

**Migrating to z/OS V2.3:
Technical Migration Actions**

IBM Systems
Technical Events
ibm.com/training/events

Z Systems Technical University
9-13 October 2017
Munich, Germany

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

IBM

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 V2.3 from either z/OS V2.2 or V2.1. It is strongly recommended that you attend both sessions for a complete migration picture.

The general availability date for z/OS V2.3 was September 29, 2017.

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®, DB2®, ESCO®, eServer®, FICON®, IBM®, IBM logo®, iSeries®, IIVS, OS/390®, pSeries®, RS6000®, S/390®, VME®, VSE/ESA®, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, 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 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.

ITIL 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 multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

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

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental tests 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.

2

© 2017 IBM Corporation

Migrating to z/OS V2.3: Migration Actions Agenda



- Scope of presentation
- Definition of a "migration action"
- Overview of migration actions for z/OS V2R3 from z/OS V2R2 or V2R1:

- ❖ General Migration Actions
- ❖ BCP
- ❖ Communications Server
- ❖ DFSMS
- ❖ ICSF dependencies
- ❖ Distributed File Service – zFS
- ❖ z/OSMF
- ❖ SDSF
- ❖ z/OS OpenSSH
- ❖ z/OS UNIX
- ❖ KC4z, HLASM, XL C/C++





3

© 2017 IBM Corporation

Scope of Presentation IBM

- This presentation is applicable to z/OS V2R3 migrations from either z/OS V2R2 or V2R1.
- Not fully inclusive of all migration actions, but rather gives you an overview of some migration actions that are:
 - Very important to understand
 - May be common to many users
- Remember: Use *z/OS V2.3 Migration Workflow* (or the *z/OS V2R3 Migration book*) for a complete list of all migration actions.
 - The latest level is **GA32-0889-30**.
 - The specific *Workflow* (or *z/OS Migration book*) 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 V2R2 *for V2.2 -> V2.3 users*
 - Chapter 4: Migration from z/OS V2R1 *for V2.1 -> V2.3 users*







4
© 2017 IBM Corporation

Migration is not Exploitation! IBM


- **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**



Means "don't overlook!"



Means some programmatic assistance is available



Means a cleanup action

5
© 2017 IBM Corporation

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 V2R3 level of code to make these changes, and the changes don't require the z/OS V2R3 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 V2R3.
- **Pre-First IPL.** These are migration actions that you perform after you've installed z/OS V2R3 but before the first time you IPL. These actions require the z/OS V2R3 level of code to be installed but don't require it to be active. That is, you need the z/OS V2R3 programs, utilities, and samples in order to perform the migration actions, but the z/OS V2R3 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 V2R3 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 V2R3 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 V2R3. You need a running z/OS V2R3 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 V2R3 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 V2R3

These elements have new V2R3 migration actions:

- BCP
- BookManager Read
- Communications Server
- Cryptographic Services
- DFSMSdfp
- Distributed File Service – zFS
- HCD
- HLASM
- IBM HTTP Server
- Integrated Security Services
- Knowledge Center for z/OS
- z/OS Management Facility
- JES3
- Language Environment
- Library Server
- OpenSSH
- RMF
- SDSF
- Security Server (RACF)
- XL C/C++
- 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 V2R3 Migration book*

6

© 2017 IBM Corporation

Migration Actions for Elements for z/OS V2R3

When migrating from z/OS V2R2 to z/OS V2R3, the specified elements in the slide above have migration actions. Refer to the *z/OS Migration Workflow* or *z/OS V2R3 Migration* for complete information on the required migration actions for all elements, and if you are on the path from z/OS V2R1. 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 V2R3



• Migration Actions Pre-First IPL: Migration and Exploitation

- **Accommodate new address spaces (Recommended)**
 - **New in V2R3:**
 - **HZR** for Runtime Diagnostics (see later)
 - **jesxEDS for JES2 Email Delivery Service**, used when JES2 email interfaces are used. Needs a user ID assigned to this address space which can same one as for JES2, security work in started procedures table or STARTED class profile.
 - **z/OSMF**, in address spaces IZUANG1 and IZUSVR1 (see later).
 - **SDSFAUX** for SDSF (see later).
 - **New in V2R2:**
 - **IBM HTTP Server Powered by Apache** (covered in Part 1)
 - **IBM Knowledge Center for z/OS**, is an address space that is an instance of the WebSphere Liberty Profile.
 - **z/OSMF**, in address spaces IZUANG1 and IZUSVR1.
 - **SDSFAUX**. SDSF APAR PI43902 (on V2.1 and V2.2)

7

© 2017 IBM Corporation

General Migration Actions for z/OS V2R3



Migration Actions Pre-First IPL:

- **Remove references to deleted syslib data sets and paths (Required)**
 - **Removed in V2R3:** Some **BCP** and **HCD** paths, **SDSF's** AISFLINK, and an **HCD** DLIB.
 - **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.
- **Add references to new syslib data sets and paths (Required)**
 - **New in V2R3:** **z/OS Liberty Embedded** path /usr/lpp/liberty_zos/IBM and BBL.SBBLEEXEC and SBBLJC, **XL C/C++'s** two CBC.SCCN* data sets, **ICSF** SCFSTUB target library.
 - **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
- **Update your health check customization for modified checks (Recom)**
 - **New in V2R3:** 5 checks. **V2R2:** 17 checks
 - **Changed in V2R3:** 3 checks **V2R2:** 10 checks (incl.RACF_SENSITIVE_RESOURCES)
 - **Deleted in V2R3:** 10 checks **V2R2:** 6 checks

8

© 2017 IBM Corporation

Migrating to z/OS V2.3: Technical Migration Actions

General Migration Actions For z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 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 V2R3 systems during your migration, and allow back out from z/OS V2R3 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 V2R3 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 V2R3 systems and specify a Fix Category (FIXCAT) value of **"IBM.Coexistence.z/OS.V2R3"**. 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 V2R3 /etc and /var subdirectories.**

Copy files from your old system to the z/OS V2R3 /etc and /var subdirectories, and then modify the files as necessary to reflect z/OS V2R3 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 files and then copy the new /var files into the copy.

The following z/OS V2R3 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 V2.3: Technical Migration Actions

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

The following z/OS V2R3 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 V2R3 and V2R2.

Add references to new data sets (Required)

Migration action: For z/OS V2R3, the following elements had data sets and paths that were added:

- IBM z/OS Liberty Embedded
- XL C/C++
- ICSF

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

Accommodate new address spaces (Recommended)

Migrating to z/OS V2.3: Technical Migration Actions

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 V2R3:

- JES2 Email Delivery Services (EDS). One new persistent address space is created by JES2 when JES2 email interfaces are used:
 - A job is submitted with the NOTIFY JCL statement that requests notification by an email message.
 - The Notify user message service (SSI 75) is called with email address as the target of a message.
 - The address space is named **jesxEDS**, where *jesx* is the name of the JES2 subsystem. You must ensure that a proper user ID is assigned to the address space. This user identifier does not have to be the same as the user identifier for JES2, but you can avoid unnecessary complexity by using the same one. The user ID is specified either by adding an entry in the started procedures table (ICHRIN03) or by creating a profile in the STARTED class that matches address space name. If you prefer, both the started procedures table and STARTED class profile can be used. This action ensures that the correct user ID is assigned to the address space.
- The following address spaces existed in prior releases, but are now started automatically in z/OS V2R3 during the IPL process:
 - **HZR** Runtime Diagnostics address space
 - **IZUANG1** IBM z/OS Management Facility (z/OSMF) angel process, and **IZUSVR1** IBM z/OS Management Facility (z/OSMF) server process

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. 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*. These address spaces have migration actions in z/OS V2.3.
- **SDSFAUX**. SDSF APAR PI43902 (also applicable to z/OS V2R1) introduces a new address space for SDSF. SDSFAUX supplies several new panels and new commands to retrieve system information.

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.)

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 were new in z/OS V2R3:

- CSAPP_FTPD_ANONYMOUS_JES
- CSAPP_MVRSHD_RHOSTS_DATA
- USS_INETD_UNSECURE_SERVICES
- USS_SUPERUSER
- ZFS_VERIFY_COMPRESSION_HEALTH

The following Health Checks were changed by IBM in z/OS V2R3:

- CSAPP_SNMPAGENT_PUBLIC_COMMUNITY
- CSVTAM_VIT_OPT_STDOPTS
- USS_KERNEL_PVTSTG_THRESHOLD

The following Health Checks were deleted by IBM in z/OS V2R3:

- CSAPP_SMTDP_MAIL_RELAY



Migrating to z/OS V2.3: Technical Migration Actions

- CNZ_CONSOLE_OPERATING_MODE
- ZOSMIGV2R2_NEXT_CS_LEGACYDEVICE
- ZOSMIGV2R2_NEXT_CS_SENDMAILCLIEN
- ZOSMIGV2R2_NEXT_CS_SENDMAILDAEMN
- ZOSMIGV2R2_NEXT_CS_SENDMAILMSA
- ZOSMIGV2R2_NEXT_CS_SENDMAILMTA
- ZOSMIGV2R2_NEXT_CS_SMTTPDDAEMON
- ZOSMIGV2R2_NEXT_CS_SMTTPDMTA
- ZOSMIGV2R2_Next_CS_TFTP

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

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 V2R3



Migration Actions Before Installing:

- **Remove INCLUDE1MAFC(NO) from LFAREA in IEASYSxx (Recommended)**
 - NO was used to override older default behavior to include frames that could be used satisfy fixed 1MB page requests in the available frame count (RCEAFC).
 - As of V2R3, this specification will be ignored. YES will be used, and NO is ignored with MSG IAR051I. The 1 M LFAREA specification is always included in the available frame count (RCEAFC).



Migration Actions Before first IPL:

- **Verify that the new default value for IEASYSxx REAL is acceptable (Required, as of V2R3)**
 - REAL controls the amount of central storage allocated for ADDRSPC=REAL (V=R) jobs. V=R is dedicated storage below 16MB line where virtual addresses are the same as real addresses.
 - As of V2R3, default is 0 (no V=R is created). Prior default was 76KB.
 - SMF Type 30 subtype 4 SMF30SFL field contains this V=R information.
 - If you not specify IEASYSxx REAL= and are unsure if run any V=R jobs:
 - Specify the old default, REAL=76 on z/OS V2.3, and
 - use the Generic Tracker on z/OS V2.3 to identify V=R jobs.



9

© 2017 IBM Corporation

BCP Migration Actions for z/OS V2R3



Migration Actions Before first IPL:

- **Remove commands or logic that start or restart Runtime Diagnostics (HZR) (Required-IF, as of V2R3)**
 - As of V2R3, HZR is started via IEACMD00, which is shipped by IBM parmlib.
 - You no longer need to start or restart it yourself via COMMNDxx or automation. Remove any manual starts of HZR that you had.
 - For instance: `COM= 'S HZR, SUB=MSTR'` in COMMNDxx.
 - As before, HZR runs under the master subsystem, and can be started and stopped. IBM recommends that you allow HZR to be classified into the SYSSTC service class, or place it into an importance 1 single period service class with a high velocity goal.
 - If you started HZR previously with another name, change IEACMD00 accordingly.
 - Perform security customization for HZR, if you hadn't used it before.
 - See *z/OS Problem Management STARTED* class profile information.



10

© 2017 IBM Corporation

BCP Migration Actions for z/OS V2R3



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, to allow latest System Logger features to be used.
 - The highest level, HBB7705, was introduced in z/OS V1R2 (in 2001).
 - Use `D XCF,C,TYPE=LOGR` to see your CDS level. `LEVEL: HBB7705`

Migration Actions Before IPL:

- **Accommodate new default log stream data set base minimum sizes (Required-IF, as of V2R3)**
 - As of V2.3 IXGCNFxx new options indicate if log stream offload and staging data sets are allocated with base minimum default sizes.
 - IBM recommended minimums: at least 1MB for offload, and 10MB for staging.
 - `USEOFFLOADADMIN` and `USERSTAGINGMIN` are YES by default.

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.

11

© 2017 IBM Corporation

BCP Migration Actions for z/OS V2R3




Migration Actions Before First IPL:

- **Ensure TVSAMCOM, TVSMMSG, and REGIONX are not used as job statement symbols (Required-IF, as of V2R3 and V2R2)**
 - As of V2R3 (for `TVSAMCOM`), APAR OA48450 (for `TVSMMSG`), 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. For example, expect problems with:
 - `//PROC1 PROC TVSAMCOM=ABC`
 - `//JSTEP1 EXEC PROC1,TVSMMSG=ABC`
 - `//STEP1 EXEC PGM=MYPROG, PARM=' ®IONX'`


12

© 2017 IBM Corporation


BCP Migration Actions for z/OS V2R3



Migration Actions Before First IPL:

- 
• Format the ARM Couple Data Set (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 at least the V2R2 level of IXCL1DSU format utility.
 - Either from a V2R3 system, or STEPLIB to V2R3 MIGLIB.



With APAR
OA46977
(IBM Coexistence z/OS V2R2)

13
© 2017 IBM Corporation

BCP Migration Actions For V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

BCP Migration Actions You Can Do Now



Ensure that your sysplex uses the distributed mode of console operation (Required-IF as of z/OS V2R3)

Required if you are using shared mode.

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



Remove INCLUDE1MAFC(NO) from the LFAREA parameter in IEASYSxx (Recommended, as of V2R3)

Not required, but the system ignores the use of the INCLUDE1MAFC(NO) specification on the LFAREA parameter in IEASYSxx. If INCLUDE1MAFC(NO) is detected, the system issues message IAR051I and uses the default of INCLUDE1MAFC(YES).

Before z/OS V2R3, you could specify INCLUDE1MAFC(NO) to override the default behavior to include frames that can be used to satisfy fixed 1 MB page requests in the available frame count (RCEAFC). The default, INCLUDE1MAFC(YES), results in less paging even when enough 1 MB frames are available to satisfy requests for fixed 1 MB pages.

Starting with z/OS V2R3, real storage is no longer reserved when the LFAREA parameter is specified. As a result, INCLUDE1MAFC(NO) is no longer applicable.

Migrating to z/OS V2.3: Technical Migration Actions

Use health check RSM_Include1MAFC to determine whether INCLUDE1MAFC(NO) was specified on the LFAREA parameter in IEASYSxx.

Migration action:

1. Check the LFAREA parameter specification in the IEASYSxx member on your pre-z/OS V2R3 system.
2. If you specified INCLUDE1MAFC(NO) on the LFAREA parameter in IEASYSxx, take one of the following actions:
 - a. Leave the INCLUDE1MAFC keyword as is. The system ignores the INCLUDE1MAFC(NO) specification.
 - b. Remove the INCLUDE1MAFC keyword, as INCLUDE1MAFC(YES) is the default and the only accepted specification.
3. 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.

Be aware that NOPASS and NODSI are no longer honored for batch jobs (Required-IF as of V2R3 and OA50215 on V2R2 and V2R1)

Required, if your installation specifies the PPT attributes NOPASS or NODSI for programs that run as batch job steps.

In the SCHEDxx parmlib member, the PPT statement allows your installation to specify a list of programs that require special attributes, or to change the attributes of the IBM-supplied default entries in the Program Properties Table (PPT). Before the system initiates a program, it uses the PPT to determine which, if any, special attributes to apply to the program.

In z/OS V2R1 and later releases with APAR OA50215, the NOPASS and NODSI PPT attributes are ignored by the system for programs that run as batch job steps (non-started tasks). Briefly, these attributes are described, as follows:

- NOPASS: The specified program can bypass security protection (password protection and RACF®).
- NODSI: The specified program does not require data set integrity. That is, the job does not hold an ENQ for the data sets it allocates.

For installations that require these options for batch jobs, APAR OA50215 introduces the new PPT attributes NOPASS_ALLOWBATCH and NODSI_ALLOWBATCH, which you can specify in SCHEDxx. However, the new attributes are not recommended and must be used with care.

