



POWER6

Next Generation in Computing

Pat O'Rourke
pmorour@us.ibm.com



IBM Systems

POWER6 System Announcement Highlights

IBM POWER6™ Processor Technology

- ▶ 5th Implementation of multi-core design
- ▶ ~100% higher frequencies

POWER6 System Architecture

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

HMC Enhancements

Enhanced Virtualization

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

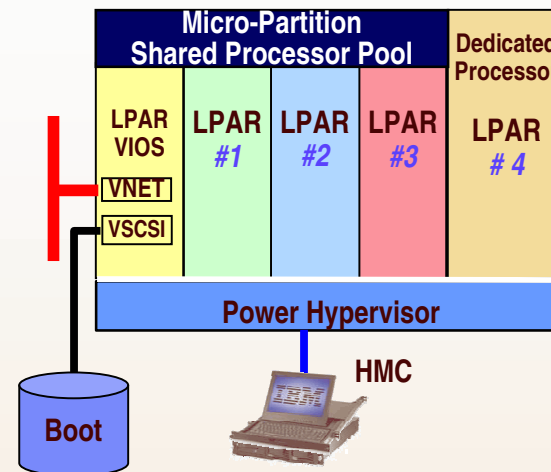

Availability


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

Operating System Preview

- ▶ AIX 6

POWER6





Processors

POWER6 Architecture

POWER Design

3.5 – 4.7 GHz

>790M transistors

.065 micron

POWER6 Characteristics

Ultra-high frequency dual-core chip: > 3.5 GHz

- ▶ 7-way superscalar, 2-way SMT core
- ▶ 9 execution units
 - 2LS, 2FP, 2FX, 1BXU, 1VMX, 1DP
- ▶ 790M transistors, 341 mm² die
- ▶ 2x4MB on-chip L2 – point of coherency
- ▶ On-chip L3 directory and controller

