IBM Systems and Technology Group

Comparing System z Virtualization to System p and System x

April 26, 2007

Reed A. Mullen
mullenra@us.ibm.com
Virtualization Strategy Manager for System z

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DB2, e-business logo, ESCON, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, System z9, VM/ESA, VSE/ESA, WebSphere, xSeries, z/OS, zSeries, z/VM.

The following are trademarks or registered trademarks of other companies

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
LINUX is a registered trademark of Linux Torvalds in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
Intel is a registered trademark of Intel Corporation.
* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Agenda

- Virtualization and hypervisor basics
- System z™ virtualization technology highlights
- System p™ virtualization overview
- System x™ virtualization overview
- System z strategic considerations

Server Partitioning and Virtualization

Affecting All Areas of Virtual Server Design and Deployment

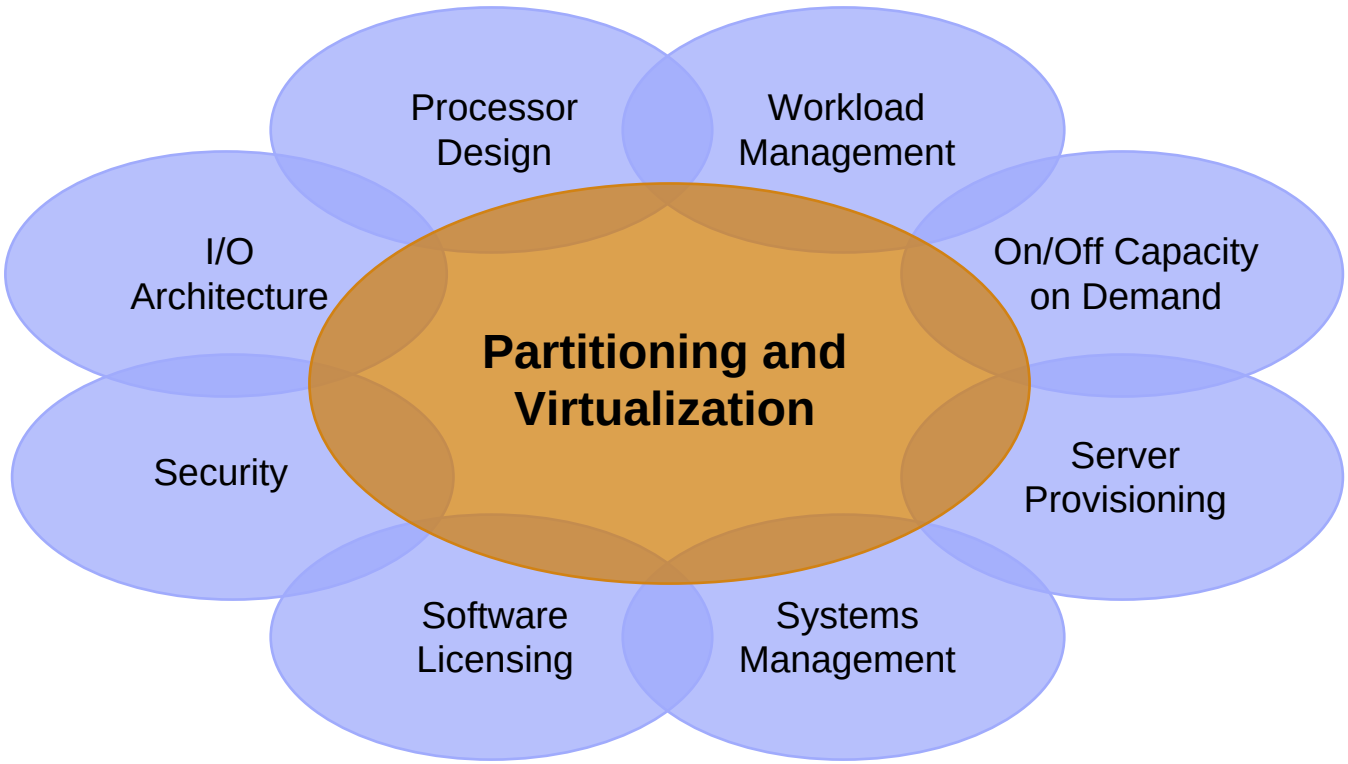

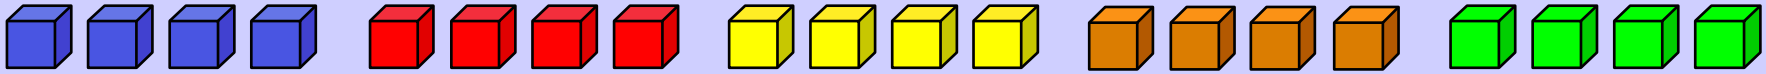

Virtual Resources

- Proxies for real resources: same interfaces/functions, different attributes.
- May be part of a physical resource or multiple physical resources.

Virtualization

- Creates virtual resources and "maps" them to real resources.
- Primarily accomplished with software and/or firmware.

Resources

- Components with architected interfaces/functions.
- May be centralized or distributed. Usually physical.
- Examples: memory, disk drives, networks, servers.

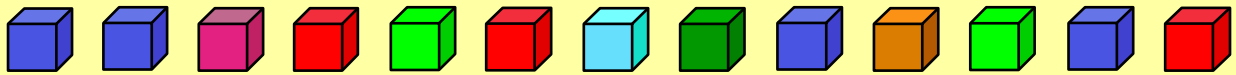

- **Separates presentation of resources to users from actual resources**
- **Aggregates pools of resources for allocation to users as virtual resources**

Hardware Partitioning

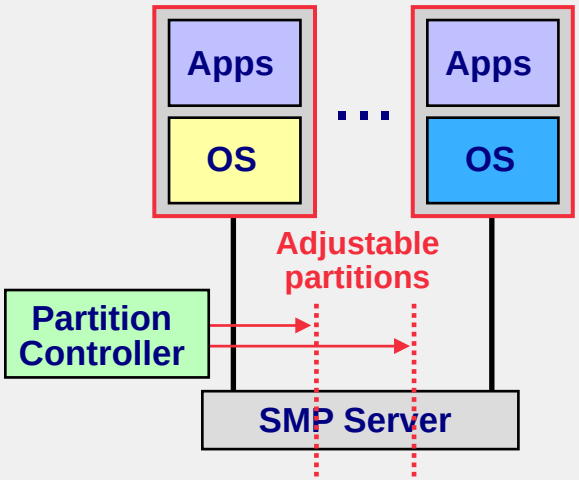

Server is subdivided into fractions each of which can run an OS

Physical partitioning
S/370™ SI-to-PP and PP-to-SI, Sun Domains, HP nPartitions

Logical partitioning
IBM eServer™ pSeries® LPAR, HP vPartitions

Bare-metal Hypervisor

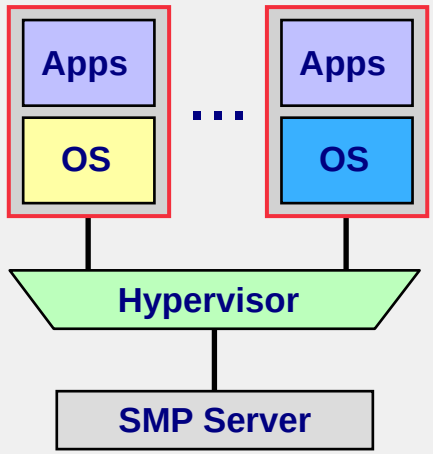

Hypervisor provides fine-grained timesharing of all resources

Hypervisor software/firmware runs directly on server
System z PR/SM™ and z/VM®
POWER™ Hypervisor
VMware ESX Server
Xen Hypervisor

Hosted Hypervisor

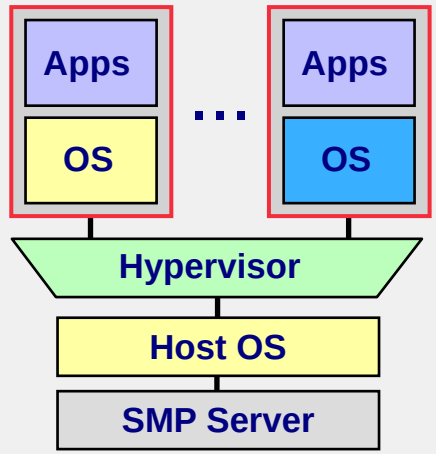

Hypervisor uses OS services to do timesharing of all resources

Hypervisor software runs on a host operating system
VMware GSX
Microsoft® Virtual Server
HP Integrity VM
User Mode Linux®

Characteristics:

- Bare-metal hypervisors offer high efficiency and availability
- Hosted hypervisors are useful for clients where host OS integration is important
- Hardware partitioning is less flexible than hypervisor-based solutions

Trap and Emulate

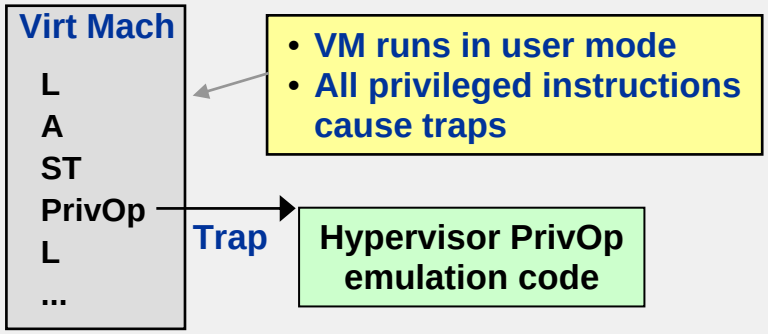

Examples: CP-67, VM/370
Benefits: Runs unmodified OS
Issues: Substantial overhead

Translate, Trap, and Emulate

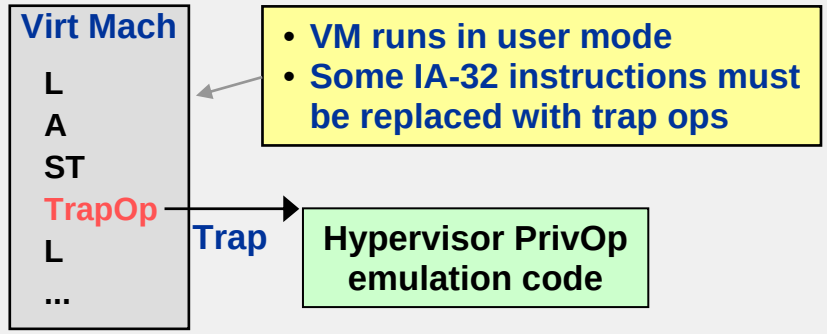

Examples: VMware, Microsoft VS
Benefits: Runs unmodified, translated OS
Issues: Substantial overhead

