

z/VM 6.4 Upgrade Installation

Richard Lewis

IBM Washington System Center

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*
DFSMS/VM*	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*	
	POWERS	Tivoli Storage Manager	
		TotalStorage*	

* Registered trademarks of IBM Corporation

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

- New installation option introduced with z/VM 6.3
- What are the objectives
- What is different
- Planning
- Stage 1
- Stage 2
- Finish Upgrade Installation
- Mixed Clusters
- Summary

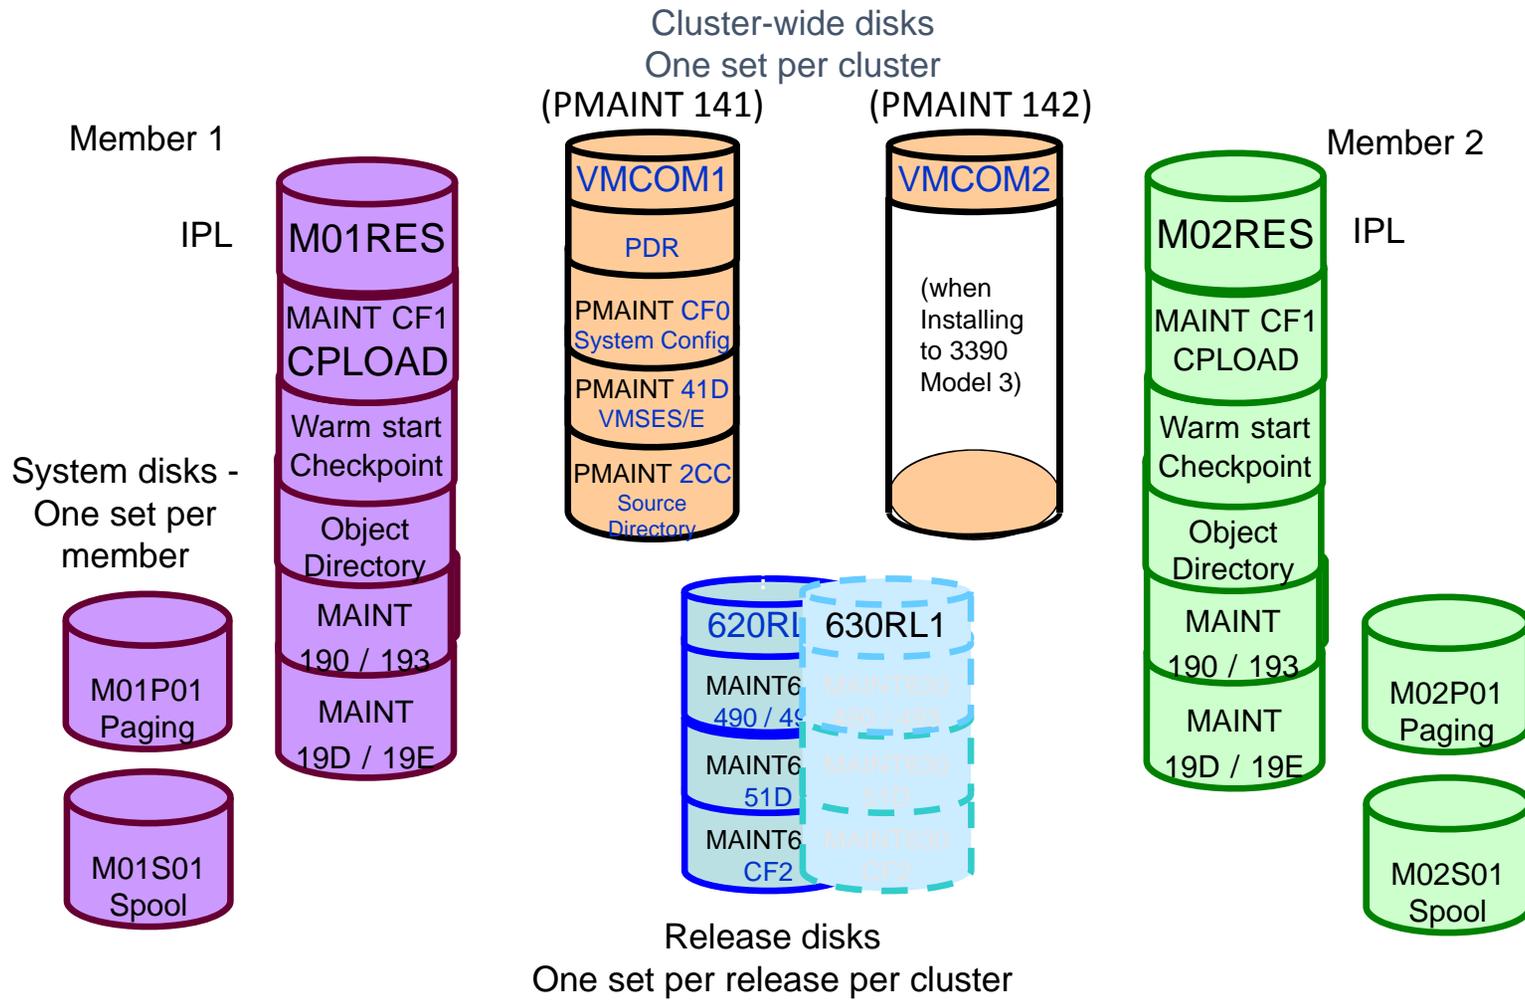
New Installation Option with z/VM 6.3

- Upgrade Installation
 - Documented in Part 4 of the z/VM 6.4 Installation Guide (GC24-6246-02)
 - Only applies to upgrading from z/VM 6.2 or 6.3
 - Traditional methods apply for migrations from z/VM 5.4 or z/VM 6.1.
 - Restrictions
 - Must not have changed identity/subconfig definitions with user entries for all entries in the initial directory shipped by IBM
 - Must not have changed IBM supplied USER or IDENTITY names
 - Must not have changed IBM supplied minidisk addresses
 - Must not have moved IBM supplied minidisks to different virtual machine definitions
 - Must not have changed the default values in VMSESE PROFILE or VMFINS DEFAULTS
 - Upgrading from a mixed-release SSI cluster is not allowed.

Objectives of Upgrade Installation

- Upgrade a z/VM 6.2 or 6.3 system with minimal impact to the current running system.
 - No manual merge of directories
 - Add new userids to current system
 - Upgrade files on existing disks under VMSES control
 - Flag local modifications for review and customized files
- z/VM 6.4 Upgrade Installation with SSI Cluster
 - All members must be at z/VM 6.2 or all members must be at z/VM 6.3
- Support upgrading single members of an SSI cluster
- Support upgrading non-SSI systems
- Support notion of mixed SSI clusters
 - Where individual members of the cluster are at different levels

