AlphaServer 4000/4100 Systems

Technical Summary

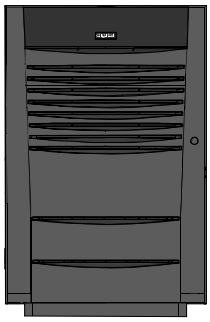

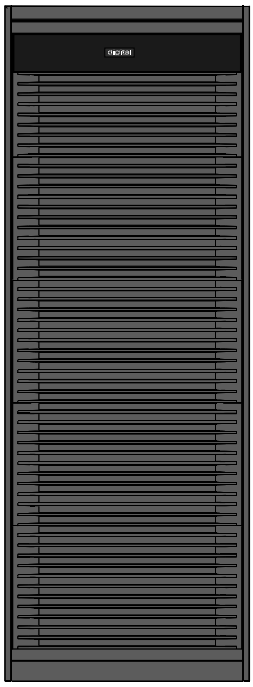

Contents

1
System Overview
Features and Benefits

2
System Packages
4100 System Drawer
4000 System Drawer

3
Pedestal System
Cabinet System

4
Architecture
CPUs
Memory
I/O

5
System Bus
AlphaServer 4100 Motherboard
AlphaServer 4000 Motherboard

6
Processor Module
Alpha Microprocessor Features
Processor Module Variants
Processor Configuration Rules

7
Memory
Memory Options
Memory Technology
Memory Configuration Rules
System I/O
PCI Bus
EISA Bus
Optional PCI Card Cage

8
I/O Implementation
I/O Bulkhead

9
Storage Architecture
Pedestal Disk Storage
Cabinet Disk Storage
RAID (Redundant Array of Independent Disks)
External Expansion
Clustering

10
PCI to MEMORY CHANNEL Interconnect
Reliability and Availability Features
Processor Features
Memory Features
I/O Features
System Features

11
Server Management
Operational Management
Platform Management
Error Reporting
Installation and Maintenance

12
Performance
Sources of Performance Information
Information for DIGITAL Partners
Service and Support
Hardware Warranty
Software Warranty

13
AlphaServer 4000/4100 Configurations

14
AlphaServer 4000/4100 Configurations

15
System Features at a Glance

16
Physical Characteristics

AlphaServer 4000/4100 Systems

The DIGITAL AlphaServer 4000/4100 products are scalable, dependable, open systems that deliver high productivity for business and technical/scientific applications at an affordable price. VLM (very large memory) power is now available in a midrange system at a midrange price. As high-capacity database servers, high-performance application servers, Network File System (NFS) servers, or Internet servers, the AlphaServer 4000/4100 products deliver exceptional value, choice, and investment protection.

The 64-bit Alpha RISC architecture at the heart of DIGITAL AlphaServer products provides fast processing and quick response for today's applications, and will run the advanced applications of tomorrow. You can choose from three popular operating environments—the DIGITAL UNIX®, Windows NT, and OpenVMS operating systems — as well as thousands of available applications. These servers integrate into your current operating environment and anticipate future needs with upgrade capabilities.

More information on AlphaServer 4000/4100 systems is on the World Wide Web: http://www.digital.com/info/alphaserver/products.html.

System Overview

Powerful symmetric multiprocessing (SMP) servers, the AlphaServer 4000/4100 systems provide for growth and configuration flexibility. The basic building block of the systems is the system drawer. The very same drawer is used in the pedestal and in the cabinet. The cabinet has space for multiple system drawers with ample room for storage.

The system drawer for the 4100 holds up to four CPUs, up to 8 gigabytes of ECC-protected memory (four memory pairs), and 64-bit PCI I/O subsystems, dual-channel with 8 slots. The system drawer for the 4000 accommodates two CPUs, 4 gigabytes of ECC-protected memory (two memory pairs), and up to 16 PCI slots. (Upgrades can be done to add the second PCI card cage, to provide the additional 8 PCI slots.)

You can configure a system for today's needs: a 5/300 CPU with an onboard 2-Mbyte cache or a 5/400 or 5/466 CPU with a 4-Mbyte cache, in a choice of two drawers, installed in a pedestal or a cabinet, with as much memory as you need today. Add CPUs and memory as needed, or upgrade as new processors become available. You can upgrade from a 4000 to a 4100 or add another eight slots of PCI I/O to a 4000. You can even "upgrade" from a 4100 to a 4000 if you need more I/O slots.

Features and Benefits

- **Performance and Price/ Performance**

The balanced system design delivers leadership performance from the scalable processor power and high-speed system bus complemented by the large memory storage capacity and high-speed 64-bit PCIs.

- **High I/O Bandwidth/Low Memory Latency**

The system bus bandwidth is 1.1 Gbytes/sec with a memory latency of 120 ns.

- **Modular Design**

The basic building block is the system drawer, which fits in both enclosures, and the same CPU and memory modules are used within the drawer. Such modularity provides for easy upgrades and easy system maintenance.

- **Reliability and Availability**

AlphaServer 4000/4100 products achieve an unparalleled level of reliability and availability through the careful application of technologies that balance redundancy, error correction, and fault management. The balance between simple error detection and error correction provides the highest possible availability for the lowest possible cost.

System Packages

The AlphaServer 4000 and 4100 systems are available in packages to suit a wide range of computing needs. Systems are built around the system drawer, designed to fit in an industry-standard 19" rack or in a pedestal. An entry-level configuration can be upgraded by adding more CPU modules and more memory to an existing system, or the processors can be turned in for higher speed processors, and higher capacity memory modules can replace the original memory modules. The system drawer can even be lifted out and moved to a different enclosure.

The basic building block of both the pedestal system and the cabinet system is the system drawer. It not only holds the system bus and the PCI I/O buses but all power and cooling. Each drawer has a CD-ROM drive and a floppy diskette drive. When the drawer is in a cabinet system, the control panel is at the bottom of the system drawer. In a pedestal system the control panel is in a tray above the system drawer.

The BA30A system drawer used for the 4100 supports up to four CPUs, 8 slots of industry-standard PCI and EISA I/O (three slots of which can be either PCI or EISA), and up to 8 Gbytes of memory.

The BA30B system drawer used for the 4000 supports up to two CPUs, 16 slots of industry-standard PCI and EISA I/O (three slots of which can be either PCI or EISA), and up to 4 Gbytes of memory.

One to three power supplies are contained in the system drawer. An optional second and third power supply may be added to provide a full N+1 power configuration. Three fans in the drawer plus a fan on each CPU provide cooling for an open office environment.

