z/VM Platform Update

March 11, 2013 Version 1.5

Brian W. Hugenbruch, CISSP
z/VM Security Design and Development
Bwhugen@us.ibm.com

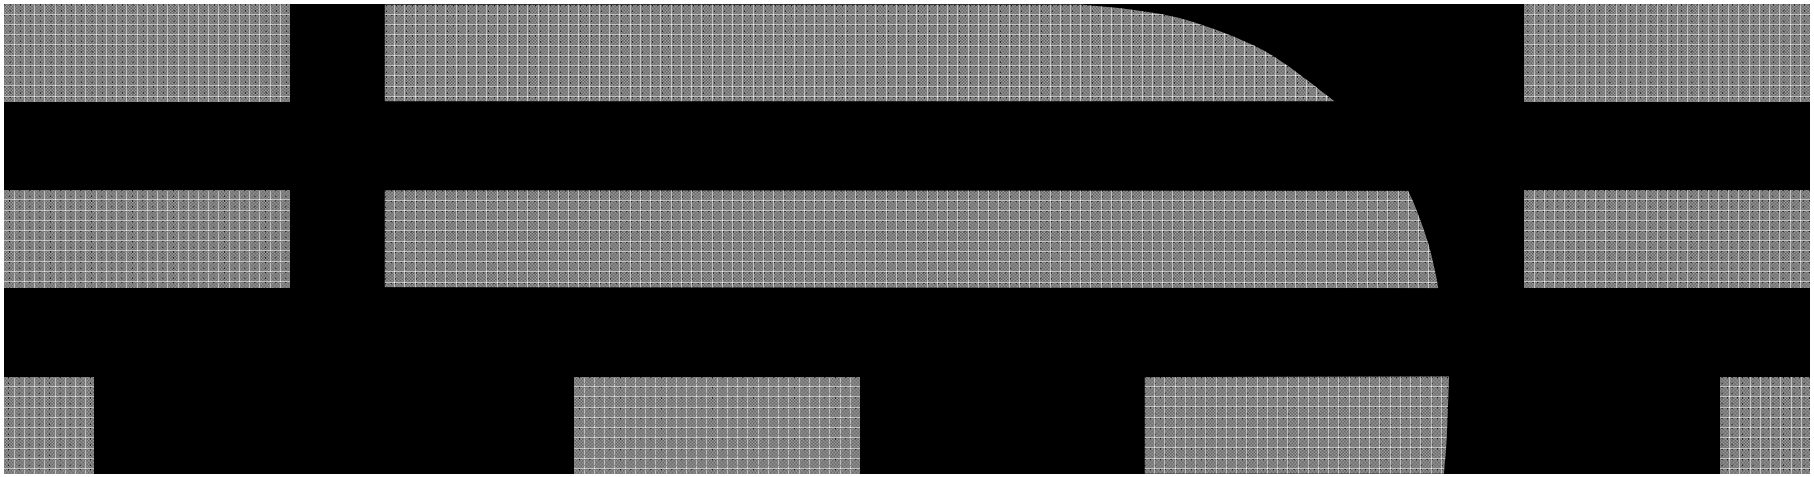

Trademarks

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

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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

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

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

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

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

Acknowledgments – Platform Update Team

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

Agenda

- Release Status and Information
- z/VM Version 6 Release 2
- z/VM Version 6 Release 3
- Futures and Statements of Direction
- Value of z/VM

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	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+ 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/LSPP

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

Extended support contracts are available.

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

z/VM Version 5 Release 4

- 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 Version 6
Security Certification Plans

- Common Criteria (ISO/IEC 15408)
  - z/VM 6.1 has been certified: BSI-DSZ-CC-0752
  - Evaluated to EAL 4+ for the Operating System Protection Profile (OSPP) with:
    - Virtualization extension (-VIRT)
    - Labeled Security extension (-LS)

- Federal Information Protection Standard (FIPS) 140-2
  - z/VM 6.1 System SSL is FIPS 140-2 Validated(TM)
  - Enablement requirements for certificate database and servers

- 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 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**

