Choosing the Right Server

Tim Arkless
Systems Architect
tim.arkless@uk.ibm.com

ON DEMAND BUSINESS™

Servers exist to help solve your Business Challenges

Business Challenges have many forms:

- Floor space
- SW License charges
- Personnel costs
- SLA Requirements
- Performance
- Budgetary Constraints
- Speed of deployment
- Flexibility of cost structure
- Power
- Which of these are most important to you?

There is no "one" server that can meet all your business challenges – just as there is no "one" mode of transportation that meets all your needs.

Let's review some specific case studies:

- What is commodity hardware -- ? (Price)
- Do all CPU architectures deliver the same results -- ? (Performance)
- Meeting SLAs – are all servers "good enough" -- ? (Reliability)
- Do all solutions offer the same flexibility -- ? (Scalability)
- Do Vendor Roadmaps lead to the same result -- ? (Future Viability)
- Oracle RAC discussion

FFDC Diagnostic Value: Impact on Availability vs. Competition

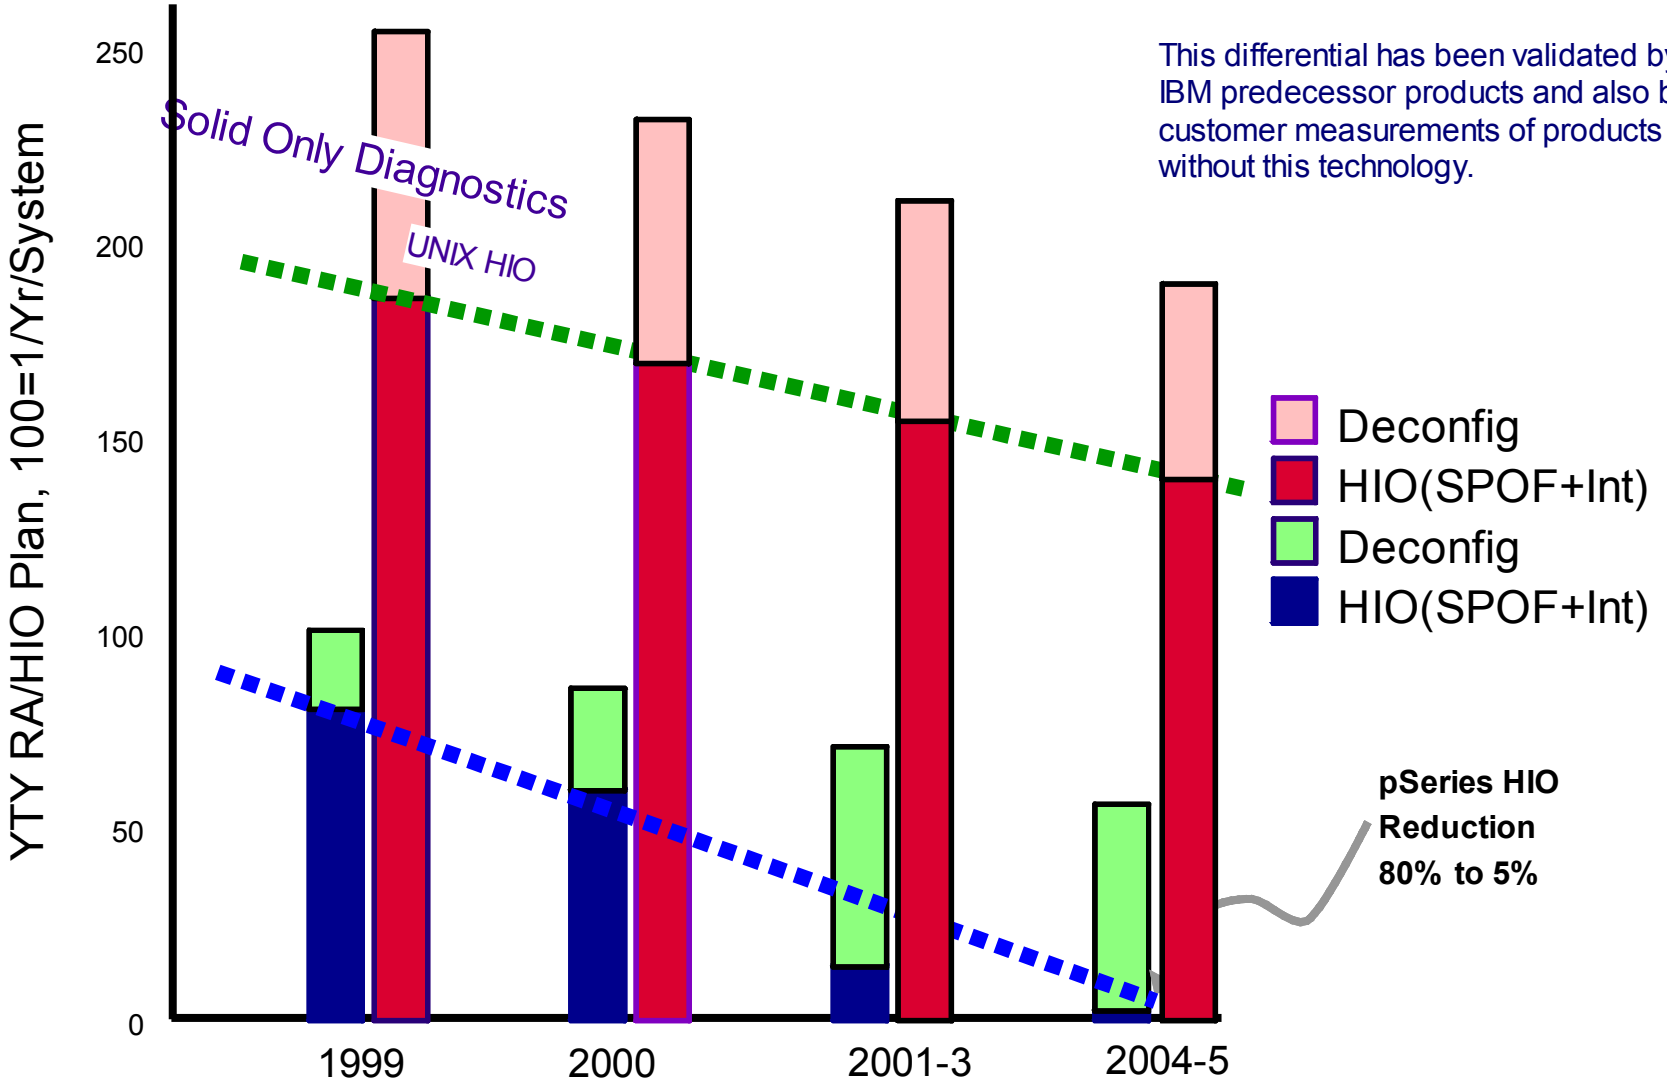

The analyses of failure probabilities and their impact on application outages and availability which are contained in this presentation are based on standard engineering methodology of accumulating failure rates based on component counts associated with various functions internal to a computer system, as well as the diagnostic effectiveness applied to the component, and the specific recovery action designed for the component. This data represents IBM measurements and projections.

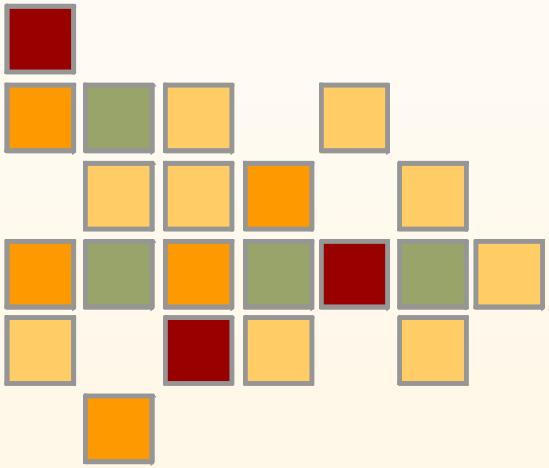

The Price Question

"What is commodity hardware?"

"Which servers offer the lowest TCO"