DASD Volumes and Minidisks



What is Different?

- Traditional Migration
 - New release installed on separate volumes
 - Users and Data from current running system migrated to system running new release
 - IPL from new volumes
 - Keep old system around for backout
- Upgrade Installation
 - Install new release as temporary work system second level
 - Move new level of code to current system with minimal impact
 - Support upgrade of single member in a cluster
 - Create system backup for reversing upgrade

Workflow

- Plan and prepare
- Backup current system
- Install z/VM 6.4 work system from MIGMAINT userid
- Generate the stage1 changes file
- Execute stage1 changes
- Finish stage1 changes
 - This may involve reworking local modifications and reviewing the user directory, and ESM
- Backup current system again
- Generate the stage2 changes file
- Execute the stage2 changes
- Review SYSTEM CONFIG, USER DIRECT, and Upgrade Warning File
- Shutdown and IPL the upgraded system or member

Planning

- Select installation media type. Supported types are:
 - Physical DVD
 - FTP server with access to content of DVD
 - CMS-formatted minidisk
- Review and Comply with Requirements
 - General
 - Hardware
 - 6.4 Architecture Level Set to z196/z114 or later
 - System Software
 - MIGMAINT
 - DVD Installation
 - Other System Requirements such as SFS space, directory and security managers
 - If service was begun on a member in the cluster, then all members must have completed put2prod before upgrade begins
 - IDENTIFY command must return a system name which matches the system being upgraded
 - If upgrading from 6.2 or 6.3 LE 6.2.0 APAR VM65718 must be applied before upgrading
 - If upgrading from 6.2 DirMaint APAR VM65297 must be applied before upgrading
 - If sharing RACF database with other z/OS or z/VM systems see RACF Program Directory for converting database templates

Planning (cont)

