



AlphaServer DS25 System **Technical Summary**



Contents

HP AlphaServer DS25 Systems 1

System Overview 1

Features and Benefits 1

Third-Generation Alpha Chip 2

Chip Operation 2

Alpha 21264C Features 2

Processor Module 2

Processor Configuration Rules 2

Architecture 3

System Block Diagram 3

Motherboard 4

Component and Connector Locations 4

Memory 5

Memory Architecture 5

Memory Options 6

Memory Arrays 6

Memory Configuration Rules 6

Memory Interleaving 6

System I/O 7

I/O Implementation 7

Bandwidth 7

Block Diagram of I/O Control 8

System Control 9

Control Panel 9

LEDs 9

Environment 9

Power 9

Storage 9

Fibre Channel 9

RAID (Redundant Array of Independent Disks) 9

Server Management 10

Operational Management 10

Platform Management 10

Error Reporting 10

Security 10

Reliability, Availability, and Maintainability 10

Processor Features 10

Memory Features 10

I/O Features 10

System Features 10

Installation and Maintenance 11

Clustering 11

PCI to Memory Channel Interconnect 11

Operating System Support 11

Performance 11

Sources of Performance Information 11

Service and Support 12

Hardware Warranty 12

Software Warranty 12

AlphaServer DS25 System Diagram 13

System Features at a Glance 14

Physical Characteristics 15

HP AlphaServer DS25 Systems

The *HP AlphaServer DS25* system delivers high performance for business, technical, and scientific applications at an affordable price. As a high-capacity database server, high-performance application server, Network File System (NFS) server, or Internet server, the *AlphaServer DS25* delivers exceptional value and investment protection.

System Overview

The *AlphaServer DS25* systems support two Alpha 21264C 1000 MHz CPU chips; up to 16 Gbytes of memory can be installed. Second-level cache is 8 Mbytes per processor. The switch-based system interconnect exploits the full potential of the Alpha chip.

The system is available as a pedestal or as a rackmount system. The components in both are identical; the rackmount variant is simply rotated to its side. The rackmount system requires 5U space and can be mounted in a 67- or 79-inch M-series cabinet along with additional disks. In the largest cabinet up to six systems and two StorageWorks shelves can be installed.

The DS25 system is available with one or two processors. The system accommodates six hard disks or tapes and a CD-R drive. The system supports up to 876 GB of internal storage using 146 GB disks.

There are six PCI slots—all support 64-bit options. Four slots run at 66 MHz, and two run at 33 MHz. Integrated on the system board are dual Ethernet controllers (10/100 port and a 10/100/1000 port), dual Ultra3 SCSI controllers, and an IDE controller, two serial ports, one parallel port, and the keyboard and mouse.

Three operating systems are supported: *Tru64 UNIX*, *OpenVMS*, and Linux. Customers can order the system that meets their present application needs.

HP AlphaServer products use the 64-bit Alpha RISC architecture that supports multiple operating systems: *Tru64 UNIX*, *OpenVMS*, and Linux.

For more information on *AlphaServer DS25* systems, see <http://h18002.www1.hp.com/alphaserver/ds25/index.html>

Features and Benefits

▪ Leadership 64-Bit Architecture

The Alpha 64-bit architecture was introduced with the Alpha 21064 chip in 1992 and now the 21264C builds upon that proven architecture.

▪ Performance

The Alpha 21264C chip when combined with a switch-based interconnect demonstrates its full potential. This switch-based system provides a memory bandwidth of up to 8 Gbytes/sec (peak) using two 256-bit memory buses running at 125 MHz.

▪ Reliability and Availability

The *AlphaServer DS25* uses the latest technologies to achieve redundancy, error correction, and fault management. The balance between simple error detection and error correction provides high availability at low cost. Hot-swappable disks, fans, and power supplies allow repairs to be done without shutting the system down.

▪ DS25 Workstation

The pedestal package can also be ordered as a workstation.

Third-Generation Alpha Chip

The third generation of the Alpha microprocessor, the Alpha 21264C, is a superscalar, superpipelined implementation of the Alpha architecture. The first offering of this chip was known as EV6; these systems are offered with the EV68 (.18 micron) chip. Each chip has over 15.2 million transistors.

Designed for performance, the Alpha 21264C achieves this goal by carefully studied and simulated architectural and circuit analysis. The 21264C memory system also enables the high performance levels. On-chip and off-chip caches provide for very low latency data access, which allows for very high bandwidth data access. (The size of the off-chip cache is 8 Mbytes per processor running at 250 MHz DDR.)

Internal to each chip is a 64-Kbyte instruction cache (I-cache) and a 64-Kbyte data cache (D-cache).

- **I-cache.** 64 Kbytes, two-way set-associative, virtually addressed cache with 64-byte blocks
- **D-cache.** 64 Kbytes, two-way set-associative, virtually indexed, physically tagged, writeback cache with 64-byte blocks

Chip Operation

Several key design choices were made in the chip architecture to maximize performance: Four instructions are fetched each cycle, and then the handling of those instructions boosts the speed of execution. Register renaming assigns a unique storage location with each write reference to a register, avoiding register dependencies that can be a potential bottleneck to processor performance.

