Session 14254

Common z/OS Problems You Can Avoid

(or explained)

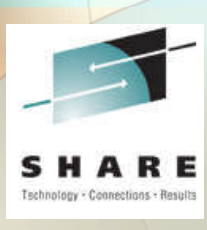

SHARE Boston, MA

August 15th, 2013

Jerry Ng **jerryng@us.ibm.com**
Patty Little **plittle@us.ibm.com**

IBM Poughkeepsie

Trademarks

- **The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**
  - **MVS**
  - **OS/390®**
  - **z/Architecture®**
  - **z/OS®**

* Corporation Registered trademarks of IBM

Contents

- A Reminder: IBM Health Checker for z/OS
- CDS Inconsistency
- OMVS Services failing
- ICSF/Crypto Master Keys
- Logger CF Structure
- PFA INI JAVAPATH
- FTP'ing problem documentation
- RASP using AUX slots
- RSU in IEASYSxx
- PROGxx REFRPROT
- SDUMP AUXMGMT issues

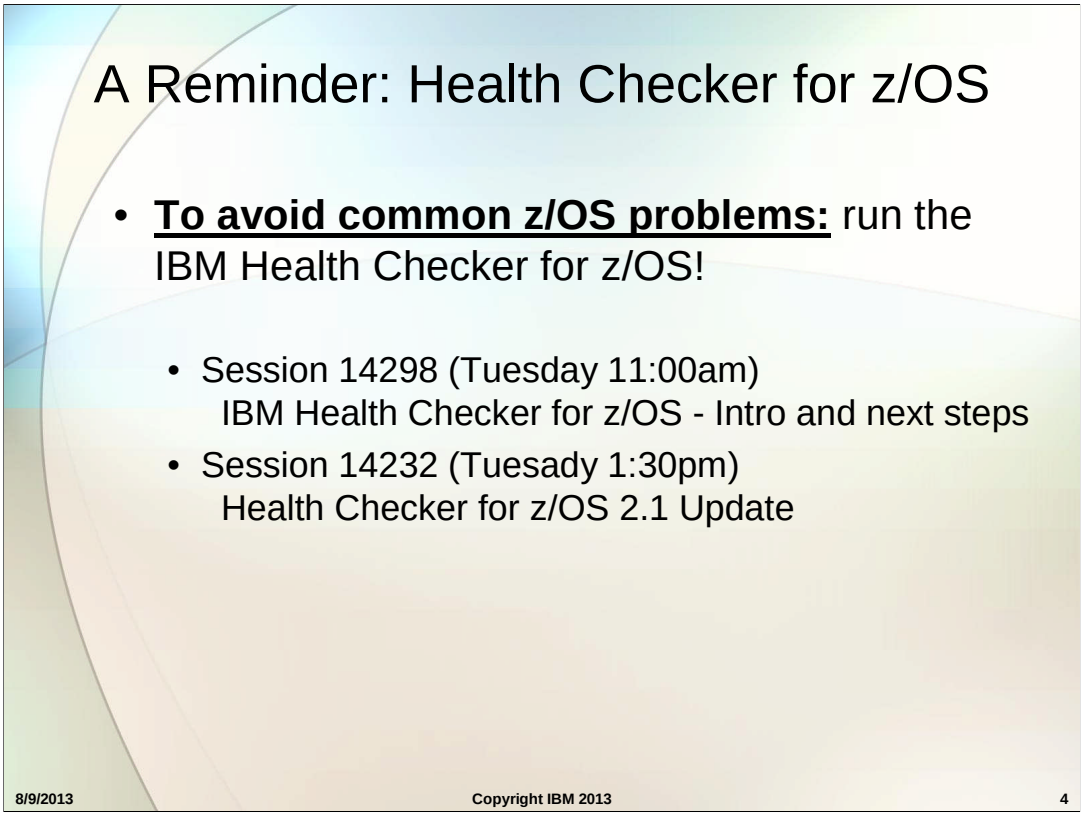

IBM Health Checker for z/OS is a component of MVS that identifies potential problems before they impact your availability or, in worst cases, cause outages. It checks the current active z/OS and sysplex settings and definitions for a system and compares the values to those suggested by IBM or defined by you. It is not meant to be a diagnostic or monitoring tool, but rather a continuously running preventative that finds deviations from best practices. IBM Health Checker for z/OS produces output in the form of detailed messages to let you know of both potential problems and suggested actions to take.

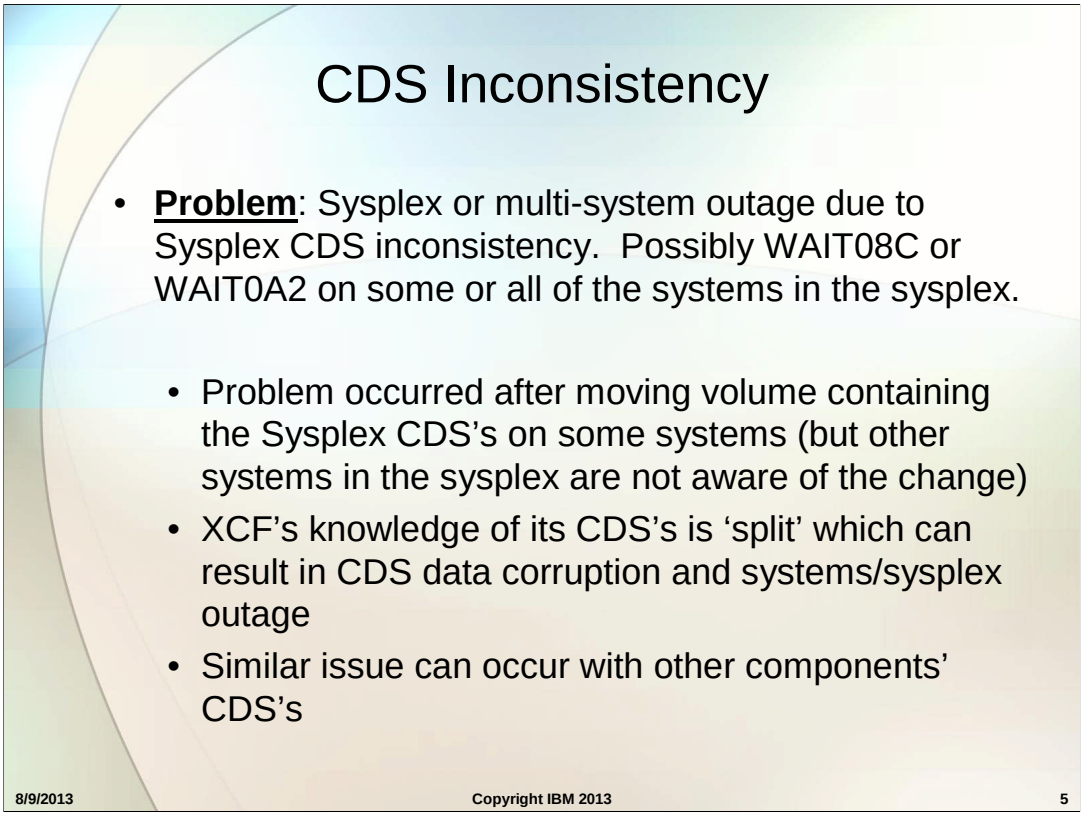

Some products such as zDMF migrate system or sysplex datasets. Extreme caution must be used when using such products against CDS's. If some systems in the sysplex start using the 'new' physical CDS's while other systems continue to write to the 'old' physical CDS's then data corruption or CDS inconsistencies may result.

The impact of 'splitting' XCF's knowledge of the CDS's varies depending on what updates are done during the timeframe of the split. Wait states on some or all the systems usually occur.

