



z/VM Platform Update & Linux and VM Program Opening

Rick Barlow - Nationwide Insurance

Richard Lewis - IBM

Bill Bitner - IBM

August 12, 2013

Session 13506



Agenda

- SHARE LVM Program Overview and Information
- Survey
- z/VM Platform Update



Welcome to SHARE: Introduction to Linux and VM Program

Rick Barlow
Nationwide Insurance



LVM Opening Agenda

- What is LVM?
- Information
 - SHARE – “It’s what we do”
 - LVM organization - what are the projects and brief description
 - LVM team; IBM reps & Special guests
- How to connect - Look for bears and penguins
- This week...
 - Highlights and Grid
 - Program Dinner announcement
- Evaluations Please!



Welcome to SHARE

“It’s what we do!”


- A volunteer, customer run organization.
- Dedicated to providing top-tier education on IBM and related technologies
- Networking with peers, sharing information and experience
- Industry influence – direct contact with IBM, vendors
- Trade Show – SHARE Technology Exchange (STE) – the latest in commercial offerings

Linux and VM Program (LVM)

More information:

<http://tinyurl.com/SHARE-LVM-info>

- Dedicated to supporting and promoting the two best operating systems in the known universe
- Provide quality education on both environments, plus the power of combining them.
- Forum for expressing concerns, problems, opinions on future of Linux, VM and System z.
- Find support and help from fellow sysprogs and admins.



Understanding the “Program”

- SHARE programs consist of “projects” dedicated to specific subjects or issues
- LVM Program has three projects:
 - Linux Project
 - VM Project
 - Virtualization Project
 - Linux and VM Technical Steering Committee

Its all about the Volunteers!



All of our Speakers and
Session Chairpersons!

- Project Managers:
 - Linux: Mark Post (Suse)
 - VM: Dan Martin (Rocket)
 - TSC Chair: Marcy Cortes (Wells Fargo)
- Deputies:
 - Neale Ferguson (SineNomine)
 - Openings available
- Scheduler:
 - Rich Smrcina (Velocity)
- Chair Wrangler:
 - Brian Jagos (CA)
- Program Dinner Host:
 - Jim Moling (Dept. of Treasury)
 - Gail Riley (EMC)

Couldn't do this without them....

- IBM Representatives:

- Richard Lewis
- Alan Altmark
- Bill Bitner
- Jim Elliott
- John Franciscovich
- Romney White




- Special Guests

- John Birtles
Director of System z Product Portfolio Management
- Steffen Thoss
System z Linux and Virtualization Product Development Team leader

How to Connect

- Many ways to connect
 - Look for the Bears and Penguins
 - SHARE Linux and VM Program web page
<http://www.share.org/p/cm/ld/fid=44>
 - SHARE Linux and VM Program community
<http://www.share.org/p/co/ly/gid=1832>
 - Requires member login



Highlights for LVM in Boston

- Direct, “hands-on” experience
 - Traditional labs (Room 202)
 - 13518: Linux for Beginners Hands-on Lab (3 parts) Monday 9:30
 - 13502: z/VM SSI Installation and Configuration *or* non-SSI to SSI Migration Hands-on Lab (3 Parts) Tuesday 1:30
 - “BYOC” labs (Room 308)
 - 13494: Introduction to REXX Workshop (2 Parts) (BYOC) Tuesday 9:30
 - *For “BYOC” – make sure you have a TN3270 emulator installed and fully charge your laptop*

Highlights for LVM in Boston

- Customer experience sessions
 - SICOOB, Boston University, Wells Fargo, Nationwide Insurance
 - Customer panels: z/VM 6.3 ESP, Experiences with Linux on System z
 - z/VM Single System Image and Live Guest Relocation Panel Discussion Wed at 9:30 in 309
 - z/VM 6.3 Early Support Program Experiences Thu at 9:30 in 309
 - Experiences With Linux and System z – Customer Panel Thu at 11:00 in 309
- “Grids” – All of the LVM sessions and some other session of interest
 - <http://tinyurl.com/SHARE121lvm>
 - Cards available with QR code

Linux and VM Program Dinner Eat, Drink and Be Merry



- Wednesday 7:00 at Durgin Park
We will gather in a lobby for the trip (Dutch treat)
- Have dinner; unwind; meet other Linux and VM people
- See Gail Riley if you didn't already sign up

Evaluations are Important!

