AlphaServer 8200/8400 Systems

Technical Summary

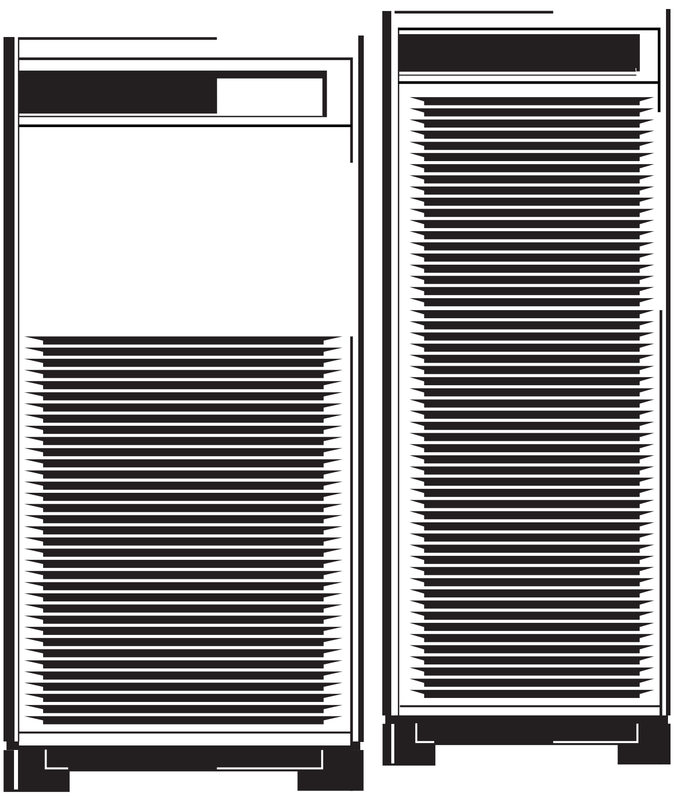

Contents

AlphaServer 8200/8400 Systems

The AlphaServer™ 8200 and AlphaServer 8400 products offer unprecedented performance, capacity, and reliability. Both systems use the same system bus, processor modules, memory modules, and I/O port modules.

Alpha processors of varying speeds (300, 350, and 440 MHz) handle applications that used to require larger, more costly systems. They provide unprecedented response times in transaction processing, database access, simulations, and file serving in an open computing environment. These are the most expandable servers on the market and can grow as your business needs increase.

More information on AlphaServer 8200/8400 systems is on the World Wide Web: http://www.digital.com/info/alphaserver/products.html.

System Overview

The AlphaServer 8200 and AlphaServer 8400 systems share the common core technologies, including the fastest CPUs, memory, and I/O in the industry. Both systems support multiple I/O channels that provide up to 1.2 Gbyte per second throughput using industry-standard PCI options. In addition, the AlphaSever 8400 also offers XMI I/O support for legacy devices.

Based on the Alpha 64-bit RISC architecture, these servers provide investment protection you can count on. You can choose either of two popular operating environments—OpenVMS or Digital UNIX®. These operating systems and thousands of leading applications run more efficiently on the Alpha platform.

AlphaServer 8200/8400 systems feature a 2.1 Gbyte per second system bus for access to multiple high-bandwidth I/O buses, very large memory capacities, up to 12 high-performance Alpha CPUs, and the reliability/availability/dependability features normally associated with mainframe systems.

Upgrades can be done from VAX/DEC 7000 machines to the AlphaServer 8400 by replacing the system card cage and installing new processor, memory, and I/O modules. In addition, in-cabinet speedup upgrades using the latest Alpha microprocessors, memory arrays, I/O systems, and operating systems have been offered and will continue to be made available on these new server products through the end of the century.

System Features

These systems provide several important features for fast application processing, high availability, low maintenance, and investment protection.

- **Unbeatable Price/Performance**

Symmetric multiprocessing with up to 12 CPUs is available now. With high-performance system and I/O buses, these servers can accommodate multiple processors without performance bottlenecks.

- **System Expansion**

Add CPUs and memory modules as your work expands. Multiple I/O system modules allow you to increase your storage capabilities. Both systems can function as Very Large Memory/Very Large Database systems. An AlphaServer 8400 supports up to 28 gigabytes of memory and 1.2 gigabytes/sec of I/O bandwidth.

- **Open Operating Systems**

Two industry-leading operating systems are supported today: Digital UNIX and OpenVMS Alpha. Such operating system support gives you a range of software development options and protects your investment in existing UNIX and OpenVMS applications.

- **Reliability and Availability**

Multiple power regulators ensure that the system keeps operating even when a power regulator fails. An integrated uninterruptible power supply is also available to protect the entire system. Power system monitoring is visible through LEDs and readable by system software. Multiple ECC checks provide for single-bit error correction to keep the system running and to provide for better failure isolation. Built-in self-tests execute on power-up and system reset.

- **Clusters**

Clustering, long available with OpenVMS systems, is now available with Digital UNIX systems. By connecting independent systems, you maximize system availability and performance and make the most of your hardware investment by sharing system resources.

- **Seamless Integration with Current Systems**

Digital's open systems, networking capabilities, and industry-standard operating environments ensure that your AlphaServer platforms work seamlessly with your existing VAX computers and AlphaGeneration products—as well as with virtually any other vendors' systems—for complete protection of your current and future hardware and software investments.

