



Technical Forum & Executive Briefing

17 al 21
Octubre
2011

Imagine PODER Imagine CAPACIDAD

Power Systems Update & Directions



Tracy Smith
Executive I/T Specialist

A mandate for change is a mandate for smart.



The world is becoming instrumented
Sensors are being embedded everywhere: in cars, appliances, cameras, roads, pipelines...even in medicine and livestock.

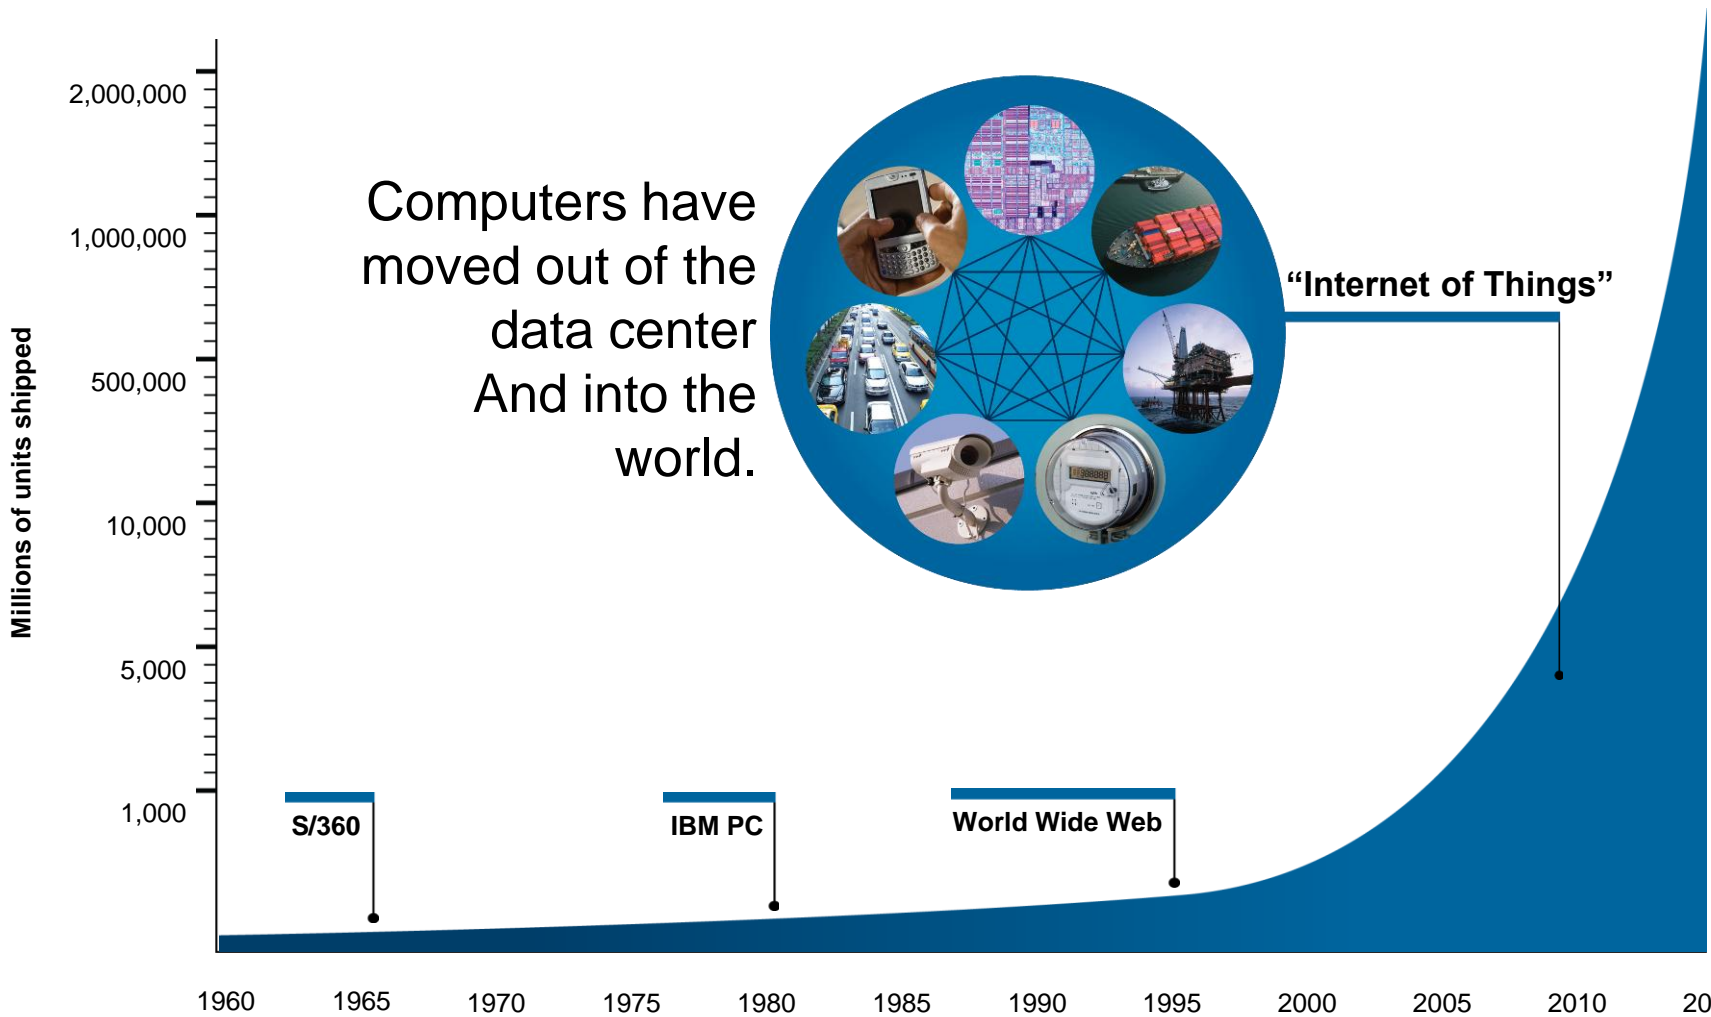


Our world is becoming interconnected
Think of a trillion connected and intelligent things, and the oceans of data they will produce.



All of those instrumented and interconnected things are becoming intelligent
They are being linked to powerful new back-end systems that can process all that data, and to advanced analytics capable of turning it into real insight, in real time.


Welcome to the decade of Smart.