1U – 2 chip/4 core systems

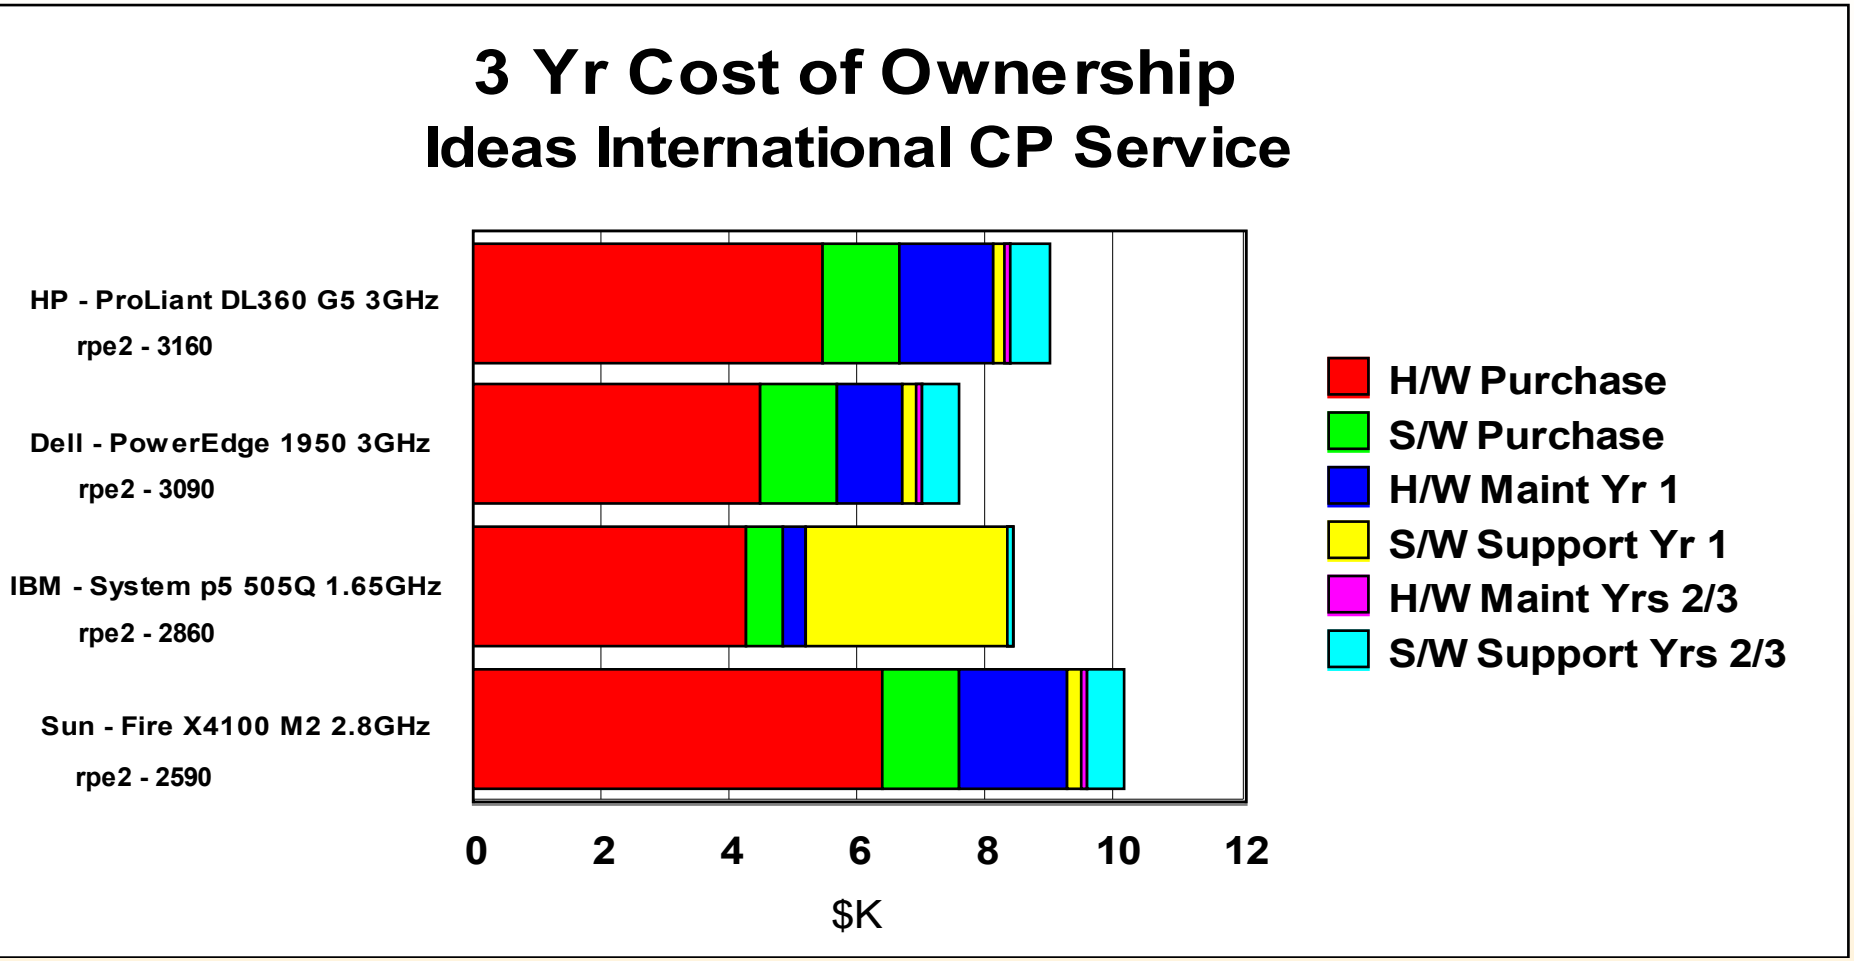

Each system has 2GB memory, 2@73GB drives, NBD Support, All except p505Q have Windows Server 2003, 10 CALS, Prices current from Ideas International as of 1/17/07

2U – 2 chip/4 core systems

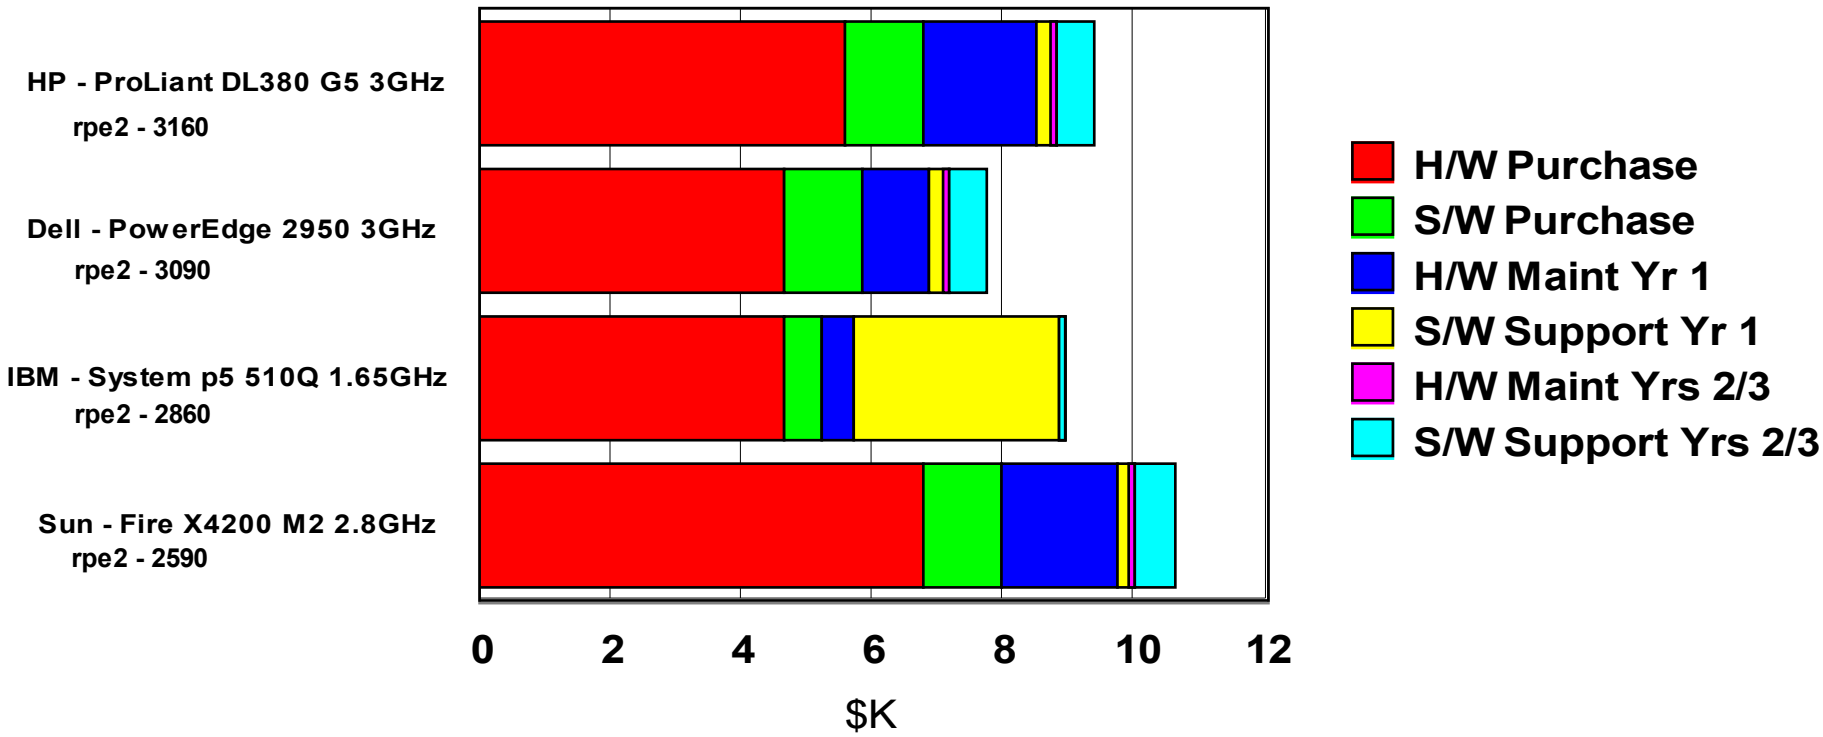

Each system has 2GB memory, 2@73GB drives, NBD Support, All except p510Q have Windows Server 2003, 10 CALS, Prices current from Ideas International as of 1/17/07

