Session 14254

# Common z/OS Problems You Can Avoid


## SHARE Anaheim, CA

## March 10th, 2014

**Evan Haruta** **haruta@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

- Runtime Diagnostics: RTD
- STP issues
- ICSF/Crypto Master Keys
- FTP'ing Problem Doc & SAD Best Practices
- WAITCCC & Quiesce
- RSU in IEASYSxx (SDUMP issue)
- PROGxx REFRPROT
- SDUMP AUXMGMT issues
- High CPU in *MASTER*

# Runtime Diagnostics: RTD

- **Problem:**
  - When z/OS runs into problems such as hangs, loops, delays, missed SLAs, numerous error messages, ENQ contention, it is often difficult for sysprog to determine the source(s) of problem

# Runtime Diagnostics: RTD

- **Solution**: Run RTD. Some of the things that RTD will examine include

- Reviewing critical messages in the operlog
- Examining address spaces with high processor usage
- Looking for an address space that might be in a loop
- Evaluating local lock conditions
- Analyzing various types of contention that include ENQ, GRS latch contention, and z/OS UNIX file system latch contention.

# Runtime Diagnostics: RTD

- **Solution**: Runtime Diagnostics is a diagnostic tool to run in your environment when your system experiences symptoms that require its use. Run it when you experience system degradation or if you want to check for potential problems; do **not** run when the system is operating normally.

- After you start Runtime Diagnostics (S HZR,SUB=MSTR), you can analyze the home system by entering the following MODIFY command.
- F HZR,ANALYZE

# Runtime Diagnostics: RTD

```
HZR0200I RUNTIME DIAGNOSTICS RESULT 568
SUMMARY: SUCCESS
REQ: 003 TARGET SYSTEM: SY1 HOME: SY1 2010/12/21 - 13:45:49
INTERVAL: 60 MINUTES
EVENTS:
FOUND: 02 - PRIORITIES: HIGH:02 MED:00 LOW:00
TYPES: HIGHCPU:01
TYPES: LOCK:01
------------------------------------------------------------------
EVENT 01: HIGH - HIGHCPU - SYSTEM: SY1 2010/12/21 - 13:45:50
ASID CPU RATE:99% ASID:002E JOBNAME:IBMUSERX
STEPNAME:STEP1 PROCSTEP: JOBID:JOB00045 USERID:IBMUSER
JOBSTART:2010/12/21 - 11:22:51
ERROR: ADDRESS SPACE USING EXCESSIVE CPU TIME. IT MIGHT BE LOOPING.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
------------------------------------------------------------------
EVENT 02: HIGH - LOCK - SYSTEM: SY1 2010/12/21 - 13:45:50
HIGH LOCAL LOCK SUSPENSION RATE - ASID:000A JOBNAME:WLM
STEPNAME:WLM PROCSTEP:IEFPROC JOBID:++++++++ USERID:++++++++
JOBSTART:2010/12/21 - 11:15:08
ERROR: ADDRESS SPACE HAS HIGH LOCAL LOCK SUSPENSION RATE.
ACTION: USE YOUR SOFTWARE MONITORS TO INVESTIGATE THE ASID.
```

# STP issues

- **Problem**: LPAR may not be able to rejoin sysplex, OR sysplex may have to be reIPL'd.
  - If a system is using SIMERTID on a CEC that is in STP timing mode
  - See OA26750/OA26800
  - If a clock jump occurs (due to "sync check" or if operator manually changes time via STP Panels), then all active LPARs will take an LPAR EPOCH (a clock offset)

# STP issues

- **Explanation**: LPAR EPOCH mismatch prevents the system rejoining the sysplex after the LPAR is re-activated.
  - Normally the EPOCH is transparent.
  - Re-IPL'ing a system is fine since EPOCH is not cleared.
  - After an LPAR is reactivated it cannot join
    - **IXC416I** SIMETRID IS SUPPORTED ONLY WHEN ALL SYSTEMS IN THE SYSPLEX RUNNING ON THE SAME CPC AND ARE USING SIMETRID CONSISTENTLY
    - **IXC404I** SYSTEM(S) ACTIVE OR IPLING: SC63 SC65
    - **IXC419I** SYSTEM(S) NOT SYNCHRONIZED: SC63 SC65
    - **IXC420D** REPLY I TO INITIALIZE SYSPLEX SANDBOX, OR R TO REINITIALIZE

