IBM System p

The POWER of 6

Introduction to System p
and System p Virtualization

IBM System Power5 - Power6 & AIX 6

Michel Marion
Certified Consulting IT Specialist
Steve Horwood
System P sales Specialist

July 2007

mmarion@ca.ibm.com

IBM Systems

Table of Contents

DFO Discussion Items

- Licence/Support Cost
- Application Software Availability
- Viruses/Spyware & Security
- Bug Fixes/Scheduled Maintenance Window
- Hardware the O/S runs on
- Memory Management
- CPU Limits
- Clustering
- Virtualization Support
- Scalability
- Multiple User Support & Privileges
- Industry Trends
- Interoperability
- Ease of Use
- Administration
- Training
- Oracle Product Features supported on each platform
- Fiber Channel Support

IBM Dominates the UNIX NMSO
Scale up. Scale out. Scale within.

- *1-core to 64-core solutions*
- *AIX 5L and Linux operating systems*
- *Using IBM POWER: Dual-Core ,Quad-Core, & MCM technologies*

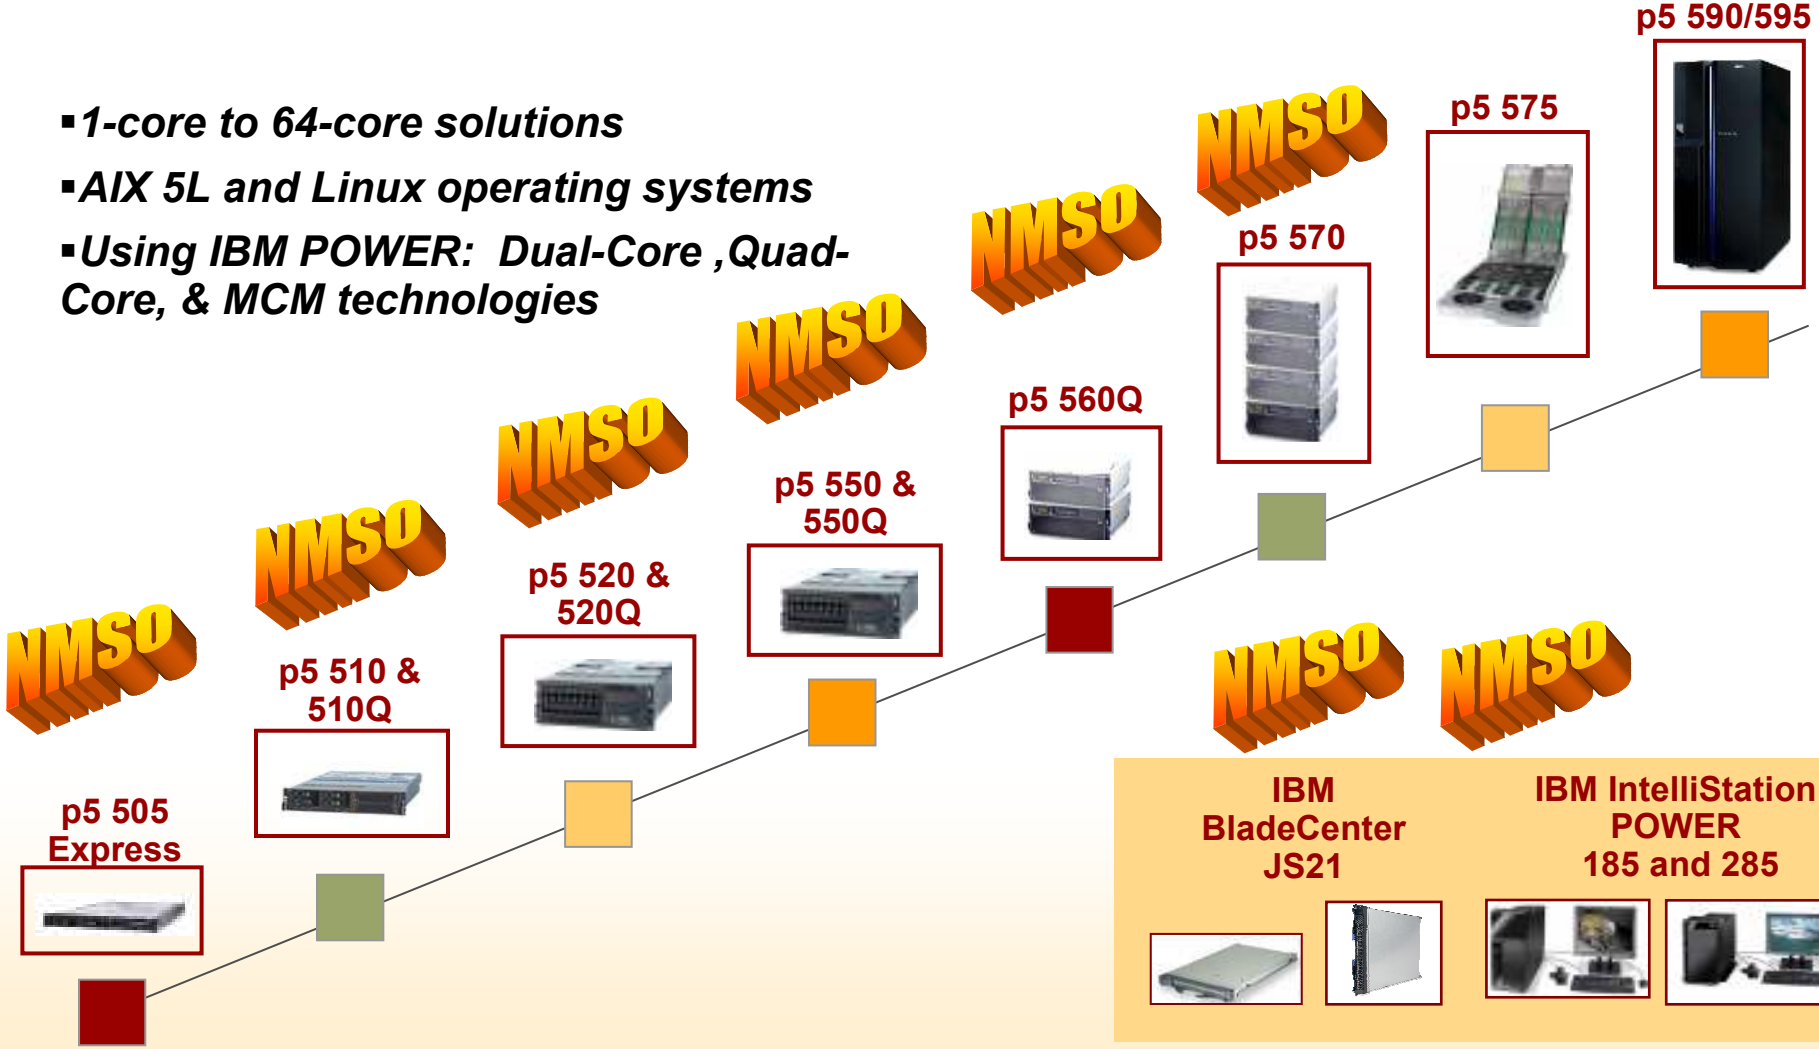

UNIX Rolling Four Quarter Revenue Share

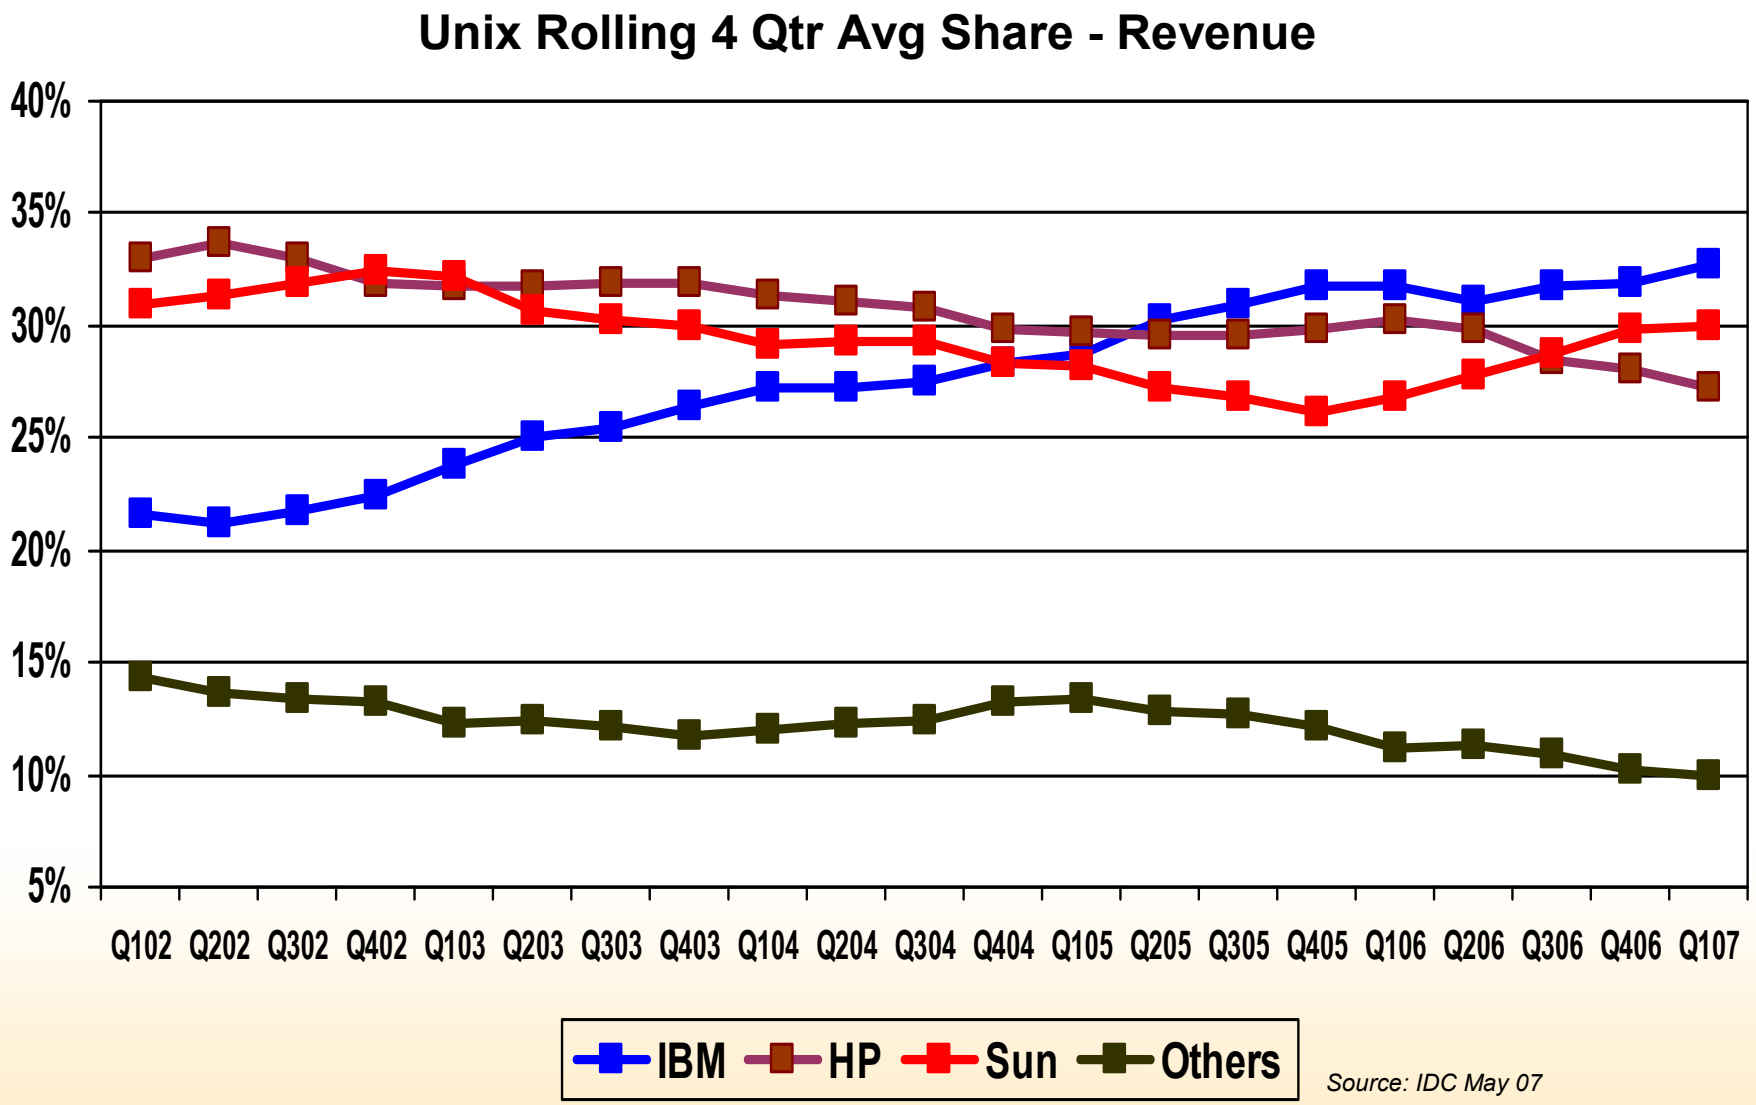

Source: IDC May 07

Only UNIX platform to grow over last five years

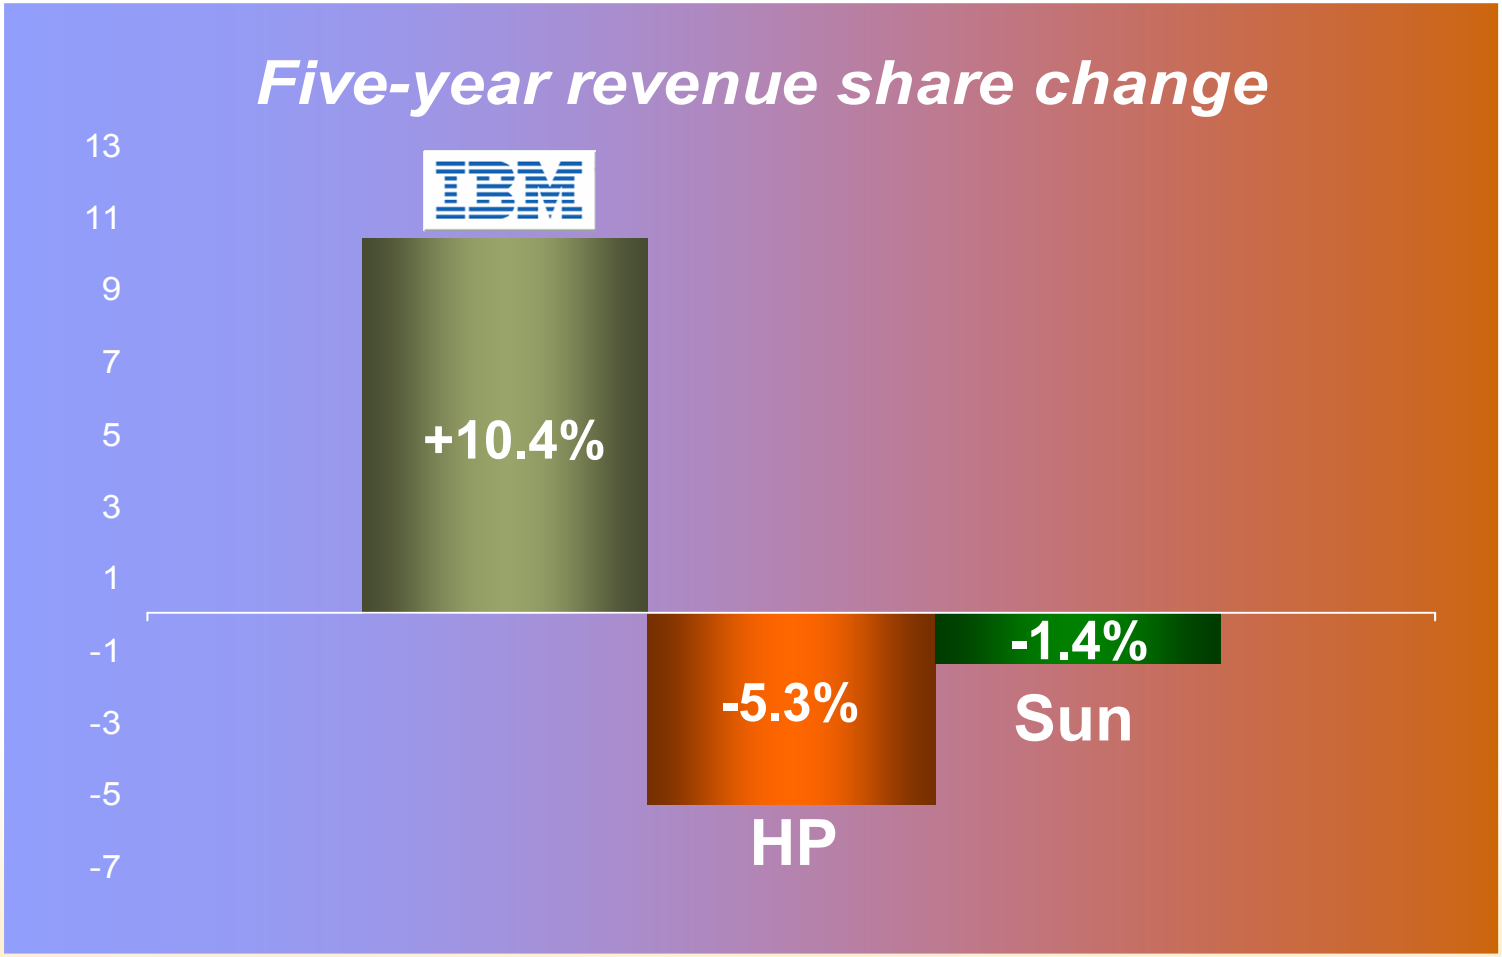

Source: IDC Server Tracker Q406 and FY06 Server Tracker, 02/24/07, rolling four quarter average

Gartner Magic Quadrant for Transaction Processing

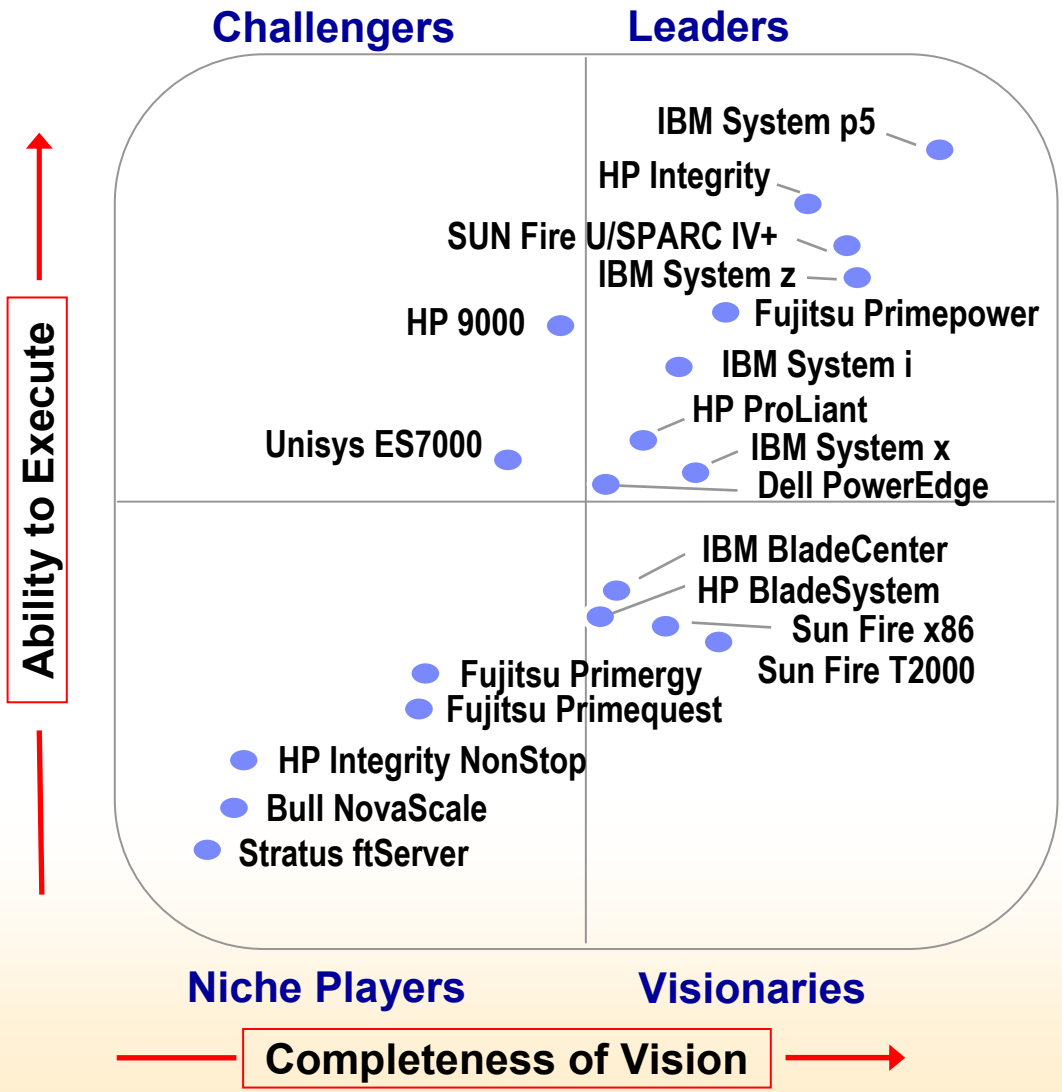

Gartner, Magic Quadrant For Enterprise Servers 2006, August 10, 2006.; Philip Dawson, Jonathon Hardcastle, Andrew Butler, Donald Feinberg, Paul McGuckin.
ID Number: G00139934

IBM Leadership

- **Technology Leadership**
  - #1 in Patents for 14 consecutive years
  - Copper, Dual core, SIO, etc.
- **Performance**
  - tpc-c, SAP, Oracle Apps, etc. leadership
- **Reliability, Availability, & Serviceability**
  - First Failure Data Capture
  - Chill Kill / Bit Steering memory
- **Roadmap**
  - Processor, Server, Virtualization, etc
  - Deliver on what we promise
- **Processor Technology**
  - POWER Roadmap plans through 2013
  - DARPA Award
  - Gaming Industry leadership
- **Virtualization**
  - Full generation ahead
- **POWER6 announced last month**
  - Pulling further ahead

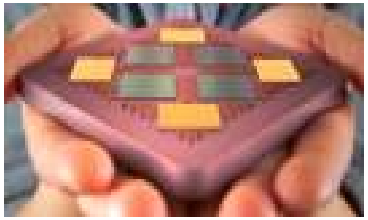

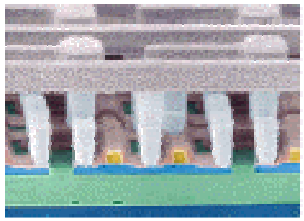

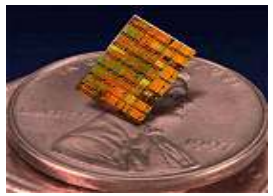

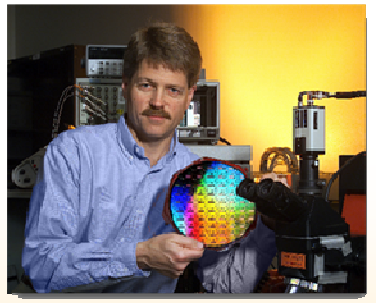

System p Servers

- Designed for mission-critical environments
- Great price/performance
- Broad range of growth options
- Outstanding RAS

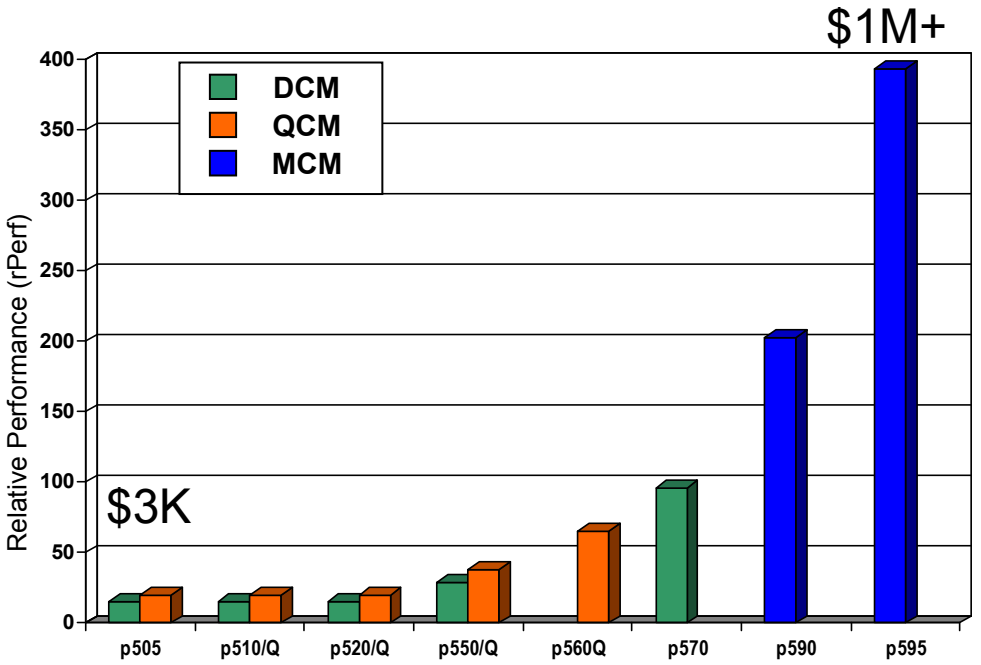

p5-595

p5-590

p5-575

p6-570

p5-560Q

p5-550 550Q

p5-520 520Q

p5-510 510Q

p5-505 505Q

Half of every dollar today is spent on energy for hardware

This is expected to increase by 54% over the next four years

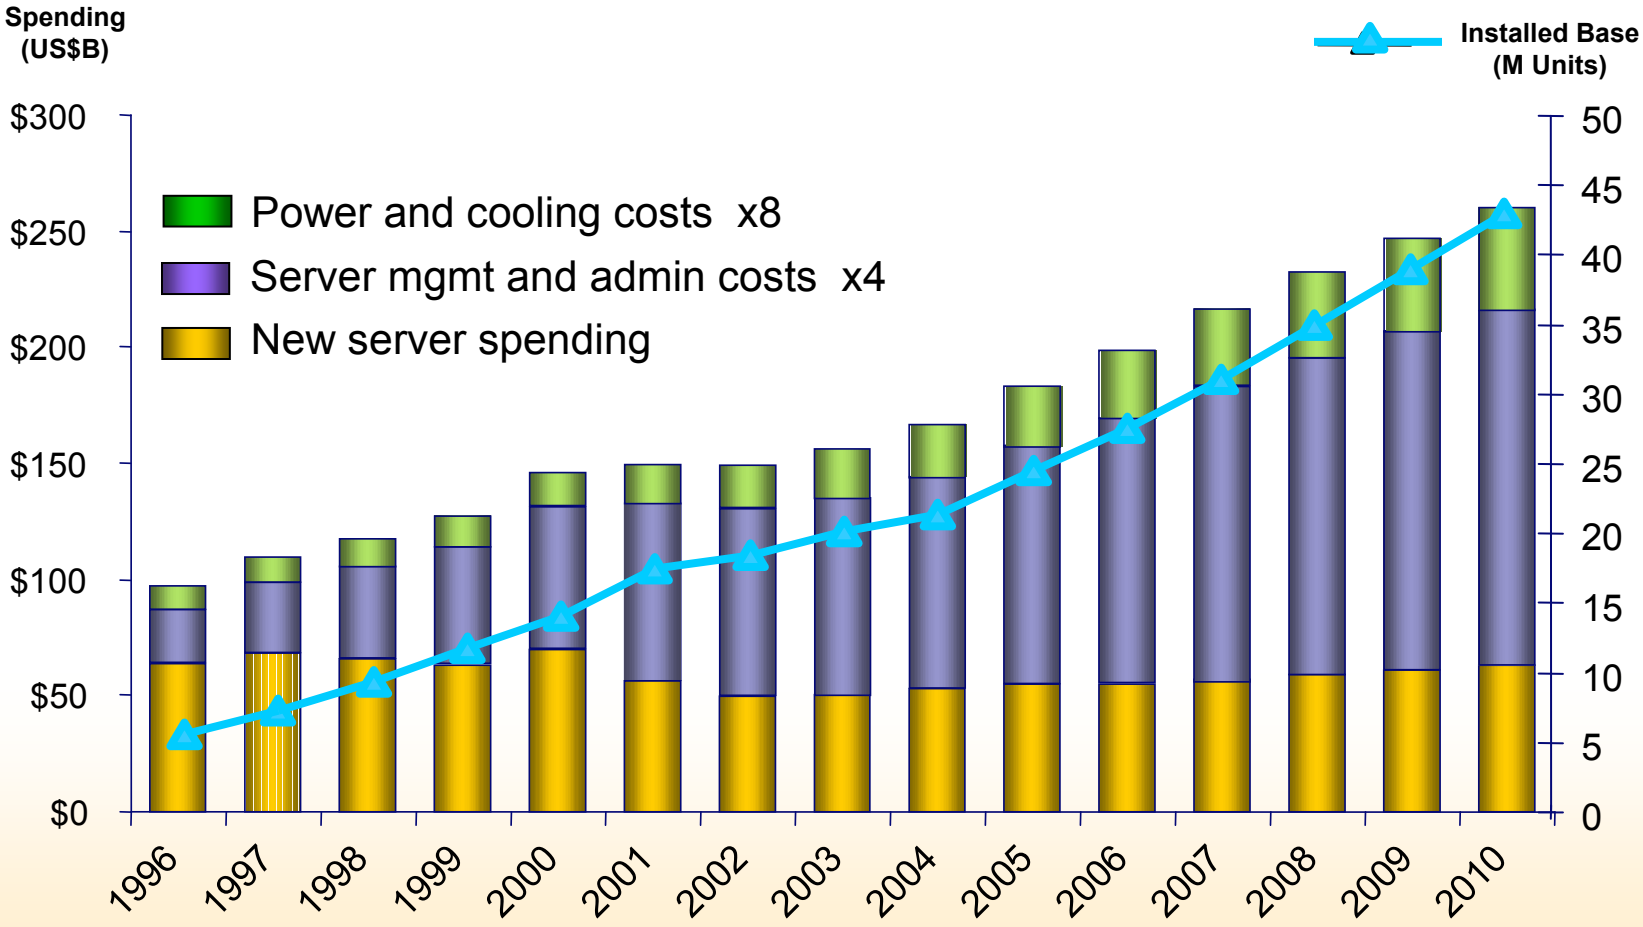

Source: IDC, Virtualization 2.0: The Next Phase in Customer Adoption, Doc #204904, Dec 2006

Source: IDC, Virtualization 2.0: The Next Phase in Customer Adoption, Doc #204904, Dec 2006

THE POWER OF SIX

IBM System p 570 . . .with POWER6™ technology!

More than twice the performance* and modular flexibility

System p Virtualization. . .with Live Partition Mobility*!

Extending the most complete virtualization offering for UNIX and Linux

IBM AIX 6 . . . with binary compatibility!

The next evolution of UNIX with new workload partitions, manageability and security

** All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.
Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

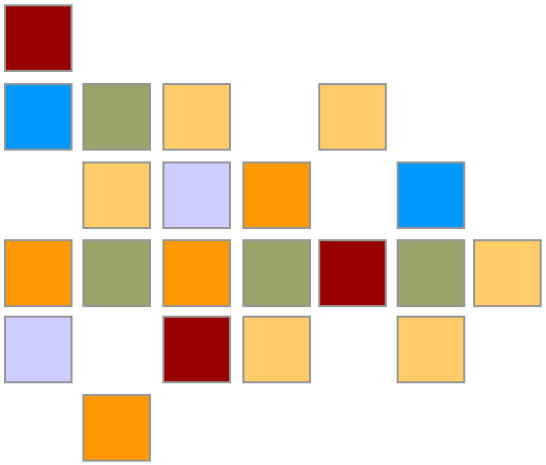

Processors

The IBM Power Architecture™ Roadmap*

System p will combine these POWER™ processor engines into systems to deliver higher-levels of customer value.

Throughput / Transaction Optimized


POWER4


PowerPC™ 970MP


POWER5 / 5+


POWER6*


POWER7*

Specialty Engines

PowerPC 405



PowerPC 440

Cell
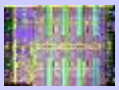

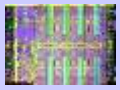

Cell +*
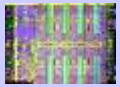

2002 2004 2006 2008 2010

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

POWER Technology

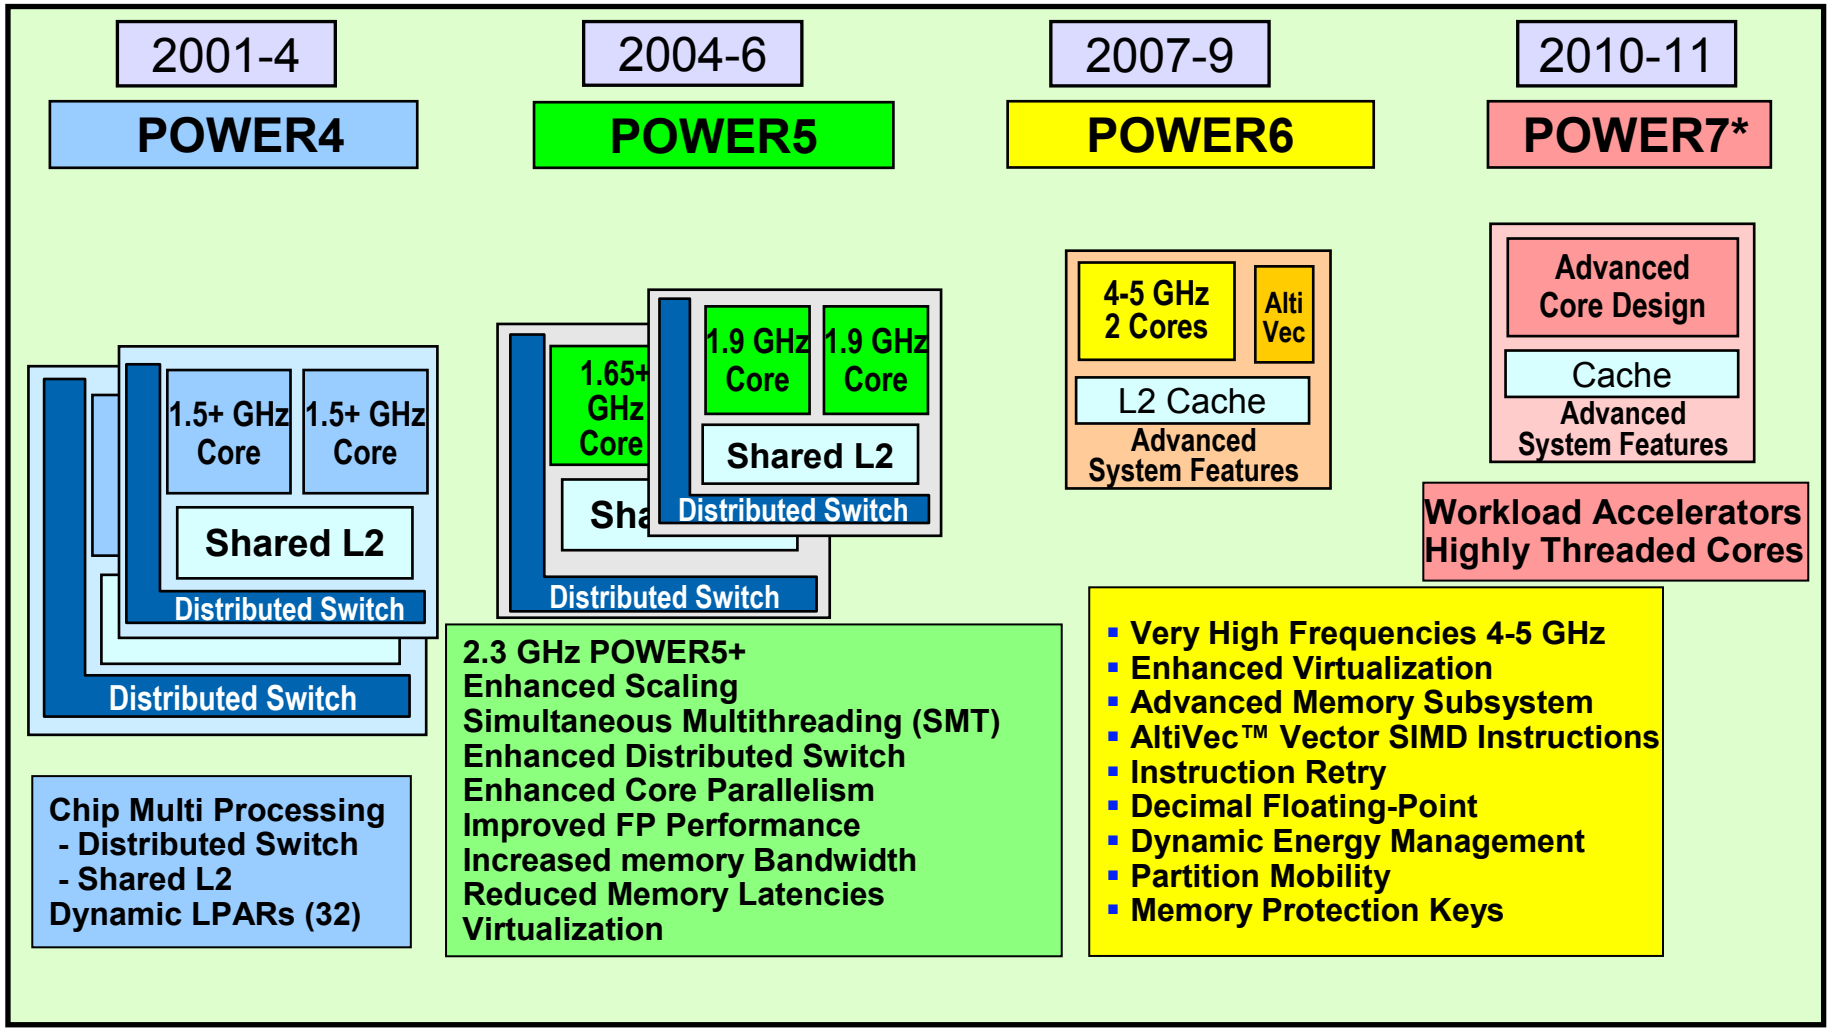

BINARY COMPATIBILITY

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

IBM POWER Technology

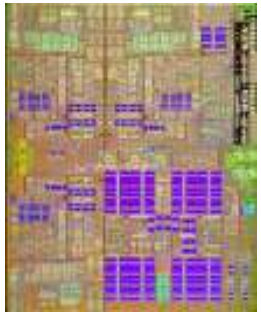

Microsoft Xbox 360®

Sony PlayStation®3

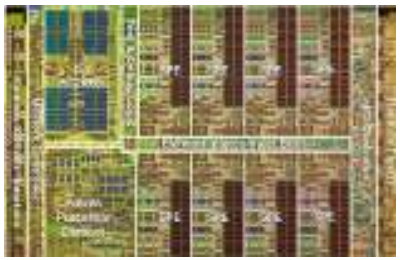

Nintendo Wii®

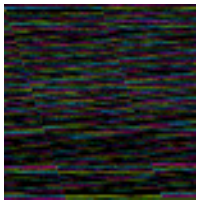

Processor History - 2001 - 2007

Dual-Core
1 GHz
Distribute Switch

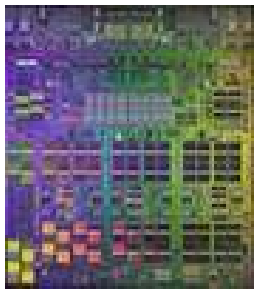

POWER4
414 mm²
1.1 – 1.3 GHz

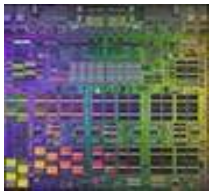

POWER4+
267 mm²
1.5 – 1.9 GHz

Multithreading
Memory Cntrl on Chip

POWER5
389 mm²
1.65 – 1.9 GHz

POWER5+
245 mm²
1.9 – 2.3 GHz

Enhanced Multithreading
Memory Cntrl on Chip
>4 GHz

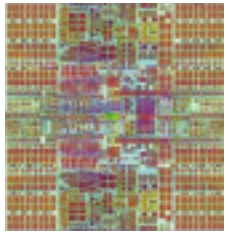

POWER6
341 mm²
3.5 – 5.0 GHz

2001 2002 2003 2004 2005 2006 2007 2008

Power Architecture Continues to Deliver for Performance

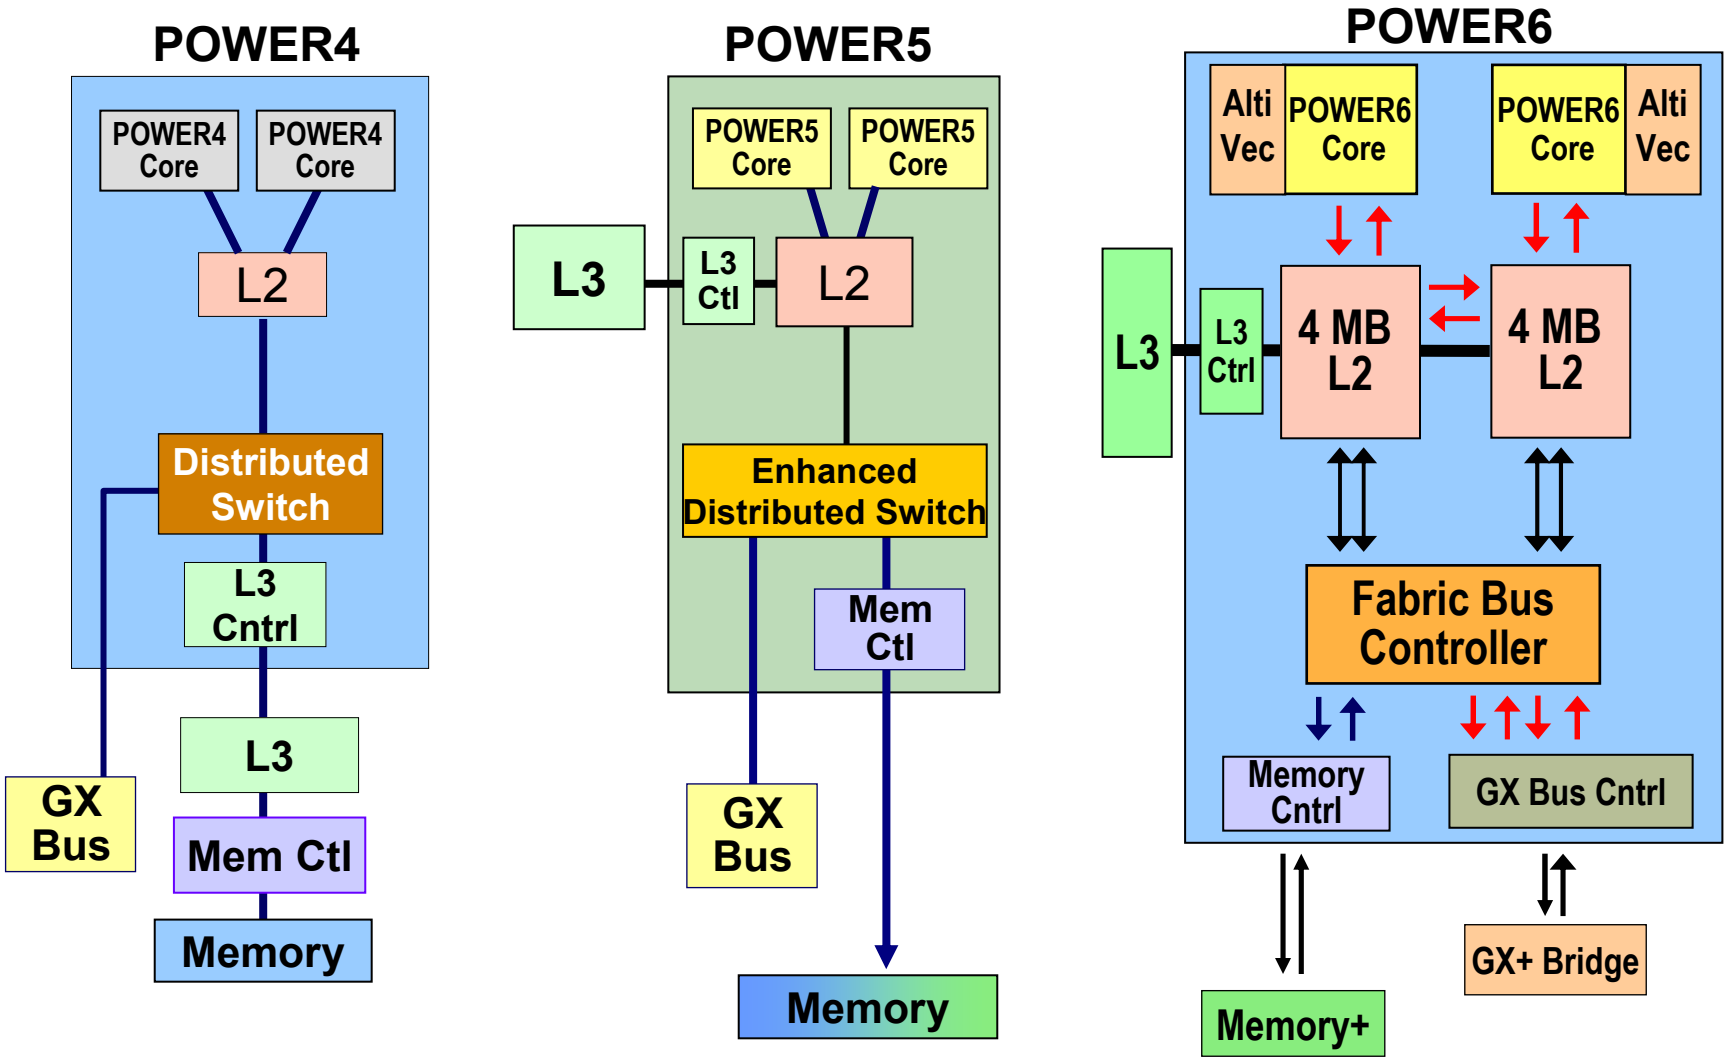

POWER6 Architecture

POWER Design	3.5 – 4.7 GHz	>750M transistors	.065 micron

POWER6 Characteristics

Ultra-high frequency Dual-Core chip: > 3.5 GHz

- 7-way superscalar, 2-way SMT core
  - Up to five instructions for one thread, up to two for other
- Nine execution units
  - 2LS, 2FP, 2FX, 1VMX, 1DP, BCX
- 790M transistors, 341 mm^2 die
- Up to 64-core SMP systems
- 2 x 4 MB on-chip L2 – point of coherency
- On-chip L3 directory and controller
- Two memory controllers on-chip

Technology

- CMOS 65nm lithography, SOI Cu

High-speed elastic bus interface at 2:1 frequency

Full error checking and recovery

Dynamic power saving

- Advanced clock gating

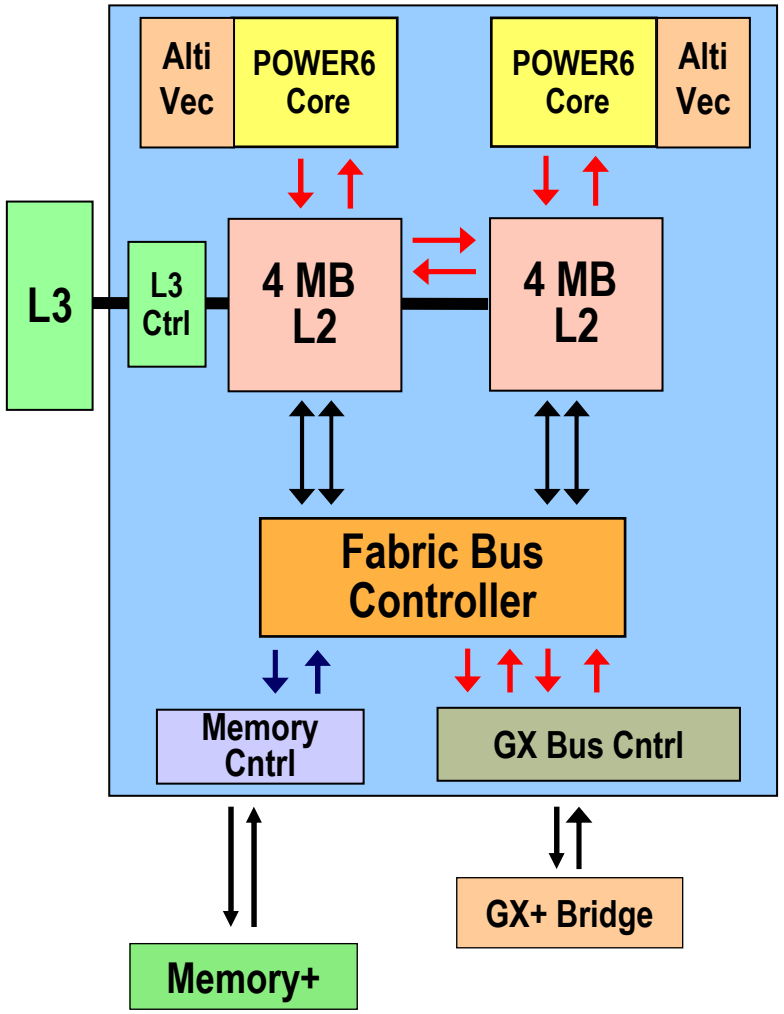

POWER5+ and POWER6 Hierarchy

	POWER5+	POWER6
L1 Cache		
ICache capacity, associativity	64 KB, 2-way	64 KB, 4-way
DCache capacity, associativity	32 KB, 4-way	**64 KB, 8-way**
L2 Cache		
Capacity, line size	1.9 MB, 128 B line	**2 x 4 MB, 128 B line**
Associativity, replacement	10-way, LRU	8-way, LRU
Off-chip L3 Cache		
Capacity, line size	36 MB, 256 B line	32 MB, 128 B line
Associativity, replacement	12-way, LRU	**16-way, LRU**
Memory	2 TB maximum	4 TB maximum
Memory bus	2x DRAM frequency	**4x DRAM frequency**

Processor Design

	POWER5+	POWER6
Style	General out-of-order execution	Mostly in-order with special case out-of-order execution
Units	2FX, 2LS, 2FP, 1BR, 1CR	2FX, 2LS, 2FP, 1BXU, 1DP,1VMX
Threading	Two SMT threads Alternate ifetch Alternate dispatch (up to five instructions)	Two SMT threads Priority-based dispatch **Simultaneous dispatch from two threads (up to seven instructions)**

POWER6 AltiVec Vector Technology

Dramatic application performance gains

SIMD (Single Instruction, Multiple Data) Extension to PowerPC Architecture™, jointly developed by Apple, Motorola, IBM

Targets High Performance Computing and Deep Computing applications

Benefit to ISVs / clients:

Provides highly parallel operations

Dramatically better performance for highly “vectorized” code

Development / test environment:

Current support: IBM BladeCenter® JS21 or IBM IntelliStation® POWER™ 185 Express

- Supported by AIX® and Linux releases

IBM XL C/C++ Enterprise Edition V8.0 for AIX (October 2005) provides

- Support for the AltiVec instruction set
- Support for the AltiVec programming model and APIs

IBM XL Fortran Enterprise Edition V10.1 for AIX (October 2005) can

- Automatically enable SIMD vectorization at higher levels of optimization

Additional compiler support for AltiVec™ vectorization extensions will be available in XL C/C++ V9.0, with Automatic SIMD vectorization

Redbook: http://www.redbooks.ibm.com/abstracts/redp3890.html

Decimal Floating-point

Description: New / more precise mechanism for Decimal Arithmetic
Markets: DB servers, Financial Sector Applications, etc.

Binary Floating-point problem….

- Add 5% sales tax to a $ 0.70 telephone call, rounded to the nearest cent
- 1.05 x 0.70 using binary double is exactly 0.73499999999999998667732370449812151491641998291015625
- (should have been 0.735) rounds to $ 0.73, instead of $ 0.74

Decimal Floating-point benefit….

