



Performance and Price/Performance for New EV7 Processors in AlphaServers

Session ID: 4086



Richard Smith, Product Manager
AlphaServer Engineering
Hewlett-Packard
August 17, 2004

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



The EV7/Marvel AlphaServer Systems



Click here to see introductory video →



2P AlphaServer ES47 tower

2-4P AlphaServer ES47 rack


2-8P AlphaServer ES80 system

2-64P AlphaServer GS1280 system

family-wide capabilities

- integrated server management
- scalable I/O
- hard partitioning & Galaxy soft partitions for OVMS
- enterprise RAS
 - RAID memory
 - Redundant power supplies and fans
 - Hot-plug disks, power supplies, fans, and platform management modules
 - Dual power feeds
 - ECC protected cache, memory, and inter-processor links, and I/O links
 - Fully isolated hard partitions
 - Pre-failure warning on hard drives, processors, and memory
- AGP graphics
- capacity on demand
- hp Tru64 UNIX®, OpenVMS and Linux® support

Strong Sales of EV7 AlphaServer Systems ...



Marvel sales would be more, but ...



“The HP AlphaServer GS1280 system will save us a lot of money because the capacity far exceeds our expectations. We originally thought we would have to double the size or add another server. Now, because of the performance we are seeing, we are quite confident that we will be able to get through this year and into next before we have to add more database capacity.”

Lee Oettinger
Manager of Technical Services
Aurora Health Care

A brief look back ...

Original Design Goals for Alpha




- 64 bit addressing
- Instruction Set Architecture targeted for Unix and VMS
- RISC design
- Enable high frequency clocks
- Good compiler target
- IEEE floating point

Alpha Chip History




	Chip	Technology, microns	Die Size, mm ²	Clock Frequency, MHz	Instructions per Cycle	SMP
1991	EV4	0.75 → 0.50	233 → 175	200	0.5	6
1995	EV5	0.50 → 0.35	300 → 210	350 → 600	0.8	8
1998	EV6	0.35 → 0.18	314 → 125	575 → 1250	1.5	32
2002	EV7	0.18	400	1150 → 1300	1.5	128

EV4 – First Alpha 1991




- 0.75 μ m 3LM CMOS, 3.3V
- 200MHz @30W
- 233 mm², 1.7M, 431-PGA
- Dual in-order issue
- Fully pipelined
- 8kB I-Cache, 8kB D-Cache
- 32 64b I-Regs, 32 64b FP-Regs
- 1-bit branch prediction
- Shared L2, system interface

EV5 - 1995



- 0.50 μ m 4LM CMOS, 3.3V
- 350MHz @60W
- 298 mm², 9.3M, 499-PGA
- Quad in-order issue
- FP latencies reduced 2 cycles
- 8KB I-Cache, 8KB D-Cache
- 96KB unified on-chip L2 cache
- 2-bit branch prediction
- Non-blocking cache scheme

EV6 - 1998



- 0.35 μ m 6LM CMOS, 2.2V
- 575MHz @90W
- 314 mm², 15.2M, 587-PGA
- Quad-issue
- Out-of-order, 80 in flight
- 64KB I-cache, 64KB D-cache
- 2 load or 2 store per cycle
- Extensive branch and jump prediction
- High bandwidth interfaces


A Look at Today's Alpha EV7 ...

Alpha EV7 Goals



- Increase
 - Single processor performance, operating frequency, and memory system
 - SMP scaling
 - System performance density (computes/ft³, I/O/ft³)
 - Reliability and availability
- Decrease
 - System complexity
 - System cost

Estimated time for Transaction Processing




EV7 – The System is the Silicon....



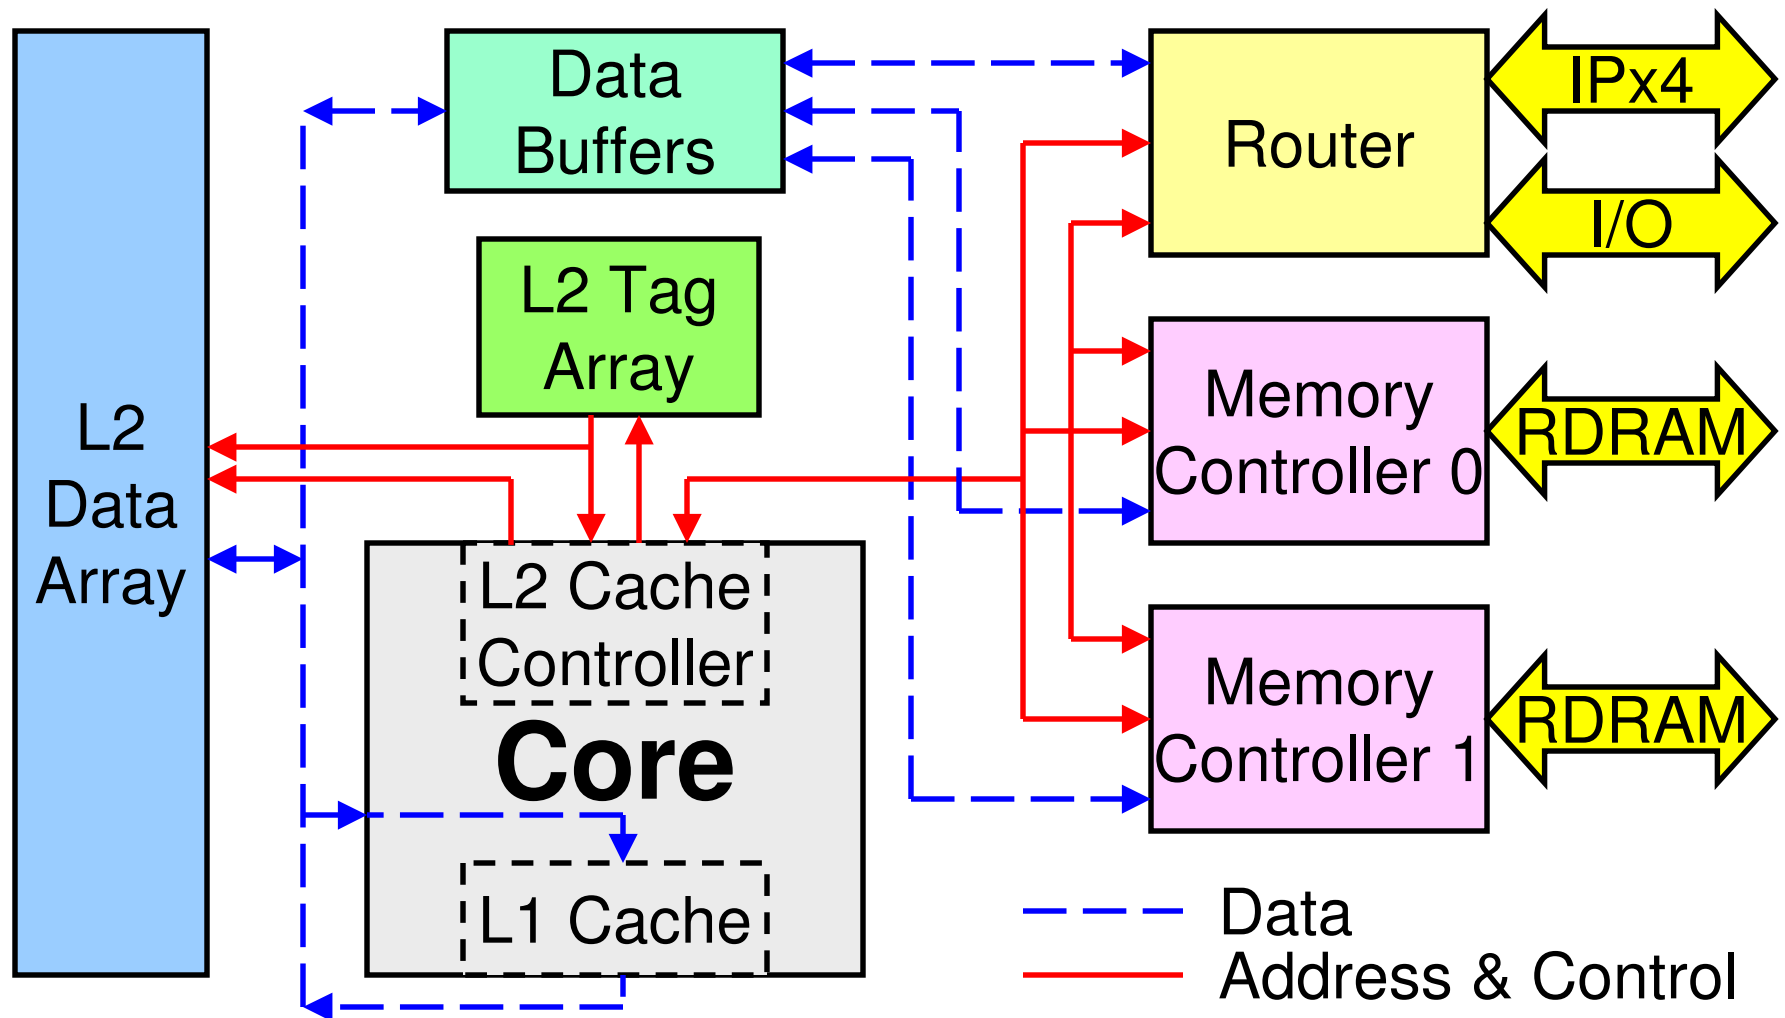
***SMP CPU interconnect was external logic..
Now it's on the chip !***




