# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

DEC 10000 Model 610 AXP Server

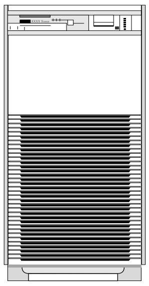

DEC 7000 Model 610 AXP Server

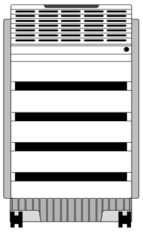

DEC 4000 Model 610 AXP Server

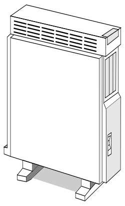

DEC 3000 Model 500S AXP Server

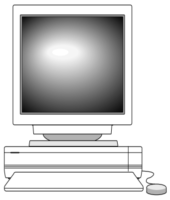

DEC 3000 Model 400S AXP Server

*INSIDE*

Benchmark results for:

- SPEC
- AIM
- Perfect
- Dhrystone
- LINPACK
- CERN
- DN&R Labs CPU2
- Livermore Loops
- FFT & Convolution
- SLALOM
- Basic Real-Time Primitives
- Rhealstone

Digital Equipment Corporation
March 1993
First Edition
EB-N0105-51

**First Printing, March 1993**

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Contents

## List of Figures



## List of Tables

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Digital's Alpha AXP Server Family

This document presents the results of industry-standard benchmarks on Digital's Alpha AXPTM servers in the DEC OSF/1™ AXPTM operating system environment. Alpha AXP servers are based on Digital's breakthrough 64-bit RISC architecture.

### DEC 3000 AXP Servers

The DEC 3000 Model 400S AXP desktop is an entry-level server or multiuser system that provides affordable power for small workgroup computing and distributed applications. The DEC 3000 Model 500S AXP server is a deskside system that provides a blend of performance and expansion to meet the challenges of the most demanding workgroup computing. The DEC 3000 Model 400S AXP server runs at a CPU clock speed of 133 MHz, and the DEC 3000 Model 500S AXP server at 150 MHz.

### DEC 4000 AXP Servers

DEC 4000 AXP servers deliver superior computing power in office-sized or rackmountable packages. Designed for symmetric multiprocessing, multiuser and technical server applications, DEC 4000 AXP servers meet the needs of technical and scientific markets. Optimized for speed and availability, DEC 4000 AXP servers will satisfy the commercial segment as well. The DEC 4000 AXP server runs at a CPU clock speed of 160 MHz and can be configured with one or two CPU boards.

### DEC 7000 AXP Servers

DEC 7000 AXP servers deliver unsurpassed data center performance and price/performance. DEC 7000 AXP servers offer single or multiprocessing capability for such commercial and technical applications as transaction processing, general ledger, securities trading, signal processing, molecular modeling, and imaging. The DEC 7000 AXP server runs at a CPU clock speed of 182 MHz and can be configured with up to six processors.

### DEC 10000 AXP Servers

DEC 10000 AXP servers are Digital's highest performance Alpha AXP systems. Designed to meet rigorous demands of compute-intensive, large enterprise applications, DEC 10000 AXP servers can be configured with up to six processors. The DEC 10000 AXP server provides the fastest enterprise server performance in the industry—at a fraction of the price of a traditional mainframe or supercomputer. The DEC 10000 AXP server runs at a CPU clock speed of 200 MHz.

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

**Table 1 Digital's Alpha AXP Server Family Benchmark Results**

| Benchmark | DEC 3000 Model 400S AXP | DEC 3000 Model 500S AXP | DEC 4000 Model 610 AXP | DEC 7000 Model 610 AXP | DEC 10000 Model 610 AXP |
|---|---|---|---|---|---|
| SPECint92 | 74.8 | 84.4 | 94.6 | 103.1 | 116.5 |
| SPECfp92 | 112.5 | 127.7 | 137.6 | 176.0 | 193.6 |
| SPECrate_int92 | 1,763 | 1,984 | 2,198 | 2,572 | 2,765 |
| SPECrate_fp92 | 2,661 | 3,023 | 3,247 | 4,179 | 4,368 |
| SPECmark89 | 111.1 | 126.1 | 137.3 | 175.5 | 192.1 |
| AIM Benchmark Suites | | | | | |
| Performance Rating | 70.3 | 82.9 | tbd | tbd | tbd |
| Maximum User Load | 485 | 649 | tbd | tbd | tbd |
| Maximum Throughput | 688.7 | 812.9 | tbd | tbd | tbd |
| Perfect Benchmarks Suite (MFLOPS-geometric mean) | 18.4 | 20.7 | 23.1 | 26.4 | 29.2 |
| Dhrystone | | | | | |
| V1.1 (instructions/sec) | 235,939 | 266,487 | 297,345 | 330,577 | 363,743 |
| V2.1 (instructions/sec) | 238,095 | 263,157 | 294,117 | 333,333 | 357,142 |
| LINPACK 64-bit Double-precision | | | | | |
| 100x100 (MFLOPS) | 26.0 | 29.6 | 35.0 | 36.9 | 40.5 |
| 1000x1000 (MFLOPS) | 91.7 | 103.5 | 110.1 | 137.8 | 151.1 |
| CERN Benchmark Suite | 18.8 | 21.3 | 23.2 | 26 | 29 |
| DN&R Labs CPU2 (MVUPs) | 185.0 | 209.1 | 225.6 | 254.6 | 279.3 |
| Livermore Loops (geometric mean) | 17.4 | 19.5 | 22.3 | 25.4 | 27.8 |
| SLALOM (patches) | 5,776 | 6,084 | 6,496 | 6,902 | 7,248 |

tbd = to be determined

***Note:*** *Digital will update this performance brief to include SPEC SFS 1.0 (097.LADDIS) benchmark results.*

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

**Benchmark Descriptions and Results:**

- **SPEC Benchmark Suites**
- **AIM Benchmark Suites**
- **Perfect Benchmark Suite**
- **Dhrystone**
- **LINPACK 100x100 and 1000x1000**
- **CERN Benchmark Suite**
- **DN&R Labs CPU2**
- **Livermore Loops**
- **FFT & Convolution**
- **SLALOM**
- **Basic Real-Time Primitives**
- **Rhealstone Benchmark**

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## SPEC Benchmark Suites

SPEC™ (Standard Performance Evaluation Corporation) was formed to identify and create objective sets of applications-oriented tests, which can serve as common reference points and be used to evaluate performance across multiple vendors' platforms.

### SPEC CINT92 and CFP92

In January 1992, SPEC announced the availability of two new benchmark suites, CINT92 and CFP92. Each suite provides performance indicators for different market segments because each has different workload characteristics. SPEC CINT92 is a good base indicator of CPU performance in a commercial environment. SPEC CFP92 may be used to compare floating-point intensive environments, typically engineering and scientific applications.

#### SPEC CINT92

CINT92, the integer suite, contains six real-world application benchmarks written in C. The geometric mean of the suite's six SPECratios is the SPECint92™ figure. CINT92 suite includes the following application classes:

- 008.espresso–Circuit theory
- 022.li–LISP Interpreter
- 023.eqntott–Logic design
- 026.compress–Data compression
- 072.sc–UNIX™ spreadsheet
- 085.gcc–GNU C compiler

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Figure 1 SPEC CINT92 Results for Alpha AXP Servers

