Session IS07

Virtualization with System z

Dr. Manfred Gnirss
Technical Marketing Competence Center Europe,
R&D Support Support Centers Boeblingen
IBM Deutschland Research & Development GmbH
gnirss@de.ibm.com

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 Storage, 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.

Topics

- **Business drivers for server consolidation**
- **Server virtualization and hypervisor basics**
- **x86 virtualization**
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Some customer examples*
  - *Summary*

All company sizes identify cost savings as a top benefit of server virtualization followed by availability and flexibility

- **Companies and their IT organizations are almost always under pressure to cut costs**
- **Servers running one or a few applications are typically underutilized**
- **Consolidating servers through virtualization is one easy way to reduce costs**

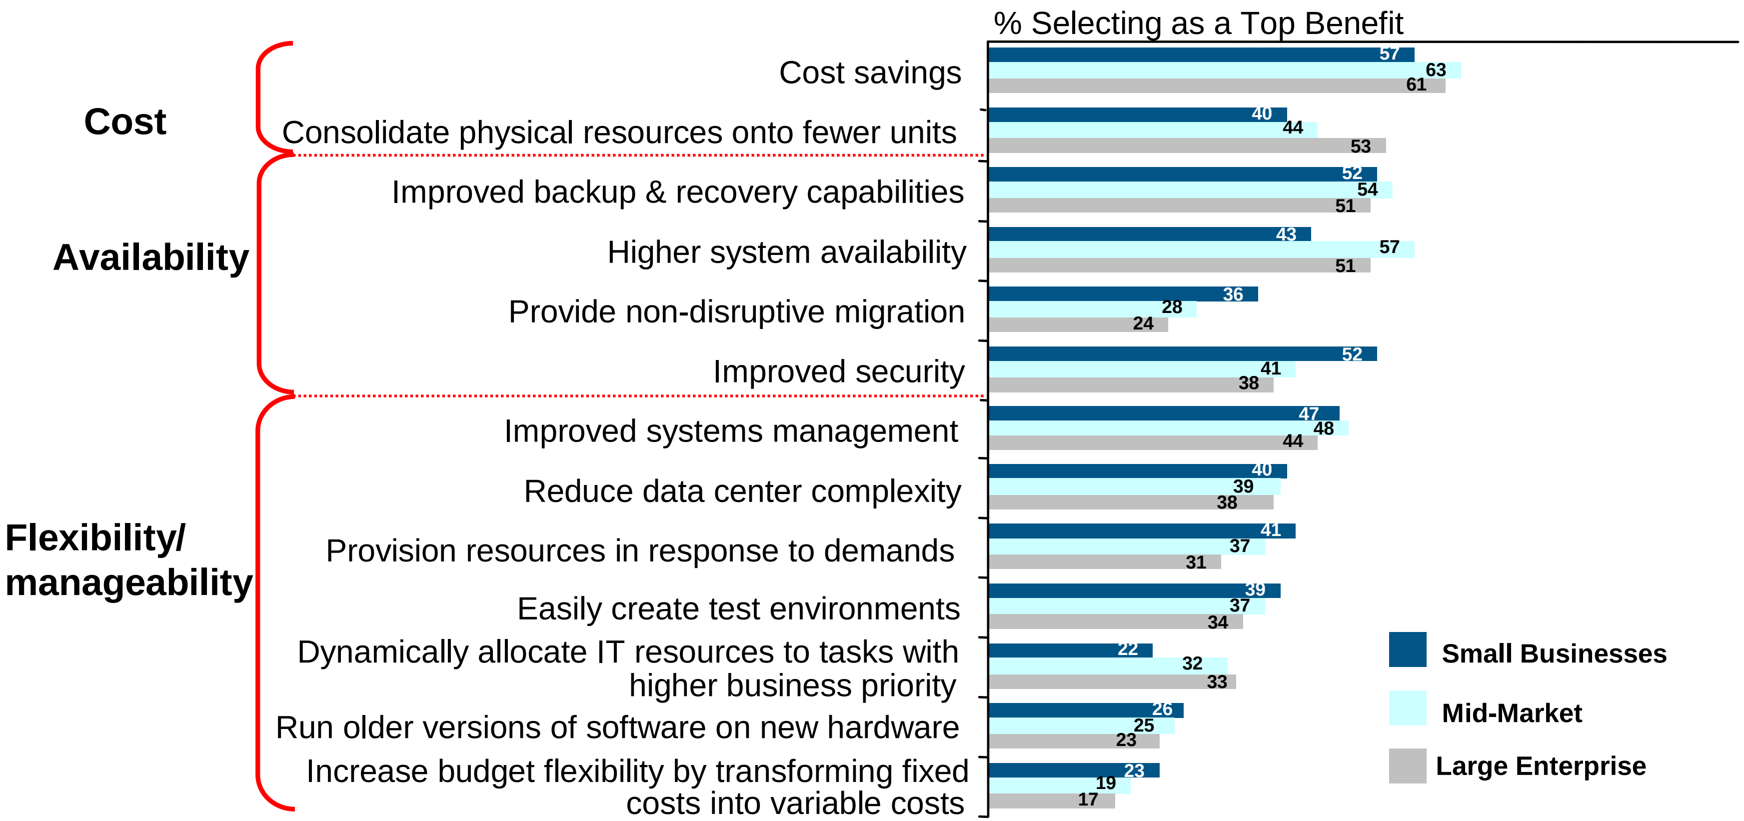

Source: IBM Study 6/08

Topics

- **Business drivers for server consolidation**
- **Server virtualization and hypervisor basics**
- **x86 virtualization**
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Some customer examples*
  - *Summary*

Server Partitioning and Virtualization
Affecting All Areas of Virtual Server Design and Deployment

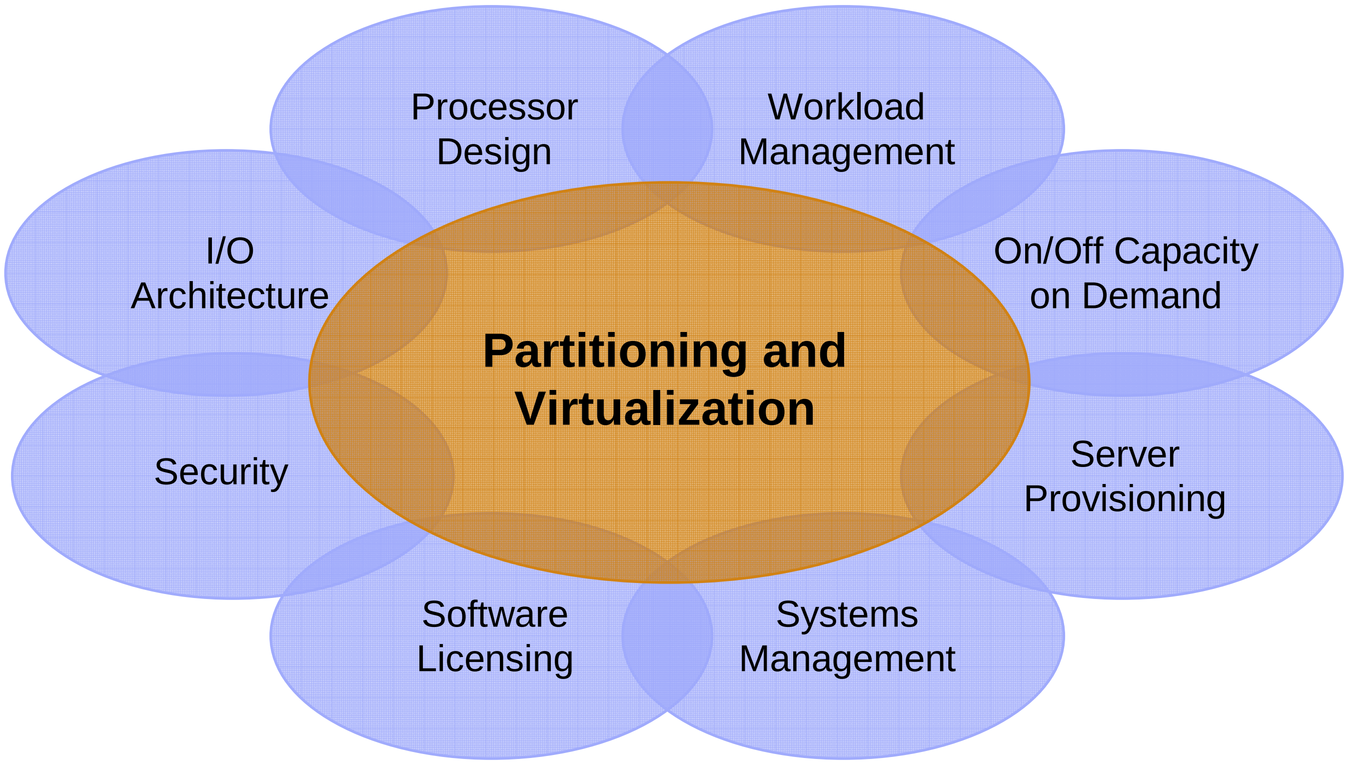

Hypervisor Implementation Methods

Trap and Emulate

Virt Mach

- **Load**
- **Add**
- **Store**
- **PrivOp** → **Trap** → **Hypervisor PrivOp emulation code**
- **Load**
- **...**

- **VM runs in user mode**
- **All privileged instructions cause traps**

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

Translate, Trap, and Emulate

Virt Mach

- **Load**
- **Add**
- **Store**
- **TrapOp** → **Trap** → **Hypervisor PrivOp emulation code**
- **Load**
- **...**

- **VM runs in user mode**
- **Some IA-32 instructions must be replaced with trap ops**

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

Hypervisor Calls (“Paravirtualization”)

Virt Mach

- **Load**
- **Add**
- **Store**
- **Hcall** → **Call** → **Hypervisor service**
- **Load**
- **...**

- **VM runs in normal modes**
- **OS in VM calls hypervisor to access real resources**

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

Direct Hardware Virtualization

Virt Mach

- **Load**
- **Add**
- **Store**
- **PrivOp** → **Exit** → **Hypervisor service**
- **Load**
- **...**

- **VM runs in normal modes**
- **Hardware does most of the virtualization (SIE architecture)**
- **Hypervisor provides control**

Hypervisor calls also supported

Examples: System z LPAR, z/VM
Benefits: High efficiency, runs unmodified OS
Issues: Requires underlying hardware support

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**

Server Virtualization Approaches

Hardware Partitioning

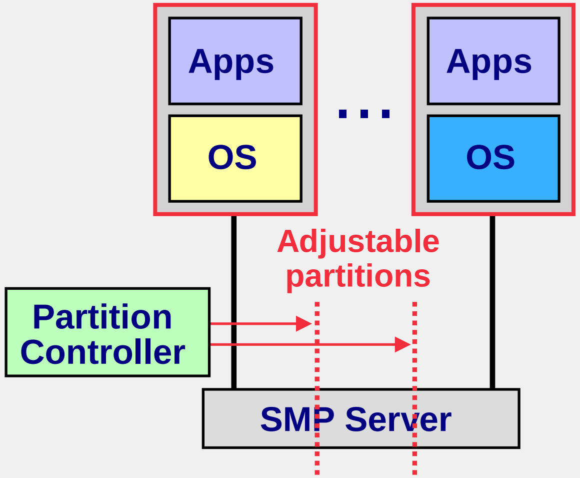

Server is subdivided into fractions each of which can run an OS

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

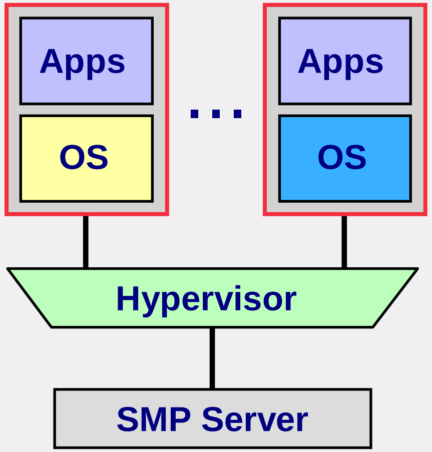

Hypervisor provides fine-grained timesharing of all resources

Hypervisor software/firmware runs directly on server
- System z LPAR and z/VM
- PowerVM Hypervisor
- VMware ESX Server
- Xen Hypervisor
- Microsoft Hyper-V

Hosted Hypervisor

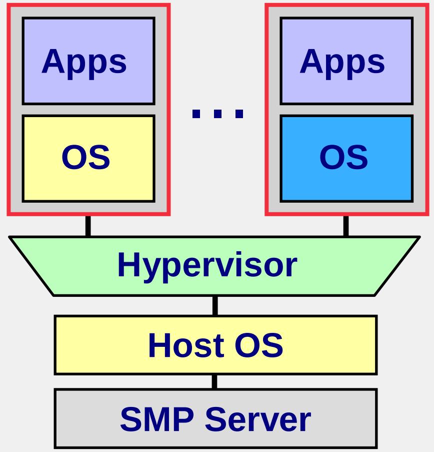

Hypervisor uses OS services to do timesharing of all resources

Hypervisor software runs on a host operating system
- VMware Server
- Microsoft Virtual Server
- HP Integrity VM

Hybrid: KVM

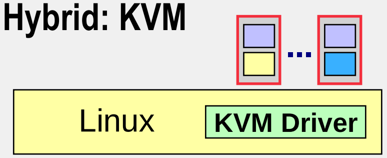

Characteristics:
- Bare-metal hypervisors offer high efficiency and availability
- Hosted hypervisors are useful for clients where host OS integration is important
- Hardware partitioning is less flexible than hypervisor-based solutions

Chip Design Affects Virtualization Capabilities

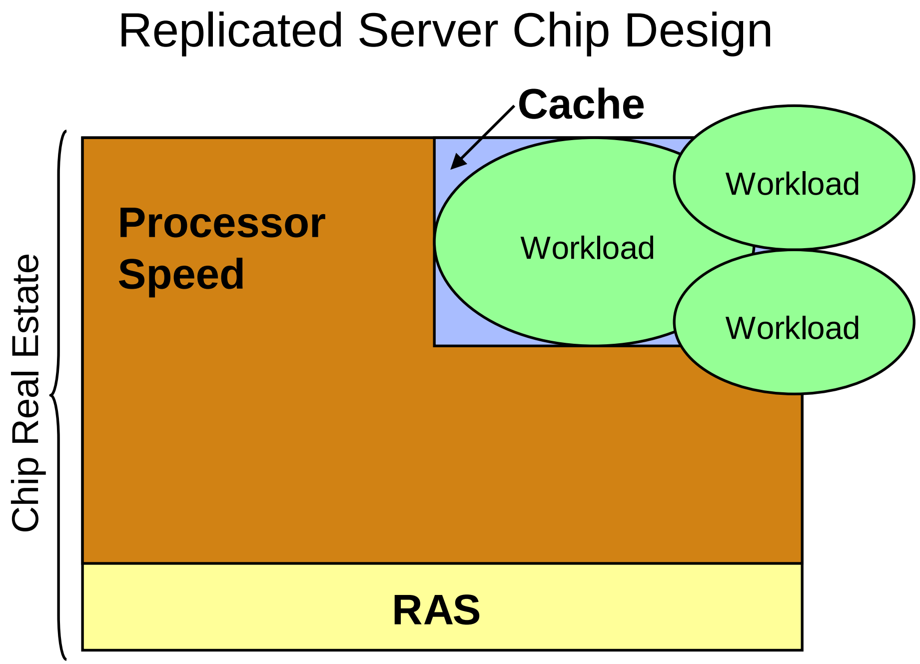

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- “Fast” processor speed is not fully realized due to cache misses

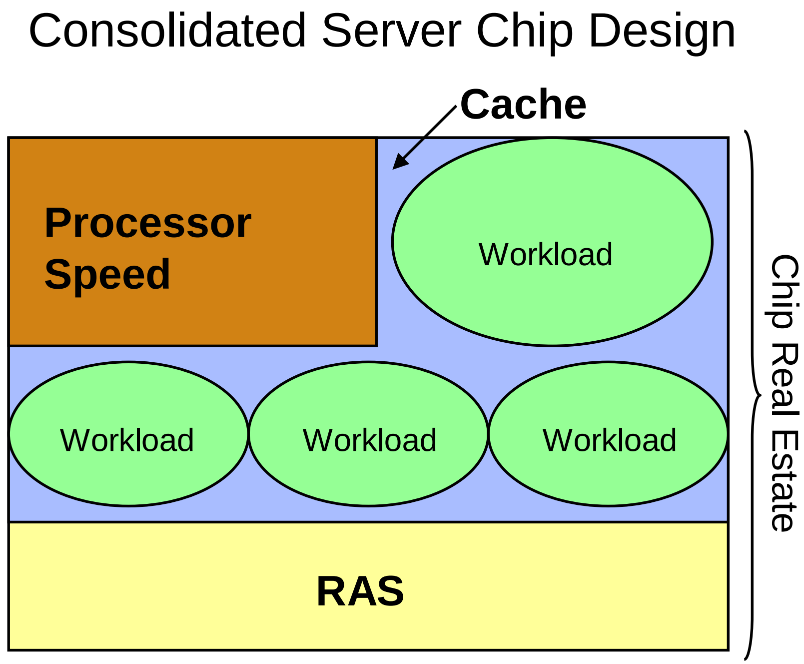

- System z cache is able to contain more working sets
- Processor speed is optimized by increased cache usage
- Additional RAS function is beneficial for mixed workloads

Server Architecture Genetics

Consider the Heritage of Today's Server Platforms

- **x86 systems**
  - Key value proposition: *end-user autonomy*
  - "Ctl-Alt-Del" not a problem for a single-user system
- **UNIX systems**
  - Key value proposition: *processor speed*
  - Engineering/scientific computing with portability of applications and skills
- **Mainframe systems**
  - Key value proposition: *mixed workloads*
  - Highest degrees of efficiency, availability, workload mgmt, security

Virtualization Essentials

Virtualization technology can be significantly constrained or compromised (or enhanced) by the underlying system architecture.

Topics

- **Business drivers for server consolidation**
- **Server virtualization and hypervisor basics**
- **x86 virtualization**
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Some customer examples*
  - *Summary*

x86 Server Virtualization Landscape
Hypervisors

So many x86 virtualization offerings makes it challenging for users to select the right technology for an x86 environment.

- **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
- **KVM (open source Linux)**
  - Integrated into the Linux 2.6.20 kernel
  - New strategic initiative for Red Hat ("Embedded Linux Hypervisor")
- **Microsoft® Windows® (proprietary)**
  - Virtual Server 2005 – Longhorn (2007) – Viridian (2007-2008)
- **Many more (Solaris Containers, Virtuozzo, HP Integrity VM, ...)**
  - So many x86 hypervisors presents a challenge for users
  - Solaris Containers, Virtuozzo, HP Integrity VM, ...

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

x86 Server Virtualization Landscape
Systems Management

- **VMware VirtualCenter**
- **Virtual Iron Virtualization Manager**
- **XenSource XenEnterprise**
- **Microsoft System Center Virtual Machine Manager**
- **IBM Systems Director**
  - Including *IBM Director* and *IBM Tivoli Provisioning Manager*

x86 Virtualization Pain Points

More Reasons for Using System z Virtualization Technology

- **Disaster recovery**
- **Speed of deployment**
- **Virtual machine sprawl and lifecycle management**
- **Chargeback and licensing**
- **Security**
- **I/O intensive workloads**

There is a downside to x86 virtualization that represents an opportunity for hosting workloads on System z virtualization technology instead.

