


**Abstract: z/OS 1.11: Migration Actions Part 2 of 2**

This is part two of a two-part session that will be of interest to System Programmers and their managers who are migrating to z/OS 1.11 from either z/OS 1.9 or z/OS 1.10. It is strongly recommended that you attend both sessions for a complete migration picture.

In part two, the speaker will cover the specific migration actions for getting to z/OS 1.11. Changes to selected elements such as the BCP, Communications Server, Distributed File Service (DFS), DFSMS, Language Environment, and z/OS UNIX System Services will be included.

*The general availability date for z/OS V1 R11 was September 25, 2009.*

Systems and Technology Group 

## Trademarks

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

AFP	DFSMSHsm	FICON	MVS/SP	PFOSM	VTAM
ADK	DFSMSHsm	HiperStrike	NetSpool	Print Services Facility	WebSphere
AspIlet	DFSORT	IBM	NetView	PrintWay	zArchitecture
Base Mail manager	ESDA	IBMLink	Open Class	RACF	zOS
CICS	ESCON	IMS	OS/390	Resource Link	zVM
DB2	Eservr	Integrity	Parallel Sysplex	RMF	zSeries
DB2 Universal Database	FICON	Language Environment		SSSD	
DFSMS/dfp		Multiple		SSSD Parallel Enterprise Server	
DFSMS/dfs		MVS		SpecIca Timer	
		MVS/ESA		SystemPac	
				Systemz9	
				ThinL	
				TotalStorage	
				VisualAge	
				VSE/ESA	

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

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

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

\* All other products may be trademarks or registered trademarks of their respective companies.

**Notes:**

Performance is an inherent throughout their ITIQ suite based on measurements and judgments using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon configurations such as the amount of data, the user's hardware, the I/O configuration, the channel configuration, and the channel processor. Therefore, an accurate view of individual user will require throughput measurements customized to the performance values stated here.


IBM hardware products are manufactured from new parts, or new and reworkable used parts. Reproduction, use, assembly, repair, and other customer activities may be affected by the availability of the parts 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 internationally, and the information may be subject to change without notice. Contact your local IBM branch 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 these products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Decisions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

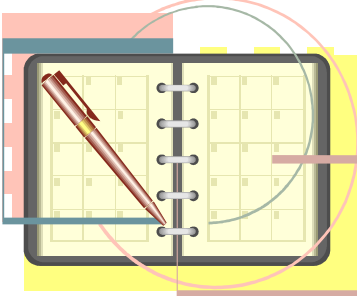
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.


 ENABLING BUSINESS. A THROUGH Z.

© 2009 IBM Corporation

## z/OS 1.11: Migration - Part 2 Agenda

- **Migration Actions**
  - General Migration Actions for Everyone
  - Element-Specific Migration Actions for z/OS R11
    - BCP
    - Communications Server
    - DFSMS
    - Distributed File Service
    - JES2 and JES3
    - z/OS UNIX
    - Various other elements (ICSF, HCM, IBM TDS)
- **Summary**








39 | z/OS 1.11: Migration - It's Up... Part 1 of 2 | © 2010 IBM Corporation

## Definition of "Migration Action":

- ▶ **Migration is**
  - the installation of a new version or release of a program to replace an earlier version or release.
- ▶ **After a successful migration, the applications and resources on the new system function the same way they did on the old system, if possible.**
- ▶ **Migration does *not* include exploitation of new functions except for new functions that are now required to use. Typically, migration is the first step, Exploitation is the second step.**
- ▶ **Migration actions are classified as:**
  - **Required:** required for all users
  - **Required-IF:** only required in certain cases
  - **Recommended:** good to do because it 1) may be required in the future, 2) resolves performance or usability problem 3) improves migration workload.
- ▶ **Migration actions are also classified as when they may be performed:**
  - **NOW, Pre-First IPL, or Post-First IPL**




**Means "don't overlook!"**


**Means some programmatic assistance is available**

### Migration Definitions and Classifications

Migration is the first of two stages in upgrading to a new release of z/OS. The two stages are:

- **Stage 1: Migration.** During this stage you install your new system with the objective of making it functionally compatible with the previous system. After a successful migration, the applications and resources on the new system function the same way (or similar to the way) they did on the old system or, if that is not possible, in a way that accommodates the new system differences so that existing workloads can continue to run. Migration does not include exploitation of new functions except for new functions that are now required.
- **Stage 2: Exploitation.** During this stage you do whatever customizing and programming are necessary to take advantage of (exploit) the enhancements available in the new release. Exploitation follows migration.

### Migration Requirement Classification and Timing

The migration actions are classified as to their requirement status:

- **Required.** The migration action is required in all cases.
- **Required-IF.** The migration action is required only in a certain case. Most of the migration actions in this presentation are in this category.
- **Recommended.** The migration action is not required but is recommended because it is a good programming practice, because it will be required in the future, or because it resolves unacceptable system behavior (such as poor usability or poor performance) even though resolution might require a change in behavior.

To identify the timing of migration actions, this presentation uses three types of headings:

- **Now.** These are migration actions that you perform on your current system, either because they require the current system or because they are possible on the current system. You don't need the z/OS V1R11 level of code to make these changes, and the changes don't require the z/OS V1R11 level of code to run once they are made. Examples are installing coexistence and fallback PTFs on your current system, discontinuing use of hardware or software that will no longer be supported, and starting to use existing functions that were optional on prior releases but required in z/OS V1R11.
- **Pre-First IPL.** These are migration actions that you perform after you've installed z/OS V1R11 but before the first time you IPL. These actions require the z/OS V1R11 level of code to be installed but don't require it to be active. That is, you need the z/OS V1R11 programs, utilities, and samples in order to perform the migration actions, but the z/OS V1R11 system does not have to be IPLed in order for the programs to run. Examples are running sysplex utilities and updating the RACF database template.

It is possible to perform some of the migration actions in this category even earlier. If you prepare a system on which you will install z/OS V1R11 by making a clone of your old system, you can perform migration actions that involve customization data on this newly prepared system before installing z/OS V1R11 on it. Examples of such migration actions are updating configuration files and updating automation scripts.

- **Post-First IPL.** These are migration actions that you can perform only after you've IPLed z/OS V1R11. You need a running z/OS V1R11 system to perform these actions. An example is issuing RACF commands related to new functions. Note that the term "first IPL" does not mean that you have to perform these actions after the very first IPL, but rather that you need z/OS V1R11 to be active to perform the task. You might perform the task quite a while after the first IPL.

Icons used in the subsequent foils:



means that you shouldn't overlook this migration action.



means that a Health Check (using the IBM Health Checker for z/OS function) or the older IBM Migration Checker for z/OS download tool can help you with this migration action.

## Elements with Migration Actions for z/OS R11

► Documented in *z/OS VIR11 Migration* (refresh is GA22-7499-16, January 2010)

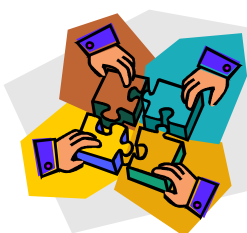
- For complete migration tasks for z/OS R11, see this book!
  - from R9 to R11, and R10 to R11, customized books

► **From z/OS R9 to z/OS R11:**

<ul style="list-style-type: none"> <li>➤ BCP</li> <li>▪ CIM</li> <li>➤ Communications Server</li> <li>➤ Cryptographic Services – ICSF and OCSF</li> <li>➤ DFSMS</li> <li>▪ DFSORT</li> <li>➤ Distributed File Service – zFS and SMB</li> <li>▪ HCD</li> <li>➤ HCM</li> <li>▪ IBM Tivoli Directory Server</li> <li>▪ ICKDSF</li> <li>• Infoprint Server</li> <li>➤ IBM Tivoli Directory Server</li> </ul>	<ul style="list-style-type: none"> <li>▪ ISPF</li> <li>➤ JES2</li> <li>➤ JES3</li> <li>➤ Language Environment</li> <li>• Library Server</li> <li>• NFS</li> <li>• RMF</li> <li>▪ SDSF</li> <li>▪ Security Server (RACF)</li> <li>▪ SMP/E</li> <li>▪ TSO/E</li> <li>▪ XL C/C++</li> <li>➤ z/OS UNIX</li> </ul>
--	---

➤ means that some of that element's migration actions will be discussed

- means that you need to look for that element's migration actions in the book




6 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

### Migration Actions for Elements Between z/OS R9 and z/OS R11

When migrating from z/OS R9 to z/OS R11, the specified elements in the slide above have required migration actions. Refer to *z/OS R11 Migration* for complete information on the required migration actions for all elements. Some migration actions for selected elements follow in this presentation.

If you are migrating from z/OS R10 to R11, there are fewer migration actions that are applicable. Use the customized *z/OS R11 Migration* book to see just the migration actions appropriate for this path.

## General Migration Actions for z/OS R11 from z/OS R9



► **Migration Actions You Can Do NOW:**


- **Install coexistence and fallback PTFs (Required)** - Refer to previous session for the list and a handy way to verify this programmatically (with FIXCAT `IBM.Coexistence.z/OS.V1R11` or using EPSPT extract files)
- **Add or change volumes to keep file systems in a single data set (Recommend)**
  - z/OS version root (and subsystems!) file systems are large. One 3390-3 is too small.

► **Migration Actions Pre-First IPL:**

- **Verify that virtual storage limits are set properly (Required)**
  - Default MEMLIMIT changed in z/OS R10 from zero to 2 GB. If you want to use the default (2GB), then you may want to remove your SMFPRMxx setting of MEMLIMIT.
  - The following programs ignore (by default) region and MEMLIMIT values specified through (or affected by) the IEFUSI exit in z/OS R10: HASJES20 (JES2), IATINTK (JES3), and IXGBLF01 (system logger initialization).
    - The NOHONORIEFUSIREGION keyword, new in z/OS V1R10, is set for those programs in the IBM-defined default program properties table (PPT).

7 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## General Migration Actions for z/OS R11 from z/OS R9



**Migration Actions Pre-First IPL:**

- **Remove references to deleted data sets and path (Required)**
  - R10 DFSORT ISPF panel removal deletes some SYSI.SICE\* data sets (may have logon proc impacts). Nothing removed from R11!
- **Add references to new data sets and paths (Required)**
  - New in R11: TCPIP.SEZAEXEC (concat to current REXX libraries), and SYSI.SCRR3BND and SCRR6BND (Metal C Runtime)
  - New in R10: SDSF data sets (ISF.SISFMOD1 needs to be in the link list and APF-authorized)!
- **Accommodate new address spaces (Recommended)**
  - New in R11: **HWIBCP11** (for BCPI1). Automatically started, can be manually restarted with a procedure. Will not initialize unless customization is complete.
- **Accommodate new SCOPE=COMMON data spaces in IEASYSxx MAXCAD (Required-IF)**
  - In R11: none. In R10: One data space for XES. Some XES common storage control blocks were moved to this new data space.
  - MAXCAD default is 50 (as of R6). IBM Health Checker for z/OS can help you determine what to specify for the MAXCAD value!

8 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## General Migration Actions for z/OS R11 from z/OS R9

### Migration Actions Pre-First IPL:


- **Verify SMP/E utility parameters to support assembling JES2 and SDSF (Required, as of R11)**  

```

REP UTILITY (ASMA90) NAME (ASMA90)
PARM (DECK,NOBJECT,GOFF,LIST(133),OPTABLE(UNI))
PRINT (ASMPRINT) /* see notes */ RC(0) .

```

  - DECK,NOBJECT must be there, GOFF,LIST(133) is for JES2, OPTABLE(UNI) is for SDSF.
  - ServerPac at Package Version 22.10.19 and higher has this done.
- **Accept the CustomPac Installation Dialog default file system of zFS (Recommended)**
  - Default has moved from HFS to zFS. Even in your saved configuration!
  - CHANGE DSNTYPE ZFS HFS dialog command – now and later.
- **Discontinue use of IBM Virtualization Engine Enterprise Workload Manager for z/OS (5655-M76 and 5655-EWM) (Required, as of R11)**
  - Deleted during installation of z/OS R11.
  - No product replacement, but can contact IBM STG Lab Services for more information about transitioning from EWLM V2.1 to an offering.



9
z/OS 1.11: Migration - ...And It's Good! Part 2 of 2
© 2010 IBM Corporation

### General Migration Actions Between z/OS V1 R9 and z/OS V1 R11

These migration actions were taken from *z/OS R11 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS R11 Migration*.

### General Migration Actions You Can Do Now



#### Install coexistence and fallback PTFs (Required)

**Migration action:** Install coexistence and fallback PTFs on your systems to allow those systems to coexist with z/OS V1R11 systems during your migration, and allow backout from z/OS V1R11 if necessary. Using the new SMP/E V3R5 FIXCAT support (the category for coexistence with z/OS R11 is **IBM.Coexistence.z/OS.V1R11**), or the EPSPT tool with the PSP extract files can make this task much easier. The handout for the prior session contains a list of required coexistence PTFs which was current at the time of the presentation.



#### Add or change volumes to keep your z/OS root file system in a single data set (Recommended)

*Recommended for ease of management if your z/OS root (or subsystem) file system resides on a 3390-3 volume (or another DASD volume that is close to the 3390-3 limit of 3339 cylinders).*

Because of enhancements and PTF service, the z/OS root file system (or “version root file system”) and subsystem file systems continue to grow in size from release to release. As of z/OS V1R11, the size of the z/OS root file system, whether HFS or zFS, was approximately 3100 3390 cylinders. This is closely approaching the 3390-3 size limit of 3339 cylinders. It is advisable to have the z/OS root file system and subsystem file systems within a single data set for ease of management.

**Migration action:** To keep the z/OS root file system in a single data set, do one of the following:

- Move your z/OS root file system to a larger DASD volume geometry.
- Use multiple volumes for the z/OS root file system data set.

If your z/OS root data set cannot fit on the volume or volumes you have defined for it, divide the z/OS root, with the smaller file systems being managed together.

Remember that all systems to which you deploy the z/OS root file system need sufficient DASD space to hold the z/OS root.

**Tips:**

- File systems for subsystems and products other than the z/OS product itself might also increase in size. When examining the volume for how much space your z/OS file system is using, check other product file system sizes too.
- Use IBM Health Checker for z/OS check CHECK(IBMUSS,ZOSMIGREC\_ROOT\_FS\_SIZE) to determine whether a volume has enough space for the z/OS root file system.
- This capability is also available in z/OS V1R9 with APAR OA28684 (PTF UA49361) and z/OS V1R10 with APAR OA28684 (PTF UA49362).

**Verify SMP/E utility parameters to support assembling JES2 and SDSF (Required, as of R11)**

The assembler utility parameters must be updated to support assembling of JES2 and SDSF source modules.

**Migration action:** Ensure you have the following parameters for your assembler utility in the SMP/E OPTION entry that will be used to install and maintain z/OS:

```
REP UTILITY (ASMA90) NAME (ASMA90)
  PARM (DECK,NOBJECT,GOFF,LIST(133),OPTABLE(UNI))
  PRINT(ASMPRINT) /* see below */
RC(0).
```

- DECK, NOBJECT must always be specified in the PARM list. For example, PARM(DECK,NOBJECT).
- As of z/OS V1R11, JES2 modules require GOFF and LIST(133) to assemble correctly. For example, PARM(DECK,NOBJECT,GOFF,LIST(133)).
- When you specify LIST(133), then you must either:
  - specify a unique output file to direct assembler output, for example, PRINT(ASMPRINT). You must then provide a DDDEF in your target and DLIB zones (or a JCL DD card) for ASMPRINT specifying SYSOUT, or
  - provide a JCL DD card such as //SYSPRINT DD SYSOUT=\*,LRECL=133,RECFM=FBA.
- SDSF modules require OPTABLE(UNI) to assembler correctly. OPTABLE(ESA) is not acceptable.

**Note:** ServerPacs produced at Package Version 22.10.19 and higher have the assembler utility parameters updated in the SMP/E OPTION provided.

**Use SOFTCAP to identify the effect of capacity changes (Recommended)**

*Not required, but is recommended to help in assessing processor capacity and available resources when migrating to new software levels, and when migrating to z/Architecture.*

**Migration action:**

- Download SoftCap from one of the following Web sites:
  - Customers: <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS268>
  - Business partners: <http://partners.boulder.ibm.com/src/atsmastr.nsf/Web/Techdocs>. Note that this requires an ID on PartnerWorld®. Run SoftCap to determine your expected increase in CPU utilization (if any) and to identify your storage requirements, such as how much storage is needed to IPL.

**Reference information:** *SoftCap User's Guide*, which is provided with the tool.

**Upgrade Windows 2000, 95, 98, and NT clients (Recommended)**

*Recommended because z/OS does not support service for client operating systems whose service is withdrawn by the operating system manufacturer.*

z/OS no longer supports service for client operating systems whose service is withdrawn by the operating system manufacturer. As a result, IBM no longer supports service for clients running Windows 2000, Windows 95, Windows 98, or Windows NT Workstation 4.xx.

**Migration action:** Use a supported follow-on to Windows 2000, Windows 95, Windows 98, or Windows NT Workstation 4.xx.

**Reference information:** For client software supported with z/OS, see *z/OS Planning for Installation*.





### **Verify that you have enough XCF groups in your CDS and enough XCF members in your XCF groups (Recommended)**

*Recommended to ensure you have an adequate number of XCF groups and members formatted in your sysplex couple data sets.*

1. Issue the DISPLAY XCF,COUPLE command on your current system. Notice the values of MAXGROUP and PEAK for your sysplex couple data sets. These values show you the maximum number of XCF groups that the couple data sets can support, and the peak number of XCF groups ever in use in the sysplex. Also notice the values of MAXMEMBER and PEAK for your sysplex couple data sets. These values show you the maximum number of members that the couple data set can support in one group, and the greatest number of members ever in use in the largest group in the sysplex.
2. If your peak member value is close to the maximum member value, you might want to reformat your sysplex couple data sets to support a larger maximum number of members to be used by any one group.

**Tip:** Use IBM Health Checker for z/OS check XCF\_SYSPLEX\_CDS\_CAPACITY, which checks the adequacy of the number of groups, members, and systems for which a sysplex CDS is formatted.

### **Discontinue use of IBM Virtualization Engine Enterprise Workload Manager for z/OS (Required-IF)**

*Required if you use either the IBM Virtualization Engine Enterprise Workload Manager for z/OS products.*

z/OS V1R11 does not support: IBM Virtualization Engine Enterprise Workload Manager for z/OS V1.1.0 (5655-M76) or IBM Virtualization Engine Enterprise Workload Manager for z/OS V2.1.0 (5655-EWM). If you are using either of these products, they will be deleted during installation of z/OS V1R11.

**Migration action:** IBM Virtualization Engine Enterprise Workload Manager (EWLM) V2.1 (5733-EWM) has no replacement. IBM is providing a transition for EWLM V2.1 clients to an IBM STG Lab Services-based offering. Contact STG Lab Services (stgls@us.ibm.com) for more information. In addition, IBM Tivoli offers workload automation and monitoring solutions with the Tivoli Workload Automation Family of products:

- Tivoli Dynamic Workload Broker and Tivoli Workload Scheduler
- Tivoli Workload Scheduler
- Tivoli Composite Application Manager for Response Time Tracking

If you want to preserve your existing copy of IBM Virtualization Engine Enterprise Workload Manager for z/OS V1R10 and EWLM V2.1, you can create a back up. However, keep in mind these products are no longer service supported.



### **Accept the CustomPac Installation Dialog default file system of zFS (Recommended)**

*Recommended if you want to continue to use HFS data sets for your CustomPac order.*

For hierarchical files that are switchable between zFS and HFS, the CustomPac Installation Dialog now sets the default to zFS. This is another step in the implementation of zFS as the strategic z/OS file system. Note that all file systems are still delivered as HFSs.

#### **Migration action:**

- For any file systems that you would like to be HFS (and not zFS) use the CHANGE DSNTYPE ZFS HFS command within the CustomPac Installation dialog to convert the file systems to HFS. Note that any HFS data sets in your saved configuration will be changed to ZFS as of z/OS V1R11, and if you want them to remain as HFS you must issue the command above.
- Subsequent saved configurations used after z/OS V1R11 will also be affected, and saved HFS file systems will be converted to ZFS unless you change the file system type to HFS.

## **General Migration Actions Pre-First IPL**

### **Set up your IPCS environment (Required)**

**Migration action:** Set up an IPCS environment. For guidance, use the documents listed in the reference information below. During setup, ensure that your logon procedure points to the target system's level of IPCS data sets, which are shown in *z/OS Migration*.

**Reference information:** For more information about IPCS, see *z/OS MVS IPCS Customization*. For more information about the correct logon procedure updates, see the *z/OS Program Directory*. For information about

setting up the JES2 IPCS environment, see *z/OS JES2 Diagnosis*. For information about setting up the JES3 IPCS environment, see *z/OS JES3 Diagnosis*.

### **Use IBM-supplied PARMLIB and PROCLIB (Required)**

**Migration action:** For parmlib, add the data set pointed to by the z/OS V1R11 PARMLIB DDDEF to your parmlib concatenation. The data set should generally be added last in the concatenation, and you should make sure that the other data sets in the concatenation don't have members with the same names as IBM-supplied members. If you place the data set on the system residence volume and use an indirect catalog entry, future migrations won't require this particular migration step.

- For proclib:
  - Ensure that the default proclib members have been copied to your default proclib to pick up the new and changed members. An example of a change is the following: the LLA startup procedure was updated in z/OS V1R10 (and z/OS V1R9 by APAR OA21932) to specify REGION=0M.
  - Update individual sample members provided and ensure they are accessible to the system, as shown in the table of proclib member updates in *z/OS Program Directory*.
  - Ensure that the procedure libraries listed in the table of libraries to be added to the proclib concatenation in *z/OS Program Directory* have been placed in the necessary procedure library concatenations and are available to the system.

**Reference information:** For lists of parmlib and proclib members that are shipped, see *z/OS Program Directory*.

### **Migrate /etc and /var system control files (Required)**

**Migration action:** The /etc and /var directories contain system control files: the /etc directory contains customization data that you maintain and the /var directory contains customization data that IBM maintains. During installation, subdirectories of /etc and /var are created. If you install z/OS using ServerPac, some files are loaded into /etc and /var due to the customization performed in ServerPac. You have to merge the files in /etc and /var with those on your previous system. If you install z/OS using CBPDO, you should copy the files from your old system to the z/OS V1R11 /etc and /var subdirectories.

Copy files from your old system to the z/OS V1R11 /etc and /var subdirectories, and then modify the files as necessary to reflect z/OS V1R11 requirements. If you have other files under your existing /var directory, then you will have to merge the old and new files under /var. The easiest way to do this is to create a copy of your current /var HFS and then copy the new /var files into the copy.

The following z/OS V1R11 elements and features use /etc:

- CIM
- Communications Server – IP
- Cryptographic Services – PKI Services and System SSL
- DCE Base Services
- DFSMSrmm
- Distributed File Service. The SMB server uses /etc/dfs.
- IBM HTTP Server
- IBM Tivoli Directory Server
- Infoprint Server uses /etc/Printsrv.
- Integrated Security Services - Firewall Technologies and Network Authentication Service
- Library Server
- z/OS UNIX System Services

The following z/OS V1R11 elements and features use /var:

- Cryptographic Services – OCSF
- DFSMSrmm
- IBM Tivoli Directory Server
- Infoprint Server
- Integrated Security Services - Network Authentication Service uses /var/skrb.

**Reference information:** For information about copying your existing /etc and /var directories, see *z/OS Migration*.



### **Verify that virtual storage (MEMLIMIT) is set properly (Required)**

**Migration action:** Determine how much virtual storage use to allow above the 2 GB bar. While there is no practical limit to the number of virtual addresses an address space can request above the bar, the system can limit the amount of virtual storage above the bar that an address space is allowed to use. The amount of virtual storage

above the bar is determined as follows. The MEMLIMIT parameter in parmlib member SMFPRMxx sets the default system-wide limit, which defaults to 2 GB as of z/OS V1R10 (and zero prior to z/OS V1R10). However, the system-wide default MEMLIMIT can be overridden by specifying REGION=0M or MEMLIMIT on JOB or EXEC statements in JCL. To set a limit on the use of virtual storage above the bar, use the SMF exit IEFUSI. For more information, see *Limiting the use of memory objects in z/OS MVS Programming: Extended Addressability Guide*.

If you want to control the use of virtual storage above the 2 GB bar, do one or more of the following:

- For MEMLIMIT prior to z/OS V1R10, you must specify a nonzero MEMLIMIT in an active SMFPRMxx member of parmlib to establish a system default other than zero for available virtual storage above 2 GB. (The default MEMLIMIT is zero prior to z/OS V1R10.) As of z/OS V1R10, the MEMLIMIT default is 2 GB. If this 2 GB default value is acceptable to you, no change to SMFPRMxx is necessary.
- You can specify MEMLIMIT explicitly in JCL to override the system default that was set (or allowed to default) in SMFPRMxx.
- You can specify REGION=0M on the job statement in JCL to implicitly set MEMLIMIT to NOLIMIT, which also overrides the system default (from SMFPRMxx).
- You can use IEFUSI both to establish a system default MEMLIMIT for different classes of work (for example, job, TSO, STC) and limit the amount of virtual storage that can be used above the bar, provided that an explicit or implicit nonzero MEMLIMIT is in effect from JCL or SMFPRMxx.

**Note:** By default, region and MEMLIMIT values specified through, or otherwise affected by, the IEFUSI exit are ignored in z/OS V1R10 for the following programs: HASJES20 (JES2), IATINTK (JES3), and IXGBLF01 (system logger initialization). The reason is that the HONORIEFUSIREGION | NOHONORIEFUSIREGION keyword, new in z/OS V1R10, is set to NOHONORIEFUSIREGION for those programs in the IBM-defined default program properties table (PPT).

**Tip:** Use IBM Health Checker for z/OS to help determine whether your virtual storage limits are set properly. The check RSM\_MEMLIMIT checks the current setting for the MEMLIMIT parameter in SMFPRMxx, which affects the amount of virtual storage above the 2 GB bar that is available to jobs. This check verifies that a nonzero MEMLIMIT value is in use.

**Reference information:** Information about how to evaluate the central storage configuration can be found in the Washington Systems Center white paper *z/OS Performance: Managing Processor Storage in a 64-bit Environment - V1* at <http://www.ibm.com/support/techdocs> (Search for "WP100269".)

### **Back virtual storage with real and auxiliary storage (Required)**

**Migration action:** As you exploit additional virtual storage by defining additional address spaces or by exploiting memory objects, ensure that you have defined sufficient real and auxiliary storage. Review real storage concentration indicators via an RMF report to evaluate if additional real or auxiliary storage is needed:

- Check UIC and average available frames.
- Check demand page rates.
- Check the percentage of auxiliary slots in use.

**Reference information:** For more information about memory objects, see *z/OS MVS Programming: Extended Addressability Guide* and Washington Systems Center flash 10165 at <http://www.ibm.com/support/techdocs>. (Search for "flash10165".)

### **Remove references to deleted data sets and path (Required)**

**Migration action:** Using the table in *z/OS Migration*, "Data sets and paths deleted from z/OS V1R11, and R10" as a guide, remove references to data sets and paths that no longer exist. Remove the references from the following places:

- Parmlib
- Proclib
- Logon procedures (especially the removed z/OS R10 DFSORT libraries, which supported the ISPF panels)
- Catalogs
- Security definitions, including program control definitions
- DFSMS ACS routines
- /etc/profile
- SMP/E DDDEF entry

- Backup and recovery procedures, as well as any references to them In the table, the high-level qualifiers in the data set names are the default qualifiers.

**Note:** Do not remove any data sets, paths, or references that are needed by earlier-level systems until those systems no longer need them, and you are sure you won't need them for fallback.

**Reference information:** *z/OS Migration* contains the list of all removed data sets and paths in z/OS R10. Note that no data sets or paths were removed in z/OS R11.

### **Add references to new data sets (Required)**

**Migration action:** Using the lists that are found in *z/OS Migration* as a guide, add references in the following places for data sets that have been added to z/OS:

- Parmlib
- Proclib
- Logon procedures
- Catalogs
- Security definitions, including program control definitions
- DFSMS ACS routines
- Any backup and recovery procedures.

Of special note are data sets:

- **TCPIP.SEZAEEXEC** (in z/OS R11, for Communications Server). This data set contains REXX samples that you can concatenate to the current REXX libraries.
- **ISF.SISFMOD1** (in z/OS R10, for SDSF). This new data set must be added to the link list and must be APF-authorized.

**Reference information:** *z/OS Migration* contains the list of all new data sets and paths in z/OS R11 and R10.

### **Accommodate new address spaces (Recommended)**

*Not required, but recommended to keep interested personnel aware of changes in the system and to ensure that your MAXUSER value in parmliib member IEASYSxx is adequate.*

There is one new address space in z/OS V1R11. Its name is HWIBCPii. The Base Control Program Internal Interface (BCPii) function allows authorized applications to query, change, and perform basic operational procedures against the installed System z hardware base. As a base BCP component, the new BCPii address space allows authorized z/OS applications to access the local support element and other support elements in the HMC network to perform HMC-like functions. Connectivity to the Hardware Management Console (HMC) network is done internally, without the requirement of connecting the HMC network to an intranet or the Internet.

The HWIBCPii address space is started automatically during z/OS initialization and can be manually restarted with a procedure. When z/OS is IPLed, HWIBCPii attempts to initialize and connect to the local CPC's support element. If the proper configuration for BCPii has been completed (as documented in *z/OS MVS Programming: Callable Services for High-Level Languages*), the address space will become active and able to accept requests from BCPii applications. However, if basic configuration steps have not been done, the address space will not be able to initialize properly. A few BCPii error messages will appear on the console when these steps have not been taken. If the BCPii address space is not required by the installation, then these error messages can be ignored. These messages include:

```
HWI016I THE BCPII COMMUNICATION RECOVERY ENVIRONMENT IS NOW ESTABLISHED.
HWI007I BCPII IS ATTEMPTING COMMUNICATION WITH THE LOCAL CENTRAL PROCESSOR
COMPLEX (CPC) .
HWI014I THE SNMP COMMUNITY NAME FOR BCPII IN THE SECURITY PRODUCT FOR THE
LOCAL CPC IS EITHER NOT DEFINED OR IS INCORRECT.
HWI006I BCPII ADDRESS SPACE HAS ENDED.
```

The MAXUSER value in parmliib member IEASYSxx specifies a value that the system uses to limit the number of jobs and started tasks that can run concurrently during a given IPL. You might want to increase your MAXUSER value to take new address spaces into account. (A modest overspecification of MAXUSER should not hurt system performance. The number of total address spaces is the sum of M/S, TS USERS, SYSAS, and INITS. If you change your MAXUSER value, you must re-IPL to make the change effective.)

### **Rework and install user modifications (Required-IF)**

*Required if you have made any user modifications that necessitate changes.*

**Migration action:** Use the z/OS SMP/E Planning Migration Assistant to help determine which user modifications need to be reworked and which just have to be reinstalled. The Top or New Intermediate Product Migration

Changes Report uses data found on your system, combined with IBM-supplied information from the Software Information Base, to show you the current levels of products available as well as product migration and functional changes using a comparison of FMIDs. You can use this report to determine the product migration impacts by reviewing the “changed” FMIDs. This can help you assess how many user modifications have to be reworked if you issued the LIST SYSMOD USERMOD FORFMID (listing the “changed” FMIDs) command. All other user modifications can be reinstalled without having to be reworked.

**Note:** IBM recommends using exit routines for any user modifications where possible, and installing the exit routines with SMP/E. By using SMP/E, it is easier to bring forward desired modifications to the z/OS release you are installing.

Several elements and features have their default options set by assembling and link editing one or more modules. These include:

- XL C/C++
- DFSORT. Consider using ICEPRMxx parmlib members, introduced in z/OS V1R10, to eliminate the assembler language installation option modules
- HLASM
- ISPF (specifically, the ISPF configuration table)
- Language Environment. Investigate using CEEROPT, which can be used to specify run-time options for CICS, IMS LRR, and other LRR users. Even better, consider using the function added in z/OS R7 to eliminate your assembler language run-time option modules in CEEPRMxx parmlib member!
- SDSF (ISFPARMS customization). Use dynamic statements for ISFPARMS to avoid reassembly!

### **Reconnect subsystems and non-IBM products (Required-IF)**

*Required if you use any ISV products and need to reconnect them after performing a ServerPac installation, or if you intend to use any subsystems with your z/OS system.*

**Migration action:** Follow the instructions for each ISV product that you use to reconnect it to your z/OS V1R11 ServerPac.

Ensure that any required service is installed prior to using the subsystem with the new z/OS V1R11 system, as well as any required SVCs, system modifications, parmlib setup, and proclib setup. Follow the instructions for the subsystem that you need to reconnect.

**Reference information:** For a list of independent software vendors (ISVs) that support z/OS, as well as announcements, testimonials, and other information, see <http://www.ibm.com/eserver/zseries/solutions/s390da/>. For a directory of ISV products that support z/OS, see the Global Solutions Directory at <http://www.ibm.com/software/solutions/isv>.

### **Update operational and other procedures (Required)**

**Migration action:** Review your operation, automation, administration, security, backup, and recovery procedures, and make any necessary changes depending on how you installed and which functions you plan to exploit. Some possible changes are:

- Allowing applicable users access to new high-level qualifiers that you may have. There are no new default high-level qualifiers introduced since z/OS R9.
- Updating and testing your backup and recovery procedures to accommodate the new target system.
- Updating and testing any disaster recovery procedures.
- Updating and testing any automation procedures to take advantage of new functions.
- Updating security system definitions, such as defining new users and resources, permitting users to use new resources, and defining new profiles in the RACF FACILITY class.

**Reference information:** For information about the new functions incorporated into z/OS V1R11, see z/OS *Introduction and Release Guide*.

### **Update your check customization for modified IBM Health Checker for z/OS checks (Recommended)**

*Recommended to ensure that your checks continue to work as you intend them to work.*

Changes that IBM makes to the checks provided by IBM Health Checker for z/OS can affect any updates you might have made.

The checks that were changed by IBM in z/OS V1R11 are: CSV\_LNKLST\_SPACE, HSM\_CDSB\_BACKUP\_COPIES, HSM\_CDSB\_DASD\_BACKUPS, HSM\_CDSB\_VALID\_BACKUPS, USS\_PARMLIB, USS\_PARMLIB\_MOUNTS, and XCF\_FDI.

The checks that were changed by IBM in z/OS V1R10 are: IBMRACF,RACF\_SENSITIVE\_RESOURCES, IBMXCF,XCF\_CDS\_SEPARATION, IBMXCF,XCF\_CF\_STR\_PREFLIST, IBMXCF,XCF\_SFM\_ACTIVE , IBMXCF,XCF\_SIG\_PATH\_SEPARATION, IBMXCF,XCF\_SIG\_STR\_SIZE, IXGLOGR\_STAGINGDSFULL, IZGLOGR\_ENTRYTHRESHOLD, RSM\_MEMPLIMIT, USS\_PARMLIB, VSM\_SQA\_THRESHOLD, and VSM\_CSA\_CHANGE .

**Migration action:**

1. Look at the updated checks in *IBM Health Checker for z/OS: User's Guide*.
2. Review changes you made for those checks, in HZSPRMxx parmlib members, for example.
3. Make any further updates for the checks to ensure that they continue to work as intended.



**Accommodate new SCOPE=COMMON data spaces (Required-IF)**

*Required if you use the MAXCAD parameter of parmlib member IEASYSxx and the value you specified is inadequate for your z/OS R10 system.*

The MAXCAD parameter of parmlib member IEASYSxx specifies the maximum number of SCOPE=COMMON data spaces to be allowed during an IPL. The new SCOPE=COMMON data spaces are:

- Added in z/OS V1R11: none.
- Added in z/OS V1R10: One data space used by cross-system extended services (XES). In support of virtual storage constraint relief, new storage management techniques were implemented for cross-system coupling facility (XCF) and XES to reduce the usage of 31-bit common storage below the 2 GB bar. Some XCF common storage control blocks were moved to XCF private storage, and some XES common storage control blocks were moved to a new SCOPE=COMMON data space owned by XES.

Your MAXCAD setting must be adequate to accommodate this new data space.

**Migration Action:** Increase the limit for the number of SCOPE=COMMON data spaces defined on the MAXCAD parameter if your specification is not adequate to cover the new data spaces that have been added. Note that in z/OS V1R6, the MAXCAD default was increased from 25 to 50. If this default is acceptable for your environment, you might want to remove your MAXCAD specification and allow it to default.

**Tip:** The IBM Health Checker for z/OS can help you determine what to specify for the MAXCAD value. Use the check IBMRSM,RSM\_MAXCADS. By running this check, you can find out:

- The MAXCAD value you specified during IPL
  - The number of SCOPE=COMMON data spaces currently in use
  - The high water mark, which is the highest usage of SCOPE=COMMON data spaces used during this IPL
- Use this information to help you set your MAXCAD specification in IEASYSxx.

**General Migration Actions Post-First IPL**

<none>

**General Hardware Migration Actions**



**Migrate from a Sysplex Timer to STP (Recommended)**

*Recommended because the Sysplex Timer (9037-002) has been withdrawn from marketing and the System z10 server is planned to be the last server to support it.*

The Server Time Protocol (STP) feature is the follow-on to the Sysplex Timer (9037-002). STP is designed to allow multiple servers and coupling facilities to maintain time synchronization with each other, without requiring a Sysplex Timer. STP is a hardware feature of the z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890.

STP is a server-wide facility that presents a single view of time to Processor Resource/Systems Manager™ (PR/SM™). STP uses a message-based protocol in which timekeeping information is passed over externally defined Coupling Links – InterSystem Channel-3 (ISC-3) links configured in peer mode, Integrated Cluster Bus-3 (ICB-3) links, and Integrated Cluster Bus-4 (ICB-4) links. These can be the same links that already are being used in a Parallel Sysplex for coupling facility message communication.

By using the same links to exchange timekeeping information and coupling facility messages in a Parallel Sysplex, STP can scale with distance. Servers exchanging messages over short distance links such as ICB-3 and ICB-4 are designed to meet more stringent synchronization requirements than servers exchanging messages over long

distance links such as ISC-3 (distances up to 100 km). This is an enhancement over the current Sysplex Timer implementation, which does not scale with distance.


The servers and coupling facilities that are capable of supporting STP are the z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890. The STP feature number is 1021. STP is a server-wide facility that is implemented in the Licensed Internal Code (LIC) of z10 ECs, z10 BCs, z9 ECs, z9 BCs, z990s, z890s, and coupling facilities, and presents a single view of time to PR/SM.

The Sysplex Timer's LIC has been upgraded to support using STP in a Mixed Coordinated Timing Network (CTN). The required Sysplex Timer LIC is shipped along with the STP feature and must be installed by the IBM Service Support Representative prior to migrating from a Sysplex Timer based External Time Reference (ETR) network to any STP Coordinated Timing Network (CTN).

**Migration action:** To implement STP, see the STP Web site and the publications and other resources that are listed there. The STP Web site is at <http://www.ibm.com/systems/z/ps0/stp.html>.

**Tip:** Use IBM Health Checker for z/OS check ZOSMIGREC\_SUP\_TIMER\_INUSE on z/OS V1R11 to determine whether the timing mode on the current system is ETR.

## BCP Migration Actions for z/OS R11 from z/OS R9




**Migration Actions You Can Do NOW:**


- **Remove SYSI.SVCLIB SYSI.NUCLEUS, and PASSWORD from GRS RNLs (Recommended, as of R11)**
  - Before R11, the GRS RNLs contained exclusion entries for ENQs issued for these data sets.
  - As of R11, these entries have been removed by default in PARMLIB's GRSRNL00. If you have RNLs that contain exclusion entries for these data sets, consider removing the exclusion entries. (Which will result in global ENQs on these data sets.)
  - If GRS ring with unique sysres per system, consider keeping the entries to reduce ENQs on the ring.
- **Use unique console names and console logon IDs within a GRS ring (Req IF, as of R10)**
  - Consoles (MCS, SMCS, subsystem, or EMCS) have never been allowed to be active on more than one system in a sysplex at one time.
  - As of R10, a new enqueue SYSZCNZ CONNAME#consolename is held when active. If enqueue cannot be obtained, console doesn't activate. (Similar enqueue is also obtained when a user logs onto an MCS or SMCS console – SYSZCNZ USERID#userid).
  - In GRS ring where the systems in the ring aren't part of the same sysplex, console can only be active on one system.
  - Note that multiple GRS complex sharing through 3<sup>rd</sup> party alternative serialization products will not cause problems, because all related ENQs are issued with RNL=NO to prevent 3<sup>rd</sup> party management.
  - If you will be affected by the new enqueues, uniquely name the consoles and console logon IDs within the GRS ring.

10 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## BCP Migration Actions for z/OS R11 from z/OS R9




**Migration Actions Pre-First IPL:**

- **Remove message flood automation exits (Required IF, as of APAR OA25602 on R9+)**

  - If you are using message flood automation, you had a command exit in MPFLSTxx and a message processing installation exit IEAVMXIT. Now, as of OA25602, message flood automation is integrated into z/OS and eliminates the use of the exit routines.
  - Update your MPFLSTxx to remove your USEREXIT(CNZZCMXT) statements.
  - If you only modified IEAVMXIT for message flood automation, update CONSOLxx's INIT statement from UEXIT(Y) to UEXIT(N).
  - If you used IEAVMXIT for other reasons than message flood automation, remove your message flood automation code from this exit. Reassemble and link edit it.
  - Failure to remove the code from IEAVMXIT might cause message flood automation to not be operable.
- **Enable changed GETMAIN and STORAGE OBTAIN behavior (Recommended)**
  - A new DIAGxx statement (introduced by APAR OA27291) is "VSM USEZOSVIR9RULES(YES|NO)". In z/OS VIRI0 with this APAR, the default setting is YES, which keeps the GETMAIN and STORAGE OBTAIN behavior the same as in previous releases.
  - Some applications (like DB2) can benefit from using the newly introduced behavior (that is, setting NO) for performance. However, use of NO may result in ABEND0Cx, overlay of storage, or other problems. (See APAR OA27291 for examples and more information.)
  - What to do: Install the PTF for APAR OA27291 to have pre-R10 behavior by default. To enable the new z/OS VIRI0 storage allocation behavior, use NO for that statement in DIAGxx (can be set dynamically). If it is necessary to back out the new storage allocation behavior, use YES (and change dynamically).

11 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation



## BCP Migration Actions for z/OS R11 from z/OS R9



### Migration Actions You Can Do Now:


- Use the new default for the system response to a status update missing condition (Recommended, as of R11)
  - SFM policy can specify how a system response to a status update missing condition. The default is used when no action is specified, or when no SMF policy is active.
  - Prior to R11, default was PROMPT, which prompted the operator.
  - As of R11, default changed to ISOLATETIME(0), which allows other systems to take immediate action to isolate the failed system.
  - There are mixed sysplex level considerations for this default.