- True decimal processing; .735

Benefits

- Performance and Accuracy

Requirement:

- AIX 5.3 TL06

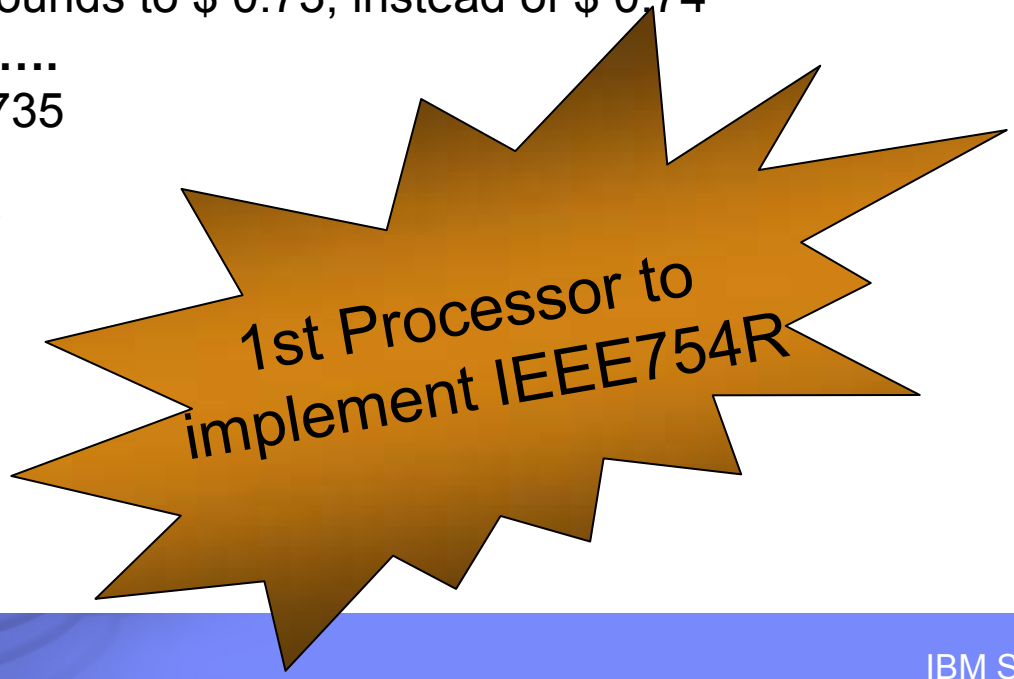

Processor Instruction Retry and Recovery

General: Hardware recovery from some non-predicted errors

Internal Checkpoint within each processor core

- Capable of preserving the state of processing operations

When certain unrecoverable faults are detected

- Previous checkpoint can be reloaded to retry operation from the last checkpoint
  - Recovers transparently from transient errors
- If retry is unsuccessful
  - Checkpoint can be reloaded into a different processor core
  - Spare processor may be used if available
  - Else Checkstop can be limited to just the processor w/error (partition fault isolation)

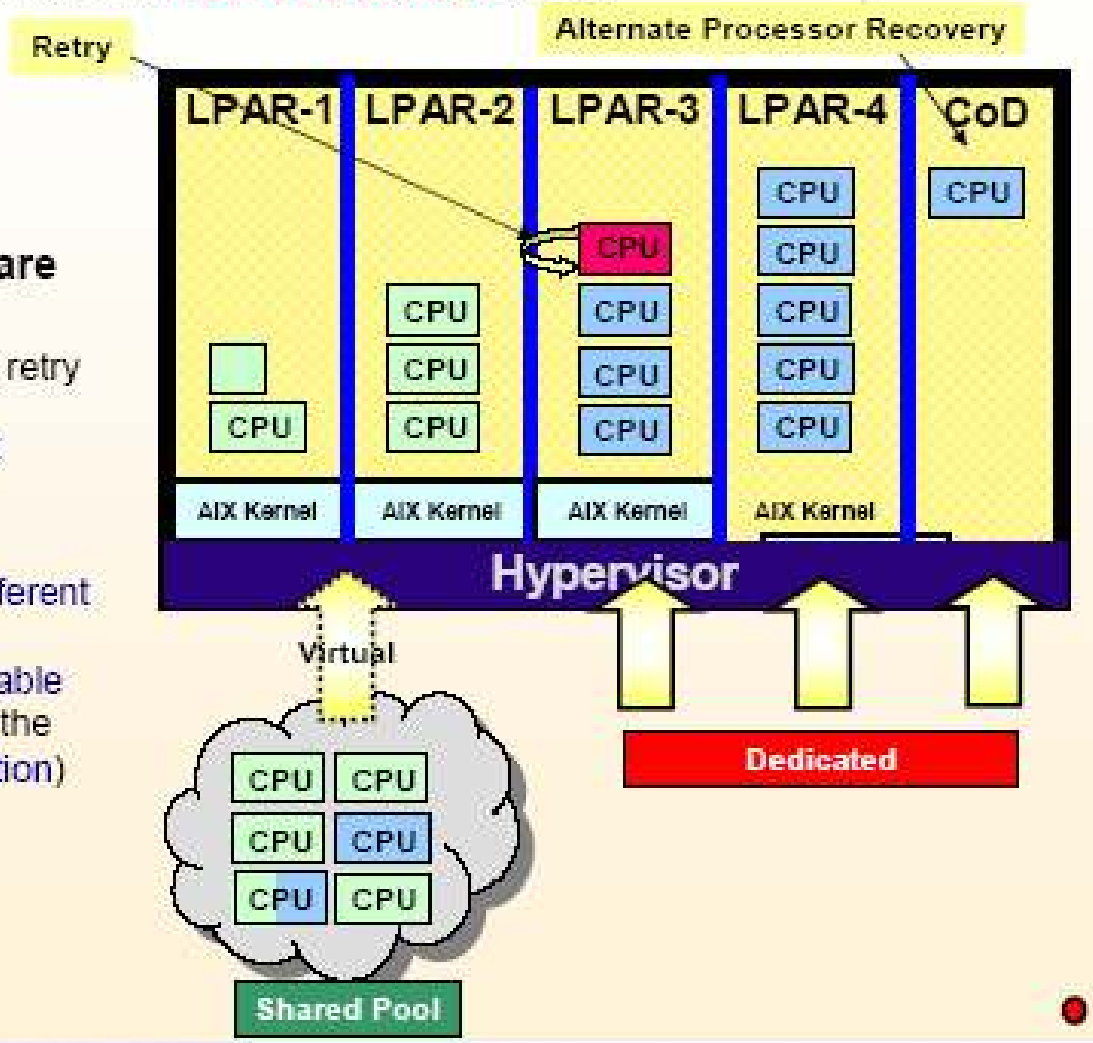

POWER6: Simultaneous Multithreading

POWER5 Simultaneous Multithreading

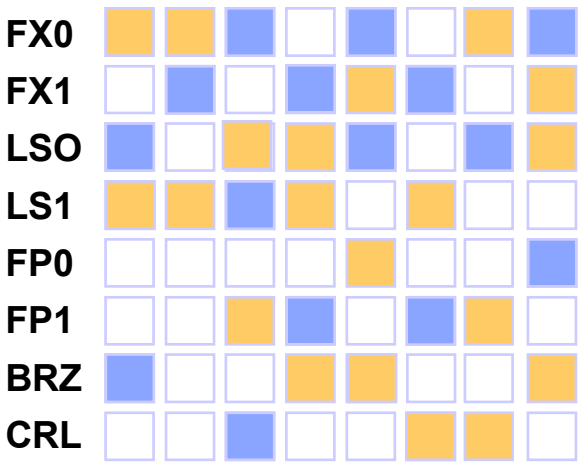

Appears as four CPUs per chip to the operating system (AIX V5.3 and Linux)

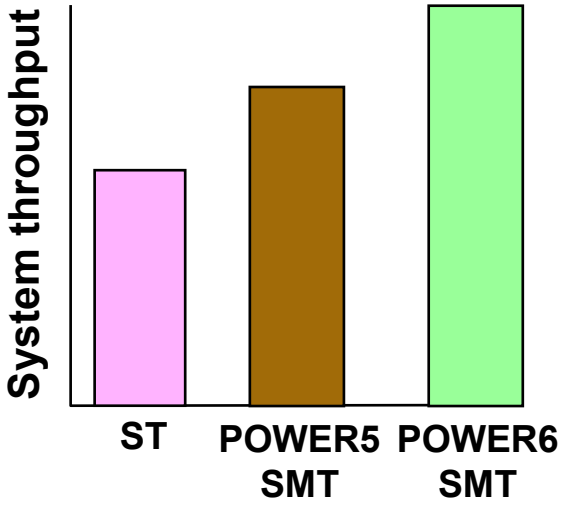

- Utilizes **unused execution** unit cycles
- **Reuse of existing transistors vs. performance from additional transistors**
- Presents symmetric multiprocessing (SMP) programming model to software
- Dispatch two threads per processor: *"It's like **doubling** the number of processors."*
- Net result:
  - ***Better performance***
  - ***Better processor utilization***

Oracle 9.2.0.4 OLTP performance gain from SMT

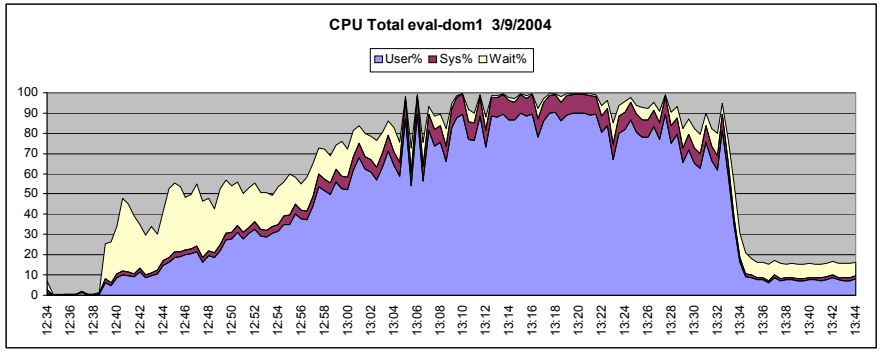

AIX 5.2 - 8-way, 64 GB p570

700 Users

180 Transactions per Second

< 0.05 sec response time

100% Sys Utilization

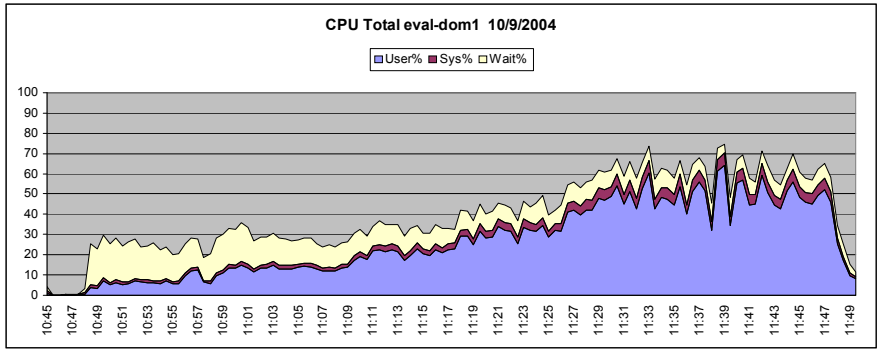

AIX 5.3 (SMT)

700 Users

~180 Transactions per Second

< 0.05 sec response time

60% Sys Utilization

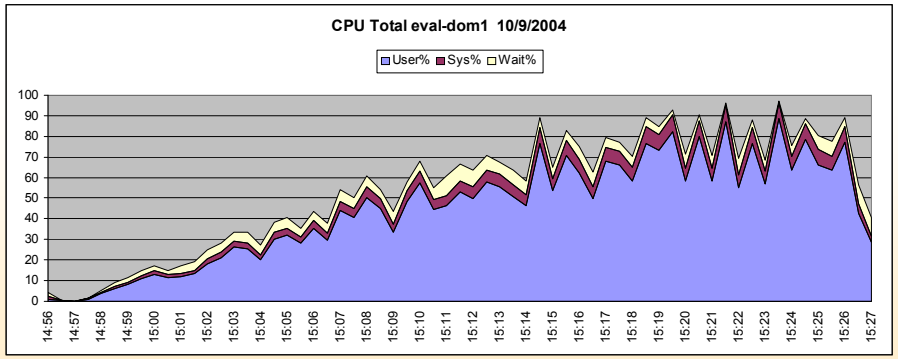

AIX 5.3 (SMT) – cached (low wait)

1000 Users

280 Transactions per Second

< 0.05 sec response time

80-90% Sys Utilization

Dual Stress Technology

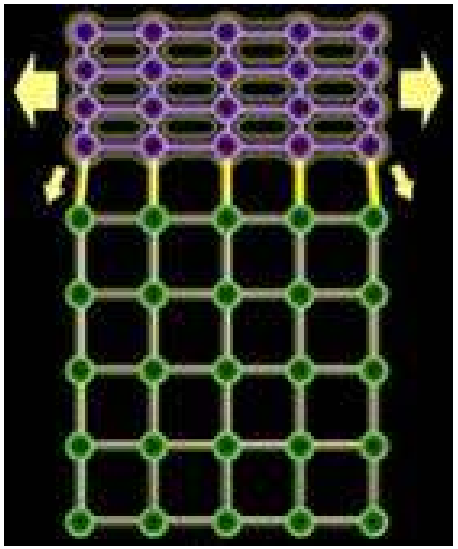

What is it?

- Incorporated into POWER5+ processor technology
- Strained Silicon on Silicon-on-Insulator technology
- Stretches and compresses transistors
  - Stress film technology
  - Provides more efficient flow of electrons
- No special materials required

Benefit

- Increase transistor speeds by up to 20% without increasing power consumption
- Reduce electric current leaks
- Greater performance without increasing power and heat

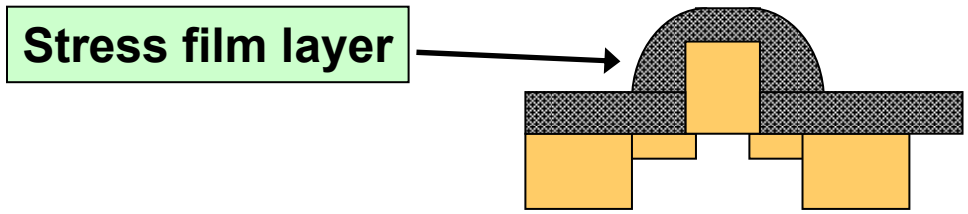

POWER5+

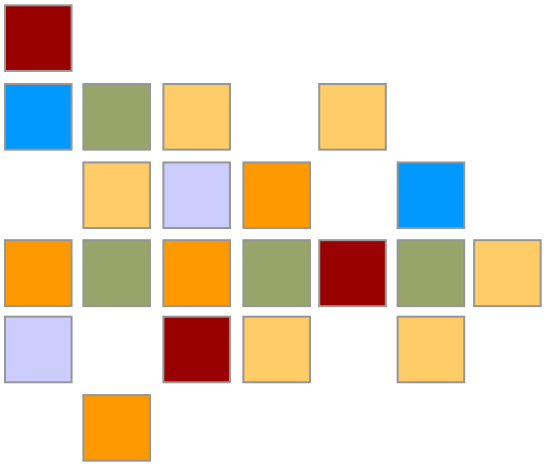

System p

POWER5+ packaging

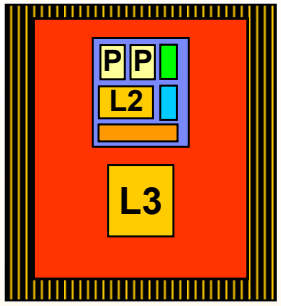

Dual-Core Module

POWER5+ Dual-Core chip + L3 cache chip
(505, 510, 520, 550, 570)
Two processor cores

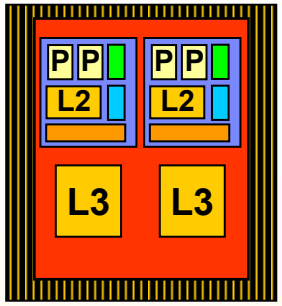

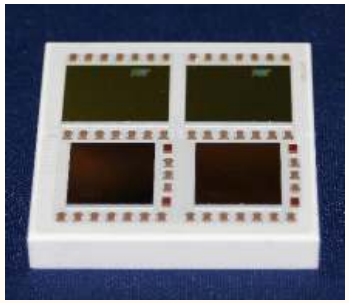

Quad-Core Module

Two POWER5+ Dual-Core chips + two L3 cache chips (505Q, 510Q, 520Q, 550Q, 560Q)
Four processor cores

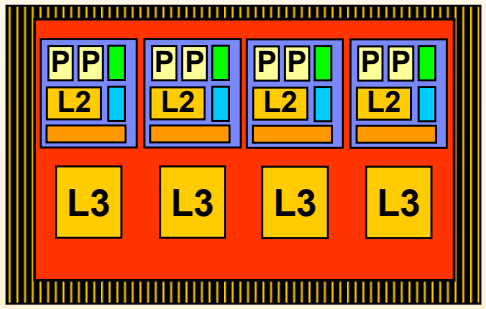

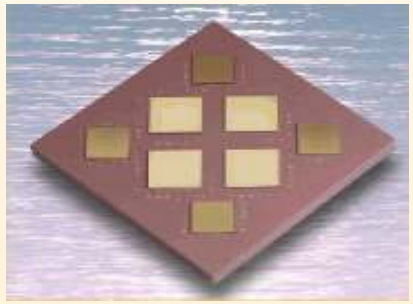

Multi-Chip Module

Four Dual-Core POWER5+ chips + four L3 cache chips
(590, 595)
Eight processor cores

IBM System p5 Express with POWER5+

	IBM System p5 185 Express	IBM System p5 505 Express	IBM System p5 510 Express	IBM System p5 520 Express	IBM System p5 550 Express
Footprint, Packaging	19-inch 5U Rack/Deskside	19-inch 1U Rack	19-inch 2U Rack	19-inch 4U Rack/Deskside	19-inch 4U Rack/Deskside
Processor	PowerPC® 970	POWER5+™	POWER5+	POWER5+	POWER5+
# of processors (# of cores)	1 , 2	1 , 2	1 , 2	1, 2	2, 4
GHz clock	2.5	1.9, 2.1	2.1	2.1	1.9, 2.1
GB memory	.5 to 8 DDR1	1 to 32 DDR2	1 to 32 DDR2	1 to 32 DDR2	1 to 64 DDR2
Internal storage	73GB – 1.2TB	73GB – 600GB	73GB – 1.2TB	73GB – 16.8TB*	73GB – 31.2TB*
Maximum rPerf	4.34	12.46	12.46	12.46	24.86
PCI-X slots GX bus slots	5 0	2 0	3 0	6 to 34 1	5 to 59 2
I/O drawers	N/A	N/A	N/A	4	8
Micro-partitions[1]	N/A	20	20	20	40
Cluster 1350[2]	No	Yes[3]	No	No	Yes[3]
Cluster 1600	No	Yes	Yes	Yes	Yes
HACMP™ (AIX 5L and Linux)	No	Yes	Yes	Yes	Yes
AIX 5L support	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2
Linux support	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10

1 Requires purchase of optional feature to support micro-partitions 2 Requires special bid configurations 3. Selected models *With maximum I/O drawers Optional

IBM System p5 with POWER5+ QCM

Updated

	IBM System p5 505Q Express	IBM System p5 510Q Express	IBM System p5 520Q Express	IBM System p5 550Q Express	IBM System p5 560Q Express	IBM System p5 560Q Express
						
Footprint, Packaging	19-inch 1U Rack	19-inch 2U Rack	19-inch 4U Rack/Deskside	19-inch 4U Rack/Deskside	19-inch 4U Rack	19-inch 4U Rack
Processor	POWER5+	POWER5+	POWER5+	POWER5+	POWER5+	POWER5+
# of processors (# of cores)	4	4	4	4, 8	4, 8, 16	4, 8, 16
GHz clock	1.65	1.65	1.65	1.65	1.5	1.8
GB memory	1 to 32 DDR2	1 to 32 DDR2	1 to 32 DDR2	1 to 64 DDR2	2 to 128 DDR2	2 to 128 DDR2
Internal storage	73GB – 600GB	73GB – 1.2TB	73GB – 16.8TB*	73GB – 31.2TB*	73GB – 3.6TB	73GB – 32.4TB*
Maximum rPerf	20.25	20.25	20.25	38.34	65.24	75.58
PCI-X slots GX bus slots	2 0	3 0	6 to 34 1	5 to 59 2	6 to 12 2	6 to 12 2
I/O drawers	N/A	N/A	4	8	N/A	8
Micro-partitions[1]	40	40	40	80	160	160
Cluster 1350[2]	No	No	No	Yes[3]	No	No
Cluster 1600	Yes	Yes	Yes	Yes	Yes	Yes
HACMP (AIX 5L and Linux)	Yes	Yes	Yes	Yes	Yes	Yes
AIX 5L support	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2
Linux support	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10	RHEL AS 4 SLES 9 or 10

1 Requires purchase of optional feature to support micro-partitions 2 Requires special bid configurations 3. Selected models *With maximum I/O drawers Optional

POWER6 System Highlights

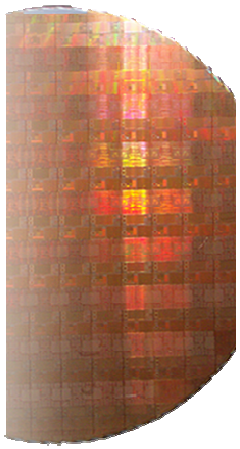

POWER6

IBM POWER6™ Processor Technology

- 5th Implementation of multi-core design
- ~100% higher frequencies
- 4X increase in L2 Cache

POWER6 System Architecture

- New generation of servers
- New IO
  - PCIe, SAS / SATA
  - New IO Drawers
- Enhanced power management

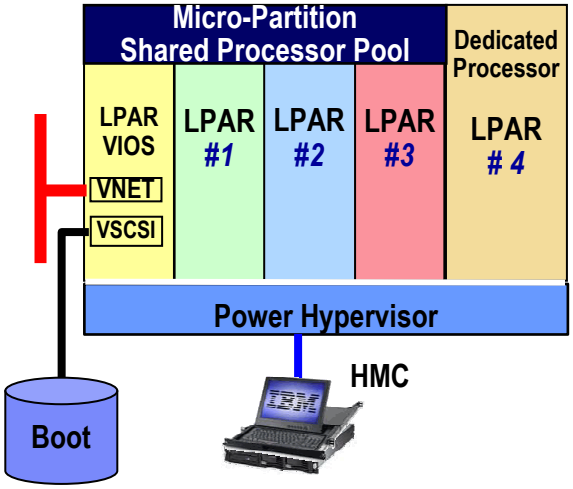

Enhanced Virtualization

- Live Partition Mobility (SoD)
- Dedicated Shared Processors
- Integrated Virtual Ethernet

Availability

- New RAS features
  - Processor Instruction Retry
- Hot-node Add support, Cold-node Repair support
  - SoD

IBM System p Enterprise Servers

	IBM System p5™ 570	IBM System p 570	IBM System p5 575	IBM System p5 590	IBM System p5 595
					
Footprint, Packaging	19-inch 4U rack	19-inch 4U rack	24-inch frame by node	24-inch frame	24-inch frame
Processor	POWER5+™	POWER6	POWER5+	POWER5+	POWER5+
# of processors (# of cores)	2, 4, 8, 12, 16	2, 4, 8, 12, 16	8 , 16	8 to 32	16 to 64
GHz clock	1.9, 2.2	3.5, 4.2, 4.7	1.9, 2.2	2.1	2.1, 2.3
DDR2 GB memory	2 to 512	2 to 768	1 to 256	8 to 1TB	8 to 2TB
Internal storage*	73GB – 79.2TB	73GB – 79.2TB	146.8GB – 2.9TB	146.8GB – 18.7TB	146.8GB – 28.1TB
Maximum rPerf	95.56	134.35	N/A	202.88	393.55
PCIe PCI-X slots PCI-X 266 slots GX bus slots	0 6 to 163 0 1 – 4	4 to 16 0 to 140 2 to 128 2– 8	0 4 to 24 0 2	0 20 to 160 0 6 – 12	0 20 to 240 0 6 – 24
Max I/O drawers	20	32	1	8	12
Max micro-partitions	160[1]	160[1]	160[1]	254	254
System Cluster 1350	No	No	No	No	No
System Cluster 1600	Yes	Yes	Yes	Yes	Yes
HACMP™ (AIX® V5.3 and Linux)	Yes	Yes	Yes	Yes	Yes
AIX support	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2	5.3, 5.2
Linux support	RHEL 4.5 SLES 9 or 10	RHEL 4.5 SLES 9 or 10	RHEL 4.5 SLES 9 or 10	RHEL 4.5 SLES 9 or 10	RHEL 4.5 SLES 9 or 10

1 Requires purchase of optional feature to support micro-partitions *With maximum I/O drawers Optional

JS21 Blade

Blade: JS21

JS21 Blade	
Architecture	2 core (2.7GHz) / 4 core (2.5GHz) PPC970 GX / MP
L3 Cache	N / A
DDR2 Memory	1GB to 16GB
DASD / Bays	0 – 2 SAS disk (36 & 73 GB)
Expansion	Daughter Card via PCI-X or PCIe
SCSI Expansion	Yes
Integrated Ethernet	Dual Port 10/100/1000 Ethernet
Fiber Support	Yes (via Blade center)
Media Bays	1 Blade Center
Remote IO Drawers	N / A
Redundant Power	Yes Blade Center
Redundant Cooling	Yes Blade Center
LPAR	Up to 40 partitions (4-way)
Virtualization	AIX V5.3 & Linux Advanced POWER Virtualization (Opt) Integrated Virtualization Manager
Certification	**RoHS Compliant**
OS Support	AIX 5.2 , 5.3 and Linux

Linux

POWER5/5+ to POWER6 Upgrade SOD

IBM is committed to enhancing their customer's investments in the IBM System p product line. Based on this commitment, IBM plans to provide an upgrade path from the current p5-570, p5-590, and p5-595 servers to IBM's next-generation POWER6 processor-based enterprise servers.

POWER5/5+

POWER6

p5-59x

p5-59x+

POWER6

**All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.*
Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

System p 570

Warranty:
1-year NBD

Base system	2- or 4-core systems @ 3.5 / 4.2 / 4.7GHz POWER6 ▸ Expandable to 16-core system 4U rack-mount modules
Functions supported	Dynamic LPAR with Shared Dedicated Processor support IBM Advanced POWER Virtualization option ▸ Micro-Partitioning support (1/10th processor granularity) ✓Maximum 160 partitions ▸ Virtual networking and storage support Integrated Virtual Ethernet Capacity on Demand for processors and memory
Features per Module	Up to 192GB DDR2 memory ▸ Max 768GB 16-core system Four PCIe& Two PCI-X slots Dual Service processor support Integrated Virtual Ethernet USB: 2; System: 2, RIO: 2 One media bay (Optional) Two GX Bus slots Up to six SAS disk drives Redundant cooling and power I/O drawers: ▸ 7311-D11, 7311-D20, or 7314-G30 ▸ Maximum of 32 drawers / 16-core system
Software support	AIX V5.2 and AIX V5.3 Red Hat Enterprise Linux 4.5 for POWER / SoD RHEL 5 SUSE Linux Enterprise Server 9 or 10 for POWER

System p 570 POWER6 Packaging

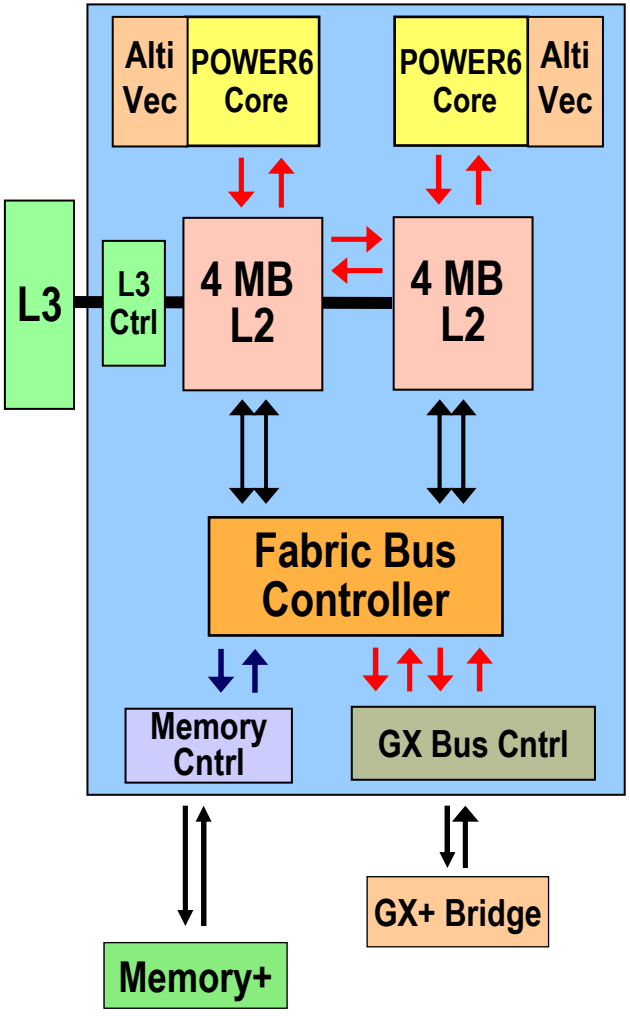

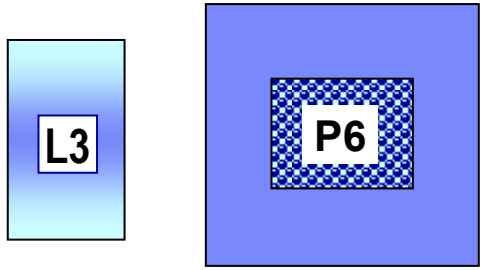

Single Chip Module POWER6 & L3

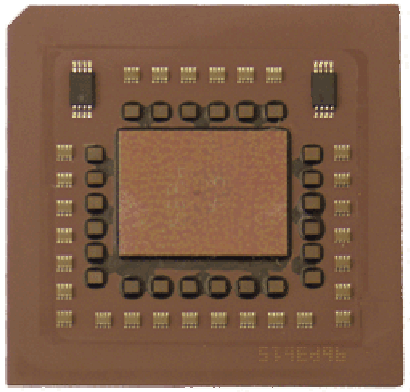

POWER6 Processor Module

POWER6 / 5+ Midrange Processor / Memory Packaging

SCM @ 3.5 – 4.7 GHz
Up to 96 DDR2+ memory
Two cores / book

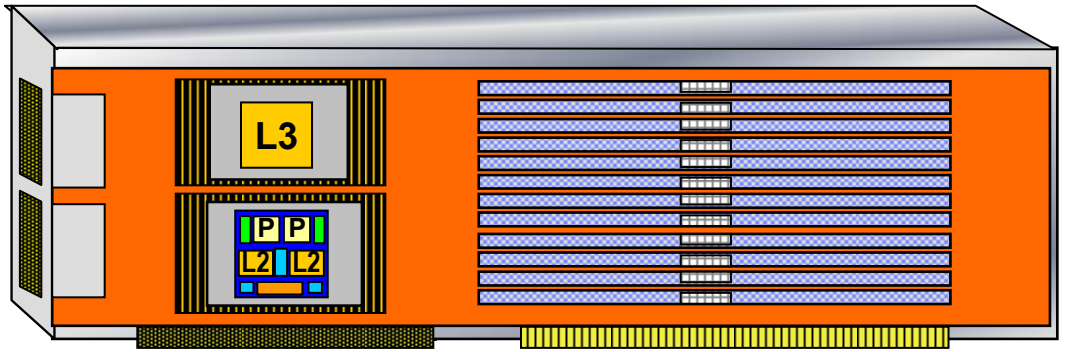

p5-570+ DCM @ 2.2 GHz
Up to 64 GB DDR2 memory
Two cores / book

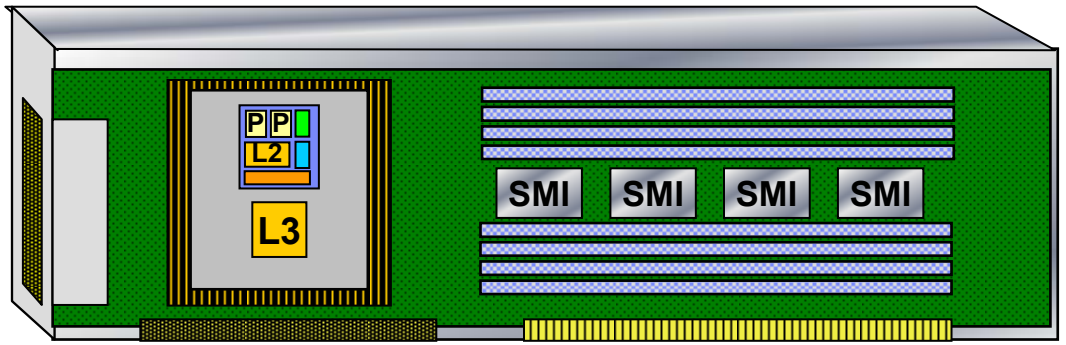

p5-570+ DCM @ 1.9 GHz
Up to 32 GB DDR2 memory
Two cores / book

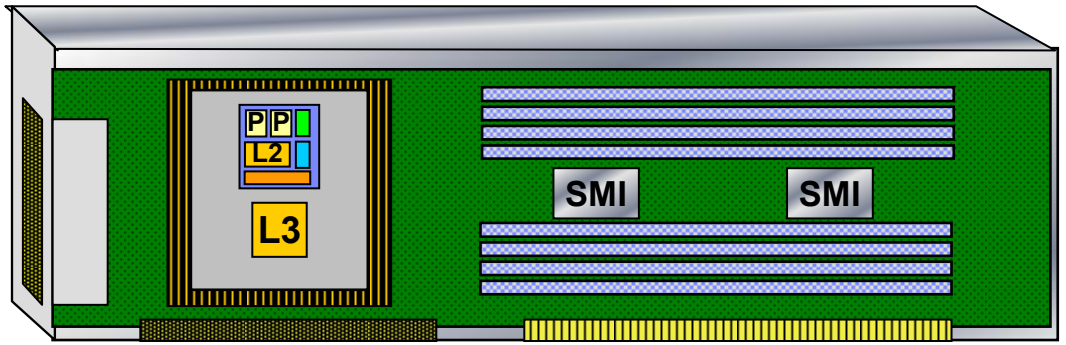

POWER6 DDR2 Memory DIMMs

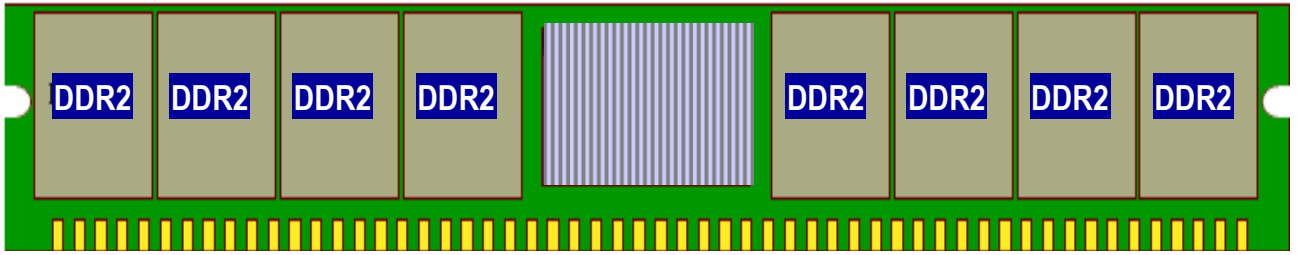

Front

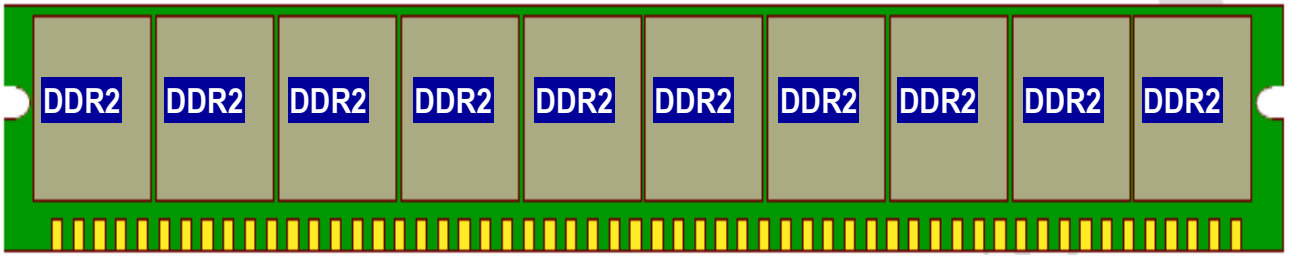

Back

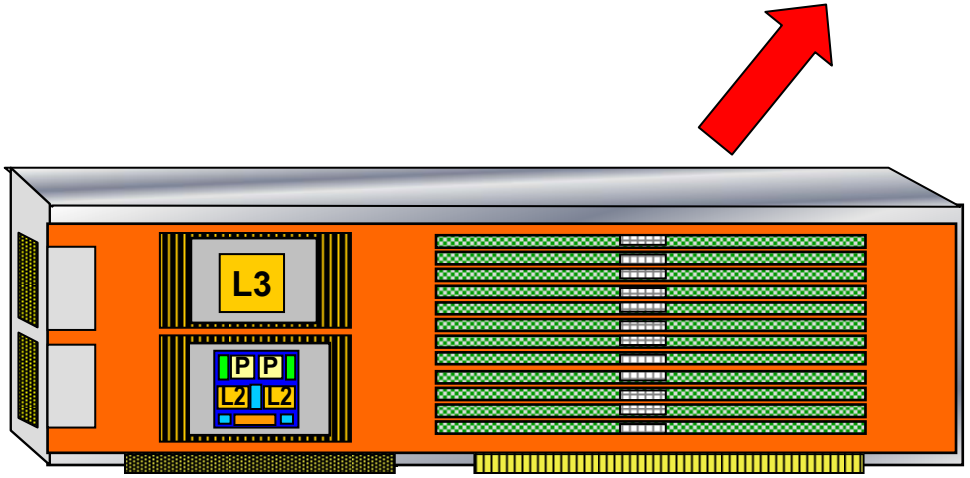

50% More DIMM Slots
Greater Memory Flexibility

Front View

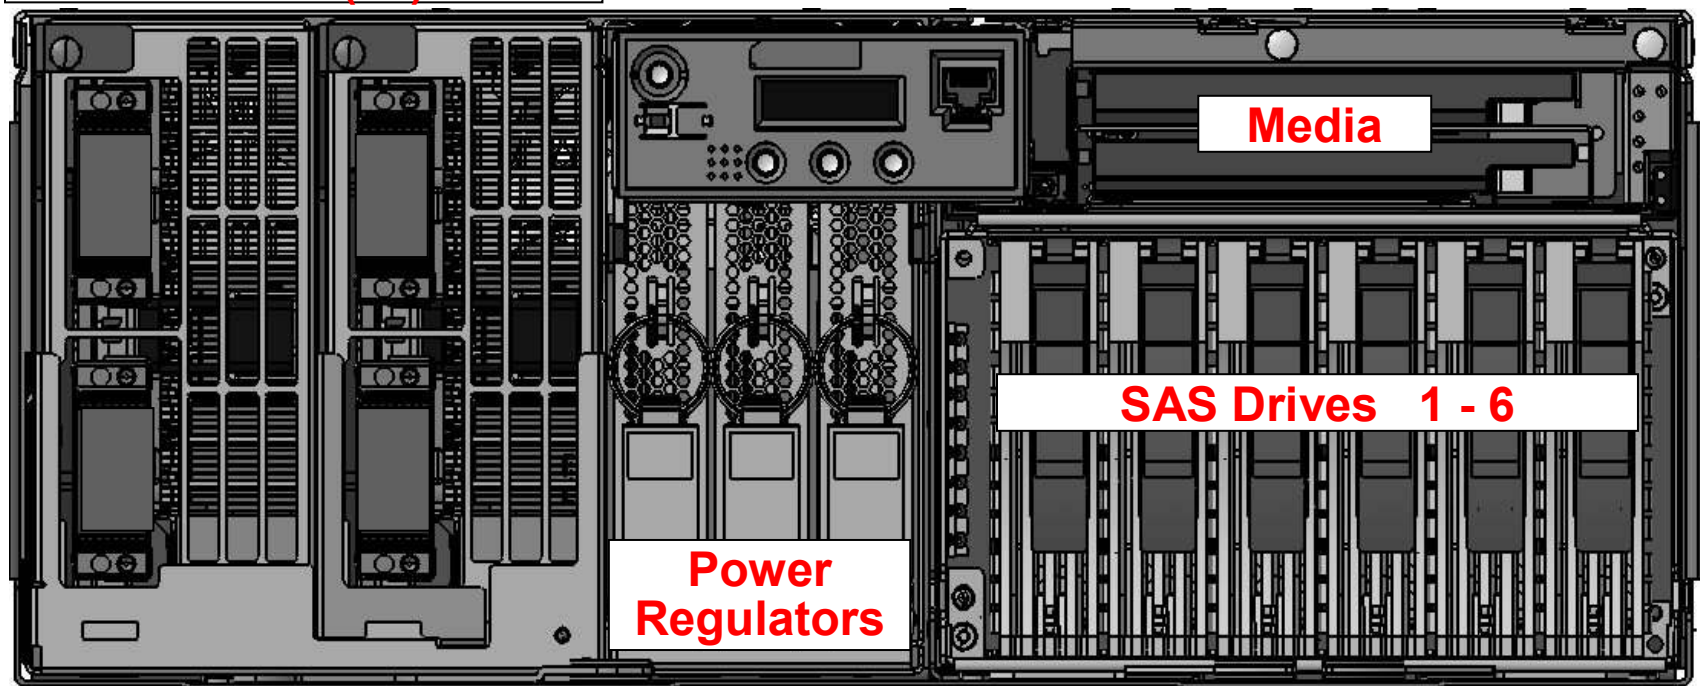

- Six 3.5” SAS disk bays
- One SAS/SATA controller
- SAS hot-swappable disk drives supported:
  - 15K rpm: 73.4, 146.8 and 300GB
- Maximum internal capacity of 1.8 TB

* 79.2 TB with optional I/O drawers in a 16-core system

Top View

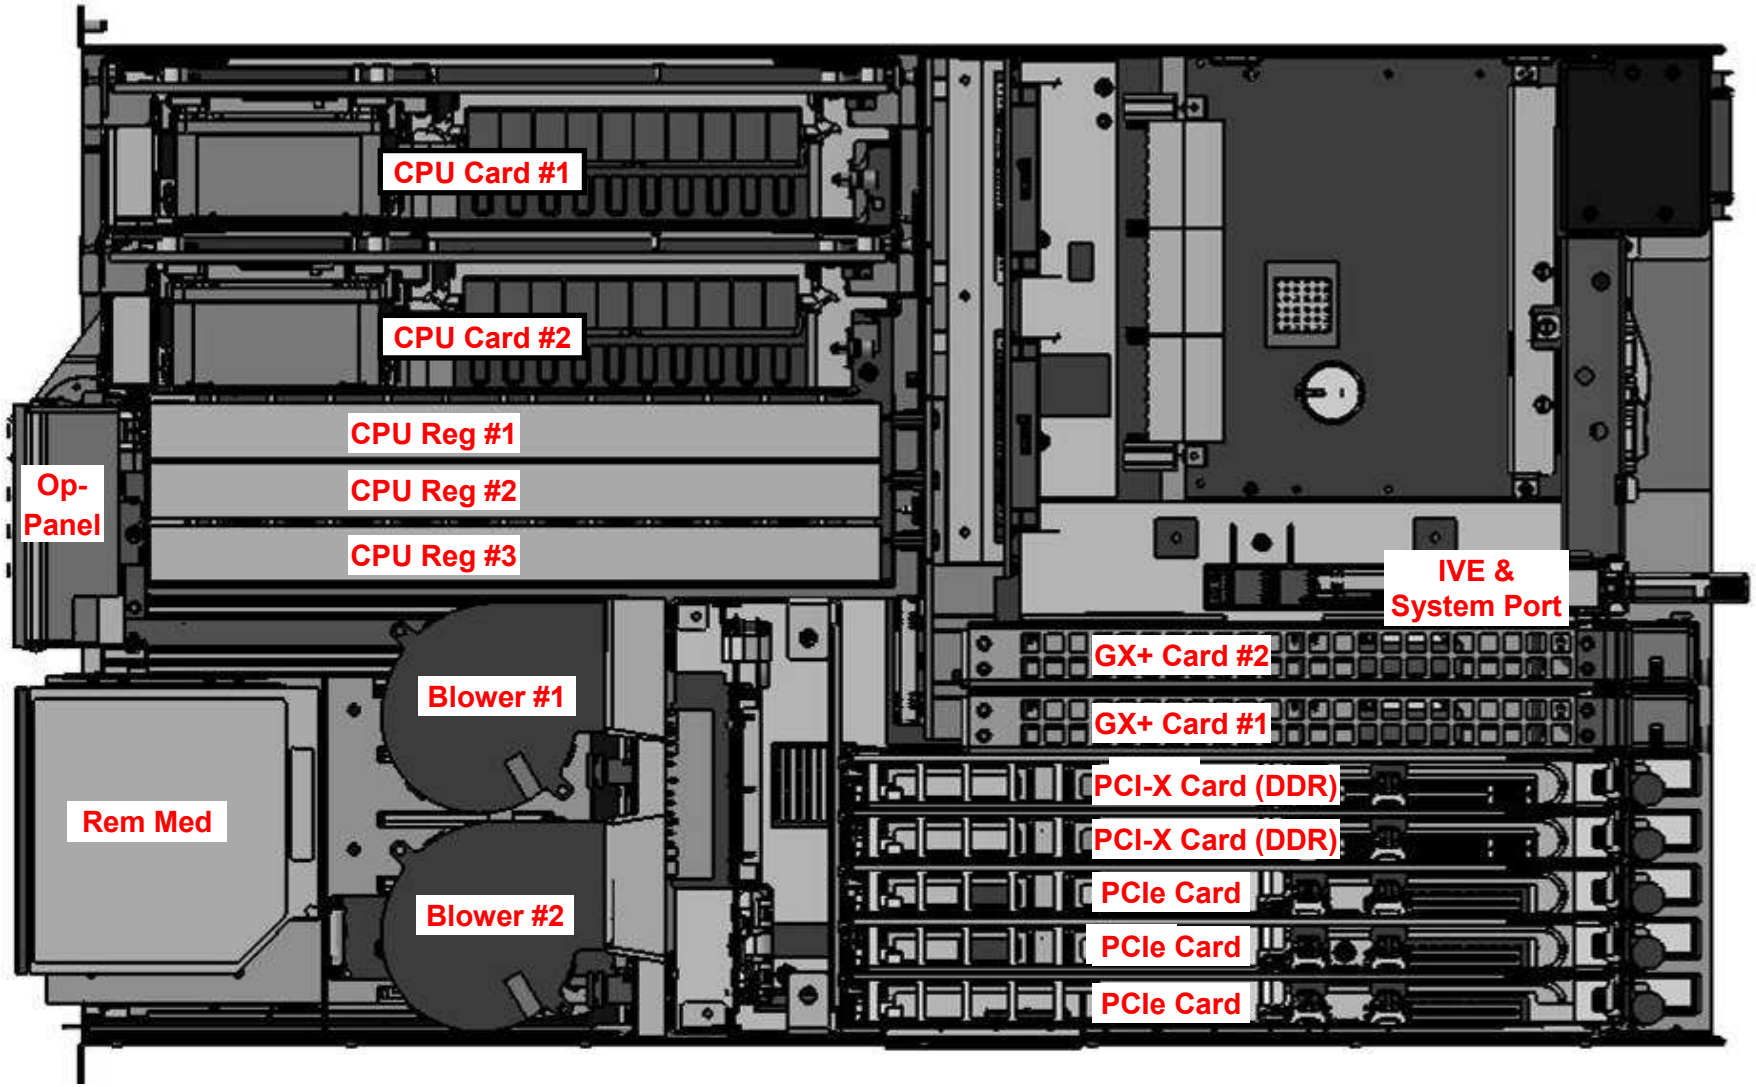

Back View

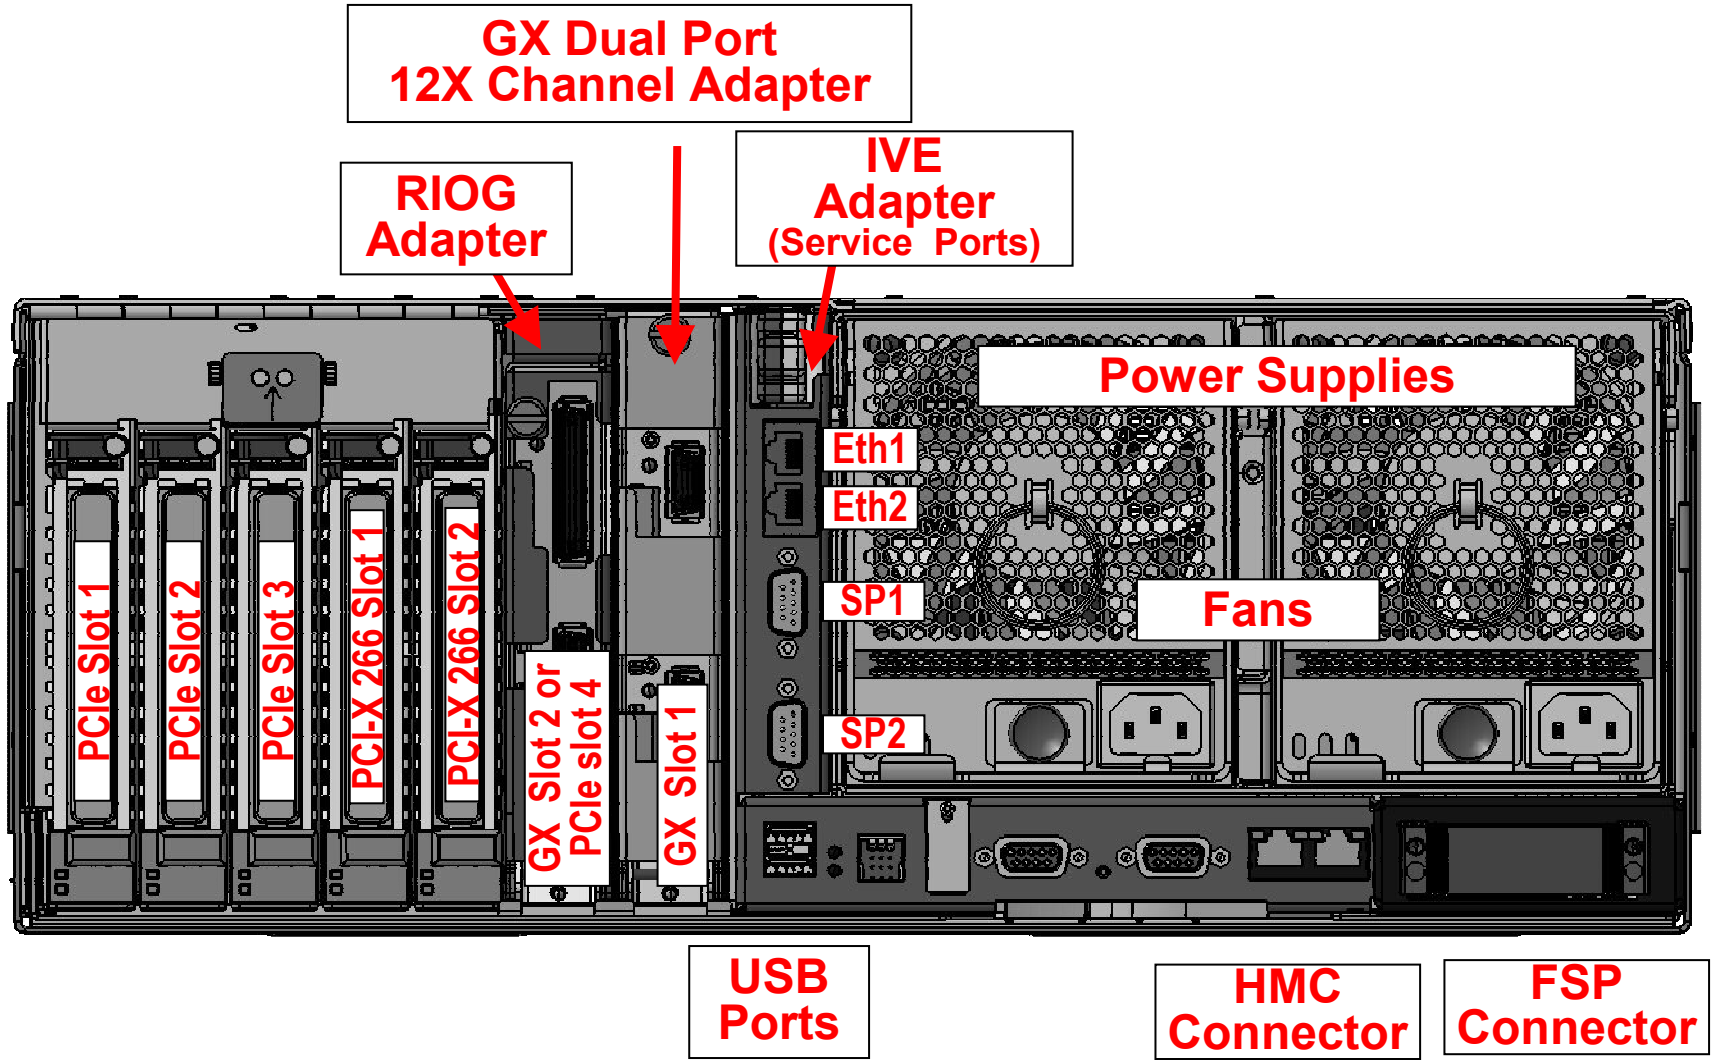

Side View

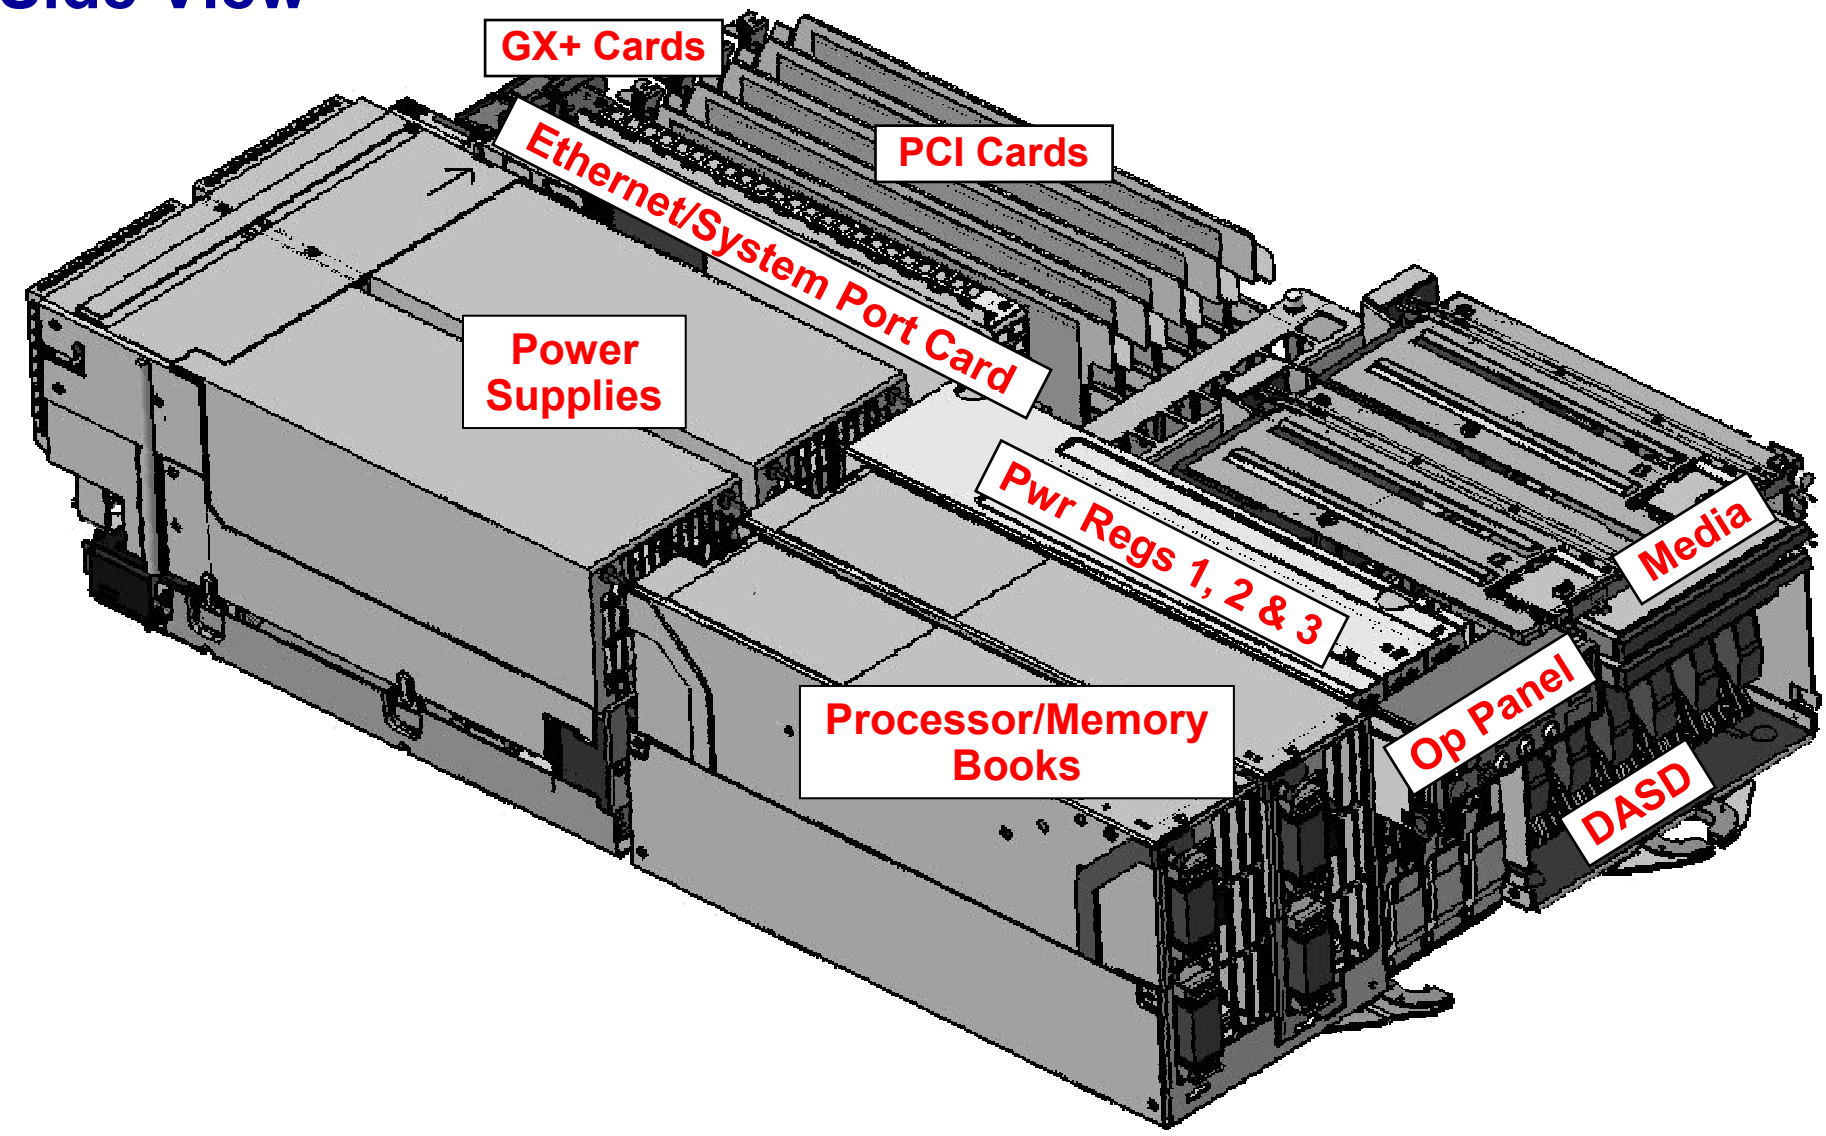

Building Block Modules: SMP Mid-range Server

Interconnect configurations of 8 / 12 / 16-core servers
Point-to-point connections

