

z/VM Platform Update

Session 12488



Alan Altmark

Senior Managing z/VM and Linux Consultant
IBM STG Lab Services and Training

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*	
GDPS*	z/OS*	
HiperSockets	z/VM*	
Parallel Sysplex*	z/VSE	
RACF*	zEnterprise*	
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.

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.

Acknowledgments – Platform Update Team

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


Agenda

- Pricing review
- Status update for older releases
- z/VM Version 6 Release 2
- Statements of Direction
- z/VM Version 6 Release 3 Preview

z/VM Pricing

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

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



z/VM Version 5 Release 4

Update

- The last release of z/VM to support IBM System z9 and older processors
 - **No longer available as of March 12, 2012**
 - Also supports the zEC12
- End of Service has been extended to **December 31, 2014** or end of IBM service for System z9, whichever is *later*
 - Was September 30, 2013
 - Later, but not *too* much later!
 - Be on the lookout for Delayed Onset Panic Syndrome



z/VM Release Status Summary

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

^[1] Currently in evaluation

^[2] Or later (Announced August 7, 2012)

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

Extended support contracts are available.

z/VM Version 6 Security Certification Plans



- Common Criteria (ISO/IEC 15408)
 - Statement of Direction issued July 22, 2010
 - **Evaluation is complete (BSI-DSZ-CC-0752)**
 - Security Target: Operating System Protection Profile (OSPP) at EAL 4+
 - Virtualization extension
 - Labeled Security extension

- Federal Information Protection Standard (FIPS) 140-2
 - z/VM 6.1 System SSL is FIPS 140-2 Validated^(TM)
 - Enablement requirements for certificate database and servers
 - <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2012.htm#1735>

- z/VM 6.2 is designed to conform to both Common Criteria and FIPS 140-2 evaluation requirements



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

z/VM Version 6 Release 2

- Generally available December 2, 2011
- Will be withdrawn from marketing 3Q 2013
 - Concurrent with z/VM V6.3 GA
- Major changes include:
 - Single System Image
 - Live Guest Relocation
 - Turnkey support for Unified Resource Manager
- Replaced z/VM V6.1
 - No longer available
 - Ends service Real Soon!



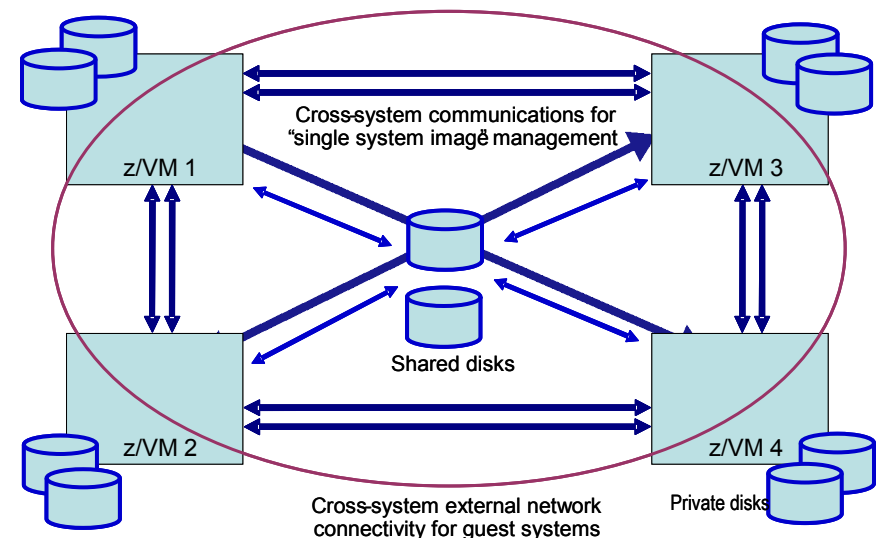
**z/VM V6.2
End of Service
April 30, 2015**

**z/VM V6.1
End of Service
April 30, 2013**

Single System Image (SSI) Feature

Clustered Hypervisor with Live Guest Relocation

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




SSI Cluster Management – Features for Greater Reliability

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

Single System Image Feature Clustered Hypervisor with Live Guest Relocation

- Dynamically move Linux guests from one member to another with Live Guest Relocation
 - Reduce planned outages
 - Enhance workload management
 - Non-disruptively move work to available system resources and non-disruptively move system resources to work
- When combined with Capacity Upgrade on Demand, Capacity Backup on Demand, and Dynamic Memory Upgrade, you will get the best of both worlds



Bring additional resources to the workload!

Move the workload to the resources!



Single System Image Feature Imbedded Relocation Wizard

- Eligibility checks done multiple times throughout the relocation process.
- Check more than just eligibility to move the virtual machine, but also check if it is “safe” to move.
 - Overrides are available
- Checks for:
 - Does virtual machine really have access to all the same resources and functions?
 - Will moving the virtual machine over commit resources to the point of jeopardizing other workload on the destination system?
- Pacing logic to minimize impact to other work in more memory constrained environments

Single System Image Feature



- xCAT 2.8 intends to support SSI and LGR
 - Already supports z/VM today
- Unified Resource Manager (zManager) *does not* support SSI and LGR
- IBM Systems Director *does not* support SSI and LGR
- Suggested best practice is to not combine SSI and LGR with zManager or Systems Director
 - Work with your IBM Sales Team, IBM Lab Services, or z/VM Development Lab to determine which technologies are most critical to your environment and business.

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

IBM System z Lab Services Offering:

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

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

z/VM Storage Support

- z/VM 6.2 supports
 - DS8000 Series (FCP or FICON)
 - DS6000 Series (FICON)
 - XIV (FCP)
 - IBM San Volume Controller (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
- Note: The System Storage Interoperation Center (SSIC) support page has some omissions of the above support
 - <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.

z/VM System Management

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



**All support
z/VM 6.2!**

z/VM 6.2 SSI and GDPS Support

- All supported GDPS releases (3.7, 3.8, & 3.9) 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

- GDPS/PPRC xDR 3.9 supports SSI configuration
 - APAR PM64211
 - Requires z/VM APAR VM65176

- No support for Live Guest Relocation of monitored Linux guests or z/VM services