Hypervisor Calls ("Paravirtualization")

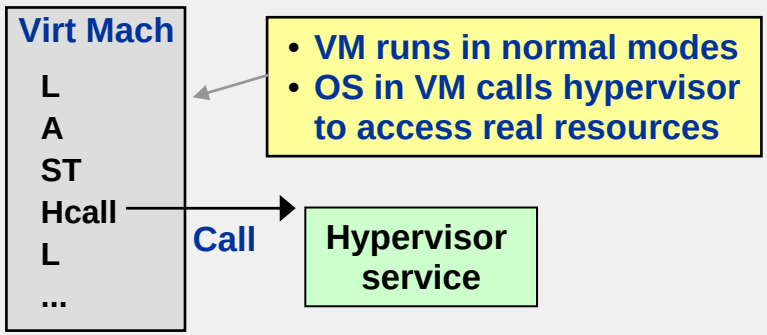

Examples: POWER Hypervisor, Xen
Benefits: High efficiency
Issues: OS must be modified to issue Hcalls

Direct Hardware Virtualization

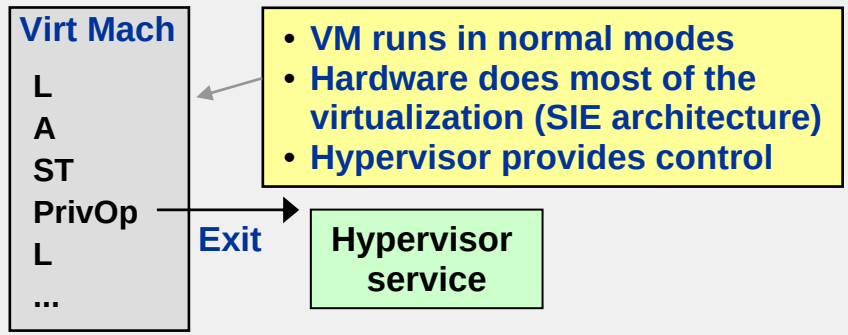

Examples: PR/SM, z/VM (also use hypervisor calls)
Benefits: High efficiency, runs unmodified OS
Issues: Requires underlying hardware support

Roles:

- Consolidations
- Dynamic provisioning / hosting
- Workload management
- Workload isolation
- Software release migration
- Mixed production and test
- Mixed OS types/releases
- Reconfigurable clusters
- Low-cost backup servers

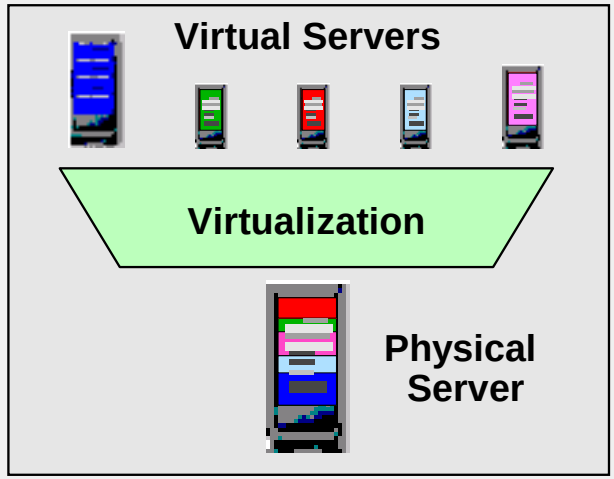

Possible Benefits:

- High resource utilization
- Great usage flexibility
- Enhanced workload QoS
- High availability / security
- Low cost of availability
- Low management costs
- Enhanced interoperability
- Legacy compatibility
- Investment protection

In the final analysis, the potential virtualization benefits take three forms:

- **Help reduce hardware costs**
  - Help increase physical resource utilization
  - Small footprints
- **Can improve flexibility and responsiveness**
  - Virtual resources can be adjusted dynamically to meet new or changing needs and to optimize service level achievement
  - Virtualization is a key enabler of on demand operating environments
- **Can reduce management costs**
  - Fewer physical servers to manage
  - Many common management tasks become much easier

Key Value Propositions for Linux on System z

Linux and z/VM – Changing the Server Landscape

- **Very large-scale, single-image virtualization technology**
  - Host hundreds of virtual Linux servers on a single copy of z/VM
  - Run multiple copies of z/VM on a single mainframe or across multiple mainframes
  - Exploit sophisticated z/VM data-in-memory techniques among server images to enable superior qualities of service (e.g., virtual networking, execute-in-place file system, virtual disks in storage, in-memory disk cache)
  - Hardware, hypervisor, and OS synergies help deliver superior levels of resource utilization
- **Infrastructure simplification and possible cost savings**
  - Large-scale consolidation of servers, networks, and data
  - Host Linux-based solutions side-by-side z/OS environments, sharing system resources and simplifying operational tasks for the integrated solution
  - Over commit virtual resources and maintain exceptional levels of performance and response times
  - Take advantage of functionally rich, built-in systems management and automated operations capabilities; complement the environment with vendor products if needed
- **On demand computing capabilities**
  - Allocate resources with a high degree of granularity and efficiency; rapidly provision or reconfigure servers even when system resources are already heavily utilized
  - Non-disruptive upgrade features (e.g., 54X CPU capacity growth within a single mainframe)

IBM System z Virtualization Architecture

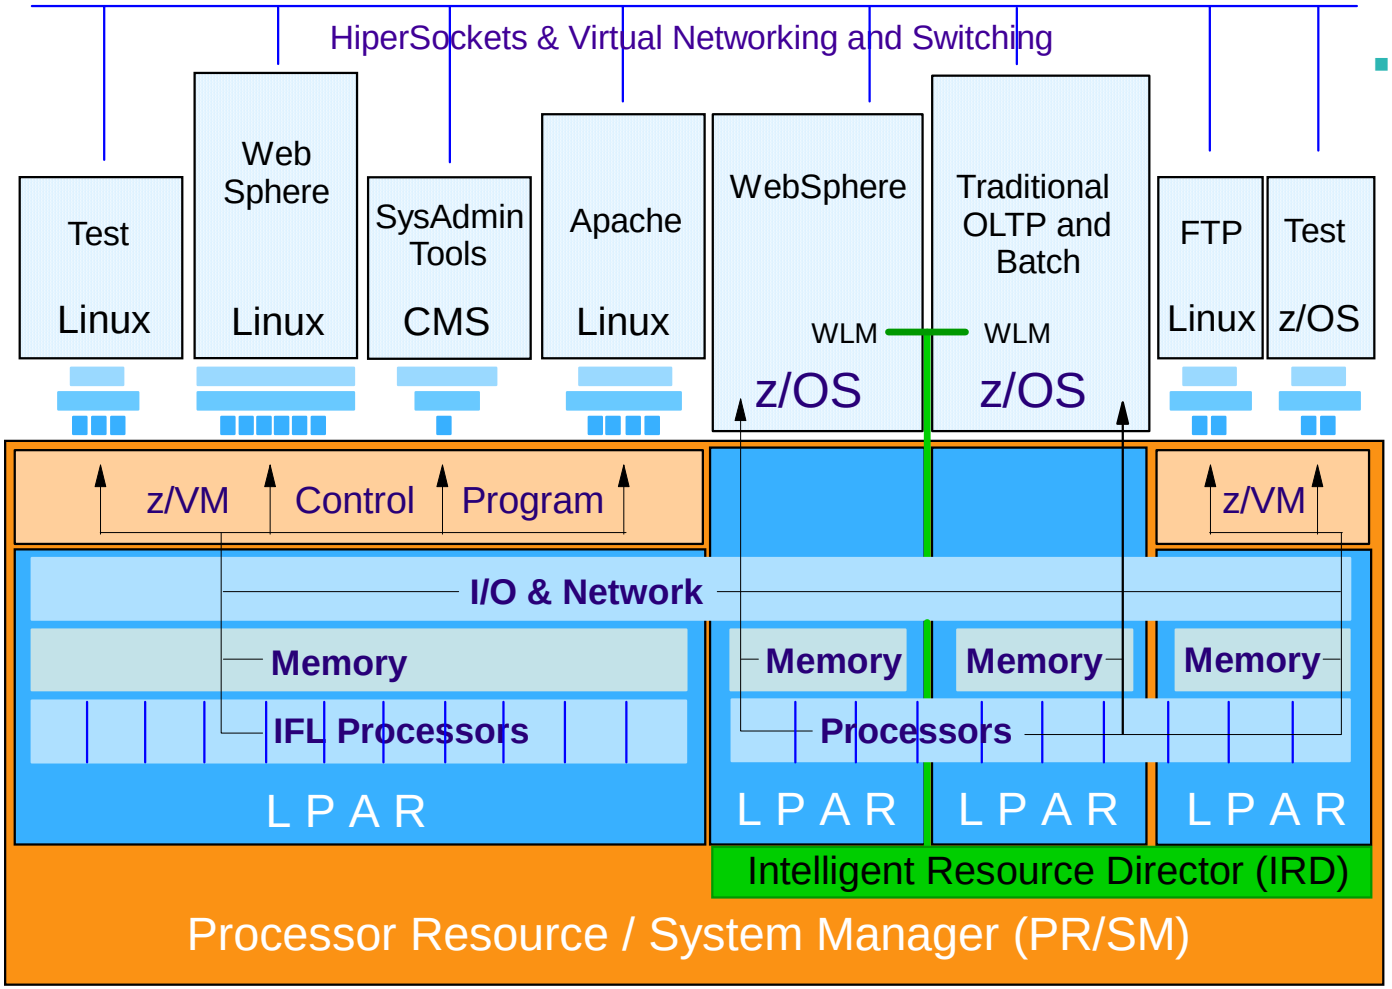

- **Multi-dimensional virtualization technology**
  - System z provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS® LPARs*

* Excluding non-shared resources like Integrated Facility for Linux processors

LPAR and z/VM
World-Class Server Virtualization

- **Mainframe Logical Partitioning (LPAR), introduced in 1988, has provided years of business-critical, high-performance server partitioning for the world's largest corporations**
- **z/VM, commercially available since 1972, has supported mixed workloads that require minimal hypervisor overhead, massive scalability, and exceptional levels of availability**
- **Both LPAR and z/VM employ hardware and firmware innovations developed over the years that make virtualization a fundamental part of IBM System z architecture**