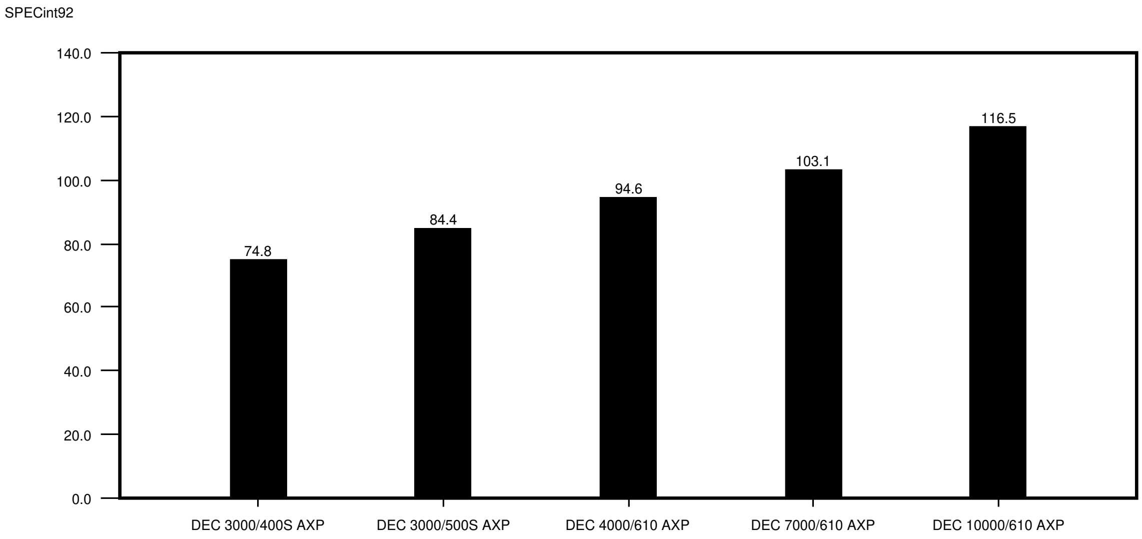


## Figure 2 SPEC CINT92 Results for Competitive Systems

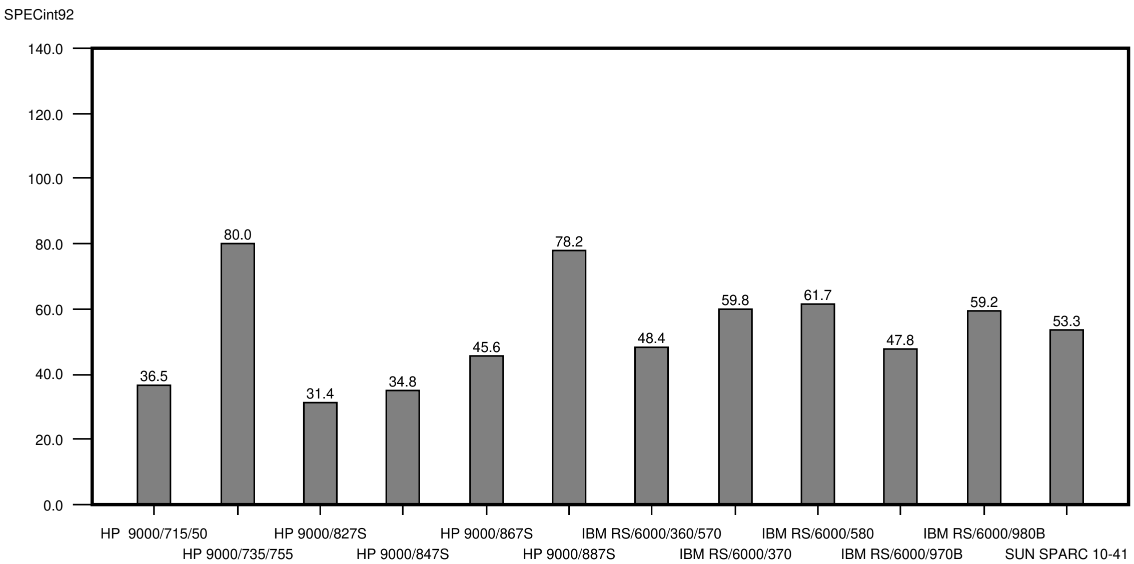

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

### SPEC CFP92

CFP92 consists of fourteen real-world applications; two are written in C and twelve in FORTRAN. Five of the fourteen programs are single precision, and the rest are double precision. SPECfp92™ equals the geometric mean of this suite's fourteen SPECratios. This suite contains the following application classes:

- 013.spice2g6–Circuit design
- 015.doduc–Monte Carlo simulation
- 034.mdljdp2–Quantum chemistry
- 039.wave5–Maxwell equations
- 047.tomcatv–Coordinate translation
- 048.ora–Optics ray tracing
- 052.alvinn–Robotics; neural nets
- 056.ear–Human ear modeling
- 077.mdljsp2–Single precision version of 034.mdljdp2
- 078.swm256–Shallow water model
- 089.su2cor–Quantum physics
- 090.hydro2d–Astro physics
- 093.nasa7–NASA math kernels
- 094.fppp–Quantum chemistry

## Figure 3 SPEC CFP92 Results for Alpha AXP Servers

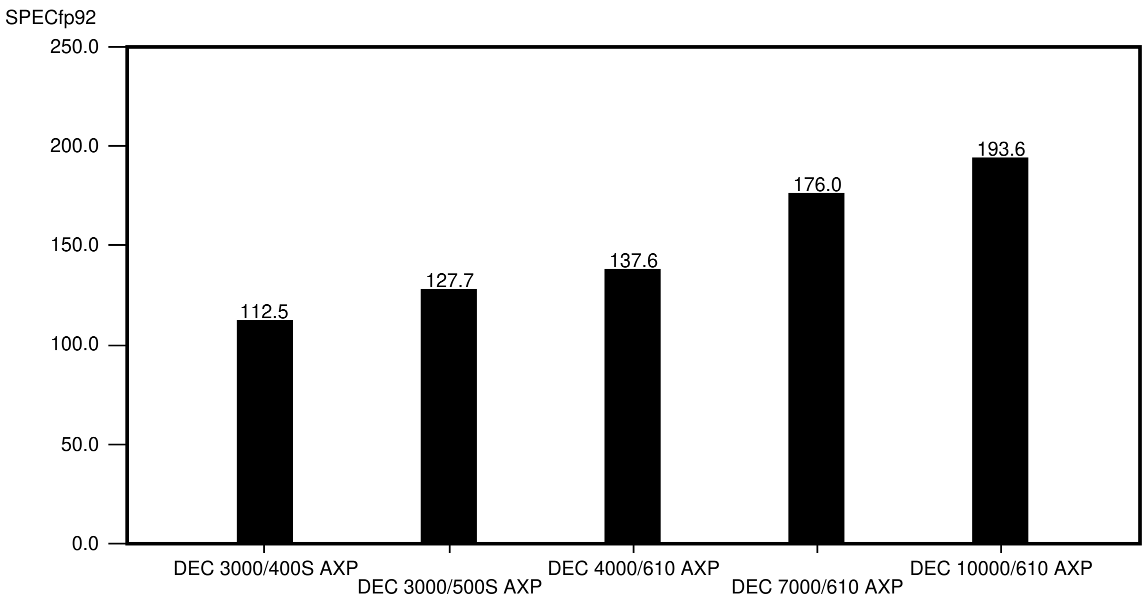


## Figure 4 SPEC CFP92 Results for Competitive Systems

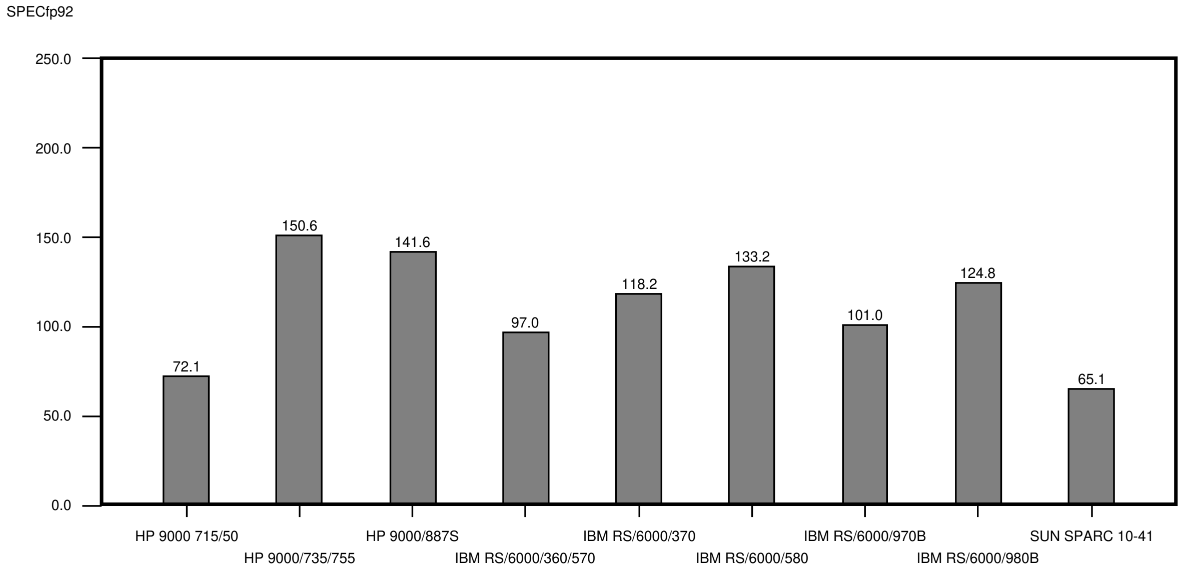

### SPEC Homogeneous Capacity Method based on SPEC CINT92 and CFP92

SPEC Homogeneous Capacity Method benchmarks test multiprocessor efficiency. According to SPEC, "The SPEC Homogeneous Capacity Method provides a fair measure for the processing capacity of a system — how much work can it perform in a given amount of time. The "SPECrate" is the resulting new metric, the rate at which a system can complete the defined tasks....The SPECrate is a capacity measure. It is not a measure of how fast a system can perform any task; rather it is a measure of how many of those tasks that system completes within an arbitrary time interval (*SPEC Newsletter,* June 1992)." The SPECrate is intended to be a valid and fair comparative metric to use across systems of any number of processors.

The following formula is used compute the SPECrate:

SPECrate = #CopiesRun * ReferenceFactor * UnitTime / ElapsedExecutionTime

SPECrate_int92™ equals the geometric mean of the SPECrates for the six benchmarks in CINT92. SPECrate_fp92™ is the geometric mean of the SPECrates of the fourteen benchmarks in CFP92.

