Migration considerations for z/OS V1.8!

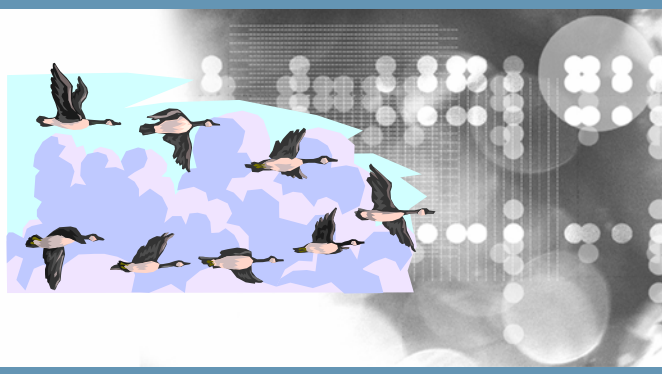

© 2007 IBM Corporation

Trademarks

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

AFP, ADX, AnyNet, BookManager, CICS, DB2, DB2 Universal Database, DFSMSdfp, DFSMShsm, DFSMSrmm, DFSORT, Enclave, ESCON, Eserver, FICON, FICON Express, HiperSockets, IBM, IBMLink, IMS, Infoprint, Language Environment, Multiprise, MVS, MVS/ESA, MVS/SP, NetSpool, NetView, Open Class, OS/390, Parallel Sysplex, PR/SM, Print Services Facility, PrintWay, RACF, Resource Link, RMF, S/390, S/390 Parallel Enterprise Server, Sysplex Timer, SystemPac, System z9, Tivoli, TotalStorage, VisualAge, VSE/ESA, VTAM, WebSphere, zArchitecture, z/OS, z/VM, zSeries

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
UNIX is a registered trademark of The Open Group in the United States and other countries.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

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

Notes:

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

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

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

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

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

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

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

ENABLING BUSINESS. A THROUGH Z.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

IBM

Migration considerations for z/OS V1.8

- **Introduction**
  - Mirko Donie, IBM z/OS Product Development Team Lead
- **Plan your migration to z/OS V1.8**
  - Marna Walle, IBM z/OS System Build and Install
- **Questions and Answers**

© 2007 IBM Corporation

IBM

Plan your migration to z/OS V1.8! Agenda

- **z/OS Coexistence-Migration-Fallback-Service Policy**
  - z/OS Coexistence-Migration-Fallback
- **Ordering z/OS V1R8 (and z/OS V1R7)**
- **Positioning for z/OS V1R8:**
  - Highlights of Removed Functions since z/OS V1R6
  - Ensuring System Requirements are Satisfied
    - Driving and Target System Requirements
  - Coexistence System Requirements
    - EPSPT Programmatic Help with Coexistence Fixes!
  - IBM Migration Checker for z/OS

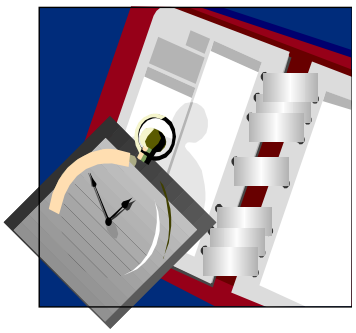

- **Migration Actions for z/OS V1R8:**
  - z/OS General Items
  - BCP
  - XL C/C++
  - Communications Server
  - HCD
  - DFSMS
  - Distributed File Service - zFS
  - JES2
  - Language Environment
  - z/OS UNIX

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Highlights of Removed Functions since z/OS V1R6

JES2 "compatibility" R4 mode (from JES2)	Base Element - to avoid a cold start, you must $ACTIVATE to convert the JES2 checkpoint to z2 mode before installing z/OS R7	as of z/OS R7
JOBCAT and STEPCAT facilities (from DFSMSdfp)	Base Element - any remaining JCL that use JOBCAT or STEPCAT must change	as of z/OS R7
OS/390 R10 level of the C/C++ compilers (from C/C++)	Priced Feature - move to the ISO 1998 Standard level of the compilers (introduced in z/OS R2)	as of z/OS R7
One-byte console IDs on macro interfaces and operator commands (and TRACK, ... commands) (from BCP)	Base Element - Use console names, instead of one-byte IDs. Use Console ID Tracking Facility to identify one-byte ID usage, supplied as of z/OS R4 Consoles Enhancement feature.	as of z/OS R7
Any remaining one-byte console ID support (from BCP)	Base Element - Use console names instead of one-byte console IDs	as of z/OS R8
zFS multi-file system aggregate shared across a sysplex (from Distributed File Service)	Base Element - zFS compatibility mode aggregates (which have a single file system per data set) will continue to be supported in all environments.	as of z/OS R8
Firewall Technologies (from Integrated Security Services)	Base Element - Many Firewall Technologies functions have been been long stabilized and can be replaced w/ Comm Server functions (IPSecurity). Some functions won't have replacements.	as of z/OS R8

= don't overlook!

007 IBM Corporation

3

Functions Planned to be Withdrawn *in the future*

Host communication between HCM and HCD with APPC (from HCM)	Priced Feature - Only TCP/IP for host communication between HCM and HCD will be allowed. (Currently both TCP/IP and APPC are supported.)	**Planned for release in 2007**
In-stack version of TN3270 Server (from Communications Server)	Base Element - use standalone TN3270 Server instead (introduced in z/OS R6)	**Planned for relase in 2007**
APPC Application Suite in Communications Server	Base Element - more full-featured alternative applications exist in modern integrated SNA/IP networks. (APPC itself remains an integral part of SNA functions, and no plans to remove APPC from z/OS.)	**Planned for release in 2007**
Run-time support for applications that use the C/C++ IBM Open Class (IOC) Dynamic Link Libraries (DLLs).	Priced Feature - Any application code that uses the IOC Library should migrate to use the Standard C++ Library	**Planned for release in 2007**
Support for VSAM data sets with IMBED, REPLICATE, or KEYRANGE attributes (from DFSMS)	Base Element - plan to redefine any affected VSAM data sets. Use tool to assist in identifying affected VSAM data sets.	**Future**
zFS multi-file system aggregates (from Distributed File Service)	Base Element - zFS compatibility mode aggregates will still be supported	**Future**

= don't overlook!

© 2007 IBM Corporation

4

© Copyright IBM Corporation, 2007

z/OS Ordering and Deliverables

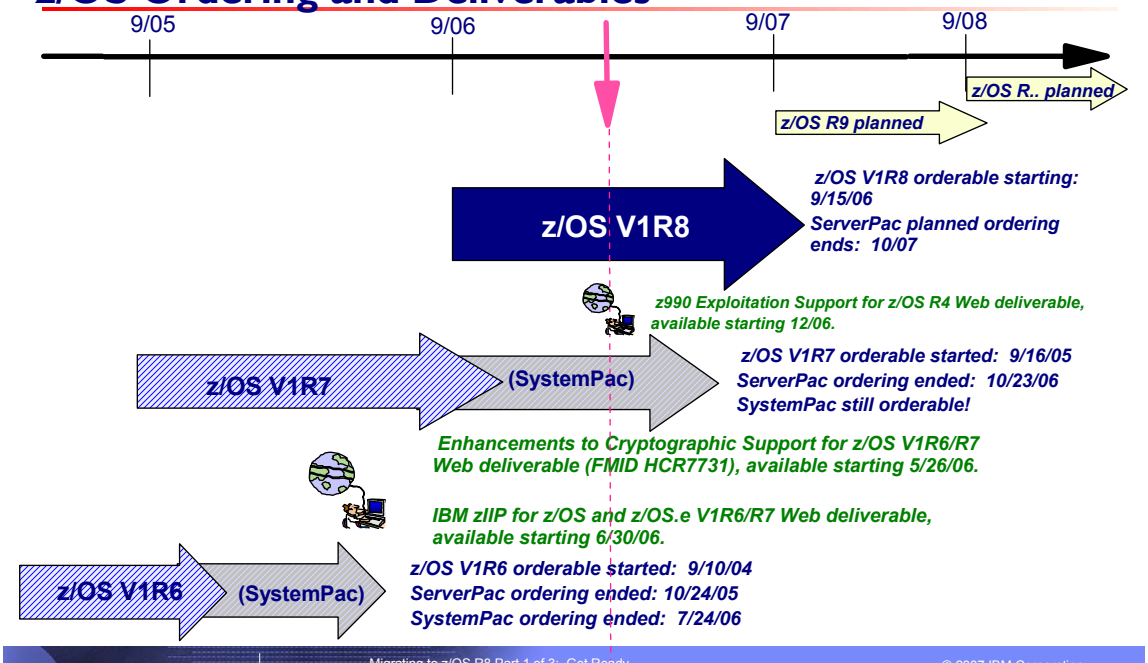

© 2007 IBM Corporation

z/OS V1R8 Coexistence-Migration-Fallback Policy

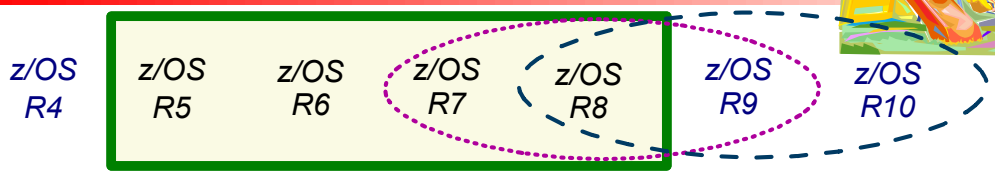

- **Starting with z/OS R6, IBM has aligned the coexistence, fallback, and migration policy with the service policy.**
  - **z/OS R5, z/OS R6, z/OS R7, and z/OS R8** are supported for coexistence, migration, and fallback
  - The lowest release for coexistence, migration, and fallback support on z/OS R8 is z/OS R5.
  - **If you are still on z/OS R4 your migration to z/OS R7 should be well underway!**
- **Only JES2/JES3 that can coexist with the shipped JES can be "staged" on z/OS. This is enforced in z/OS V1 R8. That means:**
  - z/OS R5 JES2 thru z/OS R8 JES2 are supported for coexistence, migration, and fallback
  - z/OS R5 JES3 thru z/OS R8 JES3 are supported for coexistence, migration, and fallback

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

IBM

z/OS Service Policy

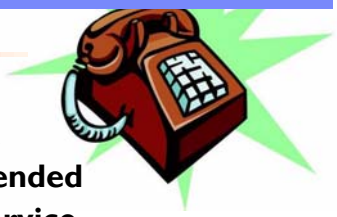

- Release serviceable for three years following GA
- Service on last release of a version might be extended
- At least 12 months notice before withdrawing service
- Handy website:
  http://www.ibm.com/servers/eserver/zseries/zos/support/zos_eos_dates.html

	General Availability	Service Expiration
z/OS V1R4	27 September 2002	Announced to be 31 March 2007
z/OS V1R5	26 March 2004	Announced to be 31 March 2007
z/OS V1R6	24 September 2004	Announced to be 30 September 2007
z/OS V1R7	30 September 2005	Planned to be September 2008
z/OS V1R8	29 September 2006	Planned to be September 2009

© 2007 IBM Corporation
7

IBM

Ensuring System Requirements are Satisfied

- Driving System Requirements to Install z/OS V1R8 ServerPac:
  - Minimally, z/OS R5 with PTFs, with other requirements (OMVS in full-function mode, UID(0) or permitted to BPX.SUPERUSER, ...)
  - Must install PTFs on z/OS V1R8 with latest Program Binder, SMP/E, and HLASM.
  - For Installing ServerPac electronically, also need:
    - SMP/E V3R3 and ICSF configured and active
    - -or- SMP/E V3R4 -and- Java 2 Technology Edition V1R4 (ICSF is not necessary in this case!)
- Target System Requirments for Running z/OS V1R8:
  - z/OS R8 must run in z/Architecture mode (...since z/OS R6!)
  - HW Server: z9 EC, z9 BC, z990, z890, z900, or z800
    - Some functions (like zAAP or zIIP) require z990 or z890 at least.
  - Minimum SW Subsystem Levels (some examples):
    - CICS TS V2 R2
    - DB2 V7 R1 with PTFs (V7 R1 announced end of service as 6/30/08)
    - IMS V8 R1 with PTFs
    - WebSphere Application Server V5 R1
    - See z/OS Planning for Installation for a more complete list.
- Coexistence PTFs
  - Programmatic assistance is now available!

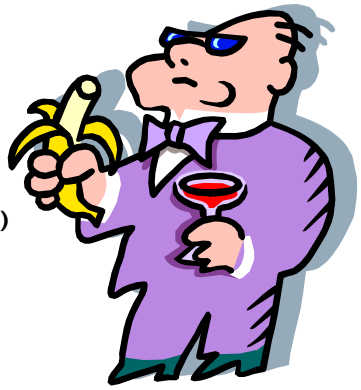

© 2007 IBM Corporation
8

IBM

EPSPT Programmatic Help With Coexistence PTFs

- **Finally, for z/OS R8! Some programmatic assistance to help you determine that you've got all the Coexistence PTFs installed for a z/OS migration!**
- **This solution uses the existing Electronic PSP Tool (EPSPT).**
  - **"Extract files" are now produced from the ZOSGEN PSP bucket that contain the Coexistence PTFs.**
  - **You should use the Extract file for the z/OS release you are migrating *FROM*.**
    - **For instance, the ZOSV1R7 ZOSGEN Extract file contains the PTFs that you need to coexist with z/OS R8.**
- **The Coexistence PTFs from the z/OS Migration book, along with any subsequently defined coexistence PTFs are included, for a complete Coexistence PTF list.**

© 2007 IBM Corporation

IBM

EPSPT Programmatic Help With Coexistence PTFs

- **Steps to follow:**

1. **If you aren't already using EPSPT, download and install the EPSPT, available from http://www14.software.ibm.com/webapp/set2/sas/f/psp/download.html.**
2. **Download the Extract file from your current release's ZOSGEN PSP bucket subset.**
   - **The list of "to" release coexistence PTFs is found in the "from" release ZOSGEN PSP bucket subset.**
   - **You may already be pulling the Extract file today, but perhaps not from the ZOSGEN subset.**
3. **Run EPSPT using the Extract file from your current release's ZOSGEN PSP bucket.**
4. **Review the EPSPT output report, and resolve any outstanding discrepancies.**

- **Periodically, rerun the steps above for any new Coexistence PTFs that may have been added.**

2007 IBM Corporation

© Copyright IBM Corporation, 2007

IBM Migration Checker for z/OS

- **Some Common Migration Problems:**
  - We have many migration actions required for each release.
  - Applicability of migration actions can be difficult to determine.
  - Determination of the applicability is (in many cases) manual, when it could be programmatic.
  - When programmatic, disparate solutions have been used.
  - Hard to verify if a migration action has been done, especially when deploying to subsequent systems.
  - Would like some way of checking pre-install and post-install "migration health" of a system.
- **An "As Is" Tool That May Help - *IBM Migration Checker for z/OS***
  - From the z/OS Download website: http://www.ibm.com/servers/eserver/zseries/zos/downloads/
  - This tool can:
    - Alert you to some migration actions that you should plan for in the future
    - Detect some migration actions that you can do now to ease your migration to z/OS V1R8
    - Verify that you've done some migration actions correctly when you're running z/OS V1R8
  - Tool verified by 52 early customers. Added 14 of their suggested enhancements.

© 2007 IBM Corporation

IBM

Programs included in the IBM Migration Checker for z/OS

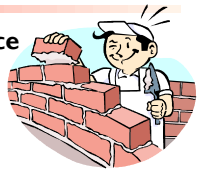

1. **TN3270AS - Run the TN3270E Telnet server as a separate address space**
2. **HFSZFS - Migrate from HFS file systems to zFS file systems**
3. **ZFSMULT - Discontinue use of multi-file system aggregates**
4. **BTSIZE - Use the new default for BLOCKTOKENSIZE in IGDSMSxx**
5. **OLDVSAM - Redefine existing VSAM data sets that contain the IMBED, REPLICATE, and KEYRANGE attributes**
6. **MASTRCON - Accommodate the removal of the master console**
7. **ONEBCON - Accommodate the removal of 1-byte console IDs**
8. **LLSTCAT - Catalog SYS1.SIEALNKE and SYS1.SIEAMIGE for the default system link list in z/OS V1R8**
9. **NEWDS - Add references to new data sets and paths**
10. **OLDDS - Remove references to old data sets and paths**
11. **SMSLVL - Ensure the integrity of SMS control data sets**
12. **ETCUSS - Update Communications Server and z/OS UNIX /etc configuration files**

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

IBM Migration Checker for z/OS Installation Instructions

1. **Download/upload Web Deliverable data sets: CLIST, JCL, and LOAD**
2. **TSO RECEIVE INDATASET(xxx) for each data set**
3. **Execute $SETUP from the CLIST data set. Do this for each system you want to run these programs from – may share the CLIST and LOAD data sets.**
4. **If desired, update $sysname in the CLIST data set, with your CSI info.**
5. **Run $MIGALL from the hlq.&sysname.JCL to execute all migration programs serially. Run an individual member for just one migration program you're interested in.**
6. **Review each member of hlq.&sysname.OUTPUT for what was found. Rerun as necessary.**

Usage Notes:

- **Evaluates the currently active system.**
- **Programs will be to detect migration actions, not to correct. Can re-run these programs as many times you want, at any point in your migration cycle.**
- **"As is" support - feedback can be provided on a forum.**

© 2007 IBM Corporation

Elements with Migration Actions for z/OS V1R8

- **Documented in *z/OS Migration***
  - *For complete migration tasks for z/OS R8, see this book!*
    - ***from R5 to R8, and R6 to R8, and R7 to R8, customized books***
    - ***"When behaviors aren't the same anymore, migration actions are called for."***
- ***From z/OS V1R7* to z/OS V1R8:**
  - **BCP**
  - **Communications Server**
  - **Cryptographic Services - ICSF, OCSF, PKI,**
  - **DFSMS**
  - **DFSORT**
  - **Distributed File Service - zFS**
  - **HCD (introduced in z/OS R7, not R8)**
  - **HCM**
  - **Infoprint Server**
  - **Integrated Security Services - Firewall**
  - **ISPF**
  - **JES2**
  - **JES3**
  - **Language Environment**
  - **Library Server**
  - **msys for Setup**
  - **NFS**
  - **RMF**
  - **SDSF**
  - **Security Server (RACF)**
  - **XL C/C++**
  - **z/OS UNIX System Services**

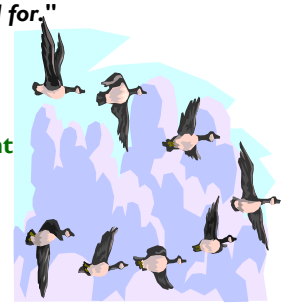

  - ***means that selected element migration actions from these elements are discussed***
  - ***If coming from z/OS R5 or z/OS R6 to R8, there are additional elements that have migration actions***

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

What exactly is a z/OS "Migration Action"?

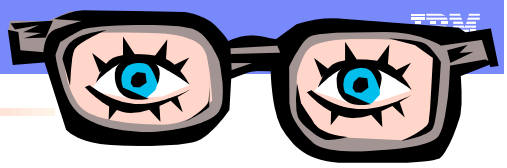

- **Migration is**
  - the installation of a new version or release of a program to replace an earlier version or release.
- **After a successful migration, the applications and resources on the new system function the same way they did on the old system, if possible. "System Equivalence"**
- **Migration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

© 2007 IBM Corporation

Some General Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - **Verify that you have enough XCF groups in your sysplex CDS and enough XCF members in your XCF groups (Required-IF, as of R7,R8)** - In R7: XCF group IOEZFS, R8 XCF group SYSXCF. Reformat if you need to.
  - **Use the "new" default for your prog mgmt binder COMPAT specification (Required-IF, as of R8)**
    - R8 intro PO5 (PO4 was intro in R3, when COMPAT=MIN was the new default). Use default, in your SMP/E UTILITY entry.

- **Migration Actions Pre-First IPL:**
  - **Set up your IPCS environment (Required)**
    - Don't forget R7 new data set SYS1.SIEAMIGE, for SYSLIB, like SYS1.MIGLIB.
  - **Migrate /etc and /var system control files (Required)** - NFS obsoletes some /etc files.
  - **Remove references to deleted data sets and path (Required)**
    - ISP.SISPSASC (ISPF now uses LE RT SCEERUN/SCEERUN2.)
    - Still need DDDEF for CNMLINK for OSA/SF PTFS! (Can point to an empty data set)
  - **Add references to new data sets and paths (Required)**
    - New in R6, PDSE linklist data set **SYS1.SIEALNKE**.
    - New in R7, PDSE linklist data sets **SYS1.SIEAMIGE and SYS1.SHASLNKE**. New APF list PDSE **SYS1.NFSLIBE** to replace SYS1.NFSLIB.
    - New in R8: **SYS1.SIEALNKE** and **SYS1.SIEAMIGE are automatically added to the default system linklist!** You **must** catalog these data sets!

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some General Migration Actions for z/OS V1R8 from z/OS V1R5

► ***Continued from previous foil...***

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

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

➡ New in R6, **SMSPDSE1**, restartable PDSE address space that allows you to recover from some PDSE problems without having to re-IPL. And **TN3270E**. By default, not created during IPL.

➡ New in R7: **DEVMAN, HZSPROC, JES2S00***1* for JES2 NJE over TCP/IP

➡ New in R8: **DSSFRDSR** for hsm. Starts automatically whenever a ds is recovered from DASD using FRRECOV DSNAME command. Terminates when hsm terminates.

➡ **Accommodate new SCOPE=COMMON data spaces in IEASYSxx MAXCAD (Required-IF)**

- In R7: 1 for *JES2*WTO (rolled back in OA05981), 1 or more for XES connections to serialized structures.
- In R6 and R8: None
- **IBM Health Checker for z/OS can help you determine what to specify for the MAXCAD value!**

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

- **<none>**

© 2007 IBM Corporation

Some BCP Migration Actions for z/OS V1R8 from z/OS V1R5

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

Accommodate the removal of 1-byte console IDs (Required, as of R8)

- **R7 removed 1-byte console ID support from macros and commands,**
- **R8 completes the 1-byte console ID removal. 1-byte console IDs might still continue to be accepted, but they will be treated differently, so programs that use them will probably behave unexpectedly.**
  - **Console ID Tracking Facility is still incorporated into z/OS R8. Use it!**
  - **New routing attribute in R8, UNKNIDS. Any consoles that request this attribute receive messages directed to 1-byte console IDs.**

Accommodate the removal of the master console (Required-IF, as of R8)

- **Before R8: one op console within the sysplex was req to be designated as the master console.**
- **As of R8: master console is removed, and functions that were unique to the master console have been made available to other consoles. Can still define multiple consoles with master authority. System enforces that the system console has master authority and LOGON(OPTIONAL).**
  - **Remove the NOCCGRP keyword from CONSOLxx and "MSTCON" keyword from the CNGRPxx. (Can keep if you share parmlib with <R8 systems, R8 warns but continues.)**
  - **Following commands are no longer supported: VARY ...,MSTCONS and DISPLAY C,MCONLY**
  - **Add INTIDS keyword to console definitions in CONSOLxx. This is a new R8 routing attribute. Any consoles that request this attribute receive messages directory to console ID zero (which used to represent the master console).**
    - **Got COM='MN JOBNAMES,T' in COMMNDxx? Use SETCON MONITOR command instead.**
  - **...**

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some BCP Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions Pre-First IPL:**
  - **Accommodate the removal of console switching (Required-IF, as of R8)**
    - In R8, console failures no longer cause a console switch. Operators cannot issue SWITCH CN command to switch consoles. Remove ALTGRP keyword from CONSOLxx. (Can keep if sharing <R8. R8 warns, and continues). Cannot issue VARY CN(...),ALTGRP command.
  - **Discontinue use of the MSGRT command (Required-IF, as of R8)**
    - In R8, MSGRTcommand was removed. Remove MSGRT keyword from CONSOLxx. (Can keep if sharing <R8. R8 warns, and continues)
  - **Update the hardcopy medium definition (Required-IF, as of R8)**
    - Hardcopy switching is no longer supported. To reduce likelihood of losing hardcopy, define both SYSLOG and OPERLOG as the hardcopy medium.
  - **Modify programs that reference unsupported console functions (Required, as of R8)**
    - In R8, many control blocks remove support for 1-byte console IDs, console switching, and master consoles. Examine your programs that reference these control blocks and ensure they will still work on R8. Clean **compile assures of successful removal of references to 1-byte fields.**
  - **GRS-related migration actions:**
    - **Use GRSCNFxx to specify the maximum number of concurrent ENQ requests (Required-IF, as of R8 and R7 APAR OA11382)**
      - Before, you could zap the GRS vector table (GVT) to increase authorized and unauthorized maximums of concurrent ENQ requests. Now, use new ENQMAXA and ENQMAXU keywords in GRSCNFxx parmlib member.
    - **Update the GRSQ setting in GRSCNFxx (Required-IF as of R8, and rolled back to R7 OA11382)**
      - GRSQ keyword is now only applicable to STAR mode and its default is CONTENTION.
    - **Ensure that GRSCNFxx is parsable for GRS=NONE (Required-IF, as of R8)**

© 2007 IBM Corporation

Some BCP Migration Actions for z/OS V1R8 from z/OS V1R5

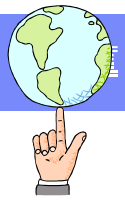

- **Migration Actions Pre-First IPL:**
  - **Add more real storage if needed (Required-IF, as of R8)**
    - **As of R8, z/OS supports up to 4 TB of real storage on single z/OS image.**
    - **z/OS R8 no longer supports the physical swapping of address spaces. (Physical swapping is the process of moving all of the virtual storage of the addr space to aux or paging storage.)**
    - **Elimination of physical swapping means that a portion of the address space must remain in real storage. Therefore, may have performance implication if real storage is in very short supply or the paging rate increases or both.**
    - **Now: Analyze RMF reports from your pre-R8 system to see if a large number of physical swaps occurred. Consider increasing the amount of real storage for z/OS R8 when physical swaps are a frequent event in your configuration, or one or more swappable addr spaces own a large amount of fixed frames that are being swapped out to aux storage on your current system.**
    - **On R8: use RMF to monitor your use of real storage by examing the Unreferenced Interval Count (UIC) value. The UIC values have changed in R8: higher the UIC, less contention in system. The new maximum (65535) means there's ample frames available in the system. Low UICs means system is storage constrained.**
    - **May update automation on msg: IRA205I 50% AUXILIARY STORAGE ALLOCATED, to add additional page data sets before a critical shortage occurs.**

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some XL C/C++ Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions Pre-First IPL:**
  - **Increase the MEMLIMIT system parameter size to avoid abends (Required-IF, as of R8)**
    - **As of R8, the IPA link step uses 64-bit virtual memory, which may cause the XL C/C++ compiler to abend if there is insufficient storage.**
    - **May need at least 3000 MB above-the-bar storage.**
    - **One way to increase the above-the-bar storage is to set your MEMLIMIT value to be at leaset 3000 MB.**

© 2007 IBM Corporation

Some Communications Server Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - **IP Services: Remove ASSORTEDPARMS and KEEPALIVEOPTIONS statements from the TCP/IP profile (Required-IF, as of R8)**
    - **Migrate to new profile statements before z/OS R8**
  - **IP Services: Migrate from OROUTED to OMPROUTE (Required-IF, as of R7)**
    - **Use OROUTED -c to make migration to OMPROUTE easier BEFORE migrating to R7!**
    - **As of R7, if you want to use a daemon to support the RIP protocols, you must use OMPROUTE.**
  - **IP Services: Update profile is still using TN3270E server legacy statements (Required-IF, as of R8)**
    - **QUEUESESSION, and three options on BEGINVTAM block: LUSESSIONPEND, MSG07, and TELNETDEVICE) are replaced by improved functions.**
    - **Equivalent functions are available.**
  - **IP Services: Ensure that UDP port 514 is available to syslogd if not started with the -i option (Required-IF, as of R8)**
    - **If syslogd is configured to process log messages from a remote system daemon, and fails to bind to UDP port 514, syslogd terminates.**
    - **If you start syslogd with the -i option, no changes to TCP/IP port reservations for syslogd.**
    - **Otherwise, ensure that UDP port 514 is reserved for OMVS or (preferably) the job name under which syslogd will run. Verify syslogd is bound to UDP port 514 with with netstat -o/allc**
  - **IP Services: Use the REXX FTP client API from customized TSO/E REXX modules (Required-IF, as of R8)**
    - **As or R8, TSO/E REXX parm modules provide the REXX FTP client API.**
    - **If you customized the TSO/E REXX parm modules in a prev rel and now want to use the new REXX FTP client, then migrate your customizations to the R8 level.**

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some Communications Server Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions Pre-First IPL:**
  - **IP Services: Update /etc configuration files (Required-IF, as of R7 and R5)**
    - Update from /usr/lpp/tcpip/samples , to:
      - for Policy Agent and Intrusion Detection Services in R8, put in a user directory.
        - pagent_IPSec.conf, pagent_CommonIPSec.conf, and pagent_IDS.conf
      - /etc/protocol : hopopt, encapsulating security payload, and authentication header protocol was added in R7
      - /etc/services : R7 APAR OA12364 provides ability for JES to communicate NJE work over TCP/IP.
      - /etc/osnmpd.data : every release sysName MIB object is updated

mig

IP Services: Run the TN3270E Telnet server as a separate address address space (Recommended, as of R6)

- Will be a requirement in the release following z/OS R8.
- Provides visibility and control over TN3270 function separate from the TCP/IP stack

1. set up a superuser with OMVS segment
2. associate user with telnet proc, using RACF STARTED class
3. change the priority of Telnet by assigning Telnet to a service class other than the default SYSSTC class in the STC subsystem.
4. remove the TN3270E Telnet server profile statements from the TCP/IP profile.
5. create Telnet proc, specify profile.

© 2007 IBM Corporation

HCD Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - <None>
- **Migration Actions Pre-First IPL:**
  - <None>
- **Migration Actions Post-First IPL:**
  - **Upgrade the IODF to V5 (Required if you want to make updates to the IODF, and it's not yet at the V5 level, as of R7)**
    - Coexistence considerations with the V5 IODF! For back-level systems:
      - Need OA07875 (HCD) to read from, IPL with, or dynamically activate a R5 IODF. CANNOT use the downlevel systems with this APAR to update the V5 IODF.
      - Need BCP Allocation APAR OA08197 installed on back-level systems to IPL. Without this compatibility code, systems below the z/OS R7 level which attempt to use a V5 IODF will load WAIT0B0 RSN02.
    - Once your IODF is at V5, you must use z/OS R7 HCD libraries to process any updates (STEPLIB or JOBLIB is fine).

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some DFSMS Migration Actions for z/OS V1R8 from z/OS V1R5

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

mig

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

- Determine which VSAM data sets were allocated with IMBED, REPLICATE, or KEYRANGE by using tool IMBDSHIP.JCL.TRSD.
- Install the PTF for APAR OA10952. The PTF fixes the problem in which an exported copy of a VSAM data set that has the IMBED or REPLICATE attribute (or both) does not remove the attribute when the IMPORT is done.
- Schedule a time for the VSAM data sets to be unavailable, and redefine them.
- When support for IMBED and REPLICATE is removed and data sets are recalled or restored by DFSMShsm or DFSMSdss, these attributes will be removed during the recall or restore and the data set will be in a usable state.

➡ Define new security profiles for DFSMShsm fast replication support (Required-IF, as of R8)

- FACILITY class resource names have been shortened for fast replication commands, so you can specify the full names of copy pools. If you don't have the new profiles, hsm commands fail!!!
  - STGADMIN.ARC.FRBACKUP.* -> STGADMIN.ARC.FB.* , for instance
- Coex consideration: mixed releases in an HSMplex need to have both older and newer policies. Don't convert on pre-R8 to newer policies until all systems are running R8, to allow you to recover indiv ds from any copy pool backup version made on any system in HSMplex.

© 2007 IBM Corporation

Some DFSMS Migration Actions for z/OS V1R8 from z/OS V1R5

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

➡ Use the new default for BLOCKTOKENSIZE in IGDSMSxx (Recommended, as of R8)

mig

- Added in R7, specifies whether applications that use certain system functions have to code BLOCKTOKENSIZE=LARGE on the DCBE macro before using large format data sets.
- Before: default was REQUIRE. Now: default is NOREQUIRE.
- If you do override the R8 default (and use REQUIRE), programs will have to have BLOCKTOKENSIZE=LARGE explicitly coded on the DCBE macro in most cases to exploit large format sequential data sets.

Use catalog search order, instead of JOBCAT and STEPCAT (Required-IF, as of R7)

- In z/OS R7, JOBCAT and STEPCAT support has been removed. You can no longer enable or disable them.

- **Migration Actions Post-First IPL:**

- Ensure the integrity of SMS control data sets (Recommended)

mig

- Update, translate, validate, and activate SMS policies from a R8 system, when in a mixed-level sysplex

Avoid issuing commands against a copy pool from a pre-R8 system after the copy pool has been converted to R8 format (Required, as of R8)

- In R8, you can dump to tape copy pool versions residing on DASD. This enh changes DFSMShsm control ds records automatically first time you issue an FRBACKUP command against a copy pool on R8. But from then on, the only FR command that can be successfully issued against a copy pool from a pre-R8 system is FRRECOV from DASD.
- Ensure you longer need to back up copy pools from any pre-R8 system before issuing an FRBACKUP command against a copy pool on R8.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some DFS Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - **Discontinue use of multi-file system aggregates (Required-IF, as of R8)**
    - In z/OS R8, cannot mount any zFS file systems contained in multi-file system aggregates shared in a sysplex. Will result in a failure, with msg IOEZ00522E.
    - Convert to zFS compatibility mode system aggregates.
- **Migration Actions Post-First IPL:**
  - ➡ **Use the F OMVS command to stop the zFS physical file system (Required-IF, as of R8)**
    - Before: to gracefully stop the zFS physical file system was STOP ZFS.
    - Now: use F OMVS,STOPPFS=ZFS. STOP ZFS has been removed.

© 2007 IBM Corporation

Some JES2 Migration Actions for z/OS V18 from z/OS V1R5

Update JES2 exit routines due to the z/OS R7 structure changes (Required-IF, as of R7)

- Some NJE and input service functions were moved from the JES2 AS to the application AS. Results in improved performance, reliability, failure isolation, and outage avoidance. This can break processing done by the following JES2 exits:
  - Exits 2, 3, 4, 7, 20, 39, 46, and 47 are not called for INTRDRs and NJE/TCP processing.
  - Exit 8 is now called for INTRDR and NJE/TCP processing.
  - Exits 36 and 37 now get control in a different address space for INTRDR and NJE/TCP processing. These are pre-SAF and post-SAF exits.
  - Exit 13 has been deleted.
  - Exit 49 has new function and is being called for $S J commands.
  - Exit 40 can now be used to control NJE mail notification (function that was in exit 13).
  - The input to exits 2, 3, 4, 20, 46, and 47 has changed significantly.
- Read "*JES2 z/OS 1.7 Exit Migration Guide*" on web for important information!
  - http://www.ibm.com/servers/eserver/zseries/zos/installation/zos17_jes2_migration.html

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some JES2 Migration Actions for z/OS V18 from z/OS V1R5

Update procedures because of changed input processing (Required-IF, as of R7)

- Input processing was rewritten, thus this change may affect your applications and procedures.
  - JCL/JECL processing: more details on errors, parsing of operands conforms to standard JCL rules (some previous errors may now run correctly)
  - SYSIN data processing: SYSIN ds properly handles carriage cntl, var recfm, and SYSIN up to 32KB long. Jobs that work on z/OS R7 may fail in prior releases. Submitted jobs from R7 will exec correctly w/ a mixed-level MAS.
  - Internal reader processing in requestor AS: Moved from JES2 AS to AS that allocated the internal reader. Higher CPU and IO now associated with the owning AS. Processing occurs at dispatching priority of the requestor and not JES2's priority, thus lower priority AS may not submit jobs as fast as higher priority AS.
  - Buffer of data in internal readers: Now, job cards are processed as soon as they are passed to the internal reader. No buffering of data, and the internal reader data space no longer exits.
  - Internal reader messages in JOB log: Input processing msgs are issued in the submitting AS and placed in the JOB log (JESMSGLG) of the submitting AS. Increases size of the JOB log. May want to spin or suppress the JOB log.
  - Internal reader processing when JES2 AS fails: JES2 AS used to discard job on an internal reader when JES2 AS abnormally terminated. Now, JES2 tries to retain jobs that are active on an internal reader when JES2 terminates.

© 2007 IBM Corporation

Some Language Environment Migration Actions for z/OS V1R8 from z/OS V1R5

- Migration Actions Pre-First IPL:
  - Determine the impact of changes to default run-time options (Required-IF, as of R8)
    - Changes to the run-time options in R8: DYNDUMP is a newly added option.
    - No changes in R6 or R7!
    - if you are using assembler modules to specify your installation-wide run-time options (for CEEDOPT, CEECOPT, or CELQDOPT), then you must change them for the new R8 DYNDUMP option.
    - If you are using CEEPRMxx (introduced in z/OS R7) and the default specification for the newly added DYNDUMP option is acceptable to you, then no migration action is required!
  - Determine the impact of the DYNDUMP run-time option (Required, as of R8)
    - New option adds the ability to write IPCS-readable dumps automatically. Default supplied by IBM does not affect the behavior when an application error results in a U4039 abend. However, the supplied default affects all other U40xx abends.
    - Before: a system dump was only written for the U40xx abends when a SYSUDUMP, SYSABEND, or SYSMDUMP ds was allocated.
    - As of R8: regardless of whether a SYSUDUMP, SYSABEND, or SYSMDUMP data set is allocated, a dump is written to the hlq specified in the DYNDUMP run-time option. You can turn off this new first failure data capture behavior by specifying NOTDUMP (whereas TDUMP is the default) as the value for the third suboption of DYNDUMP. The default hlq (*USERID) might also need to be adjusted for your installation.
    - Applicable to CEEDOPT and CELQDOPT only (ignored in CICS environment).

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Some z/OS UNIX Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - Use new COMPAT default with binding with c89 (Recommended, as of R8)
    - As of R8, c89 will no longer override the binder's COMPAT default! (btw, z/OS R8 introduces PM5 level, PM4 was introduced back in z/OS R3. Programs created with a lower COMPAT level can run in more environments.)
    - Add to /etc/profile c89 section: export ${_CMP}_OPTIONS="-WI,compat=unspec"
  - Update BPXBATCH jobsthat specify MVS data sets (Required-IF, as of R8 and back OA11699)
    - Before: only USS files could be used for STDOUT/STDERR DD - no faiures given -> /dev/null
    - As of R8: MVS data sets can be used for STDOUT/STDERR DD. Ensure that any MVS ds are intended to be used for STDOUT/STDERR.

- **Migration Actions Pre-First IPL:**
  - mig • Update config files changed by IBM (Required, as of R8) - /etc/profile, and /etc/ohelp.ENU
  - BPXPRMxx SHRLIBMAXPAGES parameter is no longer used as a limit because user shared library objects are no longer supported (Required-IF as of R8). Ignored, if specifiied.
  - mig Migrate from HFS to zFS (Recommended)
    - At R7, zFS is allowed in all levels of file system hierarchy (including root!).
    - Support for the HFS file system has been stabilized and you are encouraged to migrate to the zFS file system for better performance, but IBM has not announced removal of support for the HFS file system.
    - Use the handy ISPF-based tool, BPXWH2Z, available as of R7.

© 2007 IBM Corporation

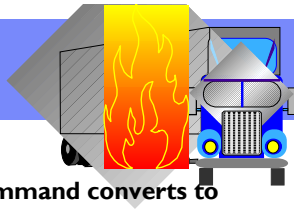

Migration Actions Not to Overlook!

- DFSMShsm copy pool R8-format coexistence- R8 FRBACKUP command converts to R8-format, then, only command on pre-R8 system is FRRECOV from DASD.
- z/OS R7 DFSMS JOBCAT/STEPCAT support removed.
  - Use the MODIFY CATALOG,ENABLE(JOBSTEPCAT) on R5 to turn support back on (for now!)
- z/OS R7 JES2 Z2 - must be in Z2 mode!
- z/OS R7 JES2 Exit and Input Processing changes
  - Internal reader processing moved from JES2 address space to address space that allocated the internal reader.
- z/OS R7/R8 BCP 1-byte Console ID support removal
  - mig Use the Console ID Tracking Facility, introduced in z/OS R4 Consoles Enhancement feature with the latest exclusion list now.
  - Final removal of 1-byte Console IDs in z/OS R8
- z/OS R7 XL C/C++ does not ship the OS/390 R10 Compilers
  - Move to the ISO C/C++ compilers now, may still use OS/390 R10 compilers from prior system.
- z/OS R7 HCD IODF V5 Coexistence considerations with lower systems
- z/OS R6 Architecture Level Set - must be z/Architecture mode
  - Use the z/Architecture checklist provided.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Future Migration Actions Not to Overlook

- **Remove your dependency on the C/C++ IBM Open Class (IOC) Dynamic Link Libraries (DLLs)**
  - Use the Standard C++ Library instead
  - Removed in the release after z/OS R8
- **Run the TN3270E Telnet server as a separate address address space**
  - Must do this before the release after z/OS R8
- **z/OS DFSMS Imbed, Replicate, Keyrange removal in a future release.**
  - Use the IMBDSHIP tool, and install PTF for APAR OA10952 now!
- **Start your conversion from HFS to zFS after you migration to z/OS R8**
  - Use z/OS R7's tool, BPXWH2Z to help.

© 2007 IBM Corporation

Summary: *Plan your migration to z/OS V1.8!*

- **z/OS Coexistence- z/OS V1R5 through z/OS V1R8 supported.**
  - z/OS V1R7 still orderable in SystemPac, if you need it
- **Removed Functions since z/OS V1R6**
  - JES2 R4 mode, JOBCAT/STEPCAT, OS/390 R10 C/C++ Compilers, 1-byte console IDs, zFS multi-file system aggregates from share fs, Firewall Technologies,
  - Future: TN3270 in TCP/IP stack, run-time support for C/C++ IBM Open Class DLLs, old VSAM attributes, zFS multi-file system aggregates.
- **Coexistence System Requirements** - use the EPSPT tool to help!
- **IBM Migration Checker for z/OS** can assist you!
- **Migration Actions for z/OS V1R8**:
  - **z/OS General** - enough XCF groups, new and old data sets, new address spaces, ...
  - **BCP** - 1-byte console id removal, master console removal, 4TB real changes, ...
  - **XL C/C++** - MEMLIMIT at least 3000MB for IPA link step
  - **Communications Server** - obsolete statements from TCP/IP profile, UDP 514 available,...
  - **HCD** - IODF V5 coexistence consideration
  - **DFSMS** - new security profiles for hsm fr commands, FRBACKUP command changes, ...
  - **DFS zFS** - use compatibility-mode aggregates, don't use STOP ZFS, ...
  - **JES2** - exit changes, changes in input processing, ...
  - **Language Environment** - new DYNDUMP run-time option
  - **z/OS UNIX** - COMPAT option for c89, HFS to zFS conversion, ...

➡ **Review the z/OS Migration book, for a complete list of z/OS V1R8 migration actions!**

© Copyright IBM Corporation, 2007

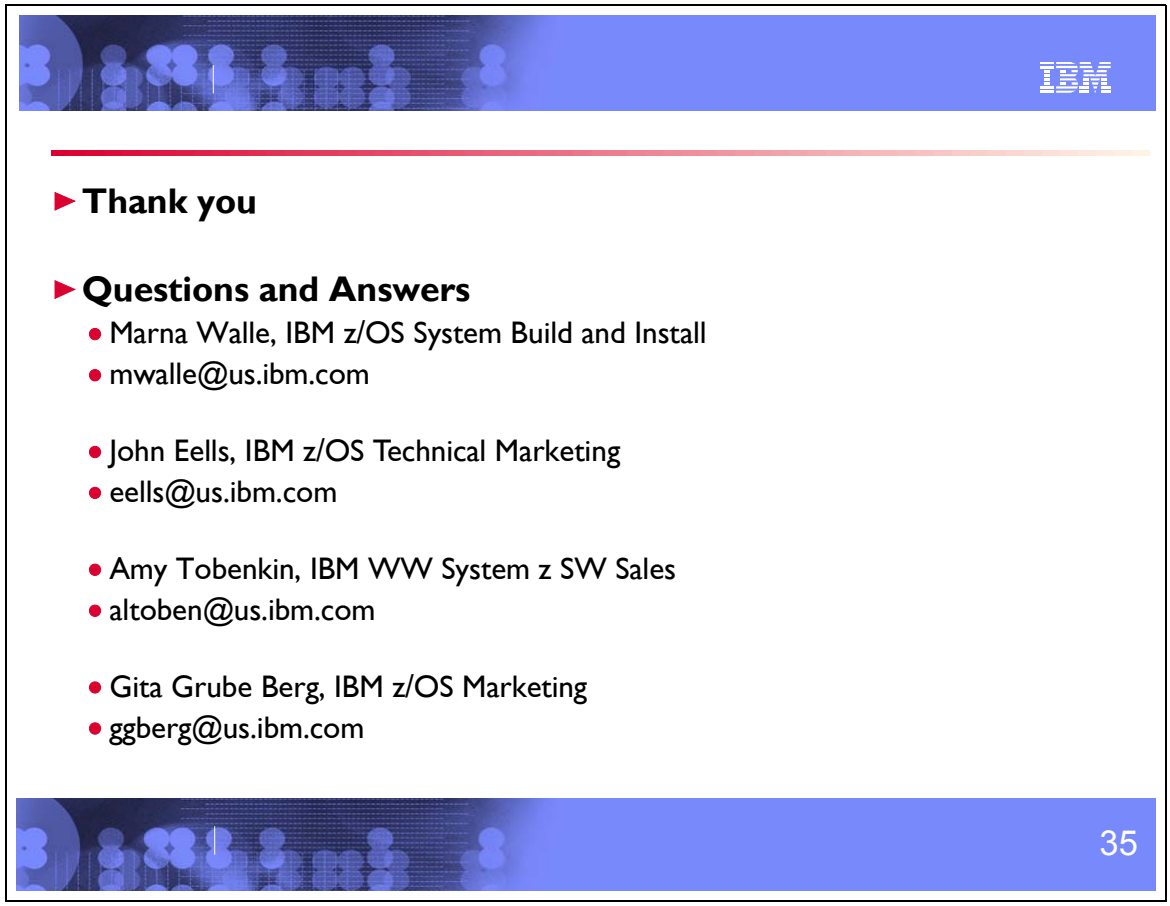

© Copyright IBM Corporation, 2007

The following portion of this document includes more details on the information presented during the teleconference.

The slides are repeated to show the context of the additional material.

© Copyright IBM Corporation, 2007

Plan your migration to z/OS V1.8!

Marna Walle
z/OS System Build and Install
mwalle@us.ibm.com

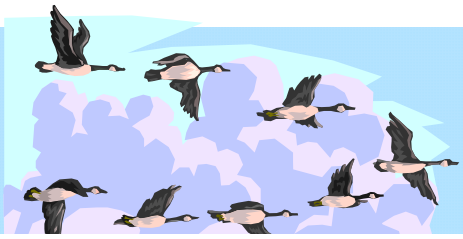

April 2007
Copyright IBM Corporation

Plan your migration to z/OS V1.8!

This discussion will cover many of the technical installation requirements for your z/OS V1R8 migration. Included will be:

- Highlights of what functions have been removed since z/OS V1R6, and some important removals in the future to prepare for now
- Ordering information for z/OS V1R8 (and also for z/OS V1R7!)
- Coexistence requirements for z/OS V1R8 and beyond
- Some helpful new tools, for coexistence checking and migration assistant
- Selected migration actions from BCP, XL C/C++, Communications Server, HCD, DFSMS, DFS zFS, JES2, Language Environment, and z/OS UNIX

The general availability date for z/OS V1 R8 occurred on September 29, 2006.

Trademarks

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

AFP	DFSMStvs	FICON	MVS/ESP	PR/SM	VTAM
AIX	DFSMStvs	HiperSockets	NetSpool	Print Services Facility	WebSphere
AnyNet	DFSORT	IBM	NetView	PrintWay	zArchitecture
BookManager	Enclave	IBMLink	Open Class	RACF	z/OS
CICS	ESCON	IMS	OS/390	Resource Link	z/VM
DB2	Eserver	Infoprint	Parallel Sysplex	RMF	zSeries
DB2 Universal Database	FICON	Language Environment		S/390	
DFSMSdfp		Multiprise		S/390 Parallel Enterprise Server	
DFSMShsm		MVS		Sysplex Timer	
		MVS/ESA		SystemPac	
				System z9	
				Tivoli	
				TotalStorage	
				VisualAge	
				VSE/ESA	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
UNIX is a registered trademark of The Open Group in the United States and other countries.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

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

Notes:

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

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

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

ENABLING BUSINESS. A THROUGH Z.

© 2005 IBM Corporation

Plan your migration to z/OS V1.8! Agenda

- **z/OS Coexistence-Migration-Fallback-Service Policy**
  - z/OS Coexistence-Migration-Fallback
- **Ordering z/OS V1R8 (and z/OS V1R7)**
- **Positioning for z/OS V1R8:**
  - Highlights of Removed Functions since z/OS V1R6
  - Ensuring System Requirements are Satisfied
    - Driving and Target System Requirements
  - Coexistence System Requirements
    - EPSPT Programmatic Help with Coexistence Fixes!
  - IBM Migration Checker for z/OS
- **Migration Actions for z/OS V1R8:**
  - z/OS General Items
  - BCP
  - XL C/C++
  - Communications Server
  - HCD
  - DFSMS
  - Distributed File Service - zFS
  - JES2
  - Language Environment
  - z/OS UNIX

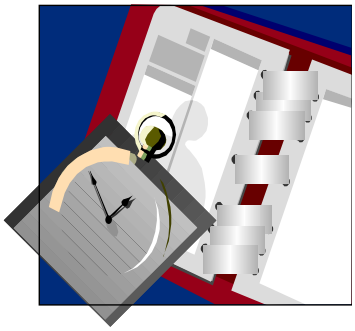

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Related Publications:

- *z/OS and z/OS.e Planning for Installation* (GA22-7504)
- *z/OS Introduction and Release Guide* (GA22-7502)
- *z/OS Program Directory* (n/a)
- *zSeries Platform Test Report for z/OS and Linux Virtual Servers* (formerly, the *z/OS Parallel Sysplex Test Report*)
- *z/OS License Program Specifications* (GA22-7503)
- *z/OS Migration* (GA22-7499)
- *z/OS Summary of Interface and Message Changes* (SA22-7505)
- *z/OS MVS Planning: Operation* (SA22-7601*)*
- *z/OS MVS Initialization and Tuning Reference* (SA22-7592)
- *z/OS UNIX System Services Planning* (GA22-7800)
- S*erverPac: Installing Your Order (no order number; custom-built to your order)*
- *ServerPac: Using the Installation Dialog,* (SA22-7815*)*

Handy Internet Web Pages:

URL	**DESCRIPTION**
http://www.ibm.com/servers/eserver/zseries/zos	IBM z/OS home page
http://www.ibm.com/servers/eserver/zseries/zos/servicetst/	Consolidated Service Test (CST) home page
http://www.ibm.com/servers/eserver/zseries/zos/installation/	z/OS Installation
http://www.ibm.com/servers/eserver/zseries/zos/support/zos_eos_dates.html	End of Service dates for z/OS releases
http://www.ibm.com/servers/eserver/zseries/zos/integtst/	zSeries Platform Test web page
http://www.ibm.com/servers/eserver/zseries/zos/bkserv/	z/OS Library web page
http://www.redbooks.ibm.com/	IBM Redbooks web page
http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/Flashes	IBM Systems Center Flashes
http://www.ibm.com/servers/eserver/zseries/zos/support/	z/OS Technical Support
http://www.ibm.com/servers/eserver/zseries/zos/wizards/	z/OS Wizards
http://www.ibmlink.ibm.com/	IBMLink web page
http://www.ibm.com/servers/eserver/zseries/software/swinfo/	US Packaged Offerings web page
http://service.boulder.ibm.com/390hdmvsinst.html	ESO Overview
http://service.boulder.ibm.com/390holddata.html	Enhanced HOLDDATA web page
http://www.ibm.com/servers/eserver/zseries/software/swinfo/pdo.htm	CBPDO web page
http://www.can.ibm.com/custompac or http://www.ibm.com/ca/custompac	CustomPac (including SystemPac) web page
http://www.ibm.com/software/shopzseries	ShopzSeries home page
http://www.ibm.com/support/operations/us/order_delivery/order_delivery.shtml	IBM Order Status homepage
http://www.ibm.com/servers/eserver/zseries/zos/software/isv18.html	Software Dev. Supporting z/OS V1.8 web page
http://www.ibm.com/servers/eserver/zseries/swprice/	IBM System z Software Pricing: home
http://www.ibm.com/servers/eserver/zseries/swprice/znalc.html	IBM System z Software Pricing: zNALC
http://www.ibm.com/servers/eserver/zseries/zos/zos_sods.html	z/OS statements of direction web page
http://www.ibm.com/servers/eserver/zseries/library/whitepapers	z/OS Technical white papers
http://www.ibm.com/servers/eserver/zseries/zos/installation/zos17_jes2_migration.html	JES2 Migration web page
http://www.ibm.com/software/network/commserver/zos/	z/OS Communications Server web page
http://www.ibm.com/s390/java	z/OS Java web page

IBM Education:

IBM courses are available for z/OS. For schedules and enrollment on the world wide web, IBM Global Campus URL: http://www.ibm.com/services/learning/ .

