



z/VM 6.2 Early Support Program Experiences

Session 10361

Richard Lewis, IBM Washington Systems Center
Marcy Cortes, Wells Fargo
Dennis O'Brien, Bank of America
Jim Vincent, Nationwide Insurance

March, 2012

Trademarks

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

| | | | |
|---|---|------------------------------------|------------------------------|
| APPN* | GDPS* | POWERPC* | Virtual Image Facility |
| CICS* | Geographically Dispersed Parallel Sysplex | PR/SM | Virtualization Engine |
| DB2* | HiperSockets | Processor Resource/Systems Manager | VisualAge* |
| DFSMSMVS | HyperSwap | QMF | VM/ESA* |
| DFSMS/VM* | IBM* | RACF* | VSE/ESA |
| DirMaint | IBM eServer | Resource Link | VTAM* |
| Distributed Relational Database Architecture* | IBM logo* | RMF | WebSphere* |
| DRDA* | IBMLink | RS/6000* | z/Architecture |
| e-business logo* | Language Environment* | S/390* | z/OS* |
| ECKD | MQSeries* | S/390 Parallel Enterprise Server | z/VM* |
| Enterprise Storage Server* | Multiprise* | System 370 | z/VSE |
| Enterprise Systems Architecture/390* | On demand business logo | System 390* | zSeries* |
| ESCON* | OS/390* | System z | zSeries Entry License Charge |
| FICON* | Parallel Sysplex* | System z9 | |
| GDDM* | Performance Toolkit for VM | Tivoli* | |
| * Registered trademarks of IBM Corporation | POWERS5 | Tivoli Storage Manager | |
| | | TotalStorage* | |

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 6.2 Early Support Program**
- **Customers Participating in the ESP**
- **General Observations**
- **Customer Experiences**

z/VM Early Support Program Objectives

- **Exercise the new features of z/VM V6.2 in customer production environments to gather feedback on product quality**
 - **SSI Feature**
 - **Live Guest Relocation**
 - **Technology Exploitation (e.g. FICON Express 8S)**
 - **Systems Management**
 - **Scalability enhancements (e.g. Disable guest page reorder)**
 - **Security enhancements (e.g. RACF support for SSI)**
- **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


- ▶ **03/26/11 z/VM 6.2.0 in production on GDLVM7**
- ▶ **03/29/11 z/VM 6.2.0 in production on GDLVME**
- ▶ **04/07/11 Shipped z/VM 6.2.0 Field Test level to vendors & IBM internals**
- ▶ **04/11/11 Field Test begins**
- ▶ **04/23/11 z/VM 6.2.0 in production on S390VM (IBM)**
- ▶ **05/05/11 Shipped z/VM 6.2.0 ESP level to ESP customers**
- ▶ **05/09/11 ESP begins**
- ▶ **10/22/11 First customer in production (medical)**
- ▶ **11/15/11 Shipped z/VM 6.2.0 GA DDR to ESP customers**
- ▶ **12/02/11 z/VM 6.2.0 General Availability**



z/VM 6.2 Worldwide ESP Customers

■ Industries

- ▶ Financial - 6
- ▶ Universities - 1
- ▶ Manufacturing - 2
- ▶ Government - 1
- ▶ Insurance - 2
- ▶ Transportation – 1
- ▶ Health Care – 1
- ▶ Information Processing – 2



■ Geographic

- ▶ N. America - 10
- ▶ Europe - 6



■ Processors

- ▶ z10s (11)
- ▶ z196s (10)

■ Plus

- ▶ Software Vendors - 16
- ▶ IBM Internal - 30



z/VM 6.2 Worldwide ESP Customers...

- **Major Hospital and Research Center**
 - ▶ Linux guests, CMS, CSE
- **2 Information Services Companies**
 - ▶ Linux guests, z/OS, CMS
- **Manufacturing**
 - ▶ SAP
- **6 Large Banks and Financial Services**
 - ▶ z/OS, Linux, TPF & CMS
- **1 University**
 - ▶ Oracle, Linux, CMS, z/OS, Domino
- **2 Major US Insurance Companies.**
 - ▶ Linux, CMS, z/OS
- **Large European Air Traffic Control**
 - ▶ Linux
- **German Manufacturing Co.**
 - ▶ SAP
- **German Government Agency**
 - ▶ z/OS, Linux, LDAP

General Observations

- **“LGR is the very best z/VM software enhancement since 64-bit support became available” – Mark Shackelford, Vice President, Information Services, Baldor**
- **Our Linux guests have been up longer than the z/VM hosts underneath. It worked exactly as intended: move the Linux guest over and service the z/VM LPAR, then move the guest back.”**
- **“We even like how LGR fails.”**
- **“VMRELOCATE TEST option gives all the reasons that a guest can’t be relocated, instead of stopping after the first one.”**
- **“We were happy when a relocation failed because the target system didn’t have enough page space, excellent error message.”**
- **LGR is “really a great function”**

z/VM 6.2 ESP Experiences

Dennis O'Brien

z/VM Chief Engineer

14 Mar 2012

Session 10361

Environment

- z/VM in multiple datacenters
- CMS applications, z/OS guests, Linux guests
- 2-node CSE clusters host production Linux guests
- Due to time constraints, ESP testing limited to guests.
- One first-level system running z/VM 6.2 now.