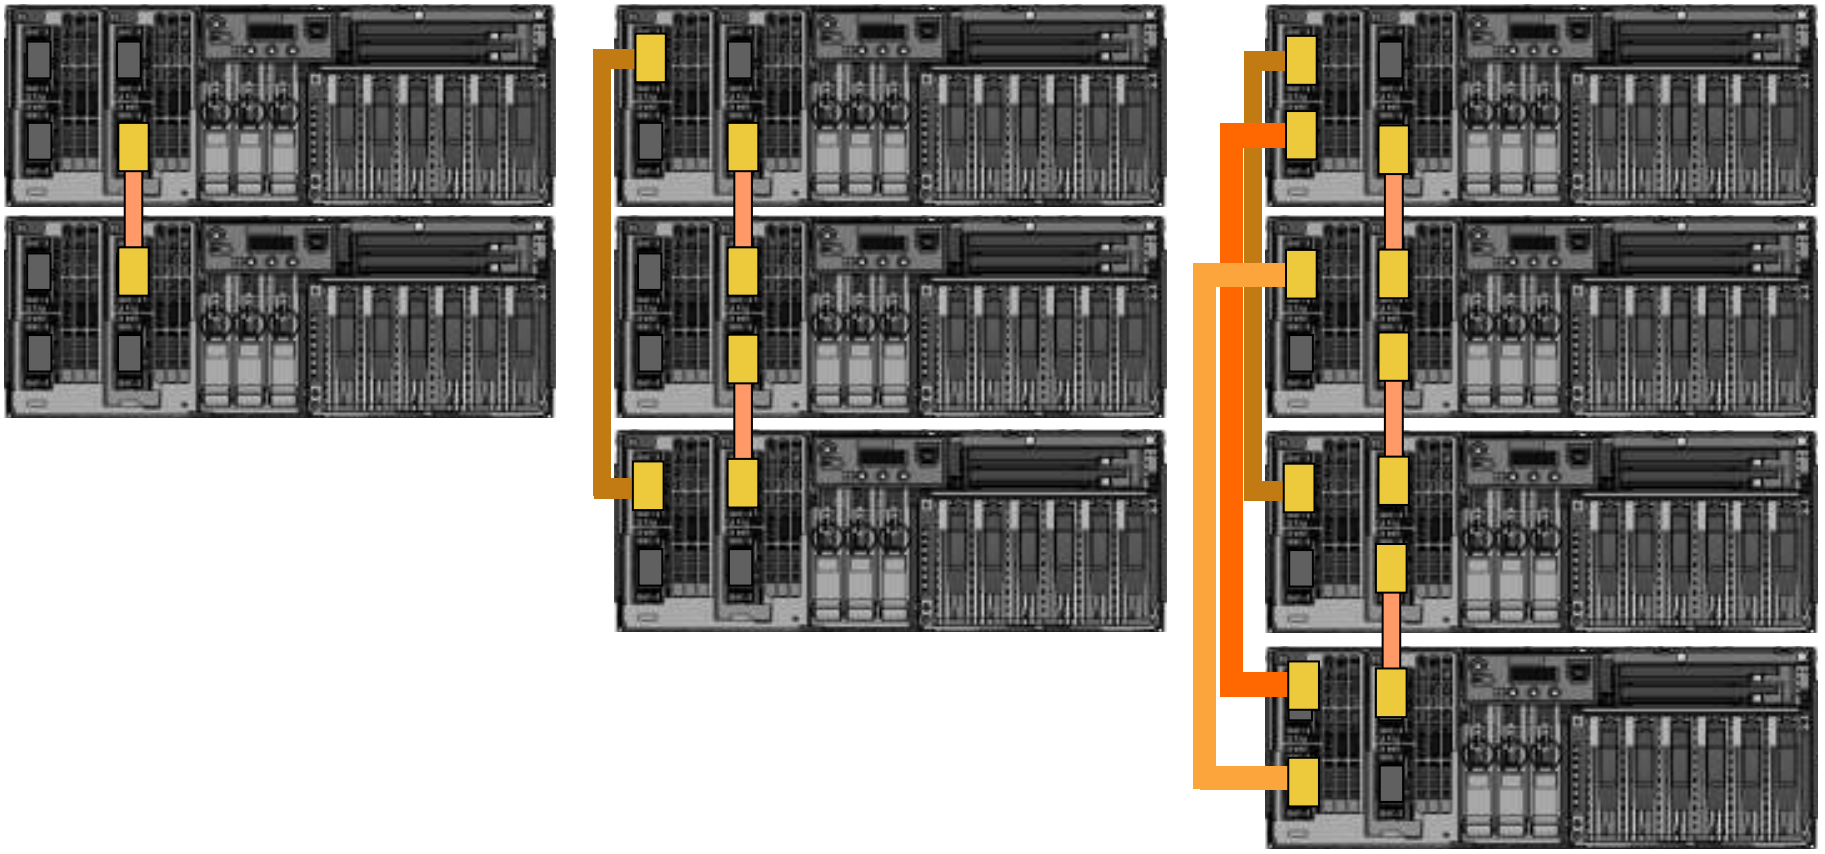

System p 570 Bandwidth

Memory	Bandwidth
L1 (Data)	**75.2** GB/sec
L2	**300.8** GB/sec
L3	**37.6** GB/sec
Memory	**32** GB/sec
Inter-Node Buses (16w)	**75.2** GB/sec
Intra-Node Buses (16w)	**100.26** GB/sec
Internal I/O Bus GX Bus Slot 1 GX Bus Slot 2 Total I/O Bandwidth	**4.7** GB/sec / node **4.7** GB/sec / node **6.266** GB/sec / node **62.6** GB/sec (16w)

Calculations for 4.7 GHz processors and 667 MHz memory

Service Processor Cabling Layouts

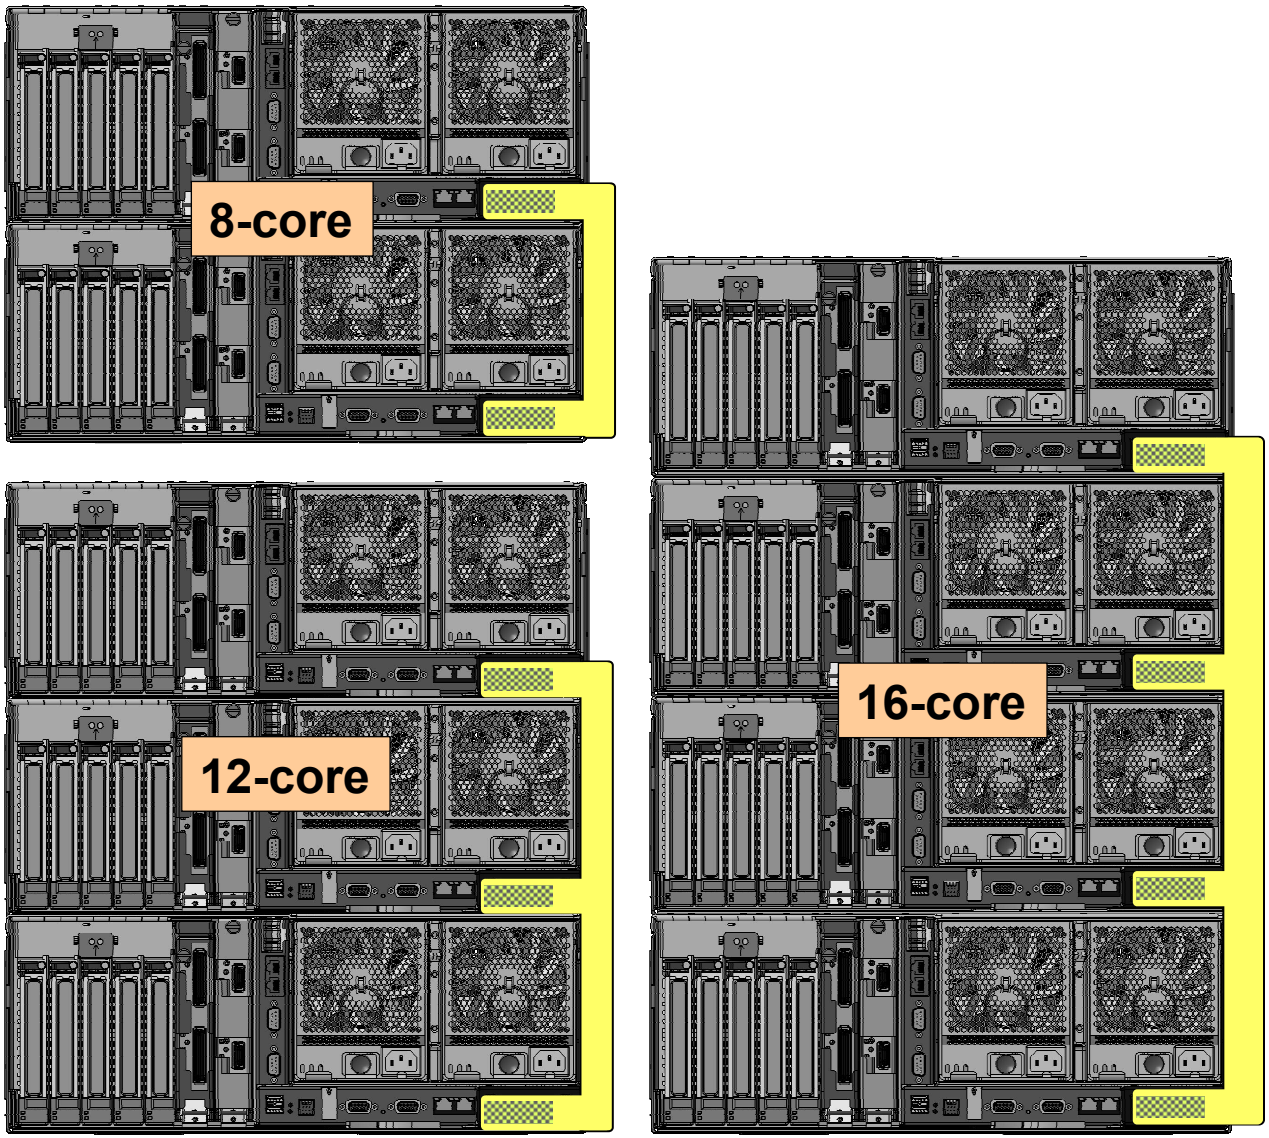

Service processor and clock use flat flex cabling to integrate individual 4-core servers into a single SMP server (rear view)

SAS (Serial Attached SCSI) DASD

Parallel SCSI vs. Serial Attached SCSI (SAS)		
	Parallel SCSI	**SAS**
Architecture	Parallel, all devices connected to shared bus	Serial, point-to-point, discrete signal paths. Often requires switches for fanout
Performance	320 Mbytes/sec (Ultra320 SCSI); performance degrades as devices added to shared bus	3.0 Gbits/sec full duplex, roadmap to 12.0 Gbits/sec; performance maintained as more drives added
Scalability	15 drives	Over 16,000 drives
Compatibility	Incompatible with all other drive interfaces	Compatible with Serial ATA (SATA)
Max. Cable Length	12 meters total (must sum lengths of all cables used on bus)	Eight meters per discrete connection; total domain cabling thousands of feet
Cable Form Factor	Multitude of conductors adds bulk, cost	Compact connectors and cabling save space, cost
Hot-plug Ability	Yes (not inherent in architecture)	Yes
Device Identification	Manually set, user must ensure no ID number conflicts on bus	Worldwide unique ID set at time of manufacture; no user action required
Termination	Manually set, user must ensure proper installation and functionality of terminators	Discrete signal paths enable devices to include termination by default; no user action required

System p 570 SoD

Redundant Service Processor

For POWER6 processor-based p570 systems with at least two CEC enclosures to have redundant service processor function no later than the end of 2007. This feature will be provided at no additional charge to existing POWER6 processor-based p570 users via a system firmware update.

Hot-node Add Support

POWER6 processor-based p570 systems will be able to add an additional CEC enclosure (node) to their p570 system, without powering down the system (hot- node add). The additional p570 enclosure would be ordered as a system upgrade and added to the original system while operations continued. The additional resources of the new additional enclosure, could then be assigned to existing applications or new applications as required. This capability will be provided at no additional charge to existing POWER6 processor-based p570 users via a system firmware update.

Planned availability: No later than the end of 2007

Integrated Virtual Ethernet

Base Offering: #5636
2 Serial, 2 1Gb Eth

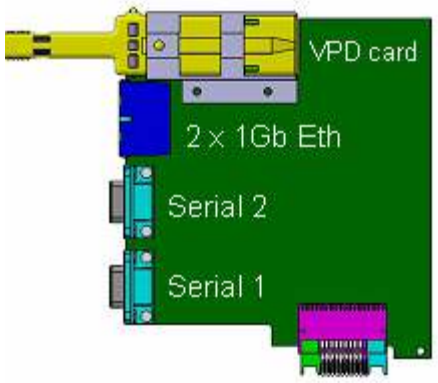

4 x 1Gb Upgrade Offering: #5639
1 Serial, 4 1Gb Eth

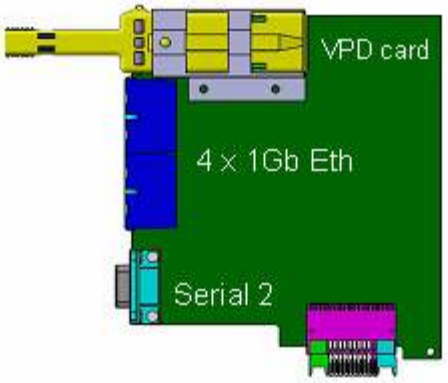

10Gb Upgrade Offering: #5637
1 Serial, 2 10Gb Eth

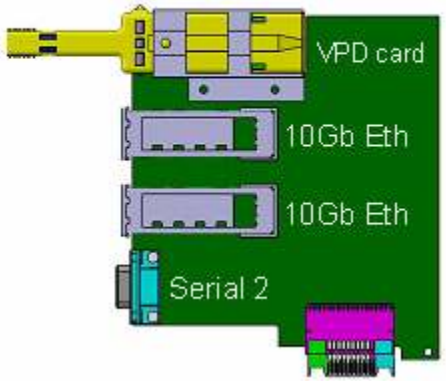

Address Sharing:

•Dual 1GB:	**16 MAC Addresses / pair**	**Total: 16**
•Quad 1GB	**16 MAC Addresses / pair**	**Total: 32**
•Dual 10GB:	**16 MAC Addresses/ port**	**Total: 32**

Non VIOS Partition: Address Sharing (MAC Addresses)

•Time Slicing "Physical" Ethernet adapter resources
VIOS Partitions: IVE ports dedicated

•Each Physical port uses 102MB of memory
A 4 port card uses 408MB of memory

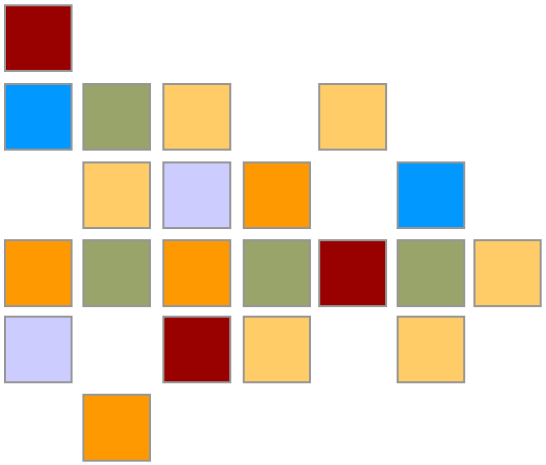

Reliability, Availability and Serviceability

Primary POWER RAS Features

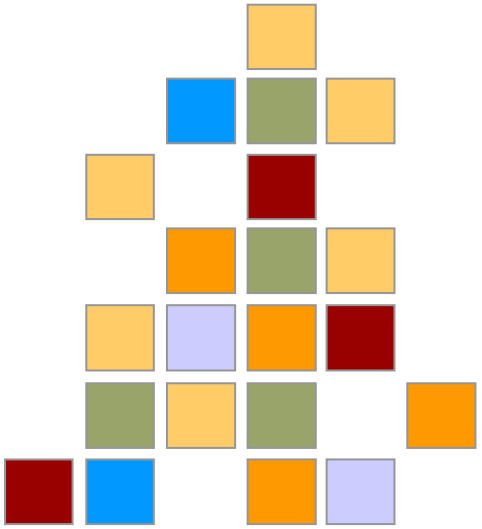

- ✓ Processor Instruction Retry
- ✓ Alternate Processor Recovery
- ✓ First Failure Data Capture
- ✓ DDR Chipkill™ memory
- ✓ Bit-steering/redundant memory
- ✓ Service Processor Failover*
- ✓ Dynamic Firmware Maintenance*
- ✓ Hot I/O Drawer Add*
- ✓ I/O error handling extended beyond base PCI adapter
- ✓ ECC extended to inter-chip connections for the fabric/processor buses
- ✓ Memory and L3 Cache soft scrubbing
  - ✓ Hardware Assisted
- ✓ L2 & L3 Cache Line Delete
- ✓ Hardware Assisted Memory Scrubbing
- ✓ Live Partition Migration
- ✓ P570 Concurrent Add & Cold Repair (SoD)

HMC required to enable these functions..

Primary POWER RAS Features con't

- ✓ Redundant power, fans
- ✓ Dynamic Processor Deallocation
- ✓ Dynamic processor sparing
- ✓ ECC memory
- ✓ Persistent memory deallocation
- ✓ Hot-plug PCI slots, fans, power
- ✓ Internal light path diagnostics
- ✓ Hot-swappable disk bays

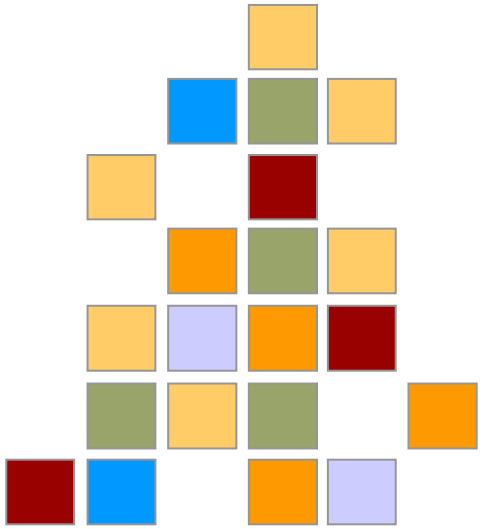

HMC required to enable these functions..

World-class Hardware RAS

Summary of key System p RAS features

Core System Design

- *High quality parts*
- *Fewer parts = Fewer failures*
- *Designed for low power consumption (less heat = fewer failures)*
- *Manufacturing methods, packaging, cooling*
- *Continuous System and Commodity Quality Actions*
- *Integrated RAS features*
- *Failure Avoidance Methodology*
- *Designed for Ease of Service*

Fault Resilience

- *N+1 Power Supplies, regulators, power cords*
- *Dual redundant fans*
- *Dynamic Processor Deallocation and sparing*
- *"Chipkill" Technology*
- *Predictive Failure Analysis*
- *Auto Path Reassignment - data paths, power*
- *Processor Instruction Retry*

System Restore

- *Deferred Repair*
- *Concurrent Repair*
- *LED Service Identification*
- *Service Consoles*
- *Migration to Guided Maintenance*

Fault Isolation & Diagnosis

- *First Failure Data Capture*
- *Run Time Self Diagnostics*
- *Service Processor*
- *Rifle-shot repairs (no "plug and pray" parts replacement approach)*

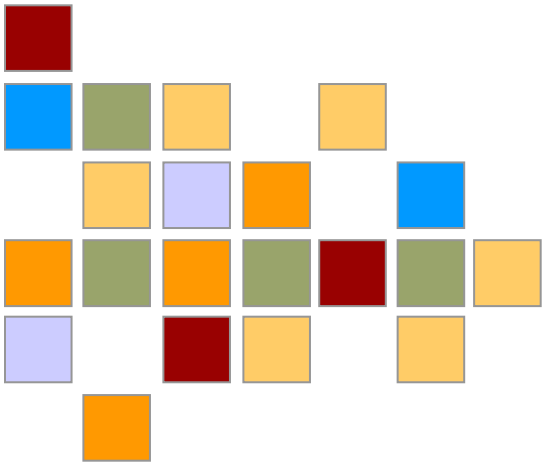

Virtualization

IBM's 40-year History of Leadership in Virtualization

1967	1973	1987	1997	2001	2004	2007
IBM develops **hypervisor** that would become VM on the mainframe	IBM announces first machines to do **physical partitioning**	IBM announces **LPAR on the mainframe**	**POWER LPAR design** begins	IBM introduces **LPAR in POWER4™ with AIX**	**Advanced POWER Virtualization** ships	IBM announces POWER6, **the first UNIX® servers with Live Partition Mobility**

"In our opinion, they [System p servers] bring mainframe-quality virtualization capabilities to the world of AIX."

- Ulrich Klenke, CIO, rku.it
January 2006

Advanced POWER Virtualization on IBM System p servers

AIX
Linux® on POWER

client quote source: rku.it case study published at http://www.ibm.com/software/success/cssdb.nsf/CS/JSTS-6KXPPG?OpenDocument&Site=eserverpseries

40% of all System p5™ Servers are Running APV

Royal Bank of Canada
BRED Banque
HSBC
AES/PHEAA
Fifth Third YeePay
Raiffeisen Bank

Acquedotto
Alstom
ISAGEN
TransAlta

Metlife

Franklin Covey
Dundee University
Central Michigan U

Whirlpool
Miele

Sinopec
Bangchak Petroleum

Campbell Soup

Colgate

Plansee
Wuerth Group
Teleflex
Pilz CmbH & Co

Audi
DaimlerChrysler

UPMC
Byram Healthcare
Mayo Clinic
St Michael's Hospital

Source: IBM Finance data. Full case studies for each named client are available at http://www.ibm.com/software/success/cssdb.nsf/advancedsearchVW?SearchView&Query=(Virtualization)+AND+[WebSiteProfileListTX]=eserverpseries&site=eserverpseries&frompage=ts&Start=1&Count=30&cty=en_us

Virtualization Changes Everything

- Deployment, maintenance, and migration of IT resources are top contributors to cost today
- With the right tools, virtualized resources can be easier to create, adjust, move, clone, checkpoint
- New complexities can emerge
  - Rapid growth of virtualized resources across multiple environments
  - Relationship of virtualized resources to underlying physical infrastructure
  - Health monitoring and problem determination across a physical and virtualized infrastructure

Server Management and Administration Costs

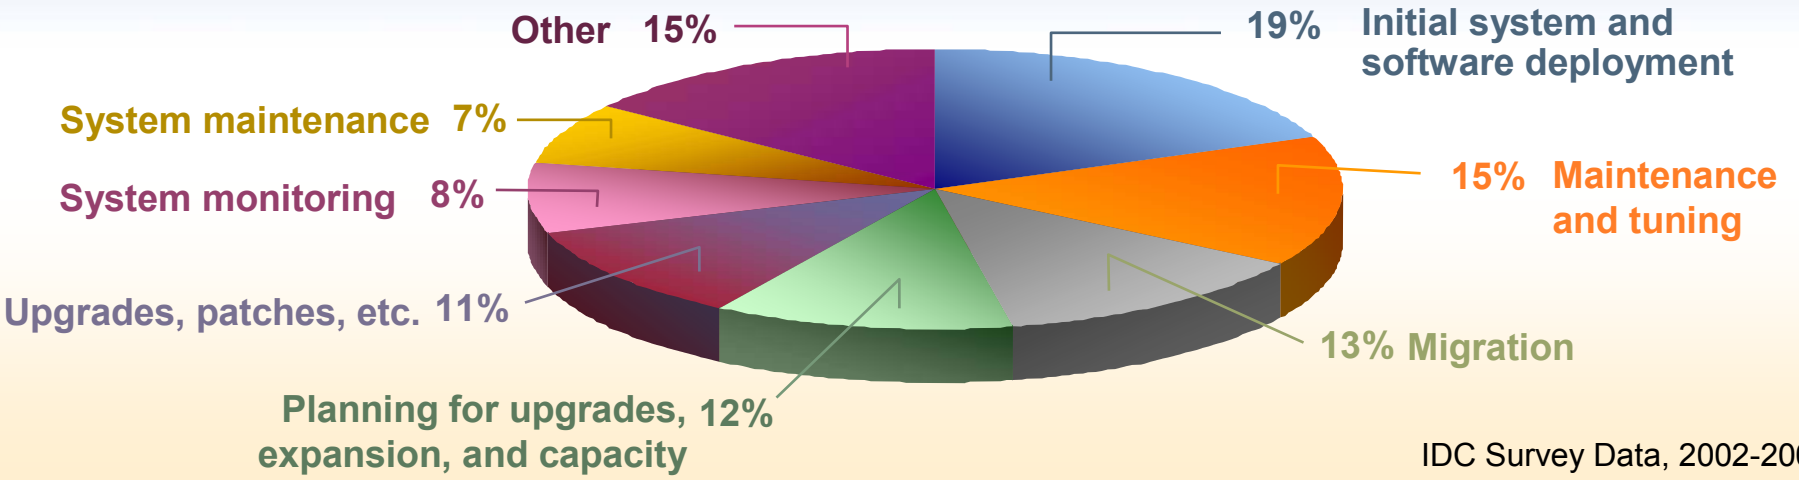

IDC Survey Data, 2002-2004

How Server Consolidation can Fuel Business Innovation

Paradigm Shift for Datacenters:
Pure cost center with 10% cut year over year ➔ Business Unit with strategic growth investments

Datacenter Example
- Pure cost center
- Management cuts costs 10% YOY
- IT cuts innovation and directs to operations

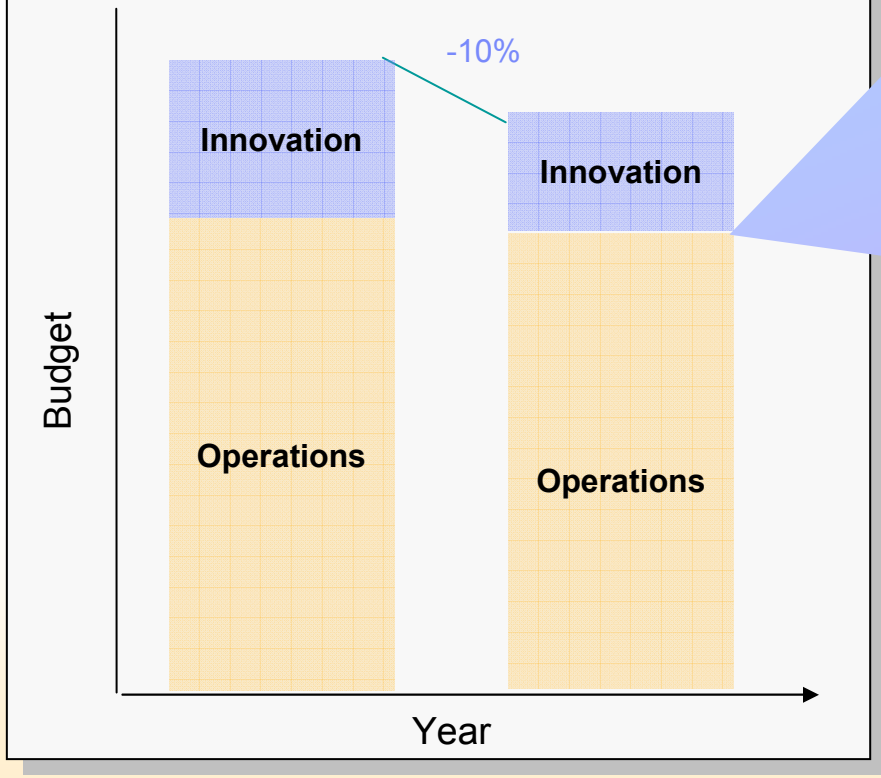

Self-Funding Model

Optimizing IT

- Energy Efficiency
- Server Consolidation and Virtualization
- Dynamic Infrastructure

POWER Virtualization Architecture

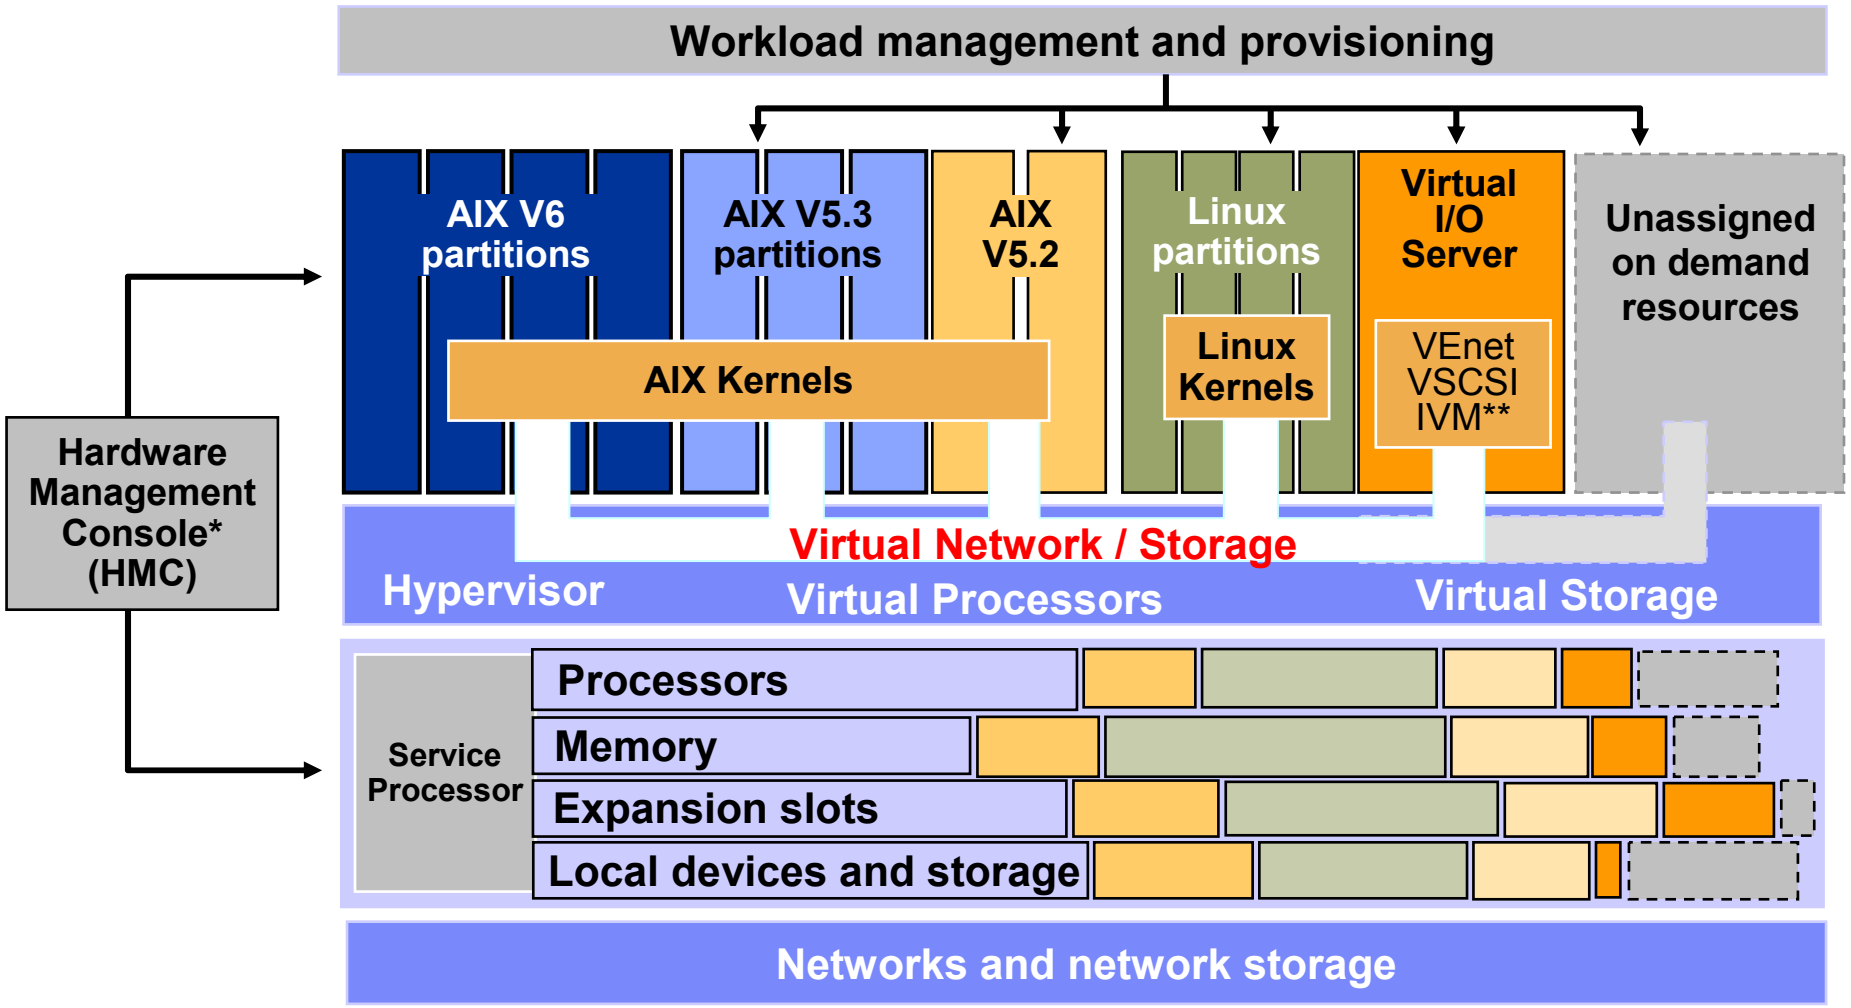

*Integrated Virtualization Manager (IVM) is disabled if HMC attached
**Available on System p5 560Q and below as well as the BladeCenter® JS21

POWER5 Advanced Power Virtualization Option

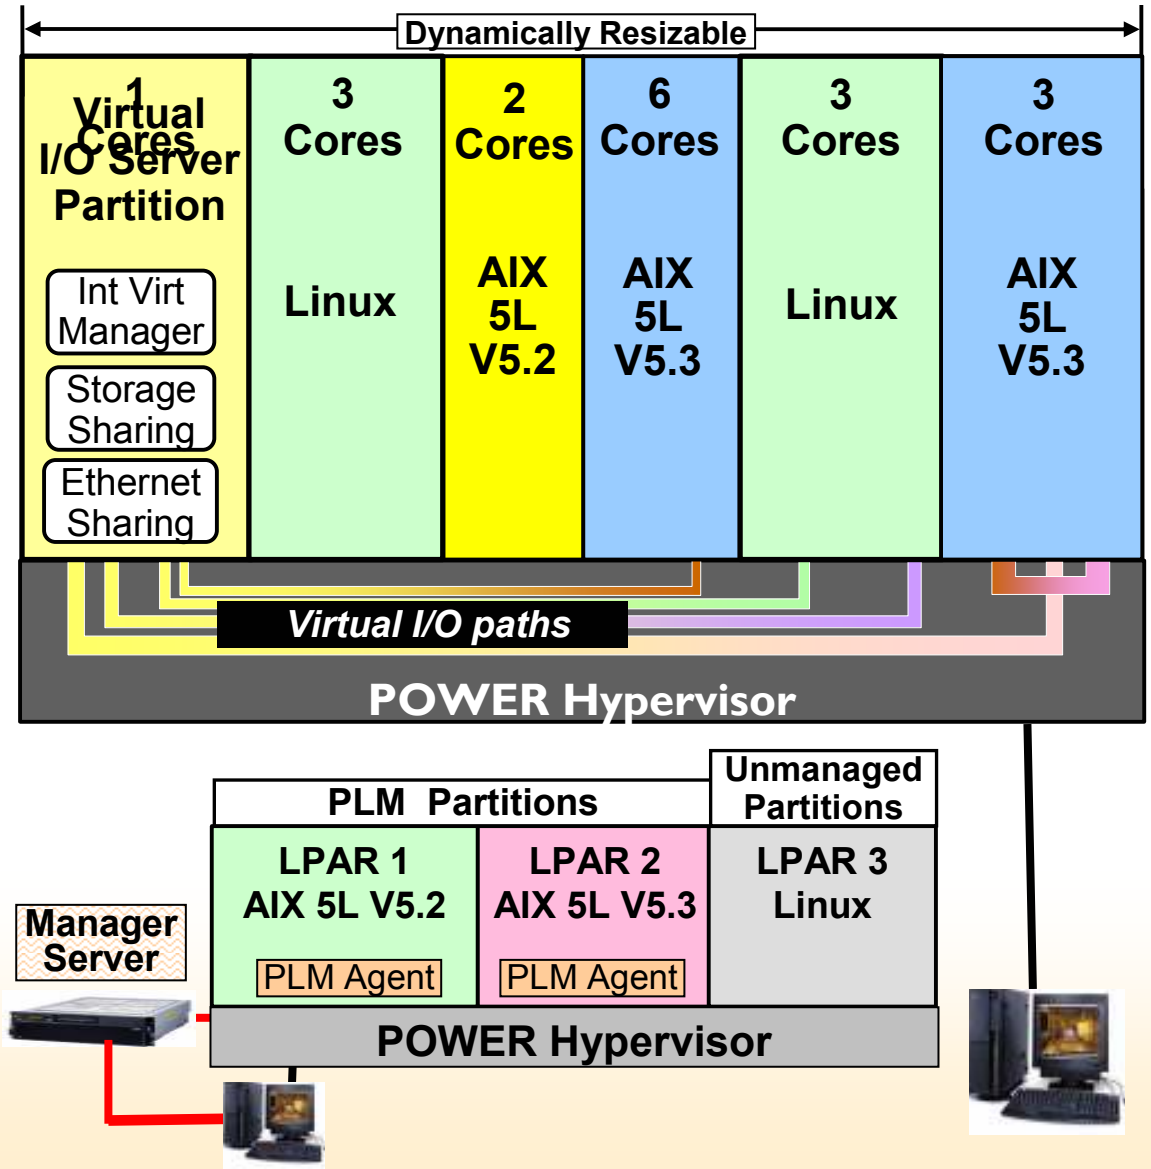

Features

- Micro-partitioning
  - Share processors across multiple partitions
  - Minimum Partition: 1/10 processor
  - AIX 5L V5.3 or Linux*
- Virtual I/O Server
  - Shared Ethernet
  - Shared SCSI & Fiber Channel
  - Int Virtualization Manager
  - AIX 5L V5.3 & Linux partitions
  - From 1 to 10 per server
- Partition LoadManager
  - AIX 5L V5.2 & V5.3 supported
  - Balances Processor & memory request
- Managed via HMC

*** = SLES 9 or RedHat v3 with update 3**

Advanced POWER Virtualization Option for POWER6

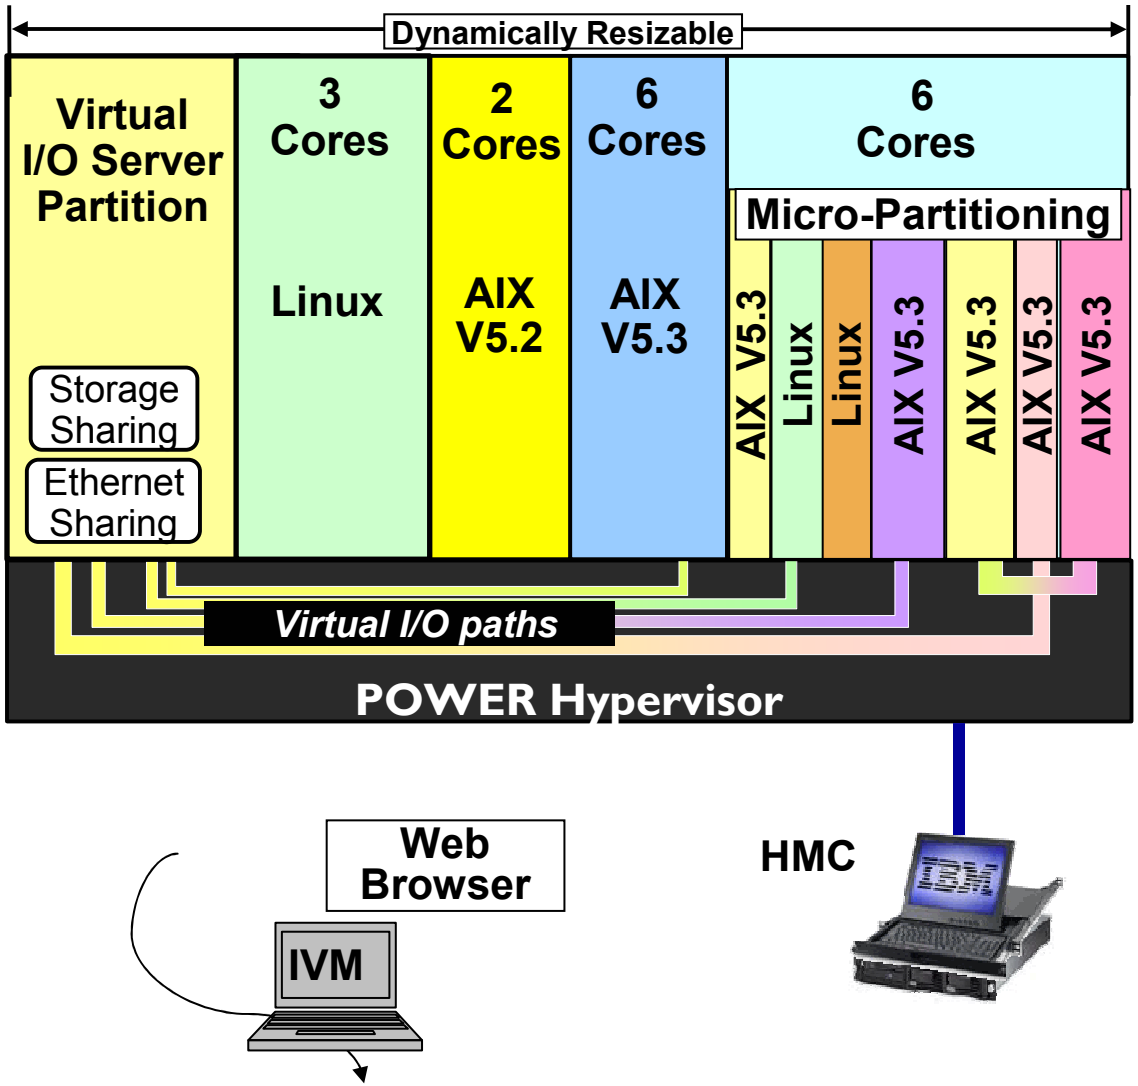

- **New Offering for POWER6**
- **Micro-Partitioning™**
  - Share processors across multiple partitions
  - Minimum Partition: 1/10 processor
  - AIX V5.3 / V6
  - Linux
- **Virtual I/O Server**
  - Shared Ethernet
  - Shared SCSI & Fiber Channel
  - Integrated Virtualization Manager
- **Live Partition Mobility (SoD)**
- **Upcoming Shared Resources**

LPARs - p4 vs p5 & p6

POWER4

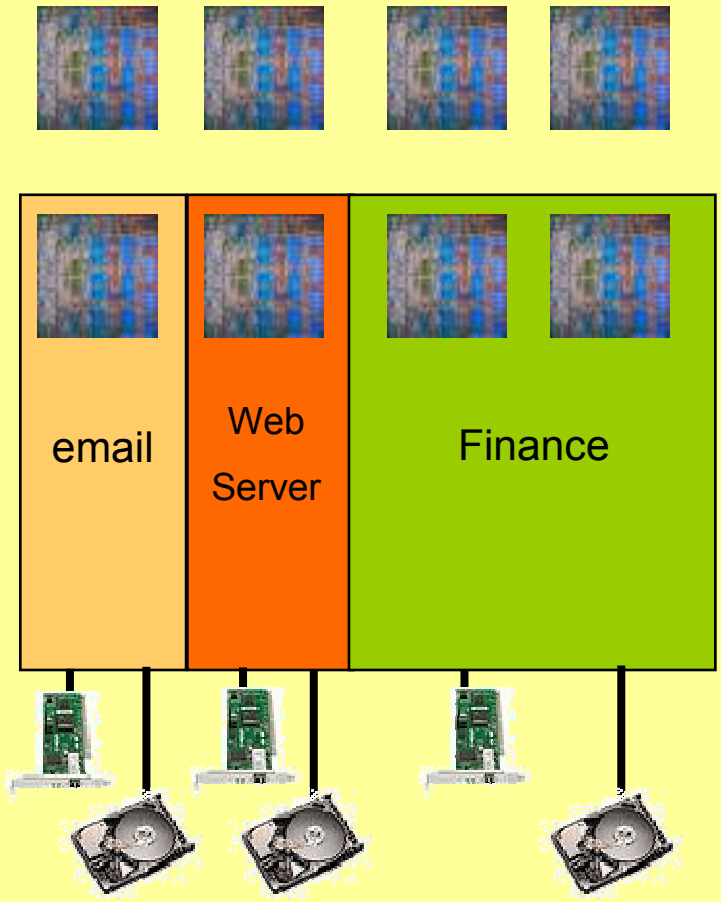

POWER5 – POWER6

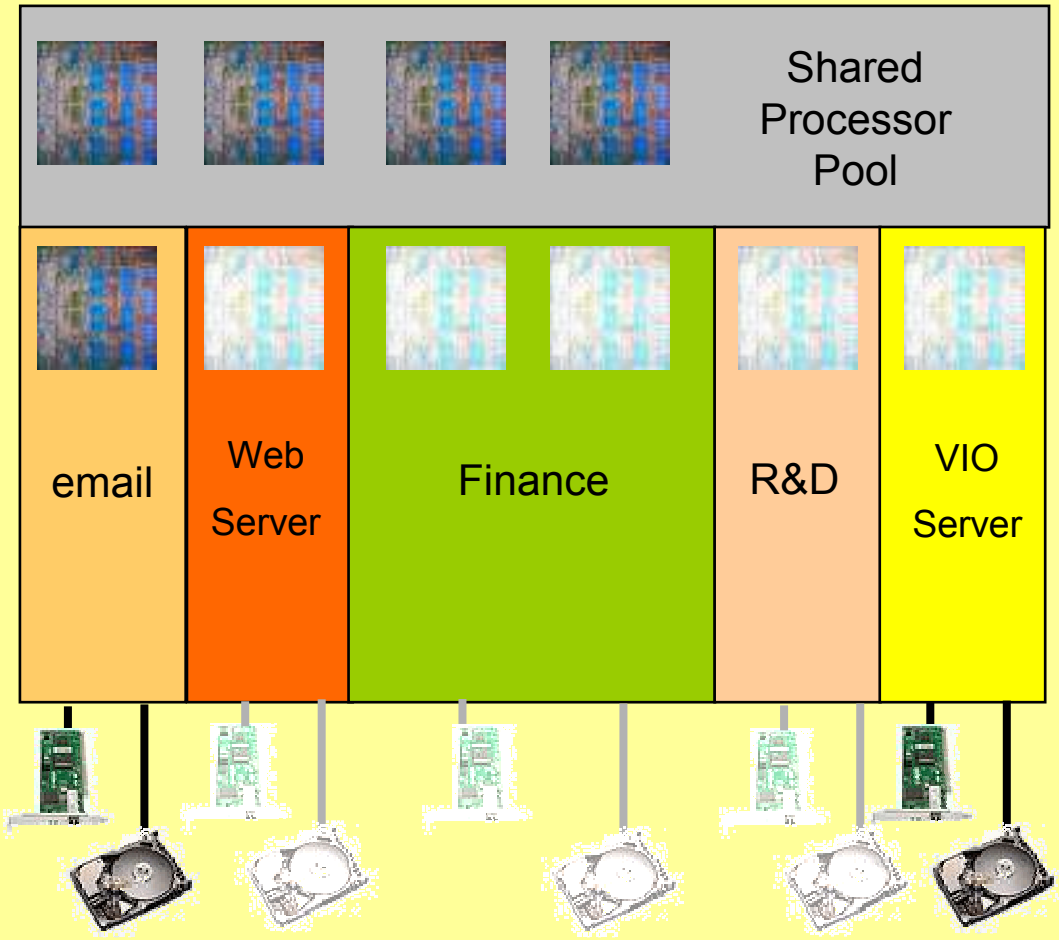

POWER6: Simultaneous Multithreading

POWER5 Simultaneous Multithreading

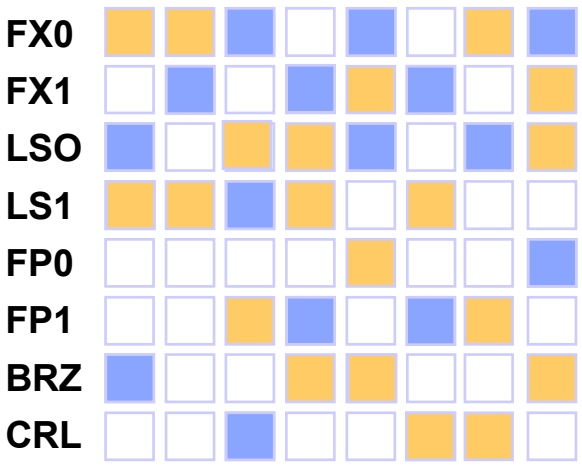

Appears as four CPUs per chip to the operating system (AIX V5.3 and Linux)

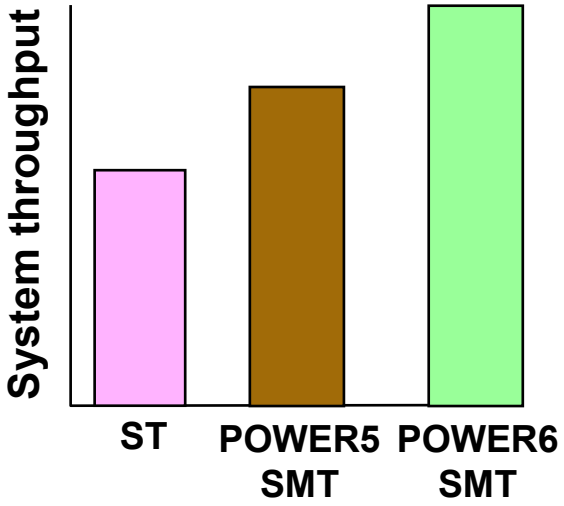

- Utilizes **unused execution** unit cycles
- **Reuse of existing transistors vs. performance from additional transistors**
- Presents symmetric multiprocessing (SMP) programming model to software
- Dispatch two threads per processor: *"It's like **doubling** the number of processors."*
- Net result:
  - ***Better performance***
  - ***Better processor utilization***

Acme's p5 Server

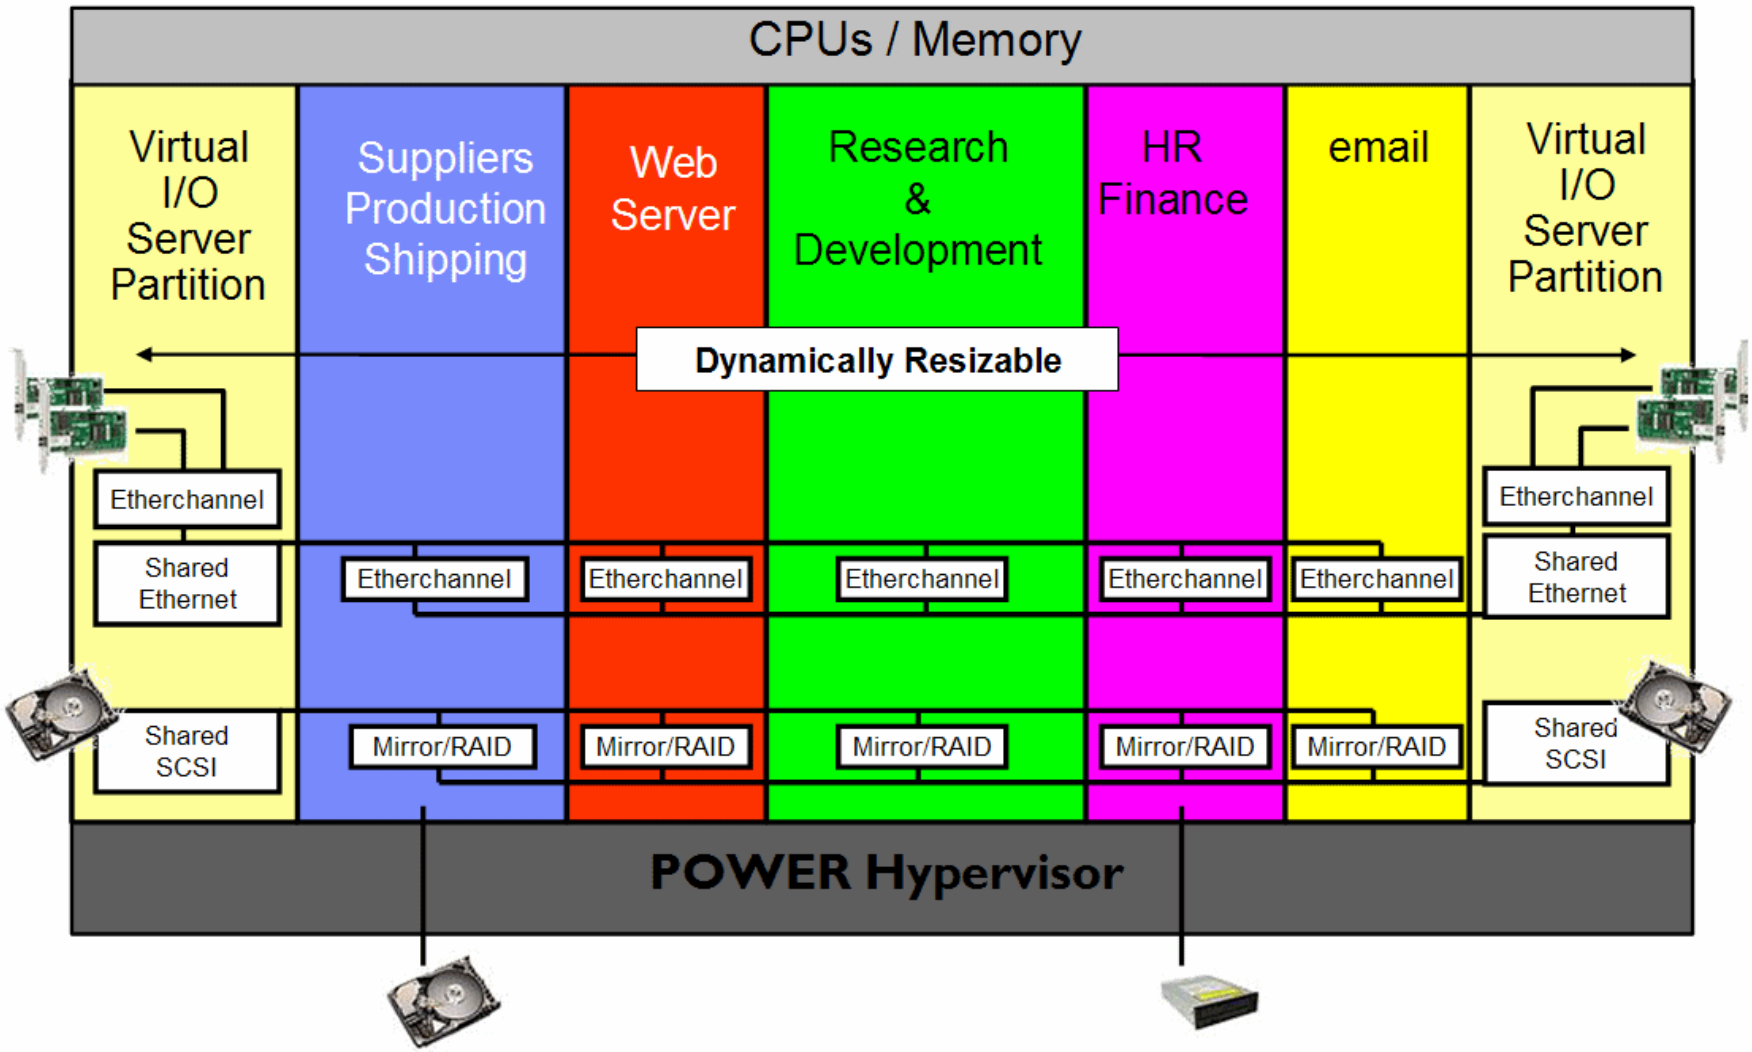

Processing Units

Specified in parts of a processor

Virtual Processors

Shared Processor Pool

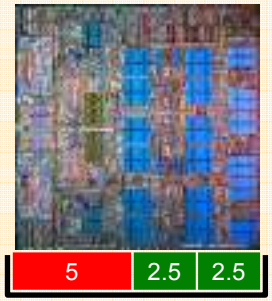

10 milliseconds

Hypervisor Dispatch Cycle

Partition #1 — Capped

.5

1 Way

1 Virtual Processor

Partition #2 — Capped

.5

2 Way

2 Virtual Processors

10 milliseconds

Partition #3 — Uncapped

1.2

4 Way

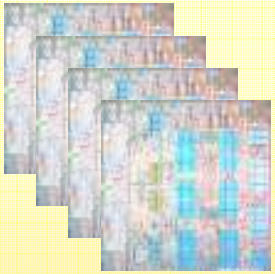

4 Virtual Processors

10 milliseconds

Virtual I/O Server

- **Special type of LPAR used to share I/O devices among partitions on a p5 Server**
- **Created like other LPARs but loaded with the VIO Server code**
- **Physical I/O devices are attached to VIO Server and then shared among other partitions as virtual devices**