4U – 4 chip/8 core systems

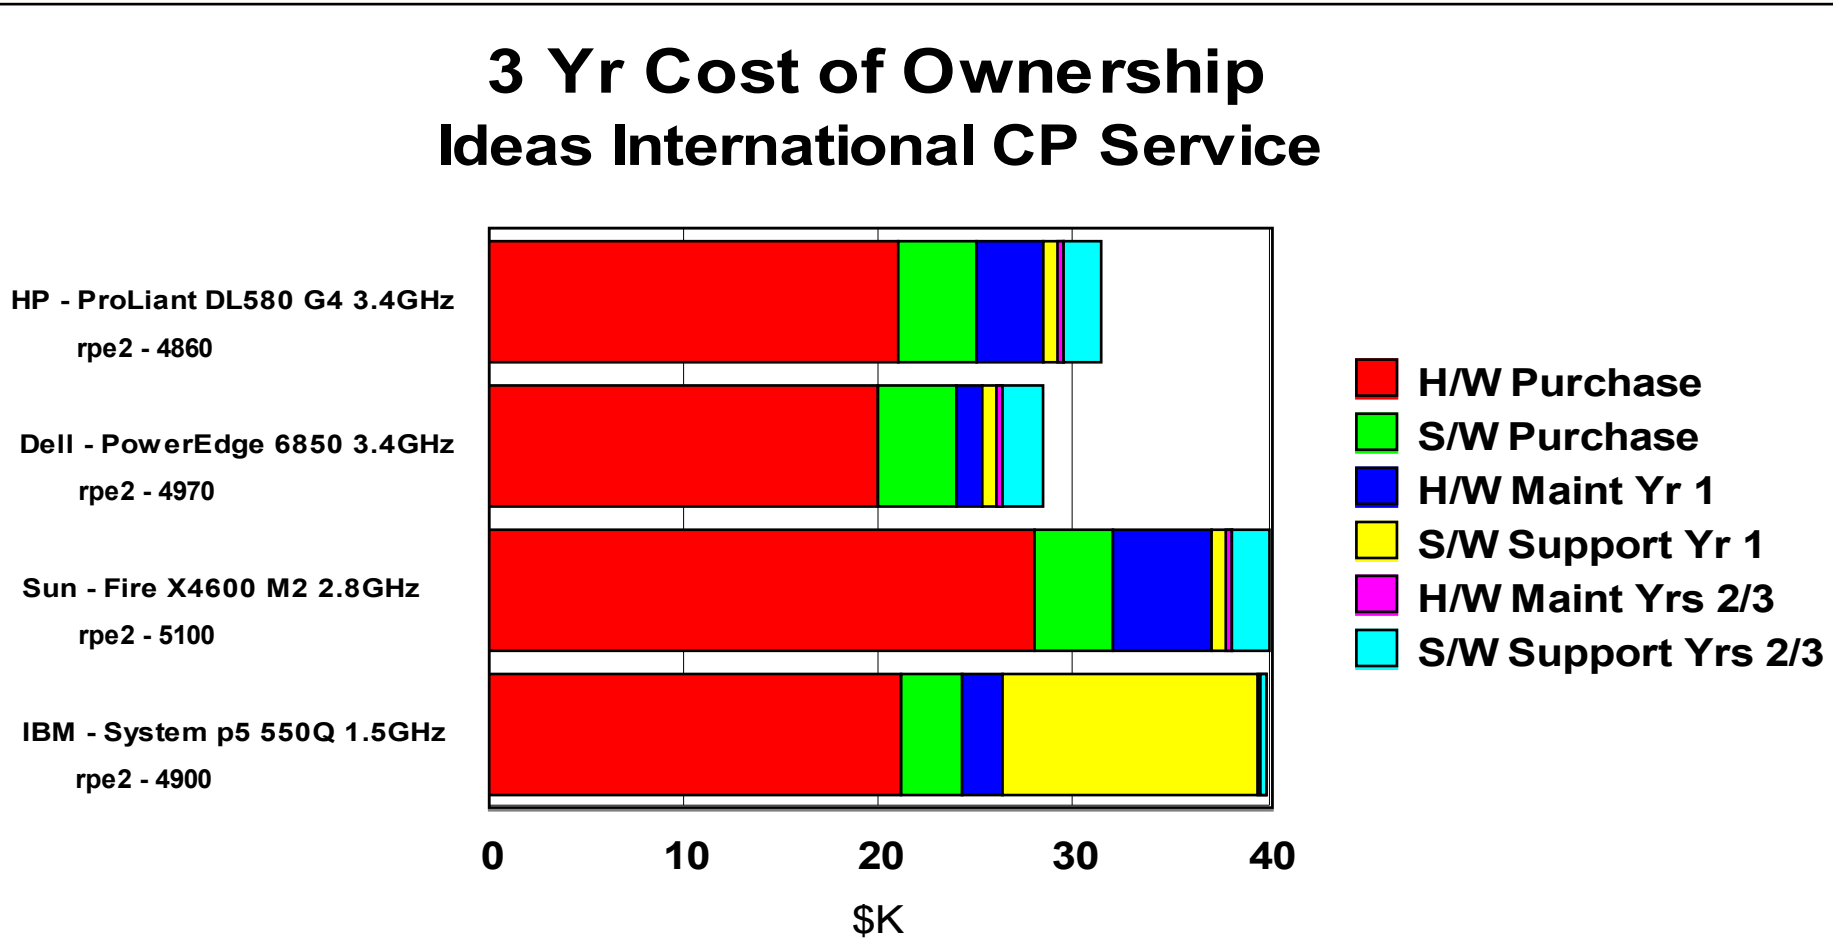

Each system has 16GB memory, 2@73GB drives, 24x7x4 Support, All except p55Q have Windows Server 2003, 25 CALS, Prices current from Ideas International as of 1/17/07

ITG: Average Three-year Server Costs for SAP Systems - Windows Server and IBM System p5

IBM System p5 - DB2 Databases, Windows Server - SQL Server

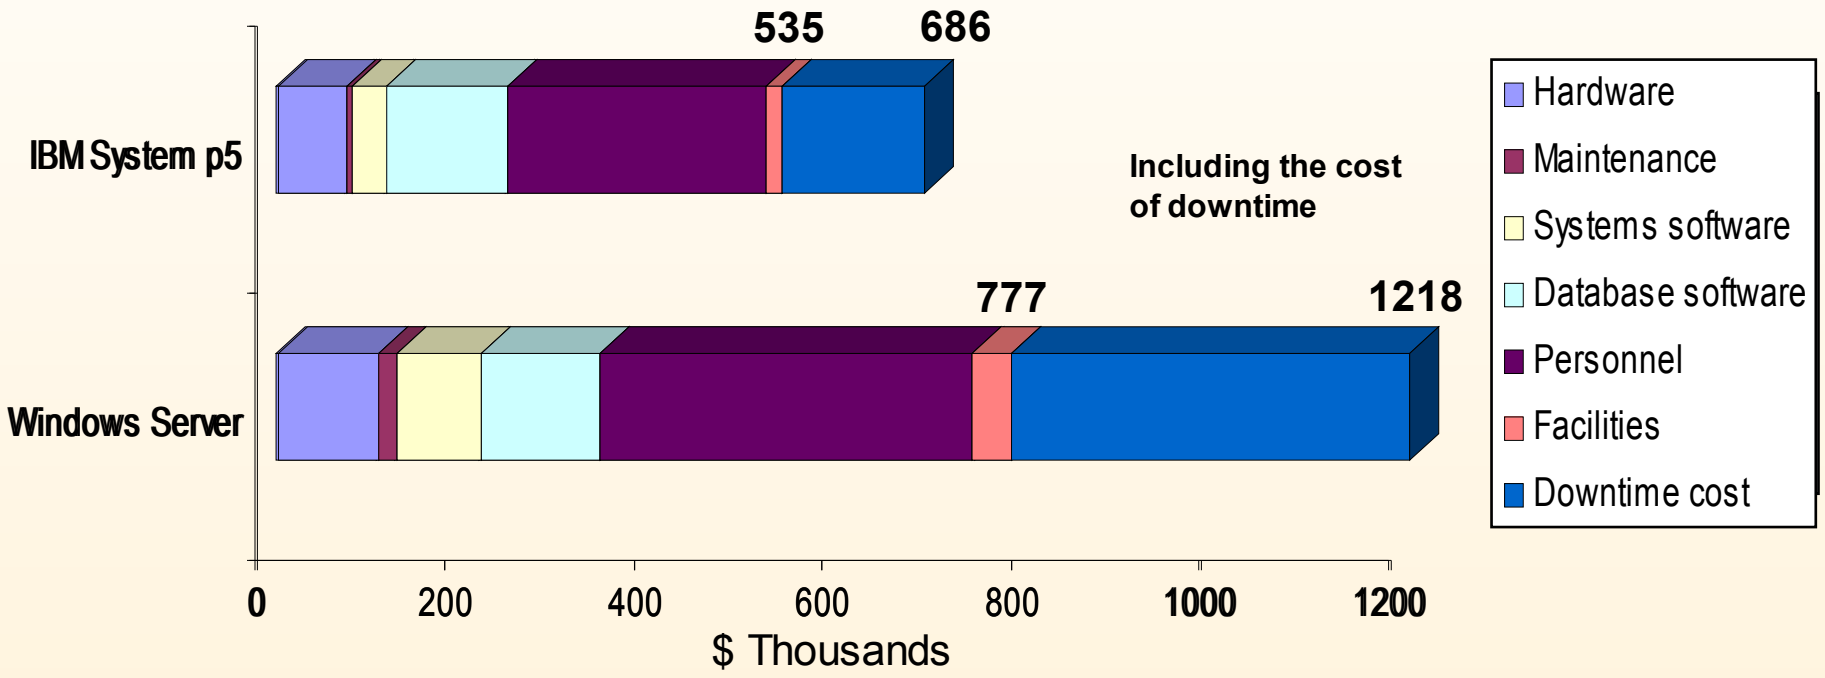

COST/BENEFIT CASE FOR IBM SYSTEM p5 FOR SAP SOFTWARE DEPLOYMENTS
Compared with Windows Servers
International Technology Group March 2006

Another Approach – Virtualization

-- Is there a better way to solve this problem (Virtualized Servers?)

Potential Savings are for System p5 virtualized compared to conventional scenarios. (**Overall average savings of 55 – 62%**)

Areas of Operating Cost Reduction for System p5 Virtualized Scenarios

Maintenance 69% - 76% Fewer, newer servers reduce maintenance contract costs.

Software 65% - 69% Fewer software copies & CPUs result in lower license, update & support costs.

Personnel 31% - 45% Fewer physical servers, reduced diversity & improved automation reduce system administration & related personnel costs.

Facilities 52% - 61% Fewer physical servers, smaller footprints, and greater energy
efficiency reduce data center occupancy, power & cooling costs.