Source: IDC, SSR and IBM Market Insights

Transformations to “smarter” solutions require smarter systems that:

- Scale quickly and efficiently
- Optimize workload performance
- Flexibly flow resources
- Avoid downtime
- Save energy
- Automate management tasks



Power Systems is helping deliver higher quality services – by impacting the way humans communicate with computers

- **IBM Watson** represents the latest in a long line of groundbreaking innovations from IBM
- Watson can understand the meaning and context of human language, and rapidly process information to find precise answers to complex questions
- What's next for Watson?
 - Project with Columbia University and Maryland School of Medicine to provide healthcare and life sciences diagnostic assistance
 - Research agreement with Nuance Communications to develop and apply Watson to healthcare
 - Other fields of investigation range from enterprise knowledge management to IT help desk




To learn more about IBM Watson: <http://www.ibm.com/watson/>






POWER7 Processor


Processor Technology Roadmap




POWER4™
180 nm




POWER5™
130 nm



POWER6™
65 nm



POWER7
45 nm



POWER8
22 nm

- Dual Core
- Chip Multi Processing
- Distributed Switch
- Shared L2
- Dynamic LPARs (32)

- Dual Core
- Enhanced Scaling
- SMT
- Distributed Switch +
- Core Parallelism +
- FP Performance +
- Memory bandwidth +
- Virtualization
- Transistors 276M

- Dual Core
- High Frequencies
- Virtualization +
- Memory Subsystem +
- Altivec
- Instruction Retry
- Dyn Energy Mgmt
- SMT +
- Protection Keys
- Transistors 790M

- Multi Core (up to 8)
- On-Chip eDRAM
- Power Optimized Cores
- Mem Subsystem ++
- SMT++
- Reliability +
- VSM & VSX (Altivec)
- Protection Keys+
- Transistors 1.2M

- Development Phase
- Core Running in Simulation

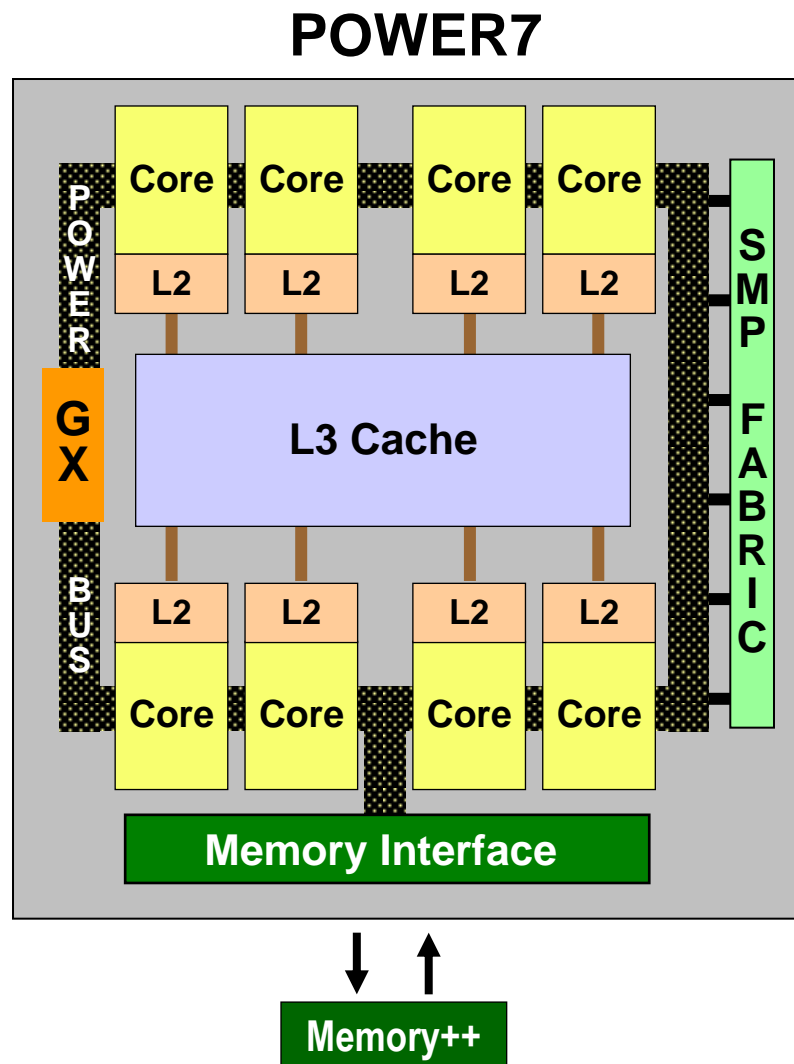
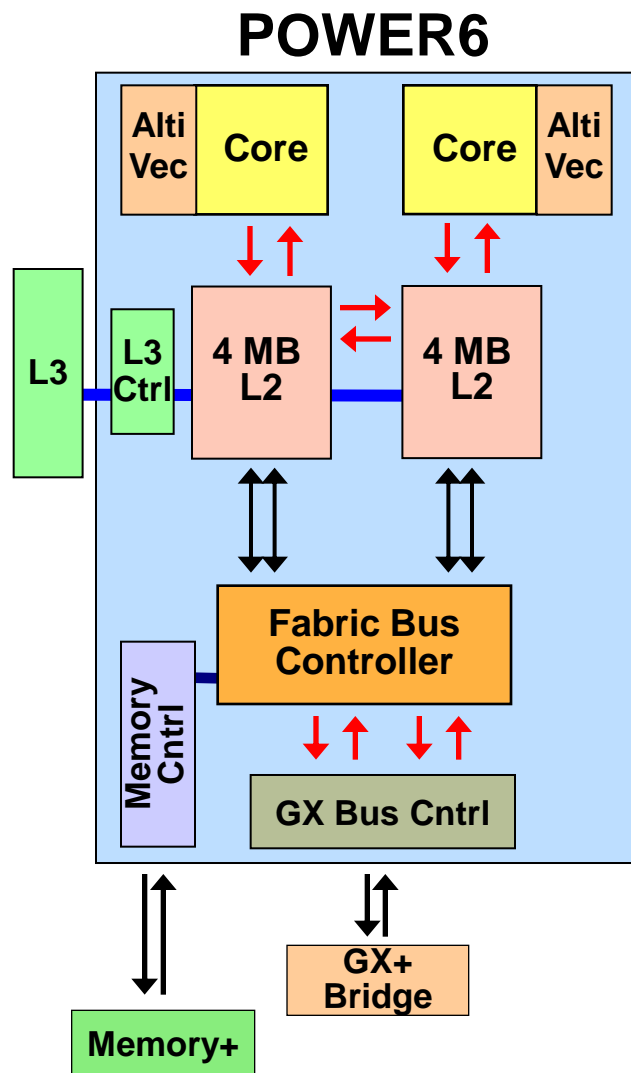
2001

2004

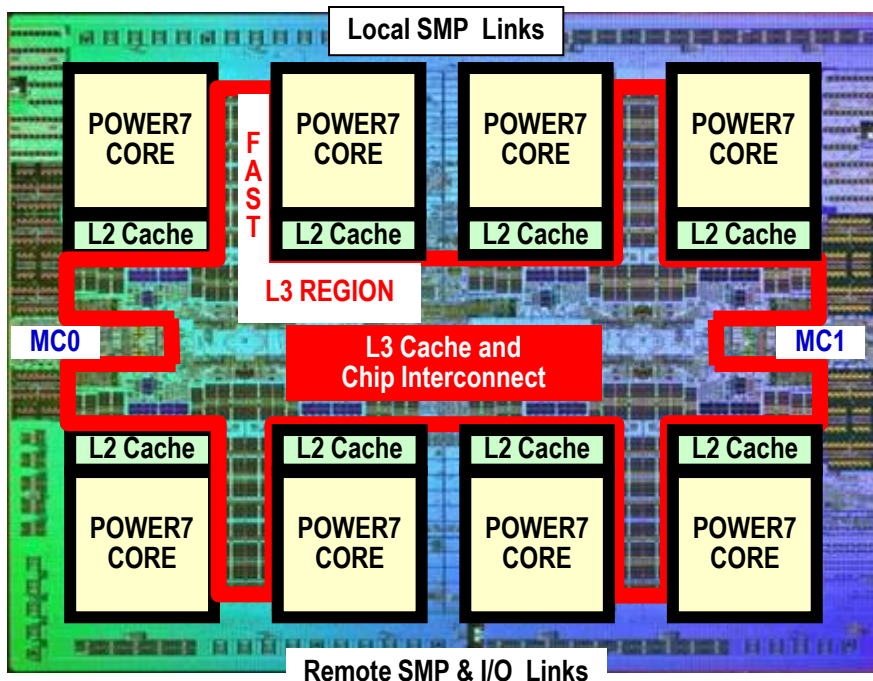
2007

2010

Transition from POWER6



POWER7 Processor Chip



**Binary Compatibility with
POWER6**

Cores : 4 / 6 / 8 core options

567mm² Technology:

- 45nm lithography, Cu, SOI, eDRAM

Transistors: 1.2 B

- Equivalent function of 2.7B
- eDRAM efficiency

Eight processor cores

- 12 execution units per core
- 4 Way SMT per core – up to 4 threads per core
- 32 Threads per chip
- L1: 32 KB I Cache / 32 KB D Cache
- L2: 256 KB per core
- L3: Shared 32MB on chip eDRAM

Dual DDR3 Memory Controllers


- 90 GB/s Memory bandwidth per chip

Scalability up to 32 Sockets


- 360 GB/s SMP bandwidth/chip
- 20,000 coherent operations in flight

Multi-threading Evolution


Single thread Out of Order




840-S80 Hardware multi-thread



POWER5 2 Way SMT



POWER7 4 Way SMT



No Thread Executing
 Thread 0 Executing
 Thread 1 Executing
 Thread 2 Executing
 Thread 3 Executing



The Power System Solution

Architecting the Power Solution



Hardware

- System Unit, Memory, I/O, Storage (integrated / external)



Virtualization

- Processors, Memory, Storage (Disk & Tape), I/O



Software

- Operating Systems, Database, Middleware, Applications (ISV and Homegrown)



Availability

- Operating System Based or Storage Based
 - ❖ Live Partition Mobility
- PowerHA, HA/DR solutions



Performance

- Monitoring, Managing and Trending



System Management

- Operations Management (Consoles, IBM Director, Tivoli)



Security and Compliance

- Tape/Storage/Network/Data Encryption



Energy

- Monitoring, Managing and Trending (Active Energy Manager)



POWER7 Systems

2011 Power Systems Portfolio

Select from the broadest system portfolio in the industry

- The highest performance, most scalable UNIX system ever
- Entry thru Enterprise Servers for IBM i, AIX and Linux



Express Servers

Power 720/740*

Power 710/730*



Power 750

Enterprise Servers

Power 795

Power 780*

Power 770*



PS Blades

i Editions Express for BladeCenter S



High Performance Computing



Power 755




Power 775

I/O Technology Transitions


RIO-2/HSL-2 to 12X (SDR and DDR)

- POWER6 supports RIO/HSL and 12X
- POWER7 supports 12X



PCI / PCI-X / PCI-X DDR to PCIe

- 2008 PCIe available in 520/550/570 System Units
- 2009 added 19" & 24" PCIe 12X DDR drawers
- 2011 PCIe Gen2 Systems



SCSI to SAS

- Disk Drives = SAS 3.5-inch moving to 2.5-inch SFF
- Solid State SFF SAS Drives
- Removable media SAS & SATA

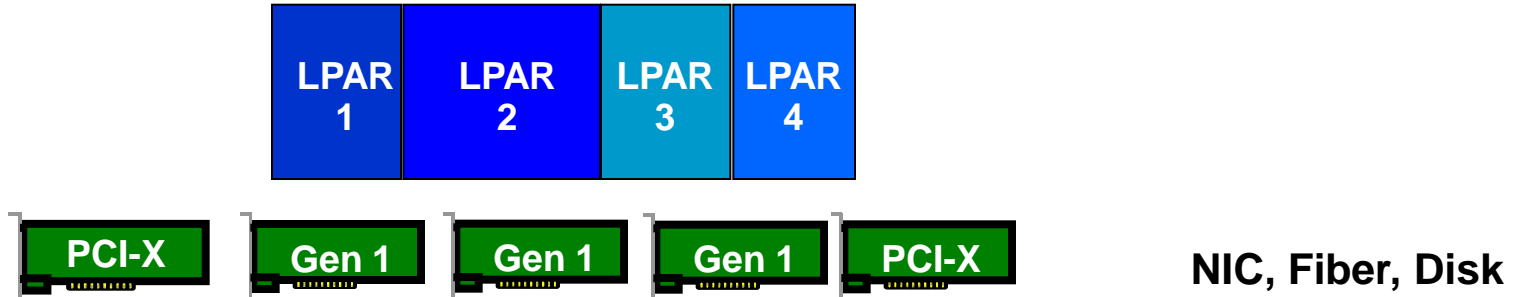


Virtualization of I/O


- POWER7 IOless only adapters
- NPIV Fiber and FCoE (CNA) adapters
- Shared Ethernet adapters