Another design feature, out-of-order execution, permits instructions to execute in an order different from the order that the instructions are fetched. In effect, instructions execute as soon as possible. This allows for faster execution since critical path computations are started and completed as soon as possible.

In addition, the Alpha 21264C employs speculative execution to maximize performance. It speculatively fetches and executes instructions even though it may not know immediately whether the instruction will be on the final execution path. This is particularly useful, for instance, when the 21264C predicts branch directions and speculatively executes down the predicted path. The sophisticated branch prediction in the 21264C coupled with the speculative and dynamic execution extracts the most instruction parallelism from applications.

For more information about the chip, see:

<http://h18002.www1.hp.com/alphaserver/download/ev6chip.pdf>

Alpha 21264C Features

- Out-of-order instruction execution
- 64 Kbyte on-chip data and instruction caches
- Improved branch prediction through intuitive execution
- Register renaming
- Increased bandwidth for high-speed access to second-level cache and system memory
- Motion video instructions
- Square root and divide instructions
- All instructions are 32 bits long and have a regular instruction format
- Floating-point unit, supports DIGITAL and IEEE floating-point data types
- 80 integer registers, 64 bits wide
- 72 floating-point registers, 64 bits wide

Processor Module

An *AlphaServer* DS25 can have one or two CPU modules.

In addition to the Alpha 21264C chip, each CPU module has an 8-Mbyte second-level cache and a DC-to-DC converter with heatsink that provides the required voltage to the Alpha chip. Power-up diagnostics are in a flash ROM on the module.

Cooling of the microprocessor chip is provided by a fan on the chip heatsink.

Processor Configuration Rules

The first CPU module is installed in CPU slot 0. The second CPU is installed in CPU slot 1.

Architecture

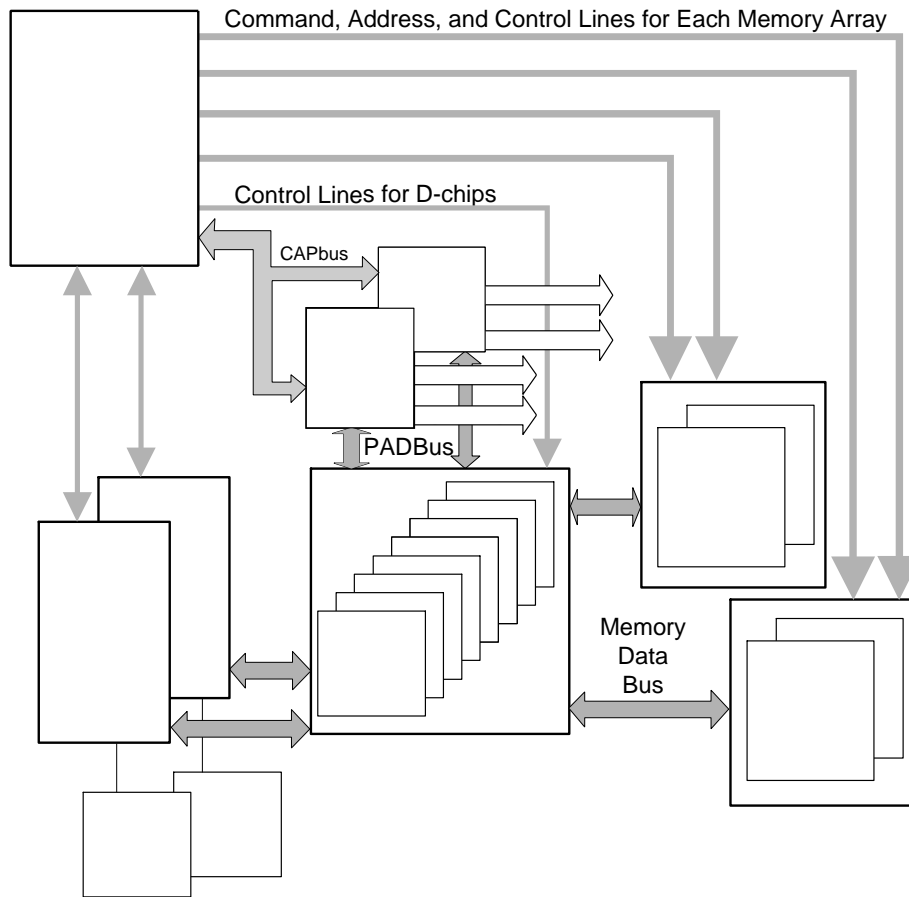
This system is designed to maximize the potential of the Alpha 21264C chip. The traditional bus interconnect has been replaced by a switch-based interconnect system. With a bus design, the processors, memory, and I/O modules must share the bus. As the number of bus users increases, the transactions interfere with one another, increasing latency and decreasing aggregate bandwidth. However, with a switch-based system there is no degradation in performance as the number of CPUs, memory, and I/O users increase. Although the users increase, the speed is maintained.

