

Box 30
Folder 20

VAX 6000, 1985-1991

Digital Equipment Corporation records
Engineers' papers: Mike Uhler papers: VAX development recor
X2675.2004, Box 30

102749989

164

THE COMPUTER HISTORY MUSEUM



1 027 4998 9

DECus 1990

SEZ358.001

90Dec06

Technical Description of VAX 6000 Enhancements

Richard Gillett

VAX Systems Engineering

digital

SEZ358.002

4DEC90 T

0DEC325

System Design Goal

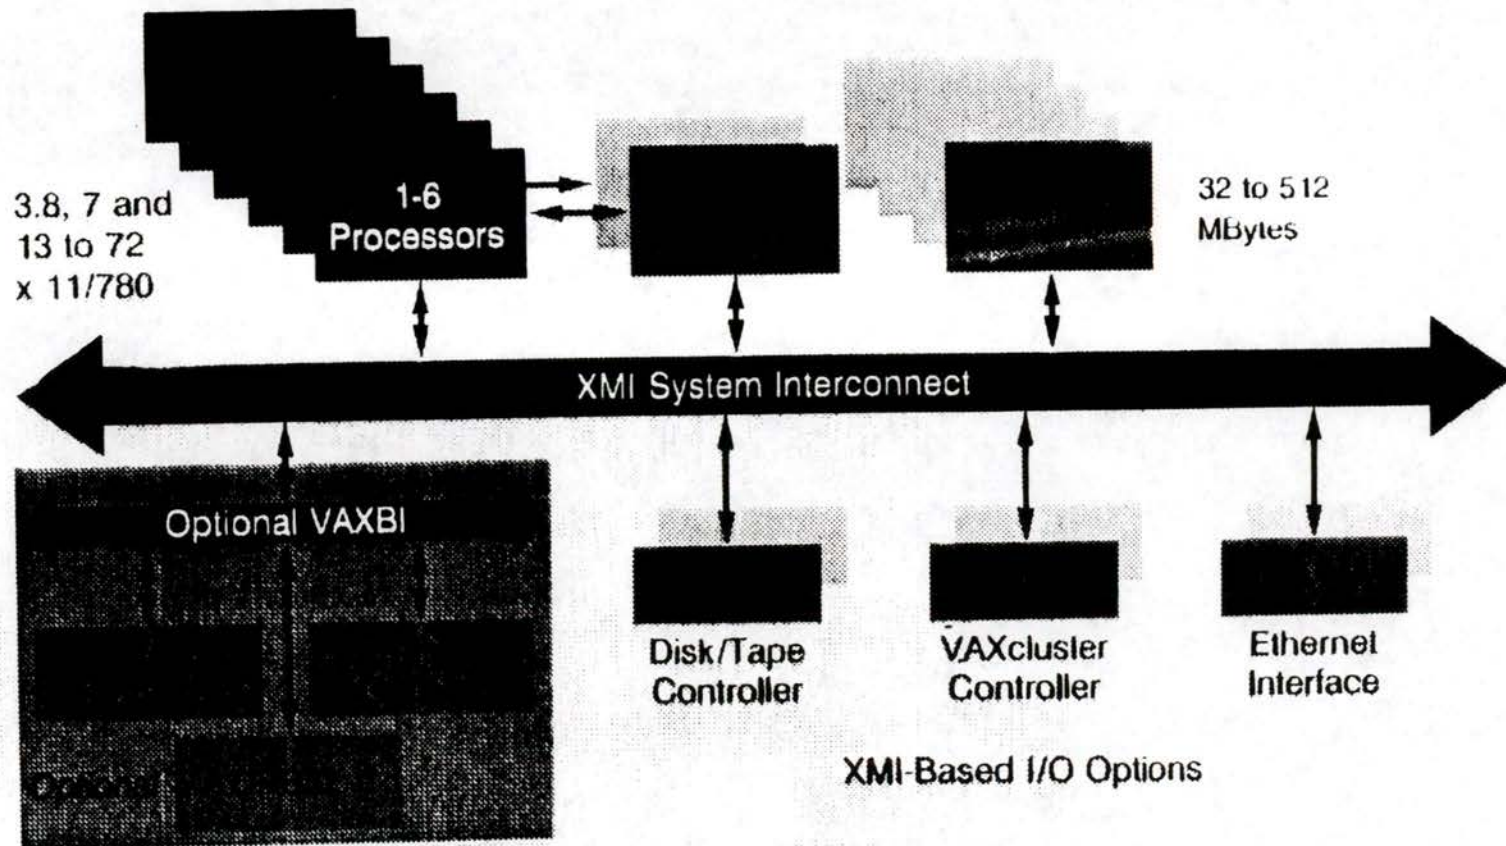
Extend the VAX 6000 family in all dimensions while maintaining traditional "just swap modules" upgradability for existing VAX 6000 customers

SEZ358.003

4DEC90 M

0DEC325

XMI Block Diagram



■ Specific Design Goals

- Increased processor performance
 - 75% to 100% more performance than Model 400
- Improved vector processor performance on real workloads

■ Specific Design Goals

- Increased memory capacity and performance
 - 2 x memory capacity - now 512MB
 - 4 x MB/XMI slot - now 128MB
 - 1.6 x memory controller bandwidth

■ Specific Design Goals

- Increased I/O capability
 - 3 x I/O controller performance
 - 4 x I/O connectivity
 - 4 x solid-state disk performance
- Increased configuration flexibility
 - More configuration options
 - New console storage options

■ The Challenges Of Increased Processor Performance

- With write-through caches, XMI bus supports about 40 VUPs
- With more complex cache implementations the XMI supports significantly more processor performance

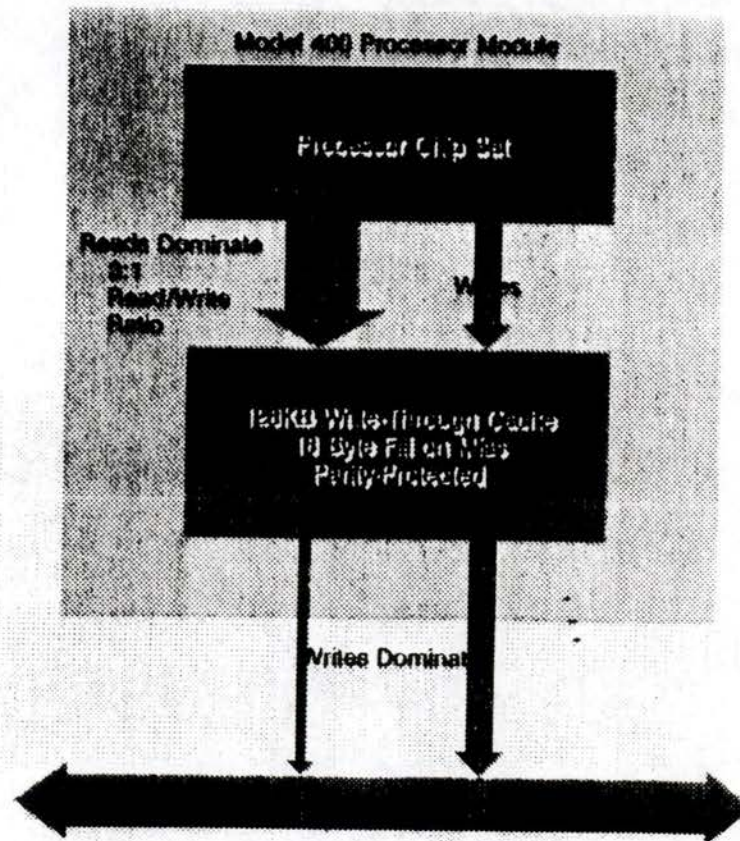
Model 500 Cache Enhancements

- Cache size increased by factor of four to 512Kbytes
 - Reduces cache miss rate
- Cache fill size was doubled to 32bytes
 - Reduces cache miss rate
 - Uses more efficient transfer size on XMI

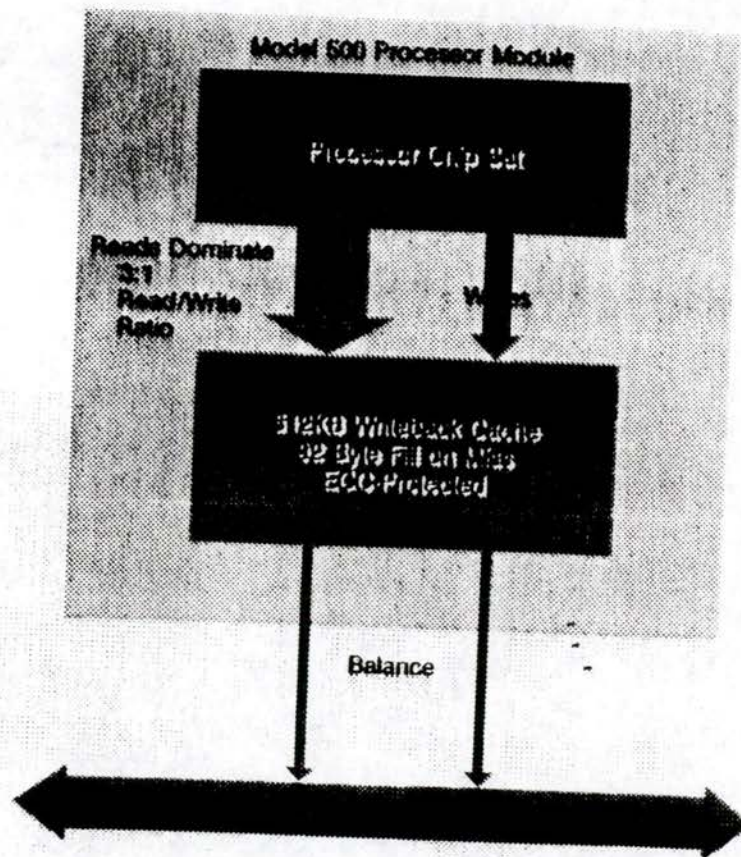
Model 500 Cache Enhancements

- Writeback protocol
 - Supports caching of writes, as well as, reads
 - Gets rid of “write bottleneck”
- These changes dramatically improve the cache subsystem performance
 - Supports higher performance single processors
 - Supports higher total system performance

Model 400 Cache System



Model 500 Cache System



■ XMI Changes To Support Writeback Caching

- No change to the XMI bus
 - Any existing VAX 6000 platform can support writeback caches
- Added two block ownership commands used only by the Model 500 processor
- Writeback cache coherency is fully maintained by hardware
- Writeback support is completely transparent to I/O devices

Model 500 Cache Performance Improvement Over The Model 400

<u>Benchmark</u>	<u>Decrease in Bus Traffic</u>
<u>Dhrystones</u>	629 x
<u>Simple Technical Workload</u>	49 x
<u>Complex Technical and Commercial Workloads</u>	2-4 x

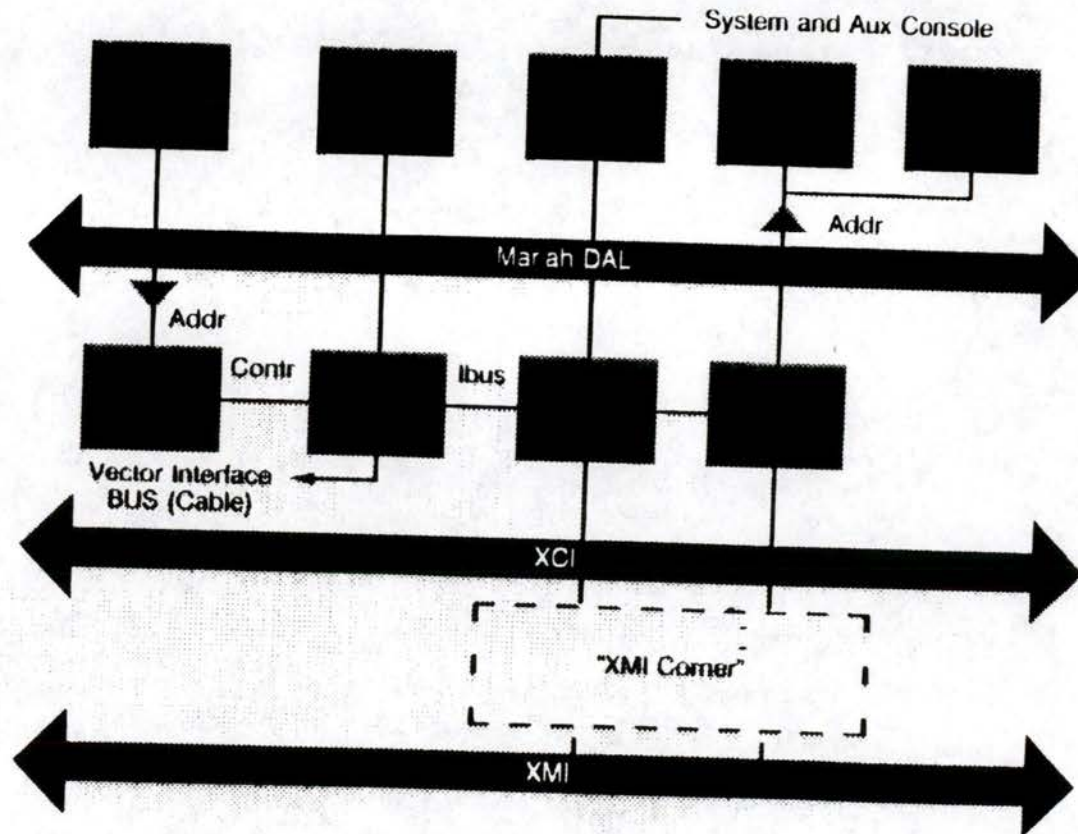
Model 500 Processor

- Processor chip set implemented in Digital's CMOS III (1u) custom CMOS process
 - 5 New chips
 - More than 1.1 million new transistors
 - This semiconductor process optimized for very fast processor cycle times
- All chips run on reduced voltage power supply (3.3v) provided by the XMI backplane
- XMI interface is implemented in two LSI Logic LL100K gate arrays (1u)

■ Model 500 Processor

- Logic design is largely a “shrink” of the Model 400 processor chip set
- Model 500 cycle time is 16ns (62.5 Mhz)
 - Fastest known CMOS cycle time in the industry
 - 75% faster than Model 400
 - 5 x faster than the VAX 6000-200

Processor Block Diagram



■ Vector Processor

- Model 400 vector processor attaches directly to Model 500 scalar processor
- Typical vector applications improve due to reduced run time on the scalar component
 - Applications like Perfect Club improve about 50%

■ Memory Improvements

- Single XMI module (called MS65A)
 - 4 x MB/slot of previous memory
 - 32 MB, 64MB, and 128MB
 - 4MB DRAMS on 64MB and 128MB variants
- Supports XMI writeback cache extensions

Memory Improvements

- Fully backwards compatible with previous XMI memory (MS62A)
 - “Mixed” configurations supported with pre-Model 500 systems
- Increased performance
 - Bandwidth (aggregate throughput) up 60%
 - Latency (read access time) down 30%
 - Improves system performance and permits more flexible interleaving rules

■ MS65A Performance Improvement Relative To MS62A

<u>Transaction</u>	<u>Bandwidth</u>	<u>Latency</u>
Hexaword Read	1.6x	1.3x
Octaword Read	1.5x	1.3x
Quadword Read	1.0x	1.3X
Hexaword Write	1.6	-
Octaword Write	1.1	-
Quadword Write	1.0x	-

■ Benefits Of Improved Memory Controller Performance

- Improved single processor performance
- Improved multiprocessor performance
- Memory configuration flexibility due to reduced memory interleaving requirements

■ Memory Interleaving

- Memory interleaving used to distribute the memory traffic between multiple memory modules
- Increased memory controller performance reduces required number of memory modules
- Customers can choose to use fewer numbers of higher density memories, preserving XMI slots for future expansion

Minimum Memory Interleaving Rules

<u>Configuration</u>	<u>Model 500</u>	<u>Model 400</u>	<u>Model 300/200</u>
X10	1-way	1-way	1-way
X20	1-way	1-way	1-way
X30	1-way	2-way	1-way
X40	2-way	2-way	2-way
X50	2-way	4-way	2-way
X60	2-way	4-way	2-way

- Why Less Interleaving For The Model 500?
- MS65A memory controller bandwidth increase is much higher for large transfer sizes
- Model 500 processor uses only large transfers
- Model 400 uses a variety of smaller size transfers
- Model 500 processor generally uses less memory bandwidth than the Model 400 even though it is running almost twice as fast

Vector Systems Require Greater Interleaving

<u>Configuration</u>	<u>Model 500</u>	<u>Model 400</u>
X10 + 1 vector	2-way	2-way
X20 + 1 vector	2-way	2-way
X20 + 2 vector	4-way	4-way
X30 + 1 vector	4-way	4-way
X40 + 1 vector	4-way	4-way

I/O Improvements

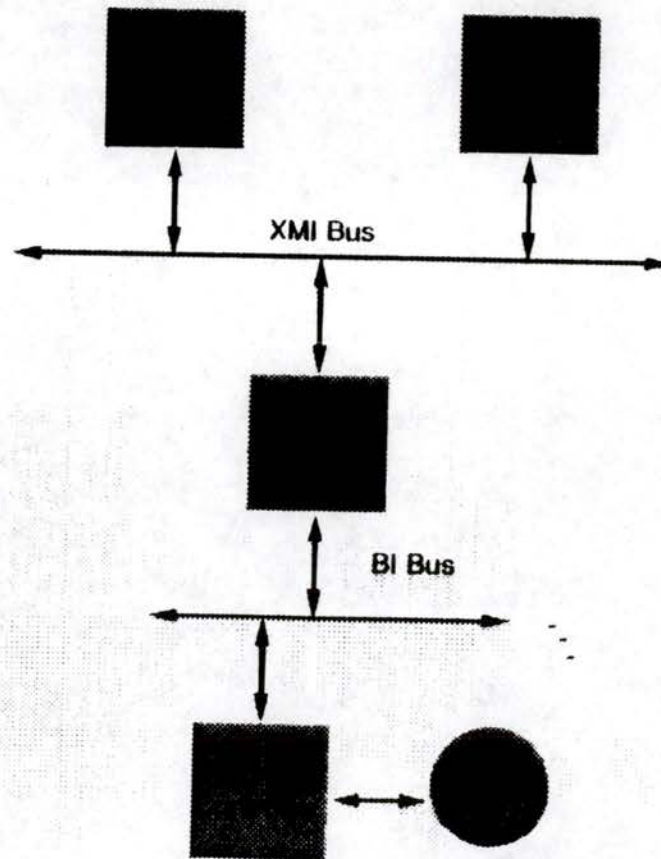
- "I/O" exercises almost all aspects of the system
- I/O performance chain includes:
 - Disk Drive Latency, I/Os, and MB per second
 - I/O Controller Latency, I/Os and MB per second
 - Main Memory Latency and MB per second
 - Processor I/O per second capability
 - Operating System

All aspects have been improved

■ I/O Performance Example

- Based on DRLABs (Digital Review) I/O load test
- Test designed to stress the complete I/O subsystem
 - Disk
 - I/O controller
 - Memory
 - Processor

