IBM z Systems

z/VM Platform Update

June 18, 2015 Version 4.11

Les Geer
z/VM Project Manager
lesgeer@us.ibm.com

© 2013, 2015 IBM Corporation

Trademarks

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

BladeCenter*	FICON*	OMEGAMON*	RACF*	System z9*	zSecure
DB2*	GDPS*	Performance Toolkit for VM	Storwize*	System z10*	z/VM*
DS6000*	HiperSockets	Power*	System Storage*	Tivoli*	z Systems*
DS8000*	HyperSwap	PowerVM	System x*	zEnterprise*	
ECKD	IBM z13*	PR/SM	System z*	z/OS*	

* 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.

© 2013, 2015 IBM Corporation

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.

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

Agenda

- Release Status and Information
- z/VM® Version 6 Release 3
  - 2014 Enhancements
  - 2015 Enhancements
- Futures and Statements of Direction

© 2013, 2015 IBM Corporation

Release Status and Information

© 2013, 2015 IBM Corporation

z/VM Release Status Summary

UPDATE

z/VM Level	GA	End of Service	End of Marktg.	Minimum Processor Level	Maximum Processor Level	Security Level
6.3	7/2013	12/2017[4]		IBM System z10®	-	EAL 4+ OSPP-LS
6.2	12/2011	12/2016[2]	7/2013	IBM System z10®	z13 [3]	-
6.1	10/2009	4/2013	12/2011	IBM System z10®	zEC12	EAL 4+ OSPP-LS
5.4	9/2008	12/2016[1]	3/2012	IBM eServer zSeries 800& 900	zEC12	-
5.3	6/2007	9/2010	9/2010	z800, z900	z196	EAL 4+ CAPP/LSPP

[1] Or later (Announced August 6, 2014)
[2] Extended from original date (Announced February 4, 2014)
[3] Announced January 14, 2015
[4] Announced February 3, 2015

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

© 2013, 2015 IBM Corporation

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
- Prices are set per value unit (based on number of engines), relative prices are illustrated below on right.
- The SSI feature includes LGR and it is priced in line with the RACF® feature
- See http://ibm.com/systems/z/swprice/zipla/zvm.html

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

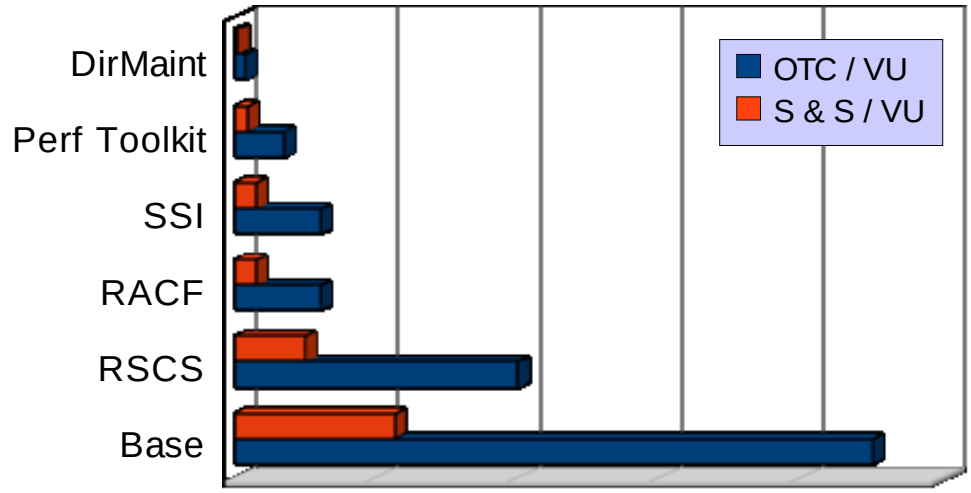

For illustrative purposes only. See IBM or Business Partner for details.

© 2013, 2015 IBM Corporation

z/VM Version 6 Release 3

© 2013, 2015 IBM Corporation

z/VM Version 6 Release 3

Making Room to Grow Your Business

July 23
Product Announcement

July 26
Product General Availability

April 30
System SSL Cryptographic Module FIPS 140-2 Certification

Feb 24
Announce additional support

July 16
z/VM support for zEDC Express and 10GbE RoCE Express features Available

June 9
CPU pooling Available

Aug 15
ILMT 9.01 support of CPU pooling Available

January 14
z13 and z/VM Enhancements Announcement

Feb 13
Base z13 & Crypto support Available

March 13
SMT and Scalability Support Available

June 26
Multi-VSwitch Link Aggregation Available

2013 | 2014 | 2015

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

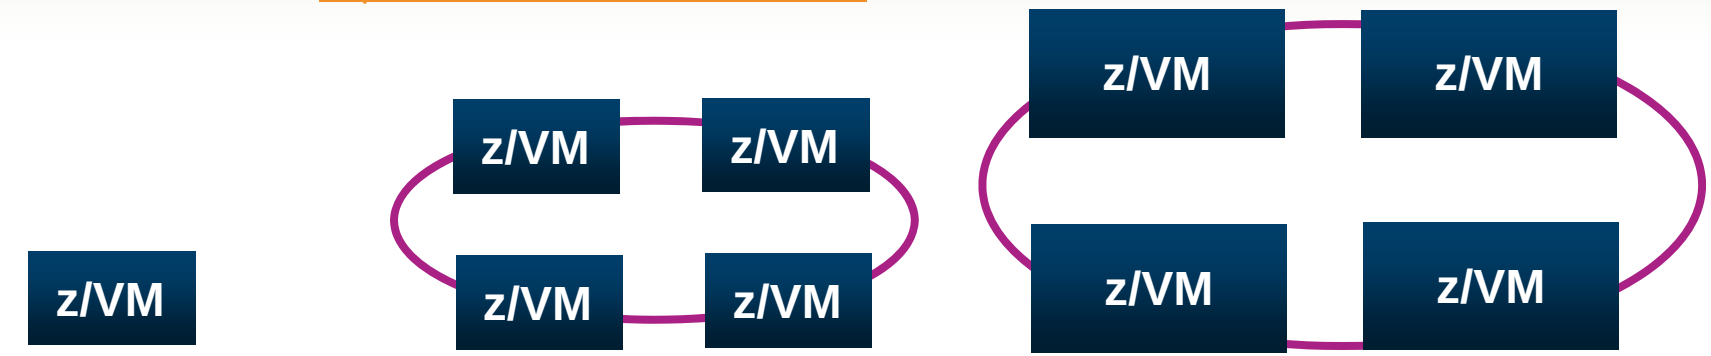

z/VM 6.1 | z/VM 6.2 | z/VM 6.3

© 2013, 2015 IBM Corporation

z/VM 6.3 Themes

- Improve Total Cost of Ownership
  - Expand z/VM systems constrained by memory or processor limitations with Large Memory Support, HiperDispatch, Scalability Enhancements, MSS for GDPS, etc.

- Enhance the Systems Management Experience
  - Enablement of OpenStack®, Upgrade in Place

- Continue Virtualization Leadership through Innovation
  - CPU Pooling, Simultaneous Multithreading (SMT), Efficiency Improvements, Virtual Switch Enhancements

© 2013, 2015 IBM Corporation

Large Memory Support

- Real memory limit raised from 256GB to **1 TB**
  - Proportionately increases total virtual memory based on tolerable overcommitment 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
- SET RESERVE enhanced to be more effective
- Convert expanded storage to real storage
  - Expanded Storage continues to be supported with a limit of **128 GB,** but is no longer recommended.

© 2013, 2015 IBM Corporation

Large Memory Support (continued)