### Migration Actions Pre-First IPL:

- Use the new default of COUPLExx parmlib member INTERVAL (failure detection interval) (Recommended, as of R11)
  - As of R11, the default failure detection interval (INTERVAL) calculation is derived from the excessive spin processing parm in EXSPATxx parmlib member:  $(N+1)*SpinTime + 5$ .
  - New default value ensures consistency with existing excessive spin parm that is specified via EXSPATxx parmlib member.
  - There are mixed sysplex level considerations for this default.

12 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## BCP Migration Actions for z/OS R11 from z/OS R9




### Migration Actions Pre-First IPL:

- Ensure your ESQA specification in IEASYSxx is adequate (**Req-IF, as of R11**)
  - As of R11, expect an additional 1608 (x'648') byte increase per address space in 31-bit ESQA.
  - Determine the number of address spaces that you have on the system. (The additional ESQA storage is allocated at address space create time and freed at address space termination time. Therefore, you might only need to have additional ESQA for the typical peak number of active address spaces.)
  - Multiply the number of address spaces by 1608 bytes to get the total additional ESQA storage required.
  - Ensure that your SQA= statement in IEASYSxx specifies an adequate amount of additional ESQA storage, as required.
  - Use the VSM\_SQA\_THRESHOLD Health Check from the prior release to examine the ESQA in use percent and high water mark percent.

13 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## BCP Migration Actions for z/OS R11 from z/OS R9



### Migration Actions You Can Do Now:


- Prepare for zAAP-eligible work to run on zIIPs **(Required-IF, as of R11 (enabled) and back to R9 with OA27495 (disabled))**
  - As of R11, zAAP-eligible work will run on zIIP when there are no zAAPs on the server and there are zIIPs on the server and defined to the LPAR.
  - Ensure sufficient zIIP capacity to contain this work without impacting workload goals. Use PROJECTCPU. You can disable this function on R11.

### Migration Actions Pre-First IPL:

- Accommodate quoted argument changes in System REXX MODIFY AXR command **(Required-IF, as of R11 and back to R9 with OA26802)**
  - Prior to APAR and R11: arguments bounded by quotes were passed through to the REXX exec unaltered (with quotes).
    - F AXR,EXEC `We love ` share -> `We love ` SHARE
  - Now: bounding quotes are removed. Single quotes will be substituted or passed as part of the argument when two single quotes are encountered in a quoted string. Arguments not in quotes continue to be upper-cased.
    - F AXR,EXEC `We love ` share -> We love SHARE

14 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## BCP Migration Actions for z/OS R11 from z/OS R9




### Migration Actions Pre-First IPL:

- Modify SMF exit routines IEFACTRT because of lengthened fields **(Required-IF, as of R11)**
  - The IBM sample IEFACTRT has lengthened some output fields to accommodate long-running job steps.
  - You may want to either rebase on the IBM sample, or verify in your own exit that output fields are long enough.
  - The field names are: TCB (9 bytes), SRB (9 bytes), SERV (13 bytes)
- Update automation to accommodate new responses to the DISPLAY CONSOLES command **(Required-IF, as of R10)**
  - Before R10, IEE889I, IEE890I, or IEE047I were issued.
  - As of R10, new messages CNZ4100I, CNZ4102I, or CNZ4130I are issued.
  - Also, when ACTIVE parm is specified (or defaulted), displays active subsystem consoles (as well as MCS and SMCS, as before), and NACTIVE displays inactive subsystem consoles.

15 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

**BCP Migration Actions for z/OS R11 from z/OS R9**




**Migration Actions Pre-First IPL:**

- Provide an appropriate HZSPDATA data set on fallback **(Required-IF, as of R10)**
  - As of R10, the format of the HZSPDATA data set has changed. Data written to the HZSPDATA by R10 or R11 cannot be read by z/OS R9.
  - If you fallback from R11 and have run Health Checker on R11, then you have to provide an appropriate HZSPDATA data set upon fallback to R9.
    - Note that R11 has no problem reading R9 HZSPDATA, but the reverse isn't true!
  - Make a copy of your current HZSPDATA before moving to R11, and use that copy upon fallback to R9.

16 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

**BCP Migration Actions for z/OS R11 from z/OS R9**




**Migration Actions Post-First IPL:**

- Reallocate the WLM couple data set **(Required-IF, as of R10)**
  - As of R10, the WLM policy structures have increased in size. Your existing WLM cds is probably not large enough for a policy activation on R10 or R11.
  - Use R10 or R11 libraries (STEPLIB) to allocate a new WLM cds. The values you provide will be used by WLM to calculate the appropriate space required as of R10.
- **Note:** On z/OS R10 and higher when you invoke the WLM Administration Application, the service definition in the partitioned data set will automatically update the ISPF tables to the z/OS R10 level.
  - (The ISPF table level in the partitioned data set is not the same as the WLM service definition functionality level. )
  - Once the ISPF table is updated, it is not possible to open or edit this service definition with the R10 ISPF tables, from lower level systems.

17 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## BCP Migration Actions for z/OS R11 from z/OS R9




### Migration Actions Post-First IPL:

- Do not specify VOLUME for DISP=UNCATLG (for dynamic or batch allocation) **(Required-IF, as of R10)**
  - Before R10, data set was uncataloged whether or not it was on the volume specified. So, wrong data set could have been uncataloged.
  - As of R10, when VOLUME and DISP=UNCATLG are given, data set is not uncataloged (and you receive IEF287I dsn NOT UNCTLG 13). This prevents accidental uncataloging of the wrong data set.
  - Remove VOLUME specification, if you have it for DISP=UNCATLG requests.
- As of OA27917 on z/OS V1R10, use SYSTEM VERIFY\_UNCAT option in the ALLOCxx parmlib member for migration assistance!
  - **FAIL**: default, meaning unallocation is not allowed (pre-OA27917)
  - **TRACK**: tracked with Console Tracker, and unallocation is allowed.
  - **MSGTRACK**: tracked with Console Tracker, msg IEF384I is given, and unallocation is allowed.

**New news!**

18 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## BCP Migration Actions for z/OS R11 from z/OS R9



### Migration Actions Post-First IPL:

- Update callers of CNZMXURF that do not hold appropriate serialization **(Required-IF, as of R10)**
  - CNZMXURF locates the console control block (UCME) that contains a specific console ID. Applications need to hold serialization while using CNZMXURF and accessing the UCME that it returns.
  - As of R10, if the application does not hold serialization in “distributed” mode, *it can fail*.
    - R10 introduces “distributed” mode (“shared” mode was the old pre-R10 mode, and is the default in R10). “Distributed” mode is where many nice R10 enhancements were introduced (such as 99 MCS, SMCS and subsystem consoles to be active per z/OS image)!
  - Use Tracking Facility to help, look “UCME-CMS” instances. This Tracking Facility support is in R10, and rolled back with OA19948.

19 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

**BCP Migration Actions Between z/OS V1 R9 and z/OS V1 R11**

These migration actions were taken from *z/OS Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

**BCP Migration Actions You Can Do Now****Evaluate your stand-alone dump data set allocations and your IPCS processing of them (Recommended)**

*Not required, but recommended because of changes to stand-alone dump processing (that reorder dump records with the intent of recording more important data early), and especially recommended if you deploy any LPARs with significantly more main storage than previously used.*

In z/OS V1R6, support was introduced for extended-format sequential data sets, a form of data set that is SMS-managed and can occupy more than 64 K tracks per volume. In z/OS V1R7, this support was supplemented with support for large format sequential data sets (DSNTYPE=LARGE), a form of data set that is essentially the same as conventional sequential data sets except that more than 64 K tracks may be spanned per volume. If your stand-alone dump data sets are spread over more volumes than you want, both types of support can help you gain better control over the number of volumes used for each stand-alone dump data set.

**Migration action:**

- Use multi-volume stand-alone dump data sets. Adjust the number of volumes and their separation to achieve tolerable stand-alone dump capture times.
- Use extended-format sequential data sets or large format sequential data sets. Copy their contents to an extended format, compressed, striped data set using the IPCS COPYDUMP subcommand prior to analysis. Use the same or a larger striping factor than you used for your stand-alone dump data sets. Dump data sets to which stand-alone dump can write may be neither compressed nor striped, but both attributes are advantageous for the target of the copy operation. Stand-alone dump data sets must remain in the track-managed space for z/OS V1R11.
- Use a large CISIZE and striping for IPCS dump directories, and blocking, striping, and compression for the stand-alone dump data set. Very large stand-alone dumps might require that you define your directory with the extended addressing attribute, allowing it to hold more than 4 GB.

**Prepare for zAAP-eligible work to run on zIIPs (Req-IF, as of R11 and back to z/OS R9 with OA27495)**

*Required if using zIIPs with no zAAPs.*

In z/OS V1R11, System z Application Assist Processor (zAAP)-eligible work running in z/OS LPARs can run on System z Integrated Information Processors (zIIP) if no zAAPs are installed on the server, and zIIPs are installed on the server and defined for the LPAR.

**Migration action:**

- For z/OS LPARs on servers that have zIIPs installed with no zAAPs installed, determine:
  - If you have zAAP-eligible work that will now be eligible to run on zIIPs.
  - If there is sufficient zIIP capacity to contain this work without impacting workload goals.
  - If the general-purpose (CP) capacity should be adjusted for any affected LPARs as a result of any additional projected offload.
- For z/VM LPARs where z/OS is running as a guest, determine:
  - If you have zAAP-eligible work that will now be eligible to run on real or virtual zIIPs defined to the guest.
  - If there is sufficient real zIIP capacity to contain this work without impacting workload goals.
  - If the general-purpose (CP) capacity should be adjusted for any affected LPARs as a result of any additional projected offload.

You can use PROJECTCPU in IEAOPTxx on your current release to determine how much zAAP-eligible work is running in each LPAR, and RMF (or equivalent) reports, to determine how busy your zIIPs are. If you want to disable zAAP-eligible work from running on your zIIPs, you can set the new ZAAPzIIP parameter in IEASYSxx to NO.

**Note:** This support is available on z/OS V1R9 and z/OS V1R10 with APAR OA27495. The behavior on z/OS V1R9 and z/OS V1R10, however, is not enabled by default. It is enabled by default on z/OS V1R11. If you wish to enable this support on z/OS V1R9 or z/OS V1R10, use the ZZ parameter in IEASYSxx. The ZZ parameter is also accepted on z/OS V1R11 where ZZ is a synonym for the ZAAPzIIP parameter.

**Remove SYS1.SVCLIB, SYS1.NUCLEUS, and PASSWORD from global resource serialization RNLs (Recommended, as of R11)**

*Not required, but recommended because by using the z/OS V1R11 default your system will have global ENQs on data sets SYS1.SVCLIB, SYS1.NUCLEUS, or PASSWORD.*

Before z/OS V1R11, the global resource serialization default resource name lists (RNLs) contained exclusion entries for ENQs issued for data sets SYS1.SVCLIB, SYS1.NUCLEUS, and PASSWORD. With z/OS V1R11, these exclusion entries have been removed from PARMLIB's GRSRNL00. If you have RNLs that contain exclusion entries for any of the three data sets, you should consider removing the exclusion entries.

The requirement for SYS1.SVCLIB and SYS1.NUCLEUS to be in the exclusion RNL became obsolete long before z/OS V1R9. The change can be made on any supported z/OS release at any time.

The PASSWORD data set contains the password for each data set that is password protected. It is strongly recommended to use SAF to invoke a security product such as RACF to protect data sets, rather than using password protection. If you are not password protecting data sets, PASSWORD can also be removed from the exclusion RNL on any supported z/OS release at any time.

### **Migration action:**

- If you have a modified version of the RNLs that still contains SYSDSN PASSWORD, SYSDSN SYS1.NUCLEUS, or SYSDSN SYS1.SVCLIB, consider removing these entries from the exclusion list to ensure that the data sets are properly serialized. But note that if your system is running in ring mode and it has a unique SYSRES volume per system, consider keeping the entries to reduce ENQs going around the ring.
- If you are using password protection (PASSWORD), use system authorization facility (SAF) instead.

### **Update procedures for vary systems out of the sysplex (Required-IF, as of R11)**

*Required if your local procedures depend on explicit reaction to message IXC102A and the installation is not running z/OS as a VM guest.*

You can use the VARY XCF,sysname,OFFLINE command to vary a system out of the sysplex. Before a system can be removed from the sysplex, it must be reset or isolated from the channel subsystem to ensure that it is no longer performing I/O against shared sysplex resources such as data sets or coupling facility structures.

Before z/OS V1R11, the VARY XCF process would only initiate automatic isolation if a Sysplex Failure Management (SFM) policy were in effect. If no SFM policy were active, a system reset would be required. The system would prompt the operator with message IXC102A to determine when the reset had been performed.

Beginning with z/OS V1R11, VARY XCF always initiates automatic isolation when physically feasible and when z/OS is not running as a VM guest, regardless of whether there is an active (SFM) policy. The system issues message IXC102A only if automatic isolation is not feasible or if no surviving system is able to complete isolation successfully. This change does not depend on the availability of the BCPii or system status detection functions.

**Migration action:** Update local procedures for removing systems from the sysplex to reflect the fact that message IXC102A may no longer be issued.

### **Use the new default for the system response to a status update missing condition (Recommended, as of R11)**

*Not required, but recommended because specifying or defaulting to ISOLATETIME(0) allows SFM to fence and partition a failed system without operator intervention and without undue delay.*

You can use the sysplex failure management (SFM) policy to specify how a system is to respond to a status update missing condition. The system default is used when no action is specified for the system or when no SFM policy is active.

Before z/OS V1R11, the system default action was PROMPT, which prompted the operator when a system entered a status update missing condition. As of z/OS V1R11, the system default action has changed to ISOLATETIME(0), which allows other systems to take immediate action to isolate the failed system. If a system enters a status update missing condition and there is no active SFM policy, the monitoring system will take the system default action:

- If both systems are at or above z/OS V1R11, and the failing system was able to identify itself to the sysplex for automatic isolation (either to a CF for fencing or through XCF System Status Detection), ISOLATETIME(0) is used as the default action against the failing system.
- If either system is prior to z/OS V1R11, and the failing system was able to identify itself to the sysplex for automatic isolation (either to a CF for fencing or through XCF System Status Detection), PROMPT is used as the default action against the failing system.

The DISPLAY XCF,C command shows the SFM action the system expects, though the monitoring system might use a different default action if no SFM policy is defined.

### **Migration action:**

- To obtain the new behavior, which is recommended, specify or default to ISOLATETIME(0) to allow SFM to fence and partition a failed system without operator intervention and without undue delay.

- To obtain the pre-z/OS V1R11 behavior, set up an SFM policy that specifies PROMPT for systems that should not be isolated automatically upon status update missing.

**Tip:** Use IBM Health Checker for z/OS check XCF\_SFM\_ACTIVE. This check reports whether SFM is active (and if it is active, whether the recommended settings are in use for the SFM policy). If the check indicates that SFM is not active, this migration action is applicable to the system on which the check ran, meaning that when this system enters a status update missing condition, the changed default from PROMPT to ISOLATETIME(0) might be in effect.

### **Enable ENF 51 listeners to listen to conditional ENQ contention (Req-IF, as of R11)**

*Required if listeners default to no qualifier at registration.*

Before z/OS V1R11, ENF 51 only allowed listeners to be notified of unconditional contention. With z/OS V1R11, previously unused byte 3 of the ENF 51 qualifier is used. Now, ENF 51 listeners that default to no qualifier at registration can be driven for conditional contention events. Listeners that specify the QUAL parameter with the QMASK parameter including the BYTE3 value on an ENFREQ ACTION=LISTEN macro should not be affected by the change. However, if listeners do not specify QUAL, or specify a QMASK without BYTE3, they will receive conditional contention events and might misinterpret them as traditional unconditional contention events.

**Migration action:** Do one of the following:

- If the QUAL parameter is specified on an ENFREQ ACTION=LISTEN macro with a QMASK parameter that includes the BYTE3 value, there is no incompatibility and no action is required.
- If QUAL is not specified or QMASK is specified without BYTE3, include BYTE3. (The QUAL specification should not need to be changed, assuming its value for the third byte is 0.)
- Listeners who are interested in getting control for the new “waitless” conditional contention events could consider excluding BYTE3 on the QMASK or using one of the newer QUAL constants defined in data area ISGE51CN.

### **Discontinue use of IPCS problem management subcommands (Required-IF, as of R11)**

*Required if you currently use the IPCS problem management subcommands to report and track problems IBM has removed.*

Interactive problem control system (IPCS) problem management subcommands were functionally stabilized in 1981 but left in IPCS. Since then, customers have been advised to migrate to other problem management tools. z/OS V1R10 was the last release that included the subcommands. As of z/OS V1R11, the subcommands have been removed from z/OS. The subcommands are:

- ADDDSN — add a data set name to a problem
- ADDPROB — add a problem to IPCS
- DELDSN — delete a data set from a problem
- DELPROB — delete a problem from IPCS
- LISTDSN — list data set attributes
- LISTPROB — list problems
- MODDSN — modify data set attributes
- MODPROB — modify problem attributes

Note that IBM plans to continue to enhance the dump and trace analysis and display facilities of IPCS.

**Migration action:** Do not use the IPCS problem management subcommands anymore. Use other problem management tools to report and track problems. For example, consider migrating to IBM Tivoli Information Management for z/OS V7 (5698-A08).

### **Use unique console names and console logon IDs within a GRS ring complex (Req-IF, as of R10)**

*Required if console names and console logon IDs are not unique within a global resource serialization ring complex where the systems in the ring are not all part of the same sysplex.*

OPERLOG activation fails on a MONOPLEX system during, or after joining, an existing GRS Ring configuration because of the ENQ contention of MAJOR NAME: SYSZCNZ and MINOR NAME: CONNAME#\*OPLOG01. With the PTF for APAR OA30757 installed, OPERLOG processing is modified to enable its activation on a MONOPLEX system during, or after joining, an existing GRS Ring configuration. The OPERLOG EMCS console name will be generated using the two (2) character System Clone value obtained from the System Name (for example, System Name = SYSTEM1 / System Clone = M1).

Old: \*OPLOG01

New: \*OPLOGxx (where 'xx' is the System Clone value)

Consoles, whether MCS, SMCS, subsystem, or EMCS, have never been allowed to be active on more than one system in a sysplex at one time. Beginning with z/OS V1R10, a new enqueue, SYSZCNZ ONNAME#*consolename*,

is held while a console is active. If the enqueue cannot be obtained, the console does not activate. If serialization is managed within a global resource serialization ring complex, the console can only be active on one system, even if the systems are not part of the same sysplex. Attempting to activate the console could result in an abend. A similar enqueue, SYSZCNZ USERID#*userid*, is obtained when a user logs on to an MCS or SMCS console. The same naming restrictions would apply to the logon ID. Note that multiple global resource serialization complex sharing through third-party alternate serialization products will not cause this problem because all the related ENQs are issued with RNL=NO to prevent third-party management.

**Migration action:** Issue D GRS,LINK to determine whether global resource serialization is managing any CTCs. If it is, do the following:

1. Issue D C,L,S and D EMCS,S from the sysplex and each standalone system in the ring to compare the defined consoles.
2. Uniquely name the consoles on the systems within the global resource serialization ring complex.

**Note:** Failure to use unique console names and console logon IDs to queue messages to the SYSLOG within the GRS ring could result in failure message CNZ3001A ACTIVATE UNSUCCESSFUL FOR SYSLOG CONSOLE *console-name* MCSOPER RETURN CODE: *retcode*, MCSOPERREASON CODE: *rsncode* SYSLOG IS NOT SUPPORTED. If your console names are generated from the SYSCONE value, ensure that those SYSCONE values are unique.

- The message IEA303W ABEND 077 REASON 060107F0 DURING INITIALIZATION UNDER RIM IEAVNPA1 is issued for this problem during the system initialization.
- This problem happens only when all the systems within the GRS Ring complex are running in z/OS V1R10 or later.
- The default value for the &SYSCONE is resolved from the last two values in the SYSNAME parameter. So, for example, "&SYSCONE=YS" is the default value for both systems "SYSNAME=ASYS" and "SYSNAME=BSYS".

### **Stop using CPU affinity (Required-IF, as of R10)**

*Required if you use CPU affinity.*

IBM has removed support for CPU affinity. CPU affinity originally provided support for assigning work to the same CP. However, this predates Processor Resource/Systems Manager (PR/SM), which introduced logical processors. CPU affinity has long been obsolete for z/OS partitions operating under PR/SM, which makes the CP on which a piece of work executes unpredictable. z/OS V1R9 is the last release to provide support for the assignment of CPU affinity to a logical processor. Beginning with z/OS V1R10, any attempt to assign CPU affinity is ignored. For example, specifications for assigning a program to a specific logical processor or processors using the program properties table (PPT) or a SCHEDxx member of parmlib are ignored.

**Migration action:** Do the following to stop using CPU affinity. If you do not make the changes indicated, the affinity specification is ignored and the result is what it would have been if affinity were not specified.

1. In parmlib member SCHEDxx, look for AFF(xxx) in the PPT section and either remove it or update it to AFF(NONE). Also, update the specified executable program so that it does not have any dependencies on affinity.
2. In programs, look for use of the SCHEDULE macro where the SRBCPAFF field is set to a value other than 0 or X'FFFF', and use of the IEAMSCHD macro with the CPUMASK keyword with an operand other than 0 or X'FFFF'. Change your programs so that they do not require a specific processor on which to run.
3. In programs, look for use of the IWMCPAFN service. Change your programs so that they do not require a specific processor on which to run, which has the effect of removing the programs' use of IWMCPAFN.

### **Migrate from the prelinker to the program management binder (Recommended)**

*Not required, but recommended because the prelinker is not planned to be enhanced. Enhancements will only be made to the program management binder..*

**Migration action:** Follow the instructions for migrating from the prelinker to binder in *z/OS MVS Program Management: User's Guide and Reference*.

## **BCP Migration Actions Pre-First IPL**

### **Create IPL text (Required)**

**Migration action:** Update and run the IPLTEXT job to write a new copy of the IPL text. If you install z/OS with a ServerPac, an installation dialog job is provided to perform this action. If you install z/OS with a CBPDO, instructions to perform this action are provided in *z/OS Program Directory*.



**Note:** When the IPLTXTEXT parameter (which was introduced by ICKDSF R17 APAR PK16403) is specified with the REFORMAT command using the IPLDD parameter, WTOR message ICK21836D is suppressed if IPL text already exists

### **Reassemble the stand-alone dump program (Required)**

**Migration action:** Reassemble the stand-alone dump program. If you install z/OS with a ServerPac, an installation dialog job is provided to perform this action. If you install z/OS with a CBPDO, instructions to perform this action are provided in *z/OS Program Directory*. Once the stand-alone dump program is properly created on a DASD residence volume, it resides in the SYS1.PAGEDUMP.Vvolser data set.

### **Ensure correct use of CIRB (Required, as of R11)**

When the CIRB function (Create Interruption Request Block) is used with SVC-entry and the IRB is to be run on a task different than the task that issued CIRB (identified within field IQETCB), and the STAB=DYN option was requested, the system is told to free storage. Before z/OS V1R11, the freemain was conditional and would always fail (with a bad return code) in the case mentioned above because the storage was not owned by the freeing task. Since you would not see this, as it was only a return code, you were not aware of the failure. Starting with z/OS V1R11, the freemain is changed to unconditional so that you will be aware of your error and can fix or change your code.

**Migration action:** The best practice is to change to use SCHEDIRB instead of CIRB. SCHEDIRB has been recommended for many years. Otherwise, take one of the following actions according to your need:

- If you do not need the system to free the storage, do not specify STAB=DYN.
- If you do need the system to free the storage, make sure that you are in PSW key 0 and place the TCB address under which the IRB will run in register 4 when you invoke CIRB and specify the BRANCH=YES option of CIRB.



### **Ensure your ESQA specification in IEASYSxx is adequate (Required-IF, as of R11)**

*Required if your ESQA specification in IEASYSxx is inadequate.*

In z/OS V1R11, each address space requires 1608 bytes of additional 31-bit Extended System Queue Area (ESQA) storage.

If ESQA storage is unavailable at any time to the system, the system will try to obtain the storage needed from Extended Common Storage Area (ECSA). When ESQA expands into ECSA, IRA104E will be issued. If ESQA and ECSA are unavailable, the system allocates space from SQA followed by CSA below 16 megabytes. As SQA storage becomes scarce, the system rejects LOGON, MOUNT, and START commands, and stops new jobs from being created and address spaces from being swapped in. ESQA/SQA storage shortages can eventually result in a system failure.

#### **Migration action:**

1. Determine the number of address spaces that you have on the system. One way to do this is to issue the DISPLAY A,L command and total the address spaces in the IEE114I and IEE115I messages. Note that the additional ESQA storage is allocated at address space create time and freed at address space termination time. Therefore, you might only need to have additional ESQA for the typical peak number of active address spaces.
2. Multiply the number of address spaces by 1608 (x'648') bytes to get the total additional ESQA storage required.
3. Ensure that your SQA= statement in IEASYSxx specifies an adequate amount of ESQA. One way to do this is to use the VSM\_SQA\_THRESHOLD check, provided by the IBM Health Checker for z/OS, from the prior release to examine the ESQA in use percent and high water mark (HWM) percent. Then add the additional ESQA storage needed from the prior step to calculate the projected ESQA in use and HWM percent. Adjust the amount of ESQA needed by your installation.

### **Remove all unnamed console definitions (Required, as of R10)**

Prior to z/OS V1R10, residual sysplex data for unnamed consoles was ignored during IPL. However, beginning with z/OS V1R10, unnamed consoles will cause an ABEND077 RSN00B03 during IPL unless they are removed. These unnamed consoles would have been created in a sysplex prior to the console naming requirements introduced in z/OS V1R4.2 and can be identified with D C,L and D C,SS commands. If present, the unnamed consoles will appear in the console name (*consname*) field of message IEE889I as two numeric characters (for example, 01,02,...99).

**Migration action:** There are two methods available for removing the unnamed console definitions:

- Follow the instructions for "Removing Console Definitions from a Configuration" located in *z/OS MVS Planning: Operations* at <http://publibz.boulder.ibm.com/epubs/pdf/iea2g391.pdf>.
- Perform a sysplex-wide IPL.

### **Use the new default of COUPLExx parmlib member INTERVAL (failure detection interval) (Recommended, as of R11)**

*Not required, but recommended because the new default value can ensure consistency with the existing excessive spin parameters that are specified through EXSPATxx*

The default handling of the COUPLExx parmlib member INTERVAL (failure detection interval) value is changed to ensure consistency with the existing excessive spin parameters that are specified via EXSPATxx.

Before z/OS V1R11, XCF used twice the default spin loop timeout value, plus 5, for the default failure detection interval. Note that XCF used the default spin loop timeout value and not what you specified in the EXSPATxx parmlib member.

With z/OS V1R11, XCF derives the default INTERVAL value from the excessive spin processing parameters in EXSPATxx. The value is computed as follows:  $(N+1) * \text{SpinTime} + 5$  where N is the number of excessive spin recovery actions, +1 indicates the implicit spin action, and SpinTime is the excessive spin loop timeout interval. The result might be higher effective failure detection intervals being used than the value that was explicitly coded in the COUPLExx parmlib member INTERVAL keyword.

If the excessive spin parameters are not explicitly set via an EXSPATxx parmlib member, the IBM-supplied default spin parameters are used. The value is computed as follows:  $(1+4) * \text{SpinTime} + 5$

**Migration action:** If the COUPLExx parmlib member codes INTERVAL, perform the following actions as needed:

- If the specified INTERVAL is greater than the default INTERVAL that will be derived from the EXSPATxx specifications, no action is needed.
- If the specified INTERVAL is less than the default INTERVAL and you want to use the specified INTERVAL, enable the USERINTERVAL switch.

If the COUPLExx parmlib member does not specify INTERVAL and the default INTERVAL value is taken, perform the following actions as needed:

- If the new default INTERVAL is acceptable, no action is needed.
- If you want to use the old default, you must code INTERVAL to set the desired value and proceed as above.
- If COUPLE does not specify INTERVAL but specifies OPNOTIFY, you should consider specifying a relative OPNOTIFY value so that OPNOTIFY can be adjusted to maintain the same relative delta between the old default INTERVAL and the new default INTERVAL.

**Note:** If a COUPLExx parmlib member is used by more than one system, the migration actions for the new system might conflict or cause different behaviors on pre-z/OS V1R11 systems.

### **Accommodate quoted argument changes in System REXX MODIFY AXR command (Required-IF as of R11, and back to R9 with OA26802)**

*Required if quotation marks are used in the MODIFY AXR,rexexecname args command and the REXX exec expects and processes the quotation marks.*

Before z/OS V1R11 (and APAR OA26802 on prior releases), arguments bounded by quotation marks were passed through to the REXX exec unaltered (including the quotation marks). With changes introduced in z/OS V1R11 (and by APAR OA26802 on prior releases), the final argument string that is passed through to the REXX exec will have bounding quotation marks removed and single quotation marks will be substituted and passed as part of the argument, when two single consecutive quotation marks are encountered within a quoted string. Arguments not within quotation marks continue to be upper-cased by system command processing.

**Migration action:** Determine if you have the new behavior (step 1), then proceed to steps 2 and 3 if quotation marks are currently being used within the argument string of the MODIFY AXR command.

1. The presence or absence of this enhancement can be detected within a program as follows:
  - a. REXX – Issue the MODIFY AXR,SYSREXX REXXLIB command from the exec by including the statement `AxrCmdRc = AXRCMD('F AXR,SR R',Msg.,4)`. When this support is not present ASA1001 SYNTAX ERROR: STATUS is returned in the Msg.1 stem variable.
  - b. ASSEMBLER – include an AXREXX macro call specifying `REQUEST=GETREXXLIB`. Return code 8 reason code `AxrBadRequest` is returned when the AXREXX GETREXXLIB support is not present. `AxrBadRequest` is defined as `x'00000846'` in `SYS1.MACLIB(AXRZARG)`.
2. Preserve the lower case of the argument. To preserve the lower case of arguments, it is necessary to bound them by quotation marks. As bounding quotation marks are removed with z/OS V1R11, this new behavior might affect argument processing within the REXX exec. Within the REXX exec, it is unnecessary to check

for bounding quotation marks when parsing the argument string.

Input argument string	Resolved argument string
'www.lakeminnewaska.org'	www.lakeminnewaska.org
'Fire declared ' out	Fire declared OUT

- Pass quotation marks as part of the argument. If quotation marks are currently used, it will be necessary to use the new syntax when entering the argument string portion of the MODIFY AXR command when running a REXX exec from the console. Bounding quotation marks are removed with z/OS V1R11. To preserve a single quotation mark in an argument string, two consecutive quotation marks must be imbedded within the argument string. For example, 'The cat"s tail. would be resolved to The cat's tail.
- Additional examples follow:

Input argument string	Resolved argument string
'Minnewaska''s Scenic Beauty'	Minnewaska's Scenic Beauty
''' Shawangunk' Mountain 'Trails'	' Shawangunk MOUNTAIN Trails
'''Lake Awosting'''	'Lake Awosting'
''Ranger Guide''	RANGER GUIDE
''''Mossy Glenn Footpath''''	''Mossy Glenn Footpath''

### **Remove message flood automation exits (Required-IF as of OA25602 on R9 and R10)**

*Required if message flood automation is used.*

Before APAR OA25602 on z/OS V1R10 and prior releases, if you wanted to implement message flood automation, you had to do so with a command exit specified in parmlib member MPFLSTxx and with message processing installation exit IEAVMXIT. Beginning with APAR OA25602, message flood automation is integrated in z/OS, eliminating the use for the exit routines. If you want to continue using message flood automation, you must remove what you have implemented with the exit routines.

#### **Migration action:**

- Update your MPFLSTxx parmlib member to remove all .CMD USEREXIT(CNZZCMXT) statements.
- If you use exit IEAVMXIT only for message flood automation, update your CONSOLxx parmlib member to change the INIT statement option of UEXIT(Y) to UEXIT(N). Do not do this if you want to continue to use IEAVMXIT for other purposes.
- If you modified your IEAVMXIT exit routine to incorporate message flood automation processing, you must remove the message flood automation code, reassemble, and use the linkage editor or program management binder to install.

If the message flood automation exits are not removed, message flood automation might not be operable. If it is operable, it will be the old message flood automation, not the new, integrated message flood automation.

**Tip:** Use IBM Health Checker for z/OS check CNZ\_OBSOLETE\_MSGFLD\_AUTOMATION introduced in APAR OA25602, available in z/OS V1R11, to validate that the obsolete version of message flood automation is not in use.

### **Accept the new lower storage limit for stopping SVC dump data capture (Recommended, as of R11)**

*Not required, but recommended because it limits DUMPSRV's use of virtual (and therefore also auxiliary) storage while reducing the tuning efforts required to support a given MAXSPACE value.*

Before z/OS V1R11, SVC dump data capture would not be prevented or stopped until auxiliary storage usage reached 85%. At this level of utilization the system might be unresponsive. For sure, no new jobs could be started. As of z/OS V1R11, the new AUXMGMT=ON default lowers the thresholds at which SVC dump capture activity is affected. New SVC dumps are not allowed when auxiliary storage usage reaches 50%, and an SVC dump that is in the process of capturing data is truncated, resulting in a partial dump, when auxiliary storage usage reaches 68%. Once either limit is reached, data capture for new dump requests is not allowed again until auxiliary storage usage drops below 35%. This new behavior makes system availability a higher priority than capturing failure data (a complete dump) when resources are restricted. It also makes tuning the MAXSPACE parameter less daunting. A very large MAXSPACE value can be specified because the amount of available auxiliary storage is a key factor in

automatically limiting what SVC dump data capture uses. The lower 68% point leaves the system in a more responsive state.

**Migration action:** Do one of the following:

- To use the new limits, there is nothing to do. The system always IPLs with AUXMGMT=ON by default. Note that you can display the current status of AUXMGMT by using the DISPLAY DUMP,OPTIONS (D D,O) command. **Note:** Once SVC dump processing has detected a shortage, the auxiliary storage utilization must drop below 35% before new SVC dump requests will be honored. The condition cannot be removed by simply changing the setting of AUXMGMT from ON to OFF.
- If you have spent considerable effort tuning the real, auxiliary, and MAXSPACE resources for a particular behavior, you might want to keep the pre-z/OS V1R11 behavior. To suppress the new behavior and use the prior limits, ensure that the CHNGDUMP SET,SDUMP,AUXMGMT=OFF command is issued out of a SYS1.PARMLIB(COMMNDxx) member that receives control during IPL. Remember that with AUXMGMT=OFF, the MAXSPACE value must be accompanied by the appropriate increase in auxiliary storage (paging) resources (at least three times the MAXSPACE value), above the system's worst case usage of auxiliary storage

### **Do not specify volume information on allocations intended to uncatalog a data set (Required-IF as of R10)**

*Required if you have allocation requests that uncatalog data sets.*

Before z/OS V1R10, a DD for a data set that specified DISP=UNCATLG for either the normal or abnormal termination disposition was allowed to be uncataloged by the job step, even if the system did not retrieve the volume information from the catalog. This would allow a job to uncatalog a different data set with the same name. Starting with z/OS V1R10, the default behavior for a DD for a data set that specifies DISP=UNCATLG will be to not allow the data set to be uncataloged if the job did not retrieve the data set information from the catalog. Instead, message IEF287I dsn NOT UNCTLGD 13 is issued. This change prevents a user from accidentally uncataloging the wrong data set.

A migration option is provided as of OA27917 on z/OS V1R10 in the ALLOCxx Parmlib member that will allow the pre-z/OS V1R10 behavior, with tracking to the Console Tracker and optionally a message as well when a data set is uncataloged without retrieving the volume information from the catalog. See the SYSTEM VERIFY\_UNCAT option in the ALLOCxx Parmlib member, described in the MVS Initialization and Tuning Reference book, for details on using the option.

Common cases where the system will not retrieve volume information from the catalog are when VOLUME is specified on the DD, or for a NEW dataset. Note that this change applies to disposition processing in dynamic allocation as well as batch allocation.

**Migration action:** Allocation requests that specify DISP=UNCATLG must be reviewed and corrected if necessary. Any request that specifies VOLUME= or DISP=NEW will likely be affected by this change. Depending on what the request is attempting to accomplish, this may require removing the VOLUME= information, changing DISP=UNCATLG to a different disposition such as DISP=DELETE, or using a utility program (such as IDCAMS) to uncatalog the data set.

With OA27917, If your installation specifies SYSTEM VERIFY\_UNCAT(FAIL) behavior in the ALLOCxx Parmlib member, or takes the default, affected jobs can be identified because message IEF287I dsn NOT UNCTLGD 13 will be issued on behalf of the DD that must be changed. The unallocation request will not be allowed.

With OA27917, If your installation specifies SYSTEM VERIFY\_UNCAT(TRACK), then affected jobs will be identified in the Console Tracking Facility, but the unallocation request will be allowed.

With OA27917, If your installation specifies SYSTEM VERIFY\_UNCAT(MSGTRACK), then affected jobs will be identified in the Console Tracking Facility and message IEF384I will be issued on behalf of the DD that must be changed, but the unallocation request will be allowed.

### **Modify your GDPS start procedures (Required, as of R10)**

*Required if you use GDPS/PPRC or GDPS/HM.*

Beginning with z/OS V1R10, if you use a currently supported release of the GDPS/PPRC or GDPS/PPRC HyperSwap Manager (GDPS/HM) products, you must update your GEOXHSWP started task JCL. Before z/OS V1R10, the GEOXHSWP started task was started using the GDPS-supplied GEOXHSWP module, which is no longer supported on z/OS V1R10 and later systems.

**Migration action:** Replace PGM=GEOXHSWP with PGM=IOSHSAPI in the GEOXHSWP started task JCL on z/OS V1R10 or later systems.

### **Update automation and exits for the modified MCS logon command forum (Required-IF as of R10)**

*Required if the MCS LOGON command is changed by automation or exit routines.*

CNZMYLGN is a new macro in z/OS V1R10. Automation and command exits that modify MCS LOGON commands must use the format mapped by the CNZMYLGN macro. Before z/OS V1R10, there was no mapping of the LOGON command format, so programs scanning the command and making modifications could reorder the command's operands. With z/OS V1R10, the operands must appear in a defined order and location. CNZMYLGN maps these locations.

**Migration action:** Change automation or command exits that modify the MCS LOGON command to ensure that the format mapped by CNZMYLGN is used.

### **Discontinue the use of console names LOGON and LOGOFF (Required-IF as of R10)**

*Required if a console name of LOGON or LOGOF is used.*

Before z/OS V1R10, LOGON and LOGOFF were accepted as valid console names for an MCS, SMCS, EMCS, or subsystem console. As of z/OS V1R10, LOGON and LOGOFF are not acceptable as console names. Queries of a console name of LOGON or LOGOFF by using the CONVCON or CnzConv service will indicate that the console name is reserved.

**Migration action:**

- Change your CONSOLxx parmlib members to remove use of console names LOGON and LOGOFF.
- Change any SEND commands that have a target console name of LOGON or LOGOFF.
- Change your code that invokes MCSOPER with consoles names of LOGON or LOGOFF.
- Change your code that invokes CnzConv or CONVCON with console names of LOGON or LOGOFF.

### **Accommodate VARY command changes (Required-IF as of R10)**

*Required if COMMNDxx or IEACMDxx parmlib members contain the VARY attribute, VARY CN(),OFFLINE with attribute, or VARY CN(),LU='NONE' commands, or if automation programs or exit routines issue MGCRE macros that execute a VARY attribute or VARY CN(),OFFLINE with attribute command.*

- VARY command attribute processing is changed in z/OS V1R10 to no longer allow attributes to be altered for inactive MCS, SMCS, and subsystem consoles. There are two exceptions: the LU and LOGON attributes can be changed for inactive SMCS consoles. Before z/OS V1R10, VARY processing allowed all attributes to be changed for inactive MCS, SMCS, and subsystem consoles. Only the attributes for inactive EMCS consoles could not be changed. Beginning with z/OS V1R10, VARY command attribute processing no longer allows attributes to be altered for inactive MCS, SMCS, and subsystem consoles (as it does for inactive EMCS consoles). The only exceptions are the LU and LOGON attributes for inactive SMCS consoles. If a VARY command requests a change to an SMCS console attribute that is not LU or LOGON, message IEE871I CN=consname NOT ACTIVE. CAN NOT CHANGE CONSOLE ATTRIBUTES is issued.
- The VARY CN(),OFFLINE command is changed in z/OS V1R10 to no longer allow attributes to be altered when specified with the OFFLINE keyword. OFFLINE processing will still deactivate the console (except if the LU or LOGON keywords are specified) but the attributes will not be changed. Before z/OS V1R10, the VARY CN(),OFFLINE command allowed all attributes to be specified and their attributes changed along with the console being deactivated. Only the attributes for inactive EMCS consoles could not be changed. In z/OS V1R10, a VARY CN(),OFFLINE command with attributes specified is rejected, as follows:
  - If the command specifies either the LU or LOGON attribute, the console is not deactivated and message ASA103I is issued indicating that the LU or LOGON keyword cannot be specified with the OFFLINE keyword.
  - If the command specifies other attributes, the console is deactivated and the following message is issued: IEE871I CN=consname NOT ACTIVE. CAN NOT CHANGE CONSOLE ATTRIBUTES.

The reason there are two different messages (ASA103I and IEE871I) is that the LU and LOGON attributes are the only attributes that can be changed for inactive SMCS consoles, and it is necessary to differentiate these attributes from the rest of the console attributes (ROUT, INTIDS, MSCOPE, and so forth). Message IEE871I does not apply to LU and LOGON because those attributes can be changed when the console is not active.

- VARY command attribute processing is changed in z/OS V1R10 to allow an SMCS console to change its LU attribute to \*NONE\* independent of the SMCS console's LOGON attribute.

Before z/OS V1R10, VARY processing required that the LOGON keyword be specified when the LU=\*NONE\* attribute was specified. The exception was when the SMCS console already had the LOGON=REQUIRED attribute, changing the LU to \*NONE\* would not require the LOGON attribute to be specified. In z/OS V1R10, it does not matter what the LOGON attribute is for an SMCS console. The LU attribute can be changed to \*NONE\* anytime.

Note that having the LOGON attribute specified with the LU=\*NONE\* attribute still works; the LOGON attribute is just no longer required.

**Migration action:**

- Check if the consoles' attributes really have to change. If they do, ensure that the consoles are active before changing their attributes. Do not issue attribute changes on a VARY CN(),OFFLINE command. If that command is required, first issue the VARY attribute command while the console is active, and then deactivate the console with a VARY CN(),OFFLINE command.
- Check the CONSOLxx parmlib member and verify that the console attributes are set correctly so that the consoles do not need to be changed by using the VARY command.
- Check if automation programs and exit routines issue MGCRE macros that execute a VARY attribute or VARY CN(),OFFLINE with attribute command.
- Notify system operators about the VARY command changes.
- If LU=\*NONE\* is specified, remove the LOGON attribute because it no longer needs to be specified. The LOGON attribute is still processed but it is no longer necessary.



