Session 15594

# Common z/OS Problems You Can Avoid

(Final Version)


## SHARE in Pittsburgh

## August 6th, 2014

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

# 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

- SADMP Best Practices
- HMC Change Logical Partition Controls
- Runtime Diagnostics: RTD
- WAITCCC & Quiesce
- High CPU in *MASTER*
- 1 Meg Large Pages
- SMF Data Loss
- Managing RACF Certificates
- System SSL Component Trace
- Dealing with a hung job

# z/OS Best Practices: SADMP

**Problem**: Unprepared for Stand-Alone Dumps

- No one wants to take a SAD, but when you do, be prepared
- z/OS Best Practices: Large Stand-Alone Dump Handling - Version 4

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 V1.13, 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

# z/OS Best Practices: SADMP

**New in V4 of the Best Practices**:

- Use of INITAPPEND option of COPYDUMP in z/OS V2.1 to append initialization dump directory to dump dataset
- This will shorten the process of dump initialization by the IBM Support Center
- Recommended for full SAD only (not for subset dump)

# z/OS Best Practices: SADMP

**Want to practice?** (highly recommended)

- Take a Standalone-Alone dump
- Open a Software Usage PMR using the component of AMDSADMP
- FTP the dump to IBM using PDUU
- Ask the support center to check the dump to make sure it is good

# HMC – Change Logical Partition Controls

- **Problem:**
  - Incorrect usage of ‘Defined Capacity’ in HMC ‘Change Logical Partition Controls’ panel
  - ‘Defined Capacity’ on this panel does not mean number of CPUs or LPAR weight
  - Instead it is the MSU soft-cap 4-hour rolling average limit of the LPAR
  - A small unintended value in Defined Capacity will most likely result in a performance issue

