



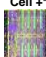


Building a Better Infrastructure With IBM Middleware on System p

POWER Hardware Improves Utilization
and Reduces Costs

At the Heart of IBM Power™ Systems is the POWER Processor

2001 POWER4™	2003 POWER4+™	2004 POWER5™	2006 POWER5+™	2007 POWER6™	Future* POWER7™	<i>General Purpose Processors</i>
						

PowerPC 405 	PowerPC 440 	Cell 	Cell 	Cell + 	<i>Specialty Engines</i>
--	--	---	---	---	--------------------------



Nintendo Wii



Game Systems

Blades

Entry Level to High-end Servers

High Performance Computers

*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
02 - POWER Hardware 2008 v2.2

The POWER Processor Is Everywhere

The POWER processor is used in everything from game systems to special purpose high performance computers.

- ▶ Largest UNIX systems market share
- ▶ 50% of automobiles worldwide
- ▶ All of the top three game consoles
 - Playstation 3, Xbox 360, Wii
- ▶ 47% of world's fastest computers (Top500)

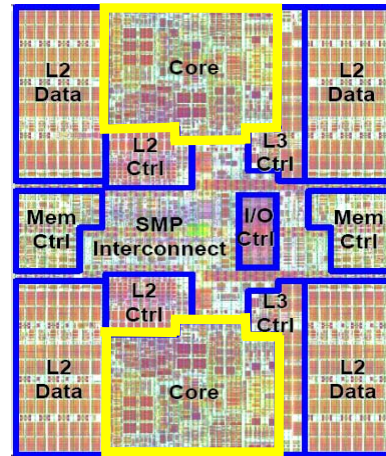
Quiz: What Does the Acronym **POWER**, Represent?

- A. **P**ackets **O**f **W**ebSphere **E**ngineered **R**equests
- B. **P**erformance **O**ptimized **W**orkloads **E**ngineered **R**esilient
- C. **P**rocessor **O**ptimized **W**ith **E**nergy **R**eduction
- D. **P**erformance **O**ptimized **W**ith **E**nhanced **R**ISC
- E. **P**rogrammatic **O**ctal **W**avelengths **E**lectromagnetically **R**efined

POWER Architectural Concepts – Multi-core

- Why multi-core?
 - ▶ Laws of physics catching up with single core processors, it is no longer a matter of simply increasing speed to increase performance
 - ▶ Chip manufacturing density enables more than one central processing unit (CPU) to be placed on a single chip or socket
- Each central processing unit is referred to as a “core”*
- Each core presents itself as an independent physical processing unit

Dual core POWER6 processor



* CPU = core

02 - POWER Hardware 2008 v2.2


5

Processor Concepts - Single Thread Execution

Four instruction threads in memory, waiting to be executed

The processor fetches instructions from a single thread and continues to execute instructions from that thread until a decision is made to execute another thread

Popular design, but branches and data fetches cause idle CPU cycles




02 - POWER Hardware 2008 v2.2

6

Processor Concepts – Symmetric Multi-processing (SMP)

Four instruction threads in memory, waiting to be executed

Each processor fetches instructions from a single thread




- Achieves higher overall throughput by utilizing all available CPUs
 - ▶ Operating system assigns threads to CPUs
 - ▶ Still single thread execution with idle cycles within each CPU
- System p is architecturally engineered to exploit full SMP function up to the largest number of processors
 - ▶ Hardware
 - ▶ Virtualization

POWER Architectural Concepts – Simultaneous Multi-threading (SMT)


Four instruction threads in memory, waiting to be executed

Each processor presents itself as two virtual processors to the operating system and fetches instructions from two threads.



- Instructions from both threads are executed in each processor concurrently
- Utilizes unused CPU cycles

Simultaneous Multi-threading (SMT) Delivers More Processing Power




- Net result
 - ▶ Better performance and processor utilization increases throughput and response time, depending on workload

02 - POWER Hardware 2008 v2.2

9

System p Architectural Concepts – Server Virtualization

- Run multiple operating system images on a single platform
 - ▶ Like having several separate servers running simultaneously
- Each image runs in a logical partition
- The POWER Hypervisor™ shares the underlying physical resources dynamically among the partitions



Dynamically shared physical resources in a single platform

We will discuss virtualization and consolidation in more detail later today.