System Architecture

The AlphaServer 8200 and AlphaServer 8400 products share the same functional components. Figure 1 and Figure 2 show the architecture of these systems.

Figure 1 AlphaServer 8400 with KFTHA

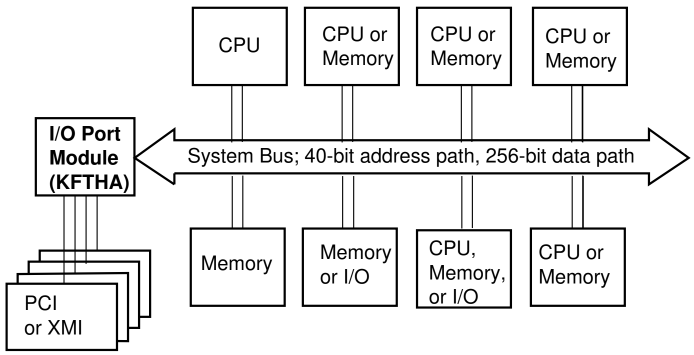

The AlphaServer 8400 has a nine-slot system bus. The minimum system configuration is one CPU module (with one or two CPUs), one memory module, and one I/O adapter. The I/O adapter shown here is the KFTHA. The remaining slots in the backplane can be used for CPU, memory, or KFTHA or KFTIA I/O adapters, within the following limits:

- Up to 6 CPU modules, for a maximum of 12 CPUs
- Up to 7 memory modules for a maximum of 28 Gbytes
- Up to 3 I/O adapter modules for a total of 12 I/O buses I/O buses can be the 12-slot PCI or 12-slot XMI, or a combination of PCI and XMI.

Figure 2 AlphaServer 8200 with KFTIA

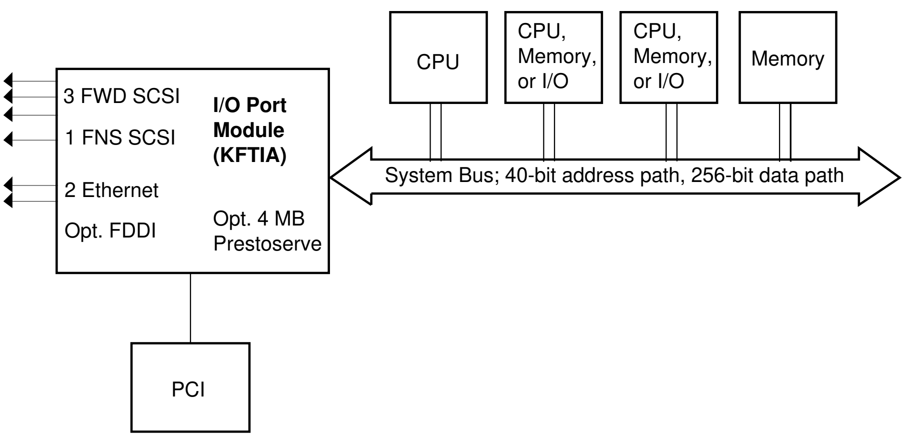

The AlphaServer 8200 has a five-slot system bus. The minimum system configuration is one CPU module (with one or two CPUs), one memory module, and one I/O adapter. The KFTIA I/O port module is shown in this block diagram. The remaining two slots in the backplane can be used for CPU, memory, or I/O adapters, either the KFTIA or the KFTHA.

Use of the KFTIA I/O module reduces system cost by providing a direct interconnect to SCSI devices, the Ethernet, and FDDI. Access to FDDI is provided by an optional daughter card on the KFTIA or through a PCI-to-FDDI adapter. An optional NVRAM (nonvolatile memory) daughter card can also be installed on the KFTIA module, which supports Prestoserve for UNIX NFS applications. The KFTIA also includes one channel for connecting to an optional PCI bus.

Expansion flexibility is possible within the following limits:

- Up to 3 CPU modules, for a maximum of 6 CPUs
- Up to 3 memory modules for a maximum of 12 Gbytes
- Up to 3 I/O adapter modules for a total of 9 PCI I/O buses (12-slot PCI bus).

