The Business Value of System z Virtualization Leadership

The future runs on System z

October 2008

Reed A. Mullen
mullenra@us.ibm.com
IBM Systems and Technology Group

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 server virtualization overview**
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Why do businesses need what System z offers?*

Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration

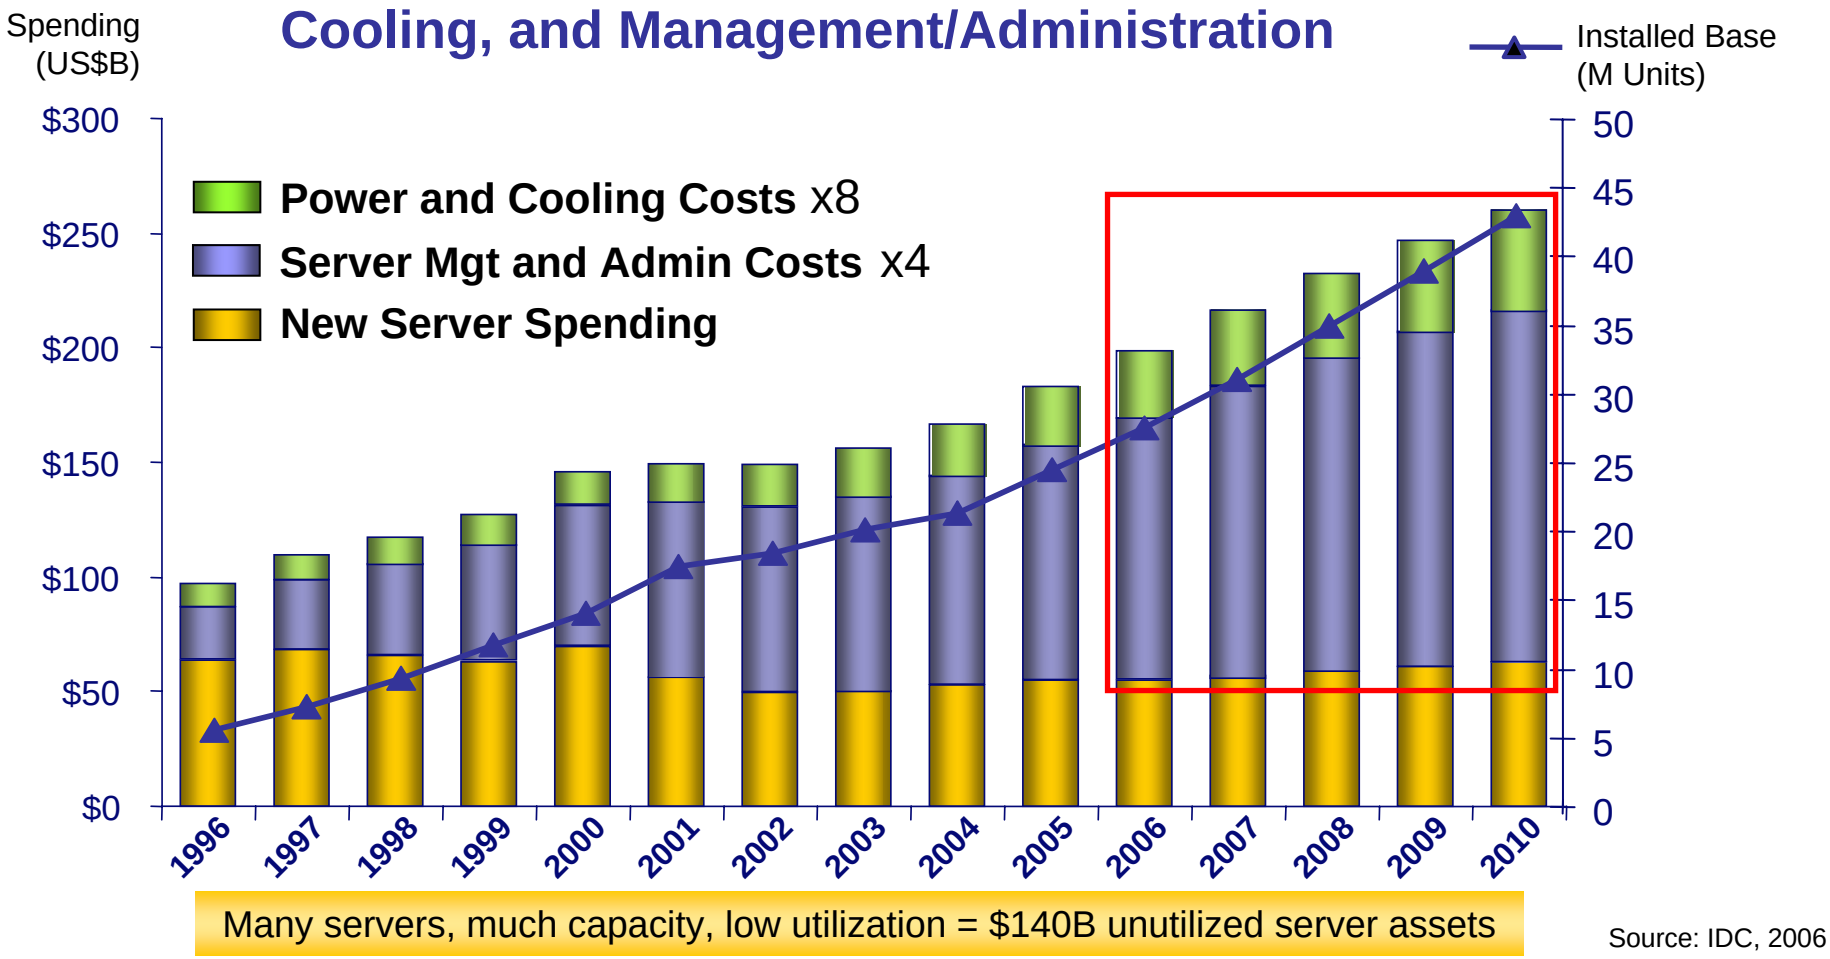

Many servers, much capacity, low utilization = $140B unutilized server assets

Source: IDC, 2006

As server volumes increase, so does complexity, making basic business requirements like availability, security, and disaster recovery more difficult to achieve.

x86 Virtualization Pain Points

System z Virtualization Marketplace Opportunities

- **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

Topics

- Business drivers for server consolidation
- **Server virtualization and hypervisor basics**
- x86 server virtualization overview
- *System z virtualization differentiation*
  - *Extreme virtualization with z/VM and System z*
  - *Why do businesses need what System z offers?*

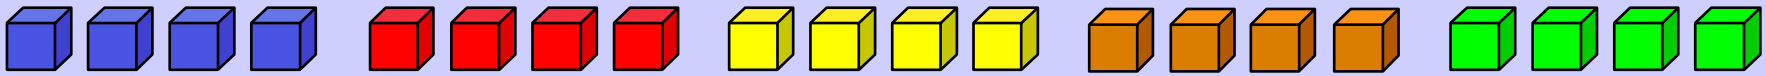

Virtual Resources

- 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.

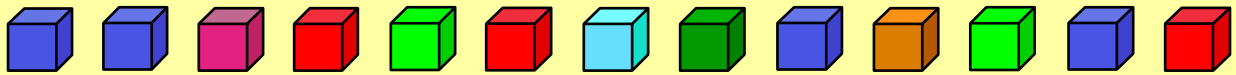

- **Separates presentation of resources to users from actual resources**
- **Aggregates pools of resources for allocation to users as virtual resources**

Server Virtualization Approaches

Hardware Partitioning

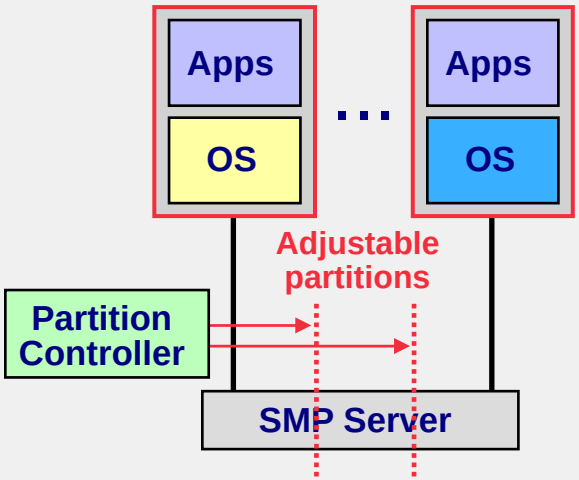

Server is subdivided into fractions each of which can run an OS

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

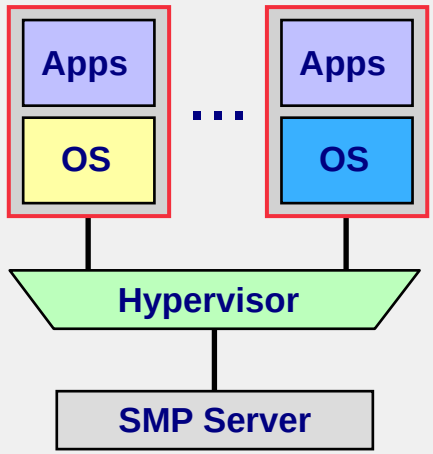

Hypervisor provides fine-grained timesharing of all resources

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

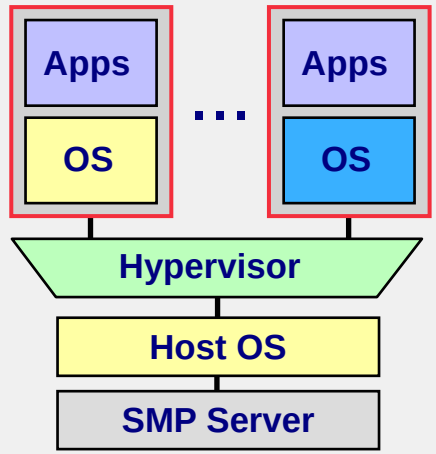

Hypervisor uses OS services to do timesharing of all resources

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

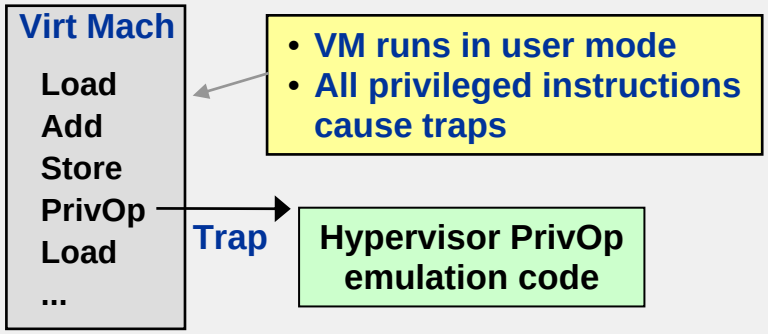

Examples: CP-67, VM/370
Benefits: Runs unmodified OS
Issues: Substantial overhead

Translate, Trap, and Emulate

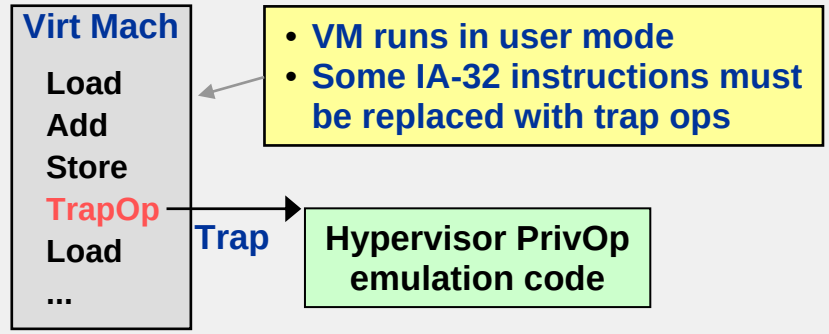

Examples: VMware, Microsoft VS
Benefits: Runs unmodified, translated OS
Issues: Substantial overhead

Hypervisor Calls ("Paravirtualization")

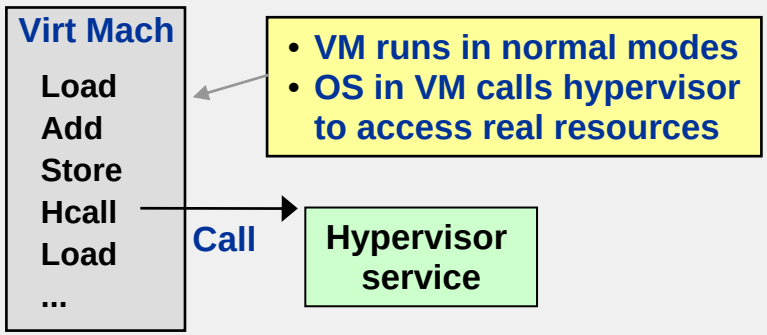

Examples: POWER Hypervisor, Xen
Benefits: High efficiency
Issues: OS must be modified to issue Hcalls

Direct Hardware Virtualization

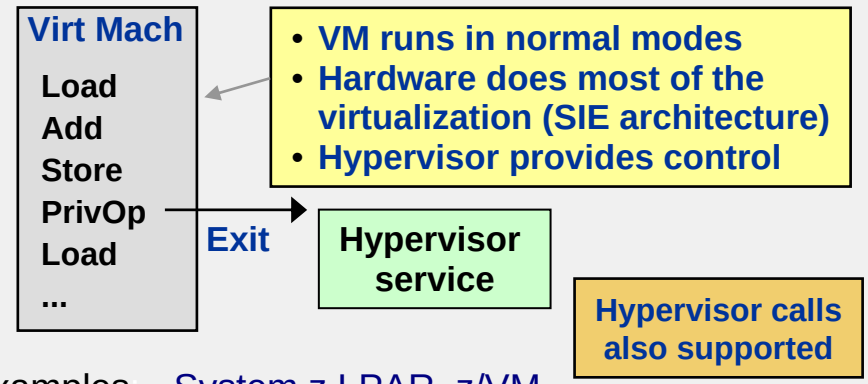

Examples: System z LPAR, z/VM
Benefits: High efficiency, runs unmodified OS
Issues: Requires underlying hardware support

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*
  - Sweet spot: engineering/scientific computing
- **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 by the underlying system architecture.

Virtualization and Security
Should IT Managers Be Concerned?

Virtualization security risks being overlooked, Gartner warns

Gartner raises warning on virtualization and security.

Companies in a rush to deploy virtualization technologies for server consolidation efforts could wind up overlooking many security issues and exposing themselves to risks, warns research firm Gartner.

“Virtualization, as with any emerging technology, will be the target of new security threats,” said Neil MacDonald, a vice president at Gartner, in a published statement.

– NetworkWorld.com, April 6, 2007

Posted September 21, 2007 at http://blogs.iss.net/archive/virtblog.html

"It is clear that with the increase in popularity, relevance and deployment of virtualization starting in 2006, vulnerability discovery energies have increasingly focused on finding ways to exploit virtualization technologies."

"...in a virtual environment all your exploitation risks are now consolidated into one physical target where exploiting one system could potentially allow access and control of multiple systems on that server (or the server itself). In total, this adds up to a more complex and risky security environment."

Known vulnerabilities across all of VMware's products

VMware Vulns by Year	Total Vulns	High Risk Vulns	Remote Vulns	Vulns in 1st Party Code	Vulns in 3rd Party Code
Vulns in 2003	9	5	5	5	4
Vulns in 2004	4	2	0	2	2
Vulns in 2005	10	5	5	4	6
Vulns in 2006	38	13	27	10	28
Vulns in 2007	34	18	19	22	12

F-Secure Vulnerability Information : VMware ESX Server Sun Java JDK / JRE Multiple Vulnerabilities

[Summary] | [Detailed Description] | [Solution] | [CVE Reference]

Security Guide
F-Secure World Map
Security Alerts
Virus Statistics
Malware Removal Tools
Malware Code Glossary
Submit Malware Sample
Select local site
Global Sites

VIRUS WORLD MAP

Global Alert Level:
- Medium -

Global Status:

> Scan My Computer Now
> Download Trial Versions

Report ID:	SA32180
Source:	Secunia
Date of Discovery:	06.10.2008
Criticality:	Urgent
Affects:	VMware ESX Server 3.x
Compromise From:	From remote
Compromise Type:	Exposure of system information Security bypass Exposure of sensitive information DoS System access

Summary

VMware has acknowledged some vulnerabilities in VMware ESX Server, which can be exploited by malicious people to bypass certain security restrictions, disclose system information or potentially sensitive information, cause a DoS (Denial of Service), or compromise a vulnerable system.

Back to the Top

Detailed Description

VMware has acknowledged some vulnerabilities in VMware ESX Server, which can be exploited by malicious people to bypass certain security restrictions, disclose system information or potentially sensitive information, cause a DoS (Denial of Service), or compromise a vulnerable system.

The vulnerabilities affect versions 3.0.1, 3.0.2, 3.0.3, and 3.5.

Back to the Top

Solution

Patches are reportedly pending release.

Do not follow untrusted links or browse untrusted websites.

Back to the Top

What is z/VM System Integrity?

- **The ability of the z/VM Control Program to operate without interference or harm, intentional or not, from guest virtual machines**
- **The inability of a virtual machine to circumvent system security features and access controls**
- **The ability of the Control Program to protect virtual machines from each other**
- **All done with the help of System z hardware and firmware**

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

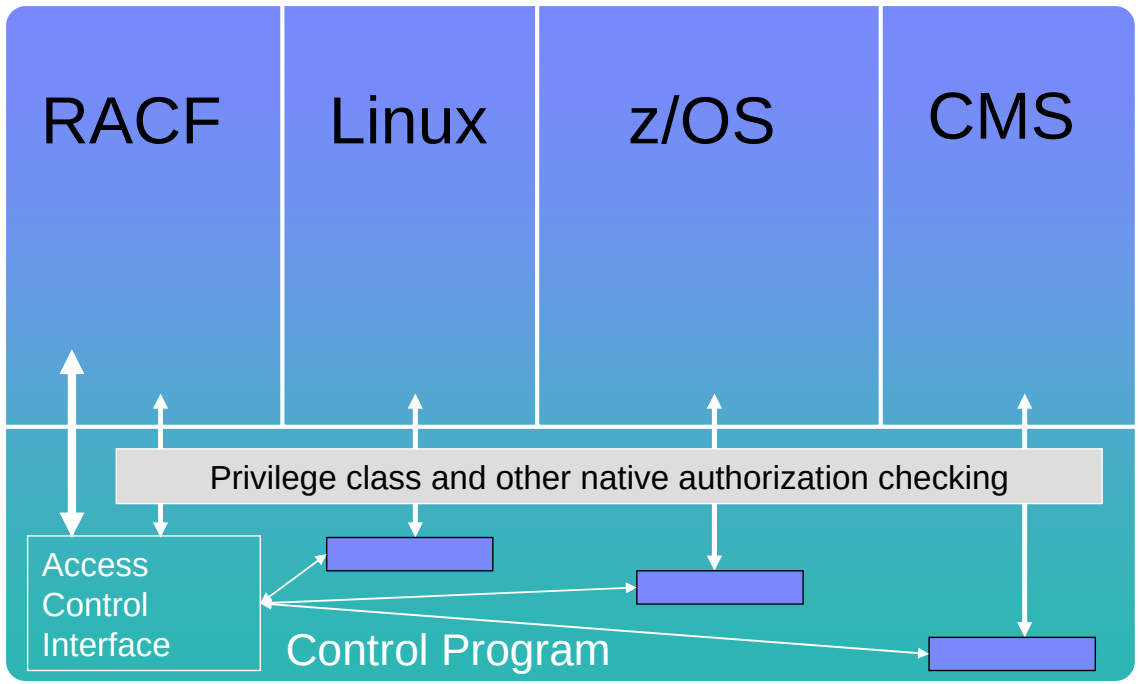

Learn more: “z/VM Security and Integrity” – ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130145.pdf

Topics

- Business drivers for server consolidation
- Server virtualization and hypervisor basics
- **x86 server virtualization overview**
- *System z virtualization differentiation*
  - *Extreme virtualization with z/VM and System z*
  - *Why do businesses need what System z offers?*

x86 Server Virtualization Landscape
Hypervisors

- **VMware ESX Server**
  - Part of VMware Infrastructure 3 offering
- **Microsoft**
  - Virtual Server 2005 and Windows Server 2008 Hyper-V
- **Xen (open source code for Linux and Solaris)**
  - Distributed by Novell (SLES 10) and Red Hat (RHEL 5)
  - Marketed by Virtual Iron Software, Citrix, Oracle (Oracle VM), Sun (xVM)
- **KVM (open source Linux)**
  - Integrated into the Linux 2.6.20 kernel
  - New strategic initiative for Red Hat (“Embedded Linux Hypervisor”)
- **Many more virtualization alternatives**
  - Solaris Containers, Virtuozzo, HP Integrity VM, ...