Why PCIe Gen2 Example: Do More With Less

With Gen 1, you might today have



And with Gen2, you might tomorrow have



Reduce adapters
Add more partitions
Improve/maintain I/O performance

Flexibility for more growth
Better price performance

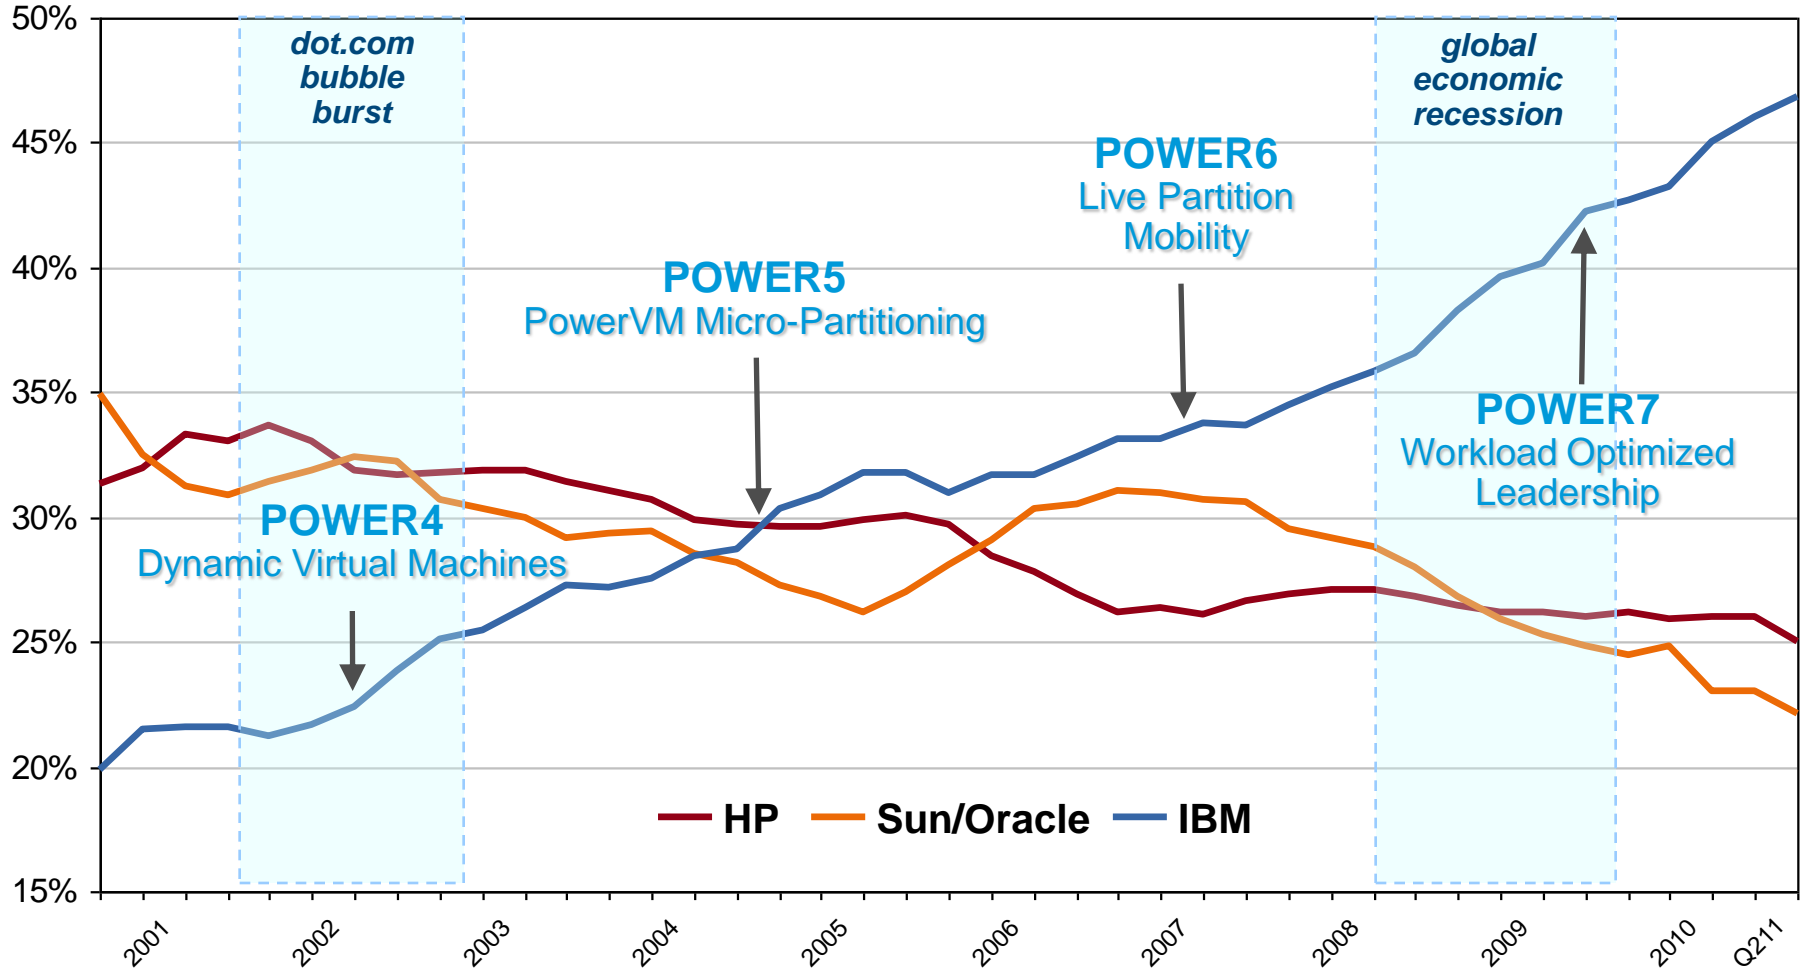


Software

IBM's Ten-Year March to UNIX Leadership

Exploiting the superior economics of Power drove the historic shift in the UNIX server market

UNIX Server Rolling Four Quarter Average Revenue Share



AIX Evolution – Over Twenty years of Progress

1986-1992 1994-1996 1997-1999 2001-2002 2004-2005 2007 2010

AIX/6000



AIX V2 & V3

Establishment in the market:

- RISC Support
- UNIX credibility
- Open Sys. Stds..
- Dynamic Kernel
- JFS and LVM
- SMIT

AIX V3.2.5

Maturity:

- Stability
- Quality

AIX V4.1/4.2

SMP Scalability:

- POWERPC spt.
- 4-8 way SMP
- Kernel Threads
- Client/Server pkg
- NFS V3
- CDE
- UNIX95 branded
- NIM
- > 2GB filesystems
- HACMP Clustering
- POSIX 1003.1, 1003.2, XPG4
- Runtime Linking
- Java 1.1.2

AIX V4.3

Higher levels of scalability:

- 24-way SMP
- 64-bit HW support
- 96 GB memory
- UNIX98 branded
- TCP/IP V6
- IPsec
- Web Sys. Mgr.
- LDAP Dir. Server.
- Workload Mgr
- Java JDT/JIT
- Direct I/O
- Alt. Disk Install
- Exp/Bonus CDs

AIX 5L V5.1/5.2

Flexible Resource Management:

- POWER4+ spt.
- Dynamic LPAR
- Dynamic CUoD
- New 64bit kernel
- 512GB mem
- JFS2
- 16 TB filesystems
- UNIX03 branded
- Concurrent I/O
- MultiPath I/O
- Flex LDAP Client
- XSSO PAM spt

AIX 5L V5.3

Advanced Virtualization:

- POWER5 support
- 64-way SMP
- SMT
- MicroPartitions™
- Virt I/O Server
- Partition Load Mgr
- NFS Version 4
- Adv. Accounting
- Scaleable VG
- JFS2 Shrink
- SUMA
- SW RAS features
- POSIX Realtime

AIX 6

Enterprise RAS:

- POWER6 support
- Workload Partitions
- Application Mobility
- Continuous Avail.
- Storage Keys
- Dynamic tracing
- Software FFDC
- Recovery Rtns
- Concurrent MX
- Trusted AIX
- RBAC
- Encrypting JFS2
- AIX Security Expert
- Director Console

AIX 7

Future of UNIX:

- 256 core/1024 tread scalability
- POWER7 Exploitation
- Domain based RBAC
- AIX Profile Manager
- WPAR enhancements
- AIX 5.2 in a WPAR
- PowerVM virtualized storage
- LVM SSD support
- Terabyte segment