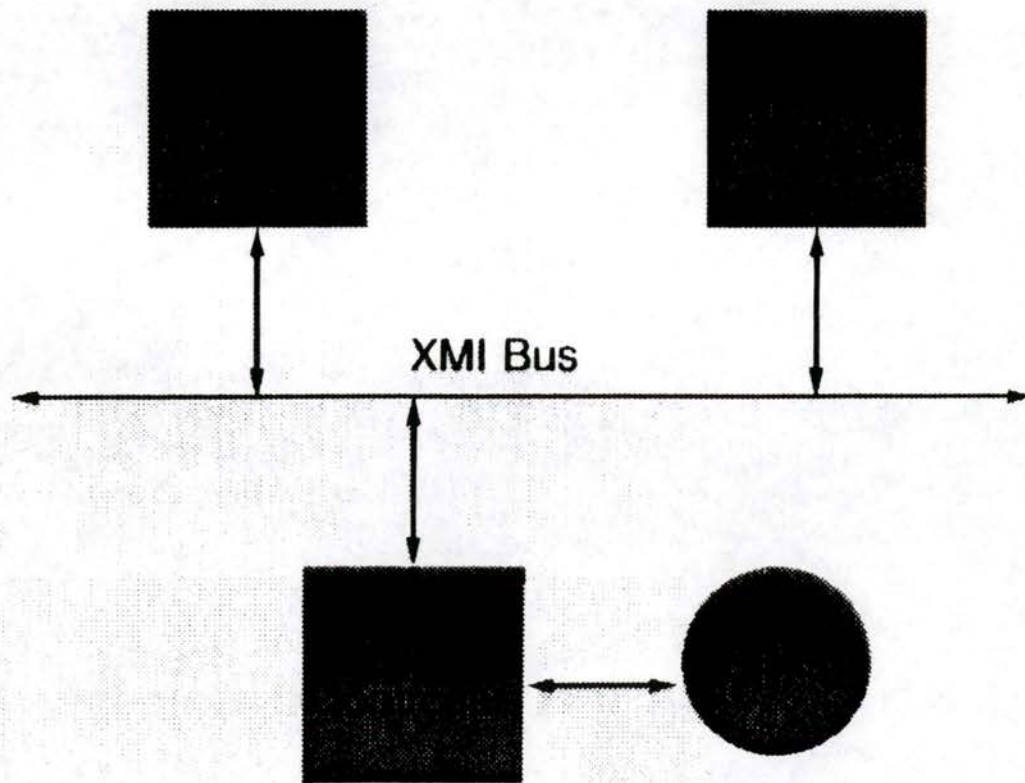
Model 400 I/O Test Configuration



Model 410 Results

- Configuration:
 - Single 410 Processor
 - 1 x CIBCA-B Cluster Interface
 - 4 x HSC70 Disk Controllers
 - 4 x ESE20 Solid-state Disk
- Maximum I/O rate: 1250 I/Os per second
- CPU was limiting factor with disk right behind

Model 500 I/O Test Configuration



Model 510 Results

- Configuration:
 - Single 510 processor
 - 2 x KDM70 disk controllers
 - 4 x ESE20 (enhanced) solid-state disk
 - Maximum I/O rate: 3388 I/Os per second
 - 2.7 x the Model 410
- CPU was the limiting factor

SEZ358.028

90Dec06

VAX 6000 Expansion

Processor Performance

1988
2.8 VUPS
Model 210



1990
72 VUPS
Model 560


SEZ358.029

90Dec06

VAX 6000 Expansion

Memory Capacity

1988
32MB
1 x MS62A




1990
512MB
4 x MS65A

SEZ358.030

90Dec06


VAX 6000 Expansion

I/O Performance



VAX 6000 Expansion

Disk Capacity



SEZ358.032

4DEC90 T

0DEC325

VAX 6000

- Expandability
- Flexibility
- Investment protection

By Design

QUANTICS

SUSAN	ALEXANDER	508 681 1533
TINA	WATSON	1547

10/91 PRE-ANNOUNCEMENT BREVARY

New VAX 6000 Model 600



**No Compromise
Datacenter Computing**

VAX/VMS Traditional Strengths




- *World's Best Distributed Computing*
- *Complete Desktop to Mainframe Compatibility*
- *Adherence to Industry Standards*
- *Investment Protection*