4100 System Drawer

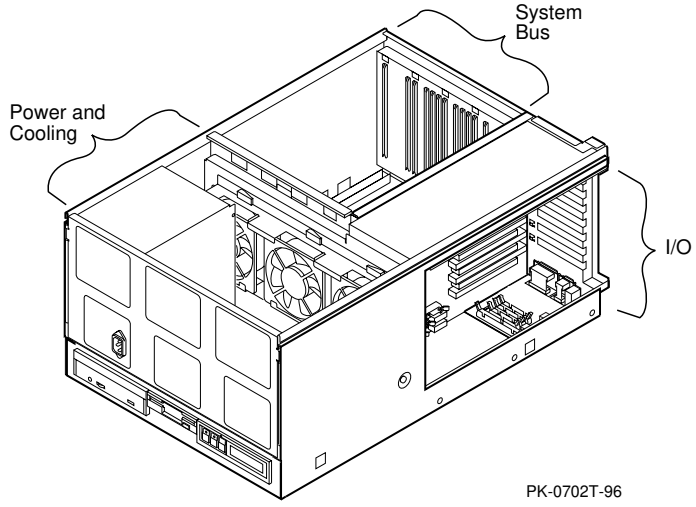

4000 System Drawer

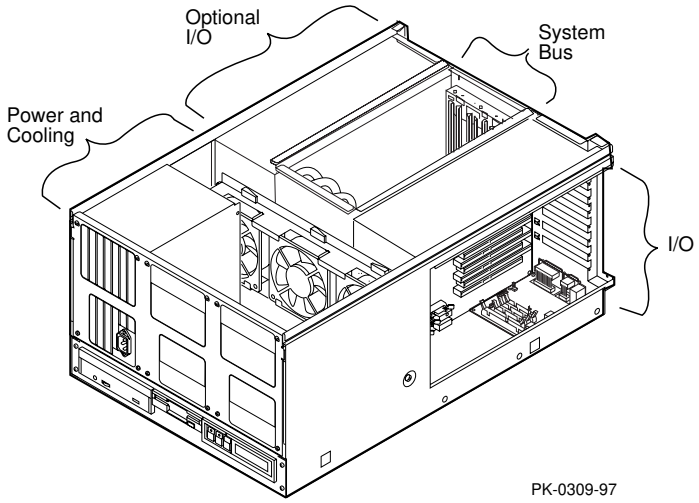

Pedestal System

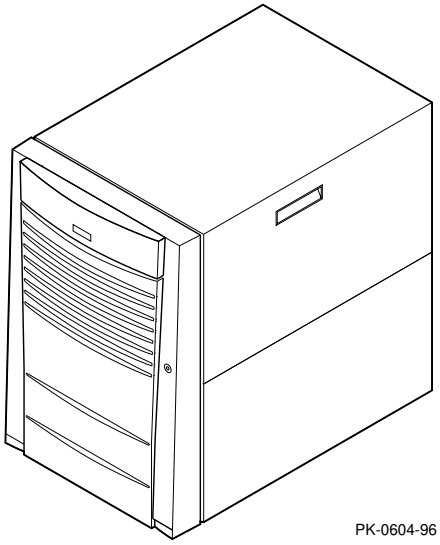

A low-cost entry SMP server, the pedestal system is designed with a small footprint for the open office environment. It is ideal for growing businesses that are ready to move up to a server that gives high productivity and cost-effective solutions. The pedestal has one system drawer and space for three StorageWorks shelves. The control panel area, above the system drawer, has space for an optional tape or CD-ROM drive. Three StorageWorks shelves provide up to 90 Gbytes of storage (21 4.3-Gbyte disks) in the pedestal.

Cabinet System

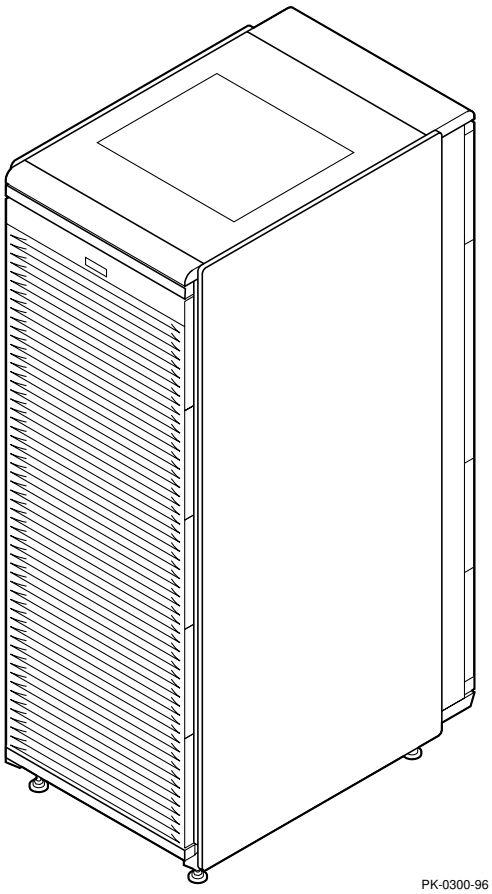

Up to four system drawers can be configured in a 19-inch-wide, industry-standard RETMA cabinet. The cabinet system has expanded storage capability, allowing a single system drawer to be in a 67-inch cabinet with up to eight StorageWorks shelves of disk storage. Each shelf supports 7 disk drives. Therefore, a cabinet system can have up to 56 disk drives. With 4.3 Gbyte drives, that's 240 Gbytes of in-cabinet storage.

The cabinet systems are designed for areas where floor space is at a premium, where large disk storage arrays and other expansion are required.

The cabinet comes with its own fans and two power supplies, with LEDs to indicate power and fan faults. Extensive testing has been conducted on the cabinet to check the shock and vibration levels that systems can sustain.

Architecture

These AlphaServer systems were designed to be modular and provide scalability. The kernel of the system is the system drawer, components of which are shown in the block diagram. The key component of each drawer is the system motherboard, which provides the system bus connecting multiple CPUs, multiple memory pairs, and bridge modules to the I/O subsystems.

The system architecture design optimizes for low memory latency (120 ns) and high delivered bandwidth (1.1 Gbytes/sec). This is accomplished by taking advantage of emerging synchronous memory technology. The I/O subsystem complements the memory subsystem; it too was designed for low target latencies. For programmed I/O and DMA writes the latencies are the same or better as what one might expect from a high-performance workstation.

CPUs

Multiple CPU modules are supported by the architecture. Modules with Alpha microprocessors must be of the same type and speed.

Memory

The memory modules must be in matched pairs (same size and same type of memory), with the largest pair in MEM 0 Low and MEM 0 High slots. Hardware detects the size of the memory in MEM 0L and 0H and then sets its memory decode logic to treat all memory pairs as being the same size as the pair in MEM 0L and 0H. Different memory sizes can be mixed, because all memory modules are assumed to be the same size, and the memory holes that result are treated as bad pages. All the operating systems have built-in support for mapping out bad pages of memory.

