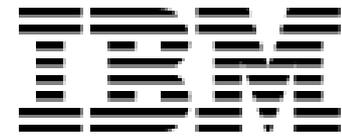




POWER8: The first OpenPOWER processor

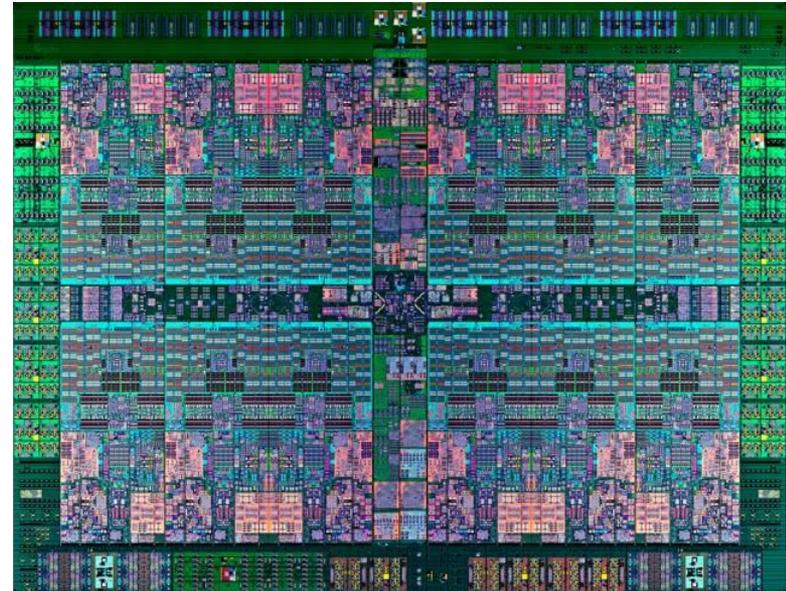
Dr. Michael Gschwind
Senior Technical Staff Member & Senior Manager
IBM Power Systems



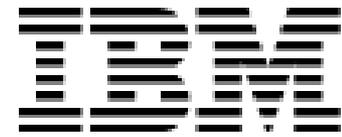


Why Power and Why Now?

- Power is optimized for server workloads
- Power8 was optimized to simplify application porting
- Power8 includes CAPI, the Coherent Accelerator Processor Interconnect
 - Building on a long history of IBM workload acceleration



POWER8 Processor



Cores

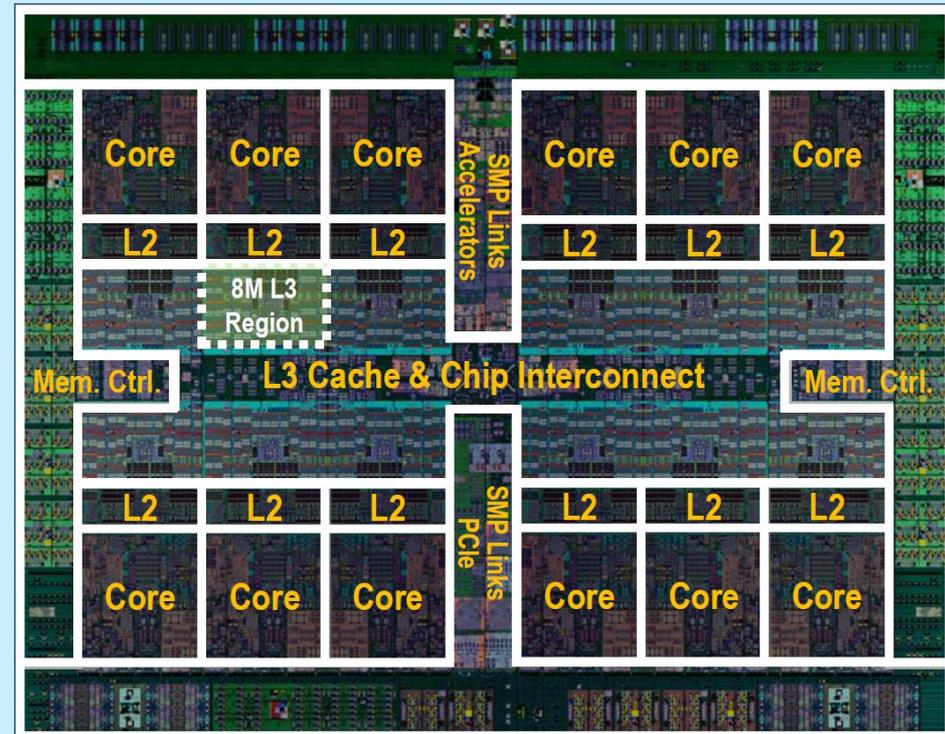
- 12 cores (SMT8) → 96 threads per chip
- 2X internal data flows/queues
- 64K data cache, 32K instruction cache

