z/VM Platform Update

March 2, 2014 Version 3.1b

Brian W. Hugenbruch, CISSP
z/VM Security Design and Development
bwhugen@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BladeCenter*	FICON*	Performance Toolkit for VM	Storwize*	System z10*	zSecure
DB2*	GDPS*	Power*	System Storage*	Tivoli*	z/VM*
DS6000*	HiperSockets	PowerVM	System x*	zEnterprise*	
DS8000*	HyperSwap	PR/SM	System z*	z/OS*	
ECKD	OMEGAMON*	RACF*	System z9*		

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.
TEALEAF is a registered trademark of Tealeaf, an IBM Company.
Windows Server and the Windows logo are trademarks of the Microsoft group of countries.
Worklight is a trademark or registered trademark of Worklight, an IBM Company.
UNIX is a registered trademark of The Open Group in the United States and other countries.

* Other product and service names might be trademarks of IBM or other 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.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”).

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Acknowledgments – Platform Update Team

- Alan Altmark
- Bill Bitner
- Miguel Delapaz
- Glenda Ford
- John Franciscovich
- Les Geer
- Susan Greenlee
- Dan Griffith
- Brian Hugenbruch
- Emily Hugenbruch
- Romney White

Agenda

- Release Status and Information
- z/VM® Version 6 Release 2
- z/VM Version 6 Release 3
  - 2014 Enhancements
- Futures and Statements of Direction

Release Status and Information

z/VM Release Status Summary

UPDATE

z/VM	Level	GA	End of Service	End of Marketing	Minimum Processor Level	Security Level
Version 6	Release 3	7/2013	4/2017		IBM System z10®	EAL 4+[2] OSPP-LS
	Release 2	12/2011	12/2016[3]	3Q/2013	IBM System z10®	-
	Release 1	10/2009	4/2013	12 2011	IBM System z10®	EAL 4+ OSPP-LS
Version 5	Release 4	9/2008	12/2014[1]	3/2012	IBM eServer zSeries 800& 900 (z800, z900)	-
	Release 3	6/2007	9/2010	9/2010	z800, z900	EAL 4+ CAPP/LSPP

[1] Or later (Announced August 7, 2012)
[2] Targeted Security Level in V6.3 SOD
[3] Extended from original date (Announced February 4, 2014)

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

Extended support contracts are available.

z/VM Version 5 Release 4

- The last release of z/VM to support IBM System z9® and older processors
  - **No longer available as of March 12, 2012**
  - Also supports the IBM zEnterprise® EC12 (zEC12) and IBM zEnterprise BC12 (zBC12)

- End of Service was been extended to **December 31, 2014** or end of IBM service for System z9, whichever is ***later***
  - Statement of Direction 2013
    - The zEC12 and zBC12 will be the last processors to support z/VM V5.4

z/VM Version 6
Security Certification Plans

- Common Criteria (ISO/IEC 15408)
  - z/VM V6.1 has been certified: BSI-DSZ-CC-0752
  - Evaluated to EAL 4+ for the Operating System Protection Profile (OSPP) with:
    - Virtualization extension (-VIRT)
    - Labeled Security extension (-LS)

- Federal Information Processing Standard (FIPS) 140-2
  - z/VM V6.1 System SSL is FIPS 140-2 Validated(TM)
  - Enablement requirements for certificate database and servers
    - http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2012.htm#1735

- z/VM V6.2 is designed to conform to both Common Criteria and FIPS 140-2 evaluation requirements

- z/VM V6.3 is **currently being evaluated** to both Common Criteria and FIPS 140-2, per IBM Statement of Direction (July 2013)

A Certification Mark of NIST, which does not imply product endorsement by NIST, the U.S. or Canadian Governments.

z/VM Pricing

- z/VM pricing consists of:
  - A one-time charge (OTC) per value unit
  - An annual charge for Service and Support, per value unit
- Number of value units is determined by number of engines, shown below on left.
- Prices are set per value unit, relative prices are illustrated below on right.
- The SSI feature includes LGR and it is priced in line with the RACF® feature

z/VM Value Unit Schedule	
Number of Engines	**Value Units per Engine**
1 to 3	**10**
4 to 6	**9**
7 to 9	**8**
10 to 12	**7**
13 to 16	**6**
17 to 20	**5**
21 to 25	**4**
26 and above	**3**

Relative Prices per Value Unit

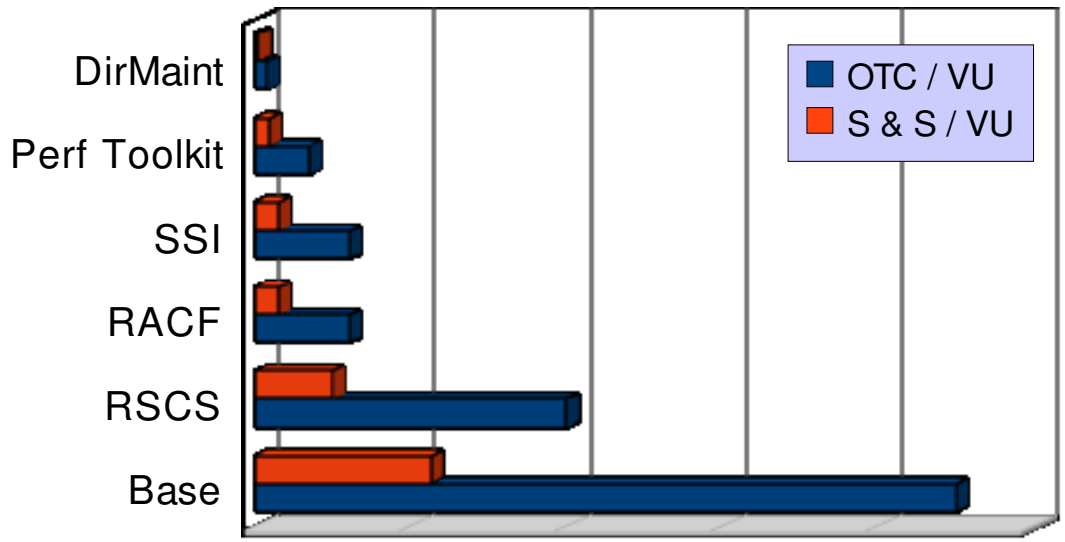

For illustrative purposes only. See IBM or Business Partner for details.

z/VM Version 6 Release 2

and other recent functional enhancements

z/VM Version 6 Release 2

- Generally available December 2, 2011
- Withdrawn from marketing in July 2013
  - Concurrent with z/VM V6.3 GA
- Major changes include:
  - Single System Image
  - Live Guest Relocation

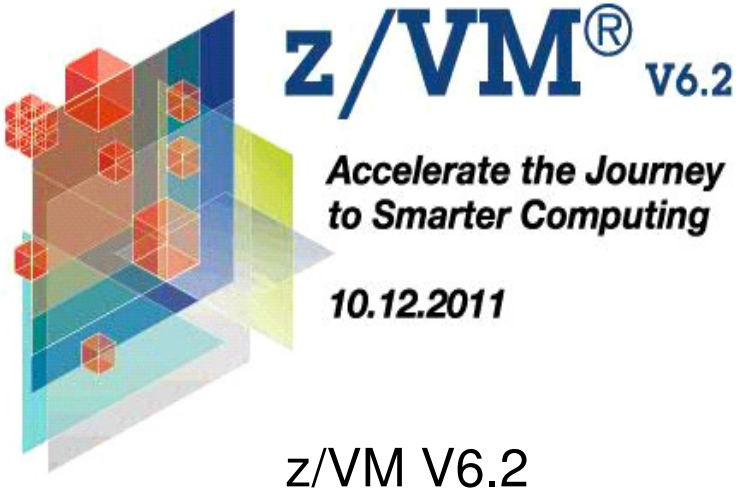

z/VM V6.2
End of Service
December 31, 2016

Single System Image (SSI) Feature
Clustered Hypervisor with Live Guest Relocation

- Optional priced feature
- Connect up to four z/VM systems as members of a Single System Image cluster
- Cluster members can be run on the same or different System z servers
- Simplifies management of a multi-z/VM environment
  - Single user directory
  - Cluster management from any member
    - Apply maintenance to all members in the cluster from one location
    - Issue commands from one member to operate on another
  - Built-in cross-member capabilities
  - Resource coordination and protection of network and disks

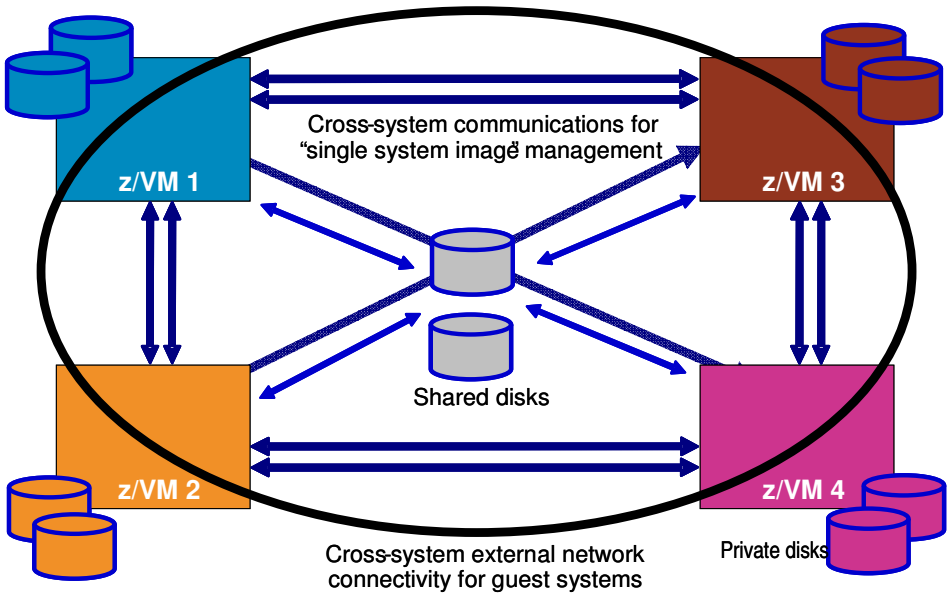

SSI Cluster Management – Features for Greater Reliability

- Cross-checking of configuration details as members join cluster and as resources are used:
  - SSI membership definition and identity
  - Consistent definition of shared spool volumes
  - Compatible virtual network configurations (MAC address ranges, VSwitch definitions)
- Cluster-wide policing of resource access:
  - Volume ownership marking to prevent dual use
  - Coordinated minidisk link checking
  - Autonomic minidisk cache management
  - Single logon enforcement
- Communications failure “locks down” future resource allocations until resolved
- Comprehensive checking for resource and machine feature compatibility during relocation:
  - Adjustment of “virtual architecture level” to support customer relocation policy

Single System Image Feature
Clustered Hypervisor with Live Guest Relocation

- Dynamically move Linux® guests from one member to another with Live Guest Relocation
  - Reduce planned outages
  - Enhance workload management
  - Non-disruptively move work to available system resources and non-disruptively move system resources to work

- When combined with Capacity Upgrade on Demand, Capacity Backup on Demand, and Dynamic Memory Upgrade, you will get the best of both worlds

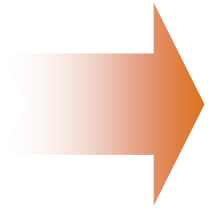

Bring additional resources to the workload!