Technology

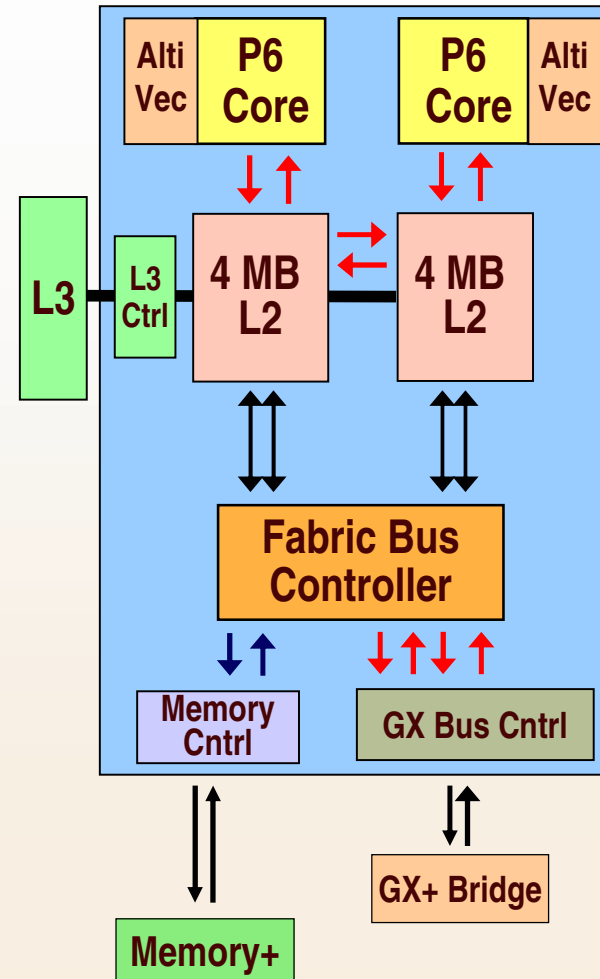
- ▶ CMOS 65nm lithography, SOI Cu

High-speed elastic bus interface at 2:1 freq

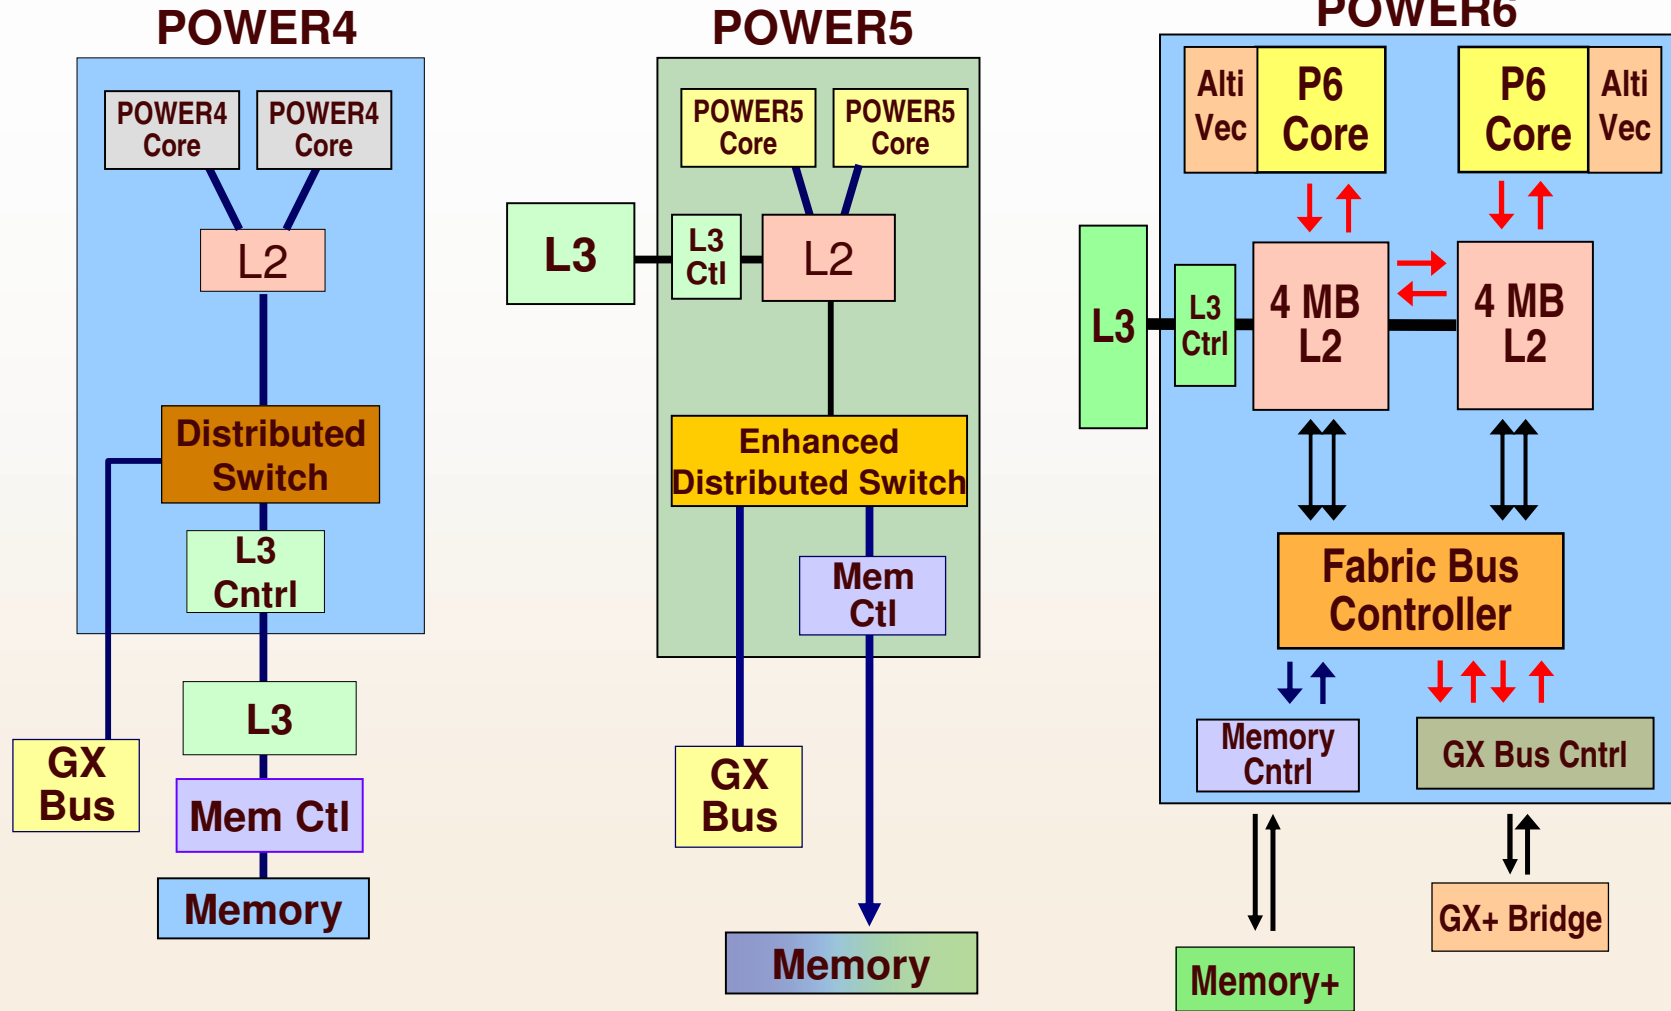
Full error checking and recovery

Dynamic power saving

- ▶ Advanced Clock gating



POWER architecture



Processor History 2001 - 2007

**Dual Core
1GHz
Distribute Switch**




POWER4
414 mm²
1.1 – 1.3 GHz




POWER4+
267 mm²
1.5 – 1.9 GHz