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

x86 Server Virtualization Landscape
Virtualization Management Offerings

- **VMware VirtualCenter**
- **Virtual Iron Virtualization Manager**
- **XenSource XenEnterprise**
- **Microsoft System Center Virtual Machine Manager**
- **OracleVM – web browser based management console**
- **Sun xVM Ops Center**
- **HP OpenView**
- **IBM Systems Director**
  - Including *IBM Director* and *IBM Tivoli Provisioning Manager*

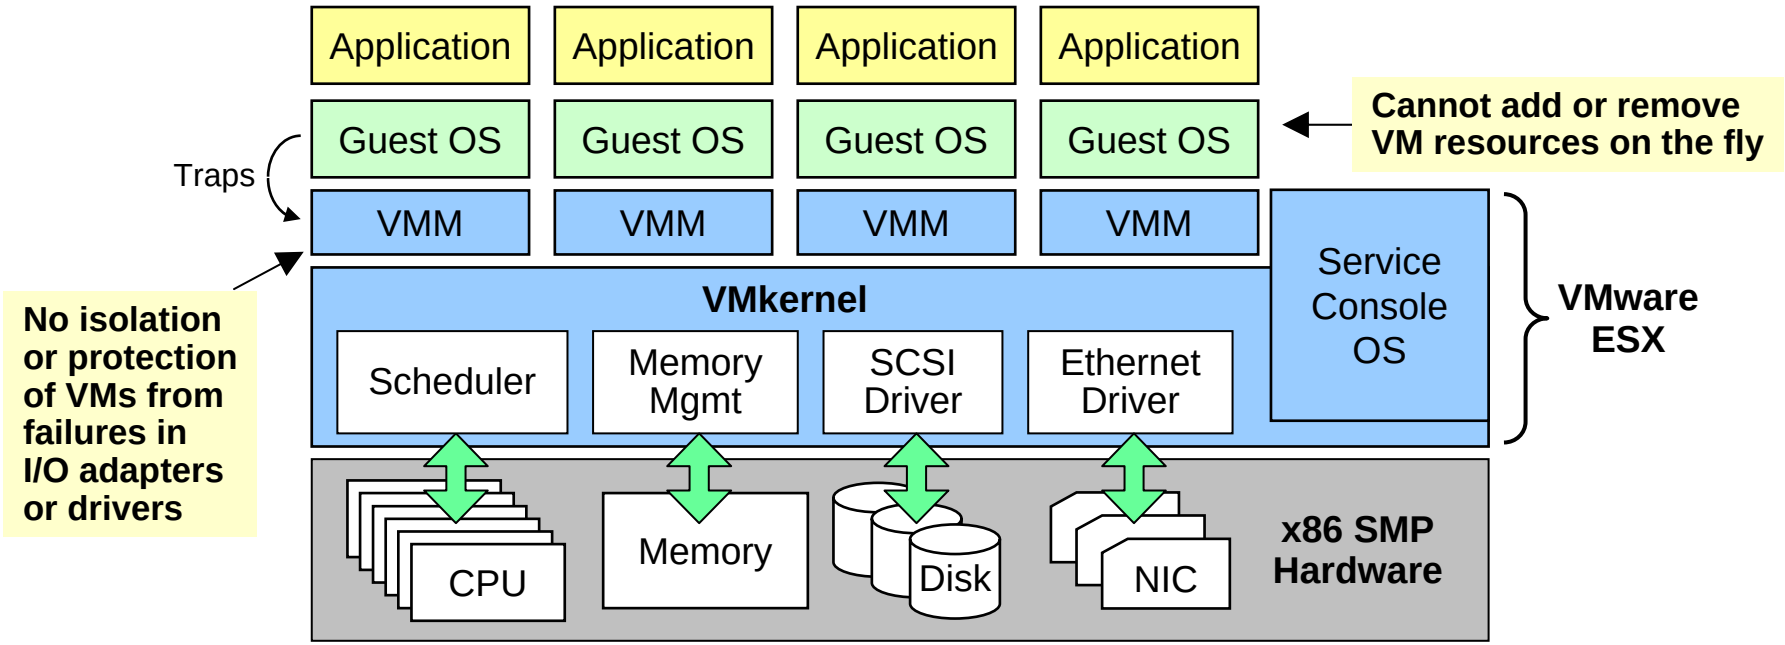

- **Supports a wide range of unmodified Windows and Linux versions**
- **Guest OS runs in User Mode; privileged instructions trap to Virtual Machine Monitor (VMM)**
  - "Trapping and mapping" is a significant source of performance overhead
- **Guest OS binary code is translated incrementally at load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

x86 Partitioning with Xen

- 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 enables hosting of unmodified guest operating systems
  - For Windows support on Xen, users need IVT-capable hardware

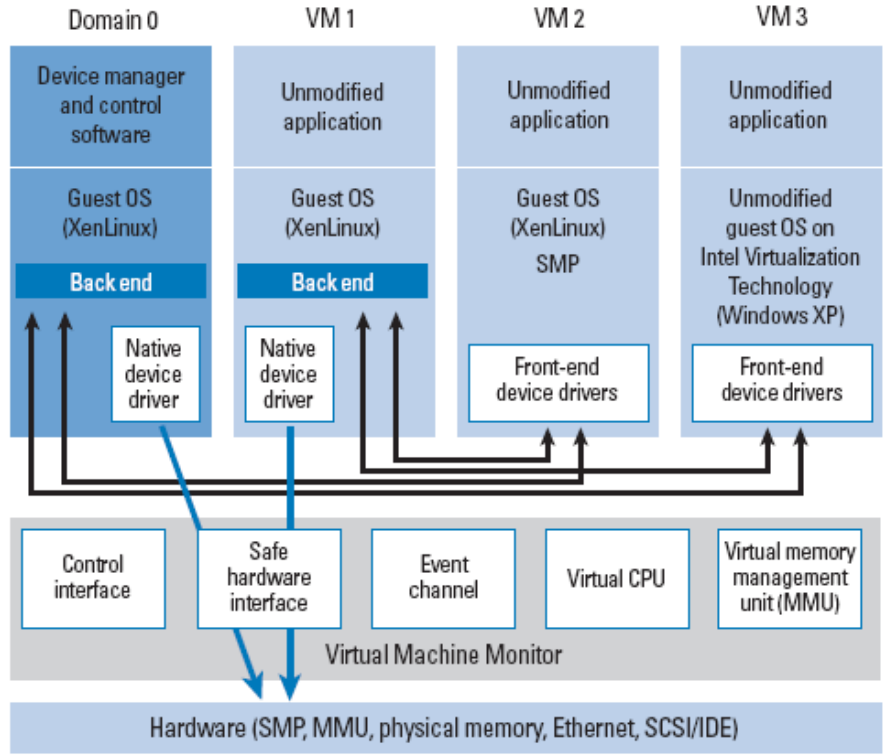

Figure 1. Xen 3.0 architecture: Hosting four VMs

x86 Partitioning with KVM

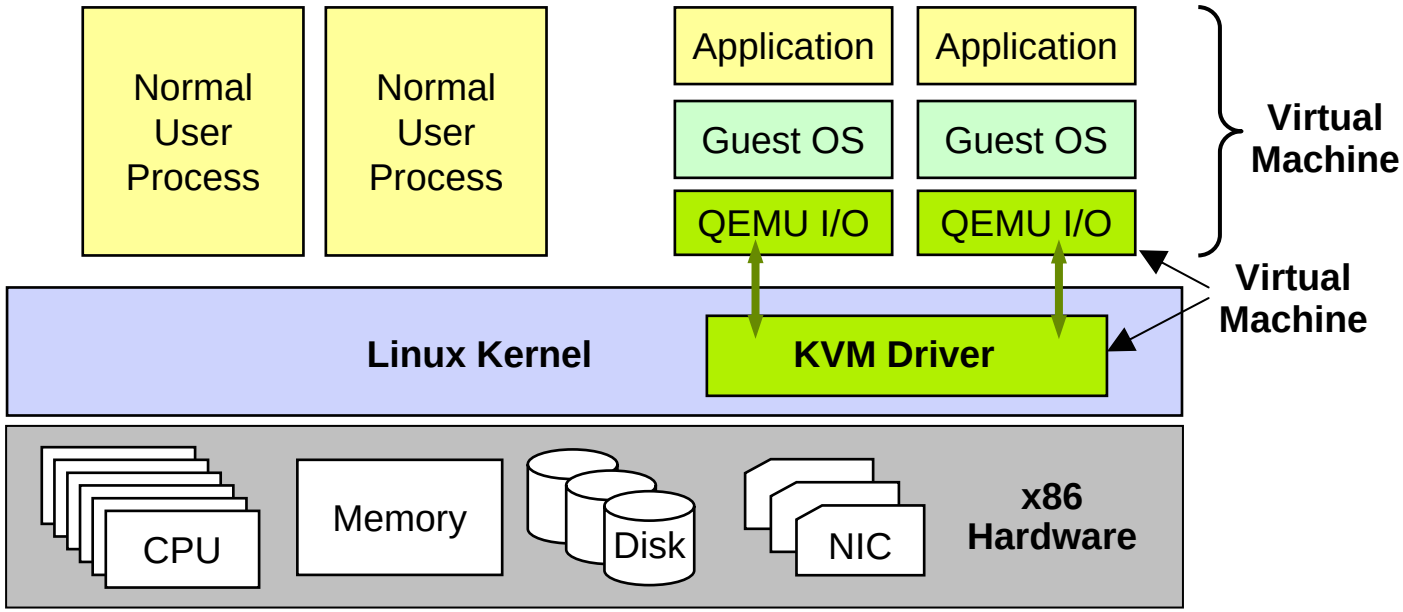

- **Normal Linux processes have two modes of execution: “kernel” and “user”**
  - KVM adds a third mode: “guest”
- **KVM virtualizes CPU, I/O Advanced Programmable Interrupt Controller (IOAPIC), and Memory Management Unit (MMU); requires IVT or AMD-V**
- **QEMU is a user space component that emulates PC hardware**
  - KVM gives QEMU near-native CPU virtualization
- **Each virtual machine is a normal Linux process**

Microsoft Virtual Server 2005

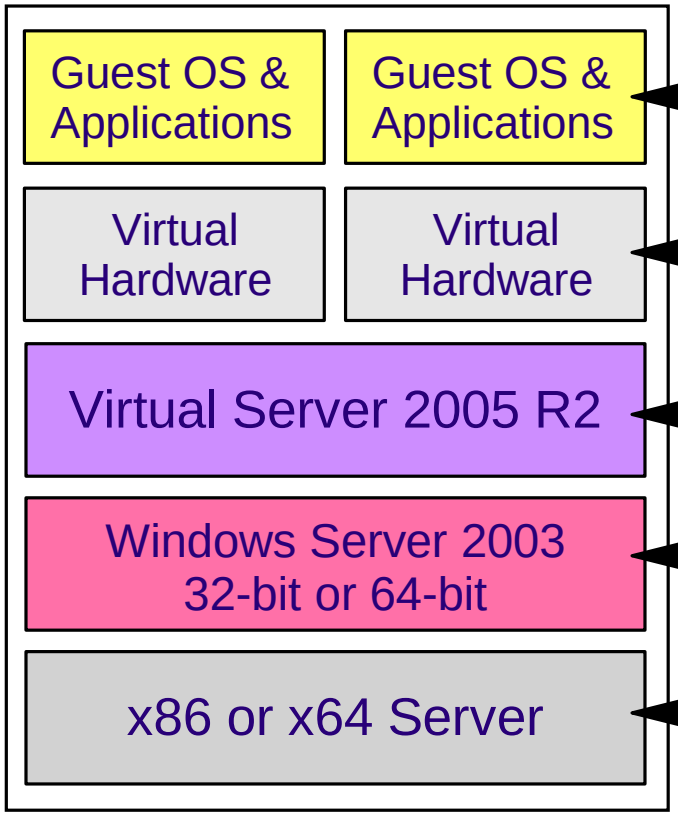

Windows NT 4.0 Server, Windows 2000 Server, Windows 2003 Server supported (32-bit)

1 CPU per virtual machine, 32-bit only
Up to 3.6 GB of memory per virtual machine

Multithreaded virtual machine monitor (VMM) provides isolation

Windows Server 2003
Broad device compatibility

Optimized for 2 to 8 way servers
Scales up to 32 CPUs and up to 256 GB

Source: “Microsoft Virtual Server 2005 R2 Technical Overview”, December 2005

Microsoft Hyper-V

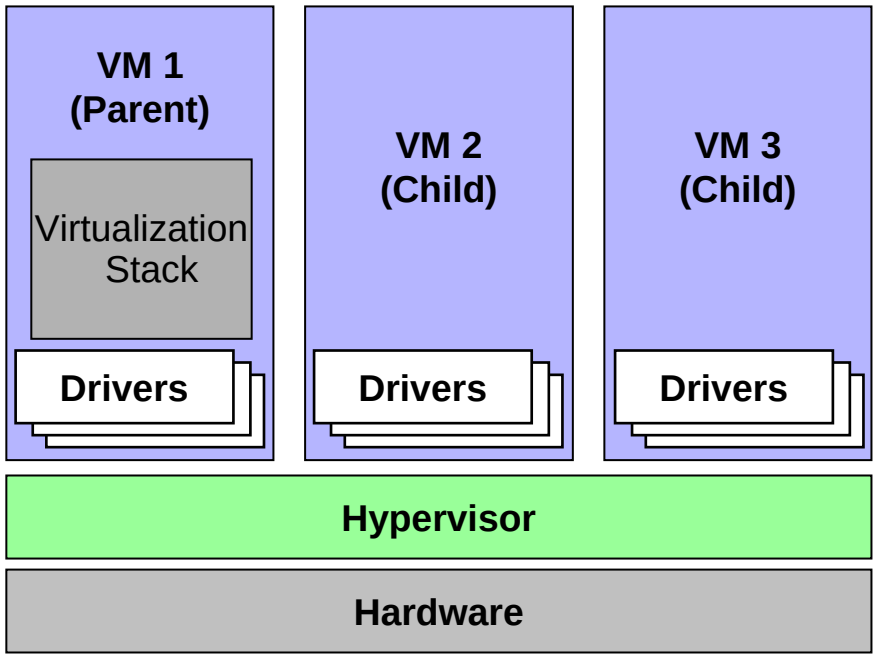

- **Hyper-V has been generally available since mid-summer 2008 as a feature of Windows Server 2008**
  - It is a thin layer of code that runs in hypervisor privileged mode using Intel VT or AMD-V hardware
  - Windows boots first, and then inserts the hypervisor beneath itself, and then becomes a parent partition
- **Para-virtualization ("hypercalls") are used for high efficiency**
- **Normal device drivers will be run in partitions**
  - Leverage the large base of Windows drivers
  - The parent Windows partition will host virtual I/O, much like the PowerVM VIOS approach
  - Unlike VMware ESX, where modified Linux drivers are run in the VMkernel
- **No guest mobility support in the first release**

Topics

- Business drivers for server consolidation
- Server virtualization and hypervisor basics
- x86 server virtualization overview
- ***System z virtualization differentiation***
  - *Extreme virtualization with z/VM and System z*
  - *Why do businesses need what System z offers?*

IBM System z Virtualization Genetics

- **System z is thoroughly architected to host applications in a virtualized environment**
- **This is accomplished with a coordinated set of investments that permeate the technology stack of hardware, firmware, hypervisors, and operating systems**
- **This means clients can maximize the utilization, scalability, and security of all system assets, including:**
  - *CPU*
  - *Memory*
  - *I/O*
  - *Networking*
  - *Cryptography*
- **All with exceptional levels of operational ease and cost efficiencies**

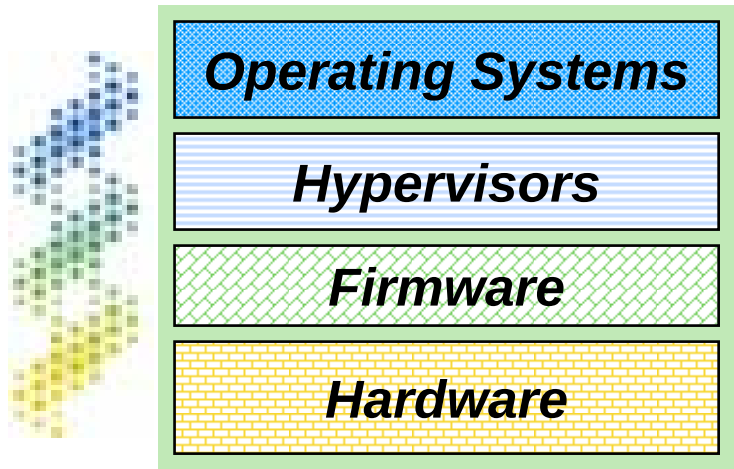

IBM System z Virtualization Genetics

Over 40 years of continuous innovation in virtualization

- Refined to support modern business requirements
- Exploit hardware technology for economical growth
- LPAR, Integrated Facility for Linux, HiperSockets
- System z Application Assist Processors
- System z Information Integration Processors

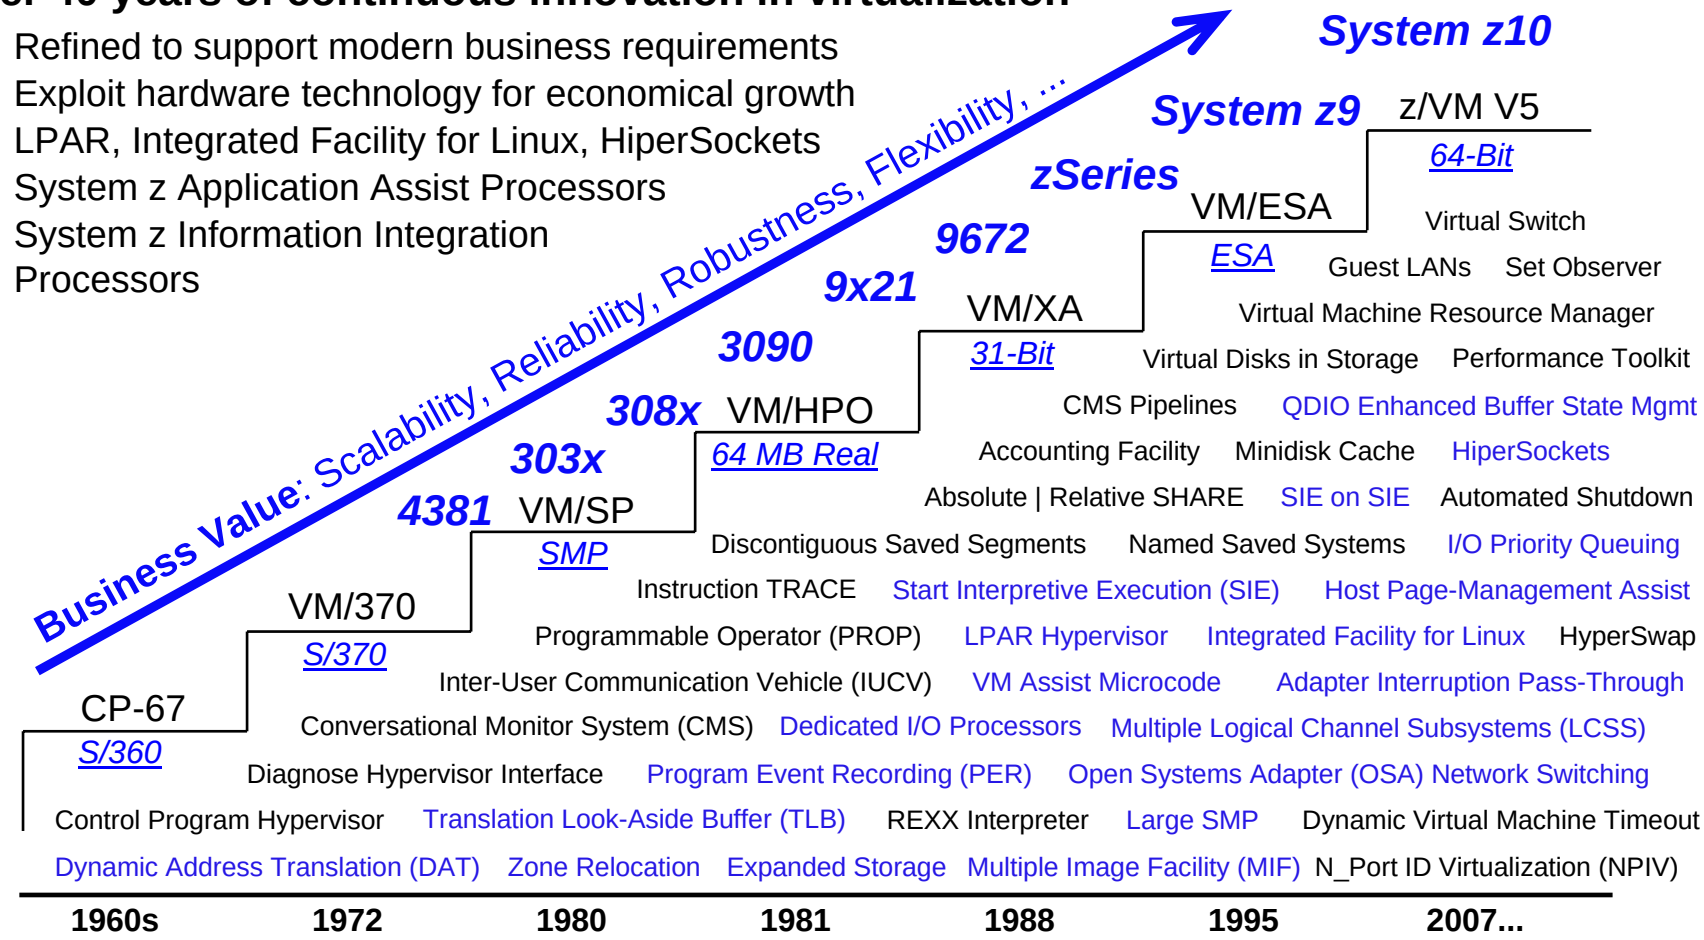

IBM System z – a comprehensive and sophisticated suite of virtualization function

IBM System z: The Ultimate Virtualization Platform

- ***Virtualize*** **everything with very high levels of utilization**
  - CPU, memory, network, I/O, cryptographic features, coupling facility, ...

Consolidate all types of workloads

- ***Massively scale*** **your workload on a single System z mainframe**
  - Host tens-to-hundreds of virtual machines on z/VM
  - Each virtual machine on z/VM can access up to 24,576 devices

Smart economics: start small and grow big in the same box

- ***Non-disruptively add*** **anything**
  - Up to 64x CPU scalability per mainframe, 32x scalability per z/VM LPAR
  - z/VM is designed to support more than 1 TB of active virtual memory

Able to respond to workload spikes

- ***Security*** **for everything**
  - Highest security classification for general purpose servers
  - System z LPAR technology is EAL 5 certified