Single System Image

- Some teething pains with mismatched service levels during ESP.
- Directory syntax takes some getting used to.
- Put SFS servers on shared DASD, just in case.

Live Guest Relocation

- Beware of constrained systems (e.g. 512M guest on 128M host)
- Problems using VSWITCH under VSWITCH. Use dedicated OSA's for second-level z/VM systems that have their own virtual switches.
- Dedicated OSA's had EQID mismatch on backup OSA's.
- Tested relocation of a disconnected CMS user. It works with shared CMS disks, but is not supported.

Changes Required for LGR

- Add OPTION CHPIDV ONE.
- Detach 190, 19D, and 19E just before Linux IPL in PROFILE EXEC.
- Replace virtual disk MDISK statements with DEFINE commands.
- Make sure VDISK USERLIM is high enough for your virtual disks.

Migration Tool

- Problem with TSAFVM. No time to open ETR, yet. Could be first-level and second-level node names being the same.
- Use caution when adding second-level system to ISFC. Even if you change the node name, the old name is in the CRR server.
- If you do change the name, don't forget owner information on CP-owned DASD.

Local Tools

- Examine impact of clusters and directory changes on tools and housekeeping jobs.
- Userid starts with either USER or IDENTITY.
- Some jobs run on every member (SPOOL file purge, audit file extract, time zone change).
- Some jobs run on one member (password expiration, password sync, directory backup, tape scratch).
- Location of SYSTEM CONFIG has changed.

VM:Secure

- Commands go to master unless LOCAL prefix is used.
- This even applies to AUDITEXT and END.
- VMSECURE MAP shows SUBCONFIG name for private minidisks, but you LINK the IDENTITY name.

VM:Tape

- Scratch tapes mounted on one system don't show up in VMTAPE LIST on other systems until later. (Maybe next day. Bouncing VMTAPE or running daily housekeeping doesn't do it). Not reported to CA, yet.

VM:Backup

- Run one server per cluster to back up shared DASD, and one per member to back up private DASD.

Reserve/Release Changes

- New rules to preserve integrity. Details in CP Planning and Admin.

Finally...

- Goodbye CSE. Clusters automatically provide link protection for all shared DASD. No need to configure link protection.

Presentation requests

Marcy Cortes
Wells Fargo Bank



- I apologize for not being able to post presentations here.
- For copies of the presentation, SHARE members should please email me directly:
 - Marcy.D.Cortes@wellsfargo.com
- Please indicate which presentation
 - GDPS
 - Linux Experience
 - ESP Experience

James Vincent Nationwide Insurance

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

z/VM 6.2 – the beginning

- Installed on LPAR and second level (in that order)
- SSI used out of the box
- Tried LGM (moving a Linux server from LPAR to LPAR)
- Worked on infrastructure third-party product migration and upgrades to support 6.2
- Initial findings
 - Large amount of thought, time and effort went into it
 - Pretty darn stable (gun and foot factor applies)
 - SSI is huge (impactful)

z/VM 6.2 - why not SSI?

- We had doubts; existing environment was the square peg into the SSI round hole
 - Long conversation with the Endicott clan while at SHARE in Orlando confirmed this
- Our Linux guests resources, what they are running, and how many we have would not lend them to use LGM very well at all

z/VM 6.2 - no SSI for you (us)!

- Some of our limiting factors:
 - NPIV SAN
 - would require FCPs on **all** the LPARs
 - “exponential” complexity on wwpn’s
 - Moving a WAS or even DB2 server was problematic
 - DB2 tends to puke when with more than a few second delay
 - *HA would kick in and move all the services anyway.*
 - WAS is also very picky about pauses or delays
 - If > 5 seconds to move then the app fails-over

z/VM 6.2 - no SSI for you (us)!

- Some of our limiting factors continue:
 - Too many servers on a given LPAR
 - Large guests; time to move all would be limiting factor
 - Size (resources) on target LPARs would have to be substantial
 - Our HA setup precludes the need to move guests
 - HA servers should kick in for LPAR outages
 - Those servers without HA already have an agreement that there will be outages in their apps.
- We have baggage
 - Long time z/VM shop and 6.2 is quite different
- **Don't let this scare you off 6.2 though!!**