**Modify applications that use unsupported subsystem console functions (Required-IF as of**

**R10)**

*Required if subsystem consoles are being used.*

Before z/OS V1R10, the following subsystem console requests were supported:

- A demand select request for an MCS console as well as a subsystem console
- A request to change the routing codes for a subsystem console
- A request to release, set master authority, or reset master authority for a subsystem console that is active on another system.

As of z/OS V1R10, the functions are no longer supported, as follows:

- A demand select request can no longer be made for an MCS console. Also, the SCSUNIT and SCSUNIT4 fields in the IEZVG100 parameter list must not be specified, meaning they must be zero.
- A request to change the routing codes for a subsystem console is ignored and returned with return code zero (assuming no other errors are detected).
- A request to release, set master authority, or reset master authority can be made only for a subsystem console that is active on the system that makes the request.

**Migration action:**

- Request for demand select:
  - – *What to look for:* Invocations of IEAVG700 in programs that run on z/OS V1R6 and later where the subsystem console service routine parameter list (SCSR as mapped by IEZVG100) flag SCSDEMSEL is being set and the console being requested is not a subsystem console, or where fields SCSUNIT or SCSUNIT4 are being used.
  - – *What to do:* Change the request to be for a subsystem console and do not set fields SCSUNIT or SCSUNIT4. Note that while it is necessary to make this change for programs that run on z/OS V1R8 and later with the PTF for APAR OA18204 installed, it is acceptable and advisable to change all such programs regardless of release level and regardless of whether the PTF is installed.
- Request to change routing codes:
  - – *What to look for:* Invocations of IEAVG700 in programs that will run on z/OS V1R10 and later where the subsystem console service routine parameter list (SCSR as mapped by IEZVG100) flag SCSRTCDF or field SCSRTCDP is being set.
  - – *What to do:* It is not necessary to make any changes because requests to change routing codes are accepted but not processed beginning with z/OS V1R10. It is, however, advisable to examine those programs to ensure that they will continue to function correctly when their requests to change routing codes are accepted but not processed. It is advisable that such an examination be done for all such programs regardless of release level.
- Request to release, set master authority, or reset master authority on nonowning system:
  - – *What to look for:* Invocations of IEAVG700 in programs that run on z/OS V1R8 and later where the subsystem console service routine parameter list (SCSR as mapped by IEZVG100) flags SCSRELSE, SCSPMSTR, or SCSNMSTR are being set and the request is for a subsystem console that is owned on some other system.
  - – *What to do:* Change the invocation so that it is executed on the system where the console is owned. Note that while it is necessary to make this change for programs that run on z/OS V1R8 and later with the PTF for APAR OA18204 installed, it is acceptable and advisable to change all such programs regardless of release level and regardless of whether the PTF is installed.

You can use the Tracking Facility to help with this migration action. In Tracking Facility output, look for violations that start with "IEAVG700".

### **Discontinue use of the DISPLAY CONSOLES,HCONLY command (Required-IF as of R10)**

*Required if you have code that either issues DISPLAY CONSOLES,HCONLY or looks for the command response from DISPLAY CONSOLES,HCONLY.*

Before z/OS V1R10, specifying the DISPLAY CONSOLES,HCONLY command resulted in the reporting of message and SYNCH loss. As of z/OS V1R10, specifying the DISPLAY CONSOLES,HCONLY command results in message IEE535I being issued to indicate that an invalid DISPLAY parameter was specified.

**Migration action:** Do not issue the DISPLAY CONSOLES,HCONLY command anymore. Update any programs that automate on the command response to the DISPLAY CONSOLES,HCONLY command to no longer automate on it.

### **Update automation to accommodate new responses to the DISPLAY CONSOLES command (Required-IF as of R10)**

*Required if you have code that looks for message IEE889I, IEE890I, or IEE047I.*

Before z/OS V1R10, specifying the DISPLAY CONSOLES command resulted in the issuance of message IEE889I, IEE890I, or IEE047I. In addition, all subsystem consoles were displayed when the NACTIVE parameter was specified. No subsystem consoles were displayed when the ACTIVE parameter was specified (or defaulted). As of z/OS V1R10, specifying the DISPLAY CONSOLES command results in the issuance of new message CNZ4100I, CNZ4102I, or CNZ4103I. In addition, the DISPLAY CONSOLES system command displays active subsystem consoles when the ACTIVE parameter is specified (or defaulted) and inactive subsystem consoles when the NACTIVE parameter is specified.

**Migration action:** Update any programs that automate on message IEE889I, IEE890I, or IEE047I to automate on CNZ4100I, CNZ4102I, and CNZ4103I instead. Also take into account the change in behavior for subsystem consoles.

### **Remove PPT entries for the GRS ENQ/RESERVE/DEQ monitor (Recommended, as of R10)**

*Not required, but recommended if you have any PPT entries for the GRS ENQ/RESERVE/DEQ monitor to avoid possible confusion in the future.*

Before z/OS V1R10, users of the global resource serialization ENQ/RESERVE/DEQ monitor were required to ensure that its address space ran at the highest possible dispatching priority. Beginning with z/OS V1R10, the monitor address space is run at the appropriate dispatching priority. Therefore, any program properties table (PPT) entries (defined in parmlib member SCHEDxx) that are related to the monitor's dispatching priority (specifically, the PRIV keyword) are no longer required and can be removed, as they could be misleading. There is no longer any need to have a PPT entry for the monitor. The suggested program name (PGMNAME in the PPT) for the monitor was ISGAUDIT but you could change this name.

Note that, because this change makes the monitor address space an extension of the global resource serialization address space, the monitor address space's CPU time is added to SMF type 30 records for the global resource serialization address space.

**Migration action:** In parmlib member SCHEDxx, remove any PRIV|NOPRIV keyword that is specified for the global resource serialization ENQ/RESERVE/DEQ monitor. If you do not remove the keyword, it will be ignored, but not removing it could confuse someone who sees it in the future and does not realize that it is no longer functional. If a SYST keyword was also specified to ensure that the monitor would not time out, specify TIME=NOLIMIT on the JOB statement instead because TIME=NOLIMIT pertains to all tasks in the address space.

### **Modify GRS ISV-oriented exit ISGNQXITQUEUED1 (Required-IF as of R10)**

*Required if your software or any of your ISV software uses the ISGNQXITQUEUED1 installation exit.*

In z/OS V1R10, the parameter list for installation exit ISGNQXITQUEUED1 was changed. The NqqpRsc\_ReturnCode in the NQQP now contains the ISGENQ return code for ISGENQ requests, and the ENQ return code for ENQ/RESERVE requests. Before z/OS V1R10, the field contained an ENQ/RESERVE style return code for both ISGENQ and ENQ/RESERVE requests.

#### **Migration action:**

1. Determine whether the ISGNQXITQUEUED1 installation exit is being used on your system by using the DISPLAY PROG,EXIT command. The target module names can help identify who the exploiters are.
2. If you are using the exit, do the following:
  - If any of the exit routines are owned by ISVs that you did not contact for new z/OS V1R10 support, contact those ISVs to ensure that you have the latest updates.

- If any of the exit routines are owned by your installation, ensure that these exit routines have been modified to enable handling of the parameter list change.

Failure to modify ISGNQXITQUEUED1 could cause a data integrity problem or system failure.

### **Update automation that handles message IEE174I (Required-IF as of R10)**

*Required if you have automation routines that examine the message.*

The content of message IEE174I has changed. Before z/OS V1R10, the message might contain the following information:

```
ID CPU VF CR SERIAL
cpuid s s s ser_no
cpuid s s s ser_no
```

```
CPC ND = tttttt.nnn.mmm.pp.ssssssssss
CPC SI = tttt.nnnnnnnnnnnnnnnn.mmmmmmmmmmmmmmmmm.pppp.ssssssssssssssss
```

Starting in z/OS V1R10, the message might contain an additional line as follows:

```
ID CPU VF CR SERIAL
cpuid s s s ser_no
cpuid s s s ser_no
```

```
CPC ND = tttttt.nnn.mmm.pp.ssssssssss
CPC SI = tttt.nnnnnnnnnnnnnnnn.mmmmmmmmmmmmmmmmm.pppp.ssssssssssssssss
      Model: oooooooooooooooooo
```

### **Use the default for MEMLIMIT (Recommended, as of R10)**

*Not required, but recommended so that there is some nonzero MEMLIMIT, by default, on the amount of virtual storage above the 2 GB bar for an address space.*

Your MEMLIMIT value affects the virtual storage limit for a single address space above the 2 GB bar. Before z/OS V1R10, the MEMLIMIT default value was zero (meaning that the MEMLIMIT default was no virtual storage available above the 2 GB bar). As of z/OS V1R10, the MEMLIMIT default value has changed to 2 GB. This changed default provides a more usable value.

#### **Migration action:**

- To accept this new behavior, do nothing.
- If you wish to revert back to the prior MEMLIMIT default of zero, override the MEMLIMIT default to specify 0. But realize that a nonzero MEMLIMIT value is recommended, which is why the default value was changed in z/OS V1R10.

**Tip:** Use IBM Health Checker for z/OS to help determine whether your virtual storage limits are set properly. The check RSM\_MEMLIMIT checks the current setting for the MEMLIMIT parameter in SMFPRMxx, which affects the amount of virtual storage above the 2 GB bar that is available to jobs. This check verifies that a nonzero MEMLIMIT value is in use.



### **Provide an appropriate HZSPDATA data set on fallback (Required-IF as of R10)**

*Required if you use IBM Health Checker for z/OS and you fall back from z/OS V1R10 to z/OS V1R9 or z/OS V1R8.*

In z/OS V1R10, the format of data written to the HZSPDATA data set used by IBM Health Checker for z/OS has changed. Data written to HZSPDATA in z/OS V1R10 cannot be read by z/OS V1R9 or z/OS V1R8. Thus, if you fall back from z/OS V1R10 to z/OS V1R9 or z/OS V1R8 after you have started IBM Health Checker for z/OS on z/OS V1R10, you have to provide an appropriate HZSPDATA data set before you can use it on z/OS V1R9 or z/OS V1R8 because z/OS V1R10 formats HZSPDATA from the older format to the newer (z/OS V1R10) format.

Note that the z/OS V1R10 version of IBM Health Checker for z/OS has no problem reading HZSPDATA data that was written on a z/OS V1R9 or z/OS V1R8 system.

Once you use the HZSPDATA data set on z/OS V1R10, you cannot use it on z/OS V1R9 or z/OS V1R8. You must either provide a different data set (presumably a copy of the original z/OS V1R9 or z/OS V1R8 HZSPDATA) or reallocate (and lose the data in) the original data set.

**Migration action:** Make a copy of your current z/OS V1R9 or z/OS V1R8 HZSPDATA data set. Upon fallback to z/OS V1R9 or z/OS V1R8, specify that copy of HZSPDATA on your IBM Health Checker for z/OS started procedure. If you do not provide an appropriate HZSPDATA data set, message HZS0015E or HZS0013A is issued to indicate that the format of the current HZSPDATA data set is invalid. You will be prompted for specification of an HZSPDATA data set so that a check can write its persistent data.



**BCP Migration Actions Post-First IPL****Modify SMF exit routine IEFACTRT because of lengthened fields (Req-IF, as of R11)**

*Required if you use SMF exit routine IEFACTRT.*

A sample IEFACTRT exit routine is provided in SYS1.SAMPLIB in member IEEACTRT. This sample exit routine puts a job summary in the log of a job using WTO with routing code 14. The summary includes both step and job information. To accurately reflect data for long-running job steps, the following fields in the output generated by the sample IEFACTRT exit routine have increased in size:

Field name	Pre-z/OS V1R11 field length	z/OS V1R11 field length
TCB	6 bytes	9 bytes
SRB	6 bytes	9 bytes
SERV	6 bytes	13 bytes

If you use IEFACTRT, whether the IBM sample or a routine that you wrote, make sure that it can handle the increased field lengths.

**Migration action:** Review your IEFACTRT exit routine to make sure that the changed fields are long enough, and modify it if necessary. If you are using the IBM sample, rebase on the sample because it accommodates the longer field lengths. If you are using your own exit routine, verify that the longer field lengths are accommodated. Also review and change, if necessary, any automation programs that are affected.

**Keep the real frames included in GRS real storage consumption for use (Recommended, as of R11)**

*Not required, but recommended to avoid confusion caused by the expected change in the report GRS real storage usage.*

Before z/OS V1R11, Resource Management Facility (RMF) reports and other monitoring tools might indicate a significant growth in Global Resource Serialization (GRS) real frame usage compared with the previous release. Starting with z/OS V1R11, GRS's QSCAN buffer processing is changed to no longer require that frames backing discarded virtual pages be immediately returned to the RSM available frame queue, but are now taken by the system if needed. This approach improves the storage management of the buffers as they usually do not need to be re-backed as QSCANs come and go. Thus, GRS is not really using any more frames, and the frames are still available but can be taken if needed by RSM due to the storage being needed elsewhere. The real frame usage count includes both frames that are in use by the address space and ones that are readily available for the system to reclaim when they are needed.

**Migration action:** No action is necessary to accept the new behavior. To avoid any confusion, inform the system administrator.

**Accommodate changed GETMAIN and STORAGE OBTAIN behavior (Recommended, as of R10)**

*Not required, but recommended for better performance of programs that issue large numbers of GETMAIN or STORAGE OBTAIN requests in a user-region private area subpool.*

z/OS V1R10 provides a new statement in parmlib member DIAGxx that allows an installation to choose the allocation method used by the GETMAIN and STORAGE OBTAIN macros to obtain user-region private area subpools. The new DIAGxx statement (introduced by APAR OA27291) is "VSM USEZOSV1R9RULES(YES|NO)". As of z/OS V1R10, the default setting is YES, which keeps the GETMAIN and STORAGE OBTAIN behavior the same as in previous releases.

Some applications, such as DB2, have experienced performance problems because descriptor queue element (DQE) and free queue element (FQE) chains have become extremely long for certain user-region private area subpools. These applications can benefit from the new z/OS V1R10 GETMAIN and STORAGE OBTAIN behavior, which merges adjacent DQEs, when possible, and thus results in shorter DQE chains. You can enable the new behavior by specifying "VSM USEZOSV1R9RULES(NO)".

The new behavior could give the perception in some cases that storage is not being cleared to zero as it previously was. Storage is cleared by the system no differently in z/OS V1R10 (or later) than it was previously. The reason for the apparent difference is because storage requests can now be satisfied partly from an FQE on a page adjacent to a newly obtained DQE. If the storage represented by that FQE had previously been obtained and then freed, the

residual data remains. Programs that assume that storage will be cleared by the system even for cases where there is no guarantee of it could be affected. Depending on the contents of the uncleared storage and how it is used, this might result in ABEND0Cx, overlay of storage, or other problems. (See APAR OA27291 for examples and more information.) Restoring to the old behavior might be necessary if programs in the installation are making unwarranted assumptions about internal VSM behavior. In this case, specify "VSM USEZOSV1R9RULES(YES)" until the applications causing problems can be corrected.

If you have existing programs that use the GETMAIN or STORAGE OBTAIN macros to obtain user-region private area storage, use the information in the table below to understand the implications of using the new behavior.

<p>z/OS V1R9 and earlier, or z/OS V1R10 with the PTF for APAR OA27291 installed and VSM USEZOSV1R9RULES(YES) specified or defaulted</p>	<p>z/OS V1R11 and VSM USEZOSV1R9RULES(NO) specified, z/OS V1R11 and VSM USEZOSV1R9RULES(YES) specified or defaulted, z/OS V1R10 with the PTF for APAR OA27291 installed and VSM USEZOSV1R9RULES(NO) specified, or z/OS V1R10 without the PTF for APAR OA27291 installed.</p>
<p>Storage is more likely to be obtained from a fresh page (which by MVS rules is guaranteed to be cleared to binary zeros).</p>	<p>Storage requests are more likely to be fulfilled from areas that were previously obtained with the GETMAIN or STORAGE OBTAIN requests (which means they might contain residual data). Depending on the contents of the uncleared storage and how it is used, programs that do not clear the storage before use might experience abend X'0Cx', overlay of storage, or other unexpected problems. (See APAR OA27291.)</p>
<p>Storage is allocated from the top (higher address) of the page to the bottom (lower address).</p>	<p>Storage is allocated from the bottom (lower address) of the page to the top (higher address). Programs that depend on the previous allocation method might be affected. (See APAR OA27291.)</p>

<p>Unless a GETMAIN request can be satisfied entirely from an existing FQE, a new DQE must be obtained for each GETMAIN request.</p> <p>Very long DQE/FQE chains can result in performance degradation.</p>	<p>Storage requests can now be satisfied partly from an FQE on a page adjacent to a newly obtained DQE.</p> <p>New DQEs obtained by GETMAIN are merged with adjacent DQEs of the same subpool and key that are owned by the same task.</p> <p>Programs that depend on the previous allocation method for DQEs might be affected.</p>
---	--

**Migration action:**

1. If you previously coded a NUCLABEL ENABLE|DISABLE(IGVGPVTN) statement in your DIAGxx parmlib member as a circumvention for HIPER APAR OA27291, delete the statement because this syntax is no longer recognized with z/OS V1R11 or with the installation of the PTF for APAR OA27291 on z/OS V1R10. If you do not delete the NUCLABEL ENABLE|DISABLE(IGVGPVTN) statement, the following message will display: ASA008I ERROR IN PARMLIB MEMBER=DIAG00 ON LINE 32, POSITION 18: VALUE IS NOT VALID - NO MATCH FOUND. DETECTING MODULE IS IGVDIR. INPUT LINE: NUCLABEL ENABLE(IGVGPVTN)

2. To enable the new z/OS V1R10 and later storage allocation behavior, use the following line in the running DIAGxx parmlib member: VSM USEZOSV1R9RULES(NO). Then issue the SET DIAG=xx command to dynamically activate the new setting.
3. If it is necessary to back out the new storage allocation behavior, specify the following line in the running DIAGxx parmlib member: VSM UseZOSV1R9RULES(YES). Then issue the SET DIAG=xx command to dynamically activate this setting.

**Note:** As with most of the parameters in DIAGxx, the current behavior for parameter USEZOSV1R9RULES can only be changed by the presence of this parameter in DIAGxx. At IPL, the current behavior and the default is as if USEZOSV1R9RULES(YES) were specified. If SET DIAG=xx changes the value to NO, this value remains the current setting until explicitly changed. SET DIAG=xx, which does not have USEZOSV1R9RULES included, does not change the value back to the default of YES. Only an IPL will change USEZOSV1R9RULES back to the default value without the parameter being included in DIAGxx.



### **Update callers of CNZMXURF that do not hold appropriate serialization (Required-IF, as of R10)**

*Required if you have applications that use the CNZMXURF macro and access the UCME that CNZMXURF returns, and you want to ensure that application errors do not occur because of a failure to hold serialization.*

The CNZMXURF macro locates the console control block (UCME) that contains a specific console ID. Applications should hold serialization while using CNZMXURF and accessing the UCME that CNZMXURF returns.

Before z/OS V1R10, when console services shared mode was the only mode available, the UCME contents did not change, so applications that did not hold serialization ran without error.

Beginning with z/OS V1R10, console services distributed mode is added, and it is now more important that applications hold serialization. When the system is in distributed mode (the default is shared mode) and a UCME is accessed without serialization, the contents of UCME storage could become cleared while being examined by the application. In addition, if the application is examining a UCME for an active console (when the system is in distributed mode) and the console becomes inactive, the UCME is cleared. If serialization is not held, the application can fail.

With z/OS V1R10 (and APAR OA19948 on z/OS V1R9 and prior releases), the Tracking Facility was enhanced so that it reports applications that call the CNZMXURF service without holding appropriate serialization. You should run the Tracking Facility to find the applications that need to be updated. In Tracking Facility output, look for violations that start with "UCME-CMS".

#### **Migration action:**

1. Obtain the Tracking Facility exclusion list at <http://www.ibm.com/eserver/zseries/zos/downloads/>.
2. Install the PTF for APAR OA19948 on z/OS V1R9 or z/OS V1R8 so that you can use the Tracking Facility to detect inappropriate serialization.
3. Activate the Tracking Facility. Use the exclusion list you obtained in the step above.
4. Periodically issue the DISPLAY OPDATA,TRACKING command to obtain the results of the tracking.
5. Forward the output to [consoles@us.ibm.com](mailto:consoles@us.ibm.com) so that the exclusion list can be updated for future users of the Tracking Facility.
6. Obtain necessary modifications for the affected programs from the program owners.

### **Modify applications that receive IOS ENF signal 10, 31, 32, or 33 for PPRC devices (Required-IF as of R10)**

*Required if you receive ENF event code 10, 31, 32, or 33.*

As of z/OS V1R10, peer-to-peer remote copy (PPRC) primary and secondary devices can be defined in subchannel set 1. Listeners for IOS ENF signals that include either device numbers or UCB addresses for PPRC primary or secondary devices need to qualify the device numbers for these devices with a subchannel set ID. The ENF event codes may contain these devices in their parameter lists. Any code directly using a device number from these ENF signal parameter lists or using UCBCHAN from the UCB addresses in the parameter lists to identify a device will also need to refer to the subchannel set ID supplied in the parameter list or UCBSSID from the UCB address to correctly identify the device.

**Migration action:** Look for a parameter list received by an ENF listen exit that is established by an ENFREQ ACTION=LISTEN for ENF event code 10, 31, 32, or 33. Based on the ENF event code, do the following:

- **10** If the UCBCHAN from the UCB addresses in the parameter list is used to identify a device, the UCBSSID is required to be used to correctly identify the device. The UCBSSID can be referenced using the IOSDUPI mapping macro based on the UCBIEXT pointer.
- **31 or 32** If the DCCD parameter list is being processed, and the DCCDEYYP=DCCDDEV entries are used, use the DCCDEYYP=DCCDSDEV as well. Use the DCCDDEVN from the DCCDSDEV entry and DCCDSSID from the DCCDDEV entry in combination to correctly identify the device. The DCCDSDEV entries are the same as the DCCDDEV entries but they have a DCCDEYYP=DCCDDEV entry immediately following them.
- **33** When processing the DACH parameter list, if the DACHQN=DACHPAV entries are being used, the DACHQN=DACHPAV\_AS entries should also be used, and if the DACHQN = DACHIORA entries are being used, the DACHQN=DACHIORA\_AS entries should also be used. Use the DACH\_PAV\_DEVN from the DACHPAV\_AS entries and DACH\_PAV\_SSIDBase in combination to correctly identify the device. The DACH\_IORA\_DEVN from the DACHIORA\_AS entries should be used in combination with DACH\_IORA\_SSID to correctly identify the device.

DACHTYPED	DACHQN if device is in subchannel set 0	DACHQN if device is not in subchannel set 0	Meaning
LPE	DACHIORA	DACHIORA_AS	I/O resource available
PAVS	DACHPAV	DACHPAV_AS	Change in PAV UCB

For any listeners of ENF signal 10, 31, 32, or 33, use the new SPECIAL(YES) keyword and the SCHSET keyword on UCBLLOOK or UCBSKAN invocations.



**Reallocate the WLM couple data set (Required-IF as of R10)**

Required if the current WLM couple data set is not significantly larger than required for the currently used WLM service definition so that the data set can accommodate the increased size of the WLM policy structures.

In z/OS V1R10, you might need a larger workload management (WLM) couple data set because the WLM service class and workload records in the WLM policy are longer, as shown below:

IWMSVPOL entry	Pre-z/OS V1R10 entry length	z/OS V1R10 entry length
Policy information	168	360
Service class information	64	144
Workload information	40	112

**Migration action:** Before you activate a WLM policy on z/OS V1R10, do the following:

1. Allocate a new WLM couple data set using the IXCL1DSU utility, as described in "Allocate a WLM Couple Data Set" in *z/OS MVS Planning: Workload Management*. Use the current NUMBER specification for each section (POLICY, WORKLOAD, SRVCLASS, and so forth).  
**Note:** By using z/OS V1R10 to allocate your WLM couple data set, the space allocated will be sufficient for policy activation on z/OS V1R10. Ensure that the values provided for number of policies, workloads, and service classes are the current values because these values will be used by WLM to calculate the space required. An overspecification to allow for growth is acceptable.
2. Switch to the new WLM couple data set or sets as described in "Make a Newly Formatted Couple Data Set Available to the Sysplex" in *z/OS MVS Planning: Workload Management*.
3. Update the COUPLExx parmlib member to specify the new WLM couple data set or sets.

During policy activation, if WLM determines that the WLM couple data set is too small, the following message is issued: IWM003I VARY WLM FAILED, WLM COUPLE DATA SET TOO SMALL. You would then have to allocate and active a new, larger WLM couple data set.

**Note:** On z/OS R10 when you invoke the WLM Administration Application, the service definition in the partitioned data set will automatically update the ISPF tables to the z/OS R10 level. (Note that the ISPF table level in the partitioned data set is not the same as the WLM service definition functionality level). Once the ISPF table is updated, it is not possible to open or edit this service definition from lower level systems. See *MVS Planning: Workload Management* book for more information about this restriction.

**Accommodate the change of directory information for PDSE program objects with authorization code zero (Required-IF, as of OA12579 on R7/R8/R9)**

Required if you use IEHLIST LISTPDS output or examine the directory entries to determine the authorization requirement of PDSE program objects.

If the authorization code for a program object (PDSE) is not specified or is specified as zero during bind, the program management binder builds a directory entry with the bit set indicating that the authorization code is not present. However, this is not consistent with load module (PDS) processing and the IEHLIST LISTPDS command shows inconsistent information for PDS modules and PDSE program objects where the authorization code is either not specified or is specified as zero (using the binder AC option or SETCODE control statement). For PDSEs, the information is blank unless a nonzero authorization code is specified. For PDSs, the Auth Req column contains NO if the authorization code is zero or unspecified. As of z/OS V1R9 and APAR OA21759, the binder always sets a valid authorization code in the directory for a program object (PDSE) with the value specified by the user on the AC option or SETCODE control statement, or the default value if none is specified.

**Migration action:** Ensure that the change of directory information for the PDSE program object with authorization code zero does not affect your existing processing. For example, if your program reads the Auth Req column listed by the IEHLIST LISTPDS command for the PDSE program object, it needs to be adjusted to handle the new display of NO, which means AC=00.





**Notes:** 1. To update the directory information and achieve the improved IEHLIST LISTPDS command output (new display of NO), existing program objects in PDSE data sets must be relinked with the program management binder after applying the PTF for APAR OA21759.

2. If there is no APF code stored in the directory entry for a program object, ISPF displays 00 in the APF authorization code column on the member list panel. This functionality was introduced by ISPF APAR OW52063, which was integrated in z/OS V1R5 and does not depend on this enhancement to the program management binder.

## Communications Server Migration Actions for z/OS R11 from z/OS R9



### Migration Actions You Can Do NOW:

- 
 • **IP Services: Migrate from Network Database function (Req-IF, as of R11)**
  - Change your processes to use the distributed data facility (DDF) provided by DB2 for z/OS, and the DB2 Runtime Client.
- 
 • **IP Services: Migrate from the Dynamic Host Configuration Protocol server (Req-IF, as of R11)**
  - Investigate using a DHCP server on Linux for System z.
- 
 • **IP Services: Migrate from the Boot Information Negotiation Layer function (Req-IF, as of R11)**
  - Consider using the product IBM Tivoli Provisioning Manager for OS Deployment V5.
- 
 • **IP Services: Migrate from BIND DNS 4.9.3 (Req-IF, as of R11)**
  - For a caching-only name server, use z/OS Resolver DSN Caching function in R11 instead.
  - Some /etc configuration files have changed for this removal.
  - BIND DNS 9.2.0 is planned to be removed in a future release.

20

z/OS 1.11: Migration - ... And It's Good! Part 2 of 2

© 2009 IBM Corporation

## Communications Server Migration Actions for z/OS R11 from z/OS R9



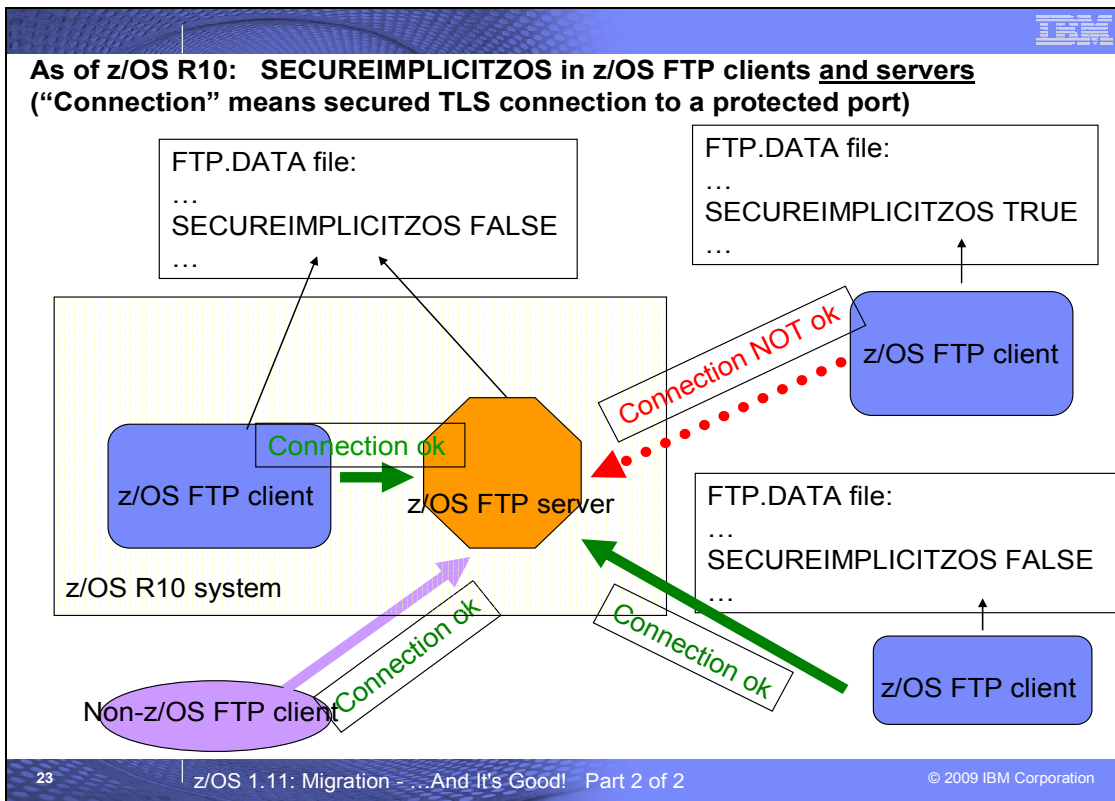
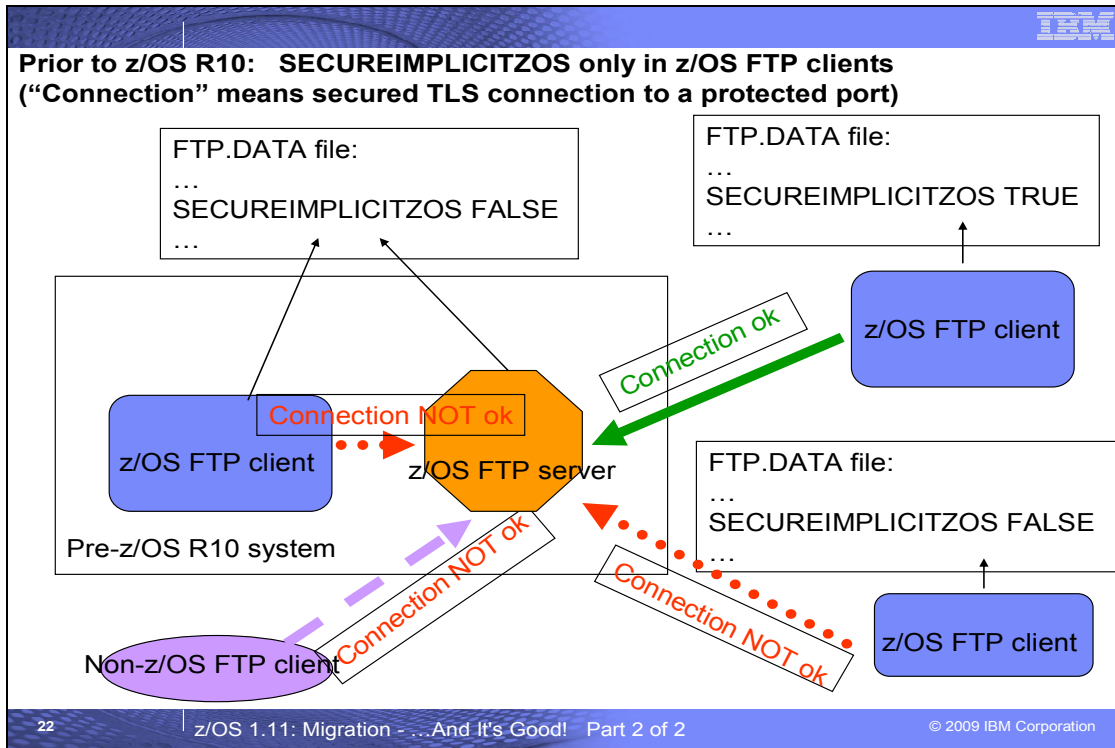
### Migration Actions You Can Do NOW:

- **IP Services: Migrate FTP servers sharing FTP.DATA with FTP clients (Required-IF, as of R10)**
  - Before R10, only FTP client supported the **SECUREIMPLICITZOS** statement.
    - This statement specified when the FTP client should negotiate the security handshake with the FTP server for **TLSPORT**, implicit connections. **TRUE** is the default.
  - As of R10, FTP server enhanced to support implicitly secured TLS connections from non-z/OS FTP clients. Thus, **SECUREIMPLICITZOS** is now supported in the FTP server.
  - Sharing (or not sharing) the **FTP.DATA** between z/OS FTP client and server may have different behaviors....see next foils.
  - **Know new rules for behavior!**
    - ❖ Same value for successful connection.
    - ❖ Non-z/OS FTP clients need **SECUREIMPLICITZOS FALSE**.

21

z/OS 1.11: Migration - ... And It's Good! Part 2 of 2

© 2009 IBM Corporation



## Communications Server Migration Actions for z/OS R11 from z/OS R9



### Migration Actions Before Install:

- Accommodate **DEFAULT** route usage changed for **CINET** environments (**Re-IF, as of R11**)
  - Before R11, the **DEFAULT** route was tracked only as a flag for each stack, indicating that the stack had a **DEFAULT** route defined. These routes were masked by selection logic.
  - As of R11, the **DEFAULT** route is treated like all other subnet and network routes, including the potential for multiple duplicate routes on a stack, having an associated cost, and selection of a stack with the least cost for that route (which is generally desirable). However, on systems where multiple stacks have a **DEFAULT** route defined, they can now be selected causing undesirable results (including connection timeouts).
  - Follow instructions in the FlashAlert (and in handout).

### Migration Actions Pre-First IPL:

- Update procedures that use the syslogd job name (**Required-IF, as of R11**)
  - Before R11, the job name of the syslog daemon was the cataloged proc name plus an additional character (usually "1"), when started from a proc and that proc was less than 8 chars in length. An additional char was also appended to the program name when started from the z/OS UNIX shell.
  - As of R11, syslog daemon job name is the same as the proc name or the program name.
  - Update any automation. If you run two copies, ensure they will be uniquely named. Update the syslogd UDP port reservation in the TCP/IP profile to use the syslogd proc name.

24

z/OS 1.11: Migration - ...And It's Good! Part 2 of 2

© 2010 IBM Corporation

## Communications Server Migration Actions for z/OS R11 from z/OS R9



### Migration Actions Pre-First IPL:

- **IP Services: Run the syslogd daemon in the background when starting from the UNIX shell (Required-IF, as of R11)**
  - As of R11, syslogd does not automatically run in the background when started from the **UNIX** shell (or from a shell script like **/etc/rc**) or in **BPXBATCH**.
    - If it is started from **/etc/rc** or another shell script, or from **BPXBATCH**, add a trailing ampersand character.
    - If you start syslogd manually from the **UNIX** shell, include a trailing ampersand on the command line.
    - If you start syslogd from a cataloged procedure that uses **BPXBATCH**, add a trailing ampersand character, and after that include a sleep command in your script to give syslogd time to initialize before the shell script ends.
    - If you start syslogd from a cataloged procedure that does not use **BPXBATCH**, no changes are needed.

25

z/OS 1.11: Migration - ...And It's Good! Part 2 of 2

© 2010 IBM Corporation



## Communications Server Migration Actions for z/OS R11 from z/OS R9



### Migration Actions Pre-IPL:

- **IP Services: Accept new resolver behavior of saving the results of DNS queries (Recommended, as of R11)**
  - Prior to R11, the z/OS system resolver did not cache results of DNS queries.
  - As of R11, the resolver automatically saves these results.
  - New behavior requires no action! This gives you performance and usability enhancements.
  - If you want prior behavior (disable resolver DNS caching function) or you want to selectively limit the caching function, you must take action.
- **SNA Services: Ensure that dump data sets are large enough to include VTAM internal trace data space (Required, as of R11)**
  - Prior to R11, VTAM INOP dump processing did not capture VIT data space in the dump.
  - As of R11, this processing automatically captures the VIT data space, when VIT is in use.
  - Ensure that the dump data set allocation will accommodate this add'l space.
  - Health Check determines whether the VIT data space table size is 50 MB (5).



26

z/OS 1.11: Migration - ...And It's Good! Part 2 of 2

© 2009 IBM Corporation

## Communications Server Migration Actions for z/OS R11 from z/OS R9



### Migration Actions Pre-First IPL:

- **IP and SNA Services: Increase region size for Telnet (Required-IF, as of R10)**
  - As of R10, ECSA usage for Telnet is reduced. However, Telnet private storage usage increased. The private storage increase depends upon the number of Telnet sessions per LPAR.
  - Determine the number of Telnet sessions that will exist on the LPAR.
  - Increase the REGION size on the Telnet started task by 1 MB for each 25,000 Telnet sessions.
  - Examine the values for your CRPLBUF pool defined in the ATCSTRxx VTAM start list, particularly the base size. If the base size is large because of a large number of Telnet sessions, you can reduce it. Failure to reduce it can minimize the amount of ECSA savings realized.
- **SNA Services: Increase region size for VTAM (Required-IF, as of R11)**
  - As of R11, ECSA usage for HPR is reduced. However, VTAM private storage usage has increased. The increase depends upon the number of RTP endpoints that are supported on the LPAR. If VTAM is using nearly all its allocated region size, you'll need to increase it.
  - Increase region size for VTAM by 1.5 MB for each 1,000 RTP pipes.

27

z/OS 1.11: Migration - ...And It's Good! Part 2 of 2

© 2010 IBM Corporation

## Communications Server Migration Actions Between z/OS V1 R9 and z/OS V1 R11

These migration actions were taken from *z/OS Migration*. Many descriptions and actions have been severely shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### Communications Server Migration Actions You Can Do Now



#### **IP Services: Migrate from Network Database function (Required-IF, as of R11)**

*Required if you use the NDB function.*

z/OS V1R10 was the last release in which z/OS Communications Server supported the Network Database (NDB) function. Starting with z/OS V1R11 you have to use the distributed data facility (DDF) provided by DB2 for z/OS, and the DB2 Runtime Client.

**Migration action:** Change your processes to use the distributed data facility (DDF) provided by DB2 for z/OS, and the DB2 Runtime Client. DDF allows client applications running in an environment that supports Distributed Relational Database Architecture™ (DRDA®) to access data at DB2 servers.

**Tip:** Use IBM Health Checker for z/OS to determine if the NDB function is in use. The check is named ZOSMIGV1R10\_CS\_NDB



#### **IP Services: Migrate from the Dynamic Host Configuration Protocol server (Required-IF, as of R11)**

*Required if you use the DHCP server function.*

z/OS V1R10 was the last release in which z/OS Communications Server supported the Dynamic Host Configuration Protocol (DHCP) server function. Starting with z/OS V1R11, you have to use a DHCP server on a system other than your z/OS system..

**Migration action:** You should implement the DHCP server on a system other than your z/OS system. If you want to continue to run the DHCP server on your System z hardware, you can implement a DHCP server inside a Linux for System z image. DHCP servers are widely available on a variety of platforms.

**Tip:** Use IBM Health Checker for z/OS to determine if the DHCP function is in use. The check is named ZOSMIGV1R10\_CS\_DHCP.



#### **IP Services: Migrate from the Boot Information Negotiation Layer function (Required-IF, as of R11)**

*Required if you use the BINL function.*

z/OS V1R10 was the last release in which z/OS Communications Server supported the Boot Information Negotiation Layer (BINL). Starting with z/OS V1R11, you have to use another product, such as IBM Tivoli Provisioning Manager for OS Deployment V5 (5724-Q99), for network-based operating system installation services.

**Migration action:** Consider using the product IBM Tivoli® Provisioning Manager for OS Deployment V5 (5724-Q99) for network-based operating system installation services.

**Tip:** Use IBM Health Checker for z/OS to determine if the BINL function is in use. The check is named ZOSMIGV1R10\_CS\_BINL.



#### **IP Services: Migrate from BIND DNS 4.9.3 (Required-IF, as of R11)**

*Required you use BIND DNS 4.9.3.*

z/OS V1R10 was the last release in which z/OS Communications Server supported BIND DNS 4.9.3.

##### **Migration action:**

- If you have been using BIND DNS 4.9.3 as a caching-only name server, use the z/OS Resolver DNS Caching function in z/OS V1R11 instead.
- If you have been using BIND DNS 4.9.3 as a primary or secondary authoritative name server, migrate to BIND DNS on Linux for System z.
- If you want to implement a temporary replacement, implement BIND DNS 9.2.0, which has been included in z/OS beginning with V1R4. Use the dnsmigrate tool to migrate a z/OS BIND 4.9.3 DNS server configuration file to a z/OS BIND 9.2.0 DNS server configuration file. Note that support for BIND DNS 9.2.0 is planned to be removed in a future release.
- If you exploit the Connection Optimization (DNS/WLM) feature of BIND 4.9.3, you have to use an alternative solution. One alternative would be using the sysplex distributor function, which is implemented using distributed dynamic VIPAs (VIPADISTRIBUTE statement in the TCPIP PROFILE). Another alternative is the z/OS Load Balancing Advisor in combination with the Automated Domain Name Registration application

(ADNR), which provides the ability to automatically update DNS with the availability status of sysplex resources.

**Tip:** Use IBM Health Checker for z/OS to determine if a BIND DNS name server function is in use. The check for the BIND DNS 4.9.3 is named ZOSMIGV1R10\_CS\_BIND4 and the check for the BIND DNS 9.2.0 is ZOSMIGV1R11\_CS\_BIND9.

**Statement of Direction:** In a future release of z/OS, the BIND 9.2.0 function will be removed from the z/OS Communications Server component. Customers who currently use or plan to use the z/OS BIND 9.2.0 function as a caching-only name server should use the Resolver cache function, which will be available in z/OS V1.11, to cache DNS responses. Customers who currently use or plan to use the z/OS BIND 9.2.0 function as a primary or secondary authoritative name server should investigate using BIND on Linux for System z.



