

# z/VM 6.4 News

Enhancements z/VM 6.3 and Preview z/VM 6.4

*Rob van der Heij*

*z/VM Development – IBM Endicott*



10<sup>th</sup> GSE / IBM Technical University for z/VSE,  
z/VM, KVM and Linux on IBM z Systems

Oct 24-26, 2016. Leipzig

# 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](http://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](http://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

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



# Agenda

- Release Status and Information
- z/VM<sup>®</sup> Version 6 Release 3
  - 2014 Enhancements
  - 2015 Enhancements
  - 2016 Enhancements
- z/VM<sup>®</sup> Version 6 Release 4 Preview



# Release Status and Information

# z/VM Release Status Summary



z/VM Level	GA	End of Service	End of Marketing	Minimum Processor Level	Maximum Processor Level	Security Level
6.4	4Q2016			IBM System z196 & z114®	-	
6.3	7/2013	12/2017 <sup>[1]</sup>		IBM System z10®	-	EAL 4+ OSPP-LS
6.2	12/2011	07/2017 <sup>[2]</sup>	7/2013	IBM System z10®	z13	-
5.4	9/2008	12/2017 <sup>[3]</sup>	3/2012	IBM eServer zSeries 800& 900	zEC12	-

<sup>[1]</sup> Announced February 3, 2015

<sup>[2]</sup> Announced February 2, 2016

<sup>[3]</sup> Announced August 2, 2016

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

Do not assume End of Service dates will be moved again. Plan your migration to z/VM 6.4

## z/VM 5.4 End of Service

- The August 2, 2016 announcement changes the z/VM 5.4 End of Service date to December 31, 2017.
  - This is regardless to the z9 End of Service date, which was relevant in earlier announcements.
  
- It is important to understand the options and restrictions associated with running z/VM 5.4 and migrating to a new release.
  - Customers running on z9 or older would need a hardware update
  - Customers running on z10 would need to either migrate to z/VM 6.3 or update hardware for new z/VM 6.4 ALS level and then migrate to z/VM 6.4
    - z/VM 6.3 can be ordered only until z/VM 6.4 is generally available
  
- For information on the value of migrating to Version 6 of z/VM see:  
<http://www.vm.ibm.com/devpages/bitner/presentations/gotovm63.pdf>



# 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<sup>(TM)</sup>
    - <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2014.htm#2139>
  - z/VM V6.1 System SSL is FIPS 140-2 Validated<sup>(TM)</sup>
    - <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.*

