
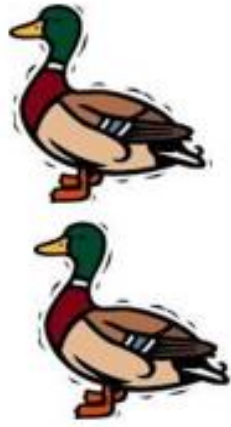



Migrating to z/OS 2.2: Part 2 of 2



Marna WALLE, mwalle@us.ibm.com
z/OS System Installation
Member of the IBM Academy of Technology
IBM z Systems, Poughkeepsie NY



#SHAREorg

SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

Copyright (C) 2015 by SHARE Inc. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or by any information storage and retrieval system, without the prior written permission of SHARE Inc.

Abstract:

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 2.2 from either z/OS 2.1 or 1.13. It is strongly recommended that you attend both sessions for a complete migration picture.

The general availability date planned for z/OS V2.2 is September 30, 2015.

Trademarks



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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml.

® AS400®; e-business®; logo®; DBE, ESDO, eServer, FICON, IBM®, IBM logo®, (Series®), iVPS, DS200®, pSeries®, RS6000®, S/390, VME548, VSEESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System®, System 8, System p, System p5, System z, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprocessing in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Migrating to z/OS V2.2: Part 2 of 2 Agenda




- **Scope of presentation**
- **Definition of a “migration action”**
- **Overview of migration actions for z/OS V2R2 from z/OS V2R1 or R13:**
 - ❖ General Migration Actions
 - ❖ BCP
 - ❖ Communications Server
 - ❖ DFSMS and ICKDSF
 - ❖ Distributed File Service – SMB and zFS
 - ❖ For FIPS 140
 - ❖ JES2
 - ❖ JES3
 - ❖ z/OSMF
 - ❖ z/OS OpenSSH
 - ❖ z/OS UNIX




Scope of Presentation




- This presentation is applicable to z/OS V2R2 migrations from either z/OS V2R1 or R13.
- Not fully inclusive of all migration actions, but rather gives you an overview of migration actions that are:
 - Very important to understand
 - May be common to many users
- Remember: Use the *z/OS V2R2 Migration* book for a complete list of all migration actions.
 - The latest level is **GA32-0889-05**.
 - The format of the *z/OS Migration* book that is targeted to your specific migration path:
 - Chapter 1: Introduction *for all users*
 - Chapter 2: General migration actions *for all users*
 - Chapter 3: Migration from z/OS V2R1 *for V2R1 -> V2.2 users*
 - Chapter 4: Migration from z/OS R13 *for R13 -> V2.2 users*

 4 © 2015 IBM Corporation

Ye Olde Definition of “Migration Action”:



- Upgrading to a new z/OS release is a two step process:
 1. **Migration:** the installation of a new version or release of a program to replace an earlier version or release.
 2. **Exploitation:** usage of new enhancements available in the new release. Not covered in this presentation.
- 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 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**

5  Means “don’t overlook!”  Means some programmatic assistance is available  Means a cleanup action on

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 V2R2 level of code to make these changes, and the changes don't require the z/OS V2R2 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 V2R2.
- **Pre-First IPL.** These are migration actions that you perform after you've installed z/OS V2R2 but before the first time you IPL. These actions require the z/OS V2R2 level of code to be installed but don't require it to be active. That is, you need the z/OS V2R2 programs, utilities, and samples in order to perform the migration actions, but the z/OS V2R2 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 V2R2 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 V2R2 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 V2R2. You need a running z/OS V2R2 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 V2R2 to be active to perform the task. You might perform the task quite a while after the first IPL.

Icons used in this presentation:



means that you shouldn't overlook this migration action.



means that an IBM Health Check (using the IBM Health Checker for z/OS function) can help you with this migration action.




means that this is a cleanup item or contains a portion that is a cleanup item. It is associated with something that is obsolete. It may cause confusion if someone thinks it does something. It is best to perform this action to avoid any confusion, since it is not needed anymore.

Elements with Migration Actions for z/OS V2R2

These elements have new V2R2 migration actions:

<ul style="list-style-type: none"> ➤ BCP ➤ BookManager Build ▪ CIM ➤ Communications Server ➤ Crypto Services – ICSF, SSL, PKI ➤ DFSMS ▪ DFSORT ➤ Distributed File Service – SMB ▪ HCD ▪ HLASM ➤ IBM HTTP Server ▪ IBM Tivoli Directory Server ➤ z/OS Management Facility 	<ul style="list-style-type: none"> ▪ ICKDSF ▪ Infoprint Server ➤ JES2 ➤ JES3 ▪ Language Environment ▪ Library Server ▪ RMF ▪ SDSF ▪ Security Server (RACF) ▪ XL C/C++ ➤ z/OS Font Collection ➤ z/OS OpenSSH ➤ z/OS UNIX
---	--

➤ means that some of that element's migration actions are discussed in these two presentations
▪ means that you need to look for that element's migration actions in the z/OS V2R2 Migration book



6
© 2015 IBM Corporation

Migration Actions for Elements for z/OS V2R2

When migrating from z/OS V2R1 to z/OS V2R2, the specified elements in the slide above have migration actions. Refer to *z/OS V2R2 Migration* for complete information on the required migration actions for all elements, and if you are on the path from z/OS R13. Some migration actions for selected elements follow in this presentation. This presentation does not cover all possible migration actions.

General Migration Actions for z/OS V2R2



• Migration Actions Pre-First IPL:

- **Accommodate new address spaces (Recommended)**
 - **New in V2R2:** **IBM HTTP Server Powered by Apache** (covered in Part 1)
 - **IBM Knowledge Center for z/OS**, in an address space that is an instance of the WebSphere Liberty Profile from z/OSMF.
 - **z/OSMF**, in address spaces IZUANG1 and IZUSR1.

- **New in V2R1:** **PCIE** (PCI Express) and **FPGHWAM** (HW Accelerator Mgr). Started automatically as of zEC12/zBC12 (regardless if zEDC or RoCE facilities are installed). They are persistent, and require no security customization. Expect either:
 - IQP002I PCIE INITIALIZATION COMPLETE , or
 - IQP031I REQUESTED SERVICE IS UNSUPPORTED BY HARDWARE
- **IBM Health Checker for z/OS** and Generic Tracker, **GTZ**, (see later slides)
- **JES2 Converter/Interpreter jesxCInn**, where jesx is the subsystem name, and nn is 01-25 for uniqueness. They are persistent, and used for interpretation process during conversion phase *only when* INTERPRET=JES is specified on JOBDEF. Default number of AS is 2, max is 25. Proper userid must be assigned to each AS, using started proc tables (ICHRIN03) or a STARTED class profile has been defined to match AS names.

7

© 2015 IBM Corporation

General Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Remove references to deleted syslib data sets and paths (Required)**
 - **Removed in V2R2:** **BookManager BUILD's** EOY.*, **IBM HTTP Server Powered by Domino's** IMW.* and /usr/lpp/internet/, **SDSF's** ISF.AISFMOD1, ISF.AISFSRC1, and ISF.AISFJCL1, **PKI's** /usr/lpp/pkiserv/samples/ihs7.
 - **Removed in V2R1:** ICLI's SYS1.FOMDATA and paths
- **Add references to new syslib data sets and paths (Required)**
 - **New in V2R2:** 8 **XL C/C++'s** CBC.* data sets, 5 **IBM HTTP Server Powered by Apache's** HAP.* data sets and path /usr/lpp/ihsa_zos/IBM, 1 **z/OSMF** dlib IZU.AIZUFS and path /usr/lpp/zosmf/IBM, and **KC** dlib HKC.AHKCKC4Z and path /usr/lpp/kc4z/IBM
 - **New in V2R1:** **z/OS Font Collection** data sets (from Part 1), path /usr/lpp/booksrv/plugins/IBM for **Library Server**, /path /usr/lpp/hzc/include/IBM and /usr/lpp/hzc/lib/IBM for **zEDC**.
- **Update your health check customization for modified checks (Recom)**
 - **New in V2R2:** 17 checks. **V2R1:** 27 checks
 - **Changed in V2R2:** 10 checks **V2R1:** 11 checks (incl.RACF_SENSITIVE_RESOURCES in both V2R2 and V2R1)
 - **Deleted in V2R2:** 6 checks **V2R1:** 2 checks 🗑️

8

© 2015 IBM Corporation

General Migration Actions For z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 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 V2R2 systems during your migration, and allow back out from z/OS V2R2 if necessary. Use the SMP/E REPORT MISSINGFIX command in conjunction with the FIXCAT type of HOLDDATA as follows:

1. Acquire and RECEIVE the latest HOLDDATA onto your pre-z/OS V2R2 systems. Use your normal service acquisition portals or download the HOLDDATA directly from <http://service.software.ibm.com/holddata/390holddata.html>. Ensure you select **Full** from the Download NOW column to receive the FIXCAT HOLDDATA, as the other files do not contain FIXCATs.
2. Run the SMP/E REPORT MISSINGFIX command on your pre-z/OS V2R2 systems and specify a Fix Category (FIXCAT) value of **"IBM.Coexistence.z/OS.V2R2"**. The report will identify any missing coexistence and fallback PTFs for that system. For complete information about the REPORT MISSINGFIX command, see *SMP/E Commands*.
3. Periodically, you might want to acquire the latest HOLDDATA and rerun the REPORT MISSINGFIX command to find out if there are any new coexistence and fallback PTFs.

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/atmastr.nsf/WebIndex/PRS268>
 - Business partners: <http://partners.boulder.ibm.com/src/atmastr.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.

General Migration Actions Pre-First IPL

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 V2R2 /etc and /var subdirectories.

Copy files from your old system to the z/OS V2R2 /etc and /var subdirectories, and then modify the files as necessary to reflect z/OS V2R2 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 V2R2 elements and features use /etc:

- BCP (Predictive Failure Analysis)
- CIM
- Communications Server – IP
- Cryptographic Services – PKI Services and System SSL
- DFSMSrmm
- Distributed File Service. The SMB server uses /etc/dfs.
- IBM HTTP Server Powered by Apache
- z/OSMF
- IBM Tivoli Directory Server – uses /etc/ldap.

Migrating to z/OS 2.2: Part 2 of 2

- Infoprint Server
- Integrated Security Services. The Network Authentication Service uses /etc/skrb.
- Library Server
- z/OS UNIX System Services

The following z/OS V2R2 elements and features use /var:

- Cryptographic Services – OCSF
- DFSMSrmm
- IBM Tivoli Directory Server – uses /var/ldap.
- z/OSMF
- 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*.

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 tables in *z/OS Migration* 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
- 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 V2R2 and V2R1.

Add references to new data sets (Required)

Migration action: For z/OS V2R2, the following elements had data sets and paths that were added:

- XL C/C++
- IBM HTTP Server – Powered by Apache
- IBM Knowledge Center for z/OS
- z/OSMF

For z/OS V2R1, the following data sets and paths were added:

- SYS1.FONTLIB for the z/OS Font Collection z/OS
- SYS1.FONTLIBB for the z/OS Font Collection z/OS
- /usr/lpp/booksrv/plugins/IBM/ for Library Server
- SYS1.SFNTILIB for the z/OS Font Collection z/OS (use in place of SYS1.FONTOLN)
- /usr/lpp/fonts/worldtype/IBM/ for the z/OS Font Collection z/OS

Migrating to z/OS 2.2: Part 2 of 2

- SYS1.SFONDLIB for the z/OS Font Collection z/OS
- /usr/lpp/hzc/ include/IBM/ in z/OS UNIX for zEnterprise Data Compression (zEDC).
- /usr/lpp/hzc/lib/IBM/ in z/OS UNIX for zEnterprise Data Compression (zEDC)
- EPH.AEPHPLIB for Library
- EPH.AEPHTAB for Library Server
- SYS1.AFNT300 for the z/OS Font Collection z/OS
- SYS1.AFNTDLIB for the z/OS Font Collection z/OS
- SYS1.AFNTILIB for the z/OS Font Collection z/OS
- SYS1.AFNTLIBB for the z/OS Font Collection z/OS
- SYS1.AFONTLIB for the z/OS Font Collection z/OS
- SYS1.AFONTHFS for the z/OS Font Collection z/OS
- SYS1.FONT300 for the z/OS Font Collection

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 parmlib member IEASYSxx is adequate.

The following elements add new address spaces for z/OS V2R2:

- **IBM HTTP Server - Powered by Apache**, which has one or more new address spaces that are associated with it. For information about setting up IBM HTTP Server - Powered by Apache, see *z/OS V2R2.0 HTTP Server - Powered by Apache User's Guide*.
- **Knowledge Center for z/OS**. This new element in z/OS V2R2 is started in an address space that is an instance of the WebSphere Liberty Profile, which is installed with z/OSMF. For information about setting up IBM Knowledge Center for z/OS, see *IBM Knowledge Center for z/OS Configuration and User Guide*.
- **IBM z/OS Management Facility (z/OSMF)**, which has the address spaces **IZUANG1** and **IZUSVR1**. For information about setting up z/OSMF, see *IBM z/OS Management Facility Configuration Guide*

The following address spaces are new for z/OS V2R1:

- **PCIE** (PCI Express) and **FPGHWAM** (Hardware Accelerator Manager). They provide the infrastructure for PCI Express I/O and hardware accelerator activities. These address spaces are started automatically during z/OS initialization, if the appropriate z/OS PCIE facilities hardware is installed. They are persistent address spaces. If the PCIE address space is successfully initialized, the message is displayed:
IQP002I PCIE INITIALIZATION COMPLETE
If the required hardware is not installed, the following message is written to the hardcopy log:
IQP031I REQUESTED SERVICE IS UNSUPPORTED BY HARDWARE .
For information about the PCIE messages, see *z/OS MVS System Messages, Volume 9 (IGF-IWM)*. For information about the FPGHWAM (Hardware Accelerator Manager) messages, see *z/OS MVS System Messages, Volume 5(EDG-GFS)*.
PCIE and FPGHWAM do not require any security customization.
- **IBM Health Checker for z/OS**. As of z/OS V2R1 the system starts IBM Health Checker for z/OS address space automatically during system initialization.
- **JES2 Converter/Interpreter**. A new persistent address space is used when the interpretation process is performed for a job during the JES2 conversion phase. The address space is only created when INTERPRET=JES is specified on JOBDEF. The number of address spaces used depends on the CISUB_PER_AS setting on JOBDEF. The number of conversion processes (PCEDEF CNVTNUM=) divided by the number of subtasks per address space (CISUB_PER_AS) gives the number of address spaces created. The default number of created address spaces is 2 and the maximum number is 25. The name of the address spaces are **jesxCInn** where jesx is the JES2 subsystem name and xx is a number (from 01 to 25) to create uniqueness. This address space accesses the PROCLIB data sets defined in the JES2 start PROC and using the JES2 dynamic PROCLIB service. You must ensure that a proper user ID is assigned to the address space (presumably the same user ID that is assigned to the JES2 address space) using entries in either the started procedures table (ICHRIN03) or that a STARTED class profile has been defined that matches each new address space names. If you prefer, both the started procedures table and STARTED class profile might be in place. This action ensures that the correct user ID is assigned.
- **Generic Tracker (GTZ)**. As of z/OS V2R1, GTZ replaces the Console Tracking Facility with enhanced capabilities to track use of system functions.

The MAXUSER value in parmlib 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.)

Verify your IFAPRDxx Product ID is correct for z/OS Version 2 (Required)

Because z/OS has changed from Version 1 (V1) to Version 2 (V2), the IFAPRDxx parmlib member Product ID should reflect the new program number for z/OS. The Feature, Release, and Modification values also found in IFAPRDxx should be specified as "*", so those values do not have to change; only the Product ID is affected by this change.



If you are sharing the IFAPRDxx parmlib member between z/OS V1 and V2 systems, keep the V1 Product ID statements (for 5694-A01). You can have both the V1 Product ID and the V2 Product ID statements in the same IFAPRDxx parmlib member. After you have completed your migration to z/OS V2, you can remove the V1 Product ID statement from your IFAPRDxx parmlib member.

Migration action:

- If you use ServerPac, the customized IFAPRDxx parmlib member has been shipped to you in CPAC.PARMLIB. Verify that you are either using that parmlib member, or have copied its contents to a parmlib member you are using.
- If you are using CBPDO, use provided job ZOSREG in the RIMLIB, to place a customized IFAPRD00 member into your parmlib.

Note: Customers must ensure that the policy in IFAPRDxx enables only the version/release at which they are licensed. Use of and enablement of z/OS features is subject to the z/OS license terms and conditions and must be done with the knowledge of your asset manager according to the terms and conditions for z/OS. For additional license terms and conditions, see the "Usage Restriction" section of the *z/OS Licensed Program Specifications*.

Update your check customization for modified IBM Health Checker for z/OS checks (Recommend)

Not required, but 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 following Health Checks are new in z/OS V2R2:

- CATALOG_ATTRIBUTE_CHECK
- CTRACE_DEFAULT_OR_MIN
- DMO_REFUCB
- ICSF_KEY_EXPIRATION (added in ICSF FMID HCR77B0)
- IOS_DYNAMIC_ROUTING
- JES3_DATASET_INTEGRITY
- JES3_DOT_POOL_USAGE
- JES3_JET_POOL_USAGE
- JES3_OST_POOL_USAGE
- JES3_SEE_POOL_USAGE
- PFA_PRIVATE_STORAGE_EXHAUSTION
- RACF_ENCRYPTION_ALGORITHM
- RACF_PASSWORD_CONTROLS
- RACF_RRSF_RESOURCES
- TSOE_OPERSEWAIT_SETTING
- USS_KERNEL_RESOURCES_THRESHOLD
- ZFS_CACHE_REMOVALS

The following Health Checks are changed by IBM in z/OS V2R2:

- ASM_PLPA_COMMON_SIZE
- ASM_PLPA_COMMON_USAGE
- CNZ_Task_Table
- RACF_SENSITIVE_RESOURCES
- RSM_HVSHARE
- USS_KERNEL_PVTSTG_THRESHOLD
- XCF_CF_STR_PREFLIST
- ZFS_VERIFY_CACHESIZE
- ZOSMIGREC_SUP_TIMER_INUSE
- ZOSMIGV2R1_ZFS_VERIFY_CACHESIZE



The following Health Checks are deleted by IBM in z/OS V2R2:

- USS_KERNEL_STACKS_THRESHOLD
- ZOSMIGREC_ZFS_RM_MULTIFS
- ZOSMIGV1R13_ZFS_FILESYS
- ZOSMIGV2R1_CS_GATEWAY
- ZOSMIGV2R1_CS_LEGACYDEVICE
- ZOSMIGV2R1_DEFAULT_UNIX_ID

The following Health Checks were new in z/OS V2R1:

- CATALOG_RNLS
- ICSF_COPROCESSOR_STATE_NEGCHANGE
- ICSF_MASTER_KEY_CONSISTENCY
- ICSFMIG_DEPRECATED_SERV_WARNINGS
- IOS_IORATE_MONITOR
- IOS_FABRIC_MONITOR
- RACF_AIM_STAGE
- RACF_CERTIFICATE_EXPIRATION
- RACF_UNIX_ID
- SUP_SYSTEM_SYMBOL_TABLE_SIZE
- SYSTRACE_MODE
- SYSTRACE_BRANCH
- OCE_XTIOT_CHECK
- USS_KERNEL_PVTSTG_THRESHOLD
- USS_KERNEL_STACKS_THRESHOLD
- VLF_MAXVIRT
- XCF_CF_STR_SCM_UTILIZATION
- XCF_CF_SCM_UTILIZATION
- XCF_CF_STR_MAXSCM
- XCF_CF_STR_MAXSPACE
- XCF_CF_STR_SCM_MAXSIZE
- ZOSMIGV2R2_Next_CS_SENDMAILDAEMN
- ZOSMIGV2R2_Next_CS_SENDMAILCLIEN
- ZOSMIGV2R2_Next_CS_SENDMAILMTA
- ZOSMIGV2R2_Next_CS_SENDMAILMSA
- ZOSMIGV2R2_Next_CS_SMTDPDAEMON
- ZOSMIGV2R2_Next_CS_SMTDPDMTA

The following Health Checks were changed by IBM in z/OS V2R1:

- ASM_LOCAL_SLOT_USAGE
- ASM_PLPA_COMMON_USAGE
- ASM_PLPA_COMMON_SIZE
- CATALOG_IMBED_REPLICATE
- RACF_classname_ACTIVE
- RACF_SENSITIVE_RESOURCES
- SLIP_PER
- VSM_CSA_LARGEST_FREE
- VSM_CSA_THRESHOLD
- VSM_SQA_THRESHOLD
- ZOSMIGV1R11_CS_DNSBIND



The following Health Checks were deleted by IBM in z/OS V2R1:

- CEE_USING_LE_PARMLIB
- PFA_FRAMES_AND_SLOTS_USAGE

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.

