IBM System z: The Ultimate Virtualization Platform

October 9, 2007

Reed Mullen
Product Manager, System z Virtualization Technology

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.

IBM System z: The Ultimate Virtualization Platform

- ***Virtualize*** **everything with up to 100% utilization rates**
  - CPU, memory, network, I/O, cryptographic features, coupling facility, ...

Consolidate all types of workloads

- ***Massively scale*** **your workload on a single System z mainframe**
  - The Linux-on-z/VM record is 97,943 virtual machines
  - Each virtual machine on z/VM can access up to 24,576 devices

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

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

Secure your virtual servers and reduce business risk

- ***Non-disruptively add*** **anything**
  - 54x CPU scalability per mainframe, 32x CPU scalability per z/VM LPAR
  - z/VM is designed to support up to 8 TB of active virtual memory

Rapidly respond to workload spikes

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

Increase staff productivity and virtualize the enterprise

IBM Consolidation Announcement Highlights

- IBM will consolidate thousands of servers onto approximately 30 System z mainframes
- We expect substantial savings in multiple dimensions: energy, software and system support costs
- Major proof point of IBM's 'Project Big Green' initiative
- The consolidated environment will use 80 percent less energy
- This transformation is enabled by the sophisticated virtualization capability of System z

IBM'S PROJECT BIG GREEN SPURS GLOBAL SHIFT TO LINUX ON MAINFRAME

Plan to shrink 3,900 computer servers to about 30 mainframes targets 80 percent energy reduction over five years

Optimized environment to increase business flexibility

ARMONK, NY, August 1, 2007 – In one of the most significant transformations of its worldwide data centers in a generation, IBM (NYSE: IBM) today announced that it will consolidate about 3,900 computer servers onto about 30 System z mainframes running the Linux operating system. The company anticipates that the new server environment will consume approximately 80 percent less energy than the current set up and expects significant savings over five years in energy, software and system support costs.

At the same time, the transformation will make IBM's IT infrastructure more flexible to evolving business needs. The initiative is part of Project Big Green, a broad commitment that IBM announced in May to sharply reduce data center energy consumption for IBM and its clients.

Why System z Now?

New HW / SW spending

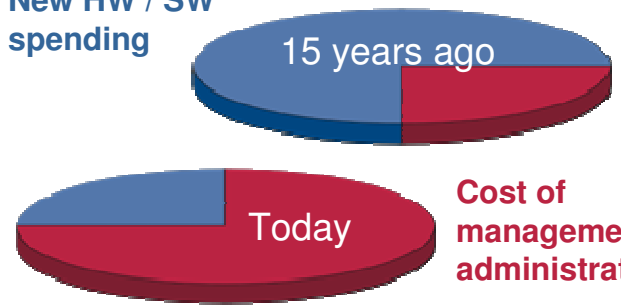

Cost of management & administration

Source: Tony Picardi, IDC
Economist.com: Make it simple. October 28th, 2004
From The Economist print edition

System z9 Managing Growth and Complexity

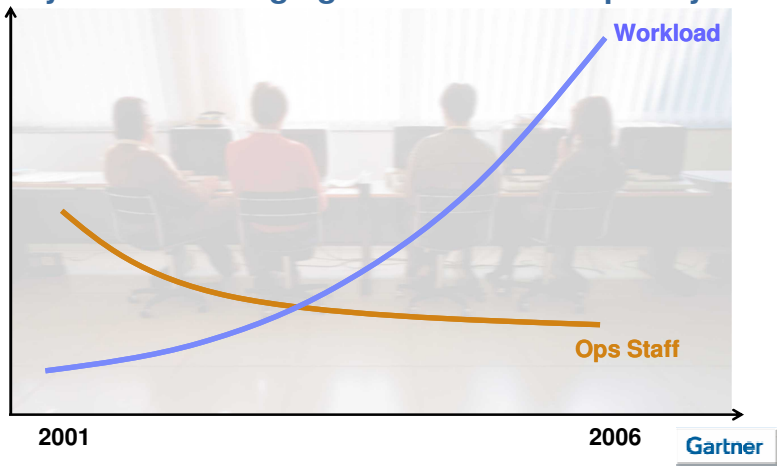

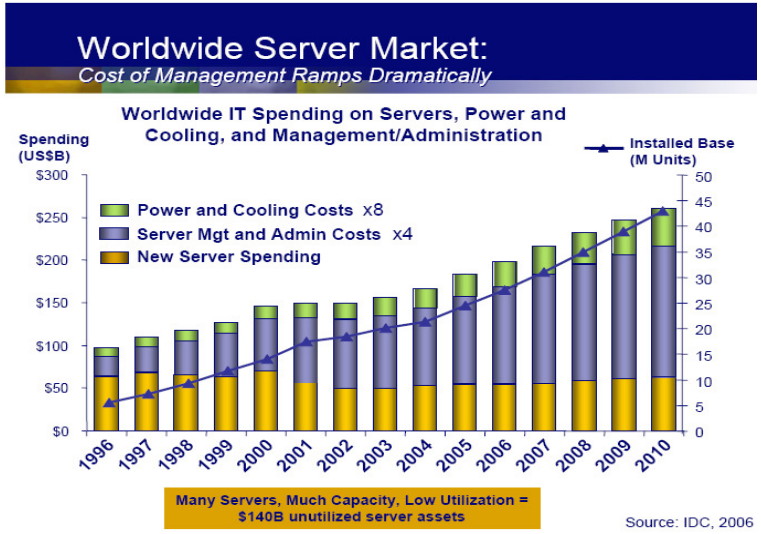

POWER:
System z9 vs. Linux on Intel

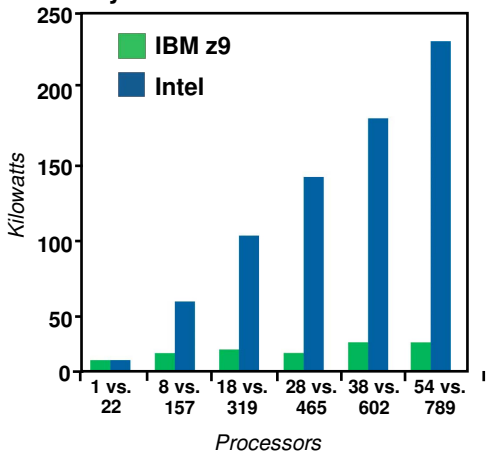

SPACE:
System z9 vs. Linux on Intel

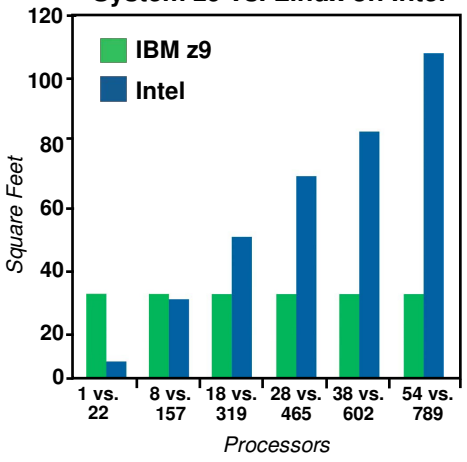

The Linux on Intel servers selected in this example are functionally eligible servers considered for consolidation to a System z running at low utilization such that the composite utilization is approximately 5%. The utilization rate assumed for System z EC is 90%. This is for illustration only actual power and space reductions, if any, will vary according to the actual servers selected for consolidation.

System z Virtualization: a Multidimensional Solution

Virtualization Support is Built In, Not Added On

With coordinated investments in the virtualization technology stack

- **Application support layer**
  - Open, reliable operating system
  - Virtual server awareness infrastructure
  - Enterprise applications
- **Hypervisor layer (z/VM)**
  - Shared-memory based virtualization model
  - Highly granular resource sharing and simulation
  - Flexible virtual networking
  - Resource control and accounting
  - Server operation continuity (failover)
  - Server maintenance tools and utilities
- **Hardware layer**
  - Legendary reliability, scalability, availability, security
  - Logical partitioning (LPAR)
  - Processor and peripheral sharing
  - Interpartition communication
  - Virtualization support at the hardware instruction level

Extreme Virtualization with z/VM V5.3

z/VM can provision virtual machines with a mix of real and virtual resources with exceptional levels of scalability, availability and security