# HMC – Change Logical Partition Controls

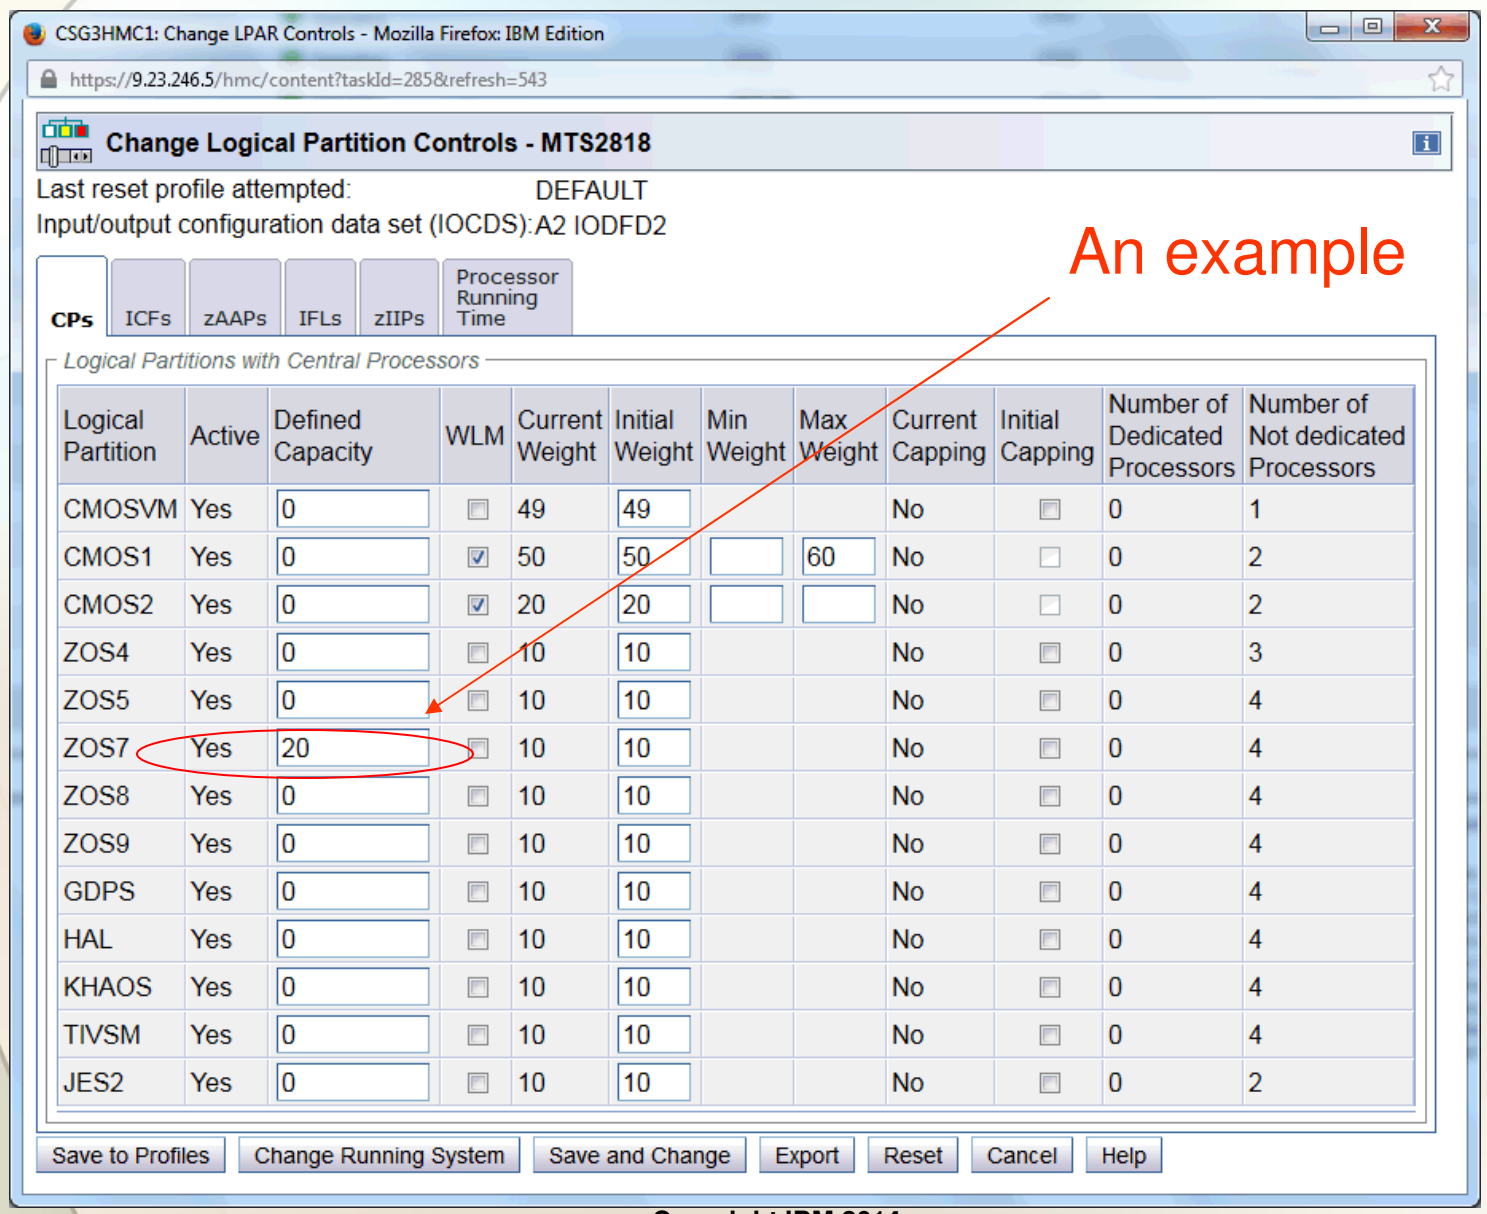

# HMC – Change Logical Partition Controls

RMF III CPC Report

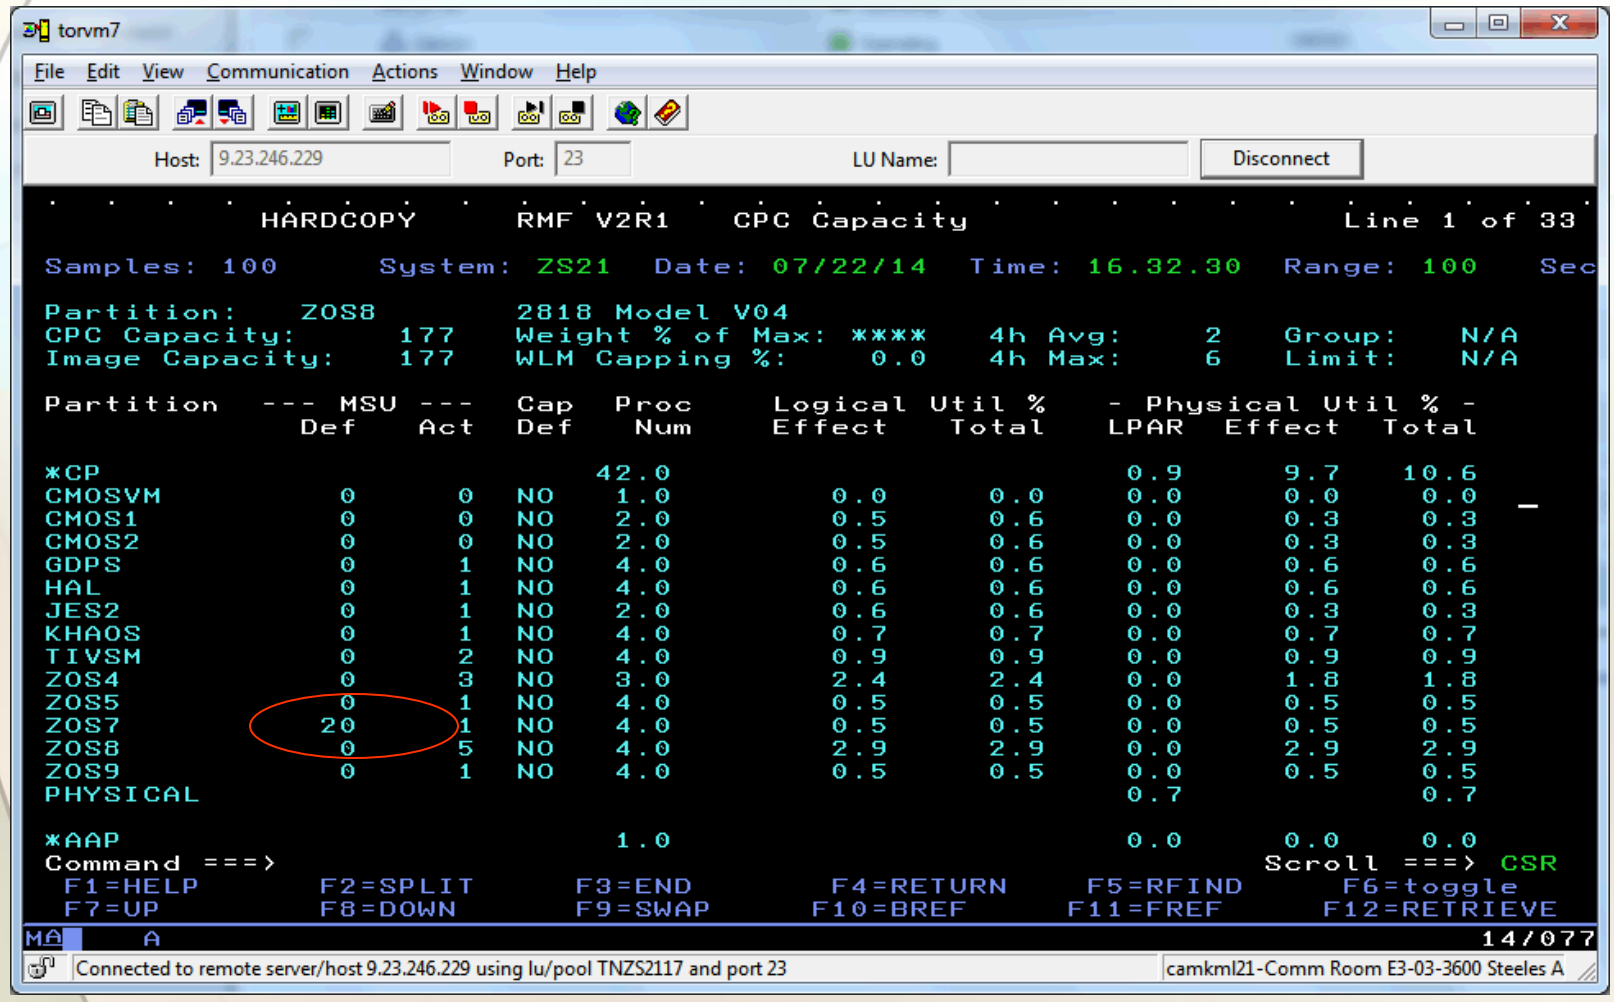

# Runtime Diagnostics: RTD

- **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**: 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 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.
```

# 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)

# 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 Web Deliverable (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

# 1Meg Large Pages

**Environment:** LPAR in which LFAREA is specified in IEASYSxx

1) The LFAREA parameter specifies the amount of online real storage available at IPL to reserve for backing 1 MB pages (and 2 GB pages).
2) Components such as DB2 and IMS, amongst others, can obtain High Virtual storage (usually during initialization) and request that it be backed with 1Meg Large Pages.
3) IARV64 GETSTOR .....,PAGEFRAMESIZE=1MEG

# 1Meg Large Pages

**Symptoms:** Displays show that either component is not using Large Pages, or far fewer total system-wide Large Pages are being used than anticipated

- Usually there will not be a message or ABEND or other indicator (except possibly performance degradation) to indicate that Large Pages were not obtained

# 1Meg Large Pages

## **Symptoms:** DISPLAY VIRTSTOR