BCP Migration Actions for z/OS V2R2



Migration Actions Before Installing:

- **Use LOGR Couple Data Set at format level HBB7705 (Recommended)**
 - Recommended to use the highest LOGR CDS supported by the lowest system in your sysplex.
 - This allows the latest System Logger features to be available.
 - The highest level, HBB7705, was introduced in z/OS V1R2 (in 2001).
 - Use `D XCF, C, TYPE=LOGR` to see your LOGR CDS level.
 - Should see:


```
... ADDITIONAL INFORMATION:
LOGR COUPLE DATA SET FORMAT LEVEL: HBB7705 ...
```
- Pre-V2R2, specify `ITEM NAME(SMDUPLEX) NUMBER(1)` to format at HBB7705 level.

Migration Actions After IPLing:

- **Plan for new default LOGR CDS format level of HBB7705 in IXCL1DSU utility (Required-IF, as of V2R2)**
 - As of V2R2, there is a new default `NUMBER()` value for the `SMDUPLEX` item when formatting a LOGR CDS.
 - `NUMBER(1)` is the default, and results in an HBB7705-level CDS.


9

© 2015 IBM Corporation

BCP Migration Actions for z/OS V2R2



Migration Actions Before First IPL:

- **Format the ARM Couple Data Set for long symbol name support (Required-IF, as of V2R2)**

 - As of V2R2, if you want to use ARM functions, you must reformat the ARM CDS to the V2R2 level (HBB77A0).
 - Otherwise, V2R2 systems can join the sysplex, but are not ARM-capable
 - Use `D XCF, C, TYPE=ARM` to see your ARM CDS level.
 - Should see:


```
... ADDITIONAL INFORMATION:
FORMAT DATA
VERSION 1, HBB77A0 SYMBOL TABLE SUPPORT
...
```
- You must use the V2R2 level of IXCL1DSU format utility.
 - Either from a V2R2 system, or `STEPLIB` to V2R2 `MIGLIB`.

With APAR
OA46977
IBM Coexistence z/OS V2R2

10

© 2015 IBM Corporation

BCP Migration Actions for z/OS V2R2



Migration Actions Before First IPL:

- **Plan for HWIREXX helper program restriction for z/OS BCPii (Required-IF, as of V2R1 APAR OA45932)**
 - As of OA45932 (RSU1411), the BCPii helper program HWIREXX requires READ authority to the FACILITY class resource HWI.HWIREXX.execname.
 - HWIREXX provides the capability to easily run BCPii API calls from basic REXX execs, without the need to code an assembler program with an AXREXX macro invocation.
 - Otherwise, expect HWIREXX return code 112 (decimal).

- **Ensure PARMDD and REGIONX are not used as job statement symbols (Required-IF, as of V2R1 and V2R2)**
 - As of V2R1 (for PARMDD), and V2R2 (for REGIONX), there are new keywords on the JCL EXEC statement and PROC statement.
 - You must ensure that you do not use symbols with these same names. Change to use another symbolic name, or expect a JCL error.
 - Meaning, no PARMDD or REGIONX on the EXEC or PROC statement as symbols.
 - E.g. `//PROC1 PROC PARMDD=ABC, //JSTEP1 EXEC PROC1, REGIONX=ABC, or //STEP1 EXEC PGM=MYPROG, PARM=' &PARMDD'`

11

© 2015 IBM Corporation

BCP Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Convert your existing IBM Health Checker for z/OS set-up for automatic start-up (Required IF, as of V2R1)**
 - As of z/OS V2R1, Health Checker starts automatically at IPL!
 - First time user: follow steps in *IBM Health Checker for z/OS User's Guide*. (You may want to set all checks' WTOTYPE to HARDCOPY.)
 - Existing user:
 - Make sure your proc can be found at IPL (default HZSPROC), or specify the name in IEASYSxx HZSPROC=.
 - Remove any existing START HZSPROC (perhaps in COMMNDxx).
 - System rejects subsequent attempts to start Health Checker.
 - Recommended to specify your HZSPRMxx members in IEASYSxx HZS=(aa,bb,...) and use HZSPRM='PREV' on HZSPROC for easy restarts. HZSPROC HZSPRM='aa,bb,...' is still supported.
 - Optionally specify HZSPDATA(dsn) in HZSPRMxx.
 - Of course, WARNUND in IEASYSxx is supported for sharing.



12

© 2015 IBM Corporation

BCP Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Move from the console tracking facility to the Generic Tracker (Required IF, as of V2R1)**
 - As of z/OS V2R1, Generic Tracker facility (GTZ) replaces the console tracking facility. Lots of enhancements in the Generic Tracker!
 - Existing users of CNZTRKR service do not need to change.
 - GTZ will auto-start and initialize at IPL;


```
GTZ0105I  GENERIC TRACKER IPL-TIME INITIALIZATION STARTING
GTZ0106I  GENERIC TRACKER IPL-TIME INITIALIZATION COMPLETE
GTZ0001I  GENERIC TRACKER INITIALIZATION STARTING
GTZ0002I  GENERIC TRACKER INITIALIZATION COMPLETE
```
 - Userid associated with GTZ needs access to read parmlib concatenation data sets.
 - Although it is started, tracking is disabled.
 - Must replace any existing controls and commands for the console tracking facility, with Generic Tracker controls and commands:
 - Support in IEASYSxx GTZ= points to GTZPRMxx members. SETGTZ and SET GTZ commands available.
 - GTZPRINT utility and GTZQUERY service available.
 - Convert CNIDTRxx to GTZPRMxx. Sample tool provided!
 - Of course, WARNUND in IEASYSxx is supported for sharing.

13

© 2015 IBM Corporation

BCP Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Consider the new default value for the LOADxx DYNCPADD keyword (Required IF, as of V2R1)**
 - Before V2R1, LOADxx could enable CPUs to be added over the life of the IPL if the HW supported it (z10 or later). Default was for all CPUs, which was the minimum supported by a z/OS release (ex. 100) or the machine (ex. 80), and which could be asked for explicitly with DYNCPADD ENABLE.
 - As of z/OS V2R1, LOADxx keyword DYNCPADD now supports 1-4 char decimal value nnnn, for how many CPUs can be dynamically added. The default has changed to 16.
 - If LOADxx DYNCPADD is omitted (defaulted), the number of CPUs z/OS can dynamically add over the life of the IPL is 16.
 - If 16 is not sufficient, specify DYNCPADD nnnn. You cannot specify a value higher than the minimum of what z/OS or the HW supports.
 - If you specify DYNCPADD ENABLE, you will get maximum flexibility, but also will result in maximum storage usage and overhead.
 - z/OS V2R1 can support 256 CPUs, earlier releases supported 100.
 - LOADxx DYNCPADD nnnn cannot be shared with a pre-V2R1 system. That is the "nnnn" part does not coexist with pre-V2R1 systems.

14

© 2015 IBM Corporation

BCP Migration Actions For V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

BCP Migration Actions You Can Do Now



Plan to move from SHARED mode to DISTRIBUTED mode for consoles (Recommend, as of V2R2)

Not required, but recommended if you are using SHARED console mode. z/OS V2R2 is the last release in which SHARED console mode is supported.

You must migrate to DISTRIBUTED mode, which was introduced in z/OS V1R10.

Migration action: Move from SHARED mode to DISTRIBUTED mode for your console environment. Note that the default changed from SHARED to DISTRIBUTED mode in z/OS V1R13.

Tip: You can check the current mode by using the command **DISPLAY OPDATA,MODE**.

Use LOGR Couple Data Set at format level HBB7705 (Recommended, as of V2R2)

Not required, but recommended to allow for the latest system logger features to be available, given the sysplex configuration. If your LOGR couple data sets are already at HBB7705 level, this migration actions does not apply to you.

IBM recommends that you use the highest format level LOGR couple data set (CDS) that can be used by the lowest system release level in your sysplex. This will allow for the latest system logger features to be available, given the sysplex configuration. Currently, the highest LOGR CDS format level is HBB7705 (introduced in z/OS V1R2). This format level is established by providing the ITEM NAME(SMDUPLEX) NUMBER(1) specification in the IXCL1DSU couple data set format utility program.

Migration action:

1. Determine your current LOGR couple data set level. Use the **D XCF,C,TYPE=LOGR** command. If your LOGR couple data sets are at level HBB7705, you do not need to perform this migration action.
2. Use the IXCL1DSU (format couple data set utility) and include the ITEM NAME(SMDUPLEX) NUMBER(1) specification to obtain at least two LOGR CDSs at the HBB7705 format level, with:
 - LOGR CDS *primarydsname* on volume *primaryvolume*
 - LOGR CDS *alternatedsname* on volume *alternatevolume*
 - **Note:** SMDUPLEX item NUMBER(1) is the default value when you run the IXCL1DSU utility on z/OS V2R2.
3. After you create the HBB7705 format-level LOGR CDSs, you can dynamically bring them into your existing sysplex with these SETXCF commands:
 - SETXCF COUPLE,TYPE=LOGR,ACOUPLE=(*primarydsname,primaryvolume*)
 - SETXCF COUPLE,TYPE=LOGR,PSWITCH
 - SETXCF COUPLE,TYPE=LOGR,ACOUPLE=(*alternatedsname,alternatevolume*)System Logger does not allow the introduction of an alternate LOGR CDS that is formatted at a lower level than the primary.
4. Remember to also specify, in your COUPLExx member of SYS1.PARMLIB, these two LOGR CDSs as the primary and alternate for any future sysplex IPLs: DATA TYPE(LOGR) PCOUPLE(*primarydsname,primaryvolume*) ACOUPLE(*alternatedsname,alternatevolume*)

Note: If you did not bring the newly formatted HBB7705 LOGR CDSs into the sysplex (with the SETXCF commands in the third step above) prior to the first z/OS system that IPLs into the sysplex using the COUPLExx member identifying the newly formatted LOGR CDSs, then there will be no persistent logger data from before the IPL. Therefore, no log stream data exists when this first system IPLs.

Consider the new COUPLExx CFRMTAKEOVERCF(NO) default (Required-IF, as of V2R2)

Required if you are currently using CFRMOWNEDCFPROMPT(YES) or if, for some reason, the new behavior of CFRMTAKEOVERCF(NO) is not desirable.

The z/OS V2R2 coupling facility (CF) gain ownership processing enhancements introduce a new COUPLExx parmlib member keyword: CFRMTAKEOVERCF. Specifying CFRMTAKEOVERCF(NO) enables CF gain ownership processing enhancements that might prevent a sysplex outage by avoiding operator errors.

CFRMTAKEOVERCF(NO) is also the default for z/OS V2R2.

CFRMTAKEOVERCF(PROMPT) can be specified to get a z/OS V2R2 system to prompt the operator as it

Migrating to z/OS 2.2: Part 2 of 2

did in prior releases of z/OS.

When the CFRMOWNEDCFPROMPT(YES) is specified by the COUPLExx parmlib member of a down-level system (prior to z/OS V2R2), the z/OS V2R2 default of CFRMTAKEOVERCF(NO) is not compatible with the configuration.

When CFRMOWNEDCFPROMPT(YES) is used by a down-level system, that system will clear the CF authorities saved in the CFRM CDS during CFRM and sysplex initialization (that is, a sysplex-wide IPL). When that occurs, an up-level system (z/OS V2R2) will reject use of any CF that has a non-zero authority. However, CFRMTAKEOVERCF(NO) is not intended to reject the use of the CF when the old CF authority in the CFRM CDS matches the CF authority in the CF. If a down-level system is no longer in the sysplex, no system will perform the desired prompting.

When the CFRMOWNEDCFPROMPT(NO) is specified (or defaulted) by the COUPLExx parmlib member of a down-level system, the z/OS V2R2 default of CFRMTAKEOVERCF(NO) is compatible with the configuration.

Migration action: Follow these steps for CFRMOWNEDCFPROMPT(YES) in the COUPLExx parmlib member:

- Create a COUPLExx parmlib member for z/OS V2R2 systems with CFRMOWNEDCFPROMPT(YES) CFRMTAKEOVERCF(PROMPT) to obtain the old default behavior. After all of the systems are on z/OS V2R2, the COUPLExx parmlib member can be changed to CFRMOWNEDCFPROMPT(YES) CFRMTAKEOVERCF(NO) if the enhanced CF gain ownership processing of CFRMTAKEOVERCF(NO) is desired.
- To get a z/OS V2R2 system to prompt the operator as it did in prior releases of z/OS:
 - Copy the existing COUPLExx parmlib member into a new COUPLExx parmlib member.
 - Add the new CFRMTAKEOVERCF(PROMPT) statement after the COUPLE statement in the new COUPLExx parmlib member.
 - Ensure that the new COUPLExx parmlib member is used when IPLing z/OS V2R2.

If you specify or default to CFRMOWNEDCFPROMPT(NO), you have no migration action.

Update a Capacity Provisioning Manager parameter to avoid a defined capacity

WTOR (Required-IF, as of V2R2)

Required if you use Capacity Provisioning for managing defined capacity or group capacity and you want to avoid a WTOR for large manual reductions of defined capacity or group capacity.

In z/OS V2R2, the Provisioning Manager can detect when a manual change to defined capacity or group capacity would interfere significantly with Capacity Provisioning management. In such cases, the Provisioning Manager suspends its management of defined capacity or group capacity and issues one of the following write-to-operator-with-reply (WTOR) messages:

CPO4218I : New DC for *systemName/sysplexName*. Previous base *previousLimit* MSU. Enter 1 to set base to *currentLimit* or 2 to set to *newLimit* MSU

CPO4219I : New GC for *groupName/CPCname*. Previous base *previousLimit* MSU. Enter 1 to set base to *currentLimit* or 2 to set to *newLimit* MSU.

In response, the operator can choose to reinitialize the capacity management by setting the management base to the new capacity value and its managed capacity to 0, or continue the capacity management by adapting the management base to the manual change.

If you want the Provisioning Manager to continue managing the defined capacity or group capacity, regardless of manual changes, you can suppress the WTOR by setting the key **DefinedCapacity.BaseToleration** to 100 in the Capacity Provisioning Manager parameter file.

In z/OS V2R2, the default value of this key is 15, meaning that while Capacity Provisioning is managing a defined capacity, any concurrent manual reduction in defined capacity by more than 15% causes the WTOR to be issued.

Migration action: Follow these steps:

- Add the entry `DefinedCapacity.BaseToleration=100` to the Capacity Provisioning parameter file.
- By default, the parameter file is named `CPO.DOMAIN1.PARM(PARM)`.

Make accommodations for RACROUTE AUTH check for SLIP Command (Required-IF, as of V2R1)

Required if you use a SLIP command with Action of REFAFTER or REFBEFOR and are using a security product that is affected by the additional RACROUTE AUTH check.

As of z/OS V2R1, a RACROUTE AUTH check is now performed for a SLIP command that is issued with action of REFAFTER or REFBEFOR. This change might affect your installation, depending on the security product you are using.

Migration action: If you use the z/OS RACF security product, there is no action to take. If you use another security product, contact your vendor to see if there is any support or changes that you need to make.

Remove the REPORTCOMPLETIONS option from the IEAOPTxx member (Recommended, as of V2R1 and with the PTF for APAR OA35428 on z/OS V1R13)

Not required, recommended to avoid the information message, if you ever used the REPORTCOMPLETIONS parameter.

As of z/OS V1R13, the support for the REPORTCOMPLETIONS=YES|NO parameter in IEAOPTxx member of parmlib has been removed. This option was introduced with OA34801 (2010/11/24) and removed with OA35428 July 7, 2011. REPORTCOMPLETIONS=YES was intended to provide a system-wide option for installations with a single large production environment, where all regions were exempted from being managed towards response time goals and if you wanted to obtain response time statistics with different granularity through service and report classes. In all other cases the usage of this parameter was not recommended.

The REPORTCOMPLETIONS=YES option has been superseded by the "Manage Region Using Goal Of Both" option on the classification rule that allows to manage CICS Terminal Owning Regions (TORs) in a very granular way.

Migration action: To clean up the IEAOPTxx parmlib member, remove the REPORTCOMPLETIONS. You can still set the REPORTCOMPLETIONS option, but it is ignored with the following message:
IRA800I OPT MEMBER IEAOPTxx KEYWORD ReportCompletions IGNORED, NO LONGER USED

Move BCPii API calls into your application instead of in BCPii ENF exits (Required-IF, as of z/OS V1R13 and z/OS V1R12, both with APAR OA37035)

Required if you have coded a BCPii API call from your ENF exit.

As stated in various IBM publications, non-SRB ENF exits need to avoid time-consuming processing. Coding an HWIEVENT ENF exit to execute BCPii APIs may result in multiple problems, such as delays with BCPii event notification processing when BCPii services are simultaneously being invoked. Starting with z/OS V1R12 and later, BCPii enforces this restriction, and BCPii API calls made from within a BCPii ENF exit are now rejected with return code HWI_UNSUPPORTED_ENVIRONMENT.

Migration action: If you have coded a BCPii API call from your ENF exit, move the BCPii API call into your application and have the BCPii ENF exit post the application when the event occurs. Your application program may now issue the BCPii API call from the user's thread. For an example of how to code a BCPii ENF exit, see the sample ENF event exit HWIXMCX1 in SYS1.SAMPLIB.

BCP Migration Actions Pre-First IPL

Plan for HWIREXX helper program restriction for z/OS BCPii (Required-IF, as of z/OS V2R1 with APAR OA45932)

Required if you use the BCPii helper program HWIREXX.

Starting with z/OS V2R1, users of the z/OS BCPii System REXX helper program HWIREXX are required to have at least READ authority to the FACILITY class resource HWI.HWIREXX.execname as defined in the security product. This function is provided in APAR OA45932 with PTF UA75120.

Migration action: To allow you to run your BCPii System REXX exec using the HWIREXX helper program, you must have at least READ authority to the FACILITY class resource HWI.HWIREXX.execname, where execname specifies a 1 to 8 character System REXX exec to be executed by the HWIREXX helper application. Also, BCPii requires the FACILITY class to be RACLIST-specified. The RACF syntax is as follows:

```
RDEFINE FACILITY HWI.HWIREXX.execname UACC(NONE)
PERMIT HWI.HWIREXX.execname CLASS(FACILITY) ID(userid) ACCESS(READ)
SETROPTS RACLIST(FACILITY) REFRESH
```

If the caller does not have sufficient SAF authorization to run the HWIREXX program, HWIREXX return code 112 (in decimal) is returned.

