



# IBM Power Systems

IBM Systems & Technology Group

## Power Systems Update

*Patrick O'Rourke*  
*pmorour@us.ibm.com*  
*Executive Briefing Center*



## Agenda

Power Systems

Processor Update

- ▶ POWER6+

New Product Announcements

- ▶ Power 520
- ▶ Power 550
- ▶ Power Blades
- ▶ IO Drawers
- ▶ IO Adapters
- ▶ SSD


Virtualization

AIX / Linux

# POWER6 Systems Technology Value...


## Technology

- ▶ Roadmap
- ▶ Processor Instruction Retry
- ▶ Green Technology built in
- ▶ Common technology from blades to high-end




## Performance

- ▶ Power Systems scalability from blades to high end systems
- ▶ Performance leadership in a variety of workloads
- ▶ Best Performance per core
- ▶ Memory and IO bandwidth




## Virtualization

- ▶ Consolidate to higher levels
- ▶ Virtualize Processors, Memory, and I/O
- ▶ Dynamic movement of Partitions and Applications
- ▶ Reduce infrastructure costs



## RAS


- ▶ Power Systems mainframe inspired RAS features
- ▶ Hot Add support / Concurrent Maintenance
- ▶ Alternate Process Recovery
- ▶ Operating Systems Availability Leadership




# Processor Update

# Power Processor History

**Dual Core  
1GHz  
Distributed Switch**




**POWER4**  
414 mm<sup>2</sup>  
1.1 – 1.3 GHz




**POWER4+**  
267 mm<sup>2</sup>  
1.5 – 1.9 GHz

**Multi-Threading  
Memory Cntrl on Chip**




**POWER5**  
389 mm<sup>2</sup>  
1.65 – 1.9 GHz




**POWER5+**  
245 mm<sup>2</sup>  
1.9 – 2.3 GHz

**Enhanced Multi-Threading  
Memory Cntrl on Chip  
>4 GHz  
Memory Protection Keys**



**POWER6**  
341 mm<sup>2</sup>  
3.5 – 5.0 GHz



**POWER6+**  
341 mm<sup>2</sup>  
4.7 – 5.0 GHz  
Refresh

2001      2002      2003      2004      2005      2006      2007      2008

## POWER6+

### POWER6+ Memory Keys:

- ▶ Adds 8 new memory keys
  - Helps prevent accidental memory over-writes that could cause critical applications to crash.
- ▶ **POWER6+ has 16 memory keys (8 Kernel, 7 User, 1 Hypervisor) while POWER6 has 8 (7 Kernel, 1 User)**
- ▶ Feature is unique to Power among UNIX, Linux, Windows, or IBM i systems

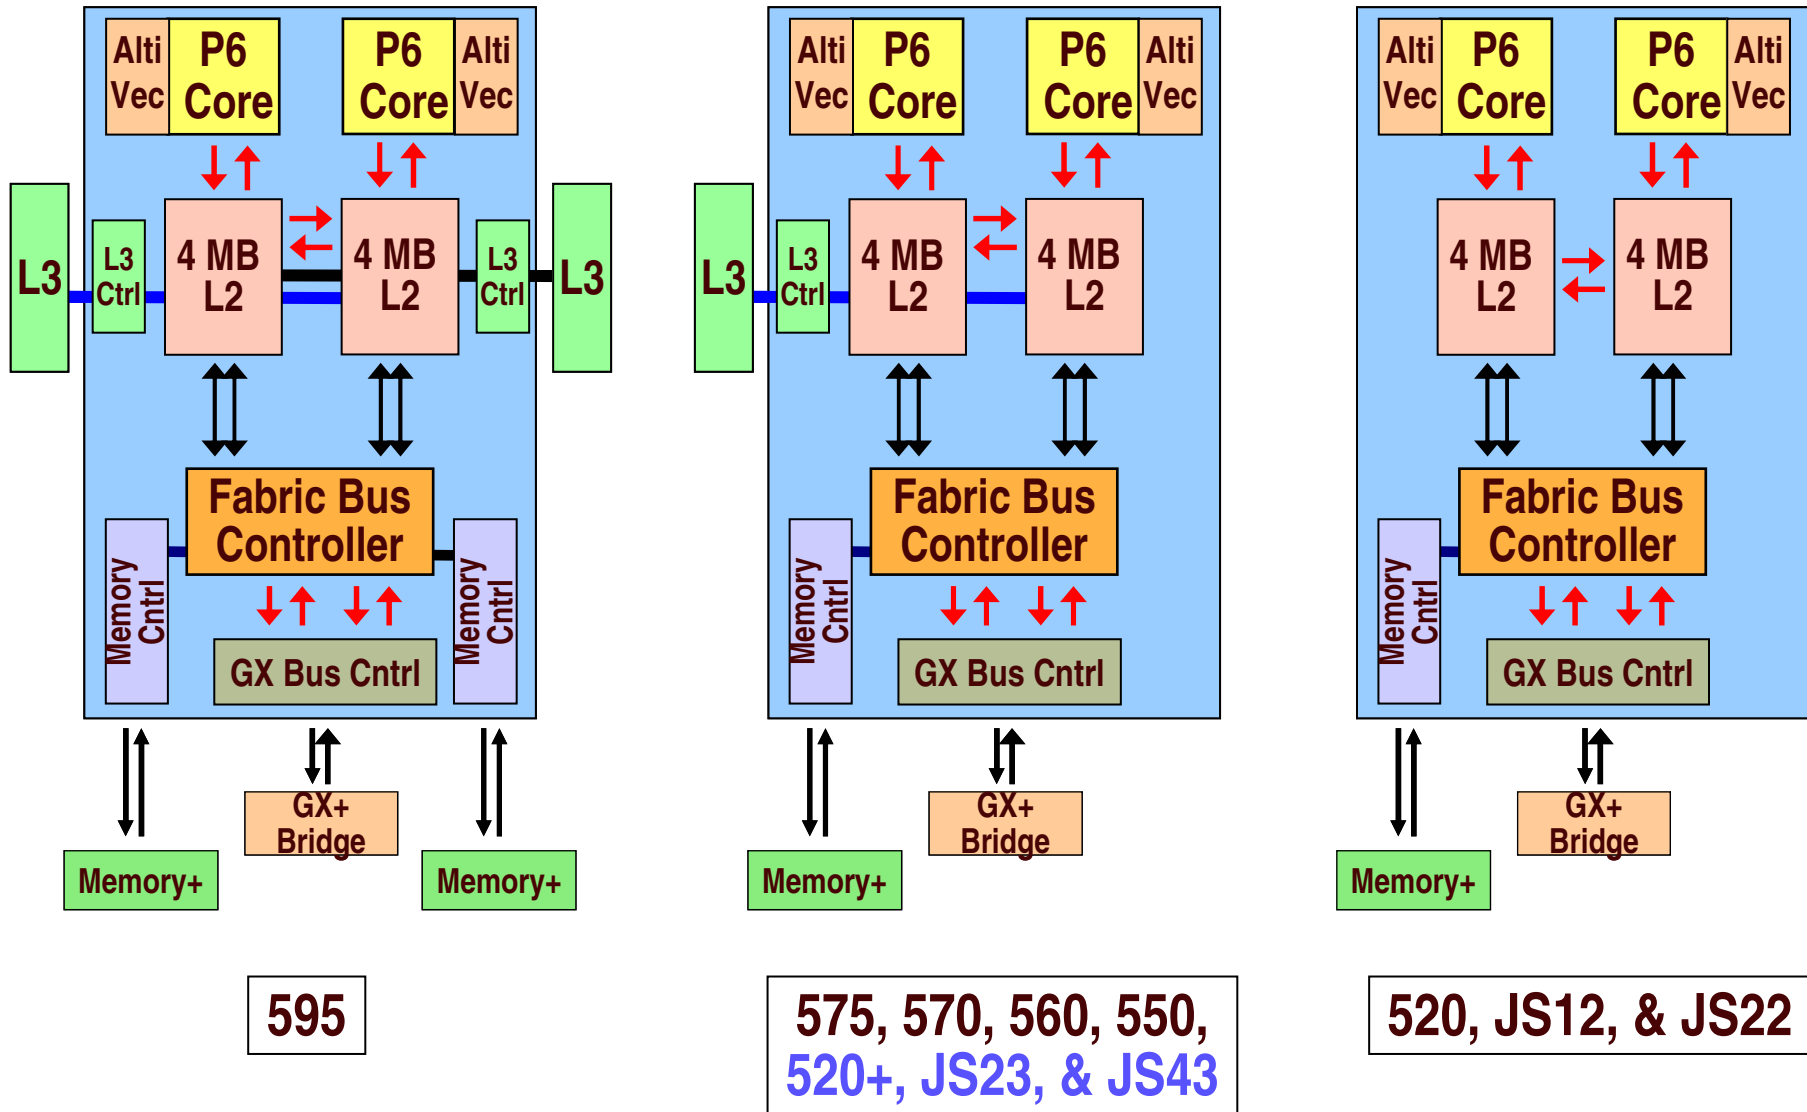
### POWER6+ Systems:

- ▶ JS23 and JS43 Blades
- ▶ Power 520 @ 4.7 GHz
- ▶ Power 550 @ 5.0 GHz
- ▶ Power 560
- ▶ Power 570/16 @ 4.4 & 5.0 GHz
- ▶ Power 570/32 @ 4.2 GHz model

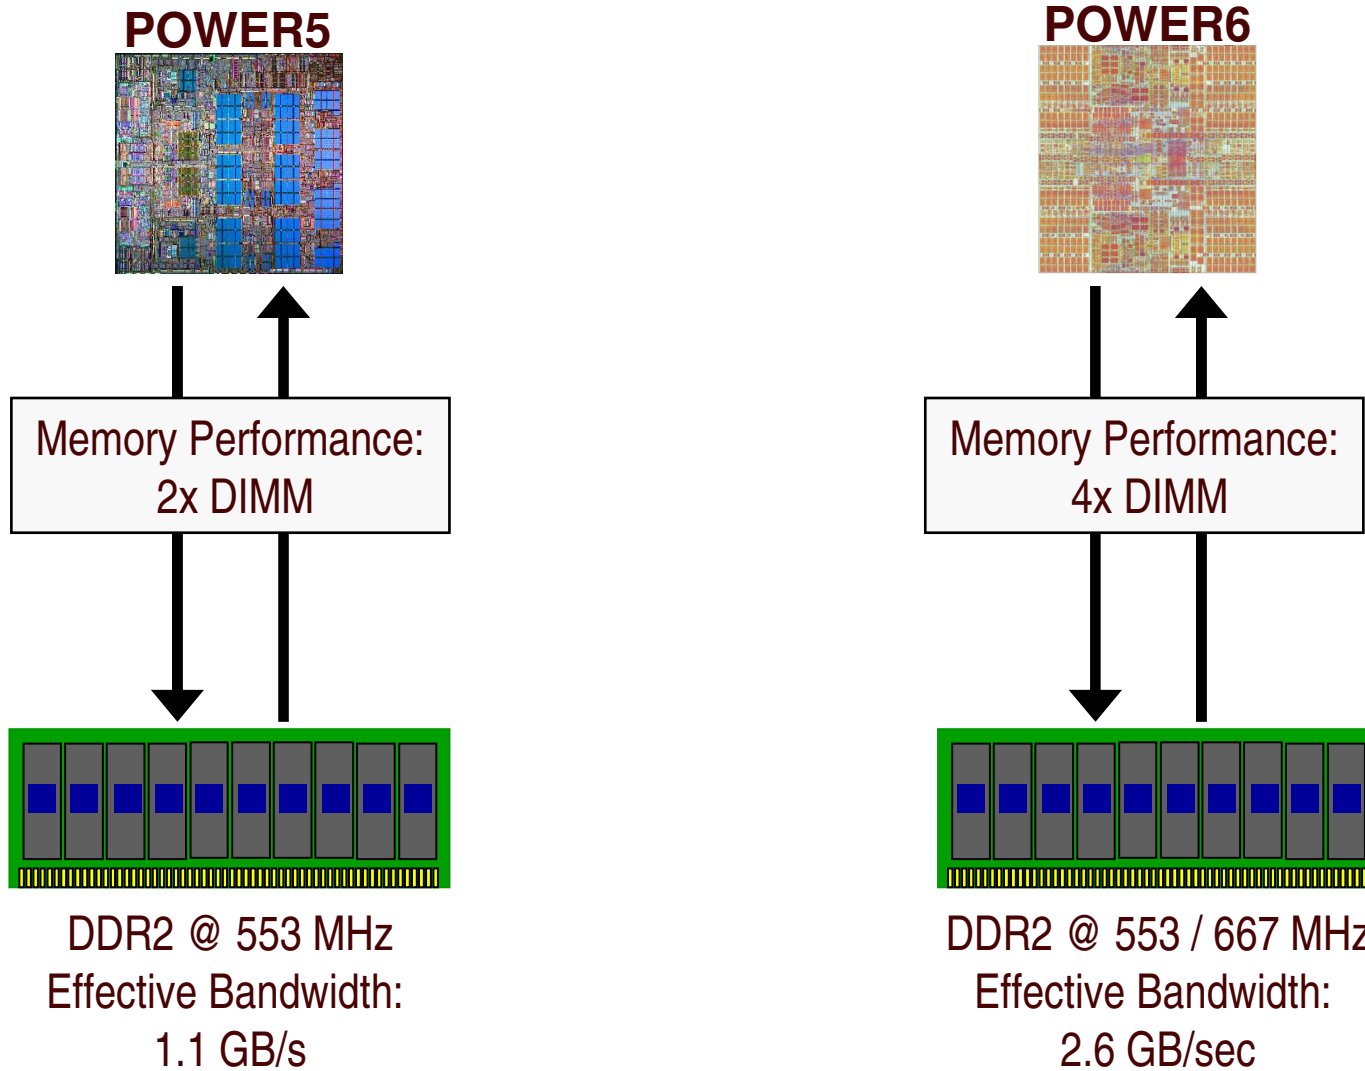
### Customer Awareness:

- ▶ Clients can choose partition environments
- ▶ Compatibility mode option available.