# 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



# z/VM Version 6 Release 3

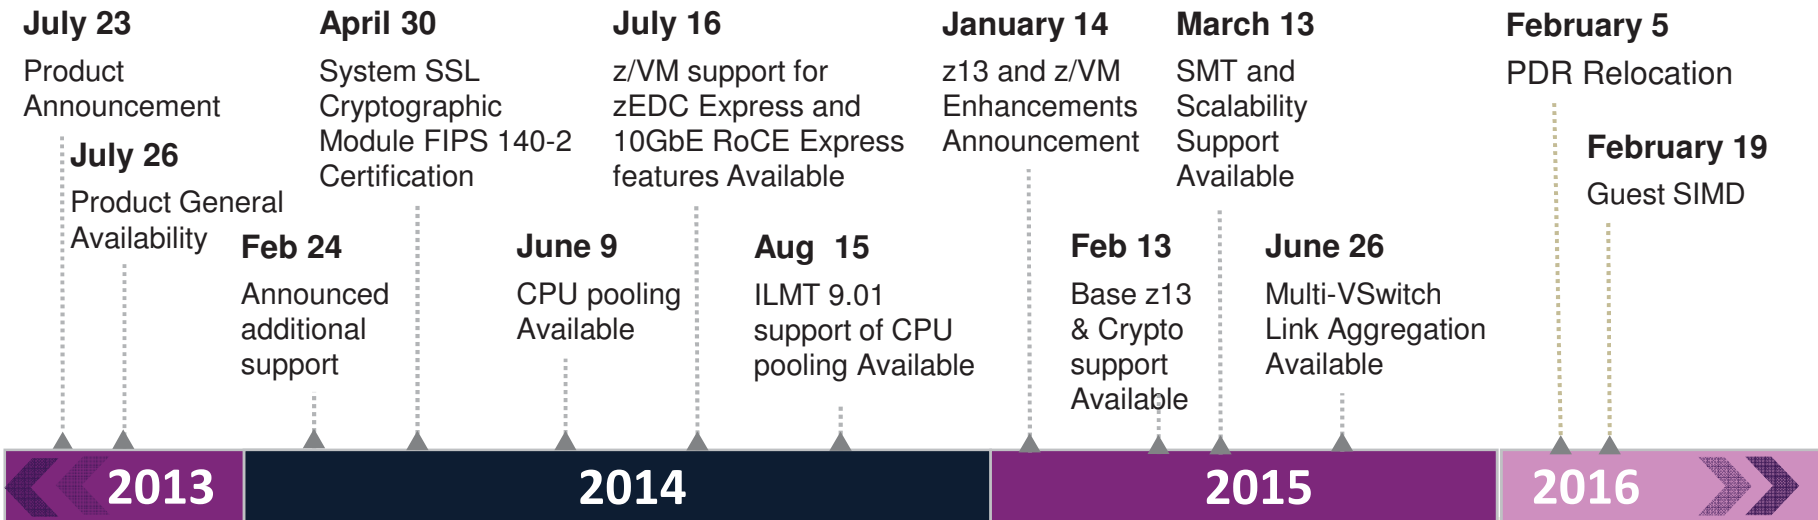
## z/VM 6.3 Themes

- 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




# z/VM Version 6 Release 3

## Making Room to Grow Your Business



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



# Large Memory Support



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

## 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

## 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

# 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



- “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




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.



# 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 <sup>1</sup>
z/VM 6.3 DS8K Synergy	Yes <sup>1</sup>	Yes <sup>1</sup>	Yes <sup>1</sup>
z/VM 6.3 SSI + LGR	No	No	Yes <sup>1</sup>

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
  - 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

# 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

## 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 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)



# What IBM Wave is NOT:

Not just for novice users

Not just a Graphical User Interface

Not a z/VM system programmer replacement

Not excluded by in-house scripts

Not just for 100s of servers

Not just a cloning tool



# IBM Wave offers value in:

## Advanced Visualization



- Shorten the learning curve needed to manage your Linux and z/V M environment
- Organize and simplify administration of virtual Linux servers;; automate and simplify management steps
- View servers and storage utilization graphically; view resource status at a glance
- Use graphical or tabular displays with layered drill down; customize and filter views
- Attach virtual notes to resources for additional policy-based management

## Simplified Monitoring





- Monitor z/VM system status through an innovative and interactive UI
- Monitor performance of CPU, paging devices, pool disks and more;
- Use agentless and lightweight discovery for a current view of your environment
- Use advanced filters, tagging, layout and layer selection to manage in a meaningful way
- Complements IBM OMEGAMON® XE used for in-depth performance monitoring and historical views

## Resource Management



- Manage your systems from a single workstation
- Assign and delegate administrative access using role based assignments
- Provision, clone, and activate virtual servers. Define and control virtual network and storage devices
- Perform complex tasks such as live guest relocation using a few keystrokes
- Execute complex scripts with a single mouse click
- Report on resources with flexible resource reporting

# A Different Better vs. A Standard Good





# 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).





# 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



## 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

## 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.

# 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



## 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.

# February 24, 2014 Announcements

# 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



# 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





# 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 9.0.1 updated August 2014- <http://ibm.biz/cpupoolimt>
  - Greater flexibility for IBM Passport Advantage products
- Support details:
  - z/VM 6.3 with APAR VM65419 – Available
    - Part of RSU 1501

# 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)

PVU Entitlements




- 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)

