#SHAREorg

SHARE

z/VM 6.3 Early Support Program Experiences

Session - 13511

- Richard Lewis – IBM Washington System Center
- Bob Bates – Wells Fargo
- Rick Barlow – Nationwide
- James Vincent – Nationwide

SHARE in Boston

Trademarks

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

APPN*
CICS*
DB2*
DFSMSMVS
DFSMS/VM*
DirMaint
Distributed Relational Database Architecture*
DRDA*
e-business logo*
ECKD
Enterprise Storage Server*
Enterprise Systems Architecure/390*
ESCON*
FICON*
GDDM*

* Registered trademarks of IBM Corporation

GDPS*
Geographically Dispersed Parallel Sysplex
HiperSockets
HyperSwap
IBM*
IBM eServer
IBM logo*
IBMlink
Language Environment*
MQSeries*
Multiprise*
On demand business logo
OS/390*
Parallel Sysplex*
Performance Toolkit for VM
POWER5

POWERPC*
PR/SM
Processor Resource/Systems Manager
QMF
RACF*
Resource Link
RMF
RS/6000*
S/390*
S/390 Parallel Enterprise Server
System 370
System 390*
System z
System z9
Tivoli*
Tivoli Storage Manager
TotalStorage*

Virtual Image Facility
Virtualization Engine
VisualAge*
VM/ESA*
VSE/ESA
VTAM*
WebSphere*
z/Architecture
z/OS*
z/VM*
z/VSE
zSeries*
zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a trademark of Linus Torvalds in the united States and other countries..

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States and other countries.

* 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

- Overview of z/VM 6.3 Early Support Program
- Customers Participating in the ESP
- Customer Experiences
- General Observations
- Wells Fargo Experiences
- Nationwide Experiences

z/VM Early Support Program Objectives

- Exercise the new features of z/VM V6.3 in customer production environments to gather feedback on product quality
  - Large Memory Support
  - HiperDispatch
  - Upgrade Installation
  - xCat and zHCP
  - Disable CSE
  - Virtual Switch Enhancements
  - Real device addresses in CP_OWNED or USER_Volume statements
  - Large memory dump support
- Provide vendors and IBM product developers early access to the new release
- Exercise the support structure for the new release

z/VM Product Introduction Schedule

- 11/09/12 z/VM 6.3.0 in production on GDLVM7
- 11/12/12 Field Test Begins
- 11/13/12 z/VM 6.3.0 in production on GDLVME
- 11/13/12 z/VM 6.3.0 in production on S390VM (IBM)
- 12/17/12 ESP begins without large memory support
- 02/18/13 Large memory code released to ESP
- 05/23/13 Code Freeze
- 07/15/13 Shipped z/VM 6.3.0 GA DDR to ESP customers
- 07/26/13 z/VM 6.3.0 General Availability

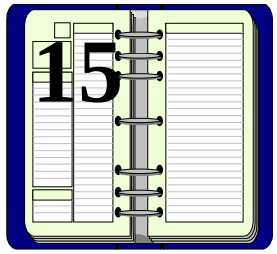

z/VM 6.3 Worldwide ESP Customers

- Industries
  - Financial - 5
  - Universities - 1
  - Manufacturing - 2
  - Government - 1
  - Insurance - 2
  - Transportation – 1
  - Health Care – 2
  - Information Processing – 2
- Plus
  - Software Vendors - 17
  - IBM Internal - 29

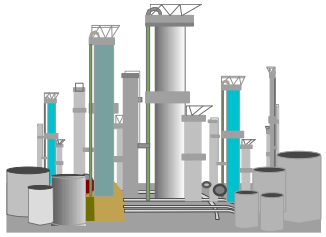

- Geographic
  - N. America - 10
  - Europe – 4
  - L. America - 2

- Processors
  - z10s (3)
  - z196s (8)
  - Z114 (2)
  - zEC12 (5)