**IP Services: Accommodate DEFAULT route usage changed for CINET environments (Required-IF, as of R11)**

*Required if your system has been identified as affected (see below). Systems running a single TCP/IP stack will not be affected.*

(This Communications Server Flash Alert can be found at: <http://www-01.ibm.com/support/docview.wss?uid=swg21410999&mysns=z000&mynp=OCSG004CE&mync=OCSWG90&mync=E>)

When running a system with multiple TCP/IP stacks active (Common INET environment), UNIX System Services keeps a table of routing definitions to determine which stack should be used for a new outbound connection (the prerouter function). After upgrading to z/OS 1.11, the potential exists for the stack selection to be different for applications that use the DEFAULT route (also called 'route of last resort', and indicated by a subnet mask of all 0s) than had been selected on earlier releases. This can cause some applications to behave differently, including having connection failures and timeouts.

**Details:** In releases prior to z/OS 1.11, the DEFAULT route was tracked only as a flag for each stack, indicating that the stack had a DEFAULT route defined. With z/OS 1.11, this has been changed to treat the DEFAULT route like all other subnet and network routes, including the potential for multiple duplicate routes on a stack, having an associated cost, and selection of a stack with the least cost for that route.

Selecting the stack with the least cost path to the requested destination address is generally desirable. However, on systems where multiple stacks have a DEFAULT route defined, these might exhibit different behavior with an upgrade to z/OS 1.11. Where before the existence of these routes was masked by the selection logic used in earlier releases, they can now be selected causing undesirable results (including connection timeouts).

The problem occurs when one or more stacks have a DEFAULT route defined, but they do not actually have network connectivity to get to all addresses. This typically occurs when the devices used by a stack have connectivity only to a subset of the total network (such as an isolated laboratory LAN), and a DEFAULT route is coded to cover associated addresses that are not directly connected to that LAN (require going through a router to other LAN segments). If this is done using a GATEWAY or BEGINROUTES section of the TCPIP PROFILE, this route will be listed with the lowest cost. If the stack that actually does have connectivity to the entire network learns its DEFAULT route using dynamic routing (OMPROUTE), its route will have a higher cost. When this occurs, the lower cost stack will always be selected in z/OS 1.11. If the desired route is on the default stack, it would have always been selected in earlier releases.

**Determining whether your system is affected:**

1. Examine the [BPXPRMxx](#) PARMLIB member, and find the [NETWORK](#) statement with DOMAINNAME(AF\_INET) specified. If this statement has TYPE(INET) specified, this system will be unaffected by the upgrade. If it has TYPE(CINET), then continue to the next step.

Alternatively, you can do a [DISPLAY OMVS,PFS](#) operator command (see the [BPXO046I](#) message for a description of the output). If the PFS TYPE for AF\_INET SOCKETS (typically the first or second entry) shows **INET**, your system is configured for a single stack and will not be affected. If it shows **CINET**, continue to the next step.

2. Do a [DISPLAY TCPIP](#) operator command with no options specified. It will list the names of the TCPIP stacks that have ever been active since the last IPL. If there is only one name listed, then your system will be unaffected as long as no other stack is started.
3. In the BPXPRMxx member, examine the list of SUBFILESYSTYPE statements with TYPE(CINET). Note which one has the DEFAULT keyword specified. If none of them do, then the new stack selection logic (while potentially different) will not have a significant effect.

If using the alternative DISPLAY OMVS,PFS output, note the FLAGS in the second section. If one has **CD**, then either no default stack has been specified or the one that is specified is not active (and will have an **SD** flag). If one has **SC**, then continue with the next step.

4. Do a [DISPLAY OMVS,CINET](#) operator command (see the [BPXO047I](#) message for a description of the output). Note which stacks in the HOME INTERFACE list have the **DRS** flag. If one of them is the default stack (determined above), then continue with the next step.

If the system is already running z/OS 1.11, then the **DRS** flag will not appear. Instead go to the end of the HOST ROUTE INFORMATION section, where there will be one or more entries with NET DESTINATION 000.000.000.000 and NET MASK 000.000.000.000. Skip step 5 and use the information from these entries for step 6.

5. Display the routing table for each active stack on the system (listed in the output from step 2). This can be done either from TSO ([NETSTAT ROUTE DETAIL TCP stackname](#)) or the operator console ([DISPLAY TCPIP\\_stackname.NETSTAT.ROUTE.DETAIL](#)). Note the entries listed with Destination Default, which will be at the beginning of the report.
6. If the default routes listed for the default stack all have a lower METRIC value than those for the other stacks, then your system will be unaffected by the upgrade as long as that does not change.
7. Test the accuracy of the default routes by using the FTP command to connect to a server whose address will not cause selection of any other route listed by step 4. From TSO enter the command [FTP -p stackname.xx.xx.xx.xx \(TIMEOUT 15](#), first using the default stack (verify the connection works) and then using all other stacks that have a default route. If the command has a timeout making the connection, it is likely that the default route through that stack is invalid.

#### Migration action:

For systems configured for CINET, review the routing definitions before upgrading to later releases to ensure that nonexistent routes are not listed. To avoid this problem, ensure that the DEFAULT route is not configured on stacks that have limited network connectivity. For those stacks, the correct configuration is to explicitly list the subnets that are directly or indirectly connected to the devices used without coding a DEFAULT route.

When an undesired DEFAULT route is being learned through dynamic routing, identify the source of that route and update the configuration of the router advertising it. For OMROUTE, the DEFAULT route will be advertised in one of the following conditions:

- Originate\_Default\_Route=YES is specified on the [AS\\_Boundary\\_Routing](#) or [IPv6\\_AS\\_Boundary\\_Routing](#) statement when using OSPF
- The [Originate RIP Default](#) or [IPv6 Originate RIP Default](#) statement is coded when using RIP

Remove this specification where this is not appropriate.

#### Background information:

The stack selection logic performed by OMVS is as follows:

- The prerouter table affects only connections outbound from the z/OS system when stack affinity has not otherwise been established. Inbound connections are always associated with the stack on

which the packets arrived. That is likely why only a subset of applications are affected; the others are servers that service inbound connections.

- If the application binds to a specific local address, then the stack owning that address is used for that socket from that point forward. Most client (outbound) applications do not do an explicit bind.
- On the `connect()` call (or `sendto()` call for UDP traffic), the requested remote address is used to search the prerouter table. It first finds the best match; a host route for the address (if one exists), then a subnet route with the most specific mask (most bits ON).
- If multiple stacks list that matching route, the one with the smallest cost (the METRIC value) is used.
- If multiple stacks have that lowest cost and one of them is the default stack, then the default stack is used. If not, the stack selected will be one of the set (it is not defined which one, probably the first encountered in the prerouter table).
- If no other route matches, then the default route is used. Here is where the difference between the release levels occurs:
  - For z/OS 1.10 and earlier, the prerouter table only tracks the existence of a default route on each stack (the DRS flag). If one of those is the DEFAULT stack, then that will be used; otherwise, one of the others is selected.
  - For z/OS 1.11 (and future releases), the prerouter table lists the default route the same way as it does all other routes (destination 0.0.0.0, mask 0.0.0.0, with a metric and associated stack). So at this point, the rules for selection become the same as described for selection of a subnet route.

### **IP Services: Modify applications to no longer add IPv6 type 0 routing headers to outgoing packets (Required-IF, as of R11)**

*Required if you run applications on the TCP/IP stack that build IPv6 type 0 routing headers in outgoing packets.*

The IPv6 type 0 routing header is deprecated, as described in RFC 5095 (<http://www.ietf.org/rfc/rfc5095.txt>).

Beginning with z/OS V1R11, an application will receive an error if it attempts to add an IPv6 type 0 routing header to outgoing packets. If your installation is affected by this change, you must update applications that run on the TCP/IP stack to remove the use of the IPV6\_RTHDR option to build IPv6 type 0 routing headers. **Note:** The IPv6 type 0 routing header is rarely used.

#### **Migration action:**

1. Review RFC 5095 (<http://www.ietf.org/rfc/rfc5095.txt>) for a description of the need for deprecating the IPv6 type 0 routing header.
2. Identify any applications you run on your TCP/IP stack that use the IPV6\_RTHDR option to build IPv6 type 0 routing headers in outgoing packets.
3. For each application identified in Step 2, determine if changes are required in order to continue using the application on a stack that does not support IPv6 type 0 routing headers.
4. For each application that requires changes, complete those changes or discontinue use of the application.

### **IP Services: Update automation to accommodate FTP output that is changed for extended address volumes (Required-IF, as of R11)**

*Required if you use FTP to display statistics about the space on any volume or to display information about MVS data sets with a common high-level qualifier, and you process the space statistics or information programmatically.*

The output of several FTP subcommands and commands has changed in support of extended address volumes:

- The FTP client LOCSITE and SITE subcommands with the QDISK parameter, and the FTP server SITE command with the QDISK parameter, display statistics about the amount of space available on a volume. In z/OS V1R11, the output format has been changed to accommodate the larger amounts of free space available on extended address volumes.
- The FTP client DIR subcommand and the FTP server LIST command display a list of directory entries or a list of files in a file group on a remote host, or a list of the members of the partitioned data set, as well as auxiliary information about the files. In z/OS V1R11, the output format has been changed to accommodate the larger data sets supported by MVS.

You must change FTP client API programs and automation tools to recognize the new formats. Note that the new formats are used for all volumes.

#### **Migration action:**

- Be aware of the changes in messages and replies that FTP uses to report space statistics for volumes:
  - The following messages have changed: EZA2192I, EZA2193I, and EZA2194I.
  - The following replies have changed:

```
200-          Percent Free Free Largest   Free
200- Volume Free    Cyls Trks Cyls--Trks Exts   Use Attr
```

- Be aware of changes in the information the z/OS FTP client reports in message EZA2284I when you use the DIR subcommand to list information about MVS data sets.
- Update any automated processing tools or FTP client API applications to recognize the format of messages and replies issued by FTP to report space statistics for volumes.

### **IP Services: Update procedures that use the syslogd job name (Required-IF, as of R11)**

*Required if you have automation or procedures that use the syslogd job name, or if you reserve the syslogd UDP port using the job name.*

Before z/OS V1R11, the job name of the syslog daemon (syslogd) was the name of the cataloged procedure plus an additional character (usually the character “1”) when syslogd was started from a procedure and when the procedure name was less than 8 characters in length. An additional character was also appended to the program name when syslogd was started from the z/OS UNIX shell. Beginning with z/OS V1R11, the syslogd job name is the same as the procedure name or program name.

#### **Migration action:**

- Update any automation or other procedures that use the syslogd job name. If you run two copies of syslogd, be sure to use a unique job name for each copy.
- Update the syslogd UDP port reservation in the TCP/IP profile to use the syslogd procedure name.

### **IP Services: Accept the new behavior of TCP receive buffer size (Required-IF, as of R11)**

*Required if you have inbound streaming TCP connections over networks with large bandwidth and high latency.*

Before z/OS V1R11, there was no automatic buffer size tuning of inbound streaming TCP connections. Starting in z/OS V1R11, the TCP stack automatically tunes the ideal window size for inbound streaming TCP connections over networks with large bandwidth and high latency. The stack enables the TCP receive buffer size to grow as large as 2 MB for such TCP connections, irrespective of the TCPMAXRCVBUFRSIZE value. The pre-V1R11 behavior is no longer possible; this is a performance enhancement.

**Migration action:** Update applications that process output from the Netstat ALL/-A command so they can handle the larger values for the reported RcvWnd, ReceiveBufferSize, and ReceiveDataQueued fields.



### **IP Services: Update IP filter policy to filter IP fragments correctly for RFC 4301 compliance (Required-IF, as of R10)**

*Required if you enable RFC 4301 compliance and your IP filter policy selectively matches routed traffic based on TCP port, UDP port, ICMP type and code, ICMPv6 type and code, or OSPF type. The migration action is also recommended if your IPsec policy is noncompliant because RFC 4301 compliance might become mandatory in a future release of z/OS.*

IP filter policy support for filtering fragments was improved in z/OS V1R10. Before z/OS V1R10, Communications Server filtered all IP fragments using a policy of first possible filter match, and filtered IPv6 fragments as protocol IPv6Frag. Beginning in z/OS V1R10, Communications Server follows rules and restrictions established by RFC 4301 (<http://www.ietf.org/rfc/rfc4301.txt>) to ensure proper classification of fragments. RFC 4301, “Security Architecture for the Internet Protocol”, specifies the base architecture for IPsec-compliant systems, including restrictions on the routing of fragmented packets.

Communications Server does not implement stateful fragment checking; therefore, restrictions are added as required by RFC 4301 to optionally require that IP filter rules applying to routed traffic not apply to specific ports, types, or codes. These restrictions are globally configured as to whether or not they are applied. Additionally, support is added for use of OPAQUE protocol selector values for IPv6 fragments in which the upper layer protocol value is not known.

**Note:** In a future release of z/OS, RFC 4301 compliance might become mandatory. Changing an IPsec policy from being noncompliant to compliant might require minor changes to IP filters for IP traffic that is routed through z/OS. The Configuration Assistant for z/OS Communications Server includes functions to assist with identifying and making such changes.

**Migration action:** If you configure the RFC4301Compliance parameter on the IpFilterPolicy statement as No, or use the default value of No, no further action is needed at this time.

If you configure the RFC4301Compliance parameter on the IpFilterPolicy statement as Yes, perform the steps that are appropriate for you as described in *z/OS Migration*.

**Tip:** Use IBM Health Checker for z/OS check ZOSMIGV1R11\_CS\_RFC4301 to determine whether IPsec filter rules are compliant with RFC4301. This check is available in z/OS V1R10 with the installation of PTF UA47191 (APAR OA28605).



### **IP Services: Remove customization of SNMP sysObjectID MIB object in OSNMPD.DATA File**

#### **(Recommended)**

*Recommended because the ability to customize the sysObjectID value will not be supported in a future release.*

The SNMP agent allows you to provide some initial settings for a small set of MIB objects by using the OSNMPD.DATA file. One of the objects for which an initial value can be provided is sysObjectID.0. The sysObjectID.0 object is the vendor's authoritative identification of the network management subsystem contained in the entity. That is, it is intended to uniquely identify the SNMP agent. Changing this value is not recommended and the ability to change it will be disabled in a future release. As of z/OS V1R4, warning message EZZ6317I is written to the syslog daemon if the object is set by using the OSNMPD.DATA file. As of z/OS V1R8, message EZZ6317I is also written to the console.

**Migration action:** Review the statements in your OSNMPD.DATA configuration file. If this file contains a statement for the sysObjectID object, remove the statement from the file.

### **IP Services: Migrate FTP servers sharing FTP.DATA with FTP clients (Required-IF, as of R10)**

*Required if your FTP server and FTP client use a common FTP.DATA file or data set, and you have coded SECUREIMPLICITZOS FALSE in FTP.DATA.*

Before z/OS V1R10, only the FTP client supported the SECUREIMPLICITZOS statement. In z/OS V1R10, the FTP server has been enhanced to support implicitly secured TLS connections from non-z/OS FTP clients. Thus, the server will start supporting the SECUREIMPLICITZOS statement. If you want to keep the behavior of past releases, you must take action.

When SECUREIMPLICITZOS FALSE is coded in FTP.DATA, the server will support implicitly secured TLS connections when non-z/OS clients connect to the server's protected port as specified by the TLSPOrt statement in FTP.DATA. The client that shares FTP.DATA with the server will now be able to log into the server's protected port. z/OS FTP clients that do not share FTP.DATA with the server could have configured SECUREIMPLICITZOS TRUE, the default. These clients will no longer be able to log into a z/OS FTP server using the protected port when the server shares FTP.DATA with an FTP client that has specified SECUREIMPLICITZOS FALSE in FTP.DATA.

**Migration action:** Do one of the following:

- Create a separate FTP.DATA file or data set for the FTP server, and code SECUREIMPLICITZOS TRUE (the default) in the server's FTP.DATA. This preserves the behavior of past releases.
- For clients that share FTP.DATA with the server, change the SECUREIMPLICITZOS setting after starting the client but prior to logging into the server's protected port, as follows:
  1. Start the FTP client with the -n option to suppress logging into the server.
  2. Use the **locsite** subcommand with the **secureimplicitzos** parameter to set SECUREIMPLICITZOS to TRUE.
  3. Use the **open** subcommand to log into the FTP server's protected port. **Tip:** Use the **locstat** subcommand to display the client's current SECUREIMPLICITZOS setting.

#### **Rules:**

- To enable a z/OS FTP client to log into a z/OS FTP server's protected port, the client and server must be configured with the same value for SECUREIMPLICITZOS.
- To enable a non-z/OS FTP client to log into a z/OS FTP server's protected port, SECUREIMPLICITZOS FALSE must be coded in the server's FTP.DATA.

### **IP Services: Update network management applications for SNMP support of RFC versions of networking MIB modules (Required-IF, as of R10)**

*Required if you have network management applications that use the changed IP-MIB MIB objects, or if you have network management applications that are using the internet draft sample files for the MIB modules.*

In z/OS V1R5 and V1R6, in order to support IPv6 network management data, the SNMP TCP/IP subagent was enhanced to support IETF Internet drafts of the standard networking MIB modules, which are: IP-MIB, IP-FORWARD-MIB, TCP-MIB, UDP-MIB, and INET-ADDRESS-MIB.

The Internet drafts provided support for IPv6 network management data in the form of version-neutral tables, that is, tables that supported both IPv4 and IPv6 data.

As of z/OS V1R10, the SNMP TCP/IP subagent has been changed to support the RFC version of these same MIB modules, as follows (RFC text is at the Web addresses in parentheses):

- RFC 4293: IP-MIB (<http://www.ietf.org/rfc/rfc4293.txt>)
- RFC 4292: IP-FORWARD-MIB (<http://www.ietf.org/rfc/rfc4292.txt>)
- RFC 4022: TCP-MIB (<http://www.ietf.org/rfc/rfc4022.txt>)
- RFC 4113: UDP-MIB (<http://www.ietf.org/rfc/rfc4113.txt>)

- RFC 4001: INET-ADDRESS-MIB (<http://www.ietf.org/rfc/rfc4001.txt>)

The RFC version of the IP-MIB renames and renumbers some of the MIB objects in the earlier Internet draft, so if you have network management applications that use the changed IP-MIB MIB objects, you might have to change the applications so they continue to access the same data. Also, the Internet draft sample files for the MIB modules are no longer shipped, so if you copied these samples for use by a network management application, you should replace them with the RFC version of the MIB modules.

**IP Services: Specify at least one valid ZIIP subparameter on GLOBALCONFIG ZIIP statements (Required-IF, as of R10)**

*Required if you are currently specifying a TCP/IP configuration statement of GLOBALCONFIG with a ZIIP parameter and no associated subparameters, and you do not want your ZIIP parameter to be ignored.*

Before z/OS V1R10, specification of a GLOBALCONFIG statement with the ZIIP parameter but without a valid ZIIP subparameter would cause the NOIPSECURITY subparameter to be assumed. Beginning with z/OS V1R10, at least one valid ZIIP subparameter is required. Lack of a valid ZIIP subparameter on a GLOBALCONFIG statement causes the ZIIP parameter to be ignored, and informational message EZZ0318I is issued.



**IP Services: Migrate from QoS TR policy to IDS TR policy (Required-IF, as of R10)**

*Required if TR policies are configured as part of the QoS policy type (PolicyScope TR).*

z/OS V1R9 was the last release that supported the Traffic Regulation (TR) policy as part of the Quality of Service (QoS) policy type. The TR policy function is still available but only as part of the Intrusion Detection Services (IDS) policy type. Note that this change is only for the TR policy configuration. The TR policy functions themselves remain unaffected and continue to run.

**Migration action:**

1. Convert QoS policies that specify the PolicyScope TR parameter on the PolicyAction statement to IDS policies. All PolicyRule statements that reference any such PolicyAction statement must be converted to IDSRule statement with the ConditionType TR parameter. The PolicyAction statements must be converted to IDSAction statements with the ActionType TR parameter. The new IDS policies must be configured in an existing or new IDS configuration file.
2. If you were not using IDS policies prior to the conversion, specify the IDSConfig statement to point to the new image-specific IDS configuration file. Optionally, configure the IDS policies in both common and image-specific configuration files, and also specify the CommonIDSConfig statement.
3. Refresh the policies by issuing the MODIFY procname,UPDATE command, waiting for the refresh interval to expire, or restarting the Policy Agent. If you are using UNIX files and you previously started the Policy Agent using the -i startup option, no action is necessary; the new policies will be refreshed as soon as the configuration files are saved.

**SNA Services: Ensure compatible levels of VTAM for HPR sessions (Required-IF, as of R10)**

*Required if you plan to enable APPN and HPR support in VTAM using the NODETYPE and HPR start options.*

In order to run z/OS V1R10 Communications Server as an HPR-capable interchange node in a mixed subarea and APPN network, you must ensure that all HPR-capable VTAMs in your APPN network (and in attached APPN networks) are running z/OS V1R8 or later Communications Server, or sessions established with or through these earlier VTAMs might fail. z/OS V1R8 and later Communications Server provides additional information on APPN session establishment flows to identify when sessions cross from APPN into subarea (or vice versa) through an interchange node. This additional information is used by z/OS V1R10 Communications Server to separate interchange node sessions from APPN-only sessions by placing them on different RTP pipes. If any of the VTAMs in your network (or in attached APPN networks) are not running z/OS V1R8 or later, then z/OS V1R10 interchange nodes might incorrectly place interchange node sessions onto the wrong RTP pipe, which could result in session setup failures. (If any of the HPR-capable VTAMs in your APPN network or in attached APPN networks are running an earlier release of z/OS Communications Server, contact your IBM representative to find out what alternatives are available.)

You should also be aware that this new function might result in interchange nodes creating more RTP pipes to adjacent APPN nodes than prior releases because separate RTP pipes are now used for interchange node sessions versus APPN-only sessions.

As a result of separating interchange node sessions from APPN-only sessions by placing them onto separate RTP pipes, there may be an increase in the number of RTP pipes created by interchange nodes to adjacent APPN nodes and a corresponding increase in the amount of storage used to represent these RTP pipes.



**Migration action:** Ensure that all HPR-capable VTAMs in your network (and any APPN-attached networks) are running z/OS V1R8 or later Communications Server. All HPR-capable VTAMs must be upgraded to this level before IPLing the first z/OS V1R10 or later Communications Server interchange node. Alternatively, you can configure your HPR-capable interchange nodes to avoid the incompatibility with earlier releases of VTAM by specifying the HPRSESLM=DISABLED start option value when the VTAM interchange node is started. The PRSESLM=DISABLED start option value is provided by APAR OA28332 for z/OSV1R10 and APAR OA28727 for z/OSV1R11.

**SNA Services: Update applications and user exits that use the VTAM version and release level in algebraic expressions (Required-IF, as of R10)**

*Required if any user-written applications or exits use the VTAM release level in algebraic comparisons to determine release support.*

The VTAM version and release level is available to applications and user exits to allow the applications and exits to react based on the level of VTAM that is active. Before z/OS V1R10, applications and exits could use the version and release fields algebraically to determine the level of support. Beginning with z/OS V1R10, the fields cannot be used algebraically.

For example, after a VTAM application program issues an OPEN macro instruction with an access method control block (ACB) to indicate it desires VTAM services, the VTAM version and release level is available to it in the access-method-support vector list. The application cannot use the two bytes containing the version and release in algebraic expressions, such as checking whether the two bytes are greater than the characters 1 and 9. The reason is that, with z/OS V1R10, the version and release characters are 1 and A, respectively, and the character A cannot be used reliably in algebraic comparisons.

**Communications Server Migration Actions Pre-First IPL**

**IP Services: Allow usage of the full range of ephemeral ports (Required-IF, as of R11, and back to R9 with PK66387)**

*Required if you have security definitions in your network that restrict ephemeral ports to various ranges.*

Before z/OS V1R11, TCP/IP used the lowest available ephemeral ports when establishing new connections.

Starting in z/OS V1R11, or in V1R10 or V1R9 with APAR PK66387, TCP/IP uses the full range of available ephemeral port values (1025 -65535) beginning with the lowest port before an available port is reused. You must take action if you have security definitions in your network that restrict ephemeral ports to various ranges.

**Migration action:** Do one of the following:

- Change the definitions in your network to allow the full range of ephemeral ports. The range of allowed values is 1025 -65535.
- Modify the TCPIP profile to restrict the ephemeral ports used so that the values of the ephemeral ports match the values that are allowed in the network. For example, to exclude the use of ports 6666-6669 for both TCP and UDP applications, you can add the following PORTRANGE statements to the TCPIP profile:

```
PORTRANGE 6666 4 TCP RESERVED
PORTRANGE 6666 4 UDP RESERVED
```



**IP Services: Run the syslogd daemon in the background when starting from the UNIX shell**

**(Required-IF, as of R11)**

*Required if syslogd is started from the UNIX shell.*

Starting in z/OS V1R11, the syslog daemon (syslogd) does not automatically run in the background when it is started from the UNIX shell. As a result, you need to include a trailing ampersand character (&) on the syslogd command line to run it in the background. If you do not include a trailing ampersand, control will not return to your shell session until syslogd ends. If syslogd is started from a shell script, in particular /etc/rc, this can cause the shell to hang, waiting for syslogd to end. For example, if you run syslogd from /etc/rc and you do not change syslogd to run in the background, OMVS will not complete initialization until OMVS times out.

**Migration action:** Determine how syslogd is started:

- If it is started from /etc/rc or another shell script, or from BPXBATCH, add a trailing ampersand character to the command line.
- If you start syslogd manually from the UNIX shell, include a trailing ampersand on the command line.
- If you start syslogd from a cataloged procedure that uses BPXBATCH, add a trailing ampersand character to the command line, and include a sleep command in your script after starting syslogd to give syslogd time to initialize before the shell script ends.
- If you start syslogd from a cataloged procedure that does not use BPXBATCH, no changes are needed.

**IP Services: Accept the new resolver behavior of saving the results of DSN queries (Recommended, as of R11)**

*Not required, but recommended if you are using the resolver to query DNS servers because the new behavior can eliminate some setup, monitoring, and administration steps.*

Before z/OS V1R11, the z/OS system resolver did not cache results of domain name server (DNS) queries. Starting in z/OS V1R11, the resolver automatically saves the results of such queries. This is a performance and usability enhancement and can eliminate some setup, monitoring, and administration steps. However, if you want to keep the behavior of past releases and disable the resolver DNS caching function, you must take action.

**Migration action:** Do one of the following:

- If you want to accept the new behavior and cache the results of previous DNS queries, no action is necessary. Optionally, you can code the CACHE resolver setup statement.
- If you have a caching-only name server defined locally, you can choose one of the following steps:
  - If you want to accept the new behavior and allow the resolver to cache results of previous DNS queries, you can remove the caching-only name server from your system, saving operation and configuration costs. Alternatively, you may choose to use resolver system-wide caching to maintain the DNS A, AAAA, and PTR record information to obtain performance improvements for those searches, but continue to use the caching-only name server for the other DNS record types.
 

**Guideline:** If the local caching-only name server is the only name server in the NSINTERADDR list of name servers to be contacted, replace the caching-only name server entry with one or more name server IP addresses to be contacted. If there is already more than one name server in the NSINTERADDR list of name servers, simply delete the IP address of the local caching-only name server.
  - If you want to keep the existing behavior and want only the caching-only name server to cache results of previous DNS queries, you must code the NOCACHE resolver setup statement. This option might be attractive to you if a significantly large percentage of the information cached by the caching-only name server involves DNS records other than the DNS A, AAAA, and PTR records that the resolver will cache.
- If you want to selectively limit the caching function, you must code the TCPIP.DATA NOCACHE statement in the TCPIP.DATA files used by applications that are not permitted to utilize the caching function.

**IP Services: Update network management applications for SNMP interface data changes (Req-IF, as of R10)**

*Required if the changes affect network management applications that retrieve the affected SNMP data..*

Starting in z/OS V1R10, the SNMP TCP/IP subagent interface data support is changed as follows:

- Changes to the standard interface tables, ifTable and ifXTable, which are defined in the IF-MIB from RFC 2233 (<http://www.ietf.org/rfc/rfc2233.txt>), are as follows:
  - The tables will support a dynamically-generated OSA-Express QDIO port entry for IPAQENET and IPAQENET6 interfaces defined only by the INTERFACE profile statement. The interface counters for the OSA-Express QDIO port entry will reflect the sum of the counters for all the IPAQENET and IPAQENET6 interfaces defined for the same port.
  - When defining interfaces with DEVICE and LINK profile statements, a device layer entry is created in the tables, with the link entry as a lower interface layer under it. When defining interfaces with INTERFACE profile statements (other than IPAQENET or IPAQENET6), only an interface entry is created in the interface tables. There is no higher-layer device entry associated with the interface.
- The standard linkUp and linkDown notifications from the IF-MIB will no longer be created for VIPA interfaces.
- The proprietary ibmTcipMvsDeviceTable (defined in the IBM TCP/IP MVS Enterprise-specific MIB module) will only support entries for IPv4 interfaces either defined with a DEVICE profile statement or dynamically defined due to dynamic XCF or dynamic VIPA configuration.
- The indexes of the proprietary osaexpEthPortTable and osaexpEthSnaTable (defined in the IBM TCP/IP MVS Enterprise-specific MIB module) are changed from the interface index of an interface defined with the LINK profile statement to the interface index of a device or OSA-Express QDIO port.
- The index of the standard dot3StatsTable, which is defined in the EtherLike-MIB from RFC 2665 (<http://www.ietf.org/rfc/rfc2665.txt>), is changed from the interface index of an interface defined with the LINK profile statement to the interface index of a device or OSA-Express QDIO port.

These changes might affect network management applications that retrieve the above SNMP data.

**IP Services: Adjust for a new client error code for the FTP client (Required-IF, as of R10)**

*Required if your FTP client application is expecting a specific client error code when it attempts a client subcommand that needs a connection to the server while the client has no connection to the server..*  
Starting in z/OS V1R10, the FTP client returns a new client error code, FTP\_NEEDS\_CONNECTION, when an interactive user or FTP client API program attempts a subcommand that needs a connection to the server. This client error code replaces more general client error codes returned by earlier releases.

**IP Services: Update automation that handles FTP job informational messages for MVS data set transfers (Required-IF, as of R10)**

*Required if your installation uses automation or software to capture FTP client messages or FTP server replies, and you use FTP to transfer MVS data sets.*

Before z/OS V1R10, when a file transfer failed because the MVS data set was in use, the FTP server reply and the FTP client message indicated that the data set was in use by another process but did not identify the job or jobs holding the MVS data set. The FTP server would send a 450 reply and the FTP client would send messages EZA2562W and EZA2799W. Beginning with z/OS V1R10, in the above situation:

- The following additional FTP client messages are issued to identify the job or jobs holding the MVS data set when the file transfer fails:
  - A sequence of EZZ9819I messages identifying the jobs that are holding the MVS data set
  - Message EZZ9820I to end the sequence of messages.
- The following additional FTP server replies are issued:
  - A sequence of 125- replies identifying the jobs that are holding the MVS data set
  - A single 125- reply to end the sequence of replies for this sequence.

**IP Services: Configure GLOBALCONFIG SEGMENTATIONOFFLOAD if TCP segmentation offload support is desired (Required-IF, as of R9 with PK47376)**

*Required if TCP segmentation offload is desired.*

Before z/OS V1R10 (and APAR PK47376 on z/OS V1R9, and APAR PK46334 on z/OS V1R8, V1R7, and V1R6), TCP segmentation was offloaded to the OSA-Express2 feature by default. Beginning with z/OS V1R10 (and APAR PK47376 on z/OS V1R9, and APAR PK46334 on z/OS V1R8, V1R7, and V1R6), the default behavior is to not offload TCP segmentation to the OSA-Express2 feature. If you want to continue to offload TCP segmentation to the OSA-Express2 feature, you must configure SEGMENTATIONOFFLOAD on the GLOBALCONFIG statement.

**Migration action:** To obtain the prior behavior, configure SEGMENTATIONOFFLOAD on the GLOBALCONFIG statement in the TCP/IP profile.

Review Washington Systems Center flash 10458, titled *OSA-Express2 Segmentation Offload*, at <http://www.ibm.com/support/techdocs>. (Search for "flash10458".)

**IP Services: Make changes for Netstat enhancements (Required-IF, as of R10 and R11)**

*Required if the changed or removed settings affect automation running off Netstat or front-end programs to Netstat.*

The Netstat command displays the status of a local host. Each release, the Netstat reports are changed in ways that can affect automation or front-end programs.

**Migration action:** Accommodate Netstat changes in your automation and front-end programs. You can begin planning your changes by reviewing the ways in which the displays are updated each release. However, you will have to execute the commands to know with certainty what changes to make.

For details about Netstat report changes, see *z/OS Summary of Message and Interface Changes*



**IP Services: Update /etc configuration files (Required-IF)**

*Required if you have customized a configuration file that IBM has changed.*

Some utilities provided by Communications Server require the use of certain configuration files. You are responsible for providing these files if you expect to use the utilities. IBM provides default configuration files as samples in the /usr/lpp/tcpip/samples directory. Before the first use of any of these utilities, you should copy these IBM-provided samples to the /etc directory (in most cases). You can further customize these files to include installation-dependent information. An example is setting up the /etc/osnmpd.data file by copying the sample file from /usr/lpp/tcpip/samples/osnmpd.data to /etc/osnmpd.data and then customizing it for the installation.

If you customized any of the configuration files that have changed, then you must incorporate the customization into the new versions of the configuration files.

**Migration action:** If you added installation-dependent customization to any of the IBM-provided configuration files listed below, make the same changes in the new versions of the files by copying the IBM-provided samples to the files shown in the table and then customizing the files.

Utility	IBM-provided sample file	Target location	What changed and when
BIND	/usr/lpp/tcpip/samples/IBM/EZANSBOO	/etc/named.boot	In V1R11, BIND DNS 4.9.3 function was removed from z/OS.
BIND	/usr/lpp/tcpip/samples/IBM/EZANSCH4	/etc/caching.boot	In V1R11, BIND DNS 4.9.3 function was removed from z/OS.
BIND	/usr/lpp/tcpip/samples/IBM/EZANSSL4	/etc/slave.boot	In V1R11, BIND DNS 4.9.3 function was removed from z/OS.
BIND	/usr/lpp/tcpip/samples/IBM/EZATDDSD	/etc/dhcpsd.cfg	In V1R11, BIND DNS 4.9.3 function was removed from z/OS.
BIND	/usr/lpp/tcpip/samples/IBM/EZATDLSS	/etc/binlsd.cfg	In V1R11, BIND DNS 4.9.3 function was removed from z/OS.

Utility	IBM-provided sample file	Target location	What changed and when
Network security services (NSS)	/usr/lpp/tcpip/samples/nssd.conf	Update as needed by your installation and put in user directory.	In z/OS V1R10, the NSS server introduced support for discipline-specific NSS clients. The two disciplines supported in V1R10 are IPSec and XMLAppliance. To account for the configurable support of each discipline, the NSS server's sample configuration file was updated to indicate that support for both disciplines is enabled by default. A new syslog level for the SAF access service within the XMLAppliance discipline was also introduced.
Policy Agent	/usr/lpp/tcpip/samples/pagent.conf	/etc/pagent.conf	In z/OS V1R11, new configuration statements are provided to support automatic starting, stopping, and monitoring for a selected set of applications.
SNMP agent	/usr/lpp/tcpip/samples/osnmpd.data	/etc/osnmpd.data	Every release, the value of the sysName MIB object is updated to the current release.  In z/OS V1R10, the SNMP agent was changed to create its default z/OS UNIX connection file in the /var file system directory instead of the /tmp directory. So the value of the dpiPathNameForUnixStream MIB object in this sample was changed from /tmp/dpi_socket to /var/dpi_socket. If you use the updated sample, the z/OS UNIX file will be created in the /var file system directory instead of in the /tmp directory. Because the Agent will no longer use the /tmp/dpi_socket file, this file can be removed.
syslog daemon	/usr/lpp/tcpip/samples/syslog.conf	/etc/syslog.conf	In z/OS V1R11, new configuration statements are provided to support automatic archiving of z/OS UNIX files.
z/OS UNIX Telnet server (otelnetd)	None	/etc/otelnetd.banner	In z/OS V1R11, this file was added to display a banner page, prior to the login prompt, when a user connects to the z/OS UNIX Telnet server (otelnetd).

**IP Services: Update automation that handles message EZZ4201I (Required-IF, as PK46221 on R9/R8/R7)**

*Required if you depend on the message EZZ4201I issued in response to the TCP/IP STOP command.*

When TCP/IP is started, the console ID used on the START command is saved. This console ID is used when TCP/IP sends messages to the console. When TCP/IP is stopped, message EZZ4201I TCP/IP TERMINATION COMPLETE FOR jobname uses the console ID from the START command rather than the STOP command, causing EZZ4201I to not be trapped. As of APAR PK46221 on z/OS V1R9, z/OS V1R8, and z/OS V1R7, STOP command processing has been changed to save the WTO CART (command and response token) and console identifier value. As a result, message EZZ4201I issued in response to the STOP command is sent to the correct console, that is, the one that issued the STOP command.

**IP and SNA Services: Increase region size for TELNET (Required-IF, as of R10)**

*Required if you use the TN3270E Telnet server.*

Starting in z/OS V1R10, extended common system area (ECSA) usage for Telnet sessions is reduced. This change requires growth in Telnet private storage usage. The amount of Telnet private storage increase depends upon the number of Telnet sessions on the LPAR. If Telnet is using most of its allotted region size, you must increase it.

**Migration action:**

1. Determine the number of Telnet sessions that will exist on the LPAR.
2. Increase the size specified on the REGION parameter of the Telnet started task by 1 MB for each 25 000 Telnet sessions.
3. Examine the values for your CRPLBUF pool defined in the ATCSTRxx VTAM start list, particularly the base size. If the base size is large because of a large number of Telnet sessions, you can reduce it; failure to reduce it can minimize the amount of ECSA savings realized.

**IP and SNA Services: Disable path MTU discovery for IPv4 and IPv6 Enterprise Extender connections (Required-IF, as of R10)**

*Required if you do not want path MTU discovery enable for IPv4 and IPv6 Enterprise Extender connections.*

Starting in z/OS V1R10, path MTU discovery is enabled for Enterprise Extender (EE) connections. It is automatically enabled for IPv6 Enterprise Extender connections, and is enabled for IPv4 when the PATHMTUDISCOVERY keyword is specified on the IPCONFIG profile statement for the EE associated TCP/IP stack.

Path MTU discovery uses ICMP fragmentation-needed errors to detect the path MTU for a path. For path MTU discovery to work, you must permit ICMP errors to flow at all hosts along the path of a connection. Path MTU discovery does not function if a firewall blocks ICMP errors.

**Migration action:** To disable path MTU discovery for IPv4 and IPv6 Enterprise Extender connections, perform one of the following actions:

- Specify PMTUD=NO in the appropriate ATCSTRxx VTAM start list or on the VTAM START command.
- When VTAM is active, issue the MODIFY *procname*,VTAMOPTS, PMTUD=NO command.

**SNA Services: Ensure that dump data sets are large enough to include VTAM internal trace data space (Required, as of R11)**

Before z/OS V1R11, VTAM INOP dump processing did not capture the VTAM internal trace (VIT) data space in the dump. Starting in z/OS V1R11, VTAM INOP dump processing automatically captures the VIT data space in the dump when the VIT data space is in use. This provides more information for problem determination in INOP dumps. In order to continue to capture all INOP dump information, you must ensure that your dump data set allocation will accommodate the additional space required for the VIT data space. The maximum possible size of the VIT data space, 50 MB, should be taken into consideration in your dump data set allocation procedure.

**Migration action:**

1. Be aware that if the VIT data space is active, VTAM INOP dumps will automatically capture the VIT data space (ISTITDS1) to provide more information in the dump for problem determination. This can result in larger INOP dumps.
2. Review the size of your current dump data sets and ensure that they are large enough to accommodate the maximum size of the VIT data space.
  - Because VIT data space tracing is not new in z/OS V1R11, your dump data sets might already be large enough for INOP dumps. VTAM SVC dumps have always dumped the VIT data space. These include dumps taken by ABEND recovery routines (FRRs and ESTAE routines) and any dumps taken with the F *vtamproc*,CSDUMP command. If your system can currently capture a complete dump of VTAM, no changes are required.

- If current VTAM dumps are incomplete, modify your dump data set allocation to accommodate the maximum size of the VIT data space.
3. VIT data space tracing is important to capture sufficient trace data for VTAM problem determination. With VIT data space tracing enabled at all times, you might be able to resolve VTAM problems on the first failure, preventing the need for you to recreate the problem, or wait for the problem to occur again, to get additional documentation. To ensure that you always have the maximum amount of VIT trace data available for problem diagnosis, do one of the following:
    - Activate the VIT data space at VTAM start by coding DSPSIZE=5 on any TRACE,TYPE=VTAM,MODE=INT start option included in the VTAM start list or on the VTAM start command.
    - If the VIT data space is not activated at VTAM start, activate it with the MODIFY *vtamproc*,TRACE,TYPE=VTAM,DSPSIZE=5 command.

**Tip:** Use IBM Health Checker for z/OS check CSVTAM\_VIT\_DSPSIZE to determine whether the VIT dataspace table size is set to 5 (50 MB).

### **SNA Services: Increase region size for VTAM (Required-IF, as of R11)**

*Required if you have large numbers of HPR connections and VTAM is using a high percentage of its allocated region size.*

ECSA usage for HPR has been reduced. To accomplish this reduction, some growth in VTAM private storage usage was required. The amount of VTAM private storage increase will depend upon the number of RTP endpoints that are supported on the LPAR. If VTAM is using nearly all of its allotted region size, you will need to increase it.

#### **Migration action:**

1. Determine the maximum number of RTP pipes that have endpoints in this VTAM.
2. Increase the region size for VTAM by 1.5 MB for each 1,000 RTP pipes.

## **Communications Server Migration Actions Post-First IPL**

### **IP Services: Ensure that IKE daemon retransmission settings are compatible (Recommended, as of R11)**

*Not require, but recommended if you have specified a KeyWait, KeyRetries, DataWait, or DataRetries parameter in your IKED configuration file to control the IKE daemon's retransmission behavior.*

Before z/OS V1R11, the following four Internet Key Exchange daemon (IKED) parameters were used to control the IKE daemon's retransmission behavior: KeyWait, KeyRetries, DataWait, DataRetries.

If you used the IBM Configuration Assistant, these parameters' values were set by selecting retry limit and retry interval values for the Phase 1 IKE key negotiation retry tuning or Phase 2 IKE data negotiation retry tuning settings on the Advanced IKE Daemon Settings dialog under the IPSec perspective. The parameters established a fixed retransmission interval and maximum number of retransmissions for a given IKE message.

Starting in z/OS V1R11, the old parameters are ignored and the following two new parameters are used to determine all IKE message retransmission: IkelnitWait and IkeRetries.