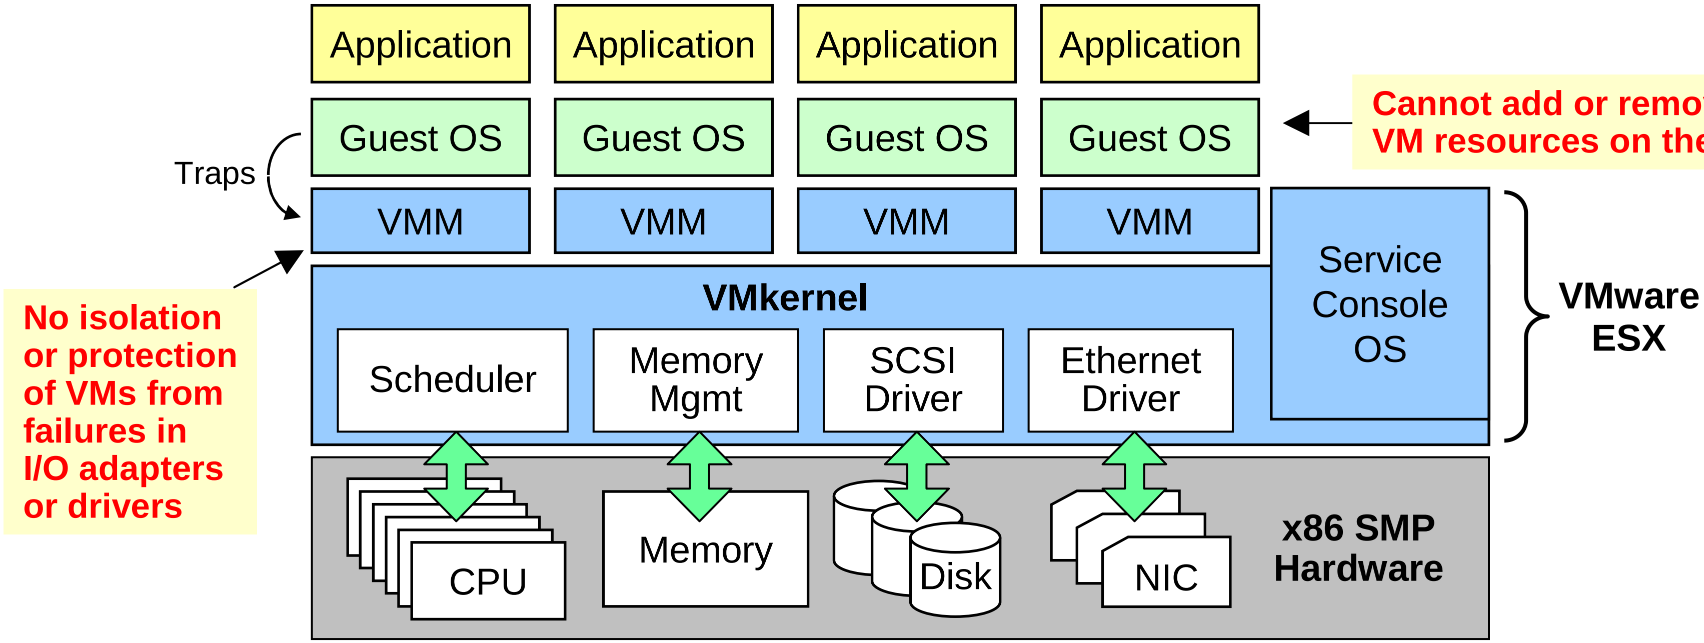

- **Supports a wide range of unmodified Windows and Linux versions**
- **Guest OS runs in User Mode; privileged instructions trap to Virtual Machine Monitor (VMM)**
  - "Trapping and mapping" is a significant source of performance overhead
- **Guest OS binary code is translated incrementally at load 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**

VMware Virtual Infrastructure

- Local, Fibre, NAS and iSCSI storage
- 8-way Virtual SMP
- 255 GB guest memory
- Unified GUI
- VMotion
- DRS
- HA

Architecture

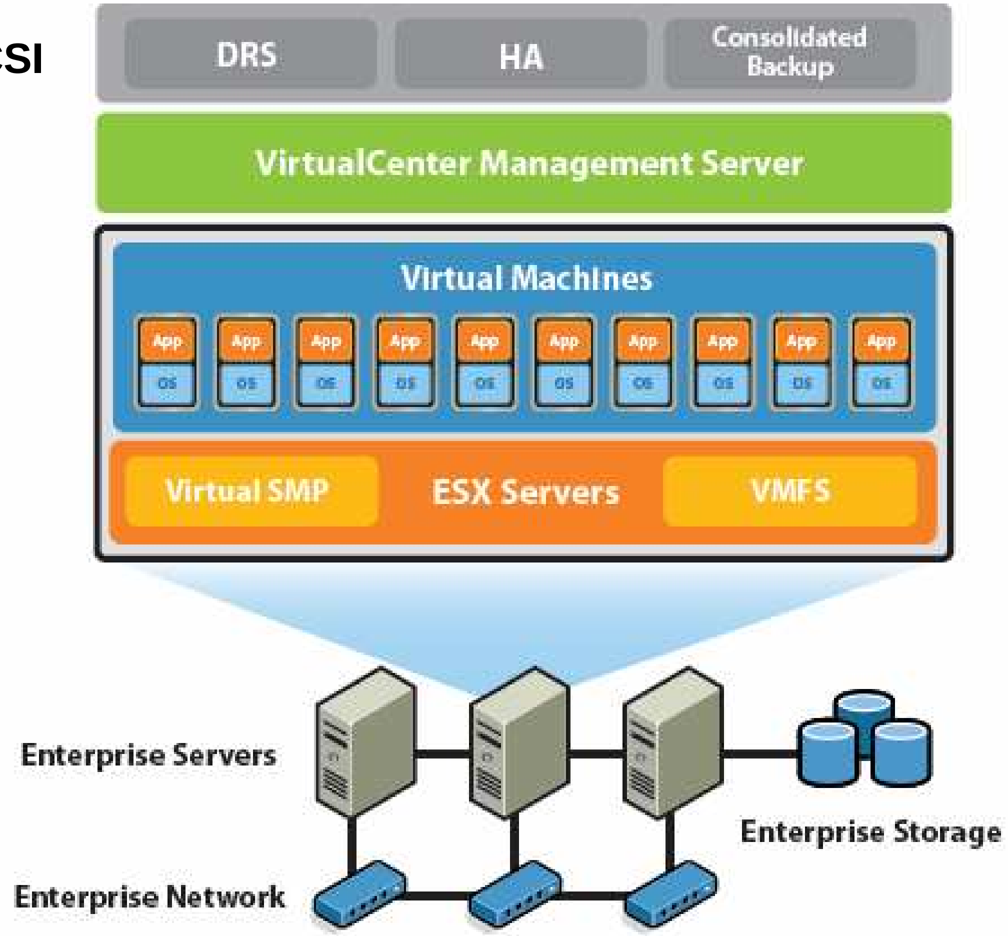

VMotion™ Technology

Instantly shift running systems across hosts often with imperceptible downtime

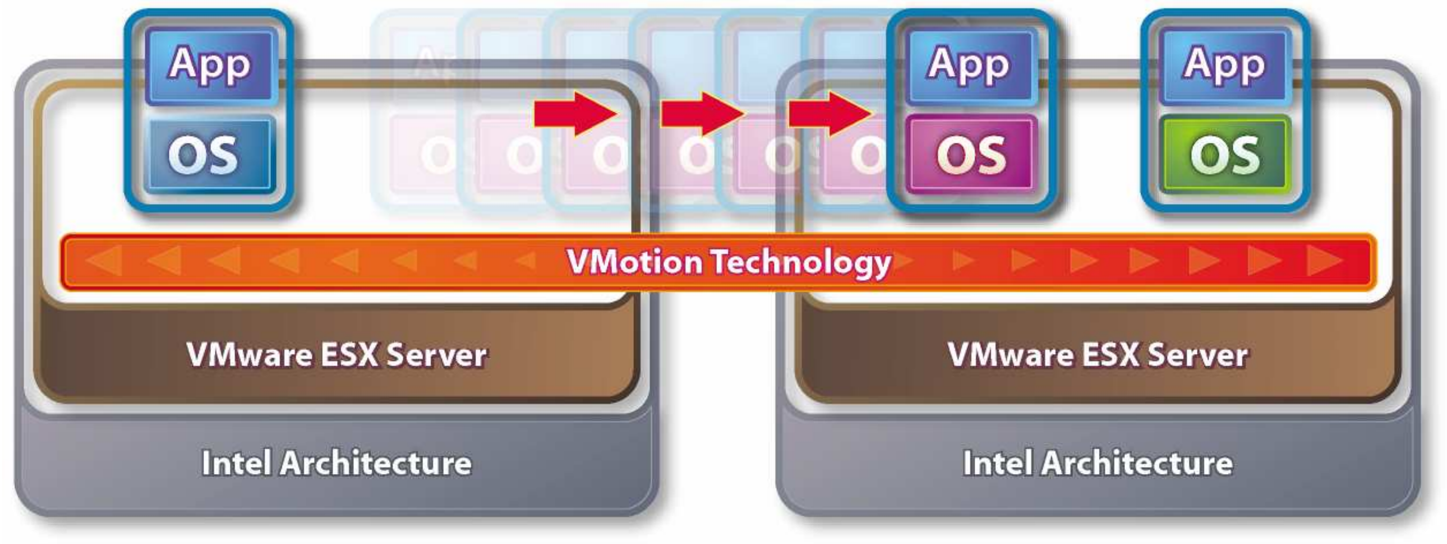

- Proactively migrate systems across an environment to optimize workloads
- Perform zero downtime, rolling hardware upgrades
- Proactively migrate VMs to new hosts in response to hardware failure

VMware Virtual Infrastructure

DRS

HA

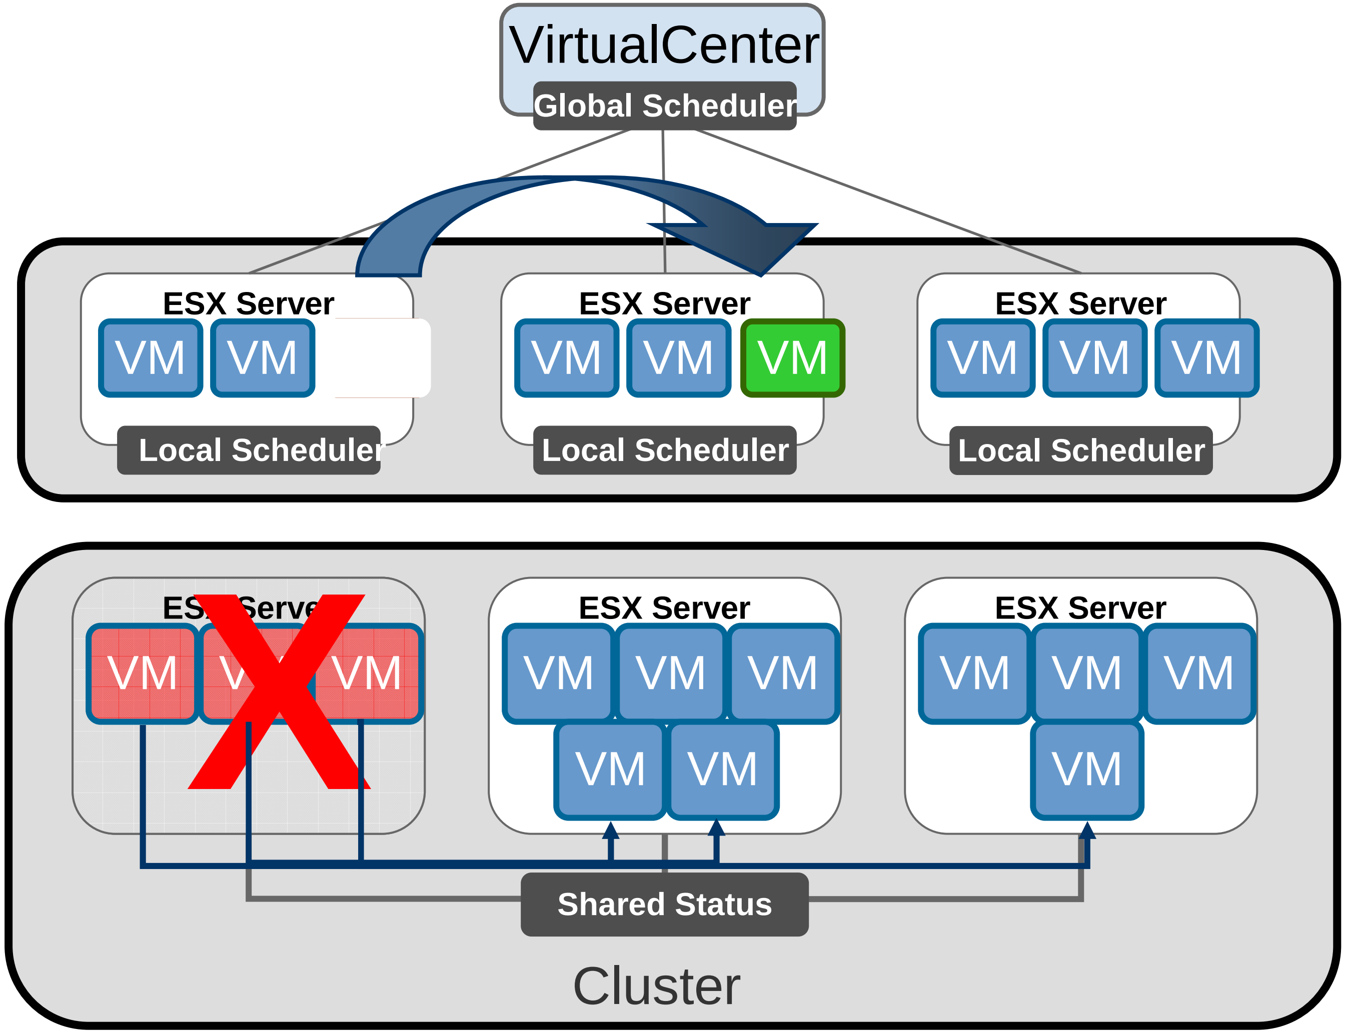

VMware Virtual Infrastructure

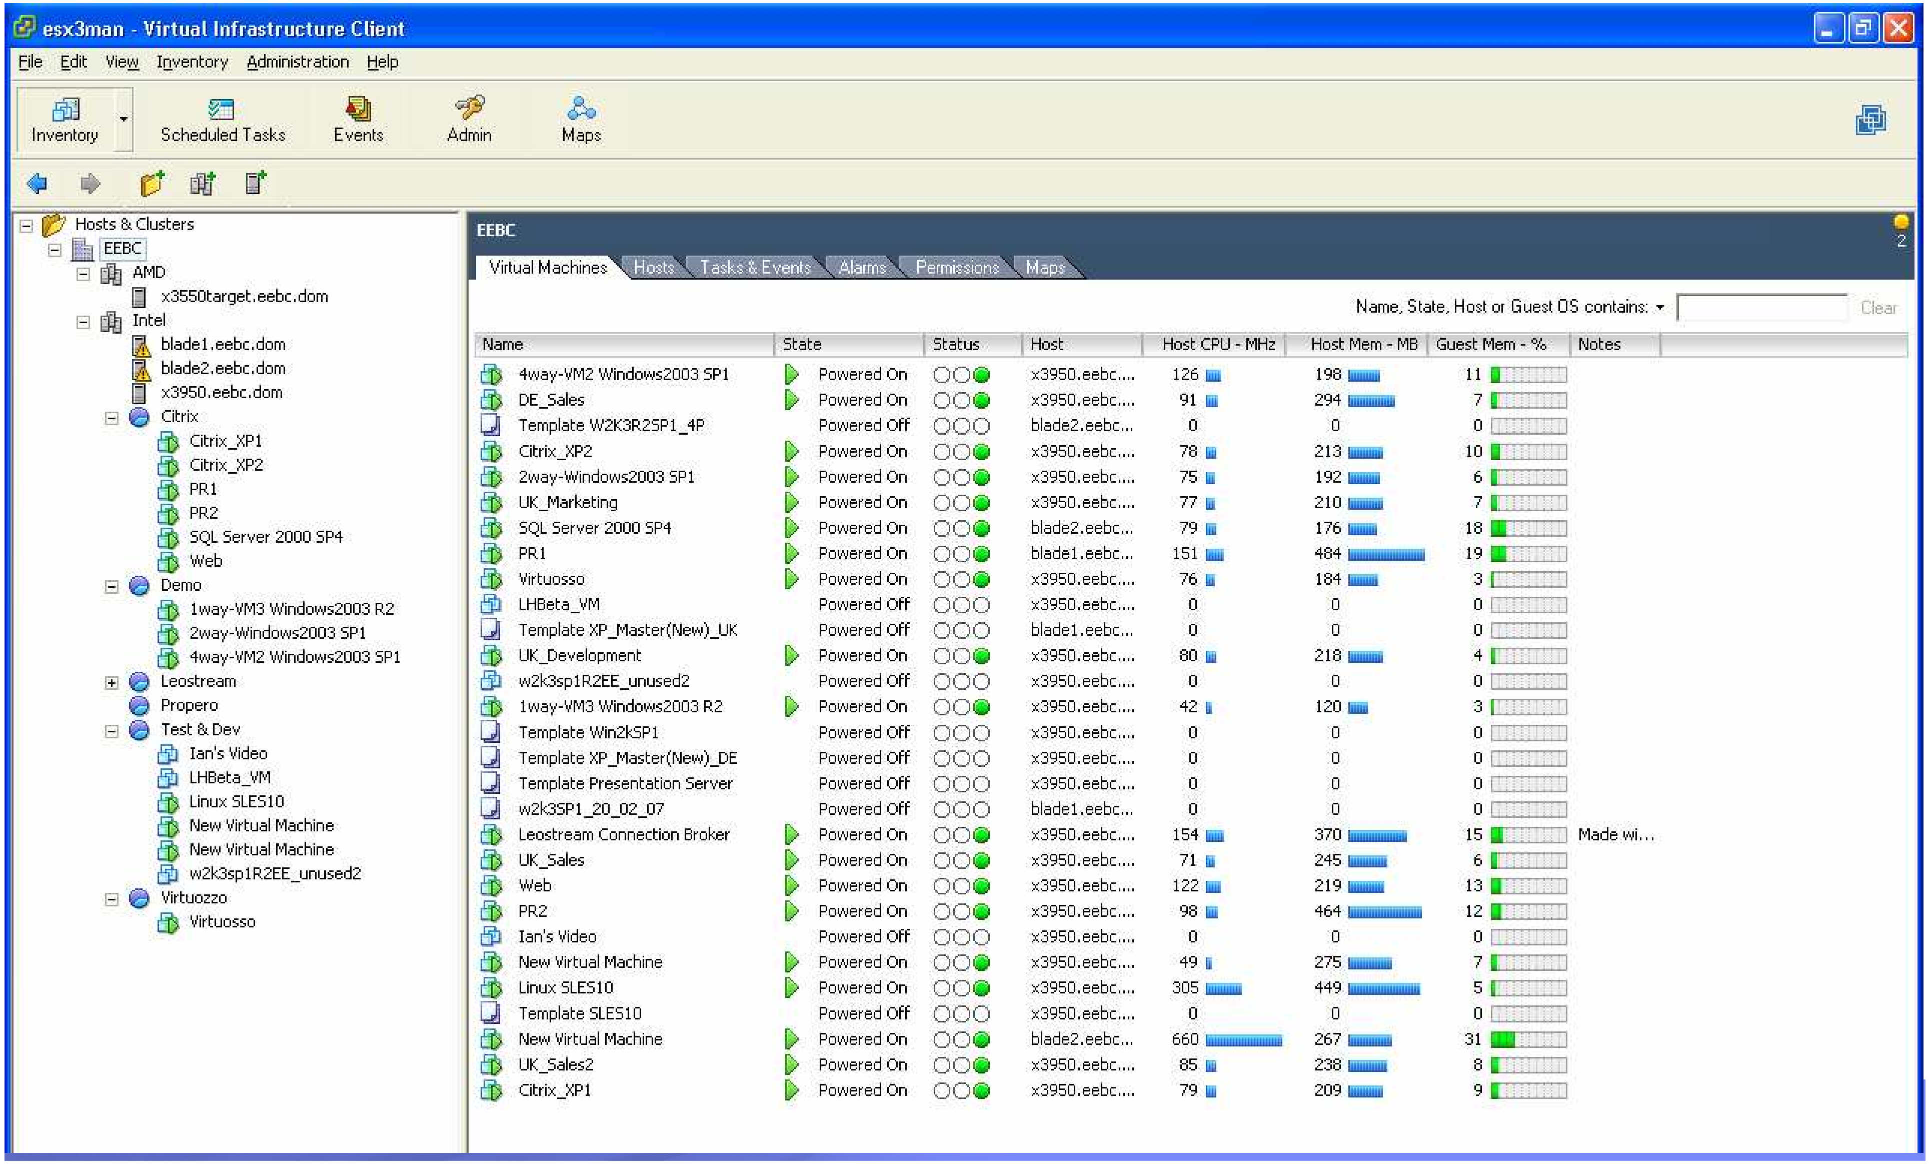

Topics

- Business drivers for server consolidation
- Server virtualization and hypervisor basics
- x86 virtualization
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Some customer examples*
  - *Summary*

