

Systems and Technology Group

Virtualization concepts & features with VMware

- 18 marzo 2009 -

Alessandro Malosio

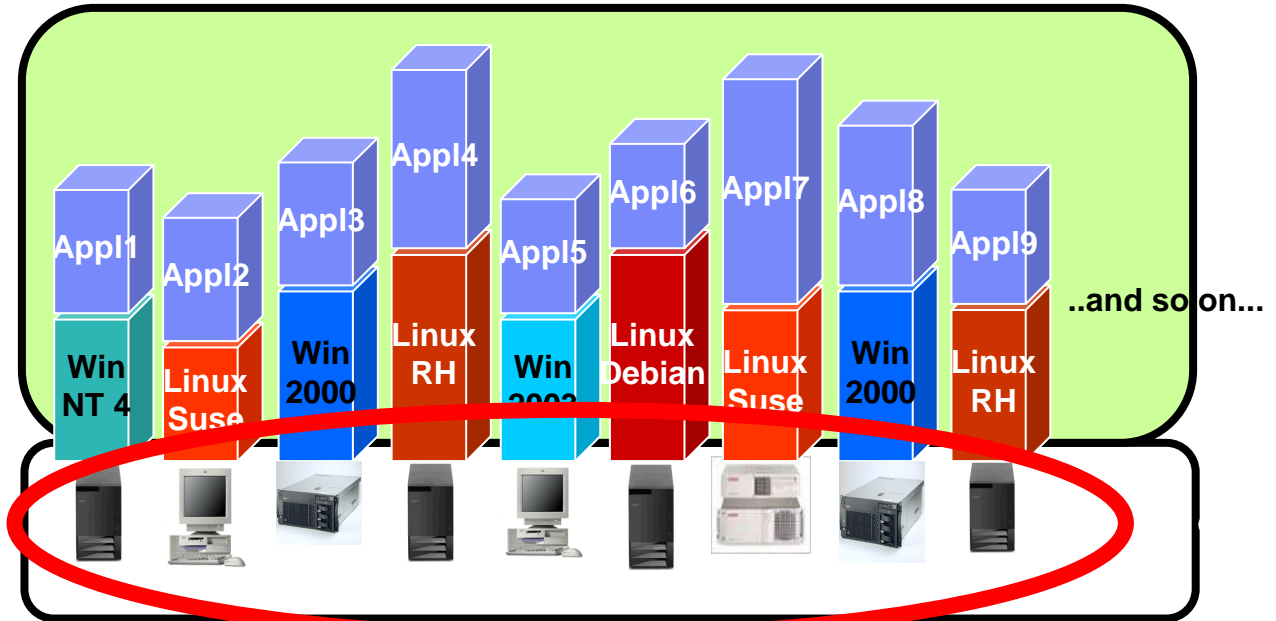
Modular Systems Technical Sales Support Team



ON DEMAND BUSINESS™


© 2007 IBM Corporation

Current Intel Server landscape



- ...too many physical server
- ...too many logical server (OS images)
- ...Wintel approach: new application to deploy, install new system
- ...to many resources involved in:
 - systems administration
 - systems integration
 - systems management