Adjust parameters to start IOSHMCTL in TPC-R Basic Hyperswap environment (Required-IF, as of V2R1 and z/OS V1R13, both with APAR OA40866)

Required if you have unsupported parameters specified to start the Basic HyperSwap address space.

With APAR OA40866 applied (with one of the following PTFs: UA71220 for z/OS V2R1, or UA71219 for z/OS V1R13), parameters supplied in the JCL to start PGM=IOSHMCTL may be rejected. APAR OA40866 introduced new function which caused JCL parameter checking to now be enforced. As a result, unsupported parameters will be rejected and IOSHMCTL will not start due to JCL failures. For example:

Migrating to z/OS 2.2: Part 2 of 2

```
ASA101I SYNTAX ERROR: MODE=EXT, WAS SEEN, WHERE ONE OF
(SOCKPORT=)
WOULD BE CORRECT.
DETECTING MODULE IS IOSHMCTL
```

Prior to APAR OA40866, unsupported parameters would be ignored and IOSHMCTL would start without any JCL failures.

Note: This problem is a documentation error where some examples of JCL to start the Basic HyperSwap address space used incorrect parameters. At this time with Basic HyperSwap Sockets Server APAR OA40866 applied, the only parameter allowed to the IOSHMCTL program is 'SOCKPORT=*port_num*'. Any other parameters are rejected.

Migration action: When applying the PTF for APAR OA40866, make sure that any HSIB procedures to start IOSHMCTL do not include unsupported parameters. For example, the following JCL may be used:

```
//HSIB JOB MSGLEVEL=(1,1),TIME=NOLIMIT,REGION=0M
//IEFPROC EXEC PGM=IOSHMCTL
```

Tip: Some older publications, including Redpapers, may have documented incorrect parameters on the PARM statement used in the examples of suggested procedures for starting IOSHMCTL. It is best to refer to the latest TPC-R publications for current supported parameters.

Reference information: APAR OA46189, which describes the potential JCL failures after installing PTFs for APAR OA40866.



Format the ARM couple data set for long symbol table support (Required-IF, as of V2R2)

Required if systems are intended to use ARM functions.

z/OS V2R2 systems require an ARM couple data set (CDS) that is formatted for long symbol table support. Otherwise, V2R2 systems can join the sysplex, but are not ARM-capable. You can establish this format level by using the z/OS V2R2 level of the IXCL1DSU format utility (either from a z/OS V2R2 system or with a STEPLIB to a z/OS V2R2 MIGLIB) to format the CDS.

Systems at lower-level releases can use an ARM CDS formatted for long symbol-table support. However, a DISPLAY XCF,COUPLE,TYPE=ARM command from a down-level system cannot indicate that the HBB77A0 level of symbol table is in use unless you apply the PTFs for OA46977. With the PTFs for OA46977 applied, the response message IXC358I contains the text "HBB77A0 SYMBOL TABLE SUPPORT".

Migration action: Follow these steps:

1. Use the IXCL1DSU (format couple data set utility) at the z/OS V2R2 level to obtain at least two ARM CDSs at the HBB77A0 format level, with:
 - a. ARM CDS *primarydsname* on volume *primaryvolume*
 - b. ARM CDS *alternatedsname* on volume *alternatevolume*
2. After you create the HBB77A0 format-level ARM CDSs, you can dynamically bring them into your existing sysplex with these SETXCF commands:
 - a. SETXCF COUPLE,TYPE=ARM,ACOUPLE=(*primarydsname,primaryvolume*)
 - b. SETXCF COUPLE,TYPE=ARM,PSWITCH
 - c. SETXCF COUPLE,TYPE=ARM,ACOUPLE=(*alternatedsname,alternatevolume*)
3. Remember to also specify, in your COUPLExx member of SYS1.PARMLIB, these two ARM CDSs as the primary and alternate for any future sysplex IPLs : DATA TYPE(ARM)
PCOUPLE(*primarydsname,primaryvolume*) ACOUPLE(*alternatedsname,alternatevolume*)

Ensure that PARMDD or REGIONX are not used as job statement symbols (Required-IF, as of V2R2 and V2R1)

Required if you used symbol names REGIONX or PARMDD on EXEC or PROC statements in jobs.

z/OS V2R1 adds PARMDD and z/OS V2R2 adds REGIONX as new keywords on the JCL EXEC statement and PROC statement. Because JCL keyword names are reserved, you must ensure that your jobs do not use symbols with these same names. That is, if a job contains the symbolic parameter name PARMDD or REGIONX on the EXEC or PROC statement, you must edit the job to use an alternatively named symbol. Otherwise, the job can fail with a JCL error.

Migration action: Search for a symbol named PARMDD or REGIONX in all libraries that contain JCL, such as procedure libraries. Specifically, search for the following occurrences:

- PROC statements that contain a symbolic parameter named PARMDD or REGIONX.

Examples:

```
//PROC1 PROC PARMDD=ABC
//PROC1 PROC REGIONX=ABC
```

- EXEC statements that contain a symbolic parameter named PARMDD or REGION

Examples:

```
//JSTEP1 EXEC PROC1,PARMDD=ABC
//JSTEP1 EXEC PROC1,REGIONX=ABC
```

- EXEC statements that contain a '®IONX' PARM or '&PARMDD' parameter value string.

Examples:

```
//STEP1 EXEC PGM=MYPROG,PARM=' &PARMDD'
//STEP1 EXEC PGM=MYPROG,PARM=' &REGIONX'
```

For any occurrences that you find, change the PROC or EXEC statement to refer to another symbolic parameter name.

Assess the use of the INCLUDE1MAFC keyword on the LFAREA parameter in IEASYSxx (Required-IF, as of V2R2)

Required if you did not specify the INCLUDE1MAFC keyword on the LFAREA parameter in IEASYSxx prior to z/OS V2R2 and you want to continue excluding 1 M fixed frames from the available frame count.

Prior to z/OS V2R2, the default was for RSM not to include 1 M fixed frames in the available frame count (RCEAFC). APAR OA41968 added the INCLUDE1MAFC keyword on the LFAREA parameter in the IEASYSxx member of parmlib to cause RSM to include 1 M fixed frames in the available frame count. This results in the following system behavior:

- RSM performs less paging when there is an abundance of available 1 M fixed frames in the system.
- RSM is more likely to break up 1 M fixed frames to satisfy 4 K page demand. Although RSM attempts to coalesce broken up 1 M fixed frames when there is fixed 1 M page demand, there is no guarantee that coalescing will be successful, especially if any of the 4 K frames making up the fixed 1 M page is fixed long term.

Starting with z/OS V2R2, RSM includes 1 M fixed frames in the available frame count by default; you no longer need to specify INCLUDE1MAFC to do this. Furthermore, the syntax of the INCLUDE1MAFC keyword has been extended to include a YES or NO option. The former INCLUDE1MAFC keyword syntax is still valid and can now also be specified as INCLUDE1MAFC(YES).

IBM recommends that you accept the default (or specify INCLUDE1MAFC or INCLUDE1MAFC(YES)). However, if you have a need to retain the old default behavior (that is, to exclude 1 M fixed frames from the available frame count), you must specify INCLUDE1MAFC(NO).

RMF support is provided in RMF APAR OA42066 for z/OS V2R1, and RMF APAR OA42510 for z/OS V1R13.

Migration action:

- Check the LFAREA parameter specification in the IEASYSxx member on your pre-z/OS V2R2 system.
 - If you specified the INCLUDE1MAFC keyword on the LFAREA parameter after APAR OA41968, and you want to continue with that setting on your z/OS V2R2 system, you can do any of these things:
 - Leave the INCLUDE1MAFC keyword as it is.
 - Remove the INCLUDE1MAFC keyword, as it is now the default.
 - Change INCLUDE1MAFC to INCLUDE1MAFC(YES).
 - If you did not specify the INCLUDE1MAFC keyword on the LFAREA parameter and you want to retain the old default behavior, you must now specify INCLUDE1MAFC(NO).
- Check any application programs that use the STGTEST SYSEVENT to determine if any changes need to be made. The STGTEST event returns information about the amount of storage available in the system, which includes the LFAREA when INCLUDE1MAFC(YES) is specified or defaulted. Application programs can check the RCEINCLUDE1MAFC bit to determine the setting of INCLUDE1MAFC in the LFAREA specification.

Plan for the use of freemained frames (Required-IF, as of z/OS V2R2, z/OS V2R1 and z/OS V1R13 (with the RSM web deliverable), both with APAR OA46291)

Required unless you disable this feature. Otherwise, installations that use software tools that monitor real storage usage must install updates to accommodate the advent of freemained frames. Applications that invoke the TPROT instruction to determine whether pages of region private storage have been GETMAIN assigned should change to use the VSMLLOC or VSMLIST services. The IARQDUMP service may also be applicable in some cases. If none of these services meet the performance requirements of the application, then the application should use the new IARBRVER and IARBRVEA services provided with APAR OA46291 and z/OS V2R2.

To enhance system performance on the IBM z13, there might be cases where the system does not free the real frame that is backing a virtual page following a FREEMAIN, that is, when the page no longer contains any GETMAIN-assigned storage ranges. If so, the system will clear or "dirty" the frame to ensure that sensitive

Migrating to z/OS 2.2: Part 2 of 2

information is removed. Such a frame is referred to as a *freemained frame*. Freemained frames do not cause the count of frames owned by the address space (RAXFMCT) to be decremented (as they would have previously), nor do they cause the count of available frames within the system (RCEAFC) to be incremented (as they would have previously). Instead, the system uses a new counter, RAX_FREEMAINEDFRAMES, to keep track of the number of frames backing freemained pages in the address space with which the RAX is associated.

This feature is active by default on the IBM z13 and only applies to region private “low” storage (below 2GB), which is defined as subpools 0-127, 129-132, 240, 244, 250-252. Storage subpools define the characteristics of virtual storage below 2 GB and are discussed in detail in *z/OS MVS Diagnosis: Reference*.

For this entire migration action, see *z/OS V2.2 Migration*.

Relocate Cross System Extended Services (XES) component trace buffers (Required-IF, as of V2R2)

Required if you use coupling facilities in your sysplex or have references to the XES CTRACE CADS ('IXLTCAD') on the DSPNAME parameter of the **DUMP** and **SLIP** commands or in automated parse routines.

In z/OS V2R2, the Cross System Extended Services (XES) buffers for component tracing are moved from a common area data space (CADS) to a 4 GB memory object in 64-bit common high virtual (HVCOMMON) storage. During system initialization, XES obtains a 4 GB memory object and manages the virtual storage for global and connection CTRACE buffers. This change allows the GLOBAL trace buffer to be increased from 16 MB to 32 MB (fixed), which reduces the possibility of buffer wrapping. It also increases the available address range for connector trace buffers, which decreases the possibility of a connector running without component tracing. In previous releases, the XES CTRACE buffers resided in a CADS object named IXLTCAD, which limited the buffers to a 2 GB range of addresses. In z/OS V2R2, XES no longer creates the IXLTCAD object.

Notes:

1. The 4 GB memory object is a fixed size area that is obtained by XES; the size cannot be modified.
2. The IXLCBCAD object is not affected by this migration action.
3. Eliminating the XES CADS decreases the number of common area data spaces that are created in the system.

Migration action:

- Ensure that enough 64-bit common storage (HVCOMMON) storage is allocated by the system, so that the additional 4 GB request by XES does not cause shortages for other components and elements. The amount of 64-bit common storage is controlled by the HVCOMMON parameter in the IEASYSxx parmlib member. Review the value that is specified on the HVCOMMON parameter to determine whether it must be increased. You can use the MVS operator command **D VIRTSTOR,HVCOMMON** to display information about the current use of the HVCOMMON storage on your system. For example:

```
IAR019I 06.55.51 DISPLAY VIRTSTOR
SOURCE = DEFAULT
TOTAL 64-BIT COMMON = 66G
64-BIT COMMON RANGE = 1982G-2048G
64-BIT COMMON ALLOCATED = 4171M
```

- To accommodate the allocation of a 4 GB XES CTRACE buffer, add 4 gigabytes (4G) to the HVCOMMON value in the IEASYSxx parmlib member.
- Check for references to the IXLTCAD object, which is no longer created in z/OS V2R2. Specifically, check for references to 'IXLTCAD' on the DSPNAME parameter of the **DUMP** and **SLIP** commands (that is, DSPNAME=('XCFAS'.IXLTCAD)) and on any automated parse routines.
- Ensure that SDATA=XESDATA is specified on any **DUMP** or **SLIP** commands where the IXLTCAD name was removed. This setting causes the XES CTRACE 64-bit common storage to be included in an SVC dump.

Failure to remove the IXLTCAD references can result in an error message, such as ASA104. This error, however, does not stop the running process. If XES cannot obtain a 4 GB memory object, message IXL017I is issued. The system continues to process XES requests normally, but SYSXES CTRACE data is not be available in dumps for analysis under IPCS.

MVS SLIP command change for MODE=HOME (Required-IF, as of V2R2)

Required if you use any IEASLPxx parmlib member or program that issues a SLIP command for a PER trap with MODE=HOME specified, but without JOBNAME and ASID specified.

When a **SLIP** command is issued for a PER trap with **MODE=HOME** specified, but without **JOBNAME** and **ASID** specified, the system issues message IEE088D to prompt the operator for a reply (continue or cancel).

The **JOBNAME**, **ASID**, and **MODE** parameters for a **SLIP** command work together to control the cross memory environment for a PER interrupt. Specifying **MODE=HOME** indicates that the SLIP trap matches only when the command is running in the home address space. When **MODE=HOME** is specified, more processing is used to monitor for the correct environment. When **MODE=HOME** is specified with the **JOBNAME** or **ASID** parameter, this monitoring is limited to the specified address spaces.

However, if **MODE=HOME** is specified without **JOBNAME** and **ASID**, the monitoring occurs for every address space on the system, which can impact performance significantly. As a warning, the system issues message IEE088D with either of the following text strings:

```
IEE088D SLIP COMMAND SPECIFIES MODE=HOME WITHOUT JOBNAME AND ASID. SPECIFY 'OK' OR 'CANCEL'
```

```
IEE088D SLIP ID=xxxx SPECIFIES MODE=HOME WITHOUT JOBNAME AND ASID. SPECIFY 'OK' OR 'CANCEL'
```

Note: When this capability was introduced in z/OS V2R1 with APAR OA45297 and even after APAR OA45912, you enabled it by specifying the option TRAPS NAME(IeaSlipConfirm) in the DIAGxx parmlib member. However, in z/OS V2R2, this capability is enabled automatically; you no longer need to specify it explicitly. If you do so in z/OS V2R2, the explicit specification is ignored without an error.

Migration action: For any IEASLPxx parmlib member or program that issues a **SLIP** command for a PER trap with **MODE=HOME** specified, but without **JOBNAME** and **ASID** specified, verify your use of **MODE=HOME** to determine whether changes are needed. Follow these steps:

- If a **SLIP** command includes **MODE=HOME**, ensure that it also includes **JOBNAME** or **ASID**, or both. Or, if you require the SLIP trap to be set this way, include the **OK** parameter to avoid message IEE088D.
- If you already include the **OK** parameter on the **SLIP** command for another purpose, verify that the **MODE=HOME** parameter is required. Previously, you might have specified the **OK** parameter to avoid receiving the following message:

```
IEE831D SLIP TRAP ID=0001 PER RANGE (00000000_00000000,00000000_01000000) EXCEEDS 1M OR WRAPS. SPECIFY 'OK' OR 'CANCEL'
```

For example, if you have a SLIP trap that is defined as follows, verify that **MODE=HOME** is required:

```
SLIP SET,IF,RA=(0,1000000),MODE=HOME,OK,END
```



Accommodate the SETLOAD xx,IEASYM command to update system symbols without initiating an IPL (Required-IF, as of V2R1)

Required if you are currently using the IEASYMUP module provided in SYS1.SAMPLIB or the SYMUPDTE routine to update system symbols.

Before z/OS V2R1, the downloadable SYMUPDTE routine and the IEASYMUP module in SYS1.SAMPLIB were provided as mechanisms to update system symbols without initiating an IPL. Starting with z/OS V2R1, the SETLOAD xx,IEASYM command is available to perform this task. In z/OS V2R1, the IEASYMUP module in SAMPLIB is updated to return with a RC=X'FFF', not having done the requested function. However, unless this IEASYMUP module is rebound, there is no way to prevent the usage of an old copy, or detect an update because of use of an old copy of the tool.

In z/OS V2R1 you should stop using the downloadable SYMUPDTE routine or the IEASYMUP module from samplib in your earlier release. Note that use of SYMUPDTE or IEASYMUP might produce incorrect results when used in conjunction with the SETLOAD xx,IEASYM command. The SYS1.LINKLIB program IEASYMU2 is instead provided via APAR OA42569 as a replacement for the function provided by IEASYMUP / SYMUPDTE on previous z/OS releases. IEASYMU2 is a supported program in z/OS, and has considerations when used with SETLOAD xx, IEASYM.

Notes:

- If you continue to use the SYMUPDTE routine or the IEASYMUP module, or if you use the IEASYMU2 program, system symbols will be updated and the system will continue to function properly until the SETLOAD xx,IEASYM command is used. Once SETLOAD xx,IEASYM has been issued, updates introduced by those other programs will be lost. You may continue to use the IEASYMU2 program again, even after a SETLOAD xx,IEASYM has been issued, but the same considerations will be in effect.
- Updates to the system symbol table through the SETLOAD xx,IEASYM command will occur only on the system issuing the command.
- When a SETLOAD command is issued and fails, messages issued by IEFPRMLB (logical parmlib service) that contain *jobname* and *stepname* will contain the master *jobname* and *stepname* of the last step that ran under master because the SETLOAD command runs under master but does not run as its own step. In this

case, the error is related to the SETLOAD processing and NOT the step whose name appears in the message.

Migration action:

- Rebind the IEASYMUP module from the z/OS V2R1 SAMPLIB to disable the code or simply remove it from LINKLIB or your LNKLST library.
- If you have used the downloadable SYMUPDTE routine, remove it from your LINKLIB or your LNKLST library. Begin using SETLOAD xx,IEASYM command instead of these obsolete modules. Or change your JCL to use IEASYMU2 instead of IEASYMUP (and remove any joblib/steplib specification). IEASYMU2 verifies access through the same profile of IEASYMUP.* in the FACILITY class that IEASYMUP did, so there are no security definition changes from using IEASYMUP to IEASYMU2.

Review the list of WTORs in parmlib member AUTOR00 (Required)

In z/OS V1R12, the DDDEF'd PARMLIB provides an AUTOR00 member. This member should be found in your parmlib concatenation during IPL and will result in auto-reply processing being activated. If the WTORs listed in AUTOR00 are automated by your existing automation product, ensure that the replies in AUTOR00 are appropriate.

Migration action: Examine the WTOR replies in the AUTOR00 parmlib member. If the replies or delay duration are not desirable, you can create a new AUTORxx parmlib member and make corresponding changes. Also compare the replies to what your automation product would reply to these WTORs. Make sure that the AUTOR00 replies are in accordance with the replies from your automation product. IBM does not recommend making updates to AUTOR00, because updates to AUTOR00 might be made by the service stream or in new z/OS releases.