- 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

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

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

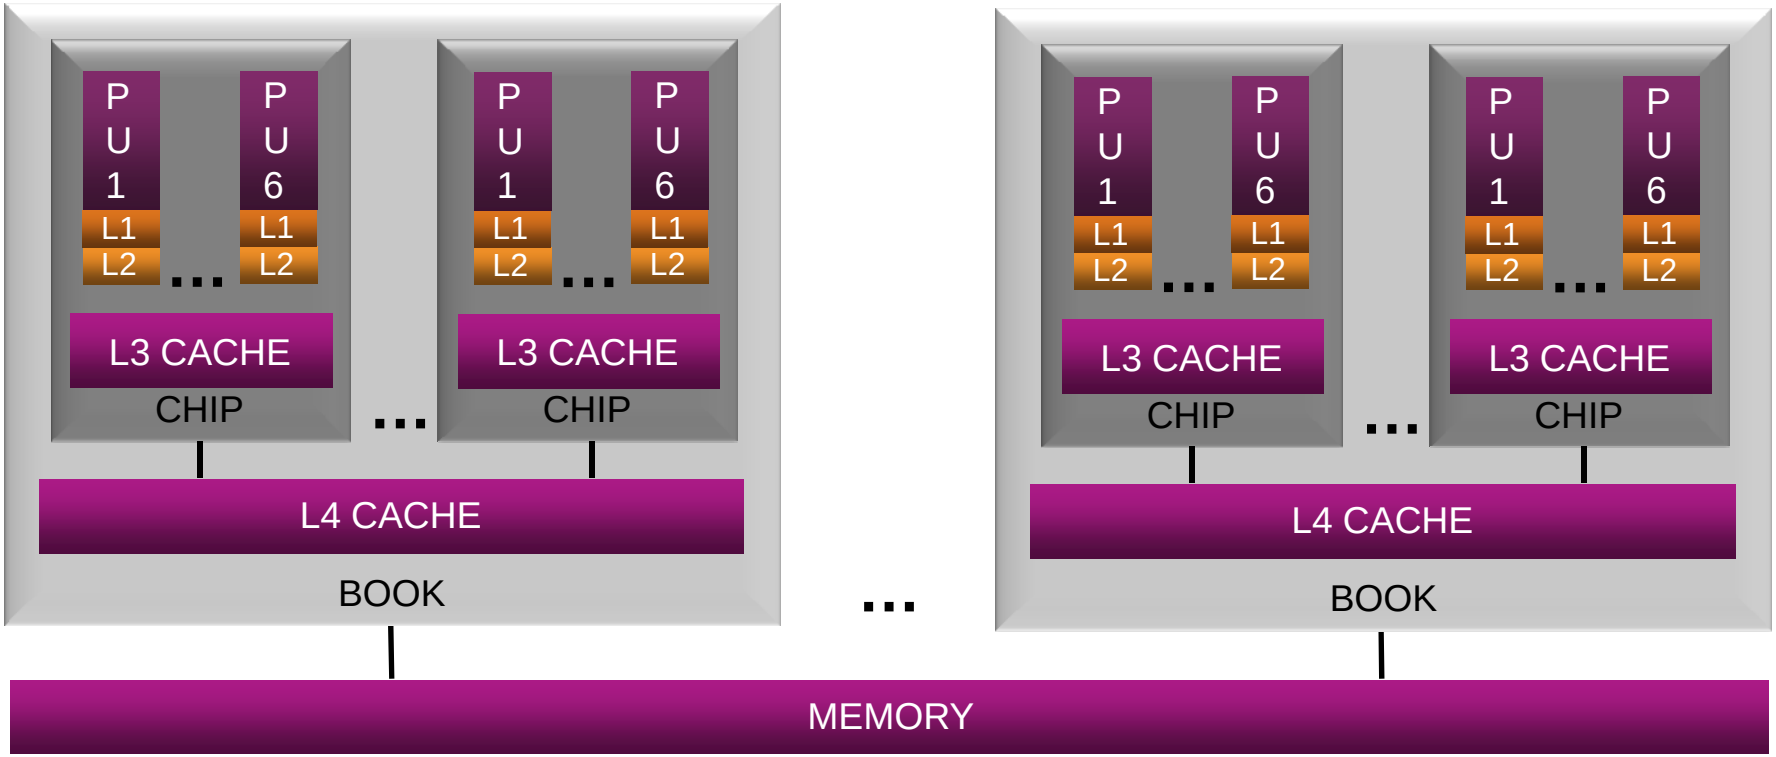

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

HiperDispatch – Vertical CPU Management

- "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)

© 2013, 2015 IBM Corporation

Horizontal vs. Vertical CPU Management

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

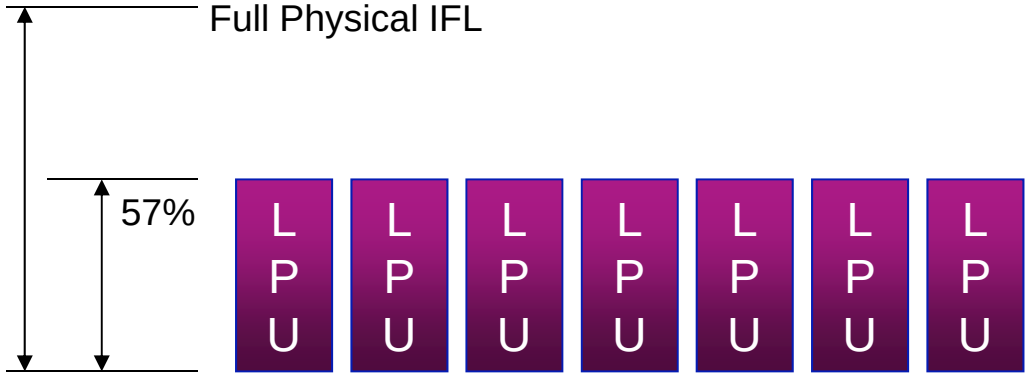

Example:
- 10 Physical IFL Processors
- 7 Logical IFL Processors
- Entitlement of 4 IFL Processors

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.

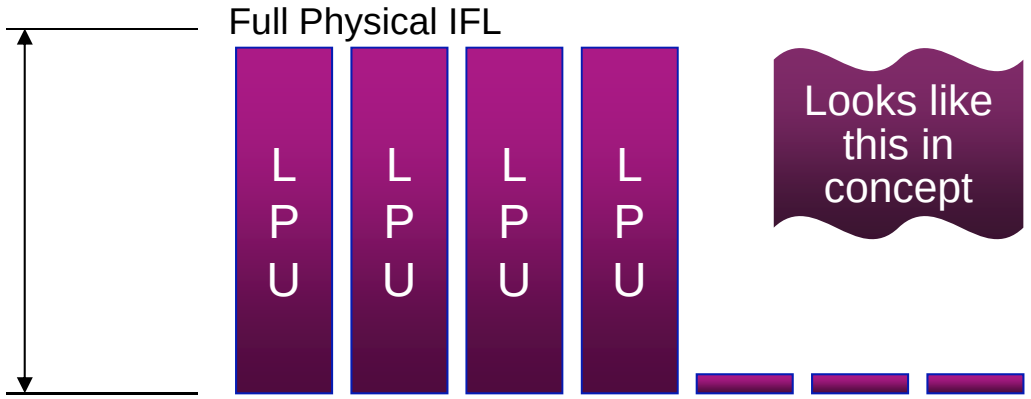

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

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
  - Client certificate validation
  - Includes support for:
    - Transport Layer Security (TLS) protocol, Version 1.2
    - SHA2 certificate support
    - TLS Protocol Selection
  - IPv6 support for SSL-enabled Telnet, FTP, and SMTP

© 2013, 2015 IBM Corporation

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
  - Preserves local mods
  - Can take advantage of guest relocation to avoid production server outages

- Upgrade Approach:
  - Install new release as temporary second level system
  - Move new level of z/VM to current system
  - IPL current system to start running with new release level

- For Single System Image (SSI) Cluster, start with single member of the cluster on new level to avoid outages

- Provides a backup to support backing out in extreme cases

- Support for local modifications

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

z/VM 6.3 Withdraws Cross System Extensions 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.

© 2013, 2015 IBM Corporation

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)

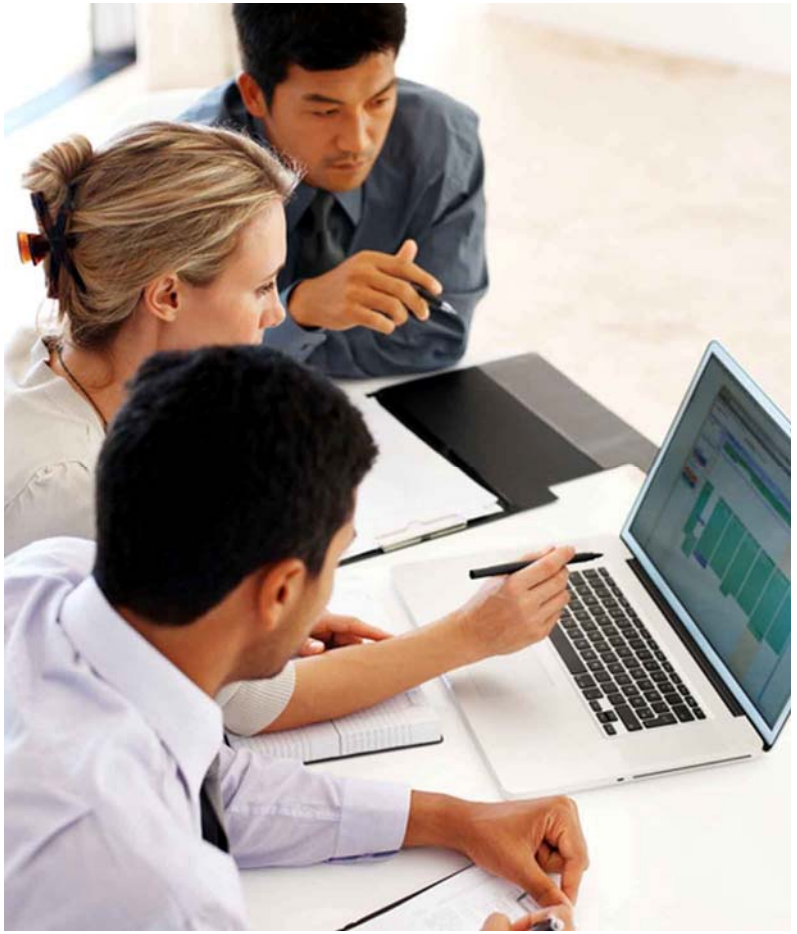

© 2013, 2015 IBM Corporation

z/VM Version 6
Security Certification Results

- Common Criteria (ISO/IEC 15408)
  - ***new*** z/VM V6.3 has been certified: BSI-DSZ-CC-0903
  - 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
  - ***new*** z/VM V6.3 System SSL is FIPS 140-2 Validated(TM)
    - http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2014.htm#2139
  - z/VM V6.1 System SSL is FIPS 140-2 Validated(TM)
    - http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2012.htm#1735
  - Enablement requirements for certificate database and servers

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

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

© 2013, 2015 IBM Corporation

IBM Wave for z/VM 1.2

NEW

- IBM Wave is a virtualization management product for z/VM® and Linux® virtual servers that uses visualization to dramatically automate and simplify administrative and management tasks
- IBM Wave for z/VM 1.2
  - Announcement May 11, 2015
  - General Availability June 19, 2015
- Strengths
  - Intelligent Visualization
  - Simplified Monitoring
  - Unified Management