If your installation specifies the PPT attributes NOPASS or NODSI for programs that run as batch job steps, these specifications are no longer honored. Instead, the following message is issued: IEF188I PROBLEM PROGRAM ATTRIBUTES ASSIGNED

Migration action: Follow these steps:

1. Check your active SCHEDxx parmlib members for usage of the NOPASS or NODSI attributes. You can display the contents of your active SCHEDxx parmlib members by using the command DISPLAY PPT.
2. For any program for which NOPASS or NODSI is specified, determine whether the program runs as a batch job step. For such programs, consider removing the NOPASS or NODSI attributes. Otherwise, if these attributes are required for the program, replace them with the NOPASS_ALLOWBATCH and NODSI_ALLOWBATCH attributes. Be aware that allowing programs that run as batch job steps to receive the NOPASS_ALLOWBATCH or NODSI_ALLOWBATCH attributes can expose potential integrity issues.

To determine whether the new attributes are in effect, enter the command DISPLAY PPT. In the command output, check for Y in column DA (NODSI_ALLOWBATCH) and Y in column PA (NOPASS_ALLOWBATCH).

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

Migrating to z/OS V2.3: Technical Migration Actions

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,ACOUPL=(*primarydsname,primaryvolume*)
 - SETXCF COUPLE,TYPE=LOGR,PSWITCH
 - SETXCF COUPLE,TYPE=LOGR,ACOUPL=(*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*) ACOUPL(*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 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)`.

BCP Migration Actions Pre-First IPL



Verify that the new default value of REAL is acceptable (Required, as of V2R3)

In parmlib member `IEASYSxx`, the `REAL` parameter controls the amount of central storage that can be allocated for `ADDRSPC=REAL (V=R)` jobs. The `V=R` (virtual=real) area is a dedicated storage area below 16 MB in which virtual addresses are the same as real addresses.

In previous releases, the `REAL` parameter had a default value of 76, which means that 76 KB of `V=R` storage was reserved on the system. In z/OS V2R3, the `REAL` parameter default is changed to 0, which means that no `V=R` area is created.

For improved performance in satisfying storage requests, IBM suggests that you do not create a `V=R` area and instead, set `REAL` to 0 (the new default). `REAL=0` is not valid if your installation runs `V=R` jobs; these jobs might abend.

Use IBM Health Checker for z/OS to determine whether a `V=R` area is defined on your system. The check `IBMRSRSM,RSM_REAL` checks the current setting for the `REAL` parameter in `IEASYSxx`.

Migration action: On a z/OS V2R2 or z/OS V2R1 system, do the following:

1. Review the current setting of the `REAL` parameter on your system by checking the `IEASYSxx` parmlib concatenation. If the `REAL=` setting is specified, you have nothing more to do. Otherwise, proceed to Step 2.
2. Search for the *Storage and Paging Section* of SMF type 30 (common address space work) subtype 4 records. Locate the `SMF30SFL` field within the record. If bit 0 is 1, the `V=R` area is used by a job. Add the setting `REAL=76` to your `IEASYSxx` member to maintain compatibility with z/OS V2R2 and z/OS V2R1.

Alternatively, you can follow these steps on your V2R3 system:

1. Before you IPL the system, review the current setting of the `REAL` parameter on your system by checking the `IEASYSxx` parmlib concatenation. If the `REAL=` setting is specified, you have nothing more to do. Otherwise, proceed to Step 2.
2. Add the setting `REAL=76` to your `IEASYSxx` parmlib member before IPLing the system. After IPL, use the Generic Tracker to identify `V=R` jobs. If any are identified, leave the setting. Otherwise, remove the setting and use the default of 0 for subsequent IPLs.

In the SMF Type 30 subtype 4 record, in the paging and storage section, check bit 0 of the byte labeled `SMF30SFL`. This bit is set to 1 to indicate `V=R` usage for the job step.



Prepare for the removal of support for user key common areas (Recommended, as of V2R3)

Not required, but recommended because this function will not be supported after z/OS V2R3.

The allocating, obtaining, or changing common areas of virtual storage, such that the storage is in user key (8-15), will not be supported after z/OS V2R3. Use IBM Health Checker for z/OS check VSM_ALLOWUSERKEYCSA to examine the setting of the ALLOWUSERKEYCSA option in the DIAGxx parmlib member.

Migration action:

1. If you are running CICS Transaction Server for z/OS, ensure that you are running V5.2 or later version.
2. Check for usage of the STORAGE, GETMAIN, or CPOOL service to obtain common ECSA/CSA storage (subpool 227, 228, 231, 241) that specify a key of 8-15. To aid in finding all instances where user key CSA/ECSA is being obtained, take one or more of the following actions:
 - a. Enable the following example SLIP trap to produce GTF trace records for the obtaining of user key CSA/ECSA storage: `SLIP SET,IF,A=TRACE,ID=UCSA,NUCEP=(IGVVSMG2,0,1),END`
 - b. Set the DIAGxx parmlib statement VSM ALLOCUSERKEYCSA to NO, which is the default. Then, IPL a test system with the updated setting. Any software on your test system that attempts to obtain user key CSA/ECSA by using the GETMAIN, STORAGE, or CPOOL service will fail. The service receives one of the following abends: B04-5C, B0A-5C, or B78-5C.
3. Check for usage of the DSPSERV service to allocate a SCOPE=COMMON data space in a key of 8-15. Do the following:
 - a. Enable the following example SLIP trap to produce GTF trace records for the allocation of user key SCOPE=COMMON data spaces: `SLIP SET,IF,A=TRACE,ID=UCAD,NUCEP=(IAXDKUKY,0,1),END`
4. Check for usage of the CHANGEKEY service to change the storage key of common storage to a key of 8-15.
5. Change the affected software to support having the user key common areas of virtual storage areas protected in a system key, or change the affected software to support the storage not be common to all address spaces. Some alternatives for sharing storage instead of having storage common to all address spaces include the following options:
 - a. Use a SCOPE=ALL data space to share data space storage with select units of work in select address spaces.
 - b. Use IARVSERV SHARE to share below the bar storage with select address spaces.
 - c. Use IARV64 GETSHARED to share above the bar storage with select address spaces.
 - d. Use z/OS UNIX shared memory to share below the bar or above the bar storage with select address spaces.

Remove commands or logic that start or restart Runtime Diagnostics (Required-IF, as of z/OS V2R3)

Required if you have any system automation that starts or restarts Runtime Diagnostics.

In z/OS V2R3, the command to start Runtime Diagnostics is added to the IBM-supplied parmlib member IEACMD00. As a result, Runtime Diagnostics is started automatically during system initialization, when you include the SYS1.IBM.PARMLIB data set in your parmlib concatenation. This change means that it is no longer necessary for you to explicitly start Runtime Diagnostics after each system IPL, whether through the COMMNDxx parmlib member, an automation program, or an operator command entered manually at the console.

For example, in previous releases, you might have placed the start command in your COMMNDxx parmlib member, as follows: `COM='S HZR,SUB=MSTR'`

As in previous releases, the Runtime Diagnostics address space (HZR) continues to run as an address space under the master subsystem and remains active until you stop it with the STOP command. IBM recommends that you allow the HZR address space to be classified into the SYSSTC service class, or place it into an importance 1 single period service class with a high velocity goal.

Migration action: Follow these steps:

1. Remove commands or logic that start Runtime Diagnostics from automation programs or the COMMNDxx parmlib member. For example, remove the statement `COM='S HZR,SUB=MSTR'` from the COMMNDxx parmlib member, if specified.
2. If your installation changed the Runtime Diagnostics default job name in SYS.PROCLIB, update the START command in IEACMD00 to associate the installation defined name with the default job name HZR. For example, if you changed the Runtime Diagnostics job name from HZR to HZRNEW, change the command in IEACMD00 from: `COM='S HZR,SUB=MSTR'` to: `COM='S HZRNEW,SUB=MSTR,JOBNAME=HZR'`

Migrating to z/OS V2.3: Technical Migration Actions

3. When Runtime Diagnostics is started, the following message might be issued if the STARTED class is active. While this message is not a change to Runtime Diagnostics processing, the message might be new to you: IRR813I NO PROFILE WAS FOUND IN THE STARTED CLASS FOR HZR WITH JOBNAME HZR. RACF WILL USE ICHRIN03.
 - a. If you receive this message, one of the following problems has occurred:
 - i. The STARTED class has the RACLIST option, but no SAF profile exists for Runtime Diagnostics (HZR)
 - ii. A SAF profile is defined for HZR, but the STARTED class is not RACLIST enabled
 - iii. A SAF profile is defined for HZR, but the STARTED class was not refreshed so that the changes could take effect.
 - iv. The RACLIST profile is not enabled on the system.
 - b. In response, RACF uses the information in the started procedures table (ICHRIN03) to assign security information for HZR . Either ensure that the proper definition exists in the ICHRIN03 for HZR, or take one of the following actions:
 - i. Define a SAF profile for HZR in the STARTED class, as described in *z/OS Problem Management*.
 - ii. Ensure that the STARTED class has the RACLIST option. For example: SETROPTS RACLIST(STARTED)
 - iii. Refresh the STARTED class. For example: SETROPTS RACLIST(STARTED) REFRESH
- If you have a security profile set up for Runtime Diagnostics (HZR) in the STARTED class before IPL, message IRR813I is not issued. Security for Runtime Diagnostics works as it did in previous releases.

Stop using the DRXRC duplex mode option for logstreams (Required-IF, as of V2R3)

Required if you intend on having a mixed-release level sysplex, and you currently use the DRXRC option.

Starting in z/OS V2R3, system logger no longer supports the DRXRC duplex mode option for logstreams. IBM recommends that you use other available mirroring options with IBM z/OS Global Mirror (zGM), also known as Extended Remote Copy (XRC), or GDPS. The withdrawal of this support has no impact on any other logstream configurations. Use IBM Health Checker for z/OS check ZOSMIGV2R2_Next_IXG_Remove_DRXRC. This migration health check was provided by APAR OA49507 for z/OS V1R13, V2R1, and V2R2.

Migration action: Follow these steps:

1. To enable health check ZOSMIGV2R2_Next_IXG_REMOVE_DRXRC to run, give the Health Checker user ID READ access to the MVS.DISPLAY.LOGGER resource in the OPERCMDS class. Also, when the FACILITY class is activated and a profile is defined that covers the MVSADMIN.LOGR resource, you must grant the user ID that is associated with the Health Checker address space READ access to this resource. For example, you might specify the following: RDEFINE FACILITY MVSADMIN.LOGR UACC(NONE)

```
PERMIT MVSADMIN.LOGR CLASS(FACILITY) ID(hcsuperid) ACCESS(READ)
```

```
SETROPTS CLASSACT(FACILITY)
```

```
SETROPTS RACLIST(FACILITY)
```

If you had already RACLISTed the FACILITY class, the last statement would have to be: SETROPTS RACLIST(FACILITY) REFRESH See the topic on LOGR parameters for administrative data utility in *z/OS MVS Setting Up a Sysplex*.

2. For any exceptions reported by the health check ZOSMIGV2R2_Next_IXG_REMOVE_DRXRC, follow the actions that are described in the following message: IXGH013E One or more log stream definitions in this sysplex has the DUPLEXMODE(DRXRC) specification. Otherwise, the following informational message is issued: IXGH012I This sysplex does not contain any log streams with the DUPLEXMODE(DRXRC) specification.

3. Alternatively:

- a. You can use the IXCMIAPU utility for DATA TYPE(LOGR) REPORT(YES), and identify any logstreams with the STG_DUPLEX(YES) DUPLEXMODE(DRXRC) attributes. Sample report job:

```
//LOGRPT JOB
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DATA TYPE(LOGR) REPORT(YES)
/*
```

If the IXCMIAPU utility identifies no logstreams with the DUPLEXMODE(DRXRC) option specified, you have no impact. You can skip the next step.

- b. For any logstreams with the DUPLEXMODE(DRXRC) specification, you must update the logstreams to use a different duplex option. You can run the IXCMIAPU utility or IXGINVNT API to update the logstream by specifying either DUPLEXMODE(COND) or DUPLEXMODE(UNCOND).

Migrating to z/OS V2.3: Technical Migration Actions

Or, set the option STG_DUPLEX(NO) to avoid having a staging data set used to duplex the logstream data.

Accommodate the new default log stream data set base minimum size (Required-IF, as of V2R3)

Required if you do not want the new logger behavior.

In z/OS V2R3, options are added to the IXGCNFxx member of parmlib so that you can indicate whether log stream offload and staging data sets are allocated on the system with base minimum default sizes. Specifically, the new options USEOFFLOADADMIN and USESTAGINGMIN are YES, by default, for the associated MANAGE OFFLOAD and MANAGE STAGING statements. This change means that the system ensures that log stream data sets are located on the system with a base minimum size of at least 1 MB and 10 MB for offload and staging data sets, respectively. These amounts are the IBM recommended minimum sizes.

A significant negative performance impact on log stream usage can occur when offload or staging data sets are sized too small. This problem can happen when base system defaults are used that result in data set sizes in the 2 - 3 track size range.

Migration action: Unless the log stream exploiter specifically recommends a size below these amounts, IBM recommends the following:

- Log stream staging data sets be sized no smaller than 10 MB
- Offload data sets be sized no smaller than 1 MB.

Check the DASD management policies for your log stream use and determine whether the new logger policy of ensuring the base minimum sizes of these data sets is appropriate for your installation. If so, no further action is necessary.

Otherwise, if the new log stream data set management behavior is not appropriate for your installation, you must provide an IXGCNFxx parmlib member that specifies MANAGE OFFLOAD USEOFFLOADADMIN(NO) or MANAGE STAGING USESTAGINGMIN(NO) for the appropriate log stream data set types.

Prepare for log stream staging data set sizes greater than 4 GB (Required, as of V2R3)

Required to ensure compatibility. Log stream staging data sets can already be allocated without any of the changes in z/OS V2R3.

Starting in z/OS V2R3, system logger allows log stream staging data set sizes greater than 4 GB. In previous releases, you could request more than 4 GB for a log stream staging data set, but system logger would allocate a smaller amount of storage for the data set (about 3.5 GB). The new support in z/OS V2R3 means that the z/OS V2R3 release systems allow log stream staging data set sizes with a size greater than 4 GB when the staging data sets are newly allocated on that system. Any z/OS V2R2 and z/OS V2R1 release levels in the same sysplex require the appropriate coexistence support to be compatible with z/OS V2R3 and correctly manage the larger data sets.

Migration action: Do the following:

- a. Install the PTFs for APAR OA49506 on all z/OS V2R2 and z/OS V2R1 systems in the same sysplex as the z/OS V2R3 systems.
- b. Understand that if you request an allocation of greater than 4 GB for a log stream staging data set, rather than receiving a smaller allocation, you now receive the requested size.

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.



Format the ARM couple data set (Required-IF, as of V2R2)

Required if systems are intended to use ARM functions.

Migrating to z/OS V2.3: Technical Migration Actions

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 TVSAMCOM, TVSMMSG, or REGIONX are not used as job statement symbols

(Required-IF, as of V2R3 and V2R2)

Required if you used symbol names **TVSAMCOM**, **TVSMMSG**, or **REGIONX** on EXEC or PROC statements in jobs.

New JCL keywords are added to the JCL EXEC statement, as follows:

- TVSAMCOM is added in z/OS V2R3
- TVSMMSG was added with APAR OA48450 for z/OS V2R1 and V2R2
- REGIONX was added in z/OS V2R2.

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 any of the symbolic parameter names TVSAMCOM, TVSMMSG, or REGIONX on the EXEC or PROC statement, you must edit the jobs to use alternatively named symbols. Otherwise, the jobs can fail with JCL errors.

Migration action: Search for a symbol that is named TVSAMCOM, TVSMMSG, 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 that is named TVSAMCOM, TVSMMSG, or REGIONX.

Example:

```
//PROC1 PROC TVSAMCOM=ABC
//PROC1 PROC TVSMMSG=ABC
//PROC1 PROC REGIONX=ABC
```

- EXEC statements that contain a symbolic parameter that is named TVSAMCOM, TVSMMSG, or REGIONX.

Examples:

```
//JSTEP1 EXEC PROC1,TVSAMCOM=ABC
//JSTEP1 EXEC PROC1,TVSMMSG=ABC
//JSTEP1 EXEC PROC1,REGIONX=ABC
```

- EXEC statements that contain a '&TVSAMCOM', '&TVSMMSG' or '®IONX' parameter value string.