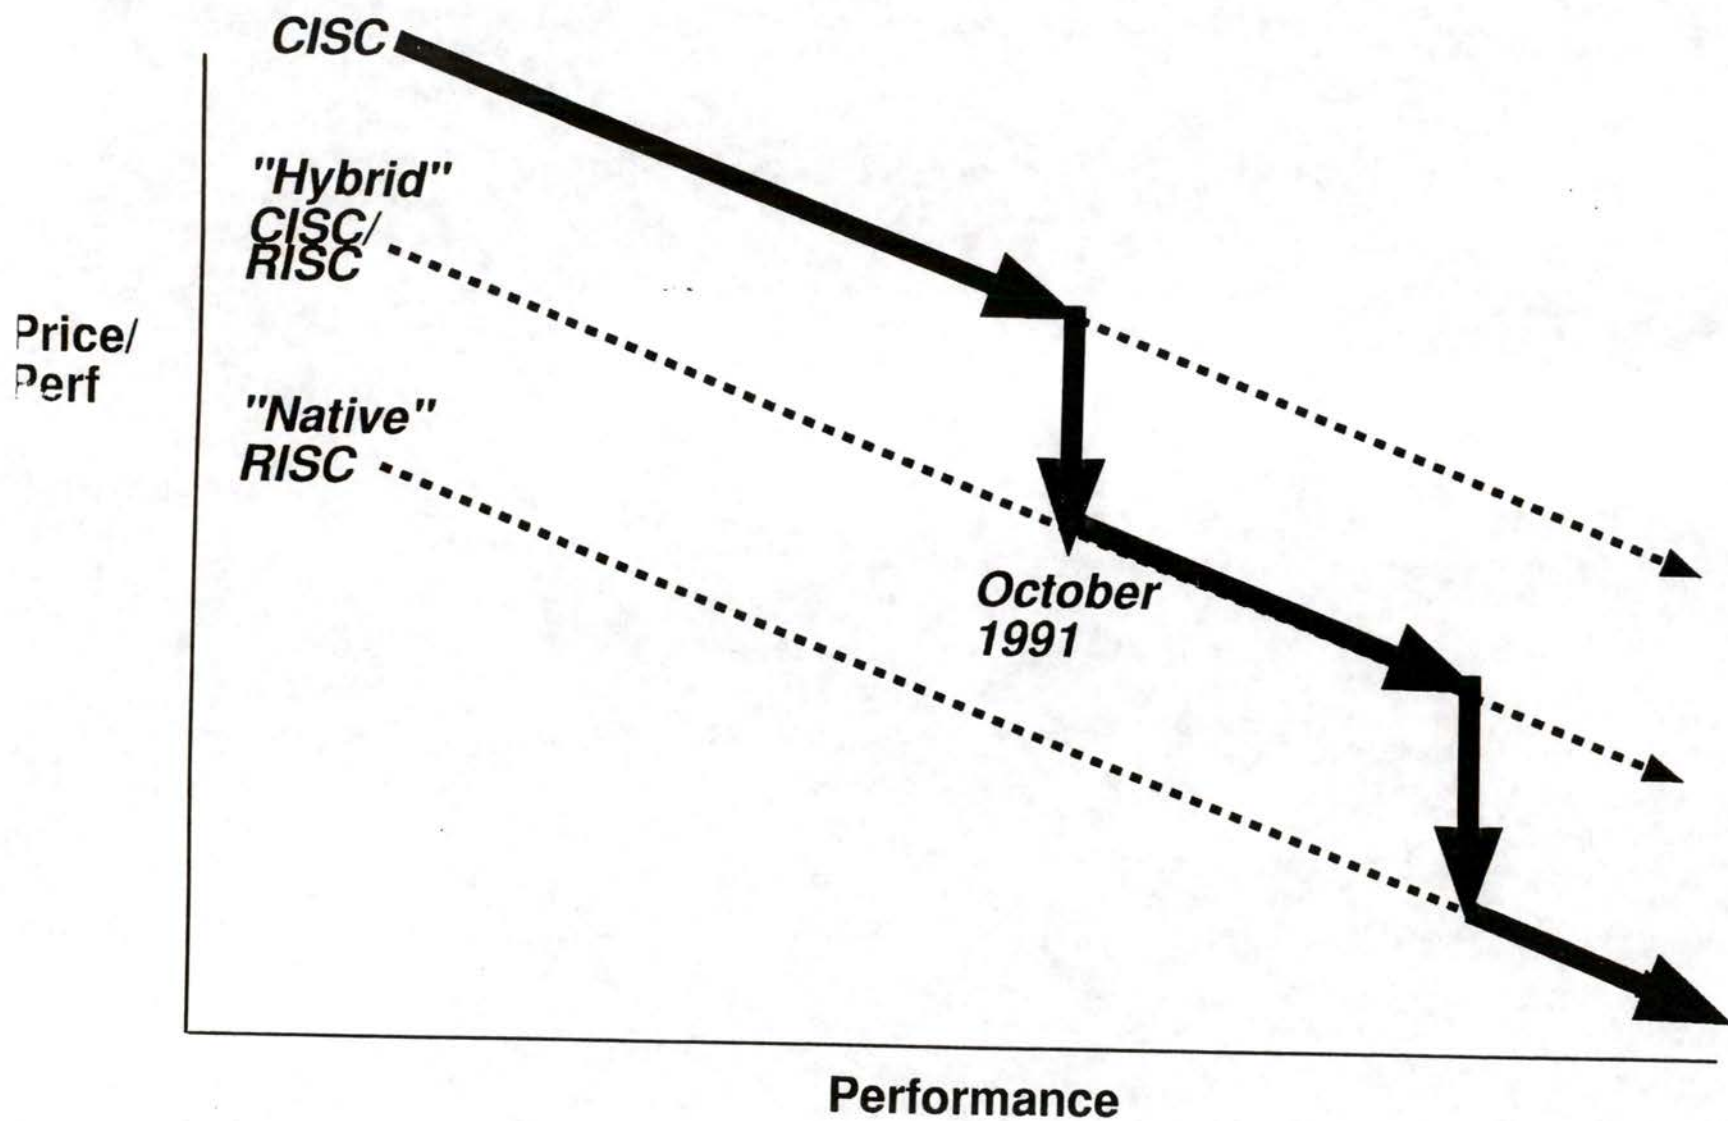
 **And Now....**

- ***Leadership Performance***
- ***Leadership Price/Performance***

New VAX Price/Performance Curve



VAX Technology Timeline



CMOS VAX Breakthrough

***VAX 9000-like
CPU
Architecture***

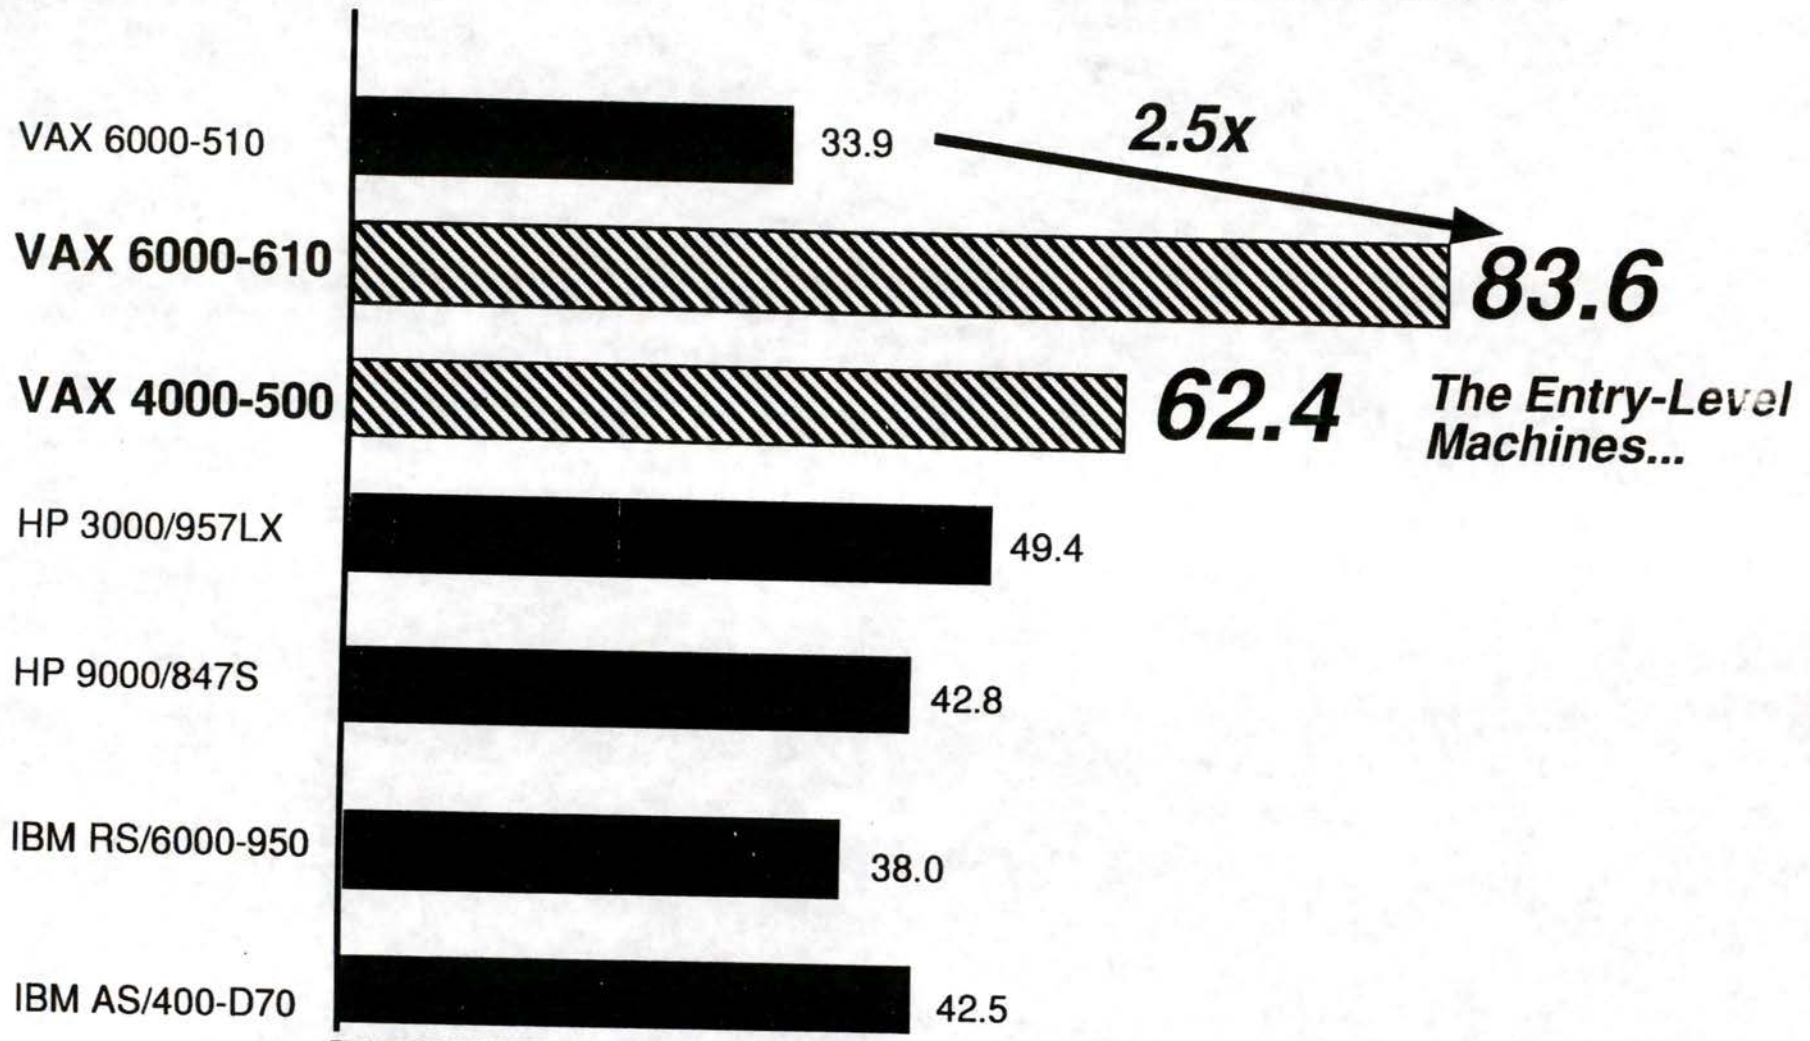


***Reads VAX CISC Instr
Internal RISC Arch.
1,300,000 Transistors
83 MHZ Clock***

***New State-of-the-Art
CMOS Process***

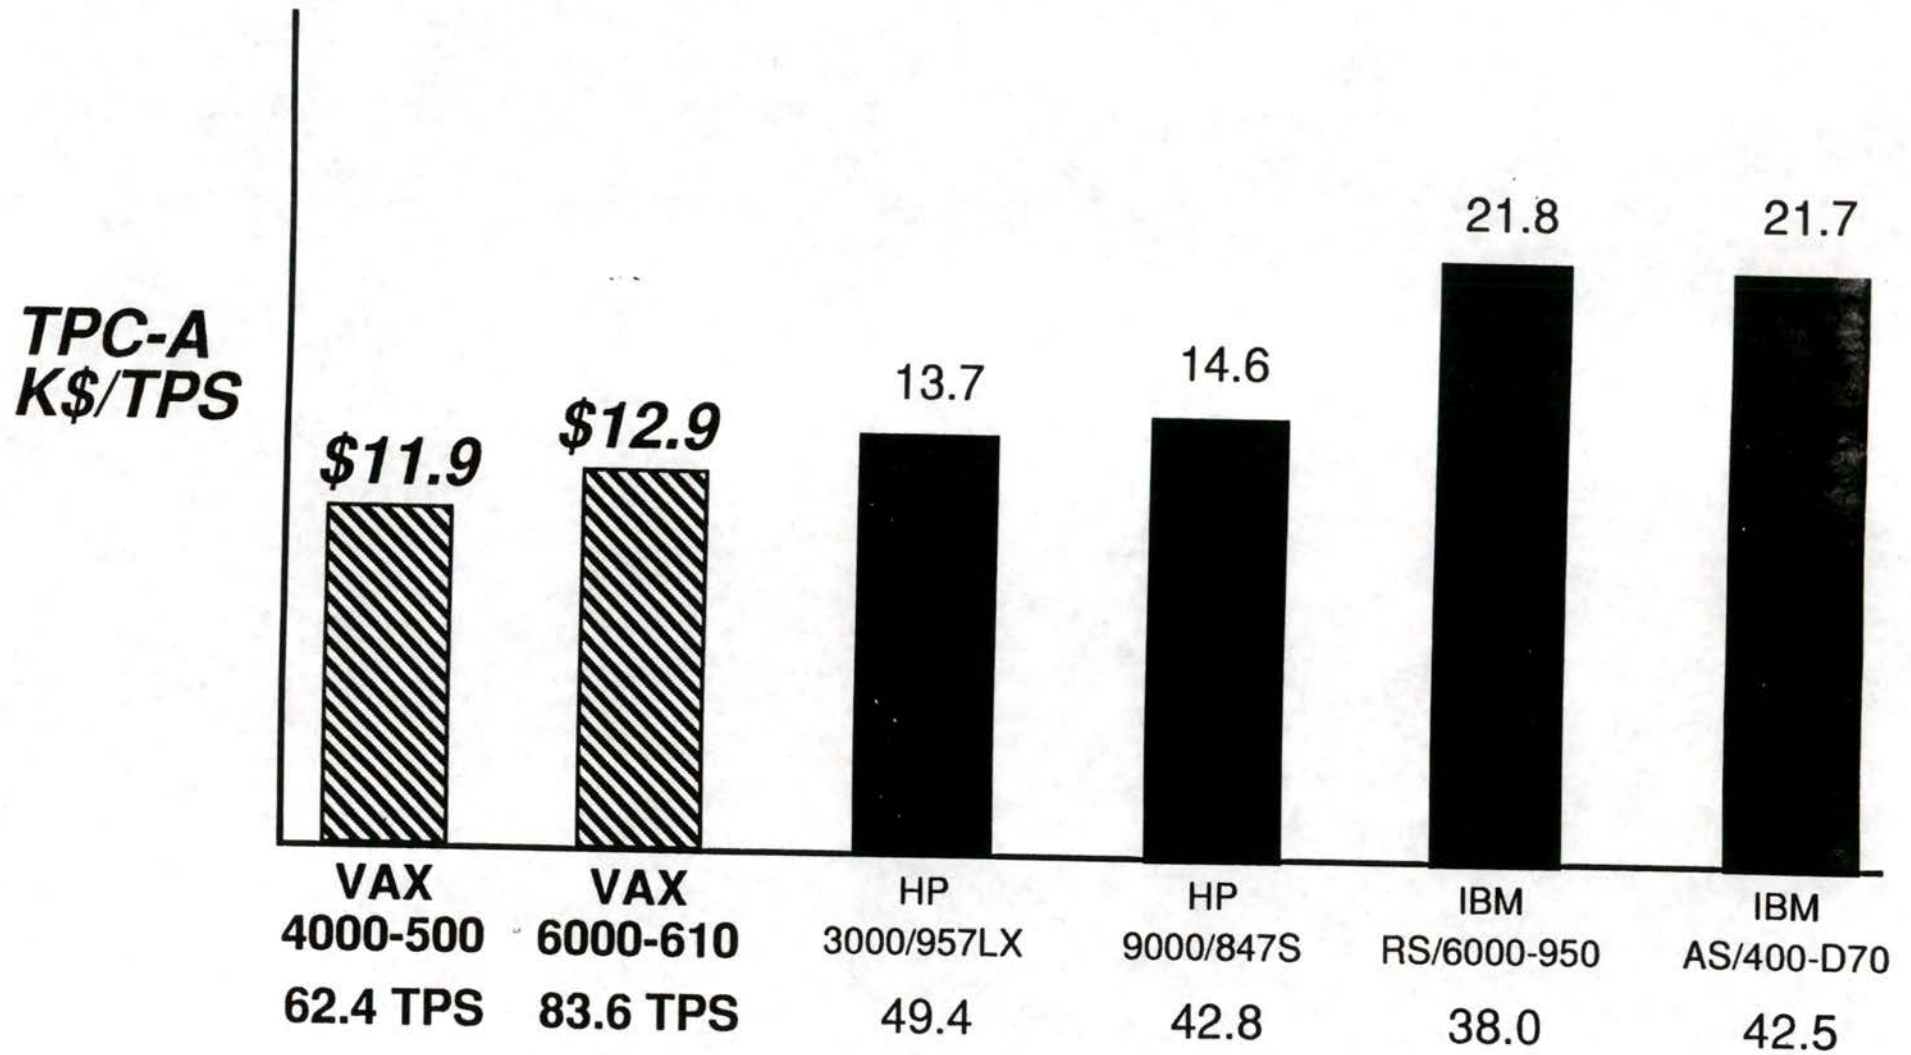


Best Commercial Performance

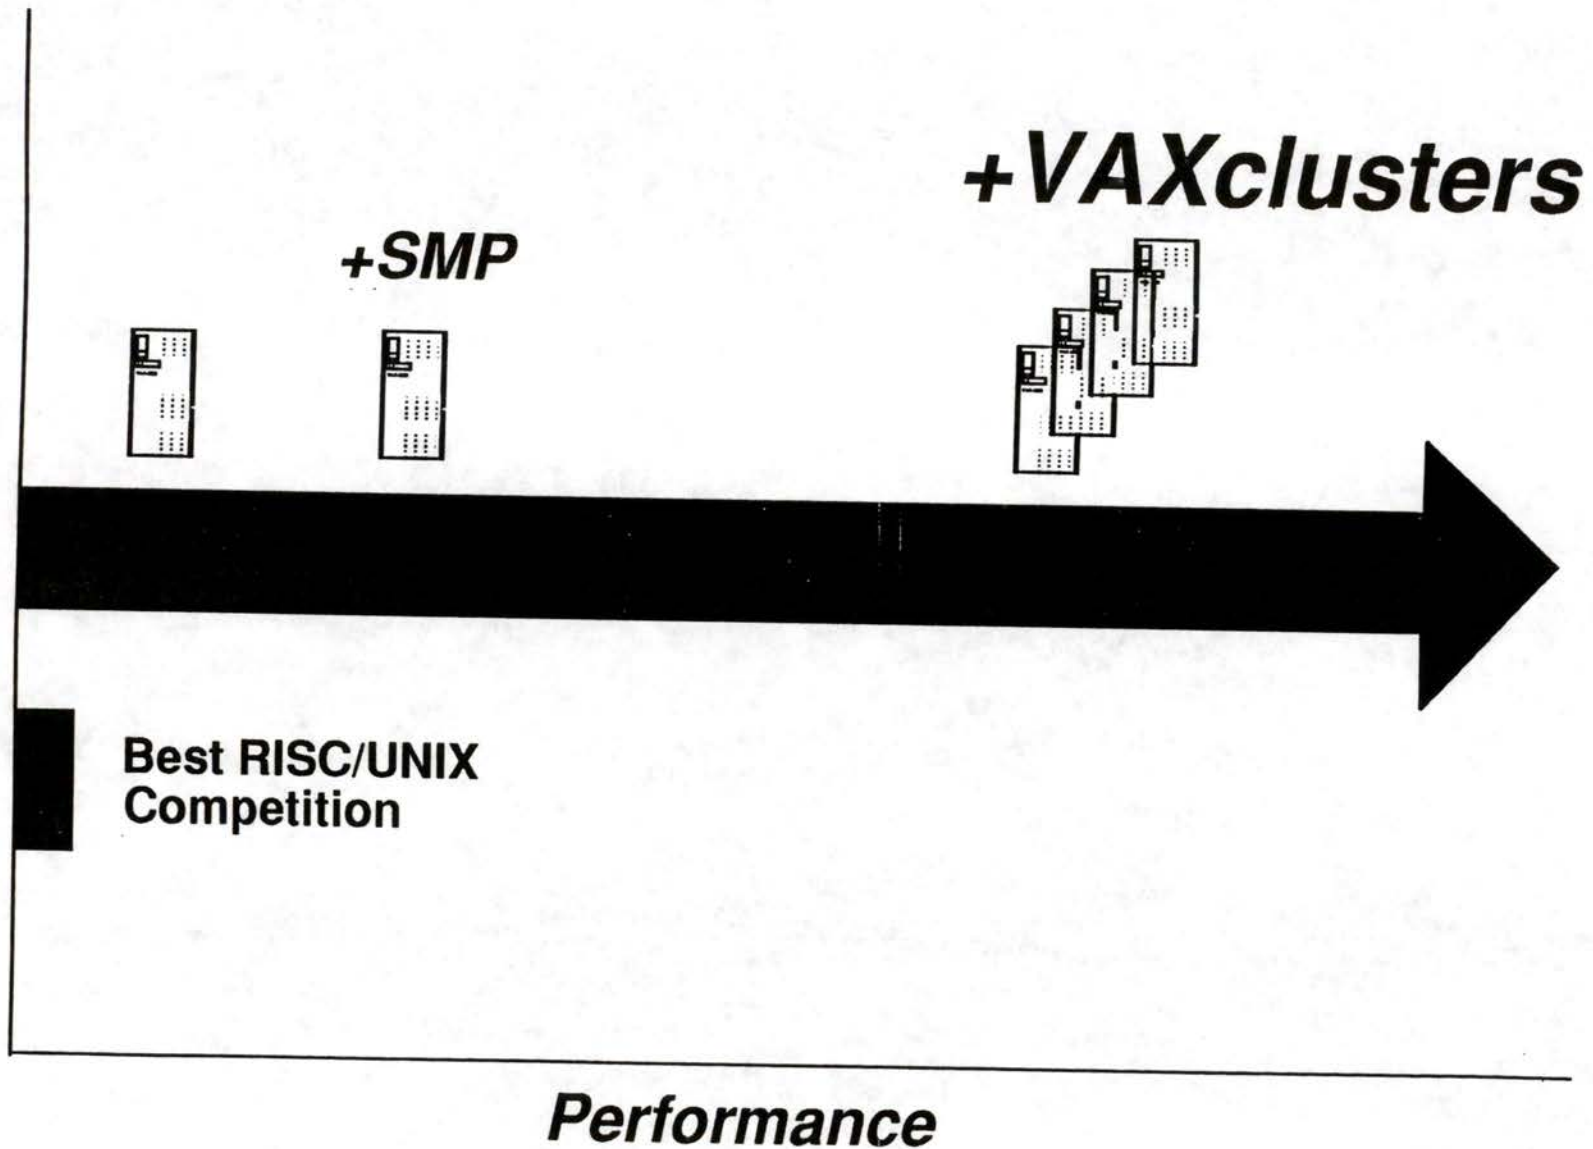


TPC-A TPS


Best Commercial Price/Performance



Leadership Expansion



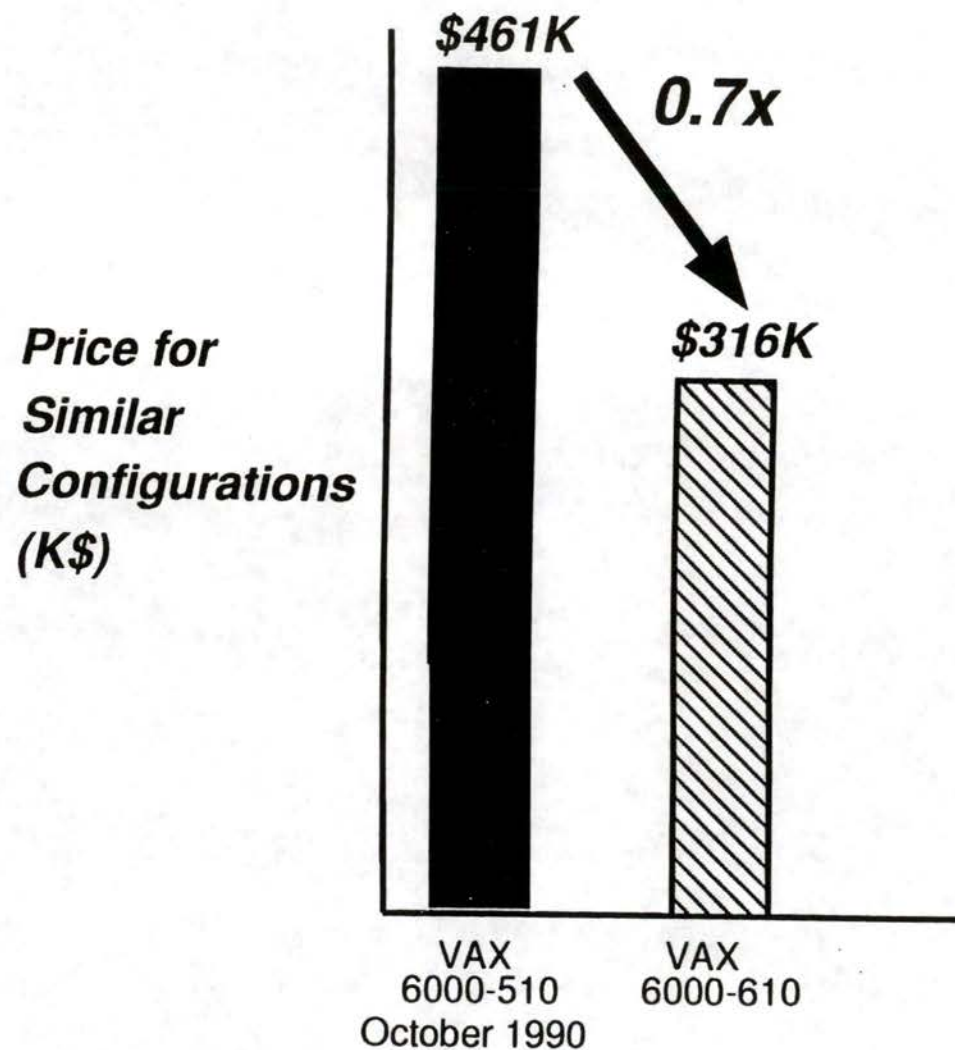
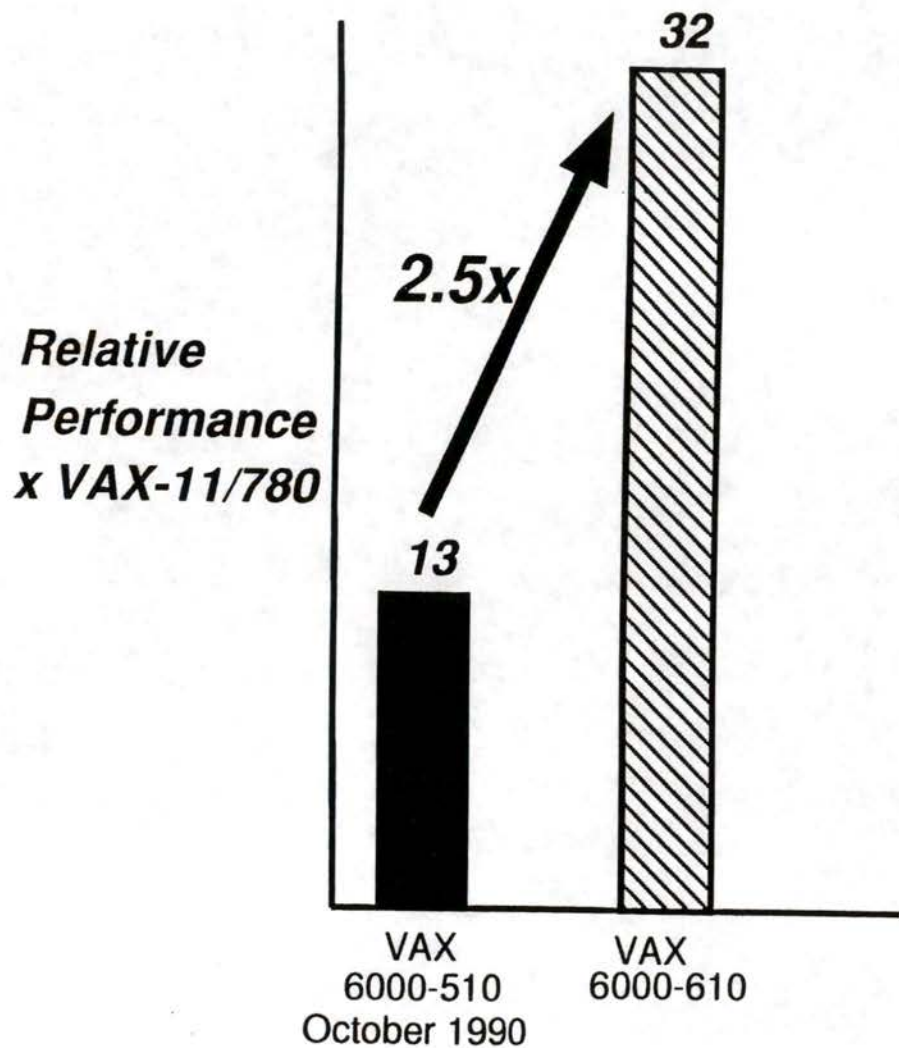
Leadership Symmetric Multiprocessing



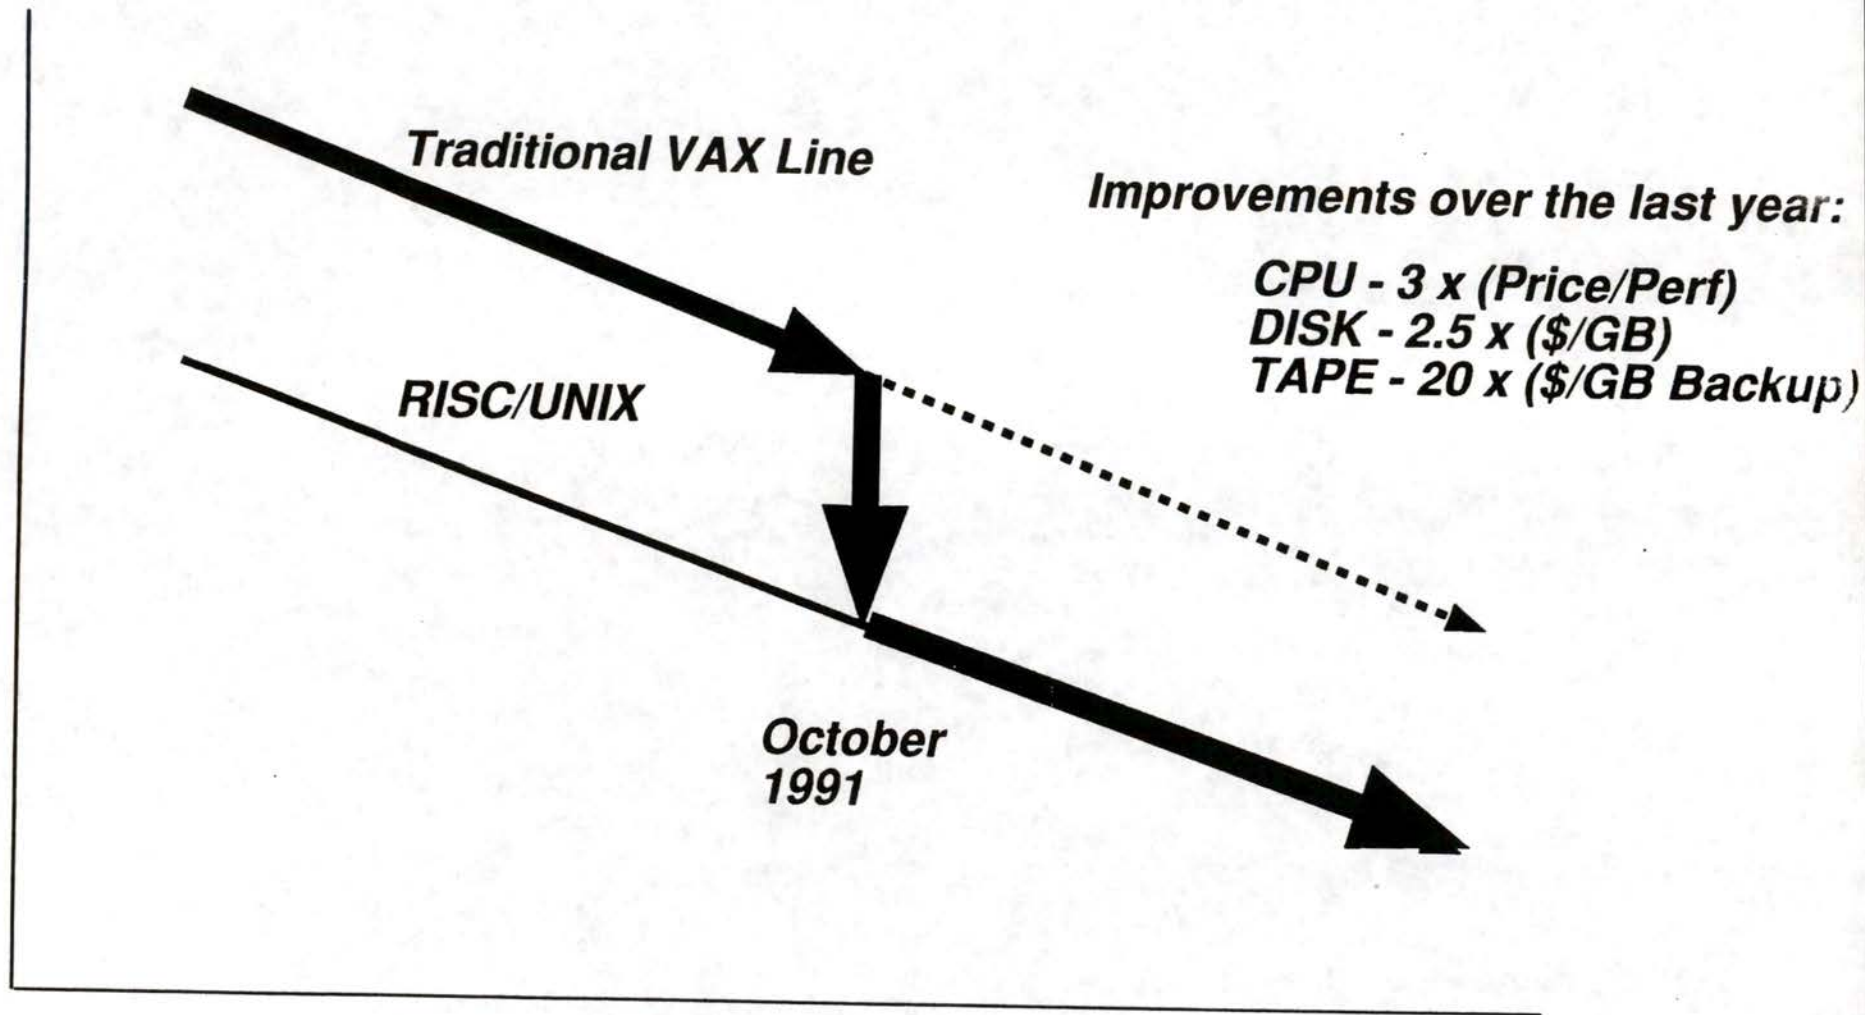
Accelerating Price/Performance

2-3 times the performance...

