Alan Altmark, z/VM Architecture and Design SHARE - August 2010



z/VM Update: z/VM 6.1 and Beyond

Session 6972





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.

© 2009, 2010 IBM Corporation



Agenda

- z/VM Timeline
- z/VM Version 5
 - Enhancements
- z/VM Version 6
 - Enhancements
 - 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. While z/VM V5.4 and V6.1 have not been officially evaluated for conformance, they are designed to meet the same standards.

© 2009, 2010 IBM Corporation



- End of Service for z/VM V5.3 is September 30, 2010 Soc
- End of Service for z/VM V5.4 has been extended to 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



z/VM Version 6.1 The Foundation for System z Virtualization Growth Previewed July 7, 2009; Available 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, and z196

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

Multi-system virtualization support (future release support)

- -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



Soon: New for "0910" – delivering Sept-Nov

- XRC timestamps
- z/OS R11 upgrade
- SSL Server
- Memory Management
- VSWITCH uplink ports
- zEnterprise

7

© 2009, 2010 IBM Corporation



XRC Timestamps

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
 - Option is included to allow use by 2nd level systems without STP synchronization
 - Cannot be used 1st level



z/OS R11 Upgrades - z/VM 6.1 only

- System SSL
 - FIPS capability
 - FIPS validation under consideration
- Binder
 - FIPS enablement for System SSL



SSL Server - z/VM V6.1 Only

- 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
- Updates to TCPIP as well



Memory Management Updates

- Coalesce adjacent free frames
 - Solves contiguous frame problem for functions like virtual SIE that require multiple pages
- Guest Page Release (diagnose 0x10)
 - Reduced contention data structures the stop guest from running

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
- Recommended only at the direction of IBM support
 - Not all performance problems are Reorder problems!



Virtual Switch Uplink Ports "It's not your grandfather's VSWITCH!"





zEnterprise System Ensemble Management

Provides management of z/VM V6.1

- Virtual server configuration
- Disk storage management
- Virtual network management
- Performance monitoring
- Virtual networks are really real...
 - Intraensemble data network IEDN OSX chpid type (you)
 - Intranode management network INMN OSM chpid type (IBM)
- ...and really virtual
 - -z/VM guest access via dedicated OSX
 - OSX and OSM VSWITCH
- Seamless administration
- Traditional network connections via OSD





z/VM Virtual Switch Configuration in support INMN (OSM)



© 2009, 2010 IBM Corporation

z/VM Platform Update



Networking in an Ensemble - OSA

- 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

 zEnterprise System - Network Architecture and Virtualization Overview Alfred B. Christensen, IBM Corporation Wednesday 1:30 p.m. - 2:30 p.m. Room 109

 IBM zEnterprise Unified Resource Manager Overview Romney White, IBM Corporation Wednesday 11:00 a.m. - 12:00 p.m. Room 305

z/VM Statements of Direction Clustered Hypervisor Support and Guest Mobility

- Clients can cluster up to four z/VM systems in a Single System Image (SSI)
- Provides a set of shared resources for the z/VM systems and their hosted virtual machines
- Users can run z/VM system images on the same and/or different System z10 servers
- Simplifies systems management of a multi-z/VM environment
 - Single user directory
 - Cluster management from any system
 - Apply maintenance to all systems in the cluster from one location
 - Issue commands from one system to operate on another
 - Built-in cross-system capabilities
 - Resource coordination and protection: network and disks
- Dynamically move Linux guests from one z/VM system to another with Live Guest Relocation
 - Reduce planned outages; enhance workload management
 - Non-disruptively move work to available system resources <u>and</u> non-disruptively move system resources to work

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

Cross-system communications for "single system image" management Z/VM 3 Shared disks Z/VM 2 Cross-system external network

connectivity for guest systems

z/VM Single System Image

and Guest Mobility Preview

Tuesday 1:30 p.m. - 2:30 p.m.

Room 208

John Franciscovich, IBM Corporation



IBM

z/VM and System z Virtualization Leadership Multi-Level Workload and Resource Management

Dynamically Add Resources to Workload and Move Workload to Resources



18 Notes: a) z/VM shutdown, required to remove LPAR memory, b) virtual machine mobility requires Live Guest Relocation support (planned)

© 2009, 2010 IBM Corporation



Sessions of interest for IBM zEnterprise

z/VM Platform Update



Today

 Managing Resources, Applications and Data in the New Dimension in Computing Jeff Frey, IBM Corporation and Mike Baskey, IBM Corporation 3:00 p.m. - 4:00 p.m. Room 306



Tuesday

Evaluating Best Fit Architecture for Customer Workloads on zEnterprise Systems
 Doris Conti, IBM Corporation
 Tuesday 9:30 a.m. - 10:30 a.m.
 Room 302

 IBM zEnterprise 196 (z196) Hardware Overview Harv Emery, IBM Corporation *Tuesday 1:30 p.m. - 2:30 p.m.* Room 302

zEnterprise Unified Resource Manager Patty Driever, IBM Corporation *Tuesday 3:00 p.m. - 4:00 p.m.* Room 302

Platform Performance Management Overview
 Hiren Shah , IBM Corporation
 Tuesday 4:30 p.m. - 5:30 p.m.
 Room 302



Wednesday & Thursday

- IBM zEnterprise Unified Resource Manager Overview Romney White, IBM Corporation Wednesday 11:00 a.m. - 12:00 p.m. Room 305
- zEnterprise System Network Architecture and Virtualization Overview Alfred B. Christensen, IBM Corporation Wednesday 1:30 p.m. - 2:30 p.m. Room 109
- Energy Management for zEnterprise Andreas Bieswanger, IBM Corporation Thursday 8:00 a.m. - 9:00 a.m. Room 313
- z/OS Software Positioning for the IBM zEnterprise System Gregory Daynes, IBM Corporation *Thursday 3:00 p.m. - 4:00 p.m.* Room 310

IBM