System z Interpretive Execution

Advanced Technology for Virtual Server Hosting

- **Start Interpretive Execution (SIE) instruction**
  - Operand is a state descriptor for an LPAR or virtual machine
  - Accommodates fixed-storage and pageable guests
  - Interception controls allow hypervisor intervention
  - Reduces context switch time
- **System z implements two levels of SIE**
  - No performance penalty for running z/VM in an LPAR
  - No shadow page tables required for DAT-on guests
  - Considerable architectural and hardware investment required
    - Potential instruction behavioral differences at each level
    - Multiple control register sets

Additional Mainframe Virtualization Facilities

- **Zone Relocation**
  - SIE capability that provides multiple zero-origin storage regions (logical partitions) on one system
  - Enables I/O subsystem to access partition memory directly, without requiring hypervisor intervention
- **Translation Lookaside Buffers (TLBs)**
  - Large allocation of microprocessor space for TLBs directly benefits virtual server scalability
  - z9 and z990 provides a TLB arrangement which advantageously uses two buffers
    - Second-level TLB feeds address translation information to the first-level TLB when the desired virtual address is not contained in the first-level TLB
- **Multiple Image Facility (MIF)**
  - Enables channel sharing among multiple LPARs
  - I/O devices on shared channel paths can be accessed simultaneously by sharing LPARs (or restricted to a subset of sharing LPARs)
- **Logical Channel Subsystem (LCSS) support**
  - Allows a z9 and z990 to be configured with up to 1024 channels (512 channels for z890)
  - 256 channels can be configured for each LPAR, with selected channel sharing among LPARs possible

Additional Mainframe Virtualization Facilities

- **I/O Priority Queuing**
  - Allows high-priority workloads to receive preferential access to I/O subsystem
  - Supported by Intelligent Resource Director and virtualized by z/VM
- **HiperSockets**
  - High-speed, security-rich TCP/IP connectivity among LPARs
  - Memory speed communications
- **Adapter Interruption Pass-Through**
  - OSA-Express (Ethernet) and FCP (SCSI) virtual machine I/O can be performed while z/VM guest image is running in SIE mode
  - “Thin” interrupt passed to z/VM Control Program when I/O operation belongs to an idle guest system
- **QDIO Enhanced Buffer-State Management (QEBSM)**
  - Two new machine instructions designed to help eliminate overhead of hypervisor interception
- **Host Page-Management Assist (HPMA)**
  - Interface to z/VM paging and storage management
  - Designed to allow hardware to assign, lock, and unlock page frames without hypervisor assistance
- **Layer 2 (MAC) and Layer 3 (IP) network switching**
  - OSA and z/VM support enables virtual IP and MAC network switching without requiring a hosting partition

System Design Affects Virtualization Capabilities

System z packs a lot of compute power into a single box

- With TCO-friendly pricing

Up to 54-way SMP

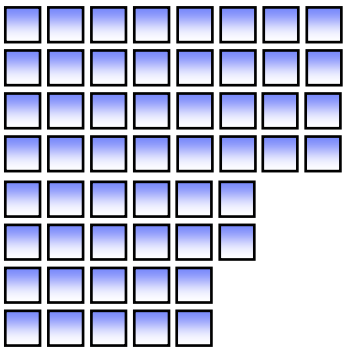

Share up to 54 processors with up to 60 LPARs

Configure these processors as CPs, IFLs, zAAPs*, zIIPs*, or ICFs*

* No software license fees

Up to 10 System Assist Processors

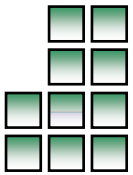

Offload system processing to dedicated CPUs with no impact to software license fees

Plus up to 336 I/O Processors

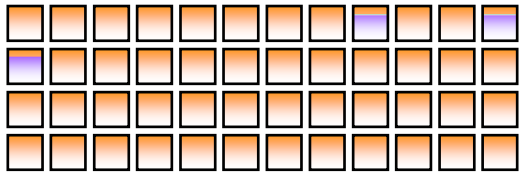

No additional charge for these processors

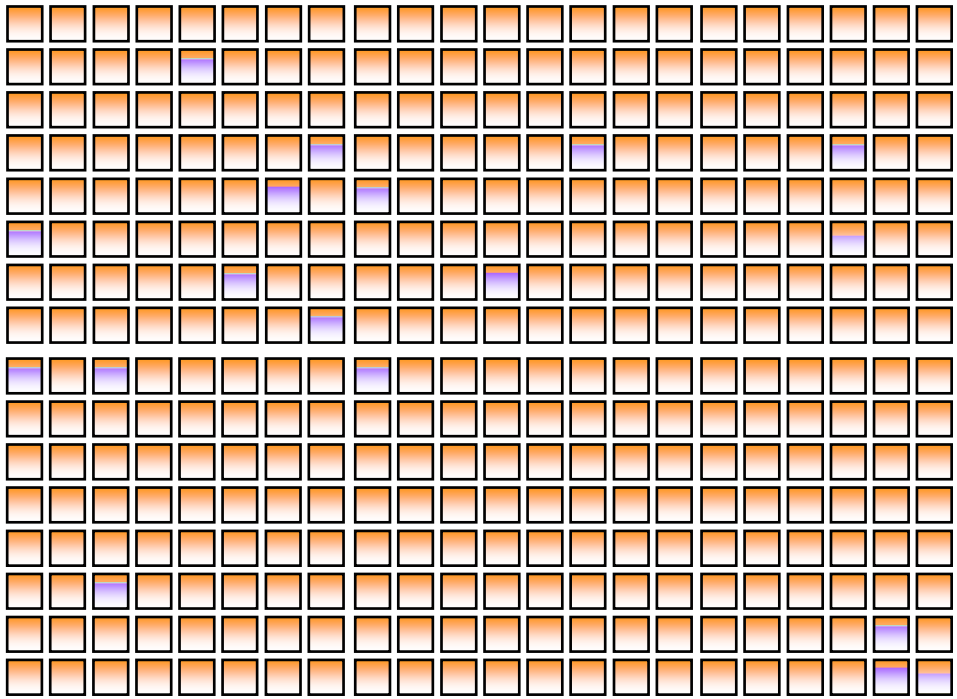

Chip Design Affects Virtualization Capabilities

Replicated Server Chip Design

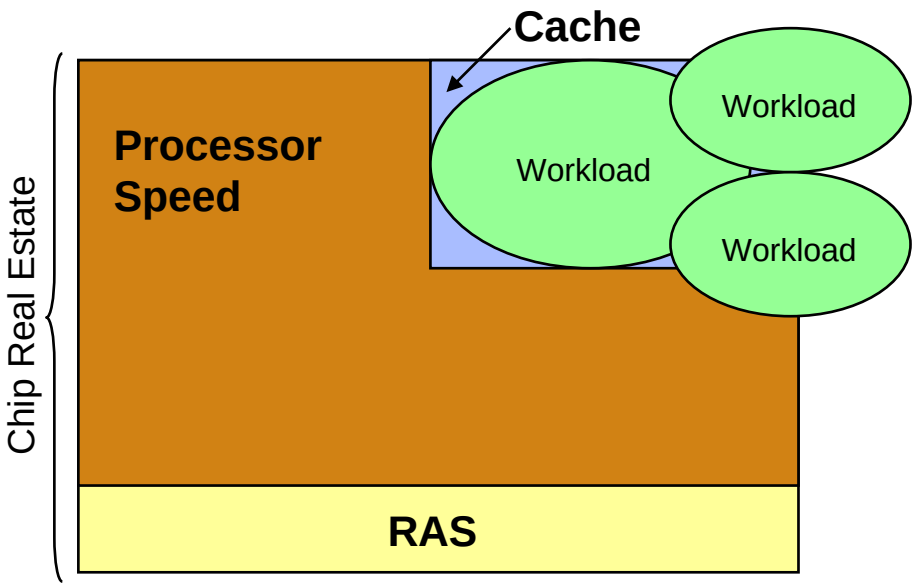

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- “Fast” processor speed is not fully realized due to cache misses

Consolidated Server Chip Design

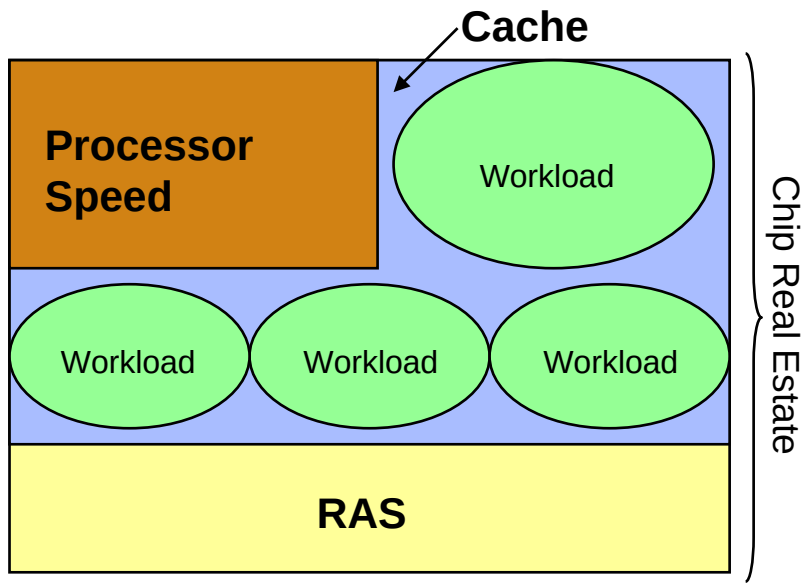

- System z cache is able to contain more working sets
- Processor speed is optimized by increased cache usage
- Additional RAS function is beneficial for mixed workloads