Move from the console tracking facility to the Generic Tracker (Required-IF, as of V2R1)

Required if you have started or customized the console tracking facility.

Before z/OS V2R1, the console tracking facility was used to track use of system functions; however, it had some limitations. Starting with z/OS V2R1, the console tracking facility is being replaced with a new enhanced Generic Tracker. As of z/OS V2R1, support for the console tracking facility parmlib member, CNIDTRxx, is no longer provided. The Generic Tracker accepts GTZPRMxx parmlib members, and a tool to convert CNIDTRxx to the Generic Tracker GTZPRMxx is also provided.

It is assumed that the retired console tracking facility commands are typically issued manually and for diagnostic purposes only, not in any scripts or automation tools. The callable service CNZTRKR, compiled into executable programs, continues to be supported and the internal routing of any data passed to CNZTRKR and the Generic Tracker will be transparent to the existing CNZTRKR caller.

Migration action:

- Replace the use of any console tracking facility commands. Use COMMNDxx, automation scripts, or manually enter the commands on the console command line with their corresponding Generic Tracker (GTZ) counterparts. You can use the following mapping as a quick reference:
 - Instead of using COMMNDxx to start the console tracking facility, use the new system parameter GTZ (in IEASYSxx) to specify a GTZPRMxx member that specifies the SETGTZ TRACKING=ON command.
 - Instead of the DISPLAY command, consider using utility GTZPRINT or a user-written program with the service GTZQUERY to retrieve, store, and process current tracking data.
 - Instead of SETCON TRACKING={ON|OFF}, use SETGTZ TRACKING={ON|OFF}
 - Instead of SETCON TRACKING=ONWITHABEND, use SETGTZ DEBUG(ACTION=ABEND...)
 - Instead of DISPLAY OPDATA,TRACKING, use DISPLAY GTZ command, the GTZPRINT tool, or the GTZQUERY macro service.
 - Instead of SET CNIDTR=xx, use SET GTZ=xx or system parameter GTZ in IEASYSxx
- Instead of having any SETGTZ commands in COMMNDxx, consider putting them into GTZPRMxx parmlib members. You can use the SET GTZ command or the GTZ system parameter in IEASYSxx to select and execute the content of those GTZPRMxx.
- For any new applications use macro GTZTRACK instead of macro CNZTRKR. Consider converting any existing use of CNZTRKR to GTZTRACK.
- Convert existing CNIDTRxx parmlib members to GTZPRMxx. See the sample GTZCNIDJ for how the GTZCNIDT conversion tool can help you automate this conversion.
- Associate a user ID with the GTZ address space and grant that user ID permission to access the PARMLIB concatenation. This is required when using new system parameter GTZ to identify GTZPRMxx parmlib members to be processed at GTZ startup. See the description of system parameter GTZ in *MVS Initialization and Tuning Reference* for details.



Convert your existing IBM Health Checker for z/OS set-up for automatic start-up (Required-IF, as of V2R1)

Required if you are currently using the IBM Health Checker for z/OS and wish to continue to use it as you have customized it. This migration action is strongly recommended for those that have not used the IBM Health Checker for z/OS on each system yet.

Before z/OS V2R1, IBM Health Checker for z/OS users had to perform a set-up procedure and start IBM Health Checker for z/OS manually. As of z/OS V2R1 the system starts IBM Health Checker for z/OS automatically. If you haven't started IBM Health Checker for z/OS before, you will probably see program exceptions. See the *IBM Health Checker for z/OS: User's Guide* for how to handle those exceptions.

Migration action: For first time users of IBM Health Checker for z/OS, follow the steps for Optimizing IBM Health Checker for z/OS in *IBM Health Checker for z/OS: User's Guide*.

For users with existing IBM Health Checker for z/OS set-ups, use the following migration actions to convert systems to the IBM Health Checker for z/OS automatic start-up:

1. Make sure the system knows the name of your hzsproc procedure if you renamed it from the default HZSPROC: The start-up procedure for IBM Health Checker for z/OS is called HZSPROC, by default. If you customized your hzsproc name, you must specify it to the system, using the new HZSPROC system parameter in IEASYSxx.



2. Remove any existing START HZSPROC invocations that start IBM Health Checker for z/OS and rely on the automatic start-up: Because IBM Health Checker for z/OS now starts automatically, you must look for instances of START HZSPROC invocations in your installation. For example, look for START HZSPROC invocations in the COMMNDxx, parmlib member. After the system brings IBM Health Checker for z/OS up automatically, the system rejects any attempt to bring up a second instance. The system issues one of the two following warning messages:

```
HZS0101I - "...HEALTH CHECKER... IS ALREADY ACTIVE"
```

or if the automatically started IPL-time instance is still initializing:

```
HZS0116I - "...HEALTH CHECKER... START PENDING"
```

3. Change the way you specify the HZSPRMxx parmlib members you want the system to use. Before z/OS V2R1, users typically specified the HZSPRMxx parmlib members for IBM Health Checker for z/OS in the HZSPROC procedure. Now starting with z/OS V2R1, IBM recommends that you do the following to tell the system which members of HZSPRMxx to use:
 - Specify the HZSPRMxx parmlib members for your installation in the new HZS system parameter of IEASYSxx. This provides the default for the automatic start of IBM Health Checker for z/OS at IPL-time.
 - In your hzsproc procedure, default to or define HZSPRM='PREV':


```
//HZSPROC PROC HZSPRM='PREV'
//HZSSTEP EXEC PGM=HZSINIT,REGION=0K,TIME=NOLIMIT,
// PARM='SET PARMLIB=&HZSPRM'
//*HZSPDATA DD DSN=SYS1.&SYSNAME..HZSPDATA,DISP=OLD
// PEND
// EXEC HZSPROC
```
 - HZSPRM='PREV' specifies the following:
 - For the initial automatic start, the system will use the HZSPRMxx suffixes listed in the HZS system parameter.
 - For manual restarts after the initial automatic start, IBM Health Checker for z/OS initially uses the HZSPRMxx parmlib members that were in effect just before the previous Health Checker instance was stopped. This action will in particular include any parmlib members specified through a MODIFY HZSPROC,ADD,PARMLIB or MODIFY HZSPROC,REPLACE,PARMLIB command, while this first instance was running.

For example, assume HZSPRM=PREV was specified when that first instance was started and system parameter HZS was set to (00,01). Then this first instance would have initially used HZSPRM00 and HZSPRM01. Now assume a MODIFY HZSPROC,ADD,PARMLIB=(02,03) was specified and then later this first instance is stopped. A manual restart, still with HZSPRM=PREV, will initially use HZSPRM00, HZSPMR01, HZSPRM02, and HZSPRM03, as in the previous instance before it was stopped.

Migrating to z/OS 2.2: Part 2 of 2

If MODIFY HZSPROC,REPLACE,PARMLIB=(02,03) is used instead, the secondary instance initially only uses HZSPRM02 and HZSPRM03. Specifying HZSPRM='PREV' makes occasional manual restarts (after applying service, for example) easy and consistent.

- Optionally specify an HZSPDATA data set for persistent data in the HZSPRMxx parmlib member: Before z/OS V2R1, you could only specify the HZSPDATA in the HZSPROC startup procedure. Now you can define your HZSPDATA data set in either the HZSPROC startup procedure or on the HZSPDATA parameter of the HZSPRMxx parmlib member.

Consider the new default value for the LOADxx DYNCPADD keyword (Required-IF, as of V2R1)

Required if the LOADxx DYNCPADD value is omitted (defaulted) and the default number of CPUs (16) z/OS will be able to dynamically add over the life of the IPL is not sufficient. Recommended, if you specify DYNCPADD ENABLE, which provides maximum flexibility but also results in maximum storage usage and overhead.

Before z/OS V2R1, through PARMLIB member LOADxx you could enable that CPUs be added to the configuration over the life of the IPL if the hardware supported such addition. The default was for all CPUs that could be configured to the LPAR, which was the minimum supported by the z/OS release (for example, 100) and the machine (for example, 80), and which could be asked for explicitly by DYNCPADD ENABLE.

In z/OS V2R1, the LOADxx keyword DYNCPADD now supports a 1-4 character decimal value nnnn that indicates how many CPUs z/OS is able to dynamically add over the life of the IPL. The default has changed to 16.

Notes:

- All system z hardware (z10 EC/BC and later hardware) supports dynamic CPU addition.
- When specifying the maximum number of CPs that z/OS can dynamically add with LOADxx DYNCPADD nnnn, this LOADxx cannot be shared with pre-V2R1 systems; that is, the nnnn parameter of DYNCPADD is not recognized by pre-V2R1 systems.
- Specifying DYNCPADD nnnn or taking the DYNCPADD 16 default allows z/OS to determine the number of CPUs that z/OS must be prepared to be dynamically added for the life of the IPL. Because z/OS can know the maximum CPU id that can be dynamically added for the life of the IPL at IPL, z/OS can obtain CPU array related storage based on the maximum number of CPUs that can be activated for the life of the IPL.

Migration action: If the default limit of 16 CPUs that can be dynamically added is not sufficient, then indicate on your LOADxx DYNCPADD the number you desire. The maximum number of CPUs that can be added over the life of the IPL will be capped by the minimum between the highest CPU id hardware and the z/OS release supports. The DYNCPADD setting cannot be changed dynamically; it requires an IPL.

Review the changed messages associated with two digit CP IDs. Update any necessary automation or operator procedures to accommodate the two digits. Before z/OS V2R1, there was only one digit used for CP IDs.

Note: If you are interesting in how the DYNCPADD controls the extent and ability to add CPs in association with your HMC setting, see the Redbook <http://www.redbooks.ibm.com/abstracts/sq248178.html?Open> (section 4.11, search on "dyncpadd" to find it directly). In short: DYNCPADD controls the extent and ability of adding CPs. It says the number of CPs you can add *above* what is the highest CPU ID at IPL to make the *new* limit. Of course, that number cannot exceed what is supported on z/OS or the HW.

Plan for the increase of the maximum number of supported CPUs to 256 (Required-IF, as of V2R1)

Required for programs using local constants or the z/OS constants for the number of CPUs the z/OS CPU infrastructure supports at the current or a specific level of z/OS

In z/OS V2R1, z/OS CPU infrastructure will support up to a maximum of 256 CPUs (CPU ids 0-255). Earlier releases of z/OS support up to 100 CPUs (CPU ids 0-99). Components or products allocating storage for CPU related arrays or bitmasks might require changes to support the V2R1 CPU infrastructure. Allocating CPU related arrays or bitmasks on a per CPU basis is done using one of the following:

- Run-time fields in the z/OS CVT (mapped by CVT) and ECVT (mapped by IHAECVT) control blocks representing the maximum CPU id a z/OS image can use for the life of the IPL. Products using run time fields will not require changes to support the V2R1 CPU infrastructure.
- Compile-time or assemble-time constants in the z/OS ECVT control block or within the product itself representing the maximum CPU id the z/OS CPU infrastructure supports. Products using compile-time or assemble-time constants will need to recompile at a minimum and may require code changes to support the V2R1 CPU infrastructure.

All products running on z/OS V2R1 must prepare to support all CPUs supported by the z/OS V2R1 CPU infrastructure (up to 256 CPUs with CPU ids 0-255). Products that support the z/OS V2R1 CPU infrastructure will be able to run on earlier z/OS releases whose CPU infrastructure supports a smaller number of CPUs.

Programs that do not support up to the maximum number of CPUs the z/OS infrastructure supports might not be able to work with all CPUs on the z/OS image. The system impacts are program dependent.

Plan for the new default TRACKDIRLOAD in PROGxx (Recommended, as of z/OS V2R1)

Recommended, not required, if you used the old default value of NOTRACKDIRLOAD and you want to retain it, you must explicitly specify NOTRACKDIRLOAD in PROGxx.

Starting with z/OS V2R1 the default in PROGxx has changed from NOTRACKDIRLOAD to TRACKDIRLOAD. With TRACKDIRLOAD, internal system processing can be enhanced. Using TRACKDIRLOAD might cause extra but generally beneficial overall processing.

Migration action: If you want to use the old default value, specify NOTRACKDIRLOAD in PROGxx.

Plan for new entries AXRINIT and AXRRXTSS in the program properties table (Recommended, as of V2R1)

Recommended to avoid an unintended override of the IBM shipped default PPT.

In z/OS V1R13 and earlier there were no entries in the program properties table for AXRINIT and AXRRXTSS to indicate that these programs needed to run privileged; system programmers had to manually add entries for AXRINIT and AXRRXTSS into SCHEDxx parmlib members. In z/OS V2R1, these entries are now included in the IBM supplied default program properties table, and you can remove the SCHEDxx specifications for AXRINIT and AXRRXTSS.

Migration action: Remove the specifications of AXRINIT and AXRRXTSS in SCHEDxx. The recommended action described in DOC APAR OA40519 is no longer needed in z/OS V2R1.

Plan for security changes to EXECIO restricting the REXX exec for allocating an internal reader (Required-IF, as of V2R1)

Required if the invoker of the System REXX exec wants to invoke EXECIO to submit JCL.

In z/OS V1R13 and earlier for a REXX exec that was running under System REXX (TSO=YES), the exec was able to allocate an internal reader and subsequently invoke EXECIO to submit JCL. As of z/OS V2R1, this function is restricted if the security product (RACF or equivalent) indicates that the invoker does not have authority to the entity JCL.

Migration action: Permit access to allow the System REXX exec that uses EXECIO to submit JCL for allocating an internal reader. The System REXX exec runs under the security environment as specified by the SECURITY keyword on the AXREXX invocation; the default is the invoker of the AXREXX macro. The invoker of the System REXX exec must have access to the JCL resource in the TSOAUTH resource class.



Accommodate increase in nucleus size for the z/OS V1R13 RSM Enablement Offering Web deliverable (Required-IF, as of R13 RSM Enablement Offering Web deliverable and V2R1)

Required if the increase in nucleus size will affect the private storage available to applications on your system above the 16MB line. The migration action is relevant on all servers.

With the introduction of the z/OS V1R13 RSM Enablement Offering Web deliverable, the nucleus will increase in size by approximately 380K above the 16MB line. Because of this increase, the available private storage might affect the way your applications run on your system.

Migration action: Review your current available private storage usage above the 16MB line using reports from RMF or an equivalent product. Ensure that an increase of 380K for the nucleus above the 16 MB line will not adversely affect your system. Adjust values accordingly.



Accommodate increase in ESQA with the z/OS V1R13 RSM Enablement Offering Web deliverable (Required-IF, as of R13 RSM Enablement Offering Web deliverable and V2R1)

Required if the increased allocation in ESQA will affect the private storage available to applications on your system. The migration action is relevant on all servers.

With the introduction of the z/OS V1R13 RSM Enablement Offering Web deliverable, there is an increased allocation of 24K (6 pages) in ESQA per CPU per LPAR. This increased allocation in ESQA per CPU includes general purpose CPs, zIIPs, and zAAPs regardless of the status (online, offline, configured, stand-by, and so forth) of the hardware. Because of this increase in ESQA, if you change the ESQA size, the available private storage may affect application execution on your system.

Migration action: Review your ESQA specification in IEASYSxx, to ensure that an ESQA increased allocation of 24K per CPU used on the LPAR will not adversely affect your system. If you need to increase your ESQA specification, you should also review the effects on your current available private storage usage above the 16 MB line using reports from RMF or an equivalent product. Adjust values accordingly.

BCP Migration Actions Post-First IPL

Plan for the new default format level of HBB7705 in the IXCL1DSU utility (Required-IF, as of V2R2)

Required if you intend on using the IXCL1DSU format couple data set utility on any z/OS V2R2 level systems.

Starting in z/OS V2R2, system logger uses a new default NUMBER() value for the SMDUPLEX item when using the XCF couple data set format utility (IXCL1DSU) to format a LOGR CDS. The SMDUPLEX item NUMBER(1) is now the default value when running the IXCL1DSU utility on a z/OS V2R2 or higher release level and results in a LOGR CDS format level of HBB7705.

The LOGR CDS format level is managed as follows:

HBB7705 : Indicates that the LOGR CDS was formatted at one of the following sets of release level ranges and options:

- z/OS V2R2 or later when the ITEM NAME(SMDUPLEX) specification is not provided
- z/OS V1R2 or later with option NUMBER(1) specified for item name SMDUPLEX

HBB6603 : Indicates that the LOGR CDS was formatted at one of the following sets of release level ranges and options:

- OS/390 V1R3 or later, through OS/390 V2R10 and z/OS V1R1
- z/OS V1R2 or later, through z/OS V2R1, without option NUMBER(1) for item name SMDUPLEX
- z/OS V1R2 or later with option NUMBER(0) specified for item name SMDUPLEX

HBB5520 : Indicates that the LOGR CDS was formatted at a release level before OS/390 V1R3.

Migration action: Plan to make use of the default HBB7705 LOGR CDS format level or take steps to format a different LOGR CDS level.

Communications Server Migration Actions for z/OS V2R2



Migration Actions Before First IPL:

- **IP Services: Modify GLOBALCONFIG SMCR PFID definitions (Required IF, as of V2R2 on V2R1 path only)**
 - At V2R1, multiple TCP/IP stacks in the same LPAR could share a RoCE Express feature by configuring and activating the same PCIe function ID (PFID) of the feature.
 - If using SMC-R* -and- using a 10GbE RoCE Express feature in a dedicated RoCE environment -and- are on a z13, then each TCP/IP stack must have a unique PFID value to represent the RoCE Express feature
 - Use HCD to define a unique FID for each TCP/IP stack.
 - Modify the GLOBALCONFIG SMCR statement to specify that PFID value.
 - * Shared Memory Communications via Remote Direct Memory Access
- **IP Services: Use the new maximum segment size adjustments if required (Required IF, as of V2R2)**
 - As of V2R2, new function automatically adjusts of the Maximum Segment Size (MSS) that is used on a TCP connection to avoid fragmentation. This new function is turned on by default.
 - Use the auto-adjusting function by default, or disable the function.
 - A performance degradation can result if the MSS is not configured to avoid fragmentation.
 - If you decide to disable the function, specify the subparameter ADJUSTDVIPAMSS NONE on the GLOBALCONFIG statement.
 - AUTO is the default.