System z Virtualization Technology
A Shared Everything Architecture

Start Interpretive Execution
- Establish architecture for guest systems
- Maintain status
- Invoke SIE assists

LPAR Zoning: each partition has a zero-origin address space, allowing I/O access to memory without hypervisor intervention

Hardware support: 10% of circuits are used for virtualization

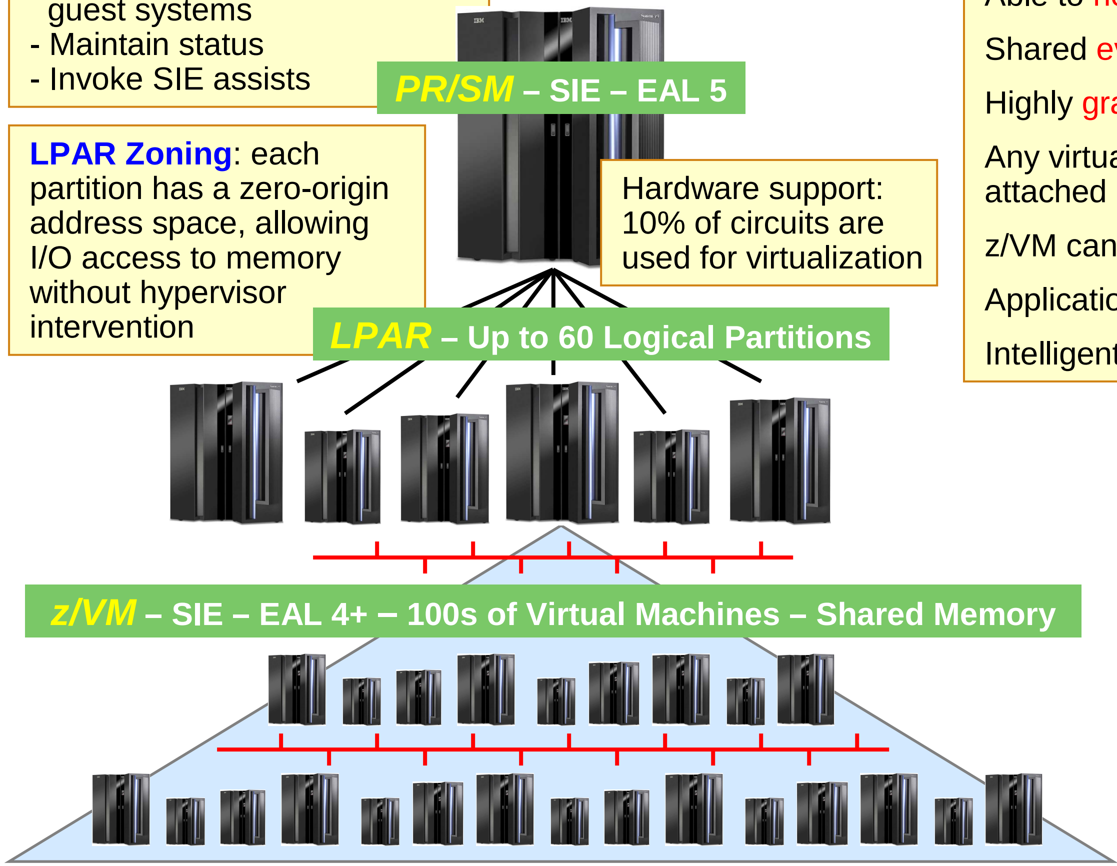

The potential performance impact of the Linux server farm is isolated from the other LPARs

Most sophisticated and functionally complete hypervisors

Able to host z/OS, Linux, z/VSE, z/TPF, and z/VM-on-z/VM

Shared everything architecture

Highly granular resource sharing (less than 1% utilization)

Any virtual CPU can access any virtual I/O path within the attached logical channel subsystem

z/VM can simulate devices not physically present

Application integration with HiperSockets and VLANs

Intelligent and autonomic workload management

Shared resources per mainframe footprint

Up to 64 OS-configurable CPUs

Up to 11 SAP processors

Up to 1.5 TB of memory

Up to 1024 channel paths

Up to 16 internal HiperSockets networks

HW (LPAR) and SW (z/VM) hypervisors

Hardware support, SIE, microcode assist

Virtualization is transparent for Op Sys execution

Hardware-enforced isolation

IBM System z Virtualization Leadership

Extreme Levels of CPU Sharing

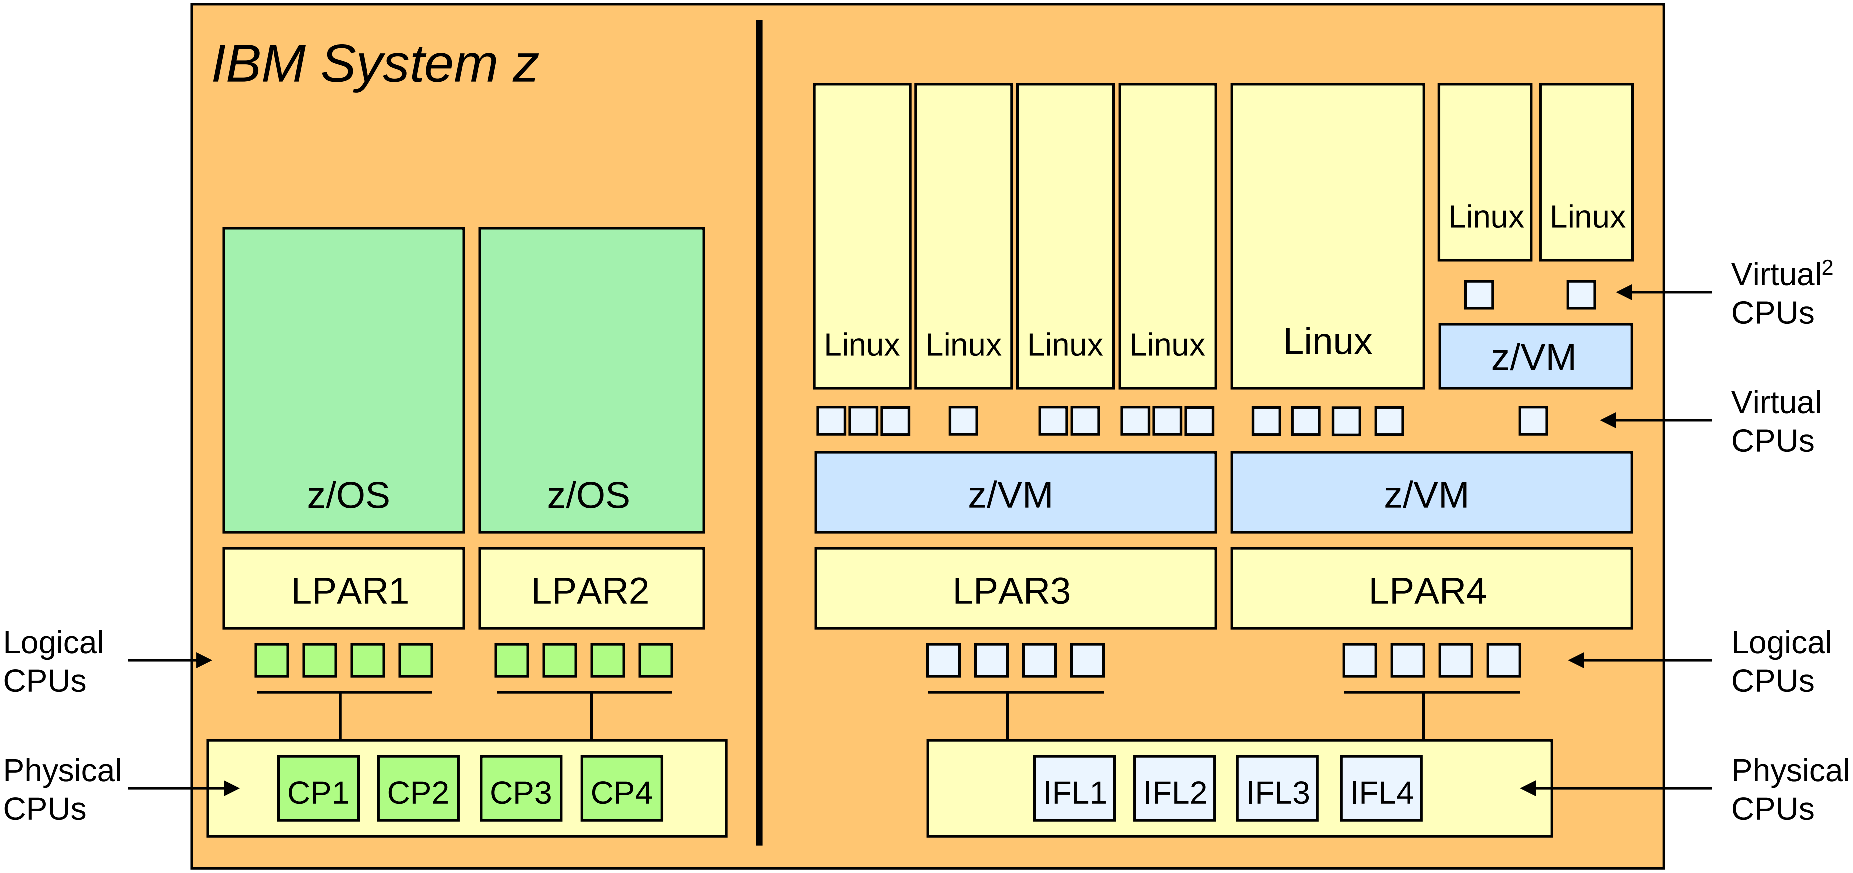

IBM System z CPU High Availability
Concurrent Processor Reassignment

- **Used to concurrently change the physical backing of one or more logical processors**
- **The state of source physical processor is captured and transplanted into the target physical processor**
- **Operation is transparent to operating systems**
- **Used for *processor sparing* and *book replacement***

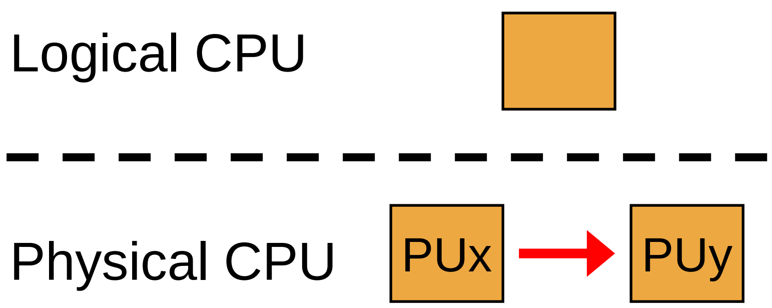

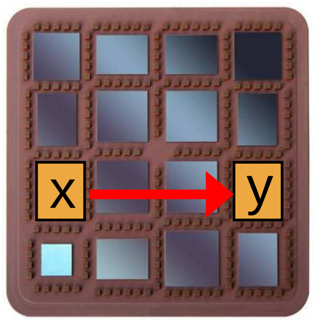

System Design Affects Virtualization Capabilities

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

➲ With TCO-friendly pricing

Up to 64-way SMP

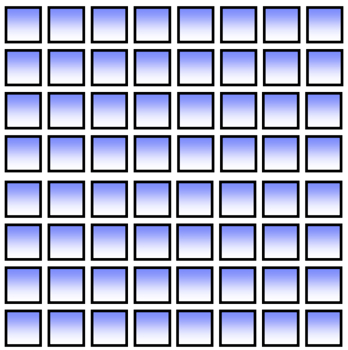

Share up to 64 processors with up to 60 LPARs

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

* No software license fees

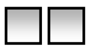

2 Standard Spare PUs

Up to 11 System Assist Processors

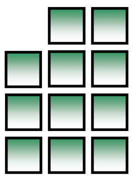

Offload system processing to dedicated CPUs (no impact to software license fees)

Up to 336 I/O Processors

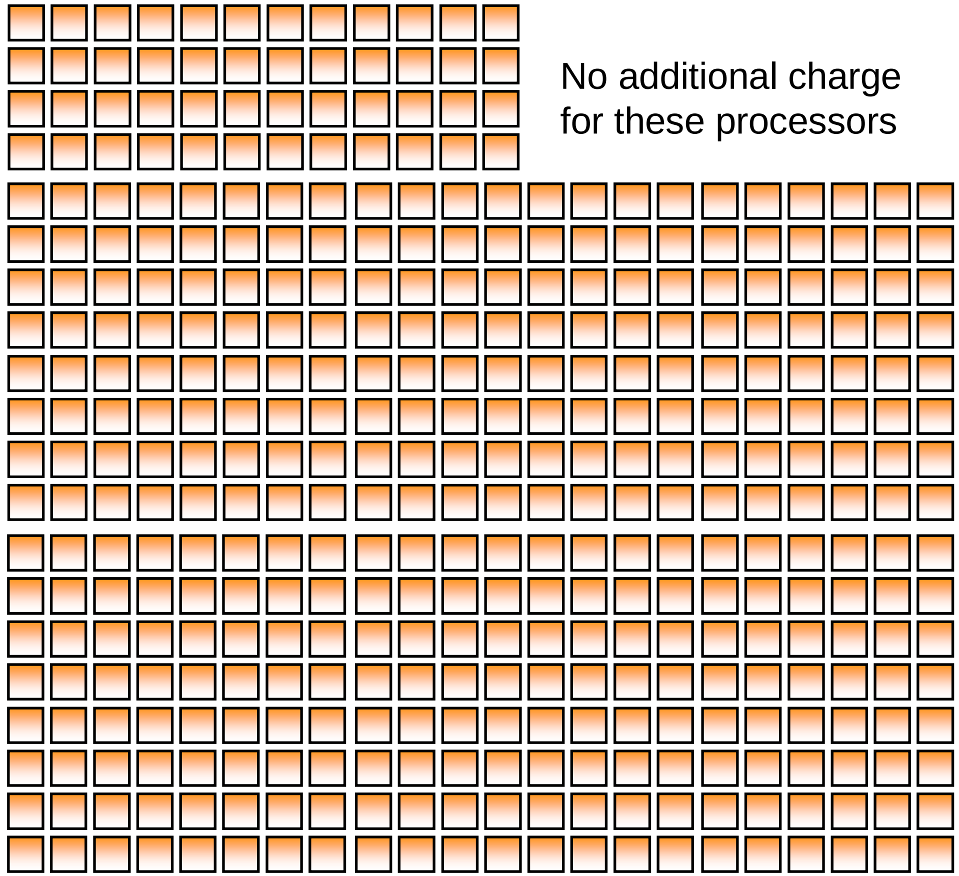

No additional charge for these processors

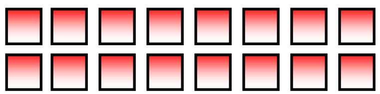

Up to 16 Crypto Express2 CPUs

High scale performance for SSL transactions

System Design Affects Virtualization Capabilities

Compare to typical UNIX system design...

Up to 128-way SMP configuration

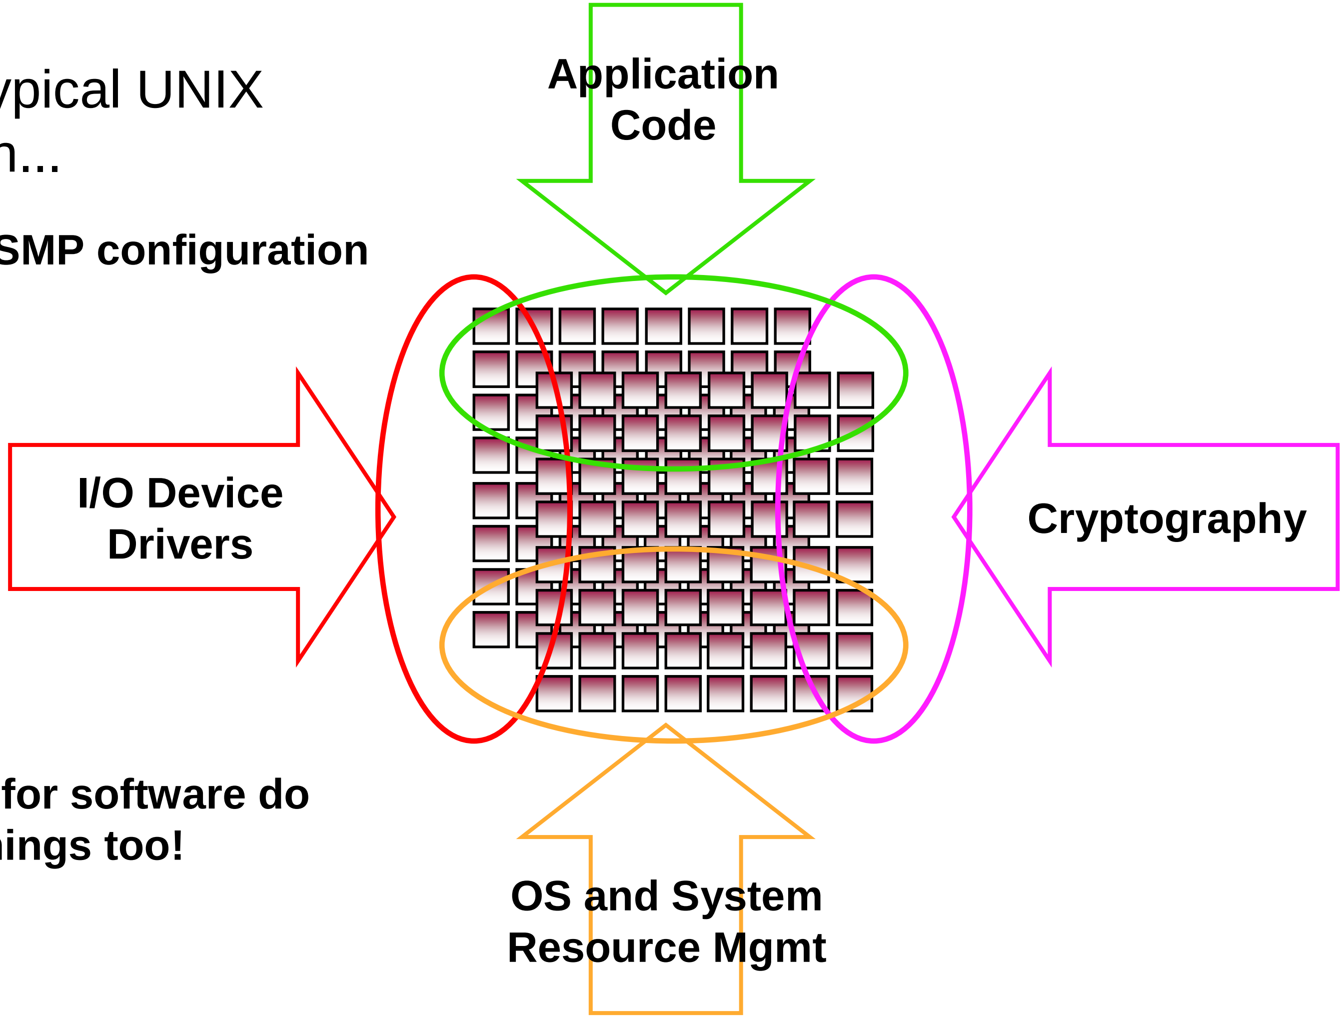

CPUs licensed for software do a lot of other things too!