Scalability Considerations: Context Switching

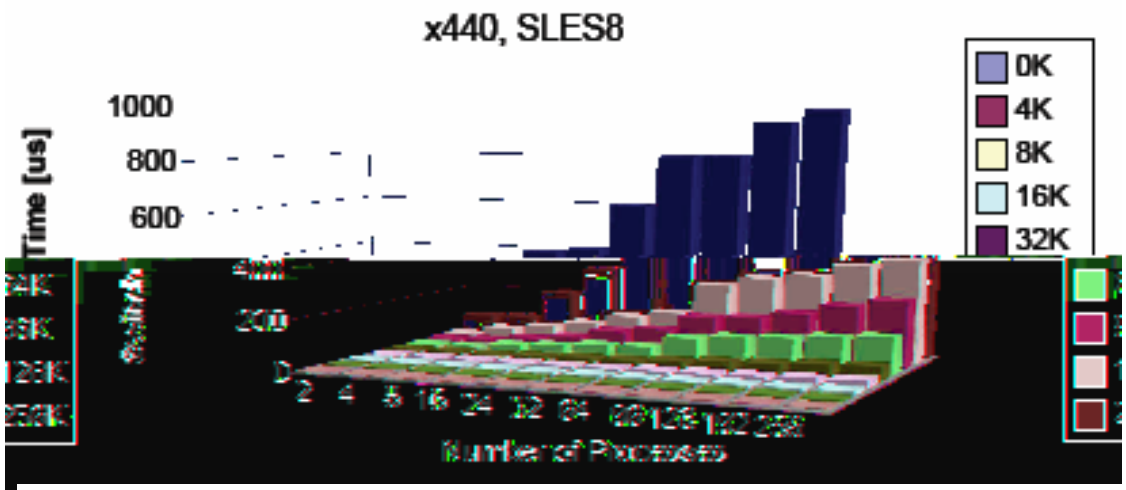

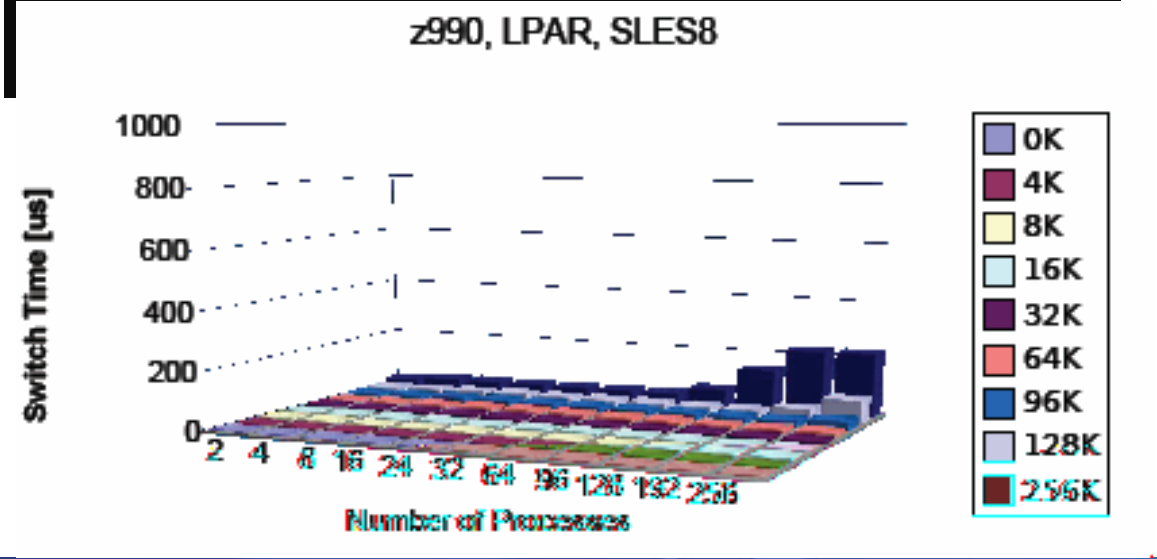

- Time required to perform context switching is an indication of memory time
- Virtualization, by definition, involves context switching

Datacenter in a Box

The Cost Benefit of Deep Integration with System z

Distributed-systems can suffer from:

- Cost and complexity (e.g., more physical servers, real network gear)
- Data silos and data synchronization
- Linear staffing costs
- Frequent outages

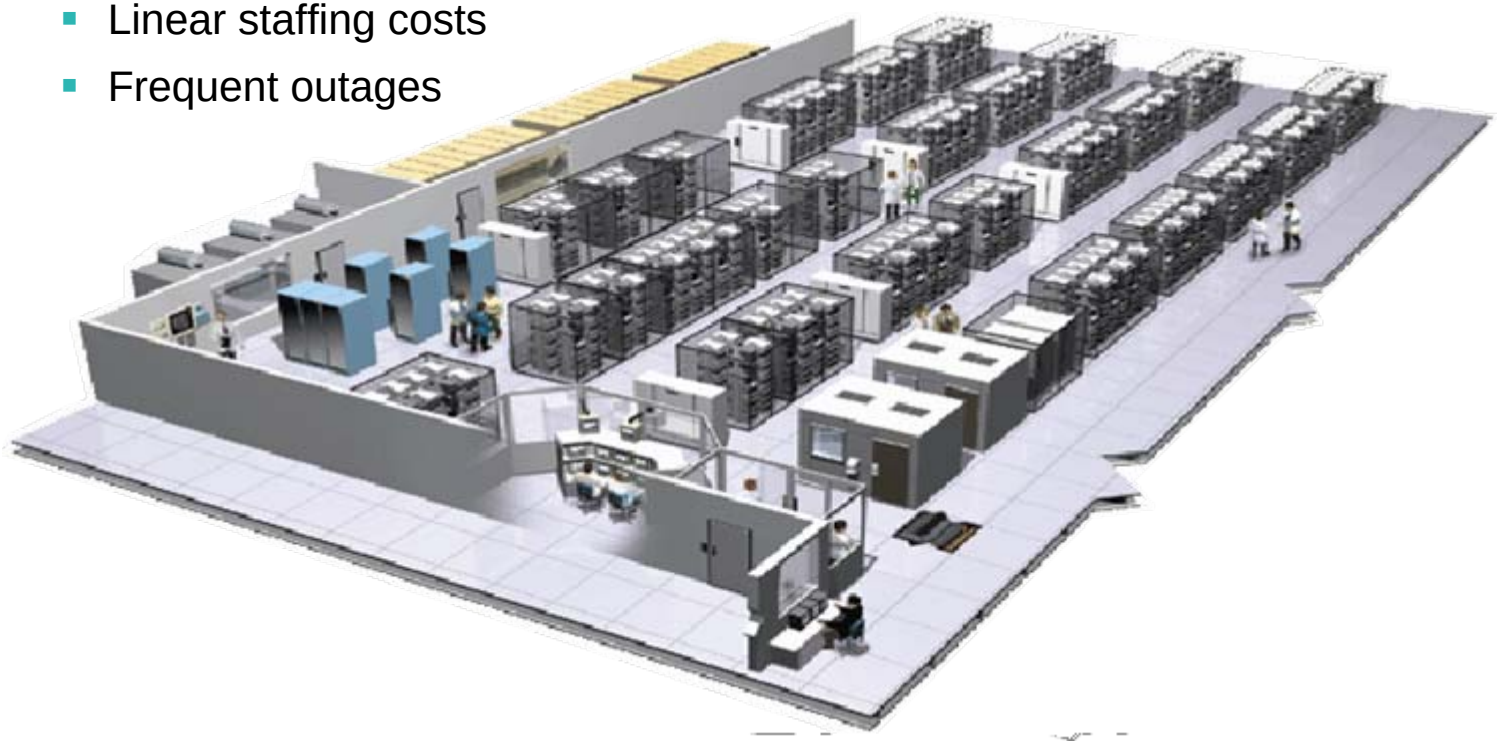

Datacenter in a Box

The Cost Benefit of Deep Integration with System z

Distributed-systems can suffer from:

- Cost and complexity (e.g., more physical servers, real network gear)
- Data silos and data synchronization
- Linear staffing costs
- Frequent outages

System z offers:

- Hundreds of processors
- Extremely large I/O bandwidth
- Built-in networking
- “Shared everything” resource model
- Billions of dollars in engineering and software development
- A pre-integrated datacenter in a box

IBM System p Virtualization Architecture

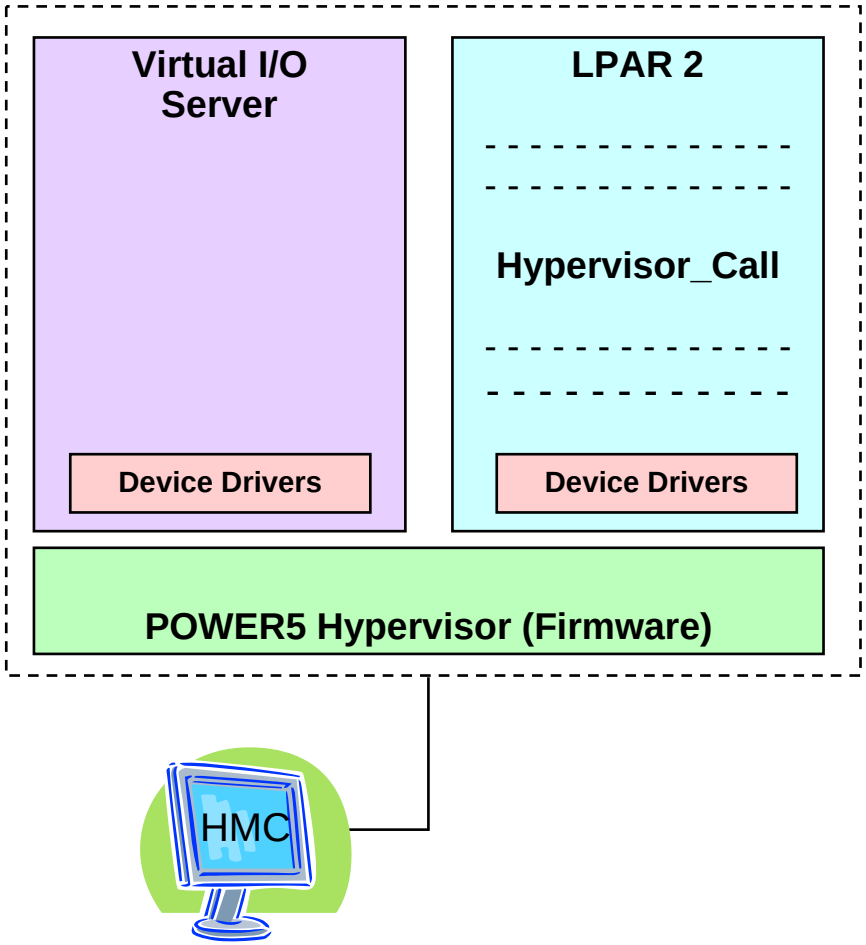

- **Performance**
  - Architected Hypervisor calls
  - Hardware assists
  - Event-driven dispatching
  - Scalable from 0.1 to 64 CPUs
- **Reliability, Availability, Serviceability**
  - Hypervisor is core firmware
  - No I/O device or 3rd party code
  - Explicit functions
  - Multiple VIO servers supported
- **Isolation and Security**
  - Each LPAR has separate O/S and patch levels
  - POWER Hardware Hypervisor mode
  - I/O, memory, and cache isolation
  - Separate HMC administrative users

