


# z/VM Platform Update

## Session 9447

Alan Altmark  
Senior Managing z/VM and Linux Consultant  
IBM Corporation



# Trademarks

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

IBM*	System z10
IBM Logo*	Tivoli*
DB2*	z10
Dynamic Infrastructure*	z10 BC
GDPS*	z9
HiperSockets	z/OS*
Parallel Sysplex*	z/VM*
RACF*	z/VSE
System z*	zEnterprise

\* 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.

---

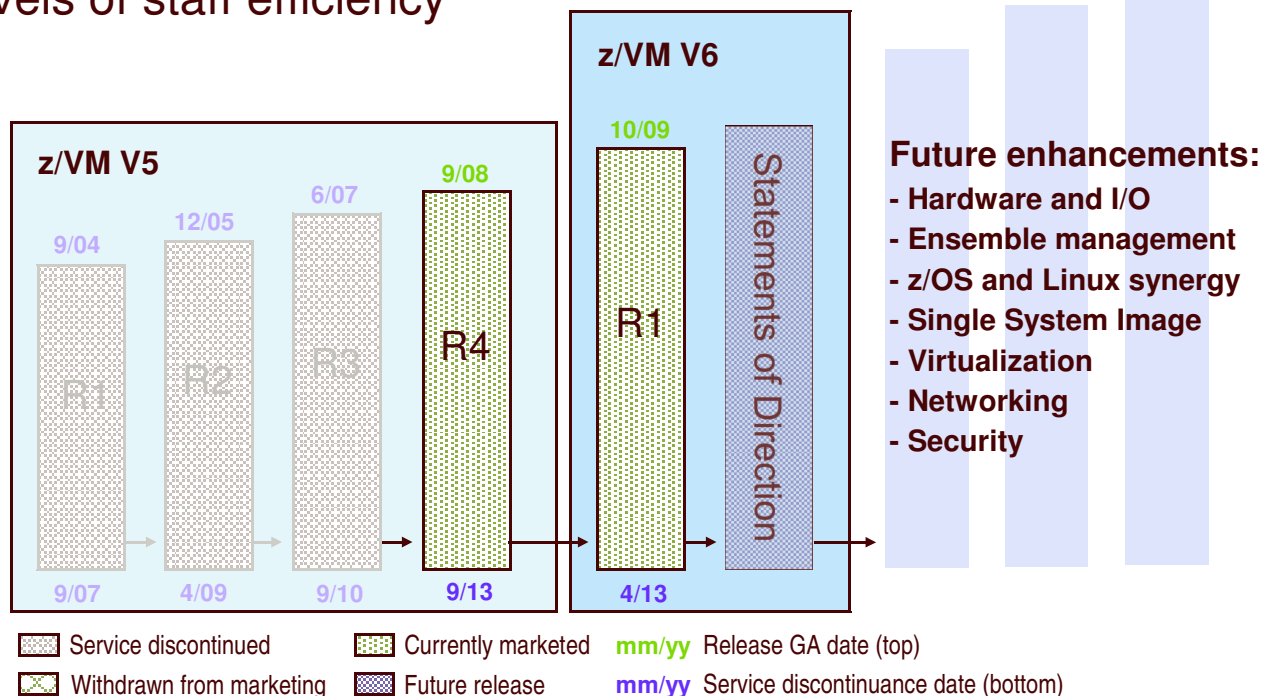
# Agenda

- z/VM Timeline
- z/VM V5.4 and V6.1 Enhancements
- The Future: IBM Statements of Direction

# z/VM Release Status

## ***z/VM: helping clients “do more with less”***

- Higher core-to-core consolidation ratios
- Higher levels of resource sharing and utilization
- Higher levels of staff efficiency



IBM received EAL 4+ certification of z/VM V5.3 from the German Federal Office of Information Security (Bundesamt für Sicherheit in der Informationstechnik) for conformance to the Controlled Access and Labeled Security protection profiles (CAPP and LSPP) of the Common Criteria standard for IT security, ISO/IEC 15408. z/VM V6.1 is currently undergoing evaluation against OSPP with the labeled security extension at EAL 4+.