I/O

The system bus to PCI bus bridge modules plug into the system bus and provide the physical as well as the logical bridge between the system bus on the motherboard and the PCI motherboards. Each bridge module implements the PCI bus bridge logic for two independent 64-bit PCI buses. Only the first (PCI 0) PCI bus provides an interface to an EISA bus. The bus options are installed on the PCI motherboards.

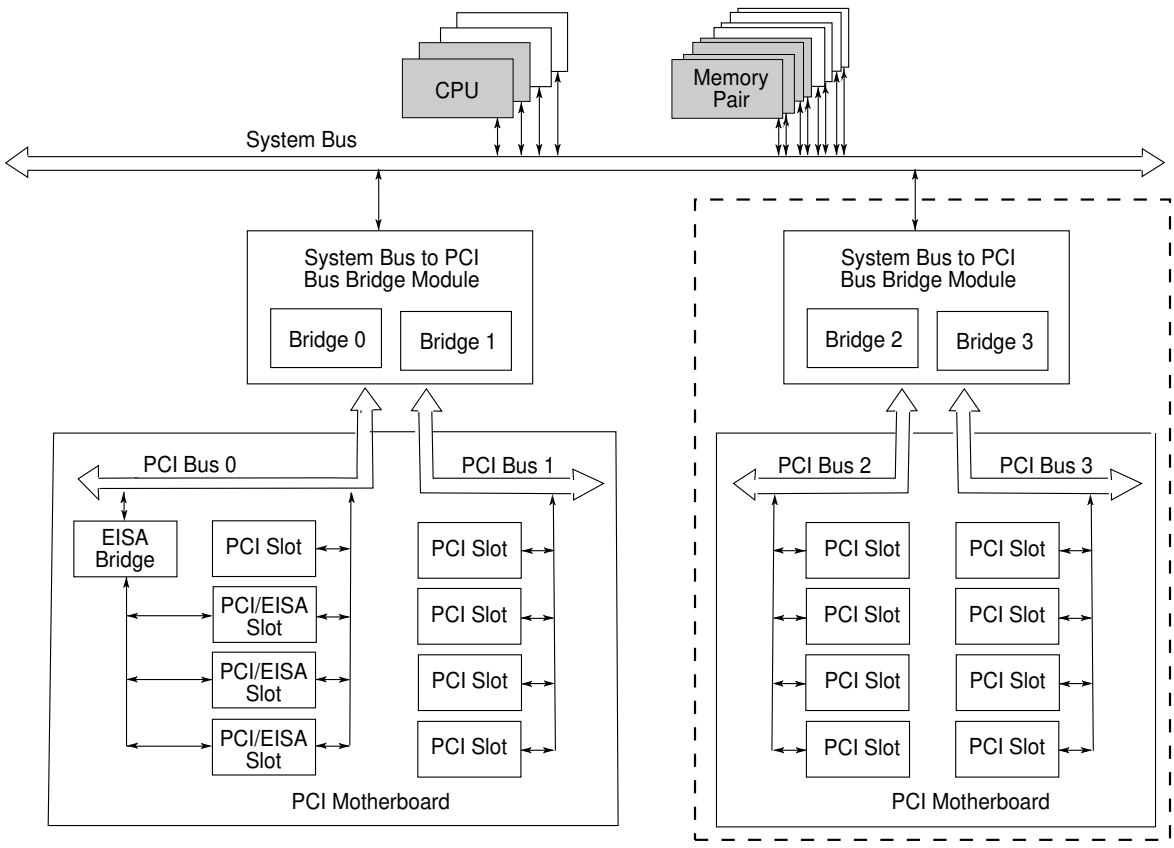

System Bus

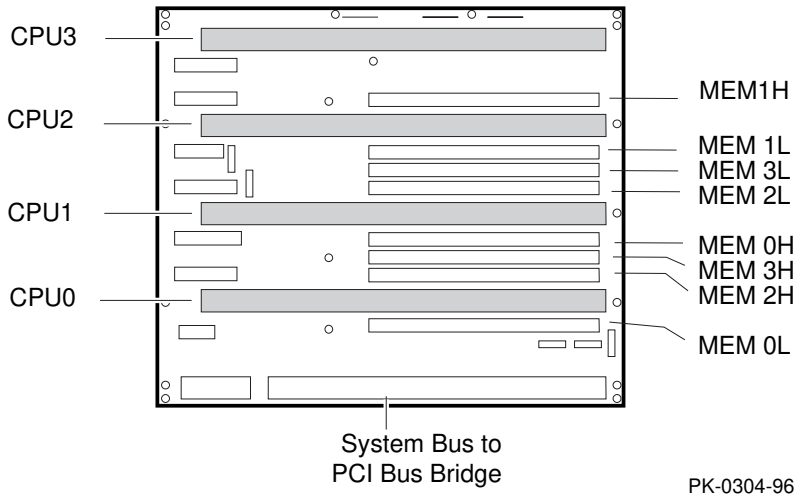

AlphaServer 4100 Motherboard

The system bus is the primary interconnect between CPU, memory, and the I/O subsystems. The AlphaServer 4100 system bus has slots for up to 4 CPUs and 4 memory module pairs, and an I/O slot that connects to the I/O subsystem. The bus consists of a 128-bit data bus and a separate 40-bit command/address bus. To accommodate a second PCI card cage, the AlphaServer 4000 motherboard is narrower than the 4100 motherboard. Slots are provided for 2 CPUs and 2 memory module pairs, and there is a second I/O slot.

The design goal for this system bus was to reduce memory latency. The bus is synchronously tied to the clock rate of the processor and memory.

As the clock speed of the processor increases, the bandwidth of the bus increases. With a 300 MHz processor, the bus bandwidth is 1 Gbyte/sec, but with a 400 or 466 MHz processor the bus bandwidth increases to 1.1 Gbyte/sec. The peak memory read bandwidth is 948 Mbytes/sec, and the peak for memory writes is 853 Mbytes/sec. The sustained mixed memory read/write bandwidth is 750 Mbytes/sec.

The bus is completely ECC protected, so that all memory, processor caches, and data are constantly checked for data integrity. Commands and addresses are parity protected.