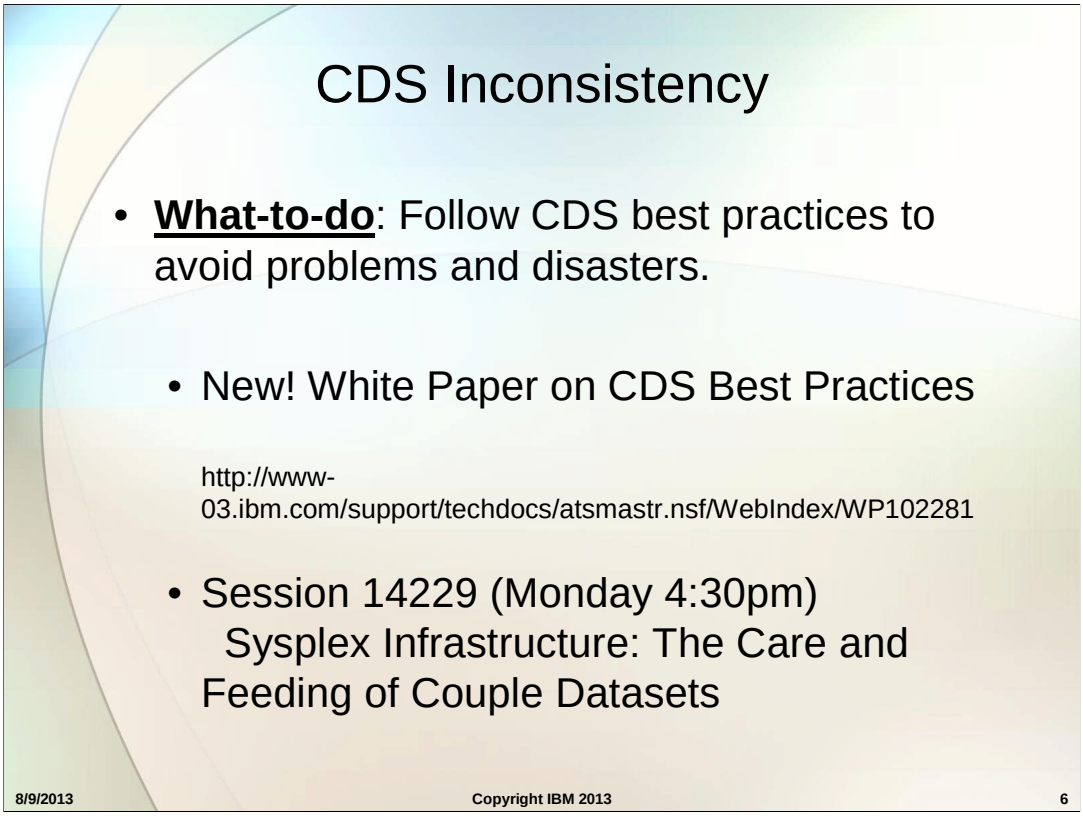

Please review the above paper and session material to follow the best practices for CDS's. This will help to avoid disasters in your systems.

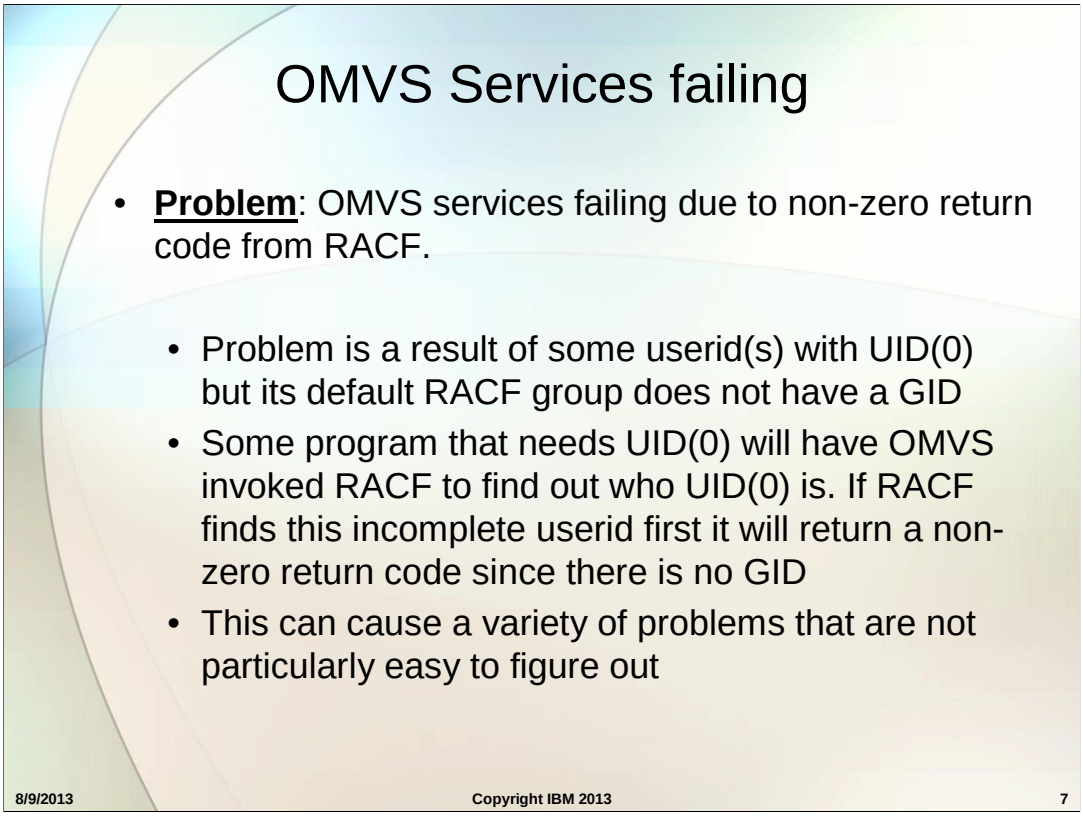

In most z/OS systems, there are always multiple userids with UID(0).

.

If a UID(0) userid is created or altered so that their RACF Default group does not have a GID, then the RACF commands will get a message like:

ICH21035I USER XXXX IS ASSIGNED AN OMVS UID, BUT DEFAULT GROUP XXXXGRP DOES NOT HAVE A GID. PROCESSING CONTINUES.

What can happen is that some program that needs UID(0) will at some point have OMVS asks RACF who UID(0) is. If this incomplete userid happens to be found first by RACF, then RACF sees that the Default Group has no GID and gives back a non-zero return code because the userid's OMVS definition is incomplete.

This can cause a variety of problems that are not particularly easy to figure out.

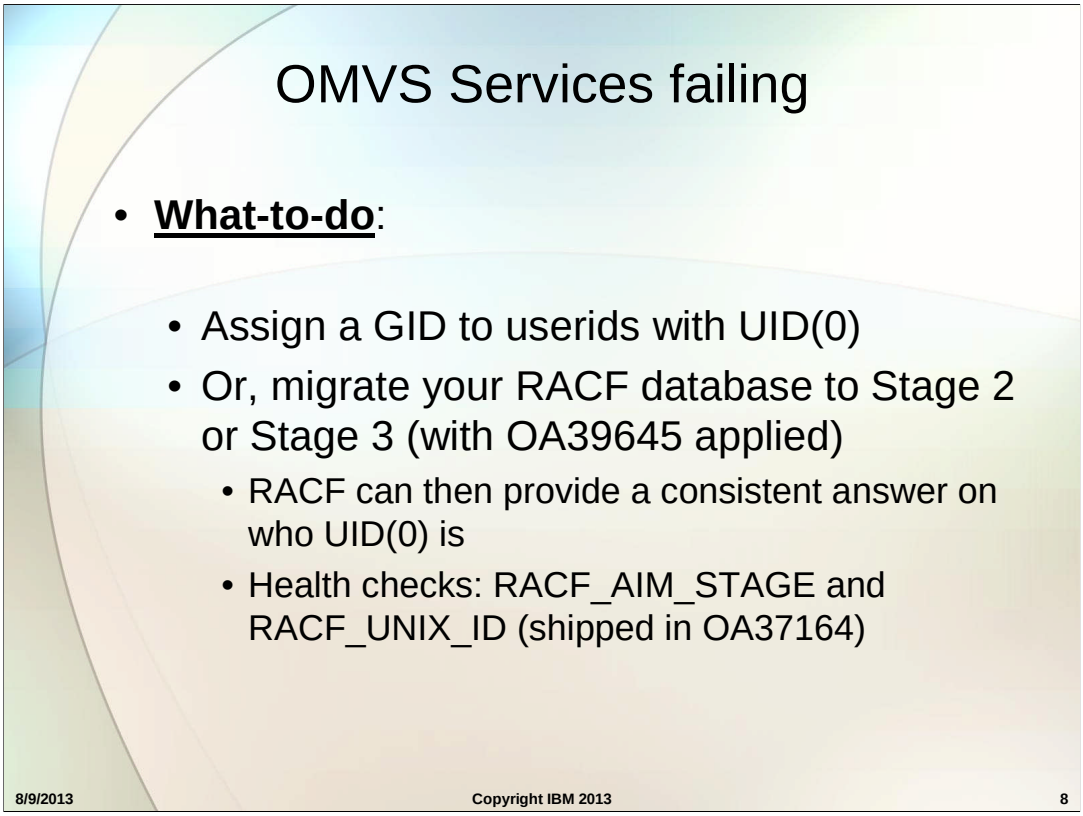

In RACF, the database has a "stage" level, 0 thru 3. You move it via the IRRIRA00 utility.