- **EV68 core with enhancements**
 - Big pages
 - 16 cache block requests at once
 - 32 victim block buffers
- **Integrated L2 cache**
 - 1.75 MB, 7-way set associative
 - 9.6ns load to use latency
 - 19.2 GB/s bandwidth
- **Integrated dual memory controllers**
 - 75ns load to use latency
 - 12.8 GB/s bandwidth
 - 64 entry directory based cache coherence engine
 - Optional RAID memory (4+1 parity)
- **Integrated 5 port switch**
 - 4 point-to-point processor links
 - 25.6 GB/s aggregate bandwidth
 - 18ns processor-to-processor latency
 - Out-of-order network with adaptive routing
 - Asynchronous clocking between processors
 - I/O interface, 2 uni-directional links, 3.2 GB/s aggregate bandwidth
- **Testability & debug features**
 - Data dump mode; L2 cache, L2 cache & memory controller address arrays
 - On-chip logic analyzer



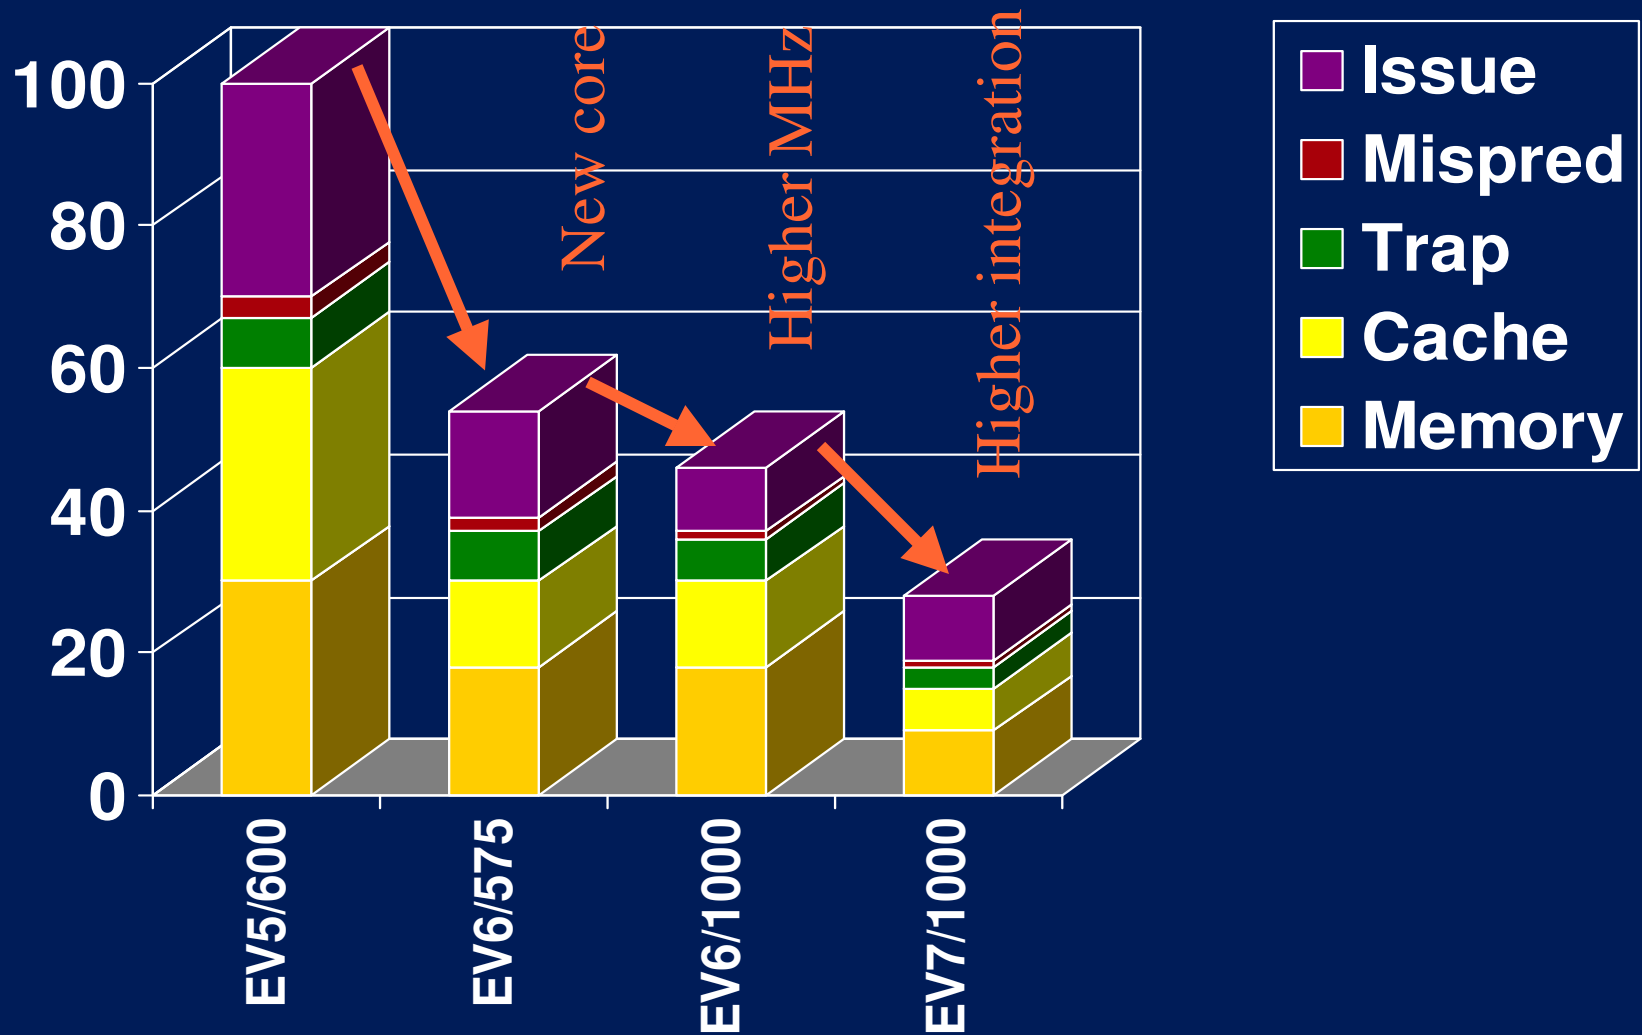
EV7 Block Diagram



EV7 - 2002



Estimated time for Transaction Processing




Design Goals Achieved ...