Caches

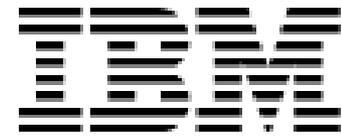
- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)

Accelerators

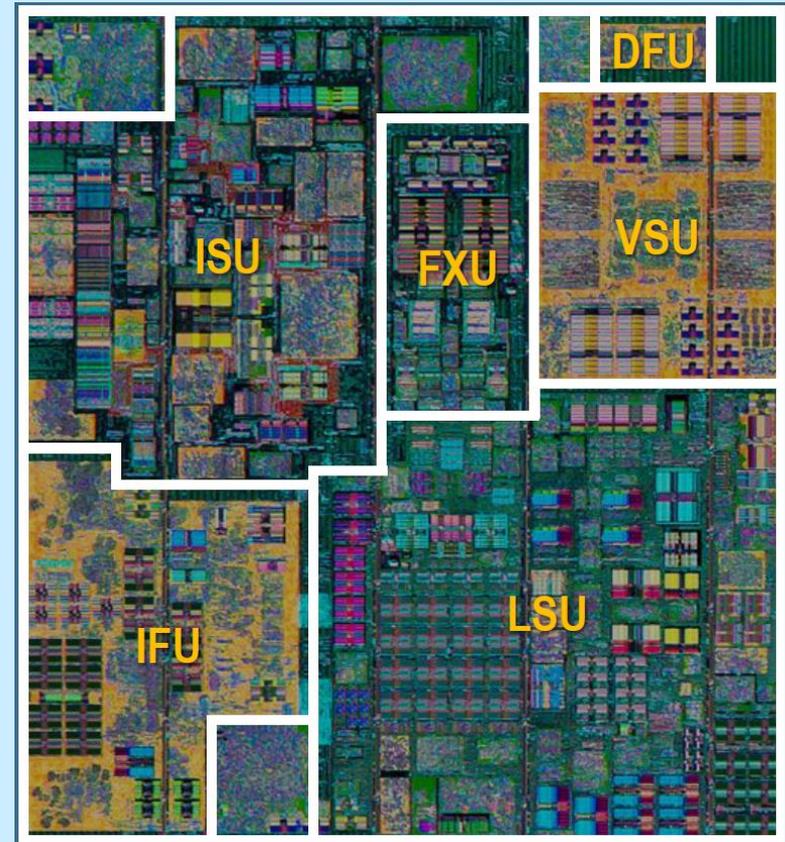
- Crypto & memory expansion
- Transactional Memory
- VMM assist
- Data Move / VM Mobility
- Coherent Accelerator Processor Interface (CAPI)