REVOLUTION



x4 architecture servers
System x
3850M2
3950M2

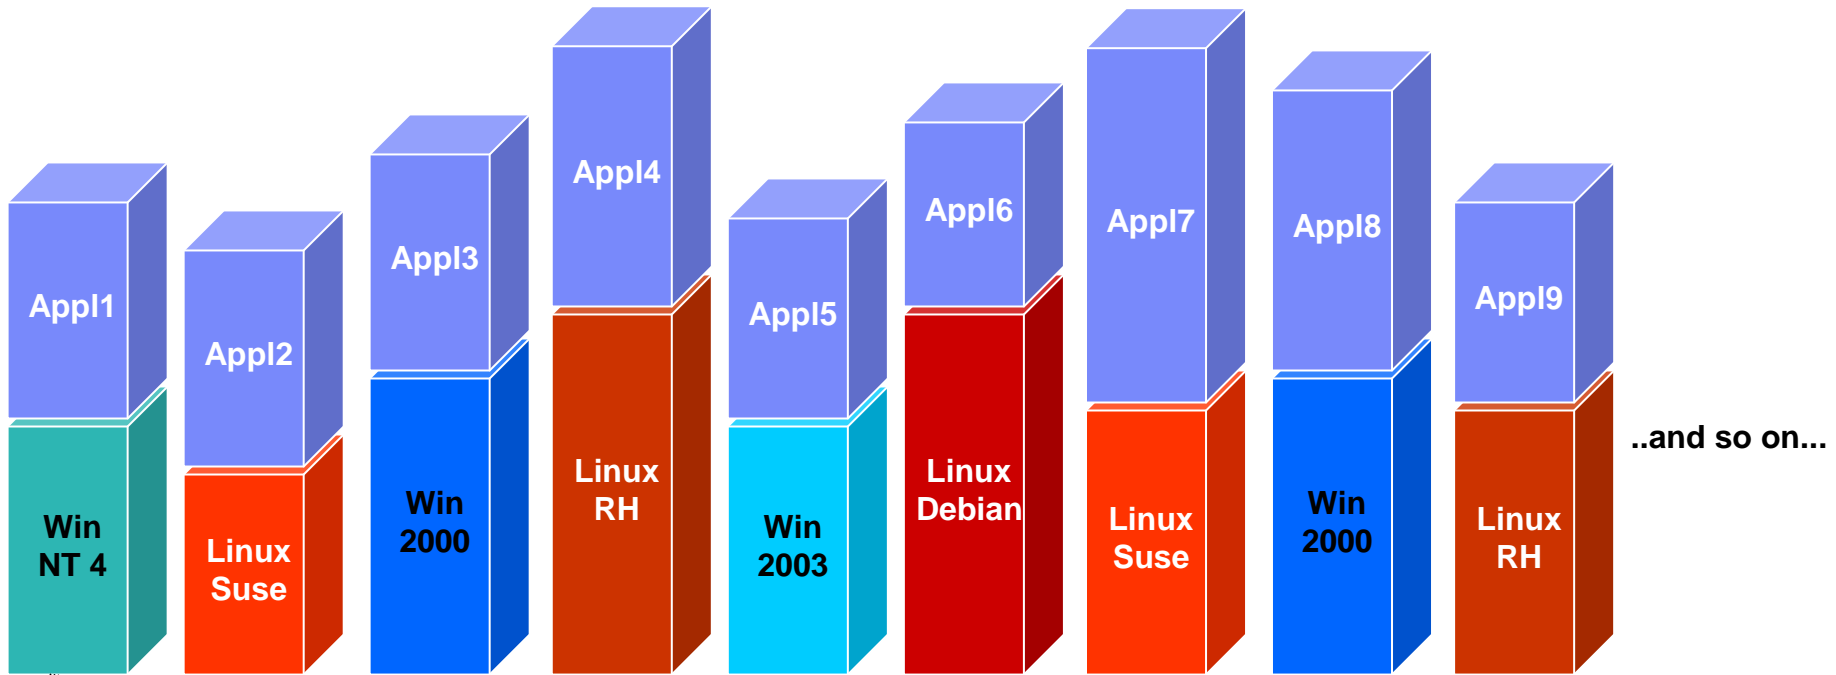
SAN

SW partitioning (i.e.VMware)



