Linux Optimization with IBM System z Virtualization

The future runs on System z

December 2008

Jon Nolting
jrnoltin@us.ibm.com
IBM System z IT Architect

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

- **IBM System z virtualization overview**
- **Advanced z/VM virtualization technologies**

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**

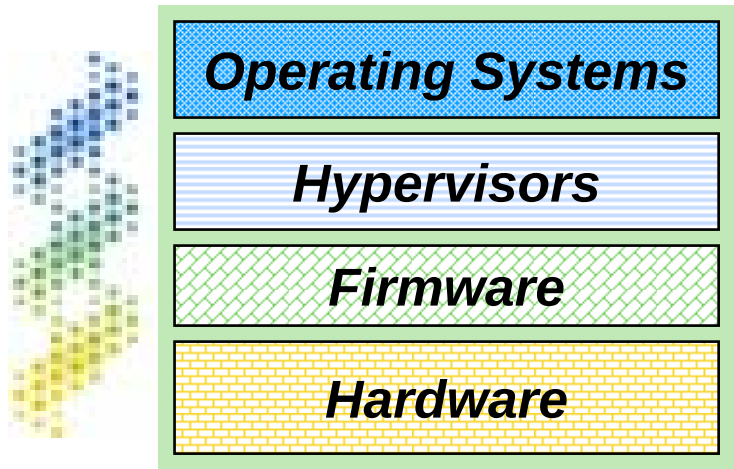

System z Virtualization Technology

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

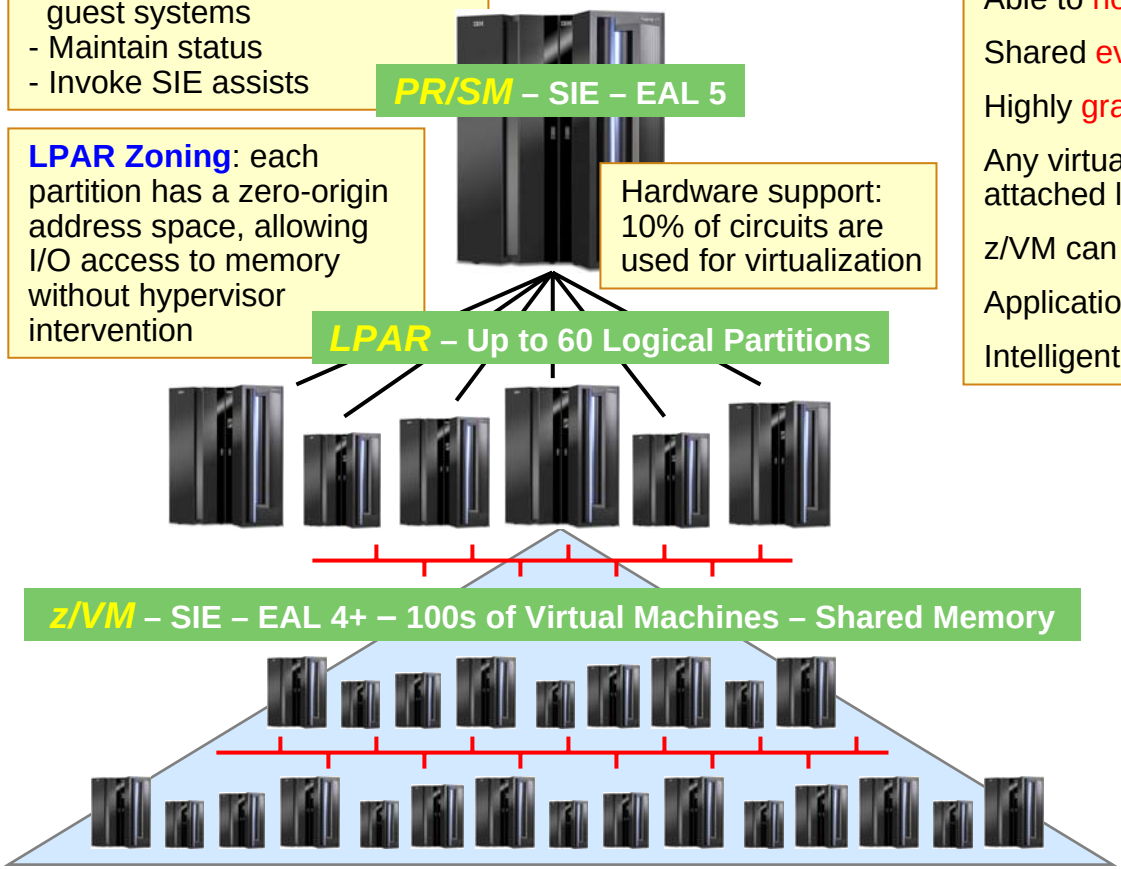

The potential performance impact of the Linux server farm is isolated from the other LPARs

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: 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

Advanced z/VM Virtualization Technologies

- **Resource sharing and scalability**
- **CPU and memory**
- **Advanced disk support**
- **Virtual communications and network consolidation**
- **Systems management, provisioning, command and control**

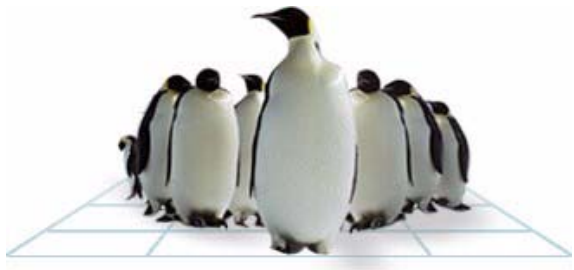

IBM System z Virtualization Leadership

Extreme Levels of CPU Sharing

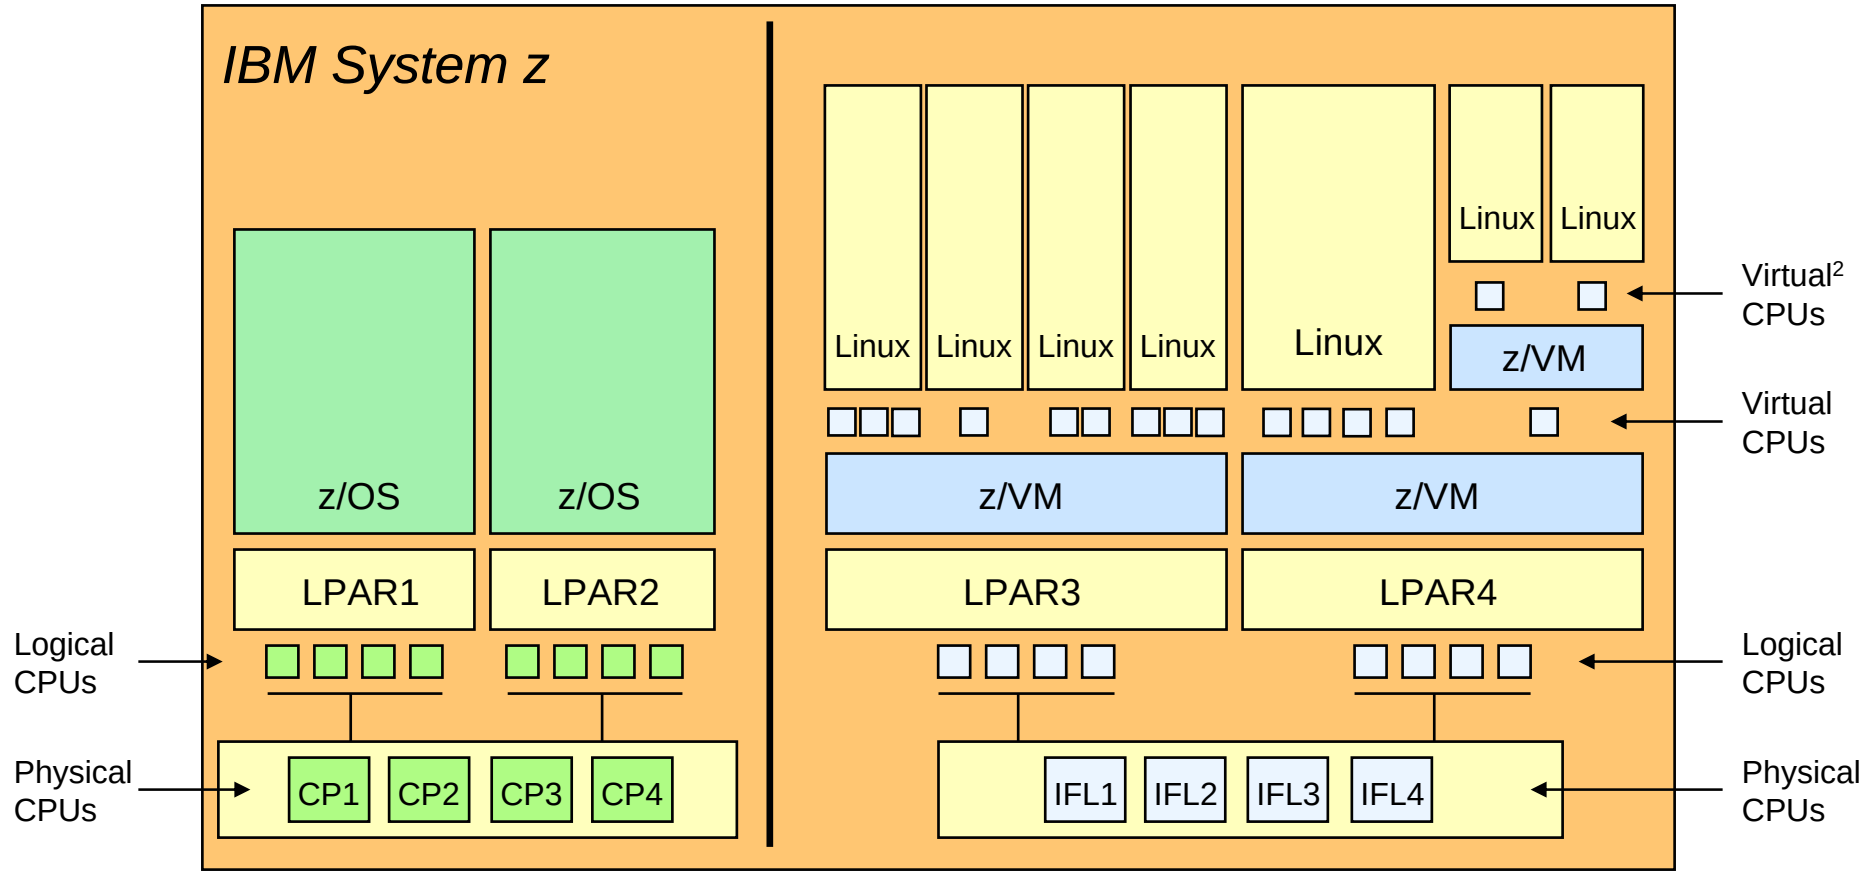

IBM System z CPU High Availability
Concurrent Processor Reassignment

- **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***

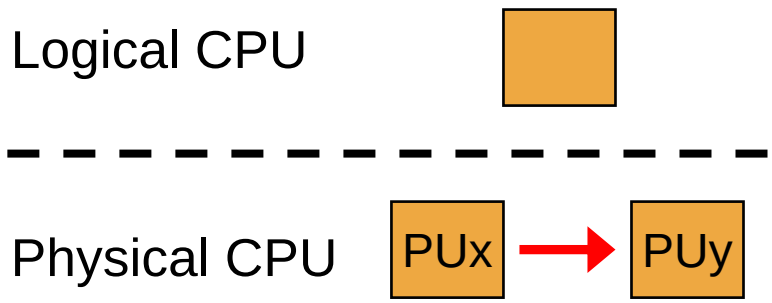

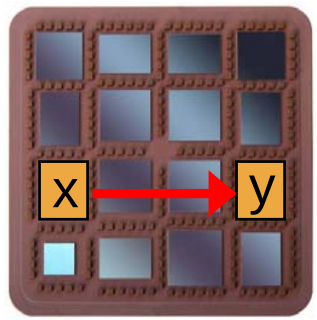

System Design Affects Virtualization Capabilities

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

- With TCO-friendly pricing

Up to 64-way SMP

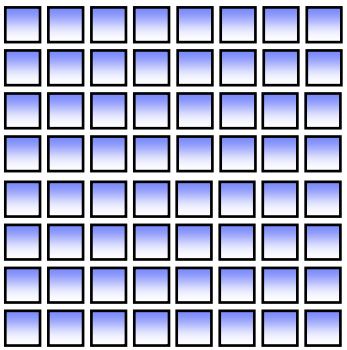

Share up to 64 processors with up to 60 LPARs

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

* No software license fees

2 Standard Spare PUs

Up to 11 System Assist Processors

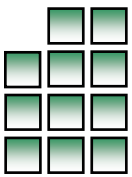

Offload system processing to dedicated CPUs (no impact to software license fees)

Up to 336 I/O Processors

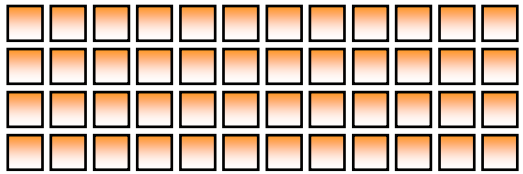

No additional charge for these processors

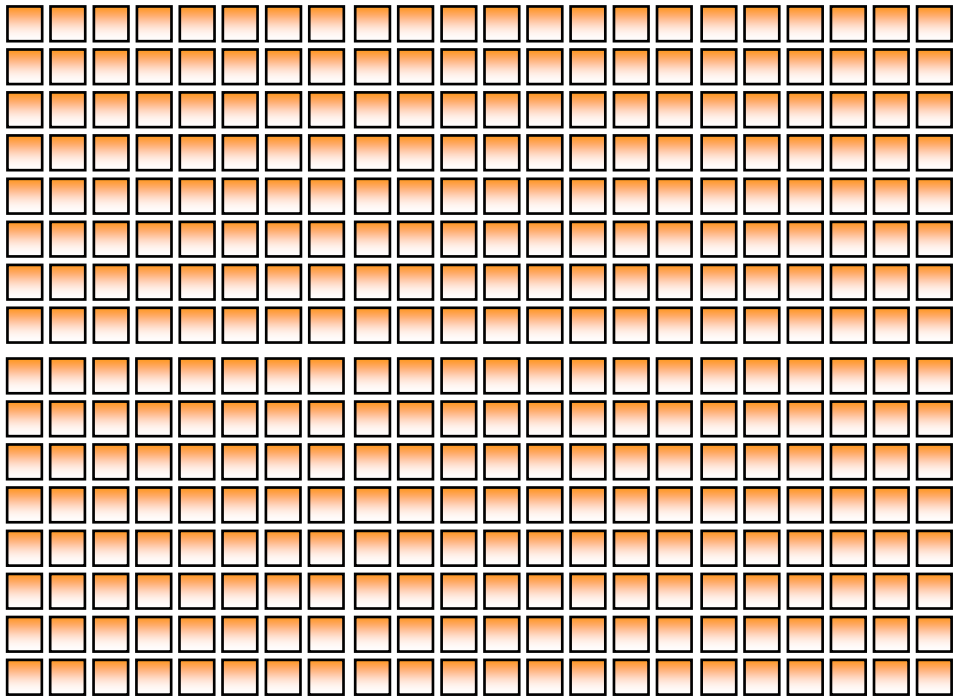

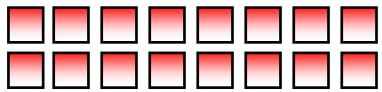

Up to 16 Crypto Express2 CPUs

High scale performance for SSL transactions

PR/SM High-Performance I/O Sharing (Multi-Image Facility)