# POWER6 System Processor Options...



# Memory Channel Bandwidth Evolution






# Power Systems

# IBM Power Systems™

## POWER6

### Common Features

- Processor technology
- Service Processor
- Virtualization



560




570/16  
570/32




575




595



**550**



**520**



**BladeCenter**  
JS12 / JS22  
JS23 / JS43



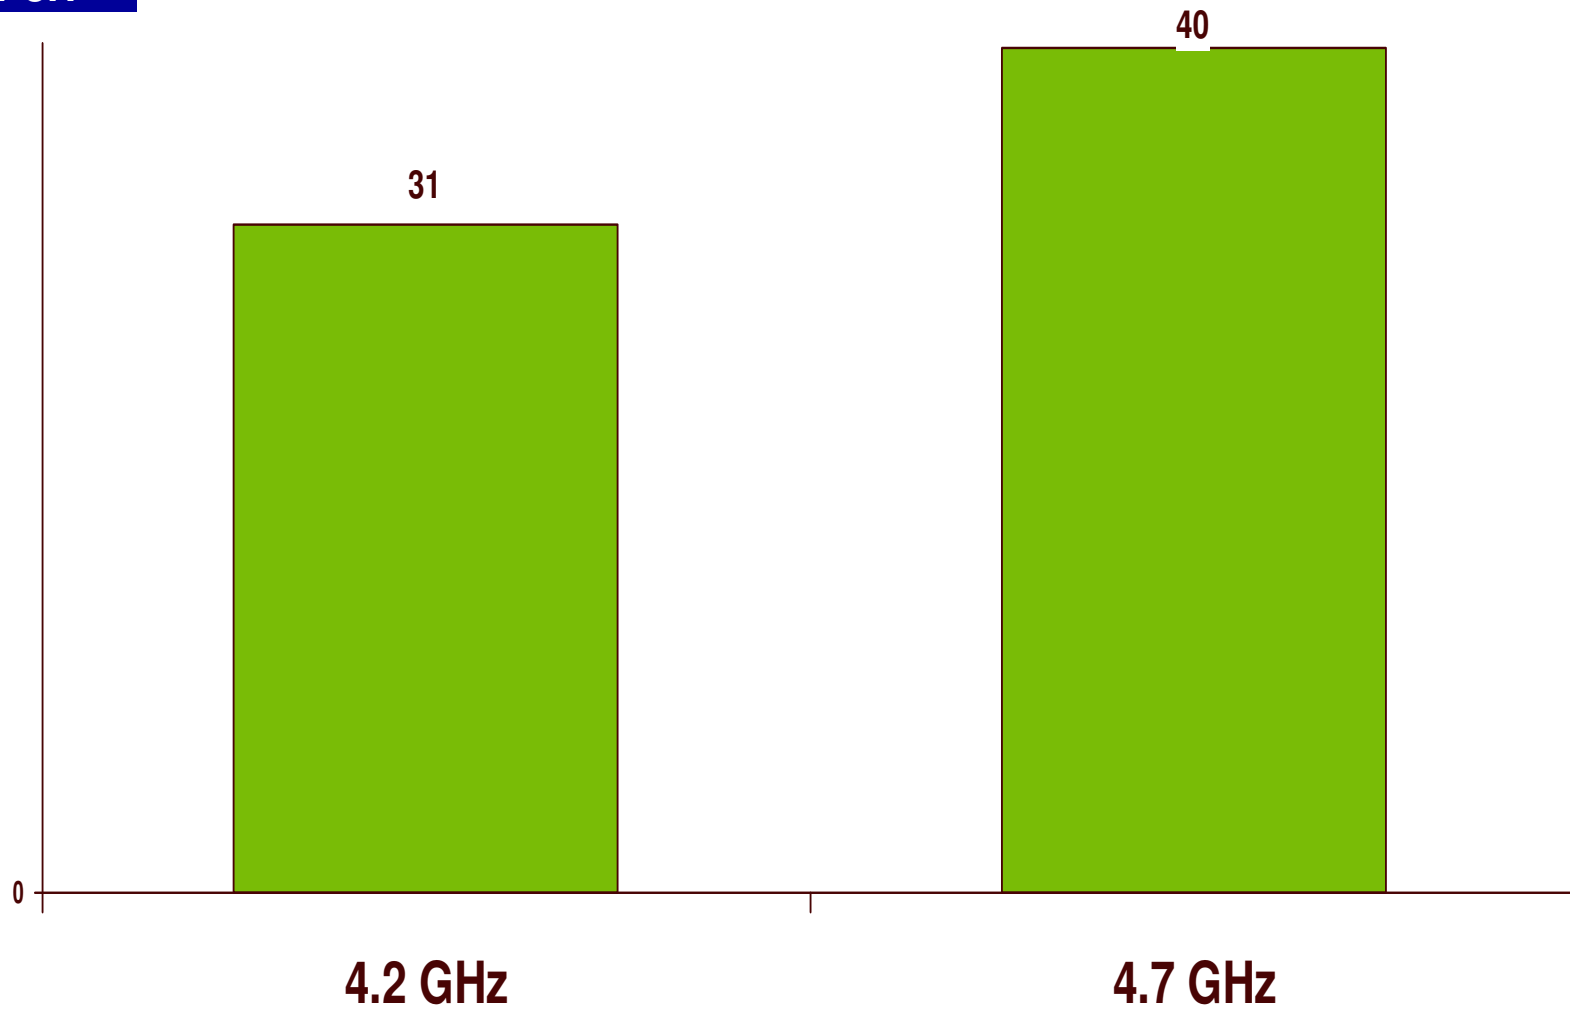
# Power 520 Rack & Tower




Power 520 Rack & Tower	
Architecture	POWER6 1 or 2 Cores 4.2 GHz POWER6+ 2 or 4 Cores 4.7 GHz L3 Cache: 32MB per chip with 4.7 GHz
DDR2 Memory	Up to 64GB (Buffered)
Internal SAS Disks	6 DASD (3.5") Optional: 8 SFF DASD (10 or 15K) Optional SSD support
Expansion	<ul style="list-style-type: none"> <li>▪ PCIe: 3 Slots</li> <li>▪ PCI-X 266: 2 Slots</li> <li>▪ GX Bus: 2 Slots</li> <li>▶ Shared with PCIe a slot</li> </ul>
Integrated SAS / SATA	Yes Optional: RAID support
Integrated Ports	3 USB, 2 Serial, 2 HMC Optional: SAS port
Integrated Virtual Ethernet	<ul style="list-style-type: none"> <li>▪ Dual Port 10/100/1000 Ethernet</li> <li>▪ Optional: Quad 1Gbt or Dual 10Gbt</li> </ul>
SATA Media Bays	1 Slim-line DVD 1 Half High Tape
Remote IO Drawers	Yes / Max: 8 GX Bus connection: RIO2 / IB / IB2
Dynamic LPAR	Up to 40 partitions
Redundant Power	Optional
Redundant Cooling	Yes

# POWER6 520 Performance


**rPerf**



# POWER6 520 4U Design



# Power 550 Rack & Tower




Power 550 Rack & Tower	
Architecture	2, 4, 6, or 8 Cores POWER6: 3.5 / 4.2 POWER6+: 5GHz L3 Cache: 32MB per chip
DDR2 Memory	Up to 256GB (Buffered)
Internal SAS Disks	6 DASD ( 3.5") Optional: 8 SFF DASD ( 10 or 15K ) Optional SSD support
Expansion	<ul style="list-style-type: none"> <li>▪ PCIe: 3 Slots</li> <li>▪ PCI-X 266: 2 Slots</li> <li>▪ GX Bus: 2 Slots                             <ul style="list-style-type: none"> <li>▶ Shared with PCIe 2 slots</li> </ul> </li> </ul>
Integrated SAS / SATA	Yes    Optional: RAID support
Integrated Ports	3 USB, 2 Serial, 2 HMC Optional: SAS port
Integrated Virtual Ethernet	<ul style="list-style-type: none"> <li>▪ Dual Port 10/100/1000 Ethernet</li> <li>▪ Optional: Quad 1Gbt or Dual 10Gbt</li> </ul>
Media Bays	1 Slim-line DVD 1 Half High Tape
Remote IO Drawers	Yes / Max: 8 GX Bus connection: RIO-2 / IB / IB2
Dynamic LPAR	Up to 80 partitions
NEBS / DC Power	Yes
Redundant Power	Optional
Redundant Cooling	Yes

# Power 550 Design

## I/O Adapters (5x)


- PCI-X, PCI-e
- 3 long adapters
- Hot Plug

## P6 processor (2x)



## Memory

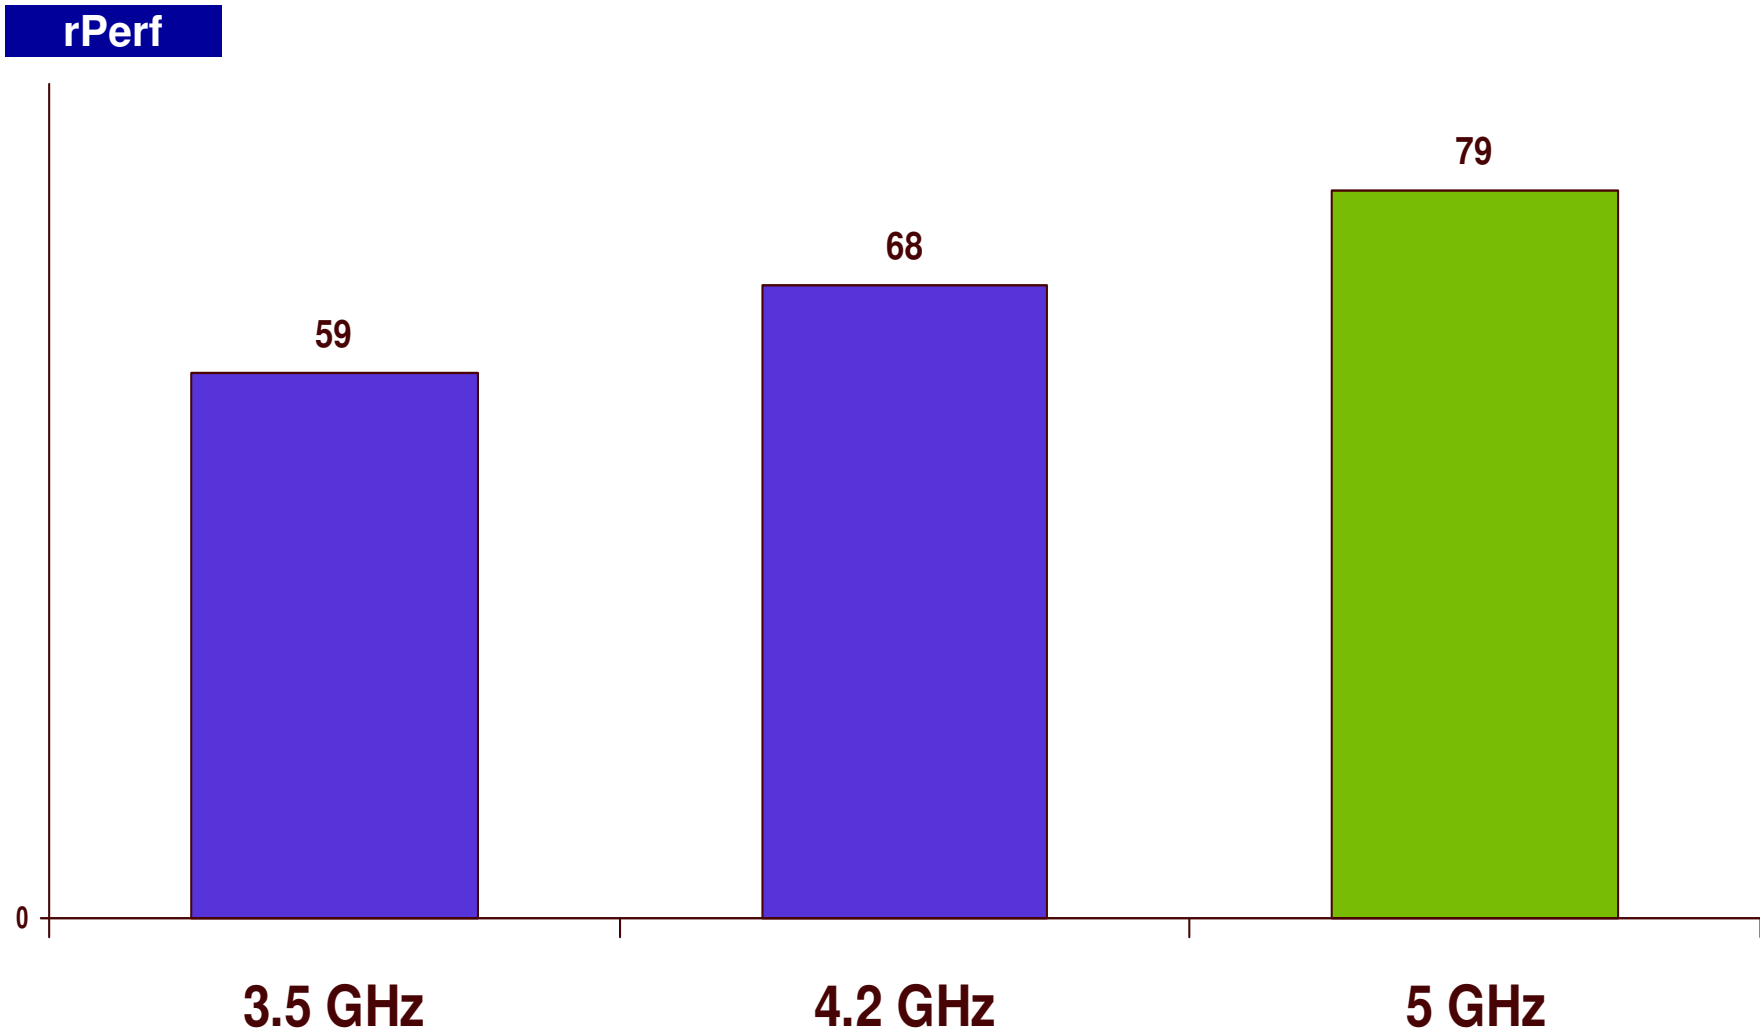
- Up to 256 GB



## SAS DASD

- 3.5 DASD Base (6x)
- 2.5 DASD Base (8x)