Open Systems Workstations
Uni-processor

Distributed Client-Server
4-8 way SMP

Network Centric Computing
24-way SMP

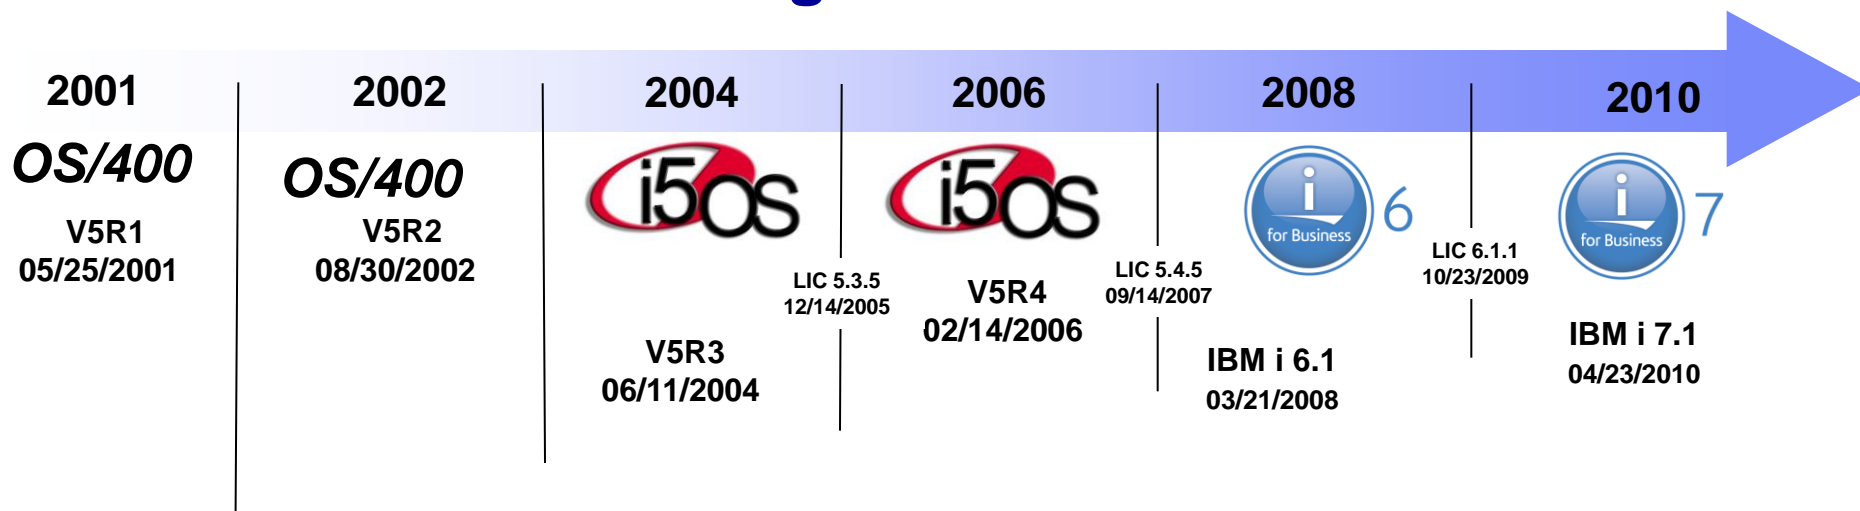
e-Business Computing
32-way SMP

On Demand Business
64/256-way SMT

New Enterprise Data Center
1024-way SMT4

Smarter Planet
1024-way SMT4

A Decade of OS/400 integrated innovations



i5/OS V5R3 May 2004	IBM i 5.4 February 2006	IBM i 6.1 IBM i 7.1 March 2008 April 2010
<ul style="list-style-type: none"> ▪ Virtual storage for AIX partitions ▪ Uncapped processor support for LPARs ▪ Cross Site Mirroring support for IASPs ▪ SQL Query Engine ▪ .NET Managed Providers ▪ iSeries Access for the Web ▪ DB2 function & performance enhancements ▪ Middleware currency ▪ ... and much more 	<ul style="list-style-type: none"> ▪ 32-bit IBM JVM ▪ Webfacing Deployment environment ▪ Web Services support for RPG ▪ TCP/IP V6 ▪ Integration of BladeCenter via iSCSI ▪ DB2 On Demand Performance Center ▪ Enterprise Extender support for SNA ▪ Intrusion detection ▪ Administrative System Domain for HA ▪ Virtual tape support for backups ▪ DB2 function & performance enhancements ▪ Middleware currency ▪ ... and much more 	<ul style="list-style-type: none"> ▪ Support for BladeCenter S and H ▪ Breakthrough performance with DS8000 ▪ Virtual storage for IBM i partitions with IBM i and PowerVM VIOS ▪ PowerHA disk clustering ▪ Improved Java and WebSphere performance ▪ 64-bit IBM JVM ▪ Systems Director Navigator for i ▪ Performance Investigator ▪ Encryption for data on disk and tape backup ▪ Intrusion prevention ▪ VMware VMotion support for iSCSI ▪ OmniFind Text Search ▪ DB2 function & performance enhancements ▪ Middleware currency

What about Linux on POWER?

Customers Seeking Linux and Open Source Solutions

Linux on Power has a wide range of open source stack options






Virtualization



Virtualization for retired Minnesotans



© 1997 P. C. Vey from The Cartoon Bank. All rights reserved.

If you consider it a sport to gather your food by drilling through 18 inches of ice and sitting there all day hoping that the food will swim by, you might live in Minnesota.

PowerVM: Virtualization without limits

Reduces IT infrastructure costs

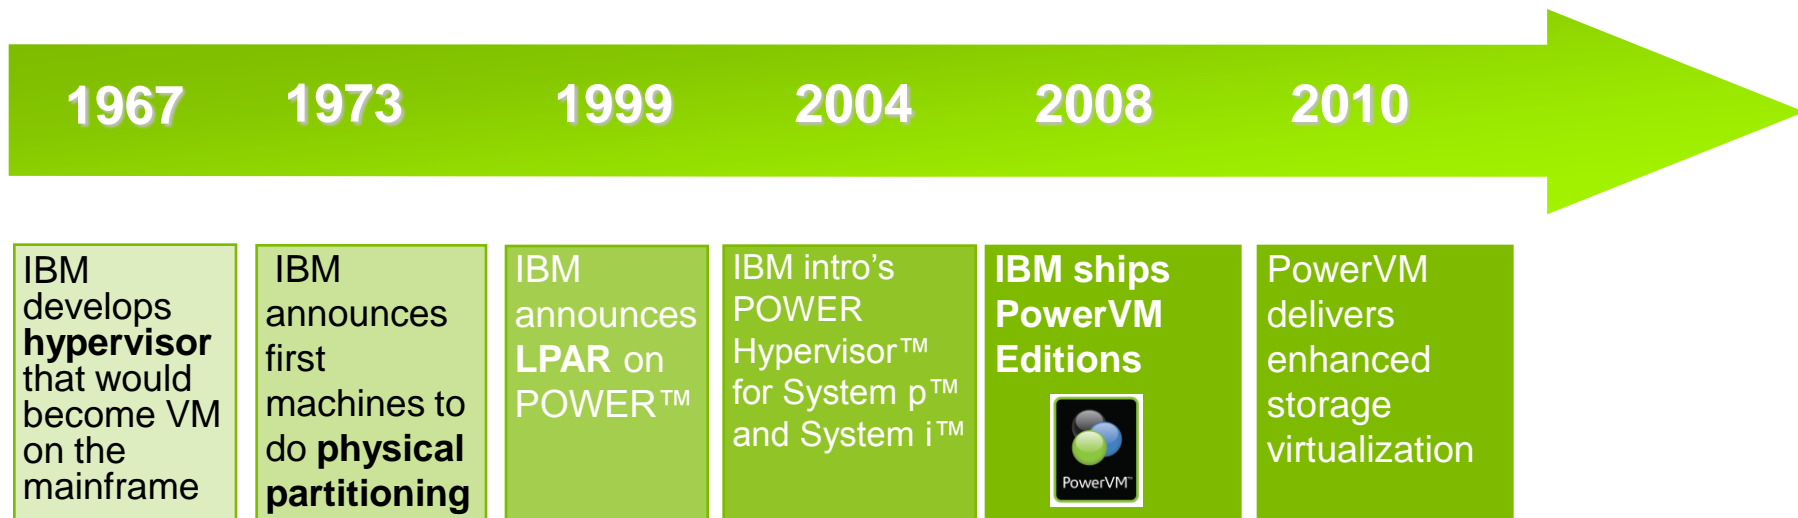
Improves service levels

Manages risk



PowerVM builds on IBM's virtualization leadership

A 40-year track record in virtualization innovation continues with PowerVM™



The Power Systems team has integrated PowerVM world-class virtualization into every server – based on best practices gained from IBMers who created the ‘gold standard’ of mainframe virtualization.





Availability

Business Resiliency Requirement Variables

Types of outages to be addressed

- **Unplanned** (for example, a hardware failure)
- **Planned** (for example, a software upgrade)
- **Backups** (for example, create copy of disk for on-line save to tape)
- **Disasters** (for example, site loss, power grid outage, etc).