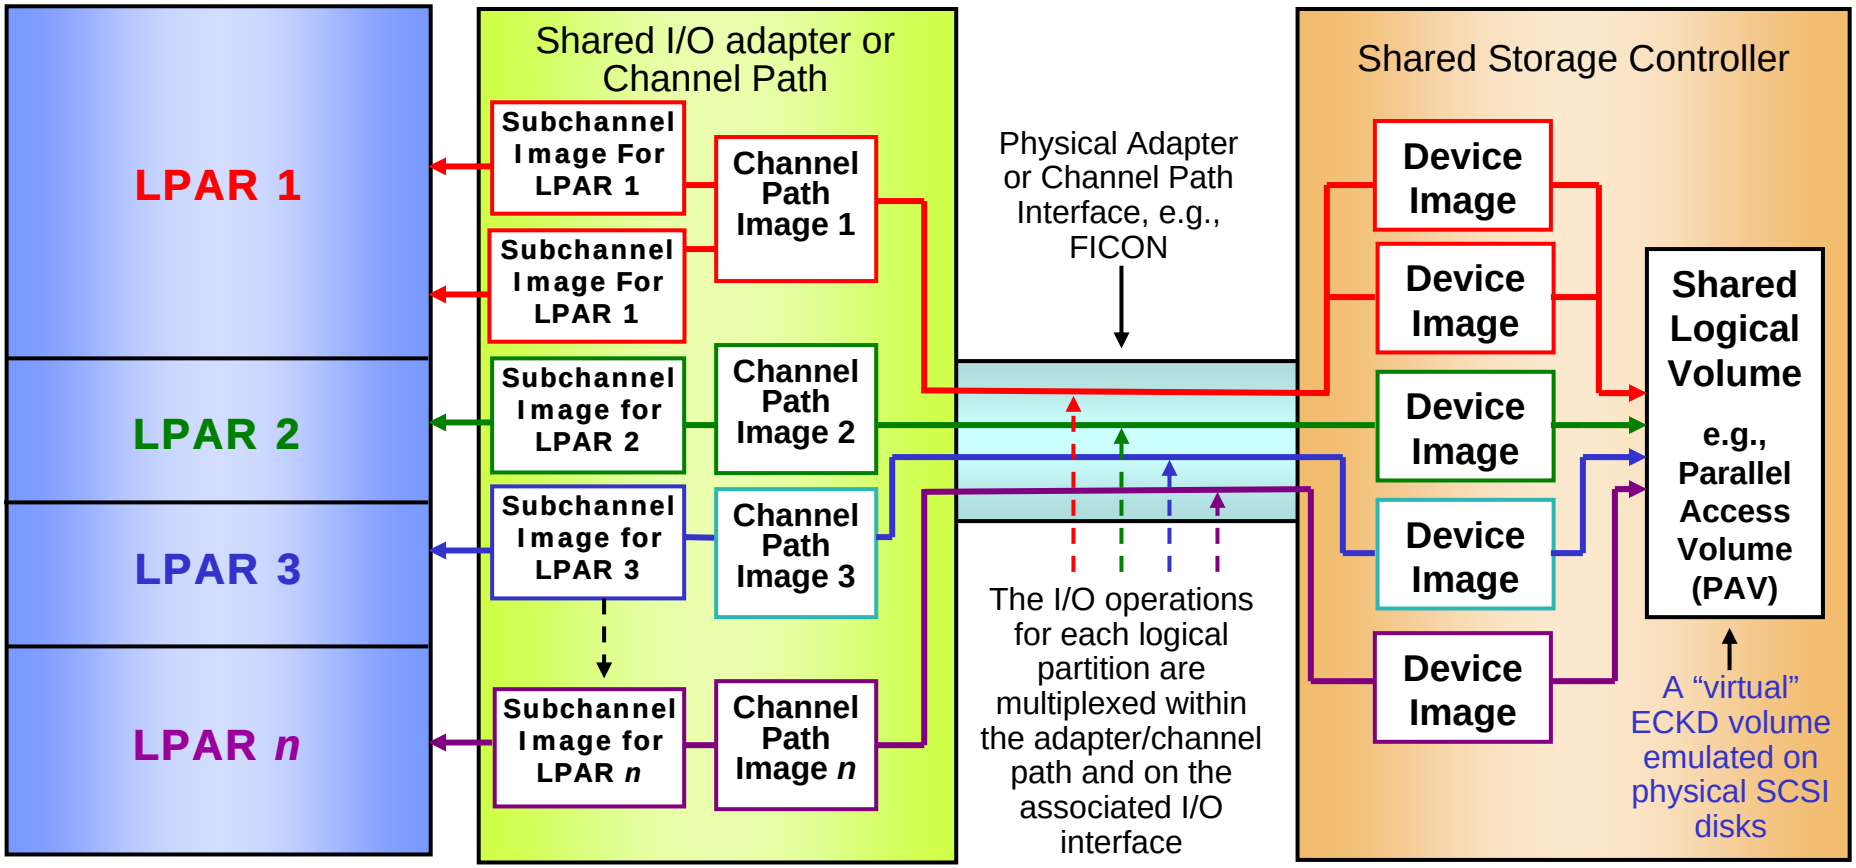

- The I/O infrastructure is shared by LPARs at native speeds, without hypervisor involvement
- Up to 8 physical channels process the I/O requests to the shared devices
  - This reduces the possibility of I/O queuing delays at the channels or at the shared storage controller

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...

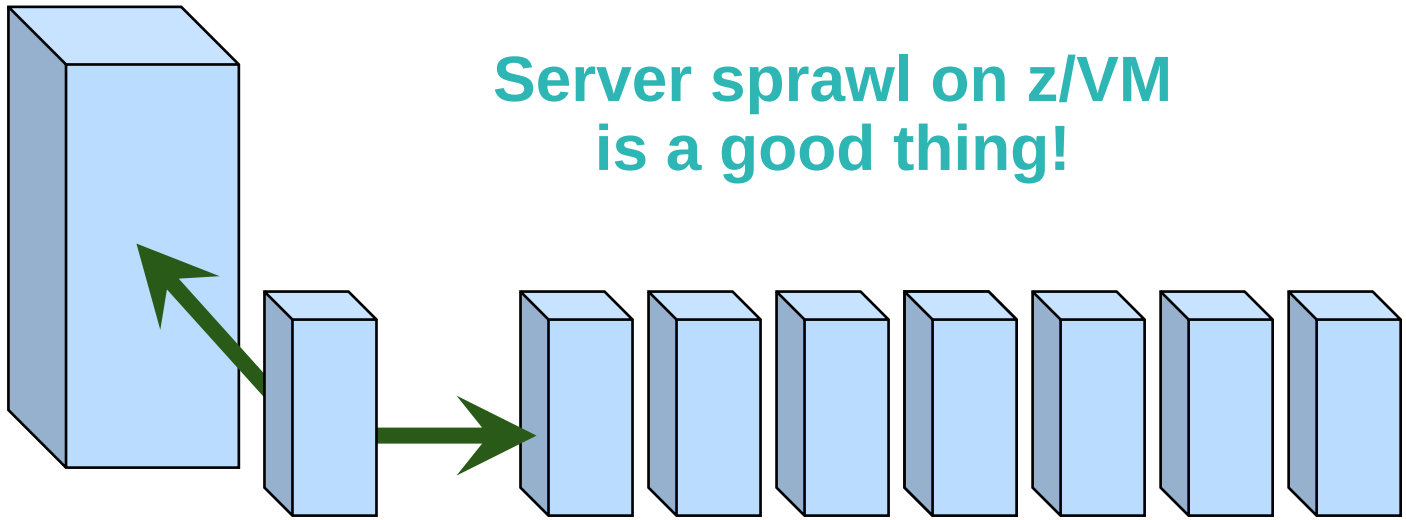

Server sprawl on z/VM is a good thing!

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

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

Additional Cost Savings and Operational Efficiencies

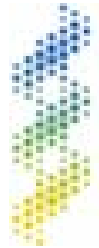

- **A fundamental strength of z/VM is its ability to share system resources to an extreme level**
- **Virtual machines can simultaneously access the I/O and networking resources available on a System z machine**
  - Both real and virtual (z/VM) resources can be used with very high levels of bandwidth 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

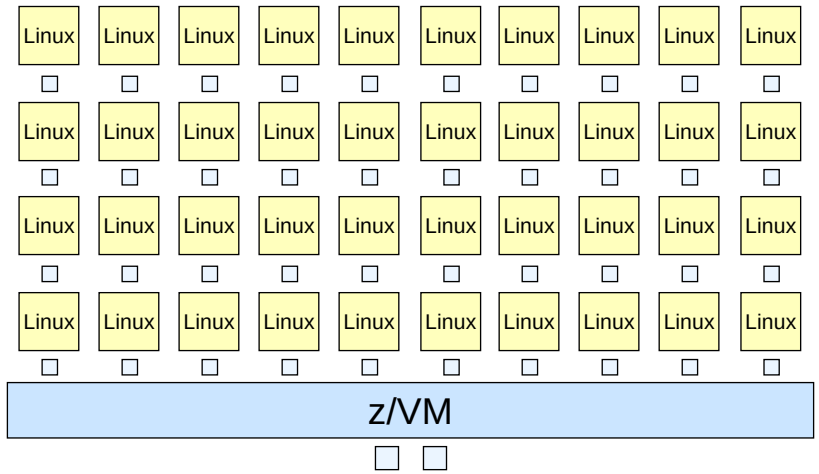

Linux-on-z/VM and *Flexible, Efficient Growth*

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**

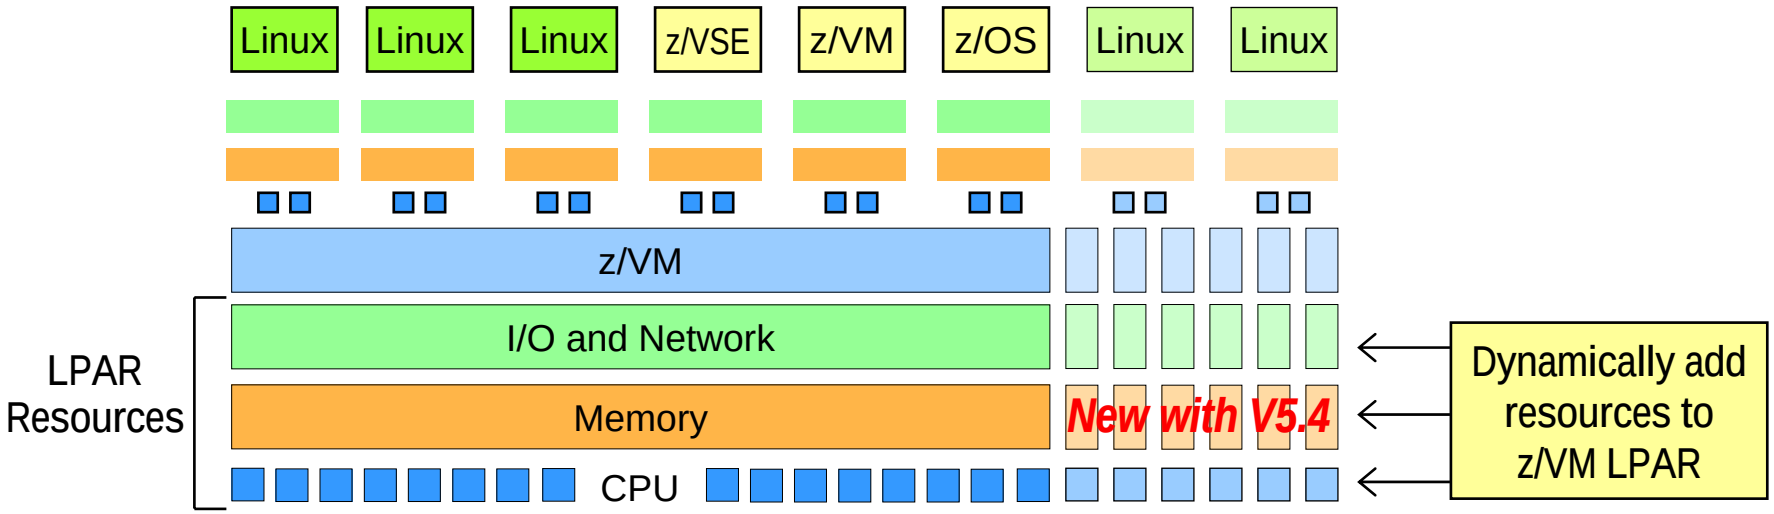

Smart economics: non-disruptively scale your z/VM environment by adding hardware assets that can be shared with *every* virtual server

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

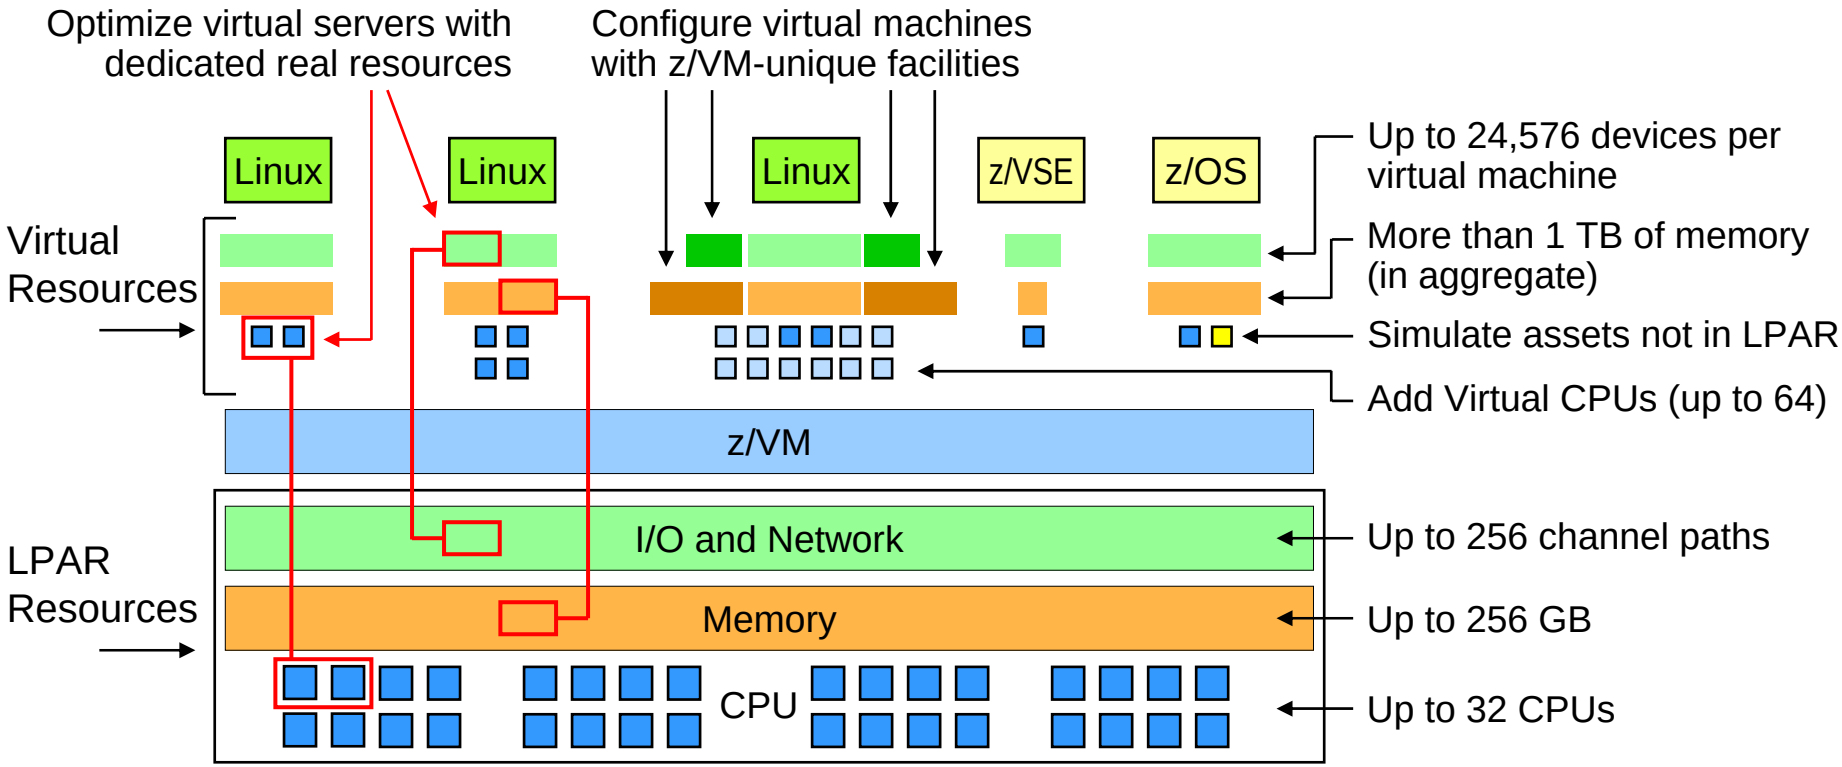

z/VM CPU Resource Controls

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"