Once at Stage 2 or Stage 3 (with APAR OA39645), RACF can guarantee that any request to translate a UID or GID will result in the same answer any time. Customers who are not in those environment are not guaranteed a consistent answer.

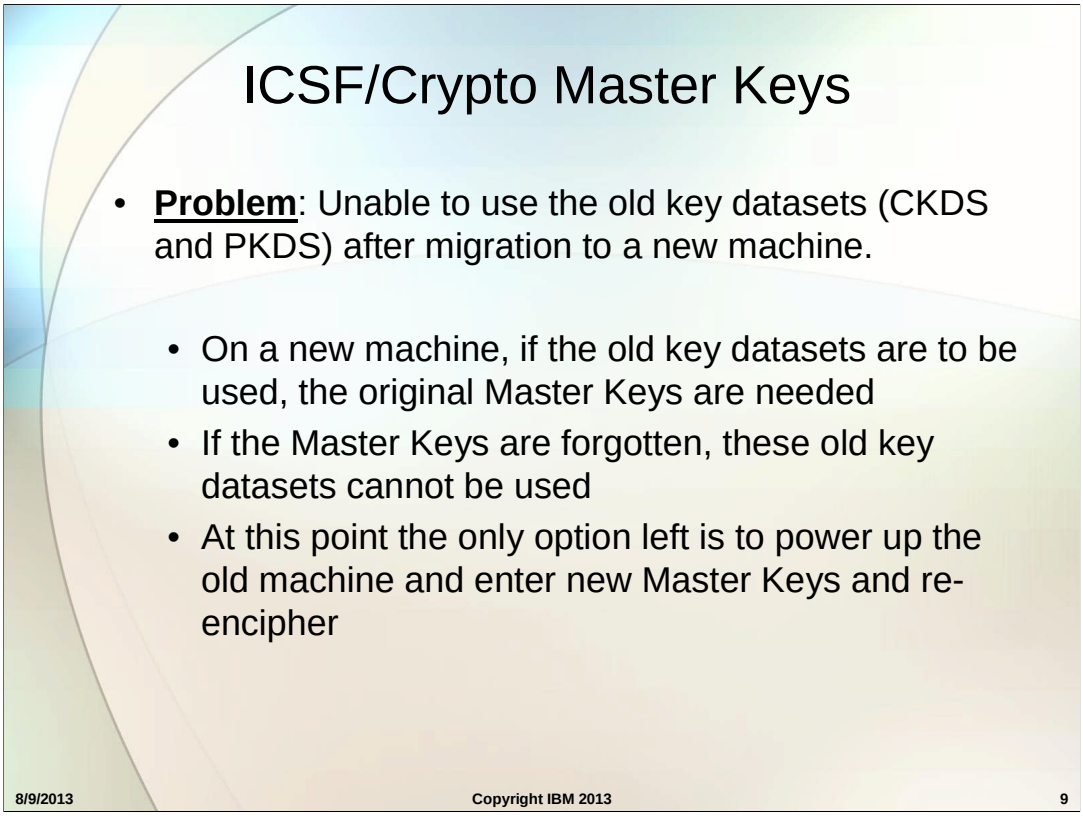

When customers using ICSF migrate to a new mainframe, they usually want to continue to use their existing Key Data Sets (ie, CKDS and PKDS). To use these datasets on the new box, you need to enter the correct Master Keys, DES & AES for the CKDS, and RSA & ECC for the PKDS. If the master keys are forgotten then these Key Data Sets can't be used and all of the keys in those Key Data Sets can't be used.

If the Master Keys are forgotten/lost the only option is to power up the old box and enter new Master Keys and re-encipher.

ICSF/Crypto Master Keys

- **What-to-do**:
  - Remember the Master Keys (or find them) prior to the migration to a new machine
  - If the Master Keys are forgotten, you need to enter new Master Keys and re-encipher on the old machine first

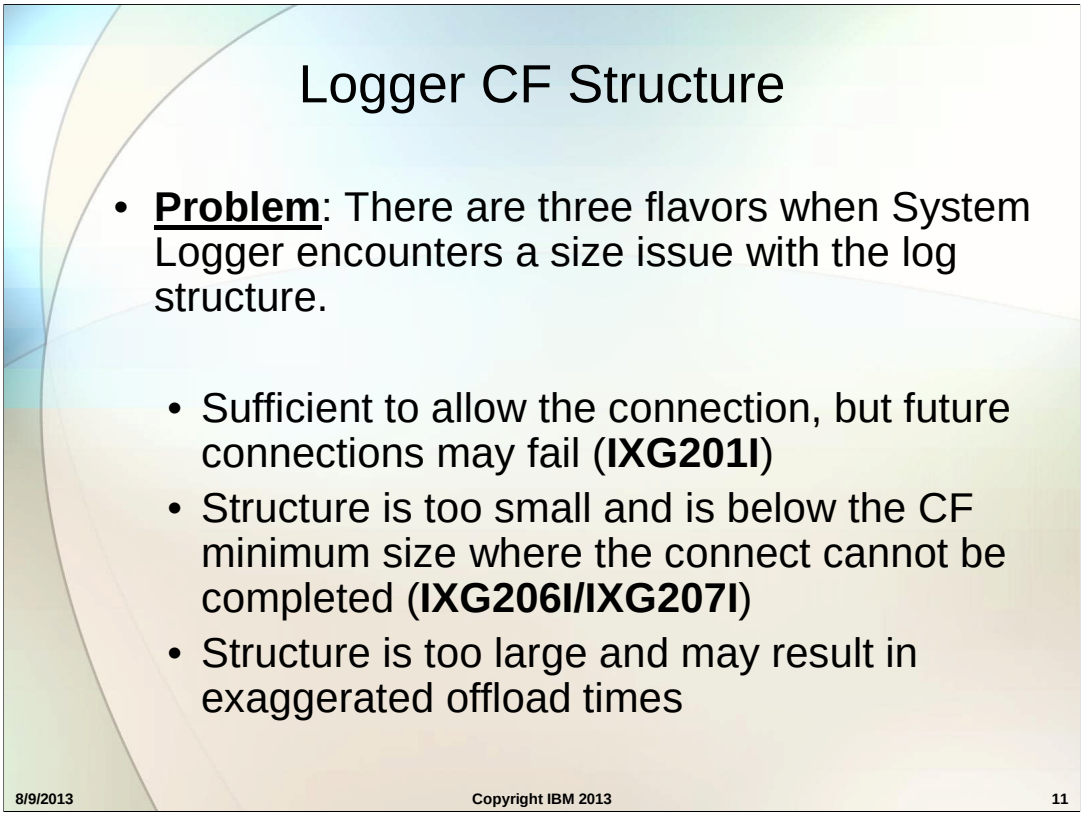

Structures which are too large may lead to various efficiency issues... Allocation or Offload may take too long to complete, backing up other work.

Logger CF Structure

- **IXG201I** REQUEST TO CONNECT TO LOGSTREAM DFHLGR32.SNEAPRMS.DFHJ00 IN STRUCTURE DFHLGR32 ACCEPTED. CONNECTION TO ADDITIONAL LOGSTREAMS MAY FAIL DUE TO INSUFFICIENT STRUCTURE STORAGE
- **IXG206I** CONNECT FAILED FOR LOGSTREAM HLM.MESSAGE.LOG IN STRUCTURE HLM_LOG. NO SUITABLE COUPLING FACILITY FOUND.
- **IXG207I** CF NAME: ST132 REASON CODE: 00000007 CF MINIMUM SIZE: 9216K BYTES.