# Power 550 Performance






# Power 550 Leadership

### Performance SPECint\_rate2006



### Performance per Watt



Source: SPECint\_rate2006. For the latest SPEC benchmark results, visit <http://www.spec.org>

## Power Blade Announcement Overview...

### **IBM BladeCenter JS23 Blade**

- ▶ Single Wide 4 core blade with two 4.2GHz P6 processors with L3 Cache
- ▶ 8 memory DIMM slots supports up to 64GB of on board memory
- ▶ Supports Scalability to 8-cores and 16 memory DIMMs

### **IBM BladeCenter JS43 Blade**

- ▶ Double Wide 8 core blade with four 4.2GHz P6 processors with L3 Cache
- ▶ 16 memory DIMM slots supports up to 128GB of on board memory


### **IBM BladeCenter S SAS RAID Controller Module**

- ▶ Supports JS blades with AIX, i and Linux
- ▶ Integrates SAS switching and RAID controller function
- ▶ SAS boot and shared storage in IBM BladeCenter S chassis

### **Virtual Tape support on JS blades with IBM i**

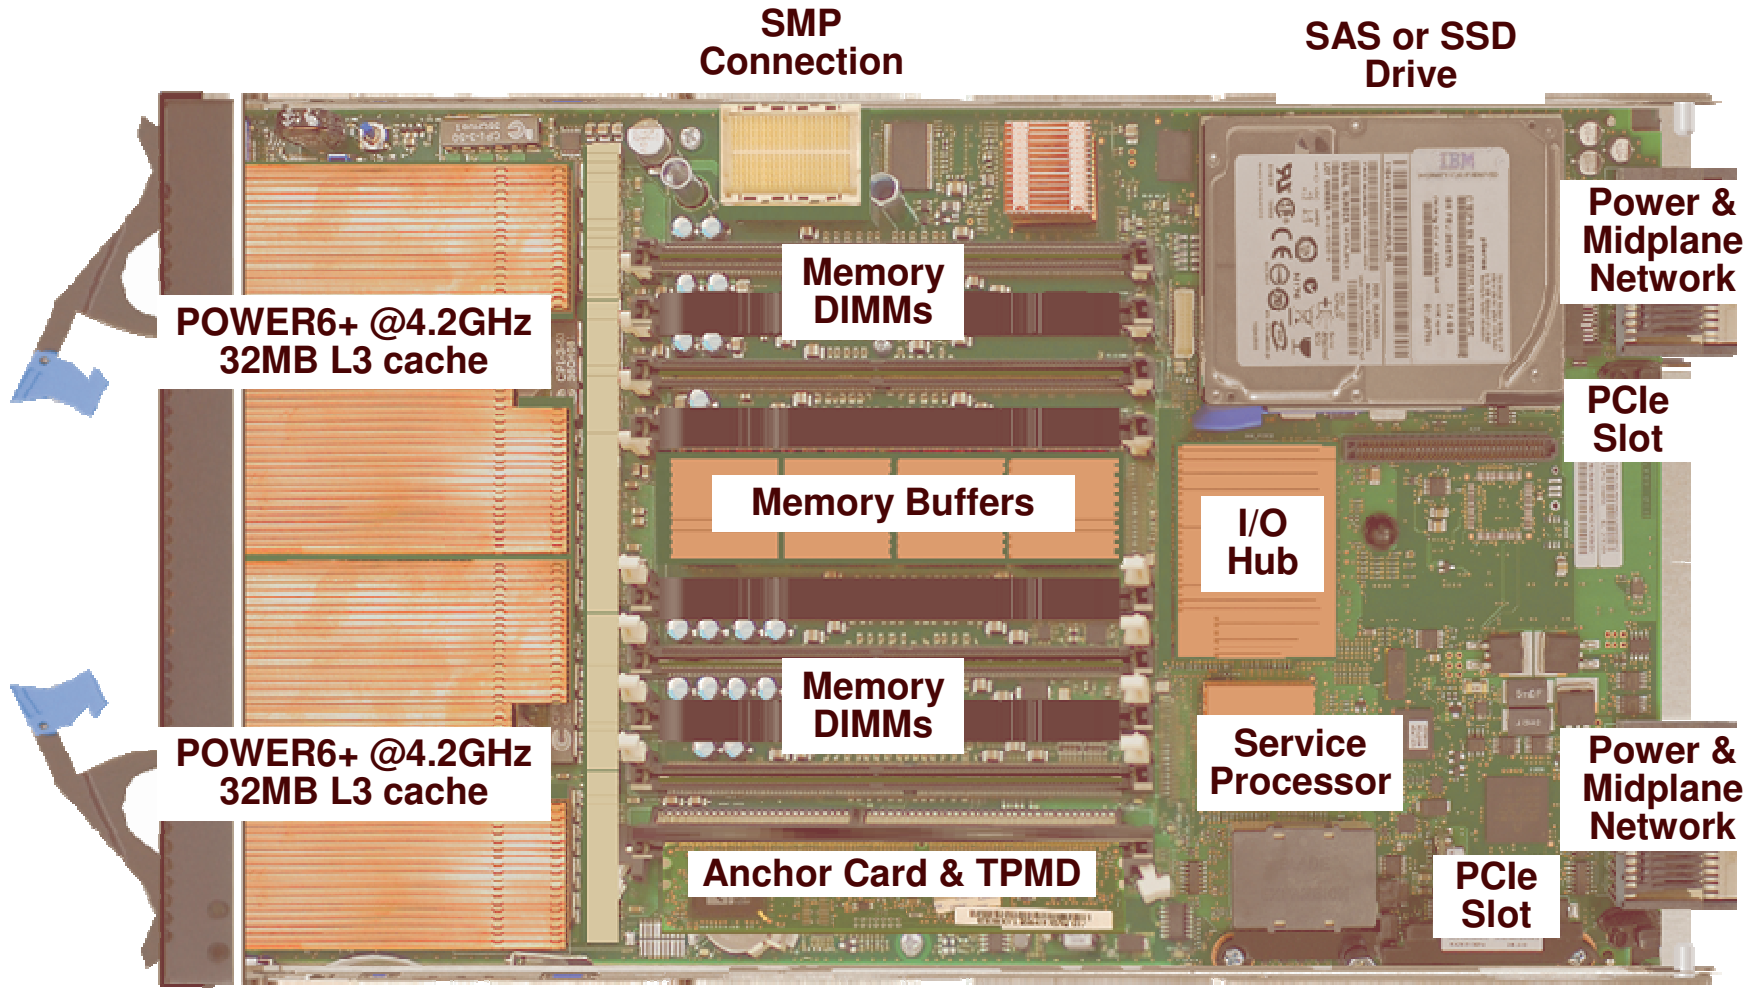
- ▶ Virtual tape support enables IBM i partitions to directly backup to PowerVM VIOS attached tape drive saving hardware costs and management time

# POWER6+ JS23 Blade




JS23 Blade	
Architecture	4 core / 2 Socket @ 4.2 GHz
L3 Cache	32MB per socket
DDR2 Memory	4GB to 64GB ( ChipKill )
DASD / Bays	SAS or Solid State: 0 – 1
Daughter Card Options	Two PCIe
Integrated Options	Dual Port 10/100/1000 Ethernet SAS Controller & USB
Fiber Support	Yes ( via Blade Center )
Media Bays	1 Blade Center
Redundant Power	Yes BladeCenter
Redundant Cooling	Yes BladeCenter
Service Processor	Yes
Virtualization	PowerVM Standard Integrated Virtualization Manager
Systems Management	IBM Director IBM EnergyScale™
OS Support	AIX 5.3 , 6.1, Linux and IBM i
BC Chassis	BC-H , BC-S, and BC-HT
Expansion	Yes

# JS23 Layout




# POWER6+ JS43 Blade



POWER6 JS43 Double Wide	
Architecture	8 core / 4 Socket @ 4.2 GHz Double wide XMP Interconnect
L3 Cache	32MB per socket
DDR2 Memory	8GB to 128GB ( ChipKill )
DASD / Bays	SAS or Solid State Disk: 0, 1, or 2
Daughter Card Options	Four PCIe
Integrated Ethernet	Dual Port 10/100/1000 Ethernet SAS Controller USB & KVM
Fiber Support	Yes ( via Blade center )
Media Bays	1 Blade Center
Redundant Power	Yes Blade Center
Redundant Cooling	Yes Blade Center
Service Processor	Yes
Virtualization	PowerVM Standard Integrated Virtualization Manager
OS Support	AIX 5.3 , 6.1, Linux, and IBM i

# JS43 Packaging



=



## JS23/43 Expansion Cards

### CIOv (Vertical) expansion cards:

- ▶ QLogic 8 Gb 2 Port Fibre Channel Expansion Card (CIOv)
- ▶ QLogic 4 Gb 2 Port Fibre Channel Expansion Card (CIOv)
- ▶ Emulex 8 Gb Fibre Channel Expansion Card (CIOv)
- ▶ SAS Connectivity Expansion Card (CIOv)

### CFFh (Horizontal) expansion cards:

- ▶ QLogic Ethernet and 4 GB Fibre Channel Expansion Card (CFFh)
- ▶ QLogic 8Gb Fibre Channel Expansion Card (CFFh)
- ▶ 4X InfiniBand DDR Expansion Card (CFFh)
- ▶ Voltaire 4X InfiniBand DDR Expansion Card (CFFh)

One combination form factor vertical (CIOv) may be installed with a combination form factor horizontal (CFFh) expansion card and one SAS HDD.

# IO Technology



## New IO Options

### 12X IO Drawers

- ▶ PCIe Support
- ▶ DDR Support

### SFF (Small Form Factor) DASD

- ▶ Power 520, Power 550, and 12X IO (PCIe) IO Drawers

### SSD (Solid State Drive) Technology

- ▶ Blades, Power 520, Power 550, Power 570, EXP 12S SAS Drawer

### New IO adapters

- ▶ PCI-X SAS Adapter
- ▶ PCIe SAS Adapter ( Power 560 / Power 570 Split Backplane Option )
- ▶ PCIe 10 Gb Ethernet adapters
  - SR optic (#5769)
  - CX4 twinax copper (#5732) cabling
- ▶ PCIe 4-Port Async EIA-232 Adapter

## Power 520/550 SAS Disk Bays/Slots



Must choose 3.5-inch or SFF bays  
 ▶ Even if empty of drives

Power 550 announced SFF in 2008, but did not GA until 2009.

Power 520 announced April 2009.

Power 520/550 SFF capability GA together.