LAN Eth

EVOLUTION + **REVOLUTION** = **INNOVATION**

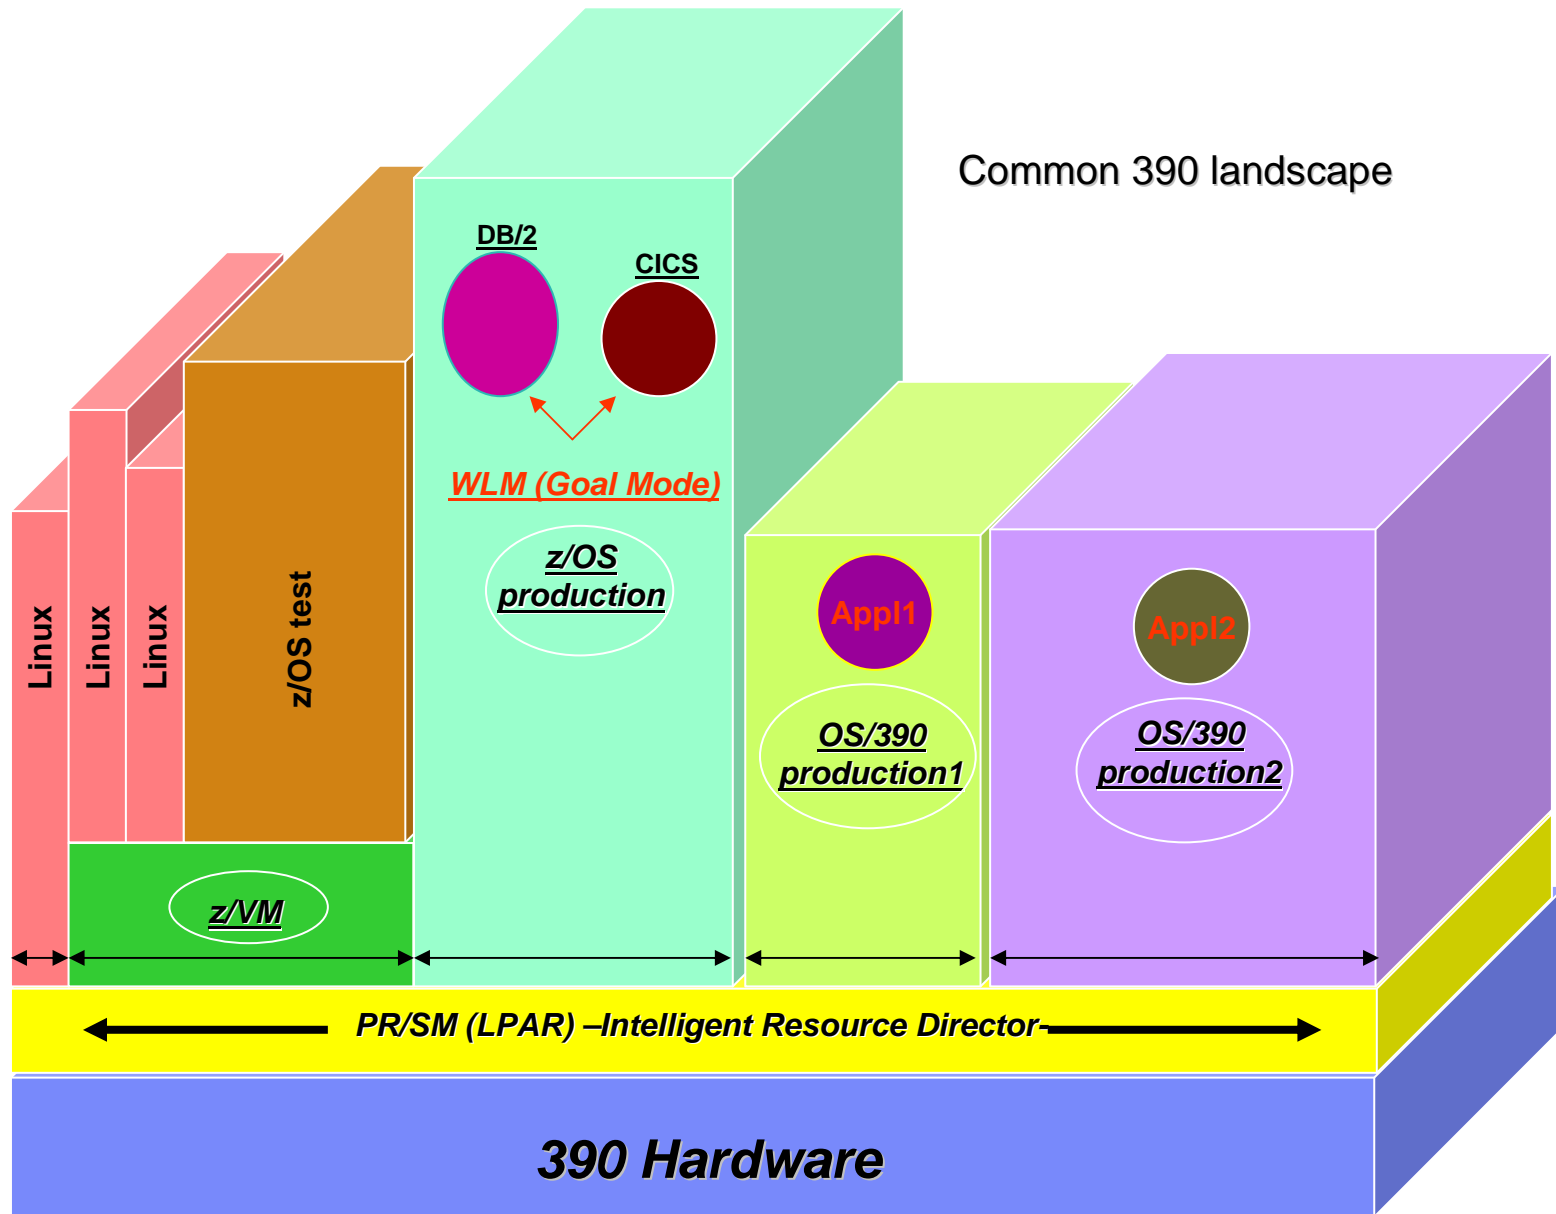


SW partitioning (i.e. VMware)

SAN




LAN Eth



New Server Consolidation approach is... Virtualization

“Virtual Machine Technology - Enables multiple operating systems to run isolated, concurrently and highly securely in virtual machines on a single server “

- VM technology was developed for the IBM Mainframe over 25 years ago
- VM technology (LPAR) available on IBM z-, p- & i-Series servers
- With VMware, VM technology is available on IBM xSeries servers


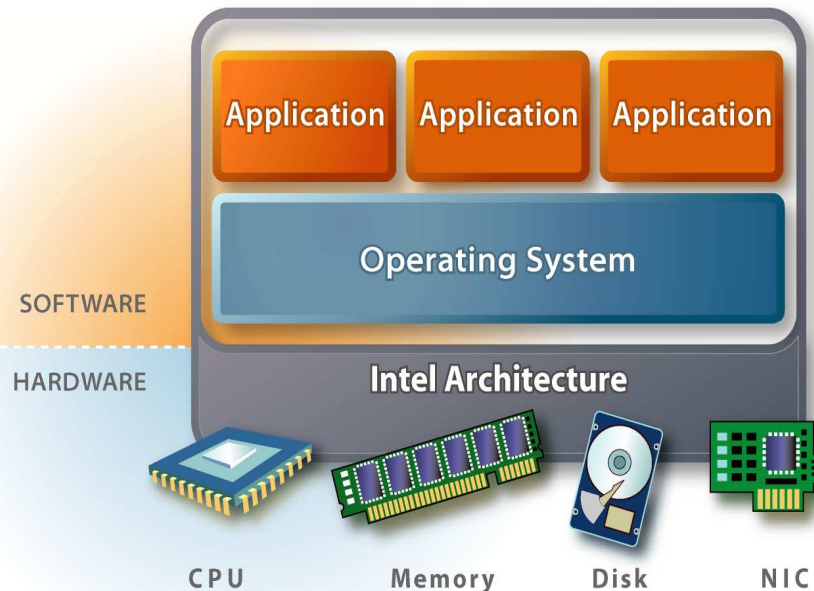


VMware ESX Architecture

"standard" computer

vs

ESX architecture



- Virtualization layer maps virtual hardware to real hardware.
- Can multiplex several virtual hardware to single real HW.
- High Performance – map directly on hardware.
- Run multiple operating systems concurrently
- Fault, performance, security isolation
- Encapsulation
- Hardware-independent


Feature: Isolation

- **Key: uses CPU hardware (protection)**
- **Fault, performance, and security isolation**
- **CPU, RAM, Disk, and network resource controls***
- **Guarantee service levels***