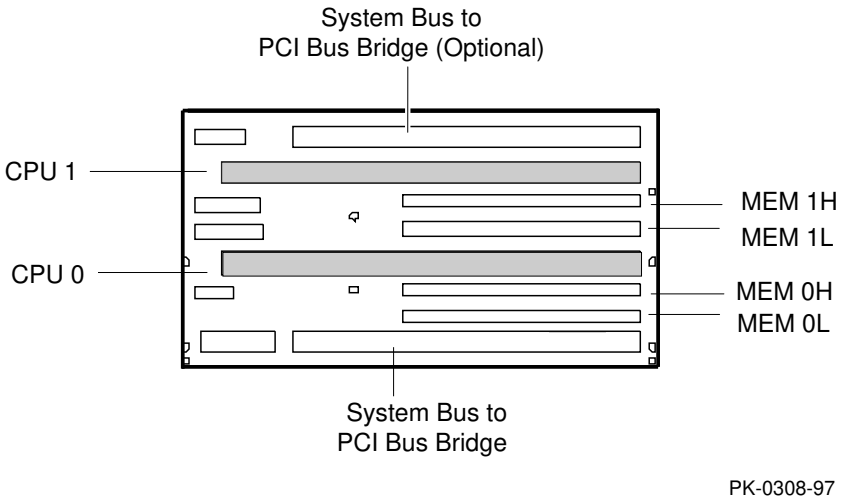

AlphaServer 4000 Motherboard

Processor Module

The AlphaServer 4000/4100 systems use an Alpha microprocessor on the processor module. The microprocessor is a superscalar (quad-instruction issue), superpipelined implementation of the Alpha architecture. The Alpha 21164 chips are manufactured using DIGITAL's state of the art CMOS-6 process, using a feature size of 0.35 micron. Over 9 million transistors are on one die. As new generations of Alpha chips become available, you can upgrade your system through a simple, cost-effective processor board swap.

Alpha Microprocessor Features

Alpha microprocessors have the following features:

- All instructions are 32 bits long and have a regular instruction format
- Floating-point unit, supports DIGITAL and IEEE floating-point data types
- 32 integer registers, 64 bits wide
- 32 floating-point registers, 64 bits wide
- On-chip, direct mapped, write-through physical data cache
- On-chip, direct mapped, read-only virtual instruction cache
- On-chip I-stream translation buffer
- On-chip D-stream translation buffer

Each chip has an 8-Kbyte instruction cache, an 8-Kbyte data cache, and a 96-Kbyte three-way set associative write-back second-level cache.

Processor Module Variants

- 300 MHz Alpha microprocessor with a 2-Mbyte cache
- 400 MHz Alpha microprocessor with a 4-Mbyte cache
- 466 MHz Alpha microprocessor with a 4-Mbyte cache

The 2-Mbyte cached module (shown below) has 16 STRAMs, 8 on each side.

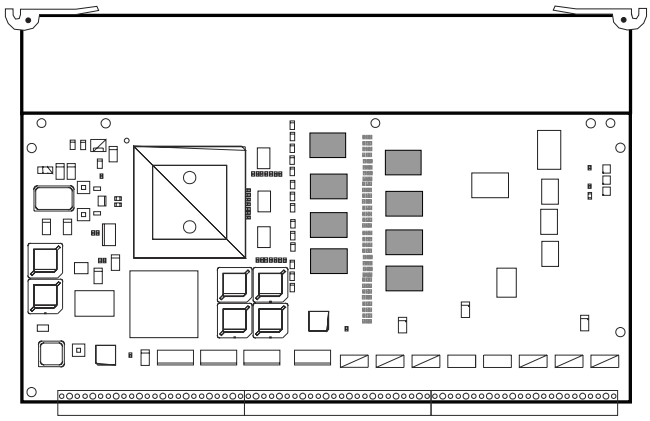

PK0422T-96

Processor Configuration Rules

- All CPUs in a system drawer must be of the same type.
- CPU0 slot must be populated, as the clock is associated with this slot.

Cooling of the microprocessor chip is provided by a fan on the Alpha microprocessor heatsink, as shown here. Space in the system card cage is optimized. The CPU modules with the fans extend beyond the shorter memory modules, making the most effective use of space in the card cage.

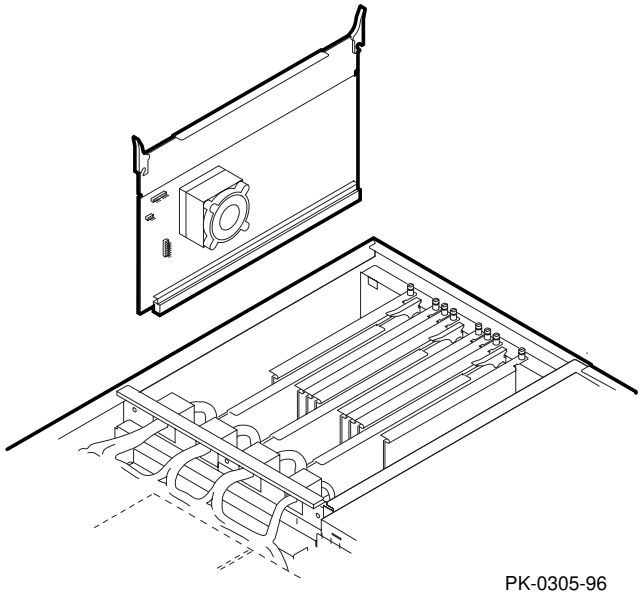

PK-0305-96

Memory

With an increase in the size of system memory comes huge gains in system performance. Since data once stored on disks can now be loaded into memory, applications can run much faster. Data warehousing and decision support require VLM (very large memory) technology that now is available in a server of this size.

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

- Very low memory latency (120 ns) with 15 ns system bus
- Numerous reliability and availability features, such as ECC memory and command/address parity

Memory Options

Each memory option consists of two identical modules. Each drawer supports up to either two or four memory options, for a total of 4 or 8 gigabytes of memory. Memory options are supported in the following sizes:

- 128 Mbytes
- 512 Mbytes
- 1 Gbyte
- 2 Gbytes

The 128-Mbyte option is implemented with synchronous DRAM technology, while the larger sizes are EDO technology.

Memory Technology

The synchronous memory used on the 128-Mbyte option is new to the industry. Synchronous DRAMs afford high performance, so even with the smallest amount of memory, the full memory bandwidth is available. For reads the bandwidth is 948 MB/sec, and for writes 853 MB/sec. All interleaving is internal to the DRAM.

The asynchronous EDO memory is a refinement of earlier DRAM technology. The bandwidth of the asynchronous memory requires one extra clock tick of latency, which results in only a 5 to 6 percent difference from the bandwidth attained with synchonous DRAMs. With EDO DRAMs the bandwidth is 853 MB/sec for both reads and writes. The highest density memory modules are implemented using the 64-megabit EDO DRAMs.

Memory Configuration Rules

- Memory modules must be in pairs.
- The largest memory modules must be in the lowest numbered slots.
- Slots must be populated in order (0, 1, 2, 3).

System I/O

Each PCI motherboard provides two separate 64-bit PCI buses. Industry-standard PCI and EISA I/O buses allow you to use inexpensive, widely available I/O options. Both 32-bit and 64-bit PCI options can be used.

PCI Bus

For options that require high performance, the systems implement a 64-bit PCI (peripheral component interconnect) bus. 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. DIGITAL is the first to offer the new 64-bit PCI in its servers.