Examples:

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

For any occurrences that you find, change the JCL statement to refer to a different symbolic parameter name.

Accommodate system symbol names that contain underscores (Required-IF, as of z/OS V2R2)

Required if you have situations in which a symbol name might be followed immediately by an underscore

z/OS V2R2 enhances the use of system symbols in the following ways:

- Longer system symbol names (up to 16 characters) and longer symbol substitution values
- Underscores (_) can be specified in any character position other than the first one.

System symbols are typically used in started procedures and jobs, parmlib members, and other objects. Information about using system symbols is provided in *z/OS MVS Initialization and Tuning Reference*.

Migrating to z/OS V2.3: Technical Migration Actions

For objects that undergo system symbol substitution, be aware that the use of underscores in symbol names can result in an incompatibility if a symbol reference is followed immediately by an underscore (that is, without a symbol-delimiting period). For example, in previous releases, specifying &SYM_A in a file would match the symbol &SYM and add "_A". In z/OS V2R2, this specification can match the symbol &SYM or the symbol &SYM_A. Here, a match is attempted first for &SYM_A (with the underscore). If no match is found, an attempt is made to match &SYM.

Note: As of z/OS z/OS 2.1, underscores can be specified in JES symbol names. Thus, this migration action does not apply to JES symbol names.

Migration action: In z/OS V2R2, IBM provides a REXX exec, ASASYMUN, to help you locate data that might encounter unexpected results if symbol names have underscores. You can run the exec on z/OS V2R2 and older systems.

ASASYMUN scans a PDS or PDSE for situations where a symbol name that contains an underscore might cause different results than expected. For example, it checks for a symbol (an ampersand followed by other characters) followed immediately by an underscore with no delimiting period. Change these lines to add a period before the underscore to delimit the symbol.

ASASYMUN is supplied by IBM in the SBLSCLI0 data set. To use ASASYMUN, you must invoke it from an ISPF environment. Follow these steps:

1. From the ISPF command line, run the exec as follows:

```
TSO EXEC EXECDSN(ASASYMUN) 'scandsn'
```

Where:

execdsn

is the cataloged data set containing the ASASYMUN exec. Use 'execdsn(ASASYMUN)' if the data set name is fully qualified. You can use alternative forms of execdsn(ASASYMUN) if the data set is in the SYSEXEC or SYSPROC concatenation of the user. For example, you might be able to specify just (ASASYMUN).

scandsn

is the PDS or PDSE you want to scan. Specify a data set contain statements subject to symbol substitution, like JCL or parmlib statements. If the data set name is fully qualified, double the quotations around the name in addition to the single quotation marks that surround the parameter. For example, specify "'fully.qual.dsn'" for a fully qualified data set.

2. For any references you find that were intended to be resolved by a symbol, add a period before the underscore to delimit the symbol. For example:

```
&SYM._A
```

In a sysplex of mixed releases of z/OS systems, if you use symbol names with underscores, you must ensure that earlier systems can handle the symbol names. Install the toleration PTF for APAR OA46739 on the earlier systems. The PTFs for this APAR are identified with the SMP/E FIXCAT IBM.Coexistence.z/OS.V2R2.

Examine your IEFUSI exit routine for possible changes (Required-IF, as of z/OS V2R2)

Recommended, to ensure that users who specify a REGIONX value in their JCL receive the correct storage allocation, according to your installation's requirements.

z/OS V2R2 includes a number of functional enhancements to support improved region management, including:

REGIONX keyword

New JCL keyword for the JOB and EXEC statements. JCL programmers can use REGIONX to explicitly request precise below-the-line and above-the-line storage amounts.

SMFLIMxx parmlib member

New parmlib member. Your installation can use SMFLIMxx to set rules for the REGION and MEMLIMIT values in job steps, or cancel job steps that violate the rules.

These changes include new parameter input fields for the IEFUSI exit routine. If your installation uses an IEFUSI exit routine to control job region size, it is recommended that you examine the routine for possible changes that you might need to make.

Note the following changes in the exit input parameter list:

- If the REGIONX keyword is used, the "region requested" value in sub-word 2 (as pointed to by word 5) is updated to contain a value based on the two values that were supplied on the REGIONX keyword. This change allows an existing exit routine to continue to receive the total amount of storage that was requested by the job step. Specifically, the value now contains either of the following values:
 - A value of 0 when the REGIONX second parameter is 0M, 0K, or 0G
 - The larger of the specified values when the REGIONX second parameter is not 0M, 0K, or 0G

Migrating to z/OS V2.3: Technical Migration Actions

If the REGIONX keyword is not used, the IEFUSI input parameter list contains one "region requested" value in sub-word 2. This behavior is the same as in previous releases of z/OS.

- For job steps that include the REGIONX keyword, sub-word 1, bit 3 (as pointed to by word 5), indicates that separate values are also available. Here, two more words are provided, sub-word 7 and 8 (as pointed to by word 5), which provide the below-the-line and above-the-line values on the REGIONX keyword.

If the total size indicated in sub-word 2 is not specific enough for your purposes, you can have your exit routine use the values in sub-word 7 and 8 to set the existing below- and above-the-line output fields in sub-words 2, 3, 4, and 5, as pointed to by word 5 of the input parameter list.

No additions or changes to the output parameter list are needed for REGIONX-related processing. The output parameter list allows for specifying separate above- and below-the-line values, as in previous releases of z/OS.

After the IEFUSI exit routine runs, the system checks the SMFLIMxx member for rules that might override region limits. This processing allows the existing IEFUSI exit to continue to set region and MEMLIMIT values for its various functions, including functions that are not supported by SMFLIMxx, such as setting limits on data space blocks. The SMFLIMxx member can be used to set values for new work or to override IEFUSI values for changes to existing work, thus reducing the need for more IEFUSI exit code changes.

In some cases, however, your installation might require that the IEFUSI exit routine make the final determination. If so, the exit routine can set a new flag (sub-word 1, bit 4, as pointed to by word 5) to bypass the SMFLIMxx rules. This bit essentially disables the SMFLIMxx processing for the current job step.

Note: A sample IEFUSI exit routine is provided in SYS1.SAMPLIB in member IEEUSI. For more information about changes that might be needed, see the commented sections in the sample IEFUSI exit routine.

Migration action: Follow these steps:

- If your installation uses an IEFUSI exit routine, examine it to determine whether changes are required for the REGIONX keyword and SMFLIMxx member processing. It is possible that no action is needed, if the decisions made by the exit routine are generic, such as setting the region above-the-line value 128 MB for all job steps. If the exit sets storage values using sub-word 2 (as pointed to by word 5) and a job uses REGIONX, the exit continues to set the desired value. Here, the REGIONX value is ignored and the IEFUSI selected storage value is used. SMFLIMxx processing, if activated by specifying SMFLIM= as part of IPL parameters or through the SET SMFLIM= operator command, overrides the IEFUSI exit-returned values, if environmental conditions match the job step's current environment.
- If you determine that your IEFUSI exit routine requires updating, you can modify your routine by using the new values in the input parameter list. For example, you can have your exit routine do the following:
 - Use the new REGIONX values for below- and above-the-line storage for region size determinations. These values are provided in sub-word 7 and 8 (as pointed to by word 5), when sub-word 1, bit 3 is set to 1.
 - Bypass SMFLIMxx processing for certain jobs by setting sub-word 1, bit 4 (as pointed to by word 5). Or, you can remove all region and MEMLIMIT processing from your routine, and create SMFLIMxx statements to limit the REGION and MEMLIMIT values. Here, you might retain your routine to perform other types of processing, such as setting limits on data space blocks.

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 VSMLOC 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 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,

Migrating to z/OS V2.3: Technical Migration Actions

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 ('IXLCTCAD') 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 IXLCTCAD, which limited the buffers to a 2 GB range of addresses. In z/OS V2R2, XES no longer creates the IXLCTCAD 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 IXLBCAD 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 IXLCTCAD object, which is no longer created in z/OS V2R2. Specifically, check for references to 'IXLCTCAD' on the DSPNAME parameter of the **DUMP** and **SLIP** commands (that is, DSPNAME=('XCFAS'.IXLCTCAD)) and on any automated parse routines.
- Ensure that SDATA=XESDATA is specified on any **DUMP** or **SLIP** commands where the IXLCTCAD name was removed. This setting causes the XES CTRACE 64-bit common storage to be included in an SVC dump.

Failure to remove the IXLCTCAD 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

Migrating to z/OS V2.3: Technical Migration Actions

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=xxxxx 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(leaSlipConfirm) 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
```

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.

DFSMSdfp Migration Actions for z/OS V2R3

• Migration Actions Before Installing:

- **Verify that the new default for CA RECLAIM is acceptable (Required-IF, as of V2R3)**
- Intro in z/OS V1R12 to reduce the need for reorganizing a VSAM KSDS, empty control area (CA) space on DASD can be reclaimed automatically, so that it can be reused later when a CA split is required. The reclaimed CAs are available to be used for new records without any processing to obtain new space.
 - IGDSMSxx default was `CA_RECLAIM(NONE)`, meaning KSDS data sets will not use Control Area Reclaim, regardless of the data class specification.
- In z/OS V2R3, the default is `CA_RECLAIM(DATACLAS)`, meaning both SMS-managed and non-SMS-managed KSDS data sets use the CA reclaim attribute in the data class.
 - Data class attribute is set to `YES` to enable CA reclaim by default.
 - Control usage at the data class level, as you desire.
 - `IDCAMS ALTER ksd_s_name RECLAIMCA | NORECLAIMCA`
 - More I/O is required to maintain reclaimed CAs, so balance that if no or very few CA splits to reuse empty CAs. Use `EXAMINE DATASET` command, which shows the number of empty CAs.

14 • Can be changed dynamically with `SETSMS`.

© 2017 IBM Corporation

DFSMSdfp Migration Actions for z/OS V2R3

• Migration Actions Before Installing:


- **Position for Data Set Encryption (Required, as of V2R3, and V2R2 and V2R1 with APAR OA50569)**
- You must control the creation of encrypted data sets and prevent loss of access to data on any system (backup systems, DR, replication target systems,...).
 - These following steps assure that encrypted data sets are not created until you are ready to encrypt and decrypt*.
 1. Restrict access to the SAF FACILITY class resource `STGADMIN.SMS.ALLOW.DATASET.ENCRYPT` until all systems have installed OA50569 and the minimum hw. UACC of `NONE` is advised.
 2. If RACF FIELD class is active, check for any profile that would allow any user without `SPECIAL` to access to `DATASET.DFP.DATAKEY`.
 - If there is such a profile, create profile `DATASET.DFP` in the FIELD class with UACC of `NONE`.
 3. Do not create `DATASET` profiles with the `KEYLABEL` field in the DFP segment until *all systems that will have access to the encrypted data* meet the sw and hw requirements to exploit data set encryption.

* Exploitation of data set encryption has several considerations, including HW and multisystem dependencies. Data set encryption can be safely tested by informed users with test data prior to making it generally available for your applications. Carefully follow instructions found in *DFSMS Using New Functions*.

15 © 2017 IBM Corporation

DFSMSdfp Migration Actions for z/OS V2R3

- **Migration Actions Before Installing:**
 - **Determine whether you need DISABLE(REFUCB) in parmlib member (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)



16 © 2017 IBM Corporation

DFSMS Migration Actions for z/OS V2 R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

DFSMS Migration Actions You Can Do Now

Verify that the new default for CA RECLAIM is acceptable (Required-IF, as of V2R3)

Required if you don't want CA Reclaim activity on your system.

To reduce the need for reorganizing a KSDS, empty control area (CA) space on DASD can be reclaimed automatically, so that it can be reused later when a CA split is required. The reclaimed CAs are available to be used for new records without any processing to obtain new space.

As of z/OS V1R12, in parmlib member IGDSMSxx, the statement CA_RECLAIM specifies whether to use CA reclaim for KSDS data sets, based on the value of the CA Reclaim attribute in the associated data classes.

In z/OS V2R3, the default for CA_RECLAIM is changed to DATACLAS, which indicates that both SMS-managed and non-SMS-managed KSDS data sets use the CA reclaim attribute in the data class, which you set with ISMF. The attribute is set to Yes to enable CA reclaim by default.

In previous releases, the default was CA_RECLAIM(NONE), which indicates that none of the KSDS data sets use CA reclaim, regardless of the data class specification.

As of z/OS V2R3, if you do not want CA reclaims to be performed, you must specify CA_RECLAIM(NONE).

Use check IBMVSAM,VSAM_CA_RECLAIM to determine whether VSAM CA reclaim is enabled. This check is invoked during initialization and whenever CA reclaim status is changed. This check was added to z/OS V2R1 and V2R2 by APARs OA51394, OA51393, and OA51002.

Migration action: For efficient DASD space usage, IBM recommends that you run your system with CA reclaim enabled.

CA reclaim can provide for improved DASD space usage, but it requires more I/O to keep track of the reclaimed CAs so that they can be reused. The cost of this I/O might not be justified if there are no or very few CA splits to reuse empty CAs. To determine whether CA reclaim is desirable for a data set, use the **EXAMINE DATASET** command, which shows the number of empty CAs in a KSDS with message IDC01728I.

If you do not want the default CA_RECLAIM(DATACLAS), you can override it by specifying CA_RECLAIM(NONE) in your IGDSMSxx parmlib member. The DATACLAS setting allows you to disable or enable CA reclaim at the data class level. You can also disable or enable the CA reclaim setting for specific data sets by using the command **IDCAMS ALTER ksd_name RECLAIMCA | NORECLAIMCA**.

Migrating to z/OS V2.3: Technical Migration Actions

After IPL, you can dynamically change the CA_RECLAIM setting by using the **SETSMS** command.



DFSMSdsp: Position for data set encryption (Required, as of V2R3, and z/OS V2R2 and V2R1 with OA50569)

Required.

The steps below are intended to assure that encrypted data sets are not created until the installation is ready to encrypt and decrypt. Until the decryption functions are available on all sharing systems (including backup systems, and disaster recovery systems), access to encrypted data can be lost at any time.

Migration action: To control the creation of encrypted data sets and prevent loss of access to data on any system that does not have the support, the following actions need to be taken before the software is installed.

1. Restrict access to the SAF FACILITY class resource STGADMIN.SMS.ALLOW.DATASET.ENCRYPT until all systems in your installation have installed the PTFs for OA50569 and the minimum hardware. To do this, you can define the STGADMIN.SMS.ALLOW.DATASET.ENCRYPT profile in the FACILITY class, and set the universal access to NONE. For example:

```
RDEFINE FACILITY STGADMIN.SMS.ALLOW.DATASET.ENCRYPT UACC(NONE)
```

2. If the RACF FIELD class is active, check for any profile that would allow any user without SPECIAL attribute access to the DATASET.DFP.DATAKEY. If there are none, no additional action is needed. If there is any profile that would allow access to DATASET.DFP.DATAKEY, create a DATASET.DFP.DATAKEY profile in the FIELD class with a UACC of NONE. For example:

```
RDEFINE FIELD DATASET.DFP.DATAKEY UACC(NONE)
```

3. Do not create DATASET profiles with the KEYLABEL field in the DFP segment until all systems in your installation have met all software and hardware minimum requirements.

Carefully follow exploitation documentation provided in *z/OS DFSMS Using New Functions*, before you start encrypting data sets.



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 FLASHCPY, 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.

Migrating to z/OS V2.3: Technical Migration Actions

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 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: Accommodate change for data set name prefix in IDCAMS ALLOCATE (Required-IF, as of z/OS V2R1 and z/OS V1R13, both with APARs OA42679 (DFSMS) and OA43330 (TSO/E))

*Required if you use **IDCAMS ALLOCATE**, do not specify the data set name in quotation marks on the DATASET keyword, and the user ID assigned to the IDCAMS batch job does not have a RACF TSO segment.*

With IDCAMS APAR OA42679 and TSO/E APAR OA43330 applied, the **IDCAMS ALLOCATE** command is changed in the way that it starts TSO/E to allocate a data set. IDCAMS processing now uses the TSO/E Service Facility (TSF) to allocate a data set, rather than running the **ALLOCATE** command under the TSO/E terminal monitor program (IKJEFT01).

With this change, the user ID assigned to the IDCAMS batch job is treated as the default data set prefix. That is, the user ID for the IDCAMS batch job is appended to the data set name as a high-level qualifier, if you specify the data set name on the DATASET keyword without quotation marks and the user ID does not have a RACF TSO segment. Previously, the **IDCAMS ALLOCATE** command used a null prefix for the allocated data set, if you specified the data set name on the DATASET keyword without quotation marks and the user ID did not have a RACF TSO segment. Assume, for example, that the user ID ZZZZZZ is defined in both UADS and in RACF without a TSO segment; note the following differences in behavior:

- Before this change, TSO runs under the UADS user. If the data set name is specified without quotation marks on the DATASET keyword, and the user has a UADS PROFILE PREFIX(*prefix*), the prefix is used as the data set prefix. Otherwise, the user ID is used as the data set prefix.
- After this change, the user ID is always used as the data set prefix. Therefore, if the user ID and UADS PROFILE PREFIX(*prefix*) are different, the high-level qualifier for the data set is changed.

Migration action: Check for JCL and programs that use the **IDCAMS ALLOCATE** command. Ensure that the data set name is specified in quotation marks on the DATASET keyword. Doing so ensures that the user ID is not

Migrating to z/OS V2.3: Technical Migration Actions

appended to the data set name as a high-level qualifier. Failure to include the data set name in quotation marks can result in allocation errors.

When the **IDCAMS ALLOCATE** command is run by a user with a RACF TSO segment defined, no change is required.

DFSMSHsm: 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 DUMCPLASS with a output target of TERM, OUTDATASET, or SYSOUT. Also, update applications as needed for the new MINSTACK field in the LIST DUMCPLASS with an output target of TERM, OUTDATASET, or SYSOUT.

Migration Actions for z/OS V2R3



Is ICSF is running on each and every system?



- **(V2R3) Network Authentication Service (Kerberos) relies on ICSF PKCS#11 callable services. These IP Services functions use Kerberos, and user IDs associated with them might therefore need access to those ICSF callable services, when you:**
 - Authenticate clients of UNIX System Services Telnet server,
 - Do z/OS FTP client or server authentication, or
 - Authenticate clients of the UNIX System Service RSH server.
- **(V2R3) PKI Services replaces an internal function with PKCS#11 Digest functions.**
 - If the CSFSERV resource class is active, ensure that the z/OS PKI Services daemon user ID has READ access to the following ICSF resources: CSFOWH, CSFIQA and CSF1TRL.

A general recommendation: **have ICSF up and running on every system* for everything that has a dependency on it.**
...and have it available on the system before the functions that need it.

17

* GDPS does not require that the K System start ICSF, and therefore this recommendation does not include GDPS.

© 2017 IBM Corporation



Migration actions for using ICSF in z/OS V2R3


Increasingly, more functions in z/OS are relying upon ICSF services and need ICSF active on your system. As of z/OS V2.1, FIPS 140 required ICSF. More functions in z/OS V2.3 also have been changed to use ICSF. If you are not running ICSF on each and every system in your enterprise, you will be limiting your functions and also forgoing secure, high-speed cryptographic service in your z/OS environment.

You should review your systems to ensure that ICSF is active is available on each z/OS LPAR.* Refer to the Cryptographic Service ICSF: System Programmer's Guide for activation instructions:


https://www.ibm.com/support/knowledgecenter/SSLTBW_2.3.0/com.ibm.zos.v2r3.csfb200/in2.htm

*As is always the case with GDPS, use GDPS requirements when configuring the K System LPAR.

DFS Migration Actions for z/OS V2R3



The only zFS migration for V2R3!

Migration Actions you can do NOW: 

zFS: Accommodate changed defaults (Required-IF, as of V2R3)

- **romount_recovery** option is changed from OFF to ON.
 - With `romount_recovery=ON`, the file system is temporarily mounted read/write to allow log recovery to run. Then, the file system is remounted read/only.
- **format_aggrversion** option is changed from 4 to 5.
 - With `format_aggrversion=5`, the default version of the aggregate is 5.
- **change_aggrversion_on_mount** option is changed from OFF to ON.
 - With `change_aggrversion_on_mount`, version 4 aggregates are changed to version 5 aggregates on a primary read/write mount.
- **honor_syslist** option is changed from OFF to IGNORED.
 - If the `honor_syslist=ignored` option is specified, it is accepted but not used.
 - This says whether to use the z/OS UNIX automove options when determining the new zFS owner.

18
© 2017 IBM Corporation

DFS Migration Actions For z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

[DFS Migration Actions You Can Do Now](#)

zFS: Accommodate the changed defaults in the IOEFSPRM configuration file (Required-IF, as of V2R3)

Yes, if the changed defaults are not appropriate for your environment.

Starting in z/OS V2R3, several default settings in the IOEFSPRM configuration file are changed, as follows:

- The default value for the `romount_recovery` option is changed from OFF to ON. With `romount_recovery=ON`, the file system is temporarily mounted read/write to allow log recovery to run. Then, the file system is remounted read/only.
- The default value for the `format_aggrversion` option is changed from 4 to 5. With `format_aggrversion=5`, the default version of the aggregate is 5.
- The default value for the `change_aggrversion_on_mount` option is changed from OFF to ON. With `change_aggrversion_on_mount`, version 4 aggregates are changed to version 5 aggregates on a primary read/write mount.
- The default value for the `honor_syslist` option is changed from OFF to IGNORED. If the `honor_syslist=ignored` option is specified, it is accepted but not used.

Migration action:

For `romount_recovery=on`:

- If you want recovery to be automatically performed for read-only mounts and `romount_recovery` is not specified in IOEFSPRM, there is no migration action.
- If you do not want recovery to be automatically performed for read-only mount, ensure that `romount_recovery=off` is in IOEFSPRM. Alternatively, you can change the `romount_recovery` value for this mount instance by using **`zfsadm config -romount_recovery off`**.

For `format_aggrversion`:

Migrating to z/OS V2.3: Technical Migration Actions

- If you want the format operation to produce version 5 aggregates when neither -version4 nor -version5 is specified, and format_aggrversion is not in IOEFSPRM, there is no migration action.
- If you want the format operation to produce version 4 aggregates when neither -version4 nor -version5 is specified, ensure that format_aggrversion=4 is in IOEFSPRM. Alternatively, you can change the format_aggrversion value for this mount instance by using **zfsadm config -format_aggrversion 4**.

For change_aggrversion_on_mount:

- If you want version 4 aggregates to be changed to version 5 aggregates on a primary read/write mount and change_aggrversion_on_mount is not in IOEFSPRM, there is no migration action.
- If you do not want version 4 aggregates to be changed to version 5 aggregates, then ensure that convert_aggrversion_on_mount=off is in IOEFSPRM. Alternatively, you can change the change_aggrversion_on_mount value for this mount instance by using **zfsadm config -change_aggrversion_on_mount off**.


For honor_syslist:

- If existing behavior for zFS ownership movement needs to change, see Dynamic movement of the zFS owner in *z/OS Distributed File Service zFS Administration*


z/OSMF Migration Actions for z/OS V2R3



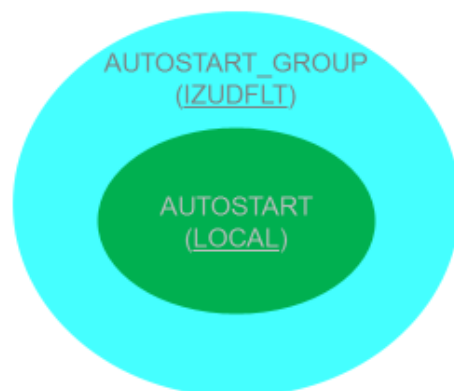
Migration Actions you MUST do NOW:

- **Prepare for autostart (Required in V2R3)**
 - By default, z/OSMF (IZUANG1 and IZUSVR1) will start by default. Several preparations can be done now to make that smoother for your first z/OS V2.3 IPL.
 - Decide your scenario, which system(s) will start z/OSMF with IZUPRMxx's AUTOSTART (LOCAL) and which will use AUTOSTART (CONNECT) *.
 - Decide what AUTOSTART_GROUP* names to use.
 - If you disable z/OSMF autostart, none of the z/OSMF capabilities is available on your system until you start a z/OSMF server manually.
 - For instance, z/OS V2.3 JES2 uses z/OSMF server to deliver email messages.
 - **Understand z/OSMF requirements (Required in V2R3)**
 - Java 8 64-bit, minimum 4GB memory. * new in IZUPRMxx for V2R3
- 19  **Sage advice:** start z/OSMF today on z/OS V2.2 or V2.1, and there are far fewer migration actions for z/OS V2.3! © 2017 IBM Corporation

z/OSMF Migration Actions for z/OS V2R3



Autostart Use Case #1: Monoplex (single system)

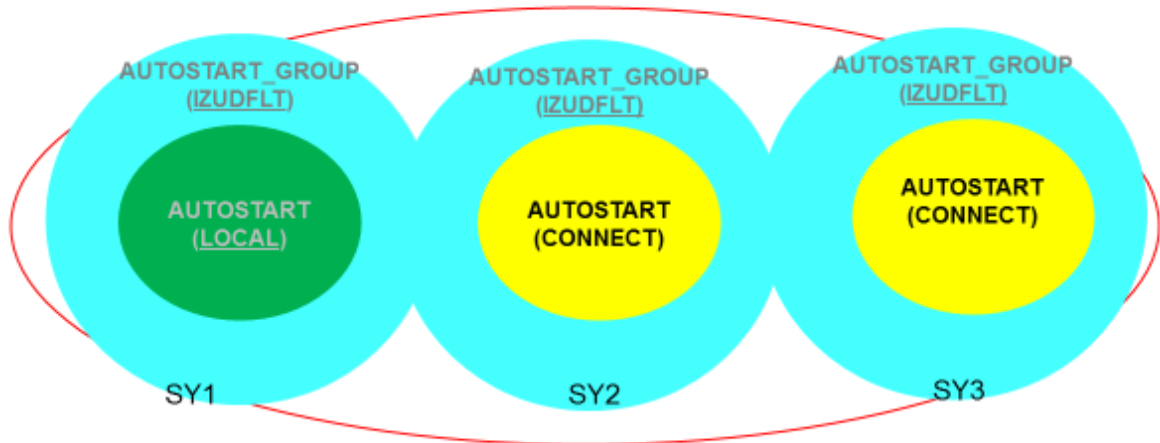


z/OSMF server started in group IZUDFLT: all defaults taken.

z/OSMF Migration Actions for z/OS V2R3



Autostart Use Case #2: Sysplex



SY1: z/OSMF server started in group IZUDFLT: all defaults taken.
 SY2: z/OSMF server **not started**: connecting to default group IZUDFLT.
 SY3: z/OSMF server **not started**: connecting to default group IZUDFLT.

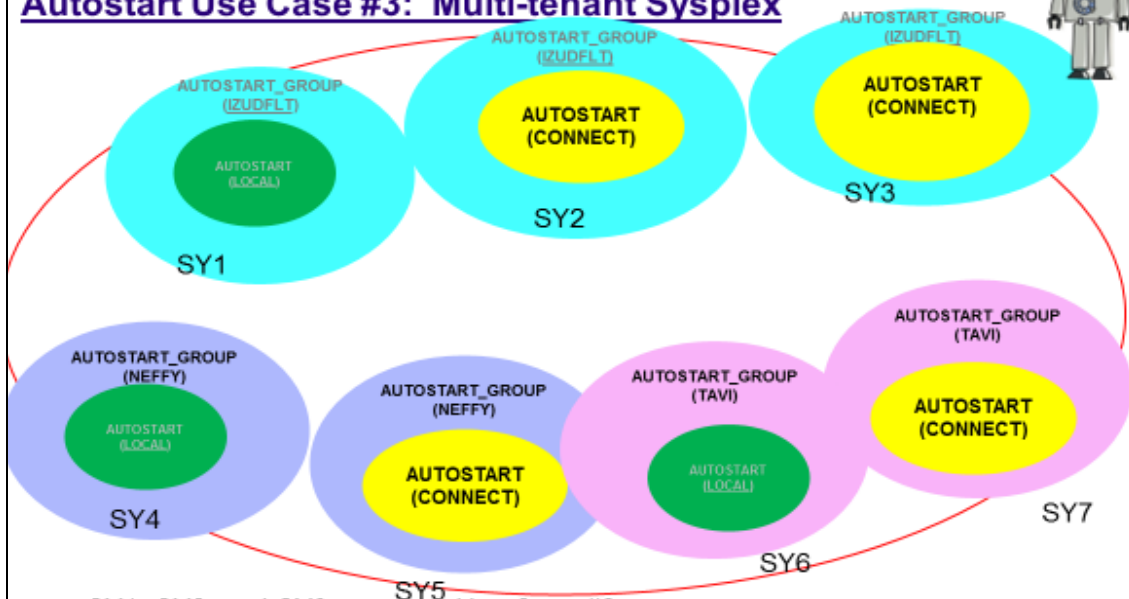
21

© 2017 IBM Corporation

z/OSMF Migration Actions for z/OS V2R3



Autostart Use Case #3: Multi-tenant Sysplex




SY1, SY2, and SY3: same as Use Case #2.
 SY4: z/OSMF server started by default, **group name is NEFFY**.
 SY5: z/OSMF server **not started**: **connecting to group NEFFY**.
 SY6: z/OSMF server started by default, **group name is TAVI**.
 SY7: z/OSMF server **not started**: **connecting to group TAVI**.

22

© 2017 IBM Corporation

z/OSMF Migration Actions for z/OS V2R3



Assuming you took the **Sage advice** and are running z/OSMF on V2R2 or V2R1 today...

Migration Actions you MUST do Before First IPL:

- **Perform security customization (Required in V2R3)**
 - A handful of additional security profiles to add: use V2R3 IZUSEC sample.
 - All the security customization done prior to z/OS V2R3 is still usable and appropriate for z/OS V2R3.
- **Remove any commands or automation that started z/OSMF (Required-IF in V2R3)**
 - ...if you choose to let z/OSMF start at the IPL itself.
 - Otherwise, your start commands will fail.

Failure to follow sage advice, means that you need to completely configure z/OSMF before the IPL of z/OS V2R3 if you want it to start. Much more work!

23

© 2017 IBM Corporation

z/OSMF Actions for z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

[z/OSMF Migration Actions You Can Do Now](#)



Prepare for z/OSMF autostart (Required, as of V2R3)

In z/OS V2R3, the base element z/OSMF is started by default at system IPL. This enhancement, which is referred to as *z/OSMF autostart*, means that z/OSMF is available for use as soon as the system is up. For new z/OSMF installations, this change means that z/OSMF is active by default. The element is available for your use after you enable the required security authorizations. For existing z/OSMF installations, this change means that it is no longer necessary for you to explicitly start the z/OSMF server after each system IPL, whether through automation or operator commands entered manually at the console.

The autostart function introduces a set of migration actions that affect all z/OS installations. The migration actions are different, depending on whether your installation is new to z/OSMF or is already using this element. Installations that currently use z/OSMF meet most of the requirements already, and thus have fewer changes to make.

It is recommended that you perform the migration actions in three phases, as follows:

1. **Planning.** To prepare for z/OSMF autostart, plan for how to deploy z/OSMF in your sysplex. For a sysplex, determine how many systems will run a z/OSMF server. Generally, it is sufficient to have one z/OSMF server active in a sysplex or monoplex, but you might choose to have more. A number of multi-system scenarios are supported. Deploying z/OSMF in a sysplex involves making updates to system libraries and your security management product, such as RACF.
2. **Security product updates.** Your external security manager, such as RACF, requires a significant number of updates. The changes are needed to protect the resources that are used by z/OSMF, and to grant users access to the z/OSMF core functions.
3. **System library updates.** To accommodate z/OSMF autostart, you might need to modify settings in proclib and parmlib. The changes are needed to establish an autostart group for z/OSMF operations in your sysplex.

Migrating to z/OS V2.3: Technical Migration Actions

If you choose not to have z/OSMF started automatically during IPL, you must explicitly disable the autostart function. If you disable z/OSMF autostart, none of the z/OSMF capabilities is available on your system until you start a z/OSMF server manually.

If this migration action is not followed, access failure messages result when the system attempts to start a z/OSMF server automatically. The messages describe required but missing SAF authorizations.

Migration action:

To prepare for z/OSMF autostart, follow the steps that are described in this section. For ServerPac users, the ServerPac configuration process performs the migration actions that are needed to establish a base z/OSMF configuration on the target system. You can use the jobs and documentation in your ServerPac order to create an initial instance of z/OSMF on your z/OS V2R3 system.

Installations that install z/OSMF from a Custom-Built Product Delivery Option (CBPDO) software delivery package, or from a ServerPac order by using the software upgrade method of installation, must perform the steps that are described in this section. The migration actions are different, depending on whether your installation currently uses z/OSMF.

For a new z/OSMF installation:

- Understand the functions and usage requirements of z/OSMF. See the IBM website z/OS Management Facility home page (www.ibm.com/systems/z/os/zos/features/zosmf).
- Ensure that the target system satisfies the minimum memory requirements. The z/OSMF server requires a minimum of 4 GB of system memory to be configured.
- To make the best use of the z/OSMF autostart capability, plan to deploy one or more autostarted z/OSMF servers in your environment. Generally, having one z/OSMF server active in a sysplex or monoplex is sufficient, but you might choose to have more, based on your workload requirements. The goal is to ensure that at least one z/OSMF server is always active in your environment. For a monoplex, little or no planning is needed. The z/OSMF server is started when you IPL the system. Or, if you do not want to use the autostart capability, you can disable it.
 - For a sysplex, more planning is required. You can choose to have one z/OSMF server autostart on a particular system and be used by the other systems in the sysplex. Or, you can select a subset of systems, or several subsets, and associate each subset with a specific z/OSMF server within an autostart group.
 - The set of systems that is served by an autostarted server is known as the *autostart group*. z/OSMF includes one autostart group by default. To have more z/OSMF servers autostarted in a sysplex, you must associate each server—and the systems it serves—with a unique autostart group name.
 - In your planning, you must decide:
 - What are the autostart groups in your sysplex.
 - Which systems autostart a z/OSMF server.
 - Which systems share the use of the autostarted server. These systems must be defined to the same autostart group as the system on which the autostarted server is running.
 - Only one z/OSMF server can be active per autostart group in the sysplex. An autostarted z/OSMF server holds an enqueue on the z/OSMF user directory file system, and handles the z/OSMF requests from other systems that are connected to the same autostart group. Based on your planning, you can enable the desired number of z/OSMF autostart groups for your sysplex.
 - Before installing z/OS V2R3, see the following publication for planning and setup considerations: *IBM z/OS Management Facility Configuration Guide*.
- After you install z/OS V2R3, but before you IPL the system, complete the following migration actions:
 - Update your security management product, as described in “Create security authorizations for z/OSMF”. This action can be done before you install z/OS V2R3, if you add the authorizations to your security database manually, rather than by running the IBM-supplied job in the z/OSMF V2R3 level of SYS1.SAMPLIB (IZUSEC).
 - Update system libraries, as described in “Define one or more z/OSMF autostart groups”.

For an existing z/OSMF installation:

- Review your plans for updating system libraries.
- Review your plans for updating the security management product with your security administrator.
- After you install z/OS V2R3, but before you IPL the system, you must complete the following migration actions:
 - “Create security authorizations for z/OSMF”. This action can be done before you install z/OS V2R3, if you add the authorizations to your security database manually, rather than by running the IBM-supplied job in SYS1.SAMPLIB (IZUSEC).
 - “Define one or more z/OSMF autostart groups”. **If you prefer not to have z/OSMF started automatically during IPL:** You can disable the autostart function. If you disable autostart, be aware that the JES2 notification service in z/OS V2R3 does not operate with full function until you enable

Migrating to z/OS V2.3: Technical Migration Actions

the z/OSMF autostart function, or connect to a valid autostart group. For instructions on disabling the autostart function, see “Define one or more z/OSMF autostart groups”.

[z/OSMF actions to perform before the first IPL of z/OS V2R3](#)



Update z/OSMF for the new minimum Java requirement (Required, as of V2R3)

In z/OS V2R3, the z/OSMF server requires the following level of Java: IBM 64-bit SDK for z/OS, Java Technology Edition, V8 (5655-DGH). If your IZUPRMxx parmlib member specifies an earlier level of Java, you must update the JAVA_HOME statement in IZUPRMxx.

Migration action:

Install IBM 64-bit SDK for z/OS, Java Technology Edition, V8 (5655-DGH).

For a new z/OSMF installation:

- By default, the SDK resides in the directory /usr/lpp/java/J8.0_64 on your system. If you install it in another location, be sure to specify the location on the JAVA_HOME statement in your IZUPRMxx parmlib member.

For an existing z/OSMF installation:

- If the JAVA_HOME statement in member IZUPRMxx specifies an earlier version of Java, update the JAVA_HOME statement to refer to the directory /usr/lpp/java/J8.0_64 on your system.

If you installed Java V8 in the default Java directory, you do not need to specify the JAVA_HOME statement in IZUPRMxx. If JAVA_HOME is not specified, the z/OSMF server searches for Java files in the directory /usr/lpp/java/J8.0_64.



Create security authorizations for z/OSMF (Required, as of V2R3)

To accommodate z/OSMF autostart, you must create the appropriate authorizations in your security management product, such as RACF. The changes are needed to protect the resources that are used by z/OSMF, and to grant users access to the z/OSMF core functions.