From ITG Study, 2006, on the Economic Benefits of Infrastructure Simplification

- ***We would welcome the opportunity to do a specific study for you…***

System p5 virtualized compared to conventional scenarios.

(**Overall average savings of 55 – 62%**)

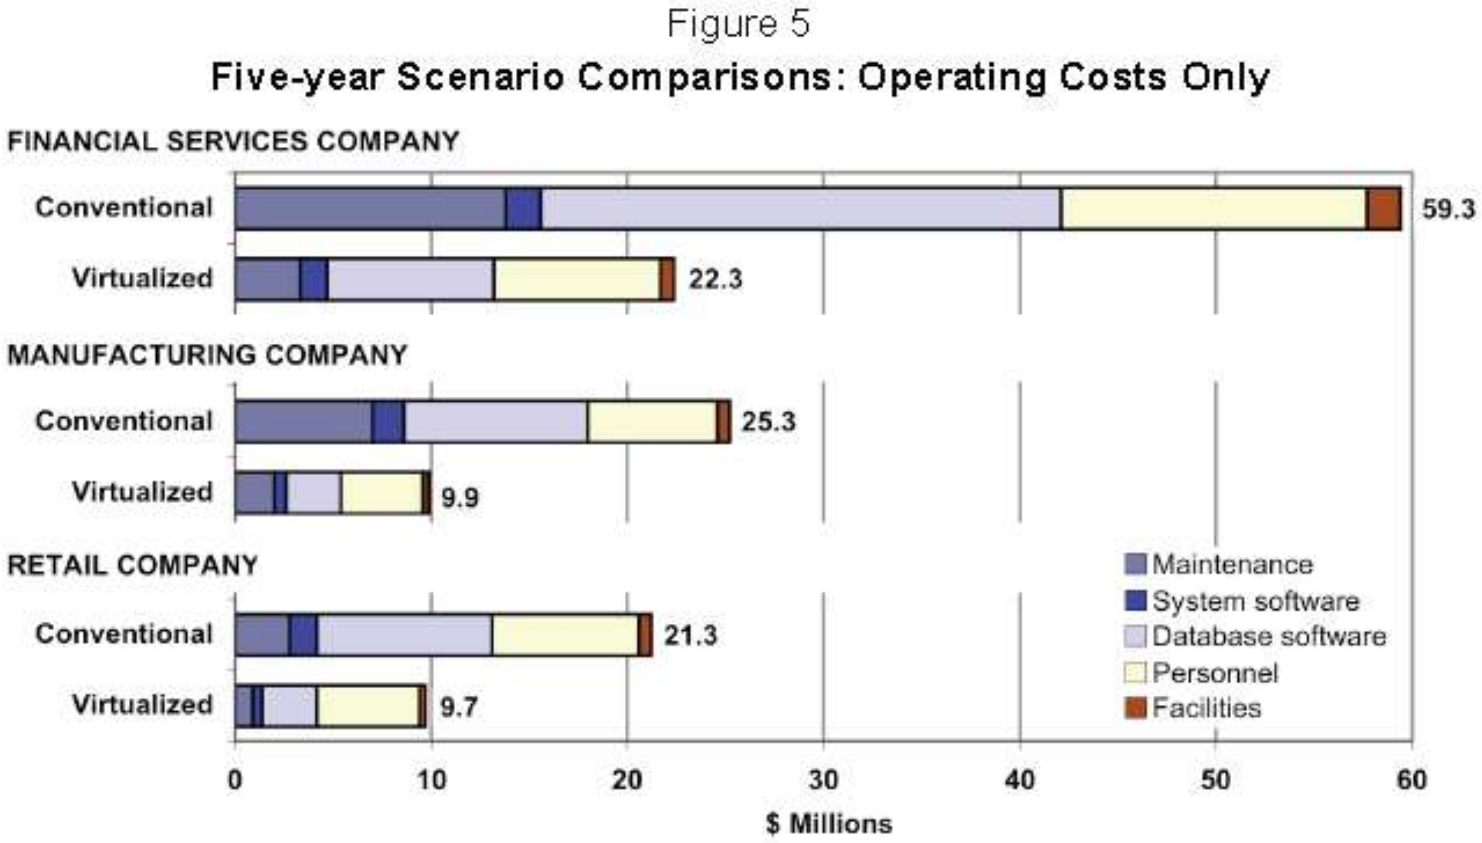

From ITG Study, 2006, on the Economic Benefits of Infrastructure Simplification

Build using today's reality – not yesterday's paradigm

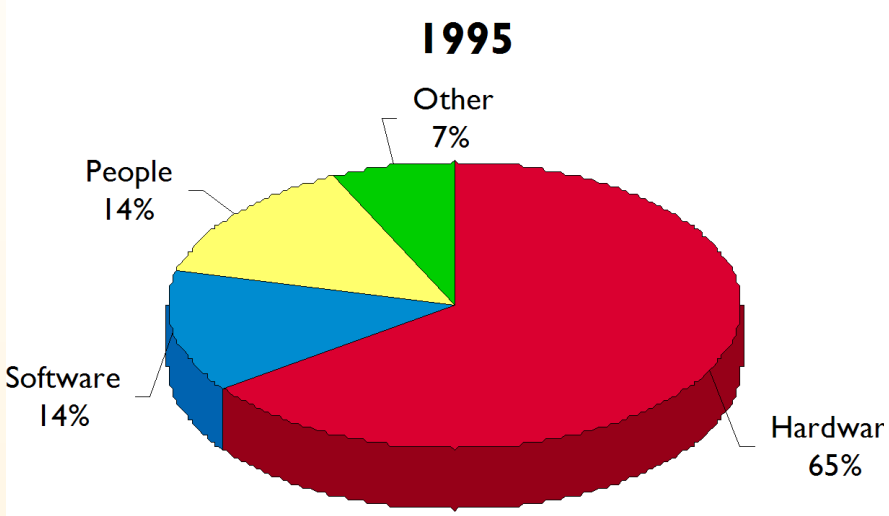

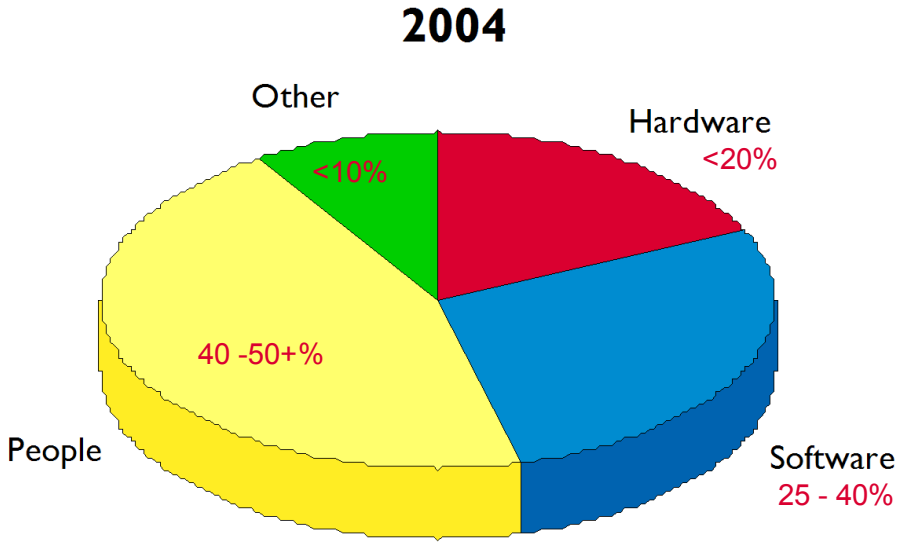

People expense has tripled as a %

Software expense has doubled as a %

Hardware is less than 1/3 of its original %

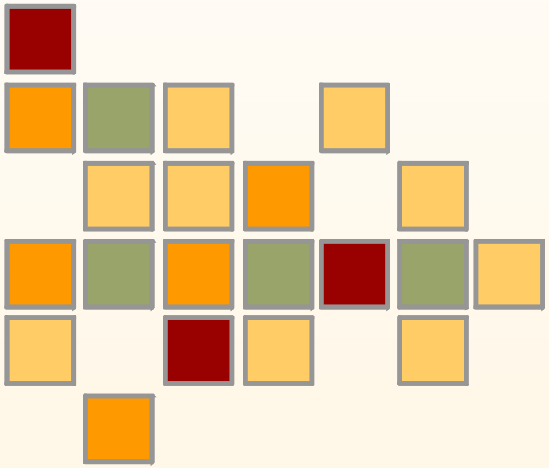

The Performance Question

"Are all CPUs created equal?"