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

**Figure 5 SPECrate_int92 Results for Alpha AXP Servers**

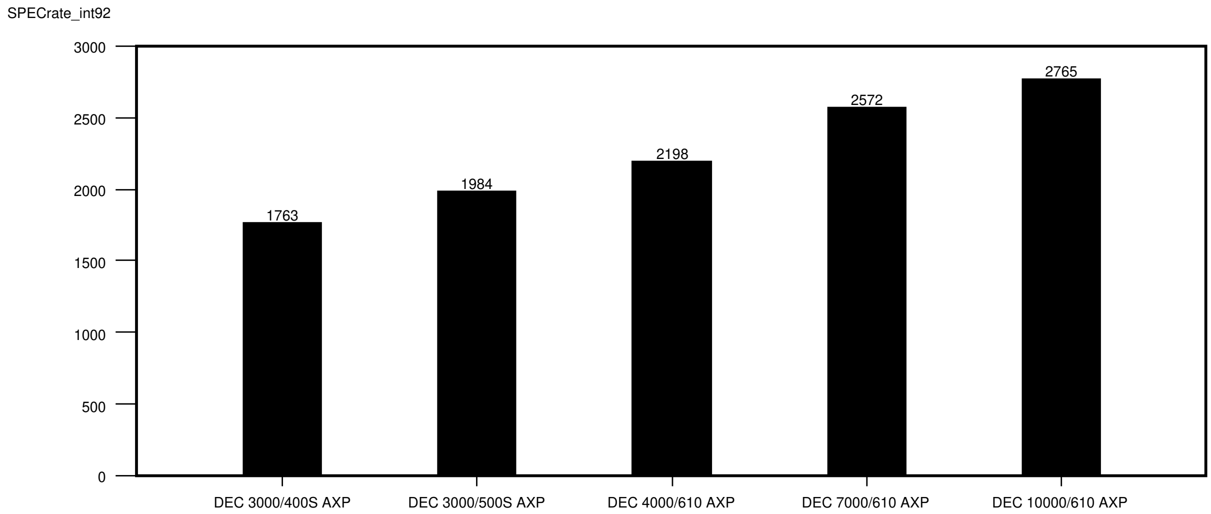


**Figure 6 SPECrate_int92 Results for Competitive Systems**

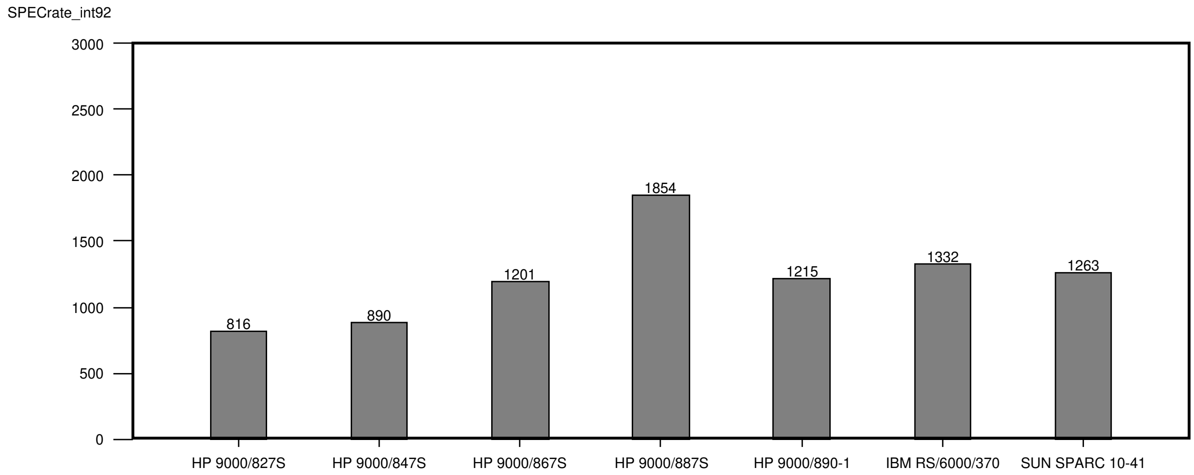

## Figure 7 SPECrate_fp92 Results for Alpha AXP Servers

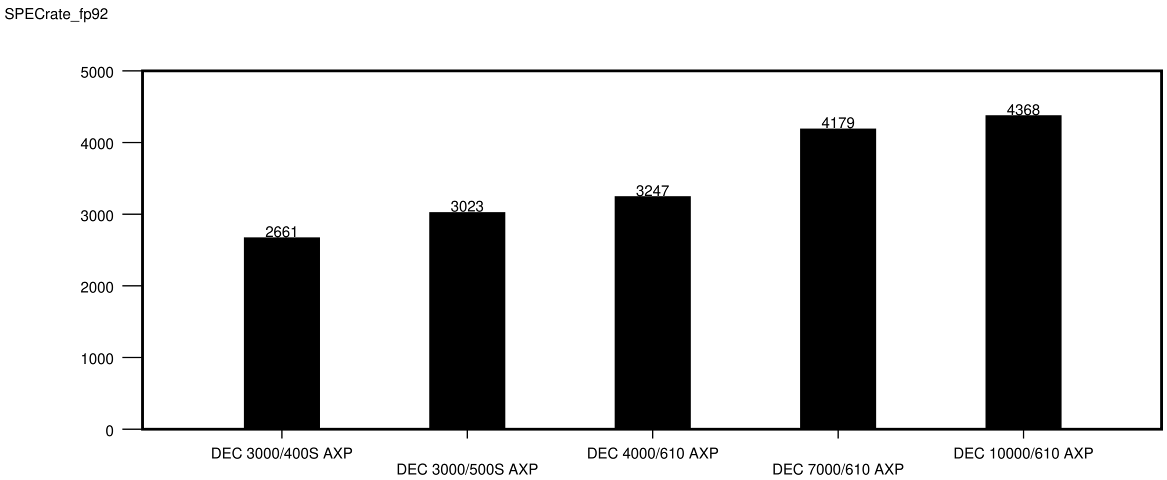


## Figure 8 SPECrate_fp92 Results for Competitive Systems

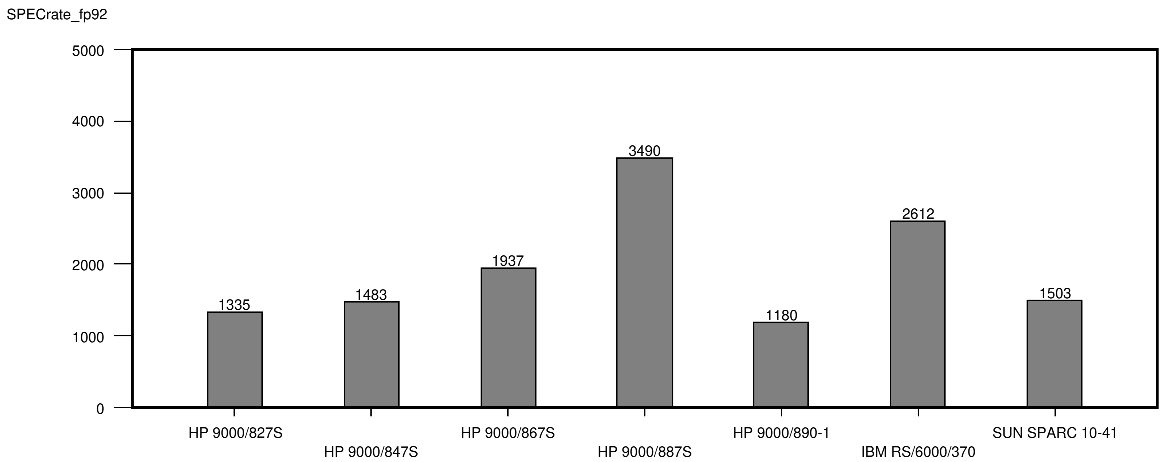

## SPEC Release 1

In October 1989, SPEC introduced SPEC Release 1, a benchmark suite that measures CPU-intensive, single stream performance of uniprocessor systems. SPEC Release 1 consists of ten portable programs similar to those found in technical environments. Four programs are written in C and primarily test integer performance. The remaining six programs are written in FORTRAN and measure floating-point performance. SPECmark89™ is the metric for this suite.

SPECmark89 represents the geometric mean of the ten benchmark SPECratios™. The SPECratio for a benchmark is the quotient derived from dividing the SPEC Reference Time by a particular machine's corresponding run time. The SPEC Reference Time is the time that it takes a DEC VAX 11/780 to run each benchmark (in seconds).

### Figure 9 SPECmark89 Results for Alpha AXP Servers

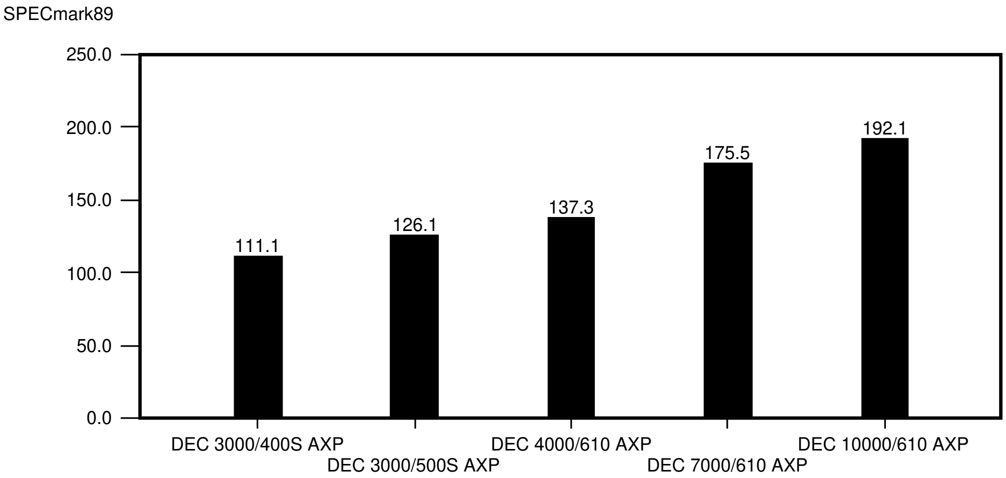


### Figure 10 SPECmark89 Results for Competitive Systems

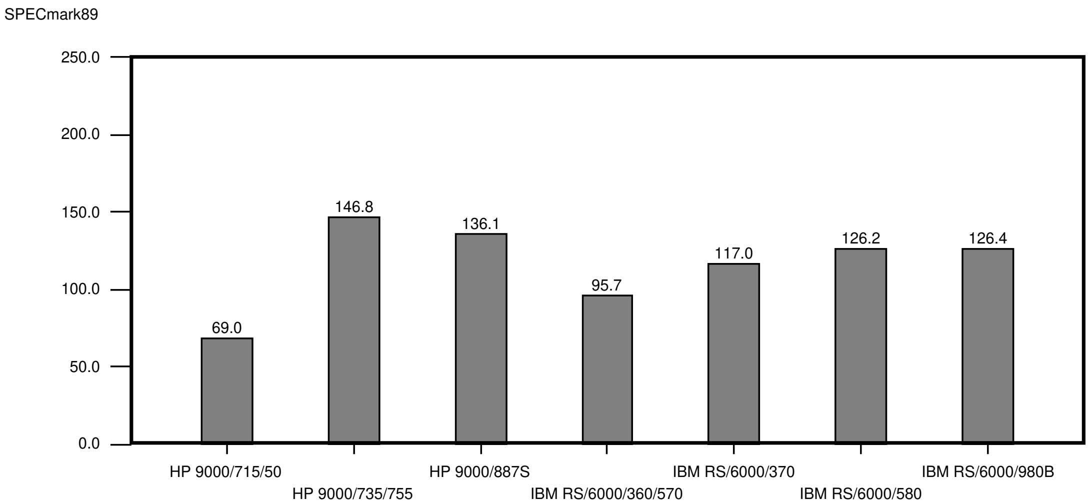

# AIM Benchmark Suites

AIM Technology developed a family of benchmarks for UNIX systems. Each benchmark evaluates systems differently and provides independent measurements. The AIM benchmarks we ran were:

- AIM Suite II Subsystem Benchmark
- AIM Suite III Multiuser Benchmark

AIM verifies the results of the Suite II and Suite III run by licensed vendors and uses the source data for their AIM Performance Report service. We discuss these benchmarks in the following sections.

### AIM Suite II Subsystem Benchmark

AIM Suite II Subsystem Benchmark provides single-user subsystem performance measurements. Each component of a UNIX computer system is exercised and timed. Forty-one subtests provide absolute processing rates in operations per second for each subsystem (i.e., I/O transfers, function calls, and UNIX system calls).

The seven basic hardware and software functional groups evaluated and reported are as follows:

- Math– mathematics
- Floating point– single- and double-precision adds, multiplies, and divides
- RAM– memory reads, writes, and copies
- Array references
- Function calls– with 0, 1, 2, and 31 arguments
- System calls– to standard UNIX system functions
- Disk– reads, writes, seeks with buffer cache enabled, and forced synchronization of writes
- Pipe– data transfer over a UNIX interprocess pipe

This suite contains application mixes, combinations of functional and subsystem tests that simulate database, scientific, software development, and other applications.

### AIM Suite III Multiuser Benchmark

AIM Suite III is designed to measure, evaluate, and predict UNIX multiuser system performance of multiple systems. AIM Suite III uses 33 functional tests, and these tests can be grouped to reflect the computing activities of various types of applications. It is designed to stress schedulers and I/O subsystems and includes code that will exercise TTYs, tape subsystems, printers, and virtual memory management. The benchmark will run until it reaches either the user-specified maximum number of simulated users or system capacity.

The 33 subsystem tests, each of which exercises one or more basic functions of the UNIX system under test, are divided into six categories based on the type of operation involved. The categories are as follows:

- RAM
- Floating Point
- Pipe
- Logic
- Disk
- Math

Within each of these six categories, the relative frequencies of the subsystem tests are evenly divided (with the exception of small biases for add-short, add-float, disk reads, and disk writes).

AIM Suite III contains no application level software. Each simulated user runs a combination of subsystem tests. The load that all simulated users put on the system is said to be characteristic of a UNIX time-sharing environment. The mix of subsystem tests can be varied to simulate environments with differing resource requirements.

AIM provides a default model as a representative workload for UNIX multiuser systems and the competitive data that AIM Technology publishes is derived from this mix of subsystem tests.

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

The AIM Performance Rating identifies the maximum performance of the system under optimum usage of CPU, floating point, and disk caching. At a system's peak performance, an increase in the workload will cause a deterioration in performance. AIM Maximum User Load Rating identifies system capacity under heavy multitasking loads, where disk performance also becomes a significant factor. Throughput is the total amount of work the system processes, measure in jobs/minute. Maximum throughput is the point at which the system is able to process the most jobs per minimum.

Digital's Alpha AXP server family's AIM benchmark results are shown in the following table:

**Table 2 AIM Benchmark Suites Results for Alpha AXP Servers**

| System | Performance Rating | Maximum User Loads | Maximum Throughput (jobs/minute) |
|---|---|---|---|
| DEC 3000/500S AXP (192MB memory, 8 disks) | 82.9 | 649 | 812.9 |
| DEC 3000/400S AXP (128MB memory, 3 disks) | 70.3 | 485 | 688.7 |

Shown below are Digital Alpha AXP servers' competitors' AIM results.

**Table 3 AIM Benchmark Suites Results for Competitive Systems**

| System | Performance Rating | Maximum User Loads | Maximum Throughput (jobs/minute) |
|---|---|---|---|
| HP 9000/755 (64MB memory, 2 disks) | 71.7 | 580 | 703.1 |
| HP 9000/735 (32MB memory, 2 disks) | 71.7 | 422 | 703.1 |
| HP 9000/887S (256MB memory, 9 disks) | 69.2 | 620 | 677.8 |
| HP 9000/750 (64MB memory, 2 disks) | 46.7 | 388 | 457.6 |
| HP 9000/715/50 (32MB memory,2 disks) | 33.7 | 252 | 330.0 |
| IBM RS/6000 580 (128MB memory, 4 disks) | 62.1 | 518 | 609 |
| SUN SPARC 10 Model 41 (64MB memory, 3 disks) | 44.3 | 203 | 433.8 |

## Perfect Benchmarks Suite

The Perfect (Performance Evaluation for Cost-effective Transformations) Benchmarks™ Suite represents an ongoing effort among several universities, research centers, and supercomputing firms to produce a benchmark package oriented to supercomputers and parallel processing. Currently, 13 FORTRAN programs totaling over 50,000 lines of code make up the Perfect Benchmark Suite. These programs include scientific and engineering applications representing four types of real applications areas: fluid flow, chemical and physical, engineering design, and signal processing.

Perfect Benchmark Suite's results are measured in millions of floating-point operations per second (MFLOPS). Figure 11 shows the geometric mean of the suite's results in MFLOPS.

**Figure 11 PERFECT Results for Alpha AXP Servers**

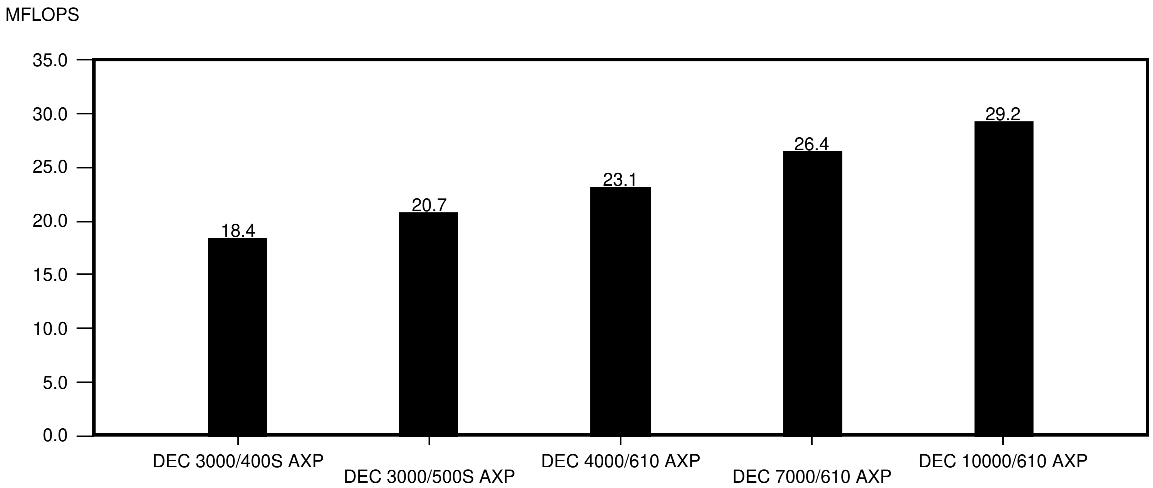

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Dhrystone Benchmarks

Developed as an Ada program in 1984 by Dr. Reinhold Weicker, the Dhrystone benchmark was rewritten in C in 1986 by Rick Richardson. Dhrystone measures processor and compiler efficiency and is representative of systems programming environments. Dhrystones are most commonly expressed in Dhrystone instructions per second.

Dhrystone V1 and V2 vary considerably. Version 1.1 contains sequences of code segments that calculate results never used later in the program. These code segments are known as "dead code." Compilers able to identify the dead code can eliminate these instruction sequences from the program. These compilers allow a system to complete the program in less time and result in a higher Dhrystones rating. Dhrystones V2 was modified to execute all instructions.