Like the old parameters, these new parameters are also set through the Advanced IKE Daemon Settings dialog of the IBM Configuration Assistant. The new parameters enforce a geometrically increasing retransmission interval (each successive interval is twice as long as the one that preceded it). The length of the first interval is specified by IkelnitWait and the number of transmissions is specified by IkeRetries.

You cannot keep the previously-existing behavior of controlling the IKED retransmissions. In most cases, the new default behavior will work well for all of your IKE peers. However, there may be some circumstances where closer control over the retransmission behavior is required (often because of low bandwidth, or slow or unreliable network connections). In these cases, you might need to specify one or both of the new parameters to tailor the behavior. In most cases, the change will be to reduce the allowed number of retransmissions to avoid very long retransmission timeouts.

#### **Migration action:**

1. If you generated the IKED configuration file using the IBM Configuration Assistant, check the Advanced IKE Daemon Settings dialog under the IKE Daemon Settings tab under the IPSec perspective for the z/OS image. If any IKE negotiation retry tuning settings are specified, you might have some site - or network-specific retransmission considerations to accommodate. If you configured your IKED configuration file manually, check it to see if any of KeyWait, KeyRetries, DataWait, or DataRetries were specified to tailor the IKE daemon's retransmission behavior. If they were, you might have some site- or network-specific retransmission considerations that must be accommodated.

2. Set the IKE retry tuning values in the IBM Configuration Assistant Advanced IKE Daemon Settings dialog or, if you manually edit your IKED configuration file, specify the IkelnitWait and/or IkeRetries parameters to accommodate the site- or network-specific considerations.

**IP Services: Use ETHERNET instead of TOKENRING with IPCS CTRACE (Recommended, as of R10)**

*Not required, but recommended if you use the TOKENRING keyword of the IPCS CTRACE subcommand because someone could be confused by ETHERNET taking effect when they see that TOKENRING was specified.*

Before z/OS V1R10, the TOKENRING keyword of the SNIFFER option of the IPCS CTRACE subcommand would format packet trace output for token-ring analysis. Because of the phaseout of token-ring technology, beginning with z/OS V1R10, the TOKENRING keyword is no longer supported.

**Migration action:** Use the ETHERNET keyword instead of the TOKENRING keyword when formatting the SNIFFER data set trace. ETHERNET formats output for Ethernet analysis and is the default. If you specify TOKENRING, the keyword is ignored and ETHERNET is in effect.

**IP Services: Modify FTPOSTR exit routines to handle a new parameter (Required-IF, as of R10)**

*Required if you have installed an FTPOSTR exit routine and it expects 18 or fewer parameters.*

The FTPOSTPR exit routine is called upon completion of the FTP commands RETR, STOR, STOU, APPE, DELE, and RNTD. In z/OS V1R10, the exit routine has a new parameter: a pointer to a buffer containing the complete text of the server reply sent to the client. If you have an FTPOSTPR exit routine that cannot tolerate this new parameter, you must modify the routine.

**Migration action:** Modify your FTPOSTPR exit routine to tolerate 19 parameters.

**SNA Services: Create a new APPN topology checkpoint data set (Required-IF, as of R11)**

*Required if you checkpoint APPN topology and restore it with the INITDB=TOPO or INITDB=ALL start option.*


In z/OS V1R11 the format of the records in the APPN topology checkpoint data set has changed. A topology checkpoint data set created on a prior release cannot be read by a z/OS V1R11 network node, and a topology checkpoint data set created in z/OS V1R11 cannot be read by a prior release. Therefore, after migrating an APPN network node to z/OS V1R11, you must create a new APPN topology checkpoint data set. When CP-CP sessions are started for the first time with other APPN network nodes, performance might be impacted temporarily because a full topology exchange is required.

**Migration action:** To create a new APPN topology checkpoint data set after migrating an APPN network node to z/OS V1R11, do the following:





1. Start all desired CP-CP sessions between network nodes in the APPN network and activate desired transmission groups (TGs). Because you cannot retrieve existing APPN topology from a topology checkpoint data set created with a prior release of VTAM when a z/OS V1R11 VTAM is started, a full topology exchange will occur when the z/OS V1R11 network node joins an APPN network for the first time. This happens when the z/OS V1R11 network node activates CP-CP sessions with any adjacent network node in the APPN network. This initial exchange, made necessary by the difference in old and new data set formats, can degrade performance until all of the network topology is transmitted to the z/OS V1R11 network node. After the z/OS V1R11 network node has received all of the network topology from the adjacent network node, there is no further negative performance impact caused by the difference in format of the new topology checkpoint data set, as long as you create a topology checkpoint data set with the new format on the z/OS V1R11 network node after the initial topology exchange. Otherwise, a full topology exchange will be required every time the z/OS V1R11 network node is recycled.
2. When your APPN network is stable, enter the following command to create a topology checkpoint data set in the new format: `MODIFY procname,CHKPT,TYPE=TOPO`. A checkpoint of APPN topology is also done automatically when you end VTAM with a `HALT` or `HALT QUICK` command but not when you end VTAM with a `HALT CANCEL` command.
3. After the topology checkpoint data set is initially created, you can restore it with `INITDB=TOPO` or `INITDB=ALL` when you restart z/OS V1R11 VTAM.

If you start VTAM as a network node with the `INITDB=TOPO` or `INITDB=ALL` start option with a version of the APPN topology checkpoint data set that is not recognized, you will receive message `IST1288I TOPOLOGY DATASET RETRIEVAL WAS NOT SUCCESSFUL, CODE = 12`.





### Various Migration Actions for z/OS R11 from z/OS R9

- 
 • **ICSF HCR7750: Increase the size of your PKDS (Required, as of HCR7750. z/OS R9 has HCR7740 incorporated.)**
  - As of HCR7750, you must allocate your PKDS with a larger record size. Failure to do this before starting ICSF will result in initialization errors.
- 
 • **ICSF HCR7770: Modify ICSF startup procedure (Required, as of HCR7770. not in z/OS R11)**
  - Switch from PGM=CSFMMAIN to CSFINIT.
- 
 • **HCM: Install the new HCM (Required, as of R11)**
  - As of R11, the installation of HCM is new – it is an MSI package and is installed via Windows Installer.
  - Uninstall previous HCM. Install R11 HCM. An upgrade of your older level is not possible.
- 
 • **IBM TDS: Determine if you will be using IBM TDS WLM support (Required, as of R11)**
  - As of R11, IBM TDS is enhanced to use WLM performance goals for work within the LDAP server (by default). It is based on client IP addr or the bound user's distinguished name assoc. with LDAP transaction names.
    - Ensure the `srvStarUpError` option is terminate (default).
    - Permit the LDAP server userid to `BPX.WLMSEVER` facility class.
    - Use WLM ISPF panels for classifications rules. *If a default service class for the LDAP subsystem is not configured, it will be discretionary!*
  - You may turn off the WLM support if you like.

28 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

### Various Migration Actions Between z/OS V1 R9 and z/OS V1 R11

Some migration actions from specific elements were selected from *z/OS Migration*. Many descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.



#### **ICSF: Increase the size of your PKDS (Required, as of HCR7750 and integrated into R10)**

With ICSF FMID HCR7750, the record size of the PKA key data set (PKDS) has increased. You must allocate a new PKDS with a larger record size so that internal I/O routines can handle 4096-bit tokens. Failure to migrate the PKDS before starting ICSF results in error messages CFC0286, CSFM406A, and CSFM407A (and ICSF not initializing.)

If the PKDS is shared with earlier systems, install the PTF for APAR OA21807 on those systems to allow continued sharing of the PKDS. ICSF will be able to interact with both sizes of the PKDS.

**Migration action:** Before initializing ICSF, do the following:

1. Make a backup copy of the PKDS you are currently using.
2. Create a larger PKDS. Use the JCL in SYS1.SAMPLIB(CSFPKDS). If the PKDS will be shared, place the VSAM data set where it can be shared.
3. Copy the old PKDS to the larger PKDS using the JCL in SYS1.SAMPLIB(CSFPKDCP). Protect the VSAM data sets from viewing by unauthorized personnel.
4. Initialize the larger PKDS and update the ICSF started procedures on all systems to reference the new PKDS.
5. Stop and restart ICSF on each system.

**Tip:** Use IBM Health Checker for z/OS to verify that the PKDS size in an ICSF pre-HCR7750 environment is sufficiently allocated to support 4096-bit RSA keys. The check is named ICSFMIG7731\_ICSF\_PKDS\_TO\_4096BIT.



### **ICSF: Modify ICSF startup procedure (Required, as of HCR7770 web deliverable – not integrated into R11)**

The program that started ICSF in earlier releases was named CSFMMAIN. In ICSF FMID HCR7770 (which was made available in the Web deliverable *Cryptographic Support for z/OS V1R9-R11*), the CSFMMAIN program is replaced by the CSFINIT program. If your ICSF startup procedure is not modified to run this new program, the procedure will not start the HCR7770 level of ICSF.

**Migration action:** In your startup procedure for ICSF:

1. Find the job step that identifies the ICSF startup program (CSFMMAIN) that was used in earlier releases. For example: CSFSTART EXEC PGM=CSFMMAIN,REGION=0M,TIME=1440
2. Modify the PGM parameter on this EXEC statement to identify the new startup program (CSFINIT): CSFSTART EXEC PGM=CSFINIT,REGION=0M,TIME=1440
3. Save your changes to the startup procedure.

Tip: Member CSF in SYS1.SAMPLIB contains a sample JCL code for an ICSF startup procedure.

#### **More details about this migration action (from WSC Flash FLASH1062):**

ICSF is being made non-swappable and non-cancelable using standard z/OS conventions. Prior versions of ICSF are non-swappable however this new support is implemented using standard conventions. Current versions of ICSF are cancelable via an operator command. If the operator issues the cancel command multiple times, the address space may be terminated before cleanup can be completed and this can sometimes cause problems to other ICSF started tasks that are sharing the key data sets (CKDS, PKDS, TKDS). By making ICSF non-cancelable, operations will have to issue an explicit command to stop the address space, driving it through normal termination routines.

Making the ICSF started task non-cancelable and non-swappable requires that the ICSF initialization module be added to the Program Properties Table (PPT). To avoid impacting prior versions of ICSF this main module name is being changed from CSFMMAIN (PGM=CSFMMAIN for HCR7751 and earlier) to CSFINIT. Simply adding CSFMMAIN to the PPT, would cause current versions of ICSF to work differently than they do today (i.e. cancel operations would fail).

The CSFINIT entry in the PPT will be automatically generated as part of the install of HCR7770, however existing users of ICSF will need to modify the ICSF started task to use the new initialization module. CSFMMAIN will still exist with HCR7770, however it will simply generate error message CSFM022E explaining that CSFINIT is required and then ICSF will terminate.

Since this version of ICSF can be installed on older versions of z/OS, the migration from CSFMMAIN to CSFINIT will need to be carefully coordinated in shops running multiple versions of z/OS and ICSF. Systems that require earlier versions of ICSF will have to continue to use //CSF EXEC PGM=CSFMMAIN,REGION=0M,TIME=1440 in the started task. However, to use HCR7770, the started task will need to be modified to specify //CSF EXEC PGM=CSFINIT,REGION=0M,TIME=1440

Installations that are running multiple LPARs with various versions of ICSF across those LPARs, might want to consider using a system variable for the startup program name. For example, the ICSF started task could be modified to specify //CSF EXEC PGM=&ICSFPROG,REGION=0M,TIME=1440 and the system symbolic ICSFPROG could be set to SYMDEF(&ICSFPROG='CSFINIT') on the HCR7770 systems and SYMDEF(&ICSFPROG='CSFMMAIN') on the earlier systems. No matter whether you change the started task directly to point to the new startup module, or use system symbols, the change will have to be made before HCR7770 can be started. If using system symbols, the system symbol will have to be refreshed before ICSF can be restarted.

Alternatively, the ICSF started task could be moved to a release specific PROCLIB. On the HCR7770 systems the version of ICSF that uses

```
//CSF EXEC PGM=CSFMINIT,REGION=0M,TIME=1440
```

would come before the library that contains the JCL with

```
//CSF EXEC PGM=CSFMMAIN,REGION=0M,TIME=1440.
```

On the pre-HCR7770 systems, the concatenation order would be reversed.



### **HCM: Install the new HCM (Required, as of R11)**

In z/OS V1R11, the installation process for HCM is new. The code is delivered as an MSI package and is installed using the Windows® Installer. If you have an older version of HCM installed, it is strongly recommended that you uninstall it before you install the new HCM, which uses the Windows Installer. An upgrade of your older level is not possible.

**Migration action:** Uninstall your previous HCM using the Windows uninstall facility or use the **Uninstall** function provided by HCM from the following path: All Programs-->IBM Hardware Configuration Manager-->Uninstall. Do this before installing z/OS V1R11 HCM.

**Note:** With z/OS V1R11, the location of the EEQHCM.INI file changes from C:\WINDOWS to the user-specific application data directory. This does not affect your migration because the first invocation of HCM automatically copies the old INI file to the new location or creates a new one if no INI file is found.

The default name of the user-specific directory, for example, for Windows XP, is C:\Documents and Settings\*user*\Application Data\IBM\Hardware Configuration Manager. Due to changes in the user authorization methods for Windows XP and later versions, the change of the location is necessary to allow users who are not administrators to update this file.



### **IBM TDS: Determine if you will be using IBM TDS WLM support (Required, as of R11)**

Before z/OS V1R11, the IBM TDS did not support using Workload Manager (WLM) to allow an installation to set performance goals for work within the LDAP server. Beginning with z/OS V1R11, the IBM TDS is enhanced to use WLM performance goals for work within the LDAP server, based on the client IP address or the bound user's distinguished name (DN) associated with LDAP requests. If a client application is spamming the LDAP server with numerous LDAP requests, a lower priority can be assigned to those requests. This support can prevent the LDAP server's resources from being consumed by spamming LDAP client applications.


**Migration action:** If you want to use the IBM TDS WLM support, which is the default:



1. Set the `srvStartUpError` option to terminate in the LDAP server configuration file. This is the default setting for this configuration option.
2. Issue the following RACF commands before starting IBM TDS:  
RDEFINE FACILITY BPX.WLMSEVER UACC(NONE)  
PERMIT BPX.WLMSEVER CLASS(FACILITY) ID (*userid*) ACCESS(READ)  
SETROPTS RACLIST(FACILITY) REFRESH  
where *userid* is the user ID of the LDAP server.
3. In the WLM ISPF panels, define a new classification rule using the subsystem type of LDAP. If you had a subsystem type of LDAP defined previously (for the Integrated Security Services LDAP Server), it will now be used for the IBM TDS. Review any classification rules that you had for the Integrated Security Services LDAP Server, and ensure that they are desirable for the IBM TDS.

Define classification rules for the LDAP subsystem type, using a transaction name of GENERAL (for LDAP work), and any other user-defined WLM transaction names used for the IBM TDS. Refer to *z/OS MVS Planning: Workload Management* for all the workload qualifiers that the LDAP subsystem type supports. **If a default service class for the LDAP subsystem is not configured in WLM, all LDAP server operations will run under the discretionary goal and receive a low priority.**

4. If you do not want to use the IBM TDS WLM support, set the `srvStartUpError` option to ignore in the LDAP server configuration file.


## DFSMS Migration Actions for z/OS R11 from z/OS R9




- **Migration Actions You Can Do NOW:**
  -  **Redefine existing VSAM data sets that contain the IMBED, REPLICATE, and KEYRANGE attributes (Recommended)**
    - IMBED, REPLICATE, and KEYRANGE support is **no longer** planned to be removed in the foreseeable future. However, it is still strongly discouraged to use these attributes.
    - Determine which VSAM data sets are affected, then schedule a time to redefine them.
    - Use tools: IMBDSHIP, and the new z/OS R11 Health Check for catalogs!
    - **Planned for z/OS R12!** These attributes will be removed automatically from VSAM data sets logically dumped using DFSMSdss and migrated using DFSMSshm when DFSMSdss is the data mover during restore and recall processing.
  -  **Migrate from VRSEL(OLD) to VRSEL(NEW) in parmlib member EDGRMMxx (Required-IF, as of R11)**
    - In R11, VRSEL(OLD) has been removed, default has been changed to NEW.
    - In R8-R10, you'd get a warning message EDG2317E and RC=4 in EDGHSKP processing.

29 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## DFSMS Migration Actions for z/OS R11 from z/OS R9



- **Migration Actions Pre-First IPL:**
  - Use **ALL, DATE, or VERSIONS** with (H)BDELETE command and ARCHDEL macro (**Required-IF, as of R11**)
      - Prior to R11: you could use (H)BDELETE *dsn* to delete all backup versions of *dsn*.
      - As of R11: you must specify the ALL keyword on this command to delete all backup versions. You must use ALL (a new keyword), DATE, or VERSIONS on the (H)BDELETE command.
  -  **Use DFSMSrmm parmlib options instead of usermod to specify how VRSEL processing handles duplicate GDGs (Required-IF, as of R11)**
    - Prior to R11: you could use a usermod to EDGVREC to influence how VRSEL handled duplicate GDGs.
    - As of R11: No more usermod! It will be ignored if you install it on z/OS R11. Use GDG operand on the OPTION command in parmlib member EDGRMMxx.

30 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## DFSMS Migration Actions for z/OS R11 from z/OS R9



### • Migration Actions Pre-First IPL:

- **Accommodate changes to DEVSERV command output (Req-IF, as of R10)**
  - DEVSERV QDASD output is changed: cylinder column expanded, new column heading ECK and is reduced to three characters.
  - DEVSERV PATHS output is changed: adds the number of cylinders in a new column called CYL, ALT column has been replaced with the CYL column.
  - Both outputs have been shifted one position to the right.
  - Change any dependency you have on this changed output.
- **Permit DEVMAN to call ICKDSF (Required-IF, as of R11)**
  - Starting with z/OS V1R11, support is added to the device manager (DEVMAN) to invoke ICKDSF for the new DEVMAN REFVTOC function. During startup of the DEVMAN address space, ICKDSF is loaded by DEVMAN regardless of whether or not you have enabled the REFVTOC function.
  - If you use program control in RACF to restrict access to ICKDSF, you must ensure the profile which protects the ICKDSF resource in the PROGRAM class allows the user ID associated with the DEVMAN address space to have READ access.
  - Failure to do this means you will not be able to exploit the new REFVTOC function. Note that DEVMAN initialization continues even if DEVMAN does not have sufficient authority.

31

z/OS 1.11: Migration - ... And It's Good! Part 2 of 2

© 2010 IBM Corporation

## DFSMS Migration Actions for z/OS R11 from z/OS R9



### • Migration Actions Pre-First IPL:

- **Accommodate changes in the CA sizes assigned to newly allocated VSAM data sets (Required-IF, as of R10)**
  - The CA size selection is done differently on z/OS V1R10 for all new VSAM data sets (not just those on EAVs).
  - The CA size is 1, 3, 5, 7, 9 or 15 tracks and is determined by the primary and secondary space allocation amounts. The index CI size may be increased when redefined on R10.
  - As before, you cannot explicitly specify a CA size. Before z/OS V1R10, other possible CA sizes existed. Adjust any primary and secondary requests to accommodate the CA size you rely upon.
  - Refer to III4458 for a tool on finding data sets that will have a larger index CI size when defining them on z/OS R10.
- **Update the storage size for jobs that invoke ADRDSSU (Required-IF, as of R11)**
  - Beginning with z/OS V1R11, jobs invoking DFSMSdss ADRDSSU might require a storage size of greater than 2M.
  - Check the JCL for jobs that invoke ADRDSSU. If the REGION parameter specifies a region of 2M or smaller, increase the region size. If the SIZE parameter on the EXEC statement specifies a size of 2M or smaller, increase the size.
  - Check your JES initialization parameters. If the REGION parameter specifies a default region size of 2M or smaller for each job step, you may want to increase this.

32

z/OS 1.11: Migration - ... And It's Good! Part 2 of 2

© 2010 IBM Corporation

## DFSMS Migration Actions for z/OS R11 from z/OS R9



### • Migration Actions Post-First IPL:

- **Update programs or procedures that depend on output in reports from the LISTDATA PINNED command (Required-IF, as of R10)**
  - Track addresses in output of IDCAMS LISTDATA PINNED are in a different format, regardless of volume size.
  - As of R10: a range is identified, and the track address is printed in 8 hex digits.
- **Accommodate the change of IDCAMS DEFINE for RECORDS involving VSAM linear data set (Required-IF, as of R11 (NOT OA21369 on R9!!!))**
  - Before R11: when space for linear data sets was calculated for the DEFINE command with the RECORDS parameter, an implied CI size of 4096 was used when the CI size was larger than 4096.
  - As of R11: in the same scenario, the actual larger CI size is used instead of 4096.
  - Consider adjusting the value of RECORDS when a linear data set is defined with CISIZE greater than 4096. Using the same IDCAMS DEFINE CLUSTER job will result in the allocated space being larger than before because the actual CISIZE (instead of 4096) has an influence on the amount of space allocated to the data set. You need to decrease the value of RECORDS to achieve an allocation amount similar to before. See OA25988.
  - Install DB2 PTF PK42840 to avoid DB2 side effects from this change.

33

z/OS 1.11: Migration - ...And It's Good! Part 2 of 2

© 2009 IBM Corporation

## DFSMS Migration Actions for z/OS R11 from z/OS R9



### • Migration Actions Post-First IPL:

- **Use the new version of the DEFRAG command in DFSMSdss (Recommended)**
  - Before R10: DEFRAG command contained a CONSOLIDATE keyword.
  - As of R10: a new DEFRAG command is introduced (for EAV usage), which drops the CONSOLIDATE keyword. Replacement for the CONSOLIDATE keyword on the new DEFRAG command is the CONSOLIDATE command. Need facility class profile access to STGADMIN.ADR.CONSolid and STGADMIN.ADR.DEFRAG to use.
  - Old DEFRAG command is still available, when specifying the VERSION1 keyword (which isn't for EAV usage).
  - Old JCL that uses CONSOLIDATE with DEFRAG (and without VERSION1) will continue to run, however CONSOLIDATE command will be executed first followed by the DEFRAG command.
- **Accommodate the change in DFSMSshm ARCBDEXT exit (Required-IF, as of R10)**
  - Before R10: ARCBDEXT was called during volume-level backup operations. Users examined input data structure to determine whether to allow backup of the data set or not.
  - As of R10 with OA28138: ARCBDEXT is called during individual data set backup operations, as well as volume-level backups. For individual data set backup, the exit's input structure differs.
  - To know whether ARCBDEXT is invoked for volume-level or individual data set level, examine offset x'04'.

34

z/OS 1.11: Migration - ...And It's Good! Part 2 of 2

© 2009 IBM Corporation

### DFSMS Migration Actions for z/OS R11 from z/OS R9

- **Migration Actions Post-First IPL:**
  - **Define facility class names to protect (H)BACKDS RETAIN DAYS commands (Required-IF, as of R11)**
    - As of R11: new hsm facility class names protect the BACKDS RETAIN DAYS and HBACKDS RETAIN DAYS commands. Only available on R11, prior levels will be able to recover backup versions created with the RETAIN DAYS keyword. New facility class names must be defined to protect from unintended use of this support by unauthorized users.
    - Define new facility class profiles to authorize to the (H)BACKDS commands:
      - STGADMIN.ARC.BACKDS.RETAIN DAYS
      - STGADMIN.ARC.ENDUSER.HBACKDS.RETAIN DAYS
  - **Define facility class names to protect (H)BACKDS NEWNAME commands (Required-IF, as of R10)**
    - As of R10: new hsm facility class names protect the BACKDS NEWNAME and HBACKDS NEWNAME commands. Also, authorization for data set name on NEWNAME is checked. New facility class names must be defined to protect from unintended use of this support by unauthorized users.
    - Define new facility class profiles to authorize to the (H)BACKDS commands:
      - STGADMIN.ARC.BACKDS.NEWNAME
      - STGADMIN.ARC.ENDUSER.HBACKDS.NEWNAME

35
z/OS 1.11: Migration - ... And It's Good! Part 2 of 2
© 2009 IBM Corporation

### DFSMS Migration Actions Between z/OS V1 R9 and z/OS V1 R11

These migration actions were taken from *z/OS Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### DFSMS Migration Actions You Can Do Now

#### DFSMSdfp: Back up SMS control data sets (Recommended)

*Recommended to ensure data integrity.*

**Migration action:** Do the following on your pre-z/OS V1R11 systems:

1. Back up SMS control data sets according to established procedures in the event that fallback is required. The control data set format is VSAM linear.
2. Install all coexistence PTFs defined for z/OS V1R11.

In addition, if you modified and activated a higher-level policy on a pre-z/OS V1R11 system, do the following to ensure that the ACDS can be accessed on z/OS V1R11:

1. On the pre-z/OS V1R11 system, save the active ACDS as an SCDS with the SETSMS SAVESCDS command.
2. On z/OS V1R11, update, translate, validate, and activate the saved SMS policy.

**Note:** With z/OS V1R11 (and prior releases by APAR OA22373), CDS uplevel checking only checks the formatted release against the updating release level. The construct size change no longer causes warning messages.



#### DFSMSdfp: Redefine existing VSAM data sets that contain the IMBED, REPLICATE, and KEYRANGE attributes (Recommended)

*Recommended to avoid degraded performance and wasted DASD space.*

No supported release of z/OS honors the IMBED, REPLICATE, and KEYRANGE attributes for new VSAM data sets. In fact, using these attributes can waste DASD space and often degrades performance. Servicing these VSAM data sets has become increasingly difficult. In some cases, unplanned outages have occurred. For these reasons, IBM recommends that you stop using IMBED and REPLICATE, and that you minimize or eliminate your use of KEYRANGE. IMBED and REPLICATE were intended as performance improvements and have been obsolete by newer, cached DASD devices. Striped data sets provide much better performance than KEYRANGE and should be viewed as a candidate for any existing KEYRANGE data sets.

The recommendation to migrate from IMBED, REPLICATE, and KEYRANGE was originally made in the z/OS V1R6 timeframe. In Software Announcement 204-180 (RFA39951), dated August 10, 2004, IBM announced its intent to withdraw support for VSAM IMBED, REPLICATE, and KEYRANGE attributes in a future release. Based on customer feedback, IBM no longer plans to remove this support from z/OS in the foreseeable future, as announced in Software Announcement 207-175 4 (RFA45594), dated August 7, 2007. IBM still recommends that you stop using these attributes.

### New function planned for z/OS R12!

In z/OS R12, the system will be designed to remove these attributes automatically from VSAM data sets logically dumped using DFSMSdss and migrated using DFSMSHsm when DFSMSdss is used as the data mover during restore and recall processing. An informational message is planned to confirm that newly restored data sets no longer retain these attributes.

### Migration action:

1. Determine which VSAM data sets were defined with the IMBED, REPLICATE, or KEYRANGE attribute. To help you perform this task, you can get a tool that reads existing VSAM data sets and reports which ones have these attributes. The tool is available from the software server (ftp.software.ibm.com) in the s390/mvs/tools directory as IMBDSHIP.JCL.TRSD. Download the file in binary format and uncompress it on your z/OS system using AMATERSE or TRSMMAIN. Instructions for using the tool are included in the downloaded JCL. See APAR II13894 for additional information. Use IBM Health Checker for z/OS check CATALOG\_IMBED\_REPLICATE on z/OS V1R11 to detect IMBED and REPLICATE attributes in your master catalog and any connected user catalogs.

### Notes:

- The IMBDSHIP tool only checks data sets that are on DASD. Data sets that were migrated by DFSMSHsm or backed up by DFSMSdss continue to retain their IMBED, REPLICATE, and KEYRANGE attributes when recalled or restored. In a future release of z/OS, when support for IMBED and REPLICATE is removed and data sets are recalled or restored by DFSMSHsm or DFSMSdss, the IMBED and REPLICATE attributes will be removed during the recall or restore and the data sets will be in a usable state. Because the data sets will be in a usable state, there is no need for a tool to identify data sets that were backed up or migrated at a time when the data sets had the IMBED or REPLICATE attribute and support for IMBED and REPLICATE still existed.
  - "AMATERSE" and "TRSMMAIN" are names for a service aid that compresses and decompresses data exchanged with IBM. "AMATERSE" is the preferred program name since its integration into z/OS V1R9. (AMATERSE is also available for z/OS V1R7 and V1R8 by PTFs UA36927 and UA36928 for APAR OA19194.) "TRSMMAIN" is the original program name and is now shipped as an alias entry point to AMATERSE. For more information about AMATERSE, including several differences with TRSMMAIN, see *z/OS MVS Diagnosis: Tools and Service Aids*.
2. Schedule a time for the affected VSAM data sets to be unavailable, and redefine them.

### **DFSMSdss: Define additional volumes to ML1 OVERFLOW and NOOVERFLOW pools for backup and migration (Required-IF, as of R11)**

*Required if you are migrating from a system that defined ML1 OVERFLOW volumes using the ADDVOL MIGRATION(OVERFLOW) command.*

The way that DFSMSHsm uses ML1 OVERFLOW volumes has changed. Before z/OS V1R11, volumes that you defined as OVERFLOW volumes using the ADDVOL MIGRATION(OVERFLOW) command were used for backup only for data sets greater than 500 tracks that encountered out of space errors. Beginning with z/OS V1R11, volumes defined with ADDVOL MIGRATION(OVERFLOW) are used by DFSMSHsm for both backup and migration of large data sets.

There is also a new command, which determines the size of data sets for which an ML1 OVERFLOW volume is preferred for migration or backup: SETSYS ML1OVERFLOW(DATASETSIZE(*dssize*) THRESHOLD(*threshold*))

The default for DATASETSIZE is 2000000 KB, so by default, DFSMSHsm backs up and migrates data sets that are 2000000 KB and larger to ML1 OVERFLOW volumes.

Note that many installations do not use ML1 OVERFLOW volumes. The default setting for ADDVOL MIGRATION is NOOVERFLOW, so if you do not specify the OVERFLOW parameter specifically now, your installation is probably not using ML1 OVERFLOW volumes and this migration action does not apply to you.

**Migration action:** Define additional volumes to ML1 OVERFLOW and NOOVERFLOW pools as follows:



1. Check to see if you are defining ML1 OVERFLOW volumes using the ADDVOL MIGRATION(OVERFLOW) command.
2. If you are, evaluate what the changed behavior will look like on your installation. See ML1OVERFLOW: Manage ML1 OVERFLOW volumes in z/OS DFSMSShsm Storage Administration for information on how different size data sets will be backed up and migrated to ML1 OVERFLOW or NOOVERFLOW volumes.
3. Define additional volumes to the ML1 OVERFLOW and NOOVERFLOW pools if needed.
4. If needed, adjust the SETSYS ML1OVERFLOW command settings to change the size of data sets that DFSMSShsm will prefer to migrate and back up to ML1 OVERFLOW volumes

**DFSMSrmm: Use the new default for the EDGHSKP XREPTTEXT DD statement (Recommended, as of R10)**

*Recommended because the new default requires less DASD space and fewer resources to create the extract.*

The default processing for the EDGHSKP XREPTTEXT DD statement is changed so that only extended records are created. Before z/OS V1R10, DFSMSrmm report extract processing used allocated ddnames to determine whether extended records were required. If the ddname was REPTTEXT, DFSMSrmm created all records except the extended records in the extract. If the ddname was XREPTTEXT, DFSMSrmm created all records in the extract and then used the D and V records to create extended records.

Starting in z/OS V1R10, by default, the XREPTTEXT DD statement receives only extended (X) records. To tailor the contents of the REPTTEXT and XREPTTEXT DD statements, DFSMSrmm report extract processing now checks for the existence of the RPTTEXT command in the SYSIN file, and if the RPTTEXT command is found, DFSMSrmm writes the selected extract records to the data set specified by either the REPTTEXT or XREPTTEXT DD statement.

The new default of only X records matches what most customers want. In addition, the new default requires less DASD space and fewer resources to create the extract.

**Migration action:** To obtain the new behavior, which is recommended, do nothing.

If you want the prior behavior, which is for the XREPTTEXT DD statement to contain all records, not just X records, update any report extract JCL used to create a report extract file containing extended records (run to create records other than the extended records) by adding the following statements to your existing JCL:

```
//SYSIN DD *
RPTTEXT RECORDS(X,O,R,B,S,P<B>,V,D,</B>
/*
```

This change must be made to your JCL before IPLing z/OS V1R11 and running extract, but it can be made on your current (pre-z/OS V1R10) release (with the coexistence PTF installed if z/OS V1R9) because it will only be used by DFSMSrmm on z/OS V1R10 and higher.



**DFSMSrmm: Migrate from VRSEL(OLD) to VRSEL(NEW) in parmlib member EDGRMMxx (Required-IF, as of R11)**

*Required if use the VRSEL(OLD) operand of the OPTION command in DFSMSrmm parmlib member EDGRMMxx.*

The VRSEL(OLD) or VRSEL(NEW) operand of the OPTION command in DFSMSrmm parmlib member EDGRMMxx specifies the type of vital record processing you want performed. In z/OS V1R11, VRSEL(OLD) has been removed. From z/OS V1R8 to V1R10, if you did not migrate from VRSEL(OLD) to VRSEL(NEW), you received warning message EDG2317E when you ran VRSEL(OLD) processing, and EDGHSKP processing ended with job step return code 4. Beginning with z/OS V1R11, you can no longer specify VRSEL(OLD) in parmlib and the default value is changed to VRSEL(NEW).

Specifying OPTION VRSEL(OLD) will result in error messages from parse:

```
EDG0208I RECORD INPUT:- OPTION VRSEL(OLD)
EDG0209E PARSE MESSAGE:- IKJ56712I INVALID KEYWORD, OLD
EDG0215D ERRORS DETECTED IN INITIALIZATION PARAMETERS
- ENTER ?Y? TO CONTINUE OR ?N? TO CANCEL
```

Reply with "Y" and the default VRSEL(NEW) is used. Reply with "N" and the response is:

```
EDG0104E DFSMSRMM SUBSYSTEM INITIALIZATION FAILED
EDG0107A ENTER SUFFIX OF INITIALIZATION MEMBER OR "CANCEL"
```

**Migration action:** To migrate from VRSEL(OLD) to VRSEL(NEW), follow the steps in z/OS V1R10.0 DFSMSrmm Guide and Reference.

**Tip:** Use IBM Health Checker for z/OS to determine whether you need to perform this migration action. The relevant check is named ZOSMIGV1R11\_RMM\_VRSEL\_OLD and is available through PTF UA45968 on z/OS V1R9 (APAR OA26947) and PTF UA45967 on z/OS V1R10 (APAR OA26947). The check determines whether you are still using OPTION VRSEL(OLD).

**DFSMSdfp: Update the storage size for jobs that invoke ADRDSSU (Required-IF, as of R11)**

*Required if the amount of storage requested for jobs that invoke ADRDSSU is 2M or smaller.*

Beginning with z/OS V1R11, jobs invoking DFSMSdss ADRDSSU might require a storage size of greater than 2M. This could require changes to the JCL for jobs that invoke ADRDSSU, or to JES initialization statements that specify a default region size for jobs.

**Migration action:**

1. Check the JCL for jobs that invoke ADRDSSU. If the REGION parameter specifies a region of 2M or smaller, increase the region size. If the SIZE parameter on the EXEC statement specifies a size of 2M or smaller, increase the size.
2. Check your JES initialization parameters. If the REGION parameter specifies a default region size of 2M or smaller, you may want to increase this.

**DFSMSdftp: Specify a control interval size for the IDCAMS DEFINE DATA component (Recommended, as of R11)**

*Recommended because specifying a control interval size other than 4096 (with the DEFINE command) can improve VSAM performance.*

Before z/OS V1R11, IDCAMS DEFINE could be used to specify CONTROLINTERVALSIZE for the CLUSTER component when calculating the space for a data set allocation. If a CONTROLINTERVALSIZE was specified for the DATA component, the specified value was ignored and the CLUSTER CONTROLINTERVALSIZE used. A default value of 4096 is used for the DATA component if a size is not specified for the CLUSTER. Beginning with z/OS V1R11, IDCAMS DEFINE now processes the CONTROLINTERVALSIZE being specified for the DATA component. Existing JCL, which specifies a CONTROLINTERVALSIZE for the DATA component, can result in a different allocation amount when executed on z/OS V1R11.

**Note:** To preserve consistency in space allocations across releases, specify CONTROLINTERVALSIZE on the CLUSTER level of the DEFINE command and apply PTF UA36500 (APAR OA21369) on systems running z/OS V1R9. Jobs that do not specify a CONTROLINTERVALSIZE for the DATA component are unaffected. Specifying a value other than 4096 can improve VSAM performance.

**Migration action:** Review production JCL libraries for instances of an IDCAMS DEFINE of a linear data set and specify a control interval size value other than 4096. The actual allocation amount will use the specified control interval size value instead of 4096. Determine if the specified value is appropriate for the data set.

**DFSMS Migration Actions Pre-First IPL**

**DFSMSdftp: Permit DEVMAN to call ICKDSF (Required-IF, as of R11)**

*Required if you use RACF program control to restrict access to ICKDSF.*

Starting with z/OS V1R11, support is added to the device manager (DEVMAN) to invoke ICKDSF for the new DEVMAN REFVTOC function. During startup of the DEVMAN address space, ICKDSF is loaded by DEVMAN regardless of whether or not you have enabled the REFVTOC function.

If you use program control in RACF to restrict access to ICKDSF, you must ensure the profile which protects the ICKDSF resource in the PROGRAM class allows the user ID associated with the DEVMAN address space to have READ access.

Message ICH408I will be issued when the DEVMAN address space is started if you use program control for ICKDSF and you have not allowed DEVMAN to call ICKDSF. You will not be able to exploit the new REFVTOC function. Note that DEVMAN initialization continues even if DEVMAN does not have sufficient authority.

**Migration action:** If you use RACF:

1. Define a user ID for the DEVMAN started task if one is not already established.
2. Define a profile for DEVMAN in the RACF STARTED class. Assign the user ID from Step 1 to the DEVMAN procedure in the STARTED class.
3. Permit the user ID for the DEVMAN started task to call ICKDSF by allowing read access to ICKDSF in the PROGRAM class. For example: PERMIT ICKDSF CLASS(PROGRAM) ID(userid) ACCESS(READ)
4. Refresh the storage RACF profiles by using the command SETR WHEN(PROGRAM) REFRESH.
5. Refresh the STARTED profiles, if necessary, using the command SETR RACLIST(STARTED) REFRESH.

**DFSMSdftp: Use a permanent VSAM data set if using REUSE with RESET (Required-IF, as of APAR OA29467, on R9, R10, and R11)**

*Required if you OPEN a temporary VSAM data set with RESET option.*

VSAM enables you to define reusable data sets to use as work files. Instead of using the DELETE- -DEFINE sequence, you can specify the REUSE parameter in the IDCAMS DEFINE CLUSTER|ALTERNATEINDEX command. The REUSE parameter lets you treat a filled data set as if it were empty and load it again and again regardless of its previous contents. In a reusable data set, you can reset to zero the high-used RBA field at OPEN by specifying MACRF=RST in the ACB at OPEN. VSAM can use this reusable data set like a newly defined data set. VSAM does not support RESET for temporary data sets. However, in z/OS V1R9, z/OS V1R10, and z/OS

V1R11, without the PTF for APAR OA29467 applied, even when a temporary VSAM data set is opened with RESET, the unsupported option is ignored with no message issued to notify the user. Therefore, there are unexpected results because the existing records in the data set are not cleared during the OPEN process.

After the PTF for APAR OA29467 is applied to z/OS V1R9, z/OS V1R10 or z/OS V1R11, OPEN will be failed with message IEC161I 084-061 if REUSE with RESET is used with temporary VSAM data sets.

**Migration action:** Use a permanent data set if you wish to use REUSE with RESET on VSAM data sets. Do not use REUSE with RESET on temporary VSAM data sets. The message IEC161I 084-061 reports a cluster without the REUSE attribute, which indicates the object of an unsuccessful reset request. OPEN processing ends for the data set and the error flag (ACBERFLG) in the ACB (access method block) for the data set is set to 232 (X'E8').

**Notes:**

1. Temporary VSAM data sets must reside in storage managed by the Storage Management Subsystem.
2. The temporary data set **does not** need to be defined with a data class specifying REUSE to get the error message IEC161I 084-061.

**DFSMSdftp: Accommodate layout change of IEHLIST LISTVTOC command output with FORMAT option (Required-IF, as of R11)**

*Required if you depend on the layout of PGM=IEHLIST LISTVTOC command output with FORMAT option.*

Beginning with z/OS V1R10, the output of PGM=IEHLIST LISTVTOC command with the FORMAT option is changed to support the Extended Address Volume (EAV) capability.

- The layout of output for the VTOC Format-4 DSCB is changed. It affects the position of header and values.
- For the components of VSAM data sets, ICF Catalogs and VVDS data sets, no value (blank) is printed in the field of LAST BLK(T-R-L). In prior releases, it contained the value of 0 0 58786, which is derived from the VTOC Format-1 DSCB field DS1TRBAL (X'E5A2').
- At the bottom of SYSPRINT, there is new information printed in the free space section as follows: THERE ARE *nnnnn* EMPTY CYLINDERS PLUS *nnnnnn* EMPTY TRACKS FROM THE TRACK-MANAGED SPACE

**Migration action:** Because the layout of SYSPRINT data set is changed, you need to review and change the post-processing program accordingly.

**Note:** After applying the PTF for APAR PK56092 for ICKDSF V1R17, the default size of the VTOC index is changed to be based on the required index records for a given volume size and the size of the VTOC. When the index size defined is smaller than 15 tracks (old default size), the number of VIRs obtained within the index will be decreased, which is reported by the PGM=IEHLIST LISTVTOC command. THERE ARE *nnnnnn* UNALLOCATED VIRS IN THE INDEX

**DFSMSdftp: Add columns ODSTATF, ODRETD, and ODINSTID to OAM object directory tables (Required-IF, as of R11)**

*Required if you use OAM object support.*

In z/OS V1R11, object access method (OAM) introduces 2 GB object support for tape media as well as archive retention enhancements. You must run the CBRSMR1B job to perform the migration from the z/OS V1R10 version of the object storage database to the z/OS V1R11 version. The job will add columns ODSTATF, ODRETD, and ODINSTID to the existing object directory tables. Even if you do not take advantage of the new function added in this release, if you use OAM object support you must run the CBRSMR1B job to add the new columns to the tables.

**Migration action:**

1. If your installation's OAM object directory tables in DB2 do not contain ODSTATF, ODRETD, and ODINSTID columns, edit the sample job CBRSMR1B to make the changes specified in the comments section, and run the job. The job adds the new columns to the OAM object directory tables and primes them with default values.
2. Run OAM DB2 BIND and GRANT jobs. To determine which BIND and GRANT jobs you need to run, see *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Object Support*.

**DFSMSdftp: Ensure that the Language Environment run-time library is available for DLLs (Required-IF)**

*Required if your installation builds or references DLLs.*

Language Environment provides common services and language-specific routines in a single run-time environment. You can use Language Environment to build and use dynamic link libraries (DLLs) for applications.

