



Welcome to:


Virtualizing the Linux Environment on POWER Systems



Ron Gordon, IBM




Virtualization Delivers significant business and IT value




Utilization

- Reduce **total cost of ownership** and make better use of IT assets by significantly **improving server utilization**



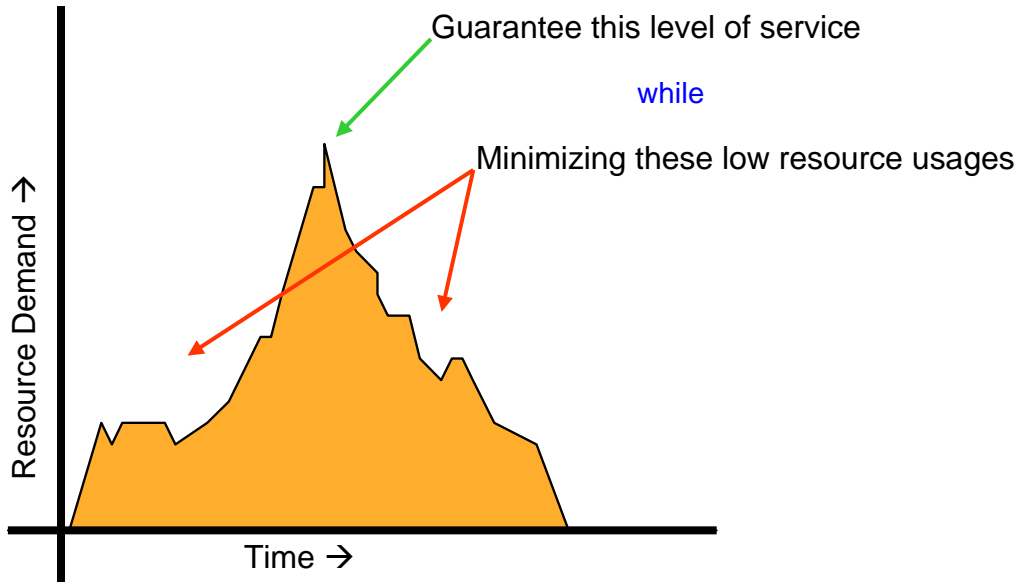
- Improve **business responsiveness** and operational speed by **dynamically re-allocating resources** to applications as needed -- to better match changing business cycles or handle unexpected surge in demand



- **Simplify the management of IT infrastructure** by making workloads **independent of hardware** resources, thereby enabling customers to make **business-driven policies** to deliver resources based on time, cost and service-level requirements

What is: Business Flexibility and Resource Utilization?

Answer: Ability to maintain workload responsiveness at the lowest system cost.



What is Advanced POWER Virtualization ?

Logical Partitions

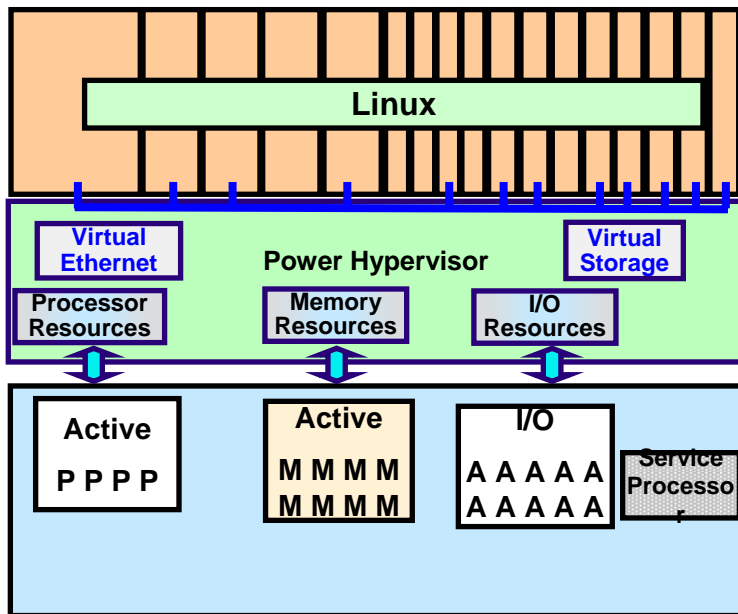
- Up to 254 active partitions
- Linux & AIX
- Isolated and secure