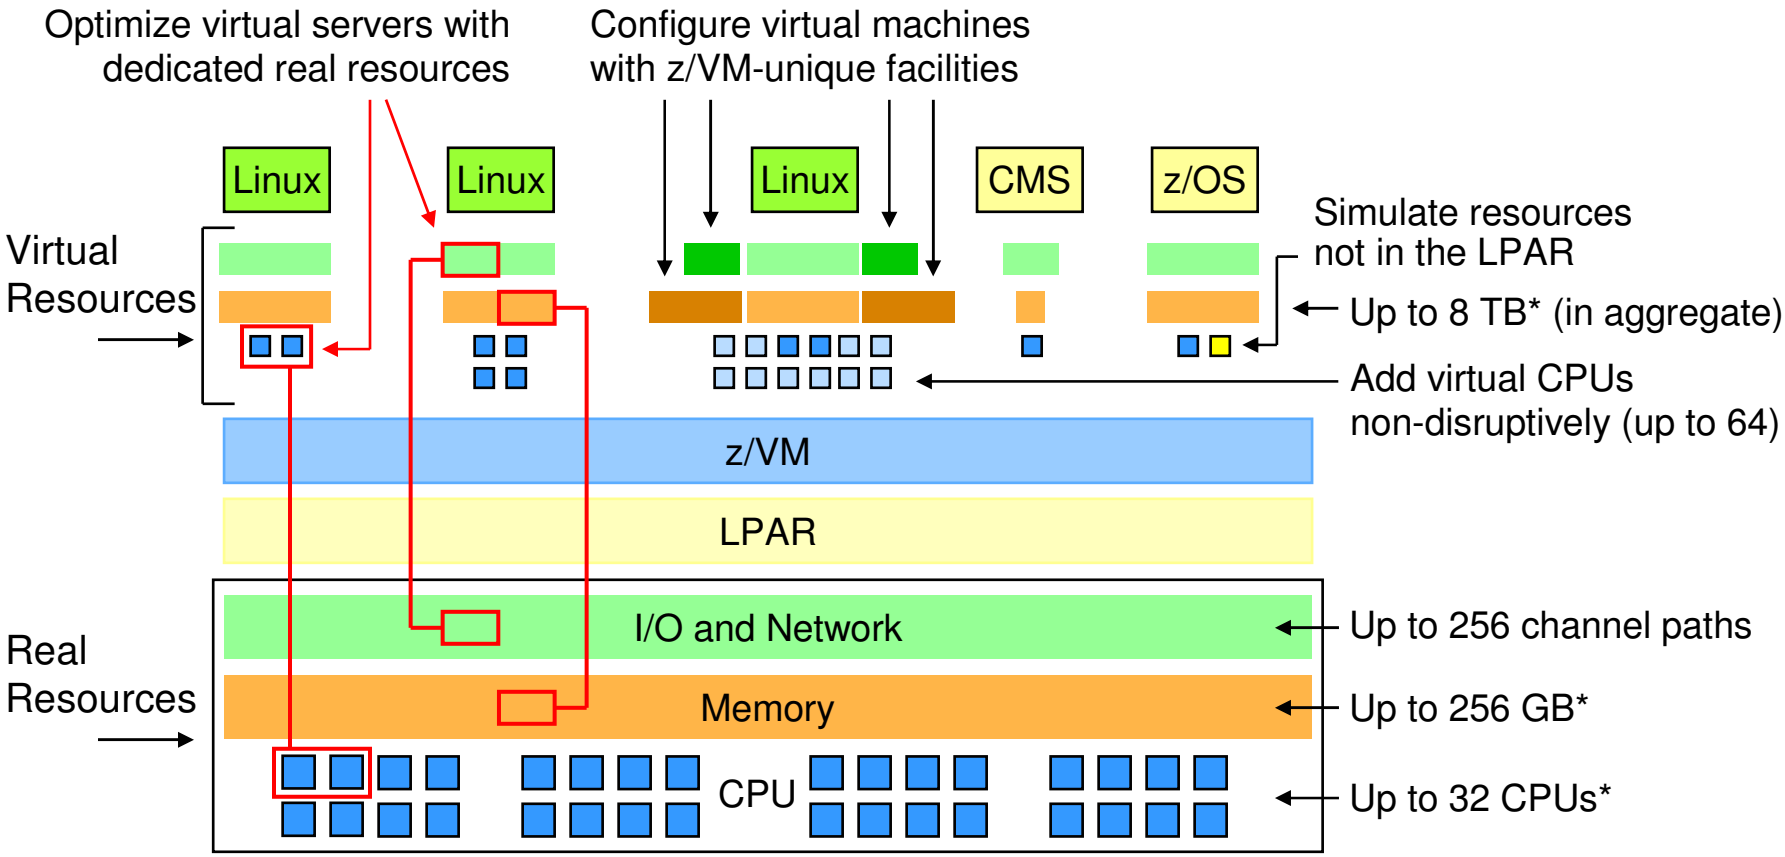

* z/VM V5.3 maximums

z/VM Virtual Switch Link Aggregation Support

Enhanced Networking Bandwidth and Business Continuance

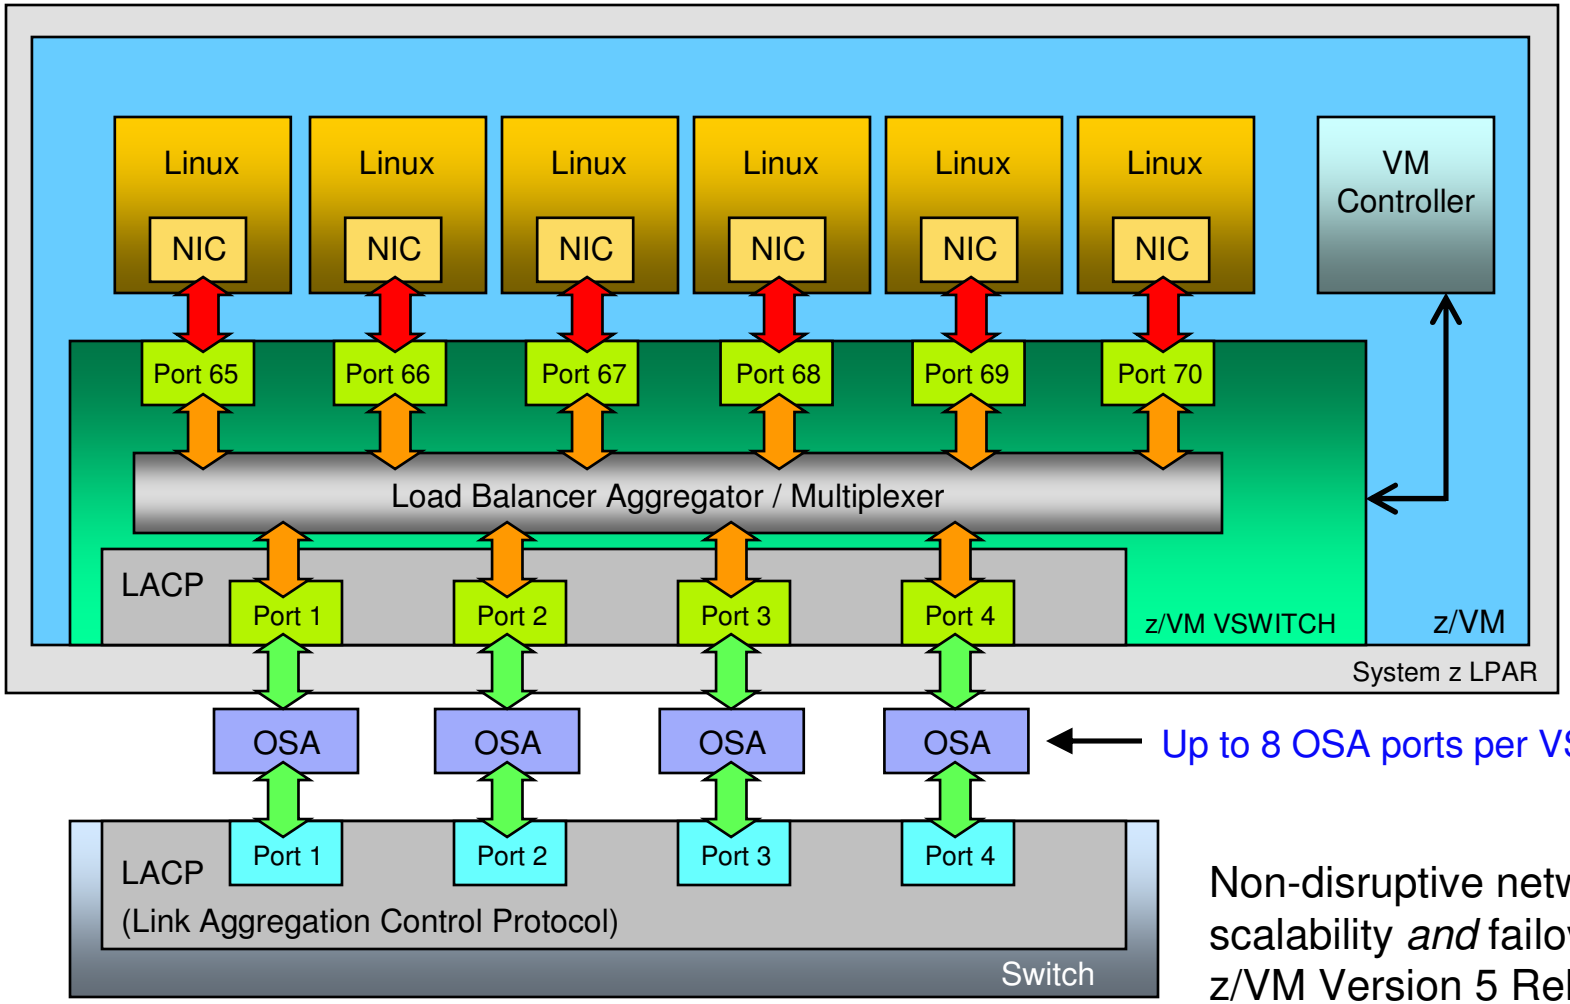

Note: Requires OSA-Express2 support available with IBM System z9 servers

Non-disruptive networking scalability *and* failover with z/VM Version 5 Release 3

Sample z/VM IFL Configuration

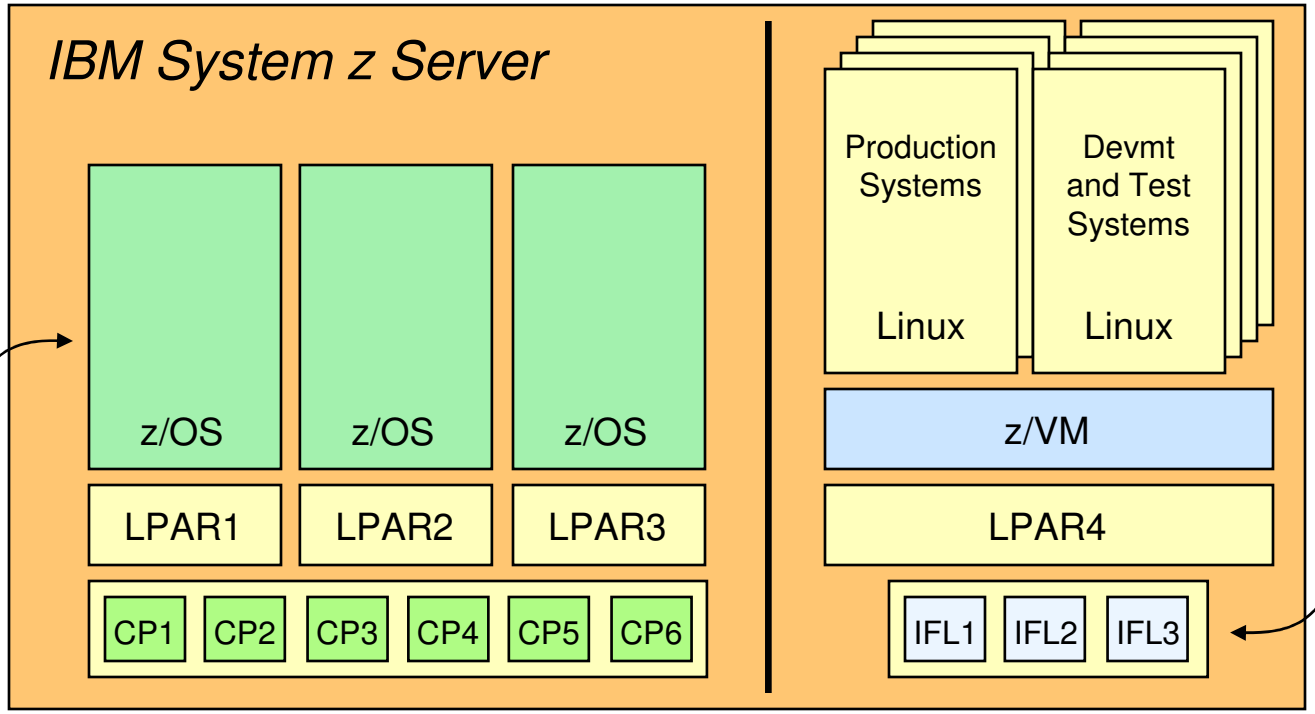

z/VM and most Linux software fees are priced on real engine capacity...

IFL engines have no impact on z/OS license fees

3-engine z/VM V5 license charges*

Year 1:	$84,390	OTC plus S&S
Year 2:	$16,890	S&S only
Year 3:	$16,890	S&S only
3-Year Total:	$118,170	

...another source of cost savings attributed to z/VM's ability to over-commit CPU capacity

*U.S. prices as of 1 Oct 2007

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**

Functional Comparison of z/VM and VMware ESX

Attribute	z/VM V5.3	VMware ESX 3	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, 8 TB of "active virtual memory"	Up to 32 CPUs, 64 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, 16 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 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; VM re-boot for 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	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*; CAPP/EAL 3+	No I/O virtualization separation; CAPP/EAL 2	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 Technology – Command and Control Infrastructure

Leveraging the IBM Software Portfolio

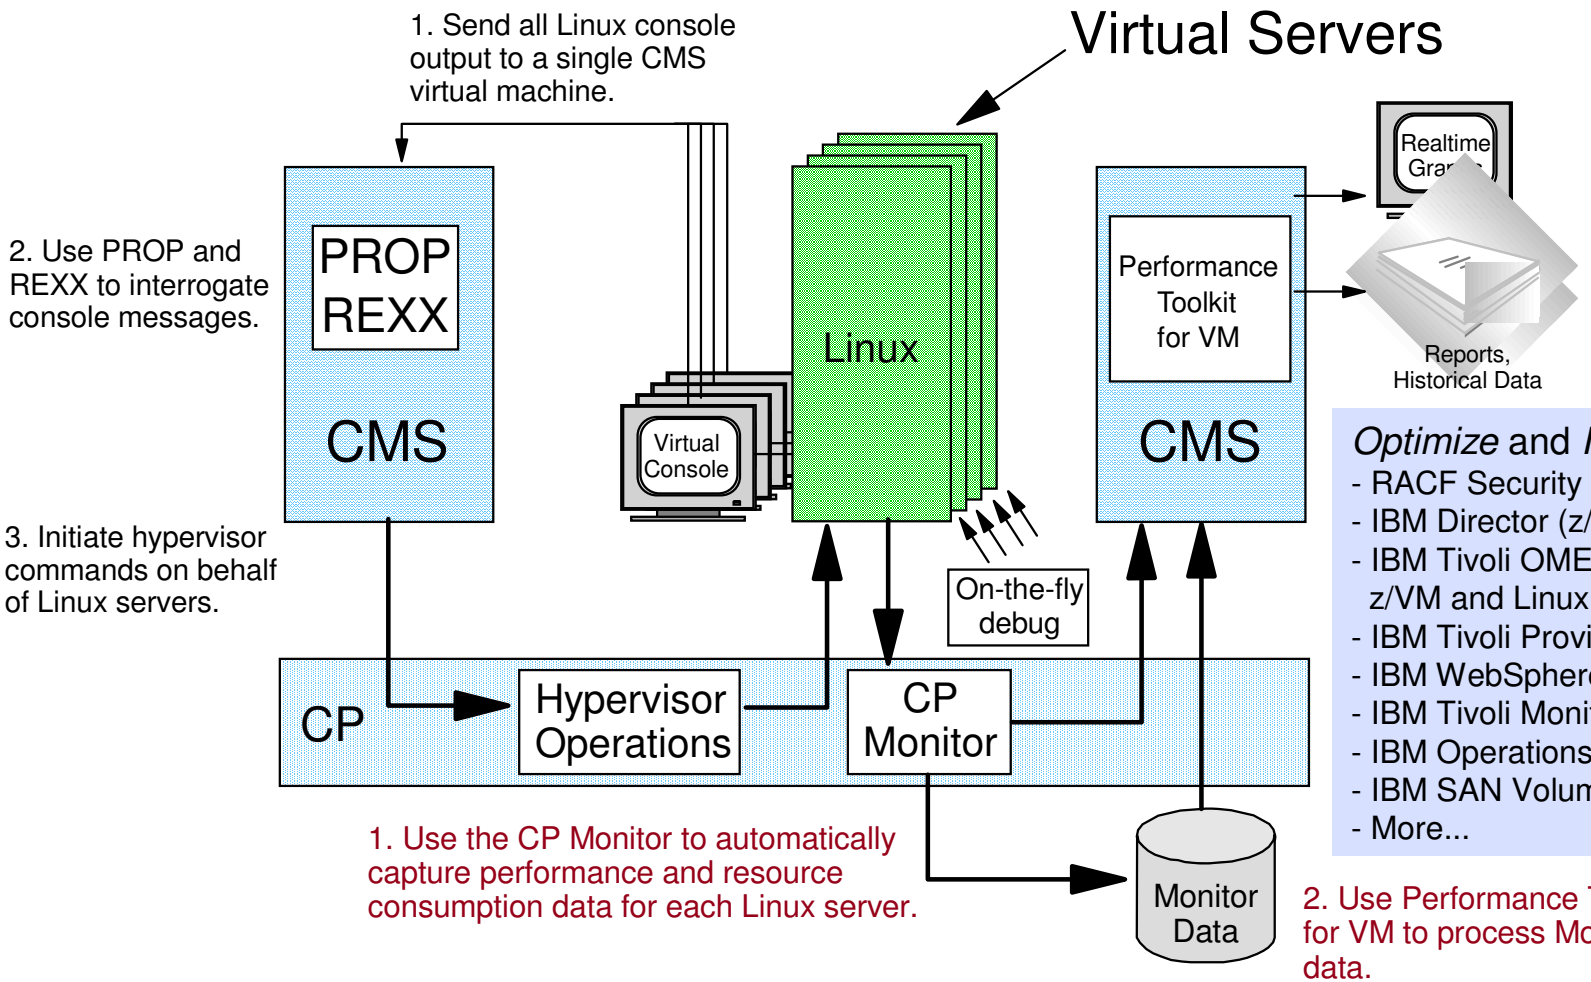

Optimize and *Integrate* with:

- 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...