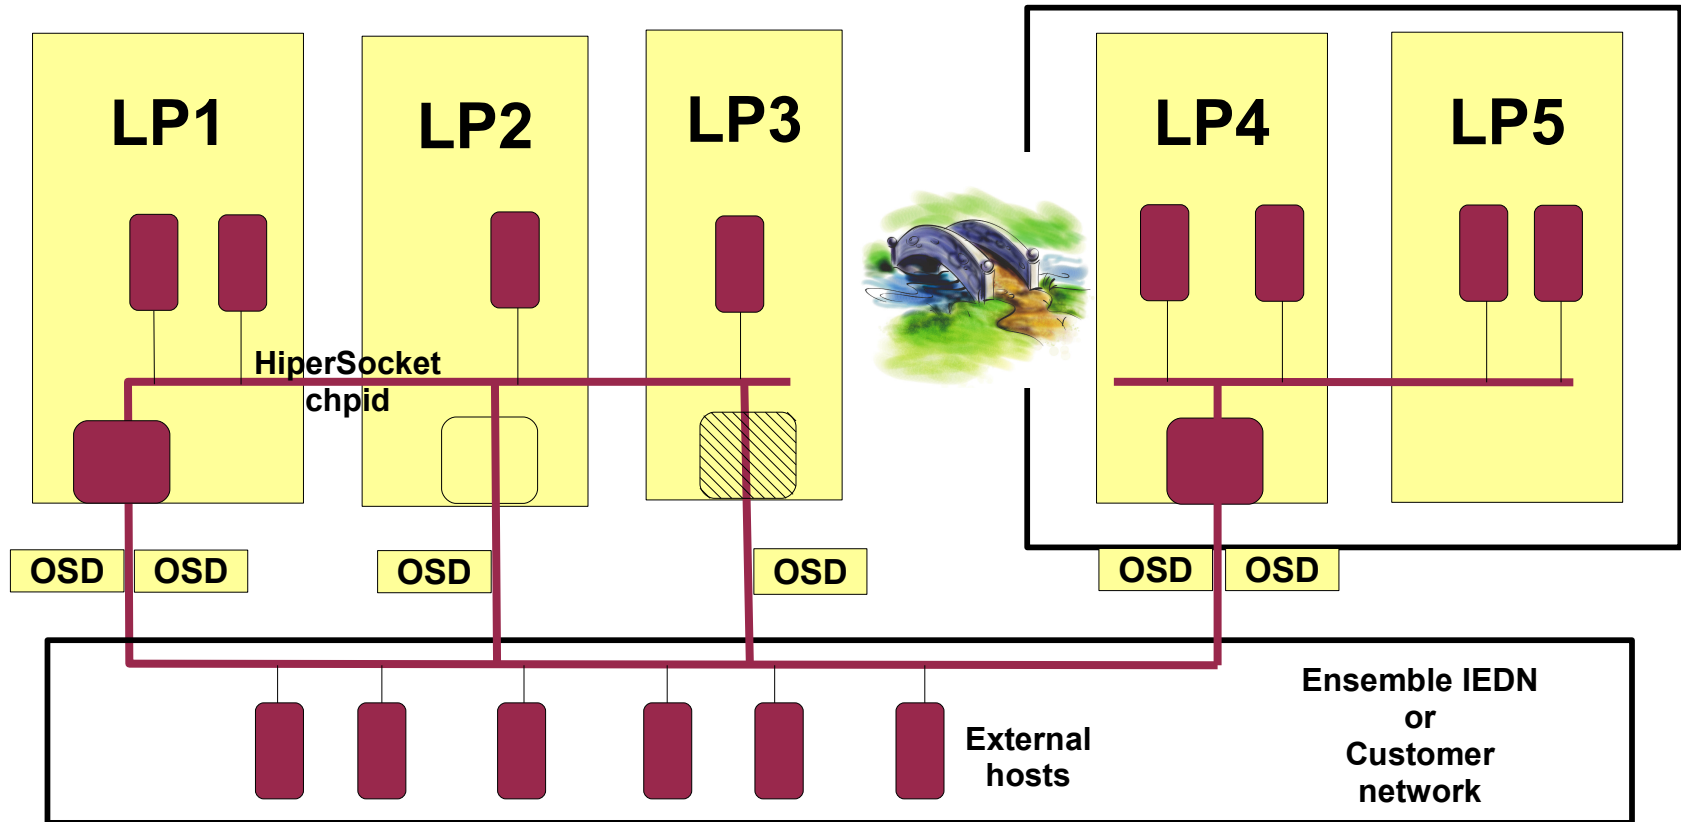


Support for IBM zEnterprise EC12

- 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** for ICSF service to support EP11 when running as a guest

HiperSocket VSWITCH Bridge

Available: April 13, 2012



- Built-in failover and failback
- CHPARM=x2 to bridge IEDN
- CHPARM=x4 to bridge customer network

- Same or different LPAR
- One active bridge per CEC
- PMTU simulation

HiperSocket VSWITCH Bridge

Available: April 13, 2012

- Virtual Switch bridge between Ethernet LAN and HiperSockets
 - zEnterprise IEDN (OSX) or customer network (OSD) connections
 - Guests can use simulated OSA or dedicated HiperSockets
 - VLAN aware
 - One HiperSocket chpid only
- Full redundancy
 - Up to 5 bridges per CEC
 - One bridge per LPAR
 - Automatic takeover
 - Optionally designate one “primary”
 - Primary will perform “takeback” when it comes up
 - Each bridge can have more than one OSA uplink
- CP: VM65042 / UM33691, TCP/IP: PM46988 / UK77220

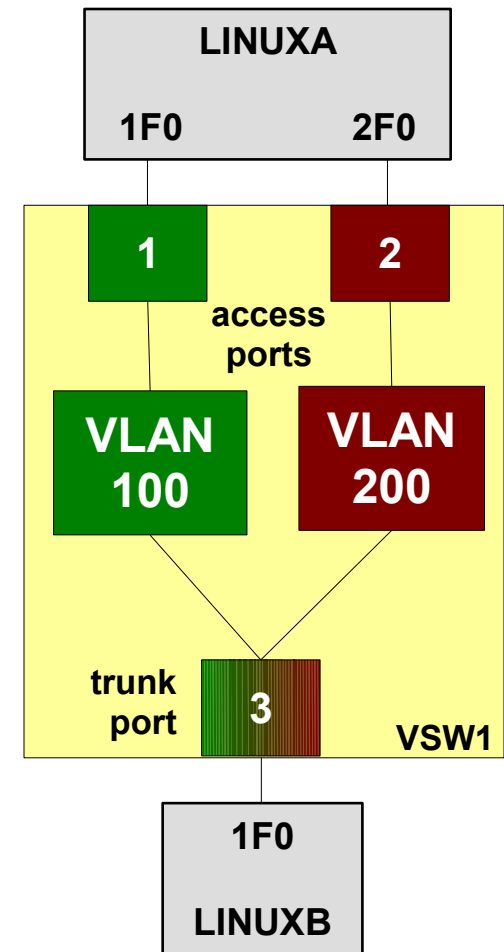
High Performance FICON

Available: April 13, 2012

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

VSWITCH: Multiple virtual access ports per guest

- New PORTBASED attribute for the VSWITCH
- VLAN and port management similar to physical switches
- One or more virtual ports on a the VSWITCH are reserved for a guest
- Each port is associated with a single VLAN
 - Each VLAN on the vswitch contains a list of associated ports
 - VLAN change takes effect immediately
- GRANT is not required
- ESM authorization requirement unchanged
 - Must detach and re-COUPLE the virtual NIC



Scalability and Performance Enhancements

- Reduction in the amount of memory and CPU required in order to manage larger memory sizes
- Integration of prior performance-related PTFs
 - SET REORDER
 - Reduce overhead of guest page release
 - Coalesce adjacent page frames
 - More accurate LIMITHARD
 - Reduce incidence of DIAGNOSE 0x9c and 0x44
 - Eliminate unexpected pauses in busy systems with large CPU over-commitment

z/VM Design for Performance Looking for Volunteers!

