COMPAQ

Compaq AlphaServer ES40 Systems

Technical Summary

Contents

Compaq AlphaServer ES40 Systems

The *Compaq AlphaServer* ES40 system is a high performance, scalable enterprise server for business, technical, and scientific applications. With its flexible design, it can go in an office environment or computer room. It performs as a high-capacity database server, high-performance application server, Network File System (NFS) server, or Internet server.

***Compaq AlphaServer* products use the 64-bit Alpha RISC architecture that supports multiple operating systems: Tru64 UNIX®, Windows NT, OpenVMS operating systems, and even Linux. The *AlphaServer* ES40 system integrates into your current operating environment and anticipates future needs with upgrade capabilities.**

More information on *Compaq AlphaServer* ES40 systems is on the World Wide Web: http://www.digital.com/info/alphaserver/products.html.

System Overview

The *Compaq AlphaServer* ES40 uses the third-generation Alpha processor, the Alpha 21264, that offers up to twice the performance of a comparable speed EV56 Alpha chip. The switch-based system interconnect exploits the full potential of the 21264 chip.

The system is available in three packages: a standalone tower, a pedestal system with expanded storage capacity, and a rack-mount system. Each system variant is available with up to four EV6 processors (500 MHz), each with a 4 MB ECC cache.

Two models are available: an entry-level system that supports 16 DIMMs (8 GB) and 6 PCI slots or a system that supports 32 DIMMs (16 GB) and 10 PCI slots.

The chassis used in the tower is rotated 90 degrees and installed in a pedestal enclosure or in a cabinet. Each chassis provides space for up to eight 1.6-inch Ultra2 SCSI disks. Up to three StorageWorks shelves can also be installed in the pedestal system; and in a cabinet up to six StorageWorks shelves are supported.

Systems can be purchased with Tru64 UNIX, OpenVMS, and Microsoft Windows NT Server operating systems installed. Or they can be purchased without any operating system, allowing customers to install Linux.

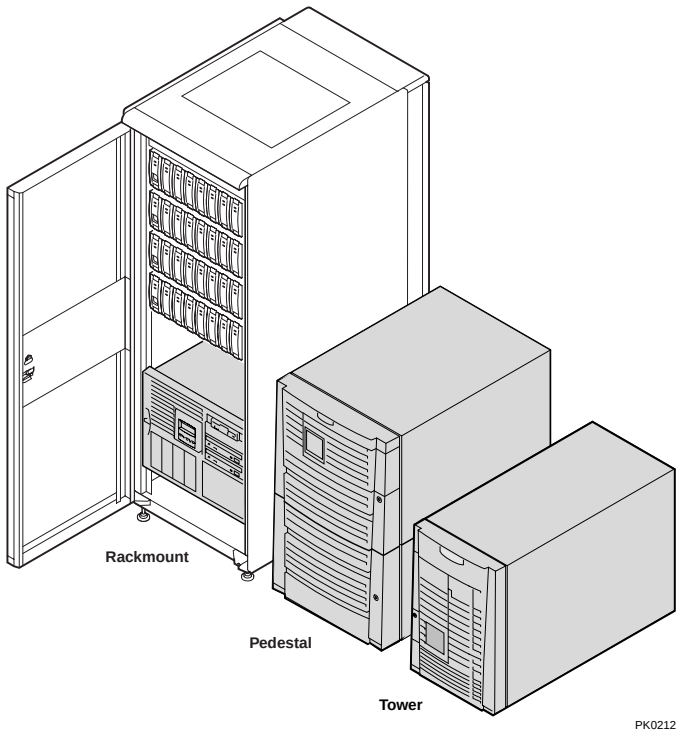

Features and Benefits

- **Leadership 64-Bit Architecture**

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

- **Performance**

The Alpha 21264 chip, the world's fastest microprocessor, is offered with a switch-based interconnect that supports four processors. This switch-based system provides a memory bandwidth of up to 5.2 Gbytes/sec (peak) using two 256-bit memory buses running at 83 MHz.

- **Scalability**

An entry-level system (Model 1) offers one processor, 512 MB memory, and 6 PCI slots. Model 2, can have 16 GB memory, and 10 I/O slots. The pedestal system provides space for another 21 disks beyond the eight in the chassis. The cabinet can house up to four systems and/or additional disks. All variants support up to four processors.

- **Reliability and Availability**

The *AlphaServer* ES40 uses the latest technologies to achieve redundancy, error correction, and fault management. The system has redundant fans and power supplies; fans, power supplies, and disks can be hot swapped. The remote management console (RMC) monitors, sends alerts, and records possible error conditions. The RMC can be accessed even if the system is completely down.

Third-Generation Alpha Chip

The Alpha 21264 microprocessor (EV6) is a superscalar superpipelined implementation of the Alpha architecture. The Alpha 21264 chips are manufactured using state of the art CMOS-6 process, with a feature size of 0.35 micron. Over 15.2 million transistors are on one die.

Designed for performance, the 21264 achieves this goal by carefully studied and simulated architectural and circuit analysis. The 21264 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 large, off-chip cache—4 MB—runs at 200 MHz.

Each chip has 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 features in the EV6 architecture maximize performance: Four instructions are fetched each cycle, and then how those instructions are handled 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 quickly as possible. This allows for faster execution, since critical path computations are started and completed as soon as possible.

In addition, the 21264 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 21264 predicts branch directions and speculatively executes down the predicted path. The sophisticated branch prediction in the 21264 coupled with the speculative and dynamic execution extracts maximum instruction parallelism from applications.

For more information about the chip, see **EV6CHIP.pdf** .