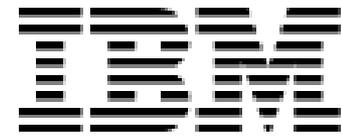
POWER8 Core



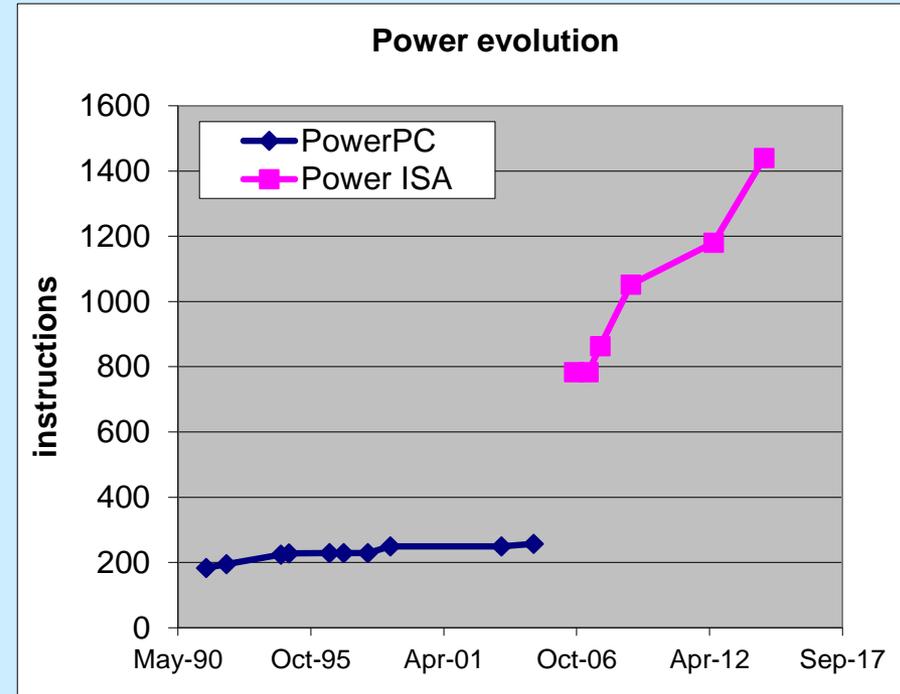
- Up to eight hardware threads per core (SMT8)
- 8 dispatch
- 10 issue
- 16 execution pipes:
 - 2 FXU, 2 LSU, 2 LU, 4 FPU, 2 VMX, 1 Crypto, 1 DFU, 1 CR, 1 BR
- Larger Issue queues (4 x 16-entry)
- Larger global completion table
- Larger Load/Store reorder queue
- Improved branch prediction
- Improved unaligned storage access
- Improved data prefetch

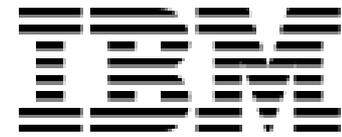


POWER8 Architecture



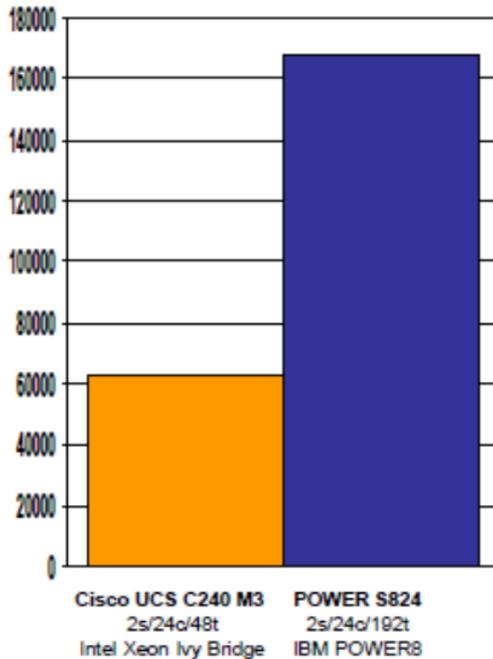
- High-performance LE support
 - Foundation for a new ecosystem
- Organic application growth
 - Instruction Fusion
- Workload optimized
 - Power SIMD Extensions
 - Coherent Accelerator Processor Interface
- Parallel programming productivity
 - Transactional Memory
- Cloud ready
 - Cloud Mode
 - Reduced context switch overhead
 - Improved Interrupt performance
- Dynamic System Optimization
 - Performance Monitoring
 - Dynamic Code Optimization