Virtualization

- Network / Storage
- Processors (1/10th Intervals)
- Memory
- I/O

Physical Resources

- Processors
 - Dedicated
 - Shared Pool
- Memory
- I/O
 - SCSI + Disk
 - Ethernet
- Service Processor / HMC
- IVM for web interface

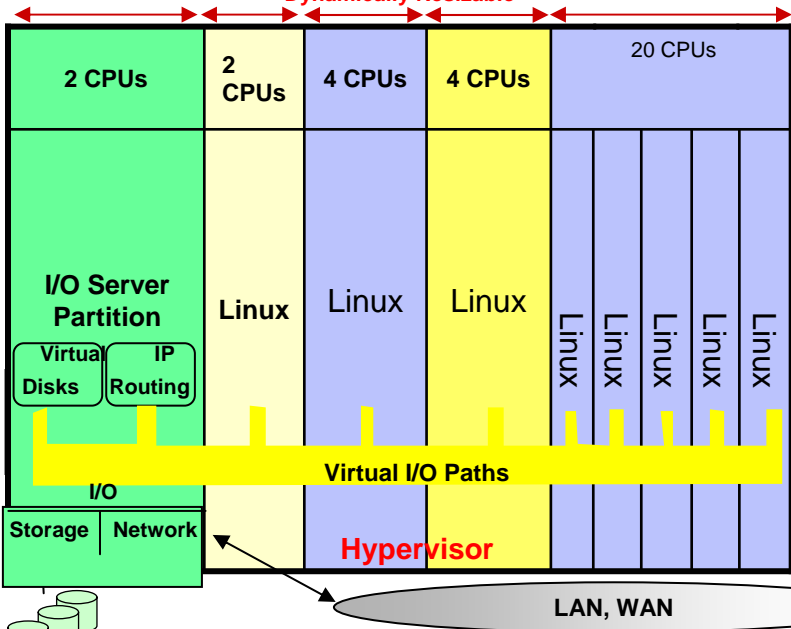


Net Effect: Business Flexibility and High system utilization at low cost

Advanced Open Power Virtualization Overview of function

Greatly improved physical resource utilization through virtualization of processors, memory, and I/O resources

Dynamically Resizable



Increased number of LPARs

AIX support for 254 partitions

Shared processor LPARs

A single processor might be shared by up to 10 partitions

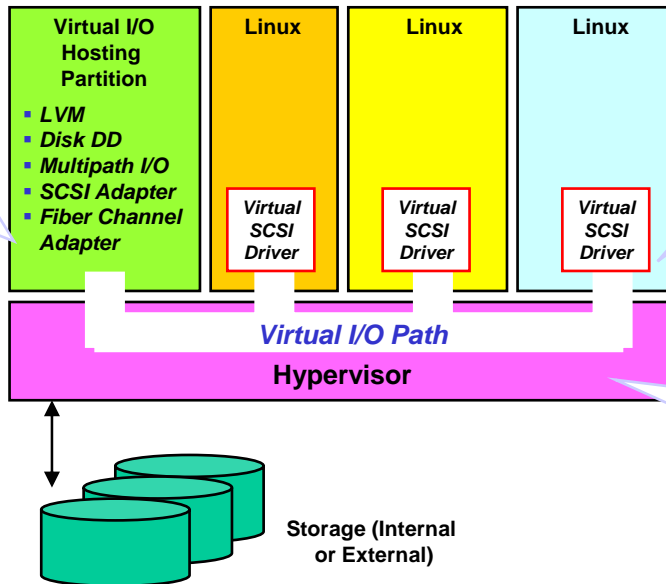
Virtual inter-partition Ethernet

LPARs can communicate without having to use physical I/O adapters

Virtual I/O Disks

Client partitions can use virtual disks hosted by another partition

Shared Fiber Channel Adapters and Virtualized Disks



I/O Hosting Partition provides Virtualized disks using Logical Volume partitioning and direct LUN access