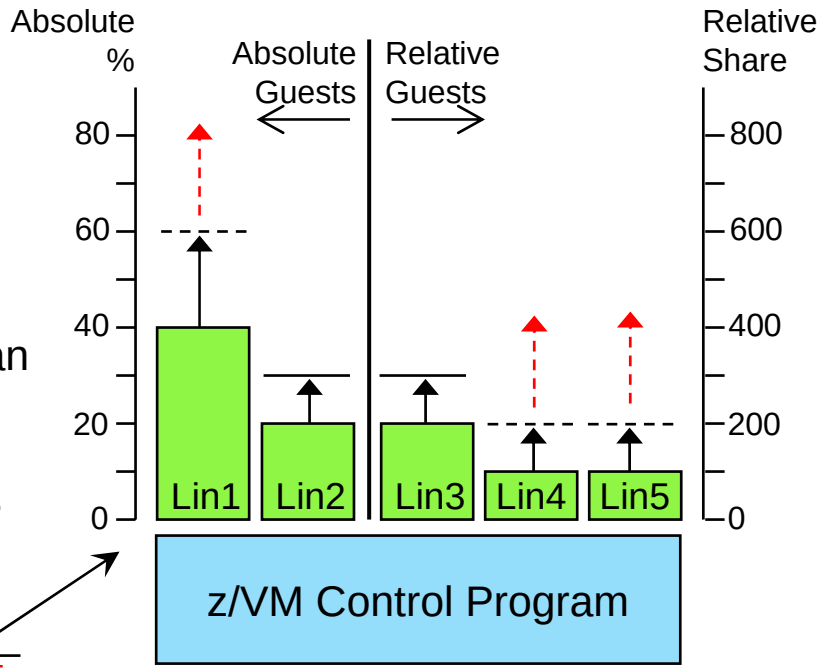


```
SHARE Lin1 ABSOLUTE 40% ABSOLUTE 60% LIMITSOFT
SHARE Lin2 ABSOLUTE 20% ABSOLUTE 30% LIMITHARD
SHARE Lin3 RELATIVE 200 RELATIVE 300 LIMITHARD
SHARE Lin4 RELATIVE 100 RELATIVE 200 LIMITSOFT
SHARE Lin5 RELATIVE 100 RELATIVE 200 LIMITSOFT
```

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

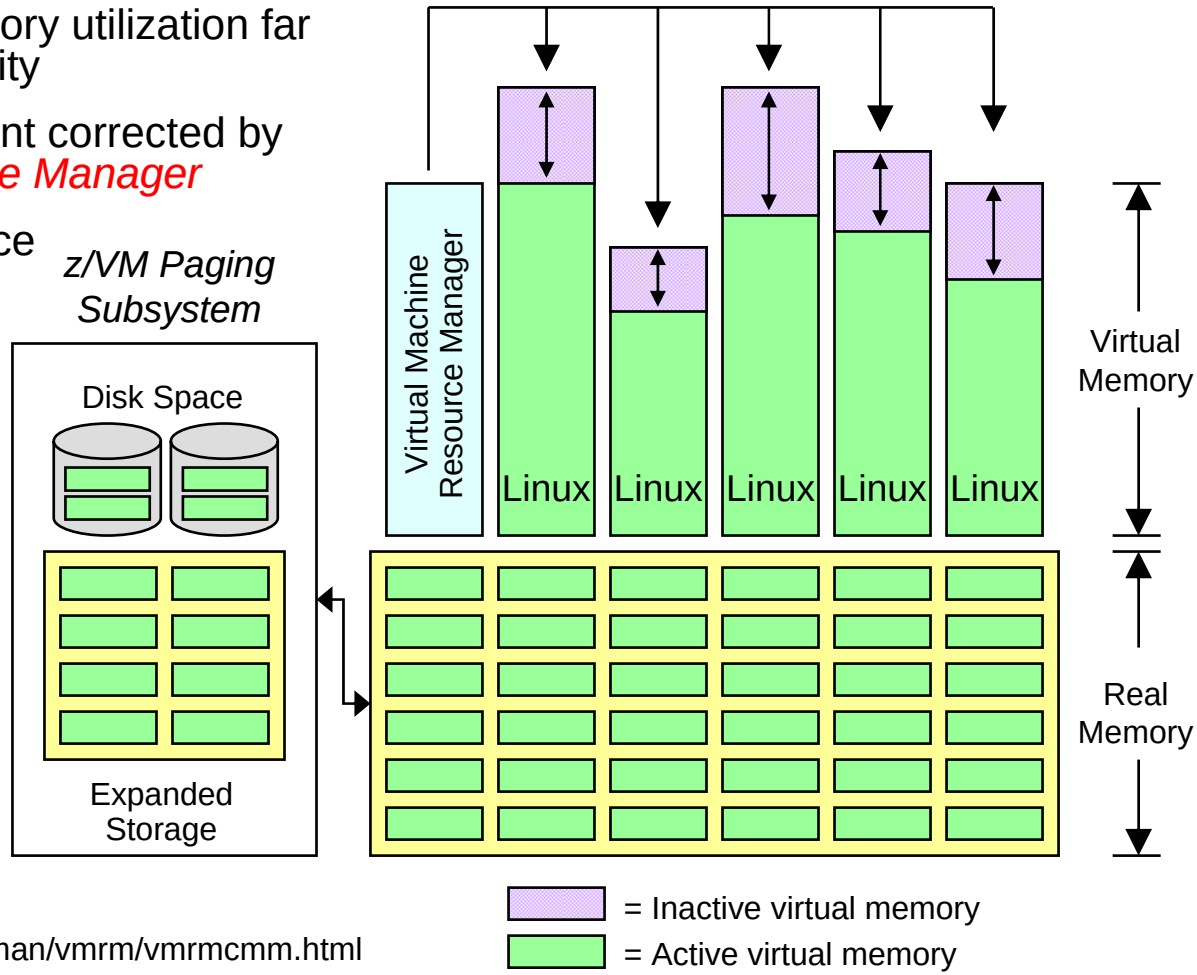

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

OLTP Database Environment with VMRM-CMM and CMMA

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

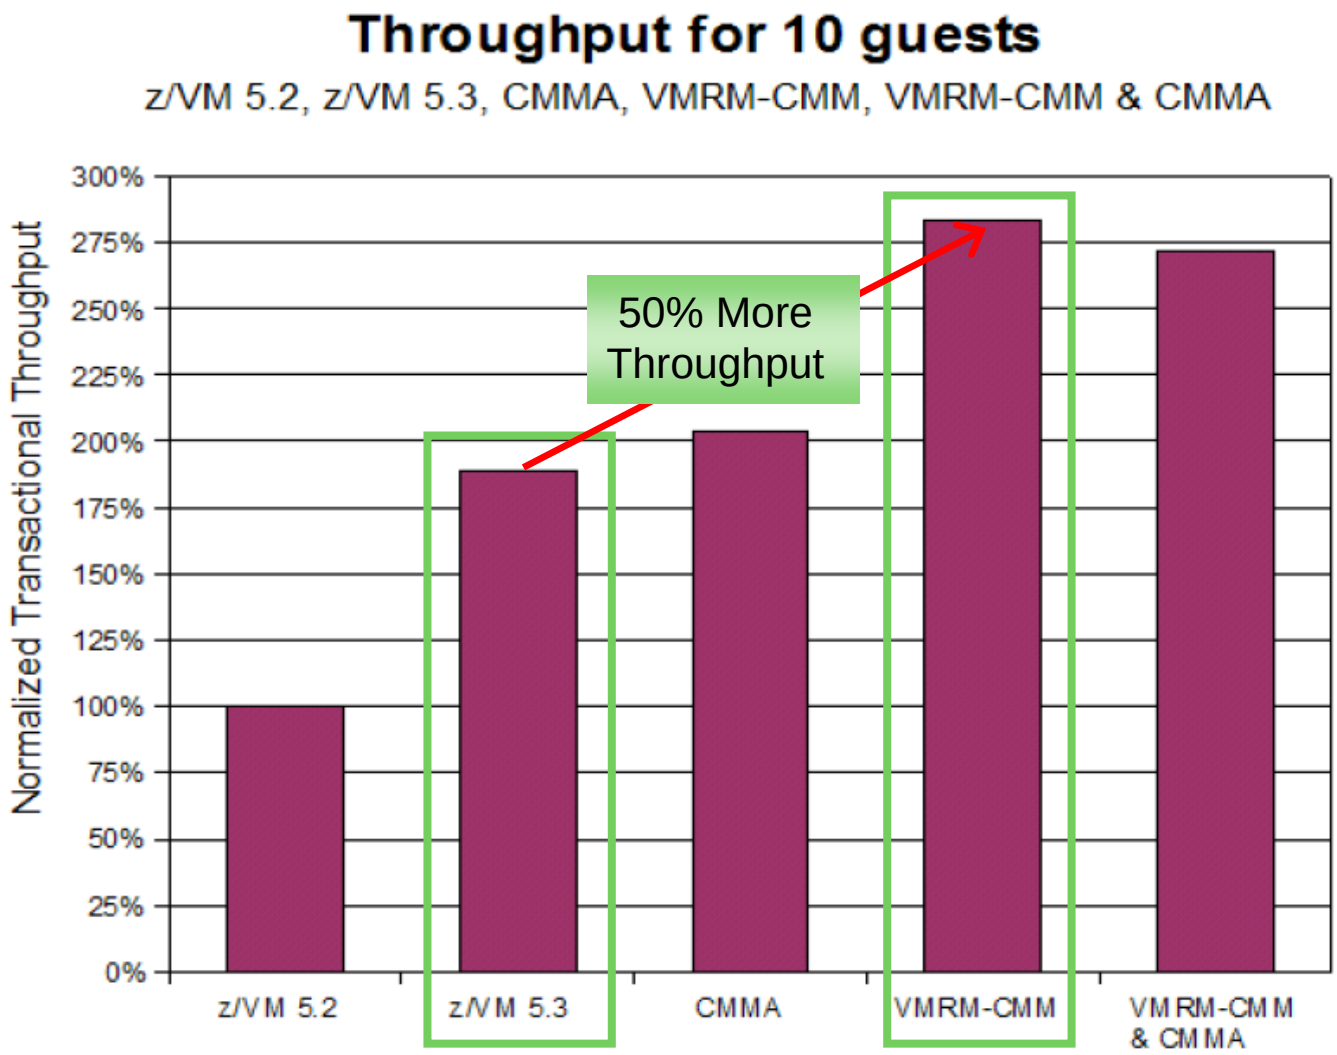

Linux and z/VM Technology Exploitation

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

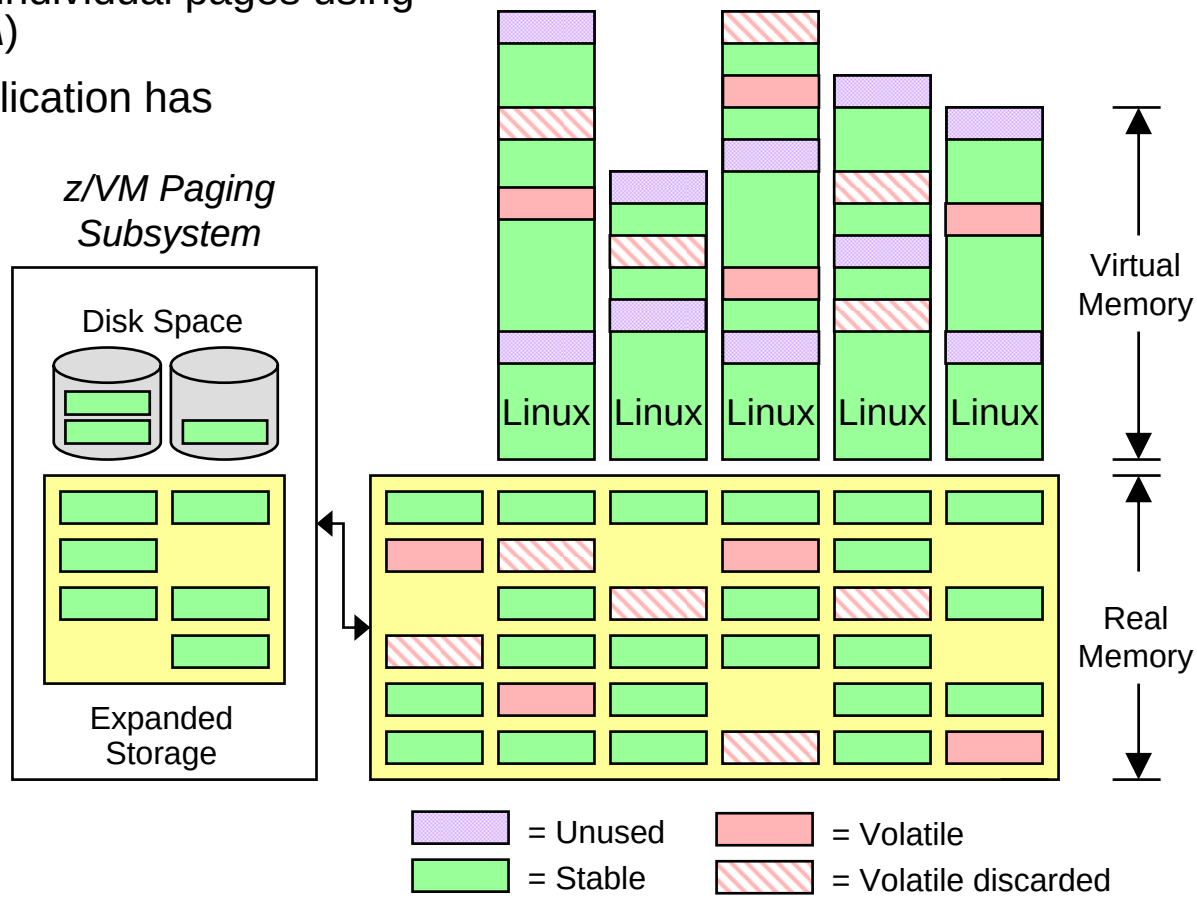

Transaction Rate versus Number of Hosted Servers

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

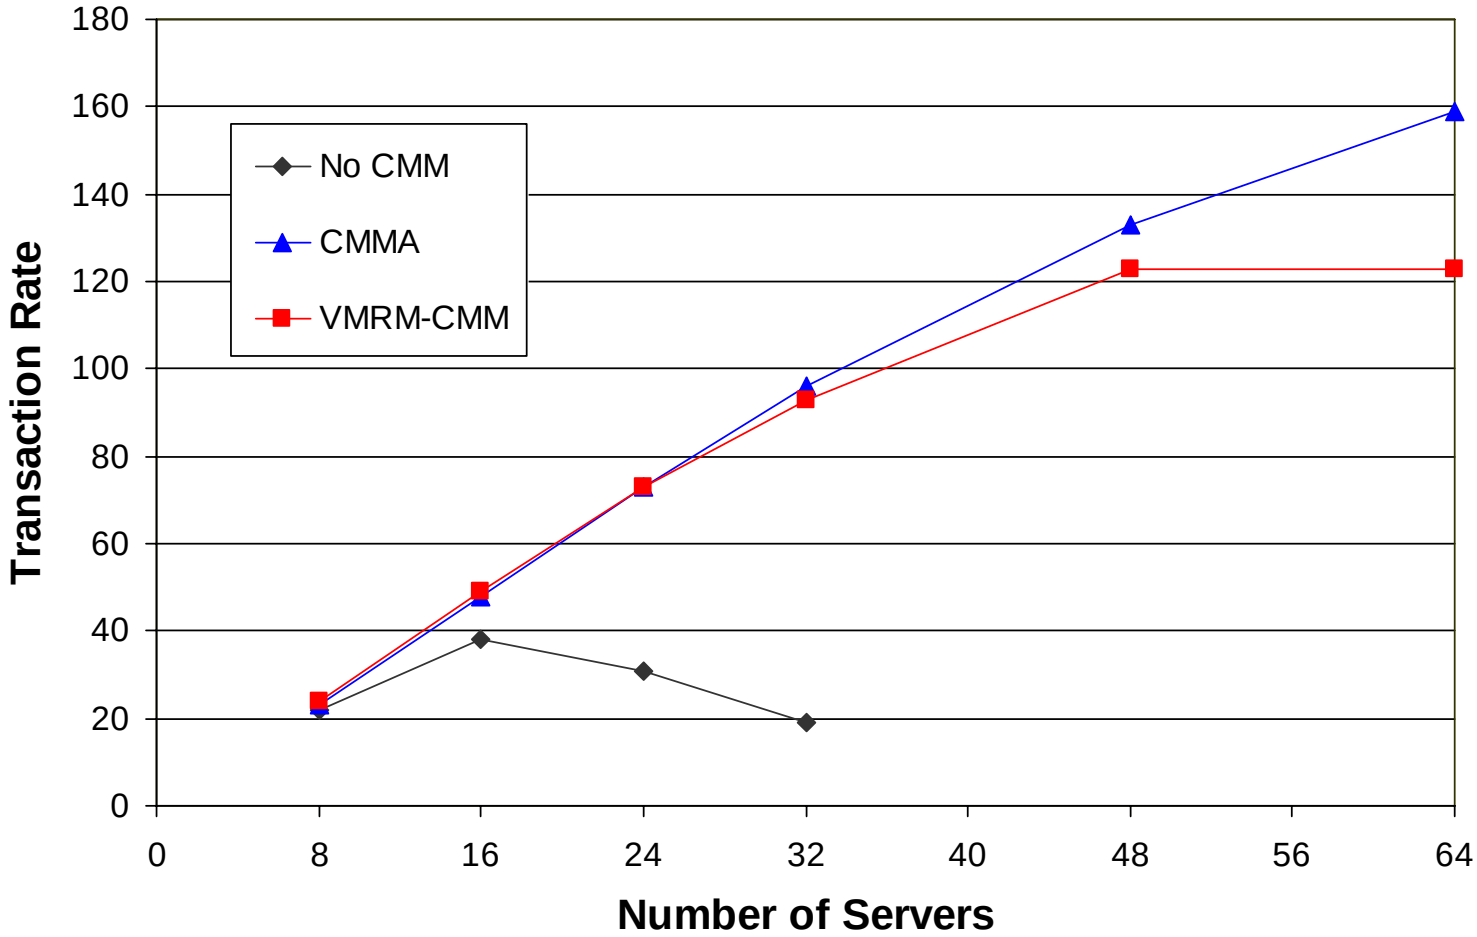

* z/VM running in IBM System z9 LPAR with 6GB of Central Storage and 2GB of Expanded Storage