- Increase single processor performance
- Increase SMP scaling
- Increase system performance density
- Increase reliability and availability
- Decrease system complexity
- Decrease system cost


Better Single CPU Performance & SMP Scaling




SPECfp_rate2000




SPECfp_rate2000



Memory Bandwidth



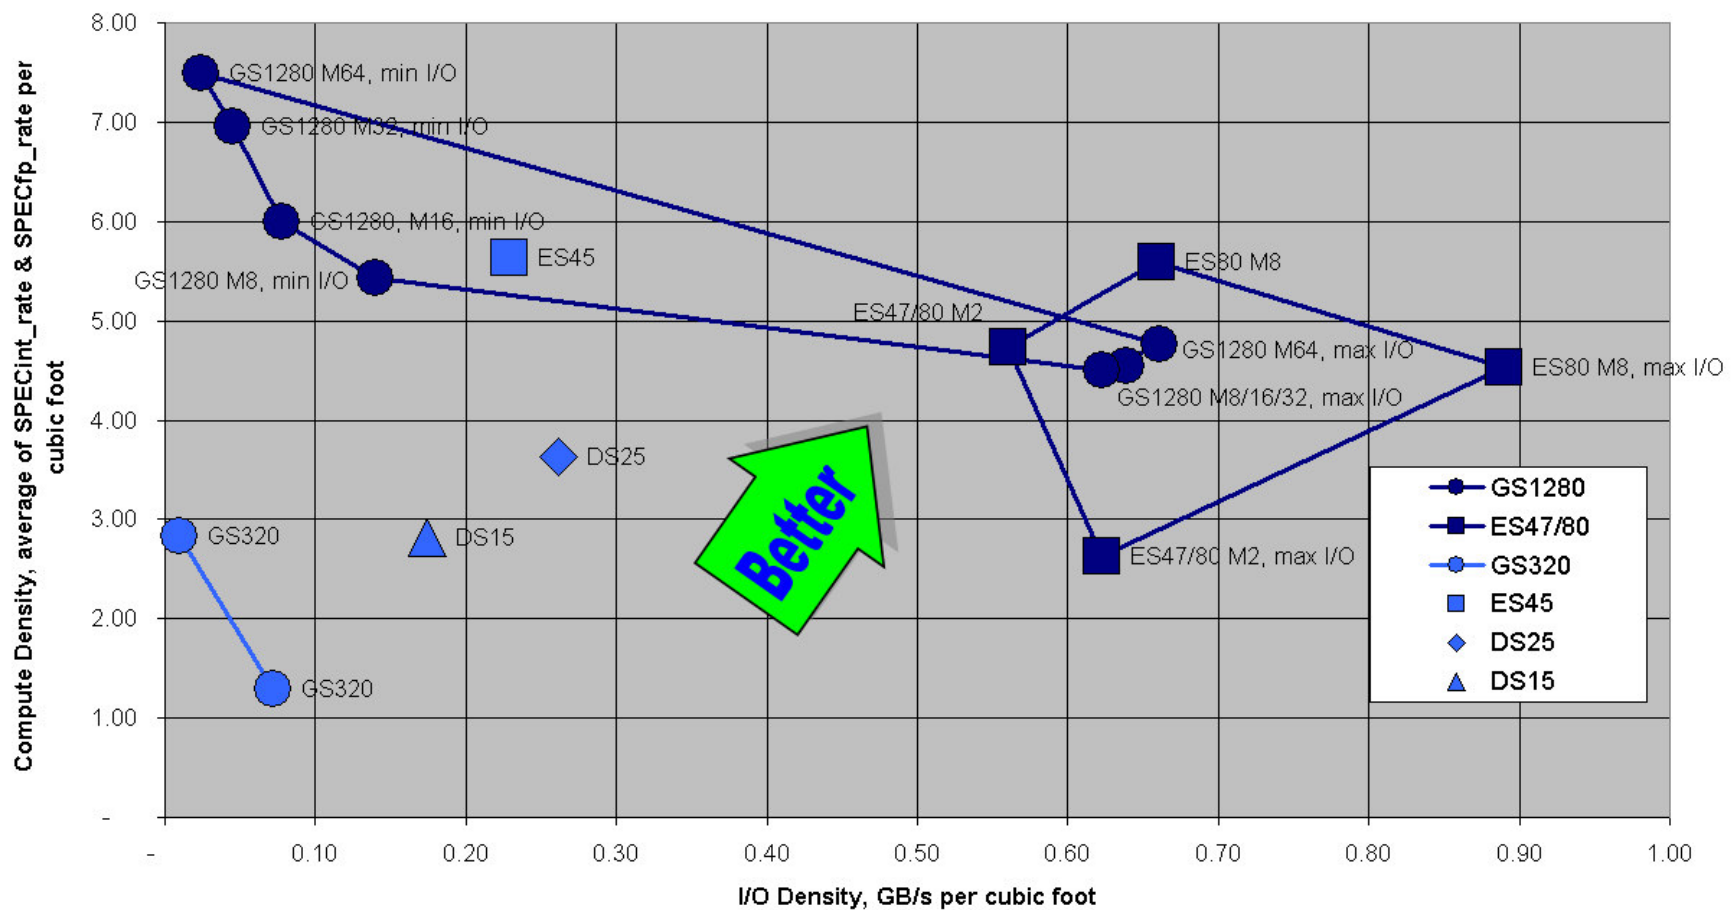
I/O Bandwidth, Delivered



Higher compute & I/O density



Comparison of Compute and I/O Density



Reliability and Availability ... Results



- Calculated

EV7/Marvel Systems reliability expected to be twice that of EV68 systems.

GS320 8P/8GB	GS1280 8P/8GB with RAID memory	Comments
5,280 Hrs MTBF	10,598 Hrs MTBF	<ul style="list-style-type: none">•GS1280 has no central system switch and infrastructure•RAID virtually eliminates memory contribution to apparent MTBF•N+1 DC/DC converters in Marvel eliminate their contribution to apparent MTBF•Both systems have the same power supply, N+1


- Actual (based on 18 months warranty data)

EV7/Marvel **four to eight times the MTBF** of EV68 systems

Reduced system complexity ... EV7 – The System is the Silicon



Building a System....



EV7 + IO7 + Memory = SYSTEM !

Reduced system complexity ...