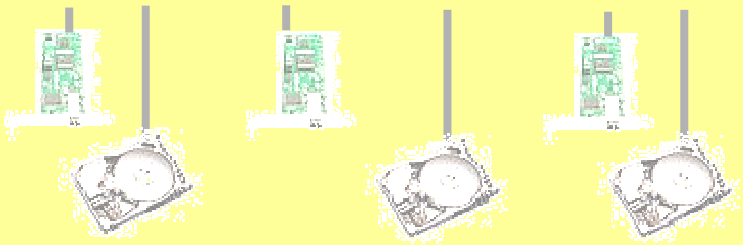

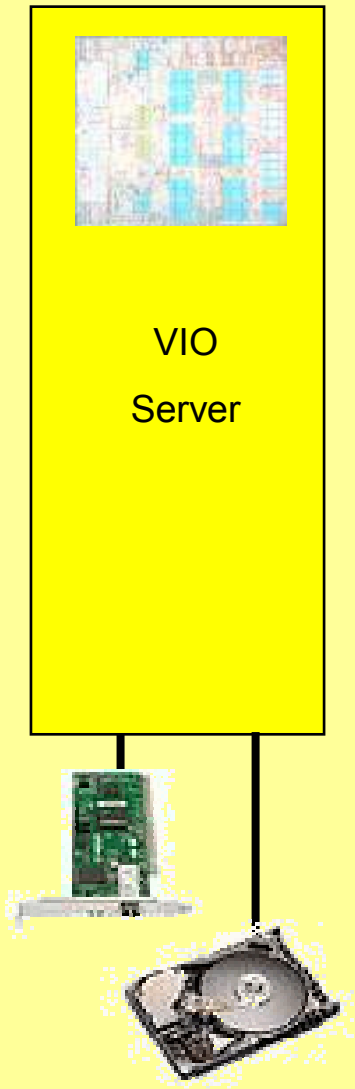

Virtual LAN / Shared Ethernet Adapter

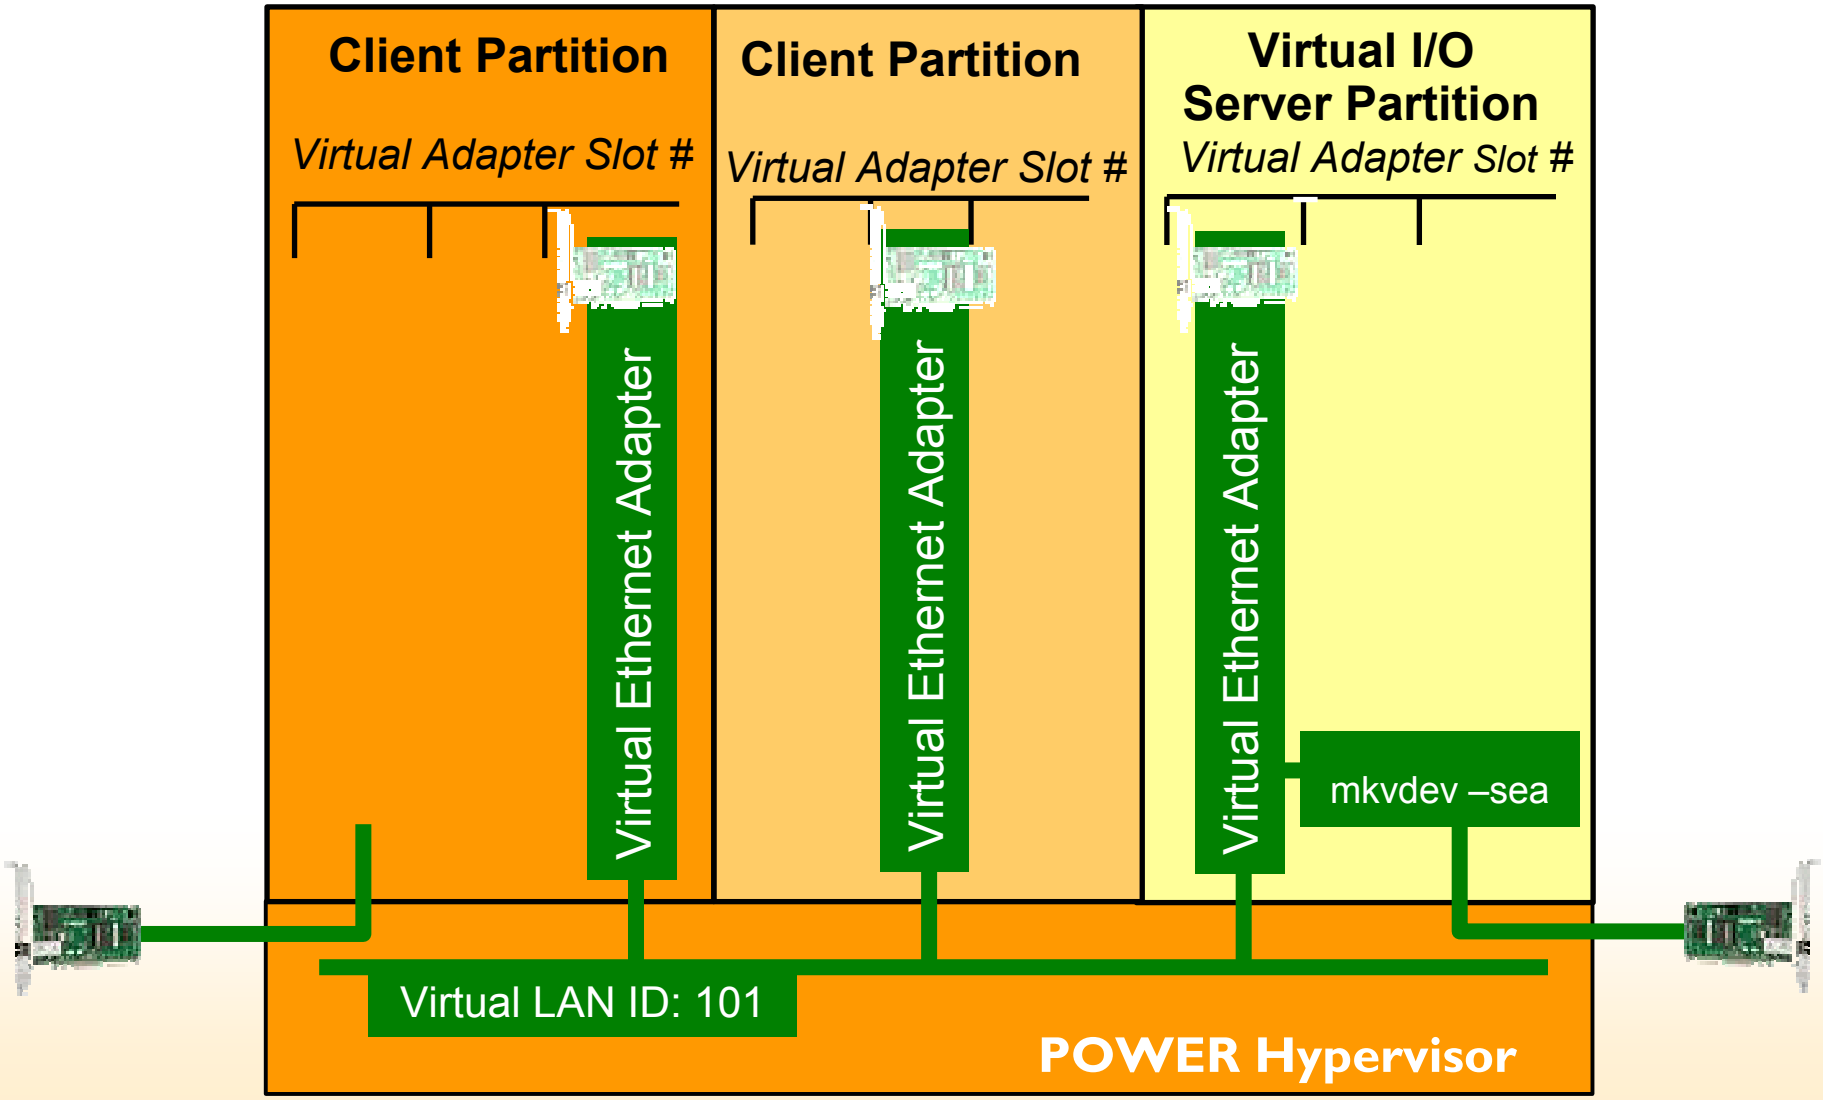

POWER5 / AIX 5.3 LPAR Virtual Ethernet

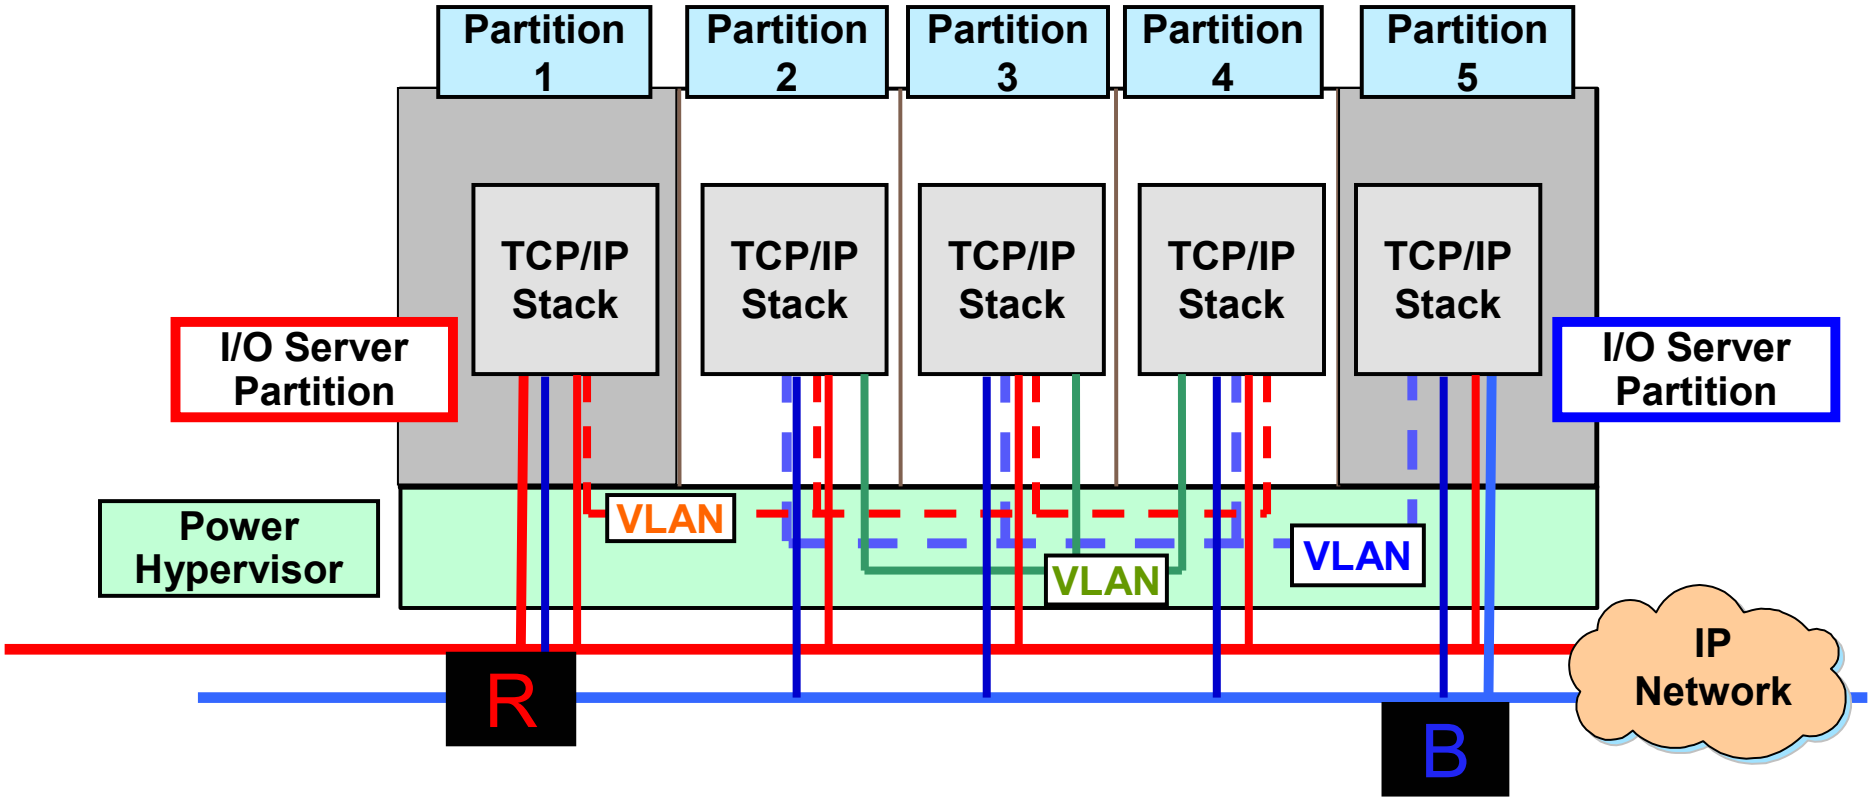

- Inter-partition communication support
- IP forwarding provided by I/O Server partition
- Can have multiple connections per partition

Virtual SCSI

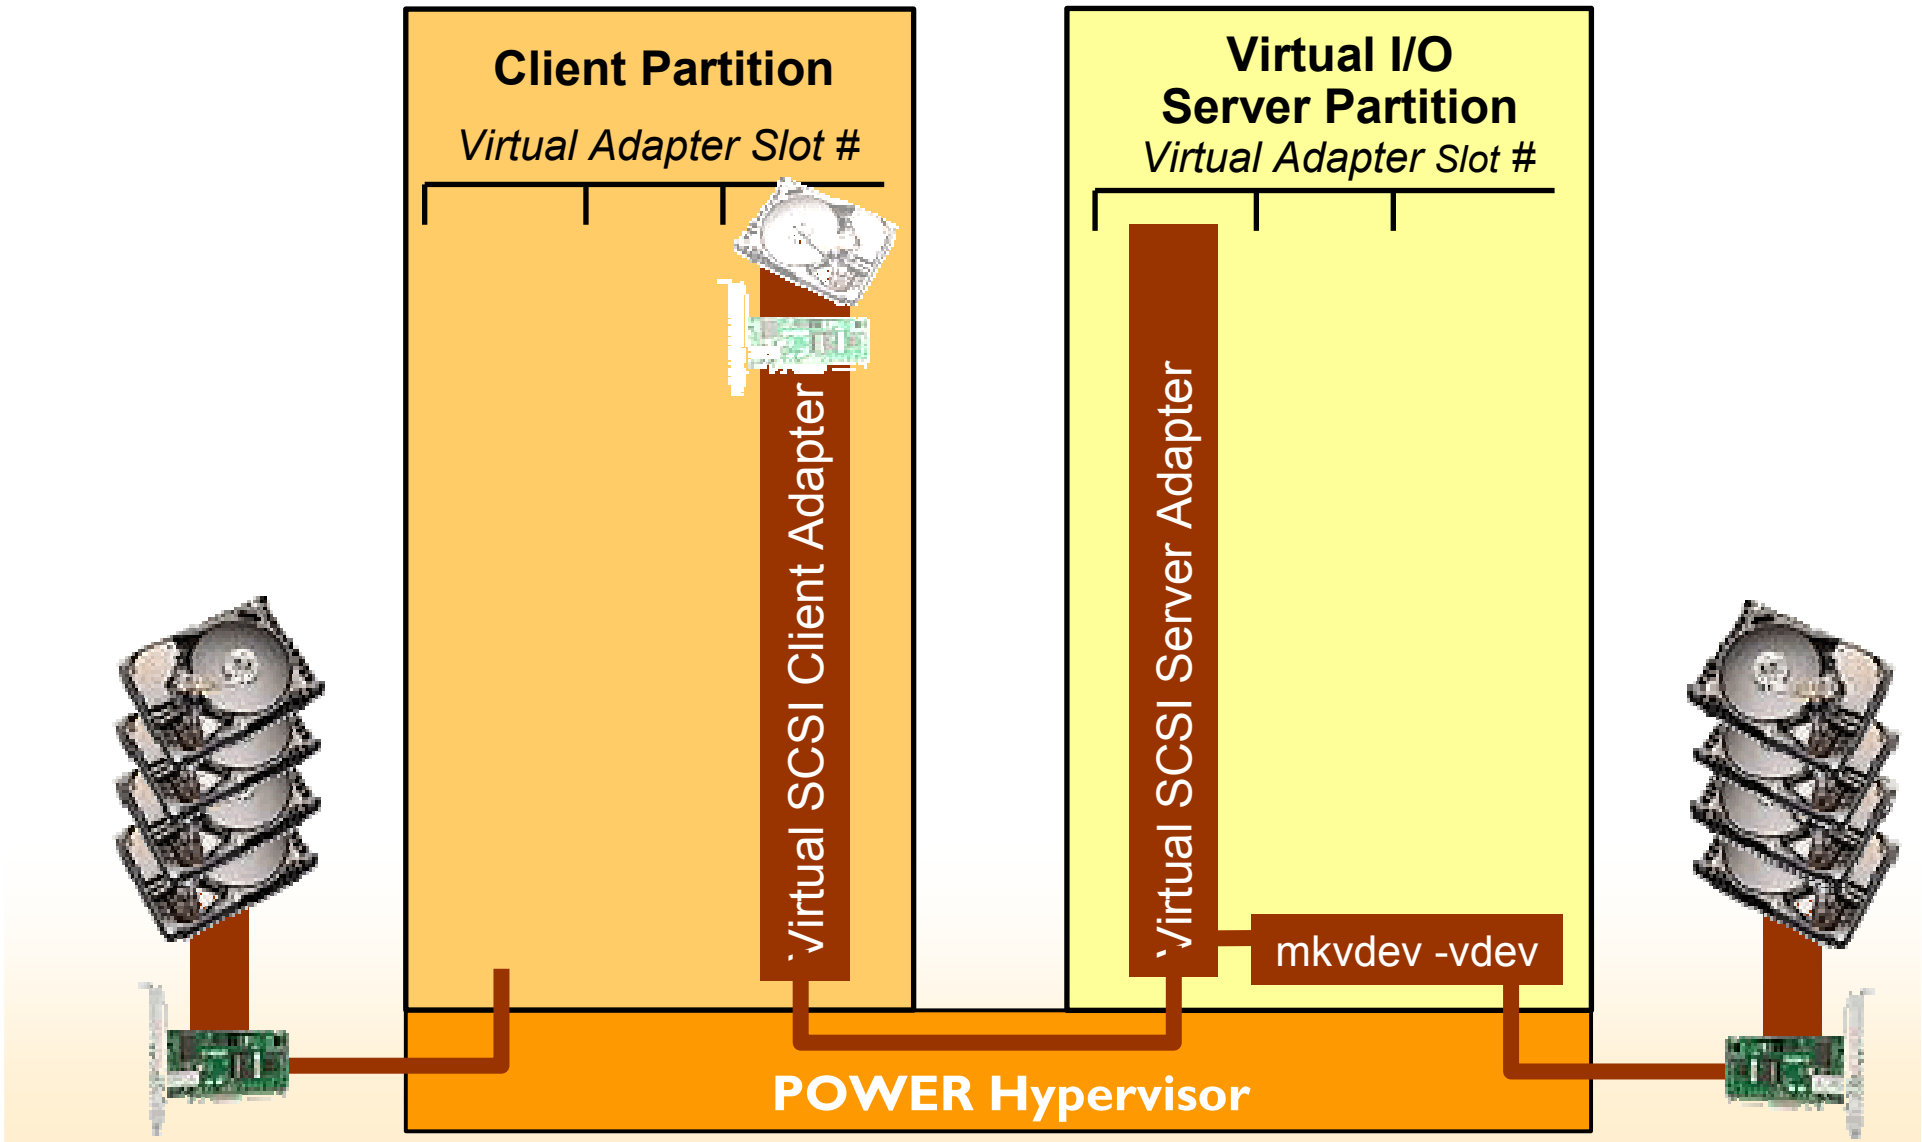

POWER5 / AIX 5.3 LPAR DISK Sharing

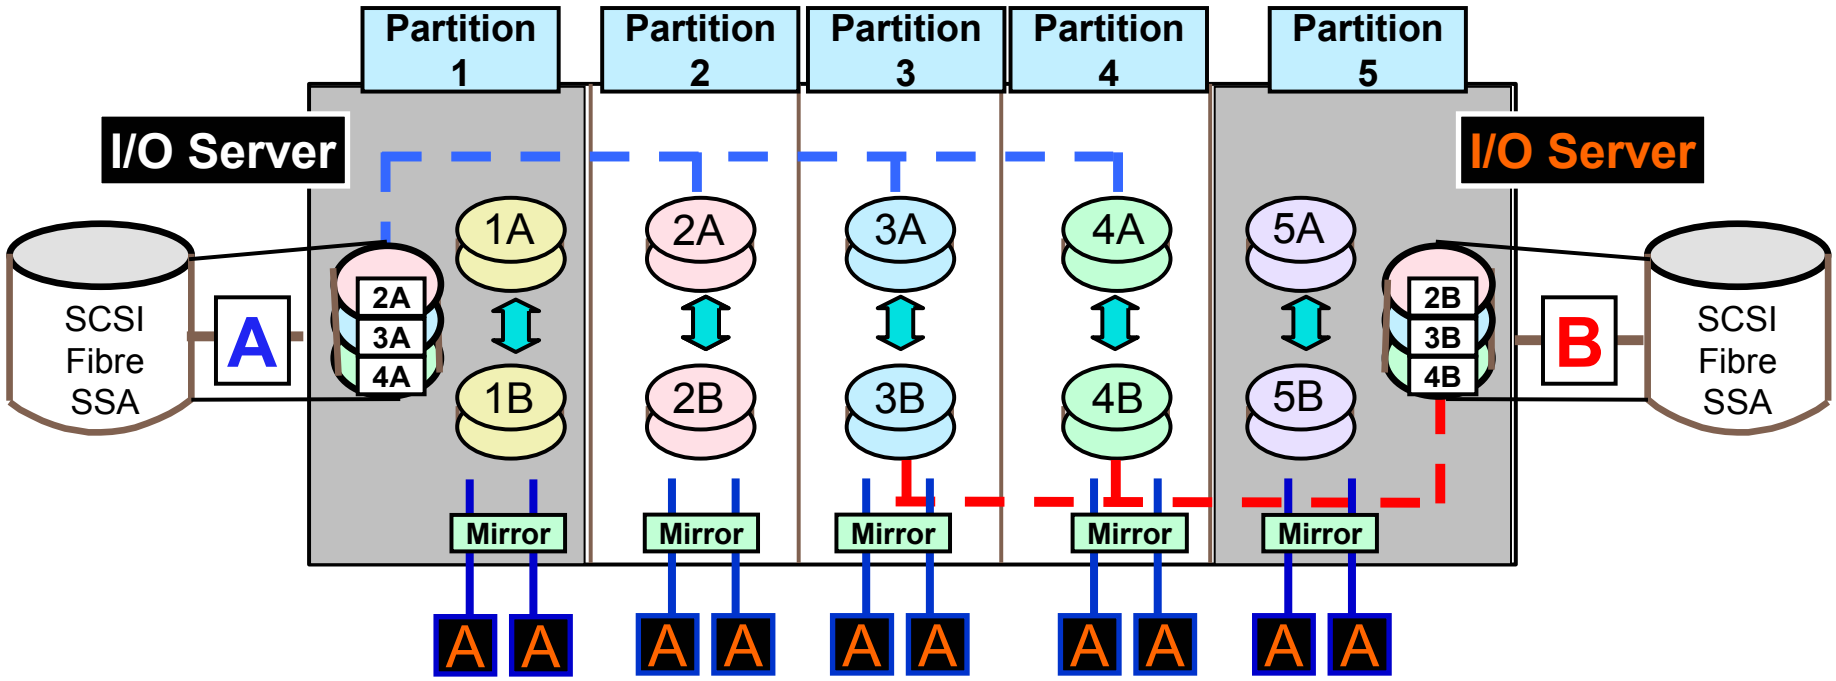

Using LVM mirroring for Availability

One physical drives appears to be multiple logical drives

- LUNs appears as individuals logical drives

Minimizes the numbers adapters

SCSI, SSA, and Fibre supported

VIO - Some Supported configurations …

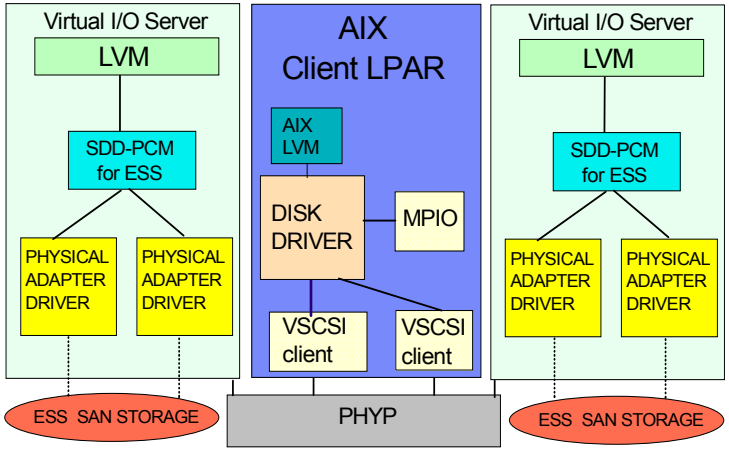

- ESS SAN Storage Configuration
- AIX Client uses MPIO to protect against Virtual I/O Server failures
- Virtual I/O Server uses SDD-PCM to protect against adapter failures

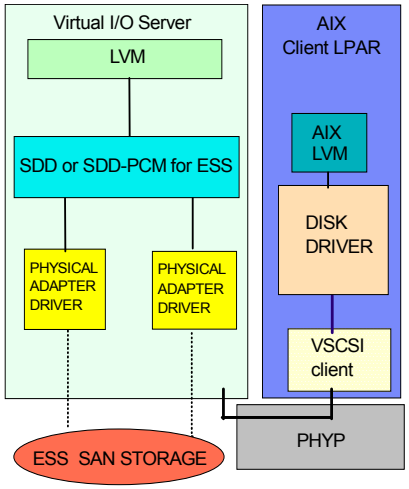

- ESS SAN Storage Configuration
- Virtual I/O Server uses SDD, or SDD-PCM to protect against adapter failures

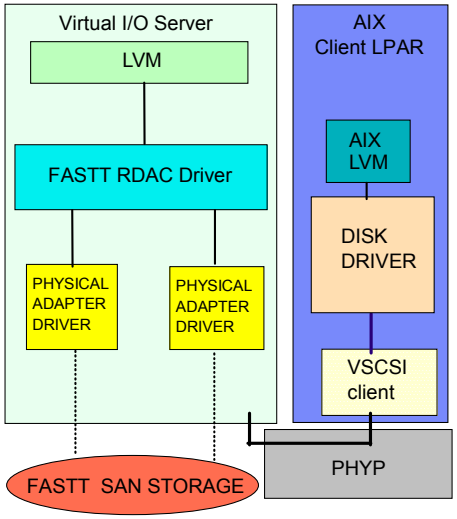

- ESS SAN Storage Configuration
- Virtual I/O Server uses FASTT RDAC Driver to protect against adapter failures

Other configurations with EMC and HDS are being tested currently

Micro-Partitioning Technology

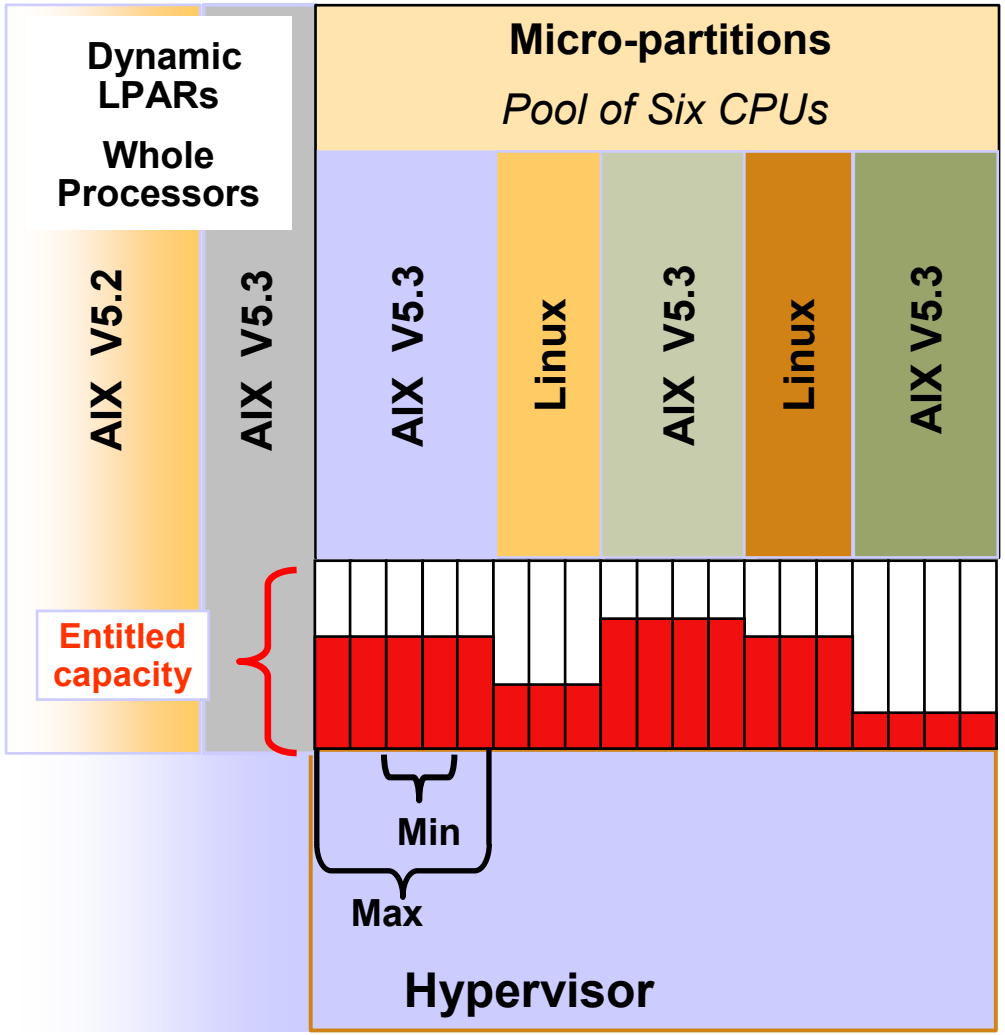

Note: Micro-partitions are available via optional Advanced POWER Virtualization or POWER Hypervisor and VIOS features.

Micro-Partitioning technology allows each processor to be subdivided into as many as 10 "virtual servers", helping to consolidate UNIX® and Linux applications.

Partitioning options

– Micro-partitions: Up to 254*

Configured via the HMC

Number of logical processors

– Minimum / Maximum

Entitled capacity

– In units of 1/100 of a CPU
– Minimum 1/10 of a CPU

Variable weight

– % share (priority) of surplus capacity

Capped or uncapped partitions

*on p5-590 and p5-595

Dedicated Shared Processors

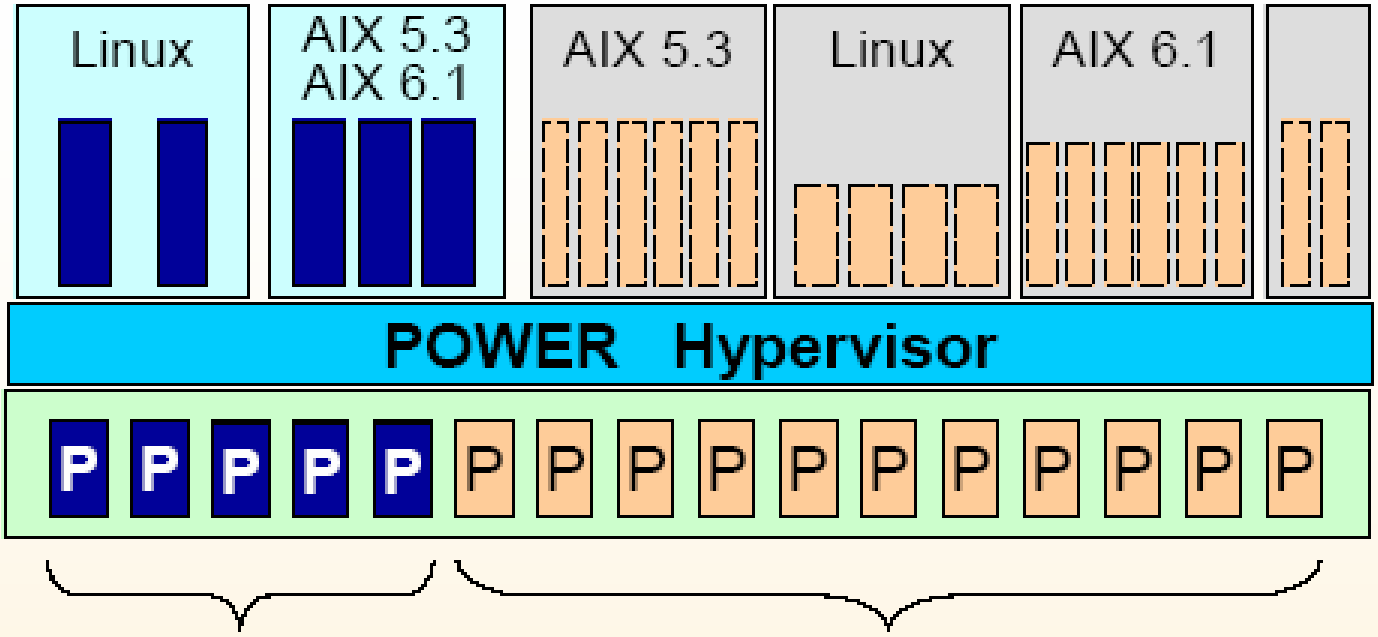

Excess Dedicated Capacity Utilization	▪ Unused capacity in dedicated processor partitions can be "Donated" to shared processor pool ▪ Excess cycles will only be utilized by uncapped partitions that have consumed all of their entitled capacity. ▪ POWER6 Servers

Shared Dedicated Capacity In Action

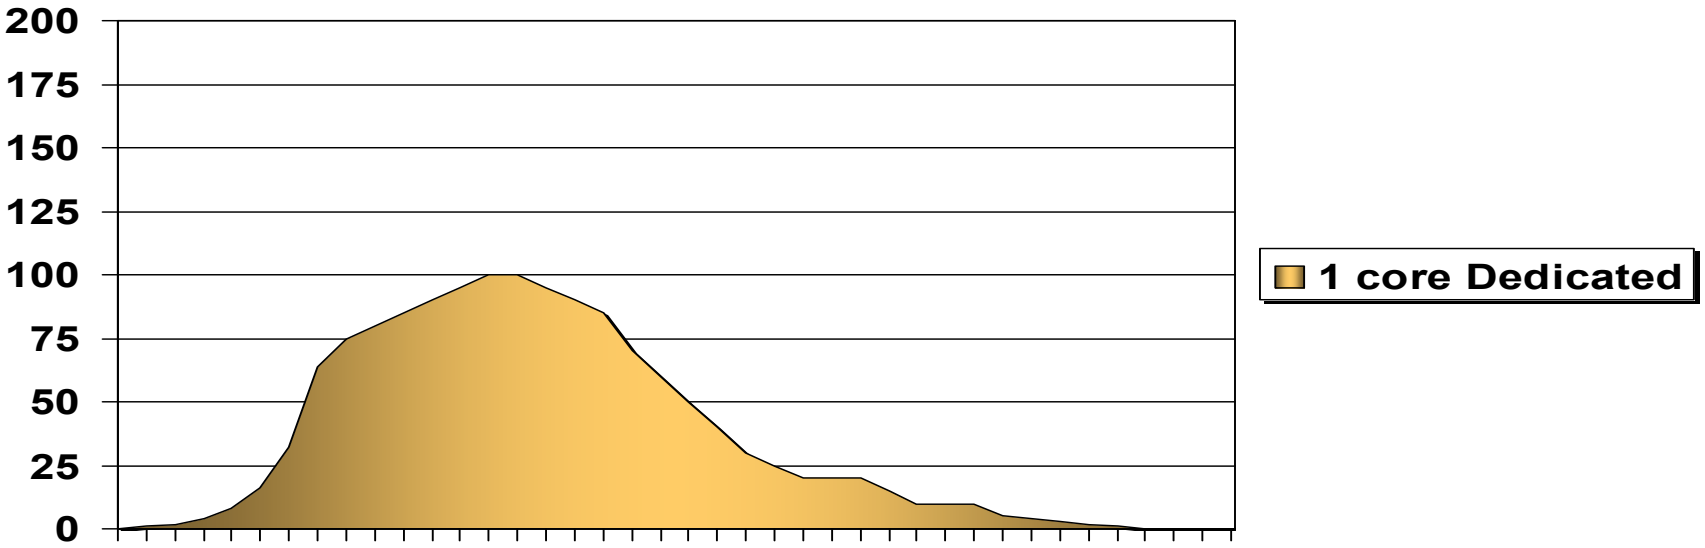

– Consider a 2-core server with 1-core dedicated partition with variable workload (between 0% and 100%)

Shared Dedicated Capacity In Action

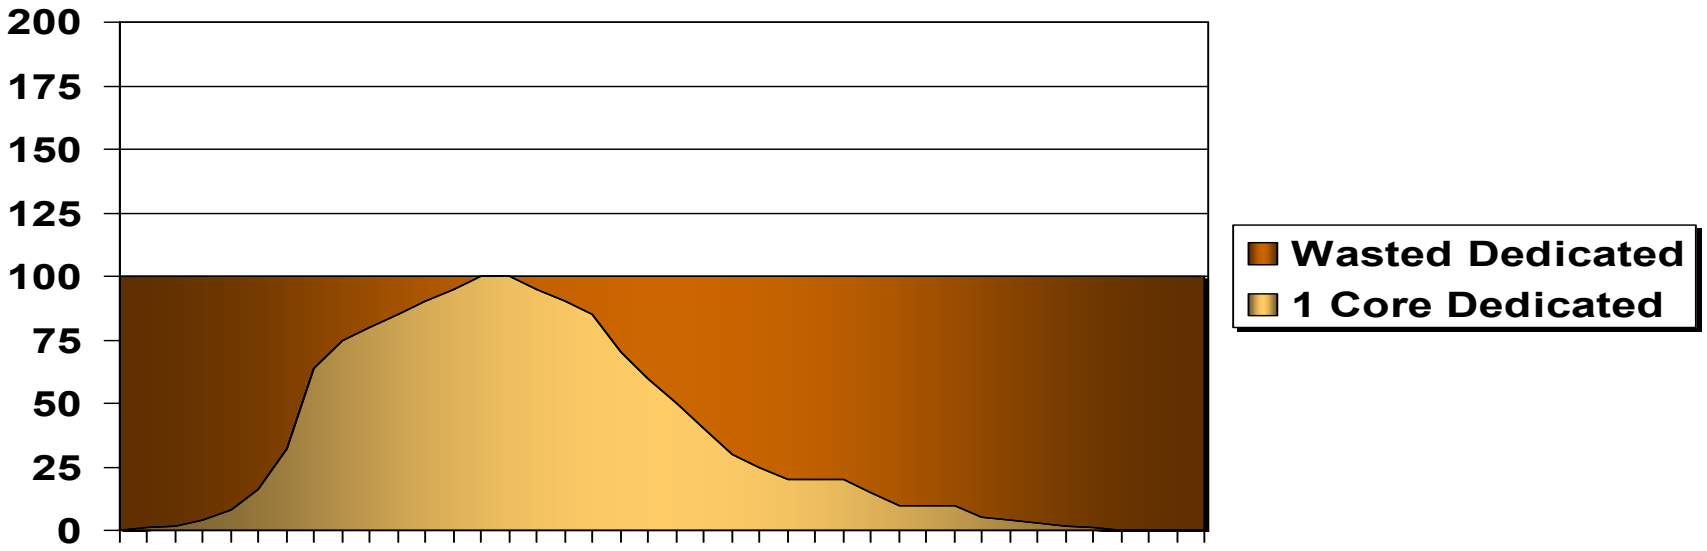

- Consider a 2-core server with 1-core dedicated partition with variable workload (between 0% and 100%)
- The excess capacity on the dedicated processor is wasted

Shared Dedicated Capacity In Action

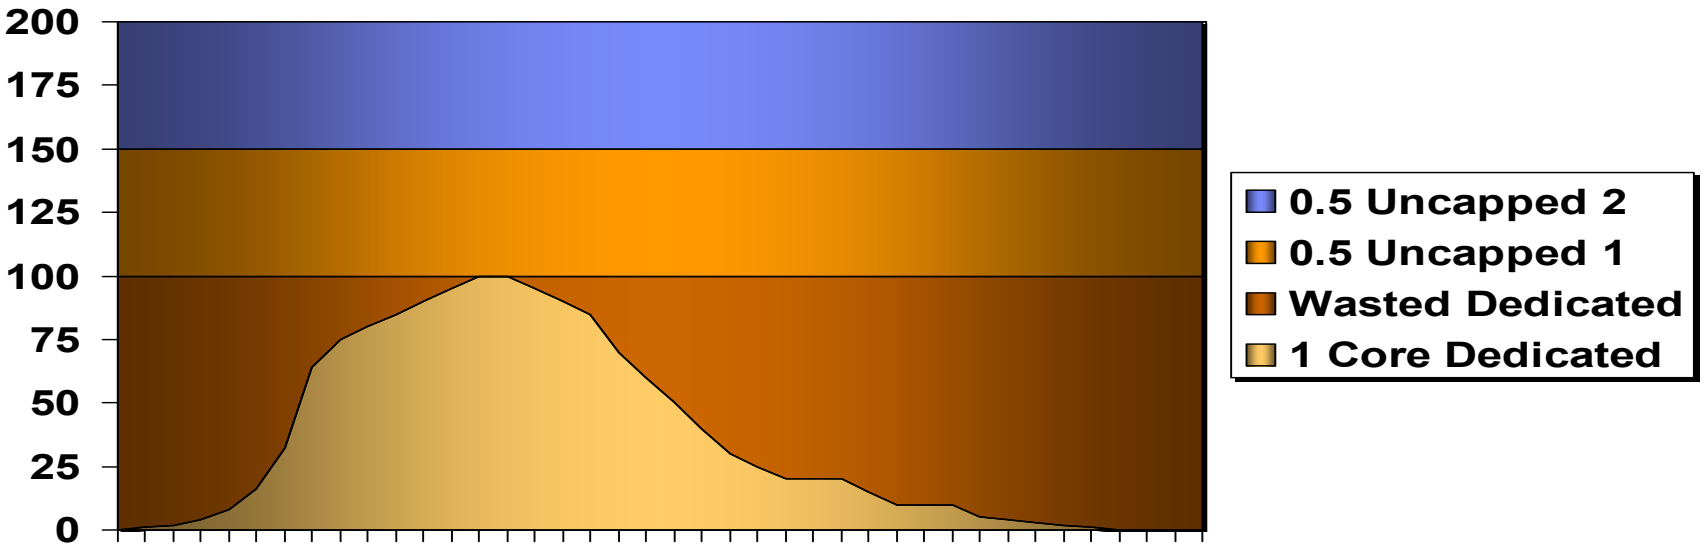

- Add two evenly weighted 1-core uncapped partitions that are CPU bound
- Each uncapped partition will share the remaining physical processor even though each can consume an entire processor

Shared Dedicated Capacity In Action

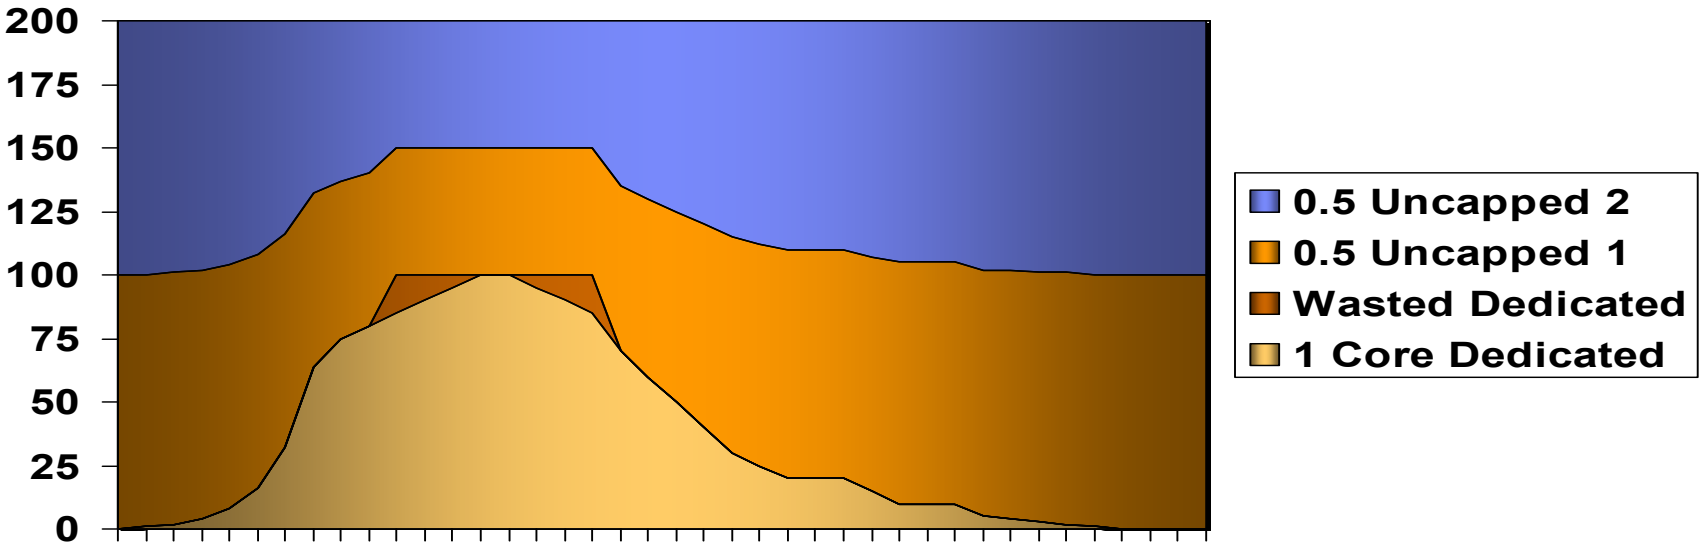

- With the new support, a dedicated partition will donate its excess cycles to the uncapped partitions
- Each uncapped partition will consume an entire processor if available (when dedicated at 0%) and will split a processor when dedicated fully utilized (when dedicated at 100%)
- The total processor capacity in the system is better utilized while the dedicated processor partition maintains the performance characteristics and predictability of the dedicated environment when under load

Integrated Virtual Ethernet - How it Works……

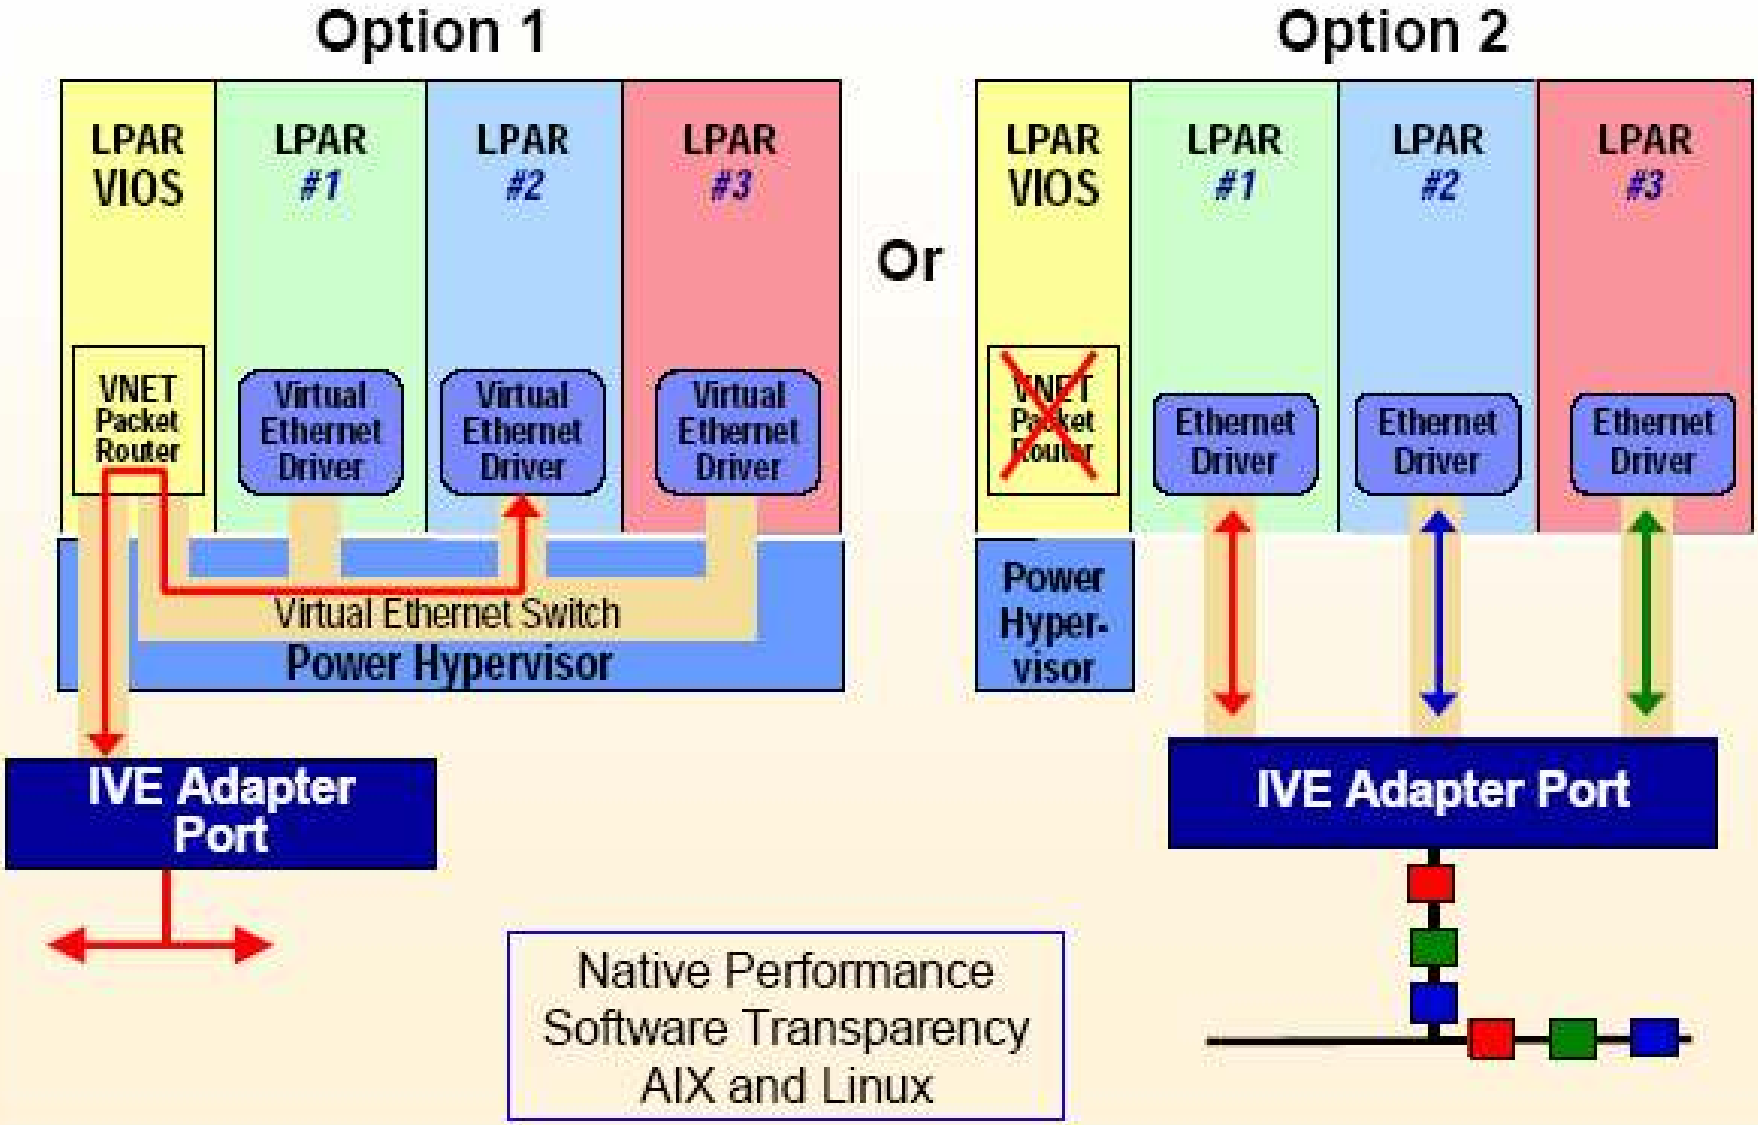

Integrated Virtual Ethernet Overview

Naming

- **Integrated Virtual Ethernet** – External name in marketing and other documentation
- **Host Ethernet Adapter (HEA)** – Name used on user interfaces

New hardware capability

- **Provides enhanced Ethernet connectivity**
  - **No TCEs, PCI protocol, etc**
  - Several integrated Ethernet adapters, called **Logical Ports**.
  - Can be assigned to one or more partitions and/or VIOS partitions.
- **Available on most POWER6 systems**
- **Multiple options of physical, external ports**
  - Dual 1 Gbit copper: 10BASE-T, 100BASE-T, 1000BASE-T
  - Quad 1 Gbit copper: 10BASE-T, 100BASE-T, 1000BASE-T
  - Dual 10 Gbit fibre: 10GBASE-SR or 10GBASE-LR
- **Logical Ports**
  - Up to 32 logical ports, but can also be configured as 1, 2, 4, 8, 16 logical ports
  - Logical port / physical dedicated when assign to VIOS partition
- **Several configuration parameters**
  - All based on tuning performance to match client configuration and environment
  - e.g.: Speed, frame size, duplex

IBM System p Flexible Resource Management

A new method of virtualization on IBM System p: AIX Workload Partitions

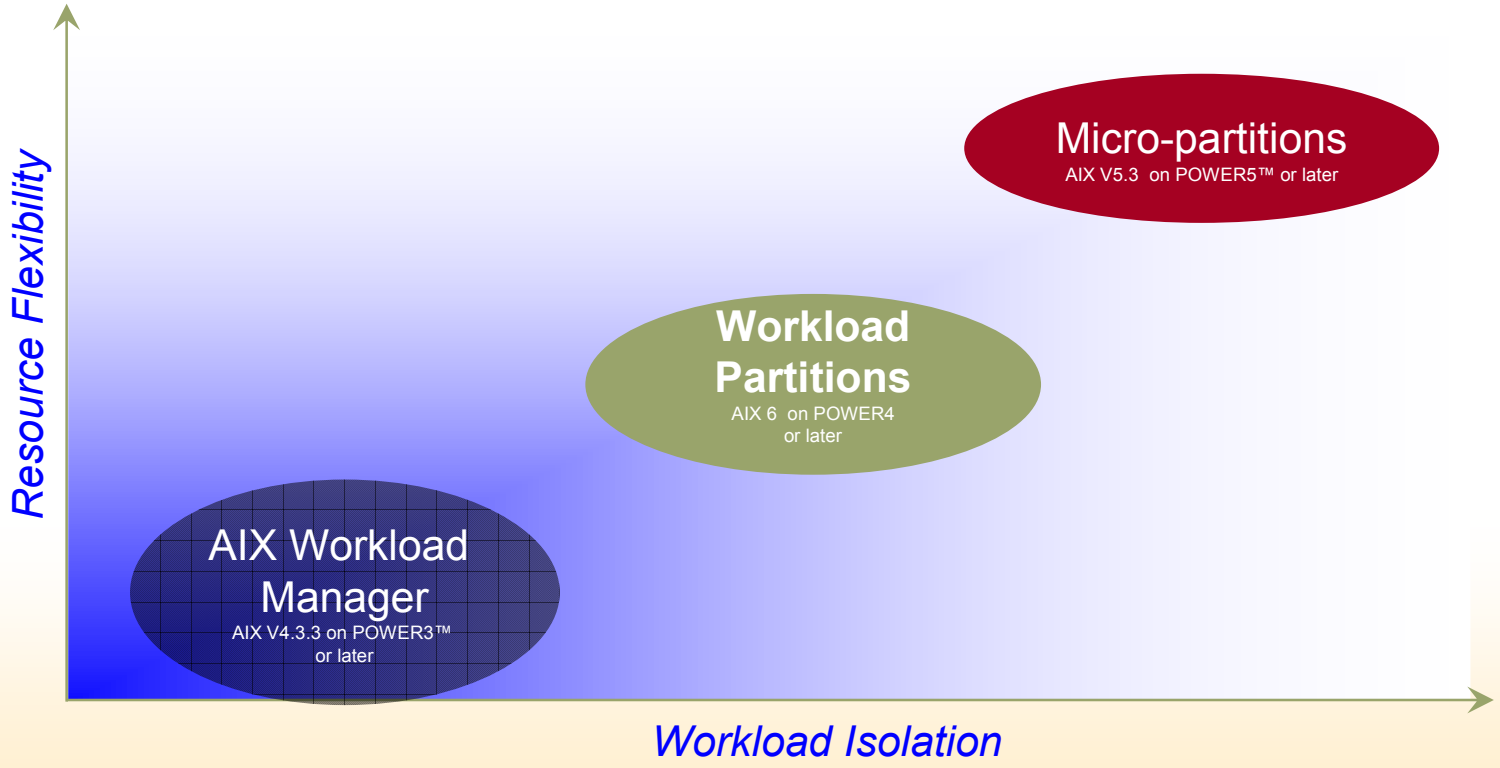

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

IBM System p Announces Two Methods of Mobility

Live Partition Mobility – move a running POWER6 partition …
Live Application Mobility – move a running AIX 6 application …
… From one server to another

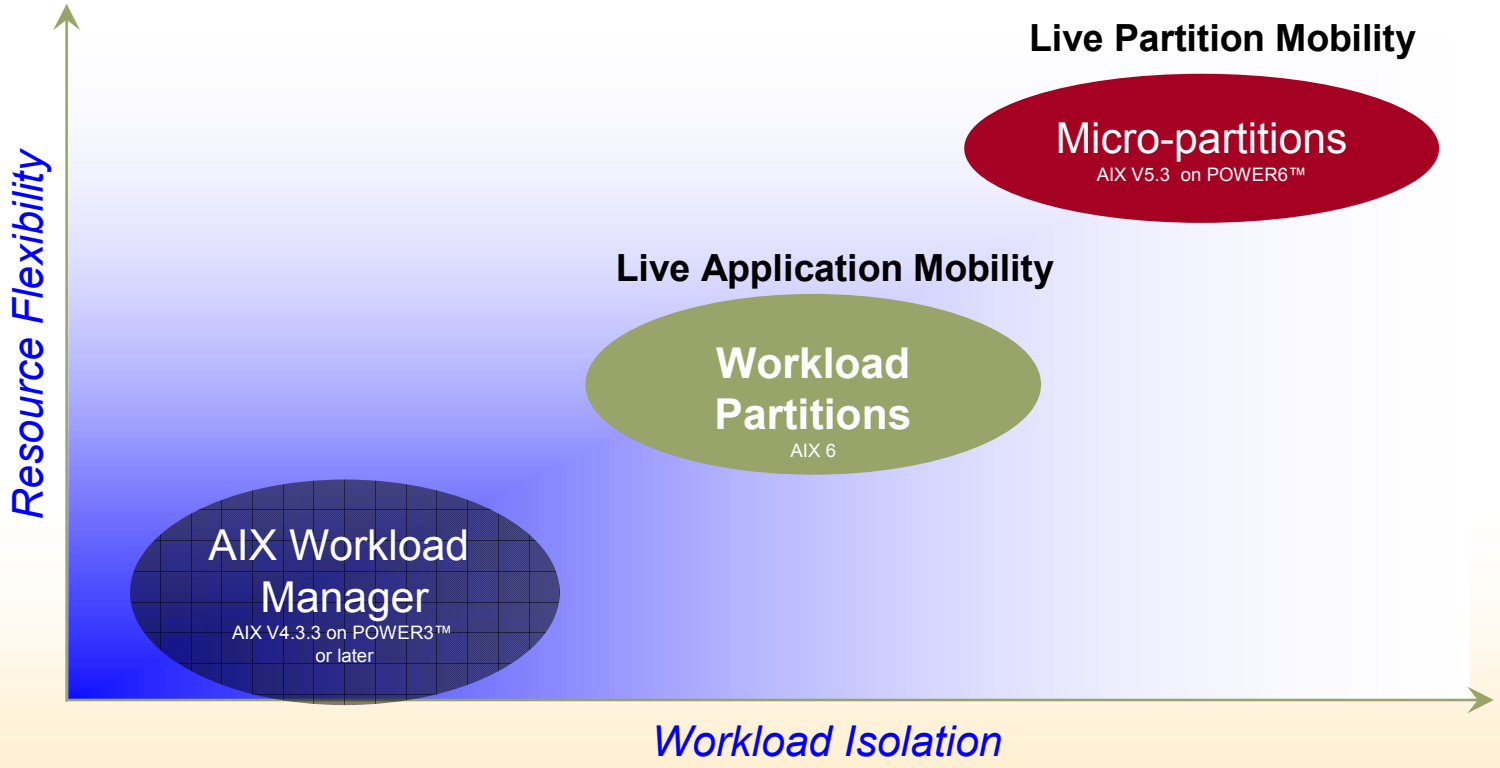

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Live Partition Mobility with POWER6*

