IBM®

z/VM Platform Update – MVMUA April 2015

March 10, 2015 Version 4.6

Bill Bitner
z/VM Client Focus and Care
bitnerb@us.ibm.com

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.

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

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

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

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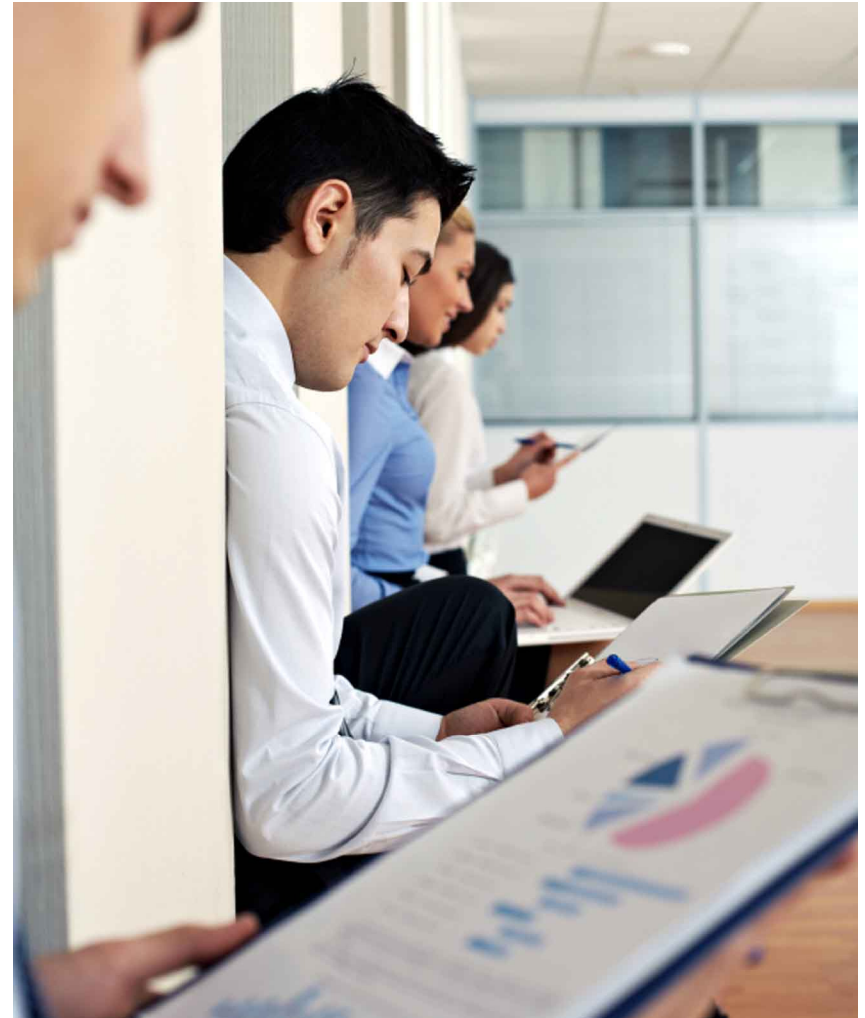
Acknowledgments – Platform Update Team