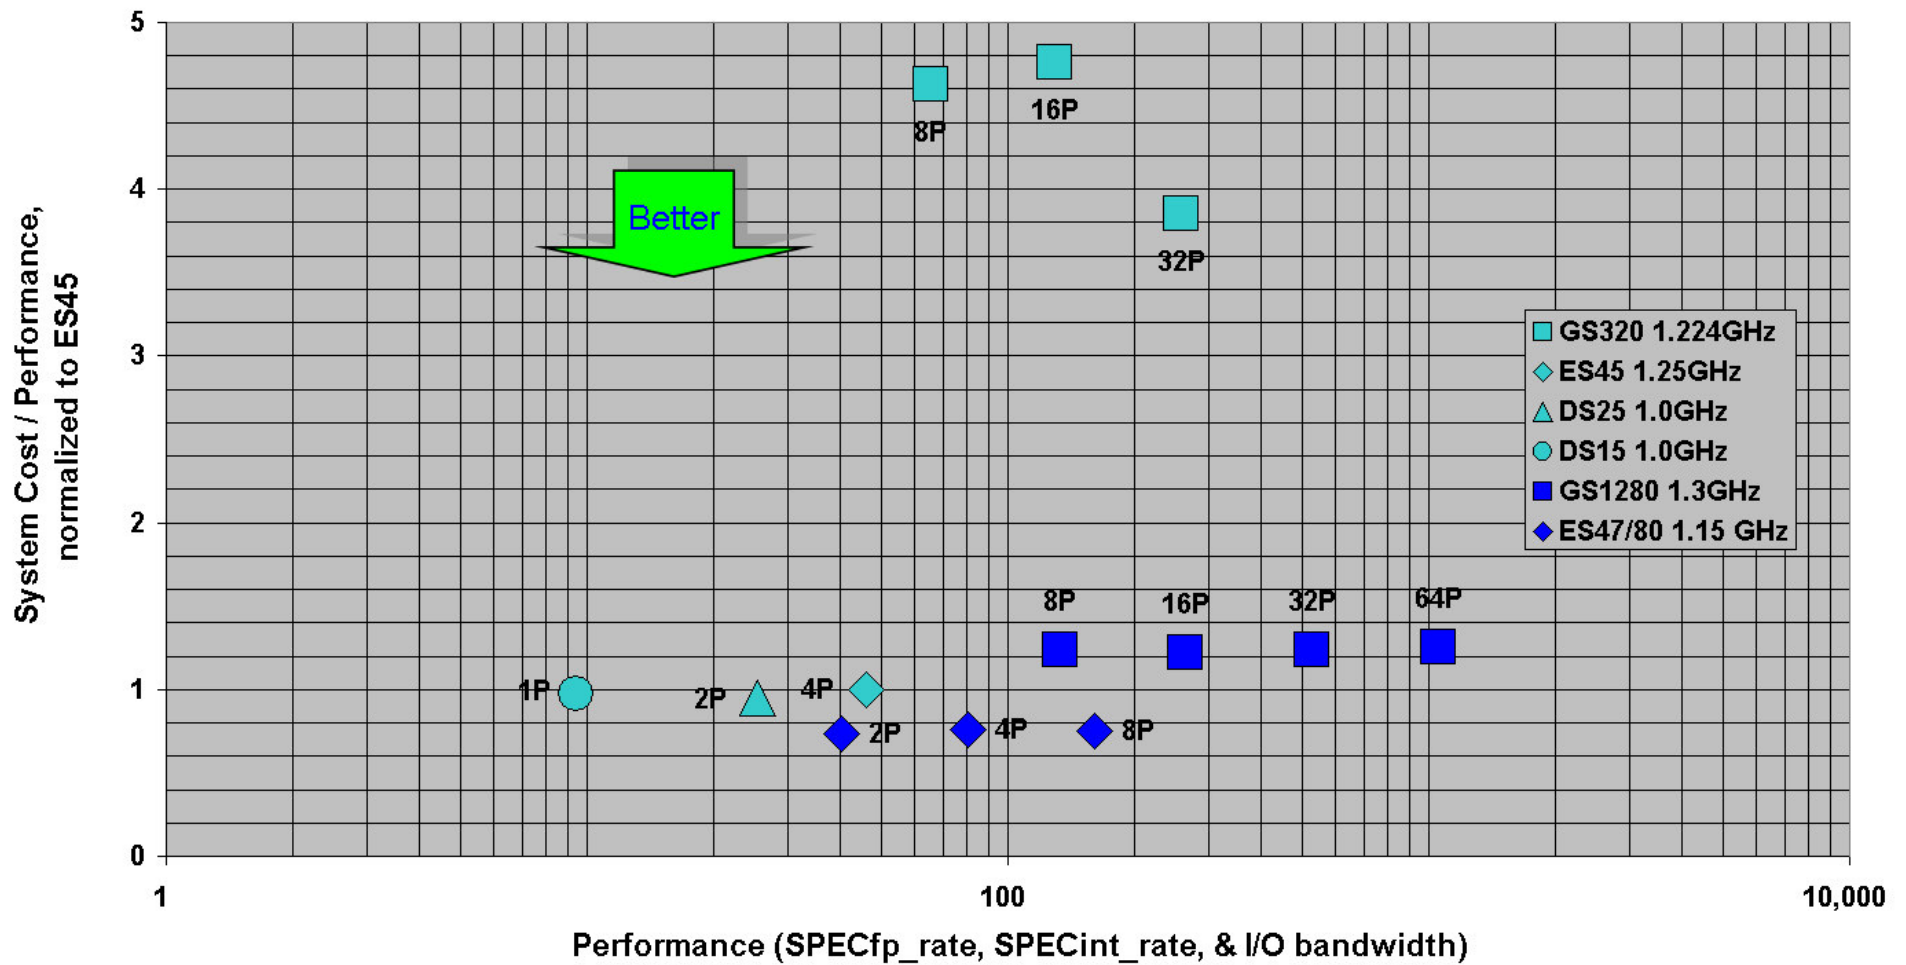
Glueless interconnect forms a two dimensional torus, mesh



Reduced System Cost



Cost/Performance



What's New ...

