European Conference on z/VSE, z/VM and Linux for System z – October, 2007

Virtualization Technology and Differentiation

Romney White
romneyw@us.ibm.com
System z Strategy and Design

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.

Abstract

IBM System z virtualization has a pedigree that spans more than 40 years of innovation, offering significant advantages for hosting virtual server workloads. This presentation will highlight those advantages and the underlying technology, with a high-level comparison to System p and System x.

Attendees will develop a better understanding of how System z virtualization capabilities offer higher levels of business value for hosting large-scale, Linux virtual server environments, and what differentiates the mainframe in the marketplace as a virtual server hosting platform.

Agenda

- Virtualization and hypervisor basics
- System z™ virtualization technology highlights
- System p™ virtualization overview
- x86 virtualization overview
- Virtualization comparisons
- System z strategic considerations

Affecting All Areas of Server Design and Deployment

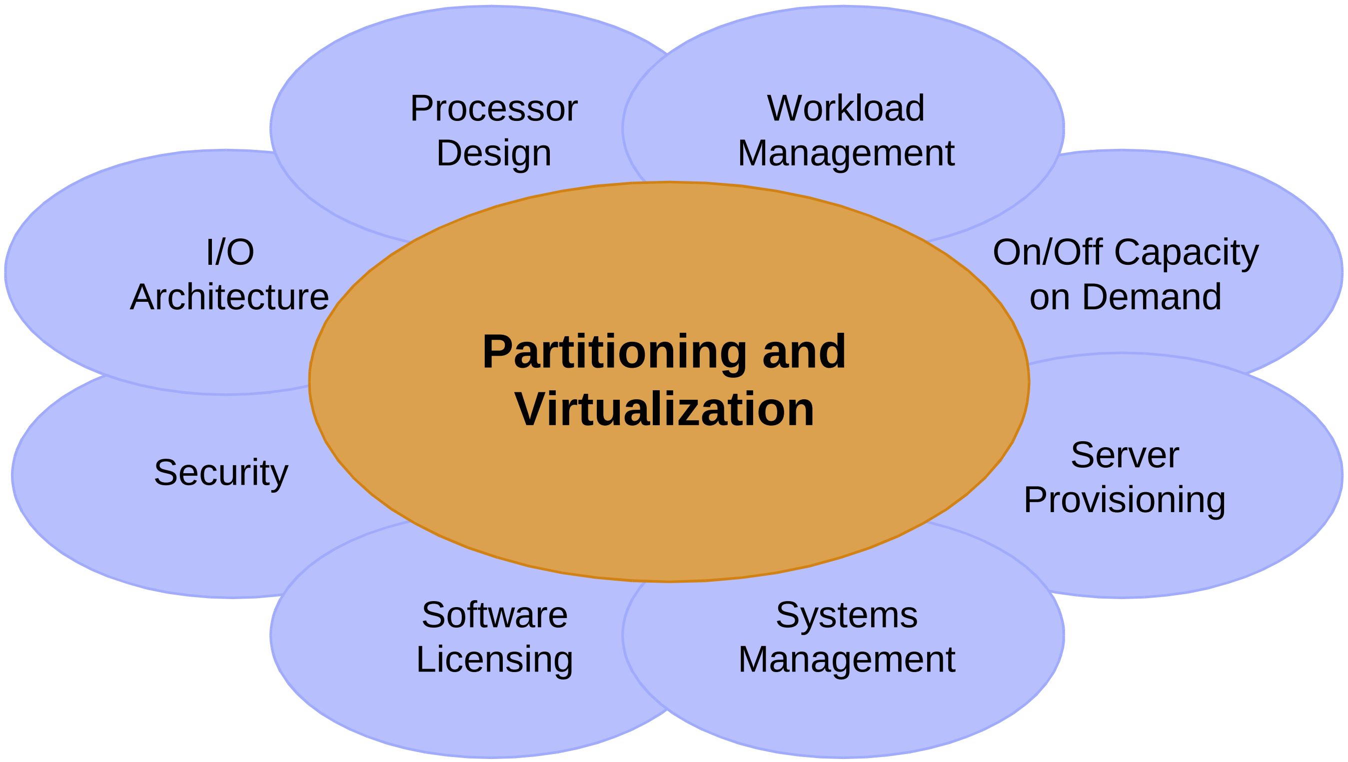

Virtual Resources

- 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

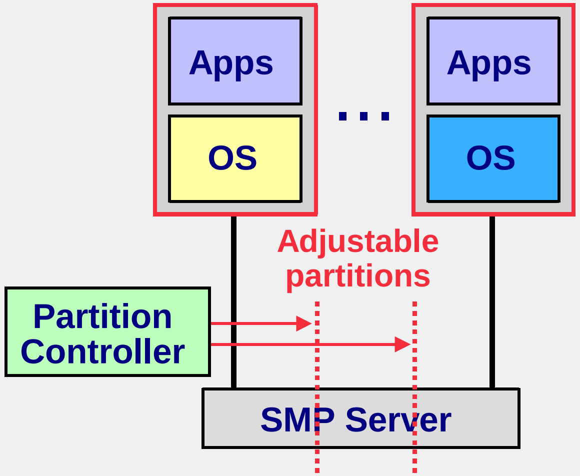

Server is subdivided into fractions each of which can run an OS

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

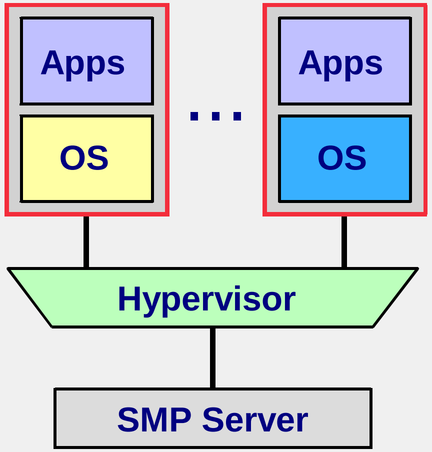

Hypervisor provides fine-grained timesharing of all resources

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

Hosted Hypervisor

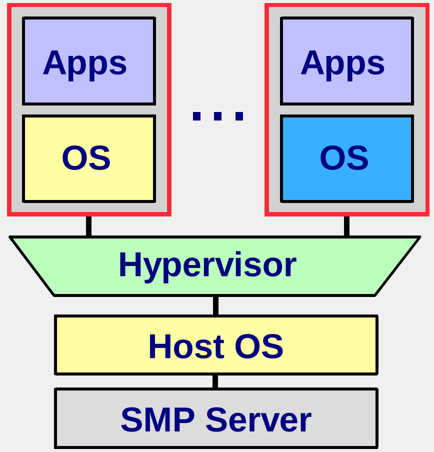

Hypervisor uses OS services to do timesharing of all resources

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

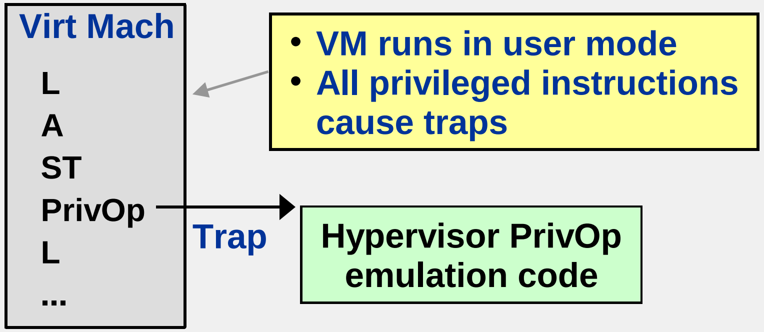

Examples CP-67, VM/370
Benefits Runs unmodified OS
Issues Substantial overhead

Translate, Trap, and Emulate

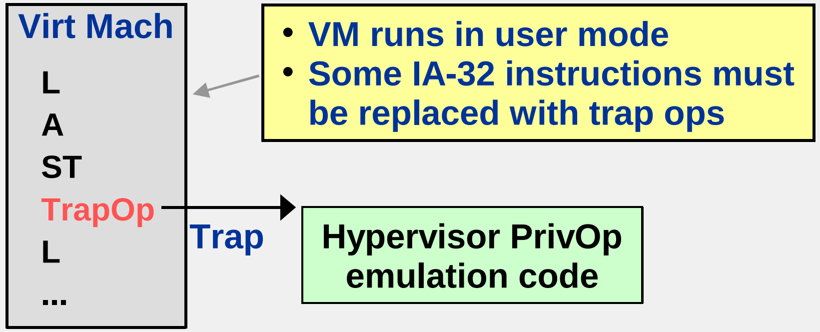

Examples VMware, Microsoft VS
Benefits Runs unmodified, translated OS
Issues Substantial overhead

Hypervisor Calls ("Paravirtualization")

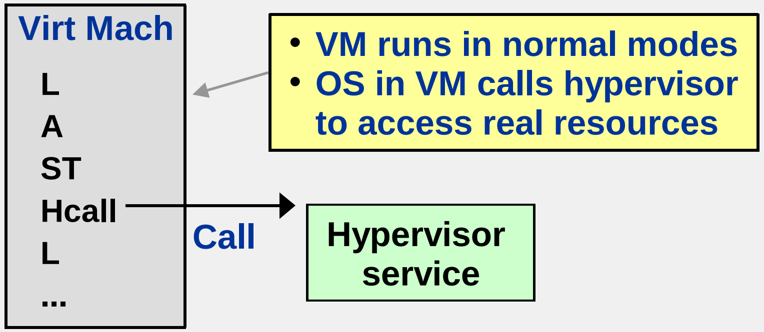

Examples POWER Hypervisor, Xen
Benefits High efficiency
Issues OS must be modified to issue Hcalls

Direct Hardware Virtualization

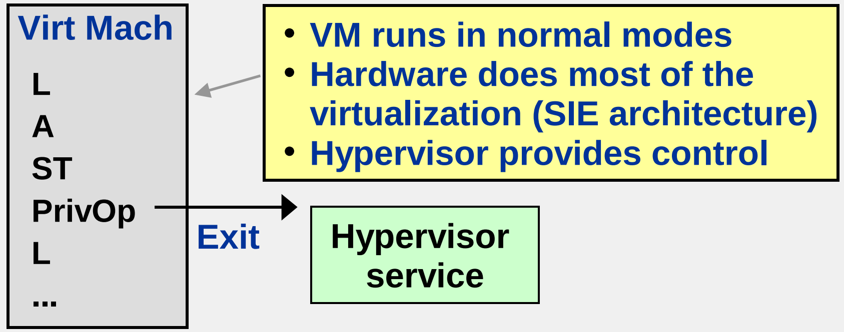

Examples PR/SM, z/VM, Xen
Benefits High efficiency, runs unmodified OS
Issues Requires underlying hardware support

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

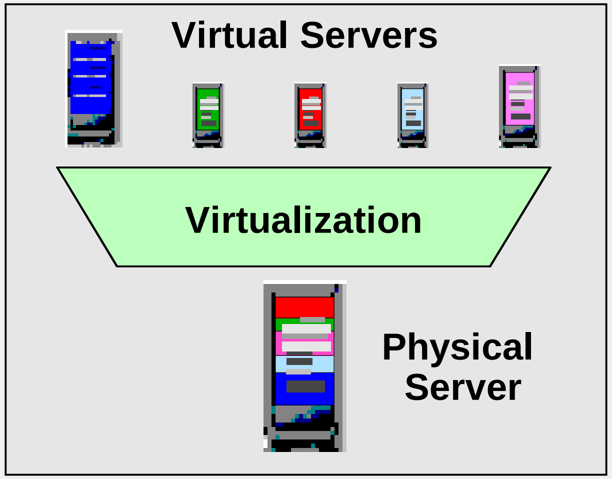

Possible Benefits:
- High resource utilization
- Great usage flexibility
- Enhanced workload QoS
- High availability / security
- Low cost of availability
- Low management costs
- Enhanced interoperability
- Legacy compatibility
- Investment protection

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

Virtualization is Built In, not Added On

On demand scale *up and out* solutions consist of multiple dimensions of function:

- **Application Support Dimension (open, stable)**
  - Open, stable operating system
  - Virtual server awareness infrastructure
  - Enterprise applications
- **Hypervisor Dimension Capabilities (powerful, flexible)**
  - Shared-memory based virtualization model
  - Granular resource sharing and simulation
  - Flexible virtual networking
  - Resource control and accounting
  - Server operation continuity (failover)
  - Server maintenance tools and utilities
- **Hardware Dimension Capabilities (robust, reliable)**
  - Legendary reliability, scalability, availability, security
  - Logical partitioning (LPAR)
  - Processor and peripheral sharing
  - Interpartition communication
  - Virtualization support at the hardware instruction level (e.g., SIE)

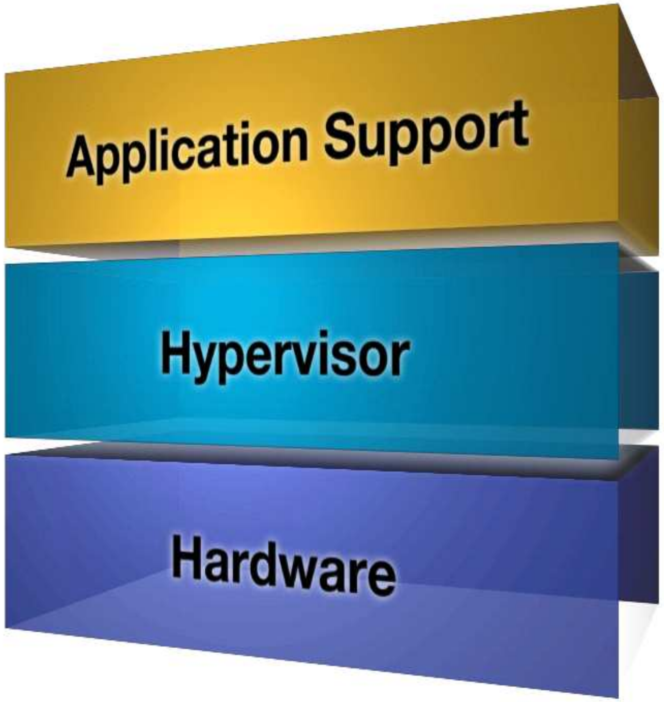

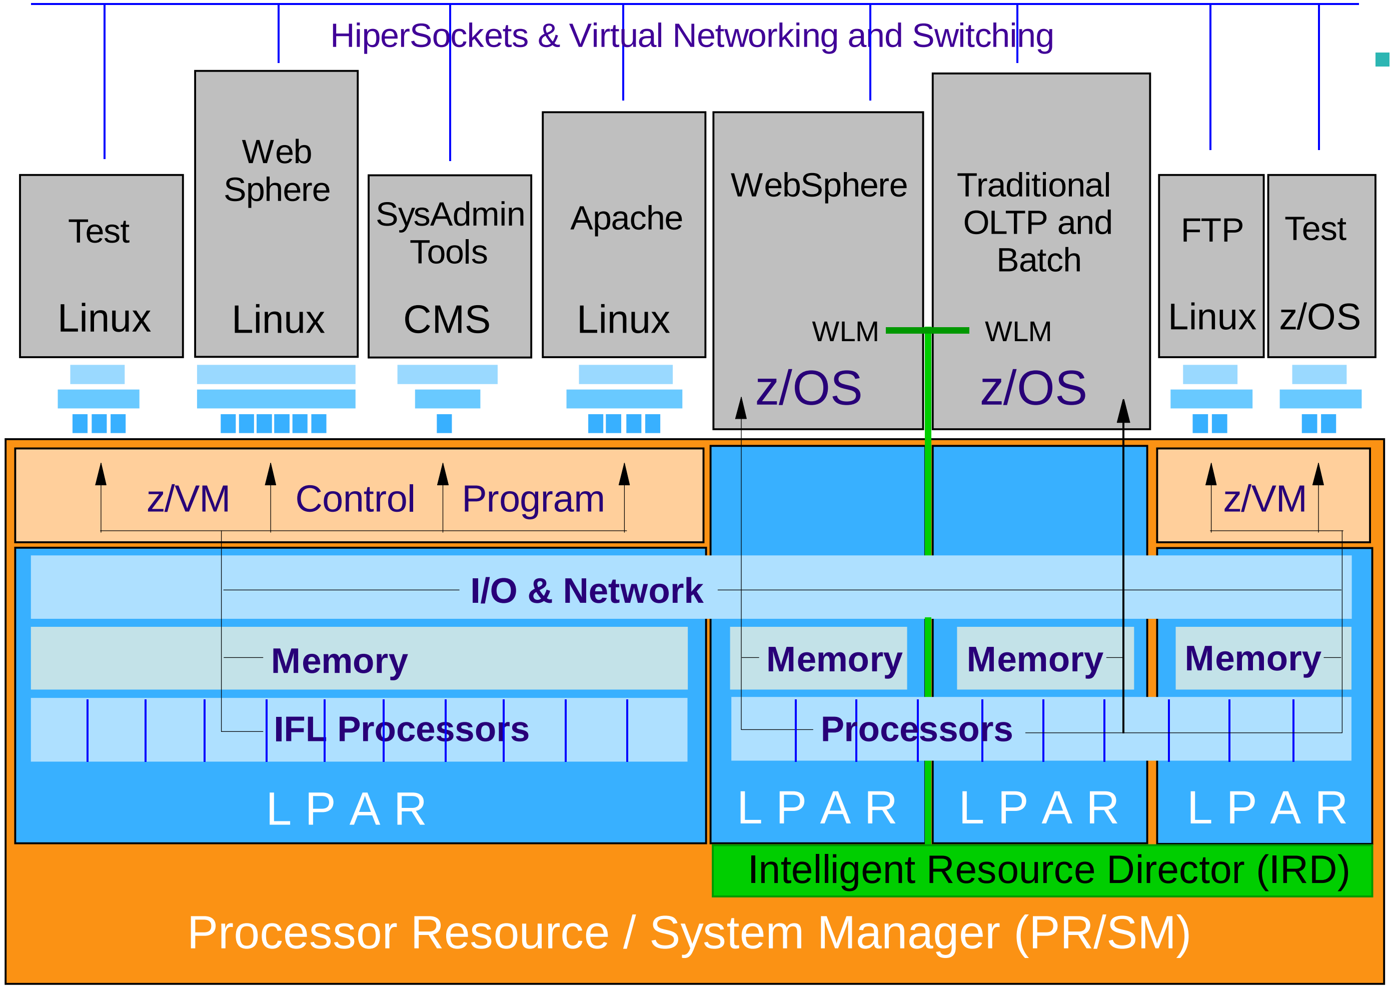

- Multi-dimensional virtualization technology
  - System z provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS®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

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

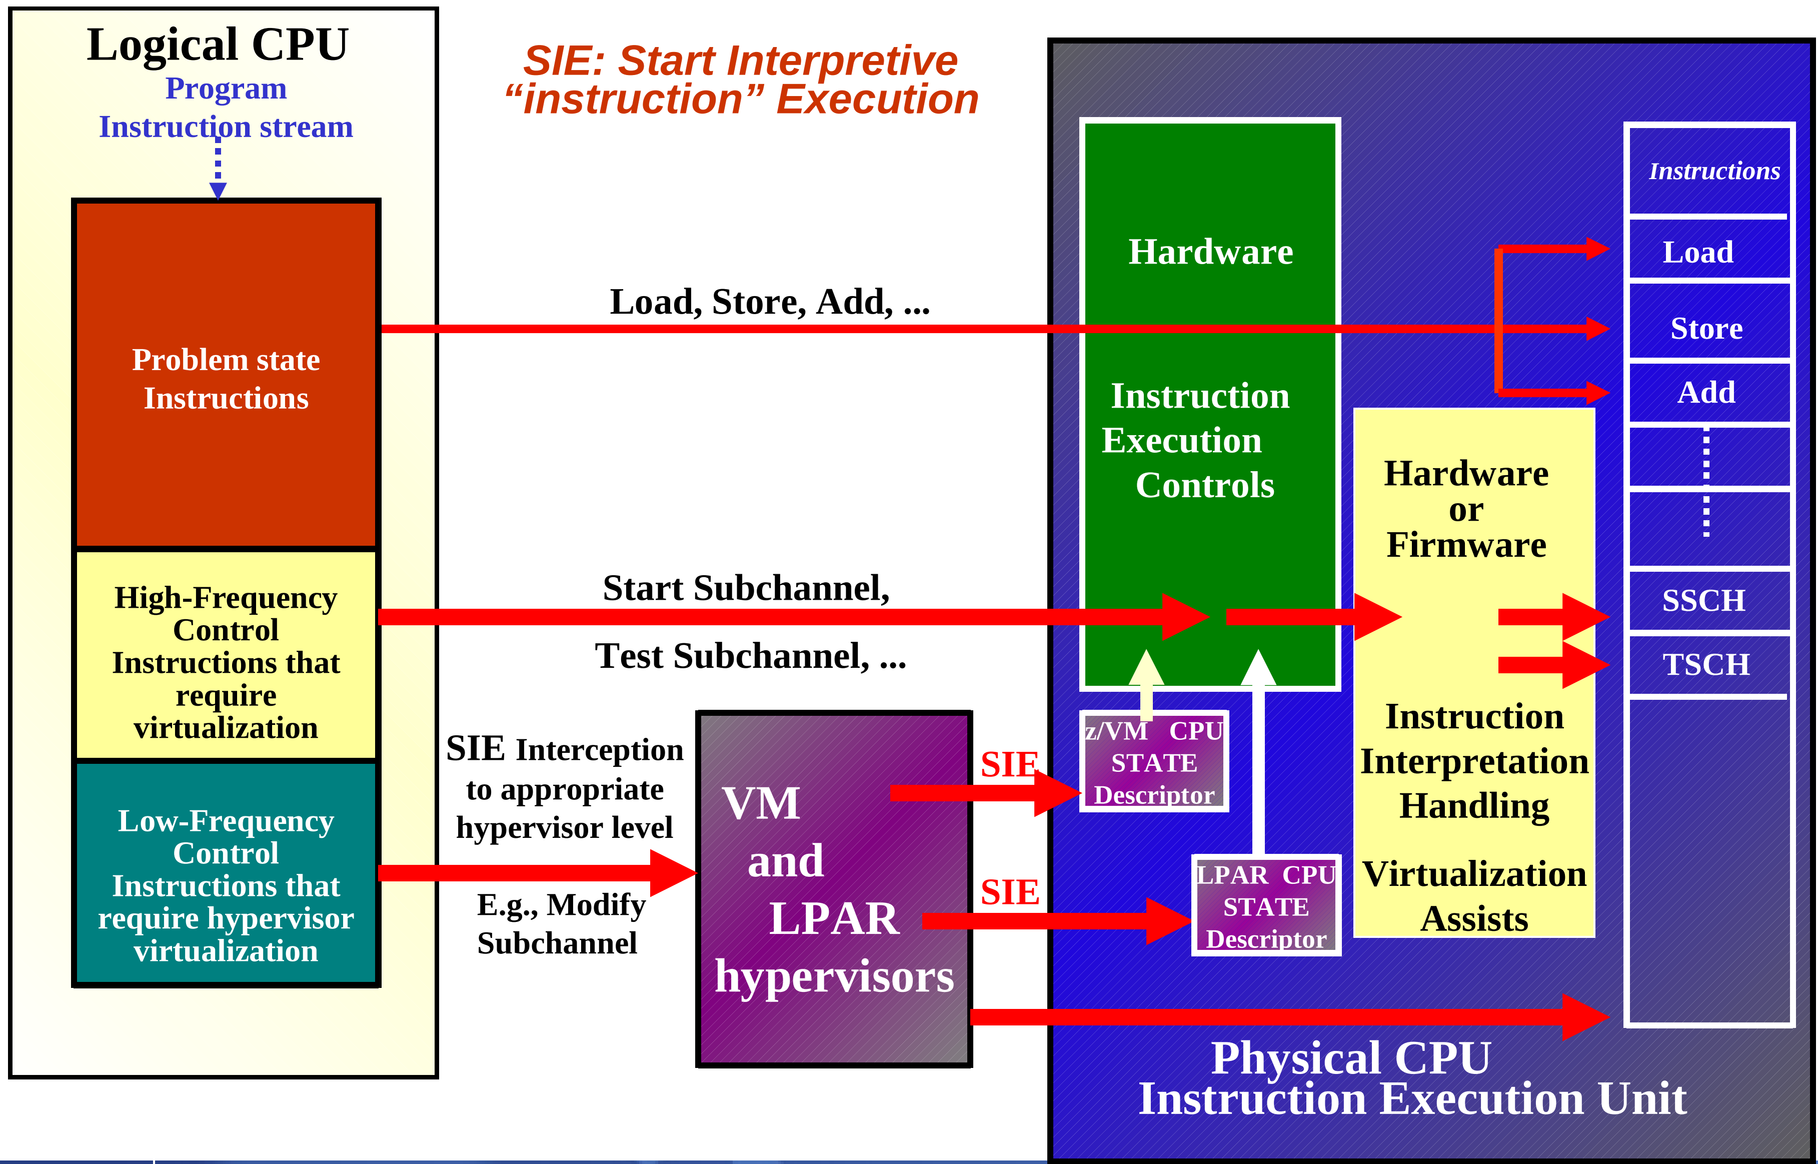
Logical CPU
Program
Instruction stream
Problem state Instructions
High-Frequency Control Instructions that require virtualization
Low-Frequency Control Instructions that require hypervisor virtualization
SIE: Start Interpretive “instruction” Execution
Load, Store, Add, ...
Start Subchannel,
Test Subchannel, ...
SIE Interception to appropriate hypervisor level
E.g., Modify Subchannel
VM and LPAR hypervisors
SIE
SIE
z/VM CPU STATE Descriptor
LPAR CPU STATE Descriptor
Hardware
Instruction Execution Controls
Hardware or Firmware
Instruction Interpretation Handling
Virtualization Assists
Instructions
Load
Store
Add
SSCH
TSCH
Physical CPU Instruction Execution Unit