IBM System p Virtualization

POWER5 CPU and Memory Virtualization

- **Processors**
  - Dedicated or shared processors
  - Fine-grained allocation of resources
  - Shared processors*
    - Entitlement
    - Virtual processors
    - Capped and uncapped
    - LPAR Weights
  - Adjustable via DLPAR
- **Memory**
  - From 128MB to all of physical memory
  - Dedicated physical memory
  - Adjustable via DLPAR (AIX and i5/OS)
- **Capacity On-Demand**

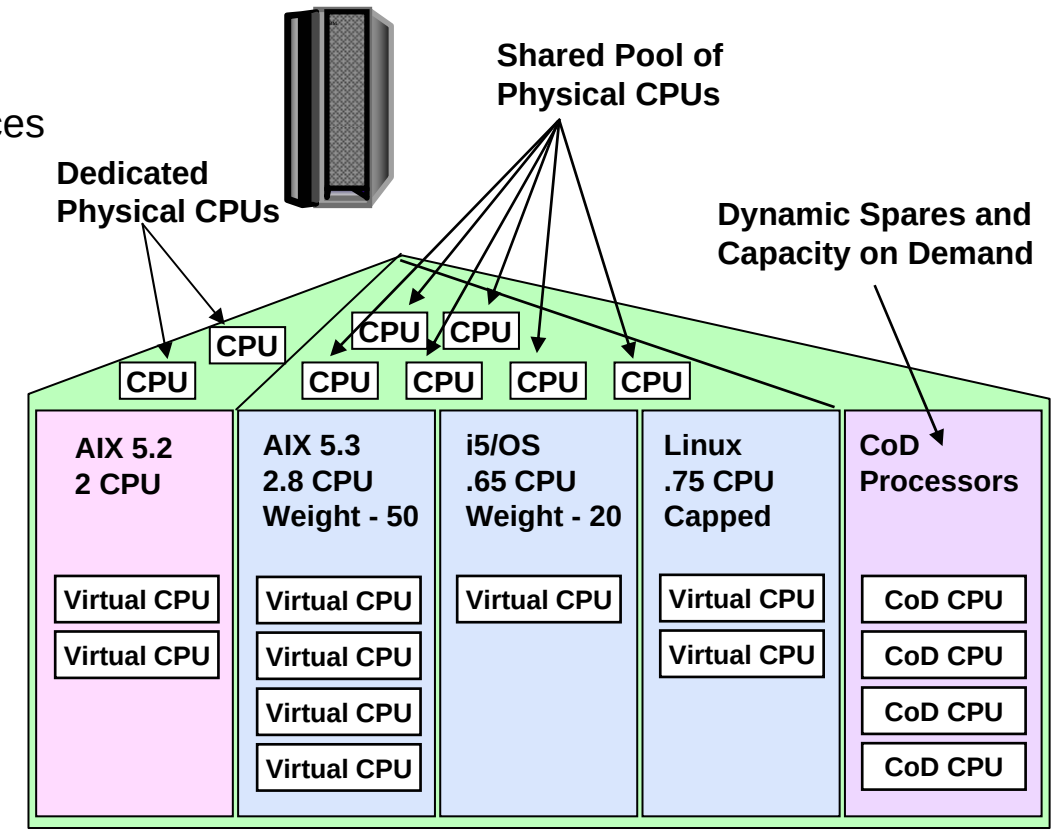

* Available on System p and OpenPower servers via the Advanced POWER virtualization features

IBM System p Virtualization
POWER5 I/O Virtualization

- **Virtual I/O Architecture**
  - Mix of virtualized and/or real devices
  - Multiple VIO Servers* supported
- **Virtual SCSI**
  - Virtual SCSI, Fibre Channel, and DVD
  - Logical and physical volume virtual disks
  - Multi-path and redundancy options
- **Benefits**
  - Reduces adapters, I/O drawers, and ports
  - Improves speed to deployment
- **Virtual Ethernet**
  - VLAN and link aggregation support
  - LPAR to LPAR virtual LANs
  - Shared Ethernet Adapter Failover

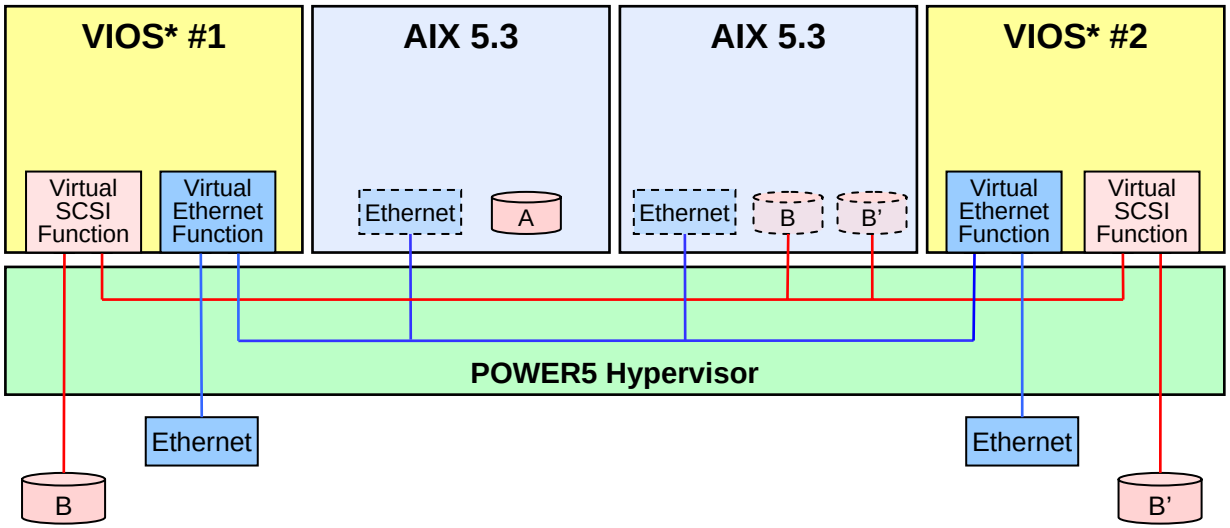

* Available on System p and OpenPower servers via the Advanced POWER virtualization features

Intel/AMD Server Virtualization Landscape
Hypervisors

- **VMware ESX Server (proprietary)**
- **Xen (open source Linux)**
  - Distributed by Novell (SLES 10) and Red Hat (RHEL 5)
  - Marketed by “Virtual Iron Software” and “XenSource” too
- **Microsoft® Windows® (proprietary)**
  - Virtual Server 2005 – Longhorn (2007) – Viridian (2007-2008)
- **Sun (open source Solaris)**
  - Containers (“zones”) – can only host copies of the underlying Solaris operating system
  - “BrandZ” – future development project (“non-native” support for Solaris zones)
- **Many more...and that can be a problem**

The clear leader in this field is VMware, but it has seen some competition from Microsoft, and now XenSource Inc. and Virtual Iron Software Inc. are pushing to make Xen the open source competitor.

– “Sun to provide Xen virtualization on Solaris”, Mark Fontecchio, SearchDataCenter.com, 05 Jul 2006

“KVM, which was presented to the public only barely two months ago, thereby easily overtakes other virtualization solutions such as Xen, OpenVZ and Vserver, which are based on other approaches, on the path toward integration into the kernel.

– Heise Online, www.heise.de/english/newsticker/news/82344, 11 Dec 2006

Intel/AMD Server Virtualization Landscape

Systems Management

- **VMware VirtualCenter**
- **Virtual Iron Virtualization Manager**
- **XenSource XenEnterprise**
- **Microsoft System Center Virtual Machine Manager**