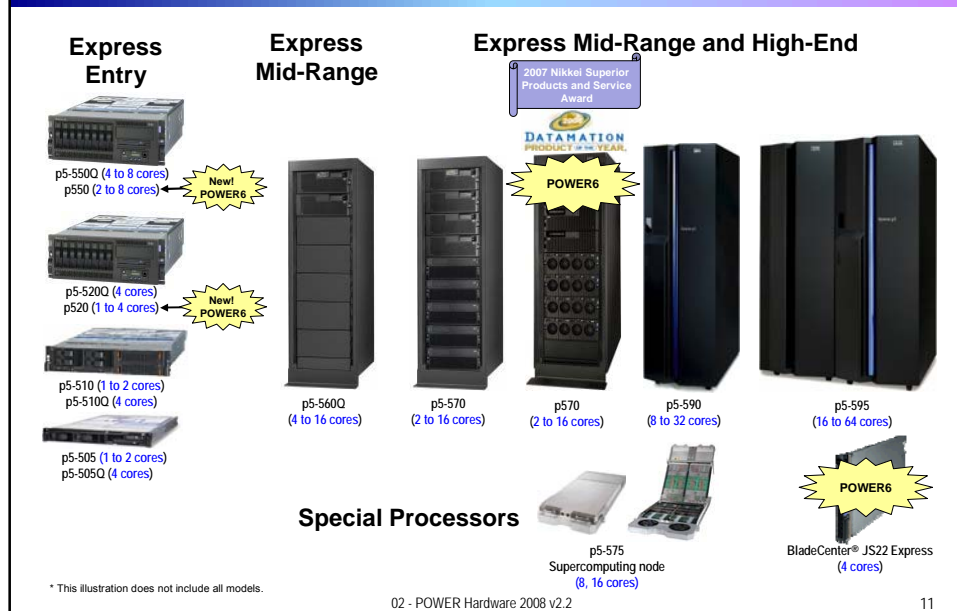


IBM

02 - POWER Hardware 2008 v2.2

10

IBM Power™ Systems Has a Complete Line of System p Servers for Your Data Center *



The diagram illustrates the IBM Power Systems product line, categorized into three main groups: Express Entry, Express Mid-Range, and Express Mid-Range and High-End. Each group contains several server models with their respective core counts. The Express Entry group includes models like p5-550Q, p550, p5-520Q, p520, p5-510, p5-510Q, p5-505, and p5-505Q. The Express Mid-Range group includes p5-560Q, p5-570, p570, and p5-590. The Express Mid-Range and High-End group includes p5-575 Supercomputing node and BladeCenter JS22 Express. The diagram also features several award logos, including the 2007 Nikkei Superior Products and Service Award and the 2007 DataMation Product of the Year award. A 'Special Processors' section shows a p5-575 Supercomputing node and a BladeCenter JS22 Express. A yellow starburst highlights 'New! POWER6' for several models. A small text note at the bottom left states: '* This illustration does not include all models.' The bottom center of the slide contains the text '02 - POWER Hardware 2008 v2.2' and the bottom right corner contains the number '11'.

Express Entry

- p5-550Q (4 to 8 cores)
- p550 (2 to 8 cores)
- New! POWER6
- p5-520Q (4 cores)
- p520 (1 to 4 cores)
- New! POWER6
- p5-510 (1 to 2 cores)
- p5-510Q (4 cores)
- p5-505 (1 to 2 cores)
- p5-505Q (4 cores)

Express Mid-Range

- p5-560Q (4 to 16 cores)

Express Mid-Range and High-End

- 2007 Nikkei Superior Products and Service Award
- 2007 DataMation Product of the Year
- POWER6
- p5-570 (2 to 16 cores)
- p570 (2 to 16 cores)
- p5-590 (8 to 32 cores)
- p5-595 (16 to 64 cores)
- BladeCenter® JS22 Express (4 cores)
- POWER6

Special Processors

- p5-575 Supercomputing node (8, 16 cores)

* This illustration does not include all models.





02 - POWER Hardware 2008 v2.2

11

Quiz: What Is a Grand Slam?


