



V52

z/VM Platform Update: What's New with z/VM V5.2

Reed A. Mullen

IBM
SYSTEM z9 AND zSERIES EXPO
October 9 - 13, 2006

Orlando, FL



IBM System z9 and zSeries Expo – Orlando, Florida

z/VM Platform Update: What's New with z/VM V5.2

October 2006

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


Topics

- **Recent z/VM Release History**
- **New z/VM Product Releases from IBM**
- z/VM Version 5 Product Information
- **z/VM Version 5 Release 2 Review**
- **z/VM V5.2 Post-GA Support**

Recent z/VM Release History

z/VM Version 5: High-Value Virtualization Technology

- ★ Generating new business with Linux on zSeries
- ★ Enabling growth for existing VM customers




z/VM Version 4 Release Highlights

- **z/VM V4.1**
 - New pricing Ts&Cs
 - Support for IFL engines
 - Linux performance support
 - Express Install (for new users)
- **z/VM V4.2**
 - Guest LAN support
 - HiperSockets support
 - PCICC / PCICA Crypto support
 - Linux performance support
 - Guest support for CF Duplexing
- **z/VM V4.3**
 - Guest FCP support
 - Virtual Machine Resource Manager
 - Guest LAN enhancements
- **z/VM V4.3 (continued)**
 - IP Wizard and “ifconfig” for z/VM
 - Virtual network accounting
 - Automated shutdown signal
 - RACF feature
- **z/VM V4.4**
 - z990 exploitation support
 - Integrated 3270 console support
 - Guest support for SCSI IPL
 - QDIO adapter interrupt passthru
 - Guest LAN IPv6 support
 - Virtual IP switch
 - IEEE VLAN support
 - System management APIs
 - Performance Toolkit feature
 - HCD and HCM support

z/VM Version 5 Release 1 New Function Highlights *Including Post-GA Support Enhancements*

- **Processor and device support**
 - IBM z990 and z890 support enhancements
 - Coupling Facility Control Code Level 14
 - FICON Express2
 - PCI Express and Crypto Express2
 - Support for 24 CPUs
 - OSA-Express Integrated Console Controller support
 - CP/CMS support for SCSI disks and FCP LUN access control
 - DS8000 and DS6000 storage subsystems
 - TotalStorage 3592 tape subsystem
- **Server hosting support**
 - Dynamic virtual machine timeout
 - HyperSwap (GDPS PPRC Multiplatform Resiliency for zSeries)
- **Networking**
 - OSA-Express2 support
 - Enhanced OSA-Express connectivity
 - Layer 2 and 3 Virtual Switch support
 - RACF authorization support for Guest LANs and Virtual Switches
 - VM TCP/IP support for IPv6
- **Systems management**
 - Capacity on Demand enhancements
 - Additional Systems Management APIs
 - Performance Toolkit for VM enhancements
 - Service support enhancements
- **General**
 - New publication: *Getting Started with Linux on zSeries*
 - EAL 3+ certification



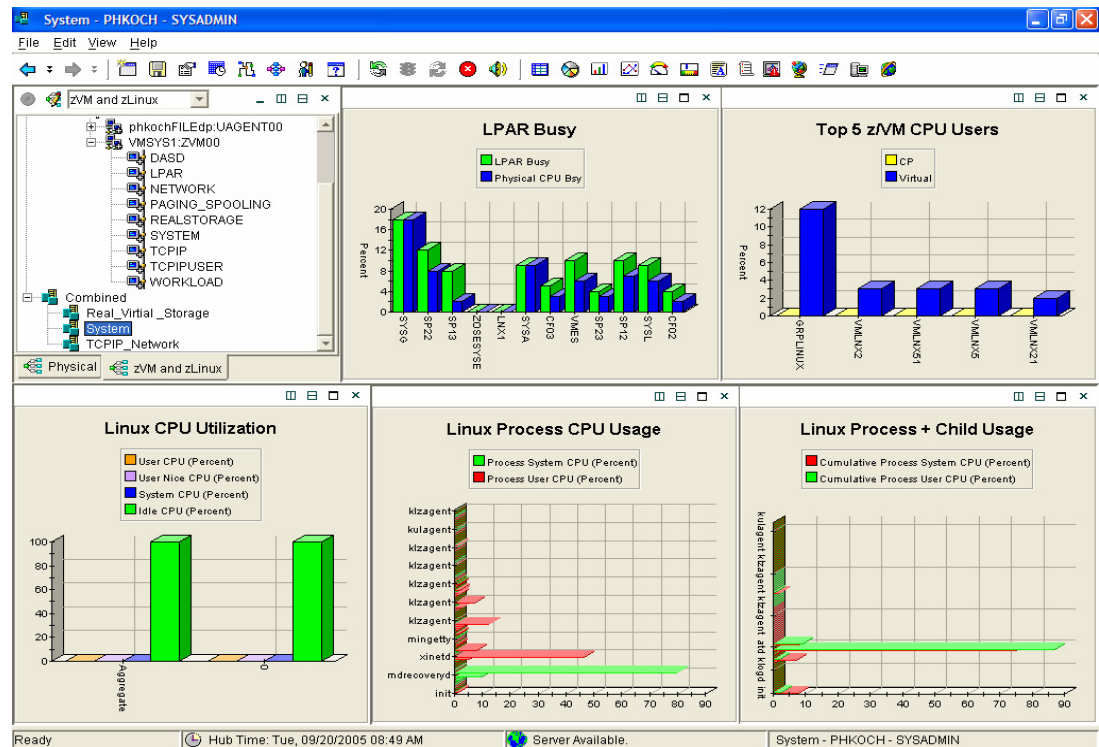
z/VM Systems Management Products from IBM

- **New z/VM product release from IBM (announced 15 Aug 2006)**
 - IBM Backup and Restore Manager for z/VM V1.2 – August 2006
- **Other z/VM Systems Management products from IBM**
 - IBM Tape Manager for z/VM V1.2 – February 2006
 - IBM Operations Manager for z/VM V1.2 – February 2006
 - IBM Archive Manager for z/VM V1.1 – August 2005
- **Same pricing model as z/VM V5**
 - Value Units based on processor engines

IBM Tivoli OMEGAMON XE on z/VM and Linux

- **Preview announcement* made August 15, 2006**
- **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**
- **Supported on z/VM V5.2**
 - Requires the z/VM V5.2 Performance Toolkit for data collection
- **Planned availability: 4Q06**

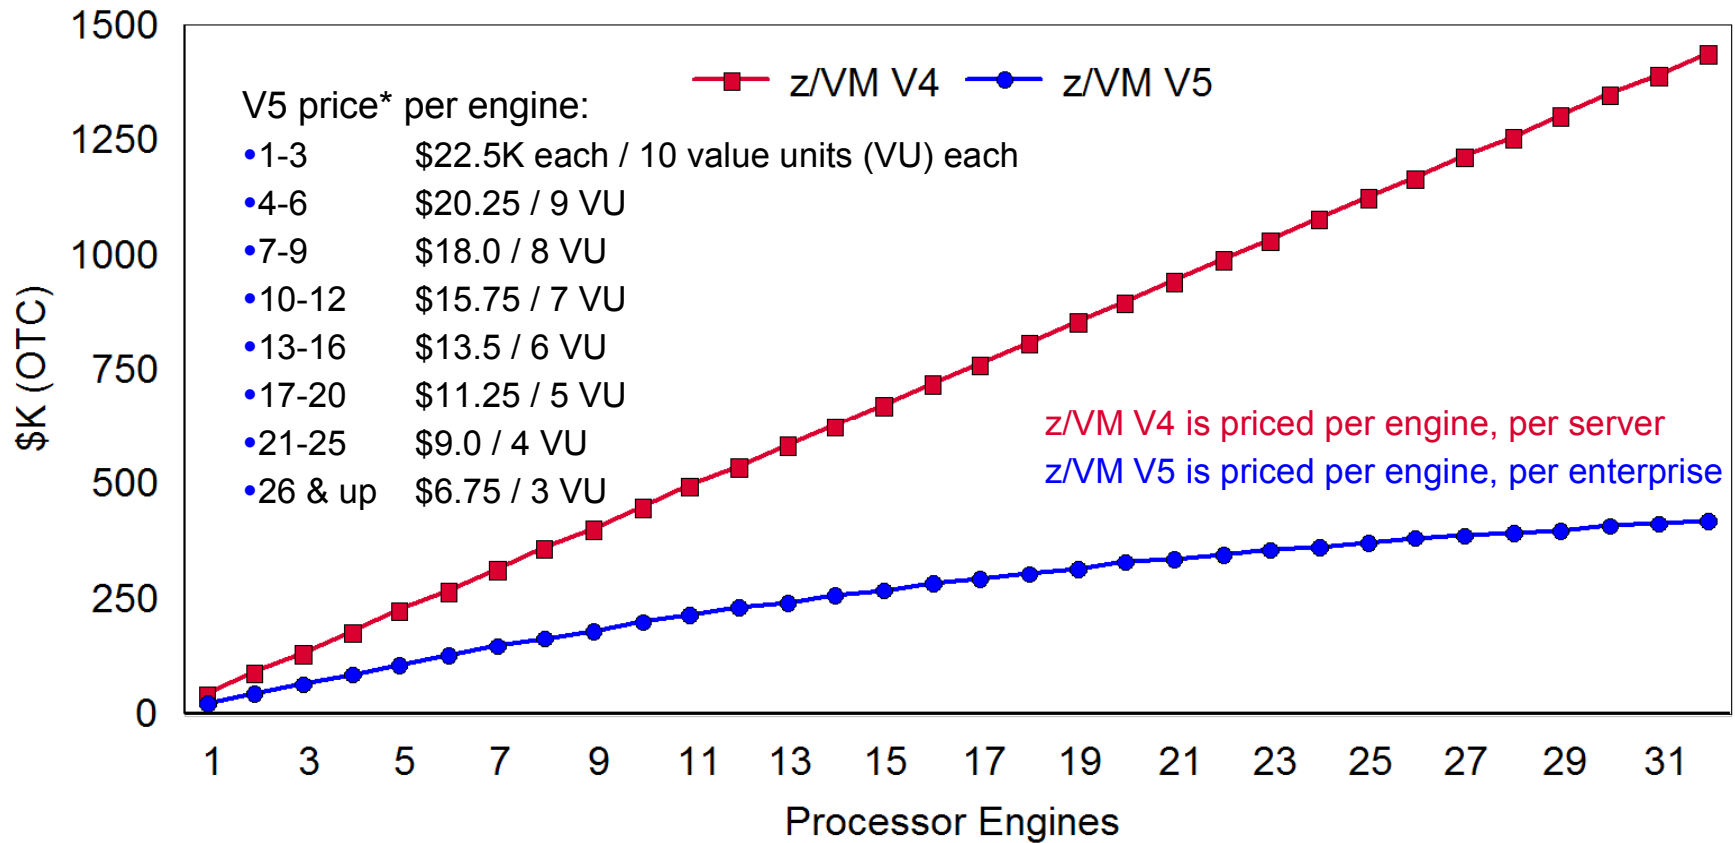
* Refer to IBM Software Announcement 206-201