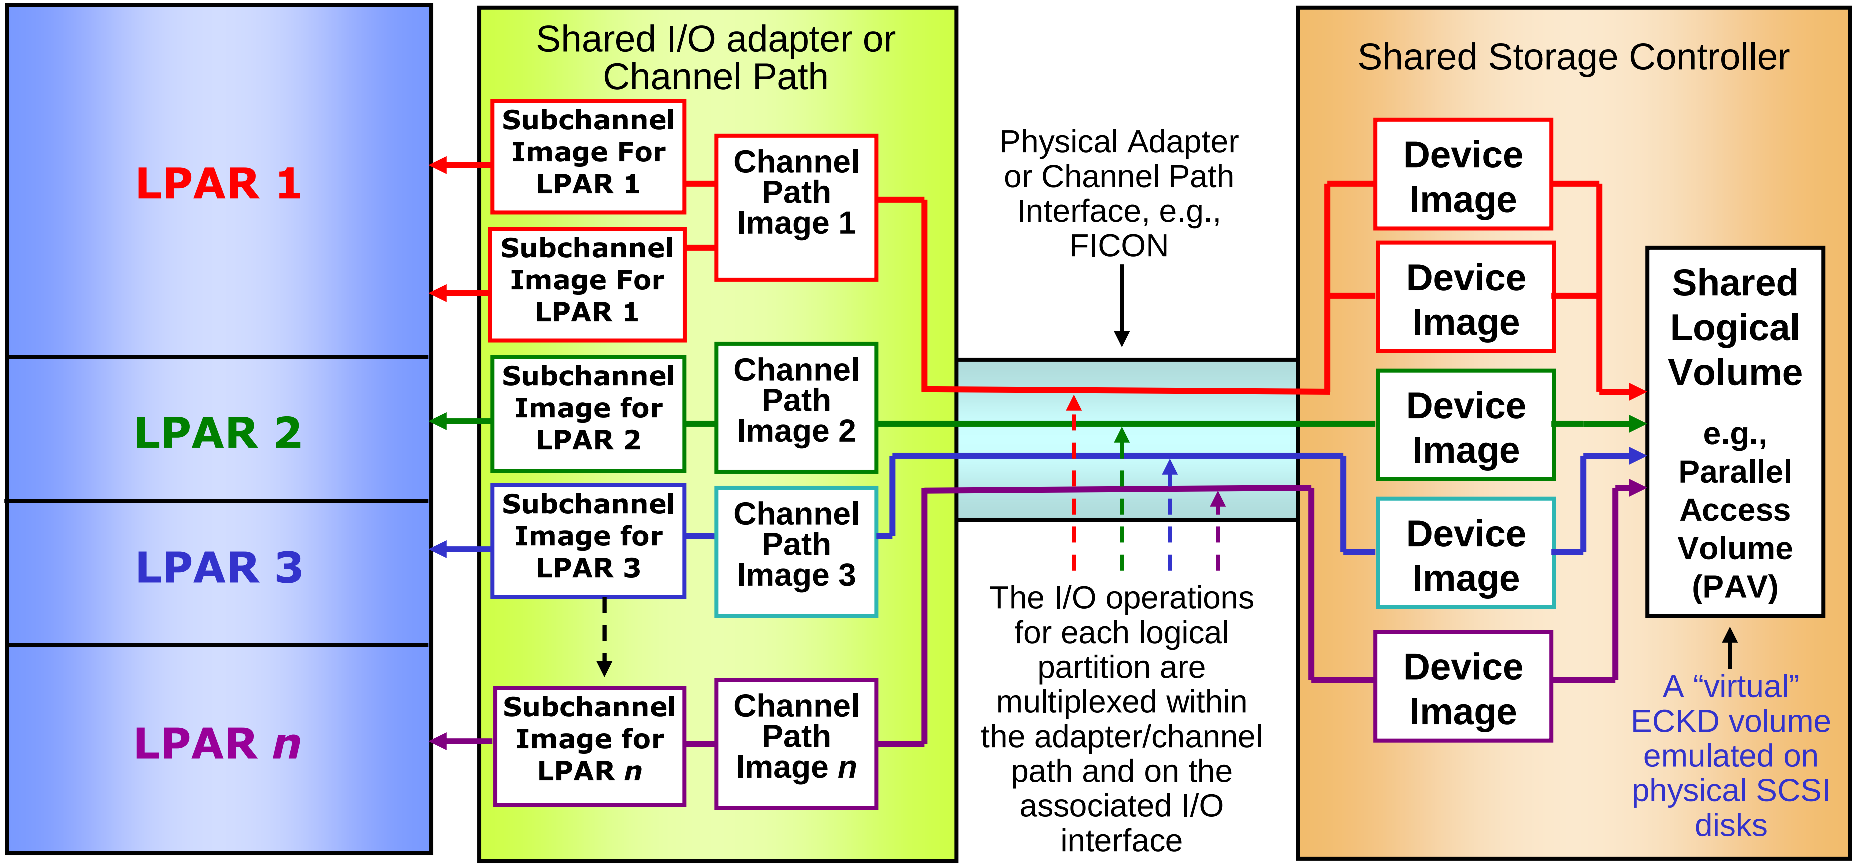

- The I/O infrastructure is shared by LPARs at native speeds, without hypervisor involvement
- Up to 8 physical channels process the I/O requests to the shared devices
  - This reduces the possibility of I/O queuing delays at the channels or at the shared storage controller

z/VM – An Exceptional Virtualization Platform

z/VM can massively scale a virtual server environment with a mix of virtual and real resources for each virtual machine

- With exceptional levels of performance, availability, and security
- Virtual and real assets can be non-disruptively added when needed
- Over 40 years of continuous innovation in virtualization

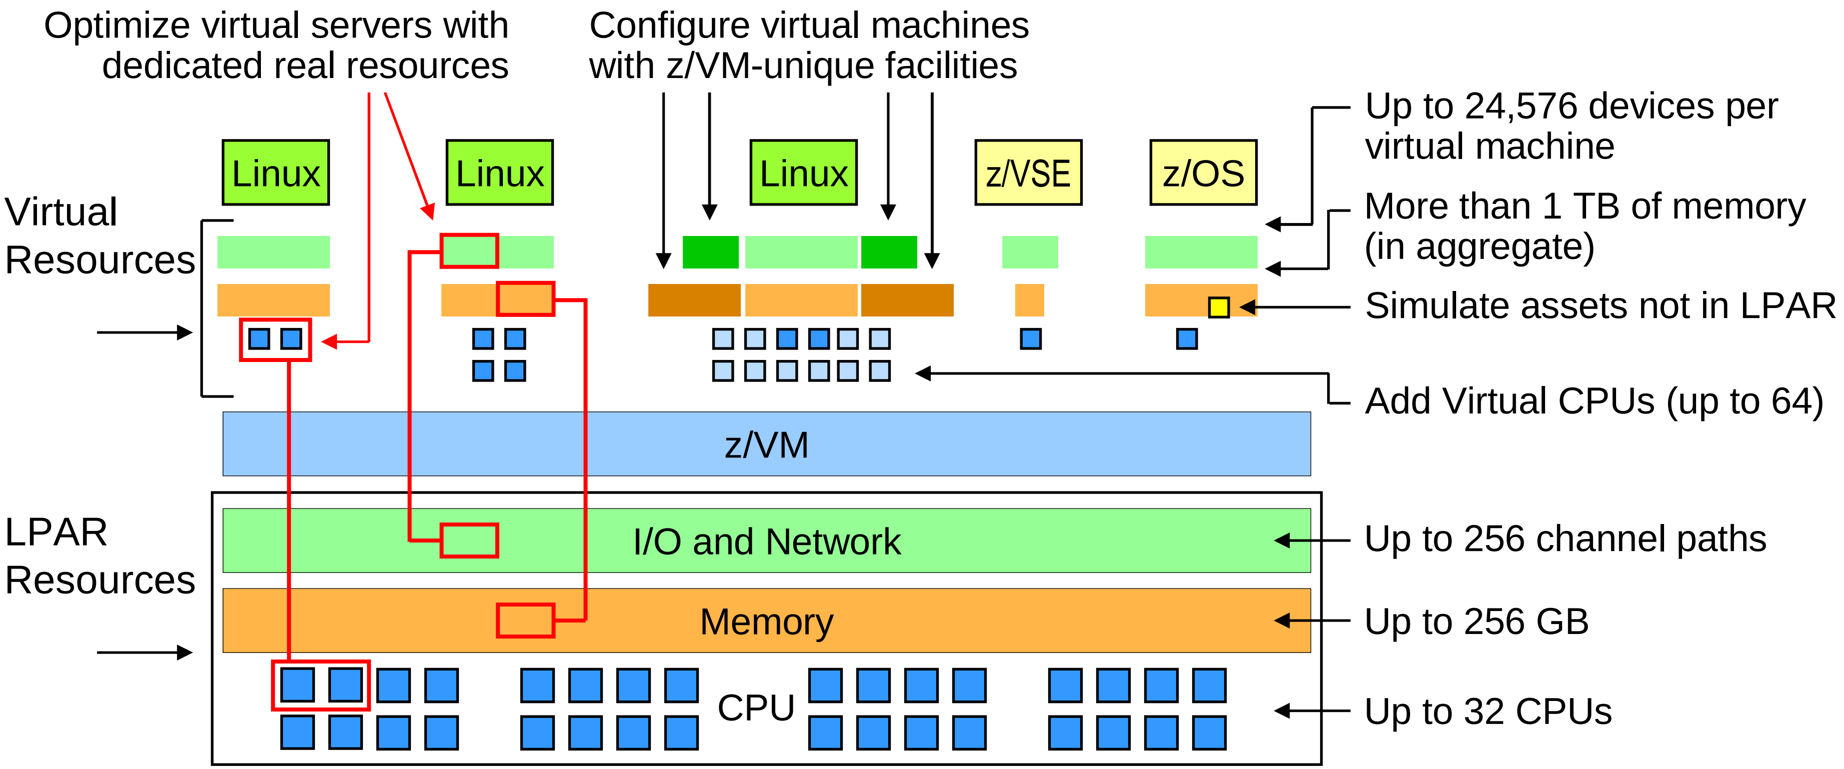

z/VM CPU Resource Controls

Highly Granular Sharing of System Resources

- Allocate system resources per guest image using SHARE command
  - This is a highly flexible and self-managed function of the z/VM Control Program
  - Reserve CPU capacity for peak usage
    - Use it when needed
    - Relinquish the processor cycles for other servers when not needed
  - "Absolute guests" receive top priority
  - The Virtual Machine Resource Manager can be used to monitor and adjust remaining capacity allocated to "Relative guests“
  - Also use VMRM to prioritize I/O operations among guest images via “I/O Priority Queuing”

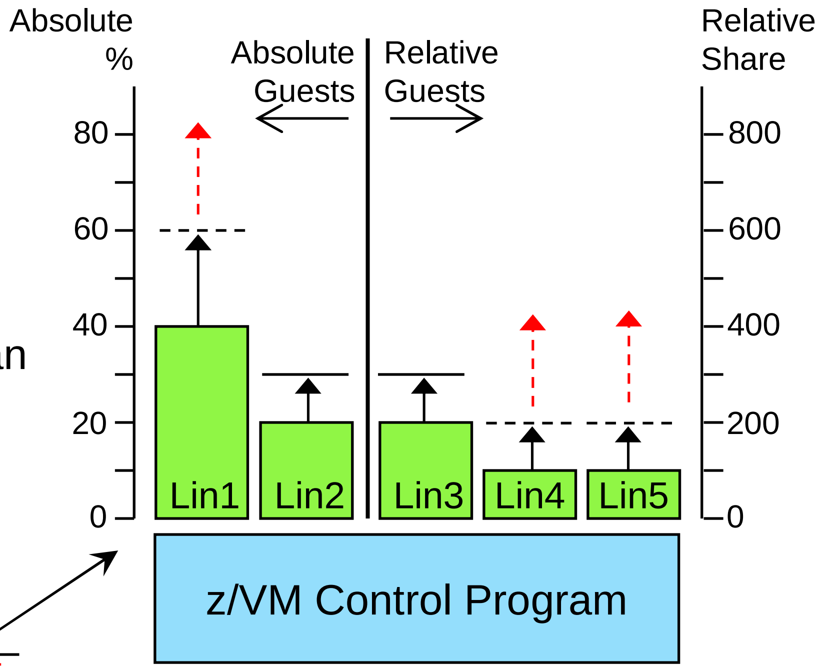


```
SHARE Lin1 ABSOLUTE 40% ABSOLUTE 60% LIMITSOFT
SHARE Lin2 ABSOLUTE 20% ABSOLUTE 30% LIMITHARD
SHARE Lin3 RELATIVE 200 RELATIVE 300 LIMITHARD
SHARE Lin4 RELATIVE 100 RELATIVE 200 LIMITSOFT
SHARE Lin5 RELATIVE 100 RELATIVE 200 LIMITSOFT
```

Notes:

- - - - - = limit can be exceeded if unused capacity is available (LIMITSOFT)

——— = limit will not be exceeded (LIMITHARD)

Resource Sharing and Scalability

Scale Up and Out with Linux on z/VM

- **With z/VM you can grow horizontally and vertically on the same System z server...dynamically**
- **Provision a virtual machine for peak utilization and allocate its resources to other servers during off-peak hours... automatically**

Server sprawl on z/VM is a good thing!

Add more resources to existing server non-disruptively...

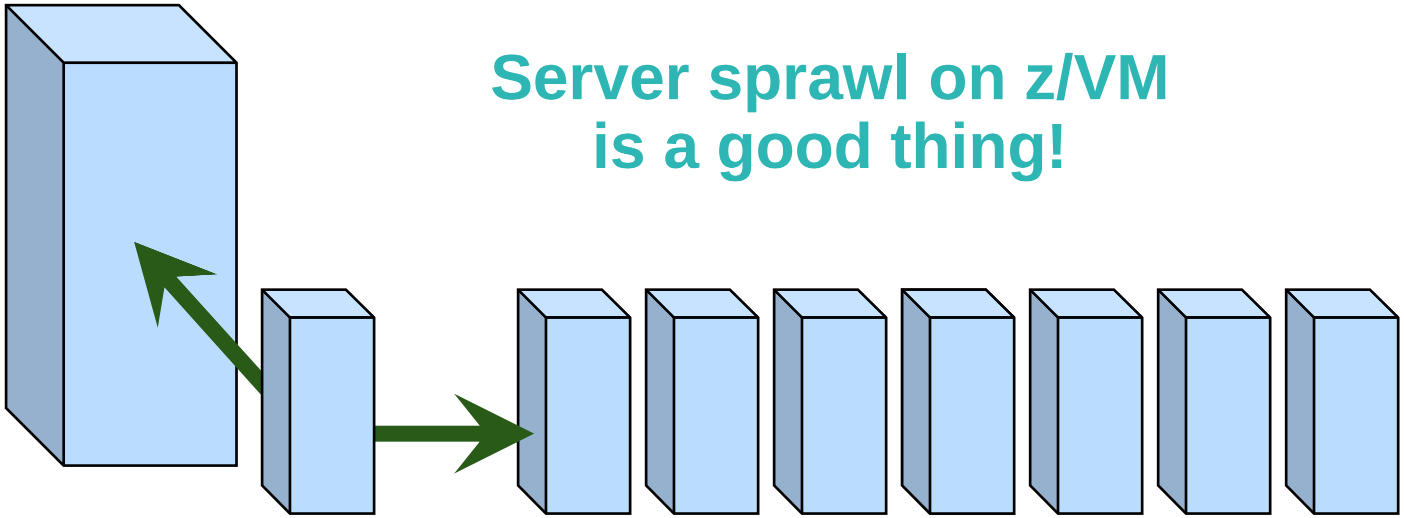

...or clone more servers *with a high degree* of resource sharing.

Single-System, Multi-LPAR, Linux-on-z/VM Environment

Maximizing Resource Utilization and System Availability

- ✓ Run multiple copies of z/VM on a single System z for enhanced scalability, failover, operations, and energy efficiency
- ✓ Share CPUs and I/O adapters across all z/VM LPARs, and over-commit memory in each LPAR for added cost effectiveness

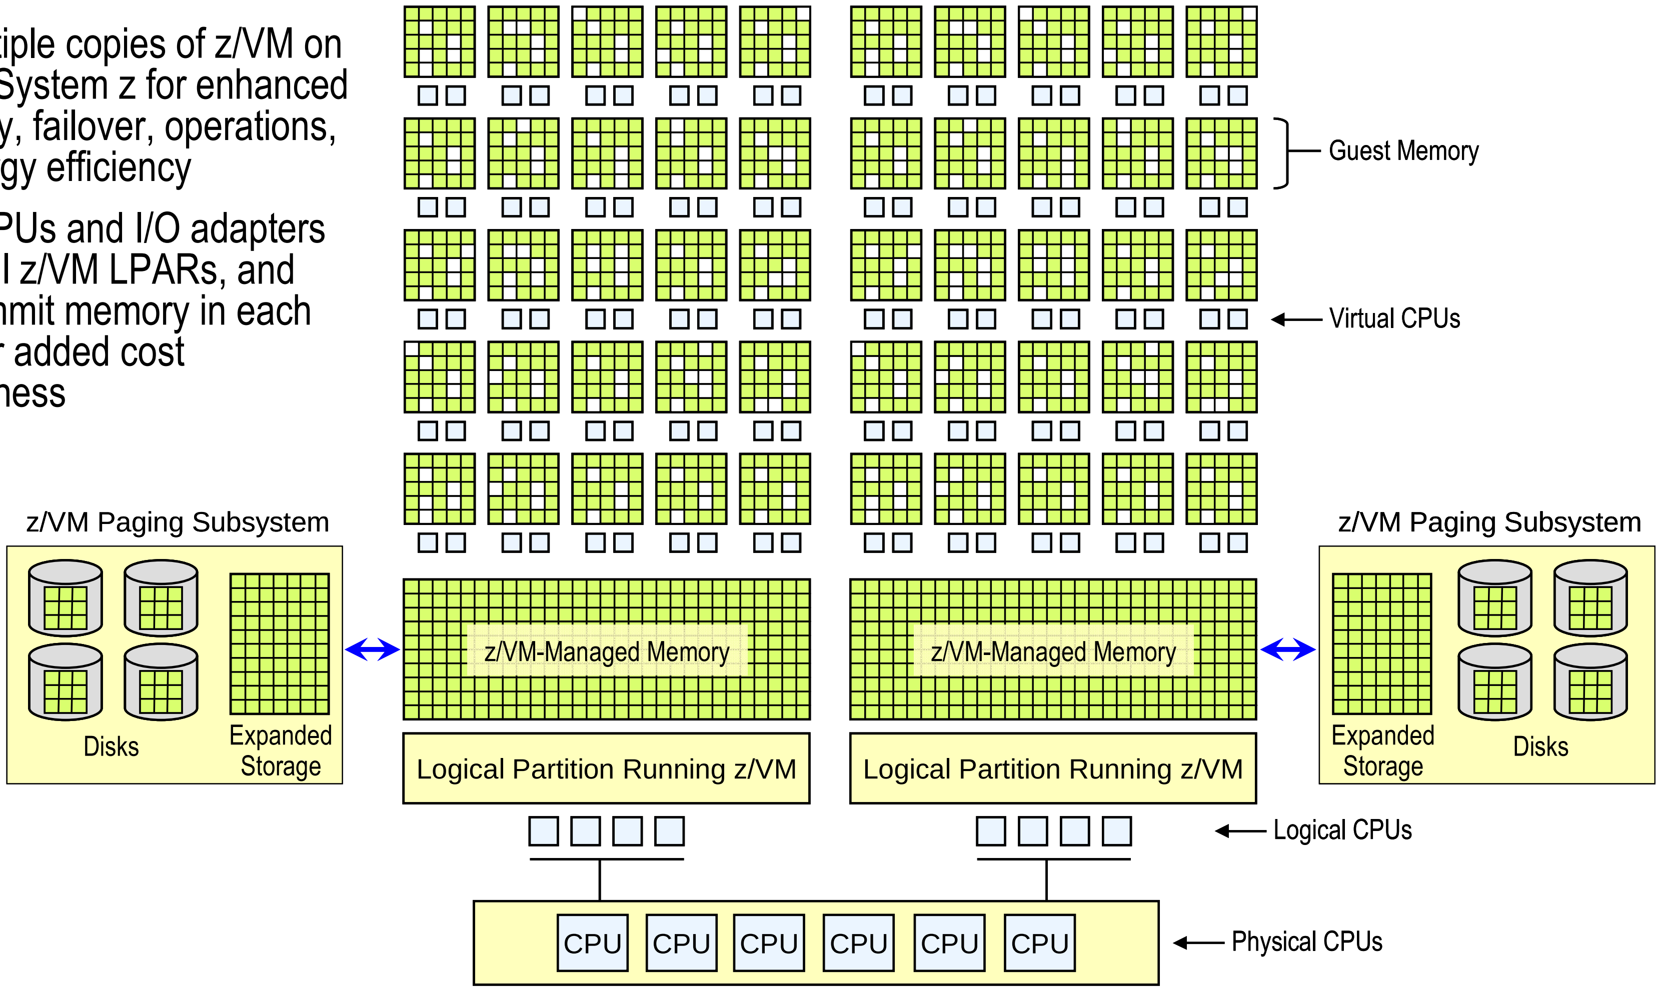

Linux-on-z/VM and *Flexible, Efficient Growth*

z/VM V5.4 Function Enhances System Availability

- **Clients can start small with Linux on System z and non-disruptively grow their environment as business dictates**
- **Users can dynamically add CPUs, memory, I/O adapters, devices, and network cards to a running z/VM LPAR**
- **z/VM virtualizes this capability for guest machines**