# STP issues

- **Explanation**: When the LPAR was re-activated the LPAR EPOCH was cleared and now there is a mismatch with the other members of the sysplex
  - ALL LPARs in sysplex must be reactivated to clear EPOCH.
  - To identify the EPOCH value, use IPCS
    - L 10?+8C?+4?+400 (CPLXEPOCH)

# STP Issues

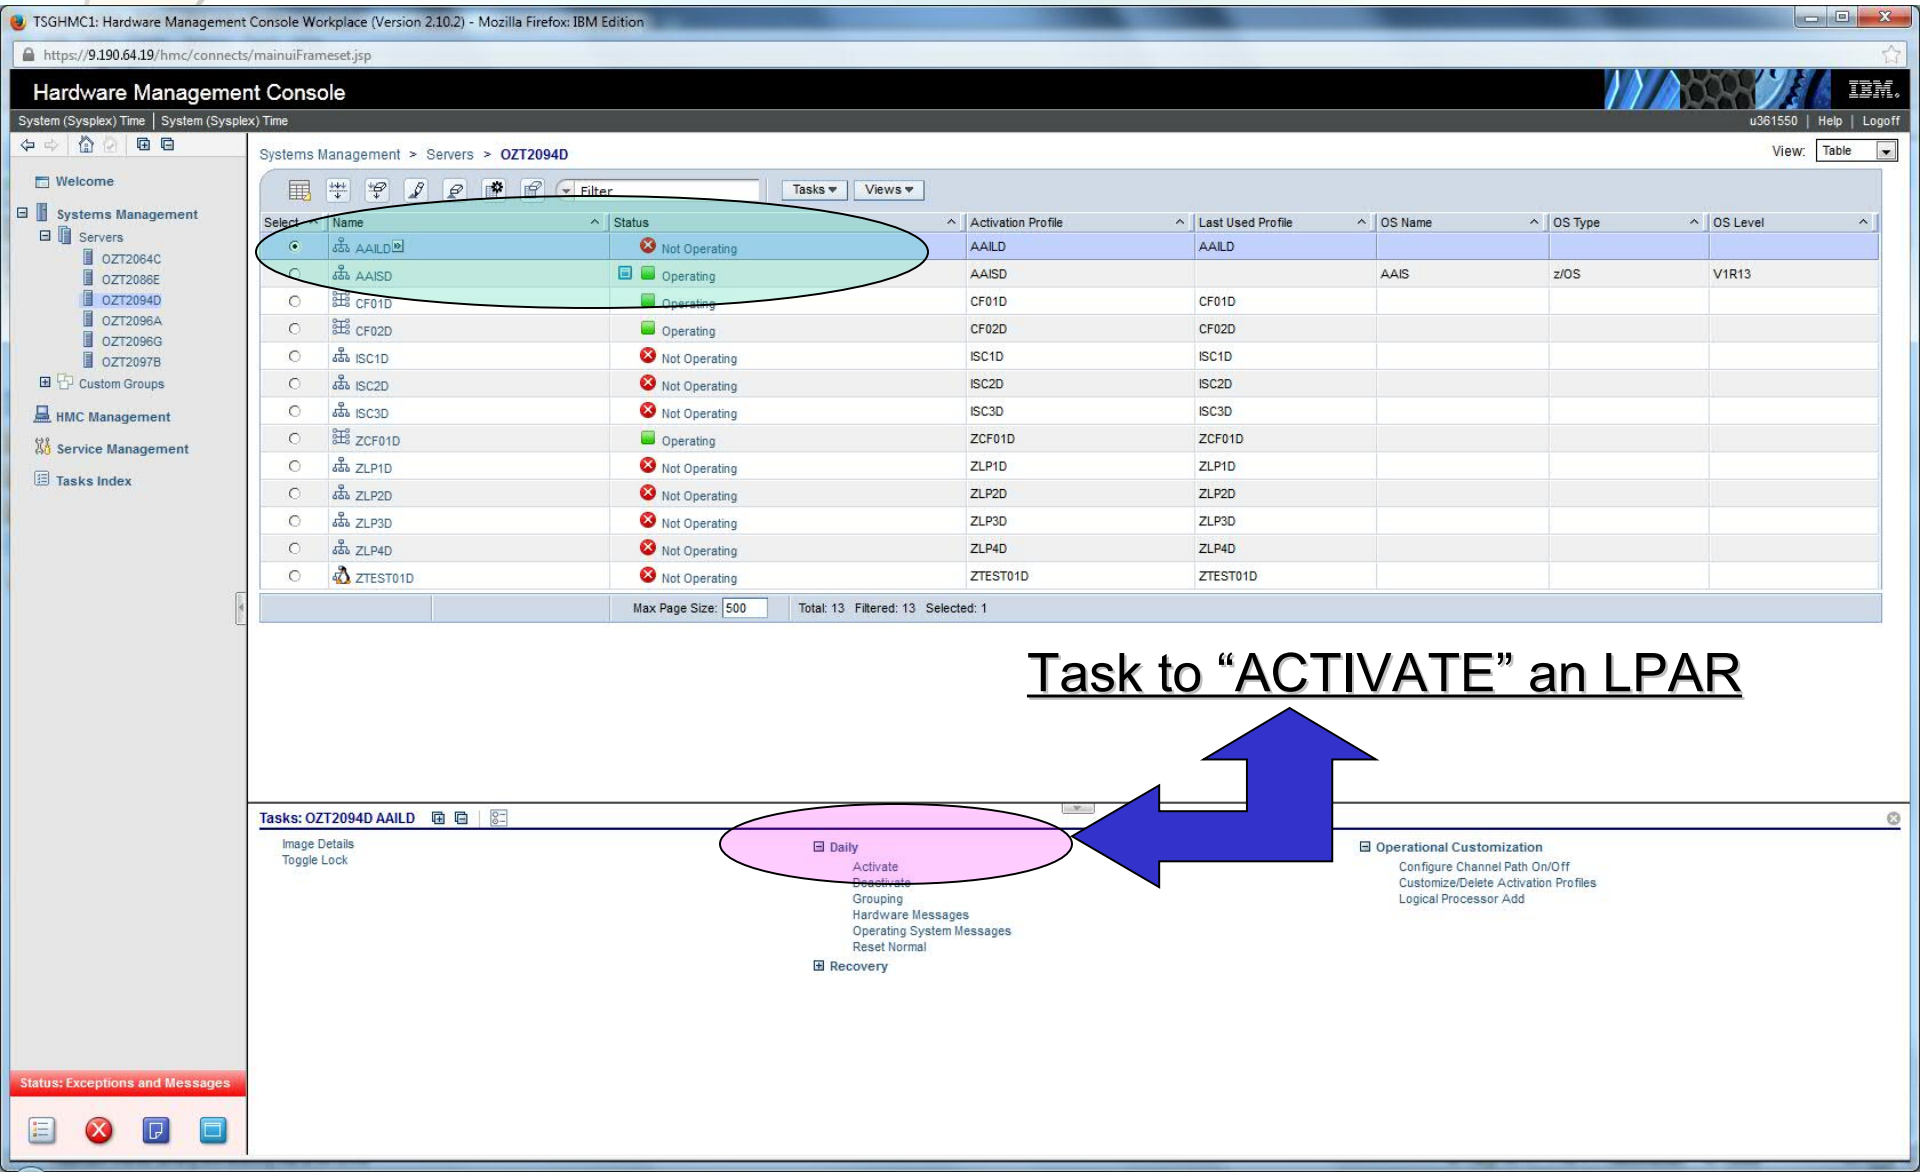
TSGHMC1: Hardware Management Console Workplace (Version 2.10.2) - Mozilla Firefox: IBM Edition
https://9.190.64.19/hmc/connects/mainuiFrameset.jsp
Hardware Management Console
IBM.
System (Sysplex) Time | System (Sysplex) Time