... At lower prices



Across the Board Improvements



VAX 6000 Family Positioning

Previous Line-up

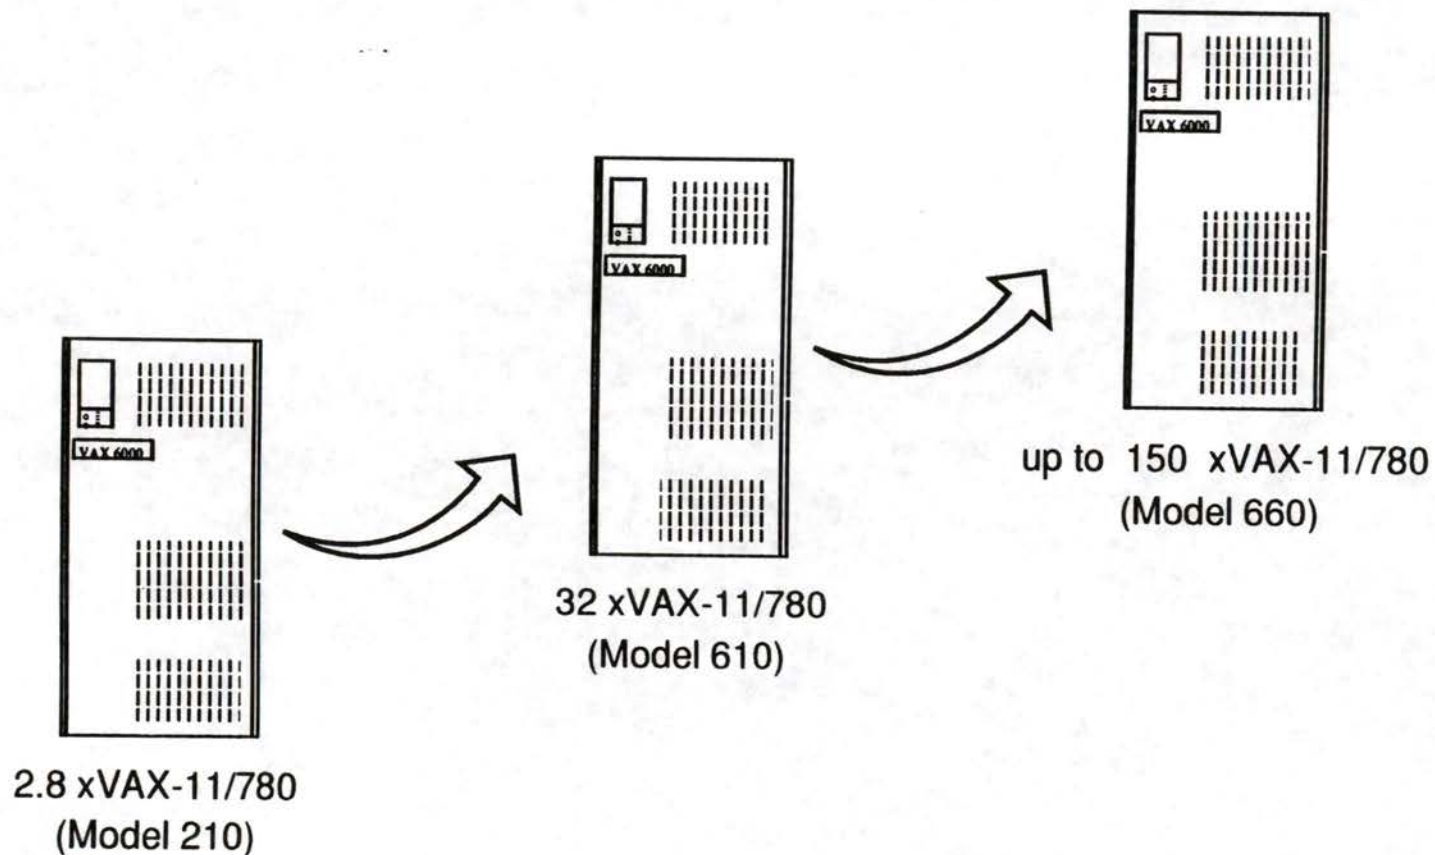
New Line-up

<u>VAX 6000 Model</u>	<u>Relative Perf.</u>		<u>VAX 6000 Model</u>	<u>Relative Perf.</u>	<u>USCLP Base Price (\$K)</u>
310	3.8	→	410	7	\$121*
410	7	→	510	13	\$155*
510	13	→	610	32	\$247*
520	up to 25	→	620	up to 58	\$359
530	up to 37	→	630	up to 84	\$434
540	up to 49	→	640	up to 106	\$509
550	up to 61	→	650	up to 128	\$584
560	up to 72	→	660	up to 150	\$659

Leadership Expansion

- Cost-effective Non-disruptive Growth

VAX 6000 systems can expand performance over 50 times in minutes



Software Performance Improvements

Database

- *RdbV4.1*


Compiler Technology

- *Cobol*
- *Fortran*
- *C*

CPU Differences - 6610 and 4500

	<i>6610</i>	<i>4500</i>	<i>Ratio</i>
<i>Frequency</i>	<i>83Mhz</i>	<i>71Mhz</i>	<i>1.17x</i>
<i>Backup Cache Size</i>	<i>2MB</i>	<i>128KB</i>	<i>16x</i>


SPEC Performance




Some Performance Comparisons

	<i>6610</i>	<i>4500</i>	<i>Ratio</i>
<i>TPC-A</i>	<i>83.6</i>	<i>62.4</i>	<i>1.34x</i>
<i>ALL-IN-1</i>	<i>348</i>	<i>264</i>	<i>1.32x</i>
<i>SPECmarks</i>	<i>40.5</i>	<i>30.5</i>	<i>1.33x</i>

VUPs Performance




Leadership Disk Arrays




- ***New RF73 Disks Provide Leadership Price/MB***
- ***2 GB Capacity per 5-1/4" Disk is First in Industry***
- ***Up to 96GB per 6000 System (>8000GB with RA92 disks)
Up to 56GB per 4000 System***
- ***High Performance and High Availability***

DSSI VAXcluster Systems




- ***Now Support Three Systems: 6000, 4000, Mixed***
- ***High Availability at Lower Cost***
- ***Another Way to Grow VAX Systems***

DSSI VAXcluster Systems



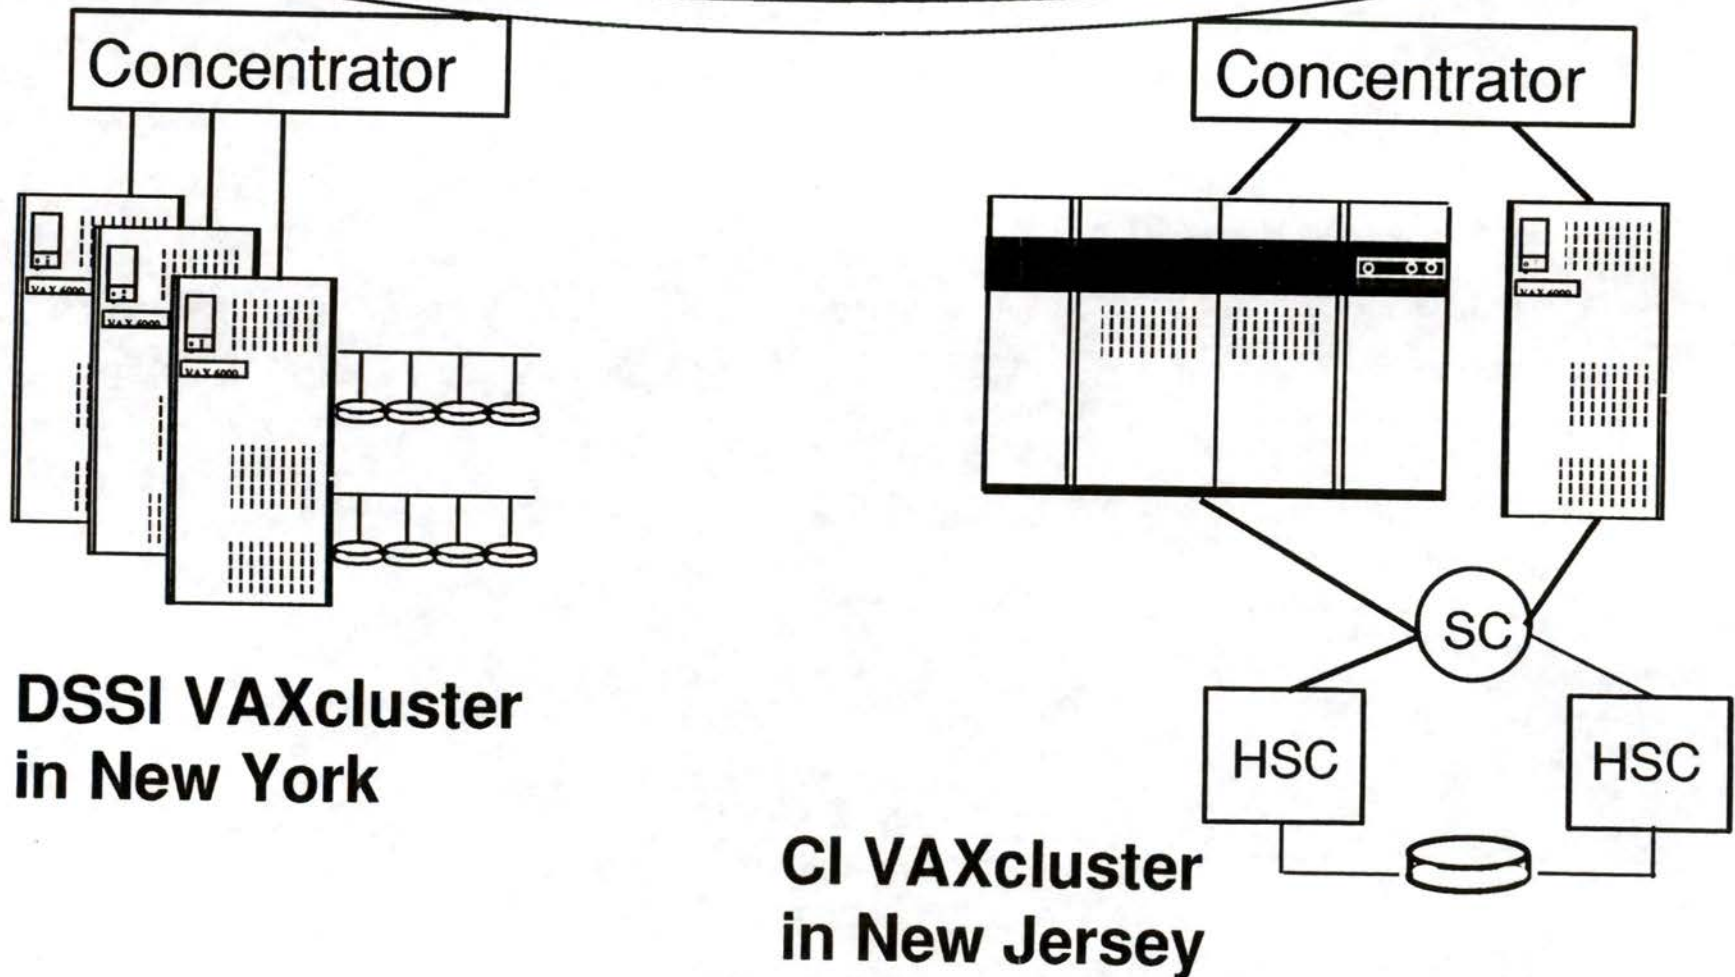
3 systems



- *Up to 450 x VAX-11/780*
- *96 GB of total disk capacity*
- *"Warm-swap" for maintenance and upgrades*
- *Aggregate DSSI bandwidth of 48 MB/sec*
- *Direct access to CI, FDDI, VME*