# 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



## 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



# January 14, 2015 Announcements


# 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/>

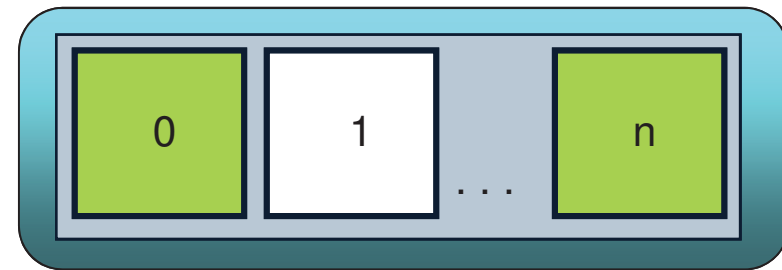


## z/VM Support for Crypto Express5S

- 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

# 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** - 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



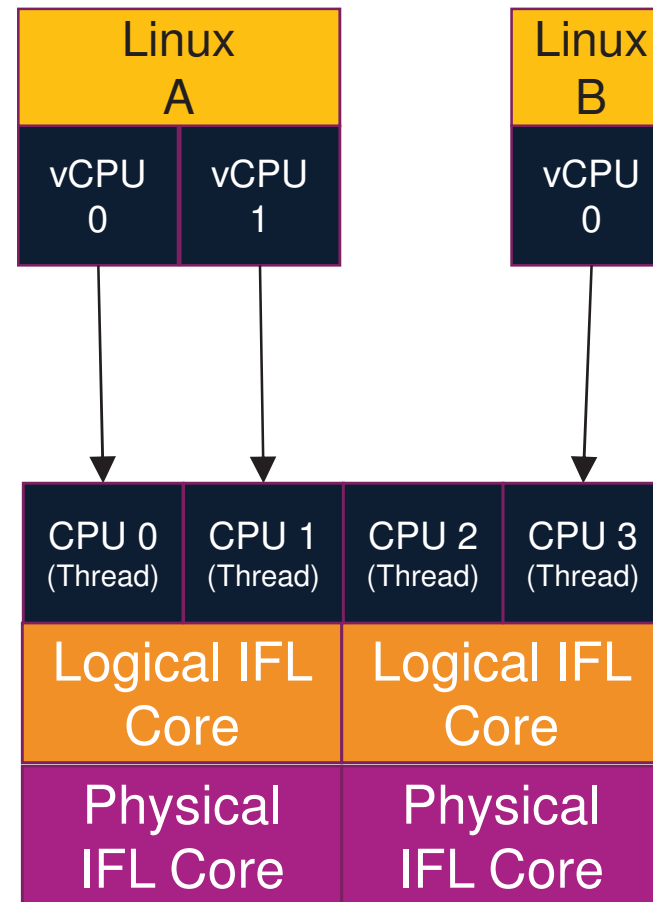
*Which approach is designed for the higher volume of traffic? Which road is faster?*

*\*Illustrative numbers only*



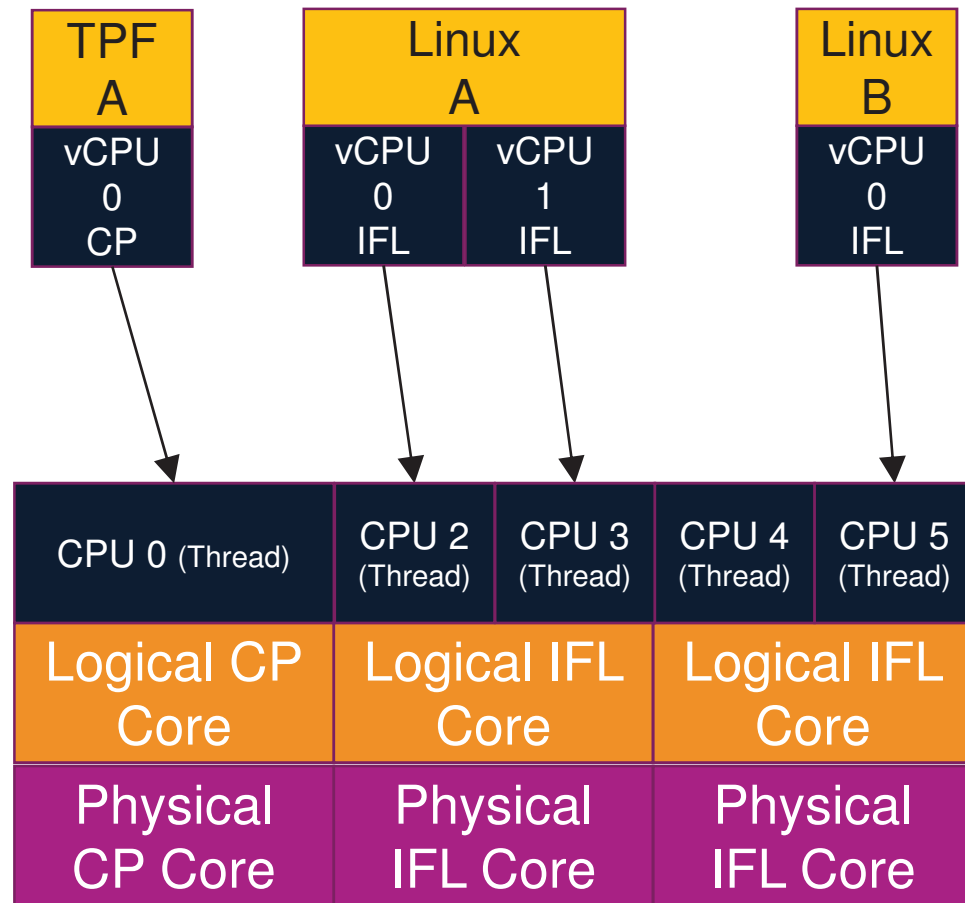
# SMT Usage

- 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.



# 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



## 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 cores (64 threads)
- For processors prior to z13
  - Limit remains at 32 cores
  - May still benefit from improved n-way curves