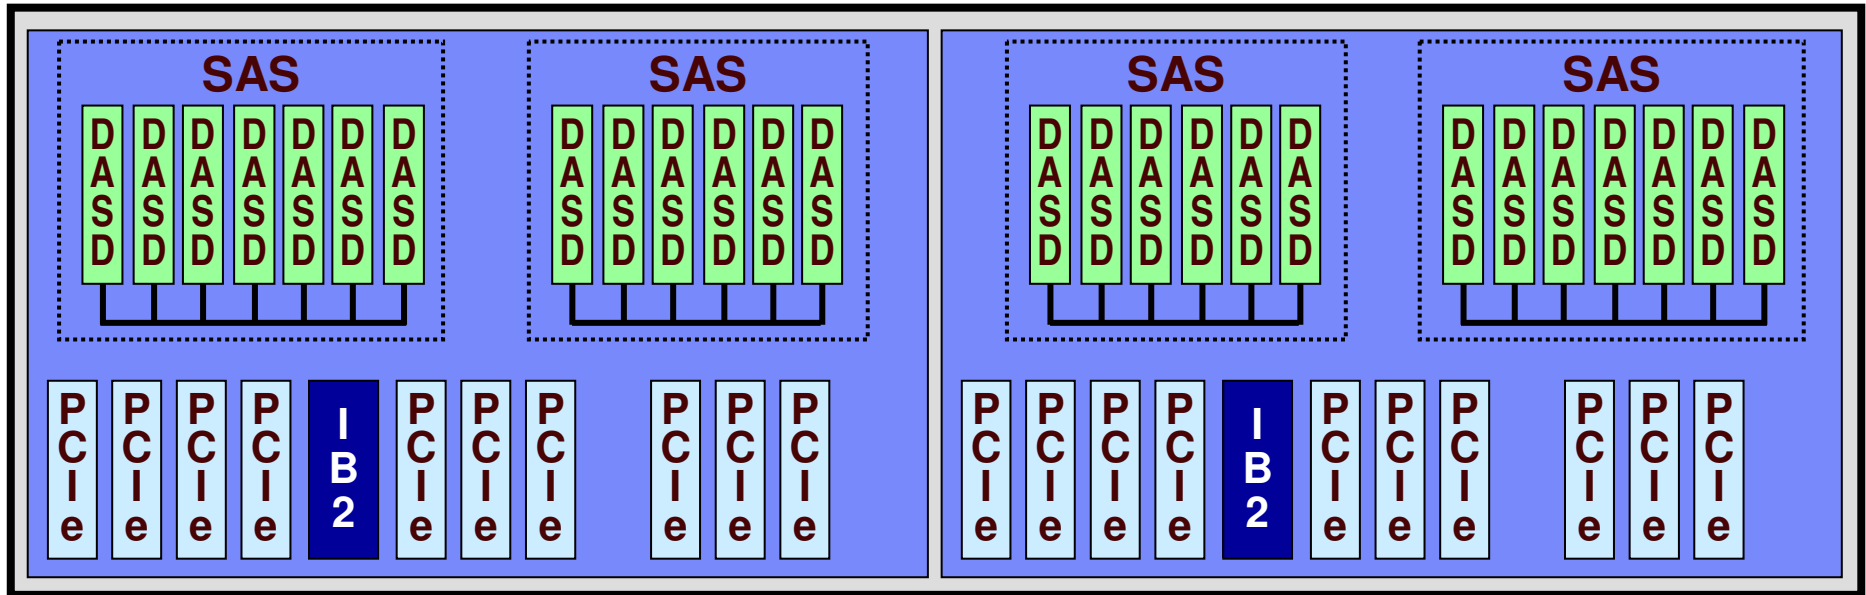
Six 3.5-inch SAS disk bays

**Eight SFF SAS disk bays**



<b>Considerations</b>	3.5-inch	SFF (2.5-inch)
Max number HDD in CEC	6	8
Split backplane option (AIX/Linux)	Y	Y
Max 10k rpm HDD capacity	N/A	146 GB
Max 15k rpm HDD capacity	428 GB 450 GB	69 GB 73 GB
SSD option	N	Y

# PCIe 24" T24 Remote IO Drawer FC: #5803



**Dimensions:** 4U x 24"

**PCIe Slots:** 20


- ▶ 2.5 GB/sec ( x8 )

**SAS Drives:** 26 ( SFF Drives )

- ▶ Adapter/s required

**Interconnect:** InfiniBand 2

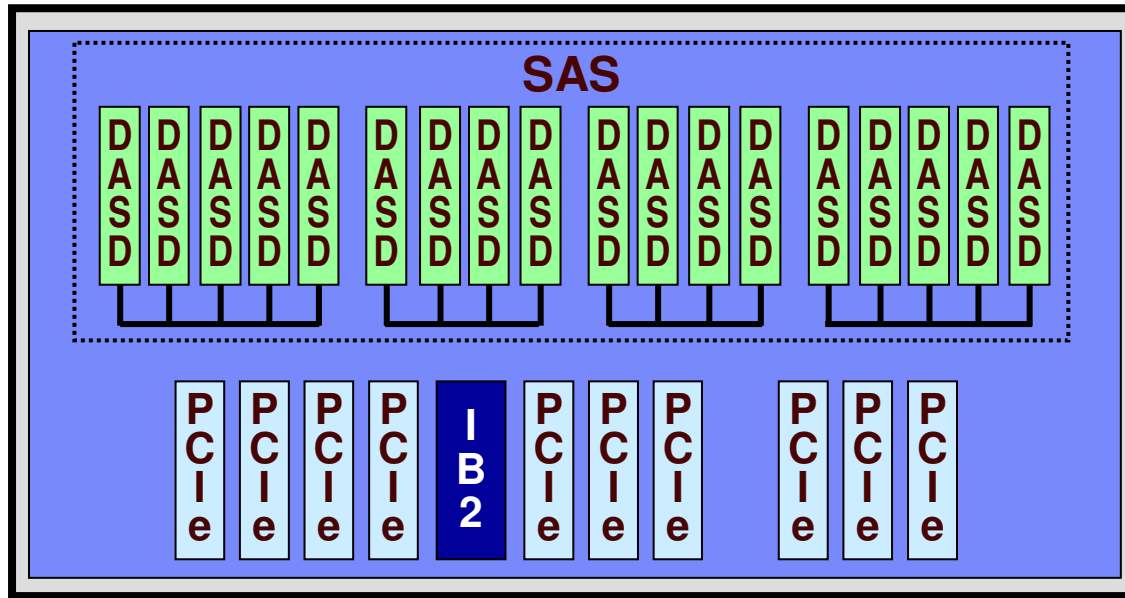
- ▶ Dual ports per controller
- ▶ 9.1 GB/sec per port ( Sustained )



**Drawer Partition Modes:**

- ▶ Single Partition: 26 Drives
- ▶ Dual Partition: 13 drives per partition
- ▶ Quad Partition: 7 / 6 / 6 / 7 Drives

# PCIe 19" T19 Remote IO Drawer FC: #5802



**Dimensions: 4U x 19"**

**PCIe Slots: 10**


- ▶ 2.5 GB/sec ( x8 )

**SAS Drives: 18 ( SFF Drives )**

- ▶ Adapter/s required

**Interconnect: InfiniBand 2**


- ▶ Dual ports per controller
- ▶ 12 GB/sec per port ( Peak )
- ▶ 9.51 GB/sec per port ( Sustained )



**Drawer Partition Modes:**


- ▶ Single Partition: 18 Drives
- ▶ Dual Partition: 9 drives per partition
- ▶ Quad Partition: 5 / 4 / 4 / 5 Drives

# IO Link Bandwidth (Simplex)




# Front and Back views 19 inch Drawer....


**Front View**



**Rear View**




## SAS Controllers



**FC: #5904 / 5906 / 5908**

**Double wide PCI-X SAS RAID adapter  
Replacement for SCSI 1.5GB write  
cache controller**

**Write cache protection included  
1.6GB read cache (like SCSI adapter)  
AIX & IBM i Support ( Linux TBD )  
POWER6 520, 550, 560, 570, 595  
POWER5 520, 525, 550, 570, 595  
Max of 60 drives per controller  
Max of 8 SSD per controller**



**FC: #5903**

**PCIe x8, Dual-x4, 3Gb SAS RAID / Cache**

**Features:**

- ▶ **Dual port adapter**
- ▶ **Dual shared Active / Active**
- ▶ **POWER 570 Split Backplane support**

**AIX 5.3, AIX 6.1; SLES10.2/11, RHEL5.2/6**

**IBM i not supported ( Later Date )**

## Networking Options

### **New Ethernet Adapters...**

FC #5769 Short Range (SR) adapter uses multi-mode (850nm) optical fiber for up to 300 meters.

FC #5732 CX4 uses twinax copper for up to 15 meters.

- ▶ Different from the AS/400 heritage twinax cabling

Feature #5769 & #5732 CX4 adapters provide Linux with additional function

- ▶ iSCSI hardware initiator support
- ▶ RDMA (Remote Direct Memory Access)

### **Current Adapter...**

FC #5773 Long Range (LR) adapter uses single-mode (1310 nm) optical fiber for up to 10 km.



# SSD Technology

## HDD History ( Seagate )

<b>Product Availability</b>	<b>2002</b>	<b>2008</b>	<b>% Improvement</b>
<b>Generation</b>	15K.3	15K.6	
<b>Capacity ( GB)</b>	73	450	<b>35%</b>
<b>Max Sustained DR (MB/sec)</b>	75	171	<b>15%</b>
<b>Read seek Time (ms) Random IOP</b>	3.6	3.4	<b>1%</b>

## Performance Comparisons

<b>Data Access</b>	<b>Computer Terms</b>
<b>DRAM</b>	<b>80 nsec</b>
<b>SSD</b>	
<b>DASD</b>	<b>5 msec</b>
<b>Tape</b>	<b>20 sec</b>

## SSD Offering

### Enterprise Grade Solid State Drive (SSD)

- ▶ Built in wear leveling

Capacity: 69 GB

### Supported Environments:

- ▶ Power 520 / 550
- ▶ Power 560 / 570
- ▶ Power 595 / 575
- ▶ EXP 12S Storage Drawer

### SAS Interface ( 3 Gb )

- ▶ 2.5 / 3.5 inch Inserts
- ▶ Dual Options: AIX & Linux and IBM i

### Performance Throughput Sustained:

- ▶ 220MB/s Read
- ▶ 122MB/s Write


### Random transactional operations (IOPS)

- ▶ 28,000 IOPS

### Average Access time:

- ▶ 20 – 120 microseconds

Power Consumption: ~ 1/2 15K HDD




**2.5 inch ( SFF ) Insert**

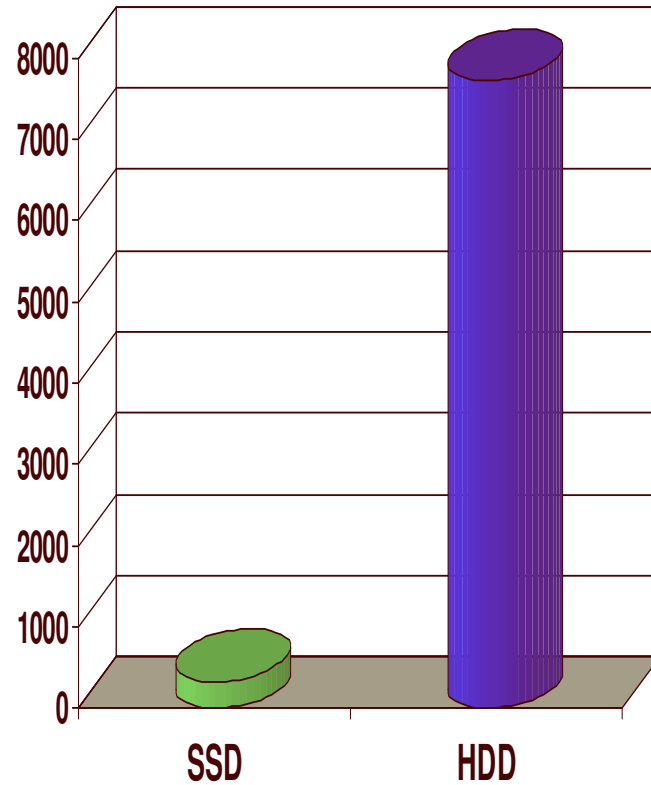


**3.5 inch Insert**

# SSD Performance




Random I/O's per second (Sustained)



Power Consumption in Watts Required for 135K IOPS performance



## Mixing SSD + HDD

It is typical for data bases to have a large percentage of data which is infrequently used (“Cold”) and a small percentage of data which is frequently used (“Hot”)





Hot data may be only 10-20% capacity, but represent 80-90% activity

SSD offers best price performance when focused on “Hot” data  
 HDD offers best storage cost, so focus it on “Cold” data sort of a hierarchal approach