Figure 3 AlphaServer 8400

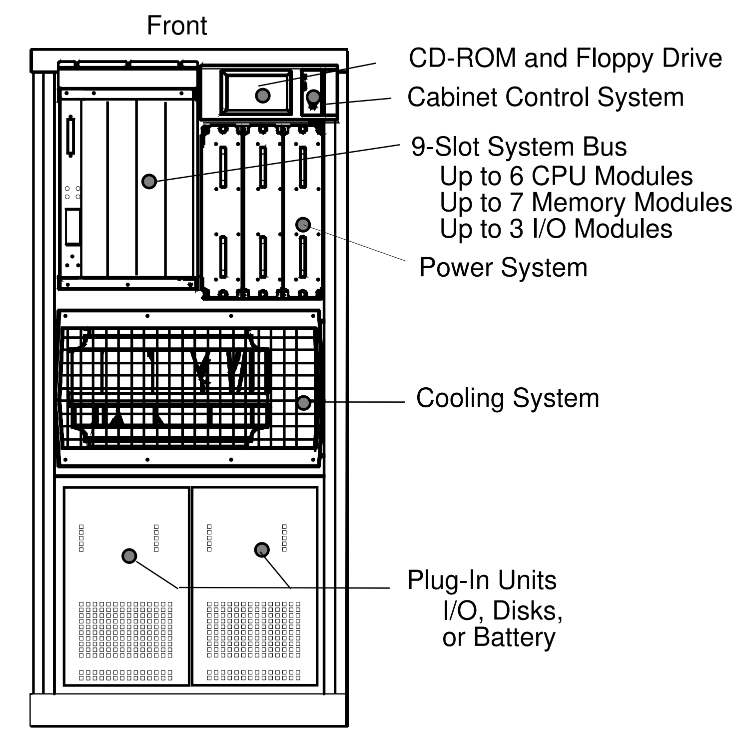

Figure 4 AlphaServer 8200

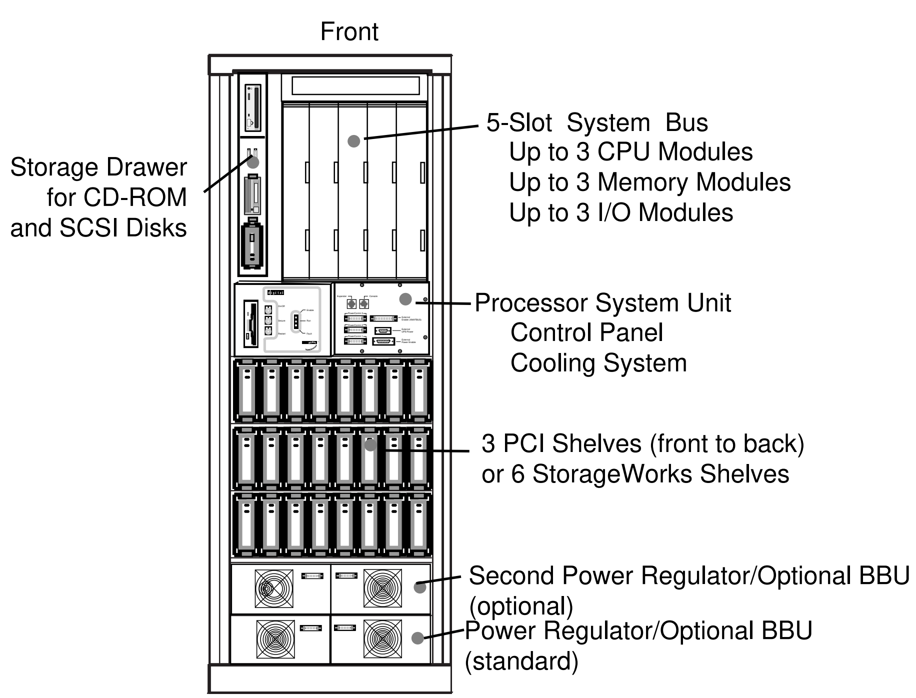

AlphaServer 8200/8400 Platform

AlphaServer 8200/8400 systems are Alpha-based high-performance SMP systems suitable for open office and data-center environments. Both systems feature a synchronous, high-speed system interconnect bus to increase system performance and lower memory latency.

The AlphaServer 8200 system is packaged in an industry-standard IEC 48D 19" rackmount cabinet enclosure with single-phase power. The system can be plugged into any standard 30 amp wall outlet without the need for special cooling requirements. It is designed for open office or satellite equipment installations where floor space is at a premium, but large disk storage densities or multiple PCI I/O are required; and laboratory/factory environments where rack-mounted application-specific equipment needs to be configured with the system.

The nine-slot AlphaServer 8400 system, designed for computer room installations, uses the same DEC/VAX 7000 cabinet, 3-phase power and cooling systems, and I/O options.

Up to two additional storage and I/O expansion cabinets can be configured with either system.

Upgrades from VAX/DEC 7000 Systems

The power, packaging, and I/O subsystem PIUs (plug-in units) of the VAX/DEC systems remain unchanged. To upgrade to an AlphaServer 8400 system, you simply change the centerplane (see Figure 5). Pull out the existing card cage and swap in the new AlphaServer system card cage. The new centerplane accommodates new processor modules, memory modules, and the new I/O port modules. Even the SIMMs from the 2-Gbyte memory modules used on VAX/DEC 7000 systems, the MS7AA-FA, can migrate over to the 2-Gbyte AlphaServer 8400 memory motherboard.