z/VM Version 5 Product Information

- **Runs on IBM System z9 (z9 EC and z9 BC) and IBM eServer zSeries (z800, z900, z890, z990) processors and other equivalent servers**
 - The z/VM V5 Control Program requires 64-bit addressing (z/Architecture)
 - 64-bit and 31-bit (ESA/390) virtual machines are supported
- **Runs on Integrated Facility for Linux (IFL) processor engines as well as standard processor engines**
- **IPLA software product (5741-A05) with new, improved pricing Ts&Cs**
 - One-time charge license fee, priced on a per-engine basis
 - Price/engine decreases (on a tiered basis) as more engines are licensed
 - Engines can be aggregated across an enterprise for licensing purposes
 - Ordered via the System Delivery Option (SDO) (5741-A06)
- **Optional Software Subscription & Support (S&S) product (5741-SNS)**
 - Required to receive IBM support center services
 - Entitles customers to future z/VM releases and versions
 - Annual, renewable license charge
- **Includes priced features**
 - DirMaint, RACF/VM, Performance Toolkit for VM
 - Pre-installed, but disabled (license required; same pricing model as base)

z/VM Version 5 Pricing



*U.S. prices as of 1 Oct 2006

z/VM Version 5 Pricing

Detailed Information

- **z/VM V5 uses a Value Unit pricing model**
 - z/VM V5 value units correspond to processor engines, not MIPS or MSUs
 - A single z/VM V5 value unit is priced at \$2,250 (U.S. pricing as of 1 Oct 2006)
 - Engines 1, 2, and 3 are priced at 10 value units each
 - Engines 4, 5, and 6 are priced at 9 value units each
 - Pricing continues on a tiered basis
- **z/VM Version 4 customers who have purchased Software Subscription and Support (S&S) are entitled to receive z/VM Version 5 at no charge**
 - No charge to run z/VM V5 on same number of V4-licensed engines
 - Subsequent S&S annual payments will be based on z/VM V5 pricing
 - Keep in mind z/VM Version 5 requires z/Architecture to operate
 - If the customer adds capacity (engines) after the migration, pricing for the added capacity will be based on the z/VM Version 5 pricing model
- **If z/VM V5 is licensed to run on an IFL engine, all IFLs must be counted to determine the z/VM V5 licensing fee**
- **If z/VM V5 is licensed to run on a standard engine, all standard engines must be counted to determine z/VM V5 licensing fee**


z/VM Version 5 Product Packaging Changes

- **DFSMS/VM is no longer *automatically* shipped with the base product**
 - It is now a no-charge feature and must be ordered via the SDO
- **3270 PC File Transfer product (5664-281) is *included* with base product**
 - Delivered with z/VM V5 as a sample program (with no support)
- **Restricted source feature and PL/X source *no longer ship* with z/VM V5**
 - Restricted source is a no-charge feature of z/VM V4
 - PL/X source is provided with the z/VM V4 installation media
 - Both will be available as no-charge downloads from IBM Resource Link for z/VM V5 customers (who register with Resource Link)
- **Tivoli Storage Manager for VM is *no longer pre-installed* with z/VM V5**
 - TSM for VM is packaged with the z/VM V4 system DDRs
 - Consider TSM for Linux on zSeries for future TSM server support
- **National Language features for ISPF have been *removed* from the SDO**
 - Features can be ordered using the standalone ordering process
- ***HCD/HCM upgraded to new level***
- ***z/VM Collection Kit publications available on DVD (supplied with order)***

z/VM V5.2-only items highlighted in *blue*

Functions Removed from z/VM Version 5

- **RTM and PRF features (replaced by Performance Toolkit for VM)**
- **SPTAPE (use SPXTAPE to backup Spool files)**
- **V=R and V=F virtual machine support**
- **CMS support for Java and NetRexx programs**
- ***System Administration Facility***
- ***Support for Server-Requester Programming Interface (SRPI)***
- **Device support**
 - DASD/Controllers: 3370, 3375, 3380(1), Multiprise Internal Disk, 9332, 9335, 9336(2), 9340, 3830, 3880
 - Optical: 3995 Optical Dataserver
 - Tape/Controllers: 2440, 3420, 3422(3), 3424, 3430, 9348, 3803
 - Communications: all SDLC, BSC, and CETO ICAs, *3705, 3720, and 3725 Communication Controllers, 8232 LAN Channel Station*
 - Terminals: *2741 and TWX Terminal Model 33/35 (TTY) as virtual consoles*
 - Refer to the z/VM V5.2 GIM for a complete listing of devices supported



Notes:

- (1) RAMAC-emulated 3380 models J and K and 3390 DASD configured for 3380-track-compatibility are supported
- (2) 9336 is a supported device geometry for Virtual Disks in Storage and emulated SCSI LUNs
- (3) OMA/2 CD-ROM emulating a 3422 is supported

z/VM V5.2-only items highlighted in *blue*

z/VM Version 5 Product Installation Support



- **z/VM V5.2 can be installed on:**
 - 3390-formatted DASD volumes (Models 3 and 9)
 - FCP-attached SCSI disks: ESS 750, ESS 800, DS8000, DS6000
- **Installation media options:**
 - DVD (the only option for installing z/VM V5 on SCSI disks)
 - 3590 and 3480 tape
 - *No CD-ROM support for z/VM V5.2*
- **System Residence (SYSRES) volume changes**
 - Spool and paging space removed from SYSRES
 - Located on separate installation volumes
- **Two installation methods available with z/VM V5**
 - Streamlined process ("Express Install") for new users
 - Traditional process for experienced systems programmers



z/VM V5.2-only items highlighted in *blue*

Key z/VM V5.2 Business Value Propositions

- **Significant “Scale out” and “Scale up” support for hosting virtual server workloads**
 - Improved utilization of large real memory configurations
 - Improved memory management for Linux guests running on a System z9 server*
 - Improved bandwidth for QDIO operations in a CPU-constrained environment (support for z9 EC, z9 BC, z990, z890 servers only)
 - Improved throughput and response time for minidisk I/O in an I/O-constrained environment*
- **Enhanced virtual networking support**
- **Improved FCP/SCSI support**
- **Support staff productivity gains with IBM Director**



* References post-GA support announced on April 27, 2006

z/VM Version 5 Release 2 New Function Highlights

Expanding IBM Mainframe Support for Virtual Server Hosting

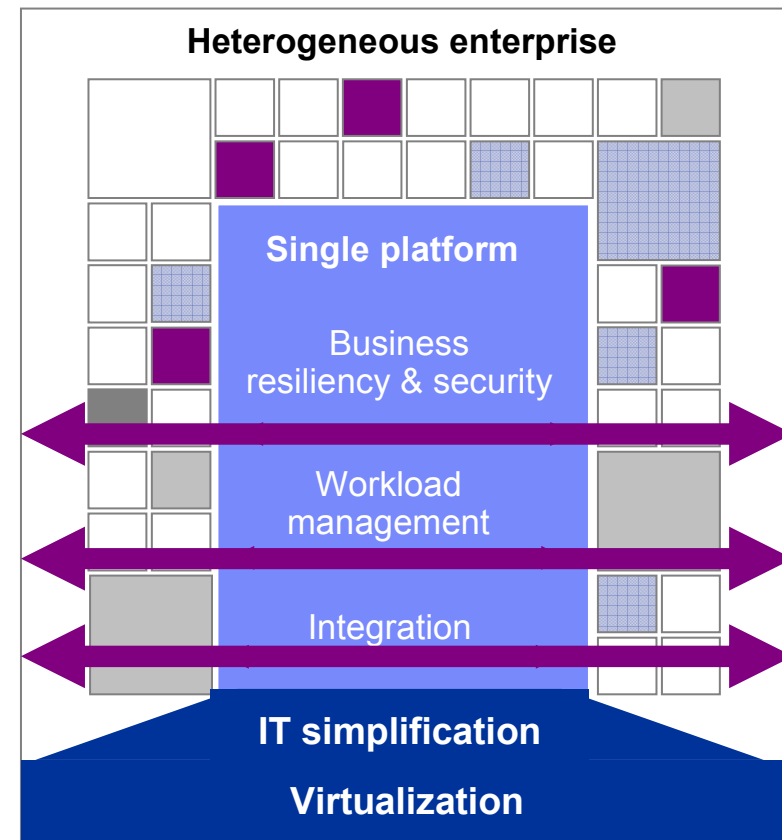
- **Processor and device support**
 - IBM System z9
 - Dynamic LPAR naming support
 - Crypto Express2 Accelerator
 - SCSI disk I/O performance improvements
 - N_Port ID Virtualization support
 - DS8000 and DS6000 storage subsystems
- **Server hosting support**
 - Enhanced z/VM support for large real memory configurations
 - Enhanced performance assists for z/VM guest images
 - z/VM Guest LAN and Virtual Switch sniffer support
- **Networking**
 - OSA-Express2 Open Systems Adapter for NCP support
 - New MPROUTE server for z/VM
 - z/VM SSL Server upgrade
- **Systems management**
 - Enhanced z/VM systems management
 - Simplified user administration: DirMaint and RACF coordination
 - Improved directory management performance
 - Performance Toolkit for VM support
 - Product service and installation enhancements