u361550 | Help | Logoff
Welcome
Systems Management
Servers
OZT2064C
OZT2086E
OZT2094D
OZT2096A
OZT2096G
OZT2097B
Custom Groups
HMC Management
Service Management
Tasks Index
Status: Exceptions and Messages
Systems Management > Servers > OZT2094D
View: Table
Filter
Tasks
Views
Select
Name
Status
Activation Profile
Last Used Profile
OS Name
OS Type
OS Level
AAILD Not Operating AAILD AAILD
AAISD Operating AAISD AAIS z/OS V1R13
CF01D Operating CF01D CF01D
CF02D Operating CF02D CF02D
ISC1D Not Operating ISC1D ISC1D
ISC2D Not Operating ISC2D ISC2D
ISC3D Not Operating ISC3D ISC3D
ZCF01D Operating ZCF01D ZCF01D
ZLP1D Not Operating ZLP1D ZLP1D
ZLP2D Not Operating ZLP2D ZLP2D
ZLP3D Not Operating ZLP3D ZLP3D
ZLP4D Not Operating ZLP4D ZLP4D
ZTEST01D Not Operating ZTEST01D ZTEST01D
Max Page Size: 500
Total: 13 Filtered: 13 Selected: 1
Task to “ACTIVATE” an LPAR
Tasks: OZT2094D AAILD
Image Details
Toggle Lock
Daily
Activate
Deactivate
Grouping
Hardware Messages
Operating System Messages
Reset Normal
Recovery
Operational Customization
Configure Channel Path On/Off
Customize/Delete Activation Profiles
Logical Processor Add