Move the workload to the resources!

Benefits and Uses of z/VM SSI Clusters

- Horizontal growth of z/VM workloads
  - Increased control over virtual server sprawl
  - Distribution and balancing of resources and workloads
- Flexibility for planned outages for service and migration
    - z/VM
    - Hardware
  - Less disruptive to virtual server workloads
- Workload testing
  - Different service/release levels
  - Various environments (stress, etc.)
  - New/changed workloads and applications can be tested before moving into production
- Simplified system management of a multi-z/VM environment
  - Concurrent installation of multiple-system cluster
  - Single maintenance stream
  - Reliable sharing of resources and data

z/VM System Management – Related Products

- **Operations Manager for z/VM V1.4**
  - Facilitates automated operations
  - Monitor, view, and interact with consoles without logging on to service machines or Linux guests
  - Take actions based on service machine console messages and other system events
  - Schedule events for immediate execution or on a regular schedule
- **OMEGAMON® XE on z/VM and Linux V4.2**
  - Performance monitoring of z/VM and Linux guests
  - Part of the Tivoli® Management Services, including Tivoli Enterprise Portal
  - Uses IBM Performance Toolkit for VM as its data source
- **Backup and Restore Manager for z/VM V1.2**
  - Backup and restore file level data for CMS minidisks and Shared File System
  - Backup and restore images of Linux guests and/or z/VM volumes
    - Use Tivoli Storage Manager for file level backup and restore of Linux data
- **Tape Manager for z/VM V1.3**
  - Manage tapes: retention, access control, data security erase
  - Manage devices: share with other z/VM and non-z/VM systems
  - Manage mount requests for ATL, VTS, and manual mount devices
- **Archive Manager for z/VM V1.1**
  - Users and administrators manage disk space more efficiently and effectively
  - Archive infrequently used or large files to tape or other disk
- **zSecure™ Manager for RACF z/VM V1.11.1**
  - Automate complex, time consuming z/VM security management tasks
  - Quickly identify and prevent problems in RACF
  - Create comprehensive audit trails

Systems Management and the Single System Image Feature

- Support for SSI and LGR by
  - xCAT 2.8 supports SSI and LGR
    - Available with March 2013 code drop

- No Support for SSI and LGR by
  - Unified Resource Manager (zManager) *does not* support SSI and LGR
  - IBM Systems Director *does not* support SSI and LGR

- Suggested best practice is to not combine SSI and LGR with zManager or Systems Director
  - Work with your IBM Sales Team, IBM Lab Services, or z/VM Development Lab to determine which technologies are most critical to your environment and business.

z/VM 6.2 SSI and GDPS Support

- GDPS® V3.10 was generally available March 29, 2013
  - Adds support for Live Guest Relocation with Standard Actions panel and Automation scripts
  - If an SSI member is being shutdown, GDPS automatically asks if any virtual machines should be relocated prior to the shutdown.
  - See http://www-03.ibm.com/systems/z/advantages/gdps/whatsnew.html for details.
- GDPS/PPRC xDR 3.9 supports SSI configuration
  - All Members of the SSI cluster must be in the same GDPS managed group.
  - APAR PM64211
  - Requires z/VM APAR VM65176
- All supported GDPS releases (3.8, 3.9, & 3.10) supported with non-SSI environment
  - See GDPS PSP buckets for required service (z/OS®, Linux, & z/VM)
  - If GDPS environment shared with older z/VM releases, z/VM service is required on them before adding z/VM 6.2

Environment	3.8	3.9	3.10
z/VM 6.2 Non-SSI	Yes	Yes	Yes
z/VM 6.2 SSI	No	Yes	Yes
z/VM 6.2 SSI + LGR	No	No	Yes

z/VM Single System Image and Live Guest Relocation Implementation Services

- IBM System z® **Lab Services Offering**:
  - In-depth **education** on the functions of Single System Image
  - Cluster **planning and deployment assistance**
  - Operational **guidance** and recommendations
  - **Migration assistance** for users of CSE
  - **Demonstrate** the technology in your own environment.
  - **Analyze** how SSI and LGR will affect your system initialization, recovery, and automation procedures
  - Early **identification of any inhibitors** to use
  - Identification of any required z/VM or Linux operating system **patches**

For more information, contact **systemz@us.ibm.com**

HiperSocket VSWITCH Bridge

Available: April 13, 2012

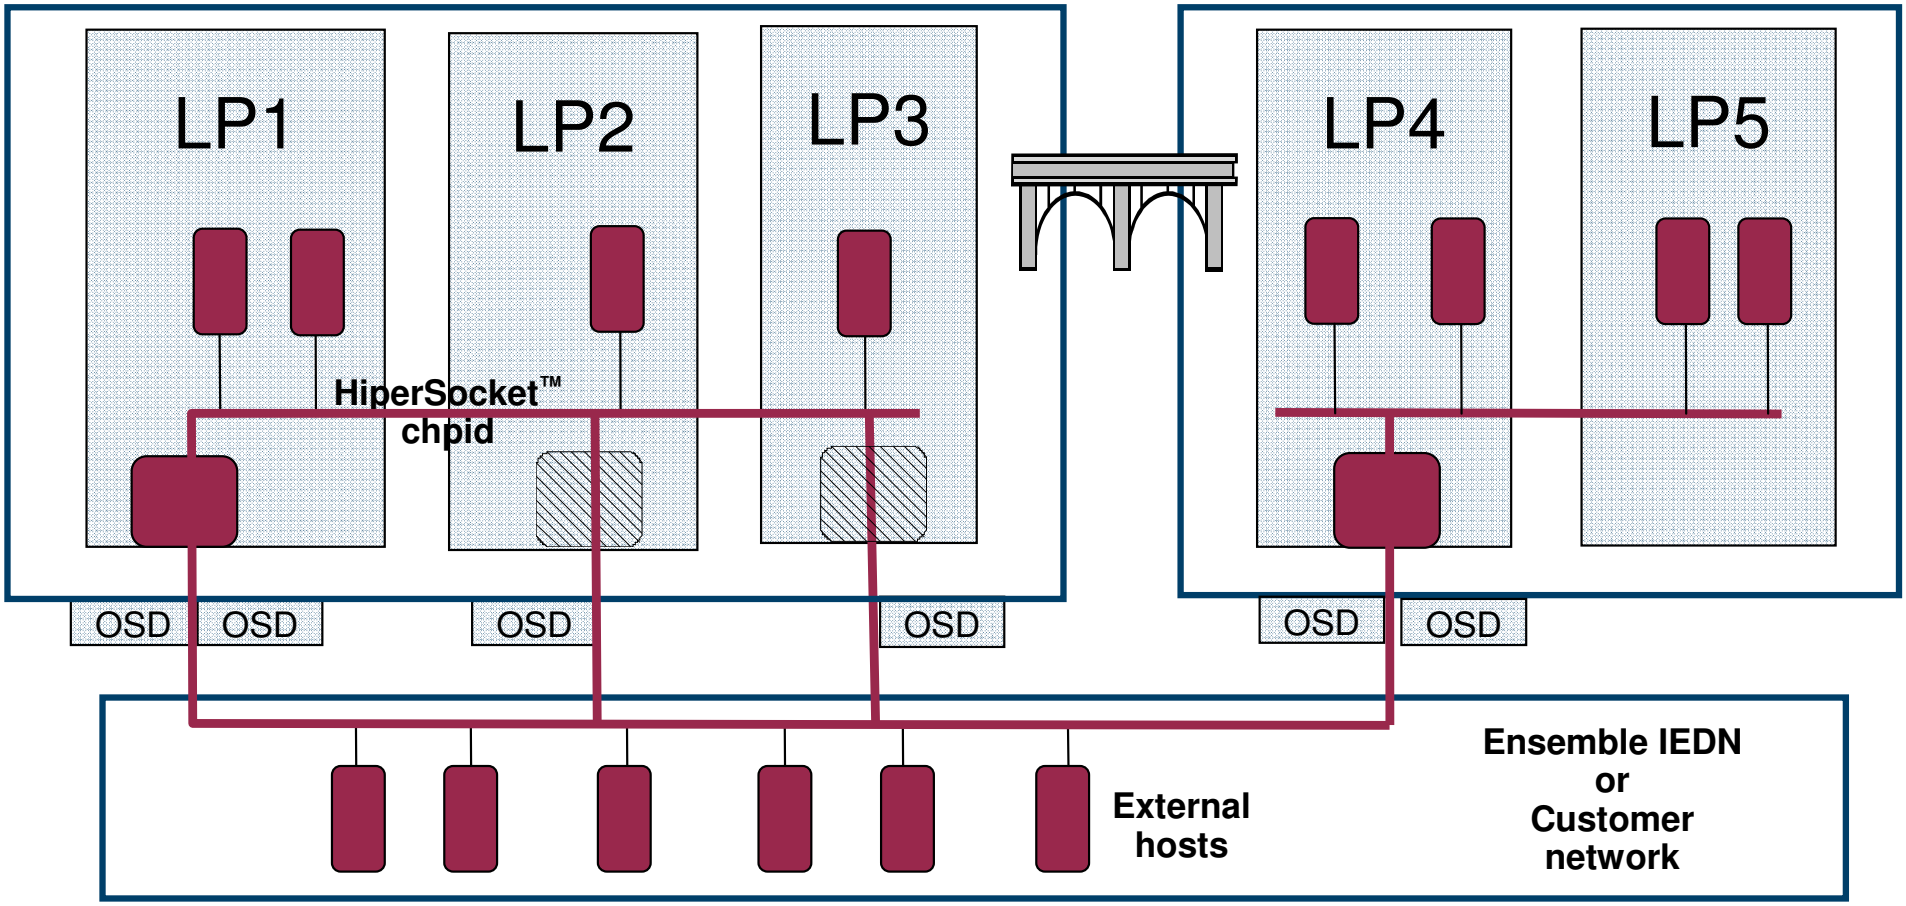

- Built-in failover and failback
- CHPARM=x2 to bridge IEDN
- CHPARM=x4 to bridge customer network
- Same or different LPAR
- One active bridge per CEC
- PMTU simulation