- Complete Worksheets
 - Backup current system before beginning and before Stage2
 - Note location of DVD content (ftp server information, CMS minidisk address)
 - Edit user directory, or use a directory manager?
 - Security manager active?
 - Name and location of SYSTEM CONFIG file
 - Products loaded into VMPSFS filepool?
 - DASD type and model for work system
 - FBA – 5 volumes if each is over 9.7, 6 volumes if each is under 9.7G
 - 3390-03 – 11 volumes (9 if all products loaded to SFS)
 - 3390-09 – 6 volumes
 - 3390-27 – 5 volumes (6.4 only)
 - Additional member specific space
 - 3390 – 7,127 cylinders for 6.3 to 6.4, 7,147 for 6.2 to 6.4. Minimum 3390-09
 - FBA – 10,262,880 FBA blocks for 6.3 to 6.4, 10,291,680 FBA blocks for 6.2 to 6.4
 - Space used for SMAPI worker machines (VSMGUARD, VSMWORK1, VSMWORK2, VSMWORK3 6.2 to 6.4 only), DTCVSW3, DTCVSW4, TCPMAINT, PERFSVM, and OPNCLOUD

Planning (cont)

- Complete Worksheets (cont).
 - Allocate additional common space
 - Must allocate additional space on DASD type used for original z/VM 6.2 or z/VM 6.3 system
 - Only allocate this additional space for first member to be upgraded in a cluster
 - Make sure additional common space available to all members of cluster
 - 3390 – need 212 additional cylinders for upgrade to z/VM 6.4
 - FBA – 305,280 for upgrade to z/VM 6.4
 - Allocate this additional space on a volume that will be considered common if upgrading a non-SSI system (in anticipation of future conversion to SSI).

Preparation

- On your existing z/VM 6.2 or z/VM 6.3 system
 - Make sure directory manager will allow MIGMAINT to make changes without specifying a password
 - For DIRMAINT ALLOW_ASUSER_NOPASS_FROM= MIGMAINT *
 - Make sure directory manager will support disk cleanup and disk operations
 - For DIRMAINT this is DISK_CLEANUP= YES and DATAMOVE_MACHINE= DATAMOVE * *
 - Make sure MIGMAINT is authorized to use directory manager
 - For DIRMAINT this is done by adding MIGMAINT to AUTHFOR CONTROL for both 140A and 150A command sets and ADGHMOPS classes.
 - Make sure MIGMAINT is able to link the MAINTvrms 4CC minidisk
 - Make sure all of the new volumes created during installation of the work system and the new space added to the z/VM 6.2 or 6.3 system is available for automatic minidisk allocation within your directory manager
 - For DIRMAINT this is done by adding REGION statements in EXTENT CONTROL for all volumes
 - Make sure directory manager excludes MAINT640 013* overlap minidisks from free space calculations.
 - For DIRMAINT this is done by adding MAINT640 013* to EXCLUDE section of EXTENT CONTROL
 - Format, label, allocate and attach to the z/VM 6.2 or 6.3 system, volumes designated as additional space allocated to current system

Work System

- Load the install tools (instpipe along with pipe commands to load 2222 and DVDPRIME to load 24CC and 2CF0)
- Run INSTPLAN with dvd and (UPGRADE
- Upgrade installation panels will be displayed
 - Different from the normal instplan panels
 - Display current z/VM 6.2 or 6.3 information
 - Prompt for license agreement for previously installed optional features that require 6.4 license also
 - Prompt for DASD type and size for installation of work system (DASD type will match DASD type for system being upgraded, e.g. 3390).
 - Specify location of SYSTEM CONFIG, how directory is maintained, and whether RACF or a different security manager is used
 - Specify labels and addresses for z/VM 6.4 work system
 - Specify label of DASD to be used for allocation of additional space on current system
 - Note: these volumes should be attached to the z/VM 6.2 or 6.3 system prior to proceeding)
- Attach DASD devices to be used for the z/VM 6.4 system to MIGMAINT
- Execute INSTALL command