Alpha 21264 Features

- Out-of-order instruction execution
- Large (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

Model Variants

AlphaServer ES40 systems are offered in two models. The entry-level model provides connectors for four DIMMs on each of the memory motherboards (MMBs) and connectors for six PCI options on the PCI backplane.

Model 1	Model 2
1–4 CPUs	1–4 CPUs
Up to 8 GB memory	Up to 16 GB memory
6 PCI slots	10 PCI slots

To upgrade from Model 1, you replace the PCI backplane and the four memory motherboards. A Model 2 system has eight DIMMs on each MMB (32 total) and 10 PCI slots.

System Packaging

All system variants use the same chassis, with space for two disk cages in the front. Each cage holds up to four 1.6-inch Ultra2 SCSI disks.

The chassis has a compartment for the CPUs and memory and another one for the PCI I/O cards. The control panel, CD-ROM and floppy drives, and two additional removable media bays are in the front, along with space for up to two disk cages. The power supplies are in the rear.

Chassis, Pedestal View

Additional storage is available in the pedestal enclosure. The chassis is mounted over a box that can house up to three StorageWorks shelves (two in the front and one in the rear).

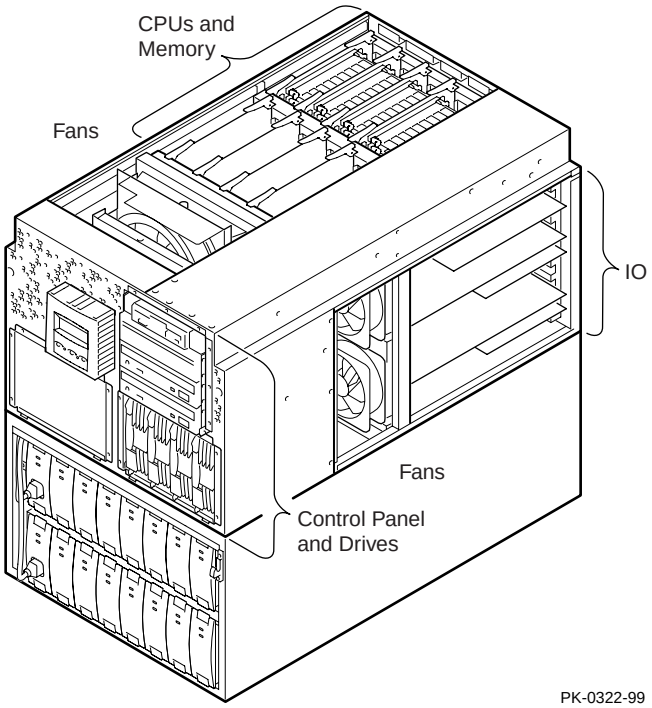

PK-0322-99

Chassis, Tower View

In the tower enclosure, the chassis is rotated 90 degrees, with the control panel and removable media bays keeping their same horizontal position as in the pedestal/rack system.

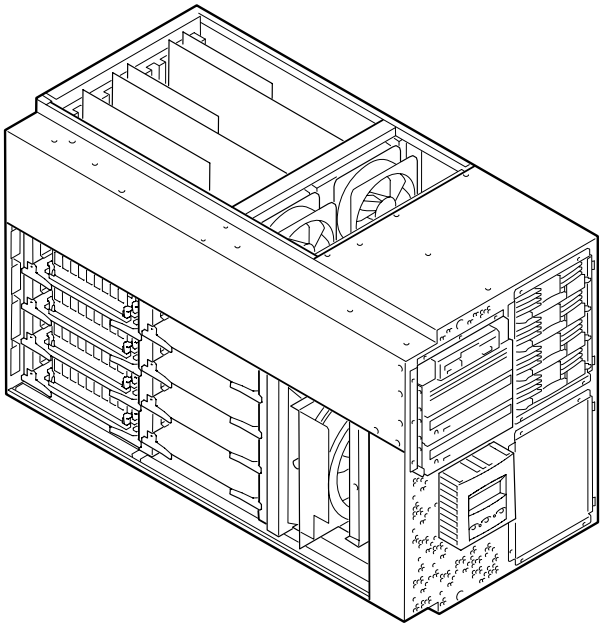

PK-0329-99

Rackmount Chassis

When installed in a cabinet, the chassis has the same orientation as that used in the pedestal. The chassis can be mounted in the 67-inch or 79-inch M-series cabinet, which is a 19-inch wide RETMA cabinet. Each system requires 14 inches (8U) of vertical space. With the maximum number of systems in a cabinet (four), there can be two StorageWorks shelves. When the maximum number of disks is desired (six shelves), there can be three systems. (These numbers are for the H9A15 79-inch cabinet.)

Architecture

This system is designed to maximize the potential of the Alpha 21264 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 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, two P-chips, and eight D-chips.

- **C-chip.** Provides the command interface from the CPUs and main memory. Although there is only one C-chip 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 two independent 64-bit 33 MHz PCI buses.

This chipset is similar to that used in the *AlphaServer* DS20 system; however, this chipset supports up to four CPUs and up to 32 Gbytes memory (16 GB available now). Interleaving occurs when at least two memory arrays are used.

Two 256-bit memory buses support four memory arrays, yielding a 5.2 Gbytes/sec system bandwidth. Transactions are ECC protected. Upon the receipt of data, the receiver checks for data integrity and corrects any errors.