HiperSocket VSWITCH Bridge

Available: April 13, 2012

- Virtual Switch bridge between Ethernet LAN and HiperSockets
  - zEnterprise IEDN (OSX) or customer network (OSD) connections
  - Guests can use simulated OSA or dedicated HiperSockets
  - VLAN aware
  - One HiperSocket chpid only
- Full redundancy
  - Up to 5 bridges per CEC
  - One bridge per LPAR
  - Automatic takeover
  - Optionally designate one “primary”
    - Primary will perform “takeback” when it comes up
  - Each bridge can have more than one OSA uplink
- CP: VM65042 / UM33691, TCP/IP: PM46988 / UK77220
- z/OS currently does not support the HiperSockets Bridge

High Performance FICON

Available: April 13, 2012

- Enable guests to use High Performance FICON for System z (zHPF)
  - Different I/O model
  - Single and multiple track I/O
  - CP APAR VM65041 PTF UM33646
  - DVF APAR VM65144 PTF UM33647
- Requires host and control unit compatibility
  - Consult a storage specialist for details
- z/OS and Linux provide exploitation
- Performance results available at:
  - http://www.vm.ibm.com/perf/reports/zvm/html/620jb.html

z/VM Version 6 Release 3

z/VM Version 6 Release 3
Making Room to Grow Your Business

February 5
Preview announcement

July 23
Product announcement

July 26
Product General Availability

Continued Evolution
- Improved TCO
- Scaling and efficiency improvements
- Ease of usability

See http://www.vm.ibm.com/zvm630/

2013

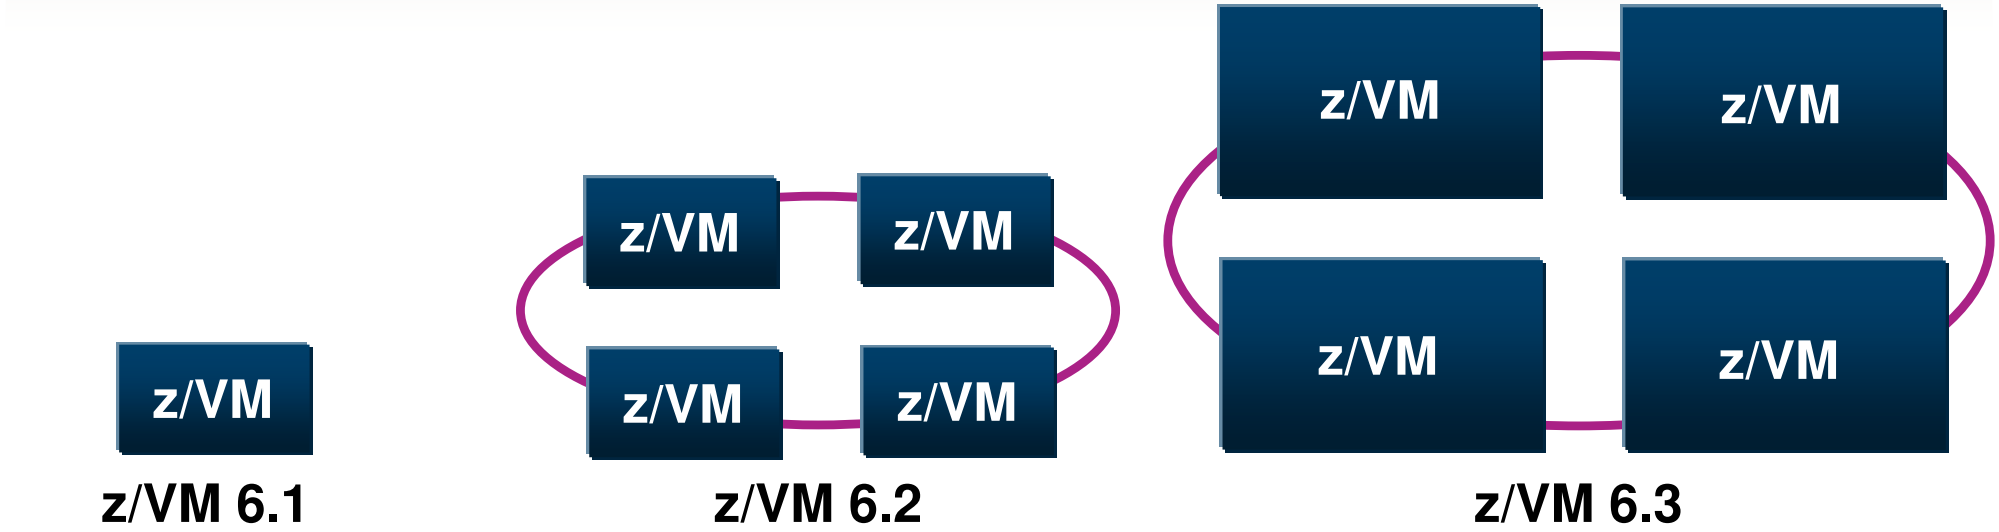

Reduce the number of z/VM systems you need to manage z/VM 6.3

- Expand z/VM systems constrained by memory up to four times (almost two times on the zBC12), thus increasing the number of Linux virtual servers in a single z/VM system
- Exploit HiperDispatch to improve processor efficiency, allowing more work to be done per IFL and therefore supporting more virtual servers per IFL, potentially requiring fewer systems for applicable workloads
- Expand the real memory used in a Single System Image Cluster up to 4 TB
  - z/VM 6.3 has the ability to fully utilize memory of a zBC12 at a maximum of 496 GB
- Exploit multiple subchannel sets in GDPS environment to place secondary Peer-to-Peer volumes in alternate subchannel set

Improved Memory Management Flexibility and Efficiency

- Benefits for z/VM systems of all memory sizes
- Prioritize virtual server use of real memory more effectively through enhanced memory reservation support
- Exploit improved management of memory on systems with diverse virtual server processor and memory use patterns
- Eliminate use of expanded storage for z/VM paging, allowing greater flexibility and avoiding some of the restrictions associated with expanded storage

Simplify z/VM Systems Management

- Managing z/VM virtual servers with xCAT (Extreme Cloud Administration Toolkit) is ready to go after z/VM V6R3 installation; nothing else needs to be installed
- Adopt a foundation to allow future extensions for open source systems management solutions, in particular through OpenStack® support
- Enable scalable support for the larger systems that z/VM V6R3 supports
- Safely migrate an existing z/VM V6R2 SSI Cluster to z/VM V6R3 in a step-wise approach, without having to shut down the cluster, using the new “Installation Upgrade In Place” capability

“z/VM 6.3 Systems Management Fundamentals” by Brian Hugenbruch, CISSP
Tuesday, 2:45pm – Room U3

Large Memory Support

- Real memory limit raised from 256GB to **1 TB**
  - Proportionately increases total virtual memory based on tolerable over- commitment levels and workload dependencies
- Virtual machine memory limit remains unchanged at **1 TB**
- Paging DASD utilization and requirements change
  - Removed the need to double the paging space on DASD
  - Paging algorithm changes increase the need to have a properly configured paging subsystem
- Expanded Storage continues to be supported with a limit of **128 GB**

Large Memory Support (cont.)

- Reorder processing removed
  - Commands remain, but have no impact
  - Improves environment for running larger virtual machines

- Improved effectiveness of the CP SET RESERVE command
  - Stronger “glue” to hold reserved pages in memory
  - Support for reserving pages of NSS or DCSS
    - Example: Use with the Monitor Segment (MONDCSS)
  - Ability to limit the overall number of reserved pages for the system

“z/VM 6.3 Performance Update” by Dr. Brian Wade
Monday, 1:30pm – Room U3

Enhanced Dump Support

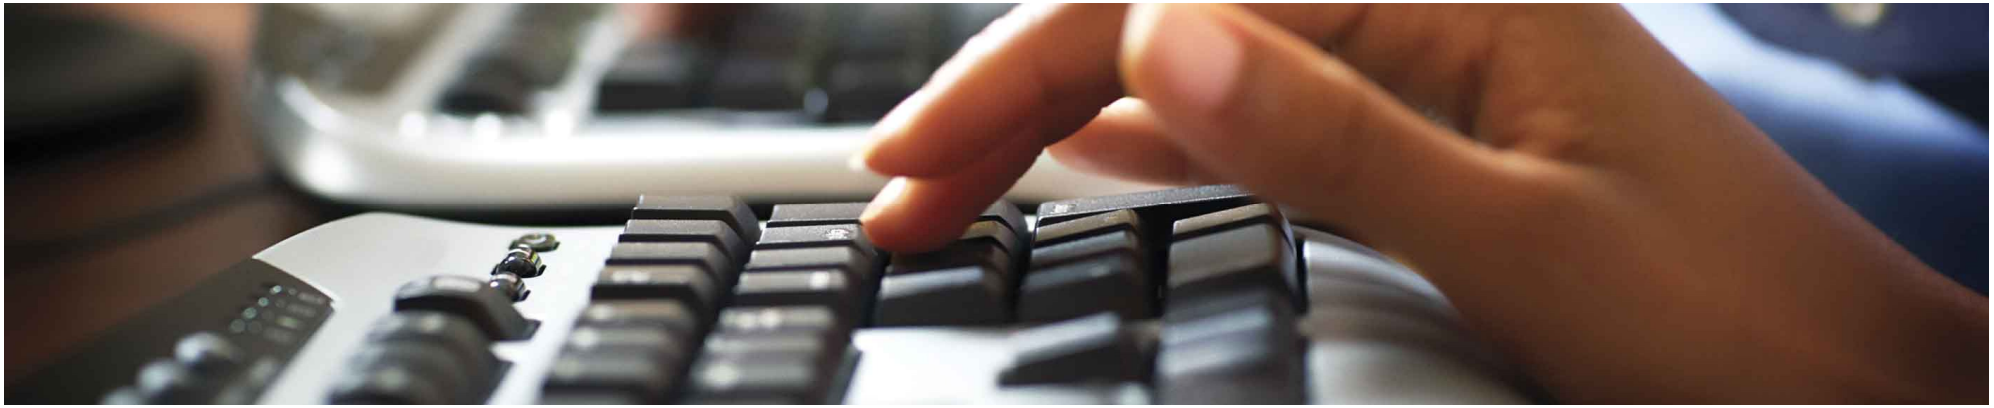

- Stand-alone Dump utility has been rewritten
  - Creates a CP hard abend format dump
  - Dump is written to ECKD™ or SCSI DASD
- Larger memory sizes supported, up to a maximum of 1 TB
  - Includes Stand-alone dump, hard abend dump, SNAPDUMP, DUMPLD2, and VM Dump Tool
- Performance improvements for hard abend dump
  - Reduces time to take a CP hard abend dump