**Multi-Threading
Memory Cntrl on Chip**




POWER5
389 mm²
1.65 – 1.9 GHz



POWER5+
245 mm²
1.9 – 2.3 GHz

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

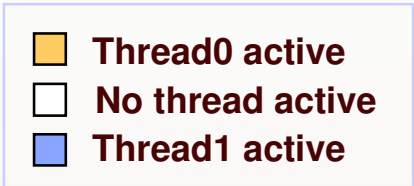


POWER6
341 mm²
3.5 – 5.0 GHz

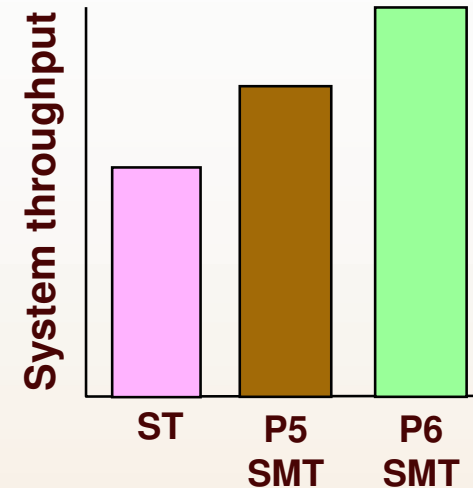
2001 2002 2003 2004 2005 2006 2007 2008

POWER6 delivers improved system utilization through enhanced 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**