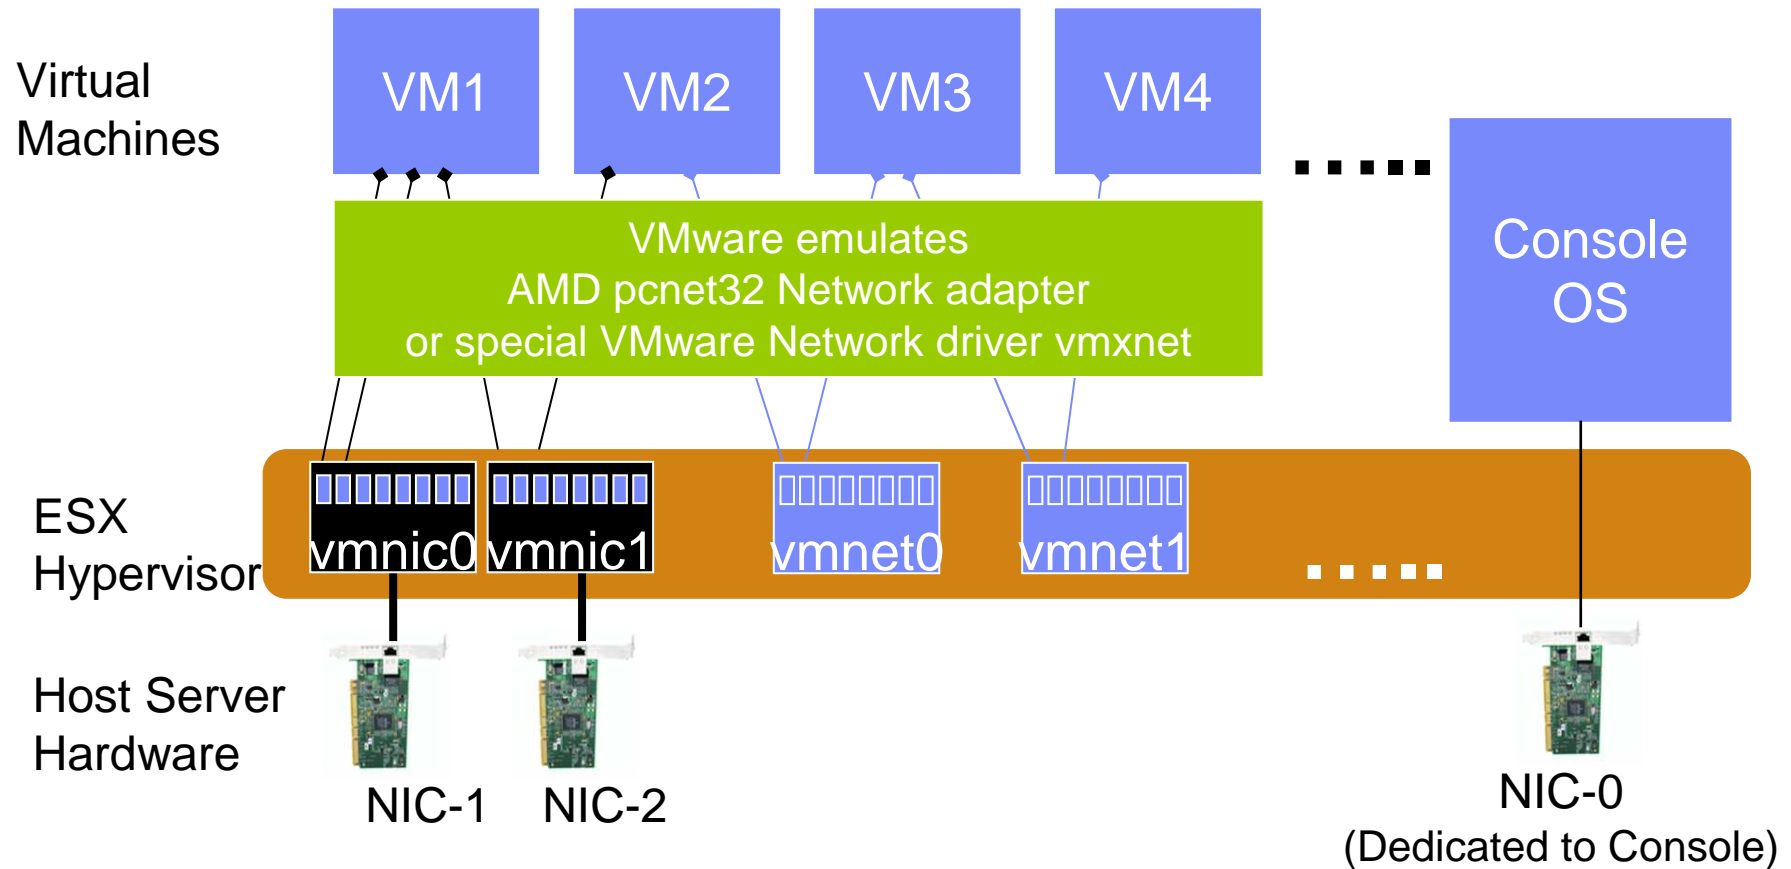


Feature: Encapsulation

- Entire state of the VM is encapsulated
 - ▶ Memory, disk images, I/O device state
- VM state can be saved to a file
- VM state can be transferred through time and space
 - ▶ Time: store in a file
 - ▶ Space: transfer over a network