**Migration action:** If your installation builds or references DLLs, either you must set up the system link list to refer to the Language Environment run-time libraries (SCEERUN and SCEERUN2), or each job that creates or uses a DLL must include a STEPLIB DD statement referencing these libraries.

**DFSMSdss: Build the IPLable stand-alone DFSMSdss image (Required-IF)**

*Required if you intend to use the Stand-Alone Services provided by DFSMSdss.*

If you intend to use the Stand-Alone Services provided by DFSMSdss, you must use the DFSMSdss BUILDSA function to create the Stand-Alone Services IPL-capable core image

**Migration action:**

1. Prepare for Stand-Alone Services by creating a Stand-Alone Services IPL-able core image with the BUILDSA command. With the BUILDSA command you can specify the device (card reader, tape drive, or DASD volume) from which Stand-Alone Services will be IPLed. You can also specify the operator console to be used for Stand-Alone Services. The BUILDSA function builds the IPL-able core image under the current operating system and determines a record size based on whether the IPL is from card, tape, or DASD.
2. Use RACF or another external security system to protect the SYS1.ADR.SAIPLD.vvolser data set and the Stand-Alone Services modules.
3. If you haven't done so already, make a backup copy of your system that can be restored by this function. For information about backing up volumes, see *z/OS DFSMSdss Storage Administration Guide*.

**DFSMSdfp: Ensure your SYS1.IMAGELIB is customized for your printing environment (Required-IF)**

*Required if you are not using your old SYS1.IMAGELIB, you are installing with ServerPac or SystemPac, and you are using line mode printers such as the 3800 or 3900.*

If you use line mode printers such as the IBM 3800 or the IBM 3900 running in line mode (not page mode), you must install library character sets, graphic character modification modules, and character arrangement tables in SYS1.IMAGELIB. This migration action does not apply if you are using IBM 3900 printers that are driven by PSF.

**Migration action:**

1. Run the LCSBLD1 job from the samplib data set to create character sets, graphic character modification modules, and character arrangement tables in SYS1.IMAGELIB.
2. Copy customized or locally-written FCBs and UCS images from your old system's SYS1.IMAGELIB data set to the new system's SYS1.IMAGELIB data set.

**DFSMSdfp: Remove user-defined trace points for VSAM record management trace (Required-IF, as of R11)**

*Required if you have previously added trace points for the VSAM record management trace function.*

Before z/OS V1R11, DFSMS included an option to add your own trace points for the VSAM record management trace function. This option has been removed in z/OS V1R11. Therefore, any applications that have user-defined trace points for the VSAM record management trace function need to have them removed.

**Migration action:** Remove user-defined trace points for the VSAM record management trace function from your applications. If a VSAM trace definition contains any trace point ID that is not listed in *z/OS DFSMSdfp Diagnosis*, remove the user-defined trace point ID from the HOOK parameter in the DD statement. Failure to remove user-defined trace points could cause an ABEND0C1 when running with VSAM record management trace. You might be able to use new IBM-defined trace options in place of a user-defined one. See the list of IBM-defined trace points in *z/OS DFSMSdfp Diagnosis*.

**DFSMSdfp: Provide a 140-byte work area when using OBTAIN (Required-IF, as of R10)**

*Required if you have programs that do not specify at least a 103-byte DSCB return area.*

Programs must pass the address of a 140-byte return area when they search for a data set name using the OBTAIN macro with the CAMLST SEARCH option or the OBTAIN SVC 27 with the SEARCH for DSNAME operation code. In releases prior to z/OS V1R10, the system returned data in only the first 101 bytes of the return area. Beginning with z/OS V1R10, the system will return data in the first 103 bytes. Therefore, you should upgrade your programs to provide the entire 140-byte return area; otherwise, you might encounter storage overlay problems or program interruptions.

**Migration action:** Upgrade programs to provide the entire 140-byte return area as stated in the reference information below.

**DFSMSdfp: Accommodate new defaults for z/OS Global Mirror (XRC) parameter values (Req-IF, as of R10)**

*Required if you rely on the previous XRC default values.*

Default values have changed for two parameters of the SHADOW statement in the XRC parmlib member. Before z/OS V1R10, the RequireUtility parameter had a default value of NO, and the UtilityDevice parameter had a default value of FLOAT. In z/OS V1R10, these default values have changed as follows:

- The default value of the RequireUtility parameter is now YES, which prevents XADDPAIR of a non-XRCUTL volume when an active storage control session does not exist for the specified (or default) SCSESSION. The previous default of NO enables XADDPAIR of a non-XRCUTL volume to proceed when an active storage control session does not exist.

- The default value of the UtilityDevice parameter is now FIX, which changes the method used for selecting a utility device. The FIX value causes XRC to use the most recently added XRCUTL volume as the utility device for the storage control session. The previous default of FLOAT allows the storage control to dynamically pick the primary volume to be used as the utility device for the storage control session. The nonreserved device with lowest I/O activity will typically be selected

**Migration action:** If your XRC sessions rely on the previous default values, you must add statements in your XRC parmlib to specify those values, thus overriding the new defaults.

#### **DFSMSdfp: Update programs that process SMF record type 19 (Required-IF, as of R10)**

*Required if you rely on the previous SMF 19 record length.*

In z/OS V1R10, the length of SMF type 19 records has increased. A type 19 record is created for each DASD volume that is online when IPL occurs, when a HALT EOD or SWITCH SMF command is issued, and when a DASD is varied offline. The type 19 record describes the space usage on the volume, and new fields have been added to the record for track-managed free space statistics and volume size information. Programs that process this record type might need to be modified to support the expanded length of the record or to utilize the new fields. SMF type 19 records are recorded by LSPACE processing when SMF=YES or SMF=TEST is specified on the LSPACE macro invocation. YES specifies that the caller wants an SMF record containing volume information to be written. TEST specifies that LSPACE is to test for an active SMF system that has been defined to request volume accounting before writing the SMF record.

**Migration action:** Check any programs that process SMF type 19 records to determine if you must make any changes. If necessary, modify the programs to use the new record length.

#### **DFSMSdfp: Accommodate changes to DEVSERV command output (Required-IF, as of R10)**

*Required if you have a dependency on the output of the DEVSERV QDASD or DEVSERV PATH command.*

z/OS V1R10 includes several changes to output from the DEVSERV command:

- DEVSERV QDASD output is changed as follows: the CYL column was expanded to accommodate a nine-digit number for cylinders. In addition, the extended function consistency check column EF-CHK is reduced to three characters. The new column heading is ECK.
- DEVSERV PATHS output is changed to add the number of cylinders, in a new column labeled CYL. To accommodate the largest possible number of cylinders for a device (nine digits) in this new column, the original ALT column is replaced with this new CYL column and unnecessary spaces around the PIN and DDC columns are removed.
- The output for both DEVSERV QDASD and PATHS has been shifted one position to the right, leaving the first character position blank.

**Note:** In z/OS V1R11, or after the PTF for APAR OA25793 is installed on z/OS V1R10 and z/OS V1R9, DEVSERV PATHS and the QDASD command display 5-digit device numbers that consist of the subchannel set number and the normal device number. The blank that was in the first character position no longer exists and is replaced by 0 or 1 (subchannel set number). The unit field will be changed to "00ddd" from " 0ddd" in case of the 3-digits device number.

**Migration action:** Change any dependency you have on the output of DEVSERV QDASD or DEVSERV PATHS. Examples of the output changes can be found in *z/OS Migration*.

#### **DFSMSdfp: Update programs that use DCOLLECT (Required-IF, as of R10)**

*Required if you use DCOLLECT to process data set (type D) records to determine whether data sets are overallocated.*

In a DCOLLECT type D record, the DCDOVERA field is a 4-byte field that reports on the amount of overallocated space for a data set. Before z/OS V1R10, the field was defined as a 31-bit signed number. Beginning with z/OS V1R10, DCDOVERA is defined as a 32-bit unsigned number in order to accommodate larger values.

**Migration action:** If you use the IDCDOUT macro to map DCOLLECT records, reassemble your programs to pick up the current mapping.

If you do not use the IDCDOUT macro to map DCOLLECT records, be aware that the field definition has changed for DCDOVERA from signed to unsigned. Use the new field definition when you create DCOLLECT reports.

#### **DFSMSdfp: Accommodate changes in the CA sizes assigned to newly allocated VSAM data sets (Required-IF, as of R10)**

*Required if you have a dependency on the CA size selected for a VSAM data set and the corresponding primary or secondary space amounts, or you have a dependency on the index control interval size (which may be increased by the CA size selection).*

z/OS V1R10 includes a change in control area (CA) size selection for all new allocations of VSAM data sets on any volume type (not only extended address volumes). The system selects a control area size of 1, 3, 5, 7, 9, or 15 tracks. Previously, the system selected from a wider variety of possible CA sizes. As previously, you cannot explicitly specify control-area size. Generally, the primary and secondary space allocation amounts determine the CA size. For requests where either the primary or secondary allocation amount is smaller than one cylinder, the system might adjust the primary and/or secondary quantity. The system might also select a CA size that is different from what was selected in a prior release. For example, a TRK(24,4) request results in a control area of 5 tracks, and primary and secondary amounts of 25 and 5 tracks, respectively.

**Migration action:** CA size is derived from the requested primary and secondary allocation. Because the CA size is computed differently in z/OS V1R10, adjust any primary and secondary allocation requests to accommodate the CA size you rely upon. If you have any programs that depend upon a specific CA size that is no longer supported, change the programs to accept a CA size of 1, 3, 5, 7, 9, or 15 tracks.

**Tip:** A tool can help you perform this migration action. The tool is available from the FTP site ftp.software.ibm.com in the directory servers/storage/support/software/dfsms/ as INDXCI10.JCL.CNTL.TRSD. Download the tool in binary mode to a data set with the following attributes: LRECL=1024, BLKSIZE=6144, RECFM=FB, and DSORG=PS. After downloading, untermse it either with AMATERSE or TRSMAIN. The tool consists of a REXX exec and JCL to execute it in batch.

The tool lists the data sets that will have their index control interval (CI) size increased when they are redefined on z/OS V1R10. (The increase can occur if the system-selected CA size is currently not 1, 3, 5, 7, 9, or 15 tracks.) You should ensure that any applications such as CICS, IMS, SMP/E, Netview, Batch Local Shared Resources (LSR), or Information Management that build LSR pools have buffers available for the increased index CI size. Otherwise, the data set will fail to open with an IEC161I 120-053 error message or other error messages issued by the application. Informational APAR II13288 describes a number of error situations that might occur if the index CI size is increased and proper preparation has not been taken. Review the comments in the JCL for more information about the tool.

**DFSMSdftp: Tell users that the ISMF SORTREC line operator is no longer available (Required-IF, as of R10)**

*Required if you use the SORTREC line operator.*

Beginning with z/OS V1R10, the DFSORT interactive storage management facility (ISMF) panels are no longer supported. Consequently, the SORTREC line operator is no longer available.

**Migration action:** Tell users that the SORTREC line operator is no longer available. Trying to use it will result in the message COMMAND SORTREC NOT FOUND. An alternative to SORTREC is to use DFSORT against the data set from a batch job.

**DFSMSdftp: Ensure that the ANTMMAIN address space has access to parmlib data sets (Recommended, as of R10)**

*Not required, but recommended to avoid security messages when ANTMMAIN tries to access parmlib data sets.*

A new parmlib member, ANTMMAIN00, was introduced in z/OS V1R10. It contains new parameters related to the virtual-concurrent-copy function of DFSMSDSS. The ANTMMAIN address space, which is started automatically at IPL, will attempt to read this member from the parmlib data set concatenation. If your installation restricts started task access to these parmlib data sets, ANTMMAIN should be given authorization to read these parmlib data sets.

Authorization failure could result in messages ICH408I, IEC150I 913-38, IEA995I for SYSTEM COMPLETION CODE=913, and ANTI1027E. Note that ANTMMAIN initialization continues if ANTMMAIN cannot read the parmlib data set and the default virtual-concurrent-copy parameters are in effect.

**Migration action:** Allow the ANTMMAIN started task read access to the data sets in the parmlib concatenation.

**DFSMSdftp: Allow the dynamic volume count (data class attribute) to exceed 1 when allocating SMS-managed data sets with partitioned organization (Required-IF, as of OA22738)**

*Required if you use data classes that specify dynamic volume counts when allocating data sets with partitioned organization, and you want to allow the dynamic volume count (data class attribute) to exceed 1 (which is the pre-z/OS R8 situation).*

Before z/OS V1R8, when allocating an SMS-managed data set with partitioned organization, you could assign to it a data class that specified a dynamic volume count greater than 1. Then, in z/OS V1R8 and V1R9, the base code was changed so that the dynamic volume count could not be greater than 1. However, APAR OA22738 on z/OS V1R8 and V1R9 restored the original situation, that is, the dynamic volume count could be greater than 1 again. Now, in z/OS V1R10, the original (pre-z/OS V1R8) situation (greater than 1 allowed) has been permanently restored. This migration action is not relevant to non-SMS data sets because the dynamic volume attribute only applies to SMS-managed data sets.

**Migration action:** If you desire, allow ACS routines to select a data class with a dynamic volume count greater than 1 for SMS-managed data sets with partitioned organization. Note that because PDSs and PDSEs can only exist on one volume (and cannot be extended to a second volume), the dynamic volume count will be ignored.

**DFSMSdfp, DFSMSdss, and DFSMShsm: Update automation for changed messages (Req-IF, as of R10)**

*Required if you have automation routines that examine the message text of the messages listed below.*

In z/OS V1R10, the text for messages with the following IDs has been changed: ADRxxxxx (DFSMSdss), ANTxxxxx (DFSMSdfp), ARCxxxxx (DFSMShsm), DMOxxxxx (DFSMSdfp), IDCxxxxx (DFSMSdfp), IEAxxxxx (DFSMSdfp), IECxxxxx (DFSMSdfp), IEExxxxx (DFSMSdfp), and IGDxxxxx (DFSMSdfp).

Some messages have had a period (.) removed from the end of the message text. Other messages have had text added, changed, or removed. These changes can affect automation programs that examine the text of the messages.

**Migration action:** Update your automation to handle the following **DFSMSdfp** message changes:

- The message text or explanation has been changed in the following messages to support extended address volumes, including new track, cylinder, or return code information:
  - ANTF073E, ANTF0401E, ANTF0443E, ANTX5117E, ANTX8030W, ANTM6004I, ANTM6005I, ANTM6006E, ANTM6007E, ANTM6013E, ANTM6015I, ANTM6016I, ANTM6020I, ANTM6021I, ANTM6022I, ANTM6023I, ANTM6024I, ANTM6025I, and ANTM8809I
  - DMO0011I, DMO0012I, DMO0013E, DMO0031E, DMO0032E, DMO0033I, DMO0040I, and DMO0041I
  - IDC3009I
  - IEA434I, IEA479E, IEA480E, and IEA480E
  - IEC142I, IEC144I, IEC602I, IEC609I, and IEC987I
  - IEE459I
  - IGD002I, IGD031I, IGD17268I, IGD17279I, IGD17269I, IGD17318I, and IGD17389I
- Message IGD17038I now provides a number of volumes (“n volumes”) in place of the words “this request”.

Update your automation to handle the following **DFSMSdss** message changes:

- The period has been removed from the end of the following messages: ADR021I, ADR032W, ADR037E, ADR045T, ADR109I, ADR146I, ADR148I, ADR204W, ADR238I, ADR239I, ADR240W, ADR226W, ADR297I, ADR299E, ADR301E, ADR325E, ADR343E, ADR344I, ADR345D, ADR359W, ADR363E, ADR365E, ADR388W, ADR390I, ADR397E, ADR401W, ADR405E, ADR444E, ADR459E, ADR499I, ADR701E, ADR702E, ADR711I, ADR715W, ADR716W, ADR719E, ADR721I, ADR731W, ADR733I, ADR734I, ADR736E, ADR737W, ADR767I, ADR785E, ADR795E, ADR796E, ADR801I, ADR804W, ADR805E, ADR865E, ADR880I, ADR901E, ADR902I, ADR903I, ADR904I, ADR906I, ADR907I, ADR908E, ADR912I, ADR913E, ADR916W, ADR940I, ADR941E, ADR942E, ADR943E, ADR944E, ADR946E, ADR969E, ADR970E, ADR974I, ADR978W, ADR979E, ADR985W, ADR990E, and ADR993E.
- The message text or explanation has been changed in the following messages to support extended address volumes, including new track, cylinder, or return code information: ADR024E, ADR037E, ADR037E, ADR040I, ADR071I, ADR145I, ADR152E, ADR201E, ADR210E, ADR226W, ADR232E, ADR234I, ADR235I, ADR247W, ADR250I, ADR251W, ADR255E, ADR256E, ADR258I, ADR260I, ADR261I, ADR262I, ADR299E, ADR309E, ADR317I, ADR336I, ADR345D, ADR364W, ADR366W, ADR373E, ADR379E, ADR433E, ADR436E, ADR441W, ADR498E, ADR520I, ADR557I, ADR558I, ADR735W, ADR817W, ADR818E, ADR887I, ADR888I, ADR932W, ADR935W, ADR938E, ADR939E, ADR945W, ADR946I, and ADR973E. (The following messages have also been changed but they will never be analyzed by automation because they are only issued during stand-alone restore when z/OS is not in operation: ADRY0500I, ADRY0700I, ADRY1500W, ADRY1524W, ADRY1711I, ADRY2500I, and ADRY3521I.)
- The ending period has been removed and additional message text has been added to messages ADR037E and ADR345D.
- Message ADR049E has been enhanced to include abend reason code information at the end.
- References to JOBCAT and STEPCAT have been removed from messages ADR380E, ADR415W, ADR470W, ADR472E, and ADR485E.

Update your automation to handle the following **DFSMShsm** message changes:

- The message text or explanation has been changed in the following messages to support extended address volumes, including new track, cylinder, or return code information: ARC0120I, ARC0157I, ARC0185I, ARC0269I, ARC0638I, ARC0742I, ARC1309I, ARC1359I, ARC1520I, and ARC1524I.

**DFSMShsm: Handle changed LIST command output (Required-IF, as of R11)**

Required if your applications depend the RACF IND field value in the output of these commands with the OUTPUTDATASET destination.

Beginning with z/OS V1R11, the LIST DSNAME(*dsname*) BCDS and LIST LEVEL(*hlq*) BCDS output will no longer display the RACF IND field when OUTPUTDATASET is specified as the destination for the output. The RACF IND field will still be displayed when TERMINAL is specified as the destination for the output.

**Migration action:** Remove any dependency on the RACF IND field on the LIST DSNAME(*dsname*) BCDS or LIST LEVEL(*hlq*) BCDS output when using OUTPUTDATASET.

### **DFSMSHsm: Use ALL, DATE, or VERSIONS with the (H)BDELETE command and ARCHBDEL macro (Required-IF, as of R11)**

Required if you use the (H)BDELETE command or the ARCHBDEL macro.

Before z/OS V1R11, you could use the (H)BDELETE *dsname* command to delete all backup versions of data set *dsname*. In z/OS V1R11, you must specify the ALL keyword on the (H)BDELETE *dsname* command to delete all backup versions of the data set.

**Migration action:** In z/OS V1R11, you must use the (H)BDELETE command with one of these mutually exclusive keywords: ALL, DATE, or VERSIONS. The new keyword, ALL, specifies that DFSMSHsm is to delete all backup versions of the specified data set, including both active and retained copies, except for the retired versions (if they exist). The DATE and TIME keyword deletes active or retained backup versions. The VERSIONS keyword deletes only active backup copies.

### **DFSMSHsm: Maintain coexistence when encountering backup copies with RETAIN DAYS value (Required-IF, as of R11)**

Required if you are using DFSMSHsm functions to process any data set backup copies with a RETAIN DAYS value made on z/OS V1R11 DFSMSHsm, you must apply toleration APAR OA26327.

In z/OS V1R11, the (H)BACKDS command is enhanced to allow users to specify a RETAIN DAYS keyword to create a backup copy with a specified retention period. DFSMSHsm maintains backup copies as active backup copies and retained backup copies. Active copies are the backup copies that have not yet rolled off. Retained copies are the backup copies that have rolled off from the active copies, but have not yet reached their retention periods.

Pre-z/OS V1R11 DFSMSHsm functions that encounter a backup copy with a RETAIN DAYS value, or a retained backup copy made on a z/OS V1R11 system, will have limited processing ability. The functions affected include (H)BACKDS, (H)RECOVER, EXPIREBV, (H)BDELETE, DELVOL, FREEVOL, AUTOBACKUP, RECYCLE, and AUDIT.

Refer to z/OS R11 Migration for coexistence considerations.

**Migration action:** Perform the following functions from a z/OS V1R11 system when data set backup versions are being created on z/OS V1R11 with the RETAIN DAYS keyword: (H)BACKDS, (H)RECOVER, EXPIREBV, (H)BDELETE, DELVOL, RECYCLE, AUTO BACKUP, FREEVOL, and AUDIT.

### **DFSMSHsm: Update the HZSPRMxx DATE parameter for HSM\_CDSB \* checks (Required-IF, as of R10)**

Required if you have policy statements in HZSPRMxx parmlib members for any of the following checks:

HSM\_CDSB\_BACKUP\_COPIES, HSM\_CDSB\_DASD\_BACKUPS, or HSM\_CDSB\_VALID\_BACKUPS.

In z/OS V1R10, DFSMSHsm made enhancements to correct a problem with the HSM\_CDSB\_\* checks. Before z/OS V1R10, if one of these checks ran while you were doing maintenance for the CDS involving the MHCR cluster number field being set to X'FF', these checks might fail continuously (even after maintenance has completed) with the following error: HZS1002E CHECK(IBMHSM,HSM\_CDSB\_ *checkname*): AN ERROR OCCURRED, DIAG: A124000\_000000FF

In z/OS V1R10, this problem was corrected so that the check fails only until maintenance is complete, with a specific error message: ARCHC0003I The MHCR indicates that the storage administrator is performing maintenance on the control data sets. The check is disabled with REASON=ENVNA until the maintenance is complete.

Note that these checks cannot run successfully during CDS maintenance.

You must update the default data parameter for these checks in any HZSPRMxx policies that reference them. In addition, you should also consider returning to default values for any other parameters you changed in HZSPRMxx policies because of the continuous HZS1002E error, such as SEVERITY or INTERVAL.

**Migration action:** Change the DATE parameter value in each policy statement referencing an HSM\_CDSB\_\* check to 20071031.

In addition, you should also consider returning to default values for any other parameters you changed in HZSPRMxx policies to make the continuous HZS1002E error less conspicuous, such as SEVERITY or INTERVAL. For example, ensure that the policy UPDATE statement for the HSM\_CDSB\_BACKUP\_COPIES check looked as follows:

```
UPDATE CHECK(IBMHSM,HSM_CDSB_BACKUP_COPIES)
SEVERITY(MED) INTERVAL(24:00) PARM('CRITVAL(4)')
DATE(20071031)
```



REASON( Ensure critical value of HSM CDS backups specified. )

Note that after you make the DATE change to the policies for this check, the policy statements will be rejected until the migration to z/OS V1R10 is made.



**DFSMSrmm: Use DFSMSrmm parmlib options instead of USERMOD to specify how VRSEL processing handles duplicate GDGs (Required-IF, as of R11)**

*Required if a usermod is used to influence duplicate GDG handling.*

Before z/OS V1R11, a USERMOD could be applied to the EDGVREC load module to influence how VRSEL processing handles duplicate generation data groups (GDGs). The USERMOD may be called RMDUPGD. Starting with z/OS V1R11, this USERMOD is no longer supported. You must now use the GDG operand of the OPTION command in parmlib member EDGRMMxx to specify how VRSEL processing handles duplicate generations.

**Migration action:**

1. If a USERMOD is used to influence duplicate GDG handling, do not install it on z/OS V1R11. Remove the USERMOD if already installed on z/OS V1R11. If the USERMOD is retained, it will be ignored and the following warning message will be issued to the MESSAGE file during VRSEL processing: EDGX004I USERMOD HAS BEEN DETECTED AND IGNORED – USE THE GDG(CYCLEBY(GEN) DUP(COUNT)) PARMLIB OPTION INSTEAD.
2. Specify the desired duplicate GDG handling by using the GDG operand of the OPTION command. To obtain consistent VRSEL processing results, you can:
  - Run VRSEL either on your existing release or exclusively on z/OS V1R11.
  - Ensure that the GDG parmlib option is specified to match the results obtained from your existing system and run VRSEL on any system. The parmlib option that closely matches the USERMOD is GDG(CYCLEBY(GENERATION) DUPLICATE(COUNT)).

**Tip:** Use IBM Health Checker for z/OS to determine whether you need to perform this migration action. The relevant check is named ZOSMIGV1R11\_RMM\_DUPLICATE\_GDG and is available through PTF UA45968 on z/OS V1R9 (APAR OA26947) and PTF UA45967 on z/OS V1R10 (APAR OA26947). The check determines whether you have duplicate GDG processing defined.

**DFSMSrmm: Update operator procedures and system automation for dynamic installation exits (Required-IF, as of R11)**

*Required if DFSMSrmm installation exits are in use.*

Before z/OS V1R11, installation exits were loaded by DFSMSrmm during initialization. Starting with z/OS V1R11, z/OS dynamic exits services is used to load and activate the default (EDGUXn00) exit modules during initialization. This change requires changes to DFSMSrmm operating procedures and system automation (if any).

**Migration action:**

- Update operator procedures to replace any mention of F DFRMM,REFRESH EXITS with use of the MVS operator commands that can be used with dynamic exits.
- If you use automation routines, ensure that any automated processing you have for the deleted messages EDG0301I, EDG0302I, EDG0303D, EDG0304I, EDG0305I, and EDG0306I is updated to instead process the EDG0311I through EDG0314I messages.
- If you use DFSMSrmm installation exits, you do not need to change them in any way; just install them as you always have. DFSMSrmm automatically exploits the dynamic exit services using your installation exits as the default exit modules at the first startup of the DFRMM procedure. You do not need to change the load module names used as exit routines with DFSMSrmm exits. However, you may change the names if desired. But if you do, you will have to update PROGxx or issue the SETPROG command to get them loaded because DFSMSrmm will not load the installation exits.

You can now have multiple exit modules associated with each of the exits. Other programs can use the CSVDYNEX macro to associate their exit modules with one or more of the DFSMSrmm exits, and can activate and deactivate as required. You can use the PROGxx member of parmlib or the operator SETPROG command to add and activate exit modules for each of the DFSMSrmm exits.



**DFSMSrmm: Replace obsolete stem variables in REXX execs (Required-IF, as of R11)**

*Required if you have written any REXX execs that are used with any of the DFSMSrmm SEARCH or LISTCONTROL subcommands.*

Before z/OS V1R11, when DFSMSrmm created REXX variables for SEARCH and LISTCONTROL subcommands, it returned these variables as stem variables and created a .0 stem variable for almost all stem variables. Starting with

z/OS V1R11, DFSMSrmm creates a .0 stem variable only for the key variable for each SEARCH and LISTCONTROL command.

**Migration action:** Use the sample JCL from SAMPLIB (EDGJSTM0) to determine whether your execs must be updated. EDGJSTM0 runs ISRSUPC with the SRCHCMP parameter with SRCHFOR 'stem\_variable' process statements to scan REXX exec libraries. If any REXX exec is identified as affected, you must replace the obsolete variable with the appropriate stem 0 variable for the subcommand processed.

**Note:** Once updated, the execs can be run on any supported release of DFSMSrmm because the key stem variables are always present on prior releases.

**Tip:** Use IBM Health Checker for z/OS to determine whether you need to perform this migration action. The relevant check is named ZOSMIGV1R11\_RMM\_REXX\_STEM and is available through PTF UA45968 on z/OS V1R9 (APAR OA26947) and PTF UA45967 on z/OS V1R10 (APAR OA26947). This check determines whether installation-written REXX execs that issue DFSMSrmm TSO/E subcommands use stem variables that are removed in z/OS V1R11 systems.



### **DFSMSrmm: Define default partitioning support (Recommended, as of R10)**

*Recommended if you have no REJECT commands in the DFSMSrmm parmlib so that you get, by default, the desired type of volumes added to the DFSMSrmm control data set when the volumes are not defined to DFSMSrmm.*

Before z/OS V1R10, the DFSMSrmm default processing for volumes that were not defined to DFSMSrmm was to add only system-managed volumes, and not non-system-managed volumes, to the DFSMSrmm control data set. Starting in z/OS V1R10, the DFSMSrmm default processing when PRTITION commands are in use, either because they are defined or because there are no REJECT commands, is to add all undefined volumes. If you do not use REJECT commands in parmlib, you must add a PRTITION command if you want to ensure that the DFSMSrmm default processing stays the same.

#### **Migration action:**

- To ensure that the DFSMSrmm default processing stays the same, add the following PRTITION command to parmlib: PRTITION VOLUME(\*) TYPE(NORMM) NOSMT(IGNORE)  
As a result, system-managed volumes are automatically added to the DFSMSrmm control data set during library entry/insert and OPEN processing, but non-system-managed volumes are not added automatically. Note that when adding volumes automatically, DFSMSrmm no longer creates a rack number but will use an empty rack number if it exists.
- To get the new behavior, do one of the following:
  - If you have no REJECT commands in parmlib, take no action. The new default is provided by the default PRTITION command created by DFSMSrmm.
  - If you have REJECT commands in parmlib, migrate to PRTITION and OPENRULE commands to exploit the new capabilities.

**Tip:** Use IBM Health Checker for z/OS check ZOSMIGV1R10\_RMM\_REJECTS\_DEFINED to determine whether your system is using REJECT commands in parmlib and whether a migration action is required. The check is available through APAR OA26608.



### **DFSMSrmm: Update volume replacement policies (Required-IF, as of R10)**

*Required if you are using the usermod to the EDGMUPD module to override the shipped value.*

If you are using the usermod to the EDGMUPD module to set LIMIT=*n*, where *n* is shipped as X'00000001', you can no longer use this usermod. Before z/OS V1R10, you could use the usermod to override the shipped value. The usermod could disable the volume replacement action setting by setting a value of X'00000000'. The usermod was first available with APAR OW43870 and may be called RMREPR.

Starting in z/OS V1R10, you must use the function of MEDINF REPLACE in parmlib to implement any override or disablement. The usermod no longer works with the EDGMUPD module, and DFSMSrmm processing is based solely on MEDINF REPLACE.

#### **Migration action:**

- To disable the replacement of IBM media, you can specify the following in the EDGRMMxx parmlib member: MEDINF NAME(IBM) REPLACE(PERM(0))
- To override the DFSMSrmm built-in default processing of PERM(1), use any of the available REPLACE operand values.

**Tip:** Use IBM Health Checker for z/OS to help with this migration action. The relevant check is named ZOSMIGV1R10\_RMM\_VOL\_REPLACE\_LIM and is available through APAR OA26608. The check determines

whether you have changed the value of the hard-coded volume replacement limit (LIMIT=1) and shows you how to set the same limit in z/OS V1R10 and later using MEDINF NAME(IBM) REPLACE(PERM(*value*)).



### **DFSMSrmm: Use DELETE disposition support for tape data sets (Required-IF, as of R10)**

*Required if you use JOBNAME(DELETED) or DSNAME('DELETED') in any data set name VRsEs.*

Before z/OS V1R10, DFSMSrmm did not consider the disposition of a tape data set. If a job step abended, and CLOSE was entered because of abend processing, DFSMSrmm tracked this in the data set record and provided VRS management by using the ABEND VRsEs.

Starting in z/OS V1R10, DFSMSrmm considers the normal disposition for a tape data set. All tape data sets, whether using system temporary data set names or not, are considered. If the normal disposition (as coded in the JCL or dynamic allocation) is DELETE, the data set record is updated to track the data set as “deleted” by disposition processing. This is a flag just like the “abend” flag. VRSEL processing is updated to process the new restricted “DELETED” VRsEs.

**Migration action:** If you use JOBNAME(DELETED) in any data set name VRsEs, you must change the VRS to use JOBNAME(DELETED\*) to avoid a conflict with the new 'DELETED' VRS support.

If you use DSNAME('DELETED') in any data set name VRsEs, you must change the VRS to use DSNAME('DELETED\*') to avoid a conflict with the new 'DELETED' VRS support.

**Tip:** Use IBM Health Checker for z/OS to help with this migration action. The relevant check is named ZOSMIGV1R10\_RMM\_VRS\_DELETED and is available through APAR OA26608. The check ensures that you do not have any VRsEs that conflict with the new 'DELETED' VRS support.

### **DFSMSrmm: Use the backward-compatible CIM provider (Required-IF, as of R10)**

*Required if you use a CIM client with the DFSMSrmm CIM agent and decide not to update your code to handle the new key formats.*

The keys used for the DFSMSrmm CIM classes have changed with z/OS V1R10. If you have CIM clients that have been working on z/OS V1R9 or z/OS V1R8 and you want to use them against z/OS V1R10, and the clients have some hard-coded or specific processing of the keys of DFSMSrmm CIM classes, you can do either of the following:

- Use the backward-compatible CIM provider.
- Update your code to handle the new key formats and then use the z/OS V1R10 CIM provider. If you choose this path, see migration action “DFSMSrmm: Replace CIM providers and CIM classes”.

The following table shows the old keys of DFSMSrmm CIM classes and the new compound keys, which have formats of concatenated strings containing the values of the old keys delimited with “+” and appended with spaces by the fix length if needed.

Class Name	Key Transformation
IBMRMM_Dataset	DataSetName,PhysicalFileSequenceNumber, VolumeSerialNumber, and CdsID keys have been replaced with Name key of format DataSetName+FileSeq+Volser+CdsID.
IBMRMM_Location	LocationName, LocationType, and CdsID keys have been replaced with Tag key of format LocationName+LocationType+CdsID.
IBMRMM_LogicalVolume	VolumeSerialNumber and CdsID keys have been replaced with DeviceID key of format Volser+CdsID.
IBMRMM_Owner	OwnerId and CdsID keys have been replaced with InstanceID key of format OrgID:OwnerID+CdsID.
IBMRMM_PhysicalVolume	VolumeSerialNumber and CdsID keys have been replaced with Tag key of format Volser+CdsID.
IBMRMM_PolicyRule	PolicyRuleType, PolicyRuleName, JobNameMask, and CdsID keys have been replaced with PolicyRuleName key of format RuleType+RuleName+JobNameMask+CdsID.
IBMRMM_Product	ProductNumber and CdsID keys have been replaced with IdentifyingNumber key of format ProductNumber+CdsID.
IBMRMM_ShelfLocation	LocationName, MediaName, ShelfLocationNumber, and CdsID keys have been replaced with Tag key of format LocationName+MediaName+Number+CdsID.

Starting in z/OS V1R10, in order to simplify processing of the compound keys, every changed class now has the KeyWithCdsIdName attribute containing the name of its compound key. Additionally, xxxFormat and xxxMask attributes are provided, where xxx is the name of the compound key. For example, TagFormat property of

IBMRMM\_PhysicalVolume is set to "Volser+CdsID" and TagMask string contains the consecutive concatenation of six blanks, symbol "+", and eight blanks.

**Migration action:** The backward-compatible CIM provider is provided as rmmcim19.tar.Z compressed tar archive within the /usr/lpp/dfsms/rmm directory. To use it, go to the Miscellaneous Functions panel provided by the rmmutil.sh script and invoke "Unpack RMM CIM Agent for V1R9". As a result, you will have the old agent unpacked in the "v1r9" path within the current directory. For example, if you use the /usr/lpp/dfsms/rmm deployment directory, the new agent will reside in /usr/lpp/dfsms/rmm/v1r9.

## **DFSMS Migration Actions Post-First IPL**

### **DFSMSdfp: Run OAM DB2 BIND jobs (Required-IF)**

*Required if you use OAM object support.*

When migrating to any new release of z/OS, you must run OAM DB2 BIND jobs if you are using OAM for object support. The BIND jobs update DB2 with new OAM DB2 code.

**Migration action:** Run the BIND jobs appropriate to your installation:

1. Update and execute the samplib job CBRPBIND (OAM DB2 Bind Package Job).
2. Do one of the following:
  - If your installation starts OAM, uses optical or tape devices, or uses the OAM storage management component (OSMC), do the following:
    - Update and execute samplib job CBRABIND (OAM DB2 Application Plan Bind for LCS and OSR).
    - Update and execute samplib job CBRHBIND (OAM DB2 Application Plan Bind for OSMC).
  - If your installation does not start OAM, use optical or tape devices, or use OSMC, update and execute samplib job CBRIBIND (OAM DB2 Application Plan Bind for OSR only).
3. For more information, see the topic "Migrating, Installing, and Customizing OAM" in *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Object Support*.

**Note:** The following list itemizes the changes to OAM BIND jobs in z/OS V1R10. If you choose to edit a previous version, you must incorporate any new changes as described in the header of each new samplib BIND job:

- CBRPBIND: CBRIEDBS added
- CBRABIND: CBRIEDBS added to CBRIDBS plan
- CBRHBIND: no changes
- CBRIBIND: CBRIEDBS added to CBRIDBS plan.

### **DFSMSdfp: Accommodate the change of IDCAMS DEFINE for RECORDS involving VSAM linear data sets (Required-IF, as of R11)**

*Required if you define a linear data set with CISIZE greater than 4096 bytes and specify the RECORDS parameter.*

Before z/OS V1R11, when space for linear data sets was calculated in response to an AMS DEFINE command with the RECORDS parameter, an implied control interval size of 4096 was used when the control interval size was larger than 4096. As of z/OS V1R11, in the same scenario, the actual (larger) control interval size is used instead of 4096.

**Note:** This migration action was previously documented, incorrectly, as being introduced by APAR OA21369 on z/OS V1R9. The correct APAR is FIN APAR OA25988 and the correct release is z/OS V1R11.

**Migration action:** Consider adjusting the value of the RECORDS specification when a linear data set is defined with CISIZE greater than 4096 bytes. If you continue to use the same IDCAMS DEFINE CLUSTER job in such a situation, the allocated space will be larger in z/OS V1R11 because the actual CISIZE (instead of 4096 bytes) has an influence on the amount of space allocated to the data set. You need to decrease the value of the RECORDS specification to achieve an allocation amount similar to before migration.

#### **Notes:**

1. DB2 allocates VSAM linear data sets using the RECORDS parameter, so DB2 allocations would be affected by this change when CISIZE greater than 4 KB is used. To avoid this side effect, apply the PTF for DB2 APAR PK42840. With this PTF, the define and extend code in DB2 has been changed to use KB when issuing the IDCAMS DEFINE command.
2. In z/OS V1R8, SMS added a new code to validate the space quantity for volume selection. The space calculation is based on the number of records multiplied by the actual CISIZE instead of the 4 KB record size. If SMS fails the allocation because the primary quantity requested is larger than the total capacity of the largest available volume, the following error message is issued: IGD17279I (n) VOLUMES WERE REJECTED BECAUSE OF INSUFF TOTAL SPACE.

**Reference information:** For details about the change, see APAR OA25988.

**DFSMSdcpf: Evaluate applications for sensitivity to changes to IDCAMS LISTCAT command output (Required-IF, as of R11)**

*Required if your application depends on the output generated by the IDCAMS LISTCAT ALL command.*

Starting with z/OS V1R11, the IDCAMS LISTCAT command output is enhanced to include a new EATTR field indicating whether or not a VSAM data set can be defined with extended attribute DSCBs.

**Migration action:** Output from the IDCAMS LISTCAT command is not an intended programming interface. Evaluate applications that parse LISTCAT output to ensure that either there is no dependency on the order of the fields in the LISTCAT output or that they can accommodate the new EATTR fields in the output. To reduce future impact and maintenance, IBM suggests migrating the parsing routines or applications to use an IBM provided programming interface such as the Catalog Search Interface. See Catalog Search Interface User's Guide in z/OS DFSMS Managing Catalogs. The following example shows the EATTR field in LISTCAT ALL output:

```
LISTCAT LEVEL(ZTAMS002) ALL
NONVSAM ----- ZTAMS002.COMPRS.DATA01
IN-CAT --- SYS1.MVSRES.MASTCAT
HISTORY
DATASET-OWNER----- (NULL) CREATION-----2004.297
RELEASE-----2 EXPIRATION-----0000.000
ACCOUNT-INFO----- (NULL)
EATTR----- (OPT)
SMSDATA
STORAGECLASS ---S1P03S12 MANAGEMENTCLASS--- (NULL)
DATACLASS -----SRX00001 LBACKUP ---0000.000.0000
VOLUMES
VOLSER-----1P0301 DEVTYP-----X'3010200F'
```

**DFSMSdcpf: Avoid errors from BDAM READ macros issued with a BSAM or BPAM DCB (Required-IF, as of R10)**

*Required if a program opened a BSAM or BPAM DCB and issues a BDAM READ macro.*

Prior to z/OS V1R10, it was possible for a program to open a BSAM or BPAM DCB and erroneously issue a BDAM READ macro without failing. In z/OS V1R10 or later, the access method takes that macro invocation as an SF64 request and the probable result will be a 002-EE ABEND and an IEC036I message.

**Migration action:**

- If you find an ABEND 002-EE or message IEC036I for reason code EE, check whether the application program is issuing a BDAM READ macro for a BSAM or BPAM DCB. Also check whether the application program is issuing a BISAM WRITE macro. In earlier releases the system did not detect these user errors.
- If the second operand on the READ macro is SF or SB, it is for BSAM. If the second operand on the READ macro is DI or DK, it is for BDAM. If the second operand on a WRITE macro begins with "K", it is for ISAM. A BDAM or BISAM macro issued with a BSAM or BPAM DCB is a user error that was not detected in earlier releases of the system.
- If the DCB has DSORG=PS, it is BSAM. If the DCB has DSORG=PO, it is BPAM. If the DCB has DSORG=DA, it is BDAM. If the DCB has DSORG=IS, it is BISAM.
- Change the READ or WRITE macro to use the documented BSAM or BPAM syntax.

**DFSMSdcpf: Examine the job and hardcopy logs for important Delete/Rename and storage constraint messages (Required-IF as of R10)**

*Required if these messages are not reaching a storage administrator now.*

Before z/OS V1R10, some important DFSMS messages could be suppressed or not issued by some components of the system. An example of such a message is IGD17054I DATA SET NOT FOUND FOR DELETE/RENAME ON VOLUME xxxxxx DATA SET IS aaaa.bbbb.cccc.

In z/OS V1R10, DFSMS writes a number of these existing messages to the job log and hardcopy log for the first time. These include messages for Delete and Rename processing, as well as messages pertaining to storage constraints that need to be addressed to prevent possible storage-related failures. This change is made to help ensure that these important messages reach the attention of a storage administrator.

**Note:** This change does not alter the functioning of the system in any way beyond externalizing messages that may have been suppressed. Also, it applies to messages for SMS-managed data sets as well as non-SMS-managed data sets.

**Migration action:** On a regular basis, just as frequently as you would today, examine the job and hardcopy logs for important messages related to Delete, Rename, and storage constraint processing.