Provisioning Linux Virtual Machines on System z

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

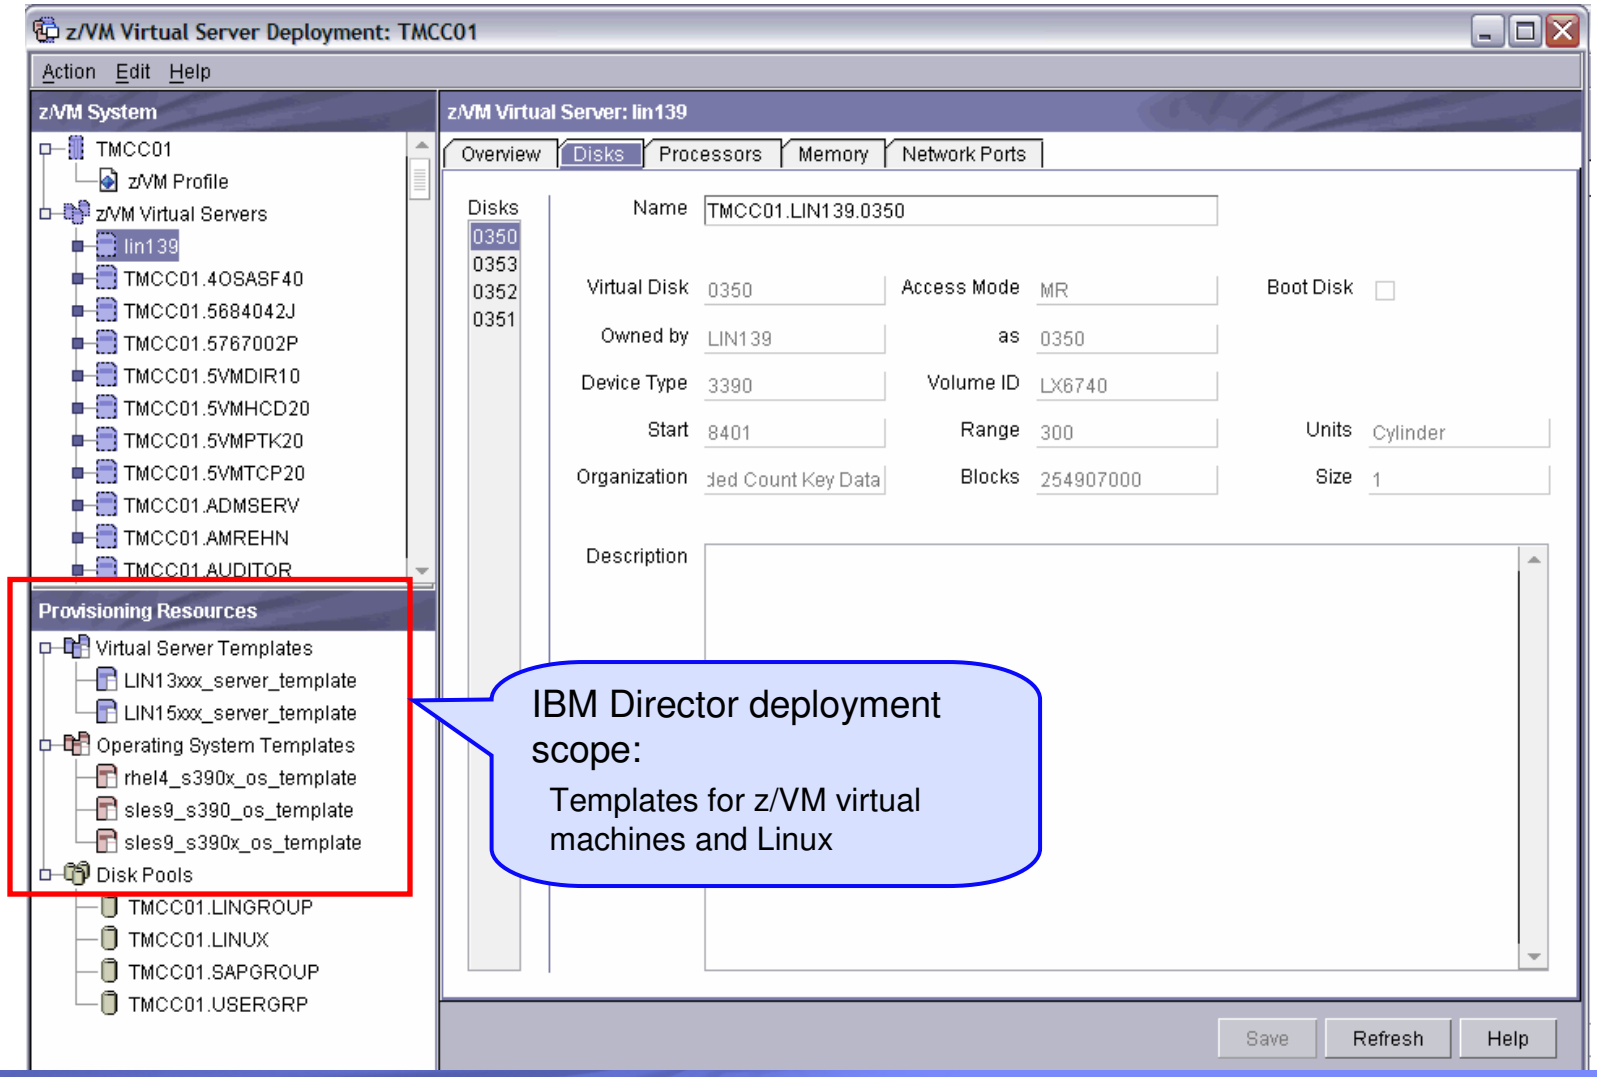

Provisioning Software in System z Virtual Linux Servers

Using IBM Tivoli Provisioning Manager

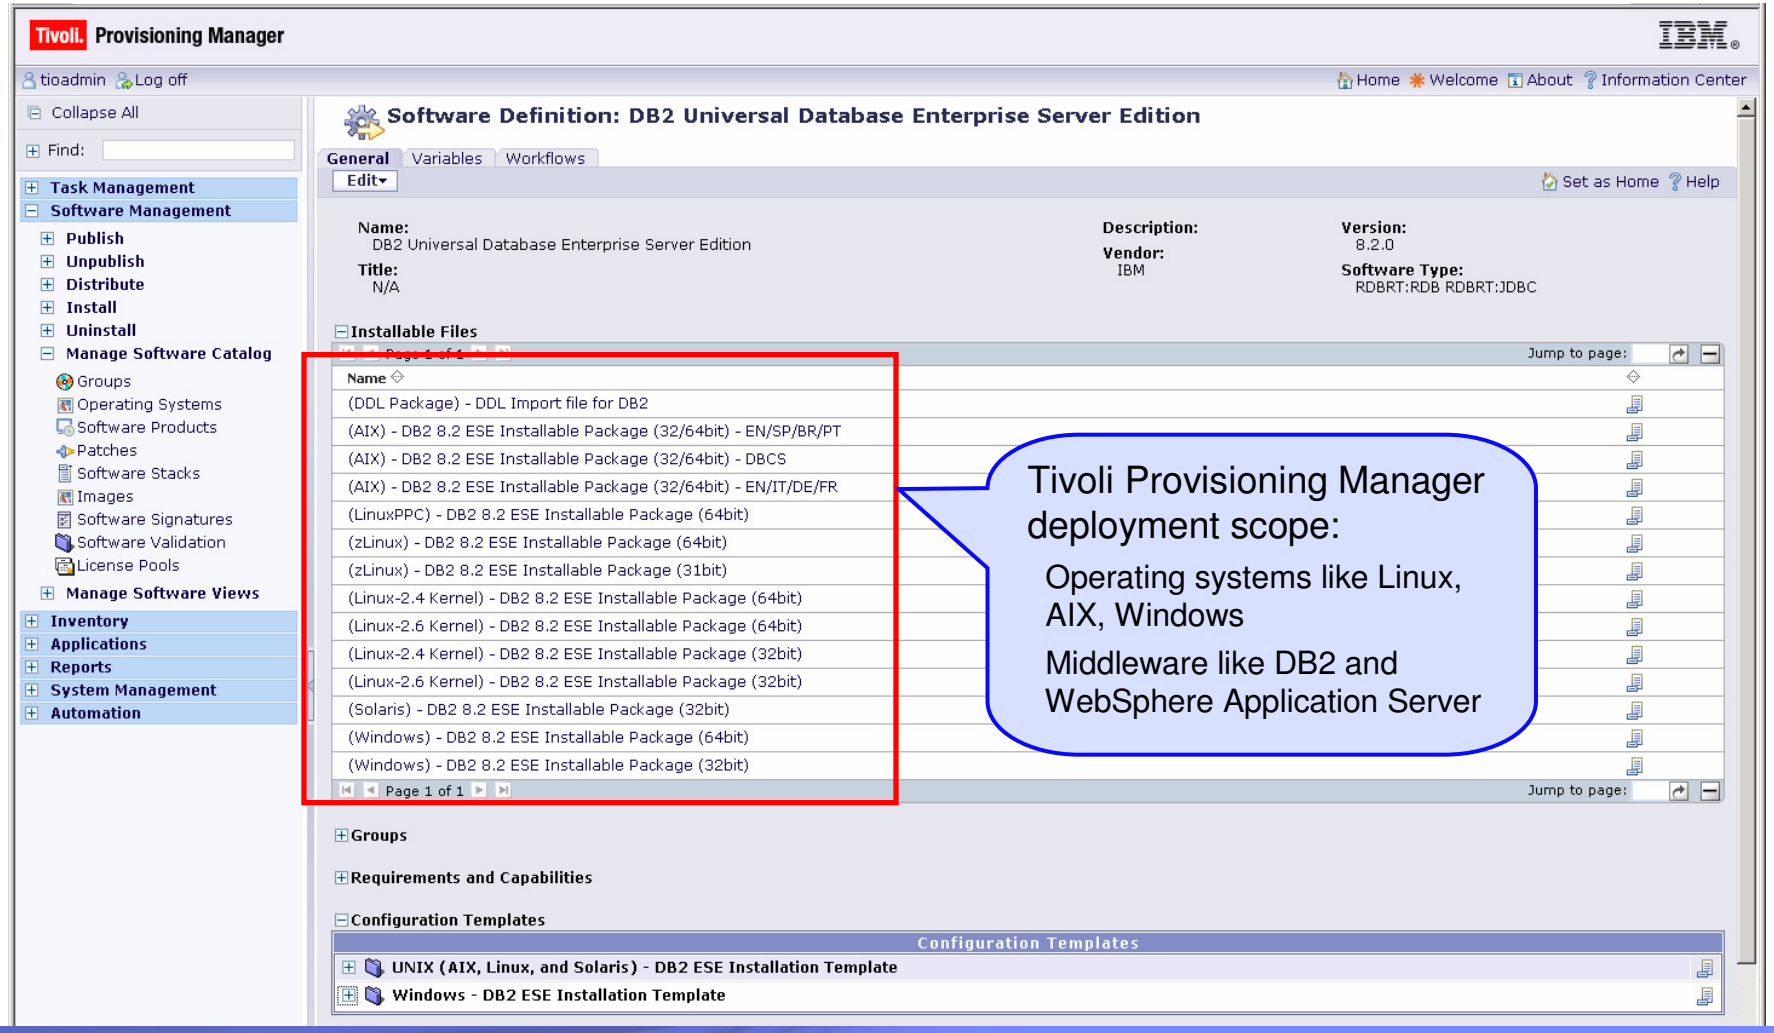

Monitoring System z Virtual Linux Servers
Using IBM Tivoli OMEGAMON XE for z/VM and Linux

- **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**

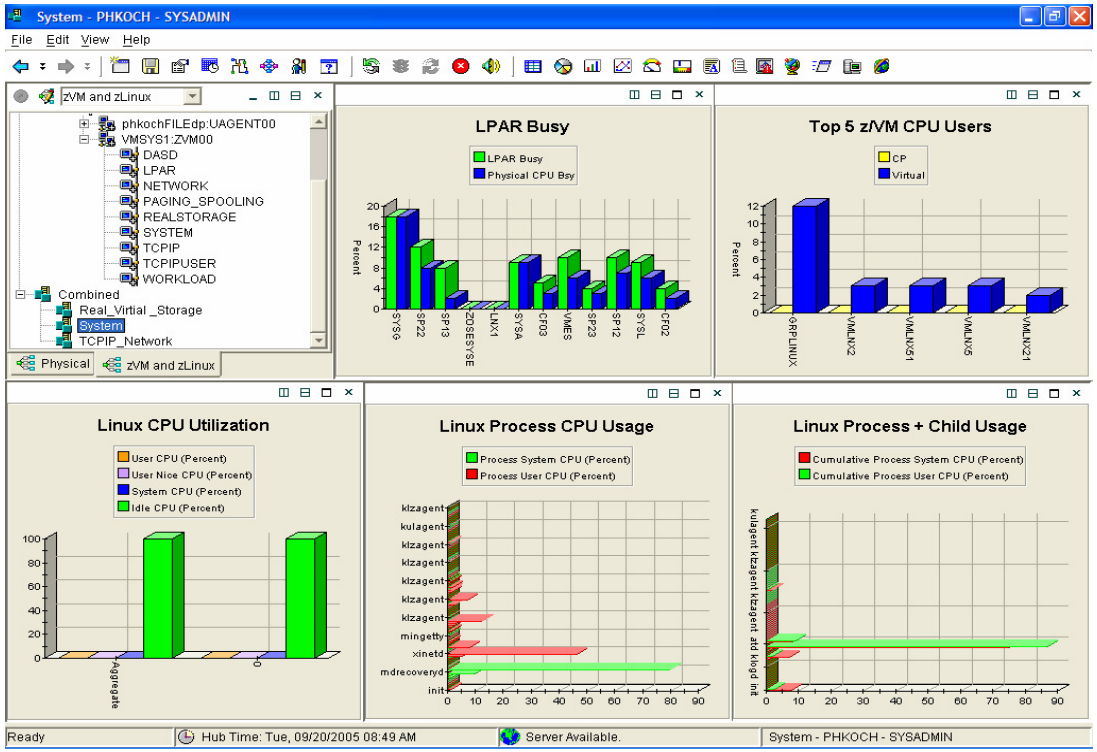

IBM System Storage SAN Volume Controller Software V4.2

- **z/VM and Linux for System z support SAN Volume Controller (SVC) V4.2**
- **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 V4.2:**
  - Multi-target FlashCopy support (up to 16 images)
  - Higher number of active FlashCopy relationships at the cluster level
  - Designed for improved cluster performance, especially when installed on IBM System Storage SVC 2145-8G4 storage engine
  - Support for additional OEM devices
- **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