The PCI bus implementation has the following characteristics:

- Fully compliant with the PCI Version 2.1 Specification
- Operates at 30 ns, delivering a peak bandwidth of 1 GB/sec; over 250 Mbytes/sec for each PCI bus
- Supports 16 option slots (with the I/O upgrade for the 4000)
- Supports peer-to-peer I/O operations
- Supports three address spaces: PCI I/O, PCI memory, and PCI configuration space
- Supports byteword, tri-byte, and longword operations
- Exists in noncached address space only

EISA Bus

The EISA bus is provided for compatibility with other AlphaServer systems. The EISA bus implementation has the following characteristics:

- Three EISA slots.
- 8.33 MHz operating speed, delivering a peak bandwidth of 33 MB/sec with no wait states.
- All slots are bus master slots.
- A two-device chipset provides a bridge from the PCI bus to EISA.

Optional PCI Card Cage

The H7150 option can be installed in the 4000 drawer to add two more PCI buses. This option consists of a system bus to PCI bus bridge module, a PCI motherboard, which has 8 PCI slots, and a power supply (required with two PCI card cages).

I/O Implementation

Refer to the figure in the Architecture section to see differences between the 4000 and 4100 systems.

The first PCI card cage holds the PCI motherboard, which interfaces to the system bus through the PCI bridge module. The PCI motherboard has a separate PCI bus on each side of the module. On the one side is PCI bus 1, which provides four PCI slots and an integrated SCSI controller that supports the system CD-ROM.

On the other side of the PCI motherboard is PCI bus 0. It too offers four PCI slots, but it provides much more.

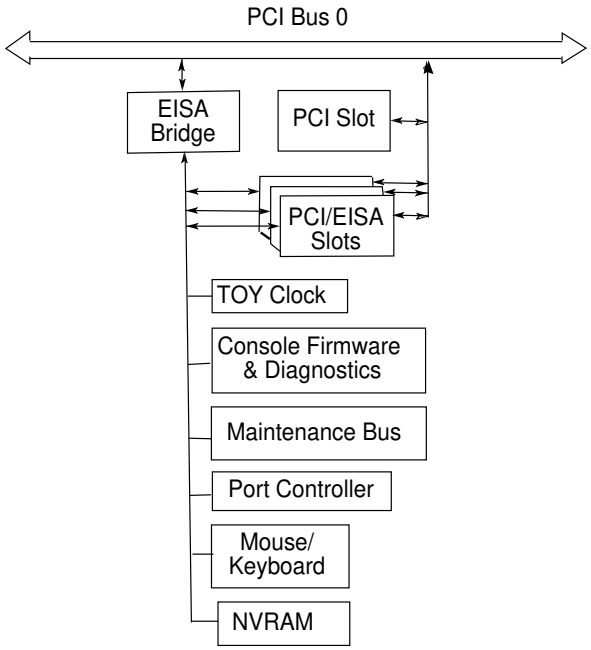

PK-0313-97

Three of those slots can instead be used for EISA options. The EISA bus is used to implement a number of system functions. Note that the system console firmware and diagnostics, which use the EISA bus, are located on the PCI motherboard. The location of this code therefore requires that much of the system hardware be functional.

Already installed in three slots of PCI 0 are the following:

- Ethernet controller
- SCSI controller
- Graphics adapter

Five slots for PCI options remain available in the first PCI card cage.

At the rear of the system drawer are connectors offering access to two serial communications ports, one parallel port, a modem connection, and ports for the keyboard and mouse.

I/O Bulkhead

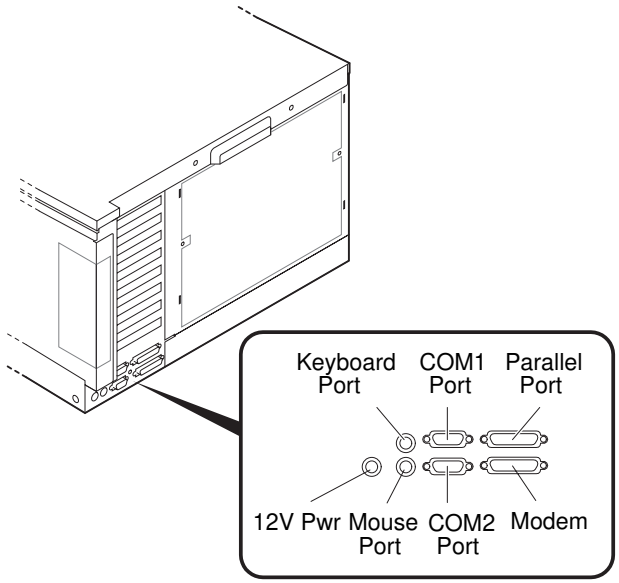

PK-0624T-96

Storage Architecture

The AlphaServer 4000/4100 products feature a modular storage architecture that is designed for maximum performance and easy operation and maintenance. All StorageWorks shelves are 16 bit (wide) shelves.

Pedestal Disk Storage

The pedestal system comes with one StorageWorks shelf and space for two more shelves. Each shelf supports up to seven 3.5-inch SCSI disk drives, so a pedestal can have up to 21 disks.

Cabinet Disk Storage

The cabinet can hold up to eight shelves for a total of 56 drives. The cabinet system supports other device formats (5.25-inch) and types (DSSI disks, SSD), which are available in the rackmountable configurations. A cabinet with only one system drawer can support up to 240 Gbytes storage.

RAID (Redundant Array of Independent Disks)

The systems can be configured with optional PCI RAID controllers to 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, 5
- Optional write cache
- Optional read cache
- Support for command queuing

Additional Expansion

The AlphaServer 4000/4100 products support optional external SCSI expansion, allowing large numbers of disks to be configured with the system. A system can support up to 7.5 TB in other enclosures with arrays (HSZ50s).

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, meaning that CPU, I/O, storage, and application resources can be added incrementally to efficiently grow 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.

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, DSSI, and MEMORY CHANNEL as the interconnect between disks and systems. OpenVMS systems can be configured into DSSI clusters using the following two options:

- The KFPSA DSSI adapter, which gives the system the capability of creating DSSI clusters.
- The HSD family of storage controllers.

Windows NT cluster systems use SCSI buses, Ethernet, or FDDI.

The primary means of clustering AlphaServer 4000/4100 systems depends on the operating system.

- CI clusters, OpenVMS only
- MEMORY CHANNEL, DIGITAL UNIX and OpenVMS
- SCSI clusters, DIGITAL UNIX and OpenVMS
- DSSI clusters, OpenVMS only

PCI to MEMORY CHANNEL™ Interconnect

Under DIGITAL 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 CCMAA 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 will result 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.

Reliability and Availability Features

The AlphaServer 4000/4100 products achieve 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

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

On multiprocessor systems, 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 the next higher working CPU, and the boot process continues. This feature ensures that a multiprocessor 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 status of failed components.