- Alan Altmark

- Bill Bitner

- Miguel Delapaz

- Glenda Ford

- John Franciscovich

- Les Geer

- Susan Greenlee

- Dan Griffith

- Brian Hugenbruch

- Emily Hugenbruch

- Romney White

Agenda

- Release Status and Information

- z/VM ® Version 6 Release 3

  - 2014 Enhancements

  - 2015 Enhancements

- Futures and Statements of Direction

Release Status and Information

z/VM Release Status Summary

z/VM Level	GA	End of Service	End of Marktg.	Minimum Processor Level	Maximum Processor Level	Security Level
6.3	7/2013	12/2017[5]		IBM System z10®	-	EAL 4+[2] OSPP-LS
6.2	12/2011	12/2016[3]	7/2013	IBM System z10®	z13 [4]	-
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] Targeted Security Level in V6.3 SOD
[3] Extended from original date (Announced February 4, 2014)
[4] Announced January 14, 2015
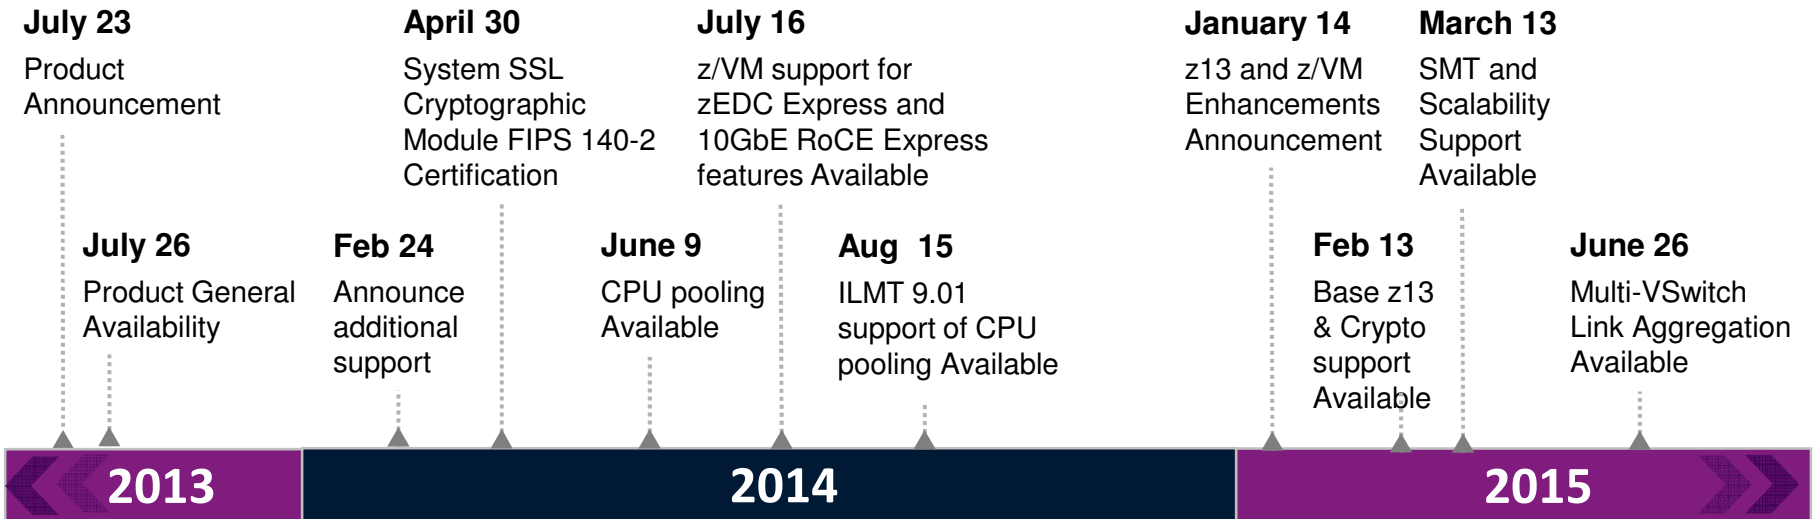[5] Announced February 3, 2015

Marketed & Serviced

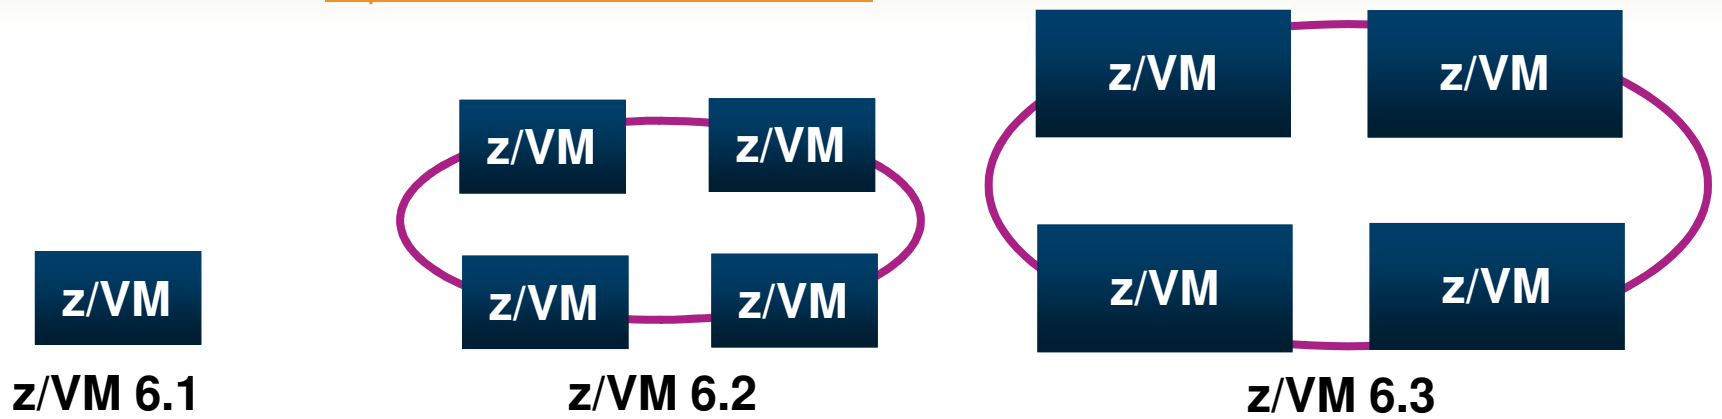Serviced, but not Marketed

End of Service & Marketing

IBM z Systems

IBM

z/VM Version 6 Release 3

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

z/VM   z/VM

z/VM   z/VM

z/VM   z/VM

z/VM   z/VM

z/VM 6.1   **z/VM 6.2**   **z/VM 6.3**

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

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

- Expanded Storage continues to be supported with a limit of **128 GB**
  - However, expanded storage is no longer recommended.