Stage 1

- When the work system has been completely installed, the INSTALL EXEC will end
 - CP Shutdown command to shutdown the work system executed after prompt for GO satisfied
 - Enter IPL CMS to restore the CMS environment to MIGMAINT
 - The work system now exists on the DASD devices attached to MIGMAINT
- The upgrade process is automated and coordinated by a new EXEC named INSTUPGR
- INSTUPGR is used to generate a table of tasks to complete
- INSTUPGR may be used to execute the tasks, or you may choose to manually execute the tasks in the table by hand.
- At each stage, the PRIME option is used to instruct INSTUPGR to generate a table of tasks.
- Tables are written to MIGMAINT's 2CF0 minidisk
- Warnings issued by INSTUPGR while building the stage1 table will be displayed and written to \$STAGE1\$ \$WRNFILE on MIGMAINT's 2CF0 minidisk
 - Review any warnings and resolve the problems before continuing
- Errors that occur are written to the INSTUPGR \$CONSLOG on MIGMAINT's 2CF0 minidisk.
 - If errors occur consult this file and resolve any errors before continuing

Stage 1 (Cont)

- The stage1 changes add new maintenance userids and updated versions of service and installation userids to your system
- Changes include:
 - Attaching the new release disks to the current system (e.g. 640RL1-3 for 3390-03)
 - Updating SYSTEM CONFIG to automatically attach these new volumes to the system during IPL
 - Add MAINT640 to the directory, and update links defined in the MAINT subconfig for this system to link to the new MAINT640 disks
 - Add 6VMPTK40, 6VMRAC40, 6VMRSC40, 6VMTCP40, and 6VMDIR40
 - DDR minidisks for new 6VMDIR40 to the additional space volume designated during INSTPLAN
 - Update LINKs for current service machines to the new install and service userids (for example PERFSVM linking to 6VMPTK40 instead of 6VMPTK20, or 6VMPTK30)
 - Enable 6VMDIR40 on z/VM 6.2 or 6.3 system
 - Add enable records to SYSTEM CONFIG file
 - Add OPNCLOUD, DTCVSW3, and DTCVSW4 to current system
 - Add new minidisks and update directory entries for SMAPI Worker Servers along with other SMAPI environment changes (migrate from 6.2 only)
 - Enroll MAINT640 as ADMIN in shared filepool servers

Stage 1 (Cont.)

- The stage 1 table is named \$STAGE1\$ \$TABLE\$ and it is stored on MIGMAINT's 2CF0
 - The table is an ordered list of actions to be taken against the current system.
 - Each entry is described by comments such that the change can be manually performed
- Execute the stage1 changes
 - Use INSTUPGR with the COMMIT option
 - INSTUPGR will generate a back out file that will allow you to reverse these changes
 - If a directory manager is being used that provides an exit to work with the IBM upgrade code the exit will be called to perform directory functions.
 - INSTUPGR is able to change the user direct file if you do not use a directory manager
 - Manually update your system with the stage1 changes
 - You must perform the changes in the order specified in the stage1 table
 - You will not have an automated back out capability
 - You must edit the stage1 table to mark each change as completed when you finish each change
 - Combination of automatic and manual changes
 - Perform the manual changes first and update the stage1 table
 - Run INSTUPGR with COMMIT to complete all changes not marked finished
 - Be careful of dependencies when following this path

Stage 1 (Cont.)

- All stage1 changes must be successfully completed before moving to Stage2.
- After completing all stage1 changes manually:
 - Run INSTUPGR stage1 (commit done)
- Post stage1 directory considerations
 - Updating a member of a cluster and directory manager is not used
 - Bring changed directory online to all cluster members with DIRECTXA command
 - Run DIRECTXA on all cluster members if directory manager did not provide an exit for INSTUPGR
 - For non-SSI, or one member SSI cluster, or first member of an SSI cluster
 - Copy configured DIRMAINT files from 6VMDIR20 or 6VMDIR30 to 6VMDIR40 (491, 492, 11F, and 41F).
 - Copy DVHPROFA DIRMSAT * files from 6VMDIR20 or 6VMDIR30 491 and 492 to 6VMDIR40 491 and 492
 - Make sure new userids such as MAINT640 are in AUTHFOR CONTROL if needed
 - Recycle all DIRMAINT servers (i.e. DIRMAINT, DIRMSATx, DATAMOVE, DATAMOVx) to begin using the 6.4 DIRMAINT code.

Stage1 (Cont.)