Allows migration of a running LPAR to another physical server

- ✓ Reduce impact of planned outages
- ✓ Relocate workloads to enable growth
- ✓ Provision new technology with no disruption to service
- ✓ Save energy by moving workloads off underutilized servers

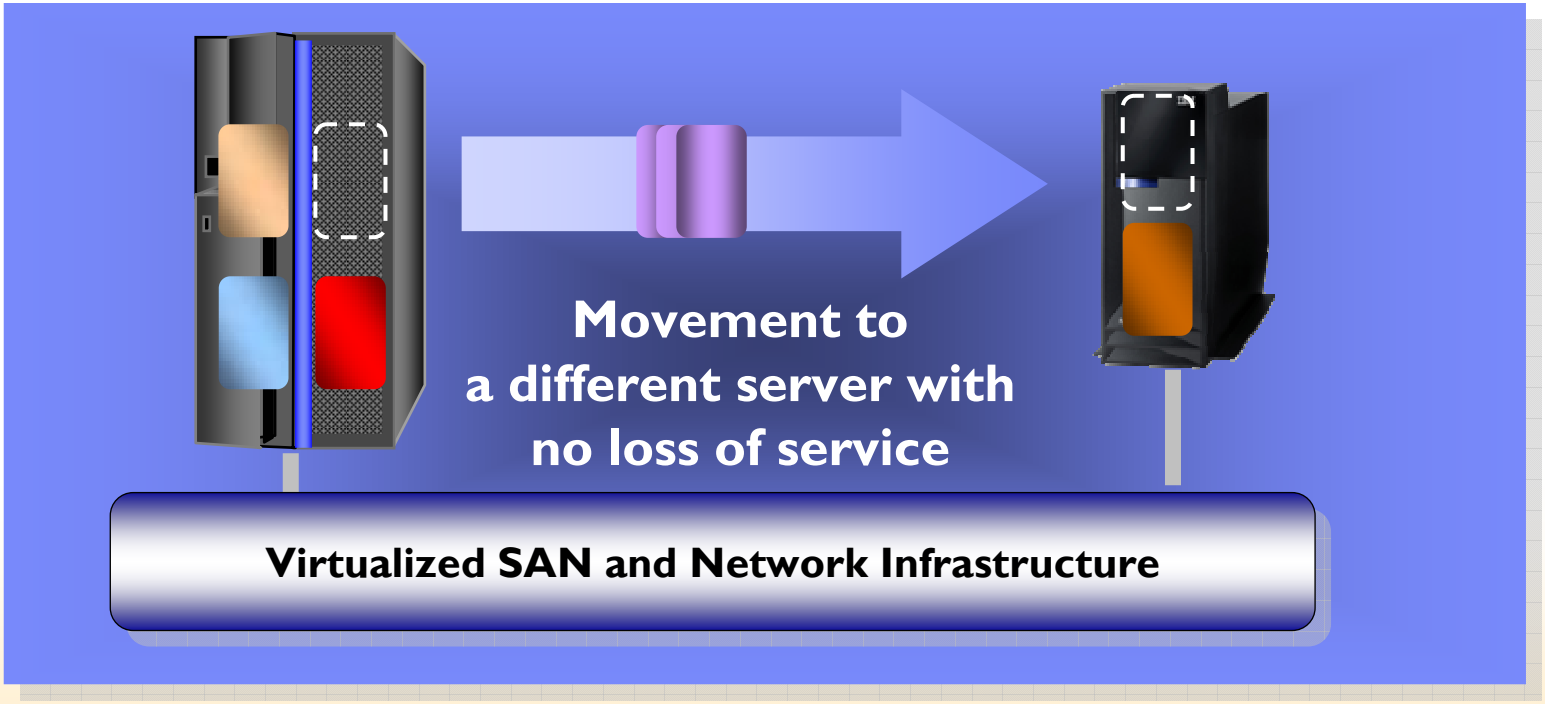

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Continuous Application Availability

With Live Partition Mobility and Live Application Mobility, planned outages for hardware and firmware maintenance and upgrades can be a thing of the past

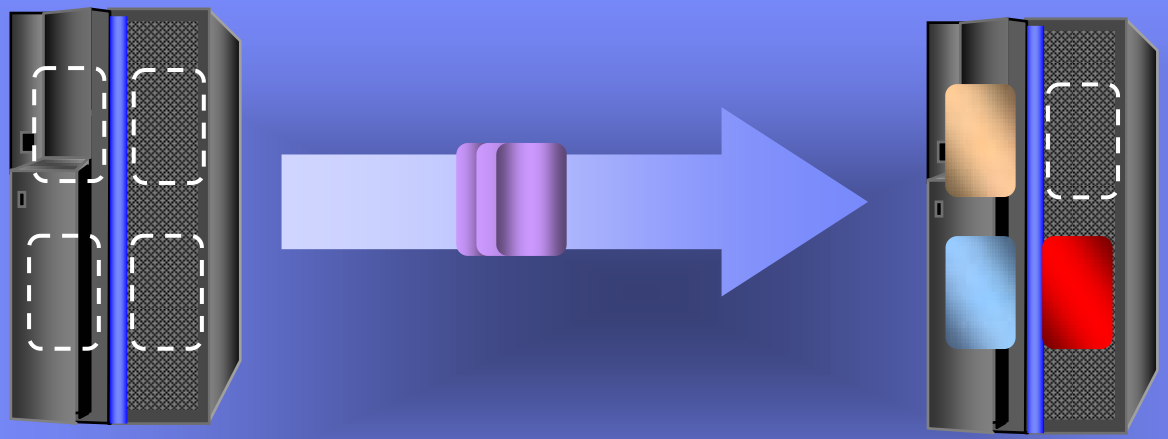

Relocate all partitions from one server to another when performing maintenance. Move the partitions back when maintenance is complete

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Energy Savings

During non-peak hours, consolidate workloads and power off excess servers

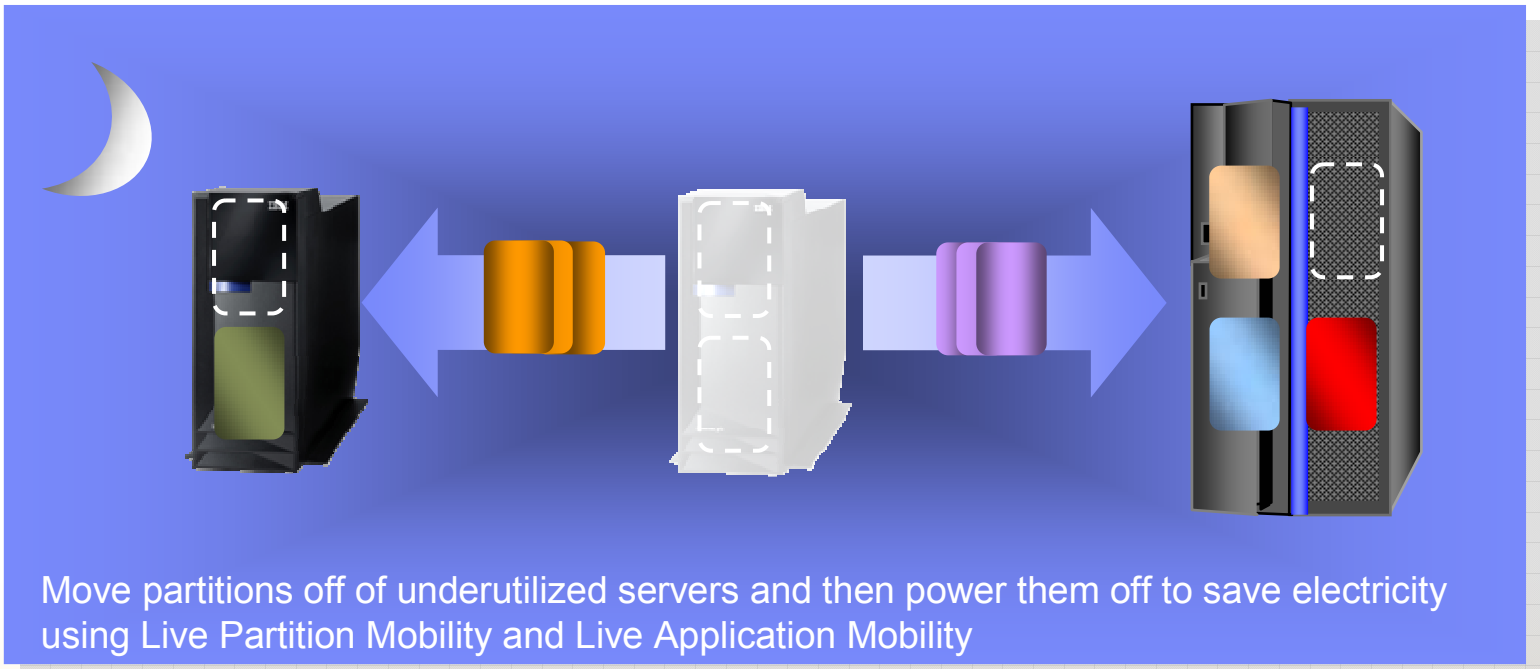

Move partitions off of underutilized servers and then power them off to save electricity using Live Partition Mobility and Live Application Mobility

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Workload Balancing with Live Partition Mobility*

As computing needs spike, redistribute workloads onto multiple physical servers without service interruption

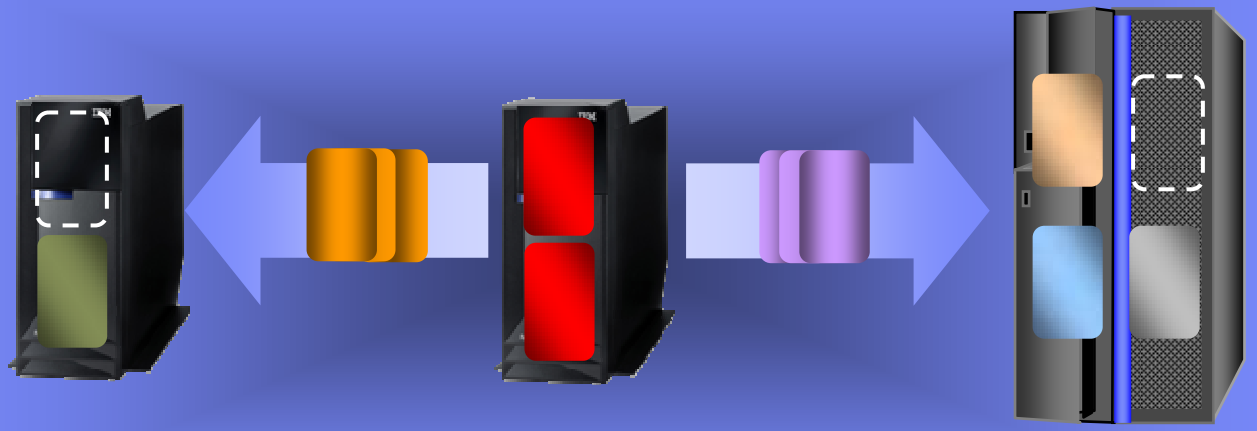

As one server gets overtaxed from a spike in demand, relocate partitions to other servers

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

AIX 6 Workload Partitions

Improved administrative efficiency by reducing the number of AIX images to maintain

- **Software partitioned system capacity**
  - Each Workload Partition obtains a regulated share of system resources
  - Each Workload Partition can have unique network, filesystems and security
- **Two types of Workload Partitions**
  - System Partitions
  - Application Partitions
- **Separate administrative control**
  - Each Workload Partition is a separate administrative and security domain
- **Shared system resources**
  - Operating system, I/O, processor, memory

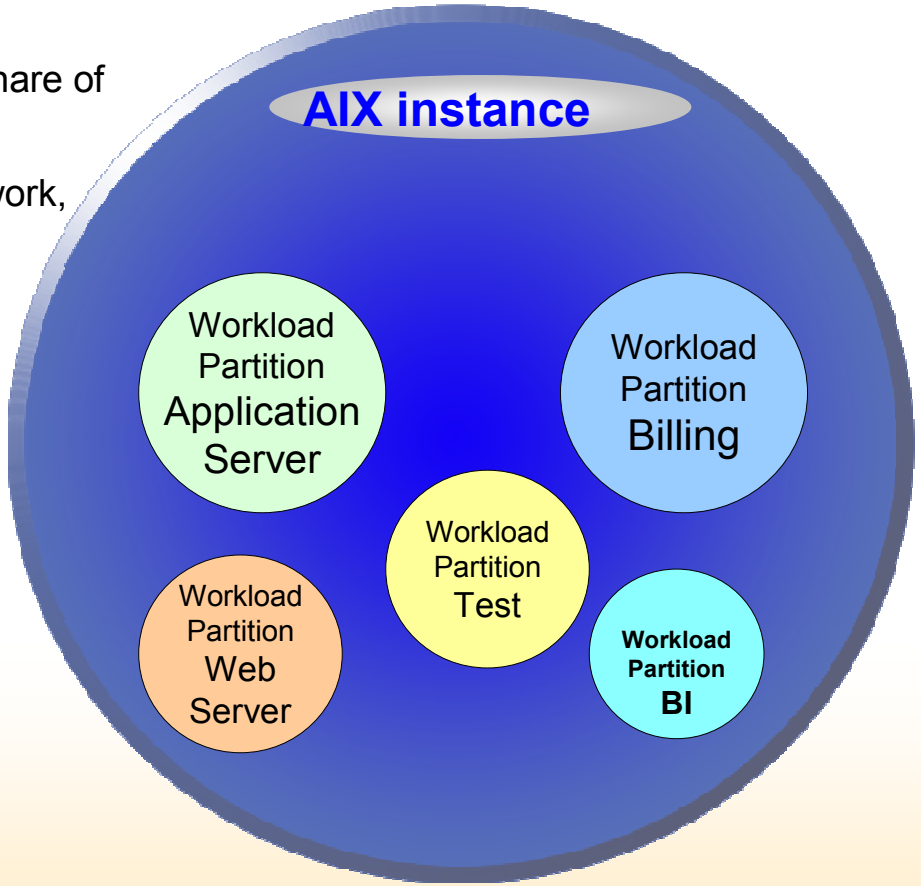

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

AIX 6 Live Application Mobility

Move a running Workload Partition from one server to another for outage avoidance and multi-system workload balancing

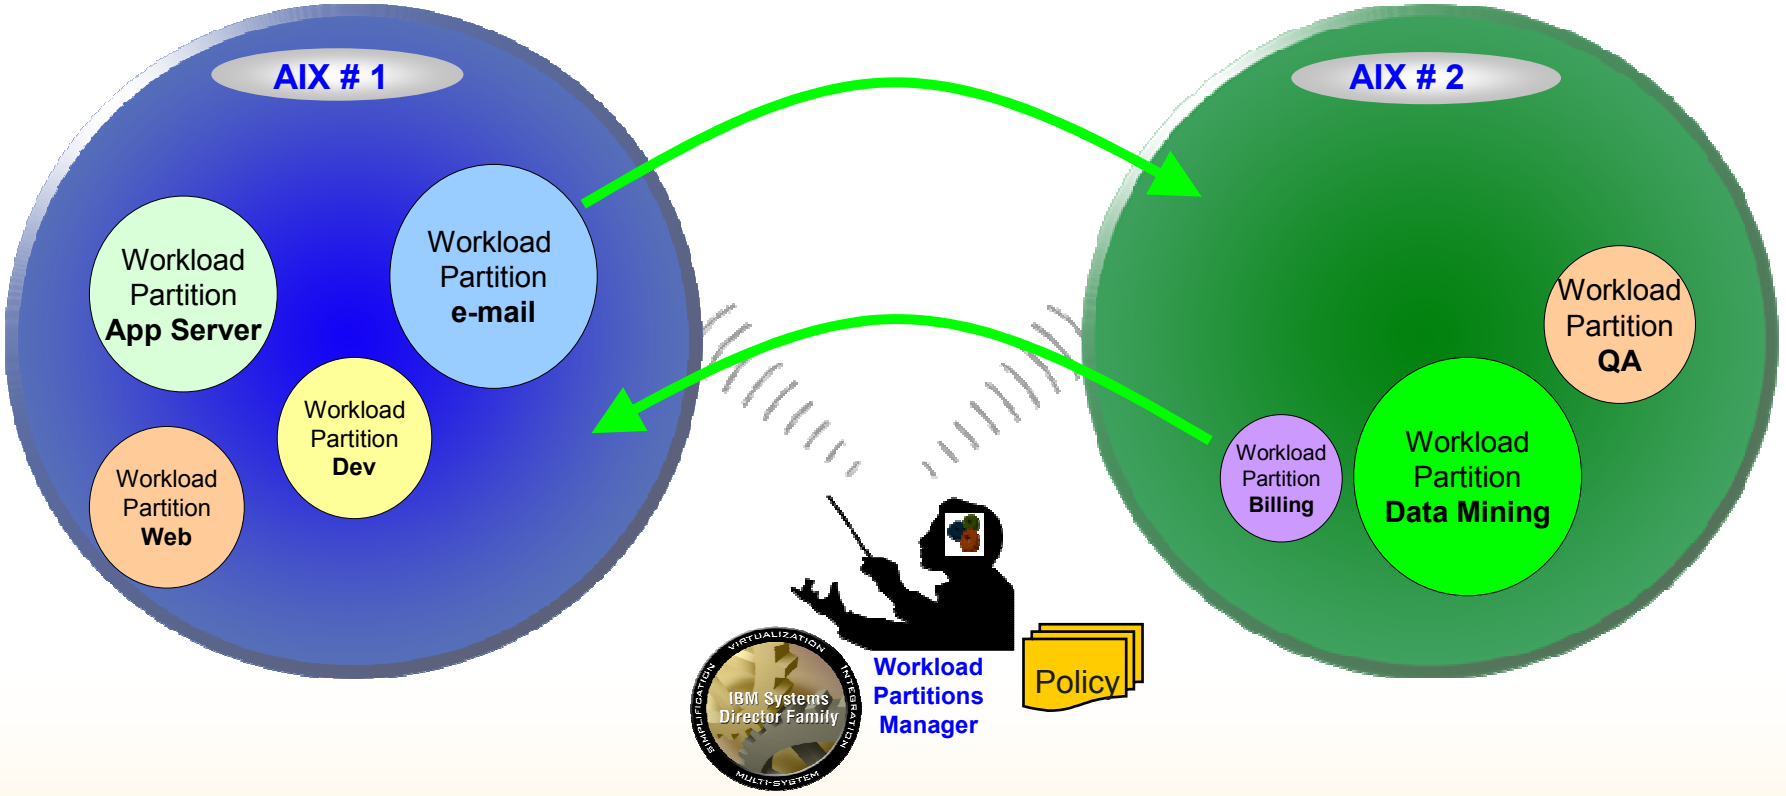

Works on any hardware supported by AIX 6 including POWER5

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

AIX Workload Partitions can be Used in LPARs

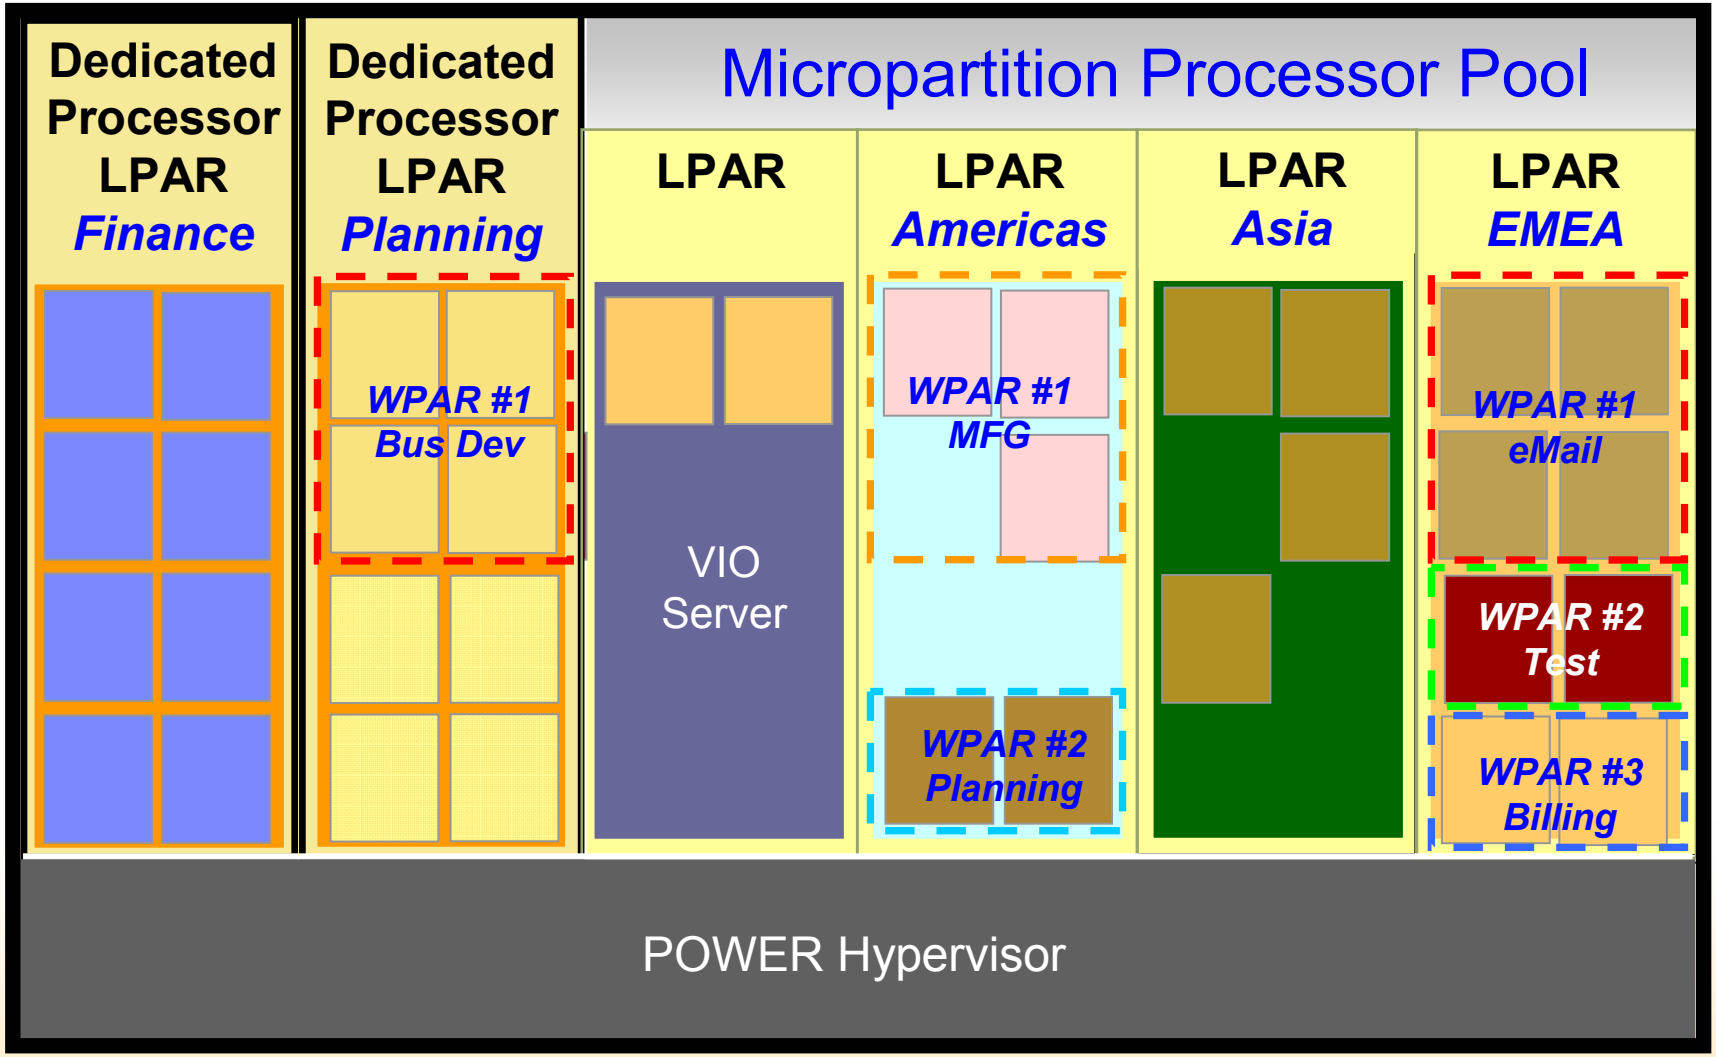

Workload Partitions Manager

- **Management of WPARS across multiple systems**
- **Lifecycle operations**
- **Single Console for:**
  - Graphical Interface
  - Create & Remove
  - Start & stop
  - Checkpoint & Restart
  - Monitoring & Reporting
  - Manual Relocation
  - Automated Relocation
  - Policy driven change
- **Infrastructure Optimization**
- **Load Balancing**

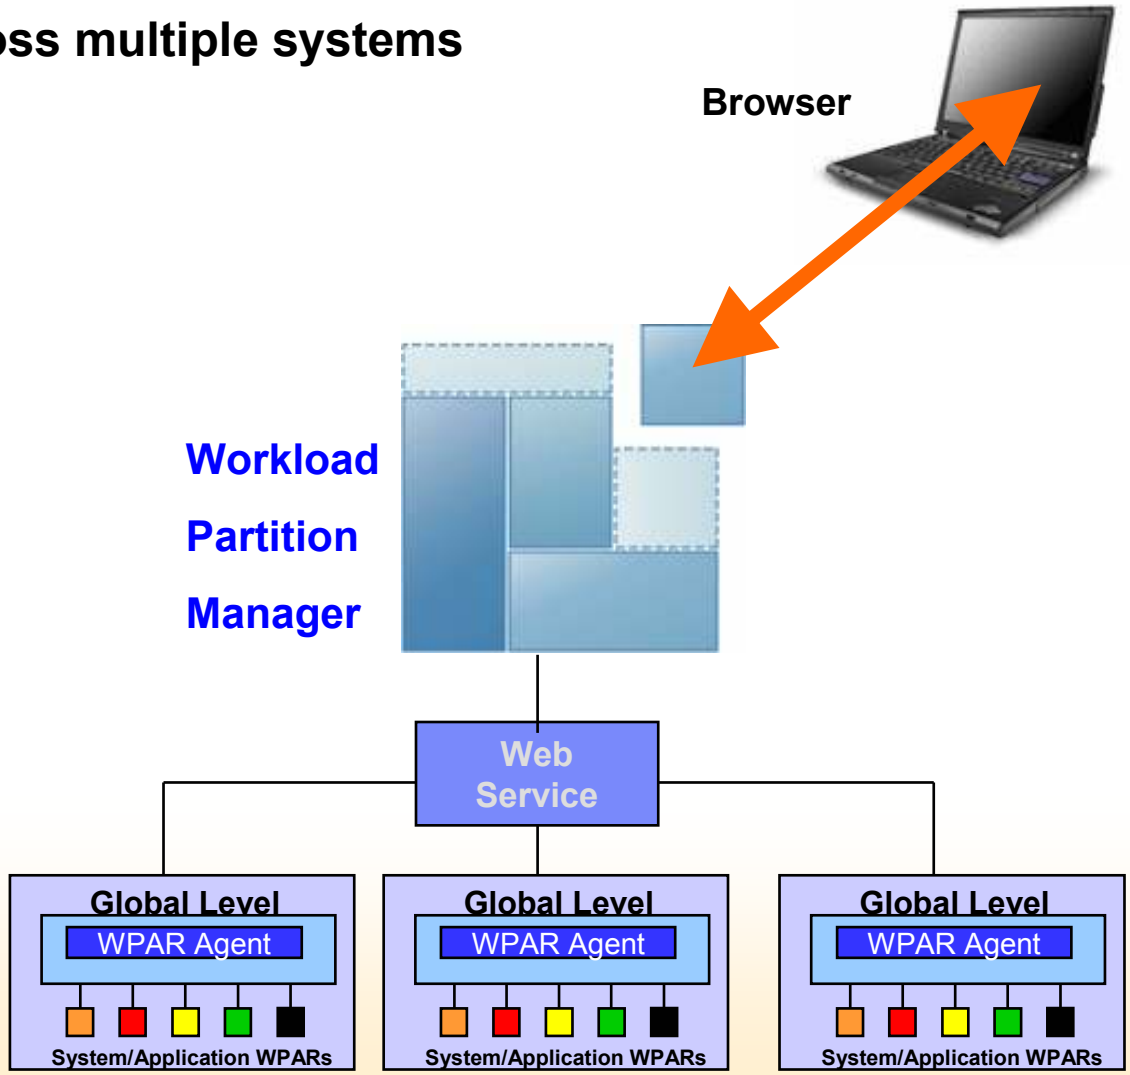

Planned CoD Offering Evolution for POWER6

Today	2007 / 2008 POWER6
Permanent	**Permanent**
CoD Processors: One processor increment Memory: 1 GB increment	**CoD** Processors: One processor increment Memory: 1 GB increment
Temporary	**Temporary**
On/Off CoD Activations: Manual Utilization Reporting Required (Contract) Post-pay Integrated into Capacity BackUp offering	**On/Off CoD** Activations: Manual Utilization Reporting Required (Contract) Post-pay Integrated into Capacity BackUp offering
Capacity BackUp (CBU) CBU offering for p5-590 / p5-595	**Capacity BackUp (CBU)** CBU offering for POWER6 High End
Reserve CoD Pre-pay usage for blocks of time Charges based on measured workload Operates within the shared pool	**Utility CoD*** Post pay or pre-pay (One Processor Minute) Charges based on measured workload Operates within the shared pool
Trial CoD Standard Exception Web-based distribution	**Trial CoD** Standard Exception Web-based distribution

* Utility CoD billing is 100 minutes

APV Enhancements – VIOS 1.4

Virtual I/O Server
- Share Ethernet, SCSI and Fibre Channel disks

Expanded support for industry standard protocols

– Centralized user management support through LDAP
– Support for monitoring and management through SNMP
– Automatic registration of VLANs to external switches through GVRP support

Expanded storage and subsystem support

–Support for nSeries and NetApp subsystems
–iSCSI and fibre channel attach
–Interconnect support for SAS (for disk attach), SATA (optical controller)

APV Enhancements – VIOS 1.4

Virtual I/O Server
- Share Ethernet, SCSI and Fibre Channel disks

Systems management enhancements

– System Planning and Deployment tool upgrades to simplify the process of planning and deploying System p, the Virtual I/O Server and virtual I/O devices.

– CLI enhancements for usability, security, maintenance, user management

Tivoli agent support

Tivoli

–Support for automatic backup and restore of VIOS partitions through IBM Tivoli Storage Manager

–Support for accounting and usage metric collection and analysis through IBM Tivoli Usage and Accounting Manager

–Support for VIOS and System p resource discovery through IBM Tivoli Application Dependency Discovery Manager

System p Advanced POWER Virtualization (APV) offers **business advantages** compared to VMware for clients doing server consolidation

Virtualization

Virtualization capability	*APV on System p**	*VMware Infrastructure 3 Enterprise on x86*	*Business Benefit of APV*
Partition scalability **System scalability**	**64 CPUs, 2 TB RAM** **64 CPUs, 2 TB RAM**	**4 CPUs,16 GB RAM**[1] **32 CPUs, 64 GB RAM**	Helps improve TCO from consolidating more workloads
Dynamic Logical Partitioning	**Yes**	**VM reboot**[2]	Adapt to changes without downtime
Security/fault isolation	**CAPP/EAL 4+**	**CAPP/EAL 2**[3]	Secured environment for mission critical applications
Support for dedicated I/O	**Yes**	**No**[4]	Superior performance for I/O intensive workloads
Capacity on Demand integration	**Reserve CoD**	**No**[1]	Add capacity when/where needed, turn it off when not required
Partitions per CPU	**10**	**8**[1]	Management flexibility
Price	**Standard on JS21, p5-590, p5-595** **$590/CPU: p5-505 - p5-550** **$990/CPU: p5-560Q** **$1130/CPU: p5-570 & p570**	**$5,750 + support: 2 sockets**[5]	Helps lower acquisition cost

*Advanced POWER virtualization is standard on the System p5 590 and 595. It is an optional feature on all other System p5 servers, except the System p5 185. Also available on the BladeCenter JS21 (1) IBM conclusion based on analysis of www.vmware.com (2) Source: http://pubs.vmware.com/vi301/wwhelp/wwhimpl/js/html/wwhelp.htm (3) Sources: VMware test at http://www.commoncriteriaportal.org/public/files/epfiles/ST_VID10056-VR.pdfIBM certification info at http://www-03.ibm.com/servers/aix/products/aixos/certifications/index.html, picture from http://www.cesg.gov.uk/site/iacs/index.cfm?menuSelected=1&displayPage=13 (4) Source: http://h20331.www2.hp.com/ActiveAnswers/downloads/vmwareESXserver_virtualize_ProLiant_1005.pdf (5) http://www.vmware.com/pdf/vi_pricing.pdf

System p Advanced POWER Virtualization (APV) offers **business advantages** compared to HP Integrity Virtual Machine (IVM) for clients doing server consolidation

Virtualization

Virtualization capability	*APV on System p**	*HP Integrity Virtual Machines 2.0*	*Business Benefit of APV*
Partition scalability	**64 CPUs 2TB RAM**	**4 CPUs (max)** [1] **1 CPU (recommended)** [2] **64GB RAM**[1]	Helps improve TCO from consolidating more workloads
Dynamic Logical Partitioning	**Yes**	**VM reboot**[3]	Adapt to changes without downtime
Security/fault isolation	**CAPP/EAL 4+**[4]	**No Certification**	Secured environment for mission critical applications
Support for dedicated I/O	**Yes**	**No**[5]	Superior performance for I/O intensive workloads
Dynamic Processor Sharing	**Yes**	**No**[6]	Highly flexible configuration, with automatic deployment of machine resources where and when they're needed
Externally Published References	**94**[7]	**1**[8]	Rest assured knowing that APV is client proven and running on 40% of all System p5 CPUs

*Advanced POWER virtualization is standard on the System p5 590 and 595. It is an optional feature on all other System p5 servers, except the System p5 185. Also available on the BladeCenter JS21 (1) Source: http://h18004.www1.hp.com/products/quickspecs/12715_div/12715_div.HTML (2) From Best Practices for Using Integrity Virtual Machines: "Scheduling, memory access, and resource contention issues are all easier if only one processor is involved ... The CPU resources of a VM Host may be prematurely exhausted by the definition of a virtual SMP." http://docs.hp.com/en/9983/BestPractices2.2.pdf p6 (3) Sources: HP states "All attributes and resources can be changed statically, so that changes take effect when the virtual machine is next restarted. Some attributes and resources can also be changed dynamically." http://docs.hp.com/en/T2767-90024/T2767-90024.pdf p 155. Although they do not provide a full listing of those resources which can be changed dynamically, they do state that change of virtual CPU cannot happen without VM Reboot http://forums1.itrc.hp.com/service/forums/questionanswer.do?threadId=1028522 (4) IBM certification info at http://www-03.ibm.com/servers/aix/products/aixos/certifications/index.html (5) Source http://h18004.www1.hp.com/products/quickspecs/12715_div/12715_div.HTML (6) "Each of virtual processor from a virtual SMP must be allocated resources from separate physical processors – two virtual processors from the same virtual machine cannot be scheduled on the same physical processor." http://docs.hp.com/en/9983/BestPractices2.2.pdf p6 (7) Complete list of IBM System p Virtualization References at http://www-306.ibm.com/software/success/cssdb.nsf/advancedsearchVW?SearchView&Query=(Virtualization)+AND+[WebSiteProfileListTX]=eserverpseries&site=eserverpseries&frompage=ts&Start=1&Count=30&cty=en_us (8) HP's single case study for Integrity Virtual Machines published at http://h71028.www7.hp.com/erc/library/GetPage.aspx?pageid=387810&audienceid=0&statusid=0&ccid=0&langid=121&ERL=true&pageTitle=Enterprise%20library:%20Royal%20London

Virtualization Can Provide Higher Availability

- **Reduce cost**
  - Significantly reduce idle high availability hardware
  - Automatically extend HA LPARs when needed
- **Cross campus disaster recovery**
  - Combine HA / Test and DR
  - Significantly reduce disaster recovery time over DR hosting site

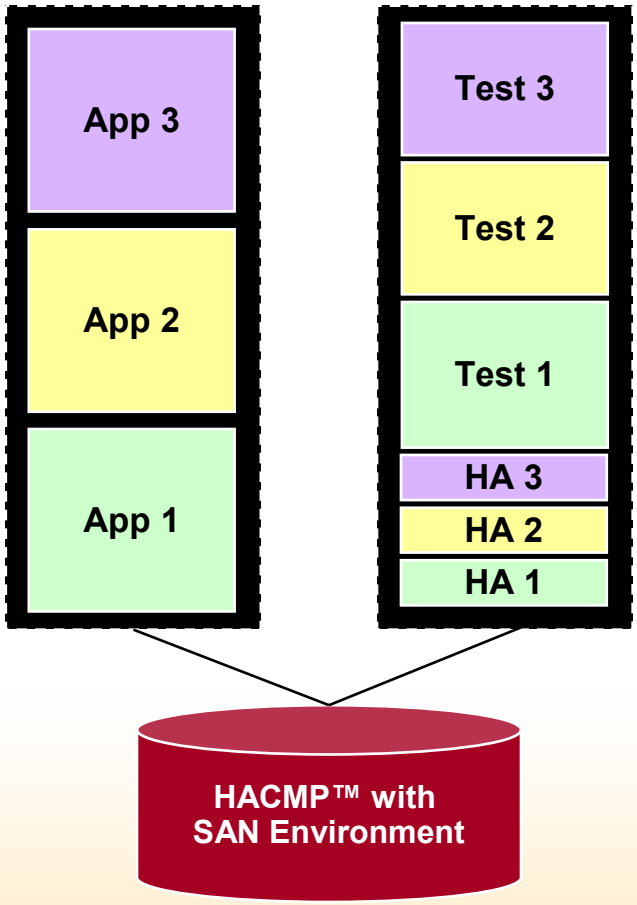

Advanced POWER Virtualization for POWER6

IBM APV Benefits

- ☑ ***Can help lower the cost** of existing infrastructure **by up to 72%**[4]*
- ☑ *Can increase **business flexibility** and **reduce** the **complexity to grow your infrastructure***
- ☑ *Deployed in production by a significant number of System p clients[5]*

Advanced POWER Virtualization[1]

Virtual I/O Server
- Share Ethernet, SCSI and Fibre Channel disks

Integrated Virtualization Manager[3]
- Manage a single system without an HMC

Live Partition Mobility 2007[3]
- Move a running partition from one POWER6 processor-based server to another with no downtime

Micro-Partitioning™
- Create up to 10 micro-partitions for each System p5 processor
- Resize without rebooting your system

AIX 6 Live Application Mobility 2007[3]
- Move a running partition from one POWER6 processor-based server to another with no downtime

Linux on POWER

AIX

1) Advanced POWER Virtualization (APV) is an optionally orderable feature on IBM System p. 3) All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM. 4) "Impact of IBM System p Server Virtualization," Transforming the IT Value Equation with POWER6 Architecture. International Technology Group, 05/2007. Study methodology: Companies in financial services, manufacturing and retail with $15 Billion+ revenues focusing on UNIX large enterprise environments with multiple, broad-ranging applications. Study compared the cost of the company's workloads running on multiple vendor servers and employing minimal virtualization to the cost of the company's workloads running on p570 (POWER6 processor-based) as well as POWER5+ processor-based servers – all using Advanced POWER Virtualization [APV]. APV is standard on System p5 590 and 595. Other System p servers have the option to add APV except the System p5 185. This cost analysis was performed for financial services, manufacturing and retail example environments with an overall average savings of up to 72% in TCO savings by virtualizing and consolidating on the System p servers. Total Cost of Ownership may not be reduced in each consolidation case. TCO depends on the specific client environment, the existing environments and staff, and the consolidation potential. 5) IBM sales Statistics

Advanced POWER Virtualization Web Site

http://www.ibm.com/systems/p/apv/index.html

Your one-stop shop for System p virtualization info:

-Discussion Forums
-Certifications
-Case Studies
-Whitepapers
-Education

"The logical partition [LPAR] capability of the System p5 server was the key factor in our decision, enabling us to run multiple independent systems on the same physical machine. In our opinion, IBM leads the market in this area." [1]

- Wolfgang Franz, IT Manager, Bionorica AG. December 2005

1) Bionorica case study published at http://www-306.ibm.com/software/success/cssdb.nsf/CS/DNSD-6KBFWW?OpenDocument&Site=eserverpseries

Advanced POWER Virtualization enhancements

*Leverage System p Virtualization and reduce server TCO by up to 60%**

Virtual I/O Server (VIOS) V1.4

VIOS Monitoring **through PTX and Topas**

Topas (part of AIX V5.3)

```
Topas CEC Monitor            Interval:  10
Partitions      Memory (GB)                Proce
Shr:  3         Mon:24.6  InUse: 2.7       Shr:1
Ded:  3         Avl:   -                   Ded:

Host         OS  M Mem InU Lp  Us Sy Wa Id  Ph
-------------------------------------shared---
ptools13     A53 c 4.1 0.4  2  20 13  5 62   0
```

PTX (AIX® LPP)

Performance Enhancements **for Virtual SCSI and Virtual Ethernet**

* "Business Case for IBM System p5 Virtualization," *Economic Benefits of IT Simplification*. International Technology Group, February 10, 2006

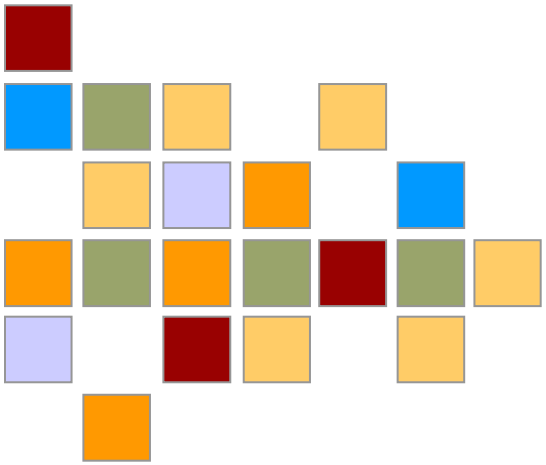

Performance

POWER5+ -- nearly 2X transaction performance and scalability per core and 50% more Java™ performance per core

Today's workloads

p5-595 and DB2®

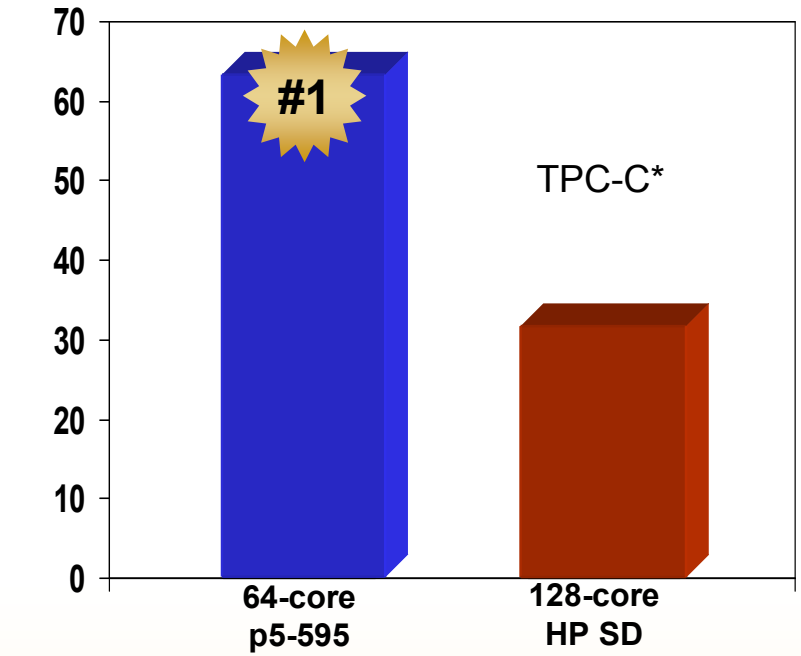

*Source: www.tpc.org/ All results as of 03/02/07

System (Processor, Chip/Core/Thread)	tpmC	Avail.	$/tpmC
IBM p5-595 (2.3 GHz POWER5+, 32/64/128)	**4,033,378**	**01/22/07**	**$2.97**
HP Superdome (1.6 GHz Itanium® 2, 64/128/256)	**4,092,799**	**8/23/07**	**$2.93**

SOA and emerging applications

p5-505Q and WebSphere®

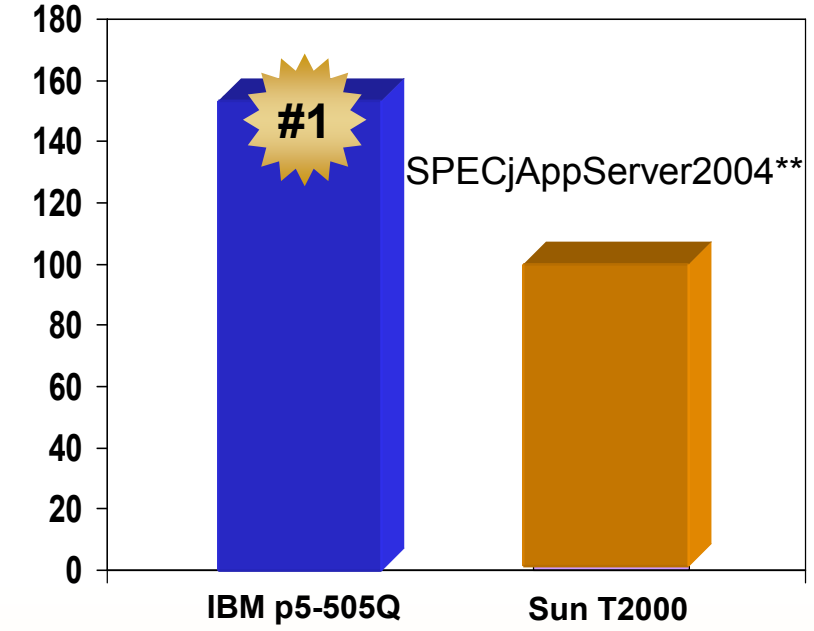

**Source: www.spec.org/ All results as of 02/15/06

System (Processor, Memory)	JOPs	Cores	Space
IBM p5-505Q (1.6GHz, 16 GB)	**618**	**4**	**1U**
Sun T2000 (1.2 GHz, 32 GB)	**733**	**8**	**2U**

Today, we are announcing…THE POWER OF SIX

IBM System p 570. . .with POWER6™ technology!

More than twice the performance* and modular flexibility

System p Virtualization. . .with Live Partition Mobility*!

Extending the most complete virtualization offering for UNIX and Linux

IBM AIX 6 . . . with full binary compatibility!

The next evolution of UNIX with new workload partitions, manageability and security

** All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.
Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Ten years ago…

Deep Blue changed the world's perception of what a computer can do

May 11, 1997

Equitable Center
New York City

IBM POWER technology: 10 years of innovation

Each core of IBM POWER6 exceeds the performance of Deep Blue

Deep Blue

1.4 tons / 1,270 kg

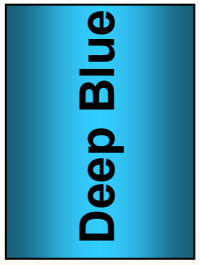

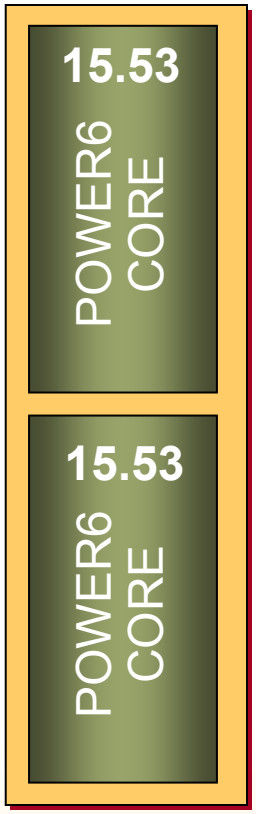

POWER6
System p 570

341mm²

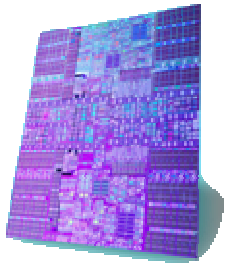

GFLOPS

Source: http://www.top500.org/list/1997/11/100 IBM DEEP BLUE(R) 1.2 GHz, 32 NODE SP2 P2SC, Rpeak: 15 GFLOPS, Rmax: 11.38 GFLOPS: IBM POWER6 CHIP, 4.7 GHZ 2 CORE, Rpeak – 37.6 GFLOPS, Rmax: 30.5 GFLOPS; to be submitted 5/21/07

System p with POWER6+

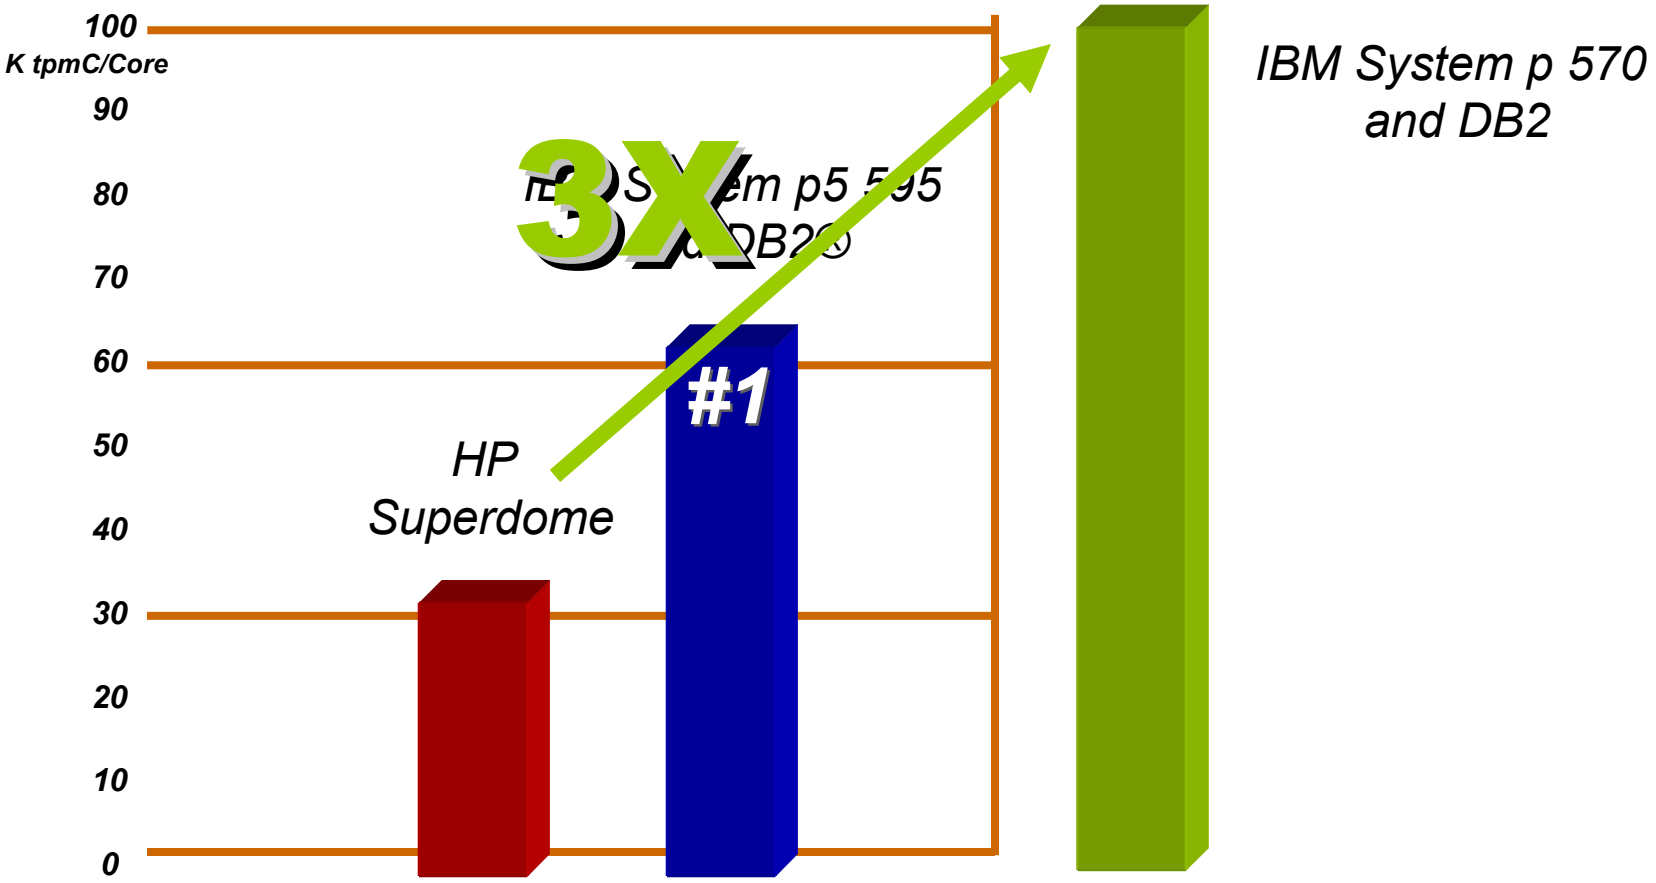

TPC-C*

System (Processor, Chip/Core/Thread)	tpmC	Avail.	$/tpmC
IBM p570 (4.7 GHz POWER6 , 8/16/32)	1,616,162	11/21/07	$3.54
IBM p5-595 (2.3 GHz POWER5+, 32/64/128)	4,033,378	01/22/07	$2.97
HP Superdome (1.6 GHz Itanium 2, 64/128/256)	4,092,799	8/23/07	$2.93

*Source: www.tpc.org/ All results as of 05.21.07

System p with POWER6+

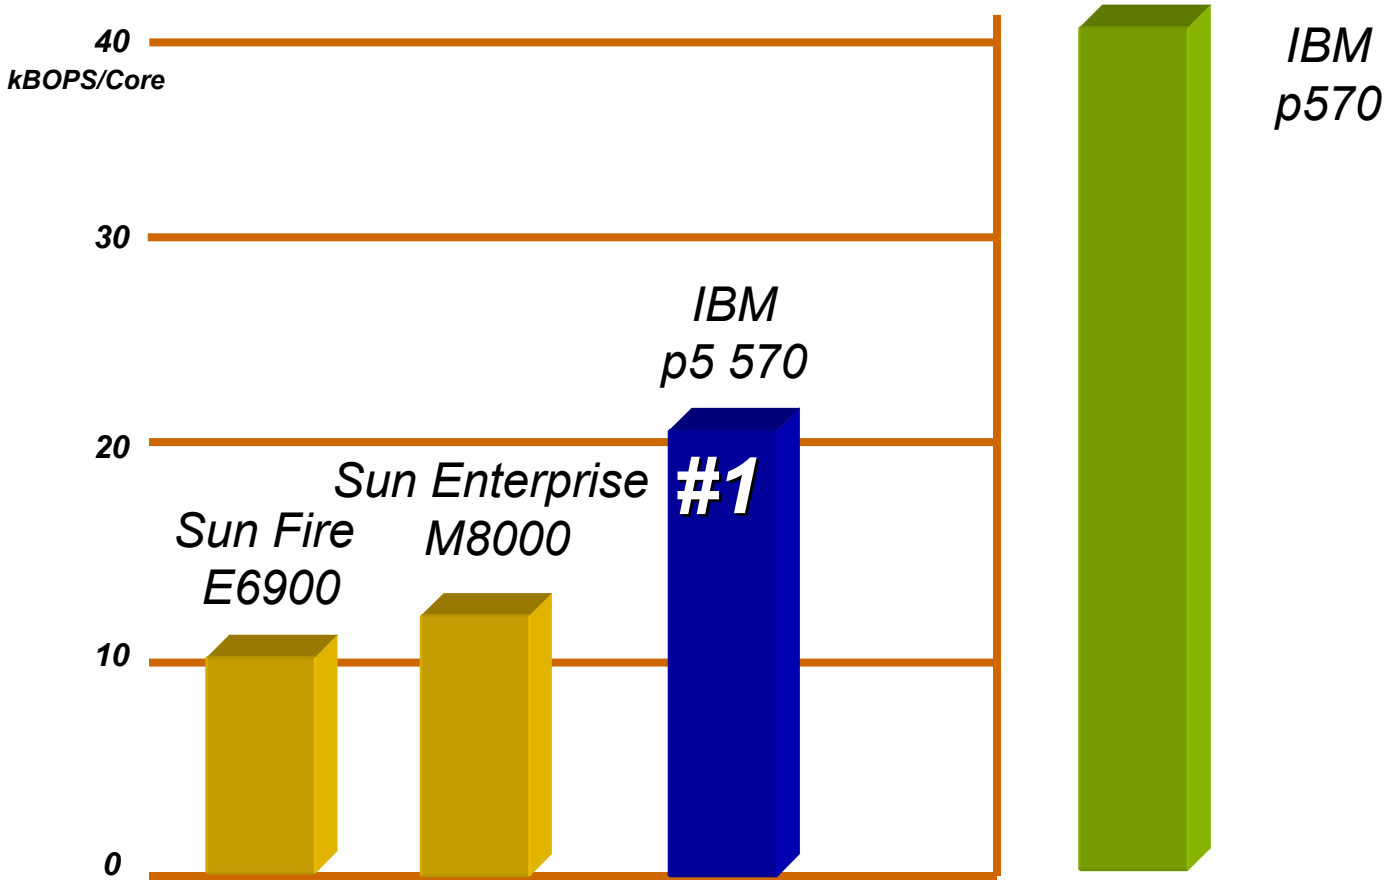

SPECjbb2005**

System (Processor, Cores)	BOPS	Date.	BOPs/JVM
IBM p570 (4.7 GHz POWER6 , 16-core)	691,975	May - 2007	86,497
IBM p5-570 (2.2 GHz POWER5+, 16-core)	326,651	March - 2006	40,831
Sun Fire E6900 (1.95 GHz UltraSPARC VI+, 48-core)	354,733	April - 2007	22,171
Sun Enterprise M8000 (2.4 GHz SPARC64 VI, 32-core)	440,207	May - 2007	27,513