- A. In tennis, a singles player or doubles team that wins all four Grand Slam event titles in the same year.
- B. In baseball, a Grand Slam is a home run with all of the bases occupied.
- C. In computing, a Grand Slam is being the performance leader in four major benchmarks at the same time.
- D. A breakfast item on the menu at Denny's restaurant.
- E. All of the above

IBM System p with POWER6 "Grand Slam" for Major Workloads



* Source: www.spec.org/
** Source: www.tpc.org/

Not Just a Grand Slam, It Is a "Slam Dunk," System p Versus The Best Competitive Result

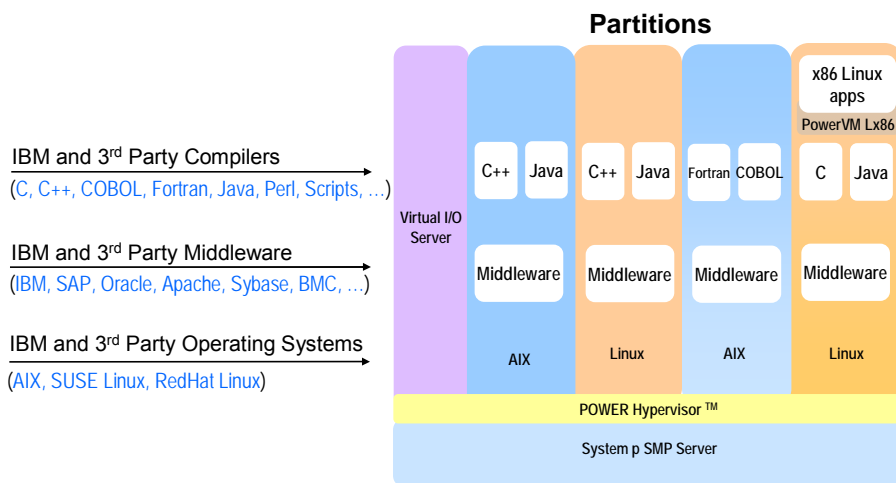


Source:
<http://www.spec.org>
<http://www.tpc.org>
<http://www.sap.com/benchmark/>
<http://www.performance.siebel.com/performanc.html>
<http://www.oracle.com/technology/products/epo/ra.htm>
 For a complete list of all IBM POWER results, go to <http://ibm.com/systems/o/benchmark>

System p with POWER6 Delivers More Than Just Lightning Fast Processor Speed

- Breaks the 4 GHz milestone within the same energy envelope as POWER5™
- On-chip energy efficiency intelligence for dynamic energy savings
- Balanced design with highest system bandwidth
- Integrated hardware accelerators for specialized performance
- Designed for continuous availability