---

## z/VM Version 5

### Marketing and Service Updates

- End of Service for z/VM V5.3 was September 30, 2010
- End of Service for z/VM V5.4 is September 30, 2013
- z/VM V5.4 is still marketed and available
  - z/VM V5.4 and z/VM V6.1 are available concurrently
  - Clients with System z9 or prior generations should acquire z/VM V5.4 now

## z/VM Version 6.1

### The Foundation for System z Virtualization

Available since October 23, 2009

- **Architectural Level Set establishes a new z/VM technology base on IBM System z10**
  - z/VM V6 operates only on z10 EC, z10 BC, z196, and z114
  
- **Allows optimization of z/VM function for greater business value on newer hardware**
  - Prefetch Data instruction improves performance of streaming network connections between guests on a VSWITCH
  
- z/VM clustering and guest mobility (statements of direction)
  - A more manageable ecosystem for cloud computing
    - add hardware to the workload
    - move workload to hardware
  - Helps clients avoid the virtual machine sprawl challenges of x86 systems: fewer real systems hosting thousands of server images

## New Function

- XRC timestamps
- Hyperswap improvements
- Imbedded z/OS component upgrade to R11
- SSL Server Reliability and Scalability
- Memory Management
- zEnterprise zManager

**Unless otherwise noted, all PTFs are available on RSU 6105 for z/VM 6.1**

**They will be on the next z/VM 5.4 RSU when it becomes available**

# XRC Timestamps

- VM64814
- VM64816

- Limited support for STP
  - CP will sync with STP at IPL (no need to deactivate LPAR)
  - CP will, optionally, obtain time zone and leap seconds from STP
  - Time will be placed in all host and guest I/O
  - Time sync checks will be observed by CP
    - No change to the host or guest TOD
    - Delta applied to I/O only
  - Enabled via SYSTEM CONFIG
  - Option to skip timestamp or delay I/O if CP is unable to sync with STP
- Does NOT include STP virtualization
  - Can be configured to allow use by 2<sup>nd</sup> level systems
    - No STP synchronization
    - Cannot be used 1<sup>st</sup> level



# Hyperswap Improvements

- VM64815
- 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

## z/OS R11 Upgrades - z/VM 6.1 only

- PM08418: Upgrade System SSL to z/OS R11
- VM64805: Add needed functions to LE
- VM64751: Upgrade Binder to z/OS R11
- PM10616: System SSL enablement of FIPS

- System SSL
  - FIPS capability
  - FIPS validation under consideration
- Binder
  - FIPS enablement for System SSL
- Configuration Information
  - <http://www.vm.ibm.com/related/tcpip/vmsslinf.html>

## SSL Server Reliability and Scalability

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

- Major rewrite
- Multiple SSL servers with 'resume' cache manager and shared database
  - Can 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
  - <http://www.vm.ibm.com/related/tcpip/tcsslspe.html>
- Updates to TCPIP as well

## Memory Management Updates

- VM64795: Coalesce free frames
- VM64774: Set / Query Reorder
- VM64715: Page release improvements (*pending*)

- Coalesce adjacent free frames
  - Solves contiguous frame problem for functions like virtual SIE that require multiple pages
  - Help avoid abends when you are using 2<sup>nd</sup> level systems in a memory-constrained environment
- SET / QUERY REORDER
  - CP periodically reorders resident guest pages in case CP needs to trim
  - Every 8 GB of resident memory results in 1 second of guest purgatory
  - Reorder can occur frequently in systems with high paging rate and a lot of guest CPU consumption
  - Can turn off reorder for any or all users
  - Not all performance problems are Reorder problems!
- Guest Page Release (diagnose 0x10)
  - Reduced contention for data structures that stop guest from running
  - Improved performance

## Performance “Must Have” Updates

- VM64927: Revised spin lock manager (z/VM 6.1 only)
- VM64887: Erratic system performance

- Revised Spin Lock Manager
  - Reduces number of diagnose 0x9C and 0x44 instructions issued by CP
  - Reduces LPAR suspend time
  - Available only for z/VM 6.1 due to use of z10+ instructions
- Erratic system performance
  - Condition: More than 14:1 over commitment of virtual to logical CPUs
  - Symptom: Brief virtual machine pauses
  - More visible when load is “bursty”

## zEnterprise System Ensemble Management via zManager

- VM64822: Base function
- VM64904: SMAPI Updates
- VM64917: SMAPI Updates (not on RSU)
- VM64956: SMAPI Updates (not on RSU)
- VM64957: CP update (not on RSU)

- VM64822 is a “Super PTF” that includes required z/VM service for:
  - z196/z114 compatibility
  - zManager support, <http://www.vm.ibm.com/service/vmrequirm.html>
  - CP, CMS, LE, TCP/IP, DIRMAINT, Performance Toolkit, HCD
  - All of the PTFs are on the RSU, but additional fixes are required
- Enables zManager to perform system and virtual server management tasks
  - Virtual server configuration
  - Disk storage management
  - Virtual network management
  - Performance monitoring