- Post stage1 ESM considerations
 - Only applies if you are running an ESM
 - If running ESM perform these steps prior to preceding DIRMAINT copy steps.
 - Ensure MAINT640 is authorized to:
 - Link to any minidisk on the system without a password
 - Perform security authorizations on behalf of other users
 - Perform all SFS admin functions
 - Define to ESM new userids added to directory
 - OPNCLOUD, MAINT640, 6VMDIR40, 6VMPTK40, 6VMRAC40, 6VMRSC40, 6VMTCP40, DTCVSW3, DTCVSW4
 - Make sure that the userids above have same disk access as the corresponding z/VM 6.2 or 6.3 userids.
 - Make sure userids such as RACMAINT, TCPMAINT, DIRMAINT, and MAINT are authorized to link to the new userid minidisks above instead of the 6.2 or 6.3 counterparts
 - Review table of indirect links in the z/VM 6.4 Installation Guide Chapter 21 to ensure that necessary ESM changes are made.
 - If your ESM manages SFS administrator authorizations make sure that MAINT640 and the SMAPI workers are authorized as admins for VMPSFS, VMSYS, VMSYSR, and VMSYSU
 - If your ESM manages CP command authorizations and access to other CP resources
 - Ensure that MAINT640 and 6VMxxx40 userids are properly authorized

Stage1 (Cont.)

- Rework local modifications
 - This step is not required for the 2nd – 4th member of an SSI cluster to be upgraded
 - Local modifications to components were copied to 6.4 service disks
 - 6.4 VM SYSLMOD table was updated with the local modifications copied to the 6.4 service disks
 - Review and update status of each local modification
 - This step must be performed from MAINT640
 - Use VMFUPDAT SYSLMOD
 - To see if any local modifications need rework
 - To mark local modifications as complete when reworked

Stage2

- The changes implemented in stage2 might affect your production workload
 - Best to stop your production workload
 - Also logoff PERFSVM and other users of MONDCSS
 - Logoff service machines such as OPERSYMP, EREP, DISKACNT, TCPIP (and client servers), VSWITCH controllers such as DTCVSW1 and 2, SMAPI virtual machines
 - Logon MIGMAINT from OSA ICC, or integrated 3270 icon to run stage 2
 - Backup the system prior to executing stage2
 - SFS servers should still be running along with directory manager and or security manager servers
- Stage2 moves the new release code into production
 - Replaces content on minidisks such as 190, 193, etc.
- Back out of stage2 requires restoring system backup
- INSTUPGR stage2 (prime)
 - Creates the \$STAGE2\$ \$TABLE\$
 - Table resides on MIGMAINT's 2CF0
 - Warnings placed in \$STAGE2\$ \$WRNFILE on MIGMAINT's 2CF0
 - Errors placed in INSTUPGR \$CONSLOG on MIGMAINT's 2CF0

Stage2 (Cont.)

- Move and copy mdisks for 6MVHCD20 to new common volume (6.2 to 6.4 only)
- Delete DHCPD, LPSEVE (6.2 to 6.4 only)
- Erase 6.2 or 6.3 content from minidisks with VMFERASE, then VMFCOPY from 6.4 minidisks
 - 190, 193, 19E, 19D
 - Copy from 6.4 alternate minidisks
- Erase 6.2 or 6.3 content from TCPMAINT 591, 592 and copy from 6VMTCP40 491 and 492
- Erase 6.2 or 6.3 content from PERFSVM 201 and copy from 6VMPTK40 200
- Copy customized files for VMSES, CP, DVF, REXX, AVS, GCS, TSAF, CMS, TCPIP, RSCS, DIRM, RACF, PERFKIT
 - If file has been customized copy new 6.4 file to new name on target alternate minidisk
 - For example copy 6.4 sample file to 6.2 or 6.3 TCPMAINT 491
- Copy updated VMSES inventory files
- VMFINS BUILD (SERVICED)
 - REXX, CMS, CP, GCS, DVF, AVS, TSAF, TCPIP, RSCS, DIRM, RACF, PERFTK
- BLDNUC
 - CMS, GCS, CP, RACF
- PUT2PROD
- PUT2PROD SEGMENTS ALL

Stage2 (Cont.)