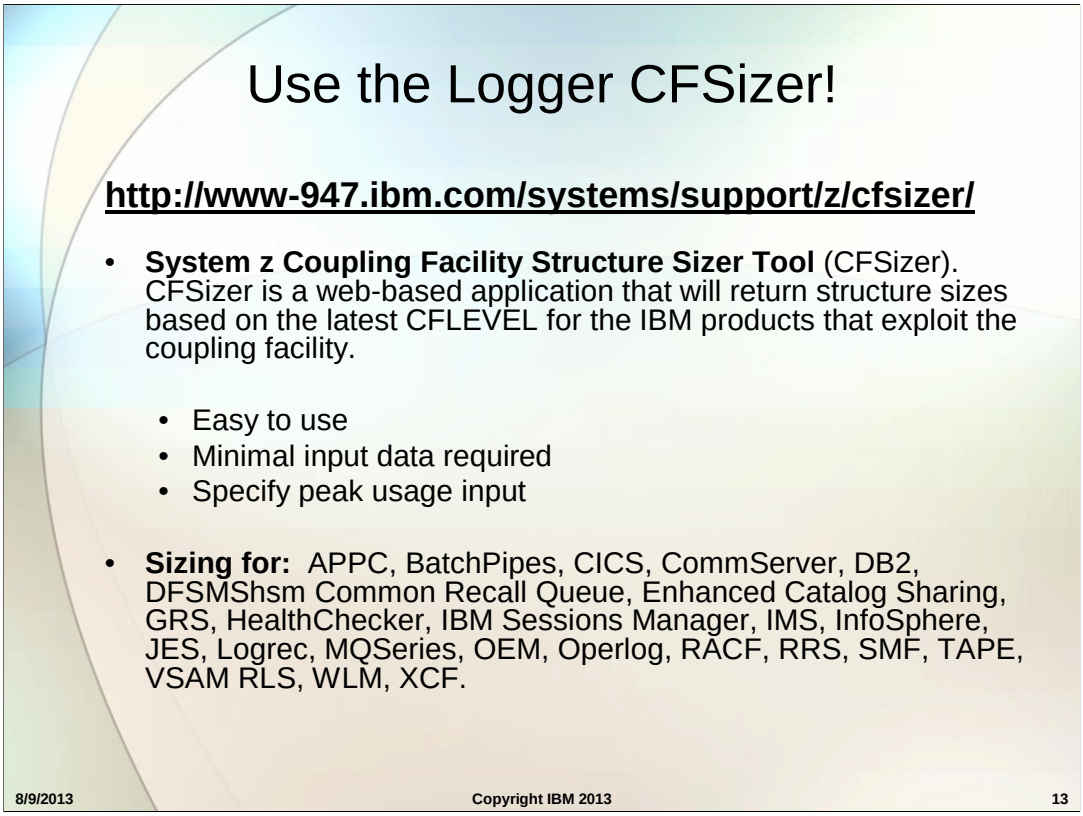

The purpose of the application is to provide an easy-to-use interface that will calculate the structure sizes for you based on some minimum input data that you provide that represents your expected usage of the structure-owning product. The inputs you supply should correspond to your expected peak usage of the product. It is generally good practice to slightly overspecify your peak values, to produce a sizing recommendation slightly larger than absolutely necessary. This will provide some room for growth, and help avoid failures caused by insufficient structure sizes.

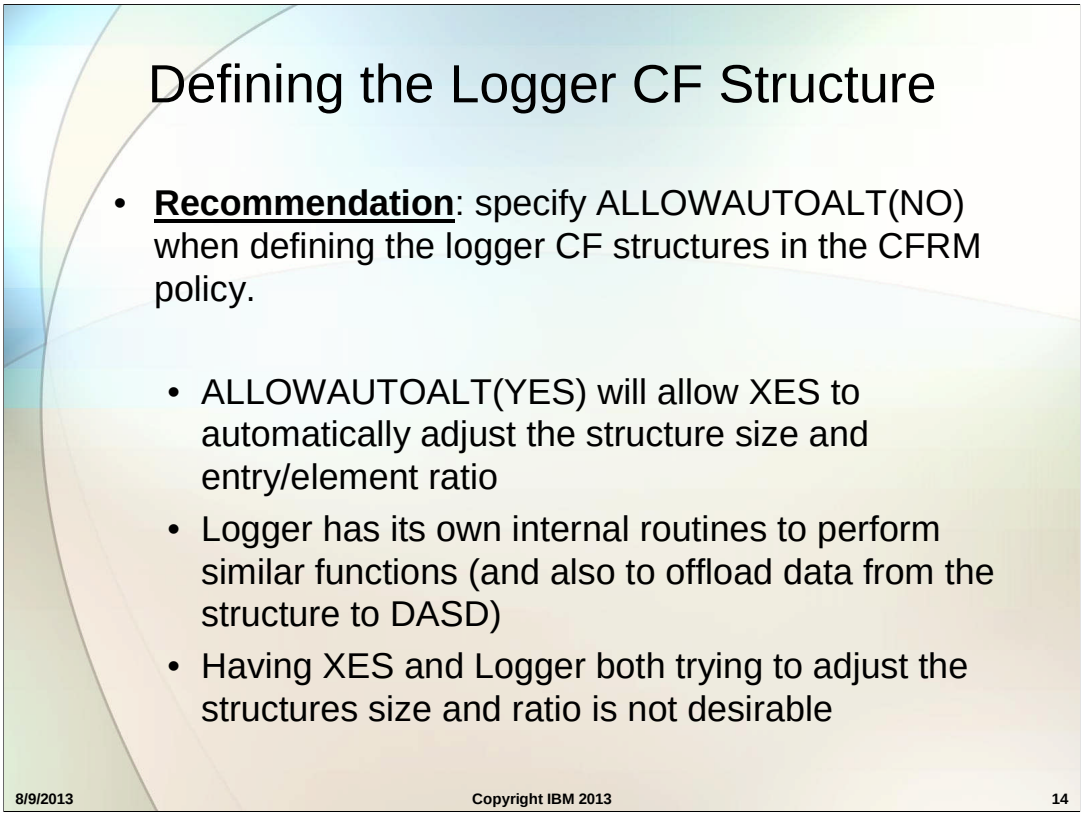

Defining a log stream structure with ALLOWAUTOALT(YES) can lead to inefficient and unexpected results. Basically, ALLOWAUTOALT is a XES parameter that allows the structure size and ratio to be adjusted by XES. However, Logger internally will manage and adjust the structure's ratio of entries and elements, as well as perform offloads to move the data from the structure to DASD (then deleting that data

from the structure). Having both XES *and* Logger trying to adjust and manage the structure is not desirable.

PFA INI JAVAPATH

Problem: Error messages issued at PFA modeling time

- **AIR022I** REQUEST TO INVOKE MODELING FAILED FOR CHECK NAME= *PFA_LOGREC_ARRIVAL_RATE* UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002
- **AIR033I** PFA has detected that SMF is not running and has stopped processing the PFA_SMF_ARRIVAL_RATE check. Processing will resume after SMF restarts

Potential error messages that may appear at PFA modeling times, depending upon which checks PFA is running

AIR022I REQUEST TO INVOKE MODELING FAILED FOR CHECK NAME= PFA_LOGREC_ARRIVAL_RATE UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002
AIR022I REQUEST TO INVOKE MODELING FAILED FOR CHECK NAME= PFA_MESSAGE_ARRIVAL_RATE UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002
AIR033I PFA has detected that SMF is not running and has stopped processing the PFA_SMF_ARRIVAL_RATE check. Processing will resume after SMF restarts.
AIR022I REQUEST TO INVOKE MODELING FAILED FOR CHECK NAME= PFA_ENQUEUE_REQUEST_RATE UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002