© Copyright IBM Corporation, 2007

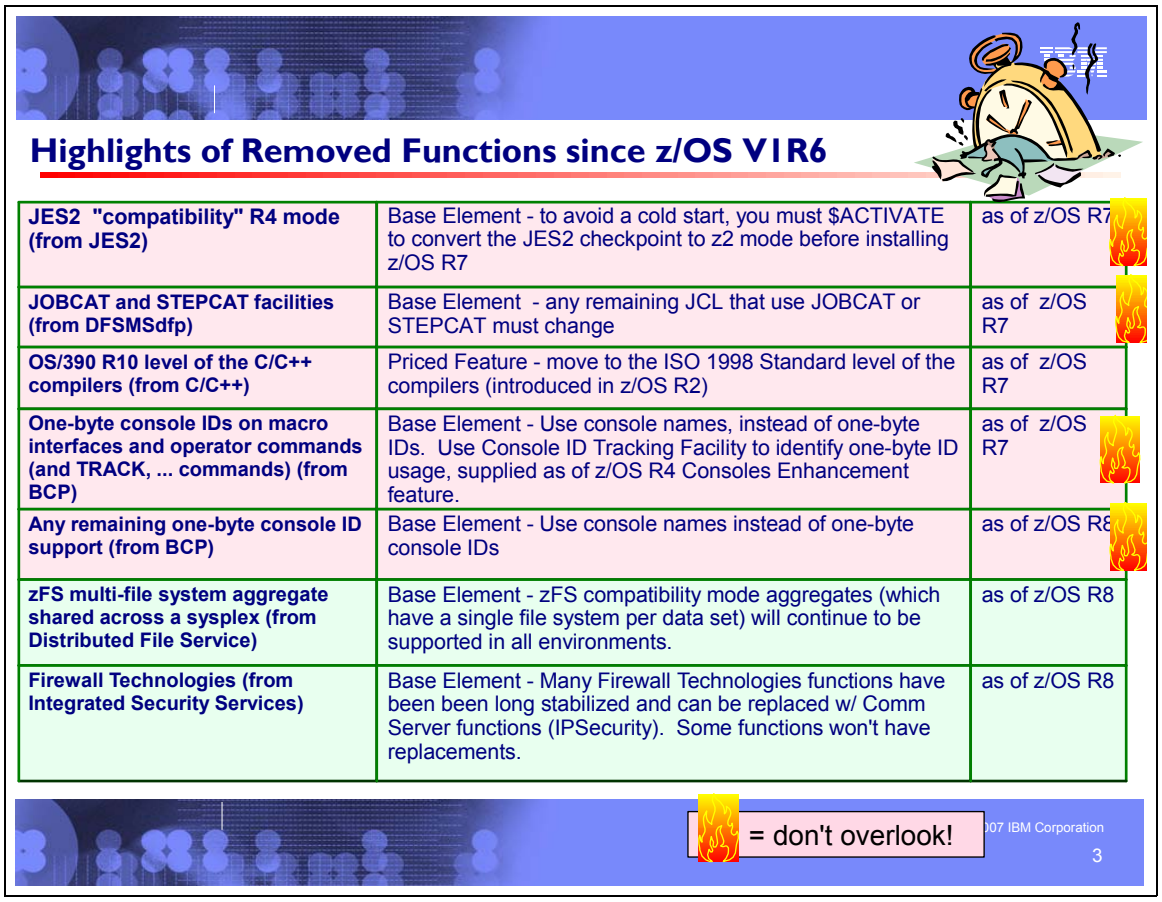

Withdrawn in z/OS R6 (last delivered in z/OS R5)

This section lists items that were withdrawn in z/OS R6. You should take this into account as you plan your migration to z/OS R8. The removal of these functions may have migration actions which you can perform now, in preparation for z/OS R8.

- The C/C++ ISPF panels, which include panels for C/C++ foreground compiles, C/C++ background compiles, and help panels for these compiles, are removed in z/OS V1.6. The z/OS C/C++ compiler can be invoked through z/OS UNIX, using JCL, and under TSO/E.
- z/OS base element Language Environment's use of run-time library services (RTLS) is withdrawn in z/OS V1.6. This function is used primarily in run-time migration. Given the stability and the upward compatibility being provided by the Language Environment run-time library in recent releases of OS/390 and z/OS, this functionality is no longer required.
- The Dynamic Link Library (DLL) Rename Utility, part of z/OS Language Environment, is removed in z/OS V1.6. This utility is used to package and redistribute IBM-supplied DLLs with applications. Since OS/390 V1R3, the C/C++ DLLs have been licensed with the OS/390 and z/OS base operating system. Therefore, the DLL Rename Utility is no longer required.
- z/OS Communications Server support for the SMIv1 version of the SNMP IBM MVS TCP/IP Enterprise-specific MIB module is eliminated in z/OS V1.6. Support will continue for the SMIv2 version of this MIB module. For customers who want to continue using SMIv1, publically available tools can be used to convert an SMIv2 MIB module to an SMIv1 MIB module.
- Effective with z/OS V1.6, IBM has removed the base element, Distributed Computing Environment (DCE) Application Support, from z/OS. DCE Application Support facilitated the interaction between DCE clients and CICS or IMS regions. With the continued evolution of technology and accompanying changes in the marketplace, there is no need for this support. If similar function is required, IBM recommends that customers use IBM WebSphere. The DCE Base Services element, which provides services for developing and running client/server applications, is planned to continue to ship with z/OS.
- Effective with z/OS V1.6, IBM has removed the base element, Encina Toolkit Executive from z/OS. Encina Toolkit Executive provided a set of tools for developing client components of distributed transactional applications. Over time,

the marketplace has moved to other technologies. This element, an enabler for DCE Application Support, is another obsolete element of z/OS V1.6 and is no longer be provided. There will be no replacement.

- With z/OS V1.6, Text Search is no longer made available as a base element of z/OS. Instead, it will be made available for Web download as z/OS and z/OS.e Text Search under a restricted license. This is a change to the statement of direction made in Software Announcement 203-266 (RFA38767), dated October 7, 2003, which stated that with z/OS V1.6, the Text Search base element would be removed and no replacement would be provided. The Text Search element has been removed from the z/OS product. However, the Text Search function is still required and licensed for use only for DB2(R) UDB Text Extender feature of DB2 Universal Database for z/OS, Program Number 5625-DB2, as a Web deliverable named z/OS and z/OS.e Text Search. Before you use the DB2 UDB Text Extender feature, you must download and install the z/OS and z/OS.e Text Search Web deliverable on z/OS V1.6.
  Notable change: Starting with z/OS V1.7, the z/OS and z/OS.e Text Search Web deliverable is no longer available as a z/OS Web deliverable. Instead, the Text Search function is provided by Web download from the DB2 UDB Text Extender Web support site:
  http://www.ibm.com/software/data/db2/extenders/text/te390/

Withdrawn in z/OS R7 (last delivered in z/OS R6)

This section lists items that were withdrawn in z/OS R7. You should take this into account as you plan your migration to z/OS R8. The removal of these functions may have migration actions which you can perform now, in preparation for z/OS R8.

- From base element JES2, IBM removes support for JES2 compatibility mode (R4). As of z/OS R7. JES2 will no longer support compatibility with pre-z/OS V1.2 systems. The $ACTIVATE command, which converts the JES2 checkpoint from z2 mode to R4 mode (compatibility mode), will be removed. Before installing z/OS R7, you must do one of the following:
  - Use the $ACTIVATE command to convert the JES2 checkpoint to z2 mode
  - Offload the SPOOL, cold start JES2 z/OS R7, and reload the jobs on the new SPOOL
- From base element DFSMSdfp, IBM removes the DFSMSdfp JOBCAT and STEPCAT facilities as of z/OS R7. The JOBCAT and STEPCAT facilities have been in existence for many years, predating the introduction of ICF (integrated catalog facility) catalogs. JOBCAT and STEPCAT were designed to address some of the functional shortcomings of VSAM catalogs, such as:
  - ☐ VSAM volume ownership, that is, all data sets on a volume having to be in the same VSAM catalog. Multiple catalogs could not point to data sets on the same volume.
  - ☐ Performance problems resulting from no multilevel alias support, as well as lack of ability to subset catalog data for recovery purposes.
  - ☐ Restrictions in the definition of the catalog SVC interface.

  The introduction of ICF catalogs in the mid-1980s and other catalog enhancements (such as the multilevel alias support) directly addressed those problems. In addition, processes were developed for system build to use system specific aliases instead of JOBCAT or STEPCAT. CBIPO introduced these processes and they are used today by offerings such as ServerPac to create data set entries in the new master catalog of the system being built.

  At the time ICF catalogs were introduced, the JOBCAT and STEPCAT facilities were functionally stabilized. Neither SMS-managed data sets nor UCBs above the 16 megabyte line may be used with JOBCAT or STEPCAT. ICF catalogs contain sufficient functional capabilities that all functions that previously could only be performed with JOBCAT or STEPCAT can now be done without them.

  Furthermore, the use of JOBCAT and STEPCAT can actually cause significant problems. Data sets are generally not cataloged according to the normal predictable search order when JOBCAT or STEPCAT is used. This impacts the ability to do comprehensive installation storage management and can increase staff requirements. For example, interval migration and recall using DFSMShsm is effectively unusable when the data sets cannot be found using the standard catalog search order. The use of JOBCAT and STEPCAT can also result in noticeable increases in the time required to perform catalog requests.
- Due to ISAM's limited functionality and the capabilities of VSAM, particularly VSAM data sets in extended format, z/OS V1.6 was the last release in which DFSMS ISAM and the utility program, IEABISAM, was available. IBM has provided the ISAM Compatibility Interface (ISAM CI) which allows users to run an ISAM program against a VSAM KSDS data set. Details on using this interface and procedures for converting ISAM data sets to VSAM data sets can be found in Appendix E of *Using Data Sets*. This compatibility interface program continues to be provided as part of DFSMS and is **not** discontinued with the ISAM removal from DFSMS.
- From optional feature C/C++ without Debug Tool, IBM has removed the OS/390 R10 level of the C/C++ compilers from z/OS R7. The OS/390 R10 C/C++ compilers were shipped as an aid to migration to the C/C++ compilers that were introduced in z/OS V1R2. The z/OS V1R2 level of the C++ compiler supports the ISO 1998 Standard level of

© Copyright IBM Corporation, 2007

C++. For information about migrating from the older to the newer level of the compilers, see *z/OS C/C++ Compiler and Run-Time Migration Guide*.

- Previously, the CMB= parameter specified the I/O device classes for which measurement data was to be collected, in addition to the DASD and tape device classes. It also allowed you to specify the number of channel measurement block (CMB) slots to reserve for adding more devices with an ACTIVATE. As of the z990 Exploitation feature on z/OS R4, the system ignores the CMB and uses instead the measurement data in the extended channel measurement block (ECMB). As a migration action since then, you must convert user-written programs that make use of the CMB to make use of the ECMB, and contact ISVs to obtain updates to ISV programs that use the CMB. The type of programs especially likely to use the CMB are monitor programs. If ISV support is not available, you could have specified ECMB=NO in parmlib member IEAOPT*xx* as a circumvention until the ISV support is available. See APAR OA06164 for additional details. The ECMB=NO circumvention has been removed in z/OS R7, so have plans in place to discontinue this usage before migrating to z/OS R7.
- One-byte console IDs are removed from macro interfaces and operator commands in z/OS R7. You are not allowed to specify one-byte console IDs on macros (such as WTO/WTOR), or on such operator commands as D C,CN= or D PFK,CN= . You should use console names instead. A service called the Console ID Tracking Facility is available to help you identify one-byte ID usage. The use of console names is already a best practice on OS/390 and z/OS**.** For information about Console ID Tracking facility, see *z/OS MVS Planning: Operations*.
  In addition, the TRACK command, as well as the commands STOPTR, CONTROL T, CONTROL D,U, CONTROL D,H, and MSGRT TR=A are removed in z/OS R7.
- The z/OS Optional Source media feature is not offered as of z/OS R7. The last release offering these materials was z/OS R6. These features contain macros and source code for some programs in the z/OS BCP, BDT base, BDT SNA NJE, BDT File-to-File, DFSMS, MICR/OCR, BCP JPN, and Security Server RACF elements.
- In z/OS R7, support for OROUTED has been removed from Communications Server. You should use OMPROUTE as your dynamic routing daemon.

Withdrawn in z/OS R8 (last delivered in z/OS R7)

This section lists items that were withdrawn in z/OS R8. You should take this into account as you plan your migration to z/OS R8. The removal of these functions may have migration actions which you can perform now, in preparation for z/OS R8.

- z/OS R7 was the last release to allow mounting zFS file systems contained in multi-file system aggregates that are to be shared across systems in a sysplex. IBM has previously recommended that these multi-file system aggregates not be shared in a sysplex environment. Attempts to mount zFS file systems contained in multi-file system aggregates will fail in a z/OS UNIX shared file system environment. Mounting zFS compatibility mode aggregates, which have a single file system per data set, will continue to be supported in all environments.
- z/OS R7 was the last release to include the Firewall Technologies component of the Integrated Security Services element. Many Firewall Technologies functions have been stabilized for some time and can be replaced using comparable or better functions provided by or planned for Communications Server, notably, IPSecurity. In addition, a functionally rich downloadable tool is planned to replace the IPSecurity and IP Filtering configuration GUI support. The following functions will be removed without replacement:
  - FTP Proxy services
  - Socks V4 services
  - Network Address Translation (NAT)
  - RealAudio (TM) support
- z/OS R7 was the last release in which z/OS Communications Server will support the following functions, after which they will be removed from the product:
  - TCP/IP configuration profile block definition statements:
    - ASSORTEDPARMS
    - ENDASSORTEDPARMS
    - KEEPALIVEOPTIONS,
    - ENDKEEPALIVEOPTIONS
      Equivalent capability is provided for the ASSORTEDPARMS statements by the GLOBALCONFIG, IPCONFIG, TCPCONFIG, and UDPCONFIG statements. Equivalent capability is provided for the KEEPALIVEOPTIONS statements by INTERVAL and SENDGARBAGE on the TCPCONFIG statement. For information on TCP profile and configuration statements, refer to *z/OS Communications Server IP Configuration Reference*.

© Copyright IBM Corporation, 2007

  - The SNMP Service Level Agreement (SLA) Version 1 MIB and the PAGTSNMP subagent. In z/OS V1.5, Communications Server provided a new SNMP SLA Version 2 MIB and NSLAPM2 subagent. IBM recommends that you migrate to the new Version 2 MIB and NSLAPM2.
  - The option of defining parallel Enterprise Extender TGs by specifying multiple SAP addresses. Beginning in z/OS V1.5, parallel EE TGs may be defined by using different EE VIPAs on one (or both) of the endpoints.
  - AnyNet. You may implement other IBM solutions such as Enterprise Extender (EE) as a replacement for AnyNet.
- In z/OS R8, the remaining support for one-byte console IDs was removed from control blocks CIB, CSCB, ORE, WQE, XSA. This completes the removal of one-byte console IDs in z/OS. Instead of using one-byte console IDs, console names should be used.
- Support for the following plug-ins for msys for Setup was withdrawn in z/OS R8: TCP/IP Services, z/OS UNIX System Services, Language Environment, Parallel Sysplex, ISPF, and RMF. You will not be able to use msys for Setup for function enablement, setup, or configuration of these areas of z/OS. The DB2 V8 msys for Setup plug-in is unaffected and remains available for setup and configuration of DB2. The TCP/IP plug-in will continue to be available for download via the Web and will no longer require msys for Setup. IBM intends to continue to deliver improvements to help with z/OS setup and configuration in the future.
- z/OS R7 is the last release to support the z/OS msys for Operations element. msys for Operations has been removed from z/OS V1.8. IBM plans to transition many of the current msys for Operations functions to a new user interface and infrastructure in a future release of z/OS.

© Copyright IBM Corporation, 2007

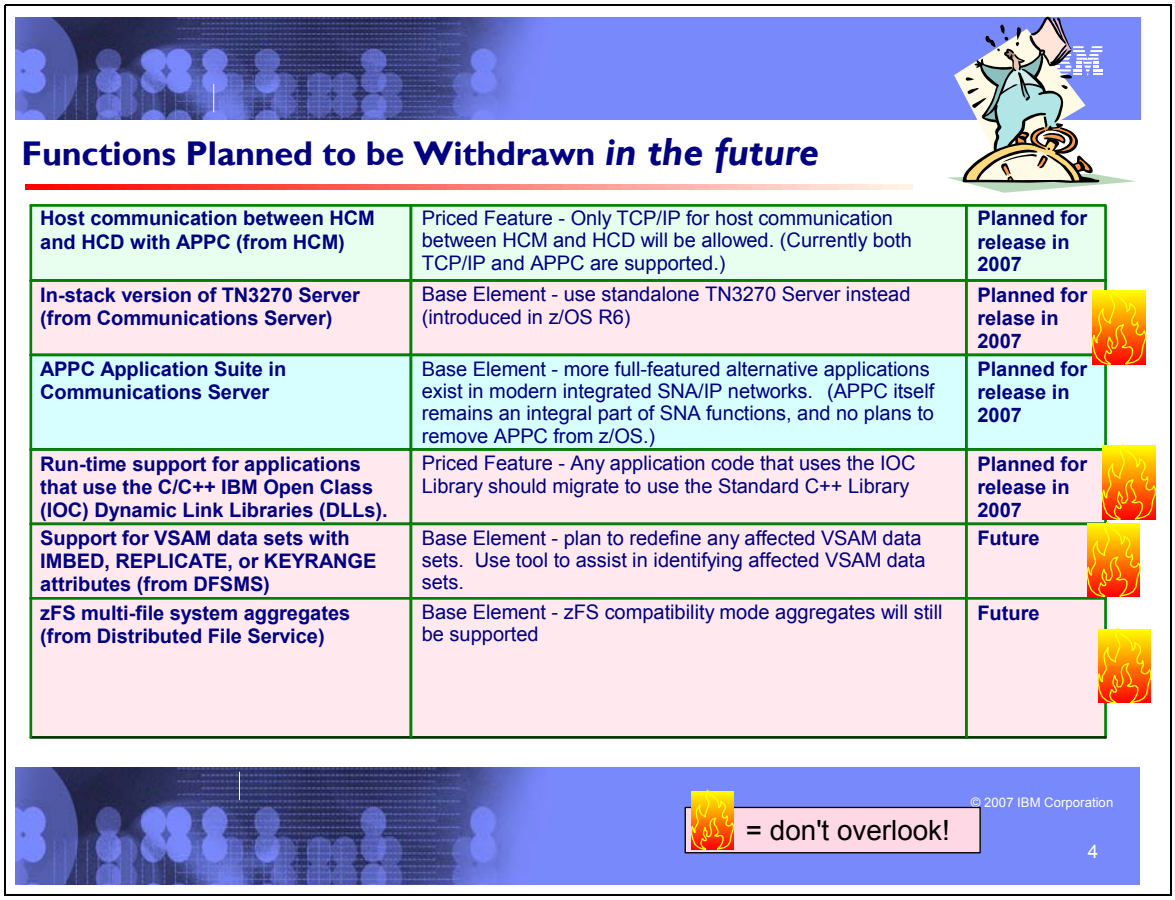

Planned for removal in the release following z/OS R8

This section lists items that IBM has announced it intends to remove in release following z/OS R8. You are encouraged to consider these removals when making your plans for system upgrades. These statements represent IBM's current intentions. IBM development plans are subject to change or withdrawal without further notice.

- z/OS R8 is planned to be the last release that supports host communication between Hardware Configuration Manager (HCM) and Hardware Configuration Dialog (HCD) elements via Advanced Program to Program Communication (APPC). Today for host communication either Transmission Control Protocol/Internet Protocol (TCP/IP) or APPC are supported. Starting with the z/OS release available in 2007, the host communication will be done exclusively via TCP/IP.
- z/OS V1R6 Communications Server and subsequent releases include a standalone TN3270 Server. This standalone TN3270 server is expected to provide increased flexibility, improved reliability, and simplified problem diagnosis as compared to the in-stack version of the TN3270 Server. z/OS R8 is planned to be the last release of z/OS Communications Server which will support the in-stack version of the TN3270 Server. After z/OS R8 this capability will be removed from the product. In preparation for that change, customers should consider implementing the standalone TN3270 Server. For more information, refer to **http://www.ibm.com/software/network/commserver/zos/**
- The Communication Server APPC Application Suite is a set of common applications originally designed to enhance the value of SNA networks for end users. Since more full-featured alternative applications exist in modern integrated SNA/IP networks, z/OS R8 is planned to be the last release of z/OS Communications Server which will include the APPC Application Suite. After z/OS R8 the APPC Application Suite will no longer be shipped with the product, and will not be supported. However, note that APPC itself remains an integral part of z/OS Communications Server's SNA functions, and there are no plans to remove APPC from z/OS. For more information please refer to http://www.ibm.com/software/network/commserver/zos/ .
  Consider the following alternatives to the Communications Server APPC Application Suite:
  - For A3270, consider migrating to TN3270. TN3270 provides a much richer capability assuming IP connectivity exists between the client and server.

© Copyright IBM Corporation, 2007

  - For APING, use the DISPLAY APING command that has been provided as a native VTAM command for many years.
  - A number of other IBM and vendor products provide SNA file transfer capability (such as NetView FTP) which can be used to replace AFTP. TCP/IP's FTP capability is also a good alternative if an IP infrastructure is in place between the client and server.
- z/OS R8 is planned to be the last release to include the C/C++ IBM Open Class (IOC) Dynamic Link Libraries (DLLs). Application development support for the C/C++ IOC Library was withdrawn in z/OS R5. The run-time support (DLLs) for applications that use the IOC Library is planned to be removed in the z/OS release available in 2007. Applications that are dependent on the IOC Library will not run starting with the z/OS release available in 2007. IBM has previously recommended that customers with application code that uses the IOC Library migrate to the Standard C++ Library. The publication *IBM Open Class Library Transition Guide* was published with z/OS 2 V1.2 C/C++ as a reference for customers migrating their code from the IBM Open Class Library to the Standard C++ Library. You can get this guide by visiting **http://www-1.ibm.com/support/docview.wss?rs=32&org=SW&doc=7001423&loc=en-us**

Planned for removal in a future z/OS release

This section lists items that IBM has announced it intends to remove in a future z/OS release. You are encouraged to consider these removals when making your plans for system upgrades. These statements represent IBM's current intentions. IBM development plans are subject to change or withdrawal without further notice.

- In a future release the support for BIND DNS 4.9.3 will be removed from Communications Server. Customers should implement BIND DNS 9.2.0 as a replacement. BIND DNS 9.2.0 is included in the product beginning with z/OS V1.4. Customers exploiting the Connection Optimization (DNS/WLM) feature of BIND 4.9.3 should investigate alternative solutions, such as the Sysplex Distributor function.
- The English and Japanese ISPF panels will be removed from DFSORT in a future release. This limited function interactive facility will no longer be provided, and there will be no replacement.
- From DFSMS, support for the VSAM IMBED, REPLICATE, and KEYRANGE attributes will be withdrawn in a future release. No supported release of z/OS allows you to define new VSAM data sets with these attributes. Using them for existing data sets can waste DASD space and can often degrade performance. When this support is withdrawn, you will not be able to process data sets with these attributes. It is best to plan for this removal now, with the aid of a tool that will help you identified affected data sets.

  A tool is available for download to help identify VSAM data sets that contain the obsolete attributes IMBED, REPLICATE, or KEYRANGE. It is available from the software server (ftp.software.ibm.com) in the s390/mvs/tools directory as IMBDSHIP.JCL.TRSD. This will need to be downloaded in binary format and untersed using TRSMAIN. Note that this tool is provided as is. Instructions for use of the tool are included in the downloaded JCL.

  Note that while no announcement has been made of the specific date or release that data sets with these attributes will no longer be able to be opened, it would be prudent to use this tool to identify the data sets in your installation and begin converting them to versions that do NOT contain these attributes. IMBED and REPLICATE were intended as performance improvements which have been obsoleted by newer cached DASD devices. Striped data sets provide much better performance than KEYRANGE, and should be viewed as a candidate for any existing KEYRANGE data sets.
- In a future release, IBM plans to withdraw support for zFS multi-file system aggregates. When this support is withdrawn, only zFS compatibility mode aggregates will be supported. (A zFS compatibility mode aggregate has a single file system per data set.)
- IBM plans to replace the RMF LDAP backend in a future release of the operating system. The RMF LDAP interface currently allows access to RMF performance data from application programs. This functionality will be replaced with a Common Information Model (CIM) Monitoring interface which has been provided since z/OS R7.

© Copyright IBM Corporation, 2007

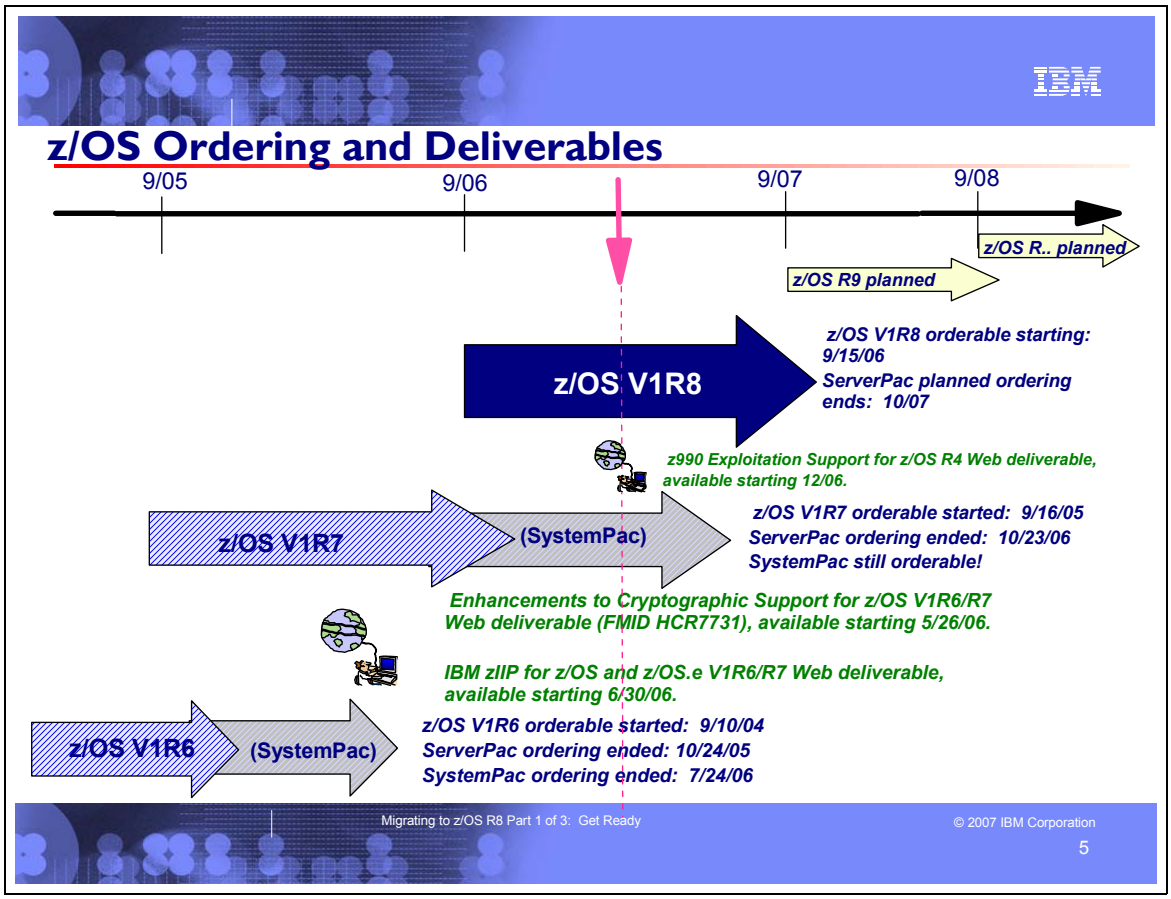

z/OS Ordering and Deliverable Key Dates

Key Dates for z/OS R7 and z/OS R8 ordering:

- **September 16, 2005**: General availability of Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 Web deliverable (FMID HCR7730). This Web deliverable supports z/OS and z/OS.e V1.6 and V1.7.
- **September 30, 2005**: General availability of IBM Health Checker for V1R4/R5/R6 of z/OS and z/OS.e Web deliverable. This Web deliverable supports z/OS and z/OS.e V1.4, V1.5, and V1.6.
- **May 26, 2006:** General availability of the Enhancements to Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 Web deliverable (FMID HCR7731). This Web deliverable supports z/OS and z/OS.e V1.6 and V1.7, and replaced the Cryptographic Support for z/OS V1R6/R7 Web deliverable (FMID HCR7730).
- **June 30, 2006:** General availability of the IBM zIIP for z/OS and z/OS.e V1R6/R7 Web deliverable.
- **September 15, 2006**: First date for ordering z/OS V1R8 ServerPac, SystemPac, CBPDO using CFSW configuration support, or ShopzSeries, the Internet ordering tool. Note that most z/OS media (executable code) is shipped only through Customized Offerings (ServerPac, SystemPac, and CBPDO).
- **September 29, 2006**: z/OS V1R8 general availability via ServerPac, CBPDO and SystemPac.
- **October 10, 2006**: Recommended last date for submitting z/OS R7 orders via the entitled Customized Offerings (ServerPac and CBPDO). This date will allow for adequate order processing time.
- **October 23, 2006**: Last date for ordering z/OS V1R7 via ServerPac and CBPDO.
- **December 2006**: Last date for ordering *z/OS V1R4 z990 Exploitation Support* feature and *z/OS V1R4 Consoles Enhancements* feature. **Note:** The Web deliverable *z990 Exploitation Support for z/OS V1.4 and z/OS.e V1.4* is available as a replacement for the z/OS V1R4 z990 Exploitation Support feature.

Web deliverables

Sometimes enhancements are provided as Web deliverables, and not integrated in your ServerPac, CBPDO, or SystemPac deliverable.

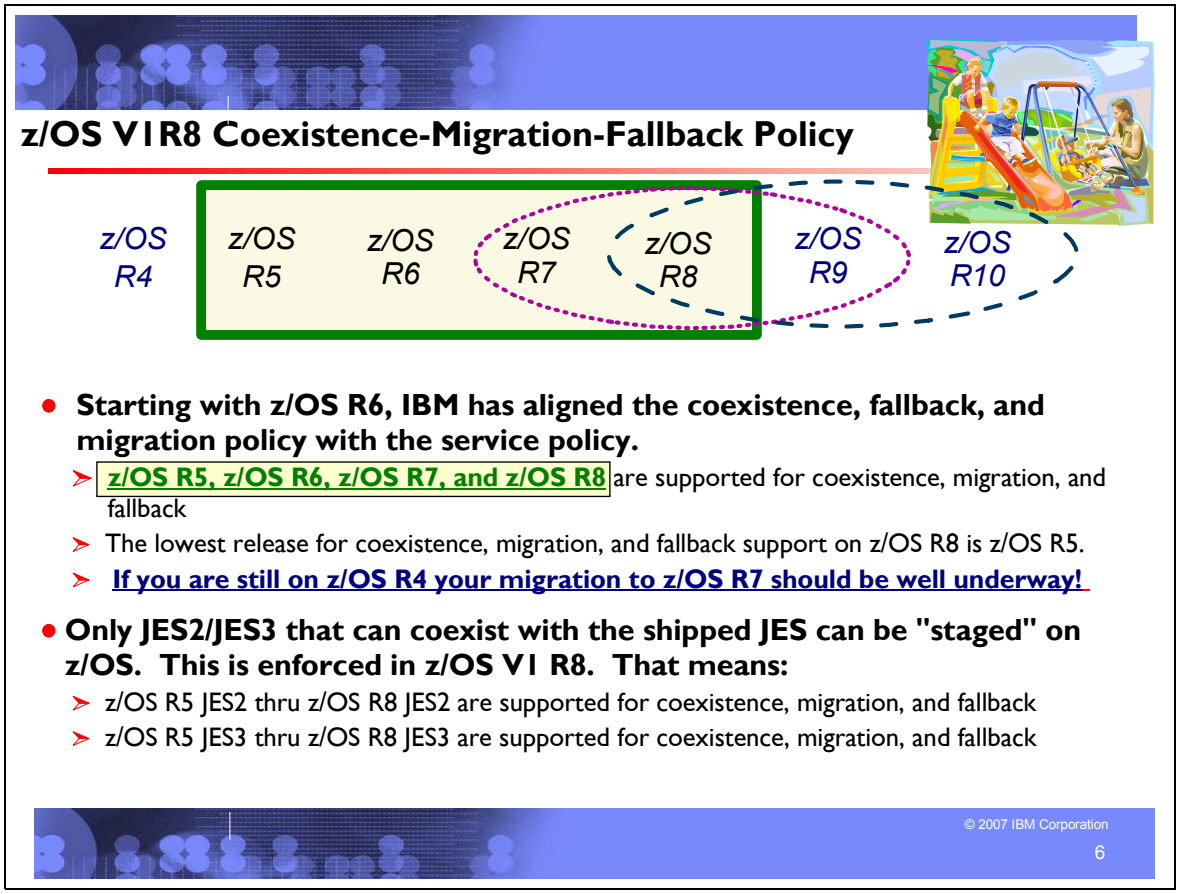

Understanding Coexistence

Coexistence occurs when two or more systems at different software levels share resources. The resources could be shared at the same time by different systems in a multisystem configuration, or they could be shared over a period of time by the same system in a single-system configuration. Examples of coexistence are two different JES releases sharing a spool, two different service levels of DFSMSdfp sharing catalogs, multiple levels of SMP/E processing SYSMODs packaged to exploit the latest enhancements, or an older level of the system using the updated system control files of a newer level (even if new function has been exploited in the newer level).

The sharing of resources is inherent in multisystem configurations that involve Parallel Sysplex implementations. But other types of configurations can have resource sharing too. Examples of configurations where resource sharing can occur are:

- A single processor that is time-sliced to run different levels of the system, such as during different times of the day
- A single processor running multiple images by means of logical partitions (LPARs)
- Multiple images running on several different processors
- Parallel Sysplex or non-Parallel Sysplex configurations

Note: The term coexistence does not refer to z/OS residing on a single system along with VSE/ESA, VM/ESA, or z/VM in an LPAR or as a VM guest.

z/OS systems can coexist with specific prior releases. This is important because it gives you flexibility to migrate systems in a multisystem configuration using rolling IPLs rather than requiring a systems-wide IPL. The way in which you make it possible for earlier-level systems to coexist with z/OS is to install coexistence service (PTFs) on the earlier-level systems. You should complete the migration of all earlier-level coexisting systems as soon as you can. Keep in mind that the objective of coexistence PTFs is to allow existing functions to continue to be used on the earlier-level systems when run in a mixed environment that contains later-level systems. Coexistence PTFs are not aimed at allowing new functions provided in later releases to work on earlier-level systems.

© Copyright IBM Corporation, 2007

Rolling z/OS across a multisystem configuration

A *rolling IPL* is the IPL of one system at a time in a multisystem configuration. You might stage the IPLs over a few hours or a few weeks. The use of rolling IPLs allows you to migrate each z/OS system to a later release, one at a time, while allowing for continuous application availability. For example, data sharing applications offer continuous availability in a Parallel Sysplex configuration by treating each z/OS system as a resource for processing the workload. The use of rolling IPLs allows z/OS systems running these applications to be IPLed one at a time, to migrate to a new release of z/OS, while the applications continue to be processed by the other z/OS systems that support the workload. By using LPAR technology, you can use rolling IPLs to upgrade your systems without losing either availability or capacity.

You can use rolling IPLs when both of the following are true:

- The release to which you're migrating falls is supported for coexistence, fallback, and migration with the releases running on the other systems.
- The appropriate coexistence PTFs have been installed on the other systems in the multisystem configuration.

Even when you're using applications that do not support data sharing, rolling IPLs often make it easier to schedule z/OS software upgrades. It can be very difficult to schedule a time when all applications running on all the systems in a multisystem configuration can be taken down to allow for a complex-wide or Parallel Sysplex-wide IPL. The use of rolling IPLs not only enables continuous availability from an end-user application point of view, but it also eliminates the work associated with migrating all z/OS systems in a multisystem configuration at the same time.

Understanding fallback

Fallback (backout) is a return to the prior level of a system. Fallback can be appropriate if you migrate to z/OS R8 and, during testing, encounter severe problems that can be resolved by backing out the new release. By applying fallback PTFs to the "old" system before you migrate, the old system can tolerate changes that were made by the new system during testing.

Fallback is relevant in all types of configurations, that is, single-system or multisystem, with or without resource sharing. As an example of fallback, consider a single system that shares data or data structures, such as user catalogs, as you shift the system image from production (on the "old" release) to test (on the new release) and back again (to the old release). The later-level test release might make changes that are incompatible with the earlier-level production release. Fallback PTFs on the earlier-level release can allow it to tolerate changes made by the later-level release.

As a general reminder, always plan to have a backout path when installing new software by identifying and installing any service required to support backout.

Fallback is at a system level, rather than an element or feature level, except for z/OS JES2 and z/OS JES3. That is, except for z/OS JES2 and z/OS JES3, you can't back out an element or feature; you can only back out the entire z/OS product. z/OS JES2 and z/OS JES3 fallback can be done separately as long as the level of JES is supported with the release of z/OS and any necessary fallback PTFs are installed.

Fallback and coexistence are alike in that the PTFs that ensure coexistence are the same ones that ensure fallback.
Note: Keep in mind that new functions can require that all systems be at z/OS V1R8 level before the new functions can be used. Therefore, be careful not to exploit new functions until you are fairly confident that you will not need to back out your z/OS V1R8 systems, as fallback maintenance is not available in these cases. You should consult the appropriate element or feature documentation to determine the requirements for using a particular new function.

Which releases are supported for coexistence, fallback, and migration?
Starting with z/OS R6, IBM has aligned the coexistence, fallback, and migration policy with the service policy. IBM intends to continue with the practice of providing service support for each release of z/OS for three years following its general availability (GA) date. IBM, at its sole discretion, may choose to leave a release supported for more than three years. This change to the coexistence, fallback, and migration policy is effective starting with z/OS R6 (that is, the first release of the new annual release cycle). **As a general rule, this means that three releases will be supported for coexistence, fallback, and migration over a period of three years.** This represents an increase of one year over the two-year period provided by the current coexistence, fallback, and migration policy of four releases under a six-month release cycle. The intention of this policy change is to simplify and provide greater predictability to aid in release migrations.
Exceptions are:

- In some cases, more than three releases may be coexistence, fallback, and migration supported if IBM at its sole discretion chooses to provide service support for greater than three years for a release.

© Copyright IBM Corporation, 2007

- Any z/OS release having three or fewer months of service remaining at general availability of a new release will not be coexistence, fallback, or migration supported with the new release.

Thus, except for JES2 and JES3, **z/OS V1R8 is coexistence, fallback, and migration supported with the following three z/OS releases: V1R7, V1R6, and V1R5.** This means that:

- Coexistence of a V1R8 system with a V1R7, V1R6, or V1R5 system is supported.
- Fallback *from* V1R8 *to* V1R7, V1R6, or V1R5 is supported.
- Migration *to* V1R8 *from* V1R7, V1R6, or V1R5 is supported.

For JES2 and JES3, the way in which four consecutive releases is determined is different than for the rest of the operating system. If a JES2 or JES3 release is functionally equivalent to its predecessor (that is, its FMID is the same), then the release is considered to be the same JES release. Thus, **z/OS V1R8 JES2 and JES3 are coexistence, fallback, and migration supported with the following JES releases: V1R8, V1R7, and V1R6-V1R5 (both are functionally equivalent).**

As of z/OS V1R2, compliance to the coexistence, fallback, and migration policy for JES2 and JES3 is enforced. A migration to a JES2 or JES3 release level that is not supported by the policy results in the following:

- For JES2: If the JES2 release level for a system that is initializing is not compatible with the other active systems in the JES2 MAS, message HASP710 is issued and the JES2 address space for the initializing system is terminated.
- For JES3: If the JES3 release level for a local is not compatible with the global in a JES3 multisystem complex, message IAT2640 is issued and the JES3 local is not allowed to connect to the global.

The z/OS coexistence, fallback, and migration policy applies to the elements and features of z/OS, not to customer-developed applications, vendor-developed applications, or IBM products that run on z/OS. IBM performs integration testing and will provide service as necessary to support the z/OS coexistence, fallback, and migration policy.

See the table below for a summary of current and planned coexistence, fallback, and migration support.
These statements represent IBM's current intentions. IBM reserves the right to change or alter the coexistence, fallback, and migration policy in the future or to exclude certain releases beyond those stated. IBM development plans are subject to change or withdrawal without further notice. Any reliance on this statement of direction is at the relying party's sole risk and does not create any liability or obligation for IBM.

Releases that are coexistence, fallback, and migration supported with z/OS V1R6 and beyond

z/OS release (see note 1)	**Releases that are coexistence, fallback, and migration supported with the release in column one (see note 1)**	**Explanation (see note 2)**
R6	R6, R5, R4, R3	The new policy starts. General availability of R6 was September 2004. R2 would be the oldest service-supported release at that time and therefore the oldest release that is coexistence, fallback, and migration supported. However, its end-of-service date (October 2004) is within three months of R6 general availability (September 2004), so R3 becomes the oldest release that is coexistence, fallback, and migration supported with R6.
R7	R7, R6, R5, R4	General availability of R7 was September 30, 2005. R4 is the oldest release that is service supported at that time and therefore the oldest release that is coexistence, fallback, and migration supported with R7.
R8	R8, R7, R6, R5	General availability of R8 was September 29, 2006. R4 is the oldest release that is service supported at that time because its end-of-service date was extended by 18 months to 31 March 2007. **However, R4 is not coexistence, fallback, and migration supported with**

© Copyright IBM Corporation, 2007

		R8. Therefore, R5 is oldest release that is coexistence, fallback, and migration supported with R8.
R9	R9, R8, R7	Planned general availability of R9 is September 2007. R7 is planned to be the oldest release that is service supported at that time and therefore the oldest release that is coexistence, fallback, and migration supported with R9.
R10	R10, R9, R8	Planned general availability of R10 is September 2008. R8 is planned to be the oldest release that is service supported at that time and therefore the oldest release that is coexistence, fallback, and migration supported with R10.

Notes:

1. For readability, the version numbers have been omitted from the releases shown. Also, release numbering of future releases (later than R7) is for illustrative purposes and is not a guarantee of actual release numbers.
2. Future general availability (GA) dates are projections based on the annual release cycle. Future end of service (EOS) dates are projections based on the 3-year service policy.

© Copyright IBM Corporation, 2007

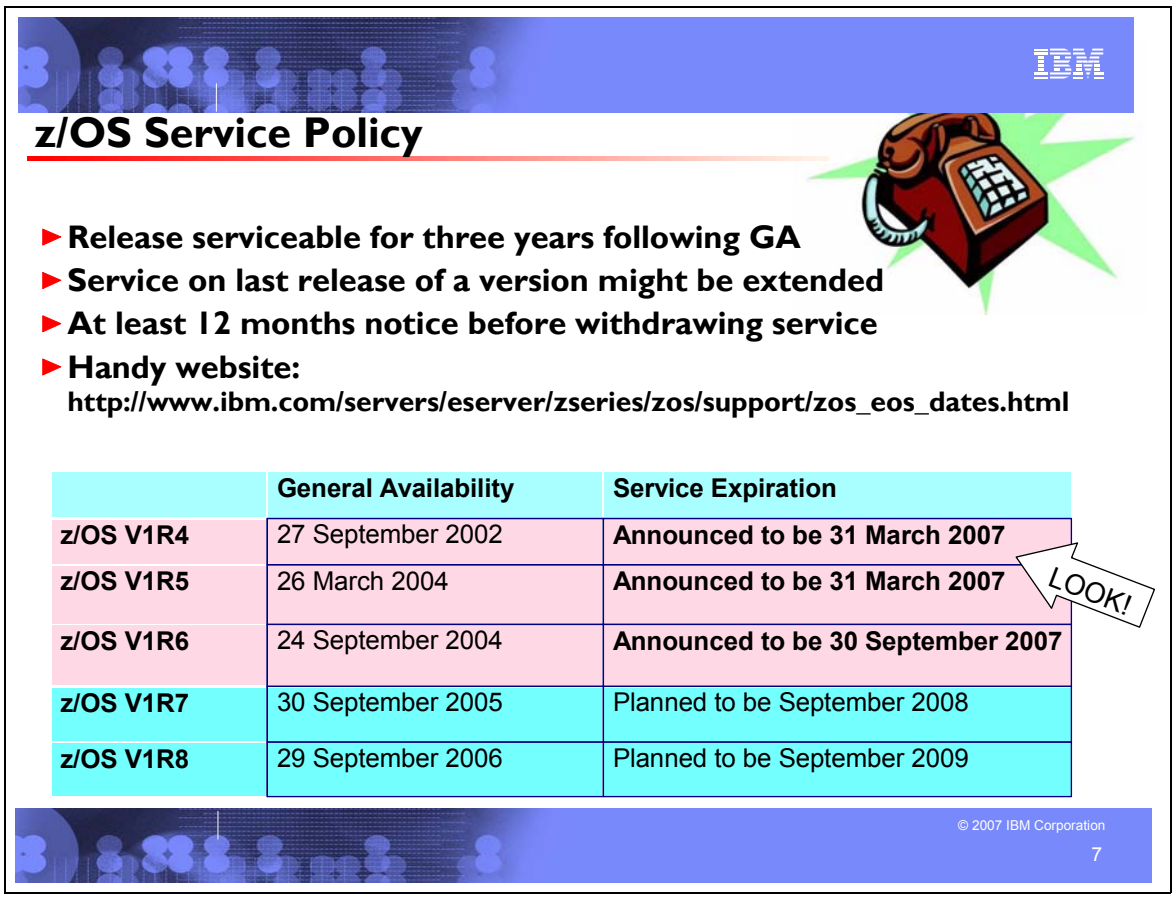

Service Policy

IBM's current policy is to provide maintenance (service) for each release of z/OS for three years following their general availability (GA) date. However, service on the last release of a version might be extended beyond the intended three-year period. Prior to withdrawing service for any version or release of z/OS, IBM intends to provide at least 12 months notice. The service policy for z/OS also applies to any enhancements (including but not limited to web deliverables), such as the z/OS V1R4 enhancements that were provided to support the z990 server.

See the table below for expiration dates for service support. Planned EOS dates are based on the 3-year service policy.

Version and release	General availability (GA)	End of service (EOS)
OS/390 V2R8	24 September 1999	Occurred 30 September 2002
OS/390 V2R9	31 March 2000	Occurred 31 March 2003
OS/390 V2R10	29 September 2000	Occurred 30 September 2004
z/OS V1R1	30 March 2001	Occurred 31 March 2004
z/OS V1R2	26 October 2001	Occurred 31 October 2004
z/OS V1R3	29 March 2002	Occurred 31 March 2005
z/OS V1R4	27 September 2002	31 March 2007 (announced) Note: This is 18 months longer than the normal three-year service period.
z/OS V1R5	26 March 2004	31 March 2007 (announced)
z/OS V1R6	24 September 2004	Announced to be 30 September 2007
z/OS V1R7	30 September 2005	September 2008 (planned)
z/OS V1R8	29 September 2005	September 2009 (planned)

Check out http://www.ibm.com/servers/eserver/zseries/zos/support/zos_eos_dates.html for a handy list of all end of service dates (back to MVS/ESA V4!).

© Copyright IBM Corporation, 2007

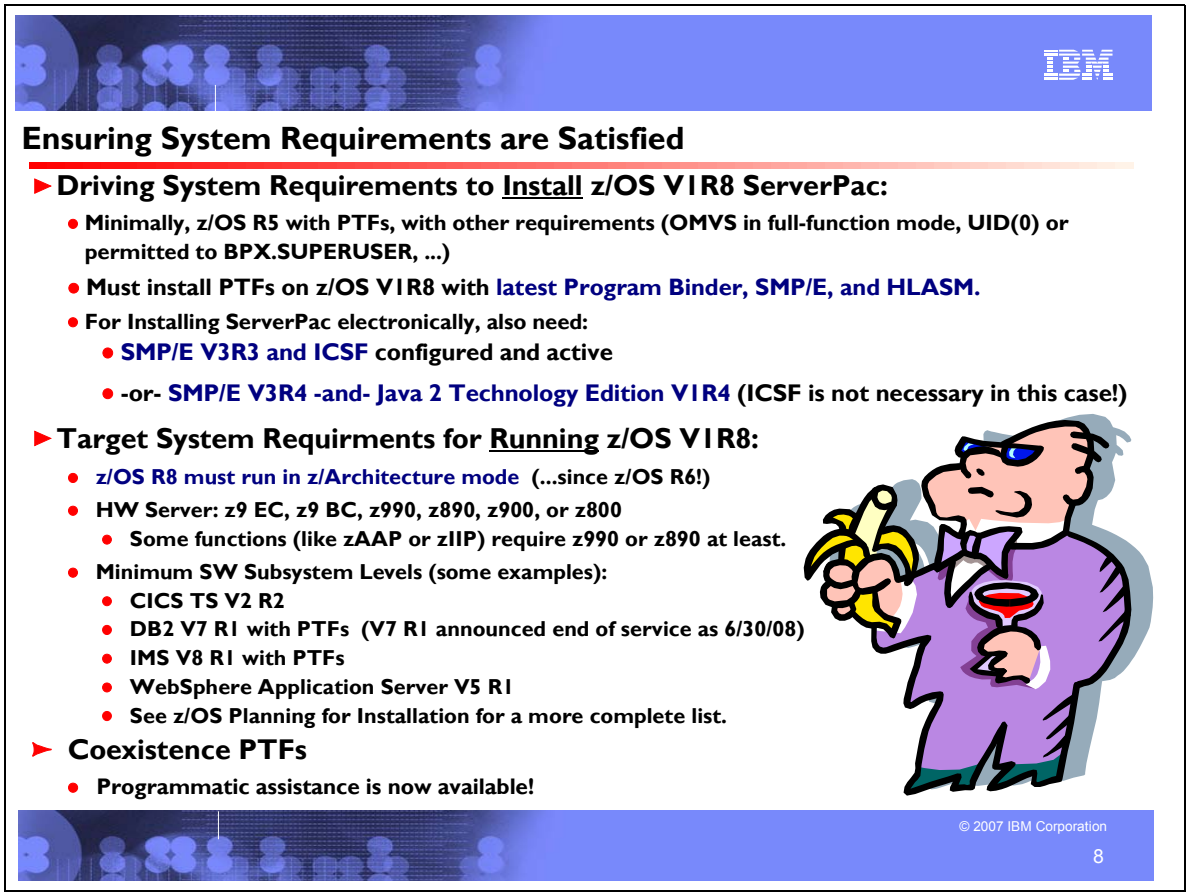

z/OS Driving System Requirements

The *driving system* is the system image (hardware and software) that you use to install the target system. The *target system* is the system software libraries and other data sets that you are installing. You log on to the driving system and run jobs there to create or update the target system. Once the target system is built, it can be IPLed on the same hardware (same LPAR or same processor) or different hardware than that used for the driving system.

If your driving system will share resources with your target system after the target system has been IPLed, be sure to install applicable coexistence service on the driving system before you IPL the target system. If you don't install the coexistence service, you will probably experience problems due to incompatible data structures (such as incompatible data sets, VTOCs, catalog records, GRS tokens, or APPC bind mappings).

Customized Offerings Driver (5665-M12)

The Customized Offerings Driver V2.1 (5665-M12) is an entitled driving system you can use if:

1. you don't have an existing system to use as a driving system, or
2. your existing system does not meet driving system requirements and you don't want to upgrade it to meet those requirements.

At z/OS R8 GA, this driver is a subset of a z/OS R5 system except that the level of SMP/E in the driver is SMP/E V3R4. The Customized Offerings Driver is in DFSMSdss dump/restore format and supports 3380 and 3390 triple-density or higher DASD devices. The Customized Offerings Driver requires a locally attached non-SNA terminal and a system console from the IBM (or equivalent) family of supported terminal types: 317x, 327x, 319x, or 348x. An IBM (or equivalent) supported tape drive is also required to restore the driver.

The Customized Offerings Driver is intended to run in single-system image and monoplex modes only. Its use in multisystem configurations is not supported. The Customized Offerings Driver is intended to be used only to install new levels of z/OS using ServerPac or CBPDO, and to install service on the new software until a copy (clone) of the new system can be made. The use of the Customized Offerings Driver for other purposes is not supported.

The Customized Offerings Driver includes an HFS and the necessary function to use Communications Server (IP Services), Security Server, and the system-managed storage (SMS) facility of DFSMSdfp, but these items are not

customized. However, existing environments can be connected to, and used from, the Customized Offerings Driver system.

