



Upgrading* to z/OS V2.4 Part 2 of 2: Technical Actions



Marna WALLE
Member of the IBM Academy of Technology
IBM System Z
Poughkeepsie, New York USA
mwalle@us.ibm.com



February 2020

* **Upgrading** is the new term for migrating!

Copyright © by SHARE Association. Except where otherwise noted, this work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 license. <http://creativecommons.org/licenses/by-nc-nd/3.0/>

© 2020 1

Abstract:

Yes, "upgrade" is the new name for these traditional "migration" sessions! This is part one of a two-part session that will be of interest to System Programmers and their managers who are upgrading to z/OS 2.4 from either z/OS 2.2 or 2.3. It is strongly recommended that you review sessions for a complete upgrade picture.

The general availability date for z/OS V2.4 was September 30, 2019.

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:

* AS/400®, e-business/Logo®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/390, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, 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. Call Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. IFL 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 multi-programming 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 environment, costs and performance characteristics will vary depending on individual customer configurations and conditions. This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

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

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

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

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

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

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

2

© 2020 IBM Corporation

Upgrade to z/OS V2.4 Part 2: Technical Actions Agenda

- Scope of presentation
- Definition of a "upgrade action"
- Overview of upgrade actions for z/OS V2R4 from z/OS V2R3 or V2R2:


- ❖ General Upgrade Actions
 - ❖ BCP
 - ❖ DFSMS
 - ❖ Various elements having security default changes
 - ❖ ICSF and ICSF dependencies
 - ❖ z/OSMF
 - ❖ SDSF
 - ❖ Security Server – RACF
 - ❖ z/OS OpenSSH
 - ❖ z/OS UNIX
 - ❖ JES2
 - ❖ Communications Server



3

© 2020 IBM Corporation

Scope of Presentation




- This presentation is applicable to z/OS V2R4 upgrades from either z/OS V2R3 or V2R2.
- Not fully inclusive of all upgrade actions, but rather gives you an overview of some upgrade actions that are:
 - Very important to understand
 - May be common to many users
- Remember: Use *z/OS V2.4 Upgrade Workflow* for a complete list of all technical upgrade actions.
 - The latest level of the workflow is at [Github for zOS](#)
 - The latest level of the *exported* workflow is at [Upgrade Abstract web page](#)

The specific *Workflow* is targeted to your specific upgrade path:

Pick one!

 {

- Upgrade from V2.3 to V2.4
- Upgrade from V2.2 to V2.4


© 2020 IBM Corporation

To use the latest version of the z/OS V2.4 Upgrade Workflow, retrieve it from this location:

<https://github.com/IBM/IBM-Z-zOS/tree/master/zOS-Workflow/zOS%20V2.4%20Upgrade%20Workflow>

IBM strongly recommends you use the z/OS Upgrade Workflow from z/OSMF in order to have a customized upgrade of your applicable steps to take.

However, if you wish to use the exported version of that z/OS V2.4 Upgrade Workflow, you can see it here (at the bottom of the page):


https://www.ibm.com/support/knowledgecenter/SSLTBW_2.4.0/com.ibm.zos.v2r4.e0zm100/abstract.htm

Upgrade is not Exploitation!



- **Upgrading to a new z/OS release is a two step process:**
 1. **Upgrade:** the installation of a new version or release of a program to replace an earlier version or release. (Formerly called "migration".)
 2. **Exploitation:** usage of new enhancements available in the new release. Not covered in this presentation.
- **After a successful upgrade, the applications and resources on the new system function the same way they did on the old system, if possible.**
- **Upgrade 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 workload.
- **Upgrade actions are also classified as when they may be performed:**
 - **NOW, Pre-First IPL, or Post-First IPL**

5



Upgrade Definitions and Classifications

Upgrade (formerly, migration) is the first of two stages in upgrading to a new release of z/OS. The two stages are:

- **Stage 1: Upgrade.** During this stage you install your new system with the objective of making it functionally compatible with the previous system. After a successful upgrade, 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. Upgrade 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 upgrade.

Upgrade Requirement Classification and Timing

The upgrade actions are classified as to their requirement status:

- **Required.** The upgrade action is required in all cases.
- **Required-IF.** The upgrade action is required only in a certain case. Most of the actions in this presentation are in this category.
- **Recommended.** The upgrade 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 upgrade actions, this presentation uses three types of headings:

- **Now.** These are upgrade 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 V2R4 level of code to make these changes, and the changes don't require the z/OS V2R4 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 V2R4.
- **Pre-First IPL.** These are upgrade actions that you perform after you've installed z/OS V2R4 but before the first time you IPL. These actions require the z/OS V2R4 level of code to be installed but don't require it to be active. That is, you need the z/OS V2R4 programs, utilities, and samples in order to perform the upgrade actions, but

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

the z/OS V2R4 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 upgrade actions in this category even earlier. If you prepare a system on which you will install z/OS V2R4 by making a clone of your old system, you can perform upgrade actions that involve customization data on this newly prepared system before installing z/OS V2R3 on it. Examples of such upgrade actions are updating configuration files and updating automation scripts.

- **Post-First IPL.** These are upgrade actions that you can perform only after you've IPLed z/OS V2R4. You need a running z/OS V2R4 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 V2R4 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 upgrade action.



means that an IBM Health Check (using the IBM Health Checker for z/OS function) can help you with this upgrade 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 Upgrade Actions for z/OS V2R4

*These elements have **new V2.3 → V2.4** upgrade actions:*

- BCP
- Communications Server
- CIM
- Cryptographic Services –ICSF and PKI Services
- DFSMS
- Distributed File Service - SMB
- HCD and HCM
- HLASM
- z/OS Management Facility
- Infoprint Server
- ISPF
- JES2
- JES3
- Library Server
- OSA/SF
- RMF
- Security Server (RACF)
- OpenSSH
- SMP/E
- TSO/E
- XL C/C++
- z/OS UNIX



➤ means that some of that element's upgrade actions are discussed in these two presentations

Only elements with new upgrade actions in V2.4 are mentioned on the list above.

6

© 2020 IBM Corporation

New Upgrade Actions for Elements in z/OS V2.4

When upgrading from z/OS V2.2 or V2.3 to z/OS V2.4, the specified elements in the slide above have new upgrade actions. Refer to the *z/OS Upgrade Workflow* for complete information on the required upgrade actions for all elements, and if you would like to see the upgrade actions introduced in V2.3. Some upgrade actions for selected elements follow in this presentation. This presentation does not cover all possible upgrade actions.

General Upgrade Actions for z/OS V2.4



• Upgrade Actions Pre-First IPL: Upgrade and Exploitation

- **Accommodate new address spaces (Recommended)**
 - New in V2R4:
 - **z/OS Container Extensions:**
 - Provides runtime support to deploy and run Linux on IBM Z applications that are packaged as Docker Container images on z/OS.
 - One new address space for each provisioned server
- New in V2R3:
- **jesxEDS for JES2 Email Delivery Service**, used when JES2 email interfaces are used. Needs a user ID assigned to this address space which can be the 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 and SDSF**, for SDSF.
- **HZR**, for Runtime Diagnostics (see later).



7

© 2020 IBM Corporation

General Upgrade Actions for z/OS V2.4



Upgrade Actions Pre-First IPL:

- **Remove references to deleted syslib data sets and paths (Required)**

 - Removed in V2R4: **BookManager Read, SMB, NLVs** (except for JPN and ENP), **Library Server, OSA/SF**
 - Removed in V2R3: Some **BCP** and **HCD** paths, **SDSF's AISFLINK**, and an **HCD DLIB**.
- **Add references to new syslib data sets and paths (Required)**
 - New in V2R4: **zCX path** /usr/lpp/zcx_zos/IBM/ and dlib.
 - 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.
- **Update your health check customization for modified checks (Recom)**

 - New in V2R3: 6 checks. V2R4: 8 checks
 - Changed in V2R3: 3 checks V2R4: 1 checks
 - Deleted in V2R3: 10 checks V2R4: 0 checks

8

© 2020 IBM Corporation

General Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

General Upgrade Actions You Can Do Now

Install coexistence and fallback PTFs (Required)

Upgrade action: Install coexistence and fallback PTFs on your systems to allow those systems to coexist with z/OS V2.4 systems during your upgrade, and allow back out from z/OS V2.4 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 V2.4 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 V2.4 systems and specify a Fix Category (FIXCAT) value of **"IBM.Coexistence.z/OS.V2R4"**. 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 upgrading to new software levels, and when upgrading to z/Architecture.

Upgrade 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 Upgrade Actions Pre-First IPL

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

Upgrade action: The /etc, /global, and /var directories contain system control files: the /etc directory contains customization data that you maintain and the /global and /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 V2.4 /etc and /var subdirectories.

Copy files from your old system to the z/OS V2.4 /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 elements and features use /etc:

- BCP (Predictive Failure Analysis).
- CIM.
- Communications Server (IP Services component).
- Cryptographic Services (PKI Services and System SSL components).
- DFSMSrmm.
- IBM HTTP Server.
- IBM Tivoli Directory Server (TDS). The LDAP server component uses /etc/ldap.
- Infoprint Server.
- Integrated Security Services. The Network Authentication Service component uses /etc/skrb.
- z/OS UNIX.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

The following elements and features use /global:

- IBM Knowledge Center for z/OS
- IBM z/OS Management Facility (z/OSMF).

The following elements and features use /var:

- Cryptographic Services (OCSF component). See OCSF: Migrate the directory structure.
- DFSMSrmm.
- IBM Tivoli Directory Server (TDS). The LDAP server component uses /var/ldap.
- Infoprint Server.
- Integrated Security Services. The Network Authentication Service component uses /var/skrb.