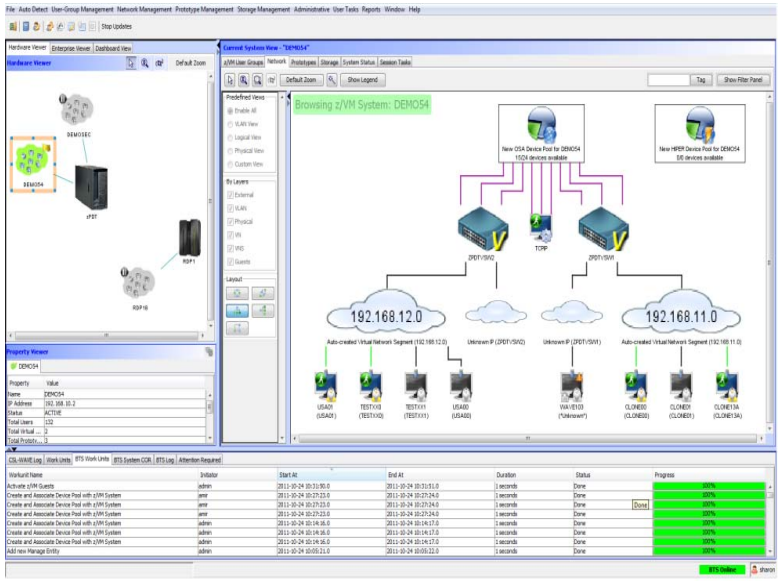

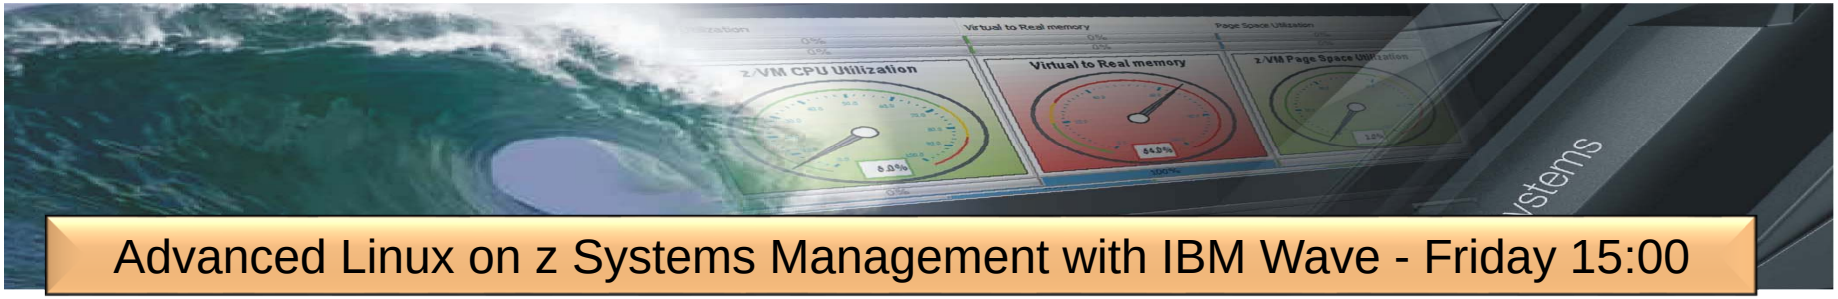

Advanced Linux on z Systems Management with IBM Wave - Friday 15:00

© 2013, 2015 IBM Corporation

A Different Better vs. A Standard Good

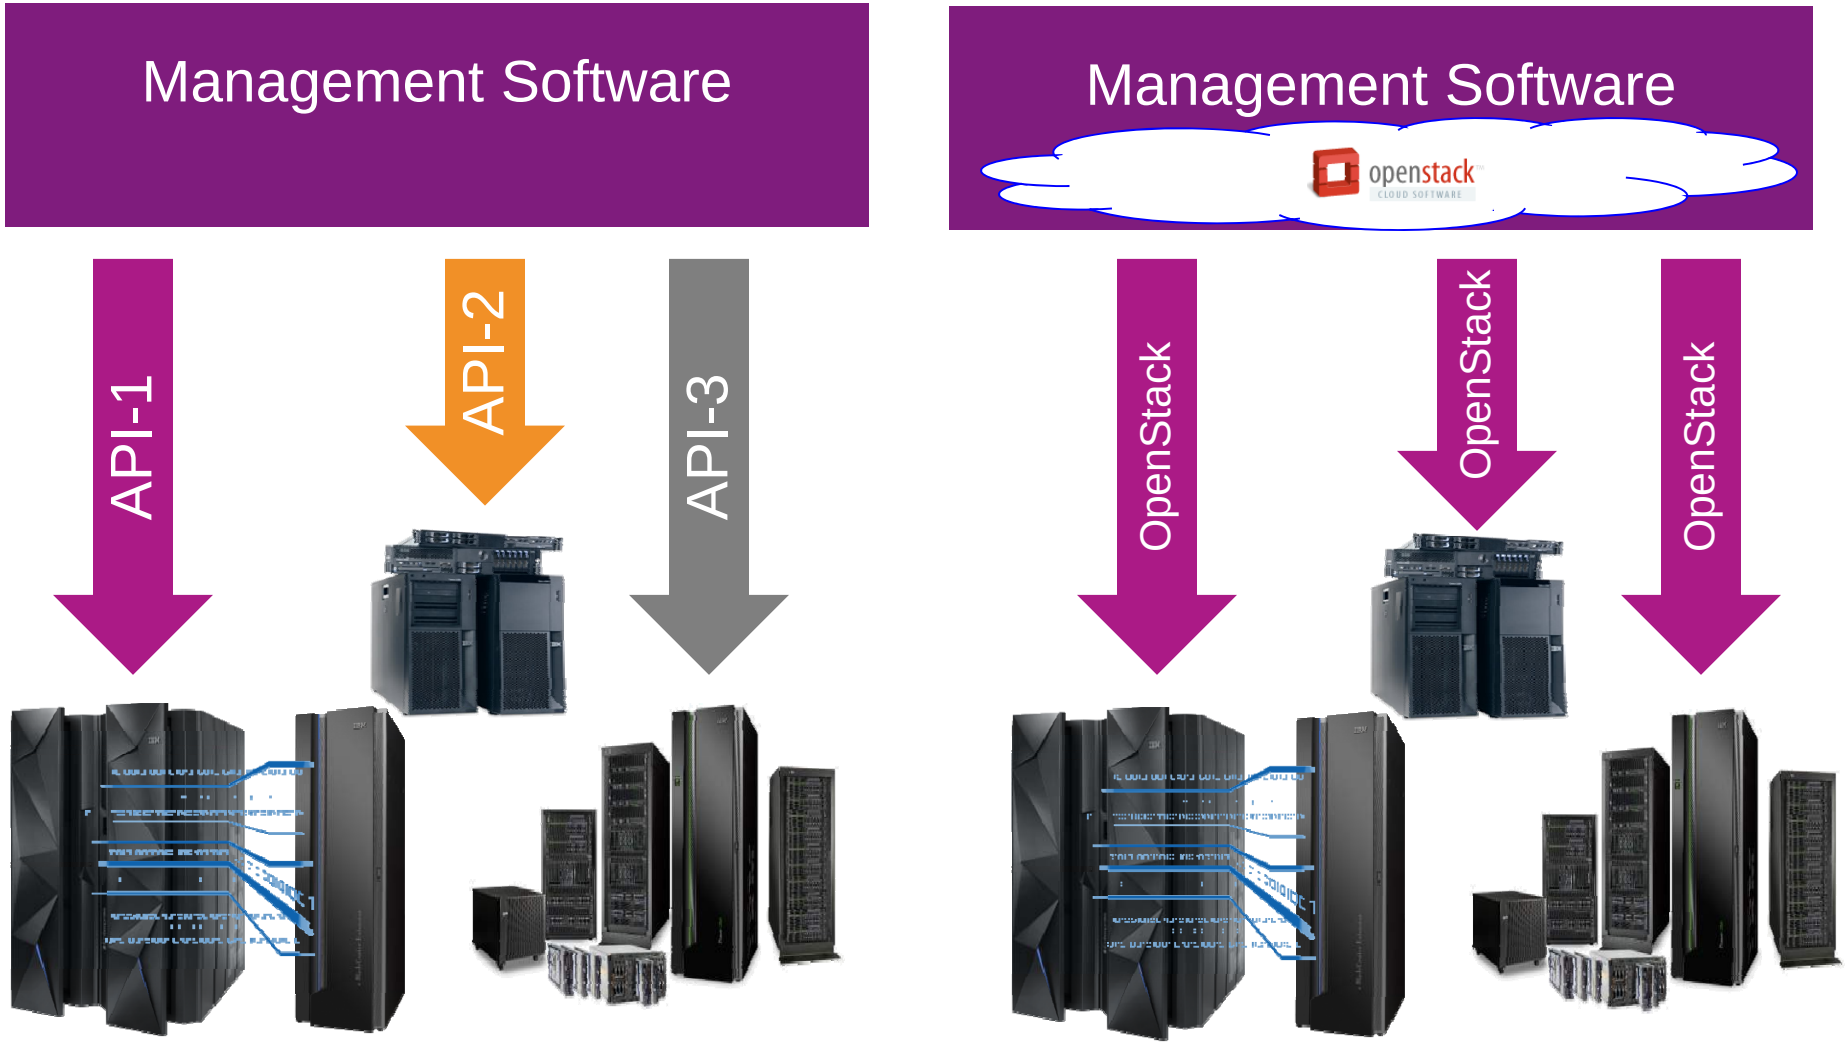

© 2013, 2015 IBM Corporation

The OpenStack Food Chain

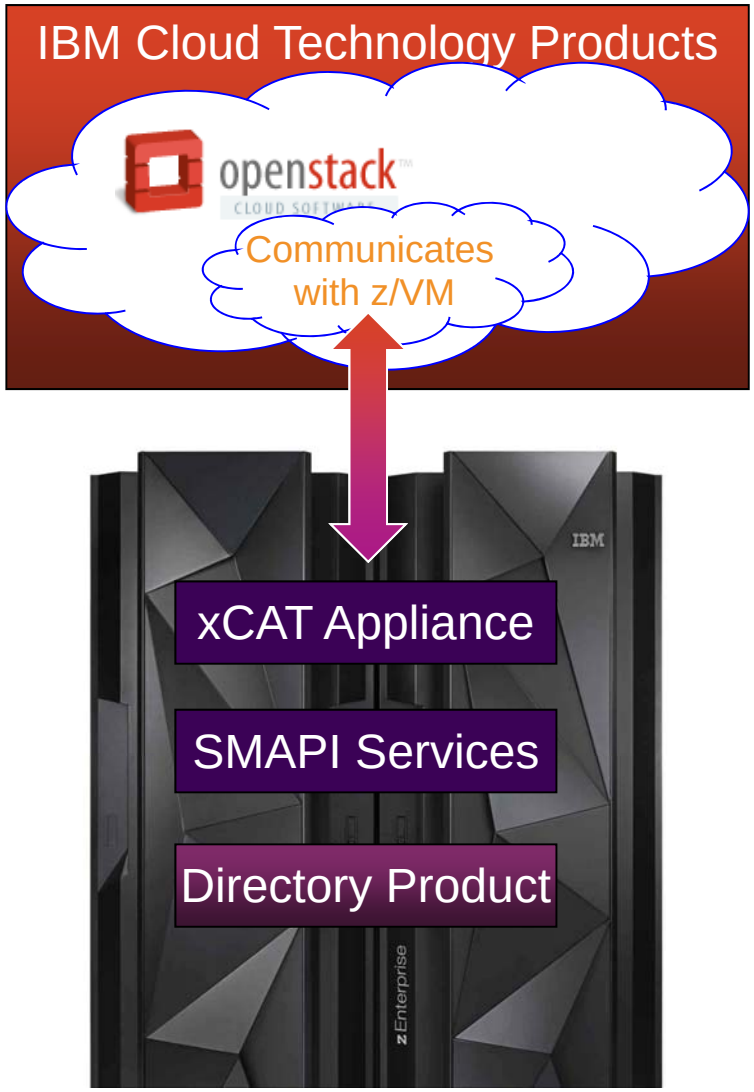

- **Top Half of the Solution:**
  - –An IBM Cloud Technology product or other vendor product will include the OpenStack support.
  - –Portions of that OpenStack support will know z/VM (i.e. code that connects and understands how to talk to z/VM).

- **Bottom Half of the Solution:**
  - –Rest APIs are used to communicate with the OpenStack code from the top half.
  - –The xCAT (Extreme Cloud Administration Toolkit) Appliance utilizes new and existing Systems Management APIs (SMAPI) to interact with the z/VM system
  - –SMAPI can interact with additional products or features (e.g. a directory manager).

Product with OpenStack Support

z/VM 6.3 Product

Additional Product or Feature

© 2013, 2015 IBM Corporation

OpenStack Related Products

- IBM Cloud Manager with OpenStack
  - Formerly known as IBM Smart Cloud Entry
  - V4.2 Announced February 24, 2015, available March 13, 2015
    - "Managed from" z Systems
  - V4.3 Announced May 11, 2015 availability of June 19, 2015
    - "Managed to" z Systems

- IBM Cloud Orchestrator
  - Formerly known as IBM Smart Cloud Orchestrator
  - "Managed to" z Systems (i.e. requires server off z)
  - V2.4 Currently Available

© 2013, 2015 IBM Corporation

IBM Infrastructure Suite for z/VM and Linux 1.1.0

- Announced and Available
  - Announced September 2, 2014
  - Available September 5, 2014
  - Announcement Letter ENUS214-350

- Includes following products:
  - IBM Tivoli® OMEGAMON® XE on z/VM and Linux V4.3
  - IBM Tivoli Storage Manager, part of IBM Spectrum Protect, Extended Edition V7.1
  - IBM Operations Manager for z/VM V1.5
  - IBM Backup and Restore Manager for z/VM V1.3
  - IBM Wave for z/VM V1.2

Customer Experiences Managing the z/VM & Linux Infrastructure – Thursday 13:45
Backup Strategies for z/VM and Linux on z Systems – Thursday 16:15
Monitoring Best Practices for z/VM and Linux – Friday 10:15

© 2013, 2015 IBM Corporation

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

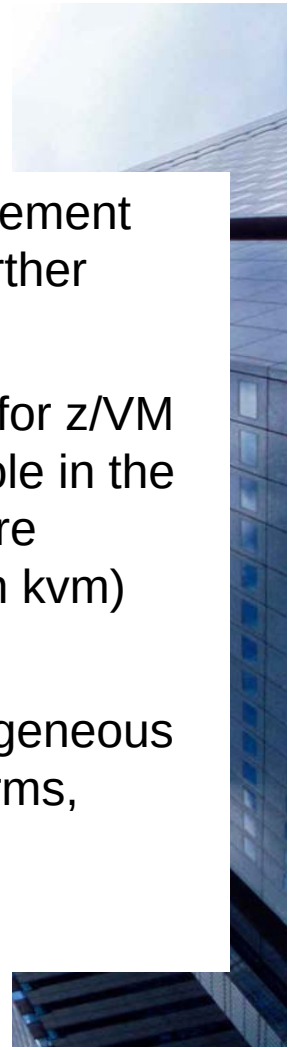