HiperDispatch

- Improved processor efficiency
  - Better n-way curves
    - Supported processor limit of 32 remains unchanged
  - Better use of processor cache to take advantage of cache-rich system design of more recent machines

- Two components:
  - Dispatching affinity
  - Vertical CPU management

"z/VM 6.3 and HiperDispatch" by Dr. Brian Wade
Monday, 9:45pm – Room U3

HiperDispatch – Dispatching Affinity

- Processor cache structures become increasingly complex and critical to performance
- Goal is to re-dispatch work close (in terms of topology) to where it last ran

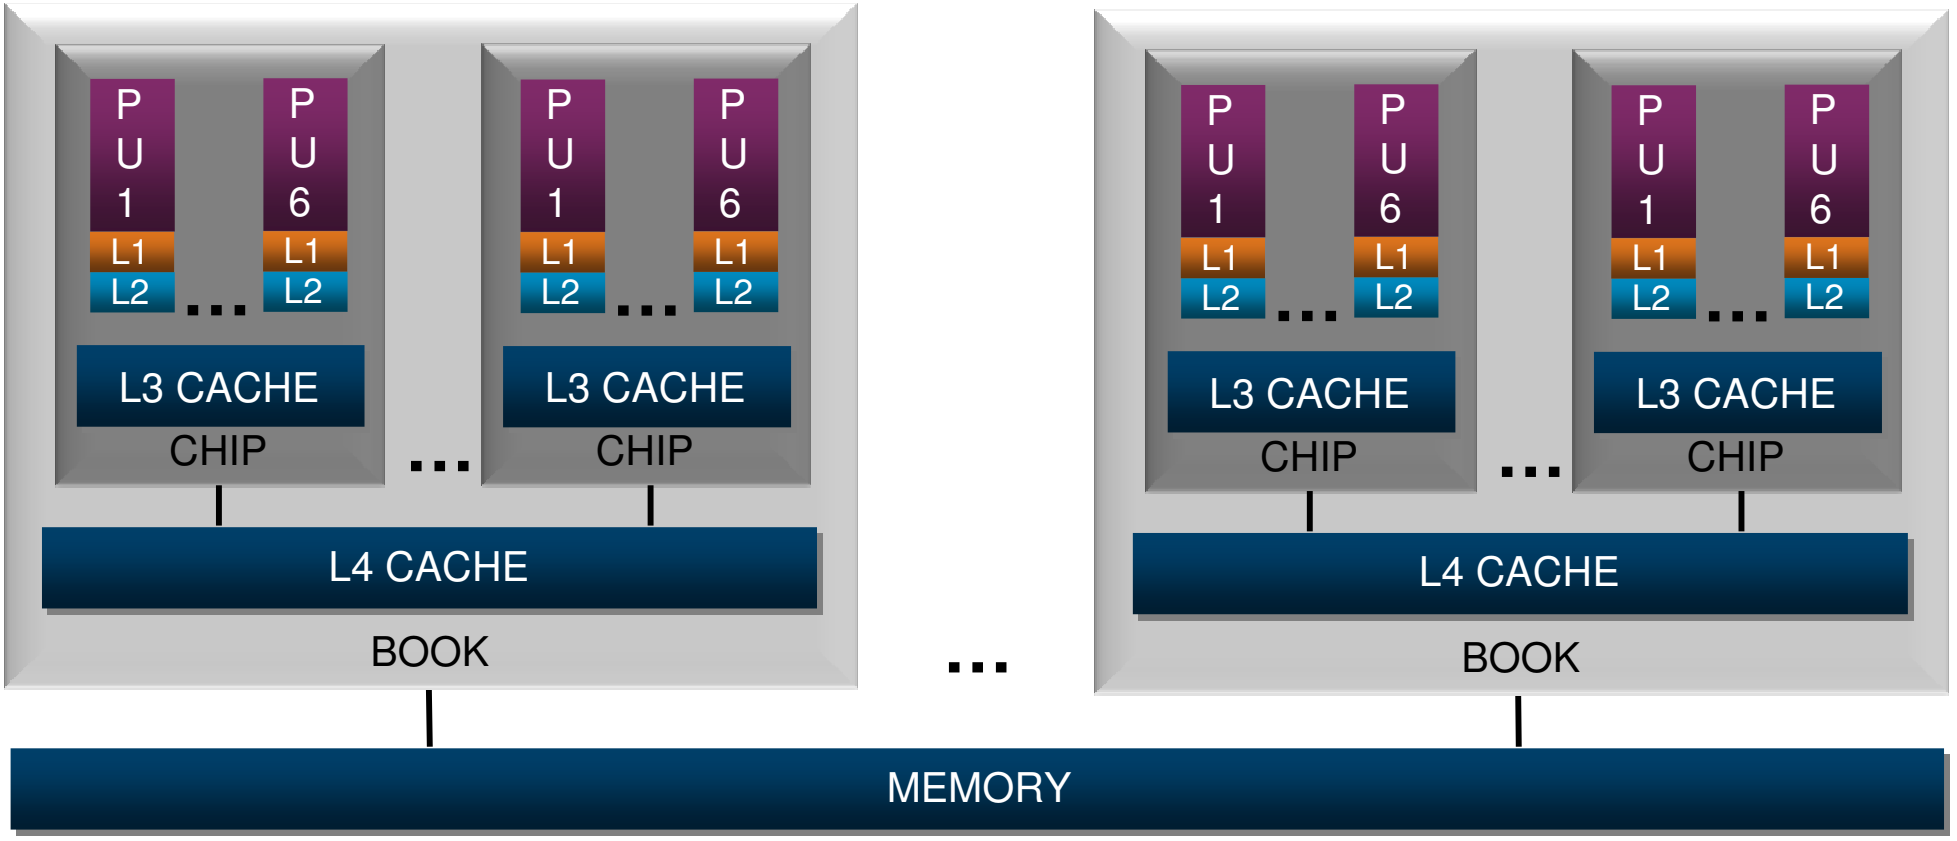

HiperDispatch – Dispatching Affinity

- Dispatcher is aware of the cache and memory topology
  - Dispatch virtual CPU near where its data may be in cache based on where the virtual CPU was last dispatched
- Better use of cache can reduce the execution time of a set of related instructions
- z/VM 6.2 and earlier uses “soft” affinity to dispatch virtual CPUs
  - No awareness of chip or book

HiperDispatch – Vertical CPU Management

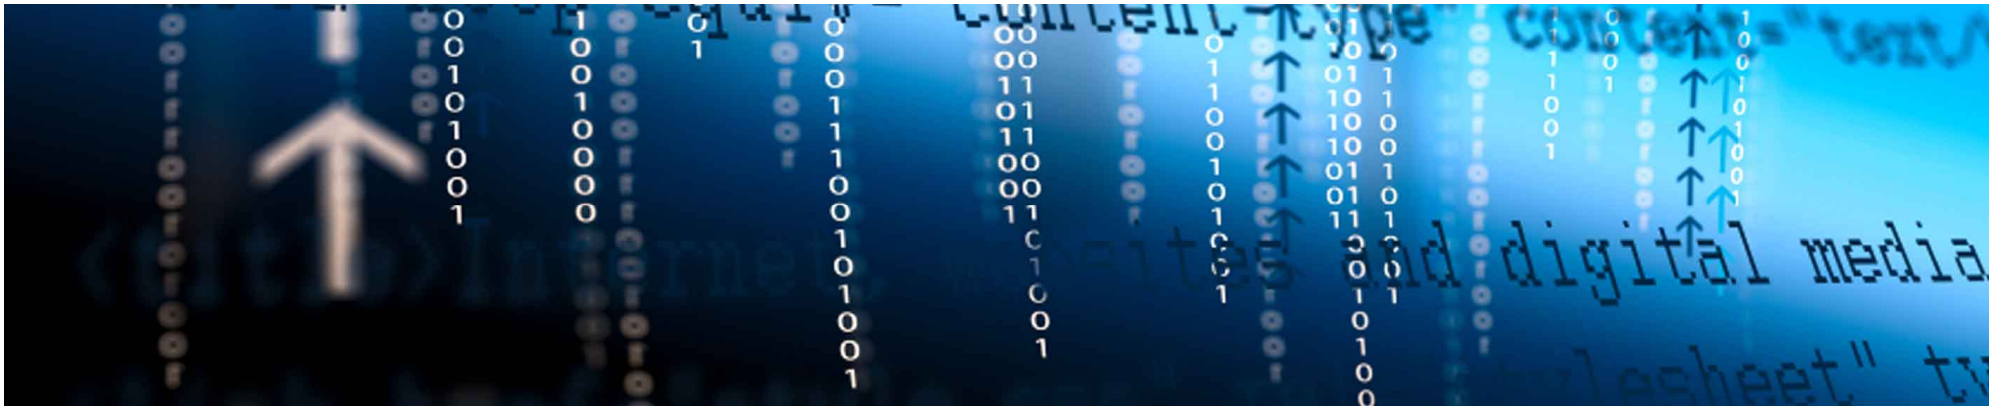

- Today's “horizontal” management distributes the LPAR weight evenly across the logical processors of the z/VM LPAR
- “Vertical” management attempts to minimize the number of logical processors, allowing LPAR to similarly manage logical CPUs

Example:

- Ten Physical IFLs, seven logical IFLs, weight of 400 out of 1000
  - Each logical IFL (LPU) entitled to 57% of an IFL
- When CEC is constrained, the LPAR’s entitlement is reduced to four IFLs, so seven is more than required
- z/VM and LPAR will cooperate
  - z/VM will concentrate the workload on a smaller number of logical processors
  - LPAR will redistribute the partition weight to give a greater portion to this smaller number of logical processors (~100% of four CPUs)

Horizontal vs. Vertical CPU Management

Horizontal:
- The logical processors are all created/treated equally.
- z/VM dispatches work evenly across the seven logical processors

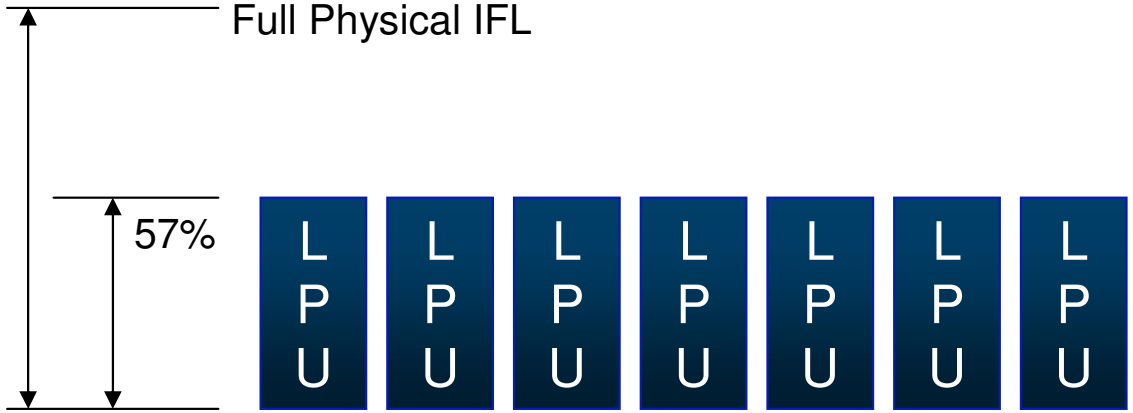

Vertical:
- The logical processors are skewed to where some get greater share of the weight.
- z/VM dispatches work accordingly to the heavier weighted workload.