## Areas Improved with Scalability Enhancements

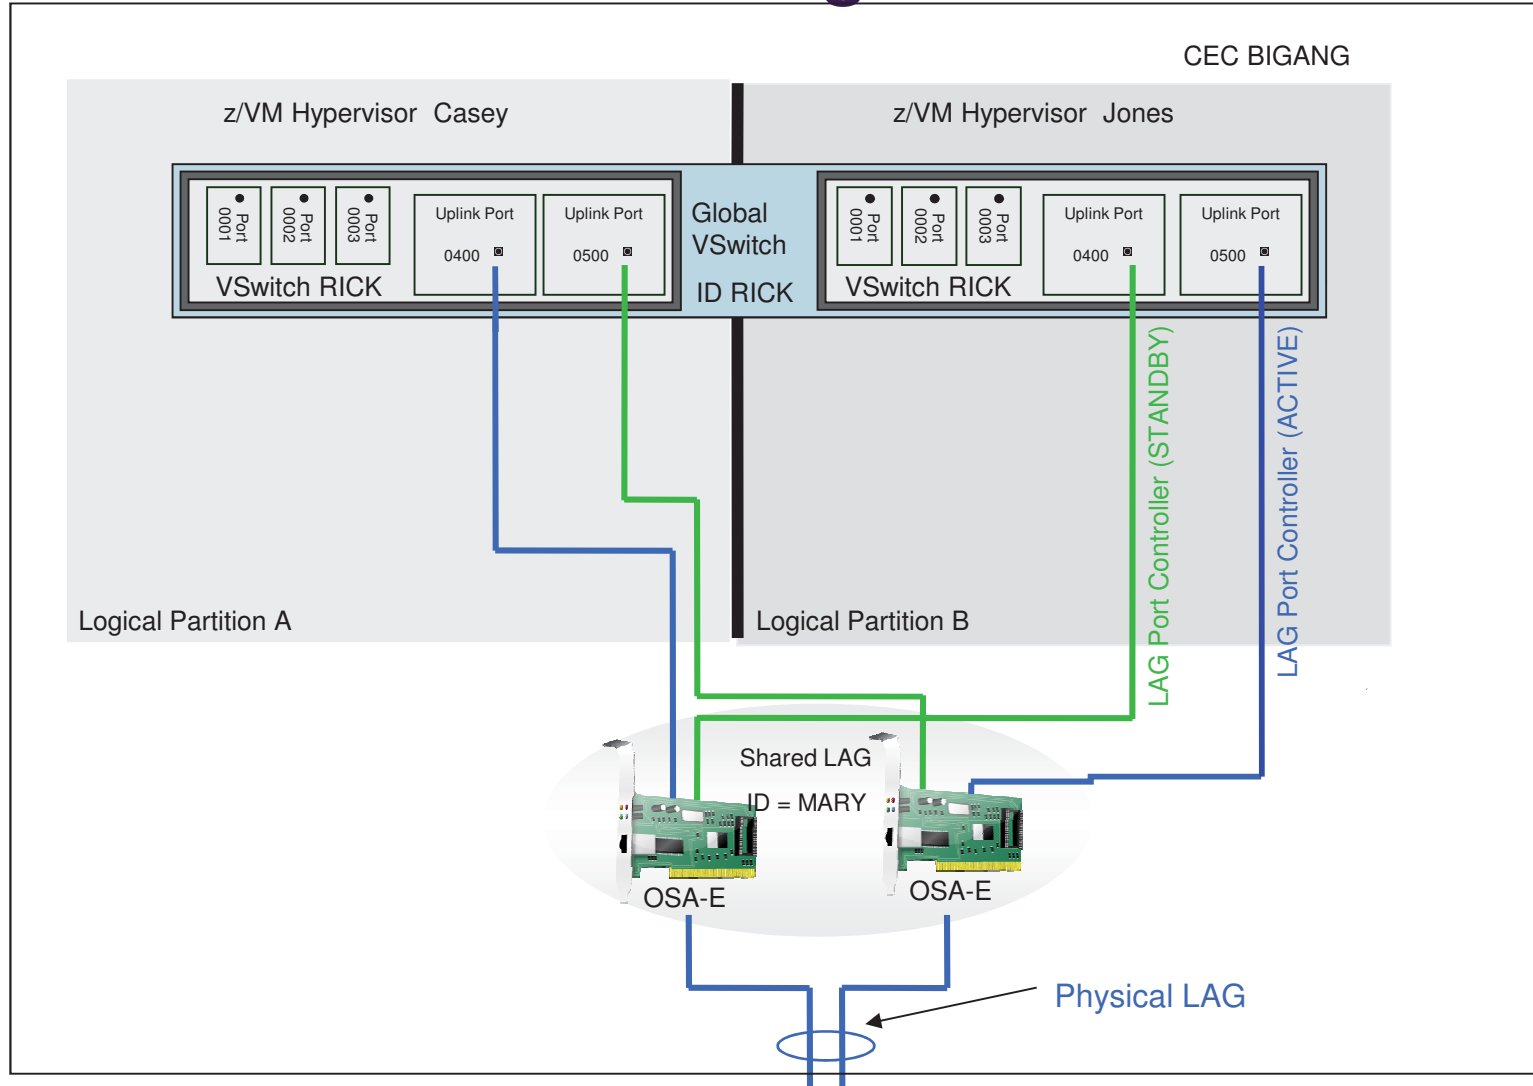
- z/VM Scheduler Lock
  - Management of internal stacked work
  - Guests going into a wait state
  
- Locking for Memory Management
  - Most benefit during system initialization and when very constrained with memory
  