Identifying Driving System Software Requirements for ServerPac for z/OS R8

Driving system requirements for installing z/OS R8 by way of ServerPac or dump-by-data-set SystemPac are:

- *An operating system:* Use any of the following:
  - z/OS V1R5 with the PTFs in the table below.
  - The Customized Offerings Driver (5665-M12).
- *A terminal:* A locally-attached or network-attached terminal that can be used to establish a TSO/E session on the IPLed system is required.
- *Proper authority:* Use the RACFDRV installation job as a sample of the security system definitions required so that you can perform the installation tasks.
- Proper security:
  - In order for you to install into the zFS, the user ID you use must have read access to the SUPERUSER.FILESYS.PFSCTL resource in the RACF FACILITY class.
  - In order for you to install into the HFS, the following is required:
    - ☐ The user ID you use must be a superuser (UID=0) or have read access to the BPX.SUPERUSER resource in the RACF facility class.
    - ☐ The user ID you use must have read access to facility class resources BPX.FILEATTR.APF, BPX.FILEATTR.PROGCTL, and BPX.FILEATTR.SHARELIB (or BPX.FILEATTR.* if you choose to use a generic name for these resources). The commands to define these facility class resources are in SYS1.SAMPLIB member BPXISEC1 and in the z/OS UNIX Customization Wizard (http://www.ibm.com/eserver/zseries/zos/wizards/).
    - ☐ Group IDs uucpg and TTY, and user ID uucp, must be defined in your security database. These IDs must contain OMVS segments with a GID value for each group and a UID value for the user ID. (For ease of use and manageability, define the names in uppercase.)
      - The group ID and user ID values assigned to these IDs cannot be used by any other IDs. They must be unique.
      - You must duplicate the required user ID and group names in each security database, including the same user ID and group ID values in the OMVS segment. This makes it easier to transport the HFS data sets from test systems to production systems. For example, the group name TTY on System 1 must have the same group ID value on System 2 and System 3. If it is not possible to synchronize your databases you will need to continue running the FOMISCHO job against each system after z/OS UNIX is installed.

      If names such as uucp, uucpg, and TTY are not allowed on your system, or if they conflict with existing names, you can create and activate a user ID alias table. For information about defining these group and user IDs to RACF and about creating a user ID alias table (USERIDALIASTABLE), see *z/OS UNIX System Services Planning*. Other sources of information are SYS1.SAMPLIB member BPXISEC1 and the z/OS UNIX Customization Wizard (http://www.ibm.com/eserver/zseries/zos/wizards/). (**Note:** You can use the RACFDRV installation job as a sample of the security system definitions required to perform the installation tasks.)
- *Language Environment run-time options:* As of z/OS R7, ServerPac requires that the following Language Environment run-time options are **not** specified as nonoverrideable (NONOVR) in the CEEDOPT CSECT: ALL31, ANYHEAP, BELOWHEAP, DEPTHCONDLIMIT, ERRCOUNT, HEAP, HEAPCHK, HEAPPOOLS, INTERRUPT, LIBSTACK, PLITASKCOUNT, STACK, STORAGE, THREADHEAP, and THREADSTACK .
- *Language Environment:* The CustomPac Installation Dialog uses the Language Environment run-time library SCEERUN. If SCEERUN is not in the link list on the driving system, you must edit the ServerPac installation jobs to add it to the JOBLIB or STEPLIB DD statements.
- *OMVS address space active:* For ServerPac only (not SystemPac), an activated OMVS address space with z/OS UNIX kernel services operating in full function mode is required.
- *SMS active:* The Storage Management Subsystem (SMS) must be active to allocate HFS and PDSE data sets, whether they are SMS-managed or non-SMS-managed. Also, the use of HFS data sets is supported only when SMS is active in at least a null configuration, even when the data sets are not SMS-managed. Do either of the following:
  - ☐ To allocate non-SMS-managed HFS and PDSE data sets, you must activate SMS on the driving system in at least a null configuration. You must also activate SMS on the target system.
  - ☐ To allocate SMS-managed HFS and PDSE data sets, you must activate SMS on the driving system in at least a minimal configuration. Then you must define a storage group, create SMS-managed volumes, and write, translate, and activate a storage class ACS routine that allows the allocation of PDSE and HFS data sets with the names in the ALLOCDS job. You must also activate SMS on the target system.

© Copyright IBM Corporation, 2007

- *SMP/E ++JAR Support:* If your ServerPac order contains any product that uses the ++JAR support introduced in SMP/E V3R2, then your driving system will require IBM Developer Kit for OS/390, Java 2 Technology Edition (5655-D35). z/OS V1R8 itself does not use the ++JAR support.
- *zFS configuration requirements (optional):* If you will specify that you will use a zFS for ServerPac installation, then you must be sure that the zFS has been installed and configured, as described in *z/OS Distributed File Service zSeries File System Administration*.
- *Electronic delivery requirements (optional):* If you intend to receive your ServerPac order by way of the Internet, you need either of the following:
  - SMP/E V3R3 **and** Cryptographic Services ICSF set up and activated
  - *-or-* SMP/E V3R4 **and** Java 2 Technology Edition V1R4 (in this case ICSF isn't necessary).
- *Proper level for service:* In order for you to install service on the target system that you're building, your driving system must minimally meet the driving system requirements for CBPDO Wave 1 and must have the current (latest) levels of the program management binder, SMP/E, and HLASM. The service jobs generated by the CustomPac Installation Dialog use the target system's (and therefore current) level of the binder, SMP/E, and HLASM. If you choose to use your own jobs, model them after the jobs provided by ServerPac or dump-by-data-set SystemPac by adding STEPLIB DD statements to access MIGLIB (for the binder and SMP/E) and SASMMOD1 (for HLASM). Be sure that the SASMMOD1 and SYS1.MIGLIB data sets are APF authorized. Another way to install service is from a copy of your target system.

Driving System PTFs for ServerPac for z/OS R8

z/OS release	PTFs
z/OS V1R5	• DFSMSdfp: UA05520, UA14837, and UA16273 • z/OS UNIX: UQ79726, UQ85020 For servicing: z/OS R8 level of the Program Management Binder, HLASM (V1R5). and z/OS R8 SMP/E (SMP/E V3R4)
z/OS V1R6	• DFSMSdfp: UA14838 and UA16274 For servicing: z/OS R8 level of the Program Management Binder, HLASM (V1R5), and z/OS R8 SMP/E (SMP/E V3R4)
z/OS V1R7	For servicing: z/OS R8 level of the Program Management Binder, HLASM (V1R5), and z/OS R8 SMP/E (SMP/E V3R4)

Identifying Driving System Software Requirements for Servicing z/OS R8 and Installing CBPDO

After installing via ServerPac, you will eventually want to install service on your z/OS R8 system. The driving system to install service on your z/OS R8 ServerPac system are the same as to install with the CBPDO method. For servicing your ServerPac system, you must meet the driving system requirements for "Wave 1" below.

When you use the CBPDO method of installing z/OS you install in three stages called *waves*. (Wave 1, in order to be more manageable, is divided into several tasks called *ripples*.) This section describes the driving system requirements for each wave.

z/OS R8 CBPDO: Driving System Wave 0

In Wave 0 you install the program management binder (part of the BCP), the HLASM base element, and the SMP/E base element. These items must be installed on (available from) the driving system for subsequent wave installations.
You can use either of the following as the driving system for installing z/OS R8 Wave 0:

- z/OS R5 or later.
- The Customized Offerings Driver (5665-M12).

z/OS R8 CBPDO: Driving System Wave 1 (and to service a ServerPac system)

In Wave 1 you install most of the elements and features. The driving system requirements for installing Wave 1 are:
An operating system: Use any of the following:

- *z*/OS R5 or later with the PTFs listed in the ServerPac table previously, except that the **program binder**, **HLASM**, and **SMP/E** must be at the current (latest) levels. To satisfy the binder and SMP/E requirements, you can use a STEPLIB DD statement to access the z/OS V1R8 program binder and z/OS V1R8 SMP/E in the Wave 0 target system's SYS1.MIGLIB data set, and to access the z/OS V1R8 HLASM's ASM.SASMMOD1 data set. Ensure that the target system's SYS1.MIGLIB and ASM.SASMMOD1 data sets are APF authorized on the driving system.

© Copyright IBM Corporation, 2007

- Since z/OS R6, failure to use the latest program binder may result in errors during APPLY of z/OS service. You may see such errors as:

```
IEW2322I 1220 5 INCLUDE SMPWRK3(CELQAWI) UQ88825 SEQ # 000025
IEW2690E 3530 ONE OR MORE FIELD DESCRIPTORS IN GOFF RECORD 3 WITHIN MEMBER CELQAWI
IDENTIFIED BY DDNAME SMPWRK3 ARE NOT VALID. ERRORID = 40.
IEW2307E 1032 CURRENT INPUT MODULE NOT INCLUDED BECAUSE OF INVALID DATA.
```

- The Customized Offerings Driver (5665-M12).
- *Proper security:* In order for you to install into the HFS, the security definitions described previously for ServerPac are required.
- *Binder entry in SMP/E UTILITY:* The SMP/E UTILITY must have an entry for the binder. You can specify any of these program names in the UTILITY entry: IEWBLINK, HEWL, IEWL, LINKEDIT, or HEWLH096. (The linkage editor, which uses the names HEWLKED, HEWLF064, IEWLF440, IEWLF880, and IEWLF128, cannot be used.)
- *Language Environment:* You must add SCEERUN (the run-time library provided by Language Environment) to your program search order because many elements and features require it. If you wish, add SCEERUN to your LNKLST concatenation. If you don't add SCEERUN to your LNKLST, you must access SCEERUN by using STEPLIB DD statements in the individual element and feature procedures that require them. The BCP's Program Management Binder is one function that requires access to SCEERUN; it has been required since OS/390 R10. This means that you must make available SCEERUN (via the LNKLST or STEPLIB, for instance) to any JCL and procedures (such as SMP/E procedures) that invoke the binder. This ensures that processing, such as conversion of long names to short names, is successful.
- *OMVS address space active:* Before you install the elements and features in Wave 1, you must activate the OMVS address space in full function mode on the driving system. To activate OMVS, complete the required customization (for example, SMS and RACF setup) as described in *z/OS UNIX System Services Planning*.
- *Target system's HFS mounted:* The target system's HFS must be mounted.

z/OS R8 CBPDO: Driving System Wave 2

In Wave 2 you install the z/OS V1R8 level of JES2 or JES3. Wave 2 is optional and can be combined with Wave 1. The driving system requirements for Wave 2 are the same as for Wave 1.

Choosing IBM Products That You Want to Run with z/OS

The functional and minimal requirements are found in the *z/OS and z/OS.e Planning for Installation* which is available on the Internet from *http://www.ibm.com/servers/eserver/zseries/zos/bkserv/*.

For a list of products available for ordering with z/OS, you can do any of the following:

- Use the self-service Internet application ShopzSeries: *http://www.ibm.com/software/shopzseries*
- Use the product catalogs (order checklists). For z/OS *http://www.ibm.com/eserver/zseries/software/swinfo/os390.htm*
- Access the software configurator used in your country, select the z/OS environment, and then select ServerPac, CBPDO, or SystemPac.
- Refer to *z/OS and z/OS.e Planning for Installation Appendix C, Minimum Releases of IBM Software Products That Run with z/OS and z/OS.e*

Many of these products can be ordered as part of your z/OS ServerPac order, z/OS SystemPac order, separate CBPDO order, or separate ProductPac order.

If you're migrating to z/OS V1R8, you can find out which products have new levels by using ShopzSeries or by using the SMP/E base element's Planning and Migration Assistant. Both tools use data found on your system as well as the latest IBM software product catalog.

Choosing ISV products that you want to run with z/OS

For a list of independent software vendors (ISVs) that support z/OS, as well as announcements, testimonials, and other information, see http://www.ibm.com/servers/eserver/zseries/zos/zos_products.html

For a directory of IBM and IBM Business partners that provide z/OS applications, tools, and services, see the Global Solutions Directory: http://www.ibm.com/software/solutions/isv
Note that you can use SystemPac to get some ISV products on system delivery vehicles with IBM products. For a list of selectable ISVs with SystemPac, see: http://www.ibm.com/ca/custompac/

© Copyright IBM Corporation, 2007

VM Guest Considerations

z/OS can run as a VM guest. Requirements are:

- Must be on a z9 EC, z9 BC, z990, z890, z900, or z800 server: VM must be z/VM (not VM/ESA). Furthermore,
  - if a z9 EC or z9 BC or z990 server, the z/VM PTF for APAR VM63124 must be installed.
  - If a z800, z890, or z9 BC server, then z/OS.e can run as a VM guest. The server must be a z800, z890, or z9 BC, VM must be z/VM, z/VM and z/OS.e must operate in z/Architecture (64-bit) mode, z/VM and z/OS.e must run in a logical partition (LPAR), and the LPAR must be named ZOSE*xxxx* (where *xxxx* is any valid combination of zero to four characters). The LPAR name restriction ensures compliance with z/OS.e license terms, which prohibit z/OS from running in a z800, z890, or z9 BC LPAR named ZOSE*xxxx*. (Linux for S/390 and Linux for IBM @server zSeries are the only other operating systems that may run as a guest under z/VM in an LPAR named ZOSE*xxxx*.) z/OS.e licensees are required to submit data to IBM that shows LPAR names. (The ZOSE*xxxx* LPAR name restriction is only applicable to z800, z890, and z9 BC servers, not other servers. The Dynamic LPAR Rename function available on z890 and z9 BC can assist you with this requirement.)

Target System Hardware Requirements

The minimal hardware requirements for z/OS, as well as additional hardware needed by specific z/OS elements and features is documented in *z/OS and z/OS.e Planning for Installation*. **Remember, z/OS V1R8 runs only in z/Architecture mode, and only on z9 EC , z9 BC, z990, z890, z900 or z800 servers.**

Identifying Processor Requirements

z/OS V1R8 runs on the following IBM @server zSeries servers (or their equivalents):

- IBM zSeries z9 EC.
- IBM zSeries z9 BC.
- IBM zSeries 990 (z990).
- IBM zSeries 890 (z890).
- IBM zSeries 900 (z900).
- IBM zSeries 800 (z800).

Identifying Coupling Facility Requirements

There are hardware and software requirements related to coupling facility levels (CFLEVELs). See http://www.ibm.com/eserver/zseries/pso/cftable.html

Selected IBM Subsystem Products that will run on z/OS R8 (and levels for zAAP and zIIP)

For a complete list of the minimum release of IBM software products that run on z/OS R8, refer to Appendix C of *z/OS and z/OS.e Planning for Installation.*

zAAP

If you are interested in exploiting (or doing capacity planning for) zAAPs, then there are minimum software levels required. In order to exploit a zAAP, the operating system must be migrated to at least z/OS 1.6 (or z/OS.e 1.6), the IBM Solution Developers Kit (SDK) for z/OS, Java 2 Technology Edition, V1.4 with PTF (or later) for APAR PQ86689 must be used.

Here is a table outlining the subsystems and minimum Java levels dependencies for determining zAAP Java execution and exploitation potential. The rows in blue identify some of the products levels that will exploit zAAPs. The rows in yellow, while not able to exploit the benefits of zAAPs, show where the zAAP Projection Tool for Java 2 Technology Edition, SDK1.3.1 can be used to assist in zAAP capacity planning. For instance, WebSphere V5.0.2 cannot exploit zAAP, however the zAAP Projection Tool can be used to determine zAAP Java execution potential if the WebSphere workloads were migrated to the required level for zAAP exploitation.

zIIP

For z/OS V1.6 or z/OS V1.7: If you will exploit the IBM System z9 Integrated Information Processor (zIIP), you will require the "IBM System z9 Integrated Information Processor Support for z/OS and z/OS.e V1R6/R7" Web deliverable plus PTFs. This Web deliverable contains FMID JBB77S9 for z/OS V1.6, and FMID JBB772S for z/OS V1.7. Install the service identified in the program directory that comes with this Web deliverable and the service which is recommended in the appropriate zOSV1Rn BCPZIIP PSP bucket. IBM DB2 UDB for z/OS V8 with PTFs, is also required for exploitation for portions of eligible workloads. For z/OS V1.8, this Web deliverable is not needed, as zIIP support is incorporated.

© Copyright IBM Corporation, 2007

EPSPT Programmatic Help With Coexistence PTFs

- **Finally, for z/OS R8! Some programmatic assistance to help you determine that you've got all the Coexistence PTFs installed for a z/OS migration!**
- This solution uses the existing Electronic PSP Tool (EPSPT).
  - "Extract files" are now produced from the ZOSGEN PSP bucket that contain the Coexistence PTFs.
  - You should use the Extract file for the z/OS release you are migrating *FROM*.
    - For instance, the ZOSV1R7 ZOSGEN Extract file contains the PTFs that you need to coexist with z/OS R8.
- The Coexistence PTFs from the z/OS Migration book, along with any subsequently defined coexistence PTFs are included, for a complete Coexistence PTF list.

© 2007 IBM Corporation

EPSPT Programmatic Help With Coexistence PTFs

- **Steps to follow:**

1. If you aren't already using EPSPT, download and install the EPSPT, available from http://www14.software.ibm.com/webapp/set2/sas/f/psp/download.html.
2. Download the Extract file from your current release's ZOSGEN PSP bucket subset.
   - The list of "to" release coexistence PTFs is found in the "from" release ZOSGEN PSP bucket subset.
   - You may already be pulling the Extract file today, but perhaps not from the ZOSGEN subset.
3. Run EPSPT using the Extract file from your current release's ZOSGEN PSP bucket.
4. Review the EPSPT output report, and resolve any outstanding discrepancies.

- **Periodically, rerun the steps above for any new Coexistence PTFs that may have been added.**

2007 IBM Corporation

© Copyright IBM Corporation, 2007

Programmatic Help with Coexistence PTFs for z/OS Release 8

You may use the Enhanced PSP Tool (EPSPT) to assist in programmatically determining which coexistence PTFs you must install on your current system in preparation for migration to a later z/OS release. When you retrieve the extract file from the ZOSGEN PSP bucket subset that is used as input to EPSPT, the file will contain the current list of coexistence PTFs for migrating to a later z/OS release. Coexistence PTFs are identified in the "Cross Product Dependencies" section of the ZOSGEN PSP bucket subset.

The steps to take to programmatically determine whether your current system has the complete list of required Coexistence PTFs for migration are:
1) Download and install the EPSPT, available from http://www14.software.ibm.com/webapp/set2/sas/f/psp/download.html.
2) Download the extract file from your current release's ZOSGEN PSP bucket subset. The list of "to" release coexistence PTFs is found in the "from" release ZOSGEN PSP bucket subset.
3) Using the extract file from your current release's ZOSGEN PSP bucket subset, run the EPSPT.
4) Resolve any outstanding discrepancies that the EPSPT has identified.

From time to time, you may want to download the extract file from your current release's ZOSGEN PSP bucket subset, and rerun EPSPT to ensure that any newly added coexistence PTFs are verified.

An example of the beginning of the ZOSGEN PSP extract file (from the ZOSV1R7 upgrade) is:

```
/* Preventive Service Planning */
/* UPGRADE: ZOSV1R7, SUBSET: ZOSGEN */
/* Updates also available at ftp site: */
/* ftp site  = ftp://ftp.software.ibm.com/s390/pspapartool/ */
/* ftp file  = ZOSV1R7_ZOSGEN.txt  */
/* Last Extract: */
/*   Time = 20:01:48 */
/*   Date = 01/20/2007 */
APAR(AA10632) FMID(HBB7720) FIX(UA26398) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA13294) FMID(HBB7720) FIX(UA25784) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA13438) FMID(HBB7720) FIX(UA24560) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA13837) FMID(HBB7720) FIX(UA90269) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA14009) FMID(HBB7720) FIX(UA24159) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA15144) FMID(HBB7720) FIX(UA27072) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA15835) FMID(HBB7720) FIX(UA25588) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA17662) FMID(HBB7720) FIX(UA29484) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA14334) FMID(HCS7720) FIX(UA90273) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA12683) FMID(HDZ11K0) FIX(UA24413) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA13177) FMID(HDZ11K0) FIX(UA25405) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA13332) FMID(HDZ11K0) FIX(UA21542) UPG(ZOSV1R7) SUB(ZOSGEN).
APAR(AA13355) FMID(HDZ11K0) FIX(UA22045) UPG(ZOSV1R7) SUB(ZOSGEN).
.....
```

Coexistence Service Requirements for z/OS Release 8

To understand the Coexistence-Fallback-Migration policy for z/OS, and which releases are supported with z/OS V1 R8, see the Coexistence section of this presentation. Listed below are the coexistence and fallback service, for installation on an z/OS V1R7 system for z/OS V1R8. For a list of the coexistence service on z/OS V1R6 and z/OS V1R5 for z/OS V1R8, see Appendix B of this presentation.

Apply on z/OS V1R7, the coexistence and fallback service (PTFs) listed in the table below:

Element or feature, and function provided by PTFs	z/OS V1R7 PTFs

© Copyright IBM Corporation, 2007

BCP: Ensures that pre-z/OS V1R8 systems tolerate the new format of the z/OS V1R8 **WLM** service definition and service policy in the WLM couple data set. Moreover, the PTF ensures the toleration of changed WLM LPAR management and JES initiator management information that is exchanged in a mixed-release (z/OS V1R8 and pre-z/OS V1R8) sysplex. If the WLM service policy in the couple data set has been activated by a z/OS V1R8 system, a pre-z/OS V1R8 system without the PTF will not be able to activate the policy and might fail during IPL with a wait state 064. If LPAR management is being used in a mixed-release sysplex and the PTF is not installed, you might experience abends and see message IWM054I FAILURE IN LPAR CPU MANAGEMENT, PROCESSING DISABLED.	UA90269 (if you do not have JBB772S installed) or UA90270 (if you have JBB772S installed)
BCP: Allows the z/OS V1R7 program management loader to properly handle module aliases produced by the S/360 linkage editor	UA29484
BCP: In z/OS V1R8, SYSIN DD * override statements are no longer required to reside in the job stream in the same order as the corresponding DD statements appear in the procedure. Because new **text units** are introduced for each SYSIN DD * statement, the PTF allows jobs to ignore the new text units when they run in a JES2 environment and convert on a z/OS V1R8 system but interpret on an earlier system.	UA24560
BCP: If an application developer compiles or assembles with the z/OS V1R8 program management **binder** version 6 API invocation macros, or codes to the new STARTD API interface, and the resulting program is run on an earlier release of z/OS, the binder STARTD API will fail with return code 12, invalid parameter list, message IEW2132S 0064 INCORRECT NUMBER OF PARAMETERS PASSED FOR CALL. The PTF allows such a calling program to execute successfully.	UA24437
BCP: If an application developer uses the z/OS V1R8 program management **binder** to create a program object with the latest format in a PDSE, the PTF is needed so that an ISPF user can browse the PDSE on z/OS V1R7.	UA25784
BCP: In z/OS V1R8, the **console restructure** enhancement eliminates the master console and the console switch function. The PTF allows earlier systems to tolerate the enhancement.	UA26398 and, if Japanese, UA26402
BCP: Allows a z/OS V1R7 system to tolerate the z/OS V1R8 **logger** log stream GROUP designation and renamed log streams.	UA24159, UA25588
BCP: Allows a z/OS V1R7 system to tolerate changes made to resource recovery services (**RRS**) display and ATRQUERY functions in z/OS V1R8.	UA27072
DFSMSdfp: Program object format 5 (PO5 format) is new in z/OS V1R8. The PTF allows PDSE program libraries that contain **PO5** program objects to be processed by **IEBCOPY** on z/OS V1R7 systems.	UA22351
DFSMSdfp: An enhancement in the z/OS V1R8 object access method (**OAM**) component introduced a new ODLOBFL column in the object directory to exploit DB2(R) binary large object (BLOB) support. The PTF enables pre-z/OS V1R8 OAMs to coexist in an OAMplex with OAMs at the z/OS V1R8 or later level.	UA24413
DFSMSdfp: In z/OS V1R8, the D SMS,SMSVSAM,DIAG(CONTENTION) command is introduced to display internal **VSAM RLS** contention. The PTF allows a z/OS V1R7 system that receives the new DIAG CONTENTION command to	UA24374

© Copyright IBM Corporation, 2007

tolerate the command by issuing a message analogous to IGW495I D SMS,SMSVSAM,DIAG COMMAND FUNCTION HAS FOUND THE DIAG TABLE HAS NO ENTRIES. If the PTF is not installed, the z/OS V1R7 system receiving the new command treats it as an irrelevant and confusing console message.	
DFSMSdfp: In z/OS V1R8, a performance improvement to **VSAM RLS** causes XES lock table cleanup to be skipped when the SMSVSAM address space is terminated, as documented in APAR OA11708. This PTF allows a z/OS V1R7 system to tolerate the enhancement.	UA21542
DFSMShsm(TM): Allows a z/OS V1R7 system to tolerate the changes made by the **fast replication tape support** enhancement in z/OS V1R8.	UA25405
DFSMShsm: In z/OS V1R8, dump volume (**DVL**) and dump class (**DCL**) control data set records were enlarged as part of the DFSMShsm dump encryption enhancement. The PTF allows a z/OS V1R7 system to tolerate the increase in record size.	UA22716
DFSMSrmm(TM): Allows a pre-z/OS V1R8 system in a client/server environment to participate in an **RMMplex** in which a z/OS V1R8 system has been started. Optionally, you can specify UTC(YES) with the EDGUTIL utility on z/OS V1R8 to enable DFSMSrmm common time support.	UA25145 if not using client/server DFSMSrmm, or both UA25145 and UA00015 if using client/server DFSMSrmm
DFSMSrmm: Allows a pre-z/OS V1R8 system to participate in an **RMMplex** in which a z/OS V1R8 system has been started. In particular, the PTF enables VRSEL processing to recognize the use of the special ABEND and OPEN vital record specifications and to correctly handle COUNT(0). It also prevents the continued use of the unsupported options STARTNUMBER and LOCATION(BOTH).	UA22045
Distributed File Service: Allows z/OS V1R7 systems to tolerate the change in the **zFS XCF** protocol that occurred in z/OS V1R8 (and was implemented in z/OS V1R7 by APAR OA13738).	UA24209
HCD: Allows z/OS V1R8 **HCM** to exploit z/OS V1R7 HCD functionality to enable new z/OS V1R8 HCM functions.	UA90273 and, if Japanese, UA90274
JES2: Allows an earlier JES2 to run with z/OS V1R8 JES2 in a multi-access spool (**MAS**). Install the PTFs that are appropriate to the JES2 level that you are using.	**z/OS V1R7 JES2:** UA24984; **z/OS V1R5-V1R6 JES2:** UA09501, UA18836, UA20902, UA90161, UA24983
JES3: Allows a pre-z/OS V1R8 JES3 to tolerate the removal of various symbols from the **IHAWQE macro**, which was done as part of the console improvements made in z/OS V1R8.	**z/OS V1R7 JES3:** UA24990; **z/OS V1R5-V1R6 JES3:** UA24989
Language Environment(R): Causes a **meaningful message** to be issued when a new z/OS V1R8 Language Environment run-time option is encountered by an earlier release. The message states that the option is ignored. If the PTF is not installed, the	UK12452

© Copyright IBM Corporation, 2007

option is still ignored but a less meaningful message, or no message at all, is issued.	
NFS: Allows a z/OS V1R7 system to read the new **NFS export file** format generated by a z/OS V1R8 system.	UA25099
RMF(TM): Ensures the toleration of new z/OS V1R8 functionality for sysplex-wide **RMF reporting** in earlier releases.	UA26170
Security Server: In z/OS V1R8, the structure of the **RACF database** was changed to allow a template block to be continued into another block. The PTF allows a z/OS V1R7 system to process templates that extend into another block. If you share a RACF database between z/OS V1R8 and z/OS V1R7, you need to install the PTF on the z/OS V1R7 system.	UA24980

© Copyright IBM Corporation, 2007

IBM Migration Checker for z/OS

- **Some Common Migration Problems:**
  - We have many migration actions required for each release.
  - Applicability of migration actions can be difficult to determine.
  - Determination of the applicability is (in many cases) manual, when it could be programmatic.
  - When programmatic, disparate solutions have been used.
  - Hard to verify if a migration action has been done, especially when deploying to subsequent systems.
  - Would like some way of checking pre-install and post-install "migration health" of a system.
- **An "As Is" Tool That May Help - *IBM Migration Checker for z/OS***
  - From the z/OS Download website: http://www.ibm.com/servers/eserver/zseries/zos/downloads/
  - This tool can:
    - Alert you to some migration actions that you should plan for in the future
    - Detect some migration actions that you can do now to ease your migration to z/OS V1R8
    - Verify that you've done some migration actions correctly when you're running z/OS V1R8
  - Tool verified by 52 early customers. Added 14 of their suggested enhancements.

© 2007 IBM Corporation

IBM

Programs included in the IBM Migration Checker for z/OS

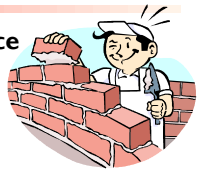

1. **TN3270AS - Run the TN3270E Telnet server as a separate address space**
2. **HFSZFS - Migrate from HFS file systems to zFS file systems**
3. **ZFSMULT - Discontinue use of multi-file system aggregates**
4. **BTSIZE - Use the new default for BLOCKTOKENSIZE in IGDSMSxx**
5. **OLDVSAM - Redefine existing VSAM data sets that contain the IMBED, REPLICATE, and KEYRANGE attributes**
6. **MASTRCON - Accommodate the removal of the master console**
7. **ONEBCON - Accommodate the removal of 1-byte console IDs**
8. **LLSTCAT - Catalog SYS1.SIEALNKE and SYS1.SIEAMIGE for the default system link list in z/OS V1R8**
9. **NEWDS - Add references to new data sets and paths**
10. **OLDDS - Remove references to old data sets and paths**
11. **SMSLVL - Ensure the integrity of SMS control data sets**
12. **ETCUSS - Update Communications Server and z/OS UNIX /etc configuration files**

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

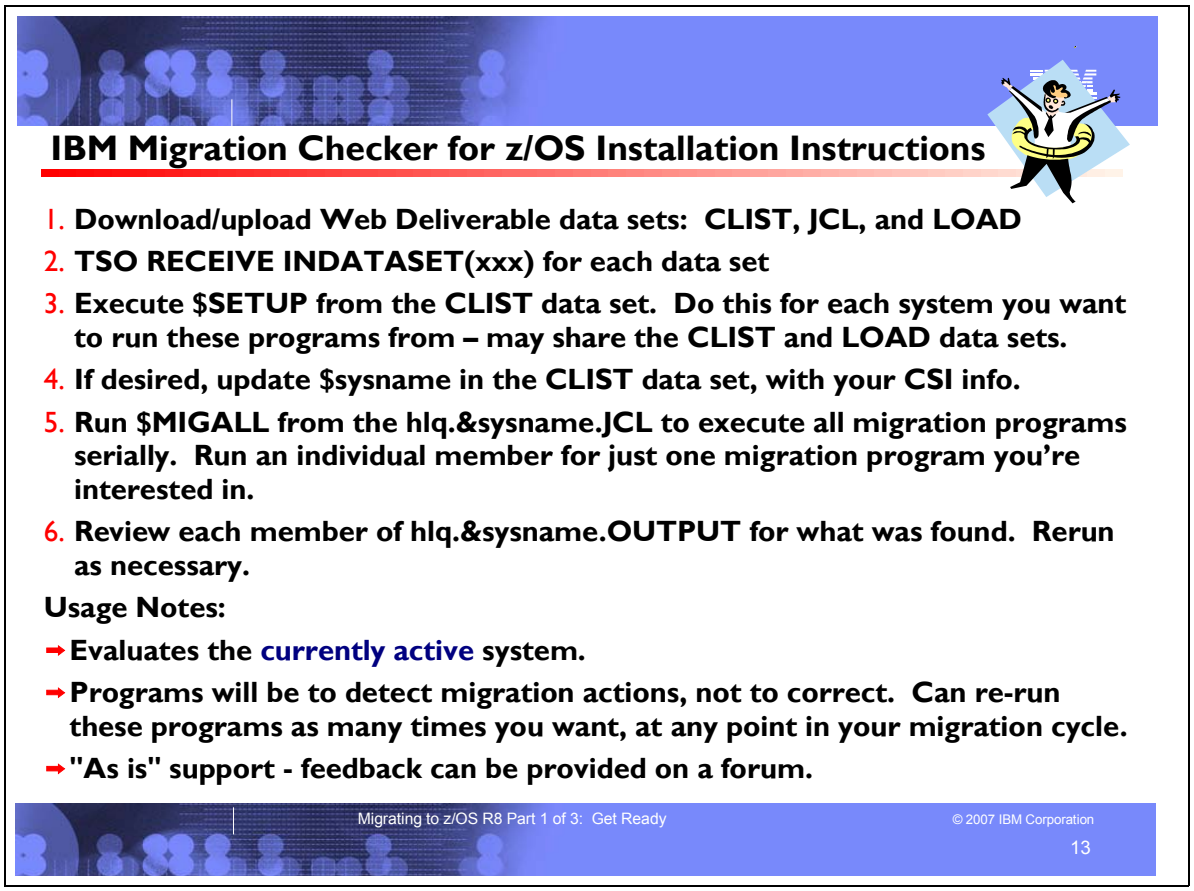

IBM Migration Checker for z/OS

The IBM Migration Checker for z/OS is a new tool that checks the applicability of certain migration actions on the currently running system. It does not **perform** migration actions on your system. Instead, the tool **reports** which migration actions you have to perform. The tool is best used to assist you in creating your migration plan in conjunction with the *z/OS Migration* book.

A selection of migration actions for this tool can assist you in your migration from **z/OS V1R7 to z/OS V1R8**. However, if you're not on the z/OS V1R7 to z/OS V1R8 migration path, most of the programs in the tool can also provide value on other migration paths such as z/OS V1R6 to z/OS V1R8, and even z/OS V1R4 to z/OS V1R7. The output reports produced from running the tool will tell you whether the tool was able to do any analysis on the currently running system, so running it outside the intended migration path does no harm, and in many cases is helpful.

The IBM Migration Checker for z/OS is an "as is" download tool available from the z/OS download Web page at http://www.ibm.com/servers/eserver/zseries/zos/downloads/. It's composed of several batch programs that can be run either independently or serially in one job to check the applicability of certain migration actions on the currently running system. The user who runs the programs must have the authority to issue the **TSO/E CONSOLE** command and have an OMVS segment defined.

You can run the tool as often as you like. On z/OS V1R7 before you've even ordered z/OS V1R8, after installing z/OS V1R8, and after each z/OS V1R8 image is deployed on your enterprise. Because the tool does not make any changes to your system and is aware of the level of the system on which it is running, it can:

- Alert you to migration actions that you should plan for in the future
- Detect migration actions that you can do now to clear the way to z/OS V1R8
- Verify that you've done the migration actions correctly when you're running z/OS V1R8.

© Copyright IBM Corporation, 2007

As configurations change, rerun the tool**.** For example, if your enterprise has multiple TCP/IP stacks, and activates and deactivates those stacks at different times, it is beneficial to run the tool multiple times to analyze each active TCP/IP stack.

You can dynamically change many of the items that the tool checks for on your currently running system. You can set up your automation to submit IBM Migration Checker for z/OS jobs on a regular schedule that fits your needs or you can manually submit the jobs. The frequencies is your choice, but keep in mind that only the latest report will be in the output data set. If you see a need to save the output from previous runs, you can include that in your automation steps too.

The IBM Service organization does not officially support this tool, but rather operates on a best-can-do basis. Please report any bugs, suggestions, or comments to the “IBM Migration Checker for z/OS” forum on IBM Customer Connect. For access to the forums, send an email to zosmig@us.ibm.com.

You must still review the migration actions in *z/OS Migration* that are not covered by the programs in the IBM Migration Checker for z/OS tool.

© Copyright IBM Corporation, 2007

Elements with Migration Actions for z/OS V1R8

► **Documented in** *z/OS Migration*

- *For complete migration tasks for z/OS R8, see this book!*
  - ***from R5 to R8, and R6 to R8, and R7 to R8, customized books***
  - ***"When behaviors aren't the same anymore, migration actions are called for."***

► ***From z/OS V1R7*** **to z/OS V1R8:**

- ➡ **BCP**
- ➡ **Communications Server**
- ◆ **Cryptographic Services - ICSF, OCSF, PKI,**
- ➡ **DFSMS**
- ◆ **DFSORT**
- ➡ **Distributed File Service - zFS**
- ➡ **HCD (introduced in z/OS R7, not R8)**
- ◆ **HCM**
- ◆ **Infoprint Server**
- ◆ **Integrated Security Services - Firewall**
- ◆ **ISPF**
- ➡ **JES2**
- ◆ **JES3**
- ➡ **Language Environment**
- ◆ **Library Server**
- ◆ **msys for Setup**
- ◆ **NFS**
- ◆ **RMF**
- ◆ **SDSF**
- ◆ **Security Server (RACF)**
- ➡ **XL C/C++**
- ➡ **z/OS UNIX System Services**

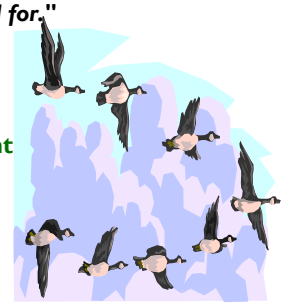

- ➡ ***means that selected element migration actions from these elements are discussed***
- ***If coming from z/OS R5 or z/OS R6 to R8, there are additional elements that have migration actions***

© 2007 IBM Corporation

Migration Actions for Elements Between z/OS R7 and z/OS R8

When migrating from z/OS R7 to z/OS R8, the specified elements in the foil above have required migration actions. Refer to *z/OS Migration* for complete information on the required migration actions for all elements. Selected elements with new z/OS R8 migration actions follow in this presentation, however, included for these selected elements are migration actions which were introduced in z/OS R6 and z/OS R7 (not just z/OS R8). This will show a more complete picture for any migration to z/OS R8.

If you are migrating from z/OS R5 or z/OS R6 to R8, there are other elements that have migration actions. For a complete list of migration actions which is applicable, refer to *z/OS Migration.*

© Copyright IBM Corporation, 2007

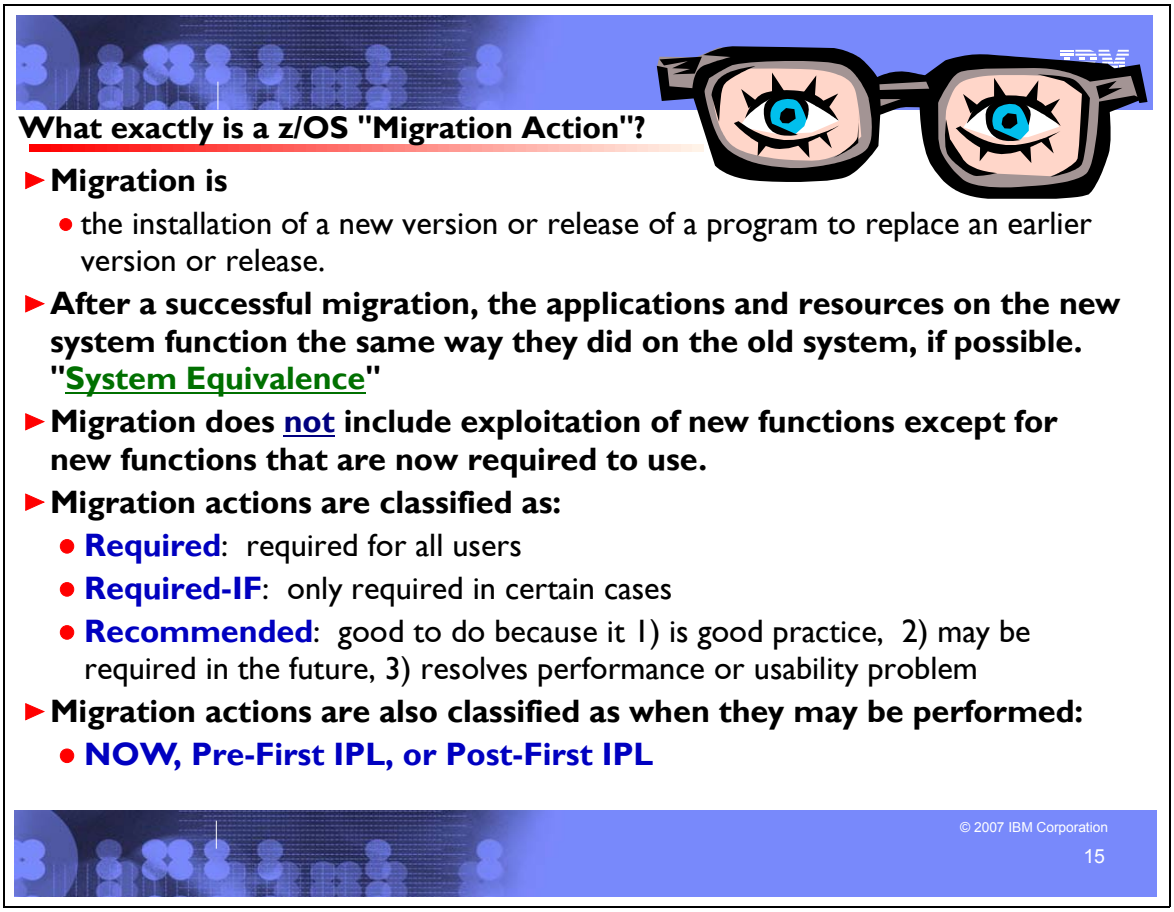

Migration Definitions and Classifications

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

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

Migration Requirement Classification and Timing

The migration actions are classified as to their requirement status:

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

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

- ***Now.*** These are migration actions that you perform on your current system, either because they require the current system or because they are possible on the current system. You don't need the z/OS V1R8 level of code to make these changes, and the changes don't require the z/OS V1R8 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 V1R8.

- ***Pre-First IPL.*** These are migration actions that you perform after you've installed z/OS V1R8 but before the first time you IPL. These actions require the z/OS V1R8 level of code to be installed but don't require it to be active. That is, you need the z/OS V1R8 programs, utilities, and samples in order to perform the migration actions, but the z/OS V1R8 system does not have to be IPLed in order for the programs to run. Examples are running sysplex utilities and updating the RACF database template.

  It is possible to perform some of the migration actions in this category even earlier. If you prepare a system on which you will install z/OS V1R8 by making a clone of your old system, you can perform migration actions that involve customization data on this newly prepared system before installing z/OS V1R8 on it. Examples of such migration actions are updating configuration files and updating automation scripts.
- ***Post-First IPL.*** These are migration actions that you can perform only after you've IPLed z/OS V1R8. You need a running z/OS V1R8 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 V1R8 to be active to perform the task. You might perform the task quite a while after the first IPL.

Icons used in the subsequent foils:

means that you shouldn't overlook this migration action

means that the IBM Migration Checker tool can help with all or part of this migration action

© Copyright IBM Corporation, 2007

Some General Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - **Verify that you have enough XCF groups in your sysplex CDS and enough XCF members in your XCF groups (Required-IF, as of R7,R8)** - In R7: XCF group IOEZFS, R8 XCF group SYSXCF. Reformat if you need to.
  - **Use the "new" default for your prog mgmt binder COMPAT specification (Required-IF, as of R8)**
    - R8 intro PO5 (PO4 was intro in R3, when COMPAT=MIN was the new default). Use default, in your SMP/E UTILITY entry.

- **Migration Actions Pre-First IPL:**
  - **Set up your IPCS environment (Required)**
    - Don't forget R7 new data set SYS1.SIEAMIGE, for SYSLIB, like SYS1.MIGLIB.
  - **Migrate /etc and /var system control files (Required)** - NFS obsoletes some /etc files.
  - **Remove references to deleted data sets and path (Required)**
    - ISP.SISPSASC (ISPF now uses LE RT SCEERUN/SCEERUN2.)
    - Still need DDDEF for CNMLINK for OSA/SF PTFS! (Can point to an empty data set)
  - **Add references to new data sets and paths (Required)**
    - New in R6, PDSE linklist data set **SYS1.SIEALNKE**.
    - New in R7, PDSE linklist data sets **SYS1.SIEAMIGE and SYS1.SHASLNKE**. New APF list PDSE **SYS1.NFSLIBE** to replace SYS1.NFSLIB.
    - New in R8: **SYS1.SIEALNKE** and **SYS1.SIEAMIGE are automatically added to the default system linklist!** You **must** catalog these data sets!

© 2007 IBM Corporation

Some General Migration Actions for z/OS V1R8 from z/OS V1R5

- ***Continued from previous foil...***

- **Migration Actions Pre-First IPL:**
  - **Accommodate new address spaces (Recommended)**
    - New in R6, **SMSPDSE1**, restartable PDSE address space that allows you to recover from some PDSE problems without having to re-IPL. And **TN3270E**. By default, not created during IPL.
    - New in R7: **DEVMAN, HZSPROC, JES2S00***1* for JES2 NJE over TCP/IP
    - New in R8: **DSSFRDSR** for hsm. Starts automatically whenever a ds is recovered from DASD using FRRECOV DSNAME command. Terminates when hsm terminates.
  - **Accommodate new SCOPE=COMMON data spaces in IEASYSxx MAXCAD (Required-IF)**
    - In R7: 1 for *JES2*WTO (rolled back in OA05981), 1 or more for XES connections to serialized structures.
    - In R6 and R8: None
    - **IBM Health Checker for z/OS can help you determine what to specify for the MAXCAD value!**

- **Migration Actions Post-First IPL:**
  - **<none>**

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

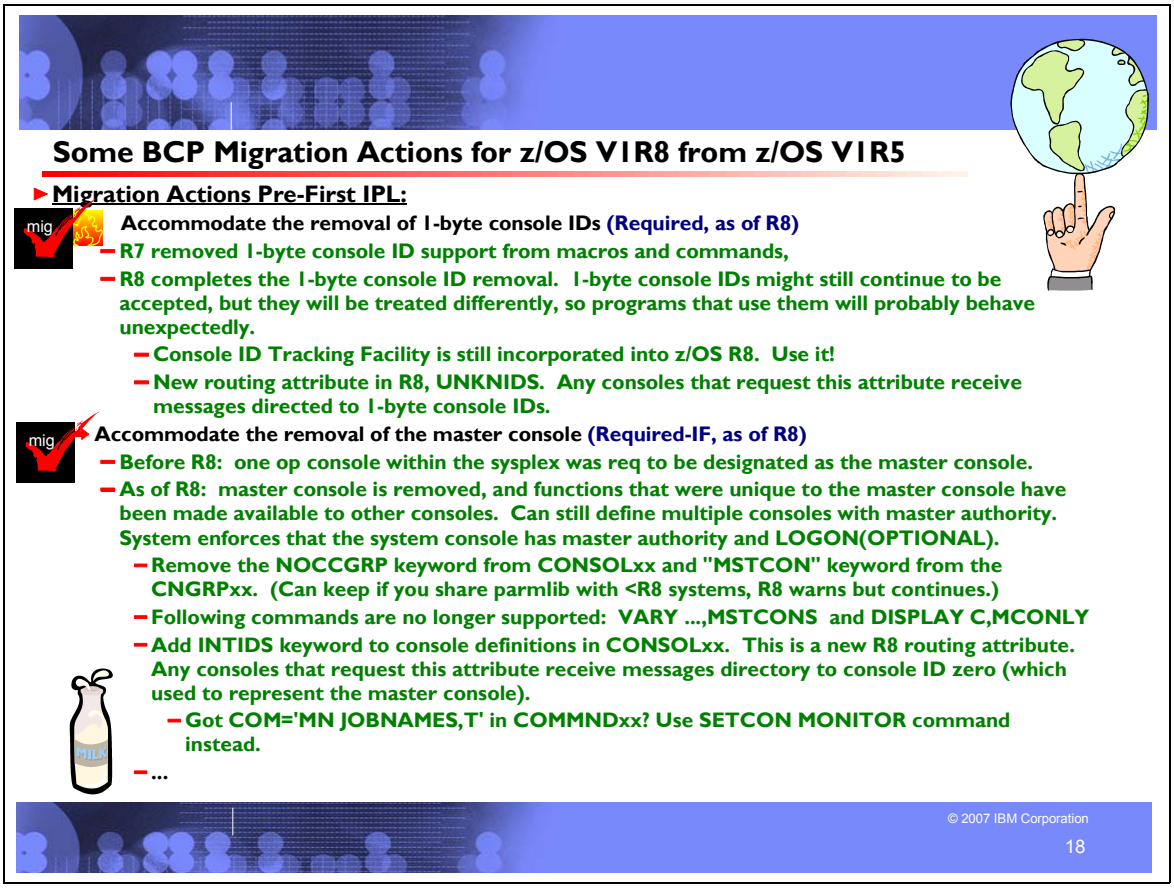

General Migration Actions Between z/OS V1 R5 and z/OS V1 R8

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

General Migration Actions You Can Do Now

Apply coexistence and fallback fixes (Required)

Migration action: Install coexistence and fallback PTFs on your systems to allow those systems to coexist with z/OS V1R8 systems during your migration, and allow backout from z/OS V1R8 if necessary. See *z/OS Migration* for a list of required coexistence PTFs, Take a look at using EPSPT to make this action even easier!

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

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

Migration action:

- Download SoftCap from one of the following Web sites:
  - Customers: http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS268
  - Business partners: http://www.ibm.com/partnerworld/sales/systems
- 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.

Find alternatives to removed elements and features (Required-IF)

Required if you use any of the removed elements and features.

Several base elements and optional features have been removed from the operating system and therefore can no longer be used. The following elements have been removed:

- DCE Application Support (in z/OS R6)
- Encina Toolkit Executive (in z/OS R6)

- Text Search (in z/OS R6, but is still available as a Web deliverable if needed for limited use)
- msys for Operations (in z/OS R8)

In addition, you should prepare for elements and functions that will be removed in future z/OS releases. For alternatives to the removed elements and features, see *z/OS* Migration.

Migration action for msys for Operations: z/OS V1R7 was the last release to support the msys for Operations base element. msys for Operations has been removed from z/OS V1R8. IBM plans to transition many of the msys for Operations functions to a new user interface and infrastructure in a future release of z/OS.

You should remove any msys for Operations customization that you performed on your pre-z/OS V1R8 system so as not to carry it forward to your z/OS V1R8 system. As a guide, use "Appendix D. msys for Operations customization checklist" in *z/OS Managed System Infrastructure for Operations Setting Up and Using*. Failure to remove customization can result in the following problems:

- JCL errors or abends trying to start the msys for Operations instance
- Redundant VTAM APPLs
- Redundant RACF customization
- Redundant system-unique data sets
- Redundant Support Element customization

Upgrade Windows 2000, 95, 98, and NT clients (Required-IF)

Required if you have clients running Windows 2000, 95, Windows 98, or Windows NT Workstation.

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

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

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

Use the new default for your program management binder COMPAT specification (Req-IF, as of R8)

Required if you invoke the program management binder specifying a compability value other than the default (COMPAT=MIN)

Since z/OS V1R3, the program management binder default for specifying the compatibility level has been MIN, meaning that the lowest level that supports the features used is selected for the current bind. In z/OS V1R8, a new COMPAT level, PO5, was introduced. If you invoke the program management binder on z/OS V1R8 and override the default of COMPAT=MIN (such as specifying COMPAT=CURRENT, which was the default before z/OS V1R3), you might have program objects produced that are at a higher format than you intended and that cannot be executed on systems earlier than z/OS V1R8.

Migration action: Make sure that any program management binder parameters you specify use the default for the COMPAT statement. This ensures that program objects can be executed from the most environments. A common place where the COMPAT statement may be specified is in your SMP/E GLOBAL zone UTILITY entry for the program management binder.

Verify that you have enough XCF groups in your CDS and enough XCF members in your XCF group (Required-IF, as of R7 and R8)

Required if you have an inadequate number of XCF groups and members formatted in your sysplex couple data sets, and functions that you are using need additional XCF groups or members.

Over time, as various z/OS functions and applications exploit XCF services, you must ensure that there is enough space in the sysplex couple data set for all the XCF groups and members that are to be defined by the exploiters. It is possible that your sysplex couple data set could contain an inadequate number of XCF groups or members.

Migration action:

1. Issue the DISPLAY XCF,COUPLE command on your current system. Notice the values of MAXGROUP and PEAK for your sysplex couple data sets. These values show you the maximum number of XCF groups that the couple data sets can support, and the peak number of XCF groups ever in use in the sysplex. Also notice the values of MAXMEMBER and PEAK for your sysplex couple data sets. These values show you the maximum number of members that the couple data set can support in one group, and the greatest number of members ever in use in the largest group in the sysplex.
2. In z/OS V1R7, one additional XCF group is used. Its name is IOEZFS and it is used by base element Distributed File Service to provide zFS support. In z/OS V1R8, one additional XCF group is used; its name is SYSXCF and it is used

© Copyright IBM Corporation, 2007

by base element BCP for XCF support. If you do not have enough groups formatted to support this additional XCF group, reformat your sysplex data sets with a larger number of groups.

3. If your peak member value is close to the maximum member value, you might want to reformat your sysplex couple data sets to support a larger maximum number of members to be used by any one group.