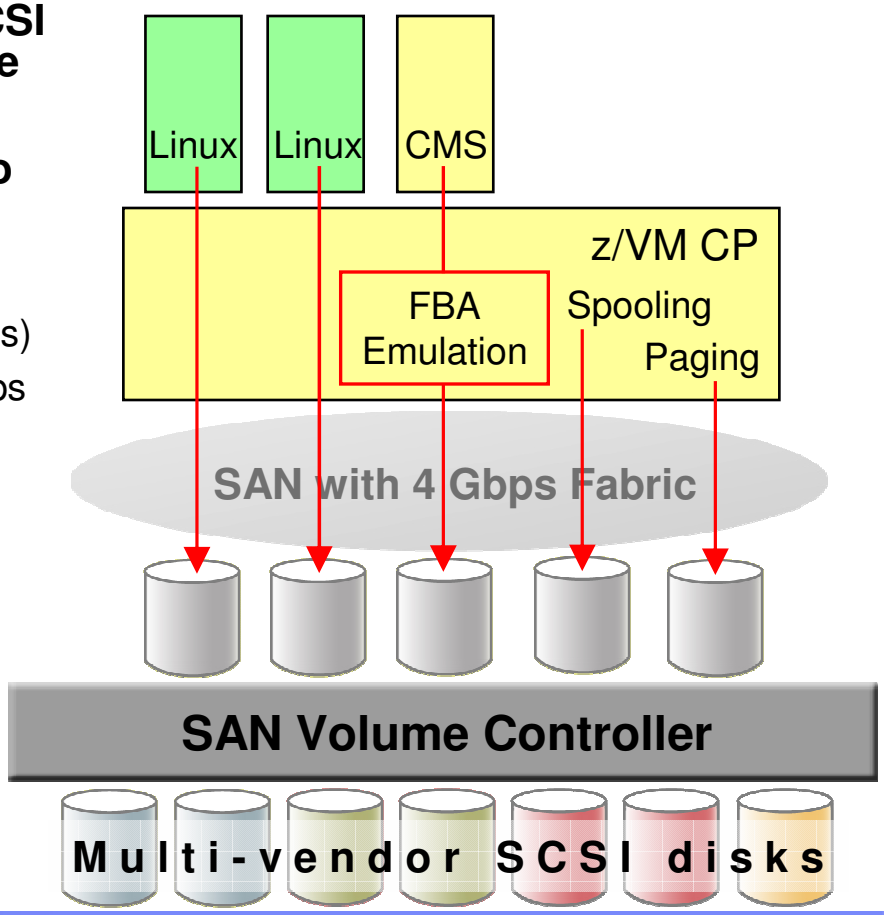

IBM Systems Director

End-to-End Management For the Enterprise

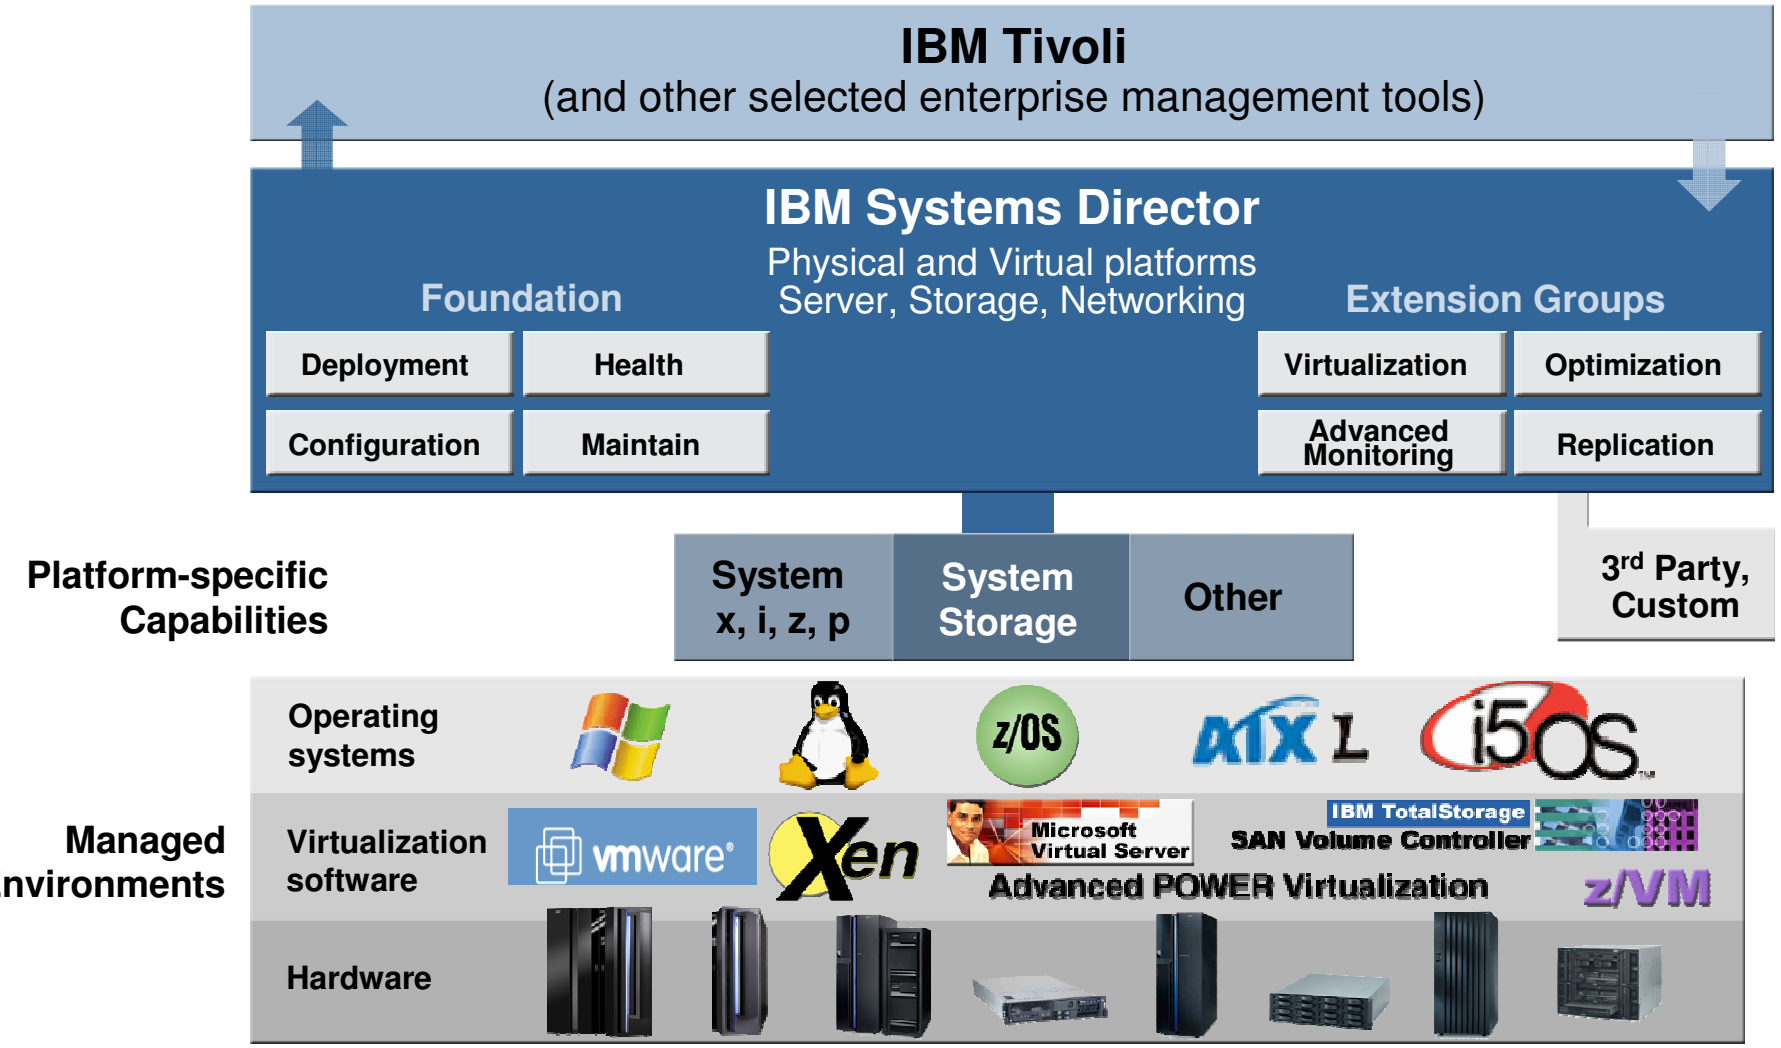

Extreme Virtualization with System z
Understanding the Value Proposition

- **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

Clients need to develop an enterprise-wide virtualization strategy that leverages the strengths of mainframe virtualization

Nationwide Saves $15M with System z

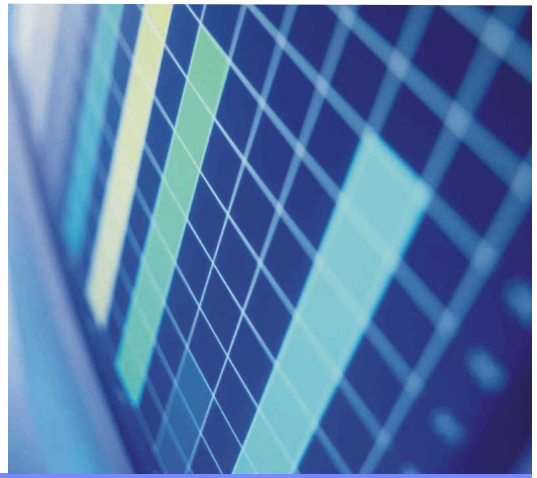

TCO: $15M savings over 3 years

- 80% reduction in data center floor space needs; power conservation
- 50% reduction in hardware and OS support efforts
- 70% average CPU utilization

Dynamic allocation of compute power

- Capacity on Demand
- Tested 22 times the capacity for 2006 Super Bowl Ad blitz traffic

Source: Guru Vasudeva, AVP & Chief Architect, Nationwide Insurance LinuxWorld August 2006 presentation

Nexxar

"The company you trust to send money worldwide"

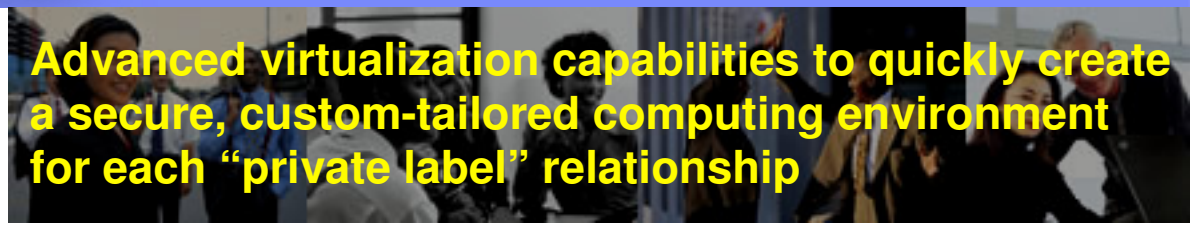

Business Need

An IT infrastructure that provides very high (24x7) availability and is able to sustain significant business growth

Key Benefits (Value Proposition)

- ✓ An architecture that suits requirements for security, manageability, reliability, availability, scalability, extensibility and flexibility
- ✓ The ability to help Nexxar's growth-by-acquisition business while staying within the same platform
- ✓ Consolidation of more than 80 x86 servers onto an IBM System z9 Business Class (BC) machine
- ✓ *A 75% reduction of headcount required to maintain the operating environment in comparison with x86 systems previously on the floor*

NEXXAR GROUP

Solution

- ❑ Hardware: IBM System Storage
- ❑ Software: z/OS, DB2, z/VM, Linux
  WebSphere Application Server
  Tivoli OMEGAMON
  Rational
- ❑ Services: GTS Infrastructure & Systems Management Services

Client Deploys Additional Workloads on System z

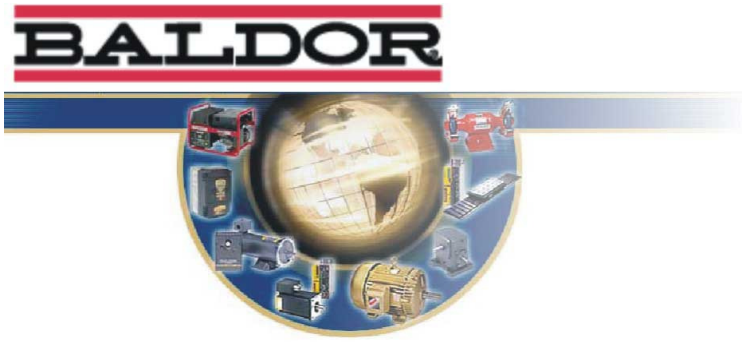

***Computerworld* published a story comparing two "growing midsize companies" ***

Baldor Electric consolidated several UNIX-based servers onto one IBM eServer™ zSeries® 990, deploying all of its SAP Enterprise Portal, Supply, and Business Warehouse solutions on zSeries and Linux.

Both have "... similar size IT departments"; Both "... use packaged ERP applications ... and want complete alignment with the business." ***However*, "when it comes to the hardware running these systems, the companies are polar opposites"**

	Baldor Electric implemented SAP using Linux on System z, z/VM, & DB2 on z/OS & is spending less than 1% of sales on I/T.	Welch's Food implemented Oracle ERP on Dell using VMware, Oracle DB, and Linux and is spending 2.5% of sales on I/T.
Supplier	**IBM**	**Dell**
Moved from:	3 Mainframes and 8 UNIX Servers	S/390® and AS/400®
Moved to:	1 z990 + Integrated Facility for Linux (IFLs)	100 Intel® Servers
Solution	DB2 database runs on IBM z/OS and SAP applications run in 24 Linux virtual machines on the same server	Oracle ERP on Dell using VMWare, and Oracle DB using Linux
Decision-to-Completion Time	Approximately 6 months	Started sometime before June 2005, project will continue into 2007
IT Staff	Down to 38	50
IT Spending	1.2% of Sales in first year of implementation	About 2.5% of Sales

* *Computerworld,* August 15, 2005

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 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 3+ certification and intends to pursue EAL 4 certification, offering system solutions that have been methodically designed, tested, and reviewed for secure operations

Questions?

Backup Material

IBM System z Virtualization Architecture

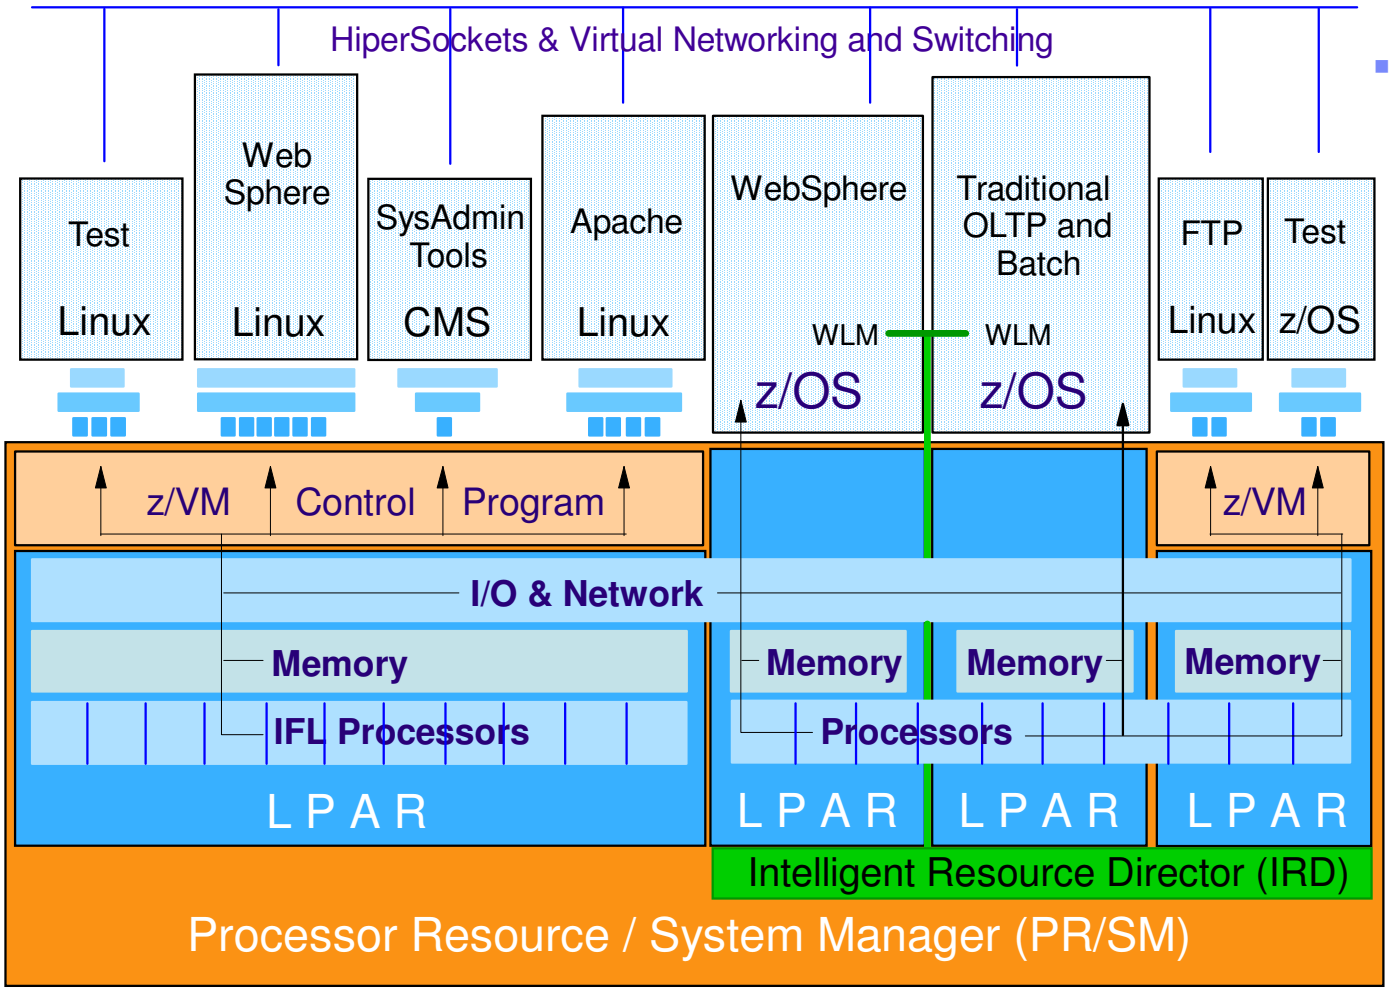

- **Multi-dimensional virtualization technology**
  - System z provides logical (LPAR) and software (z/VM) partitioning
  - PR/SM enables highly scalable virtual server hosting for LPAR *and* z/VM virtual machine environments
  - IRD coordinates allocation of CPU and I/O resources among z/OS and non-z/OS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