VMware ESX Product History

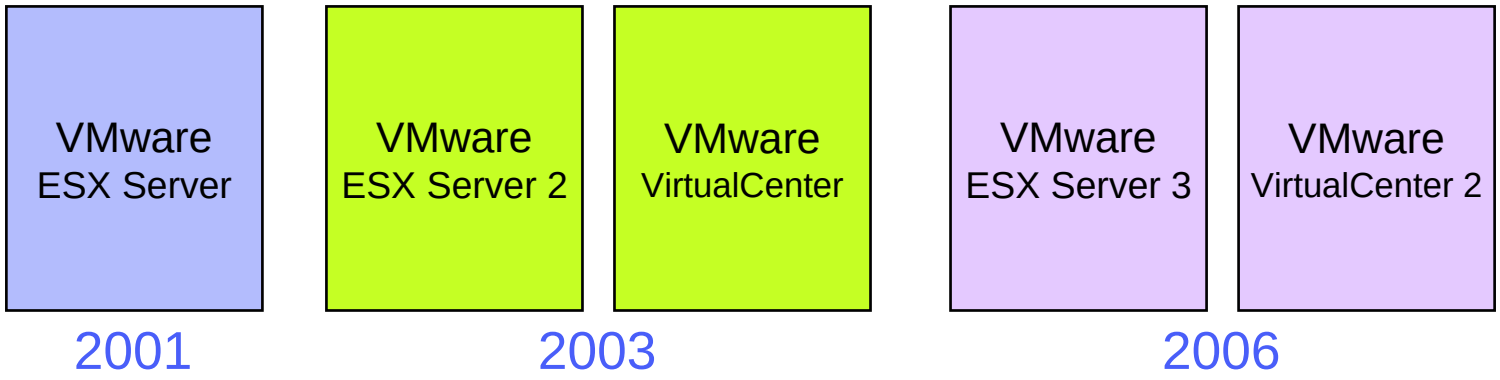

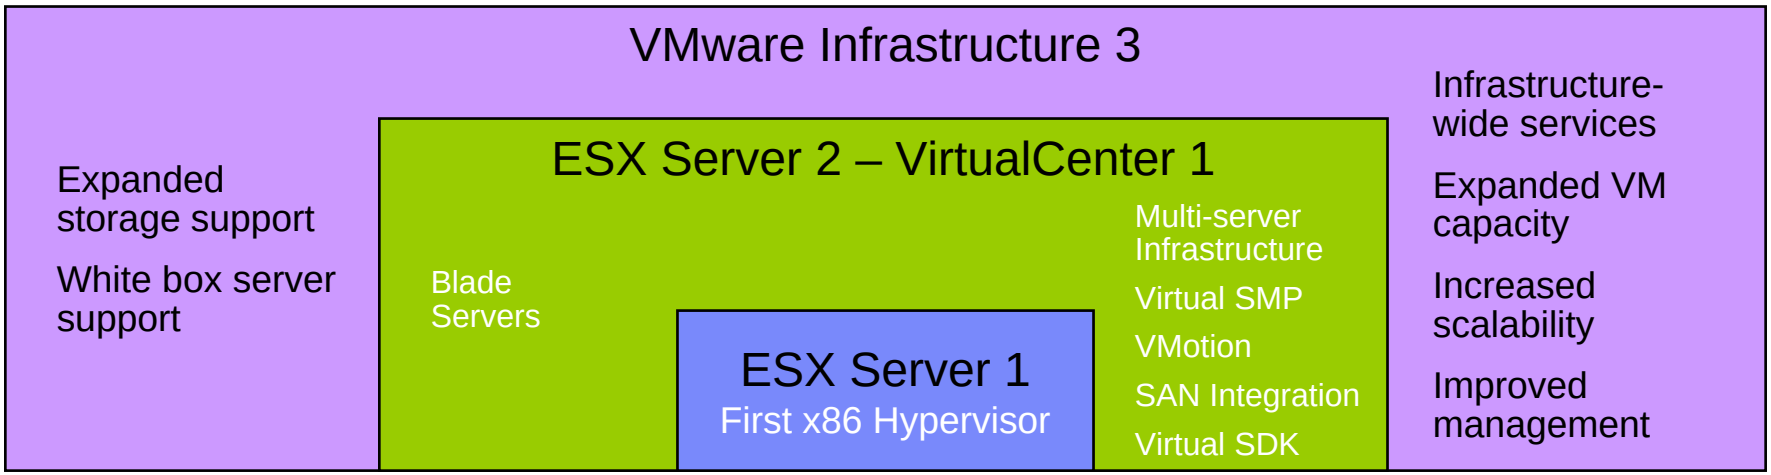

x86 Partitioning with VMware ESX Server 3

Part of VMware Infrastructure 3

- **VMware ESX Server 3 runs directly on the hardware**
  - No hosting operating system required (as is the case with VMware GSX Server)
- **Creates multiple virtual machines on a single Intel system**
  - Supports a maximum of 128 virtual machines per VMware image (depending on system resources)
- **Manages resource allocations among virtual machine images**
  - Strong fault and security isolation (uses CPU hardware protection)
  - Virtual networking support available (MAC or IP addressing)
  - Direct I/O passthrough
- **Shared data cluster-ready**
- **VMware File System (VMFS) allows multiple virtual disks to be stored on a single LUN or partition**
- **Virtual machines are encapsulated**
  - Which enables relocation of virtual machine from one copy of VMware to another (VMotion)
- **Add-on products available**
  - VMware Virtual SMP allows virtual machines to be configured with up to 4 CPUs
  - VMware VirtualCenter

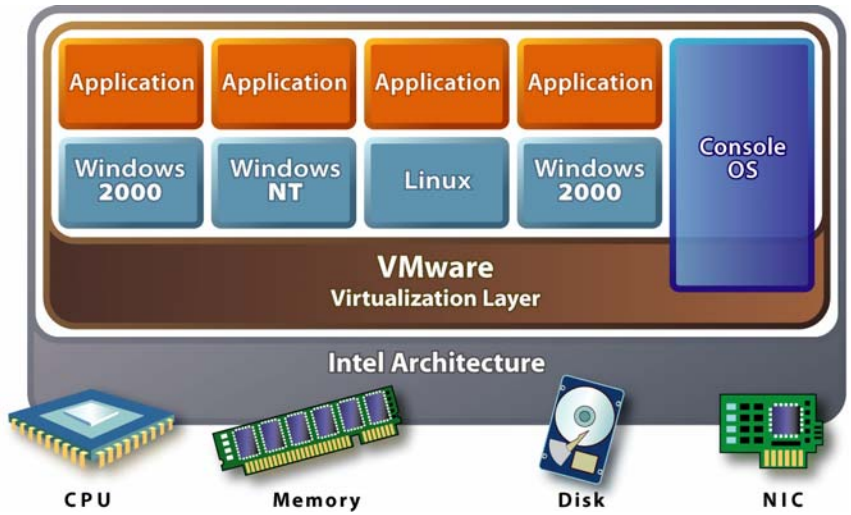

VMware VirtualCenter 2

Part of VMware Infrastructure 3

- **VirtualCenter Management Server**
  - The control node for configuring, provisioning and managing multiple VMware ESX servers
  - Runs on Windows 2000, Windows XP Pro, or Windows Server 2003
- **VirtualCenter Agent**
  - Needed to connect ESX servers with the Management Server
- **Virtual Infrastructure Client**
  - Allows users to connect to the VirtualCenter Management Server or individual ESX servers from any Windows PC
- **VirtualCenter Database**
  - Stores information used by the Management Server
  - Requires Oracle, Microsoft SQL Server, or Microsoft MSDE
- **Virtual Infrastructure Web Access**
  - Allows virtual machine management and access to graphical consoles without installing a client

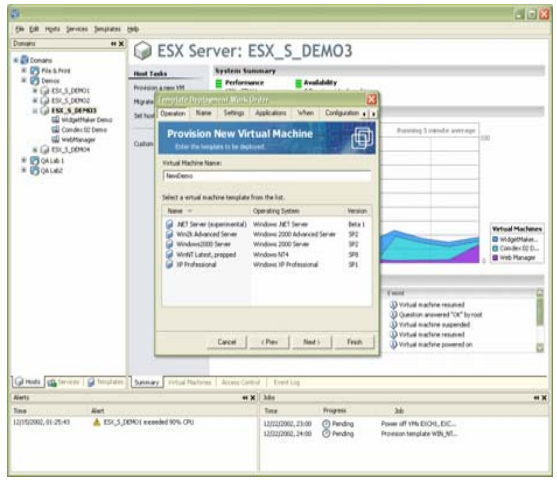

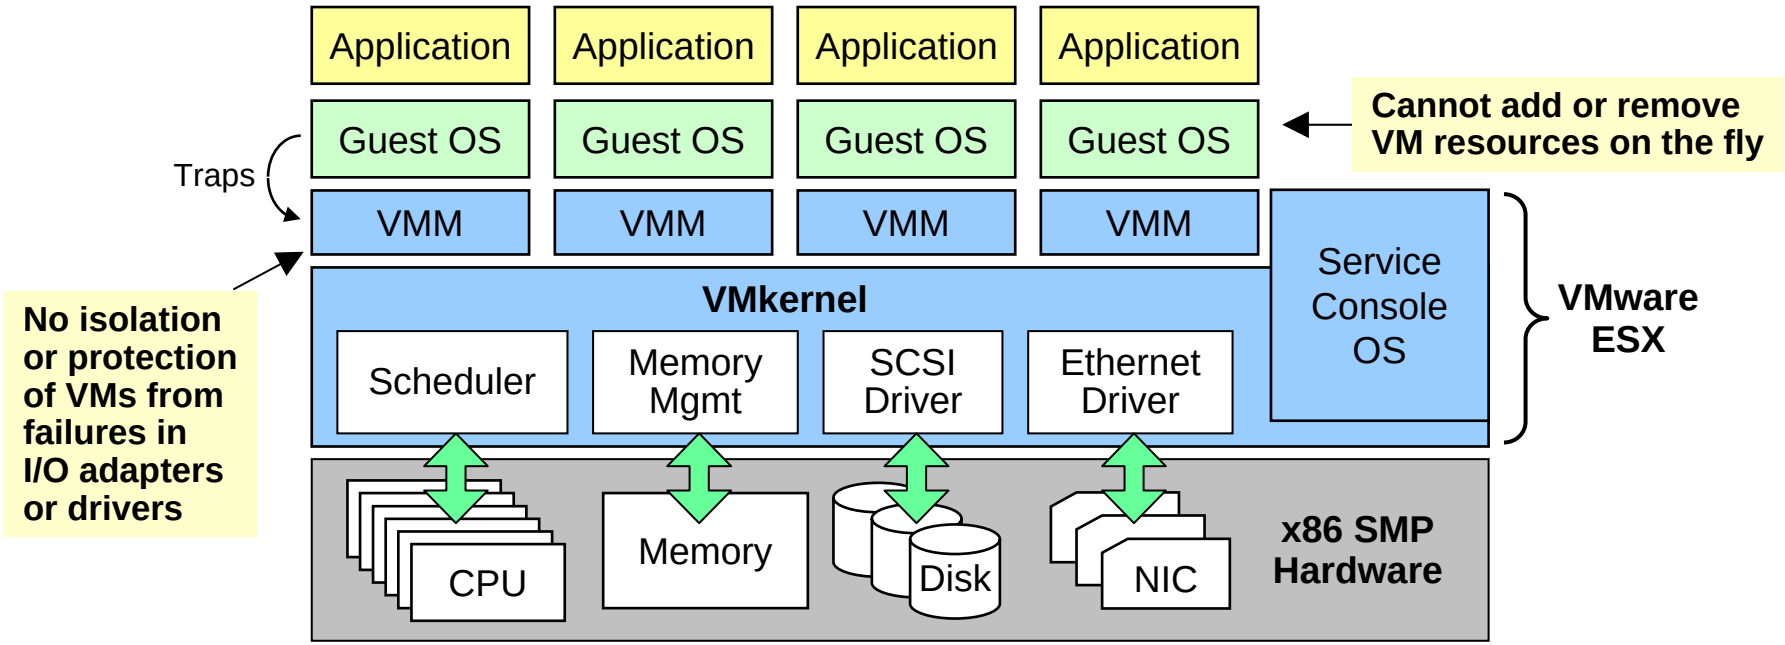

- **Supports a wide range of unmodified Windows and Linux versions**
- **Guest OS runs in User Mode; privileged instructions trap to Virtual Machine Monitor (VMM)**
  - "Trapping and mapping" is a significant source of performance overhead
- **Guest OS binary code is translated incrementally at load time**
  - Instructions that behave differently in User Mode vis-à-vis Supervisor Mode must be replaced with explicit trap instructions so the appropriate behavior will occur
- **Modified Linux device drivers run in the VMkernel**
- **Intel VT hardware feature is used only for 64-bit guests**

System z and z/VM Virtualization Leadership

Scalability, Granularity, Efficiency

- **IBM mainframes running z/VM may offer single-system scalability advantages over VMware systems**
- **z/VM enables a highly granular sharing of system resources with high levels of efficiency and resource utilization**
- **Capability of adding system resources to mainframe systems "on the fly" without disrupting hosted workloads**
- **Potential resulting business value:**
  - z/VM users can quickly and easily add virtual servers, or reconfigure existing servers, on a system that is already highly utilized
  - z/VM users can non-disruptively add system resources and grow their virtual server environment without requiring additional floor space, cabling, or previously-purchased hardware "waiting to be provisioned"
  - z/VM data-in-memory techniques can increase the operational and cost advantages of growing virtual server workloads on a single copy of z/VM
  - Resource-intensive Linux workloads can be hosted in z/VM virtual machines