Tell us how we are doing...

- LVM planners really use information from evaluations to determine what to offer
 - Which sessions are in demand?
 - Which are not?
 - What is missing?
- SHARE uses evaluations to determine best session awards. All speakers are volunteers, so this recognition is important.
- Evaluate sessions:
 - **SHARE.org/BostonEval**
 - **Online and via smart device**





The Survey!

Richard Lewis
IBM

Official Statistician and Triviaologist





z/VM Update

Bill Bitner
bitnerb@us.ibm.com
IBM



Agenda

- Release Status and Information
- z/VM Version 6 Release 3
- Hardware Support
- Futures and Statements of Direction

Release Status and Information

z/VM Release Status Summary

z/VM	Level	GA	End of Service	End of Marketing	Minimum Processor Level	Security Level
Ver 6	Rel 3	7 / 2013	4 / 2017		z10	EAL 4+ ^[2] OSPP-LS
	Rel 2	12 / 2011	4 / 2015	3Q / 2013	z10	-
	Rel 1	10 / 2009	4 / 2013	12 / 2011	z10	EAL 4+ OSPP-LS
Ver 5	Rel 4	9 / 2008	12 / 2014 ^[1]	3 / 2012	z800, z900	-
	Rel 3	6 / 2007	9 / 2010	9 / 2010	z800, z900	EAL 4+ CAPP/LSP

^[1] Or later (Announced August 7, 2012)

^[2] Targeted Security Level in V6.3 SOD

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

Extended support contracts are available.

z/VM 6.2 SSI and GDPS Support



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

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

z/VM Version 6 Release 3

z/VM Version 6 Release 3


Making Room to Grow Your Business



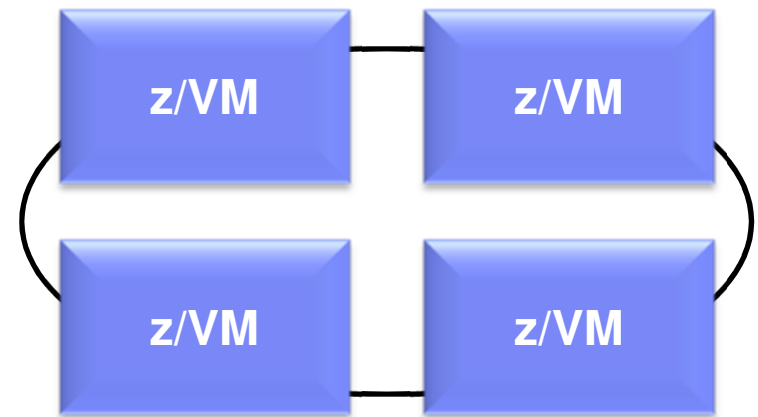
- Preview announcement: February 5, 2013
- Product announcement: July 23, 2013
- Product General Availability: July 26, 2013
- Continued Evolution:
 - Improved TCO – Reducing the number of z/VM systems you need to Manage
 - Improved Memory Management Flexibility and Efficiency
 - Simplify z/VM Systems Management
- See <http://www.vm.ibm.com/zvm630/>



z/VM 6.1




z/VM 6.2



z/VM 6.3



Reduce the number of z/VM systems you need to Manage

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

Improved Memory Management Flexibility and Efficiency

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



Simplify z/VM Systems Management

- Managing z/VM virtual servers with xCAT (Extreme Cloud Administration Toolkit) is ready to go after z/VM 6.3 installation; nothing else needs to be installed
- Adopt a foundation to allow future extensions for open source systems management solutions, in particular through OpenStack support
- Enable scalable support for the larger systems that z/VM 6.3 supports
- Safely migrate an existing z/VM V6R2 SSI Cluster to z/VM 6.3 in a step-wise approach, without having to shut down the cluster, using the new “Installation Upgrade In Place” capability