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
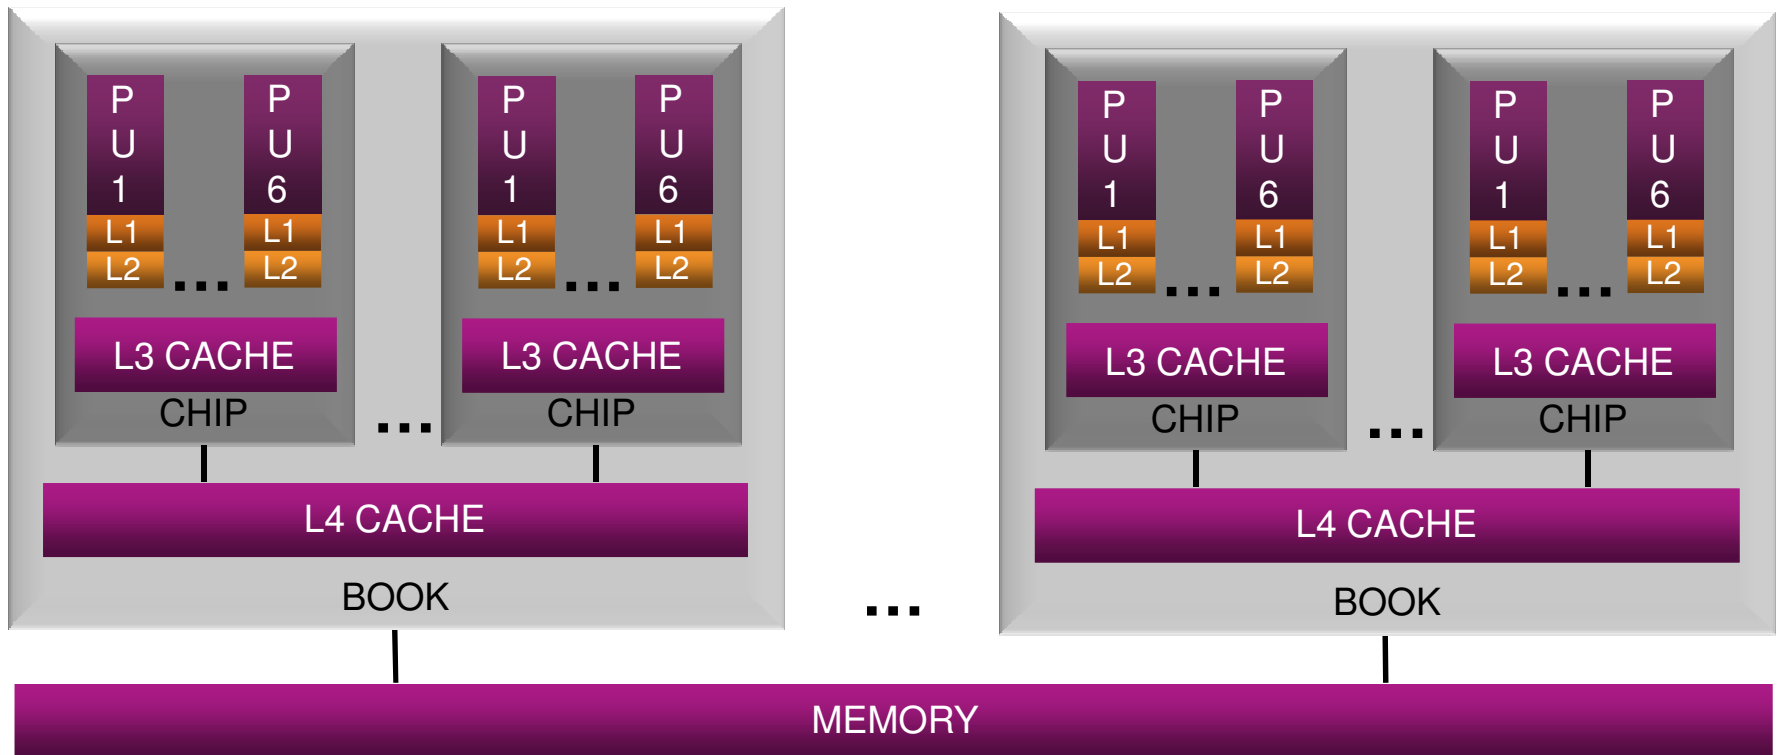
HiperDispatch – Dispatching Affinity

- 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)
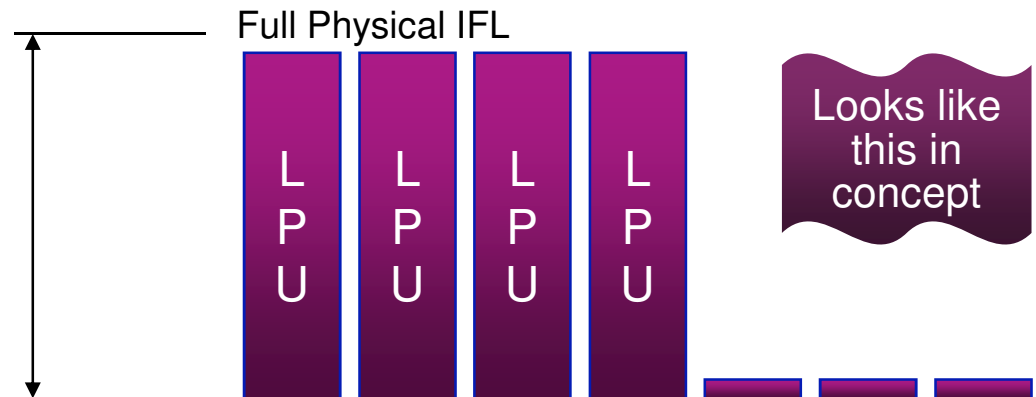
Horizontal vs. Vertical CPU Management