## Figure 12 Dhrystone V1.1 and V2.1 Results for Alpha AXP Servers

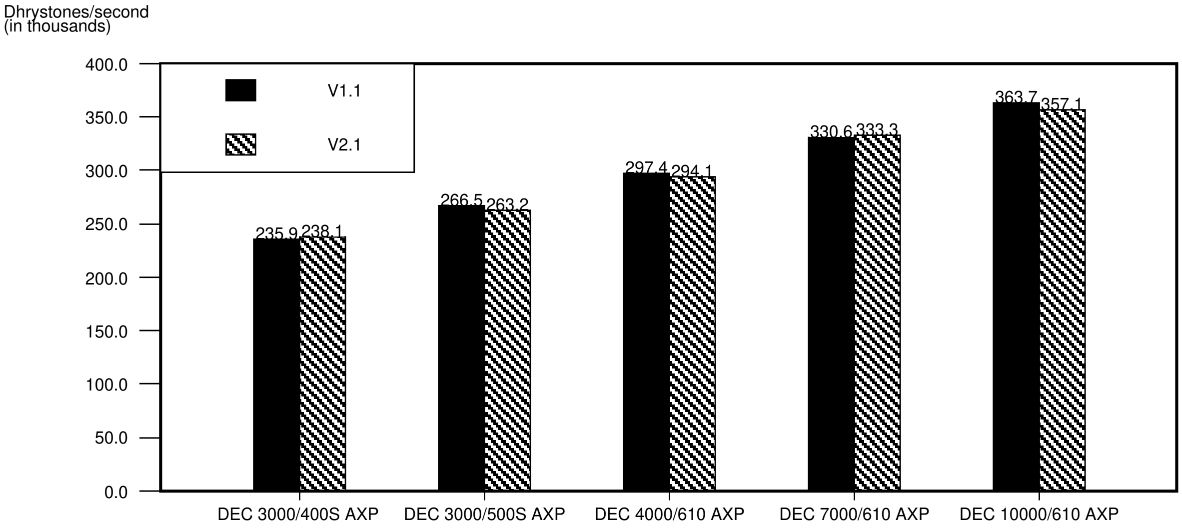


## Figure 13 Dhrystone V1.1 Results for Competitive Systems

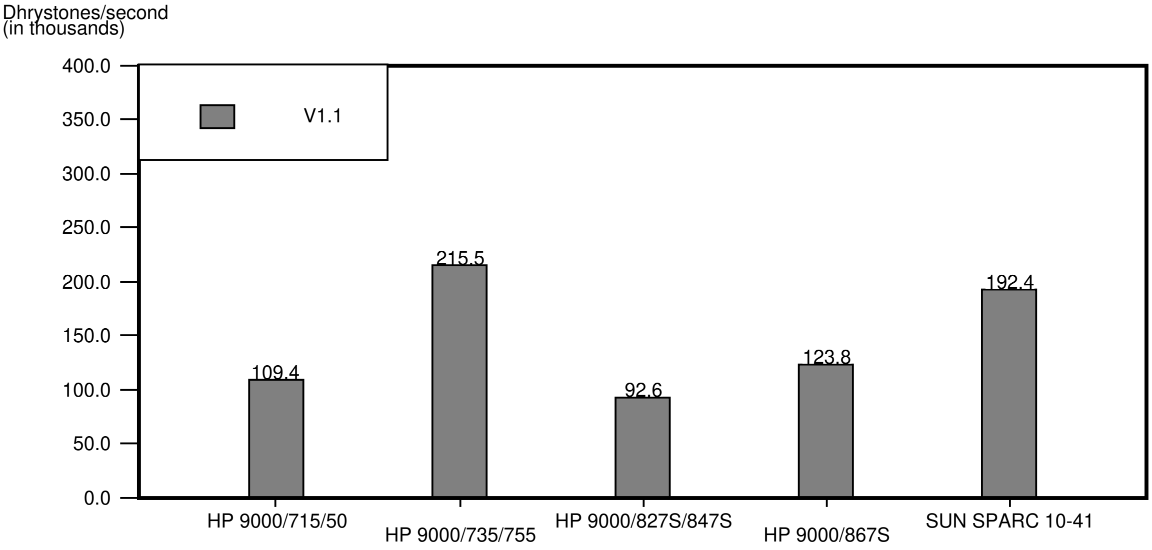


***Note:*** *Hewlett-Packard has reported Dhrystone V2.0 results, but not V2.1. SUN has not reported Dhrystone V2.0 or V2.1.*

# LINPACK 100x100 and 1000x1000 Benchmarks

LINPACK is a linear equation solver written in FORTRAN. LINPACK programs consist of floating-point additions and multiplications of matrices. The LINPACK benchmark suite consists of two benchmarks.

1. 100x100 LINPACK solves a 100x100 matrix of simultaneous linear equations. Source code changes are not allowed so that the results may be used to evaluate the compiler's ability to optimize for the target system.

2. 1000x1000 LINPACK solves a 1000x1000 matrix of simultaneous linear equations. Vendor optimized algorithms are allowed.

The LINPACK benchmarks measure the execution rate in MFLOPS (millions of floating-point operations per second). When running, the benchmark depends on memory-bandwidth and gives little weight to I/O. Therefore, when LINPACK data fit into system cache, performance may be higher.

## Figure 14 LINPACK Double-Precision Results for Alpha AXP Servers

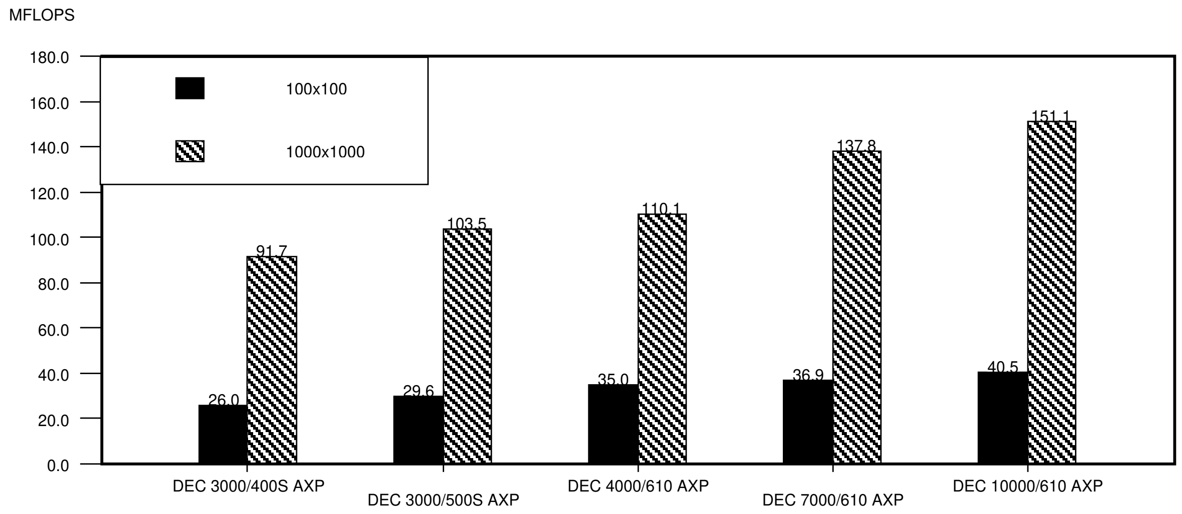


## Figure 15 LINPACK Double-Precision Results for Competitive Systems

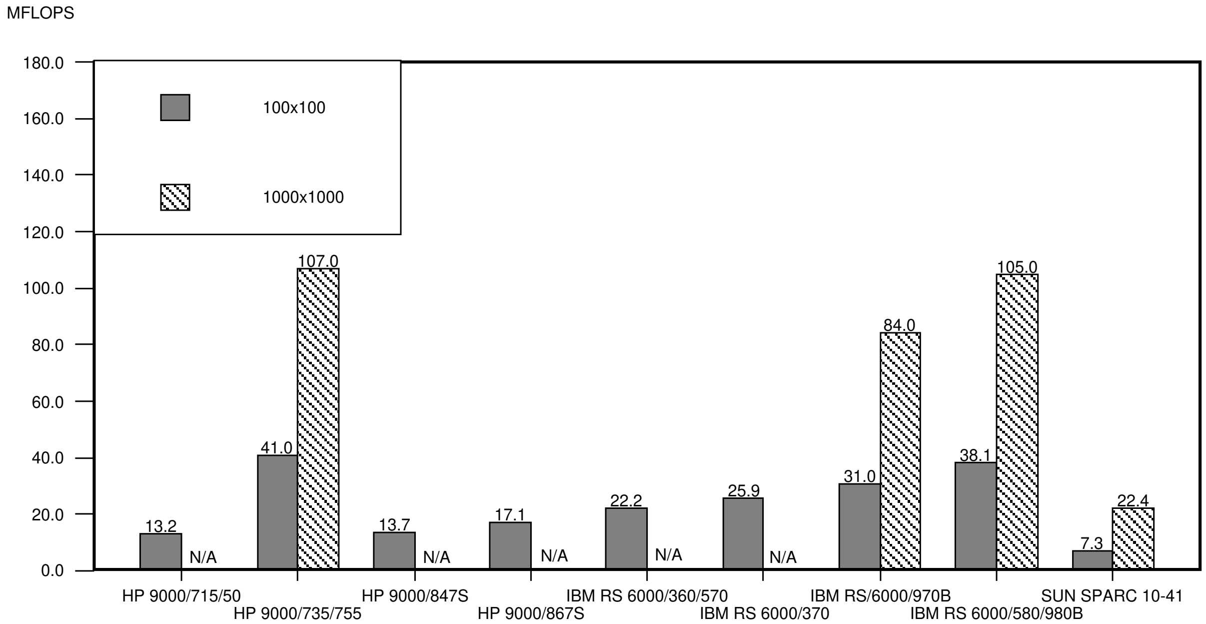

## CERN Benchmark Suite

In the late 1970's, the User Support Group at CERN (the European Laboratory for Particle Physics) collected from different experimental groups a set of typical programs for event simulation and reconstruction and created the CERN Benchmark Suite. In 1985, Eric McIntosh, system analyst, redefined the tests in order to make them more portable and more representative of the then current workload and FORTRAN 77.

Presently, the CERN Benchmark Suite contains four production tests: two event generators (CRN3 and CRN4) and two event processors (CRN5 and CRN12). These applications are basically scalar and are not significantly vectorizable nor numerically intensive. Additionally, several "kernel" type applications were added to supplement the production tests to get a feel for compilation times (CRN4C), vectorization (CRN7 and CRN11), and character manipulation (CRN6).