15

© 2015 IBM Corporation

Communications Server Migration Actions for z/OS V2R2



Migration Actions Before First IPL:

- **IP Services: Verify the new default for the QUEUEDRTT parameter (Required IF, as of V2R2)**
 - As of V2R2, there are new outbound serialization enhancements, controlled by the QUEUEDRTT parameter on the TCPCONFIG profile statement.
 - Default has changed from 20 to 0.
 - 20: only TCP/IP connections with a round-trip time (RTT) of 20 milliseconds or more are eligible to use outbound serialization.
 - 0: all connections are eligible to use outbound serialization.
 - Use new default, or update QUEUEDRTT as desired. Range is 0-50.
- **IP Services: Decide whether to accept the new FIXED CSM default (Required IF, as of V2R2)**
 - As of V2R2, communications storage manager (CSM) fixed storage for buffers is increased from 100 MB to 200 MB.
 - Use new default, or update FIXED MAX(...) in your IVTPRM00 parmlib member.
 - Use D NET, CSM to look at the FIXED MAXIMUM (on message IVT5538I):

```

IVT5536I TOTAL ALL SOURCES          21376K      8040K      29416K
IVT5538I FIXED MAXIMUM =          120M  FIXED CURRENT =    26969K
IVT5541I FIXED MAXIMUM USED =          27097K SINCE LAST DISPLAY CSM ...
    
```

16

© 2015 IBM Corporation

Communications Server Migration Actions for z/OS V2R2



Migration Actions Before First IPL:

- **IP Services: Check automation on IKE daemon syslogd messages (Required IF, as of V2R2)**
 - As of V2R2, all IKED messages written through syslogd contain a thread identifier, preceding the message ID. Messages from different IKED threads may be interleaved.
 - Verify automation.
 - Example at z/OS V2R1:

```
Jul 28 11:39:26 mvs046 IKE: E2D1061I IKE connecting to PAGENT
Jul 28 11:39:26 mvs046 IKE: E2D1062A IKE retrying connection to PAGENT
Jul 28 11:39:34 mvs046 IKE: E2D0923I IKE has received the STOP command
Jul 28 11:39:34 mvs046 IKE: Message instance 0: E2D0967I IKE release
CS V2R1 Service Level CS130924 Created on Sep 24 2013
```

- **Example at z/OS V2R2:**

```
Jul 28 15:10:47 mvs046 IKE: (00000001) E2D1061I IKE connecting to PAGENT
Jul 28 15:10:47 mvs046 IKE: (00000001) E2D1062A IKE retrying connection
to PAGENT
Jul 28 15:11:06 mvs046 IKE: (00000003) E2D0923I IKE has received the STOP
command
Jul 28 15:11:06 mvs046 IKE: Message instance 0: (00000001) E2D0967I IKE
release
17 CS V2R2 Service Level CS140728 Created on Jul 28 2014
```

© 2015 IBM Corporation

Communications Server Migration Actions for z/OS V2R2



Migration Actions You Can Do NOW:

- **IP Services: Be aware of IP Fragment IDS enhancements (Require V2R1)**
 - Prior to V2R1, When IP Fragment IDS attack type is enabled, only IPv4 traffic was monitored.
 - As of V2R1, IP Fragment attack type of the IDS will monitor both the IPv6 and IPv4 traffic for suspicious fragments. It also will check for overlays that change the data in the packet.
- **IP Services: Prepare for the additional of IPv6 support for policy-based routing (Required IF, as of V2R1)**
 - Prior to V2R1, if you are using policy-based routine on a stack running as dual-mode (IPv4 and IPv6) and didn't specify the source and destination IP address, it only applied to IPv4 packets.
 - As of V2R1, that same rule applies to both IPv4 and IPv6.
 - If you want it to continue to apply to only IPv4 packets, change the rule to have either the source or destination IP address be 0.0.0.0/0.

18

© 2015 IBM Corporation

Communications Server Migration Actions for z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

Communications Server Migration Actions You Can Do Now

IP Services: Understand the change in the support provided by the DVIPSEC parameter on the IPSEC statement in the TCP/IP profile (Required-IF, as of V2R1)

Required if the TCP/IP profile for your stack specifies the IPSECURITY parameter on the IPCONFIG6 statement and also specifies the DVIPSEC parameter on the IPSEC statement.

Before z/OS V2R1, the DVIPSEC parameter on the IPSEC statement in the TCP/IP profile enabled Sysplex-Wide Security Associations (SWSA) for IPv4 on a stack that had IPCONFIG IPSECURITY specified in the TCP/IP profile. Support for SWSA for IPv6 was not provided in these releases. Beginning with z/OS V2R1, SWSA for IPv6 is supported. The DVIPSEC parameter on the IPSEC statement in the TCP/IP profile enables SWSA for IPv6 on a stack that has IPCONFIG6 IPSECURITY specified in the TCP/IP profile.

Migration action:

- If you have both of the following specified in your TCP/IP profile, be aware that SWSA for IPv6 will be enabled on your stack:
 - The IPSECURITY parameter on the IPCONFIG6 statement
 - The DVIPSEC parameter on the IPSEC statement
- If you have IPv6 TCP traffic that is protected by an IPsec Security Association (SA) with an IPv6 DVIPA endpoint, you can see the following changes:
 - When an IPv6 DVIPA is moved during a planned or unplanned DVIPA takeover, new SAs are automatically reestablished with the same security service characteristics as the SAs that existed on the host that owned the DVIPA.
 - IPv6 TCP traffic that is protected by an IPsec SA with a sysplex-distributed DVIPA endpoint can be distributed to target hosts.
 - Ensure that you configure the appropriate IP security policy on the backup and target hosts.

IP Services: Be aware of IP Fragment attack type of the Intrusion Detection Services (IDS) enhancements to monitor both IPv4 and IPv6 traffic (Required-IF, as of V2R1)

Required if you are using IDS on a stack and the IP Fragment attack type is enabled.

Beginning in z/OS V2R1, IP Fragment attack type of the Intrusion Detection Services (IDS) is enhanced to monitor both IPv4 and IPv6 traffic for suspicious fragments. It is also enhanced further to check for overlays that change the data in the packet. Be aware that in z/OS V2R1, if you have the IP fragment IDS attack enabled, IPv6 traffic will now be monitored. In earlier releases, only IPv4 traffic was monitored.

Migration action: If you are using IDS on a stack and the IP Fragment attack type is enabled, be aware of the following information:

- Both IPv4 and IPv6 traffic are monitored for suspicious fragments.
- The IP Fragment attack type checking is enhanced to check for overlays that change the data in the packet, including changes to the length of the packet.

IP Services: Prepare for the addition of IPv6 support for policy-based routing (Required-IF, as of V2R1)

Required if you are using policy-based routing on a stack that is being run as a dual-mode stack (IPv4 and IPv6).

As of z/OS V2R1 policy-based routing is enhanced to route IPv6 traffic. In earlier releases a policy-based routing rule that did not specify the source and destination IP addresses only applied to IPv4 packets. Starting in z/OS V2R1, that same policy-based routing rule applies to both IPv4 and IPv6 packets.

Migration action: If you have a policy-based routing rule that specifies neither source IP addresses nor destination IP addresses, the rule will apply to both IPv4 and IPv6 packets. If you want the rule to continue to apply to only IPv4 packets, modify the rule to specify either a source or destination IP address of 0.0.0.0/0.



IP Services: Replace any GATEWAY statements in the TCP/IP profile with equivalent BEGINROUTES statements (Recommended, as of V2R1)

Not required, but recommended to take advantage of the latest support offered by BEGINROUTES/ENDROUTES.

As of z/OS V2R1 support for the GATEWAY statement in the TCP/IP profile is eliminated. The BEGINROUTES/ENDROUTES statement block was introduced in z/OS V1R1 Communications Server and replaces

the GATEWAY profile statement for configuration of static routes. The GATEWAY statement has been stabilized since then, and has not been updated to support enhancements such as IPv6 and replaceable static routes. Additionally, the GATEWAY syntax is error prone.

Migration action: Replace any GATEWAY statements in your TCP/IP profile with the equivalent BEGINROUTES/ENDROUTES statement block.

In order to create BEGINROUTES statements from your existing specification, you can use IPCS. Run the TCPIPCS PROFILE report on a dump of the TCP/IP address space and static routes are presented as a BEGINROUTES/ENDROUTES statement block, even if you coded them by using a GATEWAY statement. However, you will have to execute the commands to know with certainty what changes to make.

Communications Server Migration Actions Pre-First IPL

IP Services: Modify GLOBALCONFIG SMCR PFID definitions (Required-IF, as of V2R2 for V2R1 path only)

Required if all the following conditions are true: 1) You used Shared Memory Communications – RDMA in z/OS V2R1 Communications Server, 2) The 10GbE RoCE Express features operated in a dedicated RoCE environment, and 3) You are running on a z13 server.

In z/OS V2R1 Communications Server, VTAM provided physical function services for IBM 10GbE RoCE Express features used for Shared Memory Communications via Remote Direct Memory Access (SMC-R) processing. This allowed multiple TCP/IP stacks operating in the same logical partition (LPAR) to share a RoCE Express feature by configuring and activating the same Peripheral Component Interconnect Express (PCIe) function ID (PFID) representation of the feature. Starting with z/OS V2R2 on a z13 server, each TCP/IP stack must have unique PFID values to represent the RoCE Express feature.

Migration action:

Before starting your TCP/IP stacks that activate 10GbE RoCE Express features, perform the following steps:

1. Use the hardware configuration definition (HCD) to define a unique FID value for each TCP/IP stack that will be activating the 10GbE RoCE Express feature. You must also assign a virtual function number (VFN) for each potential user of the 10GbE RoCE Express feature.
2. Modify the GLOBALCONFIG SMCR statement in the TCP/IP profile to specify the PFID values that are assigned for this stack. Update the GLOBALCONFIG statements for all TCP/IP stacks that activate a given 10GbE RoCE Express before starting any of the TCP/IP stacks. The PFID values should correspond to the FID values you defined in the HCD.

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
Communications Server z/OS UNIX applications	SEZAINST(SERVICES)	/etc/services	In z/OS V2R2, the NCPROUT entry is removed because NCPROUTE is no longer supported.
DCAS	No sample provided	/etc/dcas.conf	In z/OS V2R2, a new TLSV1ONLY keyword is provided to configure SSLv3 protocol for connections secured using the DCAS SSL support.
FTP Server and	SEZAINST(FTCDATA) for the	/etc/ftp.data	In z/OS V2R2, a new SSLV3

Migrating to z/OS 2.2: Part 2 of 2

Client	client and (FTPDATA) for the server		keyword is provided to configure SSLv3 protocol for connections secured using the FTP TLS support.
Internet Key Exchange Daemon(IKED)	/usr/lpp/tcpip/samples/iked.conf	/etc/security/ iked.conf	In z/OS V2R2, a new log level is added for the IKE for the IKE daemon.
Policy agent	/usr/lpp/tcpip/samples/pagent.conf	/etc/pagent.conf	In z/OS V2R2, a new ServerSSLv3 keyword is provided to configure SSLv3 protocol for the policy client connecting to the server.
SNMP agent	/usr/lpp/tcpip/samples/snmpd.conf	/etc/snmpd.conf	In z/OS V2R2, a new privacy protocol value AESCFB128 can be specified on a USM_USER statement to request AES 128- bit encryption.
z/OS UNIX snmp command	/usr/lpp/tcpip/samples/snmpv2.conf	/etc/osnmp.conf	In z/OS V2R2, a new privacy protocol value AESCFB128 can be specified on a statement for an SNMPv3 user to request AES 128-bit encryption.
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.
File Transfer Protocol Daemon (FTPD)	SEZAINST(FTPDATA)	/etc/ftp.data	In z/OS V2R1, a new configuration statement is provided to specify that a type 119 SMF record of subtype 71 is collected for the FTP daemon configuration information when the FTP daemon starts.

IP Services: Verify the new default for the QUEUEDRTT parameter (Required-IF, as of V2R2)

Required if you do not specify the QUEUEDRTT parameter on the TCPCONFIG profile statement.

In z/OS V2R2, enhancements were made to the Communications Server outbound serialization function, which is controlled by the QUEUEDRTT parameter on the TCPCONFIG profile statement. The default value for this parameter is changed. Previously, the default setting for QUEUEDRTT was 20, meaning that only TCP/IP connections with an round-trip time (RTT) value of 20 milliseconds or more are eligible to use outbound serialization. In z/OS V2R2, the default value for QUEUEDRTT is changed to 0, meaning that all TCP/IP connections are eligible to use outbound serialization.

Note: The new default value allows more connections to be eligible for outbound serialization. This performance optimization might result in higher CPU costs at the sending host, but these costs are offset by the CPU savings on the receiving host.

Migration action: If you do not currently specify the QUEUEDRTT parameter on the TCPCONFIG profile statement, but you want to continue using a value of 20 milliseconds as the threshold to enable a TCP/IP connection to use outbound serialization, you must specify QUEUEDRTT 20 on the TCPCONFIG profile statement.

IP Services: Use the new maximum segment size adjustments if required (Required-IF, as of V2R2)

Required if you do not want the automatic reduction of MSS.

New function is introduced in z/OS V2R2 to automatically adjust the Maximum Segment Size (MSS) that is used on a TCP connection to avoid fragmentation. This function is turned on by default.

Note: A performance degradation can result if the MSS is not configured to avoid fragmentation.

Migration action: To disable the auto-adjusting function and continue using your existing setting for the MSS from the prior release, specify the subparameter ADJUSTVIPAMSS NONE on the parameter GLOBALCONFIG.

IP Services: Check code that automates on IKE daemon syslogd messages (Required-IF, as of V2R2)

Required if you use the z/OS IKE daemon and you have code that automates on IKED messages written through syslogd.

Migrating to z/OS 2.2: Part 2 of 2

Scalability enhancements in z/OS V2R2 Communications Server introduce a new internal thread pool to the Internet Key Exchange (IKE) daemon. As part of this change, all of the IKED messages that are written through syslogd will contain the thread identifier in the syslogd header, which precedes the message identifier. In addition, messages from different IKED threads might be interleaved. These changes might affect the automation code that parses these IKED messages:

- If the automation code parses individual IKED messages in syslogd destinations based on any sort of position-based logic (for example, counting blanks delimiters, relying on a specific column), you need to update to account for the new thread id field.
- If the automation code parses individual IKED messages based on message content (for example, searching the string for specific message identifiers), no change is needed.
- If the automation code depends on the order of the IKED messages, you might need to update to take the thread id field into account so that the code ignores messages from other threads that are interleaved with the messages of interest.

IKED messages are those in the ranges EZD0902I - EZD1160I, EZD1751I - EZD1800I, and EZD1901I - EZD1925I as well as EZD2017I, EZD2019I, EZD2025I and EZD2027I. The following example shows a small excerpt of messages from z/OS V2R1 and the equivalent messages from z/OS V2R2 with the imbedded thread identifiers:

V2R1:

```
Jul 28 11:39:26 mvs046 IKE: EZD1061I IKE connecting to PAGENT
Jul 28 11:39:26 mvs046 IKE: EZD1062A IKE retrying connection to PAGENT
Jul 28 11:39:34 mvs046 IKE: EZD0923I IKE has received the STOP command
Jul 28 11:39:34 mvs046 IKE: Message instance 0: EZD0967I IKE release
CS V2R1 Service Level CS130924 Created on Sep 24 2013
Jul 28 11:39:34 mvs046 IKE: Message instance 14: EZD1116I IKE detected
an NAPT in front of the remote security endpoint while initiating a new
phase 1 tunnel
```

V2R2:

```
Jul 28 15:10:47 mvs046 IKE: (00000001) EZD1061I IKE connecting to PAGENT
Jul 28 15:10:47 mvs046 IKE: (00000001) EZD1062A IKE retrying connection
to PAGENT
Jul 28 15:11:06 mvs046 IKE: (00000003) EZD0923I IKE has received the STOP command
Jul 28 15:11:06 mvs046 IKE: Message instance 0: (00000001) EZD0967I IKE release
CS V2R2 Service Level CS140728 Created on Jul 28 2014
Jul 28 15:11:06 mvs046 IKE: Message instance 14: (00000007) EZD1116I IKE detected
an NAPT in front of the remote security endpoint while initiating a new
phase 1 tunnel
```

Migration action: Change the code that automates on IKED messages written through syslogd to account for the thread identifier that is added to the header area preceding the IKED message number.

IP Services: Decide whether to accept the new FIXED CSM default (Required-IF, as of V2R2)

Required if you use the default CSM FIXED MAX value of 100M and you do not want to use the new default of 200M.

In z/OS V2R2, the default amount for communications storage manager (CSM) fixed storage for buffers is increased from 100 MB to 200 MB. Your installation can specify a value for the CSM fixed storage amount on the FIXED statement in the IVTPRM00 parmlib member.

Migration: If you did not previously code a value for FIXED in IVTPRM00 and you do not want the new default, specify FIXED MAX(100M) in your IVTPRM00 parmlib member to retain the value as formerly defaulted.

Tip: You can use the **D NET,CSM** command to display the "FIXED MAXIMUM" storage specification in message IVT5538I.

SNA Services: Update TIBUF pool size and T1BUF pool size (Required-IF, as of V2R2)

Required if you have modified the buffer pool values for the TIBUF pool or the T1BUF pool, and you might want to revisit your settings.

z/OS V2R2 reduces the number of buffers per page of storage for the following buffer pools:

TIBUF pool bufsize change : The TIBUF pool contains control information to support HPDT services for HPR or IP. It is used to contain the HPR headers and the media, IP, and UDP headers for an Enterprise Extender connection. It is also used to contain data for APPC conversations.

T1BUF pool bufsize change : The T1BUF pool contains control information to support HPDT services for HPR or IP. It is similar to the TIBUF pool, but larger. It is used as a packing buffer by HiperSockets accelerator and QDIO. It is also used to contain the HPR headers and the media, IP, and UDP headers for an Enterprise Extender connection.

Migration action: If you have tuned your systems to require a specific number of pages of TIBUF or T1BUF

Migrating to z/OS 2.2: Part 2 of 2

buffer storage, you might need to recalculate the number of pages required. For information, see Tuning Enterprise Extender specific buffer pools in *z/OS V2R2.0 Communications Server: SNA Network Implementation Guide*.

DFSMSdfp Migration Actions for z/OS V2R2



• Migration Actions Before Installing:

- **Determine whether you need DISABLE(REFUCB) in parmlib member DEVSUPxx (Required-IF, as of V2R2)**
 - REFUCB allows the system to automatically update the UCB when a DSS COPY, RESTORE, or ICKDSF REFORMAT, INIT, and FLASHCPY has changed the volser or VTOC location. It applies only to volumes shared within a sysplex.
 - If the device is ONLINE, REFUCB issues a VARY ONLINE, UNCONDITIONAL command, which updates both the volser and VTOC location in the UCB.
 - If the device is OFFLINE, no action is taken.
 - As of V2R2, REFUCB is enabled by default.
 - A new health check, DMO_REFUCB, verifies that you are using the recommended value of ENABLE (REFUCB)
 - Use new default as recommended or update as desired:
 - DEVSUPxx with DISABLE (REFUCB) , and
 - check DMO_REFUCB with ENABLE (NO)



19

© 2015 IBM Corporation

ICKDSF Migration Actions for z/OS V2R2



• Migration Before First IPL:

- **Accommodate new default for INIT and REFORMAT commands (Required-IF, as of PI25913)**
 - VERIFYOFFLINE on INIT and REFORMAT verifies device is offline before processing.
 - As of PI25913, if UNIT is specified on INIT or REFORMAT, -and- the storage control microcode and support is there, -and- the device is not a minidisk, VERIFYOFFLINE is the default (and doesn't need to be specified).
 - If the device is online during VERIFYOFFLINE processing, the command will terminate.
 - To prevent the verification from taking place, use NOVERIFYOFFLINE.



20

© 2015 IBM Corporation

DFSMSHsm Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Update applications that depend on LIST command output (Required-IF, as of V2R2)**
 - As of V2R2, LIST DUMPCLASS command output has slightly changed:
 - New MINSTACK value appears, STACK position changes.
 - Applies to all destinations: OUTDATASET, SYSOUT, or TERMINAL
 - LIST command output is not a programming interface, but update any applications as needed for these changes.

```

CLASS=COMDUMP UNIT=3490 REUSE=Y RESTORE=Y RESET=N DISABLE=N REOD=* A
AVAILABLE=Y
DAY=02 FREQ=000 RETPD=000006 TAPEEXPDT=***** VTOCCOPIES=001 RECRST=YES
ZCOMP=NO
DISP=WEEKLY COMBAT DUMPS
HWCOMP=NO ENCRYPT=NONE ENCRYPT_TYPE=***** ICOUNT=*****
ENC_KEY=*****
MINSTACK=001 STACK=001
ARC0140I LIST COMPLETED, 6 LINE(S) OF DATA OUTPUT
COMMAND REQUEST 00000023 SENT TO DFSMSHSM
***
    
```


DFSMSHsm Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Accommodate default value change for dump VTOC copies (Required-IF, as of V2R1)**
 - Prior to V2R1, the default value was 2 for VTOCCOPIES of the DEFINE DUMPCLASS command.
 - As of V2R1, the new default is 0 when hsm dumps copy pool volumes.
 - The default value of 2 remains when hsm dumps non-copy pool volumes.
 - VTOC copies are not needed in a FRBACKUP/FRRECOV environment.
 - If you want old behavior for FRBACKUP/FRRECOV environment, specify VTOCCOPIES(2) on DEFINE DUMPCLASS for each dump class used to dump copy pool volumes.