Full Physical IFL

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

57%

| L P U | L P U | L P U | L P U | L P U | L P U | L P U |

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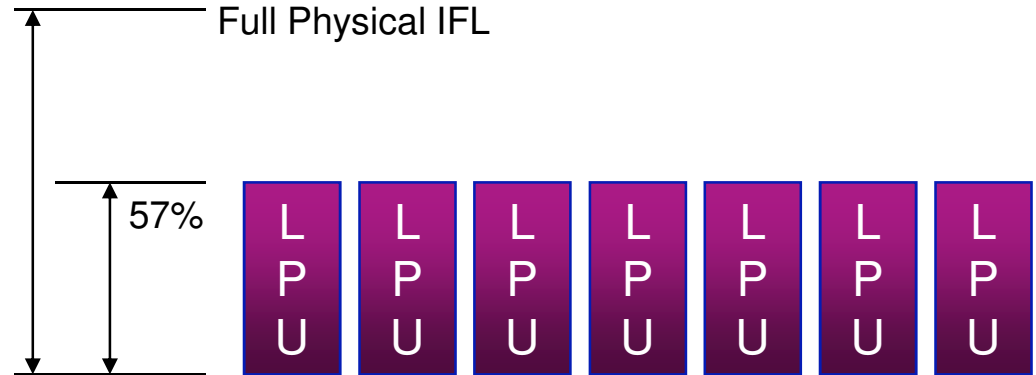
Vertical:
- The logical processors are skewed to where some get greater share of the weight.
- z/VM dispatches work accordingly to the heavier weighted workload.

| L P U | L P U | L P U | L P U |

Looks like this in concept

A Different Better vs. A Standard Good
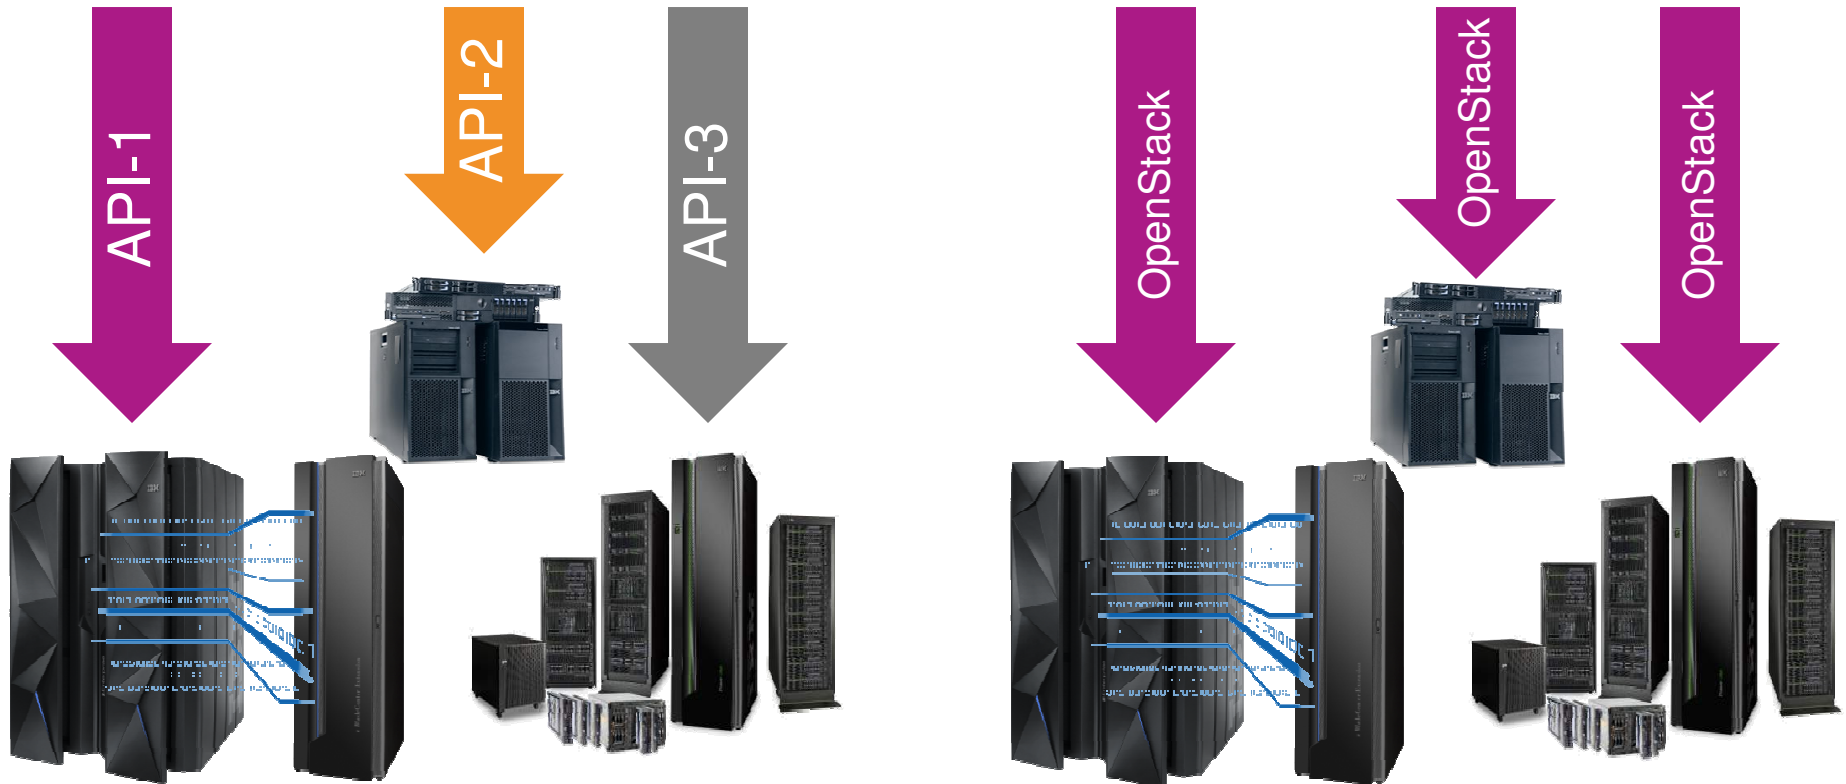
The OpenStack Food Chain