With a switch-based, or point-to-point interconnect, the performance remains constant, even though the number of transactions multiplies. The switched system interconnect uses a set of complex chips that route the traffic over multiple paths. The chipset consists of one C-chip, eight D-chips, and two P-chips.

- **C-chip.** Provides the interface from the CPUs and main memory. Although there is only one in the system, it allows each CPU to do transactions simultaneously.
- **D-chips.** Provide the data path for the CPUs, main memory, and I/O.
- **P-chips.** Provide the interface to four independent 64-bit PCI buses.

The DS25 system supports up to two CPUs and up to 16 Gbytes memory in 16 DIMM slots. There are four memory arrays, each with four slots.

Two 256-bit memory buses support the four memory arrays, yielding an 8 Gbyte/sec system bandwidth. Transactions are ECC protected. Upon the receipt of data, the receiver checks for data integrity and corrects any errors. Commands and addresses are parity protected.



MR0406

System Block Diagram

Motherboard

The interconnect switch is implemented on the motherboard by the chipset consisting of one C-chip, two P-chips, and eight D-chips. This complex chipset provides the data and address path between the CPUs, memory, and the I/O subsystem.

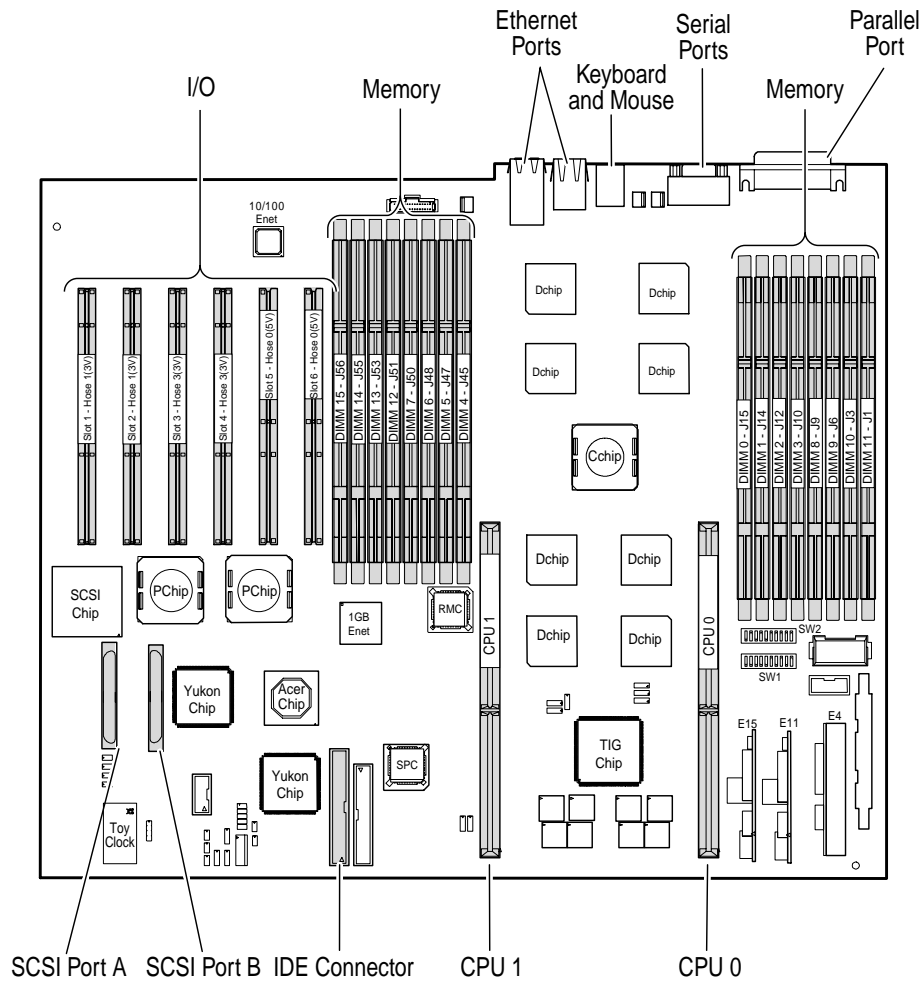
A flash ROM holds the console code and the NVRAM data. The remote management console (RMC) is implemented by the RMC processor and is accessed through the COM1 port.

Two SCSI controllers, an IDE controller, and two Ethernet controllers (one 10/100 port and one 10/100/1000 port) are integrated into the motherboard.

Connectors are provided for the following:

- 2 CPU modules
- 16 DIMM memory modules
- 6 PCI devices*
- 2 Ethernet ports
- 1 parallel port
- 2 serial ports
- Keyboard and mouse ports

* In earlier (D*-57AAA-**) models, the slot located next to the memory DIMMs can only take half-length options.