- The CPU Measurement Facility is a System z hardware facility that characterizes the performance of the CPU and nest:
 - Instructions, cycles, cache misses, and other processor related information
 - Available on System z10 and later
 - Need VM64961 (see Additional Materials)
- IBM will be using data from this facility to influence future processor design and benchmark validation of those designs.
- Increases the accuracy of future processor capacity sizing tools
- You can assist by providing sample Monwrite data containing the counters
- Contact Richard Lewis (rflewis@us.ibm.com) if you'd like to participate

TCP/IP Enhancements

- Stack
 - RFC 4191: Router selection preferences
 - RFC 5175: IPv6 router advertisement flags extension

- FTP
 - IPv6
 - Passwords suppressed in server traces
 - Wildcards supported for BFS files

- SMTP
 - IPv6
 - Includes IPv6 support in CMS NOTE and SENDFILE

TCP/IP Enhancements

OSA Diagnostics

- The NETSTAT command has been updated to provide details taken from the OSA Address Table (OAT) via new OSAINFO option.
- OSA/SF no longer required to obtain device details
- OSA-Express3 and later

```

VM TCP/IP Netstat Level 620          TCP/IP Server Name: TCPIP

Device K4L3VSW6640DEV: data as of 09/23/11 01:05:21
  OSA Generation:                    OSA-Express3
  OSA Firmware Level:                00000766
  Port Speed/Mode:                   1000 Mbs / Full Duplex
  Port Media Type:                   Multi Mode (SR/SX)
  PCHID:                             0291
  CHPID:                             0053
  Manufacturer MAC Address:          00-14-5E-78-17-F2
  Configured MAC Address:            00-00-00-00-00-00
  Data Device Sub-Channel Address:    6640
  CULA:                              00
  Unit Address:                      40
  Physical Port Number:              0
  Number Of Output Queues:            1
  Number Of Input Queues:             1
  Number Of Active Input Queues:      0
  QDIO CHPID Type:                   OSD
  QDIO Connection:                   Not Isolated
  IPv4 L3 VMAC:                      00-00-00-00-00-00
  IPv4 VMAC Router Mode:              No
  IPv4 L3 VMAC Active:                No
  IPv4 L3 VMAC Source:                n/a
  IPv4 L3 Global VLAN ID Active:      No
  IPv4 Global VLAN ID:                0
  IPv4 Assists Enabled:               00001C71
  IPv4 Outbound Checksum:             00000000
  IPv4 Inbound Checksum:              00000000

  IPv4 Address:                      IPA Flags:
  -----
  9.60.29.53                         00000002

  IPv4 Multicast Address:             MAC Address:
  -----
  224.0.0.1                          01-00-5E-00-00-01

```

Access controls for dedicated or attached devices

- The CP ATTACH and GIVE commands, as well as the DEDICATE statements in the directory will now engage ESM access controls
- Integrated ASCII console on the HMC is also managed
- Full discretionary and mandatory access controls
- RACF support included

Mandatory access controls for virtual consoles

- SET SECUSER and SET OBSERVER are now available when mandatory access controls (security labels) are active.
- Virtual security zones (“color coding” of users and resources) can now co-exist with system automation functions.
- Also applies to the user ID specified on CONSOLE directory statement.
- Users in different zones cannot see or manage each others' virtual console
 - Console cannot be given
 - Console cannot be taken
 - System administrators and automation solutions can use label SYSNONE to allow them access to all consoles

RACF Security Server

- Single System Image Support
 - Automatic propagation of most RACF commands
 - Also works with multiple RACF servers on same z/VM system
- Protected Users
 - User without a password or password phrase will not be revoked due to too many invalid password attempts or inactivity
- High Level Assembler no longer required for most common customizations
- Real device protection
 - ATTACH, GIVE, DEDICATE
 - New VMDEV class
 - Profiles: RDEV.*device.system_id*
- Support for Diagnose 0xA0 Subcode 0x48
 - Obtain information about any ESM in architected format

z/OS R12 Equivalency Upgrades

- LDAP
 - Change logging of general resources
 - Password expiry management
- Language Environment (LE) runtime libraries
- Program Management Binder
 - COMPAT supports ZOSV1R10, ZOSV1R11, ZOSV1R12
 - New suboptions on RMODE
 - Compiler parameters can be read from IEWPARMS DDNAME
 - New C/C++ API
- Support for **IBM XL C/C++ Compiler for z/VM, V1.3** (5654-A22)
 - Details can be found in announcement letter 211-369
- MPRROUTE

z/CMS

- Previously shipped with z/VM as a sample program, now supported
 - IPL ZCMS or MAINT 990
- Enables CMS programs to use z/Architecture instructions and 64-bit registers
- Existing ESA/390 architecture programs continue to run unchanged
 - CMS does not exploit memory above 2 GB
 - CMS *does* provide basic memory management API for memory above 2 GB
- Programs that examine or change architecture-sensitive memory locations (NUCON) must be updated in order to use z/CMS
- No z/Architecture extensions for DAT-off virtual machines to use access registers
 - VM Data Spaces not available

Installation Improvements

- Significant changes to system layout to support Single System Image
- Choose a non-SSI system or a complete 1- to 4-member SSI cluster
 - First or second level
- All installation information is gathered at one time
- All DASD volumes can be labeled at installation time, including the system residence volume
- Turnkey support for zEnterprise ensembles enables clients new to z/VM to easily get started with Unified Resource Manager (zManager)
 - If you have purchased a directory and/or security manager, decline this option during installation; manual enablement is required!

XEDIT – Default changed to mixed case

- For those coming to z/VM from an open system background, the folding of mixed case to upper case is surprising
- Many comments along the lines of “it hurts when you do that”
 - Linux can read CMS files
 - Often case-sensitive
- Default for other file types changed to CASE MIXED RESPECT
 - No folding
 - To get old behavior, update PROFILE XEDIT to
SET CASE UPPER RESPECT
 - Or you may wish to consider
SET CASE MIXED IGNORE

Removed Functions

- Kerberos authentication system
 - IBM Software Announcement 208-249

- CMS-based Domain Name Server (NAMESRV)
 - IBM Software Announcement 209-207

- RESOURCE option of VMSES/E VMFINS command
 - IBM Software Announcement 210-234

- z/VM Manageability Access Point (zMAP) agent and Platform agent for IBM Systems Director for Linux on System z
 - Both previously shipped with z/VM V6.1

Statements of Direction

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

Note for withdrawals: Unless otherwise stated, it is IBM's intent that z/VM V6.2 will be the last release of z/VM to support the indicated function.

HiperSockets Completion Queues z/VM Statement of Direction: New function

Available April 13, 2012

- Transfer HiperSockets messages asynchronously
- Used whenever traditional synchronous queues are full
- Automatic enablement; no z/VM configuration required
- Helpful when traffic is “bursty”
- Exploitation by CP VSWITCH only; no guest simulation

z/VM Performance Toolkit: RMFPMS agent

z/VM Statement of Direction: Stabilize existing function

- Performance Toolkit processing of the output from Linux rmfpms agent, part of the z/OS RMF PM offering, will no longer be updated
- Performance Toolkit may give incorrect results as the underlying rmfpms agent evolves
- Support for the Linux rmfpms agent has already been withdrawn, but continues to be available on an as-is basis

HMC non-ensemble z/VM System Management

z/VM Statement of Direction: Withdrawal of existing function

- z/VM V6.2 is the last release of z/VM that will be supported by the non-ensemble z/VM System Management functions of the System z10 and later
- IBM intends that z/VM V6 virtual server management will continue to be supported using the zEnterprise Unified Resource Manager on the System z196 and later.