# STP Issues

- **Explanation**: Warning is provided before Activate is performed

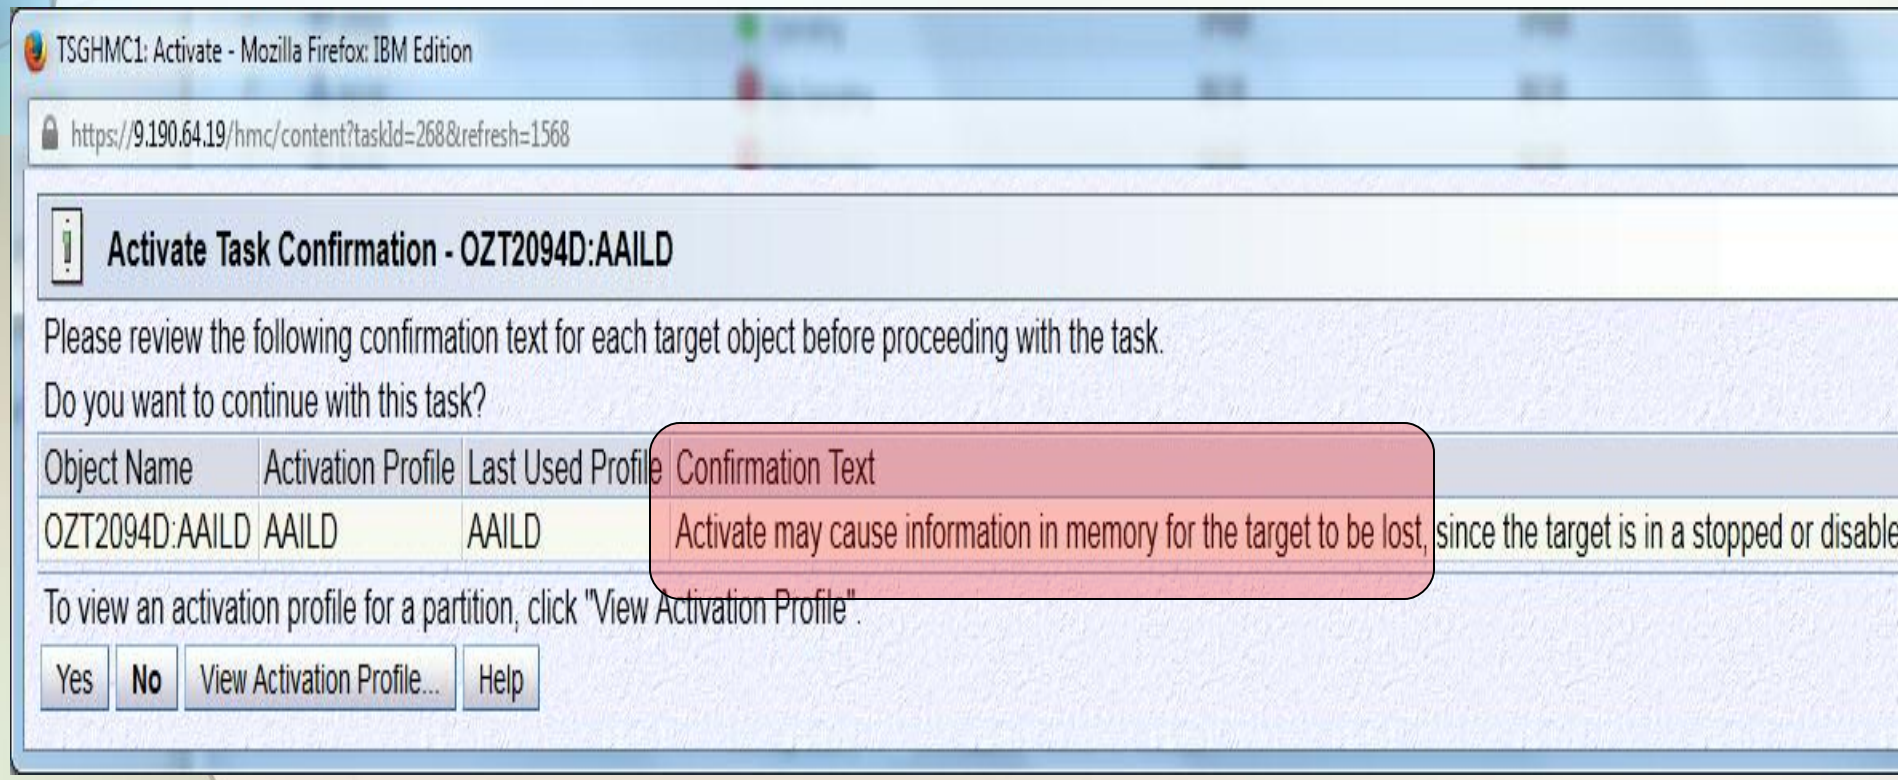

# STP issues

- **Problem**: STP configuration not saved across a POR
- **Solution**: Take advantage of an STP function that saves the STP configuration across POR. The function is activated via the Network configuration panel by selecting the "Only allow the server(s) specified above to be in the CTN". This can only be done for a one server CTN or a two server CTN with a BTS assigned.

