

z/OS SMF Logstream Mode: Update and User Experience

J.O.Skip Robinson
Southern California Edison Co.

7 February 2013
Session 13034



An Unmetered Ode to SMF

- **System **Management** Facility**
- **A powerful function unique to z/OS**
- **Framework for collecting and manipulating structured data produced by system or user components**
- **Header data (type, date/time) prescribed for all types**
- **Remainder documented for each IBM record type**
- **User records defined at will by owner/installation**
- **No action required by applications or subsystems**
- **No allocations, no opening/closing of log/trace**
- **All we have to do is collect, massage, and archive**

Your Daddy's SMF

- For more or less ever, SMF recording was to VSAM
- Multiple MANx clusters sucked up records AFAP
- Periodically each MANx file would be dumped by IFASMFDP immediately or ultimately to tape
- Triggered by indication of MANx data set full
- IEE391A SMF ENTER DUMP FOR DATA SET ON VOLSER SMFC02, DSN=SYS1.MAN2.\$SYSC0
- IEE388I SMF NOW RECORDING ON VOLSER SMFC01, DSN=SYS1.MAN1.\$SYSC0 TIME=
- Data set emptied (ZEROed) after dumping
- Switched to next available MANx data set round robin

Daddy's SMF Headaches

- Lots of problems
- MANx data sets filled up fast on a busy system
- Out of control tasks could swamp all MANx data sets
- Problems with tape could prevent normal dump/clear
- Even with emergency buffering, SMF data could be lost
- No prioritization of record types possible
- Records processed FIFO without regard to customer's needs
- If any data lost, all data lost

System Logger Rides to the Rescue

- **Solution in z/OS R9 was SMF Logger**
- **MANx replaced by system log stream**
- **Either supported by CF structure or DASD-only**
- **We use CF supported logstreams only (so far)**
- **Logstream solves many chronic problems**
- **Much faster than writing to MANx VSAM**
- **Less chance of losing data due to high rate**
- **Data directed to multiple log streams by SMF type**
- **But data must be managed somewhat differently**

Managing SMF Data Gangnam Style

- **New management program IFASMFDL**
- **Original design required start/end dates**
- **Now IFASMFDL provides ARCHIVE with no DATE parameter**
- **Simply dumps all records not previously ARCHIVED**
- **Starts with oldest non-ARCHIVED records**
- **Marks dumped records as 'ARCHIVED'**
- **Stops at current time plus...**
- **Includes 'SMART' keywords for moving end time**
- **Also adds diagnostic information for debugging**

Gangnam Style - 2

- Management of data now similar to old MANx style
- But you still get logger speed and flexibility
- I recommend...
- One day retention (minimum allowed) in log stream definition
- Multiple ARCHIVE jobs per day (at least two)
- In sum: management of MANx data sets is driven by data
- Management of Logger data is driven by time
- Current implementation similar to old IFASMFDP

Gangnam Style - 3

- **We run two dump jobs daily, midnight and noon**
- **Midnight job performs traditional SMF data processing**
- **Combines data from multiple systems, creates multiple outputs**
- **Produces reports for various subsystems, stakeholders**
- **Midday job just archives for next midnight job**
- **Purpose: make sure data is not deleted after extended outage**
- **Data marked as archived by SMF dump job**
- **Data retained/deleted by system logger**

D SMF Old School Style

- **IEE974I 15.22.16 SMF DATA SETS**
- **P-SYS1.MAN1.\$SYSX0 SMFX01 67500 84**
ACTIVE
- **S-SYS1.MAN2.\$SYSX0 SMFX02 67500 0**
ALTERNATE
- **Similar lines for each MANx data set**
- **Status changes as each one fills up or gets emptied**

SMF Logstream Status

- **IFA714I 15.22.48 SMF STATUS**
- **LOGSTREAM NAME BUFFERS STATUS**
- **A-IFASMF.DEFAULT 21524 CONNECTED**

- **ISPF 3.4 for LOGR.IFASMF shows all offload data sets**
- **LOGR.IFASMF.DEFAULT.A0096197 *VSAM***
- **LOGR.IFASMF.DEFAULT.A0096197.DATA ASMF01+**