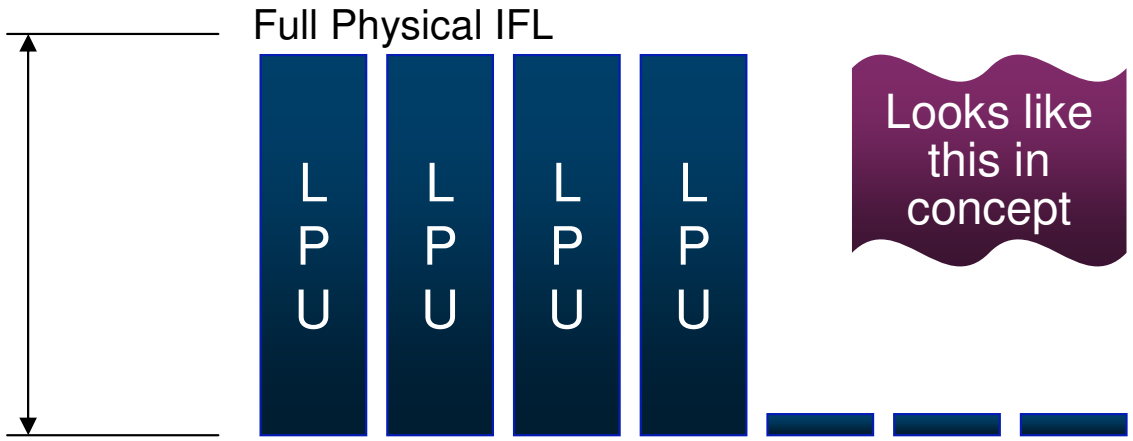

Technology Exploitation

- Fibre Channel Protocol Data Router Support
  - FCP QEBSM support enhanced for guest support use of FCP Data Router
- FICON DS8000 Series New Functions
  - Storage Controller Health message
    - New attention message from hardware providing more details for conditions in past reflected as Equipment Check.
    - Intended to reduce the number of false HyperSwap® events.
  - Peer-to-Peer Remote Copy (PPRC) Summary Unit Check
    - Replaces a series of state change interrupts for individual DASD volumes with a single interrupt per LSS
    - Intended to avoid timeouts in GDPS environments that resulted from the time to process a large number of state change interrupts.
    - Satisfies a SOD from October 12, 2011
- Multiple Subchannel Set (MSS) support for mirrored DASD
  - Support to use MSS facility to allow use of an alternate subchannel set for Peer-to-Peer Remote Copy (PPRC) secondary volumes.
  - Satisfies a SOD from October 12, 2011

z/VM 6.3 and GDPS Support

- z/VM 6.3 alternate subchannel set support
  - GDPS V3.10 prereqs the PM71447 New Function: GDPS/PPRC XDR MSS1 Support APAR
- z/VM 6.3 FICON DS8000 Series new function (DS8K synergy initiative)
  - GDPS/PPRC V3.8, V3.9, & V3.10 and prereqs the PM44141 New Function: GDPS/PPRC XDR PPRCSUM and Storage Controller Health Message APAR, and DS8K R6.2 u-code.
- Cannot mix new MSS support in an SSI environment with older z/VM systems.
- See http://www-03.ibm.com/systems/z/advantages/gdps/whatsnew.html for details.
- See GDPS PSP buckets for required service (z/OS, Linux, and z/VM)
  - Remember to check for required service for systems that share the GDPS environment.

Environment	3.8	3.9	3.10
z/VM 6.3 w/ MSS 1	No	No	Yes[1]
z/VM 6.3 DS8K Synergy	Yes[1]	Yes[1]	Yes[1]
z/VM 6.3 SSI + LGR	No	No	Yes[1]

1 – with appropriate service – Check Bucket

Virtual Networking Improvements

- Live Guest Relocation support for port-based virtual switches built on existing support:
  - Allow relocation of port-based interface
  - Prevent relocation of an interface that will be unable to establish proper network connectivity
  - Adjust the destination virtual switch configuration, when possible, by inheriting virtual switch authorization from the origin
- MPROUTE server upgraded to z/OS V1.13 OMPROUTE functional equivalency
- Support for OSA-Express5S devices
- Virtual Switch recovery and stall prevention
  - New SET VSWITCH UPLINK SWITCHOVER command
  - Change from current device to one of the configured backup devices

Security Enhancements

- Crypto Express4S
  - Guest support for Crypto Express4S which is a feature available on zEC12 and zBC12
  - Can be configured in one of three ways:
    - IBM Common Cryptographic Architecture (CCA) Coprocessor mode
    - IBM CCA Accelerator mode
    - IBM Enterprise Public Key Cryptographic Standards (PKCS) #11 (EP11) coprocessor
- SSL Server Upgrade
  - System SSL update to z/OS V1.13 equivalency
    - Transport Layer Security (TLS) protocol, Version 1.2 and SHA2 certificate support
    - TLS Protocol Selection
  - IPv6 support for SSL-enabled Telnet, FTP, and SMTP

"z/VM Security News and How To's" by Brian W. Hugenbruch, CISSP
Monday, 4:00pm – Room U3

Installation Upgrade in Place Enhancement

- Upgrade an existing z/VM 6.2 system to z/VM 6.3 with minimal impact to the current running system.
  - Fewer manual steps such as directory merging and new virtual machine creation

- Upgrade Approach:
  - Install new release as temporary second level system
  - Move new level of z/VM to current system
  - For SSI Cluster, start with single member of the cluster on new level

- Provides a backup to support backing out in extreme cases

- Support for local modifications

"z/VM 6.3 Upgrade In Place Installation" by Richard Lewis
Monday, 5:15pm – Room U3

"z/VM 6.3 Installation Lab" by Richard Lewis
Tuesday, 11:00am onward – Lab Room

Linux Disk Dump Utility can now include the NSS

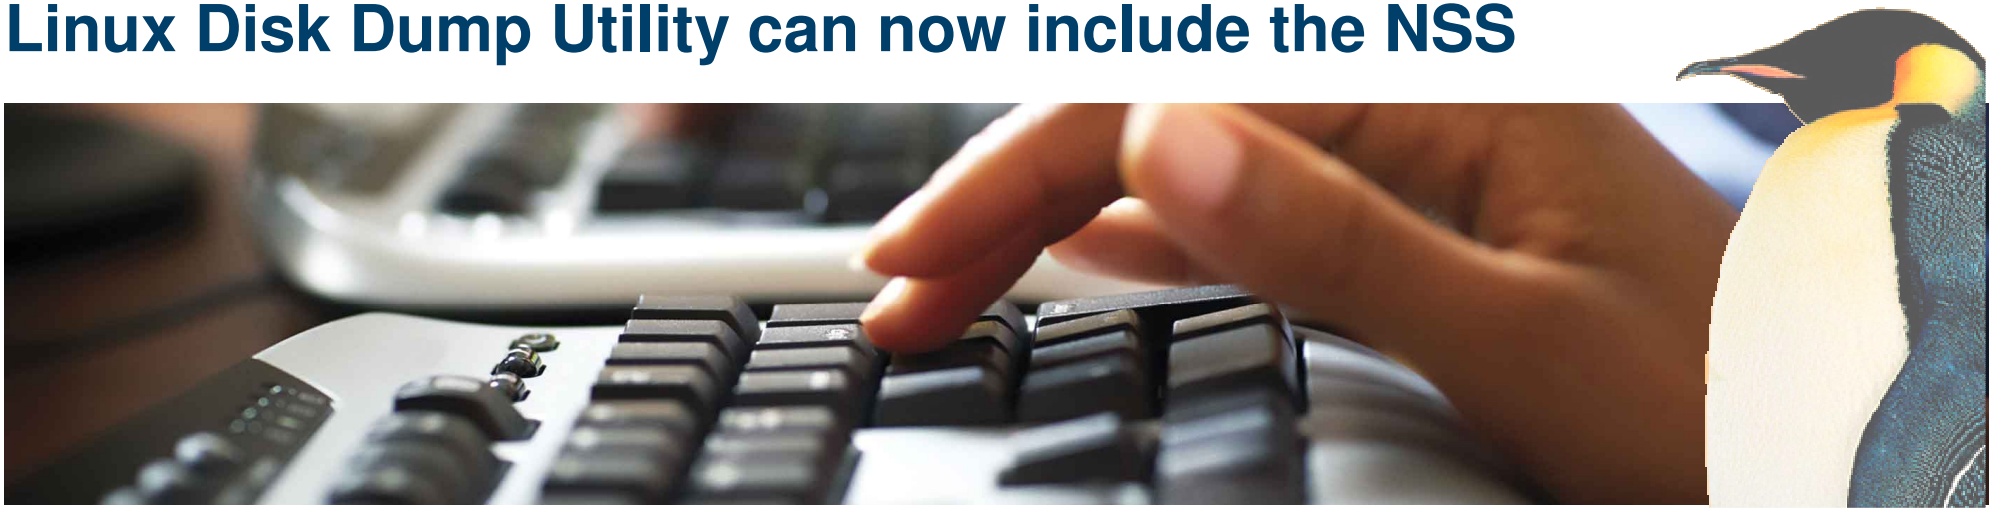

- The Linux Disk Dump utility is preferred over the CP VMDUMP command in most cases.
- Previously, the contents of an NSS could not be captured with Linux Disk Dump utility.
- Changes in IPL now allow the NSS to be included
  - New NSSDATA parameter
- For more background, see:
  - http://download.boulder.ibm.com/ibmdl/pub/software/dw/linux390/docu/l26ddt01.pdf for Linux Disk Dump utility information
  - http://www.vm.ibm.com/perf/tips/vmdump.html for information on differences between VMDUMP and Linux utility

z/VM 6.3 Withdraws Cross System Extensions (CSE) Support

- Satisfies a previous Statement of Direction
- The z/VM Single System Image (VMSSI) feature replaces the functions provided by CSE:
  - Logon once in the cluster, with exceptions
  - Cross-system MESSAGE and QUERY commands
  - Shared spool
  - Shared source directory
- VMSSI has additional value such as autonomic minidisk cache management and a single point of maintenance
- XLINK shared disk support is **not** affected.