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

IBM System z Virtualization Genetics

The virtual machine concept on the mainframe has a long history of software *and* hardware innovations

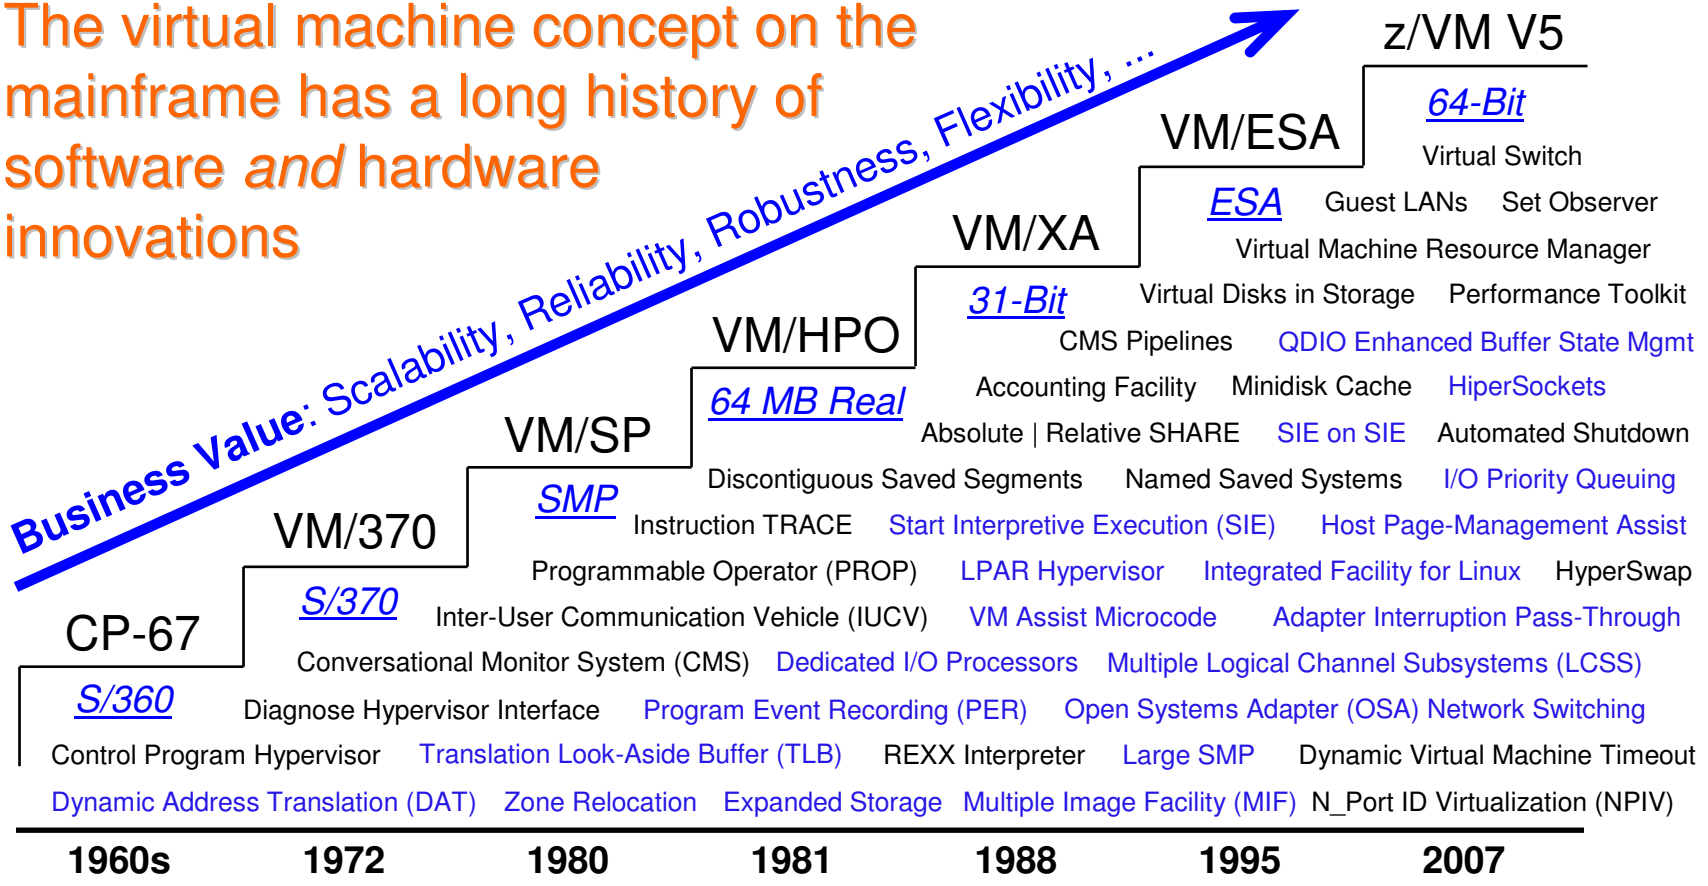

System z virtualization starts on the chip; an integration of hardware, firmware, and software functionality

Chip Design Affects Virtualization Capabilities

Replicated Server Chip Design

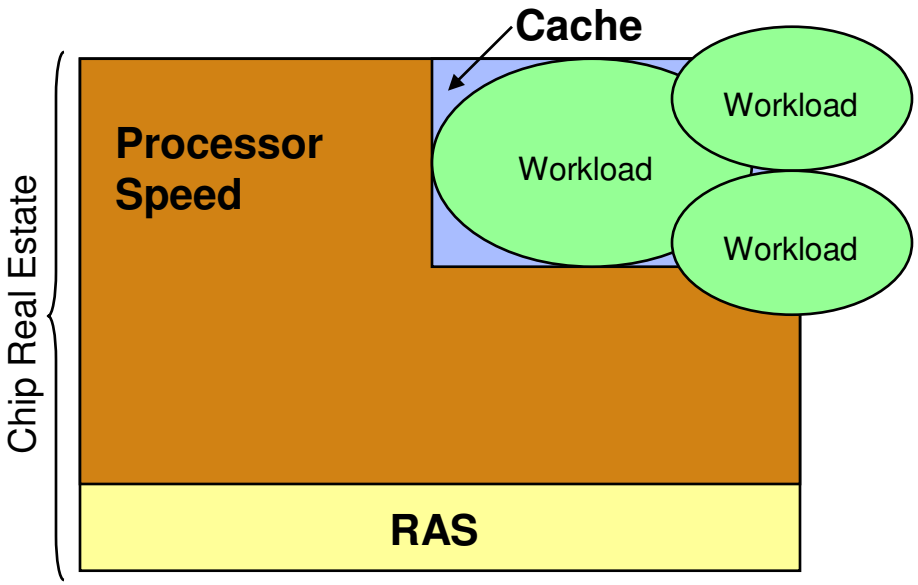

- Mixed workloads stress cache usage, requiring more context switches
- Working sets may be too large to fit in cache
- "Fast" processor speed is not fully realized due to cache misses

Consolidated Server Chip Design

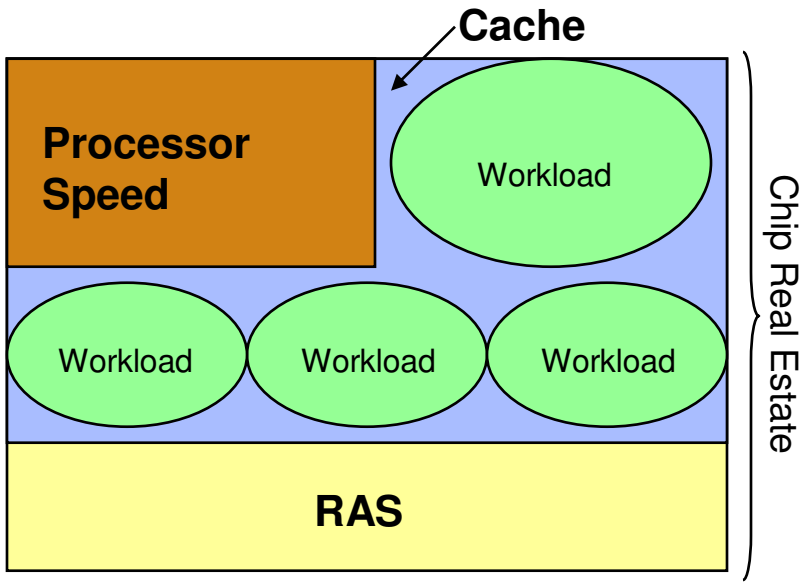

- System z cache is able to contain more working sets
- Processor speed is optimized by increased cache usage
- Additional RAS function is beneficial for mixed workloads

System Design Affects Virtualization Capabilities

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

- With TCO-friendly pricing

Up to 54-way SMP

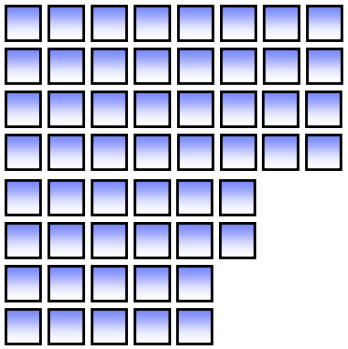

Share up to 54 processors with up to 60 LPARs

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

* No software license fees

Up to 10 System Assist Processors

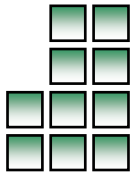

Offload system processing to dedicated CPUs with no impact to software license fees

Plus up to 336 I/O Processors

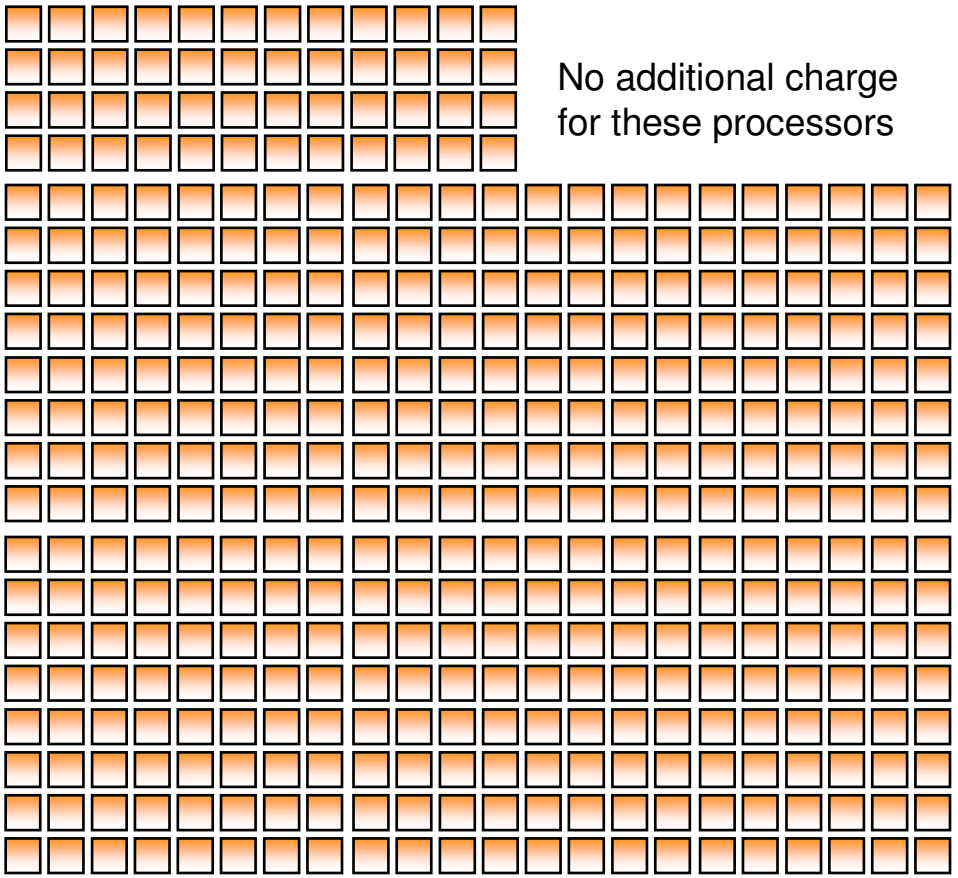

No additional charge for these processors

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"