IBM System z9

Meeting Business Needs Today, Providing a Foundation for the Future

- **Extending its leadership in**
 - Business resiliency and security
 - Intelligent workload management
 - Business integration
- **Designed with a holistic systems approach**
- **Able to intelligently manage across the enterprise**



“Leadership is not only about technology innovation, but also delivering value and actively supporting the mainframe community.”

IBM Mainframe Charter

IBM System z9 Enterprise Class* (z9 EC)

- **5 models (S08, S18, S28, S38, S54); improved performance over the z990**
- **Available Processor Units (PUs):**
 - 12 PUs per book (16 for Model S54); 2 SAPs per book; 2 spares per server
 - 8, 18, 28, 38, or 54 PUs available for CPs, IFLs, ICFs, zAAPs, zIIPs, optional SAPs
- **Up to 60 LPARs (no basic mode); up to 512 GB of central processor storage**
- **Up to 4 Logical Channel Subsystems (LCSS)**
 - Each can be configured with up to 256 channel paths (1024 per CEC)
 - Up to 15 LPARs per LCSS, one LCSS per LPAR
- **FICON Express4, FICON Express2, FICON Express**
 - Up to 84 features / 336 channels (FICON Express4 or Express2)
- **Up to 1024 ESCON channels**
- **OSA-Express2 and OSA-Express**
 - Up to 24 features; Fast Ethernet, 1000BASE-T Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet
- **N_Port ID Virtualization (NPIV)**
- **Up to 8 configurable Crypto Express2 features**
- **Up to 16 HiperSockets for high-speed interconnect**

* Formerly “IBM System z9 109”



IBM System z9 Business Class (z9 BC)

- **2 models (R07, S07) with 73 Capacity Indicators; improved performance over the z890**
- **Available Processor Units (PUs):**
 - 8 PUs per system; 1 SAP standard (5 PUs per system for Model S07)
 - Up to 7 PUs available for CPs, IFLs, ICFs, zAAPs, zIIPs, optional SAP (4 PUs for Model S07)
- **Up to 30 LPARs (no basic mode); up to 64 GB of central processor storage**
- **Up to 2 Logical Channel Subsystems (LCSS)**
 - Each can be configured with up to 256 channel paths (512 per CEC)
 - Up to 15 LPARs per LCSS, one LCSS per LPAR
- **FICON Express4, FICON Express2, FICON Express**
 - Up to 28 features / 112 channels (FICON Express4)
- **Up to 420 ESCON channels**
- **OSA-Express2 and OSA-Express**
 - Up to 24 features; Fast Ethernet, 1000BASE-T Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet
- **N_Port ID Virtualization (NPIV)**
- **Up to 8 configurable Crypto Express2 features**
- **Up to 16 HiperSockets for high-speed interconnect**



z/VM Release Support for IBM System z9

- **Support for IBM System z9 is included in the GA release level of z/VM V5.2**
- **z/VM V5.2 offers additional support for IBM System z9 via the service stream**
 - New function available with the PTFs for APARs VM63952 and VM63856
- **z/VM V5.1 and V4.4 offer compatibility support for IBM System z9**
 - Required PTFs (APARs) for z/VM V5.1 and V4.4 compatibility support:
 - CP: VM63646, VM63722, VM63744
 - CMS (IOCP): VM63740
 - EREP: VM63743
 - HCD/HCM: VM63721
 - OSA/SF: OA11650
- **z/VM V5.2, V5.1 and V4.4 offer the following support capabilities for IBM System z9:**
 - Multiple Logical Channel Subsystems support
 - Internal and external spanned channel support
 - Extended channel data measurement support
 - Support for configurations with up to 60 LPARs



z/VM Dynamic LPAR Naming Support

- **z/VM V5.2 provides the facilities to dynamically define and delete Logical Partitions (LPARs)**
 - Support is available using CP's dynamic I/O command interface and z/VM HCD/HCM support
- **Hardware support is available for z9 EC, z9 BC, z990, and z890 servers**
- **LPARs can be defined without real resource allocations and dynamically configured at a later time**
- **Capability allows customers to add meaningful LPAR names to a running configuration without requiring a Power-On Reset (POR)**

z/VM Support for IBM System z9

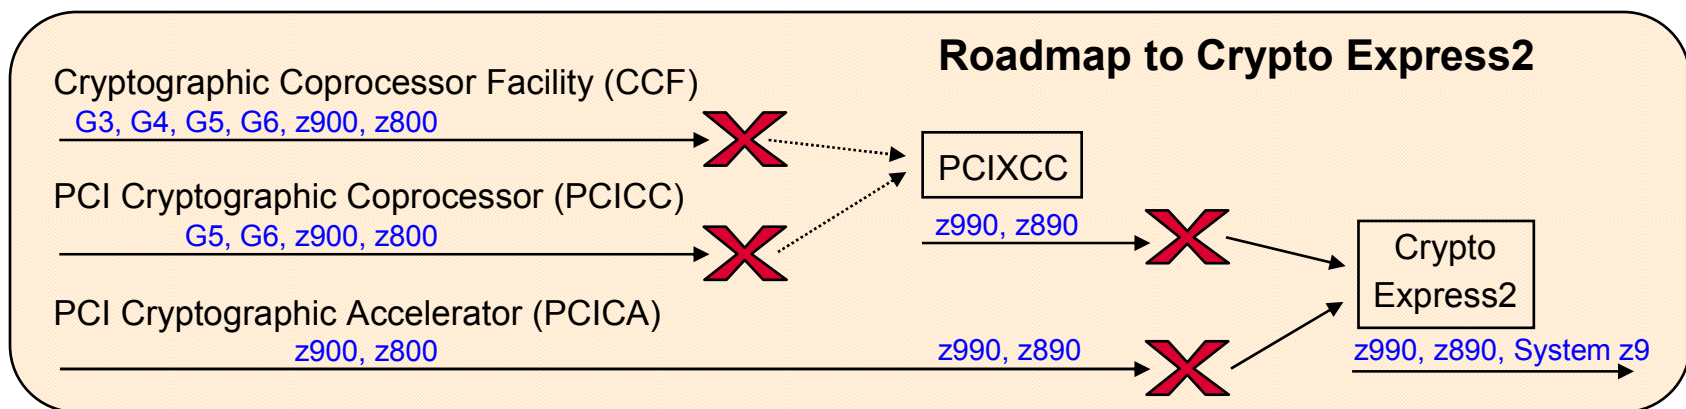

Crypto Express2 Accelerator

- **Crypto Express2 Coprocessor**

- Shared-queue and dedicated-queue support for clear key cryptographic functions for Linux guests
- Dedicated-queue support for clear-key and secure-key cryptographic functions for z/OS guests

- **Crypto Express2 Accelerator**

- Designed to support clear-key RSA operations
- Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol
- Supported by z/VM V5.2 and also V5.1 with the PTF for APAR VM63646



z/VM Version 5 Support for FCP-Attached SCSI Disks

Integrate Your z/VM Systems with Storage Area Networks

- **z/VM V5.2 and V5.1 allow FCP-attached SCSI disks to be used for both system use (CP/CMS) and guest images**
- **SCSI disks are emulated as 9336 Model 20 FBA devices for system use**
 - Enables support for install, paging, spooling, directory services, minidisks
 - Guest systems supporting FBA can also use emulated SCSI disks
 - Emulation support currently limits usable disk space to nearly 1 TB for CP volumes and 381 GB for CMS and GCS
 - Paging, spooling, and directory space must reside in first 64 GB
- **Non-emulated SCSI disks can still be attached to virtual machines**
 - For boot and/or data operations
 - Requires SCSI support in guest operating system
- **Currently supported SCSI disks:**
 - IBM TotalStorage Enterprise Storage Server Models 750 and 800
 - IBM TotalStorage DS8000 and DS6000
 - Generic SCSI driver available for other disks
- **SCSI-only disk configurations are now possible with z/VM V5**



z/VM V5.2 Improved Performance of SCSI Disk I/O


- **QDIO efficiency improvements**
 - Chaining I/O requests allows more data to be moved with a single I/O request
 - Reduces fragmentation of large channel programs; decreases CP overhead
 - Significantly lowers number of I/O operations driven through the SCSI stack
- **Paging and spooling optimization**
 - No longer uses FBA emulation
 - SCSI commands are sent directly to the SCSI stack
- **Improved FBA emulation**
 - Reduces the number of I/O requests to the z/VM SCSI stack
 - Significantly improves performance for CMS and Linux SCSI I/O
 - I/O buffers are read more efficiently by CMS
 - FBA “padding function” is handled more efficiently: reduces time required to boot a Linux guest image
 - Improves I/O performance for SCSI format functions

z/VM Support for IBM System z9 *N_Port ID Virtualization (NPIV)*

- **FICON Express features on System z9 support FCP N_Port ID Virtualization**
- **NPIV complements zoning and LUN masking**
- **Multiple operating system images can now concurrently access the same or different SAN-attached devices (LUNs) via a single, shared FCP channel**
 - Can improve channel utilization
 - Less hardware required
 - Helps reduce the complexity of physical I/O connectivity
- **Also supported by z/VM V5.1 and V4.4 with the PTF for APAR VM63744**
 - Support limited to z/VM Control Program only
- **z/VM V5.1 cannot be installed from DVD to SCSI disks when NPIV is enabled**
 - This restriction does not exist for z/VM V5.2
- **z/VM V5.2 PTF for APAR VM63952 allows guest operating systems and z/VM users to query virtual port names**