- Serialization and processing of VDisk I/Os
  
- Batching and processor-local queues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

# 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



# Multi-VSwitch LAG Configuration



# September 2015 Updates

# September 2015 Enhancements

- Security Enhancements
  - RACF/VM Update
    - Password Encryption Enhancement (AES)
    - Password Change Management
  - TLS/SSL Server Updates
    - New cipher suites
    - Changed default suites (all supported z/VM releases)
  
- SMT Prorated Core Time
  - More fair representation of consumed resources
  - Used in accounting and monitor data
  - Used in limit settings
  - Combined with CPU Pools
  
- Logoff Performance Improvement



# Securing the Path to Virtualization

- Updates to RACF for z/VM V6.3 – **APAR VM65719**
  - PTFs available 11 September, 2015
  - Security enhancements include:
    - Password encryption upgrade
    - Helpdesk support
    - Special characters in passwords
    - Minimum password change intervals
- Updates to the z/VM TLS/SSL server – **APARs PI40702, VM65717, VM65718**
  - PTFs available 11 September, 2015
  - Maintains FIPS 140-2 and NIST SP 800-131a compliance
  - Function includes:
    - System SSL V2.1 equivalency (V6.3 only)
    - AES Galois/Counter Mode encryption (V6.3 only)
  - **Changes to default cipher suites and protocols (all serviced releases)**
- More information at <http://www.vm.ibm.com/security/>



# RACF Password Encryption Upgrade

- Enables a stronger encryption mechanism for passwords and/or passphrases in a RACF database
  - *Matches support delivered by z/OS APAR OA43999*
  - *Strengthens RACF database against offline attacks*
  - Mitigates compliance issues associated with older encryption algorithms
  
- Migration to KDFAES is for an entire RACF database
  - Might cause problems if sharing this RACF database with another system!
  - Utilities available to convert databases and to clean password histories
  
- Some restrictions might apply
  - Support is for z/VM 6.3 only
  - RACF template has been updated; run RACFCONV accordingly
  - CPACF (feature 3863) must be enabled

# SMT Prorated Core Time Support

- APAR VM65680 available for z/VM 6.3 on September 2, 2015
- Applies only to z/VM systems where SMT has been enabled
- This support enforces capacity limits using core time rather than thread time so that a CPU pool will not be limited prematurely.
- Following interfaces have been updated:
  - Commands
  - Accounting records
  - Monitor records
  - Store hypervisor information (STHYI) instruction
- CPU Pooling and ILMT can now be used without the need to potentially adjust the pool values to be equivalent to non-SMT environment.
- Use QUERY CPUPOOL to determine if APAR is applied (shows limits in ‘cores’ instead of ‘CPUs’)

# Logoff Performance Improvement

- Improvement to avoid very long logoff times, particularly for large guests.
- As part of reset processing in logoff, all page management blocks (page tables) must be paged in to complete the logoff.
- New batch read process significantly improves the processing of the page management blocks (PGMBKs)
- APAR VM64770 closed September 8, 2015 as PTF UM34582
- Available for only z/VM 6.3

# February & March 2016 Updates

# February & March 2016 Updates

Updates to z/VM 6.3 only

- Dynamic Migration of the SSI PDR
  - Move the PDR in advance of a planned outage
  
- SIMD Guest Exploitation support
  - Requires z13 or z13s
  
- OpenStack Enablement changes

## Dynamically Migrate the SSI PDR

- Enhancement to be able to relocate the Single System Image (SSI) Persistent Data Record (PDR) to a different volume without a planned outage
  
- Avoids the need for a cluster-wide shutdown in order to move PDR area to a new device or new storage server
  - Facilitates moving to a new storage server
  - Does not address unplanned outage of the PDR volume
  
- New option on the CP **SET SSI** command
  - CP SET SSI PDRVOLUME valid
  
- VM65712 – Closed February 5, 2016, PTF UM34736
  - z/VM 6.3 Only

## SIMD Guest Exploitation Support

- Guests with appropriate support can now use the Vector Facility for z/Architecture, aka Single-Instruction Multiple-Data (SIMD) functions introduced with the z13 servers.
  - Access to 32, 128-bit registers.
  - Potential performance improvements for exploiting software
  
- Supported by Live Guest Relocation
  
- VM65733 – Closed February 19, 2016, PTF UM34752
  - z/VM 6.3 Only
  - Requires z13 or z13s and guest at appropriate level



# OpenStack Enablement Changes

- Support of Liberty Release of OpenStack
- Integration of xCAT function into the z/VM Cloud Manager Appliance (CMA)
  - Running a fully functional z/VM OpenStack solution in a single virtual machine per z/VM system
  - z/VM Cloud Manager Appliance 1.2.0
- Support for provisioning Red Hat RHEL 7 and SUSE SLES 12 servers
- Uses the following:
  - OpenStack authentication (Keystone) V3 API.
  - RabbitMQ as the OpenStack message queue implementation.
  - OpenStack Dashboard (Horizon) as the self service portal.
- Available March 24, 2016 via APAR VM65780