New EV7 Alpha CPU Speeds

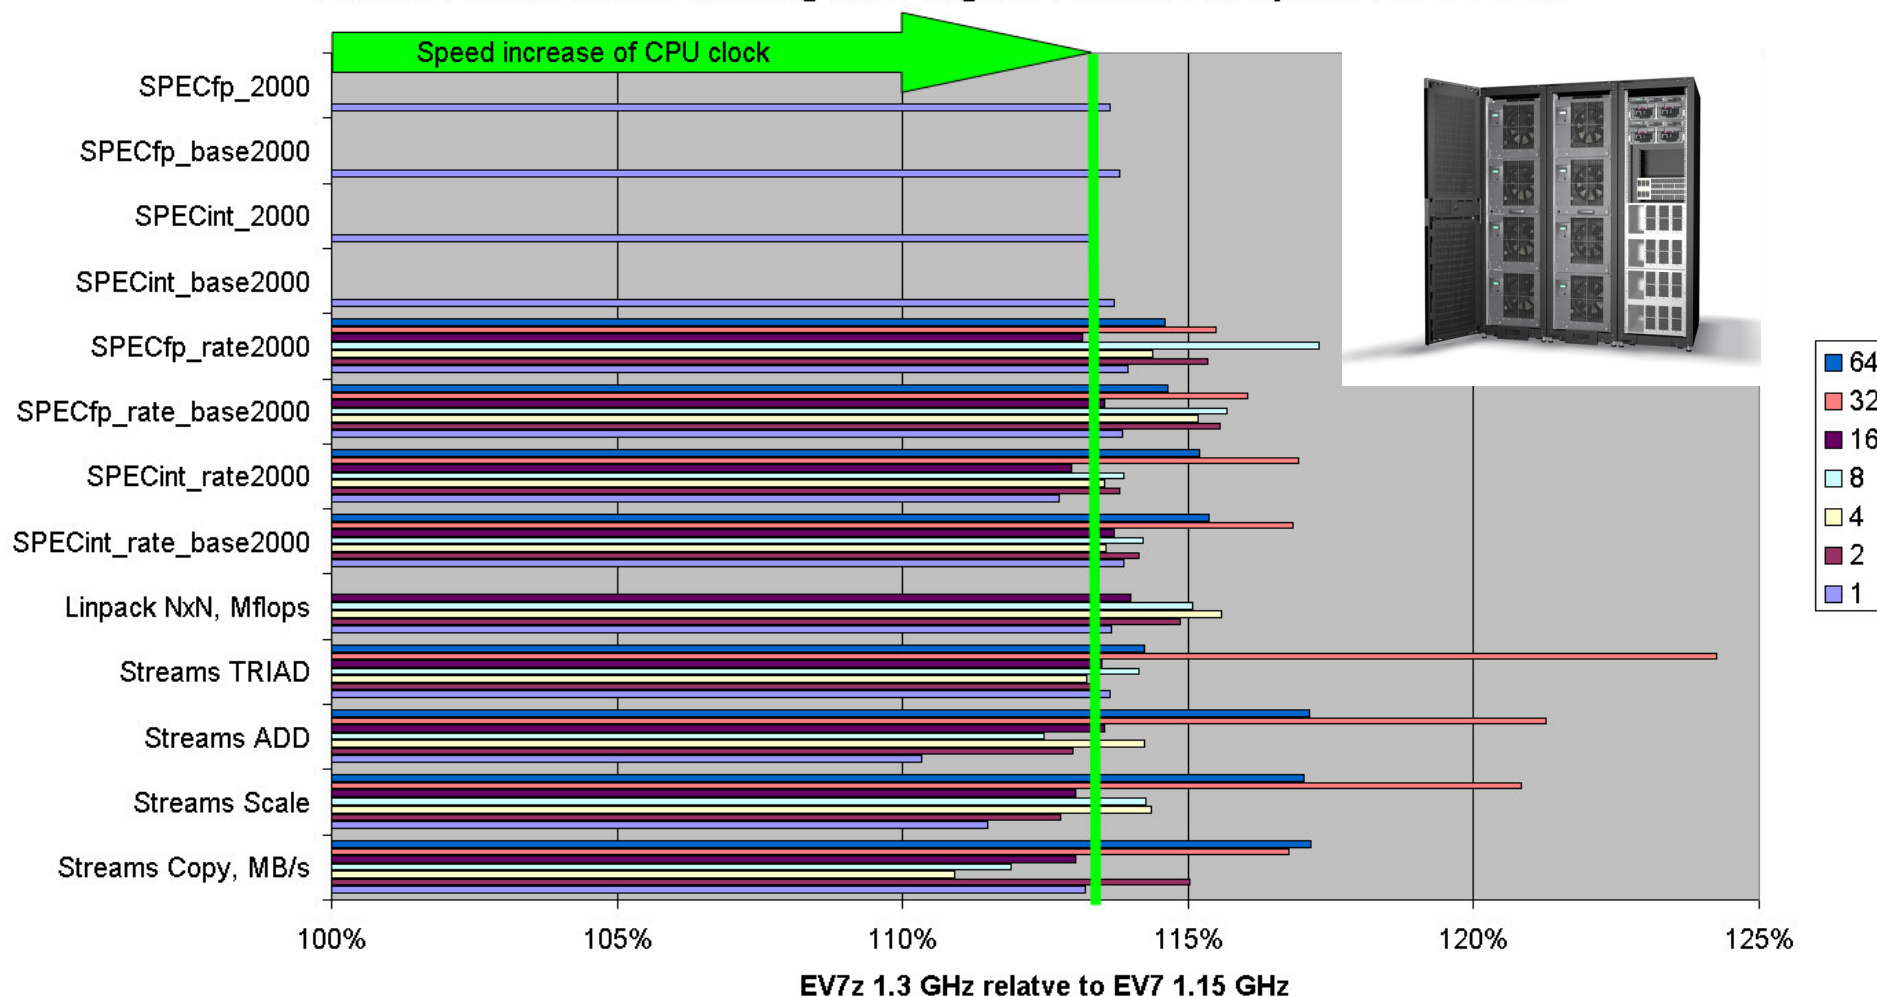


System	EV7 Speed	Jan 2003 to July 2004	August 2004
GS1280	1.3 GHz		<input checked="" type="checkbox"/>
	1.15 GHz	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
ES47/80	1.15 GHz		<input checked="" type="checkbox"/>
	1.0 GHz	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

14% More Performance for GS1280 from New EV7z CPU @1.3 GHz



Linear Performance Scaling with Higher Processor Speed for GS1280




16% More Performance for ES47/80

from New EV7 CPU @1.15 GHz




Linear Performance Scaling with Higher Processor Speed for ES47/80



Same Processor, Same Performance for ES47/80 or GS1280 with EV7 CPU @ 1.15 GHz



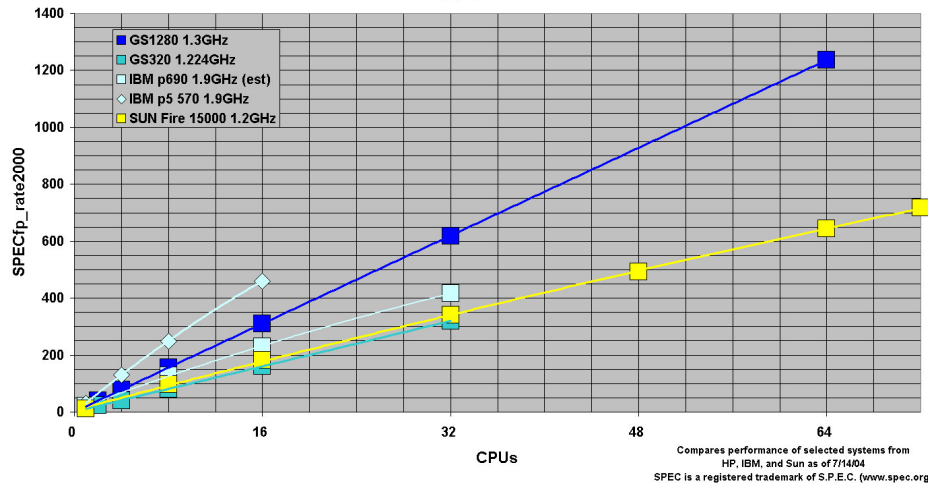
ES47/80 Compared to GS1280, both with EV7 at 1.15 GHz