TCP/IP Devices and Daemons

z/VM Statement of Direction: Withdrawal

- A220 HYPERchannel devices
- CLAW devices
- DHCP daemon
- LPSERVE (LPD)
 - RSCS LPD is provided at no charge
 - Does not affect LPR (client)

OVERRIDE utility

z/VM Statement of Direction: Withdrawal

- OVERRIDE utility is a “compiler” used to enable you to modify the privilege classes associated with CP commands and DIAGNOSE subcodes.
- Creates special spool files spool files with type UCR (“User Class Restructure”)
- Introduced in VM/SP Release 6
- Replaced in VM/ESA Release 2 by CP MODIFY
 - MODIFY COMMAND
 - MODIFY DIAGNOSE
 - CP commands or SYSTEM CONFIG statements

Cross System Extensions (CSE)

z/VM Statement of Direction: Withdrawal

- 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



Support for GDPS/PPRC

z/VM Statement of Direction: New function

- Disk subsystem preemptive HyperSwap
 - Storage controllers will notify host when failure is predicted
 - HyperSwap before I/O errors are generated

- HyperSwap scalability
 - Summary “PPRC Suspend” event notification by storage controller
 - Avoid separate notification for each disk

- Future z/VM release support for an alternate subchannel set in which to place PPRC secondary devices


z/VM Version 6 Release 3 Preview

Previews provide insight into IBM plans and direction. Availability, prices, ordering information, and terms and conditions will be provided when the product is announced.

z/VM Version 6 Release 3 Preview

Making Room to Grow Your Business

- Preview announcement on February 5, 2013
- Planned Availability: 3rd Quarter 2013
- Major Enhancements for Scalability and Performance
 - Support for larger amounts of real memory
 - Increased processor efficiency



z/VM Version 6 Release 3 Preview

Large Memory Support

- Real memory limit raised from 256GB to **1TB**
 - Proportionately increases total virtual memory based on tolerable over- commitment levels and workload dependencies
- Virtual machine memory limit remains unchanged at **1TB**
- 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 **128GB**

z/VM Version 6 Release 3 Preview

Large Memory Support

- Reorder processing removed
- 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
- Standalone dump to disk

z/VM Version 6 Release 3 Preview


HiperDispatch

- Improved processor efficiency
 - Better n-way curves
 - Supported processor limit of 32 remains unchanged
 - Better use of processor cache to take advantage of cache-rich system design of more recent machines
- Two components:
 - Dispatching Affinity
 - Vertical CPU Management

z/VM Version 6 Release 3 Preview

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



z/VM Version 6 Release 3 Preview

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
 - Keep virtual CPUs of the same virtual machine near one another
- Better use of cache can reduce the execution time of a set of related instructions
- z/VM V6.2 and earlier uses “soft” affinity to dispatch virtual CPUs
 - No awareness of chip or book

z/VM Version 6 Release 3 Preview

HiperDispatch: Vertical CPU Management

- Today's “horizontal” management distributes the LPAR weight evenly distributed 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
- 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)

z/VM Version 6 Release 3 Preview

Items of Interest

- Virtual Switch updates
- Installation enhancements

For more details....

- z/VM Single System Image
John Franciscovich
Imperial B
Monday 15:00
- z/VM Performance Update
Bill Bitner
Imperial B
Tuesday 9:30
- z/VM Virtual Switch
Alan Altmark
Franciscan D
Thursday 8:00
- z/VM Single System Image labs
Richard Lewis & Friends
Union Square 23-24
Tuesday afternoon starting at 1:30

Thanks!



Contact Information:

Alan Altmark
Alan_Altmark@us.ibm.com



Session 12488

Additional Materials

The following are provided for your information

Previously shipped Functional Enhancements Included in z/VM V6.2

- XRC timestamps
- Hyperswap improvements
- SSL Server Reliability and Scalability
- CPU Measurement Counter Facility Host support
- zEnterprise Unified Resource Manager

APAR numbers shown apply to z/VM 6.1 and z/VM 5.4 unless otherwise stated

XRC Timestamps

VM64814 and VM64816

- CP will sync with STP at IPL and, optionally, obtain time zone and leap seconds from STP
 - No need to deactivate/activate LPAR

- Correct time will be placed in all host and guest I/O
 - CP will monitor STP time signals

- Enabled via SYSTEM CONFIG with option to skip timestamp or delay I/O if CP is unable to sync with STP

- No virtualization of STP
 - Option for 2nd level systems to stamp I/O without use of STP

Hyperswap Improvements

VM64815 and VM64816

- CP HYPERSWAP command now has additional controls for missing interrupt handling
 - Do not trigger automatic quiesce (default)
 - GDPS will not be notified
 - Trigger automatic quiesce after specified number of MI detection intervals
 - GDPS will be notified

- Better management of PAV and HyperPAV devices

- Avoid unnecessary hyperswaps due to normal maintenance activities
 - Concurrent storage controller upgrade

- New wait state 9060 if abend occurs when Hyperswap is in progress
 - no checkpoint taken, no automatic dump
 - restart dump if dedicated dump volume, else standalone dump

SSL Server Reliability and Scalability

PK97437, PK97438, PK75662

- Major rewrite of SSL server
 - Updates to TCP/IP stack, as well
 - Scalable
- Multiple SSL servers with session cache manager and shared database
 - Balance total number of sessions against number of sessions per server
- Significant performance improvements
 - Interactive workloads such as telnet
 - Session establishment costs, particularly during mass reconnect
- Migration required if using pre-PTF version
 - <http://www.vm.ibm.com/related/tcpip/tcsslspe.html>
 - Not trivial; read carefully

CPU Measurement Facility Counters – Host Support

VM64961

- Sets of counters for each logical processor that count events such as cycle, instruction, and cache directory-write counts
 - Same COUNTER information as z/OS partitions
- Accumulation is a relatively low-overhead activity and is performed automatically by the machine when the counters are authorized, enabled, and activated
- Authorization controlled by a logical partition's Security settings in its activation profile
- Enablement, activation, and data collection controlled by z/VM MONITOR command

zEnterprise Unified Resource Manager

VM64822, VM64904, VM64917, VM64956, VM64957

- z/VM V6 only
 - Turn-key installation option to enable virtual server management via zEnterprise Unified Resource Manager (z/VM V6.2 only)
 - Only for “kicking the tires”
- Enables Unified Resource Manager to perform system and virtual server management tasks
 - Virtual server configuration
 - Disk storage management
 - Virtual network management
 - Performance monitoring
- CP, CMS, LE, TCP/IP, DIRMAINT, Performance Toolkit, HCD
- <http://www.vm.ibm.com/service/vmrequrm.html>

zEnterprise Unified Resource Manager Ensemble Membership

- If configured to participate in an ensemble, z/VM will automatically join the ensemble at IPL
- Configuration tasks
 - Set up OSM and OSX channel paths
 - Set up controllers for IEDN and INMN networks
 - Pre-defined controllers DTCENS1 and DTCENS2 for exclusive use by ensemble networks
 - DTCENS1 automatically creates a VSWITCH to provide SMAPI connectivity to INMN network
 - Configure directory manager (REQUIRED)
 - Configure SMAPI servers
- See chapter "Configuring z/VM for an Ensemble" in CP Planning and Administration manual

Scalability and Performance Enhancements

Available by PTF to prior releases as shown

- Control of the guest page re-ordering process, improving the performance characteristics of guests with large memory footprints (VM64774)
- Reduced system overhead of guest page release function, thereby helping to increase guest throughput (VM64715)
- Improved contiguous frame coalescing algorithms help to increase system throughput (VM64795)

Scalability and Performance Enhancements

Available by PTF to prior releases where shown

- More accurate scheduling algorithm for guests that have LIMITHARD shares (VM64721)
- Reduce LPAR suspend time by reducing the number of DIAGNOSE 0x9C and 0x44 instructions issued when obtaining system locks (VM64927 for z/VM 6.1 only)
- Eliminate erratic virtual machine pause in busy systems with more than 14:1 total virtual-to-logical CPU over-commitment (VM64887)



z/VM Platform Update

Session 12488



Alan Altmark
Senior Managing z/VM and Linux Consultant
IBM STG Lab Services and Training



Version 19b



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*	
GDPS*	z/OS*	
HiperSockets	z/VM*	
Parallel Sysplex*	z/VSE	
RACF*	zEnterprise*	
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.