In light of IBM's cloud strategy and adoption of OpenStack, the management of z/VM environments in zManager is now stabilized and will not be further enhanced.

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 Cloud offerings.

© 2013, 2015 IBM Corporation

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 V1.3**
  - 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
    - IBM TS7700: needs firmware update available as code level 8.21.0.165 (EC: M13120 / PN: 2727271 & 2727272 (DVD1&2.))
    - Oracle StorageTek automated tape libraries (ATL) and virtual tape libraries (VTL) - via either the STK VM Host Support Component or the STK VM Client
    - EMC Virtual Tape Libraries (VTL), such as the EMC DLm.
- **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

© 2013, 2015 IBM Corporation

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
  - 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.

© 2013, 2015 IBM Corporation

February 24, 2014 Announcements

© 2013, 2015 IBM Corporation

Enhancing the Foundation for Virtualization

- Release for Announcement – zBX and zEnterprise System Enhancements
  - February 24, 2014
  - http://www.vm.ibm.com/zvm630/apars.html

- Software Enhancements
  - CPU Pooling
  - Environment Information Interface

- Hardware Support
  - 10GbE RoCE Express Feature
  - zEDC Express Feature

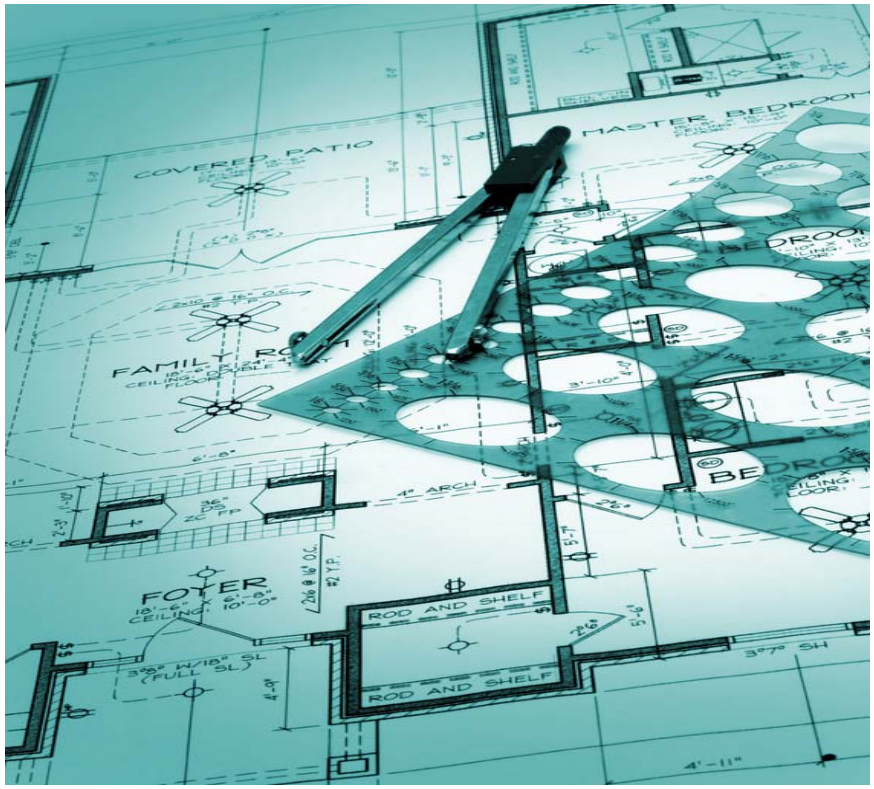

© 2013, 2015 IBM Corporation

CPU Pooling

- Fine-grained 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 – Available
    - Part of RSU 1501

z/VM CPU Pooling and ILMT Support – Friday 10:15

© 2013, 2015 IBM Corporation

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 state instruction Store Hypervisor Information (STHYI)
- Includes support for CPU Pooling enhancement
- Foundation for future software licensing tools
  - IBM License Metric Tool 9.0.1updated August 2014- http://ibm.biz/cpupoolilmt
  - Greater flexibility for IBM Passport Advantage products
- Support details:
  - z/VM 6.3 with APAR VM65419 – Available
    - Part of RSU 1501

© 2013, 2015 IBM Corporation

CPU Pooling Example

- 4 WAS production guests
  - Requires 4-engine WAS entitlement
- Create a 1-IFL pool
- Put the 2 DB2 production guests in pool
  - Requires 1-engine DB2 entitlement (avoiding the need for 2-engine DB2 entitlement without CPU pooling)