POWER6 p570 Benchmarks

Run Type	Benchmark	IBM p570 4.7GHz	Current BoB	Best of Breed (BoB) system	Leadership Claim
1-core	SPECint2006	21.6	18.5	Dell (Intel X6800, 2.93 GHz)	#1 overall
	SPECfp2006	22.3	18.1	HPrx6600 (1.6GHz/24MB Itanium2)	#1 overall
2-core	SPECint_rate200660.9	31.1	Intel EE 2.93GHz motherboard		#1 2-core
	SPECfp_rate2006	58.0	32.0	Bull Escala PL250R+	#1 2-core
	SPECjbb2005	88,089	52002	PowerEdge 840 Xeon 3.07GHz	#1 2-core
4-core	SPECint_rate2006122	56.6	Tyan (AMD Opt DC 3.0GHz)		#1 4-core
	SPECfp_rate2006	115	62.5	Bull PL450R+ (IBM p5 2.1GHz)	#1 4-core
	Linpack HPC	61.56	33.7	IBM BladeCenter (2.5GHz)	#1 4-core RISC
	SPECjbb20005	175,474	138,388	Fujitsu (woodcrest, 3.0GHz)	#1 4-core
8-core	SPECint_rate2006240	102.0	HPrx6600 (1.6GHz/24MB Itanium2)		#1 8-core
	SPECfp_rate2006 213	91.3	HP DL585 (AMD Opt DC 2.8GHz)		#1 8-core
	Linpack HPC	120.6	66.44	IBM p5 575 2.2GHz	#1 8-core RISC
	SPECjbb20005	346,742	225,042	Dell (Xeon 2.67GH, 1333MHz bus)	#1 8-core
16-core	SPECint_rate2006478	217	Bull PL1650R+ (IBM p5 2.2GHz)		#1 16-core
	SPECfp_rate2006	426	248	Bull PL1650R+ (IBM p5 2.2GHz)	#1 16-core
	Linpack HPC	239.4	111.4	IBM p5 575 1.9GHz	#1 16-core RISC
	SPECompM2001	86,624	56,211	IBM p5 570 2.2GHz	#1 16-core
	SPECjbb20005	691,975	336,653	Fujitsu (Xeon 3.5GHz)	#1 16-core
	TPC-C (KtpmC)	1,616,162	1,025,000	IBM p5 570 2.2GHz	#1 16-core

Transaction Performance - Single System TPC-C V5

System	Chip/Core /Thread	Avail.	$/tpmC	Database	OS	tpmC	tpmC/Core
IBM p570 (4.7 GHz POWER6)	2/4/8	11/26/07	$3.50	Oracle 10g	AIX V5.3	404,462	101,115
IBM p570 (4.7 GHz POWER6)	8/16/32	11/21/07	$3.54	IBM DB2 9	AIX V5.3	1,616,162	101,010
IBM Power 595 (5.0 GHz POWER6)	32/64/128	12/10/08	$2.81	IBM DB2 9.5	AIX V5.3	6,085.166	95,081
IBM p5-570 (2.2 GHz POWER5+)	8/16/32	05/31/06	$4.42	IBM DB2 v8.2	AIX V5.3	1,025,169	64,073
IBM p5-595 (2.3 GHz POWER5+)	32/64/128	01/22/07	$2.97	IBM DB2 9	AIX V5.3	4,033,378	95,081
HP rx6600 (1.6 GHz Itanium 2)	2/4/8	12/01/06	$2.63	Oracle 10g	HP-UX 11.iv2	230,569	63,021
IBM p5-570 (1.9 GHz POWER5)	8-core	09/30/04	$4.99	IBM DB2 v8.1	AIX V5.3	429,899	53,737
IBM p5-570 (1.9 GHz POWER5)	4-core	10/17/05	$3.93	Oracle 10g	AIX V5.3	203,439	50,859
IBM p5-570 (1.9 GHz POWER5)	16-core	09/30/04	$4.95	IBM DB2 v8.1	AIX V5.3	809,144	50,571
HP rx4640 (1.6 GHz Itanium 2)	2/4/8	09/01/06	$2.75	Oracle 10g	HP-UX 11.iv2	200,829	50,207
IBM p5-595 (1.9 GHz POWER5)	32-core	04/20/05	$5.05	Oracle 10g	AIX V5.3	1,601,784	50,055
IBM p5-570 (1.9 GHz POWER5)	4-core	09/30/04	$5.62	Oracle 10g	AIX V5.3	194,391	48,597
HP rx6600 (1.6 GHz Itanium 2)	4/8/16	06/11/07	$1.81	SQL EE	Windows EE	372,140	46,517
IBM p5-570 (1.9 GHz POWER5)	8-core	09/30/04	$5.26	Oracle 10g	AIX V5.3	371,044	46,380
HP rx4640 (1.6 GHz Itanium 2)	4/8/16	09/01/06	$2.71	SQL EE	Windows EE	290,644	36,330
HP ProLiant DL585 (2.6 GHz Opteron)	4-core	05/06/05	$2.80	SQL EE	Windows EE	130,623	32,655
HP Superdome (1.6 GHz Itanium 2)	64/128/256	08/06/07	$2.93	Oracle 10g	HP-UX 11iv3	4,092,799	31,974
HP ProLiant DL585 (2.4 GHz Opteron)	8-core	12/05/05	$2.02	IBM DB2 v8.2	Windows EE	236,054	29,506
HP ProLiant DL585 (2.2 GHz Opteron)	8-core	05/31/05	$2.04	SQL EE	Windows EE	187,296	23,412
HP Superdome (1.6 GHz Itanium 2)	64-core	06/05/06	$4.82	SQL EE	Windows DE	1,231,433	19,241
HP rx5670 (1.5 GHz Itanium 2)	16-core	04/15/04	$4.56	SQL DE	Windows EE	301,225	18,826
HP Superdome (1.5 GHz Itanium 2)	64-core	04/14/04	$8.33	Oracle 10g	HP-UX 11.i	1,008,144	15,752
Sun Fire	No Published Results						?

Source: http://www.tpc.org Not all results listed. Results listed with processor chip/core/thread.

All results are as of 06/12/08.

Scalability

Performance begins with the individual server…

pSeries offers…

- Industry leading SMP capacity scale-up
- Industry leading "per processor" performance
- Industry leading price/performance

Faster processors…

- Execute equivalent pathlength in shorter period of time
- Shorten transaction response time
- Minimize "concurrency management" issues
  - Dispatching queues, lock/latch contention etc...

"Use fewer, faster CPUs instead of more, slower CPUs"
Source: Oracle "Maximum Availability Architecture" whitepaper

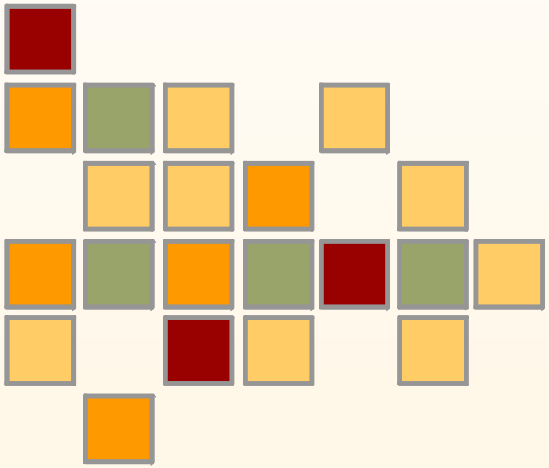

The Reliability Question

"Are all CPUs equally reliable?"

When you run a Marathon you learn a lot at mile 18!

Mainframe Class RAS

General
- Copper & SOI Chip Technology
- MCM Packaging
- FFDC & Repeat Gard
- Light Path Diagnostics
- Microcode Discovery Service
- Capacity Advantage

Cache
- Spare L1 & L2 cache bits
- Spare L2 & L3 directory bits
- L2 & L3 cache deallocation
- L3 cache line delete

Memory
- ECC Chipkill™ memory
- Spare memory chips
- Hardware memory scrubbing
- Book packaging

Hardware Management Console
- Redundant HMC
- Service Focal Point
- Service History Log
- Phone Home capabilities

Service Processor
- Boot time and operational surveillance
- Environmental monitoring
- Local / remote console

First Failure Data Capture
Monitor Errors
Determine Disposition
Predictive Failure Analysis
Notification Process

Power
- zSeries N+1 hot plug power subsystem and line cords
- Optional internal batteries

Processor
- Dynamic Processor Deallocation
- CEC Bus retry and recovery
- LPAR fast reboot

I/O Drawers
- Redundant I/O links
- Redundant power supplies

Cooling
- zSeries N+1 Hot plug fans

Disk and Adapter
- Hot swap PCI adapters
- Hot swap disk drives
- PCI bus recovery
- PCI bus deallocation
- Bad block relocation

Power5 RAS compared to x86 RAS

Reliability/Availability features	Power	x86	Comments
Automatic First-Failure Data Capture and diagnostic fault isolation capabilities	Yes	No	Used by Error Log Analysis Tool
Self-healing internal POWER5 processor array redundancy	Yes	No	ECC, bit steering, memory scrubbing, etc
Industry-first PCI bus parity error recovery	Yes	No	EEH detection: partition down vs system
Scrubbing and redundant bit-steering for self-healing in main storage	Yes	Limited	x86 not as robust
ECC and Chipkill correction in main storage	Yes	Yes	
Fault tolerance with N+1 redundancy, dual line cords, and concurrent maintenance for power/cooling	Yes	Yes	
Predictive failure analysis on processors, caches, memory, I/O and DASD	Yes	Limited	Lintel does not have predictive analysis of I/O
Processor run-time and boot-time de-allocation based on run-time errors (Dynamic Processor De-allocation and Persistent Processor De-allocation)	Yes	No	FFDC advantage
Fault avoidance through highly reliable component selection, component minimization and error mitigation technology internal to chips	Yes	No	
Service Processor is a separate, independent processor that provides hardware initialization during system IPL, operation monitoring of environmental and error events	Yes	Limited	Functions Limited on x86

High Availability…What is it worth to you?

Downtime for a typical computing infrastructure is estimated at $42,000 (USD) per hour. At this rate, a 1 percent improvement in availability can lead to millions in reduced risk and productivity losses.