z/VM 6.3 Worldwide ESP Customers…

- Major Hospital and Research Center
  - ✓ Linux guests, CMS,DB2, UDB
- Major Health Care Company
  - ✓ Linux guests, z/OS, CMS, GDPS, WAS, MQ
- 2 Information Services Companies
  - ✓ Linux guests, z/OS, Cognos
- Manufacturing
  - ✓ SAP
- 5 Large Banks and Financial Services
  - ✓ z/OS, Linux, TPF & CMS, GDPS
- 1 University
  - ✓ Oracle, Linux, CMS, z/OS
- 2 Major US Insurance Companies.
  - ✓ Linux, CMS, z/OS, WAS, Oracle
- Large European Air Traffic Control
  - ✓ Linux, Oracle
- German Manufacturing Co.
  - ✓ SAP
- German Government Agency
  - ✓ z/OS, Linux, GDPS

Customer Experiences

- zvm ppf file removed
  - Use servp2p ppf instead
- Support for user class restructure and override command removed
  - Use CP MODIFY command instead
- service and put2prod should be used instead of the more primitive receive, apply, and build VMSES commands
- DASD being used by new SDINST utility needed to be CMS formatted instead of cpfmtxa
- Upgrade Installation requires real volumes to be available matching the geometry of the DASD on which the current system was installed.
  - 13 3390-03 if everything installed on minidisk (8 given back)
  - 7 3390-09 if everything installed on minidisk (4 given back)

Customer Experiences

- Problem with Upgrade Installation when SYSTEM CONFIG had imbed files. They were not copied over with SYSTEM CONFIG resulting in CPSYNTAX errors
- Upgrading member running DIRMSATx virtual machine
  - Sometimes response too slow causing error in upgrade
  - Work around: move DIRMAINT service machine to member being upgraded
- Suggestion with large memory support is to remove xstore and add memory back into central storage
  - VMRelocate problems in multimember cluster moving from member with xSTORE defined to member without xSTORE defined.
  - Not a z/VM 6.3 problem also happens in 6.2. (VM65306)
- DASD paging volume recommendations change with new large memory support
  - Have a robust paging subsystem (lots of exposures across lots of backend resources)
  - Volumes allocated to more than 50% are now OK
  - Might have problems if try to get away with small amount of DASD paging space
    - Calculate sum of logged-on virtual machine primary address spaces + any data spaces + vdisks + total number of NSS and DCSS pages.
    - Multiply sum above by 1.01 to allow for PGMBKs and friends
    - Add to result above total number of directory pages (reported by DIRECTXA)
    - Add to result above max of (10 % of central or 4GB) for system-owned virtual pages
    - Multiply result above by 1.25 as safety factor

Customer Experiences

- VM HCD does not support designation of 3390D (which z/OS uses for mirror devices) however z/VM 6.3 with multiple subchannel set support does support the devices
- xCAT and zHCP are part of the CMS component
  - Replacement part service of large image files will cause service disks to need resizing over time
- Need to make sure the Global Performance Data setting is on in LPAR profile for HiperDispatch

General Observations

- z/VM 6.3's Upgrade in Place process greatly simplifies upgrading to new releases of z/VM. It saves hours of work and reduces the chance of errors. (Large Financial Institution)

#SHAREorg

SHARE
Technology · Connections · Results

z/VM 6.3 Early Support Program Experiences

Bob Bates
Wells Fargo Bank

August 15, 2013
Session Number 13511

SHARE in Boston

Disclaimer

- The information and opinions found herein are for informational sharing purposes only and not necessarily those of Wells Fargo Bank and should not be considered an endorsement.

Agenda

- Who we are
- Our configuration
- Our experiences

About Wells Fargo Bank

- Headquartered in San Francisco
- Acquired Wachovia Bank December 2008 (doubled) – completed migration end of 2011
- 70 Million customers, 1/3 of US Households have relationship
- 9000 Stores (branches)
- #1 Mortgage lender (commercial and residential)
- #3 ATMs (12,355)
- #4 in US in full service brokerage & wealth management
- #5 Insurance brokerage
- #4 in assets
- #1 in market cap
- 1st to do internet banking
- 270,000 employees

About me

- Systems programmer for z/VM and Linux on z Series at WF
- AVP, Operating Systems Engineer
- Enterprise Hosting Services
- At WF since 2007
- 7 team members
- Install/support/engineer Linux on z and z/VM builds
- Level 2-3 for all Linux on z servers