Memory Features

- The memory ECC scheme is designed to provide maximum protection for user data. The memory scheme corrects for single-bit errors, most double-bit errors, and total DRAM failure. It also detects numerous other types of failures, such as 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

- Parity protection on system bus, PCI, and SCSI buses.
- Extensive error correction built into disk drives.
- Optional internal RAID (redundant array of independent disks) improves reliability and data security.
- Disk hot swap on systems configured with RAID.

System Features

Auto reboot. On systems running DIGITAL 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.

The system has a built-in firmware update utility (LFU) that provides update capability for console and PCI I/O adapter firmware. A fail-safe loader provides a means of reaching 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. Three fans are in each system drawer, and each microprocessor has its own fan on the chip heatsink. If any of these fans fail, the system will shut down. In addition, the cabinet enclosure has fans installed at the top.

Error handling. Parity errors are detected on the high-speed system bus, the PCI and EISA buses, as well as in memory. Multiple ECC corrections to single-bit errors help in determining where in the system the error originated. Errors are logged for failure analysis.

Disk hot swap. The hardware is designed to enable hot swap of disks within optional RAID configurations. 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. Each system drawer has its own 450-watt power supply. One power supply is used for a system with one or two processor modules; a second power supply is needed for three or four processors and, in the 4000, with a second PCI card cage. A second power supply can be used to provide redundant power with one or two processors and one PCI card cage; a third power supply provides N+1 redundant power with three or four processors or with two PCI card cages.

Each StorageWorks shelf also has its own power supply, and a second power supply can be added to provide redundant power to the shelf. Doing this reduces the number of devices that can be installed in the shelf to six.

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

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. An integral remote console monitor (RCM) lets the administrator perform several tasks from a serial console: monitor the power supplies, temperature, and fans, and reset, halt, and power the system on or off, regardless of the operating system or hardware state. Also, the remote console monitor aids in inventory support by giving access to serial numbers and revisions of hardware and firmware.

These systems support all the management tools and features provided by the operating systems to manipulate and monitor system resources such as disks, printers, networks, and backups.

Platform Management

These 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.

Error Reporting

DECevent is a proprietary service tool that provides critical event translation and analysis for systems running the OpenVMS and DIGITAL UNIX operating systems. It provides the following functionality: translation (binary to text), reporting, analysis, notification, and graphical user interface. The analysis and notification portions of DECevent are protected functionality and require a Product Authorization Key (PAK); however, binary to text translations can be done without a PAK installed.

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.

Installation of the pedestal system is relatively simple, and may take only 30 minutes. The cabinet system installation requires two people and should be performed by qualified service technicians.

The modularity of the components contributes to ease of maintenance. The CPU, memory, PCI, and EISA options are plug-in cards that require no special switch or jumper settings. The internal mass storage elements set the SCSI IDs directly on installation. The power supplies are replaceable, modular units.

The cabinet system allows for easy access for service and maintenance. The shock and vibration rating is twice the industry standard, thus dedicated systems requirements can be easily met.

Performance

DIGITAL 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.

The AlphaServer 4000/4100 family offers record-breaking mid-range performance. The AlphaServer 4000 5/466 system is the price/performance leader for two-processor UNIX sys-tems with TPC-C results of 6,767 tpmC at $119 per tpmC using DIGITAL UNIX v4.1D4 and Sybase 11.1.

However, what really counts is application performance. These systems provide industry-leading application performance; for example:

- Lotus Notes—3,350 mail users at $31 per user on a two-processor AlphaServer 4100 5/466.
- Platinum Accounting—132,890 APVEs/hr, using AlphaServer 4100 5/300 with four processors, 512 Mbytes of memory and Windows NT.
- SPECweb96 Internet Server—1,700 operations per second, using an AlphaServer 4000 UNIX system with two processors, and 1,240 ops/sec with a single processor. And for real powerhouse performance, the AlphaServer 4100 5/466 with four processors delivers 2,304 ops/sec.

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.

See Table 1 for industry-standard benchmarks. 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

You can access performance information from DIGITAL using your fax machine as well as from several on-line sources.

- *InstaFACTS.* The InstaFACTS fax service delivers information directly to your fax machine. Call 1-800-DIGITAL (via a touch-tone phone in the U.S.A. and Canada) and 908-885-6426 (outside the U.S.A. and Canada). A catalog of documents is available from which you can order an abbreviated table of performance information, including DIGITAL performance briefs and flashes, TPC results, AIM results, and graphics results.
- *FTP.* Access performance documents from ftp://gatekeeper.dec.com/index.html. The directory is pub/DEC/DECinfo/performance/sys.
- *World Wide Web.* The document URL (Uniform Re-source Locator) is http://www.digital.com/info/performance.html.

Information for DIGITAL Partners

DIGITAL partners can access DIGITAL's Integrated Repository from DECGenisys V1.2. *DIGITAL Today, Business Partner Edition*, occasionally contains articles on performance of Alpha systems and announcements of available documents.

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

Service and Support

DIGITAL 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 DIGITAL Services, point your World Wide Web browser to http://www.service.digital.com/.

Hardware Warranty

The AlphaServer 4000 and 4100 system drawers and compo-nents installed in the drawers, including CPU, memory, PCI controllers, and power supplies, have a 3-year onsite, 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 onsite 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 DIGITAL 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.