The RACF requirements are listed in "Appendix A" of *IBM z/OS Management Facility Configuration Guide*. Sample RACF commands for setting up security for z/OSMF are supplied by IBM in the z/OSMF V2R3 level of SYS1.SAMPLIB(IZUSEC).

A new z/OSMF installation must add all of these authorizations to its security management product. An existing z/OSMF installation has most of these authorizations already, and thus has fewer changes to make.

Migration action:

To prepare for z/OSMF autostart, follow the steps that are described in this section.

For ServerPac full system replacement users, the ServerPac installation process creates the required z/OSMF security authorizations in RACF for you. You can use the ServerPac supplied jobs and documentation as a model for another external security product if you desire.

Installations that install z/OSMF from a Custom-Built Product Delivery Option (CBPDO) software delivery package, or from a ServerPac order by using the software upgrade method of installation, must perform the steps that are described in this section. The migration actions are different, depending on whether your installation currently uses z/OSMF.

For a new z/OSMF installation:

- After you install z/OS V2R3, but before you IPL the system, you must create the SAF authorizations that are described in "Appendix A" of *IBM z/OS Management Facility Configuration Guide*. This action can be done before you install z/OS V2R3, if you add the authorizations to your security database manually, rather than by running the IBM-supplied job in the z/OSMF V2R3 level of SYS1.SAMPLIB (IZUSEC).

For an existing z/OSMF installation:

- After you install z/OS V2R3, but before you IPL the system, you must create the SAF authorizations that are described below. RACF sample commands for the SAF authorizations in this table are provided in the z/OS V2R3 level of SYS1.SAMPLIB(IZUSEC).

Migrating to z/OS V2.3: Technical Migration Actions

Resource class	Resource name	Who needs access?	Type of access required	Why
FACILITY	BPX.WLMSEVER	z/OSMF server user ID (IZUSVR1, by default).	READ	Allows the z/OSMF server to use WLM functions to create and manage work requests.
SERVAUTH	CEA.SIGNAL.ENF83	z/OSMF server user ID (IZUSVR1, by default).	READ	Allows the z/OSMF server to use signal code ENF83 to indicate its status to other systems in the sysplex.
SERVAUTH	EZB.INITSTACK. <i>sysname.tcpname</i>	z/OSMF server user ID (IZUSVR1, by default).	READ	Allows the z/OSMF server to access the TCP/IP stack during TCP/IP initialization. This authorization is needed if the TCP/IP profile activates Application Transparent Transport Layer Security (AT-TLS).
SERVER	BBG.ANGEL.IZUANG1	z/OSMF server user ID (IZUSVR1, by default).	READ	Allows the z/OSMF server to use z/OS authorized services.
STARTED	IZUINSTP.IZUINSTP	z/OSMF administrator group (IZUADMIN)	READ	Defines the started task for the z/OSMF dependent address space, which is used to determine whether z/OS UNIX and TCP/IP are available. The job name must be IZUINSTP. Otherwise, the z/OSMF dependent address space is not initialized during z/OSMF autostart processing.
ZMFAPLA	<saf-prefix>.ZOSMF. VARIABLES.SYSTEM. ADMIN	z/OSMF administrator group (IZUADMIN)	READ	Allows the user to access the system variables in the Systems task.



Remove commands or code that start the z/OSMF server (Required-IF, as of V2R3)

Required if you autostart z/OSMF.

In z/OS V2R3, the z/OSMF server is started by default at system IPL. This enhancement, which is referred to as *z/OSMF autostart*, means that z/OSMF is available for use as soon as the system is up. For existing users of z/OSMF, the z/OSMF autostart capability means that it is no longer necessary for your installation to explicitly start the z/OSMF server after each system IPL, whether through automation or commands entered manually at the operations console. As part of your migration to z/OS V2R3, remove any methods you use to automate the start of the z/OSMF server. Otherwise, error messages result if **START** commands are issued against already-started z/OSMF servers.

Migration action:

If your installation does not start z/OSMF automatically, you have no migration action to take. Otherwise, you must review and, if necessary, modify or remove any methods that you currently use for starting the z/OSMF server.

For example:

- Ensure that no **START** commands are issued for the z/OSMF started procedures in the COMMNDxx parmlib member.
- Ensure that the z/OSMF started procedure names are not listed in the AUTOLOG statement in the TCP/IP profile (PROFILE.TCPIP). By default, the procedures are named IZUANG1 and IZUSVR1.
- Verify that your automation products do not start z/OSMF.

If you prefer not to have z/OSMF started automatically during IPL, you must explicitly disable the autostart function.



Define one or more z/OSMF autostart groups (Required, as of V2R3)

Based on the planning you did in "Prepare for z/OSMF autostart", you can create the desired number of z/OSMF autostart groups in your sysplex. This work involves modifying settings in the system libraries, parmlib and proclib. For any systems for which you do not want to have z/OSMF started automatically, you must explicitly disable the autostart function.

Migration action: For an overview of the autostart capability and autostart groups, see *IBM z/OS Management Facility Configuration Guide*.

If one autostart group is sufficient for your sysplex, it is recommended that you allow your systems to use the default autostart group (IZUDFLT). Otherwise, you can define more autostart groups, as needed for your environment. Doing so involves creating one or more IZUPRMxx parmlib members, setting the system parameter IZU=, and customizing the IZUSVR1 started procedure.

For ServerPac installers, if you select the ServerPac full system replacement installation type, z/OSMF is configured through a ServerPac post-installation job, using IBM defaults. The default configuration includes the parmlib and proclib member setup that is described in the steps that follow. Therefore, if you use the parmlib and proclib members from ServerPac, you do not need to perform the following steps. However, you might want to review the ServerPac provided members to ensure that they contain suitable values for your installation, or modify them, as you require.

To perform the parmlib updates, follow these steps:

1. Create an IZUPRMxx parmlib member with the following statements specified as required for your environment:

AUTOSTART(LOCAL|CONNECT)

Specifies the capability for autostarting the z/OSMF server on this system.

- AUTOSTART(LOCAL) indicates that the system can autostart a z/OSMF server.
- AUTOSTART(CONNECT) indicates that the system cannot autostart a z/OSMF server. Instead, the system uses the z/OSMF server on another system in the same autostart group.

For a monoplex, specify AUTOSTART(LOCAL). By default, AUTOSTART is set to LOCAL.

AUTOSTART_GROUP(IZUDFLT|'nnnnnnnn')

Assigns a name to the autostart group. z/OSMF includes one autostart group name by default (called IZUDFLT). To associate a group of systems with a different autostart group, ensure that the IZUPRMxx member for each system specifies the same value for AUTOSTART_GROUP.

The name can consist of 1-32 alphanumeric characters (A-Z, a-z, 0-9) or special characters (#, \$, or @). Alphabetic characters are case insensitive. By default, AUTOSTART_GROUP is set to IZUDFLT.

SERVER_PROC(proc-name)

Specifies the started procedure that is used to start the z/OSMF server on this system. It is recommended that you use the default started procedure, IZUSVR1, which should be adequate for most z/OS installations. If you specify an alternative procedure name, such as IZUSVR2, ensure that the z/OSMF user and z/OSMF administrator security groups are authorized to the started procedure name.

The name can consist of 1 to eight alphanumeric characters (A-Z, a-z, 0-9) or special characters (#, \$, or @). This value is not case-sensitive; lowercase alphabetic characters are folded to uppercase. By default, SERVER_PROC is set to IZUSVR1.

ANGEL_PROC(proc-name)

Specifies the started procedure that is used internally to start the z/OSMF angel process on this system. It is recommended that you use the default procedure, IZUANG1, which should be adequate for most z/OS installations. If you specify an alternative procedure name, ensure that the z/OSMF user and z/OSMF administrator security groups are authorized to the started procedure name.

The name can consist of 1 to eight alphanumeric characters (A-Z, a-z, 0-9) or special characters (#, \$, or @). This value is not case-sensitive; lowercase alphabetic characters are folded to uppercase. By default, ANGEL_PROC is set to IZUANG1.

2. Ensure that the IZU= specification is coded to include the suffixes of the appropriate IZUPRMxx members for each system in your sysplex. You can specify up to 38 member suffixes on the IZU= parameter in your IEASYSxx member.

To perform the proclib updates, follow these steps:

1. Ensure that you are using the z/OS V2R3 level of the z/OSMF started procedures from your installed PROCLIB data set, as they are different from prior releases. The started procedures are IZUANG1, IZUSVR1, and IZUINSTP. The IZUINSTP procedure is new in z/OS V2R3; it is used by the z/OSMF server for communicating with z/OS components. Install IZUINSTP, but do not modify it. Place the started procedures (IZUANG1, IZUSVR1, and IZUINSTP) in a data set that is in the IEFPSI concatenation that is used by the system to find started procedures before the primary subsystem (JES) initializes. It is recommended that this data set is the same data set that is used by JES to find started procedures after JES initializes. For existing z/OSMF installations, if you have older levels of the started

Migrating to z/OS V2.3: Technical Migration Actions

procedures IZUANG1 and IZUSVR1, you must remove them. Otherwise, the z/OSMF server might not start on a z/OS V2R3 system.

Note: Another started procedure, IZUSVR2, is provided in SYS1.SAMPLIB. If you choose not to autostart the z/OSMF server, the IZUSVR2 procedure can be used for starting the z/OSMF server manually.

2. Modify the started procedure for the z/OSMF server, IZUSVR1, to control its start-up behavior, as follows:
 - a. On the parameter SERVER, specify either AUTOSTART or STANDALONE, as follows:
 - i. Specify AUTOSTART to have the server started automatically.
 - ii. Specify STANDALONE if you want to start the server manually by using the START command.
 - b. If you specify SERVER=AUTOSTART, you can specify one of the following values on the parameter IZUPRM:
 - i. **PREV** Use the IZUPRMxx suffixes, if any were used by the previous instance of z/OSMF within the current IPL. IZUPRM='PREV' is used as the default in the standard IZUSVR1 procedure. IZUPRM='PREV' behaves like IZUPRM='SYSPARM' when the system encounters it during the initial IPL time (the first use of the IZUSVR1 procedure) because there is no previous instance of z/OSMF to use. This setting is not valid if the SERVER parameter is set to STANDALONE.
 - ii. **SYSPARM** Use the IZUPRMxx suffixes that are specified on the IZU system parameter in IEASYSxx. This setting is not valid if the SERVER parameter is set to STANDALONE.
 - iii. **NONE** To indicate that no IZUPRMxx members are read at server start-up.
 - iv. **xx|(xx, ..., zz)** Specify the specific suffixes for the IZUPRMxx parmlib member or members that you want the procedure to use. If you specify a suffix, the member must exist in your parmlib concatenation. Otherwise, the procedure cannot be started. Multiple suffixes must be enclosed in parentheses.

The following syntax forms are valid:

IZUPRM=(xx,yy,...)

IZUPRM=xx

IZUPRM=NONE

This setting is not valid if the SERVER parameter is set to STANDALONE. **Note:** The IZUPRMxx suffixes you specify, explicitly or implicitly, in the IZUPRM parameter of the procedure override any suffixes that are defined in the IZU system parameter in IEASYSxx.

If you prefer not to have z/OSMF started automatically during IPL: You can disable the autostart function for one or more z/OS systems, as follows:

- To prevent a z/OS system from autostarting the z/OSMF server, ensure that the system uses a IZUPRMxx member that specifies AUTOSTART(CONNECT). This setting causes the system to connect to the autostart group that is specified on the AUTOSTART_GROUP statement, rather than autostarting its own server.
- To prevent a z/OS system from connecting to an autostart group, specify the name of a group on the AUTOSTART_GROUP parameter that is not used by any autostart server in the sysplex. For example, AUTOSTART_GROUP('NEVERUSED').
- Similarly, for each system for which you want to disable z/OSMF autostart, ensure that the AUTOSTART(CONNECT) and AUTOSTART_GROUP('NEVERUSED') settings are in effect.
- In your IZU= specifications, verify that the IZU= parameter identifies the suffixes of the IZUPRMxx members that contain the desired settings.

If you disable autostart on a z/OS system, be aware that the JES2 notification service in z/OS V2R3 does not operate on that system with full function until you enable the z/OSMF autostart function, or allow the system to connect to a valid autostart group.

Note: Changing an AUTOSTART_GROUP name requires an IPL. You cannot change this option with a stop and restart of the z/OSMF server.

[z/OSMF actions to perform after the first IPL of z/OS V2R3](#)



Configure the z/OSMF optional plug-ins (Recommended)

In z/OSMF, a *plug-in* is a collection of one or more system management tasks that add function to z/OSMF. When you configure a plug-in, you make its tasks available to users in the z/OSMF navigation area. If your installation does not already use any of the z/OSMF optional plug-ins, it is recommended (not required) that you enable one or more of the following plug-ins:

Migrating to z/OS V2.3: Technical Migration Actions

- Capacity Provisioning
- Configuration Assistant for z/OS Communications Server
- Incident Log
- ISPF
- Resource Monitoring
- Software Management
- Sysplex Management
- Workload Management.

By default, z/OSMF does not include any of the optional plug-ins.



Migration action:

If your installation does not already use the z/OSMF optional plug-ins, choose one or more of the plug-ins to enable. Your decision on which plug-ins to enable depends in part on your installation's readiness to perform the various z/OS system customizations that are associated with each plug-in. When you are planning for the plug-ins, review the system setup requirements for each plug-in, as described in *IBM z/OS Management Facility Configuration Guide*.

SDSF Migration Actions for z/OS V2R3



Migration Actions Before Installing:

- **Enable the SDSF class for RACLIST processing (Req-IF, as of V2R3)**
 - If you activate the RACF SDSF class, you must also enable it RACLIST processing.
 -  • Newer SDSF functions will fail the request, although you can change this behavior with AUXSAF (NOFAILRC4) in ISFPRMxx.
- **Start SDSF and SDSFAUX address spaces at IPL (Req-IF, as of V2R3)**
 -  • For full and prior SDSF functionality, both address spaces must be started.
 - Although in V2R3 you can still have limited functionality without both address spaces (when implementing various detailed techniques), this might not be the case in future releases.

Migration Actions After First IPL:

- **Accommodate new SDSF main panel default (Req-IF, as of V2R3)**
 - Main panel is now a scrollable table.
 - Compatibility: SDSF custom property `Panel.Main.DisableTable` can be changed to true, or user can use a DDNAME ISFMIGMN.

24

© 2017 IBM Corporation

SDSF Actions for z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

SDSF Migration Actions You Can Do Now



Enable the SDSF class for RACLIST processing (Required-IF, as of V2R3)

Required if you have the RACF SDSF class active and have not yet enabled the SDSF class for RACLIST processing.

Beginning in z/OS V2.3, if you activate the RACF SDSF class, you must also enable this class for RACLIST processing. This change is needed to allow the command **RACROUTE REQUEST=FASTAUTH** to be used with the SDSF class.

If you already have the RACF SDSF class active, performing RACLIST processing on the SDSF class causes the class to be RACLISTed on the other systems that share the RACF database. If you do not enable the SDSF class for RACLIST processing:

- For older SDSF functions, return code 4 is returned for **VERIFY** requests, as is done when the SDSF class is not active. Instead, SDSF uses the non-SAF-based ISFPARMS security.
- For new SDSF functions, the default is to fail the request; the functions are not authorized.

You can change this behavior by specifying AUXSAF(NOFAILRC4) in the ISFPRMxx member, which results in all requests being allowed.

Migration action: Follow these steps:

1. If your RACF SDSF class is active, enable the SDSF class for RACLIST processing. To do so, enter the command **SETROPTS RACLIST(SDSF)**.
2. Whenever SDSF class profiles are changed, you must refresh the class. To do so, enter the command **SETROPTS RACLIST(SDSF) REFRESH**.



Verify that the SDSF address spaces (SDSF server and SDSFAUX) are started at IPL

(Required-IF, as of V2R3)

Required if you do not already start the SDSF and SDSFAUX address spaces.

As of z/OS V2R3, SDSF requires the SDSF and SDSFAUX address spaces to be active for full functionality. The SDSF address space manages connections, processes ISFPRMxx statements, handles operator commands, and starts and stops SDSFAUX. The SDSFAUX address space is used for data gathering requests.

Usually, the SDSF address space is started during IPL using COMMNDxx. During SDSF initialization, the SDSFAUX address space is started.

When a user accesses SDSF, the SDSF client program attempts to connect to the SDSF address space (also referred to as the SDSF server). To connect to the SDSF server, the user must have READ access to the ISF.CONNECT.*system* resource in the SDSF class.

If the SDSF address space is not active, SDSF provides limited functionality. The user must have READ access to the SERVER.NOPARM resource in the SDSF class so that ISFPARMS can be used instead of ISFPRMxx. Panels that require the use of the SDSFAUX data gatherers (such as APF, LPA, and LNK) are not available.

If the SDSF address is active, but no ISFPRMxx is in effect (such as a syntax error during startup), SDSFAUX is not started. The user requires access to the SERVER.NOPARM resource to fall back to ISFPARMS and requires READ access to the ISF.CONNECT.*system* resource to continue. Panels that require the use of SDSFAUX are not available.

If the SDSF address space is active, but the SDSF class is not active or not RACLISTed, the SDSF server allows requests based on the ISFPRMxx CONNECT definition. When AUXSAF(FAILRC4) is in effect (the default), the request is denied. The user cannot connect to the SDSF server and the SDSFAUX related panels are not available. SDSF falls back to ISFPARMS because access to the SERVER.NOPARM resource results in a return code 04 (indeterminate result).

When AUXSAF(FAILRC4) is in effect, the server allows the request, but access to the panel is controlled through the definitions in ISFPARMS.

IBM recommends that you start the SDSF server. Although V2R3 provides limited functionality when the server is not active, this might not be the case in subsequent releases.

Migration action: Determine whether the SDSF and SDSFAUX address spaces are started during IPL by doing either of the following:

- From SDSF, enter the **WHO** command. Verify that the response contains the SERVER=YES keyword.
- Enter the command **F SDSF,D** to verify that the SDSF address space is active.

If your installation already starts the SDSF server and SDSFAUX address spaces, no action is necessary.

Otherwise, if the **WHO** response is SERVER=NO or the **MODIFY** command results in job not found, the server address space must be started.

SDSF actions to be performed before the first IPL

Remove the entry for ISFHCTL from SCHEDxx (Required-IF, as of V2R3)

Required if you specify ISFHCTL in a SCHEDxx parmlib member.

In z/OS V2R3, the SDSF server program, ISFHCTL, is changed to run in program protect key 4. In previous releases, this program ran in program protect key 1. Your installation might currently specify a key for ISFHCTL by using parmlib member SCHEDxx to update the program properties table (PPT). If so, it is recommended that you remove the entry from SCHEDxx. Because ISFHCTL is defined in the PPT in all levels of z/OS, you no longer need to define this program in member SCHEDxx.

During system initialization, SDSF verifies that it is running in the correct key. If the key is incorrect, SDSF initialization fails with the following message:

```
ISF517E SDSF SERVER WAS NOT STARTED DUE TO INVALID  
EXECUTION ENVIRONMENT, POSSIBLE MISSING PPT ENTRY.
```

Migration action:

If your installation does not use parmlib member SCHEDxx, no action is necessary. Otherwise, check SCHEDxx for the PPT entry for program ISFHCTL. If SCHEDxx contains an entry for ISFHCTL, remove the entry.

SDSF actions to be performed after the first IPL

Modify programs that post-process SDSF panels, for new main panel (Required-IF, as of V2R3)

Required if your installation uses programs that rely on the older SDSF main panel format.

Migrating to z/OS V2.3: Technical Migration Actions

As of z/OS V2R3, the main panel of SDSF is restructured to use a scrollable table. This change allows new commands to be added, regardless of the screen depth. Entries in the table can be located, sorted, and filtered to help with selecting commands.

As part of this change, the title line of the main panel is changed. In previous releases, the main panel title contained the string "SDSF PRIMARY OPTION MENU." In z/OS V2R3, the title line contains the string "SDSF MENU" in the upper left corner.

A compatibility mode is provided, which causes SDSF to use the older format. This mode can be enabled, either by using a custom property or a special DDNAME allocated to the user's session. A new SDSF custom property Panel.Main.DisableTable is implemented. When set to false (the default), the SDSF main panel is rendered as a table.

When the special ddname ISFMIGMN is allocated (typically to a dummy data set) or the Panel.Main.DisableTable custom property is set to true, the panel is rendered in the older style two-column layout. However, only the options that fit within the older screen depth are shown.

Note: All SDSF options are available, even if not visible due to insufficient screen depth.

If your installation has programs that post-process SDSF screen output and the programs rely on the SDSF main menu title "SDSF PRIMARY OPTION MENU", you must modify your programs to check for the new SDSF main menu title line.

Migration action: If you do not have scripts that post-process SDSF screen output, no action is necessary. Otherwise, modify your scripts to check for the new SDSF main menu title line.

Tip: If you use SDSF batch or AFD scripts, convert them to SDSF/REXX, which is not sensitive to SDSF screen layouts and is thus independent of changes to the panel formats. If it is not practical to change your scripts, use the Panel.Main.DisableTable custom property or allocate special DDNAME ISFMIGMN to revert to the old main panel format.

RACF Migration Actions for z/OS V2R3



Migration Actions Before First IPL:

- **RACF only supplies three certificate-authority (CA) certificates out of the current list (Req-IF, as of V2R3)**
 - Prior to V2.3, RACF supplied 26 CA certificates.
 - At V2.3, RACF will only keep the following three CA certifications that are used by IBM components:
 1. STG Code Signing Certificate Authority certificate– issued by IBM for RACF program signature verification for V2R1 and earlier,
 2. STG Code Signing Certificate Authority – G2 certificate – issued by IBM for RACF program signature verification for V2R2 and later,
 3. GeoTrust Global Certificate Authority certificate – for the SMP/E process
 - You no longer need to remove the unused CA certificates each z/OS release.
 - If you had used the removed 23 CA certificates, they are still there and RACF will not delete them in z/OS V2.3.

25

© 2017 IBM Corporation

RACF Migration Actions For z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

RACF Migration Action Pre-First-IPL

Accommodate the CA certificates that RACF supplies (Required-IF, as of V2R3)

Required if you need additional CA certificates beyond those that RACF supplies.

Prior to z/OS V2.3 RACF supplied 26 certificate-authority(CA) certificates. Most of them are well-known CA's. Two of them are IBM code signing CA's. The rationale is to simplify the SSL/TLS set up for those who happen to trust the CA's on the list. You only need to alter the status of the certificate to TRUST since all these certificates are shipped with NOTRUST status.

But there are those who do not want to have these CA certificates in their system even they have the NOTRUST status. It might be that someone may accidentally alter the TRUST status. Even they delete them, the default list will be added again every time when the system is IPLed.

From the convenience point of view, it is difficult to maintain a list of well known CA's as the number of CA agencies is growing, and there are multiple certificates under each agency. It is hard to know which one will be commonly used in general.

In V2R3, RACF will only keep the following three CA certificates that are used by IBM components:

- 1) STG Code Signing Certificate Authority certificate– issued by IBM for RACF program signature verification for V2R1 and earlier
- 2) STG Code Signing Certificate Authority – G2 certificate – issued by IBM for RACF program signature verification for V2R2 and later
- 3) GeoTrust Global Certificate Authority certificate – for the SMP/E process

As a result of these changes, there may be migration actions when moving to z/OS V2.3 from a previous release.

Migration Action

Migrating to z/OS V2.3: Technical Migration Actions


You no longer need to remove the added RACF CA certificates anymore, if you didn't use them. See the RACF Administrator book for the 23 CA certificates which were removed.

1. If you do not want the previous shipped CA certificates in the existing RACF database, you do not need to delete them every time after IPL anymore.
2. If you want the previous shipped CA certificates in the existing RACF database, no action is needed as they are still there.
3. If you want the previous shipped CA certificates in a new v2.3 RACF Database, go to those CA websites to download them, or RACDCERT EXPORT them from the RACF DB in the previous release then RACDCERT ADD to the RACF DB in the new release.

z/OS OpenSSH Migration Actions for z/OS V2R3



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.
- Read of all changes in the *z/OS Migration* book or workflow.

26

© 2017 IBM Corporation

z/OS OpenSSH Migration Actions For z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 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 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. The z/OS V2R3 level remains at version 6.4p1.

Migration action : See the tables in the *z/OS 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 V2R3



Migration Actions Pre-First IPL:

- **Add /global directory to the sysplex root (Req-IF, as of V2R3)**
 - /global provides a location where a single copy of config files or mount points for program products can be referenced within a sysplex.
 - Add /global (7,5,5) to your sysplex root. It is already provided for you in the version root.
 -  • Also, remove the obsolete DCE dir called /... if it is empty.

- **Remove files and directories in /var/man (Required-IF, as of V2R2)**
 - As of V2R2, man no longer uses this directory. Instead, it uses /tmp on a per-user basis. Monitor /tmp space.
 -  • Remove /var/man and all subdirectories and files from your V2R2 system.

27 © 2017 IBM Corporation

z/OS UNIX Migration Actions for z/OS V2R3

These migration actions were taken from z/OS V2R3 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 V2R3 Migration.

z/OS UNIX System Services Migration Actions Pre-First IPL

Add the /global directory to the sysplex root file system (Required-IF, as of V2R3)

Required if

In V2R3, a new directory, /global, is available in the version root and sysplex root. It can be used to store and maintain a single copy of configuration files or mount points of program products that can be referenced by all members of the sysplex. Prior to V2R3, individual copies of the same configuration file had to be maintained in each member of the sysplex.

Migration action:

If you are using a sysplex root file system, take the following actions:

1. Rerun the sample REXX EXEC SYS1.SAMPLIB(BPXISYS1) to update the sysplex root, making system installation-specific adjustments. Running the EXEC will create the /global directory in the sysplex root. You do not need to run any sample jobs to create /global in the version root (or root) for a non-sysplex environment because ServerPac will provide that directory during installation.
 - a. Alternatively, a system programmer can run the MKDIR command to add the /global directory with permission bits 7,5,5 to an existing sysplex root.
 - b. At this time, you can also remove an obsolete directory called /..., after verifying that it is an empty directory. This directory was used by DCE, which is no longer shipped in z/OS.
 - c. If applicable, ensure that the same updates are also made to the alternate sysplex root.
2. Create a file system that will be mounted on the /global directory.
3. As necessary and as instructed, create additional mount points (and file systems) for exploitation by z/OS functions under /global.



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.

Knowledge Center for z/OS Migration Action

Migration Actions After First IPL:

- Use /global by default **(Required-IF, as of V2R3)**
 - KC4z uses /global for product documentation content and associated properties files. If sysplex, all member can access it.
 - Move existing configuration from /sharedapps to /global using provided scripts and JCL.

High Level Assembler Migration Action

Migration Actions Before First IPL:

- Discontinue usermod for IEV90 alias **(Req-IF, as of PI71827)**
 - IEV90 is an alias of ASMA90 within the product itself.
 - No need to have a usermod to supply that alias yourself.

XL C/C++ Migration Action

Migration Actions Before First IPL:

- Accommodate ARCH and TUNE default changes **(Req-IF, as of V2R3)**
 - ARCH (8)->ARCH(10). For instructions on zEC12/zBC12.
 - TUNE(8)->TUNE(10). Exec all models, optimiz for zEC12/zBC12.

28 © 2017 IBM Corporation

Knowledge Center for z/OS Migration Action For z/OS V2R3

Use /global by default (Required-IF, as of V2R3)

Required if you wish you take advantage of the /global directory.

In z/OS V2R3, IBM Knowledge Center for z/OS (KC4z) uses the new /global directory for KC product documentation content and associated properties files. This directory, which is available in the version root and sysplex root, can be used to store and maintain a single copy of KC4z data (KC product documentation content and associated properties files) that can be referenced by all members of the sysplex. Before V2R3, individual copies of the same KC4z data had to be maintained in each member of the sysplex.

With this change, the directory structure that is created by the KC4z 1.1 post-installation configuration scripts is changed to use the /global directory. As a result, a single copy of KC product content and associated properties files can be shared across all systems in a sysplex that use shared file system support. The entire post-installation directory structure is changed, so migration considerations, when applicable, apply to both sysplex and single system environments.

Migration action: Follow the migration steps in *z/OS V2.3 Migration*.

High Level Assembler for z/OS Migration Action For z/OS V2R3

Discontinue the usermod for the IEV90 alias (Required-IF, as of PI71827 on all releases)

Required if you previously used a usermod to define the IEV90 alias.

In z/OS V2R3, and earlier releases with APAR PI71827 applied, you no longer need to define IEV90 as an alias of ASMA90. HLASM now includes the ALIAS IEV90 statement in the JCLIN for HMQ4160J, which means that an alias for ASMA90 is automatically created. Programs that invoke ASMA90 through the name IEV90 will continue to work.

Migratino action: Review your installation procedures for instructions for applying the ASMAIEV sample usermod. You no longer have to apply the usermod.

XL C/C++ Migration Action For z/OS V2R3

Accommodate the changes to default ARCH and TUNE level (Required-IF, as of V2R3)

Migrating to z/OS V2.3: Technical Migration Actions

Required if the change in default affects you.

Starting with z/OS V2R3, the default ARCH level is changed from ARCH(8) to ARCH(10), and the default TUNE level is changed from TUNE(8) to TUNE(10).


ARCH(10) produces code that uses instructions available on the 2827-xxx (IBM zEnterprise EC12 (zEC12)) and 2828-xxx (IBM zEnterprise BC12 (zBC12)) models in z/Architecture mode. Specifically, these ARCH(10) machines and their follow-ons add instructions that are supported by the execution-hint facility, the load-and-trap facility, the miscellaneous-instruction-extension facility, and the transactional-execution facility. More information about these facilities is provided in *z/Architecture Principles of Operation, SA22-7832*.

TUNE(10) generates code that is executable on all models, but is optimized for the 2827-xxx (IBM zEnterprise EC12 (zEC12)) and 2828-xxx (IBM zEnterprise BC12 (zBC12)) models.

Migration action:

Users can still specify lower arch levels, as before. Only the default is changing. Using the TARGET option to target any OS level before V2R3 changes the default architecture level to the ALS for that OS level. For example, ARCH(7) for V2R1.

Communications Server Migration Actions for z/OS V2R3



Migration Actions Before Installing:

- **IP Services: Verify that the changed `HowToAuthMsgs` and `HowToAuth` defaults are acceptable (Required-IF, as of V2R3, and V2R1 / V2R2, both with APAR PI55022)**
 - These IPsec policy parameters changed:
 - `HowToAuthMsgs` parameter on the `KeyExchangeOffer` changes from MD5 to SHA1
 - `HowToAuth` parameter on the `IpDataOffer` changes from `HMAC_MD5` to `HMAC_SHA1`
 - MD5 is considered weak and is not recommended. Decide whether to upgrade to a more security algorithm.
 - Coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peer's policy is compatible with the default change. If the policy is incompatible with remote's peer policy, IKE daemons will not be able to negotiate IPsec tunnels.
 - Is your IPsec policy is generated by z/OSMF Config Assistant? Review and update.
 - If so: IPsec policies specify a value (and do not omit the parameter).
 - If not: Review your specification, or lack of. Notice if change in default will affect you.
- **IP Services: Verify that the changed `ANONYMOUSLEVEL` default is acceptable (Req-IF, as of V2R3)**
 - The default value for `ANONYMOUSLEVEL` parameter in `FTP.DATA` for the FTP server is changed from 1 to 3.
 - No `ANONYMOUS` statement: then it's not enabled and you are not affected.
 - If you have coded the `ANONYMOUSLEVEL` value, then you will use what you specify.
 - If you relied upon the default `ANONYMOUSLEVEL`, default of 3 enables control of individual filetype. Review the six anonymous filetype configuration settings..
 - Particularly `ANONYMOUSFILETYPEJES` setting: `FALSE` (with level 3) is recommended to prevent anonymous users from submitting jobs.



29

© 2017 IBM Corporation

Communications Server Migration Actions for z/OS V2R3



Migration Actions Before First IPL:

- **IP Services: Verify that the changed `DHGroup` default is acceptable (Required-IF, as of V2R1 and V2R2, both with APAR PI43832; R13 with APAR PI43833)**
 - As of the APAR, the default value for the `DHGroup` parameter on the `KeyExchangeOffer` statement in the IPsec policy is changed from `Group1` to `Group2`.
 - `DHGroup`: Specifies the Diffie-Hellman group used during the phase 1 key exchange.
 - `Group1`: Modular exponentiation group with a 768-bit modulus. (It is considered weak and is not recommended.)
 - `Group2`: Modular exponentiation group with a 1024-bit modulus.
 - z/OSMF Config Assistant-generated IPsec policies specify a value and do not default.
- Coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peer's policy is compatible with the default change.
- **IP Services: Ensure TLS/SSL secure connection in non-FIPS mode meets the minimum per end-entity certificate key size (Required IF, as of V2R3)**
 - System SSL raised minimum asymmetric key size for peer certificates used during the negotiation of a TLS/SSL secure connection in non-FIPS mode:
 - RSA 512 bits -> 1024.
 - DSA 512 bits -> 1024.
 - DH 512 bits -> 1024.
 - ECC 160 bits -> 192.
 - Make changes as appropriate for AT-TLS, FTP server and client, TN3270E Telnet server, Digital Certificate Access Server (DCAS), and Policy Agent Client.



30

© 2017 IBM Corporation

Communications Server Migration Actions for z/OS V2R3



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 ...
    
```