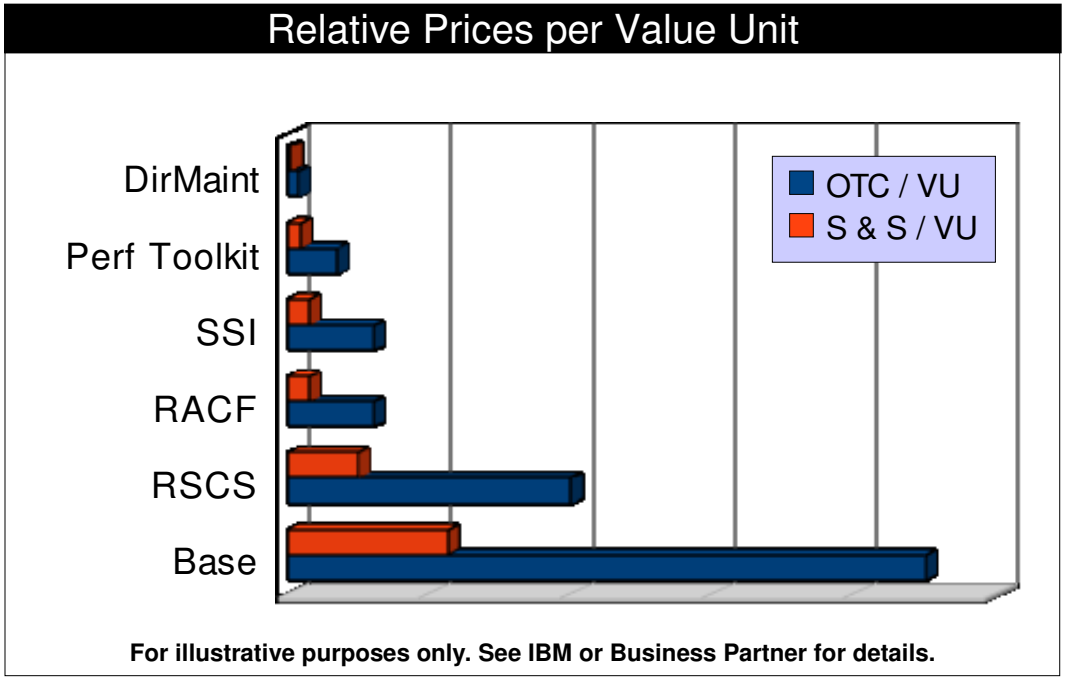

z/VM Version 6 Release 2

and other recent functional enhancements

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

- Optional priced feature
- Connect up to four z/VM systems as members of a Single System Image cluster
- 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

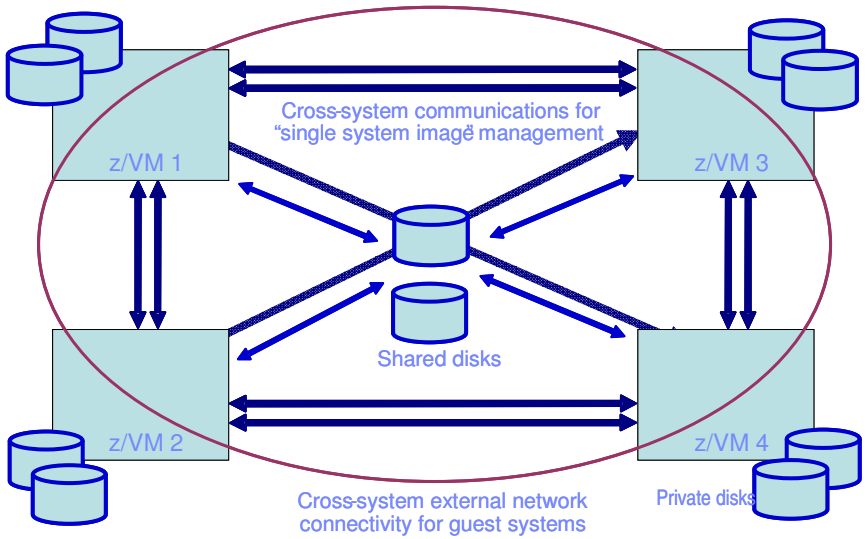

SSI Cluster Management – Features for Greater Reliability

- 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
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**