Each OS Image Implements Virtual SCSI Adapters

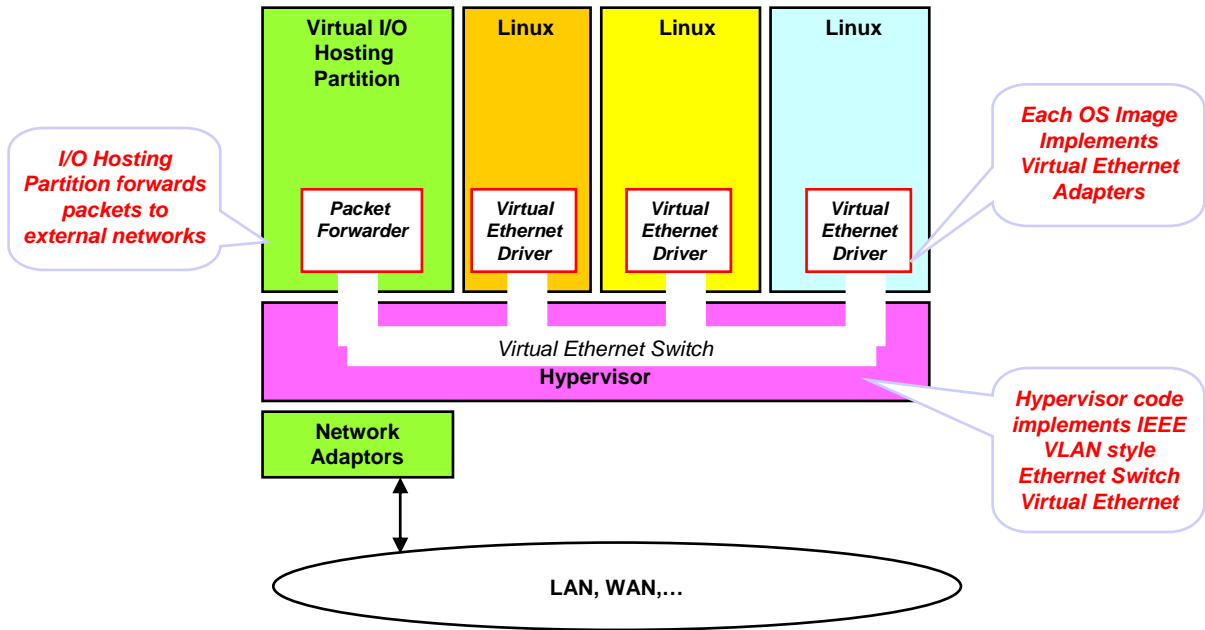
Hypervisor code implements Virtual I/O Transport

vSCSI Redundancy available:

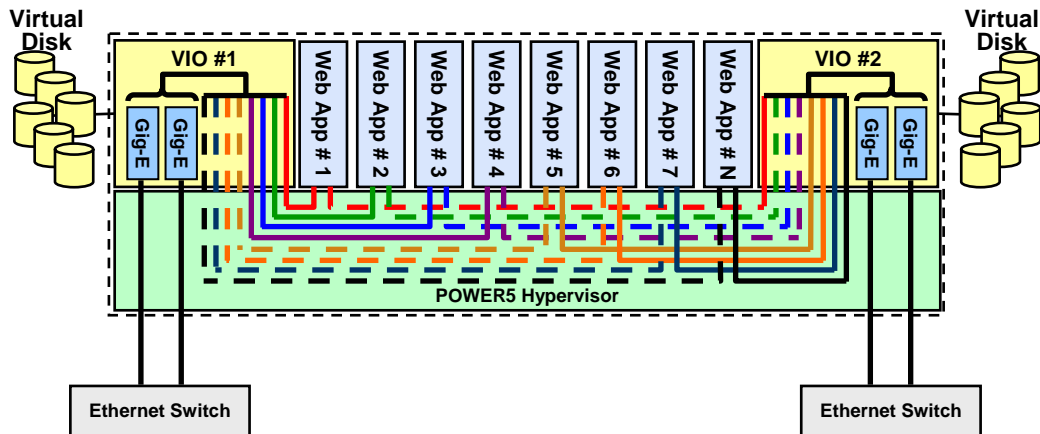
- *Multi-pathing (SDD, PowerPath, HDLM) at the I/O Server*
- *Multiple VIO Server using LVM Redundancy or MPIO at the Client*

Virtualizing I/O:

Virtual Ethernet (VLAN) and Shared Ethernet adapters



Web Application Server Example

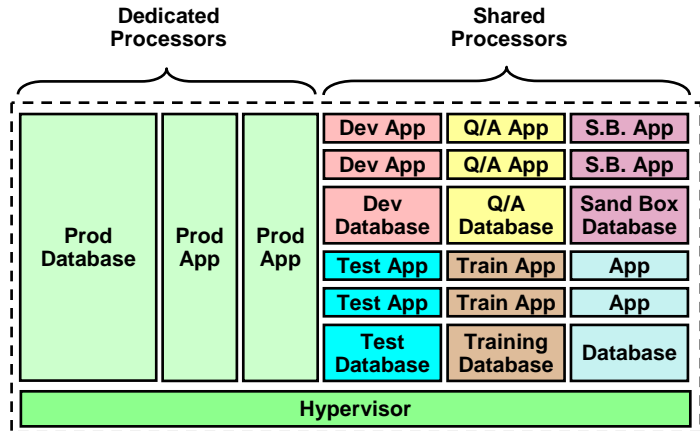


- **Reduce costs**
 - Share resources between fluctuating workloads
 - Potentially reduce software costs
 - Lower network adapter and port costs
- **Improve Service Levels**
 - Resize instances to match capacity requirements
- **Improve systems management**
 - Simplify capacity management
 - Clone new instances easily
- **Reduce time to market**
 - Deploy new instances rapidly

Consolidate Multiple Application Environments

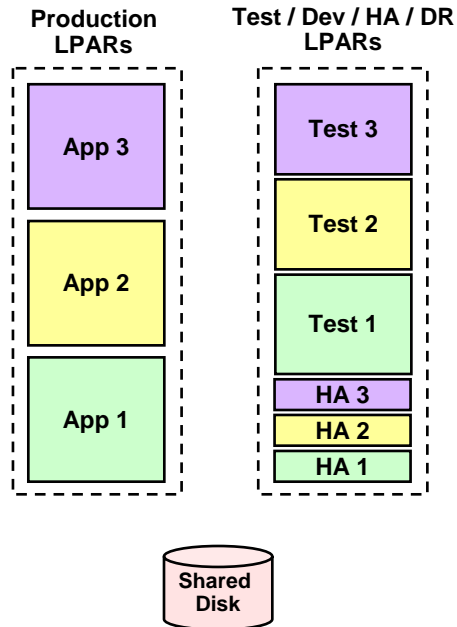
ERP Example

- **Reduce costs**
 - ▶ Reduce dedicated instance resources
 - ▶ Reduce network ports, adapters, and set up costs
 - ▶ Facilitate stress testing without excess hardware
- **Improve systems management**
 - ▶ Easily create new instances
- **Improve service levels**
 - ▶ Move resources between LPARs to handle fluctuating needs
 - ▶ Dynamically add additional application server instances
 - ▶ Utilize low latency in memory networks