MR0293A

Component and Connector Locations

Memory

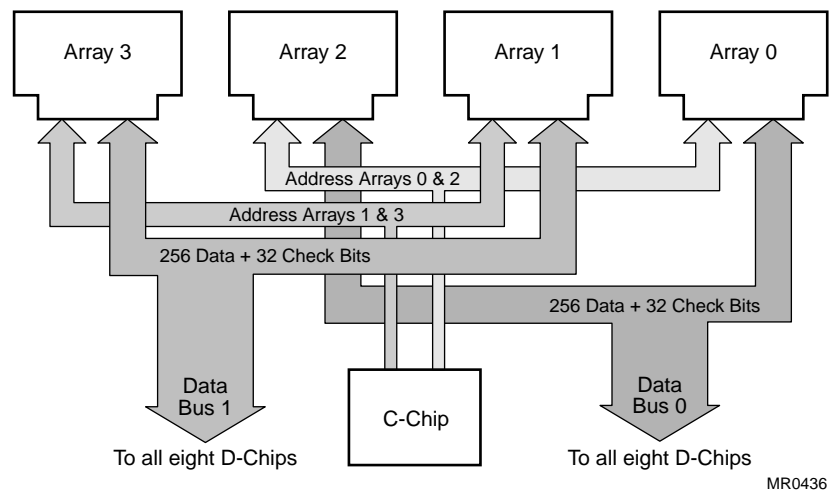
Memory throughput in this system is maximized by the following features:

- Two 256-bit wide memory data buses
- Very low memory latency (120 ns) and high bandwidth with 8 ns clock
- ECC memory

The switch interconnect can move a large amount of data over two independent memory data buses. Each data bus is 256 bits wide (32 bytes). The memory bus speed is 125 MHz. This yields 4 GB/sec bandwidth per bus ($32 \times 125 \text{ MHz} = 4 \text{ B/sec}$). The maximum bandwidth is 8 GB/sec (4×2).

The design challenge was to maximize the capabilities of the two wide data buses. Distributing the 256 data bits equally over two memory arrays was one solution: simultaneously, in a read operation, 128 bits come from one array and the other 128 bits come from the other array to make one 256-bit read. Another 256-bit read operation can occur at the same time on the other independent data bus.

In addition, two address buses allow overlapping/pipelined accesses to maximize use of each data bus. When all arrays are identical (same size and speed), the memory is interleaved; that is, sequential blocks of memory are distributed across all four arrays.



Memory Architecture

Memory Options


Each memory array has four slots that accept 200-pin industry-standard DIMMs with PECL clocks. The DIMMs are synchronous DRAMs. System memory can be from 512 Mbytes (4 x 128 MB DIMMs in one array) to a maximum of 16 Gbytes (all four arrays filled with 1 Gbyte DIMMs).

Memory options (four DIMMs) are available in the following sizes:

- 512 Mbytes (128 MB DIMMs)
- 1 Gbyte (256 MB DIMMs)
- 2 Gbytes (512 MB DIMMs)
- 4 Gbytes (1 GB DIMMs)

Memory Arrays

The illustration below shows a portion of the motherboard. Although each array must be populated with DIMMs of the same size and speed, the other arrays can be composed of DIMMs of a different size and speed.



Memory Configuration Rules

- Four DIMMs must be installed in array 0.
- Arrays are populated in numerical order.
- All DIMMs in an array must be the same size and speed.

Memory Interleaving

With one memory option (4 DIMMs) installed, memory operation interleaving will not occur. With two memory options (8 DIMMs), memory read-write operations are two-way interleaved. Interleaved operations reduce the average latency and increase the memory throughput over non-interleaved operations. With four memory options (16 DIMMs) installed, memory read-write operations are four-way interleaved, maximizing memory throughput.

System I/O

The industry-standard PCI bus is the number one choice for high-performance I/O options, such as disk storage and high-performance video applications.

I/O Implementation

The PCI bus implementation has the following characteristics:

- Fully compliant with the PCI Version 2.1 Specification
- Supports three address spaces: PCI I/O, PCI memory, and PCI configuration space
- Supports byte/word, tri-byte, and longword operations
- Exists in non-cached address space only

Four PCI buses (hoses) are implemented. One is dedicated to support of onboard controllers (Hose 2). The other three each support two option slots that take either 64-bit or 32-bit options. Three of the buses run at 66 MHz, and one bus runs at 33 MHz. Each slot is described in detail below.

Bandwidth

The total PCI throughput is 1.85 Gbytes/second.

PCI Slot Summary

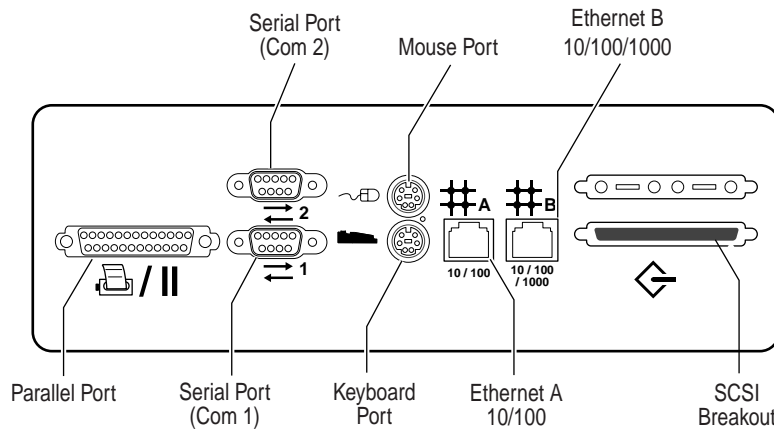
PCI Slot	PCI Bus		PCI Hot-Plug *	PCI Hose Number
	Speed (MHz)	Signal Voltage		
1	66	3.3	No	1
2	66	3.3	No	1
3	66	3.3	Yes	3
4	66	3.3	Yes	3
5	33	5.0	Yes	0
6**	33	5.0	Yes	0

* No operating system support.