IBM Cloud Technology Products

openstack™
CLOUD SOFTWARE

Communicates with z/VM

IBM

xCAT Appliance

SMAPI Services

Directory Product

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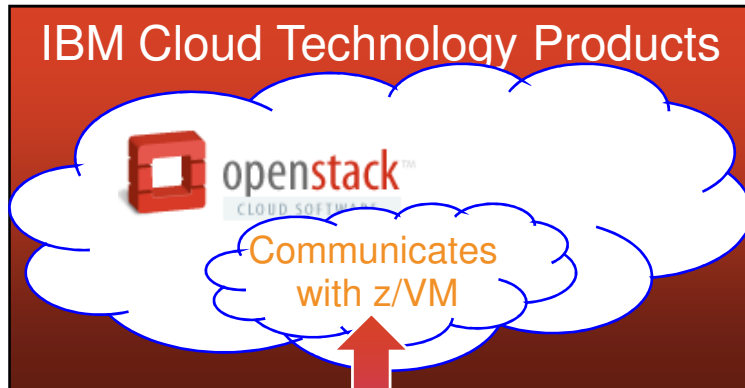
- **Top Half of the Solution:**
  –An IBM SmartCloud Technology product or other vendor product will include the OpenStack support.
  –Portions of that OpenStack support will know z/VM (i.e. code that connects and understands how to talk to z/VM).

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

  Product with OpenStack Support

  z/VM 6.3 Product

  Additional Product or Feature
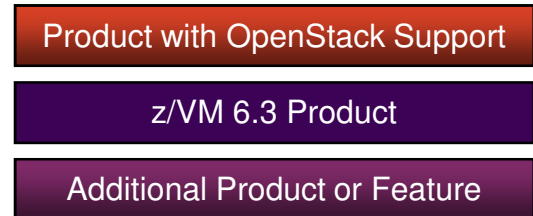
OpenStack Related Products

- IBM Cloud Manager with OpenStack
  - Formerly known as IBM Smart Cloud Entry
  - "Managed from" z Systems
  - V4.1 Currently Available
  - V4.2 Announced February 24, 2015, available March 13, 2015

- 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 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 V1.1

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

All support z/VM 6.2 and 6.3!

- **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

February 24, 2014 Announcements

Enhancing the Foundation for Virtualization

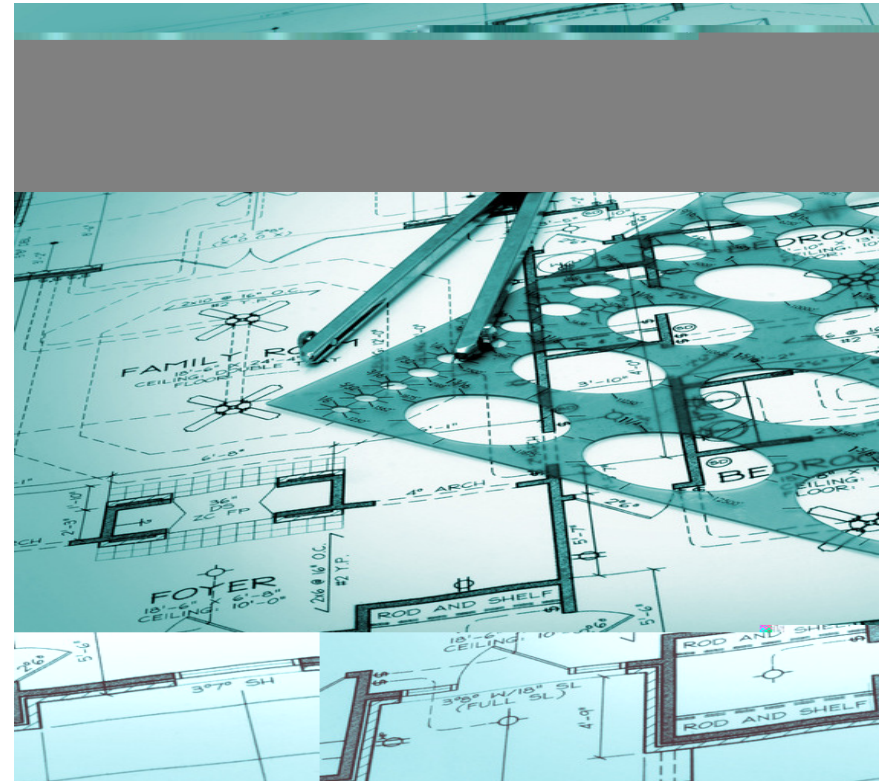- <u>Release for Announcement – zBX and zEnterprise System Enhancements</u>