Reference information: For information about formatting sysplex couple data sets with the MAXGROUP and MAXMEMBER parameters, see *z/OS MVS Setting Up a Sysplex*. For information about the DISPLAY XCF command, see *z/OS MVS System Commands*.

Update the CustomPac Installation Dialog (Required-IF, as of R6)

Required if you use ServerPac, and have not yet performed this migration action.

Migration action: If this is the first time you will install a ServerPac using the level of the CustomPac Installation Dialog supplied with z/OS V1R6 or higher, you must update the dialog before receiving your ServerPac order. Determine whether you need to update the CustomPac Installation Dialog by starting the dialog and looking at the bottom of the first panel displayed (CPPPPOLI) for the text "This dialog supports electronic delivery". Then:

- If the panel contains the text, no migration action is required.
- If the panel does not contain the text, follow the instructions for getting and running the UPDATE or EUPDATE job in *ServerPac: Using the Installation Dialog*.

Note: If you are installing an older order and want to run UPDATE or EUPDATE, use the CPPCINV sample job to convert some of the older order's data sets to VSAM before finishing its installation.

Failure to perform the update will cause an ABEND 813 during ServerPac RECEIVE processing! You'll see:

IEC149I 813-04,IFG0195H,*myjobname*,STEP01,INPUT7,*tape_unit,tape_volser*,SYS1.*orderid*.LOADLIB

Reference information: *ServerPac: Using the Installation Dialog*

Update local panels, CLISTs, and execs that invoke the CustomPac Installation Dialog (Required-IF)

Required if you use ServerPac, and have not yet performed this migration action.

Migration action: If this is the first time you will install a ServerPac using the level of the CustomPac Installation Dialog supplied with z/OS V1R6 or higher, you might need to update local panels, CLISTs, and execs used to start it. Determine whether you need to update the CustomPac Installation Dialog by starting the dialog and looking at the bottom of the first panel displayed (CPPPPOLI) for the text "This dialog supports electronic delivery". Then:

- If the panel contains the text, no migration action is required.
- If the panel does not contain the text, remove the ISPF4X parameter from the CPPCISPF command in local CLISTs, execs, and panels. For example, change:

```
PROC 1 CPPMHLQ  ISPF4X(N)
CONTROL NOMSG NOFLUSH /* LIST SYMLIST CONLIST
SET COMMAND = CPPCISPF /* COMMAND NAME, DO NOT DELETE */
```

to:

```
PROC 1 CPPMHLQ
CONTROL NOMSG NOFLUSH /* LIST SYMLIST CONLIST
SET COMMAND = CPPCISPF /* COMMAND NAME, DO NOT DELETE */
```

Also, if you have added recovery code to a local CLIST or exec to restore a user's PROFILE PREFIX in case of abnormal dialog termination, remove it. The CustomPac Installation Dialog no longer runs with NOPREFIX set.

Reference information: *ServerPac: Using the Installation Dialog*

General Migration Actions Pre-First IPL

Set up your IPCS environment (Required)

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

Note: Starting with z/OS V1R7, a new data set, SYS1.SIEAMIGE, has been added. It is a PDSE data set that complements SYS1.MIGLIB. This data set needs to be used along with SYS1.MIGLIB, when necessary, for IPCS.

Reference information: For more information about IPCS, see *z/OS MVS IPCS Customization*. For more information about the correct logon procedure updates, see the *z/OS Program Directory*. For information about setting up the JES2 IPCS environment, see *z/OS JES2 Diagnosis*. For information about setting up the JES3 IPCS environment, see *z/OS JES3 Diagnosis.*

Use IBM-supplied PARMLIB and PROCLIB (Required)

Migration action: For parmlib, add the data set pointed to by the z/OS V1R8 PARMLIB DDDEF to your parmlib concatenation. The data set should generally be added last in the concatenation, and you should make sure that the other data sets in the concatenation don't have members with the same names as IBM-supplied members. If you place the data set on the system residence volume and use an indirect catalog entry, future migrations won't require this particular migration step.

- For proclib:
  1. Ensure that the default proclib members have been copied to your default proclib to pick up the new and changed members.
  2. Update individual sample members provided and ensure they are accessible to the system, as shown in the table of proclib member updates in *z/OS Program Directory*.
  3. Ensure that the procedure libraries listed in the table of libraries to be added to the proclib concatenation in *z/OS Program Directory* have been placed in the necessary procedure library concatenations and are available to the system.

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

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

Migration action: The /etc and /var directories contain system control files: the /etc directory contains customization data that you maintain and the /var directory contains customization data that IBM maintains. During installation, subdirectories of /etc and /var are created. If you install z/OS using ServerPac, some files are loaded into /etc and /var due to the customization performed in ServerPac. You have to merge the files in /etc and /var with those on your previous system. If you install z/OS using CBPDO, you should copy the files from your old system to the z/OS V1R8 /etc and /var subdirectories.

Copy files from your old system to the z/OS V1R8 /etc and /var subdirectories, and then modify the files as necessary to reflect z/OS V1R8 requirements. If you have other files under your existing /var directory, then you will have to merge the old and new files under /var. The easiest way to do this is to create a copy of your current /var HFS and then copy the new /var files into the copy.

The following z/OS V1R8 elements and features use /etc:

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

The following z/OS V1R8 elements and features use /var:

- CIM
- Cryptographic Services - OCSF
- Infoprint Server
- Integrated Security Services - Network Authentication Service uses /var/skrb.

As of z/OS R7, NFS will no longer use the /etc directory for lock files, it will use VSAM data sets instead.

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

Update items affected by changed interfaces (Required-IF)

Required if any of the interface changes impact your installation.

Interfaces in z/OS are parts of the system that allow people, application programs, execs, or procedures to interact with the system. Examples of interfaces are commands, macros, panels, exit routines, callable services, data areas, the parameter library (SYS1.PARMLIB), the procedure library (SYS1.PROCLIB), the sample library (SYS1.SAMPLIB), configuration files, and environment variables. Various changes have been made to interfaces since z/OS R5.

In order for your application programs and resources to work the same way on your new system as they did on your old one, you'll probably have to make updates due to the interface changes. In addition, you'll probably have to notify people such as operators and end users about interface changes that affect tasks that they perform.

© Copyright IBM Corporation, 2007

Migration action: Review the book *z/OS Summary or Interface and Message Changes* and *z/OS MVS Init and Tuning Reference.* Update application programs, execs, procedures, and system parameters (in parmlib) appropriately.

Verify that virtual storage is set properly (Required)

Migration action: Determine how much virtual storage use to allow above the 2 GB bar. While there is no practical limit to the number of virtual addresses an address space can request above the bar, the system can limit the amount of virtual storage above the bar that an address space is allowed to use. The amount of virtual storage above the bar is determined as follows. The MEMLIMIT parameter in parmlib member SMFPRMxx sets the default system-wide limit, which defaults to zero (nothing) if it is not specified. However, the system-wide default MEMLIMIT can be overridden by specifying REGION=0M or MEMLIMIT on JOB or EXEC statements in JCL. To set a limit on the use of virtual storage above the bar, use the SMF exit IEFUSI. For more information, see the topic about limiting the use of memory objects in *z/OS MVS Programming: Extended Addressability Guide*.

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

- For MEMLIMIT, you must specify a nonzero MEMLIMIT in an active SMFPRM*xx* member of parmlib to establish a system default other than zero for available virtual storage above 2 GB. (The default MEMLIMIT is zero.)
- You can specify MEMLIMIT explicitly in JCL to override the system default that was set (or allowed to default) in SMFPRM*xx*.
- You can specify REGION=0M on the job statement in JCL to implicitly set MEMLIMIT to NOLIMIT, which also overrides the system default (from SMFPRM*xx*).
- You can use IEFUSI both to establish a system default MEMLIMIT for different classes of work (for example, job, TSO, STC) and limit the amount of virtual storage that can be used above the bar, provided that an explicit or implicit nonzero MEMLIMIT is in effect from JCL or SMFPRM*xx*.

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

Back virtual storage with real and auxiliary storage (Required)

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

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

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

Remove references to deleted data sets and path (Required)

Migration action: Using the table in *z/OS Migration, "*Data sets and paths deleted from z/OS V1R8, R7, and R6" 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: Especially noteworthy is DDDEF **CNMLINK**, which was used by the msys for Operations base element. Even though msys for Operations was withdrawn from z/OS V1R8, the installation of OSA/SF (FMID H0GI400) PTFs still requires DDDEF CNMLINK to be defined in the target zone. The DDDEF for CNMLINK must point to an empty data set when Tivoli NetView for OS/390 V1R4 (5697-B82) is not installed. (Part of Tivoli NetView for OS/390 V1R4, as well as part of System Automation for z/OS V2R3, was included in msys for Operations.) Removing the DDDEF for CNMLINK will result in failures when installing OSA/SF PTFs.

© Copyright IBM Corporation, 2007

Reference information: *z/OS Migration* contains the list of all removed data sets and paths in z/OS R8, R7, and R6.

Add references to new data sets (Required)

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

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

Of special note are the data sets:

- **SYS1.SIEALNKE** and **SYS1.SIEAMIGE**. These PDSE data sets are used by multiple elements and features. (SIEALNKE was introduced in z/OS V1R6 and SIEAMIGE was introduced in z/OS V1R7.) They provide a common location to store program objects that are eligible for the link list. **As of z/OS V1R8, the system automatically concatenates these two data sets to the default system link list. Therefore, you must ensure that these data sets are cataloged.**
- **SYS1.SHASLNKE**, which is a PDSE introduced in z/OS V1R7 and used by JES2. It replaces the data set SYS1.SHASLINK. When SYS1.SHASLNKE is added to the link list, by default it is APF authorized and does not need to be specified on a STEPLIB DD statement.
- **SYS1.NFSLIBE,** which is a PDSE introduced in z/OS V1R7 and used by NFS. It replaces the data set SYS1.NFSLIB. Be sure to update your APF list in parmlib with the new name.

Reference information: *z/OS Migration* contains the list of all new data sets and paths in z/OS R7, R6, and R5.

Accommodate new address spaces (Recommended)

Not required, but recommended if you use the CVAFTR procedure. If you don't use it, the migration action is not required but is recommended so that you take advantage of the benefits afforded by the new address spaces. Also recommended to use the TN3270E address space for enhanced availability and it is planned to be a future requirement (see the Communications Server portion of this presentation).

The following address spaces are new:

- **DSSFRDSR:** This address space is started by DFSMShsm to invoke DFSMSdss to recover up to 64 data sets concurrently from one or more copy pool backup versions. There is nothing for you to do. This DFSMSdss address space is started automatically by DFSMShsm whenever a data set is recovered from DASD using the FRRECOV DSNAME command. The address space terminates automatically when DFSMShsm terminates.
- ***DEVMAN:*** This DFSMSdfp address space is new in z/OS V1R7. z/OS uses this device manager address space to provide various kinds of device support. In z/OS V1R7, this support includes component trace information for DADSM creation and CVAF events. These new component trace records are formatted when viewed by IPCS using CTRACE COMP(SYSDMO). Component tracing for DADSM creation and CVAF events are activated and deactivated by DEVMAN modify commands. The CVAF trace events to GTF formerly provided by the CVAFTR procedure have been removed. The debug/trap options of CVAFTR, however, are still available.
- ***HZSPROC:*** This address space, new in z/OS V1R7, is for the IBM Health Checker for z/OS. IBM recommends the use of the IBM Health Checker for z/OS because it can identify potential problems before they impact your availability or cause outages. It checks the current active z/OS and sysplex settings and definitions for a system and compares the values to those suggested by IBM or defined by you. The open architecture of the framework supports checks written by IBM, ISVs, and users. The default ServerPac configuration provides the setup and starts the HZSPROC address space. If you do not want to start this address space, update the ServerPac generated jobs and supplied samples accordingly.
- ***JES2 NJE over TCP/IP:*** As of z/OS V1R7 JES2, NJE over TCP/IP runs in a new address space. The address space is responsible for sending data to, and receiving data from, TCP/IP. This move of processing from the JES2 main task to the user environment is expected to improve reliability, availability, and serviceability by isolating the moved function from processing in the JES2 main task, and to improve performance by taking advantage of multiple CPUs to do the work (so that not all work is done under the single JES2 main task). One address space is created for each active JES2 NETSERV. The format of the name of the address space is *jjjj*S*nnn*, where *jjjj* is the owning JES2 subsystem name and *nnn* is the subscript of the owning NETSERV. For example, for NETSERV(1) owned by subsystem JES2, the name would be JES2S001.

© Copyright IBM Corporation, 2007

This address space is started and stopped using the $S NETSERV and $P NETSERV JES2 commands. No special definitions are required. However, if desired, the characteristics of the TCP/IP connection associated with the address space can be altered in the TCP/IP profile data set.

- ***SMSPDSE1****:* This DFSMSdfp address space was new in z/OS V1R6. It is a restartable PDSE address space that allows you to recover from a PDSE problem (hang, deadlock, or out-of-storage condition) without having to re-IPL the system or sysplex, and to determine which systems require a re-IPL or repair to eliminate PDSE hangs and failures. SMSPDSE1 is related to SMSPDSE. SMSPDSE is a nonrestartable address space for PDSEs that are in the link list concatenation. SMSPDSE was introduced in z/OS V1R4 through the service stream (and rolled back to prior releases) to relieve extended common service area (ECSA) usage. SMSPDSE cannot be restarted because global connections cannot handle the interruption and reconnection that are part of an address space restart. SMSPDSE is the only PDSE address space for z/OS systems when the following conditions exist:
  – The IGDSMS*xx* initialization parameter, PDSESHARING, is set to NORMAL.
  – The IGDSMS*xx* initialization parameters in a sysplex coupled systems environment are set as follows:
    - PDSESHARING(EXTENDED)
    - PDSE_RESTARTABLE_AS(NO)

  If these conditions do not apply, the z/OS system has both the SMSPDSE and SMSPDSE1 address spaces. The SMSPDSE1 address space provides connections to, and processes requests for, PDSE data sets that are not part of the link list concatenation.
- ***TN3270E*** *Telnet server:* This Communications Server address space was new in z/OS V1R6. It allows the TN3270E server to run in a separate address space from the TCP/IP stack address space. In z/OS R9 Communications Server, support for the in-stack version of the TN3270 Server is planned to be discontinued.

Rework and install user modifications (Required-IF)

Required if you have made any user modifications that necessitate changes.
Migration action: Use the z/OS SMP/E Planning Migration Assistant to help determine which user modifications need to be reworked and which just have to be reinstalled. The Top or New Intermediate Product Migration Changes Report uses data found on your system, combined with IBM-supplied information from the Software Information Base, to show you the current levels of products available as well as product migration and functional changes using a comparison of FMIDs. You can use this report to determine the product migration impacts by reviewing the “changed” FMIDs. This can help you assess how many user modifications have to be reworked if you issued the LIST SYSMOD USERMOD FORFMID (listing the “changed” FMIDs) command. All other user modifications can be reinstalled without having to be reworked.
Note: IBM recommends using exit routines for any user modifications where possible, and installing the exit routines with SMP/E. By using SMP/E, it is easier to bring forward desired modifications to the z/OS release you are installing.

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

- XL C/C++
- DFSORT
- HLASM
- Language Environment. Investigate using CEEROPT, which can be used to specify run-time options for CICS, IMS LRR, and other LRR users. Even better, consider using the function added in z/OS R7 to eliminate your assembler language run-time option modules in CEEPRMxx parmlib member!

Reference information: For information about C/C++ customization, see z/*OS C/C++ Compiler and Run-Time Migration Guide*. For information about DFSORT customization, see *DFSORT Installation and Customization Guide*. For information about HLASM customization, see *HLASM Installation and Customization Guide*. For information about Language Environment customization, see *z/OS Language Environment Customization*.

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

Required if you use any ISV products and need to reconnect them after performing a ServerPac installation, or if you intend to use any subsystems with your z/OS system.
Migration action: Follow the instructions for each ISV product that you use to reconnect it to your z/OS V1R8 ServerPac.

Ensure that any required coexistence service is installed prior to using the subsystem with the new z/OS V1R8 system, as well as any required SVCs, system modifications, parmlib setup, and proclib setup. Follow the instructions for the subsystem that you need to reconnect.
Reference information: For a list of independent software vendors (ISVs) that support z/OS, as well as announcements, testimonials, and other information, see http://www.ibm.com/eserver/zseries/solutions/s390da/. For a directory of ISV products that support z/OS, see the Global Solutions Directory at http://www.ibm.com/software/solutions/isv.

© Copyright IBM Corporation, 2007

Update operational and other procedures (Required)

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

- Allowing applicable users access to new high-level qualifiers that you may have. There are no new default high-level qualifiers introduced in z/OS R8, z/OS R7, or z/OS R6 with the new target data sets.
- Updating and testing your backup and recovery procedures to accommodate the new target system.
- Updating and testing any disaster recovery procedures.
- Updating and testing any automation procedures to take advantage of new functions.
- Updating security system definitions, such as defining new users and resources, permitting users to use new resources, and defining new profiles in the RACF FACILITY class.

Reference information: For information about the new functions incorporated into z/OS V1R8, see *z/OS Introduction and Release Guide*.

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

Required if you use the MAXCAD parameter of parmlib member IEASYSxx and the value you specified is inadaquate for your z/OS R8 system.

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

- Added in z/OS V1R7:
  – One data space named *xxxx*WTO (where *xxxx* is the JES2 name), for example, JES2WTO. (This data space was rolled back to prior releases by APAR OA05981.)
  – One or more data spaces for XES connections to serialized structures. For details, see the BCP migration action that follows in this presentation.
- In z/OS R6 and z/OS R8: None.

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

Tip: The IBM Health Checker for z/OS can help you determine what to specify for the MAXCAD value. Use the check IBMRSM,RSM_MAXCADS. This check is available as of z/OS V1R7 and also in APAR OA09366 back to z/OS V1R4. By running this check, you can find out:

- The MAXCAD value you specified during IPL
- The number of SCOPE=COMMON data spaces currently in use
- The high water mark, which is the highest usage of SCOPE=COMMON data spaces used during this IPL

Use this information to help you set your MAXCAD specification in IEASYS*xx*.

Make z/Architecture changes (Required-IF)

Required if you have not already migrated to z/Architecture.

z/Architecture is the 64-bit architecture provided by the IBM zSystem 9 (z9) server and the IBM @server zSeries servers (z990, z890, z900, and z800). With z/OS IPLed in z/Architecture mode on a z9 or zSeries server, you can take advantage of functions such as 64-bit real and virtual storage support, Intelligent Resource Director (IRD), and HiperSockets. If you are migrating from an IBM S/390 Parallel Enterprise Server – Generation 5 (G5) or Generation 6 (G6), or from an IBM S/390 Multiprise 3000 Enterprise Server, to the z9 server or one of the zSeries servers, you need to migrate to z/Architecture. z/OS V1R7 on a z9 or zSeries server does not support ESA/390 architecture.

Migration action: Use the following checklist as a guide to the changes you have to make to take advantage of z/Architecture: (This list is kept current with WSC Flash 10185.)

1. Decide how many servers and LPARs you need. Support for 64-bit real storage might allow you to consolidate your current systems into fewer LPARs or to a single native image.
2. Determine whether any locally-developed code is affected by z/Architecture. For information about 64-bit virtual addressing support, see:
   - *z/OS MVS Programming: Assembler Services Reference ABE-HSP*
   - *z/OS MVS Programming: Assembler Services Reference IAR-XCT*
   - *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN*
   - *z/OS MVS Programming: Authorized Assembler Services Reference ENF-IXG*
   - *z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU*
   - *z/OS MVS Programming: Authorized Assembler Services Reference SET-WTO*

© Copyright IBM Corporation, 2007

3. Check with your ISVs to ensure that they support z/Architecture. See http://www.ibm.com/eserver/zseries/solutions/s390da/osnp.html
4. Check the z9 EC, z9 BC, z990, z890, z900, or z800 PSP bucket for any service needed for 64-bit addressing (or other hardware-related functions). For the z9 EC, the upgrade is 2094DEVICE, and the subset is 2094/ZOS. For the z9 BC, the upgrade is 2096DEVICE, and the subset is 2096/ZOS. For the z990, the upgrade is 2084DEVICE and the subset is 2084/ZOS. For the z890, the upgrade is 2086DEVICE and the subset is 2086/ZOS. For the z900, the upgrade is 2064DEVICE and the subset is 2064/ZOS. For the z800, the upgrade is 2066DEVICE and the subset is 2066/ZOS.
5. Install real storage manager (RSM) service: PTFs for APARs OW55209, OW55255 (and companion OW56071), OW54938, and OW55729 (and companion OW55902).
   Note: It is important to read the APAR cover letters and apply circumventions as appropriate. For example, the circumvention for OW55729 and OW55902, which is documented in OW55729, recommends that you set the IEAOPT*xx* MCCAFCTH thresholds based on the amount of central storage on your system. Once you have the fixes for OW55729 and OW55902 installed, it is recommended that you use the default MCCAFCTH thresholds (and remove the circumvention), and remove the MCCAFCTH statement from IEAOPT*xx* altogether.
6. Plan for additional resources for stand-alone dump. In z/Architecture, all of central storage is dumped rather than just the 2 GB of central storage in an ESA/390 system. In order to support this much larger requirement, review the use of multivolume stand-alone dump data set support, which was provided in z/OS V1R2. Refer to *z/OS MVS Diagnosis: Tools and Service Aids*. Also review Washington Systems Center flash 10143 at http://www.ibm.com/support/techdocs (Search for "flash10143".)
7. When you're ready to switch to z/Architecture, reconfigure the LPAR to use only real storage (formerly called central storage) rather than central and expanded storage by shutting down the system, deactivating and then activating the LPAR to pick up the new storage allocation, and then IPLing the system. (LPARs in basic mode require a power-on reset.). If you run in z/Architecture mode and still have expanded storage defined, the expanded storage is not used and the system issues these messages:
   - IAR016I THE SYSTEM WAS IPLED IN ESAME MODE WITH EXPANDED STORAGE DEFINED. THIS STORAGE WILL NOT BE USED BY THE SYSTEM.
   - IEE038E AMOUNT OF EXPANDED STORAGE EXCEEDS 0G MAXIMUM. EXPANDED STORAGE IN EXCESS OF MAXIMUM IS IGNORED. RECONFIGURATION FUNCTIONS ARE NOT AVAILABLE. This is followed by the WTOR IEE039A REPLY TO ACKNOWLEDGE MESSAGE IEE038E.

   Any expanded storage you have defined for z/Architecture mode is allocated to the partition and, even though it isn't used, is unavailable to other partitions.
   Fallback note: ESA/390 mode can still be IPLed successfully in an LPAR defined with greater than 2 GB real storage. However, only the first 2 GB of real storage is used, the system issues the following two messages, and the storage is unavailable to other partitions:
   - IAR015I THE SYSTEM WAS IPLED IN ESA/390 MODE WITH MORE THAN 2 GIGABYTES OF CENTRAL STORAGE
   - IEE020E AMOUNT OF CENTRAL STORAGE EXCEEDS 2G MAXIMUM. CENTRAL STORAGE IN EXCESS OF MAXIMUM IS IGNORED. This is followed by the WTOR IEE021A REPLY TO ACKNOWLEDGE MESSAGE IEE020E.

   The preferred fallback method is to reconfigure the LPAR to use central and expanded storage as desired, shut down the system, deactivate and then activate the LPAR to pick up the new storage allocation, and then IPL the system.
   System services that used expanded storage in ESA/390 mode (such as hiperspaces) have been changed to use real storage in z/Architecture mode. Programs that use these system services should not require any changes.
   The amount of central storage you should need is the sum of your current central and expanded storage. No additional processor storage is required by z/Architecture itself. You should need more storage only if your workloads grow.
   Slightly more real storage might be required to IPL in 64-bit mode. Use the SoftCap tool to determine how much.
8. Understand how to evaluate an all central storage environment. See Washington Systems Center white paper WP100269 at http://www.ibm.com/support/techdocs
9. Understand Load Real Address (LRA) instruction considerations: LRA issued by a program running in AMODE 24 or AMODE 31 cannot return an address larger than 2 GB and it can be used if the virtual address is known to be backed below 2 GB.
   Some additional information: When z/OS is running in z/Architecture mode, authorized applications that issue the LRA instruction against unfixed storage might receive an abend 0D3, reason code 13. The results of an LRA instruction against unfixed storage has always been unpredictable. In z/Architecture mode, the LRA instruction might require a 64-bit result when only a 32-bit result can be returned; in this case, the hardware causes a program interrupt. If the program is using the LRA instruction to validate that the virtual address is backed by real storage, then use the TPROT instruction instead. If a valid real address is required, the storage must be properly fixed in real

© Copyright IBM Corporation, 2007

storage below 2 GB before issuing the LRA instruction. The STRAG or LRAG instruction can be used to obtain a 64-bit real address if the storage can be backed anywhere. You should review your usage before migrating to z/Architecture, and make any required changes.

10 Understand that z/Architecture provides a different format for architected parts of the first page of storage (the PSA). The ESA/390 format is mapped by macro IHAPSA. The z/Architecture format is mapped by macro IHAPSAE. In particular, the "old PSW", "new PSW" and "interruption parameter" fields have moved. Programs should not use any of these fields. But, if they do, they will not work correctly when running in z/Architecture (64-bit) mode. Note that none of the programming interface fields in the PSA have changed. For information about IHAPSA and IHAPSAE, look under the name PSA in *z/OS MVS Data Areas, Vol 3 (IVT-RCWK)*.

11. Possibly increase the size of the extended system queue area (ESQA). Before z/Architecture, the system determined the size of the ESQA automatically during IPL/NIP to allow for the storage required to manage expanded storage. A portion of this ESQA storage was typically available for system use. However, when running in z/Architecture mode, the system does not set aside as much ESQA because none is needed to manage expanded storage, and thus there is also less ESQA available for other system uses. The effects of this reduction in ESQA will vary from system to system, ranging from:
    - No negative impact at all, if the previous total ESQA size was large enough
    - Increased conversion of ECSA/CSA to ESQA/SQA
    - An inability to IPL
    - Outages after IPL with abend X'878', abend X'80A', or a wait state

    To avoid the more serious impacts, such as inability to IPL and outages, you might need to increase either:
    - The ESQA value in the SQA= parameter of the IEASYSxx parmlib member, or
    - The INITSQA value in the LOADxx parmlib member.

    If you do specify additional ESQA on one of the two parmlib parameters, keep in mind that the reduction in ESQA available after migrating to z/Architecture mode will be about 8 MB less for every 1 GB of expanded storage before the migration. However, because most of this storage was actually used for expanded storage management, you might need to specify only a fraction of this amount of additional ESQA. The amount of additional ESQA required will vary on each system.
12. Determine if you want to control (by using the API64R start option) whether Communications Server - IP Services application programs should be passed 64-bit backed storage. TCP/IP 64-bit real addressing support is automatically enabled. TCP/IP exploits real storage in excess of 2 GB by allowing z/OS to back most fixed communications storage manager (CSM) data space pages above the 2 GB real storage line. For information about API64R, see *z/OS Communications Server: SNA Resource Definition Reference*.
13. Understand the capacity impacts of migrating to 64-bit mode. Washington Systems Center flash 10086 outlines the capacity planning methodology. (Search for "flash10086" at http://www.ibm.com/support/techdocs) Use the SoftCap tool, available through the same Web address, to evaluate the effect on z/Architecture and processor capacity when migrating to newer levels of software.
14. Evaluate and plan your paging environment. Review Chapter 2 of Washington Systems Center white paper *z/OS Performance: Managing Processor Storage in a 64-bit Environment - V1*.1 at http://www.ibm.com/support/techdocs (Search for "WP100269".)
15. Understand enhancements in RMF. All metrics related to expanded storage are obsolete.
16. Understand 64-bit virtual storage restrictions:
    - Data spaces greater than 2 GB are not supported.
    - Hiperspaces greater than 2 GB are not supported.
    - Subspace capability is not extended to virtual storage above 2 GB.
    - DIV is not extended to virtual storage above 2 GB.
    - Sharing virtual storage above 2 GB is not supported.
    - Copy-on-Write is not supported for virtual storage above 2 GB.
    - DAT tables for virtual storage above 2 GB are not pageable
17. Review the IEFUSI exit routine for 64-bit virtual considerations. In the past you could limit program storage below 16 MB in virtual storage by using IEALIMIT. IEALIMIT can still be used to limit program storage in the nonextended region; however, IEFUSI is the preferred exit routine. For details on the advantages over IEALIMIT, see *z/OS MVS Installation Exits*.
18. Use the MEMLIMIT parameter of parmlib member SMFPRM*xx*. This new parameter establishes a system default other than zero for available virtual storage above 2 GB. It is used for jobs that do not explicitly specify MEMLIMIT (or REGION=0) in their JCL. You must set up the system MEMLIMIT default according to individual system environment requirements. You need to take into account the overall central storage and paging space requirements in assessing an individual system's suitable MEMLIMIT default value. See Washington Systems Center flash 10165 at http://www.ibm.com/support/techdocs. (Search for "flash10165".) We recommend that you use the IEFUSI exit routine

© Copyright IBM Corporation, 2007

to establish a MEMLIMIT value to protect the system from runaway tasks using too much storage above the bar, thus exhausting central and auxiliary storage.

19. Review APAR OW54022, which describes changes to the thresholds at which SRM detects a common storage shortage below the 16 MB line.
20. Specifying ARCHLVL 1 is not supported as of z/OS V1R6, so remove the ARCHLVL parameter from parmlib member LOAD*xx* once your migration to z/Architecture is complete, letting ARCHLVL default to 2.
21. Determine if you want to control whether DFSMSdss obtains I/O buffers backed by 64-bit real storage. By default as of z/OS V1R5, DFSMSdss obtains 64-bit backed I/O buffers when running in z/Architecture mode. If you want DFSMSdss to obtain 31-bit backed I/O buffers, specify EXEC PARM='ZBUFF64R=NO' in the JCL or API invocation. The DFSMSdss ZBUFF64R parameter is new in z/OS V1R5. **Note:** Because of APAR OA11139 on z/OS V1R5 and z/OS V1R6, and APAR OA12033 on z/OS V1R7, DFSMSdss has been modified to obtain I/O buffers backed below the 2 GB bar when a tape device does not support I/O buffers backed above the bar. You can still use ZBUFFR64 for tape devices if you would like to force buffers below the bar, but it is no longer necessary; if the tape device support is not present, DFSMSdss automatically obtains the buffers below the bar.

General Migration Actions Post-First IPL

<none>

© Copyright IBM Corporation, 2007

Some BCP Migration Actions for z/OS V1R8 from z/OS V1R5

Migration Actions You Can Do NOW:

- Evaluate your stand-alone dump data set allocations and your IPCS processing of them (Recommended)
  - R6 introduced support for ext format ds >64k trks/volume, install OA08339, ...
- Migrate from the prelinker to the program management binder (Recommended)
  - Prelnker has been stabilized. Use *z/OS MVS Program Management: User's Guide and Reference*
- Remove ECMB=NO from IEAOPTxx (Required-IF, as of R7)
  - Affected programs (monitor programs) must use the Extended Channel Measurement Block.
- Prepare to use synchronous reserves in GRSCNFxx (Recommended)
  - Default value for SYNCHRES was NO, in R6 it's YES.
  - For preventing deadlocks between systems sharing DASD.
  - Checked in IBM Health Checker for z/OS to warn you if default is not used.
- Decide on the number of CPUs to bring online at IPL (Required, via OA09649)
  - If you want # CPUs/IFAs brought online at IPL = # defined in the initial processor counts for the LPAR, do nothing.
  - If you want #CPUs/IFAs brought online at IPL = same CPUs/IFAs that were online at end of last IPL, instal OA09649 and PRESCPU parameter in IEASYSxx.

Continued on next foil...

© 2007 IBM Corporation

Some BCP Migration Actions for z/OS V1R8 from z/OS V1R5

Migration Actions Pre-First IPL:

- Accommodate the removal of console switching (Required-IF, as of R8)
  - In R8, console failures no longer cause a console switch. Operators cannot issue SWITCH CN command to switch consoles. Remove ALTGRP keyword from CONSOLxx. (Can keep if sharing <R8. R8 warns, and continues). Cannot issue VARY CN(...),ALTGRP command.
- Discontinue use of the MSGRT command (Required-IF, as of R8)
  - In R8, MSGRTcommand was removed. Remove MSGRT keyword from CONSOLxx. (Can keep if sharing <R8. R8 warns, and continues)
- Update the hardcopy medium definition (Required-IF, as of R8)
  - Hardcopy switching is no longer supported. To reduce likelihood of losing hardcopy, define both SYSLOG and OPERLOG as the hardcopy medium.
- Modify programs that reference unsupported console functions (Required, as of R8)
  - In R8, many control blocks remove support for 1-byte console IDs, console switching, and master consoles. Examine your programs that reference these control blocks and ensure they will still work on R8. Clean **compile assures of successful removal of references to 1-byte fields.**
- GRS-related migration actions:
  - Use GRSCNFxx to specify the maximum number of concurrent ENQ requests (Required-IF, as of R8 and R7 APAR OA11382)
    - Before, you could zap the GRS vector table (GVT) to increase authorized and unauthorized maximums of concurrent ENQ requests. Now, use new ENQMAXA and ENQMAXU keywords in GRSCNFxx parmlib member.
  - Update the GRSQ setting in GRSCNFxx (Required-IF as of R8, and rolled back to R7 OA11382)
    - GRSQ keyword is now only applicable to STAR mode and its default is CONTENTION.
  - Ensure that GRSCNFxx is parsable for GRS=NONE (Required-IF, as of R8)

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

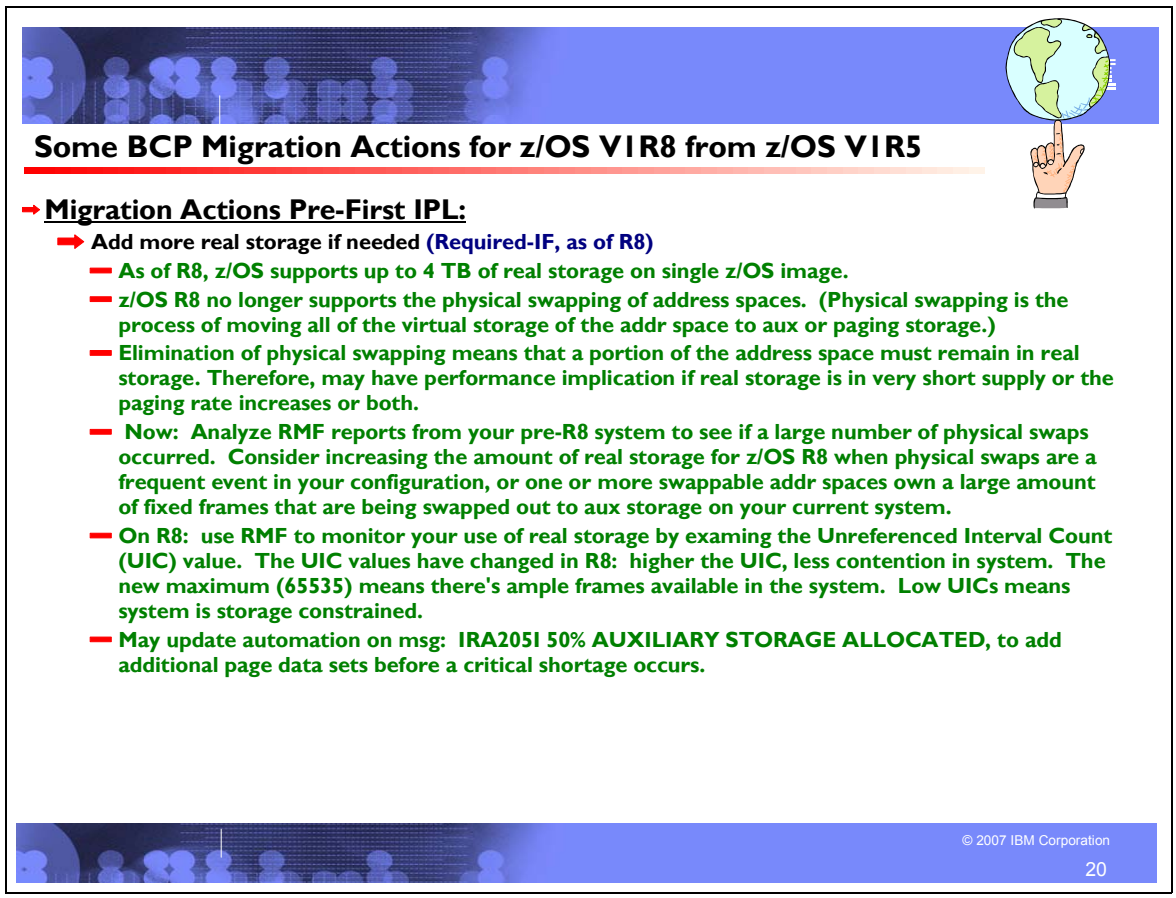

BCP Migration Actions Between z/OS V1R5 and z/OS V1R8

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

BCP Migration Actions You Can Do Now

Evaluate your stand-alone dump data set allocations and your IPCS processing of them (Recommended)

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

Migration action:

- Use multi-volume stand-alone dump data sets. Adjust the number of volumes and their separation to achieve tolerable stand-alone dump capture times.
- Use extended format data sets with more than 64 K tracks per volume (after first IPL on z/OS V1R6 or later). Copy their contents to an extended format, compressed, striped data set using the IPCS COPYDUMP subcommand prior to analysis. Use the same or a larger striping factor than you used for your stand-alone dump data sets.
- Install the PTF for APAR OA08339 for performance enhancements and an enhanced COPYDUMP subcommand.
- Use a large CISIZE and striping for IPCS dump directories, and blocking, striping, and compression for the stand-alone dump data set. Very large stand-alone dumps might require that you define your directory with the extended addressing attribute, allowing it to hold more than 4 GB.

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

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

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

Remove ECMB=NO from parmlib member IEAOPTxx (Required, as of R7)

Required if you use the ECMB=NO parameter of parmlib member IEAOPTxx.

Migration action: If you have ECMB=NO specified in parmlib member IEAOPT*xx*, convert user-written programs that use the CMB to use the ECMB, and contact ISVs to obtain updates to ISV programs that use the CMB. The type of programs that are most likely to use the CMB are monitor programs. Finally, remove your ECMB=NO specification.

Prepare to use synchronous reserver for GRSCNFxx (Recommended, as of R6)

Not required, but recommended to prevent deadlocks when reserving volumes.

Prior to z/OS V1R6, the default value for the SYNCHRES (synchronous reserve) statement on the GRSCNF*xx* (global resource serialization configuration) parmlib member was NO. Beginning with z/OS V1R6, the SYNCHRES statement defaults to YES to ensure that synchronous processing is activated at all times. This new default prevents deadlocks between systems sharing DASD, which could occur when there is a delay between the time a request is made to reserve a volume and the time when the reserve of the volume is obtained. Because of the importance of synchronous reserve processing, a check for SYNCHRES=YES has been added to the current level of IBM Health Checker for z/OS and Sysplex to warn you if the default is not used.

Migration action: Examine your current settings to determine what is specified in the GRSCNF*xx* parmlib member, then take the following action:

- If SYNCHRES=YES is specified, remove it because YES is now the default.
- If SYNCHRES=NO is specified, determine why and take the necessary action to use the new default of YES.
  - If it was the prior default, remove the SYNCHRES statement so the system defaults to YES.
  - If it was specified for an application's performance purposes in a STAR environment, convert the application's RESERVE requests to global ENQs. For RESERVE conversion, see *z/OS MVS Planning: Global Resource Serialization*.
  - If it was specified for an application's performance purposes in a RING environment, remove SYNCHRES=NO because the synchronous reserve overhead is minimal. For RESERVE conversion, see *z/OS MVS Planning: Global Resource Serialization*.
  - If it was specified for other application purposes, change the application to use ISGENQ SYNCHRES=NO. This affects only the RESERVE requests for this application. For more information about ISGENQ, see *z/OS MVS Programming: Assembler Services Reference IAR-XCT*.

To change the default for all systems before migrating to z/OS V1R6, use the SETGRS operator command. Specify SYNCHRES=YES and then update shared parmlib member GRSCNF*xx* to specify SYNCHRES=YES.

Decide on the number of CPUs to bring online at IPL (Required, as of OA09649)

Before z/OS V1R7, the number of CPUs and zAAPs brought online at IPL was the number defined in the initial processor counts for the logical partition profile. This was the case in z/OS V1R6 and z/OS V1R5, and in z/OS V1R4 with PTF UA03339 installed (for APAR OA03335), when running on z990, z890, z900 (GA-3), and z800 servers.

As of z/OS V1R7, you have the option of bringing online at IPL all CPUs and ZAAPs that were online at the end of the prior IPL (provided that specific actions are taken, as described below). This option is also furnished back to z/OS V1R4 by APAR OA09649.

Migration action: If you want the number of CPUs and ZAAPs brought online at IPL to be the number defined in the initial processor counts for the logical partition profile, no action is required.

If you want the CPUs and ZAAPs brought online at IPL to be the same CPUs and ZAAPs that were online at the end of the last IPL, then install the PTF for OA09649 (which is integrated into z/OS V1R7) and use the PRESCPU parameter for parmlib
member IEASYS*xx*. The PRESCPU parameter causes the system to ignore the number of initial CPUs and ZAAPs that PR/SM provides and bring logically online all CPUs and ZAAPs that the machine presents as physically online.

BCP Migration Actions Pre-First IPL

Create IPL text (Required)

Migration action: Update and run the IPLTEXT job to write a new copy of the IPL text. If you install z/OS V1R7 with a ServerPac, an installation dialog job is provided to perform this action. If you install z/OS V1R7 with a CBPDO, instructions to perform this action are provided in *z/OS Program Directory*.

© Copyright IBM Corporation, 2007

Reference information: For a sample IPLTEXT job, see *z/OS Program Directory*. ServerPac provides a similar job for accomplishing this task; see *ServerPac: Installing Your Order*.

Reassemble the stand-alone dump program (Required)

Migration action: Reassemble the stand-alone dump program. If you install z/OS V1R7 with a ServerPac, an installation dialog job is provided to perform this action. If you install z/OS V1R7 with a CBPDO, instructions to perform this action are provided in *z/OS Program Directory.* Once the stand-alone dump program is properly created on a DASD residence volume, it resides in the SYS1.PAGEDUMP.Vvolser data set.
Reference information: *ServerPac: Installing Your Order, z/OS Program Directory*, and *z/OS MVS Diagnosis: Tools and Service Aids*

Accommodate the removal of 1-byte console IDs (Required)

In z/OS V1R7, support for 1-byte console IDs and migration console IDs was removed from z/OS macros and commands. In z/OS V1R8, the removal of 1-byte console IDs is complete. In some cases, z/OS V1R8 might continue to accept 1-byte console IDs but will treat them differently, so programs that use 1-byte console IDs might behave unexpectedly. The Console ID Tracking Facility is still available to help identify 1-byte console ID usage.
Migration action:

- Obtain the latest product and service levels for IBM and vendor products that used 1-byte console Ids.
- Consider adding the UNKNIDS keyword to console definitions specified in your CONSOL*xx* parmlib member. z/OS V1R8 has a new routing attribute, UNKNIDS, and any consoles that request this attribute receive messages directed to 1-byte console Ids.
- Because some control blocks and macros have been changed to remove 1-byte console ID support, avoid using 1-byte console ID fields when referencing these control blocks. For a list of changed control blocks and macros, see "**Modify programs that reference unsupported console functions**".

Accommodate the removal of the master console (Required-IF, as of R8)

Required if you have operational procedures that are dependent on the master console.
Before z/OS V1R8, one operator console within a sysplex was required to be designated as the master console. As of z/OS V1R8, the concept of the master console is removed and functions that were unique to the master console have been made available to other consoles. You can still define multiple consoles as having master authority. With the removal of the master console, the system console takes on more significance. z/OS V1R8 enforces that the system console has master authority and LOGON(OPTIONAL).

When the first z/OS V1R8 system joins a sysplex consisting of z/OS V1R7 (or earlier) systems, support for the master console is removed for all systems in the sysplex. When the last z/OS V1R8 system leaves the sysplex, a master console may again be selected.
Migration action:

- Remove the NOCCGRP keyword from the CONSOL*xx* parmlib member and the *MSTCON* keyword from the CNGRP*xx* parmlib member. You can keep these keywords if you plan to share these members with z/OS levels earlier than V1R8. z/OS V1R8 issues warning messages about unsupported keywords but the system continues with no errors.
- Notify operators that the following commands are no longer supported:
  - – VARY ...,MSTCONS
  - – DISPLAY C,MCONLY
- Examine your automation procedures for commands, functions, and messages that no longer exist, and make appropriate changes. Also, determine if any new messages warrant automation.
- Because some control blocks and macros have been changed to remove master console support, avoid using master console fields when referencing these control blocks. For a list of changed control blocks and macros, see "**Modify programs that reference unsupported console functions**".
- Consider adding the INTIDS keyword to console definitions in your CONSOL*xx* member of parmlib. z/OS V1R8 has a new routing attribute, INTIDS, and any consoles that request this attribute receive messages directed to console ID zero. (Console ID zero originally represented the master console.)
- Determine if your operational procedures refer to the use of the external interrupt key. If so, update your procedures accordingly. Use of the external interrupt key to switch the master console is no longer supported.
- Examine your automation procedures for the no-master-console and no-consoles conditions, and make appropriate changes. These conditions no longer exit.

Have you got COM='MONITOR JOBNAME,T' or something like that in your COMMNDxx?

© Copyright IBM Corporation, 2007

If so, then take notice of this migration action:
In z/OS R7 and R8 *Init and Tuning Reference* for COMMNDxx, it states: Do not use COMMNDxx to enter MONITOR commands. Instead, use the SETCON MONITOR command to enable or disable Monitor messages. The SETCON MONITOR command is used to control the monitoring of messages in your installation. The MONITOR option allows you to receive monitored messages without requiring that the messages be queued to a console or be written to SYSLOG or OPERLOG.

Since all the 1-byte console ID support has been removed as of z/OS R8, if COM='MN JOBNAMES,T' is specified in COMMNDxx, then during ipl the message CNZ0005I MONITOR REJECTED.REASON=CONSOLE ID ZERO NOT SUPPORTED. is issued. Review your COMMNDxx specification to see if you use the MONITOR command, and change it to the SETCON MONITOR command.

Accommodate the removal of the console switching (Required-IF, as of R8)

Required if you have automation or operational procedures that process console failures.
Before z/OS V1R8, when a console failed, z/OS would attempt to switch the console's function to an alternate console. This facility was originally designed to ensure that a master console was always active. As of z/OS V1R8, this console switching function is removed, and console failures no longer cause a console switch. Operators cannot issue the SWITCH CN command to switch consoles. In addition, use of ALTGRP to specify an alternate console is no longer possible, and the VARY CN(...),ALTGRP command is removed.

When a z/OS V1R8 system joins a sysplex consisting of z/OS V1R7 (or earlier) systems, support for console switching is removed for all systems in the sysplex. When the last z/OS V1R8 system leaves the sysplex, console switching becomes operative.
Migration action:

- Explicitly configure backup consoles because console switching is not supported.
- Notify operators that the SWITCH CN and VARY CN(...),ALTGRP commands cannot be issued.
- Modify any automated procedures to eliminate use of SWITCH CN and VARY CN(...),ALTGRP commands.
- Remove the ALTGRP keyword from the CONSOL*xx* parmlib member. You can, however, keep the ALTGRP keyword if you plan to share CONSOL*xx* members with z/OS levels earlier than V1R8. z/OS V1R8 issues warning messages about unsupported keywords but the system continues with no errors.
- Because some control blocks and macros have been changed to remove console switching support, avoid using console switching fields when referencing these control blocks. For a list of changed control blocks and macros, see "**Modify programs that reference unsupported console functions**".

Discontinue use of the MSGRT command (Required-IF, as of R8)

Required if you have operational procedures that suggest using the MSGRT command.
Before z/OS V1R8, the MSGRT command was used to establish and change message routing instructions for certain options of DISPLAY, CONFIG, CONTROL, MONITOR, and STOPMN commands. As of z/OS V1R8, the MSGRT command is removed. The MSGRT command was originally developed for OS/360 and, over the years, has lost its usefulness.
Migration action:

- Notify operators that the MSGRT command can no longer be used.
- Modify any automated procedures to eliminate the use of the MSGRT command. If you want, update the commands (with the L= operand) that MSGRT was going to process.
- Remove the MSGRT keyword from the CONSOL*xx* parmlib member. You can, however, keep the MSGRT keyword if you plan to share the CONSOL*xx* member with z/OS levels earlier than V1R8. z/OS V1R8 issues warning messages about unsupported keywords but the system continues with no errors.

Update the hardcopy medium definition (Required-IF, as of R8)

Required if you require that a hardcopy medium always be active.
Before z/OS V1R8, if the specified hardcopy medium (SYSLOG or OPERLOG) failed, the system would attempt to switch hardcopy processing to the other log. As of z/OS V1R8, hardcopy switching is no longer supported. Therefore, if only one hardcopy medium is defined and fails, the system suspends hardcopy processing. To avoid this problem, define both SYSLOG and OPERLOG as the hardcopy medium.
Migration action: To reduce the likelihood of losing hardcopy, define both SYSLOG and OPERLOG as the hardcopy medium. In your CONSOL*xx* parmlib member, specify HARDCOPY DEVNUM (SYSLOG,OPERLOG). The default is

© Copyright IBM Corporation, 2007

HARDCOPY DEVNUM (SYSLOG). When specifying OPERLOG, ensure that you have set up DASD-only or coupling facility log streams.

Modify programs that reference unsupported console functions (Required)

As of z/OS V1R8, many control blocks have been changed to remove the support for 1-byte console IDs, console switching, and master consoles. Examine your programs that reference these control blocks and ensure that the programs continue to work with the z/OS V1R8 changes.

Migration action: Examine programs that reference the control blocks below and ensure that the programs continue to work now that 1-byte console ID support has been removed from the control blocks. Because the 1-byte fields have been removed, a clean assembly or compilation will assure you that you were successful in removing all references to 1-byte fields.

Control block	Mapping macro
MDB	IEAVM105
CSCB	IEECHAIN

UCM	IEECUCM
XSA	IEEXSA
CMPL	IEEZB810
PLIBAREA	IEEZB822
VDRSARR	IEEZB834
ECDM	IEEZB887
CIB	IEZCIB
CXSA	IEZVD002 and IHACTM
SCSR	IEZVG100
MCSOPPRM	IEZVG111
CONV	IEZVG200
CTXT	IEZVX100
CMDX	IEZVX101
CXSA	IHACTM and IEZVD002
ORE	IHAORE
WQE	IHAWQE

Accommodate the subsystem console service interface change (Required-IF, as of R8)

Required if you have applications that use subsystem consoles.

Callers allocate, release, query, and modify subsystem consoles by invoking module IEAVG700 with an SCSR parameter list (mapped by IEZVG100). As of z/OS V1R8, the parameter list for IEAVG700 has been changed, resulting in different subsystem console interface behavior. If a subsystem console request requires a console ID as input or output, a 4-byte ID field (SCSCONID) must be provided. A demand select request (SCSDEMSL) must now request only a subsystem console, and IEAVG700 can no longer be used to change an MCS console into a subsystem console. IEAVG700 no longer supports the changing of a subsystem console's routing codes.

Migration action: Modify IEAVG700 as follows:

- Reassemble it with the current length and version level. The parameter list's version level must also be at or above a level (z/OS V1R6 or later) that supports the 4-byte console ID field.
- Make sure that the demand select request field (SCSDEMSL) is being used only for a subsystem console. If the unit name is specified, remove the request to change an MCS console.
- If the SCSRTCDF field is being set, remove the request for the routing code service.

Use GRSCNFxx to specify the maximum number of concurrent ENQ requests (Required-IF, as of R8 and rolled back to R7 in OA11382)

Required if you have any programs that use the GVTCREQA or GVTCREQ field of the GVT.

Before z/OS V1R8, you could zap the GVTCREQA and GVTCREQ fields of the global resource serialization vector table (GVT) to increase the authorized and unauthorized maximums of concurrent ENQ requests. Starting with z/OS V1R8, and

© Copyright IBM Corporation, 2007

rolled back to V1R7 by APAR OA11382, you cannot use this method. Instead, achieve the same system-wide updates using the new ENQMAXA and ENQMAXU keywords in the GRSCNF*xx* parmlib member.

Migration action:

1. Remove any zaps that update GVTCREQ and GVTCREQA and make corresponding updates to the ENQMAXU and ENQMAXA keywords of GRSCNF*xx* before the next IPL.
2. Determine which subsystems and applications require greater authorized or unauthorized maximums than the defaults. To help in this investigation, use ISGQUERY REQINFO=ENQSTATS to query address space peaks.
3. Ensure that these programs use the new ISGADMIN REQUEST=SETENQMAX service to set address space specific maximums so that lower system-wide maximums can be maintained, providing better reliability, availability, and serviceability (RAS).
4. Once all the necessary programs exploit ISGADMIN appropriately, reset the ENQMAXU and ENQMAXA keywords to their default values by using the new SETGRS command support. This command improves reliability, availability, and serviceability because global resource serialization can better detect "ENQ runaways" (address spaces that loop on ENQ requests). Make corresponding updates to GRSCNF*xx* for future IPLs.