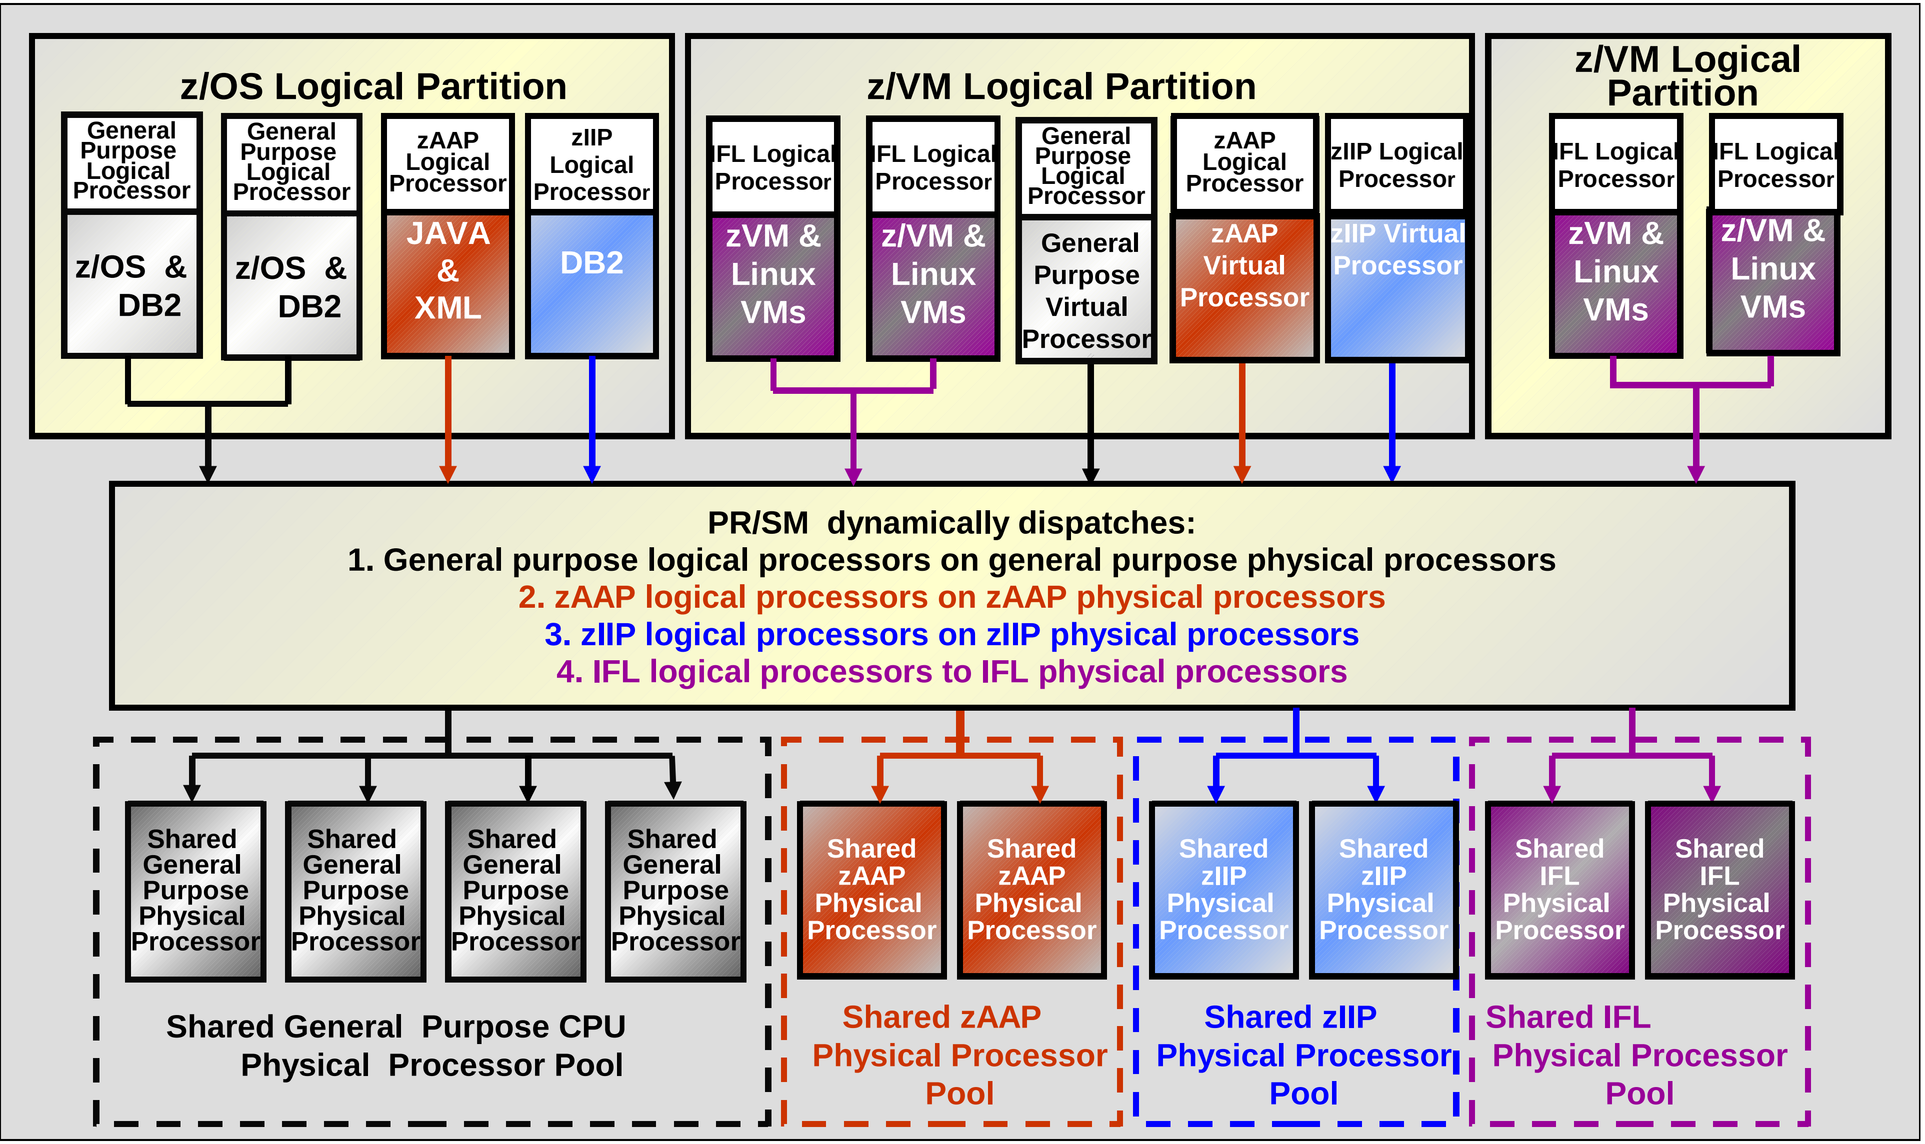
z/OS Logical Partition
General Purpose Logical Processor
z/OS & DB2
General Purpose Logical Processor
z/OS & DB2
zAAP Logical Processor
JAVA & XML
zIIP Logical Processor
DB2
z/VM Logical Partition
IFL Logical Processor
zVM & Linux VMs
IFL Logical Processor
z/VM & Linux VMs
General Purpose Logical Processor
General Purpose Virtual Processor
zAAP Logical Processor
zAAP Virtual Processor
zIIP Logical Processor
zIIP Virtual Processor
z/VM Logical Partition
IFL Logical Processor
zVM & Linux VMs
IFL Logical Processor
z/VM & Linux VMs
PR/SM dynamically dispatches:
1. General purpose logical processors on general purpose physical processors
2. zAAP logical processors on zAAP physical processors
3. zIIP logical processors on zIIP physical processors
4. IFL logical processors to IFL physical processors
Shared General Purpose Physical Processor
Shared General Purpose Physical Processor
Shared General Purpose Physical Processor
Shared General Purpose Physical Processor
Shared General Purpose CPU Physical Processor Pool
Shared zAAP Physical Processor
Shared zAAP Physical Processor
Shared zAAP Physical Processor Pool
Shared zIIP Physical Processor
Shared zIIP Physical Processor
Shared zIIP Physical Processor Pool
Shared IFL Physical Processor
Shared IFL Physical Processor
Shared IFL Physical Processor Pool