  – 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
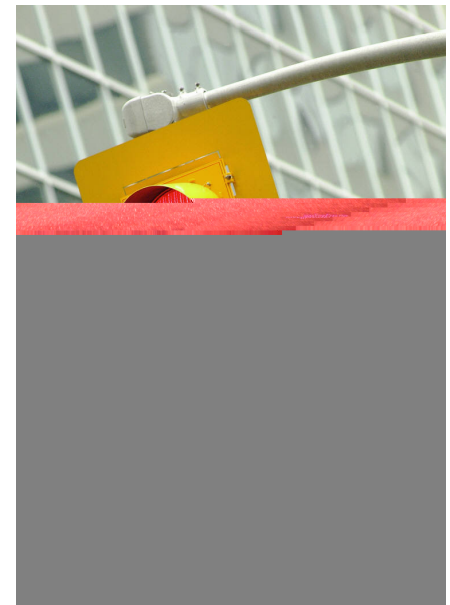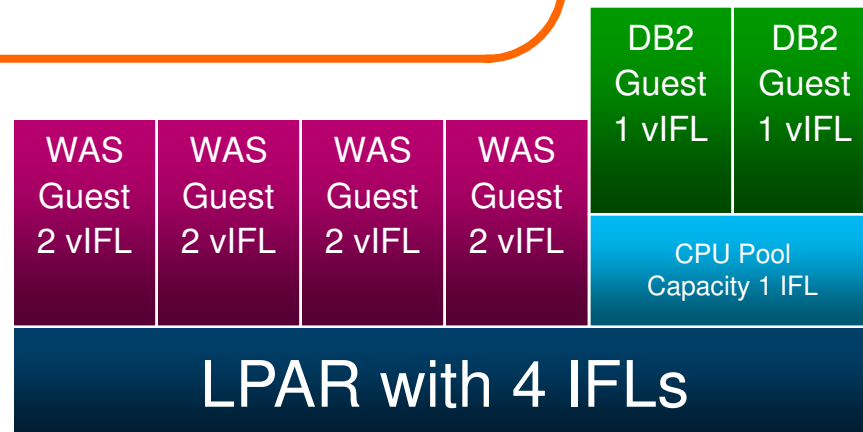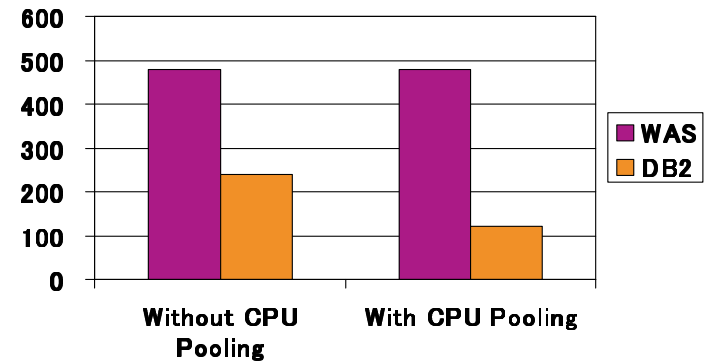
Environment Information Interface

- 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 9.0.1updated 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

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

	WAS	DB2
Without CPU Pooling		
With CPU Pooling		

DB2 Guest 1 vIFL	DB2 Guest 1 vIFL

| WAS Guest 2 vIFL	WAS Guest 2 vIFL	WAS Guest 2 vIFL	WAS Guest 2 vIFL	CPU Pool Capacity 1 IFL

LPAR with 4 IFLs

- 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 zEC12 or zBC12 with appropriate updates – see support buckets
    - z/VM 6.3 with APAR VM65417 – Available
        - 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
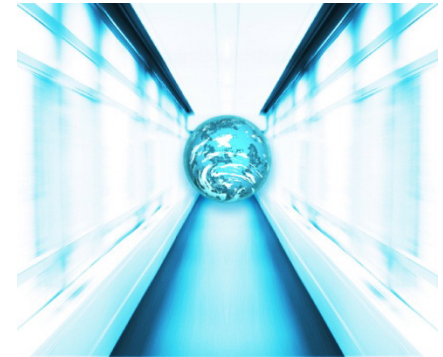
zEDC Express Feature

- 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
    - 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
  - Refer to bucket for full list

- 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)

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 planned to be available March 13, 2015