Change from SoD

z/VM 6.3 Withdraws support for TCP/IP Devices and Daemons

- Satisfies a previous Statement of Direction
- A220 HYPERchannel devices
- CLAW devices
- DHCP daemon
- LPSERVE (LPD)
  - RSCS LPD is provided at no charge
  - Does not affect LPR (client)

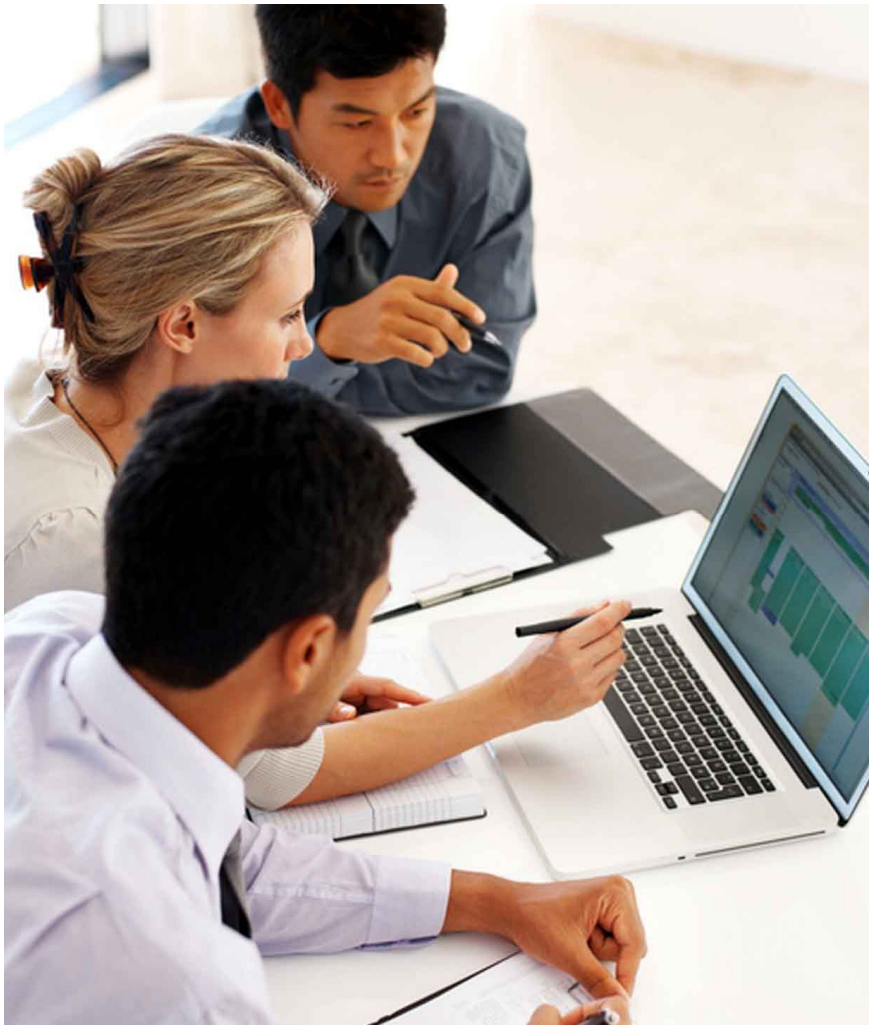

Hybrid computing model integrated and enabled for Cloud

Service Management Layer (provided by SmartCloud technologies)

openstack CLOUD SOFTWARE

z/VM | zManager | kvm | Hyper-V | Power VC | FSM | VMware

Datawarehousing IBM DB2® Analytics Accelerator Solution

zManager for z/OS and IBM zEnterprise BladeCenter® Extension (zBX)

Systems Director for Power® IBM System x® and storage

FSM for Intel® and Power ITEs

Third party Managers and Servers

The OpenStack Food Chain

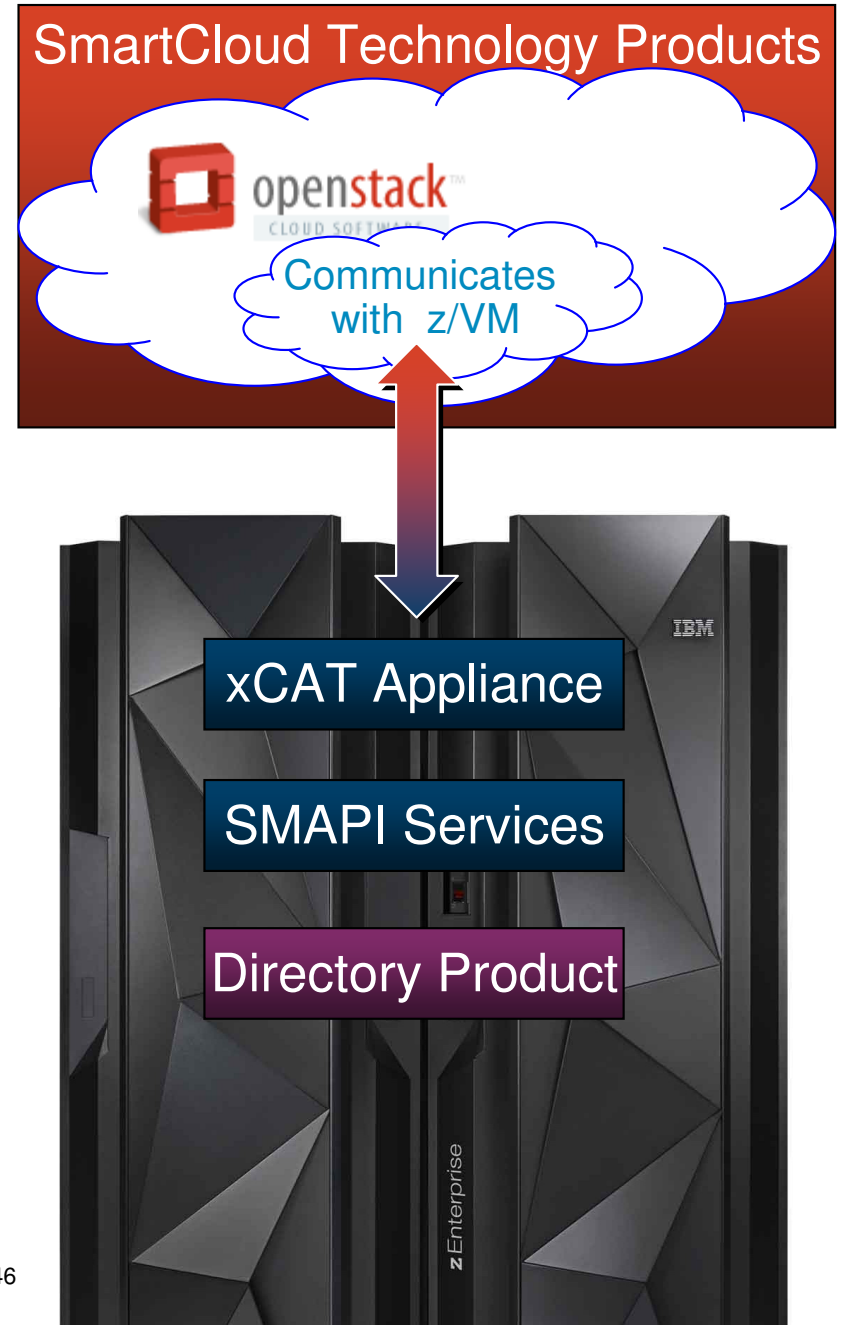

- **Top Half of the Solution:**
  - An IBM SmartCloud Technology product or other vendor product will include the OpenStack support.
  - Portions of that OpenStack support will know z/VM (i.e. code that connects and understands how to talk to z/VM).
- **Bottom Half of the Solution:**
  - Rest APIs are used to communicate with the OpenStack code from the top half.
  - The xCAT Appliance utilizes new and existing Systems Management APIs (SMAPI) to interact with the z/VM system
  - SMAPI can interact with additional optional products or features (e.g. a directory manager).

Product with OpenStack Support

z/VM 6.3 Product

Optional Product or Feature

IBM Wave for z/VM V1.1 (IBM Wave)

- IBM Wave is a new virtualization management product for z/VM® and Linux® virtual servers that uses visualization to dramatically automate and simplify administrative and management tasks
- Enhanced Enterprise Linux Server (ELS*) solution is also available with IBM Wave for z/VM
- New! Jumpstart Services to help customers get started with IBM Wave
- Read the announcement here!
- General availability - February 28th, 2014

Supported IBM System z® processors: IBM System z10® Enterprise Class (z10 EC™), IBM System z10 Business Class™ (z10 BC™) and later

Supported z/VM versions/releases:
- z/VM 6.3
- z/VM 6.2
- z/VM 5.4

"System z Virtualization Management with IBM Wave for z/VM"
Tuesday, 9:45am – Vendor Room (see agenda)

Enterprise Linux Server is an integrated solution comprised of Hardware, Hypervisor, Memory, Easy to Use Virtualization Management and 3 years Service and Support designed to get customers including FIEs started with low cost scalable computing environment

IBM Wave for z/VM (IBM Wave) Overview

IBM Wave simplifies and helps automate management and administration of z/VM and Linux virtual servers, jumpstarting the steps needed to get to cloud. With its content rich interface IBM Wave extends the reach of your staff and lets you manage z/VM and Linux intuitively and cost effectively, reducing reliance on deep expert skills.