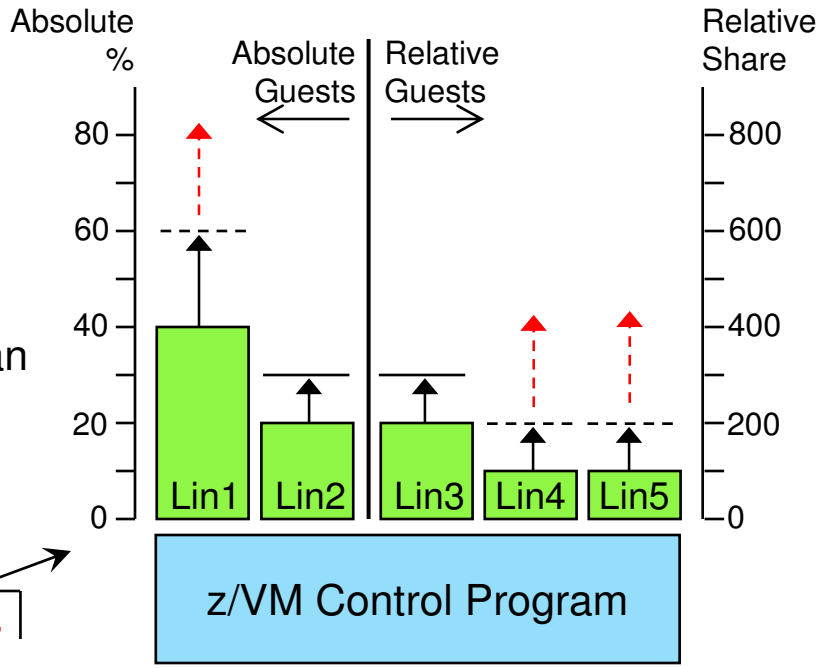

z/VM Directory Entries (or "on-the-fly" commands)

```
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)

PR/SM High-Performance I/O Sharing (MIF)

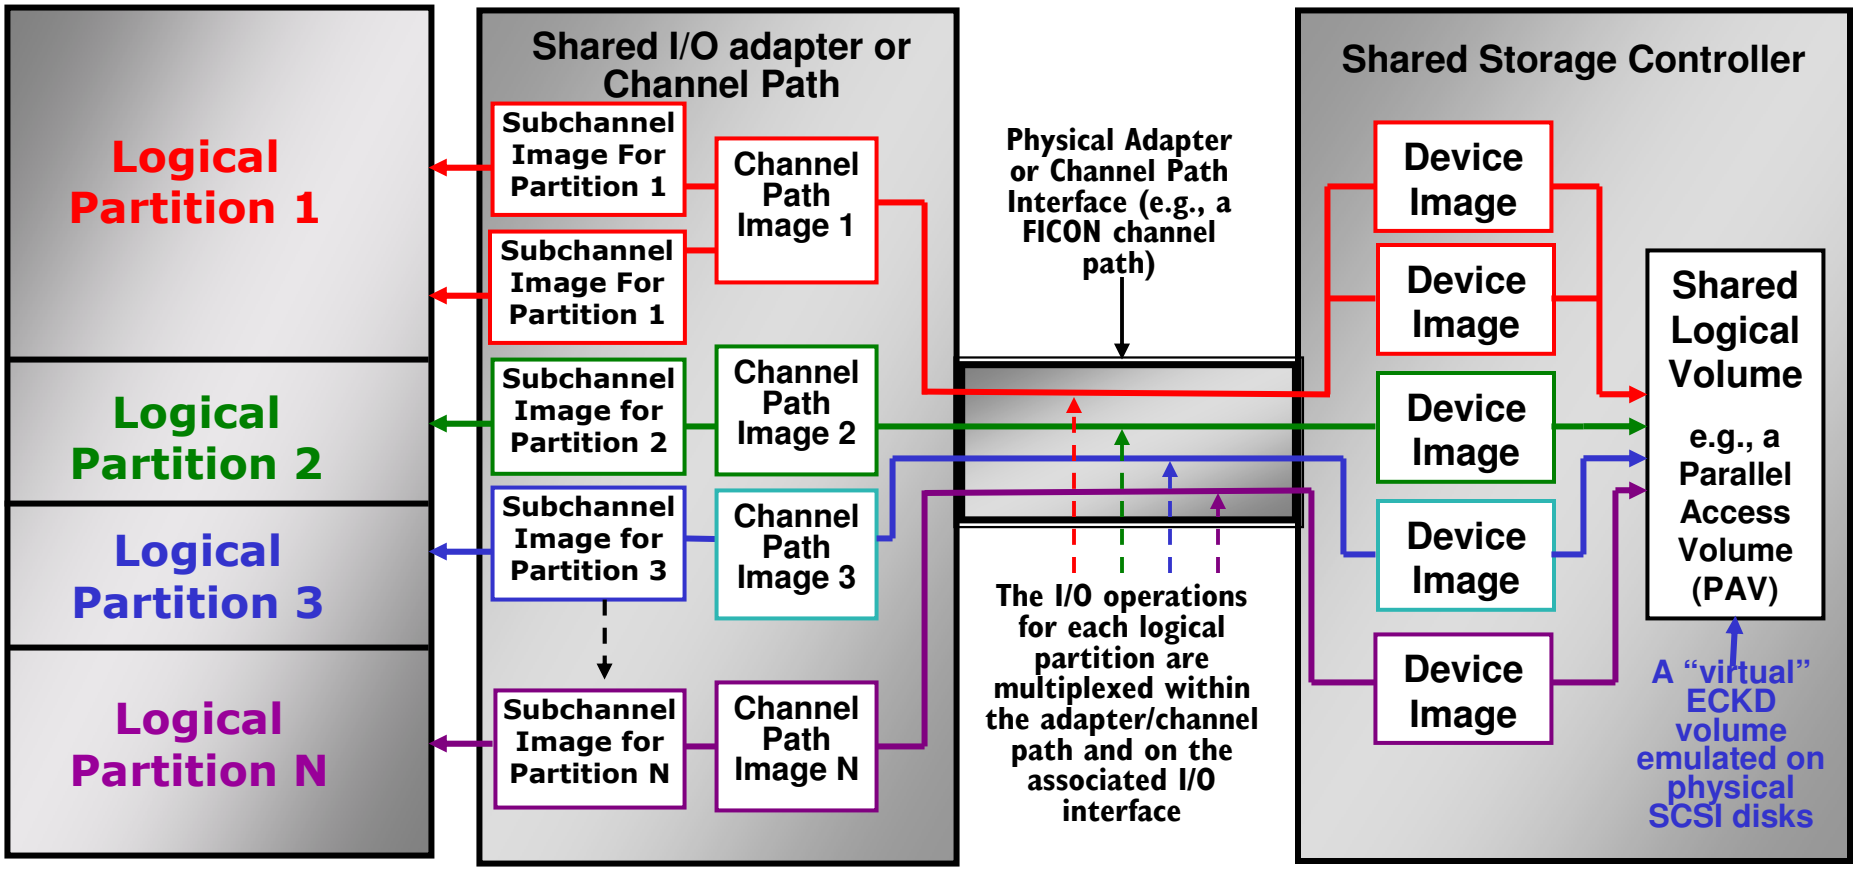

- **The I/O infrastructure (adapters/channels, their transmission links, and attached I/O resources) are shared by logical partitions at native speeds (without hypervisor involvement)**
  - I/O requests, their associated data transfers and I/O interruptions flow between each logical partition OS instance and the shared I/O components just as if the I/O components were physically dedicated to a single logical partition
  - Dynamic paths enables up to 8 physical channels (either dedicated or shared) to process the I/O requests to the shared devices; reduces possibility of I/O queuing delays at the channels or at the shared storage controller

System z and N_Port ID Virtualization (NPIV)

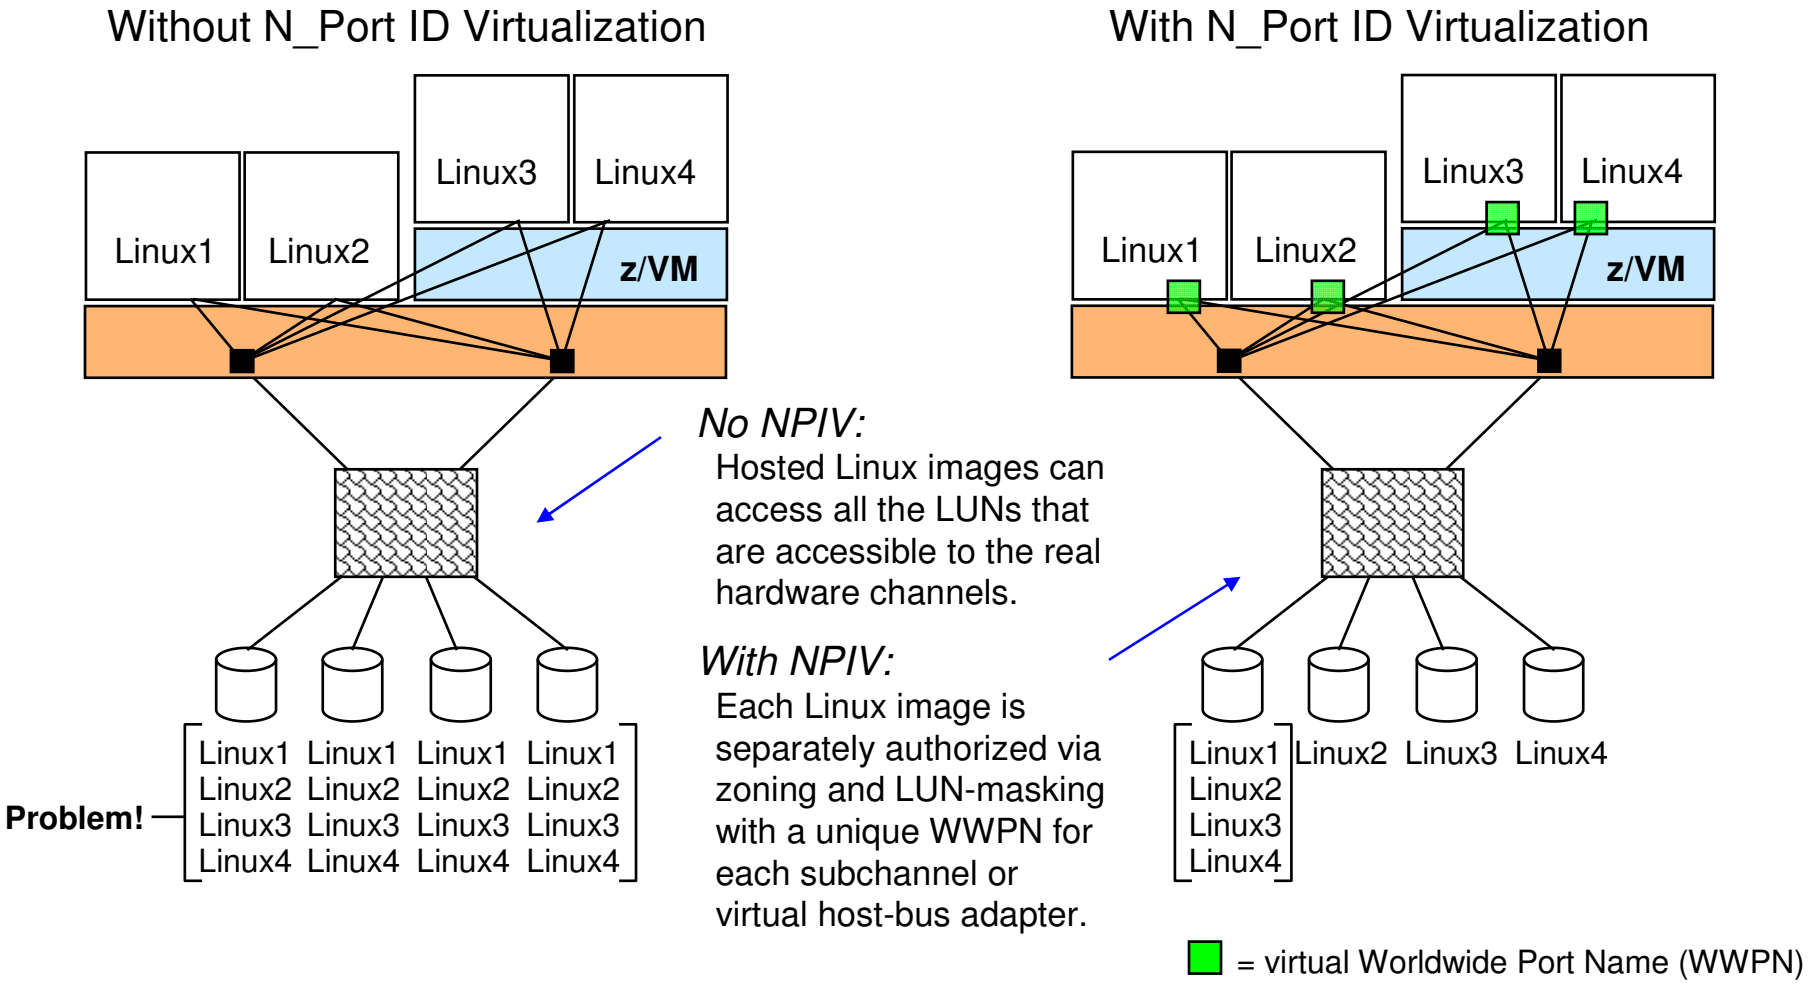

Linux and z/VM Technology Exploitation
Cooperative Memory Management (CMM)

- Problem scenario: virtual memory utilization far exceeds real memory availability
- z/VM Control Program paging operations become excessive
- Overall system performance and guest throughput suffers

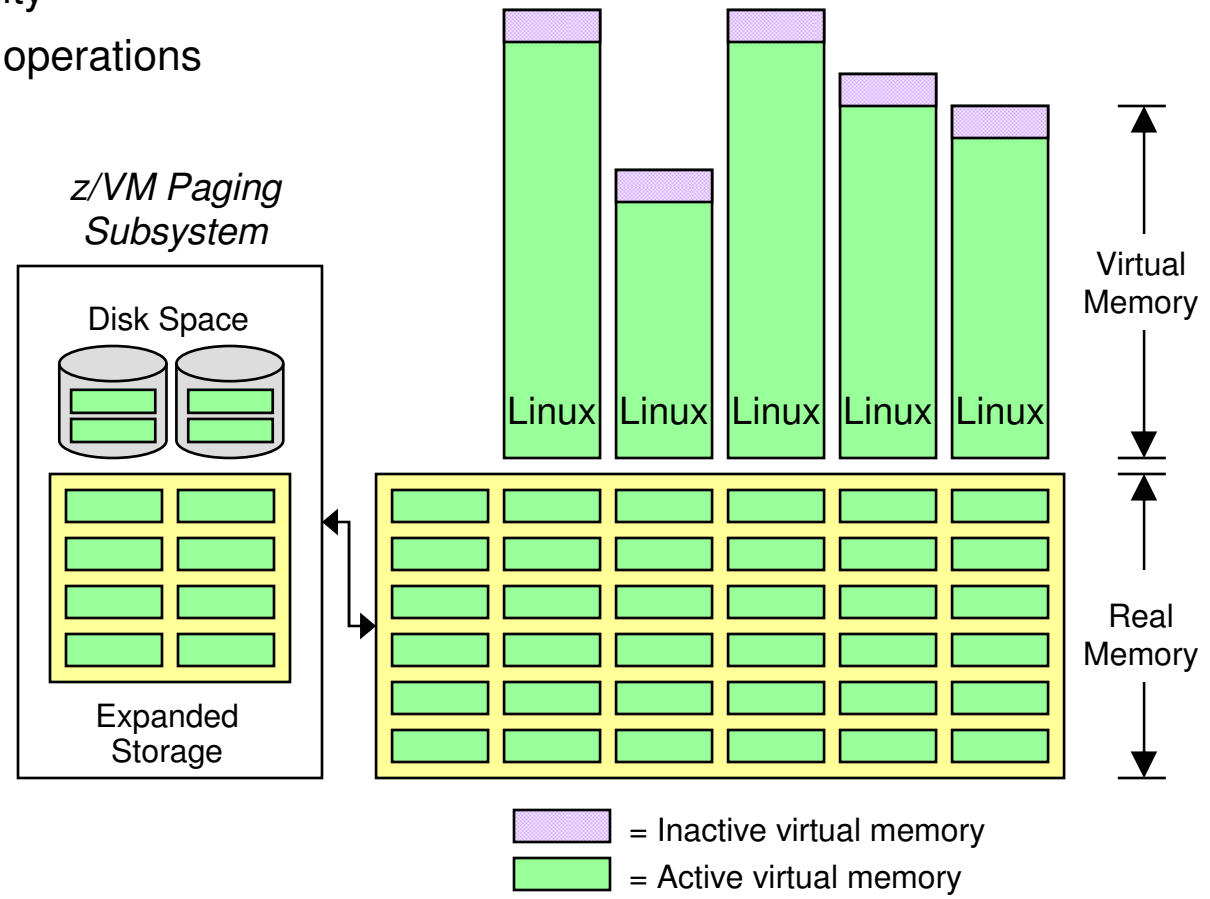

Linux and z/VM Technology Exploitation