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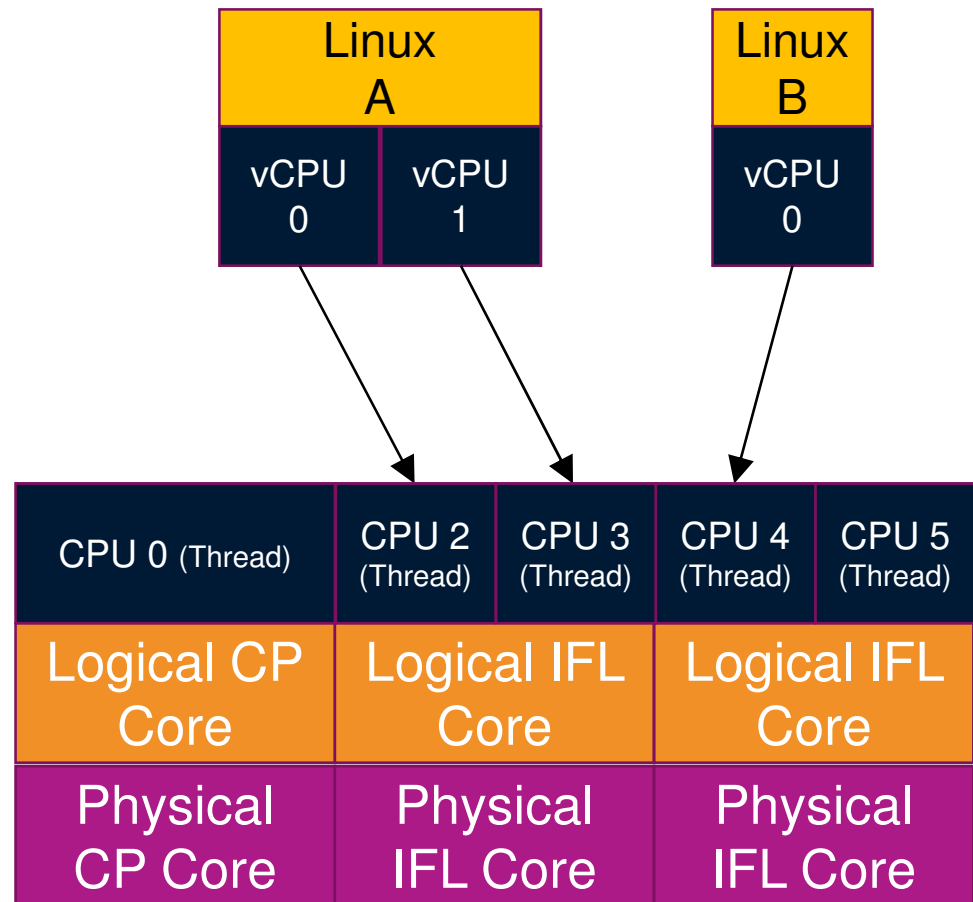 - Workload dependent
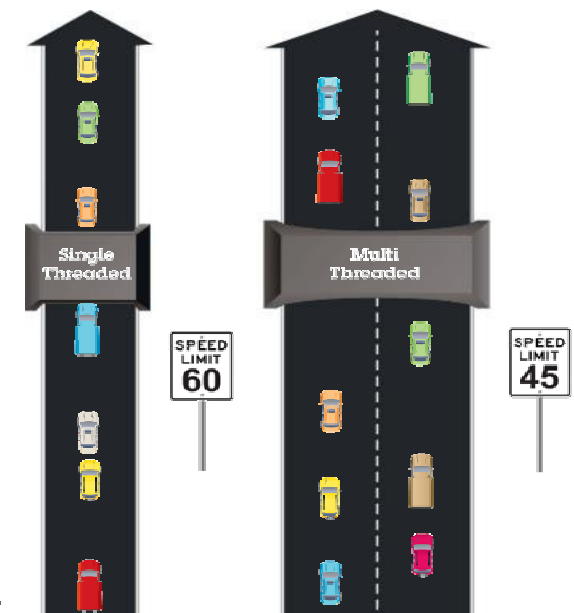
Which approach is designed for the higher volume of traffic? Which road is faster?

*Illustrative numbers only

SMT Dispatching

- Physical IFLs (or Cores) with SMT allow up to two threads to be used.

- Logical IFLs are presented to z/VM as in the past.

- z/VM creates a CPU associated with each thread for it to use.

- The virtual CPUs of guests can then be dispatched on different threads intelligently, based on topology information.

- In a mixed-engine environment, general purpose processors can not do threading, but a second CPU address is consumed (CPU 1 in example)

Linux A		Linux B
vCPU 0	vCPU 1	vCPU 0

CPU 0 (Thread)	CPU 2 (Thread)	CPU 3 (Thread)	CPU 4 (Thread)	CPU 5 (Thread)
Logical CP Core	Logical IFL Core		Logical IFL Core	
Physical CP Core	Physical IFL Core		Physical IFL Core	

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

- APAR VM65586 for z/VM 6.3 **only**
  - PTFs planned to be available March 13, 2015

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

- For processors prior to z13
  - Limit remains at 32
  - 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

- Makes it possible to do Link Aggregation with VSwitches without 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 or z/VM systems without support

- APARs VM65583 and PI21053 for z/VM 6.3 **only**
  - PTFs planned to be available June 26, 2015