Unplanned Downtime (Mission Critical)	Typical Uptime	Hours Down per Year	Cost per Unplanned Downtime Hour	Downtime Risk
Worse than average	**98.000%**	**174.72**	**$42,000**	**$7,338,240**
Average	**99.000%**	**87.36**	**$42,000**	**$3,669,120**
Better than average	**99.500%**	**43.68**	**$42,000**	**$1,834,560**
Good	**99.900%**	**8.736**	**$42,000**	**$366,912**
Best in class	**99.999%**	**.09**	**$42,000**	**$3,780**

Estimated downtime impact—Source: Alinean (http://www.alinean.com/)

AIX is “Most Reliable”

According to a recent Yankee Group study* of 400 Windows, Linux and UNIX users, AIX was the most reliable server operating system:

“IBM’s AIX achieved the highest level of reliability, with corporate enterprises reporting an average of only 36 minutes of downtime per server in a 12-month period”

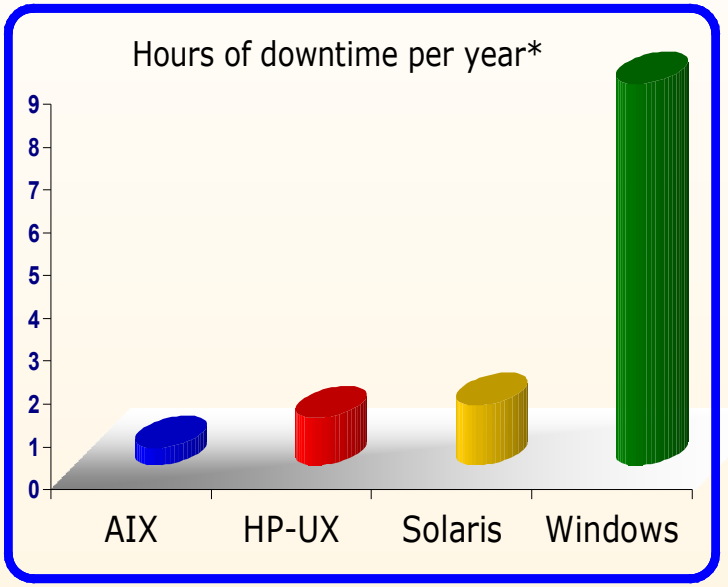

*** Source: “Unix, Linux Uptime and Reliability Increase; Patch Management Woes Plague Windows”**

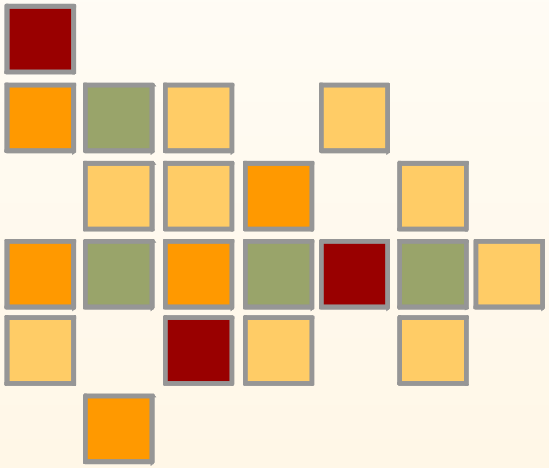

The Virtualization Question

"Are all approaches the same?"

Flexibility: Partitioning Technologies

Hard Partitions with multiple nodes	Software Virtual Machines	OS Partitions	Hardware Virtual Machines
Clusters Blades Physically Partitionable E-25K	EMC VMWare	IBM WLM & eWLM Solaris 10 Containers HP WLM (PRM)	IBM System z9 IBM System p5
- LAN in a Can - Hardware Board based Partitioning - Multiple OS images - Limited Configurations Slow Interconnect	- Software emulation layer - CPU cycle granularity 2 CPU pool limit High overhead Emulation has high overhead low perf	- Dynamic resource allocation - automatic goal-based resource allocation via set SLOs - share (%) granularity - Single OS image All Apps have to run at same patch level Hardware usually more reliable than OS.	- Complete isolation - Highly Flexible - CPU cycle granularity - 100's of OS images - High Reliability - Partition Load Mgr - Low Overhead Enterprise proven on MFs for decades.

Goal: Decrease Costs, Increase Utilization, Maximize ROI

OS and Virtualization

	Linux cost	Windows cost	VMware GSX	VMware ESX	Advanced O.P. Virt
2 way 710	RH $295 SUSE $495	n/a	n/a	n/a	$1372
4 way 720	RH $595 SUSE $1095	n/a	n/a	n/a	$1744
2 way Dell "intel"	RH $899 SUSE na	$3295 EE $799 SE	$1694	$4688	n/a
4 way Dell "intel	RH $899 SUSE na	$3295 EE $799 SE	$3388	$9375	n/a

AIX cost per CPU: p505: $150 -- p505, p510, p520; $385 – p550

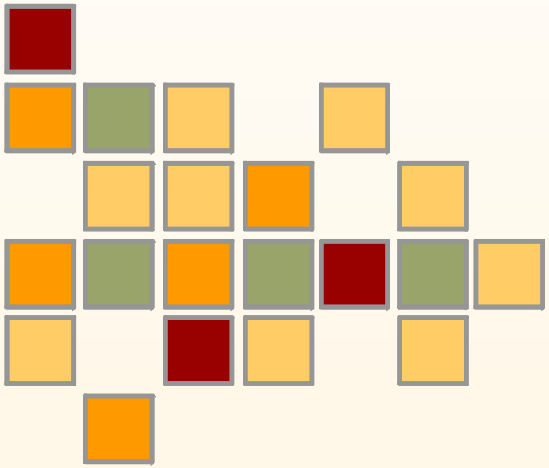

AIX and POWER Roadmaps

"a future you can count on"