The CERN Benchmark Suite metric is CPU time. Results are normalized to a DEC VAX 8600, and the geometric mean of the four production tests' ratios yields the number of CERN units. CERN units increase with increasing performance.

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Figure 16   CERN Results for Alpha AXP Servers

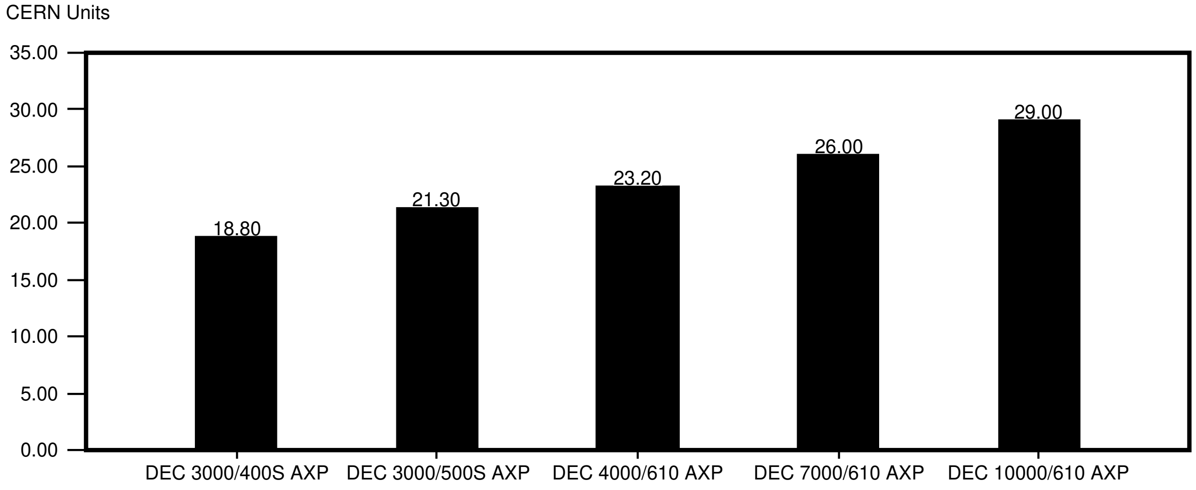


## Figure 17   CERN Results for Competitive Systems

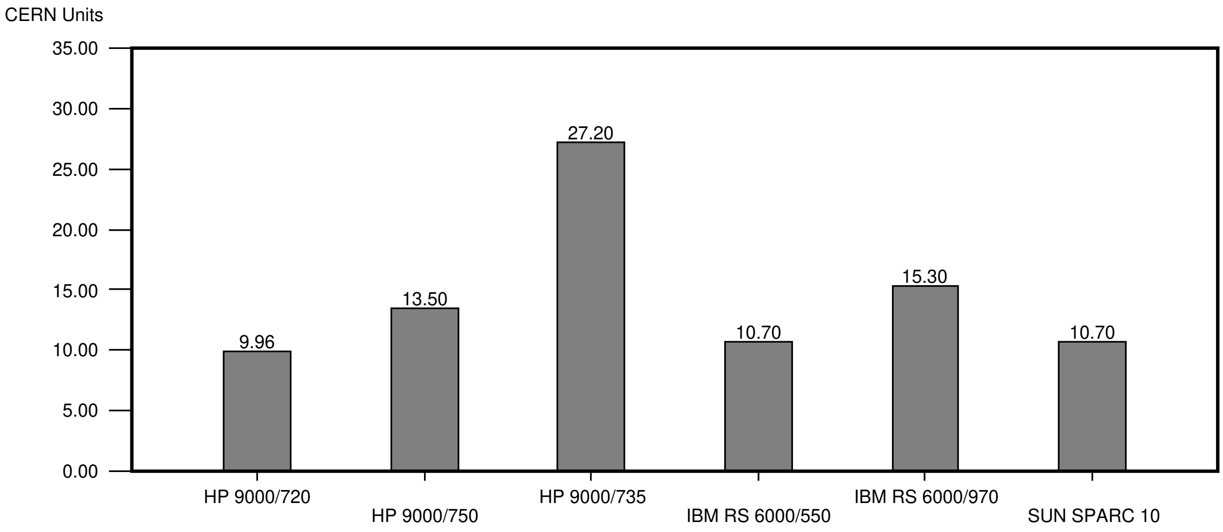


Source: "CERN Report" (11/23/92).

## DN&R Labs CPU2 Benchmark

DN&R Labs CPU2, a benchmark from *Digital News & Review* magazine, is a floating-point intensive series of FORTRAN programs and consists of thirty-four separate tests. The benchmark is most relevant in predicting the performance of engineering and scientific applications. Performance is expressed as a multiple of MicroVAX II Units of Performance (MVUPs).

**Figure 18 DN&R Labs CPU2 Results for Alpha AXP Servers**

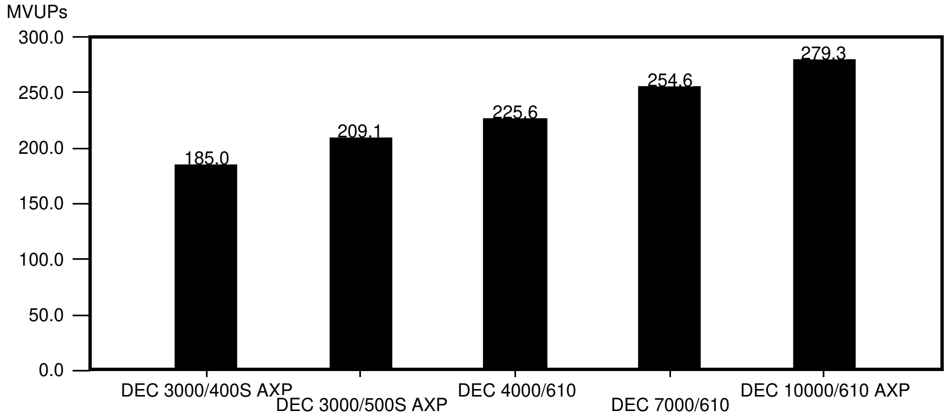


**Figure 19 DN&R Labs CPU2 Results for Competitive Systems**

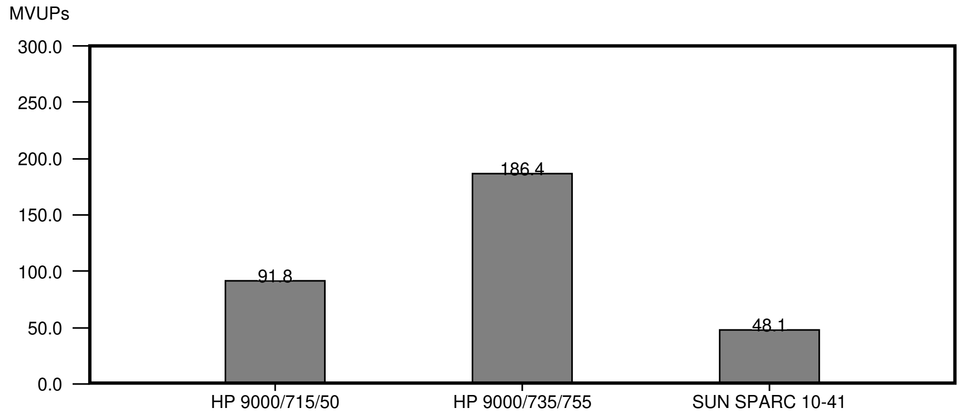

# Livermore Loops

This benchmark, also known as Livermore FORTRAN Kernels, was developed by the Lawrence National Laboratory in Livermore, CA. The laboratory developed this benchmark to evaluate large supercomputer systems. Computational routines were extracted, 24 sections of code in all, from programs used at the laboratory in the early 1980's to test scalar and vector floating performance.

The routines (kernels) are written in FORTRAN and draw from a wide variety of scientific applications including I/O, graphics, and memory management tasks. These routines also inhabit a large benchmark driver that runs the routines several times, using different input data each time. The driver checks on the accuracy and timing of the results.

The results of the 24 routines, one for each kernel, are reported in millions of floating-point operations per second (MFLOPS). Shown in this report are the calculated geometric means.

**Figure 20 Livermore Loops Results for Alpha AXP Servers**

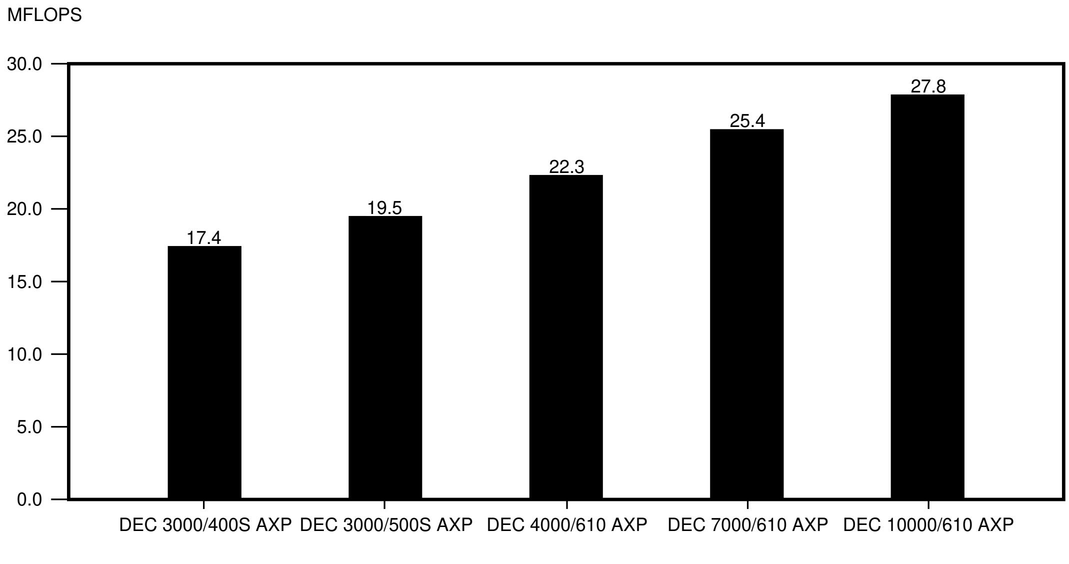

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Fast Fourier Transform and Convolution

### Fast Fourier Transform

Fast Fourier Transform (FFT) is the basis for a mathematical technique called Fourier Analysis that breaks down a function in space or time into sinusoidal components that have varying frequencies, amplitudes, and phases. A discrete Fourier Transform decomposes a collection of data into component sine and cosine representations.