- Only available on z13
  - Requires OSA enhancements introduced with the z13

- Allows better consolidation and availability while improving TCO

Multi-VSwitch LAG Configuration
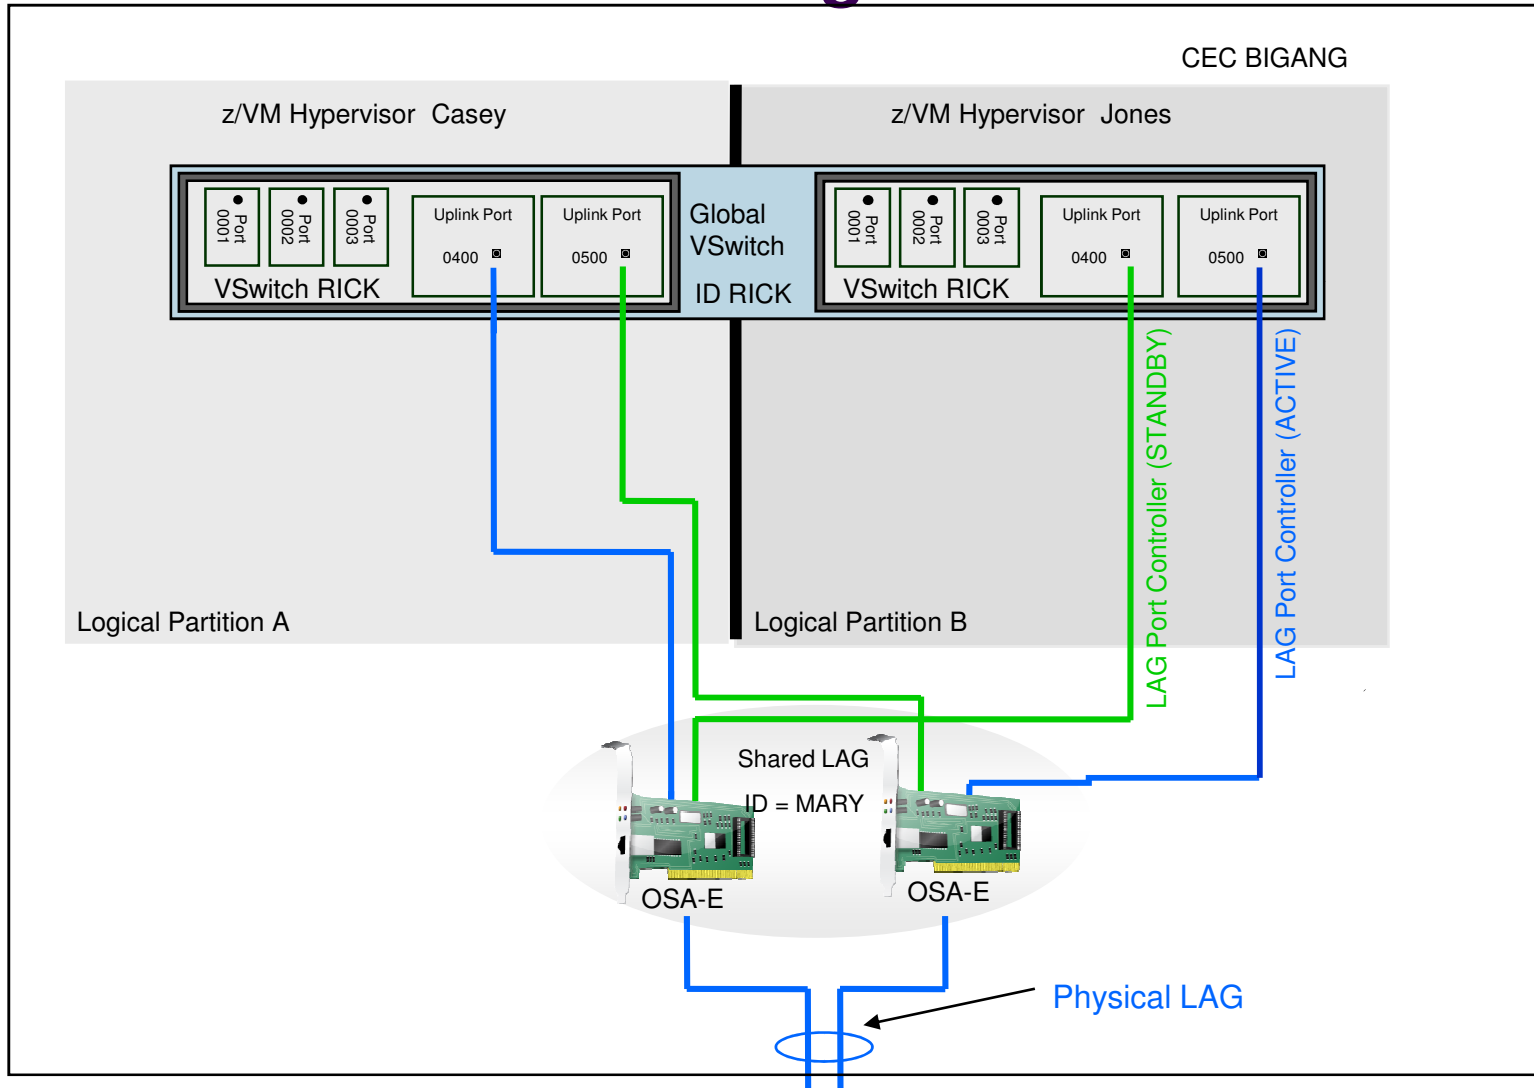
Hardware Support

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

- 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.**

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.

- 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

UPDATE

WARNING!

Statements of Direction
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.

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.

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.

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.

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.

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.

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.

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.

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.

Summary
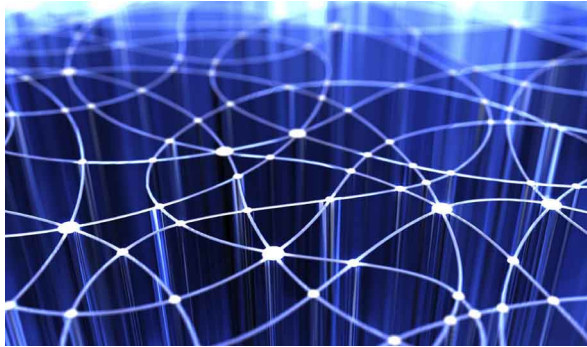
Leadership

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. Now the 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.