Helps secure your virtual servers and reduce business risk

- ***Optimize and integrate*** **it all with the IBM software portfolio**

Increase staff productivity and virtualize the enterprise

IBM System z Virtualization Leadership

Extreme Levels of CPU Sharing

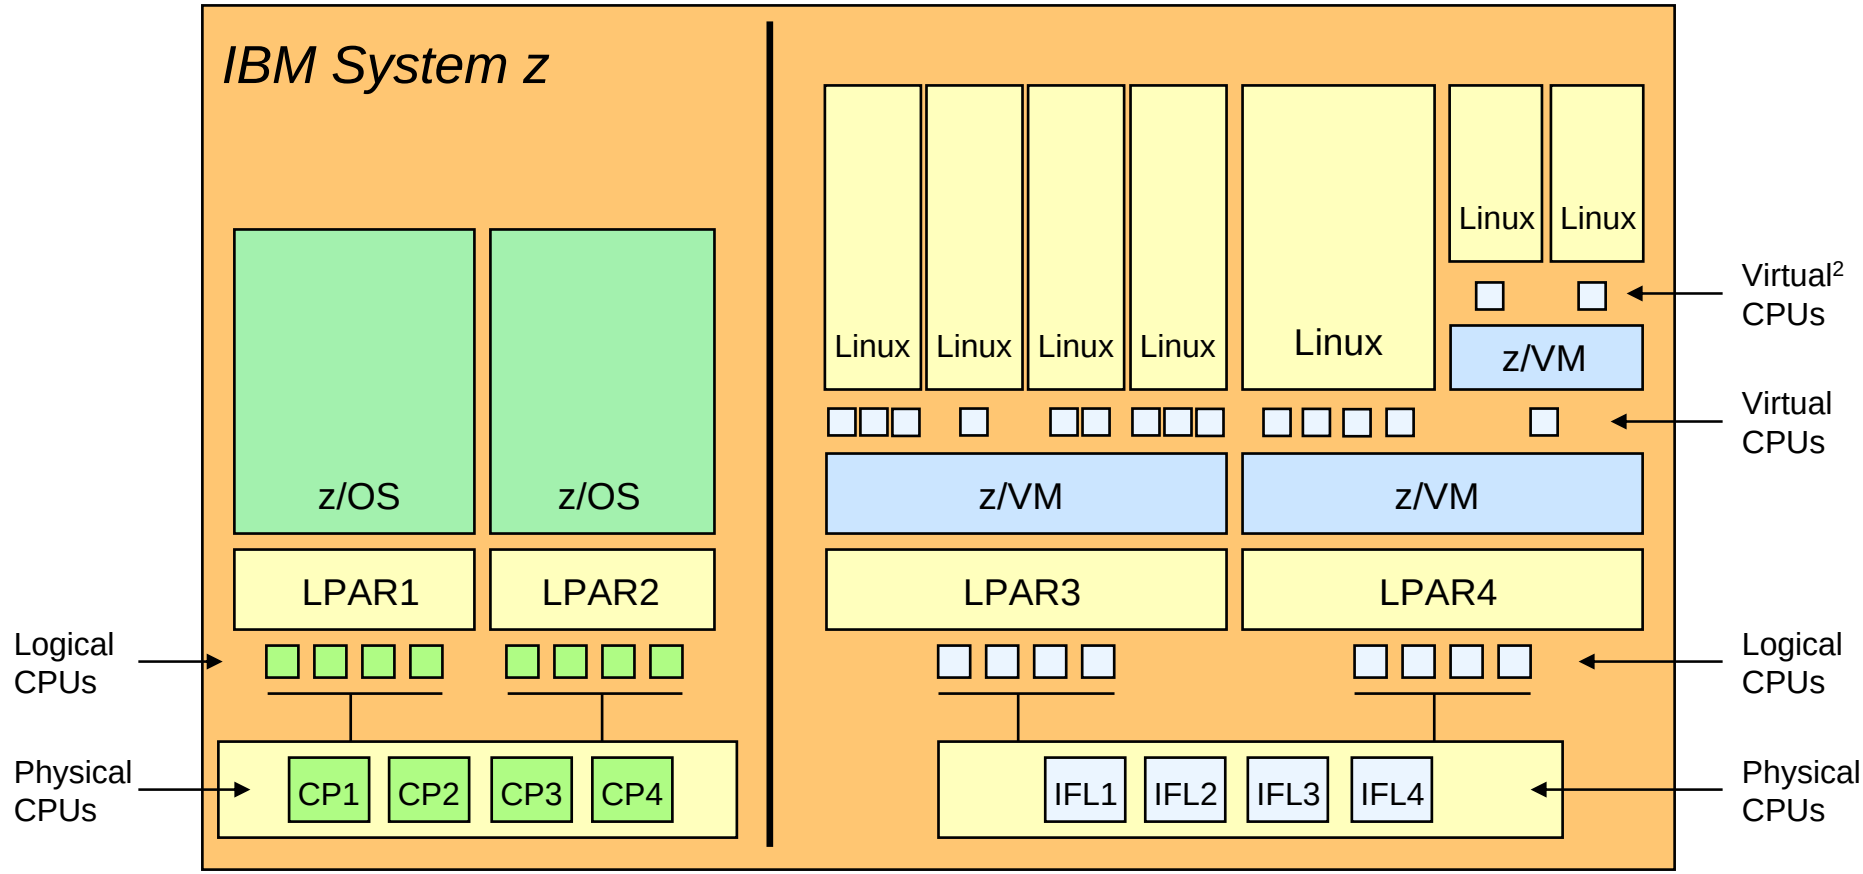

IBM System z CPU High Availability

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***

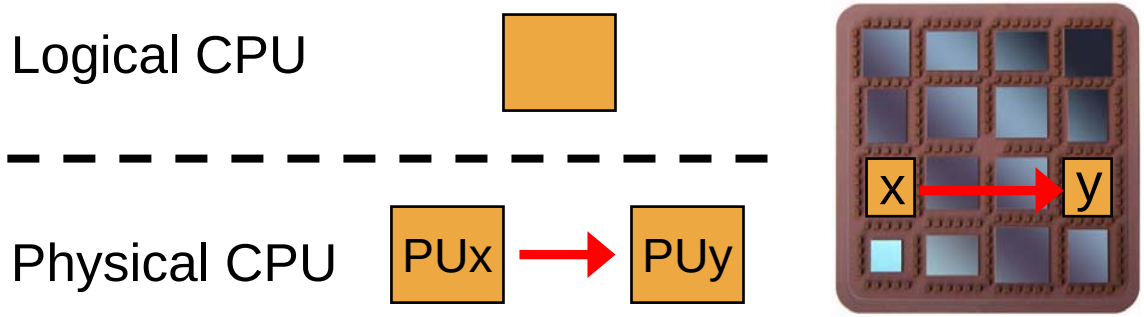

System z Virtualization Technology

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

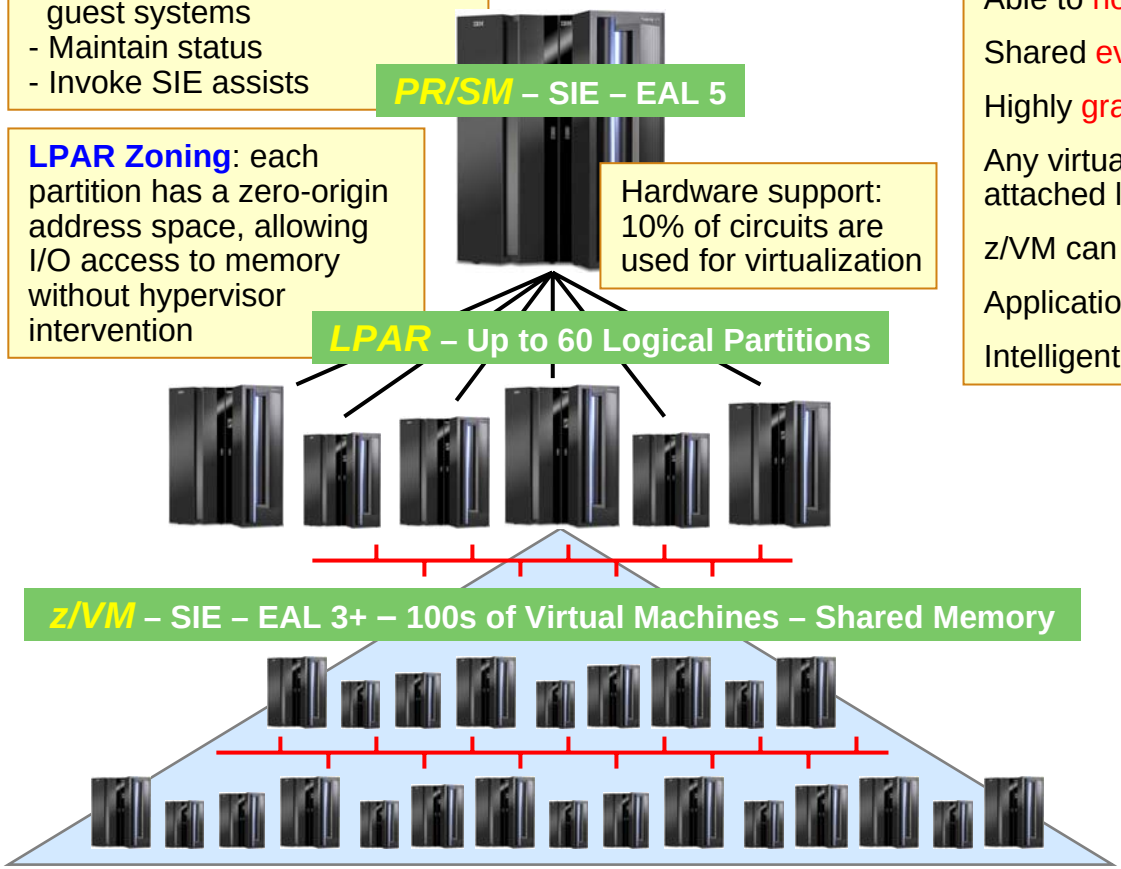

The potential performance impact of the Linux server farm is isolated from the other LPARs

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 Architecture

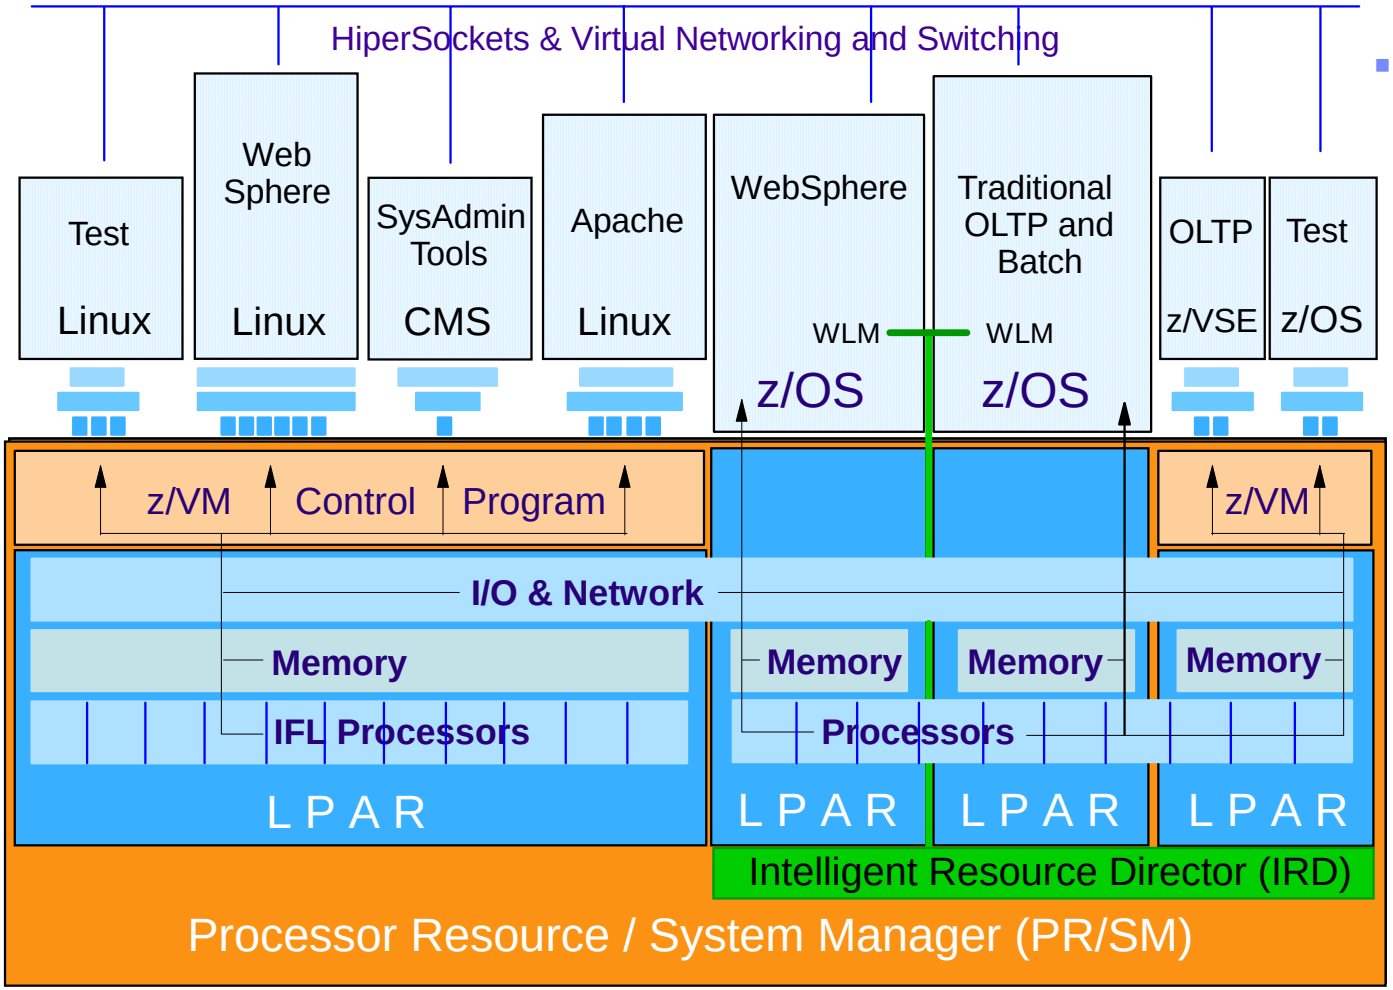

- **Multi-dimensional virtualization technology**
  - System z provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS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

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

Sample Linux-on-z/VM IFL Configuration

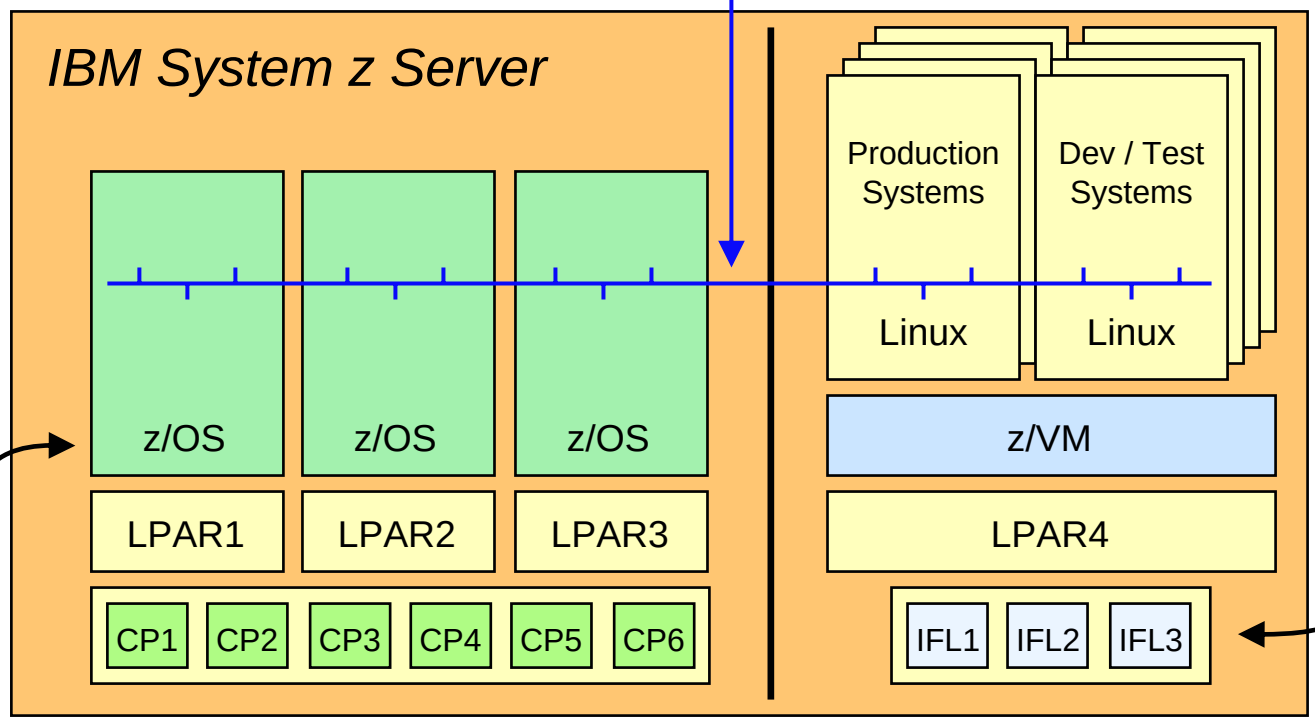

System Design Affects Virtualization Capabilities

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

- With TCO-friendly pricing

Up to 64-way SMP

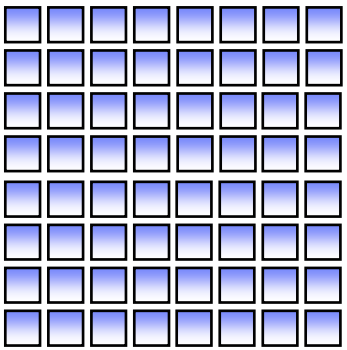

Share up to 64 processors with up to 60 LPARs

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

* No software license fees

2 Standard Spare PUs

Up to 11 System Assist Processors

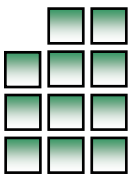

Offload system processing to dedicated CPUs (no impact to software license fees)

Up to 336 I/O Processors

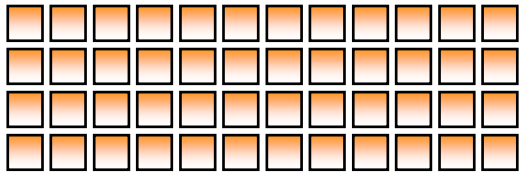

No additional charge for these processors

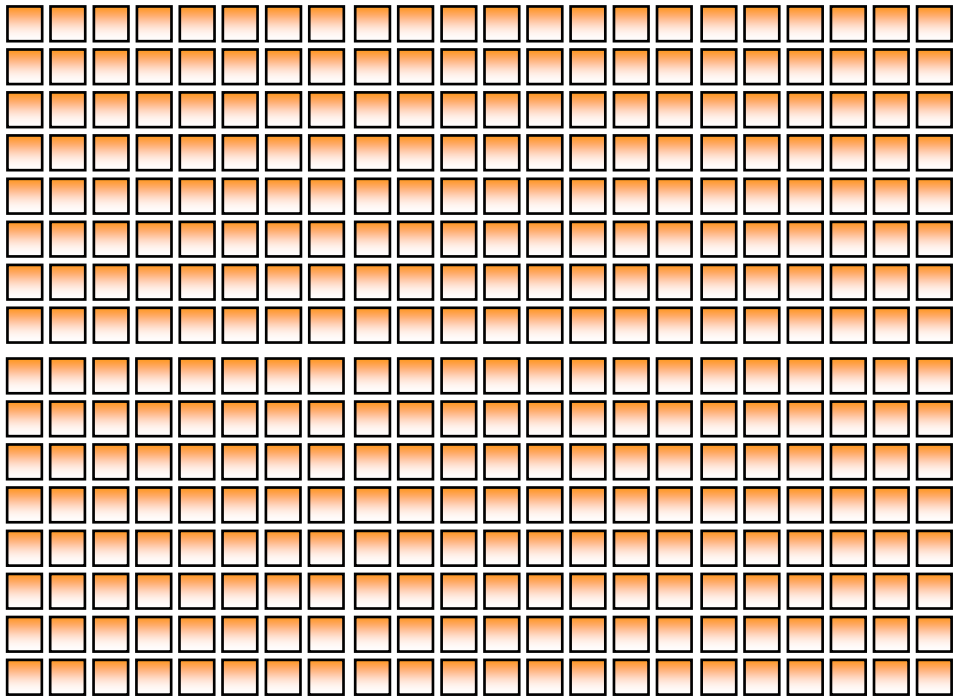

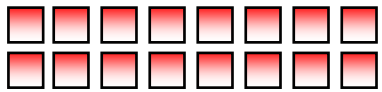

Up to 16 Crypto Express2 CPUs

High scale performance for SSL transactions

System Design Affects Virtualization Capabilities

Compare to typical UNIX system design...

Up to 128-way SMP configuration

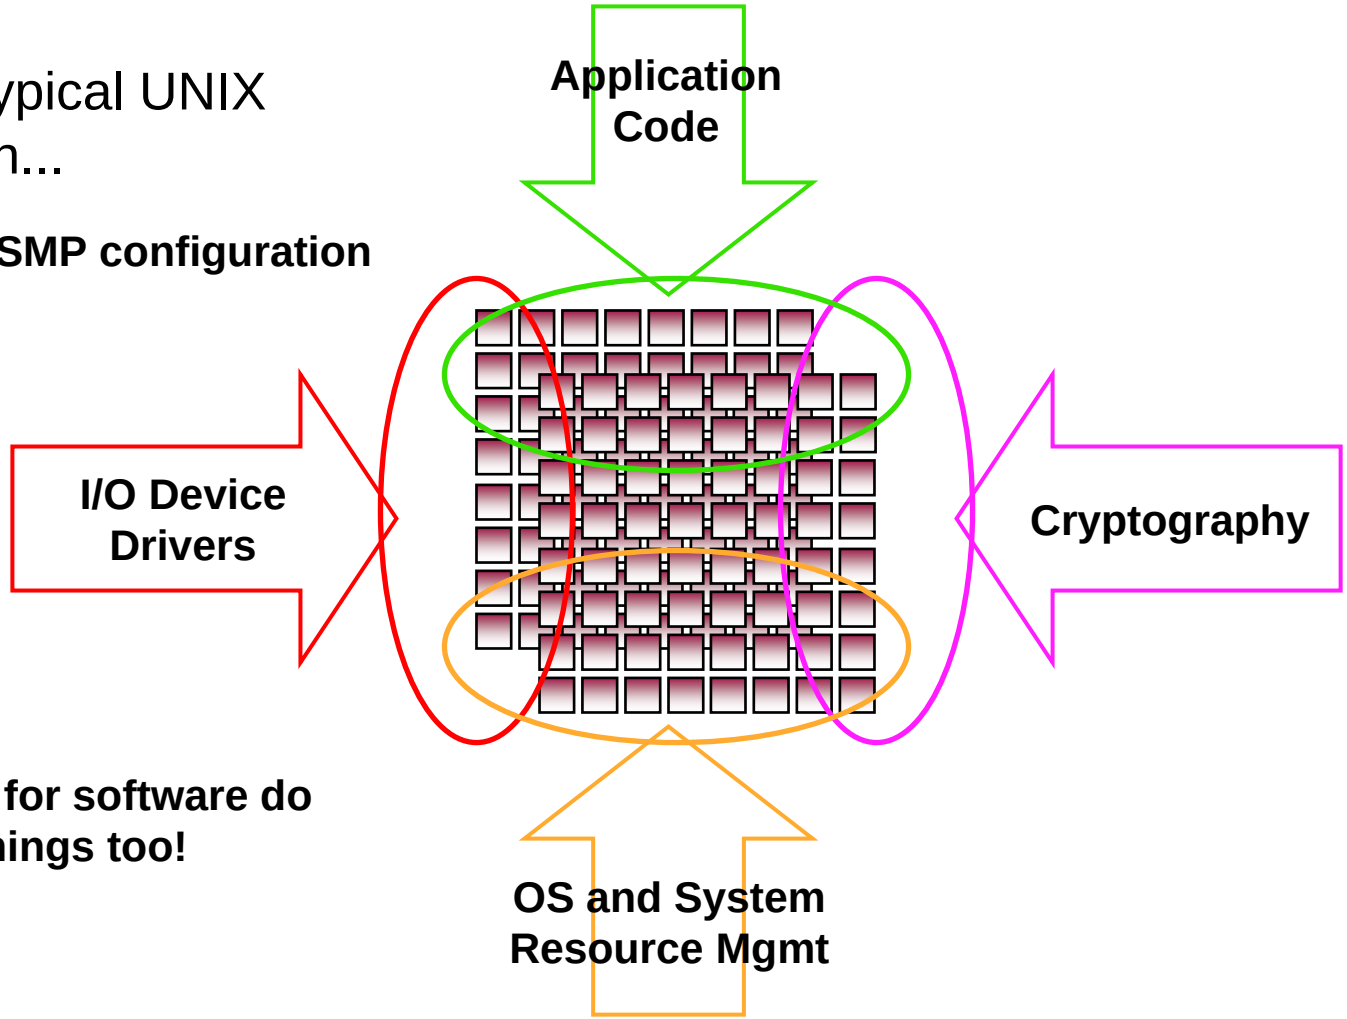

CPUs licensed for software do a lot of other things too!