**DFSMSdftp: Update programs or procedures that depend on output in reports from the LISTDATA PINNED command (Required-IF as of R10)**

*Required if users or programs depend on output from the LISTDATA PINNED command.*

The track addresses in the output of the IDCAMS LISTDATA PINNED command are in a different format, regardless of the volume size. Before z/OS V1R10, the report listed each pinned track and its associated data set name.

Beginning with z/OS V1R10, the report identifies a range of pinned tracks associated with a data set. Each track address is printed in eight hex digits in the native format of the device, using the form CCCCcccH. In addition, instead of one track being listed per line, consecutive tracks for each data set are gathered into one line and the range of track addresses is shown.

The meaning of "CCCCcccH" is as follows: "CCCC" is the low-order 16 bits of the cylinder address, "ccc" is the high-order 12 bits of the cylinder address, and "H" is the 4-bit track address. When a 28-bit cylinder address is greater than 65535 (X'FFFF'), the F column in the report contains an asterisk.

**Migration action:** SYSPRINT and TSO terminal output from IDCAMS are not intended programming interfaces but the changes are stated here as a convenience to help you identify any changes to a program that uses this output. Change any affected program appropriately to parse the new format.

**DFSMSdftp: Remove or correct multivolume VSAM data sets whose volumes are out of order (Required-IF, as of R10, and OA24116 on R9, and OA16372 on R8)**

*Required if you have affected VSAM multivolume data sets.*

Beginning with z/OS V1R10 (and APAR OA24116 on z/OS V1R9, and APAR OA16372 on z/OS V1R8 and z/OS V1R7), non-RLS VSAM OPEN detects volumes out of order. Informational message IEC161I 071(005)-026 or IEC161I 071(006)-026, indicating extents out of sequence because volumes are out of order, is issued for a multivolume VSAM data set in this situation. Despite the volumes being out of order, OPEN processing continues successfully with ACBERFLG=00 and RC=00.

**Note:** Out-of-sequence volumes could have existed before z/OS V1R7, but at that time DFSMSdftp did not expose this situation.

**Migration action:** Correct the out-of-sequence volumes indicated by message IEC161I. Issue IDCAMS DELETE NOSCRATCH and DEFINE RECATALOG commands with the volumes specified in the correct order. Issue a LISTCAT command for the data set before doing the DELETE so that you can determine the correct order for the volumes.

**Tip:** Informational APAR II14289 documents an as-is tool (VOLSRCHT) that can help identify multivolume VSAM data sets where the volumes are out of order in the catalog. The tool scans specified catalogs and provides the correct volume sequence for data sets.

**DFSMSdss: Use the new version of the DEFrag command (Recommended, as of R10)**

*Recommended for improved efficiency, improved debugging, and compatibility with extended address volumes.*

Before z/OS V1R10, the DFSMSdss DEFrag command contained a CONSOLIDATE keyword. In z/OS V1R10, a new version of the DEFrag command is added that is compatible with extended address volumes. This new version of the command drops the CONSOLIDATE keyword and adds other keywords to improve efficiency and debugging. To replace the CONSOLIDATE keyword, a new DFSMSdss CONSOLIDATE command is added for use in conjunction with the new version of the DEFrag command. This new CONSOLIDATE command allows the consolidate function to be used without any additional DEFrag processing of a volume.

The pre-z/OS V1R10 version of the DEFrag command with its CONSOLIDATE keyword is still available if you specify it with the new VERSION1 keyword. Note that the previous version is not compatible with extended address volumes.

**Migration action:**

- To use the z/OS V1R10 version of the DEFrag command, specify the DEFrag command without specifying the CONSOLIDATE keyword, and use the new, separate CONSOLIDATE command. Using both the CONSOLIDATE and DEFrag commands consecutively in the same job will give the same information and do essentially the same thing as the old DEFrag command with the CONSOLIDATE keyword. But in addition, the new commands are compatible with extended address volumes, you can use the new MAXTIME and MMOVPCT keywords to limit the amount of time a job runs, and the new CONSOLIDATE command reduces the number of extents of a data set as much as possible rather than only when the extents can be reduced to one.

To use the z/OS V1R10 CONSOLIDATE command, you must have access to the STGADMIN.ADR.CONSolid profile in the FACILITY class. (Using the z/OS V1R10 DEFrag command continues to require access to the existing STGADMIN.ADR.DEFRAG profile.)

- To use the pre-z/OS V1R10 version of the DEFrag command with its CONSOLIDATE parameter intact, specify the new VERSION1 keyword. Note that this previous version is not compatible with extended address volumes, and data sets will only have extent reduction when they can be reduced to one extent. Existing (created on z/OS V1R9 or earlier) JCL that specifies the DEFrag command with the CONSOLIDATE keyword on z/OS V1R10 (without VERSION1) will cause the CONSOLIDATE command to be executed first, followed by the DEFrag command using the new version of DEFrag.

#### **DFSMSHsm: Define facility class names to protect (H)BACKDS RETAIN DAYS commands (Req-IF as of R11)**

*Required if you use the BACKDS RETAIN DAYS and HBACKDS RETAIN DAYS commands.*

Starting in z/OS V1R11, new DFSMSHsm facility class names protect the BACKDS RETAIN DAYS and HBACKDS RETAIN DAYS commands. The new facility class names must be defined to protect from unintended use of this support by unauthorized users. This support is only available on z/OS V1R11. Earlier systems will be able to recover backup versions created with the RETAIN DAYS keyword.

The RETAIN DAYS keyword is added to the DFSMSHsm (H)BACKDS command. You can use the RETAIN DAYS keyword to specify the number of days to retain a specific backup version of a data set.

**Migration action:** Define new facility class name STGADMIN.ARC.BACKDS.RETAIN DAYS to authorize any BACKDS RETAIN DAYS command and STGADMIN.ARC.ENDUSER.HBACKDS.RETAIN DAYS to authorize any HBACKDS RETAIN DAYS command.

#### **DFSMSHsm: Accommodate the change in ARCBDEX exit (Req-IF as of R11 APAR OA28948 and R10 APAR OA28136)**

*Required if you want to allow command backup but disallow autobackup for some data sets.*

Prior to DFSMSHsm z/OS V1R10, ARCBDEX was called during volume-level backup operations (AUTOBACKUP for automatic and BACKVOL for command-based operations). Users examined information in the exit's input data structure to determine whether to allow or disallow backup of a data set. Starting with DFSMSHsm z/OS V1R10, ARCBDEX is called during individual data set backup operations, as well as volume-level backup operations. When called for individual data set backup, the exit's input data structure differs, but the return code values and meanings remain the same as for volume-level backups. Users must examine the level information in the input data structure, at offset x'04', to determine whether ARCBDEX is being invoked for volume-level or for individual data set backup.

**Migration action:** Examine the level information in the input data structure at offset x'04' to determine whether ARCBDEX is being invoked for volume-level or for individual data set backup. The level field will contain '\*EXPAND1' for volume-level requests and '\*EXPAND2' for individual data set backup requests. For individual data set backup requests, the field at offset x'0C' contains additional status information about whether the backup request is the result of a backup data set command or the result of retry from a volume command.

#### **DFSMSHsm: Define facility class names to protect (H)BACKDS NEWNAME commands (Req-IF as of R10)**

*Required if you use the BACKDS NEWNAME and HBACKDS NEWNAME commands.*

Starting in z/OS V1R10, new DFSMSHsm facility class names protect the BACKDS NEWNAME and HBACKDS NEWNAME commands. DFSMSHsm also checks for RACF authorization to the data set name specified as the NEWNAME parameter. The new facility class names must be defined to protect from unintended use of this support by unauthorized users. This support is only available on z/OS V1R10. Earlier-level systems will be able to recover backup versions created with the NEWNAME keyword.

The NEWNAME keyword is added to the DFSMSHsm (H)BACKDS command. When users specify the (H)BACKDS NEWNAME command, DFSMSHsm creates a backup version of the data set specified on the command and makes the backup version look like it is a backup of the data set specified with the NEWNAME keyword.

**Migration action:** Define new facility class name STGADMIN.ARC.BACKDS.NEWNAME to authorize any BACKDS NEWNAME command and STGADMIN.ARC.ENDUSER.HBACKDS.NEWNAME to authorize any HBACKDS NEWNAME command.

#### **DFSMSHsm: Use the generic rather than the esoteric unit name for duplex generated tape copies (Required-IF as of R10)**

*Required if you need to use the generic unit name.*

Before z/OS V1R10, when a duplex failure generated a tape copy, DFSMSHsm passed the generic equivalent of a SETSYS unit name (such as 3590-1) to the tape copy. Beginning with z/OS V1R10, DFSMSHsm instead passes a SETSYS-specified unit name (such as SETSYS MIGUNITNAME (*esoteric-unitname*)).

This change ensures that the same output devices and unit name that are targeted for the failed duplex copy are targeted for the tape copy. This change affects the DFSMSHsm backup, migration, and recycle functions. The ABARS function is not changed.

**Migration action:** To change the behavior back to always passing the generic unit for the tape copy, use the DFSMSHsm PATCH command to set the MCVTF\_GENERIC\_TCN\_UNIT flag on. For example:

```
PATCH.MCVT.+196 BITS(...1...)
```

To return to the default behavior of using the SETSYS specified unit name, use the following command:

```
PATCH.MCVT.+196 BITS(...0...)
```



### **DFSMSrmm: Replace DFSMSrmm CIM providers and CIM classes (Required-IF as of R10)**

*Required if migrating from z/OS V1R9: Yes, if you have chosen to use the z/OS V1R10 CIM provider (not the backward-compatible CIM provider) as described in migration action “DFSMSrmm: Use the backward-compatible CIM provider”.*

*Required if migrating from z/OS V1R8: Yes, if you have chosen to use the z/OS V1R10 CIM provider (not the backward-compatible CIM provider) as described in migration action “DFSMSrmm: Use the backward-compatible CIM provider”, or if you previously set up a CIM server and agent for DFSMSrmm and registered the DFSMSrmm CIM classes.*

In z/OS V1R10, the keys used for the DFSMSrmm CIM classes were changed. In z/OS V1R9, DFSMSrmm CIM providers and CIM classes were changed. You must unregister the previous CIM providers and CIM classes, and register the z/OS V1R10 CIM providers and CIM classes.

#### **Migration action:**

1. If you are planning to use the z/OS V1R10 CIM provider rather than the backward-compatible CIM provider, update your code to handle the new key formats in z/OS V1R10. For the new formats, see migration action “DFSMSrmm: Use the backward-compatible CIM provider”.
2. Using the rmmutil.sh tool, unregister all the pre-z/OS V1R10 CIM providers and unload all the pre-z/OS V1R10 CIM classes.
3. Using the same rmmutil.sh tool, register the complete set of z/OS V1R10 CIM providers and load the z/OS V1R10 CIM classes.

### **DFSMSrmm: Modify programs that process the scratch list reports (SCRLIST and NEWSOCR) (Required-IF as of R10, and OA23142 on R9 and R8)**


*Required if your program depends on the DCB information in scratch lists (SCRLIST or NEWSOCR).*

When DFSMSrmm utility program EDGRPTD creates scratch list reports, the output data set DCBs for SCRLIST and NEWSOCR reports have the hardcoded block size equal to the logical record size (LRECL=133, BLKSIZE=133, RECFM=FBA.) This causes the data sets to use more DASD space than they really need. For instance, if the data set for the SCRLIST report using 781 tracks on a 3390 volume is copied with IEBGENER to another data set with BLKSIZE=0, the resulting data set uses 134 tracks on the 3390 — an 83% saving of DASD space.

As of z/OS V1R10 (and APAR OA23142 on z/OS V1R9, V1R8, and V1R7), to make the most of system determined blocksize (SDB) functionality, the utility program EDGRPTD has been changed to remove BLKSIZE and change LRECL from 133 to 132 for the SCRLIST and NEWSOCR reports.

**Migration action:** After the PTF for APAR OA23142 is installed, the DCB information for SCRLIST and NEWSOCR is changed. For example, the new DCB will have LRECL=132, BLKSIZE=27984, and RECFM=FBA. Therefore, if you have any programs that process the output of SCRLIST or NEWSOCR, you need to remove the hardcoded LRECL of 133.





### DFS Migration Actions for z/OS R11 from z/OS R9

- **Migration Actions You Can Do NOW:**
  - Ensure that `sysplex_admin_level=2` is available on all systems in a shared file system environment (**Required-IF, as of R11**)
    - Before starting R11 zFS you **MUST** specify this level in your IOEFSPRM file on your R9 and R10 zFS systems, otherwise R11 zFS will **not** initialize in a shared environment. This level is found in OA25026 (and an IPL with this APAR will put you at level 1 by default).
      - First rolling IPL: OA25026 (level 1). Second rolling IPL: level 2. Then ready for R11!
- **Migration Actions Pre-First IPL:**
  - **Modify programs or procedures that issue an unmount of zFS file systems (Required-IF, as of R11)**
    - Prior to `sysplex_admin_level=2`: zFS unmount would succeed on a quiesced, growing, or cloning zFS file system.
    - As of `sysplex_admin_level=2`: zFS unmount fails on a quiesced, growing, or cloning zFS file system, unless unmounted with force.
    - Handle the failure on an unmount, or issue unmount with the force option.
  - **Permit the SMB server ID to BPX.DAEMON and BPX.SERVER (Required, as of R11)**
    - As of R11, SMB server must have READ access to these FACILITY class profiles. Otherwise SMB startup will fail.

36 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

### DFS Migration Actions Between z/OS V1 R9 and z/OS V1 R11

These migration actions were taken from *z/OS Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### DFS Migration Actions You Can Do Now



#### **zFS: Discontinue use of multi-file system aggregates (Required-IF, as of R11)**

*Required if you are using zFS multi-file system aggregates.*

Beginning in z/OS V1R11, you can no longer attach zFS multi-file system aggregates that are shared across systems in a sysplex. IBM has previously recommended that multi-file system aggregates not be shared in a sysplex environment. Any attempts to attach zFS multi-file system aggregates will fail in a z/OS UNIX shared file system environment with message IOEZ00046E. Attaching zFS compatibility mode aggregates, which have a single file system per data set, will continue to be supported in all environments. In a future release, IBM plans to withdraw support for zFS multi-file system aggregates. When this support is withdrawn, only zFS compatibility mode aggregates will be supported.

**Migration action:** Determine whether you are using zFS multi-file system aggregates by scanning your zFS IOEFSPRM configuration options file for `define_aggr` statements. You can also scan your `/etc/rc` file for any `zfsadm attach` commands. If you are using zFS multi-file system aggregates, copy the data from each file system into its own zFS compatibility mode aggregate.

**Tip:** Use IBM Health Checker for z/OS check ZOSMIGV1R11\_ZFS\_RM\_MULTIFS (or ZOSMIGREC\_ZFS\_RM\_MULTIFS in a single-system or monoplex environment) to determine whether you have any multi-file system aggregates attached on your system.



#### **zFS: Ensure that `sysplex_admin_level=2` is available on all systems in a shared file system environment (Required-IF, as of R11)**

*Required if you have a shared file system environment with more than one system in that environment.*

z/OS V1R11 adds a new XCF protocol between sysplex members for zFS administration functions. The new protocol requires the installation of a coexistence PTF on coexisting systems so that zFS can function properly in z/OS V1R11 in a shared file system environment.

If you have a problem on a z/OS V1R9 or z/OS V1R10 system running with zFS `sysplex_admin_level=2`, you can remove the specification (or specify `sysplex_admin_level=1`, which is equivalent to the default) and perform a rolling IPL or restart zFS on each system. This cannot be done after zFS on a z/OS V1R11 system has joined the sysplex. Also, if you try to start zFS in z/OS V1R11 on another system after you have backed out to `sysplex_admin_level=1` on zFS V1R9 or z/OS V1R10, zFS on z/OS V1R11 will not come up because it requires all other systems to be at zFS `sysplex_admin_level=2`. If you try to bring in zFS z/OS V1R11 when `sysplex_admin_level=2` is not active on all systems, the message "IOEZ00614A zFS has detected an incompatible interface level *IntLevel* for member *Sysname* " will appear. (*Intlevel* will be 1 if APAR OA25026 is installed; 0 if it is not.) If you must go back to z/OS V1R8, remove z/OS V1R11 from the sysplex and restart any systems that are running zFS `sysplex_admin_level=2` by removing the `sysplex_admin_level` specification (or by specifying `sysplex_admin_level=1`) and performing a rolling IPL or restart zFS on each of these systems. (The `sysplex_admin_level` specification is ignored in z/OS V1R8 and z/OS V1R11.)

The new zFS toleration support uses additional ENQs to keep track of zFS ownership of zFS aggregates. The new ENQs all have a qname of SYSIOEZ. Also, the behavior of unmount changes for anyzFS file systems that are owned on a system running with `sysplex_admin_level=2`.

#### Migration action:

1. Install the PTF for APAR OA25026 on all z/OS V1R9 and V1R10 systems. This is a conditioning function for zFS in z/OS V1R11. Make the PTF available on all systems through a rolling IPL. You are now running with zFS `sysplex_admin_level=1`.
2. Specify the `sysplex_admin_level=2` configuration option in your IOEFSPRM file and make this level available on all z/OS V1R9 and V1R10 zFS systems through another rolling IPL or by restarting zFS. This allows zFS on z/OS V1R9 and z/OS V1R10 to tolerate running with zFS on z/OS V1R11. (The default for `sysplex_admin_level` is `sysplex_admin_level=1`).

#### Notes:

1. You cannot specify the `sysplex_admin_level` configuration option dynamically through the **zfsadm config** command.
2. You cannot skip Step 1 and do only Step 2 in "Steps to take" above if you are running z/OS V1R9 or V1R10 on any systems that do not have the PTF applied and are activated by IPL.

When you configure `sysplex_admin_level=1` or 2, the MODIFY ZFS,QUERY,LEVEL operator command returns a new line, at the end of the output, that shows the sysplex admin interface level. The following example shows z/OS V1R10 displaying interface level 1:

```
f zfs,query,level
IOEZ00639I zFS kernel: z/OS zSeries File System
Version 01.10.00 Service Level OAnnnnn - HZFS3A0.
Created on Tue Jan 6 20:02:25 EST 2009.
sysplex(admin-only) interface(1)
```

When you run zFS on z/OS V1R11, the `sysplex_admin_level` option is ignored and the MODIFY ZFS,QUERY,LEVEL operator command displays interface level 3. When zFS initializes in a shared file system environment, the following message is displayed (as of z/OS V1R10): IOEZ00617I zFS is running sysplex admin-only with interface level 1. This is normal. It shows the sysplex admin interface level that zFS is running.

**Tip:** Use IBM Health Checker for z/OS to help determine whether you are running zFS at the correct interface level. You need level 2 in order to start zFS on z/OS V1R11. The check is available through APAR OA27198 and is named IBMZFS,ZOSMIGV1R11\_ZFS\_INTERFACELEVEL.

#### **zFS: Stop using zfsadm setquota of compatibility mode aggregates (Required-IF, as of R10)**

*Required if you normally use the zfsadm setquota command for compatibility mode aggregates.*

Before z/OS V1R10, using the zfsadm setquota command on a compatibility mode aggregate was discouraged but accepted. Starting in z/OS V1R10, if you attempt to set the file system quota for a compatibility mode aggregate, zFS issues return code 121 (EINVAL) and reason code EFxx69B8. (The maximum size of the file system is known as its **quota**.)

**Migration action:** Do not use the zfsadm setquota command for compatibility mode aggregates.

### **DFS Migration Actions Pre-First IPL**

**zFS: Modify programs, execs, or procedures that issue an unmount of zFS file systems (Required-IF, as of R11, and APAR OA25026 with sysplex\_admin\_level=2)**

*Required if you have programs that assume that unmount is successful.*

Beginning with z/OS V1R11, and for z/OS V1R9 and V1R10 systems that are running zFS with `sysplex_admin_level=2`, zFS fails an unmount of a quiesced, growing, or cloning zFS file system unless unmount force is specified. In previous releases (without `sysplex_admin_level=2`), zFS successfully returned in these cases. Here is an example of the failure message that is issued when a file system is quiesced:

```
unmount filesystem(plex.jms.aggr006.lids0006)
RETURN CODE 00000072, REASON CODE EF0969B5. THE UNMOUNT FAILED FOR FILE SYSTEM
PLEX.JMS.AGGR006.LDS0006.
```

Return code X'72' indicates that the file system is busy with an administration command.

**Migration action:** Check your programs, execs and procedures to see if they can handle a failure on unmount. If they cannot, modify them to handle a possible unmount failure or to issue unmount with the FORCE option.

**Note:** An unmount force can cause a grow or clone to be terminated, causing a partial grow or a partial clone. The partial grow should not cause a mount failure. The grow can be issued again and completed. A partial clone must be deleted and issued again. An unmount force of a quiesced aggregate could cause a backup to fail if the quiesce was issued by the backup program. (It will fail when it tries to issue the unquiesce.)

**zFS: Permit the SMB server user ID to BPX.DAEMON and BPX.SERVER (Required, as of R11)**

In z/OS V1R11, the SMB server requires access to BPX.DAEMON and BPX.SERVER. The new settings provide greater security.

**Migration action:** Define the BPX.SERVER and BPX.DAEMON FACILITY classes in RACF (if not already defined) using the following steps:

1. Create the BPX.SERVER and BPX.DAEMON FACILITY class profile:
 

```
RDEFINE FACILITY BPX.SERVER UACC(NONE)
RDEFINE FACILITY BPX.DAEMON UACC(NONE)
SETROPTS RACLIST(FACILITY) REFRESH
```
2. Permit the DFS user ID to the BPX.SERVER and BPX.DAEMON RACF FACILITY class profiles:
 

```
PERMIT BPX.SERVER CLASS(FACILITY) ACCESS(READ) ID(DFS)
PERMIT BPX.DAEMON CLASS(FACILITY) ACCESS(READ) ID(DFS)
SETROPTS RACLIST(FACILITY) REFRESH
```
3. Follow your installation's procedure for allowing the SMB daemon control task (IOEPDCT) to run if your system is set up for RACF program control protection. If your system is not set up for RACF program control protection, no action is necessary.

If this security setup is not done prior to SMB startup, the following message occurs: IOEN00514A ThisProgram: Not permitted to BPX.DAEMON. Severity: svc\_c\_sev\_fatal.


**DFS Migration Actions Post-First IPL****Fix syntax errors now reported during SMB startup (Recommended, as of R10)**

*Recommended to avoid getting SMB startup errors and to ensure that SMB uses the configuration options you intended.*

Before z/OS V1R10, when you started SMB, the syntax of SMB environment variables was not validated. Beginning with z/OS V1R10, when you start SMB, a new command automatically validates the syntax of the `/opt/dfslocal/home/dfskern/envar` file or another environment variable file you specify. SMB starts as usual, but if errors are not corrected, SMB might not start with the intended parameters.

**Migration action:** Before starting SMB, run the `dfssyntax` command against the file `/opt/dfslocal/home/dfskern/envar` or another environment variable file that your installation uses, and then correct any syntax errors that are reported.

## JES2 Migration Actions for z/OS R11 from z/OS R9




---

**Migration Actions Pre-First IPL:**

- **Update uses of JES2 macros and exit routines (Required-IF, as of R11)**
  - In R11, a variable extension is added to the JOE data area. This extension is stored in the BERT data area in the JES2 checkpoint. Use \$DOGJOE to manage JOE data areas.
  - Various macros changes as a result of the \$DOGJOE service. Some macro interfaces have changed, some rules for using these macros have changed. These changes are required regardless of JES2 checkpoint level (z2 or z11).
  - Refer to *z/OS R11 Migration*, *JES2 Macros*, and *JES2 Installation Exits* for more details.

**Migration Actions Post-First IPL:**

- **Migrate to z11 mode (Recommended, as of R11)**

  1. Determine your z11 checkpoint activation readiness with the \$D ACTIVATE command. This command indicates if activation to z11 mode will succeed.
    - Review your current utilization of BERT data to determine if there are sufficient BERTS. Additional BERTs are needed for each SYSOUT data set that has transaction data associated with it.
  2. Run the JES2 \$ACTIVATE command to activate z11 mode following the considerations for this command found in *z/OS JES2 Commands*.

37
z/OS 1.11: Migration - ...And It's Good! Part 2 of 2
© 2010 IBM Corporation

### JES2 Actions Between z/OS V1 R9 and z/OS V1 R11

These migration actions were taken from *z/OS Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### JES2 Migration Actions You Can Do Now

<none>

### JES2 Migration Actions Pre-First IPL

#### Update uses of JES2 macros and exits (Required-IF, as of R11)

*Required if using various JES2 macros or exits mentioned in z/OS Migration.*

With z/OS V1R11 enhanced transaction SYSOUT selection for the SYSOUT API and extended status SSIs, a variable extension is added to the JOE data area. This extension is stored in the BERT data area in the JES2 checkpoint. To manage this data area in a release independent manner, a check out and check in service (\$DOGJOE) is available to manage JOE data areas. This service must be used when JOE data areas are being modified. It is also recommended that this service be used to examine fields in a JOE data area. Installation exits that update JOE data areas need to be updated to use the \$DOGJOE service to check out (fetch) an update mode JOA. The checked out JOA (also referred to as an artificial JOE) is a composite of the work JOE, the characteristics (char) JOE, and the JOE extension in the BERTs. The \$DOGJOE service also insulates the exit code from the current \$ACTIVATE level of JES2.

Various z/OS V1R11 JES2 macros changed as a result of the \$DOGJOE service. Certain macro interfaces now require a JOA instead of a work JOE or characteristics JOE. In other scenarios, the rules for using these macros have changed.

**Migration actions:** Before upgrading to z/OS V1R11 JES2, certain conditions must be met if your installation is using any of the following JES2 exits and/or macros. The actions described below must be considered regardless of checkpoint activation level.

- Be aware of the \$DOGJOE macro if you reference JOE fields in any of your exits. This new macro enables creation of read and update mode artificial JOEs termed JOAs.
- Determine which fields need to be replaced for specific exits. Before z/OS V1R11 JES2, a real work and sometimes a characteristics JOE were passed to JES2 exits 1,15,38,46, or 56. Starting with z/OS V1R11

JES2, an artificial JOE (JOA) will be passed to each of these exits. The table in *z/OS Migration* explains which fields need to be replaced for specific exits.

- Make required code updates if you use JES2 macros `##ADD`, `##ALCHK`, `##BLD` and `##BUSY`. In many cases, the interface has changed to require that a JOA be passed into the macro instead of a work JOE or the work/characteristics JOE combination. In the case of `##BUSY` and `##ALCHK`, additional rules must be followed.
- Be aware that the default for `##JOE` has been changed to return a read mode JOA. It is the responsibility of the user to return a JOA upon early exit from the `##JOE` processing loop. The `##JOE` macro returns a real JOE or a read mode JOA. Before z/OS V1R11 JES2, this macro only returned a real JOE.
- Change exit routines that examine field CBMTTR to examine field CBMQTR, which is now a 6-byte MQTR. Prior to z/OS V1R11 JES2, `##CBIO` used 4-byte MTTTRs internally to determine what record is being read or written. Starting with z/OS V1R11 JES2, `##CBIO` uses 6-byte MQTRs internally. Though not formally passed to exit 7 (`##CBIO` exit), some exits are known to locate the track address being processed.

### **Accommodate the removal of message \$HASP301 (Required-IF, as of R10, and as of R9 and R8 with OA22291)**

*Required if you depend on the message \$HASP301.*

JES2 is designed to prohibit the execution of multiple jobs with the same name, except in the case of started task (STC) jobs or time-sharing user (TSU) classes. A callable routine can be used to determine if the name of a candidate job duplicates the name of an already executing job. Prior to z/OS V1R8, when a job was being considered for selection, a scan of the entire queue was made to see if the jobname matched that of any already executing jobs. If so, the job was rejected and a flag was set in the JQE (JQE1HLD2).

JES2 processing for duplicate job name detection and blocking was changed in z/OS V1R8. Starting in z/OS V1R8, jobs with the same name are linked together. The linked group is called a duplicate job name family. The family is controlled by a structure called the DJB (duplicate job block). When a job is being considered for selection, an examination of the DJB is made to see if a member of the family is already executing. In this background, the flag JQE1HLD2 is no longer set on or off, which affects the Exit 14 code.

**Note:** The new duplicate job name processing is in effect as soon as all the JES2s sharing SPOOL are at system level z/OS V1R8 or later. If one or more members of the MAS are on z/OS V1R7 or earlier, the duplicate job processing is done the way it was done prior to the changes made in z/OS V1R8.

Beginning with z/OS V1R10 (and z/OS V1R9 and z/OS V1R8 with OA22291), when the default specification of `DUPL_JOB=DELAY` in JES2PARM JOBDEF statement is in effect, you no longer receive the message \$HASP301, even if a job is delayed for processing because a job with the same name is already running.

```
$HASP301 jobname - DUPLICATE JOB NAME - JOB DELAYED
```

#### **Migration action:**

- Review the impact of the message automation processing.
- Determine the impact by displaying the Status "DUP" on the SDSF I (INPUT QUEUE) panel and ST (STATUS) panel.
- Also, the command `$DJ(*),HOLD=DUP` is useful for this purpose. It displays the message \$HASP890 to show the STATUS=(AWAITING EXECUTION) and HOLD=(DUP) for the affected job.

### **Accommodate JES2 dynamic exits function (Required-IF, as of R10)**

*Required if you use JES2 exit routines or load installation load modules using the JES2 LOADMOD (or LOAD) initialization statement.*

The dynamic exits function of JES2, integrated in z/OS V1R10 and rolled back to prior releases by APARs, makes the following possible:

- Installation load modules (exit modules) can be dynamically loaded, refreshed, or deleted. Note that this function does not apply to modules that are part of the base JES2 product, meaning that you cannot use this function to apply IBM service.
- The list of routines associated with an exit can be dynamically altered.

Because of this new function, you should review and change (if necessary) any JES2 exit routines that use the following data areas and macros:

- The `##XRT` and `##XIT` data areas. The structure of the `##XRT` has been significantly altered. In addition, some fields in the `##XIT` have been moved to the `##XRT`. Reassembly of your installation code will, in general, fail if it references any fields in the `##XRT` or `##XIT` that have changed.
- The `##LMT` data area. Before z/OS V1R10, a `##LMT` can be active (LMT1INVD bit off) or invalid/free (LMT1INVD bit on). Beginning with z/OS V1R10, a `##LMT` can be in two new states: it can represent a

module that has been logically deleted (LMT2DELT bit on) or a module whose storage has been freed (LMT3FREE bit on). Code that examines \$LMTs might need to take these two new states into consideration.

- The \$MODLOAD and \$MODCHK macros. Message processing for these macros has changed. Beginning with z/OS V1R10, if MESSAGE=NO is coded or defaulted on the macro invocation, the message area returned in register 1 must be freed by the caller (using \$FREMAIN with SP=0, KEY=1, and LV=128). If the caller does not require the message to be returned, a new MESSAGE=SUPPRESS option can be used.

In addition, not all exit routines can tolerate the fact that installation load modules (ones loaded using the JES2 LOADMOD initialization statement) can now be deleted and refreshed (a new copy loaded and the old copy deleted) using an operator command. You can use the new DYNAMIC=NO option on the \$MODULE macro to prevent unwanted processing.

#### Migration action:

1. **Update exit routines that refer to the JES2 \$XRT and \$XIT data areas, if necessary.** The \$XIT is a 256-element array that represents all possible JES2 exits. \$XRTs are pointed to by the \$XIT. Each \$XRT is an array in which each entry represents a routine associated with a JES2 exit. Before z/OS V1R10, the \$XRT was a simple array of elements, each one representing a routine. With the dynamic exits function in z/OS V1R10, the \$XRT starts with a header that describes the list, followed by an array of XRTEs in which each entry represents a routine.

The change was needed to ensure predictability when routines are invoked by the JES2 \$EXIT facility.

When the \$EXIT facility is invoked, it selects the appropriate \$XRT and increments the use count in that \$XRT. The list of routines in that instance of the \$XRT are used by the \$EXIT facility. Operator changes to the list of routines while the \$EXIT facility is in control do not affect the currently active exit environments.

When the list of routines for an exit is altered, a new \$XRT is built and, if the existing one is in use, the new one is chained in ahead of the existing \$XRT and the old \$XRT is marked as logically deleted. When the old \$XRT is no longer in use, it is deleted.

Note that the first \$XRT associated with an exit (pointed to by XITXRTAD) can be logically deleted (XRTFDEL bit on) if there are no longer any routines associated with the exit. If there is a chain of \$XRTs, the second and beyond are always logically deleted. (Only the first \$XRT pointed to by XITXRTAD can be active.) One consequence of this change is that some fields that were in the \$XIT had to be moved to the \$XRT header. In particular, the fields XITBSPL (SPLEVEL check bypassed), XIT#RTNS (number of routines associated with this exit), and XITUSCNT (exit use count) have been moved to the \$XRT.

2. **Update exit routines that refer to the JES2 \$LMT data area, if necessary.** The \$LMT is used to track modules loaded by the JES2 \$MODLOAD service, including modules loaded by the LOADMOD initialization statement and the new \$ADD LOADMOD command. Code that searches \$LMTs needs to consider the state of the \$LMT before examining fields (such as the module address) in the \$LMT. Before z/OS V1R10, the only state needing examination was invalid/free (bit LMT1INVD on). However, the dynamic exits function in z/OS V1R10 introduces two new states that need to be considered: logically deleted \$LMTs and force freed \$LMTs.

Logically deleted \$LMTs (bit LMT2DELT on) represent modules that are not being used for new requests (they are logically deleted) but are in use by existing processing. These \$LMTs should not be examined when looking for the current version of a module but can be used if trying to resolve what module a particular address is located in.

Force freed \$LMTs (bit LMT3FREE on) represent load modules where the actual module is freed before JES2 is done using it. This can only occur if dynamic LPA services is used to delete the module from LPA (something IBM does not recommend doing). Code cannot access the actual load module because the storage is no longer valid. Force freed \$LMTs will eventually be logically deleted (LMT2DELT set on), although the LMT3FREE bit will always be on.

If an exit routine examines \$LMTs to find a particular routine address by name or to find the load module name associated with an address, there are JES2 services to do this that isolate the code from the internal JES2 data structures. In the JES2 environment, the LOCLMOD service obtains the module name associated with an address (including searching for MVS module names). Outside the JES2 environment, FINDLMOD performs the same function. To locate the address of a routine within a JES2 load module, you can use the LOCENTRY service. However, this is only available within the JES2 environment.

3. **Update exit routines that use the \$MODLOAD and \$MODCHK macros, if necessary.** Message processing for the \$MODLOAD and \$MODCHK macros has changed. Before z/OS V1R10, if MESSAGE=NO was coded or defaulted on the macro invocation, an error message could be returned to the caller in a data area pointed to by register 1. The caller could choose to use or ignore the returned message. The message text was in a work area chained off the HCT and did not need to be freed. However, with the introduction of the dynamic exits function, the data area is obtained by the service and must be freed by the caller using \$FREMAIN with SP=0, KEY=1, and LV=128. If the caller does not need

the message text, change the macro invocation to specify MESSAGE=YES (message is issued via WTO) or MESSAGE=SUPPRESS (no message is issued or returned).

4. **Enable load modules to support dynamic processing, if necessary.** With the dynamic exits function in z/OS V1R10, JES2 supports loading, refreshing, and deleting installation load modules using an operator command. JES2 will process all dynamic tables (TYPE=DYNAMIC) that exist in the module just as they were processed before z/OS V1R10 by initialization. Exit routine addresses are cleared when a module is deleted and replaced when a module is refreshed. (\$ADD LOAD will not affect any exit routine address.) However, even with this support, not all load modules will function properly if loaded, deleted, or refreshed after initialization. You must review all your installation load modules (ones loaded with the LOADMOD initialization statements) to determine whether they can be dynamically loaded or deleted and, if they cannot be dynamically loaded or dynamically deleted, decide what needs to be done.

Examples of code that cannot be dynamically loaded or deleted without changes are:

- Modules that have assembled tables that are pointed to by JES2 data areas, for example, a \$UCT data area pointed to by the \$UCT field in the \$HCT. Other tables can be pointed to by other user pointers in JES2 data areas, or the tables can be pointed to by a named token (created using the \$TOKENSR macro).
- Modules that front-end JES2 services.
- Modules that implement their own \$PCE or \$DTE.
- Modules with \$BERTTABS that were not loaded at initialization.
- Modules that define installation \$WSTABS.

Many of these modules will have an exit 24 (post initialization) associated with them. The module was loaded as part of normal initialization processing and then hooked to JES2 data structures using exit 24. Alternatively, they might have code that tests a user field for zero and, if it is zero, they do one-time initialization processing. However, if the initialization code involves obtaining storage and saving that storage address in a JES2 data area, it might not have a problem.

These exit routines might not function properly if they are processed with an operator command. In some cases there will be residual pointers in JES2 data areas that do not get updated. In other cases (such as user defined PCEs or DTEs) there might be no way to stop code that is running in the module. Modules with \$BERTTABS need special considerations. At JES2 initialization, the \$BERTTABS are processed into a table that is used while JES2 is running. Updating this table while JES2 is active could result in unexpected side effects. As a result, JES2 does not process new \$BERTTABS when a module is dynamically loaded. Only \$BERTTABS in modules that were loaded at JES2 initialization (including hot starts) are used. There is no restriction on refreshing modules with \$BERTTABS except that new \$BERTTABS will not be processed. Be careful with any module that contains \$WSTABS. For performance reasons, the address of the \$WSTAB is stored in various JES2 data areas. These addresses are only updated when the WS= parameter associated with a device is updated. New \$WSTAB entries are not processed until the WS= values on any device that use the new work selection criteria are updated (in general by operator command).

There are two ways to address modules that cannot be dynamically updated:

- Prohibit all dynamic processes for these modules. This can be done by setting the DYNAMIC=NO parameter on the \$MODULE macro in the modules.
- Use the new \$\$\$\$LOAD and \$\$\$\$DEL service to perform the initialization processing normally done by exit 24. This generally eliminates the need for exit 24 even for the case when the module is loaded by initialization processing. The \$\$\$\$LOAD and \$\$\$\$DEL routines can also be used to dynamically create and detach PCEs and DTEs. Furthermore, for the case where you want new code loaded but you also want the table that exists in the old module to remain, a function in \$\$\$\$DEL prevents JES2 from ever deleting the load module.

### **Update automation that handles \$D NODE, \$D EXIT, and \$D LOADmod command responses (Required-IF, as of OA21346, and as of R10)**

*Required if you have automation that examines these command responses.*

**For \$D NODE:** Before z/OS V1R10, the command response to \$D NODE would return all parameters. Beginning with z/OS V1R10, the command response is changed to display only an abbreviated list of parameters. For the complete list, you can issue the \$D NODE, LONG command. For more details about the parameters in the abbreviated list, see the description of the LONG parameter on the \$D NODE command.

**For \$D EXIT:** Before z/OS V1R10 (and APAR OA21346 on z/OS V1R9 and z/OS V1R8), the command response to \$D EXIT would return all parameters. Beginning with z/OS V1R10, the command response is changed to display only an abbreviated list of parameters. For the complete list, you can issue the \$D EXIT, LONG command. For more

details about the parameters in the abbreviated list of \$D EXIT command parameters, see the description of the LONG parameter on the \$D EXIT command.

**For \$D LOADmod:** Before z/OS V1R10 (and APAR OA21346 on z/OS V1R9 and z/OS V1R8), the command response to \$D LOADmod would return all parameters. Beginning with z/OS V1R10, the command response is changed to display only an abbreviated list of parameters. For the complete list, you can issue the \$D LOADmod, LONG command. For more details about the parameters in the abbreviated list of \$D LOADmod command parameters, see the description of the LONG parameter on the \$D LOADmod command.

**Migration action:** Review and update automation that examines these commands responses.

## JES2 Migration Actions Post-First IPL

### Migrate to z11 mode (Recommended, as of R11)

*Not required, but recommended to activate the full-function level of z/OS V1R11 JES2 processing.*

If you wish to take advantage of the full-function level of z/OS V1R11 JES2, you must be in z11 mode. Activating z11 mode upgrades the JES2 checkpoint and enables JES2 functionality that is introduced in z/OS V1R11, including JOE data area extensions supported by BERTs. For more information on the JES2 functionality introduced in z/OS V1R11, see the reference links below.

#### **Migration action:**

- After migrating to z/OS V1R11 JES2 on all systems in your MAS, determine your z11 checkpoint activation readiness:
  1. Use the \$D ACTIVATE command. This command indicates if activation to z11 mode will succeed.
  2. Review your current utilization of BERT data to determine if there are sufficient BERTs, as detailed in “**Check BERT utilization**” below.
- Run the JES2 \$ACTIVATE command to activate z11 mode following the considerations for this command found in *z/OS JES2 Commands*.

By default, JES2 restarts in the same mode (z2 or z11) as other members of the MAS (if any are active) or the mode the last active JES2 member was in when it came down. To restart JES2 in z2 mode, specify UNACT on PARM=. On a cold start JES2 starts in z11 mode unless overridden by OPTSDEF COLD\_START\_MODE.

**Tip:** Use IBM Health Checker for z/OS health check JES2\_Z11\_Upgrade\_CK\_JES2. This check determines if the system is ready to upgrade the JES2 checkpoint to z11 mode.

### **Check BERT utilization**

Before issuing the \$ACTIVATE, LEVEL=z11 command, review the current utilization of BERT data to determine whether there are sufficient BERTs. Additional BERTs are needed for each SYSOUT data set that has transaction data associated with it. These SYSOUT data sets can be seen using SDSF by setting APPC ON and examining SYSOUT data sets on the H and O panels; SYSOUT data sets with transaction data have nontraditional JES2 job IDs. Consider increasing the number of BERTs to correspond to two times the maximum number of transaction SYSOUT data sets on the system. BERT utilization should be monitored after the \$ACTIVATE to z11 mode to ensure there are sufficient BERTs for the jobs and SYSOUT in the MAS. There are several ways to determine your current BERT usage:

- The \$D CKPTSPACE, BERTUSE command displays a table of the types of control blocks in BERTs and how many BERTs are used by each control block type. The example below shows the output of the command:
 

```

$HASP852 CKPTSPACE CURRENT BERT UTILIZATION
$HASP852 TYPE COUNT CB COUNT
$HASP852 -----
$HASP852 INTERNAL 11 1,
$HASP852 JQE 211 108,
$HASP852 CAT 114 38,
$HASP852 WSCQ 1 1,
$HASP852 DJBQ 0 0,
$HASP852 JOE 0 0,
$HASP852 FREE 763 0
```

In the example, there are 108 JQEs that have a total of 211 BERTs associated with them. This example is for a system in z2 mode and does not have any BERTs associated with JOEs.