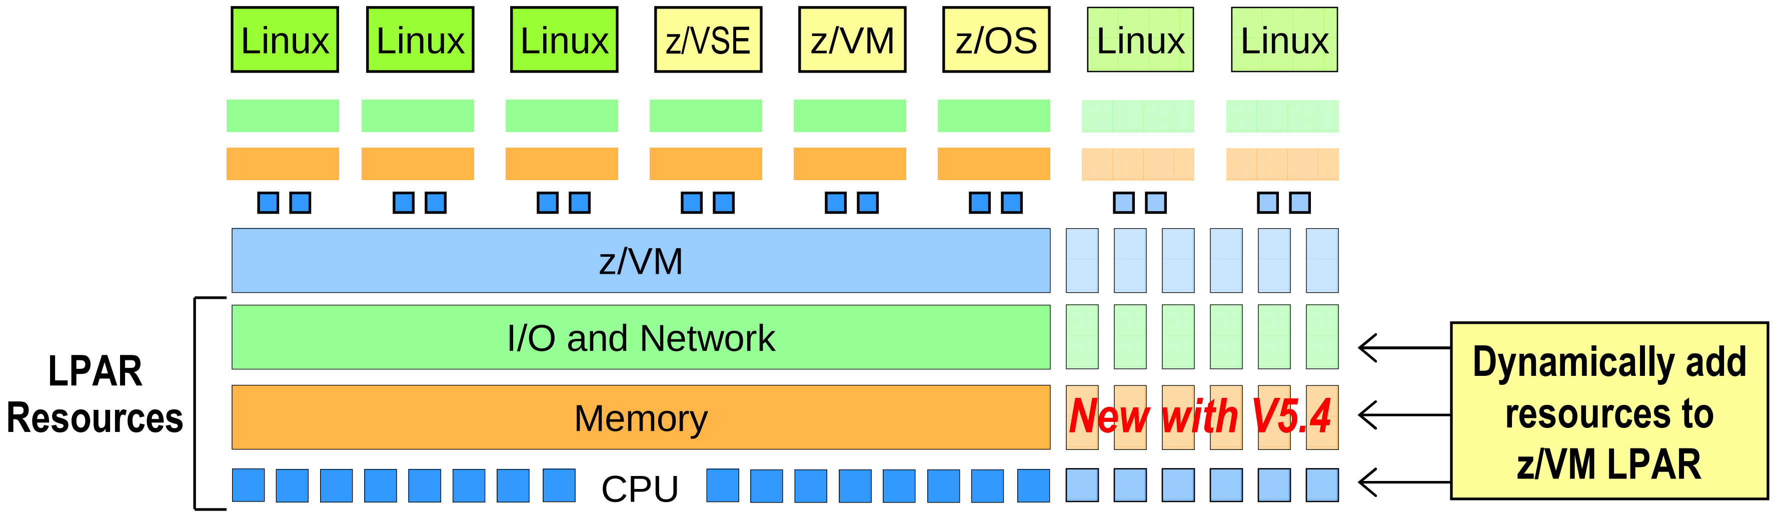

Smart economics: non-disruptively scale your z/VM environment by adding hardware assets that can be shared with *every* virtual server

Extreme Virtualization with Linux on z/VM
VMRM Cooperative Memory Management (VMRM-CMM)

- Problem scenario: virtual memory utilization far exceeds real memory availability
- Solution: real memory constraint corrected by z/VM *Virtual Machine Resource Manager*
- Linux images signaled to reduce virtual memory consumption
- Demand on real memory and z/VM paging subsystem is reduced
- Helps improve overall system performance and guest image throughput

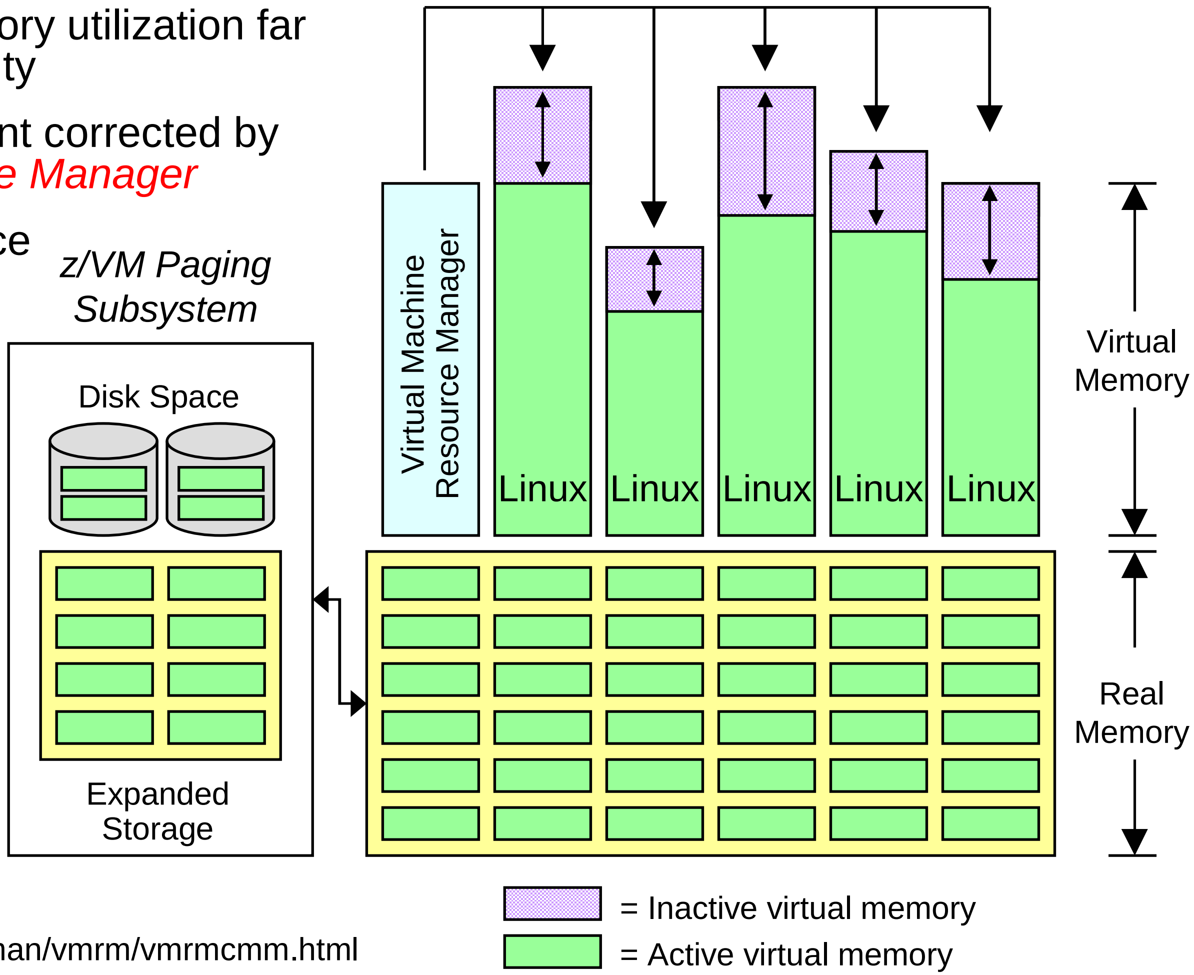

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

Linux and z/VM Technology Exploitation

Collaborative Memory Management Assist (CMMA)

- Extends coordination of memory and paging between Linux and z/VM to the level of individual pages using a new hardware assist (*CMMA*) on z9 and z10 systems
- z/VM knows when a Linux application has released a page of memory
- Host Page-Management Assist (*HPMA*), in conjunction with CMMA, further reduces z/VM processing needed to resolve page faults
- Can help z/VM host more virtual servers in the same amount of memory
- Supported by z/VM V5.3 and later and Novell SLES 10
- Full support not accepted by Linux community
- Future distributions likely to implement a subset of function
  - Supporting "Stable" and "Unused" states

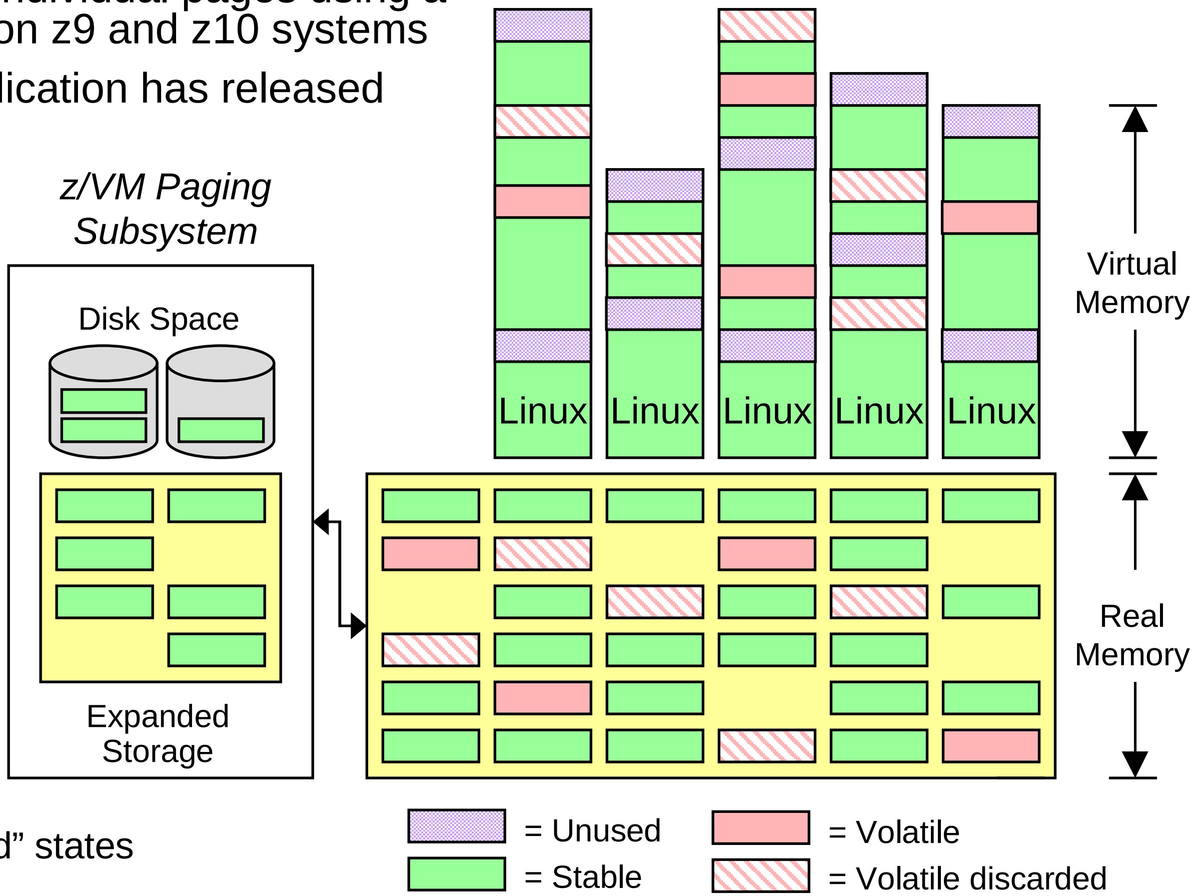

Transaction Rate versus Number of Hosted Servers

*Apache Servers with 1GB of Memory Each – z/VM with 8GB of Memory**

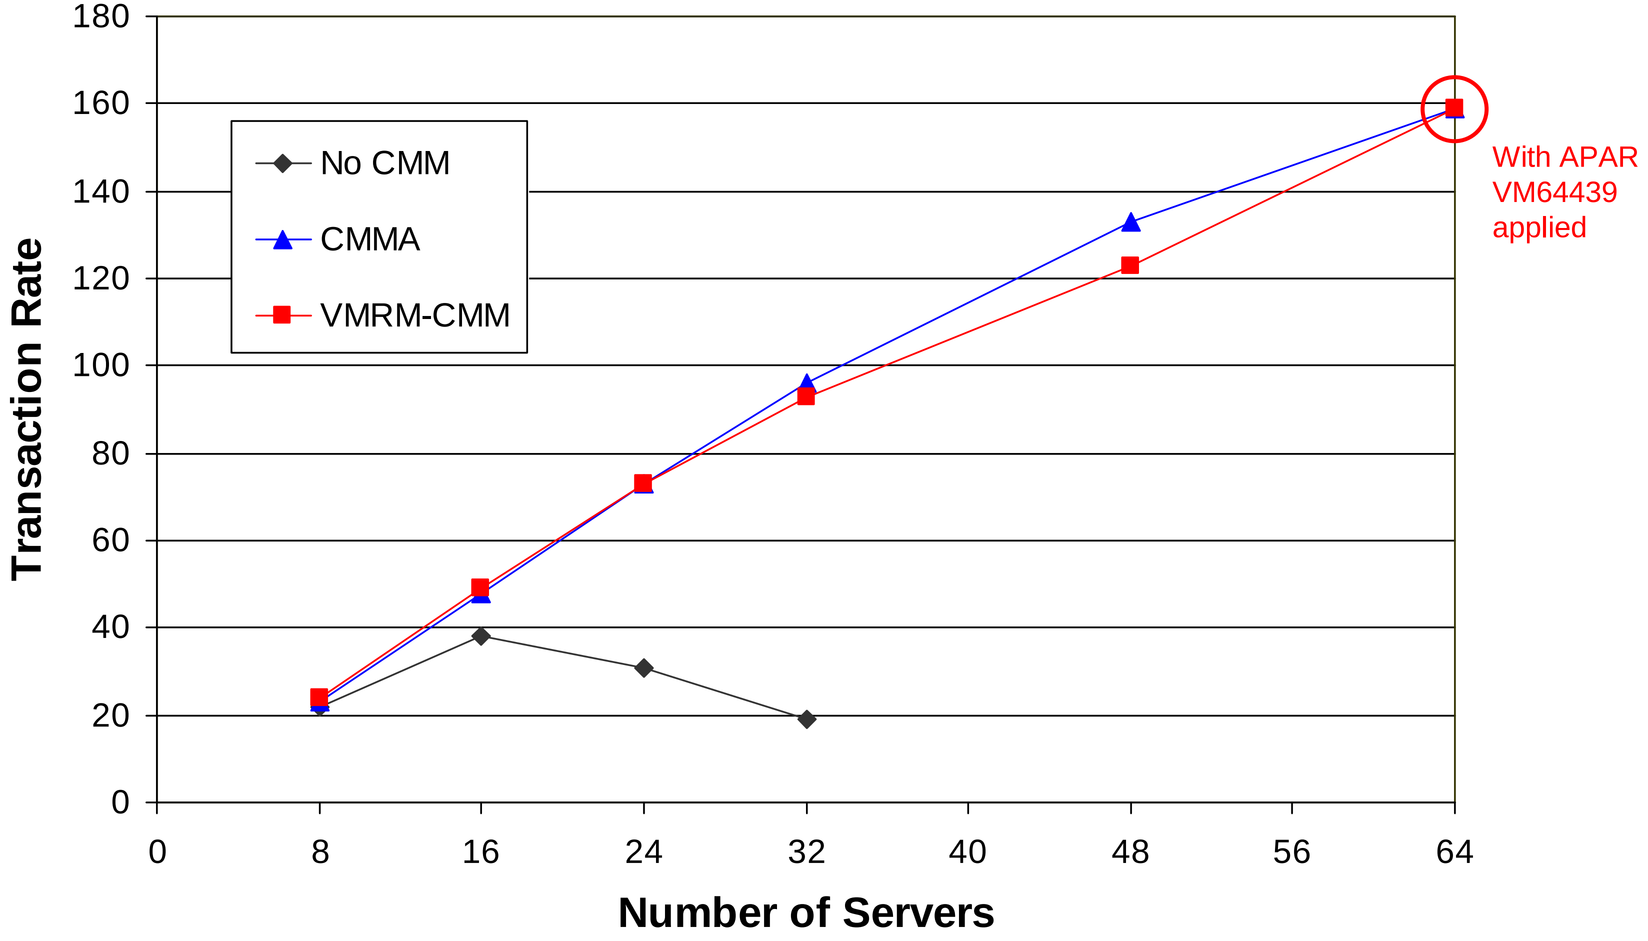

* z/VM running in IBM System z9 LPAR with 6GB of Central Storage and 2GB of Expanded Storage