## zEnterprise System Ensemble Management via zManager

- Choose ensemble management or traditional management for z/VM – do not use both
- If configured to participate in ensemble management, z/VM will automatically join the ensemble at IPL after configuration tasks are performed
- 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

## Networking in an Ensemble - OSA

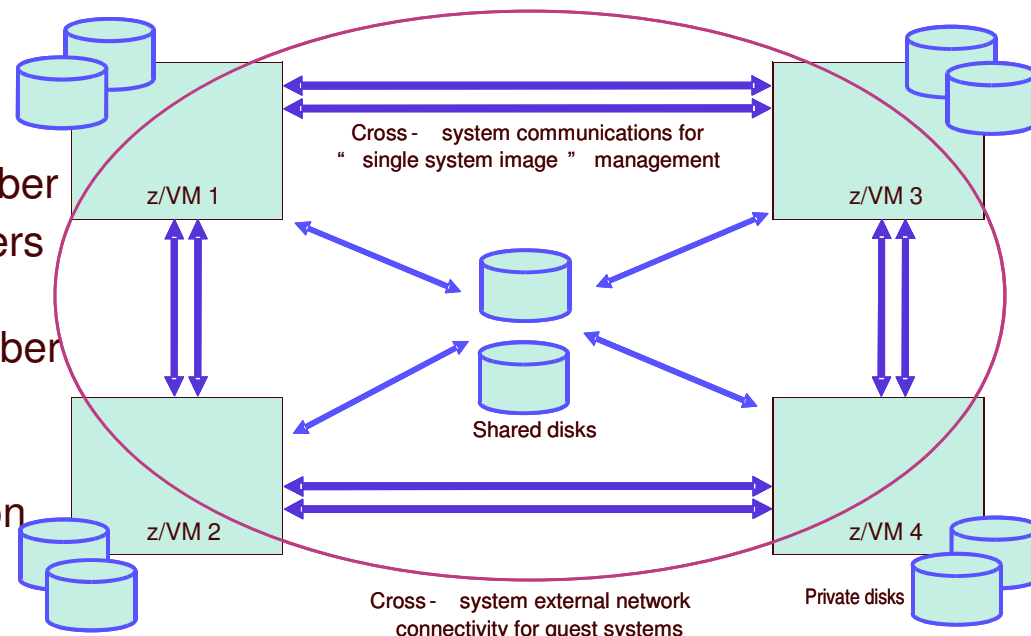
- Virtual networks are really real...
  - IntraEnsemble Data Network – IEDN – OSX chpid type (you)
  - IntraNode Management Network – INMN – OSM chpid type (IBM)
  
- ...and really virtual
  - IEDN VSWITCH
    - z/VM guest access via dedicated OSX
  - INMN VSWITCH
    - Defined when SMAPI server is started
  
- Traditional network connections via OSD
  
- When z/VM is not part of an ensemble, OSX and OSM chpids will not come online
  - Their use requires active participation of OSA firmware in cooperation with zManager



# z/VM Statement of Direction

## Clustered Hypervisor with Guest Mobility


- Clients can connect up to four z/VM systems as members of a Single System Image (SSI) 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 systems 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



Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

## z/VM Statement of Direction Clustered Hypervisor with Guest Mobility

- 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!