Back virtual storage with real and auxiliary storage (Required)

Upgrade 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)

Upgrade action: Using the tables in *z/OS Upgrade Workflow* 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 Upgrade Workflow* contains the list of all removed data sets and paths in z/OS V2R.4 and V2.3.

Add references to new data sets (Required)

Upgrade action: For z/OS V2.4, the following element had a DLIB data set and a path (with a file system) that were added:

- z/OS Container Extensions

For z/OS V2.3, the following elements had data sets and paths that were added:

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

Accommodate new address spaces (Recommended)

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

The following element adds new address spaces for z/OS V2.4:

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

- **z/OS Container Extensions.** This a new element in z/OS V2.4. It provides the runtime support to deploy and run Linux on IBM Z applications that are packaged as Docker Container images on z/OS. zCX creates one address space for each provisioned server that is started by the user.

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 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 are new in z/OS V2.4:

- IBMCS,ZOSMIGV2R4_NEXT_CS_OSIMGMT
- IBMCS,ZOSMIGV2R4PREV_CS_IWQIPSEC_tcpipstackname
- IBM_JES2, JES2_UPGRADE_CKPT_LEVEL_JES2
- IBMICSF,ICSF_PKCS_PSS_SUPPORT
- IBMINFOPRINT, INFOPRINT_CENTRAL_SECURE_MODE
- IBMINFOPRINT, ZOSMIGV2R3_NEXT_INFOPRINT_IPCSSL
- IBMISPF,ISPF_WSA
- IBMVSM,ZOSMIGV2R3_NEXT_VSM_USERKEYCOMM

The following health checks are changed by IBM in z/OS V2.4:

- IBMUSS,ZOSMIGV2R3_NEXT_USS_SMB_DETECTED

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
- ZOSMIGV2R3_RSM_Minimum_Real (and RSM_Minimum_Real)

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_SMTPD_MAIL_RELAY
- 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_SMTPDDAEMON
- ZOSMIGV2R2_NEXT_CS_SMTPDMTA
- ZOSMIGV2R2_Next_CS_TFTP

Upgrade 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 Upgrade Actions for z/OS V2.4



Upgrade Actions Before First IPL:

- **Stop referencing the ETR parameters in CLOCKxx (Recommended, as of V2.4)**
 - In shipped CLOCK00 parmlib member, ETRMODE and ETRZONE parameter defaults are changed from YES to NO.
 - z10 was the last IBM mainframe to support the ETRMODE and ETRZONE parameters in CLOCKxx.
 - Because z/OS V2R4 requires a zEC12 or later processor to run, the ETRMODE and ETRZONE parameters are now obsolete. They are deactivated (set to NO) by default.
- If your CLOCKxx specifies ETRMODE YES or ETRZONE YES, set these parameters to NO, or remove them so that they default to NO.
- Specifying YES can cause unexpected behavior if you also specify STPZONE of NO. For details, see APAR OA54440.
 - OA54440: concerns MSGIEA598I not displaying the seconds correctly based on combination of STPZONE of NO, ETRMODE and ETRZONE (defaulting to YES).



9

© 2020 IBM Corporation

BCP Upgrade Actions for z/OS V2.4



Upgrade Actions Before First IPL:

- **Evaluate the changed default for system logger's use of IBM zHyperwrite (Required-IF, as of OA57408)**
 - IBM zHyperWrite data replication by system logger is changed from enabled to disabled by default.
 - In IXGCNFx:


```
MANAGE . . . HYPERWRITE ALLOWACCESS (YES|NO)
```
 - DISPLAY LOGGER, IXGCNF, MANAGE to see whether system logger on your system uses IBM zHyperwrite data replication
 - To enable system logger use of IBM zHyperwrite data replication, you can do either of the following:
 - Enter the command SETLOGR
MANAGE, HYPERWRITE, ALLOWUSE (YES)
 - Update the IXGCNFxx parmlib member with ALLOWUSE (YES) .
 - zHyperwrite is used for Staging data sets for DASDONLY type log streams, and Offload data sets for both the DASDONLY type and Coupling Facility structure-based type log streams.
 - System logger does not use zHyperwrite for the staging data sets for Coupling Facility structure-based type log streams.

10

© 2020 IBM Corporation

BCP Upgrade Actions for z/OS V2.4



Upgrade 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 as of z/OS V2.3 to identify V=R jobs.
- **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 USESTAGINGMIN are YES by default.

11

© 2020 IBM Corporation

BCP Upgrade Actions for z/OS V2.4



Upgrade 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.

12

© 2020 IBM Corporation

BCP Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

BCP Upgrade 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.

Upgrade 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.

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

Upgrade 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.

BCP Upgrade Actions Pre-First IPL

Stop referencing the ETR parmeters in CLOCKxx (Recommended, as of V2.4)

In z/OS V2R4, the default values of the CLOCKxx parmlib member ETRMODE and ETRZONE parameters are changed from YES to NO.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

The IBM z10 was the last IBM mainframe to support the ETRMODE and ETRZONE parameters in CLOCKxx. Because z/OS V2R4 requires a zEC12 or later processor to run, the ETRMODE and ETRZONE parameters are now obsolete. They are deactivated (set to NO) by default.

It is recommended that you set these parameters to NO, or remove them from CLOCKxx. Specifying YES can cause unexpected behavior if you also specify STPZONE NO. For details, see APAR OA54440.

Upgrade action: Review the CLOCKxx member that you use to IPL your system, and take one of the following actions:

- If CLOCKxx does not specify ETRMODE or ETRZONE, you have no action to take.
- If CLOCKxx specifies ETRMODE YES or ETRZONE YES, set these parameters to NO, or remove them so that they default to NO.
- If CLOCKxx specifies ETRMODE NO and ETRZONE NO, you have no action to take.

Evaluate the changed default for system logger's use of IBM zHyperwrite (Required-IF, as of V2.4)

Required if you want system logger to use IBM zHyperwrite data replication.

With the installation of APAR OA57408, the use of IBM zHyperWrite data replication by system logger is changed from enabled to disabled by default. The IBM zHyperWrite function was originally provided by APAR OA54814.

In parmlib member IXFCNFx, this option is specified, as follows:

```
MANAGE . . . HYPERWRITE ALLOWACCESS(YESNO)
```

When IBM zHyperwrite is enabled for use with system logger, it provides data replication for the following types of log stream data sets:

- Staging data sets for DASDONLY type log streams
- Offload data sets for both the DASDONLY type and Coupling Facility structure-based type log streams.

System logger does not use zHyperwrite for the staging data sets for Coupling Facility structure-based type log streams.

Upgrade action: If you do not want system logger to use IBM zHyperwrite data replication, you have no action to take. When you install APAR OA57408, the use of IBM zHyperWrite data replication by system logger is disabled by default. To determine whether system logger on your system uses IBM zHyperwrite data replication, you can query the current setting by entering the following command:

```
DISPLAY LOGGER,IXGCNF,MANAGE
```

In response, message IXG6071 identifies the current setting. To enable system logger use of IBM zHyperwrite data replication, you can do either of the following:

- Enter the command SETLOGR MANAGE,HYPERWRITE,ALLOWUSE(YES)
- Update the IXGCNFxx parmlib member with ALLOWUSE(YES).



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 IBMRSM,RSM_REAL checks the current setting for the REAL parameter in IEASYSxx.

Upgrade action: On a z/OS V2R2 system, do the following:

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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 or V2R4 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 (Required, as of V2.4)

IBM strongly recommends that you eliminate all use of user key common storage. Allowing programs to obtain user key (8-15) common storage creates a security risk because the storage can be modified or referenced, even if fetch protected, by any unauthorized program from any address space. Therefore, without the restricted use common service area (RUCSA) optional priced feature, the obtaining of user key (8-15) storage is not supported in z/OSV2R4.

RUCSA is more secure because it can be managed as a SAF resource. However, it does not prevent two or more different SAF-authorized applications from altering or referencing another application's RUCSA storage. RUCSA became available as a part of the BCP base element with APAR OA56180 on earlier z/OS releases, but it is only available as a priced feature in z/OS V2R4.

Related to this change:

- Support to change the key of common ESQA storage to a user key (via CHANGKEY) is removed regardless of RUCSA exploitation.
- NUCLABEL DISABLE(IARXLUK2) is no longer a valid statement in the DIAGxx parmlib member. (Note: This statement only exists in z/OS V2R3.)
- Support of user-key (8 - 15) SCOPE=COMMON data spaces is removed regardless of RUCSA exploitation.
- YES is no longer a valid setting for the following statements in the DIAGxx parmlib member: VSM ALLOWUSERKEYCSA - controls the allocation of user key CSA. ALLOWUSERKEYCADS - controls the allocation of user key SCOPE=COMMON data space

Upgrade action:

- See Part 1 (Planning) information for the details on this removal.

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.

Upgrade action: Follow these steps:

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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'
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 IRR8131 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.