Figure 5 Simple Upgrade to AlphaServer 8400

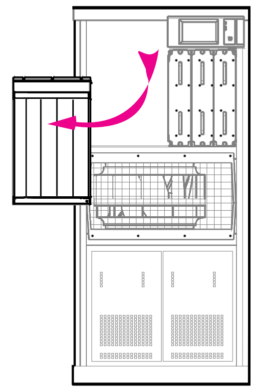

Table 1 Comparison of AlphaServer 8400/8200 System Features

Feature	8400 Cabinet	8200 Cabinet
CPUs	Up to 12 on 6 modules	Up to 6 on 3 modules
Memory	Up to 28 Gbytes on 7 modules	Up to 12 Gbytes on 3 modules
I/O slots Bandwidth	Up to 3 I/O modules (KFTHA) 1260 Mbytes/sec	Up to 3 I/O modules (1 KFTIA, 2 KFTHA) 945 Mbytes/sec
System bus With 300 MHz CPUs Peak bandwidth Sustainable	9 slots 2.1 Gbytes/sec 1.6 Gbytes/sec	5 slots 2.1 Gbytes/sec 1.6 Gbytes/sec
With 350/440 MHz CPUs Peak bandwidth Sustainable	 2.1 Gbytes/sec 1.87 Gbytes/sec	 2.1 Gbytes/sec 1.87 Gbytes/sec
Internal storage 3.5" disks 5.25" FH storage	 48 16	 40 12
Cabinets	1 system cabinet 0–2 expander Up to 2 battery cabinets	1 system cabinet 0–2 expander
Power	Three-phase Optional N+1 redundant power regulator	Single-phase Optional N+1 redundant power regulator
Battery backup	Optional	Optional
PIUs/shelves	PCI XMI SCSI Battery	PCI SCSI

Processor Module

The processor module can have one or two Alpha 21164 microprocessors on-board running at 300 MHz, 350 MHz, or 440 MHz (the module with the 440 MHz CPU is only available with two). Each microprocessor has its own independent data and address path and its own independent cache. Inside the chip is an 8-Kbyte instruction cache, an 8-Kbyte data cache, and a 96-Kbyte write-back second-level cache. A 4-Mbyte third-level cache is on the module for each CPU chip.

The latest Alpha 21164 chips are manufactured using Digital's state of the art CMOS-6 process. This process uses a feature size of 0.35 micron. The Alpha 21164 contains over 9 million transistors on one die.

An AlphaServer 8400 system can have up to 6 processor modules, for a total of 12 CPUs, while the AlphaServer 8200 can have up to 3 processor modules with up to 6 CPUs.

System Bus

The system bus used in the AlphaServer 8200/8400 systems is the fastest ever offered by Digital. This bus runs synchronously to the CPU chips at a sub-multiple of the CPU chip clock rate, and can operate up to 100 MHz (a 10 nanosecond cycle time). With the 350 and 440 MHz CPUs, the bus operates at 87.4 MHz to provide a bandwidth of 1.87 Gbytes/sec. With 300 MHz CPUs, the bus operates at 75 MHz, for a bandwidth of 1.6 Gbytes/sec.

The system bus has separate paths for the address and data. Data is now moved on a 256-bit bus, double the size of the DEC 7000 bus, which was 128 bits. In addition, there are 32 ECC bits. The command/address bus is a 40-bit bus.

The system bus is a synchronous bus. The address and commands on one bus are linked to the data on a separate bus by a sequence number. The sequence number guarantees that data is driven onto the data bus in the same order as the command/address are driven onto the address bus.

The greater performance of the system bus makes it possible to add more I/O port modules. The I/O bandwidth triples the bandwidth possible on the DEC 7000 system bus simply by adding I/O port modules. With up to 12 I/O channels, these systems now provide what amounts to a four-lane superhighway for I/O.

The AlphaServer 8400 system bus provides 9 slots for CPU, memory, and I/O modules. The AlphaServer 8200 system bus provides 5 slots.

Memory Options

Up to 28 Gbytes of main memory can be configured using "industry-available" single in-line memory module (SIMM) technology. Up to seven memory modules can be installed on the system bus in the 8400, and up to three in the 8200 system. Memory modules are available in these sizes: 128, 256, 512 Mbytes, 1 Gbyte, and 2 Gbytes, now, and 4 Gbytes in early 1997.

Today's standard memory module provides 800 Mbytes/sec memory bandwidth, independent of interleaving between modules. The 4-Gbyte module greatly increases the bandwidth provided. The 4-Gbyte module uses a new motherboard with 64-megabit SIMMs, which cannot be used on the other memory modules. However, the 4-Gbyte module can be used on the system bus alongside existing memory modules.

Interleaving

Each 2-Gbyte or less memory module supports onboard 2-way interleaving. With multiple memory modules installed, you get a minimum of 2-way interleaving and a maximum of 8-way interleaving, depending on the modules installed. The 4-Gbyte memory module is mode-selectable between 4 or 2 memory banks. The 4-Gbyte memory module supports onboard 4-way interleaving, so one 4-Gbyte memory module provides essentially the same memory bandwidth as two 2-Gbyte memory modules when they are 4-way interleaved.

Two "like-sized" memories will give you 4-way or 8-way interleaving (2-way or 4-way onboard each module times 2-way between modules). Modules of different densities can be interleaved together, provided certain rules are followed. For example:

1 X 128 MB = 128 MB	2-way interleaving
2 X 128 MB = 256 MB	4-way interleaving
2 X 128 MB + 256 MB = 512 MB	4-way interleaving
4 X 128 MB + 512 MB = 1 GB	4-way interleaving
4 X 128 MB = 512 MB	8-way interleaving
1 X 4 GB = 4 GB	4-way interleaving
4 X 4 GB = 16 GB	16-way interleaving
7 X 4 GB = 28 GB	16-way interleaving

The system console "looks" at the memory installed at power-up, determines the best possible interleave configuration, and then configures the memory.

High Bandwidth/High Performance

Design decisions relating to the memory modules were made in support of quality and speed, characteristics that undergird the reliability of these AlphaServer systems that depend upon Very Large Memory. And because all memory is shared, the investment in memory delivers better price/performance than that of competing systems in which memory is not shared.