Recovery Objectives

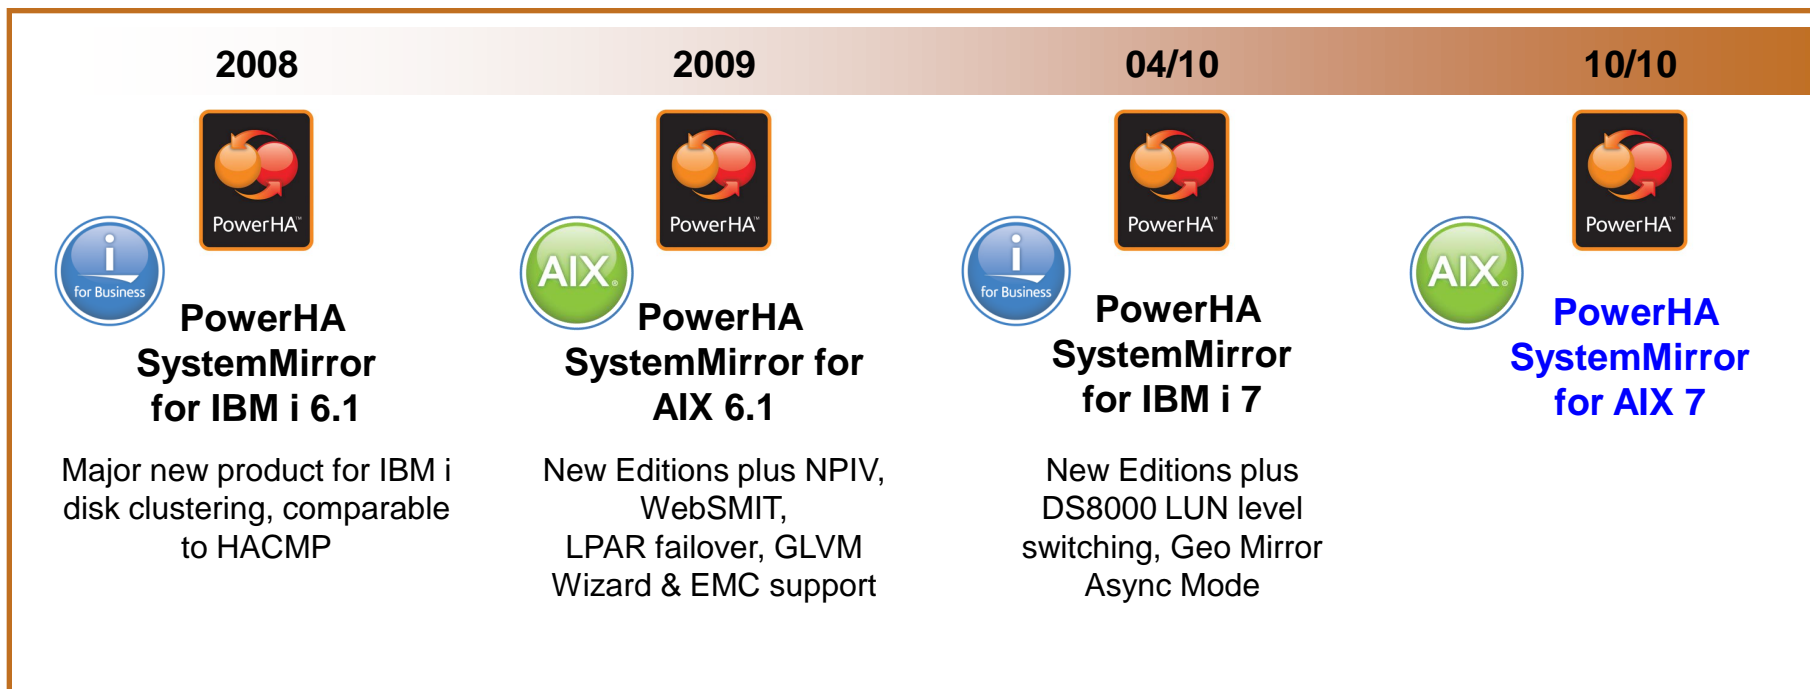
- Recovery Time Objective – **RTO**: Time to recovery from an outage
- Recovery Point Objective – **RPO**: The amount of tolerable data loss

Recovery tiers by application (for example)

- **Tier-0**: minimal downtime needed
- **Tier-1**: 4 hours downtime tolerable
- **Tier-2**: Extended downtime of a day or more tolerable
- etc.

Resiliency without Downtime

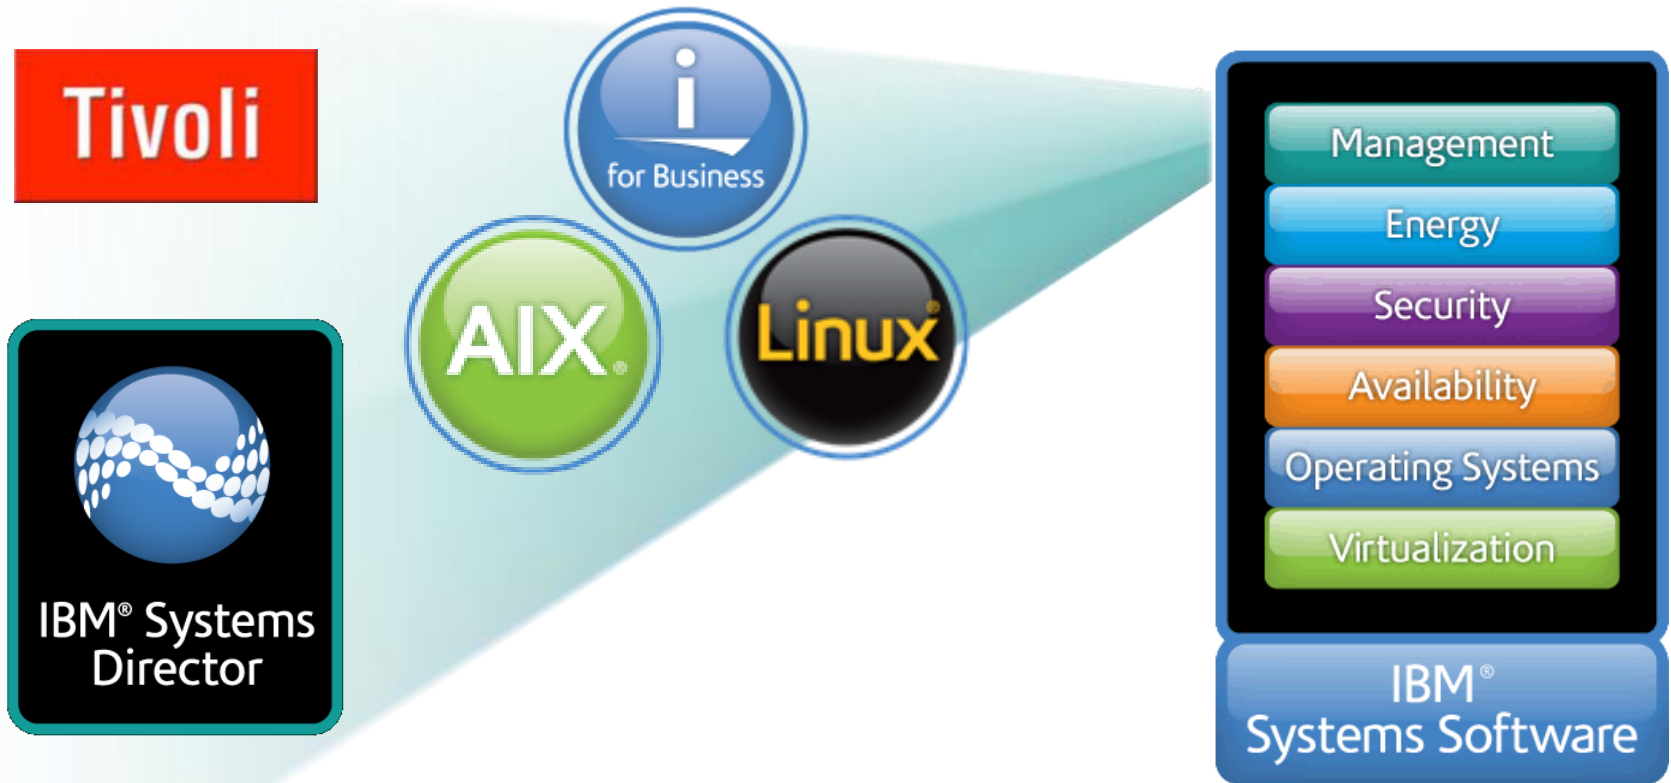
- **PowerHA SystemMirror**
 - Active/standby high availability / disaster recovery clustering
 - Strategy to differentiate via deep integration with AIX and IBM i
- **PowerHA pureScale**
 - Active/active cluster management interconnect
 - Included with DB2 pureScale