Large Memory Support

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

Large Memory Support

- Reorder processing removed
 - Commands remain, but have no impact
- 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 hardabend 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, hardabend dump, SNAPDUMP, DUMPLD2, and VM Dump Tool.
- Performance improvements for hardabend dump
 - Reduces time to take a CP hardabend 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

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

Example:

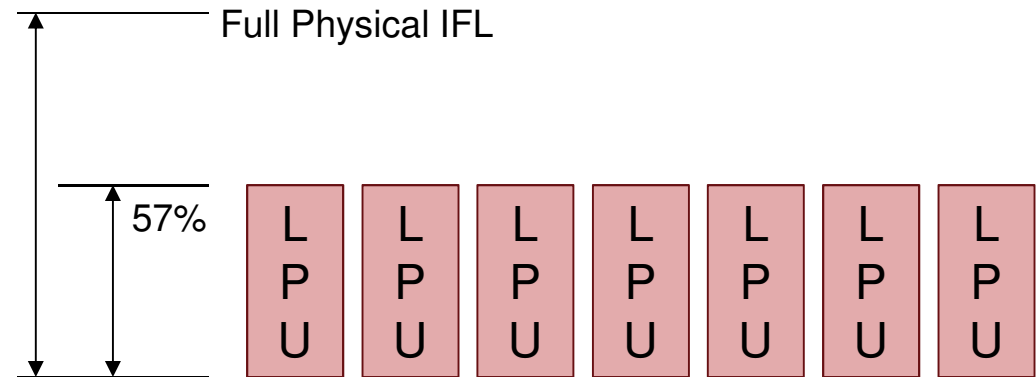
- 10 Physical IFLs, 7 logical IFLs, weight of 400 out of 1000
 - Each logical IFL (LPU) entitled to 57% of an IFL
- When CEC is constrained, the LPAR’s entitlement is reduced to 4 IFLs, so 7 is more than required
- z/VM & LPAR will cooperate
 - z/VM will concentrate the workload on a smaller number of logical processors
 - LPAR will redistribute the partition weight to give a greater portion to this smaller number of logical processors (~100% of 4 CPUs)

Horizontal vs. Vertical CPU Management




Horizontal:

- The logical processors are all created/treated equally.
- z/VM dispatches work evenly across the 7 logical processors



Vertical:

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



Technology Exploitation

- Fibre Channel Protocol Data Router Support
 - FCP QEBSM support enhanced for guest support use of FCP Data Router
- FICON DS8000 Series New Functions
 - Storage Controller Health message
 - New attention message from HW 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, & z/VM)
 - Remember to check for required service for systems that share the GDPS environment.

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

1 – with appropriate service – Check Bucket



Virtual Networking Improvements

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

Security Enhancements



- Crypto Express4S
 - Guest support for Crypto Express4S which is a feature available on zEC12 and zBC12
 - Can be configured in one of three ways:
 - IBM Common Cryptographic Architecture (CCA) Coprocessor mode
 - IBM CCA Accelerator mode
 - IBM Enterprise Public Key Cryptographic Standards (PKCS) #11 (EP11) coprocessor
- SSL Server Upgrade
 - System SSL update to z/OS V1.13 equivalency
 - 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

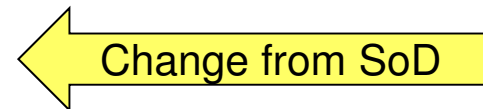
Linux Disk Dump Utility can now include the NSS

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

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



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



Hybrid computing model integrated and enabled for Cloud



Datawarehousing IDAA Solution

zManager for z/OS® and zBX


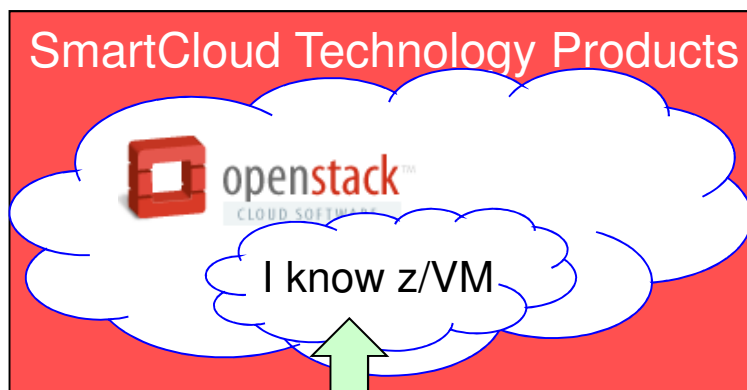
Systems Director for Power System x and storage