Upgrade 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 *
```

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

DATA TYPE(LOGR) REPORT(YES)

/*

If the IXCMIAPI 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 IXCMIAPI utility or IXGINVNT API to update the logstream by specifying either DUPLEXMODE(COND) or DUPLEXMODE(UNCOND). 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 USEOFFLOADMIN 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.

Upgrade 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 USEOFFLOADMIN(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.

Upgrade action: Do the following:

- a. Install the PTFs for APAR OA49506 on all z/OS V2R2 systems in the same sysplex as the z/OS V2R3 and V2R4 systems (PTFs have been marked with the appropriate coexistence FIXCAT.)
- 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.

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.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

Upgrade action: Search for a symbol that is named TVSAMCOM, TVSMSG, 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, TVSMSG, or REGIONX.

Example:

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

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

Examples:

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

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

Examples:

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

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

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

As of 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.

Upgrade 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.

DFSMS Upgrade Actions for z/OS V2,4



• Upgrade Actions Before Installing:

- **DFSMSdfp: Increase storage for SMS ACDS data sets (Req-IF, as of OA52913)**
 - Length of the management class definition (MCD) in the SMS active control data set (ACDS) was increased from 328 to 760 bytes.
 - To support of automatic migration support for transparent cloud tiering (TCT).
 - With this change, the values that are used to estimate the size of storage that is needed for an active control data set are out of date.
- Check your utilization of the ACDS. To prevent the ACDS from reaching full, consider reallocating the ACDS to allow for more space.
 - When you allocate the SCDS and ACDS, specify secondary space allocations. This action helps to ensure that extends can be performed when:
 - New classes, groups, and other structures are added
 - Sizes of these structures increase in size.
 - See DOC APAR OA55563 for the most recent calculation.

13

© 2020 IBM Corporation

DFSMS Upgrade Actions for z/OS V2,4



• Upgrade Actions Before First-IPL:

- **DFSMSdss: SHARE keyword is ignored for COPY and RESTORE of PDSE (Required-IF, as of V2.4)**
 - Your programs must close PDSE data sets before these data sets are overwritten by a COPY or RESTORE operation. Otherwise, could see serialization errors, such as an ADR412E.
 - TOLERATE(ENQFAILURE) can be used, but the program is responsible for ensuring data consistency – use at your own risk.
- **DFSMSdfp: Review ANTXINxx XRC parmlib members for default changes (Recommended, as of V2.4)**
 - Several defaults have been changed to match recommended settings. Review settings to ensure the desired behavior will happen.

Value	Old	New	Reason
SCDumpType	STATESAV	NDSS	Uses less disruptive mechanism to take storage control statesaves for diagnosis
ReadDelay	1000	250	XRC record sets are read from primary storage control more frequently.
SuspendOnLongBusy	NO	YES	XRC is suspended immediately if the cache fills up on primary disk, resulting in less application impact.
TracksPerRead	3	12	(Match..) Volume initialization and resync operations are faster and more reliable
TracksPerWrite	3	12	(Match) Volume initialization and resync operations are faster and more reliable

14

DFSMSdfp Upgrade Actions for z/OS V2.4



• Upgrade 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.
- 15 • Can be changed dynamically with `SETSMS`.



© 2020 IBM Corporation

DFSMSdfp Upgrade Actions for z/OS V2.4



• Upgrade 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*.

16

© 2020 IBM Corporation

DFSMS Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

DFSMS Upgrade Actions You Can Do Now

DFSMSdfp: Increase storage for SMS ACDS data sets (Required-IF, as of OA52913)

Required if you use the ACDS and it is reaching full capacity.

With APAR OA52913, the length of the management class definition (MCD) in the SMS active control data set (ACDS) was increased from 328 to 760 bytes. This change was made in support of automatic migration support for transparent cloud tiering (TCT). With this change, the values that are used to estimate the size of storage that is needed for an active control data set are out of date.

If the ACDS data set is full, the follow errors could occur:

- IEC070I 203: Extend was attempted but no secondary space was specified
- IGD058I 6068 (X'17B4'): Problem could be caused by insufficient space to extend the data set

Check your utilization of the ACDS. To prevent the ACDS from reaching full, consider reallocating the ACDS to allow for more space. When you allocate the SCDS and ACDS, specify secondary space allocations. This action helps to ensure that extends can be performed when:

- New classes, groups, and other structures are added
- Sizes of these structures increase in size.



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.

Upgrade 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_s_name RECLAIMCA | NORECLAIMCA**.

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.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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.

Upgrade 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.

DFSMS UpgradeActions Pre-First IPL

DFSMSdss: SHARE keyword is ignored for COPY and RESTORE of PDSE (Req-IF, as of V2.4)

Required if your programs open PDSE data set when DFSMSdss is writing to them.

Starting in z/OS V2R4, your programs must close PDSE data sets before these data sets are overwritten by a DFSMSdss COPY or RESTORE operation. Otherwise, the DFSMSdss COPY or RESTORE operation encounters serialization errors, such as an ADR412E.

Note: This behavior occurs regardless of whether the SHARE keyword is specified.

For programs that cannot close the PDSE data sets while a COPY or RESTORE operation is in progress, these programs can specify the TOLERATE(ENQFAILURE) keyword. If so, the program is responsible for ensuring data consistency. IBM recommends against using the TOLERATE(ENQFAILURE) keyword; use it at your own risk.

Upgrade action: Applications must now close PDSEs before a copy or restore is to overwrite it. If you do not close down the application, serialization errors such as an ADR412E occurs on the DFSMSdss copy or restore.



Note: Changing DFSMSdss to ignore the SHARE keyword for output data sets ensure integrity and prevent DFSMSdss from restoring a PDSE that an application has open. It also prevents an application from opening a PDSE that DFSMSdss is restoring.



DFSMSdfp: Review XRC parmlib members for use of default values (Recommended, as of V2.4)

Required if your programs open PDSE data set when DFSMSdss is writing to them.

In z/OS V2R4, some XRC (Extended Remote Copy, aka Global Mirror), default settings are changed to match IBM recommendations. If you depend on the old default, you must explicitly specify the old values in the appropriate parmlib members.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

Value	Old	New	Reason
SCDumpType	STATESAV	NDSS	Uses less disruptive mechanism to take storage control statesaves for diagnosis
ReadDelay	1000	250	XRC record sets are read from primary storage control more frequently.
SuspendOnLongBusy	NO	YES	XRC is suspended immediately if the cache fills up on primary disk, resulting in less application impact.
TracksPerRead	3	12	(Match..) Volume initialization and resync operations are faster and more reliable
TracksPerWrite	3	12	(Match) Volume initialization and resync operations are faster and more reliable

Upgrade action: Examine parmlib members that are related to XRC for settings with changed defaults:

If the setting is not specified in any of the XRC-related parmlib members, determine whether the old default value needs to be retained according to the following table. If the value needs to be the old default, it must be specified in the member. Otherwise, the new default takes effect when XRC is started on the new z/OS release.

If the setting is specified in an XRC-related parmlib member (with any value, but especially with the old default value), consider whether the new default value is appropriate to use.

Upgrade Actions for z/OS V2.4



Various element default changes for more secure communications (V2R4)



CIM: Accommodate the default change from HTTP to HTTPS

- By default, CIM server listens on the HTTPS port (5989) rather than the HTTP port. New configuration defaults are:
 - `enableHttpConnection=false`
 - `enableHttpsConnection=true`
 - A client connection to this port must be secured with AT-TLS.
- **PKI Services: Ensure that users have the CA root certificate for the PKI web interfaces.**
 - Before, the first PKI web page presented by PKI Services for downloading the CA root certificate of the web server cert into the web browser, used HTTP.
 - As of V2.4, this first web page uses HTTPS, so you need to use alternate method to distribute the CA root certificate to the appropriate users' web browsers.

These changes are intended to allow for more secure communication, by default.

17

© 2020 IBM Corporation

Upgrade Actions for z/OS V2.4



Various element default changes for more secure communications (V2R4)



Infoprint Central requires SSL connections by default

- Previously, Infoprint Central GUI worked with or without SSL.
 - As of V2.4, SSL is required by default.
 - This means enabling SSL for IBM HTTP Server Apache.
- **RMF: Configure AT-TLS to enable secure communication with the RMF distributed data server**
 - Before, GPMSRVxx parmlib member option **HTTPS (NO)** was the default
 - As of V2.4, GPMSRVxx parmlib member option **HTTPS (ATTLS)** is the new default.
 - As a result, the RMF Distributed Data Server (DDS) ensures incoming connections are secured by AT-TLS, or it is refused.

These changes are intended to allow for more secure communication, by default.

18

© 2020 IBM Corporation



Various element default changes for more secure communications

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.



CIM: Accommodate the default change from HTTP to HTTPS (Required, as of V2.4)

In z/OS V2R4, the common information model (CIM) server is changed to use HTTPS connections by default. On start-up, the CIM server listens on the HTTPS port, rather than the HTTP port, as was done in previous releases. The change from HTTP to HTTPS is intended to allow for more secure communication with the CIM server.

The following CIM server configuration defaults are affected:

- `enableHttpConnection=true`
- `enableHttpsConnection=false`

In z/OS V2R4, these defaults are changed to:

- `enableHttpConnection=false`
- `enableHttpsConnection=true`

Based on this change, the CIM server opens a listener on port 5989 by default. If a client connection on this port is not secured by AT-TLS, the connection is closed and an appropriate error message is issued on the operations console.

A new migration health check is planned to be delivered on z/OS V2R2 and z/OS V2R3.

It is recommended that you use HTTPS instead of HTTP, which allows the CIM server to use Application Transparent Layer Security (AT-TLS) functions. Here, the communication between the CIM client and the CIM server is encrypted.

You must also configure the CIM client applications to use HTTPS. Ensure that the applications run as CIM clients connected to the CIM server on HTTPS port 5989.

Fall back to HTTP is possible, though it is not secure and therefore not recommended. If you need the CIM server to use the HTTP protocol on start-up, follow these steps:

- Start the CIM server with the following settings specified in the configuration properties file:
 - `enableHttpConnection=true`
 - `enableHttpsConnection=false`
- Alternatively, you can dynamically change the `enableHttpConnection` value and restart the CIM server by entering either of these commands on the operations console:
 - `cimconfig -s enableHttpConnection=true -p`
 - `F CFZCIM,APPL=CONFIG,enableHttpConnection=true,PLANNED`

PKI Services: Ensure that users have the CA root certificate for the PKI web interfaces (Required-I, as of V2.4)

Required if you are a first-time user of PKI Services web page interfaces. A web browser must have the root CA of the web server, otherwise, the user cannot access the PKI page.

As of z/OS V2R4, the PKI Services component of z/OS uses the HTTPS protocol for its web page interfaces. In previous releases, PKI Services presented the first web page by using the HTTP protocol. Doing so allowed the user to use the link on this page to download the certificate authority (CA) root certificate of the web server certificate into the web browser. Thereafter, all the subsequent web pages used HTTPS.

To help ensure strong security, web pages should use HTTPS rather than HTTP. To use HTTPS, the CA root certificate of the web server must be preinstalled in the user's web browser.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

Starting with z/OS V2R4, the PKI page that previously used HTTP is now updated to use HTTPS. The link that is used to deliver the CA root certificate on the first page is removed. Therefore, you must use an alternative method to distribute the CA root certificate to the appropriate users at your installation.

Upgrade action: For any web browser that is planned to be used to access the PKI web page interface for the first time, your PKI administrator must distribute the CA root certificate of the web server (HTTP server, WebSphere, or WebSphere Liberty) to users who require access to the PKI web page interfaces. A possible method is to distribute the CA root certificate is by using the same communication channel that you use to provide users with the URI of the PKI web page. For example, if you send email, you can

1. Add the CA root certificate as an attachment or include it as Base64 encoded text in the first email.
2. Send a separate email with the root certificate fingerprint to help users ensure that the correct CA certificate was received in the previous email.

Refer to *PKI Services Guide and Reference* topic "Steps for accessing the user web pages" for more options to distribute the CA root certificate.

Infoprint Central requires SSL connections by default (Required, as of V2.4)

Required if you are running Infoprint Central

Starting with z/OS V2R4, SSL connection is required by default for the Infoprint Central web browser GUI. In previous releases, the GUI worked with or without SSL.

This change requires your installation to specify enabled SSL for your IBM HTTP Server (IHS) powered by Apache 31-bit configuration. Doing so causes the **httpd.conf.updates** file to be updated with a rewrite directive that converts the HTTP links to HTTPS links. This action saves users from having to update their Infoprint Central bookmarks.

For users who do not want to change their IHS configurations to use SSL, an Infoprint Server configuration attribute and new ISPF field are provided. The new option has an attribute name of `use-unencrypted-connection` and an ISPF panel field name of "Use unencrypted connection". Note that this connection type is less secure than HTTPS.

You are encouraged to set up SSL before upgrading to z/OS V2R4.

Upgrade action: If SSL is not enabled in IHS and you are using Infoprint Central, follow these instructions. To use an SSL connection for Infoprint Central, do the following:

1. Follow the instructions to enable SSL in *IBM HTTP Server on z/OS Migrating from Domino-powered to Apache-powered*, which is available online:<http://www.redbooks.ibm.com/redpapers/pdfs/redp4987.pdf>.
2. Uncomment the Rewrite directive in the `httpd.conf.updates` file provided by Infoprint Server before copying the Infoprint Central directives to the `httpd.conf` file.
3. Test an Infoprint Central link or bookmark to ensure that Infoprint Central loads in the browser.

If you need to use an unencrypted connection, do the following:

1. Use the Infoprint Server System Configuration ISPF panel or PIDU to set the `use-unencrypted-connection` attribute to yes in the printer inventory.
2. Test an Infoprint Central link or bookmark to ensure that Infoprint Central loads in the browser.

RMF: Configure AT-TLS to enable secure communication with the RMF distributed data server (Required, as of V2.4)

Required if you rely on the current default of HTTPS(NO).

With RMF V2R4, the GPMSRVxx parmlib member option HTTPS(ATTLS) is specified by default. As a result, the RMF distributed data server (DDS) ensures that incoming connections are secured by AT-TLS. If an incoming connection is not secured by AT-TLS, the connection is refused.

Upgrade action: To allow insecure communication for the DDS, you can specify the option HTTPS(NO) in the GPMSRVxx parmlib member. However, this setting is not recommended because it allows the DDS to accept insecure connections.

Before you can configure the DDS, you must enable the Policy Agent for AT-TLS. Information about how to set up AT-TLS communication is provided in *z/OS Communications Server: IP Configuration Guide* and *z/OS Security Server RACF Security Administrator's Guide*. For other security management products, refer to the corresponding security product documentation.

The following example shows a rule that enables secure communication with the DDS:

```
# RMF Distributed Data Server Rule TTLSRule DDSServerRule { LocalPortRange 8803
Jobname GPMSERVE Direction Inbound Priority 1 TTLSGroupActionRef DDSServerGRP
```

```
TTLSEnvironmentActionRef DDSServerENV } TTLSGroupAction DDSServerGRP { TTLSEnabled On
Trace 1 } TTLSEnvironmentAction DDSServerENV { HandshakeRole Server TTLSEnvironmentAction
{ Keyring DDSServerKeyring } TTLSEnvironmentAdvancedParms { ServerCertificateLabel
RMFDDS } }
```

The example rule is described, as follows:

TTLSEnvironmentAction: Jobname

The name value specifies the job name of the application. GPMSEVE is the job name of the DDS

TTLSEnvironmentAction: LocalPortRange

The local port the application is bound to for this rule's action to be performed. 8803 is the default HTTP Port of the DDS.

TTLSEnvironmentAction: Direction

Specifies the direction from which the connection must be initiated for this rule's action to be performed. In this example, Inbound is specified, which means that the rule applies to connection requests that arrive inbound to the local host.

TTLSEnvironmentAction: Priority

An integer value in the range 1 -2000000000 represents the priority that is associated with the rule. The highest priority value is 2000000000. If you use multiple rules for the DDS server, the more specific a rule is, the higher its priority should be. Generic rules without detail specifications of the incoming connections should have a low priority.

TTLSEnvironmentAction: HandshakeRole

Specifies the SSL handshake role to be taken for connections in this AT-TLS environment. In this example, Server is specified which means that the SSL hand shake is performed as a server.

TTLSEnvironmentAction: Keyring

Specifies the path and file name of the key database z/OS UNIX file, the ring name of the SAF key ring, or the name of the z/OS PKCS #11 token. In this example, the RACF key ring DDSServerKeyring is specified.

TTLSEnvironmentAction: ServerCertificateLabel

Specifies the label of the certificate for a server application to authenticate the server. In this example, the DDS Server certificate with the label RMFDDS is used.

Upgrade Actions for z/OS V2.3



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.

19

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

© 2020 IBM Corporation



Upgrade actions for using ICSF as of z/OS V2.3

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.

ICSF Upgrade Actions for z/OS V2.4



Upgrade Actions you can do NOW

Review your CSFPARM DD statement in your ICSF procedure

(Required, as of V2.4 – HCR77D0)

- As of ICSF FMID HCR77D0, the ICSF procedure will not accept a CSFPARM DD that is a sequential data set.
 - If you have this:

```
//ICSF PROC
//ICSF EXEC PGM=CSFINIT,TIME=NOLIMIT,MEMLIMIT=NOLIMIT
//CSFPARM DD DSN=SYS1.CSFPARM,DISP=SHR
```

- Use a partitioned data set instead, such as this:

```
//ICSF PROC
//ICSF EXEC PGM=CSFINIT,TIME=NOLIMIT,MEMLIMIT=NOLIMIT
//CSFPARM DD DSN=USER.PARMLIB(CSFPRM00),DISP=SHR
```

- SYS1.SAMPLIB(CSF), can be used to model after.

20

© 2020 IBM Corporation

ICSF Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

ICSF Upgrade Actions You Can Do Now

Plan for the removal of sequential data sets from the CSFPARM DD statement (Required, as of V2.4)

Not required, but recommended because this change will be required in a future release.

In an upcoming ICSF release, sequential data sets will no longer be accepted on the CSFPARM DD statement in the ICSF procedure.

Upgrade action:

Check for the specification of any sequential data sets specified on the CSFPARM DD statement.

Consider using a partitioned data set in place of a sequential data set on the CSFPARM DD statement, for example:

```
//ICSF PROC PRM=XX
//ICSF EXEC PGM=CSFINIT,TIME=NOLIMIT,MEMLIMIT=NOLIMIT
//CSFPARM DD DSN=SYS1.ICSFLLIB(CSFPRM&PRM),DISP=SHR
```

You can find a sample ICSF procedure in SYS1.SAMPLIB(CSF), which is:

```
//* Licensed Materials - Property of IBM
//* 5650-ZOS Copyright IBM Corp. 2009, 2018
//CSF PROC
//CSF EXEC PGM=CSFINIT,REGION=0M,TIME=1440,MEMLIMIT=NOLIMIT
/** When using CSFPARM DD, the installation options data set must be
/** a partitioned data set on systems running HCR77D0 or later.
//CSFPARM DD DSN=USER.PARMLIB(CSFPRM00),DISP=SHR
```

z/OSMF Upgrade Actions for z/OS V2.4



Upgrade Actions you can do NOW:

Prepare for the removal of the policy data import function of Network Configuration Assistant (Recommended, as of V2.4)

- z/OS V2.4 is the last release in which the Network Configuration Assistant (NCA) plug-in of z/OSMF supports the policy data import function.
 - This function allows the user to import existing Policy Agent configuration files into the Network Configuration Assistant.
- After z/OS V2.4, it will not be possible to import policy configuration files for AT-TLS, IPSec, PBR, and IDS technologies.
- Import of TCP/IP profiles into Network Configuration Assistant is not affected.



Easy upgrade for z/OSMF from V2.3!


21

© 2020 IBM Corporation

z/OSMF Upgrade Actions for z/OS V2.3



Upgrade 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 V2.4 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 with automation or 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



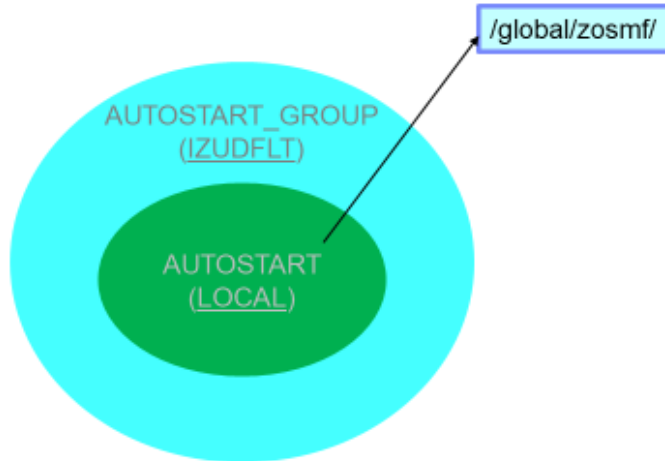
Sage advice: start z/OSMF today on z/OS V2.2, and there are far fewer upgrade actions for z/OS V2.3+!

22

© 2020 IBM Corporation

z/OSMF Upgrade Actions for z/OS V2.3

Autostart Use Case #1: Monoplex (single system)



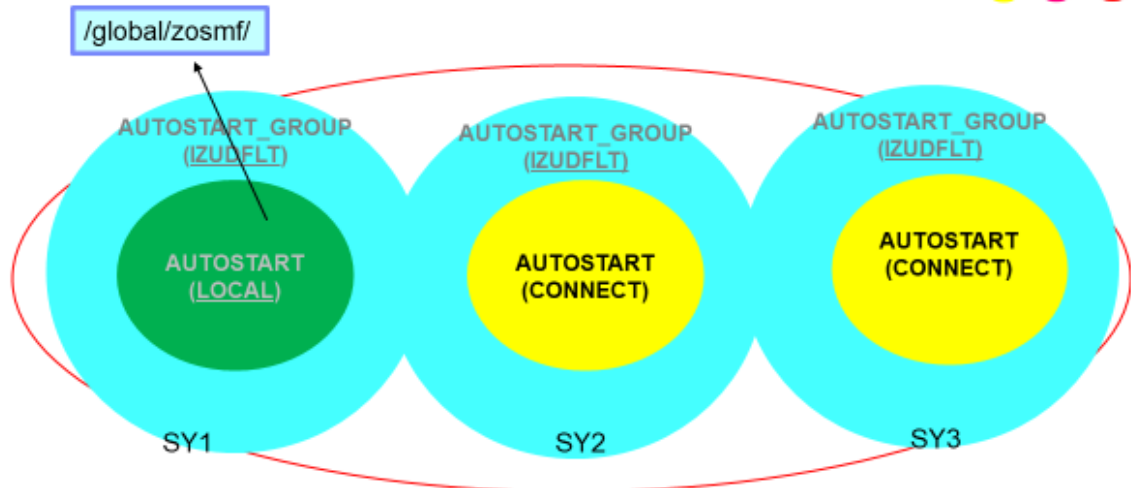
z/OSMF server started in group IZUDFLT: all defaults taken.

23

© 2020 IBM Corporation

z/OSMF Upgrade Actions for z/OS V2.3

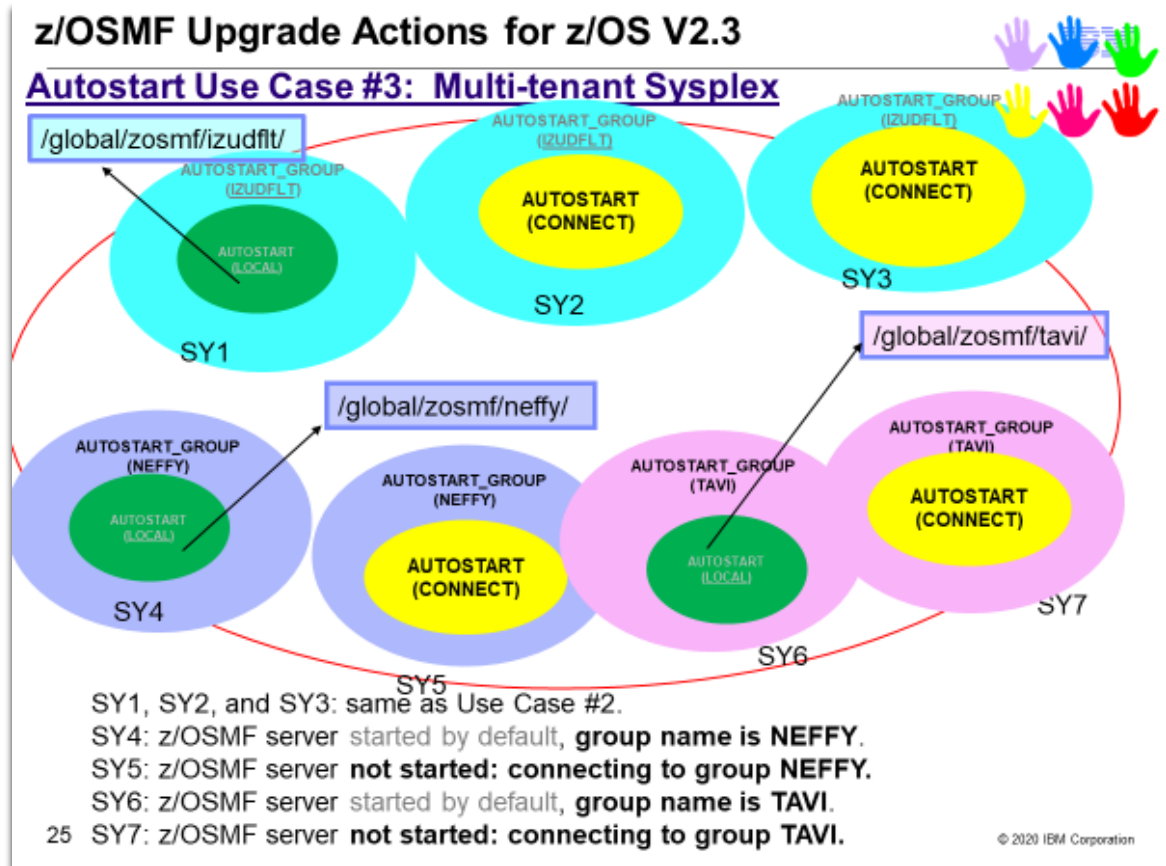
Autostart Use Case #2: Sysplex



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.

24

© 2020 IBM Corporation



z/OSMF Upgrade Actions for z/OS V2.3



Assuming you took the **Sage advice** and are running z/OSMF on V2.2 today...



Upgrade Actions you MUST do Before First IPL:

- **Perform security customization (Required in V2R3)**
 - A handful of additional security profiles to add: use IZUSEC sample to see the V2.3 additions.
 - All the security customization done prior to z/OS V2R3 is still usable and appropriate for z/OS V2R3!
- **Automation or AUTOSTART? You decide on one. (Required-IF in V2R3)**
 - **If automation: set all to CONNECT (with proper group)**
 - **If AUTOSTART: then set one LOCAL in the group, and the rest to CONNECT. Ensure automation doesn't try to start it a second time, or the commands will fail.**



Failure to follow sage advice, means that you need to configure z/OSMF before the IPL of z/OS V2.3 or V2.4 if you want it to start. Much more work!

z/OSMF Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

z/OSMF Upgrade Actions You Can Do Now

Prepare for the removal of the policy data import function of Network Configuration Assistant (Recommended, as of V2.4)

Recommended because importing Policy Agent configuration files into Network Configuration Assistant will not be supported in a future release.

z/OS V2R4 is the last release in which the Network Configuration Assistant (NCA) plug-in of z/OSMF supports the policy data import function. This function allows the user to import existing Policy Agent configuration files into the Network Configuration Assistant.

After z/OS V2R4, it will not be possible to import policy configuration files for AT-TLS, IPsec, PBR, and IDS technologies.

Import of TCP/IP profiles into Network Configuration Assistant is not affected.

Upgrade action: If you plan to import Policy Agent configuration files into the Network Configuration Assistant, perform this work in the current release of z/OS (V2.2, V2.3, or V2.4). This function will not be supported in a future release.

Otherwise, you have no action to take.



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 upgrade actions that affect all z/OS installations. The upgrade 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 upgrade 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.

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 upgrade 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.

Upgrade 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 upgrade 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

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

described in this section. The upgrade 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 upgrade 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 upgrade 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 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.

Upgrade action:

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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.

Upgrade 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 upgrade 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).

Upgrading to z/OS V2.4 Part 2 of 2: Technical 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	< <i>saf-prefix</i> >.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 upgrade 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.

Upgrade action:

If your installation does not start z/OSMF automatically, you have no upgrade 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.

Upgrade 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

Upgrading to z/OS V2.4 Part 2 of 2: Technical 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. *This setting is required if you are going to use JES2 Email Delivery Service. Do not confuse it with the AUTOSTART setting in IZUPRMxx!*
 - 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, and choose to use your automation: 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

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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:

- 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.



Upgrade 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 Upgrade Actions for z/OS V2.4



Upgrade Actions Before Installing:

- **Enable the SDSF class for RACLIST processing (Req-IF, as of V2R3)**
 -  • If you activated 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. To connect to the SDSF server, the user must have READ access to the ISF.CONNECT.system resource in the SDSF class.
 - 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.

Upgrade 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.

27



No new V2.4 upgrade actions from SDSF!

© 2020 IBM Corporation

SDSF Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

SDSF 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.

Upgrade 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.

Upgrade 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.

Upgrade 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)

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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

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.

Upgrade 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 Upgrade Actions for z/OS V2.4



Upgrade Actions Before First IPL:

- **Ensure that the ECC master key is activated in the CCA coprocessor (Req-IF, as of V2.4)**
 - The RACDCERT command is used to manage RACF digital credentials. When you use the RSA(PKDS) keyword for an RSA private key, RACF stores the RSA private key in the ICSF PKDS.
 - Prior to V2.4, the RSA private key was stored in PKDS under a 24-byte master key called the RSA master key.
 - As of V2.4, the RSA private key is stored in PKDS under a more secure master key that is called the ECC master key (32 bytes).
 - The ECC master key must be activated in the CCA coprocessor.



Upgrade Actions After First IPL:

- **RACF only supplies three certificate-authority (CA) certificates out of the current list (Req-IF, as of V2.3)**
 - Prior to V2.3, RACF supplied 26 CA certificates.
 - At V2.3, 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
 - You no longer need to continually remove the unused CA certificates. Clean them out one last time, after all V2.3. (RACF will not delete them.)



28

© 2020 IBM Corporation

RACF Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

RACF Action Pre-First IPL

Ensure that the ECC master key is activated in the CCA coprocessor

Required if your installation uses the RACF RACDCERT command with the RSA(PKDS) keyword and you did not activate the ECC master key in the CCA coprocessor.

In RACF, the RACDCERT command is used to manage RACF digital certificates. When you specify RACDCERT with the RSA(PKDS) keyword for an RSA private key, RACF stores the RSA private key in the ICSF PKA key data set (PKDS).

Starting in z/OS V2R4, the RSA private key is stored in PKDS under a more secure master key that is called the ECC master key (32 bytes). In previous releases, the RSA private key was stored in PKDS under a 24-byte master key called the RSA master key. With this change, you must ensure that the ECC master key is activated in the CCA coprocessor. Otherwise, the RACDCERT command that attempting to store an RSA key in PKDS fails with an error. This change is intended to help maintain strong security in your enterprise. It also supports the generation of a new signature algorithm on certificates that are used for TLS1.3, which is introduced in z/OS V2R4.

To detect an active ECC master key and the availability of a coprocessor CCA-5.3 or above, use health check IBMICSF,ICSF_PKCS_PSS_SUPPORT. Based on this status, the health check indicates whether PKCS-PSS algorithms can be used under current conditions. This health check is available with APAR OA56837.

Upgrade action:

Look for uses of the RACDCERT command specified with the RSA(PKDS) keyword. For example:

```
RACDCERT GENCERT...RSA (PKDS (...))
RACDCERT REKEY...RSA (PKDS (...))
RACDCERT ADD ...PKDS (...)
```

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

If so, verify that the ECC master key is activated before these commands are used on a z/OS V2R4 system. You can use either of the following approaches:

- Open ICSF panel option 1 and check the CCA coprocessor ECC master keys status.
- Run the ICSF ECC master key health check.

Failure to do this upgrade action will result in RACDCERT command failures with reason code x'00002B08' and the following message:

```
IRRRD117I Unexpected ICSF CSNDPKG return code x'00000008 ' and reason code x'00002B08'. The request is not processed.
```

RACF Action Post-First IPL

Accommodate the removal of certain CA certificates from RACF (Required, as of V2.3)

Required if you depend on certain CA certificates that RACF added at each system IPL.

Prior to z/OS V2.3 RACF supplied 26 certificate-authority(CA) certificates. With APAR OA54748 on z/OS V2.2 and V2.3, an additional CA certificate has been added. 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 upgrade actions when moving to z/OS V2.3 from a previous release.

Upgrade Action

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. Once you are fully at z/OS V2.3, the 23 removed certificates will no longer be added (nothing will be removed by RACF). You may clean out any or all of the 23 removed certificates if you wish. They will no longer be added once you are fully at V2.3.
2. If you want the previous shipped CA certificates in the existing RACF database, no action is needed as they are still there. 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 Upgrade Actions for z/OS V2.4



Upgrade Actions After First IPL:

- **Accommodate the OpenSSH ported level (Required-IF, as of V2.4)**



- z/OS V2.3 and V2.2 were open source version OpenSSH level 6.4p1.
- z/OS V2.4 contains open source version OpenSSH level 7.6p1.



- Several differences in the ported levels, which may cause upgrade actions. Some functions no longer available are:
 - SSH Version 1 protocol (also called SSH-1).
 - Running without privilege separation for sshd (SSH Daemon).
 - Support for legacy V00 OpenSSH certificate format.
 - Support for pre-authentication compression by sshd.
 - Accepting RSA keys smaller than 1024 bits (RSA1).
 - ...
- Several features will no longer be enabled by default:
 - 1024-bit Diffie Hellman key exchange.
 - DSA (ssh-dss, ssh-dss-cert-*) host and user keys.
 - MD5-based and truncated MD5 and SHA1 HMAC algorithms.
 - ...
- Read of all changes in your handout, or the *z/OS V2.4 Upgrade Workflow*.

29

© 2020 IBM Corporation

z/OS OpenSSH Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

z/OS OpenSSH Upgrade Action Post-First-IPL



Accommodate the OpenSSH ported level (Required-IF, as of V2.4)

Required if you are reliant upon any of the changes in the newer ported level..

With z/OS V2.4, OpenSSH is upgraded to include a new level of open source version OpenSSH 7.6p1. (The last change was in z/OS V2.2 for open source version OpenSSH 6.4p1. With z/OS OpenSSH V2R4, significant new features are included:

- Elliptic-curve DSA keys are now supported in Key Rings and FIPS.
- Support new key algorithms: ssh-ed25519, ssh-ed25519-cert-v01@openssh.com
- Support new key exchange algorithms: diffie-hellman-group14-sha256, diffie-hellman group16-sha512, diffie-hellman-group18-sha512, curve25519-sha256 and curve25519-sha256@libssh.org.
- Support new cipher algorithms: chacha20-poly1305@openssh.com
- The SMF Type 119 subtype 94 and 95 (ssh / sshd connection started) records will include a section that identifies the IP addresses and ports for the connection.
- A new command ssh-proxyc is added, which can be used by the ssh client to connect through SOCKS5 proxy servers.

Also with z/OS OpenSSH V2R4, following features are no longer available (since they are deprecated by the Open Source community in OpenSSH 7.6p1). This had been previously announced as a Statement of Direction in the

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

4Q2017 z/OS V2.3 Enhancement Announcement (<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=g pateam&supplier=897&letternum=ENUS217-536>):

- SSH Version 1 protocol (also referred to as SSH-1).

These options from **ssh_config** have been removed: Cipher, CompressionLevel, RhostsAuthentication, RhostsRSAAuthentication, RSAAuthentication.

These options from **sshd_config** have been removed: KeyRegenerationInterval, RhostsAuthentication, RhostsRSAAuthentication, RSAAuthentication, ServerKeyBits, UseLogin and PAMAuthenticationViaKbdInt.

- Running without privilege separation for sshd (SSH Daemon).
- Support for the legacy v00 OpenSSH certificate format.
- Support for pre-authentication compression by sshd (SSH Daemon). SSH clients will either need to support delayed compression mode or otherwise compression will not be negotiated.
- Support for Blowfish and RC4 ciphers and the RIPE-MD160 HMAC (Hash Message Authentication Code), specifically: blowfish-cbc, cast128-cbc, arcfour, arcfour128, arcfour256, hmac-ripemd160, hmac-ripemd160@openssh.com, and hmac-ripemd160-etm@openssh.com.
- Accepting RSA keys smaller than 1024 bits. (RSA1)

With z/OS OpenSSH V2R4, the following features will no longer be enabled by default. This had been previously announced as a Statement of Direction in the 4Q2017 z/OS V2.3 Enhancement Announcement (<https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=g pateam&supplier=897&letternum=ENUS217-536>):

- Support for the 1024-bit Diffie Hellman key exchange, specifically diffie-hellman-group1-sha1
- Support for DSA (ssh-dss, ssh-dss-cert-*) host and user keys
- Support for MD5-based and truncated MD5 and SHA1 HMAC algorithms, specifically: hmac-md5, hmac-md5-96@openssh.com, hmac-sha1-96, hmac-sha1-96@openssh.com, hmac-md5-etm@openssh.com, hmac-md5-96-etm@openssh.com, hmac-sha1-96-etm@openssh.com,
- Support for the Triple DES cipher and cbc cipher, specifically 3des-cbc, aes128-cbc, aes192-cbc and aes256-cbc.

z/OS UNIX Upgrade Actions for z/OS V2.4

Upgrade Actions Post-First IPL:

- **Remove FORKCOPY(COW) and KERNELSTACKS statements from BPXPRMxx (Recommended, as of V2.4)**
 - FORKCOPY (COW) – copy-on-write – will be ignored, and FORKCOPY (COPY) will always be used.
 - COW was the default prior to V2.4.
 - KERNELSTACKS (BELOW) will be ignored, and KERNELSTACKS (ABOVE) will always be used.
 - BELOW was the default.

Upgrade 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.
 - **Using it for z/OSMF is recommended!**
- **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.**

30

© 2020 IBM Corporation

z/OS UNIX Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

z/OS UNIX Actions Pre-First IPL

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

Required if you wish to avoid additional customization work, and follow some elements default customization.

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 .

Upgrade 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.



Optionally delete the FORKCOPY(COW) option of BPXPRMxx (Recommended, as of V2.4)

Not required, but recommended for clarity in the BPXPRMxx parmlib member because FORKCOPY(COPY) will always be used even if FORKCOPY(COW) is specified.

Previously, with the FORKCOPY option in the BPXPRMxx parmlib member, copy-on-write (COW) mode could be specified for fork processing. The default was FORKCOPY(COW). In z/OS V2R4, the COW option is disabled. FORKCOPY(COPY) will always be used even if FORKCOPY(COW) is specified.

Upgrade action:

Determine whether your BPXPRMxx parmlib member specifies FORKCOPY(COW) mode. If so, remove it. FORKCOPY(COPY) is always used, even when FORKCOPY(COW) is specified.

Use of FORKCOPY(COPY) avoids any additional ESQA use in support of fork. Use of FORKCOPY(COW) caused the system to use the ESQA to manage page sharing.



Optionally delete the KERNELSTACKS option of BPXPRMxx (Recommended, as of V2.4)

Not required, but recommended for clarity in the BPXPRMxx parmlib member because KERNELSTACKS(BELOW) will not be used.

Previously, with the KERNELSTACK option in the BPXPRMxx parmlib member, stacks could be allocated either above or below the bar. The default was KERNELSTACKS(BELOW). As of z/OS V2R4, stacks are always allocated above the bar. If KERNELSTACKS(BELOW) is specified, it is ignored and the stacks are allocated above the bar.