# STP Issues

- **What to do**: Sysplex time panel

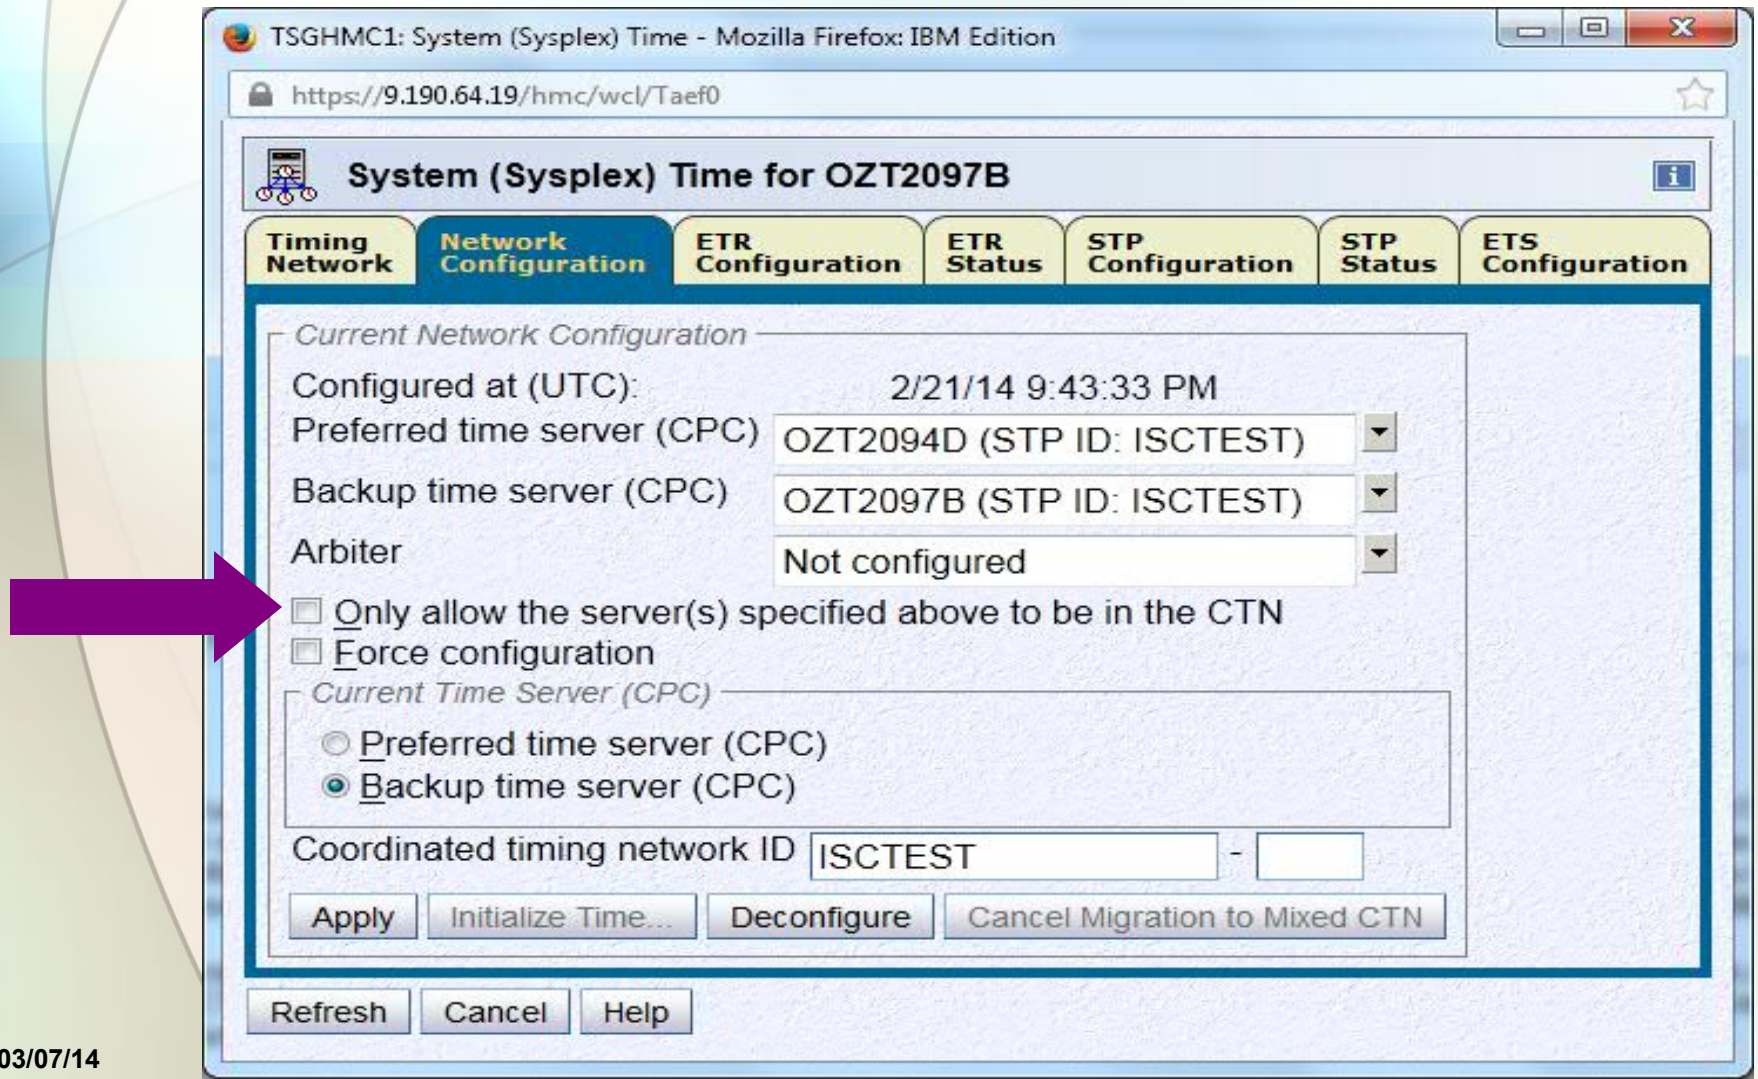

# ICSF/Crypto Master Keys

- **Problem**: Unable to use the old key datasets (CKDS and PKDS) after migration to a new machine.
  - On a new machine, if the old key datasets are to be used, the original Master Keys are needed
  - If the Master Keys are forgotten, these old key datasets cannot be used
  - At this point the only option left is to power up the old machine 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

# FTP'ing Problem Doc

- **Problem**: L2 is not able to readily find the problem documentation that you FTP.
  - File name of ppppp.bbb.ccc.short.desc should be used (ppppp=problem number, bbb=branch, ccc=country)
    - Automation tools look for it and update PMR with doc arrival information
  - Use of file names like pmrxxxxx.bbb.ccc.short.desc will be an anomaly and your doc will not be found readily

# 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
```

# z/OS Best Practices: SADmp

- **Problem**: Stand Alone Dumps
- No one wants to take a SADmp, but when you do, be prepared
- z/OS Best Practices: Large Stand-Alone Dump Handling - Version 3

- http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD103286