IBM System z Channel and I/O Architecture

The Big Picture

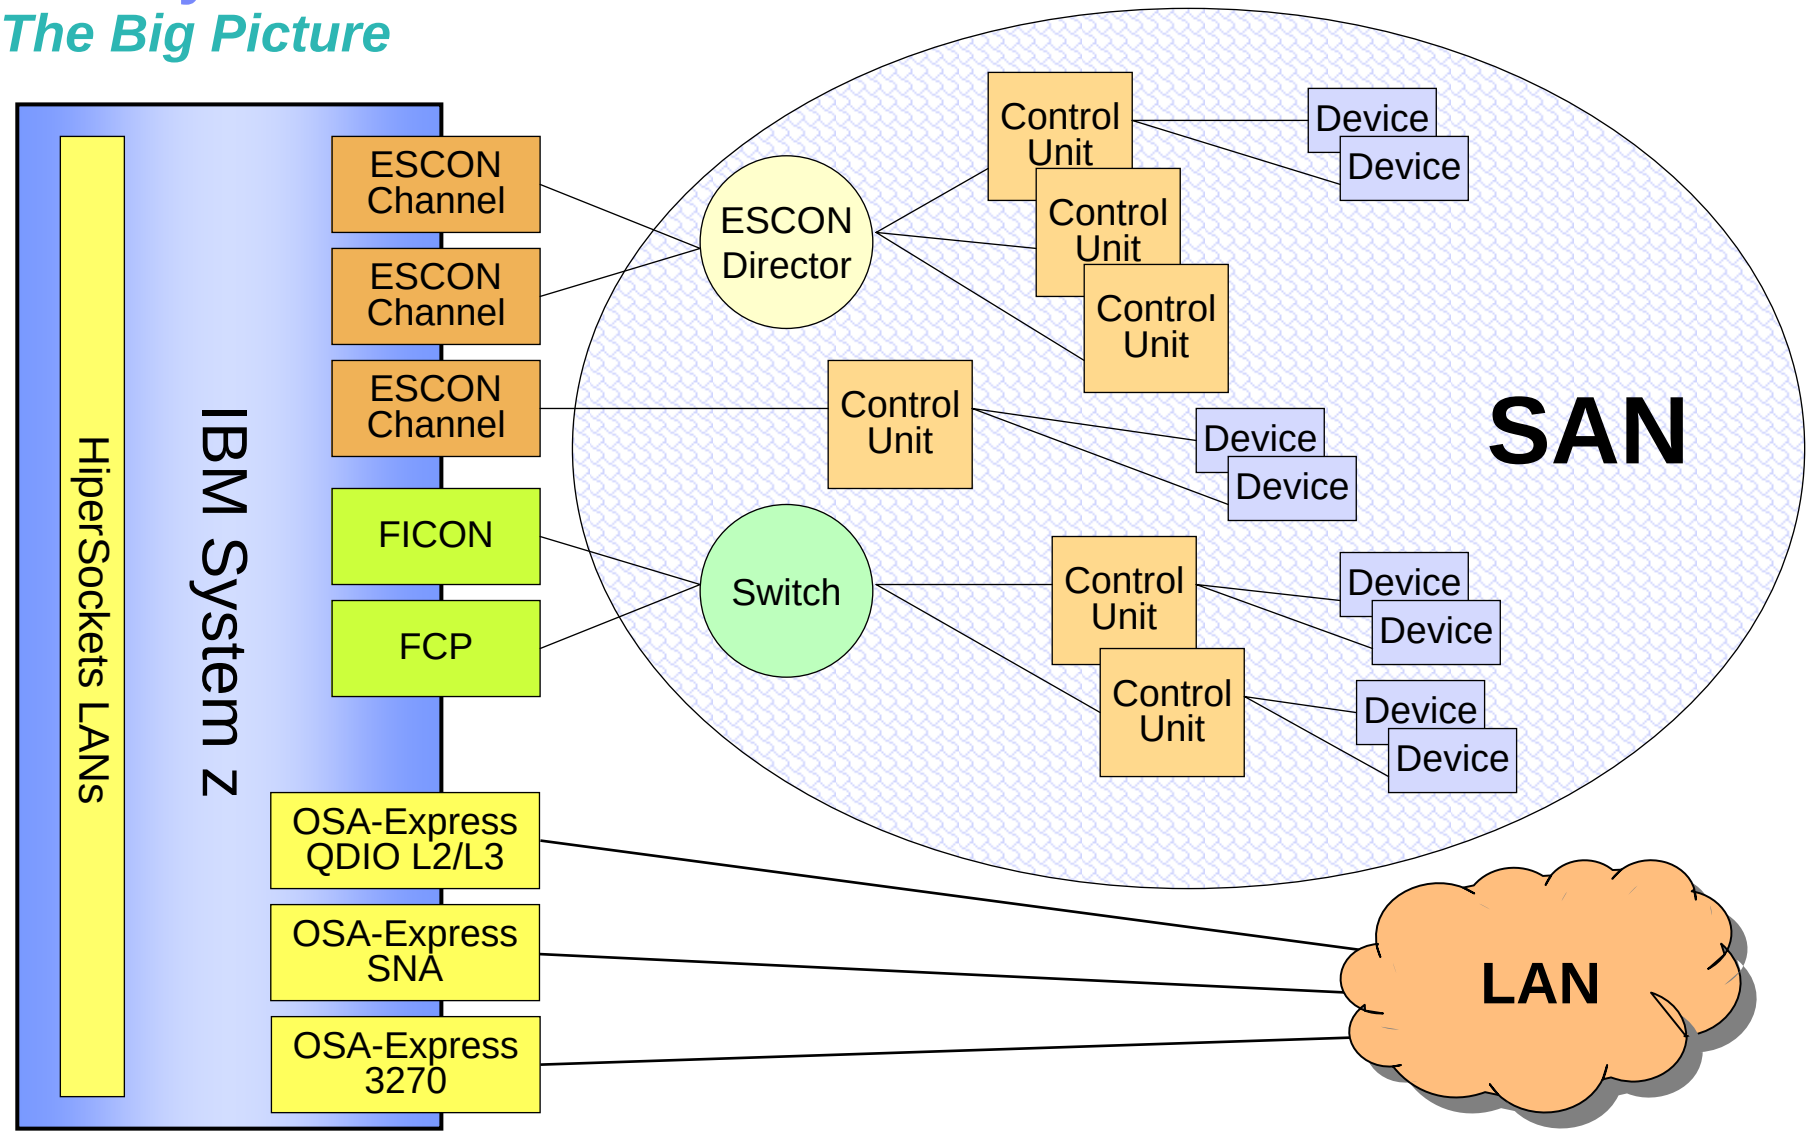

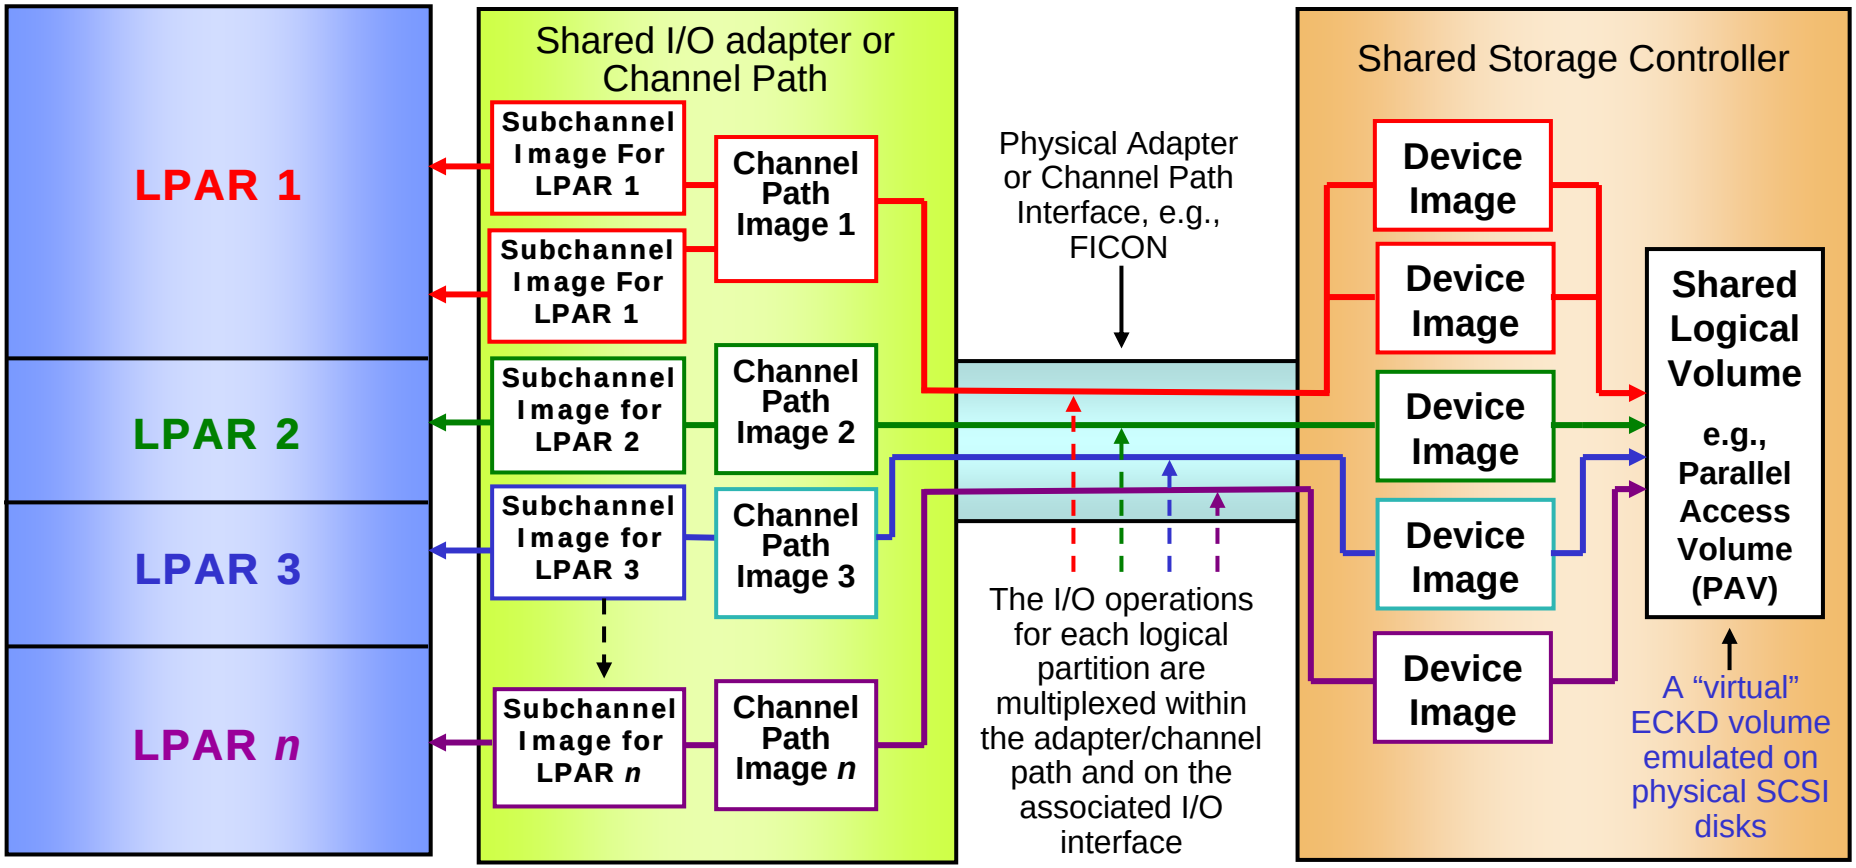

- The I/O infrastructure is shared by LPARs at native speeds, without hypervisor involvement
- Up to 8 physical channels process the I/O requests to the shared devices
  - This reduces the possibility of I/O queuing delays at the channels or at the shared storage controller

Chip Design Affects Virtualization Capabilities

Replicated Server Chip Design

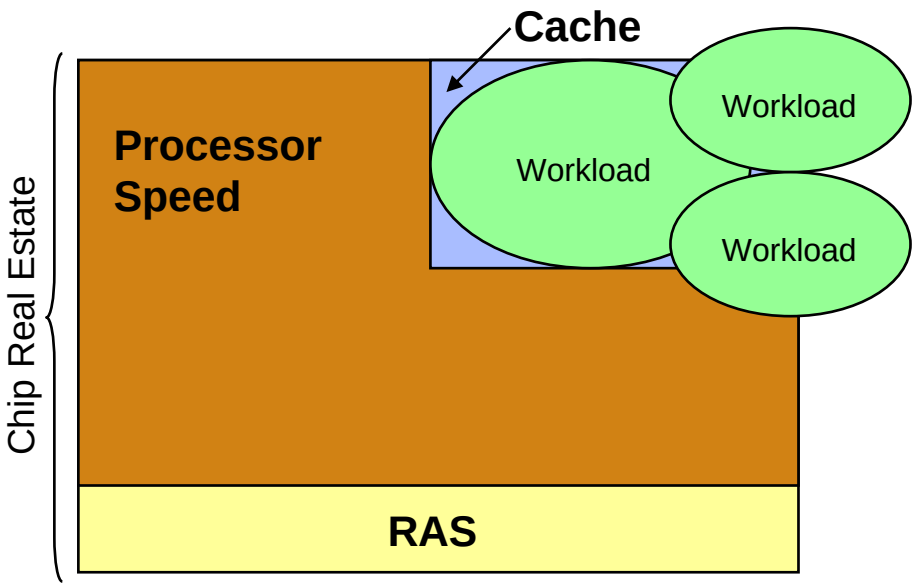

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- “Fast” processor speed is not fully realized due to cache misses

Consolidated Server Chip Design

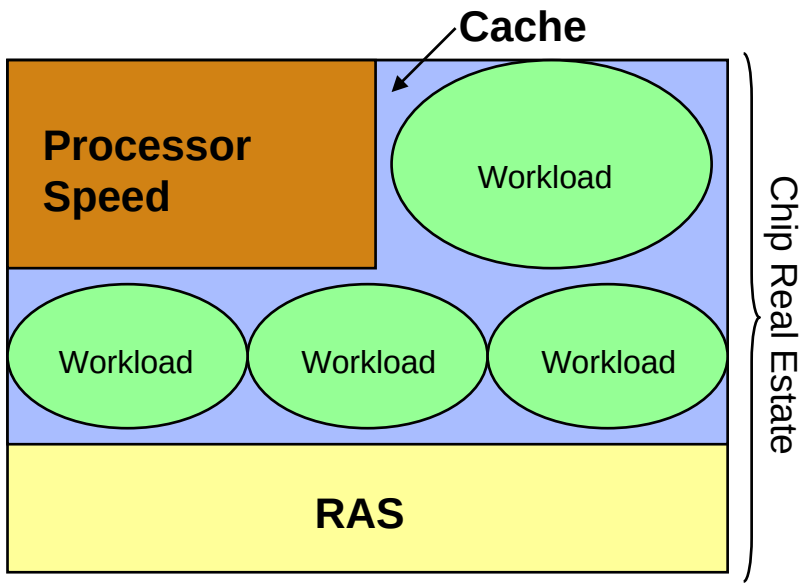

- System z cache is able to contain more working sets
- Processor speed is optimized by increased cache usage
- Additional RAS function is beneficial for mixed workloads

Note: System representations are not to scale, proportions may vary based on generation of chip and model

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**

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

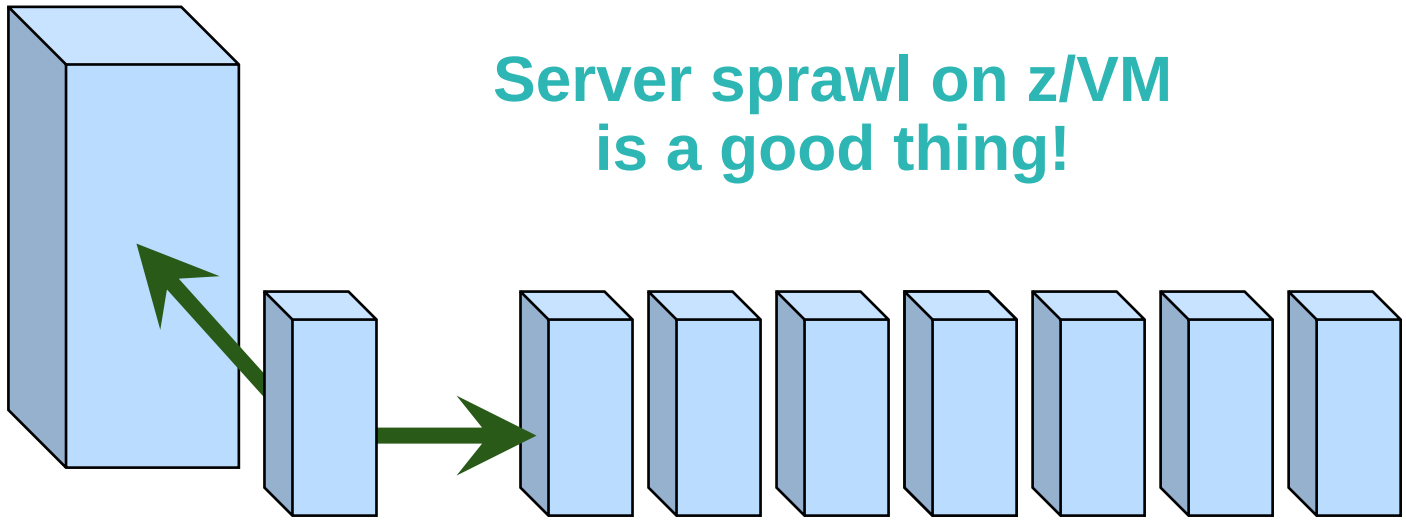

Server sprawl on z/VM is a good thing!

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

Linux-on-z/VM and *Resource Sharing*

Additional Cost Savings and Operational Efficiencies

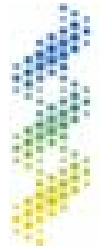

- **A fundamental strength of z/VM is its ability to share system resources to an extreme level**
- **Virtual machines can simultaneously access the I/O and networking resources available on a System z machine**
  - Both real and virtual (z/VM) resources can be used with very high levels of bandwidth and reliability for enhanced workload throughput
- **Linux can exploit z/VM-unique facilities for even higher levels of resource utilization and operational efficiencies**
  - Increase staff productivity and reduce memory consumption by sharing Linux program executables with z/VM DCSS technology
  - Improve memory utilization with Virtual Disks in Storage and Cooperative Memory Management
  - Enhance virtual networking bandwidth and availability using Link Aggregation and the z/VM Virtual Switch

Linux-on-z/VM and *Resource Overcommitment*

A Key Aspect of Cost Savings When Running Linux on System z

- **A fundamental strength of z/VM is its ability to overcommit system resources: "Do more with less"**
- **Users can host an environment that consumes considerably more CPU and memory, in aggregate, than what is configured in the z/VM LPAR**
  - This can translate into cost savings for hardware *and* software
  - Consider a Linux-on-z/VM environment with a 20-to-1 overcommitment of CPU capacity:

Software licensed for two CPUs can run in 40 virtual machines

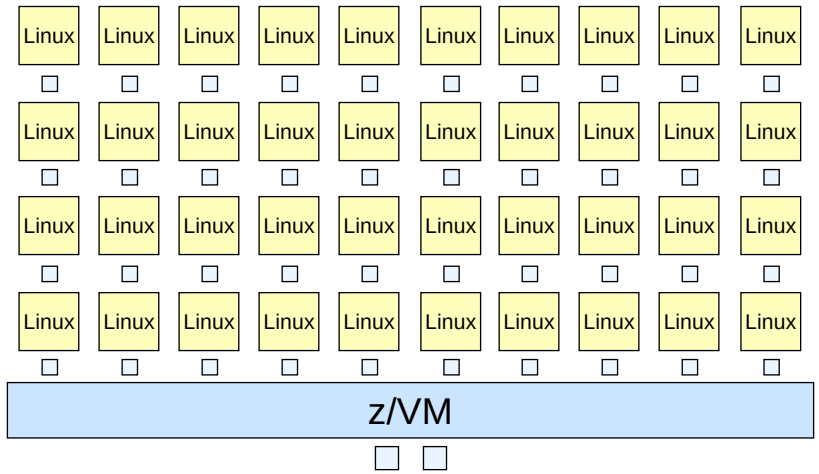

Linux-on-z/VM and *Flexible, Efficient Growth*

New 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**

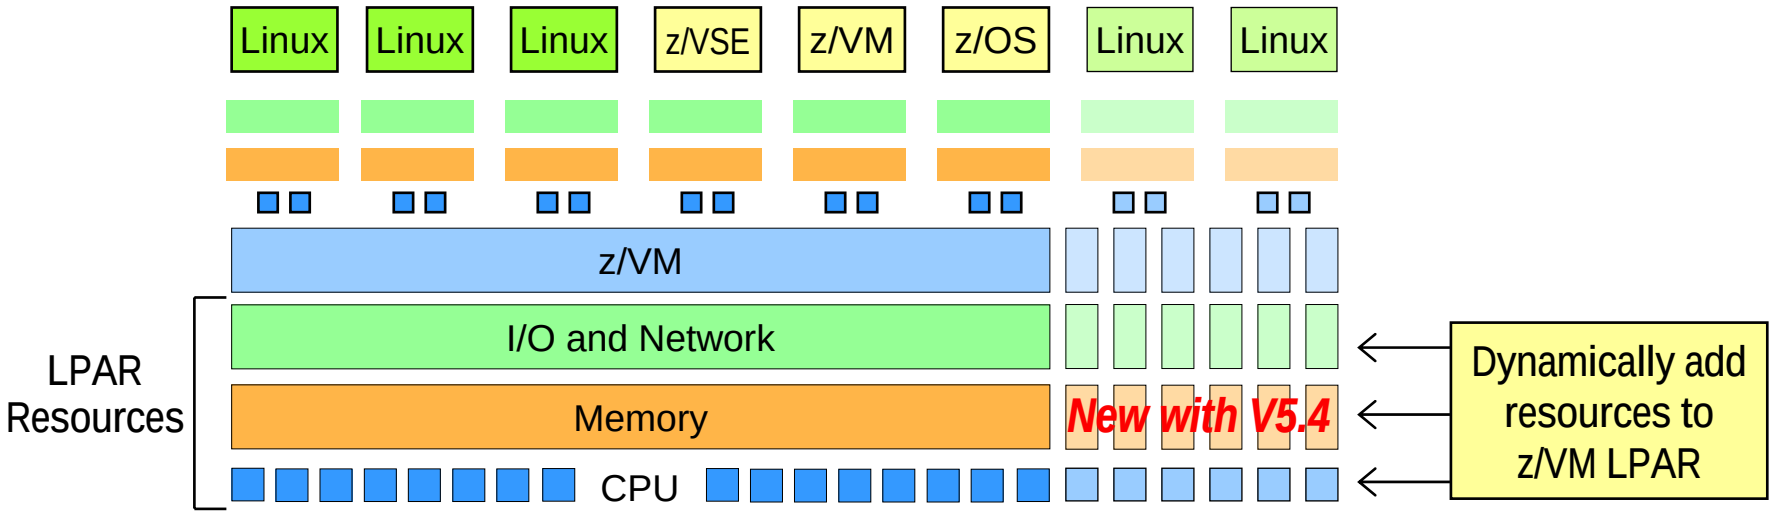

Smart economics: non-disruptively scale your z/VM environment by adding hardware assets that can be shared with *every* virtual server

z/VM V5.4 – 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

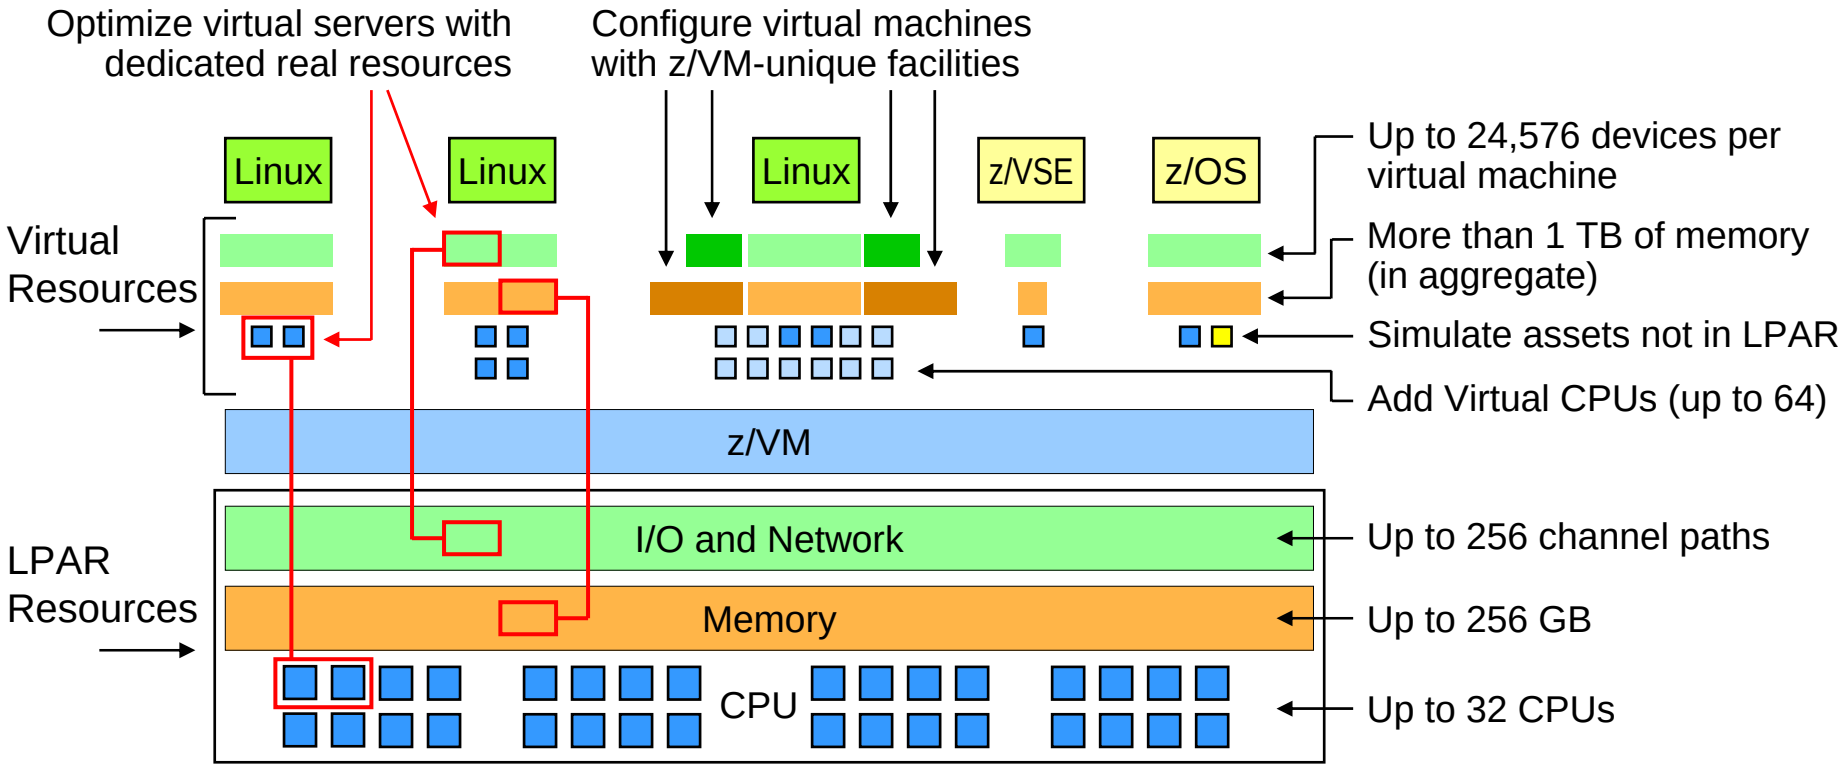

z/VM CPU Resource Controls

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"