Transaction Rate versus Number of Hosted Servers

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

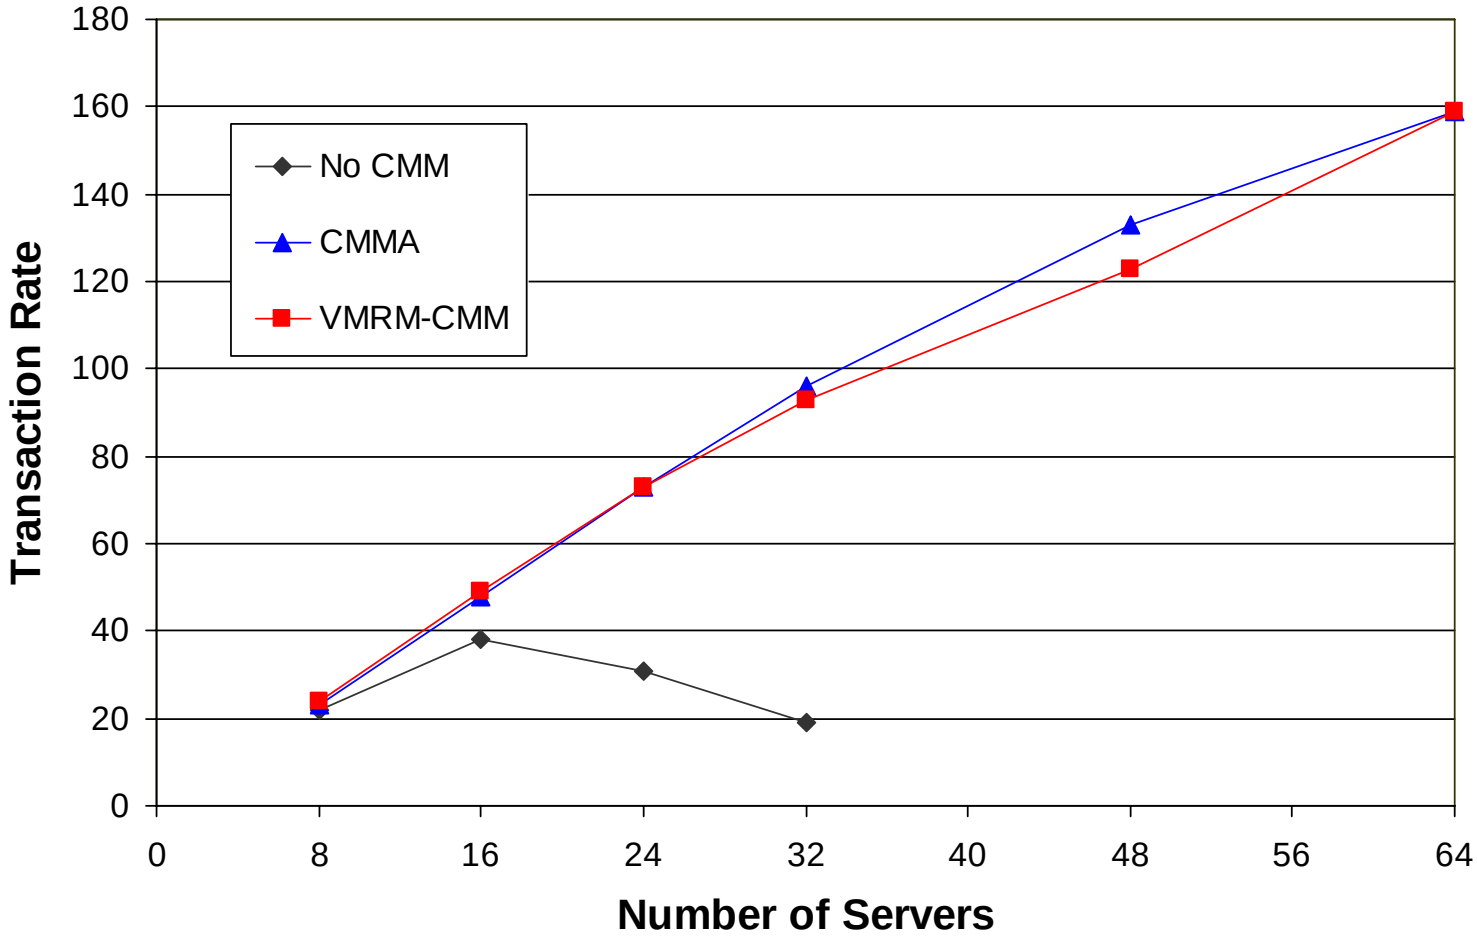

With APAR VM64439 applied

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

Paging Space Utilization versus Number of Hosted Servers

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

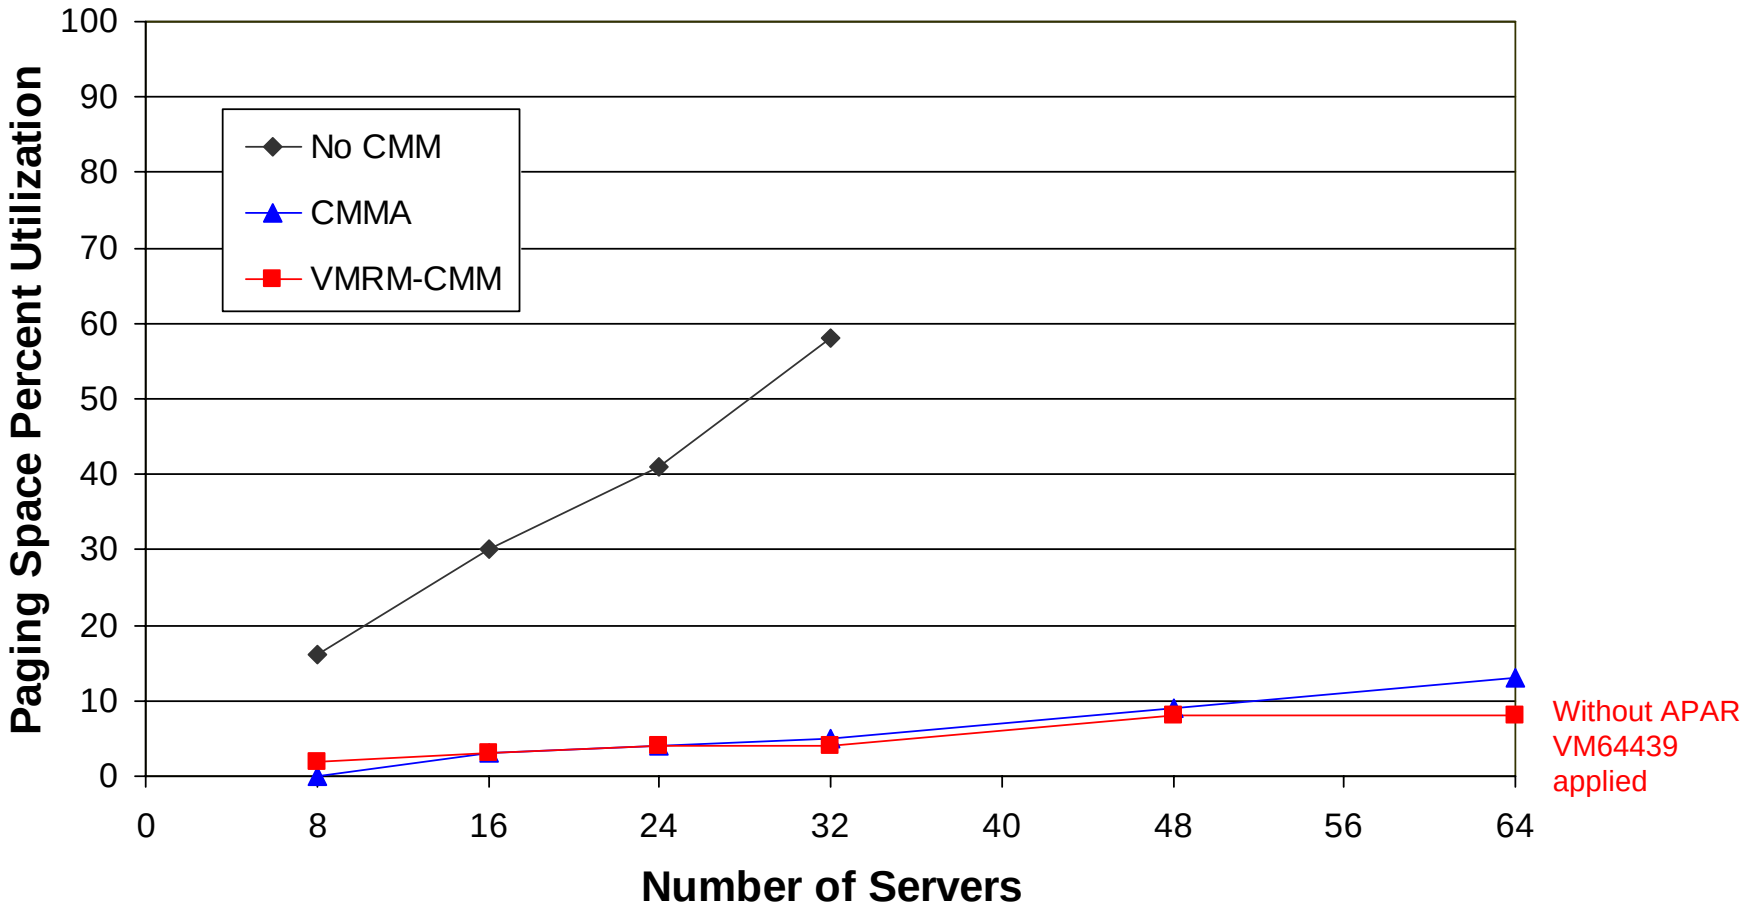

* z/VM running in IBM System z9 LPAR with 6GB of Central Storage and 2GB of Expanded Storage