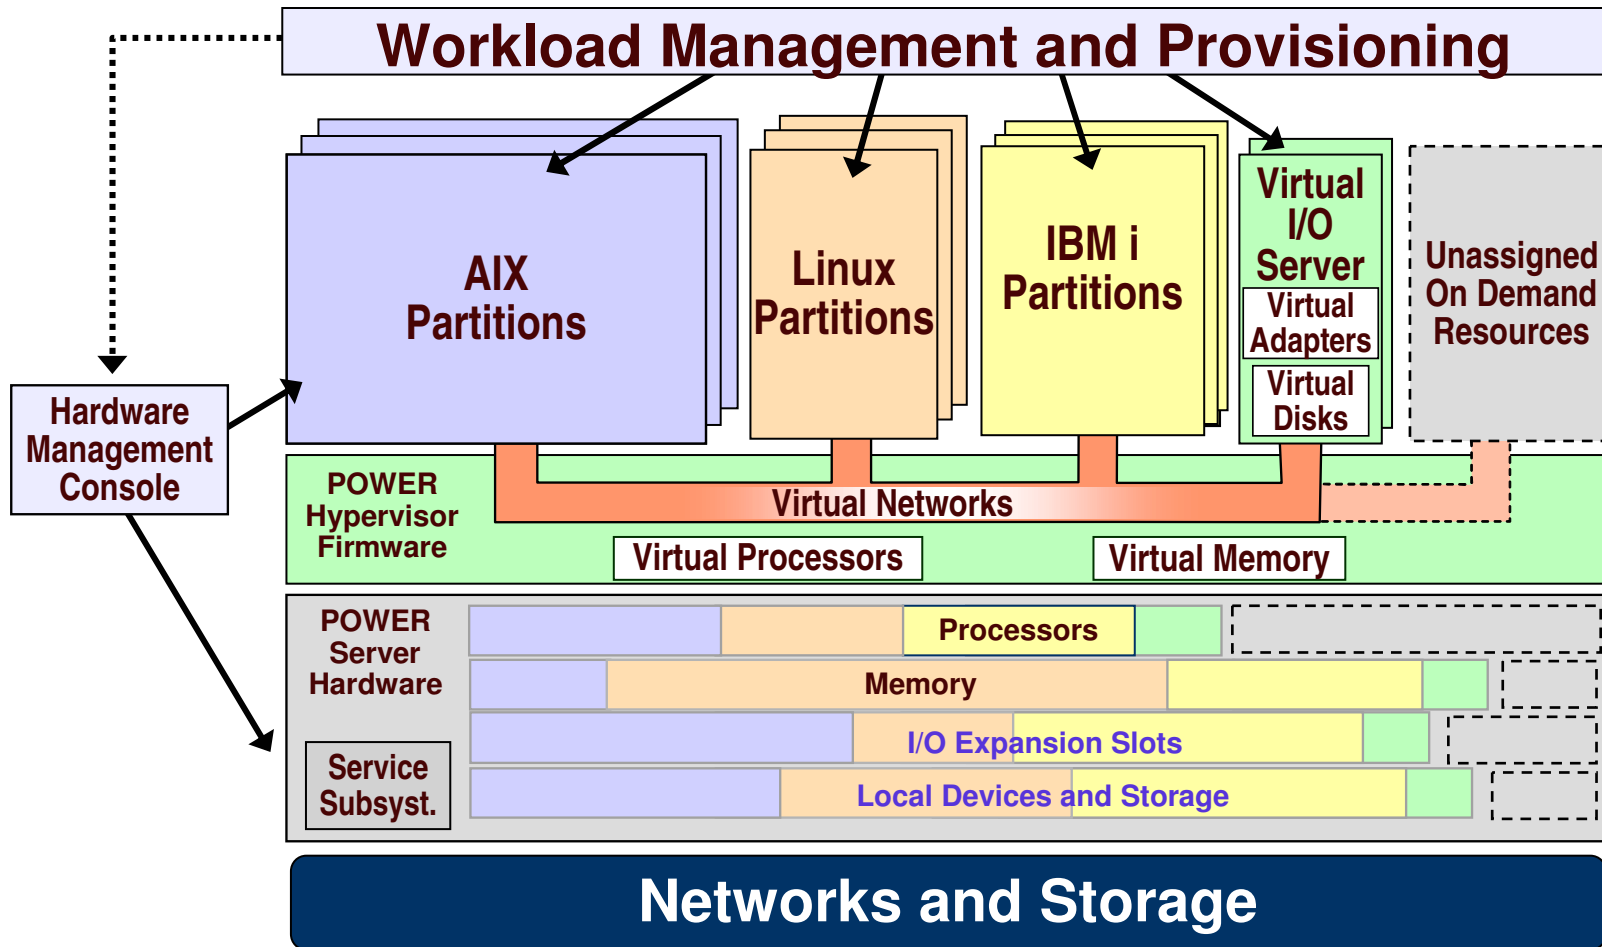
Can run SSD closer to 100% capacity



May be able to use larger HDD and/or a larger % capacity used

# Virtualization

# PowerVM Virtualization Architecture



Virtualization of P5 & P6 servers is accomplished using two layers of firmware:

- A thin core hypervisor that virtualizes processors, memory, and local networks
- One or more Virtual I/O Server partitions that virtualize I/O adapters and devices



## Customer Acceptance..

### Significant leap in the uptake of virtualization in 2008

- ▶ PowerVM revenue grew 73% ( Compared to 2007 )
- ▶ 66+% on POWER6-based systems in 2008

### High End Systems:

- ▶ Adoption rate on the high end Power 595 server: 93%

### Entry Systems:

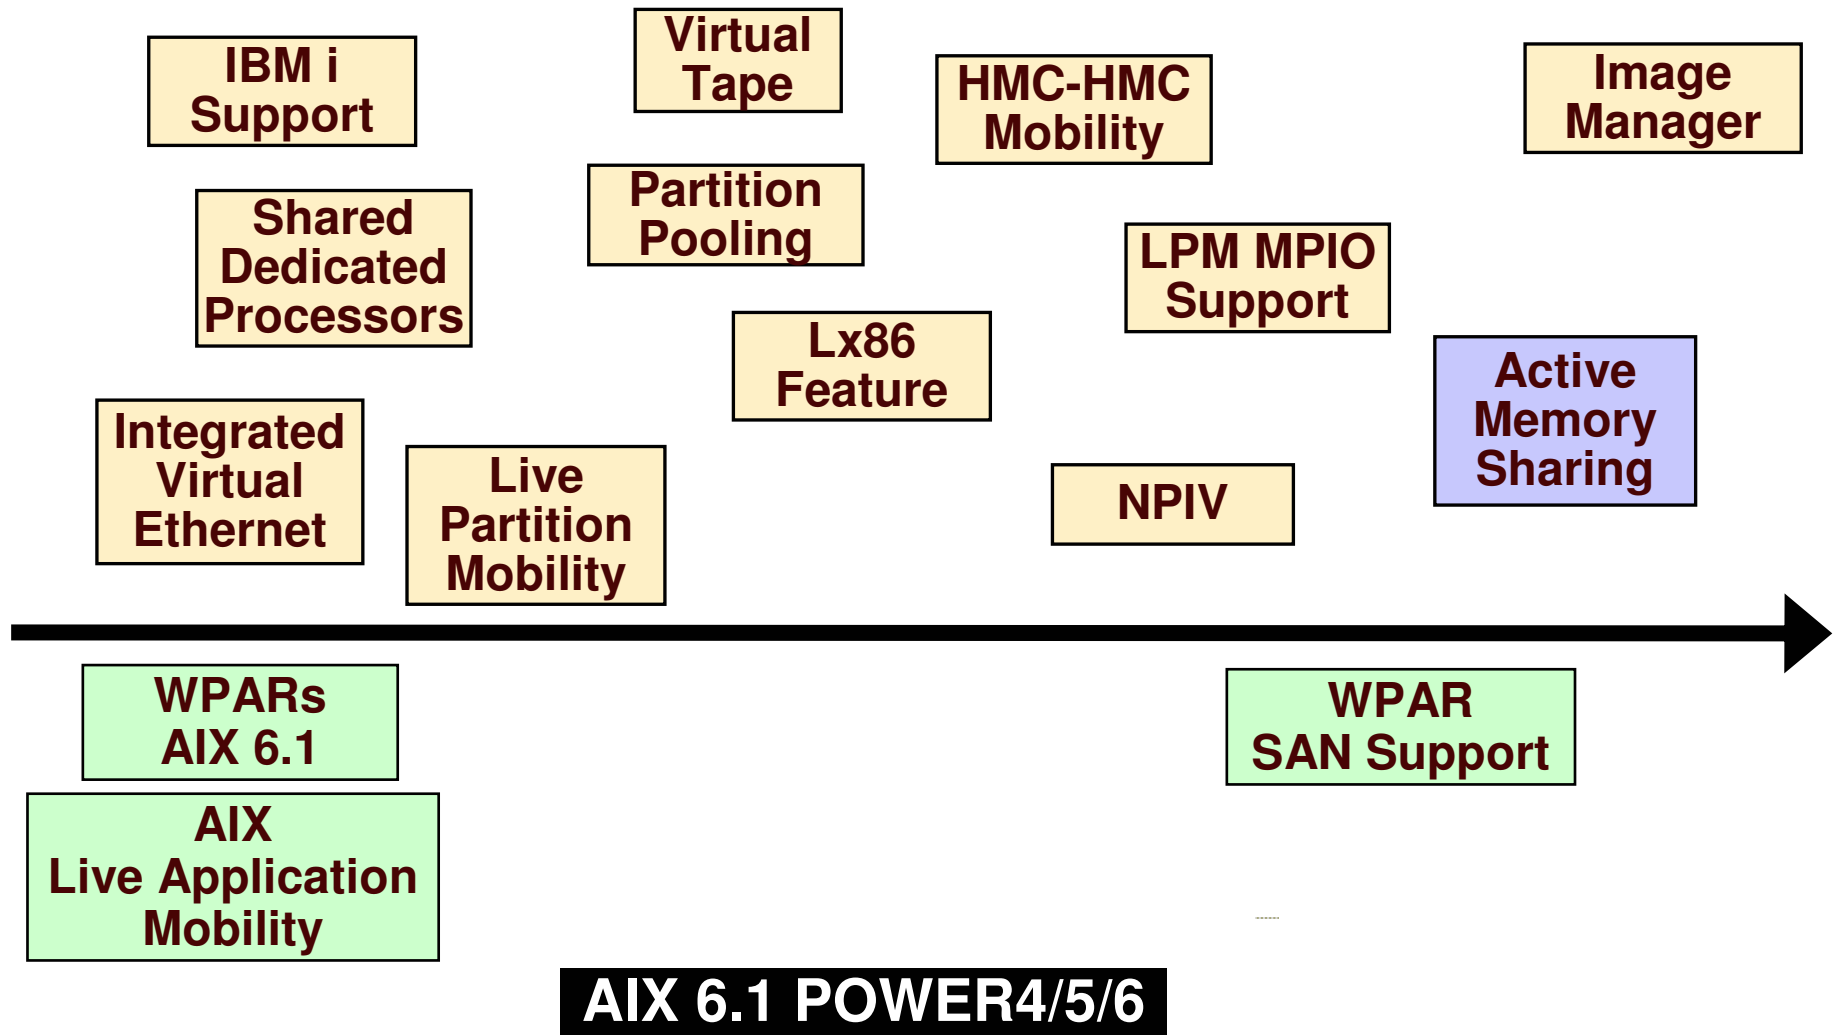
- ▶ Adoption rate on the low-end Power 520 server: 53 %
- ▶ 11% in 2007

### Reasons:

- ▶ Server consolidation, Provisioning ability, Floor space reduction, Infrastructure flexibility, Reduced costs

# PowerVM Rollout

## POWER6 AIX 5.3 / 6.1 / Linux



# Partition Pooling: Multiple Shared Processor Pools:

**POWER6 16 Core System**

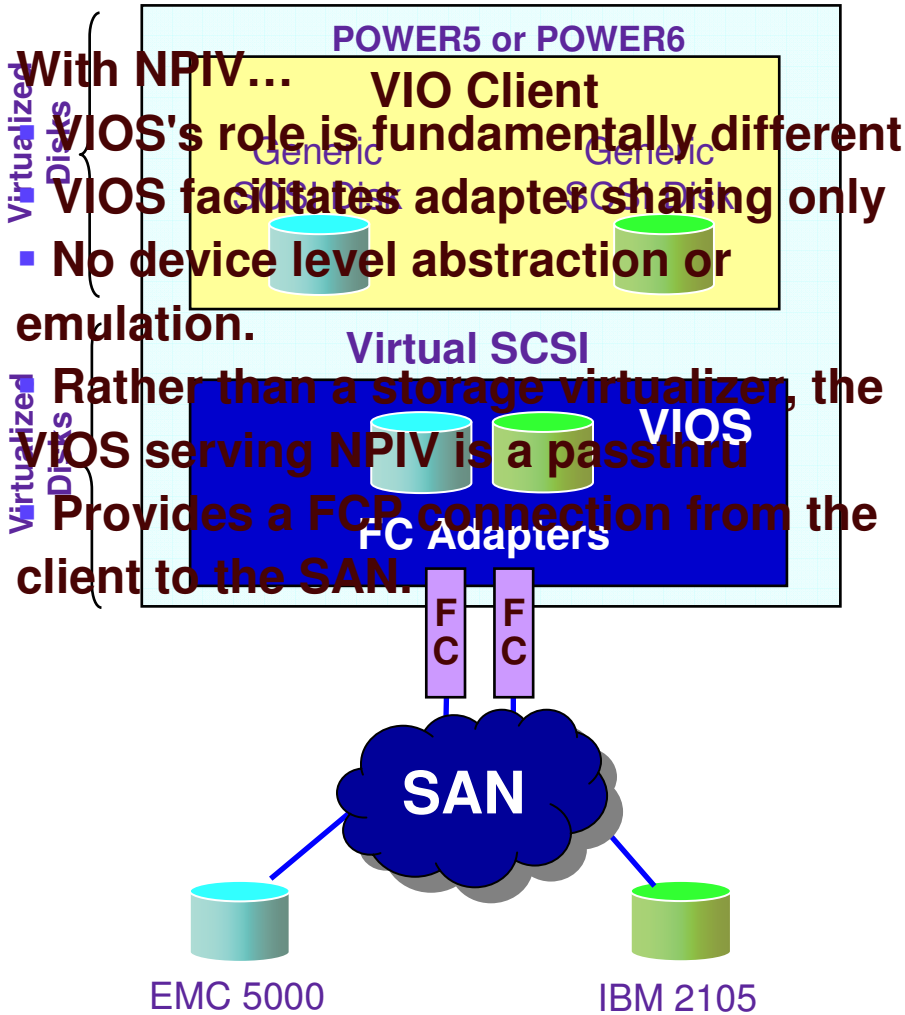
P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15
AIX	AIX	AIX	Linux	AIX	AIX	AIX	AIX	Linux	IBM i	IBM i	Linux	AIX	AIX	Linux
		2	1	4	4	1	2	3	2	2		7	1	1
		0.75	0.25	1.5	0.5	0.25	0.5	0.5	0.25	0.25	0.5	0.5	0.25	0.25
		V Pool: 0 Max Cap: 2 Ent Cap: 1		V Pool: 1 Max Cap: 10 Ent Cap: 3.25				V Pool: 2 Max Cap: 3 Ent Cap: .5		V Pool: 3 Max Cap: 4 Ent Cap: 2				
2 Core	1 Core	13 res ( Physical Shared Processor Pool )												