AlphaServer 4000/4100 Configurations

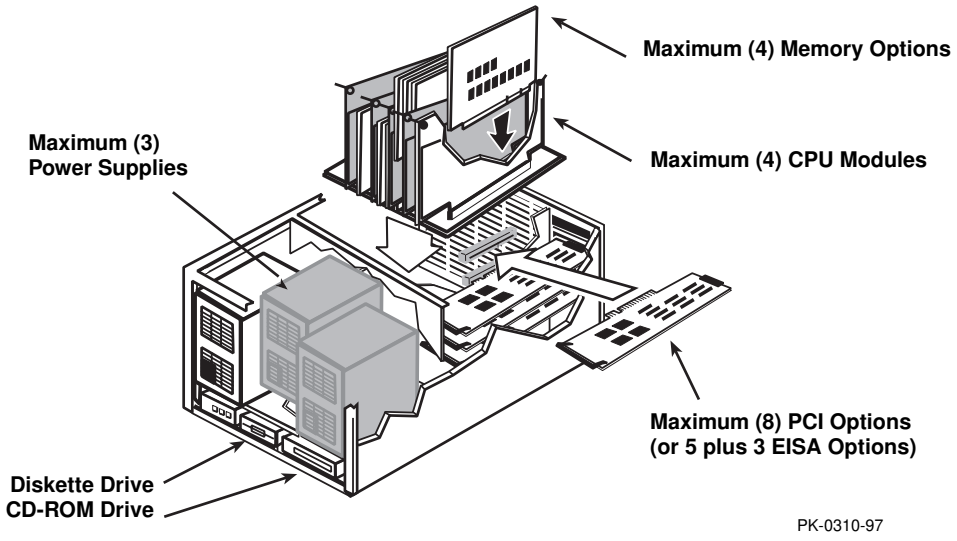

AlphaServer 4100 System Drawer

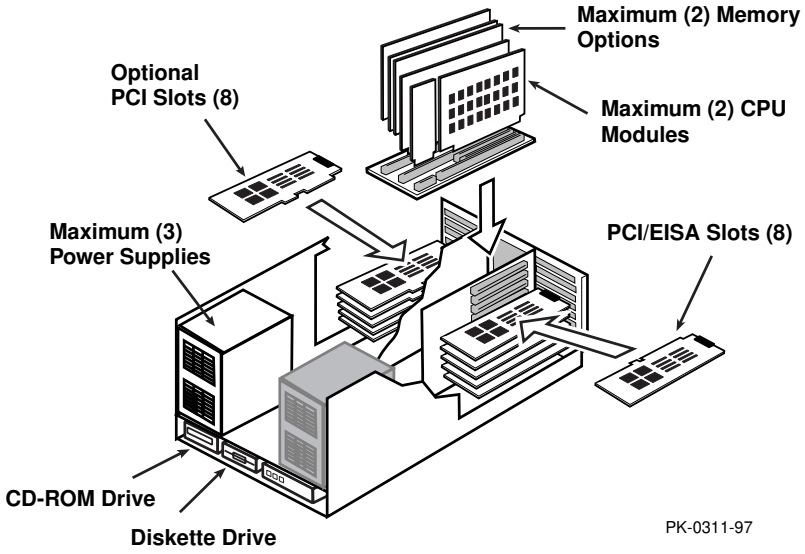

AlphaServer 4000 System Drawer

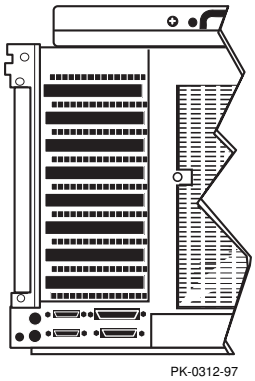

PCI Bulkhead

AlphaServer 4000/4100 Configurations

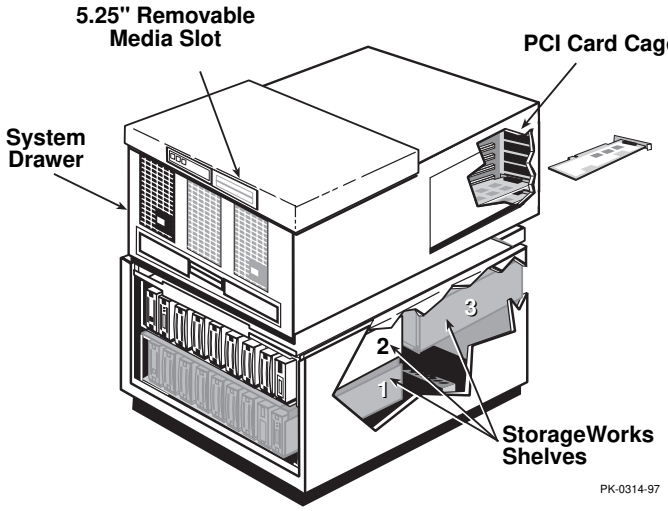

AlphaServer 4000/4100 Pedestal System

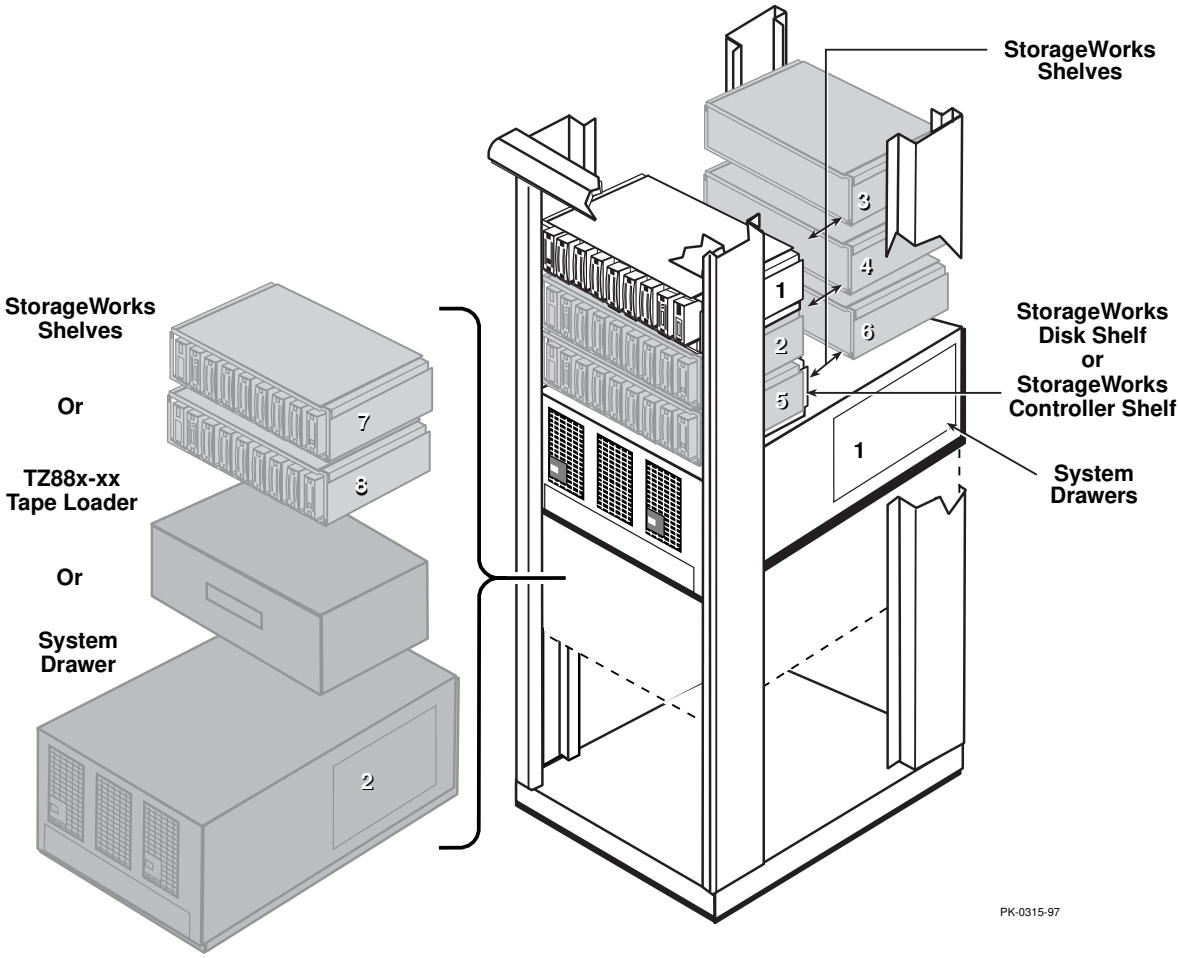

AlphaServer 4000/4100 Cabinet System

= Optional

System Features at a Glance

Table 1 provides a quick reference to features of the AlphaServer 4000/4100 systems.