# z/OS Best Practices: SADmp

- **What is discussed**:
  - Stand-Alone dump data set definition and placement
  - IPCS performance considerations
  - Preparing documentation to be analyzed
  - Sending documentation to IBM support
  - Testing your stand-alone dump setup

# z/OS Best Practices: SADmp

- **What is discussed**: PDUU, Copydump, AMDSADDD, compression, encryption, transmission, etc.…
  - Most recently, in z/OS V1R13, the Problem Documentation Upload Utility allows you to transmit large files much more quickly and encrypts the data all in the same process.
  - See *MVS Diagnosis: Tools and Service Aids (GA22-7589)*, Chapter 4

# WAITCCC and QUIESCE

- **Problem**: System stops with a WAITCCC for no apparent reason.

```
Wait State Message Issued at 17:22:28 on Day 032 of 2014:

 *BLW002I SYSTEM WAIT STATE 'CCC'X -
       QUIESCE FUNCTION PERFORMED
```

  - This system Wait State indicates that a system QUIESCE command has been issued
    - This is a **restartable** wait state
      - Most customers are taken by surprise and IPL
      - To restart the system following a WAITCCC, perform Restart function on waiting processor

# WAITCCC and QUIESCE

- **Problem**: QUIESCE command may be entered from console or JCL