Our z/VM environment

- Production
  - 4 LPARs across 2 sites with failover between them
  - z196 with 48 IFLs total
  - z/VM 6.2
- Dev test
  - 3 LPARs on 2 boxes, 6 IFLs on each
    - z10 2 LPARs 6 IFLs 6.2, 2 IFLs 6.3
    - z196 1 LPAR 6 IFLs 6.3

Development/Test environment

z10 LPAR F	z10 LPAR E	z196 LPAR 6
2 IFLs 4G Sandbox zVM 6.3	6 IFLs 230G 132 Linux test dev zVM 6.2 Memory about 1:1 virtual to real	6 IFLS 208G 77 Linux test dev zVM 6.3 Memory about 1.5:1 virtual to real

Why z/VM 6.3?

- Upgrade in place
  - Wanted simplified rollout
- Large memory

The Experience

- Mid December 2012 downloaded files for DVD and built install disks
  - Used 2 system 6.2 test 2nd level system
  - Basic packages installed. (VM:Secure, VM:Spool, VM:Operator
- Early February Update in Place available. But VM:Secure not ready for us (Oy, legal stuff).
- Finally got everything going and began UIP late February

Problem!

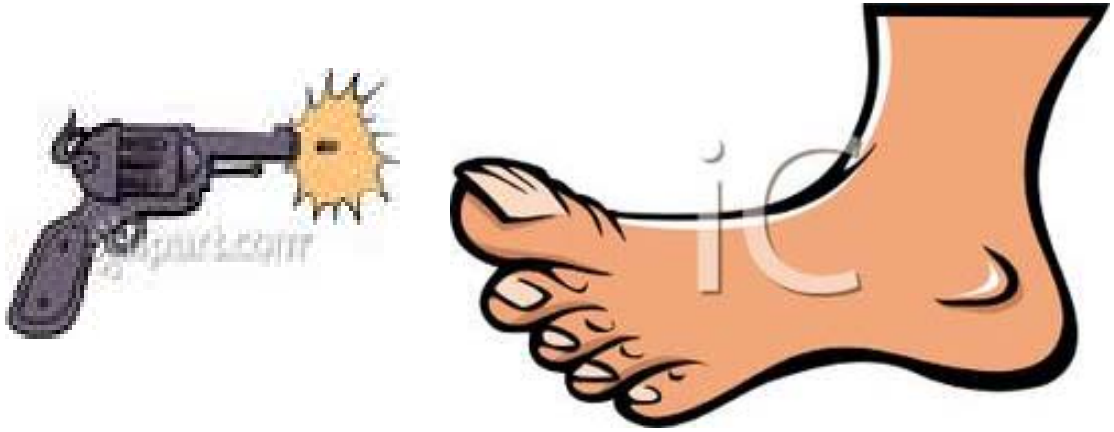

Maintenance

- Year ago an error occurred trying to test special fix
- Ended up copying files from 1st level to 2nd level to recover.
- HCD was installed in MDISK 1st level, SFS 2nd level
- INSTUPGR got confused.
- Finally got things sorted out, moved HCD to MDISK and cleaned out all the HCDSFS stuff.

Issues during test

- 6.2 requires fixes to be able to apply maintenance in a mixed cluster.
  - VM65317 (VMSES)
  - VM65318 (CP)
  - VM65319 (CMS)
  - VM65320 (RACF)
- Changes to prefix page map changes HCPCALL (VM:Secure had to make adjustments)
- XCAT made CMS maintenance file get very large.
- Converting from ESP to GA got complicated.

ESP to GA

- Process for getting away from the ESP was slow in coming but we got there.

Summary

- ESP experience is extremely valuable to our company
- 6.3 continues to build on the cluster environment and simplifies our support processes.

