

Session 16633

Common z/OS Problems You Can Avoid



SHARE in Seattle
March 5th, 2015

Patty Little John Shebey
IBM Poughkeepsie
plittle@us.ibm.com jshebey@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**
- **Runtime Diagnostics: RTD**
- **Dealing with a hung job**
- **zFS Large Directories**
- **System SSL Component Trace**
- **HMC Change Logical Partition Controls**
- **SMF Data Loss**
- **Managing RACF Certificates**

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/atmastr.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

z/OS Best Practices: SADMP

Other considerations

- Problem with using PDUU to FTP SAD?
 - Check for any networking issues in your system
- FTP initially anything that is smaller (and therefore faster)
 - Subset dump
 - Any SVC dumps taken prior to the failure? (check for IEA794I or IEA611E)
 - They may show the problem!

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

Runtime Diagnostics does many of the same tasks you might typically do when looking for a failure, such as:

- Reviewing critical messages in the log
- 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

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.

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

zFS Large Directories

Problem: Poor performance when accessing zFS version 1.4 aggregate (file system) with large directory

- Slow response time or elongated job run time when reading or writing to zFS aggregate
- High CPU in user or system address space (running PRIMARY in zFS) when reading or writing to zFS aggregate

zFS Large Directories

Environment: Following conditions together can lead to problem

- zFS version 1.4 aggregate (v4 directory)
 - Fast Lookup Cache (FLC) buffers not implemented
- Directory being accessed in zFS file system has large number of objects (hundreds of thousands or millions) or is large in size
 - 10,000+ objects in directory
 - or -
 - Size of directory \geq 1024K (1 MB)

zFS Large Directories

Symptoms: Check for evidence of v4 directory size problem

- MODIFY ZFS,QUERY,KN operator command or z/OS UNIX **zfsadm query -knpfs** command
 - Look at **Avg Time** field on lines for operations that require zFS to search through directory objects (i.e., zfs_lookup or zfs_create)
 - If times are larger than a few milliseconds (ten to a hundred times larger), could be dealing with large directory problem

zFS Large Directories

Symptoms (continued):

- **largedir.pl** command (perl script) can be downloaded from z/OS UNIX Tools & Toys (<http://www.ibm.com/systems/z/os/zos/features/unix/bpxa1ty2.html>)
 - Uses **find** command to determine large directories
 - Directories over 1 MB in size
- **NOTE: Space not reclaimed when objects removed from v4 directory, so must look at size, rather than current number of objects**

zFS Large Directories

Solution: At z/OS 2.1, can convert zFS file system to version 1.5 aggregate (v5 directory)

- Systems at earlier z/OS release levels cannot access version 1.5 aggregate
 - **Only use version 1.5 if all systems that access aggregate are running z/OS V2R1 (or higher release)**
- If unable to convert to version 1.5 aggregate, remove large directory, and try to spread objects across multiple directories

zFS Large Directories

Solution (continued):

- Otherwise, performance may be improved for version 1.4 aggregate by creating Fast Lookup Cache (FLC)
 - IOEPRMxx configuration option 'flc'
 - Implemented via zFS APAR OA40530
 - Valid for z/OS V1R13 and z/OS V2R1
- Refer to **z/OS V2R1.0 Distributed File Service zFS Administration** manual for more details about this issue and its mitigation

zFS Large Directories

Problem: zFS dynamic grow delays with similar symptoms to zFS large directory

- Delayed response to users/applications accessing zFS aggregates.
 - Correlated with various zFS aggregate dynamic grow messages (i.e., IOEZ00312I)
- NFS server delays (GFSA1030E and GFSA1033E)
- Problem most predominant with very large zFS
 - 10's or 100's of GB in size
- **HIPER APAR OA46665**

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

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

CSG3HMC1: Change LPAR Controls - Mozilla Firefox: IBM Edition
https://9.23.246.5/hmc/content?taskId=285&refresh=543

Change Logical Partition Controls - MTS2818

Last reset profile attempted: DEFAULT
Input/output configuration data set (IOCDs): A2 IODFD2

CPS ICFs zAAPs IFLs zIIPs Processor Running Time

Logical Partitions with Central Processors

Logical Partition	Active	Defined Capacity	WLM	Current Weight	Initial Weight	Min Weight	Max Weight	Current Capping	Initial Capping	Number of Dedicated Processors	Number of Not dedicated Processors
CMOSVM	Yes	0	<input type="checkbox"/>	49	49			No	<input type="checkbox"/>	0	1
CMOS1	Yes	0	<input checked="" type="checkbox"/>	50	50		60	No	<input type="checkbox"/>	0	2
CMOS2	Yes	0	<input checked="" type="checkbox"/>	20	20			No	<input type="checkbox"/>	0	2
ZOS4	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	3
ZOS5	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
ZOS7	Yes	20	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
ZOS8	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
ZOS9	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
GDPS	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
HAL	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
KHAOS	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
TIVSM	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	4
JES2	Yes	0	<input type="checkbox"/>	10	10			No	<input type="checkbox"/>	0	2

Save to Profiles Change Running System Save and Change Export Reset Cancel Help

An example

HMC – Change Logical Partition Controls

RMF III CPC Report

Host: 9.23.246.229 Port: 23 LU Name: Disconnect

```

HARDCOPY      RMF V2R1      CPC Capacity      Line 1 of 33
Samples: 100   System: ZS21  Date: 07/22/14   Time: 16.32.30   Range: 100   Sec
Partition:    ZOS8      2818 Model V04
CPC Capacity: 177      Weight % of Max: ****   4h Avg: 2   Group: N/A
Image Capacity: 177   WLM Capping %: 0.0     4h Max: 6   Limit: N/A

Partition --- MSU --- Cap Proc Logical Util % - Physical Util % -
           Def  Act  Def  Num  Effect Total  LPAR  Effect Total
*CP                               42.0          0.9          9.7          10.6
CMOSVM           0    0  NO    1.0          0.0          0.0          0.0
CMOS1            0    0  NO    2.0          0.5          0.6          0.3
CMOS2            0    0  NO    2.0          0.5          0.6          0.3
GDPS             0    1  NO    4.0          0.6          0.6          0.6
HAL              0    1  NO    4.0          0.6          0.6          0.6
JES2             0    1  NO    2.0          0.6          0.6          0.3
KHAOS            0    1  NO    4.0          0.7          0.7          0.7
TIVSM            0    2  NO    4.0          0.9          0.9          0.9
ZOS4             0    3  NO    3.0          2.4          2.4          1.8
ZOS5             0    1  NO    4.0          0.5          0.5          0.5
ZOS7             0    1  NO    4.0          0.5          0.5          0.5
ZOS8             0    5  NO    4.0          2.9          2.9          2.9
ZOS9             0    1  NO    4.0          0.5          0.5          0.5
PHYSICAL                               0.7          0.7          0.7

*AAP                               1.0          0.0          0.0          0.0
Command ===>
F1=HELP      F2=SPLIT      F3=END      F4=RETURN      F5=RFIND      F6=toggle
F7=UP        F8=DOWN       F9=SWAP     F10=BREF       F11=FREF      F12=RETRIEVE
MA A
14/077
Connected to remote server/host 9.23.246.229 using lu/pool TNZS2117 and port 23
camkml21-Comm Room E3-03-3600 Steeles A
    
```

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):

```
RACDCERT ID(cert-owner) LISTCHAIN(LABEL('label-name'))
```

- A new health check:
RACF_CERTIFICATE_EXPIRATION