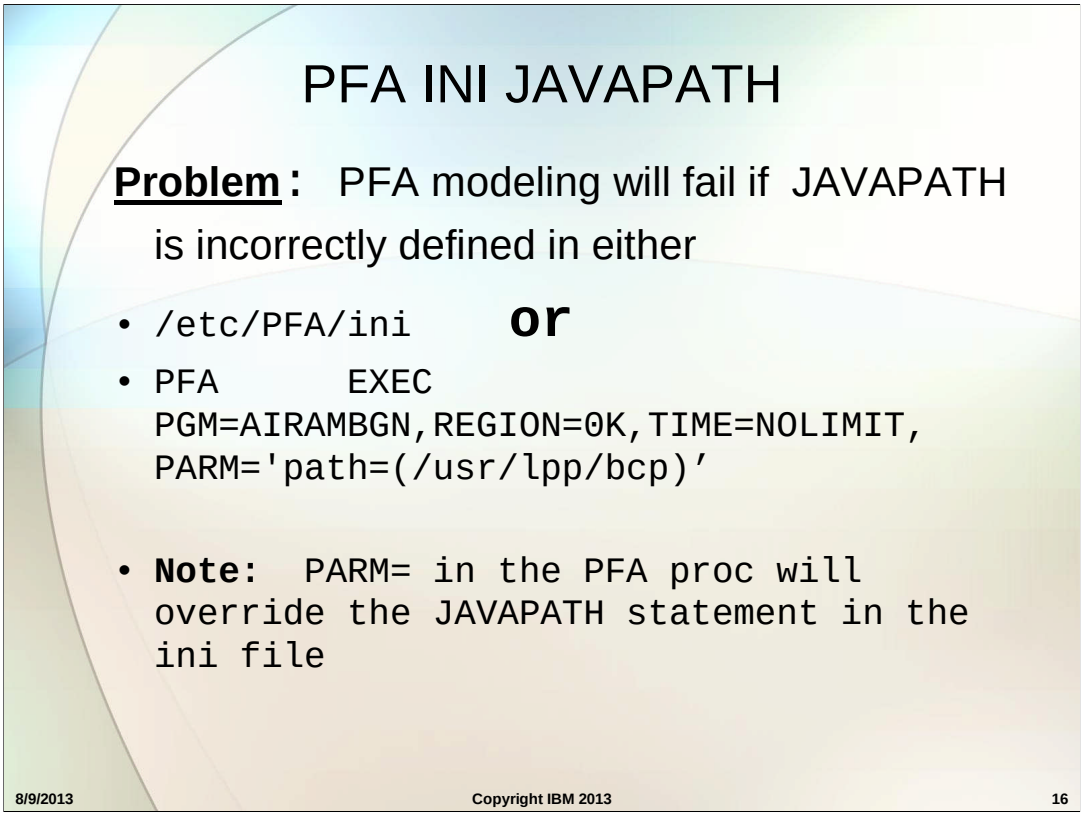

Potential error messages that may appear at PFA modeling times, depending upon which checks PFA is running

AIR022I REQUEST TO INVOKE MODELING FAILED FOR
CHECK NAME= PFA_LOGREC_ARRIVAL_RATE
UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002
AIR022I REQUEST TO INVOKE MODELING FAILED FOR
CHECK NAME= PFA_MESSAGE_ARRIVAL_RATE
UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002
AIR033I PFA has detected that SMF is not running and has stopped processing the PFA_SMF_ARRIVAL_RATE check. Processing will resume after SMF restarts.
AIR022I REQUEST TO INVOKE MODELING FAILED FOR
CHECK NAME= PFA_ENQUEUE_REQUEST_RATE
UNIX SIGNAL RECEIVED= 00000000 EXIT VALUE= 00000002

PFA INI JAVAPATH

Problem: Example of ini file:

```
/*This file customized 14Sep2011 09:19:22 by Serverpac Job */
   VERSION=01010101
/* NLSPATH  = path to NLS files                              */
/* LIBPATH  = path to JNI library using libpath              */
/* JAVAPATH = path to JAVA code used for PFA                 */
PATH= /usr/lpp/java/J5.0/bin/classic:/usr/lpp/java/J5.0/bin
NLSPATH= /usr/lpp/nls/msg/%L/%n:/usr/lib/msg/%L/%n.catxlc/bin
LIBPATH=/usr/lpp/java/J5.0/bin:/usr/lpp/java/J5.0/bin/
        classic:/lib:/usr/lib:
LANG= C
JAVAPATH= /usr/lpp/bcp
```

PFA INI JAVAPATH

Explanation :

- The JAVAPATH statement identifies the location of where PFA's Java code used for modeling resides
- It does NOT represent where JAVA 6.0 code resides

PFA INI JAVAPATH

Solution:

- Check the ini file to ensure that it does not point to Java 6.0 code, but rather PFA's Java modeling code

- Check the PARM= value in the PFA PROCLIB EXEC statement to ensure it does not point to Java 6.0 code, but rather PFA's Java modeling code

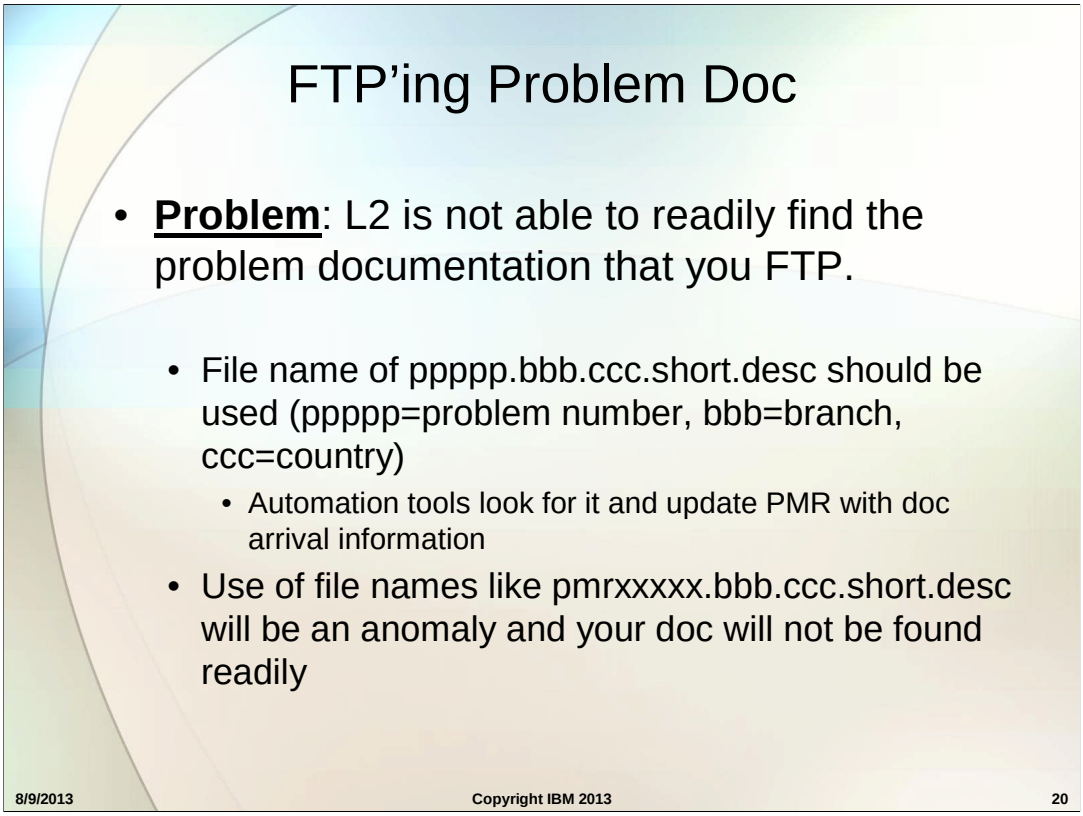

Remember to use the recommended file name of ppppp.bbb.ccc.short.desc. It will be found readily by the automation tools and will be processed right away.

FTP'ing Problem Doc

- **What-to-do**: Use the recommended file name of ppppp.bbb.ccc.short.desc
  - Example (doc for MVS L2)

```
Please send your documentation using the z/OS Problem Documentation
 Upload Utility (MTFTPS prior to R13). Place the files in directory
 /toibm/mvs                    on the geographically closest server:
    Americas: testcase.boulder.ibm.com (or 170.225.15.31)
    Europe    ftp.ecurep.ibm.com (or 192.109.81.7)
    AP        ftp.ap.ecurep.ibm.com (or 210.143.141.69)
OR 1.Compress your dataset using AMATERSE (TRSMAIN replacement).
   2.FTP to the server and directory above using userid:anonymous
   3.Specify BINary mode for transfer of the dataset.
   4.PUT the file using your PMR number as the start of the file name
      ppppp.bbb.ccc.short.desc[.TRS]
Small files ( <2Gb ) can be sent as an attachment through SR.
For more information and FAQ's on transferring documentation to IBM
 see url http://www-05.ibm.com/de/support/ecurep/index.html
```

RASP using AUX Slots

- **Problem**: Growth in Auxiliary paging slots owned by the RSM Address Space (RASP ASID 3)
  - IRA206I showing RASP with relatively small real frames, but very large amount of AUX slots

```
*IRA200E AUXILIARY STORAGE SHORTAGE
*IRA206I JMONDB2A ASID 0089 FRAMES 0002307735 SLOTS 0000555697 % OF AUX  13.2
*IRA206I DBNGDBM1 ASID 01E1 FRAMES 0000592191 SLOTS 0000323722 % OF AUX   7.7
*IRA206I DBNGDBM1 ASID 00D7 FRAMES 0000139593 SLOTS 0000219606 % OF AUX   5.2
*IRA206I DBNGDBM1 ASID 01A0 FRAMES 0000111173 SLOTS 0000200963 % OF AUX   4.7
*IRA206I RASP     ASID 0003 FRAMES 0000000468 SLOTS 0000192729 % OF AUX   4.5
```


When you are in an auxiliary storage shortage condition, you will receive messages indicating the top users of aux slots. RASP (ASID 3) may own few real frames but a large amount of aux slots. This is not an indication that RASP has a problem.