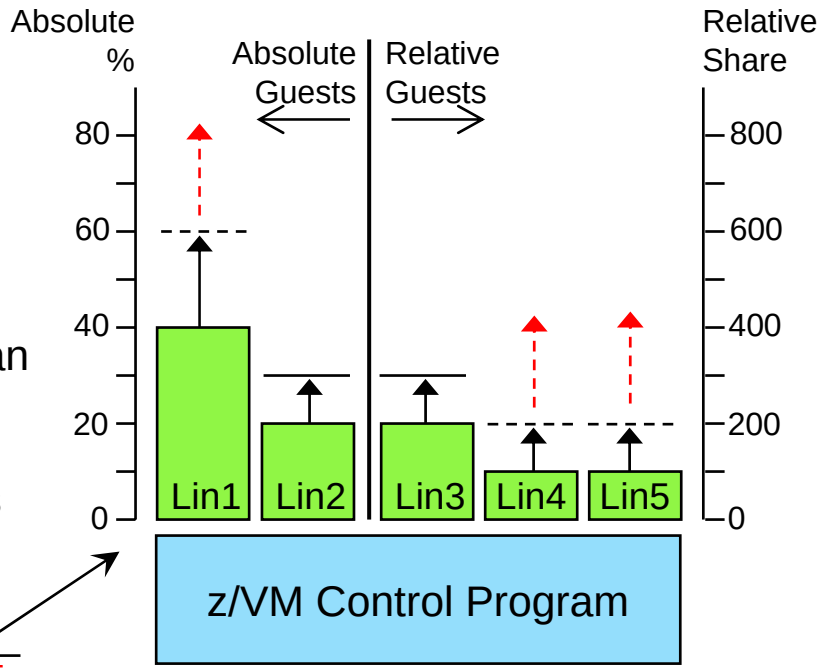


```
SHARE Lin1 ABSOLUTE 40% ABSOLUTE 60% LIMITSOFT
SHARE Lin2 ABSOLUTE 20% ABSOLUTE 30% LIMITHARD
SHARE Lin3 RELATIVE 200 RELATIVE 300 LIMITHARD
SHARE Lin4 RELATIVE 100 RELATIVE 200 LIMITSOFT
SHARE Lin5 RELATIVE 100 RELATIVE 200 LIMITSOFT
```

Notes:

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

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

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

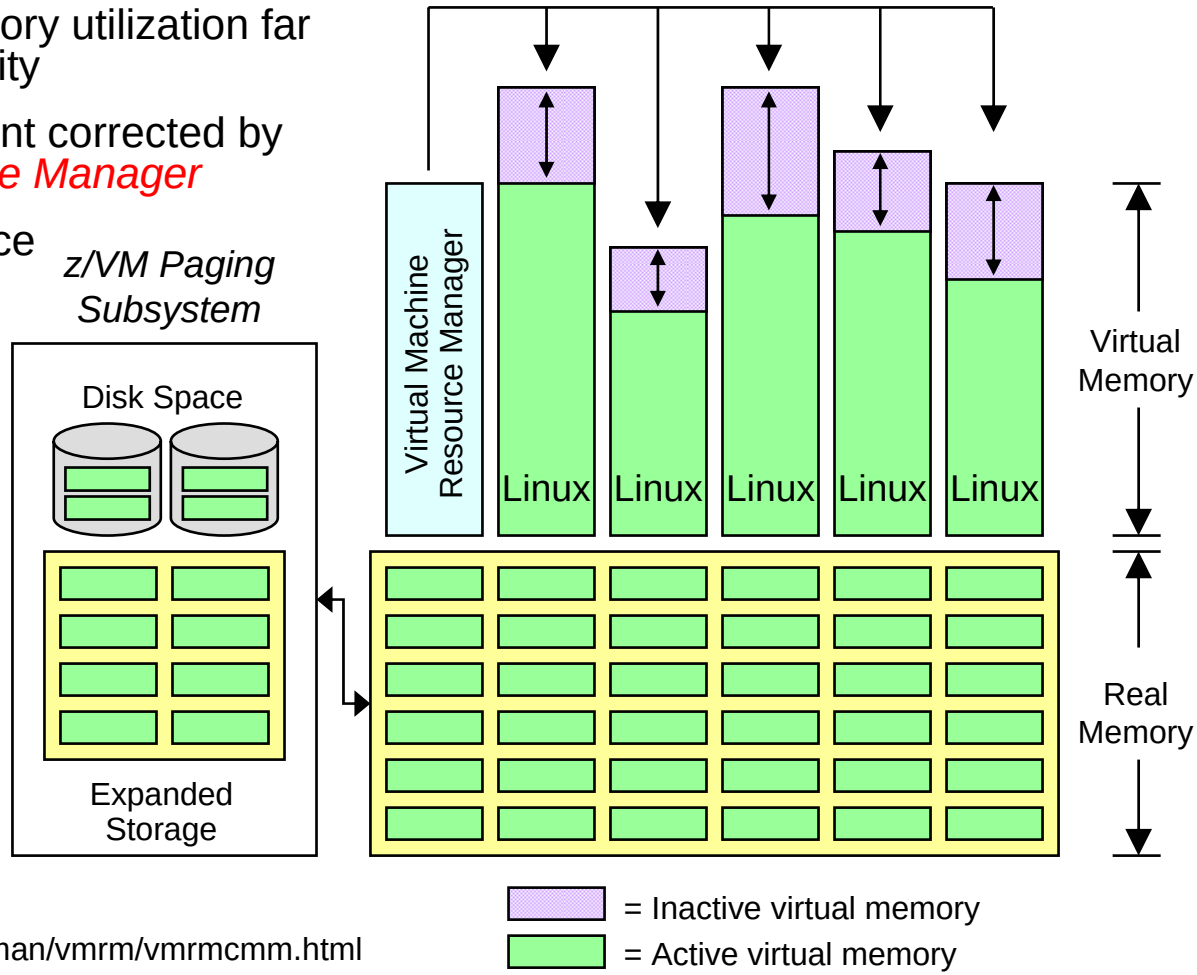

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

OLTP Database Environment with VMRM-CMM and CMMA

Excerpt from "z/VM Large Memory – Linux on System z" Whitepaper

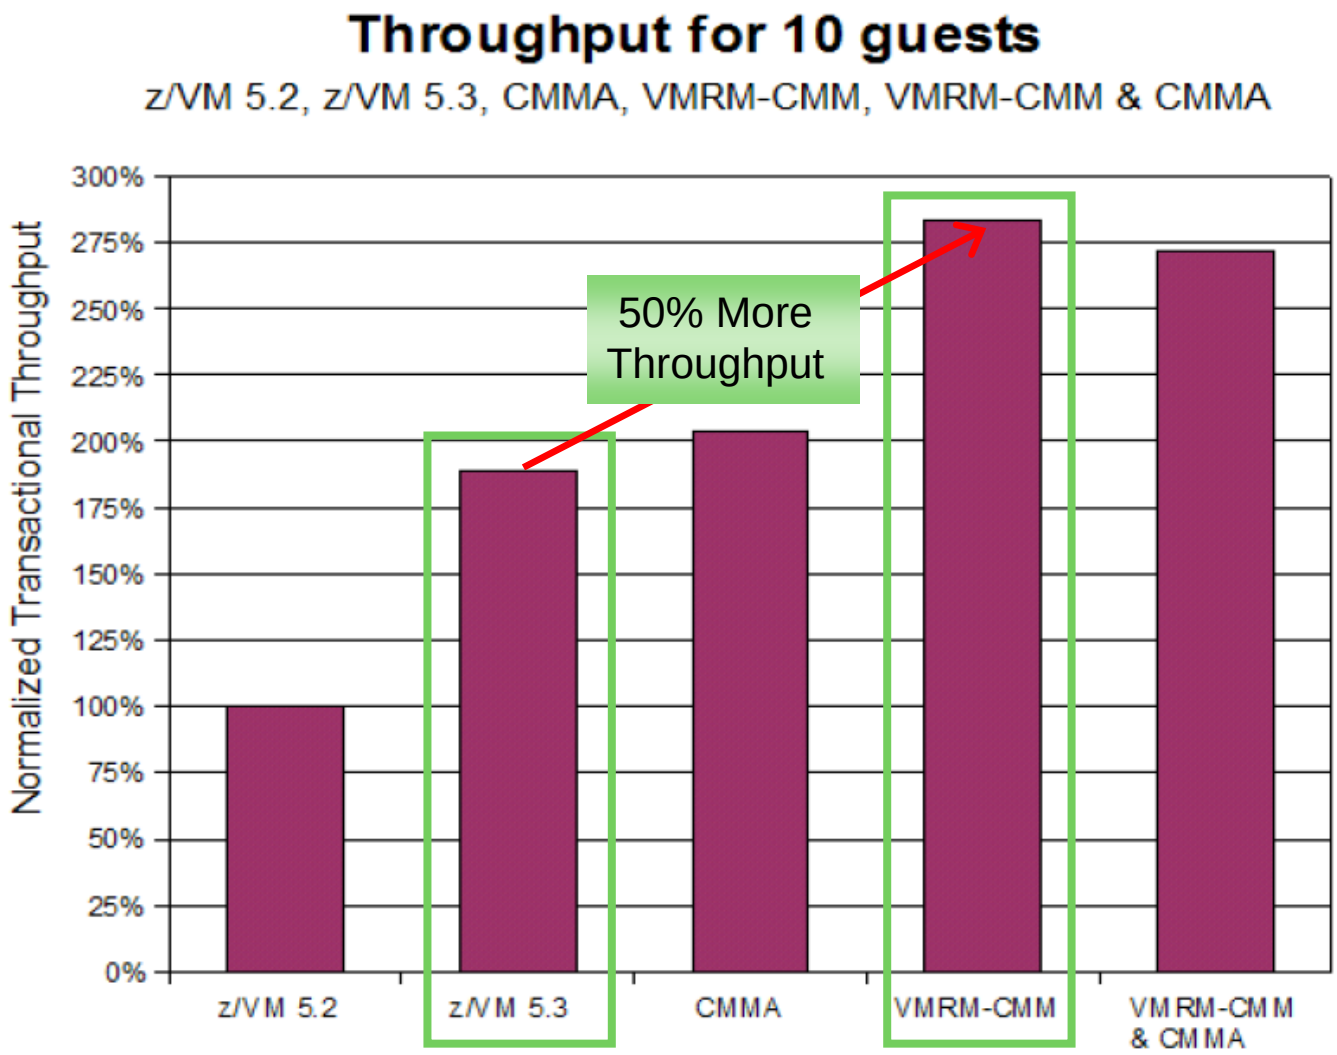

Linux and z/VM Technology Exploitation

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*)
- 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 System z9 and z/VM V5.3 and later
- Linux support available with Novell SLES 10 SP1

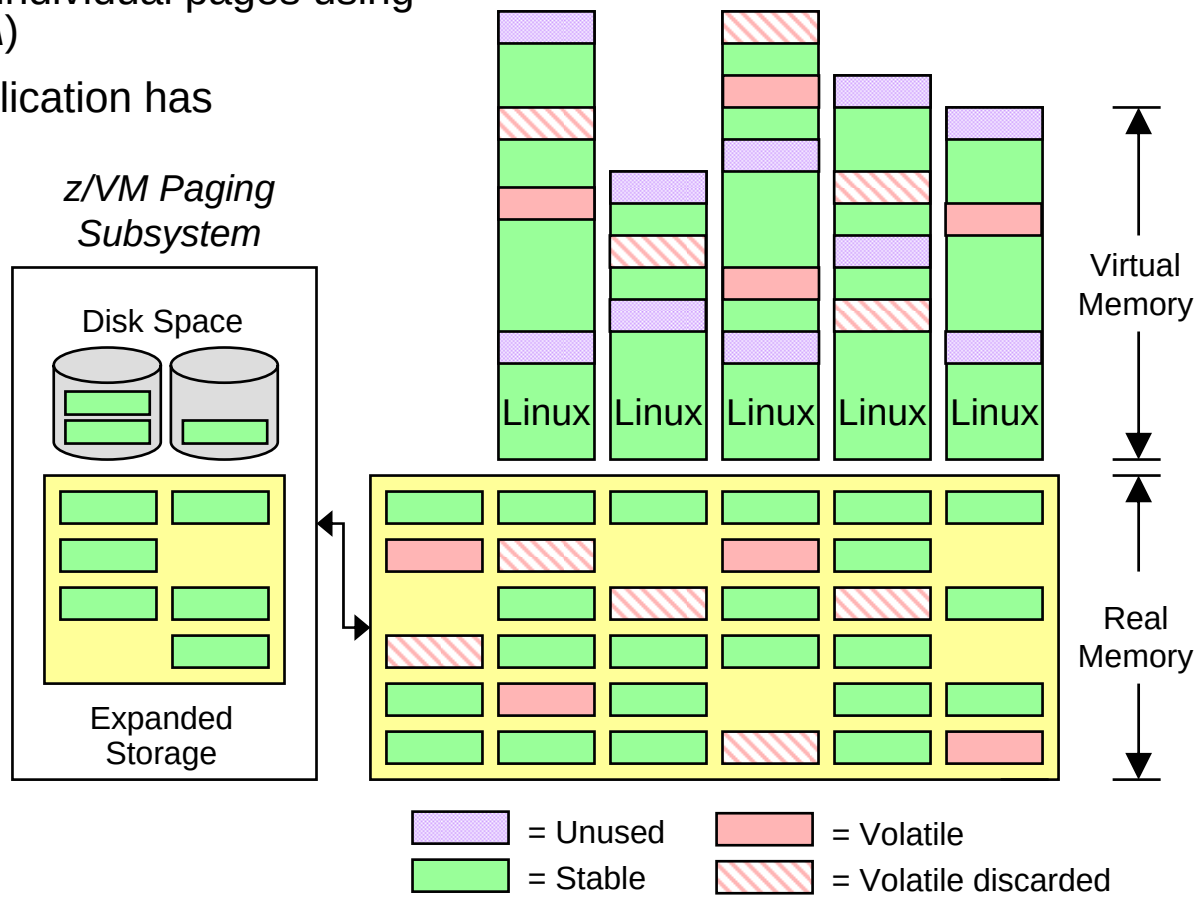

z/VM Technology: Advanced Disk Support

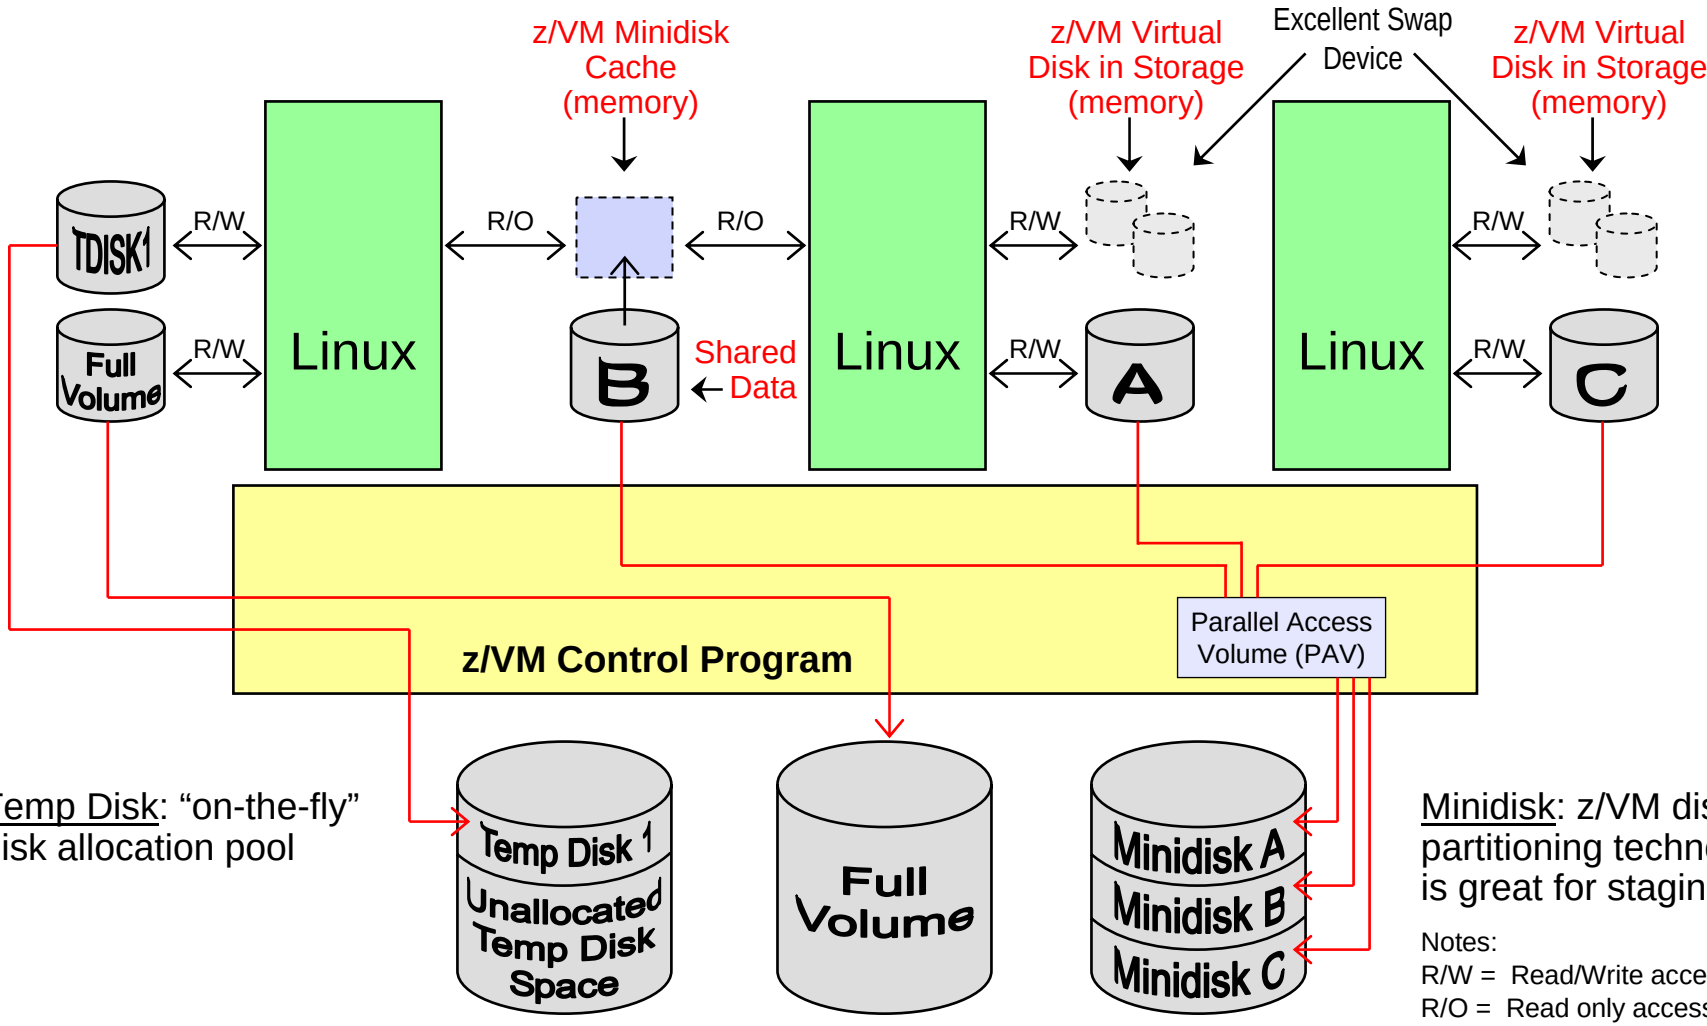

Temp Disk: "on-the-fly" disk allocation pool

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

- Discontinguous 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

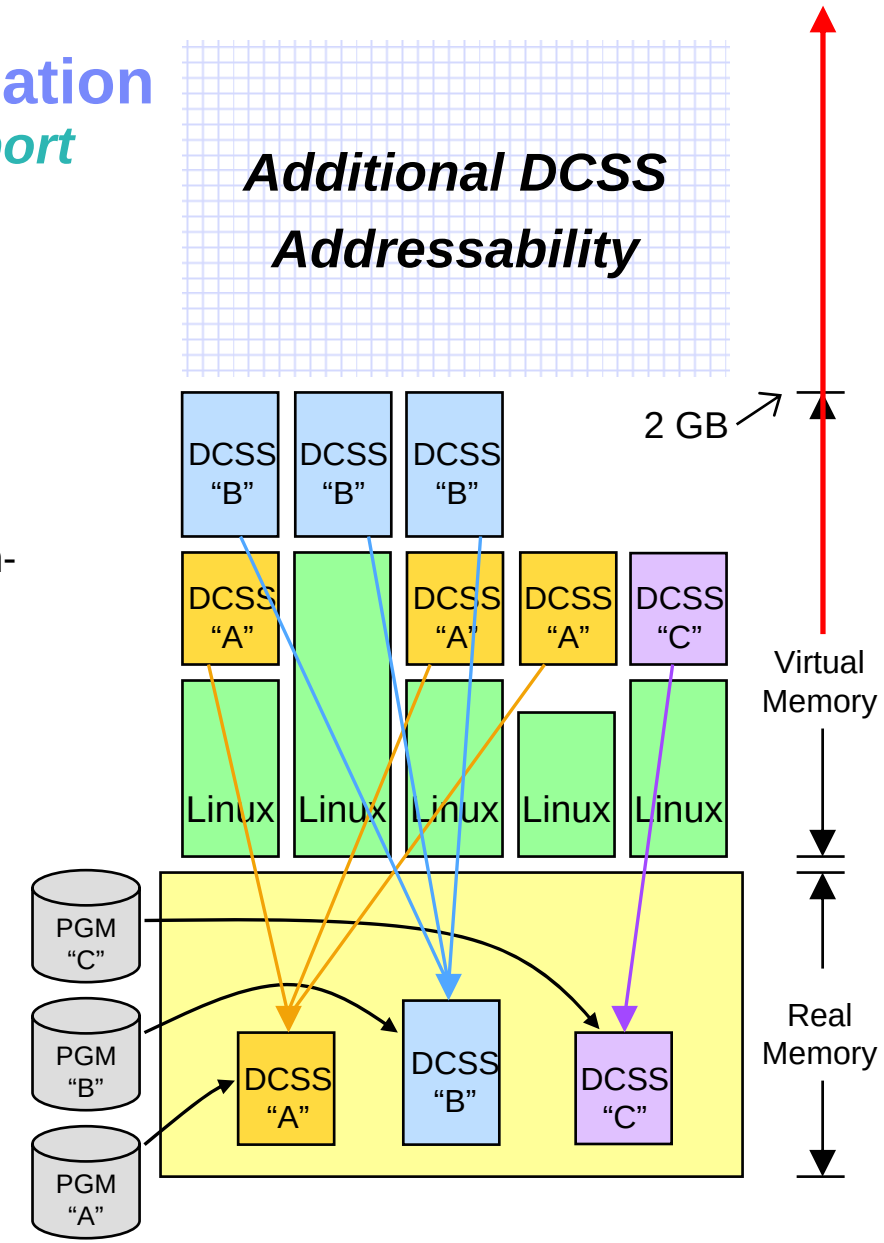

Extreme Virtualization with Linux on z/VM

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

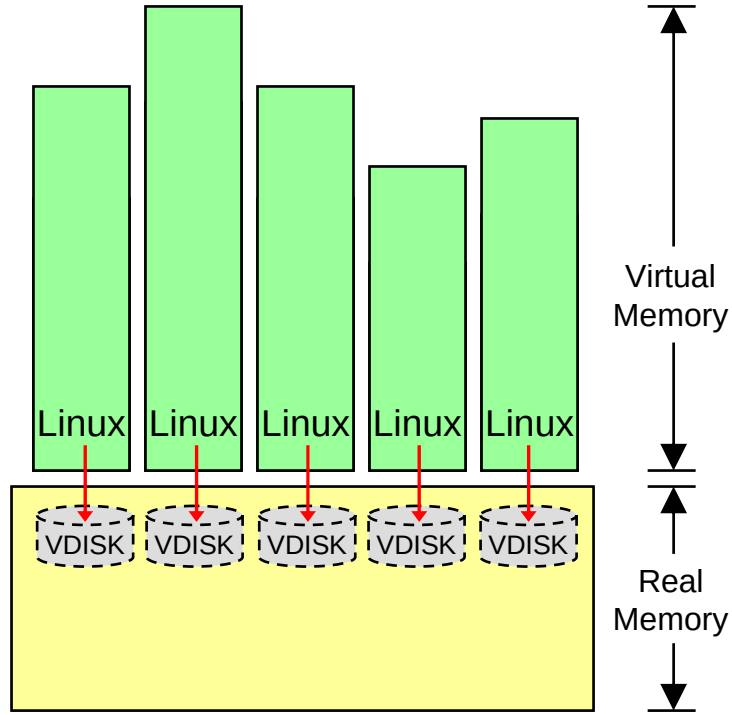

System z and N_Port ID Virtualization (NPIV)