DFSMSdss Migration Actions for z/OS V2R2



- **Migration Actions Post-First IPL:**
 - **Accommodate new default behavior for full-volume and track restore operations (Required-IF, as of V2R1)**
 - The data-set-changed indicator (DS1DSCHA) in the VTOC indicates whether or not the data set has changed since its last backup. Before V2R1:
 - During full-volume restore, dss unconditionally turned off DS1DSCHA for each ds restored to the target volume.
 - During a tracks restore, if any of the VTOC track was restored, DS1DSCHA might be turned off for all ds on the volume.
 - Based on numerous requirements, as of V2R1, the default has changed. Now, dss turns off DS1DSCHA only if the RESET keyword was specified on the DUMP.
 - Also, a RESET keyword is added to the RESTORE FULL and RESTORE TRACKS commands, to specify whether DS1DSCHA is to be reset (turned off).

23
© 2015 IBM Corporation

DFSMS Migration Actions for z/OS V2 R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

DFSMS Migration Actions You Can Do Now



DFSMSdss: Determine whether you need DISABLE(REFUCB) in parmlib member DEVSUPxx (Required-IF, as of V2R2)

Required if you require the automatic refresh UCB function to be disabled. Note that the automatic refresh UCB function is only applicable to volumes that are shared with other systems in the same sysplex.

With z/OS V2R2, the automatic refresh UCB (REFUCB) function of the Device Manager is enabled by default. Previously, it was disabled by default. If you want the function to be disabled, you must explicitly disable it by using the statement DISABLE(REFUCB) in parmlib member DEVSUPxx.

ICKDSF FLASHCOPY, INIT, and REFORMAT commands, and DFSMSdss full volume COPY and RESTORE functions, might update the volume serial and location of the volume table of contents (VTOC). This behavior can present a problem when the device is online to other systems. To address this problem, the automatic refresh UCB function (REFUCB) was introduced in z/OS V1R13, and is controlled by DISABLE | ENABLE(REFUCB) in parmlib member DEVSUPxx. DISABLE(REFUCB) was the default.

Note: For each system that has enabled the REFUCB function, an unconditional VARY ONLINE to the device is performed when the system is notified that the volume serial, the VTOC location, or both, has changed since the device was last varied online. This action updates fields in the UCB, including the volume serial (UCBVOLI) and the start location of the VTOC (UCBVTOC). If the VARY ONLINE, UNCOND fails for the device, or if the VARY ONLINE, UNCOND is not performed because the REFUCB function is not enabled on a system in the sysplex, the following write-to-operator (WTO) message is written to the system console: DMO0063E dddd.volser, UCB NOT UPDATED REFUCB=[Y/N],USERS=xxxx.

Migration action: If you determine that you need the automatic refresh UCB function to be disabled, review your current parmlib member DEVSUPxx. If DISABLE(REFUCB) is not present, add DISABLE(REFUCB). If ENABLE(REFUCB) is present, delete it.

DFSMS SDM: Prepare for the removal of TSO copy services commands (Required-IF, as of V2R2)

Required if you use the TSO commands.

z/OS V2R2 is planned to be the last release to include a number of TSO/E-based System Data Mover (SDM) related commands. Except for the query commands (CQUERY, FCQUERY, RQUERY, XQUERY, XSTATUS), and the XSET command, which will remain, IBM recommends that users migrate to the REXX versions of these commands.

Specifically, the following commands will no longer be supported in TSO:

- FCESTABL
- FCWITHDR
- CDELPAIR
- CDELPATH
- CESTPAIR
- CESTPATH
- CGROUP
- CRECOVER
- CSUSPEND
- RSESSION
- RVOLUME
- XADDPAIR
- XADVANCE
- XCOUPLE
- XDELPAIR
- XEND
- XRECOVER
- XSTART
- XSUSPEND

Migration action: Convert existing non-query TSO commands to the REXX version using the programs provided in SYS1.DGTCLIB (ANTFREXX for FlashCopy, ANTPREXX for PPRC, ANTXREXX for XRC). Some of the command keywords are slightly different than the TSO version, and might need to be modified. For example, for full volume FlashCopy establish, you might enter the TSO command, as follows:

```
FCESTABL SDEVN(0F60) TDEVN(0F61)
```

To use the REXX interface, you can enter:

```
ANTFREXX FCESTABLISH SDEVN(0F60) TDEVN(0F61) SRCEXTNA() andTGTEXTNA()
```

DFSMS Migration Actions Pre-First IPL

DFSMSdfp: Adjust parameters to control whether offline device discovery is Performed (Required-IF, as of z/OS V2R1 and z/OS V1R13, both with APAR OA43706)

Required if you need offline device discovery to be performed during IPL.

Beginning in APAR OA43706 (PTFs: UA74649 for z/OS V2R1 and UA74648 for z/OS V1R13), the System Data Mover (ANTAS000 address space) is changed to conditionally invoke offline device discovery during IPL. Prior to this APAR, offline device discovery was always performed.

APAR OA43706 introduced a new keyword called OfflineDiscovery to the ANTXIN00 parmlib member to allow you to control whether offline device discovery is performed during IPL when ANTAS000 first starts up, and thereafter when ANTAS000 restarts following a CANCEL. A value of YES causes offline discovery to take place. A value of NO bypasses offline device discovery. If OfflineDiscovery is not specified, the default is OfflineDiscovery(NO).

Tip: The Offline Device Discovery was introduced by APAR OA36570 (PTFs for z/OS V1R12 and V1R13) and incorporated into z/OS V2R1 to reduce the XSTART time following an IPL.

SYSLOG messages example:

```
ANTB8000I XRC INITIALIZATION STARTED
ANTB8002I OFFLINE DEVICE DISCOVERY COMPLETE; RC=0 REAS=606
ANTB8001I XRC INITIALIZATION COMPLETED
```

Migration action: You should only need to perform offline device discovery during IPL (that is, specify OfflineDiscovery(YES) in ANTXIN00) if you are using one of the following:

- GDPS MzGM (aka XRC) in HYPER-PPRC mode (Incremental Resynch)
- GDPS MzGM under z/OS with XRC Offline Primary devices

Other XRC configurations or non-XRC environments will not directly benefit from having offline discovery performed at IPL. Therefore, if it is not required, do not specify the keyword and avoid discovery, which may reduce overall IPL elapsed time.

Tip: The value of OfflineDiscovery takes effect when the ANTAS000 address space is started during IPL, or with the automatic restart of ANTAS000 after it has been cancelled. To activate a new value for OfflineDiscovery without an IPL, do the following:

1. Update the value in PARMLIB member ANTXIN00
2. Issue the system command CANCEL ANTAS000

The OfflineDiscovery keyword is ignored and will not be processed if it is specified in the ALL or session member of hlq.XCOPY.PARMLIB.

DFSMSdss: Accommodate ADREID0 EIREC14 expansion (Required-IF, as of z/OS V2R1 and z/OS V1R13, both with APAR OA42238)

Required if you use the DFSMSdss application programming interface.

With APAR OA42238, the length of the EIREC14 structure in the ADREID0 macro is increased by 16 bytes. This expansion, which is related to DFSMSdss support for dump data sets compressed with zEDC services, may require changes for applications that invoke DFSMSdss using its application programming interface.

Migration: Check for use of the Function Ending user interaction module exit option (EIOPTION14) in applications that exploit the DFSMSdss application programming interface. If EIOPTION14 is used, determine if the UIM copies the structure pointed to by EIRECPTR. In this exit the structure pointed to by EIRECPTR is named EIREC14. If the application copies the area of storage pointed to by EIRECPTR, then action may be required.

If the application copies the area of storage using EIRECLEN for the amount of data to copy, then the application must ensure the area that is the destination of the copy is large enough to accommodate the larger EIREC14 structure.

The length of the EIREC14 structure is not provided in the ADREID0 data area. Prior to the application of the PTFs for OA42238, determine the length of the EIREC14 structure with an equate of this form:

```
EQU EI14CPUT+L'EI14CPUT-EIREC14
```

After the application of the PTFs for OA42238, determine the length of the EIREC14 structure with an equate of this form:

```
EQU EI14ZCSV+L'EI14ZCSV+15-EIREC14
```

DFSMSshm: Update applications that depend on LIST command output (Required-IF, as of V2R2)

Required if your application depends on the output of the LIST DUMPCLASS command.

Beginning in z/OS V2R2, the output of the LIST DUMPCLASS(name) command includes a new MINSTACK value, and the position of the existing STACK value in the output has changed. This new and changed output is displayed when OUTDATASET, SYSOUT (the default), or TERMINAL is specified as the destination for the output.

Migration action: Remove any dependency on the STACK field location in the LIST DUMCPCLASS with a output target of TERM, OUTDATASET, or SYSOUT. Also, update applications as needed for the new MINSTACK field in the LIST DUMCPCLASS with a output target of TERM, OUTDATASET, or SYSOUT.

DFSMSdfp: Do not use IEBCOPYO (Required-IF, as of V2R1)

Required if you used IEBCOPYO directly or through the alias IEBDSCP Y

In z/OS V1R13, the IEBCOPY utility was enhanced, and one of the enhancements was that the utility was no longer APF-authorized. For users who needed to fallback from the new IEBCOPY to the previous APF-authorized form, IEBCOPY was retained in z/OS V1R13 under the name IEBCOPYO. The IEBDSCP Y alias name for IEBCOPY was also revised to apply to IEBCOPYO for z/OS V1R13.

In z/OS V2R1, the IEBCOPYO form of the utility is eliminated and the IEBDSCP Y alias now applies to the standard non-APF-authorized form of IEBCOPY.

Migration action:

- Any programs or jobs that called IEBCOPYO in z/OS V1R13 must be changed to call IEBCOPY in V2R1. It is not expected that you used IEBCOPYO on z/OS V1R13, unless you had a problem with the z/OS V1R13 level of IEBCOPY and had to fall back to the z/OS V1R12 level of IEBCOPY.
- Any programs or jobs that call the IEBDSCP Y alias name in z/OS V2R1 will invoke the non-APF-authorized IEBCOPY. Make any appropriate changes if this affects your programs or jobs.

DFSMSshm: Accommodate default value change for dump VTOC copies (Required-IF, as of V2R1)

Required if you want the old default value of VTOCCOPIES(2) in a FRBACKUP/ FRRECOV environment.

Before z/OS V2R1, the system used a default value of 2 for the VTOCCOPIES parameter of the DFSMSshm DEFINE DUMPCLASS command. Starting with z/OS V2R1, if the VTOCCOPIES parameter is not specified, the

system uses a default value of 0 when DFSMSShsm dumps copy pool volumes. The default value of 2 will continue to be used when DFSMSShsm dumps non-copy pool volumes. If the VTOCCOPIES parameter is never specified, the LIST DUMPCLASS output displays a VTOCCOPIES value of *** to indicate that no value is specified. VTOC copies are not needed in a FRBACKUP/FRRECOV environment.

Migration action: If you want to keep the default value of VTOCCOPIES(2) used before z/OS V2R1 in a FRBACKUP/FRRECOV environment, specify VTOCCOPIES(2) on the DEFINE DUMPCLASS command for each dump class used to dump copy pool volumes.

DFSMSShsm: Update applications that depend on LIST command output (Required-IF, as of V2R1)

Required if your application depends on the output of the LIST DUMPCLASS command

The output of the LIST DUMPCLASS command now includes the value of the RECOVERRESET parameter. RECOVERRESET is a new parameter of the DEFINE DUMPCLASS command that specifies whether or not the data-set-changed indicator in the VTOC entry is reset for all data sets that are restored during full volume recover processing.

Migration action: Update applications that depend on the output of the LIST DUMPCLASS command to accommodate the new value for RECOVERRESET.

DFSMS Migration Actions Post-First IPL

DFSMSdss: Accommodate new default behavior for full-volume and track restore operations (Required-IF, as of V2R1)

Required if you want DFSMSdss to unconditionally reset the data-set-changed indicator during full-volume and tracks restore operations.

The data-set-changed indicator (DS1DSCHA) in the VTOC indicates whether or not the data set has changed since its last backup. Before z/OS V2R1, during a full-volume restore operation, DFSMSdss unconditionally reset (turned off) the data-set-changed indicator for each data set restored to the target volume. During a tracks restore operation, if any VTOC track was restored, DFSMSdss might reset the data-set-changed indicator for all data sets on the volume. This applies to all VSAM and non-VSAM data sets and all SMS and non-SMS data sets.

With z/OS V2R1, the default behavior for full-volume and tracks restore operations has changed. By default, DFSMSdss now resets the data-set-changed indicator only if the RESET keyword was specified on the DUMP command. Along with this change, a RESET keyword has been added to the RESTORE FULL and RESTORE TRACKS commands, which allows you to specify whether the data-set-changed indicator is to be reset. In addition, you can use the options installation exit routine, ADRUIXIT, to control the resetting of the data-set-changed indicator.

You can use RESET on the RESTORE command for any FULL or TRACKS dump taken with V2R1, or any previous releases. RESET(YES) and RESET(NO) will work as expected. However, RESET(DUMP) will act as RESET(NO) because there is no way to determine if the RESET keyword was specified at DUMP time on a down-level dump. The new RESET(DUMP) functionality is only available in V2R1. Note that the default behavior for DFSMSdss DUMP is unchanged when DFSMSShsm invokes DFSMSdss to perform a full volume restore. By default, DFSMSShsm will pass RESET(YES) to DFSMSdss. The other parameters might be specified through the DEFINE DUMPCLASS command.

On a down-level dump, RESET(DUMP) will act as RESET(NO) because there is no way to determine if the RESET keyword has been specified at DUMP time on a down-level dump. The new RESET(DUMP) function is only available in V2R1.

Migration action: To obtain the behavior in previous releases to allow DFSMSdss to unconditionally reset the data-set-changed indicator during full-volume and tracks restore operations, use RESET(YES) with the RESTORE FULL and RESTORE TRACKS commands.

ICKDSF Migration Actions for z/OS V2 R2

[ICKSSF Migration Actions After First IPL](#)

ICKDSF: Accommodate new default for INIT and REFORMAT commands (Required-IF, as of APAR PI25913 (on ICKDSF R17) for both z/OS V2R1 and z/OS V1R13)

Required if you require that the new default not be used.

The VERIFYOFFLINE parameter of the INIT and REFORMAT commands verifies that the device is offline to all other systems before the command is attempted. Beginning with ICKDSF APAR PI25913, if the UNIT parameter is specified for the INIT or REFORMAT command, and the required storage control microcode support and software support exists, and the device is not a minidisk, the VERIFYOFFLINE parameter is the default, meaning that verification that the device is offline occurs even if VERIFYOFFLINE is not specified. Prior to ICKDSF APAR PI25913, verification occurred only when VERIFYOFFLINE was specified.

If the VERIFYOFFLINE parameter is processed as the default, and the device is online to another system, the command terminates, with message ICK31306I in the job output. To request that verification not take place, you can use a new NOVERIFYOFFLINE parameter.

Note: The VERIFYOFFLINE parameter of the INIT and REFORMAT commands was introduced by ICKDSF APAR PM76231

Target system hardware requirements: DS8870 R7.1 GA, Bundle 87.10.87.0, LIC 7.7.10.287

DS8700 R6.3 SP6, Bundle 76.31.79.0, LIC 6.6.31.670

DS8800 R6.3 SP6, Bundle 86.31.95.0, LIC 7.6.31.1150

Target system software requirements: The software support is provided in APAR OA40719 for z/OS V2R1 and V1R13, and incorporated into z/OS V2R2.

Note: If the VERIFYOFFLINE parameter is specified either explicitly or by default, and the device is online to any other systems, message ICK31306I VERIFICATION FAILED: DEVICE FOUND TO BE GROUPED is issued and the command is terminated. If the NOVERIFYOFFLINE parameter is specified and the device is online, command processing continues.

Migration action: If you do not want INIT and REFORMAT commands with the UNIT parameter to verify that the device is offline to all other systems before the command is attempted, add the new NOVERIFYOFFLINE parameter to those commands.

Migration Actions for z/OS V2R2



If you use FIPS 140, these migration actions may be applicable to you:



- **IP Services: Ensure ICSF is active before starting the NSS daemon in FIPS 140 mode (V2R1)**
- **IP Services: Ensure ICSF is active before starting the IKE daemon in FIPS 140 mode (V2R1)**
- **IP Services: Ensure ICSF is active before starting the Policy Agent when AT-TLS groups are configured in FIPS 140 mode (V2R1)**
- **System SSL: Ensure ICSF is available when running System SSL in FIPS 140-2 mode (V2R1)**

Something in common: As of z/OS V2R1, FIPS 140 support now requires ICSF services.



24



Migration actions for FIPS 140 users in z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

IP Services: Ensure ICSF is active before starting the NSS daemon in FIPS 140 mode (Required-IF, as of V2R1)

Required if the NSS daemon is configured in FIPS 140 mode.

As of z/OS V2R1 FIPS 140 support now requires ICSF services. If the NSS daemon is configured in FIPS 140 mode, the daemon will fail to activate if ICSF is not active. Ensure ICSF is started before starting the NSS daemon if it is configured in FIPS 140 mode.

Migration action: If the NSS daemon is configured in FIPS 140 mode, ensure ICSF is active prior to starting the NSS daemon.

IP Services: Ensure ICSF is active before starting the IKE daemon in FIPS 140 mode (Required-IF, as of V2R1)

Required if the IKE daemon is configured in FIPS 140 mode

As of z/OS V2R1 FIPS 140 support now requires ICSF services. If the NSS daemon is configured in FIPS 140 mode, the daemon will fail to activate if ICSF is not active. Ensure ICSF is started before starting the NSS daemon if it is configured in FIPS 140 mode.

Migration action: If the IKE daemon is configured in FIPS 140 mode, ensure ICSF is active prior to starting the IKE daemon.

IP Services: Ensure ICSF is active before starting the Policy Agent when AT-TLS groups are configured in FIPS 140 mode (Required-IF, as of V2R1)

Required if AT-TLS groups are configured in FIPS 140 mode.

As of z/OS V2R1, FIPS140 support now requires ICSF services. Ensure ICSF is started before starting AT-TLS groups with FIPS140 support enabled. ICSF services will be used for random number generation and for Diffie Hellman support for generating key parameters, key pairs and key exchanges.

Migration action:

Migrating to z/OS 2.2: Part 2 of 2

1. Ensure ICSF is active before starting AT-TLS groups configured to support FIPS140-2
2. If the CSFSERV class is defined, give READ access to the userid associated with the TCPIP stack and any application userid using the TTLSGroup to the CSFRNG resource within the RACF CSFSERV class.
3. If the CSFSERV class is defined and Diffie Hellman is being used, give READ access to the application userid to the CSF1TRC, CSF1DVK, CSF1GKP, CSF1GSK, CSF1GAV, and CSF1TRD resources within the RACF CSFSERV class.