Update the GRSQ setting in GRSCNFxx (Required, as of R8)

The GRSQ keyword was added to the GRSCNF*xx* parmlib member in z/OS V1R7 (and rolled back to z/OS V1R4 by APAR OA07975), where it was applicable to all global resource serialization modes (STAR, RING, and NONE) and its default was ALL. Starting in z/OS V1R8 (and rolled back to z/OS V1R7 by APAR OA11382), the GRSQ keyword is only applicable to STAR mode and its default is CONTENTION.

Migration action:

- If global resource serialization **is not** running in STAR mode, ensure that the GRSQ parameter is not in parmlib member GRSCNF*xx*.
- If global resource serialization **is** running in STAR mode and you still want GRSQ set to ALL (rather than the new default, CONTENTION), specify GRSQ(ALL) in parmlib member GRSCNF*xx*.

Ensure that GRSCNFxx is parsable for GRS=NONE (Required-IF, as of R8)

Required if you use global resource serialization in NONE mode.

Before z/OS V1R8, the GRSCNF*xx* parmlib member was not parsed when global resource serialization was operating in NONE mode. Starting with z/OS V1R8, GRSCNF*xx* is parsed when IEASYS*xx* keyword GRS is set to NONE. You therefore need to have GRSCNF*xx* accessible and parsable.

In addition, the following keywords are now relevant for all modes including GRS=NONE: SYNCHRES, ENQMAXA, ENQMAXU.

Migration action: To avoid warning messages, ensure that the GRSCNF*xx* member is accessible. If global resource serialization is not active and GRSCNFxx is not accessible, the following messages can be issued:

- The system issues ISG313I when mode-irrelevant keywords are found in GRSCNF*xx* parsing.
- The system issues ISG372E when GRSCNF*xx* cannot be accessed.

Modify programs to properly handle the new ASCBIOSC maximum value (Required-IF, as of R8 and rolled back to R6 in OA14340)

Required if you have any programs that treat ASCBIOSC, SMF30TEP, SMF32EXP, or SMF33EXP as signed integers.

The ASCBIOSC field in the ASCB data area contains an ongoing I/O count. The field is 4 bytes in length and is updated by the SMFIOCNT service. The ASCBIOSC count is propagated to SMF record type 30 (field SMF30TEP), SMF record type 32 (field SMF32EXP), and SMF record type 33 (field SMF33EXP).

The maximum ASCBIOSC value has been increased. In z/OS V1R7 without the PTF for APAR OA14340 installed, z/OS V1R6 without the PTF for APAR OA14340 installed, and z/OS V1R5, the maximum ASCBIOSC value is 2 147 483 647 (2 GB, or X'7FFFFFFF'). In z/OS V1R8, z/OS V1R7 with the PTF for APAR OA14340 installed, and z/OS V1R6 with the PTF for APAR OA14340 installed, the maximum ASCBIOSC value is 4 294 967 294 (4 GB, or X'FFFFFFFE').

Because of the new maximum ASCBIOSC value, when ASCBIOSC has reached its maximum capacity, the SMF fields to which ASCBIOSC is propagated (SMF30TEP, SMF32EXP, and SMF33EXP) are now cleared, and new bit indicators are set to identify that the fields are invalid. The new bits are defined in fields SMF30DCF,SMF32DSF, and SMF33DSF. The ASCBIOSC field and the SMF fields to which it is propagated are now unsigned 4-byte fields, and must be treated as unsigned by any programs that use them.

Migration action: Modify programs so that they treat the ASCBIOSC field and the SMF fields to which it is propagated (SMF30TEP, SMF32EXP, and SMF33EXP) as *unsigned* fields. Use logical instructions when using these fields for arithmetic. If the values of these fields are displayed, be aware that the high-order bits might be on, which does *not* indicate a negative value.

© Copyright IBM Corporation, 2007

Modify programs so that they check the new bits SMF30TEF in SMF30DCF, SMF32EXF in SMF32DSF, and SMF33EXF in SMF33DSF before using the corresponding SMF30TEP, SMF32EXP, and SMF33EXP fields because these fields will be cleared and invalid once the ASCBIOSC has reached its maximum capacity.

Modify installation exits to ignore new or unknown C/I text units (Required-IF, as of R8)

Required if you have any installation exits that are not coded to ignore new or unknown C/I text units.
Before z/OS V1R8, converter/interpreter (C/I) text unit key X'14' was not defined. In z/OS V1R8, C/I text unit key X'14' is introduced to be used by JES2 and JES3 to associate SYSIN data with the correct file when a SYSIN statement is coded as an override to a DD statement that is in a procedure. If you have any IEFUJV or other exit routines that process C/I data, you might have to change them if they are not coded to ignore new or unknown text units.
Migration action: Verify that your installation exits that process C/I data are coded to ignore new or unknown text units. If they are not, make the appropriate modifications to the exit routines.

Check for reserved console names INTERNAL and INSTREAM (Required-IF, as of R8)

Required if you depend on the CONVCON service returning RC=x'8', RSN=x'C' for console names of INTERNAL or INSTREAM.
Before z/OS V1R8, the CONVCON service returned RC=X'8' (invalid console name) and RSN=X'C' (console name reserved) for a console name of INTERNAL or INSTREAM. As of z/OS V1R8, CONVCON returns RC=X'0' with the appropriate value in the 4-byte console ID field (CONVID). Therefore, if you depend on CONVCON returning RC=X'8' and RSN=X'C', you must make appropriate changes to your code.
Migration action:

- Verify that your code does not depend on the CONVCON service returning RC=X'8' (invalid console name) and RSN=X'C' (console name reserved) for console names of INTERNAL and INSTREAM. If it does, update your code to check for these console names and take appropriate action.
- Verify that your code can handle receiving a CONVID of X'00000000' for a console name INTERNAL and X'00000080' for a console name INSTREAM. If it does not, make appropriate changes to your code.

Put a blank character in column 19 of NLS skeleton lines (Required-IF, as of R8)

Required if you have any NLS skeleton lines, other than comment lines, that contain a non-blank character in column 19.
Before z/OS V1R8, column 19 of national language support (NLS)
skeletons was documented as "not part of the intended programming interface", but the message services compiler did not enforce that it contained a blank. As of z/OS V1R8, the message services compiler uses column 19 for a specific purpose when the value is 1. The compiler is changed to validate that column 19 is either blank or 1. Any other value results in a nonzero completion code from the NLS skeleton compile process. Therefore, if you have a nonblank character in column 19, you need to change it to a blank. Note that column 19 in a comment line can contain any value, including nonblank.
Migration action: Verify that column 19 in NLS skeleton lines contains a blank character. If it does not, modify the skeleton line so that column 19 contains a blank.

Use the changed defaults for HOTIO in IECIOSxx (Recommended, as of R8)

Not required, but recommended unless the installation can properly deal with responding to messages issued using DCCF.
Before z/OS V1R8, the defaults for the DFLT112 and SDFT112 parameters of the HOTIO statement in the IECIOS*xx* parmlib member were: DFLT112=(CHPK,OPER) and SDFT112=(CUK,OPER) .
These defaults allowed the operator to specify the recovery action. As of z/OS V1R8, the DFLT112 and SDFT112 defaults are changed as follows: DFLT112=(CHPK,BOX) and SDFT112=(CUK,BOX) .
These defaults specify that the device is to be forced offline without operator intervention. There are conditions when the message issued to prompt the operator might be done using the Disabled Console Communication Facility (DCCF). IBM recommends that operator involvement not be requested for HOTIO actions.
Migration action: Use the new HOTIO defaults if your installation cannot properly deal with responding to messages issued using DCCF.

Ensure that IXGWRITE in unauthorized programs correctly handles reason code x'0867' (Required-IF, as of R8 and rolled back in OA14125)

Required if you issue the IXGWRITE macro in unauthorized programs.
The IXGWRITE macro allows a program to write a log block to a log stream. IXGWRITE returns a unique identifier for each log block written to the log stream.

© Copyright IBM Corporation, 2007

Before z/OS V1R8, IXGWRITE reason code X'0867' indicated that available local buffer space for the system logger address space was full. With z/OS V1R8, the X'0867' reason code is updated to indicate also that the IXGWRITE request is rejected when an unauthorized caller attempts to write log data while the outstanding asynchronous write activity for this log steam connection is considered too high. Therefore, ensure that the X'0867' reason code is handled appropriately.

Migration action:

- Review the use of IXGWRITE in unauthorized programs.
- Ensure that the existing X'0867' reason code is managed appropriately. In the IXGANSAA answer area, Ansaa_Diag1 contains 1 for this error and Ansaa_Diag2 contains the total number of outstanding asynchronous write requests for this connection. The unauthorized writer should consider waiting for a short interval and then reissue the IXGWRITE request. If subsequent write attempts continue to fail for an unacceptable period, consider notifying operations or disconnecting from the log stream.

Limit the number of browse sessions initiated by IXGBRWSE in unauthorized programs (Required-IF, as of R8 and rolled back in OA08661)

Required if you use IXGBRWSE in unauthorized programs.

The IXGBRWSE macro is used to read and browse a log stream for log block information.

Before z/OS V1R8, the number of browse sessions initiated by IXGBRWSE REQUEST=START was not limited. As of z/OS V1R8, IXGBRWSE REQUEST=START can be rejected if an unauthorized caller attempts to start a browse session when 100 or more browse sessions already exist for this connection, or if more than 20 consecutive unused browse sessions are detected for this connection. Therefore, you should review your use of IXGBRWSE in unauthorized programs to determine whether it is possible that your programs use too many browse sessions.

Migration action: Review the use of IXGBRWSE in unauthorized programs to determine whether it is possible that your programs use too many browse sessions, and modify your programs if necessary. Reason code X'0845' for IXGBRWSE REQUEST=START now indicates that the request can be rejected if an unauthorized caller attempts to start a browse session when 100 or more browse sessions already exist for this connection, or if more than 20 consecutive unused browse sessions are detected for this connection. Field ANSAA_BROWSESTARTSLIMITED in answer area IXGANSAA indicates to users of IXGCONN REQUEST=CONNECT whether IXGBRWSE REQUEST=START requests are limited for unauthorized users.

Add more real storage if needed (Required-IF, as of R8)

Required if insufficient real storage is configured to support the workload without physical swapping.

Before z/OS V1R8, the amount of real storage that z/OS allowed to be configured to a single z/OS image was 128 GB (where GB equals 1 073 741 824 bytes). The IBM hardware systems that are now available, however, support more than 128 GB. As of z/OS V1R8, z/OS supports up to 4 TB (where TB equals 1 099 511 627 776 bytes) of real storage to be configured to a single z/OS image.

With this support, z/OS no longer supports the physical swapping of address spaces. Physical swap is the process of moving all of the virtual storage of the address space to auxiliary or paging storage. Elimination of physical swap means that a portion of the address space must remain in real storage. There can be performance implications for some systems if real storage is in very short supply or the paging rate increases, or both.

In some systems, physical swapping releases real storage that will no longer be released because physical swapping is no longer supported. Examination with RMF can determine the average number of physically swapped address spaces and the average number of pages in the swap working set. All of the pages of the physically swapped address space working set will not be backed in real storage; generally, one third to one half will reside in real storage, with the remaining paged to auxiliary storage. If the total is significant, the paging rate may increase. If the increased paging is not advisable, you can increase the real storage to the logical partition.

Currently, the hardware supports the following maximum amounts of real storage for a z/OS image (where GB equals 1 073 741 824 bytes): z9 EC: 512 GB , z9 BC: 64 GB , z990: 256 GB , z890: 32 GB , z900: 64 GB , z800: 32 GB .

Migration action:

- Analyze RMF reports from your pre-z/OS V1R8 system to determine whether a large number of physical swaps occurred. (Physical swaps to auxiliary storage have been eliminated in V1R8. Therefore, you need to increase the amount of real storage to accommodate the fixed frames owned by an address space.) The RMF CPU Activity Report specifies how many physical swaps have occurred in the OUT READY and OUT WAIT Queue Type fields.
  For each address space listed on the OUT READY or OUT WAIT queue, use the RMF STORF Report to find out the number of fixed frames that each address space owns when the address space is Swapped-IN.
  Consider increasing the amount of real storage on your z/OS V1R8 system when:
  – Physical swaps are a frequent event in your configuration.

© Copyright IBM Corporation, 2007

– One or more swappable address spaces own a large amount of fixed frames that are being swapped out to auxiliary storage on your current system.

- After migrating to z/OS V1R8, use RMF to monitor your use of real storage by examining the Unreferenced Interval Count (UIC) value. The UIC values have changed in V1R8, in that the maximum value has increased from 2540 to 65535. The higher the UIC value, the less contention there is for storage in the system. A UIC value of 65535 indicates that ample frames are available in the system. The lower the UIC value, the more contention there is for storage in the system. A very low UIC value indicates that the system is storage constrained.
  Changes to the UIC value appear in both RMF reports and SMF records. The "old" UIC values are maintained for compatibility by RMF Monitor I in SMF record type 71. Additional fields are added to SMF record type 71 to contain the "new" UIC values. RMF replaces the "old" UIC value with the "new" UIC value in Monitor II and Monitor III reports, as well as in SMF record type 79 subtype 3 (storage/processor data) and subtype 4 (paging activity data) fields.
- Update automation for message IRA205I 50% AUXILIARY STORAGE ALLOCATED. This message is issued when more than 50 percent of the auxiliary storage slots are in use. This allows an automation program to add additional page data sets before a critical shortage occurs. The message is repeated every two hours as long as the auxiliary slot usage is above 50 percent.

Remove CUNUNIxx parmlib members (Recommended, as of R7, with additional function in R7 APAR OA14231)

Not required, recommended to help avoid confusion in the future by anyone who encounters Unicode parmlib members and is unaware that they are no longer necessary.

As of z/OS V1R7, enhancements made to z/OS support for the Unicode Standard no longer require you to run the image generator program to create a conversion image and to customize the CUNUNI*xx* parmlib members with that conversion image name. Initially, the z/OS V1R7 enhancements were only for callers that run in TCB mode. APAR OA14231 (which was made available on z/OS V1R7 and integrated in z/OS V1R8) has added function for SRB mode callers. Now callers of Unicode Services that run in either TCB mode or SRB mode will have their conversion tables loaded when needed.

All tables needed for character conversion, case conversion, normalization, and collation services are now loaded automatically when they are required and not already present in storage. The only failure that will occur is when an invalid or unsupported conversion or service is specified.

This support is only available starting with z/OS V1R7 with APAR OA14231. If a parmlib is shared between z/OS V1R7 and prior levels, the CUNUNI*xx* parmlib members still need to be made available to the pre-z/OS V1R7 systems.

Migration action: You no longer need to provide parmlib members CUNUNI*xx*. You can delete them. Remember also to remove your UNI= references in any IEASYS*xx* parmlib members. New operator commands (SETUNI ADD, SETUNI DELETE, SETUNI REPLACE, and SETUNI REALSTORAGE) are provided to manipulate conversion images and Unicode settings.

Manage zAAP dispatching with IFAHONORPRIORITY, not IFACROSSOVER, in IEAOPTxx (Required-IF, as of the Web Deliverable for zIIP, and integrated into R8)

Required if you use zAAPs.

Before z/OS V1R8, the IEAOPT*xx* parmlib member contained the IFACROSSOVER=YES|NO parameter. This parameter was introduced in support of the IBM System z9 and zSeries Application Assist Processors (zAAPs). It specified whether work eligible for zAAPs should be examined by standard processors when the standard processors would otherwise enter a wait state. The default was YES. A related parameter, IFAHONORPRIORITY=YES|NO, was made available at the same time as IFACROSSOVER.

As of z/OS V1R8 (and the IBM zIIP Support for z/OS and z/OS.e V1R6/R7 Web deliverable), the IFACROSSOVER parameter is no longer honored because of a general lack of interest and the increased complexity it caused. Now, only the IFAHONORPRIORITY parameter is available to manage the zAAP dispatching algorithms. The default is YES, which specifies that zAAP-eligible work should run on standard central processors (CPs) only when the zAAP is "overloaded". This scheme is designed to use the speciality engine fully while not significantly impacting overall response time. If z/OS determines that response time is being impacted, zAAP work flows back to standard CPs, if they are available, at the priority assigned to the unit of work. If you explicitly specify NO, zAAP-eligible work never overflows to a standard CP. In this case, the overall response time might be impacted if not enough zAAPs are defined to the configuration.

Note that the setting of the IFAHONORPRIORITY parameter has no effect on System z9 Integrated Information Processor (zIIP) speciality engines. zIIPs always have the equivalent of IFAHONORPRIORITY=YES processing.

Migration action: Make sure that the IFAHONORPRIORITY parameter is set to YES (or allowed to default to YES) if you want zAAP work to flow back to standard Cps whenever the zAAPs are fully utilized.

© Copyright IBM Corporation, 2007

Remove the IFACROSSOVER parameter from IEAOPT*xx*. If you do not remove it, it is ignored and no error message is generated. Nevertheless, removing it can help avoid confusing anyone who encounters it in the future and is unaware that it is no longer functional.

Update programs that use SMF record type 88, subtype 1 (Required-IF, as of R7)

Required if you run any programs that use explicit offsets from the beginning of the SMF88 record to the section in the table below.

Migration action: Because the Events section (SMF88ESD) of SMF record type 88, subtype 1 has increased in size, offset values to some sections have changed. Any programs that use an explicit offset from the beginning of the SMF88 record to a field in the following section might need to be updated:

Section	Related offset	Related DSECT
Structure (interim storage)	SMF88SOF	SMF88SSD

Check any programs that use the section listed above to determine if you must make any changes. If necessary, modify the programs using the updated offsets.

Accommodate the change from GMT to UTC for timer related messages (Required-IF, as of R7)

Required if you have automation routines that examine message text or if you want operators to use UTC instead of GMT.

Before z/OS V1R7, timer related WTO and WTOR messages included both the local time and Greenwich Mean Time (GMT). Starting with z/OS V1R7, the local time continues to be included but Greenwich Mean Time is replaced by Coordinated Universal Time (UTC). For example, here is message IEE136I from the D T command:

- **z/OS V1R6**: IEE136I LOCAL: TIME=17.47.56 DATE=2006.186 GMT: TIME=08.47.56 DATE=2006.186
- **z/OS V1R7** and later: IEE136I LOCAL: TIME=17.42.16 DATE=2006.186 UTC: TIME=08.42.16 DATE=2006.186

As another example, here is WTOR message IEA888A from IPL/NIP:

- **z/OS V1R6**: IEA598I TIME ZONE = W.00.00.00
  IEA888A GMT DATE=2005.273,CLOCK=09.52.20
  IEA888A LOCAL DATE=2005.273,CLOCK=09.52.20 REPLY U, OR GMT/LOCAL TIME
  R 00,'[DATE=yyyy.ddd][,CLOCK=hh.mm.ss][,GMT]'
- **z/OS V1R7** and later: IEA598I TIME ZONE = E.00.00.00
  IEA888A UTC DATE=2006.177,CLOCK=03.49.17
  IEA888A LOCAL DATE=2006.177,CLOCK=03.49.17 REPLY U, OR UTC/LOCAL TIME
  R 00,'[DATE=yyyy.ddd][,CLOCK=hh.mm.ss][,UTC|GMT]'

Migration action: Change your automation to look for "UTC" instead of "GMT". Tell operators to use UTC instead of GMT. "GMT" is still acceptable, for example, in the reply to WTOR message IEA888A, but using the new term ("UTC") is good practice.

Update programs that use SMF record type 90, subtypes 5,9,13, and 15 (Required-IF, as of R6)

Required if you run any programs that use explicit offsets from the beginning of the SMF90 record to any of the sections affected.

Migration action: Because the "IPL SMF/SET SMF/SETSMF" section of SMF record type 90, subtypes 5, 9, 13, and 15 has increased in size, offset values to some sections have changed. Any programs that use an explicit offset from the beginning of the SMF90 record to a field in any of the following sections might need to be updated:

Section	Related offset	Related DSECT
SMF data set	SMF90ODA	SMF90DSE
Subsystem record	SMF90OWK	SMF90WCH
Subsystem parameter	SMF90OOT	SMF90SUB

Check any programs that use any of the sections listed above to determine if you must make any changes. If necessary, modify the programs using the updated offsets.

Discontinue use of the TRACK and related commands (Required-IF, as of R7)

Required if you use the TRACK command or any of the related commands listed below.

Migration action: As of z/OS V1R7, the TRACK command has been removed from z/OS. Operators can no longer issue the TRACK command to display job information on an MCS or SMCS console. In addition, other commands that deal exclusively with the TRACK command have also been removed. These commands are:

- TRACK
- STOPTR
- CONTROL T (K T)
- CONTROL D,H (K D,H)

© Copyright IBM Corporation, 2007

- CONTROL D,U (K D,U)
- MR TR={a,name-a}

Because CONTROL T was removed, there is no longer a need to specify the UTME value for an MCS or SMCS console in CONSOL*xx*. Therefore, the UTME keyword is no longer accepted in CONSOL*xx*.

Update ENF listen exits that make use of logger ENF48 signals (Required-IF, as of R7)

Required if you use IXGENF mapping or listen for ENF signal 48.

Installation programs and ISVs that exploit the following system services are impacted if they need to process logger ENF48 log stream events.

Migration action: Programs that listen to ENF signal 48 (ENFPC048) for any of the following logger events need to be changed to additionally listen for the new ENF information:

- *Background:* After the last disconnect to a logstream on a system that is not disconnected due to other reasons (such as Force or XES recommendation), an IxgenfLogStreamConnDisc event is issued with new event reason IxgEnfSystemLevelDisc, and the connection count of remaining systems connected to the logstream is given (IxgEnfConnDiscCount). The affected logstream no longer has conections on this system.
  What to do: Change programs to take into account new event reason IxgenfSystemLevelDisc and the connection count IxgenfConnDiscCount.
- *Background:* After a successful SETLOGR FORCE,DISCONNECT command operation, existing event IxgEnfLogstreamsNotAvailable is issued with new event reason IxgEnfSetLogrForceDisconnect along with existing event specific information IxgEnfLogstreamDisconnected. This means that the affected logstream has lost all of its connections on this system.
  What to do: Change programs to take into account new event reason IxgenfSetLogrForceDisconnect for specific reason IxgenfLogstreamDisconnected.
- *Background:* After a successful SETLOGR FORCE,DELETE command operation, existing event IxgenfLogStreamDelete is issued with new event reason IxgEnfSetLogrForceDelete. This means that the logstream referred to by field IxgenfInventoryDelLogStreamName has been deleted from the LOGR couple data set as a result of the command request.
  What to do: Change programs to take into account new event reason IxgenfSetLogrForceDelete.

Evaluate the MAXMSG value in COUPLExx (Required-IF, as of R7)

Required if you specified a value for MAXMSG that is less than the default value of 2000.

The MAXMSG parameter in parmlib member COUPLE*xx* specifies a value that XCF uses to determine the allotment of message buffers when the MAXMSG parameter is not specified on any one of following:

- The CLASSDEF statement
- The PATHIN statement
- The SETXCF START,CLASSDEF command
- The SETXCF START,PATHIN command

Before z/OS V1R7, the default value for MAXMSG was 750. As of z/OS V1R7, the default value for MAXMSG has been increased to 2000.

Migration action: Examine the MAXMSG setting in the COUPLE*xx* parmlib member. If a value of less than 2000 is specified, either change it to a value greater than 2000 or accept the default value. A check for the value of MAXMSG has been added to the IBM Health Checker for z/OS to warn you if the MAXMSG value is less than the default.

Migrate to, or start using, the IBM Health Checker for z/OS (Recommended)

Not required, but is recommended in order to prevent potentional problems that could impact system availability and cause outages.

Migration action:

- If you are currently using the prototype, you should migrate to the IBM Health Checker for z/OS that is integrated as of z/OS V1R7. There are no incompatibilities between the prototype and the IBM Health Checker for z/OS framework; they can run in parallel. However, implementation of the IBM Health Checker for z/OS is strongly recommended and provides significant improvements in functionality compared to the prototype. For instructions on migrating from the prototype to the IBM Health Checker for z/OS, see *IBM Health Checker for z/OS: User's Guide*.
- If you did not use the prototype, it is recommended that you use the IBM Health Checker for z/OS that has been integrated as of z/OS V1R7. If you install with ServerPac, the setup is performed for you. If you do not install with

© Copyright IBM Corporation, 2007

ServerPac, you must perform all the setup steps. See *IBM Health Checker for z/OS: User's Guide* for instructions on setting up IBM Health Checker for z/OS.

- If you choose not to use the IBM Health Checker for z/OS, review the ServerPac-supplied jobs and procedures to remove the COM='START hzsproc' statement from the sample COMMND*xx* member provided with your order.

Update JCL that refers to the AMDSADDD utility (Required-IF, as of R7)

Required if you have any JCL that refers to the AMDSADDD utility.
In z/OS V1R7, as part of the enhancement to the IPCS 3.6 option that helps you create stand-alone dump data sets, AMDSADDD (the utility that allocates and initializes stand-alone dump data sets) has been moved to SYS1.SBLSCLI0.
Migration action: Update JCL that points to the AMDSADDD utility in SYS1.SAMPLIB to now point to SYS1.SBLSCLI0.

BCP Migration Actions Post-First IPL

Check the maximum number of XES connectors per address space (Required-IF, as of R7 w/ rollback via OA03194)

Required if you support connections to serialized coupling facility structures.
Before z/OS V1R7, the number of XES connections to a serialized coupling facility structure (lock or serialized list) was limited to 64 per address space. With z/OS V1R7, the limit is reduced to approximately 6 lock structure connections per address space. This reduction is due to the increase in the number of data spaces owned by connectors to coupling facility lock structures. The limit for connections to serialized list structures is not changed; however, if an address space uses connections to both lock and serialized list structures, the extra use of data spaces by the lock structure will reduce the availability of data spaces for the serialized list structure. Also, connections that support user-managed structure rebuild use double the number of data spaces, thus reducing the maximum connector limit even more.
Migration action: Review the use of serialized coupling facility structures per address space. If the number of connectors to lock structures from any one address space exceeds about six, reduce the number of connectors. If the connectors are
associated with applications written by you, restructure the applications to connect from multiple address spaces. If the connectors are associated with vendor applications, contact the owning vendors.

Accept the UCB overlay detection default change (Recommended, as of R7)

Not required, but is recommended to detect UCB overlays before they happen, which could have serious system impacts.
The unit control block (UCB) is a critical control structure in z/OS that designates the state and definition of an I/O device. Since MVS/ESA SP 5.2.0, applications could capture a UCB into a private virtual address. This implies that application programs could inadvertantly overlay this storage and cause the real UCB contents to be overlayed, resulting in system outages. In z/OS V1R5, the input/output supervisor (IOS) was enhanced to provide read-only views for all captured UCB requests. However, you had to use an IBM internal option in DIAG*xx* to enable IOS to provide the views. In z/OS V1R7, this overlay protection mechanism is enabled by default to help prevent UCB overlays.
Migration action: You need do nothing to take the recommended action, which is to accept the UCB overlay detection default change. By default, z/OS V1R7 systems come up with UCB overlay protection enabled. During operation of your system, if you encounter UCB overlays and dumps, report them to the software vendor or program owner. If you need to disable the system level trap (and thereby stop detecting UCB overlays), you can do so with either of the following methods:

- By command: SETIOS CAPTUCB,PROTECT=NO
- By IECIOS*xx* parmlib member: CAPTUCB,PROTECT=NO

If a program attempts to write into storage that represents a captured UCB, a protection check abend occurs. It is the responsibility of the abending program to recover from this appropriately. Programs that wish to store into a UCB must do so using the actual UCB address and not a captured view. Obtaining the actual UCB address can be done using services provided by IOS.
Note: The UCBLOOK macro and the UCBSCAN ADDRESS macro obtain an actual, not a captured, UCB address. Therefore, the addresses obtained by these macros are not affected by the UCB overlay detection default change.

Review the IXLLOCK hash algorithm (Required-IF, as of R7 with rollback via OA03194)

Required if you use coupling facility lock structures.

© Copyright IBM Corporation, 2007

Before z/OS V1R7, the number of XES connections to serialized structures (lock and serialized list structures) was limited to 64 structures per address space. This limit related to the number of connector-owned data spaces that can reside on an address space's PASN access list. With z/OS V1R7, the system obtains additional data spaces for each lock structure and spreads resource information across the data spaces based on the lock hash value specified on IXLLOCK requests.
Migration action: Review the hash algorithm used by your application to ensure that the algorithm produces hash values that are evenly distributed on the low-order digit of the result.

Update IPCS subcommands that require the ERROR criterion to select ASIDs (Required-IF, as of R7)

Required if you have IPCS procedures that require the ERROR criterion to select ASIDs.
Before z/OS V1R7, most IPCS subcommands that support ASID selection criteria defaulted to both CURRENT and ERROR ASIDs. As of z/OS V1R7, the default for address space selection has been changed to solely the CURRENT option. If you have specific IPCS procedures that benefit from using the ERROR option to select ASIDs, update those procedures to explicitly request the option.
Migration action: Update appropriate IPCS procedures to explicitly request the ERROR option to select ASIDs.

Update operator procedures for certain changed MVS commands (Required-IF, as of R7)

Required if you expect certain results from issuing any of the following MVS commands: START DEALLOC, VARY OFFLINE, and UNLOAD.
Before z/OS V1R7, certain MVS commands produced results that are different from what occurs in z/OS V1R7:

- The START DEALLOC command was used to redrive pending offline processing, often as an automated procedure. As of z/OS V1R7, the pending offline process is a timer-driven process and a redrive procedure is not required by the installation.
- Before z/OS V1R7, when message IEF238D ("*Reply device_name,WAIT OR CANCEL"*) was issued, VARY OFFLINE processing for a device could not complete until the operator responded to the message. The device entered a pending offline state (message IEF524I, "device PENDING OFFLINE") and remained so for 15 minutes, at which point message IEF525E ("device STILL PENDING OFFLINE") was issued and repeated at 15 minute intervals until either a response was provided or the job was canceled. As of z/OS V1R7, when message IEF238D is issued, VARY OFFLINE processing is not initiated until the message is replied to. Messages IEF524I and IEF525E are not issued until the VARY OFFLINE processing is started. Devices that were already in a pending offline state before message IEF238D was issued can be processed.
  Also, prior to the improvements in VARY OFFLINE processing in z/OS V1R7, when an operator responded to message IEF238D (Recovery Allocation) with a device number that was "pending offline", the device would remain "pending offline" after the device was unallocated (that is, when the job was finished with it). As of z/OS V1R7, it is likely that the device will go offline before the Recovery Allocation process can bring it online. Therefore, the device will no longer be marked "pending offline" after the allocation completes. You must reissue the VARY command to bring the device offline.
- Previously, when a device was "pending unload" because of being allocated when an UNLOAD command was issued, the device would remain "pending unload" after the device was unallocated (that is, when the job was finished with it). As of z/OS V1R7, because of changes in UNLOAD processing, if a volume is mounted on a device and the device is allocated, the system passes responsibility for unloading the device to the job that allocated it. The pending unload process will complete with message IEF415I.

Migration action: Update your operator procedures to allow for the changed behavior of these commands

© Copyright IBM Corporation, 2007

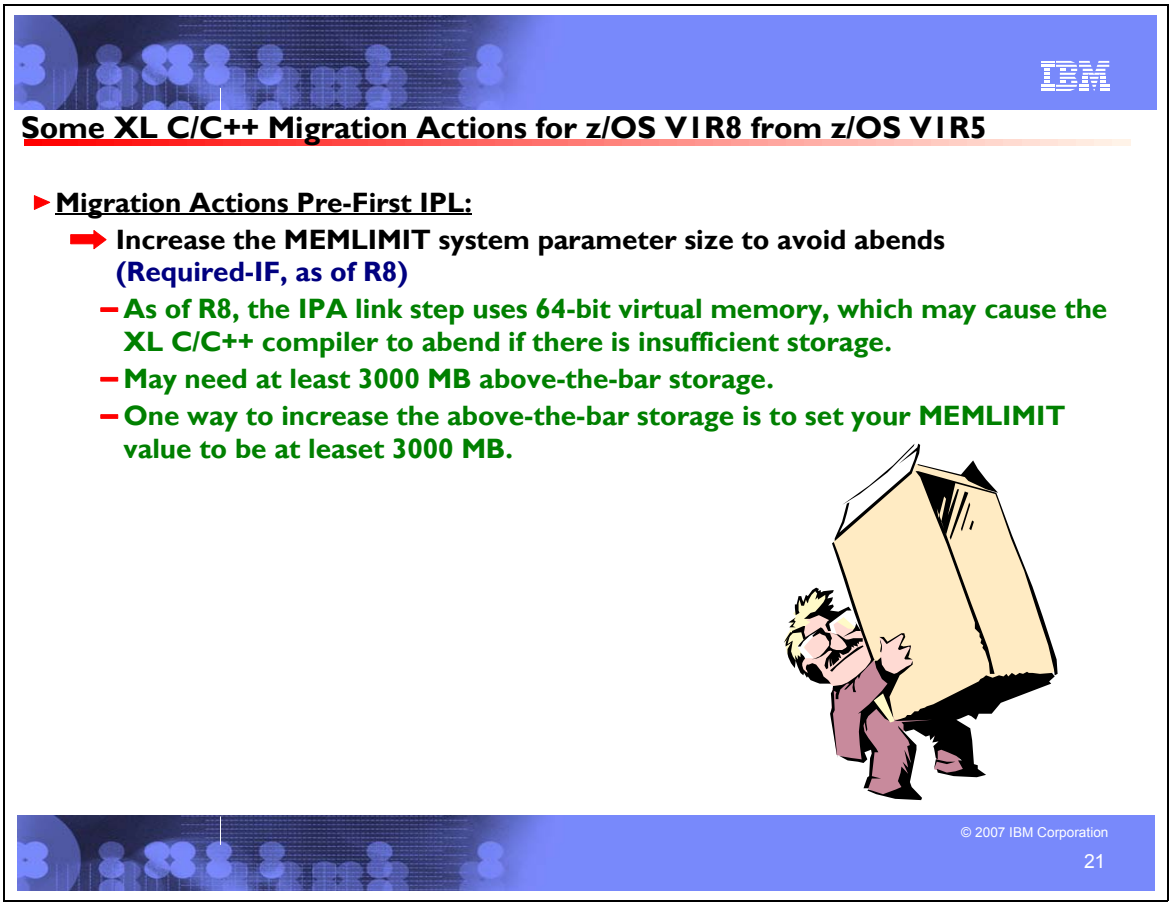

XL C/C++ Migration Actions Between z/OS V1R5 and z/OS V1R8

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

XL C/C++ Migration Actions You Can Do Now

Review the *XL C/C++ Migration Guide for the Application Programmer* (Recommended)

Not required, but recommended if you use XL C/C++.

The publication *z/OS XL C/C++ Compiler and Run-Time Migration Guide for the Application Programmer* is written for application programmers, whereas *z/OS Migration* is written for system programmers. However, in some customer locations, job scope could overlap such that system programmers might find information in the XL C/C++ publication that is relevant to their responsibilities. For example, migration information related to the c89 utility in the XL C/C++ publication could be of interest. Therefore, you ought to review this XL C/C++ publication if you use XL C/C++.

Migrate from use of the C/C++ ISPF panels (Required-IF, as of R6)

Required if you use the C/C++ ISPF panels.

In z/OS V1R6, IBM has removed the C/C++ ISPF panels from z/OS. z/OS V1R5 was the last release that includes the panels. These panels are shipped with the C/C++ without Debug Tool feature and include panels for C/C++ foreground compiles, C/C++ background compiles, and help panels for these compiles.

Migration action: Instead of using the ISPF panels, you can continue to invoke the C/C++ compiler via z/OS UNIX commands, using JCL, and under TSO/E.

Reference Information: *z/OS C/C++ User's Guide*.

Remove your run-time dependency on the C/C++ IBM Open Class Library (Recommended)

Not required, recommended because it is planned to become a requirement in the release following z/OS R8.

In z/OS V1R5, development support for the C/C++ IBM Open Class® Library, specifically, the Application Support Class and Collection Class libraries, was removed. Since z/OS V1R5, only run-time support has been provided for the C/C++ IBM Open Class Library.

z/OS V1R8 is planned to be the last release that includes the run-time support. If you have applications that depend on this support, they will not run after the run-time support is removed.

Migration action: Stop using the following dynamic link libraries (DLLs), which provide the run-time support:

- CBC.SCLBDLL, which contains member CLB3CLAS (aliases APPSUPP and ASCCOLL)
- CBC.SCLBDLL2, which contains members CLBECOLL (alias COLL) and CLBEIOC (alias IOC)

If you are using the CBC.SCLBSID library, which contains side-decks, ensure that you are using only the IOSTREAM, IOSX64 (64-bit IOSTREAM), and COMPLEX members, as they will be the only members that will remain when the run-time support has been removed.

For C/C++ run-time library support, use the C++ Standard Library, shipped with base element Language Environment, instead of the C/C++ IBM Open Class Library.

Remove your dependency on the OS/390 R10 level of the C/C++ compilers (Required-IF, as of R7)

Required if you are currently using the OS/390 V2R10 C/C++ compilers to compile your programs.

Starting with z/OS V1R2, the OS/390 V2R10 C/C++ compiler was shipped in addition to the strategic ISO XL C/C++ compiler that is provided with the z/OS optional feature C/C++ without Debug Tool. As of z/OS V1R7, the OS/390 V2R10 C/C++ compiler is removed from the C/C++ without Debug Tool feature, leaving only the ISO C/C++ compiler..

Migration action: The ISO C++ Standard (also known as ISO/C++98 or ANSI/C++98) introduces new features to the C++ language. Semantics are changed, keywords are added, and new facilities are added to the Standard Library. Highlights are:

- The **for**-loop scoping rule is changed.
- Implicit **int** is no longer supported.
- The semantics of **friend** class declaration is changed.
- The semantics of the **throw** operator is changed.
- Exception handling is added to the **new** and **delete** operators.

The following keywords are added to the compiler. If your existing code uses any of these keywords as identifiers, the recommended practice is to change your code. If that is not possible or not practical, use the NOKEYWORD(keyword) option to disable the individual keyword.

- bool
- explicit
- export
- false
- mutable
- namespace
- true
- typename
- using

Use the LANGLVL option, which controls the language level of the compiler, to help in your migration:

- LANGLVL(COMPAT92). This option group instructs the z/OS C/C++ compiler to follow the OS/390 V2R10 compiler's semantics whenever possible. Use this as a tactical transition if your code compiles with OS/390 V2R10 and you want to move quickly to the new compiler with minimal immediate changes.
- LANGLVL(STRICT98) or LANGLVL(ANSI). These two are identical; they instruct the z/OS C/C++ compiler to follow the ISO/C++98 standard.
- LANGLVL(EXTENDED). This is the extended language level; it includes all IBM extensions on top of ISO/C++98. The OS/390 V2R10 C/C++ compiler is shipped along with the new z/OS C/C++ compiler. If the source changes required by the new C++ Standard cannot be applied to your existing code all at once, you can also use the OS/390 V2R10 C/C++ compiler to help the transition. Setup considerations in the environment to select the respective compiler are described below:

Replace the following OS/390 V2R10 compiler options, which are not supported with the newer compilers:

- DECK. The replacement for the DECK functionality that routes output to DD:SYSPUNCH is OBJECT(DD:SYSPUNCH).
- GENPCH.
- HWOPTS. The replacement for HWOPTS is ARCHITECTURE.
- LANGLVL(COMPAT).
- OMVS. The replacement for OMVS is OE.
- SRCMSG.
- SYSLIB. The replacement for SYSLIB is SEARCH.

© Copyright IBM Corporation, 2007

- SYSPATH. The replacement for SYSPATH is SEARCH.
- USEPCH.
- USERLIB. The replacement for USERLIB is LSEARCH.
- USERPATH. The replacement for USERPATH is LSEARCH.

Note that GENPCH and USEPCH served as a tactical initiative to reduce compile times. IBM is concentrating instead on improving compile times by building the compilers with successively higher levels of optimization from release to release, exploiting the ongoing optimization improvements made in the compilers.

Note: If you are not yet able to migrate to the ISO XL C/C++ compiler, you can copy the OS/390 V2R10 C/C++ compiler (which is not shipped with z/OS V1R7) from your old z/OS system for use on z/OS V1R7, with limitations. See informational APAR PK05324 for details.

z/OS UNIX System Services
c89, cc, and cxx default to z/OS C/C++. You can override the compiler selection with the use of environment variables, shown below:

Environment variables	Value as of z/OS V1R2 C/C++ (default)	Value for OS/390 V2R10 C/C++
C89_CVERSION	"0x41020000"	"0x220A0000"
CC_CVERSION	"0x41020000"	"0x220A0000"
CXX_CVERSION	"0x41020000"	"0x220A0000"
C89_CNAME	"CCNDRVR"	"CBCDRVR"
CC_CNAME	"CCNDRVR"	"CBCDRVR"
CXX_CNAME	"CCNDRVR"	"CBCDRVR"
STEPLIB	include SCCNCMP	include SCBCCMP

JCL jobs and cataloged procedures
z/OS C/C++ cataloged procedures are in the hlq.SCCNPRC library. There are new cataloged procedures for z/OS C/C++, aliased to old procedure names where applicable. OS/390 V2R10 C/C++ cataloged procedures are in the hlq.SCBCPRC library. You can switch compilers by switching proclibs.

TSO/E REXX execs and logon procedures
z/OS C/C++ TSO/E REXX execs are in the hlq.SCCNUTL library. OS/390 V2R10 C/C++ TSO/E REXX execs are in the hlq.SCBCUTL library. You can switch compilers by switching SYSPROC and STEPLIB allocations in the logon procedure.

Reference information: For more information about the compiler changes that were made between OS/390 V2R10 and z/OS V1R2, and about migrating from the older to the newer level, see *z/OS C/C++ Compiler and Run-Time Migration Guide*. The new language standard is described in publications from the International Standards Organization (ISO) and the American National Standards Institute (ANSI). They make them available for purchase on the Internet at http://www.iso.ch and at http://www.ansi.org. For details about the LANGLVL option, see *z/OS C/C++ User's Guide*.

XL C/C++ Migration Actions Pre-First IPL

Increase the MEMLIMIT system parameter size to avoid abends (Required-IF, as of R8)

Required if you will do IPA links and you do not have at least 3000 MB of virtual storage available above the 2GB bar.
As of z/OS V1R8, the IPA Link step makes use of 64-bit virtual memory, which could cause the XL C/C++ compiler to abend if there is insufficient storage. Increasing the default MEMLIMIT size in parmlib member SMFPRM*xx* can overcome the problem. The default value from SMFPRM*xx* takes effect if a job does not specify MEMLIMIT on the JOB or EXEC statement and does not specify REGION=0M. Note that the MEMLIMIT specified in an IEFUSI exit routine overrides all other MEMLIMIT settings.
Migration action: If a compile that uses IPA Link abends because of insufficient storage, increase the amount of above-the-bar virtual storage available for the IPA Link to at least 3000 MB. One way to accomplish this is to change your MEMLIMIT value to at least 3000M.

Accommodate the change to the euro as the default currency (Required-IF, as of R6)

Required if your applications use the default currency and need to use the local currency (that is, not the euro currency symbol).
Migration action: Prior to z/OS V1R6, the default currency for European Union countries was set to the local currency in the LC_MONETARY category of the locale. If you wanted to set the euro as the currency, you would set the @euro locales using setlocale(). As of z/OS V1R6, the LC_MONETARY information in the base locale is set to use the euro. If

© Copyright IBM Corporation, 2007

you set the base locale, the euro is now the default currency. If you want your applications to continue using the old (local) currency, you must now issue setlocale() with the new @preeuro locale as the parameter. Behavior of the @euro locales has not changed.

To set the euro as the currency, do nothing. But if you want to restore the prior situation, which is to use the local currency as the default, use the setlocale() function to do so.

XL C/C++ Migration Actions Post-First IPL

<none>

© Copyright IBM Corporation, 2007

Some Communications Server Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions You Can Do NOW:**
  - **IP Services: Remove ASSORTEDPARMS and KEEPALIVEOPTIONS statements from the TCP/IP profile (Required-IF, as of R8)**
    - Migrate to new profile statements before z/OS R8
  - **IP Services: Migrate from OROUTED to OMPROUTE (Required-IF, as of R7)**
    - Use OROUTED -c to make migration to OMPROUTE easier BEFORE migrating to R7!
    - As of R7, if you want to use a daemon to support the RIP protocols, you must use OMPROUTE.
  - **IP Services: Update profile is still using TN3270E server legacy statements (Required-IF, as of R8)**
    - QUEUESESSION, and three options on BEGINVTAM block: LUSESSIONPEND, MSG07, and TELNETDEVICE) are replaced by improved functions.
    - Equivalent functions are available.
  - **IP Services: Ensure that UDP port 514 is available to syslogd if not started with the -i option (Required-IF, as of R8)**
    - If syslogd is configured to process log messages from a remote system daemon, and fails to bind to UDP port 514, syslogd terminates.
    - If you start syslogd with the -i option, no changes to TCP/IP port reservations for syslogd.
    - Otherwise, ensure that UDP port 514 is reserved for OMVS or (preferably) the job name under which syslogd will run. Verify syslogd is bound to UDP port 514 with with netstat -o/allc
  - **IP Services: Use the REXX FTP client API from customized TSO/E REXX modules (Required-IF, as of R8)**
    - As or R8, TSO/E REXX parm modules provide the REXX FTP client API.
    - If you customized the TSO/E REXX parm modules in a prev rel and now want to use the new REXX FTP client, then migrate your customizations to the R8 level.

© 2007 IBM Corporation

IBM

Some Communications Server Migration Actions for z/OS V1R8 from z/OS V1R5

- **Migration Actions Pre-First IPL:**
  - **IP Services: Update /etc configuration files (Required-IF, as of R7 and R5)**

- Update from /usr/lpp/tcpip/samples , to:
  - for Policy Agent and Intrusion Detection Services in R8, put in a user directory.
    - pagent_IPSec.conf, pagent_CommonIPSec.conf, and pagent_IDS.conf
  - /etc/protocol : hopopt, encapsulating security payload, and authentication header protocol was added in R7
  - /etc/services : R7 APAR OA12364 provides ability for JES to communicate NJE work over TCP/IP.
  - /etc/osnmpd.data : every release sysName MIB object is updated

IP Services: Run the TN3270E Telnet server as a separate address address space (Recommended, as of R6)

- Will be a requirement in the release following z/OS R8.
- Provides visibility and control over TN3270 function separate from the TCP/IP stack

1. set up a superuser with OMVS segment
2. associate user with telnet proc, using RACF STARTED class
3. change the priority of Telnet by assigning Telnet to a service class other than the default SYSSTC class in the STC subsystem.
4. remove the TN3270E Telnet server profile statements from the TCP/IP profile.
5. create Telnet proc, specify profile.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Communications Server Migration Actions Between z/OS V1R5 and z/OS V1R8

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

Communications Server Migration Actions You Can Do Now

IP Services: Migrate from Firewall Technologies to Integrated IPSec/VPN Support (Required, as of R8)

Required if Firewall Technologies is currently being used.

Before z/OS V1R8, Firewall Technologies could be used to provide IP filtering and IPSec support. This support is discontinued in z/OS V1R8. As a result, you must remove the FIREWALL and accompanying DVIPSEC parameters from the IPCONFIG statement in the TCP/IP profile after you have completed your **successful** migration from Firewall Technologies to Integrated IPSec/VPN.

Furthermore, starting in z/OS V1R8, you must migrate to the Integrated IPSec/VPN support if you want to continue using IP filtering and IPSec support. Integrated IPSec/VPN support provides IPv4 and IPv6 support for IP filtering, IP Security/Virtual Private Network (IPSec/VPN), and Internet Key Exchange (IKE). This support provides easier configuration, greater scalability, improved performance, and enhanced serviceability over the Firewall Technologies version available in z/OS V1R7.

Migration action:

1. Code IPSECURITY on the IPCONFIG statement in the TCP/IP profile.
2. If SWSA processing is desired, code DVIPSEC on the IPSEC statement in the TCP/IP profile.
3. Use the IBM Configuration Assistant for z/OS Communications Server tool to configure IP Security policy. You can download the tool from **http://www.ibm.com/software/network/commserver/zos/.**
   Note: You can configure IP Security policy manually without the use of the tool.
4. Remove the FIREWALL and DVIPSEC parameters from the IPCONFIG statement in the TCP/IP profile after you have completed your successful migration from Firewall Technologies to Integrated IPSec/VPN.

Migrate from using the SNMP SLA subagent (pagtsnmp) to the SNMP Network SLAPM2 subagent (nslapm2) (Required, as of R8)

Required if you use the SNMP SLA subagent (pagtsnmp).

The SLA subagent supports the Service Level Agreement Performance Monitor (SLAPM) MIB. Before z/OS V1R8, the SNMP SLA subagent (pagtsnmp) allowed network administrators to retrieve data and to determine if the current set of SLA policy definitions was performing as needed or if adjustments needed to be made.

Starting in z/OS V1R8, the support for pagtsnmp is discontinued. As a result, you need to use the SNMP Network SLAPM2 subagent (nslapm2) instead. It has the following extensive improvements:

- Better and more relevant information on network and TCP-related performance data (for example, loss, delays, and 64-bit count values).
- System overhead reduction due to restructure of the MIB and the interface for the subagent to get performance data.
- Transparency in supporting either IPv4 or IPv6 because the MIB no longer keeps track of IP addresses.

Migration action: Move management applications from the SNMP SLA subagent (pagtsnmp) to the SNMP Network SLAPM2 subagent.

IP Services: Migrate from LDAP protocol version 2 to LDAP protocol version 3 for QoS and IDS policies in Policy Agent (Recommended)

Not required, but recommended because it will be a requirement in the release following R8, if you use LDAP protocol version 2 for QoS or IDS policies.

z/OS V1R8 Communications Server is planned to be the last release to support the LDAP protocol version 2 for QoS and IDS policies in Policy Agent. After z/OS V1R8, you will have to use the LDAP protocol version 3 for QoS and IDS policies in Policy Agent. The LDAP protocol version 3 has improvements in internationalization, authentication, referral, and deployment.

Migration action:

1. Install the LDAP protocol version 3 schema definition files on the LDAP server instead of the LDAP protocol version 2 schema definition files.
2. Change the protocol version (LDAP_ProtocolVersion parameter) on the Policy Agent ReadFromDirectory configuration statement to use protocol version 3.

IP Services: Remove ASSORTEDPARMS and KEEPALIVEOPTIONS statements from TCP/IP profile (Required-IF, as of R8)

© Copyright IBM Corporation, 2007

Required if you use the ASSORTEDPARMS and KEEPALIVEOPTIONS statements.
Support for the ASSORTEDPARMS and KEEPALIVEOPTIONS statements in the TCP/IP profile has been removed. These statements have been deprecated since z/OS V1R2. The parameters on these statements have equivalent parameters on the GLOBALCONFIG, IPCONFIG, IPCONFIG6, TCPCONFIG, and UDPCONFIG statements. Current users of these statements must migrate their profiles to use the equivalent statements or risk losing intended function.
Migration action:

- Remove the ASSORTEDPARMS statement from the TCPIP profile. For each parameter specified on the ASSORTEDPARMS statement, specify the equivalent parameter on another statement. For replacement function, specify:
  - CLAWUSEDOUBLENOP on the IPCONFIG statement.
  - IGNOREREDIRECT on the IPCONFIG and/or IPCONFIG6 statement.
  - NOFWD on the IPCONFIG and/or IPCONFIG6 statement as NODATAGRAMFWD.
  - NOUDPQUEUELIMIT on the UDPCONFIG statement.
  - RESTRICTLOWPORTS on the TCPCONFIG and/or UDPCONFIG statement.
  - SOURCEVIPA on the IPCONFIG and/or IPCONFIG6 statement.
  - STOPONCLAWERROR on the IPCONFIG statement.
  - TCPIPSTATISTICS on the GLOBALCONFIG statement.
- Remove the KEEPALIVEOPTIONS statement from the TCPIP profile. For each parameter specified on the KEEPALIVEOPTIONS statement, specify the equivalent parameter on the TCPCONFIG statement. For replacement function, specify:
  - INTERVAL on the TCPCONFIG statement.
  - SENDGARBAGE on the TCPCONFIG statement.