**Source: www.spec.org/ All results as of 05/21/07

The IBM POWER6 *"Grand Slam"* for major workloads

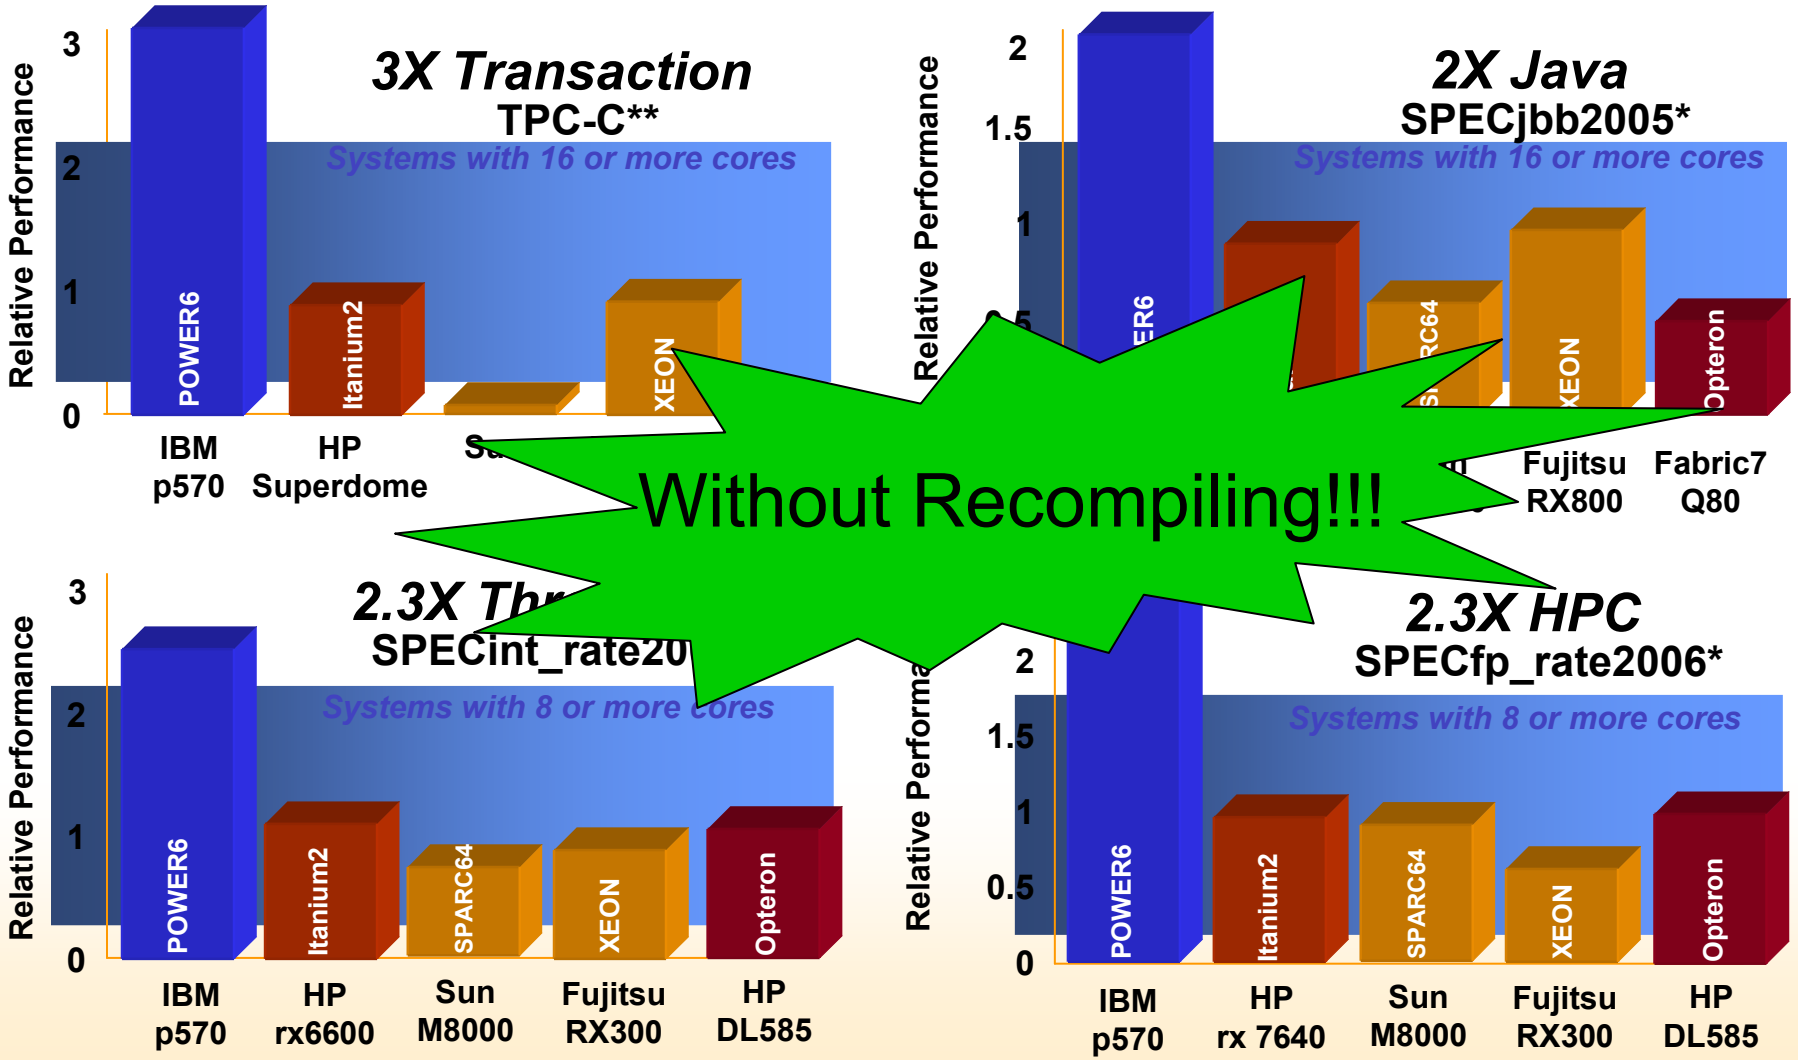

* Source: http://www.spec.org/ IBM p570 POWER6 results to be submitted on 5/21/07: All other results as of 04/27/07;
** Source: www.tpc.org/ IBM p570 POWER6 result to be submitted on5/21/07; All other results as of 04/27/07
See next page for full detail

The IBM POWER6 *"Grand Slam"* for major workloads

- SPECjbb2005 comparisons (Source: http://www.spec.org/ IBM p570 POWER6 results to be submitted on 5/21/07: All other results as of 04/27/07)
  - **IBM POWER6 p570 (8 chips, 16 cores) @ 4.7 GHz with 691,975 bops (86,497 bops/JVM) and 43,125 bops per core**
  - **Fujitsu PRIMEQUEST 580 (32 chips, 64 cores) @ 1.6 GHz with 1,214,251 bops (75897 bops/JVM) and 18,974 bops per core**
  - **Sun Fire M8000 (16 chips, 32 cores) @ 2.4 GHz with 440,207 bops (27,513 bops/JVM) and 13,756 bops per core**
  - **Fujitsu RX800 (8 chips, 16 cores) @ 3.5 GHz with 336,653 bops (42,082 bops/JVM) and 21,041 bops per core**
  - **Fabric7 Q80 (8 chips, 16 cores) @ 2.6 GHz with 180,418 bops (22,552 bops/JVM) and 11,276 bops per core**
- TPC-C comparisons (Source: www.tpc.org/ IBM p570 POWER6 result to be submitted on 5/21/07; All other results as of 04/27/07)
  - **IBM POWER6 p570 (8 chips, 16 cores, 32 threads) @ 4.7 GHz with tpmC of 1,616,162 @ $3.54 $/tpmC with availability of 11/20/07 and 101,010 tpmC per core**
  - **HP Integrity Superdome (64 chips, 128 cores, 256 threads) @ 1.6 GHz with tpmC of 4,092,799 @ $2.93 $/tpmC with availability of 8/23/07 and 31,953 tpmC per core**
  - **Unisys ES7000 (8 chips, 16 cores, 32 threads) @ 3.4 GHz with tpmC of 520,467 @ $2.73 $/tpmC with availability of 5/1/07 and 32,529 tpmC per core**
- SPECint_rate2006 (Source: http://www.spec.org/ IBM p570 POWER6 results to be submitted on 5/21/07: All other results as of 04/27/07)
  - **IBM POWER6 p570 (4 chips, 8 cores) @ 4.7 GHz with 242 and 30.25 per core**
  - **HP rx6600 (4 chips, 8 cores) @ 1.6 GHz with 102 and 12.75 per core**
  - **Sun Fire M8000 (16 chips, 32 cores) @ 2.4 GHz with 298 and 9.3125 per core**
  - **Fujitsu RX300 (4 chips, 8 cores) @ 2.66 GHz with 91.2 and 11.4 per core**
  - **HP ProLiant DL585 (4 chips, 8 cores) with 98.3 and 12.29 per core**
- SPECfp_rate2006 (Source: http://www.spec.org/ IBM p570 POWER6 results to be submitted on 5/21/07: All other results as of 04/27/07)
  - **IBM POWER6 p570 (4 chips, 8 cores) @ 4.7 GHz with 224 and 28 per core**
  - **HP rx7640 (4 chips, 8 cores) with 90.8 and 11.35 per core**
  - **Sun Fire M8000 (16 chips, 32 cores) @ 2.4 GHz with 313 and 9.7813 per core**
  - **Fujitsu RX300 (4 chips, 8 cores) @ 2.66 GHz with 60.9 and 7.61 per core**
  - **HP ProLiant DL585 (4 chips, 8 cores) with 91.3 and 11.41 per core**

* Source: http://www.spec.org/ IBM p570 POWER6 results to be submitted on 5/21/07: All other results as of 04/27/07;
** Source: www.tpc.org/ IBM p570 POWER6 result to be submitted on5/21/07; All other results as of 04/27/07

IBM WebSphere® optimized for System p

Performance Leadership

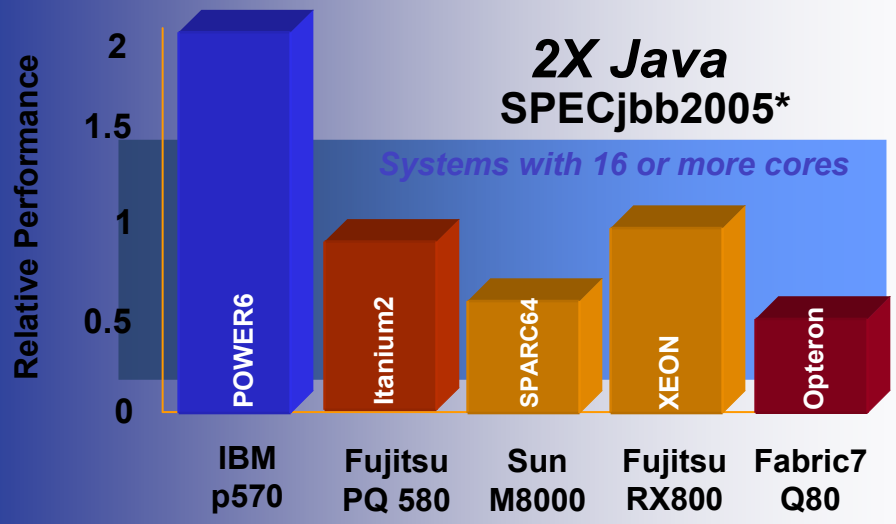

SOA Affinity

IBM System p Configurations for SOA Entry Points

*Source: http://www.spec.org/ IBM p570 POWER6 results to be submitted by 5/21/07: All other results as of 04/27/07

PERFORMANCE OF SIX -- for server sprawl

"Grand Slam" leadership provides more than 2X better business application performance

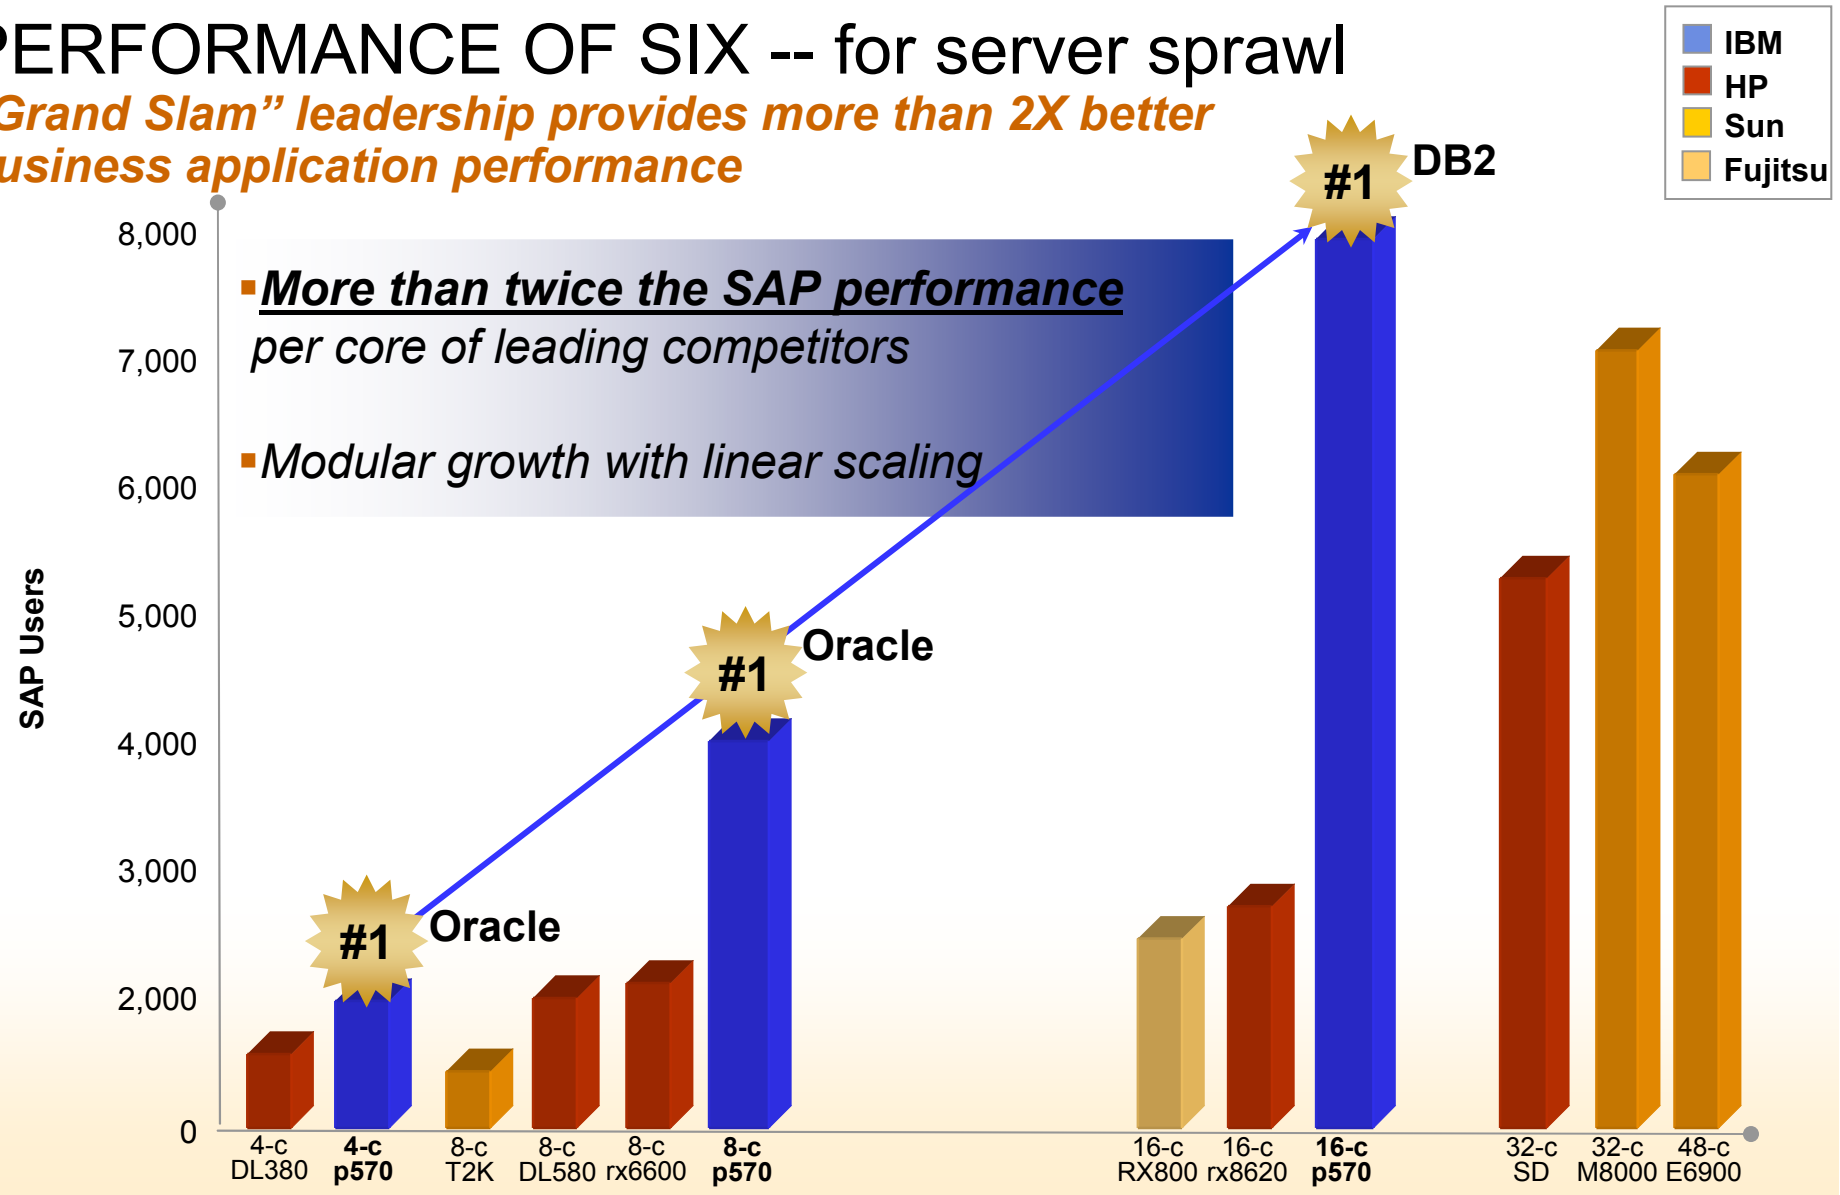

Source: www.sap.com/benchmark/ IBM results submitted on 5/21/07; All other results as of 04/27/2007; see accompanying table for additional detail

SAP SD 2-tier detailed results

Configuration	Chips, Cores, Threads	Software	Certification #	Benchmark users	Response Time
IBM p570 4.7 GHz POWER6	8, 16, 32	SAP ECC 6.0 AIX 5L V5.3, DB2 9.1	New *	8,000	1.98 sec
Sun Microsystems M8000 2.4 GHz SPARC64 VI	16, 32, 64	SAP ECC 6.0 Solaris 10, Oracle 10g	2007026	7,300	1.98 sec
Sun Microsystems E6900 1.95 GHz USIV+	24, 48, 48	SAP ECC 6.0 Solaris 10, Oracle 10g	2007023	6,160	1.98 sec
HP Integrity Superdome 1.6 GHz Itanium2	16, 32, 64	SAP ECC 6.0, Windows Server 2003 DE, SQL Server 2005	2006090	5,600	1.91 sec
IBM p570 4.7 GHz POWER6	4, 8, 16	SAP ECC 6.0 AIX 5L V5.3, Oracle 10g	New *	4,010	1.96 sec
HP Integrity rx8620 1.5 GHz Itanium2	8, 16, 32	SAP R/3 4.70 HP-UX 11i, Oracle 9i	2003062	2,880	1.95 sec
Fujitsu Siemens PRIMERGY RX800 3.0 GHx Xeon	8, 16, 32	SAP ECC 5.0, Windows Server 2003 EE, SQL Server 2005	2006022	2,600	1.94 sec
HP Integrity rx6600 1.6 GHz Itanium2	4, 8, 16	SAP ECC 6.0, HP-UX11/V3, Oracle 10g	2006082	2,150	1.97 sec
HP Proliant DL580 G4 3.4 GHz XEON	4, 8, 16	SAP ECC 5.0, Windows Server 2003 EE, SQL Server 2005	2006060	2,127	1.99 sec
IBM p570 4.7 GHz POWER6	2, 4, 8	SAP ECC 6.0 AIX 5L V5.3, Oracle 10g	New *	2,035	1.99 sec
HP Proliant DL380 G5 3.0 GHz XEON	2, 4, 4	SAP ECC 5.0, Windows Server 2003 EE, SQL Server 2005	2006039	1,216	1.99 sec
Sun T2000 1.2 GHz US T1	1, 8, 32	SAP ECC 5.0 Solaris 10, MaxDB 7.5	2005047	950	1.91 sec

* The SAP certification number was not available at press time and can be found at the following web page: http://www50.sap.com/benchmarkdata/sd2tier.asp
Source: www.sap.com/benchmark/
New IBM results submitted on 5/21/07
All other results as of 04/27/07

Save more than $100K per year on energy and space costs!

Based on new IBM p570 performance and efficiency*

30 Sun Fire V890 systems at 20% utilization

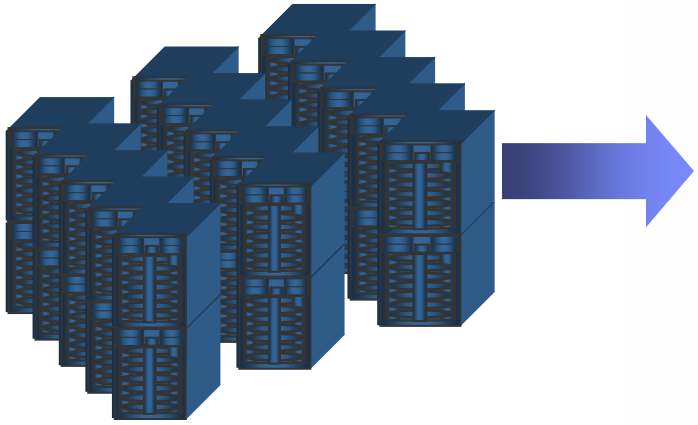

- **480 total cores @ 1.5 GHz**
- **$5,625 annual space costs @ $62.50 sq ft**
- **$113,607 annual energy costs @ $0.09 / kWhr**

Two IBM System p 570 systems at 60% utilization

- *Save up to 90% cost of floor space*
- *Save up to 90% cost of energy*
- *Save up to 90% on per core SW costs*

- **32 total cores @ 4.7 GHz**
- **$375 annual space costs @ $62.50 sq ft**
- **$13,667 annual energy costs @ $0.09 / kWhr**

*Datacenter floor space cost was estimated as of 5/3/2007 based on Alinean, Inc.'s ROI Analyst software. Energy cost of $.0928 per kWh is based on 2007 YTD US Average Retail price to commercial customers at $.0928 per US DOE at http://www.eia.doe.gov/cneaf/electricity/epm/table5_6_b.htm as of 05/18/2007
The reduction, if any, in floor space, power, cooling and software costs depends on the specific customer, environment, application requirements, and the consolidation potential. Sun system power requirements based on http://www.sun.com/products-n-solutions/hardware/docs/html/817-3956-12/system_specs.html#pgfId-1001301
Air conditioning power requirement estimated at 50% of system power requirement. SPEC® results source: www.spec.org as of 05/22/2007: System p 570 (16-core, 8 chips, 2 chips per core, 4.7 GHz): SPECjbb2005 691,975 bops 86497 bops/JVM; Sun Fire v890 (16-core, 8 chips, 2 chips per core) 1.5 GHz SPECjbb2005 117,986 bops, 29,497 bops/JVM

Why buy fifteen HP systems when you can get better performance AND reduce costs with only two IBM p570s?

FIFTEEN 16-core HP Integrity rx7640 systems

- 240 total cores @ 1.6 GHz
- $38,538 annual energy costs @ $0.09 / kWhr
- $1,500 annual space costs @ $62.50 sq ft

TWO 16-core IBM p570 systems

- *Gain up to 14% performance advantage, and*
- *Save up to 85% cost per core in software fees*
- *Save up to 75% cost of floor space*
- *Save up to 65% cost of energy costs*

- **32 total cores @ 4.7 GHz**
- **$13,254 annual energy costs @ $0.09 / kWhr**
- **$375 annual space costs @ $62.50 sq ft**

*Datacenter floor space cost was estimated as of 5/3/2007 based on Alinean, Inc.'s ROI Analyst software. The reduction, if any, in floor space, power, cooling and software costs depends on the specific customer, environment, application requirements, and the consolidation potential. SPEC® results source: www.spec.org as of 05/22/2007: System p 570 (16-core, 4.7 GHz): SPECint_rate2006: 478, HP Integrity rx7640 16-core, 1.6 GHz SPECint_rate2006: 167 Energy cost of $.0928 per kWh is based on 2007 YTD US Average Retail price to commercial customers at $.0928 per US DOE at http://www.eia.doe.gov/cneaf/electricity/epm/table5_6_b.htm as of 05/18/2007

Now you may be able to save even more!

The German company Voith consolidated on System p5 570s saving*:

- *70-80% on energy costs*
- *30% on SW costs*

Twice the performance at almost the same energy

30 HP systems
Capacity = 50,000 SAPs

Four IBM System p5 570 servers
Capacity = 80,000 SAPs

Two POWER6 System p 570s
Capacity =80,000 SAPs

http://www-306.ibm.com/software/success/cssdb.nsf/CS/STRD-72NM7N?OpenDocument&Site=corp&cty=en_us

Get 39% more Java application performance by migrating from a 32-core Sun Fire E6900 to an 8-core System p 570

And potentially save. . .

- *87% cost of floor space*
- *50% cost of energy*
- *75% cost per core SW license*

Sun Fire E6900
- 32-core, 1.5 GHz (UltraSPARC IV+)
- 28 Rack Units
- 5,945 Volt-Amps
- SPECjbb2005: 248,075 bops (bops/JVM: 31,009)

System p 570
- 8-core, 4.7 GHz (POWER6)
- 8 rack Units
- 2,800 Volt-Amps
- SPECjbb2005: 346,742 bops (bops/JVM: 86,686)

*The reduction, if any, in floor space, power, cooling, and software costs depends on the specific customer, environment, and application requirements. SPECjbb2005 bops source: www.spec.org as of 05/22/2007 **System power requirements based on VA ratings at http://www.powerware.com/UPS/selector/BuildByDevice.asp used with permission of Eaton Corporation. Normalized power factor = 1.0

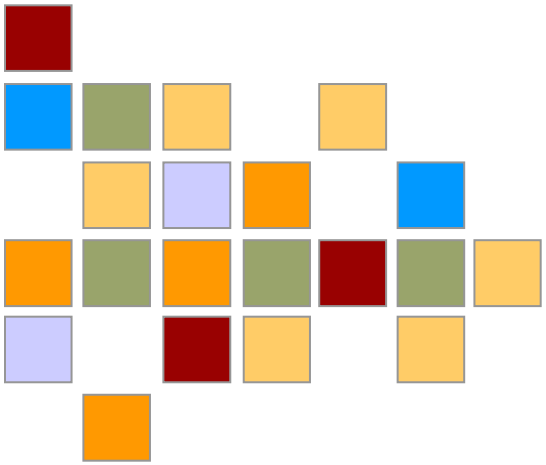

Additional Hardware

POWER6 Remote I/O Drawer 19” Rack

7314-G30

- ½ x 19” 4U Rack Mount
  - Optional enclosure for two individual drawers
- Six PCI Adapter Slots
  - PCI-X 2.0 (DDR)
  - 64-bit @ 266 MHz (2 GB/s)
- “InfiniBand” Host Interface (Loop Architecture)
  - 12 x 2.5Gb/s Full Duplex (30 Gb/sec)
  - Short Run (Intra-rack) or Long Run (Inter-rack)
  - Four cable lengths: 0.6, 1.5, 3.0, & 8.0m
- Customer Setup
- Concurrently Maintainable Redundant Power
- Concurrently Maintainable Redundant Cooling
- Hot Drawer Add
- Blind Swap Cassettes for Adapter Cards
  - Hot-pluggable

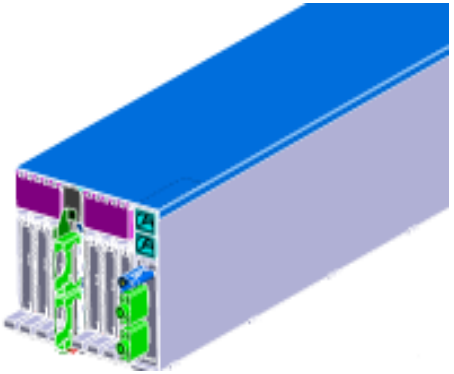

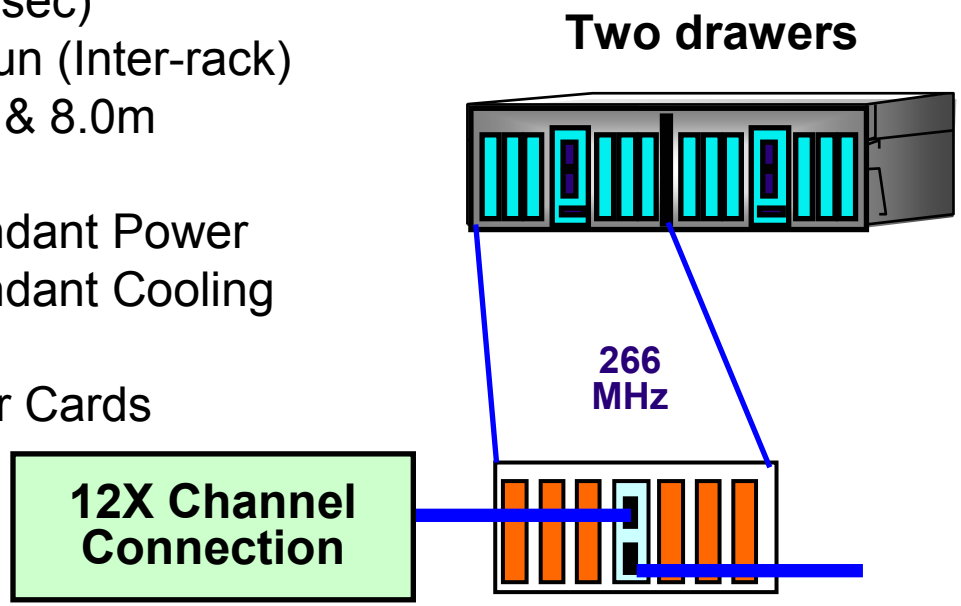

New POWER6 I/O Adapters

Adapter	Fibre Channel	Ethernet	Host Enet Adap. Integ. Virt. Enet	2D Graphics
Host Bus	PCIe x4 1 GB/s FDX	PCIe x4 1 GB/s FDX	GX+ (P5IOC2) ~2-3 GB/s FDX	PCIe x1 250 MB/s FDX
Technology	4 Gb/s	1 Gb/s TX/SX	1 Gb/s TX 10 Gb/s SR	Analog & Digital
Ports	1 & 2	2	2 & 4 (1 Gb/s) 2 (10 Gb/s)	2
Feature Codes	5773 / 5774	5767 / 5768	5636 / 5639 / 5637	5748
AIX	5.3..6	5.2.10, 5.3.6	5.2.10, 5.3.6	5.2.10, 5.3.6
Linux	SLES 10 SP1 RHEL 4.5	SLES 10 SP1 RHEL 4.5	SLES 10 SP1 RHEL 4.5	SLES 10 SP1 RHEL 4.5
Virtualization	VIOS	VIOS	Stand-alone	NA

PCI Express Overview

PCI Express is available in several sizes

- x1, x2, x4, x8 & x16
- x4 is what most adapters are being designed to

PCI Express uses less pins then does PCIx

- PCIx = 188 pins, PCIe = 64 pins

(x4) theoretical bandwidth is 10 Gb in each direction (20 Gb aggregate bandwidth), where as PCI-X 1.0 (133 MHz) aggregate bandwidth is 8 Gb.

Overall, we recommend all the new installations to take advantage of PCI Express slots, due to the inherent advantages of PCI Express technology (e.g. less I/O pins, lower power, lower latency etc.) and protect their investment for foreseeable future.

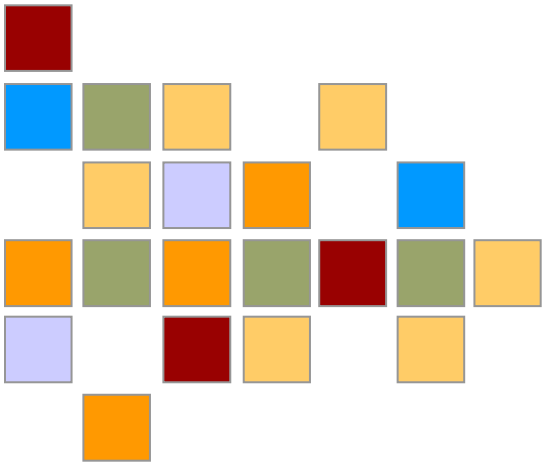

Hardware Management Console & Tools

Hardware Management Console (HMC)

Models available:	7310/7042-C06 (desktop) 7310/7042-CR4 (rack-mount)
Ethernet support:	POWER6 and POWER5 Systems
Requirements:	Required for: ▸ Partition Management ▸ CoD ▸ Virtualization activities Optional for APV on standalone servers
Licensed Machine Code Version 7:	Supports POWER6 and POWER5/5+ processor-based servers only

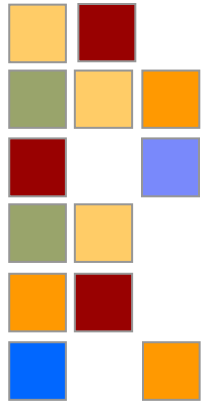

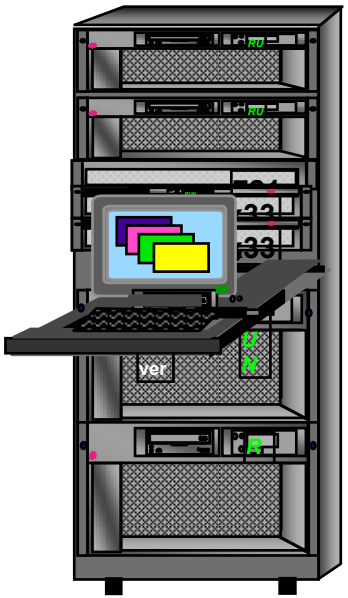

POWER6 HMC Enhancements

Support for POWER5 and POWER6 on same HMC

Updated HMC hardware (Intel® technology refresh)

No change: Hardware scaling support

- 7342-CR4 & 7342-C06
- 32 physical systems
- Up to 254 LPARS

Native browser access; WebSM no longer required

- Firefox 1.5.0.6 or later.
- Microsoft® Internet Explorer 6.0 or later

Support for modified CSM on HMC

Upgrade support for POWER5 HMC to POWER6

- 7310 will support POWER6 environment
- New model type for POWER6: 7042

Internal modem support for rack models

- Available with CR5 models
- Support will vary by geography

MES on Existing p5 HMC to get HMC v7

HMC Enhancements

- New Web-based User Interface
  - Improved task discovery – improve visibility of important tasks
  - Reduced task depth - reduce number of clicks and views to find a task
  - Consistent navigation and categorization
  - No special remote client installation required, supports both Firefox and Internet Explorer
  - Unifies user experience across HMC, IVM and IBM Web UI family
- Enhanced Management and Deployment Wizards

New HMC Layout

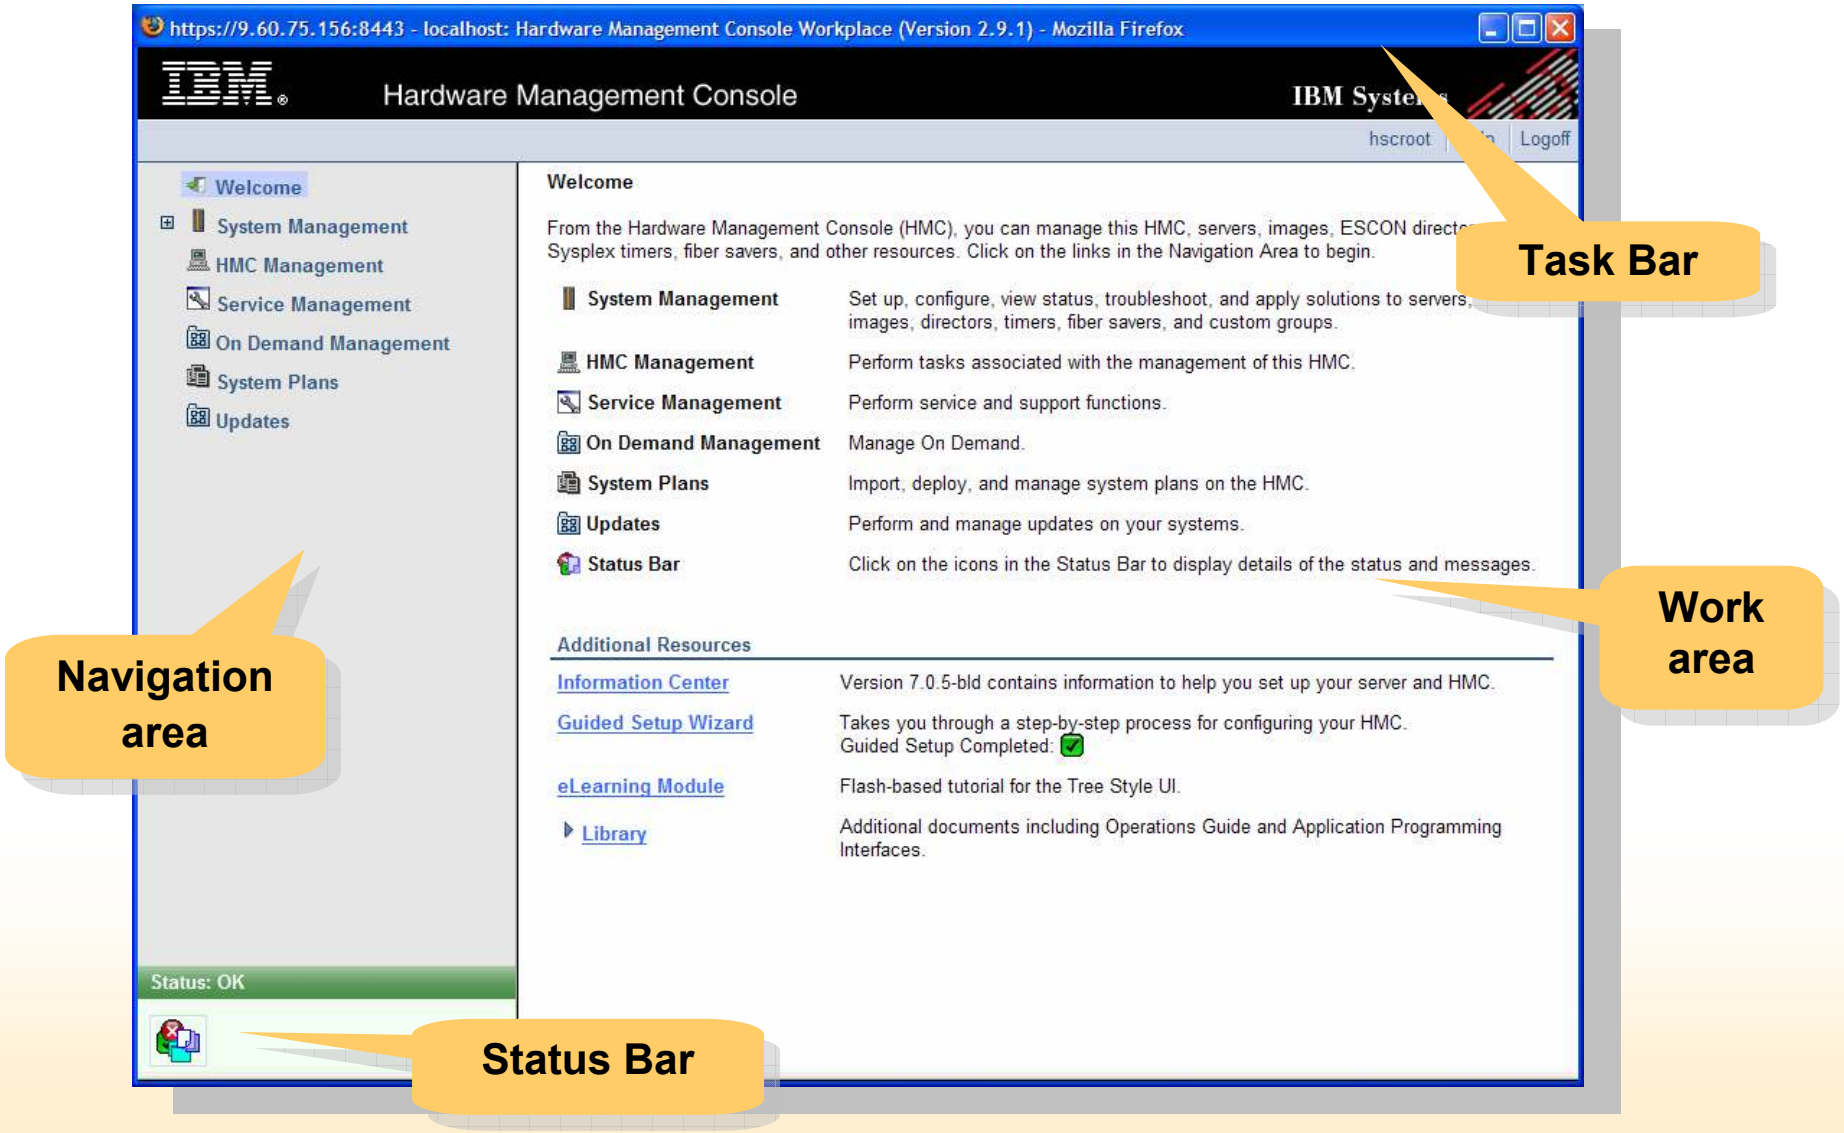

Simplify Virtualized Systems Planning

With the free IBM Systems Planning and Deployment Tool

- **Seamlessly plan, order, and deploy your system using this free tool**

- **Enhancements include**
  - Virtual I/O Server Partition Planning and Deployment Support
    - Support for primary and backup VIOS instances
    - Redundancy, Shared Ethernet, SAN connections and internal storage
    - MPIO attachments
    - VIOS configuration recovery support
    - Configuration validation
  - Integrated Virtualization Manager Support
    - For automated configuration of partitions and deployment through IVM
  - Enhanced Reporting
    - HTML reporting for partition hardware inventory and system attributes

Free download from http://www.ibm.com/systems/support/tools/systemplanningtool/

IBM Systems Director Value Propositions for System p

- Consistent platform management of both **physical** and **virtual** System p resources (AIX, Linux, HMC, IVM, VIOS)
- IBM Director on System p provided at no charge for core management capabilities
- SWMA available for $120 / CPU

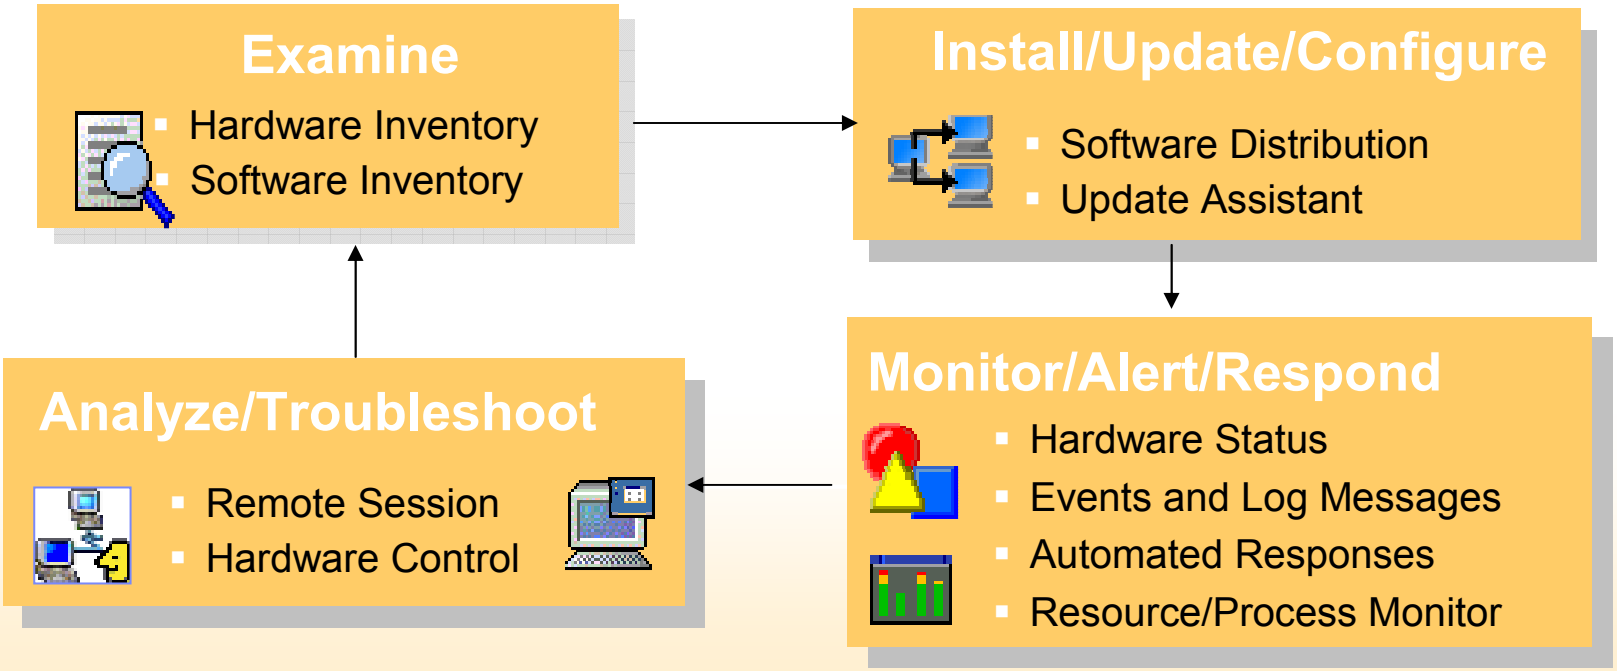

HMC Display and Management

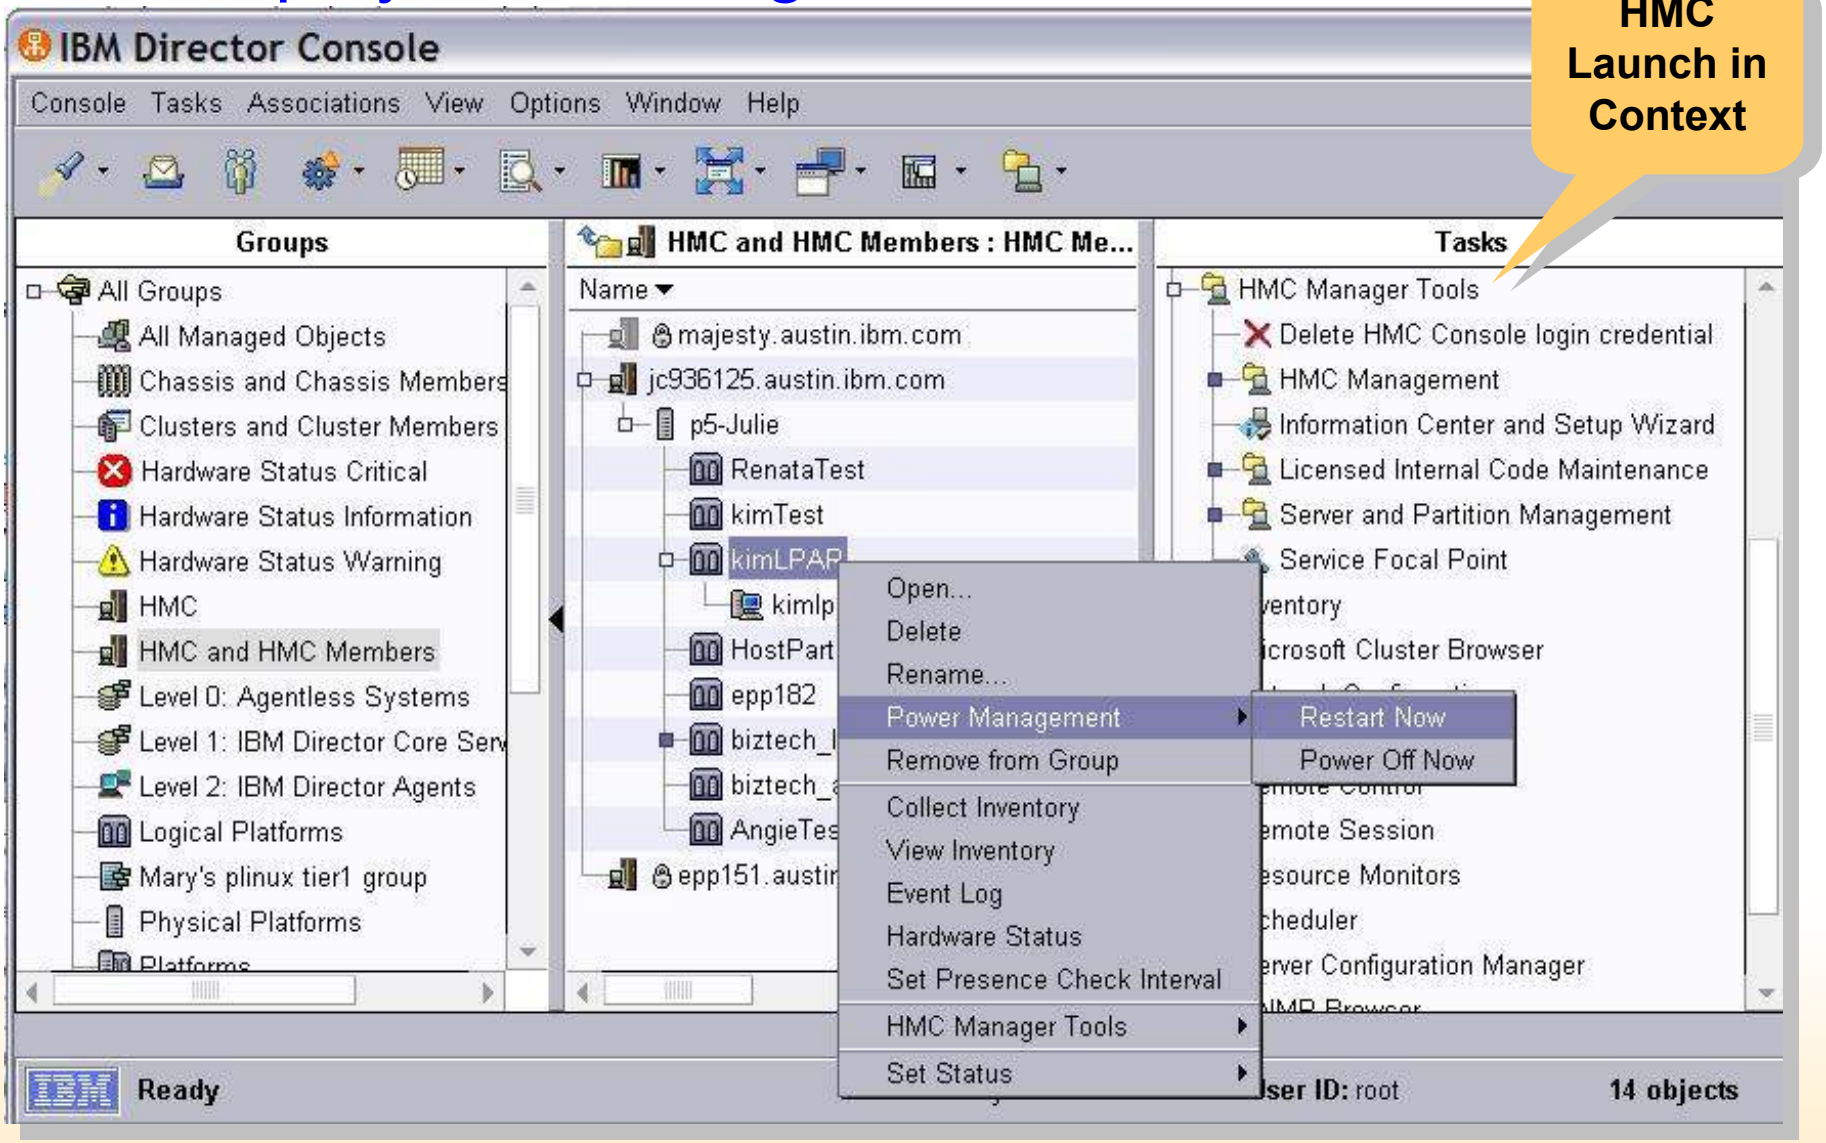

IBM Tivoli Monitoring System Edition for System p

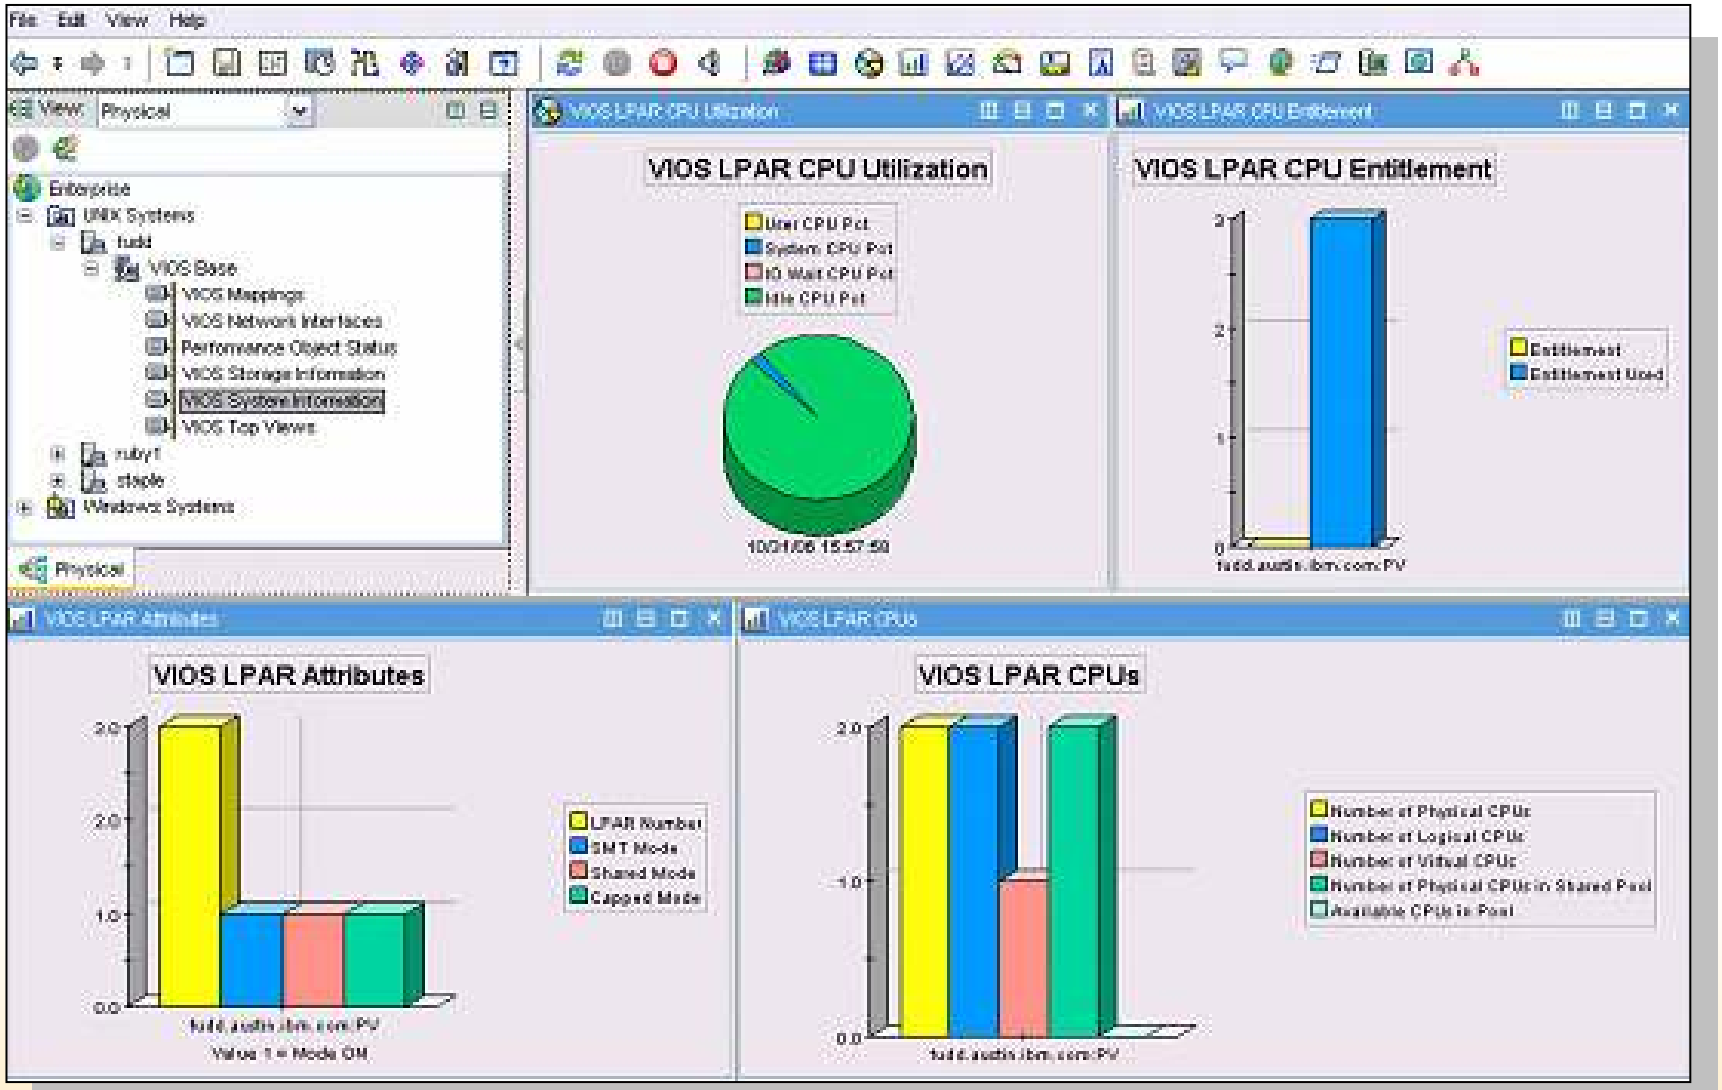

IBM Tivoli Monitoring Free vs. Fee Products

	Free	Fee
Topology and Navigation ▪ HMC, IVM, VIOS, CEC, LPARs, VIOS Server and Client, WPARs	X	X+
Availability Monitoring ▪ HMC, IVM, VIOS, LPAR, WPAR Status ▪ AIX and VIOS System Level CEC, LPAR, CPU, Memory Metrics	X	X+
Health ▪ AIX and VIOS Checks, Alert Messages, Expert Advise, Actions ▪ CPU, Memory, Disk, and Network Thresholds, File System Status, Paging Space, Status of Daemons and Services (i.e. NFS, http), Top Resource Consumers, etc.	X	X+
Client Customizable Workspaces, Navigators, Eventing, Situations	X	X+
Performance and Throughput ▪ AIX and VIOS ▪ Existing ITM Metrics (i.e. CPU, Memory, I/O, Network, File System) ▪ AIX PTX Metrics (i.e. CPU, Memory, LAN, TCP, UDP, IP , WLM, Process, LPAR, Disk, I/O, LVM, Paging Space, IPC, NFS, CEC) ▪ Customized VIOS and WPAR Metrics		X
Data Warehouse ▪ Historic Performance Data		X
Workflows ▪ Client Configurable		X

HACMP provides leadership High Availability and Disaster Recovery for AIX 5L and

Now available for Linux!