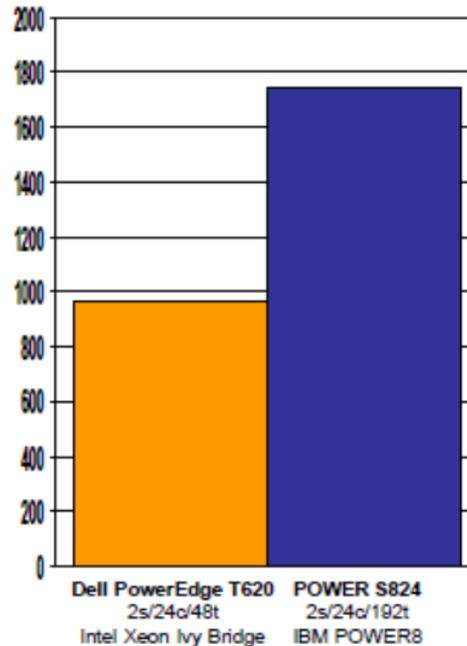


Up to 2.7x performance across key workloads vs. other 24-core Scale-Out Systems

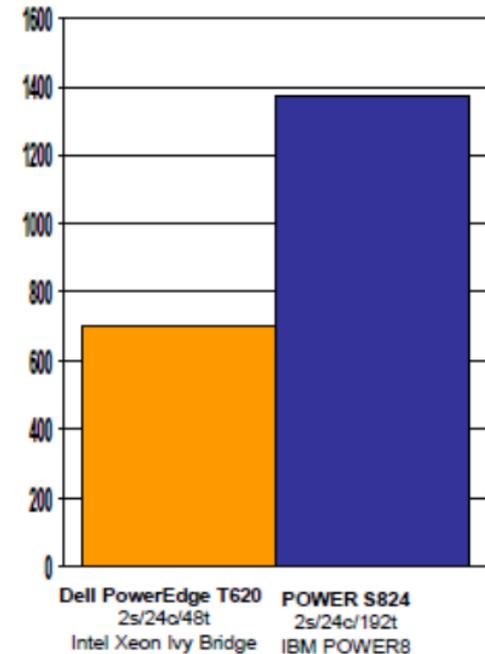
Java – SPECjbb2013 (Max-jOPS)
2.7x Performance



SPECint_rate2006
1.8x Performance



SPECfp_rate2006
2x Performance



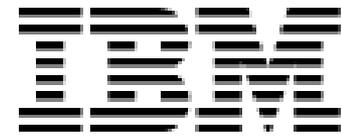
Results are based on best published results on Xeon E5-2697 v2 from the top 5 Intel system vendors.
SPECjbb2013 results are valid as of 7/7/2014. For more information go to <http://www.specbench.org/jbb2013/results>
SPECcpu2006 results are submitted as of 4/22/2014. For more information go to <http://www.specbench.org/cpu2006/results/>