RASP using AUX Slots

Explanation:

- When High Virtual Shared or Common storage (above the 2Gig bar) is used by any job, frames used to back this area of storage are owned by the job which obtained the storage area
- When REAL storage is low enough to drive paging, these High Virtual pages that are paged to AUX slots are given to and owned by the RSM address space (RASP)
- Need to find out which jobs are using High Virtual Shared or Common storage (and if the amount is higher than normal)
  - One way is to get a dump and use IPCS RSMDATA (see next pages)

The RASP aux slot counts are actually slots used for high virtual shared or common pages belonging to jobs in the system. One way to identify the owner of these pages is to take a dump and issue the IPCS RSMDATA command. See next 2 pages for examples.

IPCS RSMDATA HVSHRDATA Example

```
S START VSA         END VSA           ST K F VT JOBNAME  ASID CREATE TIME         REQUESTOR RQAS
- ----------------- ----------------- -- - - -- -------- ---- ------------------- --------- ----
L 00000200_50700000 00000200_551FFFFF S  2 Y -  J7DTA1   018E 06/23/2013 08:26:20 A1AAED78  000F
                                               J7D42    0190
                                               J7D09    0154
L 00000200_55200000 00000200_583FFFFF S  2 Y -  J7DDM    018F 06/23/2013 08:26:20 A1AAED78  000F
L 00000200_58400000 00000200_5CEFFFFF S  8 Y -  J7D45S   0197 06/23/2013 08:27:07 A1AAED78  000F
                                               J7D09S   0198
L 00000200_5CF00000 00000200_600FFFFF S  8 Y -  J7DDMS   019C 06/23/2013 08:30:46 A1AAED78  000F
G 00000200_80000000 00000220_7FFFFFFF R  7 Y -  TCPIP    0043 06/23/2013 07:17:55 9F062AAC  0041
                                               DB2DIST  005E
                                               DB2DBM1  005C
                                               DB2MSTR  0041
```


The output of the IPCS RSMDATA HVSHRDATA command shows the high virtual shared pages owned by jobs in the system. Please see *z/OS MVS Diagnosis: Reference* for details.

IPCS RSMDATA HVCOMMON Example

```
START VSA         END VSA           Size St T K F L JOBNAME  JOBID    CREATE TIME         REQUESTOR RQAS
----------------- ----------------- ---- -- - - - - -------- -------- ------------------- --------- ----
0000017F_82600000 0000017F_829FFFFF 0004 AC J 2 Y N PFA      STC57270 07/21/2013 02:29:33 AC119DF8  003D
0000017F_82A00000 0000017F_82DFFFFF 0004 AC J 6 N N ACFNET   STC57367 07/21/2013 02:30:56 AC5C98D2  009A
0000017F_82E00000 0000017F_82EFFFFF 0001 AC J 1 Y N JESFAUX  ........ 07/21/2013 02:30:57 A544BDFA  00A1
0000017F_83000000 0000017F_830FFFFF 0001 AC J 1 Y N JESFAUX  ........ 07/21/2013 02:30:57 A544BDFA  00A1
0000017F_83100000 00000180_243FFFFF 0A13 AC J 6 Y N TCPIP    STC57382 07/21/2013 02:30:58 A4B93968  00A7
```


The output of the IPCS RSMDATA HVCOMMON command shows the high virtual common pages owned by jobs in the system. Please see *z/OS MVS Diagnosis: Reference* for details.

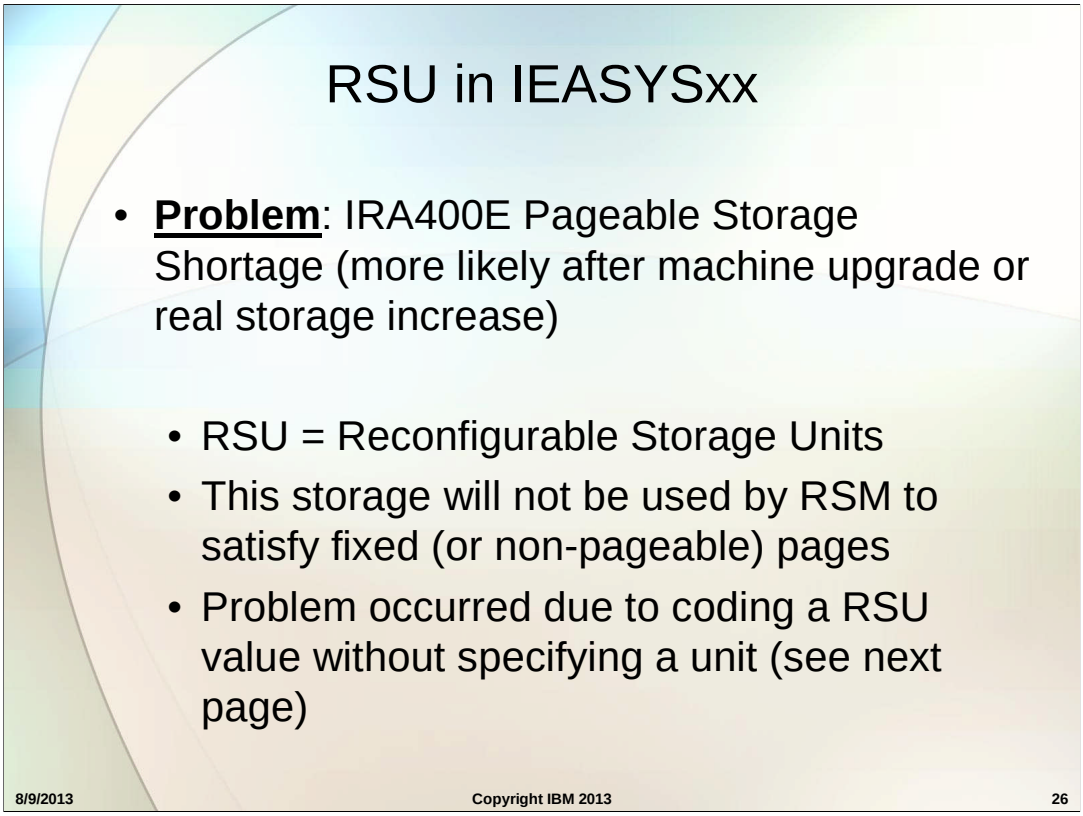

The RSU parameter in IEASYSxx specifies the amount of central storage to be made available for storage reconfiguration. The frames in these storage increments are not to be used for long-term pages and will be designated the non-preferred area. (Long-term pages include SQA pages, common area fixed pages and LSQA or private area fixed pages associated with non-swappable address spaces.)

If you specify a value of 1-9999 without a qualifier (M, G, T, or %), the value is considered to be the number of the units, and the default storage increment size is used. For example, if your machine has a storage increment size of 64 megabytes, specifying 20 causes 20 units of 64M (1.25G in total) to be set aside for storage reconfiguration. Note that the storage increment size is entirely hardware dependent, based not only on the hardware model, but possibly also on the amount of real storage installed on the physical machine (not the LPAR). This means using an unqualified value of 1-9999 can have unexpected results, because its meaning can change dramatically with a simple upgrade to the amount of real storage on the system.

RSU in IEASYSxx

Explanation:

- For best performance, it is recommended that RSU=0 is coded (Healthcheck: RSM_RSU)
- If you need to code a RSU value, use units of M, G or %, instead of a number (which means storage increments)
- Storage increments size can change after a machine upgrade or increase in real storage (see PR/SM Planning Guide)

PROGxx REFRPROT to protect code

Problem:

- Overlays to code are difficult to debug and can cause serious system impact.

Example:

- Recently a customer experienced a 1-bit overlay to authorized code living in Key0 private storage in a CICS region.
- This 1-bit code overlay led to a 5-word overlay of code in Key0 CSA storage.
- Recurring ABEND0C1 errors in the CSA-resident code had significant system impact.