```
IAR019I  07.44.41 DISPLAY VIRTSTOR
    SOURCE                        = CC    (IEASYSCC)
    TOTAL LFAREA                  = 18M
    LFAREA AVAILABLE              = 10M
    LFAREA ALLOCATED (1M)         =  8M   (Currently used)
    LFAREA ALLOCATED (4K)         =  0M
    MAX LFAREA ALLOCATED (1M)     =  8M
    MAX LFAREA ALLOCATED (4K)     =  0M
```

**Why isn't the full 18M LFAREA being used?** **Assuming that the sysprog correctly sized LFAREA, it could be that applications have not yet requested backing by 1Meg large pages. It could be that they requested 1Meg large pages, but not enough space was available.**

# 1Meg Large Pages

**Problem:** Application requests 1Meg large pages and request is failed by RSM. Either the LFAREA was not sized correctly based on current users, or unanticipated user(s) requesting 1Meg Large Pages made requests prior to anticipated users.

**Example**: Based on DB2 sizings, sysprog defines LFAREA to be 18Meg size. LFAREA properly gets sized to 18Meg. Then during address space initialization, IMS initializes earlier than DB2, and requests 8Meg of Large pages to back IMS Log Buffers.

At this point there will be 10Meg Available LFAREA. But when DB2 initializes and requests 18Meg it will be failed by RSM since there is not enough ("all or nothing"). Control returns to DB2 which then falls back to traditional 4K buffers instead of 1Meg large pages.

# 1Meg Large Pages

## **Symptoms:** DB2 DISPLAY BUFFERPOOL