Performance Comparisons




SPECfp_rate2000




SPECfp_rate2000



Memory Bandwidth



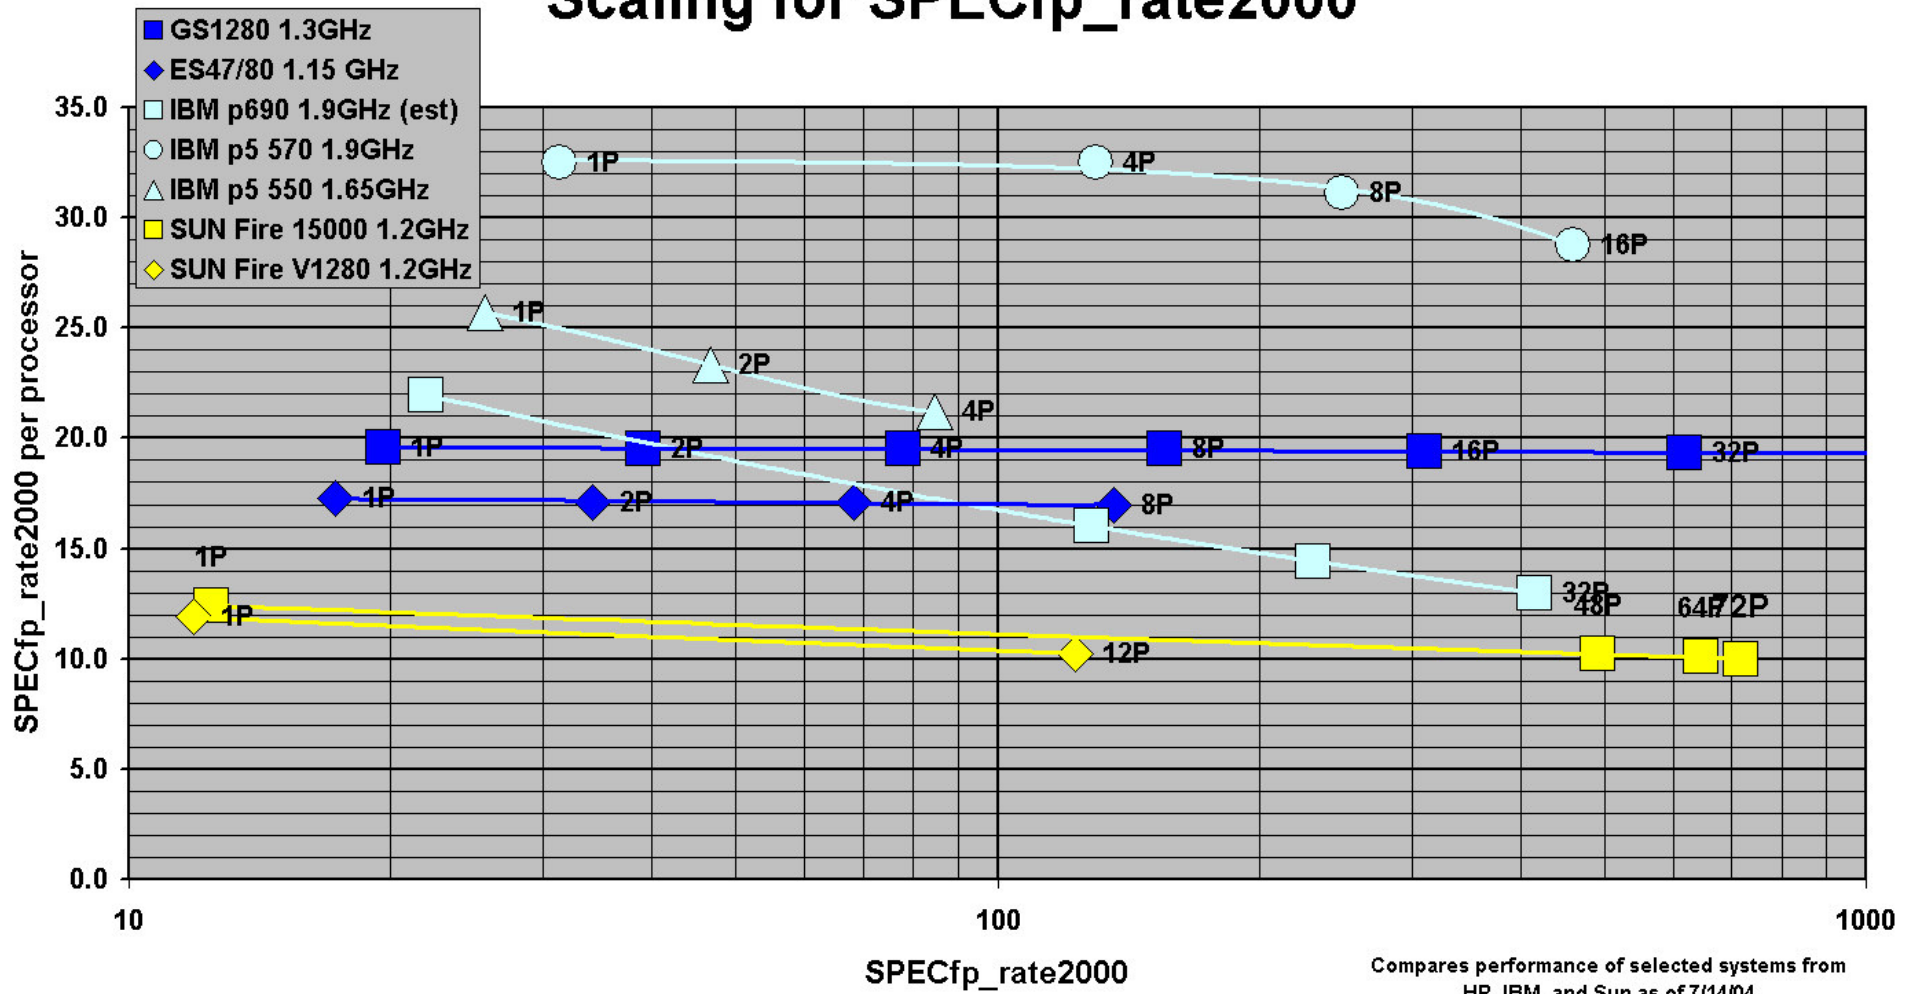
I/O Bandwidth, Delivered



Performance Comparisons



Scaling for SPECfp_rate2000



Compares performance of selected systems from HP, IBM, and Sun as of 7/14/04
SPEC is a registered trademark of S.P.E.C. (www.spec.org)

Pricing



- GS1280
 - New 1.3 GHz CPU, same price as current 1.15 GHz
 - Current 1.15 GHz CPU price **reduced** 25%
- ES47/80
 - New 1.15 GHz CPU priced the same as current 1.0 GHz
- Marvel Memory list prices **reduced** up to 40%

Up To 25% Price/Performance Improvement



Based on US List Prices	Entry Level as of Aug 1, 2004		Typical Configuration prior to Aug 1, 2004		Typical Configuration as of Aug 1, 2004		Performance Change	Price Change	Change in Price/Performance
AlphaServer ES47 Tower system	1 EV7 CPU @ 1.15 GHz 1 iCoD EV7 CPU @ 1.15GHz 512 MB RDRAM 5 PCI-X & 1 AGP I/O slots 18 GB disk	\$29,200	2 EV7 CPU @ 1.00 GHz 4 GB RDRAM 5 PCI-X & 1 AGP I/O slots 18 GB disk	\$55,039	2 EV7 CPU @ 1.15 GHz 4 GB RDRAM 5 PCI-X & 1 AGP I/O slots 18 GB disk	\$48,450	+16%	-12%	-24%
AlphaServer ES47 system	1 EV7 CPU @ 1.15 GHz 1 iCoD EV7 CPU @ 1.15GHz 512 MB RDRAM 5 PCI-X & 1 AGP I/O slots 18 GB disk	\$37,950	4 EV7 CPU @ 1.00 GHz 8 GB RDRAM 10 PCI-X & 2 AGP I/O slots 18 GB disk	\$157,211	4 EV7 CPU @ 1.15 GHz 8 GB RDRAM 10 PCI-X & 2 AGP I/O slots 18 GB disk	\$135,950	+16%	-14%	-25%
AlphaServer ES80 system	1 EV7 CPU @ 1.15 GHz 1 iCoD EV7 CPU @ 1.15GHz 512 MB RDRAM 5 PCI-X & 1 AGP I/O slots 18 GB disk	\$49,300	3 EV7 CPU @ 1.00 GHz 6 GB RDRAM 15 PCI-X & 3 AGP I/O slots 18 GB disk	\$245,887	3 EV7 CPU @ 1.15 GHz 6 GB RDRAM 15 PCI-X & 3 AGP I/O slots 18 GB disk	\$214,050	+16%	-13%	-25%
AlphaServer GS1280 system	2 EV7 CPU @ 1.15 GHz 512 MB RDRAM 8 PCI-X, 3 PCI & 1 AGP I/O slots 18 GB disk	\$96,100	16 EV7 CPU @ 1.15 GHz 32 GB RDRAM 16 PCI-X, 6 PCI & 2 AGP I/O slots 18 GB disk	\$790,667	16 EV7 CPU @ 1.30 GHz 32 GB RDRAM 16 PCI-X, 6 PCI & 2 AGP I/O slots 18 GB disk	\$705,950	+14%	-11%	-22%
					16 EV7 CPU @ 1.15 GHz 32 GB RDRAM 16 PCI-X, 6 PCI & 2 AGP I/O slots 18 GB disk	\$593,950	-	-25%	-25%