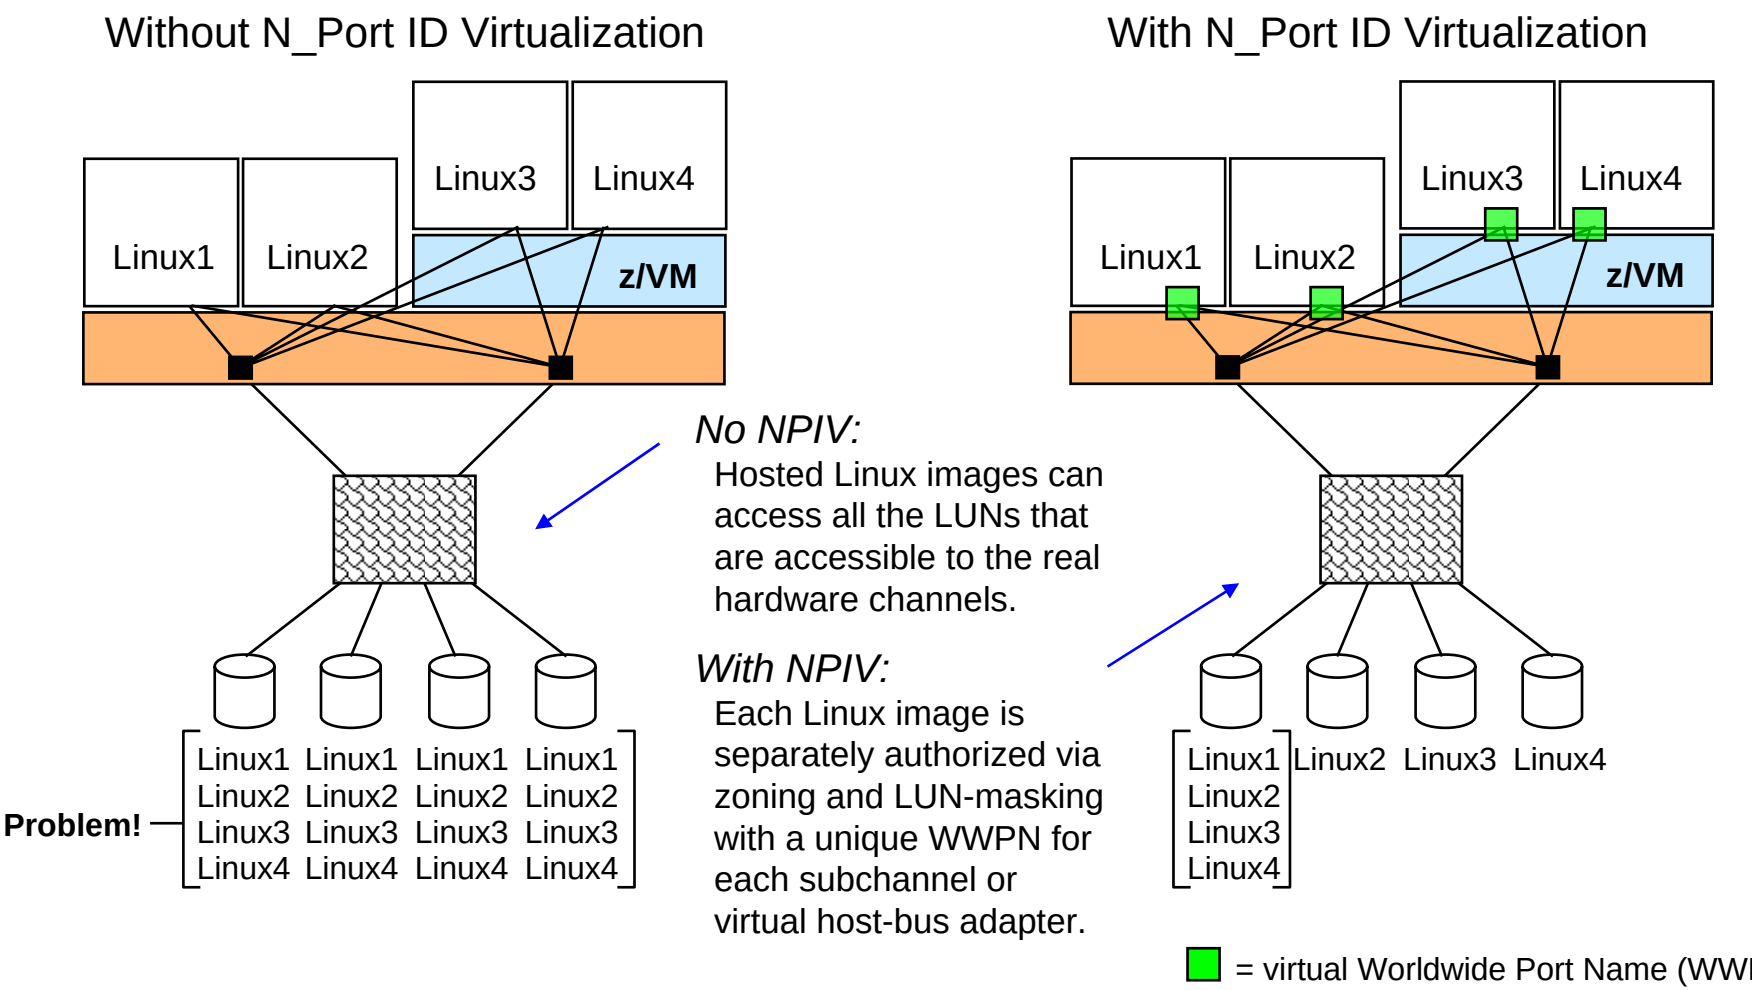

No NPIV:

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.

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 base product**
  - z/VM V5.2 support available with PTF for APAR VM64128

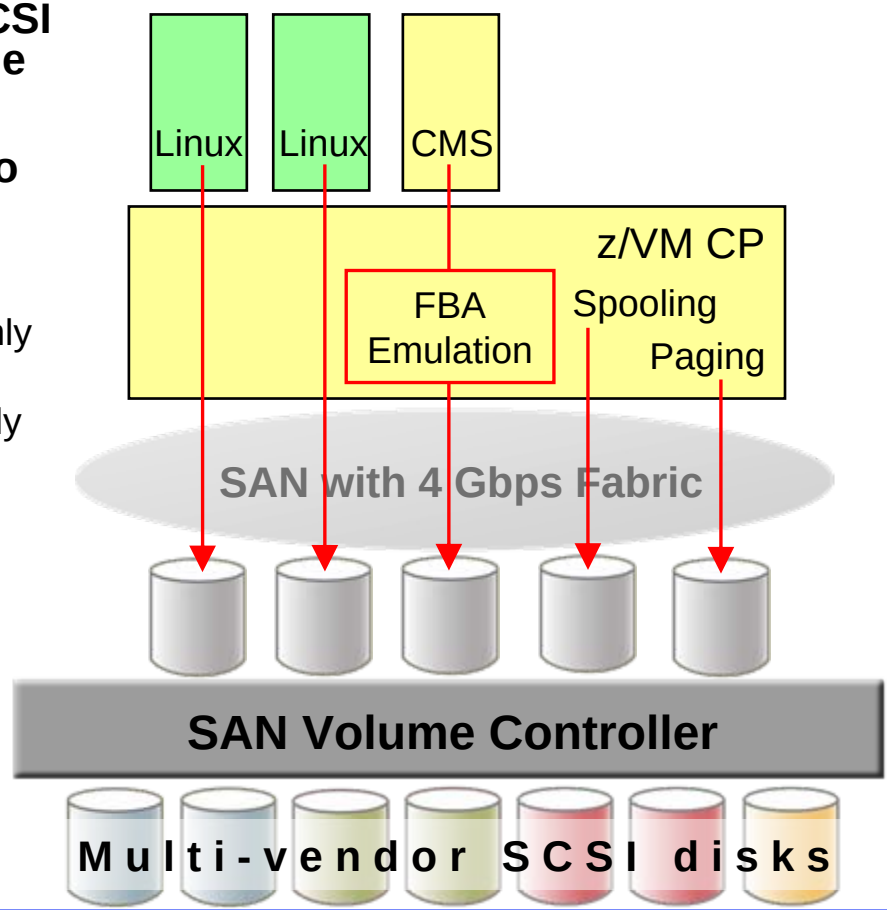

Learn more at: **ibm.com**/storage/support/2145

z/VM Virtual Switch Link Aggregation Support

Enhanced Networking Bandwidth and Business Continuance

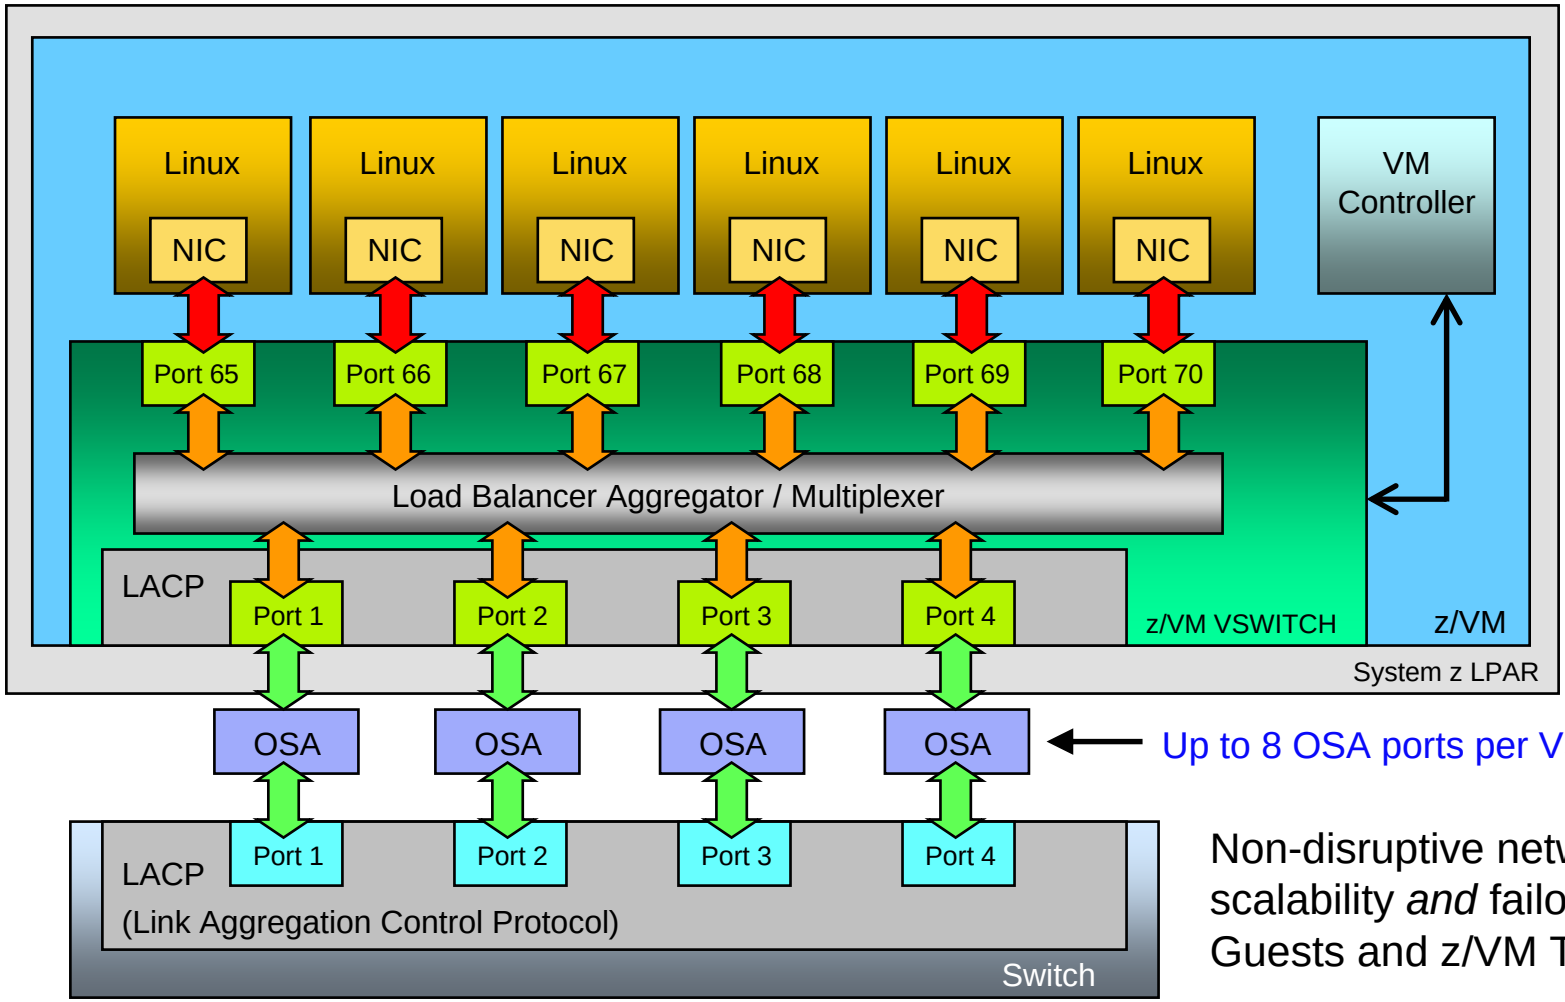

Non-disruptive networking scalability *and* failover for Guests and z/VM TCP/IP.

z/VM Virtual Switch Link Aggregation

With z/VM TCP/IP Stack Connectivity Support in z/VM V5.4

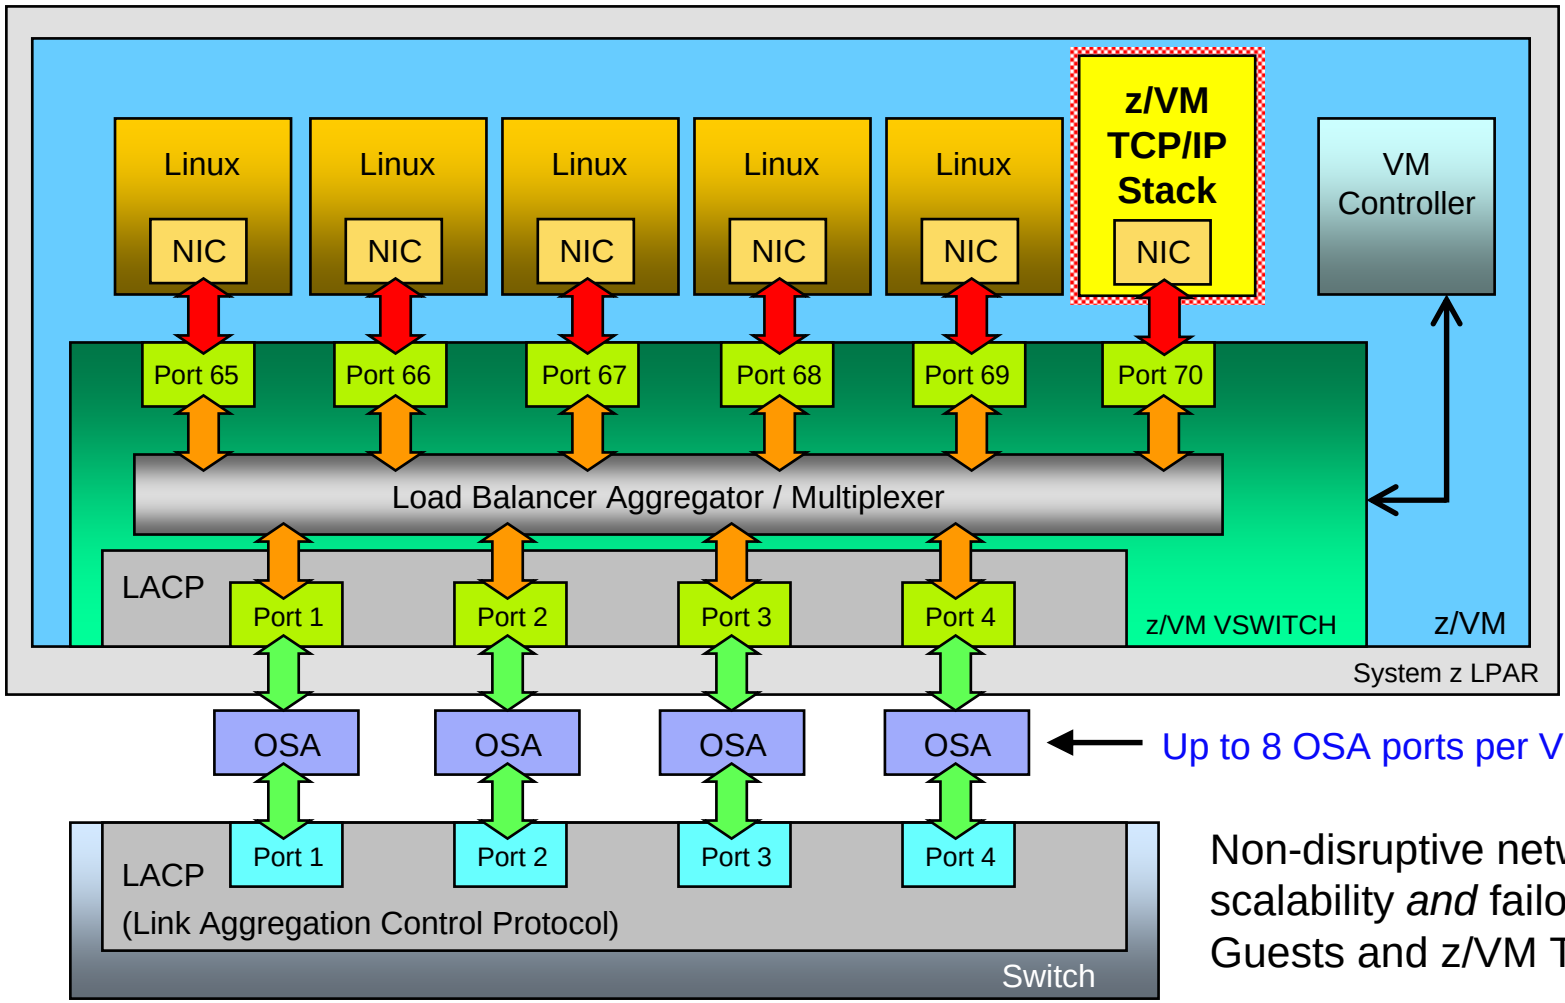

Non-disruptive networking scalability *and* failover for Guests and z/VM TCP/IP.

z/VM Command and Control Infrastructure

- **Built-in z/VM facilities enable cost-effective command and control**
  - Performance data collection and reporting for every Linux image
  - Log accounting records for charge-back
  - Automate system operations with CMS, REXX, Pipelines, virtual console interrogation using PROP (VM programmable operator)
  - Dynamic I/O reconfiguration (e.g., dynamically add more disks)
  - Run EREP on z/VM for system-level hardware error reporting
  - Priced z/VM features:
    - DirMaint – simplifies task of adding/modifying/deleting users
    - Performance Toolkit for VM – performance recording and reporting
    - RACF Security Server for z/VM – security services (including LDAP)
    - RSCS – provides NJE connectivity support for Linux systems
- **Samples, examples, downloads available**
  - IBM Redbooks
  - z/VM web site (www.vm.ibm.com/download)
- **Extensive suite of solutions available from ISVs**
  - Visit: **ibm.com**/systems/z/os/linux/apps/all.html

z/VM Technology – Command and Control Infrastructure

Leveraging the IBM Software Portfolio

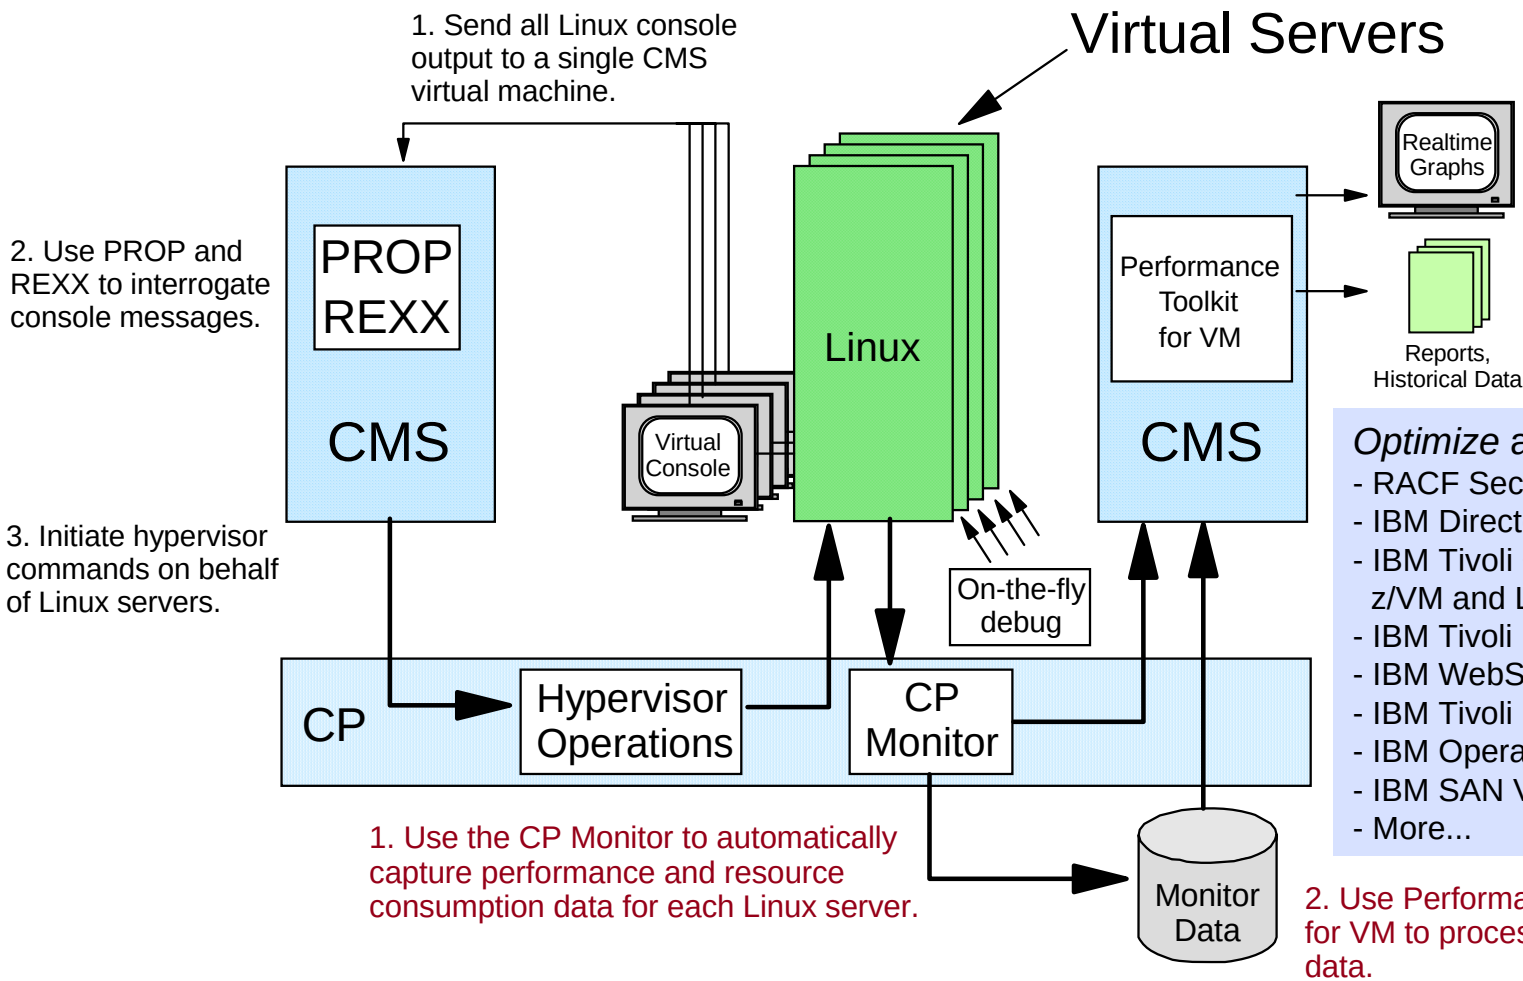

Optimize and *Integrate* with:

- RACF Security Server for z/VM
- IBM Director (z/VM Center)
- IBM Tivoli OMEGAMON XE for z/VM and Linux
- IBM Tivoli Provisioning Manager
- IBM WebSphere solutions
- IBM Tivoli Monitoring
- IBM Operations Manager for z/VM
- IBM SAN Volume Controller
- More...