N_Port ID Virtualization (NPIV)

Without N_Port ID Virtualization




No NPIV

- Image access to shared FCP channel allows read-write access to all LUNs not masked
- No concurrent LUN sharing



With NPIV

- Image-defined LUN access on shared FCP channel
- Read-only LUN sharing is possible

Problem!



With N_Port ID Virtualization



IBM TotalStorage DS6000 and DS8000

Advanced Disk Storage Solutions

- **DS6000 Series:**

- Designed to provide exceptional price & performance in a modular package that redefines enterprise-class storage
 - 16 drives / 3U package; scalable from 4 to 128 drives; up to 38.4 TB
- Shares advanced software features with ESS and DS8000



- **DS8000 Series:**

- Designed to deliver massive scalability using IBM POWER5 processors
- Up to 6 times the performance of a base ESS 800 for increased response time and usable capacity (up to 192 TB of storage) with 20% smaller footprint
- Storage system LPARs in select DS8300 models allow two virtual storage systems within a single array (reducing price/megabyte)



- **z/VM support:**

- z/VM V5.2 includes support for SCSI disk capacities nearing 1TB* for CP volumes and 381 GB for CMS and GCS volumes
- z/VM V5.1 PTFs – *DS8000*: VM63534, VM63653; *DS6000*: VM63535, VM63653; *near-1 TB support*: VM63700, VM63664
- z/VM V4.4 PTFs – *DS8000*: VM63534; *DS6000*: VM63535
- VM63535 includes preferred pathing support for DS6000

* 2,147,483,640 512-byte blocks

Understanding z/VM Support for SCSI Disks

- **z/VM ECKD I/O can achieve a higher level of performance than SCSI disk I/O**
 - Continue to use ECKD disks for CP/CMS I/O if it is an option
- **Increased pathlength of z/VM SCSI disk I/O can be offset**
 - Reduce over commitment of virtual-to-real memory (i.e., reduce paging)
 - Use minidisk cache for read-mostly I/O
 - FCP/SCSI channels are faster than ESCON/ECKD channels
 - Additional processor cycles will offset increased SCSI I/O pathlength
- **Sharing FBA-emulated SCSI disks among Linux images can offer disk and administrative savings**
 - Allows partitioning of SCSI disks using z/VM minidisk support (includes exploitation of minidisk cache support)
 - Allows use of tuning options like “Set Throttle” and “Set IOPriority”
 - Performance monitoring of emulated disks is functionally richer than SCSI disks accessed via dedicated FCP-subchannels
- **IBM TotalStorage DS6000 offers a low-cost ECKD option for z/VM data**

Enhanced z/VM Support for Large Real Memory

Constraint Relief for Memory-Intensive Virtual Server Environments

- **z/VM V5.2 Control Program (CP) offers improved performance and scalability for environments with high demand on storage below 2 GB**
 - I/O data can now be transferred from buffers located anywhere in memory
 - QDIO structures may now reside above the 2 GB address line
 - Most CP control blocks may now reside above the 2 GB address line
- **z/VM V5.2 can scale a mainframe server to 128 GB of real memory**
- **Storage above 2 GB address line is included in dumps**
 - CP hard and soft abend dumps and SNAPDUMPs
 - Standalone z/VM dumps or VMDUMPs of z/Architecture virtual machines
- **TCP/IP for z/VM exploitation of 64-bit Diagnose 98**
 - Enhanced QDIO device driver uses I/O buffers above 2 GB when possible
 - Helps reduce chance of server failure due to lack of buffer space
- **Block I/O (Diagnose 250) support**
 - Virtual machines can specify parameter addresses and I/O buffers above the 2 GB address line
 - Linux exploitation support is available with SUSE Linux Enterprise 10




Enhanced z/VM Support for Large Real Memory

A Closer Look at the Numbers

- **z/VM V5.2 provides considerable scalability advantages compared to V5.1**
 - z/VM V5.2 users can now deploy large-memory Linux virtual machines that cannot run on z/VM V5.1, a significant “**scale up**” improvement for Linux on z/VM
 - In memory-constrained environments, clients are able to increase the number of virtual Linux servers that can run on a single copy of z/VM, a significant “**scale out**” improvement for Linux on z/VM
 - These benefits can be realized on any z/Architecture-capable server
- **New Limits...**
 - z/VM V5.2 cannot support a virtual machine that has 256 GB of active memory
 - The effectiveness of hosting workloads in an environment with more than 128 GB of active virtual memory will vary; this potential constraint is due to the fact that page management control blocks must still reside in the first 2 GB of real memory

Enhanced z/VM Support for Large Real Memory

Large-Memory Linux Guest Configurations

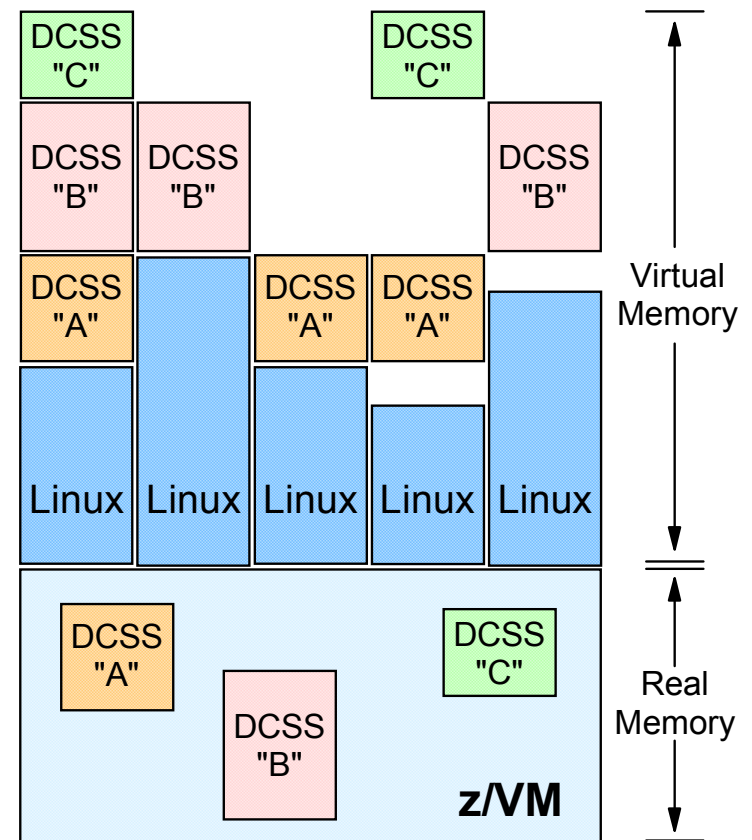


- ➔ Large Linux guest systems can now run on z/VM without special treatment in the Linux disk device driver
- ➔ System scalability is effective up to 128 GB of real memory

Linux and z/VM Technology Exploitation

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


- **DCSS support is a z/VM exclusive**
 - Share a single, real memory location among multiple virtual machines
 - High-performance data access
 - Can reduce real memory utilization
- **Linux exploitation support available today**
 - Execute-in-place (xip2) file system
 - 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
 - Enables throughput benefits for Linux guest images and enhances overall system performance and scalability