I/O Architecture

These systems offer access to multiple high-bandwidth I/O buses: the Peripheral Component Interconnect (PCI) bus and Digital's XMI bus. See Table 2.

Table 2 I/O Capabilities

Subsystem	Mbytes/Sec	Slots Available
PCI	132	12 per bus
XMI	100	12 per bus

PCI I/O Subsystem

The PCI adapter provides connection to PCI devices and also to the EISA bus. The same 32-bit PCI module, the DWLPB, is used in both systems.

The PCI adapter is implemented electrically as three 4-slot PCI buses, but these appear logically to software as one 12-slot PCI bus, sharing the same address space. The PCI bus supports peer-to-peer PCI transactions (direct data transfer between two PCI modules), but only between PCI modules on the same electrical 4-slot segment.

The EISA bus (Extended Industry Standard Architecture) has 32 Mbyte/sec bus bandwidth, an 8 MHz 16-bit/32-bit data-path, with 8 slots. One EISA bus is supported.

The PCI adapter provides 12 option slots, plus a special slot for a bridge module. The bridge module is required to interface with the EISA bus. With a bridge module installed, the number of I/O slots is limited to 10: 2 EISA, 6 PCI/EISA, and 2 PCI slots.

XMI I/O Subsystem

The XMI bus options connect to the XMI I/O subsystem. An XMI I/O adapter resides on the XMI backplane and is the interface between the I/O port and the XMI I/O subsystem. One adapter is required for each 12-slot XMI I/O subsystem; up to six XMI I/O subsystems are supported on AlphaServer 8400 systems.

I/O Port Modules

The interface from the system bus to I/O is provided by two types of I/O adapter modules:

- KFTHA module (4 channels to external I/O)
- KFTIA module (1 channel to external I/O)

Up to three I/O adapter modules (KFTHA, KFTIA, or a combination) can be installed on the system bus.

These systems can handle large amounts of data at very high speeds. The I/O port module multiplexes that data between the high-speed I/O buses and the system bus. The I/O port module resides on the system bus and interfaces between it and the I/O subsystems.

The KFTIA module integrates the following:

- 3 fast wide differential SCSI ports
- 1 single-ended SCSI port for internal CD-ROM and tape
- 2 Ethernet ports (802.3 twisted-pair)
- 1 FDDI port (optional daughter card)
- 1 4-Mbyte NVRAM port (optional daughter card) for support of the Prestoserve option in UNIX configurations to accelerate NFS transactions
- 1 channel to PCI I/O

The KFTHA module provides four channels to I/O subsystems. For the AlphaServer 8400, bus adapters to the I/O subsystems are as follows:

- DWLPB — PCI
- DWLMA — XMI

For the AlphaServer 8200, only PCI I/O is supported, using the DWLPB PCI bus adapter.

Figure 6 KFTIA Module: One Channel with Integrated I/O

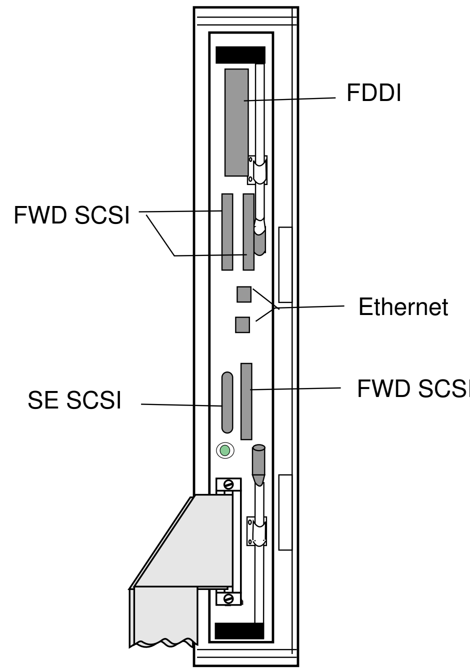

Figure 7 KFTHA Module with Four Channels

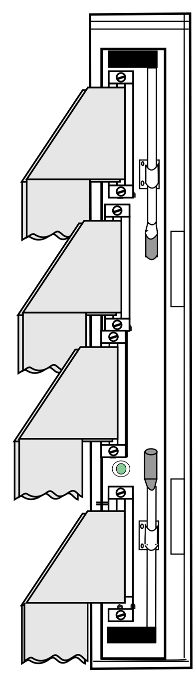

Clustering for UNIX and OpenVMS Environments

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.

Under OpenVMS, you can build CI, SCSI, DSSI, FDDI, MEMORY CHANNEL in November, and Ethernet-based clusters (for this class of system, we do not recommend Ethernet-based clusters) using the following hardware:

- CI clusters using the CIXCD on the XMI or the CIPCA on the PCI bus
  The HSJ controller is a bidirectional converter that takes CI signals in and outputs SCSI signals or takes in SCSI signals and outputs DSSI signals.
- SCSI clusters using the KZPSA on the PCI bus
- DSSI clusters using the KFMSB on the XMI or the KFPSA on the PCI bus
- FDDI clusters using one of the following:
  — KFTIA PCI-based FDDI daughter card
  — DEFPA PCI option
  — DEMFA XMI adapter
- MEMORY CHANNEL clusters using the CCMAA PCI adapter

The primary means of clustering AlphaServer 8200/8400 systems depends on the operating system.