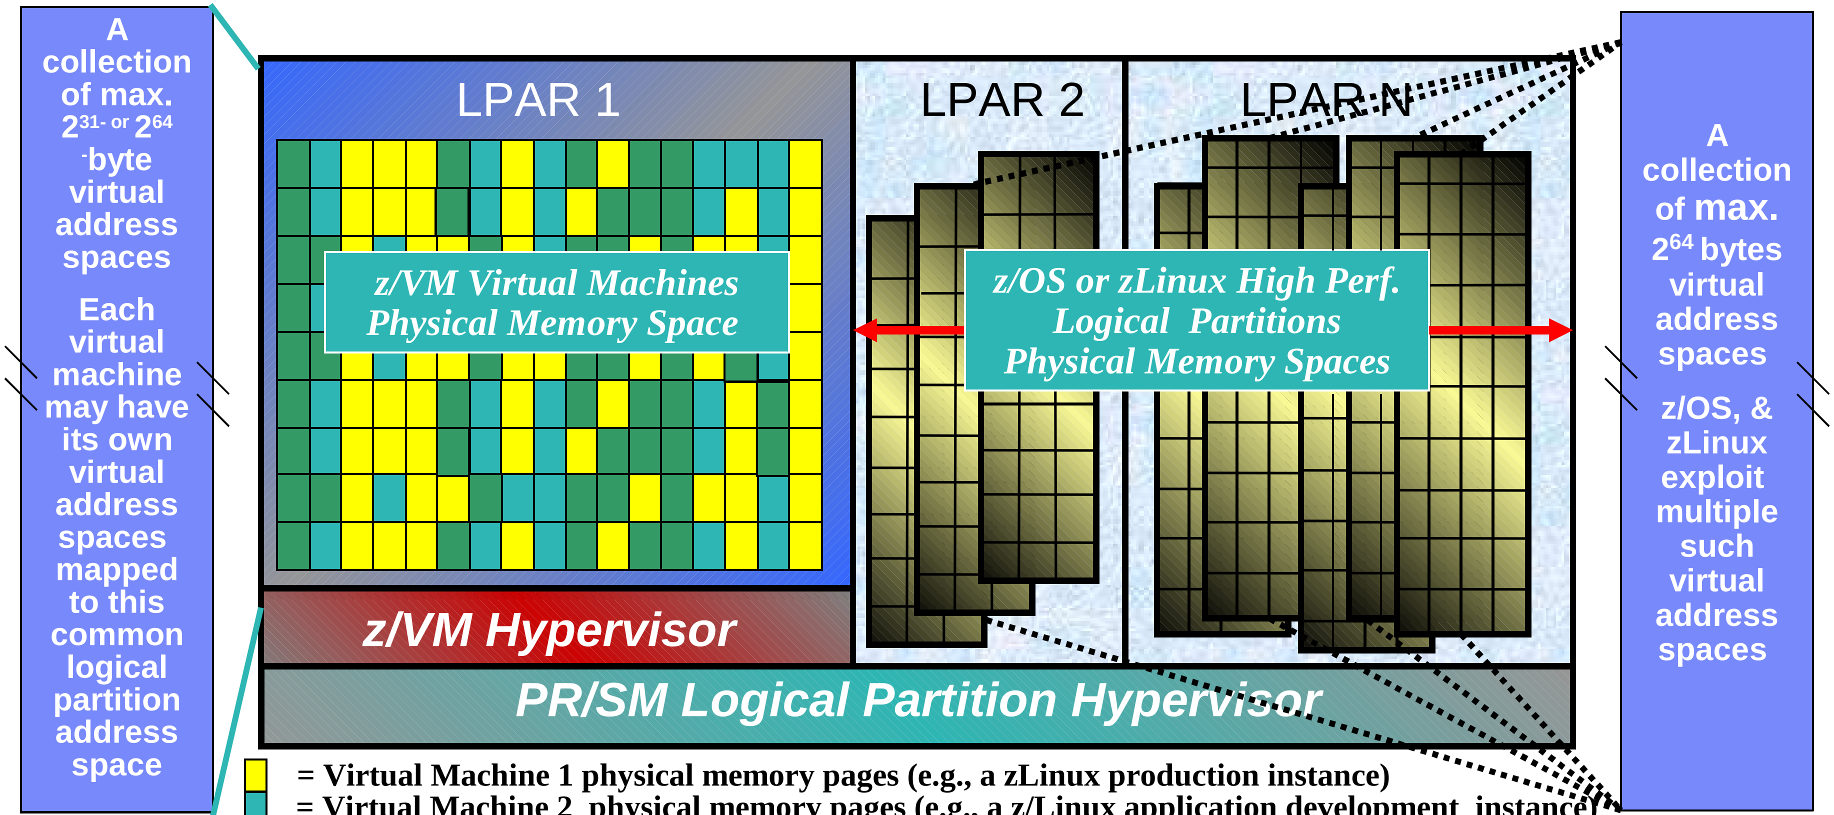

= Virtual Machine 1 physical memory pages (e.g., a zLinux production instance)
= Virtual Machine 2 physical memory pages (e.g., a z/Linux application development instance)
= Virtual Machine N physical memory pages (e.g., a z/OS test instance)

LPAR1: A virtual machine's transparently shared physical memory space = A logical partition physical memory address space used by z/VM for multiple virtual machine physical memory address spaces that are concurrently and transparently mapped to the partition's physical memory address pages (to allow the deployment of large numbers of virtual machines)

= the sets of z/OS or zLinux DAT physical memory pages associated with each configured virtual address space. That is, the sets of dynamically-allocated physical memory pages necessary to operate the z/OS or z/Linux image with high performance and minimum paging overhead most/much of the time (i.e., the OS image's working page sets)

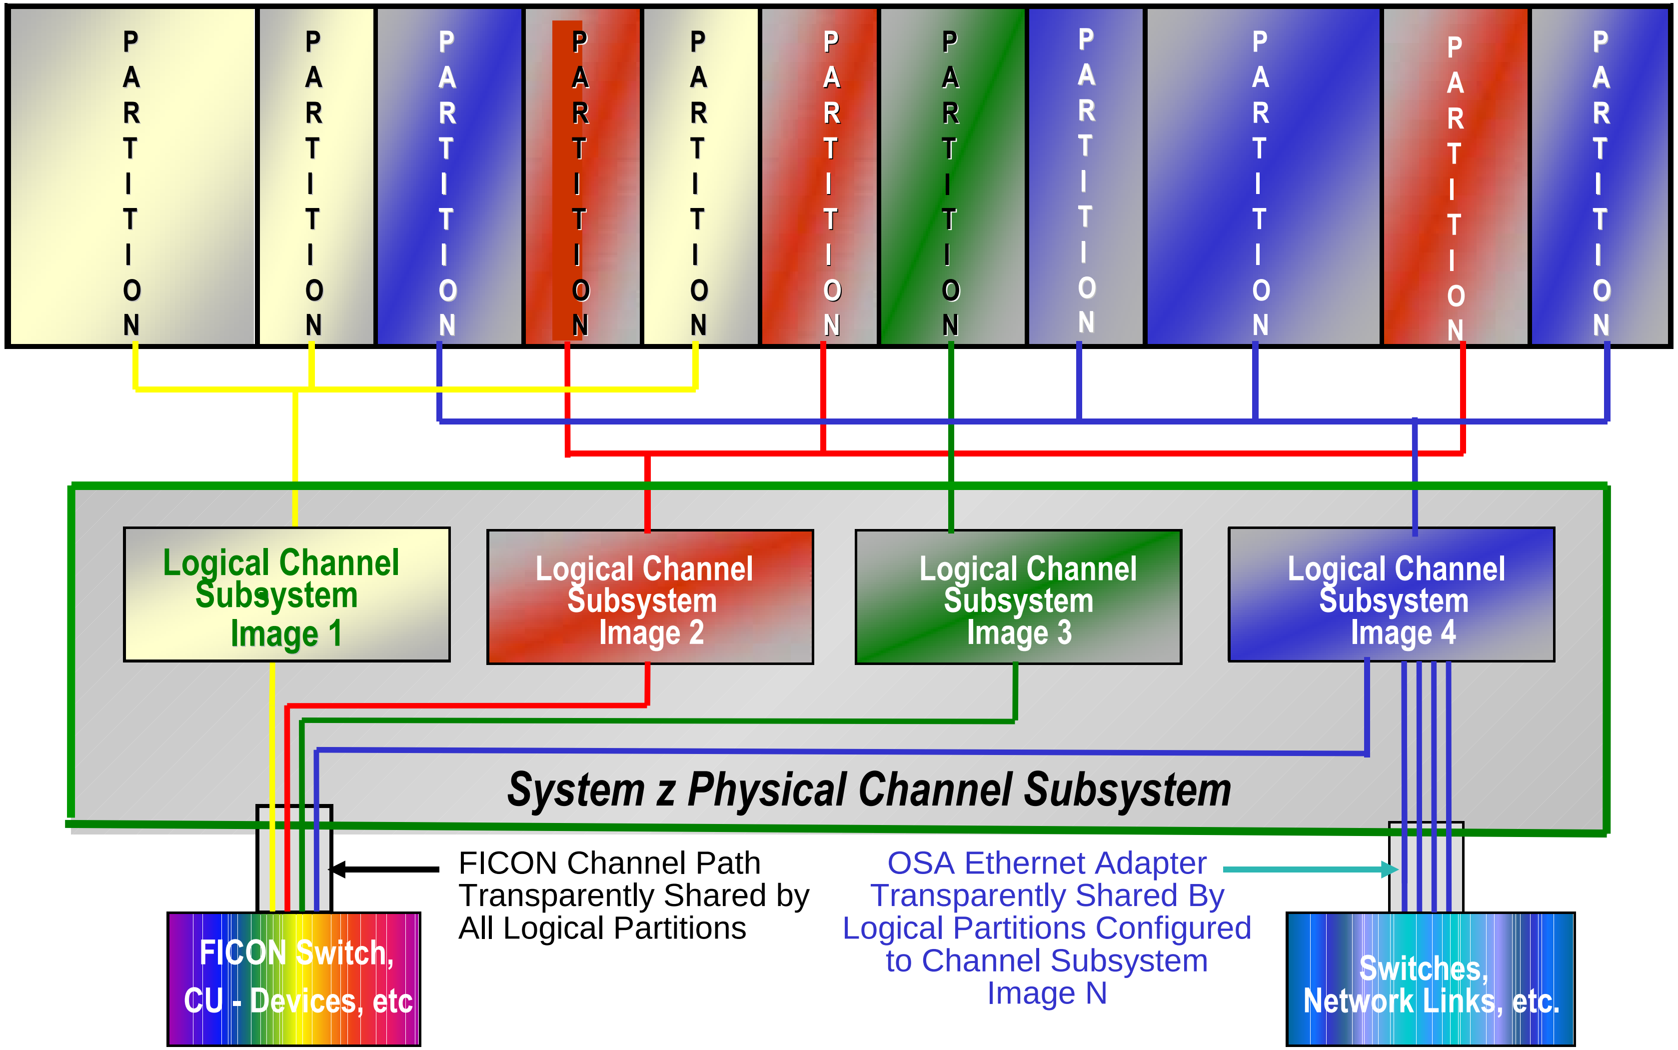
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
PARTITION
Logical Channel Subsystem Image 1
Logical Channel Subsystem Image 2
Logical Channel Subsystem Image 3
Logical Channel Subsystem Image 4
System z Physical Channel Subsystem
FICON Channel Path Transparently Shared by All Logical Partitions
OSA Ethernet Adapter Transparently Shared By Logical Partitions Configured to Channel Subsystem Image N
FICON Switch, CU - Devices, etc
Switches, Network Links, etc.