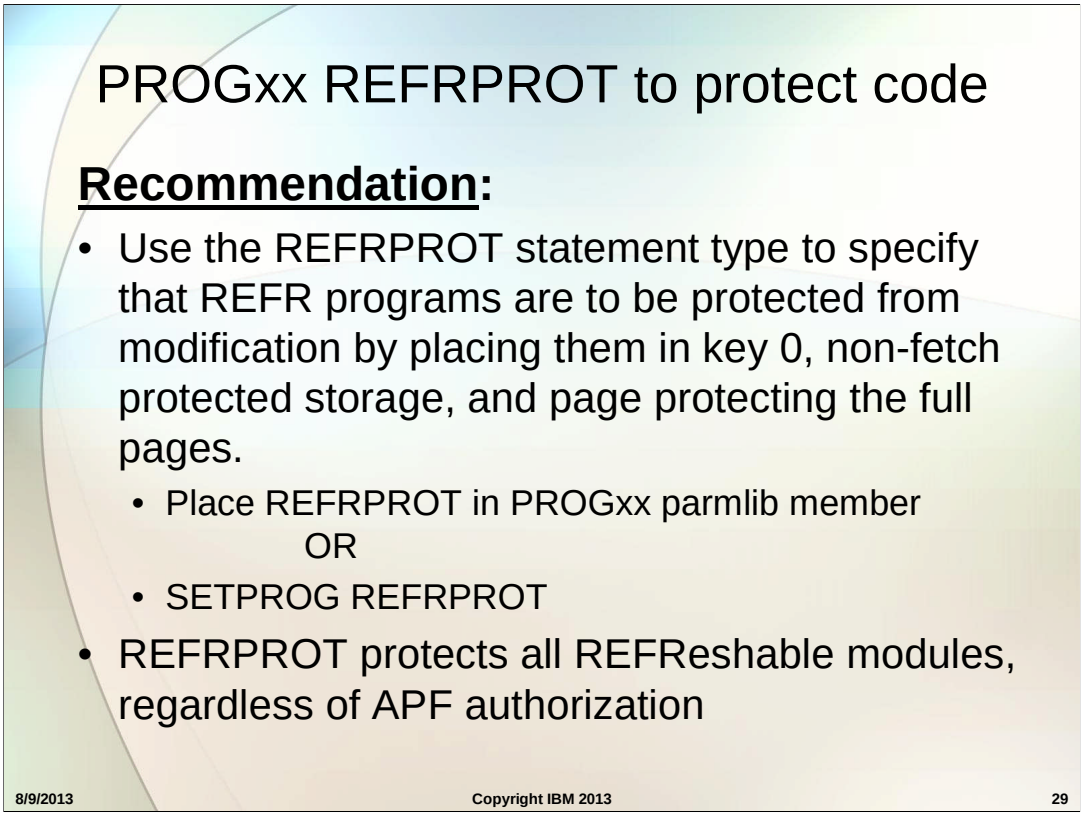

For more information on protection of REFR programs, see *z/OS MVS Program Management: User's Guide and Reference*.

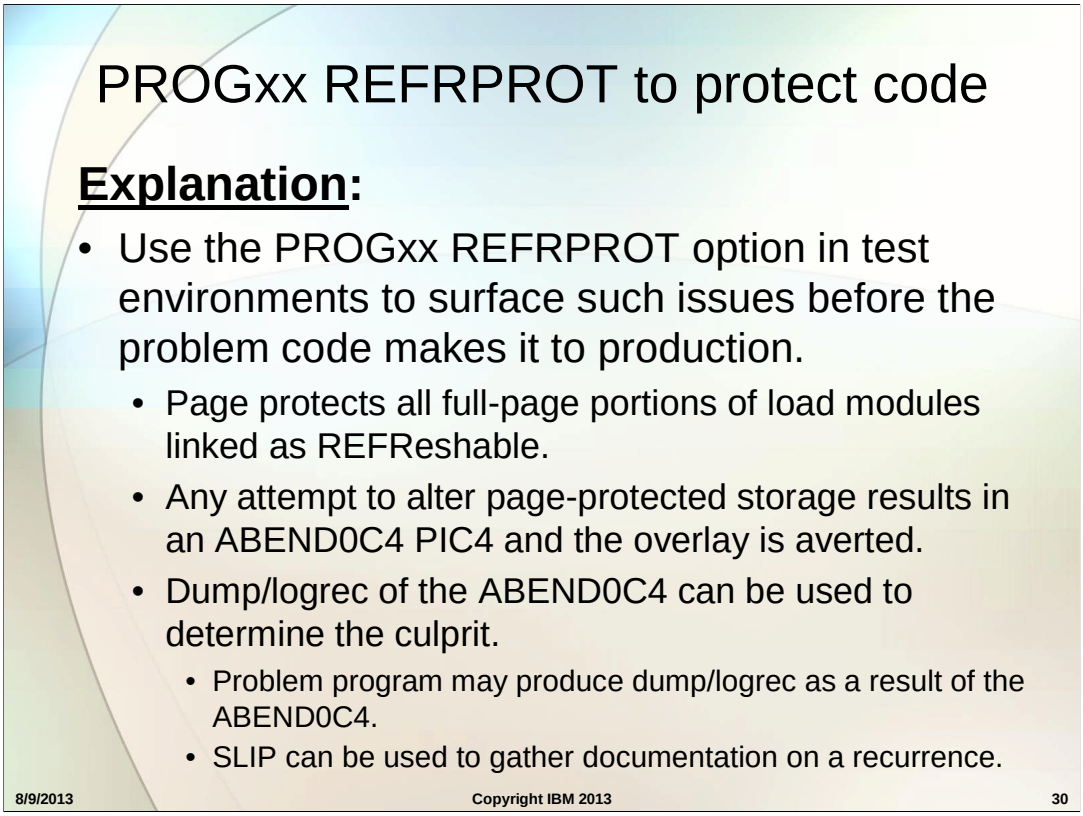

About the REFR link edit attribute:

The module is refreshable. It can be replaced by a new copy during execution without changing the sequence or results of processing. A refreshable module cannot be modified during execution. A module can only be refreshable if all the control sections within it are refreshable. The refreshable attribute is negated if any input modules are not refreshable. Refreshable modules are also reentrant and serially reusable.

The refreshable attribute can be specified for any non-modifiable module.

If ***REFRPROT*** has been specified on the SETPROG command or in parmlib member PROGxx, the module is protected from modification by placing it in key 0, non-fetch protected storage, and page protecting the whole pages.

REFRPROT can be used on production systems as well as test systems, but be aware that ABEND0C4 can occur if a module is link edited as REFR but turns out to be modified.

SDUMP MAXSPACE

Problem:

DB2 dump was partial due to reaching MAXSPACE. What should I set MAXSPACE to?

- **IEA043I** SVC DUMP REACHED MAXSPACE LIMIT - MAXSPACE=xxxxxxxx MEG
- **IEA611I** {COMPLETE|PARTIAL} DUMP ON dsname. MAXSPACE LIMIT REACHED WHILE CAPTURING DUMP

Since dump processing will write captured storage to a dump data set on DASD as soon as the dump data capture completes, the presence of captured data for multiple dumps would imply an issue with obtaining the storage needed to allocate dump data sets.