System Block Diagram

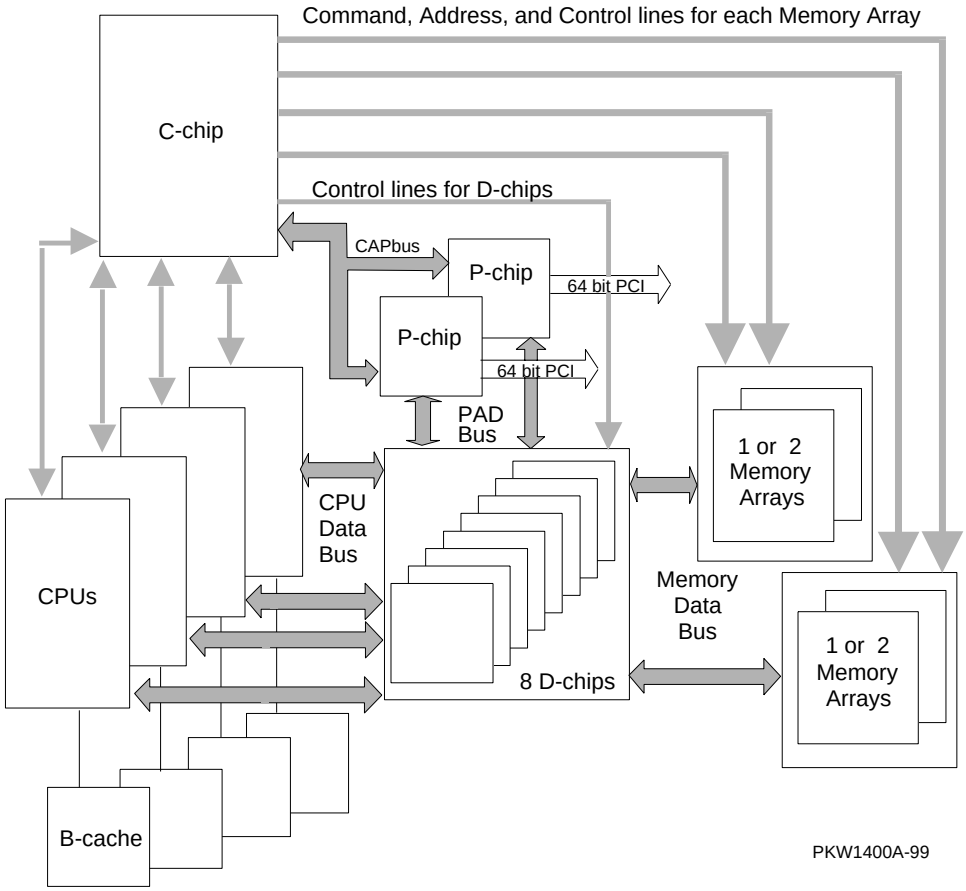

System Board

The interconnect switch is implemented on the system board 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 AlphaBIOS console code, the SRM console code, and the NVRAM data.

Connectors are provided for four CPU modules and four memory motherboards.

One corner of the board holds logic for the remote management console (RMC).

On the back side of the module are connectors for the power supplies. The I/O connector provides both signals and power to the PCI backplane.

Component and Connector Locations

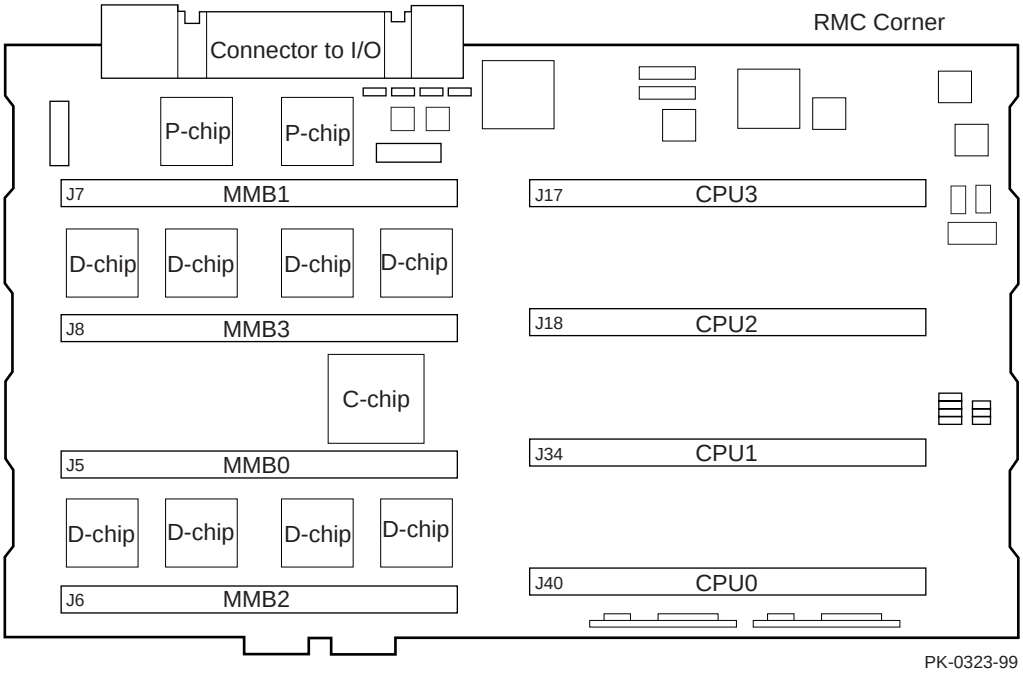

Processor Module

An *AlphaServer* ES40 can have up to four CPU modules. In addition to the Alpha 21264 chip, the CPU module has a 4-Mbyte second-level cache and a 2.2V DC-to-DC converter with heatsink that provides the required voltage to the Alpha chip. Power-up diagnostics are stored in a flash ROM on the module.