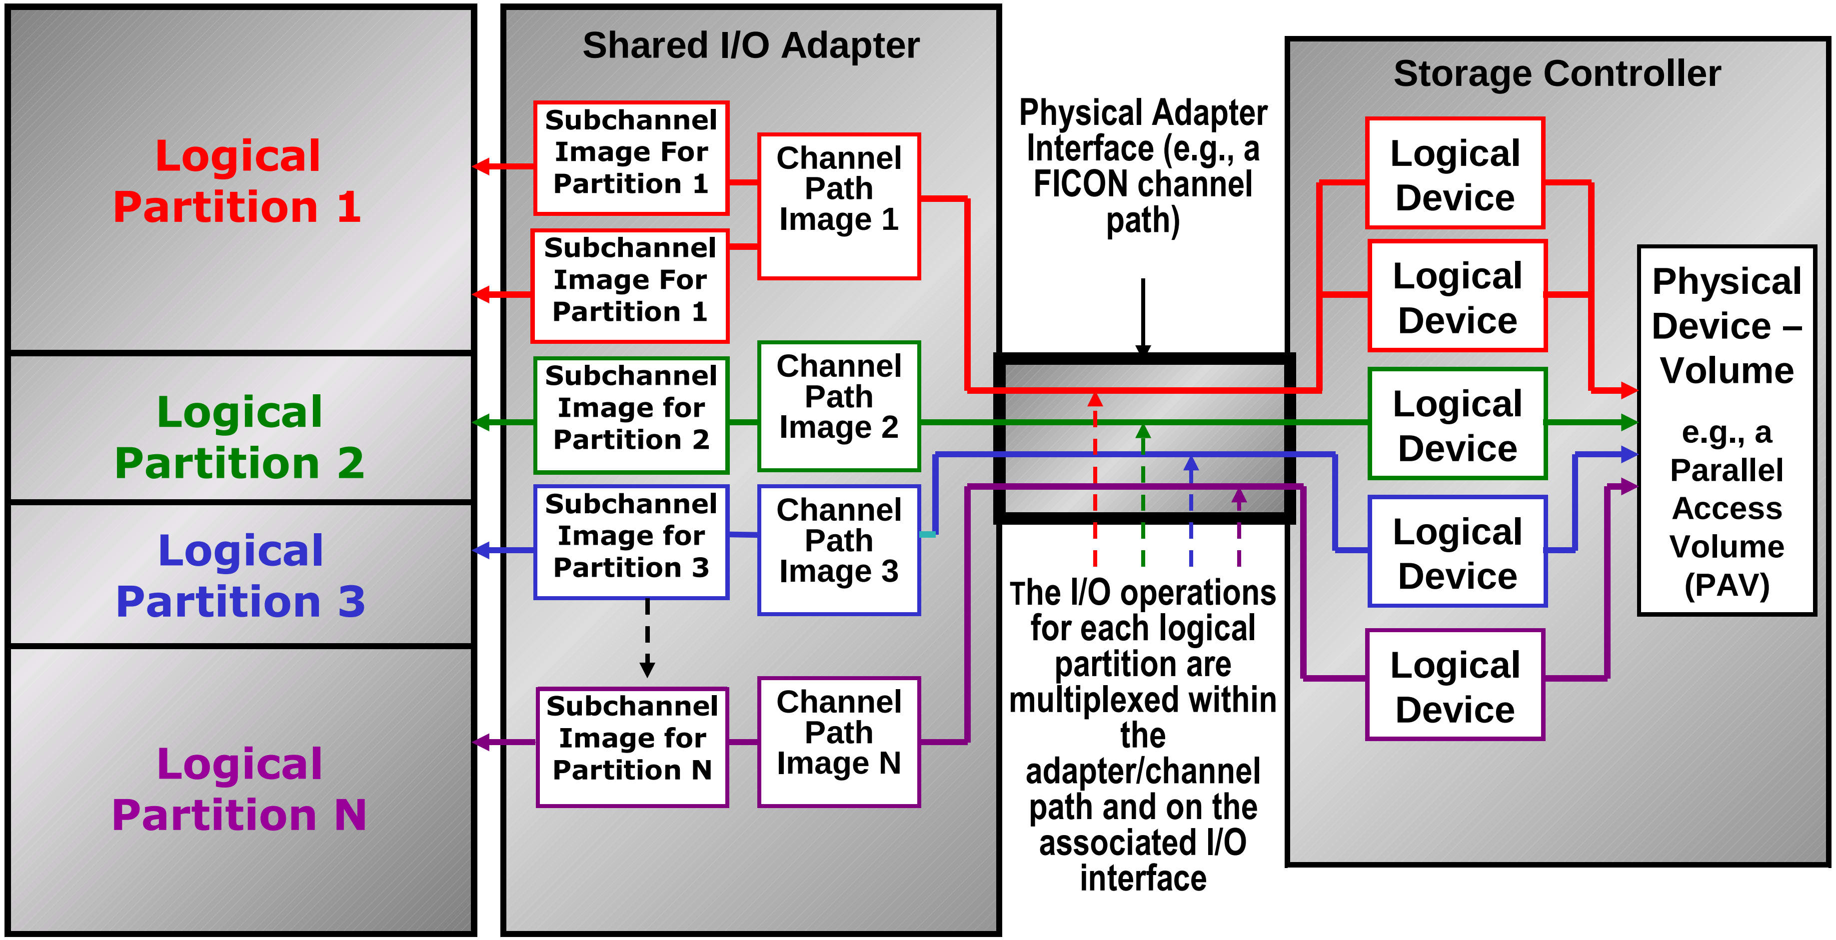

- The I/O infrastructure (adapters/channels, their transmission links, and attached I/O resources are shared by LPARs at native speeds (without hypervisor involvement)
  - I/O requests, their associated data transfers , and I/O interruptions flow between each OS instance and the shared I/O components just as if the I/O components were physically dedicated to a single OS instance

z/VM

- **Virtualizes almost everything**
  - CPUs up to architectural limit of 64
  - Memory up to machine-defined maximum (e.g, 1TB on z9)
  - I/O
    - Dedicated
    - Minidisks
    - Virtual disks in storage
    - Unit-record
    - LANs and switches
    - CTCAs
    - Coupling channels
    - Consoles
    - Cryptographic facilities

- **Provides added value**
  - Resource accounting
  - Performance monitoring
  - Autonomic CPU resource control
  - Minidisk cache
  - Shared memory
  - Collaborative memory management
  - GDPS/PPRC multi-platform resiliency
  - Integrated performance management
  - IBM Systems Director
  - QDIO Pass-through

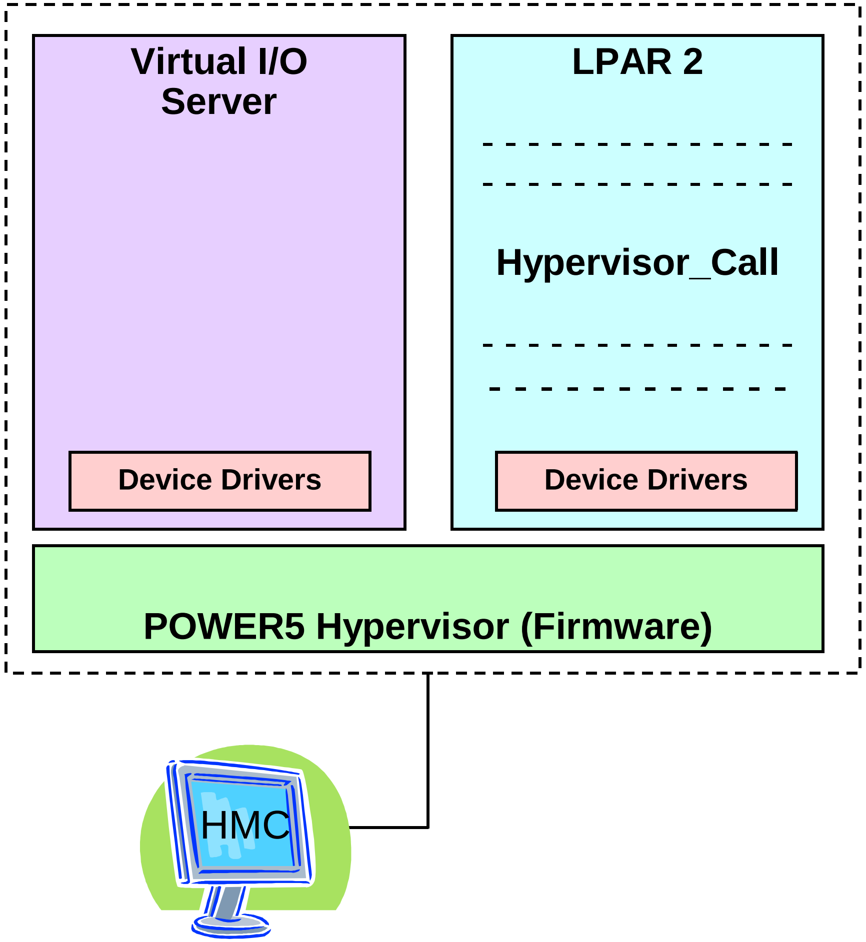

- Performance
  - Architecturally-defined Hypervisor calls
  - Hardware assists
  - Event-driven dispatching
  - Scalable from 0.1 to 64 CPUs
- Reliability, Availability, Serviceability
  - Hypervisor is core firmware
  - No I/O device or 3rd party code
  - Explicit functions
  - Multiple VIO servers 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

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

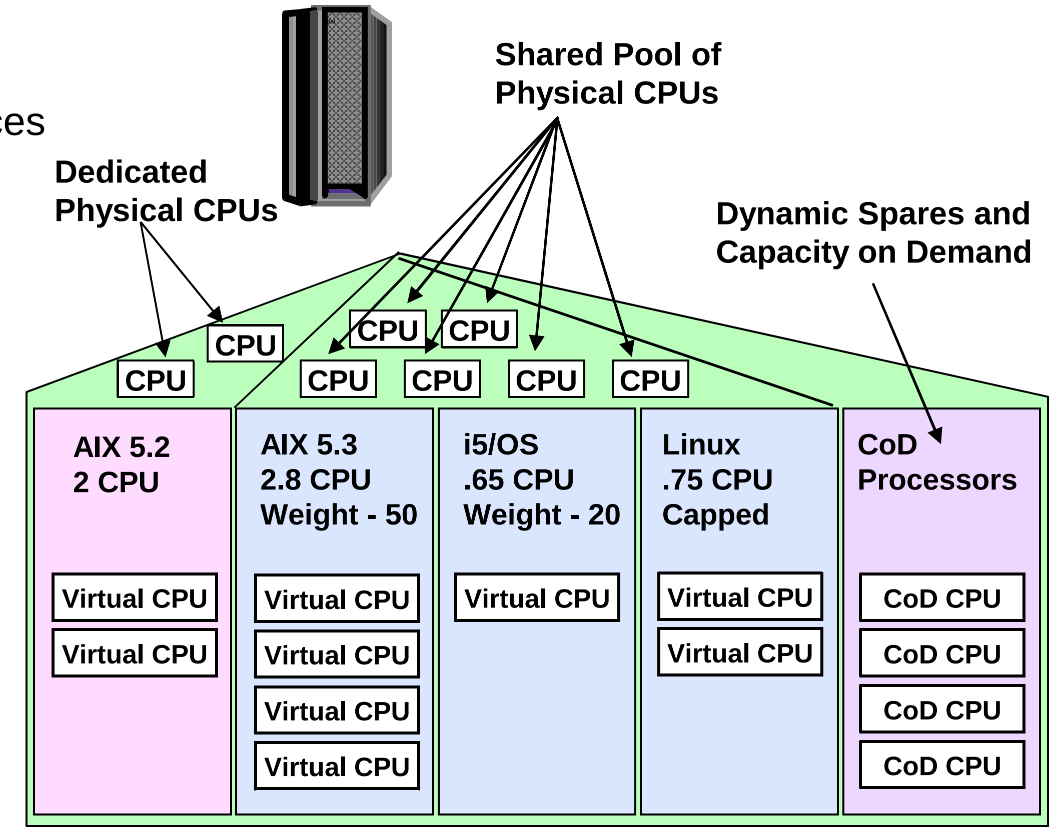

* Available on System p and OpenPower servers via the Advanced POWER virtualization features

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

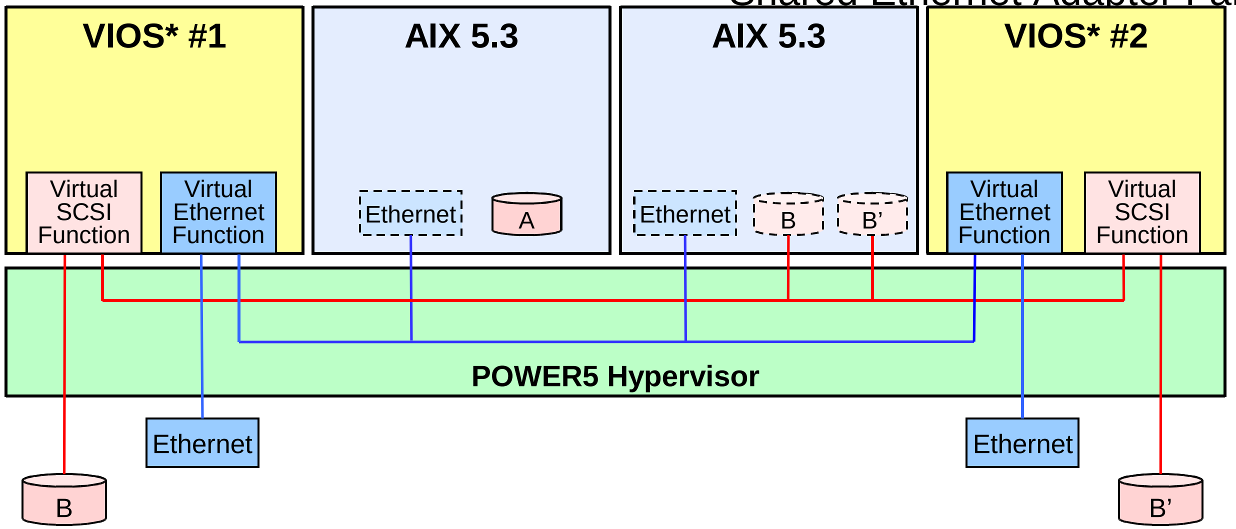

* Available on System p and OpenPower servers via the Advanced POWER virtualization features