System p 570

POWER6 System p 570

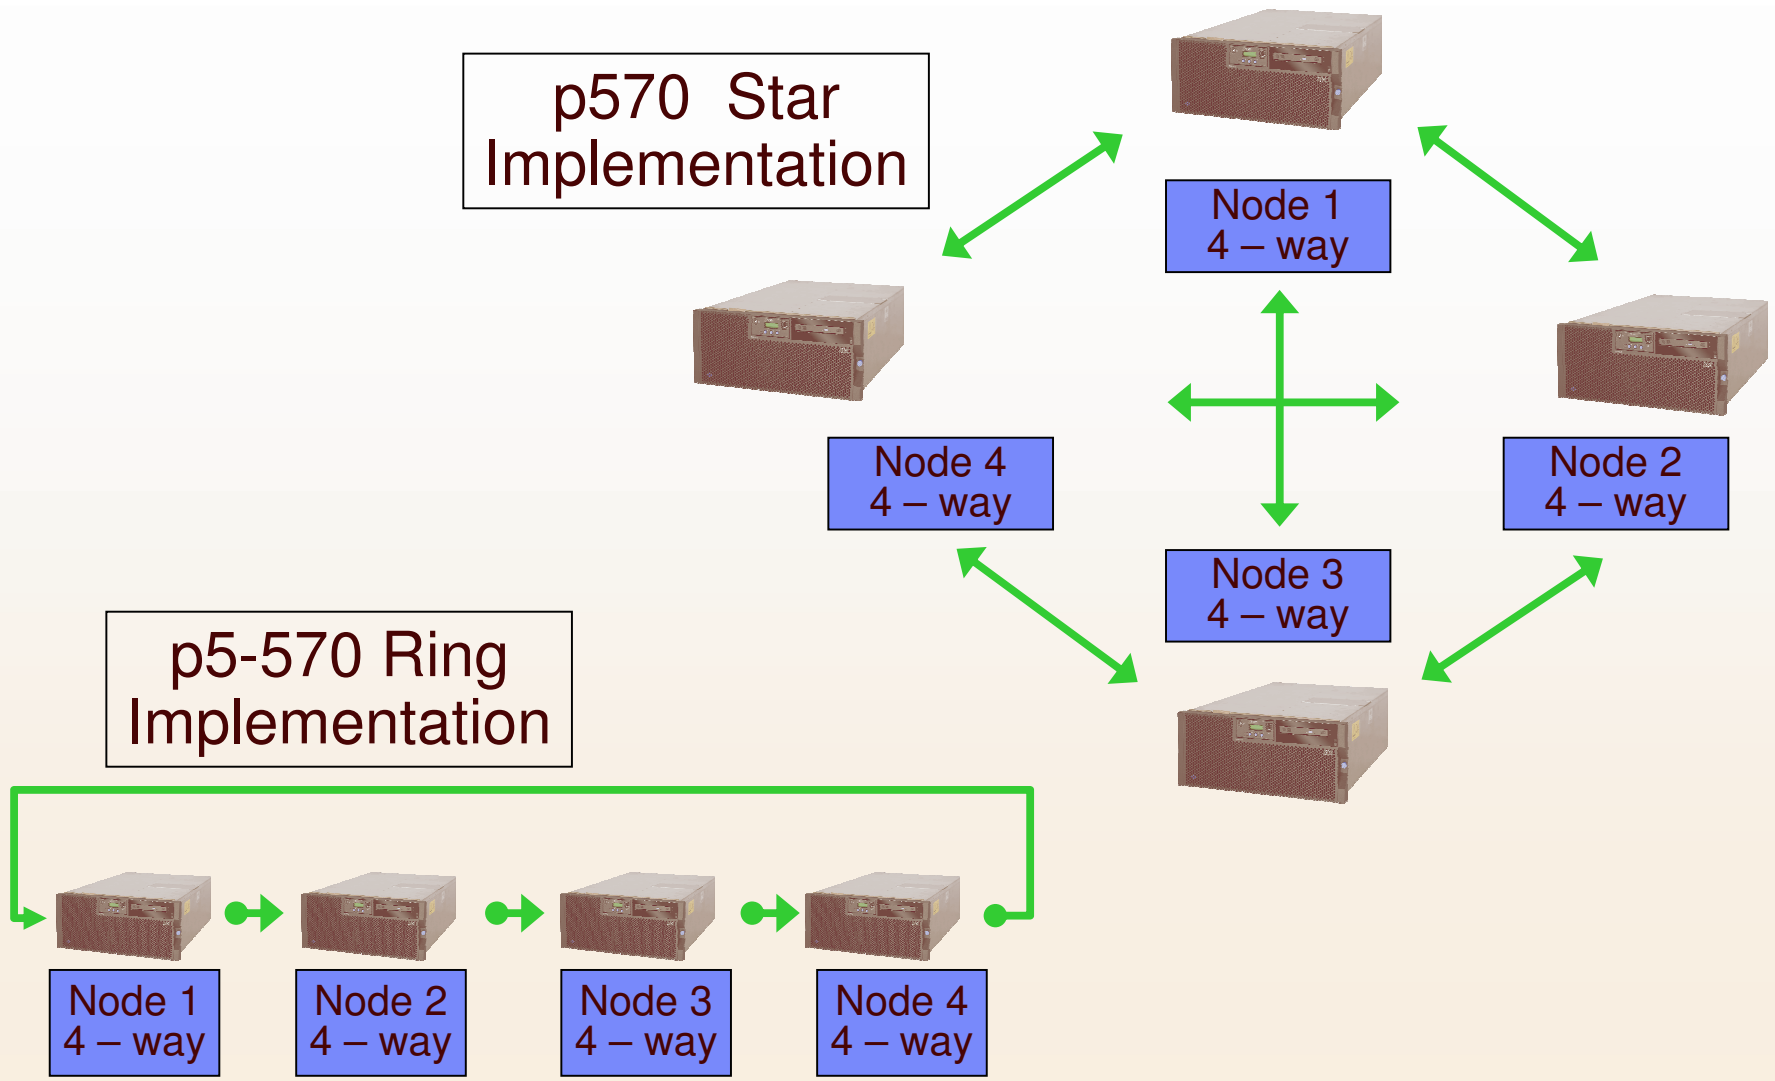


Processor	3.2 / 4.2 / 4.7 GHz POWER6	
L3 Cache	32 MB / Chip	
Redundant Power & Cooling	Yes	
Redundant Server Processor	Optional (Two Nodes minimum)	
Redundant Clock	Yes (Cold Detect) (Two Nodes minimum)	
Warranty	1 Year, NBD	
Operating System Support	AIX 5.2 & AIX 5.3 RHEL 4.5 for POWER, SoD RHEL5 SLES 9 or 10 for POWER	
	Per Building Block	System Maximum
Processors	4	16
DDR2 Memory (Buffered)	2-192	768
SAS Disk Bays (3.5")	6	24
Media Bays	1 Slim-line	4 Slim-line
SAS/SATA Controller	1	4
PCI (Internal)	4 PCIe, 2 PCI-X 266	16 PCIe, 8 PCI-X 266
GX Bus Slots	2	8
Integrated Virtual Ethernet	Dual Port 10/100/1000 Quad Port 1 Gb (Optional) Dual Port 10 Gb (Optional)	4 Dual Port 10/100/1000 4 Quad Port 1 Gb (Optional) 4 Dual Port 10 Gb (Optional)
USB	2	8
Remote I/O Drawers	8 Max	32 Max
Dynamic LPARs	40	160

Midrange Functional Differences


POWER5 p570	POWER6 Midrange
Ring Cabling Architecture	Point to Point Cabling Architecture
2GB to 128 GB per node	2GB to 192 GB per node
8 DIMM sockets / CPU Card	12 DIMM sockets / CPU Card
DDR2 DIMMs	Fully Buffered DIMMs
SCSI DASD / 6 Drives	SAS DASD / 6 Drives
Six PCI-X slots	4 PCIe & 2 PCI-X slots
Media devices = 2	Media Devices = 1
Separate DASD & Media Controllers	Same controller for DASD & Media
Dual DASD Controllers	One SAS & SATA DASD / Media Controller
Two 1Gb Ethernet ports	Integrated Virtual Ethernet <i>One of the following:</i> <ul style="list-style-type: none"> ▶ Two 1Gb Ethernet ports ▶ Four 1Gb Ethernet ports ▶ Two 10Gb Optical SR Ethernet ports
GX Bus Slots: 1	GX Bus Slots: 2
Integrated RIO ports: 2	Integrated RIO ports: 0