System SSL: Ensure ICSF is available when running System SSL in FIPS 140-2 mode (Required-IF, as of V2R1)

Required if your installation runs System SSL in FIPS mode.

In z/OS V2R1, System SSL, when running in FIPS 140-2 mode, uses ICSF's random number generation and Diffie-Hellman support. Before running System SSL in FIPS 140-2 mode you must ensure that ICSF is running and that all user IDs that start SSL applications in FIPS 140-2 mode, invoke the gskkyman utility to manage FIPS 140-2 key database files, or invoke the GSKSRVR started task in FIPS mode have access to certain CSFSERV classes.

When it is running in non-FIPS mode, System SSL uses its own implementation of Diffie-Hellman and does not require ICSF. In non-FIPS 140-2 mode, however, System SSL attempts to use ICSF's random number generation as it would when running in FIPS 140-2 mode. If ICSF or the required resource is unavailable, System SSL uses its own random number generation capabilities as in earlier releases.

Migration action: To run System SSL in FIPS 140-2 mode, you must now make sure that ICSF is running and that all user IDs that start SSL applications in FIPS 140-2 mode, invoke the GSKSRVR started task in FIPS 140-2 mode, or invoke the gskkyman utility to manage FIPS 140-2 key database files can access the necessary ICSF callable services.

DFS Migration Actions for z/OS V2R2



Migration Actions you can do NOW:



- **SMB: Add `_BPX_UNUSEDTASKS=KEEP` to the DFSCNTL environment file (Required-IF, as of V2R2, before with OA39909)**
 - Before V2R2 and OA39909, SMB users did not have to specify that medium weight threads were to be kept.
 - With V2R2 and OA39909, SMB users must keep the original threading behavior by setting the `_BPX_UNUSEDTASKS=KEEP` environment variable.
 - Add this value to the SMB DFSCNTL environment file
 - Usually found in `/opt/dfslocal/home/dfscntl/envar`
 - You can stop and restart SMB to pick up this change
 - Ensure you use `C DFS` (not `STOP DFS`) and then `S DFS`
 - Failure to do this means SMB server will not start

```
DFSKERN S069 422 SA03
BPXP023I IOEP01100I
IOEP01123A RESTART SIGKILL
      or
SYSTEM ABEND S069 - REASON CODE 04
```

25

© 2015 IBM Corporation

DFS Migration Actions for z/OS V2R2



Migration actions new for zFS in V2R2 ?

- **None! Thank you zFS team!**





- **zFS does have many excellent *exploitation* items in V2R2, which do not cause any migration actions:**
 - Providing support for 64-bit addressing (AMODE64)
 - With 64-bit support, some caches now support a larger value range
 - Improved monitoring of zFS performance
 - `fsinfo` query command (with `-exceptions` keyword!)
 - zFS can be run in the OMVS address space
 - The IOEFSPRM configuration options file has new options
 - ...

26


© 2015 IBM Corporation

DFS Migration Actions for z/OS V2R2

Migration Actions you can do NOW:

- 
zFS: Determine whether to accept the new default values for certain zFS variables (Required-IF, as of V2R1)
- 
 - As of z/OS V2R1, IOEPRMxx parmlib member variables `meta_cache_size`, `metaback_cache_size`, `user_cache_size`, (for better performing caches) or `convert_auditfid` have new defaults.
 - If `meta_cache_size` and `metaback_cache_size` are not specified, it is typically 10% of the real storage available during zFS initialization*, **
 - If `user_cache_size` is not specified, it is typically 10% of real storage available during zFS initialization**, **
 - If `convert_auditfid` is not specified, the default is on. so that all files and directories in zFS file systems can be uniquely identified in SMF audit records.

*See zFS Admin book for calculations when real is <64M, 64M-100M, 100M-(200G+100M), and >200G+100M.
 **See zFS Admin book for calculations when <256M, 256M-2G, and >2G.
 ***Use health check for specific calculations to help know what the default values will be pre-V2R1.





27
© 2015 IBM Corporation

DFS Migration Actions For z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

DFS Migration Actions You Can Do Now

Determine whether to accept the new default values for certain zFS variables in the zFS IOEFSPRM configuration file (Required-IF, as of V2R1)

Required if you do not want the new default values to take effect or if you want to change existing values to the new default values.

Before z/OS V2R1, certain default values were used for the IOEFSPRM files or IOEPRMxx parmlib member variables `meta_cache_size`, `metaback_cache_size`, `user_cache_size`, or `convert_auditfid`. Starting in z/OS V2R1, new default values are created for them. In z/OS V2R1, the zFS IOEFSPRM configuration file variable `convert_auditfid` default value was changed to ON so that all files and directories in zFS file systems can be uniquely identified in SMF audit records. The zFS IOEFSPRM configuration file variable `user_cache_size` default value will be changed to a value that is calculated based on the amount of real storage in the system. The zFS IOEFSPRM configuration file variables `meta_cache_size` and `metaback_cache_size` default values will be changed when both values are not specified to also be calculated based on the amount of real storage in the system. These are so that a system that has the capacity for more storage use and has sufficient space in the ZFS address space can have better performing caches. If enough real storage exists, zFS can use more storage by default than it would have used in prior releases.

Migration action: Look for these:

- IOEFSPRM files or IOEPRMxx parmlib members that do not specify both `meta_cache_size` and `metaback_cache_size` options.
- IOEFSPRM files or IOEPRMxx parmlib members that do not specify the `user_cache_size` option.
- IOEFSPRM files or IOEPRMxx parmlib members that do not specify `convert_auditfid` settings.
- Programs that use **zfsadm format** commands where unique auditfids are not desired.

Migrating to z/OS 2.2: Part 2 of 2

- JCL that contains calls to **ioeagfmt** that create aggregates for which unique auditfids are not desired.
- Programs that use zFS format API where unique auditfids are not desired.

Take these actions:

- For `meta_cache_size`, `metaback_cache_size`, or `user_cache_size`, if the old default values are desired, specify these values in your IOEFSPRM files or IOEPRMxx parmlib members.
- For `auditfid`, if you want the previous defaults, specify `-nonewauditfid` on calls to **ioeagfmt** or **zfsadm format** and `convert_auditfid=OFF` in your IOEFSPRM files or IOEPRMxx parmlib members.

DFS Migration Actions After First IPL



SMB: Add `_BPXK_UNUSEDTASKS=KEEP` to the DFSCNTL environment file (Required-IF, as of z/OS V2R1 and z/OS V1R13, both without APARs OA39909 and APAR OA43472)

Required if you use SMB.

Before z/OS V2R2 and without the z/OS UNIX PTF for APAR OA39909, which changed threading behavior, SMB users did not have to specify that medium weight threads were to be kept. As of APAR OA39909 and z/OS V2R2, SMB users must keep the original threading behavior by setting the `_BPXK_UNUSEDTASKS=KEEP` environment variable, as described in DFS SMB APAR OA45175.

Note: Failure to implement this migration action prevents the DFS/SMB server from starting or restarting because z/OS UNIX System Services issues a SIGKILL to the DFS/SMB server. The following symptoms might be encountered if this migration action is not implemented:

```
DFSKERN S069 422 SA03
```

```
BPXP023I IOEP01100I
```

```
IOEP01123A RESTART SIGKILL
```

or

```
SYSTEM ABEND S069 - REASON CODE 04
```

Migration action: Because SMB is dependent on the threading model that was in effect before the PTFs for APAR OA39909 and APAR OA43472 were installed, you need to take the following steps:

1. To allow SMB to use the original threading behavior, add the environment variable `_BPXK_UNUSEDTASKS=KEEP` to the SMB DFSCNTL environment file, which is typically found in the `/opt/dfslocal/home/dfscntl/envar` directory.
2. If SMB is started, it must be stopped and restarted using the C DFS and S DFS commands.

After you complete the steps, SMB will use the original threading behavior that was in place prior to the PTF for APAR OA39909.

zFS: Ensure that the zFS kernel is active when using the batch utility `ioeagfmt` (Required-IF, as of V2R1)

Required if you do not run `ioeagfmt` after the kernel was initialized. However, it would be unusual not to have zFS active when running `ioeagfmt`.

Before z/OS V2R1, the batch utility `ioeagfmt` did not require that the zFS kernel be active. Starting in V2R1, `ioeagfmt` requires that the zFS kernel be active.

Migration action: Look for JCL or vendor applications that might be creating or submitting JCL that contain `ioeagfmt` calls that are run when the zFS kernel is not active. Run them when the zFS kernel is active.

If you are formatting file systems to be used by another system, you should be setting the `-version` parameter on `ioeagfmt` because the local zFS kernel would not have the information for a remote zFS system. You can also use the IOEFSUTL program to format aggregates. IOEFSUTL does not require the zFS kernel to be active to format a file system if you specify the `-version` parameter. But it will require the zFS kernel to be active if the `-version` parameter is left out.

JES2 Migration Actions for z/OS V2R2



Migration Actions you can do NOW:



• **Activate z11 mode (Required in z/OS V2R2)**



- Determine your z11 checkpoint activation readiness with the **\$D ACTIVATE** command. This command indicates if activation to z11 mode will succeed.
- Run the **JES2 \$ACTIVATE** command to activate z11 mode following the considerations for this command found in *z/OS JES2 Commands*.
- After V2R2, you may **not** fallback to z2 mode.

Migration Actions Before First IPL:

- **Review changes applicable to JES2 exits**

(Required)

- Review the *z/OS JES2 Installation Exits* book (section "JES2 exit migration considerations") for changes that may impact you.
- Generally, assembly errors in JES2 exits indicate that you were affected by a JES2 data area change.

28

© 2015 IBM Corporation

JES2 Actions for z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

JES2 Migration Actions You Can Do Now



Activate z11 mode (Required, as of V2R2)

JES2 z/OS V2R1 is the last release level that will support z2 checkpoint mode. JES2 must be activated to z11 mode before migrating to z/OS V2R2. Activating z11 mode upgrades the JES2 checkpoint and enables JES2 function that is introduced in z/OS V1R11, including JOE data area extensions supported by BERTs.

Note: You are able to fall back to z2 mode, if necessary, prior to z/OS V2R2. After z/OS V2R2, you cannot fall back to z2 mode.

Migration action: Follow these steps:

- After migrating to z/OS V1R11 JES2, or later, 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" in *z/OS V2.2 Migration*.
 3. If you issue the **\$ACTIVATE,LEVEL=z11** command, activation of LARGEDS support is required.
 4. An additional *nnn* 4K records for CKPT1 is required for z11 mode.
- Run the **JES2 \$ACTIVATE** command to verify non-configuration changes that must be accommodated before going to z11, and to activate z11 mode following the considerations for this command found in *z/OS JES2 Commands*.

Note: The **SPOOLDEF LARGEDS=FAIL** (default value) in **JES2PARM** parmlib member is not supported in z11 mode. In z11 mode, on a COLD start, JES2 defaults to **LARGEDS=ALLOWED**. However, you cannot issue the **\$ACTIVATE,LEVEL=z11** command in the environment of **SPOOLDEF LARGEDS=FAIL**.

By default, JES2 V1R13 and V2R1 restart 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 or UNACT parameter.

Change JESJOBS profiles (Required-IF, as of V2R1)

Required if JESJOBS profiles could grant unintended authority to modify the job.

In z/OS V2R1, the new Job Modify SSI 85 defines additional resource names to the JESJOBS class. If any existing JESJOBS entity name matches one of the new resource names, it could result in unintended user authorization to modify job attributes. If necessary, update your JESJOBS profiles to avoid granting unintended user authority to modify jobs.

Migration action:

1. Search for JESJOBS profile names which match any of the new JESJOBS resource names:
 - SUBMIT.localnodeid.jobname.userid
 - HOLD.nodename.userid.jobname
 - RELEASE.nodename.userid.jobname
 - PURGE.nodename.userid.jobname
 - CANCEL.nodename.userid.jobname
 - START.nodename.userid.jobname
 - RESTART.nodename.userid.jobname
 - SPIN.nodename.userid.jobname
 - MODIFY.nodename.userid.jobname
 - REROUTE.nodename.userid.jobnameFor more information, see *z/OS JES2 Initialization and Tuning Guide*.
2. Ensure that your existing JESJOBS profiles grant the intended authority, given the new use of JESJOBS by the Job Modify SSI 85.
3. If any JESJOBS profile inadvertently allows user authority to Job Modify SSI 85 actions, update the profile or create a new profile, if necessary.

JES2 Migration Actions Pre-First IPL

Review changes applicable to JES2 exits (Required-IF)

Required if you have any affected exits..

In z/OS V2R2, there may be exit changes necessary depending on what exits you are using, and what JES2 data areas those exits are referencing.

Migration action: Refer to the "JES2 exit migration consideration" section found in the *JES2 Installation Exits*. This section details the changed information which you may need to accommodate. In z/OS V2R1, the changes include (but are not limited to) Exit 6 considerations, with the macro \$DTECNV containing fields which have changed.

Generally, assembly errors in JES2 exits are an indication that you are affected by a change.



Remove BROADCAST= from the OUTDEF initialization statement (Recommended, as of V2R1)

Not required, but recommended to keep your initialization deck clean of outdated and obsolete specifications.

In z/OS V2R1, JES2 ignores any specifications of the BROADCAST= parameter on OUTDEF. If a TSO user is not currently logged on, notify messages will always be issued locally (processing as with BROADCAST=YES) instead of being routed to the member that submitted the job (processing as with BROADCAST=NO). BROADCAST= never affected message processing when the TSO user is logged on. You can no longer use BROADCAST= to specify which member issues TSO notify messages when a user is not logged on.

Migration action:

1. Check the OUTDEF statement in the JES2 initialization deck for BROADCAST= specifications.
2. Remove BROADCAST= specifications from the OUTDEF statement. Any value you provide on BROADCAST= will be ignored. The external behavior will be to issue notify messages for a TSO user that is not logged on, on the member that requested the message. The message will be placed in the broadcast data set by TSO on that member.



Remove JCLERR= from the JOBDEF initialization statement (Recommended, as of V2R1)

Not required, but recommended to keep your initialization deck clean of outdated and obsolete specifications.

Before z/OS V2R1, for JCL job card errors that were detected during the input phase that the converter also detected, you could either have JES2 fail the job during input phase (JCLERR=YES) and the job never was sent to conversion, or you could have JES2 ignore the errors (JCLERR=NO) and send the job to the conversion phase where the converter could detect them. Starting with z/OS V2R1, input phase still detects errors, but jobs are always queued to the conversion phase and the input errors are added to those found by conversion and reported in the same way. The message id associated with the message indicates where the error occurred.



In z/OS V2R1, JES2 ignores any specifications of the JCLERR= parameter. If an error on the JOB card is encountered during the INPUT phase, the job is sent to the converter for INPUT phase error message processing. You can no longer use JCLERR=YES to specify for the job to be failed under these conditions.

Migration action:

1. Check the JOBDEF statement in the JES2 initialization deck for JCLERR= specifications.
2. Remove any JCLERR= specifications from the JOBDEF statement.

JES3 Migration Actions for z/OS V2R2 IBM

Migration Actions Pre-First IPL:

Be aware JES3 honors the DSI setting in SCHEDxx (Required-IF, as of V2R2)

- Before V2R2, if you specified Data Set Integrity (DSI) for IATINTK and IATINTKF in SCHEDxx, ENQs occurred only for batch allocations of JES3 data sets in the JES3 proc with DD statements.
- As of V2R2 in the same situation, JES3 will *additionally* obtain ENQs for all dynamic allocations of JES3 data sets.
- Recommended to specify DSI for IATINTK and IATINTKF in SCHEDxx.
 - To prevent accidental damage to your ds due to lack of ENQ.
 - New health check will verify setting.
 - IEFSDPPT still has coded NODSI for IATINTK and IATINTKF.
- If you have set DSI in SCHEDxx, note the new ENQs:
 - ENQs might have unintended consequences.
- Any DD allocations in your JES3 proc (if any), should have DISP=SHR. Not DISP=OLD.
- Verify use of unique spool data set names.

29
© 2015 IBM Corporation

JES3 Migration Actions For z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

JES3 Migration Action Pre-First-IPL



Be aware that JES3 honors the DSI setting in SCHEDxx (Required-IF, as of V2R2)

Required if you have previously set DSI in the SCHEDxx member, JES3 has been ignoring it for dynamic allocations. When z/OS V2R2 is started, an ENQ results from the allocations, which might lead to unexpected consequences. It is recommended that you use DSI to prevent accidental damage to your data sets due to the lack of an ENQ.

Prior to z/OS V2R2, if you specified Data Set Integrity (DSI) for IATINTK or IATINTKF in your SCHEDxx, ENQs occur only for batch allocations of JES3 data sets identified in the JES3 start procedure with DD statements. If you did not specify IATINTK or IATINTKF in SCHEDxx, then no ENQs would occur. That is because the value of NODSI for IATINTK and IATINTKF is coded in the IBM-shipped Program Properties Table (PPT, or IEFSDPPT). (NODSI remains coded in IEFSDPPT in z/OS V2R2, as it had before.)

Starting with z/OS V2R2, if you override the IBM-shipped PPT with a SCHEDxx entry of DSI for IATINTK or IATINTKF, ENQs occur for batch allocations of JES3 data sets identified in the JES3 start procedure with DD statements (as before), and in addition, JES3 will now obtain ENQs for all dynamic allocations of JES3 data sets.

It is recommended that you take advantage of the protection provided by specifying DSI for IATINTK and IATINTKF in your SCHEDxx. Notice that the SCHEDxx default when you specify a program name, is that the program name will have DSI. Meaning, if you do not specify any DSI value (DSI or NODSI) for a SCHEDxx program name, you will receive DSI. For that reason, it is recommended that you explicitly code a value of DSI in your SCHEDxx for IATINTK and IATINTKF so that there is no confusion on which value is being used.

Migration action: Follow these steps:

Migrating to z/OS 2.2: Part 2 of 2

- If you use DD allocations in your JES3 start procedure (not recommended), ensure that the DISP keyword specifies SHR, not OLD. Otherwise, the first JES3 that comes up will lock out other JES3 address spaces in the same sysplex. Note that all dynamic allocations by JES3 use DISP=SHR.
- Because an ENQ is held on the data set names, it is recommended that you use unique spool data set names to allow for offline allocation, unallocation, and formatting. Check your JES3 start procedure and your DSN specification on the DYNALLOC statement in your inish deck for duplicate names.



Remove DUMP=JES and DUMP=MVS from the OPTIONS initialization statement

(Recommended, as of V2R1)

Not required, but recommended to avoid warning messages. Any references to unsupported values are ignored, DUMP=PRDMP is used, a message is issued and initialization continues.

For enhanced problem determination and availability in z/OS V2R1, JES3 has dropped support for the DUMP=JES and DUMP=MVS parameter specifications on the OPTIONS initialization statement. In z/OS V2R1, DUMP=PRDMP is the default value and the only supported DUMP parameter specification.

Migration action:

1. Check the OPTIONS statement in the JES3 initialization deck for DUMP=JES and DUMP=MVS specifications.
2. Remove DUMP=JES and DUMP=MVS specifications from the OPTIONS statement, or change the values to DUMP=PRDMP.

Remove SDI from the OPTIONS initialization statement (Recommended, as of V2R1)

Not required, but recommended to avoid warning messages and command errors.