Enhanced Performance Assists for z/VM Guests

Improved Performance for Guest QDIO Operations

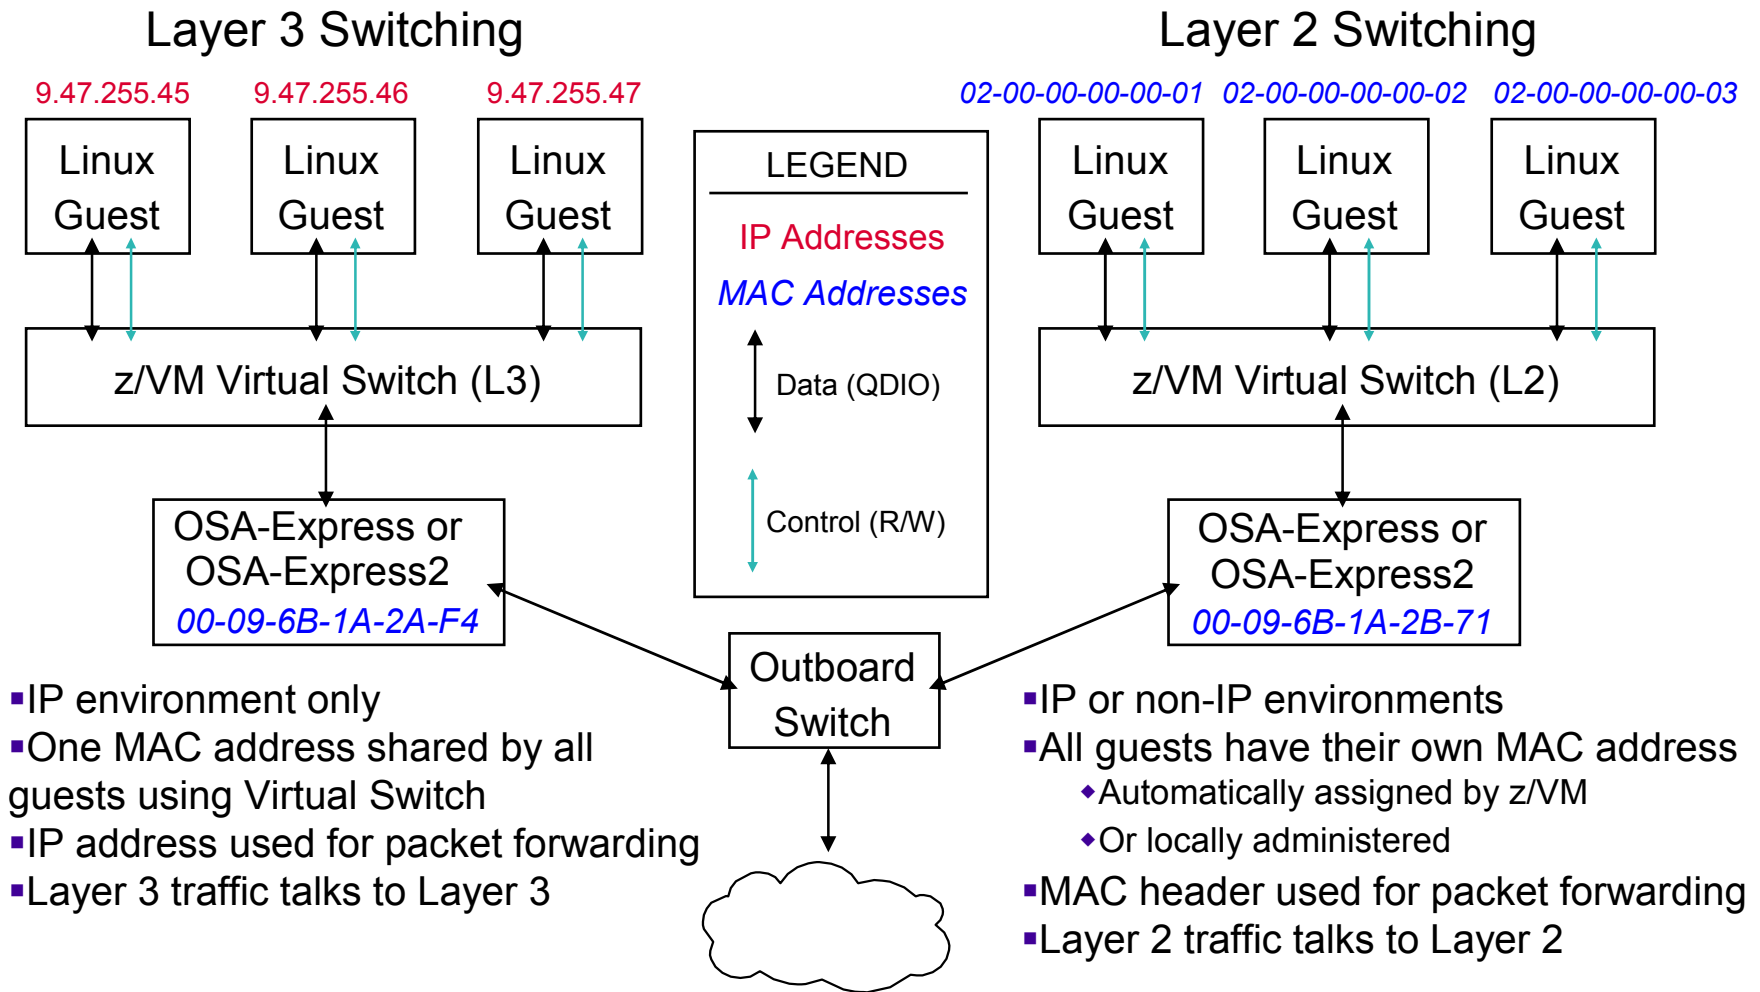
- **QDIO Enhanced Buffer-State Management (QEBSM)**
 - Two new machine instructions designed to help eliminate overhead of hypervisor interception
- **Host Page-Management Assist (HPMA)**
 - Interface to z/VM paging and storage management
 - Designed to allow hardware to assign, lock, and unlock page frames without hypervisor assistance
- **Assists are applicable to the following environments:**
 - First-level guests of z/VM V5.2
 - HyperSockets (CHPID type IQD)
 - All OSA features (CHPID type OSD)
 - All FICON features (CHPID type FCP)
- **Complements performance assists introduced in z/VM V4.4**
- **Hardware/firmware support available with z9 EC, z9 BC, z990, and z890**



Note: The z990 and z890 servers will require MCL updates for these new assists. MCL112 in the J13484 stream must be installed on the z990 and z890 prior to IPLing z/VM V5.2 to prevent system outages from occurring.

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
- Layer 3 traffic talks to Layer 3

- 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
- Layer 2 traffic talks to Layer 2

z/VM Guest LAN and Virtual Switch Sniffer Support

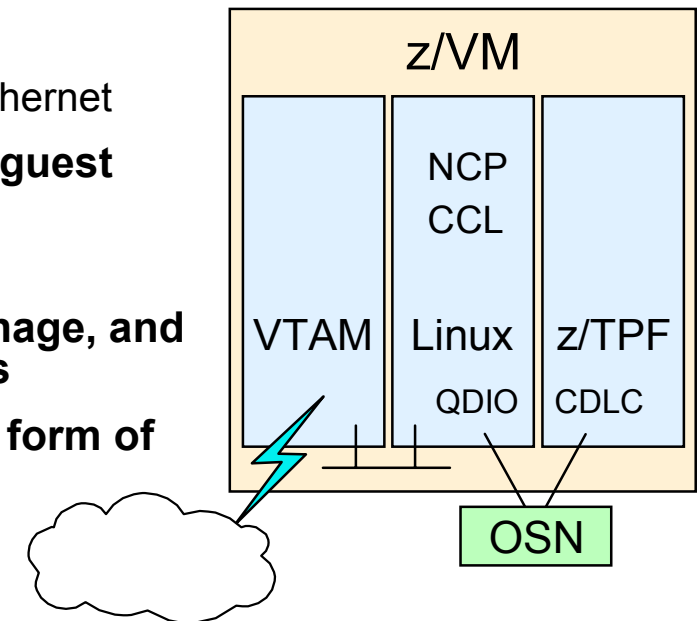
- **LAN sniffer captures network traffic on a z/VM Guest LAN or Virtual Switch**
- **Useful for resolving virtual networking problems**
 - Authorization for use is under the control of the system administrator, not users or guest images
- **Linux tracing capabilities:**
 - Authorized Linux guest can use CP commands (or the Linux device driver, when available) to put the guest NIC in “promiscuous mode”
 - Network traffic can be traced, recorded, & analyzed using existing Linux tools
 - Collected data can be printed or written to disk using tools such as tcpdump
- **Native z/VM tracing capabilities:**
 - Network traffic can be traced, recorded, & analyzed using z/VM facilities
 - Only authorized to users with Class C privileges
 - CP TRSOURCE command has been enhanced to trace and record data
 - New IPFORMAT tool can analyze data to determine cause of problems
- **RACF for z/VM feature provides ability to control promiscuous mode authorizations for Guest LANs and Virtual Switches**
- **IBM is working with its distributions partners to provide this function in future Linux on System z9 and zSeries distributions or service updates**



OSA-Express2 Open Systems Adapter for NCP

New CHPID Type OSN

- **Designed to help eliminate dependencies on hardware**
 - 3745/3746, ESCON, Token-Ring
- **Provides support for IBM Communications Controller for Linux (CCL)**
- **Appears to operating systems as an ESCON-attached channel connected to a 3745 device type**
- **Communication traffic can flow from LPAR to LPAR**
- **OSN support is exclusive to IBM System z9**
 - Requires OSA-Express2 Gigabit or 1000BASE-T Ethernet
- **z/VM support enables VTAM on z/VM and other guest images to access OSA-Express CDLC channels**
 - Support available in z/VM V5.2 and V5.1
- **Allows system administrators to configure, manage, and operate their CCL NCPs just like real 374x NCPs**
- **Can help eliminate the requirement to have any form of external medium (and all related hardware) for communications between the host operating system and the CCL image**