High Availability Example

- **Reduce cost**
 - ▶ Significantly reduce idle high availability hardware
 - ▶ Automatically extend HA LPARs when needed
- **Cross campus disaster recovery**
 - ▶ Combine HA / Test and DR
 - ▶ Significantly reduce disaster recovery time over DR hosting site



Mixed Production and Test Environment

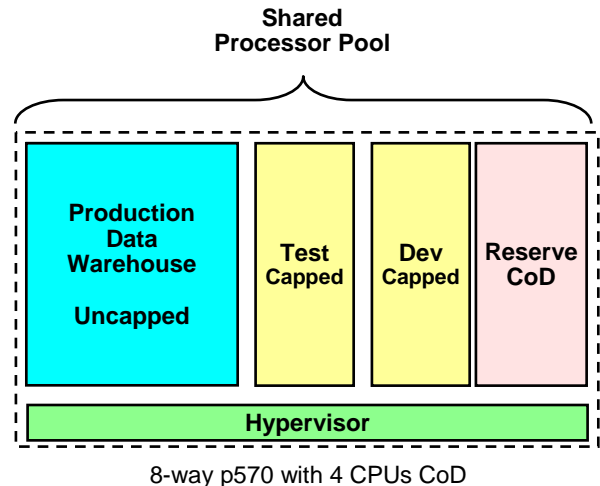
Data Warehouse Customer Example

■ Reduced costs

- ▶ Reduced the number of physical machines to manage
- ▶ Used idle test and development resources for production data warehouse

■ Improved service levels

- ▶ Eliminated older less reliable hardware
- ▶ Significantly reduced the nightly batch window without additional hardware
- ▶ Utilized reserve capacity on-demand for month or year end processing




Test Environment Example

- **Reduce Costs**

- ▶ Eliminate the need for low volume dedicated test environments

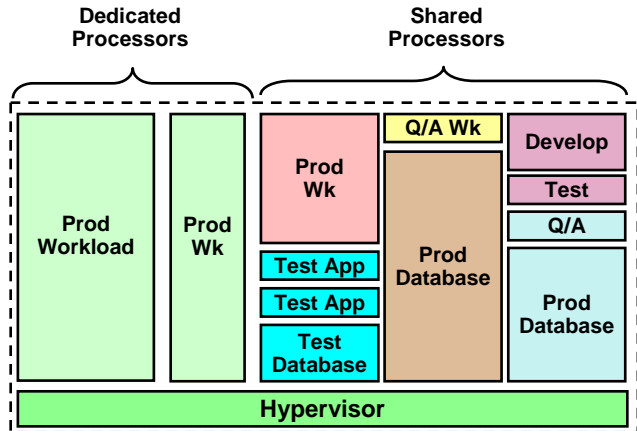
- **Improve systems management**

- ▶ Deploy test instances on the same hardware type as production
- ▶ Utilize the same hardware to do functional and performance stress testing
- ▶ Test different O/S and software combinations on the same hardware
- ▶ Rapid reboots