Multi-Datcenter VAXcluster Systems

- 100 Mb/s open long-distance network based on FDDI

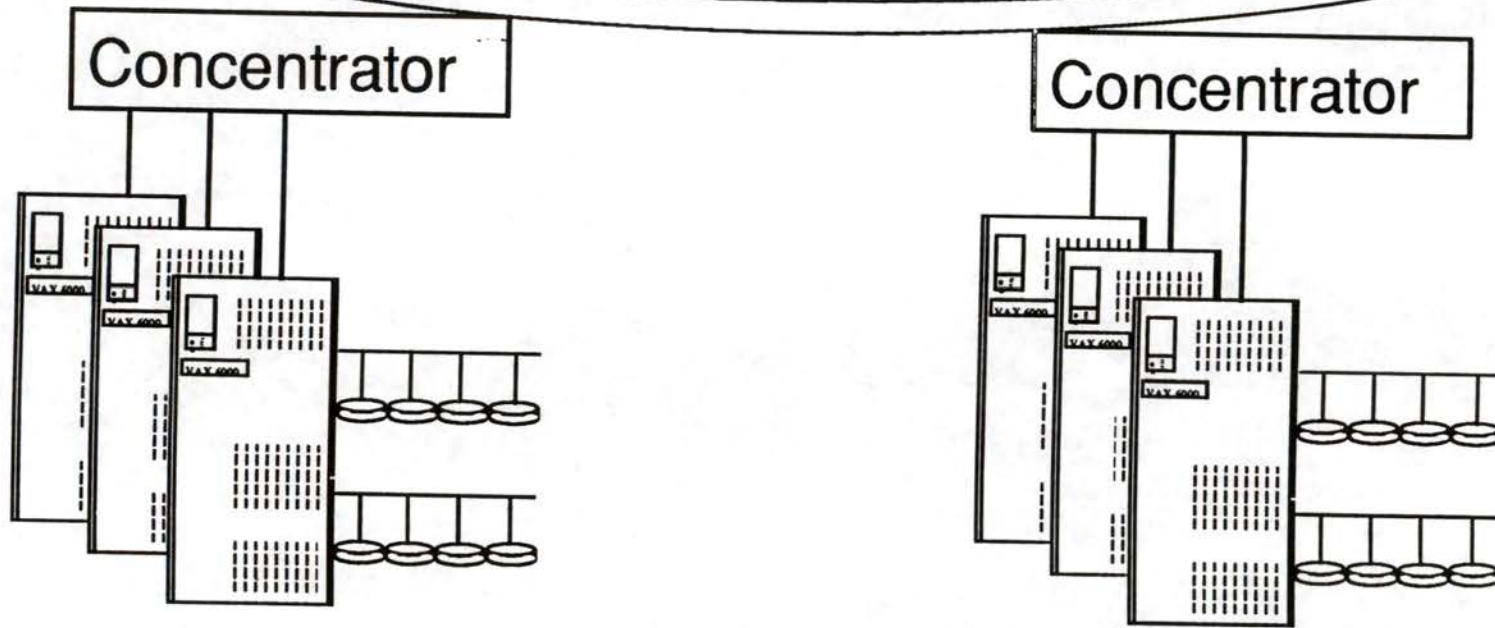


**DSSI VAXcluster
in New York**

**CI VAXcluster
in New Jersey**

Disaster Tolerant Systems

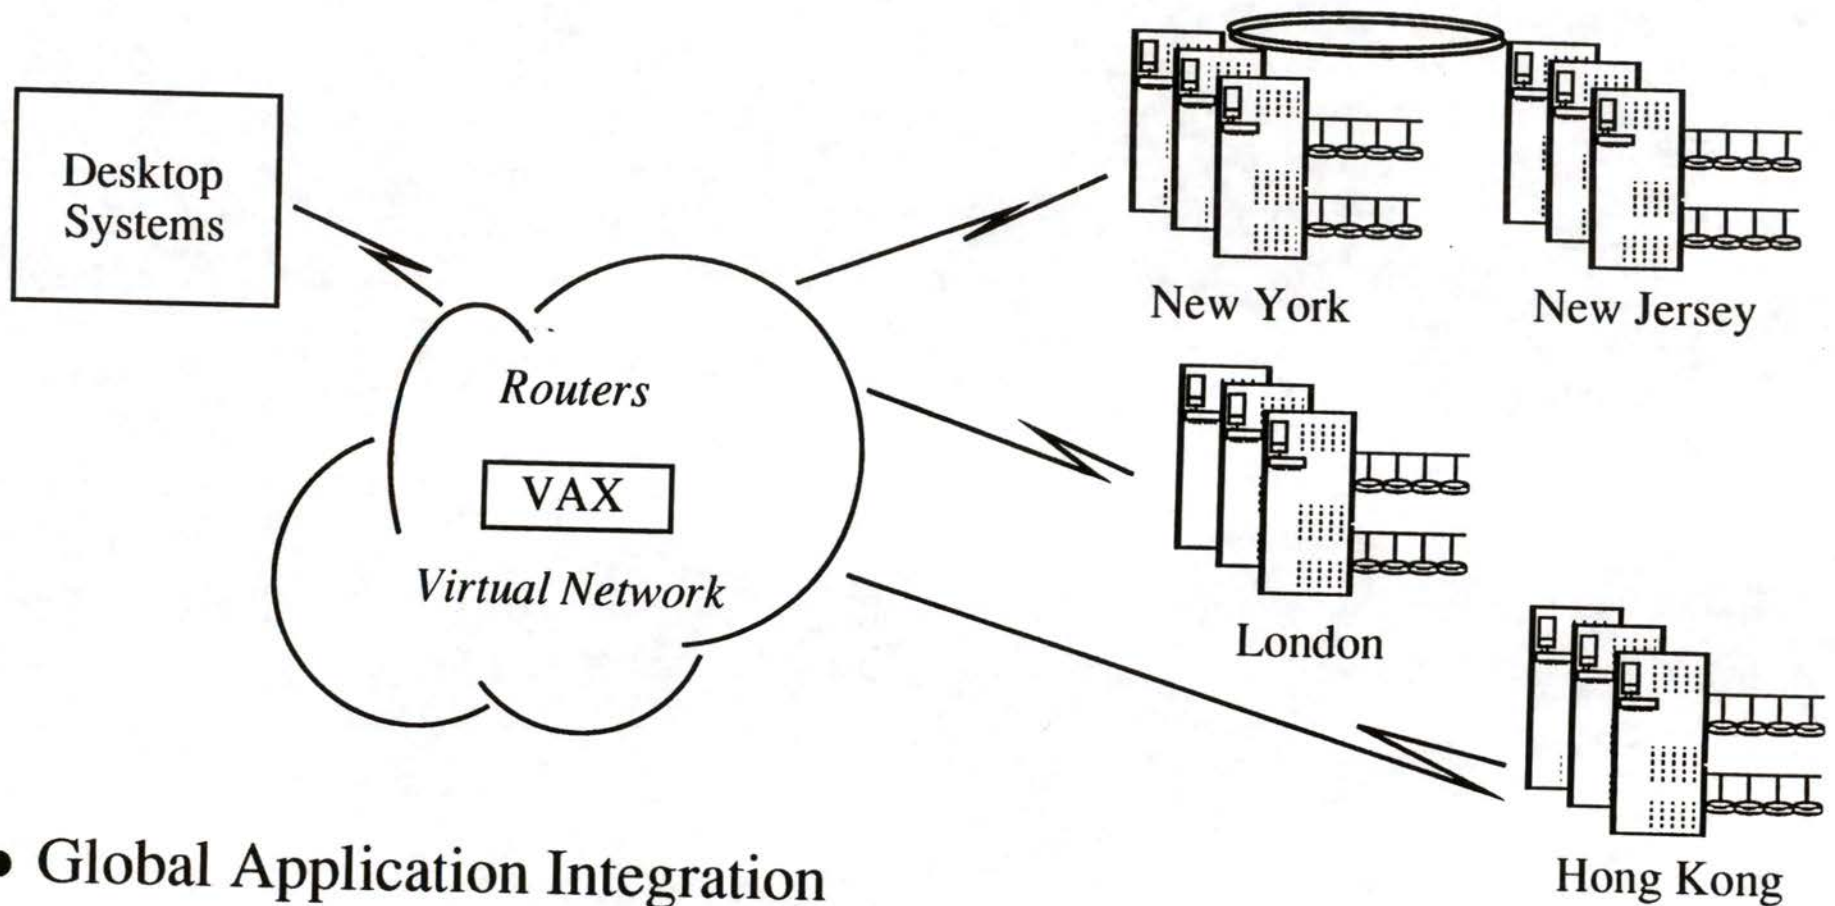
- 100 Mb/s open long-distance network based on FDDI



**VAXcluster
with Volume Shadowing
in Hong Kong**


**VAXcluster
with Volume Shadowing
in Kowloon**

DEC Reliable Transaction Router™



- Global Application Integration
- Assures transaction completion across widely distributed applications
- Transparent Failover/recovery of routers and servers


Open VME Bus Support



VME - More Industry-
Standard Options

- *Support for ANSI IEEE Standard for open access*
- *Over 4000 VME boards available*

ULTRIX Systems



Support for Both Computing Environments

Continuing Investment Protection

Every VAX 6000 Platform can:

- ↪ Upgrade to most powerful processors***
- ↪ Attach newest low-cost disks and I/O devices***
- ↪ Use latest VAXcluster technology
(DSSI, FDDI VAXclusters)***
- ↪ Take advantage of new user licensing and
new applications based on NAS***


No Compromise Datacenter Computing

● Price/Performance

- Leadership Performance
- Leadership Price/Performance
- Symmetric Multiprocessing
- User Licensing
- RF73 Disk Arrays

● Functionality and Investment Protection

- CPU Expansion
- 1 GB Memory
- DSSI/FDDI VAXclusters
- Disaster Tolerant Systems




● Openness

- VME Bus
- ULTRIX
- Open VMS with NAS

● Datacenter Solutions


- Desktop Integration
- Client/Server
- Business-Critical

No Compromise Datacenter Computing




- Solutions for New Applications

- Competitive Migrations
- Datacenter Solutions



VAX 6000

**Performance
Price/Performance
Functionality and Quality
Openness**



- Upgrades for Existing VAX VMS Applications

- VAXclusters
- CPU Upgrades

FALL DECUS 1991

TECHNICAL DESCRIPTION OF VAX 6000 ENHANCEMENTS

Richard Gillett
VMS Systems and Servers Engineering

Rick Gillett

DECUS

December 1991

Presentation Outline

- ***Engineering Goals / Technology Context***
- ***Technical Challenges***
- ***Technical Solutions***
- ***Actual Results***

System Design Goal

- ***Extend the VAX 6000 Family in all dimensions while maintaining traditional "just swap modules" upgradability for existing VAX 6000 customer***
- ***Firmly establish VAX/VMS as the commercial performance and price/performance leader***

XMI Block Diagram- October 1990

Rick Gillett

DECUS

December 1991


Specific Design Goals

- ***2-3 times increase in processor performance***
- ***Only swap CPU module***
- ***Add direct XMI FDDI and VME connections***
- ***Extend DSSI technology (tri-host, RF73)***

VAX Technology Plans

- ***Long-term RISC-VAX strategy***
- ***Compliment with a parallel strategy for a new high-performance CMOS VAX processor***

CISC to RISC Continuum



VAX 9000 is a hybrid architecture

Mike's #10

Rick Gillett

DECUS

December 1991

Perfect Combination

- ***VAX 9000 Architecture***
- ***CMOS Technology***

Increasing Processor Chip Performance

- ***Increase clock rate***
- ***Increase work done per clock cycle***
- ***Do Both***

Clock Rate Opportunities

- ***Faster CMOS process (CMOS IV)***
- ***Integration reduces off-chip critical paths***
- ***Faster designs through better CAD tools***
- ***...Improvements of 30-40% possible***

Work Done Per Cycle Opportunities

- ***Jump to macro-pipelined "RISC-like" design***
- ***Use integration to reduce chip-to-chip communication overhead***
- ***Reduce memory stalls by dedicating more devices to on-chip caching***
- ***...Improvements of over 200% possible (variable)***

What is Pipelining ?

- ***A technique to increase parallel execution***
- ***Performance vs complexity tradeoffs***
- ***VAXes are tough to pipeline...***

Micro vs Macro-pipelining

- ***Micro-pipelining permits parallel execution **within** a given VAX macro instruction***
- ***Macro-pipelining permits parallel execution **across** multiple VAX macro instructions***

Macro-pipelining

- ***Autonomous functional units (I, E/F, M/C)***
- ***Each maintains states needed to operate across instruction boundaries***
- ***Micro-pipelining is used within each functional unit***

History of VAX Chip Pipelining

- ***MicroVAX*** ***No pipelining***
- ***CVAX (6000-200/300)*** ***Partial micro-pipelining***
- ***RIGEL (6000-400/500)*** ***Full micro-pipelining***
- ***NVAX (VAX 6000-600)*** ***Macro-pipelining***

Managing Architecture Complexity in CMOS VAXes

- ***Copy ECL design concepts whenever possible***
- ***Patchable microcode is a requirement***
- ***Implement very thorough design verification plan***
- ***Balance complexity/time-to-market/performance requirements***

Implementation (CHIP/WITH HEAT SINK PICTURE)

- ***Fastest CMOS clock rate in the industry***

Platform Challenges of Increased Processor Performance

- ***Impact of main memory latency (single processor)***
- ***Bus bandwidth requirements (multiprocessor)***
- ***Memory capacity***
- ***I/O capacity***

Cache Background - Write-through

Rick Gillett

DECUS

December 1991

Cache Background - Writeback

Rick Gillett

DECUS

December 1991

Cache Improvements (CHART) cache_miss.chart

- ***Plot Memory Bandwidth Consumed vs Cache Size***
- ***2x reduction in memory bandwidth for each 4x increase in cache size***
- ***Write-through flattens out after 128KB.....***

Increased Integration

- ***BLOCK DIAGRAM OF 6000-500 with NVAX and NEXMI functions outlined***

Logic Partitioning Changes

Function	6000-500	6000-600
Instruction Decode	P-chip \	
Instruction Execution	P-chip \	---- NVAX-chip
Cache Control	C-chip /	
Floating-Point	F-chip /	
Bcache Tag Store	C-chip	64K x 4
Bcache Data Store	64K x 4	256K x 4

Logic Partitioning Changes

<i>Function</i>	<i>6000-500</i>	<i>6000-600</i>
<i>Console Support</i>	<i>SSC-chip \</i>	
<i>XMI Control</i>	<i>MCA-chip >---</i>	<i>NEXMI-chip</i>
<i>XMI Datapath</i>	<i>MDA-chip /</i>	
<i>Total VLSI Chips</i>	<i>6</i>	<i>2</i>

Block Diagram of 6000-600 CPU Module

Rick Gillett

DECUS

December 1991

VAX 6000-600 CPU Module

- ***Simply swap CPU modules***

Commercial Performance

- ***Highest Single Processor TPC-Atm TPS***
- ***83.6 TPC-A TPS***
- ***\$12.9K/TPS***
- ***SMP Test Results are coming***

Commercial Performance -

- ***Highest Single Processor TPC-Btm TPS***
- ***153 TPC-B TPS***
- ***\$4.4K/TPS***

Technical Performance

- ***40.5 SPECmarks***
- ***196 SPECthruput***

XMI Block Diagram - October 1991

Rick Gillett

DECUS

December 1991

VAX 6000 Performance History

- ***6000_VUPS_HISTORY.CHART***
- ***11.5x over 3.5 years ==> >100% annually***

1988 to 1991 - VUPS

Rick Gillett

DECUS

December 1991

1988 to 1991 - TPS

Rick Gillett

DECUS

December 1991

1988 to 1991 - Memory

Rick Gillett

DECUS

December 1991

1988 to 1991 - Disk Storage

Rick Gillett

DECUS

December 1991


VAX 6000

- ***Expandability***
- ***Flexibility***
- ***Investment Protection***
- ***....By design***

VAX 6000 Expansion Guide

A Guide to Increasing the Power and Performance of
Your VAX 6000 System

digital



Higher Performance Processors,
Symmetric Multiprocessing,
and Vectors



Networking



Memory and I/O
for a Balanced
System



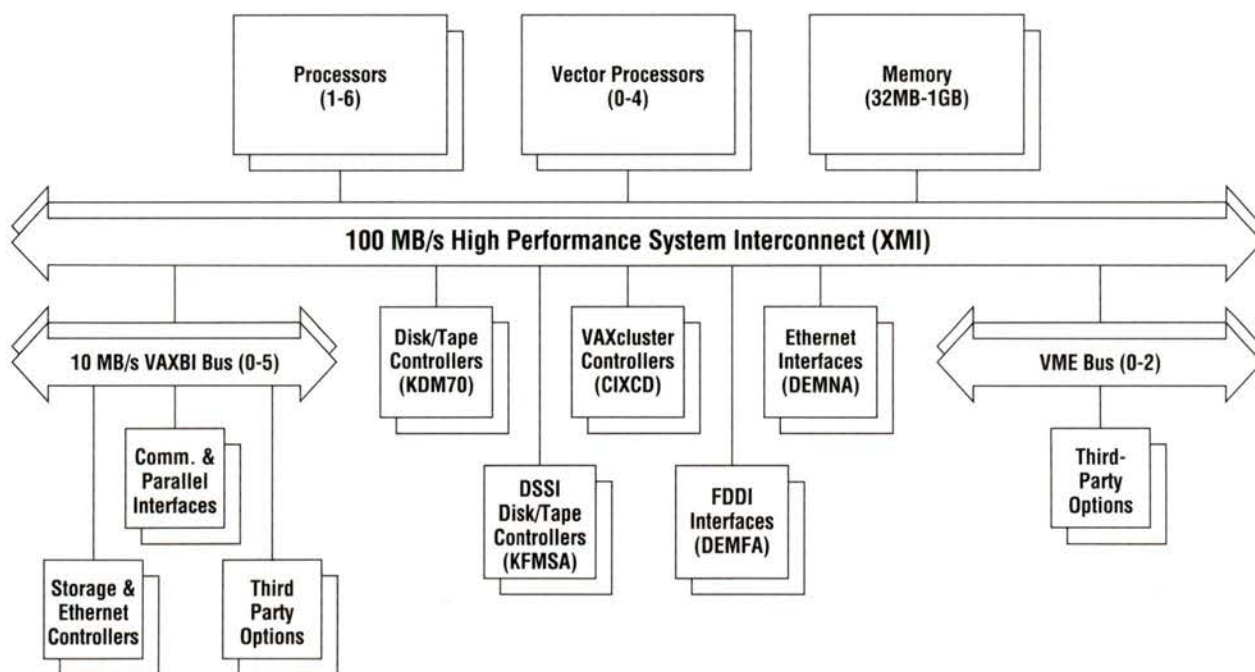
VAXclusters

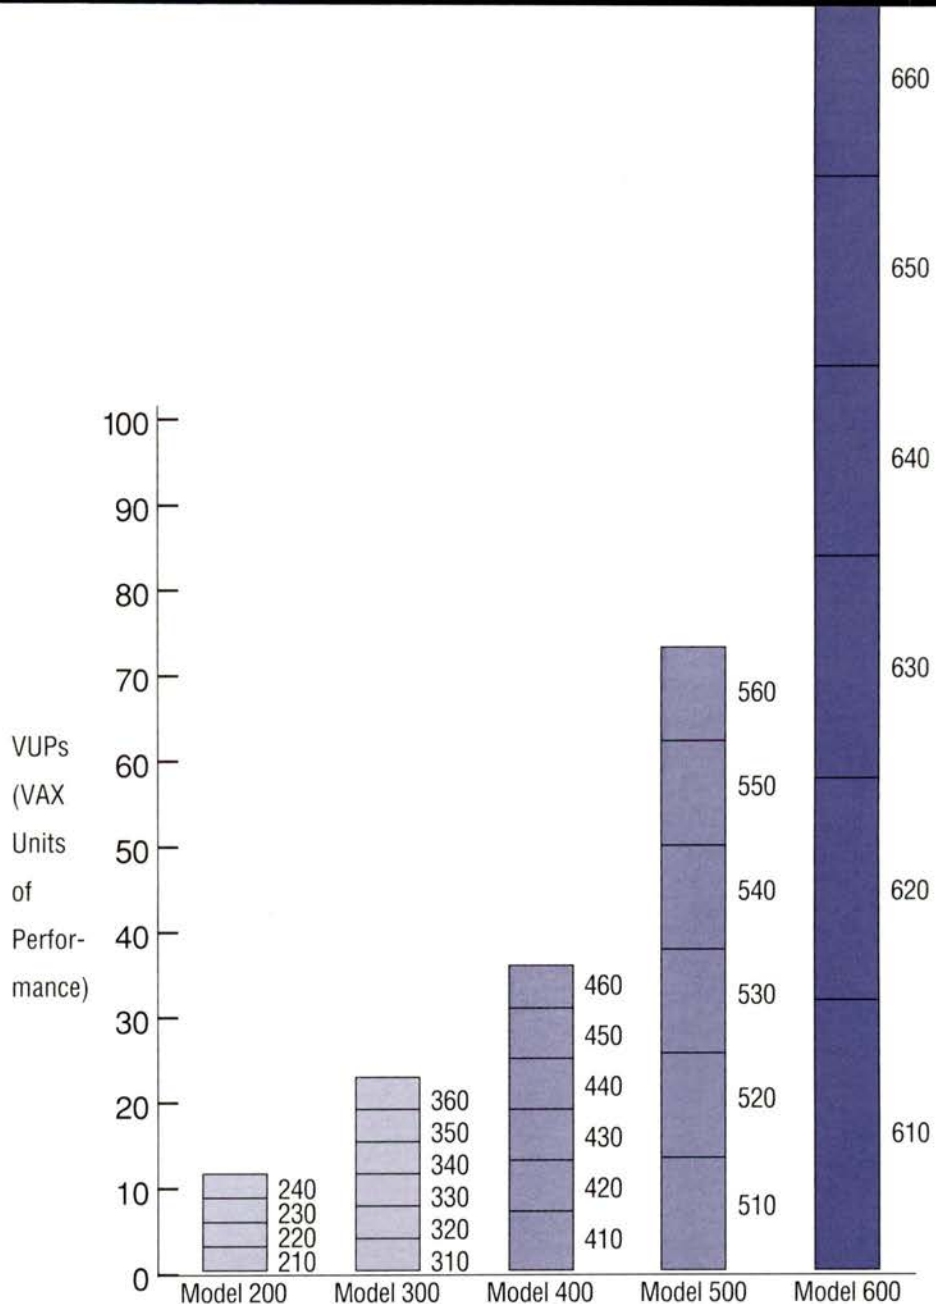
As a current owner of one or more VAX 6000 systems, you are no doubt well aware of its modular design and of the easy and economical expansion capabilities this design makes possible. However, you may not be aware of just how much more power and performance is now available to VAX 6000 owners.

The introduction of the most recent models of the VAX 6000 series and the VAX 6000 DSSI VAXcluster configuration with the economical DSSI storage options represent some of the latest expansion possibilities available for this versatile system. Also available is a Fiber Distributed Data Interface (FDDI) controller that allows the VAX 6000 system to connect directly to a high-performance fiber optic network. With FDDI, you can configure a VAXcluster with systems in different locations, or even link two entire VAXclusters that are located 40 km apart. In addition, you have a choice of several load devices offering improved software distribution and data access capabilities.

Today, a VAX 6000 Model 660 is capable of 150 VUPs, or 150 times the processor performance of the VAX-11/780.

The purpose of this guide is to bring you up-to-date on these expansion options, and to help you make the right expansion decisions. It is designed for use after you have made a thorough assessment of your organization's needs, and have determined that expanding your VAX 6000 system is necessary based on current and projected system demands. We encourage you to ask your Digital sales representative for assistance and use Digital's VAX Performance Advisor (VPA) and System Performance Monitor (SPM) to help you make the important decisions around when and how to expand your system capabilities.





VAX 6000 Expansion Overview

The VAX 6000 system was introduced in 1988 as an expandable system platform based on a modular design. This platform strategy has allowed VAX 6000 system owners to constantly accelerate the performance of their systems while maintaining the value of their initial hardware investment.

The first model VAX 6000 system had a performance rating of 2.8 VUPs (VAX Units of Performance), or 2.8 times the processor performance of the VAX-11/780. Today, a VAX 6000 Model 660 system is capable of 150 VUPs, or 150 times the processor performance of the VAX-11/780. This dramatic performance increase has been made possible through the development of faster processors and the use of symmetric multiprocessing (SMP), which allows up to six processors to reside in the same system cabinet.

Vector processing is an option for owners of VAX 6000 Model 400 and 500 series systems. An integrated VAX vector processor can offer 3 to 20 times faster execution time over a scalar processor for operations running compute-intensive FORTRAN or APL applications.