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.

Acknowledgments – Platform Update Team

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


Agenda

- Pricing review
- Status update for older releases
- z/VM Version 6 Release 2
- Statements of Direction
- z/VM Version 6 Release 3 Preview

z/VM Pricing

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

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





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

- End of Service has been extended to **December 31, 2014** or end of IBM service for System z9, whichever is **later**
 - Was September 30, 2013
 - Later, but not *too* much later!
 - Be on the lookout for Delayed Onset Panic Syndrome



5.3 is out of Service!

Extended support contracts

5.4 will stay in service as long as the z9 stays in service

- May be extended beyond current date

For those not moving off of z/VM 5.4 due to running on older hardware, the new z114 offers some good alternatives. A lot of the z9 configurations, even up to around 6 IFLs of the Enterprise Class machines will fit on a z114. Bill Bitner can provide more info on moving up to newer machines.

z/VM Release Status Summary

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

^[1] Currently in evaluation

^[2] Or later (Announced August 7, 2012)

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

Extended support contracts are available.

z/VM Version 6 Security Certification Plans

- Common Criteria (ISO/IEC 15408)
 - Statement of Direction issued July 22, 2010
 - **Evaluation is complete (BSI-DSZ-CC-0752)**
 - Security Target: Operating System Protection Profile (OSPP) at EAL 4+
 - Virtualization extension
 - Labeled Security extension
- Federal Information Protection Standard (FIPS) 140-2
 - z/VM 6.1 System SSL is FIPS 140-2 Validated^(TM)
 - Enablement requirements for certificate database and servers
 - <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2012.htm#1735>
- z/VM 6.2 is designed to conform to both Common Criteria and FIPS 140-2 evaluation requirements



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

The last Common Criteria Certification was z/VM Version 5 Release 3, which certified to the CAPP and LSPP Protection Profiles at EAL 4+. Those Protection Profiles have since expired – hence the new PP being used for z/VM Version 6 Release 1.

The Target of Evaluation for Common Criteria includes RACF. IBM makes no claims about the relative security of your configuration if you are using a different External Security Manager (ESM), because those products have not been evaluated.

FIPS APARs to note for z/VM 6.1:

- PM10616: System SSL enablement of FIPS
- PM43382: System SSL Self-Defense

z/VM Version 6 Release 2




- Generally available December 2, 2011
- Will be withdrawn from marketing 3Q 2013
 - Concurrent with z/VM V6.3 GA
- Major changes include:
 - Single System Image
 - Live Guest Relocation
 - Turnkey support for Unified Resource Manager
- Replaced z/VM V6.1
 - No longer available
 - Ends service Real Soon!

z/VM V6.2
End of Service
April 30, 2015

z/VM V6.1
End of Service
April 30, 2013

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

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



Talking about it earlier than we normally do for new z/VM function.

Cluster will be "normal" way to run z/VM.

Don't have to rely on workload to do maintenance. Can move virtual machines and do maintenance on your own schedule.

Move workload to another server and let it keep running.

- without triggering High Availability flags/alarms.

SSI Cluster Management – Features for Greater Reliability

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

Single System Image Feature Clustered Hypervisor with Live Guest Relocation

- Dynamically move Linux guests from one member to another with Live Guest Relocation
 - Reduce planned outages
 - Enhance workload management
 - Non-disruptively move work to available system resources and non-disruptively move system resources to work
- When combined with Capacity Upgrade on Demand, Capacity Backup on Demand, and Dynamic Memory Upgrade, you will get the best of both worlds



Bring additional resources to the workload!

Move the workload to the resources!



This is good slide to add pointer if given in a conference where the SSI Overview will be given as well.

Single System Image Feature Imbedded Relocation Wizard

- Eligibility checks done multiple times throughout the relocation process.
- Check more than just eligibility to move the virtual machine, but also check if it is “safe” to move.
 - Overrides are available
- Checks for:
 - Does virtual machine really have access to all the same resources and functions?
 - Will moving the virtual machine over commit resources to the point of jeopardizing other workload on the destination system?
- Pacing logic to minimize impact to other work in more memory constrained environments

Single System Image Feature



- xCAT 2.8 intends to support SSI and LGR
 - Already supports z/VM today
- Unified Resource Manager (zManager) *does not* support SSI and LGR
- IBM Systems Director *does not* support SSI and LGR
- Suggested best practice is to not combine SSI and LGR with zManager or Systems Director
 - Work with your IBM Sales Team, IBM Lab Services, or z/VM Development Lab to determine which technologies are most critical to your environment and business.

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

IBM System z Lab Services Offering:

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

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

z/VM Storage Support

- z/VM 6.2 supports
 - DS8000 Series (FCP or FICON)
 - DS6000 Series (FICON)
 - XIV (FCP)
 - IBM San Volume Controller (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
- Note: The System Storage Interoperation Center (SSIC) support page has some omissions of the above support
 - <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.

z/VM System Management

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



**All support
z/VM 6.2!**

z/VM 6.2 SSI and GDPS Support

- All supported GDPS releases (3.7, 3.8, & 3.9) 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
- GDPS/PPRC xDR 3.9 supports SSI configuration
 - APAR PM64211
 - Requires z/VM APAR VM65176
- No support for Live Guest Relocation of monitored Linux guests or z/VM services

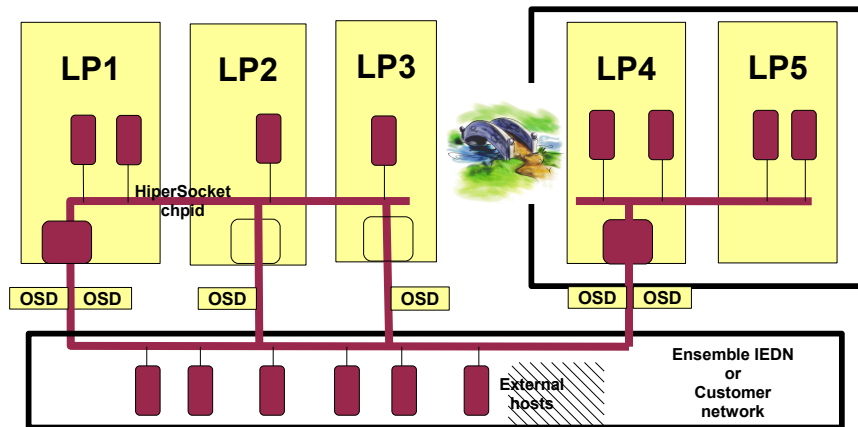
**Update**

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** for ICSF service to support EP11 when running as a guest

HiperSocket VSWITCH Bridge

Available: April 13, 2012



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

z/OS or other native OSes can talk to z/VM guests on same CEC, as usual, but HiperSockets will not forward any of their traffic to the bridge.

To perform similar function, native OS must provide its own “bridge” capability.

Will it support OSD? Not part of SOD. So, maybe.

CP generates Path MTU Discovery response packets automatically if frame size is too large for OSA.

- Allows large MTU for efficient local HiperSocket communications.