z/VM Technology: Advanced Disk Support

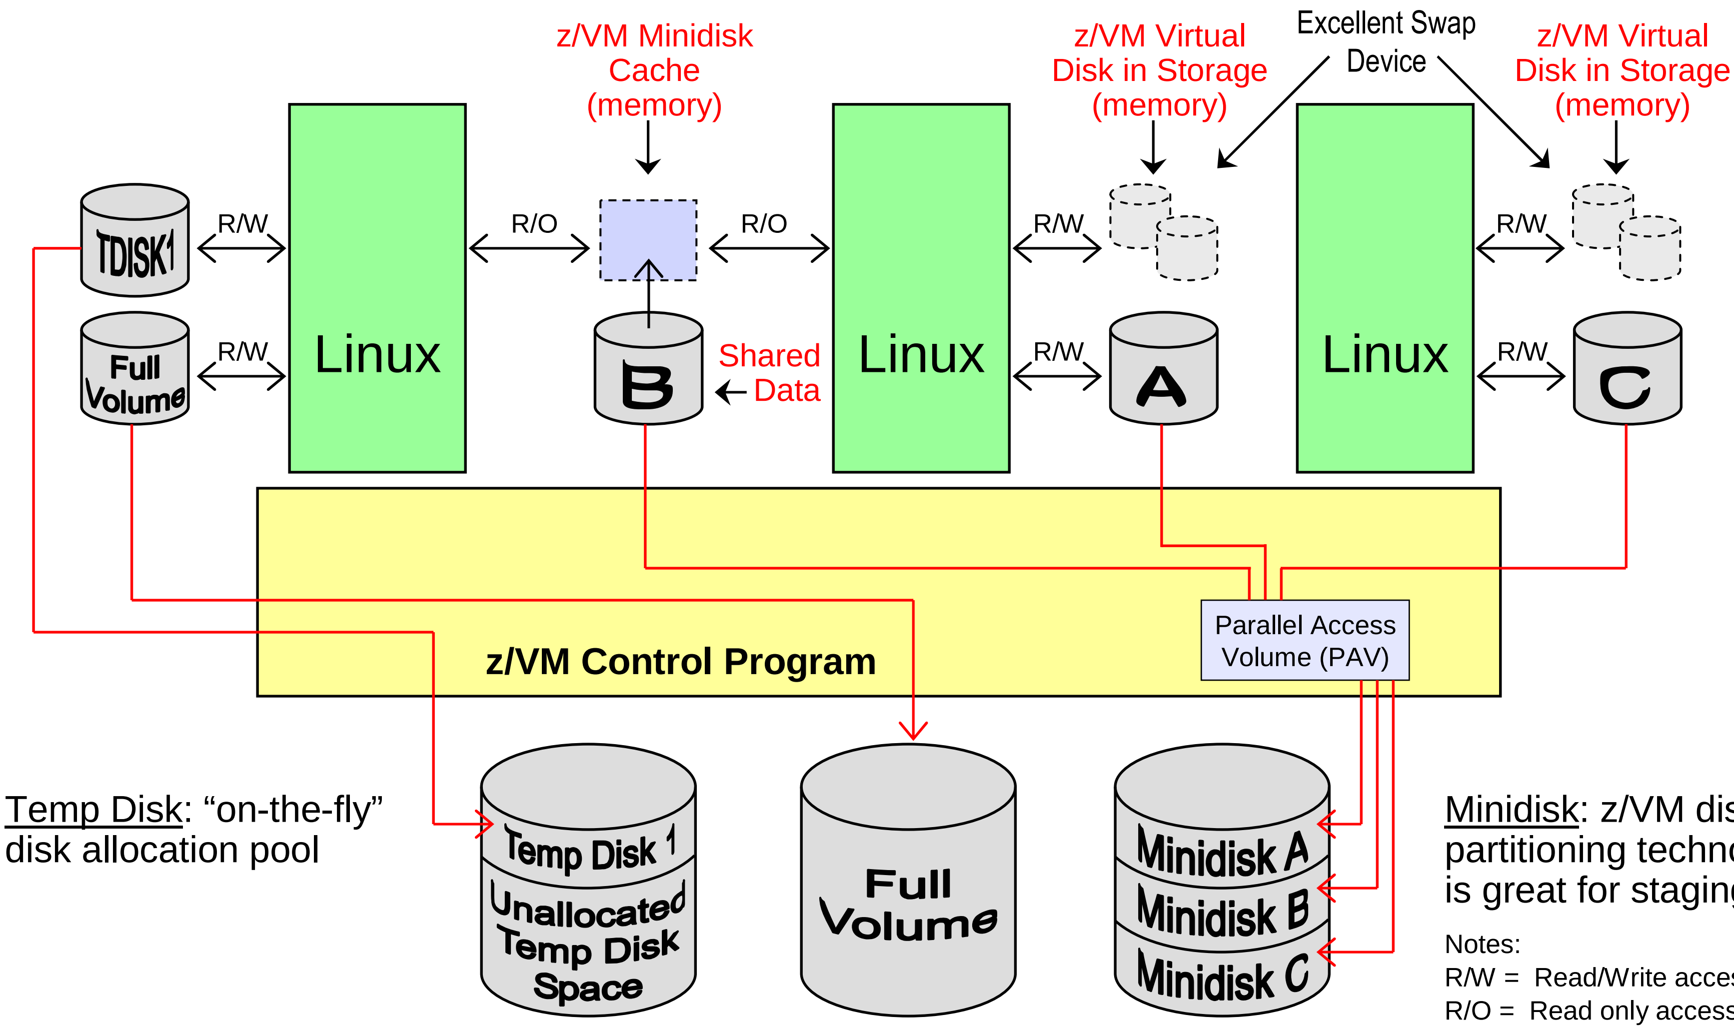

Temp Disk: “on-the-fly” disk allocation pool

Minidisk: z/VM disk partitioning technology is great for staging

Notes:
R/W = Read/Write access
R/O = Read only access

Extreme Linux-on-z/VM Virtualization

Linux Exploitation of z/VM DCSS Support

- Discontiguous Saved Segments (DCSS)
  - Share a single, real memory location among multiple virtual machines
  - Can reduce real memory utilization
- Linux exploitation: shared program executables
  - Program executables are stored in an execute-in-place file system, then loaded into a DCSS
  - DCSS memory locations can reside outside the defined virtual machine configuration
  - Access to file system is at memory speeds; executables are invoked directly out of the file system (no data movement required)
  - Avoids duplication of virtual memory
  - Helps enhance overall system performance and scalability
- **z/VM V5.4 support enhancements:**
  - Segments can reside above 2 GB address line
  - Enables even greater system scalability
  - New addressing limit is 512 GB

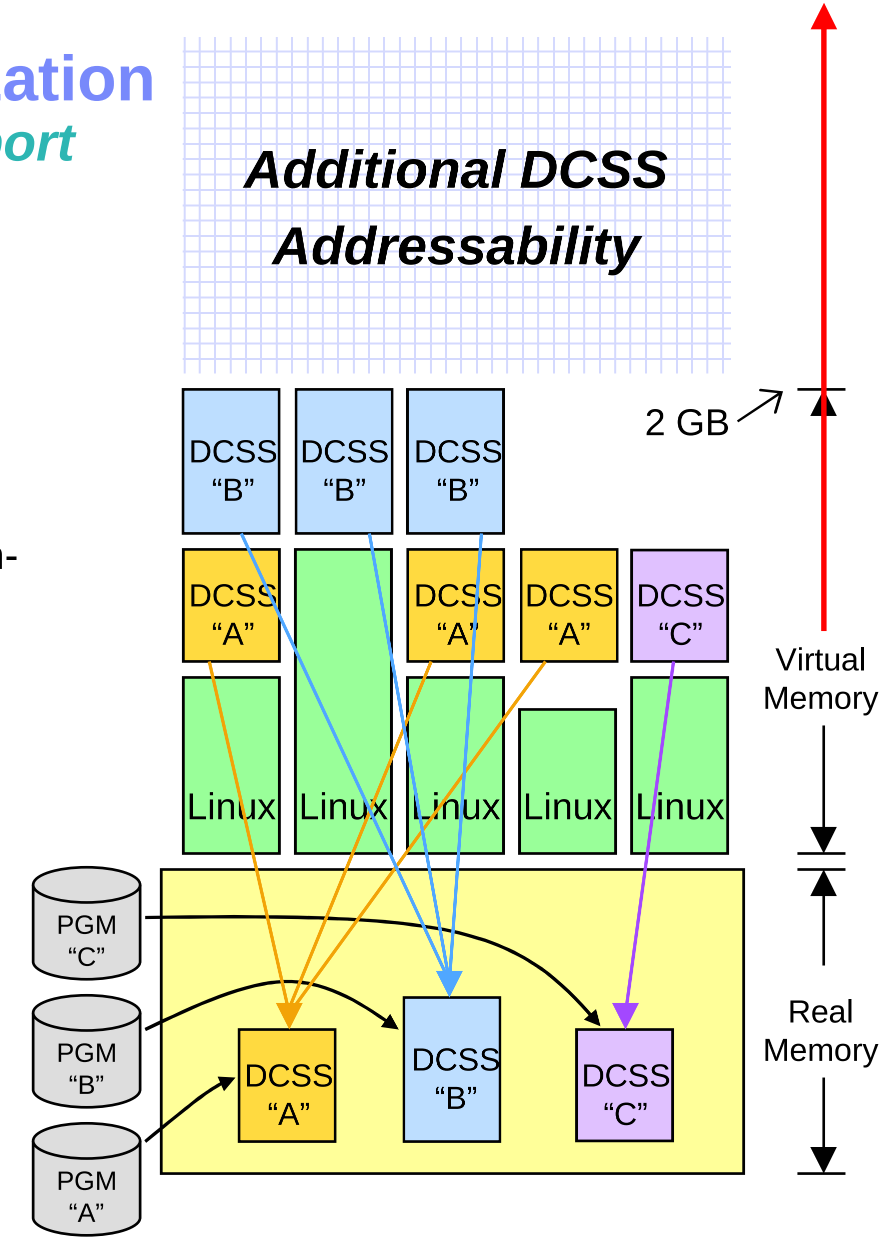

Note: Maximum size of a single DCSS is 2047 MB

Extreme Virtualization with Linux on z/VM
Linux Exploitation of z/VM Virtual Disks in Storage (VDISK)

- VDISK support is Data-in-Memory technology
  - Simulate a disk device using real memory
  - Achieve memory speeds on disk I/O operations
  - VDISKs can be shared among virtual machines
- Linux exploitation: high-speed swap device
  - Use VDISKs for Linux swap devices instead of real disk volumes
  - Reduces demand on I/O subsystem
  - Helps reduce the performance penalty normally associated with swapping operations
  - An excellent configuration tool that helps clients minimize the memory footprint required for virtual Linux servers
  - Helps improve the efficiency of sharing real resources among virtual machines

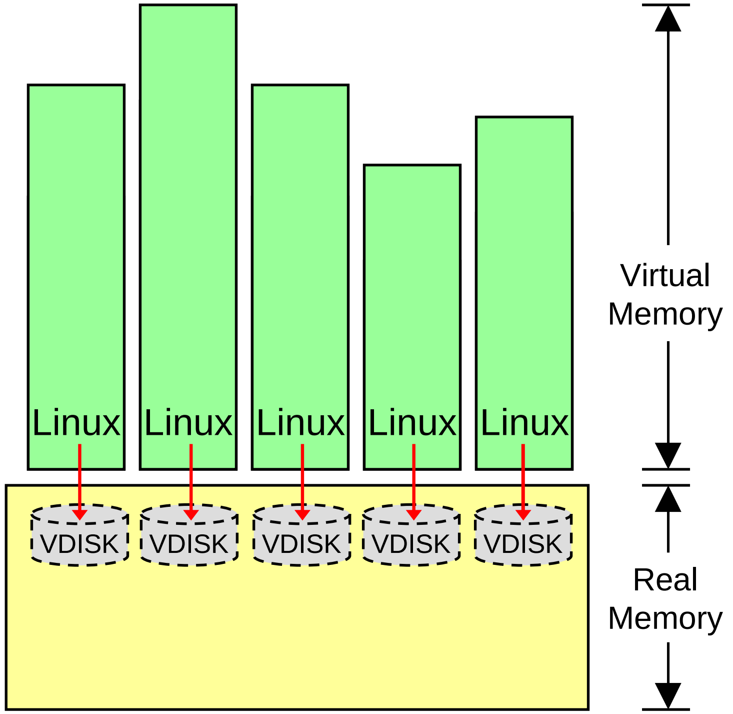

System z and N_Port ID Virtualization (NPIV)

Without N_Port ID Virtualization

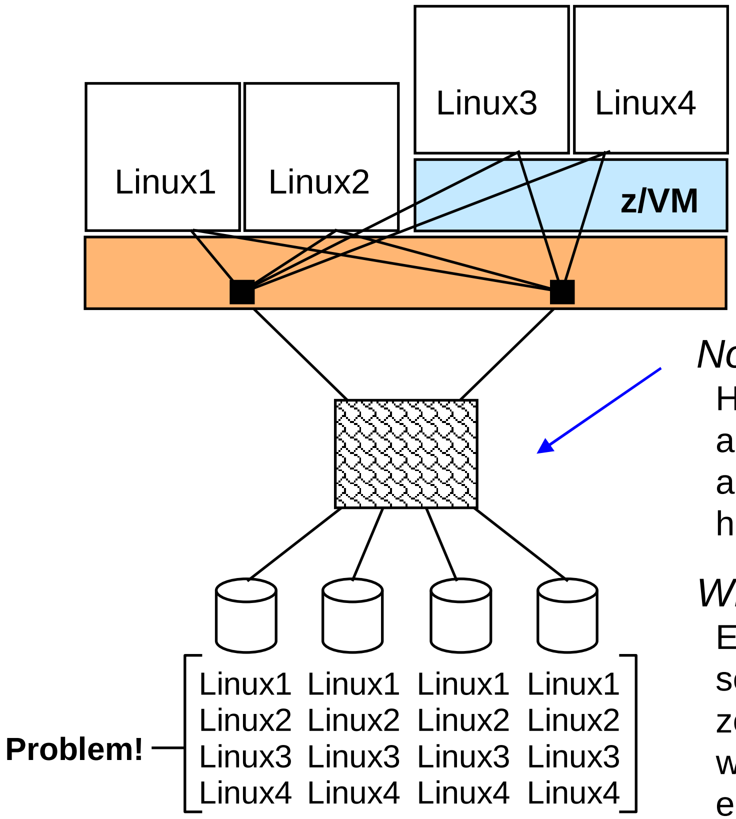

No NPIV:
Hosted Linux images can access all the LUNs that are accessible to the real hardware channels.

With NPIV:
Each Linux image is separately authorized via zoning and LUN-masking with a unique WWPN for each subchannel or virtual host-bus adapter.

With N_Port ID Virtualization

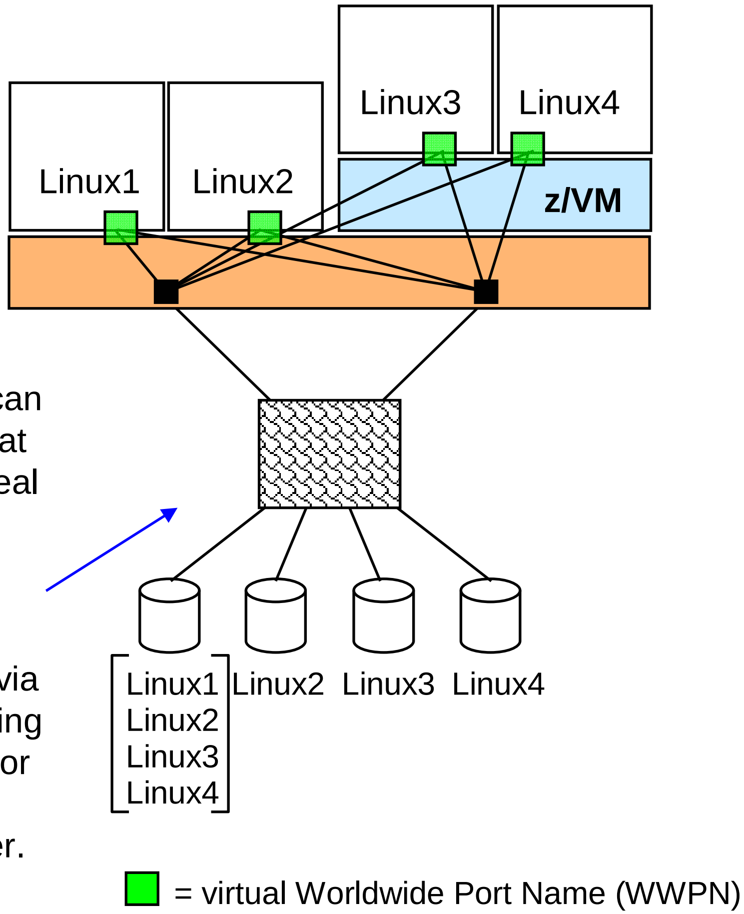

z/VM Support for N_Port ID Virtualization

- **FICON Express features on System z9 and z10 support FCP N_Port ID Virtualization (NPIV)**
- **NPIV enables zoning and LUN masking on a virtual machine basis**
- **Multiple operating system images can now concurrently access the same or different SAN-attached devices (LUNs) via a single, shared FCP channel**
  – Can increase channel utilization
  – Less hardware required
  – Helps reduce the complexity of physical I/O connectivity
- **Supported by z/VM V5.4, V5.3, and V5.2**

IBM System Storage SAN Volume Controller Software V4.3

- **z/VM and Linux for System z support SAN Volume Controller (SVC) V4.3**
- **SVC allows z/VM and Linux to access SCSI storage from multiple vendors as a single pool of disk capacity**
- **z/VM FBA emulation allows CMS users to access SVC-managed disk space**
- **New function in SVC V4.3:**
  - Space-Efficient Virtual Disks use disk space only when data is written
  - Space-Efficient FlashCopy uses disk space only for changes between source and target data
  - Virtual Disk Mirroring helps improve availability for critical applications by storing two copies of a virtual disk on different disk systems
- **Supported in z/VM V5.3, V5.4, and V6.1**
  - z/VM V5.2 support available with PTF for APAR VM64128

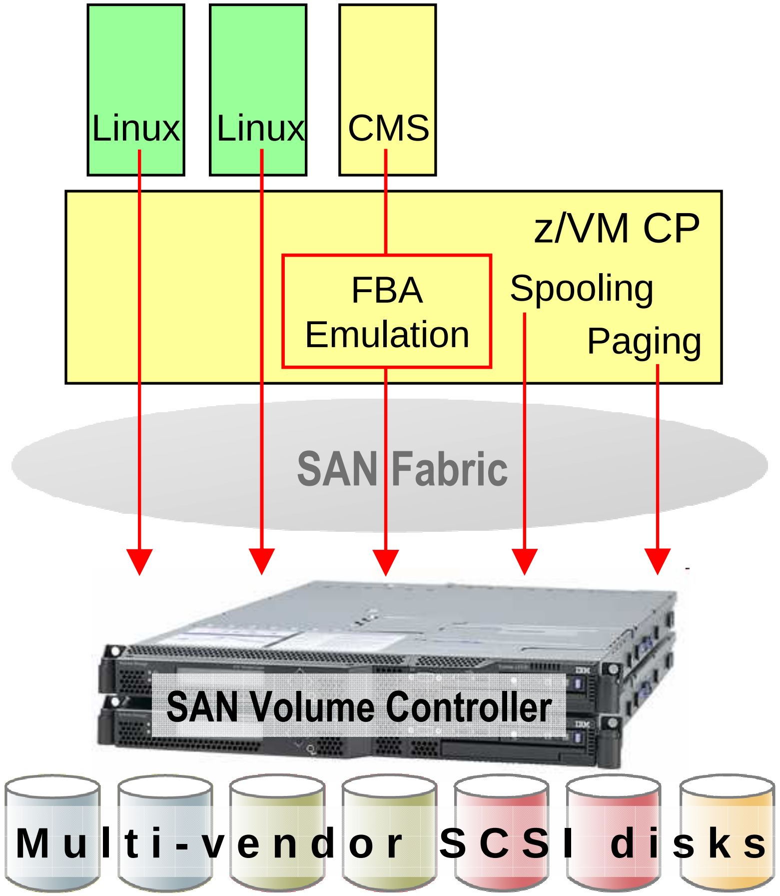

Learn more at: **ibm.com**/storage/support/2145

z/VM Virtual Networking
Using the z/VM Virtual Switch

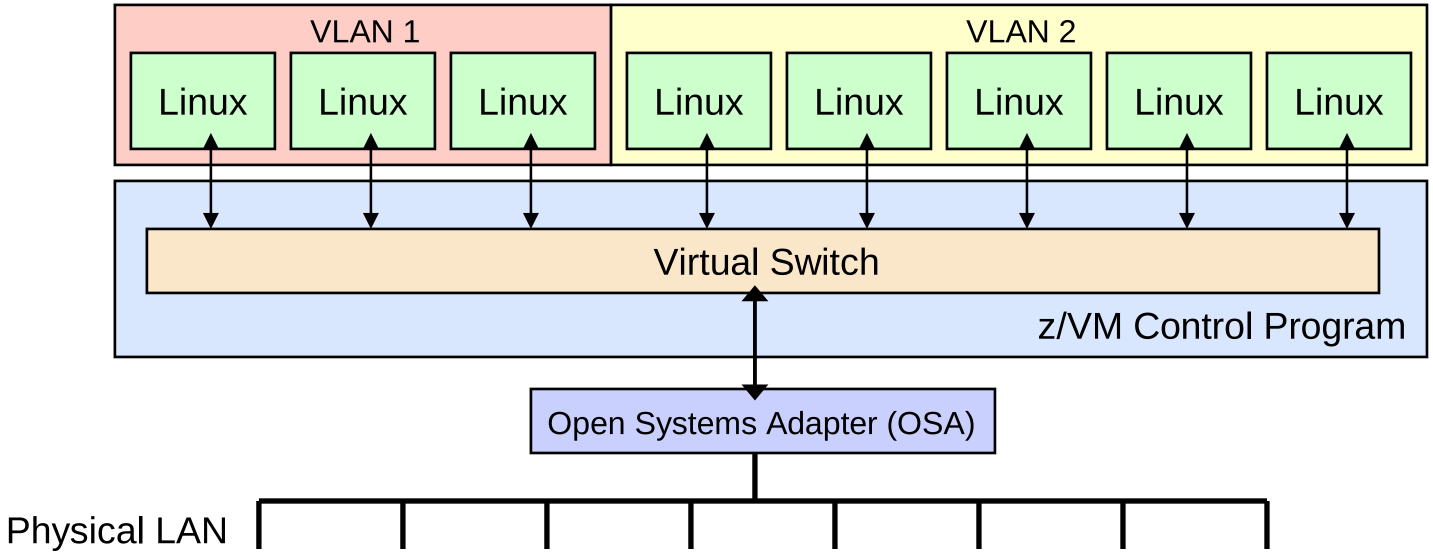

- **Eliminates need for router to connect virtual servers to physical LAN segments**
  - May reduce overhead associated with router virtual machines
  - Allows virtual machines to be in the same subnet with the physical LAN segment
- **Supports Layer 2 (MAC) and Layer 3 (IP) switching**
  - Includes support for IEEE VLAN
  - Provides centralized network configuration and control
  - Easily grant and revoke access to the real network
  - Dynamic changes to VLAN topology can be made transparent to virtual servers

z/VM Integrated Systems Management

Using the System z Hardware Management Console (HMC)