Customers' vote of confidence in HP AlphaServers and the roadmap



"HP AlphaServers are a real powerhouse for our various bioinformatics projects, especially when configured as a Tru64 UNIX cluster... AlphaServers just can't be beat for this mix of applications."



"We have had a long standing history with HP AlphaServer Technology. With each AlphaServer update, the servers continue to offer better price performance and scalability. These are major benefits to our success in offering exceptional healthcare solutions to our clients."



HP AlphaServer systems... provide the Korea Stock Exchange with the high availability, scalability, and rock-solid reliability demanded by our business...very pleased that OpenVMS will also be available on HP Integrity servers."


What's Next ...

Integration of Products has begun ...


Dual AC Power Supply from Superdome in GS1280, since March



- Dual AC, standard (previously an extra cost option)
- Less space than previous single feed AC power

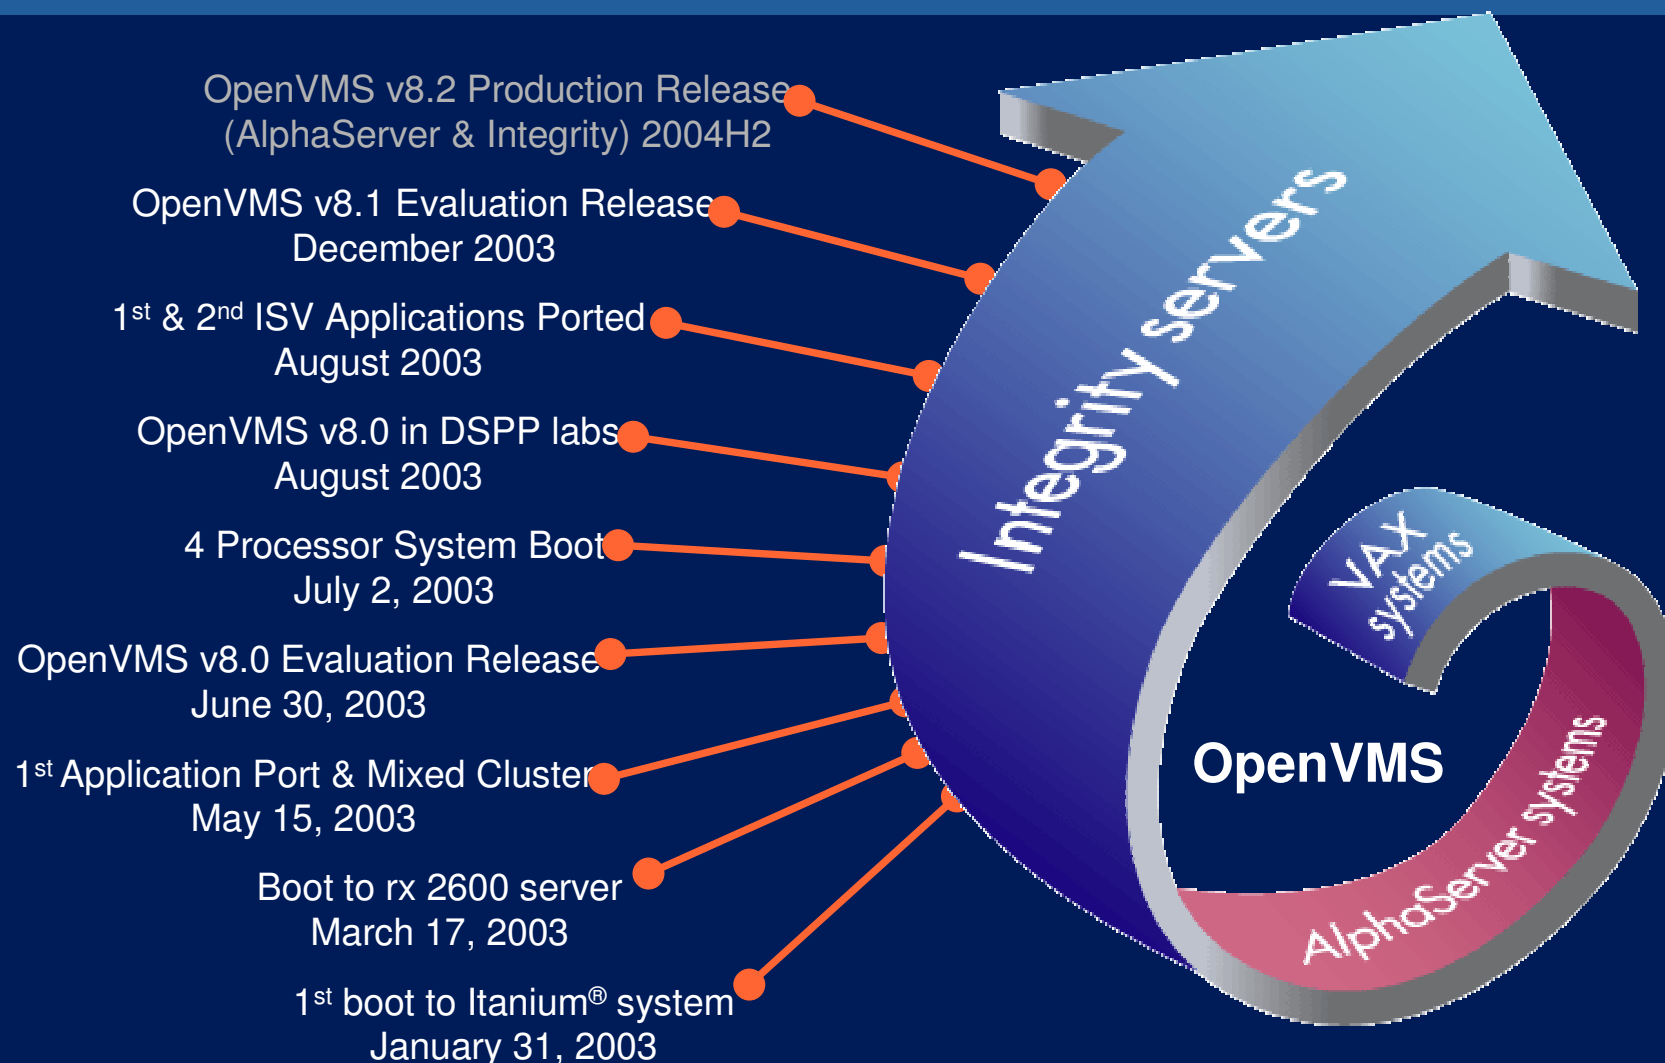


GS1280 Model 32



GS1280 Model 64

HP OpenVMS – Itanium® port status exceeding expectations



A common HP OpenVMS operating environment across Integrity and Alpha servers



OpenVMS V8.2

FRS: Q4 2004

Major New

Platform Release: Alpha & Integrity

- First production release for Integrity servers
- Mixed Alpha & Integrity clusters with shared fibre channel storage, up to 16 nodes
- Integrity server support, up to 8 CPU's (rx1600, rx2600, rx4640-8)
- Supports Itanium 2 6M and Itanium 2 9M processors
- Infrastructure changes to support future new features
- Host Based MiniMerge