HiperSocket VSWITCH Bridge

Available: April 13, 2012

- Virtual Switch bridge between Ethernet LAN and HiperSockets
 - zEnterprise IEDN (OSX) or customer network (OSD) connections
 - Guests can use simulated OSA or dedicated HiperSockets
 - VLAN aware
 - One HiperSocket chpid only
- Full redundancy
 - Up to 5 bridges per CEC
 - One bridge per LPAR
 - Automatic takeover
 - Optionally designate one “primary”
 - Primary will perform “takeback” when it comes up
 - Each bridge can have more than one OSA uplink
- CP: VM65042 / UM33691, TCP/IP: PM46988 / UK77220

LPARs and z/VM Guests are separate entities. Only z/VM bridge can forward to an outside entity.

Provides a nice way to put HiperSocket connections on an IEDN or external subnet


High Performance FICON

Available: April 13, 2012

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

VSWITCH: Multiple virtual access ports per guest

- New PORTBASED attribute for the VSWITCH
- VLAN and port management similar to physical switches
- One or more virtual ports on the VSWITCH are reserved for a guest
 - Each VLAN on the vswitch contains a list of associated ports
 - VLAN change takes effect immediately
- GRANT is not required
- ESM authorization requirement unchanged
 - Must detach and re-COUPLE the virtual NIC



Prior to z/VM 6.2 connection to Virtual Switch ports was granted by z/VM userid.

Therefore if multiple virtual NICs were defined for a user for the same Vswitch then all the virtual NICs had the same attributes.

With z/VM 6.2, guests will be able to have multiple **unique** Access Ports which allows each virtual NIC to be in a different VLAN.

Vswitch Port numbers can be defined for connections that remain constant. Existing Vswitch definitions do not have to be changed when migrating to z/VM 6.2, the older non-port based Vswitch definitions remain valid.

Advantages:

Guest can have multiple unique access ports connected to a single virtual switch.

Allows a guest to access multiple VLANs without being VLAN aware. Each NIC uses an access port with a unique VLAN.

Customers can define port numbers for connections that remain static. That is disconnect and reconnect and port number remains the same. This will make SNMP monitoring of switch status more useful to customers, since they know the mapping of systems to switch ports ahead of time.

Scalability and Performance Enhancements

- Reduction in the amount of memory and CPU required in order to manage larger memory sizes
- Integration of prior performance-related PTFs
 - SET REORDER
 - Reduce overhead of guest page release
 - Coalesce adjacent page frames
 - More accurate LIMITHARD
 - Reduce incidence of DIAGNOSE 0x9c and 0x44
 - Eliminate unexpected pauses in busy systems with large CPU over-commitment

z/VM Design for Performance Looking for Volunteers!

- The CPU Measurement Facility is a System z hardware facility that characterizes the performance of the CPU and nest:
 - Instructions, cycles, cache misses, and other processor related information
 - Available on System z10 and later
 - Need VM64961 (see Additional Materials)
- IBM will be using data from this facility to influence future processor design and benchmark validation of those designs.
- Increases the accuracy of future processor capacity sizing tools
- You can assist by providing sample Monwrite data containing the counters
- Contact Richard Lewis (rflewis@us.ibm.com) if you'd like to participate

TCP/IP Enhancements

- Stack
 - RFC 4191: Router selection preferences
 - RFC 5175: IPv6 router advertisement flags extension
- FTP
 - IPv6
 - Passwords suppressed in server traces
 - Wildcards supported for BFS files
- SMTP
 - IPv6
 - Includes IPv6 support in CMS NOTE and SENDFILE

“Ipv6: The Journey Continues...”

TCP/IP Enhancements OSA Diagnostics

- The NETSTAT command has been updated to provide details taken from the OSA Address Table (OAT) via new OSAINFO option.
- OSA/SF no longer required to obtain device details
- OSA-Express3 and later

```

VM TCP/IP Netstat Level 620      TCP/IP Server Name: TCPIP

Device K4L3VSW6640DEV: data as of 09/23/11 01:05:21
OSA Generation:                  OSA-Express3
OSA Firmware Level:              00000766
Port Speed/Mode:                 1000 Mbs / Full Duplex
Port Media Type:                 Multi Mode (SR/SX)
PCPID:                           0291
CHPID:                           0053
Manufacturer MAC Address:        00-14-5E-78-17-F2
Configured MAC Address:          00-00-00-00-00-00
Data Device Sub-Channel Address: 6640
CULA:                            00
Unit Address:                    40
Physical Port Number:             0
Number Of Output Queues:         1
Number Of Input Queues:          1
Number Of Active Input Queues:   0
QDIO CHPID Type:                 OSD
QDIO Connection:                 Not Isolated
IPv4 L3 VMAC:                    00-00-00-00-00-00
IPv4 VMAC Router Mode:           No
IPv4 L3 VMAC Active:             No
IPv4 L3 VMAC Source:             n/a
IPv4 L3 Global VLAN ID Active:   No
IPv4 Global VLAN ID:             0
IPv4 Assists Enabled:            00001C71
IPv4 Outbound Checksum:          00000000
IPv4 Inbound Checksum:           00000000

IPv4 Address:                    IPA Flags:
9.60.29.53                      00000002

IPv4 Multicast Address:          MAC Address:
224.0.0.1                      01-00-5E-00-00-01

```

OAT contains lots of information about the OSA, including registered IP and MAC addresses.

Access controls for dedicated or attached devices

- The CP ATTACH and GIVE commands, as well as the DEDICATE statements in the directory will now engage ESM access controls
- Integrated ASCII console on the HMC is also managed
- Full discretionary and mandatory access controls
- RACF support included

Mandatory access controls for virtual consoles

- SET SECUSER and SET OBSERVER are now available when mandatory access controls (security labels) are active.
- Virtual security zones (“color coding” of users and resources) can now co-exist with system automation functions.
- Also applies to the user ID specified on CONSOLE directory statement.
- Users in different zones cannot see or manage each others' virtual console
 - Console cannot be given
 - Console cannot be taken
 - System administrators and automation solutions can use label SYSNONE to allow them access to all consoles

In prior releases, enabling SECLABELs caused SET SECUSER and SET OBSERVER to issue an error message.

RACF Security Server

- Single System Image Support
 - Automatic propagation of most RACF commands
 - Also works with multiple RACF servers on same z/VM system
- Protected Users
 - User without a password or password phrase will not be revoked due to too many invalid password attempts or inactivity
- High Level Assembler no longer required for most common customizations
- Real device protection
 - ATTACH, GIVE, DEDICATE
 - New VMDEV class
 - Profiles: *RDEV.device.system_id*
- Support for Diagnose 0xA0 Subcode 0x48
 - Obtain information about any ESM in architected format

The PROTECTED attribute – Protected user ids cannot be used for normal logon to the system and they will not be revoked through inactivity or invalid password attempts.