FSM for Intel® and Power ITEs


3rd party Managers and Servers



The OpenStack Food Chain



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



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

Other Considerations with z/VM 6.3



- You need to plan for Large Memory and for HiperDispatch. z/VM 6.3 changes some of the rules of thumb and planning guidelines from previous releases.
- DUMP Considerations
 - At time of publishing the calculations for dump space was not complete for the largest systems: <http://www.vm.ibm.com/service/zvmpladm.pdf>
- 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
- The integrated xCAT currently does not have support for the Command Line Interfaces
 - Will need to use the download version if CLI is required for now
 - Support available through service planned for December 2013

Sessions with z/VM 6.3 Content

- Monday
 - 11:00 13593 What's New in the z/VM 6.3 Hypervisor
 - 4:30 13564 z/VM Upgrade in Place Installation
- Tuesday
 - 9:30 13505 z/VM Performance Update
- Thursday
 - 9:30 13511 z/VM Early Support Program Experiences
- Friday
 - 8:00 13508 z/VM Virtual Switch – Advancing the Art
 - 9:30 13512 z/VM System Limits

Hardware Support

Support for IBM zEnterprise EC12



- Updates for z/VM 6.2, 6.1, and 5.4
 - VM65007 CP
 - VM65131 IOCP
 - VM65046 Performance Toolkit
 - 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**



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
- Update for z/VM 6.2 and in base of z/VM 6.3
 - PM83966: TCP/IP support
- PSP Bucket
 - Upgrade **2828DEVICE**
 - Subset **2828/ZVM**



z/VM Storage Support



- z/VM 6.3 and 6.2 supports
 - DS8000 Series (FCP or FICON)
 - DS6000 Series (FICON)
 - XIV (FCP)
 - IBM San Volume Controller (FCP)
 - IBM FlashSystem when behind an SVC (FCP)
 - IBM Storwize V7000 (FCP)
 - See http://www.ibm.com/support/docview.wss?uid=ssg1S1003703#_zvm
 - As well as many of the older storage devices
- The System Storage Interoperation Center (SSIC) support page:
 - <http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>
- The z/VM 6.2 General Information Manual has additional information, but had not been updated for Storwize, see URL above for requirements.

Statements of Direction

July 23, 2013

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

Security Evaluation of z/VM 6.3



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+).

- We continue the practice of taking every other release through certification.
- Evaluation is with inclusion of RACF Security Server optional feature.
- See <http://www.vm.ibm.com/security/> for current z/VM Security information.

FIPS Certification of z/VM 6.3

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^(TM)
 - 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.

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

Support of the 10GbE RoCE Express Feature

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

Support of the zEDC Express Feature

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

Stabilization of z/VM 5.4 Support

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

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

Withdrawal of Support for Expanded Storage

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

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

Summary

Leadership

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

Innovation

z/VM 6.2 introduced horizontal scalability and guest mobility through Single System Image clustering and Live Guest Relocation with RAS in the forefront of the design. z/VM 6.3 protects that investment with upgrade in place.

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.

Have a great week!



SHARE.org/BostonEval

Trademarks



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

IBM*	System z10*	System z196
IBM Logo*	Tivoli*	System z114
DB2*	z10 BC	System zEC12
Dynamic Infrastructure*	z9*	System zBC12
GDPS*	z/OS*	FlashSystem
HiperSockets	z/VM*	FICON*
Parallel Sysplex*	z/VSE	HyperSwap*
RACF*	zEnterprise*	DS8000*
System z*		

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

OpenSolaris, Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.
INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
OpenStack is a trademark of OpenStack Foundation

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

Notes:

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

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

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

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

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

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

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

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.