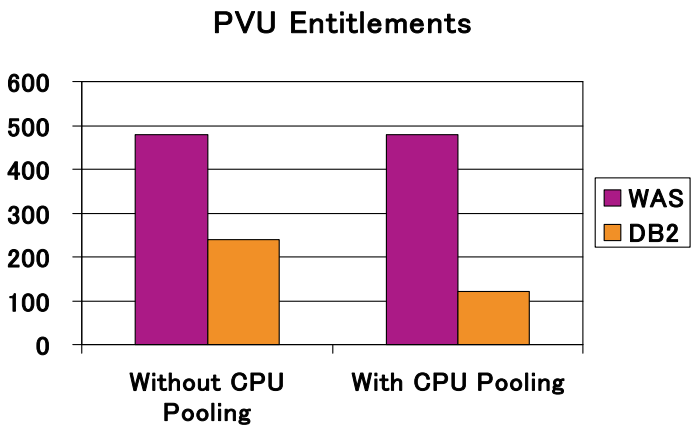

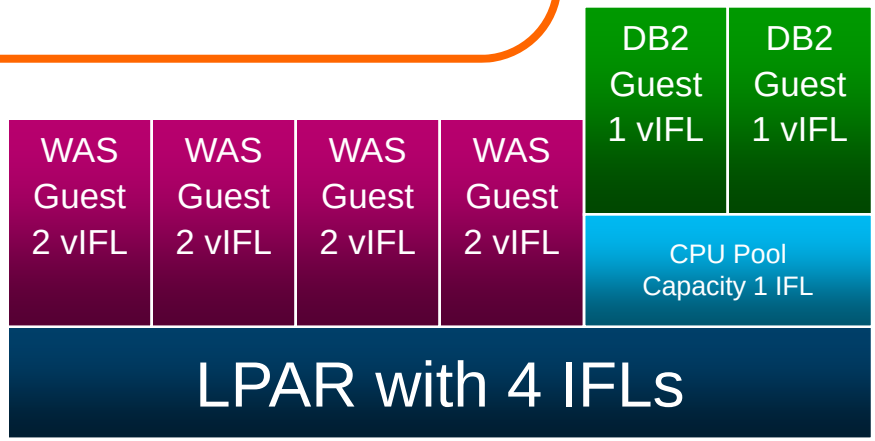

- Allows new workloads and additional workload consolidation to be more cost effective

Note: All PVU Entitlement examples based on zEC12 (120 PVU per IFL) – will look proportionally the same on zBC12 (100 PVU per IFL)

© 2013, 2015 IBM Corporation

10GbE RoCE Express Feature

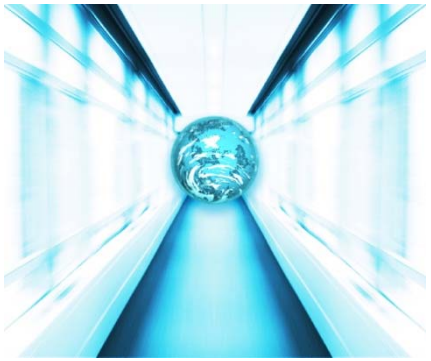

- Support for RDMA over Converged Ethernet for guests
- Based on new hypervisor PCIe support
- Designed to support z/OS's Shared Memory Communications-Remote Direct Memory Access (SMC-R) in z/OS V2.1
- Support details:
  - IBM z13, zEC12, or zBC12 with appropriate updates – see support buckets
  - z/VM 6.3 with APAR VM65417 – Available – RSU 1501
    - System Config option – disabled by default.
    - You need to have required millicode fixes applied prior to enabling in system config
  - z/OS 1.12, z/OS 1.13, z/OS 2.1 with APAR OA43256
  - Fulfills 2013 Statement of Direction

© 2013, 2015 IBM Corporation

zEDC Express Feature

- Guest support for zEDC Express Feature

- High performance, low CPU consumption compression
- Possible disk utilization reduction
- Support details:
  - IBM zEC12 or zBC12 with appropriate updates – see support buckets
  - z/VM 6.3 with APAR VM65417 – Available – RSU 1501
    - System Config option – disabled by default.
    - You need to have required millicode fixes applied prior to enabling in system config
  - 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

© 2013, 2015 IBM Corporation

January 14, 2015 Announcements

© 2013, 2015 IBM Corporation

Expanding the Horizon of Virtualization

- Release for Announcement – The IBM z13™
  - January 14, 2015
  - Announcement Link

- z/VM Compatibility Support
  - PTFs available February 13, 2015
  - Also includes Crypto enhanced domain support
  - z/VM 6.2 and z/VM 6.3
  - No z/VM 5.4 support

- Enhancements and Exploitation Support - only on z/VM 6.3
  - IBM z13 Simultaneous Multithreading
  - Increased Processor Scalability
  - Multi-VSwitch Link Aggregation Support (Link Aggregation with Shared OSAs)

- Performance Report at http://www.vm.ibm.com/perf/reports/zvm/html/

© 2013, 2015 IBM Corporation

z/VM Support for Crypto Express5S

- z/VM supports the z13 and Crypto Express5S feature
  - z/VM 6.2 and z/VM 6.3 only
  - APAR VM65577

- Expanded domain selection for dedicated domains
  - z/VM supports architected limits for CryptoExpress domains
  - CryptoExpress5S supports 85 domains per feature, with a maximum of 16 features

- Selection of APVIRT domains in System Configuration
  - Avoid collisions when reassigning domains in user directory
  - Minimize need for LPAR restart

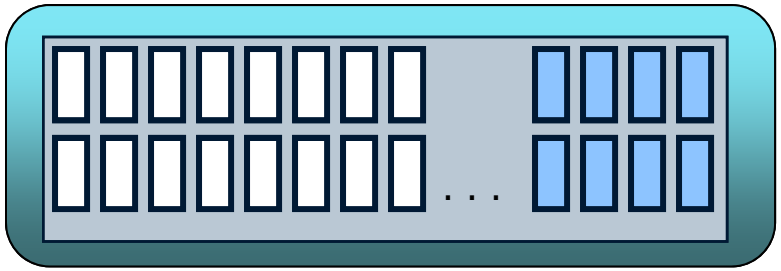

CEX5C 0

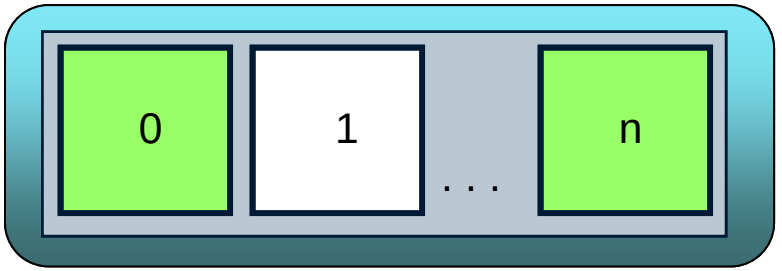

CEX5A 1

Crypto Session of Interest

Keys to the Virtual Kingdom – Friday12:30

© 2013, 2015 IBM Corporation

Simultaneous Multithreading (SMT)

- Objective is to improve capacity, not performance.
- Allows z/VM to dispatch work on up to two threads of a z13 IFL
- VM65586 for z/VM 6.3 **only**
  - PTFs available March 13, 2015
- At least z13 millicode bundle 11
- Transparent to virtual machine
  - Guest does not need to be SMT aware
  - SMT is not virtualized to the guest
- z13 SMT support limited to IFLs and zIIPs
  - z/VM support is only for IFLs
- SMT is disabled by default
  - Requires a System Configuration setting and re-IPL
  - When enabled, applies to the entire system
- Potential to increase the overall capacity of the system
  - Workload dependent

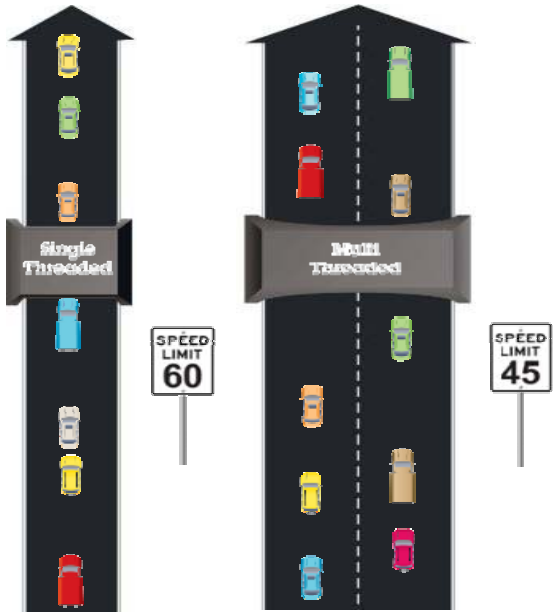

Which approach is designed for the higher volume of traffic? Which road is faster?

**Illustrative numbers only*

© 2013, 2015 IBM Corporation

SMT Usage

- Physical IFL Cores with SMT allow up to two threads to be used. You purchase these.
- Logical IFL Cores are presented to z/VM as in the past. You define these in the logical partition definition on HMC.
- z/VM creates a CPU or logical processor associated with each thread for it to use. Reflected in commands like QUERY PROCESSORS.
- The virtual CPUs of guests can then be dispatched on different threads intelligently, based on topology information.