HiperSocket VSWITCH Bridge

Available: April 13, 2012

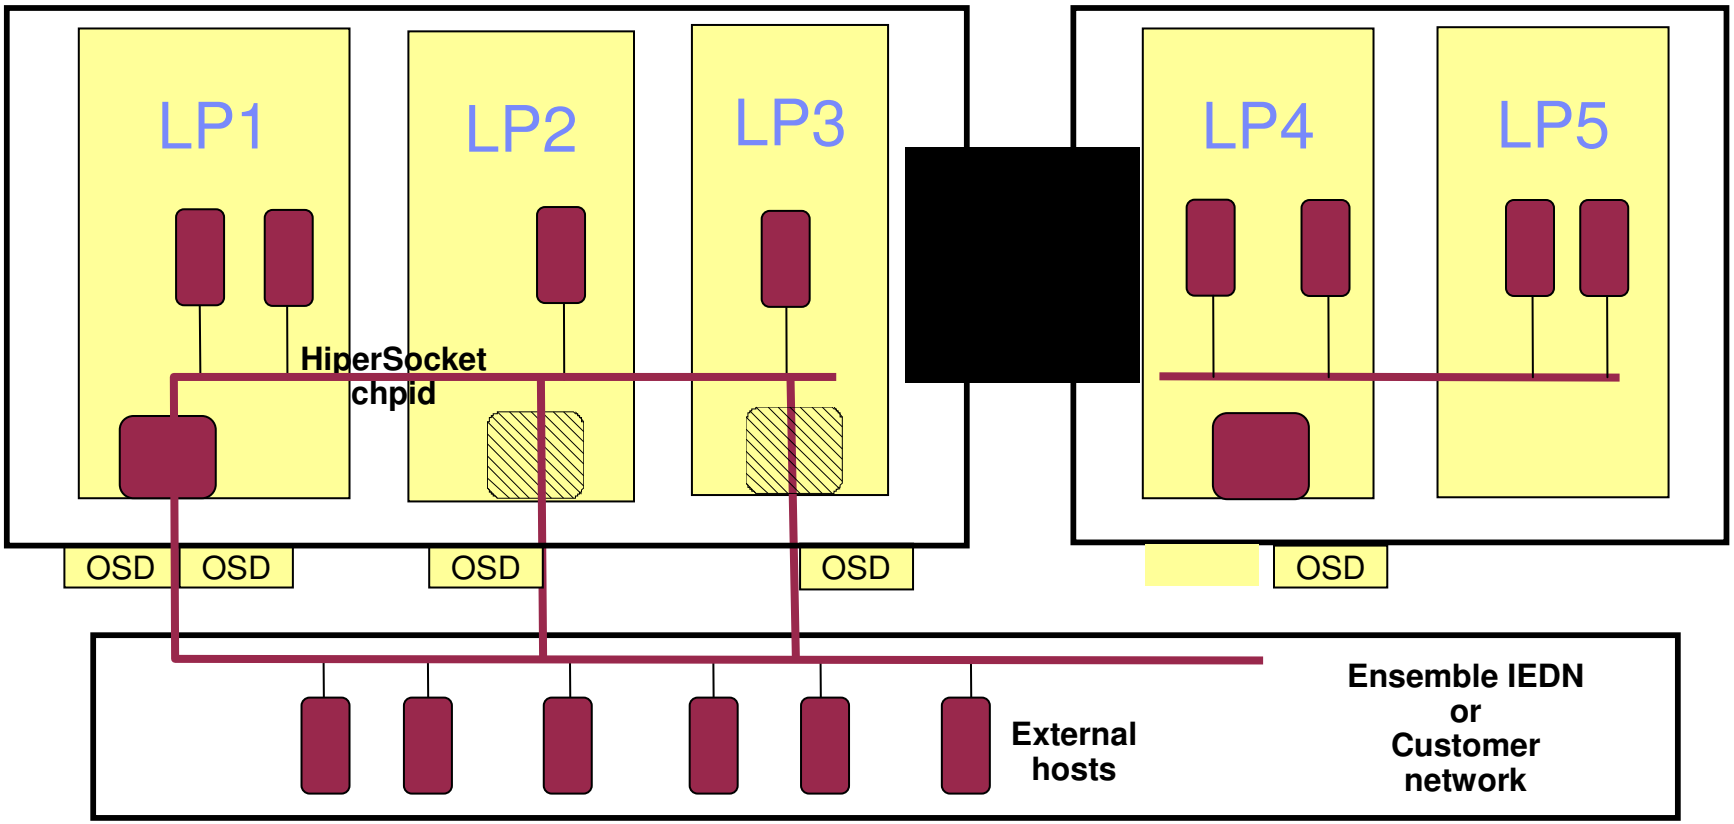

- Built-in failover and failback
- CHPARM=x2 to bridge IEDN
- CHPARM=x4 to bridge customer network
- Same or different LPAR
- One active bridge per CEC
- PMTU simulation