** Supports only half-length options in earlier (D*-57AAA-**) models.

I/O Ports

At the rear of the system are connectors offering access to two serial communication ports, one parallel port, two Ethernet ports, one SCSI port, and ports for the keyboard and mouse. The COM1 serial port is used for the system console and for the remote management console.

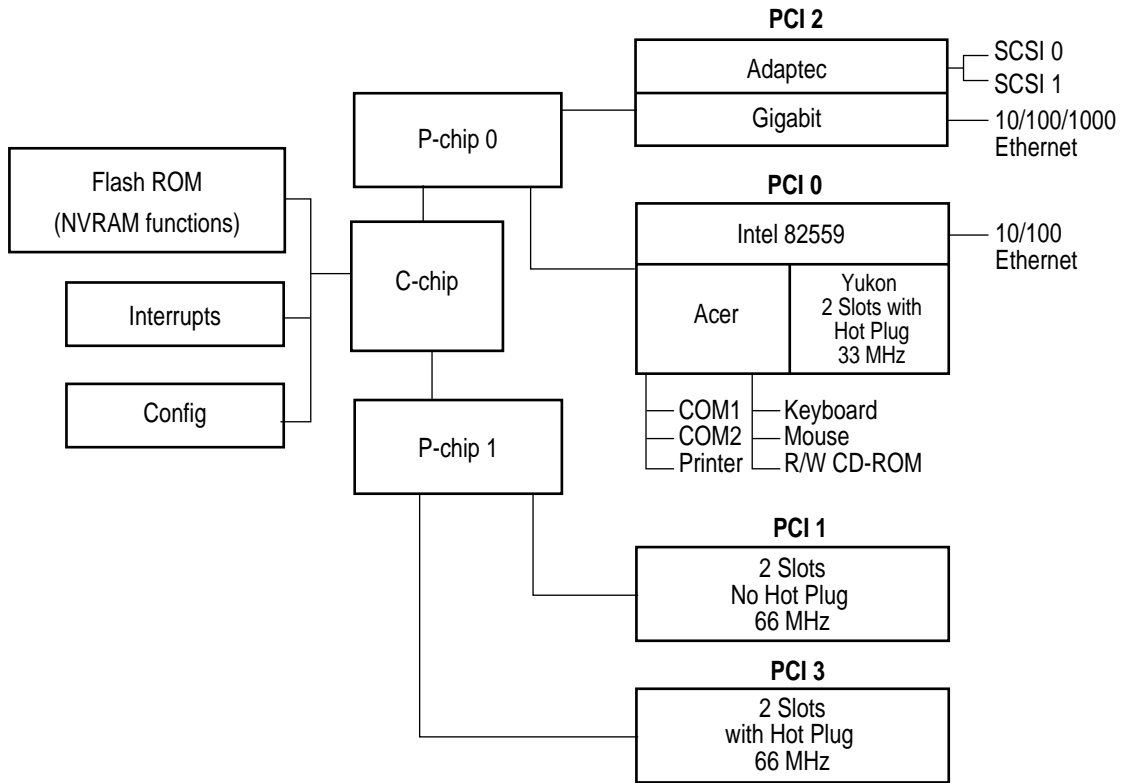


MR0445A

As seen in the block diagram, each P-chip controls two PCI buses. The C-chip controls accesses to memory on behalf of both P-chips.

The block diagram shows the other chips that are used to implement the onboard controllers.

Block Diagram of I/O Control



MR0407A


System Control

Control Panel

On the control panel are three pushbuttons:

- Power button
- Halt button
- Reset button

LEDs beside the Power and Halt buttons show status when those buttons are pushed. In addition, four LEDs on the control panel can be used in diagnosing system problems such as identifying failing CPUs, fans, or power supplies.



LEDs

In addition to the LEDs on the control panel, LEDs on the PCI slots indicate hot-swap capability. The two network connectors also have status LEDs.

Environment

In addition to providing for remote management of the system, the RMC chip performs environmental monitoring.

Power

The system is powered by two or three 500-watt power supplies. An optional third power supply provides N+1 redundancy. Previous (D*-57AAA-**) models used 375-watt power supplies. In those models the optional third supply provided N+1 redundancy for systems up to 8 Gbytes of memory. (The third supply was required for systems with more than 8 Gbytes of memory, and N+1 was not yet supported.)

Storage

The DS25 includes a six-device storage cage for Ultra3 SCSI one-inch disks or tapes. Disks can be 18.2, 36.4, 72.8, or 146 Gbytes. The tapes can be AIT or DAT.