New MPROUTE Server for z/VM

Enhanced Dynamic Routing Support

- **OMPROUTE V1.7 module has been ported from z/OS to z/VM V5.2**
- **New MPROUTE server includes support for:**
 - IPv6 dynamic routing including RIPng and OSPF
 - IPv4 OSPF authentication using MD5 (cryptographic authentication)
 - Improved IPv4 VIPA support
 - Receiving RIPv1 and RIPv2 on same link
 - Up to 16 equal cost paths to a single destination
- **ROUTED server will be removed from a future release of z/VM**
- **Utility (RTD2MPR EXEC) is supplied to assist in migration from ROUTED to MPROUTE**
- **Support benefits include:**
 - Additional supported protocols
 - Greater efficiency may be achieved within an IP network
 - Manual network routing table updates may be reduced or eliminated

z/VM SSL Server Upgrade


Support for More-Current Linux Distributions

- **z/VM V5.2 SSL server supports the following Linux distributions:**
 - SUSE SLES8 Service Pack 3 (31-bit)
 - SUSE SLES9 Service Pack 2 (31-bit)
 - SUSE SLES9 Service Pack 2 (64-bit)
 - Red Hat Enterprise Linux AS V3 (31-bit)
 - Red Hat Enterprise Linux AS V3 (64-bit)
 - Red Hat Enterprise Linux AS V4 (31-bit)
 - Red Hat Enterprise Linux AS V4 (64-bit)
- **Additional support in upgraded server:**
 - Industry-standard encryption algorithms, including DES, triple-DES, RC2, and RC4 with keys up to 128 bits in length; hashes provided by SHA-1 and MDS
 - Certificate activation and removal without server shutdown/restart
 - Add or deactivate a certificate while SSL-secured sessions are active
 - Federal Information Processing Standard (FIPS) 140-2 support
 - FIPS 140-2 support allows connections to be restricted to those that employ FIPS-approved cipher suites



Enhanced z/VM Systems Management Functions *For Allocating and Managing Guest Resources*

- **z/VM V5.2 implements Version 3 (V3) of the systems management server**
 - z/VM V4.4 and V5.1 functions continue to operate with the new V3 server
 - API support enables solution providers to more easily help administrators manage a large number of virtual images running on a z/VM system
- **New APIs provided for the following systems management functions:**
 - Creating/updating the LOADDEV directory statement for a virtual machine
 - Responding to queries of a virtual machine's LOADDEV settings
 - Obtaining the time when a virtual image was activated
 - Scanning (searching) the directory for a specified pattern
 - Defining Local Tags by the directory manager
- **Enhancements to existing APIs:**
 - Support for larger disk block sizes when creating disks and defining disk space on virtual image volumes
 - Capability to create persistent Virtual Switch definitions
 - New definition parameters for Virtual Switch APIs
- **Requires a directory manager**
 - IBM DirMaint FL510 supports the new/enhanced APIs





The screenshot shows the IBM Director Console interface. The main window is titled "All Systems and Devices : Server Complexes Me...". It features a menu bar (Console, Tasks, Associations, View, Options, Window, Help) and a toolbar with various icons. The interface is divided into three main sections:

- Groups:** A tree view on the left showing categories like "All Groups", "All Systems and Devices", "Chassis and Chassis Mem", "Clusters and Cluster Memb", "Hardware Status Critical", "Hardware Status Informatio", "Hardware Status Warning", "IBM Director Systems", "Platforms and Platform Men", "Systems with Linux", "z/VM Server Complexes", and "z/VM Systems".
- Table:** A central table listing system details. The header includes "Status and Name" and "TCP/IP". The table contains several rows representing different server groups and individual machines, such as "Free guests", "Production", "Print Servers", "Web Servers", "Test", "T1", and "Not Associated".
- Tasks:** A list of management tasks on the right, including "Event Action Plans", "Event Log", "External application launch", "File Transfer", "Hardware Status", "Inventory", "Microsoft Cluster Browser", "Network Configuration", "Process Management", "Remote Control", "Remote Session", "Resource Monitors", "Scheduler", "SNMP Browser", "Software Distribution", "All Software Distribution Packag", "System Accounts", "z/VM Center", "z/VM Server Complexes", and "z/VM Virtual Server Deployment".

A red arrow points to the "z/VM Center" task in the right-hand pane.

IBM Director is part of the IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and zSeries

Learn more at: ibm.com/servers/eserver/xseries/systems_management/ibm_director/extensions/zvm.html



Ready


Host: ps-biran User ID: PS-BIRAN\biran

11:00 AM GMT

IBM Director V5.10 – Extensions

z/VM Center – Fee-based Extension

- **Virtual Server Deployment – easy deployment of Linux servers on z/VM**
 - Creation of templates for virtual server and Linux operating system provisioning
 - Creation of virtual server from template
 - Applying Linux into a virtual server from template
- **Server Complexes – one-step provisioning of multiple Linux virtual servers**
 - Exploits Virtual Server Deployment task
 - Manages configuration settings of z/VM Linux virtual servers in order to reduce customization steps (user interaction) needed by VSD
 - Definition of virtual networking, minidisk attachments
 - Supports configuration scripts
 - Controls virtual server resource assignment (CPU and IO velocity goals) via z/VM VMRM
 - Manages configuration consistency:
 - Shows deviation
 - Applies desired configuration



Simplified User Administration Support

Coordination of DirMaint and RACF Changes

- **z/VM V5.2 can integrate the directory management functions of DirMaint with the security management functions of RACF**
 - DirMaint can be configured to notify RACF whenever important changes are made to user definitions and the resources they own
- **Functions that are coordinated by DirMaint with RACF include:**
 - User creation, deletion, and changes
 - Password management
 - POSIX segment management
 - Access Control Interface (ACI) group management
 - Profile creation and deletion for selected VM functions
- **Benefits:**
 - Reduces the administration effort and skills needed to deploy and manage users and their resources when DirMaint and RACF are used together
 - Eliminates the need to manually define and manage z/VM resources in RACF
 - Helps reduce the chance of incomplete or incorrect RACF configuration data

Improved Directory Management Performance

Beneficial for z/VM Systems with Large User Directories

- **Performance improvement in DIRECTXA with exploitation by the z/VM V5.2 DirMaint feature**
- **Changes made to the z/VM user directory using DirMaint should be processed faster than previous z/VM releases**
 - New function allows a change to the directory without requiring reprocessing of the entire directory
 - Directory updates complete in less time
- **Performance benefit depends on the type of directory changes being made and the size of the z/VM directory**
 - The larger the directory size, the more beneficial the performance improvement



Performance Toolkit for VM Enhancements

- **The Performance Toolkit for VM feature is a performance and reporting tool for the z/VM system and its guest images**
 - Realtime and historical reporting
 - Offers threshold monitoring and user loop detection
 - Can monitor remote z/VM systems
 - Results can be viewed graphically with a web browser
- **Replaces RTM and PRF**
 - z/VM V4 RTM and PRF customers with current S&S subscriptions are entitled to a no-charge upgrade to the Performance Toolkit for VM
- **z/VM V5.2 support enhancements:**
 - Handle changes to offsets in CP control blocks
 - Handle fields removed in CP control blocks
 - New System Execution Space report
 - New System Storage reporting
 - Eliminate Detailed User Storage report
 - Expand QDIO report based on new Monitor data
- **“IBM Tivoli OMEGAMON XE on z/VM and Linux V4.1” requires the Performance Toolkit for VM FL520 for data collection**



z/VM Service and Installation Support Enhancements


- **Further automation of the local modification process**
 - New capability to rework local mods and provide support for local service
- **Simplified migration of pre-installed z/VM products**
 - Allows disks associated with the pre-installed products on your first-level z/VM V5.1 system to be made available to your second-level z/VM V5.2 system
 - Transfers the following file types to your second-level system:
 - Customized files
 - Local modifications
 - Service
 - User-created files residing on selected disks



New Book in the z/VM Version 5 Product Library

"Getting Started with Linux on zSeries"

- **Introduced with z/VM V5.1; intended for new z/VM users**
- **Provides an explanation of z/VM basics, including how to configure and use z/VM functions and facilities**
- **Focus is on creating and managing Linux virtual machines**
- **Subject material includes:**
 - Configuring, administering, and servicing a z/VM system
 - Configuring TCP/IP for z/VM
 - Creating and cloning Linux virtual machines
 - Setting up basic system automation
 - Monitoring performance and capacity
 - Diagnosing z/VM and Linux problems
- **A PDF version of the book is available at: ibm.com/zseries/zvm**
- **Other good resources to help you get started with Linux on z/VM:**
 - “z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9” – available at: www.redbooks.ibm.com/abstracts/sg246695.html
 - “z/VM and Linux on IBM System z: The Virtualization Cookbook for Red Hat Enterprise Linux 4” – available at: www.redbooks.ibm.com/abstracts/sg247272.html



z/VM V5.2 Post-GA Support

- **Support for new processor instructions and facilities***
 - CP exploitation of Program-Event-Recording 3 (debug) and Store-Clock-Fast facility (overhead reduction)
 - Guest support for Signal Processor (SIGP) instruction Conditional-Emergency-Signal, Sense-Running-Status orders, and Program-Event-Recording 3
 - CP Trace support for new instructions
- **FICON Express4 (4 Gbps FICON/FCP)**
 - Improved capacity and performance with faster channel link data rates
 - More manageable migration to higher performance with 1/2/4Gb auto-negotiating links
- **Performance Toolkit enhancements***
 - Support for the Open Systems Adapter for NCP (CHPID type OSN)
 - Shutdown and restart of the Performance Toolkit is no longer required to add VM systems within the enterprise for performance-data retrieval
- **DS8000 and DS6000 support enhancements***
 - Recognition of unique real controller types (2105, 2107 and 1750)
 - Virtualization support to identify the level of controller support for z/VM guests
- **Query support for FCP N_Port ID virtualization***
 - New query capability allows z/VM users and guest operating systems to query hardware-defined virtual port names

*Requires the PTF for APAR VM63952

z/VM V5.2 Post-GA Support (continued)