- Monitors and manages virtual servers and resources from a single interface
- Simplifies and automates administration and management tasks
- Provisions virtual resources (Guests, Network, Storage)
- Supports advanced z/VM capabilities such as Single System Image and Live Guest Relocation
- Allows delegation of administrative capabilities to the appropriate teams

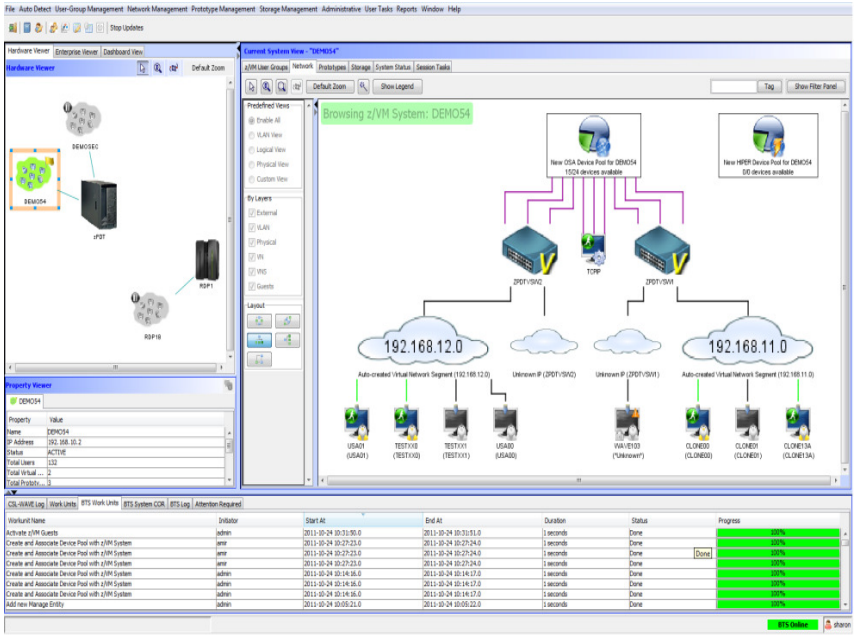

A simple, intuitive graphical management, provisioning, and automation tool to help you fully leverage the power of System z virtualization on z/VM.

Unified Resource Manager (zManager) and z/VM 6.3 Announcement

In light of IBM's cloud strategy and adoption of OpenStack, the management of z/VM environments in zManager is now stabilized and will not be further enhanced.

Accordingly, zManager will not provide systems management support for z/VM 6.3. However, zManager will continue to play a distinct and strategic role in the management of virtualized environments created by integrated firmware hypervisors (PR/SM™, PowerVM ™, and System x hypervisor based on kvm) of zEnterprise.

Looking ahead, IBM's vision is to enable OpenStack to provide heterogeneous systems management across zEnterprise, z/VM and distributed platforms, which in turn can be exploited by IBM's future SmartCloud offerings.

z/VM System Management – Related Products

- **Operations Manager for z/VM V1.5**
  - Facilitates automated operations
  - Monitor, view, and interact with consoles without logging on to service machines or Linux guests
  - Take actions based on service machine console messages and other system events
  - Schedule events for immediate execution or on a regular schedule
- **OMEGAMON® XE on z/VM and Linux V4.3**
  - Performance monitoring of z/VM and Linux guests
  - Part of the OMEGAMON and IBM Tivoli Monitoring infrastructure, including Tivoli Enterprise Portal
  - Uses IBM Performance Toolkit for VM as its data source
- **Backup and Restore Manager for z/VM V1.2**
  - Backup and restore file level data for CMS minidisks and Shared File System
  - Backup and restore images of Linux guests and/or z/VM volumes
    - Use Tivoli Storage Manager for file level backup and restore of Linux data
- **Tape Manager for z/VM V1.3**
  - Manage tapes: retention, access control, data security erase
  - Manage devices: share with other z/VM and non-z/VM systems
  - Manage mount requests for ATL, VTS, and manual mount devices
    - Supports IBM and Oracle STK libraries
- **Archive Manager for z/VM V1.1**
  - Users and administrators manage disk space more efficiently and effectively
  - Archive infrequently used or large files to tape or other disk
- **zSecure™ Manager for RACF z/VM V1.11.1**
  - Automate complex, time consuming z/VM security management tasks
  - Quickly identify and prevent problems in RACF
  - Create comprehensive audit trails

Other Considerations with z/VM 6.3

- You need to plan for Large Memory and for HiperDispatch. z/VM 6.3 changes some of the rules of thumb and planning guidelines from previous releases.

- DUMP Considerations
  - At time of publishing the calculations for dump space was not complete for the largest systems. This information has been made available on the z/VM Home Page
    - http://www.vm.ibm.com/techinfo/ or http://www.vm.ibm.com/service/zvmpladm.pdf
  - Should learn DUMPLD2 which replaces DUMPLOAD and has ability to segment a dump into multiple files.

- The size of CMS component grew significantly as a result of including an appliance server for xCAT, LOHCOST, and Stand-alone dump
  - Two additional install volumes

- If using z/VM 6.3 Upgrade in Place installation ensure required service is applied to z/VM 6.2 system being upgraded.

February 24, 2014 Announcements

Enhancing the Foundation for Virtualization

- Release for Announcement – zBX and zEnterprise System Enhancements
  - February 24, 2014
- Software Enhancements
  - CPU Pooling
  - Environment Information Interface
- Hardware Support
  - 10GbE RoCE Express Feature
  - zEDC Express Feature
- Available June 27, 2014

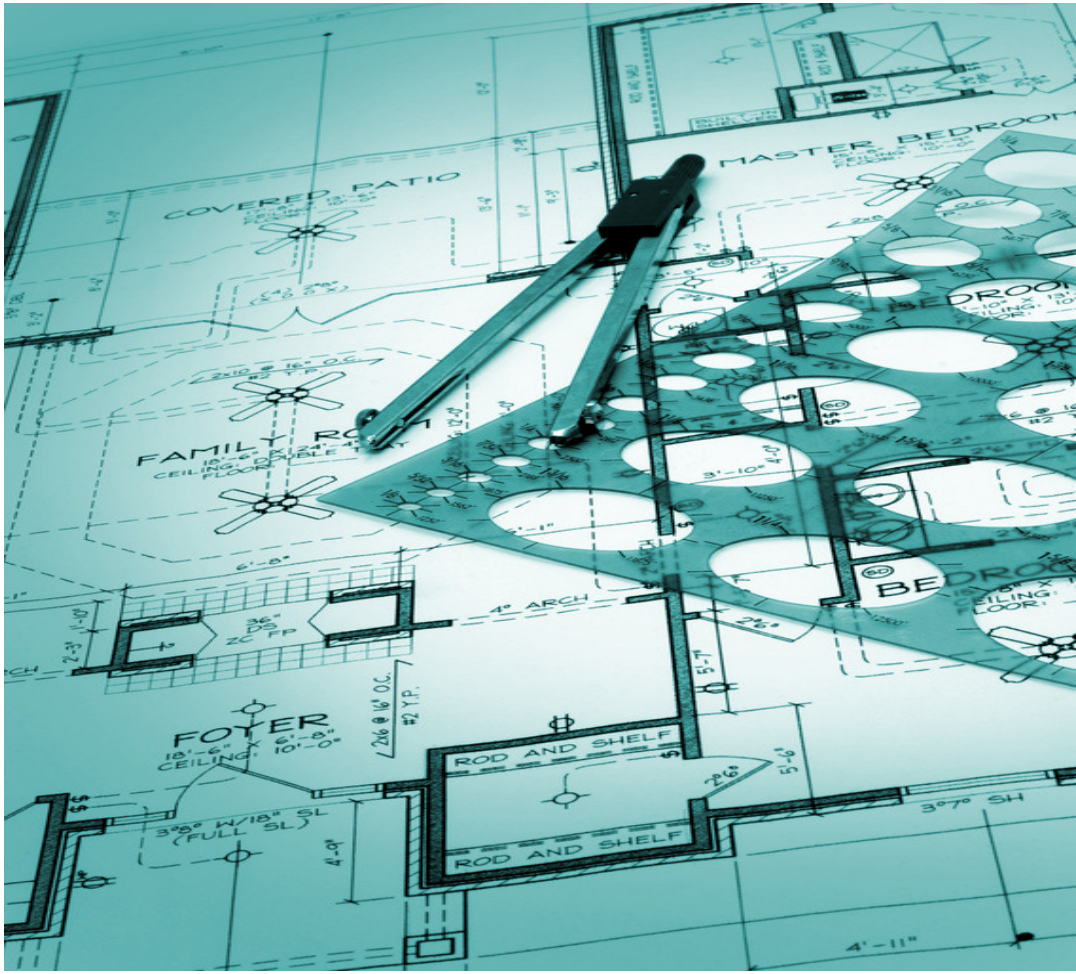

CPU Pooling

- Fine grain CPU limiting for a group of virtual machines
- Define one or more pools in which a limit of CPU resources is set.
- Two flavors of limits:
  - LIMITHARD - Percentage of system
  - CAPACITY – Number of CPUs
- Coexists with individual limit shares
  - More restrictive limit applies
- Support Details
  - z/VM 6.3 with APAR VM65418 - June 27, 2014

Environment Information Interface

- New interface allow guest to capture execution environment
  - Configuration and Capacity information
  - Various Levels:
    - Machine, logical partition, hypervisor, virtual machine
- New problem statement instruction Store Hypervisor Information (STHYI)
- Includes support for CPU Pooling enhancement
- Foundation for future software licensing tools
- Support details:
  - z/VM 6.3 with APAR VM65419 – June 27, 2014

10GbE RoCE Express Feature

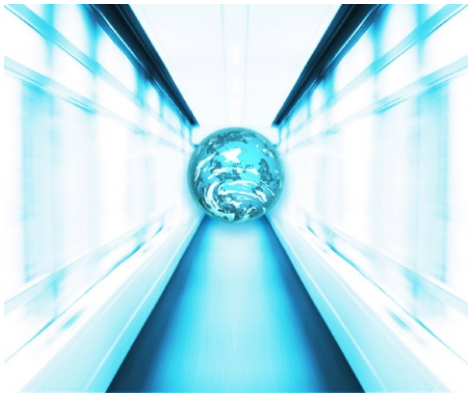

- Support for RDMA over Converged Ethernet for guests
- Based on new hypervisor PCIe support
- Designed to support z/OS's Shared Memory Communications-Remote Direct Memory Access (SMC-R) in z/OS V2.1
- Helps reduce CPU resource consumption
- Support details:
  - IBM zEC12 or zBC12 with appropriate millicode (driver 15)
  - z/VM 6.3 with APAR VM65417 – June 27, 2014
  - z/OS 1.12, z/OS 1.13, z/OS 2.1 with APAR OA43256
  - Fulfills 2013 Statement of Direction

zEDC Express Feature

- Guest support for zEnterprise Data Compression (zEDC) Express Feature
- High performance, low latency, low CPU consumption compression
- Possible disk utilization reduction
- Support details:
  - IBM zEC12 or zBC12 with appropriate millicode (driver 15)
  - z/VM 6.3 with APAR VM65417 – June 27, 2014
  - z/OS 1.12, z/OS 1.13, z/OS 2.1 with APAR OA43256
  - z/OS 1.12, z/OS 1.13, z/OS 2.1 with APAR OA44482
  - Fulfills 2013 Statement of Direction

Hardware Support

Support for IBM zEnterprise EC12

- **Updates for z/VM 6.2, 6.1, and 5.4**
  - VM65007 CP
  - VM65131 IOCP
  - VM65046 Performance Toolkit for VM™
  - VM65047 HCD
  - VM64747 HCM (z196 support: 6.1 and 5.4 only)
  - VM65130 EREP
  - OA38418 OSA/SF for OSA-Express4S
  - PM49761 High Level Assembler (new instructions)
- **PSP Bucket**
  - Upgrade **2827DEVICE**
  - Subset **2827/ZVM**
  - Subset **2827/ZOS** for ICSF service to support EP11 when running as a guest