Fourier analysis is widely used in all scientific and engineering work including biology, seismic engineering, electromagnetic analysis, and image processing. The DXML (Digital eXtended Math Library) Fourier transform subroutines take advantage of certain computational techniques to reduce the time required to compute the discrete Fourier Transform.

### Convolution

Convolution is a mathematical operation that is used to complement the signal processing abilities of the Fourier Transform.

This operation allows for modification of signal sequence by having a function or additional sequence of numbers selectively weight the chosen sequence. This weighting may be used to obtain properties from a signal source, to selectively enhance the signal source or to transform a signal source to another domain. The basis of most signal processing resides in the application of some type of convolution operation.

### Table 4 FFT Benchmark Results for Alpha AXP Servers

| Metric DEC OSF/1 AXP (MFLOPS) | DEC 3000 Model 500 AXP | DEC 4000 Model 610 AXP | DEC 7000 Model 610 AXP | DEC 10000 Model 610 AXP |
|---|---|---|---|---|
| FFT-1D | | | | |
| SP- Complex FFT | 91.0 | 106.6 | 115.8 | 127.2 |
| Real FFT | 78.4 | 89.8 | 94.8 | 104.4 |
| DP- Complex FFT | 57.5 | 69.8 | 77.1 | 80.6 |
| Real FFT | 56.5 | 72.5 | 82.6 | 84.6 |

# SLALOM Benchmark

Developed at Ames Laboratory, U.S. Department of Energy, the SLALOM (Scalable Language-independent Ames Laboratory One-minute Measurement) benchmark solves a complete, real problem (optical radiosity on the interior of a box). SLALOM is based on fixed time rather than fixed problem comparison. It measures input, problem setup, solution, and output, not just the time to calculate the solution.

SLALOM is very scalable and can be used to compare computers as slow as 104 floating-point operations per second to computers running a trillion times faster. You can use the scalability to compare single processors to massively parallel collections of processors, and you can study the space of problem size versus ensemble size in fine detail.

The SLALOM benchmark is CPU-intensive and measures, in units called patches, the size of a complex problem solved by the computer in one minute.

**Table 5 SLALOM Results**

| System | Patches |
|---|---|
| DEC 10000 Model 610 AXP Server | 7,248 |
| DEC 7000 Model 710 AXP Server | 6,902 |
| DEC 4000 Model 610 AXP Server | 6,496 |
| DEC 3000 Model 500S AXP Server | 6,084 |
| DEC 3000 Model 400S AXP Server | 5,776 |

## Basic Real-Time Primitives

Measuring basic real-time primitives such as process dispatch latency and interrupt response latency enhances our understanding of the responsiveness of the DEC OSF/1 AXP Real-Time kernel.

**Process Dispatch Latency** is the time it takes the system to recognize an external event and switch control of the system from a running, lower-priority process to a higher-priority process that is blocked waiting for notification of the external event.

**Interrupt Response Latency** (ISR latency) is defined as the amount of elapsed time from when the kernel receives an interrupt until execution of the first instruction of the interrupt service routine.

The DEC 3000 Model 500 AXP's Basic Real-Time Primitives results appear in the following table.

### Table 6 Basic Real-Time Primitives Results

| Metric | Minimum (µsec) | Maximum (µsec) | Mean (µsec) |
|---|---|---|---|
| Process Dispatch Latency | 62.5 | 187.7 | 68.9 |
| Interrupt Response Latency | 7.0 | 50.4 | 8.4 |

Test configuration: DEC 3000 Model 500 AXP, 256 MB memory, DEC OSF/1 AXP RT operating system, Ver 1.2.
Test conditions: Single user mode, no network.

Shown next are histograms of the process dispatch latency times and the interrupt response latency times of a DEC 3000 Model 500 AXP system running in the DEC OSF/1 AXP Real-Time operating system environment.

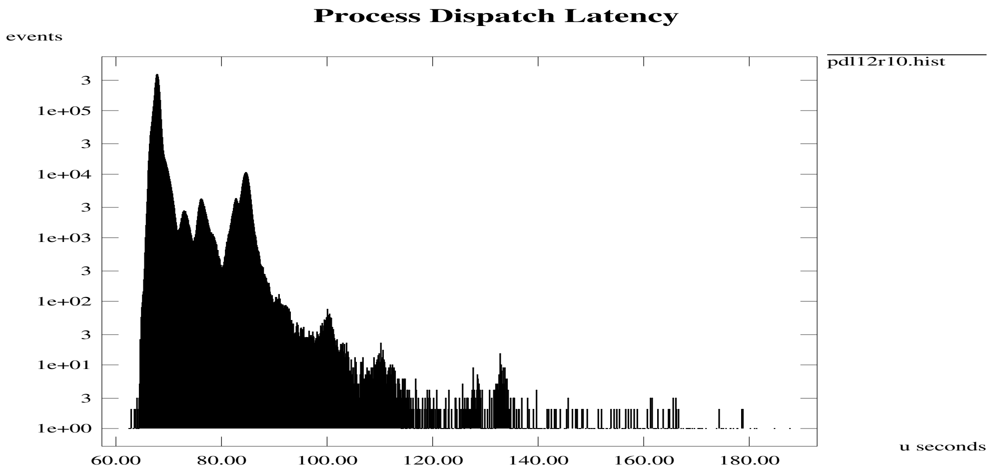


**Figure 21 Basic Real-Time Primitives Process Dispatch Latency Results**

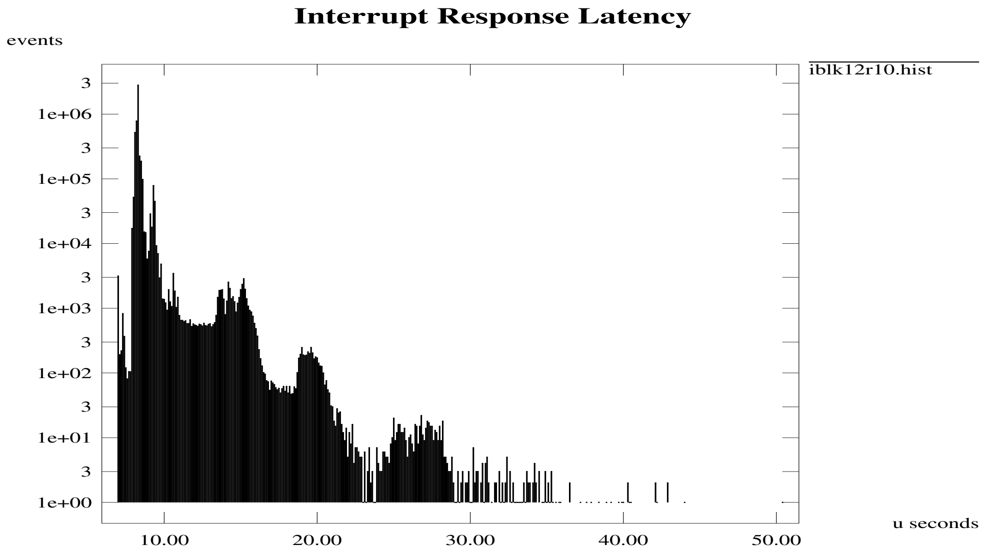


**Figure 22 Basic Real-Time Primitives Interrupt Response Latency Results**

## Rhealstone Benchmark

The Rhealstone Real-time Benchmark is a definition for a synthetic test that measures the real-time performance of multitasking systems. It is unique in that "the verbal and graphical specifications, not the C programs, are the essential core of the benchmark" (Kar 1990). We implemented the following components of the Rhealstone Benchmark, which measure the critical features of a real-time system:

1. Task switch time—the average time to switch between two active tasks of equal priority.
2. Preemption time—the average time for a high-priority task to preempt a running low-priority task.
3. Interrupt latency time—the average delay between the CPU's receipt of an interrupt request and the execution of the first instruction in an interrupt service routine.
4. Semaphore shuffle time—the delay within the kernel between a task's request and its receipt of a semaphore, which is held by another task (excluding the runtime of the holding task before it relinquishes the semaphore).

   ***Note:*** *We report the delays associated with* `sem_wait` *as well as* `sem_post` *plus the two corresponding context switches.*
5. Intertask message latency —the delay within the kernel when a non-zero-length data message is sent from one task to another.

Deadlock-break time is not applicable to DEC OSF/1 AXP Real-Time.

The following figures show the four components of the Rhealstone benchmark we implemented.