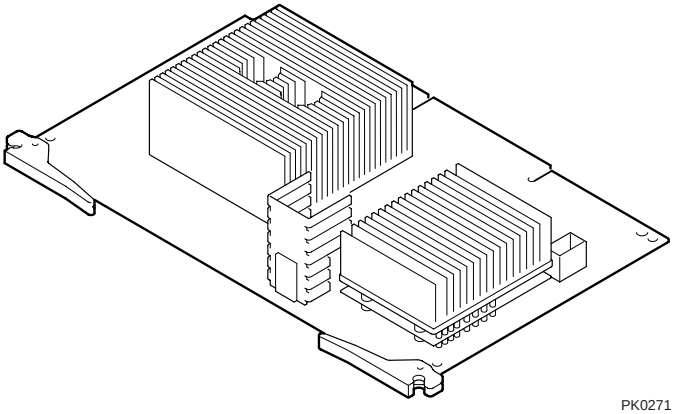

PK0271

Processor Configuration Rules

- The first CPU module is installed in CPU slot 0.
- Additional CPUs are installed in the next available slot.
- CPUs must be identical in speed and cache size.

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 12 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 83 MHz. This yields 2.67 GB/sec bandwidth per bus (32 x 83 MHz = 2.67 GB/sec). The maximum bandwidth is 5.2 GB/sec (2.67 x 2).

The design challenge was to maximize the capabilities of the two wide data buses. Distributing the 256 data bits equally over two memory motherboards (MMBs) was one solution: simultaneously, in a read operation, 128 bits come from one MMB and the other 128 bits come from another MMB, 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 per MMB (one for each array) 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.

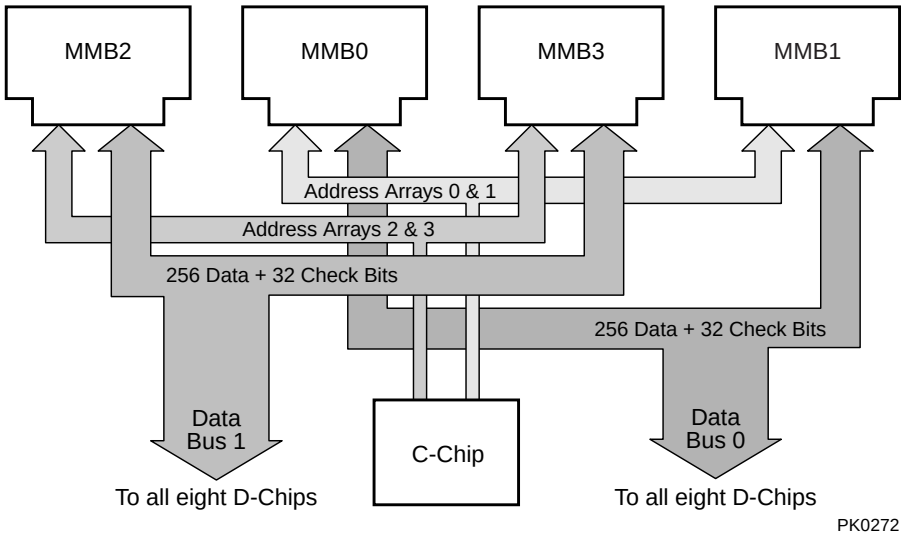

PK0272

Memory Options

Each memory option consists of four 200-pin industry-standard DIMMs. The DIMMs are synchronous DRAMs. The Model 1 system supports up to four memory options (16 DIMMs), and the Model 2 system supports up to eight options (32 DIMMs). Memory options are available in the following sizes:

- 256 Mbytes (64 MB DIMMs)
- 512 Mbytes (128 MB DIMMs)
- 1 Gbyte (256 MB DIMMs)
- 2 Gbytes (512 MB DIMMs)

With the 2 GB option, Model 1 supports 8 GB memory and Model 2 supports 16 GB memory. In the future, when 1 GB DIMMs become available, the maximums will be 16 and 32 GB, as supported by the architecture.

Memory Motherboard

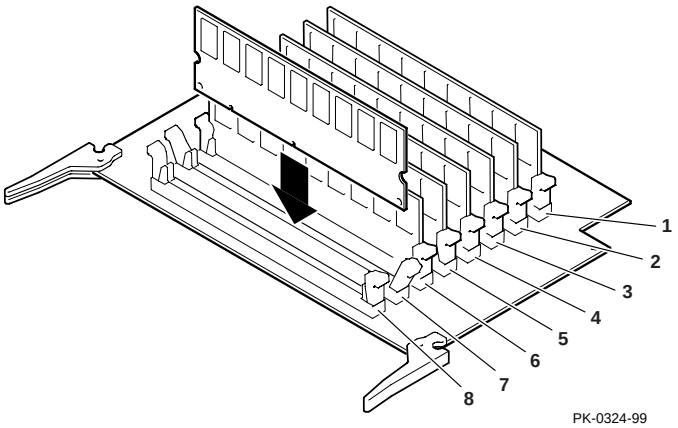

Memory Configuration

A memory option (or set) consists of four DIMMs, and all four DIMMs must be the same size.

- Fill sets in numerical order. Populate all 4 slots in Set 0, then populate Set 1, and so on.
- An array is one set for systems that support 16 DIMMs and two sets for systems that support 32 DIMMs.
- DIMMs in an array must be the same size and type. For example, suppose you have populated Sets 0, 1, 2, and 3. When you populate Set 4, the DIMMs must be the same size and type as those installed in Set 0. Similarly, Set 5 must be populated with the same size and type of DIMMs as are in Set 1, and so on, as indicated in the following table.

Array	Model 1 System	Model 2 System
0	Set 0	Set 0 and Set 4
1	Set 1	Set 1 and Set 5
2	Set 2	Set 2 and Set 6
3	Set 3	Set 3 and Set 7

Memory Performance Considerations

With one memory option (4 DIMMs) installed in either Model 1 or Model 2, 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 noninterleaved operations. With four memory options (16 DIMMs) installed, memory read-write operations are four-way interleaved, maximizing memory throughput.