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**

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

Systems Management

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

VMware ESX Evolution

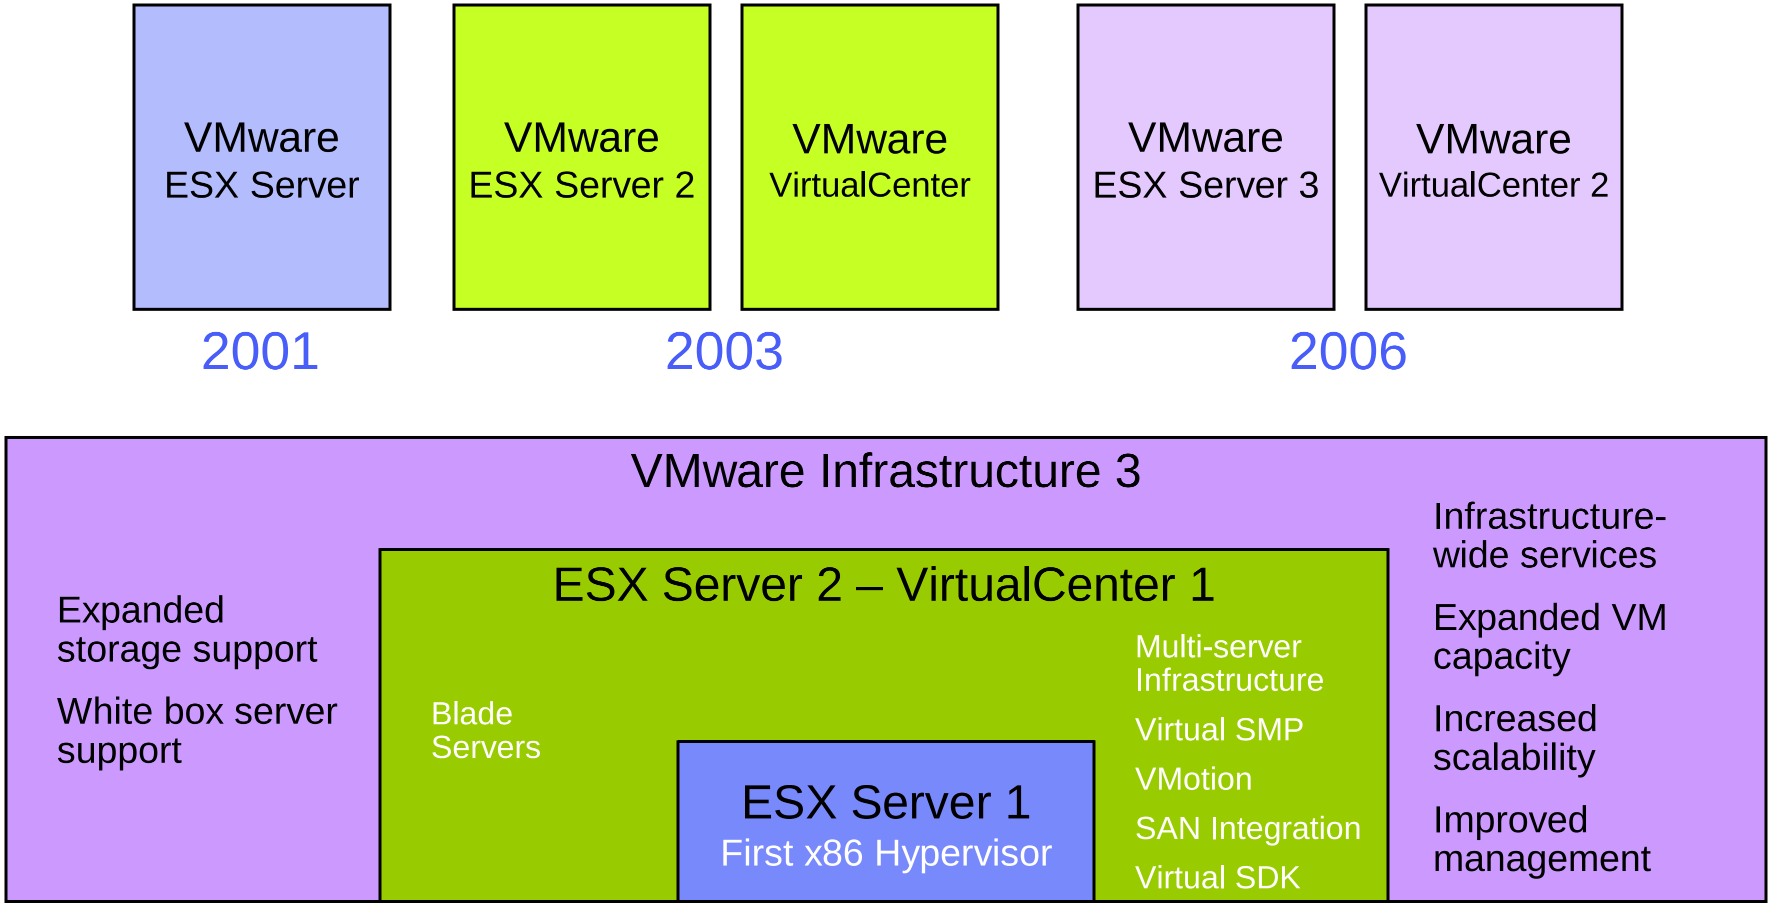

Part of VMware Infrastructure 3

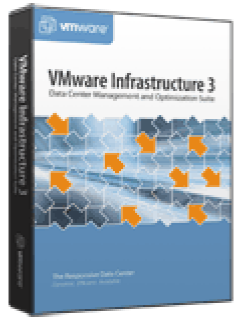

- **VMware ESX Server 3 runs directly on the hardware**
  - No hosting operating system required (as is the case with VMware GSX Server)
- **Creates multiple virtual machines on a single Intel system**
  - Supports a maximum of 128 virtual machines per VMware image (depending on system resources)
- **Manages 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

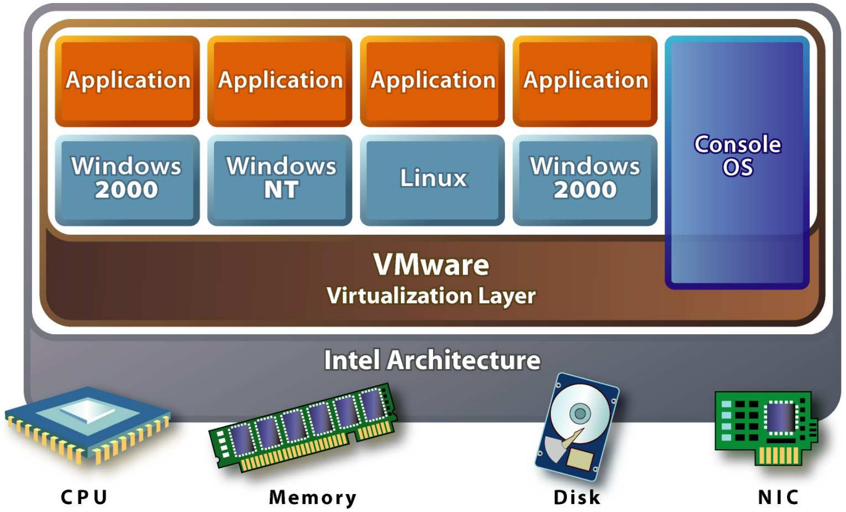

Part of VMware Infrastructure 3

- **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

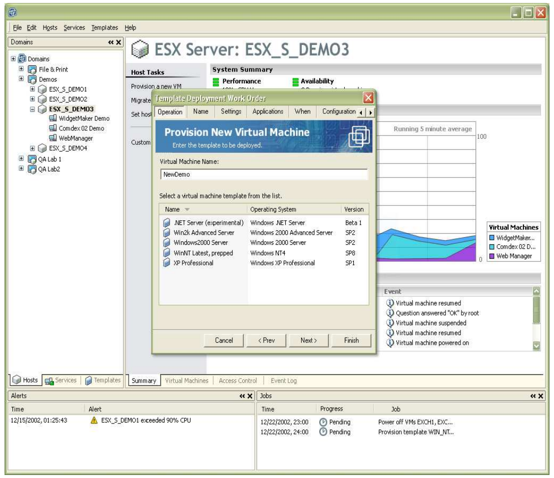

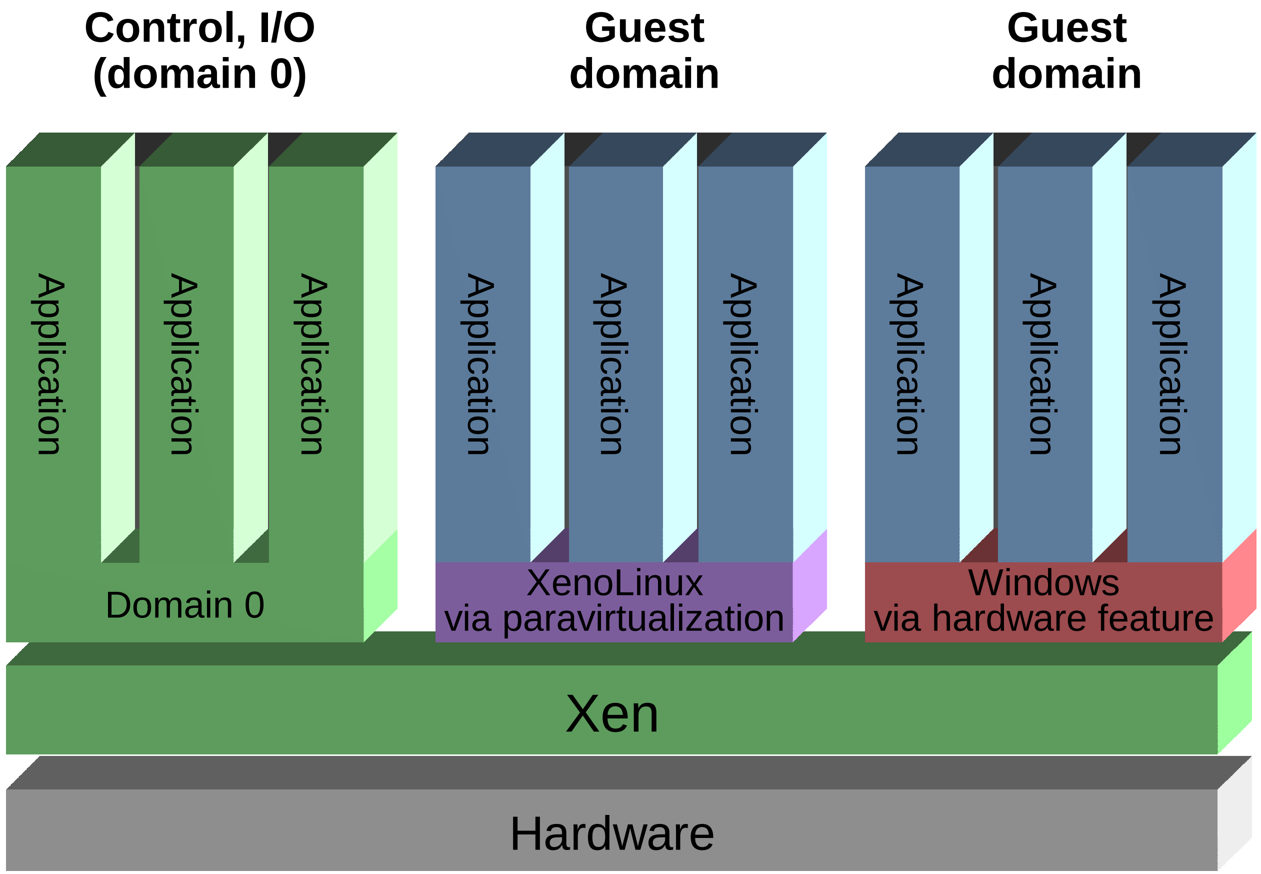

- **Like the POWER Hypervisor, Xen uses "paravirtualization" to run Linux, and uses a special type of partition (Domain 0) to virtualize I/O above the core hypervisor**
- **To run Windows, Xen will use new hardware virtualization support features from Intel (VT-x, VT-i) and AMD (SVM)**