The *AlphaServer* DS25 system includes integrated controllers for the following drives:

- Dual Ultra3 SCSI – one for the system disk cage and the other for an external storage shelf.
- IDE CD-Record drive

Ultra3 StorageWorks shelves can also be installed in the cabinet. Each Ultra3 shelf holds 14 one-inch disks.

Fibre Channel

The DS25 supports Fibre Channel, which eliminates issues with today's SCSI interfaces such as distance, bandwidth, scalability, and reliability. Fibre Channel (FC) is the answer to not only server-to-storage connections but also to server-to-server networking, because multiple protocols are supported. SCSI, TCP/IP, video, or raw data can all take advantage of high-performance, reliable Fibre Channel technology.

With the KGPSA PCI Fibre Channel adapter, the DS25 systems provide a storage interconnect that is 2.5 times as fast as UltraSCSI: 100 vs 40 Mbytes/sec data throughput. The KGPSA adapter allows users to manage storage, including the HSG80 RAID controller in a switched FC topology.

RAID (Redundant Array of Independent Disks)

With optional PCI RAID controllers users can organize disk data cost-effectively, improve performance, and provide high levels of storage integrity.

The optional RAID controllers have the following features:

- Support for hot-swap drives
- Automatic rebuild after hot swap
- Console support for booting system from RAID
- RAID levels 0, 1, 0+1, 3, 5
- Optional write cache
- Optional read cache
- Support for command queuing

Server Management

The *AlphaServer* products support important operational and platform management requirements.

Operational Management

Server/Network Management. *Compaq Insight Manager* is included with each system. This software tool allows users to monitor and control Alpha-based servers. *Insight Manager* consists of two components: a Windows-based console application and server- or client-based management data collection agents. Management agents monitor over 1,000 management parameters. Key subsystems are instrumented to make health, configuration, and performance data available to the agent software. The agents act upon that data, by initiating alarms in the event of faults and by providing updated management information, such as network interface or storage subsystem performance statistics.

Remote Server Management. The integrated remote management console (RMC) lets the operator perform several tasks from a serial console: monitor the system power, temperature, and fans, and reset, halt, and power the system on or off. The monitoring can be done locally or remotely.

Platform Management

The *AlphaServer* DS25 systems support platform management tasks such as manipulating and monitoring hardware performance, configuration, and errors. For example, the operating systems provide a number of tools to characterize system performance and display errors logged in the system error log file.

In addition, system console firmware provides hardware configuration tools and diagnostics to facilitate quick hardware installation and troubleshooting. The system operator can use simple console commands to show the system configuration, devices, boot and operational flags, and recorded errors. Also, the console aids in inventory support by giving access to serial numbers and revisions of hardware and firmware.

Error Reporting

Compaq Analyze, a diagnostic service tool used to determine the cause of hardware failures, is installed with the operating systems. It provides automatic background analysis, as it constantly views and reads the error log file. It analyzes both single error/fault events and multiple events. When an error condition is detected, it collects the error information and sends it and an analysis to the user. The tool requires a graphics monitor for its output display.

Security

- The front doors can be locked to prevent access to the disk drives and the rest of the system.
- An interlock sensor switch shuts down power if the chassis cover is removed while power is on.
- Password protection is offered by the SRM console and RMC.

Reliability, Availability, and Maintainability

Reliability and availability features are built into the CPU, memory, and I/O, and implemented at the system level.

Processor Features

- CPU data cache provides error correction code (ECC) protection.
- Parity protection on CPU cache tag store.
- Multi-tiered power-up diagnostics to verify the functionality of the hardware.

With two processors, when users power up or reset the system, each CPU, in parallel, runs a set of diagnostic tests. If any tests fail, the failing CPU is configured out of the system.

Responsibility for initializing memory and booting the console firmware is transferred to the other CPU, and the boot process continues. This feature ensures that a system can still power up and boot the operating system in case of a CPU failure.

Memory Features

- The memory ECC scheme is designed to provide maximum protection for user data. The memory scheme corrects single-bit errors and detects double-bit errors and total DRAM failure. It also detects RAM address errors.
- Memory failover. The power-up diagnostics are designed to provide the largest amount of usable memory, configuring around errors.

I/O Features

- ECC protection on the switch interconnect and parity protection on the PCI and SCSI buses.
- Extensive error correction built into disk drives.
- Optional external RAID improves reliability and data security.

System Features

Auto reboot. A firmware environment variable lets users set the default action the system takes on power-up, reset, or after an operating system crash. For maximum system availability, the variable can be set to cause the system to automatically reboot the operating system after most system failures.

Software installation. The operating systems are factory installed. The system is also available with no operating system, so that the customer can install Linux.

Diagnostics. During the power-up process, diagnostics are run to achieve several goals:

- Provide a robust hardware platform for the operating system by ensuring that any faulty hardware does not participate in the operating system session. This maximizes system uptime by reducing the risk of system failure.
- Enable efficient, timely repair.

Audible beep codes report the status of diagnostic testing. The system has a firmware update utility (LFU) that provides update capability for console and PCI I/O adapter firmware.