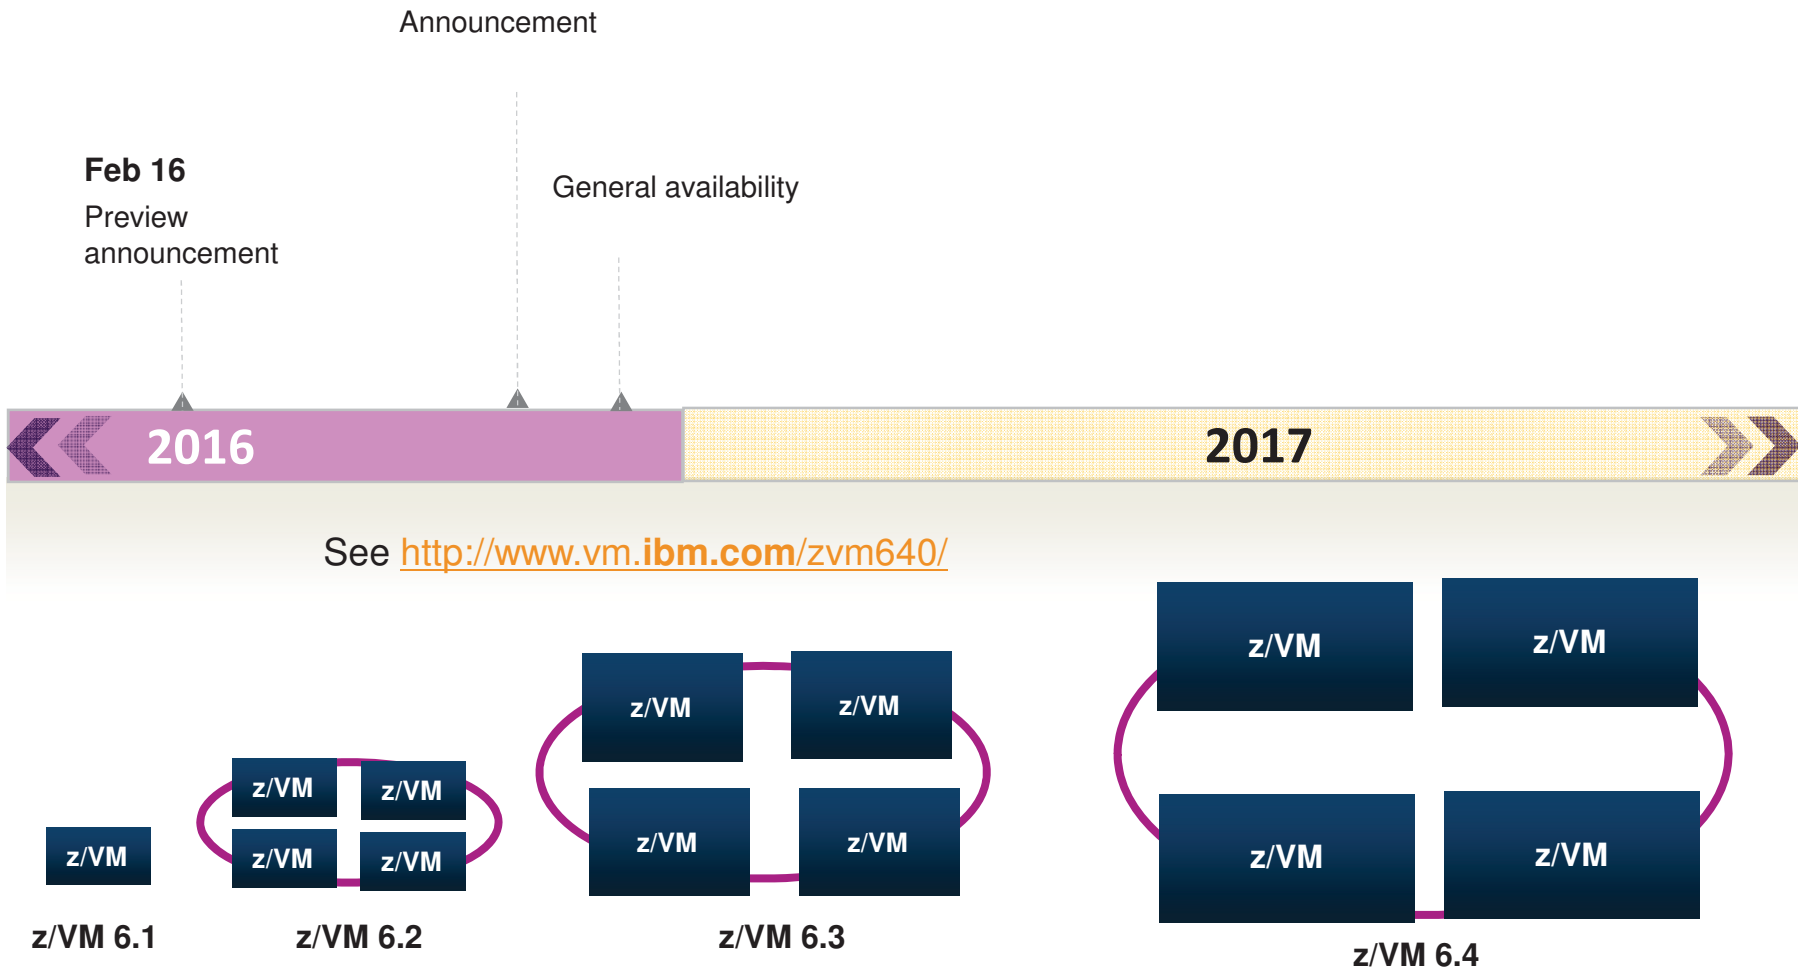
## Excessive Simulation In Certain Circumstances