- CI clusters, OpenVMS only
- MEMORY CHANNEL, Digital UNIX; OpenVMS in November
- SCSI clusters, Digital UNIX and OpenVMS
- DSSI clusters, OpenVMS only

PCI to MEMORY CHANNEL™ Interconnect

Under Digital UNIX, 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.

Storage Capabilities

With AlphaServer 8200/8400 systems, you can build enormous amounts of storage using the system cabinet and expander cabinets. Over 39 terabytes of storage can be configured in Digital UNIX configurations. OpenVMS Alpha configurations support over 20 terabytes—more than enough capacity for the largest data center applications.

StorageWorks building blocks enable you to configure the amount of storage you need. Each shelf can hold up to seven SCSI disks (typical configurations require that a DWZZA or DWZZB bus adapter be installed in each shelf, taking the place of one disk).

Up to eight StorageWorks shelves can be mounted in the AlphaServer 8400 system; in the AlphaServer 8200 system up to six StorageWorks shelves can be mounted back-to-back. In addition, in the AlphaServer 8200 system, you can use the optional integrated storage drawer to configure even more internal storage. The integrated storage drawer can hold up to four additional 3.5" hard disk drives plus two removable media devices (CD-ROM, tape drive). Therefore, in the system cabinets you can have:

- Up to 48 internal 3.5" disk drives (8400 system) (192 Gbytes total disk storage)
- Up to 40 internal 3.5" disk drives (8200 system) (160 Gbytes total disk storage)

See Figure 8 and Figure 9 for the storage available in each system cabinet and in an expander cabinet.

The AlphaServer 8200 and 8400 systems also support DSSI, which is available on VAX/DEC systems. DSSI subsystems are configured by using one of these optional devices:

- A PCI DSSI adapter, the KFPSA
- An XMI DSSI adapter, the KFMSB
- The HSDxx controller, which converts DSSI to single-ended SCSI. The HSDxx converts the StorageWorks bus to a DSSI bus.

The cabling for the DSSI controller is attached to one of the DSSI ports on the HSDxx. The other port can be terminated or fed to another DSSI system. The StorageWorks disks are a single node on the DSSI bus.

Figure 8 AlphaServer 8400 Storage

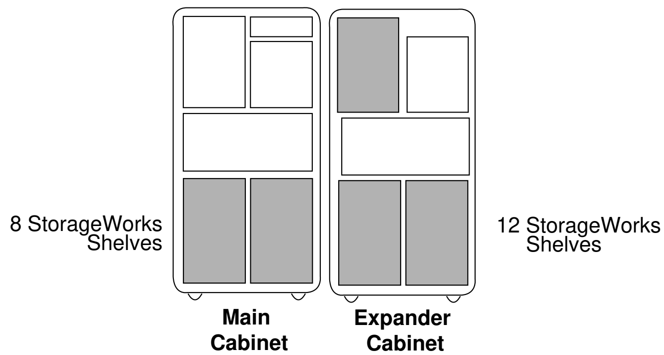

Figure 9 AlphaServer 8200 Storage

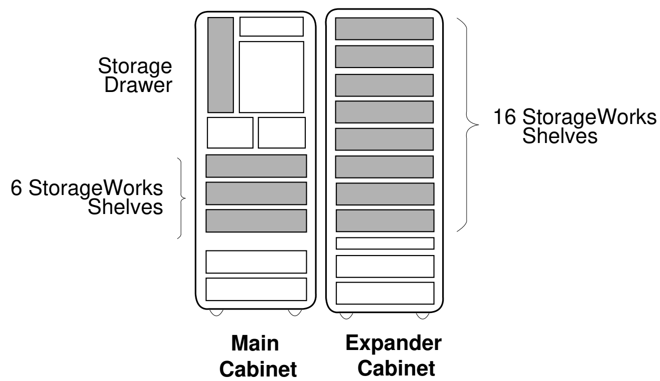

System Reliability, Availability, and Maintainability

AlphaServer 8200/8400 systems have numerous features that improve the reliability and availability of the system. The overall system reliability benefits from extensive use of CMOS technology in the design. The improvements are gained by having high circuit density, less interconnect and overall less heat dissipation than other technologies. Availability is improved by having more error detection and retry of error conditions.

System Features

- Built-in self-test and console ROM-based diagnostics at system level
- Console messages reflecting the status of booting
- Parity and error correction (ECC) on the system bus, all secondary caches, and memory
- Test-directed diagnostics and symptom-directed diagnostics
- System fault management
- Ease of repair
- Online system exercisers
- Modular power components

High Availability

AlphaServer 8200/8400 systems offer many features to address high availability:

- Clustering provides continuous availability, to storage and to computational abilities and applications, in spite of failure of a complete system. Digital invented clustering, first on VMS and OpenVMS systems, and now on Digital UNIX.
- With the UNIX operating system, TruCluster Available Server (DECsafe) software provides for application failover. By monitoring the performance of cluster members and automatically initiating recovery procedures in the event of system or component failures, a TruCluster system ensures system availability.
- RAID (redundant array of independent disks) may be deployed to enhance availability of storage. We offer RAID level 0 which improves performance by spreading I/O over a number of disks, RAID level 1 (a 1:1 high-performance redundancy technique), and RAID levels 3 and 5 (an N+1 redundancy technique) under OpenVMS and UNIX, in a variety of implementations: layered product software, internal controllers on the PCI I/O bus, and external controllers for clusters.
- Disks may be hot swapped to eliminate a source of system downtime. When used with RAID configurations, hot swapping is transparent to applications and to users. External RAID controllers may be configured to automatically replace failed disks in RAID sets with hot spare units.
- N+1 redundant power regulators decrease system failure rates. UPS eliminates downtime caused by external power outages. And, the wide allowable voltage range decreases sensitivity to brownouts.
- Systems with multiple CPU or memory modules automatically recover from failures of those modules, by rebooting to exclude those failed modules. Thus, a hard fault is transformed into a transient outage, followed by continued operation with degraded performance.
- Both UNIX and OpenVMS recover from errors in the I/O subsystem by a variety of retry schemes, including retrying failed transfers and reinitializing I/O adapters and controllers. Both UNIX and OpenVMS support extensive error logging and error reporting for I/O events.

Parity and Error Correction

Multiple ECC checks provide for better failure isolation. Each checkpoint within the system preserves error information, assisting software in determining where in the system the error originated.

Parity protection is used on the address bus, and on the data bus an 8-bit ECC check code protects each 64 bits of data. Single-bit errors are corrected. The ECC check code detects double-bit errors and some 4-bit errors in each 64 bits.

For optimal performance and integrity, the memory modules do not correct the data when single-bit errors are detected. Only CPU and I/O port modules correct single-bit ECC errors. Any errors are logged in the system error log, and the console program can then identify a failing SIMM, so that it can be replaced if the same error occurs repeatedly.

Diagnostics

Processor modules and I/O modules include an extensive self-test to verify their functionality. Testing is executed on system power-up and can be executed on every system boot if desired. The console also provides a command for users to execute the self-tests.

Since self-test is normally executed concurrently on most system components, extensive testing is done quickly. Typical UNIX configurations on any 8200/8400 system will boot up to the single-user prompt in 1 to 2 minutes, with another minute or so to get to the multi-user prompt. Very large memories increase testing times by approximately 30 seconds. Booting from remote storage servers, particularly booting into active OpenVMS clusters from an HSJ40, can increase boot time by another minute or two.

The Loadable Firmware Update Utility (LFU) is distributed on CD-ROM. It is used to check the revision of firmware on all modules and to upgrade firmware as new revisions are required.

Power

The AlphaServer 8200 has a single-phase power system, and the 8400 has three-phase. Because the three-phase regulators used in the 8400 are variants of those used in VAX/DEC 7000 systems, on-site upgrades from those systems require no change to the power system. Battery backup for both systems is optional; details are given in Table 3.

Table 3 Power Systems

Feature	Three-Phase 8400	Single-Phase 8200
N+1 regulators	Optional (3 max.)	Optional (2 max.)
Battery backup	Optional (60 mins.)	Optional (5 mins.)
Power monitoring (via console)	Yes	No

Both systems offer optional uninterruptible power supply (UPS) subsystems. In the AlphaServer 8200 system, batteries are installed in the power regulator, so components in each cabinet, whether the main cabinet or an expander cabinet, are backed up from the batteries in the cabinet power regulator. The AlphaServer 8400 system, on the other hand, has separate battery plug-in units (PIUs). These PIUs are installed in the system cabinet to back up components in that cabinet, or in a battery cabinet to back up components in an expander cabinet.

In both systems, a redundant power regulator guarantees system operation should one power regulator fail. With the H7263 power regulators in 8400 systems, heavily configured systems may require more than one regulator. A third regulator can be used as a backup.

Installation and Upgrades

AlphaServer 8200/8400 system hardware is not customer installable or upgradable. System installations and upgrades must be performed by qualified Digital customer service technicians. An upgrade to an AlphaServer 8400 system can be done in a few hours.

The operating systems are factory installed; upgrades and other software are customer installable.

Server Management

The AlphaServer products support important operational and platform management requirements.

Operational 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.

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. For example, system managers can use the POLYCENTER suite of tools to manage an enterprise-wide system. These tools are usable in a highly distributed environment.

Platform Management

The 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. Most console firmware features can be accessed remotely using the POLYCENTER console manager product.

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.

Performance Monitoring

A system monitoring tool called Monitoring Performance History (MPH) collects error log entries, crash dump footprints and configuration information from the monitored systems. The information is collected weekly and is sent back to Digital by either Internet mail or the Digital Services Network Link (DSNLink) transport mechanisms.

Performance and Benchmarking

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 or close CPU performance, coupled with comparable memory and disk expandability.

Industry-standard benchmarks run on AlphaServer 8200/8400 systems show that these systems deliver unsurpassed computing performance and price/performance. See Table 7 for details.

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 several online sources.

- *InstaFACTS*. The InstaFACTS fax service delivers information directly to your fax machine. Call 1-800-723-4431 (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's performance briefs and flashes, TPC results, AIM results, and graphic results.
- *FTP*. Access performance documents from gatekeeper.dec.com. The directory name is pub/DEC/DECinfo/performance/sys.
- *CompuServe*. Type GO VAXFORUM and look in the "hardware" library. For more information contact Doyle Myers at Internet address doyle@wrq.com or 76703.4403@compuserve.com.
- *World Wide Web*. The document URL (Uniform Resource 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.