31

© 2017 IBM Corporation

Communications Server Migration Actions for z/OS V2R3

These migration actions were taken from *z/OS V2R3 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 V2R3 Migration*.

Communications Server Migration Actions You Can Do Now

IP Services: EZZ6044I and EZZ6045I descriptor codes are changed (Required-IF, as of V2R3)

Required if your automation is sensitive to descriptor codes of messages.

In z/OS V2R3, the descriptor codes for EZZ6044I and EZZ6045I messages are changed from 4 to 5. These messages are generated when either the Telnet Server is started or the command **VARY OBEY** is issued for the Telnet Server.

Only the descriptor codes are changed. The message text is unchanged from previous releases:

EZZ6044I jobname PROFILE PROCESSING BEGINNING FOR FILE dataset_name

EZZ6045I jobname PROFILE PROCESSING COMPLETE FOR FILE dataset_name

Migration action: If your automation processing is sensitive to message descriptor codes, ensure that the automation is updated for the change in the descriptor code for EZZ6044I and EZZ6045I. These messages are now issued as a command response (descriptor code=5).



IP Services: Verify that the changed HowToAuthMsgs and HowToAuth defaults are acceptable (Required-IF, as of V2R3, and V2R2 and V2R1 with APAR PI55022)

Required if you use an IPSec policy.

In z/OS V2R3, the default values for the following IPSec policy parameters are changed:

- The HowToAuthMsgs parameter on the KeyExchangeOffer statement is changed from MD5 to SHA1.
- The HowToAuth parameter on the IpDataOffer statement is changed from HMAC_MD5 to HMAC_SHA1.

If you have an IPSec policy, determine whether this change affects your policy. If you use the IBM Configuration Assistant for z/OS Communications Server to configure your IPSec policy, an explicit HowToAuthMsgs value is generated on every KeyExchangeOffer statement and an explicit HowToAuth value is generated on every

Migrating to z/OS V2.3: Technical Migration Actions

IpDataOffer statement, so default values are not used. If you manually configure your IPsec policy, default values might be used.

Note: MD5 is considered a weak algorithm and is not recommended. Regardless of whether you use IBM Configuration Assistant for z/OS Communications Server or manually configure your policies, you should evaluate your usage of the MD5-based algorithms and decide whether to upgrade to a more secure algorithm.

Migration action: If your policy is not generated by IBM Configuration Assistant for z/OS Communications Server:

- Search your IPsec policy files for any KeyExchangeOffer statements that do not specify a HowToAuthMsgs parameter. If you find such a KeyExchangeOffer statement, your policy is affected. If you require the HowToAuthMsgs value to continue to use MD5, update your policy to explicitly set the HowToAuthMsgs parameter to MD5. If you want to use the new default of SHA1, you must coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peer's policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with the remote peer's policy, the IKE daemons will not be able to negotiate IPsec tunnels.
- Search your IPsec policy files for any IpDataOffer statements that do not specify a HowToAuth parameter. If you find such an IpDataOffer statement, your policy is affected. If you require the HowToAuth value to continue to use HMAC_MD5, update your policy to explicitly set the HowToAuth parameter to HMAC_MD5. If you want to use the new default of HMAC_SHA1, you must coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peer's policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with the remote peer's policy, the IKE daemons will no longer be able to negotiate IPsec tunnels.
- During this exercise, you should note any cases where your policy explicitly uses MD5-based algorithms. If you find any, consider changing your policy to use a stronger authentication algorithm. Before you make changes, you must coordinate with the owners of each remote IKE peer associated with the z/OS policy changes to ensure that the remote peers' policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with a remote peer's policy, the IKE daemons will not be able to negotiate IPsec tunnels.

If your policy is generated by IBM Configuration Assistant for z/OS Communications Server, evaluate your existing policy to determine whether MD5 is configured on any data offers or key exchange offers. If so, consider changing your policy to use a stronger authentication algorithm. Before you make changes, you must coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peers' policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with a remote peer's policy, the IKE daemons will not be able to negotiate IPsec tunnels.



IP Services: Verify that the changed HowToEncrypt default is acceptable (Required-IF, as of V2R3, and V2R2 and V2R1 with APAR PI74383)

Required if you use an IPsec policy.

In z/OS V2R3, the default value for the HowToEncrypt parameter on the KeyExchangeOffer and IpDataOffer statements in the IPsec policy is changed from DES to AES_CBC Keylength 128. If you have an IPsec policy, determine whether this change affects your policy. If you use the IBM Configuration Assistant for z/OS Communications Server to configure your IPsec policy, an explicit HowToEncrypt value is generated on every KeyExchangeOffer and IpDataOffer statement, so default values are not used. If you manually configure your IPsec policy, default values might be used.

Note that DES is considered a weak algorithm and is not recommended. Therefore, regardless of whether you use IBM Configuration Assistant for z/OS Communications Server or manually configure your policies, you should evaluate your usage of the DES and 3DES algorithms and decide whether to upgrade to a more secure algorithm.

Migration action: If your policy is not generated by IBM Configuration Assistant for z/OS Communications Server, do the following:

- Search your IPsec policy files for any KeyExchangeOffer statements and any IpDataOffer statements that do not specify a HowToEncrypt parameter. If you find such a KeyExchangeOffer statement or IpDataOffer statement, your policy is affected. If you require the HowToEncrypt value to continue to use DES, update your policy to explicitly set the HowToEncrypt parameter to DES. If you want to use the new default of AES_CBC Keylength 128, you must coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peer's policy is compatible with the z/OS changes. If

Migrating to z/OS V2.3: Technical Migration Actions

the z/OS policy changes so that it is incompatible with the remote peer's policy, the IKE daemons will not be able to negotiate IPsec tunnels.

- During this exercise, note any cases where your policy explicitly uses DES and 3DES algorithms. If you find any, consider changing your policy to use a stronger authentication algorithm. Before you make changes, you must coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peers' policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with a remote peer's policy, the IKE daemons will not be able to negotiate IPsec tunnels.

If your policy is generated by IBM Configuration Assistant for z/OS Communications Server, you should evaluate your existing policy to determine whether DES or 3DES are configured on any data offers or key exchange offers. If so, consider changing your policy to use a stronger authentication algorithm. Before you make changes, you must coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peers' policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with a remote peer's policy, the IKE daemons will not be able to negotiate IPsec tunnels.



IP Services: Verify that the changed DHGroup default is acceptable (Required-IF, as of V2R1 and V2R2, both with APAR PI43832, R13 with APAR PI43833)

Required if you use an IPsec policy.

In z/OS V2R2, the default value for the DHGroup parameter on the KeyExchangeOffer statement in the IPsec policy is changed from Group1 to Group2. If you have an IPsec policy, determine whether this change effects your policy. If you use the IBM Configuration Assistant for z/OS Communications Server to configure your IPsec policy, an explicit DHGroup value is generated on every KeyExchangeOffer statement. A default value is not used.

Migration action: If your policy is not generated by IBM Configuration Assistant for z/OS Communications Server, search your IPsec policy files for any KeyExchangeOffer statements that do not specify a DHGroup parameter. If you find such a KeyExchangeOffer statement, your policy is effected. If you require the DHGroup value to continue to use the previous default of Group1, update your policy to explicitly set the DHGroup parameter to Group1. If you want to use the new default, you need to coordinate with the owners of each remote IKE peer that is associated with the z/OS policy changes to ensure that the remote peer's policy is compatible with the z/OS changes. If the z/OS policy changes so that it is incompatible with the remote peer's policy, the IKE daemons will no longer be able to successfully negotiate IPsec tunnels.

Note: Diffie-Hellman group 1 is considered a weak algorithm and is not recommended.



IP Services: Verify that the changed ANONYMOUSLEVEL default is acceptable

Required if ANONYMOUS logon is enabled.

In z/OS V2R3, the default value for the ANONYMOUSLEVEL parameter for the FTP server is changed from 1 to 3. If you have ANONYMOUS login enabled, determine whether this change affects your configuration.

IBM suggests that ANONYMOUSLEVEL is set to 3 and ANONYMOUSFILETYPEJES is set to FALSE when ANONYMOUS is configured on the FTP server. Specifying ANONYMOUSLEVEL less than 3 or ANONYMOUSFILETYPEJES TRUE allows anonymous users to submit jobs.

With the new default of ANONYMOUSLEVEL 3, anonymous access is controlled by the following FTP.data statements:

- ANONYMOUSFILETYPESEQ
- ANONYMOUSFILETYPEJES
- ANONYMOUSFILETYPESQL
- ANONYMOUSFILEACCESS
- ANONYMOUSHFSFILEMODE
- ANONYMOUSHFSDIRMODE

Application health check CSAPP_FTPD_ANONYMOUS_JES can help you determine whether anonymous FTP users can submit jobs. This check is available in z/OS V2R3. It is also available in z/OS V2R1 and V2R2 with TCP/IP APAR PI47637 and SNA APAR OA49668 applied.

Migration action: Examine all instances of your FTP server configuration files (FTP.DATA) for an ANONYMOUS statement.

- If you do not have an ANONYMOUS statement configured, anonymous access is not enabled and ANONYMOUSLEVEL is ignored. No action is required.