AIX Release Plan

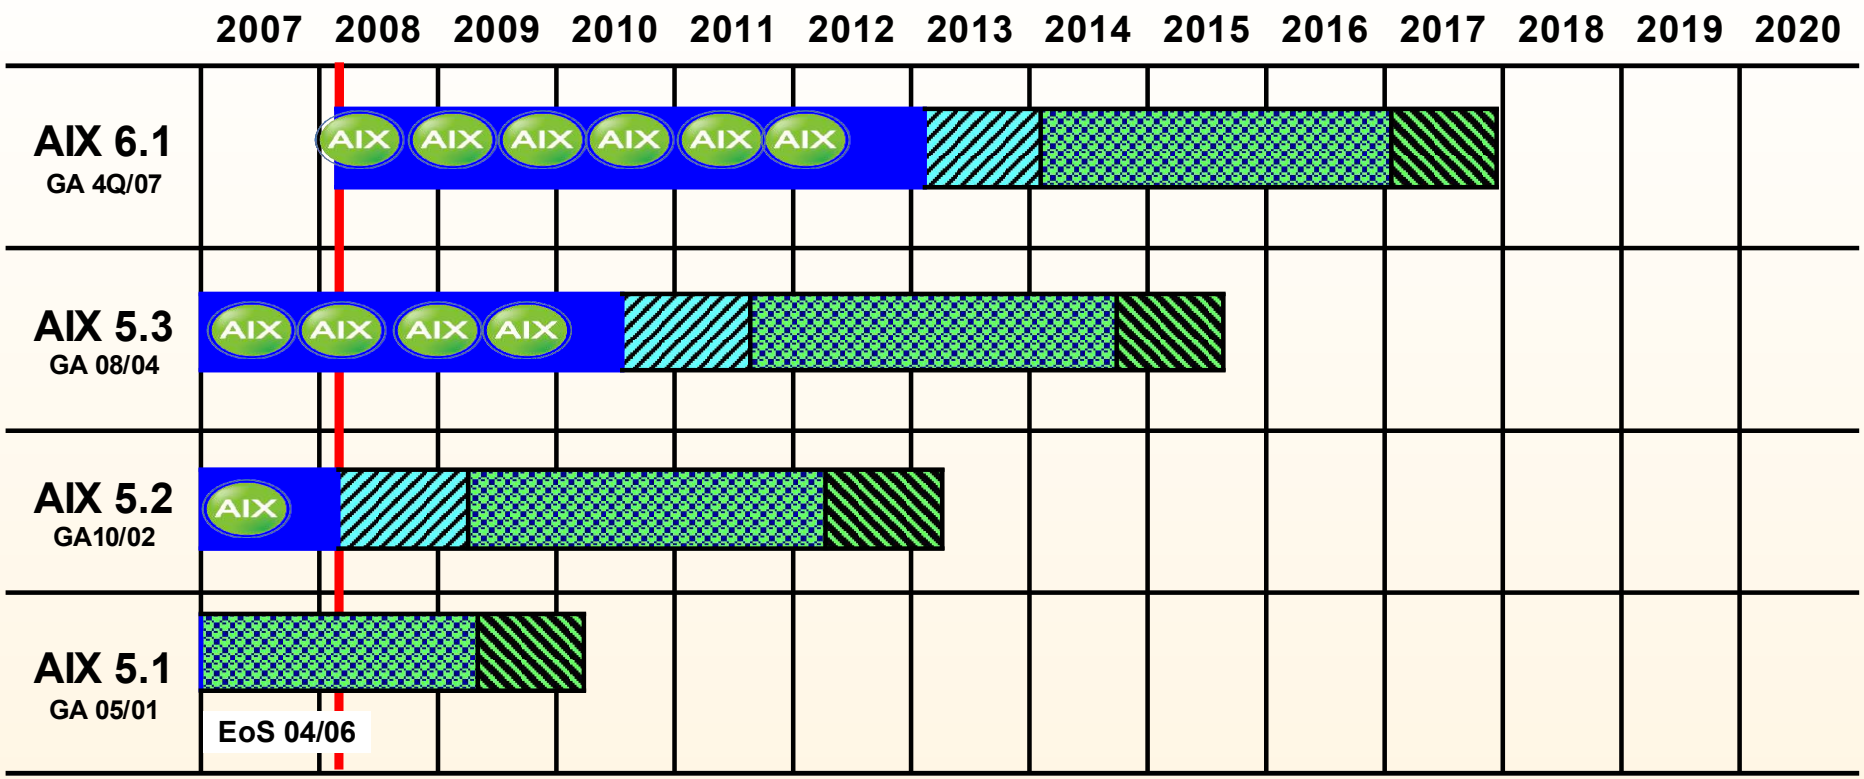

*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

POWER Processor Roadmap

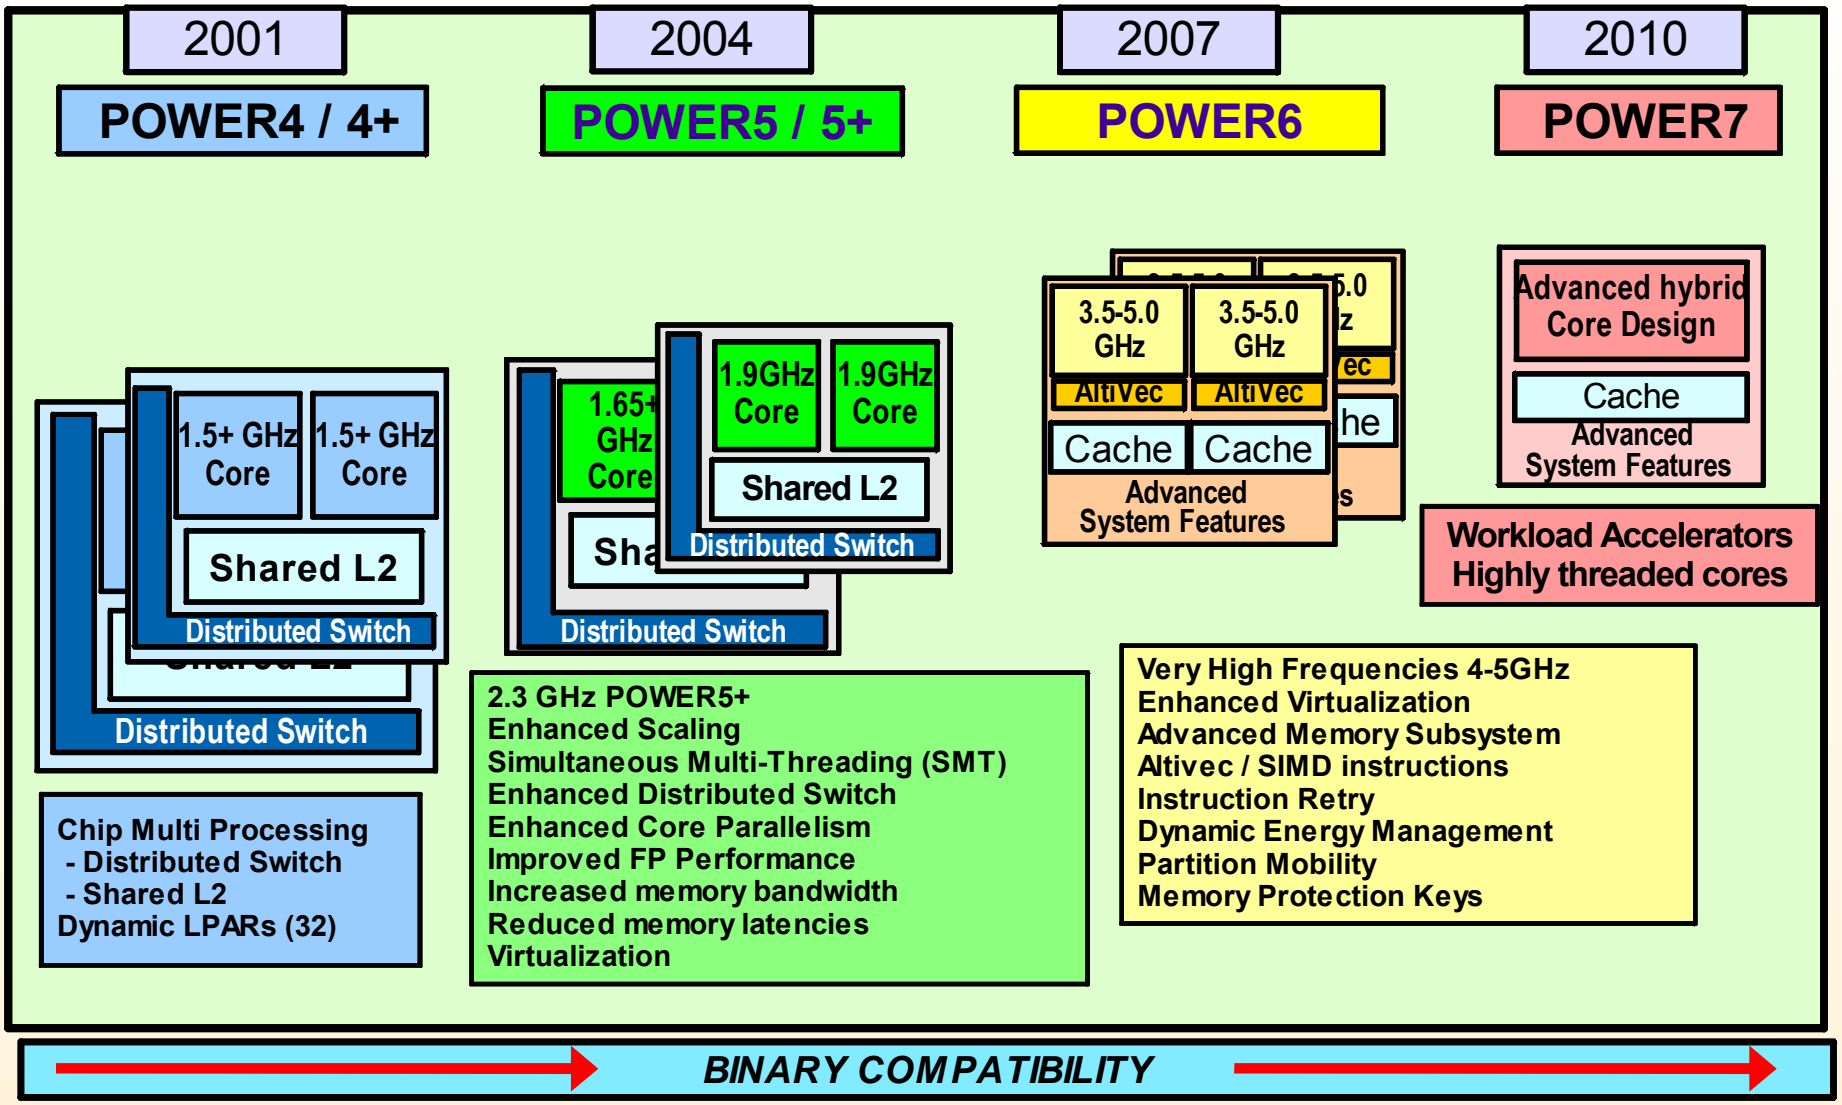

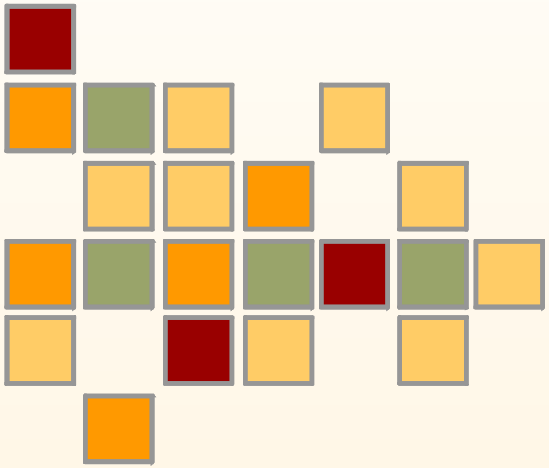

Oracle RAC Deployments and Scalability

Joint Development - The IBM/Oracle Development team

- **30 IBM engineers on-site at Oracle (primarily AIX), approximately 500 Servers on loan**
- **On-site team's mission :**
  - Advanced Product Development
    - Future joint offerings (Oracle on AIX 5.3, for example)
  - Performance testing
    - Generic and IBM-specific Oracle product improvements
  - Technical assistance and platform-specific training to Oracle
  - On-site team assists in tough debugging and critical customer situations

pSeries with AIX is a"concurrent development platform"at Oracle

- **Concurrent development**
  - Simultaneous release 9i, 10g and 11i
  - 8i, 9i, 10g and 11i all available on AIX 5.3 at GA
- **Daily builds on AIX during development cycle**
  - For Oracle Database Servers
  - For Oracle E-Business Suite
  - For Oracle Application Servers
- **IBM pSeries/AIX currently utilized as a development platform for Oracle RAC technology**

IBM AND ORACLE®

"*Although many assume that the fiercely competitive Oracle and IBM would make strange bedfellows, in fact, Oracle's strongest development relationship is with IBM, Perkins reported. "Our engineers are happiest working with IBM because they're interested in making things work," he said. About 30 IBM engineers work at Oracle helping Oracle developers optimize the software for IBM's hardware, adding functions such as dynamic tuning of processors, performance analysis and testing, and adjustments to cache size that benefit all of Oracle's customers.*"

-- Barry Perkins,
Oracle Vice President Global Strategy and Solutions.
(http://www.as400network.com/news/nwn/story.cfm?ID=19433)

Understand RAC's Advantages / Limitations

About 1,000 production (9i and 10g) RAC known implementations worldwide out of an estimated 250,000 Oracle installs – less than 1%

- Overwhelming majority of RAC installs are 2-node and 3-node clusters
- VERY FEW implementations beyond 4-node clusters.

One published industry benchmark result involving a RAC cluster of more than 4 nodes.

- TPC-C result for a 16-node cluster of HP Integrity 4-way servers
- p595 64-way result was nearly 3 times the tpmC of the HP RAC result (with the same number of processors) and showed similar $/tmpC numbers
- Compared to non-clustered Itanium results, the RAC cluster exhibited rather poor scale-out efficiency

Tangible RAC advantages depending on particular customer requirements

- Faster failover in the event of a server or software failure on one node
- Ability to scale capacity beyond the capacity of a single SMP

From a TCO standpoint (as opposed to hardware acquisition costs only), RAC VERY RARELY SAVES MONEY due to increased software licensing costs, increased configuration complexity (SAN switches, Network switches, adapter cards...) and increased administration/management complexity.