Cooperative Memory Management (CMM)

- Solution: real memory constraint detected by z/VM *Virtual Machine Resource Manager*
- Linux images signaled to reduce virtual memory consumption
- Linux memory pages are released
- Demand on real memory and z/VM paging subsystem is reduced
- Helps improve overall system performance and guest image throughput

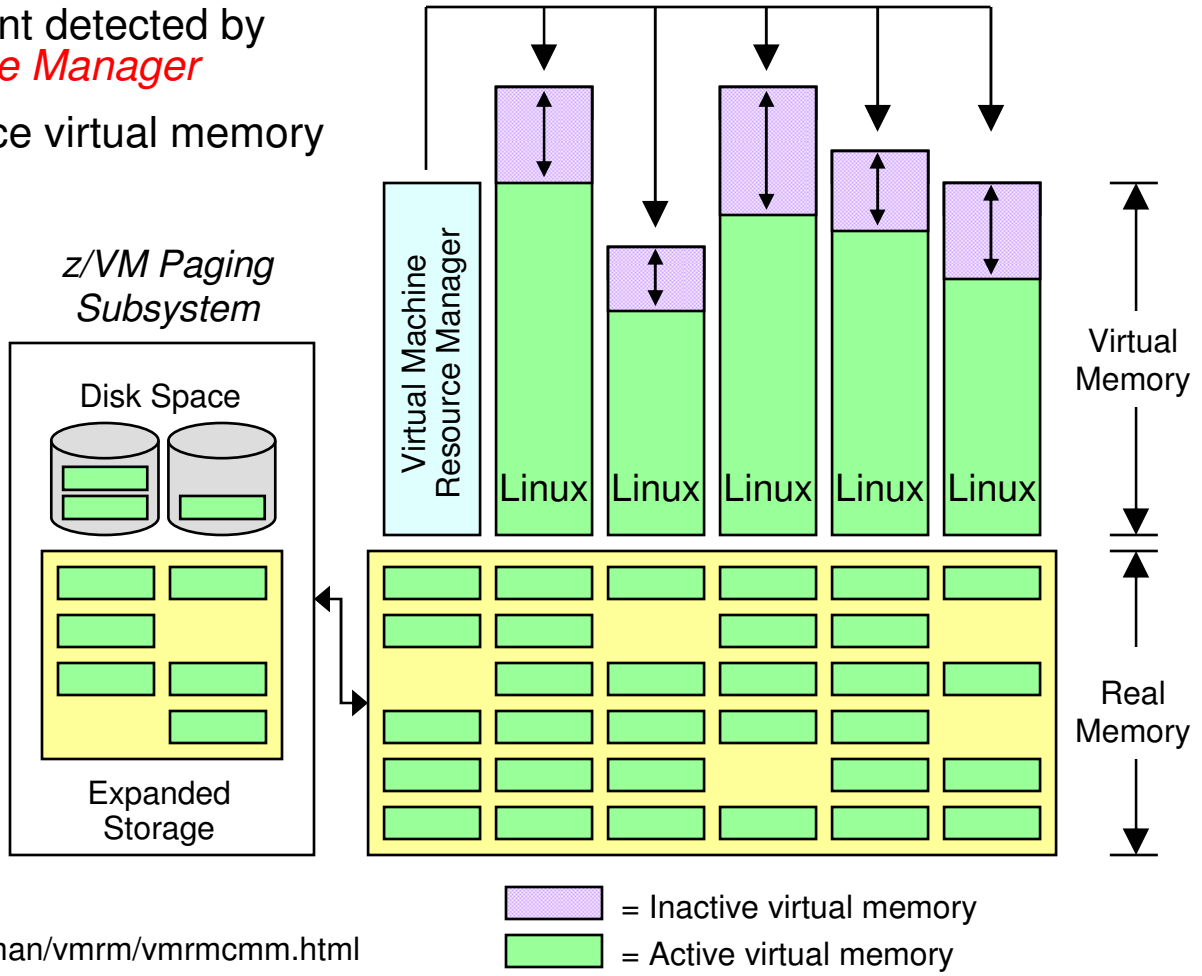

Learn more at:
ibm.com/servers/eserver/zseries/zvm/sysman/vmrm/vmrmcmm.html

Cooperative Memory Management with Linux on z/VM

Apache Web Serving Example

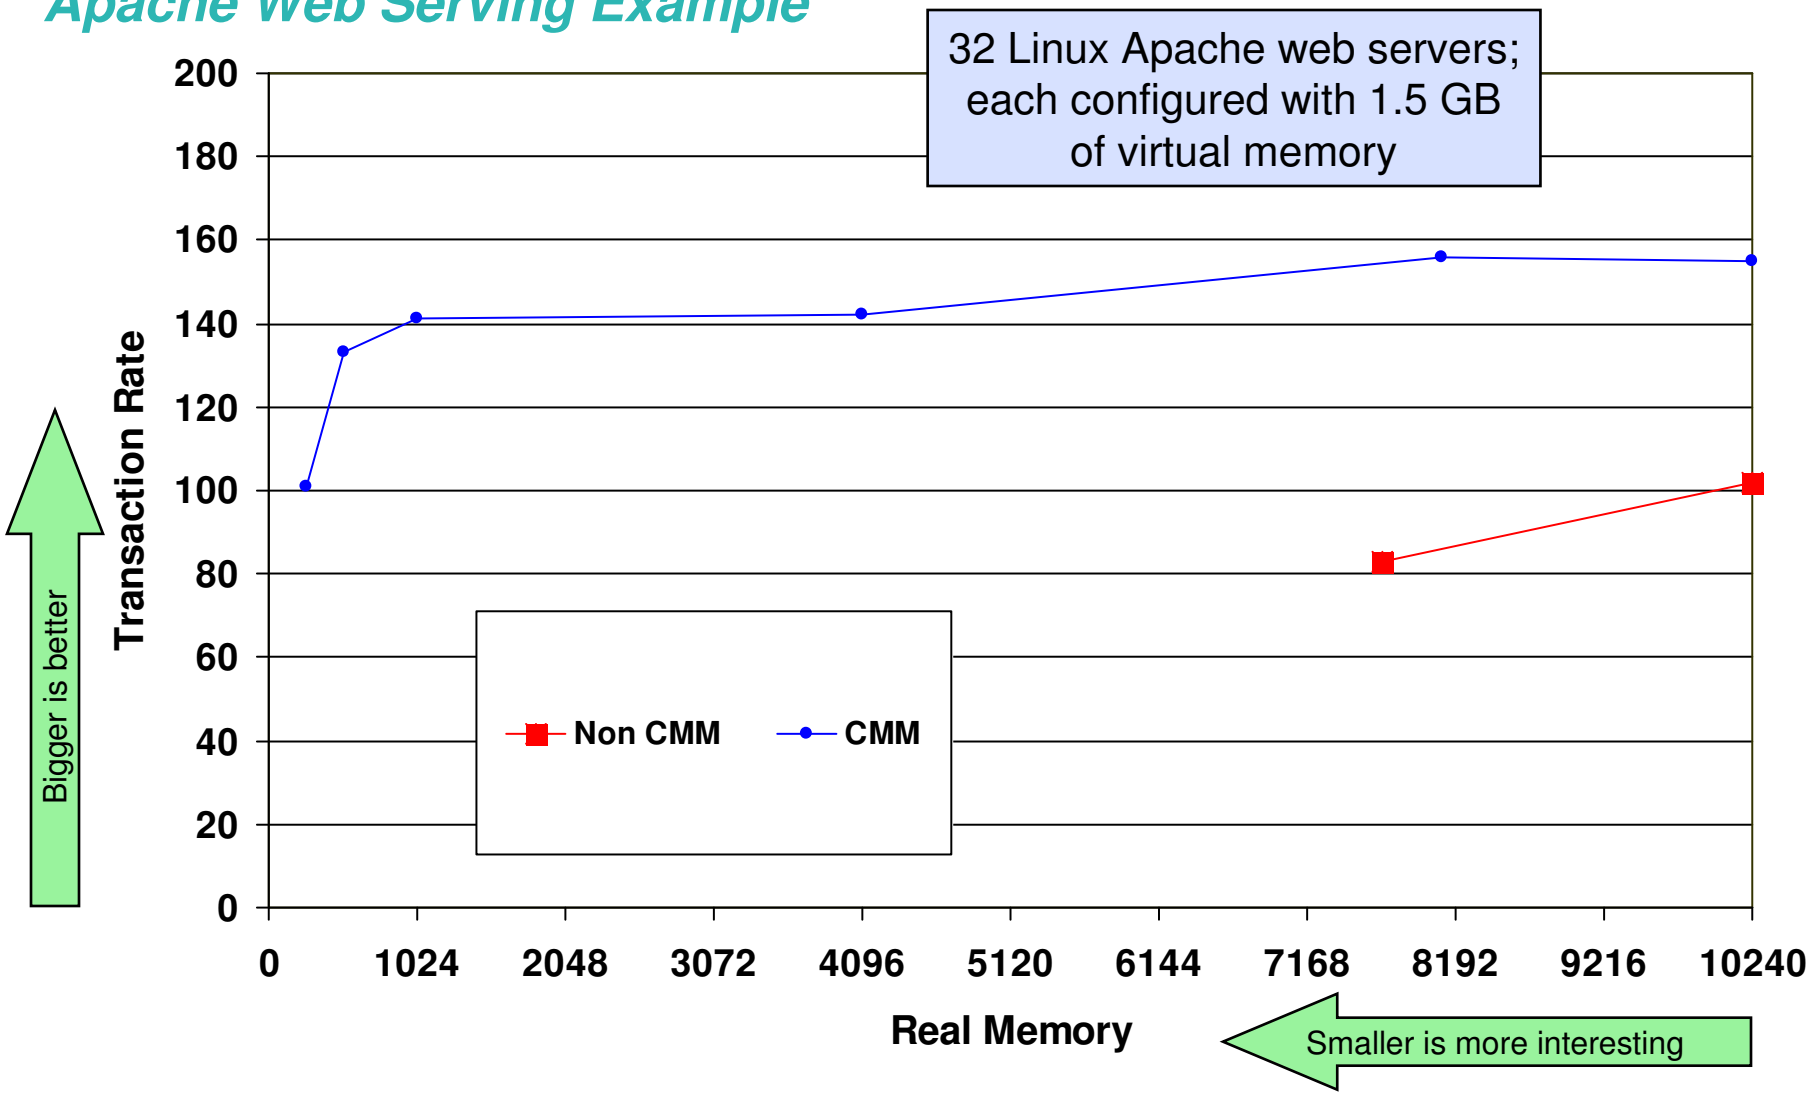

Linux and z/VM Technology Exploitation

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
- IBM is working with its Linux distribution partners for exploitation support

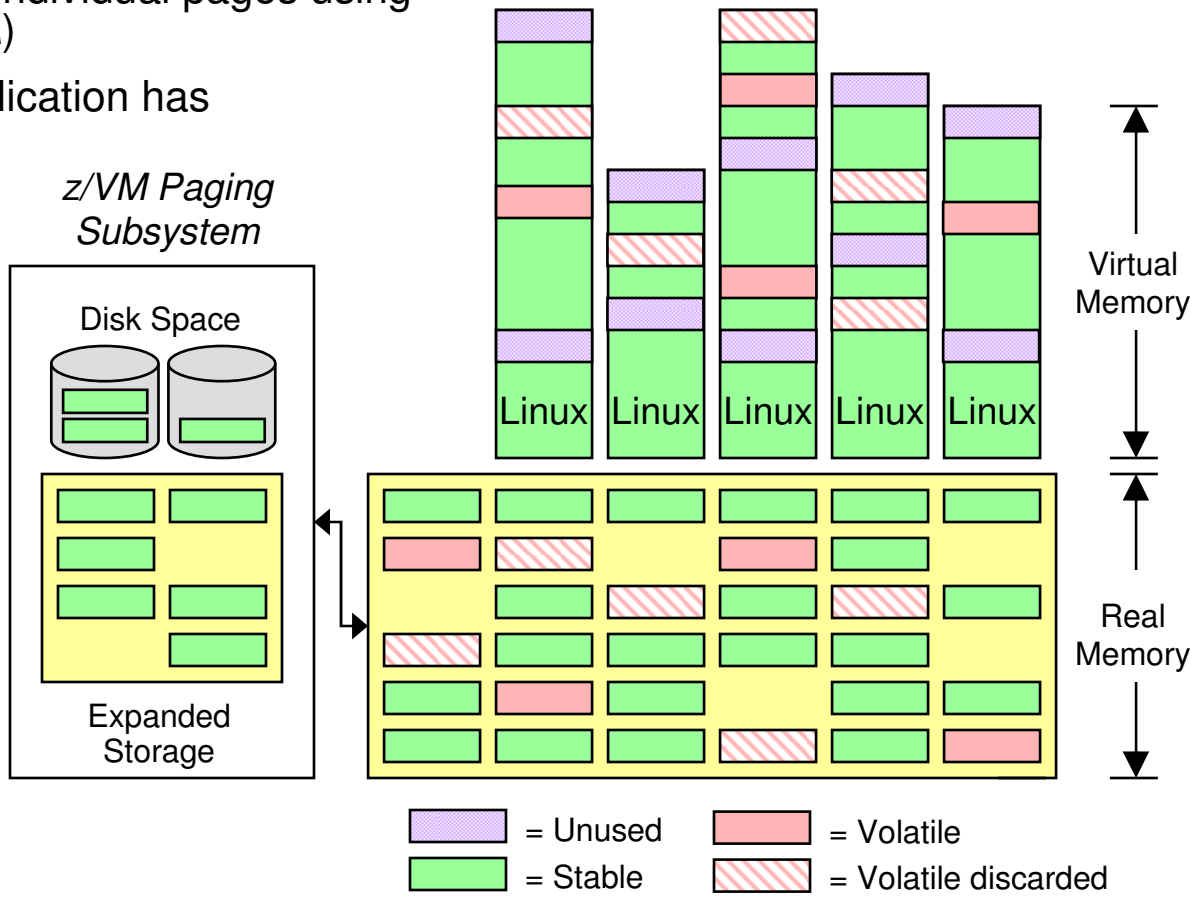

Linux and z/VM Technology Exploitation

Linux Exploitation of z/VM Discontiguous Saved Segments (DCSS)

- DCSS support is Data-in-Memory technology
  - Share a single, real memory location among multiple virtual machines
  - High-performance data access
  - 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 and data stored on disks
  - Helps enhance overall system performance and scalability

Learn more:
"Using DCSS/XIP with Oracle 10g on Linux for System z"
www.redbooks.ibm.com/redpieces/abstracts/sg247285.html