Source: Hot Chips 26

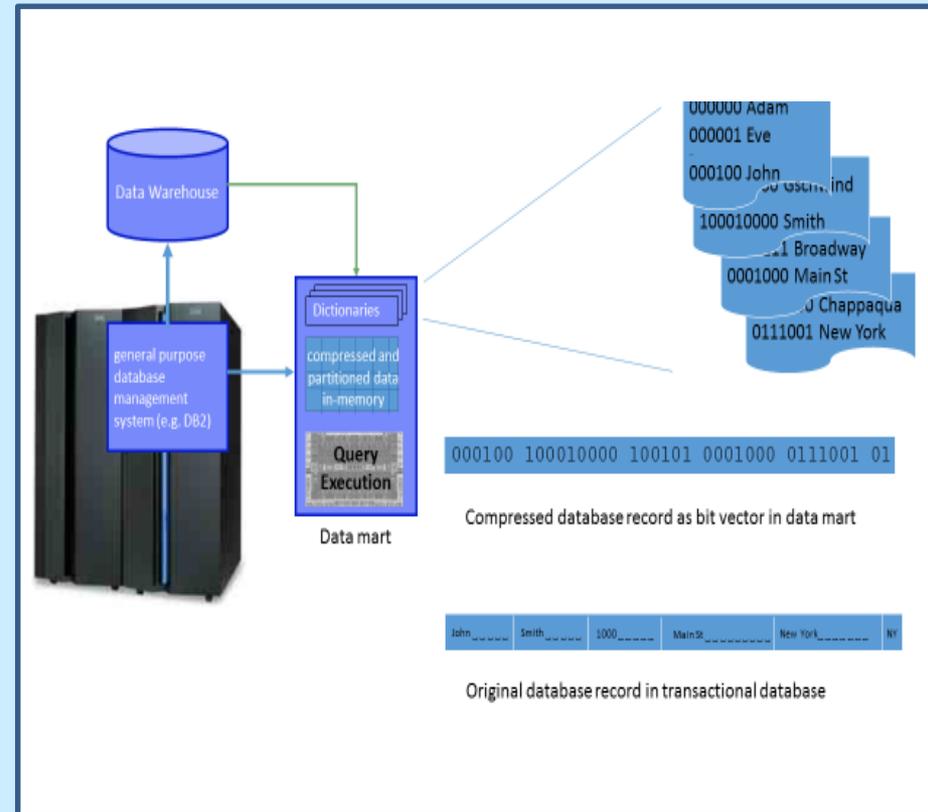


Join the conversation at #OpenPOWERSummit

Acceleration with Vector-Scalar Execution

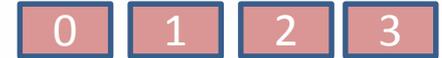


- Expand VSX beyond HPC
- Memory to memory primitives
 - Exploit system bandwidth in apps & system
- Coding support
 - AES, CRC, RAID,...
- Business analytics
 - Numeric analytics algorithms
 - Structured & Unstructured Text Analytics
- Database acceleration
 - Accelerate in-memory columnar DB



POWER8: expanding choice for data formats

- Big-Endian: Most Significant Byte First



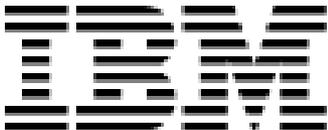
- Traditional enterprise server data format
- The standard internet data format (“network byte order”)
- Process five decades of database content
- Protect investment into data storage and curation

- Little-Endian: Least Significant Byte First



- Traditional small system format (PC, mobile)
- Participate in rich standard ecosystem
- Simplify application portability for those seeking to upgrade to Power
- High-volume hardware components often little-endian

Workload-optimized acceleration with coherent accelerators



- Attached accelerators
 - Accelerate workloads with functions that are not a good fit for traditional general purpose CPU
- Coherent integration in system architecture
 - Data sharing
 - Programming
 - Performance

Typical I/O Model Flow



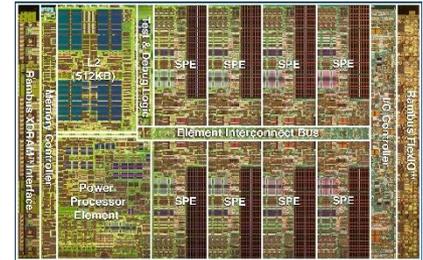
Flow with a Coherent Model



Join the conversation at #OpenPOWERSummit

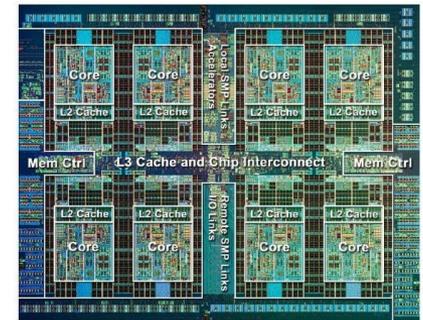
Workload-optimized acceleration

- On-chip integrated accelerators (SoC design)
 - Compute accelerator (Cell BE)
 - Compression (P7+)
 - Encryption (P7+)
 - Random number generation (P7+)
 - ...