- Note: For customers using SSI clusters, appropriate service is required to avoid excessive processor usage due to extra simulation involved with managing relocation domains.
- Ensure that you have APAR VM65692 installed (RSU 1601), especially if you have APAR VM65638 (RSU 1501)
- So, just be aware: if you are running SSIs and have installed VM65638, also install VM65692

# February 16, 2016 Announcement Preview z/VM 6.4

# z/VM Version 6 Release 4

Designed for Clients of Today and Tomorrow



## Preview IBM z/VM 6.4

- Preview announcement 216-009, dated February 16, 2016
  - <http://www.vm.ibm.com/zvm640/index.html>
- Planned availability date Fourth Quarter 2016
- A release born from customer feedback
- Key components:
  - Enhanced technology for improved scaling and total cost of ownership
  - Increased system programmer and management capabilities
- New Architecture Level Set (ALS) of z196 and higher



## z/VM 6.4 Preview Topics\*

- Efficiency and Elasticity
  - HyperPAV Paging
  - Guest Large Page Support
  - Guest Transactional Execution
  - Memory Scalability Improvements
  - SoD: Dynamically Managed Simultaneous Multithreading (SMT)
- Ease of Use Improvements
  - CP Environment Variables
  - Query Shutdown
  - CMS Pipelines
- Hardware Currency
  - ESA/390 Removal
- Migration Enhancements
  - Single System Image member upgrade
  - Upgrade from z/VM 6.2 or z/VM 6.3

\* Subset of published preview items

# Efficiency and Elasticity

# HyperPAV Paging

- Traditional S/390 I/O limited to one concurrent I/O operation per volume
  - Single cache-miss locks out cached data during back-end I/O
  - Many small volumes to allow for parallelism
    - Large number of page volumes for large z/VM images with memory overcommit
- Parallel Access Volumes (PAV)
  - Alias devices provide an alternative path to access the data
  - Multiple devices associated with the same base volume
  - HyperPAV – Pool of alias devices associated with the base for a single I/O
- HyperPAV Paging in z/VM 6.4
  - Exploit DS8000 features to improve z/VM paging throughput
  - HyperPAV to use alias devices for paging (and other I/O to CP Owned devices)
  - High Performance FICON for increased efficiency
  - Tuning option to retain or release page slot after page-in
- Expected Advantages
  - Increased paging bandwidth because of PAV
  - Opportunity to use larger volumes for paging
  - Reduced CPU overhead for paging



# Guest Large Page Support

- Page size of virtual memory remained 4K over decades
  - Typical memory size increased with several orders of magnitude
  - Overhead of memory management becomes a concern
  - With virtual memory guests like Linux the effect multiplies
    - Translation Lookaside Buffers become less efficient
    - Memory overhead required for memory management gets large
      - Oracle guest with 500 connections uses half of the memory for page tables
  
- Application subsystems like databases and JVM manage their own memory
  - No need for the Linux operating system to add the 4K granularity
  
- Guest Large Page Support in z/VM 6.4
  - Enables Enhanced DAT support for the guest to use large pages (1 MB per page)
  - Use `hugepages=` kernel parameter in Linux to activate

# Guest Transactional Execution Support

- Multi-processing overhead increases when number of processing threads increases
  - Contention between threads when accessing shared resources
  - Scalability concern – diminishing benefit from adding more processors
  - Primary solution in the design of the operating system and application
  
- Hardware implementation of transactional memory
  - Introduced with the zEC12
  - Builds on memory cache architecture
  - Allows for atomic operations on shared data without explicit locking
    - Instruction sequence with escape path to retry or abort
  
- Guest Transactional Execution Support in z/VM 6.4
  - Allows a guest to exploit Transactional Execution
  - Performance improvement for multi-processing systems and applications
  - Exploited by Java Virtual Machine