Memory Arrays and Sets on Memory Motherboards

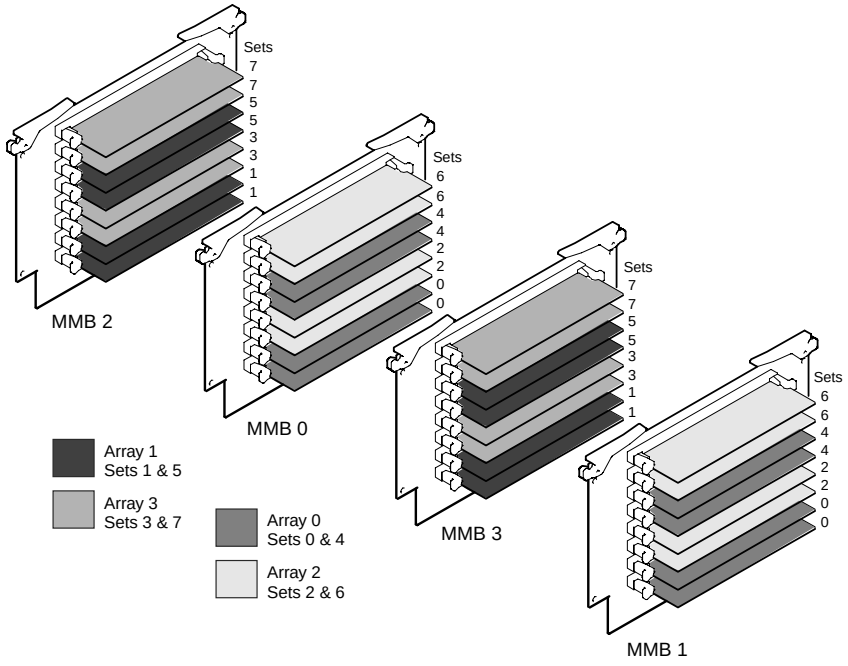

System I/O

Two industry-standard PCI I/O buses allow you to use inexpensive, widely available I/O options. Both 32-bit and 64-bit PCI options can be used; 3.3V and 5V options are supported.

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.

The PCI bus implementation has the following characteristics:

- Fully compliant with the PCI Version 2.1 Specification
- Operates at 33 MHz, delivering a peak bandwidth of 500 MB/sec; over 250 Mbytes/sec for each PCI bus
- 6 option slots (Model 1) or 10 option slots (Model 2)
- Supports three address spaces: PCI I/O, PCI memory, and PCI configuration space
- Supports byte/word, tri-byte, and longword operations
- Exists in noncached address space only

Block Diagram of I/O Control

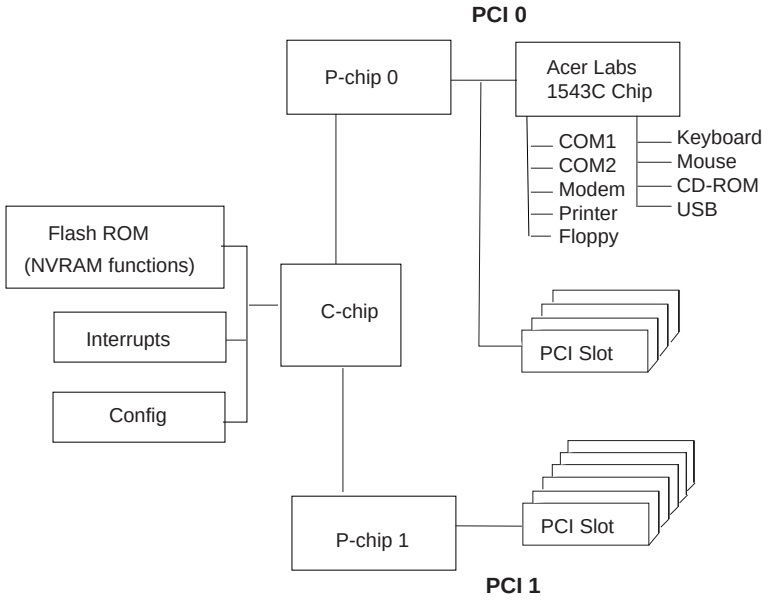

PK-0319A-98

I/O Implementation

In a Model 2 system that has 10 I/O slots, PCI 0 has 4 slots, and PCI 1 has 6 slots. In a Model 1 system with 6 slots, each PCI has 3 slots.

The Acer Labs 1543C chip provides the bridge from PCI 0 to ISA. The C-chip controls accesses to memory on behalf of both P-chips.

I/O Configuration Rules

A VGA controller, if present, must be installed in PCI 0.

I/O Ports

At the rear of the system are connectors offering access to two serial communications ports, one modem port for the remote management console, one parallel port, ports for the keyboard and mouse, and two USB ports. Two SCSI breakouts are also on the back panel.