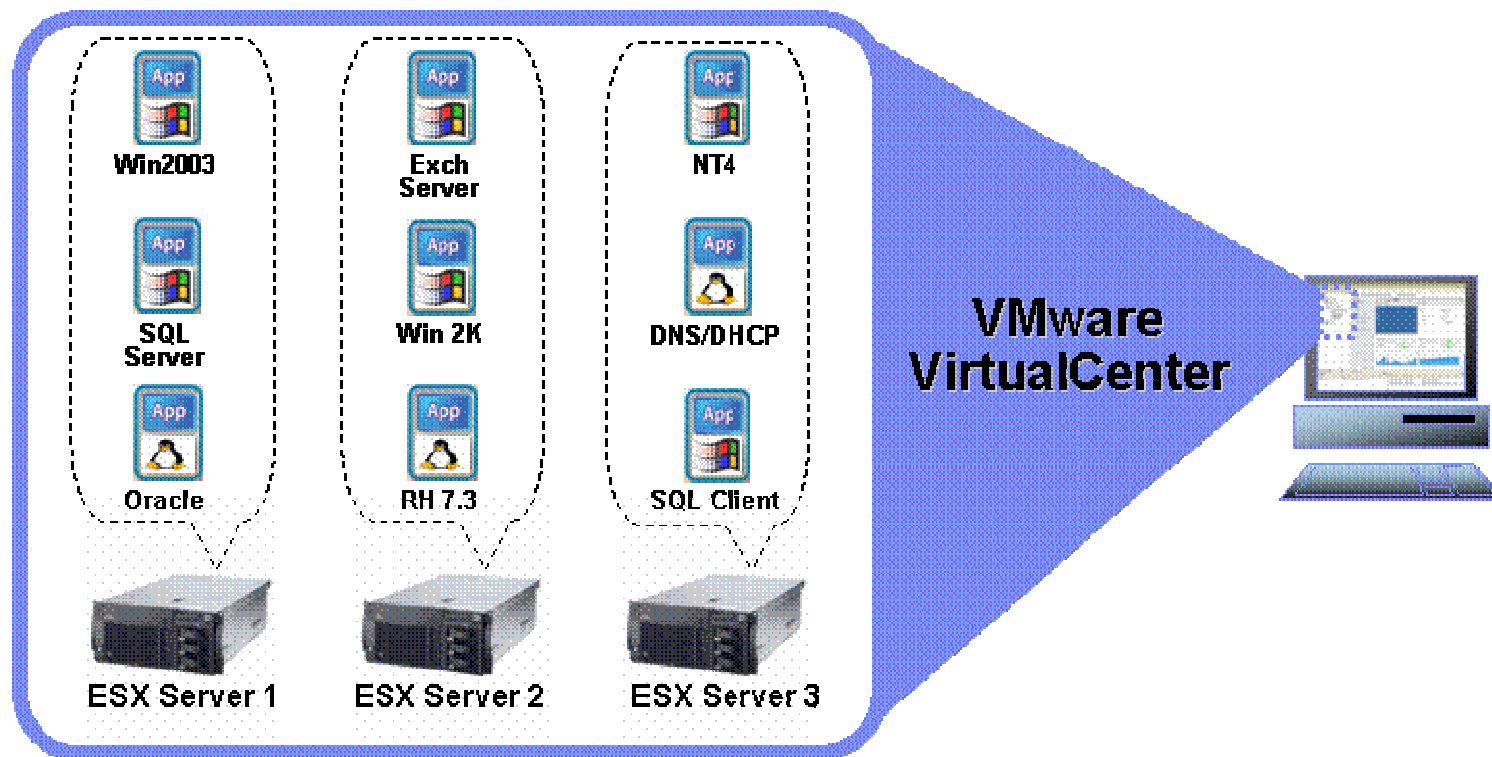


Virtual Machine Networking



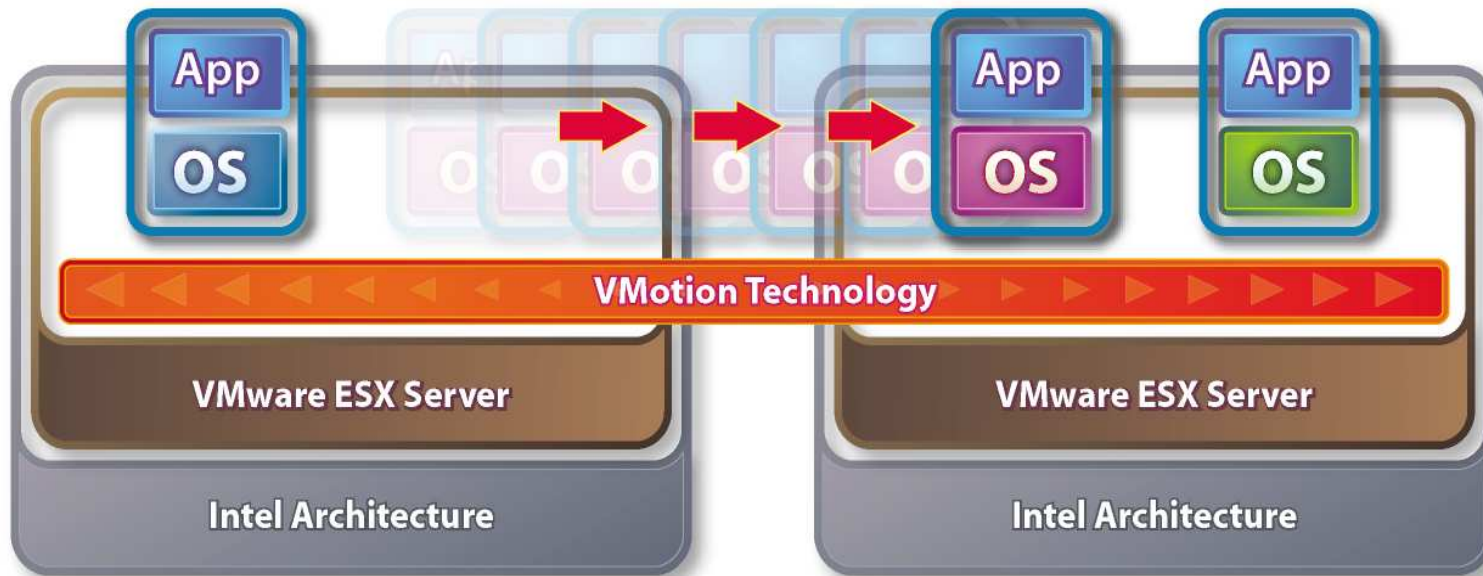
Centralized Management Console

Centrally manage a heterogeneous computing environment from a single GUI



VMotion™ Technology


Instantly shift running systems across hosts often with imperceptible downtime



- High application availability
- High transaction integrity
- High data availability
- High transparent to end users