RAC Scalability

"Perfect" RAC scalability occurs when:

- Data structures are regular and evenly divisible
- Data can be partitioned equally

Then:

- RAC will scale very well
- incremental 0.8 – 0.9
- shown to scale to 16 nodes and beyond

"Real" RAC scalability occurs when:

- Data structures are irregular and not evenly divisible
- Data can't be partitioned equally

Then:

- RAC will not scale well
- 2nd node: incremental 0.6 – 0.7
- 3rd node: incremental 0.5 – 0.6
- 4th node: incremental 0.4 – 0.5
- 5th node: incremental zero - to - negative

✓Oracle RAC has scalability issues because of its distributed lock manager architecture. The overhead grows larger and larger as more servers are added to the cluster. The result is that customers get less and less benefit from each server added. Thus software costs can be much higher in a multi-node configuration.

✓We Recommend using a 3-node RAC Configuration with vertical scalability.

Scalability Continued

System	Cores	Avail.	$/tpmC	Database	OS	tpmC	tpmC/Core
HP Superdome (1.5 GHz Itanium 2)	64	04/14/04	$8.33	Oracle 10g	HP-UX 11.i	1,008,144	15,752
HP rx5670 (1.5 GHz Itanium 2)	4	12/31/03	$7.25	Oracle 10g	HP-UX 11.i	131,639	32,909
IBM p5-595 (1.9 GHz POWER5)	64	05/14/05	$5.07	IBM DB2 v8.1	AIX 5L V5.3	3,210,540	50,164
IBM p5-570 (1.9 GHz POWER5)	4	10/17/05	$3.93	Oracle 10g	AIX 5L V5.3	203,439	50,859

Note the issues HP has (any NUMA architecture has same problem) scaling. The Performance per CPU decreases by a factor of 2 when you get to 64 CPUs. Cache Coherency and the speed of the interconnect fabric (system bus in this case) limit Scalability. Just think how this problem is magnified when using many small servers Over a Ethernet interconnect. Notice that IBM performance per CPU stays constant.

IBM offers a clear system performance advantage because we are the only server Company to offer true SMP servers all the way to 64 CPUs. **(The 900M bet!)**

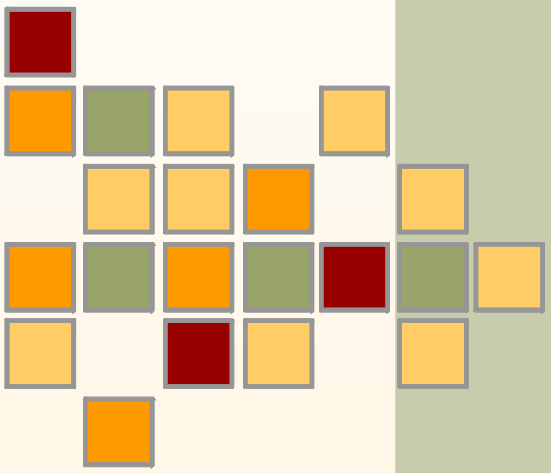

Summary

More customers are choosing System p

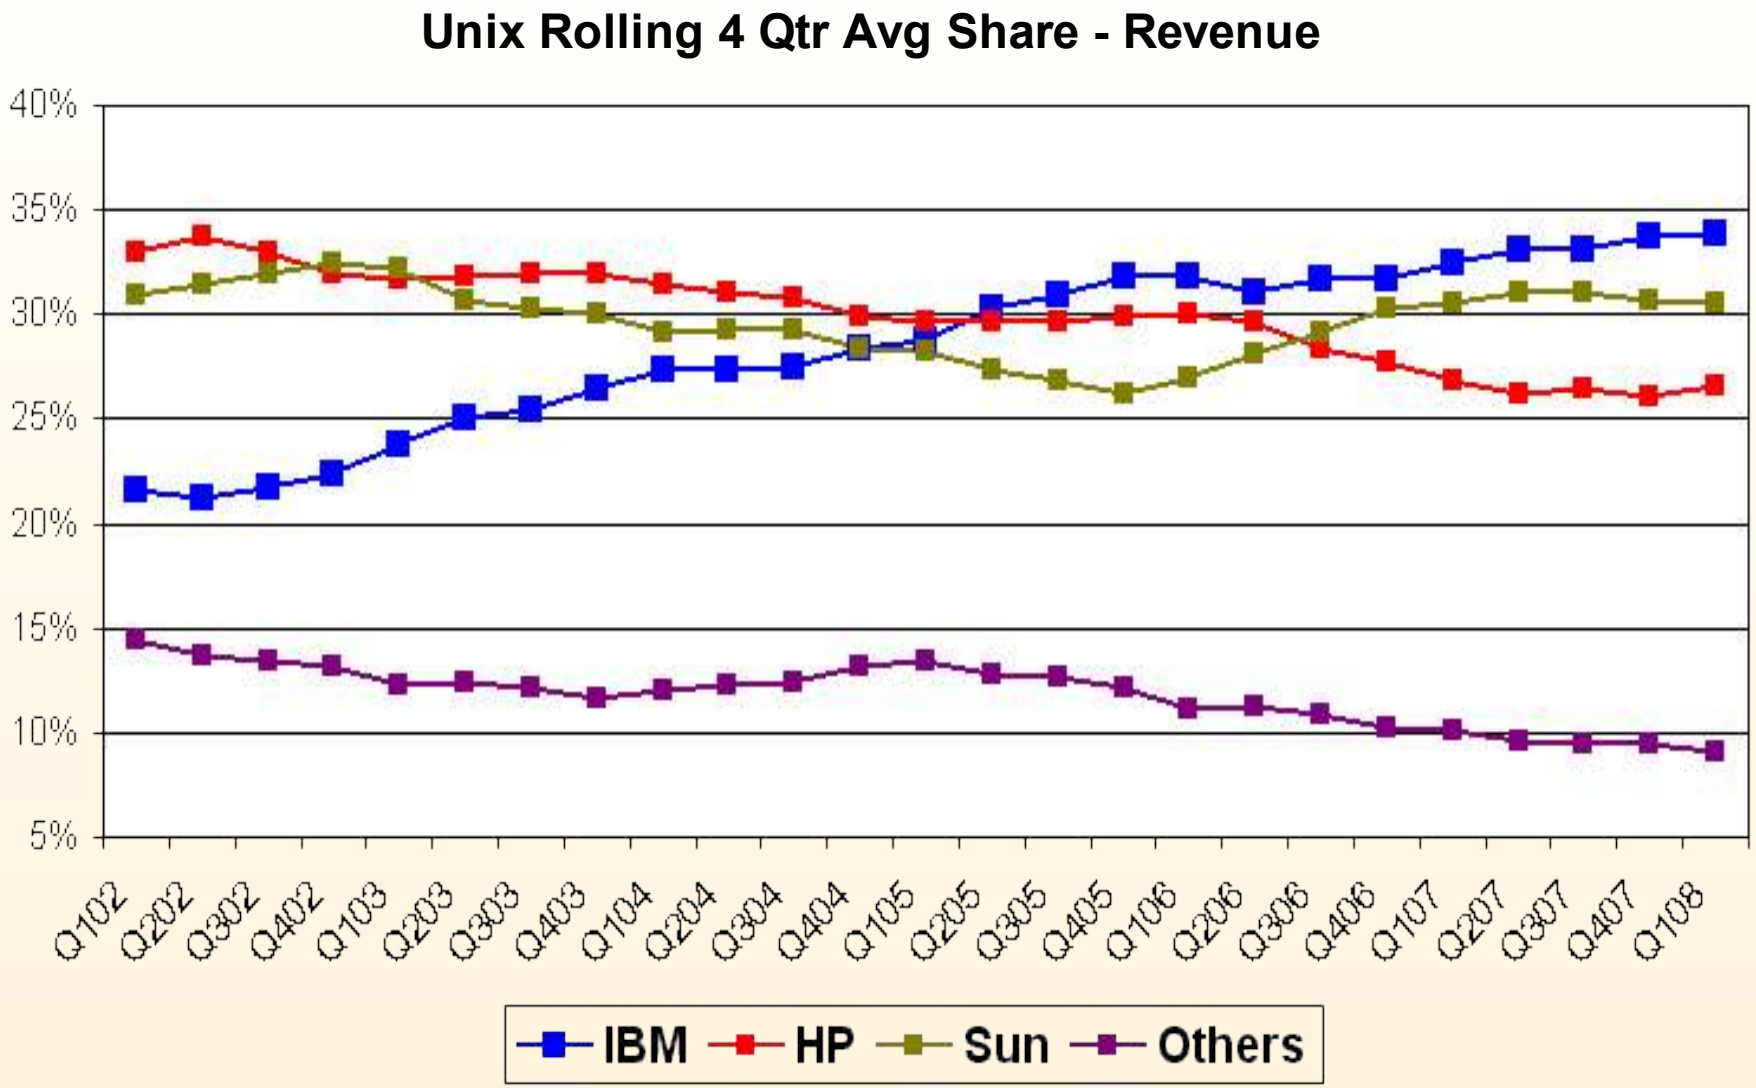

Source: IDC Server Tracker Q108 Release, May 2008

Server Acquisition vs. Total Cost

Commodity hardware may reduce server hardware acquisition cost

- Many Multiple low-end servers vs. 1 or 2 high-end servers

However, Server hardware savings may be offset by:

- Increased software license and support costs
- Increased cost of network and storage adapters/switches
- Increased administration complexity
  - Multiple OS images to maintain
  - Shared disk & Network interconnects
  - Failover/Fallback strategies
- Loss of flexibility/functionality
  - LPAR & DLPAR support, Capacity Upgrade on Demand etc…
  - Workload Consolidation opportunity (WLM)
  - Higher System utilization and flexibility with larger CPU counts
  - RAS features

If price/performance, reliability, scalability, and flexibility are important to you – then I believe IBM System p6 deserves consideration.

Thank you for your time….

When you succeed … we succeed… thanks for the feedback.

The Real Answer: (42)