POWER6 p570 SMP Fabric


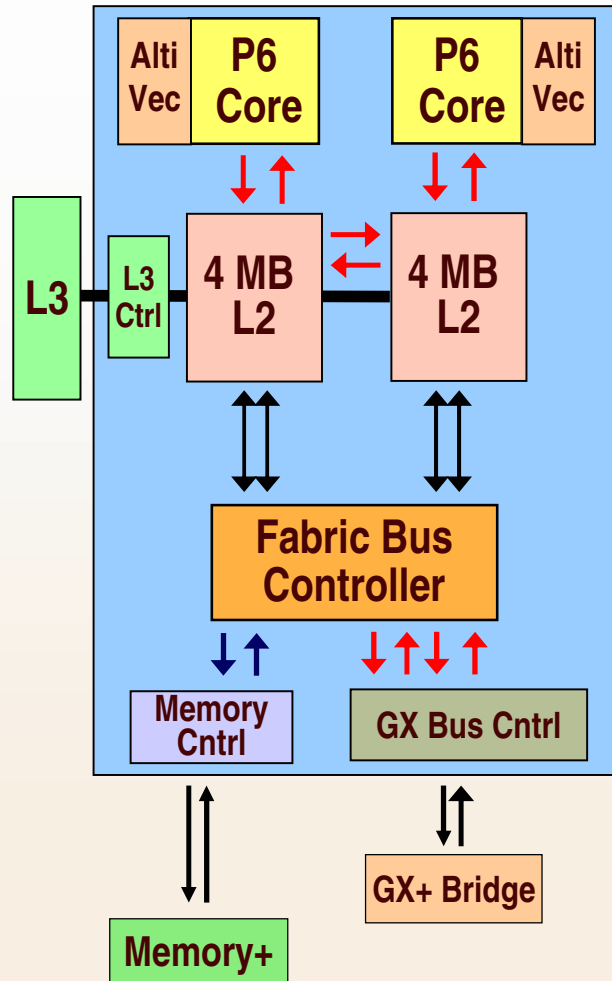


Building Block Modules: SMP Mid-range Server


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



System p 570 POWER6 Processor Packaging



Single Chip Module
POWER6 & L3




POWER6 Processor
Module

System p 570 Bandwidth

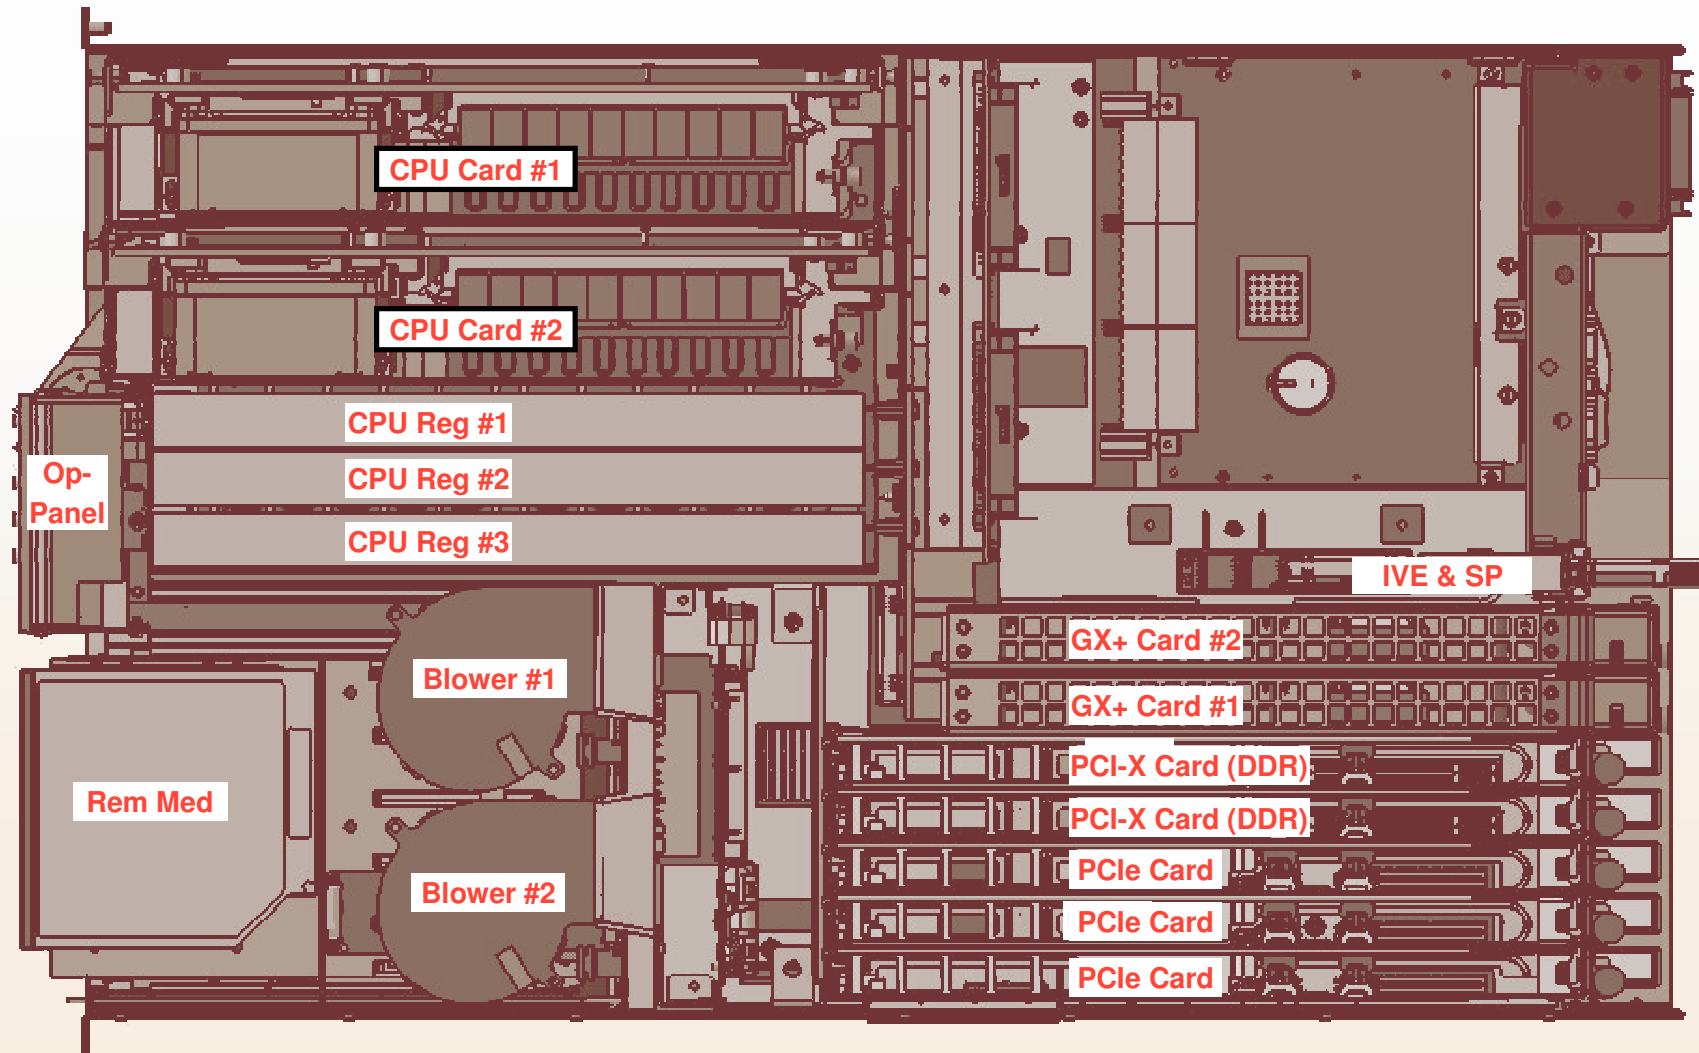
Memory	Bandwidth
L1 (Data)	75.2 GB/sec
L2	300.8 GB/sec
L3	37.6 GB/sec
Memory	32 GB/sec
Inter-Node Buses (16w)	75.2 GB/sec
Intra-Node Buses (16w)	100.26 GB/sec
Internal I/O Bus	4.7 GB/sec / node
GX Bus Slot 1	4.7 GB/sec / node
GX Bus Slot 2	6.266 GB/sec / node
Total I/O Bandwidth	62.6 GB/sec (16w)