- **System and guest exploitation of real HiperSockets IPv6 support***
 - Includes support for simulated HiperSockets devices (HiperSockets Guest LANs)
 - z/VM TCP/IP support for IPv6 transmissions over real and simulated HiperSockets
- **OSA GARP VLAN Registration Protocol (GVRP) support***
 - Allows registration of VLAN IDs; OSA-Express can become a participant or end station in a GVRP network
 - Provides VLAN filtering for z/VM Layer 2 and Layer 3 switches (can reduce unwanted inbound VLAN traffic)
 - Support details: VLAN tagging for z/OS, z/VM and Linux; VLAN priority queuing for z/OS and Linux
 - Also supported by z/VM V5.1 with PTFs for APARs VM63784 and PK08444
- **Parallel Access Volume (PAV) support***
 - Provides full support for 2105/2107 PAV feature for minidisk I/O
 - z/VM I/O dispatcher will exploit real PAV hardware features for balancing workload across PAV alias volumes
- **Collaborative memory management assist (CMMA)****
 - Extends coordination of memory and paging between Linux and z/VM to the level of individual pages
 - Improves efficiency of managing virtual memory page allocations
 - IBM is working with its Linux distribution partners for exploitation support

* Requires the PTF for APAR VM63952

** Requires the PTF for APAR VM63856 (planned availability 4Q06); also requires PTF for APAR VM63952

z/VM Support for Parallel Access Volumes

- **PAVs allow:**
 - Multiple concurrent I/Os to the same volume by one or more users or jobs
 - Automatic coordinated Read and Write I/O referential integrity when needed
- **z/VM V5.2 with PTF for APAR VM63952:**
 - Supports PAVs as minidisks for guest operating systems that exploit the PAV architecture (e.g., z/OS and Linux for System z)
 - Provides the potential benefit of PAVs for I/O issued to minidisks owned or shared by guests that do not support native exploitation of PAVs, such as z/VSE, z/TPF, CMS, or GCS
- **IBM TotalStorage DASD volumes must be defined to z/VM as:**
 - 3390 Model 2, 3, or 9 on a 3990 Model 3 or 6 Controller
 - Or...2105, 2107, or 1750 Storage Controller
 - Note: 3380 track-compatibility mode for the 3390 Model 2 or 3 is also supported.
- **Potential benefit:**
 - Improve I/O response times by reducing device queuing delays

Linux and z/VM Technology Exploitation

Collaborative Memory Management

- 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

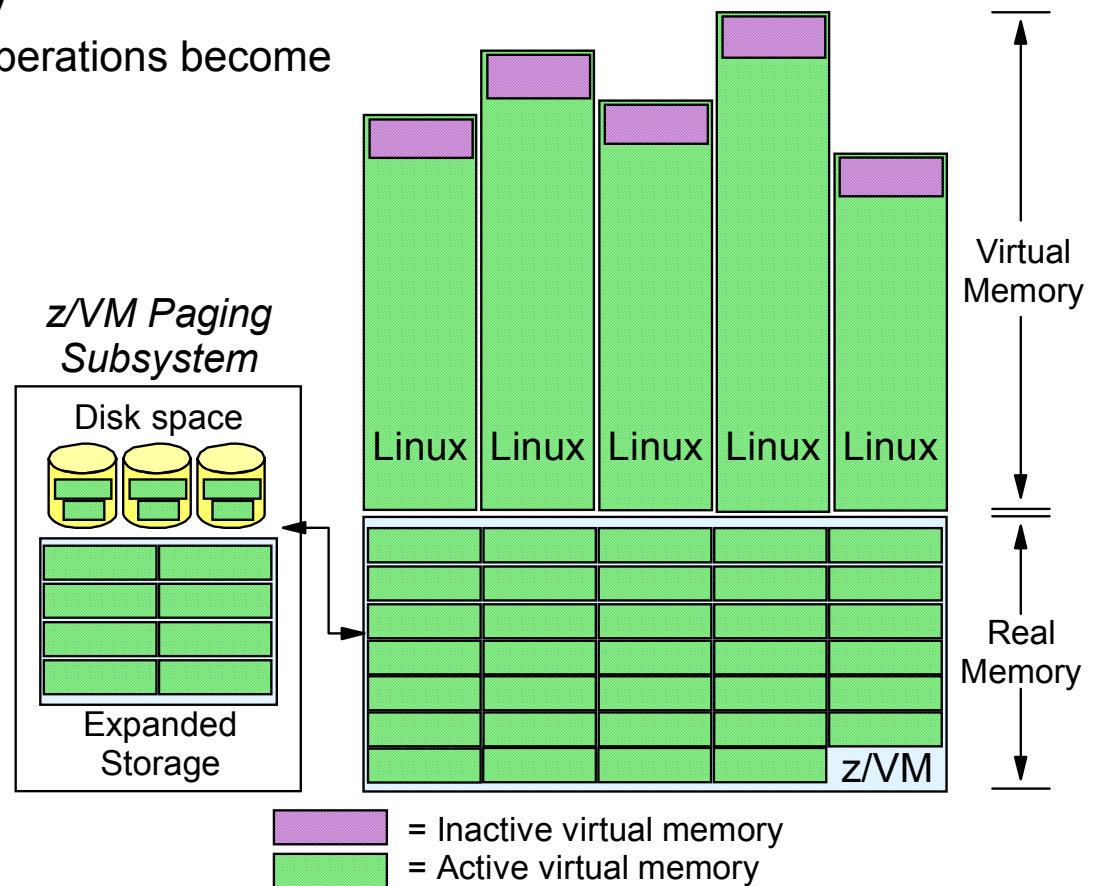


Chart 1 of 2

Linux and z/VM Technology Exploitation

Collaborative Memory Management

- Solution: real memory constraint detected and Linux images signaled to reduce virtual memory consumption
- Linux memory pages are released
- Demand on real memory and z/VM paging subsystem is reduced
- Overall system performance and guest image throughput improves
- z/VM V5.2 support available with PTF for APAR VM64085


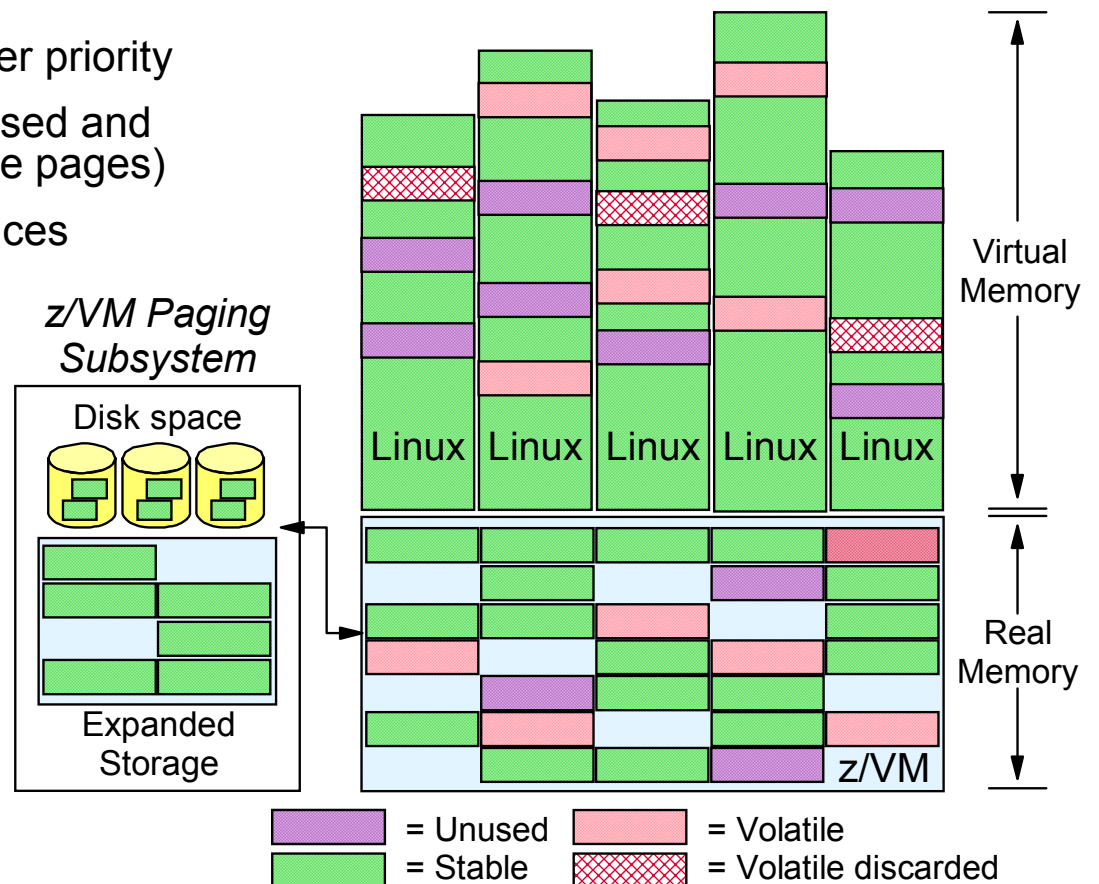


Chart 2 of 2

Linux and z/VM Technology Exploitation

Collaborative Memory Management Assist

- Solution: exchange page usage information between Linux guests and z/VM
- Reclaim “unused” pages at higher priority
- Bypass host page writes for unused and “volatile” pages (clean disk cache pages)
- Signal exception if guest references discarded volatile page
- Use host page management assist to re-instantiate pages for next use
- Supported by System z9 only