IP Services: Change syslogd configuration files to use the kernel facility instead of the user facility for messages from remote syslog daemons (Required-IF, as of R8)
Required if you receive syslog log messages from the network that have an origin facility of kernel and want to continue to process these messages.
If you have configured syslogd to process log messages from a remote syslog daemon (received over the network) using the user facility, you might need to change your syslogd configuration file. Before z/OS V1R8, log messages received from the network with a facility of kernel were changed by syslogd to have a facility of user. With z/OS V1R8, log messages received from the network with a facility of kernel will not have this facility changed. This change does not affect log messages received from the local system, which will continue to have the kernel facility converted to the user facility.
In this example, the following rule is coded in /etc/syslog.conf: user.info /tmp/syslog.info
Messages are received from the network that are sent with a facility of kernel. Before z/OS V1R8, the system converted kernel to user, so the rule worked as expected. With z/OS V1R8, you must change the rule to include the kernel facility explicitly: user.info;kernel.info /tmp/syslog.info.
Migration action:

1. Determine whether you use syslogd for processing messages received from remote syslog daemons (over the network). If you are currently starting syslogd with the -i option, you do **not** need to make any changes to your syslogd configuration file.
2. Examine your syslogd configuration file. If you have any rules that use the user facility, you might need to change them or add new rules.
3. If you want to maintain the same function as before, the simplest change is to add to the existing rules that use the user facility. For example, if you have a rule that looks similar to this: user.info /tmp/syslog.info
   Consider changing it to this: user.info;kernel.info /tmp/syslog.info
   This rule causes messages received with the facility of user and priority of info to be logged to /tmp/syslog.info. Also, any message received with a facility of kernel and a priority of info is logged to the same file.

IP Services: Ensure TN3270E SMF type 119 termination record parser can accept a record with two additional new sections (Required-IF, as of R8)
Required if you monitoring performance and want the information reported on the TN3270E SMF type 119 termination record.
Performance monitoring data is now optionally reported in two new sections of the TN3270E SMF type 119 termination record. If you already have monitoring set up, you should make sure it can handle the two additional SMF sections.

© Copyright IBM Corporation, 2007

Migration action: Ensure that your TN3270E SMF type 119 termination record parser can accept a record with up to two additional sections. If your parser was written to read sections based on the SMF section header, you should not have to make any changes. For additional information about these Telnet changes, see TN3270 SMF record enhancement in *z/OS Communications Server: New Function Summary*.

IP Services: Ensure that Policy Agent LDAP IDS attack policies have the intended result (Required-IF, as of R8)

Required if you are using z/OS Policy Agent IDS policy for restricted IP protocol (ibm-idsAttackType RESTRICTED_IP_PROTOCOL) or raw restrictions (ibm-idsAttackType OUTBOUND_RAW) and if your policy does not specify the ibm-idsProtocolRange attribute in the attack rule.

The ibm-idsProtocolRange attribute specifies a list of protocols for intrusion detection services (IDS) rules defined in LDAP. Before z/OS V1R8, the accepted values were 1 thru 255. The default was 0 and indicated none. Starting in z/OS V1R8, the accepted values are 0 thru 255. The default is still 0; however, 0 now indicates protocol 0 instead of none.

Migration action: Consider removing the attack type from the policy. Otherwise, protocol 0 will be restricted.

IP Services: Rename duplicate policy objects defined in the configuration file and LDAP server of the same type (Required-IF, as of R8)

Required if you are using z/OS Policy Agent IDS and you have QoS or IDS policies configured with duplicate names in the configuration file and LDAP.

Policy objects defined in the configuration file and LDAP of the same type (such as QoS or IDS) with duplicate names are discarded by the Policy Agent. Starting in z/OS V1R8, these duplicate policies log an error instead of a warning and issue the following console message: EZZ8438I PAGENT POLICY DEFINITIONS CONTAIN ERRORS FOR *tcpImage :type*

Migration action: Look for duplicate policy objects and rename them to avoid the error.

IP Services: Ensure that some UDP ports are not reserved (Required-IF, as of R8)

Required if all UDP ports are reserved and you have an application that needs to bind to an ephemeral UDP port.

Before z/OS V1R8, a UDP port that was reserved with the PORT or PORTRANGE statement in the TCPIP PROFILE could be assigned to an application as an ephemeral port if the application was authorized to use the port. Starting in z/OS V1R8, ports that were reserved by the UDP PORT statement will no longer be given out as ephemeral ports even if the job name of the application that is doing the bind matches the job name on the PORT configuration statement.

Migration action: Change the PORT configuration statement so that some UDP ports are not reserved.

IP Services: Update profiles if still using TN3270E server legacy statements (Required-IF, as of R8)

Required if you use QUEUESESSION, LUSESSIONONPEN, MSG07, or TELNETDEVICE in the BEGINVTAM block.

In z/OS V1R8, the QUEUESESSION statement and three options of the BEGINVTAM block (LUSESSIONPEND, MSG07, and TELNETDEVICE) have been replaced by improved functions. These removals were announced several releases ago in the TN3270E server configuration publications and messages.

Equivalent functions are available using other options or by moving the existing statement to its more appropriate Telnet configuration block. Specifically:

- Instead of the QUEUESESSION statement, use the QSESSion parameter on the ALLOWAPPL or RESTRICTAPPL statement.
- Instead of the LUSESSIONPEND option of the BEGINVTAM block, use the LUSESSIONPEND statement in the TelnetGlobals, TelnetParms, or ParmsGroup block.
- Instead of the MSG07 option of the BEGINVTAM block, use the MSG07 statement in the TelnetGlobals, TelnetParms, or ParmsGroup block.
- Instead of the TELNETDEVICE option of the BEGINVTAM block, use the TELNETDEVICE statement in the TelnetGlobals, TelnetParms, or ParmsGroup block.

Migration action:

- For QUEUESESSION migration:
  1. Remove the QUEUESESSION statement from the BEGINVTAM block. This statement affects all DefaultAppl statements.
  2. Make a list of all DEFAULTAPPL applications. Note if there is already an ALLOWAPPL or RESTRICTAPPL statement for this application. If neither statement exists for the application, create an ALLOWAPPL statement for the application and add the QSESSion parameter.
  3. If an ALLOWAPPL or RESTRICTAPPL statement exists, add the QSESSion parameter to that statement.
- For LUSESSIONPEND, MSG07, and TELNETDEVICE migration:
  1. Remove the statements from the BEGINVTAM block.

© Copyright IBM Corporation, 2007

2. Add the statements to either TelnetGlobals or TelnetParms, depending on the scope necessary.

IP Services: Migrate from BIND DNS 4.9.3 (Recommended)

Recommended if you use BIND DNS 4.9.3 function because it is planned to become a requirement in the future.

In a future release of z/OS, support for BIND DNS 4.9.3 is planned to be discontinued. If you are using this name server, you are encouraged to find a replacement as soon as you can.

Migration action: You should implement BIND DNS 9.2.0 as a replacement. BIND DNS 9.2.0 is included in z/OS beginning with V1R4. If you exploit the Connection Optimization (DNS/WLM) feature of BIND 4.9.3, you should investigate alternative solutions, such as sysplex distributor and the z/OS Load Balancing Advisor for load balancing, and the Automated Domain Name Registration application (ADNR) for its ability to automatically update DNS with the availability status of sysplex resources. Note that implementing ADNR would have to be done after first IPL because ADNR is new in z/OS V1R8.

IP Services: Remove customization of SNMP sysObjectID MIB object in OSNMPD.DATA File (Recommended)

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

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

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

IP Services: Migrate from OROUTED to OMPROUTE (Required-IF, as of R7)

Required if you use OROUTED as your routing daemon..

Before z/OS V1R7, the OROUTED routing daemon could be used to support the RIP1 and RIP2 protocols. However, the support for OROUTED is discontinued in z/OS V1R7. As a result, you must remove OROUTED configuration statements if you have them.

In addition, starting in z/OS V1R7, if you want to use a daemon to support the RIP protocols, you must use OMPROUTE. OMPROUTE is another routing daemon and it supports RIP1 and RIP2, as well as IPv6 RIP and Open Shortest Path First (OSPF) protocols. The *z/OS Migration* book explains how to migrate to using OMPROUTE to implement the RIP protocol as you are doing with OROUTED.

An OROUTED start parameter is available to assist with migration from OROUTED to OMPROUTE. This function is invoked by specifying -c on the OROUTED startup parameters or by using the following modify command: f orouted,parms=-c

Note: The -c parameter must be used before installing z/OS V1R7 or higher. (It must be used on your current z/OS system, not on V1R7 or higher.) Both routing daemons use syslogd, resolver configuration files (for DATASETPREFIX and TCPIPJOBNAME), the services file (route 520/udp router routed), and have optional cataloged procedures.

Migration action: Refer to the *z/OS Migration* book for the detailed procedures for migrating from OROUTED to OMPROUTE.

IP Services: Verify the correct VLANID value on LINK and INTERFACE profile statements (Required-IF, as of R6)

Required if you have specified a VLANID value of 4095.

Prior to z/OS V1R6, TCP/IP allowed a VLANID value on the LINK statement for IPAQENET, or INTERFACE statement for IPAQENET6, to be from 1–4095. As of z/OS V1R6, a value of 4095 is no longer supported. Therefore, if you have specified a VLANID value of 4095 on a LINK IPAQENET profile statement or an INTERFACE IPAQENET6 profile statement, you must change the value to a number from 1–4094.

SNA Services: Update automated VTAM internal trace analysis tools to recognize format changes in GBLK, FBLK, and COPY entries (Required-IF, as of R8)

Required if you use automated tools to analyze the VTAM internal trace and these tools depend on the format of the GBLK, FBLK, and COPY entries.

In z/OS V1R8, the VTAM internal trace entries GBLK, FBLK, and COPY have been changed. Certain fields have been changed or moved within the entries. Also, new GBL3 and COP2 continuation records have been added for GBLK and COPY, respectively.

Migration action: For details about the changes to the entries, see *z/OS Communications Server: SNA Diagnosis Vol 2, FFST Dumps and the VIT*.

© Copyright IBM Corporation, 2007

1. Be aware of the changes in format for the following VIT entries: GBLK, FBLK, and COPY.
2. Update any automated tools you use to analyze the VTAM internal trace that are dependent on the format of these entries.

SNA Services: Change your method of defining parallel EE TGs (Required-IF, as of R8)

Required if you define parallel EE TGs by specifying multiple SAP addresses.

As of z/OS V1R8, the definition of parallel Enterprise Extender (EE) Transmission Groups (TGs) by specifying multiple SAP addresses is no longer supported.

Migration action: If you need parallel EE TGs, define them by using different EE VIPAs on one (or both) of the endpoints.

SNA Services: Migrate from AnyNet (Required-IF, as of R8)

Required if you use AnyNet configured with your current definitions.

In z/OS V1R8, support for AnyNet® has been discontinued. All code supporting the AnyNet function has been removed from the z/OS product.

Migration action:

1. Remove any TCP major nodes from VTAMLST libraries.
2. Remove or disable the activation of the TCP major node from configuration files or any CLISTs.
3. Remove any port statements from your TCPIP profile.
4. Remove any FILESYSTYPE definition statements used to define AnyNet Sockets-over-SNA to z/OS UNIX.
5. For AnyNet SNA over IP, consider implementing alternate SNA/IP integration technologies such as Enterprise Extender (EE). For AnyNet Sockets over SNA, there is no replacement function.

Communications Server Migration Actions Pre-First IPL

IP Services: Verify your use of loopback addresses with multiple instances of TCP/IP in a CINET environment (Required-IF, as of R8)

Required if you use loopback addresses with multiple instances of TCP/IP in a CINET environment.

Before z/OS V1R8, loopback addresses were not reported by the TCP/IP stack to Common INET (CINET). As a result, when multiple instances of TCP/IP existed in a CINET environment, binds to a loopback address from an application that had not established affinity to a specific TCP/IP stack were always passed to the default stack. To ensure that all binds to loopback addresses can be successfully processed by the TCP/IP stack to which they are passed, loopback addresses are now reported to CINET, and CINET may now select a stack other than the default stack.

Migration action: If you use applications that do not establish affinity to a specific stack and that bind to a loopback address, and you require that the applications be bound to the default stack, make at least one of the following modifications:

- Have the applications establish affinity to the default stack.
- Have the applications bind to a nonloopback address on the default stack.
  Tip: Consider using a virtual IP address (VIPA).
- Create a unique loopback address (not 127.0.0.1) on the default stack and bind the applications to it.

IP Services: Ensure that UDP port b14 is available to syslogd if not started with the -i option (Req-IF, as of R8)

Required if you have configured syslogd to process log messages from a remote syslog daemon (received over the network).

If you have configured syslogd to process log messages from a remote syslog daemon (received over the network), you need to verify that the TCP/IP PORT reservation statement does not exclude syslogd from binding to port 514. Before z/OS V1R8, if the syslog daemon was not started in local-only mode (-i option) and failed to bind to UDP port 514 because that port was reserved by an application other than syslogd or was already in use, syslogd would start but would not process messages from the network. With z/OS V1R8, if syslogd is configured to allow messages from the network and fails to bind to UDP port 514, syslogd terminates.

Migration action:

1. Determine whether you use syslogd for processing messages received from remote syslog daemons (over the network). If you are currently starting syslogd with the -i option, you do **not** need to make any changes to your TCP/IP PORT reservations for syslogd.
2. Examine your TCP/IP PORT reservation table in your TCPIP PROFILE. Ensure that UDP port 514 is reserved for OMVS or (preferably) for the job name under which syslogd will run. The job name under which syslogd will run can be set with the _BPX_JOBNAME environment variable.

© Copyright IBM Corporation, 2007

3. Start syslogd and ensure that it starts properly and binds to the correct UDP port. Use netstat -o/allc to verify that syslogd is bound to UDP port 514.

With z/OS V1R8 you may now run two instances of syslogd. However, only one instance may process the UDP network port and a separate instance may only process the local UDP UNIX socket.

IP Services: Update diagnostic procedures for changed failure notification from certain TCP/IP-authorized programs (Required-IF, as of R8)

Required if you have any diagnostic procedures that are affected by this change.

Before z/OS V1R8, z/OS did not validate whether TCP/IP-authorized programs were installed in appropriate libraries and had appropriate attributes. With z/OS V1R8, the validation is performed. Now, when a program fails validation, a system abend of S4C5-77A53217 (or U4087 for some load modules) occurs. This change will ensure that selected TCP/IP programs are executing in the correct installed environment. If you were using program-specific failures to identify improper executing environments, you should update your diagnostic procedures to take into consideration the changed failure notification.

The following TCP/IP programs are affected:

ADNR, DCAS, EZAZSSI, EZBREINI, EZBTNINI, FTPD, FTPDNS, IOBSNMP, LBADV, LBAGENT, LPQ, LPR, LPRM, MODDVIPA, MVPMAIN, MVPXDISP, NAMEDXFR, NAMED4, NSUPDATE/NSUPDAT4, OMPROUTE, OPING, OSNMPD, OTRACERT, PING, POPPER, RSHD, SMTP, ,SNTPD, TRACERTE, TRAPFWD, TRMD.

Migration action: Update your diagnostic procedures to reflect the changed behavior. For the TSO commands in the list above, ensure that the command names are specified in the AUTHCMD NAMES section of your SYS1.PARMLIB IKJTSO*xx* member.

IP Services: Update automation that handles message EZZ4345I (Required-IF, as of R8)

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

In the text of message EZZ4345I, "NED" has been changed to "ND". Before z/OS V1R8, the text for message EZZ4345I was INTERFACE *takeover_interface_name* HAS TAKEN OVER NED RESPONSIBILITY FOR INACTIVE INTERFACE *inactive_interface_name*.

Starting with z/OS V1R8, the text is INTERFACE *takeover_interface_name* HAS TAKEN OVER ND RESPONSIBILITY FOR INACTIVE INTERFACE *inactive_interface_name*. If you are automating on the message text (not the message number), your automation might be affected.

IP Services: Use the REXX FTP client API from customized TSO/E REXX modules (Required-IF, as of R8)

Required if you customized the TSO/E REXX parameter modules in a previous release and now want to use the new new REXX FTP client.

Before z/OS V1R8, the default TSO/E REXX parameter modules did not include the REXX FTP client API. Starting in z/OS V1R8, the modules provide the REXX FTP client API. If you customized any of the following modules, you must also customize the z/OS V1R8 versions: IRXTSPRM IRXISPRM IRXPARMS

Migration action: Migrate your customizations to the z/OS V1R8 level.

IP Services: Make changes for Netstat enhancements (Required-IF)

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

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

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

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

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

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

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

© Copyright IBM Corporation, 2007

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

Migration action: If you added installation-dependent customization to any of the IBM-provided configuration files listed in Table 16, make the same changes in the new versions of the files by copying the IBM-provided samples to the files shown in the table and then customizing the files.

Note: Some /etc files were removed because of discontinuance of support for OROUTED.

Utility	IBM-provided sample file	Target location	What changed and when
Policy Agent	• /usr/lpp/tcpip/samples/pagent_IPSec.conf • /usr/lpp/tcpip/samples/pagent_CommonIPSec.conf	Update as needed by your installation and put in user directory.	In z/OS V1R8, IPv6 support for IPSec was added. Sample files were updated to include example definitions for IPv6 IPSec. In addition, in z/OS V1R8, support for IPSec was added to traverse an NAPT. Sample files were updated to include an additional traffic descriptor for IKE traffic to be used when traversing an NAPT. Note that this Policy Agent entry only applies to migrating from z/OS V1R7; IPSec was introduced in that release.
Intrusion Detection Services	/usr/lpp/tcpip/samples/pagent_IDS.conf	Update as needed by your installation and put in user directory.	In z/OS V1R8, support for IDS policies in configuration files was added. A new sample was added to provide examples of configuration file IDS policies.
NJETCP	/usr/lpp/tcpip/sample/services	/etc/services	z/OS V1R7 APAR OA12364 provides the ability for JES to communicate NJE work over TCP/IP.
getprotocolbyname() API	/usr/lpp/tcpip/samples/protocol	/etc/protocol	In z/OS V1R7, support for hopopt, encapsulating security payload, and authentication header protocol was added.
SNMP agent	/usr/lpp/tcpip/samples/osnmpd.data	/etc/osnmpd.data	Every release, the value of the sysName MIB object is updated to the current release.

IP services: Change FTP general resource profiles in the CSFSERV and CSFKEYS classes (Required-IF, as of R7)

Required if the conditions mentioned are true.

z/OS V1R7, the z/OS FTP server supports delegated general resource profiles defined in the CSFSERV and CSFKEYS classes to protect ICSF cryptographic hardware keys and resources. Because of this new support, you may need to make some changes in order for your system to continue working as it did prior to z/OS V1R7.

Specifically, you will need to make changes if **all** of the following conditions are true:

- You are using the z/OS FTP server
- Your server is configured to allow or require FTP users to log in using TLS security
- You are using an SAF compliant security product, such as RACF
- You have applied APAR PQ80574 so that you do not need to allow FTP client user IDs access to general resource profiles in the CSFSERV and CSFKEYS classes that protect ICSF cryptographic keys and services **OR IF** You are migrating from z/OS V1R6, and you use the general resource profiles in the CSFSERV and CSFKEYS classes that protect ICSF cryptographic keys and services.

If the above conditions are true, follow the steps outlined in *z/OS Migration.*

© Copyright IBM Corporation, 2007

IP Services: Use the SO_CLUSTERCONNTYPE option of the GETSOCKOPT API (Recommended, as of R7)

Not required, but recommended if you are using the SO_CLUSTERCONNTYPE option on the GETSOCKOPT socket API.

The SO_CLUSTERCONNTYPE option of the GETSOCKOPT API allows applications using TCP communications on z/OS to determine the location of the partner application (for example, at the other end of the TCP connection). The information returned includes indicators on whether the partner application is running in the same image or in the same sysplex, allowing the applications to optimize their communications based on this knowledge. In the case where the partner application is located in the same sysplex, an additional indicator is returned that describes the types of communication links/interfaces used for the TCP connection. If the communications links/interfaces are not exposed outside the sysplex, then the communication path for the connection is described as internal and applications can use this information to further optimize their processing (such as to avoid encryption of data).

Starting with z/OS V1R7, this internal communications indicator will no longer be set if the destination IP address for a connection (such as the partner's IP address) is a Dynamic VIPA or Distributed Dynamic VIPA residing in the sysplex. The reason is that traffic destined to these IP addresses can now be forwarded to the target TCP/IP stacks over links or interfaces that are external to the sysplex. Applications exploiting this internal indicator should continue to function properly from a communications perspective, but they may no longer optimize their processing when the destination address being used is a Dynamic VIPA or a Distributed Dynamic VIPA.

Migration action: If you have applications that exploit this feature with Dynamic VIPAs or Distributed Dynamic VIPAs, you should consider modifying their configuration to use real IP addresses or static VIPAs as the destination addresses if optimized processing is an issue.

IP Services: Use IPV6_PKTINFO in a multilevel secure environment (Required-IF, as of R7)

Required if running in a multilevel secure environment and any applications enable the IPV6_PKTINFO socket option.

In z/OS V1R7, the IPV6_PKTINFO socket option authorization checks in a multilevel secure environment have been changed to require READ access to the EZB.SOCKOPT.sysname.tcpname.IPV6_PKTINFO SERVAUTH resource. Prior to V1R7, the IPV6_PKTINFO socket option did not require READ access to this resource.

Migration action: Define the EZB.SOCKOPT.sysname.tcpname.IPV6_PKTINFO SERVAUTH resource using the RACF definition sample that is shipped in EZARACF. Permit the appropriate users or programs access to the EZB.SOCKOPT.sysname.tcpname.IPV6_PKTINFO resource using the RACF permit sample that is shipped in EZARACF. EZARACF is shipped in data set TCPIP.SEZAINST.

IP Services: Modify the FTPOSTPR exit routine (Required-IF, as of R7)

Required if you are migrating the FTPOSTPR exit routine.

In z/OS V1R7, the FTP user exit routine FTPOSTPR has an additional parameter, increasing the total number of parameters to 18. The new parameter is a 1-byte confidence level descriptor. If you do not activate FTP confidence checking in the FTP server by coding CHKCONFIDENCE TRUE in FTP.DATA, the confidence level is always set to X'04' — Confidence level checking is not active.

Migration action:

1. Verify that your exit routine will tolerate the additional parameter.
2. Modify FTPOSTPR if needed.

IP Services: Update programs for the TCP segmentation offload performance enhancement (Req-IF, as of R7)

Required if you have programs that analyze formatted packet trace output or if you have programs that format packet trace data using the TCP/IP real-time packet trace interface.

Before z/OS V1R7, the TCP/IP stack performed all outbound TCP segmentation processing. Starting in z/OS V1R7, if you are using an OSA-Express2 Ethernet feature in QDIO mode that supports TCP segmentation offload, the stack offloads most IPv4 outbound TCP segmentation processing to the OSA.

If you have programs that analyze formatted packet trace output, your programs might need changes. If you have programs that format packet trace data using the TCP/IP real-time packet trace interface, your programs might need additional changes.

Note: These improvements are automatically enabled for an OSA-Express2 Ethernet feature in QDIO mode that supports the TCP segmentation offload function. You cannot disable them.

Migration action: Change any programs affected by the TCP segmentation offload packet trace changes by following the steps outlined in the Communications Server publications.

IP Services: Allocate buffer size for Network Management Interface real-time packet trace data (Rec, as of R7)

Not required, but recommended to improve storage usage.

© Copyright IBM Corporation, 2007

Before z/OS V1R7, users of the Network Management Interface (NMI) for TCP/IP real-time packet tracing (SYSTCPDA) were expected to allocate a 64 KB buffer when copying the trace buffer into their storage. Starting in z/OS V1R7, the TMI initialization record sent to the NMI user provides the size of the data buffer that the user must allocate to receive the trace data.

Migration action: Although the default size of the data buffer is 64 KB, the same as the default in prior releases, IBM recommends that applications use the TMI initialization record field to determine the size of the buffer required to collect the trace data, and then allocate the buffer accordingly. For NMI applications that work with previous z/OS releases, the buffer size in the initialization record is zero, meaning that a 64 KB buffer is required.

IP Services: Remove customization of PPT entries in SCHEDxx parmlib members previously added for z/OS Load Balancing Advisor (Recommended, as of R7)

*Not required, but recommended to remove the SCHED*xx *entries that are no longer needed after you've installed z/OS V1R7.*

The z/OS Load Balancing Advisor function, which is integral to z/OS V1R7, was first shipped as new function PTFs for z/OS V1R6, z/OS V1R5, and z/OS V1R4. The PTF documentation for this function instructed you to add PPT entries to the SCHEDxx parmlib members for the z/OS Load Balancing Advisor (Advisor) and Agent (Agent).

z/OS V1R7 ships with the PPT entries already present. If you installed the PTFs and manually added the PPT entries on any of the back-level systems, you should remove the entries to avoid duplication when z/OS V1R7 is installed. (You must restore the PPT entries if falling back from z/OS V1R7 or if PPT entries are shared with pre-z/OS V1R7 systems.)

Migration action: Remove the PPT entries with PGMNAME of EZBLBAGE and EZBLBADV from the SCHED*xx* parmlib member.

IP Services: Modify the TN3270E server profile to continue SSLV2 support (Required-IF, as of R7)

Required if you have Telnet clients that only support the SSLV2 protocol for secure connections.

Before z/OS V1R7, Telnet SSLV2/NOSSLV2 parameter statements did not exist and the SSLV2 protocol was always available. Starting in z/OS V1R7, Telnet's default setting is to prohibit the use of the SSLV2 protocol and to require that the SSLV2 parameter be specified to allow that protocol. This might affect migration for older Telnet clients that only support the SSLV2 protocol for secure connections.

Migration action: Determine if your installation has any Telnet clients that only support the SSLV2 protocol. If so, upgrade those clients or add the SSLV2 Telnet parameter statement to allow the use of this protocol.

IP Services: Stop inferring TCP/IP stack status or system capability from certain GetAddrinfo results (Required-IF, as of R6 with PTF UQ92871)

Required if you used the results of GetAddrinfo invocations in the situations described above as a means to infer the capability of the system in terms of IPv6 or as a means to infer the status of the TCP/IP stack.

With APAR PQ92830, the GetAddrinfo resolver API was modified in three specific situations to return slightly different results depending on whether IPv6 capability is available in the system. This change has been integrated into the base code for z/OS V1R7. Refer to the migration action in the *z/OS Migration* book.

IP Services: Remove the backslash from double quotation marks when submitting REXEC and TSO requests in batch (Required-IF, as of R5 with PQ96308)

Required if you use double quotation marks on REXEC and RSH coammand to be executed on the remote host and you did not install the APAR.

Starting in z/OS V1R5, the REXEC and RSH TSO client commands are written in C. This changes the way special characters are passed to the remote execution server for processing. However, there was a problem in z/OS V1R6 and z/OS V1R5 such that if the command to be executed on the remote host contained double quotation marks, you had to precede the quotation marks with a backslash (\) in order for the quotation marks to be passed to the remote execution server. For example, you had to specify \"value\" if you wanted the character string "value" tobe passed to the remote server. If you did not add the backslash, then the quotation marks would be replaced by blanks.

APAR PQ96308 fixed the problem. If you installed the appropriate PTF (PTF UQ95549 on z/OS V1R6 or PTF UQ95548 on z/OS V1R5), you did not have to use the backslash. But if you did not install the appropriate PTF, and you therefore used the backslash, remove the backslash in z/OS V1R7 or later because it is no longer needed and, if not removed, will cause an error.

IP Services: Ensure the correct value for TRUNCATE in FTP.DATA (Required-IF, as of R6)

Required if you use the EZAFTPAC or EZAFTPAS sample files without alternation.

The TRUNCATE statement in the FTP.DATA configuration data set controls whether records are truncated when the input LRECL is longer than the output LRECL.

© Copyright IBM Corporation, 2007

Before z/OS V1R6, the sample FTP.DATA files, EZAFTPAC and EZAFTPAS, did not contain a TRUNCATE statement. Therefore, the default value of TRUE was in effect unless you added a TRUNCATE FALSE statement to your FTP.DATA file. Starting in z/OS V1R6, TRUNCATE FALSE is coded in the sample FTP.DATA files. TRUNCATE FALSE specifies that truncating of records is not allowed. With TRUNCATE FALSE, if records would be truncated because the input record LRECL is longer than the destination LRECL, the file transfer fails. If the FTP.DATA statement WRAPRECORD TRUE is coded, long records are wrapped, not truncated; no error is set.
Migration action: Verify that FTP.DATA is coded correctly for your installation or that the default, TRUNCATE TRUE, is acceptable.

IP Services: Use sock option access control (Required-IF, as of R6)
Required if either of the following is true: You are running a multilevel-secure environment -or- your security product distinguishes between a resource profile not defined and a user not permitted to that resource.
Socket option access control gives system administrators the ability to assign permission for z/OS users to set selected socket options using an SAF security server. Access control is provided for the SOL_SOCKET level, SO_BROADCAST option.
You may need to define a profile in the SERVAUTH class for the EZB.SOCKOPT.sysname.tcpname.SO_BROADCAST resource. Programs that send datagrams to broadcast addresses may fail if the user is not permitted at least READ authority to this profile.
TCP/IP's default behavior in a *nonmultilevel*-secure environment is to *allow* all users to set the SO_BROADCAST socket option so they may send broadcast datagrams when this profile is not defined. Some security products do not distinguish between a resource profile not defined and a user not permitted to that resource. If your security server does not make this distinction, you must define the socket option resource profile and permit users to it whenever the SERVAUTH class is active.
TCP/IP's default behavior in a *multilevel*-secure environment is to *disallow* all users to set the SO_BROADCAST socket option so they may send broadcast datagrams when this profile is not defined. If you plan to run in a multilevel-secure environment, you must define the socket option resource profile and permit users to it.
You will need to determine the TCP/IP programs that set the socket option SO_BROADCAST and send datagrams to a broadcast address. The users of these programs will require permission to this profile. A list of TCP/IP programs that would require users to have permission to this new profile can be found in the security chapter of *z/OS Communications Server: IP Configuration Guide*.
Migration action: 1. Determine if your security product distinguishes between a resource profile not defined and a user not permitted to that resource.
2. Determine if you are running a multilevel-secure environment (that is, if the RACF MLACTIVE option is active).
3. If step 1 or 2 is TRUE, then you must define a profile in the SERVAUTH class for the EZB.SOCKOPT.sysname.tcpname.SO_BROADCAST resource and permit programs that send datagrams to broadcast addresses with at least READ authority to this profile.

IP Services: Define a profile in the SERVAUTH class for VIPARANGE (Required-IF, as of R6)
Required if your security product does not distinguish between a resource profile not defined and a user not permitted to that resource.
A RACF SERVAUTH class profile can control which applications may bind() to define a Dynamic Virtual IP Address (DVIPA) using VIPARANGE. To restrict applications that can bind() to a DVIPA within a VIPARANGE, define the EZB.BINDDVIPARANGE.sysname.tcpname profile in RACF, and permit the applications that are allowed to bind(). If the profile is not defined, then the bind() is processed and no checking is performed.
Migration action: Define a profile in the SERVAUTH class for the EZB.BINDDVIPARANGE.sysname.tcpname resource.

IP Services: Remove duplicate VIPADYNAMIC statements for same dynamic VIPA address (Required-IF, as of R6)
Required if your TCPIP profile has multiple VIPADEFINE or VIPABACKUP statements for the same dynamic VIPA address (DVIPA).
The processing of VIPADYNAMIC statements is changed to work the same whether the statements are specified within a single profile or with separate profiles. In previous releases, if a Dynamic Virtual IP Address (DVIPA) was specified in more than one VIPADEFINE or VIPABACKUP statement within a single profile, only the last VIPADEFINE or VIPABACKUP statement containing the DVIPA would be processed. Now all statements within a single profile are processed.
Migration action: Edit your profile and remove all VIPADEFINE and VIPABACKUP statements for the same dynamic VIPA address so that only the one desired definition remains.

© Copyright IBM Corporation, 2007

IP Services: Ensure that TCPSTACKSOURCEVIPA is coded correctly (Required-IF, as of R6)

Required if you have a TCPSTACKSOURCEVIPA configuration option coded with an address that is neither a static VIPA nor an active dynamic VIPA.

The types of addresses allowed for the TCPSTACKSOURCEVIPA configuration option are now restricted to the following:

- Static VIPAs
- Active dynamic VIPAs

Migration action: Ensure that your TCPSTACKSOURCEVIPA configuration option is coded with an address that is a static VIPA or an active dynamic VIPA.

IP Services: Update handling policy statistics because of changed FLUSH processing (Required-IF, as of R6)

Required if you have applications or automation procedures that deal with policy statistics and that use the MODIFY REFRESH command to perform FLUSH processing.

In releases prior to z/OS V1R6, the MODIFY REFRESH command for the Policy Agent performs FLUSH processing if both of the following are true: (1) the FLUSH parameter is specified on the TcpImage statement and (2) the HFS configuration is changed. Starting in z/OS V1R6, only the first condition needs to be true; the second condition is irrelevant. That is, starting in z/OS V1R6, the MODIFY REFRESH command for Policy Agent performs FLUSH processing if FLUSH is specified on the TcpImage statement, regardless of whether the HFS configuration is changed.

Policy statistics are affected by this change. They might not accumulate over as great a period of time as they did prior to z/OS V1R6. If you have applications or automation procedures that rely on policy statistics to be collected over a certain period of time, you might need to alter them to account for the fact that a MODIFY REFRESH command resets those statistics. For example, you might have to delay issuing a MODIFY REFRESH command in your automation procedures until after policy statistics are collected.

Migration action: Update applications and automation procedures to take into consideration that policy statistics being collected in the TCP/IP stack are now always reset by MODIFY REFRESH with FLUSH because the command now always deletes and reinstalls all policies.

IP Services: Update the TCP/IP sample profile to start the SNMP TCP/IP subagent (Required-IF, as of R6)

Required if you use the sample profile and want the SNMP TCP/IP subagent to be started during stack initialization.

The TCP/IP sample profile is installed in data set *hlq*.SEZAINST(SAMPPROF). Prior to z/OS V1R6, the SACONFIG profile statement, which governs the SNMP TCP/IP subagent, appeared in the sample profile as follows:

```
SACONFIG ENABLED COMMUNITY public AGENT 161
```

This caused the SNMP TCP/IP subagent to be started during TCP/IP stack initialization. Because the TCP/IP subagent requires other configuration and other functions to be started outside of stack initialization in order to successfully initialize, as of z/OS V1R6 the SACONFIG profile statement in the sample profile is changed to only include the DISABLED keyword. If you use the sample profile as the base for your TCP/IP profile, and don't change the SACONFIG profile statement, the TCP/IP subagent is not started during stack initialization.

Migration action: Do either of the following:

- Remove the SACONFIG statement. If no SACONFIG statement is present in the profile, then the SNMP TCP/IP subagent is started with the default parameter values of public for the COMMUNITY keyword and a port of 161 for the AGENT keyword.
- Change the DISABLED parameter to ENABLED. The SNMP TCP/IP subagent will be started with the default parameter values of public for the COMMUNITY keyword and a port of 161 for the AGENT keyword.

IP Services: Migrate from the SMIv1 version of the SNMP IBM MVS TCP/IP Enterprise-specific MIB module (Required-IF, as of R6)

Required if you use the SMIv1 version of the MIB module.

Until z/OS V1R6, two versions of the SNMP IBM MVS TCP/IP Enterprise-specific MIB module were shipped and installed in HFS directory
/usr/lpp/tcpip/samples, with the following filenames:

- mvstcpip.mib — the SMIv1 language version of the MIB module
- mvstcpip.mi2 — the SMIv2 language version of the MIB module

As of z/OS V1R6, the SMIv1 version is no longer shipped. If you want to continue using SMIv1, you can use publicly available tools to convert the SMIv2 MIB module to an SMIv1 MIB module.

Migration action: If you want to continue using SMIv1, you can use publicly available tools to convert the SMIv2 MIB module (mvstcpip.mi2) to an SMIv1 MIB module.

© Copyright IBM Corporation, 2007

SNA Services: Remove switch major node definitions that define parallel EE definitions using the same local and remove IP addresses (Required-IF, as of R8)

Required if you have switch major nodes that define parallel EE connections between the same local and remove IP address pair.
As of z/OS V1R8, Communications Server restricts the activation of parallel Enterprise Extender (EE) connections between two IP addresses. An outbound request to activate an EE connection will fail if an existing connection is found that uses the same local and remote IP addresses. When the activation fails, message IST2123I is issued. An inbound request to start a parallel EE connection will be accepted.

Migration action:

1. Review your VTAMLST members for switch major node definitions that define PUs to the same adjacent EE endpoint with the same local and remote IP address.
2. Remove all but one of these definitions.

SNA Services: Ensure POAQLIM is adequate (Required-IF, as of R7)

Required if POAQLIM is not definited on the VTAM APPL definition statements for program operator applications.
POAQLIM is a parameter on the APPL definition statement for a program operator to limit the number of messages that can be queued to a program operator. In releases before z/OS V1R7, the POAQLIM default is 0, which means there is no limit. Starting in z/OS V1R7, the POAQLIM default is 5000. This default change has been made to help avoid a storage shortage. Previously, if the program operator quit receiving messages and too many messages queued up, a storage shortage could occur. Now, if POAQLIM is exceeded for a program operator, error message IST983E is issued and no more messages are sent to that program operator.
You should ensure that POAQLIM is greater than the largest number of messages that a program operator expects to receive in response to any single VTAM command.

Migration action:

1. Determine whether POAQLIM is already defined on the VTAM APPL definition statements for program operator applications. If it is, then no action is required.
2. If POAQLIM is not defined:
   a. Determine whether any program operator issues a command that can receive greater than 5000 messages in response to the command. One way to do this is to review VTAM command response output in the network log.
   b. If necessary, define POAQLIM greater than 5000 on the VTAM APPL definition statements for program operator applications. Be sure to define a high enough POAQLIM value that a program operator's use of commands and expected response messages falls within the value.

SNA Services: Accept the VTAMMAP CSMCMPID ALLCMPID default change (Required-IF, as of R7)

Required if you want the display to show the details of the CSM buffers.
The default value of VTAMMAP CSMCMPID ALLCMPID has changed. Before z/OS V1R7, the default displayed the details of all the addresses of all CSM buffers currently used by all component IDs. As of z/OS V1R7, the default of ALLCMPID is to display the summary.

Migration action: If you want the display to show the details of all the CSM buffers (as VTAMMAP CSMCMPID ALLCMPID did by default before z/OS V1R7), use the new DETAIL keyword with ALLCMPID. For example:
VTAMMAP CSMCMPID ALLCMPID DETAIL

SNA Services: Make changes to PPO, SPO, or CLISTs (Required-IF, as of R7)

Required if your existing PPO, SPO, or CLISTs automate on the affected output displays.
z/OS V1R7, several VTAM display enhancements were introduced in regard to Enterprise Extender. These changes may affect operations of your Primary Program Operator (PPO), Secondary Program Operator (SPO), or existing CLISTs._The following display commands are affected:

- **DISPLAY NET,RTPS**
- **DISPLAY NET,SESSIONS,SID=**
- **DISPLAY NET,ID=rtppuname,HPRDIAG=YES**
- **D NET,ID=generic_resource,IDTYPE=GENERIC**
- **D NET,AUTOLOG,SCOPE=ALL**

Migration action: Make any necessary changes to your automation procedures to handle the display changes.

SNA Services: Exclude new session failure RSCV messages (Required-IF, as of R7)

Required if your existing PPO, SPO, or CLISTs automate on the affected output displays, or if the new messages are not wanted.

© Copyright IBM Corporation, 2007

Prior to z/OS V1R7, several session failure problems could have been more easily resolved if the Route Selection Control Vector (RSCV) that was computed and used for the session was displayed before the session information was freed. Therefore, z/OS V1R7 introduces a new VTAM start option, RSIRFSMG, that indicates that the RSCV data should be displayed (when available) for sessions that fail to reach the ACTIVE state.
In addition to introducing the new start option RSIRFSMG, z/OS V1R7 also changes the behavior of two existing start options. You can expect to see some new messages and existing messages issued in new cases. These changes may affect operations of your Primary Program Operator (PPO). Specifically, the behavior of the following two start options are changed:

- DSIRFMSG — New messages are optionally issued with the DSIRFMSG message group based on the new RSIRFMSG start option. The DSIRFMSG message group starting with IST663I is optionally issued when a search fails.
- SIRFMSG — New messages are optionally issued with the SIRFMSG message group based on the new RSIRFMSG start option. The SIRFMSG message group starting with IST663I is optionally issued when a session initiation request fails.

Migration action: If you decide you want the new behavior, you should set the RSIRFMSG start option value appropriately for your system in your start list and you should be aware of the DSIRFMSG and SIRFMSG changes. If, however, you want to maintain the behavior of your previously existing system and have the same output as prior releases, you must explicitly specify RSIRFMSG=NONE on the START command or use the MODIFY VTAMOPTS,RSIRFMSG=NONE command.

SNA Services: Prevent VTAM subarea nodes from automatically joining the ISTXCF sysplex group (Required-IF, as of R7)

Required if you want to prevent an APPN node that specifies an HPR value other than RTP, or a pure subarea node, from joining the ISTXCF sysplex group.

z/OS V1R7, XCFINIT is supported as a start option on pure subarea VTAM nodes (nodes without the NODETYPE start option specified) and the default value for XCFINIT on a pure subarea node is XCFINIT=DEFINE. Therefore, if you have pure subarea VTAM nodes in your sysplex when you upgrade to z/OS V1R7, those nodes will automatically join the ISTXCF sysplex group. Furthermore, as they learn of other nodes in the sysplex group, they will add definitions for XCF connectivity to those nodes.
The XCF connectivity definitions can be used by TCP/IP to communicate across XCF to other stacks in the sysplex. If you do not want this to happen for a specific pure subarea VTAM node, you must add XCFINIT=NO to your start option list for that pure subarea VTAM node.
In z/OS V1R7, APPN nodes that specify an HPR start option value other than RTP and have the XCFINIT start option either specified as or defaulted to YES will have XCFINIT forced to DEFINE. Prior to z/OS V1R7, it had been forced to NO. If you do not want these APPN nodes to connect to the XCF group, you must add XCFINIT=NO.
Migration action: If you have a subarea VTAM node that you want to prevent from automatically joining the ISTXCF sysplex group, add XCFINIT=NO to your start option list.
If you have an APPN node specifying an HPR value other than RTP and you do not want it to automatically join the ISTXCF sysplex group, add XCFINIT=NO to your start option list.

SNA Services: Eliminate messages caused by improved APPN search diagnostics (Recommended, as of R6)

Recommended to reduce the large number of messages that could be issued.

A new VTAM start option, LSIRFMSG, is introduced to provide additional information to help diagnose the reason for APPN locate search failures using the IST663I message group. LSIRFMSG processing is affected by the SIRFMSG= and CPNAME= operands on CDRSC definitions. Therefore, you might see an increase in the number of messages that are displayed for resources that you currently have defined with these operands. If you have existing CDRSC definitions with SIRFMSG=ALLSSCP or SIRFMSG=OLUSSCP and CPNAME=, consider whether you still want SIRFMSG coded.
Migration action: If you no longer want to see messages displayed as a result of SIRFMSG, DSIRFMSG, and LSIRFMSG processing, remove the SIRFMSG coding from CDSRC definitions.

SNA Services: Adjust automation for changed RTP PU display (Required-IF, as of R6)

Required if the previous release you were using one or both of the display commands in the description below and automating on the responses.

Problems in the data link control (DLC) layer used by a High-Performance Routing (HPR) PU might cause data flow to stall. Prior to z/OS V1R6, the operator was not informed of the stall and there was no automatic recovery. Starting in z/OS V1R6, a stalled HPR pipe solution is provided to detect the stall, inform the operator, and attempt to recover. As a result of the stalled HPR pipe solution, there is a difference in some of the display RTP PU command responses. Specifically, stall information fields are included in the DISPLAY NET,RTPs and DISPLAY NET,ID=*name*,HPRDIAG=YES command

© Copyright IBM Corporation, 2007

responses. Because of these new fields, the message formats are changed and might adversely affect automation. The changed messages are IST1695I, IST1858I, and IST1859I. In addition, message IST1696I is replaced by message IST1960I.

Migration action: Change any programs that automate on IST1695I, IST1696I, IST1858I, and IST1859I to take into consideration the new format of the messages.

Communications Server Migration Actions Post-First IPL

IP Services: Run the TN3270E Telnet server as a separate address space (Recommended, as of R6)

Recommended in order to take advantage of the benefits described in below and because it will become a requirement in the release after z/OS R8.

z/OS V1R8 is planned to be the last release that supports the in-stack version of the TN3270E server (that is, the TN3270E server running as part of the TCP/IP address space). After z/OS V1R8, the only way to run the TN3270E server will be in its own address space.

The TN3270E server has been able to run in its own address space since z/OS V1R6. This enhancement provides visibility and control over the TN3270E function separate from the TCP/IP stack. For example, users can run the TN3270E server at a different priority than the TCP/IP stack. This enhancement also provides the ability to stop and restart the TN3270E server without stopping the TCP/IP stack. This makes it easier to reset the server or install maintenance. Overall, problem diagnosis is easier and better when the TN3270E server and the TCP/IP stack are separate.

In z/OS V1R8, users can continue to run the TN3270E server as part of the TCP/IP address space. But, as stated above, this support is planned to be discontinued in the release following V1R8, and the only option will be to run the TN3270E server in its own address space.

Migration action:

1. Set up a superuser user ID with an OMVS segment, which may include adding a RACF user ID. See the Telnet chapter in *z/OS Communications Server: IP Configuration Guide* for details.
2. Assign the Telnet procedure name to the RACF STARTED class or other equivalent security measure and associate it with the user ID.
3. In z/OS V1R6, the default program property table (PPT) has been updated for the EZBTNINI program. Ensure that your SCHED*xx* parmlib member does not specify EZBTNINI or modify the IBM defaults for EZBTNINI (NOCANCEL, KEY(6), PRIV, NOSWAP, and SYST).
4. Change the priority of Telnet by assigning Telnet to a service class other than the default SYSSTC class in the STC subsystem.
5. Remove the TN3270E Telnet server profile statements from the TCP/IP profile.
6. Create JCL to run Telnet in its own address space. The Telnet profile statements data set is specified in the JCL. Samples of each are in hlq.SEZAINST. Member EZBTNPRC contains the JCL sample and member TNPROF contains a sample profile.
7. Update any automated operator scripts to specify the name of the Telnet address space as the PROCNAME field in any DISPLAY or VARY commands for TELNET options.
   Restriction: The obsolete options described for the DISPLAY TCPIP,,TELNET command are not supported by the Telnet server address space. Any use of these will also need to be converted to the equivalent ObjectID or ClientID forms.
   Once these steps are done, the stand-alone Telnet JCL can be started.

IP Services: Ensure that the translated OMPROUTE message catalog timestamp matches (Required-IF, as of R8)

Required if you are translating the OMPROUTE message catalog.

Starting in z/OS V1R8, OMPROUTE checks the actual timestamp on the message catalog against the timestamp variable that is generated by the gencat process. If the two timestamps do not match, the default message table is used. If you are going to translate the OMPROUTE message catalog, you must ensure that the two timestamps match.

Migration action:

1. Run dspcat -t -g omprdmsg.cat to retrieve the output needed to be supplied to the gencat command to generate your translated message catalog, as well as the timestamp of the current message catalog. The output looks like this:
   The time stamp of catalog omprdmsg.cat is: 2004 112 00:11 UTC
   $delset 1
   $set 1
   $quote "

© Copyright IBM Corporation, 2007

2. Remove the first line of the output of the dspcat and replace it with $timestamp so that the file now looks like this:
   $timestamp 2004 112 00:11 UTC
   $delset 1
   $set 1
   $quote "
3. Make any other appropriate changes to customize your message catalog and run gencat to generate the new message catalog.

IP Services: Rejoin a sysplex group (Required-IF, as of R7)

Required if you want the stack to automatically rejoin the sysplex group (without operator intervention) after all detected problems are cleared.

z/OS V1R7, a new keyword, AUTOREJOIN/NOAUTOREJOIN, was added to the SYSPLEXMONITOR parameter on the GLOBALCONFIG statement in the TCP/IP profile to control whether the TCP/IP stack should automatically rejoin the sysplex group when all detected problems (that caused the stack to leave the group) have been relieved. The default value of the parameter is to disable this function and not to automatically rejoin the sysplex group. However, even if you choose to not use the new function and you do not change the default to enable it, you will still see some changes. Specifically, starting in z/OS V1R7, when a stack leaves the sysplex group (either by the VARY TCPIP,,SYSPLEX,LEAVEGROUP command, or by the Sysplex Autonomics function), you will notice the following changes:

- The stack's VIPADYNAMIC configuration information is no longer discarded when the stack leaves the sysplex group. Instead, it is saved and restored when the stack rejoins the group.
- The way to cause the stack to rejoin the sysplex group has changed.
- DYNAMICXCF and VIPADYNAMIC statements are ignored when the stack has left the sysplex group. However, as with the previous release, existing DYNAMICXCF links remain active but new DYNAMICXCF links cannot be activated while the stack is out of the sysplex group.

In addition, the Netstat VIPADCFG/-F report has changed in z/OS V1R7. .

Note: Although the default keeps the stack from automatically rejoining the sysplex group, it is recommended that you enable this function; it is intended to benefit you. To enable it, specify the existing keyword RECOVERY and the new keyword AUTOREJOIN.

IP Services: Use multilevel security configuration consistency check (Recommended, as of R6)

Recommended in a multilevel-secure environment to remove inconsistencies found in your TCPIP.PROFILE and security server resource profiles.

Secure communication in a multilevel-secure environment requires configuration of several statements in the TCPIP.PROFILE and security server resource profiles in the SERVAUTH, SECLABEL, and STARTED classes. Inconsistencies in this configuration can allow unintended communication or prevent intended communication.

When the RACF MLACTIVE option is set, TCP/IP automatically checks the TCPIP.PROFILE and security server resource profiles for consistency. Consistency checking occurs at TCP/IP initialization when a VARY,TCPIP,,OBEYFILE command is processed and when a RACLIST REFRESH is done on the SERVAUTH or SECLABEL class. TCP/IP writes an informational message to the job log for each inconsistency detected. If inconsistencies are found, a final message, EZD1217I, summarizing the number of problems found is written to the system console.

Migration action: While running with the RACF MLACTIVE option active:

1. Specify NOMLSCHECKTERM on the GLOBALCONFIG statement in the TCPIP.PROFILE to continue running the TCP/IP stack when inconsistent configuration information is detected.
2. Check the job log for messages in the range EZD1218I–EZD1234I whenever message EZD1217I appears on the system console.
3. Correct your configuration as indicated by the job log messages until TCP/IP no longer detects any errors.
4. Specify MLSCHECKTERM on the GLOBALCONFIG statement in the TCPIP.PROFILE to bring down the TCP/IP stack when inconsistent configuration information is detected.

IP Services: Make changes to continue using old X Windows and Motif libraries (Required-IF, as of R6)

Required if you want to continue using the old X Windows and Motif libraries (X11R6.1 and Motif 1.2).

z/OS V1R6 Communications Server provides new versions of the X Window System and Motif libraries. The new libraries are based upon X Window System X11R6.6 and Motif 2.1.30. The old X Windows and Motif libraries (X11R6.1 and Motif 1.2) are still provided and are compatible with existing programs. They are compiled with 31-bit non-XPLINK and use IBM Hexadecimal Floating Point. These DLLs do not contain new function. Existing applications that use the old libraries will

© Copyright IBM Corporation, 2007

continue to run unchanged. However, if you want to compile and link with the old libraries, some changes are necessary to use the correct headers and libraries. See the *z/OS Migration* book for details.

SNA Services: Update automation routines that involve MODIFY TOPO or DISPLAY TOPO (Required-IF, as of R8)

Required if you use EE Connection Network Reachability Awareness and you hvae automation that looks for specific message in response to MODIFY TOPO,FUNCTION=CLRUNRCH or DISPLAY TOPO,LIST=UNRCHTIM commands.

In z/OS V1R8, the MODIFY TOPO and DISPLAY TOPO commands have changed as follows:

- To improve flexibility, the ORIG, VRN, and DEST operands replace the ID operand of the following:
  - The MODIFY TOPO,FUNCTION=CLRUNRCH command.
    Error message IST2060I is no longer issued in response to the MODIFY TOPO,FUNCTION=CLRUNRCH command.
  - The DISPLAY TOPO,LIST=UNRCHTIM command.
    Error messages IST1308I and IST2060I are no longer issued in response to the DISPLAY TOPO,LIST=UNRCHTIM command.
- To improve the accuracy of indicating that a virtual node has exceeded its unreachable partner limit, messages IST2058I and IST2112I are no longer displayed when a DISPLAY TOPO,LIST=ALL command is issued for an Enterprise Extender virtual routing node (VRN). Message IST2151I now displays when a DISPLAY TOPO,LIST=UNRCHTIM command is issued.

Migration action: Follow the instructions in the *z/OS Migration* book for automation changes.

SNA Services: Rework network automation CLISTs that process output from a DISPLAY EE or DISPLAY EEDIAG command that specifies HOSTNAME filters (Required-IF, as of R8)

Required if network automation CLISTs parse output from a DISPLAY EE or DISPLAY EEDIAG command that specifies host name filters.

In z/OS V1R8, the processing of host name filters specified on a DISPLAY EE or DISPLAY EEDIAG command has been enhanced. In previous releases, VTAM simply matched the host name filters to the known host names that were resolved when activating the Enterprise Extender (EE) connection. In certain instances, the remote host names were not always known to VTAM, making the display results inconsistent. With z/OS V1R8, VTAM performs TCP/IP name-to-address resolution to acquire the defined IP addresses.