System z and z/VM Virtualization Leadership
Integrated Technology Stack

- **Mainframe virtualization enhancements can be delivered where they belong**
  - In hardware and/or firmware
  - In the hypervisor, operating system, or application layers
  - All of the above for some technology advances
  - Enables a coordinated and timely delivery of support for new hardware (e.g., servers, storage, networking)
- **System z virtualization users can receive support from a single vendor (IBM) for resolving issues and addressing functional requirements**
- **The combination of LPAR and z/VM on System z offers powerful options for:**
  - Workload management and isolation
  - System resource sharing
  - Business continuance (e.g., failover)
  - Application integration with z/OS
- **Co-residency of Linux-on-z/VM and z/OS can mean...**
  - Linux and z/VM systems can exploit z/OS technologies for enhanced qualities of service (e.g., GDPS/PPRC Multiplatform Resiliency for System z)
  - z/OS users can deploy applications and services on Linux that may more easily integrate with their existing mainframe operations and infrastructure

System z and z/VM Virtualization Leadership

Scale Out and Scale Up Support with z/VM

- **Some Linux workloads may be constrained when running on VMware, or may not be supported due to resource requirements**
  - VMware ESX Server 3 allows no more than 4 CPUs per guest image, for an added fee, and only if 4 real CPUs exist on the underlying server
    - ☑ z/VM V5.2 supports up to 64 CPUs per virtual machine and 24 CPUs per z/VM LPAR (z/VM V5.3 supports up to 32 CPUs per LPAR)
  - VMware ESX Server 3 limits virtual machines to 16 GB of memory
    - ☑ z/VM V5.2 supports virtual machines with up to 128 GB of memory (z/VM V5.3 will offer considerable growth beyond 128 GB of memory)
  - I/O-intensive workloads may suffer considerable overhead with VMware
    - Timely interrupt processing may be a problem for VMware
    - ☑ Mainframe I/O bandwidth in a z/VM environment is very large and hardware assisted
  - VMware ESX support for 64-bit guests is new
    - ☑ z/VM has supported a mix of 32-bit and 64-bit guest images since 2001
  - VMware guests cannot add or remove resources without being re-booted
    - ☑ When properly configured, Linux-on-z/VM guests can non-disruptively add system resources such as CPUs, I/O devices, and network adapters (as can z/VM itself)

System z and z/VM Virtualization Leadership

Highly Balanced Utilization and Sharing of System Resources

- **The vast majority of installed Intel and AMD systems are not designed to support virtualization**
  - VMware has to use a “Trap and Map” method of hosting guest operating systems on these systems
  - This impacts the efficiency of the VMware hypervisor and introduces some degree of reliability concerns
- **New IVT and AMD-V chips improve CPU virtualization support**
  - But you have to purchase new hardware* to enjoy these advantages
  - And I/O virtualization support may still be an issue for IVT and AMD-V systems
- **In general, VMware users are advised to be careful over committing virtual memory to real memory, as well as CPU**

☑ **z/VM and System z are *architected* to host a large number of virtual machines that enable a balanced consumption of system resources**

* Intel/AMD users have to (re-)purchase new systems to exploit new CPU technologies; System z IFL features do not have to be (re-)purchased when clients migrate to new mainframes.

System z and z/VM Virtualization Leadership

Flexible and Functionally Rich Systems Management Tooling

- **Users can run z/VM as a guest of z/VM**
  - Added flexibility for test and verification, release-to-release migration support, user education and training
  - Beneficial for hosting disaster recovery solutions
- **z/VM offers an extensive suite of built-in tools and utilities for debug, problem determination, and system automation**
  - Trace and trap at the instruction level, with no modification of the guest operating system expected
  - Record and report resource utilization with a high degree of selectivity and frequency
  - Sniff virtual network traffic among guest systems
  - Automate system operations with PROP and SET OBSERVER
- **The Performance Toolkit for VM™ allows users to monitor system activity and capture resource consumption data to enable chargeback and capacity planning**
  - The Performance Toolkit also works with “OMEGAMON® XE for z/VM and Linux” for enterprise-level performance monitoring
- **The “z/VM Center” task of IBM Director provides fast and easy provisioning of virtual Linux server images on z/VM**

System z and z/VM Virtualization Leadership

Virtualization Adoption Issues for x86 Systems

- **Getting started with x86 virtualization is the easy part**
  - There is plenty of “head room” on a typical x86 server to run a proof of concept
  - Even on a small scale, x86 consolidations can save money
  - But large-scale adoption of x86 virtualization technology has issues...
- **x86 virtualization solutions are relatively immature**
  - Chargeback can be a challenge
  - Capacity planning is difficult to accomplish
  - There is little opportunity to consolidate I/O-intensive workloads
  - There is no LPAR alternative for performance-critical workloads
- **x86 virtualization does not eliminate the need for discrete servers or the need to manage them**
  - x86 virtualization actually adds a layer of complexity to managing distributed servers
- **Given the limited scalability of x86 systems, server sprawl will continue**
  - Further fueled by the mindset of: “all I want is another virtual machine”

System z and z/VM Virtualization Leadership

Virtualization Adoption Issues for x86 Systems (continued)

- **What technology is the right technology?**
  - ESX? Virtual Server? Xen? KVM? Linux-VServer? QEMU? Integrity Virtual Machines? Virtuozzo? Containers?
- **Will your x86 virtualization technology provider be the right strategic partner three years from now?**
  - If not, then what?
- **Are you ready to re-purchase your x86 servers when the next "must have" chip technology is released?**
- **Will you discover software licensing issues for technologies like VMotion?**
  - Is my software licensed for the machine I want to move it to?
- **Will Microsoft restrict its software from supporting your virtualization solution?**
  - See: www.vmware.com/solutions/whitepapers/msoft_licensing_wp.html as an example

z/VM Virtualization Leadership Support

- **High levels of RAS built into the hardware**
- **Non-disruptive On/Off Capacity on Demand capability**
- **Linux and z/OS application integration**
- **Highly granular allocation of hardware assets**
  - Add “small” server images to existing configuration with minimal impact to other server images expected
- **Large-scale server hosting**
  - Potentially hundreds of server images
- **Resource consumption recording / reporting**
  - Capture data at hypervisor level (CP Monitor)
  - Useful for charge-back, capacity planning, problem determination, and fix verification
- **Hot stand-by without the hardware expense**
  - Idle backup images ready to run (or be booted) if primary servers fail
- **Autonomic, non-disruptive disk failover to secondary storage subsystem capability**
- **Architecture simulation**
  - Help satisfy configuration requirements without necessarily suffering expense of real hardware
- **In-memory application sharing**
  - Share program executables among multiple server images
- **Server-memory-cached disk I/O**
  - High-speed read access to files on disk
- **Virtual Disks in Storage**
  - High-speed read and write access to files in memory (excellent swap devices for Linux)
- **Built-in console message routing**
  - Route messages from all virtual servers to a single virtual machine (system automation)
- **Virtual Machine Resource Manager**
- **“Hands free” auto-logon of server images**
  - Using z/VM “Autolog” support
- **Initiate operating system shutdown from “outside” the server image**
  - Without requiring agent running on guest operating system
- **Up to 256 Linux servers can share a single System z cryptographic card using z/VM**
- **Clone, patch, and “go live” with easy rollback**

When Do You Need More Than "Good Enough"?

Making the Case for Mainframe Virtualization

- **When workload growth and decline is difficult to predict (be it production, development, or test and assurance systems)**
- **When business results suffer as a result of IT resources not matching customer demand**
- **When your IT staff wants to optimize their productivity for deploying and managing virtual servers**
- **When innovation is stifled because your staff cannot experiment or develop new solutions using existing resources**
- **When speed to market affects your business results**
- **When your server applications need fast and flexible access to z/OS data and applications**
- **When business continuance is a high priority**
- **When you want more control over your environmental expenses (e.g., floor space, cooling)**

Linux and z/VM on System z

Business Value Propositions for Linux Workloads

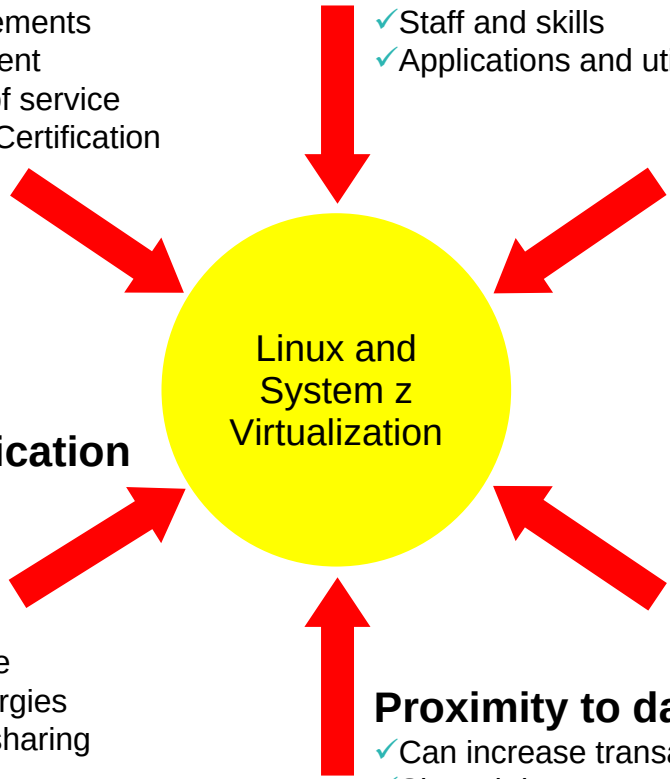

Security capabilities:
- ✓Privacy
- ✓Regulatory requirements
- ✓Identify management
- ✓zSeries qualities of service
- ✓Common Criteria Certification
- ✓Ethical hacking
- ✓HiperSockets

Consolidation capabilities:
- ✓Server, network, storage
- ✓Staff and skills
- ✓Applications and utilities

Business resiliency capabilities:
- ✓High availability
- ✓Disaster recovery
- ✓Serviceability
- ✓Reliability
- ✓HyperSwap
- ✓XRC and PPRC