System Management

Management with Automation



Systems Management Hierarchy

- Base management – no Director



Systems Management Hierarchy

Optional

Upward Integration & Service Management Software

Service Management

End-to-End Management

- Base management – no Director
- Add optional Director & Tivoli

IBM Systems Director


VMControl

Storage Control

AEM

Network Control

Optional



POWER servers

Power is performance redefined

BladeCenter

© 2011 IBM Corporation

Systems Management Hierarchy

Optional

Upward Integration & Service Management Software

Service Management

End-to-End Management

- Base management – no Director
- Add optional Director & Tivoli
- Replace HMC with SDMC
- Function can run in Blade chassis
- SDMC also manages Power Blades
- Systems Director hierarchical management (Advanced Managers not available in first release)

IBM Systems Director


VMControl

Storage Control

AEM

Network Control

Optional




POWER servers

Power is performance redefined

BladeCenter

© 2011 IBM Corporation

Managing client transitions from HMC to SDMC



The SDMC is the next-generation management console for Power Systems, launched in 2011 and phased in as an HMC replacement over several years

An SDMC can be used alongside the HMC during client trials, until such a time as the client is ready to switch over permanently to SDMC management

During the transition phase, the HMC and IVM will be kept current with new Power Systems models, but will not incorporate future advanced management capabilities

IBM can provide services and support to help ensure seamless client transitions



Energy Management

EnergyScale™

EnergyScale is IBM Trademark. It consists of a built-in Thermal Power Management Device (TPMD) card and Power Executive software.

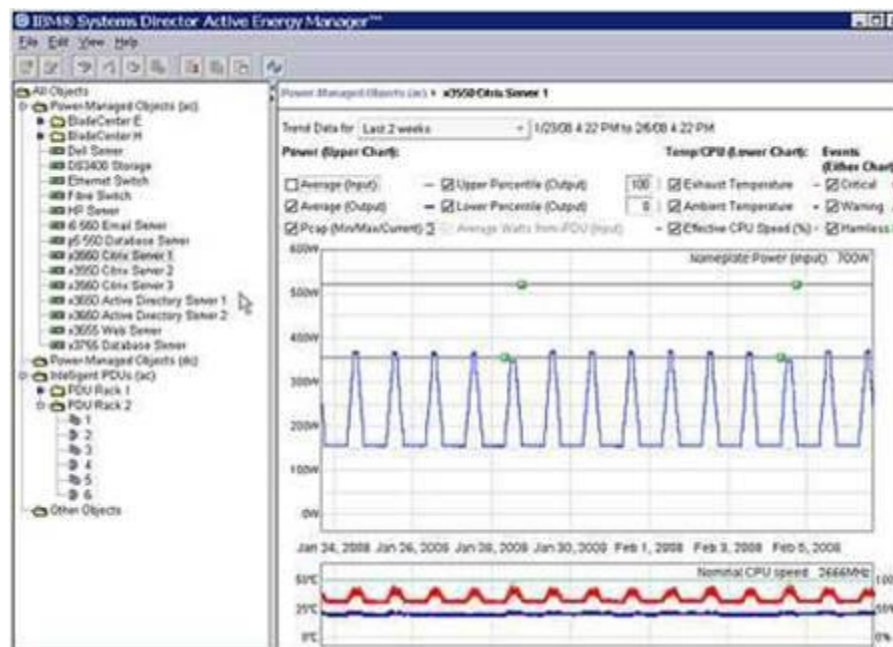


IBM Systems Director is also required to manage Energy-Scale functions.

EnergyScale is used to dynamically optimize the processor performance versus processor power and system workload.

IBM Systems Director is also required to manage AEM functions and supports the following functions:

- Power Trending
- Thermal Reporting
- Static Energy Saver Mode
- Dynamic Energy Saver Mode
- Energy Capping
- Soft Energy Capping
- Processor Nap
- Energy Optimized Fan Control
- Altitude Input
- Processor Folding



Active Energy Manager Trend Data

Energy trends in upper chart

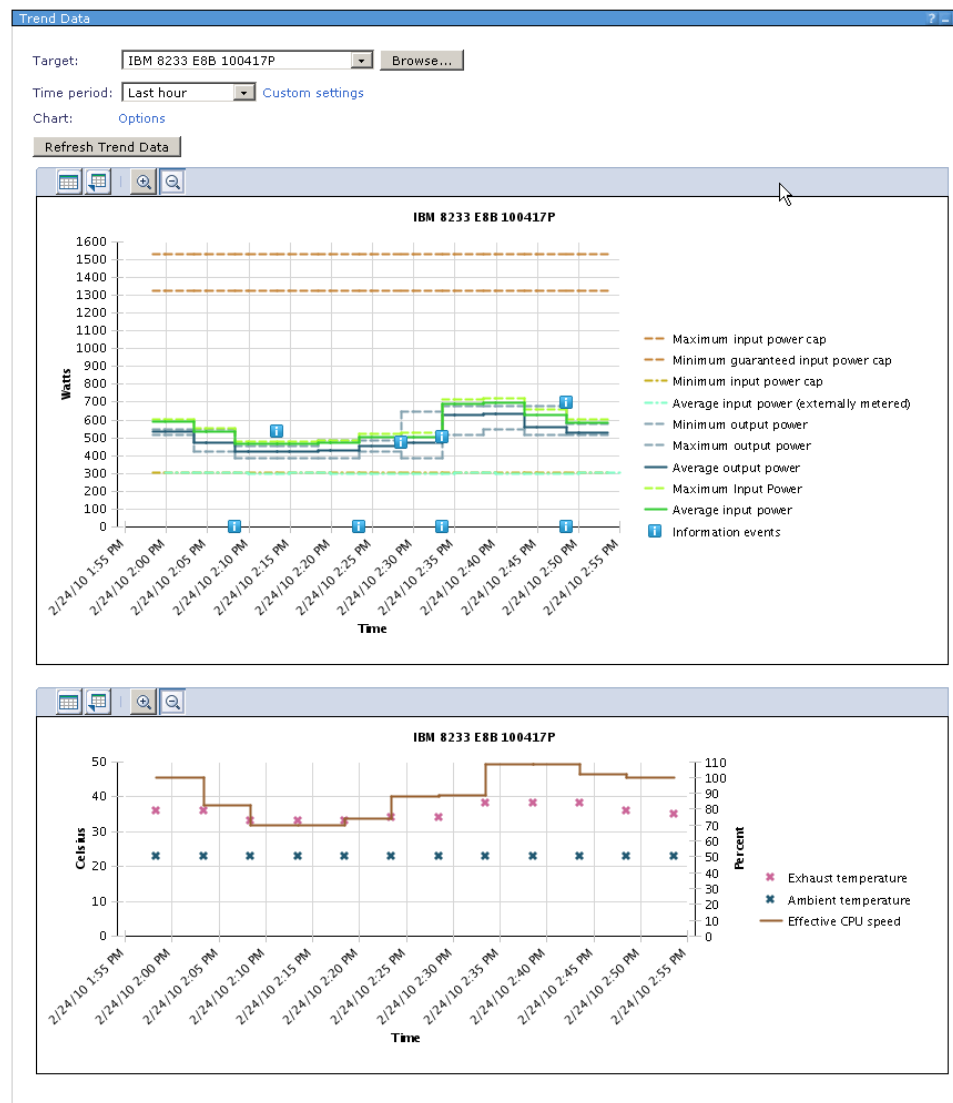
- Data are averages over intervals
- Input & output power
- Maximums, averages & minimums
- Power caps
- Energy events at points in time