Upgrade action:

Determine whether your BPXPRMxx parmlib member uses the KERNELSTACK option. If so, remove it. The stacks will always be allocated above the bar.

When the kernel stacks were allocated from kernel private storage that is below the bar, the number of threads running in the kernel was limited to about 30,000. KERNELSTACKS(ABOVE) increases thread limit to a maximum of 500,000; the actual amount will vary, depending on work load and system configuration. Usage of real storage in the kernel address space is also increased.

JES2 Upgrade Actions for z/OS V2.4



Upgrade Actions Before First IPL:

Accommodate the changed default for the EXECUTABLE keyword on \$GETMAIN (Required-IF, as of V2.4)

- V2.3 added the EXECUTABLE= keyword to \$GETMAIN to indicate whether obtained storage is marked executable on systems running on z14. The default was YES.
- V2.4 changes the default to NO. Meaning, any storage obtained when running that system on z14 will not support executable code. If this code attempts to run, it will result in 0C4 abend.
- Review your use of \$GETMAIN. Particularly check for the DCB exit stubs that might be obtained and pointed to by the EXLST area.

Accommodate changes in NJE input phase processing for multi-object job streams (Required-IF, as of V2.4)

- Before V2.4, if an NJE job stream contained more than one job or job group, the stream and all jobs within it would fail input phase processing.
- As of V2.4, JES2 supports multiple job and job group objects in NJE job transmission stream.
- Subtle differences may be seen input phase processing.
 - Multiple jobs and job groups can be marked busy on a job receiver at the same time.
 - Final processing for the jobs is delayed until the job trailer is received.

31 ---

© 2020 IBM Corporation

JES2 Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

z/OS UNIX Actions Pre-First IPL

Accommodate the changed default for the EXECUTABLE keyword on \$GETMAIN (Required-IF, as of V2.4)

Required if you have any affected JES2 exit routines or user modifications.

z/OS V2R3 added the EXECUTABLE= keyword on the \$GETMAIN macro to indicate whether the obtained storage is marked as executable for systems that run on an IBM z14 server or later. For more information about non-executable memory, see the description of the EXECUTABLE= keyword on the STORAGE macro.

When the EXECUTABLE= keyword was added to the \$GETMAIN macro in z/OS V2R3, the default was YES. Starting with z/OS V2R4, the default on \$GETMAIN is changed to NO. This change implies that, by default, any storage that is obtained in a supported subpool does not support executable code on a z14 server. Any attempt to run code in the storage results in an 0C4 abend.

Upgrade action:

Review your exits for instances of the \$GETMAIN macro that obtains storage for executable code. In particular, check for DCB exit stubs that might be obtained and pointed to by the EXLST area. When appropriate to do so, convert the \$GETMAIN and its associated \$FREMAIN to specify EXECUTABLE=YES. On a z/OS V2R3 system, you can perform this change before you upgrade to z/OS V2R4.

If in doubt, you can change all \$GETMAIN and \$FREMAIN macros to specify EXECUTABLE=YES.

Checkpoint versions are moving to 64-bit storage (Required-IF, as of V2.4)

Required if your JES2 exit routines or applications use the IAZDSERV macro with an SSI call 71 or the \$DSERV macro.

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

Applications and exits that do not run in the JES2 address space can access checkpoint version data from the IAZDSERV data area, which provides a stable copy of the data. To access the IAZDSERV data area, applications can use subsystem interface call (SSI) 71, subfunction SSJIFOBT, and JES2 exits can use the \$DSERV macro.

Before z/OS V2R4, the data returned by these services was in a 31-bit storage data space. Starting in z/OS V2R4, the returned data might reside in 64-bit storage.

z/OS V2R4 adds version 10 of the IAZDSERV data area, which supports 64-bit pointers and ALETs for use in accessing the checkpoint version data. Version 10 is the only IAZDSERV version that is supported by z/OS V2R4.

Upgrade action:

Check for applications and exits that access checkpoint version data from the IAZDSERV data area. Determine whether this code must access the checkpoint data blocks directly.

As an alternative to using the IAZDSERV data area directly, your code can use an SSI call to access checkpoint version data. IBM recommends that you evaluate the use of SSI calls, such as SSI 80 (extended status) or SSI 82 (JES property SSI). Determine whether you can use these services to obtain the data that you require. Otherwise, you must upgrade your code to accept the 64-bit data pointers that are returned in the version 10 IAZDSERV data area.

If you use the \$DSERV macro in an exit, be aware that the IAZDSERV data area it returns is at the version 10 level and therefore contains 64-bit pointers. All JES2 services that use the \$DSERV macro as input are upgraded in z/OS V2R4 to accept the version 10 of the IAZDSERV data area.

In general, unless your code must reference individual fields in the \$DSERV mapping, it should not require changes. However, if your code references fields in the \$DSERV, you must upgrade your code to use the 64-bit fields and run in 64-bit addressing mode.

Accommodate changes in the NJE input phase processing for multi-object job streams

(Required-IF, as of V2.4)

Required if you have JES2 exit routines or processes that are impacted.


Before z/OS V2R4, if an NJE job stream contained more than one job or job group, the stream and all jobs within it would fail input phase processing. Starting with z/OS V2R4, JES2 now supports multiple job and job group objects in an NJE job transmission stream. This is done by keeping all the jobs in the stream busy in the INPUT phase of processing until the job trailer (end of the job stream) is received. When received, the jobs complete input phase processing and are queued to the next phase. If an error occurs on the connection and the job trailer is not received, all of the jobs are purged from the receiving system. When the NJE connection is reestablished, the entire job stream is sent again.

This change implies a number of subtle changes in how input phase processing works:

- Multiple jobs and job groups can be marked busy on a job receiver at the same time
- Final processing for the jobs is delayed until the job trailer is received. This includes exits 20 and 50 (end of input exits) and exit 51 (queue change). As a result, the end of input exit for the first job in the stream is not given control until all other input exits (job statement, accounting string, JCL statement) are called for all of the jobs. The environment in which the end of input exit is called is the same as in prior releases, however, the order is changed.
- When a connection drops, because multiple jobs can exist on the NJE job receiver, all of these jobs must be purged. Prior releases would only have at most one job that needed to be purged.

Upgrade action: Review these changes to NJE input phase processing and evaluate suitable changes to your system. For example, if you have any exits or procedures that rely on only one job or job group being active on a particular NJE job receiver at a time, review and update those exits and procedures. Similarly, if you have an end of input exit or a queue change exit that relies on being called immediately after other input exits for a particular job, review those exits and change them appropriately.

Communications Server Upgrade Actions for z/OS V2.4



Upgrade Actions Before Installing:

IP Services: Decide whether to accept the new FIXED CSM default (Req-IF, as of V2.4)



- As of V2.4, the default for CSM fixed storage for buffers is increased to 512M.
 - Review your HVCOMMON setting in IVTPRM00.
 - Use D NET,CSM to look at the "FIXED MAXIMUM" storage in use, on IVT5538I.



```

IVT5532I -----
IVT5536I TOTAL    ALL SOURCES                23032K    5296K    28328K
IVT5538I FIXED    MAXIMUM =                120M    FIXED    CURRENT = 27165K
IVT5541I FIXED    MAXIMUM USED =           27189K    SINCE LAST DISPLAY CSM
IVT5594I FIXED    MAXIMUM USED =           27189K    SINCE IPL
IVT5539I ECSA     MAXIMUM =                120M    ECSA     CURRENT = 2035K
    