These IP addresses are then used to identify the EE connections to process for the display output. The IP addresses are always known to VTAM for active EE connections, making the display results consistent. As a result, the outputs for the DISPLAY EE and DISPLAY EEDIAG commands using host name filters have been changed. These changes may affect operations of your Primary Program Operator (PPO), Secondary Program Operator (SPO), or existing CLISTs.

Migration action: Update any automation CLISTs that parse the output from a DISPLAY EE or DISPLAY EEDIAG command that specifies host name filters.

© Copyright IBM Corporation, 2007

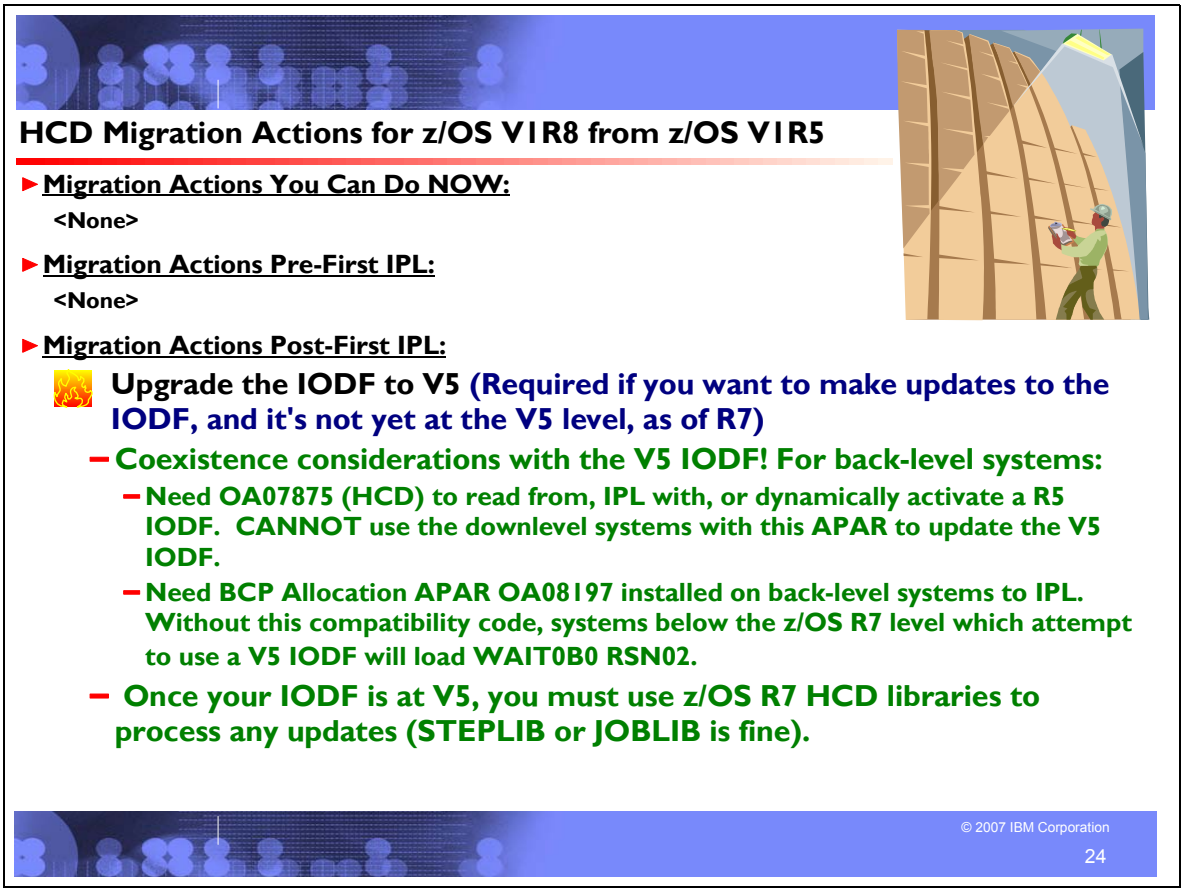

HCD Migration Actions Between z/OS V1R5 and z/OS V1R8

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

HCD Migration Actions You Can Do Now

<none>

HCD Migration Actions Pre-First IPL

<none>

HCD Migration Actions Post-First IPL

Upgrade the IODF to V5 (Required-IF, as of R7)

Required if you need to update the IODF but the IODF is not yet at the z/OS V1R7 level (that is, IODF V5).
A new IODF level, called V5, is introduced in z/OS V1R7. This new IODF level reduces the size of the IODF and improves the processing performance of large configurations. Once z/OS V1R7 is running, updates are only allowed to an IODF that has been upgraded to the V5 level.

Coexistence consideration: To read from, IPL with, and dynamically activate an IODF at the V5 level, the PTFs for coexistence APARs OA07875 and OA08197 are required on back-level systems. If you attempt to IPL with a V5 IODF from a back-level z/OS system that does not have the PTF for APAR OA08197 installed, a wait state occurs. The coexistence PTF for APAR OA07875 allows you to view and dynamically activate the V5 IODF, but does not allow you to update the V5 IODF from back-level systems. Once the IODF has been upgraded to V5, the z/OS V1R7 HCD libraries must be used to process updates to it. (A STEPLIB or JOBLIB from a back-level system to the z/OS V1R7 HCD libraries, or to a copy of the libraries, is acceptable.)

Note: APAR OA07875 is available for z/OS V1R4 systems that have the z/OS V1R4 z990 Compatibility Support feature installed (that is, HCD FMID HCS7708). z/OS V1R4 HCD **without** the z990 Compatibility Support feature (that is, HCD FMID HCS6091) does **not** have an applicable coexistence PTF.
(**Tip:** HCD FMID HCS7708 shows up as "z/OS V1.4 HCD" on its primary panel and is described as "z/OS V1.4 HCD" in the documentation. HCD FMID HCS6091 shows up as "OS/390 Release 9 HCD" on its primary panel.)

If you are running z/OS V1R4 **without** the z990 Compatibility Support feature, an additional coexistence and fallback requirement applies: you cannot read or dynamically activate a V5 IODF from the z/OS V1R4 system. If you want to read or dynamically activate a V5 IODF from the z/OS V1R4 system, install the z/OS V1R4 z990 Compatibility Support feature (or its replacement, the z/OS V1R4 z990 Exploitation Support feature,) and install the PTF for APAR OA07875.

Other methods to satisfy this coexistence requirement are:

- Do not share the IODF. Use EXPORT/IMPORT or COPY to create a second IODF.
- Delay making IODF updates from the z/OS V1R7 or later system until all systems are at z/OS V1R7 or later.

Restrictions: Once z/OS V1R7 is running, updates are only allowed to an IODF that has been upgraded to the V5 level.

Migration action: Upgrade your IODF to the V5 level in order to allow updates to device configurations, following the instructions in *z/OS HCD User's Guide*.
If you wish to make updates to your V5 IODF from back-level systems, ensure that you invoke the z/OS V1R7 HCD libraries to make the updates. A STEPLIB or JOBLIB from the back-level system to the z/OS V1R7 HCD libraries, or to a copy of the libraries, is acceptable.

© Copyright IBM Corporation, 2007

Some DFSMS Migration Actions for z/OS V1R8 from z/OS V1R5

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

mig

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

- Determine which VSAM data sets were allocated with IMBED, REPLICATE, or KEYRANGE by using tool IMBDSHIP.JCL.TRSD.
- Install the PTF for APAR OA10952. The PTF fixes the problem in which an exported copy of a VSAM data set that has the IMBED or REPLICATE attribute (or both) does not remove the attribute when the IMPORT is done.
- Schedule a time for the VSAM data sets to be unavailable, and redefine them.
- When support for IMBED and REPLICATE is removed and data sets are recalled or restored by DFSMShsm or DFSMSdss, these attributes will be removed during the recall or restore and the data set will be in a usable state.

➡ **Define new security profiles for DFSMShsm fast replication support (Required-IF, as of R8)**

- FACILITY class resource names have been shortened for fast replication commands, so you can specify the full names of copy pools. If you don't have the new profiles, hsm commands fail!!!
  - STGADMIN.ARC.FRBACKUP.* -> STGADMIN.ARC.FB.* , for instance
- Coex consideration: mixed releases in an HSMplex need to have both older and newer policies. Don't convert on pre-R8 to newer policies until all systems are running R8, to allow you to recover indiv ds from any copy pool backup version made on any system in HSMplex.

© 2007 IBM Corporation

Some DFSMS Migration Actions for z/OS V1R8 from z/OS V1R5

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

➡ **Use the new default for BLOCKTOKENSIZE in IGDSMSxx (Recommended, as of R8)**

mig

- Added in R7, specifies whether applications that use certain system functions have to code BLOCKTOKENSIZE=LARGE on the DCBE macro before using large format data sets.
- Before: default was REQUIRE. Now: default is NOREQUIRE.
- If you do override the R8 default (and use REQUIRE), programs will have to have BLOCKTOKENSIZE=LARGE explicitly coded on the DCBE macro in most cases to exploit large format sequential data sets.

Use catalog search order, instead of JOBCAT and STEPCAT (Required-IF, as of R7)

- In z/OS R7, JOBCAT and STEPCAT support has been removed. You can no longer enable or disable them.

- **Migration Actions Post-First IPL:**

- **Ensure the integrity of SMS control data sets (Recommended)**

mig

- **Update, translate, validate, and activate SMS policies from a R8 system, when in a mixed-level sysplex**

Avoid issuing commands against a copy pool from a pre-R8 system after the copy pool has been converted to R8 format (Required, as of R8)

- In R8, you can dump to tape copy pool versions residing on DASD. This enh changes DFSMShsm control ds records automatically first time you issue an FRBACKUP command against a copy pool on R8. But from then on, the only FR command that can be successfully issued against a copy pool from a pre-R8 system is FRRECOV from DASD.
- Ensure you longer need to back up copy pools from any pre-R8 system before issuing an FRBACKUP command against a copy pool on R8.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

DFSMS Migration Actions Between z/OS V1R5 and z/OS V1R8

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

DFSMS Migration Actions You Can Do Now

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

Recommended because it will become a requirement in a future release.

No supported release of z/OS honors the IMBED, REPLICATE, and KEYRANGE attributes for new VSAM data sets. In fact, using these attributes can waste DASD space and often degrades performance. Servicing these VSAM data sets has become increasingly difficult. In some cases, unplanned outages have occurred. To provide an environment in which outages are less likely, IBM has announced the removal of the support in existing VSAM data sets for the IMBED, REPLICATE, and KEYRANGE attributes. IMBED and REPLICATE were intended as performance improvements and have been obsoleted by newer, cached DASD devices. Striped data sets provide much better performance than KEYRANGE and should be viewed as a candidate for any existing KEYRANGE data sets.

Migration action:

1. Determine which VSAM data sets were defined with the IMBED, REPLICATE, or KEYRANGE attribute. To help you perform this task, you can get a tool that reads existing VSAM data sets and reports which ones have these attributes. The tool is available from the software server (ftp.software.ibm.com) in the s390/mvs/tools directory as IMBDSHIP.JCL.TRSD. Download the file in binary format and unterse it on your z/OS system using TRSMAIN. Instructions for using the tool are included in the downloaded JCL. See APAR II13894 for additional information.

   Important note: The tool only checks data sets that are on DASD. Data sets that were migrated by DFSMShsm or backed up by DFSMSdss will continue to retain their IMBED, REPLICATE, and KEYRANGE attibutes when recalled or restored. When support for IMBED and REPLICATE is removed and data sets are recalled or restored by DFSMShsm or DFSMSdss, the IMBED and REPLICATE attributes will be removed during the recall or restore and the data sets will be in a usable state. Because the data sets will be in a usable state, there is no need for a tool to identify data sets that were backed up or migrated at a time when the data sets had the IMBED or REPLICATE attribute and support for IMBED and REPLICATE still existed.

2. Install the PTF for APAR OA10952. The PTF fixes the problem in which an exported copy of a VSAM data set that has the IMBED or REPLICATE attribute (or both) does not remove the attribute when the IMPORT is done. Note that this change helps further remove, through normal processing, data sets that have the IMBED or REPLICATE attribute. The PTFs for APAR OA10952 are UA17230 for z/OS V1R6, UA17229 for z/OS V1R5, and UA17237 for z/OS V1R4.
3. Schedule a time for the affected VSAM data sets to be unavailable, and redefine them.

Please note that while no announcement has been made of the specific date or release that data sets with these attributes will no longer be able to be opened, it would be prudent to use this tool to identify the data sets in your installation and begin converting them to versions that do NOT contain these attributes.

Replace the STARTNUMBER and LOCATION(BOTH) options in vital record specifications (Required-IF, as of R8 coexistence APAR OA13355 back to z/OS R5)

Required if you use the unsupported options.

The STARTNUMBER and LOCATION(BOTH) options in vital record specifications have been unsupported but tolerated for many releases. One of the functions of APAR OA13355 on z/OS V1R5, R6, and R7 (and integrated in z/OS V1R8) is to remove the toleration.

Before the APAR, DFSMSrmm ignored STARTNUMBER so that the latest generation or copy was sent to the first storage location. For LOCATION(BOTH), DFSMSrmm added the two STORENUMBER values together. The data set was kept in its current location until the total store number was reached. DFSMSrmm kept the correct number of data sets or volumes but not necessarily in the location you wanted. Inventory management vital record processing issued message EDG2221E, set a return code of 4, and continued. With the APAR, DFSMSrmm issues message EDG2221E for each vital record specification with the unsupported options and ends inventory management with a return code of 12.

If you use the unsupported options, you should replace them to avoid abnormally ending inventory management.

© Copyright IBM Corporation, 2007

Migration action: Check the message file of your last run of inventory management vital record selection. If you see message EDG2221E and return code 4 during inventory management, EDGHSKP, VRSEL processing, replace the vital record specification according to the system programmer response for message EDG2221E.

Convert ISAM data sets and programs to VSAM (Required-IF, as of R7)

Required if you have any ISAM data sets in use or programs that use ISAM.

In z/OS V1R7, IBM removed most support for the indexed sequential access method (ISAM) and indexed sequential data sets. Users are no longer able to process (create, open, copy, or dump) ISAM data sets. Current ISAM data sets can still exist, and users can display information about them and delete them. Before migrating to z/OS V1R7, any ISAM data sets still in use should be copied or converted to VSAM data sets. Subject to existing restrictions, programs that use ISAM will continue to work using the ISAM interface to VSAM.

Migration action:

1. Find out if you have ISAM data sets or programs that use ISAM:

   For data sets: If you know your system has no ISAM data sets, you do not have to take migration steps. If you are not sure whether your system has ISAM data sets in use, you can take one of the following actions to find out:
   - Check for SYS1.ISAMLPA — if SYS1.ISAMLPA is not in your existing LPA concatenation or LINKLST concatenation, then programs are not using ISAM to access ISAM data sets (they could still be using the ISAM interface to VSAM). Note: the system data set SYS1.ISAMLPA and the distribution library SYS1.AOSD8 no longer exist in z/OS V1R7. If SYS1.ISAMLPA is in your LPA concatenation, you should remove it. This will save you about 177,528 bytes below the 16 MB line in LPA. Removing this will not affect programs that use the ISAM interface to VSAM data sets.
   - Use the DCOLLECT function of IDCAMS to search online volumes for a data set that has IS or ISU indicated for the DSORG value.
   - Check system logs for message IEC134I – starting in OS/390 V2R10, this message has been issued whenever ISAM data sets are opened, to warn that ISAM support is being discontinued.

   For programs: To find programs that use ISAM services, do the following:
   - Search for DSORG=IS in assembler source code and JCL. It appears on the DCB macro for BISAM and QISAM and on the DD statements. Search for equivalent statements in COBOL and PL/I programs. Then find the data sets used by these programs or JCL.
   - It is possible that you have an assembler language program that uses certain internal IBM macros that are no longer distributed. They were previously shipped in the SYS1.AMODGEN macro library. Their names, and any alternatives to use, are: IGGDEBD, IGGIOBD, IEZIOB, IGGLDCP, IGGLOAD, IGGSCAN, IGGWKNCP.
   - Look for any invocations of the IEBISAM utility program and remove them or ensure they are not used.
2. For ISAM data sets and the programs that use them, take one of the following actions before migrating to z/OS V1R7:
   - Delete the data sets or programs if no longer needed.
   - Copy the data to VSAM data sets and switch to using the ISAM interface to VSAM.
   - Convert the programs and data sets to VSAM.

Remove the CATALOG parameter from DEFINE PAGESPACE commands unless RECATALOG is specified (Required-IF, as of R7)

Required if you specified the CATALOG parameter on a DEFINE PAGESPACE command without the RECATALOG parameter.

Before z/OS V1R7, it was possible to specify a catalog on the access method services command to define a new page space. However, IBM recommended against the use of this CATALOG parameter, since it could conflict with the allocate specifications for the main page space data set. As of z/OS V1R7, the CATALOG parameter is no longer supported on the AMS DEFINE PAGESPACE command, unless it is used together with the RECATALOG parameter. Commands specifying the CATALOG parameter without RECATALOG will fail with the message IDC3171I.

Migration action: In jobs that issue the access method services DEFINE PAGESPACE command with the CATALOG parameter, delete the CATALOG parameter unless the RECATALOG parameter is also specified.

Ensure that the DFSMShsm journal is allocated correctly (Required, as of R7)

The DFSMShsm journal data set must be allocated with the following attributes:

- Single volume
- Single extent
- Contiguous space
- No secondary space allocation

© Copyright IBM Corporation, 2007

- Non-striped

Before z/OS V1R7, definition of the attributes was not enforced by the DFSMShsm code, even though the attributes were documented in *z/OS DFSMShsm Implementation and Customization Guide*. Starting with z/OS V1R7, their definition is enforced by the code. Therefore, you should make sure that you've defined your journal data set correctly. Failure to define the journal with the attributes above causes DFSMShsm to shut down. These attributes are required even if you use a large format data set for the journal.

DFSMS Migration Actions Pre-First IPL

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

Required if your installation builds or references DLLs.

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

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

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

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

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

Migration action:

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

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

Required if you are not using your old SYS1.IMAGELIB, you are installing with ServerPac or SystemPac, and you are using line mode printers such as the 3800 or 3900.

If you use line mode printers such as the IBM 3800 or the IBM 3900 running in line mode (not page mode), you must install library character sets, graphic character modification modules, and character arrangement tables in SYS1.IMAGELIB. This migration action does not apply if you are using IBM 3900 printers that are driven by PSF.

Migration action:

1. Run the LCSBLD1 job from the samplib data set to create character sets, graphic character modification modules, and character arrangement tables in SYS1.IMAGELIB.
2. Copy customized or locally-written FCBs and UCS images from your old system's SYS1.IMAGELIB data set to the new system's SYS1.IMAGELIB data set.

Add column ODLOBFL to OAM object directory tables (Required-IF, as of R8)

Required if you have any OAM object directory tables.. (The minimum level of DB2 required by z/OS V1R8 OAM is DB2 V6R1. However, because DB2 V6 is no longer service supported, you should use DB2 V7 or later.)

In z/OS V1R8, object access method (OAM) DB2 binary large object support enables objects larger than 32 KB (32640 bytes) to be stored using DB2's large object (LOB) support and the binary large object (BLOB) data type.

Before z/OS V1R8, for each object storage group, there was an object directory table and an object storage table hierarchy consisting of a 4 KB and a 32 KB table. With the new support in z/OS V1R8, a LOB storage structure is added to the object storage table hierarchy, consisting of a LOB base table and a LOB auxiliary table. A new column, ODLOBFL, is needed in the object directory table to track whether the object resides in a LOB storage structure.

Even if you do not take advantage of the large object support enhancement, you must add the new ODLOBFL column to the OAM object directory tables.

© Copyright IBM Corporation, 2007

Migration action: Modify and run the CBRSMR18 job in SYS1.SAMPLIB to add the new ODLOBFL column to the DB2 object directory tables for all object storage groups.

In addition: When migrating to any new release of z/OS, you must run OAM DB2 BIND jobs if you are using OAM for object support. The BIND jobs update DB2 with new OAM DB2 code. The BIND jobs always include CBRPBIND and typically include CBRHBIND and CBRABIND. To determine which jobs to run, see the chapter "Migrating, Installing, and Customizing OAM" in *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Object Support*.

Migrate from VRSEL(OLD) to VRSEL(NEW) in parmlib member EDGRMMxx (Recommended, as of R8)

Not required, but recommended because it will become a requirement in a future release and also because the EDGHSKP return code changes from 0 to 4.

The VRSEL(OLD) or VRSEL(NEW) operand of the OPTION command in DFSMSrmm parmlib member EDGRMM*xx* specifies the type of vital record processing you want performed. In a future release, VRSEL(OLD) is planned to be removed. Starting in z/OS V1R8, a warning message is issued when you run VRSEL(OLD) processing. If you do not migrate from VRSEL(OLD) to VRSEL(NEW) before moving to z/OS V1R8, you will receive warning message EDG2317E each time VRSEL is run and EDGHSKP processing will end with job step return code 4.

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

Migrate from JOBNAME(OPEN) or JOBNAME(ABEND) to JOBNAME(OPEN*) or JOBNAME(ABEND*) in data set name vital record specifications (Required-IF, as of R8)

Required if you have real job names of OPEN or ABEND and you want to continue to use existing vital record specifications to match real job names of either OPEN or ABEND.

Starting in z/OS V1R8, any vital record specification that includes JOBNAME(OPEN) or JOBNAME(ABEND) in its data set name vital record specification is used as a policy for retaining data sets that are either OPEN or have been closed by ABEND processing. You must change the vital record specification to use JOBNAME(OPEN*) or JOBNAME(ABEND*) if you want to continue to use existing vital record specifications to match real job names of either OPEN or ABEND. If you do not do this, DFSMSrmm uses the existing vital record specifications for OPEN and ABEND retention policies.

Migration action: To migrate from JOBNAME(OPEN) or JOBNAME(ABEND) to JOBNAME(OPEN*) or JOBNAME(ABEND*), follow the steps in Special ABEND and OPEN retention in *z/OS DFSMSrmm Guide and Reference*.

Stop using the path to bypass duplex alternate tape SYNCDEV operations (Recommended, as of R8)

Not required, but recommended because the patch is no longer beneficial.

Description: Before z/OS V1R8, the following patch was suggested as a method to potentially decrease elapsed recycle processing time during duplex alternate tape SYNCDEV operations:

```
PATCH .YGCB.+C8 BITS(.....1.)
```

Beginning in z/OS V1R8, this patch is no longer beneficial and no longer recommended. If you are using it, you should turn it off. If this patch is issued when SETSYS DUPLEX(MIGRATION(Y ERRORALTERNATE(MARKFULL))) is in effect, the patch is ignored and message ARC0824I is issued.

Migration action: Turn off the patch with the following command: `PATCH .YGCB.+C8 BITS(.....0.)`

Define new security profiles for fast replication commands (Required-IF, as of R8)

Required if you use DFSMShsm fast replication commands.

z/OS V1R8 changes the FACILITY class resource names for DFSMShsm fast replication commands. The names are shortened to allow you to specify the full names of copy pools (the static portions of the class names are shortened by six characters to allow you to specify the entire copy pool name). Along with this change, the maximum length of copy pool names is decreased in z/OS V1R8 from 30 characters to 23.

In mixed releases of z/OS in an HSMplex, you will need to define both the older profiles and the newer profiles. In z/OS V1R8, DFSMShsm checks for the newer profiles, but your pre-z/OS V1R8 systems will continue to use the older profiles. For pre-z/OS V1R8 systems in an HSMplex, do not actually convert your existing FACILITY class profiles until you have migrated all such systems to z/OS V1R8. Instead, retain the older profiles until all of your systems are at the z/OS V1R8 level. Saving this task until after the migration will allow you to recover individual data sets from any copy pool backup version made on any supported system in the HSMplex (z/OS V1R5 and later).

Migration action: To define new security profiles for the fast replication commands, do the following:

© Copyright IBM Corporation, 2007

1. Define the profiles shown in the table below:

Command name	New resource name (as of z/OS V1R8)	Old resource name (pre-V1R8 systems)
FRBACKUP	• STGADMIN.ARC.FB.*, or • STGADMIN.ARC.FB.*cpname*	• STGADMIN.ARC.FRBACKUP, or • STGADMIN.ARC.FRBACKUP.*cpname*
FRRECOV	• STGADMIN.ARC.FR.*, or • STGADMIN.ARC.FR.*cpname*	• STGADMIN.ARC.FRRECOV.*, or • STGADMIN.ARC.FRRECOV.*cpname*
FRDELETE	• STGADMIN.ARC.FD.*, or • STGADMIN.ARC.FD.*cpname*	• STGADMIN.ARC.FRDELETE.*, or • STGADMIN.ARC.FRDELETE.*cpname*

2. Activate the FACILITY class before starting DFSMShsm. When started, DFSMShsm checks to determine whether the FACILITY class is active. If you have not defined the new profiles, the DFSMShsm commands fail.

Recompile your user interaction module (Required-IF, as of R8)

Required if you invoke DFSMSdss through the API and use EIOPTION 20.

The data area mapped by EIREC20 has changed in z/OS V1R8. If you invoke DFSMSdss through the application programming interface (API) and use the Volume Notification exit option (EIOPTION 20), you must recompile your user interaction module (UIM).

Migration action: Recompile and link-edit your UIM using the z/OS V1R8 macro libraries.

Replace the INDD and INDYNAM keywords in COPY DATASET and DUMP DATASET jobs (Recommended, as of R8)

Recommended if your jobs issue the COPY DATASET or DUMP DATASET commands and you do not want to receive informational messages stating that DFSMSdss has replaced INDD and INDYNAM with the correct keywords.

The INDD or INDYNAM keyword on a COPY DATASET or DUMP DATASET command is used to perform either a logical data set copy or a physical data set dump, depending on the command.

Starting in z/OS V1R8, the INDD and INDYNAM keywords are being replaced as follows:

- For uses of the COPY DATASET command, the INDD and INDYNAM keywords are replaced with the LOGINDD and LOGINDYNAM keywords, respectively.
- For uses of the DUMP DATASET command, the INDD and INDYNAM keywords are replaced with the PHYSINDD and PHYSINDYNAM keywords, respectively.

You can still use INDD and INDYNAM in z/OS V1R8. However, if you do, you will receive an informational message stating that DFSMSdss has replaced INDD or INDYNAM with the correct keywords. To avoid receiving such messages, review and update any DFSMSdss jobs that issue these commands to use the correct keywords.

Migration action: Check your DFSMSdss jobs for usage of the COPY DATASET and DUMP DATASET commands. Update the JCL, as follows:

- For uses of the COPY DATASET command, replace any occurrences of the INDD or INDYNAM keywords with the LOGINDD or LOGINDYNAM keywords, respectively.
- For uses of the DUMP DATASET command, replace any occurrences of the INDD or INDYNAM keywords with the PHYSINDD or PHYSINDYNAM keywords, respectively.

Update automation that handles messages ARC1801I, ARC1820I and ARC1821I (Required-IF, as of R8)

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

The text for messages ARC1801I, ARC1820I, and ARC1821I has changed.

Before z/OS V1R8, the message text was as follows:

- ARC1801I FAST REPLICATION {BACKUP | BACKUP DUMP | BACKUP DUMPONLY | PREPARE | RECOVERY | RECOVERY FROM DUMP | DATA SET RECOVERY} IS STARTING FOR {COPY POOL *cpname* | VOLUME *volser* | DATA SET *dsname*}, AT *time* ON *date*
- ARC1820I THE FOLLOWING VOLUMES IN COPY POOL *cpname* ARE IN AN ACTIVE FLASHCOPY RELATIONSHIP
  ARC1821I NONE OF THE VOLUMES IN COPY POOL *cpname* ARE IN AN ACTIVE FLASHCOPY RELATIONSHIP

Starting with z/OS V1R8, the message text is as follows:

© Copyright IBM Corporation, 2007

- ARC1801I FAST REPLICATION {BACKUP | BACKUP DUMP | BACKUP DUMPONLY | PREPARE | RECOVERY | RECOVERY FROM DUMP | DATA SET RECOVERY} IS STARTING FOR {COPY POOL *cpname* | VOLUME *volser* | DATA SET *dsname*}, AT *time* ON date[, TOKEN='*token*']
- ARC1820I THE FOLLOWING VOLUMES IN COPY POOL *cpname*, VERSION *vernum*, HAVE AN ACTIVE FLASHCOPY BACKGROUND COPY
- ARC1821I NONE OF THE VOLUMES IN COPY POOL *cpname*, VERSION *vernum*, HAVE AN ACTIVE FLASHCOPY BACKGROUND COPY

If you are automating on the message text (not the message number), your automation might be affected.
Migration action: Change your automation to expect the new message text format.
For message ARC1801I, TOKEN='token' is appended to the message whenever the TOKEN keyword has been specified on the FRBACKUP command.
Message ARC1820I or ARC1821I will be issued for each DASD backup version that exists for the copy pool. Before z/OS V1R8, these messages were only issued for the most recent DASD backup version.

Use the new default for BLOCKTOKENSIZE in IGDSMSxx (Recommended, as of R8)

Not required, but recommended in order to exploit large format sequential data sets in applications without having to explicitly specify in IGDSMSxx the ability to use them. By using the system default BLOCKTOKENSIZE of NOREQUIRE, it is easier for applications to exploit large format sequential data sets.
When large format sequential data sets were introduced in z/OS V1R7, the BLOCKTOKENSIZE parameter in the IGDSMS*xx* parmlib member was introduced. The IBM-supplied member contains no setting of this parameter, so the system assumes the default value of REQUIRE. This value requires applications that use certain system functions to code BLOCKTOKENSIZE=LARGE on the DCBE macro before using large format data sets.
Starting in z/OS V1R8, the BLOCKTOKENSIZE default in IGDSMS*xx* is changed to NOREQUIRE, meaning applications do not need to code BLOCKTOKENSIZE=LARGE on the DCBE macro in most cases. Because the default has changed, if you were using the default (REQUIRE) value in z/OS V1R7 and want to continue using that value in V1R8, you must explicitly code BLOCKTOKENSIZE(REQUIRE) in the IGDSMS*xx* member of SYS1.PARMLIB.
Migration action: To use the new default value, which is BLOCKTOKENSIZE(NOREQUIRE), do nothing. But if you want to use the value that was the default in z/OS V1R7, which is BLOCKTOKENSIZE(REQUIRE), code this value in parmlib member IGDSMS*xx*. If you do specify REQUIRE, remember that programs will have to have BLOCKTOKENSIZE=LARGE explicitly coded on the DCBE macro in most cases to exploit large format sequential data sets.
Note that if you want to change the BLOCKTOKENSIZE value after you have IPLed z/OS V1R8, you can do so by modifying IGDSMS*xx* and issuing the SET SMS=*xx* command. If the parmlib member contains no setting of the BLOCKTOKENSIZE parameter, the system does not change the current value in effect.

Use the catalog search order instead of JOBCAT and STEPCAT (Required-IF, as of R7)

Required if you use the JOBCAT or STEPCAT facilities.
As a first step toward removal of the DFSMSdfp JOBCAT and STEPCAT facilities, in z/OS V1R5 the JOBCAT and STEPCAT DD statements were disabled by default, but you could still enable or disable them. **Now, as of z/OS V1R7, they have been removed and you can no longer enable or disable them.** If you continue to specify the JOBCAT or STEPCAT DD statement, you will receive message IEFC034I JOBCAT OR STEPCAT NOT PERMITTED.

By way of background, the JOBCAT and STEPCAT facilities have been in existence for many years, predating the introduction of ICF (integrated catalog facility) catalogs. JOBCAT and STEPCAT were designed to address some of the functional shortcomings of VSAM catalogs, such as:

- VSAM volume ownership, that is, all data sets on a volume having to be in the same VSAM catalog. Multiple catalogs could not point to data sets on the same volume.
- Performance problems resulting from no multilevel alias support, as well as lack of ability to subset catalog data for recovery purposes.
- Restrictions in the definition of the catalog SVC interface.

The introduction of ICF catalogs in the mid-1980s and other catalog enhancements (such as the multilevel alias support) directly addressed those problems. In addition, processes were developed for system build to use system specific aliases instead of JOBCAT or STEPCAT. CBIPO introduced these processes and they are used today by offerings such as ServerPac to create data set entries in the new master catalog of the system being built.
At the time ICF catalogs were introduced, the JOBCAT and STEPCAT facilities were functionally stabilized. Neither SMS-managed data sets nor UCBs above the 16 megabyte line may be used with JOBCAT or STEPCAT. ICF catalogs

© Copyright IBM Corporation, 2007

contain sufficient functional capabilities that all functions that previously could only be performed with JOBCAT or STEPCAT can now be done without them.
Furthermore, the use of JOBCAT and STEPCAT can actually cause significant problems. Data sets are generally not cataloged according to the normal predictable search order when JOBCAT or STEPCAT is used. This impacts the ability to do comprehensive installation storage management and can increase staff requirements. For example, interval migration and recall using DFSMShsm is effectively unusable when the data sets cannot be found using the standard catalog search order. The use of JOBCAT and STEPCAT can also result in noticeable increases in the time required to perform catalog requests.
Migration action: Change the jobs that use JOBCAT and STEPCAT to use the normal catalog search order, as described in *z/OS DFSMS Access Method Services for Catalogs*.

Tune DFSMShsm for the SSM multiple tasks enhancement (Recommended, as of R6)

Recommended if you use secondary space management.
The secondary space management (SSM) function is enhanced to run with multiple tasks in a single DFSMShsm host, allowing the function to complete more quickly. As a result of this enhancement, the following performance improvements have been made, and you might want to retune DFSMShsm because of them:

- The default number of tasks for SSM has increased from one to three: two cleanup tasks and one tape movement task.
- The default number of recall tasks has increased from 5 to 15 to aid in the processing of deletions that SSM initiates.
- Small data set packing (SDSP) data sets are no longer allocated for the entire SSM window.
- Deletion requests for expired migrated data sets are given a low priority so that user recalls and user-initiated deletions are processed before SSM-generated deletions.

In addition, a new SETSYS parameter, MAXSSMTASKS, allows you to specify the number of SSM tasks that you want to run concurrently.
Migration action:

- Determine the maximum number of SSM tasks that should be processed concurrently. Prior to z/OS V1R6, the default was 1. Now the default is 3: two cleanup tasks and one tape movement task. You can set the maximum by using the SETSYS MAXSSMTASKS command. Note that each SSM tape movement task requires two tape drives if you are using DUPLEXing for migration; this might affect the number of SSM tasks that you want to allow.
- Determine the maximum number of recall tasks that can be processed concurrently. Prior to z/OS V1R6, the default was 5. Now the default is 15 to accommodate SSM deletion requests. You can set the maximum by using the SETSYS MAXRECALLTASKS command. If a DFSMShsm host will not be performing SSM, you will probably want to reduce the maximum number of recall tasks.
- Because SDSPs are no longer allocated for the entire SSM window, you might want to reconsider your SDSP reorganization procedures.
- Reconsider any procedures you have in place to reduce or eliminate user-initiated recalls or deletes during SSM processing because SSM-initiated deletes are processed only after all user-initiated recalls and deletes have been processed.
- Determine the number of SDSP data sets that should be defined to your DFSMShsm system, now that the number of SSM tasks has increased from one to three (or more). It is important to plan the number of SDSP data sets in relation to the number of concurrent migration tasks and the amount of processing performed by functions with a higher usage priority for the SDSP data sets.

DFSMS Migration Actions Post-First IPL

Ensure the integrity of SMS control data sets (Recommended)

Recommended when using SMS in a mixed-level sysplex.
In a multisystem Storage Management Subsystem (SMS) complex, operating systems share a common set of SMS classes, groups, ACS routines, and a configuration base, which make up the storage management policy for the complex. This storage management policy is maintained in a source control data set (SCDS). When this policy is activated for SMS, the bound policy is maintained in processor storage and on DASD in an active control data set (ACDS). Systems in the complex communicate SMS information through a common communications data set (COMMDS).
It is strongly recommended that to successfully share SMS control data sets in a multisystem environment where there are mixed levels of DFSMS, you update, translate, validate, and activate SMS policies on the system with the latest level of DFSMS. If you do not translate your ACS routines and validate your SCDS on the latest level of DFSMS, the translation and validation might fail. This failure occurs because some constructs and definitions are known only to later levels of

© Copyright IBM Corporation, 2007

DFSMS. It can also occur due to changes in validation rules between releases. Once the control data sets are formatted by the later-level system, the SMS control blocks reflect the new rather than the back-level lengths and control information.

Migration action:

1. Install all coexistence PTFs, which you can find in the z/OS Migration book, or from the EPSPT.
2. Save the active ACDS as an SCDS with the SETSMS SAVESCDS command.
3. Update, translate, validate, and activate SMS policies using a z/OS V1R7 system. Remember, once you've translated using a z/OS V1R7 system, you risk failure if you update and translate using a back-level system.

Avoid issuing commands against a copy pool from a pre-R8 system after the copy pool has been converted to the R8 format (Required- as of R8)

Required, as of z/OS R8.

In DFSMShsm, a *copy pool* represents one or more storage groups that are copied collectively and can be recovered at the copy pool or volume level. Starting in z/OS V1R8, you can dump to tape copy pool versions residing on DASD. This enhancement changes DFSMShsm control data set (CDS) records. The records are converted to the new, z/OS V1R8 format automatically the first time you issue an FRBACKUP command against a copy pool. But from then on, the only fast replication (FR) command you can issue successfully against a copy pool from a pre-z/OS V1R8 system is FRRECOV from DASD. Therefore, before issuing an FRBACKUP command against a copy pool on z/OS V1R8, ensure that your installation no longer needs FRBACKUP processing for copy pools from any pre-z/OS V1R8 systems.

Migration action: Ensure that your installation no longer needs to back up copy pools from any pre-z/OS V1R8 systems before issuing an FRBACKUP command against a copy pool on z/OS V1R8. Notify operations personnel that after an FRBACKUP command has been issued against a copy pool, the only FR command that can be successfully issued against a copy pool from a pre-z/OS V1R8 system is FRRECOV from DASD.

Rename existing copy pools (Required-IF, as of R8)

Required if you plan to use the new FACILITY class resource names to protect individual copy pools names.

z/OS V1R8 restricts the names of new copy pools to a maximum of 23 characters. Before z/OS V1R8, copy pool names could be up to 30 characters in length. The shorter names will allow security administrators to specify full names of copy pools in the new FACILITY class resource names introduced in z/OS V1R8 (also see migration action "Define new security profiles for fast replication commands").

z/OS V1R8 continues to support the use of longer names (up to 30 characters) for your existing copy pools.

Migration action: Your actions to rename existing copy pools will depend on whether your installation has enough additional target DASD storage available to create a new copy pool backup version without using your existing target DASD storage.

If you have enough additional target DASD storage available, do the following:

1. Define the new copy pool with the same characteristics as the copy pool to be replaced.
2. Add enough additional target DASD storage to the copy pool backup storage group to create one additional version.
3. Enter the command FRBACKUP COPYPOOL(*new_cpname*) for the new copy pool to create a backup version.
4. Enter the command FRDELETE COPYPOOL(*old_cpname*) VER(*oldvernum*) for the oldest version of the existing copy pool to be replaced. This step will free enough target DASD storage to create the subsequent backup version for the new copy pool name.
5. Repeat Steps 3 and 4 until you have replaced all of the versions for the existing copy pool.
6. You can now remove the extra DASD storage from the copy pool backup storage group.

If your installation must use its existing target DASD storage to create a new copy pool backup version, do the following:

1. Define the new copy pool with the same characteristics as the copy pool to be replaced.
2. Enter the command FRDELETE COPYPOOL(*old_cpname*) VER(*oldvernum*) for the oldest version of the existing copy pool to be replaced. This step will free enough target DASD storage to create the subsequent backup version for the new copy pool name.
3. Enter the command FRBACKUP COPYPOOL(*new_cpname*) for the new copy pool to create a backup version.
4. Repeat Steps 2 and 3 until all of the versions for the existing copy pool have been replaced.

© Copyright IBM Corporation, 2007

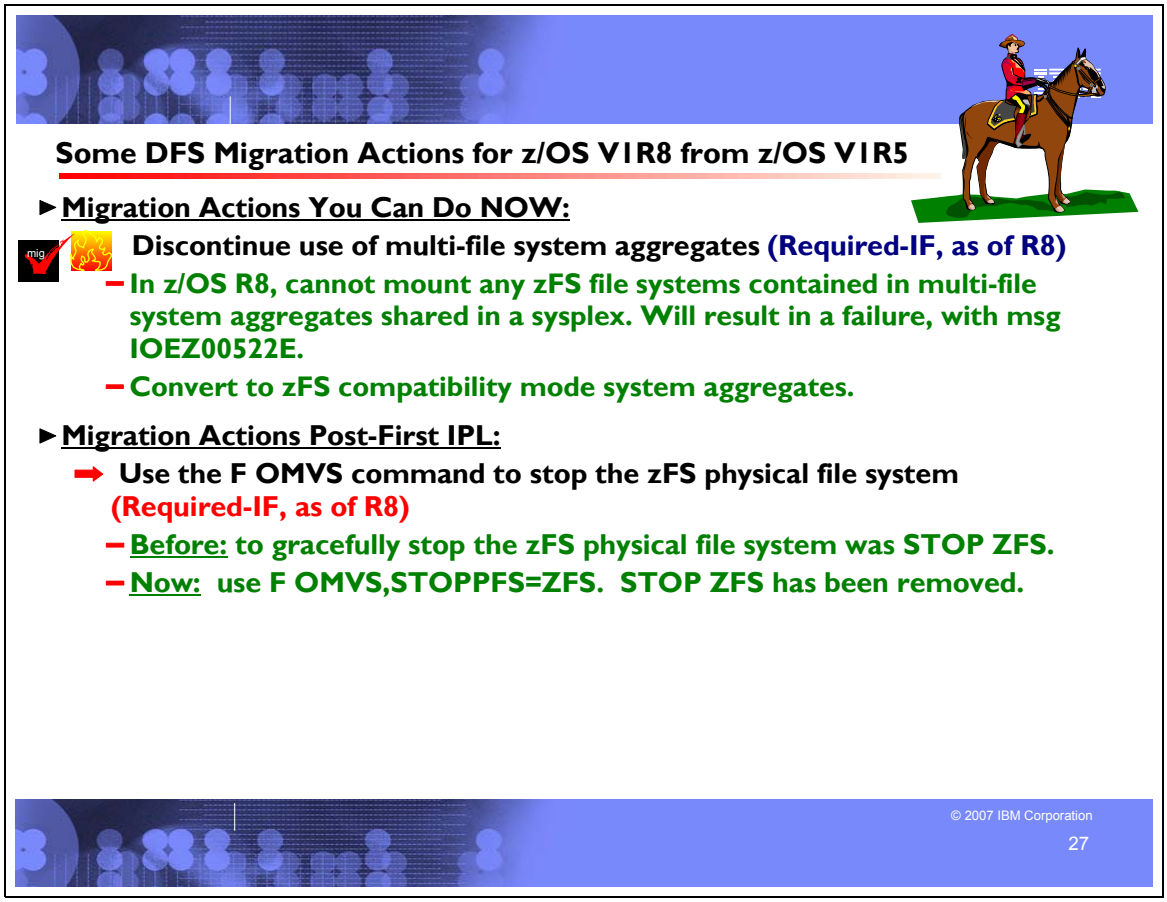

Distributed File Service Migration Actions Between z/OS V1R5 and z/OS V1R8
These migration actions were taken from *z/OS Migration.* Some descriptions and actions have been shortened for inclusion in this presentation. For the complete descriptions and actions, refer to *z/OS Migration.*

DFS Migration Actions You Can Do Now

Discontinue use of multi-file system aggregates (Required-IF, as of R8)

Required if you are using zFS multi-file system aggregates.
Beginning in z/OS V1R8, you can no longer mount zFS file systems contained in multi-file system aggregates that are shared across systems in a sysplex. IBM has previously recommended that these multi-file system aggregates not be shared in a sysplex environment. Any attempts to mount zFS file systems contained in multi-file system aggregates will fail in a z/OS UNIX shared file system environment with message IOEZ00522E. Mounting zFS compatibility mode aggregates, which have a single file system per data set, will continue to be supported in all environments.
In a future release, IBM plans to withdraw support for zFS multi-file system aggregates. When this support is withdrawn, only zFS compatibility mode aggregates will be supported. (A zFS compatibility mode aggregate has a single file system per data set.)
Migration action: Determine if you are using zFS multi-file system aggregates. You can determine this by scanning your zFS IOEFSPRM configuration options file for define_aggr statements. You should also scan your /etc/rc file for any zfsadm attach commands. If you are using zFS multi-file system aggregates, you should copy the data from each file system into its own zFS compatibility mode aggregate.

DFS Migration Actions Pre-First IPL

Remove extraneous characters in the FILESYSTYPE PARM option for zFS (Required-IF, as of R6)

Required if you use zFS parmlib support (that is, if you do not have an IOEZPRM DD statement in your ZFS PROC). Even if you are not using zFS parmlib support, this action is recommended to avoid future problems.

Before z/OS V1R6, any value specified in the PARM option of the FILESYSTYPE statement for zFS (in BPXPRM*xx*) was ignored. Starting in z/OS V1R6, the zFS FILESYSTYPE PARM option is interpreted as the zFS parmlib search specification. Therefore, any non-parmlib specification in the zFS FILESYSTYPE PARM option causes the initialization of the zFS physical file system to fail. (If you are not using zFS PARMLIB support, that is, if you have an IOEZPRM DD statement in your ZFS PROC, this type of failure does not occur. However, you should still remove any non-parmlib support specification from your zFS FILESYSTYPE PARM option so that you can use zFS parmlib support in the future.)
Migration action: Ensure that you do not have any extraneous characters in your zFS FILESYSTYPE PARM specification. Your zFS FILESYSTYPE statement should either have no PARM option, such as:
FILESYSTYPE TYPE(ZFS) ENTRYPOINT(IOEFSCM) ASNAME(ZFS,'SUB=MSTR')
Or it should have the parmlib support type specification, such as:
FILESYSTYPE TYPE(ZFS) ENTRYPOINT(IOEFSCM) ASNAME(ZFS,'SUB=MSTR') **PARM('PRM=(00)')**

DFS Migration Actions Post-First IPL

Use the F OMVS command to stop the zFS physical file system (Required-IF, as of R8)

Required if you need to stop the zFS physical file system.
Before z/OS V1R8, the command to stop the zFS physical file system gracefully was STOP ZFS. In z/OS V1R8, z/OS UNIX introduced a new F OMVS operator command to be used for stopping physical file systems: F OMVS,STOPPFS=*pfsname*. Use this new command to stop the zFS physical file system gracefully (with *pfsname* = ZFS).
Migration action: To stop the zFS physical file system gracefully, use the following command:
f omvs,stoppfs=zfs
Discontinue using the previously recommended command, stop zfs, because it has been removed from z/OS.

© Copyright IBM Corporation, 2007

Some JES2 Migration Actions for z/OS V18 from z/OS V1R5

Update JES2 exit routines due to the z/OS R7 structure changes (Required-IF, as of R7)

- Some NJE and input service functions were moved from the JES2 AS to the application AS. Results in improved performance, reliability, failure isolation, and outage avoidance. This can break processing done by the following JES2 exits:
  - Exits 2, 3, 4, 7, 20, 39, 46, and 47 are not called for INTRDRs and NJE/TCP processing.
  - Exit 8 is now called for INTRDR and NJE/TCP processing.
  - Exits 36 and 37 now get control in a different address space for INTRDR and NJE/TCP processing. These are pre-SAF and post-SAF exits.
  - Exit 13 has been deleted.
  - Exit 49 has new function and is being called for $S J commands.
  - Exit 40 can now be used to control NJE mail notification (function that was in exit 13).
  - The input to exits 2, 3, 4, 20, 46, and 47 has changed significantly.
- Read "*JES2 z/OS 1.7 Exit Migration Guide*" on web for important information!
  - http://www.ibm.com/servers/eserver/zseries/zos/installation/zos17_jes2_migration.html

© 2007 IBM Corporation

Some JES2 Migration Actions for z/OS V18 from z/OS V1R5

Update procedures because of changed input processing (Required-IF, as of R7)

- Input processing was rewritten, thus this change may affect your applications and procedures.
  - JCL/JECL processing: more details on errors, parsing of operands conforms to standard JCL rules (some previous errors may now run correctly)
  - SYSIN data processing: SYSIN ds properly handles carriage cntl, var recfm, and SYSIN up to 32KB long. Jobs that work on z/OS R7 may fail in prior releases. Submitted jobs from R7 will exec correctly w/ a mixed-level MAS.
  - Internal reader processing in requestor AS: Moved from JES2 AS to AS that allocated the internal reader. Higher CPU and IO now associated with the owning AS. Processing occurs at dispatching priority of the requestor and not JES2's priority, thus lower priority AS may not submit jobs as fast as higher priority AS.
  - Buffer of data in internal readers: Now, job cards are processed as soon as they are passed to the internal reader. No buffering of data, and the internal reader data space no longer exits.
  - Internal reader messages in JOB log: Input processing msgs are issued in the submitting AS and placed in the JOB log (JESMSGLG) of the submitting AS. Increases size of the JOB log. May want to spin or suppress the JOB log.
  - Internal reader processing when JES2 AS fails: JES2 AS used to discard job on an internal reader when JES2 AS abnormally terminated. Now, JES2 tries to retain jobs that are active on an internal reader when JES2 terminates.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

JES2 Migration Actions Between z/OS V1R5 and z/OS V1R8

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

JES2 Migration Actions You Can Do Now

Migrate to z2 mode (Required-IF, as of R7)

Required if you use R4 mode.

Before z/OS V1R7, JES2 could operate in either of two modes, as specified on the $ACTIVATE command: *compatibility (R4) mode* and *full function (z2) mode*. As of z/OS V1R7, JES2 can operate only in z2 mode; R4 mode is withdrawn. z2 mode enables larger limits for jobs, output elements, trackgroup space, and job number range, and improves trackgroup usage tracking. You must migrate to z2 mode before IPLing z/OS V1R7 JES2. Do not activate z2 mode until all systems in the MAS are at z/OS V1R2 or later JES2.

Migration action: In preparing for your migration to JES2 z2 mode, analyze the following:

- Exit routines and unique code that would be affected by activating JES2 in z2 mode. If your exit routines use JQE or JOE chaining fields, or reference job number fields, then z2 mode changes require you to update the exit routines. Field names have been changed to help identify where logic changes are needed.
  You might want to enable your exit routines to support both R4 and z2 modes to facilitate migration. If your exit routines must support systems running in both modes, be aware that fields within control blocks such as the JQE are mode sensitive. For example in R4 mode, use JQEJOBNO_R4 for the job number in a real JQE; in z2 mode, use JQEJBNUM for the job number. However if you are referencing the job number in a JQA (artificial JQE), use JQEJBNUM.
- How your JES2 system processes the JOE, JQE, JQA, and JOT control blocks that are significantly changed in z2 mode.
- Control blocks and fields used to store job numbers (such as the $SJB, $JIB, $FSAXB, and $COMWORK).
- How you use the following JES2 macro services:
  - *$QJQE and $#JOE:* These macros process the JQE and JOE control block chains using proper chaining fields based on the current checkpoint processing mode (either R4 mode or z2 mode). Therefore, users of these macros do not have to be sensitive to which JES2 processing mode is currently being used. $#JOE is changed to process chains (such as the CHAR JOE chain). Both $QJQE and $#JOE have improved their loop control and have usability enhancements.
  - *$DOGJQE:* The JQA returned by $DOGJQE always reflects the z2 mode processing format. Code that processes an artificial JQE (JQA) need not examine the JES2 processing mode (z2 mode or R4 mode) to examine JQA fields.
  - *$JQEJNUM:* $JQEJNUM is a new macro used to obtain the binary job number for a particular JQE. This macro operates in z2 mode and R4 mode and processes the JQA or JQE that is passed.
  - *$JBIDBLD:* This macro converts a binary job number to a printable job ID. $JBIDBLD is updated to accept a JQE (or JQA) as input and operates in z2 mode or R4 mode. $JBIDBLD formats the job ID, in the correct way, based on the current job number range.
- Your changes made to support JES2 in R4 mode, to ensure that they are compatible with z2 mode.

After your analysis is complete, perform the following tasks as appropriate for your installation:

- Convert any locally-developed code to use the new fields and field names.
- Process JOE and job queue element (JQE) control blocks as much as possible with the macros provided ($QJQE, $#JOE, $DOGJQE, $JQEJNUM, $JBIDBLD) instead of writing and using your own code.
- Ensure that all the field and control block changes identified with z/OS V1R2 JES2 are incorporated, including updates to the surrounding logic.

You can use the $ACTIVATE command ($ACTIVATE,LEVEL=R4) to revert to R4 mode from z2 mode. Before deactivating z2 mode, z/OS checks that no features exploiting z2 mode are active. If functions are exploiting z2 mode, the $ACTIVATE,LEVEL=R4 command fails with message HASP445. By default, JES2 restarts in the same mode (R4 or z2) as other members of the MAS (if any are active) or the mode the last active JES2 member was in when it came down. To restart JES2 in R4 mode, specify UNACT on PARM=.

Note: As of z/OS V1R7 JES2, the $ACTIVATE command has been deleted.