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

VSWITCH: Multiple virtual access ports per guest

Available: z/VM 6.2 base

- 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

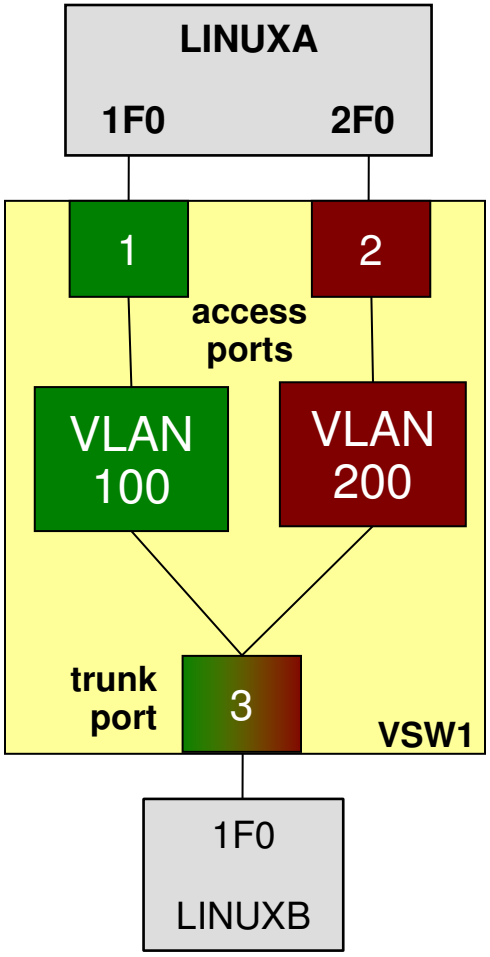

High Performance FICON

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

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.

Security Enhancements: RACF Security Server
Available: z/VM 6.2 base

- 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

Scalability and Performance Enhancements

Available: z/VM 6.2 base (some available for earlier releases)

- 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) or z/VM 6.2
- 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