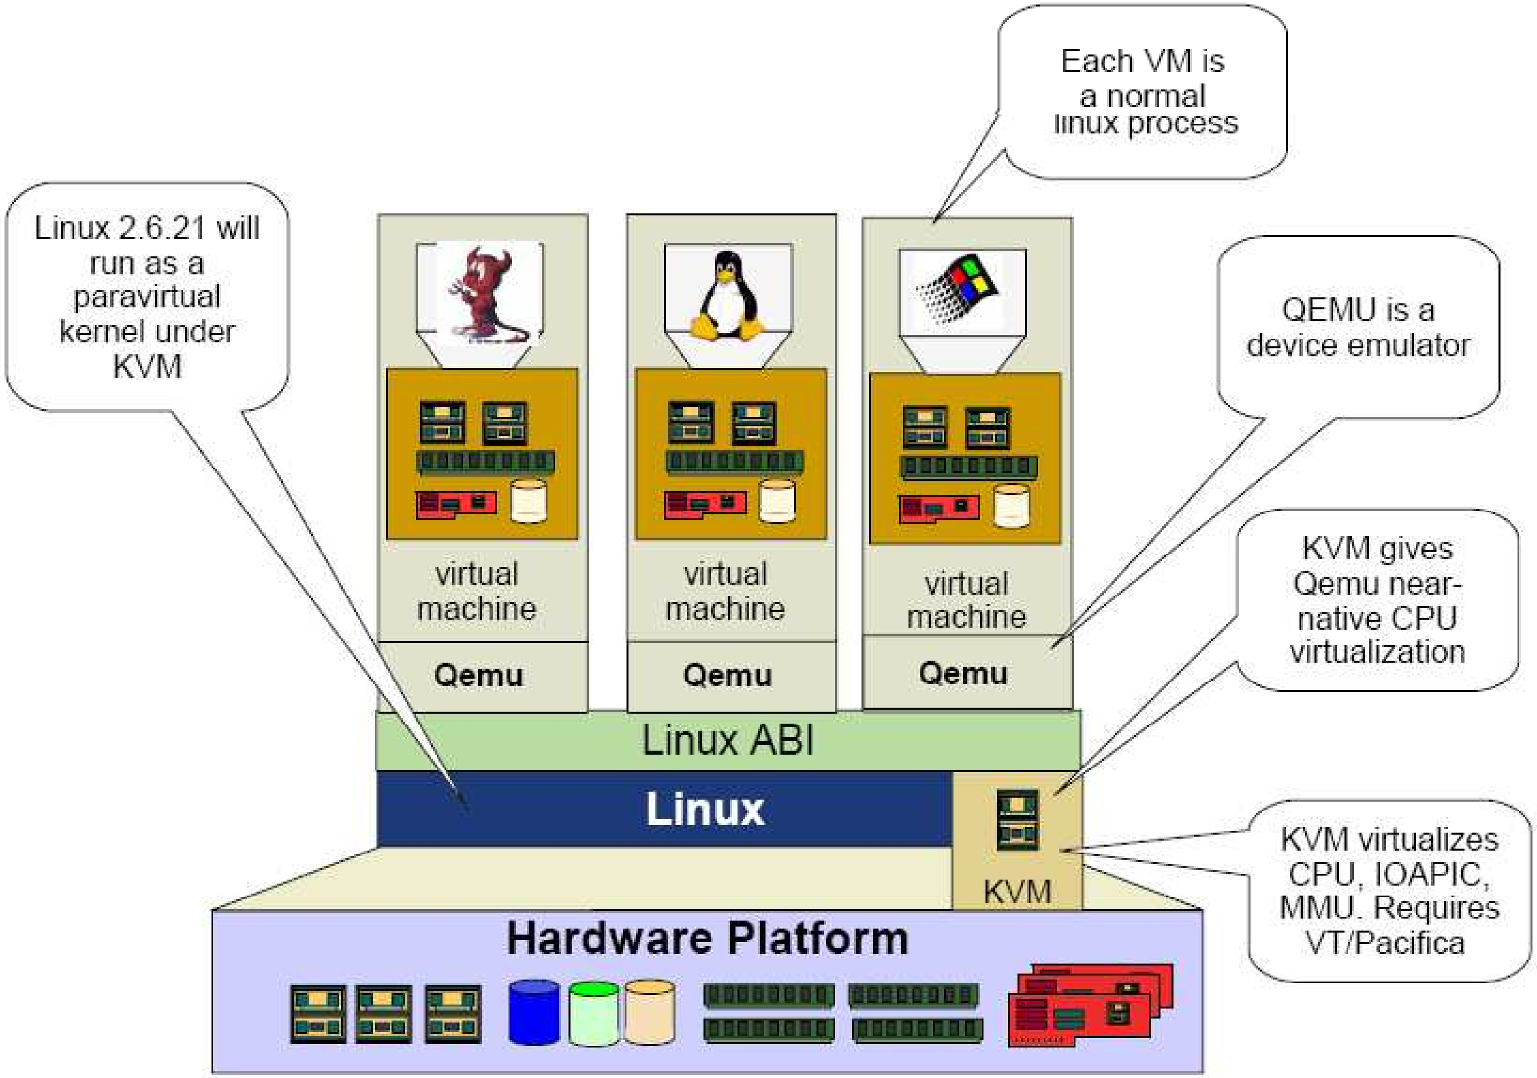
Each VM is a normal linux process
Linux 2.6.21 will run as a paravirtual kernel under KVM
QEMU is a device emulator
KVM gives Qemu near-native CPU virtualization
KVM virtualizes CPU, IOAPIC, MMU. Requires VT/Pacifica
virtual machine
virtual machine
virtual machine
Qemu
Qemu
Qemu
Linux ABI
Linux
KVM
Hardware Platform

Chip Design Affects Virtualization Capabilities

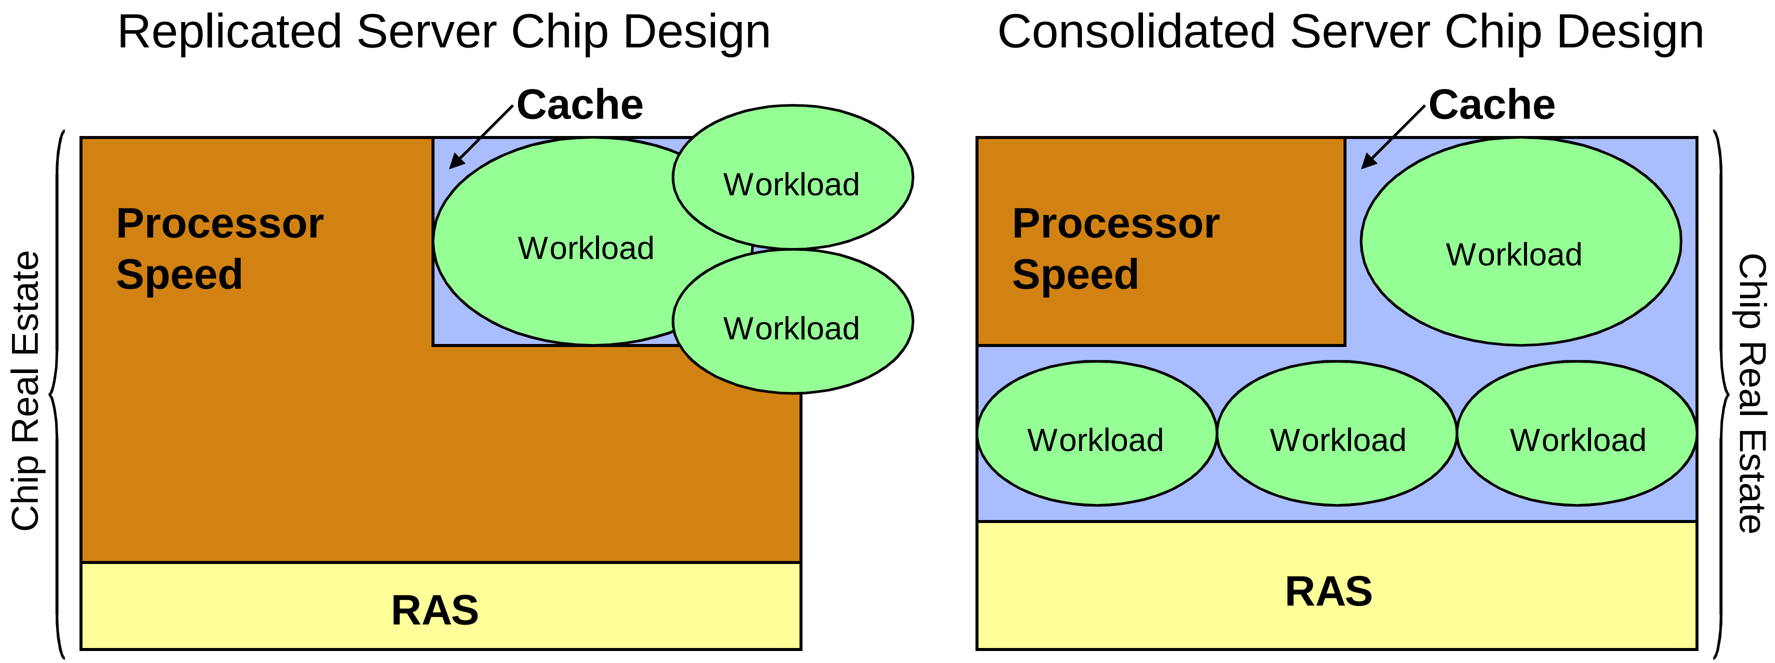

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- "Fast" processor speed is not fully realized due to cache misses

- 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

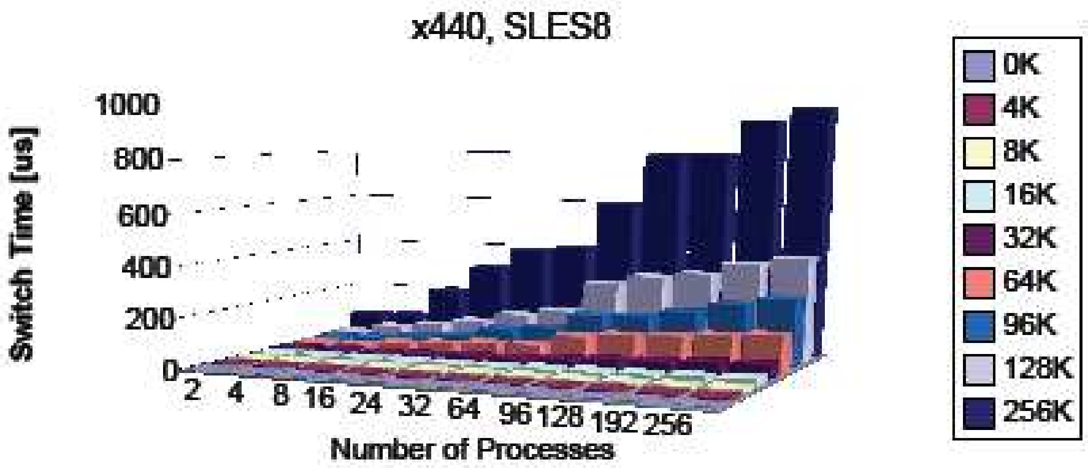

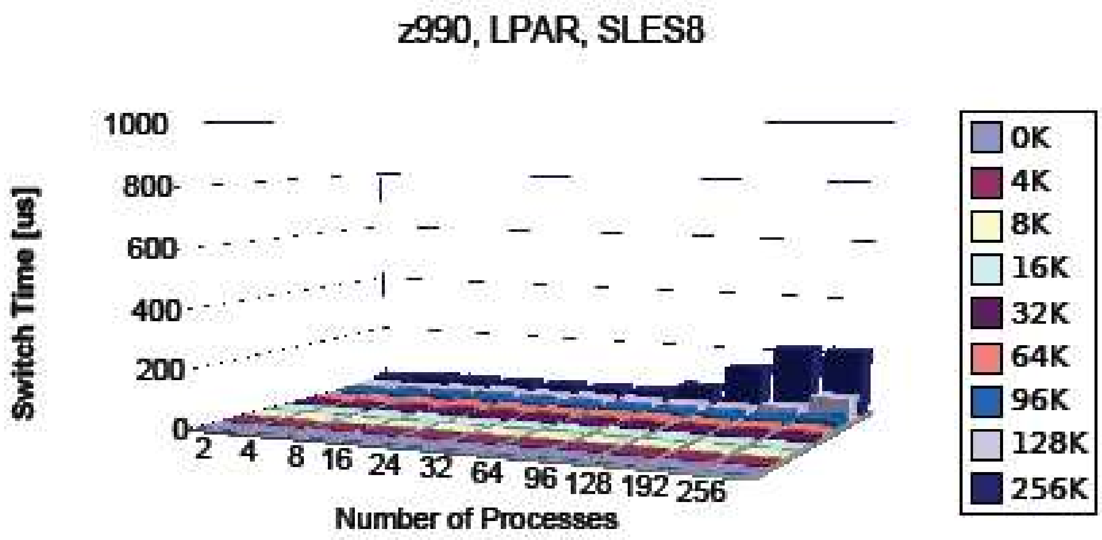

- Time required to perform context switching is an indication of memory time
- Virtualization, by definition, involves context switching