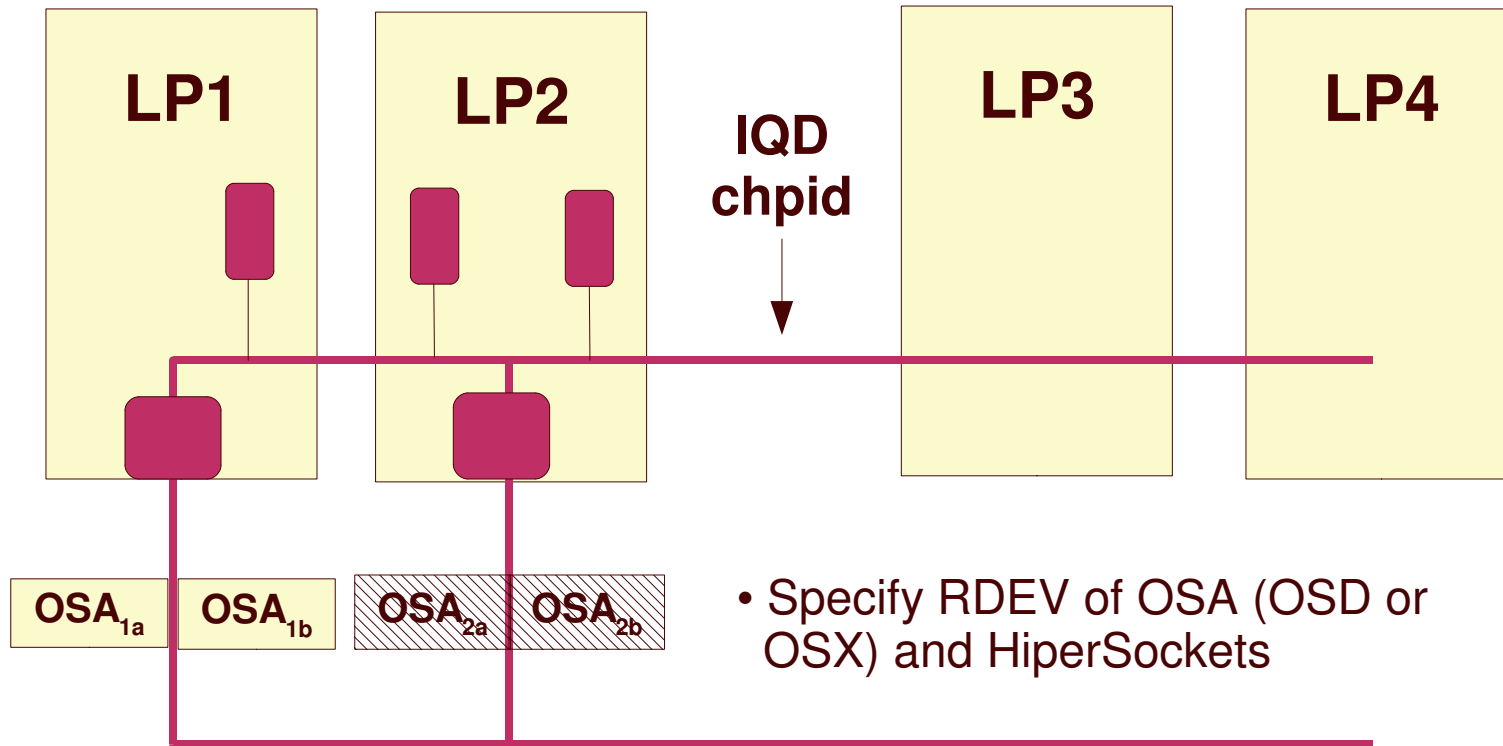
## z/VM Statement of Direction HiperSocket Bridge



- Virtual Switch bridge between layer 2 ethernet and HiperSockets
  - zEnterprise IEDN
  - OSA (QDIO) to external network
  - HiperSocket chpid
  - Guests can use virtual OSA or dedicated HiperSockets
- Firmware routes Guest-Guest and Guest-LPAR traffic via real HiperSocket
  - All other traffic flows to VSWITCH
- VSWITCH forwards outbound traffic to guest or HiperSocket
- Up to four VSWITCHes can act as bridge
  - 1 active, 3 standby
  - Automatic takeover
  - Optionally designate one “primary”
    - Primary will perform “takeback” when it comes up

# z/VM Statement of Direction

## HiperSocket Bridge



- Specify RDEV of OSA (OSD or OSX) and HiperSockets

- LP1 or LP2 will handle bridge traffic
- Built-in failover and failback

- Same or different LPAR
- Up to 4 bridges

# New Releases of IBM z/VM Management Software

since September 2010

## ▣ **Operations Manager for z/VM V1.4**

- Improved support for repeating events
- Can define unique RACF or other ESM profiles per Operations Manager instance
- Improve data sharing among action routines
- Usability improvements when viewing consoles, spool files, and the Operations Manager log
- Support for IPv6

## ▣ **Tape Manager for z/VM V1.3**

- Share a single tape catalog among multiple Tape Manager instances
- Usability enhancements
- Support for RACF or other ESM control of commands and tape pool access
- Manage and use StorageTek® tape devices

# Thanks!

*Contact Information:*

Alan Altmark  
IBM Senior Managing IT Consultant  
[Alan\\_Altmark@us.ibm.com](mailto:Alan_Altmark@us.ibm.com)