Thermal management. The air temperature and fan operation are monitored to protect against overheating and possible hardware destruction. The two system fans are at the rear of the system, and if the temperature rises, the system fans speed up; or if necessary shut down the system to prevent damage. Each power supply has two fans that also speed up when more cooling is needed. In addition, a fan cools the PCI area, and each CPU chip has a fan.

Error handling. Parity and other error conditions are detected on the PCI bus. The memory checking scheme corrects single-bit errors and detects double-bit errors. Multiple ECC corrections to single-bit errors detected by the operating systems help determine where in the system the error originated. Errors are logged for analysis.

Hot-plug disks. The hardware is designed to enable hot plug of disks. Disks can be removed and replaced while the rest of the system remains powered on and continues to operate. This feature contributes significantly to system availability. Since many disk problems can be fixed without shutting down the entire system, users lose access only to the disks that are removed.

Uninterruptible power supply. An external UPS can be used to support critical customer configurations. Because power is maintained for the entire system (CPU, memory, and I/O), power interruptions are completely transparent to users.

Installation and Maintenance

The systems are designed for easy hardware, software, and option installation. Options ordered with a system are preinstalled and tested at the factory. The *Tru64 UNIX* and *OpenVMS* operating systems are also installed at the factory.

Clustering

A cluster is a loosely coupled set of systems that behaves (is addressed and managed) like a single system, but provides high levels of availability through redundant CPUs, storage, and data paths. Clusters are also highly scalable; that is, CPU, I/O, storage, and application resources can be added incrementally to efficiently increase capacity. For customers, this translates to reliable access to system resources and data, and investment protection of both hardware and software.

Clustering allows multiple computer systems to communicate over a common interface, share disks, and spread the computing load across multiple CPUs. Clustering is implemented using our traditional interconnects and using the newest technology.

PCI to Memory Channel Interconnect

Under *Tru64 UNIX* and *OpenVMS*, users can build high-availability clusters using the PCI to Memory Channel interconnect. The Memory Channel interconnect is a high-bandwidth, low-latency PCI-based communications interconnect for up to eight *AlphaServer* systems. Data written to one computer's memory is shared by other computers on the Memory Channel bus.

The PCI adapter is the interface between a PCI and a Memory Channel bus. This bus is a memory-to-memory computer system interconnect that permits I/O space writes in one computing node to be replicated into the memories of all other nodes on the Memory Channel bus. A write performed by any CPU to its reflected address region results in automatic hardware updates to memory regions in other nodes. One node's write is "reflected" to other nodes as a direct side effect of the local write. This provides a memory region with properties similar to a high-performance shared memory across a group of nodes.

Operating System Support

For clustered *Tru64 UNIX* systems, TruCluster Software solutions allow users access to network services and provide further failover recovery from server, network, or I/O failures. *Tru64 UNIX* cluster systems use the SCSI bus, Ethernet, and/or PCI to Memory Channel interconnect bus between disks and systems.

OpenVMS cluster systems use the SCSI bus, Ethernet, FDDI, and/or Memory Channel as the interconnect between disks and systems.

Performance

HP has an ongoing program of performance engineering, using industry-standard benchmarks that allow comparisons across major vendors' systems. These benchmarks against competitive systems are based on comparable CPU performance, coupled with comparable memory and disk expandability.

See Table 1 for *AlphaServer* DS25 performance numbers. System performance, however, is highly dependent upon application characteristics. Thus, benchmark information is one helpful "data point" to be used in conjunction with other purchase criteria such as features, service, and price.

Sources of Performance Information

Performance information is available at:

<http://h18002.www1.hp.com/alphaserver/performance/index.html> and

http://www.ideasinternational.com/benchmark/spec/specfp_s2000.html

Also, see the following:

HP's Developer and Solution Partner Portal (DSPP) at:

<http://h21007.www2.hp.com/>

Service and Support

Hewlett-Packard provides a comprehensive set of services that range from migration, consulting, and training, to direct support of Alpha systems, software, and applications. For information on services, see <http://www.hp.com/hps/>

Hardware Warranty

The *AlphaServer* DS25 system and components, including CPU, memory, PCI controllers, and power supplies, have a 1-year on-site, 5-day per week, 9-hour per day hardware warranty with next business day response time.

StorageWorks components are supported by the comprehensive StorageWorks warranty: five years for disks, three years for controllers, two years for tape devices, and one year for other components. The first year includes on-site next-day response time. Network products carry the network products warranty.


Users can upgrade to higher levels of service through a variety of hardware supplemental services.

Software Warranty

The warranty for *Tru64 UNIX* and *OpenVMS* is conformance to SPD with advisory telephone support for a period of 90 days. Users can upgrade to higher levels of service through a variety of software supplemental services.