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

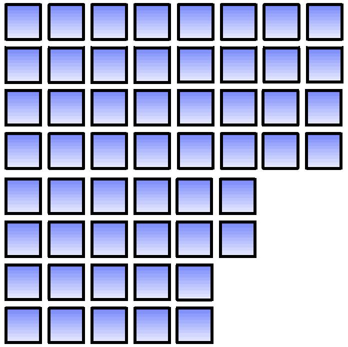

Share up to 54 processors among up to 60 LPARs

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

Up to 10 System Assist Processors

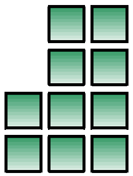

Offload system processing to dedicated CPUs with no effect on software license fees

Plus up to 336 I/O Processors

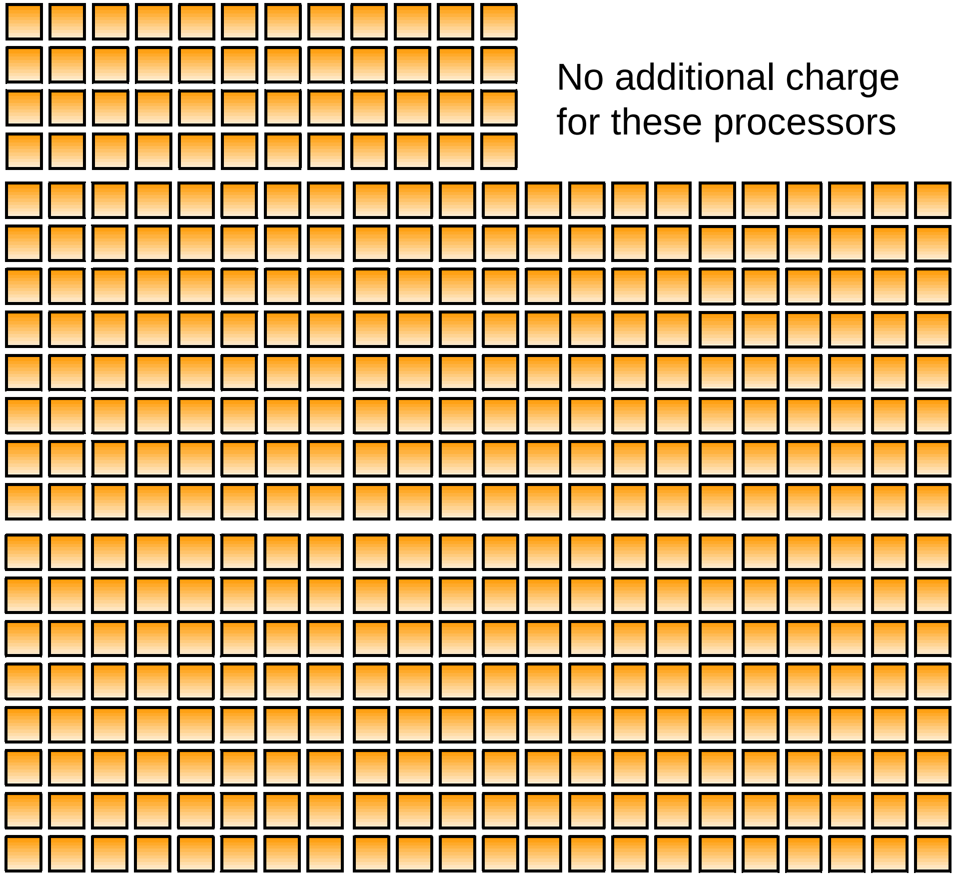

No additional charge for these processors

Datacenter in a Box: Cost Benefits 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

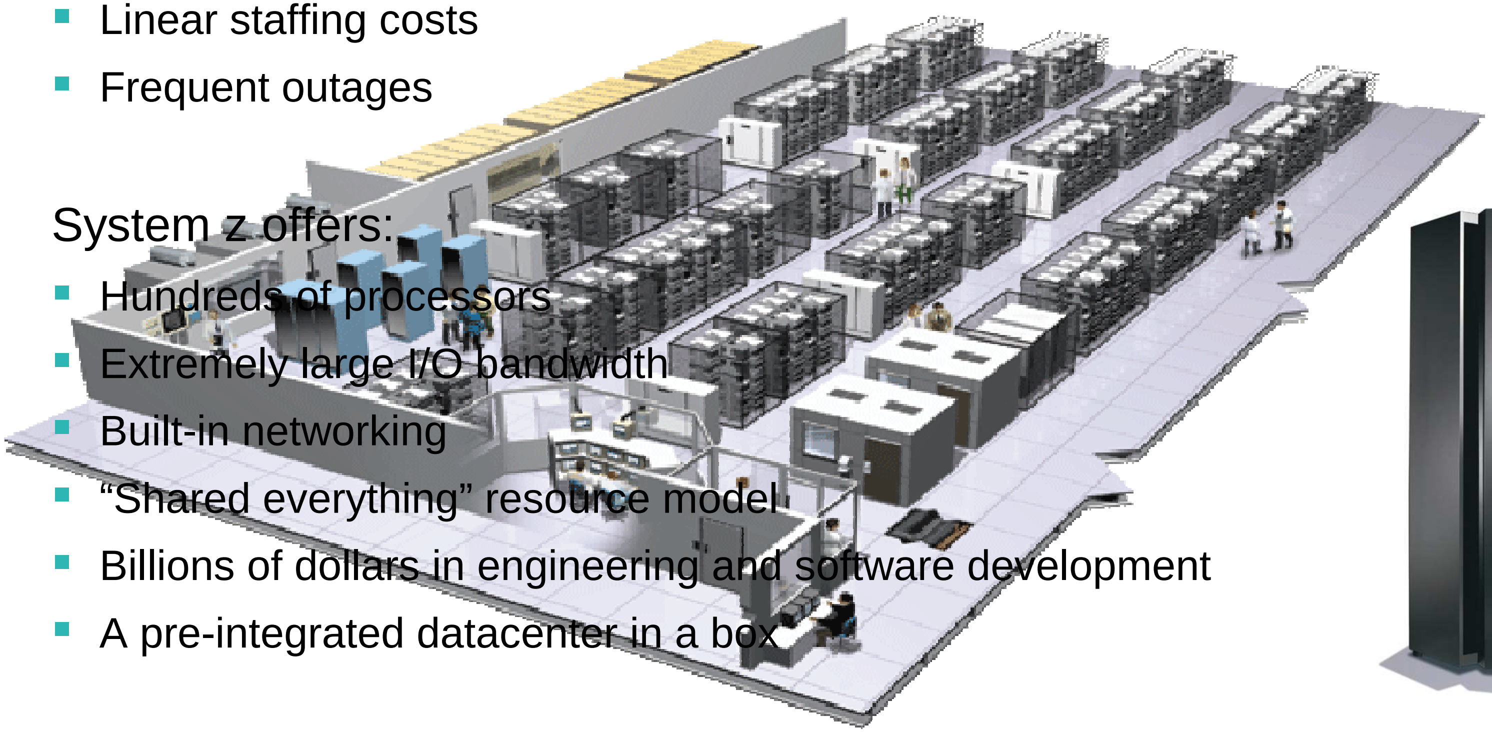

Complexity Issues

- **Getting started with x86 virtualization**
  - Easy to do proof-of-concept
    - Most servers heavily underutilized
    - Potential for significant savings
- **After x86 virtualization is adopted**
  - Adds a layer of complexity
  - Does not eliminate discrete servers or the need to manage them
  - Virtualization technology immature
    - Chargeback
    - Capacity planning
    - No LPAR alternative for performance-critical workloads
    - Small hope of consolidating I/O intensive workloads
  - Too many options
    - Xen, VMware, kvm, Microsoft VS, Virtuozzo, …
    - What is the strategic choice?
  - Windows licensing issues for VMotion and similar technologies
  - VMware adoption wide but shallow
  - Virtualizaion on x86 may help reduce rate of server sprawl
    - But sprawl will continue

Scale Out and Scale Up Support

- **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 V5.3 supports up to 64 CPUs per virtual machine and 32 CPUs per z/VM LPAR
  - VMware ESX Server 3 limits virtual machines to 16 GB of memory
    - ☑ z/VM V5.3 supports virtual machines with up to 1TB of memory on System z9
  - 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)

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 reduces 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
  - I/O virtualization support may still be an issue for IVT and AMD-V systems
- **In general, VMware users are advised to be careful overcommitting real memory and CPU resources**

☑ **z/VM and System z are *designed* from the ground up 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.

Scalability, Granularity, Efficiency

- **z/VM offers single-system scalability advantages**
  - 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

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**
  - 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

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**

Leadership Technology

- **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**

Key Value Propositions for Linux on System z

- **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)

System z Virtualization Technology Evolution

- **Over 40 years of continuous innovation in virtualization**
  - Refined to support modern business requirements
  - Exploit hardware technology for economical growth

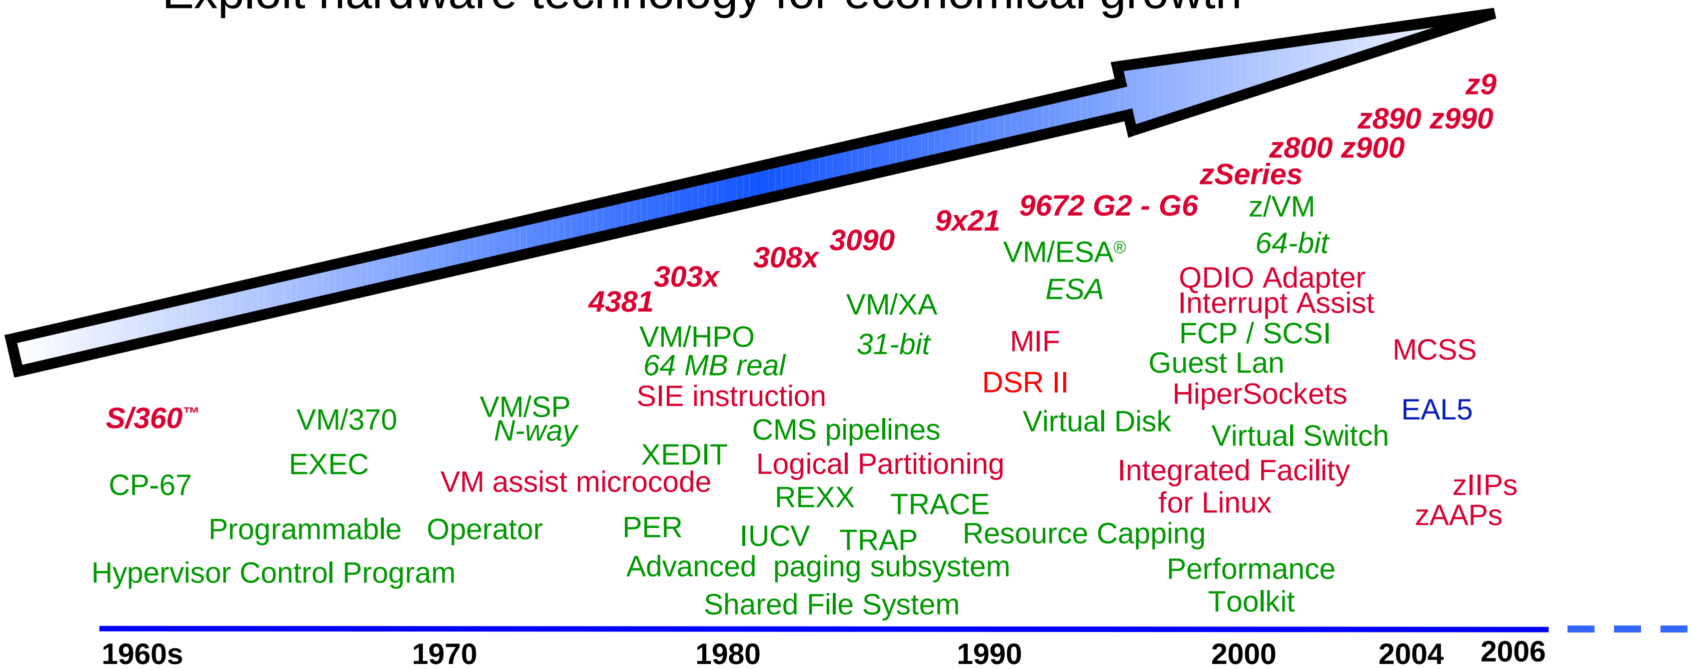

System z – A comprehensive and sophisticated suite of virtualization functions