z/VM Technology: Advanced Disk Support

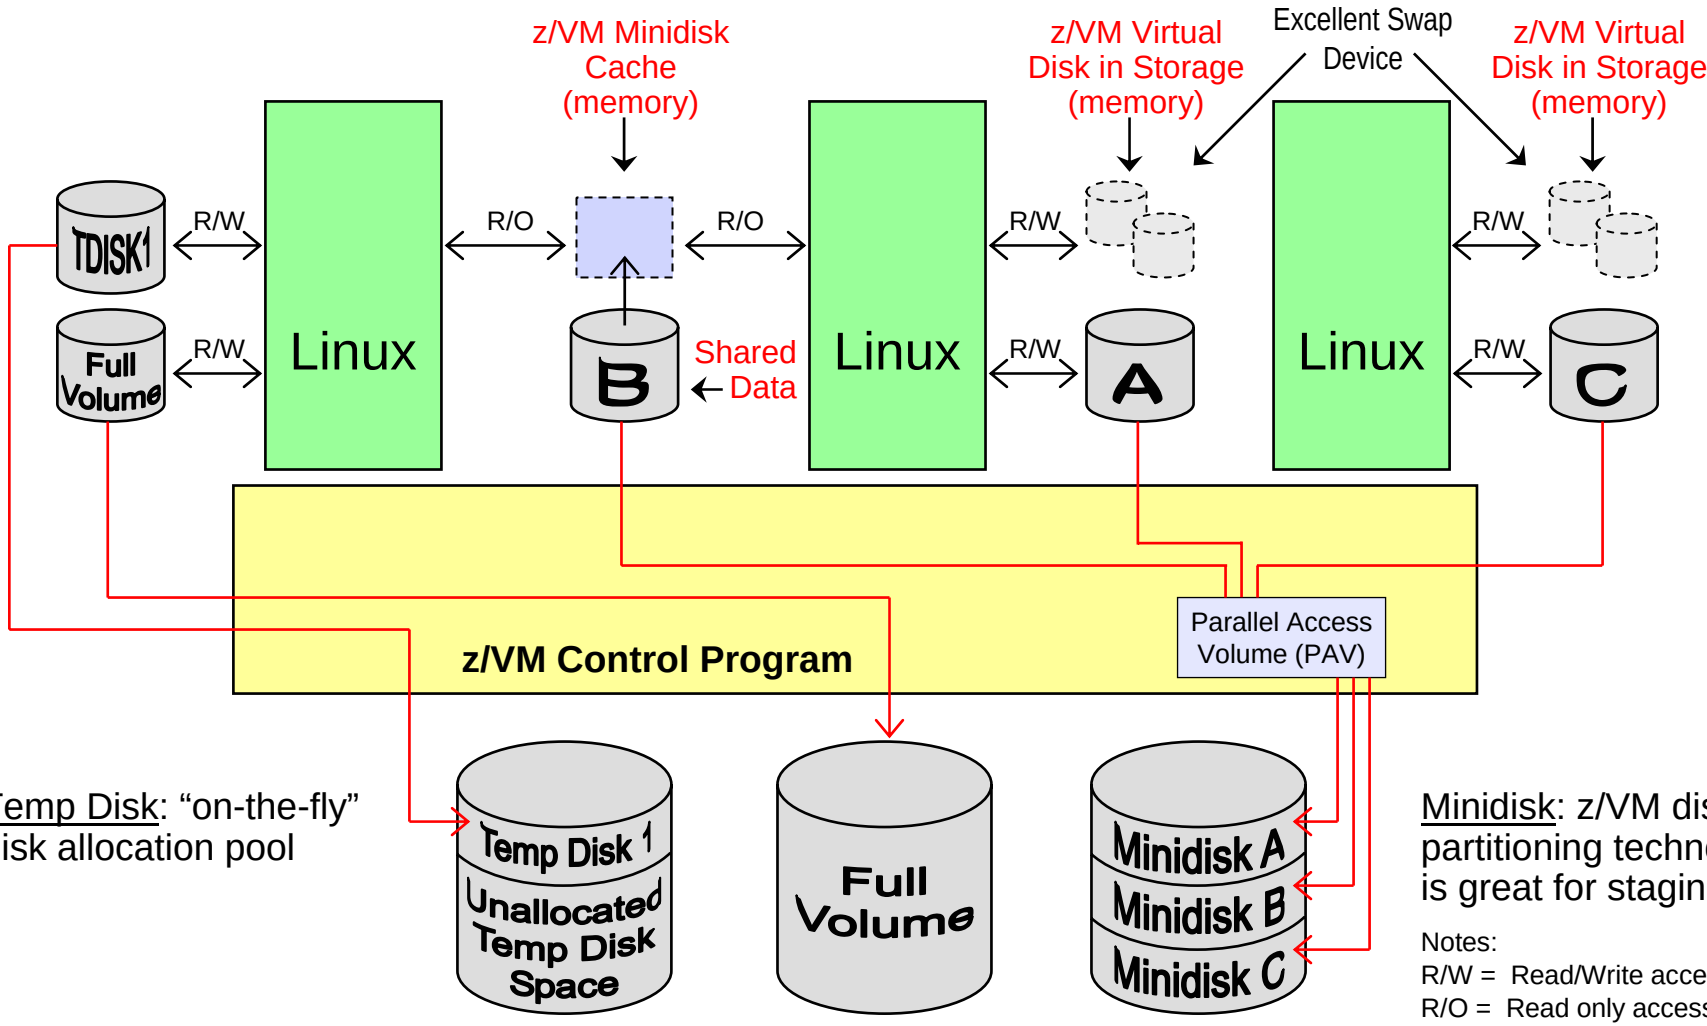

Temp Disk: "on-the-fly" disk allocation pool

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

- 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

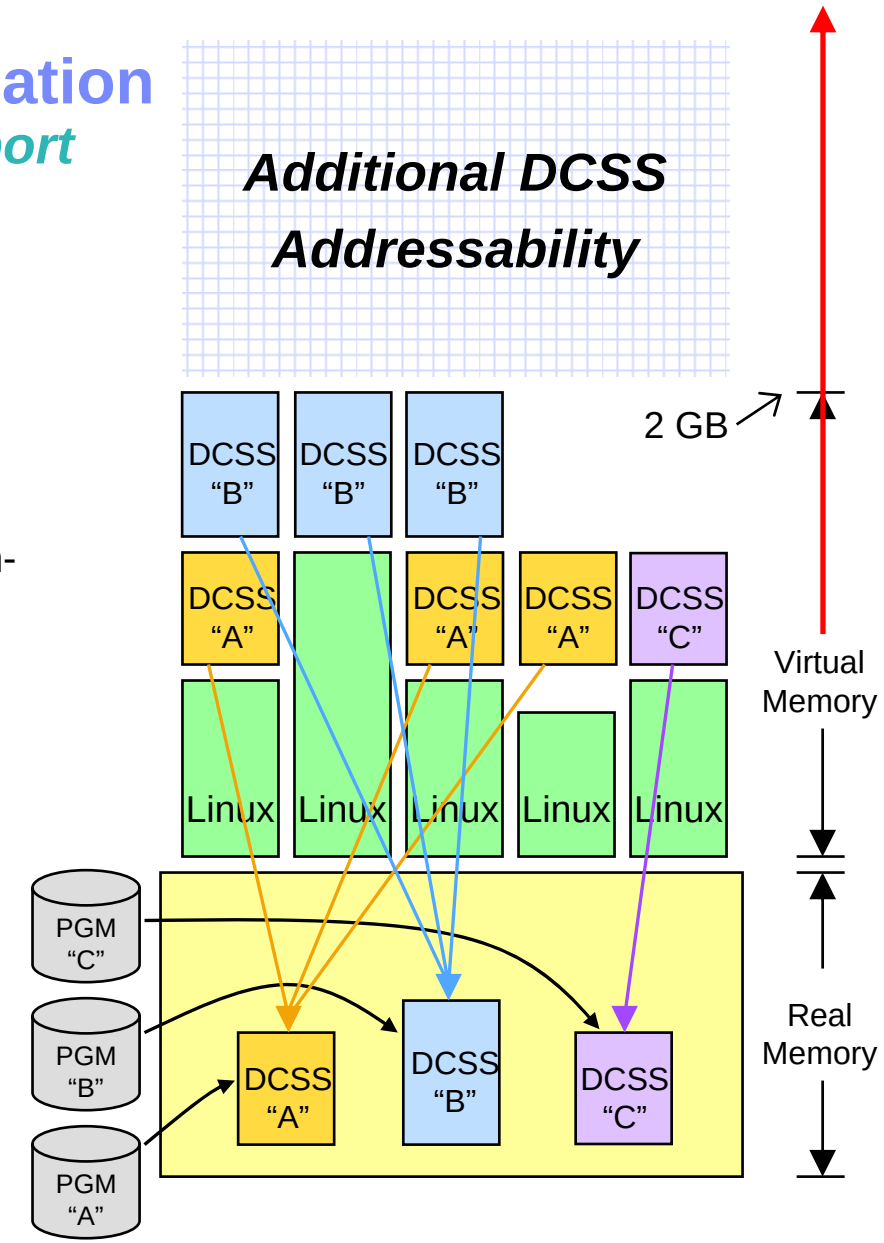

Extreme Virtualization with Linux on z/VM

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

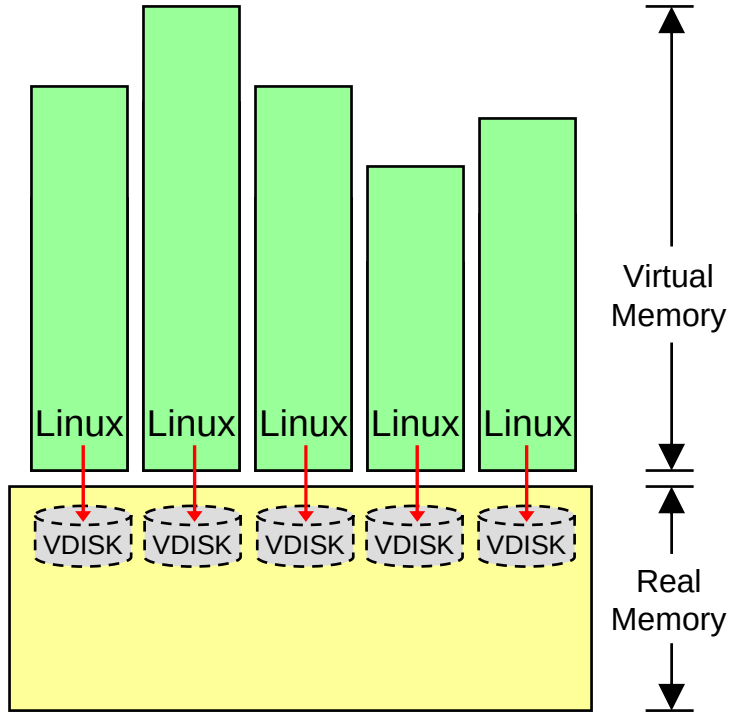

System z and N_Port ID Virtualization (NPIV)

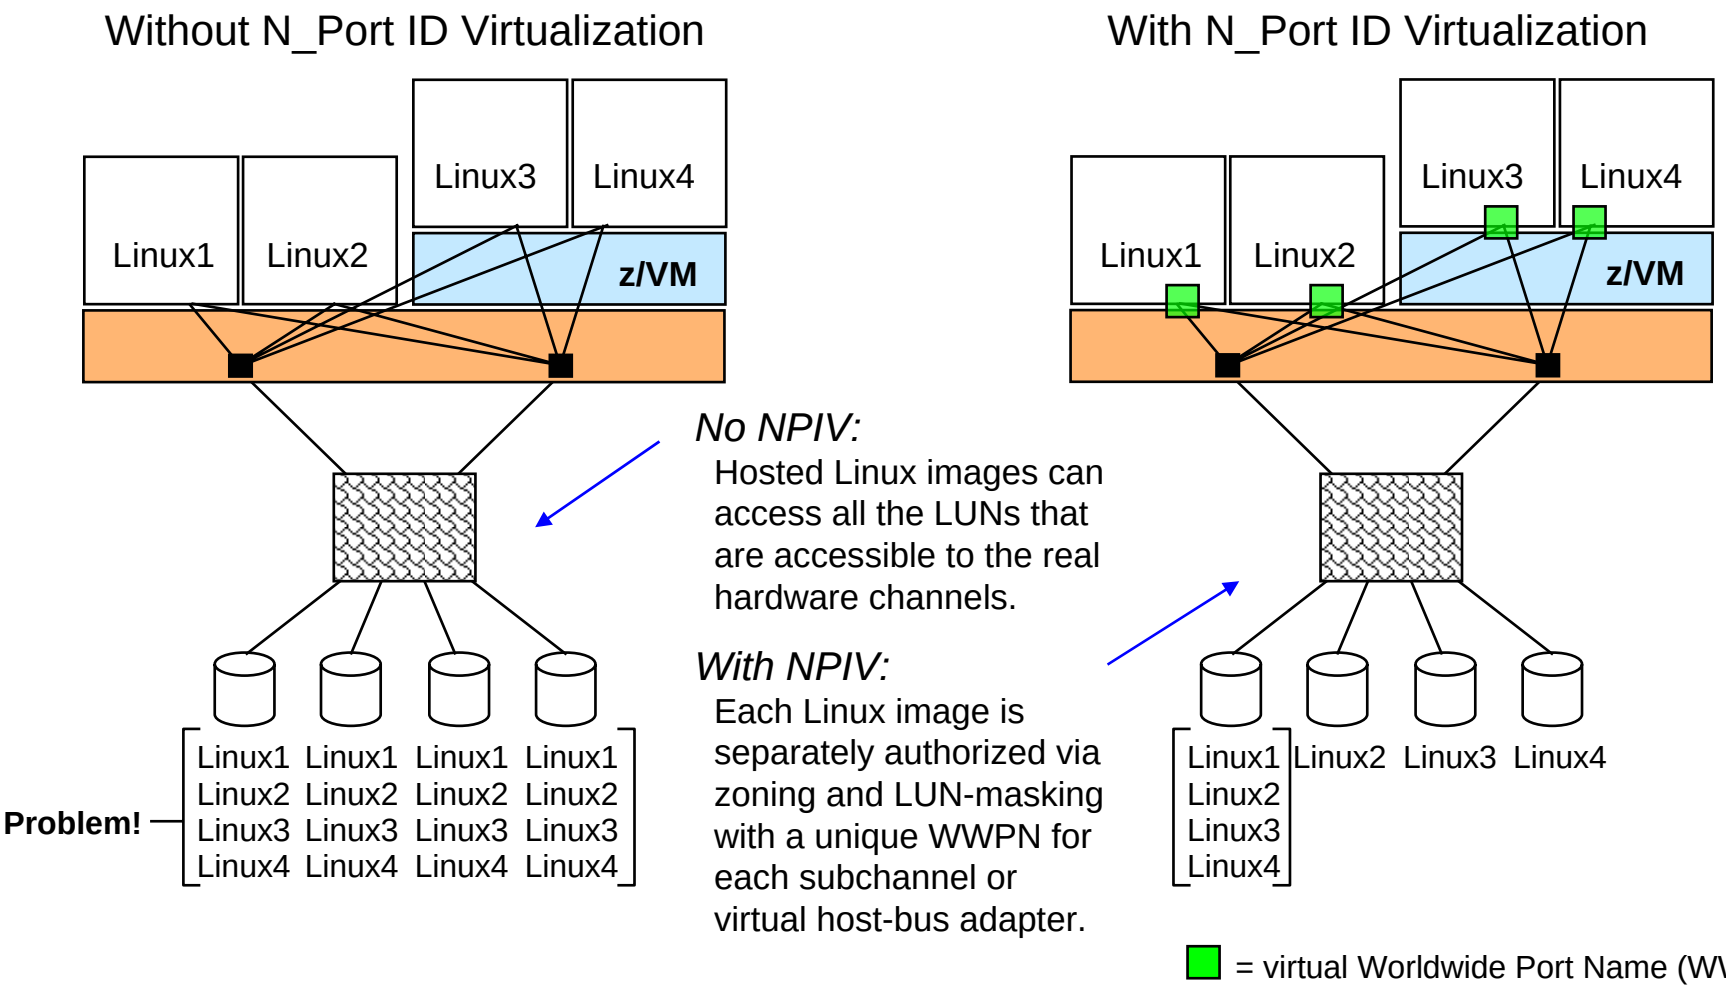

No NPIV:

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 and V5.4**
  - z/VM V5.2 support available with PTF for APAR VM64128

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

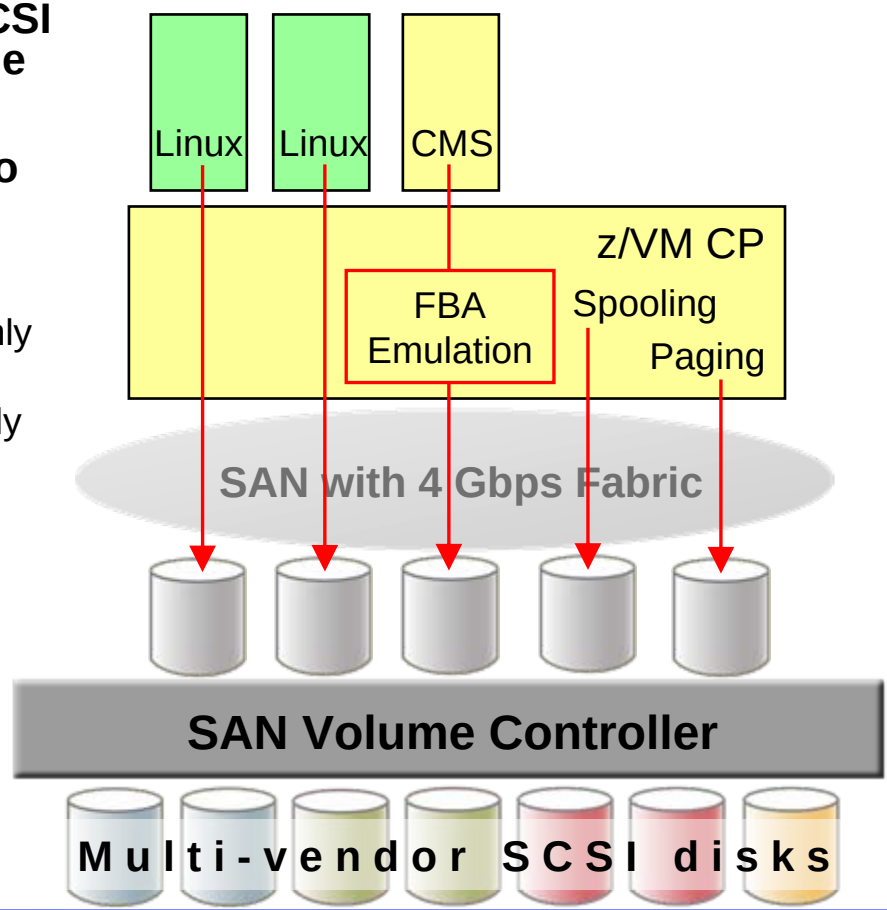

z/VM Virtual Switch Link Aggregation Support

Enhanced Networking Bandwidth and Business Continuance

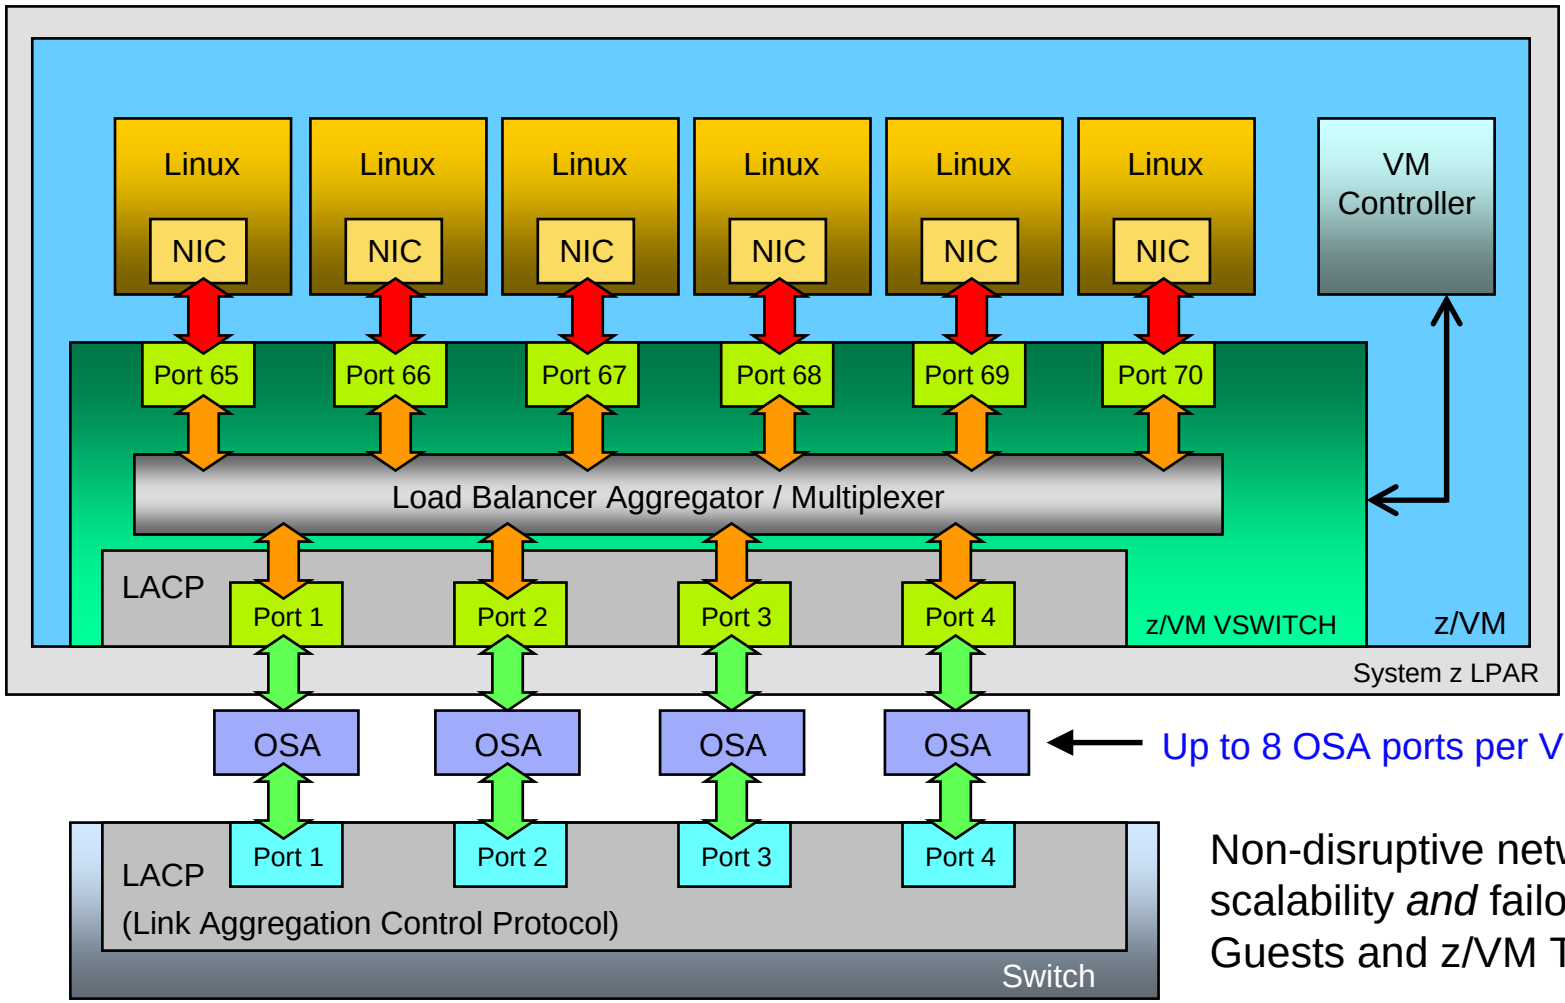

Non-disruptive networking scalability *and* failover for Guests and z/VM TCP/IP.