Support for IBM zEnterprise BC12

- **Updates for z/VM 6.3, 6.2 and 5.4**
  - VM65239: VMHCD support
  - VM65236: VMHCM support
  - VM65279: EREP support
  - VM65278: IOCP support
  - VM65360: SYSEVENT QVS support
    - VM65356: SYSEVENT QVS support (pre-req to VM65360)
- **Update for z/VM 6.2 and in base of z/VM 6.3**
  - PM83966: TCP/IP support
- **PSP Bucket**
  - Upgrade: **2828DEVICE**
  - Subset: **2828/ZVM**

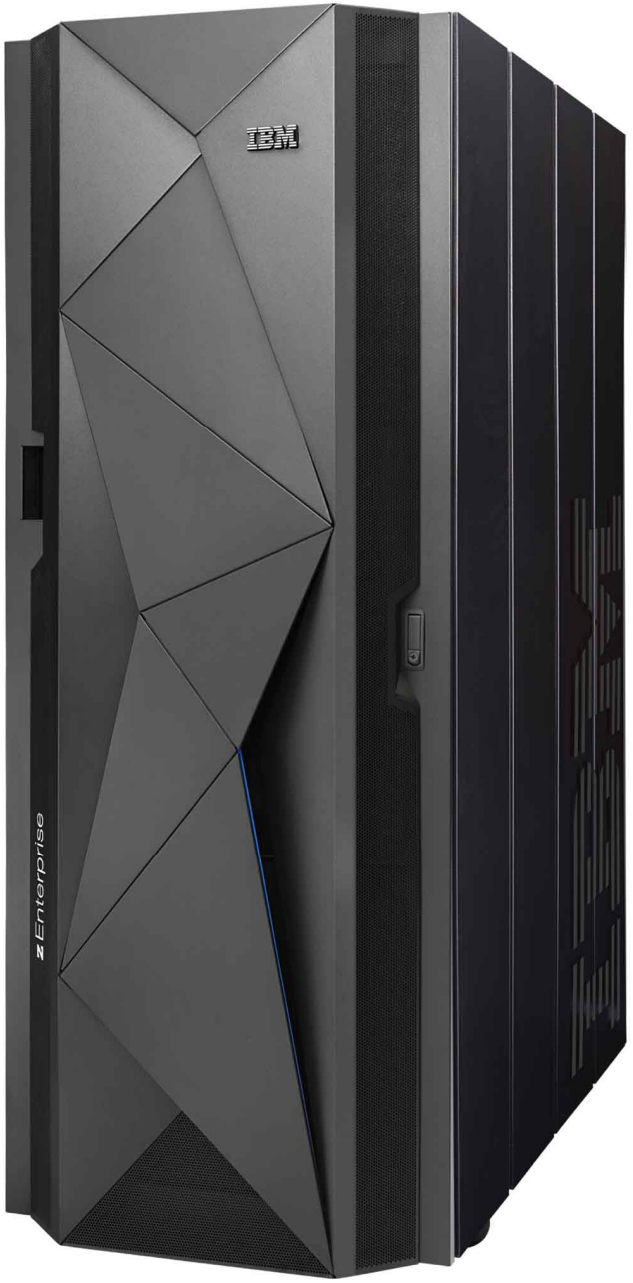

z/VM Disk Storage Support

- z/VM 6.3 supports
  - DS8000® Series (FCP or FICON®)
  - DS6000® Series (FICON)
  - XIV (FCP)
  - IBM San Volume Controller (FCP)
  - IBM Storwize® V7000 (FCP)
    - See **ibm.com**/support/docview.wss?uid=ssg1S1003703#_zvm
  - IBM FlashSystem when behind an SVC (FCP)
  - As well as many of the older storage devices

- The IBM System Storage® Interoperation Center (SSIC) support page:
  - **ibm.com**/systems/support/storage/ssic/interoperability.wss

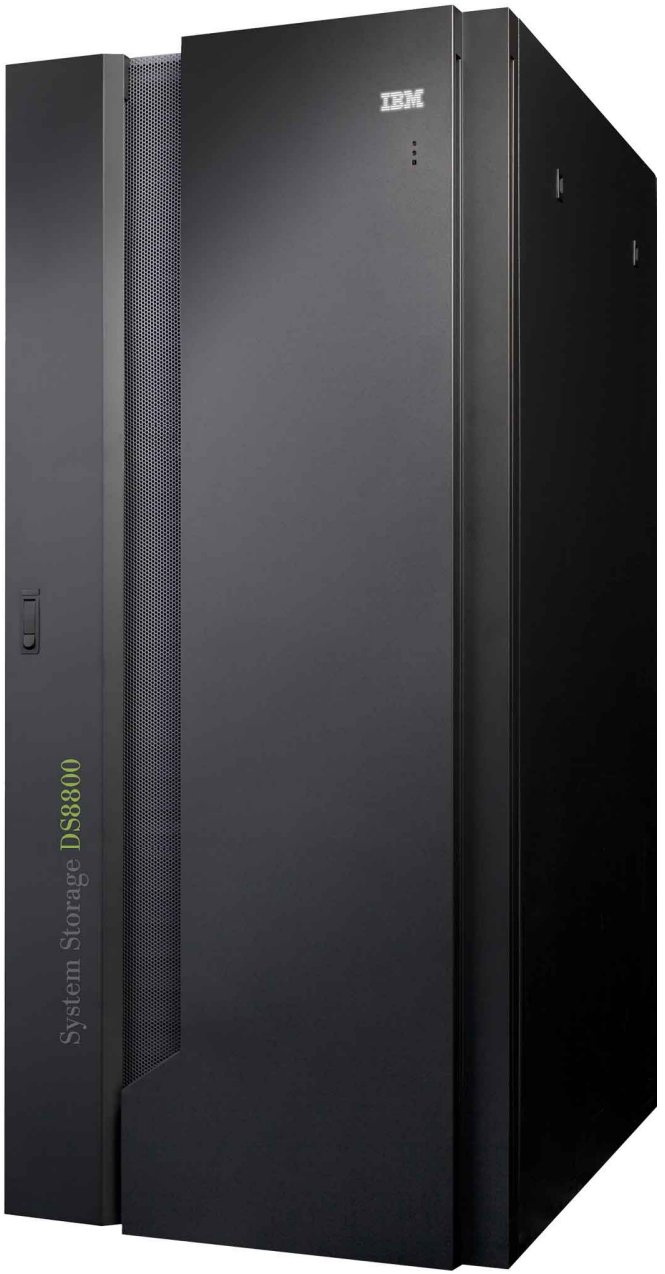

z/VM Tape Storage Support

- z/VM 6.3 Supports:
  - 3494 Virtual Tape Server (VTS) Library
  - TS3500 (3584) Tape Library
  - Virtualization Engine TS7700 (7720,7740) Tape Library
  - TS3400 Tape Auto-Stacker
  - Emulated 3490 Tape Subsystems
  - 3590, 3592, TS1120, TS1130, & TS1140 Enterprise Tape Subsystems

- z/VM provides CP native support for FICON only
  - FCP attachment supported by Linux guests via FCP subchannels
  - FICON supported by Linux for stand-alone tape only; no FICON library support

- The IBM System Storage® Interoperation Center (SSIC) support page:
  - **ibm.com**/systems/support/storage/ssic/interoperability.wss

Statements of Direction

July 23, 2013

Subject to change or withdrawal without notice, representing IBM goals and objectives only.

Statements of Direction We've Already Covered

- Security Certifications of z/VM 6.3 (in progress)
- Support of the 10GbE RoCE Express Feature
- Support of the zEDC Express Feature

Stabilization of z/VM 5.4 Support

The IBM zEnterprise EC12 and IBM zEnterprise BC12 are planned to be the last System z servers supported by z/VM V5.4 and the last System z servers that will support z/VM V5.4 running as a guest (second level). z/VM V5.4 will continue to be supported until December 31, 2014, or until the IBM System z9® Enterprise Class (z9 EC) and IBM System z9 Business Class (z9BC) are withdrawn from support, whichever is later. Refer to Withdrawal Announcement 912-144, (RFA56762) dated August 7, 2012.

- While support will continue to the later date of December 31, 2014 or until the z9 processors are withdrawn from future, support for new function and processors is being stabilized.
- z/VM 5.4 will not be supported on processors after the zEC12 and zBC12.
  - This includes running as a guest of a supported z/VM Version 6 release.
- Plan now to avoid a migration which would involve both hardware and software at the same time.

Withdrawal of Support for Expanded Storage

z/VM 6.3 will be the last release to support expanded storage (XSTOR) as part of the paging configuration. With the enhanced memory management support added in z/VM V6.3, expanded storage is no longer recommended as part of the paging configuration. z/VM can run efficiently in a configuration using only central storage

- In z/VM 6.3, it is recommended to configure all processor memory as central storage.
  - Support remains to use expanded storage in z/VM 6.3, but is suggested for use only in special cases.

Summary

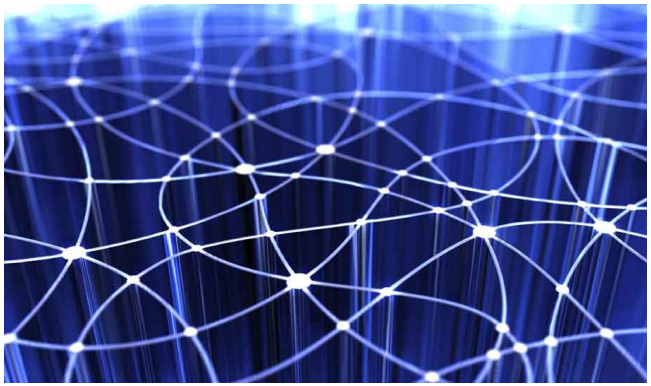

Leadership

z/VM continues to provide additional value to the platform as the strategic virtualization solution for System z.

Innovation

z/VM 6.2 introduced horizontal scalability and guest mobility through SSI Clustering and Live Guest Relocation with RAS in the forefront of the design. z/VM 6.3 continues the innovation with improved algorithms for memory and processor management.

Growth

z/VM 6.3 increases the vertical scalability and efficiency to complement the horizontal scaling introduced in z/VM 6.2, because we know our customers' systems continue to grow.

Thanks!!

Contact Information:

Brian W. Hugenbruch, CISSP
z/VM Security Design and Development
bwhugen at us dot ibm dot com
+1 607.429.3660
@Bwhugen