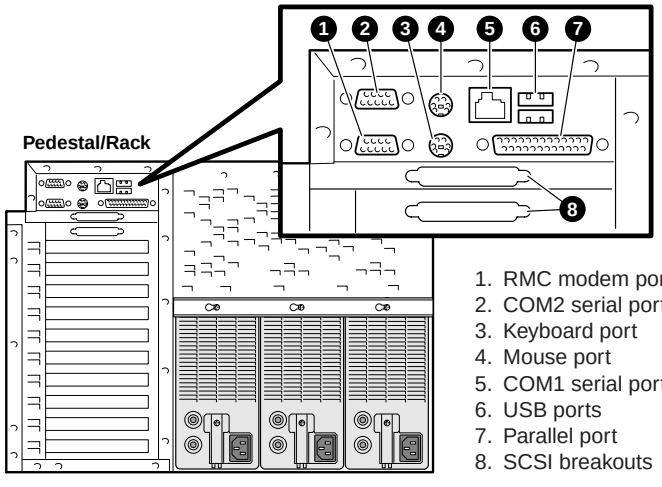

1. RMC modem port
2. COM2 serial port
3. Keyboard port
4. Mouse port
5. COM1 serial port
6. USB ports
7. Parallel port
8. SCSI breakouts

PK-0325-99

System Control

Close monitoring and control of the system environment and hardware is done by the remote management console (RMC). This logic also allows the system operator to monitor the system from a remote location. The RMC logic is implemented using a PIC17C44 microprocessor on the system board. The RCM firmware code resides on the microprocessor and in flash memory. The RMC is powered by an auxiliary 5V supply, so even when the system is powered off at the control panel the RMC can be accessed—so long as the system is plugged in.

The RMC provides the following monitoring and control functions:

- Monitors thermal sensors on the CPUs, the PCI backplane, and the power supplies
- Monitors voltages, power supplies, and fans
- Handles hot swapping of power supplies and fans
- Controls the control panel display and writes status messages on the display
- Detects alert conditions such as excessive temperature, fan failure, and power supply failure and sends an alert
- Shuts down the system if any fatal conditions exist
- Records error log information in nonvolatile RAM on each failing device

Storage

Removable media in the chassis includes a CD-ROM drive, floppy diskette drive, and two 5.25 inch half-height bays, which can be combined for one 5.25 inch full-height bay.

A disk cage in the front of the chassis supports up to four Ultra2 SCSI disk drives (9.1 or 18.2 GB). The Ultra2 SCSI drives are 1.6-inch 10,000 RPM. Another disk cage can be added so that up to eight disks can be installed. Each cage requires its own SCSI adapter.

The pedestal system can include up to three StorageWorks shelves, each of which holds seven disks. In the rackmount cabinet up to six StorageWorks shelves can be installed. The StorageWorks shelves are 16-bit (wide) UltraSCSI; UltraSCSI disks run at 7,200 RPM.

RAID (Redundant Array of Independent Disks)

The system can be configured with optional PCI RAID controllers to organize disk data cost-effectively, improve performance, and provide high levels of storage integrity. Today, RAID is only available with StorageWorks shelves.

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, 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. ServerWORKS Manager software is included with each system. This software utilizes the Simple Network Management Protocol (SNMP) environment to assist the network or server administrator by constantly monitoring the network for problems, thus avoiding expensive downtime. The software monitors vital server information, such as CPU and file system utilization, as well as the condition of the network supported by the management console.

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 through a modem.

Platform Management

The *AlphaServer* ES40 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 provides inventory support and configuration management 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

- 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 top cover to the CPU/memory area is removed while power is on.
- Password protection is offered by the SRM console, AlphaBIOS, and RMC.

Reliability and Availability Features

The *AlphaServer* ES40 system achieves an unparalleled level of reliability and availability through the careful application of technologies that balance redundancy, error correction, and fault management. 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.

When you 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 another 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. Messages on the operator control panel power-up/diagnostic display indicate the test status and component failure information.

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.
- 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 internal RAID improves reliability and data security.
- Disk hot swap.

System Features

Auto reboot. On systems running Tru64 UNIX or OpenVMS, a firmware environment variable lets you 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. Windows NT auto reboots by default, but lets you specify a countdown value so you can stop the system from booting if you need to carry out other tasks from the console firmware.

Software installation. The operating systems are factory installed. Factory installed software (FIS) allows you to boot and use your system in a shorter time than if you install the software from a distribution kit.

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. A fail-safe loader provides a means of reloading the console in the event of corrupted firmware.

Thermal management. The air temperature and fan operation are monitored to protect against overheating and possible hardware destruction. Six fans provide front to back cooling, and the power supplies, in the rear, have their own fans. If the termperature rises, the system fans speed up; or if necessary to prevent damage, the system shuts down. If the main fan, which cools the system card cage, fails, a redundant fan takes over.

Error handling. Parity and other error conditions are detected on the PCI buses. 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 in determining where in the system the error originated. Errors are logged for analysis.

Disk hot swap. The hardware is designed to enable hot swap of disks. Hot swap is the removal of a disk or disks from any of the storage compartments 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.

N+1 power redundancy. A second or third power supply can be added to provide redundant power to the chassis. A second power supply is needed for more than two CPUs or if a second disk cage is installed. In this case the third supply provides redundancy. Power supplies are 735 watts (DC). Each has two LEDs to indicate the state of power to the system.

An external UPS can be purchased 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 operating systems are also installed at the factory.

Additional CPUs, memory, power supplies, and disks can be added to the tower and pedestal systems by anyone with appropriate technical training and experience. Installation of components in a rackmount system is reserved for service providers and self-maintenance customers.

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

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

Windows NT cluster systems use SCSI buses and Ethernet.

The primary means of clustering *AlphaServer* ES40 systems depends on the operating system.

- CI clusters, OpenVMS only
- Memory Channel, Tru64 UNIX and OpenVMS
- SCSI clusters, Tru64 UNIX, OpenVMS, and Windows NT

Performance

Compaq 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 the record-breaking performance numbers of the *AlphaServer* ES40 systems. For example:

- SPECweb96 Internet Server — 11,159 operations per second, using an *AlphaServer* ES40 UNIX system with four processors

The SPECweb96 benchmark focuses on server performance for static Web pages, measuring the ability of the server to service HTTP requests or "gets." One or more clients are used by SPECweb96 to send the HTTP requests to the Web server. The software then measures the response time for each request. At the end of the benchmark run, SPECweb96 calculates a metric based on overall throughput, measured as maximum benchmark operations per second.