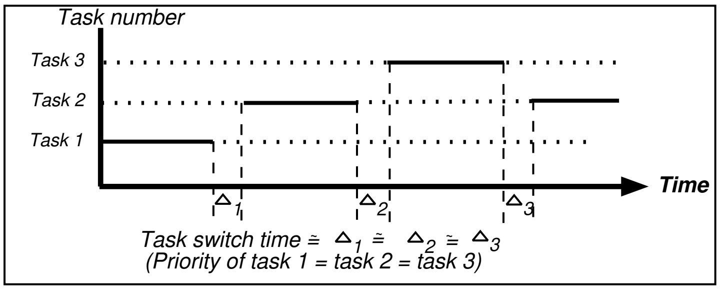


**Figure 23 Rhealstone Component–Task-switch Time**

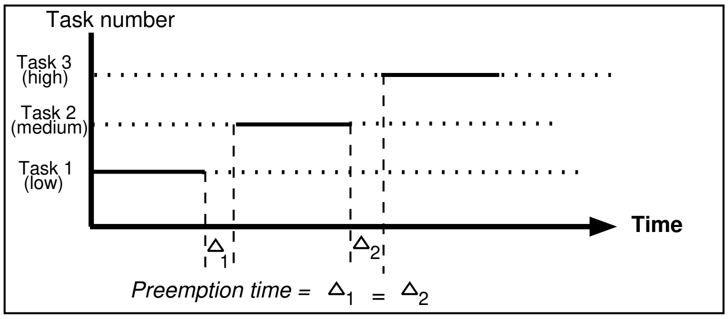


**Figure 24 Rhealstone Component–Preemption Time**

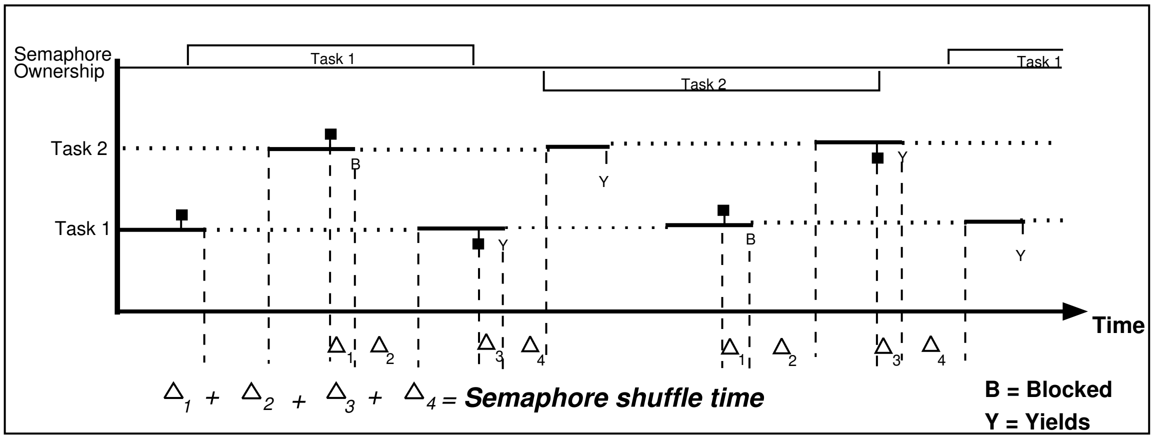


**Figure 25 Rhealstone Component–Semaphore-shuffle Time**

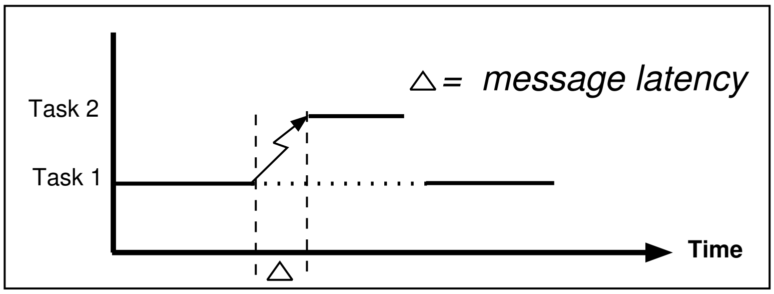


**Figure 26 Rhealstone Component–Intertask Message Latency**

**Source:** Figures 23, 24, and 26 largely taken from Kar, Rabindra, P., "Implementing the Rhealstone Real-Time Benchmark" *(Dr. Dobb's Journal*, April 1990).

Rhealstone results measured on a DEC 3000 Model 500 AXP are shown in the following table.

**Table 7 Rhealstone Benchmark Results**

| Metric | Mean (µsec) |
|---|---|
| Task Switch Time | 18.1 |
| Preemption Time | 33.1 |
| Interrupt Latency Time | 8.4 (Range: 7.0–50.4) |
| Intertask Message Time | 92.8 |
| Semaphore Shuffle Time | 152.2 |

Test configuration: DEC 3000 Model 500 AXP, 256 MB memory, DEC OSF/1 AXP RT operating system, Ver 1.2.
Test conditions: Single user mode, no network.

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## Information about Performance

The performance of the Alpha AXP server family was evaluated using industry-standard benchmarks. These benchmarks allow comparisons across vendors' systems. Performance characterization is one "data point" to be used in conjunction with other purchase criteria such as features, service, and price.

We chose the competitive systems (shown in the preceding charts and tables) based on comparable or close CPU performance, coupled with comparable expandability capacity (primarily memory and disk). Although we do not present price comparisons in this report, system price was a secondary factor in our competitive choices.

***Note:*** *The performance information in this report is for guidance only. System performance is highly dependent upon application characteristics. Individual work environments must be carefully evaluated and understood before making estimates of expected performance.*

This report simply presents the data, based on specified benchmarks. Competitive information is based on the most current published data for those particular systems and has not been independently verified (except as noted). The Alpha AXP performance information presented in this brief is the latest measured results as of the date published. Digital has an ongoing program of performance engineering across all products. As system tuning and software optimizations continue, Digital expects the performance of its servers to increase. As more benchmark results become available, Digital will publish reports containing the new and updated benchmark data.

For more information on Digital's Alpha AXP server family, please contact your local Digital sales representative.

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

## References

| System and Vendor | Sources |
|---|---|
| DEC 3000 Model 400S and 500S AXP Servers | All benchmarking performed by Digital Equipment Corporation. |
| DEC 4000 Models 610 AXP Servers | All benchmarking performed by Digital Equipment Corporation. |
| DEC 7000 Models 610 AXP Servers | All benchmarking performed by Digital Equipment Corporation. |
| DEC 10000 Models 610 AXP Server | All benchmarking performed by Digital Equipment Corporation. |
| HP 9000 Models 715/50, 735, and 755 | SPEC, LINPACK, and Dhrystone benchmark results reported by Hewlett-Packard in "HP Apollo 9000 Series 700 Workstation Systems" (11/92). 735 LINPACK results from Dongarra, J., "Performance of Various Computers Using Standard Linear Equations Software" (3/6/93).<br>AIM results reported in *UNIX System Price Performance Guide,* AIM Technology (Winter 1992-1993).<br>DN&R Labs CPU2 results reported in Workstation Laboratories, Volume 19, Chapters 20 & 21 (1/93). |
| HP 9000 Model 730 | SPEC benchmark results reported in *SPEC Newsletter* (12/92).<br>Dhrystone and LINPACK 100x100 benchmark results reported in "HP Apollo 9000 Series 700 Performance Brief" (1/92).<br>LINPACK 1000x1000 benchmark results reported in Dongarra, J., "Performance of Various Computers Using Standard Linear Equations Software" (9/28/92).<br>DN&R Labs CPU2 results reported in Workstation Laboratories, Inc., Volume 16, Chapter 12 (3/1/92). |
| HP 9000 Model 750 | SPEC benchmark results reported by HP (11/10/92).<br>Dhrystone and LINPACK benchmark results reported in "HP Apollo 9000 Series 700 Performance Brief" (1/92).<br>AIM results reported in *UNIX System Price Performance Guide,* AIM Technology (Winter 1992-1993). |
| HP 9000 Models 827S, 847S, and 867S | SPEC benchmark results reported in *SPEC Newsletter* (3/92).<br>Dhrystone and LINPACK benchmark results reported in"HP 9000 Series 800 Business Server System Performance" (1992).<br>AIM results reported in *UNIX System Price Performance Guide,* AIM Technology (Winter 1992-1993). |
| HP 9000 Model 887S | SPEC benchmark results reported in *SPEC Newsletter* (9/92). |
| IBM RS 6000 Models 360 and 370 | SPEC and LINPACK 100x100 results reported by IBM (2/2/93). |
| IBM RS 6000 Models 580, 970B, and 980B | SPEC benchmark results reported in *SPEC Newsletter* (9/92) and by IBM (9/92).<br>AIM results reported in *UNIX System Price Performance Guide,* AIM Technology (Winter 1992-1993).<br>LINPACK 100x100 benchmark results reported by IBM (9/92) . LINPACK 1000x1000 results reported in Dongarra, J., "Performance of Various Computers Using Standard Linear Equations Software" (3/6/93). |

# Alpha AXP Server Family
# Performance Brief - DEC OSF/1 AXP

| | |
|---|---|
| SUN SPARC 10 Model 41 | SPECmark, LINPACK 100x100 dp, and DN&R Labs CPU2 results reported in Workstation Laboratories, Volume 19, Chapter 19 (1/93).<br>AIM results reported in *UNIX System Price Performance Guide*, AIM Technology (Fall 1992).<br>SPEC92 benchmark results reported in *SPEC Newsletter* (9/92).<br>LINPACK and Dhrystone benchmark results reported by SUN (11/10/92). |

### Article

Kar, Rabindra P. "Implementing the Rhealstone Real-Time Benchmark" (*Dr. Dobb's Journal,* April 199).