Table 1 AlphaServer 4000/4100 Features

CPU Features	4000 5/300, 5/400, 5/466			4100 5/300, 5/400, 5/466		
Symmetric multiprocessing	1–2 processors			1–4 processors		
Processor	Alpha 21164			Alpha 21164		
CPU clock speed	300, 400, 466 MHz			300, 400, 466 MHz		
Cache on chip						
I-cache/D-cache	8 KB/8 KB			8 KB/8 KB		
Secondary cache	96 KB, 3-way set associative			96 KB, 3-way set associative		
On-board cache	2 MB, 4 MB, 4 MB			2 MB, 4 MB, 4 MB		
Upgradable in pedestal and cabinet	CPU, memory, I/O, storage			CPU, memory, storage		
Memory						
Memory (maximum)	4 GB			8 GB		
Performance	**5/300**	**5/400**	**5/466**	**5/300**	**5/400**	**5/466**
tpmC @ $/tpmC	-	-	6,767.80 @ $119 (2 CPU)[1]	-	-	-
SPECint95	8.11	12.1	14.1	8.11	12.1	14.1
SPECfp95 1 CPU	12.7	17.2	19.2	12.7	17.2	19.2
SPECfp95 SMP	16.7 (2 CPU)	24.2 (2 CPU)	26.5 (2 CPU)	21.8 (4 CPU)	33.4 (4 CPU)	36.1 (4 CPU)
SPECint_rate95	144 (2 CPU)	211 (2 CPU)	250 (2 CPU)	278 (4 CPU)	422 (4 CPU)	485 (4 CPU)
SPECfp_rate95	200 (2 CPU)	275 (2 CPU)	301 (2 CPU)	303 (4 CPU)	439 (4 CPU)	466 (4 CPU)
SPECweb96	-	1,240 (1 CPU)	1,700 (2 CPU)	-	-	2,304 (4 CPU)
LINPACK	775[2] (2 CPU)	1,019[2] (2 CPU)	1,119 (2 CPU)	1,380[2] (4 CPU)	1,841[2] (4 CPU)	1,934 (4 CPU)
Standard Features	1.44 MB diskette drive, CD-ROM drive, 10/100 Mbit Ethernet controller, FWSE SCSI-2 controller, integral FNSE SCSI-2 bus for removable media (CD-ROM and tape), S3 TRIO graphics adapter, 2 serial ports, 1 parallel port, keyboard and mouse, integral remote system console, operating system license and customer documentation, Internet software					
Internal Storage						
Removable media	CD-ROM, floppy diskette (pedestal has space for optional drive)					
Maximum StorageWorks shelves	3 (pedestal); 8 (cabinet)					
Maximum internal storage						
Pedestal	21 hot swap disks (90 GB)					
Cabinet	56 hot swap disks (240 GB)					
Total storage	7.5 TB					
I/O slots	8 PCI slots (3 of these can be used for EISA); 4000 can be upgraded to 16 PCI slots					
Maximum I/O throughput						
system throughput	1.1 GB/sec					
PCI	500 MB/sec (1.0 GB/sec when upgraded to 16 PCI slots)					
EISA	33 MB/sec					
Availability Features						
System	System auto reboot, thermal management, optional redundant power system (N+1), remote system management, RAID 0, 1, 0+1, 5, disk hot swap, memory failover, ECC memory, ECC cache, ECC system bus, SMP CPU failover, error logging, optional uninterruptible power supply					
OpenVMS clusters	CI, Ethernet, DSSI, SCSI, FDDI, PCI to MEMORY CHANNEL Interconnect					
UNIX TruClusters Solutions	SCSI, PCI to MEMORY CHANNEL Interconnect					
Windows NT cluster	Ethernet, FDDI, SCSI					
Operating Systems	DIGITAL UNIX, OpenVMS, Microsoft Windows NT Server					
Warranty						
Hardware	3-year, onsite					
Software	90-day telephone advisory support for OpenVMS and DIGITAL UNIX					

1. 6,767.80 using two 5/466 MHz CPUs, 4 GB, SYBASE v11.1, and DIGITAL UNIX v4.1 D4.
2. Linpack 1000x1000.

Physical Characteristics

Table 2 details basic physical characteristics of the system drawer and the pedestal and cabinet systems.

Table 2 AlphaServer 4000/4100 Physical Characteristics

Characteristic	System Drawer	Pedestal	Cabinet
Height	30 cm (11.8 in.)	75 cm (29.5 in.)	170 cm (67.0 in.)
Width	45 cm (17.7 in.)	49 cm (19.3 in.)	60 cm (23.6 in.)
Depth	69 cm (27.2 in.)	90 cm (35.4 in.)	90 cm (35.4 in.)
Weight	45.5 kg (100 lbs) fully configured	113.6 kg (250 lbs)	428.2 kg (932 lbs)
Power	Up to three 450 watt power supplies	Single phase	Single phase
4100 max. heat dissipation	1100 watts, 3754 Btu/hr	1800 watts, 6143 Bth/hr	2700 watts, 16,382 Btu/hr
4000 max. heat dissipation	700 watts, 3754 Btu/hr	1450 watts, 6143 Bth/hr	2700 watts, 16,382 Btu/hr
Operating temperature	10°C–35°C (50°F–95°F)	10°C–35°C (50°F–95°F)	10°C–35°C (50°F–95°F)
Operating humidity	20%–90% noncondensing	20%–90% noncondensing	20%–90% noncondensing
Acoustics, operating	5.8 LNPEc (Bels)	6.2 LNPEc (Bels)	6.7 LNPEc (Bels)

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

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

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

The following are trademarks of Digital Equipment Corporation: DIGITAL, the DIGITAL logo, AlphaServer, OpenVMS, ServerWORKS, StorageWorks, and TruCluster.

MEMORY CHANNEL is a trademark of Encore Computer Corporation. Microsoft is a registered trademark and Windows NT is a trademark of Microsoft Corporation. SPECint95, SPECfp95, SPECfp95 SMP, SPECint_rate95, SPECfp_rate95, and SPECweb96 are registered trademarks of the Standard Performance Evaluation Corporation. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company.

Copyright © 1997 Digital Equipment Corporation.
All rights reserved.