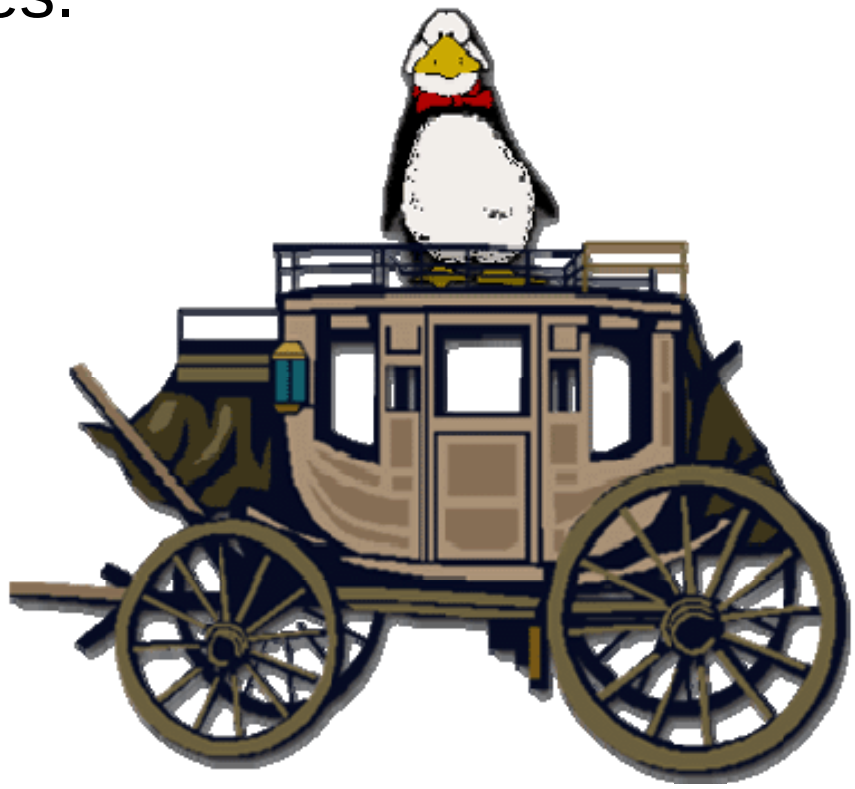

- My info
  - Robert.Bates@wellsfargo.com
  - 469-892-6660

#SHAREorg

SHARE
Technology · Connections · Results

z/VM 6.3
Early Customer Experiences

James Vincent

Rick Barlow

Nationwide Insurance

SHARE in Boston

A bit of history & why participate

- Nationwide has a long history of ESPs, beta and even research & design teams with IBM
- z/VM ESP program is well structured *and* flexible to fit work interrupts
- Gives the opportunity to…
  - Hear and use the newest advancements
  - Talk to lead developers and architects of the z/VM code
  - Influence direction and function
- Why Nationwide? Reasonably large z/VM footprint and zLinux usage

Installing

- Installed second level, then “copied” to two LPARs
- Worked on infrastructure third-party product migration and upgrades to support 6.3
- Quite a few updates along the way, then a re-install when GA came in
- Initial findings
  - Large amount of thought, time and effort went into it
  - Pretty darn stable
- Some packaging decisions were hard to accept

Large Memory Support

- IBM offered us about 500GB of memory to test with
- Built five 200GB RHEL 6 Linux servers (2:1 ratio)
- First test: start them up, shut them down
- Second test: a mean memory app, random use and sizes
  - Oops – 4G of page space wasn't enough!
  - Sustained paging at 40,000 - 50,000/second
  - Performance seemed fine
  - Then, we shut them all down at once
    - Ouch
  - VM65368
- Reorder processing is moot now

Improved Paging

- The old 50% rule for page space is no more (yea!)
- We have hundreds of Mod 9 page volumes – half sitting empty!
- You HAVE to watch page space calculations though
  - Heavier use of paging will demand a little more attention to allocating page space

VSWITCH Recovery / Stall Prevention

- Allows a controlled failover to a backup OSA port
  - Network interruption is barely noticeable
  - Much cleaner than previous methods

Summary

- z/VM 6.3 has some serious internal CPU and Memory management enhancements
  - Although most changes are not external bells-and-whistles, the updates are a Big Deal
    - HyperDispatch internals are geeky-cool!
- Be diligent in watching for APARS and remain very current
  - GA RSU plus other PTFS now available

Thank you!