AlphaServer DS25 System Diagram

Side View



MR0444

Rear View



MR0290A

System Features at a Glance

Table 1: AlphaServer DS25 Features

CPU Features		68/1000	
Processor	Up to 2 Alpha 21264C chips		
CPU clock speed	1000 MHz		
Cache on chip	64 KB I-cache 64 KB D-cache		
On-board cache	8 MB per processor		
Internal Storage			
Hard disk	Ultra3 SCSI drive (18.2 GB, 36.4 GB, 72.8 GB, 146 GB); up to 876 GB		
Removable media	CD Record drive		
Memory (maximum)		16 GB	
Performance		1 CPU	2 CPUs
SPECint2000	678	–	
SPECint_rate2000	7.86	15.5	
SPECfp2000	985	–	
SPECfp_rate2000	11.4	21.5	
Standard Features		CD Record drive, 2 serial ports, 1 parallel port, integrated remote system management console, operating system license and customer documentation, Internet software	
I/O System			
I/O slots	6 64-bit slots		
Maximum PCI throughput	1.85 GB/s		
High Availability Features			
System	System auto reboot, thermal management, remote system management, RAID 0, 1, 0+1, 5, hot plug of disks, power supplies, fans, memory failover, ECC memory, ECC cache, N+1 power supply, SMP CPU failover, error logging, optional uninterruptible power supply		
<i>OpenVMS</i> clusters	Ethernet, SCSI, FDDI, PCI to Memory Channel Interconnect		
<i>Tru64 UNIX</i> TruCluster Solutions	Ethernet, SCSI, PCI to Memory Channel Interconnect		
Operating Systems		<i>Tru64 UNIX</i> , <i>OpenVMS</i> , and Linux	
Warranty			
Hardware	1-year, on-site, 5 day x 9 hour warranty with next business day response time		
Software	90-day telephone advisory support for <i>Tru64 UNIX</i> and <i>OpenVMS</i>		

Physical Characteristics

Table 2: AlphaServer DS25 Physical Characteristics

Dimensions	Pedestal	Rackmount	
Height	47.0 cm (18.5 in.)	22.2 cm (8.75 in.)	
Width	22.5 cm (8.85 in.)	44.5 cm (17.5 in.)	
Depth	69.9 cm (27.5 in.)	66.0 cm (26 in.)	
Weight	36 kg (80 lbs) typical; 40 kg (88 lbs) maximum	36 kg (80 lbs) nominal; 39 kg (86 lbs) maximum With brackets, slides, and cables 38 kg (84 lbs) nominal; 40 kg (88 lbs) maximum	
Dimensions	H9A10 M-Series cabinet	H9A15 M-Series cabinet	
Height x width x depth	67 x 23.6 x 41.5 in. 170 x 60 x 100 cm	79 x 23.6 x 41.5 in 200 x 60 x 100 cm	
Weight	Configuration dependent	Configuration dependent	
Environmental			
Temperature	Operating	50 – 95° F (10 – 35° C)	
	Nonoperating	41 – 122° F (5 – 50° C)	
	Storage (60 days)	–40 to 151° F (–40 to 66° C)	
	Rate of change (operating)	20° F /hr / 11° C /hr	
Relative humidity	Operating	20 – 90 % (maximum wet bulb: 82° F/28° C)	
	Nonoperating	10 – 90 % (maximum wet bulb: 90° F/32° C)	
	Storage (60 days)	10 – 95 % (maximum wet bulb: 115° F/46° C)	
Maximum altitude	Operating	10,000 ft (3,048 m)	
	Nonoperating	40,000 ft (12,192 m)	
Shock	Operating	10 G, 11ms half-sine	
Vibration	Operating	10 to 500 Hz at 0.1 G peak (limited by 0.02"/0.5 mm double amplitude)	
	Non-operating	1.03 Grms 5–300 Hz	
<i>Acoustics</i> (Declared values per ISO 9296 and ISO 7779)			
Current values for specific configurations are available from HP, note that 1 B = 10 dBA.			
Acoustics	L_{Wad}, B	L_{pAm}, BA (bystander positions)	
Idle	6.3	45	
Operating	6.4	46	
<i>Electrical</i> (Power supplies are universal, PFC, auto ranging, 100/240 Vac)			
Nominal voltage (Vac)	100	120	200 – 240
Voltage range (Vac) (temporary condition)	90 to 100	110 to 128	180 – 250
Power source phase	Single	Single	Single
Nominal frequency (Hz)	50/60	50/60	50/60
Frequency range (Hz)	49 to 51/59 to 61	49 to 51/59 to 61	49 to 51/59 to 61
RMS current (max. steady state)			
<i>Pedestal and Rackmount</i>			
Each cord, two power supplies	5.2A	4.2A	2.5A
Each cord, three power supplies	3.6A (3x for rack)	3.0A (3x for rack)	1.75A (3x for rack)
Maximum VA	1035	1010	975
<i>M-Series Cabinet (configuration dependent)</i>			
Nominal voltage (Vac)	100	120	220–240
Each nominal voltage (Vac)	24A	24A	16A



© Copyright Hewlett-Packard Company 2003.

Linux is a registered trademark of Linus Torvalds in several countries. UNIX is a trademark of The Open Group in the United States and other countries. All other product names mentioned herein may be trademarks of their respective companies.

Hewlett-Packard shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Revision 2.0 – February 2003