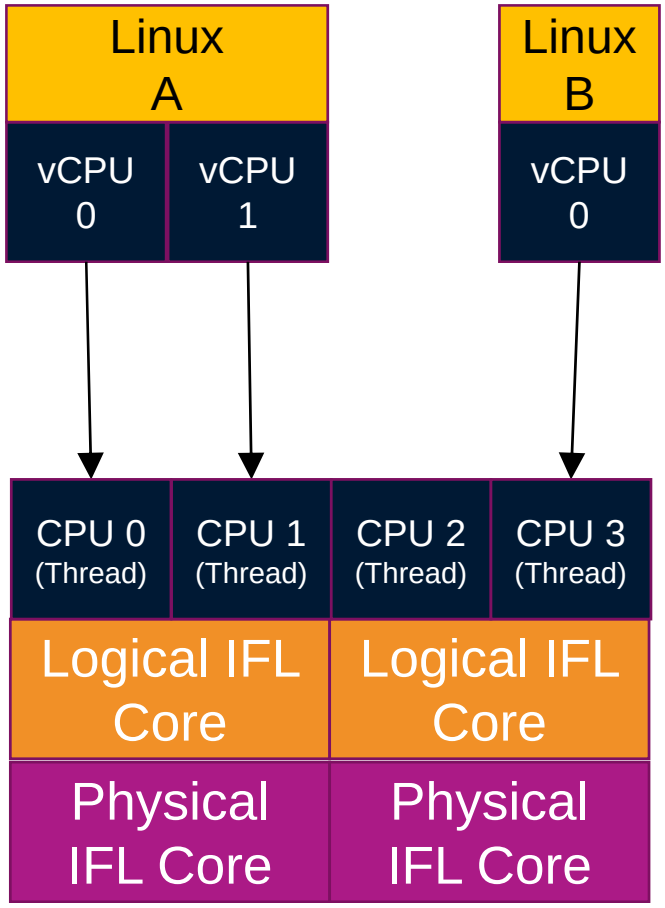

© 2013, 2015 IBM Corporation

SMT Sessions of Interest

How do you spell SMT on z Systems? – Thursday 15:00

A look at z/VM Performance with SMT – Saturday 9:00

Highly encourage attending first session if you're going to the SMT Performance one.

© 2013, 2015 IBM Corporation

SMT Usage – Mixed Engine Environment

- In a mixed-engine environment, general purpose processors cannot do threading, but a second CPU address is consumed (CPU 1 in example).
- Virtual IFL CPUs would get dispatched to the logical IFLs and virtual CP CPUs would get dispatched on the logical CPs

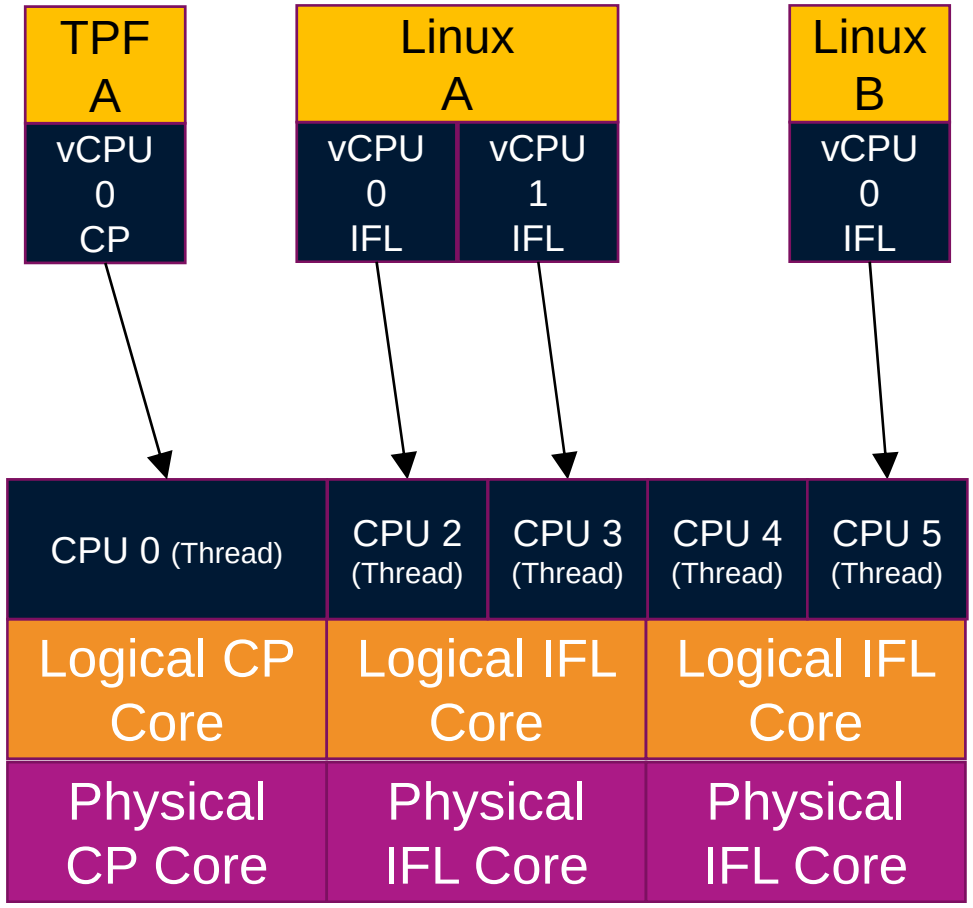

© 2013, 2015 IBM Corporation

Increased CPU Scalability

- Various improvements to allow z/VM systems to be larger in terms of processors and more efficient, improving the n-way curve

- APARs VM65586 & VM65696 for z/VM 6.3 **only**
  - PTFs available March 13, 2015

- For z13
  - With SMT disabled, increases logical processors supported from 32 to 64
  - With SMT enabled, the limit is 32 IFL cores (64 threads)

- For processors prior to z13
  - Limit remains at 32 cores
  - May still benefit from improved n-way curves

© 2013, 2015 IBM Corporation

Areas Improved with Scalability Enhancements

- z/VM Scheduler Lock
  - Management of internal stacked work
  - Guests going into a wait state
- Locking for Memory Management
  - Most benefit during system initialization and when very constrained with memory
- Serialization and processing of VDisk I/Os
- Batching and processor-local queues for VSWITCH buffers

© 2013, 2015 IBM Corporation

Multi-VSwitch Link Aggregation

- Link Aggregation is ability to combine or aggregate up to 8 OSAs to increase the bandwidth available to a VSwitch
- This enhancement makes it possible to do Link Aggregation with VSwitches with shared OSAs rather than previous the requirement for dedicated OSAs
- Allows a port group of OSA-Express features to span VSwitches within a single or multiple z/VM systems.
  - Cannot be shared with non-z/VM logical partitions
- APARs VM65583 (CP), PI21053 (TCP/IP), VM65528 (Performance Toolkit), and VM65670 (SMAPI) for z/VM 6.3 **only**
  - PTFs planned to be available June 26, 2015
- Available only on the z13
  - Requires OSA enhancements introduced with the z13
- Allows better consolidation and availability while improving TCO