Dedicated

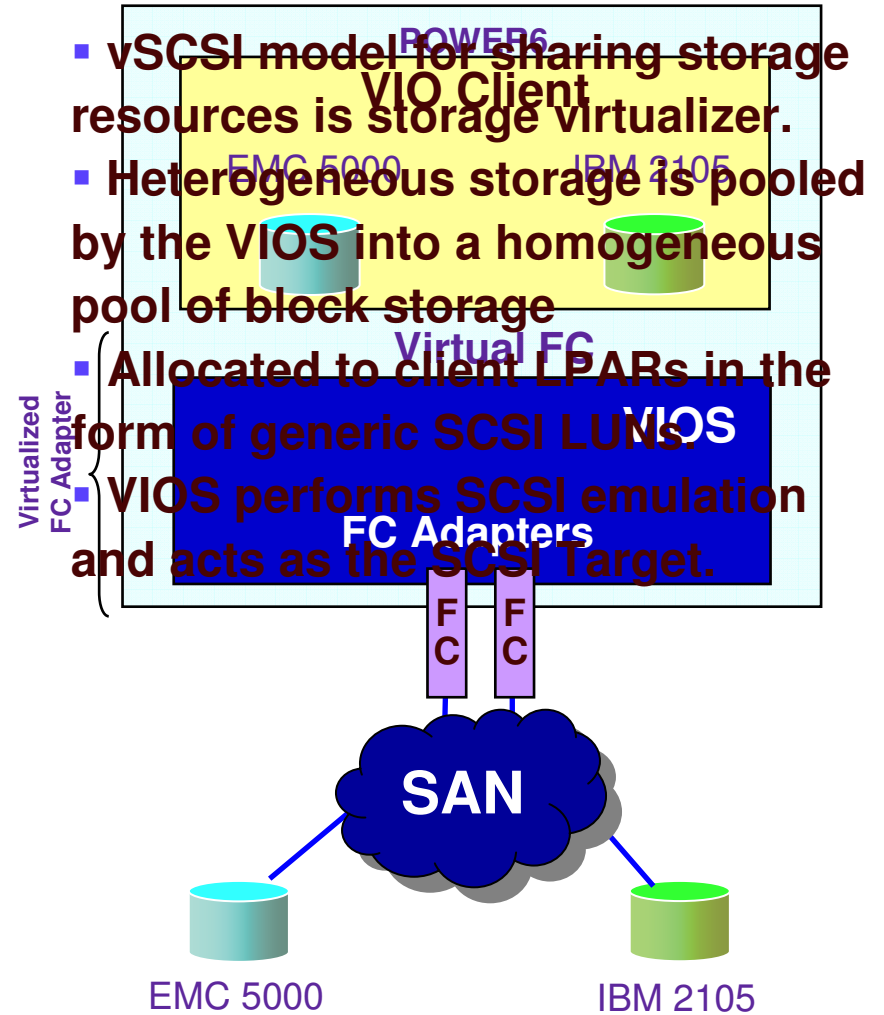
- ▼ Capped Partition
- # Number of VP's
- # Entitled Capacity

**POWER6 technology**  
 Ability to cap uncapped partitions  
 Helps manage software costs  
 Segment production / development / test / etc.  
 Mobility of partitions is supported  
 Maximum: 64 Pools

## vSCSI model



## N-Port ID Virtualization



## Active Memory Sharing Overview

***Active Memory Sharing intelligently flows memory from one partition to another for increased utilization and flexibility of memory usage***

### **Memory virtualization enhancement for Power Systems**

- ▶ Memory dynamically allocated based on partition's workload demands
- ▶ Contents of memory written to a paging device
- ▶ Improves memory utilization

### **Designed for partitions with “Variable Memory” requirements**

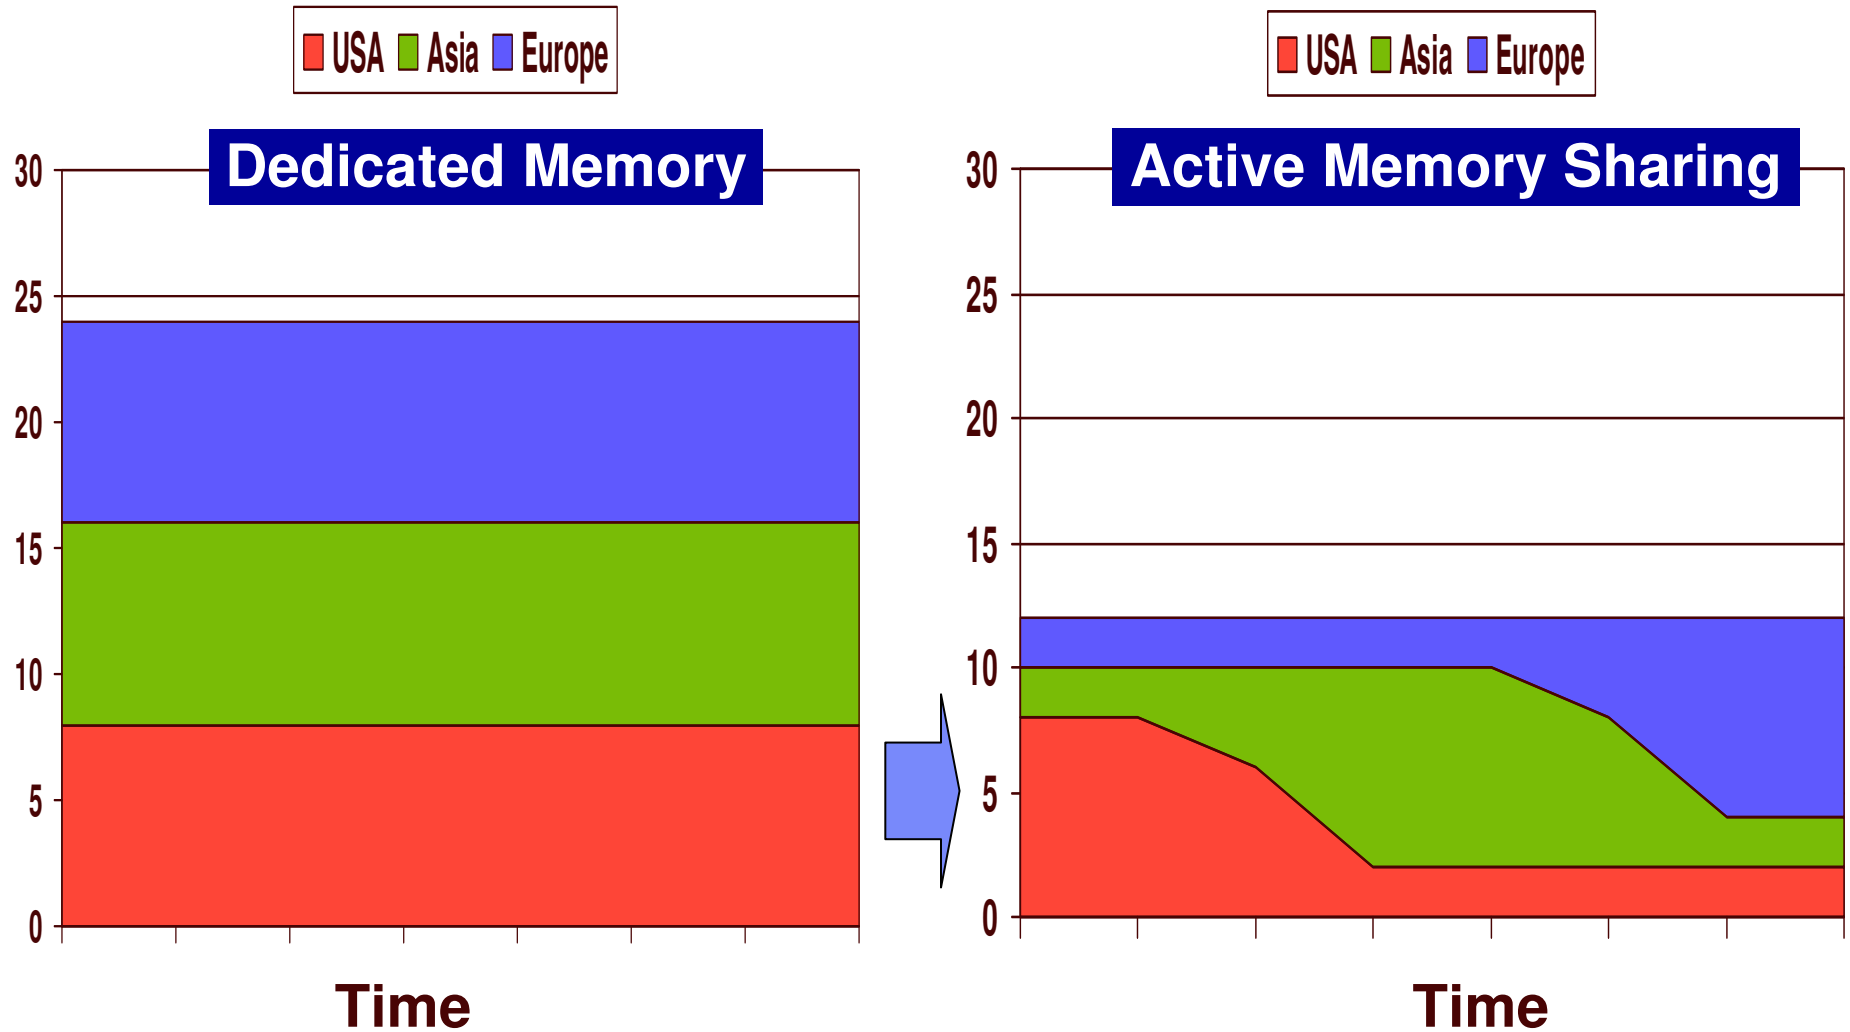
- ▶ Low average memory requirements
- ▶ Active / Inactive environments
- ▶ Workloads that peak at different times across partitions

### **Available with PowerVM Enterprise Edition**

- ▶ AIX 6.1, Linux, and IBM i 6.1 partitions that use VIOS and shared processors
- ▶ POWER6 processor-based systems
- ▶ Must use Shared Processor and have Virtual IO ( VIOS managed )

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

# Dedicated vs Active Memory Sharing Environment

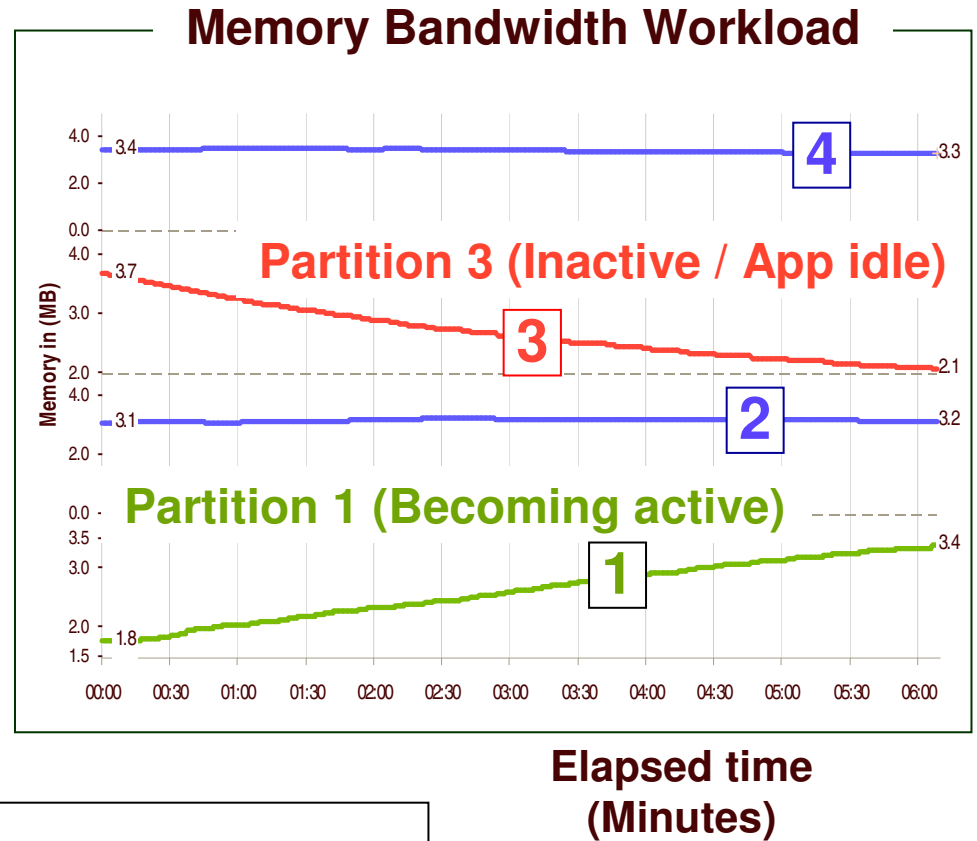


## An example of Activity Memory Sharing

Partition 1 started, memory removed from partition 3 until performance full speed

Partition 3 workload idle but, memory not released by application

Partition 2 and 4 workload performance protected



**Performance dependencies:**

- Level of activity in shared memory partitions
- Application memory behavior
- Amount of memory to be moved
- Paging device performance

# AIX



## AIX 6 Features....

POWER6 Exploitation

Software Reliability Availability Serviceability

**Enhancements to existing Virtualization Technologies**

**Workload Partitions** ( Software based Virtualization )

**Application Mobility** ( Cross system Workload Mobility )

64 bit Kernel only

Integrated Multilevel Security

**Role Base Access Control** ( Partial Root base )

Encrypted File system

CAPP EAL4+ and LSPP Security Certification

Solution Performance Tuning


**AIX Kernel Hot-Patching**

Dynamic Tracing for AIX

Ease of Use

- ▶ Portal base SMT, LPAR Simplification

# AIX Release Plan



	-Marketed & serviced		-Technology Level Update		-Fee-based service extension
	-Serviced only				-Web support only

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



# AIX / PowerVM updates for May 2009

## AIX 6 Technology Level 3

- **Workload Partition SAN support**

- Workload Partitions will be able to own SAN devices which will allow WPAR administrators to directly manage their own storage. This will provide the opportunity to reduce administrative effort and increase flexibility.