System p Supports Standard Programming Interfaces In a Flexible Ecosystem




Survey of 277 Enterprise Customers Rank AIX the Best


- AIX has a better history of binary compatibility

Compatibility guarantee: <http://www-03.ibm.com/systems/p/os/aix/compatibility/>

Operating System Features



Operating System Quality




Source: Unix Vendor Preference Survey, December 2006, Gabriel Consulting Group

System p User Interfaces


Command Interface on ASCII Terminal



Command Interface on ASCII Emulator



X Windows Server with System p



ASCII Terminal

ASCII Emulator

Desktop OS

X Windows Server

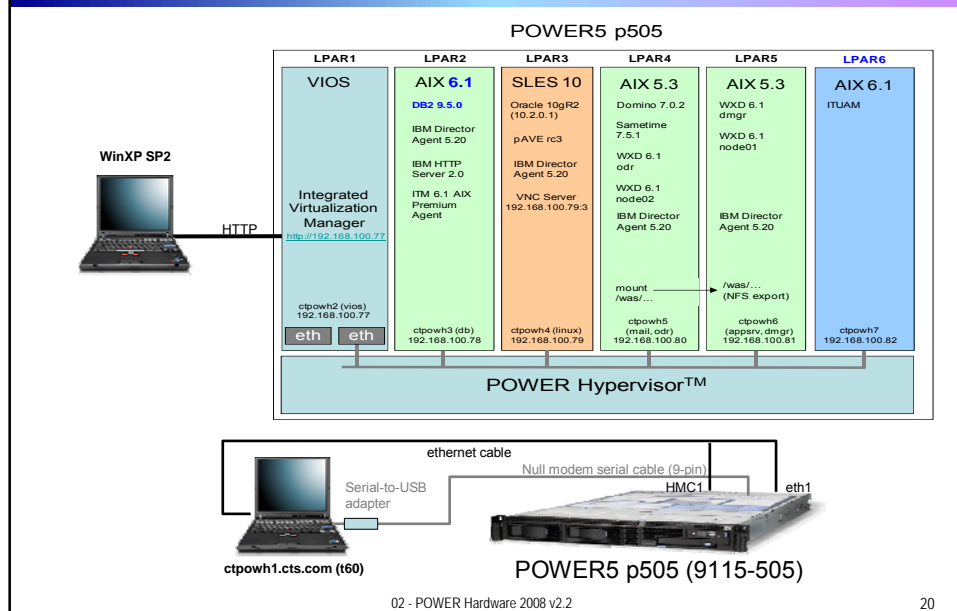
Desktop OS

Server Application



p570

DEMO: Demonstrations in Partitions on One System p



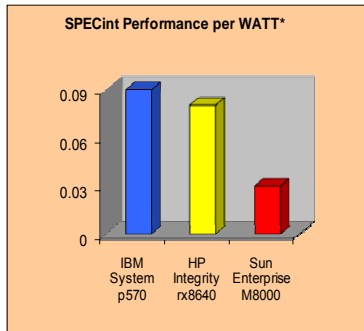
Cool Blue System p is Environmentally Friendly

- Energy efficiency management functions for POWER6 chips
 - ▶ Power reduction
 - ▶ Nap mode
 - ▶ Thermal tuning
- Energy efficiency management and control functions in the new System p
 - ▶ Enhanced design and implementation
 - ▶ Variable fan speed
 - ▶ Power off PCI slots when not in use
 - ▶ Rear door heat exchanger
 - ▶ System energy efficiency management software
- Consolidation and virtualization inherently reduce energy requirements



Cool Blue System p Is Green

- IBM Project Green - The IBM Energy Efficiency Initiative
 1. Best practices and services
 2. Technology and management innovations
 3. Environmental responsibility
- System p570's innovative "green" design saves energy



The top 26 most energy efficient computers are Power processor based systems (Green500).**

* Performance per WATT based on dividing SPECint_rate2006 performance by vendor recommended maximum power requirement.
 Source: www.spec.org
 ** Only systems where measured power is available.
 Source: <http://www.green500.org/>


What Is System p Capacity Upgrade on Demand?



p570
2-16 cores
2-768 Gb RAM


Customer orders p570* with maximum cores and memory, but with only half activated

Only pay in full for active resources and get 65% discount to cover the cost of unactivated processors



Customer determines they need additional capacity for their workload and activates** remaining processors and memory

Additional payment required for newly activated resources, total cost is the same as if all hardware was activated at the time of the original purchase



* System optional capacity options do not require purchase of a new system. Assumes appropriate storage, IO drawers, software licensing, etc. for the configuration. Only processors and memory are shown here as an illustration.

** Activation occurs via one of the optional capacity options available for System p. Not all processors or memory need to be activated at one time. This was done as an illustration

Also Four Other Capacity on Demand Options

- Capacity Upgrade on Demand (processors, memory)
 - ▶ Activate capacity permanently for non-disruptive growth
- On/Off Capacity on Demand (processors, memory)
 - ▶ Temporary capacity for fluctuating workloads
- Trial Capacity on Demand (processors)
 - ▶ Temporary capacity for workload testing or any one time need
- Utility Capacity on Demand (processor minutes)
 - ▶ Autonomic, charges based on measured workload
- Capacity Backup (processors, memory)
 - ▶ Only pay for installed disaster recovery capacity when used
 - ▶ Can be used with IBM HACMP™ V5 and HACMP/XD software (5765-F62) in failover scenarios

System p Reliability


- Error detection
 - ▶ Predictive failure analysis on processors, caches, memory, I/O and DASD
- Built-in redundancy
 - ▶ Redundant power supplies, fans, service processor
- Processor instruction retry
 - ▶ Transparent retry of soft errors on application fault
- Storage protection keys
 - ▶ POWER6 storage keys will isolate data and protect against memory overlay that can cause subtle, intermittent problems
- First Failure Data Capture (FFDC)
 - ▶ Automatic capture of diagnostic information

System p Serviceability

- AIX supports concurrent maintenance
 - ▶ Non-disruptive fixes to AIX
 - ▶ No downtime (reboot) required to apply fix and make it active
- Hot swappable parts replacement
- Electronic Service Agent (ESA)
 - ▶ Automatically contact IBM if problem




ESA runs automatically in the background!