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

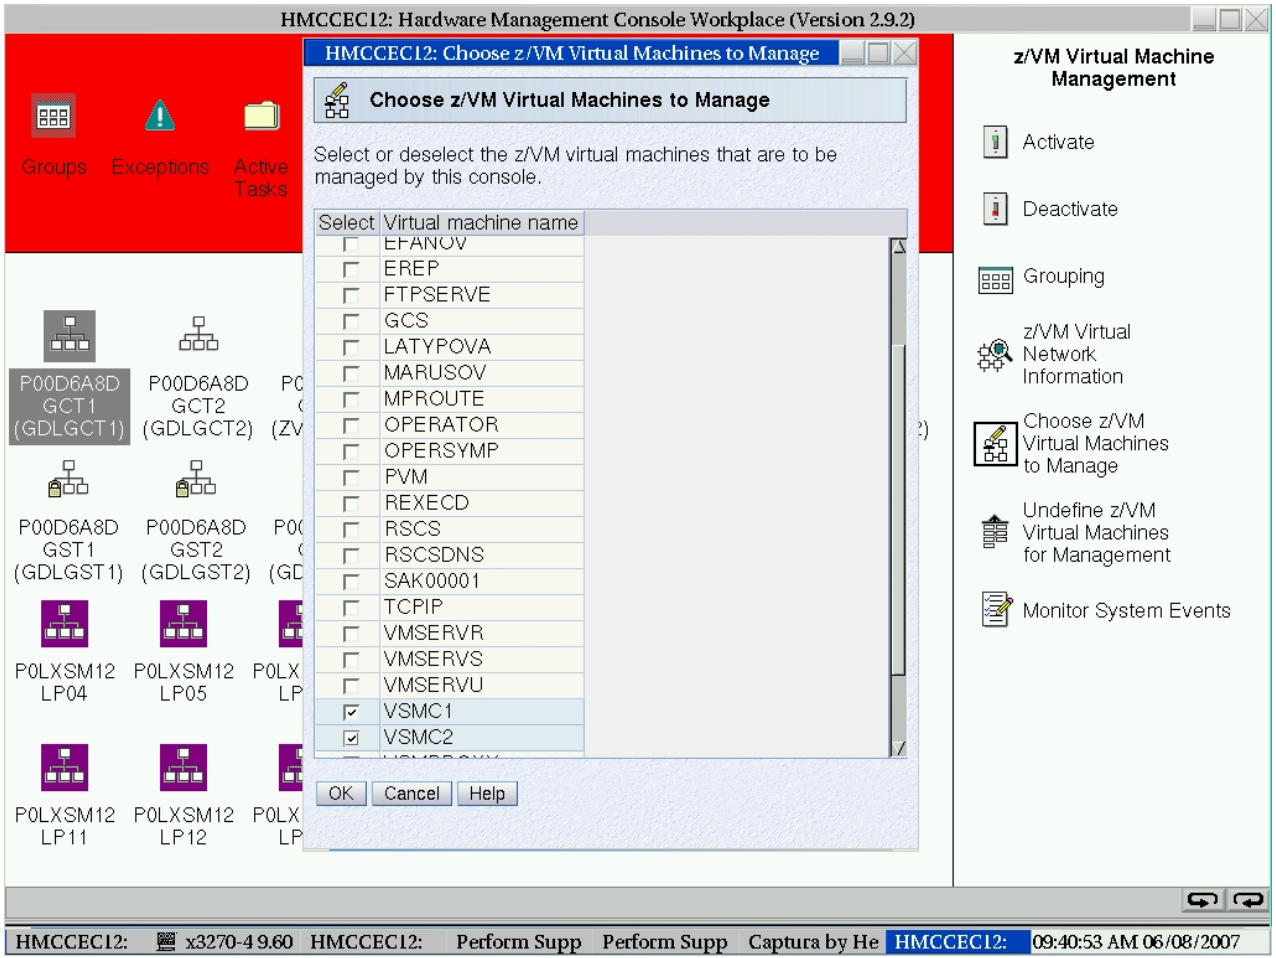

Provisioning Linux Virtual Machines on System z

Using IBM Director for Linux on System z with z/VM Center

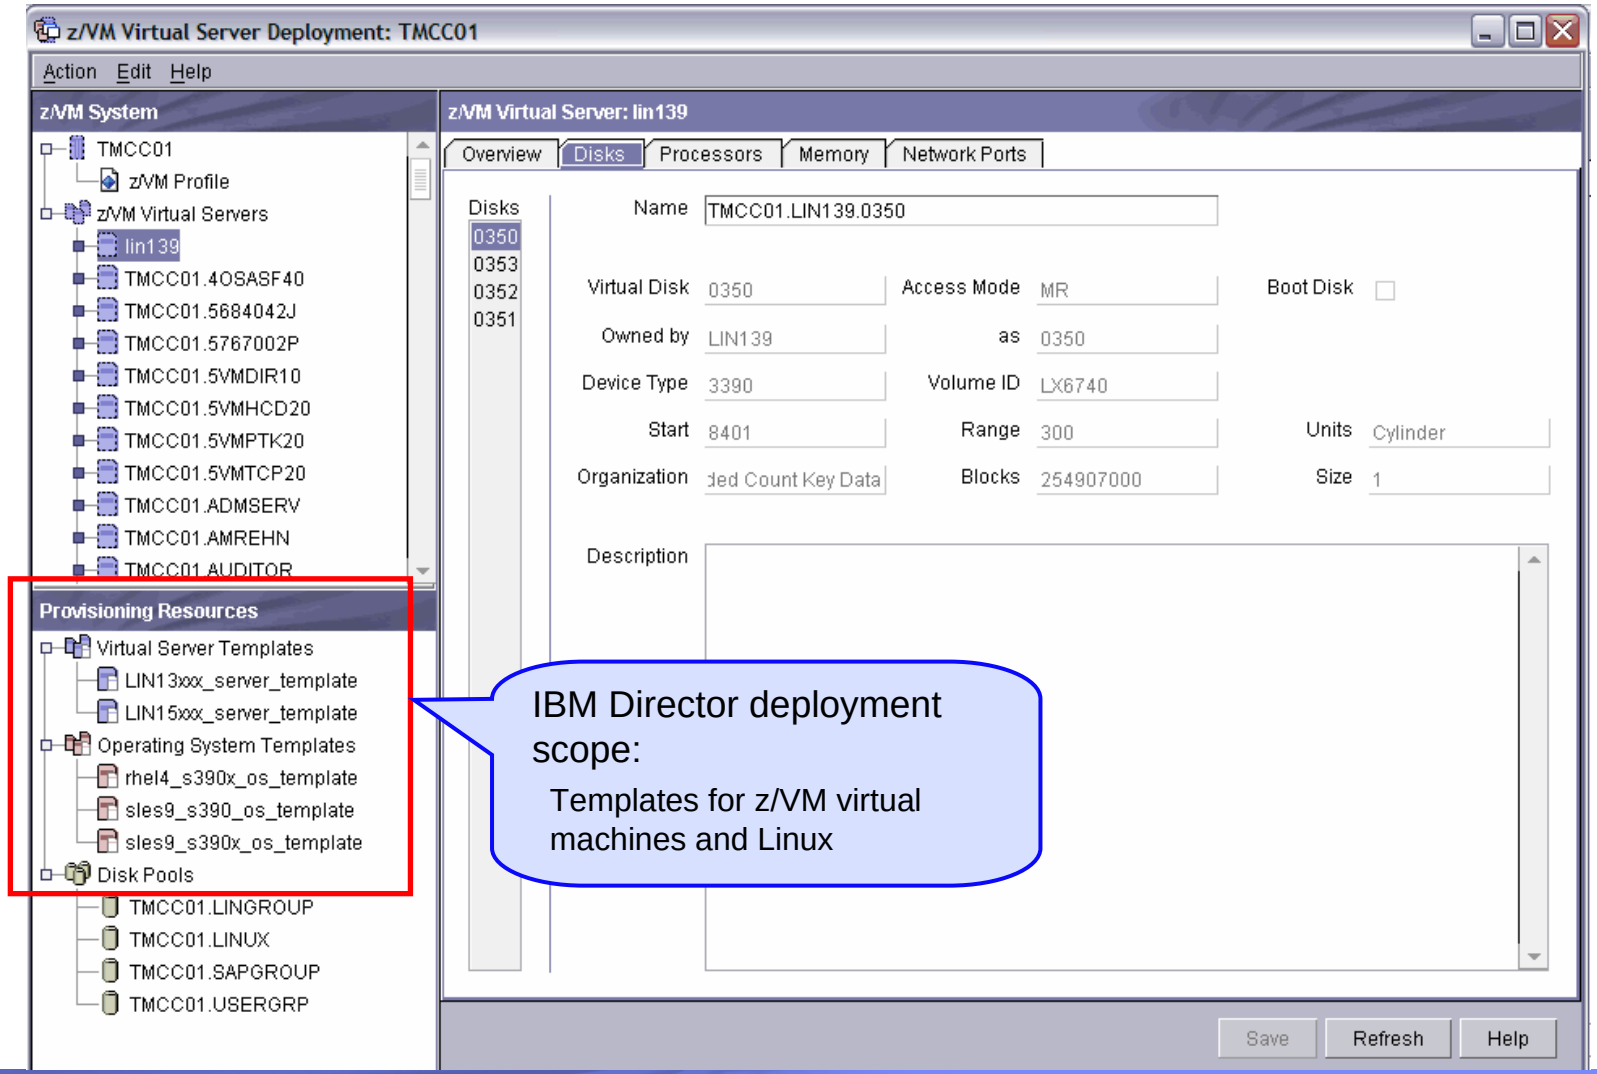

IBM Systems Director

End-to-End Management Approach

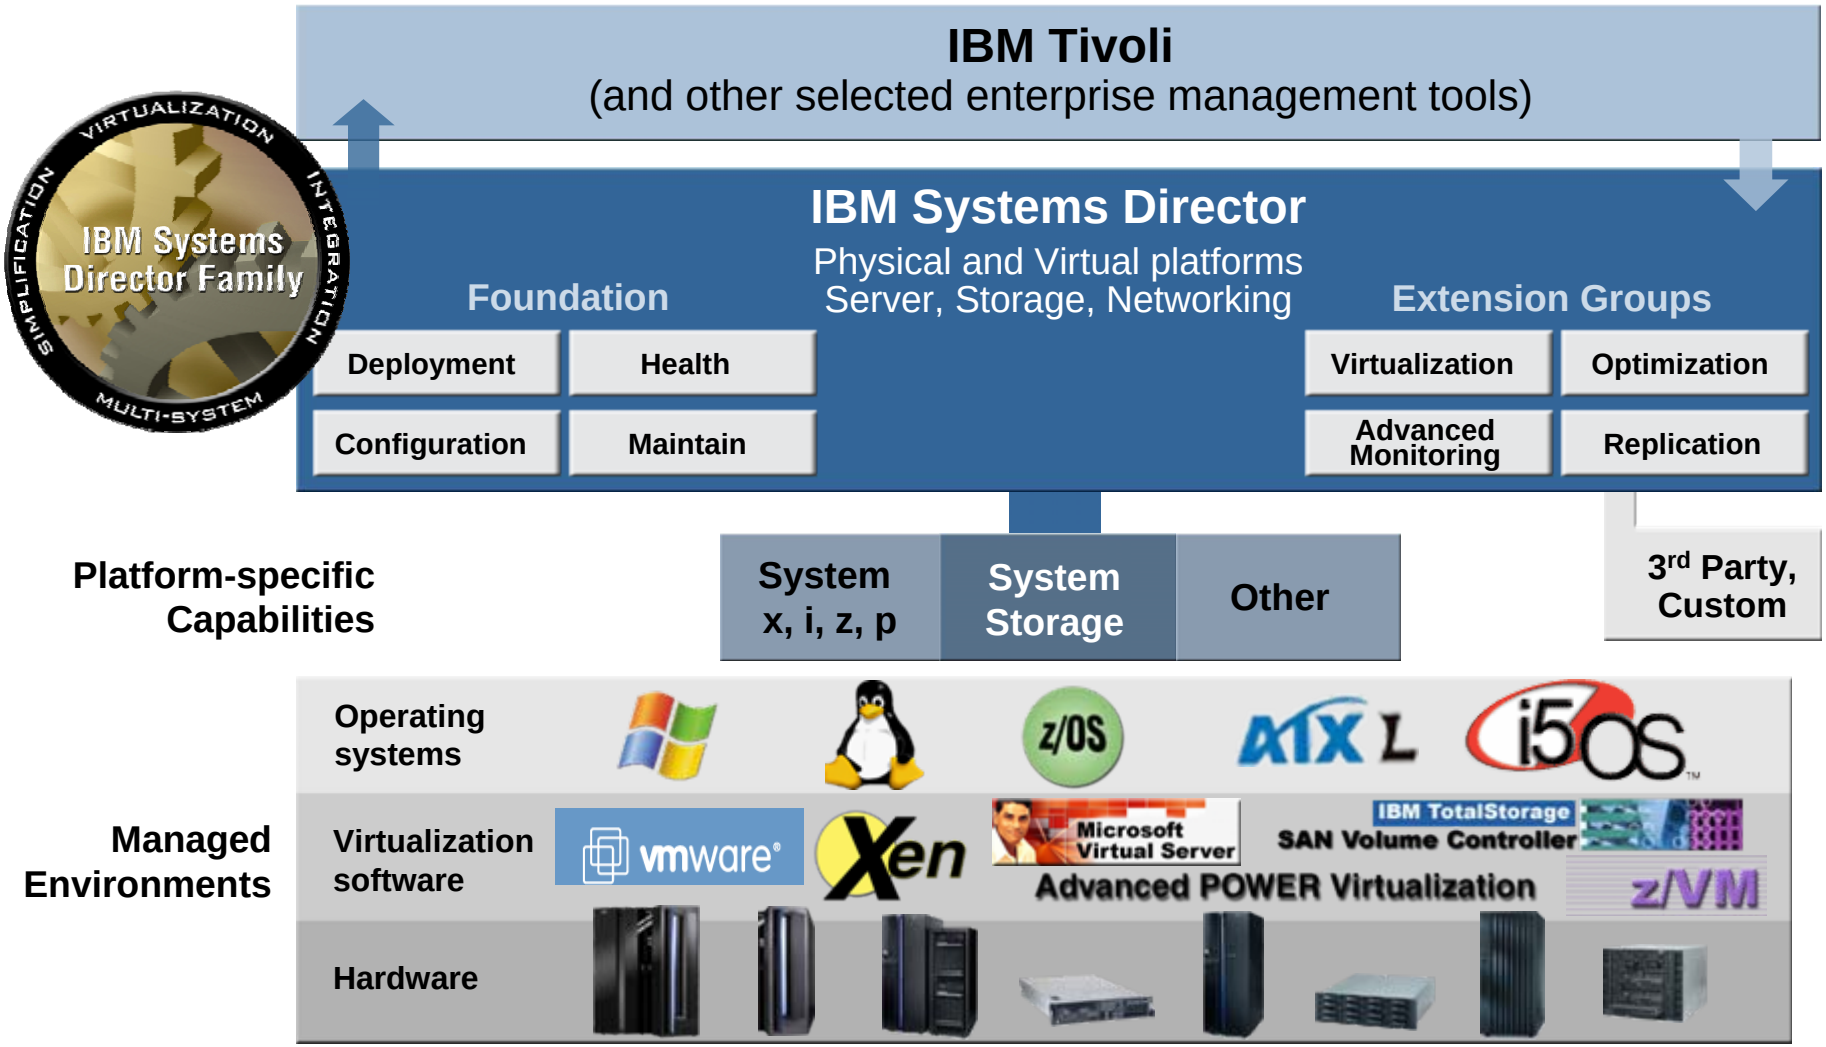

Provisioning Software in System z Virtual Linux Servers

Using IBM Tivoli Provisioning Manager

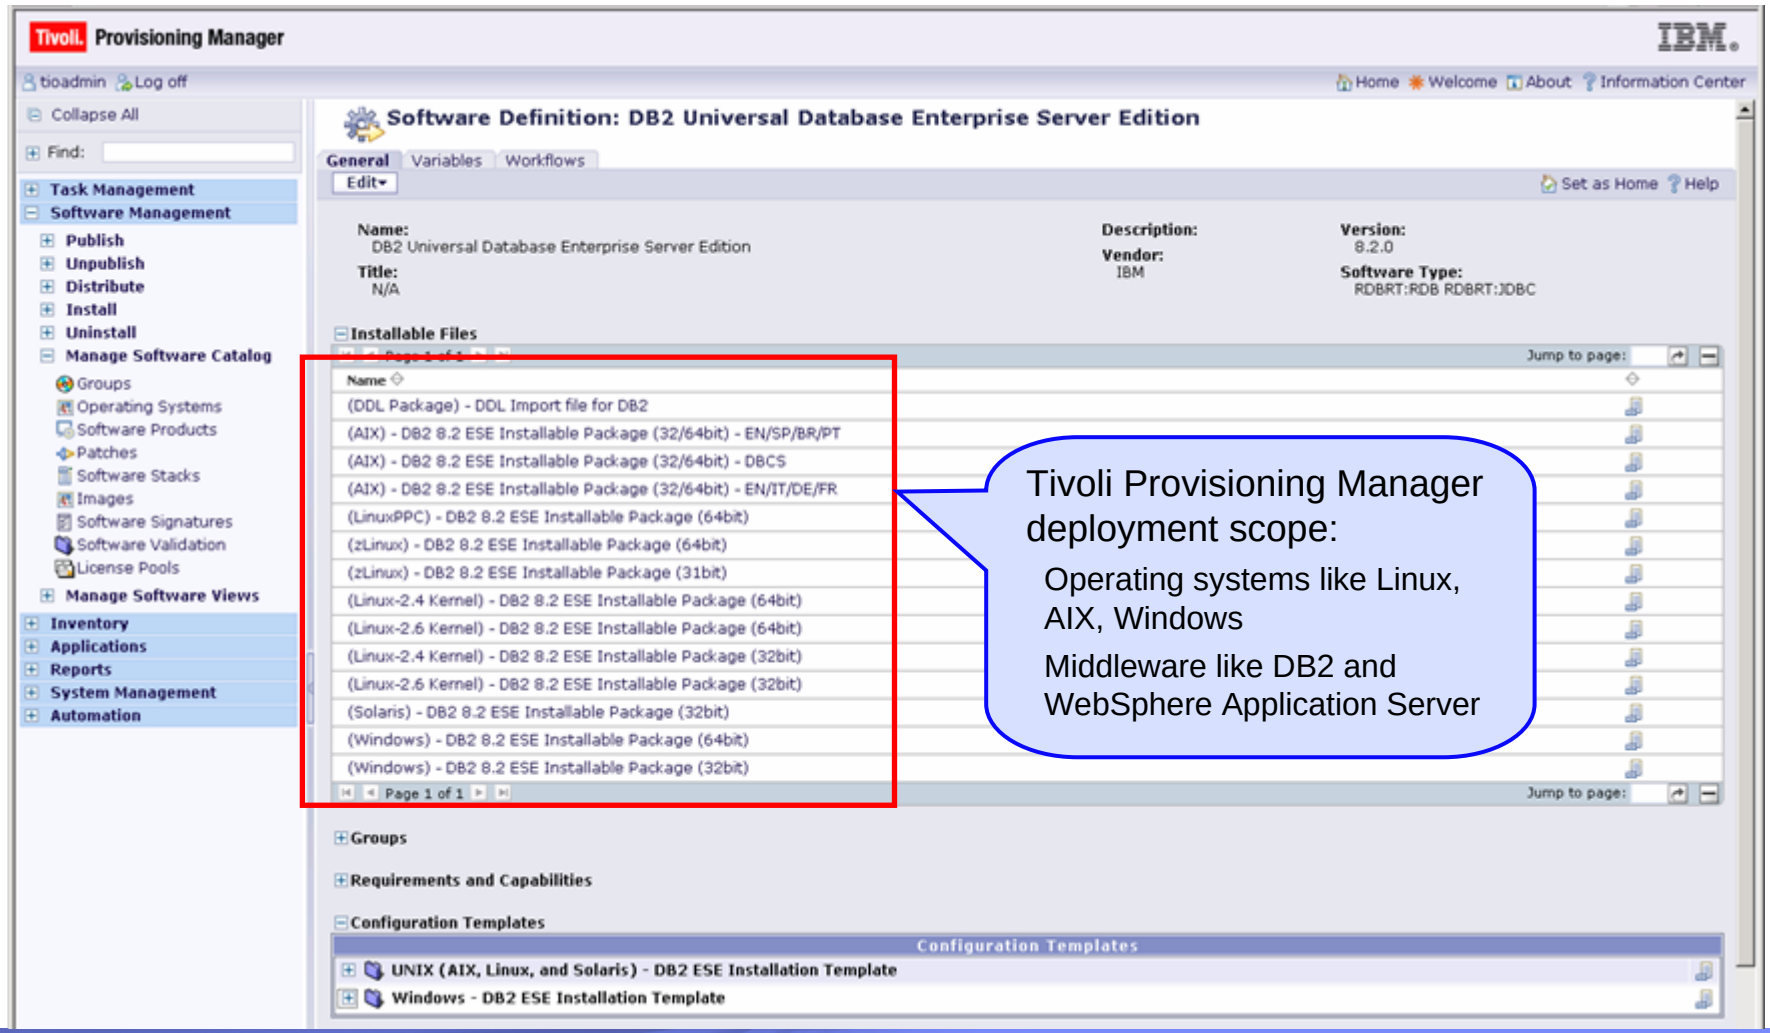

Monitoring System z Virtual Linux Servers

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

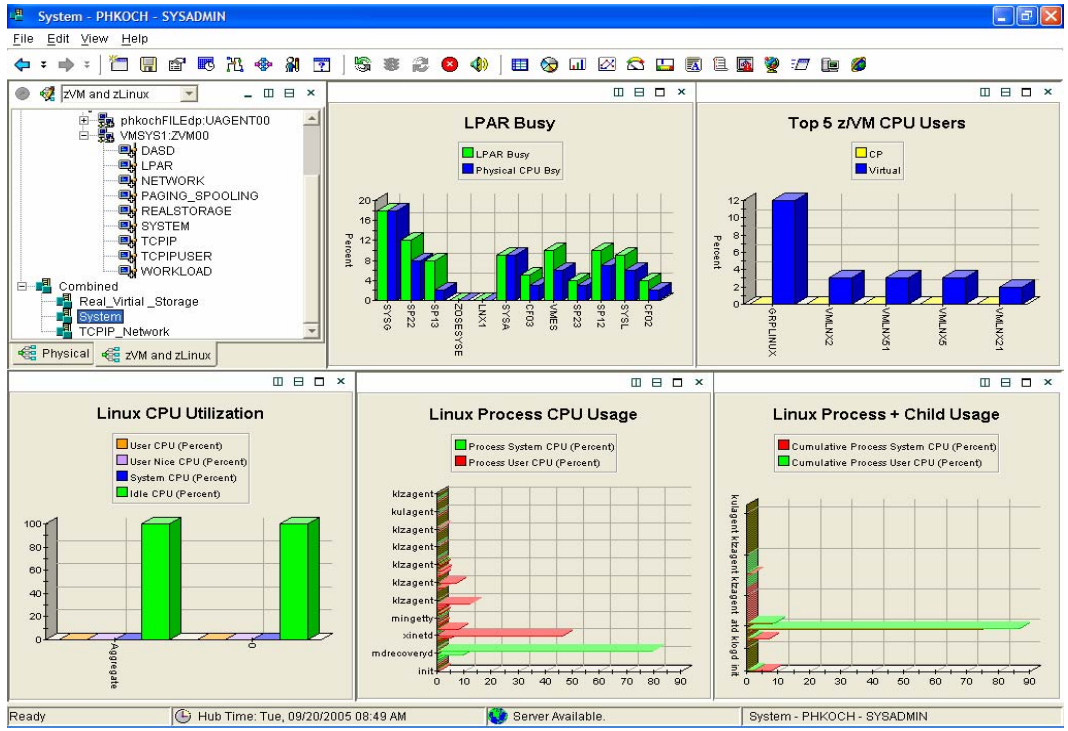

Learn more at: **ibm.com**/software/tivoli/products/omegamon-xe-zvm-linux

IBM Tivoli Virtualization Management for System z

Helping Clients Manage and Control Their Virtualized IT Infrastructure

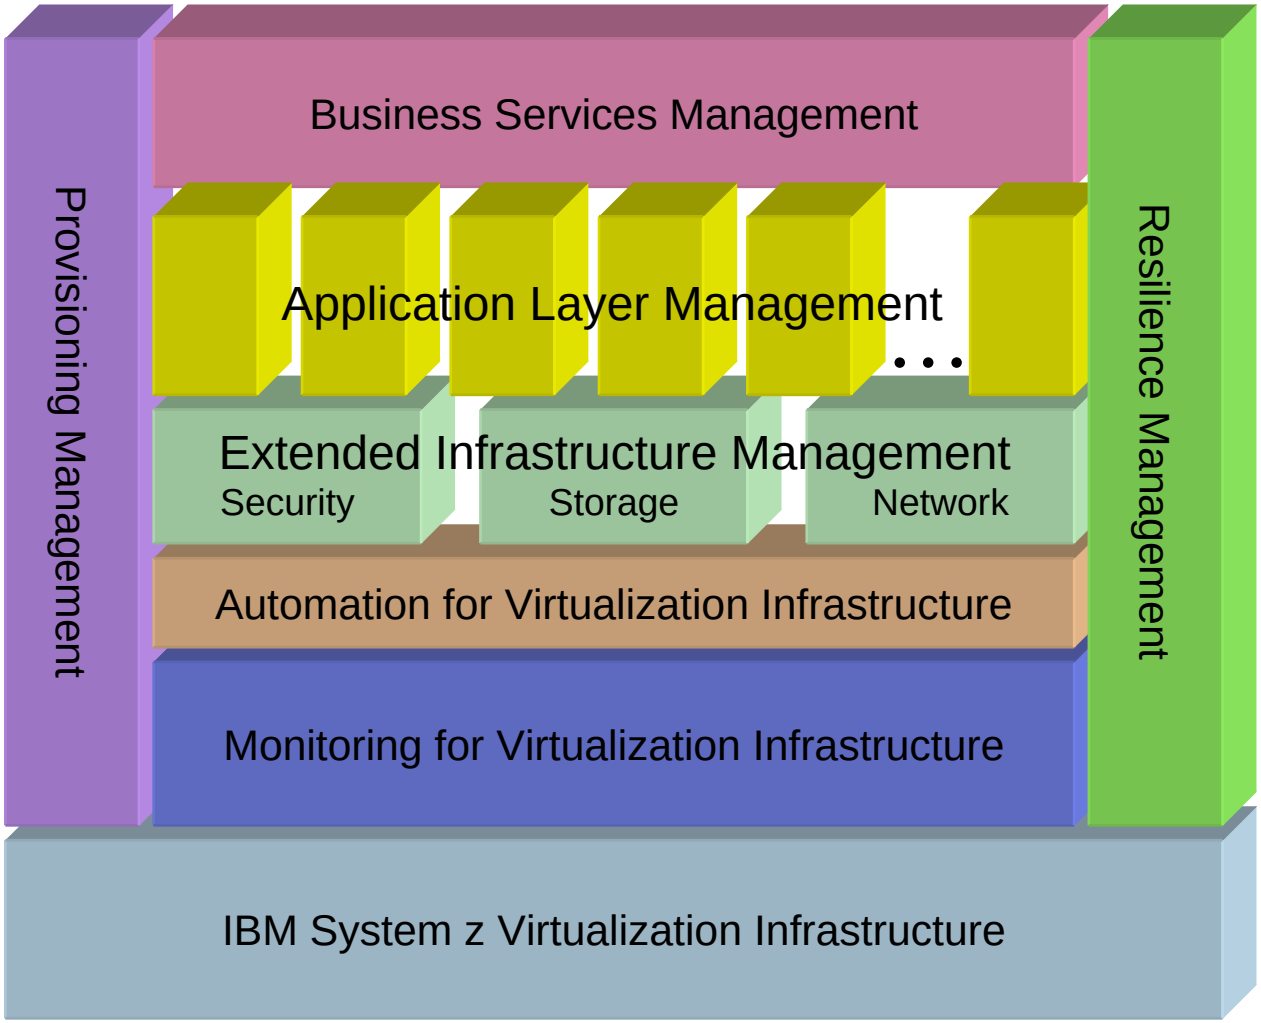

IBM System z Virtualization Infrastructure
- IBM System z hardware (including LPAR hypervisor)
- IBM z/VM Version 5

Monitoring for Virtualization Infrastructure
- z/VM Virtual Machine Resource Manager (included with z/VM)
- IBM z/VM Performance Toolkit for VM (z/VM priced feature)
- IBM Director
- IBM Tivoli OMEGAMON XE on z/VM and Linux
- IBM Tivoli Monitoring
- IBM Tivoli Composite Application Manager for SOA
- IBM Tivoli Usage and Accounting Manager

Automation for Virtualization Infrastructure
- IBM Operations Manager for z/VM
- IBM Tivoli Netcool OMNIbus
- IBM Tivoli Workload Scheduler

Provisioning Management
- IBM z/VM DirMaint (z/VM priced feature)
- z/VM Center task of IBM Director
- IBM Tivoli Provisioning Manager

Resiliency Management
- IBM Tivoli System Automation for Multiplatforms

Application Layer Management
- IBM Tivoli Application Dependency Discovery Manager
- IBM Tivoli OMEGAMON XE for Messaging
- IBM Tivoli Composite Application Manager for Response Time
- IBM Tivoli Composite Application Manager for Web Resources
- IBM Tivoli Composite Application Manager for Transactions
- IBM Tivoli License Compliance Manager

Extended Infrastructure Management *(Security)*
- IBM z/VM RACF Security Server (z/VM priced feature)
- IBM Tivoli zSecure
- IBM Tivoli Access Manager for e-business
- IBM Tivoli Access Manager for OS
- IBM Tivoli Federated Identity Manager
- IBM Tivoli Identity Manager
- IBM Directory Server
- IBM Directory Integrator

Extended Infrastructure Management *(Storage)*
- IBM SAN Volume Controller (SVC)
- IBM Tivoli Storage Manager
- IBM TotalStorage Productivity Center
- IBM Backup and Restore Manager for z/VM
- IBM Tape Manager for z/VM
- IBM Archive Manager for z/VM

Extended Infrastructure Management *(Network)*
- IBM z/VM RSCS (z/VM priced feature)
- IBM Tivoli Network Manager IP Edition

Business Services Management
- IBM Tivoli Business Service Manager
- IBM Tivoli Service Request Manager
- IBM Change and Configuration Management Database (CCMDB)

For specific releases, refer to Tivoli Platform Support Matrix at: **ibm.com**/software/sysmgmt/products/support/Tivoli_Supported_Platforms.html

IBM Tivoli Service Management Center for System z

Enabling clients to strategically use their System z as an integrated, enterprise-wide hub for the efficient management of business and IT services

New on System z

IBM Tivoli Service Management Center for System z

- Best Practices and Services
- Process Mgt
- Service Mgt Platform
- Operational Mgt
- Optimized Infrastructure

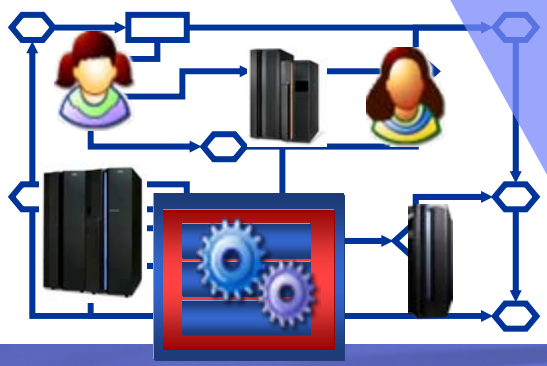

- *IBM Tivoli Service Request Manager*
- *IBM Tivoli Change & Release Management*
- *IBM Tivoli Business Continuity Process Manager*
- *IBM Tivoli Business Service Manager*
- *IBM Tivoli Application Discovery and Dependency Manager (TADDM)*
- *IBM Tivoli Change and Configuration Management Database (CCMDB)*
- *IBM Tivoli Service Request Manager*
- *IBM Tivoli Netview for z/OS, OMEGAMON, IBM Tivoli Composite Application Manager, DFSMS*
- *IBM Tivoli System Automation & IBM Tivoli Workload Automation*
- *IBM Tivoli Usage and Accounting Manager, Tivoli Decision Support for z/OS*
- *IBM Tivoli Identify Manager, IBM Tivoli Access Manager, zSecure, Security Information and Event Mgmt, RACF*

z/VM Systems Management Products from IBM

- **IBM Operations Manager for z/VM**
  - Helps improve the monitoring and management of z/VM virtual machines by *automating* routine maintenance tasks
  - Enables users to *automatically respond* to predictable situations that require intervention
  - Assists with monitoring and problem determination by allowing authorized users to view and interact with *live consoles* of z/VM service machines or Linux guests
- **IBM Backup and Restore Manager for z/VM**
  - Provides z/VM system administrators and operators the ability to efficiently and effectively backup and restore files and data on *z/VM systems*
  - Can also backup and restore images of *non-z/VM guest systems* such as Linux
- **IBM Tape Manager for z/VM**
  - Manages and monitors tape resources; helps increase data availability and improve *operator efficiency*
  - *Automates* common daily tape operations and helps eliminate tedious, often error-prone, manual tasks
- **IBM Archive Manager for z/VM**
  - Addresses storage and data management concerns by allowing users to archive historical or other infrequently used data to *increase data availability*
  - Helps companies *comply* with data storage requirements mandated by fiscal or legal regulations and policies

GDPS/PPRC Multiplatform Resiliency for System z

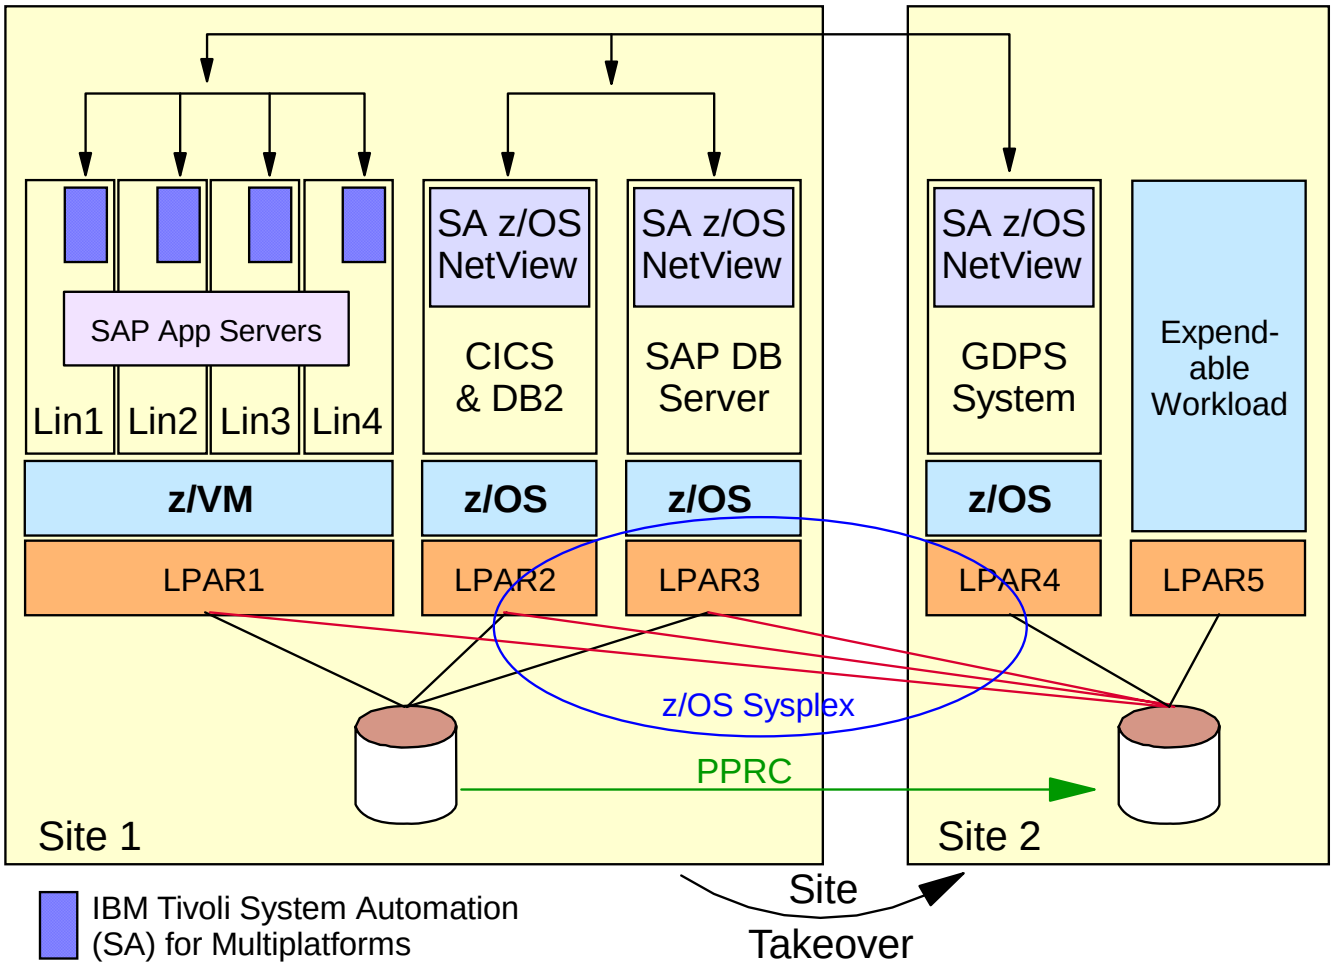

- Designed for customers with distributed applications
- SAP application server running on Linux for System z
- SAP DB server running on z/OS
- Coordinated near-continuous availability and DR solution for z/OS, Linux guests, and z/VM
- Uses z/VM HyperSwap function to switch to secondary disks
- Sysplex support allows for site recovery

z/VM Virtualization Leadership

The Value of Scaling on a Single Hypervisor

- **Grow virtual server workloads without linearly growing energy costs**
- **Enhance staff productivity with a single point of control at the hypervisor level**
- **Dynamically add and remove physical resources in a single machine to optimize business results**
- **Exploit hypervisor automation tools with higher degrees of integration and optimization**

System z Virtualization Leadership

Offering Virtual Server Solutions the IT Industry Demands

- **Highly scalable, granular, and efficient virtual server hosting**
  - Capable of running thousands of virtual servers on a single mainframe
  - Designed to run memory-rich and I/O-intensive (disk and network) workloads with data integrity
  - Able to achieve extremely high levels of physical CPU, memory, networking, and disk resource sharing
  - Allows significant over commitment of real resources, resulting in higher utilization while processing peak business demands and maintaining service levels – “doing more with less”
- **Infrastructure simplification and flexible operations**
  - Can improve the efficiency of your IT staff with robust and powerful systems management capabilities, allowing staff to quickly provision and manage more virtual servers
  - Provides non-disruptively adding and removing of physical resources (CPU, memory, I/O, networking) to satisfy virtual server requirements in response to changing business demands