Calculations for 4.7 GHz processors and 667 MHz memory


Front View



Top View




Back View




POWER6 / 5+ Midrange Processor / Memory packaging


SCM @ 3.5 – 4.7GHz
 Up to 96 DDR2+ Memory
 2 Cores / book




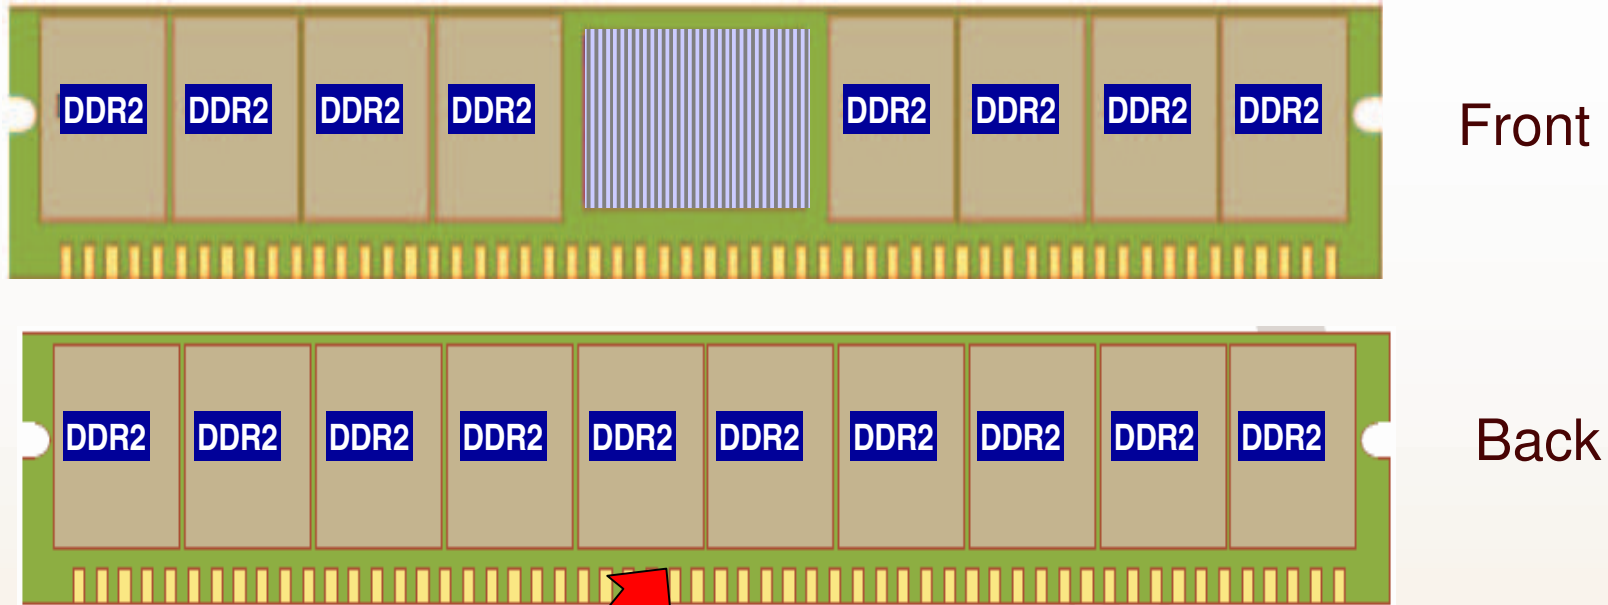
p5-570+ DCM @ 2.2GHz
 Up to 64GB DDR2 Memory
 2 Cores / book



p5-570+ DCM @ 1.9GHz
 Up to 32GB DDR2 Memory
 2 Cores / book



POWER6 Buffered Memory DIMMs...



**50% More DIMM Slots
Greater Memory Flexibility**

System p 570 SoD

Redundant Service Processor

For POWER6 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 based p570 users via a system firmware update.


Hot Add Support

POWER6 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 based p570 users via a system firmware update.

Planned availability: No later than the end of 2007

Service Processor Cabling Layouts





8 Core



12 Core

16 Core





Performance

POWER6 Leadership

Number 1 Leadership


- TPC-C
 - 1,616,612 @ \$3.42

- SAP
 - 8000 Users

- Oracle Apps
 - 3000 Users

- Other #1 benchmarks
 - SPECjbb2005 (Java)
 - SPECint_rate2006
 - SPECfp_rate2006

- Benchmarks for 4 core, 8 core, and 16 core configurations




System p 570 Remote I/O Drawer


POWER6 Remote I/O Drawer 19" Rack


7314-G30

- ½ x 19" 4u Rack Mount
 - Optional enclosure for 2 individual drawers
- 6 PCI Adapter Slots
 - PCI-X 2.0 (DDR)
 - 64-bit @ 266MHz (2GB/s)
- "InfiniBand" Host Interface (Loop Architecture)
 - 12x 2.5Gb/s Full Duplex (30 Gb/sec)
 - Short Run (Intra-rack) or Long Run (Inter-rack)
 - 4 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



Two drawers






RAS

Primary POWER RAS Features

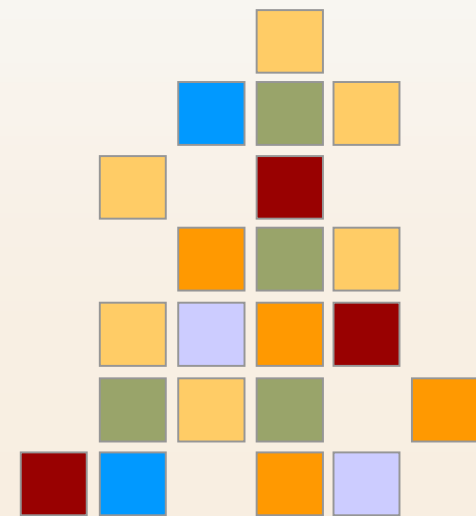
- ✓ **Processor Instruction Retry**
- ✓ **Alternate Processor Recovery**
- ✓ **First Failure Data Capture**
- ✓ **DDR Chipkill™ memory**
- ✓ **Bit-steering/redundant memory**
- ✓ **Service Processor Failover***
- ✓ **Dynamic Firmware Maintenance***
- ✓ **Hot I/O Drawer Add***
- ✓ **I/O error handling 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..