z/VM Virtual Switch Support

Layer 3 Compared to Layer 2 Switching

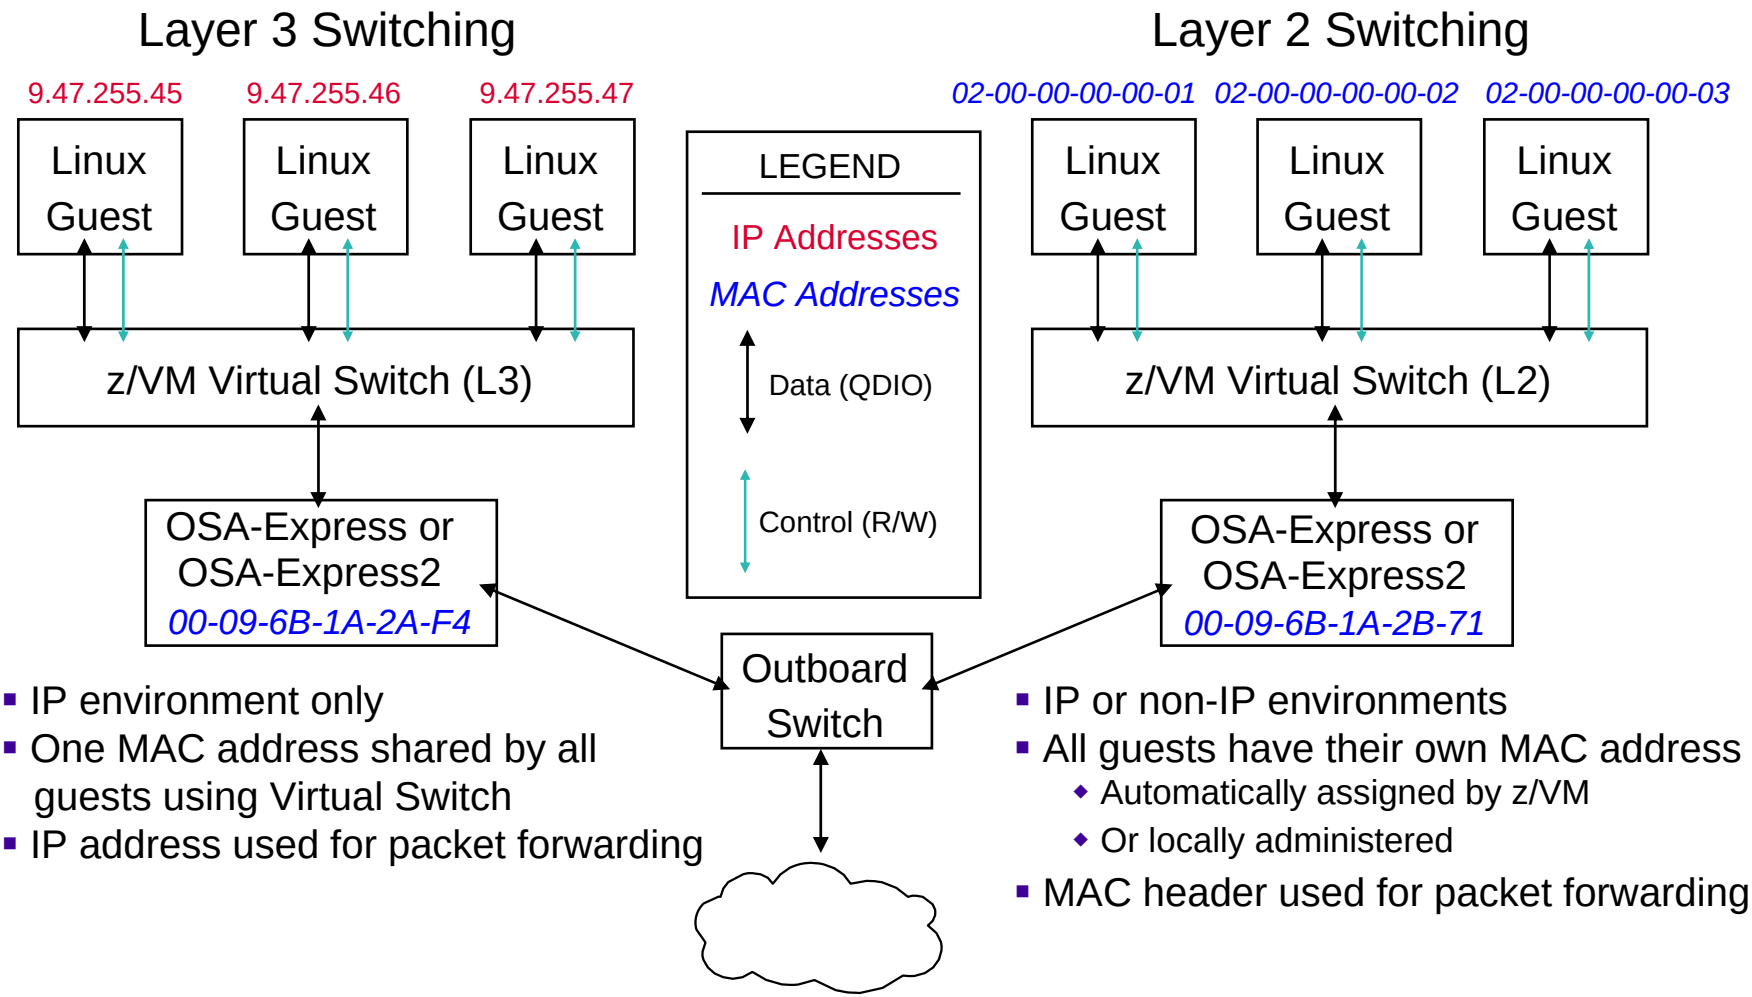

- IP environment only
- One MAC address shared by all guests using Virtual Switch
- IP address used for packet forwarding

- IP or non-IP environments
- All guests have their own MAC address
  - Automatically assigned by z/VM
  - Or locally administered
- MAC header used for packet forwarding

z/VM Virtual Switch and OSA-Express Port Isolation

- **Allows users to restrict guest-to-guest communications within a Virtual Switch by exploiting OSA-Express QDIO data connection isolation**
- **Provides a mechanism to isolate a QDIO data connection on an OSA port**
  - Enables network isolation for operating systems sharing physical network connectivity

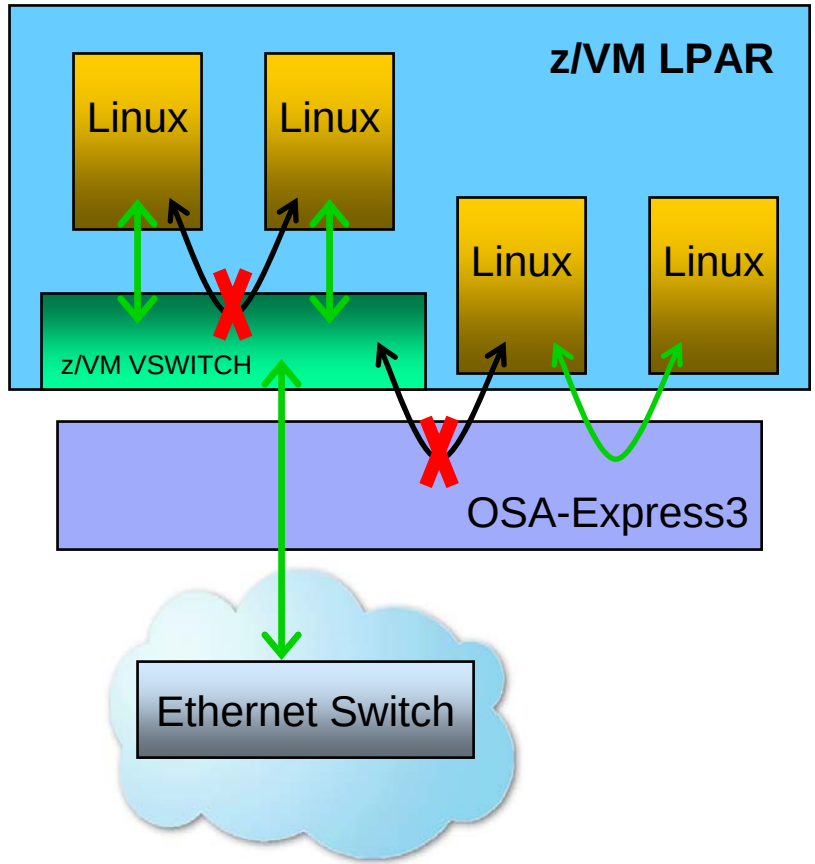

z/VM Command and Control Infrastructure

- **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

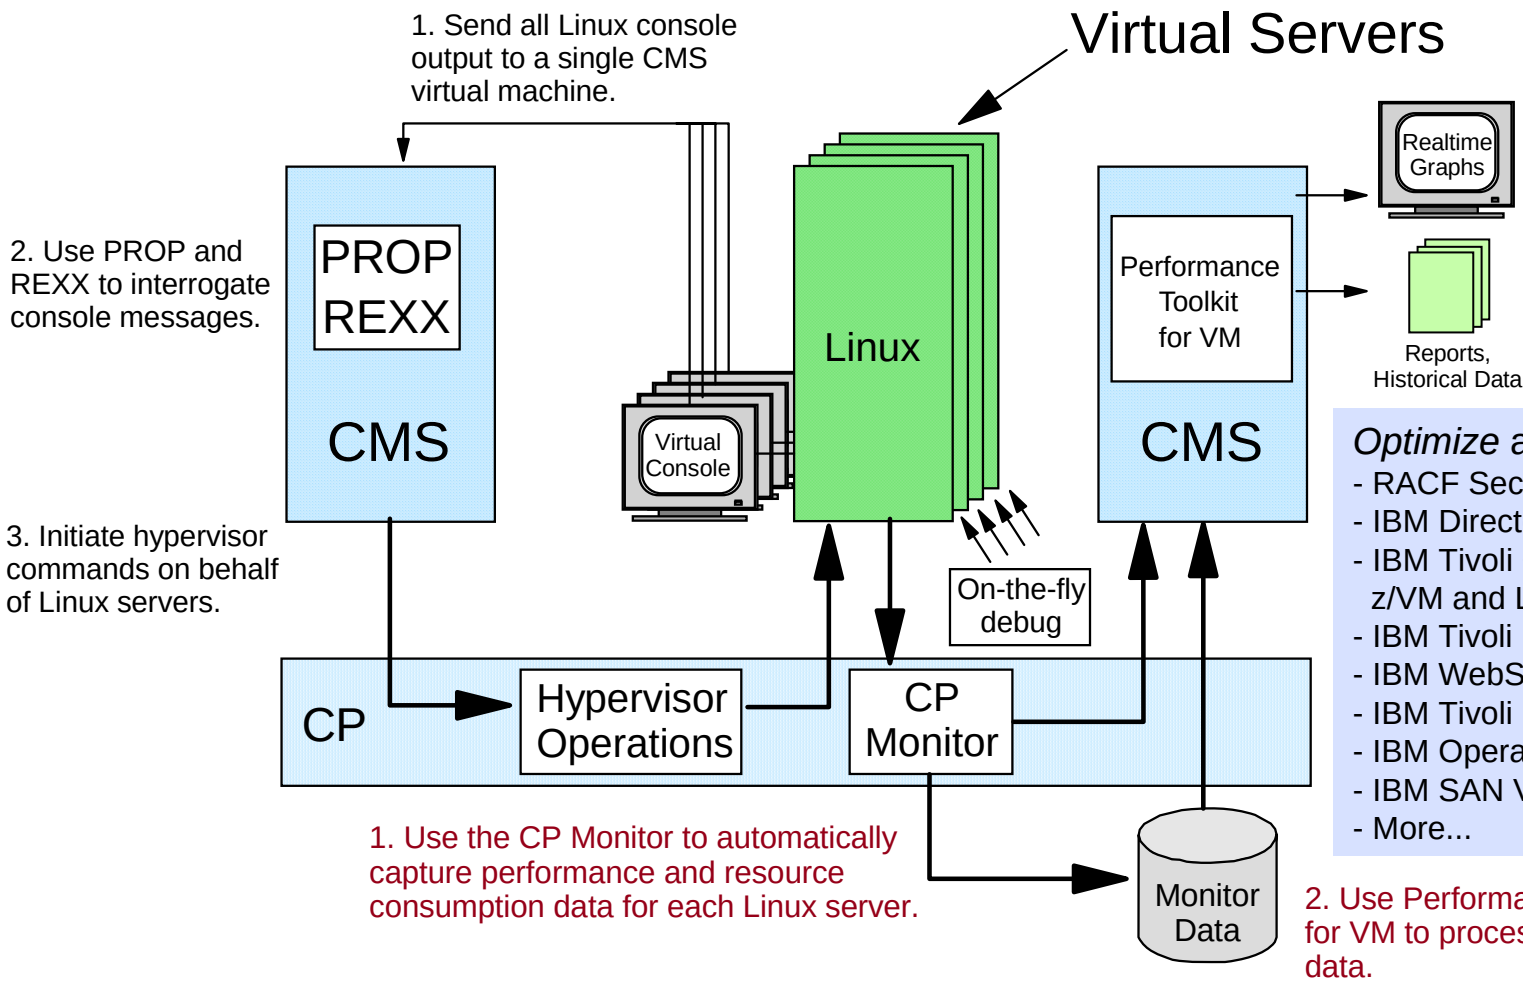

Optimize and *Integrate* with:

- 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

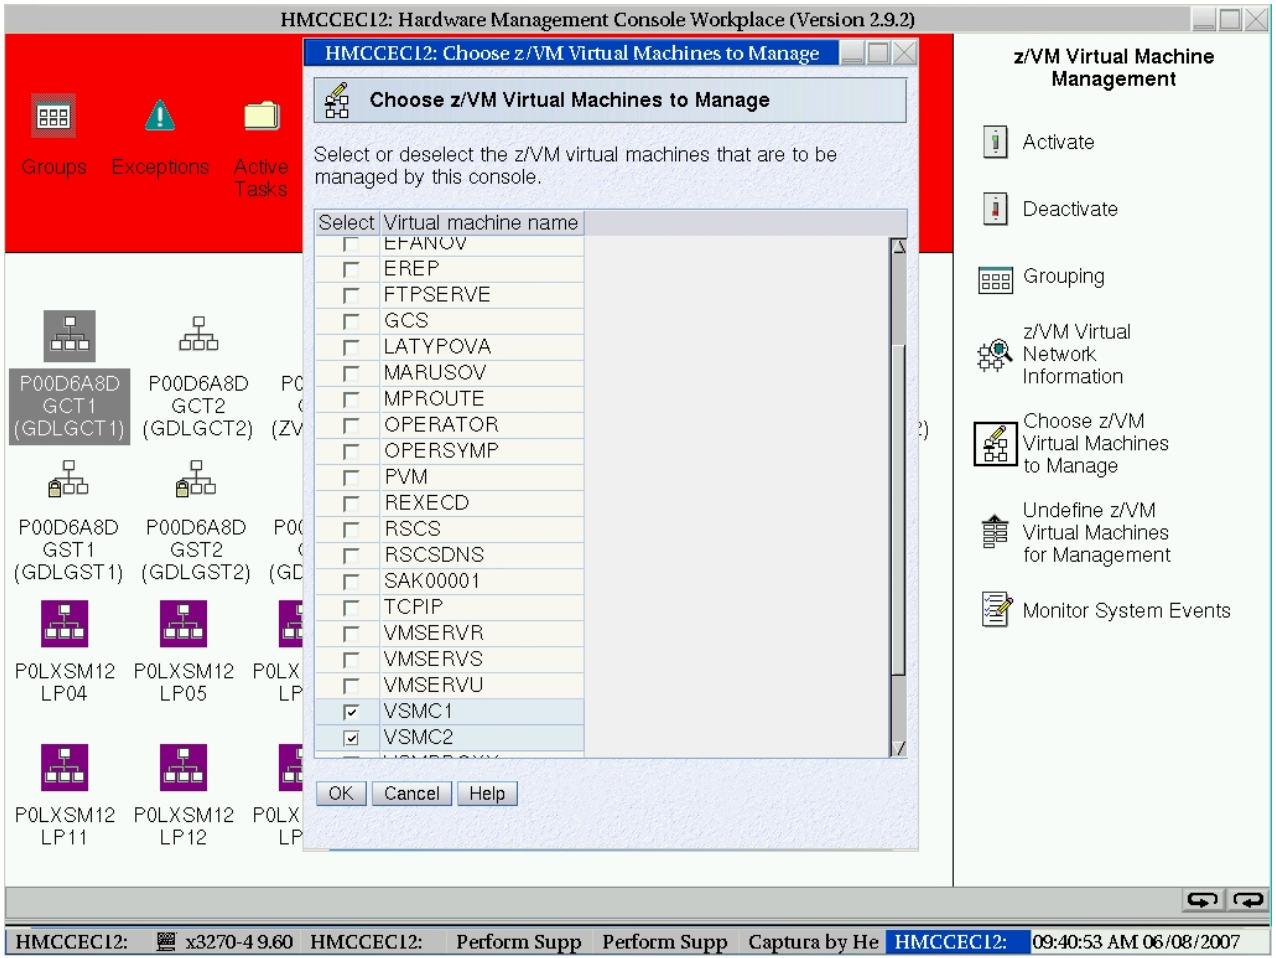

GDPS/PPRC Multiplatform Resiliency for System z

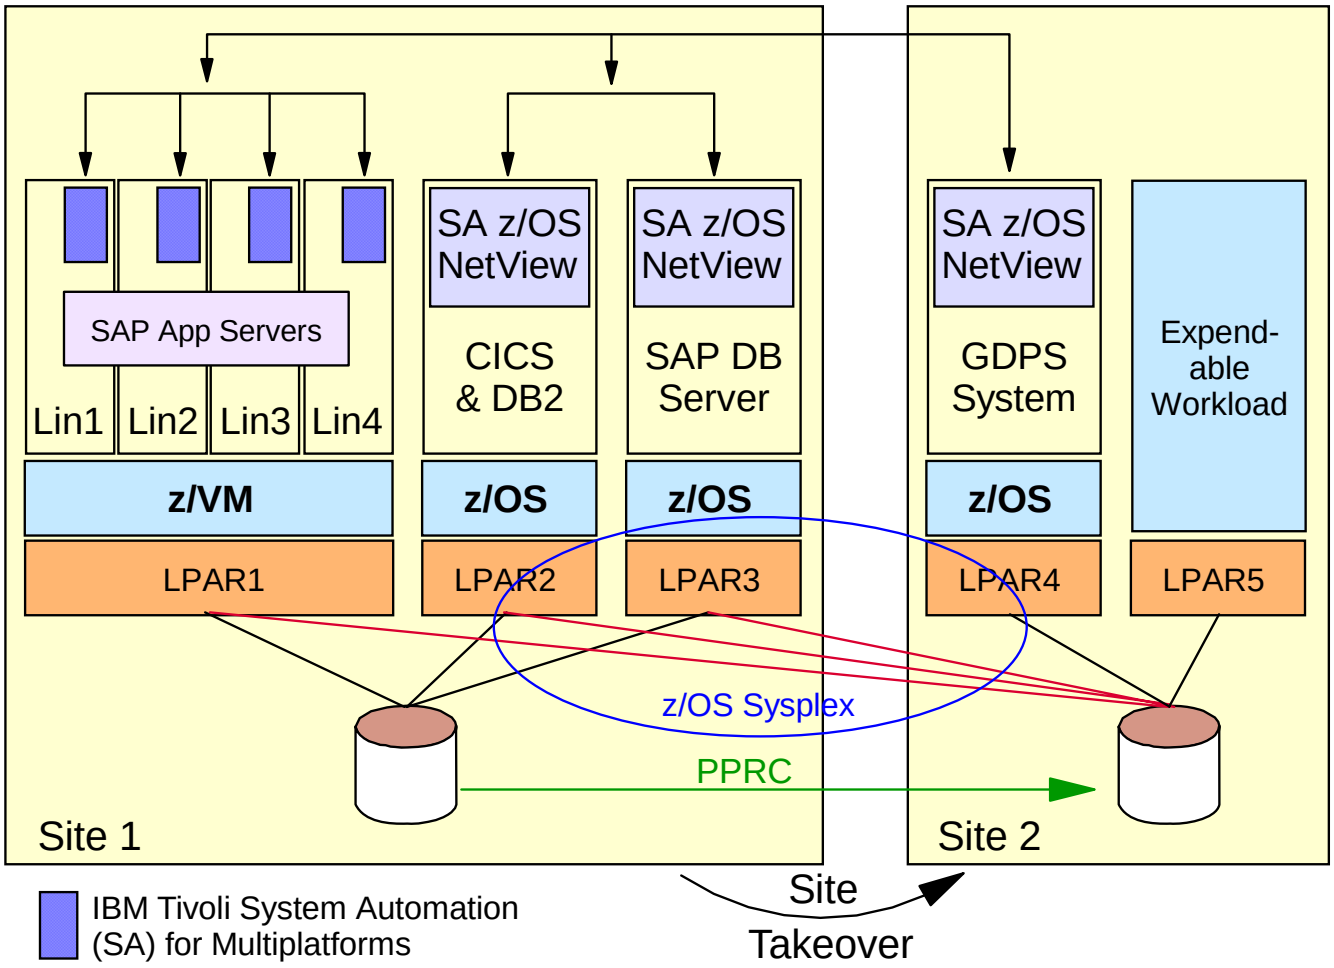

- Designed for customers with distributed applications
- SAP application server running on Linux for System z
- SAP DB server running on z/OS
- Coordinated near-continuous availability and DR solution for z/OS, Linux guests, and z/VM
- Uses z/VM HyperSwap function to switch to secondary disks
- Sysplex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

Questions

Backup Material

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

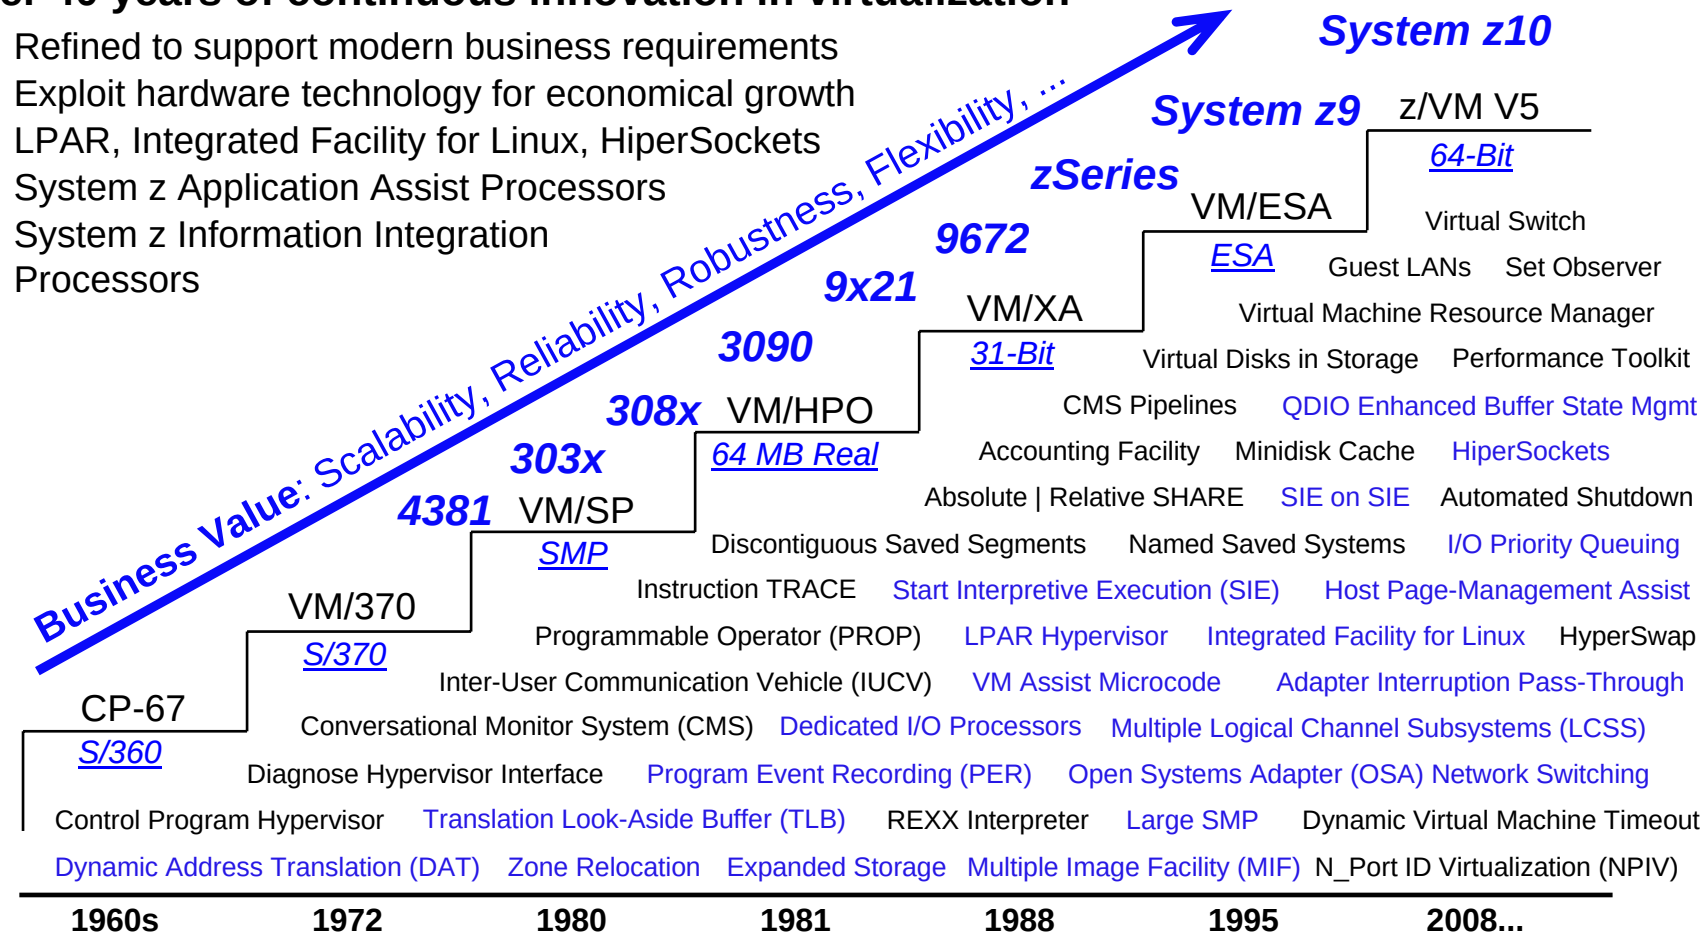

IBM System z – a comprehensive and sophisticated suite of virtualization function

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

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**

The Power and Flexibility of System z Virtualization

- Over 40 years of continuous innovation in virtualization technologies
- Multiple images concurrently share all physical resources
- Resources delivered as required, automatically, based on business-oriented goals
- New OS images can be started without affecting ongoing work
- Hardware assists used to accelerate virtualization operations (e.g., SIE)

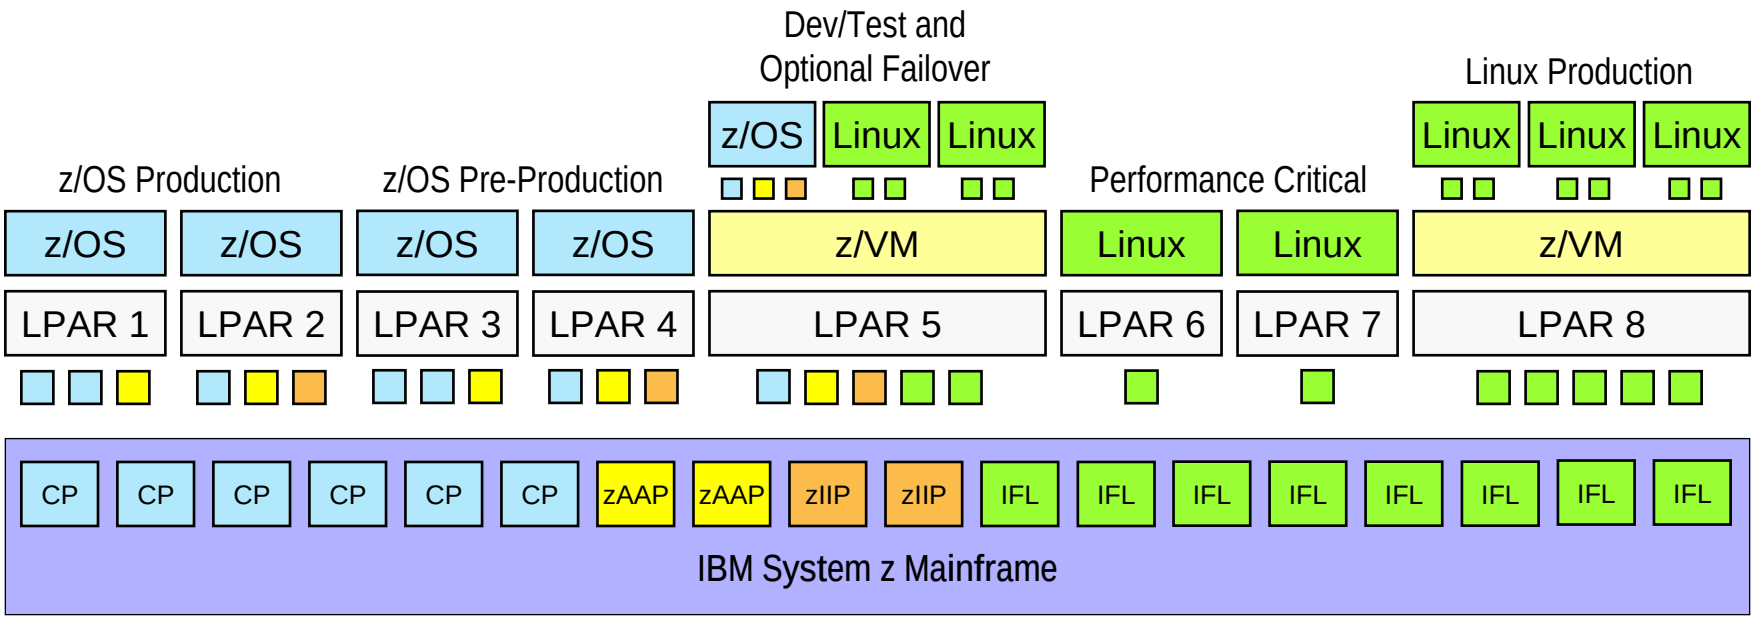

z/VM-Mode LPAR Support for IBM System z10

- **New LPAR type for IBM System z10: *z/VM-mode***
  - Allows z/VM V5.4 users to configure all CPU types in a z10 LPAR
- **Offers added flexibility for hosting mainframe workloads**
  - Add *IFLs* to an existing standard-engine z/VM LPAR to host Linux workloads
  - Add *CPs* to an existing IFL z/VM LPAR to host z/OS, z/VSE, or traditional CMS workloads
  - Add *zAAPs* and *zIIPs* to host eligible z/OS specialty-engine processing
  - Test integrated Linux and z/OS solutions in the same LPAR
- **No change to IBM software licensing**
  - Software continues to be licensed according to CPU type