VMDEV is a new general resource class that includes qualification by system ID. This allows for authorization across multiple systems. When the same device number is used, a generic resource can be defined: *RDEV.0480.***

RDEV.SYSASCII.systemid is used to protect the integrated ASCII console.

z/OS R12 Equivalency Upgrades

- LDAP
 - Change logging of general resources
 - Password expiry management
- Language Environment (LE) runtime libraries
- Program Management Binder
 - COMPAT supports ZOSV1R10, ZOSV1R11, ZOSV1R12
 - New suboptions on RMODE
 - Compiler parameters can be read from IEWPARMS DDNAME
 - New C/C++ API
- Support for **IBM XL C/C++ Compiler for z/VM, V1.3** (5654-A22)
 - Details can be found in announcement letter 211-369
- MPROUTE

z/CMS

- Previously shipped with z/VM as a sample program, now supported
 - IPL ZCMS or MAINT 990
- Enables CMS programs to use z/Architecture instructions and 64-bit registers
- Existing ESA/390 architecture programs continue to run unchanged
 - CMS does not exploit memory above 2 GB
 - CMS *does* provide basic memory management API for memory above 2 GB
- Programs that examine or change architecture-sensitive memory locations (NUCON) must be updated in order to use z/CMS
- No z/Architecture extensions for DAT-off virtual machines to use access registers
 - VM Data Spaces not available

Installation Improvements

- Significant changes to system layout to support Single System Image
- Choose a non-SSI system or a complete 1- to 4-member SSI cluster
 - First or second level
- All installation information is gathered at one time
- All DASD volumes can be labeled at installation time, including the system residence volume
- Turnkey support for zEnterprise ensembles enables clients new to z/VM to easily get started with Unified Resource Manager (zManager)
 - If you have purchased a directory and/or security manager, decline this option during installation; manual enablement is required!

Turnkey: In this mode, a free version of DIRMAINT and Performance Toolkit are used. Customer cannot directly access DIRMAINT or PERFKIT services. Used to support zManager. Client can convert from turnkey to purchased version.

XEDIT – Default changed to mixed case

- For those coming to z/VM from an open system background, the folding of mixed case to upper case is surprising
- Many comments along the lines of “it hurts when you do that”
 - Linux can read CMS files
 - Often case-sensitive
- Default for other file types changed to CASE MIXED RESPECT
 - No folding
 - To get old behavior, update PROFILE XEDIT to
SET CASE UPPER RESPECT
 - Or you may wish to consider
SET CASE MIXED IGNORE

Removed Functions

- Kerberos authentication system
 - IBM Software Announcement 208-249

- CMS-based Domain Name Server (NAMESRV)
 - IBM Software Announcement 209-207

- RESOURCE option of VMSES/E VMFINS command
 - IBM Software Announcement 210-234

- z/VM Manageability Access Point (zMAP) agent and Platform agent for IBM Systems Director for Linux on System z
 - Both previously shipped with z/VM V6.1

Statements of Direction

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

Note for withdrawals: Unless otherwise stated, it is IBM's
intent that z/VM V6.2 will be the last release of z/VM to
support the indicated function.

HiperSockets Completion Queues z/VM Statement of Direction: New function

Available April 13, 2012

- Transfer HiperSockets messages asynchronously
- Used whenever traditional synchronous queues are full
- Automatic enablement; no z/VM configuration required
- Helpful when traffic is “bursty”
- Exploitation by CP VSWITCH only; no guest simulation

z/VM Performance Toolkit: RMFPMS agent

z/VM Statement of Direction: Stabilize existing function

- Performance Toolkit processing of the output from Linux rmfpms agent, part of the z/OS RMF PM offering, will no longer be updated
- Performance Toolkit may give incorrect results as the underlying rmfpms agent evolves
- Support for the Linux rmfpms agent has already been withdrawn, but continues to be available on an as-is basis

HMC non-ensemble z/VM System Management

z/VM Statement of Direction: Withdrawal of existing function

- z/VM V6.2 is the last release of z/VM that will be supported by the non-ensemble z/VM System Management functions of the System z10 and later
- IBM intends that z/VM V6 virtual server management will continue to be supported using the zEnterprise Unified Resource Manager on the System z196 and later.

The Announcement Letter refers to the “HMC z/VM Tower”. This is the “z/VM Virtual Machine Management” task group on the HMC.

TCP/IP Devices and Daemons

z/VM Statement of Direction: Withdrawal

- A220 HYPERchannel devices
- CLAW devices
- DHCP daemon
- LPSERVE (LPD)
 - RSCS LPD is provided at no charge
 - Does not affect LPR (client)

OVERRIDE utility

z/VM Statement of Direction: [Withdrawal](#)

- OVERRIDE utility is a “compiler” used to enable you to modify the privilege classes associated with CP commands and DIAGNOSE subcodes.
- Creates special spool files spool files with type UCR (“User Class Restructure”)
- Introduced in VM/SP Release 6
- Replaced in VM/ESA Release 2 by CP MODIFY
 - MODIFY COMMAND
 - MODIFY DIAGNOSE
 - CP commands or SYSTEM CONFIG statements

Cross System Extensions (CSE)

z/VM Statement of Direction: Withdrawal

- 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



Support for GDPS/PPRC

z/VM Statement of Direction: New function

- Disk subsystem preemptive HyperSwap
 - Storage controllers will notify host when failure is predicted
 - HyperSwap before I/O errors are generated
- HyperSwap scalability
 - Summary “PPRC Suspend” event notification by storage controller
 - Avoid separate notification for each disk
- Future z/VM release support for an alternate subchannel set in which to place PPRC secondary devices


If the I/O Updates part 1 & 2 are given at same conference, this is a good place for pointer to that.

z/VM Version 6 Release 3 Preview

Previews provide insight into IBM plans and direction. Availability, prices, ordering information, and terms and conditions will be provided when the product is announced.

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

- Preview announcement on February 5, 2013
- Planned Availability: 3rd Quarter 2013
- Major Enhancements for Scalability and Performance
 - Support for larger amounts of real memory
 - Increased processor efficiency



z/VM Version 6 Release 3 Preview
Large Memory Support

- Real memory limit raised from 256GB to **1TB**
 - Proportionately increases total virtual memory based on tolerable over- commitment levels and workload dependencies
- Virtual machine memory limit remains unchanged at **1TB**
- 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 **128GB**

z/VM Version 6 Release 3 Preview
Large Memory Support


- Reorder processing removed
- 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
- Standalone dump to disk

z/VM Version 6 Release 3 Preview
HiperDispatch

- Improved processor efficiency
 - Better n-way curves
 - Supported processor limit of 32 remains unchanged
 - Better use of processor cache to take advantage of cache-rich system design of more recent machines
- Two components:
 - Dispatching Affinity
 - Vertical CPU Management

z/VM Version 6 Release 3 Preview
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



z/VM Version 6 Release 3 Preview
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
 - Keep virtual CPUs of the same virtual machine near one another
- Better use of cache can reduce the execution time of a set of related instructions
- z/VM V6.2 and earlier uses “soft” affinity to dispatch virtual CPUs
 - No awareness of chip or book

z/VM Version 6 Release 3 Preview

HiperDispatch: Vertical CPU Management

- Today's "horizontal" management distributes the LPAR weight evenly distributed 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
- 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)


z/VM Version 6 Release 3 Preview Items of Interest

- Virtual Switch updates
- Installation enhancements

[For more details....](#)

- z/VM Single System Image
John Franciscovich
Imperial B
Monday 15:00
- z/VM Performance Update
Bill Bitner
Imperial B
Tuesday 9:30
- z/VM Virtual Switch
Alan Altmark
Franciscan D
Thursday 8:00
- z/VM Single System Image labs
Richard Lewis & Friends
Union Square 23-24
Tuesday afternoon starting at 1:30

Thanks!