HMC

POWER6 Hardware Management Console Enhancements

Support for POWER5 and POWER6 on same HMC
Updated HMC hardware (Intel® technology refresh)
No change: Hardware scaling support

- ▶ 7310/7042-C06 (desktop)
- ▶ 7310/7042-CR4 (rack-mount) 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



New HMC Release – Version 7

To be released as Version (7.3.1.0)

- ▶ Provides synergy with System Firmware, i.e. 01_SF310_010

Compatibility


- ▶ Command line compatibility preserved
- ▶ DHCP usage (essentially) the same as POWER5
 - Required for 24" systems (590 / 595 / 575 and POWER6 follow-ons)
- ▶ Upgrade from HMC Version 6 (only)

HMC Hardware Supported

- ▶ **NOTE: Oldest Machine Type 7315 ("POWER4" HMC) is not supported**
- ▶ Older Machine Type 7310 ("POWER5" HMC) is supported
- ▶ New Machine Type is 7042 ("POWER6" HMC)

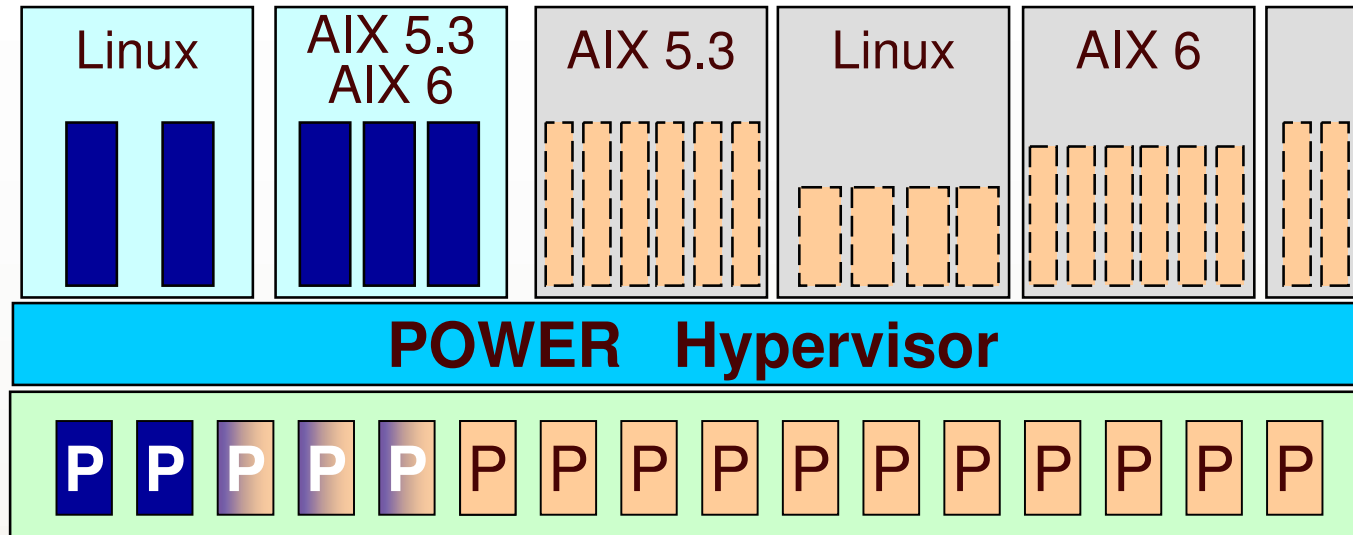
Servers Supported

- ▶ Version 7 required for POWER6 Servers
- ▶ Also supports all POWER5 Servers
 - Server must be at release 240
 - Testing statement, no enforcement



Virtualization

Dedicated Shared Processors



Dedicated / Shared Processors


Shared (Non-Dedicated) 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

Utilizing Unused Dedicated Processor Capacity

■ 1-way Dedicated □ Wasted Dedicated ■ 0.5 Uncapped 1 ■ 0.5 Uncapped 2



■ 1-way Dedicated □ Wasted Dedicated ■ 0.5 Uncapped 1 ■ 0.5 Uncapped 2



- ❖ 2-way server with:
 - ▶ Two evenly weighted 1-way uncapped partitions (100% Utilized)
 - ▶ 1-way dedicated partition (Variable workload)
- ❖ Unused processor cycles in dedicated partition are wasted

❖ Donation option enabled, a dedicated partition donates its excess cycles to the uncapped partitions. Unused processor cycles are utilized.

Result: Better System Utilization

Integrated Virtual Ethernet Overview




Naming

- ▶ **Integrated Virtual Ethernet** – External name in marketing & 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 fiber: 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 customer configuration and environment
 - E.g. Speed, frame size, duplex

Integrated Virtual Ethernet Adapter



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

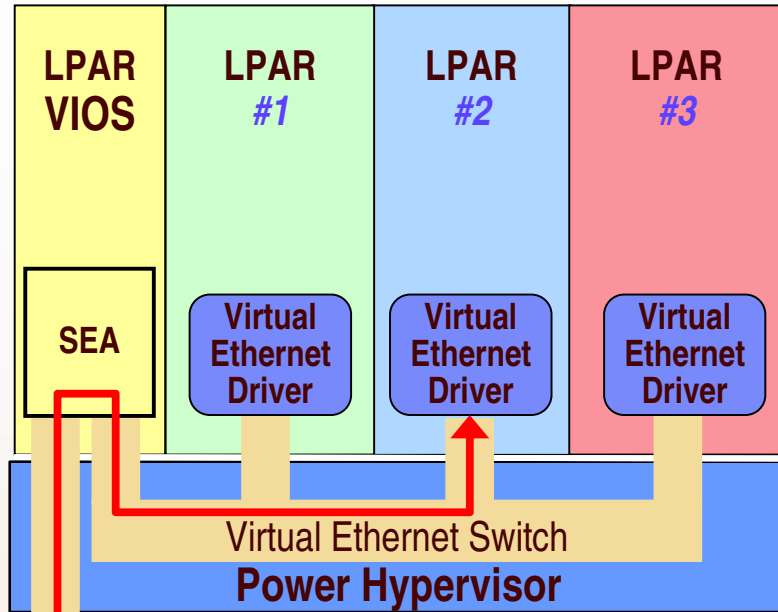
Non VIOS Partition: Address Sharing (MAC Addresses)