Temperature data in lower chart

- Data are point-in-time values
- Plotted against left axis

Effective CPU speed in lower chart

- Data are averages over intervals
- Plotted against left axis
- May exceed 100%



POWER7 on Cloud Delivers Time to Value

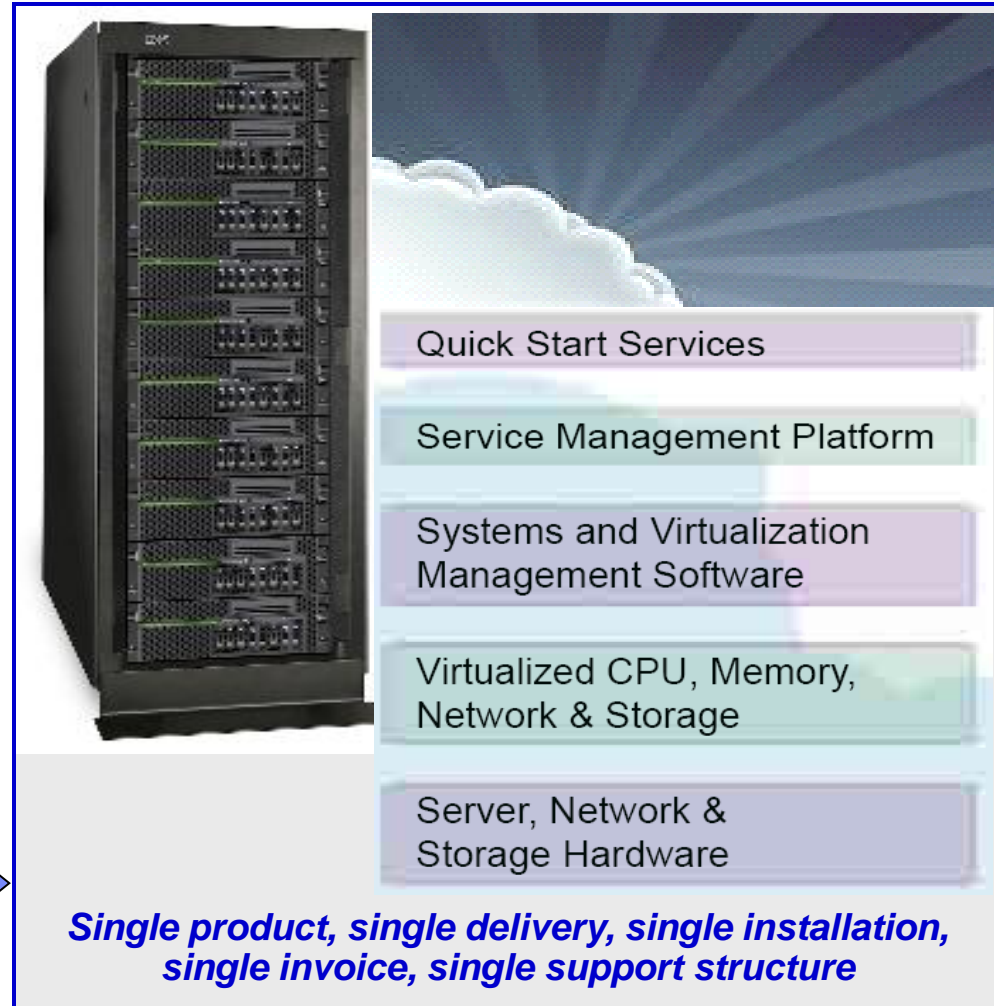
An integrated service management platform with network, servers, storage, quickstart services that enables the fastest Private Cloud Deployment Today

Customer Benefits

- ✓ **Improved time to value**- Quickly deliver a private cloud using a preloaded and integrated system
- ✓ **Improved innovation**- Dramatically improve business value and IT's effect on time-to-market by delivering services faster
- ✓ **Decrease capital expenses** – Maximize capital usage and reduce added capital expense.
- ✓ **Reduce complexity and risk**- With automation and standardization the human error factor is minimized.
- ✓ **Scales to the enterprise** – Able to scale and manage additional Platforms and Workloads

PowerVM &
VMControl

IBM Service
Delivery Manager



POWER7 Key Feature Summary

Power Systems offers balanced systems designs that automatically optimize workload performance and capacity at either a system or VM level

- ✓ **MaxCore or 6 Core** for incredible parallelization and high capacity
- ✓ **TurboCore™** on the 780 & 795 for max per core performance for databases
- ✓ **Processor Overclocking** for increased performance when workload is heavy
- ✓ **Intelligent Threads** utilize more threads when workloads benefit
- ✓ **Intelligent Cache** technology optimizes cache utilization flowing it from core to core
- ✓ **Intelligent Energy Optimization** maximizes performance as thermal conditions allow
- ✓ **Active Memory Expansion™** provides more memory for SAP and other Applications
- ✓ **Binary Compatibility** with POWER6 and POWER7 modes
- ✓ **Partition and Workload Mobility** from POWER6 to POWER7
- ✓ **Power Flex Offering** for high availability at a reduced cost
- ✓ **Low Cost Capacity on Demand** for Rapid Resource Provisioning
- ✓ **TPMD**: Granular Thermal and Energy Management
- ✓ **Balanced System Design**: Complete System Performance and Linear Scaling


Workload-Optimizing Features make POWER7 processors
#1 in both Transaction and Throughput Computing

**Unique
POWER7
Features**



Gracias from Power Systems

Smarter systems for a Smarter Planet.



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 5L, 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, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel 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, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, 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 full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

AltiVec is a trademark of Freescale Semiconductor, Inc.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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.

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

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.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

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

Revised December 2, 2010

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.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, 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 for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

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
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
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 2, 2010

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.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, 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 for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

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	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm
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. http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mssoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	http://hmmer.janelia.org/ http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised December 2, 2010

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

Revised April 2, 2007