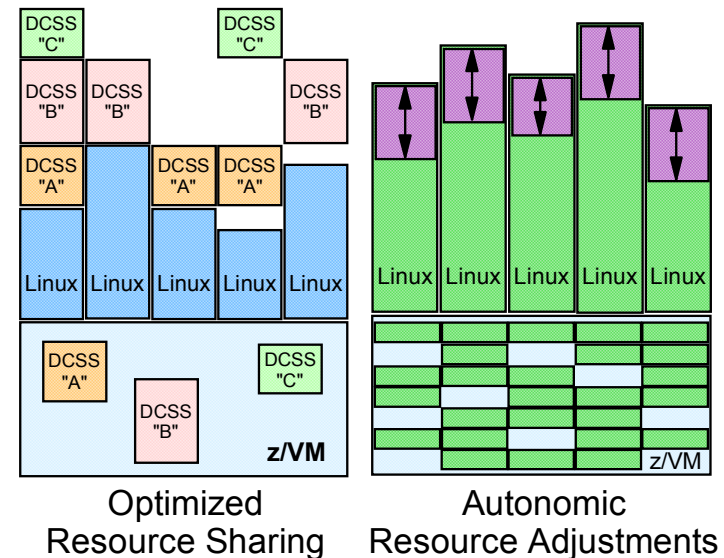
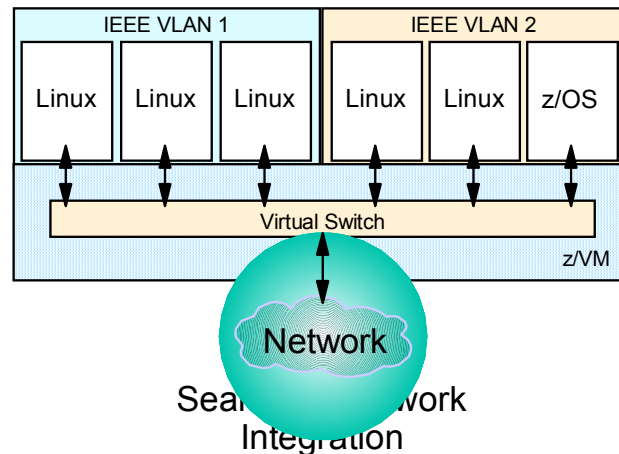
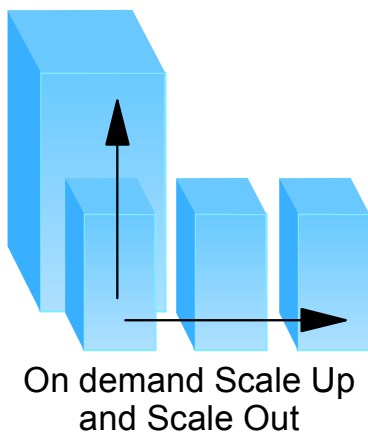

z/VM Statements of Direction

- IBM plans to remove the **ROUTED** and **BOOTP** servers from a future release of z/VM. z/VM V5.2 is planned to be the last release in which these servers will be available.

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

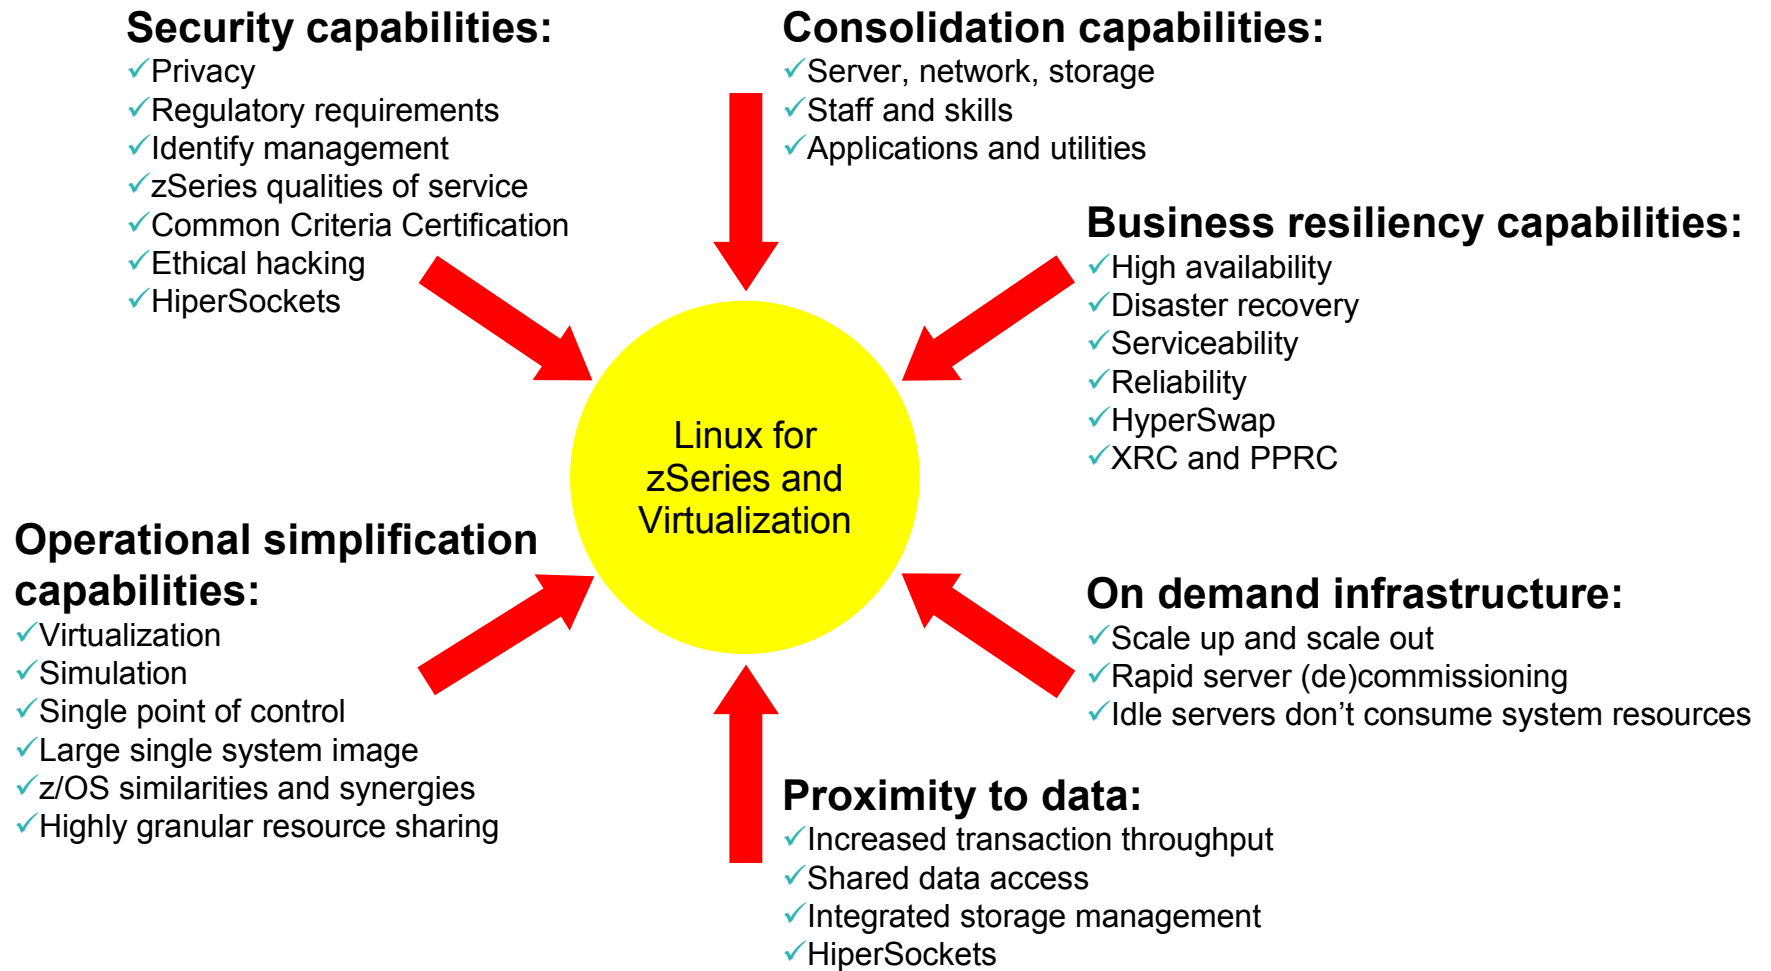
z/VM Vision and Investment Strategy

- **Make z/VM virtualization technology pervasive on zSeries**
 - Improved performance, scalability, operational ease of use
 - Complementary investments in LPAR (PR/SM) and z/VM
- **Deliver platform solutions that exploit Linux and z/VM synergies**
 - On demand scale up and scale out
 - Seamless network integration
 - Optimize resource sharing
 - Autonomic resource adjustments



Linux and z/VM on System z9 and zSeries

Providing Unmatched Value Propositions for Linux Workloads



Thank you

For more information, please contact
Reed A. Mullen
mullenra@us.ibm.com
+1 607 429 3824

ibm.com/zSeries



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DB2, e-business logo, ESCON, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, System z9, VM/ESA, VSE/ESA, WebSphere, xSeries, z/OS, zSeries, z/VM.

The following are trademarks or registered trademarks of other companies

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.

LINUX is a registered trademark of Linux Torvalds in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

Intel is a registered trademark of Intel Corporation.

* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.