Contact Information:

Alan Altmark
Alan_Altmark@us.ibm.com



Session 12488

Additional Materials

The following are provided for your information

Previously shipped Functional Enhancements Included in z/VM V6.2

- XRC timestamps
- Hyperswap improvements
- SSL Server Reliability and Scalability
- CPU Measurement Counter Facility Host support
- zEnterprise Unified Resource Manager

APAR numbers shown apply to z/VM 6.1 and z/VM 5.4 unless otherwise stated

XRC Timestamps

VM64814 and VM64816

- CP will sync with STP at IPL and, optionally, obtain time zone and leap seconds from STP
 - No need to deactivate/activate LPAR
- Correct time will be placed in all host and guest I/O
 - CP will monitor STP time signals
- Enabled via SYSTEM CONFIG with option to skip timestamp or delay I/O if CP is unable to sync with STP
- No virtualization of STP
 - Option for 2nd level systems to stamp I/O without use of STP

CP will sync with correct time from STP at IPL

STP is not virtualized to 2nd level guests

CP will track time changes (drift) – will be put in timestamps for I/O

Strictly for CP to timestamp I/O

– CP will not do STP synchronization for virtual machines

FEATURES statement

STP_Timestamping

tells CP to enable the STP protocol (if the STP facility is installed) and apply timestamps to all XRC-capable DASD devices.

STP_TimeZone / STP_TZ

tells CP to enable the STP protocol (if the STP facility is installed) and obtain timezone information automatically from the STP server.

XRC_OPTIONal

when STP_Timestamping is also enabled, this will allow non-timestamped I/O to be issued whenever STP is in an unsynchronized state, as opposed to deferring I/O until STP synchronization completes.

XRC_TEST

tells CP to timestamp I/O regardless of STP availability. This option is meant only for vendor and testing purposes, and can only be specified for systems running within a virtual machine.

Hyperswap Improvements

VM64815 and VM64816

- CP HYPERSWAP command now has additional controls for missing interrupt handling
 - Do not trigger automatic quiesce (default)
 - GDPS will not be notified
 - Trigger automatic quiesce after specified number of MI detection intervals
 - GDPS will be notified
- Better management of PAV and HyperPAV devices
- Avoid unnecessary hyperswaps due to normal maintenance activities
 - Concurrent storage controller upgrade
- New wait state 9060 if abend occurs when Hyperswap is in progress
 - no checkpoint taken, no automatic dump
 - restart dump if dedicated dump volume, else standalone dump

Trigger is by device.

SSL Server Reliability and Scalability

PK97437, PK97438, PK75662

- Major rewrite of SSL server
 - Updates to TCP/IP stack, as well
 - Scalable
- Multiple SSL servers with session cache manager and shared database
 - Balance total number of sessions against number of sessions per server
- Significant performance improvements
 - Interactive workloads such as telnet
 - Session establishment costs, particularly during mass reconnect
- Migration required if using pre-PTF version
 - <http://www.vm.ibm.com/related/tcpip/tcsslspe.html>
 - Not trivial; read carefully

- PK97437: SSLADMIN, TCPRUN and Related Packaging Changes
- PK97438: SSLSERV Module Updates
- PK75662: TCPIP Module Updates

SSL Server now CMS based

- Old server was not scaling (telnet sessions)

Network-free SSL server administration.

The SSL server can be managed without requiring a network connection between the SSL server administrator and the SSL server.

Capacity and performance issues prevented scaling (from field)

Multiple SSL servers under a single TCP/IP stack

When number of sessions reaches max, spills over to another server

Migration required

CPU Measurement Facility Counters – Host Support

VM64961

- Sets of counters for each logical processor that count events such as cycle, instruction, and cache directory-write counts
 - Same COUNTER information as z/OS partitions
- Accumulation is a relatively low-overhead activity and is performed automatically by the machine when the counters are authorized, enabled, and activated
- Authorization controlled by a logical partition's Security settings in its activation profile
- Enablement, activation, and data collection controlled by z/VM MONITOR command

zEnterprise Unified Resource Manager

VM64822, VM64904, VM64917, VM64956, VM64957

- z/VM V6 only
 - Turn-key installation option to enable virtual server management via zEnterprise Unified Resource Manager (z/VM V6.2 only)
 - Only for “kicking the tires”
- Enables Unified Resource Manager to perform system and virtual server management tasks
 - Virtual server configuration
 - Disk storage management
 - Virtual network management
 - Performance monitoring
- CP, CMS, LE, TCP/IP, DIRMAINT, Performance Toolkit, HCD
- <http://www.vm.ibm.com/service/vmrequirm.html>

z/VM can take direction from z/Manager

z/VM is different than p and x

- Will discover that it is running in a CEC
- System will automatically join Ensemble
- Can be managed by Ensemble Manager (Unified Resource Manager) (NOT REQUIRED)
 - (uses SMAPI interfaces)
- Directory Manager is required
- If customer chooses Ensemble Management, just use Ensemble management – do not manage natively also – choose one way or the other

Example: Using both Perf Toolkit and Ensemble Perf Manager, they do not know about each other and could conflict

OSX dedicated to guests via z/Manager

Vswitch dedication handled by CP

z/Manager not required for all z/VM Management – but choose one or the other

- can choose to have virtual serves not managed – they are not managed, do not manage them separately

Requires many steps to implement – Red Alert

- Requires Linux servers
- SMAPI servers

If VM64904 is applied VM64917 should also be applied (SMAPI server updates)

zEnterprise Unified Resource Manager Ensemble Membership

- If configured to participate in an ensemble, z/VM will automatically join the ensemble at IPL
- Configuration tasks
 - Set up OSM and OSX channel paths
 - Set up controllers for IEDN and INMN networks
 - Pre-defined controllers DTCENS1 and DTCENS2 for exclusive use by ensemble networks
 - DTCENS1 automatically creates a VSWITCH to provide SMAPI connectivity to INMN network
 - Configure directory manager (REQUIRED)
 - Configure SMAPI servers
- See chapter "Configuring z/VM for an Ensemble" in CP Planning and Administration manual

z/VM can take direction from z/Manager

z/VM is different than p and x

- Will discover that it is running in a CEC
- System will automatically join Ensemble
- Can be managed by Ensemble Manager (Unified Resource Manager) (NOT REQUIRED)
(uses SMAPI interfaces)
- Directory Manager is required
- If customer chooses Ensemble Management, just use Ensemble management – do not manage natively also – choose one way or the other

Example: Using both Perf Toolkit and Ensemble Perf Manager, they do not know about each other and could conflict

OSX dedicated to guests via z/Manager

Vswitch dedication handled by CP

z/Manager not required for all z/VM Management – but choose one or the other

- can choose to have virtual serves not managed – they are not managed, do not manage them separately

Requires many steps to implement – Red Alert

- Requires Linux servers
- SMAPI servers

Scalability and Performance Enhancements

Available by PTF to prior releases as shown

- Control of the guest page re-ordering process, improving the performance characteristics of guests with large memory footprints (VM64774)
- Reduced system overhead of guest page release function, thereby helping to increase guest throughput (VM64715)
- Improved contiguous frame coalescing algorithms help to increase system throughput (VM64795)

Scalability and Performance Enhancements

Available by PTF to prior releases where shown

- More accurate scheduling algorithm for guests that have LIMITHARD shares (VM64721)
- Reduce LPAR suspend time by reducing the number of DIAGNOSE 0x9C and 0x44 instructions issued when obtaining system locks (VM64927 for z/VM 6.1 only)
- Eliminate erratic virtual machine pause in busy systems with more than 14:1 total virtual-to-logical CPU over-commitment (VM64887)