Operational simplification capabilities:
- ✓Virtualization
- ✓Simulation
- ✓Single point of control
- ✓Large single system image
- ✓z/OS similarities and synergies
- ✓Highly granular resource sharing

On demand infrastructure:
- ✓Scale up and scale out
- ✓Rapid server (de)commissioning
- ✓Idle servers don't consume system resources

Proximity to data:
- ✓Can increase transaction throughput
- ✓Shared data access
- ✓Integrated storage management
- ✓HiperSockets

Reference Material

Xen 3.0

- Open Source virtualization software solution based on Linux
- Uses paravirtualization to abstract CPU, memory, and I/O resources
- Guest operating systems are responsible for allocating and managing page tables
- Management and control software runs in Domain 0
- IVT and AMD-V will enable hosting of unmodified guest operating systems
  - For Windows support on Xen, users need IVT-capable hardware
- IBM is actively contributing to the Xen open source project

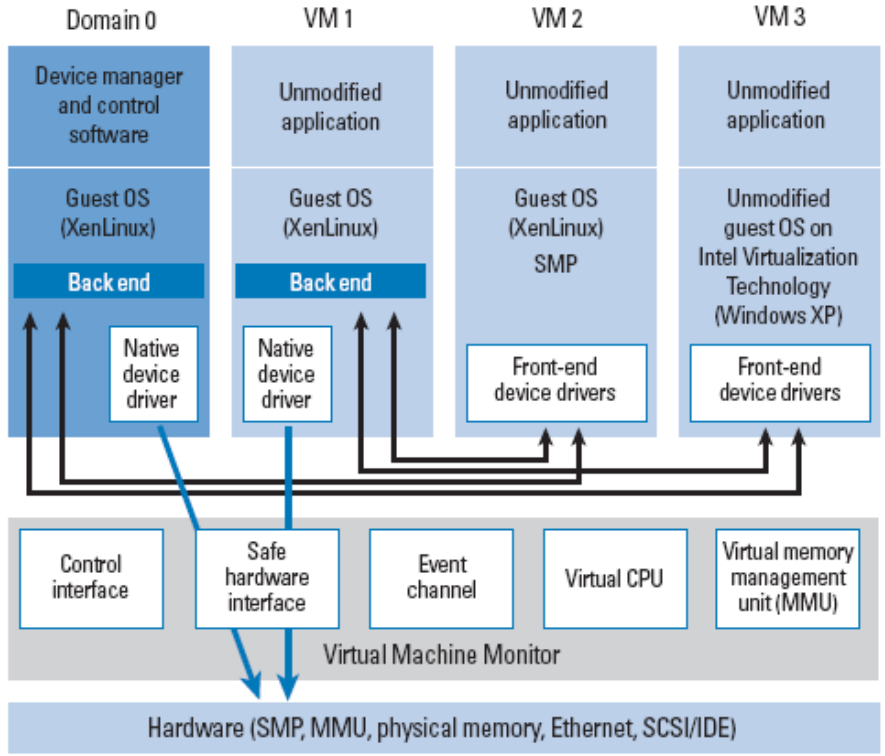

Figure 1. Xen 3.0 architecture: Hosting four VMs

XenSource XenEnterprise

"Delivering the Power of Xen"

- **XenEnterprise is a robust and easy-to-use bare metal virtualization platform for running Windows and Linux virtual machines**
- **The foundation of XenEnterprise is the open source Xen hypervisor**
  - Support for Windows virtual machines requires Intel VT or AMD-V hardware
- **XenEnterprise comes with a comprehensive management, monitoring and virtual machine deployment console**
  - Easy to use graphical administrator tool runs on Windows and Linux
- **Real time and trended reporting of virtual machine and total server performance metrics including CPU, memory, plus disk and network I/O**
- **Suspend and resume virtual machines on the fly**
- **Bundled and integrated physical to virtual (P2V) migration tools allow you to consolidate existing servers with ease**
- **Clone existing virtual machines for rapid system deployment**

Source: "XenEnterprise 3.1 Data Sheet", V11282006

Virtual Iron 3.0

- **Requires Intel Virtualization Technology (IVT) and AMD-V CPUs**
- **Replaces the proprietary hypervisor used in Virtual Iron 2.0 with Xen**
- **Virtual Iron Virtualization Services runs on top of Xen hypervisor**
  - Provides dynamic infrastructure management with hot-pluggable resources, server partitioning, and live migration of guest images
  - Supports Linux and Windows guest images (32- and 64-bit systems)
  - Up to 8 CPUs, 128GB RAM, 8 virtual NICs per virtual server
- **Virtual Iron Virtualization Manager provides command and control**
  - Java application that runs in any browser
  - Manage OS and application images (install, deploy, clone, etc.)
- **Virtual Iron 3.0 beta**
  - Available for Linux guests since July 2006
  - Available for Windows guests since September 2006

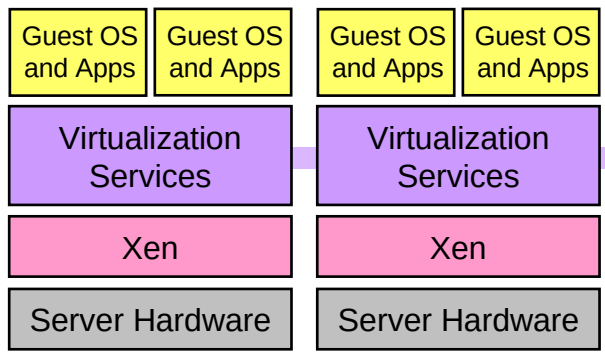

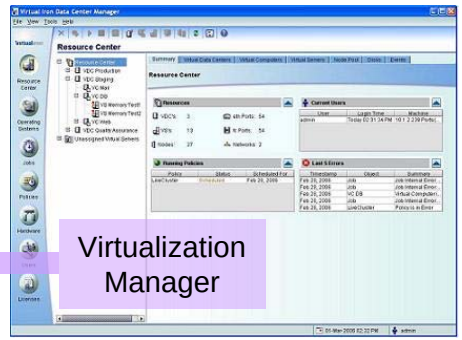

Microsoft Virtual Server 2005

- **Allows multiple operating systems to run simultaneously on the same processor**
- **Each independent virtual machine functions as a self-contained computer**
- **Virtual machines are encapsulated in portable Virtual Hard Disks (VHDs)**
  - Up to 32 VHDs can connect to a single virtual machine
  - VHDs can expand as data is added and “differenced”
  - VHD format is available royalty free
- **Virtual networking options are available**
- **Requires a hosting operating system (e.g., Windows Server 2003)**
- **Initially intended for hosting unsupported Windows NT environments**
- **Native Windows Server 2003 environment still recommended for many workload deployments**
- **Add-ins are available to support Linux**
  - Must be installed in the Linux guest images
  - Designed to improve guest-host interactions
- **Virtual Server 2005 is available as a free download**

Sun Containers (“Zones”)

- **Available with Solaris 10**
- **Similar in concept to z/OS address spaces**
- **Provides ability to host multiple OS instances, not virtual machines**
  - Private name, IP address, and port range
  - Can boot / re-boot zone image
  - Can create a new zone in minutes
  - Separate security, resource management, and failure scopes
  - Private file system options
    - “whole root” – zone owns private, R/W file systems
    - “sparse root” – zone inherits R/O /usr, /sbin, /platform, /lib
- **Instances share the same *Solaris* kernel**
- **Intra-zone networking (memory speeds)**
- **Can share binaries in RAM across zones**

HP Virtual Server Environment (VSE)

- **HP VSE consists of several product components**
  - HP Integrity Essentials Capacity Advisor
  - HP Integrity Essentials Virtualization Manager
  - HP-UX Workload Manager
  - HP Integrity Essentials Global Workload Manager
  - HP Systems Insight Manager
  - HP Integrity Virtual Machines

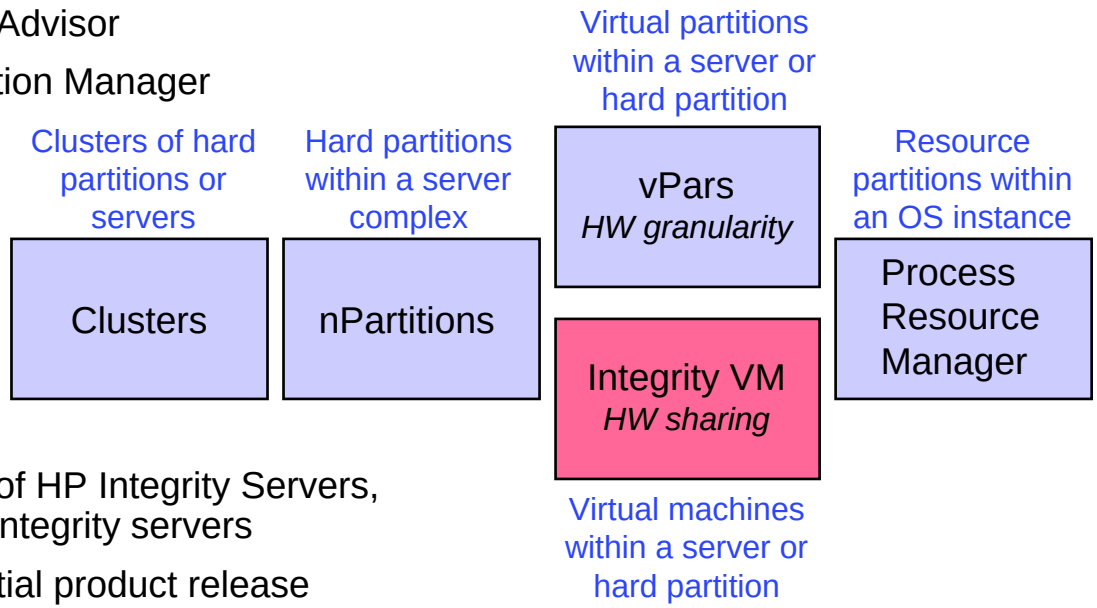

- **HP Integrity Virtual Machines**
  - Supports low- to high-end range of HP Integrity Servers, Integrity Superdome, and future Integrity servers
  - Support limited to HP-UX with initial product release
  - Supports Uni and SMP guest images with CPU sharing
  - Able to host different guest operating systems
- **"For the most demanding mission-critical workloads, the HP VSE for HP-UX 11i is unmatched"**

Sources: "HP Virtual Server Environment for HP Integrity and HP 9000 servers", August 2, 2005
"Introducing HP Integrity Virtual Machines", white paper from HP, September 6, 2005