**Within minutes,
you can expand a
VAX 6000 Model 410
to a Model 660.**

As your need for processor power increases, so too does your need for additional memory, storage capacity, and I/O speed in order to achieve a truly balanced system. With the VAX 6000 system, you can have up to 1 Gbyte of memory and a full range of storage products – both internal and external to the CPU cabinet – to meet increased system performance and capacity requirements. You can increase I/O speed by adding high-performance controllers or up to five VAXBI buses to the high-speed VAX 6000 XMI system interconnect. Up to two VME bus interface adapter modules are also optionally available, if you need to connect industry-standard peripherals.

VAXcluster configurations – DSSI-based, CI-based, and FDDI-based VAXcluster systems – are three additional ways to expand the capabilities of the VAX 6000 system. Adding a VAX 6000 system to a VAXcluster enables efficient, economical resource sharing, plus higher availability and higher storage capacity.

And finally, you can connect your VAX 6000 system to an Ethernet network, a high-performance FDDI network, or a wide area network that can grow as large as you need it to for access to nationwide, or even worldwide, resources and information.

VAX 6000 Platform Strategy in Action

Perhaps you now have a VAX 6000 Model 410 system that you use for basic processing. Within minutes, you could expand it to a Model 610 system and get increased CPU performance, I/O, and memory. Then, for even more growth within the cabinet, you could add an additional CPU board to upgrade to a symmetric multiprocessing Model 620 system.


If high availability became a requirement, you could add one or two additional nodes to create a DSSI VAXcluster configuration, then grow to a CI VAXcluster. And as your needs grew to nationwide proportions, you could connect the cluster to a Wide Area Network (WAN).

And these are just some of the many growth paths you could take. As your needs change and grow, the VAX 6000 system is designed to keep up with them with a wide range of expansion options.

The following pages will cover each of the following expansion possibilities in detail.

• CPU Upgrades	Page 4
• Vector Processing Options	Page 6
• Memory Expansion	Page 8
• I/O Expansion	Page 10
• Storage Expansion	Page 12
• VAXcluster System Options	Page 16
• Load Device Options	Page 18

CPU Upgrades



The VAX 6000 system's modular design allows expansion of up to 50 times the CPU power of your original system within the same cabinet.

Perhaps you have recently added more users or applications to your VAX 6000 system, changed your desktop topology by replacing terminals with workstations, or moved to a windowing environment. If so, you are probably going to need to enhance performance to keep your VAX 6000 system running smoothly and your users happy and productive.


When you need more performance, adding CPU power is often where you start. The VAX 6000 system's modular design allows expansion of up to 50 times the CPU power of your original system within the same cabinet.

You merely add or replace processor boards, depending upon the applications you are running. If your operation demands superior multistream performance, you can use the symmetric multiprocessing feature, and add more processors. If superior single-stream performance is what you need, you can replace your existing CPU boards with faster processors. With models available in five basic processor families, VAX 6000 systems offer a wide range of performance—from the Model 210's 2.8 VUPs to the 150 VUPs possible on the Model 660.

Upgrade Paths for VAX 6000 Systems

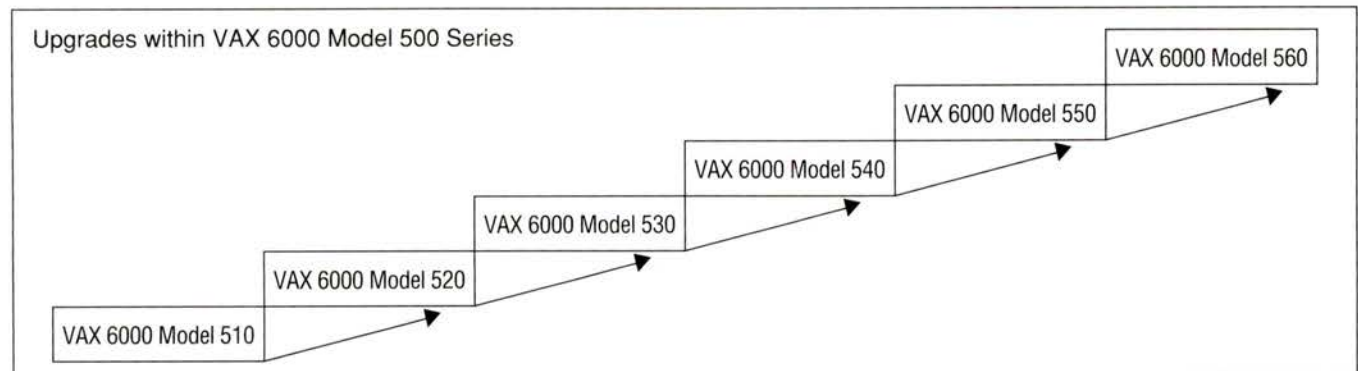
Upgrades to VAX 6000 Model 500 Series Systems

Upgrading to the VAX 6000 Model 500 Series from the Model 200, 300, or 400 provides performance increases of up to 72 VUPs. Like the Model 400 series, VAX 6000 Model 500 systems also offer vector processing options.



Upgrades within VAX 6000 Model 500 Series Systems

You can boost your processing power from 13 to 72 VUPs within the VAX 6000 Model 500 series platform by adding up to five additional processors.




Upgrades to VAX 6000 Model 600 Series Systems

When you upgrade your VAX 6000 Model 200, 300, 400, or 500 system to a VAX 6000 Model 600, you can take advantage of the most powerful single-stream performance in the VAX 6000 series. You can expand your processing power up to 150 VUPs, and your memory up to 1 Gbyte. You can also take advantage of write-back cache that bolsters the efficiency of the system bus to make the whole system run faster.

There are only two simple steps involved in most upgrades to the VAX 6000 Model 600 (and Model 500) system:

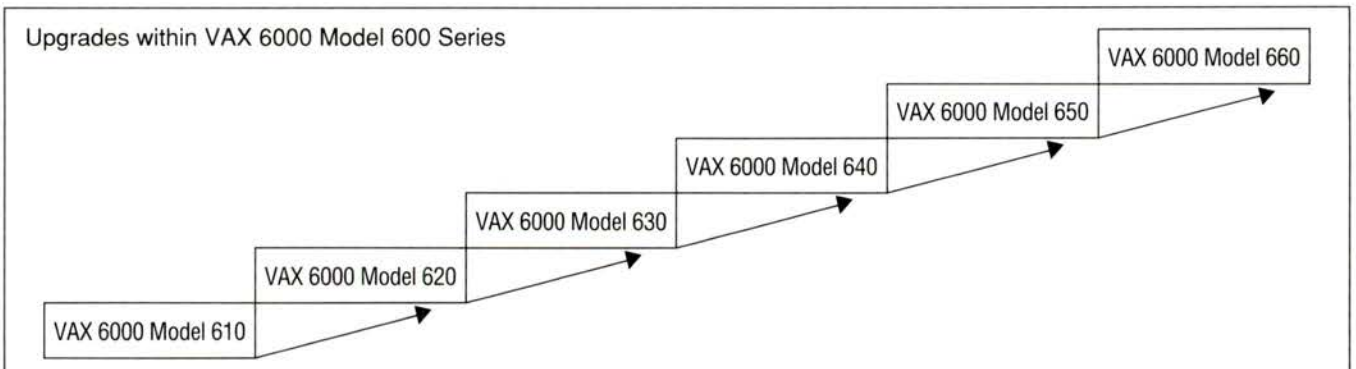
1. Select the CPU upgrade.
2. Replace all MS62A memory with MS65A memory. See page 8 of this guide for additional information on MS65A memory options.

If your system requires battery backup, more than 12 XMI slots, or more than four CPUs, additional ordering steps are necessary. Consult the latest edition of the *VAX Systems/DECsystems Systems and Options Catalog* for specific ordering instructions for these requirements.




Upgrades within VAX 6000 Model 600 Series Systems

VAX 6000 Model 600 series systems let you expand from 32 to 150 VUPs, reap the benefits of faster I/O throughput, and significantly expand memory through the use of MS65A memory modules.



Vector Processing Options



If you have Fortran or APL applications that involve floating point or intensive arithmetic operations, and you need to speed them up, you should consider expanding your VAX 6000 system to include the vector processor option. Available on VAX 6000 Models 400 and 500, the addition of a vector processor can increase the processing speed of your VAX 6000 system up to 20 times for applications that involve repeated operations on groups of floating point numbers.

Examples of such applications include linear algebra, weather simulation, finite element analysis, financial and econometric modeling, and computational fluid dynamics. Many leading software vendors' applications are available in versions that can be optimized by VAX vector processors.

Vector processing extends the VAX architecture by adding 63 additional instructions. This enables a vector processor to work in concert with a scalar processor to enhance the performance of VAX FORTRAN and APL applications. (A software compiler is also available to automatically vectorize these applications. See VAX FORTRAN Compiler information below.)

The VAX 6000 Model 400 and 500 integrated vector processing module fits into the slot in the XMI backplane – just like a memory board. You can also upgrade any VAX 6000 Model 200 or Model 300 system to a Model 400 or 500 system to take advantage of the vector processing option.

Supported Vector Configurations

System	Model 410	Model 420	Model 510	Model 520	Model 530	Model 540
Maximum Vectors	2	2	4	2	3	4

**A VAX 6000 with vectors
can process applications
at three to twenty times
the speed of a non-vector
system.**

VAX 6000 Systems with Vectors Performance Benchmarks

The following table summarizes the performance of VAX 6000 Models 410 and 510 for various industry-standard benchmarks in both MFLOPS and VUPs. These benchmarks measure the performance of vector-intensive algorithms common among numerically intensive applications. Due to their high vector content, they represent a best-case performance scenario, achieving peak performance levels of 555 VUPs.

Double-Precision VAX 6000 Series Benchmark Performance

	Model 410 MFLOPS	Model 410 VUPs*	Model 510 MFLOPS	Model 510 VUPs*
FFT	14.6	104	14.9	106
Convolution	30.1	177	31.0	182
Matrix Multiply	38.3	384	39.9	400
Livermore Loops Average	5.3	33.1	6.3	39.8
Livermore Loops Peak	36.5	152	38.4	160

*One VUP equals the performance of one VAX-11/780 system.

Double Precision LINPACK

	Model 410		Model 510		Model 520*		Model 530*		Model 540*	
	MFLOPS	VUPs	MFLOPS	VUPs	MFLOPS	VUPs	MFLOPS	VUPs	MFLOPS	VUPs
LINPACK										
1000x1000	27	205	28	213	50	380	85	494	73	555

*LINPACK numbers represent a single job decomposed to run across multiple CPUs.


VAX FORTRAN Compiler and the VAX FORTRAN High-Performance Option

The VAX FORTRAN compiler and the VAX FORTRAN High-Performance Option automatically vectorize existing VAX FORTRAN applications. The following tools are also available to help you fully utilize the vector processing capabilities of the VAX 6000 system.

- The VAX Vector Instruction Emulator Facility supports the development of vector applications on any VAX system running the VMS operating system. The application can then be ported and executed on the VAX 6000 Models 400 and 500.
- The Language Sensitive Editor (LSE) handles diagnostic statements generated by the VAX FORTRAN compiler, allowing for vector optimization.
- The Digital Extended Math Library enables programmers to easily access optimized versions of commonly used mathematical subroutines. This tool can be called from any VAX language to enhance the speed of applications written in that language.
- The VAX Performance Coverage Analyzer (PCA) identifies application program bottlenecks.

If your supercomputing application doesn't require the boost of vector processors, consider upgrading to a VAX 6000 Model 600 system.

Memory Expansion



Increased memory is important to maintain a truly balanced VAX 6000 system.

As you increase processor power, you often need to increase memory in order to maintain a truly balanced VAX 6000 system. In fact, configuring the correct amount of memory is crucial to system performance. Insufficient memory can result in excessive paging and swapping, long memory I/O command queues, and dissatisfied users. Memory is especially important for vector processors because they typically generate a high level of memory activity. (See page 6 of this guide for more information on the vector processing option.)

The MS65A memory modules provide enough memory for the most demanding applications – and users. These modules offer 32, 64 and 128 Mbytes of memory per board, up to a total of 1 Gbyte.† Advanced, high-density 4-Mbit DRAM technology gives the MS65A memory modules their larger memory capacities.

You can mix and match the three module sizes, thus allowing a more flexible approach to memory configuration than with the original 32-Mbyte MS62A* memory module. The various-sized modules allow you to custom-configure memory capacities to suit your applications, and to use the VAX 6000 system's expansion capabilities more efficiently. For example, you will be able to achieve the 1-Gbyte maximum memory using only eight expansion slots through the use of the 128-Mbyte MS65A memory module. On Model 500 and 600 systems, the MS65A memory modules also support write-back cache that bolsters the efficiency of the system bus to make your whole system still faster.

Memory Interleaving

Mixed-mode interleaving of the 32, 64, and 128-Mbyte modules increases the performance of the memory subsystem by distributing the memory traffic across multiple memory modules. It also provides additional configuration flexibility and ensures maximum system performance in higher performance VAX 6000 systems; for example, systems with more than two processors, or systems with an integrated vector processor.

*MS62A modules cannot be used on VAX 6000 Model 500 or 600 systems.

†512 MB available now, with an upcoming release of VMS, 896 MB for the Model 650, 768 MB for the Model 550, and 1 GB for the others will be available.

New densely-packed memory boards increase memory greatly while using little space on the system backplane.

The table below lists minimum interleave requirements for various VAX 6000 system models. (For more information on interleaving, refer to the System Expansion section of the latest edition of the *VAX Systems/DECsystems Systems and Options Catalog*.)

Recommended Minimum Interleave

Number of Processors	Model 300	Model 400	Model 500	Model 600
One	One	One	One	One
Two	One	One	One	One
Three	One	Two	One	One
Four	Two	Two	Two	Two
Five	Two	Four	Two	Two
Six	Two	Four	Two	Two
One with one vector	n/a	Two	Two	n/a
Two with one vector	n/a	Two	Two	n/a
Two with two vectors	n/a	Four	Four	n/a
Three with one to three vectors	n/a	Four	Four	n/a
Four with one to four vectors	n/a	Four	Four	n/a

Configuring Your System Memory

VAX 6000 systems can be configured with up to eight memory arrays. For optimum interleaved performance, arrays are configured as 1, 2, 4, 6, or 8 memory arrays.


All MS65A modules can be used on VAX 6000 Model 200/300/400 systems when appropriate system ROM upgrades are installed. They may be mixed and matched with MS62A memory modules on these systems. Memory interleaving is also supported with these mixed-mode memory configurations.

NOTE: VMS V5.4 (or later) is required with the MS65A memory options. Expansion beyond 512 Mbytes will be supported on a VMS release beyond V5.5.

DECram for Increased I/O Performance

Another way to take advantage of large memory capacity is to further speed up the system. The DECram memory-resident disk device driver is a facility that resides in memory, but behaves like a disk. It allows you to create a "pseudo disk unit" in VAX main memory. DECram performs at greater I/O rates than an actual disk. In fact, you can access DECram files at memory speeds without having to modify the applications or system software.

I/O Expansion



Now you can add up to 2 VME buses for use with third-party business-specific devices such as laboratory instruments, simulators, and medical equipment.

The VAX 6000 system's I/O design begins with the 100-Mbyte-per-second, XMI high-speed system bus.

The speed and efficiency of I/O are vital components of the overall performance of any computer system. No matter how powerful the CPU engine, or how great the system's storage and memory capacity, overall system performance suffers if data transfer is less than efficient.

The VAX 6000 system's I/O design complements the high-performance CPU and superior storage and memory capacity to eliminate the bottlenecks that can cause processing delays. It begins with the high-speed 100-Mbyte-per-second XMI system interconnect, which offers 14 slots to provide for every area of VAX 6000 system expansion.

When expanding the I/O capacity of your VAX 6000 system, your first decision involves whether to use the high-performance native XMI I/O controllers that allow you to take advantage of the bandwidth of the XMI system bus directly, or to use the original, and now optional, VAXBI technology. (In many cases, your choice of storage options will have an impact on the type of storage controller you need. Please see the the Storage Expansion section on page 12 of this guide, for information on storage expansion options.)

XMI I/O Controllers

Native XMI controllers include a VAXcluster controller, storage controllers for both SDI and DSSI storage options, an Ethernet network controller, a Fiber Distributed Data Interface (FDDI) network controller that features a data transmission rate of 100 Mb/s for the most demanding I/O requirements, and a VME bus to connect industry-standard devices. Most of these XMI controllers include one or more RISC CPU chips that help to speed I/O throughout the system.

The chart below offers at-a-glance performance information on the various XMI controllers, as well as configuration information and performance comparisons relative to the controllers' VAXBI counterparts where applicable.

As you can see, performance gains are significant with the native XMI controllers. Also note that there are no VAXBI counterparts to the KFMSA DSSI storage controller or the DEMFA FDDI network controller.

XMI Controllers At A Glance

	SDI Storage	DSSI Storage	VAXcluster	Ethernet	FDDI	VME
Option	KDM70	KFMSA	CIXCD	DEMNA	DEMFA	DWMVA
Ports	8 disk/tape	2 DSSI	1 CI	1 Ethernet	1 FDDI	6(6U) 3(9U)
Max. I/Os/sec.	1200	1600	2400	10,400	n/a	n/a
Max. Throughput	4.0 MB/s	5.5 MB/s	8.0 MB/s	10 Mb/s	100 Mb/s	n/a
Performance Comparison vs. VAXBI Controllers	3 x KDB50	n/a	3 x CIBCA	3.75 x DEBNI	n/a	n/a
Max. No./System	3	6	4	6	2	2
System	Any VAX 6000 System					

When you need additional I/O capabilities, add up to five VAXBI bus adapters for use with Digital-developed synchronous and asynchronous communications controllers, as well as with third-party devices.

VAXBI I/O Controllers

Originally the only I/O options available for the VAX 6000 systems, the VAXBI controllers remain a cost-effective alternative to the higher-performance XMI devices. You can add up to five VAXBI buses to the XMI system interconnect for use with Digital-developed controllers including storage controllers, VAXcluster controllers, Ethernet controllers, synchronous and asynchronous communications controllers, and various third party devices.

The following charts offer performance and configuration information on the various VAXBI controllers.