Included in z/VM V5.4

- Allows basic z/VM functions to be performed from HMC
- Network connection not required
- Uses SCLP hardware interface to access z/VM systems management APIs

Supported operations:

- View z/VM guests
- Activate z/VM guests
- Deactivate z/VM guests
- Display guest configuration and status

z/VM V5.3 also supported

- Requires PTFs for APARs VM64233 and VM64234

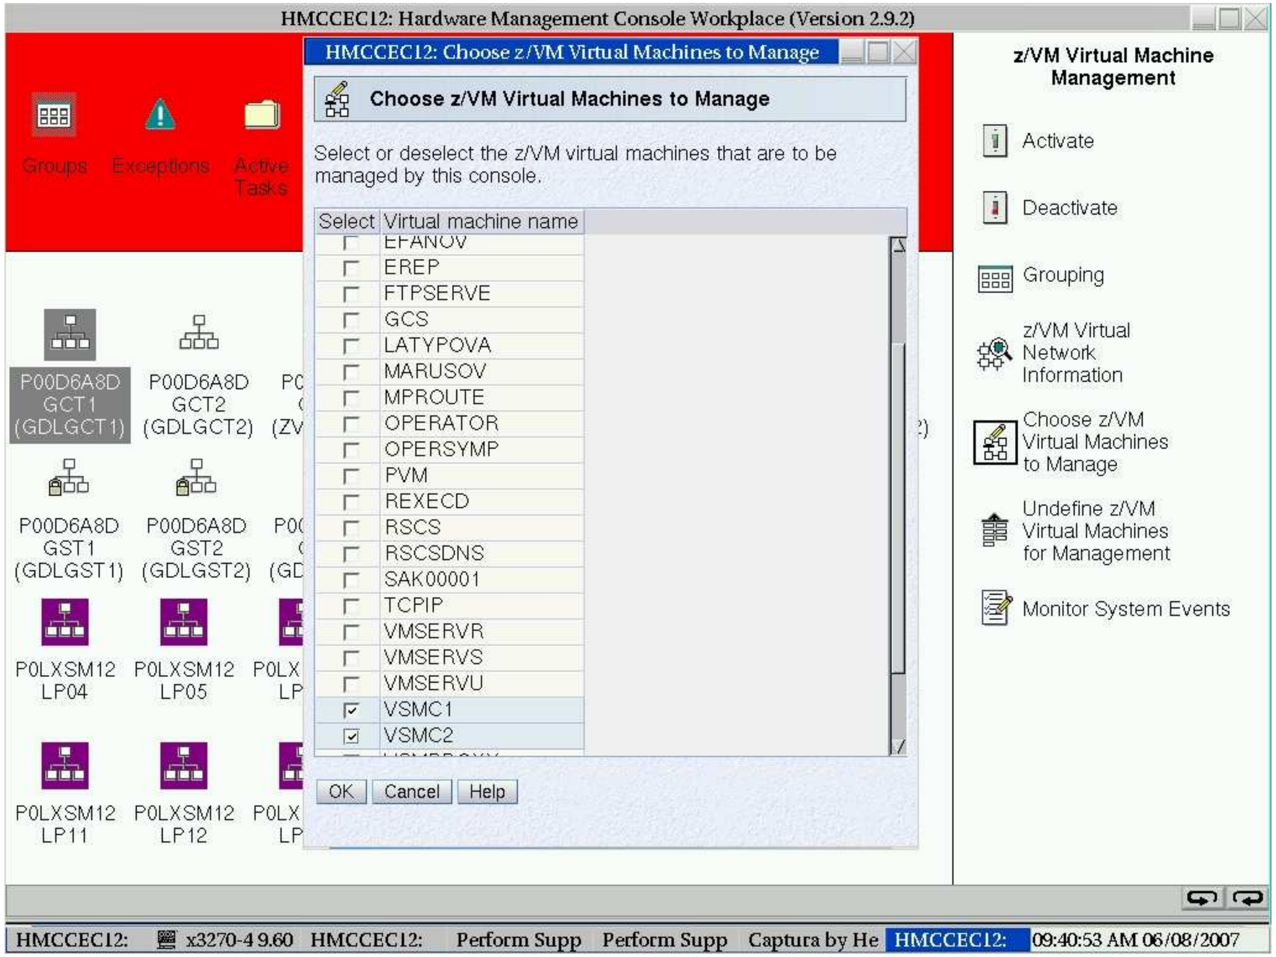

IBM Systems Director VMControl Image Manager
for Linux on System z Version 2.1 – Available July 24, 2009

- **VMControl Image Manager is a plug-in to IBM Systems Director V6.1**
  - Effectively replaces the “z/VM Center” extension of IBM Director V5.20
- **Provides support to manage and automate the deployment of virtual images from a centralized location**
  - A virtual image consists of an operating system instance and the software stack, such as middleware and applications, running on that operating system
- **VMControl Image Manager provides a graphical interface to create and deploy Linux images on z/VM and AIX images on Power systems**
  - Definition of these system images is based on the industry-standard Open Virtualization Format (OVF) specifications – facilitates importation of virtual images
  - Deploy an all-in-one solution instead of OS, middleware, and application piece parts
  - Clone already-tested system configurations
  - Propagate virtual image updates to all instances
- **IBM Systems Director and VMControl Image Manager help support a Dynamic Infrastructure**
  - Helps improve responsiveness to changing business needs
  - May increase operational productivity
  - Can help reduce service and support costs

Provisioning Software in System z Virtual Linux Servers

Using IBM Tivoli Provisioning Manager

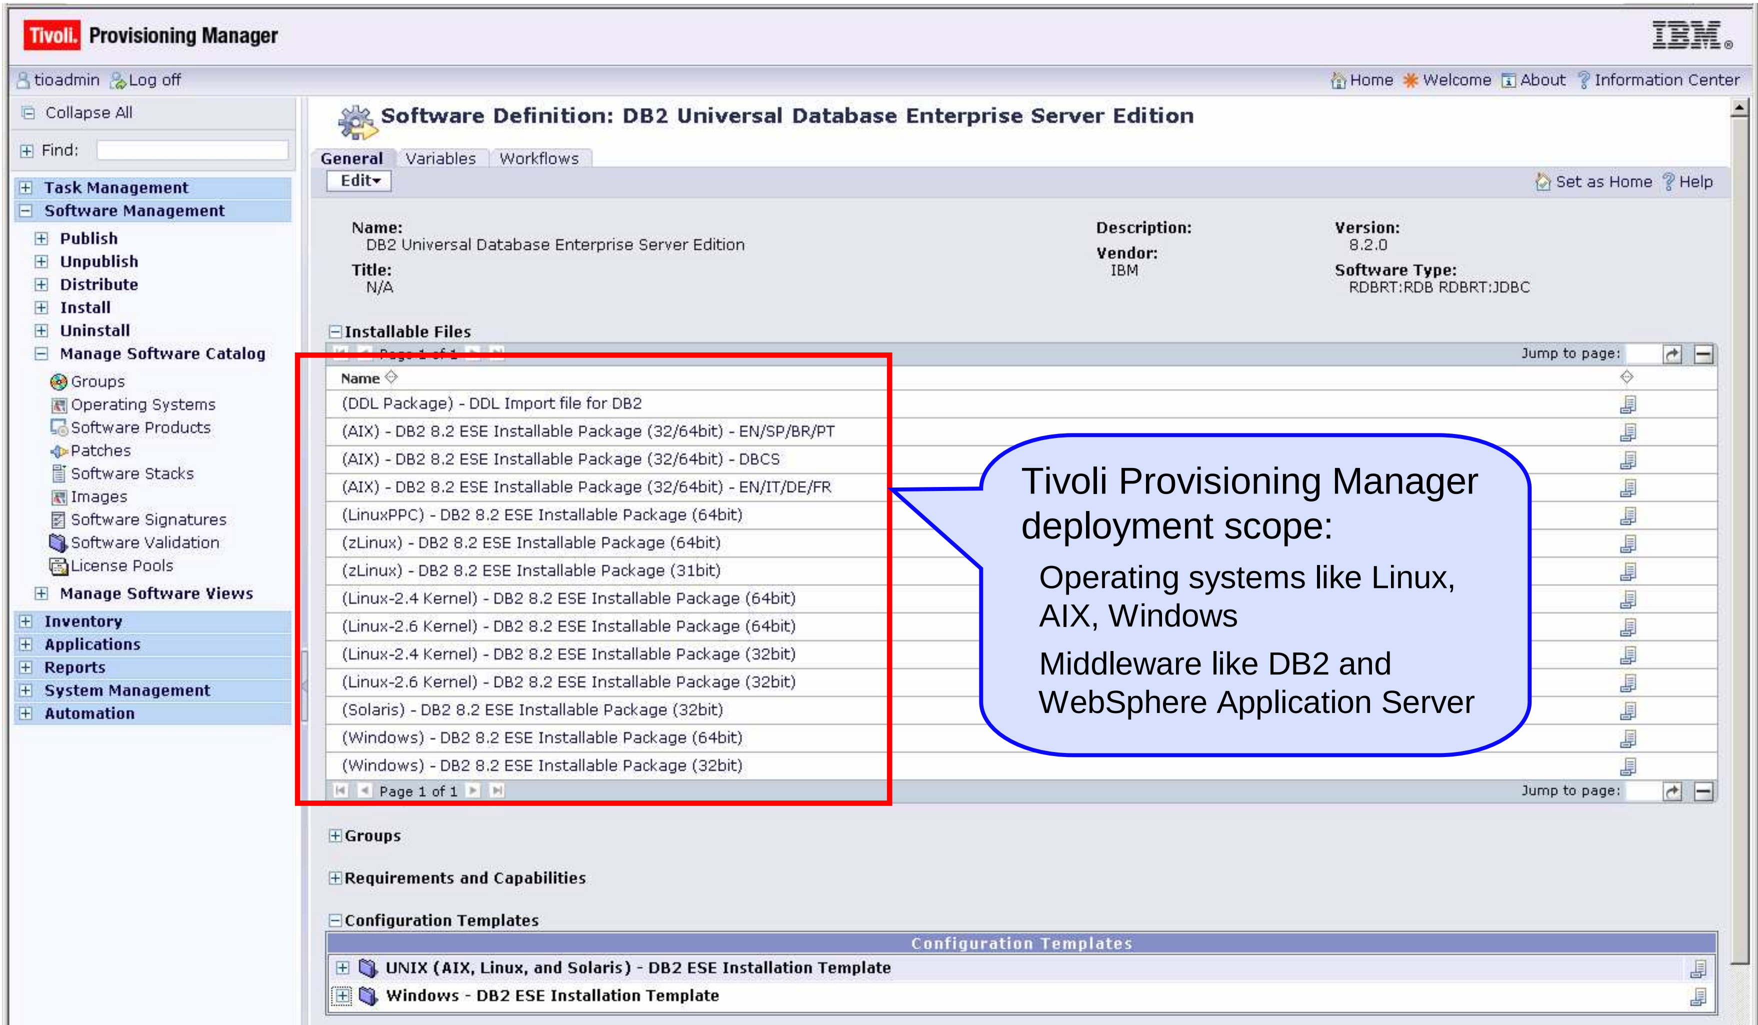

Monitoring System z Virtual Linux Servers

Using IBM Tivoli OMEGAMON XE on z/VM and Linux V4.1.2

- **Combined product offering that monitors z/VM and Linux for System z**
- **Provides work spaces that display:**
  - Overall system health
  - Workload metrics for logged-in users
  - Individual device metrics
  - LPAR data
- **Provides composite views of Linux running on z/VM**
- **New function in V4.1.2:**
  - Additional monitoring to help identify bottlenecks in the I/O subsystem
  - Processor spin lock wait statistics

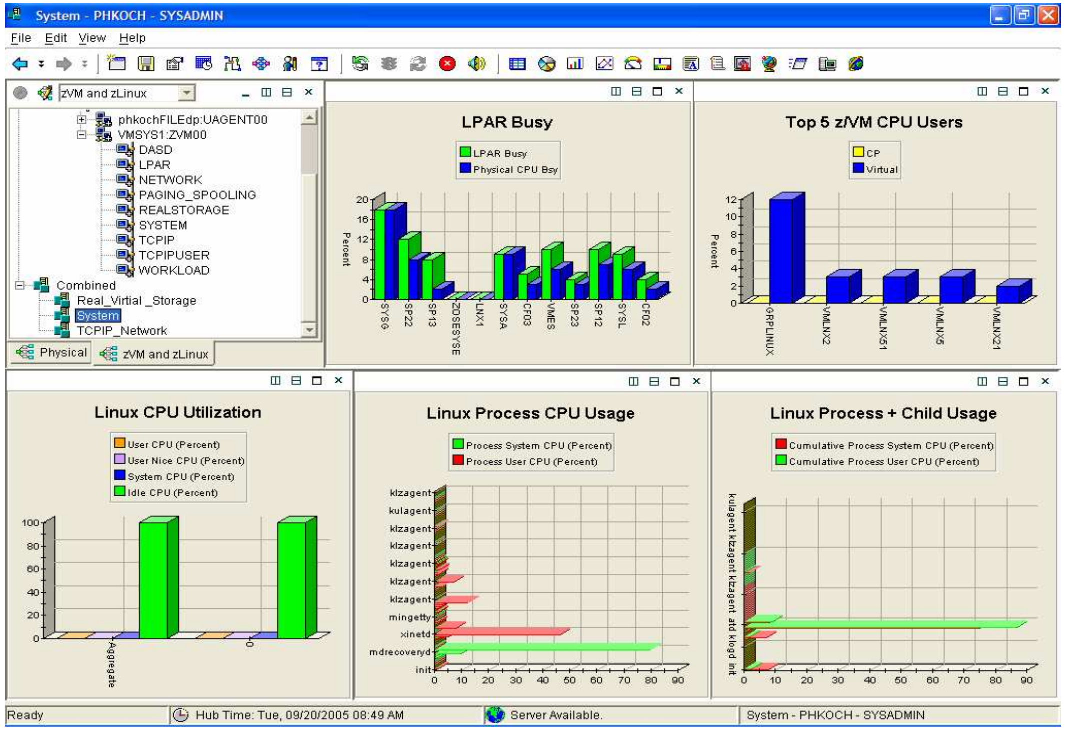

Learn more at: **ibm.com**/software/tivoli/products/omegamon-xe-zvm-linux

IBM Tivoli Virtualization Management for System z

Helping Clients Manage and Control Their Virtualized IT Infrastructure

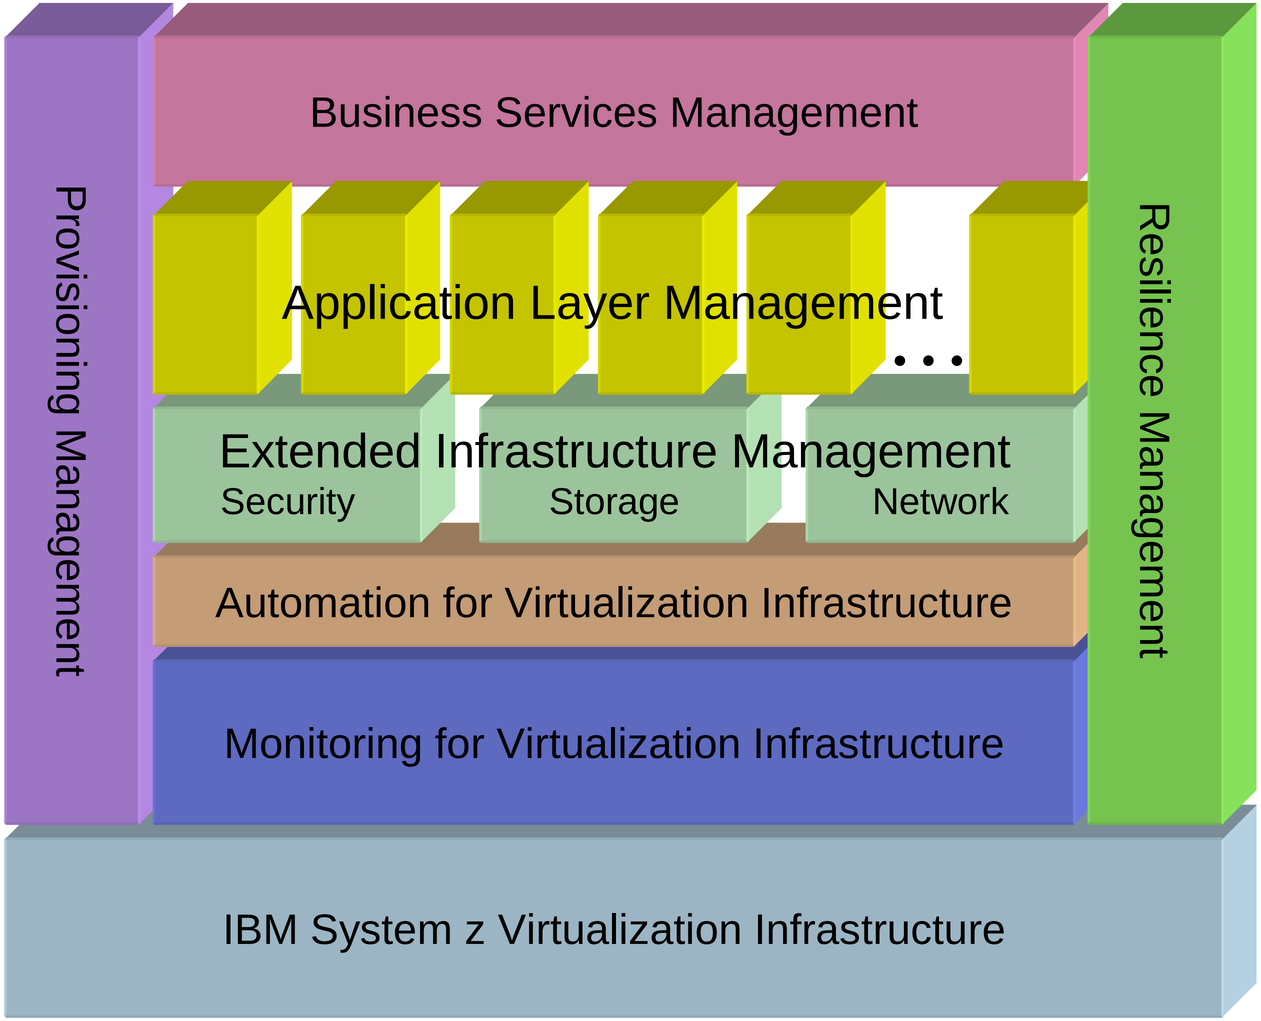

z/VM Statements of Direction
Clustered Hypervisor Support and Guest Mobility
Overview of Planned New Function