© 2013, 2015 IBM Corporation

Multi-VSwitch LAG Configuration

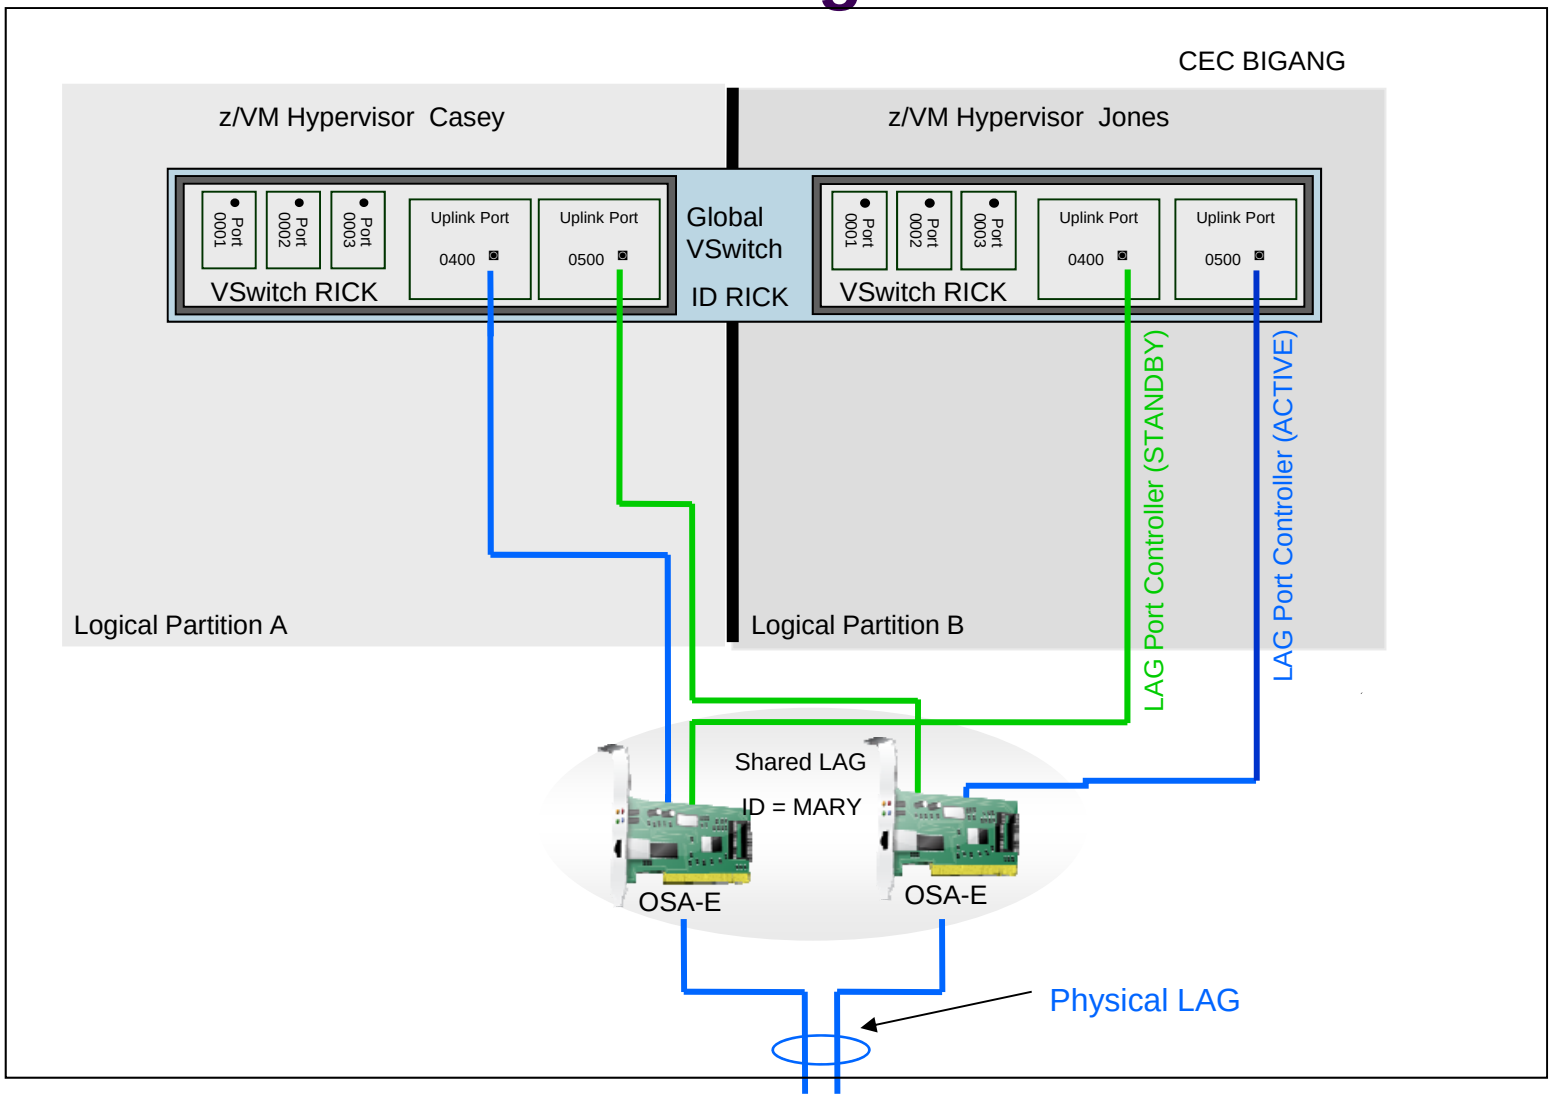

Hardware Support

© 2013, 2015 IBM Corporation

Support for IBM z13

- **Updates for z/VM 6.2 and 6.3**
  - http://www.vm.ibm.com/service/vmreqz13.html
  - Many components affected
  - Note: Directory space requirements increased slightly.
- No z/VM 5.4 Support as previously announced
- No z/VM 6.1 Support even if you have extended support contract.
- **PSP Bucket**
  - Upgrade **2964DEVICE**
  - Subset **2964/ZVM**
- **If running Linux, please also check for required updates prior to migration.** **WARNING!**

© 2013, 2015 IBM Corporation

Tested Linux Platforms

http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

Distribution	z13	zEnterprise - zBC12 and zEC12	zEnterprise - z114 and z196	System z10 and System z9
RHEL 7	✓ (1,3)	✓ (4)	✓ (4)	✗
RHEL 6	✓ (1,3)	✓ (5)	✓	✓
RHEL 5	✓ (1,3)	✓ (6)	✓	✓
RHEL 4 (*)	✗	✗	✓ (9)	✓
SLES 12	✓ (2,3)	✓	✓	✗
SLES 11	✓ (2,3)	✓ (7)	✓	✓
SLES 10 (*)	✗	✓ (8)	✓	✓
SLES 9 (*)	✗	✗	✓ (10)	✓

© 2013, 2015 IBM Corporation

Support for IBM zEnterprise EC12

- **Updates for z/VM 6.2 and 5.4 (In base of z/VM 6.3)**
  - 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
- http://www.vm.ibm.com/service/vmreqzeh.html

© 2013, 2015 IBM Corporation

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**
- **http://www.vm.ibm.com/service/vmreqzeh.html**

© 2013, 2015 IBM Corporation

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

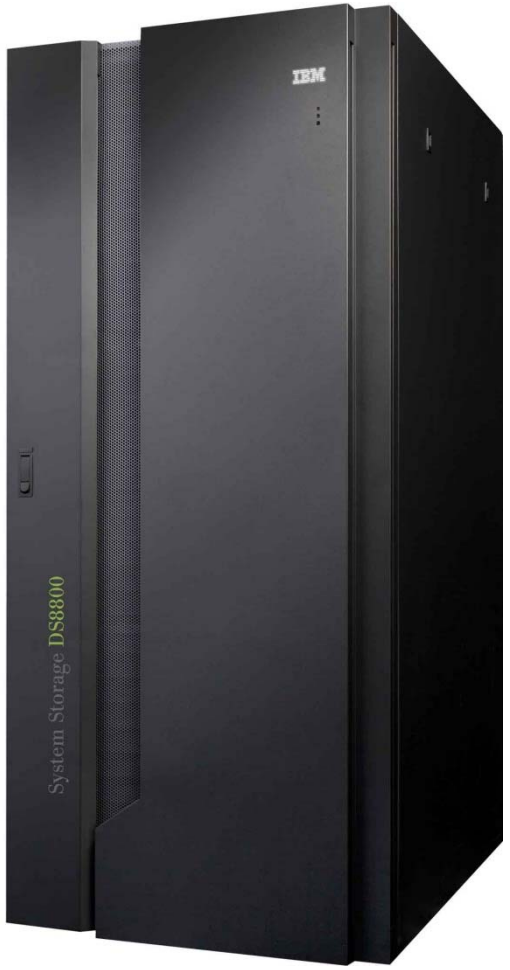

© 2013, 2015 IBM Corporation

Multiple Target Peer-to-Peer Remote Copy Support

- Multiple Target Peer-to-Peer Remote copy (MT-PPRC) Support
  - Allows two PPRC relationships on a single primary volume.

UPDATE

- IBM DS8870 systems
  - Microcode level 7.4 required
  - Announced October 6, 2014
  - Available December 5, 2014

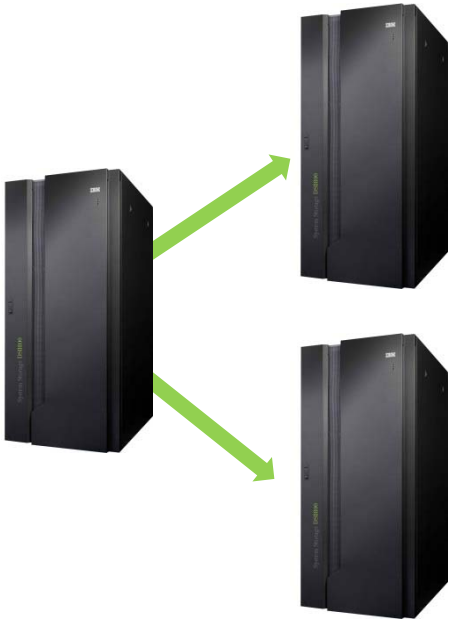

- Device Support Facilities (ICKDSF)
  - APAR PM99490

- z/VM Support
  - APAR VM65544
    - Closed November 13, 2014
  - Primary in subchannel set 0
  - Does not support a multiple target secondary in the alternate subchannel set
  - APAR **must** be applied prior to storage server upgrade to microcode level 7.4
    - APAR is required even if not exploiting new function
    - See Red Alert http://www.vm.ibm.com/service/redalert/#VM65544

WARNING!

© 2013, 2015 IBM Corporation

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

© 2013, 2015 IBM Corporation

IBM Enterprise Cloud System

Trusted Cloud. Simply Delivered.

Build It

Open Linux Environment
- *Red Hat/SUSE*
- *3000+ Applications*

openstack

Fully Automated Cloud Orchestration & Monitoring

Hypervisor and Virtualization Management

Utility Pricing and MSP Flexible Financing

Trusted, 24/7 IBM Support

CRN TECH INNOVATOR AWARDS 2013

Award Winning Hardware Design

LINUX JOURNAL 2013 Readers' Choice

- **Integrated**
- **Delivered in 30-45 Days**
- **Production Ready in Hours**

- **99.99%+ Availability**
- **EAL4 Server Security**
- **Available June 20, 2014**

© 2013, 2015 IBM Corporation

Enterprise Cloud System- Offering Components

- **Server:**
  - zEC12, zBC12, or z13
- **Storage:**
  - IBM DS8870 **or** Storwize® V7000
- **Software:**
  - z/VM® 6.3 with following features:
    - Directory Maintenance (DirMaint™) Feature
    - Resource Access Control Facility (RACF®)
    - Performance Toolkit for VM™ Feature
    - Single System Image (SSI) Feature – (Requires ECKD DASD)
  - IBM Wave for z/VM
  - OMEGAMON® XE on z/VM and Linux
  - Tivoli Storage Manager
  - IBM Cloud Orchestrator
  - Operations Manager for z/VM
  - Backup and Restore Manager for z/VM

© 2013, 2015 IBM Corporation

Statements of Direction
July 23, 2013
January 14, 2015

- Subset of IBM Statements of General Direction that are most important to the z/VM environment. See announcement materials for additional statements.
- Subject to change or withdrawal without notice, representing IBM goals and objectives only.