VMotion™ – Potentially Eliminate Planned Downtime

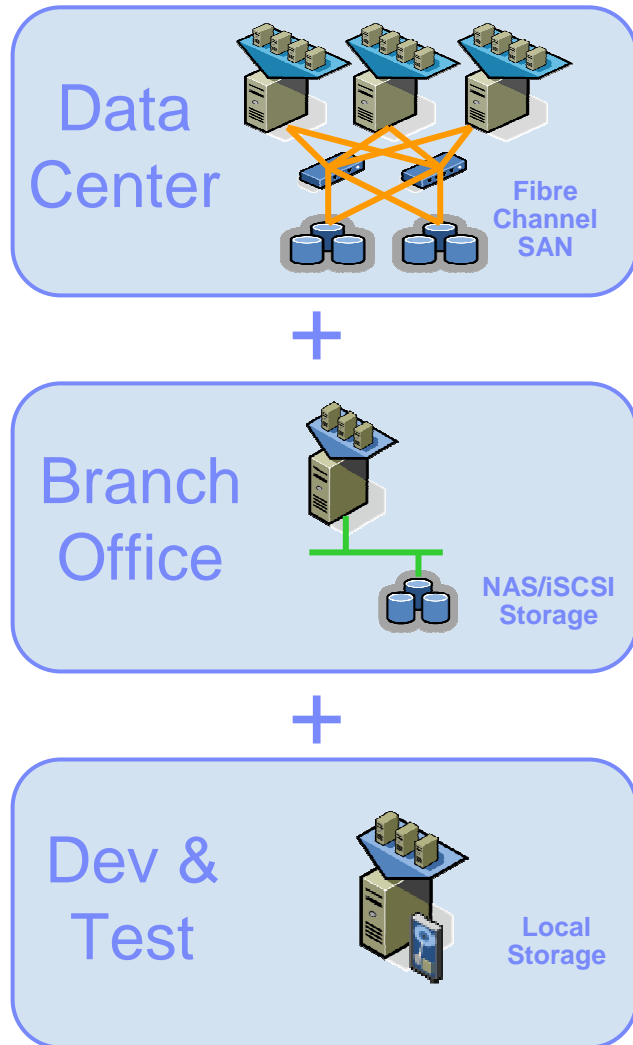
Upgrade and service production hardware using VMotion™ with near zero downtime and high customer transparency



Call For Upgrade

(e.g. replacing defective
PCI card or BIOS
upgrade)

ESX Server 3 (Virtual Infrastructure 3) Features




- NAS and iSCSI storage
- Expanded hardware compatibility list
- 4-way Virtual SMP (..from previous 2-way)
- 16GB guest memory (..from previous 3.6GB)
- Hot-add virtual disks
- Red Hat Enterprise Linux 4 guests
- Multiple snapshots
- Up to 128 powered on VMs per server
- Updated Service Console (based on RHEL 3)
- Flexible virtual switches
- 64-bit guest technology preview

Available from
June 2006

DRS (Distributed Resource Scheduler) Add-On

Creating a Unified Compute Resource

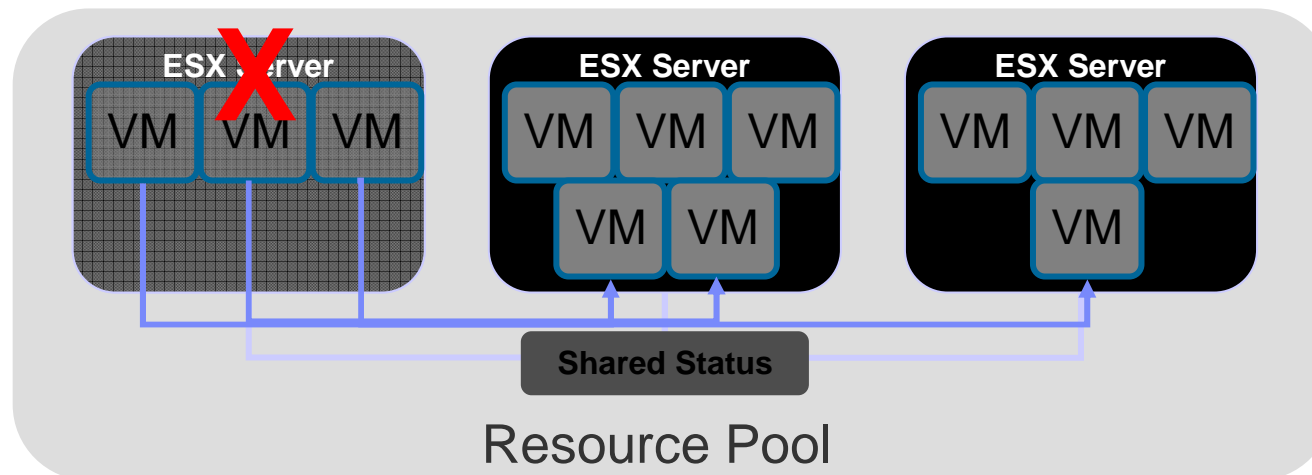
- **Global scheduler**
 - ▶ Automates initial virtual machine placement
 - ▶ Uses VMotion to continuously optimize based on current workload
 - ▶ Reacts to adding or removing hosts from the cluster
- **Achieve >80% utilization**



HA (High Availability) Add-On

High availability for all your servers

- **Losing a host in a cluster means fewer resources, not lost virtual machines**
 - ▶ Impacted virtual machines are restarted on remaining hosts
 - ▶ Placement optimized by global scheduler
 - ▶ VirtualCenter handles all setup and configuration automatically
- **None of the cost and complexity of clustering**



What's New with VMware Infrastructure 3.5

- ESX Server 3.5 and ESX Server 3i v3.5 -



- ESX Server 3i
- VMware Update Manager
- VMware Storage VMotion (only CLI, on same host)
- VMware Site Recovery Manager (G.A. Q1-08)
- VMware Distributed Power Management (Experimental)
- ESX Server Performance Optimizations
- ESX Server Scalability Enhancements (max 128GB for each host; max 64GB for each V.M.)
- Expanded Storage and Networking Choices (10GbE, IB, local SATA)