Version V8.2

2004

- rx1600-2
- rx2600-2
- rx4640-8 (1-8 way platform)
- ❖ CPU: Itanium 2 9M

Version V8.3

2005

- rx1600-2
- rx2600-2
- rx4640-8
- Rx7620-16
- rx8620-32
- Superdome
- ❖ CPU: Montecito
Dual Core


Subject to change without notice - current as of March, 2004

Advanced File System for HP-UX 11i v3

Online storage management for maximum business agility




- Unlimited growth
 - Supports 16 TB files, 512 TB file systems
 - Direct I/O for high performance database access
- Single management
 - Integrated with storage management features for hardware lun expansion, closing, and snapshots
 - On-line defragmentation policy based on dynamic policy management
 - Undelete facility – trashcans
 - Journaling file system –fast restart in case of failures
- Self-tuning
 - Dynamically resize – expand and shrink as needed



The gold standard in UNIX® clustering: HP-UX 11i v3 Serviceguard with TruCluster technology

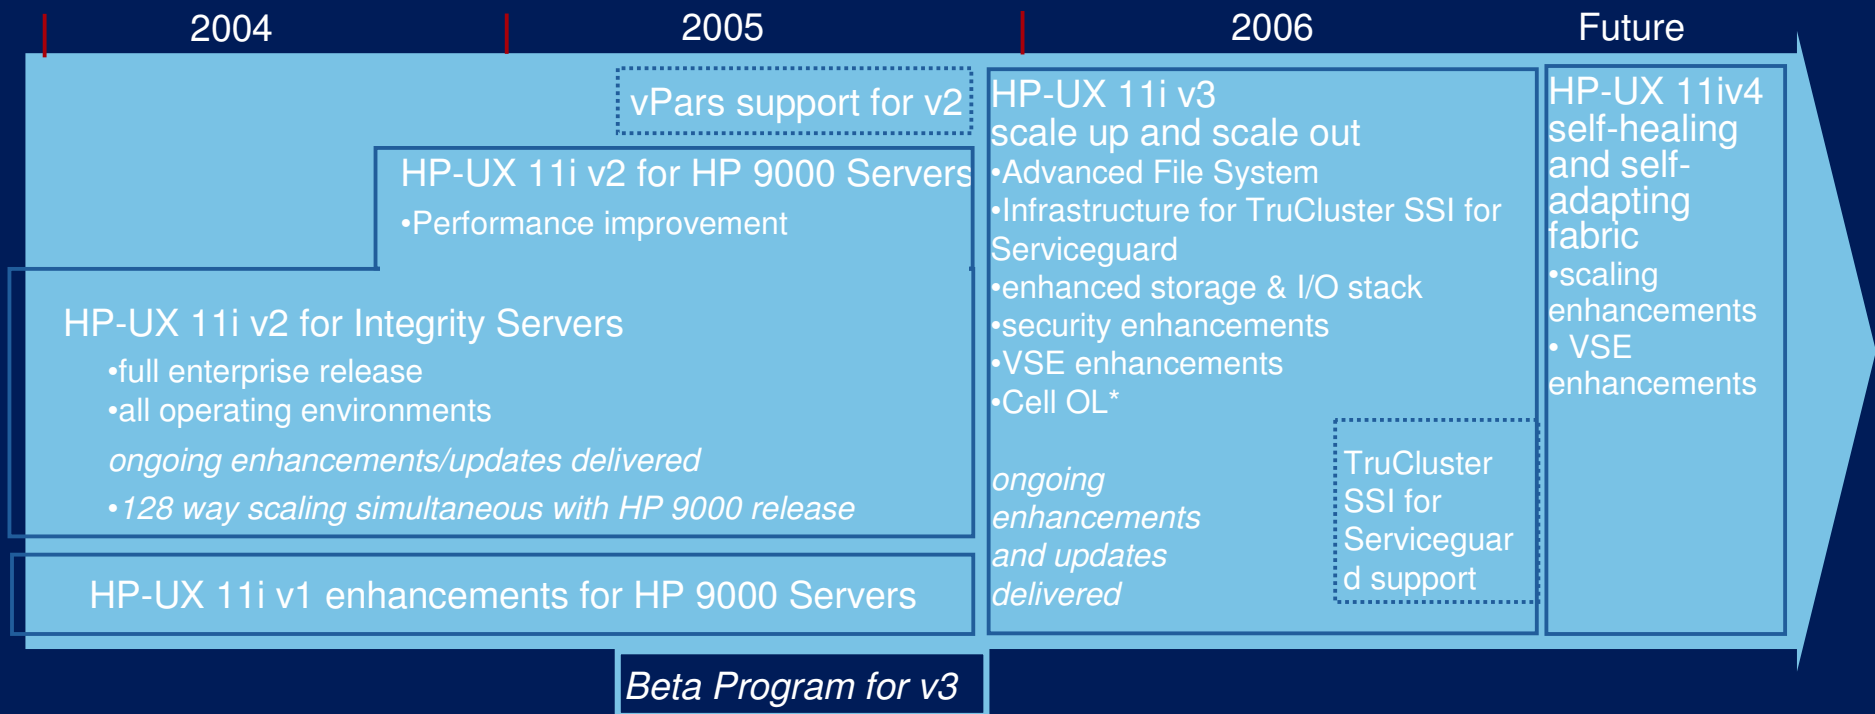


- Cluster file system with shared root
 - Manage multiple systems as one
 - More automation, fewer tasks
 - More productivity, fewer errors
 - Improved TCO
- Single unified storage management
 - All storage devices are shared and consistently named across all members, automatically
 - Configuration flexibility (direct access or not, SSI handles it)
 - Choose what to add, we just use it
 - As long as one path exists to the device, we can get there
- Scalability built in
 - Applications scale out of the box
 - Scale across all dimensions
 - Uses all wires, paths to resources
 - No superhuman effort to tune



HP-UX 11i roadmap

The UNIX foundation of the Adaptive Enterprise



HP-UX 11i is your operating environment for IT enterprise virtualization – flexibility is key

HP-UX 11i v2 Full ecosystem accelerated making it the version of choice

- Simplicity with the common release for HP 9000 and Integrity servers

HP-UX 11i v3 Advances leadership in scale-up and scale-out

- HP remains committed to Tru64 UNIX customers bringing best technology from Tru64 UNIX into HP-UX 11i & Serviceguard (AdvFS and TruCluster Single System Image)

Where to look for information



- AlphaServers

<http://h18002.www1.hp.com/alphaserver/index.html>

- AlphaServer system performance comparison

http://h18002.www1.hp.com/alphaserver/performance/perf_by_perf.html

- QuickSpecs

<http://h18000.www1.hp.com/products/quickspecs/Division/10410.html>

- System Documentation

<http://h18002.www1.hp.com/alphaserver/technology/index.html>

- Firmware

<http://ftp.digital.com/pub/DEC/Alpha/firmware/>

AlphaServers



Dependable ...
for the long haul.

HP WORLD 2004

Solutions and Technology Conference & Expo



RECOMMENDED TRAINING VENUE FOR THE
HP Certified Professional





i n v e n t