In z/OS V2R1, JES3 has removed the option to turn off Spool Data Integrity (SDI) functionality. If JES3 encounters the SDI keyword, command processing stops and a message is issued. If SDI is specified in the inish deck, the system issues the following message, and JES3 initialization continues: IAT3256 SDI KEYWORD IGNORED, NO LONGER SUPPORTED.

If SDI is used on a command (*F Q,SDI=), the system issues the following message, and the command is aborted: IAT8050 SDI IS AN INVALID FIELD.

Migration action:

1. Remove the SDI keyword from the OPTIONS statement in the JES3 initialization deck.
2. Remove or update any automated instances of the *F Q and *I Q commands that specify the SDI parameter.

Change JES3 release level format (Required-IF, as of V2R1)

Required if code is dependent on the JES3 release level returned by SSI 54, SSI 82 or SSI 83.

In z/OS V2R1 SSI 54, SSI 82 and SSI 83 report the JES3 release level using the JES2 format z/OS 2.1. In previous versions of z/OS, JES3 reported the release level using the format z 1.13.0. The SSI fields which are affected by the format change are:

- SSI 54: SSVIVERS field from macro IEFSSVI
- SSI 82: JPXVERSN field from macro IAZJPLEX
- SSI 83: JPSYVERN field from macro IAZJPLXI.

Migration action:

1. Search for fields SSVIVERS, JPXVERSN and JPSYVERN in invocations of SSI 54, SSI 82 and SSI 83.
2. For any code that requires the z/OS 2.1 JES3 release level, change the expected format to 'z/OS 2.1' (z/OS #.#).

z/OSMF Migration Actions for z/OS V2R2



Migration Actions After First IPL:

- **Migrate to the new release of z/OSMF (Required, as of V2R2)**
 - If you were a previous z/OSMF user, you can run the script `izumigrate.sh` to create a customized IZUPRMxx parmlib member from your current settings.
 - IZUPRMxx replaces the interactive scripts you used previously to configure z/OSMF.
 - If you want to start using z/OSMF, follow the *z/OSMF Configuration Guide*.
 - Summary is: 1) security set up 2) create a USERDIR file system (`/var/zosmf/`) and IZUPRMxx, 3) start z/OSMF server.
 - Use ServerPac full-system replace jobs and doc to create an initial instance of z/OSMF. Then add the plug-ins following the *z/OSMF Configuration Guide*.
 - ServerPac software-only and CBPDO users follow the *z/OSMF Configuration Guide* and samples.
- **Recreate all table filters in the z/OSMF user interface (Required-IF, as of V2R2)**
 - Before V2R2, table filters were saved and reused for subsequent table views.
 - As of V2R2, z/OSMF has a more modern interface consistent with other IBM products. This interface doesn't support the filters that were previously saved.
 - Recreate the filters you set in previous z/OSMF releases.



z/OSMF Migration Actions For z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

z/OSMF Migration Action Post-First-IPL

Migrate to the new release of z/OSMF (Required, as of V2R2)

As in previous releases, migrating to the new release of z/OSMF includes running the script `izumigrate.sh`. In z/OS V2R2, this script is enhanced to create a customized IZUPRMxx parmlib member, based on the configuration settings from your current (old) system. The parmlib member replaces the interactive scripts that you used to configure z/OSMF in previous releases.

Migration action: For ServerPac users, use the jobs and documentation supplied with your ServerPac order to create an initial instance of z/OSMF.

Installations that install z/OSMF from a Custom-Built Product Delivery Option (CBPDO) software delivery package, or from a ServerPac order using the software upgrade method of installation, should follow the instructions in the migration chapter in *IBM z/OS Management Facility Configuration Guide*. This work includes running the `izumigrate.sh` script on the z/OS V2R2 system.

After you complete the ServerPac installation, you can add plug-ins to z/OSMF through the Configuration Workflow, as described in *IBM z/OS Management Facility Configuration Guide*.

Remove the most-generic profile for z/OSMF authorizations (Required-IF, as of z/OSMF V2R1 with APAR PI20091)

Required if you have used a generic profile and want to use discrete profiles for more granular control.

In previous releases of z/OSMF, the generated program `izuconfig1.cfg.rexx` included RACF commands for defining the following generic profile in the ZMFAPLA class, and for creating permissions to it:

```
<SAF-prefix>.ZOSMF.**
```

where `<SAF-prefix>` is the SAF profile prefix that was defined for your configuration (by default, `IZUDEFLT`). If you used `izuconfig1.cfg.rexx` in a previous release, this generic profile was created for your configuration.

Migrating to z/OS 2.2: Part 2 of 2

As of z/OS V2R2, the z/OSMF configuration process no longer creates authorizations based on this most-generic profile. Instead, authorizations are now created based on the following discrete generic profile: <SAF-prefix>.ZOSMF

In this release, a sample job is provided in SYS1.SAMPLIB(IZUSEC) to help you create the resource names and security authorizations for z/OSMF. Your security administrator can edit and run this job to secure various resources on the z/OS system. For the profiles and permissions that are created, see the contents of the IZUSEC sample job.

Migration action:

1. Create the discrete profiles that you need for z/OSMF. In a RACF installation, you can use the commands in job IZUSEC to define profiles and create authorizations.
2. Ensure that your existing user authorizations are converted to user authorizations based on the discrete profiles. To preserve your existing authorizations, you can copy the access list from an existing profile into another existing profile, using the FROM operand on the PERMIT command.
3. Remove the most-generic profile for z/OSMF authorizations. In a RACF installation, you can use the following commands to remove the generic profile:

```
RDELETE ZMFAPLA IZUDFLT.ZOSMF.**  
SETROPTS RACLIST(ZMFAPLA) REFRESH
```

Recreate all table filters in the z/OSMF user interface [\(Required-IF, as of V2R2\)](#)

Required if users want to view only a subset of the data included in a table.

To display a subset of the items in a table in the IBM z/OS Management Facility (z/OSMF) user interface, you must define filter rules that show only the items in which you are interested. Before z/OS V2R2, the filter rules you set were saved and reapplied for subsequent views of the table.

Starting with z/OS V2R2, z/OSMF provides a modern, intuitive interface that is consistent with other IBM products. The enhanced user interface does not support the filter rules that were preserved in previous z/OSMF releases; therefore, when you display a table in z/OSMF V2R2, you might see more data or different data than you are accustomed to seeing. To show only the items in which you are interested, recreate the filter rules you set in previous releases of z/OSMF.

Migration action: Follow these steps:

1. Log into the z/OSMF V2R2 user interface.
2. Navigate to the tables for which you previously set filters.
3. In the filter row in the table, click the *Filter* link. The Build Filter window is displayed.
4. In the columns field (the first field), select the name of the column to be filtered, if needed. This field lists all the columns in the table that are visible and filterable, and it includes an *Any Column* option so that you can apply the filter rule to any column in the table.
5. In the filter condition field (the second field), select the filter condition.
6. In the value field, select or specify the value for which to filter. All the fields that follow the filter condition field are value fields. The number of value fields displayed depends on the type of data being filtered and the filter condition. If a value field is enabled, a value is required. To filter a single column for multiple values, create a separate filter rule for each value.
7. To create additional rules, click the add filter rule icon -- a plus (+) sign -- and specify the column, condition, and value for the new rule.
8. To remove a rule, click the remove rule icon -- a minus (-) sign. If the remove rule icon is disabled, the rule cannot be removed.
9. If you specified more than one filter rule, in the **Match** field, indicate whether all the rules (**All rules**) or at least one rule (**Any rule**) must be satisfied for an item to be displayed in the table. By default, **All rules** is selected. If the **Match** field is disabled (grayed out), you cannot change the relationship between the rules. In this case, all rules must be satisfied.
10. If you want the filter criteria to be case sensitive for all the rules you specified, select **Match case**. Otherwise, case is ignored.
11. Click **Filter** to filter the table.

z/OS OpenSSH Migration Actions for z/OS V2R2



Migration Actions After First IPL:

- **Accommodate the OpenSSH ported level (Required-IF, as of V2R2)**

- If you were a previous Ported Tools OpenSSH user (5.0p1), use the z/OS OpenSSH level (6.4p1).
- Several differences in the ported levels, which may cause migration actions. Some are:



- **ssh-rand-helper command:** not supported. OpenSSH requires a working `/dev/random` device, which requires that ICSF is configured to support `/dev/random/`.
- **sftp -P option:** previously it specified a path, now it specifies a port. Use `-D` to specify the `sftp_server_path`.
- **sftp ln and symlink subcommands:** previously, they created a symbolic link from `oldpath` to `newpath` on the remote host. Now, if `-s` is specified, it is a symbolic link. Otherwise it is a hard link.
- **sshd_config file RhostsAuthentication keyword:** Previously supported for protocol Version 1. Now, not supported for protocol Version 1. Update the application.

31

• Read of all changes in the *z/OS V2.2 Migration* book.

© 2015 IBM Corporation

z/OS OpenSSH Migration Actions For z/OS V2R2

These migration actions were taken from *z/OS V2R2 Migration*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to *z/OS V2R2 Migration*.

z/OS OpenSSH Migration Action Post-First-IPL

Accommodate the OpenSSH ported level (Required-IF, as of V2R2)

Required if any of the changes in the z/OS V2.2 Migration book are applicable to your environment.

Before z/OS V2R2 and IBM Ported Tools for z/OS V1R3, OpenSSH was version 5.0p1. Starting with z/OS V2R2 and IBM Ported Tools for z/OS V1R3, OpenSSH is version 6.4p1. Before z/OS V2R2, OpenSSH was available from IBM Ported Tools for z/OS. Starting with z/OS V2R2, it is now available as a base element of z/OS.

Migration action : See the tables in the *z/OS V2.2 Migration* book for the following is a list of potential migration actions for the OpenSSH base element:

- “Changes to the ssh-rand-helper command that might require a migration action”. Included below.
- “Changes to the sftp command that might require a migration action”
- “Changes to the ssh command that might require a migration action”
- “Changes to the ssh, sftp or scp client commands that might require a migration action”
- “Changes to the ssh_config file that might require a migration action”
- “Changes to the sshd command that might require a migration action”
- “Changes to the ssh, sftp or scp client commands that might require a migration action”
- “Changes to the sshd_config file that might require a migration action”
- “Changes to the ssh-keygen command that might require a migration action”
- “Changes to the ssh-keyscan command that might require a migration action”
- “Changes to the users running ssh, sftp or scp client commands that might require a migration action”
- “Changes to `/samples/ssh_smf.h` and `FOTSMF77` in `SYS1.MACLIB` that might require a migration action”



Changes to the ssh-rand-helper command that might require a migration action:

What changed Migration action needed? The `ssh-rand-helper` command. Now, the `ssh-rand-helper` is not supported.

Migration action needed? Yes. If no migration action, the following message is returned: `FOTS1949 PRNG is not seeded. Please activate the Integrated Cryptographic Service Facility (ICSF).`

Migration action: The new OpenSSH requires that a working `/dev/random` device be available to all OpenSSH client and server jobs. This requires that ICSF be configured to support `/dev/random` and that users have SAF authority to the CSFRNG service.

z/OS UNIX Migration Actions for z/OS V2R2



Migration Actions Pre-First IPL:

- **Remove files and directories in /var/man (Required-IF, as of V2R2)**
 - Before V2R2, /var/man was used by the `man` command.
 - As of V2R2, `man` no longer uses this directory. Instead, it uses /tmp on a per-user basis.
 - Remove /var/man and all subdirectories and files from your V2R2 system.
 - Because /tmp is used for caching man pages for each user, it might increase the size of your /tmp directory depending on the `man` command usage on the system.
 - Monitor the /tmp space usage to ensure adequate space for users.

32

© 2015 IBM Corporation

z/OS UNIX Migration Actions for z/OS V2R2



Migration Actions You Can Do Now:

Accommodate the new Shell and Utilities version of the zlsf utility (Required-IF, as of V2R1)

- As of V2R1, the z/OS UNIX Tools and Toy utility `zlsf` is incorporated. It differs from the Tools and Toy version in several ways. (`zlsf` lists users or processes using files or file systems, and lists files or file systems in use by users or processes.)
- If you use the Tools and Toy version, then decide if you want to continue to use it.
- z/OS' `zlsf` installs into /bin. Make sure that does not interfere if you wish to continue to keep your Tools and Toy `zlsf`.

Migration Actions Pre-First IPL:

Update applications that use SMF type 92 subtype 11 close records (Required-IF, as of V2R1)

- Before V2R1, SMF type 92, subtype 11 close records were written when sockets or char special records were closed. They have a very high close rate and do not provide useful information to determine how often files and directories are used.
- As of V2R1, they are written as subtype 16 close records instead, and will not appear in type 92 subtype 11.
- If you want to process this information, you need to include type 92 subtype 16 records in SMFPRMxx to collect them.

33

© 2015 IBM Corporation

z/OS UNIX Migration Actions for z/OS V2R2

These migration actions were taken from z/OS V2R2 Migration. Some descriptions and actions have been shortened for inclusion in this presentation. Not all migration actions have been included. For the complete descriptions and actions, refer to z/OS V2R2 Migration.

[z/OS UNIX System Services Migration Actions You Can Do Now](#)

Accommodate the new Shell and Utilities version of the `zlsnf` utility (Required-IF, as of V2R1)

Required if you currently use the Tools and Toys version of the `zlsnf` utility.

Before z/OS V2R1, the `zlsnf` utility was obtained from the Tools and Toys section of the z/OS UNIX website.

Starting in z/OS V2R1, Shell and Utilities support of the `zlsnf` utility has been added. The supported version differs from the Tools and Toys version in a number of ways. For example, the new `zlsnf` version includes support for displaying file lock holders and waiters when the byte range lock manager is used.

Migration action: Look for current use of the Tools and Toys version of `zlsnf`. If there is no current use of the Tools and Toys version of `zlsnf`, then no actions or changes are required.

If there is current usage of the Tools and Toys version of `zlsnf`, determine if the command is in `/bin`, or in another directory. Also, determine if you want to preserve the Tools and Toys version in addition to the officially shipped version. Note that `zlsnf` can also reside in data sets where rexx execs can be run.

1. If you want to preserve the Tools and Toys version, ensure that you save it into a directory that z/OS V2R1 will not install into. z/OS V2R1 provides `zlsnf` in the `/bin` directory.
2. If you do not want to preserve the Tools and Toys version and it is in `/bin`, then the installation of z/OS V2R1 automatically replaces the Tools and Toys version with the new officially supported version. If the Tools and Toys version is not in `/bin`, remove it from its current location.

[z/OS UNIX System Services Migration Actions Pre-First IPL](#)

Remove files and directories in `/var/man` (Required-IF, as of V2R2)

Required if you use the `man` command to view the man pages.

Before z/OS V2R2, the `/var/man` directory was used by the `man` command. Starting with z/OS V2R2, the `man` command does not use this directory. Instead, it uses the `/tmp` directory on a per-user basis.

Migration action:

1. Remove the `/var/man` directory and all subdirectories and files on your z/OS V2R2 system. These files are no longer used for the z/OS V2R2 `man` command.
2. The z/OS V2R2 `man` command uses the temporary directory for caching man pages. Each user has their own man page cache in the temporary directory. This change might increase the size of the temporary directory, depending on how much man command usage there is on your system. Monitor space usage for the temporary directory to ensure that there is adequate space for users to issue the `man` command. The temporary directory can be the directory referred to by the `TMPDIR` environment variable, or `/tmp` if `TMPDIR` is not defined.

Update applications that use SMF type 92 subtype 11 close records (Required-IF, as of V2R1)

Required if you want continue to get close records for sockets and character special files.

Before z/OS V2R1, SMF type 92, subtype 11 close records were written when sockets or character special records were closed. Sockets and character special files have very high close rates and do not provide useful information in determining how often files and directories are used. Starting in z/OS V2R1, they are written as subtype 16 close records instead. The sockets and character special file close records will no longer appear in type 92 subtype 11 records. Instead, you need to include type 92 subtype 16 records in the SMFPRMxx parmlib member to collect them.

Migration action:

1. Determine whether you have applications that use SMF type 92 subtype 11 close records. For those applications, `SMF92TYP` is set to `SMF92#CLOSE (11)` for subtype 11. `SMF92CTY` is set to `FT_SOCKET (7)` for sockets and `FT_CHARSPEC (2)` for character special files.
2. Change the application to look at subtype 16 records. `SMF92TYP` will be set to `SMF92#CLSOCCHARSPEC (16)`.

[z/OS UNIX System Services Migration Actions Post-First IPL](#)



of V2R

Use the BPX.UNIQUE.USER profile instead of BPX.DEFAULT.USER (Required-IF, as

Required if you are using BPX.DEFAULT.USER. 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. As of z/OS V2R1 BPX.DEFAULT.USER has been removed.

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
```

RACF APAR OA42554 provides assistance with the conversion to BPX.UNIQUE.USER on z/OS V1R13 and z/OS V1R12. With this APAR you can model the user's home directory path by specifying &racuid in the model user's OMVS segment. Then, when the user's OMVS segment is automatically created, RACF will substitute the correct user ID. For more information on this capability, see the information in APAR OA42554.

Marna's "Big Migs" for Migrating from V2R1 to V2R2



Migration actions as of V2R2 you should not overlook:

-  1. **HTTP Server: Move from Domino to Apache**
2. **BCP: Format the ARM Couple Data Set for long symbol name support**
3. **DFS SMB: Add `_BPX_UNUSEDTASKS=KEEP` to the DFSCNTL environment file**
-  4. **JES2: Activate z11 mode**
5. **JES3: Be aware JES3 honors the DSI setting in SCHEDxx**
6. **z/OS OpenSSH: Accommodate the OpenSSH ported level**

34




© 2015 IBM Corporation

Marna's "Big Migs" for Migrating from R13 to V2R2



Migration actions from the prior slide, plus...

Migration actions as of V2R1 you should not overlook:

-  1. **RACF: Remove BPX.DEFAULT.USER**
-  2. **HTTP Server: Move from Domino to Apache (Recommended)**
3. **BCP: Convert your existing IBM Health Checker for z/OS set-up for automatic start-up**
4. **zFS: Determine whether to accept the new default values for certain zFS variables**
-  5. **FIPS 140 Users: Must use ICSF (5 migration actions)**
6. **DFS: Determine whether to accept the new default values for certain zFS variables in the zFS IOEFSPRM configuration file**

35

© 2015 IBM Corporation

Migrating to z/OS V2.2: Part 2 of 2 Summary



- **General Migration Actions:**
 - New address spaces, new and old data sets, changed checks.
- **BCP Migration Actions:**
 - LOGR CDS at HBB7705, NUMBER(1) default for SMDUPLEX for LOGR CDS, HWIBCPII needs READ authority, PARMDD and REGIONX on JCL EXEC and PROC.
- **Communications Server Migration Actions:**
 - Unique PFID for each stack for RoCE Express, automatic maximum segment size adjustments, new QUEUERTT default, fixed CSM default of 200M, IKE daemon syslogd messages,
- **DFSMS Migration Actions:**
 - REFUCB is enabled, LIST DUMPCLASS output.

36

© 2015 IBM Corporation

Migrating to z/OS V2.2: Part 2 of 2 Summary



- **ICKDSF Migration Actions:**
 - VERIFYOFFLINE default for INIT and REFORMAT.
- **JES2 Migration Actions:**
 - Review applicable JES2 exit changes.
- **z/OSMF Migration Actions:**
 - Migrate configuration, recreate table filters.
- **z/OS UNIX Migration Actions:**
 - /var/man removals.



37

© 2015 IBM Corporation