for future direct control of internal power-supply signals.

The D-class server team also accommodated “consoleless” systems, whereby a D-class server can be completely managed remotely without a local console at all. In addition, graphics console customers can still take advantage of remote console mirroring (formerly reserved strictly for RS-232 consoles) by merely flipping a switch on the product.

* LVM is software that allows the number of file systems to be relatively independent of the number of physical devices. One file system can be spread over many devices, or one device can have many file systems.

Conclusion

The system partitioning design for the first release of the D-class servers helped to achieve all of our introduction goals. We were able to introduce both PA 7100LC-based and PA 7200-based processor modules, integrate the industry standard EISA I/O bus into the Series 800 hardware for the first time, and achieve our cost and schedule goals without any investment in new VLSI ASICs.

In the end, the D-class design had leveraged from all current entry-level and midrange Series 800 servers and many Series 700 workstations. Because of the care taken during the adaptation process, performance enhancements made to the original design sources were made available in the latest D-class module quickly and with very little investment. As an example, only two weeks after a larger-cache, higher-speed PA 7200 K-series processor module was released, the corresponding D-class PA 7200 processor module had been modified and released for prototype. This module provides four times the cache and a 20% increase in frequency over the initial D-class PA 7200 module.

Table IV summarizes the leverage sources for the various subsystems that make up the D-class servers.

Table IV
Leverage Sources for the D-Class Subsystems

D-Class Subsystem	Leverage Source
EISA I/O	HP 9000 Series 700 Workstations
HP-HSC I/O	K-class server and J-class and HP 9000 Series 700 Workstations
Clocking	E-class and K-class Servers
PA 7100LC Processor Modules	E-class Server
PA 7200 and PA 8000 Processor Modules	K-class Server
PA 7300LC Processor Modules	C-class Server
Power Supply	Industry-Standard Suppliers
Hot Swap Disk Array	HP's Disk Memory Division

[▶ Go to Next Article](#)

[▶ Go to Journal Home Page](#)