```
-DB2A DIS BPOOL=BP1
 DSNB401I  -DB2A BUFFERPOOL NAME BP1, BUFFERPOOL ID 1, USE COUNT 1
 DSNB402I  -DB2A BUFFER POOL SIZE = 2000 BUFFERS  AUTOSIZE = NO
            VPSIZE MINIMUM = 0 VPSIZE MAXIMUM = 0
              ALLOCATED       =     2000   TO BE DELETED   =        0
              IN-USE/UPDATED  =        0
 DSNB406I  -DB2A PGFIX ATTRIBUTE -
              CURRENT = YES
              PENDING = YES
            PAGE STEALING METHOD = LRU
DSNB546I  -DB2A PREFERRED FRAME SIZE 1M
        0 BUFFERS USING 1M FRAME SIZE ALLOCATED
DSNB546I  -DB2A PREFERRED FRAME SIZE 1M
       2000 BUFFERS USING 4K FRAME SIZE ALLOCATED
```

# 1Meg Large Pages

## **Symptoms:** IMS Log Buffers (BUFSTOR)

**BUFSTOR=**

Specifies whether to obtain 31-bit or 64-bit virtual storage for OLDS log buffers. The default value is 31.

Recommendation: When running with 64-bit log buffers, back your log buffers by large (1 MB) pages, so that IMS can improve performance, due to more efficient dynamic address translation.

To set up your system to have large pages available, you must specify the LFAREA= parameter in the z/OS® IEASYSxx PARMLIB member. Make sure that you request enough storage on LFAREA= to contain all of your log buffers, plus any other large page usage in your system.

Restriction: IMS obtains log buffers in 64-bit virtual storage only when the OLDS block size is a multiple of 4096, and when the OLDS are allocated as DFSFMS extended-format data sets. Otherwise, IMS obtains log buffers in 31-bit virtual storage, even if BUFSTOR=64 is specified.

.

If IMS does not get what it needs, IMS will issue **ABENDU0070** and IMS message **DFS2205I** at initialization time.

# 1Meg Large Pages

**Solution:** Ensure proper LFAREA sizing is done by evaluating ALL applications and subsystems which will be using 1Meg Large Pages

- Current users include, but not limited to, DB2, IMS, WAS - this list will grow
- Review migration changes on release boundaries for applications/subsystems
- Also check IAR028I at IPL time to verify that LFAREA requested was obtained, or IAR041I-IAR048I for potential errors with LFAREA

# 1Meg Large Pages

## **Documentation**:

**z/OS**: IEASYSxx (LFAREA= )

z/OS MVS Initialization and Tuning Reference

IARV64 Macro

z/OS MVS Programming: Authorized Assembler Services Reference EDT-IXG

**IMS**: Defining DASD logging initialization parameters

http://www-01.ibm.com/support/knowledgecenter/SSEPH2_13.1.0/com.ibm.ims13.doc.sdg/ims_dfsvsmxx_proclib_dasd_logging.htm?lang=en

**DB2**: DB2 10 for z/OS Performance Topics

3.1 1 MB page frame support

http://www.redbooks.ibm.com/redbooks/pdfs/sg247942.pdf

# SMF Data Loss

**Problem:** Seeing the following messages

- IEE986E SMF HAS USED 100% OF AVAILABLE BUFFER
- IEE979W - SMF DATA LOST - NO BUFFER SPACE AVAILABLE TIME=hh.mm.ss

- Not only does suffering SMF data loss compromise the ability to do accounting and problem analysis, but because SMF uses the CMS SMF lock which is a global system resource, concentrated periods of SMF activity such as may be experienced when recovering from a “data lost” condition could have temporary but unsettling system impact.

**Solution**: Proactively manage your SMF environment

# SMF Data Loss

## **Recommendations**:

- Take advantage of the SMF option of writing to a logstream and let logger take care of the file management.
- If using SYS1.MANx data sets, automate dumping of the full data set on SMF data set switch:
  - Start SMF dump program SMFDUMP (IFASMFDP) from IEFU29
  - Automate start of SMF dump program on SMF message
    - IEE362A SMF ENTER DUMP FOR SYS1.MANn ON ser
  - Automate start of SMF dump program on WTO issued by IEFU29

# SMF Data Loss

**Troubleshooting:** Possible reasons for an SMF buffer shortage

- Too much SMF data coming in, i.e. faster than can be processed
- SMF does not have enough processing power to keep up
- DASD has slowed down
- SMF data not getting dumped

# SMF Data Loss

**Troubleshooting**: What to check

- DASD-related checks
  - Hardware errors on the device
  - I/O delays due to contention on volumes due to high volume of SMF records
  - No other data sets on volume, including other SYS1.MANx data sets
- SMFDUMP job not running to offload SMF data
  - Ensure IEFU29 is specified in the SUBSYS STC option in the SMFPRMxx parmlib member
  - Ensure IEFU29 exit and SMFDUMP job are running to successful completion
  - Manually run IEFU29 job if needed

# SMF Data Loss

**Conclusion**: Why it matters!

- Sometimes “little things” can become “big things”
- **Example**:
  - Customer’s SMFDUMP job was failing repeatedly
  - 1Gig SMF in-core buffer filled up and over time got paged to aux
  - Upon manual run of SMFDUMP, all pages had to be brought back in from aux
    - CMS SMF lock held (global resource) which caused system-wide delay for over 5 minutes

*** SMF APAR OA45637 ***

# Managing RACF Certificates

**Problem**: Time and effort spent in the following tasks:

- The RACDCERT command is used to manage certificates. However, many times one may need to follow a certificate chain, such as:

  Server certificate->signed by intermediate certificate->signed by root certificate

- One may not realize that a certificate is about to expire

# Managing RACF Certificates

**Solution**: The following enhancements in z/OS V2.1 will help:

- The following RACDCERT command will list out the certificate chain (and expiration date too):

  RACFDCERT ID(cert-owner) LISTCHAIN(LABEL('label-name')

- A new health check:
  RACF_CERTIFICATE_EXPIRATION

# SSL Component Trace

**Problem**: When capturing System SSL component trace, the trace dataset is empty

- This problem occurs if the SSL connection is secured by AT-TLS
- For this type of connection the work is performed under TCPIP instead of the application

# SSL Component Trace

**Solution:** When the TRACE CT command is used to capture component trace for System SSL secured by AT-TLS, reply with the jobname of TCPIP

- For native SSL (non AT-TLS), the usual reply is with the jobname of the application calling System SSL

# Dealing with a hung job

**Problem**: what to do if a job is hung?

- Take a dump of the job first (else we would not know why)
- If needed, use Cancel/Force to terminate the job
  - be prepared for an IPL if FORCE is used
- Only use this with the guidance of IBM support:
  - There is a new way in z/OS V2.1 to terminate a TCB in an address space (if this TCB is related to the hang)
  - *See **FORCE** command with new **TCB*** parameter in z/OS V2.1 System Commands
  - Some address spaces do not support the FORCE TCB command

Session 15594

# Common z/OS Problems You Can Avoid

## Thank You!