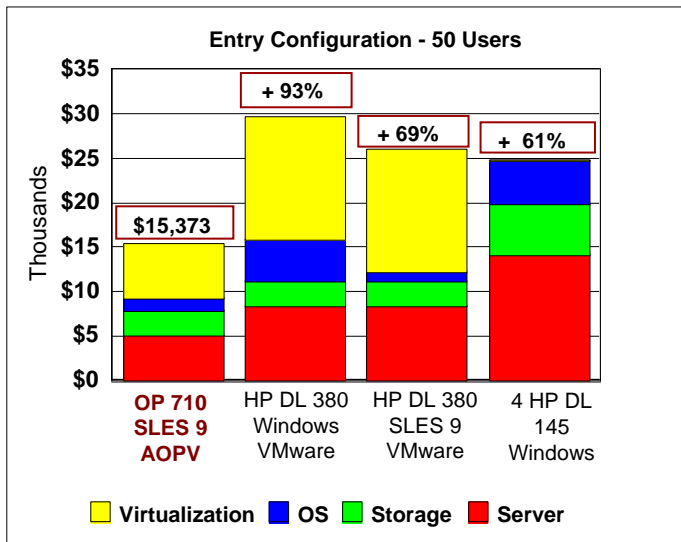


Mixed Workload - Server Consolidation Example

- **Reduce costs**
 - Consolidate underutilized servers
 - Share resources between workloads
 - Reduce floor space costs
- **Respond to market conditions**
 - Rapidly provision / de-provision servers to respond to business needs
 - Change server capacity to match business requirements
- **Improve systems management**
 - Improve testing environments
 - Ease software migrations
 - Support different software levels to match application needs
 - Simplify capacity management
 - Simply disaster recovery



OpenPower Consolidation Pricing Comparison



- OpenPower with virtualization:
 - Lower price than HP with VMware
 - Lower price than four (4) dedicated low-end servers
 - Higher performance
 - Fewer servers to manage than dedicated servers
 - More adaptable/flexible than dedicated servers

The solutions represented by the bar graph assume a client is consolidating 4 servers based on 15% utilization using Advanced OpenPower Virtualization or VMware. The OpenPower Consolidation Express solution entry configuration with 50 users was used. Derived from the WebBench benchmark (see footnote 2), one OpenPower 710 can collapse 4 1-way HP DL 145 servers with 65%+ reserve capacity. IBM Advanced OpenPower Virtualization is assumed to be 95% efficient. The last bar represents a solution which contains 4 dedicated servers. It is for illustrative purposes only, and actual solutions for the same or similar replacements will vary. Prices for the DL380 hardware, the DL140, and the VMware for Virtual Infrastructure Node (for two processors) are U.S. prices as listed at <http://www.hp.com> on January 4, 2005. Prices for the OpenPower Consolidation Express solution are IBM U.S. list prices as of January 25, 2005. Reseller prices may vary. Linux operating system prices for the OpenPower Consolidation Express solution are applicable only when licensed through IBM

Summary

Linux on POWER technology — performance and reliability for an on demand world

Linux

- Fastest-growing operating environment
- Provides unmatched versatility
- A cost-effective, security-rich environment, powerful enough to run business-critical applications +
- Hundreds of applications available

POWER

- A unifying architecture featuring the most innovative chip available
- Exceptional reliability, availability and serviceability — and the scalability you need to power your business
- [IBM Virtualization Engine](#) enables consistent management of heterogeneous environments

Linux on POWER

Great today, even better tomorrow

A clear path for the future, backed by an industry leader



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.

Many of the pSeries features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html

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 February 6, 2004

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, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DFDSM, DirectTalk, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, IBM, IBM(logo), ibm.com, IBM TotalStorage Proven, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, POWERparallel, PowerPC, PowerPC(logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RISC System/6000, RS/6000, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SP1, SP2, SpaceBall, System/390, The Engines of e-business, THINK, ThinkPad, 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/L(logo), AIX 5L, AIX PVMe, AS/400e, BladeCenter, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, Domino, e-business(logo), e-business on demand, eServer, GigaProcessor, HACMP, HACMP/6000, Hypervisor, i5/OS, IBMLink, IBM Virtualization Engine, IMS, Intelligent Miner, Micro-Partitioning, iSeries, NUMACenter, OpenPower, POWER, Power Architecture, Power Everywhere, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, S/390 Parallel Enterprise Server, ThinkVision, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Visualization Data Explorer, X-Architecture, z/Architecture.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

UNIX is a registered trademark in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

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

Intel, Itanium and Pentium are registered trademarks and Intel Xeon and MMX are trademarks of Intel Corporation in the United States and/or other countries

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 and/or other countries.

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

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

Revised August 23, 2004

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

Other company, product and service names may be trademarks of their respective owners. © Copyright IBM Corporation 2004