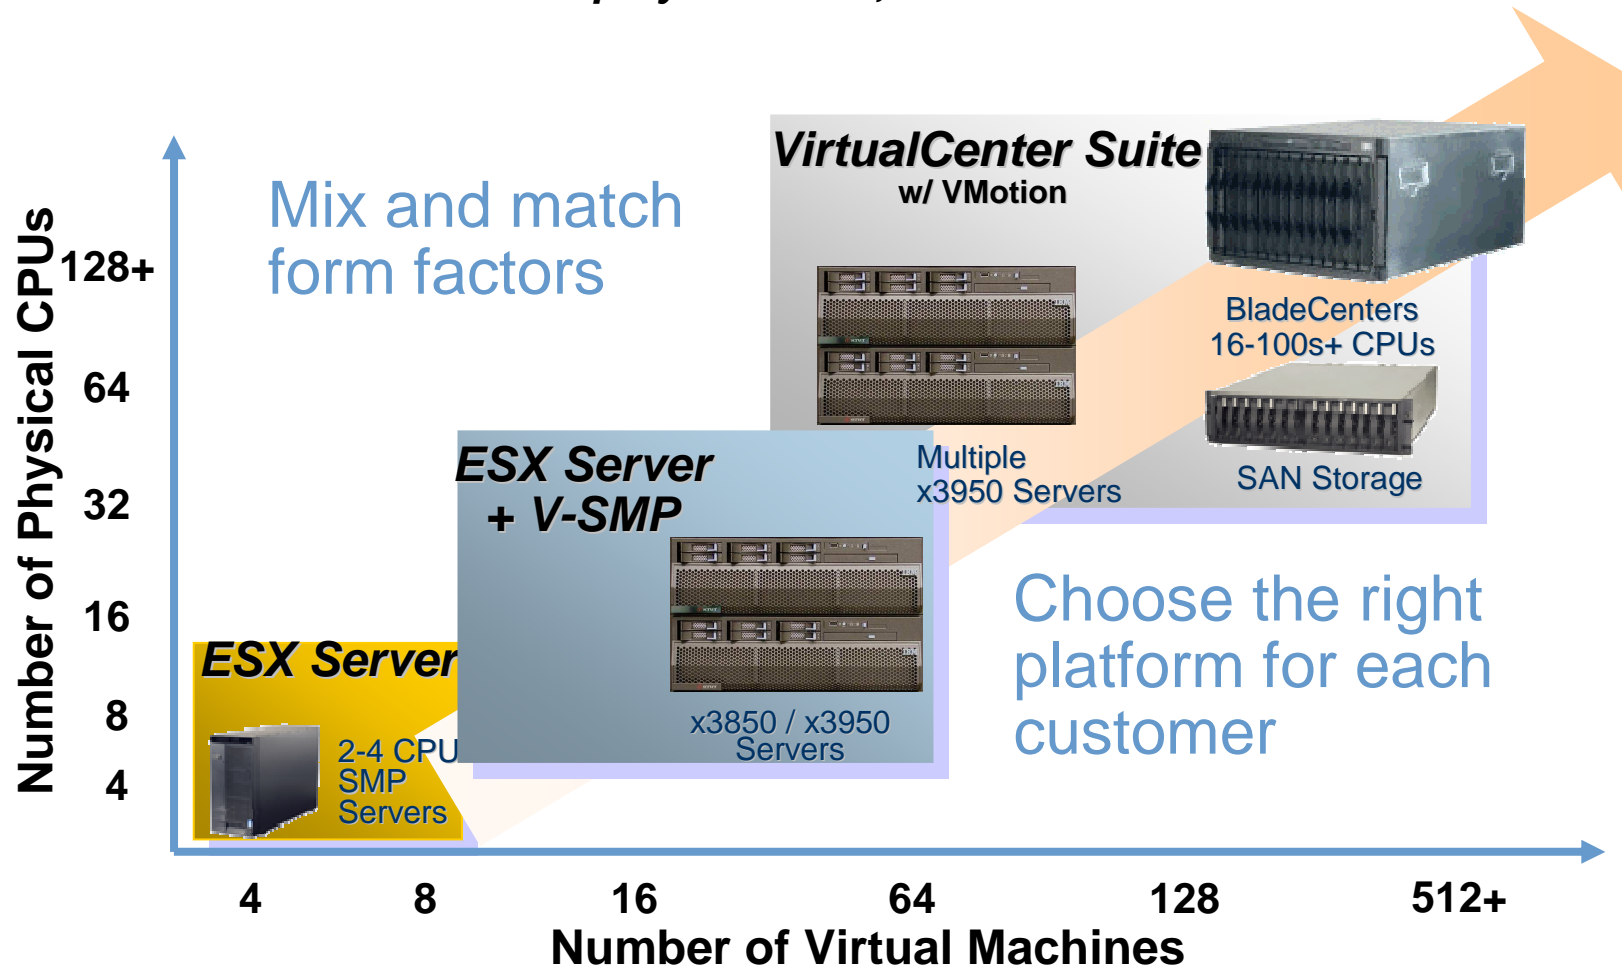
- **Announced at VMworld 2007**
- **Planned availability date Q4 2007**

<http://www.vmware.com/products/vi/whatsnew.html>

Virtualization: why with IBM ?