```

• A brief history of this default change: V2.1 → 100M, V2.2 → 200M, V2.4 → 512M.

IP Services: Ensure storage availability for IWQ IPsec traffic (Recommended, as of APAR PI77649)



- The processing of IPAQENET and IPAQENET6 INTERFACE statements is enhanced when you use OSA-Express6S (running in QDIO mode on z14).
- If you enabled QDIO inbound workload queuing (WORKLOADQ) and you have IPsec traffic, an additional ancillary input queue (AIQ) is established for IPsec inbound traffic.
 - Additional storage is allocated for this input queue.
- Each ancillary input queue increases storage utilization in both: ECSA by approx. 36KB, 64-bit CSM HVCOMMON for READSTORAGE.
- Extensive instructions on how to do the calculation are provided in handout, and in the *z/OS Upgrade Workflow*. Use the workflow for some assistance.

32

© 2020 IBM Corporation

Communications Server Upgrade Actions for z/OS V2.4



Upgrade Actions Before Installing:

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.

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.

33

© 2020 IBM Corporation

Communications Server Upgrade Actions For z/OS V2.4

These upgrade actions were taken from *z/OS V2.4 Upgrade Workflow*. Some descriptions and actions have been shortened for inclusion in this presentation. Not all upgrade actions have been included. For the complete descriptions and actions, refer to *z/OS V2.4 Upgrade Workflow*.

[Communications Server Actions You Can Do Now](#)



IP Services: Ensure storage availability for IWQ IPsec traffic (Recommended, as of APAR PI77649)

Not required, but recommended if you have the WORKLOAD parameter that is specified on the OSA IPAQENET and IPAQENET6 INTERFACE statements, you have IPsec traffic and you have concerns about using additional ECSA or real (fixed) storage.

As of z/OS V2R3 with TCP/IP APAR PI77649, or z/OS V2R2 with TCP/IP APAR PI77649 and SNA APAR OA52275, the processing of IPAQENET and IPAQENET6 INTERFACE statements is enhanced when you use OSA-Express6S. If you enabled QDIO inbound workload queuing (WORKLOADQ) and you have IPsec traffic, an additional ancillary input queue (AIQ) is established for IPsec inbound traffic. Additional storage is allocated for this input queue.

Each AIQ increases storage utilization in the following two areas:

- Approximately 36 KB of fixed ECSA
- 64-bit CSM HVCOMMON for READSTORAGE

If you are using IPsec, when the first IPsec tunnel is activated (protocols ESP or AH), the new AIQ is backed with 64-bit CSM HVCOMMON fixed storage. The amount of HVCOMMON storage that is used is based on the specification of the INTERFACE READSTORAGE parameter.

If you configured QDIO inbound workload queuing (WORKLOADQ), ensure that sufficient fixed ECSA and fixed (real) 4 KB CSM HVCOMMON storage is available for the AIQ for IPsec traffic.

This upgrade action concerns OSA-Express6S Ethernet features or later in QDIO mode running on IBM z14 and ZR1. To determine whether sufficient fixed storage is available for IWQ IPsec enabling, use the following IBM health checks:

- IBMCS,ZOSMIGV2R4PREV_CS_IWQSC_tcpipstackname issues a message if the TCP/IP stack has inbound workload queuing (IWQ) and IPsec enabled, but the OSA does not support IWQ for IPsec. Only OSA-Express6S and later support IWQ for IPsec. This health check is shipped INACTIVE and set to run once.
- IBMCS,CSTCP_IWQ_IPSEC_tcpipstackname issues a message if the TCP/IP stack has IWQ and IPsec enabled, and the OSA does support IWQ for IPsec. This health check is shipped ACTIVE and is set to run once.

These health checks are available with PH11837 and OA57525 for V2R3, and PH12005 and OA57560 for V2R4.

Upgrade action:

1. If you are using or plan to use OSA-Express6S or later, verify that the following conditions are true:
 - a. WORKLOADQ is specified on the IPAQENET and IPAQENET6 INTERFACE statements.
 - b. Have IPsec traffic; protocols ESP or AH.
2. If step 1 is applicable, IWQ IPsec uses additional storage. Continue with step 3 - 6. Otherwise, there is no increase in storage usage, and no further action is required.
3. To calculate the total storage increase, count the total number of IPAQENET and IPAQENET6 INTERFACE statements that are coded with the WORKLOADQ parameter that are associated with OSA-Express6S or later. Make a note of the number.
4. Verify that sufficient ECSA is available. To calculate this, multiply the total INTERFACE statements that are counted in step 3 by 36 KB. The resulting number indicates how much additional ECSA is required. To determine whether sufficient ECSA is available to enable this function, verify the following ECSA definitions:
 - a. D CSM usage (PARMLIB member IVTPRM00, ECSA MAX value)
 - b. VTAM (Start Options CSALIMIT and CSA24)
 - c. TCPIP (GLOBALCONFIG ECSALIMIT statement in the TCPIP PROFILE)
5. Verify that sufficient real (fixed) storage is available. 64-bit fixed (CSM HVCOMMON) storage is used for the IPsec AIQ read buffers. To calculate this value, multiply the total number of INTERFACE statements that are counted in step 3 by your configured READSTORAGE value (for example, 4 MB). You can verify how much

storage is being used for READSTORAGE by using the D NET, TRLE command. The resulting number indicates how much additional fixed (64-bit) storage is required. To verify that the additional fixed storage is not a constraint for your system, take the following actions:

- a. Use the DISPLAY CSM command to verify that sufficient fixed storage is available (CSM FIXED MAXIMUM defined in IVTPRM00).
 - b. Verify the actual amount of real storage available to this z/OS system by using D M=STOR or D M=HIGH.
6. If sufficient ECSA or real storage is not available, increase the available real storage or consider defining some of the OSA-Express6S (or later) INTERFACE statements with the NOWORKLOADQ parameter. If your CSM FIXED MAXIMUM is too low, increase this value in IVTPRM00.



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

V2.4)

Required if you use the default CSM FIXED MAX value of 200M and you do not want to use the new default of 512M.

In z/OS V2R4, the default amount for communications storage manager (CSM) fixed storage for buffers is increased from 200 MB to 512 MB. Your installation can specify a value for the CSM fixed storage amount on the FIXED statement in the IVTPRM00 parmlib member.

Upgrade action: Review your HVCOMMON setting in IVTPRM00 and determine whether you need to increase this value to account for the fact that the system reserves a larger portion of this storage by default for CSM buffers. If you did not previously code a value for FIXED in IVTPRM00 and you do not want the new default, specify FIXED MAX(200M) 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.

A brief history of FIXED CSM defaults, for those interested: V2.1 – 100M, V2.2 – 200M, V2.4 – 512M.

IP Services: Determine the storage impact if QDIOSTG=126 is in effect (Required, as of V2.4)

If you specify QDIOSTG=126 in your VTAM start options, each OSA-Express QDIO interface that uses the default READSTORAGE setting of GLOBAL on the INTERFACE or LINK statement gets 8 MB of fixed CSM for read storage. For any OSA-Express QDIO interfaces that have a bandwidth of at least 10 GbE and use 8 MB of read storage, z/OS V2R4 allocates additional fixed CSM 4K HVCOMMON storage for work element processing.

The system allocates additional fixed CSM HVCOMMON storage for work element processing for each OSA-Express QDIO interface that is using 8 MB of read storage.

Upgrade action: See [Additional fixed storage for OSA interfaces using 8 MB of read storage in z/OS Communications Server: IP Configuration Guide](#) to understand how much additional storage z/OS V2R4 allocates for work element processing.

If you do not want the system to allocate this extra storage for a specific interface, update your INTERFACE or LINK statement to specify a READSTORAGE value other than GLOBAL.

IP Services: EZZ6044I and EZZ6045I descriptor codes are changed (Required-IF, as of V2.4)

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

Upgrade 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 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.

Upgrade 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.

Upgrade 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 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 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.

Upgrade 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.
- 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 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.

Upgrade 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 TTLRule 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 TTLRule statement.
3. Locate the TTLSEnvironmentAdvancedParms statement that is referenced by or contained in the TTLSEnvironmentAction statement.
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 TSSLSPORT 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 TSSLSPORT.

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.

Upgrade 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 (otelnetsd), check your inetd configuration file (/etc/inetd.conf) for invocations of otelnetsd that specify the -a user parameter. If you find such an invocation, your otelnetsd server is using Kerberos version 5 authentication. If not, your otelnetsd 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

Upgrading to z/OS V2.4 Part 2 of 2: Technical Actions

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 upgrade 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_etype = des-cbc-crc,des-cbc-md5
```

```
default_tkt_etype = 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: 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.

My "Big Migs" for Upgrading from V2.2 to V2.4



Upgrade actions in V2.3 you should not overlook:

-  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.

plus...





34

© 2020 IBM Corporation


My "Big Migs" for Upgrading from V2.3 to V2.4



Upgrade actions in V2.4 you should not overlook:

-  1. BCP: Removal of support for user key common areas
-  2. Use Network File System (NFS) instead of DFS/SMB
-  3. Various actions related to HTTP → HTTPS for CIM, PKI Services, RMF, and Infoprint Central.
-  4. OpenSSH higher ported level of 7.6p1.

Future upgrade actions to do now:

-  A. HFS removal planning for release after V2.4 (in 2021).
- B. JES3 removal planning for 2023 release.

35

© 2020 IBM Corporation

Upgrade to z/OS V2.4: Technical Actions Summary



- **General:**
 - New address spaces, new and old data sets, changed checks.
- **BCP:**
 - CLOCKxx ETRMODE and ETRZONE defaults are NO. System logger use of zHyperwrite is disabled by default.
 - IEASYSxx REAL default is 0 (nothing). Logstream size defaults.
- **DFSMS:**
 - SMS ACDS size increase, DSS SHARE is ignored, XRC default changes, CA_RECLAIM(DATACLAS) default.
- **ICSF:**
 - ICSF proc CSFPARM DD cannot be sequential data set
- **z/OSMF:** Prepare for removal of import for Policy Agent configuration files into Network Configuration Assistant.

36

© 2020 IBM Corporation

Upgrade to z/OS V2.4: Technical Actions Summary



- **SDSF:** New main panel.
- **RACF:** ECC master key in activated in the CCA coprocessor. Only three CA Certificates shipped
- **z/OS UNIX:** FORKCOPY and KERNELSTACKS behavior, /global sysplex root
- **JES2:** \$GETMAIN EXECUTABLE default of NO. NJE input phase processing for multi-object job streams changes.
- **Communications Server:** Another FIXED CSM default change, more storage for IWQ IPsec traffic.



37

© 2020 IBM Corporation