- **Clients can cluster up to four z/VM systems in a** Single System Image **(SSI)**
- **Provides a set of shared resources that can be used by both z/VM and hosted virtual machines, with full awareness of sharing by the clustered z/VM systems – be they on the same and/or different z10 servers**
  - Directory, minidisks, spool files, Virtual Switch MAC addresses
- **Helps simplify systems management for a multi-z/VM environment**
  - Single user directory
  - Cluster management from any system
    - Apply maintenance to all systems in the cluster from one location
    - Issue commands from one system to operate on another
  - Built-in cross-system capabilities
  - Service consolidation: run one copy of service virtual machines for the cluster
  - Resource coordination and protection: network and disks

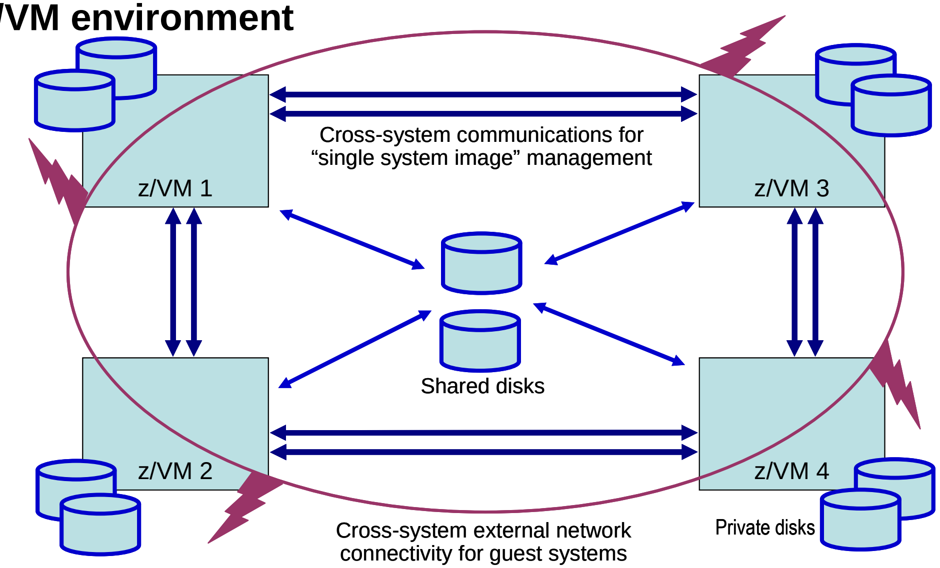

- **Dynamically move Linux guests from one z/VM system to another in the cluster via** Live Guest Relocation
  - Helps reduce planned outages; enhances workload management
  - With z/VM: dynamically move work to available resources **and** dynamically move resources to work

Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Topics

- **Business drivers for server consolidation**
- **Server virtualization and hypervisor basics**
- **x86 virtualization**
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Some customer examples*
  - *Summary*

System z LPAR and z/VM: World-Class Server Virtualization

Helping clients reduce costs and improve control of their IT infrastructure

✓ ***Virtualization*** ✓ ***Consolidation*** ✓ ***Workload management*** ✓ ***Automation***

- ***Logical Partitioning (LPAR) and z/VM are complementary technologies***
  - Both employ great hardware and firmware (PR/SM) innovations developed over the years
  - Virtualization is a part of the *basic componentry* of the System z platform

- ***LPAR***
  - Host a relatively small number of very high-performance virtual servers
  - Very low overhead, hardware-based virtualization through partitioning

- ***z/VM***
  - Host large numbers of high-performance virtual servers
  - Low overhead, hardware-based, true virtualization with extreme levels of software augmentation

Together, System z LPAR and z/VM technology provide:

- High performance “on the metal” virtual servers for larger, performance-critical workloads
- The ability to provision 1000s of additional virtual servers flexibly and on demand

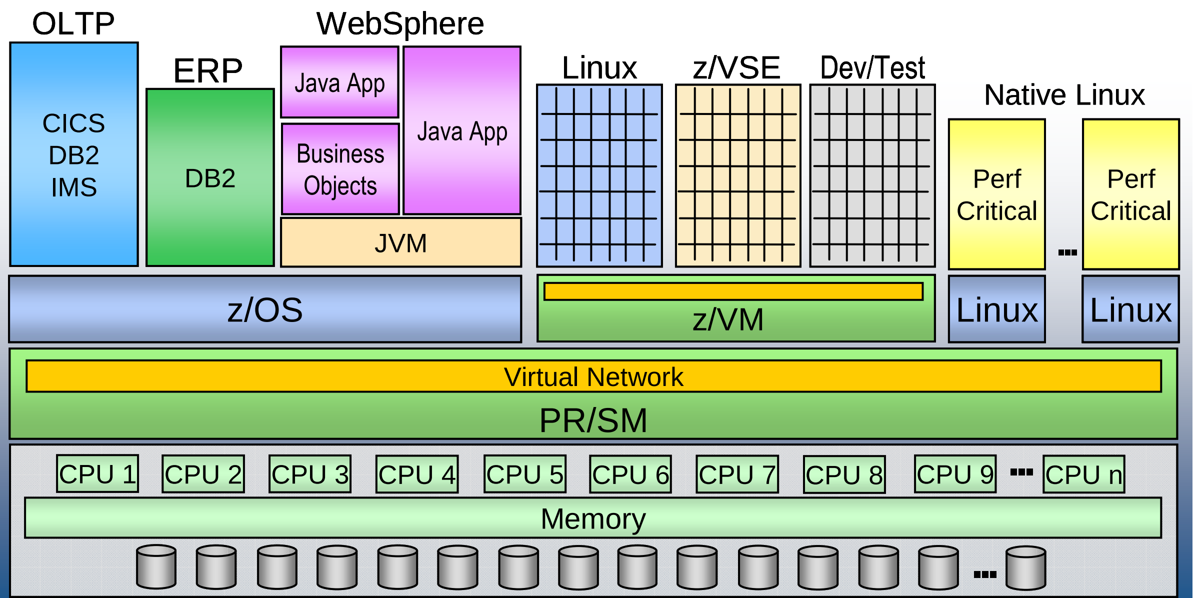

System z Consolidation and Virtualization Success Stories

Maximizing your IT Competitiveness

1	First National Bank Omaha	*Consolidated more than 500 distributed Sun and x86 systems to Linux on z/VM.*	• 7x increase in system utilization • Improved system reliability • 67% reduction in staffing
2	BALDOR MOTORS • DRIVES • GENERATORS	*Consolidated SAP infrastructure to IBM System z: DB2 database runs on z/OS and SAP applications run in z/VM virtual machines.*	• Enhanced SAP availability and performance • 60% performance improvement for batch work • Reduced IT spending from 1.8 to 1 percent of sales • Easier to manage and maintain systems
3	Nationwide® On Your Side℠	*Made strategic decision to deploy J2EE applications on Linux and z/VM to avoid building a new data center.*	• Streamlined application development & operations • On schedule to save more than $16M over 3 years • 50% reduction in web infrastructure costs • 80% reduction in floor space
4	Hannaford	*Consolidated 300 Linux store servers onto a single mainframe. Introduced new web portal for business partners.*	• Running 62 virtual servers instead of 300 servers • Significant labor savings across IT organization • Improved customer and partner satisfaction

IBM System z Achieves High Core-to-Core Ratios When Consolidating from Distributed Environments

- A fundamental strength of z/VM is its ability to overcommit system resources: “Do more with less” leads to cost savings for hardware and software
- Demonstrates unique IBM value proposition against competitive scale-out solutions (e.g., Sun, HP, Others)
- **Real customers, real workloads!**

Customer	Distributed Cores	System z Cores	Ratio of Distributed to System z cores*
Nationwide	450	21	21 to 1
Large Bank	111	4	27 to 1
Government Agency	292	5	58 to 1

* Client results will vary based on each specific customer environment including types of workloads, utilization levels, target consolidation hardware, and other implementation requirements.

Actual Customer Saves $1.5M with Oracle on System z

Versus 45 Oracle x86 Servers

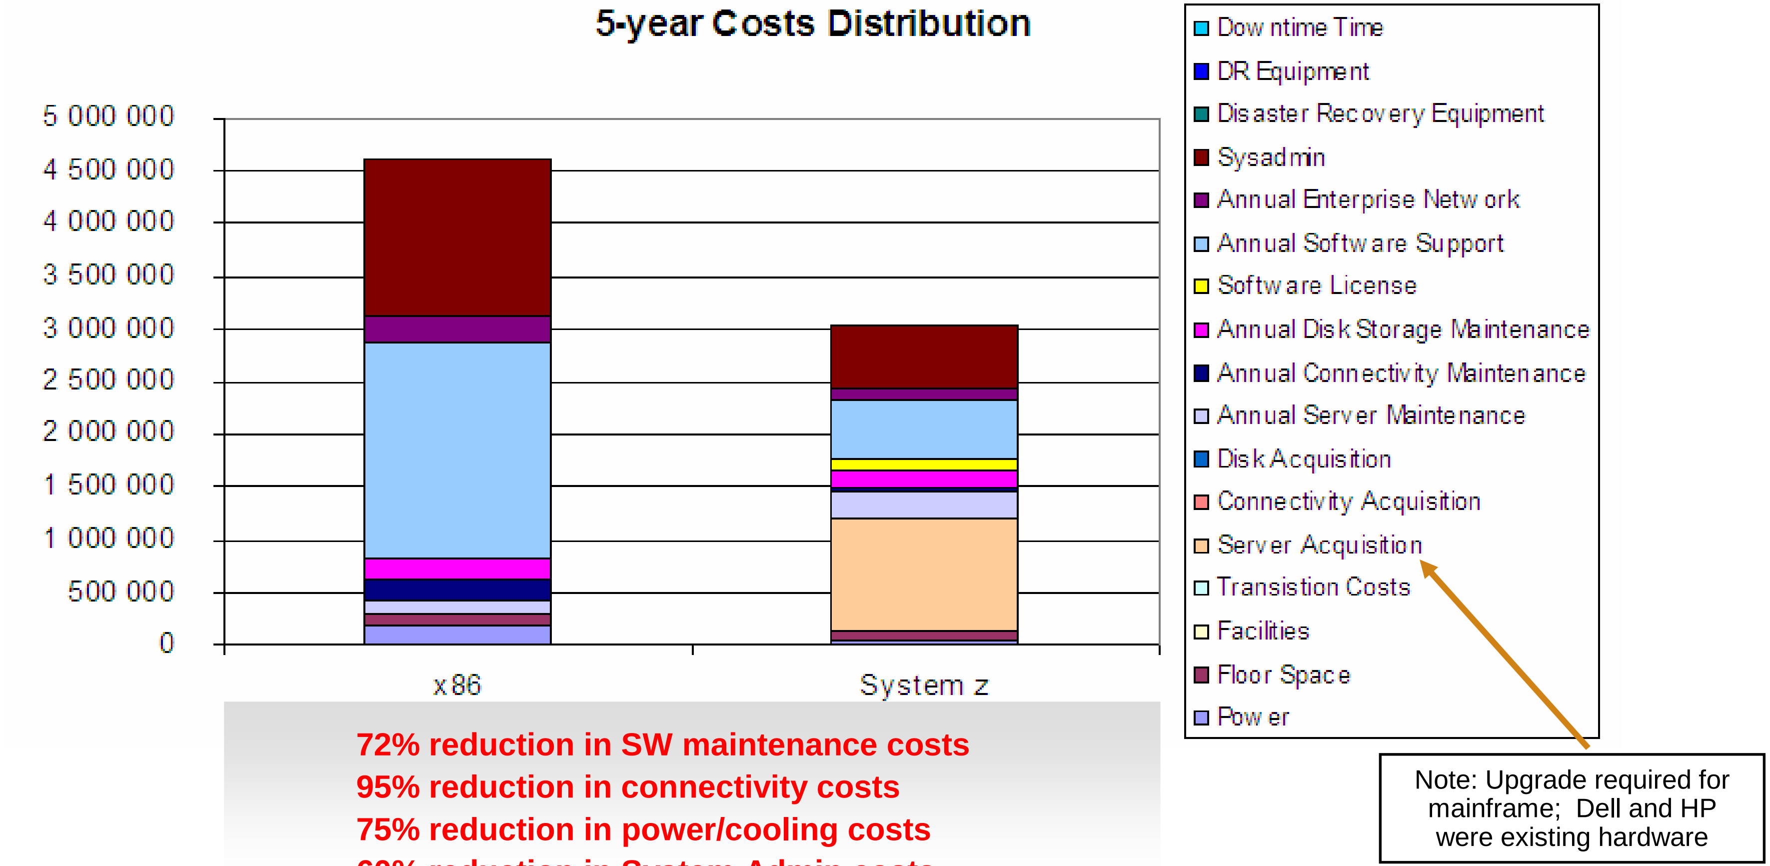

72% reduction in SW maintenance costs
95% reduction in connectivity costs
75% reduction in power/cooling costs
60% reduction in System Admin costs

5-year Cost Comparison	1st Year	2nd Year	3rd Year	4th Year	5th Year
x86	923 625	1 847 250	2 770 874	3 694 499	4 618 124
System z	1 482 559	1 871 822	2 261 085	2 650 348	3 039 611
Delta	558 934	24 572	-509 789	-1 044 151	-1 578 513

Details on next slide

Large Bank Reduces Space and Energy Requirements

Saves More Than $1.5M (detailed information for previous slide)

	FROM…	TO…
Current Hardware Infrastructure	45 x86 (HP + Dell)	IBM System z10 Enterprise Class (z10 EC)
Footprints	45	1
Cores	111	4 IFLs
Average utilization	Less than 10%	60%
Peak utilization	35%	85%
# of DBs; size of DB	111 Oracle DB	111 Oracle DB
Application	Oracle 10G databases	Oracle 10G databases
OS	Linux	Linux Red Hat RHEL5 + z/VM
Energy usage		75% less
Floor Space usage		28% less
TCO: 5 years	$4.62M	$3.04M / Savings: $1.58M

Summary of Benefits:

- 111-to-4 core reduction; 45:1 footprint reduction
- Up to 72% software maintenance cost reduction
- Improved application reliability, and efficient Disaster Recovery capabilities

Topics

- **Business drivers for server consolidation**
- **Server virtualization and hypervisor basics**
- **x86 virtualization**
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Some customer examples*
  - *Summary*

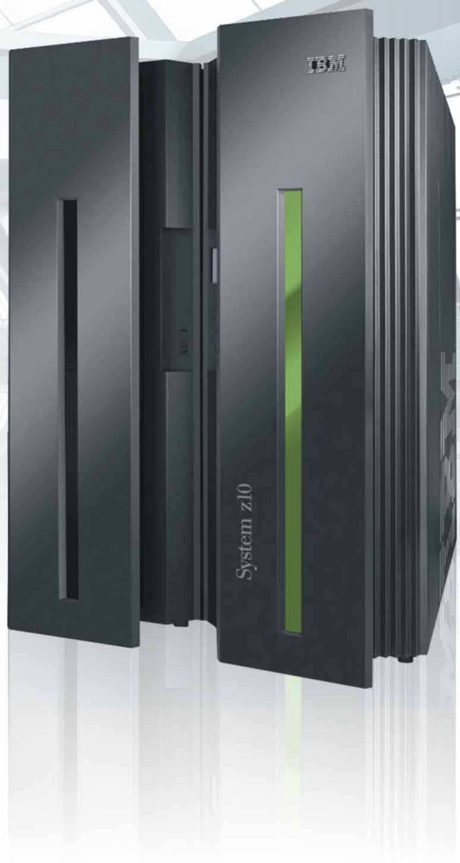

System z and z/VM Virtualization Leadership
Virtualization Adoption Issues for x86 Systems

- **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**
- **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 in 3 years?**
- **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?**
- **Will Microsoft restrict its software from supporting your virtualization solution?**
  - See: www.vmware.com/solutions/whitepapers/msoft_licensing_wp.html as an example

When Do You Need More Than "Good Enough"?

Business Drivers – Making the Case for Mainframe Virtualization

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

Questions

The future runs on System z

Appendix

Moving Linux Images

- **VMware VMotion (and similar x86 technologies)**
  - Moving a running Linux virtual server is how workload management is accomplished with virtualized x86 systems
  - Can also be used to move virtual machines off a “hot” server
  - Helps customers manage planned downtime
- **System z and z/VM**
  - Virtual server workload management is very robust with System z and z/VM
    - Very granularity allocation of real resources helps absorb workload spikes
    - Capacity on demand can quickly satisfy the need for more hardware
  - System z is a very “green” machine (power consumption is not a problem)
  - Live guest migration indeed helps avoid planned outages
    - Running multiple copies of z/VM (even on the same CEC) can help mitigate this issue
    - Tivoli Systems Automation and clustered Linux solutions also offer improved availability
    - Even so, live guest migration is a z/VM investment candidate, focusing on a) reducing planned outages, b) helping customers manage virtual servers deployed across multiple z/VM systems

What is z/VM System Security?

- **Knowing who is accessing the system or its resources**
- **Ensuring a user only has access to system resources specifically permitted**
- **Knowing who is accessing (or failing to access) what resources**
- **Security is only meaningful in the presence of system integrity**
  - Integrity prevents bypass of security controls
  - Audit trail confirms conformance

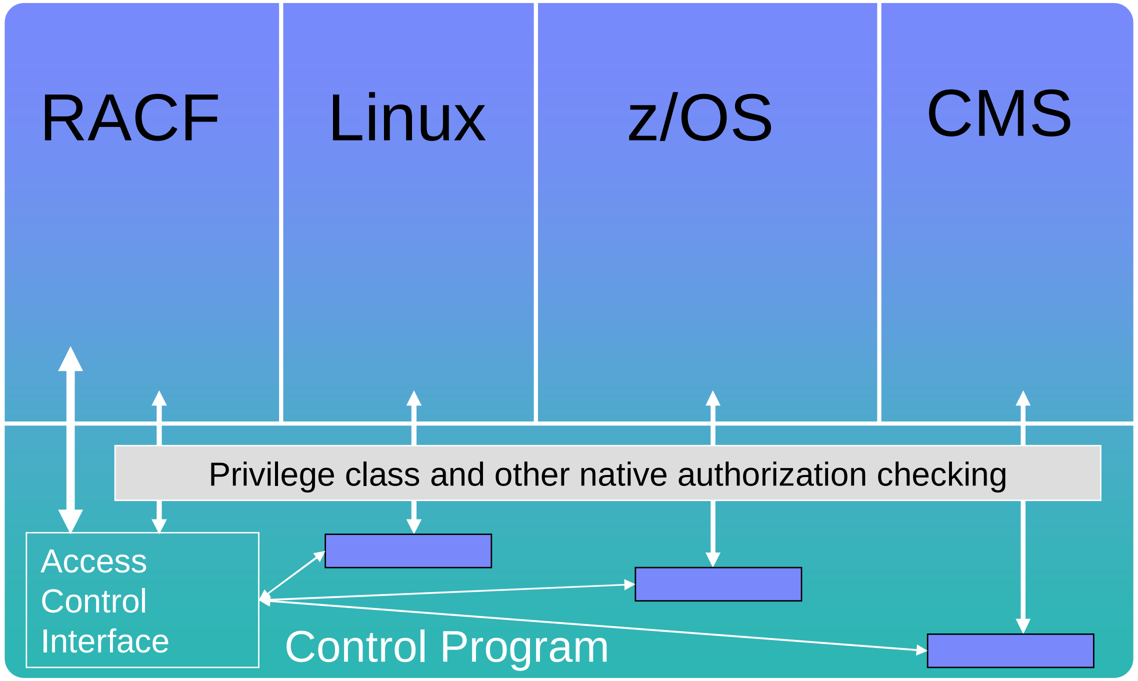

Learn more: "z/VM Security and Integrity" – ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130145.pdf

IBM System z Ecosystem – the Growing z Community

Dramatic Growth Responding to Market Demand

Expanding System z application portfolio driven by ISV growth:

- 2008: 130 new ISVs; 600+ new applications/tools
- More than 5,000 applications available
- More than 2,450 applications for Linux on System z
- Over 1,400 ISVs building applications

IBM Academic Initiative is driving skills growth on System z

- 50,000 students attended mainframe education
- 481 schools registered offering 29 separate courses, with more to come
- Student Mainframe Contests:
  - 10 contents
  - More than 1,000 schools
  - More than 8,000 students