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 on the Internet.

- *World Wide Web*
  http://www.digital.com/alphaserver/products/perform3.html
- *FTP*. Access performance documents from
  ftp://gatekeeper.dec.com/index.html.
  The directory is pub/DEC/DECinfo/performance/sys.

Information for Compaq Partners

Compaq and its partners and customers can register with DIGITAL Business Link to access information needed to purchase and sell Compaq products and services; including access to pricing, product, sales, and marketing information. http://www.businesslink.digital.com/.

Also see the Alliances and Partners Web site located at http://www.digital.com/other-servers.html

and the Compaq Partner Network (CPN): http://CPN.compaq.com

Service and Support

Compaq 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 Compaq Services, point your Web browser to http://www.service.digital.com/.

Hardware Warranty

The *AlphaServer* ES40 system and components, including CPU, memory, PCI controllers, and power supplies, have a 3-year on-site, 5-day per week, 9-hour per day hardware warranty with next-day response time.

StorageWorks components contained in the pedestal or cabinet systems 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 in the pedestal or cabinet systems 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. The warranty for Windows NT is conformance to the written material accompanying the software. Users can upgrade to higher levels of service through a variety of software supplemental services.

Compaq AlphaServer ES40 System Diagrams

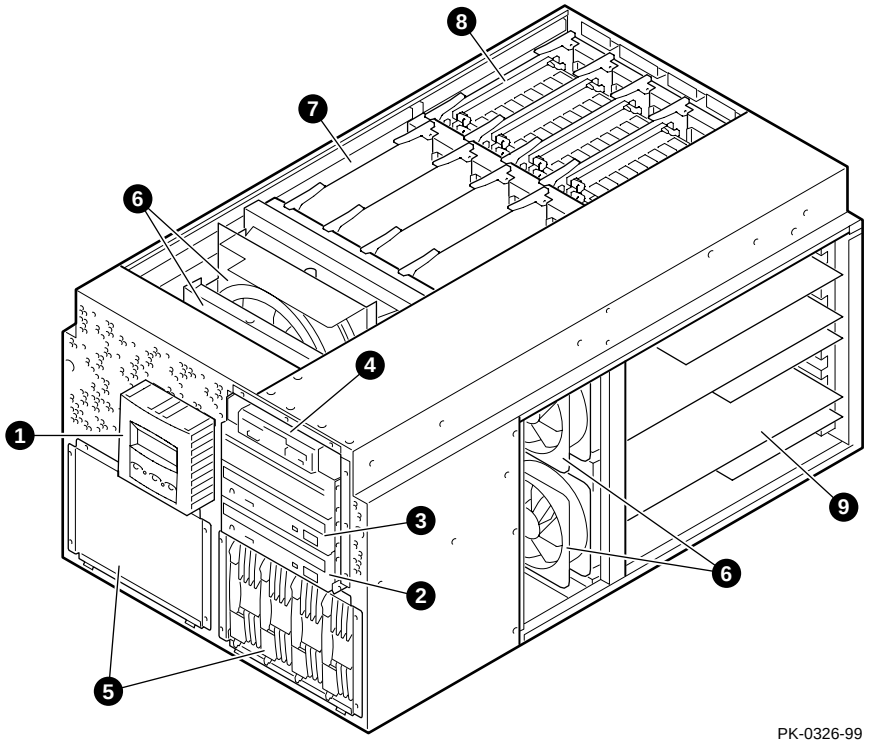

PK-0326-99

1. Control panel
2. CD-ROM drive
3. Removable media bays
4. Floppy diskette drive
5. Storage drive bays
6. Fans
7. CPUs
8. Memory
9. PCI slots

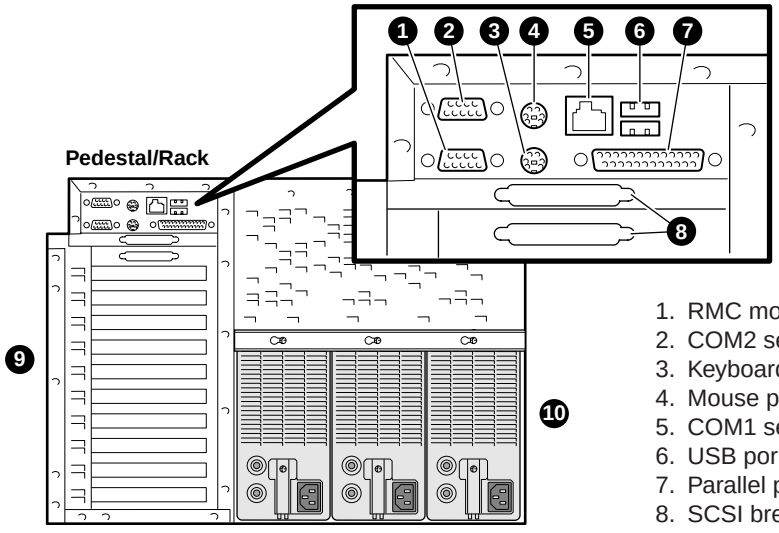

PK-0327-99

1. RMC modem port
2. COM2 serial port
3. Keyboard port
4. Mouse port
5. COM1 serial port
6. USB ports
7. Parallel port
8. SCSI breakouts
9. PCI slots
10. Power supplies

System Features at a Glance

Table 1 provides a quick reference to features of the *Compaq AlphaServer* ES40 systems.