- **Each data set is 2,745 tracks (183 cylinders)**
- **Size determined by logstream definitions**
- **Count rises and falls throughout the day**

Set Recording Mode in PARMLIB

- **PARMLIB SMFPRMxx**
- **ACTIVE**
- **DSNAME(SYS1.MAN1.\$SYS&SYSCLONE,**
- **SYS1.MAN2.\$SYS&SYSCLONE,**
- **SYS1.MAN3.\$SYS&SYSCLONE)**
- **RECORDING(LOGSTREAM)**
- **/* RECORDING(DATASET) /***
- **DEFAULTLSNAME(IFASMF.DEFAULT)**

- **Note: SMFPRMxx members cannot be concatenated**

Defining SMF Logstream - 1

- **//LOGRSMF EXEC PGM=IXCMIAPU**
- **//SYSPRINT DD SYSOUT=***
- **//SYSIN DD ***
- **DATA TYPE(LOGR)**
- **DEFINE STRUCTURE**
- **NAME(IFASMF_DEFAULT)**
- **LOGSNUM(5)**
- **AVGBUFSIZE(32767)**
- **MAXBUFSIZE(65532)**

Defining SMF Logstream - 2

- **DEFINE LOGSTREAM**
- **NAME(IFASMF.DEFAULT)**
- **HLQ(LOGR)**
- **STRUCTNAME(IFASMF_DEFAULT)**
- **LS_SIZE(32768)** ← Determines size of offload data set
- **LS_DATACLAS(SMFLOG)**
- **STG_DATACLAS(SMFLOG)**
- **STG_DUPLEX(YES)**
- **DUPLEXMODE(UNCOND)**
- **AUTODELETE(YES)**
- **RETPD(1)** ← Determines number of offload data sets

Logstream Archive Job

- **//ARCHLOGR EXEC PGM=IFASMFDL**
- **//OUTSMF DD DSN=LOGR.SMFDATA(+1),**
- **// DISP=(CATLG,DELETE),UNIT=TAPE,**
- **// LRECL=32760,RECFM=VBS,BLKSIZE=4096,**
- **//SYSPRINT DD SYSOUT=***
- **//SYSIN DD ***
- **LSNAME(IFASMF.DEFAULT,OPTIONS(ARCHIVE))**
- **OUTDD(OUTSMF,TYPE(0:255))**
- **SMARTENDPOINT /* R13 AND UP */**
- **SMARTEPOVER(0100) /* TWICE THE VALUE OF MAXDORM */**
- **//PROCSMF EXEC ... Process LOGR.SMFDATA() as before**

Whither Offload Data Sets?

- Location of offload data sets determined by SMS rules
- 1. Restricted: go only to defined pool (one or more specific volumes)
- 2. Contained: go to defined pool but allow overflow elsewhere
- 3. Free range: go to any available volume, i.e. SYSALLDA
- With (1), you risk losing data (see below)
- With (3), you risk flooding the whole DASD farm
- I recommend option (2) unless space problems occur
- When you're all done, you get back most of the old MANx space
- Keep a few smallish MANx guys for emergency fallback to DS
- If you do have to fall back temporarily, be sure to dump MANx

What Could Possibly Go Wrong?

- Early on, IFASMFDL would choke on 'bad data'
- Especially invalid time stamp, other errors possible
- IFASMFDL would end badly with no diagnostics
- SMF data continued to accumulate in offload data sets
- Eventually DASD would fill up, causing buffering
- Finally SMF disconnected from log stream ;-(((
- IFA714I 15.22.48 SMF STATUS
- LOGSTREAM NAME BUFFERS STATUS
- A-IFASMF.DEFAULT 9999999 **DISCONNECTED**
- Try switching to MANx, then back to logstream
- **SMF Logstream has been very reliable for some time**

Questions?