© 2013, 2015 IBM Corporation

Completed Statements of Direction

July 23, 2013

- The following Statements of Direction from z/VM 6.3 have been **met**:
  - Security Evaluation of z/VM 6.3
  - FIPS 140-2 Validation of z/VM 6.3
  - Support of 10 GbE RoCE Express Feature
  - Support of zEDC Express Feature
  - Stabilization of z/VM 5.4 Support
- Requires support from hardware and/or guests operating systems as appropriate
- Refer to www.vm.ibm.com or www.vm.ibm.com/security for more information

Satisfied

© 2013, 2015 IBM Corporation

Withdrawal of Support for Expanded Storage

July 23, 2013

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.

© 2013, 2015 IBM Corporation

KVM offering for IBM z Systems

January 14, 2015

In addition to the continued investment in z/VM, IBM intends to support a Kernel-based Virtual Machine (KVM) offering for z Systems that will host Linux on z Systems guest virtual machines.

The KVM offering will be software that can be installed on z Systems processors like an operating system and can co-exist with z/VM virtualization environments, z/OS, Linux on z Systems, z/VSE, and z/TPF.

The KVM offering will be optimized for z Systems architecture and will provide standard Linux and KVM interfaces for operational control of the environment, as well as providing the required technical enablement for OpenStack for virtualization management, allowing enterprises to easily integrate Linux servers into their existing infrastructure and cloud offerings.

- An additional option for virtualization on z Systems.
- The IBM commitment to z/VM remains steadfast.

© 2013, 2015 IBM Corporation

GDPS/PPRC Multiplatform Resiliency Capability

January 14, 2015

In the first half of 2015, IBM intends to deliver a **GDPS/Peer to Peer Remote Copy (GDPS/PPRC) multiplatform resiliency capability** for customers who do not run the z/OS operating system in their environment. This solution is intended to provide IBM z Systems customers who run z/VM and their associated guests, for instance, Linux on z Systems, with similar high availability and disaster recovery benefits to those who run on z/OS. This solution will be applicable for any IBM z Systems announced after and including the zBC12 and zEC12.

- Lower the skill expense of running a GDPS environment, particularly for those customers with little, or no, z/OS background.

Satisfied

© 2013, 2015 IBM Corporation

Enhanced RACF Password Encryption Algorithm for z/VM

January 14, 2015

Enhanced RACF® password encryption algorithm for z/VM: In a future deliverable an enhanced RACF/VM password encryption algorithm is planned. This support will be designed to provide improved cryptographic strength using AES-based encryption in RACF/VM password algorithm processing. This planned design is intended to provide better protection for encrypted RACF password data in the event that a copy of RACF database becomes inadvertently accessible.

- z/OS support for this currently exists.
- Lack of this support in z/VM complicates sharing RACF databases with z/OS where the support is used.

© 2013, 2015 IBM Corporation

z/VM Support for Single Instruction Multiple Data (SIMD)

January 14, 2015

In a future deliverable IBM intends to deliver support to enable z/VM guests to exploit the Vector Facility for z/Architecture (SIMD).

- The Single Instruction Multiple Data (SIMD) was introduced as part of the z13, allowing use of the new Vector Facility.
- The initial z/VM support for z13 does not contain the virtualization of SIMD, which would allow guests to exploit it and gain potential performance benefits.

© 2013, 2015 IBM Corporation

Removal of Support for Expanded Storage

January 14, 2015

z/VM V6.3 is the last z/VM release that will support Expanded Storage (XSTORE) for either host or guest usage. The IBM z13 server family will be the last z Systems server to support Expanded Storage (XSTORE).

- The previous SoD spoke of removal of paging to expanded storage, but there is more.
- All z/VM support for expanded storage will be removed in future release
  - Attaching to guests
  - Minidisk Cache
  - Paging
  - etc.
- This SoD also goes on to speak to hardware support being removed as well, after the z13 server family.

© 2013, 2015 IBM Corporation

Removal of ESA/390 Architecture Mode

January 14, 2015

The IBM z13 will be the last z Systems server to support running an operating system in ESA/390 architecture mode; all future systems will only support operating systems running in z/Architecture mode. This applies to operating systems running native on PR/SM as well as operating systems running as second level guests. IBM operating systems that run in ESA/390 mode are either no longer in service or only currently available with extended service contracts, and they will not be usable on systems beyond IBM z13. However, all 24-bit and 31-bit problem-state application programs originally written to run on the ESA/390 architecture will be unaffected by this change.

- While a hardware statement, there are potentially changes required for z/VM.
- Note implication of older operating systems.

© 2013, 2015 IBM Corporation

Stabilization of z/VM 6.2 Support

January 14, 2015

The IBM z13 server family is planned to be the last z Systems server supported by z/VM V6.2 and the last z systems server that will be supported where z/VM V6.2 is running as a guest (second level). This is in conjunction with the statement of direction that the IBM z13 server family will be the last to support ESA/390 architecture mode, which z/VM V6.2 requires. z/VM V6.2 will continue to be supported until December 31, 2016, as announced in Withdrawal Announcement 914-012, dated February 04, 2014.

- While z/VM 6.2 will be supported until the end of 2016, there will **not** be support for the next server family.
- Similar to the statement of direction with z/VM 5.4 not supported on z13.

© 2013, 2015 IBM Corporation

Product Delivery of z/VM on DVD/Electronic Only

January 14, 2015

Product Delivery of z/VM on DVD/Electronic only: z/VM V6.3 will be the last release of z/VM that will be available on tape. Subsequent releases will be available on DVD or electronically.

- No more tapes for z/VM product delivery for future z/VM releases.
- Allows testing resources to be spent else where.

z/VM Installation, Service, and You – A Developer Dialogue – Thursday 16:15

© 2013, 2015 IBM Corporation

CMS Pipelines

- IBM is investigating ways to bring the latest version of CMS Pipelines into the z/VM Product.
- Various approaches are being considered to balance resources available for this effort with customer value.
- If you have thoughts or ideas, or just want to share your view on this effort, please see Rob Van Der Heij this week, or contact him via email (robvdheij@nl.ibm.com).

IBM plans and directions are subject to change without notice.

© 2013, 2015 IBM Corporation

Summary

© 2013, 2015 IBM Corporation

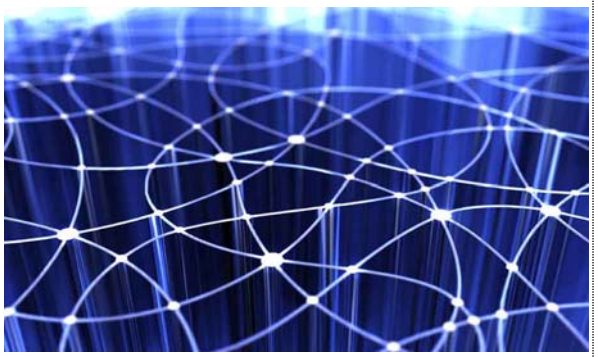

Leadership

z/VM continues to provide additional value to the platform as the strategic virtualization solution for z Systems. Virtual Switch technology in z/VM is industry leading.

Innovation

z/VM 6.3 added HiperDispatch, allowing greater efficiencies to be realized. Adding SMT with topology awareness raises the bar again.

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. This year we continue to extend the limits with processor scalability improvements.

© 2013, 2015 IBM Corporation

Backup Slides

© 2013, 2015 IBM Corporation

Security Evaluation of z/VM 6.3

July 23, 2013

IBM intends to evaluate z/VM V6.3 with the RACF Security Server feature, including labeled security, for conformance to the Operating System Protection Profile (OSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4 (EAL4+).

- Evaluation is with inclusion of **RACF Security Server** and **Single System Image** priced features enabled.
  - Evaluated configuration supports clusters of *1-n* z/VM systems.
  - No claims made about standalone systems, or systems without RACF for VM.

- See http://www.vm.ibm.com/security/ for current z/VM Security information.

Satisfied

© 2013, 2015 IBM Corporation

FIPS Certification of z/VM 6.3

July 23, 2013

IBM intends to pursue an evaluation of the Federal Information Processing Standard (FIPS) 140-2 using National Institute of Standards and Technology's (NIST) Cryptographic Module Validation Program (CMVP) for the System SSL implementation utilized by z/VM V6.3.

- Federal Information Protection Standard (FIPS) 140-2
  - Target z/VM 6.3 System SSL is FIPS 140-2 Validated*
  - Enablement requirements for certificate database and servers
    - http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2012.htm#1735

- See http://www.vm.ibm.com/security/ for current z/VM Security information.

Satisfied

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

© 2013, 2015 IBM Corporation

Support of the 10GbE RoCE Express Feature

July 23, 2013

In a future z/VM deliverable IBM plans to offer support for guest exploitation of the 10GbE RoCE Express feature (#0411) on the IBM zEnterprise EC12 and IBM zEnterprise BC12 systems. This is to allow guests to utilize Remote Direct Memory Access over Converged Ethernet (RoCE) for optimized networking.

- RoCE is high bandwidth, low latency link layer protocol
- Guest support for devices dedicated to z/VM guests that support RoCE
- Requires 10GbE RoCE Express feature on either the IBM zEC12 or IBM zBC12

Satisfied

© 2013, 2015 IBM Corporation

Support of the zEDC Express Feature

July 23, 2013

In a future z/VM deliverable IBM plans to offer z/VM support for guest exploitation of the IBM zEnterprise Data Compression (zEDC) Express feature (#0420) on the IBM zEnterprise EC12 and IBM zEnterprise BC12 systems.

- New data compression hardware feature to improve ability to do compression by offloading to zEDC
- Support is planned for guest usage
- Requires zEDC Express feature on either the IBM zEC12 or IBM zBC12

Satisfied

© 2013, 2015 IBM Corporation

Stabilization of z/VM 5.4 Support

July 23, 2013

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, 2016, 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 for z/VM 5.4, 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.

Satisfied

© 2013, 2015 IBM Corporation