Table 1 *AlphaServer* ES40 Features

CPU Features			
Symmetric multiprocessing	1–4 processors		
Processor	Alpha 21264		
CPU clock speed	500 MHz		
Cache on chip	64 KB I-cache		
	64 KB D-cache		
On-board cache	4 MB		
Memory (maximum)	8 GB in Model 1 and 16 GB in Model 2		
Performance	**1 CPU**	**2 CPUs**	**4 CPUs**
SPECint95	27.3	–	–
SPECfp95	57.7	76.3	104
SPECint_rate95	244	481	957
SPECfp_rate95	514	972	1658
SPECweb96	–	–	11,159 ops/sec @ 23.9 ops
TPC-C	–	–	25,903 tpmC @ $37
Standard Features	1.44 MB diskette drive, CD-ROM drive, 1 RMC dedicated serial port, 2 serial ports, 1 parallel port, keyboard and mouse, integrated remote system management console, operating system license and customer documentation, Internet software		
Internal Storage			
Removable media	CD-ROM, floppy diskette, two HH 5.25” removable media bay		
System chassis storage	8 disks (145 GB with 18.2 GB disks) Ultra2 SCSI		
Additional pedestal storage	21 disks (382 GB with 18.2 GB) UltraSCSI		
Cabinet storage, maximum	66 disks – 1.2 TB		
I/O System			
I/O slots	6 PCI slots in Model 1 or 10 PCI slots in Model 2		
Maximum PCI throughput	Over 500 MB/sec with two 256 MB/sec buses		
High Availability Features			
System	System auto reboot, thermal management, remote system management, RAID 0, 1, 0+1, 5, hot swap of disks, power supplies, and fans, memory failover, ECC memory, ECC cache, N+1 power supply, SMP CPU failover, error logging, optional uninterruptible power supply		
OpenVMS clusters	CI, Ethernet, SCSI, FDDI, PCI to Memory Channel Interconnect		
UNIX TruClusters Solutions	SCSI, PCI to Memory Channel Interconnect		
Windows NT cluster	Ethernet, FDDI, SCSI		
Operating Systems	Tru64 UNIX, OpenVMS, Microsoft Windows NT Server, Linux		
Warranty			
Hardware	3-year, on-site, 5 day x 9 hour warranty with 24-hour response		
Software	90-day telephone advisory support for OpenVMS and Tru64 UNIX		

Physical Characteristics

Table 2 details basic physical characteristics of the system, and Table 3 gives the electrical characteristics.

Table 2 *AlphaServer* ES40 Physical Characteristics

Dimensions	Tower	Pedestal	Rackmount
Height	50.8 cm (20.0 in.)	78.2 cm (30.8 in.)	35.2 cm (13.87 in.)
Width	38.7 cm (15.25 in.)	50.8 cm (20.0 in.)	44.7 cm (17.6 in.)
Depth	78.7 cm (31.0 in.)	80.6 cm (31.75 in.)	76.5 cm (30.1 in.)
Weight	65 kg (143 lb) typical; 96 kg (211 lb) maximum	127 kg (280 lb) typical; 159 kg (350 lb) maximum	With brackets, slides, & cables Nominal 59 kg (130 lb) Maximum 92 kg (202 lb)

Environmental		
Temperature	Operating	10–35° C (50–95° F)
	Storage (60 days)	–40 to 66° C (–40 to 151° F)
	Rate of change	11° C/hr (20° F/hr)
Relative humidity	Operating	20–80%
	Nonoperating	20–80%
	Storage (60 days)	10–95%
	Rate of change	20%/hr.
Maximum wet bulb temperature	Operating	28° C (82° F)
	Storage (60 days)	46° C (115° F)
Minimum dew point temperature	Operating	2° C (36° F)
Heat dissipation	Nominal	Maximum
Tower and Rack	900 w, 3074 BTU/hr	1300 w, 4440 BTU/hr
Pedestal	1480 w, 5054 BTU/hr	2400 w, 8196 BTU/hr
H9A10/H9A15	Configuration dependent	4800 w, 16392 BTU/hr
Airflow and quality	Intake location	Front
	Exhaust location	Rear–Tower, Pedestal, Rack Rear/Top–H9A10/H9A15
Altitude	Operating	3037 m (10,000 ft)
	Nonoperating	12,190 m (40,000 ft)
Vibration	Operating	10–500 Hz .1 G peak
Mechanical shock	Operating	
	Tower/Pedestal	7.5 G, 10 +/– 3ms
	M-Series Cabinet	5.0 G, 10 +/– 3ms

Acoustics—Declared values per ISO 9296 and ISO 7779

Current values for specific configurations are available from Compaq. 1 B = 10 dBA

Acoustics	L_{WAd3} B	L_{pAm5} dBA (bystander position)
Idle	6.6	48
Operating	6.6	48

Table 3 *AlphaServer* ES40 Electrical Characteristics

Nominal voltage (Vac)	100	120	200–240
Voltage range (Vac) temporary condition)	90–110	110–128	180–250
Power source phase	Single	Single	Single
Nominal frequency (Hz)	50/60	50/60	50/60
Frequency range (Hz)	49–51/59–61	49–51/59–61	49–51/59–61
RMS current (max. steady state)			
Tower and Rackmount			
Single power cord	11.0 A	8.5 A	5.0 A
Multiple power cords	6.5 A	5.3 A	3.0 A
Pedestal			
Each power cord	12.0 A	10.5 A	7.0 A
M-Series cab config.-dependent			
Nominal voltage (Vac)	100	120	220–240
Each power cord	24 A	24 A	16 A

NOTE: Power supplies are universal, PFC, auto ranging, 100 / 120 / 200–240 Vac.

Features may differ among operating environments. Performance may vary depending on configuration, application, and operating environment.

Compaq believes the information in this publication is accurate as of its publication date; such information is subject to change without notice. Compaq is not responsible for any inadvertent errors.

Compaq conducts its business in a manner that conserves the environment and protects the safety and health of its employees, customers, and the community.