System p High Availability Cluster Multi-Processing (HACMP) Increases Availability

- HACMP manages Resource Groups
- Resource Groups move from one node to the other when an event happens
- Resources in a group could be one of the following
 - ▶ Applications, Disk Drives, Volume Groups, File Systems, NFS File Systems, Tape Drives, IP Addresses, Replicated Resources
- AIX and Linux support
- Ease of use enhancements in HACMP 5.4
 - ▶ Configure or upgrade an HACMP cluster on a node without disrupting the target application
 - ▶ Resource dependency graphs

System p HACMP with Failover



02 - POWER Hardware 2008 v2.2

28

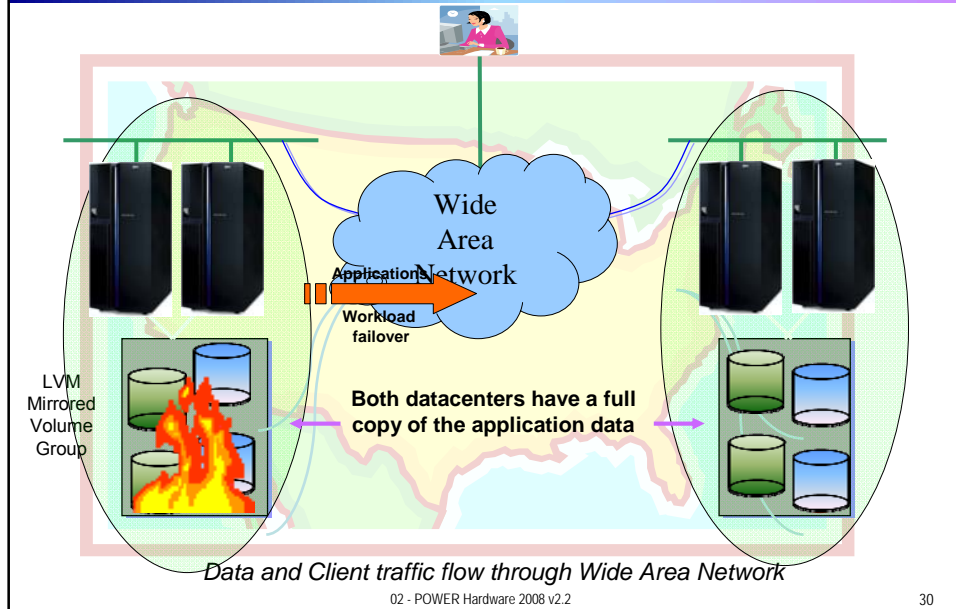
System p HACMP Extended Distance (XD) Provides Disaster Recovery

- Provides high availability and disaster recovery solutions
 - ▶ Recovers locally or moves entire application to backup site
- Integrates HACMP with unique data replication code to provide a fully automated solution
- Resources automatically failover to surviving node, no manual steps required
- Clusters
 - ▶ Campus wide – resources in multiple buildings, customer owns network
 - ▶ Metro wide – separate datacenters within same metro area, local network provider
 - ▶ Unlimited – datacenters in different states or countries, leased line networks
- Setup typically requires IBM services

02 - POWER Hardware 2008 v2.2

29

System p HACMP XD Can Provide Data Center Level Disaster Recovery



Compare the Competition

	IBM ★ System p	Sun Families	HP Families
Processor speed within an energy envelope	✓	Slower	Slower
Integrated energy efficiency functions	✓	Future	Future
Overall, price/performance	✓	Higher Cost	Higher Cost
Integrated hardware decimal floating point	✓	No	No
Integrated RAS functions	✓	Partial	Partial
Single unified processor architecture	✓	No (SPARC, UltraSPARC, SPARC64)	No (Itanium, PA-RISC)
Integrated virtualization assists	✓	Only UltraSPARC T1,T2	Only With Itanium (Limited Capability)
Has gone to Mars	✓	No	No

POWER Processor Has Gone To Mars!

NASA's New Phoenix Lander



The Phoenix Lander is powered by a processor based on IBM's POWER Architecture, similar to the one used in Sony's PlayStation 3.