- The \$D ACTIVATE command displays the number of BERTs that are needed for activation to z11 mode. This is the number of BERTs that will be associated with JOEs after the \$ACTIVATE. The example below shows the output of the \$D ACTIVATE command.



```

$HASP895 $DACTIVATE
$HASP895 JES2 CHECKPOINT MODE IS CURRENTLY Z2
$HASP895 THE CURRENT CHECKPOINT:
$HASP895 -- CONTAINS 1100 BERTS AND BERT UTILIZATION IS 30
$HASP895 PERCENT.
$HASP895 -- CONTAINS 158 4K RECORDS.
$HASP895 z11 CHECKPOINT MODE ACTIVATION WILL:
$HASP895 -- EXPAND CHECKPOINT SIZE TO 165 4K RECORDS.
$HASP895 -- REQUIRE 22 ADDITIONAL BERTS AND UTILIZATION
$HASP895 WOULD REACH 32 PERCENT.
$HASP895 z11 ACTIVATION WILL SUCCEED IF ISSUED FROM THIS MEMBER.

```

In the example, there are 22 additional BERTs that will be used after the \$ACTIVATE to z11 mode, for transaction data associated with JOEs.

- A general history of BERT usage can be obtained by using the \$JD HISTORY(BERT) command or by using the SDSF RM panel. This displays the usage of BERTs after the system was IPLed. The example below shows the output of the \$JD HISTORY(BERT) command:

```


$HASP9130 D HISTORY
$HASP9131 JES2 BERT USAGE HISTORY
DATE TIME LIMIT USAGE LOW HIGH AVERAGE
-----
2009.086 16:00:00 1100 337 337 337 337
2009.086 15:50:09 1100 337 125 337 192

```

#### Reference information:


- For a list of the enhancements introduced in z/OS V1R11 for z11 mode, see *z/OS Introduction and Release Guide*.
- For \$ACTIVATE, \$D ACTIVATE, \$D CKPTSPACE and \$JDHISTORY command details, see *z/OS JES2 Commands*.

## JES3 Migration Actions for z/OS R11 from z/OS R9



---

**Migration Actions Pre-First IPL:**



**Identify the JES3 auxiliary address space to your security product (Required-IF, as of R11)**

- Verify the name of the address space (JES3AUX, for instance). Check for a profile defined in the STARTED class that covers the name of the JES3AUX address space. If none is defined, create a profile for JES3AUX. Make sure JES3AUX has the trusted attribute.

- **Update exit IATUX72 (Required-IF, as of R11)**
  - As of R11: when you use SSI 70 to move a data set from the hold queue to the TCP or BDT queues, SSI 70 calls exit 72. SSI 70 is a new caller of exit 72, so the parameter list is changed.
  - Reassemble your IATUX72 against the new parameter list, possibly update your exit to handle the possibility of SSI 70 calling this exit.
- **Update exit IATUX63 (Required-IF, as of R10)**
  - Before R10: information strings returned by SSI 54 (the “Who Am I” SSI) were built during JES3 initialization.
  - As of R10: They are also built during CONNECT processing.
  - IATUX63, the exit that builds the installation-defined SSI 54 string, is likewise called during CONNECT processing. Decide when this exit is called during CONNECT processing if you want to: do nothing, create a new installation-defined string, or update your existing installation-defined string.

38
z/OS 1.11: Migration - ...And It's Good! Part 2 of 2
© 2010 IBM Corporation

### JES3 Migration Actions Between z/OS V1 R9 and z/OS V1 R11

These migration actions were taken from *z/OS Migration*. Many descriptions and actions have been severely shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### JES3 Migration Actions You Can Do Now

#### **Modify code that depends on the value of STVSTOD returned by extended status (Required-IF, as of R11)**

*Required if the installation has a code dependency on the value of STVSTOD corresponding to the time that output is available and not before. It is expected that this dependency is unlikely.*

Extended status includes a timestamp, STVSTOD, for each dataset returned. Prior to z/OS V1R11, the timestamp value represented when that output dataset became available. Beginning with z/OS V1R11, the value is changed to represent when the dataset is created. This new value reflects when a dataset is available for spool browse.

#### **Migration action:**

Ensure that any dependency on STVSTOD will tolerate a timestamp corresponding to when the dataset is created instead of when output is available

### JES3 Migration Actions Pre-First IPL



#### **Identify the JES3 auxiliary address space to your security product (Required-IF, as of R11)**

*Required if your security product is configured to monitor started tasks.*

The JES3 auxiliary address space provides critical functional support for JES3 and is started during JES3 initialization. Before z/OS V1R11, a limited-function address space for JES3AUX was created by using a link to IEEMB881. Beginning with z/OS JES3 V1R11, the ASCRE macro is used to create the JES3 auxiliary address space. Therefore, you need to identify this address space to your security product (for example, RACF) if it is configured to monitor started tasks. JES3 will not start unless this action is taken.

#### **Migration action:**

1. Determine the JES3 auxiliary address space name. This name consists of the primary subsystem name combined with AUX. The primary subsystem name is defined in the IEFSSNxx member of SYS1.PARMLIB. For example, if the primary subsystem name is JES3, then the address space will be JES3AUX.

- Verify that the JES3AUX address space is identified to your security product. For RACF, check for a profile defined in the STARTED class that covers the name of the JES3AUX address space. If none is defined, create a profile for JES3AUX. Make sure JES3AUX has the trusted attribute. For example:

```
RDEFINE STARTED JES3AUX.* STDATA( USER(JES3ID) GROUP(SYS1) TRUSTED(YES) )
```

### **Update JES3 exit IATUX72 (Required-IF, as of R11)**

*Required if you use exit IATUX72.*

Beginning with z/OS V1R11, you can use SSI 70 to change the characteristics of a SYSOUT data set. If your changes are to move a data set from the hold queue to the TCP or BDT queues, SSI 70 calls exit 72 (IATUX72). Because SSI 70 is a new caller of exit 72, the parameter list of the exit, IATYUX72, is changed. This change consists of a new bit, YUX73SJF, which identifies SSI 70 as the caller, and a new pointer, UX72SMW, which addresses a work area, IATYSMW, that is used by SSI 70.

**Migration action:** Reassemble your IATUX72 exit against the new IATYUX72 parameter list. In addition, if your IATUX72 exit bases its actions either on its caller or on the caller specific data that is provided in the IATYU72 parameter list, update your exit to handle the possibility of SSI 70 calling the exit.

### **Update JES3 exit IATUX63 (Required-IF, as of R10)**

*Required if you use exit IATUX63.*

Before z/OS V1R10, information strings returned by SSI 54, also known as the "Who-Am-I" SSI, were built during JES3 initialization. Beginning with z/OS V1R10, the information strings are also built during CONNECT processing. This allows information to be updated when information about the global changes without the local having been restarted. The global information can change in two ways: (1) the global is IPLed and the JES3 release is changed during this IPL, and (2) a Dynamic System Interchange (DSI) is performed, which changes the name of the global and possibly the release. IATUX63, the exit that builds the installation-defined SSI 54 string, is likewise called during CONNECT processing.

**Migration action:** Update your IATUX63 code to test the flag YUX63CON. This flag is set during connect processing. When this bit is set, you have the following options:

- Return without changing anything.
- Free the storage for the old installation-defined string and create a new one.
- Update your existing installation-defined string without changing its location or length.

### **Check extended status applications for the need for testing STSTHTCP and STSTHBDT (Req-IF, as of R10)**

*Required if you have an application that looks specifically at the existing STSTHSTA flags.*

Before z/OS V1R10 JES3, terse and verbose output information returned by SSI 80 (extended status) included output on the HOLD and WTR queues. Beginning with z/OS V1R10, output on the BDT and WTR queues is returned in addition to output on the HOLD and WTR queues. The flags that JES3 sets to indicate the TCP and BDT queues are STSTHTCP and STSTHBDT, respectively, in the flag byte STSTHSTA.

**Migration action:** Look at your application code that uses the extended status interface (SSOBFUNC=SSOBESTA) and answer the following questions to determine what to do:

- Do you use the STATTYPE of STATOUTT or STATOUTV?
  - No: No action needed.
  - Yes: Proceed to the next question.
- Do you examine flag byte STSTHSTA?
  - No: No action needed.
  - Yes: Proceed to the next question.
- Do you test the hold type flags STSTHOPR, STSTHUSR, STSTHSYS, STSTHTSO, and STSTHXWT within STSTHSTA and take some actions if none of them are set?
  - No: No action needed.
  - Yes: Determine what action your code takes in this case and evaluate the need for adding code to test STSTHTCP and STSTHBDT.

### **Delete the EDIT keyword from the CONSTD statement (Required-IF, as of R10)**

*Required if you have any automated procedures or disk readers that specify the backspace, line delete, or escape character.*

In z/OS V1R10 JES3, the CONSTD EDIT= keyword is no longer supported. Message IAT3256 EDIT KEYWORD IGNORED, NO LONGER SUPPORTED is issued. JES3 initialization continues.

The characters previously defined as backspace, line delete, and escape are no longer special characters and are now treated by JES3 console services as literal characters. The MVS command delimiter from the INIT CMDDELIM() parameter in SYS1.PARMLIB(CONSOLxx) continues to be honored by JES3.

**Migration action:**


1. Remove the EDIT= keyword from the CONSTD statement to clean up the warning messages.
2. Remove automation or disk reader commands if they are using the backspace, line delete, or escape character. For example, before z/OS V1R10 JES3, if you had the pound sign defined as the backspace character and an automation routine performing an action on a job with a pound sign as part of its name, you had to enter the escape character before the pound sign. Otherwise, the command would treat the pound sign as a backspace and change the command to a job name that you did not want the command to operate on. With z/OS V1R10 JES3, because the pound sign is interpreted literally, you can remove the escape character.

**JES3 Migration Actions Post-First IPL**

<none>

## Language Environment Migration Actions for z/OS R11 from z/OS R9


### Migration Actions Pre-First IPL:

- 
 • **Determine the impact of added and changed runtime options (Required-IF, as of R11)**
  - As of R11: four suboptions were added to HEAPCHK for serviceability of heap pools. These suboptions allow you to specify a heap pools trace table size that is greater or less than the current hardcoded size of 1024, and to limit heap pools tracing to a single pool instead of always tracing all pools.
  - As of R10, and rolled back: HEAPPOOLS option was added to the CELQDOPT options group (was already valid in CEEDOPT and CEECOPT)
- If you use CEEPRMxx: ***no migration action*** if you want to use the defaults of the newly added options. The newly added options do not change default behavior.
- If you use assembler language modules: update and reinstall your usermods.

39 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## Language Environment Migration Actions for z/OS R11 from z/OS R9

### Migration Actions Pre-First IPL:

- 
 • **Use CEEPRMxx's CEEROPT instead of making copies of Language Environment initialization routines (Required-IF, as of R10)**
  - Many installations have copied LE init routines (making updates to the copies for specific runtime options), and STEPLIB'd to the copies for different applications. **This method will no longer work as of R10!**
  - Note! No effect on CEEBXITA users! Link CEEBXITA as you always have!***
  - Now in CEEPRMxx, CEEROPT(ALL) indicates that CEEROPT should be loaded for all application execution environments, and that the options found in that module should be used. **This method does not required any updates to applications!**
    - Turn on CEEROPT(ALL) in CEEPRMxx, which is not the default.
    - Create a separate CEEROPT module with your customization similar to how you made the changes before, and put in a data set. Use sample CEEWROPT.
    - STEPLIB to that CEEROPT data set you customized for your application.
    - Ensure you remove, and longer have the LE init modules in the STEPLIB concat.

40 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2009 IBM Corporation

## Language Environment Migration Actions Between z/OS V1 R9 and z/OS V1 R11

Some migration actions from specific elements were selected from *z/OS Migration*. Many descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### Language Environment Migration Actions You Can Do Now

<none>

### Language Environment Migration Actions Pre-First IPL

#### Update the CICS CSD based on the newest CEECCSD sample (Required)

Each release, Language Environment adds or deletes load modules in the CICS system definition (CSD) file. Thus, you should update the file each release using the program definitions found in member CEECCSD and, if using CICS Transaction Server (TS) for z/OS V3 (5655-M15), in member CEECCSDX..

**Migration action:** Update the CSD file using the program definitions in member CEECCSD (and member CEECCSDX if using CICS TS V3) found in the *hlq.SCEESAMP* data set.

**Note:**

1. The group containing the Language Environment runtime routines must be in the group list used during CICS startup.
2. The size of the CEECCSD decreased dramatically in z/OS V1R9. The decrease was intentional; it occurred because iconv() started using the Unicode converters instead of the converters supplied with Language Environment.

#### Review Language Environment load modules in LPA (Required-IF)

*Required if you need to make modules accessible through the link pack area (LPA).*

Each release you must update the Language Environment load modules that you make accessible through the link pack area (LPA). In addition, each release you should review your list of load modules in the LPA to determine if it's still suitable.

**Migration action:** Review Language Environment load modules in the LPA. To move load modules into the LPA, use the following sample members in the CEE.SCEESAMP data set:

- *AFHWMLP2*: This is a sample of all Language Environment Fortran component modules eligible for the LPA.
- *CEEWLPA*: This is a sample of a PROGxx member of SYS1.PARMLIB that includes all Language Environment CEE-prefixed runtime modules eligible for the LPA (that is, all Language Environment base modules) except the callable services stubs.
- *CELQWLPA*: This is a sample for AMODE 64 runtime support.
- *EDCWLPA*: This is a sample of a PROGxx member of SYS1.PARMLIB that includes all Language Environment EDC-prefixed and CEH-prefixed runtime modules eligible for the LPA (that is, all XL C/C++ component modules) except locales and code page converters.
- *IBMALLP2 (or IBMPLPA1 for Enterprise PL/I for z/OS)*: This is a sample of all Language Environment PL/I component modules eligible for the LPA.
- *IGZWMLP4*: This is a sample of all Language Environment COBOL component modules eligible for the LPA.

If you are considering placing the modules listed in *z/OS Language Environment Customization* in the LPA or the ELPA, then IBM recommends that you place the SCEELPA data set in the LPA list (LPALSTxx). SCEELPA contains Language Environment load modules that are reentrant, that reside above the 16 MB line, and that are heavily used by z/OS.

In *z/OS Language Environment Customization* you will also see tables of modules eligible for the LPA and the ELPA above and beyond what is found in the SCEELPA data set. You will need to use the dynamic LPA or MLPA approach to move these modules into the LPA or ELPA. You do not need to include recommended modules if they contain functions your installation does not use. Language Environment modules not listed in these tables can be moved into the LPA or ELPA at your discretion.



#### Determine the impact of added and changed run-time options (Required-IF, as of R11)

*Required you are using assembler modules to specify your installation-wide runtime options (CEEDOPT, CEECOPT, or*

CELQDOPT). If you are using CEEPRMxx (introduced in z/OS V1R7) and the default specification for the newly added options is acceptable to you, then no migration action is required.

In z/OS V1R11:

- No options were added.
- Four suboptions were added to HEAPCHK to enhance the serviceability of heap pools. The suboptions allow you to specify a heap pools trace table size that is greater or less than the current hardcoded size of 1024, and to limit heap pools tracing to a single pool instead of always tracing all pools.. The suboptions apply to the main user heap in an application and the 31-bit heap (`__malloc31()`) in a 64-bit environment:
  - Number of Entries: Specifies the number of entries to be recorded in one heap pool trace table for the main user heap in the application. Each pool has its own trace table. If the number of entries is 0, the heap pool trace table is not generated.
  - Pool Number: Specifies which pools are traced for the main user heap in the application. You can either trace one pool or all pools. The value should be a valid pool number from 1 to 12. If the pool number is 0, all pools will be traced.
  - Number of Entries 31: Specifies the number of entries to be recorded in one heap pool trace table when an AMODE 64 application is using heap storage from 31-bit addressable storage `__malloc31()`. Each pool has its own trace table. If the number of entries is 0, the heap pool trace table is not generated. This value is only supported in an AMODE64 environment.
  - Pool Number 31: Specifies which pools are traced when an AMODE 64 application is using heap storage from 31-bit addressable storage (`__malloc31()`). You can trace either one pool or all pools. The value should be a valid pool number from 1 to 12. If the pool number is 0, all pools will be traced. This value is only supported in an AMODE64 environment.

In z/OS V1R10:

- One runtime option, HEAPPOOLS, was added to the CELQDOPT options group. (HEAPPOOLS was already valid in both CEEDOPT and CEECOPT.) This support is integrated in z/OS V1R10; it was initially made available on z/OS V1R9 by APAR PK49427 and on z/OS V1R8 and V1R7 by APAR PK57579.
- No suboptions were added and no default changes were made to existing runtime options.

Because of these changes, you might have to perform migration actions.

#### Migration action:

If you are using CEEPRMxx to specify your installation-wide runtime options:

- The four suboptions added to HEAPCHK in z/OS V1R11 do not change the default behavior. You have no migration action.
- The HEAPPOOLS runtime option, which was added to CELQDOPT in z/OS V1R10, does not change the default behavior. You have no migration action.

**Tip:** Use IBM Health Checker for z/OS to check that default Language Environment runtime options are set within a CEEPRMxx parmlib member. The check is named CEE\_USING\_LE\_PARMLIB.

If you are using the assembler modules to specify your installation-wide runtime options, compare your existing source for the installation-wide runtime options CSECT, CEEDOPT (non-CICS environment), CEECOPT (CICS environment), or CELQDOPT (AMODE 64) with the new samples in *hlq.SCEESAMP* to determine whether you need to change the defaults. Then, update (and reinstall) your usermods with the following new options:

- The four suboptions added to HEAPCHK in z/OS V1R11.
- The HEAPPOOLS runtime option, which was added to CELQDOPT in z/OS V1R10. But note that this function, which is integrated in z/OS V1R10, was rolled back to prior releases (APAR PK49427 on z/OS V1R9 and APAR PK57579 on z/OS V1R8 and V1R7), so you might have already updated this usermod.



### **Use CEEPRMxx's CEEROPT instead of making copies of Language Environment initialization modules (Required-IF, as of R10)**

*Required if you have copied any Language Environment initialization module (CEEBINIT, CEEBINSS, CEEPIPI, or CEEBPICI) for the purpose of manually linking in a separate installation-wide runtime option CSECT, CEEDOPT, for use by a specific set of applications. Note: If you have used only SMP/E to maintain a customized CEEDOPT CSECT (by using a sample usermod CEEWD01), this migration action does not apply to you. Using the CEEWD01 usermod to specify your runtime options in CEEDOPT will still result in Language Environment correctly using your specified runtime options. This migration action is applicable only if you have linked CEEDOPT outside of SMP/E into separate libraries that were not SMP/E-controlled (and were STEPLIBed to by the applications).*

Many installations have made copies of Language Environment's initialization modules CEEBINIT, CEEBINSS, CEEPIPI, and CEEBPICI (updating each copy with specific runtime options specified in the CEEDOPT CSECT), and then used STEPLIB to access those modules for different applications. This method allows applications to individually have their own runtime options while not requiring any updates to the applications themselves. However,

this method has the disadvantages of being difficult to manage when service updates are made in Language Environment, having many copies of Language Environment modules to keep track of, and creating synchronization problems between the initialization modules and other Language Environment modules, which are hard to debug.

This method of customizing Language Environment's initialization modules no longer works as of z/OS V1R10. As of z/OS V1R10, the CEEDOPT CSECT is no longer part of the Language Environment's initialization modules CEEBINIT, CEEBINSS, CEEPIPI, and CEEBPICI. Therefore, any customization for runtime options added to these initialization modules will not be used by Language Environment for the application's runtime environment. To easily allow applications to continue to use customized runtime options, and to eliminate the need to make separate copies with updates to the initialization modules, support was added to the CEEPRMxx parmlib member. The CEEPRMxx CEEROPT statement (for non-AMODE 64 applications and CELQROPT for AMODE 64 applications) controls this support and CEEROPT(ALL) indicates to Language Environment that CEEROPT should be loaded for all application execution environments, and that the options found in that module should be used. Using CEEPRMxx's CEEROPT statement does not require any updates to the applications themselves. Because of these changes, you might have to perform migration actions.

**Migration action:**


1. CEEROPT processing can only be specified in your CEEPRMxx parmlib member (as CEEROPT is not an option in the CEEDOPT assembler language module). Turn on CEEROPT processing for all non-AMODE 64 environments by setting the keyword CEEROPT(ALL) in your active CEEPRMxx member. (CEEROPT(COMPAT) is the default, which indicates CEEROPT is to be loaded and used only for CICS and LRR environments. This is the behavior prior to z/OS V1R10.)
2. Create a separate CEEROPT module with runtime option overrides similar to those you are currently manually linking into a copy of one of the Language Environment initialization modules. You can use sample CEE.SCEESAMP(CEEWROPT) for the assembly and link-edit. Place each CEEROPT module in a separate data set.  
**Note:** You only need to set those runtime options that are different from the ones set in your installation defaults or system-wide defaults.
3. Place this CEEROPT data set into the STEPLIB concatenation for the application.
4. Review the other modules you have in the STEPLIB concatenation. If you had any Language Environment initialization modules in the concatenation for the sole purpose of customizing the CEEDOPT CSECT, you must remove those modules from the concatenation.

**Language Environment Migration Actions Post-First IPL**



<none>



## z/OS UNIX Migration Actions for z/OS R11 from z/OS R9



### Migration Actions You Can Do Now:


- 
**• Migrate from HFS to zFS file systems (Recommended)**
  - HFS has been stabilized, and zFS is strategic.
  - Convert your file systems from HFS to zFS.
- 
**• 1) Use ServerPac (during install, for the shipped file systems), 2) use BPXWH2Z (for after install, and other file systems you have), and 3) look at R10's F OMVS,NEWROOT= support for the sysplex root**
  - **Planned for z/OS R12:** zFS data sets to be recataloged with an indirect volser or system symbol!

### Migration Actions Pre-First IPL:

- Use the BPX.UNIQUE.USER profile instead of BPX.DEFAULT.USER (Recommended, as of R11)**
  - Before R11: BPX.DEFAULT.USER were assigned a default OMVS segments for the length of the user session. They used the same UID/GID.
  - As of R11: If BPX.UNIQUE.USER is used, users will have unique UID/GID saved into their profiles in the RACF data base.
  - Follow steps in *z/OS UNIX System Services Planning*.
- Update z/OS UNIX configuration files (Required-IF, as of R10)**
  - /etc/rc was changed to clear man page cache for R10.

41 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## z/OS UNIX Migration Actions for z/OS R11 from z/OS R9



### Migration Actions Post-First IPL:

- Use the Shell and Utilities version of the submit command (Recommended, as of R10)**
  - Before R10: submit command was a Tools and Toys download.
  - As of R10: it is part of the Shell and Utilities support in z/OS.
  - Start using the R10 submit command, remove the "old" submit.
- Use man instead of OHELP (Required-IF, as of R11)**
  - Before R11: OHELP command was used to display online reference information for many commands and services.
  - As of R11: Use the z/OS UNIX **man** command instead. You do not need to delete the IBM-supplied /samples/ohelp.ENU file or the local /etc/ohelp.ENU file.

42 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## z/OS UNIX Migration Actions Between z/OS V1 R9 and z/OS V1 R11

Some migration actions from specific elements were selected from *z/OS Migration*. Many descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration*.

### z/OS UNIX System Services Migration Actions You Can Do Now



#### **Migrate from HFS file systems to zFS file systems (Recommended)**

*Not required, recommended because the action is planned to become a requirement in a future release. zFS is the strategic file system for z/OS UNIX and continues to be enhanced to provide superior performance, reliability, and data integrity. IBM has not announced removal of support for the HFS file system.*

Before z/OS V1R7, the HFS file system was the primary hierarchical file system. As of z/OS V1R7, you can use any combination of HFS and zFS file systems. Because zFS has higher performance characteristics than HFS and is the strategic file system, you should migrate your HFS file systems to zFS.

The HFS and zFS file system types in mount statements and command operands are now generic file system types that can mean either HFS or zFS. Based on the data set type, the system will determine which is appropriate. But note that you must still specify a type (HFS or zFS and it cannot be defaulted), and if the type you specify is not correct for the file system being mounted, any associated parameter string setting in the mount statement or command is ignored, even though the system sets the type correctly and processes the mount.

Understand the zFS recommendations and limits. For more information, see Minimum and maximum file system sizes in *z/OS Distributed File Service zSeries File System Administration*

#### **Migration action:**

1. Before beginning the migration, do the following:

- Establish backout procedures.
- Decide on naming conventions.
- Decide on unavailability.
- Understand any cloning or deployment changes required by zFS systems being linear data sets. Considerations would include any copy utility invocations, BPXPRMxx specifications for symbolics, and placement of zFS file systems on system volumes.

2. Perform the conversion from an HFS to zFS file system.

**Tip:** Use the BPXWH2Z tool to perform the conversion. It is an ISPF-based tool that migrates HFS file systems to zFS file systems. Using its panel interface, you can alter the space allocation, placement, SMS classes, and data set names. A HELP panel is provided. With this tool, you can:

- Migrate HFS file systems (both mounted and unmounted) to zFS file systems. If the HFS being migrated is mounted, the tool automatically unmounts it and then mounts the new zFS file system on its current mount point.
- Define zFS aggregates by default to be approximately the same size as the HFS. The new allocation size can also be increased or decreased.
- Have the migration run in TSO foreground or UNIX background.

You can use the JCL sample ISPBTCH in SYS1.SAMPLIB to invoke BPXWH2Z as an ISPF batch job. Before you run the job, be sure to read the Notes section. When you run BPXWH2Z on your z/OS V1R11 system, it uses the z/OS V1R11 level of the **pax** command. This level was enhanced in z/OS V1R7 for sparse file support and other characteristics that are of concern when migrating from an HFS to zFS file system. You can manually migrate from an HFS to zFS file system without using the tool. However, you would need to allocate and format the target zFS file systems.

**Requirement:** The BPXWH2Z tool requires the zFS address space to be operational. Therefore, before attempting to migrate existing HFS to zFS file systems using BPXWH2Z, make sure that the zFS address space has been successfully configured and initialized.

**Tip:** Consider migrating the sysplex root file system to zFS after you IPL z/OS V1R11. You can now dynamically migrate the HFS sysplex root in a shared file system configuration to zFS while the root is in use, without disrupting workloads. Before z/OS V1R10, you could not migrate without disrupting active workloads. Although the shared file system configuration is required, the sysplex can be a single system. For the steps involved in migrating the sysplex root file system from HFS to zFS, see below.

**Planned for z/OS R12!** DFSMSdftp is planned to allow a zFS data set to be recataloged with an indirect volume serial or system symbol. This is designed to allow the zFS file systems used for z/OS system software files (called version root file systems) to be cataloged using an indirect volume serial or a system symbol the same way as non-VSAM data sets to make cloning and migration easier.

3. Change policies and scripts, and so forth, to reflect the change from the HFS file system to zFS file system.

**Tip:** Use the RMF Monitor III option to report on zFS activity.

### Migrating the sysplex root file system from HFS to zFS after IPLing z/OS V1R11:

Before you begin the migration:

- Ensure that the following requirements have been met:
  - All systems in the sysplex are at the V1R11 level.
  - The current sysplex root file system PFS, and the new sysplex root file system PFS, are up in all the systems in shared file system configuration.
- Be aware of the following restrictions:
  - The current sysplex root file system must be mounted as a read-only file system.
  - The systems that do not meet the requirements for this migration action cannot join the sysplex during the sysplex root file system migration processing, but they can join the sysplex after the sysplex root migration has completed.
  - The current sysplex root and the new sysplex root must be either HFS or zFS in any combination. If the new sysplex root is zFS, then it must be HFS-compatible.
  - The sysplex root or any directories on it cannot have been exported by the DFS or SMB server.
- Note the following:
  - Remote NFS mounts of the sysplex root or any directories on it are considered active use of the current sysplex root file system.
  - During the migration, the new zFS sysplex root file system must not be HSM-migrated, mounted, or in use.
  - Mount parameters are preserved during the migration or replacement of the sysplex root file system of the same file system type (PFS). They are dropped if the file system type is different.
  - Directories, data, files, and links are not copied from one file system to another.

Perform the migration as follows:

1. Ensure that a file system has been mounted read-only as the current sysplex root file system. When the root is mounted read-only, there are no function-shipping clients as long as physical paths to the DASD are available to each system. To verify that there are no function-shipping clients, issue: `D OMVS,F,NAME=root_file_system_name`. You should see `CLIENT=N` on each system.

2. Allocate and set up the new zFS sysplex root file system:

- Create a new zFS file system to be used as the new sysplex root file system. *z/OS Distributed File Service zSeries File System Administration* discusses creating and managing zFS file systems.

**Rules:**

- The UID, GID and the permission bits of the root directory in the new sysplex root file system must be same as the root directory in the current sysplex root file system.
- If the SECLABEL class is active and the MLFSOBJ option is active, the security label for the new zFS file system must match the assumed security label of the current sysplex root file system.
- On the new sysplex root file system, set up the active mount points and the symbolic links. The mount points and symbolic links must be the same as the ones on the current sysplex root file system. You can set them up either (1) manually or (2) by using the **pax** shell command to populate the new sysplex root file system using the existing sysplex root as a source. To do it manually, create a mount point in the existing sysplex root (for example, `/newroot`) and mount the new sysplex root file system in the `MODE(RDWR)` on that mount point. After mounting the new sysplex root file system, manually issue `MKDIRs` and `ln -s` to create the mount point directories and symbolic links similar to the existing sysplex root file system. Note that the new sysplex root file system must contain all active mount points and symbolic links exactly as on the existing sysplex root file system.
- Use the **pax** shell command to populate the new file system, using the existing sysplex root as a source.

**Example:**

```
cd /
pax -wr -pe -XCM ./ /newroot
```

For more information about using **pax** to copy data from an HFS file system to a zFS file system, see *z/OS Distributed File Service zSeries File System Administration*.

- Unmount the new zFS file system.

3. Before proceeding with this step, ensure PTF UA49782 for APAR OA29567 is installed on your z/OS V1R10 systems. Then, on any system in the shared file system configuration, issue:

F OMVS,NEWROOT=new.root.file.system.name,COND=<Yes|No>

**YES** Proceed conditionally. The system checks for active usage in the current sysplex root file system and reports the active usage in a BPXF245I message. If file activity is found, the command fails with EBUSY return code and JrActivityFound reason code. If file activity is not found, the command continues processing to replace the sysplex root. YES is the default.

**NO** Proceed unconditionally. The system checks for active usage in the current sysplex root file system and reports the active usage in a BPXF245I message. Replacement of the sysplex root file system will continue.

The migration of the sysplex root file system will begin. During the migration, active connections to files and directories in the current sysplex root file system are broken.

After the migration completes:

- The root CWD('/') is updated on all systems in the sysplex to point to the new sysplex root file system.
- New opens go to the new sysplex root file system. The current sysplex root for the root directory is replaced for all processes in all systems. The current directory for root directory is replaced for any processes using it.
- Old connections in the previous sysplex root file system might get EIO errors.

4. Update the TYPE parameter and name of the sysplex root file system in the BPXPRMxx member of SYS1.PARMLIB.

### **z/OS UNIX System Services Migration Actions Pre-First IPL**

#### **Update z/OS UNIX configuration files changed by IBM (Required)**

Some utilities provided by z/OS UNIX require the use of certain configuration files. You are responsible for providing these files if you expect to use the utilities. IBM provides default configuration files as samples in the /samples directory. Before the first use of any of these utilities, you must copy these IBM-provided samples to the /etc directory (in most cases). You can further customize these files to include installation-dependent information. An example is setting up the /etc/rc file by copying the sample file from /samples/rc to /etc/rc and then customizing it for the installation.

Whenever IBM ships a changed sample configuration file, you must copy the changed sample to its target location and, if you customized the previous version of the sample, incorporate the customization into the new version.

**Migration action:** For each of the files listed in the table below, copy the IBM-provided sample file to the listed target location. If you added installation-dependent customization to any of the IBM-provided files, make the same changes in the new versions of the files.

Utility	IBM-provided sample file	Target location	What changed and when
initialization	/samples/rc	/etc/rc	In z/OS V1R10, the initialization script was changed to clear out the man page cache. This eliminates the migration action "Update the date and time of book files used for man pages".

#### **Update automated actions for z/OS UNIX messages (Required-IF, as of R11 and R10)**

*Required if you have automation that is affected by the message changes.*

The following z/OS UNIX message have changed:

- Before z/OS V1R11, message BPXO040I was issued in response to one of the following DISPLAY OMVS commands: D OMVS,A, D OMVS,U, D OMVS,PID=  
Beginning with z/OS V1R11, message BPXO070I is issued instead.
- Before z/OS V1R11, message BPXO046I was issued in response to the DISPLAY OMVS,PFS command. As of z/OS V1R11, message BPXO068I is issued.

- Before z/OS V1R10, message BPXF034I was issued if a file system was quiesced for more than 10 minutes. As of z/OS V1R10, message BPXF083I is issued instead. The new message includes the file system name, quiescing system, job name, and process ID (PID).

## **z/OS UNIX System Services Migration Actions Post-First IPL**

### **Use man instead of OHELP(Required-IF, as of R11)**

*Required if you use the OHELP command.*

Previously, the TSO/E OHELP command was used to display online reference information for shell commands, TSO/E commands, C functions, callable services, and messages issued by the shell and **dbx**. It required the use of bookshelf provided by the z/OS *Collection Kit*. In V1R11, the OHELP command is no longer available.

**Migration action:** Use the z/OS UNIX **man** command instead. You do not need to delete the IBM-supplied /samples/ohelp.ENU file or the local /etc/ohelp.ENU file.

### **Use the BPX.UNIQUE.USER profile instead of BPX.DEFAULT.USER (Recommended, as of R11)**

*Recommended because in the future, support for the BPX.DEFAULT.USER profile might be removed. The use of BPX.UNIQUE.USER increases security because shared IDs are not used.*

Before z/OS V1R11, if the BPX.DEFAULT.USER profile in the FACILITY class was defined, users who accessed z/OS UNIX services who did not have an OMVS user or group segment were assigned the default OMVS segments for the length of the user session. All users of the default OMVS segments shared the same UID and GID. As of z/OS V1R11, if BPX.UNIQUE.USER has been defined, users who access z/OS UNIX services who do not have an OMVS user or group segment are automatically assigned an OMVS segment with a unique UID and GID. The new OMVS segments are added to the user and group profiles in the RACF database.

**Migration action:** Follow the steps in *z/OS UNIX System Services Planning* to set up the BPX.UNIQUE.USER profile. If BPX.DEFAULT.USER has not been deleted, BPX.UNIQUE.USER takes precedence when default OMVS segments are used.

To remove the BPX.DEFAULT.USER profile, use the following RACF commands:

```
RDELETE FACILITY BPX.DEFAULT.USER
SETROPTS RACLIST(FACILITY) REFRESH
```

### **Update automation that handles ISPF Edit and ISPF Browse (Required-IF, as of R11)**

*Required if you have automation that depends on the display of OEDIT and OBROWSE panels.*

Before z/OS V1R11, the TSO/E commands OEDIT and OBROWSE displayed panels that were different than the conventional ISPF Edit and ISPF Browse panels and used the ISPF EDIF and BRIF services. As of z/OS V1R11, OEDIT and OBROWSE will no longer display their unique panels and will directly invoke the ISPF EDIT and BROWSE commands.

**Migration action:** Update any automation that depends on the display of panels by OEDIT and OBROWSE.

### **Use the Shell and Utilities version of the submit command (Recommended, as of R10)**

*Recommended because the new supported version of the submit command is enhanced compared to the Tools and Toys version.*

Before z/OS V1R10, the **submit** command was available from the Tools and Toys section of the z/OS UNIX Web site. As of z/OS V1R10, Shell and Utilities support of the **submit** command has been added. The supported version differs from the Tools and Toys version in a number of ways:

1. Multiple files can be used as input.
2. After the **submit** command completes, a message that indicates success or failure is sent. If the command succeeds, the resulting message contains the job IDs of the submitted jobs.
3. Formatting for submitting jobs from a partitioned data set (PDS) is different. Because the supported **submit** command is now consistent with the other subcommands, the // characters are required as a prefix to the PDS name. For example: "'/' full.pds.name(member)'"
4. The Tools and Toys version of **submit** had the following conventions, which have been changed:
  - **?** was used as a way to display usage but is now treated as a valid file name.
  - **-** was used as a file name but is now used to specify standard input.

**Migration action:**

1. If you are using the Tools and Toys version of the **submit** command, delete it or move it to a location that is not in the system's search path. If you leave it in the search path, it might be picked up instead of the supported version that is in the **/bin** directory.

2. Begin using the supported version in the `/bin` directory.

### **Update cksun and sum invocations (Required-IF, as of R10)**

*Required if any automated tasks rely on cksun or sum exiting when an automatic conversion error occurs.*

Before z/OS V1R10, the `cksum` and `sum` commands exited if they could not turn off automatic conversion on a file stream. As of z/OS V1R10, if this error is encountered, the file stream (for example, `stdin`) is skipped and the remaining files are processed. The message associated with this error is FSUMF182. This change should not affect most users.

**Migration action:** Modify any `sum` or `cksum` invocations (for example, in shell scripts) that rely on the utility exiting when automatic conversion errors occur.

### **Update ln invocations to accommodate changes made to the ln shell command (Req-IF, as of R10)**

*Required if `ln -e` is used in either of the following ways: 1) in conjunction with the `-s`, `-r`, or `-R` option, or 2) with an existing target directory.*

Before z/OS V1R10, if you specified the `-e` option with the `-s`, `-r`, or `R` option when issuing the `ln` shell command, the results were not valid.

- If *source* was a directory containing symbolic links or external links and *target* was an existing directory, and you issued `ln -e` with `-s`, `-r`, or `-R`, one of the following events occurred:
  - Message FSUM6245 was issued.
  - One or more incorrect symbolic links was created.
  - One or more incorrect external links was created.
- If *source* was a file or a PDS or PDSE member, `-e` was specified with `-s`, and *target* was not a directory, `ln` created a symbolic link.
- If `-e` was specified with an existing target directory, `ln` created an incorrect external link.

As of z/OS V1R10, the `-e` option is mutually exclusive with the `-s`, `-r`, and `-R` options. You will receive an error message if you use `-e` in conjunction with `-s`, `-r`, or `-R`, or with an existing target directory.

#### **Migration action:**

- Modify any `ln` invocations (for example, in shell scripts) so that `-e` is not used in conjunction with the `-s`, `-r`, and `-R` options.

If you want to create . . .	Use . . .
A symbolic link	<code>-s</code>
An external link	<code>-e</code>
Symbolic links or external links recursively	<code>-s</code> with <code>-r</code> or <code>-R</code>

- Modify `ln` invocations so that `-e` is not used if the target specified on the command line is an existing directory.


### **Update scripts to accommodate changes made to the locale shell command (Req-IF, as of R10)**

*Required if automated scripts evaluate strings based on error messages FSUM6701 or the default locale output.*

Before z/OS V1R10, the `locale` shell command did not support the `LC_TOD` category and associated keywords. The default locale command (specified without any options or operands) did not display the `LC_TOD` environment variable, and using `LC_TOD` (or associated keywords) with or without options generated error message FSUM6701. Beginning with z/OS V1R10, the default locale command displays `LC_TOD` information along with other `LC_*` environment variables. Also, using `LC_TOD` or associated `LC_TOD` keywords is honored.

**Migration action:** Modify automated scripts if they evaluate strings based on error message FSUM6701 or the default locale output.

## Marna's "Big Migs" for z/OS R11 from z/OS R9




**Migration Actions You Should NOT Overlook:**

- Install coexistence and fallback PTFs (Required)
- Add or change volumes to keep file systems in a single data set (Recommended)
- Accept CustomPac default file system of zFS (Recommended)
- Ensure your ESQA specification in IEASYSxx is adequate (Req-IF, R11)
- Provide an appropriate HZSPDATA data set on fallback (Req-IF, as of R10)
- Reallocate the WLM couple data set (Required-IF, as of R10)
- Update callers of CNZMXURF that do not hold appropriate serialization (Required-IF, as of R10)
- Accommodate DEFAULT route usage changed for CINET environments (Required-IF, as of R11)
- Run the syslog daemon in the daemon in the background when starting from UNIX shell (Required-IF, as of R11)
- ICSF HCR7750: Increase the size of your PKDS (Required, as of HCR7750)
- ICSF HCR7770: Modify ICSF startup procedure (Required, as of HCR7770. not in z/OS R11)
- Determine if you will be using the IBM TDS WLM support (Req-IF, R11)

43 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## Marna's "Big Migs" for z/OS R11 from z/OS R9



**Migration Actions You Should NOT Overlook:**

- Install the new HCM (Required-IF, as of R11)
- Redefine existing VSAM data sets that contain the IMBED, REPLICATE, and KEYSRANGE attributes (Recommended)
- Ensure that sysplex\_admin\_level=2 is available on all systems (Required-IF, as of R11)
- Discontinue use of zFS multi-file system aggregates to compatibility mode aggregates (Required-IF, as of R11 and Recommended)
- Permit the SMB server ID to BPX.DAEMON and BPX.SERVER (Required-IF, as of R11)
- Identify the JES3 auxiliary address space to your security product (Required-IF, as of R11)
- Use CEEPRMxx's CEEROPT instead of making copies of Language Environment initialization routines (Required-IF, as of R10)
- Migrate from HFS to zFS file systems (Recommended)

44 | z/OS 1.11: Migration - ...And It's Good! Part 2 of 2 | © 2010 IBM Corporation

## z/OS 1.11: Migration Part 2 Summary

- **General Migration Actions for Everyone:**
  - TCPIP.SEZAEXEC and ISF.SISFMOD I are new, HWIBCP II new address space, SMP/E utility parameters to support assembling, VE EVLM is removed during installation.
- **BCP Migration Actions:**
  - Remove MFA automation exits, Use new defaults in sysplex, Prepare for zAAP on zLIP, Accommodate quoted arg changes in System REXX, Ensure IEFACTRT field lengths are ok, Use VOLUME with DISP=UNCATLG migration aid.
- **Communications Server Migration Actions:**
  - Removed functions, Shared FTP.DATA between FTP server and client, Syslogd job name is proc name, More space in dump ds for VTAM Internal Trace data space, More region size for Telnet and VTAM.



80

z/OS 1.11: Migration - It's Up... Part 1 of 2

© 2010 IBM Corporation

## z/OS 1.11: Migration Part 2 Summary

- **DFSMS Migration Actions:**
  - Move from VRSEL(OLD) to NEW in EDGRMMxx, Use ALL, DATE, or VERSIONS of (H)BDELETE, Use rmm parmlib options instead of usermod, Permit DEVMAN to use ICKDSF, DEVSERV command output, CA sizes may change, Use new DEFRAG command, New facility class names for (H)BACKDS NEWNAME and (H)BACKDS RETAINDDAYS commands. Larger REGION for ADRDSSU
- **DFS Migration Actions:** Unmount will fail on quiesced, growing, or cloning file systems.
- **JES2 Migration Actions:** Changes for variable extension for JOE data area. Migrate to z I I mode (more BERTs needed?)
- **JES3 Migration Actions:** IATUX63 and IATUX72 changes.
- **LE Migration Actions:** New runtime options and sub-options (no change in behavior, though).
- **z/OS UNIX Migration Action:** Use BPX.UNIQUE.USER instead of BPX.DEFAULT.USER, Submit command, Use man instead of OHELP.



81

z/OS 1.11: Migration - It's Up... Part 1 of 2

© 2010 IBM Corporation