z/OS R12 Equivalency Upgrades

Available: z/VM 6.2 base

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

Installation Improvements
Available: z/VM 6.2 base

- 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!*

CMS Enhancements: z/CMS

Available: z/VM 6.2 base

- 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

CMS Enhancements: XEDIT – Default changed to mixed case

Available: z/VM 6.2 base

- **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

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

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

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

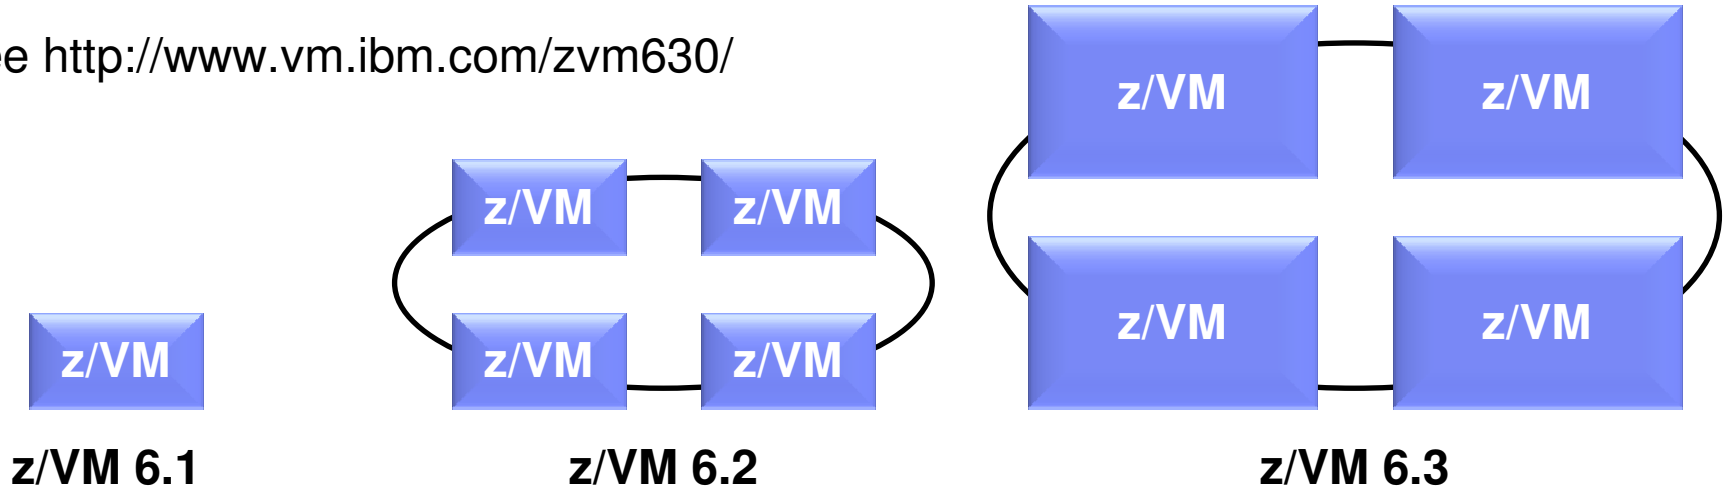

z/VM Version 6 Release 3 Preview
Large Memory Support

- 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: 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