VAXBI Controllers At A Glance

	SDI Storage	VAXcluster	Ethernet
Option	KDB50	CIBCA	DEBNI
Ports	4 disk	1 CI	1 Ethernet
Max. I/Os/sec.	215	n/a	n/a
Max. Throughput	1.2 MB/s	2.6 MB/s	10 Mb/s
Max. No./VAXBI	2	1	2
Max. No./System	10	4	4
System	Any VAX 6000 System		

VAXBI Communications Controllers At A Glance

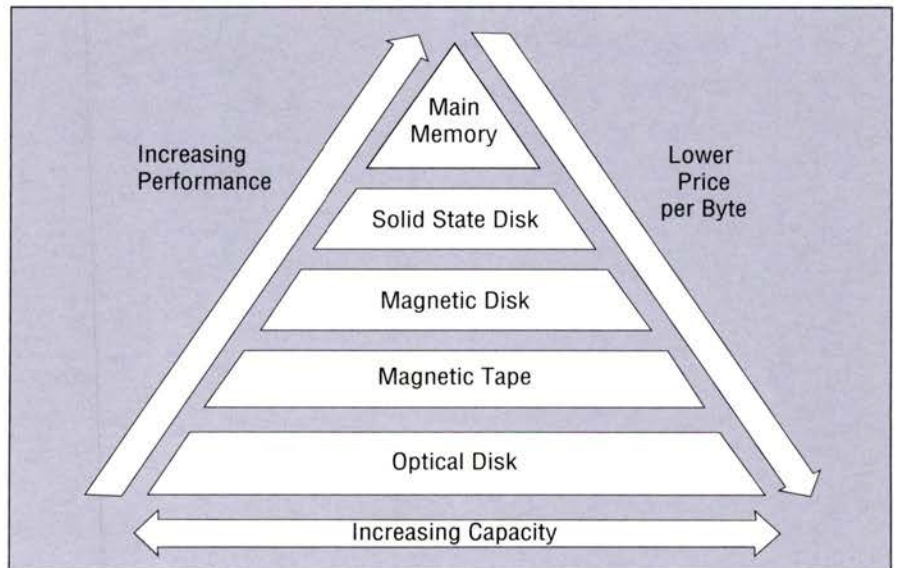
Option	DMB32	DSB32	DHB32	DRB32
Ports	8 asynch 1 synch 1 line printer port	2 synch	16 asynch	1 parallel interface
Max. No./VAXBI	4	4	4	2
Max. No./System	8	8	8	8
System	Any VAX 6000 System			

Storage Expansion



As your organization grows, so grows the need for high-capacity data storage. Business expansion, a corporate merger, or more stringent government regulations requiring you to store certain types of data longer can all significantly strain the capacity of your existing storage system.

Digital offers a wide range of storage products to allow you to expand the storage capacity of your VAX 6000 system to meet increasing demands, as well as products that protect your stored data, speed access to it, and allow you to manage it more efficiently.



With 5¹/₄-inch RF73 disk drives you can get 16 Gbytes of storage capacity right in the system cabinet.

Internal Storage Alternatives

You can get up to 16 Gbytes of storage in the system cabinet to form a complete, compact system where space requirements are critical. A number of alternative in-cabinet storage configurations are possible using RA70, RA92, RF72, and RF73 disk drives. For example, you could house the maximum 16 Gbytes storage inside the VAX 6000 cabinet using eight RF73 disk drives. Internal RA Family disk drives can be attached to either KDM70 or HSC storage controllers. (See the section on the Digital Disk Family on page 13 for more information on the RA family of disk drives.) These alternative configurations allow you the flexibility to add in-cabinet storage in increments, as needed.

Adding Storage Outside the System Cabinet

When your storage needs expand beyond the capacity of the VAX 6000 system cabinet, there are also many options available. Disk and tape storage arrays offer efficient use of floor space and economical, volume-pricing advantages. The TF800 tape family offers unattended backup. The ESE20 solid-state disk provides superior response time for time-critical applications, and optical storage disks offer optimal storage capacity for huge amounts of data that must be stored for long periods of time.

Digital gives you a choice between Standard Disk Interconnect/Standard Tape Interconnect (SDI/STI) and Digital Storage Systems Interconnect (DSSI) storage devices.

SDI/STI Storage and DSSI Storage: A Comparison

You probably already have SDI/STI (Standard Disk Interconnect/Standard Tape Interconnect) storage devices, either within or outside your system cabinet. These SDI/STI devices are implementations of the Digital Storage Architecture (DSA), a framework for Digital's storage products that provides standardized interfaces to govern the interactions of all Digital host systems and storage subsystems.

There is another implementation of DSA available called the DSSI (Digital Storage Systems Interconnect) product family. Both SDI/STI and DSSI storage products offer high-quality, reliable storage solutions. However, there are differences in configuration and the amount of storage capacity available with each implementation. The table below provides an easy decision matrix for selecting between DSSI or SDI/STI storage options. Consult your Digital sales representative for details.

SDI/STI and DSSI Decision Matrix

System Description	SDI/STI	DSSI
Standalone system or VAXcluster system expected to grow beyond 96 GB	X	
VAXcluster system requiring more than 3 nodes	X	
System requires ESE20 solid-state disk	X	
System requires 3480/3490 IBM tape compatibility	X	
Standalone system or DSSI VAXcluster system requiring less than 96 GB		X
Volume shadowing supported	X	X
Disk striping supported	X	X
System requires RV20 or RV64 optical disk	X	X
VAXsimPLUS supported	X	X

The Digital Disk Family

Disks available for expanding the storage capacity of your VAX 6000 system include the RF family of disks (DSSI) and the RA family of disks (SDI). The table below provides capacity and performance comparisons.

RF and RA Disks

	RF72	RF73	RA72	RA90	RA92
DSA Implementation	DSSI	DSSI	SDI	SDI	SDI
Capacity	1.0 GB	2.0 GB	1.0 GB	1.2 GB	1.5 GB
Form Factor	5.25 in	5.25 in	5.25 in	9 in	9 in
Average Seek Time	13.3 ms	13.0 ms	12.5 ms	18.5 ms	16.5 ms
Average Access Time	21.7 ms	21.3 ms	20.8 ms	26.8 ms	24.9 ms
I/Os/s	50	50	58	46	53
Peak Transfer Rate	2.0 MB/s	4.0 MB/s	2.0 MB/s	2.8 MB/s	2.8 MB/s
VAXsimPLUS Support	yes	yes	yes	yes	yes

If efficient use of floor space and economical volume pricing are concerns, Digital offers disks in multiple configurations called disk arrays. Examples of disk arrays, their capacity, and performance data appear in the table below.

Disk Arrays

	SA600	SA800	SF200	SF210
DSA Implementation	SDI	SDI	DSSI	DSSI
Disk Type Supported	RA90	RA92	RF72	RF73
Tape Type Supported	n/a	n/a	TF800(2)	TF800(2)
No. of Disks Supported per Array	8	8	24	24
Max. aggregate I/O Requests/s (@100ms)	368	424	1200	1200
Max. Array Capacity (Formatted)	9.6 GB	12.0 GB	24.0 GB	48.0 GB
VAXsimPLUS Support	yes	yes	yes	yes
Average Access Time	26.8 ms	24.9 ms	21.7 ms	21.3 ms
Average Seek Time	18.5 ms	16.5 ms	13.3 ms	13.0 ms
Peak Transfer Rate	2.8 MB/s	2.8 MB/s	2.0 MB/s	2.8 MB/s

Special Disk Features

The following is a listing of special features available with Digital's disk storage options. These special features offer additional data protection and improved response time for business- and time-critical applications.

Digital's disk devices offer special features for additional data protection and improved response time.

- Volume shadowing copies vital data onto more than one disk to provide greater data availability. In the event that one disk goes down, your data is protected on the other disk.
- Disk striping spreads data over multiple disks for increased I/O performance.
- HI SPEED mode is available on DSSI disks. It provides faster seek times by using only half of the disk surface area, thereby leading to better response time.

The Digital Tape Family

Digital offers a wide spectrum of tape products for backing up or transferring your data. These range from the TK70, which fits inside the CPU cabinet, to the highest performance TA91, which offers data interchange with IBM 3480 and 3490 devices. Unattended backup on some models allows your operation to run more smoothly.

Digital Tape Drives

	TK70	TF85	TF857	TA81	TA79	TA91
Capacity	296 MB	2.6 GB	18.2 GB	145 MB	145 MB	1.2-4.8 GB
Transfer Rate	90 MB/s	800 MB/s	800 MB/s	468 MB/s	781 MB/s	2700 MB/s
Data Interchange*	TF837 TZ30 TK50	TF857 TF837 TK70 TK50 TZ30	TF837 TF85 TK70 TK50 TZ30	TA79	TA81	IBM 3480 IBM 3490 TA90

*Most devices are read compatible only. TF837 and TK70 are read/write compatible. TA79 and TA81 are read/write compatible

Special Tape Features

With Digital's tapes, you can get many gigabytes of unattended back-up, as well as data interchange with IBM systems.

- The TF800 series tapes bring affordable unattended tape backup capabilities to VAX 6000 systems. The TF857 transfers data at 800 Kbytes per second and can back up 18.2 Gbytes of data in one shift.
- Data interchange capabilities enable the movement of data from one VAX system to another, or movement to or from an IBM 3480 or 3490 device. In addition, tapes generated by a TK50 or TK70 can be read by the newer TF85, TF837 and TF857 devices.


ESE20 Solid-state Disk

The ESE20 SDI solid-state disk has been designed to increase the performance of the VAX 6000 system supporting I/O intensive applications such as transaction processing, large population timesharing, and I/O-intensive scientific analysis. The ESE20 provides dramatic improvements in system throughput, response time, batch run completion time, and maximum user count via the KDM70 storage controller, or via a Hierarchical Storage Controller (HSC) in a VAXcluster configuration. For example, when utilized with the KDM70 or HSC60/90 controller, the ESE20 can service more than 1200 I/O requests per second.

Optical Storage

Certain applications generate large amounts of data that must be reliably stored for long periods of time. VAX 6000 systems support the RV64, which is a VAXBI-based SDI optical disk subsystem containing the mechanical loaders, 1 to 4 RV20 drive units, and up to 64 optical disks for a total of 128 GB of storage. The RV64 is an excellent device for automating the long-term storage of such data as customer records and medical images.

VAXcluster System Options



Now you can mix DSSI and CI VAXcluster systems.

Is your business increasingly dependent on your computing systems and business applications?

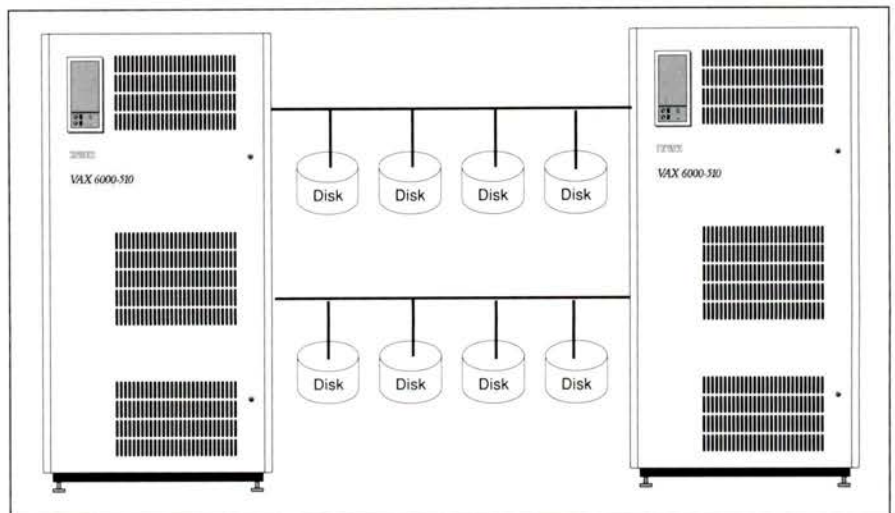
- Are you looking for ways to provide increased system uptime and application availability without having to rewrite your existing applications?
- Are you looking for ways to improve data access for your users to increase their productivity?
- Are cost concerns driving an increased effort to share your computer resources in the most economical way possible?
- Are you looking for ways to extend the life of your VAX 6000 system?

If you answered yes to any of these questions, moving to a VAXcluster configuration may be the solution.

A VAXcluster system is a highly integrated configuration of two or more VAX systems operating together like a single, powerful system. Users may have access to one, some, or all of the systems, depending on the access privileges assigned to them by the system manager. And a system in the VAXcluster can fail without impacting the user's access to other computing resources and stored data, reducing downtime and increasing user productivity.

A VAXcluster configuration protects your investment by enabling you to add new VAX systems as your needs grow, while you continue to use the first VAX system you purchased. And it offers a highly cost-effective growth path, providing centralized system management capability and the ability to share expensive computer resources. These features make the VAXcluster system the ideal choice for your data center, whether you're supporting centralized, business-critical applications or managing a distributed application environment, with a growing number of desktop systems and workstations.

Once you have made the decision to expand beyond your single VAX 6000 system to a VAXcluster configuration, your next decision involves the choice among the DSSI-, CI-, and FDDI-based VAXcluster options.



DSSI VAXcluster System

DSSI VAXclusters are economical, entry-level, high-performance configurations of two or three systems.

CI VAXclusters offer more storage, capacity, and higher expansion.

FDDI VAXclusters make remote system management simple. They also provide a high level of disaster tolerance.

DSSI VAXcluster Systems

Three-system DSSI configurations are the latest type of VAXcluster available to the VAX 6000 user. With the DSSI bus connecting systems and storage cabinets, DSSI VAXclusters offer a cost-effective, entry-level option for the user whose business and computing needs require VAXclustering capability.

The DSSI-based VAXcluster provides a maximum of 96 GBytes of shareable disk storage for up to three VAX 6000-class systems. (See the Storage Expansion section on page 12 for more information on DSSI.) Should your availability, online storage, and performance needs grow in the future, the DSSI VAXcluster system can be integrated with a high-end CI-based VAXcluster system.

CI VAXcluster Systems

CI-based configurations offer extended growth capabilities for VAXcluster-wide resources, supporting up to 32 VAX 6000 and higher class-systems or more than 8 Terabytes of shareable disk storage.

And CI VAXcluster configurations offer the highest levels of availability. New VAX systems and additional mass storage devices, including IBM-compatible tape cartridge subsystems and the high performance ESE20 solid-state disks, can be added to the VAXcluster while the VAXcluster remains online!

The table below provides an at-a-glance comparison of the attributes of the DSSI and CI VAXcluster configurations.

VAXcluster Attribute Positioning Table


	DSSI VAXcluster	CI VAXcluster
VAXcluster Interconnect	DSSI	CI
Interconnect Bandwidth	4 MB/s	10 MB/s
Maximum number of VAX/Systems	3	32
Max. Shared Storage (unshadowed)	96 GBytes	8 Tbytes
Disk/Tape Family	SF/RF/TF	SA/RA/TA
ESE20 Solid-state Disk Support	No	Yes
IBM Tape Cartridge Compatibility	No	Yes

Expanding Beyond Your VAXcluster...

By combining VAXcluster technology with an FDDI network and Digital's new Multi-Datcenter Facility, two or more VAXcluster systems which are in data center facilities separated by up to 40 kilometers can be joined into a single VAXcluster system. The resulting configuration makes it possible to manage multiple sites with one operations staff for significant potential cost savings.

For even higher availability, the addition of Volume Shadowing technology to your Multi-Datcenter VAXcluster system provides you with site redundancy – and a disaster tolerant system. A VAXcluster Multi-Datcenter system can continue to operate even if an entire data center is lost due to a disaster. With disaster planning, your Multi-Datcenter VAXcluster will enable your business-critical applications to be disaster tolerant.

Load Device Options



Digital load devices make it quick and efficient to distribute software to your growing user base.

As the number of users supported by your VAX 6000 system grows, and your system and application software library expands to meet your organization's needs, your need to distribute software quickly and efficiently to your users grows too.

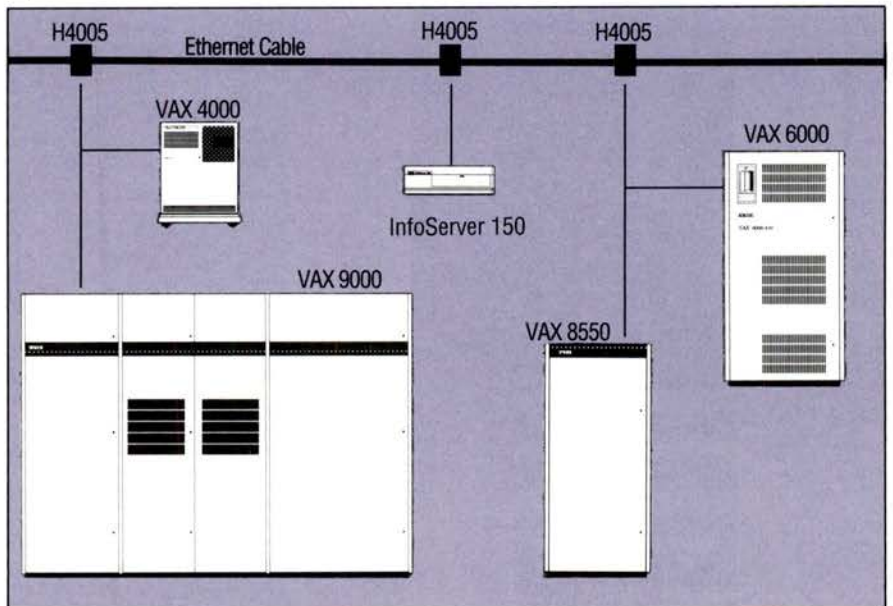
The VAXBI-based TK50 load device served this purpose on the original VAX 6000 systems, and was later replaced by the higher-performance, higher-capacity TK70 load device. However, there are now more advanced software load device options available to VAX 6000 system owners: the InfoServer 150 and TF85/TF857 high-capacity tape drives.

InfoServer 150

The InfoServer 150 is a network storage server. This dedicated hardware/software product attaches directly to Ethernet, and provides shared read-only access to locally attached SCSI storage devices.

The InfoServer 150 supports VMS operating system clients as well as MS-DOS®-based clients. It permits online distribution of VMS and nearly 200 VMS layered software products and accompanying software documentation to multiple VMS systems across the network. MS-DOS users can access the InfoServer 150 using either PATHWORKS for DOS (InfoServer) or InfoServer clients for DOS client software. The first serves Digital's PC network users and the latter serves Novell NetWare® operating environments. You simply purchase the appropriate software kits for whichever operating environment you require. VMS client software is incorporated in VMS V5.4 and above.

InfoServer 150 on the network



TF85/TF857 Tape Drives

The TF85/TF857 tape drives are designed for high performance, high data integrity, software distribution, and unattended backup. The TF85 is a 5.25-inch tape cartridge, while the TF857 is a tape subsystem with room for seven TF85 tape cartridges. The TF85 and TF857 also accept and read cartridges from TZ30, TK50, TK70, TF70, and TF837 drives. Therefore, TK50 cartridges can be used to distribute software to systems with TF85 and TF857 tape drives. If you want to order a TF85 for an existing system, you need the tabletop version.

The VAX 6000 system supports up to four TF857 tape drives, each providing unattended backup capacity of up to 18 Gbytes in an 8-hour shift, for a total of 72 Gbytes of unattended backup.

TK70 Load Device

You may want to pursue an upgrade path within the TK family. If so, you can upgrade your TK50 to the higher-performance TK70 load device.

The 296-Mbyte TK70 provides three times the capacity of the 96-Mbyte TK50 to increase the efficiency of disk backup, software distribution, archiving, and data interchange on your system.