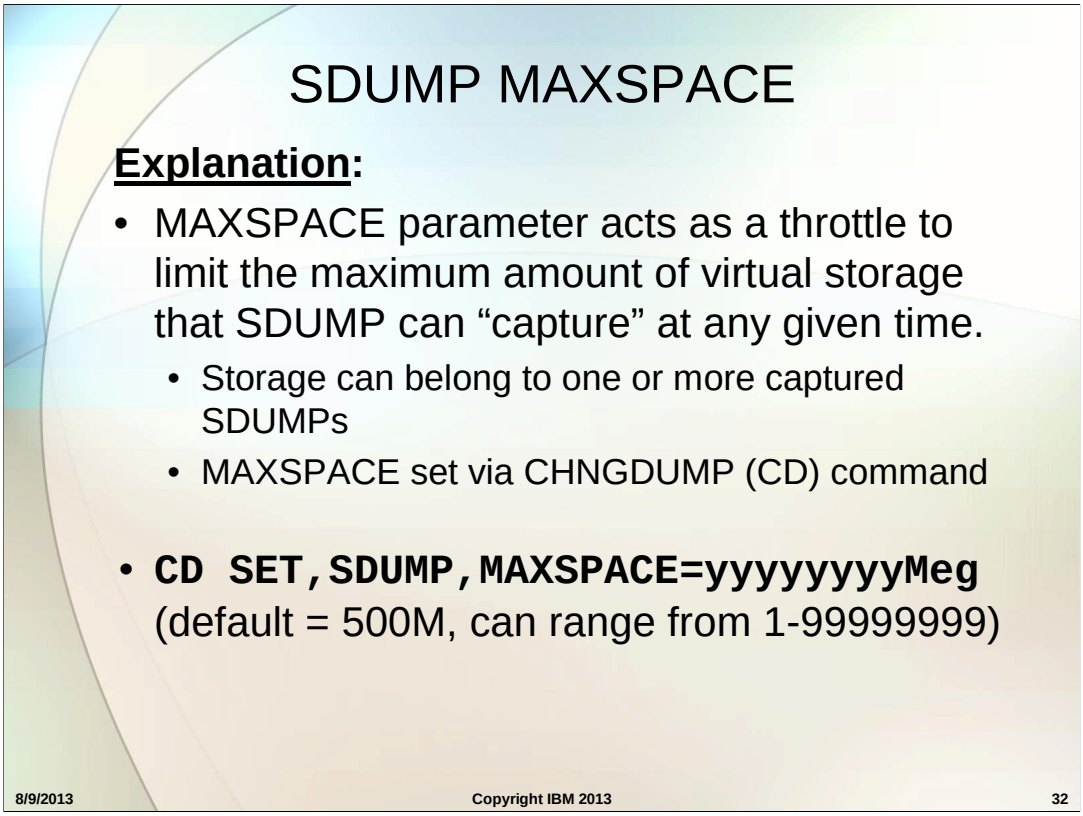

Since dump processing will write captured storage to a dump data set on DASD as soon as the dump data capture completes, the presence of captured data for multiple dumps would imply an issue with obtaining the DASD storage needed to allocate dump data sets.

SDUMP MAXSPACE

Solution:

1. Check sizes of your largest dumps. Given these sizes, what seems like a reasonable value for MAXSPACE?
2. Examine your AUX storage definitions. How much is 1/3rd of your AUX?
3. If Answer1 <= Answer2, then choose a MAXSPACE value in between the two. This will protect your system, while giving you the greatest probability of obtaining a complete dump.

DB2 and WAS tend to produce the largest SVC dumps.

If you have no dumps to use for comparison, see Diagnosis: Tools and Service Aids 2.1.2.1 Allocating SYS1.DUMPxx data sets with secondary extents.

[DB2 guideline is add up the DB2 address spaces + CSA (including above the bar CSA) + up to 800Meg (for buffer pools).]

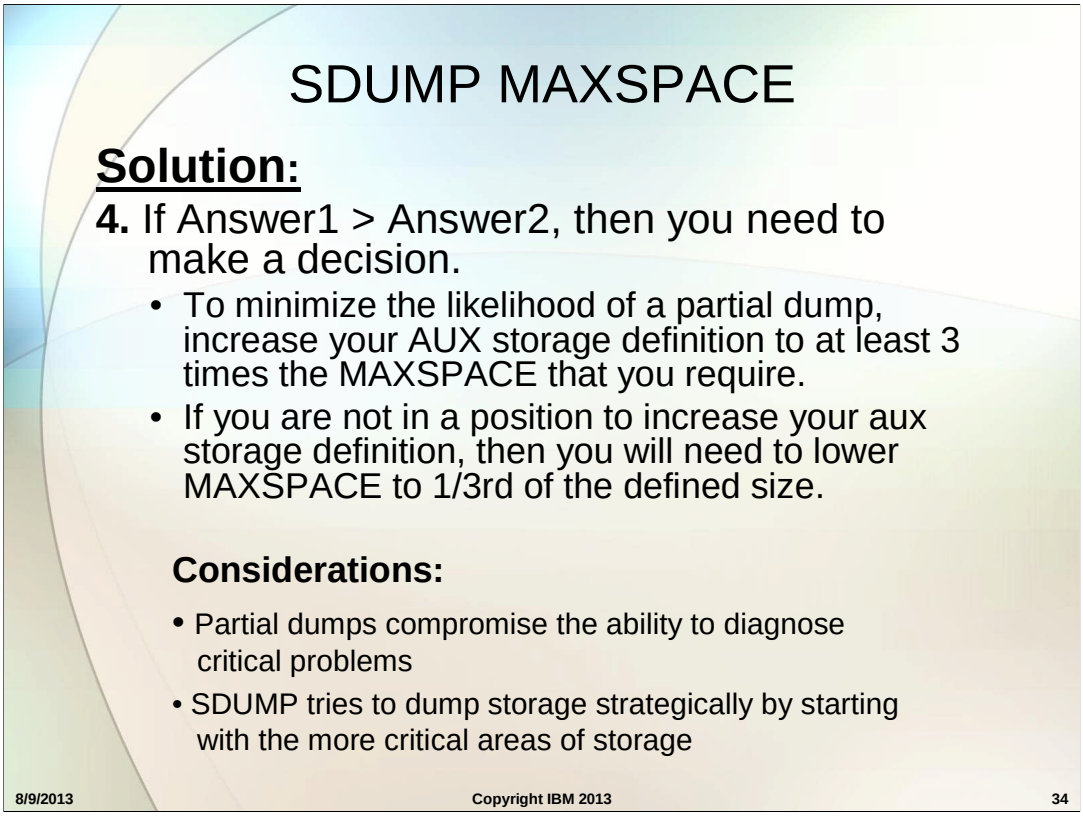

DB2 and WAS tend to produce the largest SVC dumps.

If you have no dumps to use for comparison, see Diagnosis: Tools and Service Aids 2.1.2.1 Allocating SYS1.DUMPxx data sets with secondary extents.

[DB2 guideline is add up the DB2 address spaces + CSA (including above the bar CSA) + up to 800Meg (for buffer pools).]

SDUMP AUXMGMT

Problem:

I ran into AUX storage issues when taking an SVC dump. I'm using a reasonable MAXSPACE. Why did this happen ?

- **IRA205I** 50% AUXILIARY STORAGE ALLOCATED
- **IRA200E** AUXILIARY STORAGE SHORTAGE
- **IRA201E** CRITICAL AUXILIARY STORAGE SHORTAGE
- **IEE711I** [SYSTEM UNABLE TO DUMP|SYSTEM DUMP NOT TAKEN. A CRITICAL AUXILIARY STORAGE SHORTAGE EXISTS]

SDUMP AUXMGMT

Explanation:

Even with a properly set MAXSPACE, SDUMP can still trigger an AUX storage condition if the overall system is using a sizeable amount of AUX storage. The AUXMGMT parameter offers additional system protection.

SDUMP AUXMGMT

Solution:
Use AUXMGMT parameter!

- SDUMP AUXMGMT acts as a safety net **for systems exceeding recommended AUX utilization** (=30%).
  - **CD SET,SDUMP,AUXMGMT=ON** (the default)
- New SDUMPs are prevented when AUX storage usage reaches 50%
- SDUMPs in the process of being captured are stopped when AUX usage reaches 65%.
- If AUXMGMT=OFF, then SDUMP function is not affected until AUX usage goes to 85% (critical)

SDUMP AUXMGMT

Problem:

AUXMGMT protection detected aux storage usage greater than 50% and is preventing any new SVC dumps from being taken. How do I recover my system's ability to take a dump?

- **IEA611I** {COMPLETE|PARTIAL} DUMP ON dsname. A CRITICAL AUXILIARY STORAGE SHORTAGE EXISTS

Note: SDUMP's critical storage indication means the AUXMGMT threshold has been reached, but doesn't mean the system has 70%-85% AUX storage used.

SDUMP AUXMGMT

Explanation:

A low threshold of 35% must be attained (35%) before SDUMP processing is allowed to resume.

- Resetting AUXMGMT=OFF after AUX storage utilization has reached the 50% threshold will ***not*** relieve the above low threshold requirement! Once you hit the AUXMGMT ON limit you **MUST** hit the low limit (35%) before SDUMPs will again be allowed.

SDUMP AUXMGMT

Solution:

There are two ways to attain the low limit:

1. CANCEL or wait for the address spaces that have pages on AUX to free the storage or the job to end **OR**
2. Add page datasets such that the percentage of overall available AUX slots is then below 35%. If you hit a AUXMGMT limit, and **cannot** add additional page datasets, you will have to revert to option 1.

- If set correctly, MAXSPACE and AUXMGMT work hand in hand to protect the system.