- **Automatically detect failure in servers, software, storage or network and recover fast**
- Over 60,000 licenses worldwide
- Recent releases have greatly improved ease of use
- Unlimited distance geographic clusters for remote disaster recovery
  - XD option supports IP and hardware data mirroring

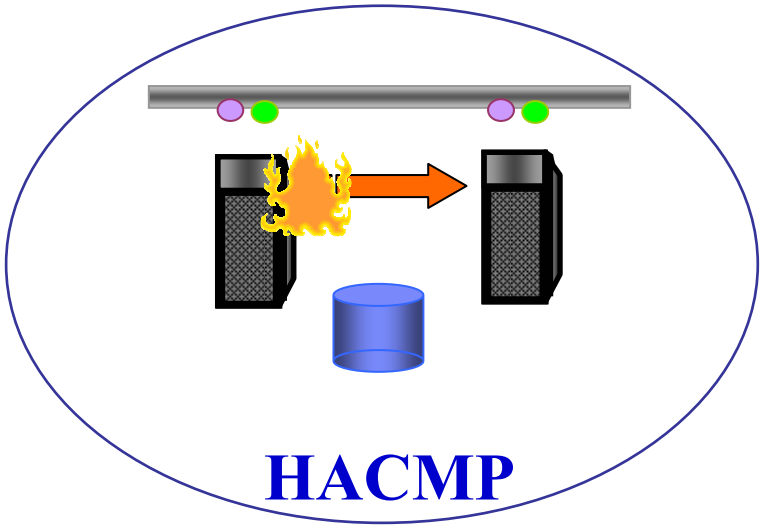

- Configurable recovery policies and options
- Scalable from 2 to 32 server nodes
- Single point-of-control administration
- Dynamic cluster reconfiguration
- Automatic configuration discovery
- Supports clustered file system

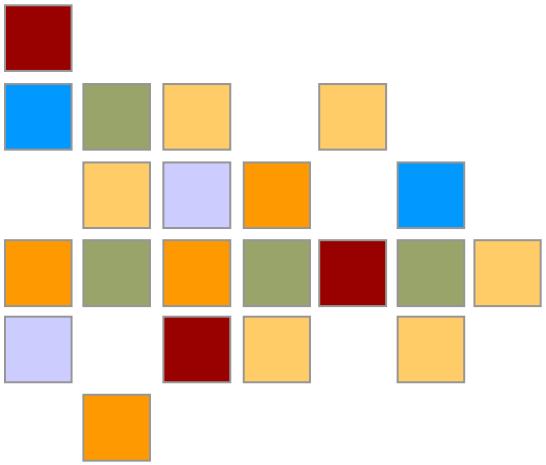

Operating System

AIX Release Plan

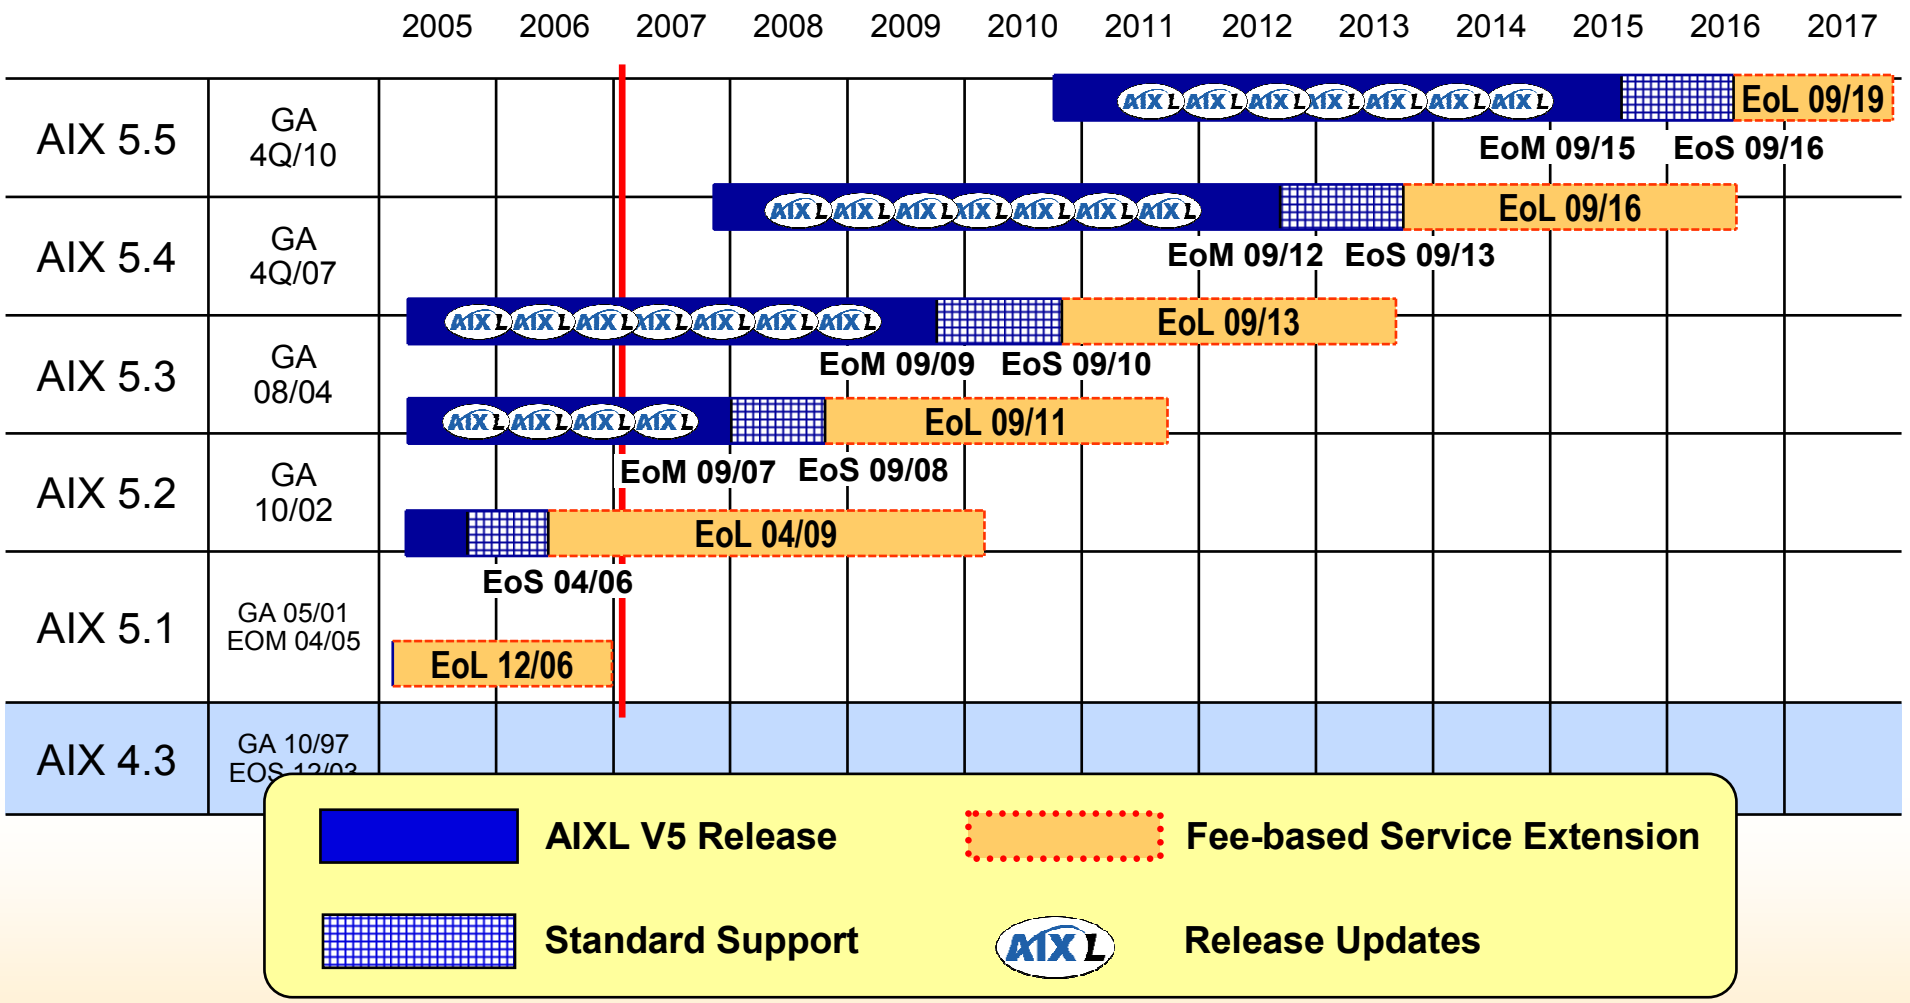

Plan is subject to changes

AIX Features Roadmap

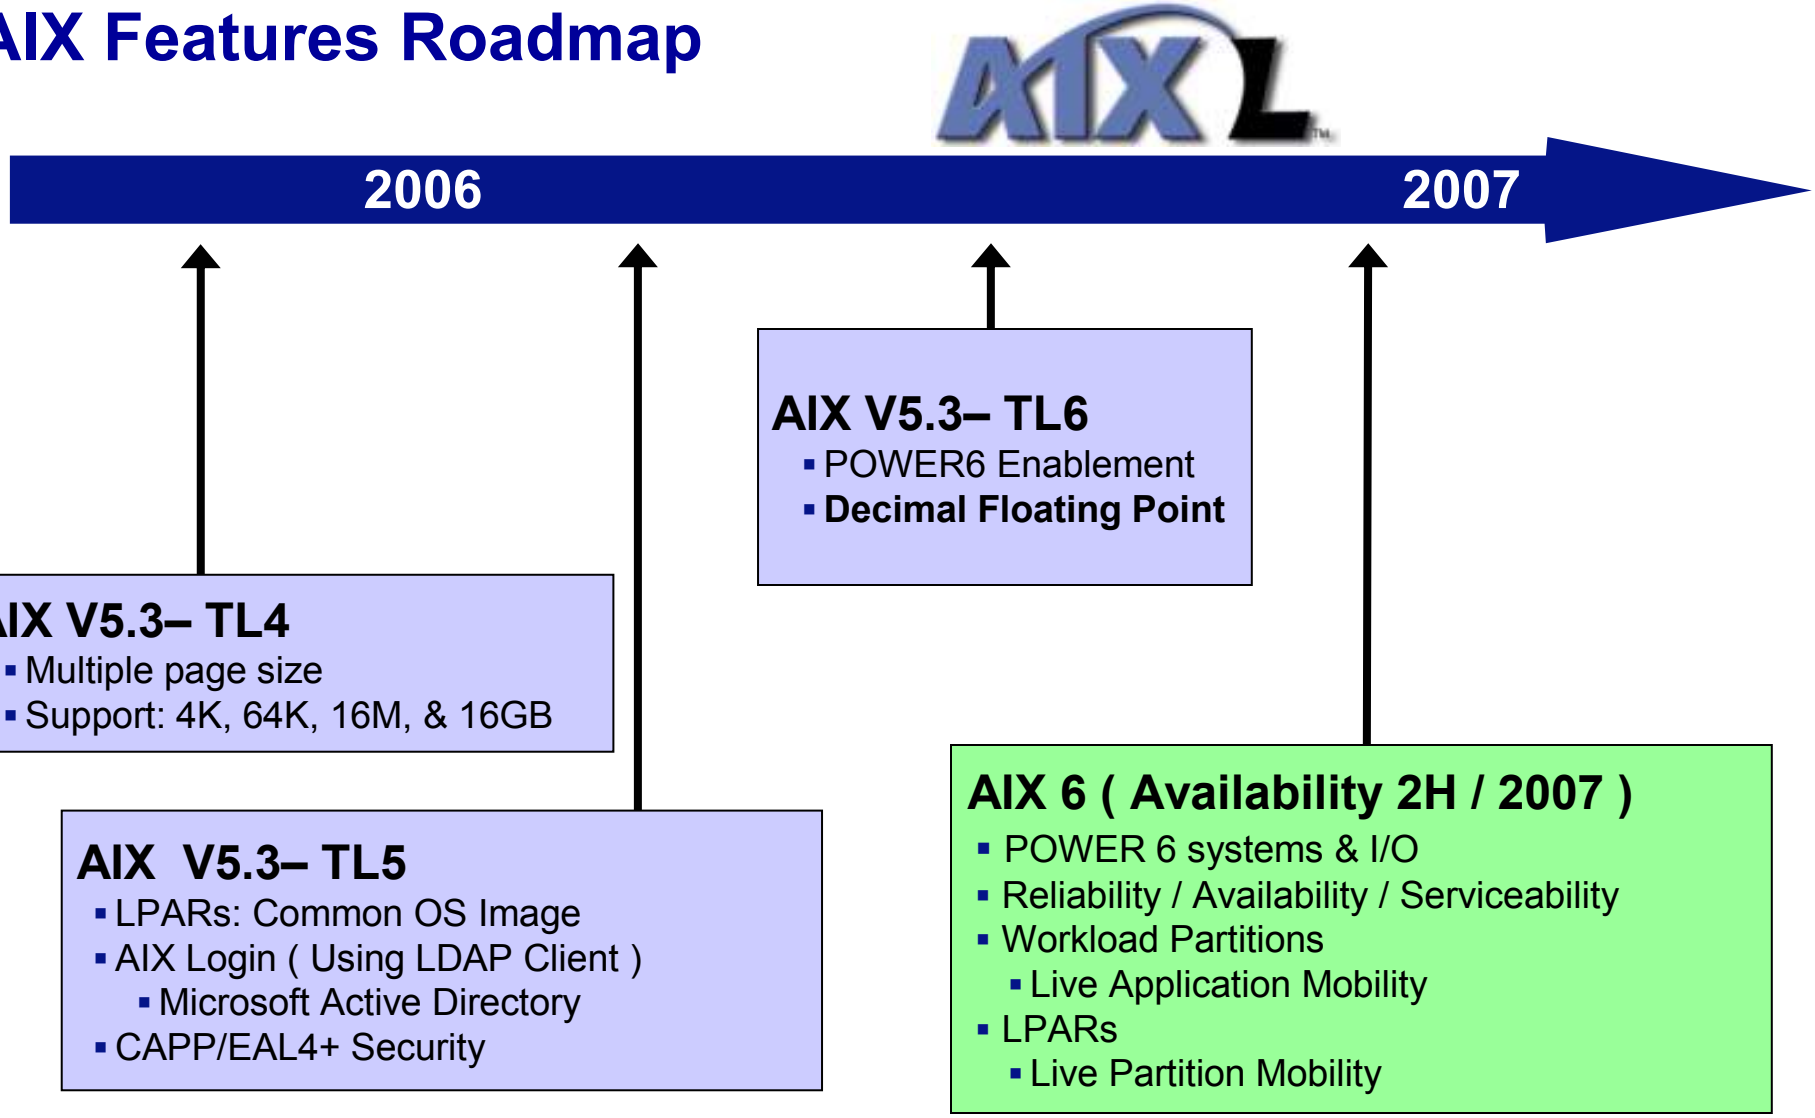

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice

AIX Version 5.3 Features

Flexible resource management	IBM Advanced POWER Virtualization ▸ Micro-Partitioning ▸ Virtual I/O : Networking and Storage ▸ Partition Load Manager ▸ Integrated Virtualization Manager Advanced accounting JFS2 file system shrink Cross partition performance monitoring
Enterprise scalability	POWER5+ support Simultaneous multithreading processor 1,024 disk volume group NFSv4
Entreprise reliability	AIX First Failure Data Capture instrumentation Component level tracing Parallel dump, Minidump
Data center management	AIX Security Expert Shared Boot Image management SUMA integration with NIM Network Data Administration Facility (NDAF) Virtualization management Tivoli® integration

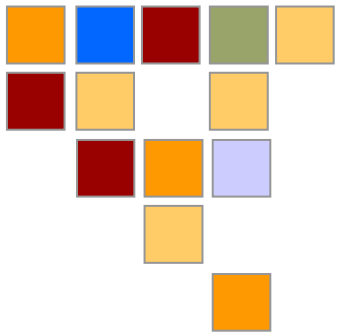

AIX V5.3 August 2006 Update Highlights

AIX Security Expert

- Template-based security wizard
- Security templates can be distributed throughout an enterprise

AIX LDAP client support for Microsoft Active Directory

- AIX systems can use Active Directory for user and group repository

Shared Boot Disk Management

- Boot multiple LPARs/systems from a single boot disk

LPAR Historical Performance Reports

Enterprise Reliability

- Parallel system dump processing
- Component level system trace
- AIX First Failure Data Capture

Tivoli Integration

- Tivoli Usage and Accounting Mgr
- Tivoli Access Manager (U.S.)

NFS Proxy Caching system

Ideal for WAN environments

NDAF – NFSv4 management

Network Data Administration Facility

Manage large, replicated filesystems

NIM and installation enhancements

New command *geninv* & *niminv*

Installation support for >1TB LUN

AIX 6: The Next Step in the Evolution of UNIX®

Introducing AIX 6 (a.k.a AIX 5.4)

AIX 6 is **binary compatible*** with AIX 5L™
It is *named* to reflect it's unity with POWER6™

The POWER of SIX – AIX 6 and POWER6

- **Workload Partitions**
- **Live Application Mobility**

} **AIX 6**

- Live Partition Mobility
- Storage Keys
- Hardware Decimal Floating-Point
- Dynamic Variable Page Size

} **POWER 6**

IBM System p Innovation and
Advanced POWER™ Virtualization
Provide Unique Features for ISV and Customer Exploitation

*Complete details on AIX binary compatibility can be found at http://www.ibm.com/servers/aix/os/compatibility/

Planned Smooth Upgrade to AIX 6

- **AIX 6 is binary compatible with AIX 5L**
  - Current applications will continue to run
  - Runs on POWER4™, POWER5™, POWER6 systems
  - Open beta will provide early access to AIX 6
  - Other activities planned to assure ISVs
- **No charge upgrade for current AIX 5L clients with SWMA**
  - No additional out of pocket expense for clients
- **Upgrade process**
  - Tools like alt disk installation and multi-bos minimize client risk

Two Planned WPAR AIX Offerings in 2007

- AIX 6
  - Workload Partitions (WPAR) included in base AIX 6
  - Element (single system) WPAR Management

- Workload Partitions Manager
  - Enablement for Live Application Mobility
  - Cross System Management for Workload Partitions
  - Automated, Policy-based Application Mobility
  - Part of the IBM System Director Family

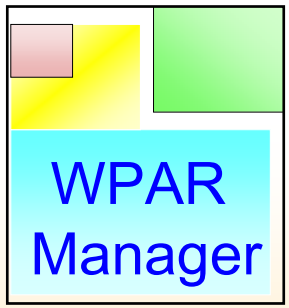

*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Workload Partitions provide Simplified System Management

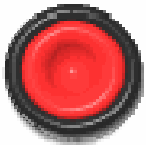

One button install and configuration

- *Installs in minutes*
- *Highly configurable and flexible*
- *Command line, SMIT interface, or template*

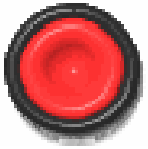

One button startup/boot

- *Workload Partition boots in seconds*

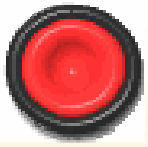

One button stop/shutdown

Planned Role Based Access Control

- **Improved Administrative Security**
  - Improved security by reducing the need for many root users
  - Reduced administration cost thought delegation

aix

- device
- fs
- network
- proc
- ras
- security
- **system**
- wpar

boot

- config
- install
- stat

create	“create boot image”
Halt	“halt the system”
Info	“display boot information
Reboot	“reboot the system”
Shutdown	“shutdown the system”

auth = aix.system.boot.create

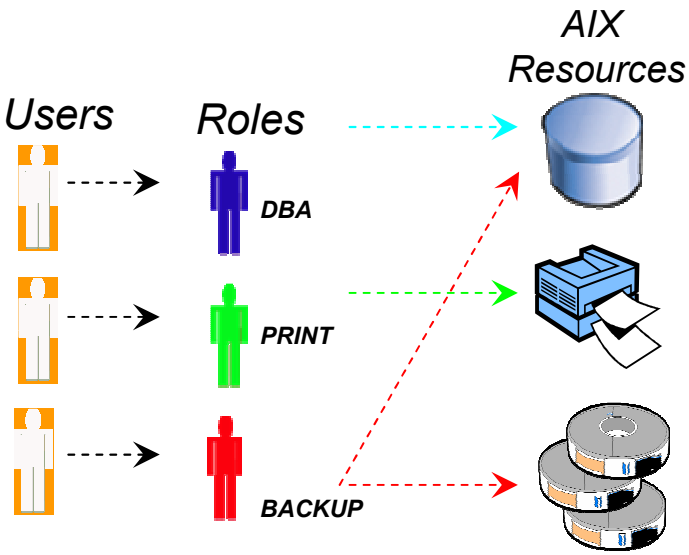

- **Improved Program Security**
  - Allows programs to do system level operations without running as root or having setuid root capability
  - Only allow program to perform restricted set of needed operations

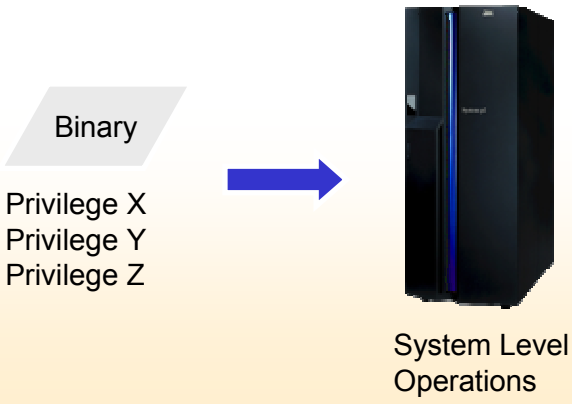

Encrypted File System

- **Backup in encrypted or clear formats**
- **Can be protected from root access to encrypted data**
- **Integrated into user and group administration**
- **Automatic key store creation on user creation**
- **Key store open on login, integrated into AIX security Loadable Authentication Module**
- **Each file encrypted with unique key**
- **No keys stored in clear in kernel memory**
- **Key stores in PKCS12 format.**
- **AES, and RSA Cryptography**

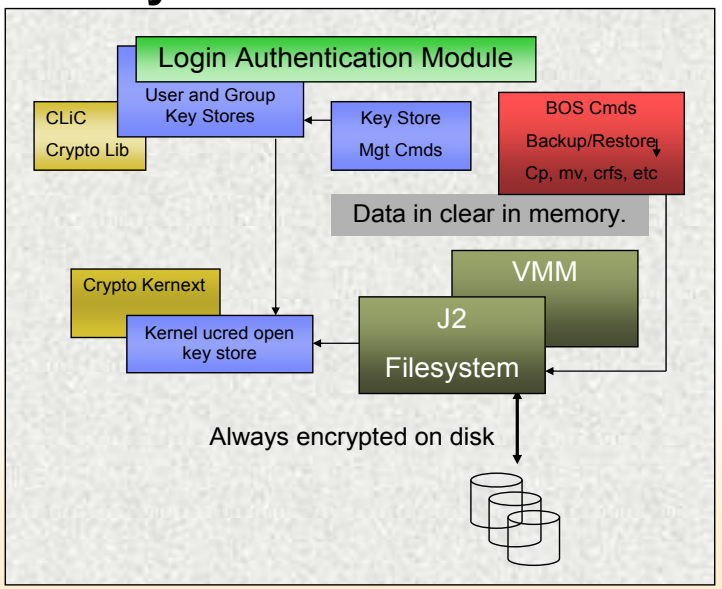

AIX Planned Concurrent Maintenance

Fix selected AIX kernel problems without a service outage

- **Non-disruptive fixes to executable code in a running AIX kernel**
  - Base AIX Kernel (/unix), kernel extension, or device driver
- **No downtime (reboot) required to apply fix and make it active**
- **Concurrent updates will be packaged as Interim Fixes**

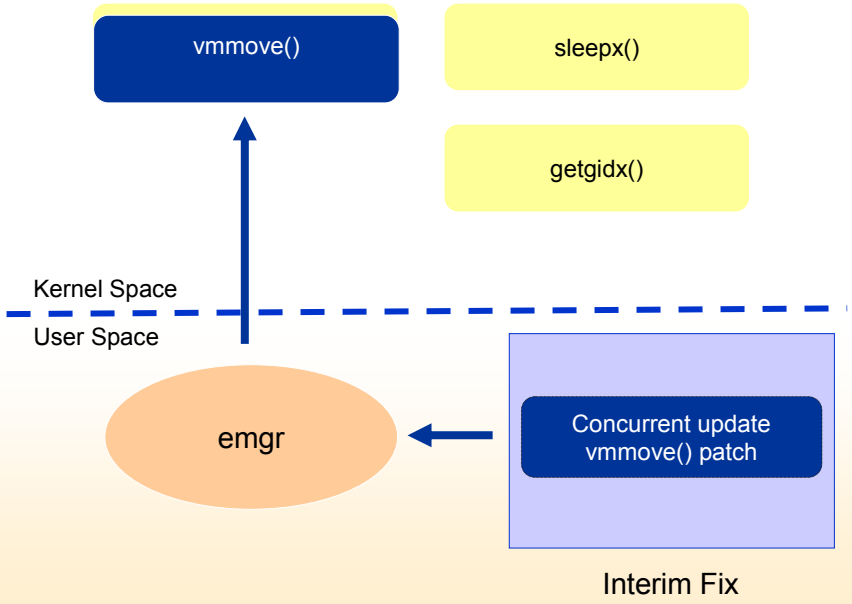

AIX Planned Storage Keys

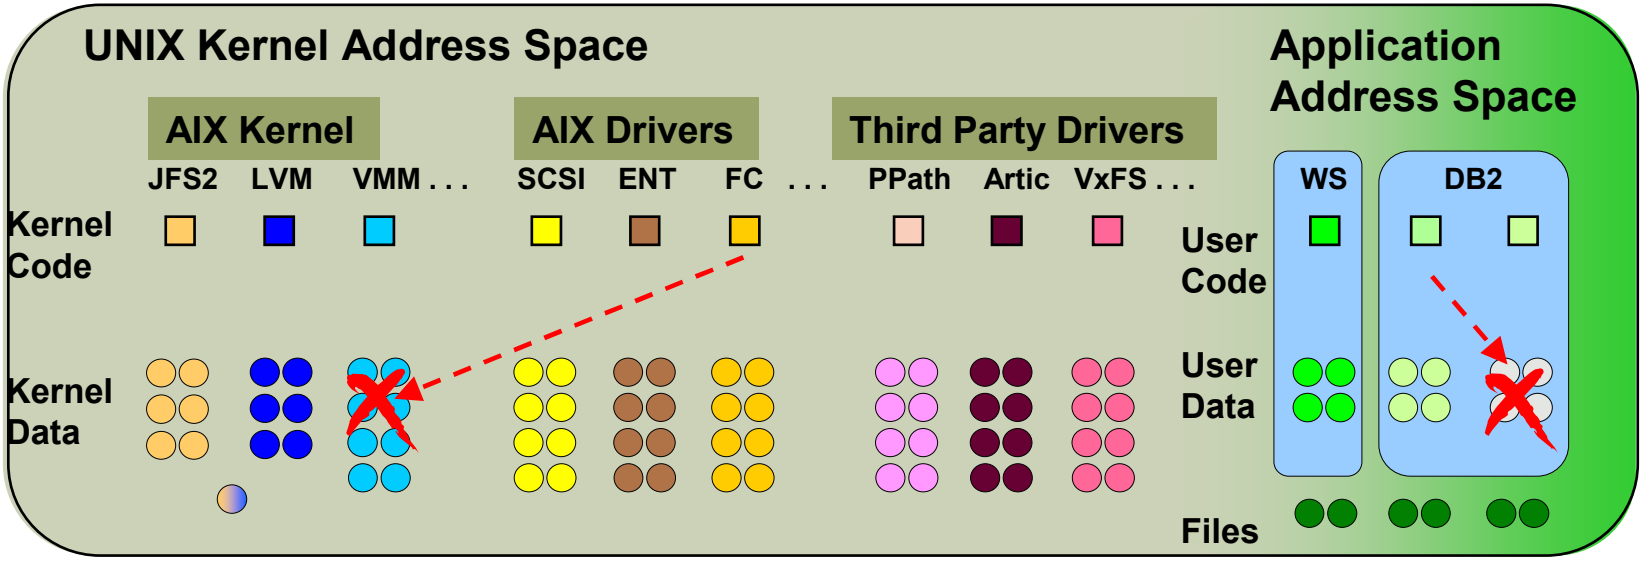

- **In current UNIX implementations, any kernel routine can overwrite any kernel memory**
  - **Memory overlay can cause subtle, intermittent problems**

AIX Planned Storage Keys

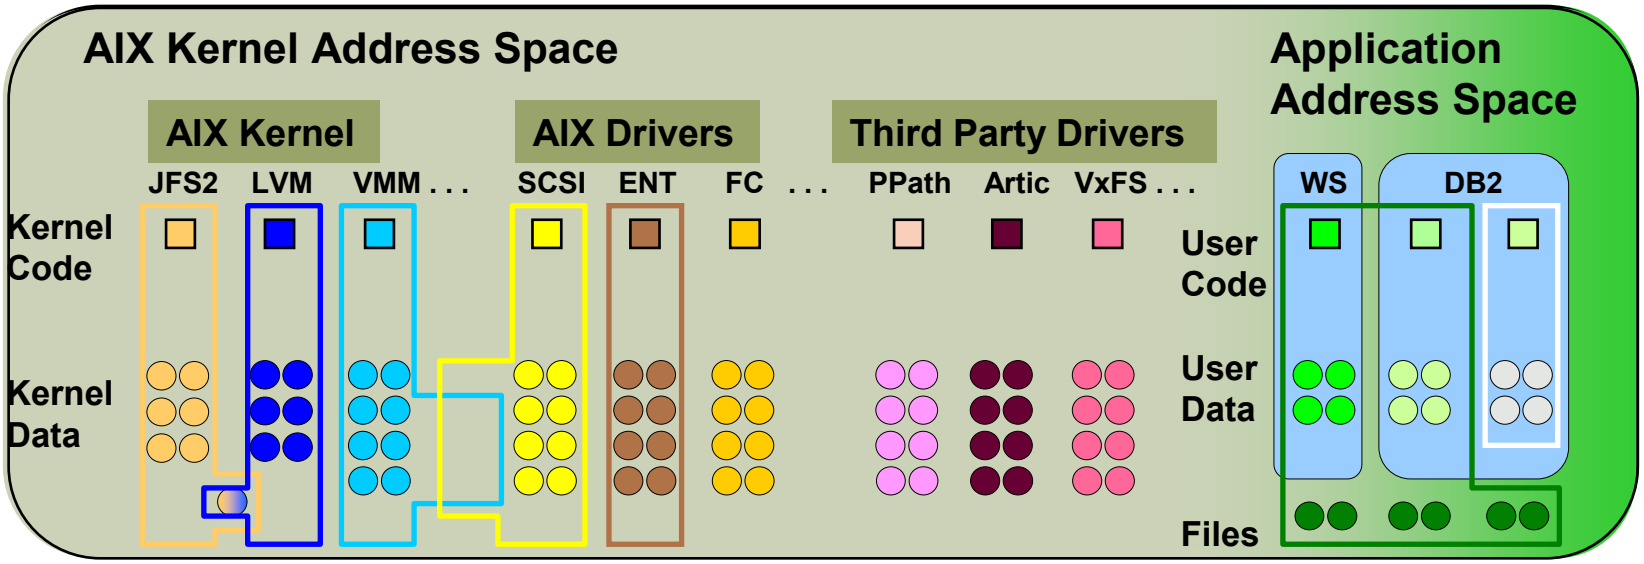

- **In current UNIX implementations, any kernel routine can overwrite any kernel memory**
  - Memory overlay can cause subtle, intermittent problems
- **POWER6 Storage Keys will isolate data and protect against corruption**
  - Enabled through POWER6 H/W & provides isolation between subsystems or subsystems classes
- **Initially provide eight keys w/POWER6**
  - More keys brings finer-grain isolation and better protection
- **Extensible to applications to protect against corruption within the application**
  - AIX will provide enablement to allow applications to exploit keys
- **Application keys with AIX V5.3 - Kernel key exploitation with AIX 6**

Planned Dynamic Tracing With *probevue*

- **Trace existing programs without recompiling**
- **Dynamic placement of trace probes**
- **For debugging and performance analysis**
- **AIX system calls, application functions, and application calls to library functions traceable**
- **Dynamic tracing language called Vue**
- **Initial support for "C" programs**

"Vue" probe code example

```
#!/usr/bin/probevue
/* countreads.v */

@@syscall.$1.read.entry
{
  count++;
}
@@interval.*.clock.100
{
  printf("Number of reads = %d\n", count);
  count = 0;
}
```

```
# countreads.v 404
Number of reads = 22
Number of reads = 0
Number of reads = 1
Number of reads = 17
.....
```

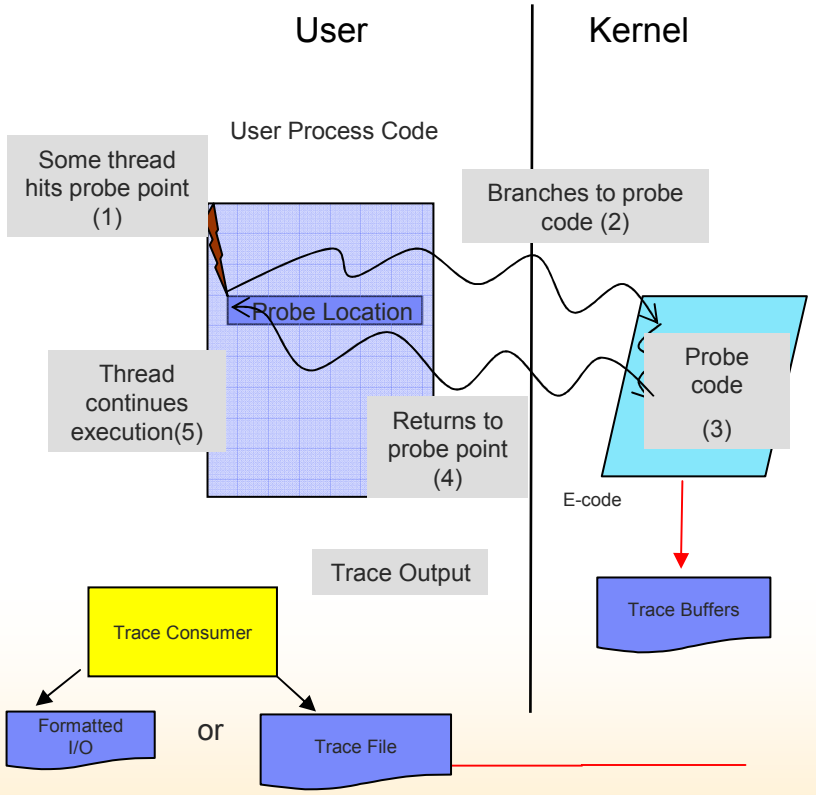

Planned AIX Enhanced Manageability

- **Systems Director Console for AIX**
- **WPAR Management**
- **Integrated Filesystem Snapshot**
- **IBM Director enablement**
- **Tivoli Integration**

Planned Systems Director Console for AIX

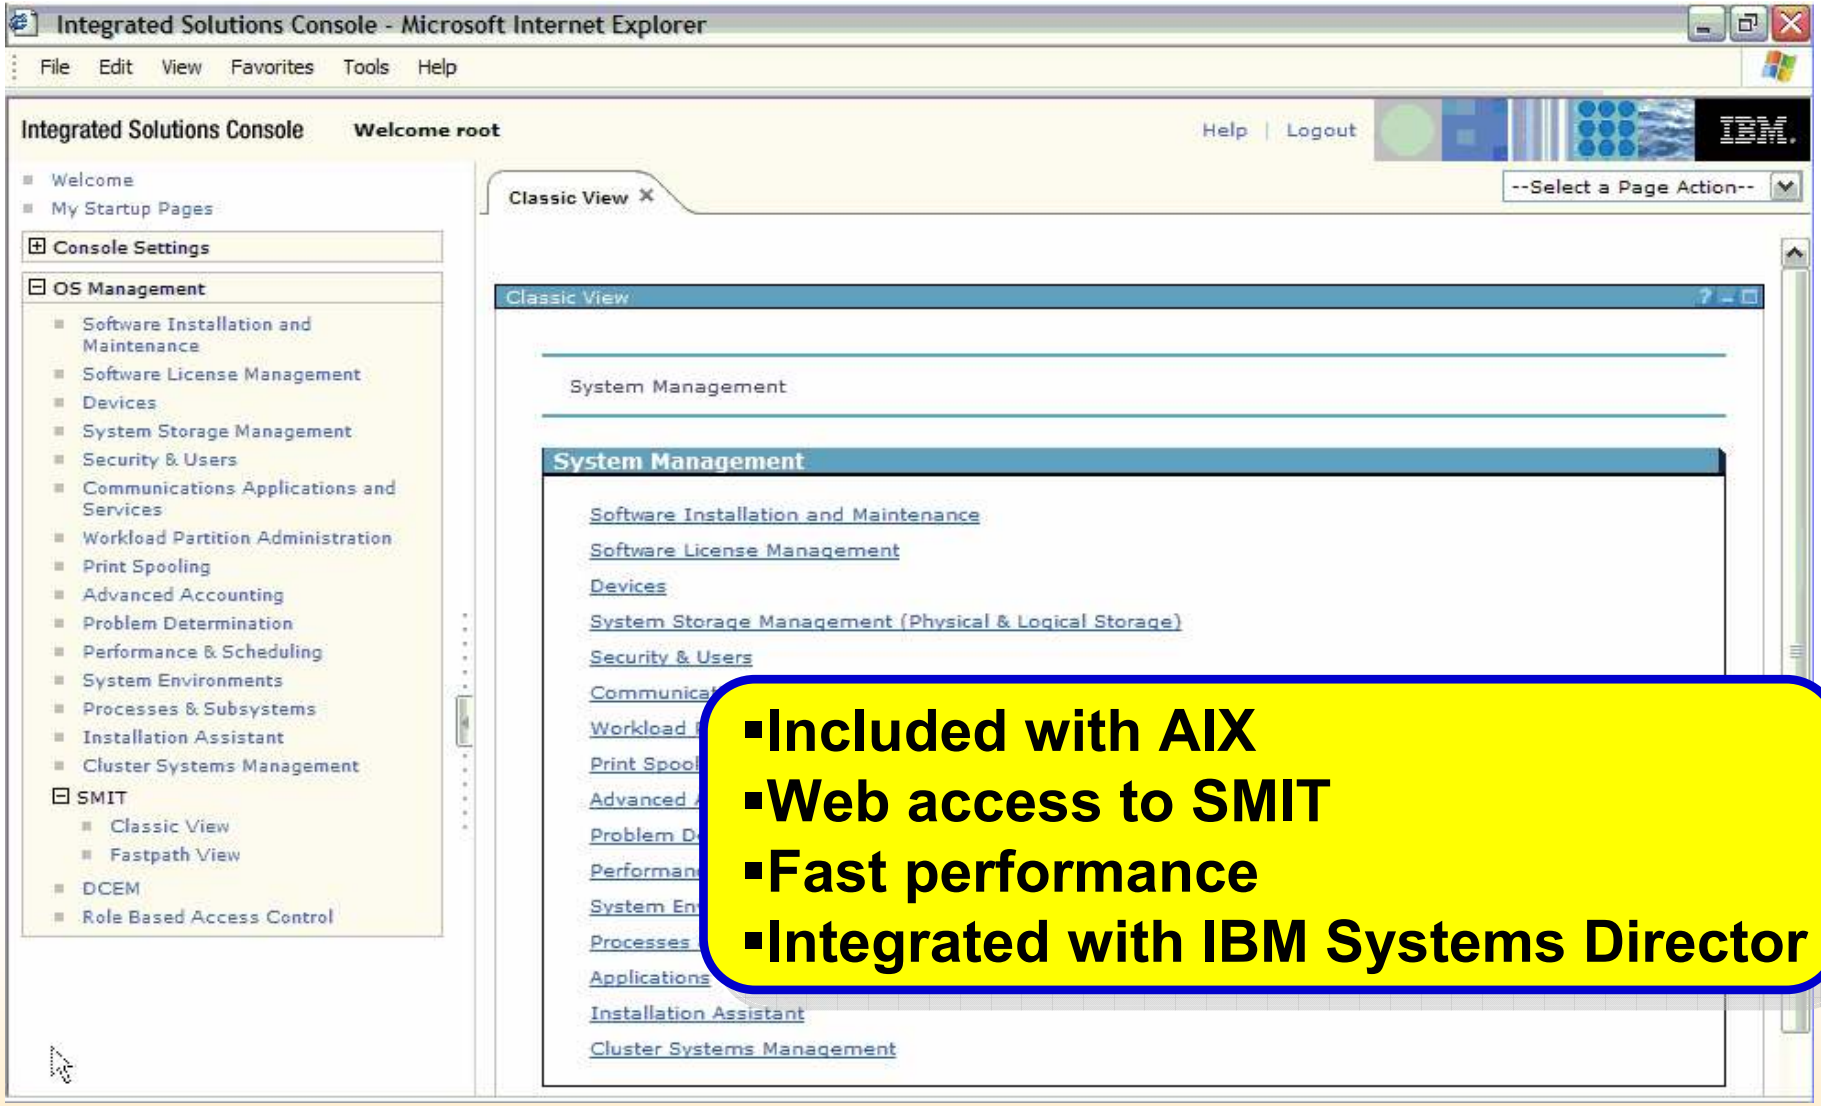

AIX - Planned POWER6 Support

- **2Q07 support POWER6 by AIX 5L V5.2 and V5.3**
  - Technology Level 6 for AIX 5L V5.3
  - Technology Level 10 for AIX 5L V5.2 **(this is the last planned update for AIX V5.2)**
- **POWER6 is Binary Compatible* with previous POWER processors**
- **Exploitation of some features of POWER6 will require AIX 6**
  - **Kernel Storage Keys**
  - **Dynamic, variable page size**
- **But many features of POWER6 will be supported by AIX 5L V5.3**
  - **Live Partition Mobility**
  - **Shared Dedicated processor**
  - **Hardware Decimal Floating-Point**
  - **Application Storage Keys**

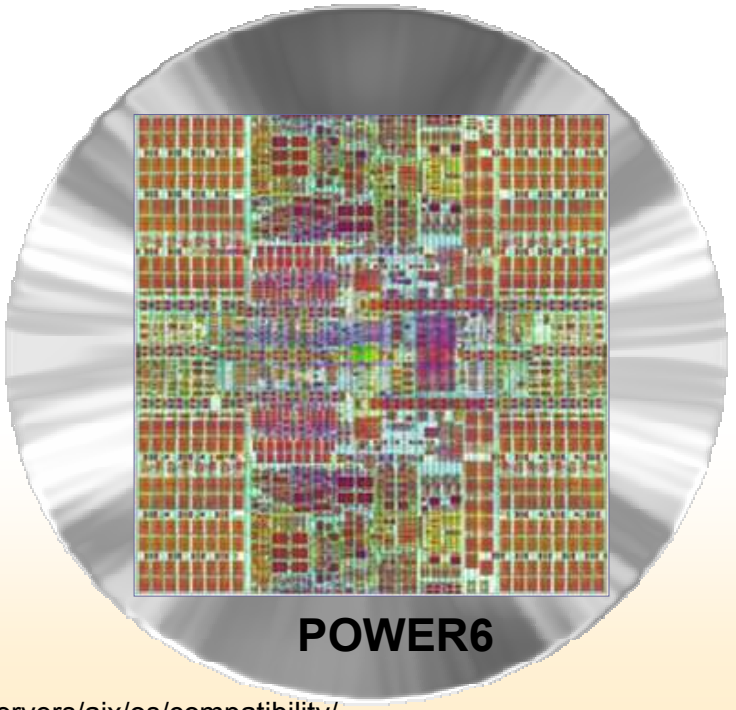

*Complete details on AIX binary compatibility can be found at http://www.ibm.com/servers/aix/os/compatibility/

AIX 6 Hardware Support

- **Systems based on POWER4, PPC970, POWER5 and POWER6 processors will be supported**
- **32- and 64-bit applications will continue to run unchanged on AIX 6**
- **64-bit Kernel only**

*Complete details on AIX binary compatibility can be found at http://www.ibm.com/servers/aix/os/compatibility/

Planned AIX Beta Programs

First ever 'Open Beta' for an AIX release

AIX Beta "Traditional"		AIX "Open Beta"
		
Select customers & ISVs	*Participants*	Open to all
Tens	*Number of participants*	Thousands
Physical Media	*Distribution method*	Web download only
Feedback on functionality	*Goal*	Mind share
Beta Support team	*Support*	Self help via forum
Questionnaire, Support interaction	*Feedback*	Web feedback only

AIX Planned Open Beta Program

Focus Areas

- **Workload Partitions**
- **WPAR Manager**
- **Application Mobility**
- **Role Based Access Control**
- **AIX Security Expert**
- ***probevue* dynamic trace**
- **Director for AIX management**
- **Host Name Caching**

Open Beta Overview

- **Open to everyone**
- **Planned availability – Early 3rd quarter 2007**
- **Simple "Click to accept" license**
- **Not for production use**
- **Limited support – Q&A and Self Help via forum**
- **AIX Developers will monitor forum**
- **Web download only – no physical media distribution**
- **Image will be delivered as multiple ISO CD images**
- **Documentation: "Quick Start Guide" and early pubs**
- **No translation – English only**

Performance Proof Points for Linux on POWER

Linux on POWER holds 24 leadership benchmark results across a range of commercial and HPC workloads – with three #1 overall results and twenty-one #1 Linux results

Current Linux Leadership Results Last updated 10/11/06								Results submitted to SPEC on 10/11/06					
Benchmark	p5 570	p5 570	OP 720	OP 710	p5 575	p5 595	p5 595	p5 505	p5 510	p5 550	p5 575	p5 575	
	1.9GHz	1.9GHz	1.65GHz	1.65GHz	1.9GHz	1.9GHz	1.9GHz	2.1GHz	2.1GHz	2.1GHz	2.2GHz	1.9GHz	
SPECfp 2000											8w SLES		
SPECint_rate 2000											8w SLES	16w SLES	
SPECfp_rate 2000					8w RHEL	32w SLES		2w SLES	2w SLES	4w SLES	8w SLES	16w SLES	
SPECompM 2001	4w SLES				8w RHEL								
Linpack HPC					8w RHEL	32w SLES	64w SLES						
Stream Triad Tuned			4w SLES	2w RHEL									
SPECjbb2000	16w SLES					32w SLES							
SPECweb99	4w RHEL	8w RHEL (Overall											
SPECsfs97_R1.v3		4w SLES	4w SLES										
TPC-C	4w RHEL												Totals
Subtotal	4	2	2	1	3	3	1	1	1	1	3	2	24
#1 Leadership n-way	1	1	0	0	0	1	0	0	0	0	0	0	3
Best of Linux n-way	3	1	2	1	3	2	1	1	1	1	3	2	21

Source: spec.org, November 2006

Turbocharge Linux with POWER6 platform

Performance
Flexibility
Availability

Performance Of Six

- Breaking the 4 GHz barrier: up to 2X the performance of POWER5
- Decimal Floating-point
- Altivec support in GCC toolchain with autovectoring

Availability of Six

- New reliability features: processor instruction retry & storage keys
- Processor recovery
- Live Partition mobility SoD [11/07]*
- Concurrent maintenance SoD [11/07]*

Efficiency Of Six

- More performance than POWER5 in same energy envelope
- Lower energy at same performance as POWER5
- Reduce energy costs through leadership virtualization
- Energy monitoring and savings capabilities

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

System p Application Virtual Environment for x86 Linux

- **What is System p Application Virtual Environment for x86 Linux (System p AVE - x86)**
  - Supports installation and running of existing 32-bit x86 Linux applications[1,2]
  - Creates an x86 Linux application environment running on Linux on System p
  - Extends value of IBM System p and BladeCenter JS21 to x86 Linux apps

- **How does it work?**
  - Dynamically translates and maps x86 Linux instructions to POWER
  - Mapping and caching techniques are used to enhance application performance within the System p AVE-x86 environment

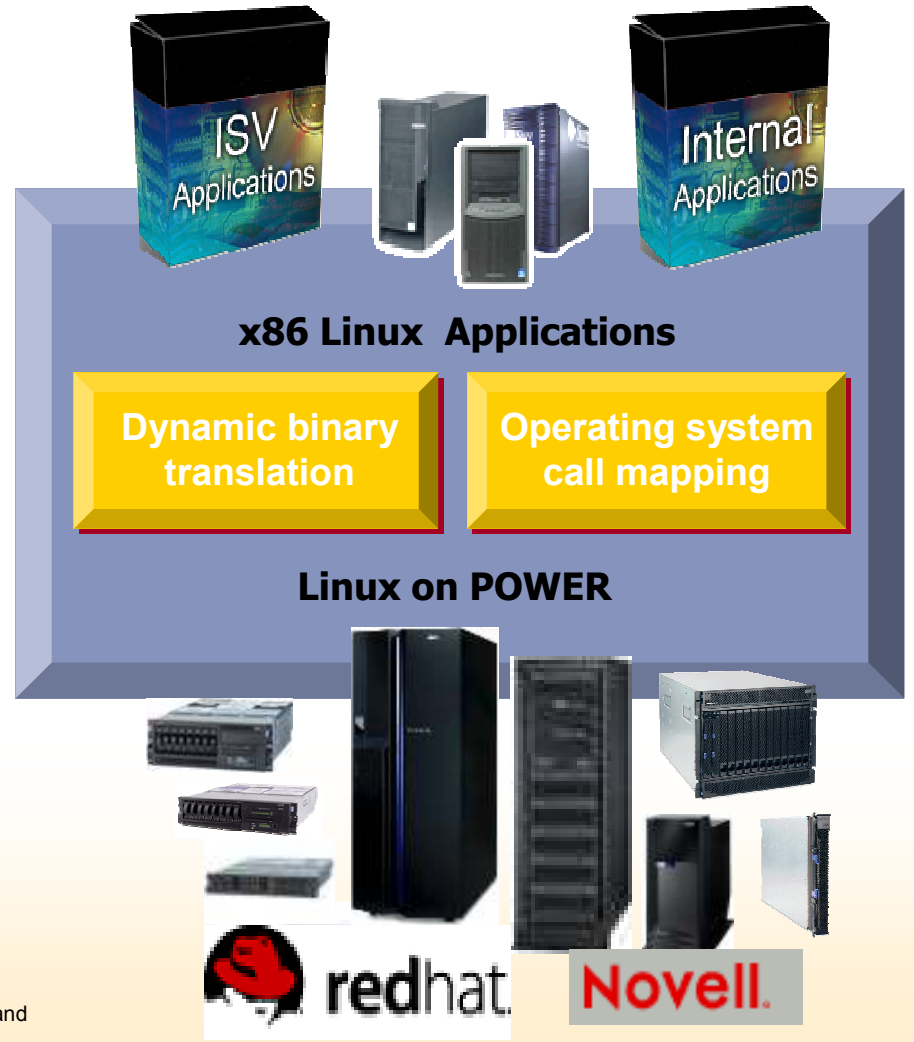

(1) No direct hardware access and no kernel access
(2) IA-32 instruction set architecture (x86) *
* As defined by the 1997 Intel Architecture Software Developer's Manual consisting of Basic Architecture (Order Number 243190), Instruction Set Reference Manual (Order Number 243191) and the System Programming Guide (Order Number 243192) all dated 1997.