- Time Slicing "Physical" Ethernet adapter resources

VIOS Partitions: IVE ports dedicated

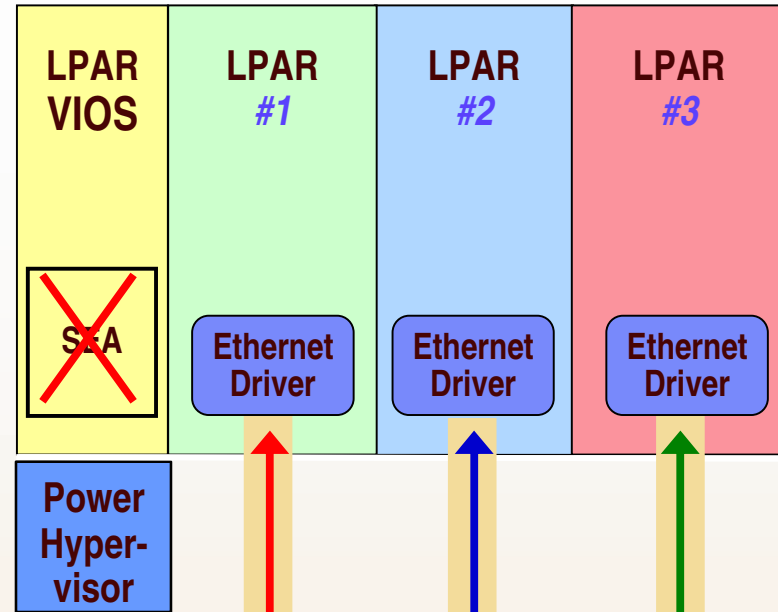
Integrate Virtual Ethernet How it works.....

Option 1



Or

Option 2




Integrated Virtual Adapter



Native Performance
Software Transparency
AIX 5.3 / 5.4 & Linux
P6 Midrange & Rack

Integrated Virtual Adapter



SoD: 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





Capacity On Demand

CoD Offering Evolution for POWER6

Today

Permanent
<p>CoD Processors: 1 processor increment Memory: 1 GB increment</p>
Temporary
<p>On/Off CoD Activations: Manual Utilization Reporting Required (Contract) Post-pay Integrated into Capacity Back Up offering</p>
<p>Capacity Back Up (CBU) CBU offering for p590 / p595</p>
<p>Reserve CoD Pre-pay usage for blocks of time Charges based on measured workload No Contracts Operates within the shared pool</p>
<p>Trial CoD Standard Exception Web based distribution</p>

2007 / 2008 POWER6

Permanent
<p>CoD Processors: 1 processor increment Memory: 1 GB increment</p>
Temporary
<p>On/Off CoD Activations: Manual Utilization Reporting Required (Contract) Post-pay Integrated into Capacity Back Up offering</p>
<p>Capacity Back Up (CBU) CBU offering for POWER6 High End</p>
<p>Utility CoD Post pay or pre-pay (1 Processor Minute) Charges based on measured workload No Contracts Operates within the shared pool</p>
<p>Trial CoD Standard Exception Web based distribution</p>

Utility CoD: What is it?

Utility CoD

Automatically provides additional processor capacity on a temporary basis within the Shared Processor Pool.


Usage is measured in **processor minute increments** and is reported via a web interface. Billing is based on the reported usage.

Customer Value

Utility CoD will automatically provides additional processor capacity when needed for a short, temporary basis within the shared processor pool.

Charges will only occur when workload requirements require the need for the Utility CoD processors.

Usage is measured in processor minute increments and is cost effective for environments where it is critical to be responsive to short, temporary workload spikes.



AIX 6


Preview: AIX 6

AIX® 6: Binary compatible* with AIX 5L™
Reflect unity with POWER 6



The P O W E R of SIX – AIX 6 and POWER6™

- Workload Partitions
 - Live Application Mobility
 - Live Partition Mobility
 - Storage Keys
 - Hardware Decimal Floating-Point
 - Dynamic Variable Page Size
- } **AIX 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/>

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

AIX6 Processor Support....

AIX6 will not support 32 bit hardware and older 64 bit hardware

Processors no longer supported:

- All 32 bit processors i.e. F50, E30, etc.
- RS64 family of servers i.e. S80, H80, M80, etc.
- POWER3 servers i.e. p610, p640, etc.

Processors supported:

- PPC970
- POWER4
- POWER5
- POWER6

AIX 6 will support both 32 bit & 64 bit libraries & applications

- Full support for 32 bit applications



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

AIX 6 Beta Program

AIX 6 Beta



Select customers & ISVs

Tens

Physical Media

Feedback on functionality

Beta Support team

Questionnaire, Support interaction



Participants

Number of participants

Distribution method

Goal

Support

Feedback

AIX 6 “Open Beta”



Open to all

Hundreds

Web download only

Generate mind share

None

Web feedback only

AIX Open Beta Program

<http://www-03.ibm.com/servers/aix/6/beta.html>



Open Beta Overview

Open to everyone

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

Summary

POWER6



New Processor Technology

- ▶ POWER6

POWER6 System Architecture

- ▶ New generation of servers
- ▶ New IO
- ▶ New IO Drawers



Enhanced Virtualization

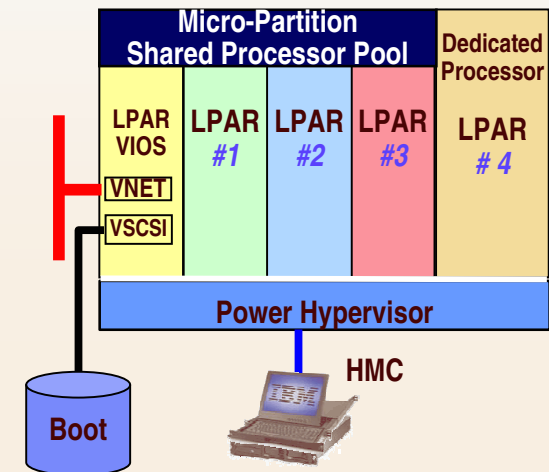
- ▶ Partition Mobility (SoD)
- ▶ Dedicated Shared Processors
- ▶ Integrated Virtual Ethernet

Availability

- ▶ New RAS features

Operating System

- ▶ AIX 6



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 5L, 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, 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, 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.

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 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.

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.shtml
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 5L 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