- Implement stage2 changes
- Automatically
 - INSTUPGR stage2 (COMMIT
- Manually
 - Follow implementation steps written as comments in stage2 table
 - Order of items in table must be followed
 - Update table items when complete
- Manual and Automatic
 - Manually implement some tasks in table
 - Must follow order of table
 - INSTUPGR stage2 (COMMIT to automatically process all entries not marked complete

Finish Upgrade Installation

- Directory
 - If upgrading a member of an SSI cluster make sure directory is put online on all other members
 - Verify definitions added to directory comply with your local standards (i.e. minidisk passwords, etc)
- SYSTEM CONFIG
 - Review changes under “Upgrade Statements” comment at end of file
 - User_Volume statements
 - Product enable records
 - Make sure new volumes added to the 6.2 or 6.3 system before upgrade began are in appropriate User_Volume statement
 - New volume for OPNCLOUD
 - New common volume for 6VMDIR40
- \$WRNFILE Messages
 - Created on MIGMAINT 2CF0 minidisk by INSTUPGR
 - Edevice warnings
 - Customizable file warnings
 - Warnings for parts residing on PMAINT 551
 - HCD warnings

Finish Upgrade Installation (Cont).

- Delete obsolete saved segments (migration from 6.2 only)
 - helpseg
 - nslameng
 - These were shipped with 6.2 and are no longer used
 - Use purge nss name helpseg and purge nss name nlsameng commands
- Shutdown and IPL upgraded system
 - If running ESM other than RACF consult vendor documentation prior to shutdown for any special requirements before IPLing
 - If using RACF
 - Test disks updated with new level of code
 - On IPL bring up test RACF/VM server (RACMAINT)
 - Move updated code to production RACF server
- Use MIGR51D for licensed products that are not preinstalled on the system DDR
 - E.g. High Level Assembler
 - System software inventory files not updated for these products
 - Involves also merging segmap information and rebuilding segments

Finish Upgrade Installation (Cont).

- Change default passwords for new virtual machine definitions added to directory
- If you have created stand alone dump in the past it will need to be recreated
 - Create a tape using the z/VM 6.4 program
 - Use the new utility SDINST to dump to DASD
 - z/VM 6.2 release stand alone dump should not be used with z/VM 6.4
- Review links in directory for non-IBM virtual machines that might link to release specific minidisks
 - E.G. links to 6VMPTK20 or 6VMPTK30 200 and 29D for perffit usage
- Review new program directories for any additional configuration steps needed for new levels of the enabled pre-installed features
- Return work volumes to DASD pool
- If DIRMAINT is used disable old level (i.e. 6VMDIR20, or 6VMDIR30)
 - If this is first member of an SSI cluster consider manually moving the 630 or 640 level of DIRMAINT help files to other members not yet upgraded
- Create final backup of system

Remove Old Release

- Only do this step if upgraded non-SSI 6.2 or 6.3 system or completed upgrading last member of an SSI cluster
- Remove old release level userids no longer used on the 6.4 system
 - MAINT620, 6VMDIR20, 6VMPTK20, 6VMRAC20, 6VMRSC20, 6VMTCP20 or corresponding xxxxxx30 userids
 - Do not delete 6VMLEN20 or 6VMHCD20 these continue to be used in the 6.4 system
 - Remove from the user directory and shared filepool servers (if applicable)
- When old release volumes no longer contain minidisks they can be returned to DASD pool
- Update SYSTEM CONFIG
 - Remove product enable records for old pre-installed features no longer used
 - 6VMPTK20, 6VMRAC20, 6VMRSC20 or xxxxxx30 versions
 - Set old features to DISABLED on the system
 - Note: DIRMAINT (if used) was already disabled

Life In a Mixed Release Cluster

- Need to make sure that common utilities are at the highest release level in the cluster
 - Utilities maintained on PMAINT 550/551 mdisks
- Install new service on z/VM 6.2 members to support highest release level processing
 - VM65317 - VMSES
 - UM34006 – not currently on any RSU
 - <http://www.vm.ibm.com/service/tips/vm65317.pdf> (for pubs changes)
 - VM65318 - CP
 - UM34029 – on RSU 1302
 - VM65319 - CMS
 - UM34030 – on RSU 1302
 - VM65320 – RACF VM
 - UV61196 – not currently on any RSU

Summary

- SSI clusters with more than one member represent a challenge for upgrading using traditional methods
- New Upgrade Installation path provides capability to upgrade an existing member of a cluster
 - Minimal disruption to existing system
 - Ensures that cluster wide utilities and directory manager are at new release level for entire cluster
- Upgrade Installation path provides a quick and easy method of upgrading non-SSI and single member SSI clusters as well
- Upgrade Installation currently only applies to upgrading z/VM 6.2 or z/VM 6.3 systems to z/VM 6.4
- Minimal downtime of production workload