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/.

Physical Characteristics and Operating Environment

Table 4 lists the physical characteristics and the operating environment for the systems.

Table 4 Physical and Environmental Specifications

Cabinet Type	AlphaServer 8400	AlphaServer 8200
Physical Dimensions		
Height	170.0 cm (67.0 in)	170.0 cm (67.0 in)
Width	80.0 cm (31.5 in)	60.0 cm (23.6 in)
Depth	87.5 cm (34.4 in)	92.5 cm (36.4 in)
Approx. weight		
without batteries	408 kg (900 lbs)	272 kg (600 lbs)
with batteries	545 kg (1200 lbs)	330 kg (728 lbs)
Service clearance, front	1.5 m (59 in)	1.0 m (40 in)
Service clearance, rear	1.0 m (40 in)	.75 m (29.5 in)
Environmental Requirements		
Temperature	15°–28°C (59°–82°F)	10°–35°C (50°–95°F)
Humidity	20–80%	10–90%

Power Requirements

Table 5 and Table 6 summarize the power requirements for the AlphaServer 8400 and AlphaServer 8200 systems.

Table 5 AlphaServer 8400 Power Requirements

Three-Phase	U.S./Canada	Europe/APA	Japan
Nominal voltage	120/208 V	380–415 V	202 V
Frequency range	50–60 Hz	50–60 Hz	50–60 Hz
Phases	3-phase star 4-wire N-GND	3-phase star 4-wire N-GND	3-phase delta 4-wire mid-GND or 3-wire junction GND
Max. input current/phase	24 A rms	12.8 A rms	24 A rms
Surge current	50 A peak	50 A peak	50 A peak
Rating	30 A	16 A	30 A

Table 6 AlphaServer 8200 Power Requirements

Single-Phase	U.S./Canada/Japan	Europe/APA
Nominal AC input line voltage	202–240 (208) V Japan (202) V	202–240 (240) V
Frequency range	50–60 Hz	50–60 Hz
Phases	Single-phase line-to-line or line-to-neutral	Single-phase line-to-line or line-to-neutral
Maximum input current	16 A rms	16 A rms
Surge current	80 A peak	80 A peak
Rating	16 A	16 A

System Features at a Glance

Table 7 provides a quick reference to the features of the AlphaServer 8200 and AlphaServer 8400 systems.

Table 7 System Features at a Glance

System Features	**AlphaServer 8400**			**AlphaServer 8200**		
CPU Features	**5/300**	**5/350**	**5/440**	**5/300**	**5/350**	**5/440**
Symmetric multiprocessing	Up to 12	Up to 12	Up to 12	Up to 6	Up to 6	Up to 6
CPU clock speed	300 MHz	350 MHz	440 MHz	300 MHz	350 MHz	440 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		
Cache on-board/per CPU	4 MB			4 MB		
Performance	**5/300**	**5/350**	**5/440**	**5/300**	**5/350**	**5/440**
SPECint95 (1 CPU)	7.43	10.1	13.6	7.43	10.1	13.6
SPECfp95 (1 CPU)	12.4	14.2	16.2	12.4	14.2	16.2
SPECfp95 SMP	–	38.5[1]	42.6[1]	–	31.4[2]	34.3[2]
SPECint_rate95	767[3]	1,004[3]	1,358[3]	388[2]	506[2]	701[2]
SPECfp_rate95	919[3]	1,039[3]	1,118[3]	420 [2]	505[2]	588[2]
tpmC @ $/tpmC	11,014 @ $222	14,176.61 @ $198.37[4]	–	–	7,426 @ $235[4]	–
tpc-D @100 GB						
QppD	–	–	864.3	–	–	–
QthD			445.4			
$/QphD			$1,863			
Internal Storage						
3.5-inch disks: system cabinet	48			40		
5.25-inch FH storage: system cabinet	16			12		
I/O Features						
Maximum I/O throughput						
system throughput	1.2 GB/sec			1.2 GB/sec		
subsystems	PCI: 132 MB/sec EISA: 33 MB/sec XMI: 100 MB/sec			PCI: 132 MB/sec EISA: 33 MB/sec		
PCI I/O bus	12 I/O slots per PCI bus Up to 12 PCI buses 144 PCI slots			12 I/O slots per PCI bus Up to 11 PCI buses 132 PCI slots		
EISA I/O bus	One 8-slot EISA bus			One 8-slot EISA bus		
XMI I/O bus	12 I/O slots per XMI bus Up to 6 XMI buses					
High Availability Features						
System	ECC on critical data and memory paths Built-in self-tests and system fault management Optional N+1 redundant power system Power and cooling system monitoring Optional uninterruptible power supply Disk hot swap Optional RAID levels 0, 1, 0+1, 3, and 5					
OpenVMS clusters	Yes			Yes		
TruCluster Solutions (UNIX)	Yes			Yes		
Warranty (hardware)	One-year, on-site			One-year, on site		

[1] Eight CPUs [2] Six CPUs [3] Twelve CPUs [4] Ten CPUs

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 errors in the information given in this publication.

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