IBM and VMware – Positioning

Massive Scalability with VMware and IBM xSeries
Benefits scale across deployment size, server form factors



HS21 XM – VMware 3i Preload




HS21 XM – 3i preload
Announce: 25 march 2008
Availability: 9 may 2008

<u>MTM</u>	<u>GAV</u>	<u>CPU Speed</u>	<u>CPU Power</u>	<u>FSB</u>	<u>CPU Cache</u>	<u>CPUs Std</u>	<u>Memory Std</u>	<u>DIMM Slots</u>	<u>HDDs</u>	<u>Blade Width</u>	<u>Additional</u>
7995-HVx	7995-HVY	2 x 3.0GHz (E5450)	80W	1333MHz	12MB	2	2 x 2GB	8	4GB Modular Flash Drive	30mm	VMware3i Preload

- Dedicated model of HS21 XM (7995-HVY)
 - ▶ Contains embedded hypervisor on 4GB Modular Flash Drive



IBM Modular Flash Drive

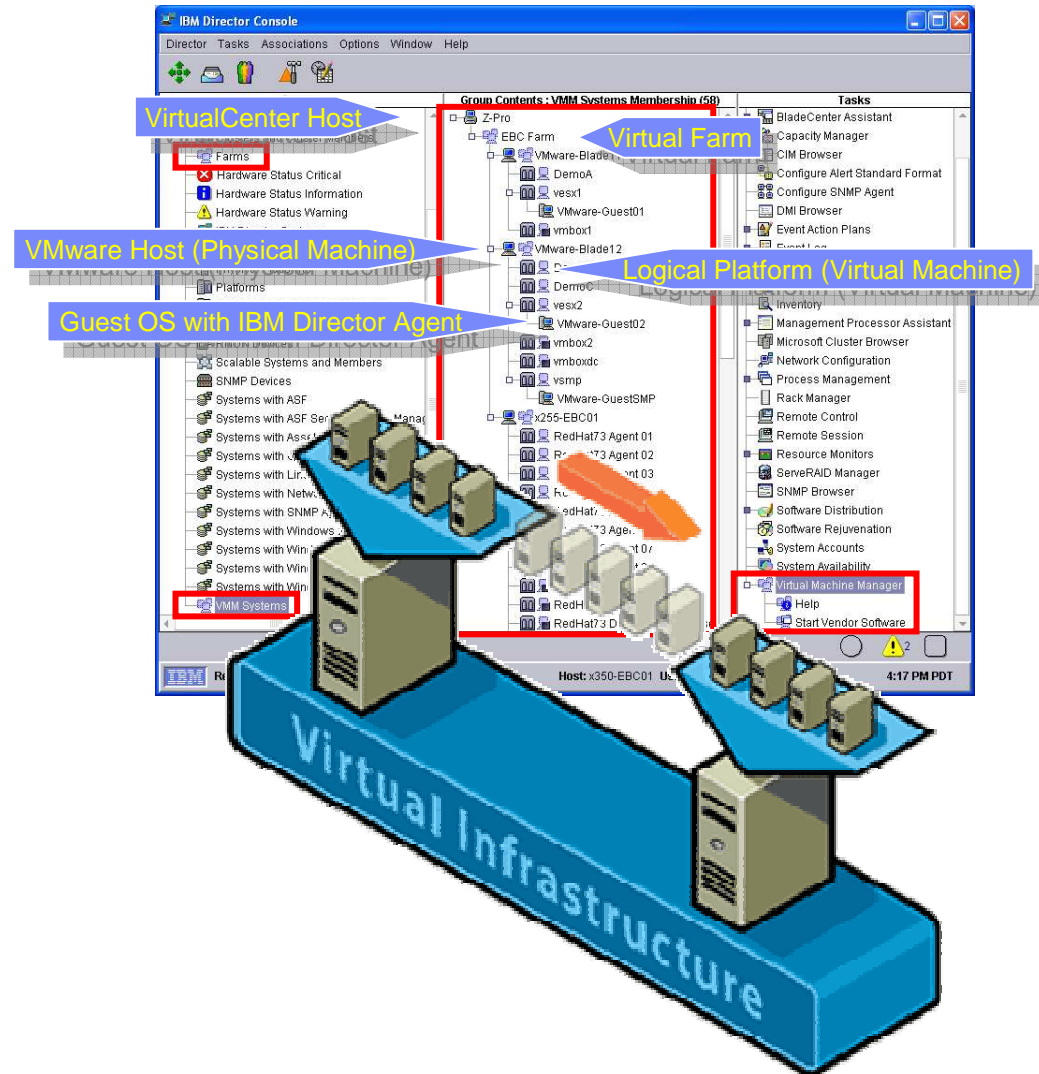


IBM Director & VMware: Differentiation Through Integration

- **New! Virtual Machine Manager**
 - ▶ VMM free add-on to IBM Director
 - ▶ “Single glass management” of virtual/physical machines

- **Improves core management of VMs inside of IBM Director**
 - ▶ Create/Delete VMs and Virtual Farms
 - ▶ Manage resource allocations to VMs
 - ▶ Scheduled static migration of VMs (VMware and MS Virtual Server)

- **VMM functions driven by event action plans**
 - ▶ PFA Alert on pending hardware failure triggers VM migration using VMotion
 - ▶ Eliminates downtime due to server subsystem failures



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 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/systems/p/software/whitepapers/linux_overview.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 January 19, 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, 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), PowerPC 601, 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 PVM, AS/400e, Chipkill, Chiphopper, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, e-business(logo), e-business on demand, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, GigaProcessor, HACMP, HACMP/6000, I5/OS (logo), IBMLink, IBM TotalStorage Proven, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, ON DEMAND BUSINESS logo, OpenPower, POWER, 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+, 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>.

UNIX is a registered trademark 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 and/or other countries.

Intel, Intel Xeon, Itanium and Pentium are registered trademarks or 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).

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.

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

Revised January 19, 2006

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 p5, ~ p5, pSeries, OpenPower and IBM RS/6000 Performance Report at http://www.ibm.com/servers/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.shtm
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 January 19, 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 @server 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 January 19, 2006