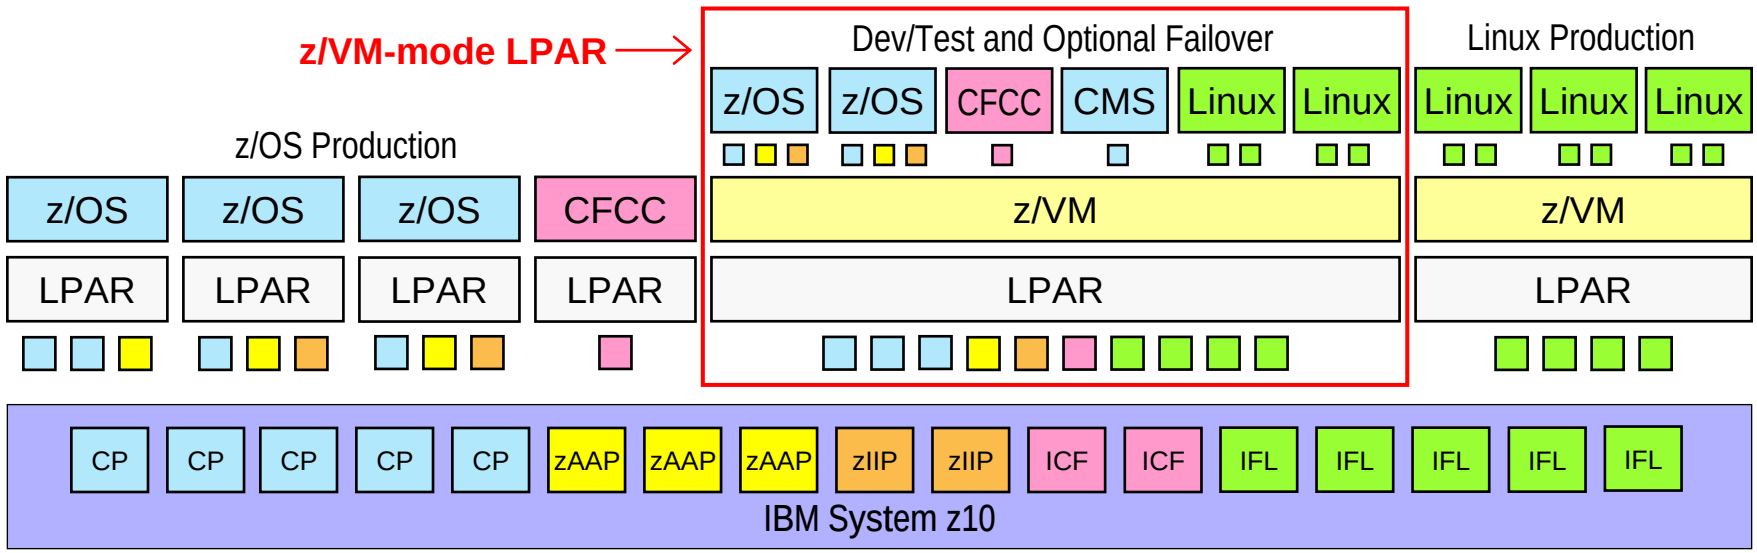

Provisioning Linux Virtual Machines on System z

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

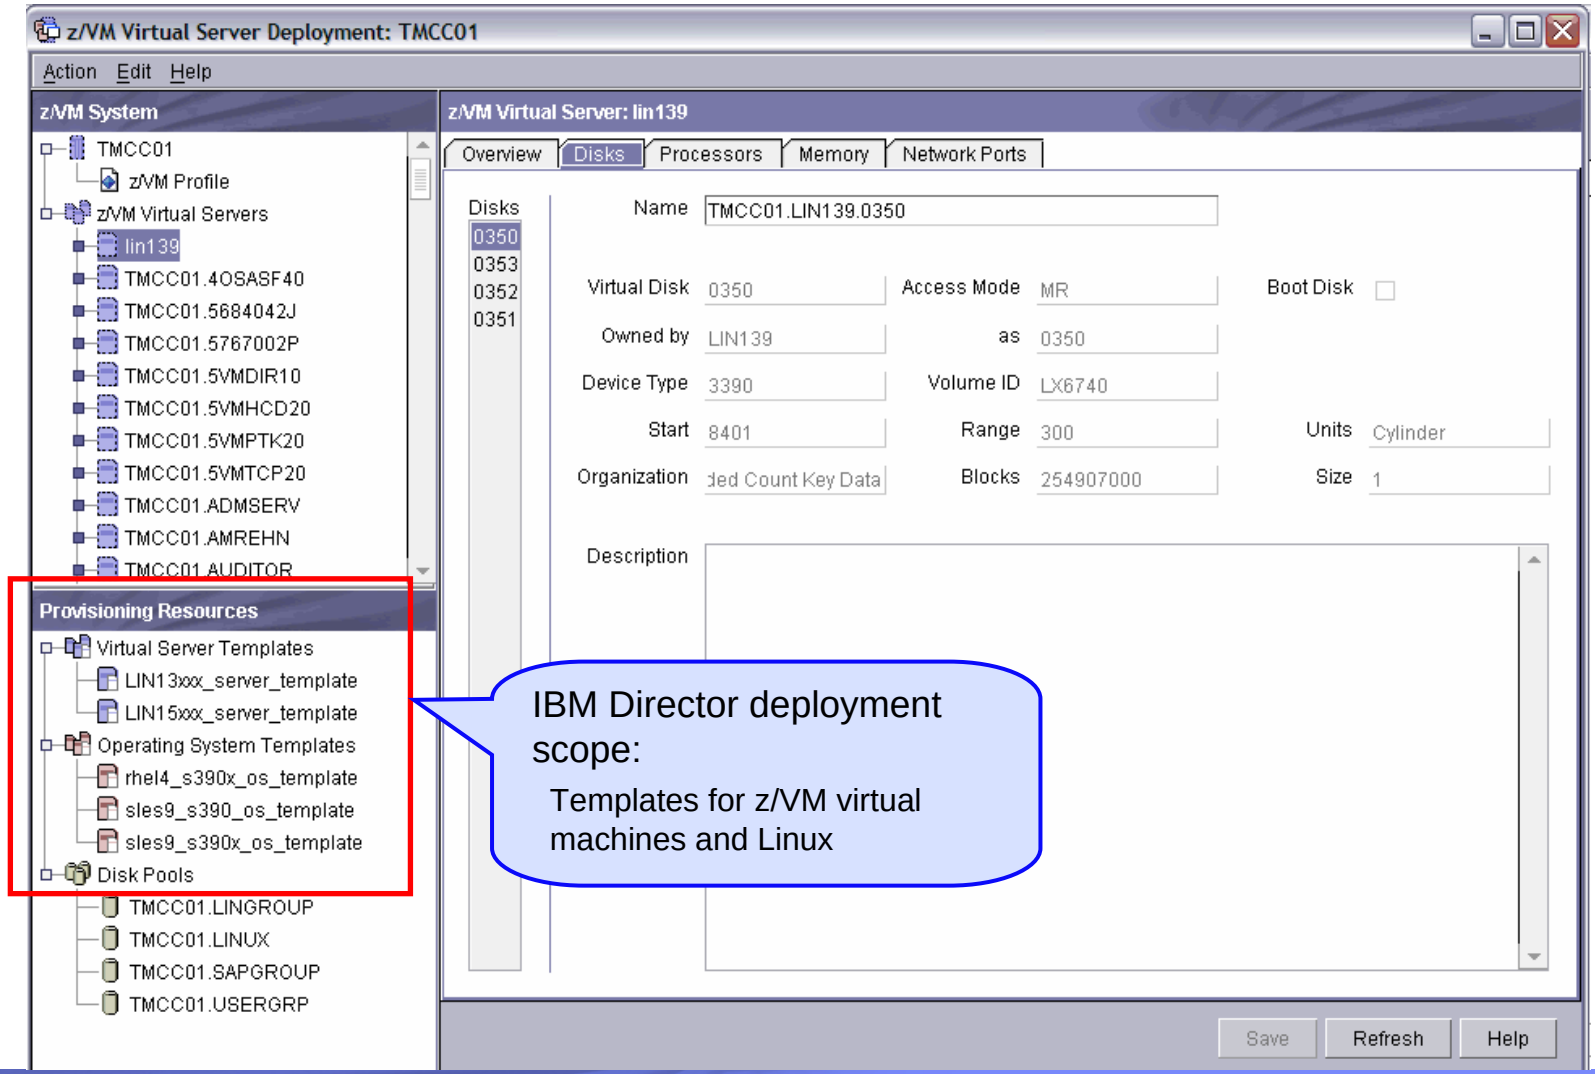

Monitoring System z Virtual Linux Servers

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

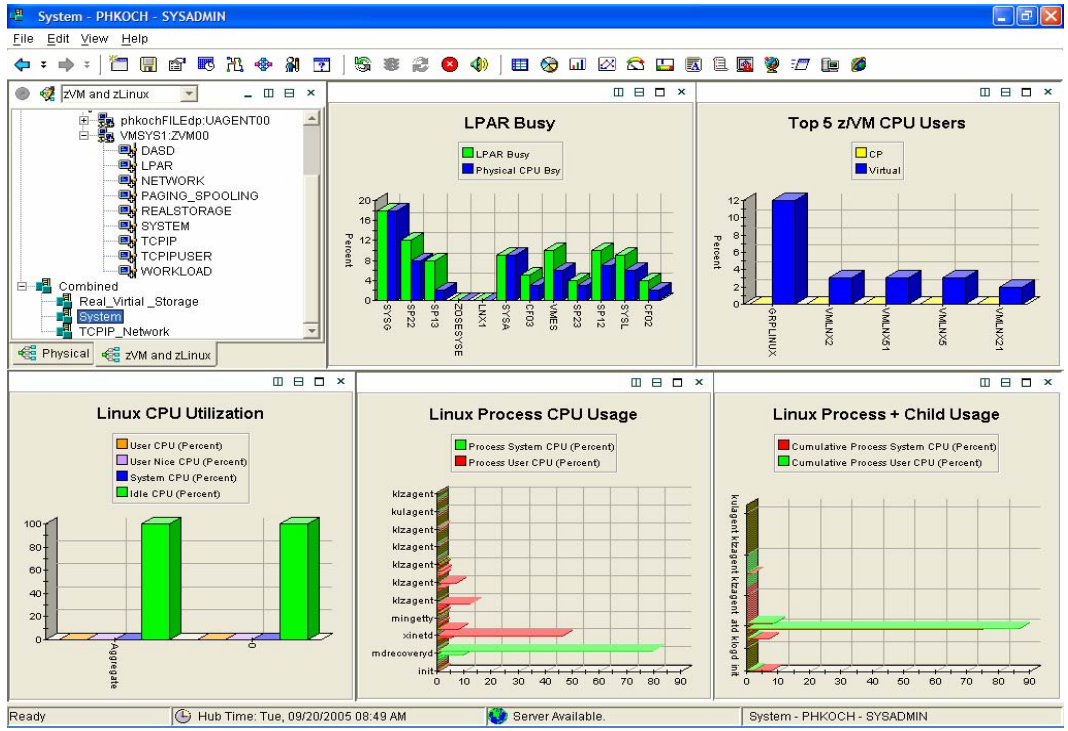

Learn more at: **ibm.com**/software/tivoli/products/omegamon-xe-zvm-linux

IBM Tivoli Virtualization Management for System z

Helping Clients Manage and Control Their Virtualized IT Infrastructure

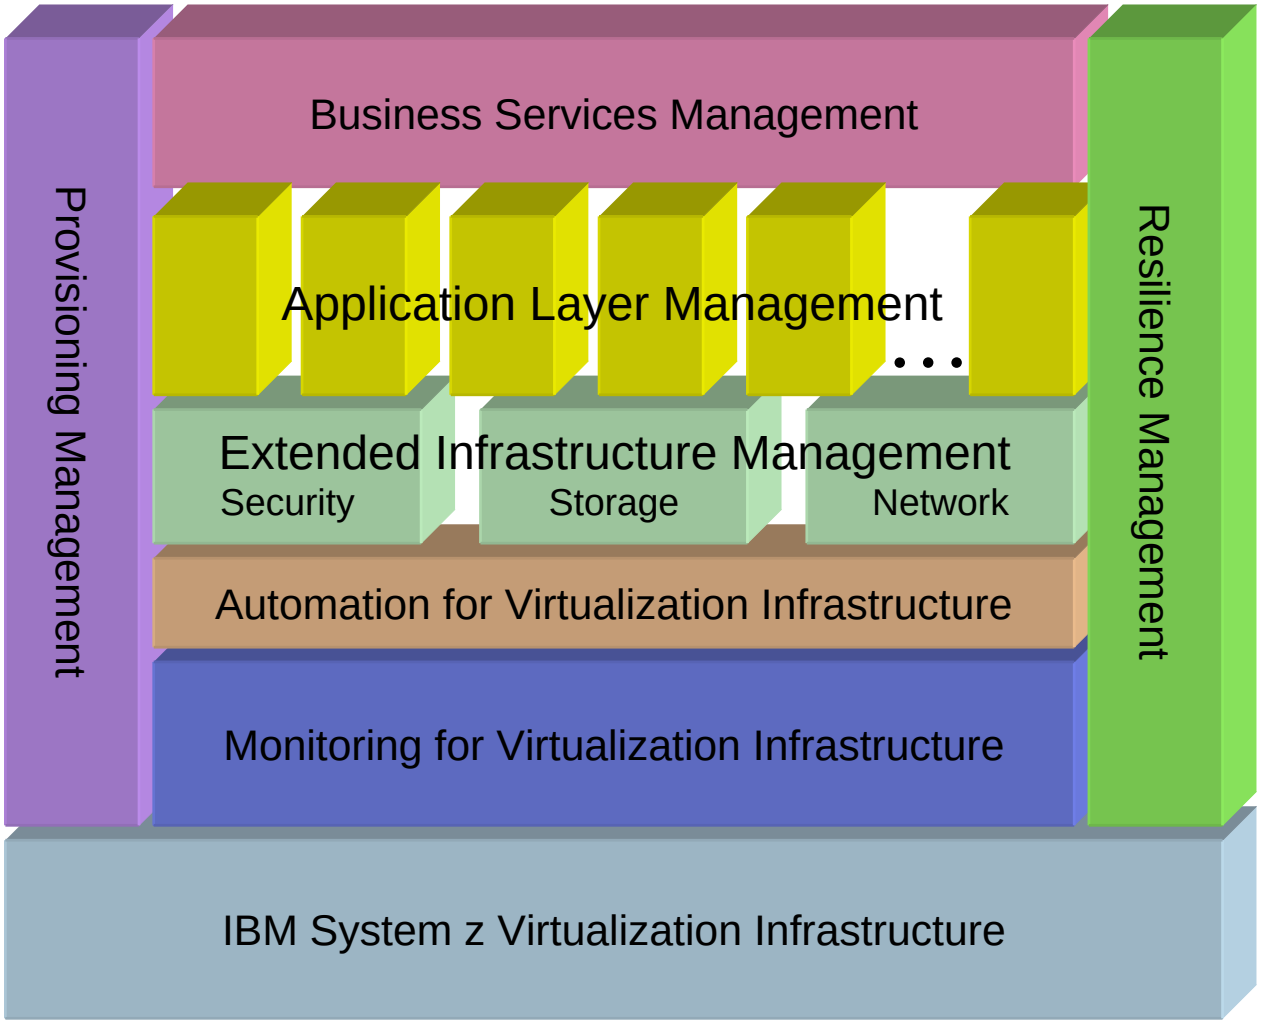

IBM System z Virtualization Infrastructure
- IBM System z hardware (including LPAR hypervisor)
- IBM z/VM Version 5

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

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

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

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**

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)**