Migrating to z/OS V2.3: Technical Migration Actions

- If you have an ANONYMOUS statement configured and an ANONYMOUSLEVEL statement is configured with an explicit value, no action is required.
- If you have an ANONYMOUS statement configured and are allowing ANONYMOUSLEVEL to default, evaluate what ANONYMOUSLEVEL is needed and take the corresponding action.
 - The new default of 3 for ANONYMOUSLEVEL along with the default value of FALSE for ANONYMOUSFILETYPEJES, help prevent job submissions by anonymous users.

Note: ANONYMOUSLEVEL 3 enables control of individual filetypes. If you choose to let ANONYMOUSLEVEL to default to 3, evaluate all the following filetype controls to ensure the required access is allowed. The anonymous filetype configuration statements are listed here with the default and allowed values.

ANONYMOUS TYPE	DEFAULT	ALLOWED VALUES
ANONYMOUSFILETYPESEQ	TRUE	FALSE TRUE
ANONYMOUSFILETYPEJES	FALSE	FALSE TRUE
ANONYMOUSFILETYPESQL	FALSE	FALSE TRUE

ANONYMOUS TYPE	DEFAULT	ALLOWED VALUES
ANONYMOUSFILEACCESS	HFS	BOTH MVS HFS
ANONYMOUSHFSFILEMODE	000	nnn
ANONYMOUSHFSDIRMODE	333	nnn

To get the pre-V2R3 behavior, explicitly configure ANONYMOUSLEVEL 1 in the relevant FTP server configuration data set (FTP.DATA).

Note: Specifying ANONYMOUSLEVEL less than 3 or ANONYMOUSFILETYPEJES TRUE allows anonymous users to submit jobs. Optionally, disable anonymous access by removing the ANONYMOUS keyword.

Communications Server Migration Actions Pre-First IPL



IP Services: Ensure TLS/SSL secure connection in non-FIPS mode meets the minimum peer end-entity certificate key size (Required-IF, as of V2R3)

Yes, if you use any affected functions.

System SSL is raising the minimum asymmetric key size for peer certificates used during the negotiation of a TLS/SSL secure connection in non-FIPS mode. The minimum key sizes are as follows:

- RSA changed to 1024 bits from 512 bits
- DSA changed to 1024 bits from 512 bits
- DH changed to 1024 bits from 512 bits
- ECC changed to 192 bits from 160 bits

Any of the listed Communications Server components can generate error messages or trace entries that indicate the specific error returned by System SSL during a TLS/SSL handshake. The new System SSL return code GSK_ERR_KEY_IS_SMALLER_THAN_MINIMUM (508) or GSK_ERROR_KEY_IS_SMALLER_THAN_MINIMUM (-127) is returned during the negotiation of the connection, if the peer provides an RSA, a DSA, or a DH certificate with a key size smaller than 1024 or an ECC certificate with a key size smaller than 192.

Migration action : Review each of the Communications Server components that follow to determine whether you are affected. Make changes as directed.

AT-TLS

To update the Application Transparent Transport Layer Security (AT-TLS) policy files manually, take the following steps:

1. Locate the TTLSSRule statement that applies to the traffic that still requires the weak key length.
2. Locate the TTLSEnvironmentAction statement that is referenced by or contained in the TTLSSRule statement.
3. Locate the TTLSEnvironmentAdvancedParms statement that is referenced by or contained in the TTLSEnvironmentAction statement.

Migrating to z/OS V2.3: Technical Migration Actions

4. In the `TTLSEnvironmentAdvancedParms` statement, code the `PeerMinRsaKeySize`, `PeerMinDsaKeySize`, `PeerMinDHKeySize`, or `PeerMinECCKeySize` parameters as appropriate with the required minimum key size.
5. Save the updated policy file and refresh Policy Agent according to your local site procedures to put the changes into effect.

FTP server and FTP client

When the FTP client or server is configured with `TLSMECHANISM ATTLS` in the `FTP.DATA` data set, AT-TLS is used to provide the TLS/SSL protection. Therefore, you must follow the steps for AT-TLS applications. When the FTP client or server is configured with `TLSMECHANISM FTP`, it calls System SSL directly. To enable the weaker key length in this case, set the appropriate `GSK_PEER_RSA_MIN_KEY_SIZE`, `GSK_PEER_DSA_MIN_KEY_SIZE`, `GSK_PEER_DH_MIN_KEY_SIZE`, or `GSK_PEER_ECC_MIN_KEY_SIZE` environment variable to specify the required minimum key size before starting the FTP server or FTP client program. As an alternative, you can choose to enable your FTP client or server for AT-TLS and then use the procedure described above for AT-TLS applications. See *z/OS Communications Server: IP Configuration Guide* for more information about converting FTP from using `TLSMECHANISM FTP` to `TLSMECHANISM ATTLS`.

TN3270E Telnet server

When the TN3270E Telnet server is configured with `TTLSPORT` in the Telnet profile, AT-TLS is used to provide the TLS/SSL protection. Therefore, you must follow the steps for AT-TLS applications. When the TN3270E Telnet server is configured with `SECUREPORT`, the related GSK environment variables cannot be passed to the TN3270E server. Therefore, the only way to enable the weaker key lengths is to enable the TN3270 for AT-TLS and then use the procedures described above for AT-TLS applications to enable the weaker key lengths. See *z/OS Communications Server: IP Configuration Guide* for more information about converting TN3270E from using `SECUREPORT` to `TTLSPORT`.

DCAS

When the Digital Certificate Access Server (DCAS) server is configured with `TLSMECHANISM ATTLS` in the DCAS configuration file, AT-TLS is used to provide the TLS/SSL protection. Therefore, you must follow the steps for AT-TLS applications.

When the DCAS server is configured with `TLSMECHANISM DCAS`, it calls System SSL directly. To enable the weaker key length in this case, set the appropriate `GSK_PEER_RSA_MIN_KEY_SIZE`, `GSK_PEER_DSA_MIN_KEY_SIZE`, `GSK_PEER_DH_MIN_KEY_SIZE`, or `GSK_PEER_ECC_MIN_KEY_SIZE` environment variable to specify the required minimum key size before starting the DCAS server. As an alternative, you can choose to enable your DCAS server for AT-TLS and then use the procedure described above for AT-TLS applications. See *z/OS Communications Server: IP Configuration Guide* for more information about converting DCAS from using `TLSMECHANISM DCAS` to `TLSMECHANISM ATTLS`.

Policy Agent Client

When the Policy Agent is configured with the `DynamicConfigPolicyLoad` statement in the main Pagent configuration file, it acts as a policy server and can be protected using AT-TLS. AT-TLS is used to provide the TLS/SSL protection. Therefore, you must follow the steps for AT-TLS applications. When the Policy Agent is configured with the `PolicyServer` and `ServerConnection` statement, it acts as a policy client. If the `ServerSSL` parameter is specified on the `ServerConnection` statement, the connection between the client and the remote policy server is protected by using System SSL directly.

To enable the weaker key length in this case, set the appropriate `GSK_PEER_RSA_MIN_KEY_SIZE`, `GSK_PEER_DSA_MIN_KEY_SIZE`, `GSK_PEER_DH_MIN_KEY_SIZE`, or `GSK_PEER_ECC_MIN_KEY_SIZE` environment variable to specify the required minimum key size before starting the Policy Agent.

IP Services: Permit Communications Server components to ICSF resources required by Network Authentication Service (Kerberos) (Required-IF, as of V2R3)

Required if you use Kerberos in certain conditions below.

In z/OS V2R3, Kerberos relies on ICSF PKCS#11 callable services for encryption, decryption, and hashing. As a result of this change, ICSF is required to be running before any Kerberos components or applications are running on the z/OS system. The following z/OS Communications Server components use Kerberos in certain situations and might therefore require access to the ICSF callable services that Kerberos uses:

- With the UNIX System Services Telnet server, clients can support Kerberos version 5, as described in RFC 1416, to log in to the shell environment by using Kerberos as the authentication protocol.
- FTP server and FTP client support connections to or from other servers and clients that support Kerberos version 5 authentication for the FTP protocol, as described in RFC 2228.
- UNIX System Services RSH server can be configured to support client authentication by using Kerberos from RSH clients that support Kerberos version 5.

In addition, the default encryption and checksum (hash) types that are used when not explicitly set in the Kerberos configuration files are being changed from weak and non-collision proof types to stronger, more secure types.

Migrating to z/OS V2.3: Technical Migration Actions

Migration action: Determine whether any of z/OS Communications Server components on your system are using Kerberos in the following conditions:

- If you use the UNIX System Services Telnet server (otelnetd), check your inetd configuration file (/etc/inetd.conf) for invocations of otelnetd that specify the -a user parameter. If you find such an invocation, your otelnetd server is using Kerberos version 5 authentication. If not, your otelnetd server is not affected.
- If you use the FTP server, check the FTP.DATA data set of your server for the EXTENSIONS AUTH_GSSAPI parameter. If you find this parameter, your FTP server supports Kerberos version 5 authentication. If not, your FTP server is not affected.
- If you use the FTP client command or API, check the FTP.DATA data set of each FTP client for the SECURE_MECHANISM GSSAPI parameter. If you find this parameter in any of those FTP.DATA data sets, those FTP clients use Kerberos version 5 authentication. Additionally, search for any invocations of the FTP command or FTP Client API that use the -a GSSAPI or the -r GSSAPI parameter, which also result in the use of Kerberos version 5 authentication. Any clients that use any of the above mechanisms are affected.
- If you use the UNIX System Services RSH daemon (orshd), check your inetd configuration file (/etc/inetd.conf) for orshd with the -k KRB5 or the -k GSSAPI parameter. If you find these parameters, your orshd daemon uses Kerberos version 5 authentication. If not, your RSH daemon is not affected.

If any of the above components use Kerberos, you must perform the associated migration actions, as follows:

1. Ensure that ICSF is started and completes initialization before starting the z/OS KDC or any Kerberized applications on the system. ICSF needs to be running for the duration of use of all Kerberos functions, including KDC, application servers, application clients, commands, and utilities.
2. If the CSFSERV class is active, ensure that the z/OS KDC user ID and all user IDs that use Kerberos commands or Kerberized application running on the z/OS system have read access to the following ICSF resources:
 - a. When Kerberos is enabled for FIPS 140 mode: CSFRNG, CSFOWH, CSF1TRC, CSF1TRD, CSF1SKD, and CSF1SKE
 - b. When Kerberos is not enabled for FIPS 140 mode: CSFRNG and CSFOWH For the UNIX System Services Telnet server, the KDC user ID is the user ID under which the otelnetd -a user command is started.

For the FTP server that is configured with the EXTENSIONS AUTH_GSSAPI parameter, the KDC user ID is the user ID under which the FTP server runs.

For FTP clients, the KDC user ID is any user ID that starts the FTP client command or API with the -a GSSAPI or -r GSSAPI parameter or that starts the FTP client with the SECURE_MECHANISM GSSAPI parameter specified in their FTP.DATA data sets.

For the UNIX System Services RSH server, the KDC user ID is the user ID under which the orshd -k KRB5 or -k GSSAPI command is issued.

3. The default encryption types for Kerberos applications have changed from DES encryption types to stronger encryption types, AES and 3DES. If the former default encryption types are still required, they must be explicitly set in the Kerberos configuration files, /etc/skrb/krb5.conf by default, by issuing the following commands:

```
default_tgs_enctypes = des-cbc-crc,des-cbc-md5
```

```
default_tkt_enctypes = des-cbc-crc,des-cbc-md5
```

4. The default checksum types for Kerberos applications have changed from obsolete checksum types to more modern and secure checksum types. If the former defaults values are still required, they must be explicitly set in the Kerberos configuration files, /etc/skrb/krb5.conf by default, by issuing the following commands:

```
ap_req_checksum_type = rsa-md5
```

```
kdc_req_checksum_type = rsa-md5
```

```
safe_checksum_type = rsa-md5-des
```

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:

Migrating to z/OS V2.3: Technical Migration Actions

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	Target location	What changed and when
Communications Server z/OS UNIX applications	/etc/services	In z/OS V2R2, the NCPROUT entry is removed because NCPROUTE is no longer supported. If you update /etc/services, ensure that you also update the ETC.SERVICES data set.
DCAS	/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 Client	/etc/ftp.data	In z/OS V2R2, a new SSLV3 keyword is provided to configure SSLv3 protocol for connections secured using the FTP TLS support.
Internet Key Exchange Daemon(IKED)	/etc/security/iked.conf	In z/OS V2R2, a new log level is added for the IKE daemon.
Policy agent	/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.
sendmail	/etc/mail/sample.cf	In z/OS V2R3, this sample daemon configuration file for the sendmail application is removed. sendmail is no longer supported.
sendmail	/etc/mail/submit.cf	In z/OS V2R3, this sample daemon configuration file for the sendmail application is removed. sendmail is no longer supported.
sendmail	/etc/mail/zOS.cf	In z/OS V2R3, this sample daemon configuration file for the sendmail application is removed. sendmail is no longer supported.
sendmail bridge	/etc/mail/ezatmail.cf	In z/OS V2R3, this sample file was added for the z/OS sendmail to CSSMTP bridge.
SNMP agent	/etc/osnmpd.data	Every release, the value of the sysName MIB object is updated to the current release.
SNMP agent	/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.

Migrating to z/OS V2.3: Technical Migration Actions

Utility	Target location	What changed and when
z/OS UNIX snmp command	/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.

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 a 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.

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:

Migrating to z/OS V2.3: Technical Migration Actions

```
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.

CSM buffer pools are in 31-bit backed data space, 64-bit backed data space, and in ECSA. Data space storage is a common area data space and is associated with the master scheduler address space. This association results in a data space that persists for the life of the system.

Data space storage is either 31-bit backed or 64-bit backed. 31-bit backed data space, when fixed, resides below the 2 GB real storage bar. 64-bit backed data space, when fixed, can reside below or above the 2 GB bar. Where the storage is backed is a concern for only those products performing I/O into or out of the storage.

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

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**
-  3. **JES2: Activate z11 mode**
- 4. **z/OS OpenSSH: Accommodate the OpenSSH ported level**


Marna's "Big Migs" for Migrating from V2R2 to V2R3



Migration actions from V2.2, plus...

-  1. **8 GB memory requirement for z14**
- 2. **z/OSMF Autostart**
- 3. **DFSMSdfp positioning for data set encryption**
- 4. **ICSF configured and running, everywhere.**
- 5. **SDSF/SDSFAUX and ensure SDSF class is RACLISTed.**

Future migration actions to do now:

- A. **BookManager READ removal after V2.3.**
-  B. **HFS removal planning for 2021 release.**
- C. **Future tape removal: use DVD or electronic deliveries.**

Migrating to z/OS V2.3: Migration Actions Summary



- **General Migration Actions:**
 - New address spaces, new and old data sets, changed checks.
- **BCP Migration Actions:**
 - INCLUDE1MAFC(NO) removal, IEASYSxx REAL=0 default, HZR starting, Default logstream ds minimum sizes, new JCL keywords no as system symbols, LOGR CDS at HBB7705, NUMBER(1) default for SMDUPLEX for LOGR CDS.
- **Communications Server Migration Actions:**
 - Defaults changed for security: HowToAuthMsgs, HowToAuth, ANONYMOUSLEVEL, DHGROUP, minimum asym key size for TLS/SSL secure connection.
- **DFSMS Migration Actions:**
 - CA_RECLAIM(DATACLAS) default, REFUCB is default is enabled.

34

© 2017 IBM Corporation

Migrating to z/OS V2.3: Migration Actions Summary



- **zFS Migration Actions:**
 - IOEFSPRM default changes for romount_recovery , format_aggrversion, change_aggrversion_on_mount, and honor_syslist
- **SDSF Migration Actions:**
 - New main panel.
- **RACF Migration Actions:**
 - Only three CA Certificates shipped
- **z/OS UNIX Migration Actions:**
 - /global in sysplex root, /var/man removals.
- **KC4z Migration Actions:** /global usage.
- **HLASM Migration Actions:** IEV90 alias so no usermod.
- **XL C/C++ Migration Actions:** ARCH(10) and TUNE(10) defaults.



35

© 2017 IBM Corporation