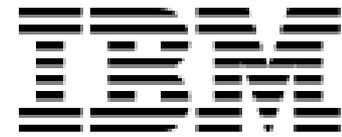
Cell BE

- SoC design offers highest integration, but...
 - New chip design to introduce new accelerator
 - Long time to market
 - Requires very high volumes

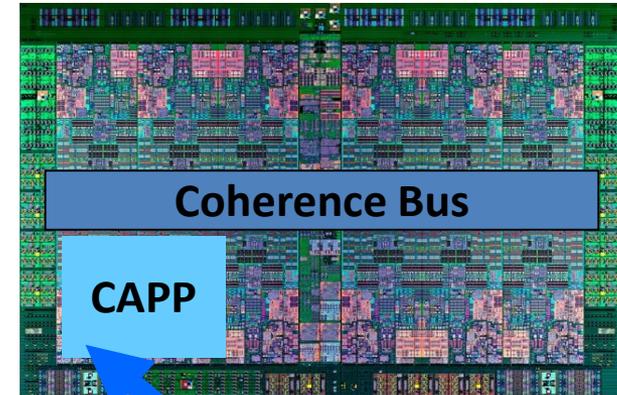


POWER7+

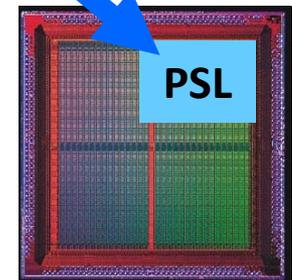
CAPI: Coherent Accelerator Processor Interface



- Integrate accelerators into system architecture with standardized, coherent protocol
 - CAPI functional protocol
 - PCIe signaling protocol
- Modular interface enables third parties to provide high value-add components
- Create workload-optimized innovative solutions
 - Faster time to market
 - Lower bar to entry
 - Variety of implementation options
 - FPGAs, ASICs



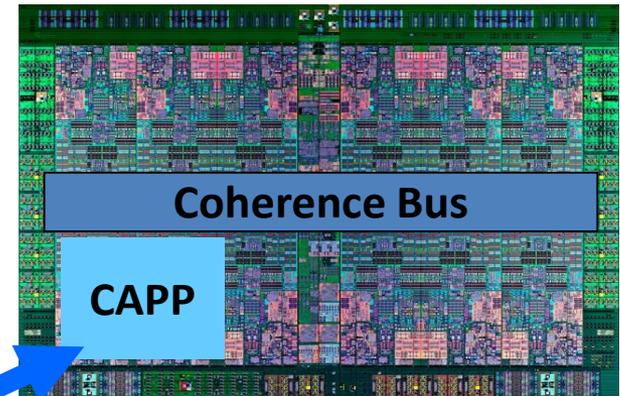
POWER8



CAPI accelerator programming

Virtual Addressing

- Accelerator can work with same virtual memory addresses that the processors use
- CAPI shares page tables and provides address translation
- Pointers de-referenced same as the host application
- Peer-to-peer programming between CPU and accelerator with in-memory data sharing



Hardware Managed Coherence

- No need for memory pinning
- Data fetched by accelerator based on accelerator application flow
- Accelerator participates in locks
- Low latency communication

Summary

- POWER8 introduce traditional microarchitecture improvements
 - More Instruction level parallelism
 - More Data level parallelism
 - Larger caches and queues
 - ⇒ Best-in-class systems performance for scale-out and scale-up
- New foundational capabilities for OpenPOWER
 - High-performance little- and big-endian
 - Improved integration of industry standard ecosystem hardware
 - Improved application portability to OpenPOWER
 - CAPI enables system architects to extend processor capabilities
 - Create a broad range of workload optimized offerings

Implementation / HPC / Research



Software



System / Integration



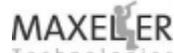
I/O / Storage / Acceleration

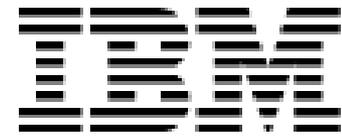


Boards / Systems



Chip / SOC





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.

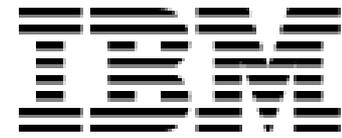
IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and 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.

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.





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 Watson, 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, PowerLinux, 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, POWER7+, POWER8, Systems, 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.

PowerLinux™ uses the registered trademark Linux® pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the Linux® mark on a world-wide basis.

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, SPECjobb, 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.