- **Probevue support for C++**

- Extends the code development and debugging capabilities of the probevue dynamic tracing facility to applications written in C++

- **IBM Systems Director Agent included on AIX Base installation media**


- Installed by default starting with AIX 6 TL3 and AIX 5.3 TL10

- **Active Memory Sharing support**

- New capability of PowerVM Enterprise



# 2008 – 2009 AIX Release Roadmap

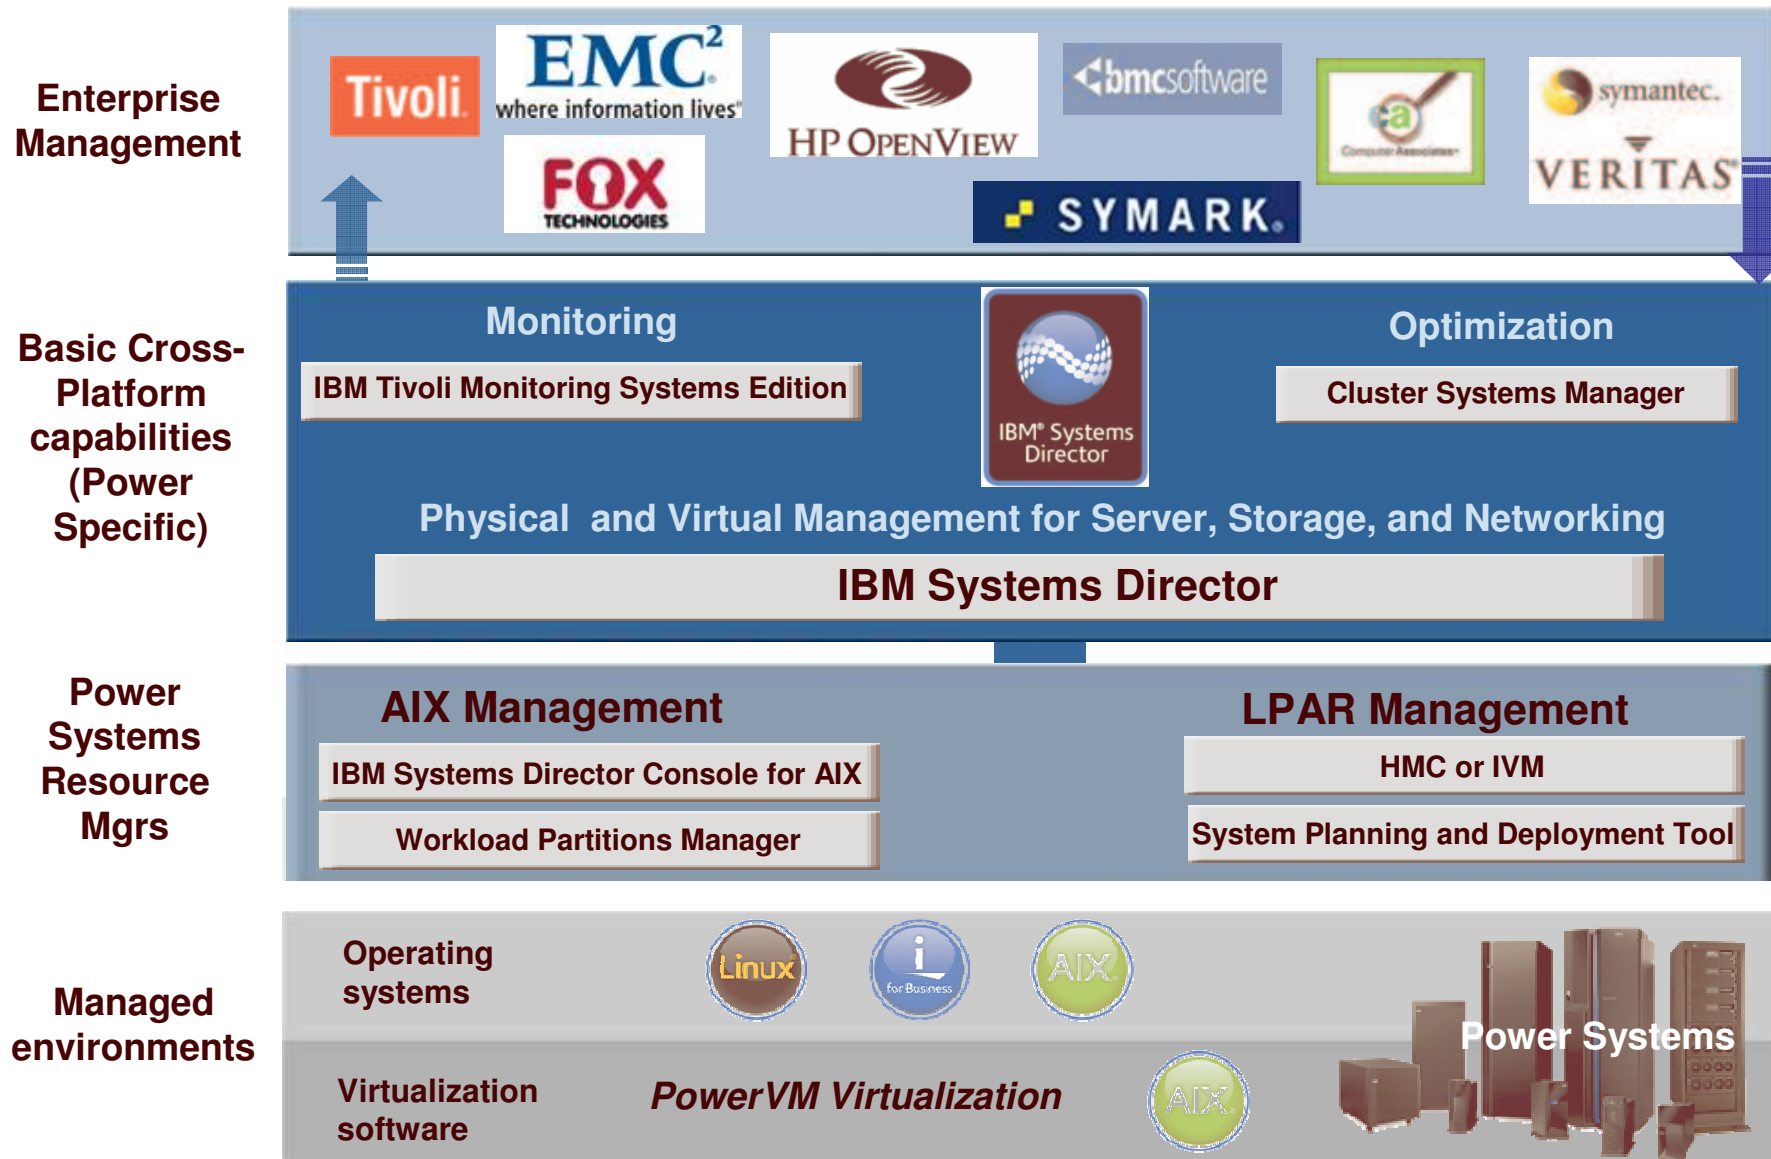


HW Support Release  
Via Service Packs

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

# IBM Director

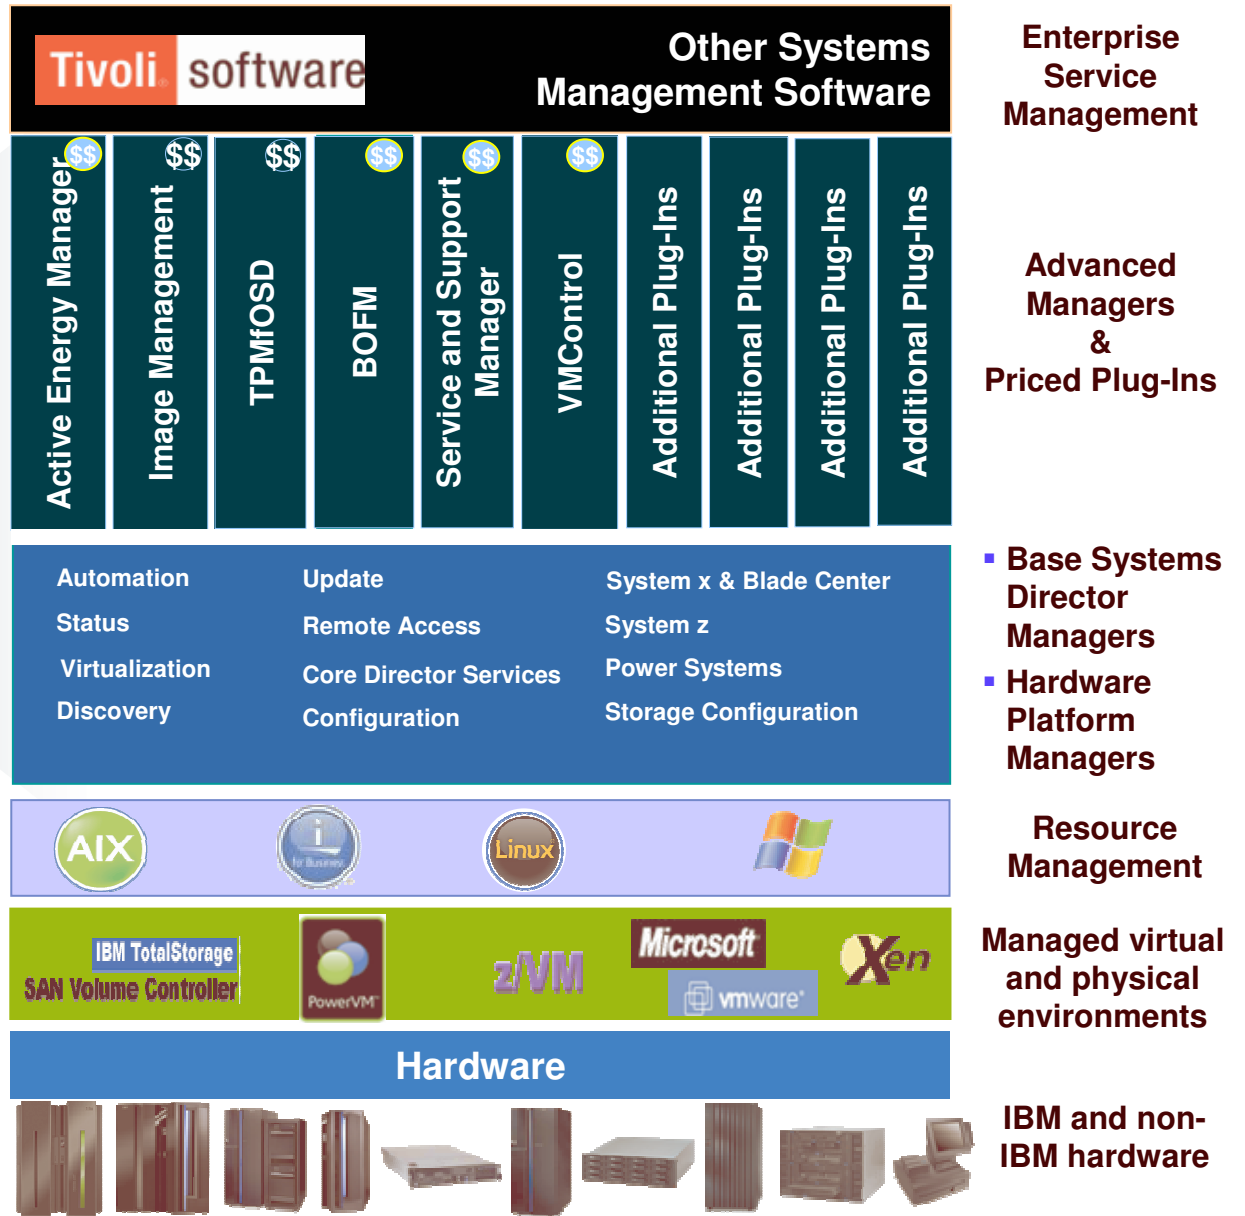
# End-to-end Management Approach for Power Systems...