JES2 Migration Actions Pre-First IPL

Review automation that handles message $HASP003 (Required-IF, as of R8)

© Copyright IBM Corporation, 2007

Required if you have automation routines that examine message $HASP003.
Before z/OS V1R8 JES2, message $HASP003 displayed either the keywords being scanned when the error occurred (*trace-of-keywords*) or the number of the message in error (*nnn*). As of z/OS V1R8 JES2, *trace-of-keywords* lists the command verb (VERB from the $SCAN call) and list of keywords being scanned when the error occurs.
For example, the command $ADD NETSERV(1),TRACE=ALWAYS results in $HASP003 RC=(03),ADD NETSERV2 TRACE ALWAYS - INVALID PARAMETER STATEMENT, where ADD is the verb and NETSERV, TRACE, and ALWAYS are the operands where the error was detected. If the $HASP003 message is being issued because of an error building a message, *trace-of-keywords* is the number (*nnn*) of the message being constructed.
Also in z/OS V1R8 JES2, reason codes 106 and 107 have been added:
106 CANNOT SET DYNAMIC IF NOT ACTIVE ON ALL MEMBERS
107 SERVICE CLASS NOT DEFINED OR WLM ERROR
Migration action: Review and update automation that handles message $HASP003 to ensure that it can handle the new reason codes and text.

Review automation that handles message $HASP007 (Required-IF, as of R8)

Required if you have automation routines that examine message $HASP007.
Message $HASP607 displays the reason why a $PJES2 command is rejected. Before z/OS V1R8 JES2, RC 03 meant there were either active processors (including networking devices) or active address spaces that caused the $PJES2 command to be rejected. As of z/OS V1R8 JES2, RC 03 is split into the following three distinct reason codes:

- RC 03, active processors (excluding the line manager and execution processors)
- RC 10, active address spaces
- RC 11, active networking devices

Also, the text related to RC 06 is changed. Before z/OS V1R8 JES2, $HASP607 displayed additional text detailing the ASID and job name. As of z/OS V1R8 JES2, the additional text has been removed from the $HASP607 message and moved to the new $HASP608 message with more information.
Migration action: Review and update automation that handles message $HASP607 to ensure that it can handle the updated reason codes and text.

Review automation that handles message $HASP621 (Required-IF, as of R7 and rolled back via OA08768)

Required if you have automation routines that examine message $HASP621.
Message $HASP621 displays the number of work elements in a class. Before z/OS V1R7 JES2, if the number exceeded 99999, ***** is displayed. Be sure that your automation can handle the new display.
In z/OS R7, the format of the command is: OUT R=..... CLASS class= *nnnnn,...* where **class=***nnnnn,...* are the classes and the number of work elements in each class. **Note:** If the number of work elements exceeds 99999 then ***** will display.
Migration action: Review automation routines that handle message $HASP621 to make sure they can handle ***** for the number of work elements, and update any routines that can't handle *****.

Modify SAPI applications that set SSS2DUPJ (Required-IF, as of R7 and rolled back via OA04745)

Required if you have SAPI applications that select sysout data sets by specific job namde and sett SSS2DUPJ.
SAPI applications that select sysout data sets by specific job name and set SSS2DUPJ now work differently. Before z/OS V1R7, without the PTF for APAR OA04747 installed, if two jobs in the system matched the name specified in SSS2JOBN and one of them had sysout data sets destined for another node, SSS2DUPJ was not set. As of z/OS V1R7, if two jobs in the system match the name specified in SSS2JOBN and one of them has sysout data sets destined for another node, SSS2DUPJ **is** set.
Migration action: Modify any applications that are affected by the change in processing.

Remove the RELADDR= parameter from SPOOLDEF initialization statements (Required-IF, as of R7)

Required if SPOOLDEF RELADDR= is set to NEVER.
z/OS V1R2 (and rolled back to OS/390 V2R10 by an APAR), JES2 added initial support for large SPOOL volumes by allowing you to place a SPOOL data set anywhere on a volume, with the size of the SPOOL data set limited to 64K tracks. Also in z/OS V1R2, the keyword RELADDR= was added to the SPOOLDEF initialization statement to specify relative track addressing. Because older releases did not support relative track addressing, this keyword specified that relative addressing was allowed when a new SPOOL volume was started.

© Copyright IBM Corporation, 2007

In z/OS V1R7, JES2 introduces support for SPOOL volumes greater than 64K tracks. Because of this larger SPOOL volume support and because compatibility with older releases is no longer an issue, the RELADDR= keyword is no longer supported and you should discontinue using it. All new SPOOL volumes started in z/OS V1R7 use relative track addressing.

Migration action: Remove RELADDR= from the SPOOLDEF initialization statement in your z/OS V1R7 initialization deck. If you don't remove it, RELADDR=ALWAYS or RELADDR=ASNEEDED is ignored, but RELADDR=NEVER results in an invalid statement error and the entire SPOOLDEF statement is not processed.

Specify SPOOLDEF RELADDR=ALWAYS on pre-z/OS V1R7 members that share SPOOL with z/OS V1R7 by updating their initialization decks and by issuing the command $TSPOOLDEF,RELADDR=ALWAYS on one pre-z/OS V1R7 MAS member.

z/OS V1R7 JES2 continues to support existing SPOOL volumes that use absolute track addressing. No changes to your existing SPOOL volumes are needed. Only new volumes that are started by a z/OS V1R7 JES2 are affected by removal of support for the RELADDR= keyword.

Update JES2 exit routines because of z/OS V1R7 structure changes (Required-IF, as of R7)

Required if you use any of the affected JES2 exits.

In z/OS V1R7 JES2, some NJE and input service functions were moved from the JES2 address space to application address spaces. This was done to improve performance and reliability as well as reduce system outages. In particular, the new NJE over TCP/IP (NJE/TCP) function and existing internal reader (INTRDR) processing now occur outside the JES2 address space. This can break processing done by the following JES2 exit routines: 2, 3, 4, 7, 8, 13, 20, 36, 37, 39, 40, 46, 47, and 49. The specific changes are:

- Exits 2, 3, 4, 7, 20, 39, 46, and 47 are not called for INTRDRs and NJE/TCP processing.
- Exit 8 is now called for INTRDR and NJE/TCP processing.
- Exits 36 and 37 now get control in a different address space for INTRDR and NJE/TCP processing.
- Exit 13 has been deleted.
- Exit 49 has new function and is being called for $S J commands.
- Exit 40 can now be used to control NJE mail notification (function that was in exit 13).
- The input to exits 2, 3, 4, 20, 46, and 47 has changed significantly.

If you are using any of these changed exits, examine your routines and modify them if necessary.

Migration action: Modify your exit routines to take into account the changes above. See *z/OS JES2 Installation Exits* for details about the functions and interfaces of the changed exits.

Some services have been added to z/OS V1R7 JES2 to help you rework your exit routines to support the new environment. The services are a new USER,ANY environment, local $JCT extensions, $BLDMSG and $SCAN support, table pair enhancements, and multiple exit 0s. These services are discussed in *z/OS JES2 Macros* and *z/OS JES2 Installation Exits*.

Tip: An excellent reference that contains important information is available from this website: http://www.ibm.com/servers/eserver/zseries/zos/installation/zos17_jes2_migration.html. Read the "JES2 z/OS 1.7 Exit Migration Guide". This website also contains other helpful information available from the z/OS JES2 development team to assist you in performing the z/OS R7 JES2 migration actions.

Update automation and procedures because of changed input processing (Required-IF, as of R7)

Required if any of the changes impact your local procedures.

z/OS V1R7 JES2, input processing was rewritten to support NJE over TCP/IP and internal reader changes. If any of your applications are affected by the changed processing, you have to modify them. The changes are:

- **JCL/JECL processing:** In previous releases, error messages for erroneous statements simply stated that the statement was incorrect (for example, HASP112 jobname -- INVALID /*JOBPARM CARD). In z/OS V1R7 JES2, these messages include details about the operand in error (for example, HASP112 jobname -- INVALID /*JOBPARM CARD - VALUE OF CARDS= EXCEEDS 99999999). You might need to update processing that examines error messages.
  In z/OS V1R7 JES2, the parsing of operands has been modified to conform to standard JCL rules. Some JCL that would result in errors in earlier releases now runs correctly.
  For example, SECLABEL='SYSHIGH' is allowed in z/OS V1R7 but is not processed correctly in previous releases. Pre-z/OS V1R7 JES2 code does not support the use of apostrophes with SECLABEL. Because of this, be careful when developing JCL in z/OS V1R7 JES2 that is to be run on previous releases. Test JCL on the target release to ensure that it is processed correctly.

© Copyright IBM Corporation, 2007

Another example is that in releases earlier than z/OS V1R7, JES2 flags the coding of USER=(myuserid) as an error. If JCLERR=YES is specified on a JOBDEF statement, the job fails. If JCLERR=NO is specified, the submitter user ID may propagate and the job may run under the submitter user ID. As of z/OS V1R7 JES2, USER=(myuserid) is honored.

- **SYSIN data processing:** In z/OS V1R7 JES2, the processing of SYSIN data sets has been updated to properly handle carriage control, variable record formats, and SYSIN records up to 32 KB in length. Because of this, job streams that work in z/OS V1R7 might fail in prior releases of JES2. However, jobs submitted on z/OS V1R7 will run correctly on other MAS members, even if the member they run on is pre-z/OS V1R7.
- **Internal reader processing in requestor address space:** In z/OS V1R7 JES2, internal reader processing has been moved from the JES2 address space to the address space that allocated the internal reader. This was done to allow multiple job streams to be processed in parallel and to improve error isolation. However, because the processing is now in the owner address space, CPU and I/O activity that was previously associated with the JES2 address space is now associated with the owning address space. In addition, because the processing occurs in the address space that owns the internal reader, the processing occurs at the dispatching priority of the requestor instead of JES2's priority. Thus, in a busy system, low priority address spaces are not able to submit jobs as fast as higher priority address spaces.
- **Buffering of data in internal readers:** Before z/OS V1R7 JES2, jobs submitted to an internal reader were buffered in 8 KB buffers in a data space. The buffers were not processed until the internal reader was closed, an ENDREQ was issued, or a /*EOF or /*DEL card was processed. In z/OS V1R7 JES2, the cards are processed as soon as they are passed to the internal reader. There is no buffering of data and the internal reader data space no longer exists.
- **Internal reader messages in JOB log:** In z/OS V1R7 JES2, internal reader processing has been moved to the submitting address space. This causes input processing messages (such as $HASP100) to be issued in the submitting address space and to be placed in the JOB log (JESMSGLG) of the submitting address space. This is useful for verifying the status of jobs that are submitted. However, it can increase the size of the JOB log, especially for address spaces that submit a significant number of jobs. If the size of the JOB log becomes a concern, consider using the JESLOG JOB card parameter to spin or suppress the JOB log data set.
- **Internal reader processing when JES2 address space fails:** Before z/OS V1R7 JES2, JES2 discarded any job that was on an internal reader when the JES2 address space abnormally terminated. With z/OS V1R7 JES2, JES2 attempts to retain jobs that are active on an internal reader when JES2 terminates. However, in some cases, JES2 still discards a job because of lost checkpoint updates.

Migration action: Review your automation and procedures, and update them as needed based on the changes discussed above.

Adjust debugging procedures because SVC 111 is no longer used for PUTs to an internal reader (Required-IF, as of R7)

Required if you debug problems using system trace.

In z/OS V1R7, many uses of SVC 111 were converted to stacking Program Call (PC) instructions. As a result, when you debug problems related to internal readers, you must look for the JES2 PC trace entries rather than SVC 111 trace entries to understand what JES2 is doing.

Migration action: Evaluate your debugging procedures for the system trace impacts.

JES2 Migration Actions Post-First IPL

Adjust accounting practices for changed duplicate jobname processing (Required-IF, as of R8)

Required if you are concerned with duplicate job names and processing time is significant to your accounting practices.

SMF record type 26 has a bit that indicates whether a job was held at least once during its life because of duplicate job processing. The byte is SMF26IX2 and the bit is SMF26JDL.

Before z/OS V1R8 JES2, the bit was set only if the job selection process rejected the job because the job duplicated an already running job. Queue delay time used by WLM counted only the time that the job was not selectable. Starting with z/OS V1R8 JES2, jobs are counted as duplicates as soon as they enter the system, regardless of whether the job is rejected during job select processing. This results in the SMF 26 record having the bit set on, indicating that the job was held at least once during its life because of duplicate job name processing. In addition, the time that the job is not selectable because a job by the same name is running is accounted for during the whole life of the job, not just the time that the job was rejected. This in turn can affect WLM algorithms that rely on queue delay computations. The queue delay will be less (because the job is held "longer").

© Copyright IBM Corporation, 2007

The change in behavior could affect accounting or throughput analysis that your installation performs. The WLM algorithm could (favorably) affect the number of initiators that WLM makes active (that is, there could be fewer initiators started).
Migration action: Change accounting or analysis programs, based on your installation's policies, to take into account the change in SMF 26 record data.

Remove the RDINUM= parameter from INTRDR initialization statements (Recommended, as of R7)

Not required, but recommended to avoid confusion in the future.

z/OS V1R7 JES2, internal reader processing was updated to move processing from the JES2 address space to the address space that allocates the internal reader. With this change, the table of internal readers that was maintained in ECSA is no longer required. The data areas that represent the internal reader are now in the private storage of the allocating address space. Because the table of internal readers no longer exists, you no longer need to specify the size of the table. As a result, the RDINUM= parameter, which specified the number of internal readers, has been removed from the INTRDR initialization statement. You should remove the obsolete RDINUM= parameter from your INTRDR initialization statements.
Migration action: Delete the RDINUM= parameter from the INTRDR statement in your initialization stream. If you don't delete it, it is ignored and no error message is generated. Nevertheless, deleting it can help avoid confusion by anyone who encounters it in the future and is unaware that it is no longer functional. Moreover, although the parameter is ignored now, it might not be ignored in a future release.

© Copyright IBM Corporation, 2007

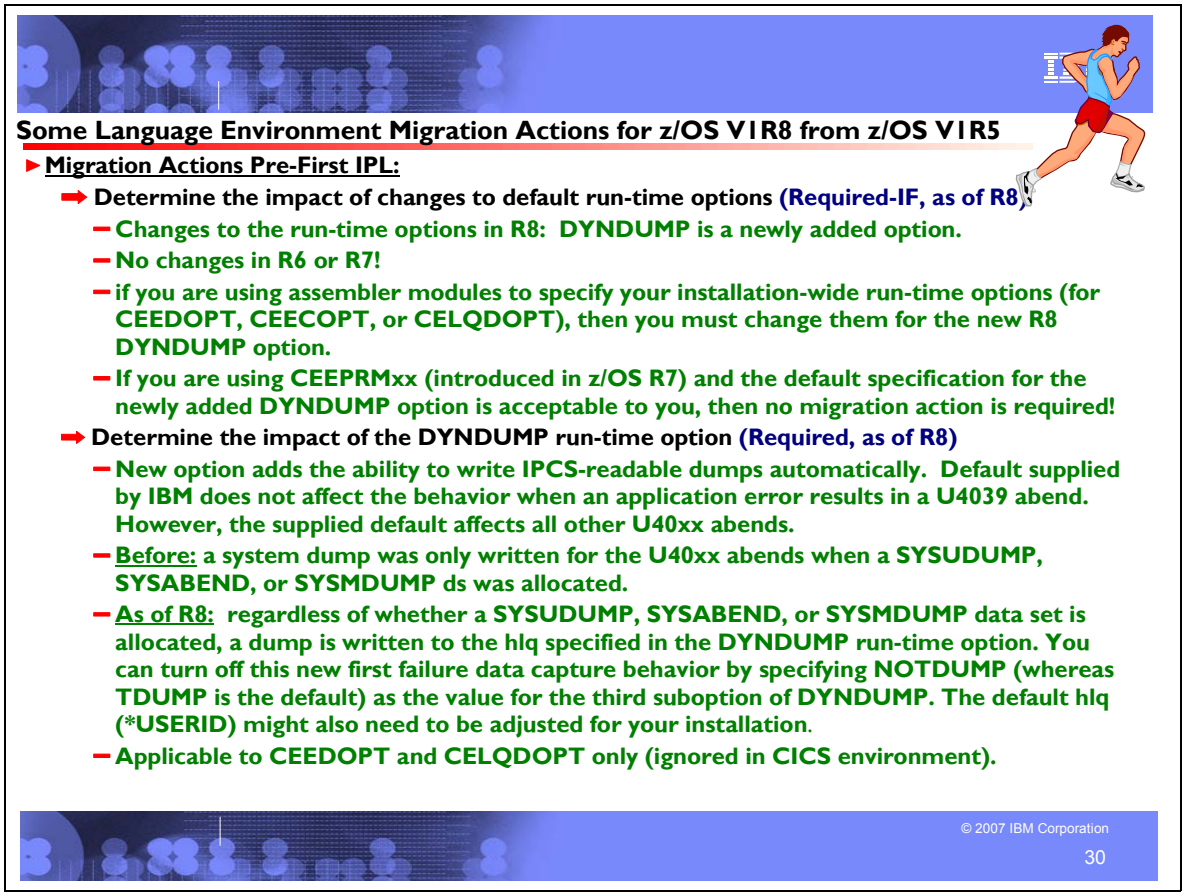

Language Environment Migration Actions Between z/OS V1R5 and z/OS V1R8

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

Language Environment Migration Actions You Can Do Now

Migrate from use of RTLS (Required-IF, as of R6)

Required if you use RTLS or if you have any CEEUOPT or CEEROPT modules that contain RTLS, VERSION, or LIBRARY run-time options.

Effective with z/OS V1R6, Language Environment's use of run-time library services (RTLS) is discontinued. z/OS V1R5 is the last release that included the support. This function has been used primarily in run-time migration. Given the stability and the upward compatibility being provided by the Language Environment run-time library in recent releases of OS/390 and z/OS, this functionality is no longer required.

Migration action: Remove any references that your applications make to the LIBRARY, RTLS, and VERSION run-time options. All references to these run-time options are removed from the CEEDOPT, CEECOPT, CEEROPT, and CEEUOPT modules, and from all options reports. Note that existing CEEUOPT and CEEROPT modules might contain RTLS, VERSION, or LIBRARY options. If so, you can use them unchanged on z/OS V1R6. Informational message CEE3611I is issued to indicate that the run-time options are ignored. To eliminate the message, modify the CEEUOPT or CEEROPT module.

Note: Even if you explicitly specify OFF on the RTLS option (OFF is the default value), you will receive message CEE3611I ("The run-time option *RTLS was an invalid run-time option or is not supported in this release of Language Environment*").

Migrate from use of DLL Rename Utility (Required-IF, as of R6)

Required if you use the DLL Rename Utility.

Effective with z/OS V1R6, IBM has removed the Dynamic Link Library (DLL) Rename Utility, which was in Language Environment. z/OS V1R5 was the last release that included the utility. This utility was used to package and redistribute

IBM-supplied DLLs with applications. Since OS/390 V1R3, the C/C++ DLLs have been licensed with the OS/390 and z/OS base operating systems. Therefore, the DLL Rename Utility is no longer required.

Language Environment Migration Actions Pre-First IPL

Update the CICS CSD based on the newest CEECCSD sample (Required)

Each release, Language Environment adds or deletes load modules in the CICS system definition (CSD) file. Thus, you should update the file each release using the program definitions found in member CEECCSD and, if using CICS Transaction Server (TS) for z/OS V3 (5655-M15), in member CEECCSDX.

For example, in z/OS V1R2, the CLER CICS(R) transaction was added to display run-time options under CICS and alter a subset of options. Therefore, you must define the CLER transaction in the CICS CSD. The CEECCSD sample job in the *hlq*.SCEESAMP data set has been updated to include the CSD entries for the CLER transaction.
Migration action: Update the CSD file using the program definitions in member CEECCSD (and member CEECCSDX if using CICS TS V3) found in the *hlq*.SCEESAMP data set.
Note: The group containing the LE run time routines must be in the group list used during CICS startup.

Review Language Environment load modules in LPA (Required-IF)

Required if you need to make modules accessible through the link pack area (LPA).
Each release you must update the Language Environment load modules that you make accessible through the link pack area (LPA). In addition, each release you should review your list of load modules in the LPA to determine if it's still suitable.
Migration action: Review Language Environment load modules in the LPA.

- To move load modules into the LPA, use sample members are provided in the CEE.SCEESAMP data set.
- To see which modules are eligible for the LPA, refer *to z/OS Language Environment Customization*. The modules listed there can be put in the LPA or extended LPA (ELPA) depending on their RMODE value:
- If you are considering placing the modules listed in z/OS Language Environment Customization in the LPA or the ELPA, then IBM highly recommends that you place the SCEELPA data set in the LPA list (LPALST xx). SCEELPA contains Language Environment load modules that are reentrant, that reside above the 16 MB line, and that are heavily used by z/OS.
- In *z/OS Language Environment Customization* you will also see tables of modules eligible for the LPA and the ELPA above and beyond what is found in the SCEELPA data set. You will need to use the dynamic LPA or MLPA approach to move these modules into the LPA or ELPA. You do not need to include recommended modules if they contain functions your installation does not use. Language Environment modules not listed in these tables can be moved into the LPA or ELPA at your discretion.

Update applications that do not use the standard interfaces to process a utmpx database (Required-IF, as of R6)

Required if you have applications that do not use the standard intefaces to process a utmpx database.
The /etc/utmpx file contains a user accounting database. In z/OS V1R6, the format of the utmpx records is extended to support longer remote host names (ut_host) of up to 1024 bytes. The new format also supports the natural growth of the 64-bit timeval structure, providing a separate ut_tv element for applications compiled as 64-bit. Finally, an internal element, ut_version, replaces space formerly taken up for alignment and is used by the library to distinguish the utmpx record format.
Existing applications (compiled prior to z/OS V1R6) that use the standard interfaces to process a utmpx database will continue to run as before, with the same mappings. The offsets of existing elements in the old format have not changed. Applications will simply not "see" beyond the first 32 bytes of ut_host. (Of course, if you want your applications to reference any of the new elements, you will have to modify and recompile the applications.)

Existing applications that *do not* use the standard interfaces to process a utmpx database must be rewritten for use on a z/OS V1R6 system.
Migration action: Rewrite applications that use nonstandard interfaces to process a utmpx database. For details, see the utmpx.h header shipped with z/OS V1R6.
If possible, convert to use of standard interfaces so that future changes do not impact you. The standard interfaces include the following C functions: setutxent(), getutxent(), getutxid(), getutxline(), pututxline(), and endutxent().

© Copyright IBM Corporation, 2007

Determine the impact of changes to default run-time options (Required-IF, as of R8)

*Required if you are using assembler modules to specify your installation-wide run-time options (CEEDOPT, CEECOPT, or CELQDOPT). If you are using CEEPRM*xx *(introduced in z/OS V1R7) and the default specification for the newly added option is acceptable to you, then no migration action is required.*

Periodically, Language Environment introduces new run-time options, adds new suboptions to existing run-time options, and changes the defaults of run-time options. Additions and changes since z/OS V1R6 are:

- In z/OS V1R8, one Language Environment run-time option was added and none were changed. The new run-time option is DYNDUMP, which controls whether certain dumps are automatically produced.
- In z/OS V1R7, no Language Environment run-time options were added or changed.
  In z/OS V1R7, a new parmlib member, CEEPRM*xx*, was added for Language Environment. You can use it to specify Language Environment run-time options for the system. Operator commands are also provided to allow you to query and update the active run-time options for the system. This simplifies the management of Language Environment options, particularly in multisystem environments, and makes it possible to move Language Environment customization out of assembler language modules maintained using SMP/E usermods. However, specifying Language Environment options using CEEDOPT, CEECOPT and CELQDOPT modules continues to be supported.
- In z/OS V1R6, no Language Environment run-time options were added or changed.

Migration action:

- If you are using CEEPRM*xx* to specify your installation-wide run-time options, and the z/OS V1R8 DYNDUMP option default of TDUMP (which specifies that you want certain dumps) **is** desirable, then no migration action is required.
- If you are using CEEPRM*xx* to specify your installation-wide run-time options, and the z/OS V1R8 DYNDUMP option default of TDUMP (which specifies that you want certain dumps) **is not** desirable, then update your CEEPRM*xx* parmlib member accordingly.
- If you are using assembler modules to specify your installation-wide run-time options, add the new DYNDUMP run-time option to your CEEDOPT, CEECOPT, and CELQDOPT CSECTs. Compare your existing source for the installation-wide run-time options CSECT, CEEDOPT (non-CICS environment), CEECOPT (CICS environment), or CELQDOPT (AMODE 64), with the new samples in *hlq*.SCEESAMP to determine whether you need to change your defaults. If changes are necessary, you must make them and apply the corresponding usermods.

Determine the impact of the DYNDUMP run-time option (Required, as of R8)

Required if you updated the installation-wide run-time option CSECT (CEEDOPT or CEECOPT) in a previous release.

In z/OS V1R8, Language Environment, to assist with first failure data capture, added the ability to write an IPCS readable dump without requiring you to supply a SYSMDUMP DD statement. This activity is controlled by the new DYNDUMP run-time option.

The default supplied by IBM does not affect the behavior when an application error results in a U4039 abend. However, the supplied default affects all other U40xx abends. Before this support, a system dump was only written for the U40xx abends when a SYSUDUMP, SYSABEND, or SYSMDUMP data set was allocated. With this support, regardless of whether a SYSUDUMP, SYSABEND, or SYSMDUMP data set is allocated, a dump is written to the high-level qualifier specified in the DYNDUMP run-time option. You can turn off this new first failure data capture behavior by specifying NOTDUMP as the value for the third suboption of DYNDUMP. The default high-level qualifier might also need to be adjusted for your Installation.

Note: LE's DYNDUMP uses the IATDUMP service in the BCP, and will produce *no* symptoms for those dumps, so they will not be eligible for suppression by DAE

Migration action:

1. Understand the new DYNDUMP run-time option and its default suboption values.
2. Determine whether the default values supplied by IBM are acceptable for your installation, and adjust if needed.
   Note: For U40xx abends (other than U4039), IBM service usually requests an IPCS-readable system dump to assist with problem determination. Therefore, the third suboption of the DYNDUMP run-time option is set to TDUMP by default.

Reference information: For details about specifying the DYNDUMP run-time option, see *z/OS Language Environment Customization*.

Language Environment Migration Actions Post-First IPL <none>

© Copyright IBM Corporation, 2007

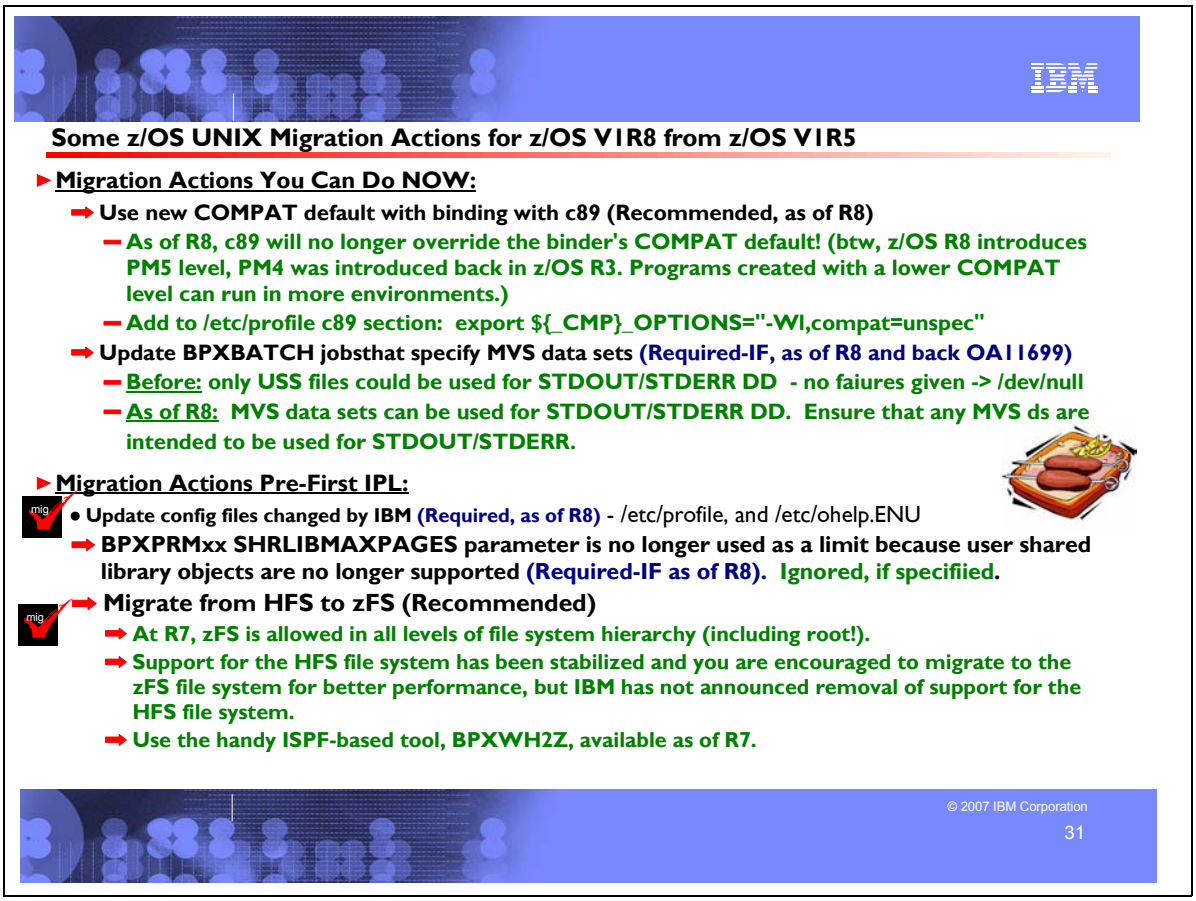

z/OS UNIX System Services Migration Actions Between z/OS V1R5 and z/OS V1R8

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

z/OS UNIX System Services Migration Actions You Can Do Now

Use the new COMPAT default when binding with c89 (Recommended, as of R8)

Not required, but recommended because using the program management binder option COMPAT=MIN allows programs to be executable on more systems than COMPAT=CURRENT does.

The c89 utility can be used to invoke the program management binder. The c89 utility will no longer override the program management binder default value for the COMPAT option. This change affects the following values:

- COMPAT=CURRENT, which specifies that the output is to be as defined for the current level of the program management binder.
- COMPAT=MIN, which specifies that the program management binder should select the minimum level that supports the features actually in use for the current bind. Programs built with COMPAT=MIN can run on any system that has a program management binder that can do the following: produce a program at the PROGRAM OBJECT level that the binder has selected as the minimum. For example, if the program management binder selects PROGRAM OBJECT(FORMAT 4) as the minimum, the programs can run on z/OS V1R3 and all later levels of z/OS.

Before z/OS V1R3, the COMPAT default for both the program management binder and c89 was COMPAT=CURRENT. In z/OS V1R3, the program management binder default was changed to COMPAT=MIN but the c89 default was left unchanged, that is, COMPAT=CURRENT. (The c89 value would override the program management binder value if they were different.) With z/OS V1R8, the c89 COMPAT option specification has been removed, allowing the program management binder default to always be COMPAT=MIN.

You will see some changes in program management binder LIST output in z/OS V1R8. In the PROCESSING OPTIONS section, you will see the COMPAT option with the value UNSPECIFIED, whereas COMPAT=CURRENT with the z/OS V1R8 program management binder would have shown the value PM5. In the SAVE OPERATION SUMMARY section, you will see the PROGRAM TYPE with the level of program management binder object that the binder determines to be the minimum required (perhaps PROGRAM OBJECT(FORMAT 3), whereas COMPAT=CURRENT with the z/OS V1R8 program management binder would have shown PROGRAM OBJECT(FORMAT 5 OS COMPAT LEVEL z/OS V1R8).

The benefit of the new default is that programs created with a lower COMPAT level can run in more environments. You are encouraged to start using the new default on your pre-z/OS V1R8 system.
In order that program objects with the latest format can be viewed and copied on earlier releases, install the following coexistence PTFs if you have not already done so: UA25784 and UA22351 on z/OS V1R7, UA25783 and UA22350 on z/OS V1R6, and UA25782 and UA22349 on z/OS V1R5.
Restrictions: The coexistence support for earlier z/OS releases does not include support for execution of the higher COMPAT level programs.
Migration action: There is nothing for you to do to use the new c89/cc/c++ default on your z/OS V1R8 system. However, to start using it on your pre-z/OS V1R8 system, update the "c89/cc/c++ customization" of /etc/profile, and within the "for" loop add the statement:

```
export ${_CMP}_OPTIONS="-Wl,compat=unspec"
```

This method of customization is chosen because of the way in which the change is implemented. Instead of c89 specifying COMPAT=MIN, c89 picks up the binder default, which is shipped as MIN. This allows c89 to pick up the installation-defined default, which could be different than the IBM-defined default.

Update BPXBATCH jobs that specify MVS data sets (Required-IF, as of R8 and rolled back to R5 with OA11699)

Required if any BPXBATCH jobs specify an MVS data set for the STDOUT or STDERR DD statement.
Before z/OS V1R8 and APAR OA11699, only z/OS UNIX files could be used on the STDOUT and STDERR DD statements for BPXBATCH jobs. If MVS data sets were specified, a failure message was not sent. Instead, the output would go to **/dev/null**. Starting in z/OS V1R8 or with the application of APAR OA11699, the BPXBATCH utility allows for the use of MVS data sets to be specified for BPXBATCH jobs.
With this new support, BPXBATCH attempts to write to an MVS data set specified on STDOUT or STDERR, if it is a supported type and is accessible to the job. If you do not update those BPXBATCH jobs, output may be written to the specified data sets or an unrecoverable I/O error may occur because the data sets are not accessible and the BPXBATCH job could fail. If the MVS data sets are not intended to be written to, you should update those BPXBATCH jobs that specify MVS data sets for the STDOUT or STDERR DDs to not specify STDOUT and/or STDERR Dds.
Migration action: Ensure that all your BPXBATCH jobs that specify MVS data sets for STDOUT and STDERR are data sets that are intended to receive output.

Modify BPXPRMxx because SET OMVS=xxx now executes MOUNT, FILESYSTYPE, and NETWORK statements (Required-IF, as of R7)

*Required if you have MOUNT, FILESYSTYPE, or NETWORK statements in your BPXPRM*xx *parmlib member and you do not want these statements to be executed when the SET OMVS=*xx *operator command is issued.*
Before z/OS V1R7, the SET OMVS=*xx* operator command did not execute the MOUNT, FILESYSTYPE, and NETWORK statements from the BPXPRMxx parmlib member. As of z/OS V1R7, SET OMVS=*xx* executes MOUNT, FILESYSTYPE, and NETWORK statements from BPXPRM*xx*. (If BPXPRM*xx* is used in a sysplex, the MOUNT, FILESYSTYPE, and NETWORK statements are ignored on pre-z/OS V1R7 systems but executed on z/OS V1R7 systems.)
Migration action: If you do not want SET OMVS=*xx* to execute the BPXPRM*xx* MOUNT, FILESYSTYPE, or NETWORK statements, then edit your BPXPRM*xx* parmlib member to remove these statements and copy them into another BPXPRM*xx* parmlib member. If you want to use that new BPXPRM*xx* parmlib member for your next IPL, add it to your OMVS= list in IEASYS*xx*. If you want SET OMVS=*xx* to execute the BPXPRM*xx* MOUNT, FILESYSTYPE, or NETWORK statements, then do nothing.

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

Change NOAUTOMOVE to AUTOMOVE for sysplex-aware file systems (Recommended, as of R6)

Recommended to avoid confusion in the future.
Previously, you could set either NOAUTOMOVE or an automove system list (SYSLIST) for both sysplex-aware and sysplex-unaware file systems. However, they would be ignored for sysplex-aware file systems during dead system recovery. Now you cannot set NOAUTOMOVE or SYSLIST for sysplex-aware file systems. If you do, AUTOMOVE is used instead and you will get an informational message telling you of the change.
What are sysplex-aware file systems? The owner of a file system is the first system that processes the mount. This system always accesses the file system locally, that is, the system does not access the file system through a remote system. Other nonowning systems in the sysplex access the file system either locally or through the remote owning system, depending on the physical file system (PFS) and the mount mode. If a PFS allows a file system to be locally accessed on all systems in a sysplex for a particular mode, the PFS is *sysplex-aware* for that mode. If a PFS requires that

© Copyright IBM Corporation, 2007

a file system be accessed through the remote owning system from all other systems in a sysplex for a particular mode, then the PFS is *sysplex-unaware* for that mode.
To find out if the file system is sysplex-aware, you can use the D OMVS,F command to display each mounted file system. Those file systems for which this system is not the owner, and for which it says Client=N, are sysplex-aware (they are not clients and are locally-mounted to the PFS). Those file systems for which this system is not the owner, and for which it says Client=Y, are sysplex-unaware.
Migration action: For sysplex-aware file systems, replace NOAUTOMOVE in the BPXPRM*xx* parmlib member with AUTOMOVE. Also, when mounting file systems, do not use a SYSLIST. On MOUNT statements in BPXPRM*xx*, we recommend that you use either AUTOMOVE or UNMOUNT for sysplex-aware file systems.

Migrate from central BRLM support (Required-IF, as of R6)

Required if you use central BRLM support.
Prior to z/OS V1R6, central byte range lock manager (BRLM) was the default and distributed BRLM was an option. Now, central BRLM is no longer supported once all systems in a sysplex are at the z/OS V1R6 level. Distributed BRLM is the only supported byte range locking method when all systems in a sysplex are at the z/OS V1R6 level. With distributed BRLM, the lock manager is initialized on every system in the sysplex.
Migration action: If you are already running with a z/OS UNIX couple data set that enables distributed BRLM, there is no change required to activate distributed BRLM for z/OS V1R6. Likewise, if your sysplex only has systems at the z/OS V1R6 level, there is no change required because distributed BRLM is the default. z/OS V1R6 systems ignore the DISTBRLM setting in the couple data set.
If you run the IXCL1DSU utility to create a z/OS UNIX couple data set, distributed BRLM is set up as the default.
To enable distributed BRLM if you are running mixed levels in a sysplex, follow these steps before bringing the first z/OS V1R6 system into the sysplex:

1. Use the IXCL1DSU utility to update the BPXMCDS data set.
   - DISTBRLM NUMBER(0) is no longer supported; NUMBER(1) is the default
   - DISTBRLM NUMBER is ignored in a sysplex that consists of all z/OS V1R6 systems; NUMBER(1) is recommended if you are running mixed levels in a sysplex.
     Example: DEFINEDS SYSPLEX(PLEX1) ...
     DATA TYPE(BPMXCDS) ITEM NAME(DISTBRLM) NUMBER(1) /*default */
2. Activate the BPXMCDS data set.

If you migrate to z/OS V1R6 by running mixed levels in a sysplex, you must enable distributed BRLM before IPLing the z/OS V1R6 system because a z/OS V1R6 system may attempt to activate distributed BRLM when the central BRLM server leaves the sysplex, regardless of the setting in the z/OS UNIX couple data set. The inconsistency between distributed BRLM being active and central BRLM being defined in the couple data set can cause an EC6-BadOmvsCds abend on downlevel systems. It is a notification-only abend indicating that the CDS should be updated. z/OS UNIX will still operate normally, and distributed BRLM will be active in the sysplex.

Update z/OS UNIX configuration files changed by IBM (Required)

Some utilities provided by z/OS UNIX require the use of certain configuration files. You are responsible for providing these files if you expect to use the utilities. IBM provides default configuration files as samples in the /samples directory. Before the first use of any of these utilities, you must copy these IBM-provided samples to the /etc directory (in most cases). You can further customize these files to include installation-dependent information. An example is setting up the /etc/rc file by copying the sample file from /samples/rc to /etc/rc and then customizing it for the installation.
Whenever IBM ships a changed sample configuration file, you must copy the changed sample to its target location and, if you customized the previous version of the sample, incorporate the customization into the new version.
Migration action: For each of the files listed in the table below, copy the IBM-provided sample file to the listed target location. If you added installation-dependent customization to any of the IBM-provided files, make the same changes in the new versions of the files.

© Copyright IBM Corporation, 2007

Utility	IBM-provided sample file	Target location	What changed and when
sh	/samples/profile	/etc/profile	Before z/OS V1R8, exec sh -L was used to rerun the shell in the current address space with the environment variables that were just defined. Starting in z/OS V1R8, exec -a $0 $SHELL is used instead to allow alternate shells to work with the same /etc/profile.
OHELP	/samples/ohelp.ENU	/etc/ohelp.ENU	Every release, bookshelf and book data set names are updated. See "Install books and a bookshelf for the OHELP command" on

Use the changed cp and mv defaults for files with different source and target tags (Required-IF, as of R7, and rolled back to prior releases with OA10314)

Required if you use the ***cp*** *or* ***mv*** *command for files with different source and target tags.*

In z/OS V1R7 and later (and prior releases with the appropriate PTF for APAR OA10314 installed), the **cp** and **mv** commands have new default behaviors when copying or moving files with different source and target tags.

For a UNIX-to-UNIX copy or move of files to a target file system that does not support setting file tags (such as NFS):

- The previous default behaviors (in z/OS V1R6 without PTF UA17740 installed, in z/OS V1R5 without PTF UA17739 installed, and in z/OS V1R4 without PTF UA17738 installed) are:
  - Automatic conversion is disabled for both the source and target files.
  - When the target file has a different tag than the source file, the target file is tagged the same as the source file.
- The new default behaviors (in z/OS V1R7 and later, in z/OS V1R6 with PTF UA17740 installed, in z/OS V1R5 with PTF UA17739 installed, and in z/OS V1R4 with PTF UA17738 installed) are:
  - Automatic conversion is enabled for both the source and target files.
  - Existing target tags are not changed or a new target file is created with a tag according to file system attributes. (The MOUNT parameter can specify TAG.)

For an MVS-to-UNIX copy or move of files to a target file system that does not support setting file tags (such as NFS), if the MVS source file contains text data:

- The previous default behaviors (in z/OS V1R6 without PTF UA17740 installed, in z/OS V1R5 without PTF UA17739 installed, and in z/OS V1R4 without PTF UA17738 installed) are:
  - Automatic conversion is disabled for the target file.
  - An existing tagged target file is untagged.
- The new default behaviors (in z/OS V1R7 and later, in z/OS V1R6 with PTF UA17740 installed, in z/OS V1R5 with PTF UA17739 installed, and in z/OS V1R4 with PTF UA17738 installed) are:
  - Automatic conversion is enabled for the target file.
  - Existing target tags are not changed or a new target file is created with a tag according to file system attributes. (The MOUNT parameter can specify TAG.)

If you use the **cp** or **mv** command in any of the situations described above, the new default behaviors occur. There is no way to return to the prior default behaviors.

Migration action: Notify users of the **cp** and **mv** commands about the new default behaviors.

BPXPRMxx SHRLIBMAXPAGES parameter is no longer used as a limit because user shared library objects are no longer supported (Required-IF as of R8)

The BPXPRMxx **SHRLIBMAXPAGES(nnnnn)** parameter is intended to control the maximum number of pages that can be allocated in the system to contain user shared library modules. This value, in conjunction with MAXSHAREPAGES, was used to control the amount of ESQA consumed by user shared library modules. As of z/OS R8, this parameter is no longer used as a limit because user shared library objects are no longer supported.

Migration action: Remove your SHRLIBMAXPAGES specfication to avoid confusion, because it is no longer suported. If it is specified, it will be ignored.

© Copyright IBM Corporation, 2007

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

Migrate from HFS file systems to zFS file systems (Recommended)

Not required, recommended. Support for the HFS file system has been stabilized and you are encouraged to migrate to the zFS file system for better performance, but IBM has not announced removal of support for the HFS file system.

Before z/OS V1R7, the HFS file system was the primary hierarchical file system. As of z/OS V1R7, you can use any combination of HFS and zFS file systems. Because zFS has higher performance characteristics than HFS and is the strategic file system, HFS might no longer be supported in future releases and you will have to migrate the remaining HFS file systems to zFS.

The HFS and zFS file system types in mount statements and command operands are now generic file system types that can mean either HFS or zFS. Based on the data set type, the system will determine which is appropriate.

Migration action: Before beginning the migration, follow these steps:

1. Establish backout procedures.
2. Decide on naming conventions.
3. Decide on unavailability.
4. Perform the conversion from HFS to zFS file system.

   Tip: You can use the BPXWH2Z tool to perform the conversion. It is an ISPF-based tool that migrates HFS file systems to zFS file systems. It has a panel interface that enables you to alter the space allocation, placement, SMS classes, and data set names. A HELP panel is provided. With this tool, you can:

   - Migrate HFS file systems (both mounted and unmounted) to zFS file systems. If the HFS being migrated is mounted, the tool automatically unmounts it and then mounts the new zFS file system on its current mount point.
   - Define zFS aggregates by default to be approximately the same size as the HFS. The new allocation size can also be increased or decreased.
   - Have the migration run in TSO foreground or UNIX background.

   BPXWH2Z is located in partitioned data set SBPXEXEC (usually called SYS1.SBPXEXEC). When you run the BPXWH2Z tool on your z/OS V1R7 system, it uses the z/OS V1R7 level of the **pax** command. This level has been enhanced for sparse file support and other characteristics that are of concern when migrating from an HFS to zFS file system. You can manually migrate from an HFS to zFS file system without using the tool. When doing that, the z/OS V1R7 level of the **pax** command is recommended. However, you would need to allocate and format the target zFS file systems.
5. Change policies and scripts, and so forth, to reflect the change from the HFS file system to zFS file system.

   Tip: Make sure you're using the RMF Monitor III option to report on zFS activity.

The following is sample JCL to run the BPXWH2Z conversion tool as an ISPF batch job. Before you run the job, be sure to read the Notes section:

```
//*************************************************************
//*** Licensed Materials - Property of IBM *
//*** 5694-A01 *
//*** (C) Copyright IBM Corp. 2002, 2006 *
//*** *
//*** This JCL will run the BPXWH2Z(HFS to zFS migration) *
//*** conversion tool as an ISPF Batch Job. *
//*** CAUTION: This is neither a JCL procedure nor a complete *
//*** job. Before using this job step, please inspect it *
//*** carefully for any changes that are necessary to make it *
//*** fit for your installation's needs. *
//*** *
//*** Notes: *
//*** *
//*** 1. Add the job parameters to meet your system *
//*** requirements. *
//*** *
//*** 2. Run only the steps that are applicable to your *
//*** installation. *
//*** *
//*** 3. This job will complete with a completion code 0 if *
```

© Copyright IBM Corporation, 2007

```
*** successful, or with a completion code 8 if it *
*** encounters any error while running the tool for *
*** single HFS data set migration. *
*** *
*** 4. For multiple HFS data set migrations, this job will *
*** complete with a completion code 0 if successful, or *
*** with a completion code 8 if no migrations were *
*** successful, or with a completion code 4 if one or *
*** more migration is successful, but atleast one failed.*
*** *
*** 5. Check the job output for execution information *
*** including error details. *
*** *
*** 6. The zFS data sets cannot be pre-allocated when ISPF *
*** batch job method is used for BPXWH2Z conversion tool.*
*** *
*** 7. SYSEXEC DD specifies the BPXWH2Z REXX EXEC location. *
*** *
*** 8. SYSTSPRT DD specifies the output log. *
*** *
*** 9. SYSTSIN DD specifies the ISPSTART cmd with parms. *
*** *
************************************************************
//ISPBTCH EXEC PGM=IKJEFT01,
// DYNAMNBR=40,TIME=1440
//SYSTSIN DD *
ISPSTART CMD(BPXWH2Z -VC /HFS/ZFS/ OMVS.MYHFS)
//SYSTSPRT DD SYSOUT=*
//ISPLIST DD SYSOUT=A,DCB=(LRECL=121,BLKSIZE=1210,RECFM=FBA)
//ISPLOG DD SYSOUT=A,DCB=(LRECL=125,BLKSIZE=129,RECFM=VA)
//SYSEXEC DD DSN=POSIX.AHILAN.TAR,DISP=SHR,VOL=SER=BPXLK1,
// UNIT=3390
//ISPPLIB DD DSN=POSIXBLD.PRIMARY.PANELS,DISP=SHR
// DD DSN=ISP.SISPPENU,DISP=SHR
// DD DSN=ISP.SISPPENU,DISP=SHR
//ISPMLIB DD DSN=ISP.SISPMENU,DISP=SHR
//ISPSLIB DD DSN=ISP.SISPSENU,DISP=SHR
//ISPTLIB DD DSN=SYS1.SBPXTENU,DISP=SHR
// DD DSN=ISP.SISPTENU,DISP=SHR
//ISPPROF DD DISP=NEW,UNIT=SYSVIO,SPACE=(CYL,(1,1,5)),
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=3120)
```

Install books and a bookshelf for the OHELP command (Required-IF)

Required if you use the OHELP command.

The TSO/E OHELP command uses the BookManager READ base element to display online reference information about shell commands, TSO/E commands, C functions, callable services, and messages issued by the shell and dbx. Before this online help facility can be used, you must set it up by installing books and a bookshelf.

Migration action: Copy the IBM-supplied sample file /samples/ohelp.ENU to /etc/ohelp.ENU. If you choose not to use the IBM-supplied default file, you can create your own file called /etc/ohelp.*language_i*d, where language_id is the TSO/E user's primary language code.

© Copyright IBM Corporation, 2007

Migration Actions Not to Overlook!

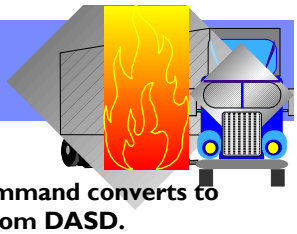

- DFSMShsm copy pool R8-format coexistence- R8 FRBACKUP command converts to R8-format, then, only command on pre-R8 system is FRRECOV from DASD.
- z/OS R7 DFSMS JOBCAT/STEPCAT support removed.
  - Use the MODIFY CATALOG,ENABLE(JOBSTEPCAT) on R5 to turn support back on (for now!)
- z/OS R7 JES2 Z2 - must be in Z2 mode!
- z/OS R7 JES2 Exit and Input Processing changes
  - Internal reader processing moved from JES2 address space to address space that allocated the internal reader.
- z/OS R7/R8 BCP 1-byte Console ID support removal
  - Use the Console ID Tracking Facility, introduced in z/OS R4 Consoles Enhancement feature with the latest exclusion list now.
  - Final removal of 1-byte Console IDs in z/OS R8
- z/OS R7 XL C/C++ does not ship the OS/390 R10 Compilers
  - Move to the ISO C/C++ compilers now, may still use OS/390 R10 compilers from prior system.
- z/OS R7 HCD IODF V5 Coexistence considerations with lower systems
- z/OS R6 Architecture Level Set - must be z/Architecture mode
  - Use the z/Architecture checklist provided.

© 2007 IBM Corporation

Future Migration Actions Not to Overlook

- **Remove your dependency on the C/C++ IBM Open Class (IOC) Dynamic Link Libraries (DLLs)**
  - Use the Standard C++ Library instead
  - Removed in the release after z/OS R8
- **Run the TN3270E Telnet server as a separate address address space**
  - Must do this before the release after z/OS R8
- **z/OS DFSMS Imbed, Replicate, Keyrange removal in a future release.**
  - Use the IMBDSHIP tool, and install PTF for APAR OA10952 now!
- **Start your conversion from HFS to zFS after you migration to z/OS R8**
  - Use z/OS R7's tool, BPXWH2Z to help.

© 2007 IBM Corporation

© Copyright IBM Corporation, 2007

Summary: *Plan your migration to z/OS V1.8!*

- **z/OS Coexistence- z/OS V1R5 through z/OS V1R8 supported.**
  - z/OS V1R7 still orderable in SystemPac, if you need it
- **Removed Functions since z/OS V1R6**
  - JES2 R4 mode, JOBCAT/STEPCAT, OS/390 R10 C/C++ Compilers, 1-byte console IDs, zFS multi-file system aggregates from share fs, Firewall Technologies,
  - Future: TN3270 in TCP/IP stack, run-time support for C/C++ IBM Open Class DLLs, old VSAM attributes, zFS multi-file system aggregates.
- **Coexistence System Requirements** - use the EPSPT tool to help!
- **IBM Migration Checker for z/OS** can assist you!
- **Migration Actions for z/OS V1R8**:
  - **z/OS General** - enough XCF groups, new and old data sets, new address spaces, ...
  - **BCP** - 1-byte console id removal, master console removal, 4TB real changes, ...
  - **XL C/C++** - MEMLIMIT at least 3000MB for IPA link step
  - **Communications Server** - obsolete statements from TCP/IP profile, UDP 514 available,...
  - **HCD** - IODF V5 coexistence consideration
  - **DFSMS** - new security profiles for hsm fr commands, FRBACKUP command changes, ...
  - **DFS zFS** - use compatibility-mode aggregates, don't use STOP ZFS, ...
  - **JES2** - exit changes, changes in input processing, ...
  - **Language Environment** - new DYNDUMP run-time option
  - **z/OS UNIX** - COMPAT option for c89, HFS to zFS conversion, ...

➡ **Review the z/OS Migration book, for a complete list of z/OS V1R8 migration actions!**

© Copyright IBM Corporation, 2007