```
JOB37109 00000090  IEFC165I // QUIESCE MQ375J
INSTREAM 00000290  QUIESCE MQ375J
```

  - Loss of command prefix character (CPF) for a subsystem QUIESCE command such as that used for DB2 or for MQ can lead to the command being interpreted as a z/OS *system* QUIESCE
  -
    - Check SYSLOG or MTRACE in the SADmp for evidence of who issued QUIESCE

# WAITCCC and QUIESCE

- **What-to-do**: An ounce of prevention is worth a pound of cure....

  **Don't allow a system QUIESCE to get accidentally issued on your system!!**

  - Create a RACF profile in the OPERCMDS class to restrict/prevent users from issuing the QUIESCE command
    - Profile name is MVS.QUIESCE
    - Define it with UACC(NONE)

# RSU in IEASYSxx

**Problems**:

- ✓ IRA400E Pageable Storage Shortage
- ✓ Job performance issues (missed SLAs)
- ✓ Long SVC dump times

  - RSU = Reconfigurable Storage Units
  - This storage will not be used by RSM to satisfy fixed (or non-pageable) pages
  - Problem occurred due to coding an RSU value without specifying a unit (see next page)

# 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)
- SVC Dump times can be seen to increase

# 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.

# PROGxx REFRPROT to protect code

**Recommendation:**

- Use the REFRPROT statement type to specify that REFR programs are to be protected from modification by placing them in key 0, non-fetch protected storage, and page protecting the full pages.
  - Place REFRPROT in PROGxx parmlib member
    OR
  - SETPROG REFRPROT
- REFRPROT protects all REFReshable modules, regardless of APF authorization

# PROGxx REFRPROT to protect code

**Explanation:**

- Use the PROGxx REFRPROT option in test environments to surface such issues before the problem code makes it to production.
  - Page protects all full-page portions of load modules linked as REFReshable.
  - Any attempt to alter page-protected storage results in an ABEND0C4 PIC4 and the overlay is averted.
  - Dump/logrec of the ABEND0C4 can be used to determine the culprit.
    - Problem program may produce dump/logrec as a result of the ABEND0C4.
    - SLIP can be used to gather documentation on a recurrence.

# 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

# SDUMP MAXSPACE

**Explanation:**

- MAXSPACE parameter acts as a throttle to limit the maximum amount of virtual storage that SDUMP can “capture” at any given time.
  - Storage can belong to one or more captured SDUMPs
  - MAXSPACE set via CHNGDUMP (CD) command

- **`CD SET,SDUMP,MAXSPACE=yyyyyyyyMeg`** (default = 500M, can range from 1-99999999)

# 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.

# SDUMP MAXSPACE

**Solution:**

4. If Answer1 > Answer2, then you need to make a decision.
   - To minimize the likelihood of a partial dump, increase your AUX storage definition to at least 3 times the MAXSPACE that you require.
   - If you are not in a position to increase your aux storage definition, then you will need to lower MAXSPACE to 1/3rd of the defined size.

**Considerations:**

- Partial dumps compromise the ability to diagnose critical problems
- SDUMP tries to dump storage strategically by starting with the more critical areas of storage

# 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 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 an 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.

# High CPU in *MASTER*

**Problem:** Spike in CPU consumption in *MASTER* and/or RASP Address Space

- Repetitive short-lived spikes in CPU, which are hard to capture in SVC dumps
- SYSTRACE PERFDATA (or various monitors) may show SRB time in either IAXUR or IAXUO
- Spikes seen on occasions as high as 30-50%

# High CPU in *MASTER*

**Environment:** All of the following conditions must be present

- z/OS 1.13.1 (JBB778H) or
  z/OS 2.1 (HBB7790)
- zEC12 (d/t 2827) or zBC12
  - SCM (aka Flash memory) capable
- No PAGESCM statement in IEASYSxx
  - Or PAGESCM=NONE not specified
- **No** SCM is actually configured for use

# High CPU in *MASTER*

**Symptoms:** Issued by RSM at IPL time

```
IAR031I USE OF STORAGE CLASS
MEMORY FOR PAGING IS ENABLED -
PAGESCM=ALL, ONLINE=00000000M
```

Note: PAGESCM=ALL is the default

# High CPU in *MASTER*

**Solution:** Specify PAGESCM=NONE

- Default value is PAGESCM=ALL which reserves some real storage for the Pageable Large Frame area *
- Using PAGESCM=NONE at IPL will eliminate the high CPU completely

*Not to be confused with Large Frame Area (LFAREA)

# High CPU in *MASTER*

**Solution:** Specify PAGESCM=NONE in IEASYSxx parmlib

```
IAR032I USE OF STORAGE-CLASS
  MEMORY FOR PAGING IS NOT
ENABLED - PAGESCM=NONE
```

In this environment no Pageable Large Frame area will be built