  - Can host Linux applications side-by-side LPARs on the same mainframe with fast and secure connectivity, leveraging z/TPF, z/VSE, and z/OS secure data serving
- **Virtual server integrity and security**
  - For decades z/VM and the mainframe have been architected for secure processing, offering high levels of integrity and security
  - System z servers have achieved EAL 5 certification; z/VM has achieved EAL 4+ certification, offering system solutions that have been methodically designed, tested, and reviewed for secure operations

Functional Comparison of z/VM and VMware ESX

Attribute	z/VM V5.4	VMware ESX 3.5	System z Value
Supported operating systems	Linux, z/OS, z/VSE, z/TPF, z/VM itself	Linux, Windows, Netware, Solaris 10	z/VM-on-z/VM = added flexibility
Scalability and Performance			
Hypervisor scalability	Up to 32 CPUs, 256 GB of memory, More than 1 TB of active virtual memory	Up to 32 CPUs, 256 GB of memory	Cost-saving, extreme scalability of virtual server environment
Virtual Machine (VM) scalability	Up to 64 CPUs, 1 TB of memory, extensive I/O bandwidth	Up to 4 CPUs, 64 GB of memory, modest I/O bandwidth	Virtualizes servers on z/VM that cannot run on VMware
CPU sharing	No limit	Up to 8 VMs per CPU	Add servers without adding HW
Architected (practical) VM limit	Thousands (hundreds) per copy of z/VM	128 (singles) per copy of VMware	Avoid real server sprawl
CPU/memory capacity on demand	Yes, non-disruptively	No	Fast, easy capacity growth
In-memory support	Minidisk cache; Virtual Disks in Storage; DCSS (shared program executables)	Shared virtual memory pages (detected via background operation)	Enhanced resource utilization
Logical Partition (LPAR) support	Yes	No	Secure Linux access to z/OS
Flexible Operations			
Resource over-commitment support (memory, CPU, network, I/O)	Extensive	Modest	Absorb workload spikes; add more servers to a “full” system
Reconfiguration of Virtual Machines	Non-disruptive re-config for CPU, I/O, networking, and memory	VM reboot required for re-config of CPU, memory, ethernet, disk	Higher server and application availability; staff productivity
Command and control, monitoring, automation infrastructure	Extensive, robust, time-tested	Modest, yet easy to use	Cost-optimized systems management support
Virtual Machine mobility support	No; single-image scalability of z/VM does not require mobility for mgmt	Yes; essential for workload mgmt across multiple copies of VMware	Can dynamically add or remove resources to meet demand
Integrity and Security			
Fault isolation / hypervisor security	Hardware-assisted isolation*; EAL 4+ (CAPP/LSPP)	No I/O virtualization separation; EAL 4+ (No protection profile)	Helps to avoid security breaches; data security and integrity
Run multiple copies of hypervisor on single server	Yes; share CPU, I/O, and networking resources among z/VM systems	No	Workload isolation; lower-cost failover (using same hardware)

* z/VM runs in System z LPARs, which have achieved EAL 5 certification; System z HiperSockets provide high-speed, secure connectivity among LPARs.

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"?

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 data and applications
- When innovation is stifled because your staff cannot experiment or develop new solutions using existing resources

Extreme Virtualization with System z
Understanding the Value Proposition

- **Business pain points addressed by server virtualization:**
  - Underutilized IT assets
  - Environmental costs
  - Linear software costs per server image
  - Staff inefficiencies managing multiple real servers
  - Spiraling people costs
- **x86 virtualization pain points addressed by System z**
  - Virtual server workload management
  - Reliable high-bandwidth I/O virtualization
  - Virtual server and total system performance reporting and planning
  - Virtual server reconfiguration outages
  - Virtual machine security and integrity
  - Server sprawl with added complexity

"Scale-out distributed servers are like the 'crack-cocaine of the IT systems business', highly addictive and extremely damaging." – Ian Bramley, Software Strategies, "Mass Distributed Server Consolidation" White Paper

Extreme Virtualization with System z
Opportunities for Cost Savings

- **Energy and floor space savings**
- **Reduced software license fees via CPU over-commitment**
- **Enhanced staff productivity via large-scale virtual server deployment on a single z/VM hypervisor**
- **Reduced application outages – reliability and redundancy of System z infrastructure**
- **Flexible configuration options for business continuance (e.g., Capacity Backup on Demand)**
- **Cost-attractive economic model for technology refreshes (e.g., specialty engines carry forward to next generation)**

Questions

Backup Material

z/VM on System z Virtualization

Technology Leadership – Maximizing Hardware Asset Value

- **Resource sharing**
  - All hardware assets can be virtualized for all types of workloads
  - Very high levels of resource utilization
  - Highly granular (efficient) sharing of resources
  - High levels of resource over-commitment
  - Relatively low levels of hypervisor overhead in a highly shared resource environment
- **Scalability**
  - Individual virtual machines can scale to considerable size
  - Overall hosted environment can scale to considerable size
- **Flexible, cost-efficient growth and technology refresh**
  - Hardware assets can be added to z/VM LPAR quickly and non-disruptively
  - Added resources are available for new *and/or* existing virtual servers
  - Hosted virtual machines can be re-configured quickly *and* non-disruptively
  - Relatively fast migration of large virtual server environment to new mainframe systems
- **Energy and floor space efficiency**
  - Low-cost, large-scale virtual server hosting with z/VM on System z
  - Very dense virtual server packaging without corresponding energy / cooling expenses

z/VM on System z Virtualization
Technology Leadership – Operational Advantages

- **Command and Control**
  - Extensive, robust, proven infrastructure for automated operations
  - Single point of control at hypervisor level for managing a very large number of virtual servers
  - Performance reporting, monitoring, and capacity planning for individual virtual machines and the z/VM LPAR itself
  - Fast virtual machine provisioning
  - Highly granular controls for workload management
  - Enhanced staff productivity
- **Security**
  - Hardware-enforced system integrity (e.g., SIE instruction, storage protection keys)
  - RACF authorization and authentication for z/VM hypervisor and virtual machines
  - Integrated cryptographic hardware for guest use
- **Disaster recovery**
  - Robust heritage of failover and business continuance options
  - Capacity Backup Upgrade (CBU)
  - GDPS/PPRC Multiplatform Resiliency for System z

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

Highly Balanced Utilization and Sharing of System Resources

- **The vast majority of installed x86 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, x86 virtualization 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**

* x86 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**