# Memory Scalability Enhancements

- Major redesign of memory management with z/VM 6.3
  - Removed infamous reorder processing based on reference bits
  - Implementation of the age list with early paging
  - Eliminated the need for expanded storage
  
- Memory Scalability Enhancements in z/VM 6.4
  - Address spin lock contentions observed in the field
  - Further enhancements to improve memory management efficiency
  - Provide a framework for future increase in supported real memory size for z/VM

# Dynamic Simultaneous Multithreading Level

- Simultaneous Multithreading (SMT) introduced with z/VM 6.3
  - Performance improvement for some workloads
  - Enabled or disabled through system configuration option
  - IPL to disable it again was too disruptive to encourage experimentation

**Dynamically managed thread activation levels.** IBM intends to provide support in a future z/VM deliverable that will allow clients to dynamically manage the number of activated threads per configured core that can be enabled for simultaneous multithreading (SMT) without requiring an IPL of the z/VM system.

- Statement of Direction in z/VM 6.4 Preview
  - Intent to deliver the function with z/VM 6.4
- Dynamic SMT
  - Enabling SMT still requires an IPL
  - Maximum 32 cores with SMT enabled (be aware when mixing CP and IFL)
  - Number of active threads per IFL core can be changed without an IPL
    - Switch between SMT-1 and SMT-2
  - Setting is system-wide – all IFL cores either with one or two active threads
  - Allows installations to experiment and measure impact
  - Disabling SMT still requires an IPL

# Ease of Use Improvements

# CP Environment Variables

- System wide configuration settings for customer usage
  - Useful to guide automation and operations processes
    - Distinguish between disaster recovery test and actual disaster recovery
    - Control startup of guests after IPL
    - Replace hidden assumptions on CPUID or node name
  - Avoids the need for creative tricks in systems management solutions
  
- CP Environment Variables in z/VM 6.4
  - Global system wide variables names with a value
    - Privilege required to set the value
    - Privilege class G to view the variables
    - Long names provide a fair amount of flexibility
    - Some reserved names for use by z/VM itself
  - Likely exploited in Linux (through vmcp command)
  - Available in programs on CMS (REXX or CMS Pipelines)

# Query Shutdown

- Provide automation and systems management tools information about system status
  - Determine whether shutdown is in progress
  - Determine the reason for the system shutdown
- Eliminate the need for creative alternative solutions
  
- Query Shutdown in z/VM 6.4
  - Users can tell whether shutdown is in progress
    - Distinguish between SIGNAL SHUTDOWN and system shutdown
  - Privileged users can tell who issued the shutdown
  - System operator messages to report the details

# CMS Pipelines

- CMS Pipelines has been functionally frozen for two decades
  - Unsupported “plastic pipes” for download on non-IBM web site
  - Increased concerns about production use of unsupported software
  - New installations are confused about software and documentation
  - Tools developers are limited in what can be expected
  
- CMS Pipelines in z/VM 6.4
  - Built-in version of CMS Pipelines is upgraded
  - Includes most function provided by Runtime Library Distribution (Marist download)
  - Additional bug fixes and improvements included
  - Supported and serviced through usual processes
  - Future enhancements through service process and releases
  - Single source of product documentation
  - On-line help updated to match product documentation





# Hardware Currency

## ESA/390 Removal

- The z13 will be the last z Systems server to support an ESA/390 operating system
  - Future machines require z/Architecture mode operation
  - Avoid spending hardware resources on features that are not needed anymore
  
- Implications for z/VM 6.4
  - IPL process and stand-alone utilities run in z/Architecture mode
  - New facility to simulate a z/Architecture-only environment
    - Through CP command or directory (GLOBALOPTS)
  
- CMS can continue to run in ESA/390 mode
  - CMS runs without Dynamic Address Translation (DAT-off)
  - Still wise to look at ZCMS to run CMS in z/Architecture mode
    - Ability to address more than 2 GB of virtual memory
    - Access to wide 64-bit registers
    - New instructions that (e.g. SIMD vector instructions and new crypto support)
  - More components will start using ZCMS

# Upgrade in Place


- Upgrade in place introduced with z/VM 6.3
  - Enhanced to upgrade Single System Image members as well as stand-alone systems
  - Upgrade to z/VM 6.4 from z/VM 6.2 or z/VM 6.3
  
- For upgrade of an SSI cluster all members must be on the same previous release
  - Either all z/VM 6.2 or all z/VM 6.3 – no mix

## z/VM 6.4 Supported Hardware

- Following z Systems servers:
  - z13
  - z13s
  - LinuxONE Emperor
  - LinuxONE Rockhopper
  - IBM zEnterprise EC12
  - IBM zEnterprise BC12
  - IBM zEnterprise 196
  - IBM zEnterprise 114
  
- Electronic and DVD install
  - No tapes
  
- Installation drops 3390-3 (continues 3390-9, etc)



# Summary



## 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.