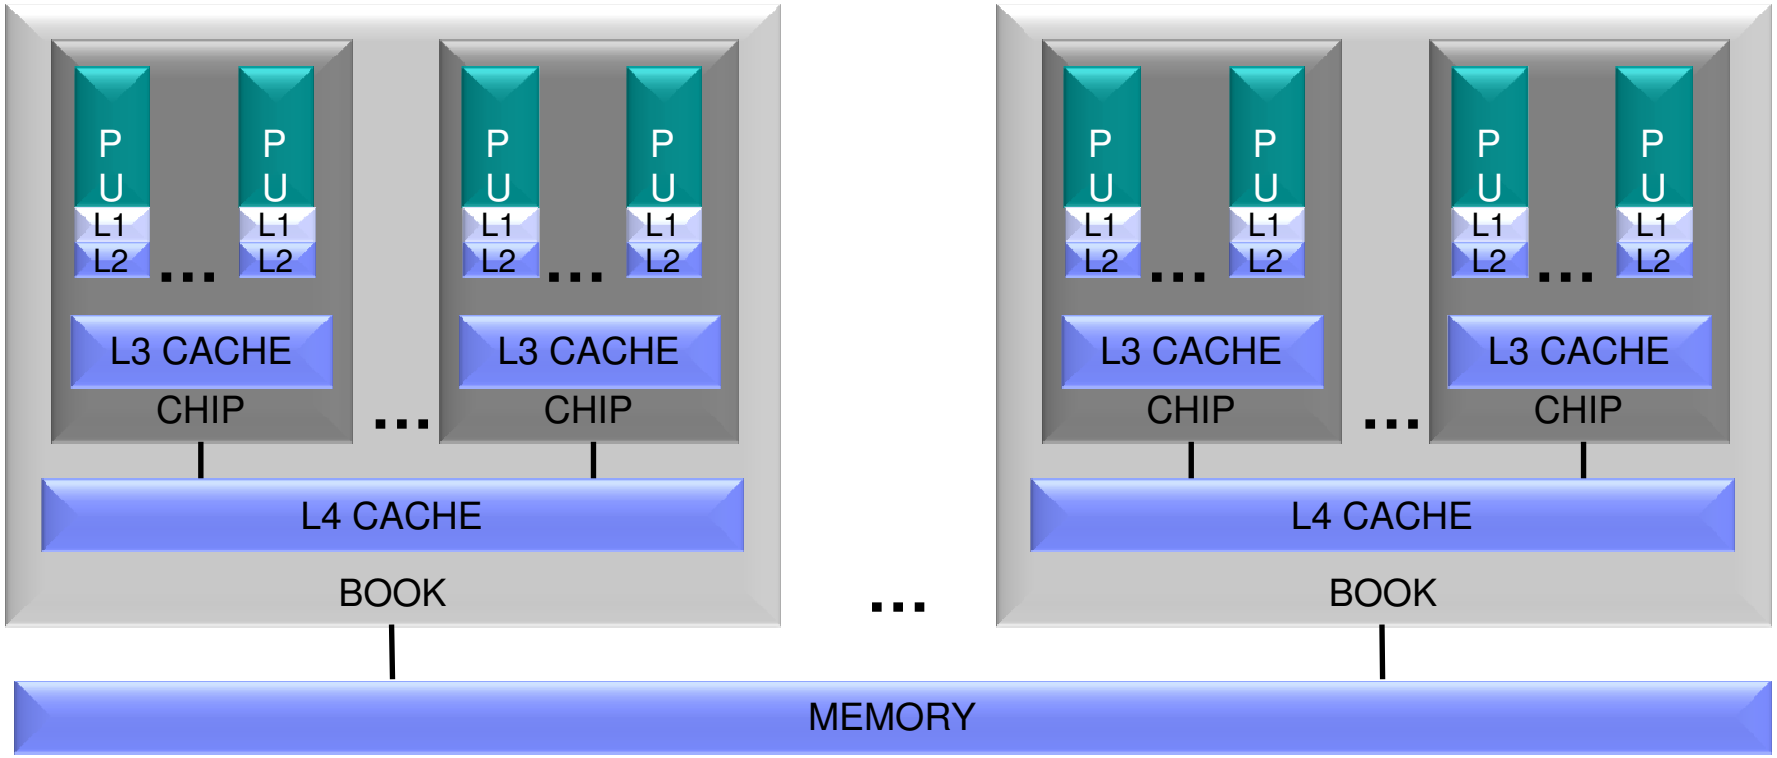

z/VM Version 6 Release 3 Preview
HiperDispatch: Dispatching Affinity

- 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

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)

z/VM Version 6 Release 3 Preview
Horizontal vs. Vertical CPU Management

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

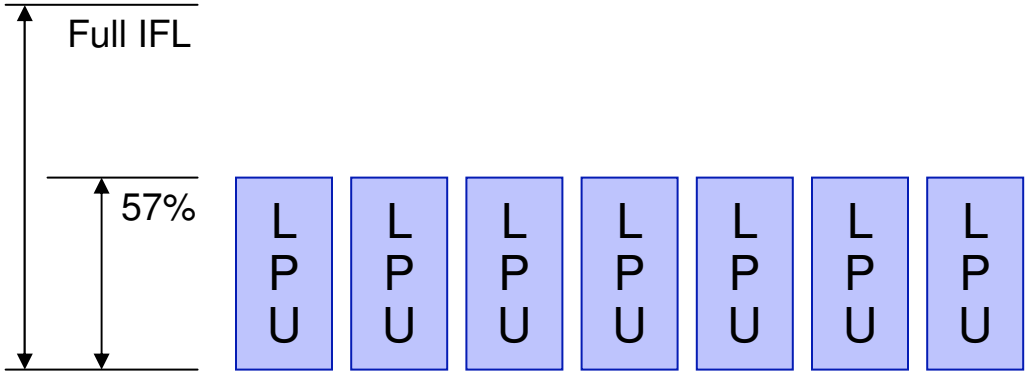

- Today's “horizontal” management distributes the LPAR weight evenly distributed across the logical processors of the z/VM LPAR

z/VM Version 6 Release 3 Preview
Horizontal vs. Vertical CPU Management

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.

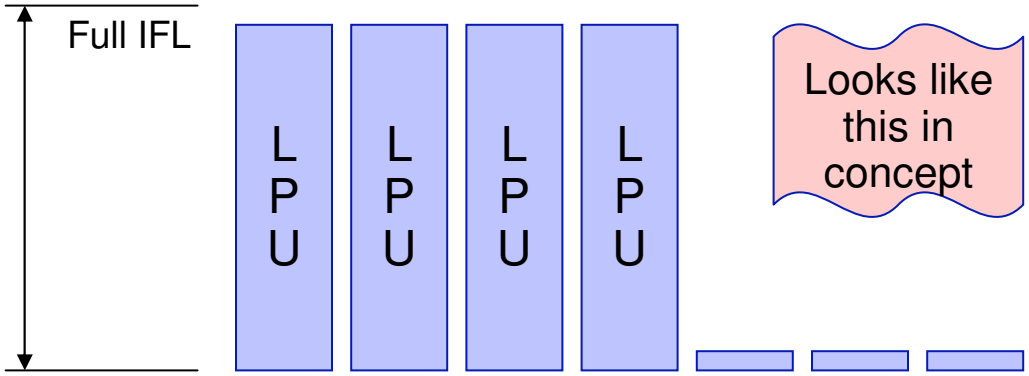

- "Vertical" management attempts to minimize the number of logical processors, allowing LPAR to similarly manage logical CPUs

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 – Available April 13, 2012
z/VM Statement of Direction: New function

- 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 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

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

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

Thanks!!

Contact Information:

Brian W. Hugenbruch, CISSP
bwhugen@us.ibm.com

Twitter: @Bwhugen