Linux Distribution Support of POWER6™ platform*

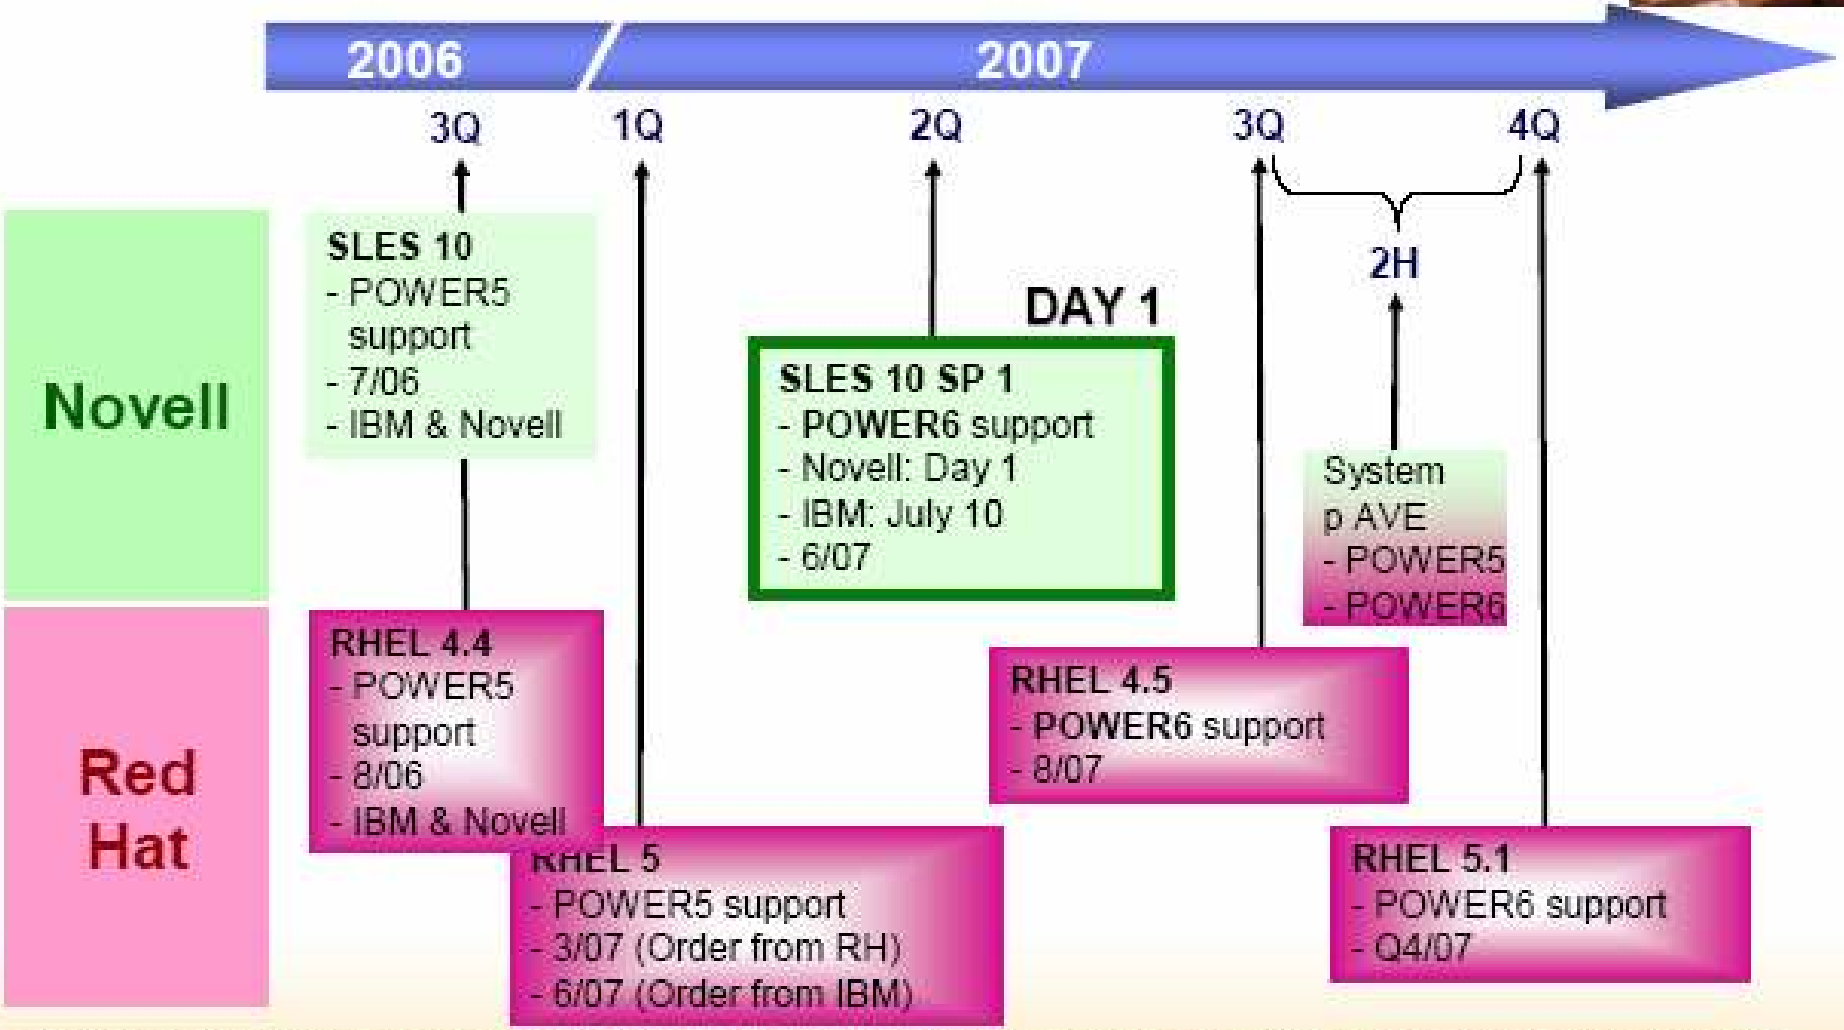

*All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Linux on POWER benefits

Virtualization
Exceptional Virtualization for Linux
Partitions 1/10 with .01 Increments
Low Overhead Hypervisor
Shared Dedicated Capacity* *(NEW)*
Integrated Virtual Ethernet *(NEW)*

Security
Native Linux Security
Hardware Virtualization
Virtual Storage Protection Keys
(NEW!)

Manageability
Installation Tool Kit for POWER
HMC Partition Management
IBM Director/Tivoli®
New Management Tools (NEW!)

Reliability, Availability, Serviceability
Hardware Service Processor
Self Managing/Self Healing
High Availability Solutions
Processor Recovery *(NEW!)*
Live Partition Mobility SoD [11/07]** *(NEW!)*

Scalability

Performance

* Available on Novell SLES 10 SP1 and Red Hat Enterprise Linux (RHEL) 5.1. Not available on RHEL 4.5

** All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Linux on POWER*

- **Enterprise Class Computing with Linux on POWER (LoP)**
  - **Flexibility with LPAR, dynamic LPAR, and virtualization features**
  - **Reliability with built in self-healing capabilities**
  - **Power Architecture/Servers include POWER5, POWER6 and PowerPC97(JS20, JS21 blades)**
- **Linux distributions available for LoP:**
  - **SUSE LINUX Enterprise Server 9 or 10 for POWER (SLES 9, SLES 10)**
  - **Red Hat Enterprise Linux 4 for POWER (RHEL)**
- **Technical support available through IBM SupportLine contract.**
  - **SUSE LINUX and Red Hat, Inc. also provide support, upgrades and maintenance**
- **Orderable from IBM or directly from Linux distributors**
- **For more information about Linux running on IBM System p servers:**
  - **http://www.ibm.com/systems/p/linux/**
  - **http://www.redhat.com/rhel/server/**
  - **http://www.novell.com/products/server/**

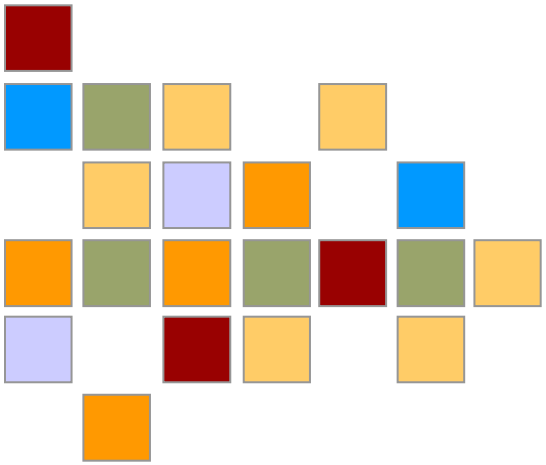

Operating System Release / Service Strategy

New AIX Service Strategy

IBM has made significant enhance the AIX Release and Service Delivery strategy in 2007 for AIX V5.3

The principal changes planned are:

- **Twenty four months of support for each Technology Level (TL)****
- **Service for entire period is provided by Service Packs, PTFs and Interim Fixes**
- **New hardware within the same family will be supported on previous Technology Levels for ease of migration.**

****Due to variations in the release dates of Technology Levels from year to year, some Technology Levels will be supported for slightly more that two years and some will be supported for slightly less than two years. A two year service life for each Technology Level is an objective, not an absolute limit.**
The service life of Technology Levels will also be limited by the end of service life for the underlying AIX release

Release Strategy Transition*

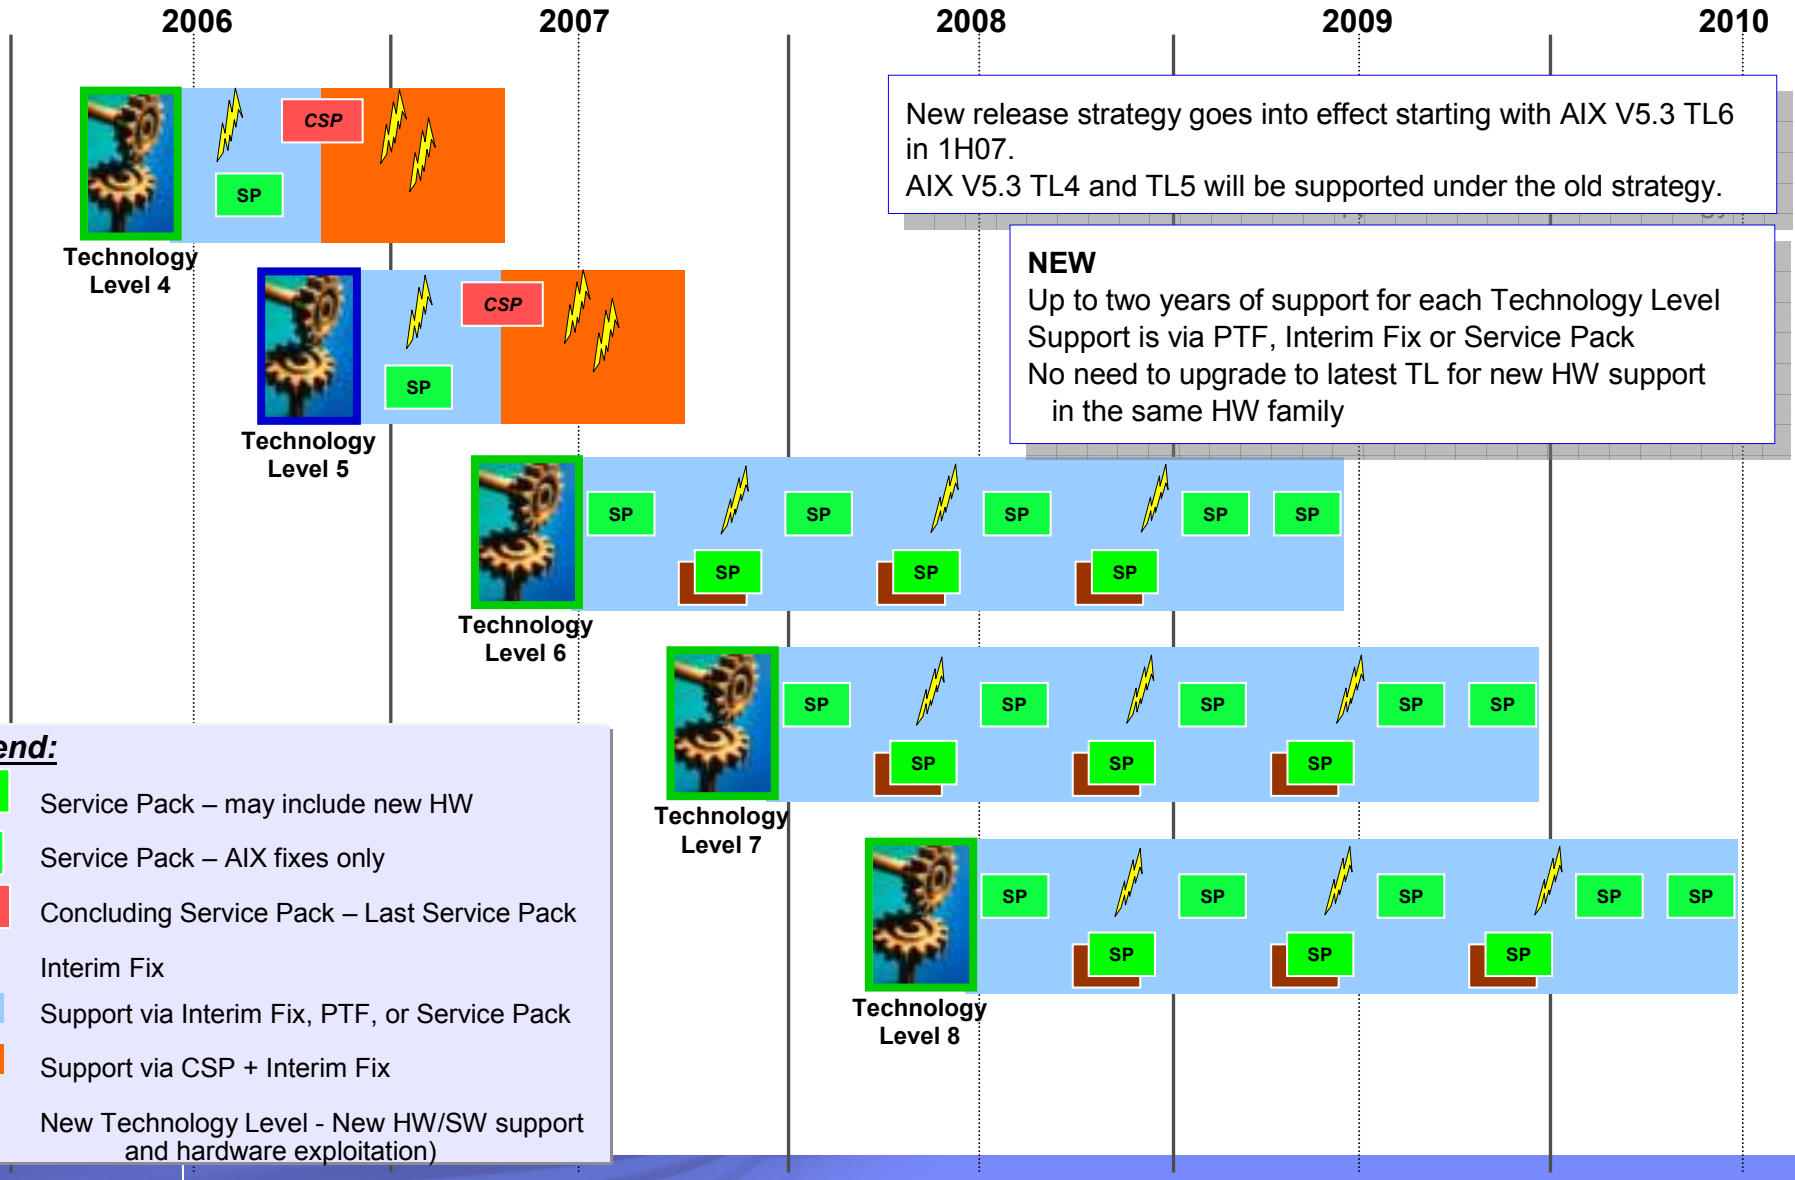

*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

New Hardware Supportability*

Degree of hardware change	On prior TLs plus latest Service Pack	On the latest Technology Level
Processor Speed Increase Only **(No AIX Code Changes)**	**Yes**	Supported
New Processor in Compatibility Mode **(No AIX Code Changes)**	**Yes**	Supported
New Processor in Family **(Recognize New Processor)**	**Yes**	Supported
New I/O **(New Device Drivers)**	**Yes**	Supported
New Technology **(Significant/pervasive)**	**No**	**Yes**

Note: Exploitation of new hardware features will require moving up to the latest TL or in some cases, moving up to the next AIX release

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

Differences From Previous Release Strategy

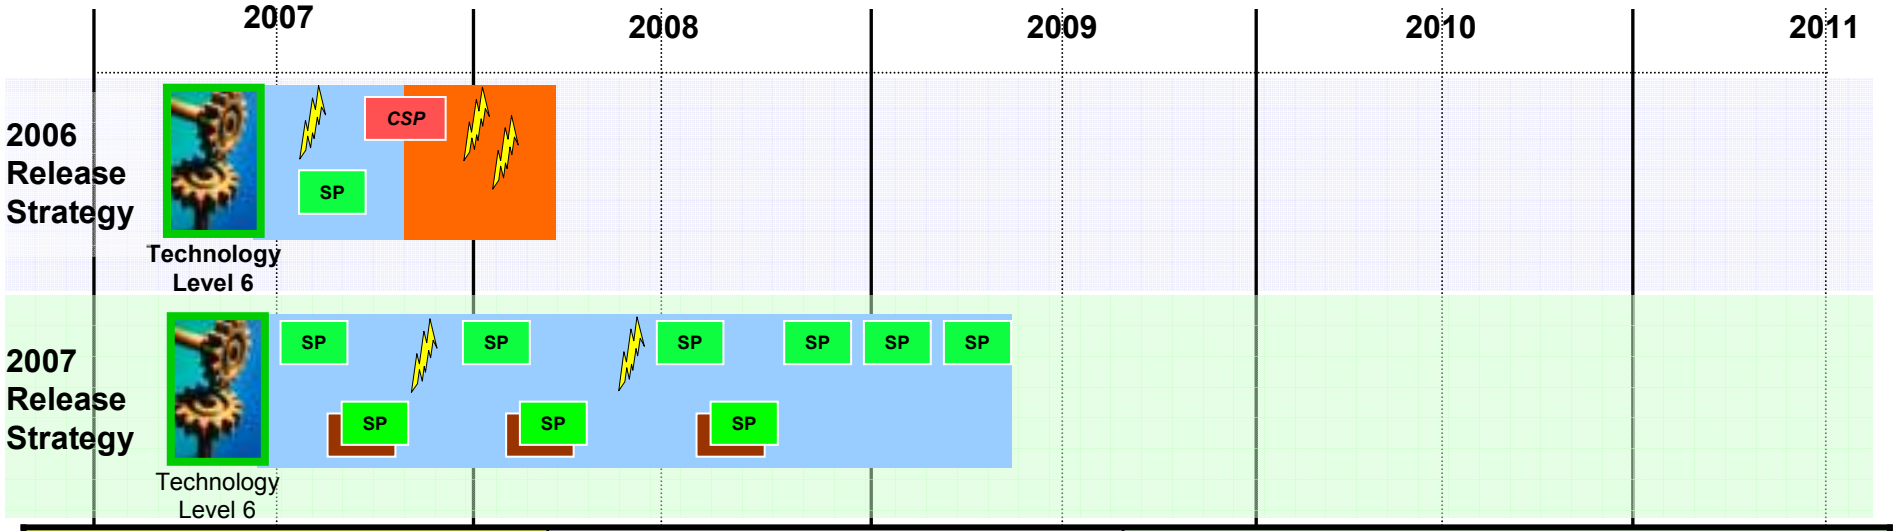

Difference	2006	2007*
Length of Service for a TL	12 months	24 months
Extended service via …	CSP + Interim Fix only	PTF, Interim fix or Service Pack
Concluding Service Pack?	Yes, start of extended service	No longer used
Service Packs include…	Only fixes	Fixes + new HW support within same family
Service Packs ship every…	4-6 weeks	8-12 weeks
AIX releases supported…	AIX V5.2 and AIX V5.3	AIX V5.3 and future releases
Policy started with . . .	AIX V5.2 TL8, AIX V5.3 TL4	AIX V5.3 TL6
Hardware support via…	Latest Technology Level only	Exploitation via latest TL. Some hardware support available via prior TLs plus a SP
Version Release Mod. Fix	5.3.0.41	5.3.7.1

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

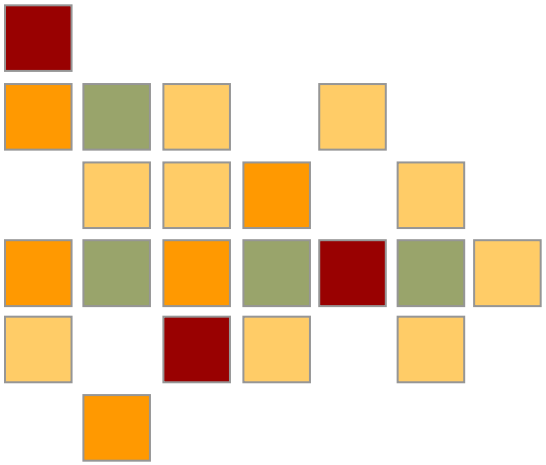

Total Cost of Ownership & Customer Experiences

Oracle introduces socket based licensing

for Standard Edition and Standard Edition One Technology Products

- Socket based licensing significantly advantages p5-505Q, p510Q, p5-520Q, p5-550Q & p5-560Q vs. comparably configured Sun & HP servers running Oracle s/w
- Short window to exploit System p "Q advantage" with Oracle vs. Sun and HP
- 3-year TCA examples: p5-560Q vs. HP rx7640 & 2 x p5-550Q vs. 2 x Sun V490
  - p5-560Q and 2-node p5-550Q RAC cluster running Oracle Standard Edition
  - Sun & HP comparably configured systems running Enterprise Edition – not eligible to run Std. Ed.

HP rx7640 is up to 331% more than p5-560Q

	IBM 560Q	HP rx7640
Max Sockets	4	8
Max Cores	16	16
Nodes/Cluster	n/a	n/a
Oracle Software	Standard Ed.	Enterprise Ed.
3YR TCA HW+OS+SWMA+ Oracle DB + Lic/Upd/Spt	$190,594	$821,941

Sun V490 two-node cluster is up to 617% more than two-node p5-550Q cluster

	IBM 550Q x 2	Sun V490 x 2
Max Sockets	2	4
Max Cores	8	8
Nodes/Cluster	2	2
Oracle Software	Standard Ed.	Enterprise Ed.
3YR TCA HW+OS+SWMA +Oracle DB + Lic/Upd/Spt	$186,586	$1,337,698

All prices are U.S. list price, Oracle Standard Edition and Enterprise Edition pricing and License Update & Support for from Oracle are taken from Oracle Technology Global Price List dated February 16, 2007 at http://www.oracle.com/corporate/pricing/technology-price-list.pdf IBM, Sun, and HP 3yr TCA from IBM MI as of 12/13/06. All compares are 16-core servers for p5-560Q and 8-core servers for p5-550Q with 2 GB memory/core, 2x73 GB drives, 3yr 9x5 NBD & 3yr SWMA NBD Support

System p Products Eligible for Oracle Standard Edition & Standard Edition One Programs

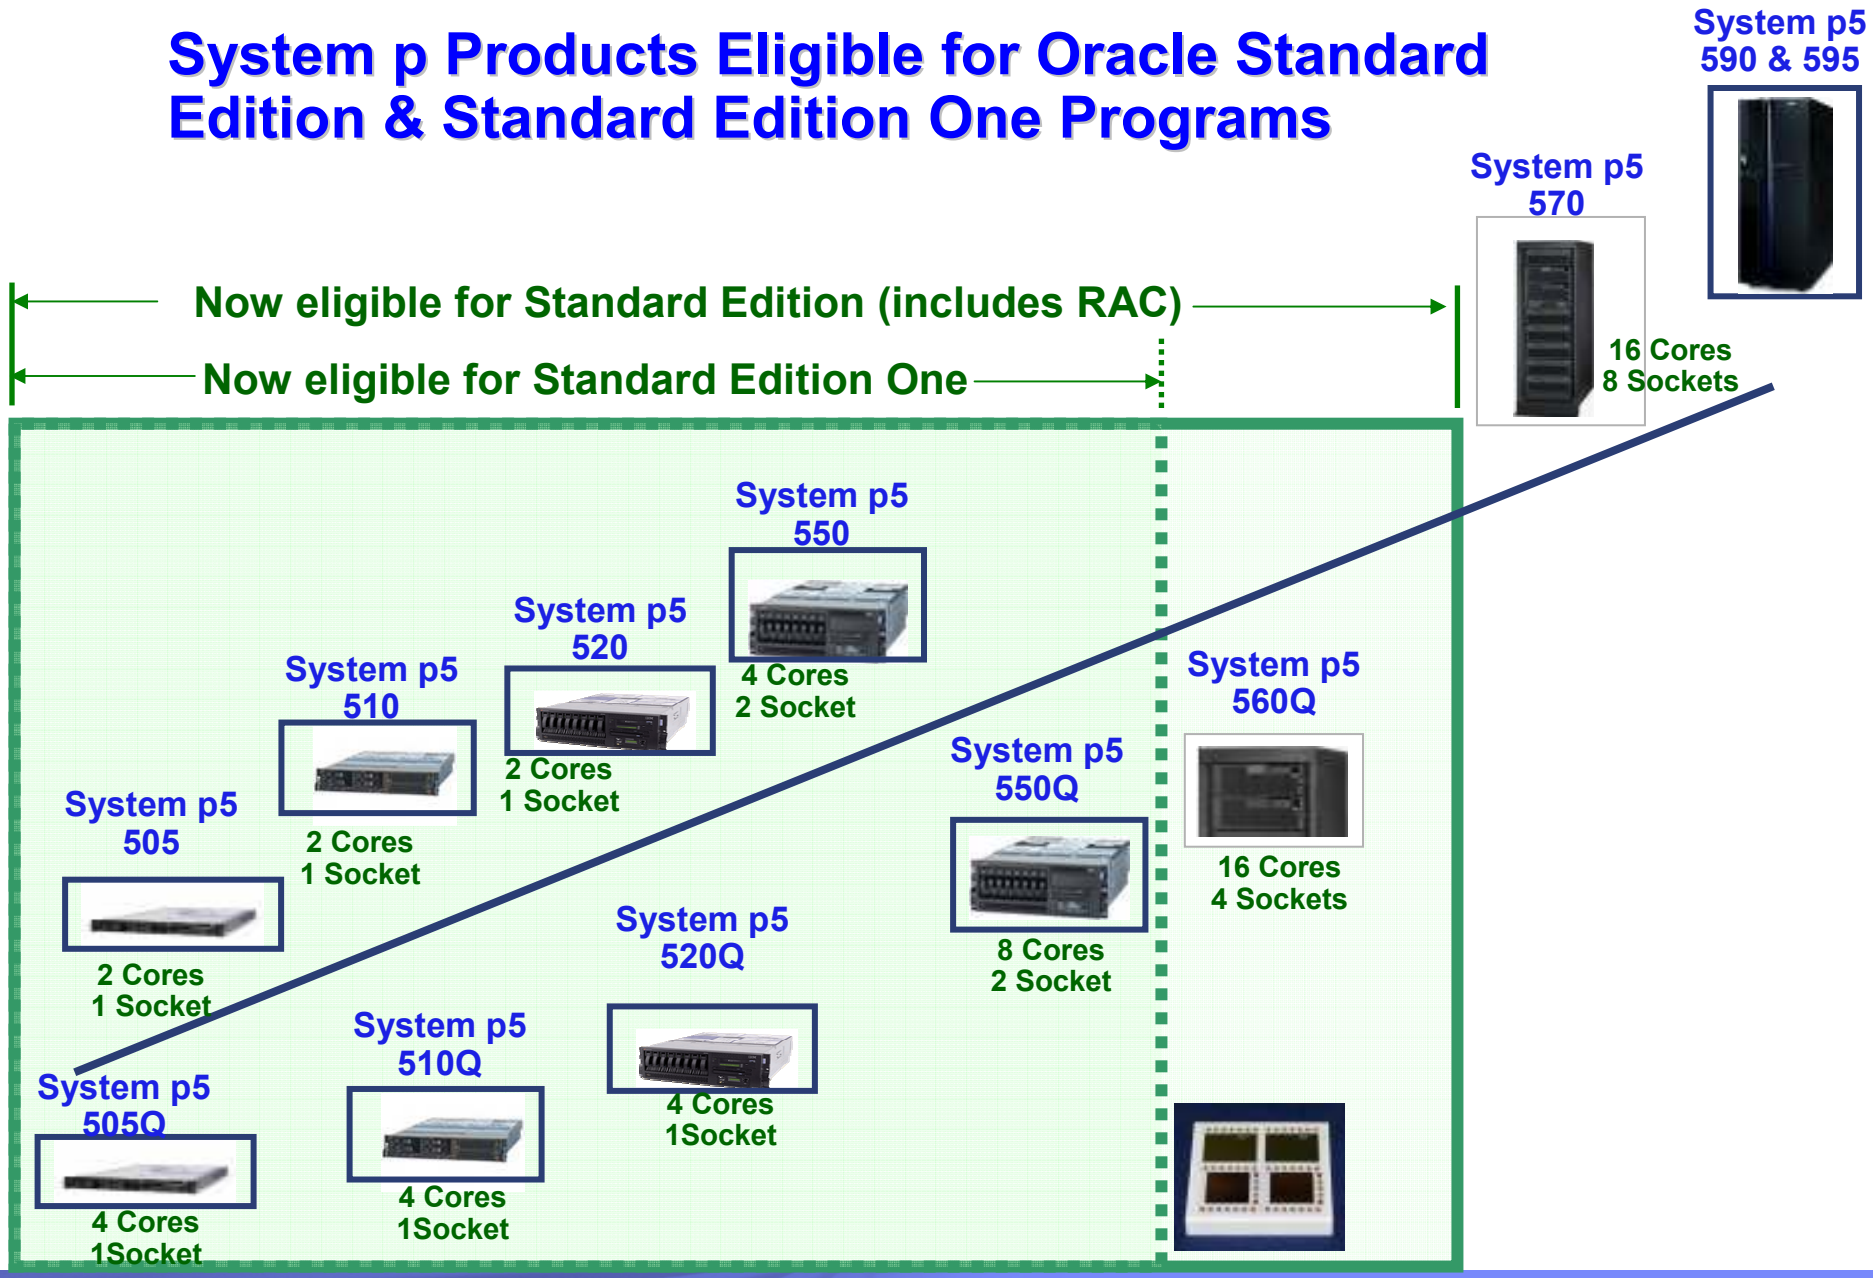

p5-560Q

Oracle TCA Example - single system

- Standard Edition now supports 16-core 560Q

- Sun and HP still require Enterprise Edition for 16-core support

	IBM p5-560Q	Sun V890	Sun X4600	HP rx7640
Oracle SW	SE	EE	EE	EE
Max Sockets	4	8	8	8
Max Cores	16	16	16	16
Nodes/Cluster	n/a	n/a	n/a	n/a

3YR HW/SW/SWMA Comparisons

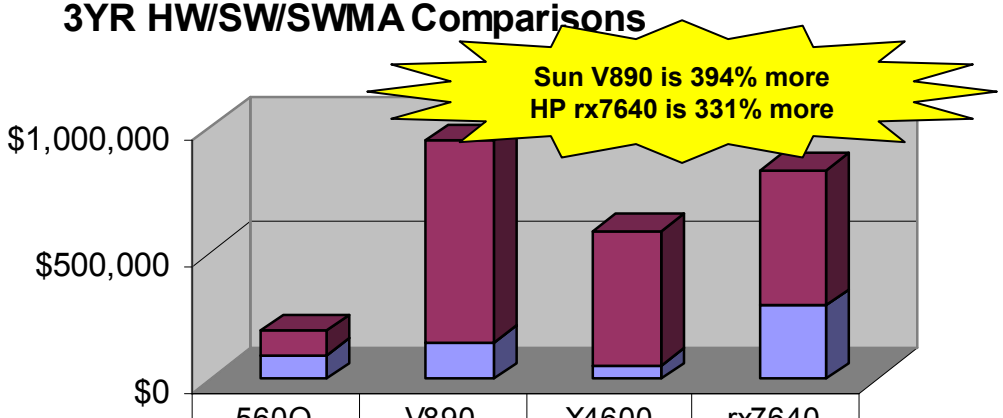

	560Q	V890	X4600	rx7640
3YR Oracle + Lic/Upd/Spt	$99,600	$796,800	$531,200	$531,200
3YR HW+OS+SWMA	$90,994	$144,265	$53,027	$290,741
3YR TCA	**$190,594**	**$941,065**	**$584,227**	**$821,941**

IBM p5-560Q 3YR TCA for HW + OS + SWMA + Oracle + Oracle Lic Upd & Spt

- ✓ Oracle SW plus 3yr License Update and Support cost for Sun and HP is **over 5x 560Q**
- ✓ Oracle SW in each case is greater than the HW + OS TCA
- ✓ Total 3yr TCA: Sun X4600 is up to 207% and Sun V890 is up to 394% **more than 560Q**
- ✓ Total 3yr TCA: HP rx7640 is up to 331% **more than 560Q**

All prices are U.S. list price, Oracle Standard Edition and Enterprise Edition pricing and License Update & Support for from Oracle are taken from Oracle Technology Global Price List dated February 16, 2007 at http://www.oracle.com/corporate/pricing/technology-price-list.pdf IBM, Sun, and HP 3yr TCA from IBM MI as of 12/13/06. All compares are 16-core servers for p5-560Q and 8-core servers for p5-550Q with 2 GB memory/core, 2x73 GB drives, 3yr 9x5 NBD & 3yr SWMA NBD Support

p5-550Q

Oracle TCA Example – single system

-Standard Edition now supports 8-core p5-550Q, 2 licenses req'd
-Sun & HP require 4 licenses for Standard Edition, 8-core

	IBM p5-550Q	Sun V490	Sun V40Z	HP rx4640	HP rx6600
Oracle SW	SE	SE	SE	SE	SE
Max Sockets	2	4	4	4	4
Max Cores	8	8	8	8	8
Nodes/Clust er	n/a	n/a	n/a	n/a	n/a

3YR HW/SW/SWMA Comparisons

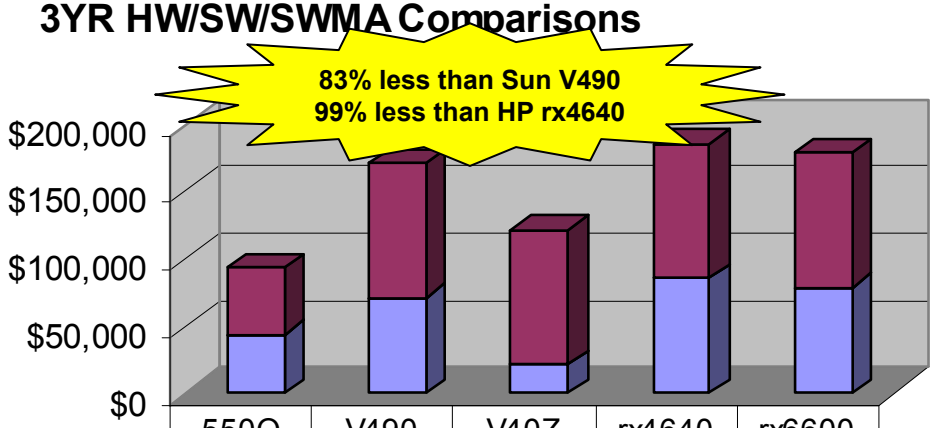

	550Q	V490	V40Z	rx4640	rx6600
3YR Oracle + Lic/Upd/Spt	$49,800	$99,600	$99,600	$99,600	$99,600
3YR HW+OS+SWMA	$43,493	$71,249	$21,280	$85,791	$78,807
3YR TCA	**$93,293**	**$170,849**	**$120,880**	**$185,391**	**$178,407**

IBM p5-550Q 3YR TCA for HW + OS + SWMA + Oracle + Oracle Lic Upd & Spt

- ✓ Oracle SW plus 3yr License Update and Support cost for Sun/HP is **2x 550Q**
- ✓ Total 550Q 3yr TCA is up to 30% less than Sun V40z and 83% less than Sun V490
- ✓ Total 550Q 3yr TCA is up to 91% less than HP rx6600 and 99% less than HP rx4640

All prices are U.S. list price, Oracle Standard Edition and Enterprise Edition pricing and License Update & Support for from Oracle are taken from Oracle Technology Global Price List dated February 16, 2007 at http://www.oracle.com/corporate/pricing/technology-price-list.pdf IBM, Sun, and HP 3yr TCA from IBM MI as of 12/13/06. All compares are 16-core servers for p5-560Q and 8-core servers for p5-550Q with 2 GB memory/core, 2x73 GB drives, 3yr 9x5 NBD & 3yr SWMA NBD Support

Oracle Database Pricing – Enterprise Edition

To determine # of Processor Licenses required, multiply the number of physical cores by the applicable Processor Factor and round up

Server Technology	Processor Factor	Examples
Single-Core Chips	**1.00**	**A 4-processor (4x Single Core Itanium) : 4 x 1.00 = 4 processor licenses**
Multi-Core (Sun Sparc III, IIIi, IV, IV+, IBM Power 4 & 5, HP PA-RISC,)	**0.75**	**An 8-core IBM p570 or a SUN V490 requires : 8 x 0.75 = 6 processor licenses**
UltraSparc T1 *be careful with non-web workloads*	**0.25**	**A "2-processor" (8-core) Sun T2000 requires: 8 x 0.25 = 2 processor licenses**
AMD/Intel Multi-Core	**0.50**	**A "4-processor" (8-core) IBM x366 with Dual-Core Xeon processors requires: 8 x 0.50 = 4 processor licenses**

http://oraclestore.oracle.com/OA_HTML/ibeCCtpSctDspRte.jsp?media=os_g_english_help_licensing&section=11365&minisite=10021&respid=22372&grp=STORE&language=US

Oracle 3 Year License & Support Costs/tpmC

Power5 has the lowest throughput adjusted ($/tpmC) Oracle Licensing & Support costs in the industry

Based on:

- Enterprise Edition
- Per Processor
- Perpetual License
- 3 Yr Software Updates & Support
- E-Business Discount

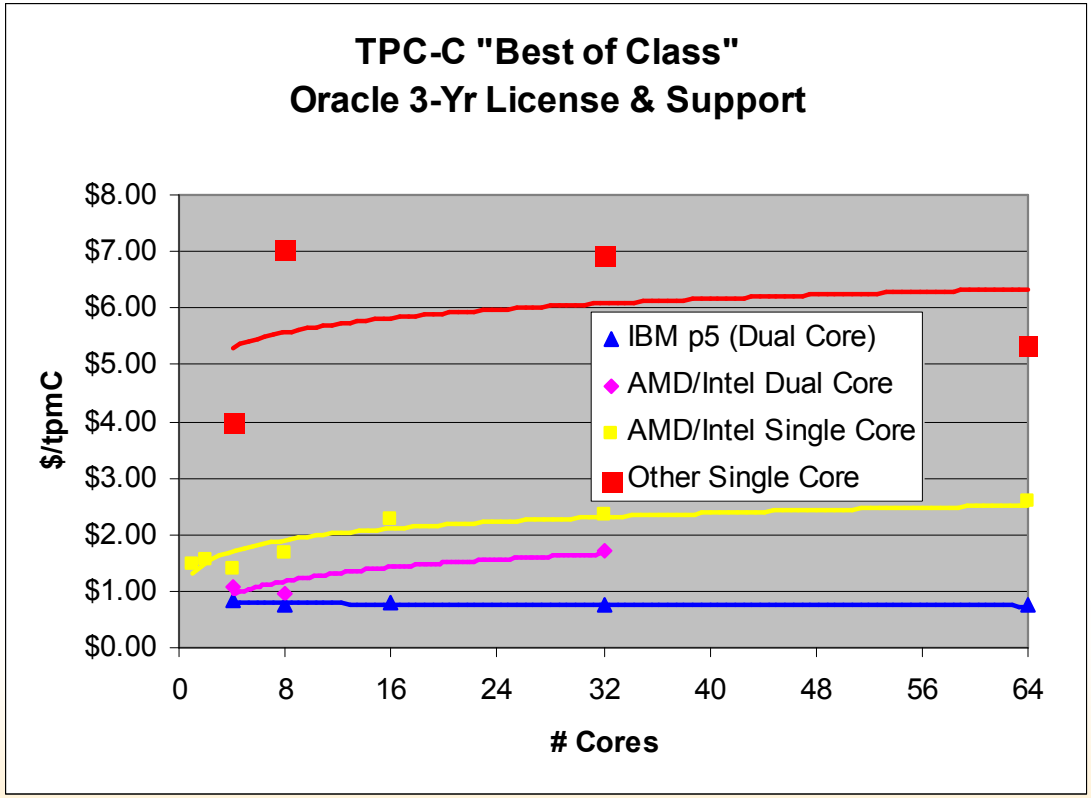

Source: http://oraclestore.oracle.com (as of Jan 17th, 2006)

ITG: IBM System p5 with Virtualization delivers the lowest TCO for SAP, Peoplesoft and Oracle ERP on Unix

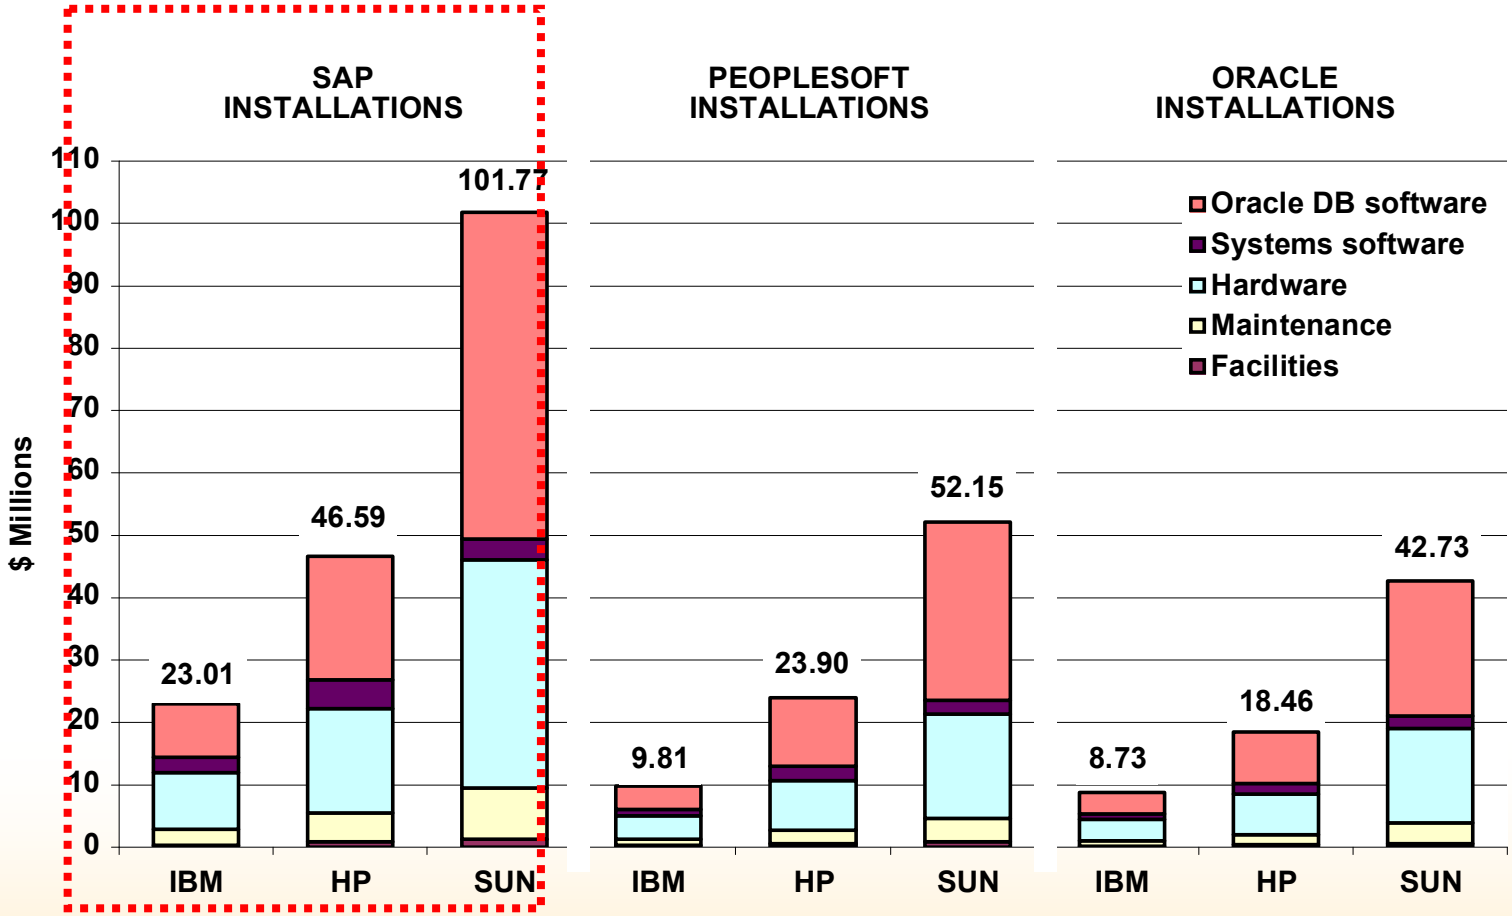

International Technology Group
Los Altos, California

Source: *Business Case for IBM System p5 Virtualization, Economic Benefits of Infrastructure ;* International Technology Group, Feb 2006

Florida Customer: p550 vs. Sun T2000

Customer benchmark of Oracle workload

Server / Workload	p550 4 x 1.65Ghz (4 cores)	TWO T2000 (2) in an Oracle RAC Configuration (16 Cores)	Comments
Oracle Report Gen	17Min	1Hr 43Min	6X better
Oracle Batch	37Min	2+Hrs	3.2X better
Load time	19Min	3 Hrs	9.5X better
Oracle DB licences required	4 x .75 = 3	16 x .25 = 4	Lower Oracle TCO with IBM

Two T2000's each with 8 cores & 32GB in a RAC configuration vs.
One p550 16GB (a second p550 was available for HA but was not used)

Sun later told customer they didn't recommend the T2000 for database loads

POWER5 Sizing Example
Micro-partitioned vs. Standalone

Potential Savings: 25% Hardware, 60% Software, 80+% Infrastructure

	p5-510 (10) Standalone	p5-550 (1) Micro-partition	Savings
Processors	1 x 1.65 GHz	4 x 1.65 GHz	
Memory	1GB	12GB	
Boot disks	2	8	
Ethernet	2 -100 Mb	2 (Gb)	
rPerf	5.24	19.66	
# servers	10	1	
Average utilization	20%	60%	
Peak	1 CPU	4 CPU	
	Totals		
Total list price	$72,370	$55,528	-25%
Processors	10	4	-60%
Power (watts)	6,000	1,100	-82%
Cooling (BTU)	20,460	2,557	-87%
Rack (inches)	35	7	-80%
Ethernet ports	20	2	-90%

All price information is based on IBM U.S. list prices effective Sept.7, 2006. Prices are subject to change without notice and reseller prices may vary.

Insurance Industry

Oracle Test/Development Environment Consolidation

- **Solution**
  - Consolidated seven p650s (27 CPUs) onto a single p5-570 with seven LPARs and 12 CPUs
- **Benefits**
  - Significant reduction in Oracle software licensing costs
  - Able to share resources between fluctuating workloads
  - Reduced floor space and inventory management costs
  - Able to dynamically provision databases as required
  - Able to do full scale performance testing without dedicated low use hardware

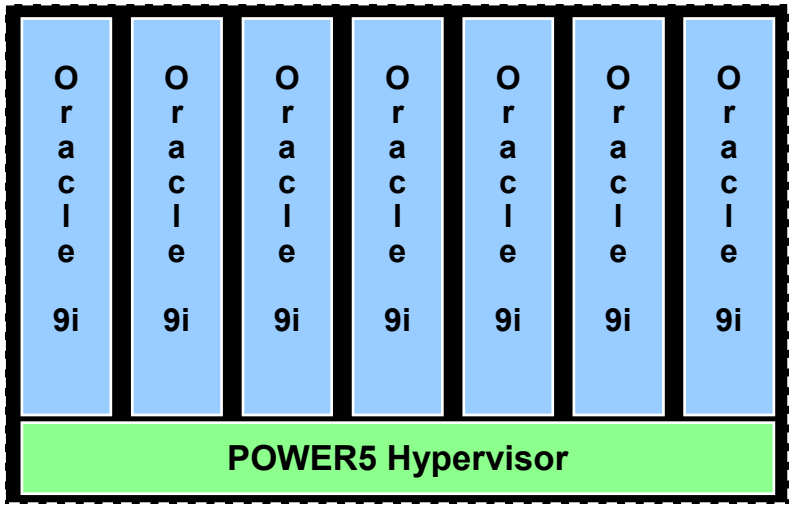

1.65 GHz 12-core p5-570

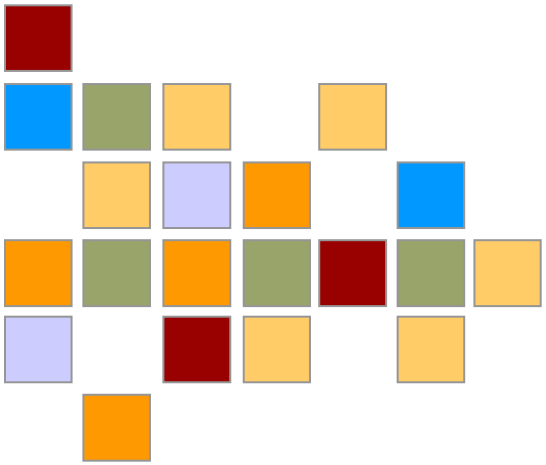

Summary

New IBM System p 570

Modular design for scale-out economics and scale-up performance

- ***Modular building blocks***
  - *--Start with four cores, grow to 16*
- ***All 3 speeds are faster than competition***
  - *--3.5, 4.2 and 4.7 GHz POWER6*
- ***More memory per core – than anyone***
  - *--Up to 768 GB max, 48GB/core!*
- ***Full binary compatibility for investment protection***
  - *-- Existing Apps and AIX® 5L V5.2 & V5.3 run on POWER6*
- ***Price-reduced POWER5+ 570s for customer flexibility***
  - *--IBM System p5 570 is now 20-25% more affordable*
  - *--Upgrades to POWER6 – buy now, upgrade when ready*

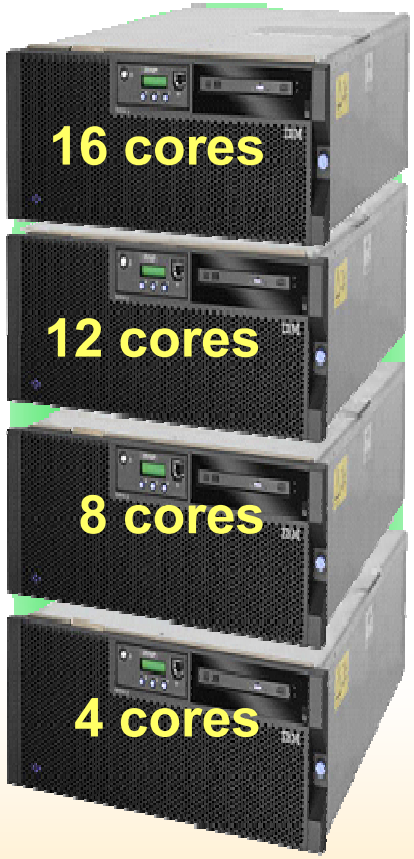

System p Virtualization

Advancing the Most Complete Virtualization Offering for UNIX and Linux

Optimize your IT Infrastructure and Respond to Rapidly Changing Business Needs

- *By balancing workloads across multiple servers*

Eliminate Planned Downtime

- *By moving workloads from one server to another without application disruption*

Improve Power Efficiency

- *By consolidating workloads as business demands change and powering down/off underutilized servers*

Securely Share Systems Resources

- *Through industry-standard certified security capabilities*

New Capabilities

- Live Partition Mobility
- Workload Partitions
- Live Application Mobility
- Shared Dedicated Capacity
- Integrated Virtual Ethernet
- New Management Tools
- IBM System p AVE

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Previewing AIX 6

Investing in the future of the #1 UNIX

Just a few examples of dozens of new features:

Virtualization
Workload Partitions
Live Application Mobility between systems

Security
Encrypting Filesystem

Continuous Availability
Concurrent AIX Updates

Manageability
System Director Console for AIX

A first for AIX! Open Beta: Downloadable AIX coming this summer for clients and ISVs

POWER6 and AIX 6 New Functionality

Feature	Licensed Via			Supported OS			Supported Hardware			GA Date*
	Firm Ware	APV	AIX V6.1	AIX V5.3	AIX V6.1	Linux	POWER4™	POWER5™	POWER6	
Dedicated processor sharing		✓		✓	✓	✓			✓	6/07
Hardware Decimal FP	✓			✓	✓	✓			✓	6/07
Integrated Virtual Ethernet	✓			✓	✓	✓			✓	6/07
Storage keys - application	✓			✓	✓				✓	6/07
Storage keys – kernel*			✓		✓				✓	4Q07
Live Partition Mobility*		✓		✓	✓	✓			✓	4Q07
WPARs*			✓		✓		✓	✓	✓	4Q07
Live Application Mobility*			✓		✓		✓	✓	✓	4Q07

* All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM. Planned availability is 4Q07.

System p APV vs. HP/Sun/VMware

	System p APV	HP vPar	HP IVM	Sun Containers	VMware Infrastructure 3 Standard (EMC)
OS supported	AIX, Linux, i5/OS	HP-UX	HP-UX 11i	Solaris 10	Windows, Linux, Netware
VM granularity	10 LPAR / CPU	1 Partition / CPU	4 partitions / CPU	8,192 containers per Solaris image	8 virtual machines per physical CPU
Security/fault isolation	Hardware & firmware (Hypervisor) - based	Not hardware enforced	Single OS kernel exposed	Single OS kernel exposed	Hypervisor-based
Scalability (Max CPUs per VM)	64	64-128	4 (requires Virtual SMP)	72 (144 cores)	4 (requires Virtual SMP)
CPU Capacity sharing	Yes	No	Yes	Yes	Yes
I/O Sharing	Yes	No	Yes	Yes	Yes
Support for dedicated I/O	Yes (supported)	Yes	No – All I/O must be virtualized	No	No – All I/O must be virtualized
CoD integration	Reserve CoD	Yes (integrated utility pricing)	Yes, GWLM adds temporary instant capacity	No	No

“Checkmate!” with the POWER OF SIX

IBM Innovations

- **IBM System p 570 with POWER6 Technology**
- **System p Virtualization**
- **IBM AIX 6**

IBM Statements of Direction

IBM is committed to enhancing its clients' investments in IBM System p servers. Based on this commitment, IBM plans to provide future enhancements as identified here.

1. **It is IBM's intention to support a future version of Red Hat Enterprise Linux 5 for POWER starting in the second half of 2007 on POWER6 processor-based servers.**

2. **IBM plans to provide redundant service processor capability to existing POWER6 processor-based p570 systems with two or more CEC enclosures. This capability is planned to be provided via a firmware upgrade at no additional charge by the end of 2007. Once this upgrade is applied, the appropriate HMC connections made and the system rebooted, service processor functions can maintain operational status in the unlikely event that one of the service processors fails.**

3. **IBM plans to provide the capability for POWER6 processor-based p570 systems to add an additional CEC enclosure (node) to a p570 system, without powering down the system ("hot-node add"). The additional p570 enclosure would be ordered as a system upgrade and added to the original system while operations continued. The additional resources of the new additional enclosure could then be assigned to existing applications or new applications as required. This capability is planned to be provided at no additional charge to existing POWER6 processor-based p570 users via a system firmware upgrade by the end of 2007. (Additional CEC enclosures are ordered and charged as per normal business practices.)**

4. **In addition, IBM plans to provide the capability for POWER6 processor-based p570 systems that have experienced a failure and rebooted without one of the CEC enclosures active to be able, in certain cases, to have the de-activated enclosure repaired and reintegrated into the active system without powering down the system ("cold-node repair"). The additional resources of the repaired CEC enclosure could then be assigned to existing applications or new applications as required. This capability is planned to be provided at no additional charge to existing POWER6 processor-based p570 users via a system firmware upgrade by the end of 2007.**

5. **IBM plans to provide a new feature, "Live Partition Mobility", as part of IBM System p Advanced POWER Virtualization for POWER6 processor-based servers in late 2007. Live Partition Mobility will allow clients to move a running partition from one physical System p POWER6 processor-based server to another System p POWER6 processor-based server without application downtime helping clients to avoid application interruption for planned system maintenance, provisioning and workload management.**

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

Special Notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised: September 26, 2006

Special Notices (Cont.)

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L(logo), alphaWorks, AS/400, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DirectTalk, Domino, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, i5/OS, IBM, IBM(logo), ibm.com, IBM Business Partner (logo), Informix, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, OpenPower, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, Passport Advantage, POWERparallel, Power PC 603, Power PC 604, PowerPC, PowerPC(logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SpaceBall, System/390, The Engines of e-business, THINK, Tivoli, Tivoli(logo), Tivoli Management Environment, Tivoli Ready(logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Advanced Micro-Partitioning, AIX , AIX PVMe, AS/400e, Chiphopper, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, DS4000, DS6000, DS8000, e-business(logo), e-business on demand, EnergyScale, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, General Purpose File System, GigaProcessor, GPFS, HACMP, HACMP/6000, IBM TotalStorage Proven, IBMLink, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, On Demand Business logo, POWER, PowerExecutive, Power Architecture, Power Everywhere, Power Family, Power PC, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, pure XML, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, System i, System i5, System p, System p5, System Storage, System z, System z9, S/390 Parallel Enterprise Server, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, X-Architecture, z/Architecture, z/9.

A full list of U.S. trademarks owned by IBM may be found at: http://www.ibm.com/legal/copytrade.shtml.
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
UNIX is a registered trademark of The Open Group in the United States, other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
Intel, Itanium, Pentium are registered tradas and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
AltiVec is a trademark of Freescale Semiconductor, Inc.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
Other company, product and service names may be trademarks or service marks of others.

Revised April 17, 2007

Notes on Benchmarks and Values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM System p and BladeCenter Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3 or AIX were used. All other systems used previous versions of AIX. The SPEC CPU2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
NotesBench	http://www.notesbench.org
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle Applications	http://www.oracle.com/apps_benchmark/
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	http://www.siebel.com/crm/performance_benchmark/index.shtm
Baan	http://www.ssaglobal.com
Microsoft Exchange	http://www.microsoft.com/exchange/evaluation/performance/default.asp
Veritest	http://www.veritest.com/clients/reports
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 12, 2006

Notes on Performance Estimates

rPerf

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

Revised April 27, 2006