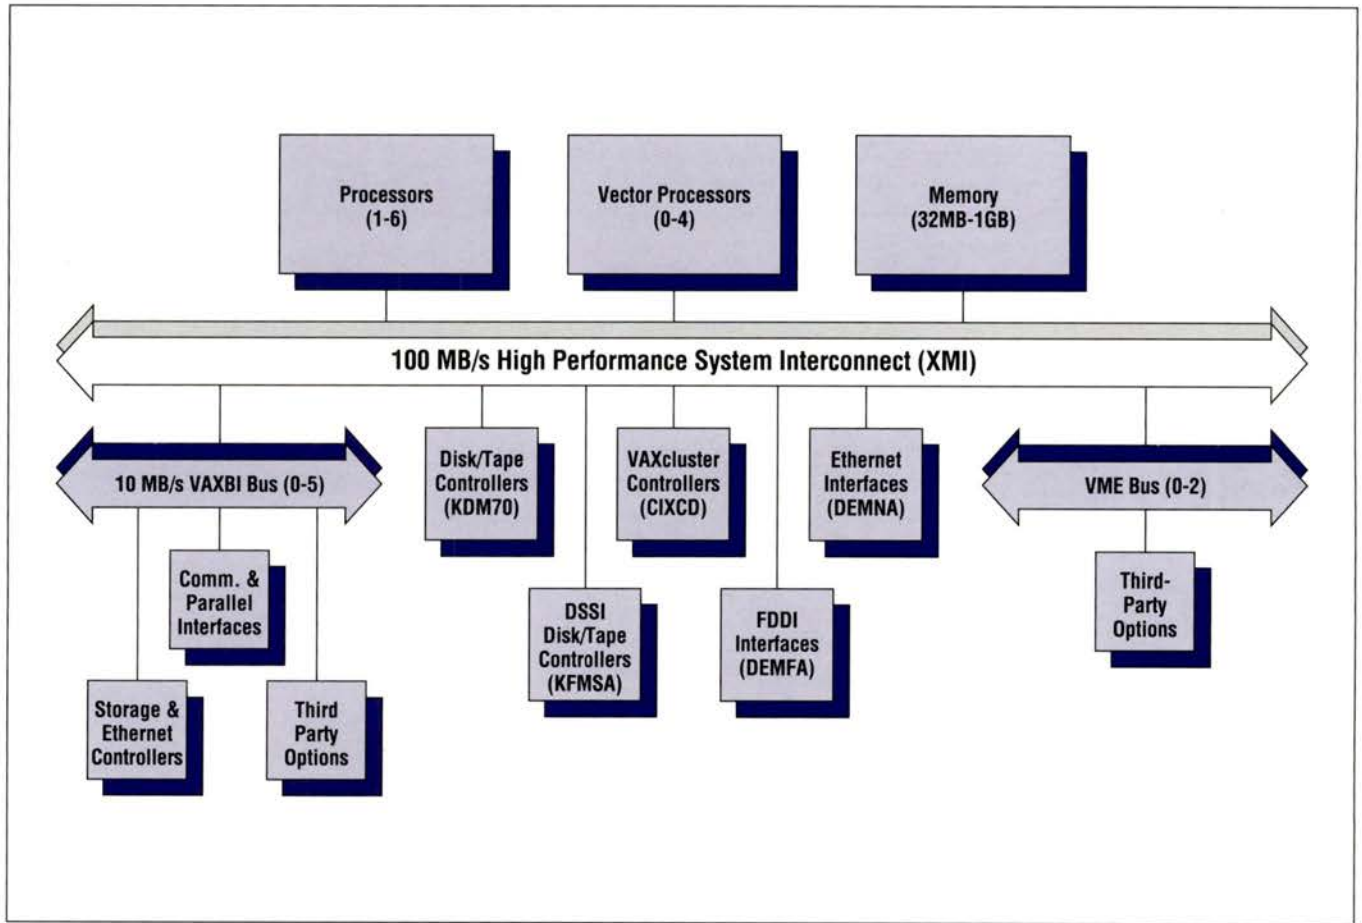
The high-performance TK70 reduces backup time by almost 50 percent over the original TK50 load device. It also transfers formatted data at 90 KB/s, twice as fast as the TK50 transfer rate of 45 KB/s.

The TK70 can read cartridges written on a TK50 load device. Therefore, the TK50 cartridges can be used to distribute software to both TK50 and TK70 load devices.

Summary

The modular design of the VAX 6000 system, and its membership in the VAX family of compatible computer systems, offer numerous options for expanding your initial system purchase to meet your growing needs. This guide is meant to provide you with helpful information about the many expansion options available to meet the increasing demands on your system. However, it is not all-inclusive.

There are several sources of additional information on managing the growth of your VAX 6000 system and help on configuring the best VAX 6000 system for your needs. Consult the latest edition of the *VAX Systems/DECsystems Systems and Options Catalog*, or contact your local Digital sales office. Ask your sales representative for a copy of *Successful Growth Planning* (EE-EM618-78), a guide to planning the growth of your system, and his or her personal assistance in planning the expansion of your VAX 6000 system.




Digital believes the information in this publication is accurate as of its publication date; such information is subject to change without notice.

The following are trademarks of Digital Equipment Corporation: CI, the DIGITAL logo, DRB32, DSSI, HSC, RA, RV20, RV64, SDI, SPM, STI, TA, TK, VAX, VAXBI, VAXcluster, VAX FORTRAN, VAXsimPLUS, VPA, VMS, and XMI.

IBM is a registered trademark of International Business Machines Corporation. NetWare is a registered trademark of Novell, Inc.

digital



VAX 6000 Datacenter Systems Performance Report

Included in this report:

- TPC Benchmark™ A
- ALL-IN-1™ /OEU
- Relative Performance
- Commercial Batch Benchmark
- SPEC® Release 1.2

VAX 6000 Product Description

VAX 6000 systems provide datacenter solutions across a wide range of applications. By this, we mean that VAX 6000 systems provide the necessary capabilities to support your most critical data or applications -- your key information assets. Datacenter managers do not have to compromise performance, price/performance, functionality, quality, or openness with VAX 6000 systems. New VAX 6000 Model 600 systems now deliver leadership performance and price/performance -- even better than RISC systems from IBM or HP.

VAX 6000 Performance

In order to evaluate the performance of the VAX 6000 systems, both commercial and technical benchmarks/workloads were used. VAX 6000 systems provide excellent performance under several commercial applications. The new VAX 6000 Model 610 takes a leadership position with TPC Benchmark A, an industry-standard transaction-processing workload, used to simulate an on-line transaction processing banking network environment. The ALL-IN-1 Office End-User (OEU) workload, developed by Digital, consists of a number of general office tasks, and is used to determine the maximum number of users supported by VAX systems. This workload clearly demonstrates that VAX 6000 datacenter systems can support large, demanding office environments.

The outstanding processing power available on VAX 6000 systems makes them equally attractive in the technical environment. SPEC Release 1.2 benchmark suite is used to evaluate processor power and compiler efficiency. The programs in the suite represent both electronic engineering and scientific applications. These benchmarks are single-stream and compute-intensive. In addition, a Digital-developed suite of compute-intensive benchmarks is used to compare relative performance among different VAX systems.

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. The report simply presents the data, based on specified benchmarks and workloads.

For more information on VAX 6000 systems, please contact your local Digital sales representative.

TPC Benchmark™ A

Transaction processing is found in many industries such as finance, manufacturing, telecommunications, retailing and healthcare. On-line Transaction Processing (OLTP) environments emphasize update-intensive database services characterized by multiple terminal sessions, significant disk input/output, a specified execution time and transaction integrity.

Benchmark Description:

In order to test these variables, the TPC Benchmark A, which represents a banking network, was developed by the Transaction Processing Performance Council (TPC) as a vendor neutral, standard specification for measuring the number of transactions a system is capable of executing in a given period of time. The transaction, which is performed by a bank teller at a branch, represents the work done when a customer makes a deposit or withdrawal against an account. The benchmark requires that 90% of all transactions complete in less than two seconds. These results serve as useful criteria for businesses who depend upon providing timely service to their customers in order to gain a competitive edge.

Performance Metrics:

- tpsA-Local™ is defined as the number of transactions executed per second under the TPC Benchmark A.
- \$ per tpsA-Local is the associated price-per-tps metric, which is based on a five-year cost of ownership.



Results:

Using VAX Rdb V4.1, the VAX 6000 Model 610 clearly delivers leadership performance and price/performance with independently-audited benchmark results of 83.6 tpsA-Local at \$12,922 per tpsA-Local. As shown below, the VAX 6000 Model 610 system provides leadership capabilities compared to models with the highest-announced TPC-A results.

	tpsA-Local	\$ per tpsA-Local
• VAX 6000 Model 610	83.6	12,922
• HP3000-957LX	49.4	13,720
• HP9000-847S	42.8	14,560
• IBM RS/6000-950	38.0	21,823
• IBM AS/400-D70	42.5	21,720

VAX SYSTEMS AND SERVERS

The charts below demonstrate the high performance and low cost of the VAX 6000 Model 610 system compared to the highest-reported tpsA-Local results available from other vendors. The source of this information is the September 30, 1991 issue of the *Complete TPC Results*, published by the Transaction Processing Performance Council. Starting with the Model 610, up to five additional processors can be added to achieve six-way, symmetric multiprocessing, providing even higher levels of performance. VAX 6000 systems can then be grouped together in VAXcluster configurations for even higher performance. These are the performance levels required for today's demanding datacenter applications.



VAX SYSTEMS AND SERVERS

ALL-IN-1™ Multiuser, Multiprocessor Performance


Benchmark Description:

Banks, branch offices, retail businesses, manufacturing sites, health care facilities, software development houses, and other businesses, all include office end users. These office users typically perform a variety of tasks such as word processing, electronic mail and time management functions.

The ALL-IN-1 Office End-User (OEU) workload, which is based on Digital's ALL-IN-1 office software product, is comprised of the same mix of general office tasks, and is used to determine the maximum number of users and subscribers supported by various VAX systems. Using the results of the ALL-IN-1/OEU workload, businesses can determine the appropriate VAX 6000 system required to support their office users.

Results:

The VAX 6000 Model 610 uniprocessor supports almost 700 subscribers under the ALL-IN-1/OEU workload. The dual processor, Model 620, supports even more subscribers (936) using SMP. Supporting over 900 subscribers in this rigorous environment demonstrates the powerful performance and commercial strength of VAX 6000 systems. The number of subscribers indicates how many people can be supported by a system while 50% are actively using their accounts.



VAX SYSTEMS AND SERVERS

Commercial Batch Benchmarks



Workload Description:

The commercial batch suite consists of four COBOL programs developed by Digital and is based on corporate financial applications from Dun and Bradstreet, Ross Systems, and Cyborg Systems. General ledger, accounts receivable/payable, and payroll are typical financial functions required by businesses. In the Commercial Batch Benchmarks, a substantial portion of the computing required consists of large batch jobs to support these functions.

- Dun and Bradstreet Accounts Payable - processes 545 Accounts Payable Vouchers against 240,000 Master File Records, 77,000 Vendor Master File Records, and 73,000 Current Invoice File Records.
- Dun and Bradstreet General Ledger - simulates each transaction's read, edit, and posting of 24,000 records to a 26,000 record General Ledger Master File.
- Ross General Ledger - simulates the reporting functions during the end-of-month update to a general ledger master file containing 37,000 records. It does not simulate the processing of detail journals or posting and update operations.
- Cyborg Payroll - simulates the calculation of a weekly payroll consisting of 60,000 employees with a timecard-to-employee ratio of 2:1. Overall, the Cyborg payroll simulates a 120,000 transaction workload. It does not include printing the checks.
- Dun and Bradstreet Millenium General Ledger - simulates quarter-ending financial report of a Fortune 10 company with two million journal entries. Tasks include sort, print, and copy.

Results:

The results show that the new VAX 6000 Model 610 datacenter system cuts the time required to complete these demanding financial tasks. For example, the Millenium batch job, running on the Model 610, now completes overnight -- so other activities can take place promptly the next day. Note: The data represented in this analysis is meant solely to position Digital platforms. The data is not meant to position the applications, especially since the databases used for each application are different.



VAX SYSTEMS AND SERVERS


Relative Performance

Benchmark Description:

Many VAX users like the simplicity of a single performance number (times the VAX 11/780), which can be used to estimate relative performance among different VAX systems. A suite of benchmarks is used to calculate the relative CPU performance of VAX systems. Relative positioning is established by comparing the performance of the system under test with that of the VAX 11/780. Relative performance is the mean of those results and can be used as a relative measure in predicting CPU performance.

Results:

The relative performance numbers shown below represent the performance of VAX 6000 systems in a mixed commercial and technical application environment. Note that Models 620 through 660 performance results show the compute power available with symmetric multiprocessing (SMP) -- Digital's unique way to grow system performance in minutes by simply installing additional processors in the system cabinet. Up to six processors are supported in a VAX 6000 system, capable of delivering up to 150 times the processing power of the VAX 11/780.



VAX SYSTEMS AND SERVERS

SPEC[®] Release 1.2

Benchmark Description:

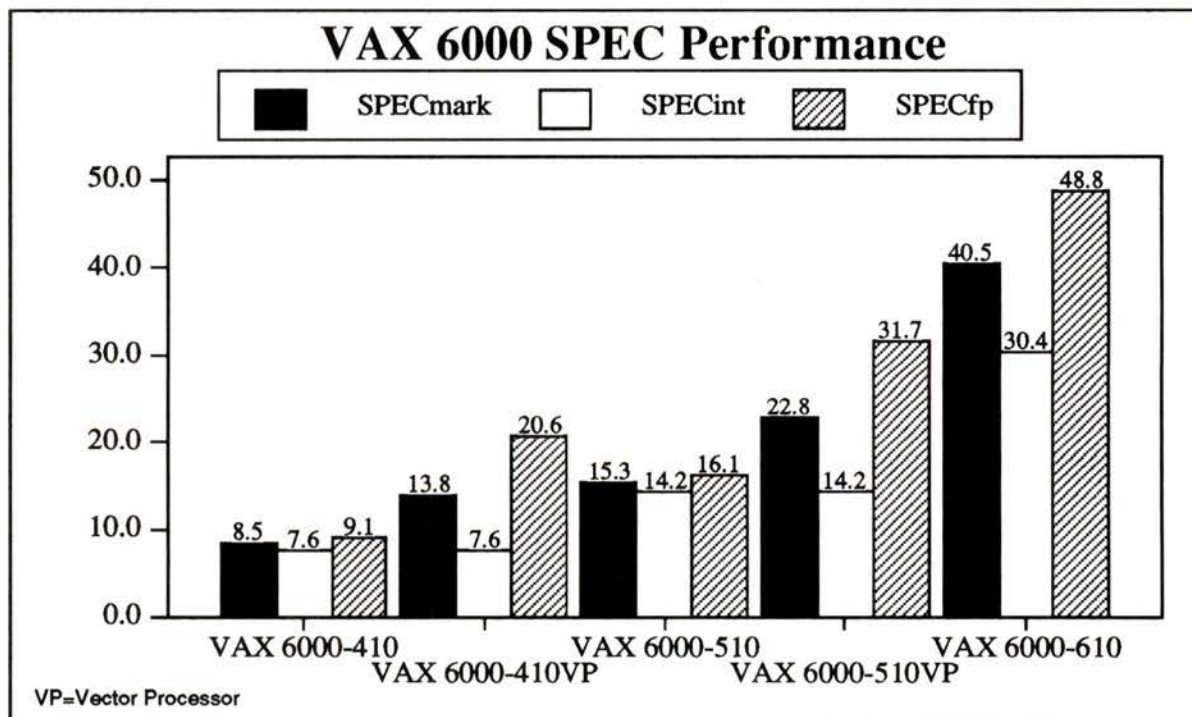
The System Performance Evaluation Cooperative (SPEC) consists of 23 vendors who work together on developing performance benchmarks that can be run on multiple vendors' platforms. The SPEC benchmark suite (Release 1.2) is used to evaluate the performance of systems in single-stream, engineering/scientific compute-intensive application environments. The suite consists of ten benchmarks -- six written in FORTRAN, using double-precision floating point data format, and four non-floating point benchmarks written in C, referred to as the integer benchmarks.

Performance Metric:

SPECmark is the geometric mean of the performance of the ten SPEC benchmarks relative to the VAX 11/780. The geometric mean, using six FORTRAN benchmarks, is defined as SPECfp. SPECfp represents the performance of the machine where FORTRAN double-precision floating point operations are a "non-trivial portion of the computation". Similarly, the geometric mean of the four C benchmarks is defined as SPECint. SPECint is intended to represent the non-floating point intensive portion of the machine's performance.

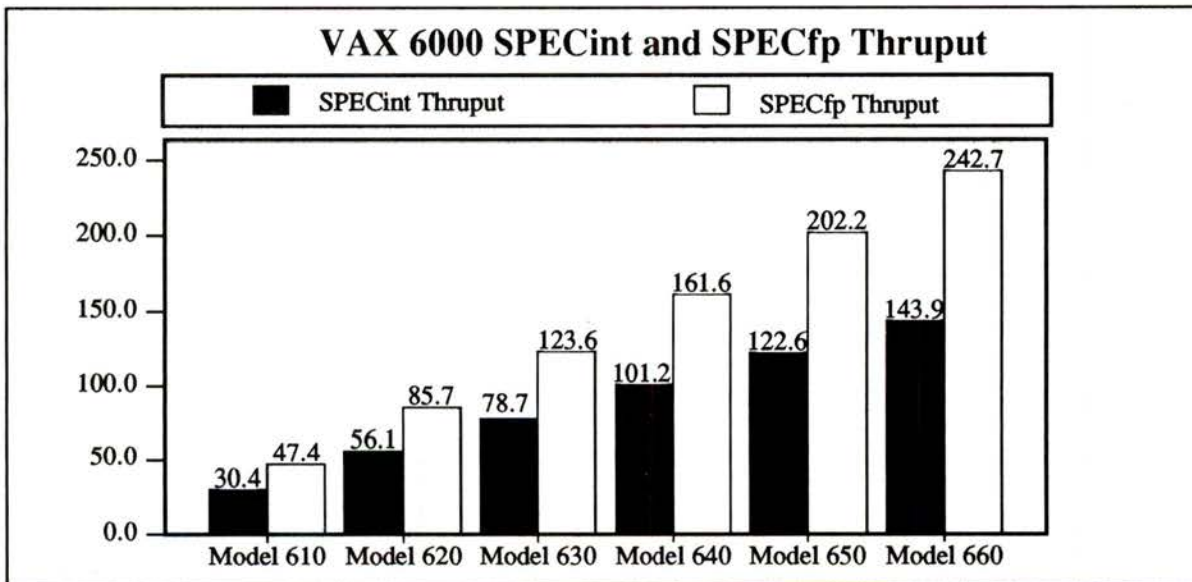

Results:

The following SPEC results have been achieved. Refer to the *SPEC Newsletter* for continuous reporting of SPEC data. Note: The performance numbers for the VAX 6000 Models were measured using advanced VAX compiler technologies.



VAX SYSTEMS AND SERVERS

SPECthruput is a metric used by SPEC to measure the batch throughput on a multiprocessor system. Two streams of each benchmark are run on each processor of a multiprocessor system and the elapsed time is normalized to one stream and to a reference system (VAX 11/780). A geometric mean of these numbers is the SPECthruput. SPECthruput multiplied by the number of processors provides SPEC Aggregate Thruput.



Digital believes the information in this publication is accurate as of its publication date; such information is subject to change without notice.

The following are trademarks of Digital Equipment Corporation: ALL-IN-1, CI, DECnet, DECsystem, the DIGITAL logo, DSSI, Micro-VAX, PATHWORKS, RA, Rdb/VMS, RX, SDI, ULTRIX, VAX, VAXBI, VAXserver, VAXstation, VAXcluster, VAXft, VAX FORTRAN, VAX Performance Advisor, VAX Volume Shadowing, VMS, VT, and XMI.

HP is a registered trademark of Hewlett-Packard Company. IBM is a registered trademark of International Business Machines Corporation. Millennium is a trademark of Dun and Bradstreet Software. SPEC, SPECratio, and SPECmark are registered trademarks of the Standard Performance Evaluation Corporation. TPC Benchmark, TPC-A, TPC-B, and tpsA-Local are trademarks of the Transaction Processing Performance Council.

Printed in U.S.A. EC-N0149-46/91 11 32 15.0 MKO Copyright 1991 Digital Equipment Corporation All Rights Reserved.