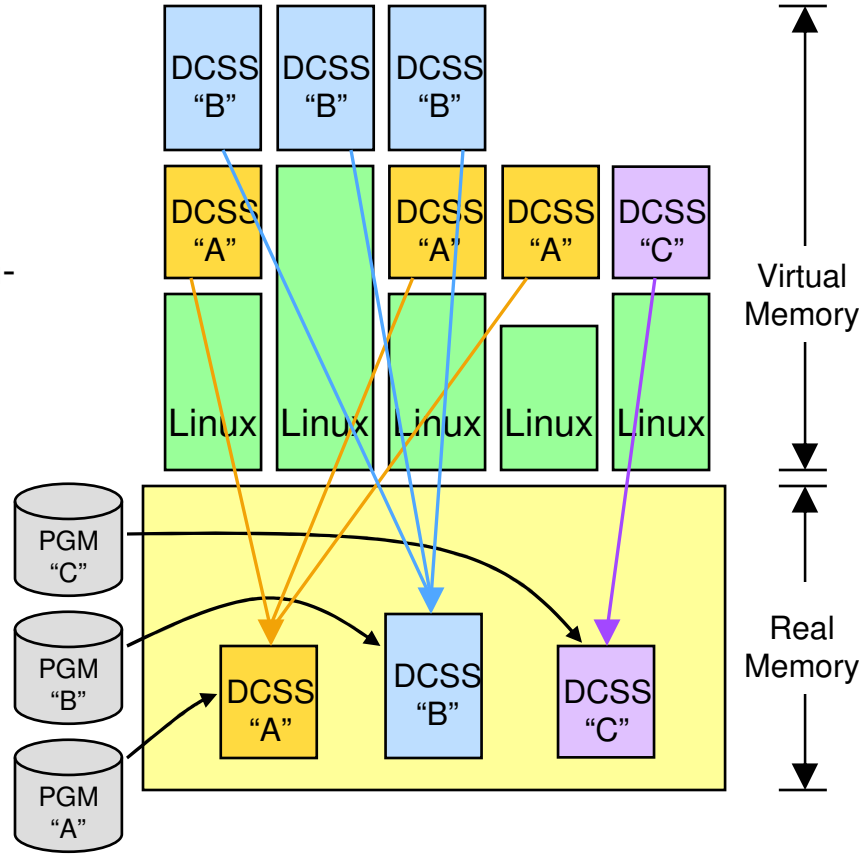

Linux and z/VM Technology Exploitation

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

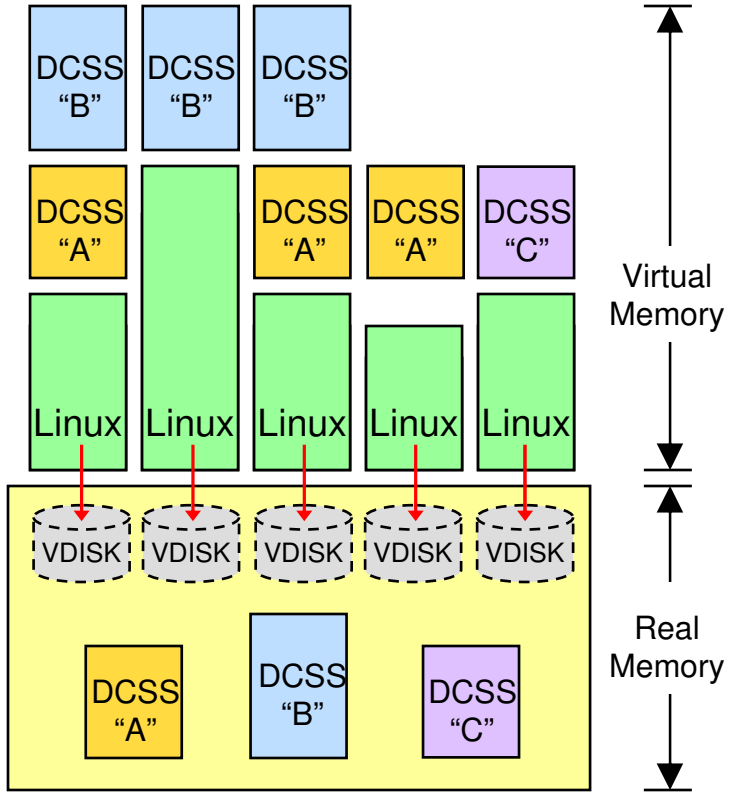

GDPS/PPRC Multiplatform Resiliency for System z

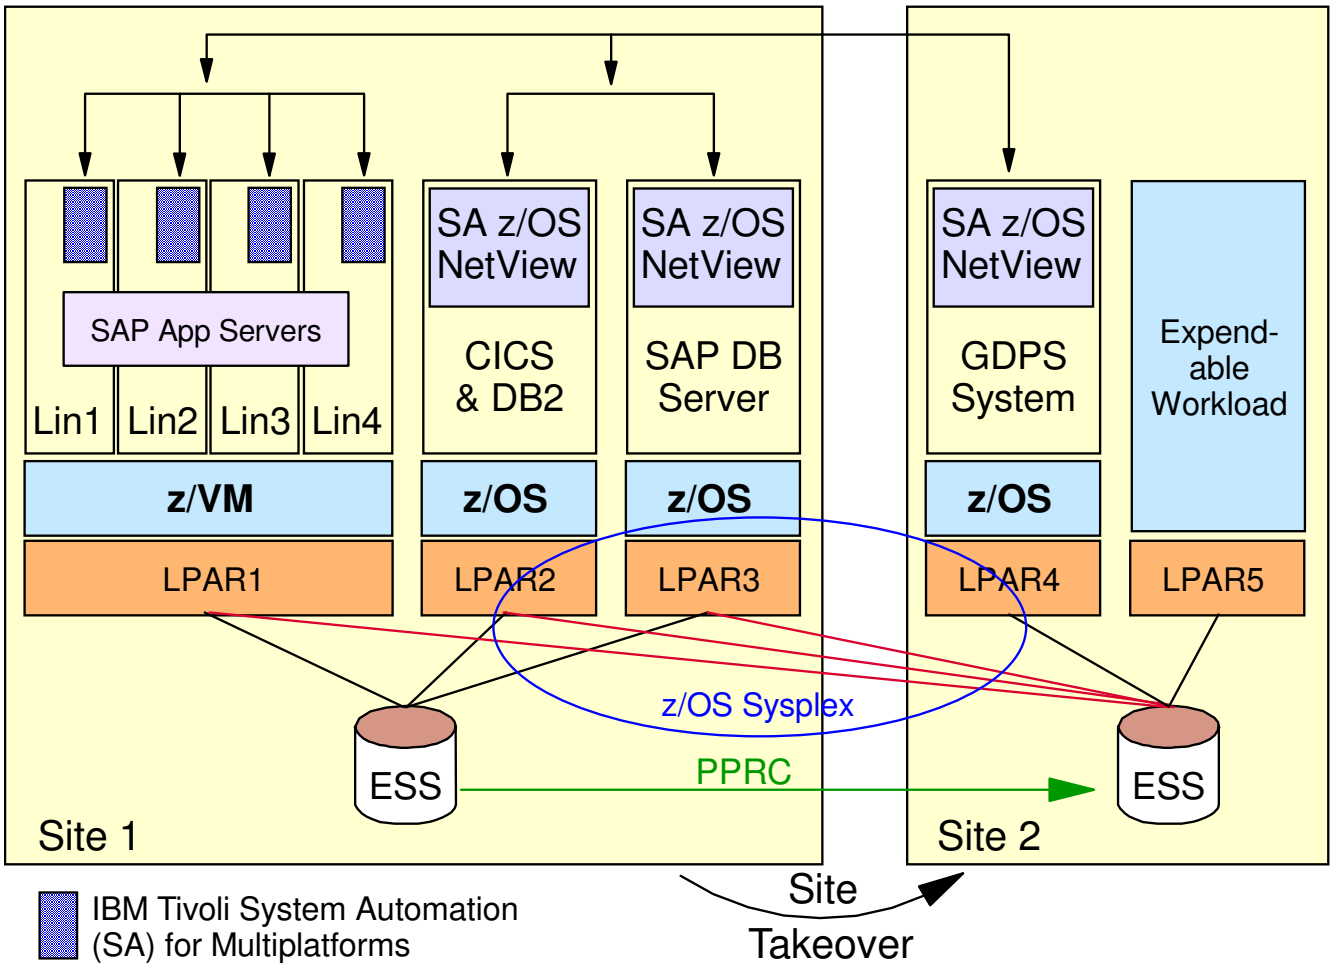

- Designed for customers with distributed applications
- SAP application server running on Linux for System z
- SAP DB server running on z/OS
- Coordinated near-continuous availability and DR solution for z/OS, Linux guests, and z/VM
- Uses z/VM HyperSwap function to switch to secondary disks
- Sysplex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

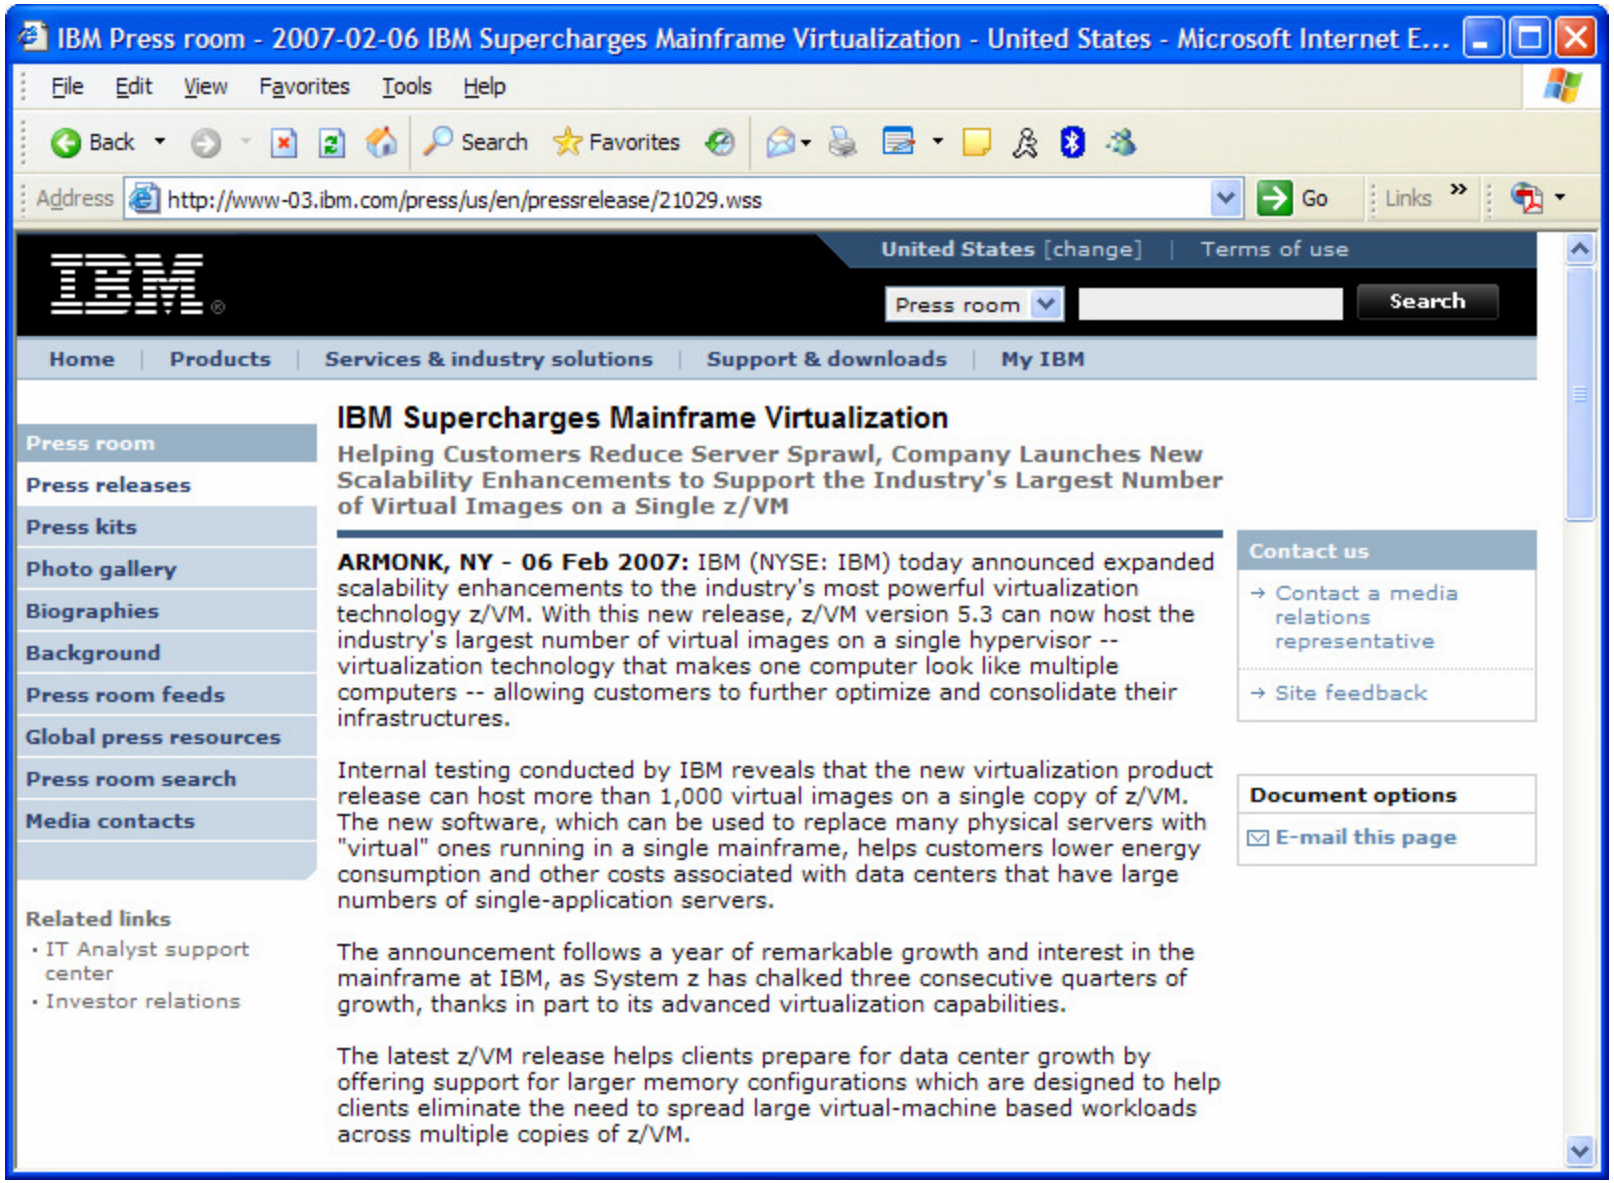

Refer to IBM Software Announcements 207-019 (February 6, 2007) and 207-135 (June 12, 2007)

Provisioning Virtual Linux Servers on System z

Using IBM Director for Linux on System z V5.20 with z/VM Center

IBM Director Base Functions

- Discovery
- Group Management
- Inventory
- Basic Resource Monitor
- Event Action Plan (EAP)
- Process Management
- Remote Session
- File Transfer
- Network Configuration
- Software Distribution
- SNMP Browser

z/VM Center

- Utility Service Configuration Manager New
- z/VM Virtual Server Deployment
- z/VM Server Complexes

Software Distribution Premium Edition New

- Software package distribution