# Director 6.1



**End-to-end Management**



# Linux



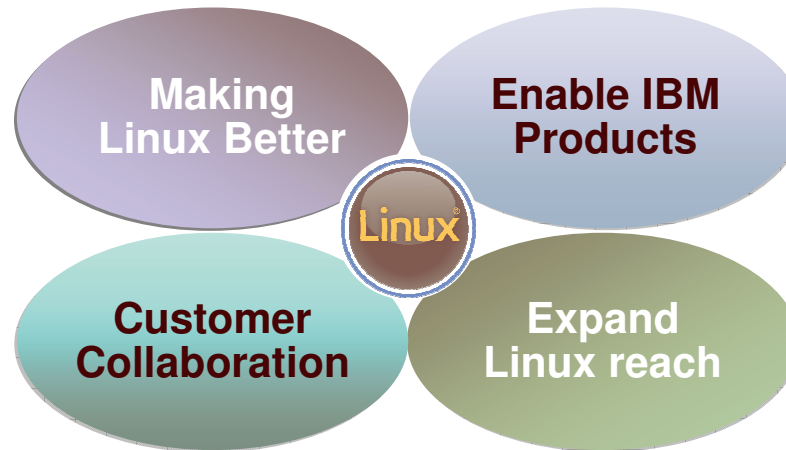
# Enhancing Linux Capabilities & Driving Linux Adoption

## IBM Contributes to the Community

- Over 600 IBM developers contributing to 100+ Linux and Open Source projects
- Develop closely with Red Hat and Novell
- Developers sharing technical knowledge on <http://planet-ltc.org>

## IBM Supports Linux as a Tier 1 OS

- All IBM Systems, SW, and Middleware run on and are certified for Linux
- Driving performance parity with IBM's own operating systems
- Making contributions in security, RAS, scalability, performance, management



## IBM Collaborates with Customers

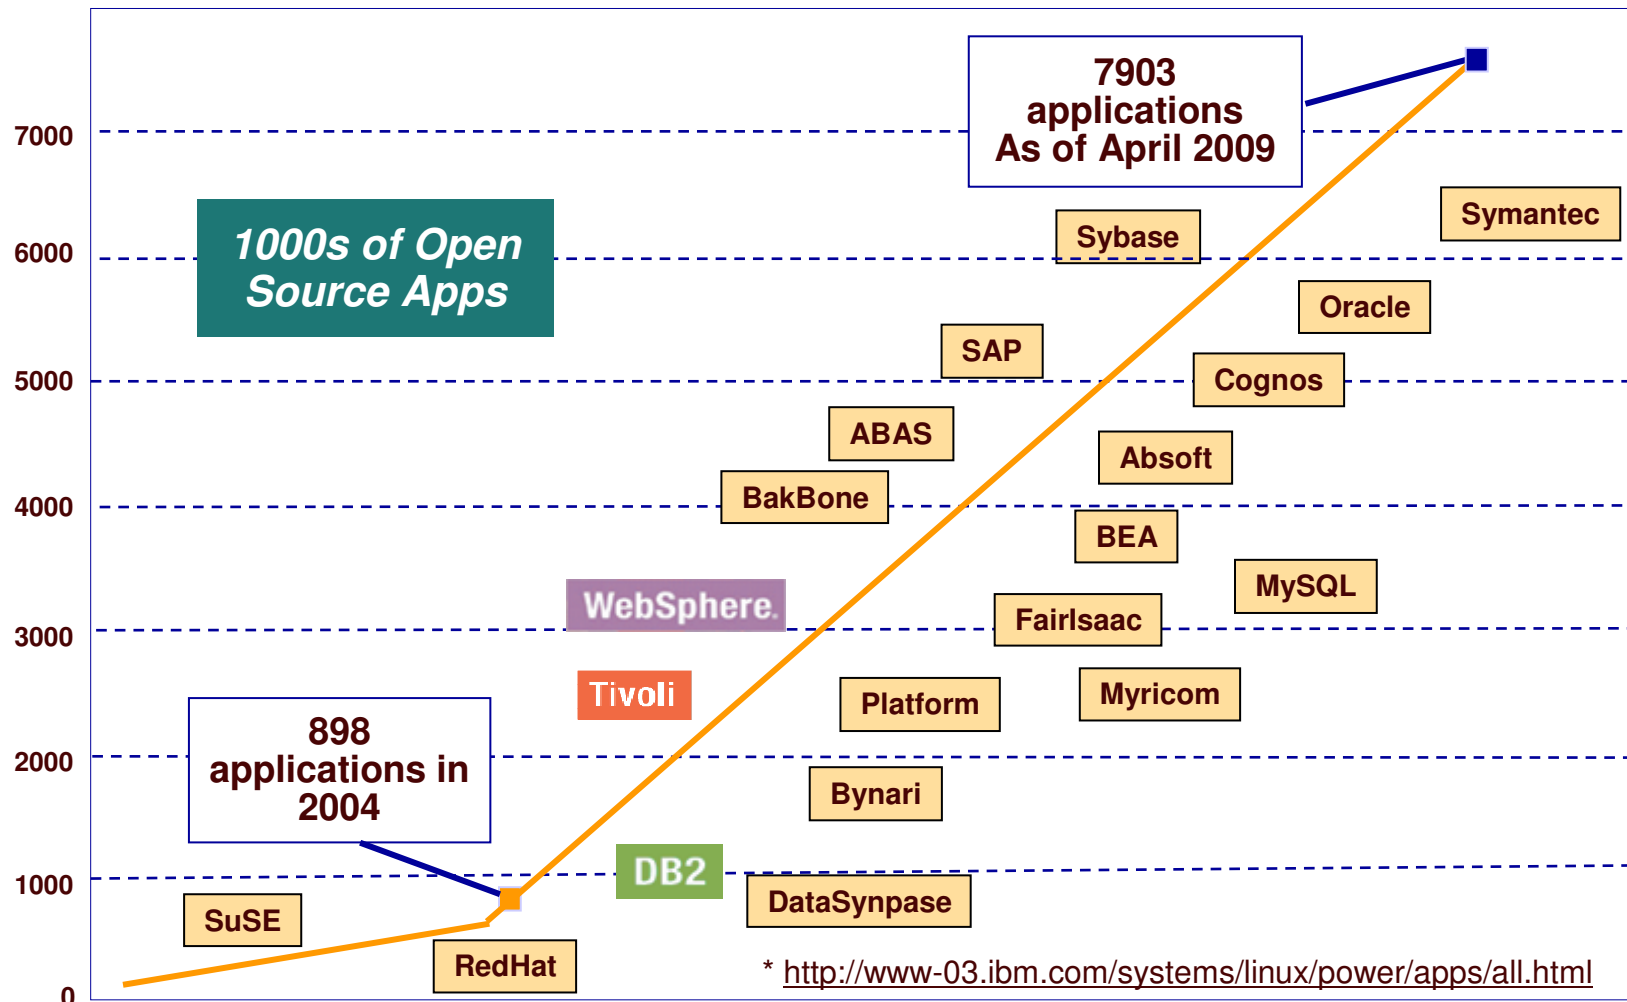
- Specialized and very detailed knowledge of IBM Systems and Software
- The LTC works with customers on unique proof of concept projects
  - Scale Out File Systems (SOFS)
  - Real Time Linux and Java

## IBM Enables Linux for New Markets

- Working with groups such as the Linux Foundation to address new workloads
- Expanding and providing leadership in:
  - Blue Cloud Computing
  - SOA / Web 2.0 / SaaS
  - Distributed computing and HPC

# Linux on POWER Apps

*7,900+ applications\* now available for Linux on POWER*



## Summary

Power Systems

Processor Update

- ▶ POWER6+ 560 / 570

New Product Announcements

- ▶ Power 520 4.7 GHz
- ▶ Power 550 5.0 GHz
- ▶ Power Blades JS23 / JS43
- ▶ New IO Drawers T19 / &24
- ▶ IO Adapters SAS
- ▶ SSD 69GB

Virtualization

Active Memory Sharing

AIX / Linux

# Questions????

## 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.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, AIX 5L, Chiphopper, Chipkill, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Purpose File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, System i, System p, System p5, System Storage, System z, Tivoli Enterprise, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](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 registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

Intel, Itanium, Pentium are registered trademarks 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, SPECchpc, 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.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Other company, product and service names may be trademarks or service marks of others.

Revised April 24, 2008

## 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 Power Systems Performance Report at [http://www.ibm.com/systems/p/hardware/system\\_perf.html](http://www.ibm.com/systems/p/hardware/system_perf.html).

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3, AIX 5L or AIX 6 were used. All other systems used previous versions of AIX. The SPEC CPU2006, SPEC2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L 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	<a href="http://www.tpc.org">http://www.tpc.org</a>
SPEC	<a href="http://www.spec.org">http://www.spec.org</a>
LINPACK	<a href="http://www.netlib.org/benchmark/performance.pdf">http://www.netlib.org/benchmark/performance.pdf</a>
Pro/E	<a href="http://www.proe.com">http://www.proe.com</a>
GPC	<a href="http://www.spec.org/gpc">http://www.spec.org/gpc</a>
VolanoMark	<a href="http://www.volano.com">http://www.volano.com</a>
STREAM	<a href="http://www.cs.virginia.edu/stream/">http://www.cs.virginia.edu/stream/</a>
SAP	<a href="http://www.sap.com/benchmark/">http://www.sap.com/benchmark/</a>
Oracle Applications	<a href="http://www.oracle.com/apps_benchmark/">http://www.oracle.com/apps_benchmark/</a>
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	<a href="http://www.siebel.com/crm/performance_benchmark/index.shtml">http://www.siebel.com/crm/performance_benchmark/index.shtml</a>
Baan	<a href="http://www.ssaglobal.com">http://www.ssaglobal.com</a>
Fluent	<a href="http://www.fluent.com/software/fluent/index.htm">http://www.fluent.com/software/fluent/index.htm</a>
TOP500 Supercomputers	<a href="http://www.top500.org/">http://www.top500.org/</a>
Ideas International	<a href="http://www.ideasinternational.com/benchmark/bench.html">http://www.ideasinternational.com/benchmark/bench.html</a>
Storage Performance Council	<a href="http://www.storageperformance.org/results">http://www.storageperformance.org/results</a>

Revised March 12, 2009

## Notes on HPC 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 Power Systems Performance Report at [http://www.ibm.com/systems/p/hardware/system\\_perf.html](http://www.ibm.com/systems/p/hardware/system_perf.html).

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3 or AIX 5L 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 5L 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.

SPEC	<a href="http://www.spec.org">http://www.spec.org</a>
LINPACK	<a href="http://www.netlib.org/benchmark/performance.pdf">http://www.netlib.org/benchmark/performance.pdf</a>
Pro/E	<a href="http://www.proe.com">http://www.proe.com</a>
GPC	<a href="http://www.spec.org/gpc">http://www.spec.org/gpc</a>
STREAM	<a href="http://www.cs.virginia.edu/stream/">http://www.cs.virginia.edu/stream/</a>
Fluent	<a href="http://www.fluent.com/software/fluent/index.htm">http://www.fluent.com/software/fluent/index.htm</a>
TOP500 Supercomputers	<a href="http://www.top500.org/">http://www.top500.org/</a>
AMBER	<a href="http://amber.scripps.edu/">http://amber.scripps.edu/</a>
FLUENT	<a href="http://www.fluent.com/software/fluent/fl5bench/index.htm">http://www.fluent.com/software/fluent/fl5bench/index.htm</a>
GAMESS	<a href="http://www.msg.chem.iastate.edu/gamess">http://www.msg.chem.iastate.edu/gamess</a>
GAUSSIAN	<a href="http://www.gaussian.com">http://www.gaussian.com</a>
ANSYS	<a href="http://www.ansys.com/services/hardware-support-db.htm">http://www.ansys.com/services/hardware-support-db.htm</a>
ABAQUS	Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results. <a href="http://www.simulia.com/support/v68/v68_performance.php">http://www.simulia.com/support/v68/v68_performance.php</a>
ECLIPSE	<a href="http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&amp;">http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&amp;</a>
MM5	<a href="http://www.mmm.ucar.edu/mm5/">http://www.mmm.ucar.edu/mm5/</a>
MSC.NASTRAN	<a href="http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm">http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm</a>
STAR-CD	<a href="http://www.cd-adapco.com/products/STAR-CD/performance/320/index/html">www.cd-adapco.com/products/STAR-CD/performance/320/index/html</a>
NAMD	<a href="http://www.ks.uiuc.edu/Research/namd">http://www.ks.uiuc.edu/Research/namd</a>
HMMER	<a href="http://hmmmer.janelia.org/">http://hmmmer.janelia.org/</a> <a href="http://powerdev.osuosl.org/project/hmmmerAltivecGen2mod">http://powerdev.osuosl.org/project/hmmmerAltivecGen2mod</a>

Revised March 12, 2009



## Notes on performance estimates

### rPerf for AIX

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. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.

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.

=====

### CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: [www.ibm.com/systems/i/solutions/perfmgmt/resource.html](http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html)

Revised April 2, 2007