



SMFPRMxx Parameters – Which can Help; Which can Hurt

Cheryl Watson (cheryl@watsonwalker.com) &
Frank Kyne (frank@watsonwalker.com)

Watson & Walker, Inc.
www.watsonwalker.com

Session 18034
August 12, 2015



#SHAREorg



SHARE is an independent volunteer-run information technology association
that provides **education, professional networking and industry influence.**

Copyright (c) 2015 by SHARE Inc. Except where otherwise noted, this work is licensed under
<http://creativecommons.org/licenses/by-nc-sa/3.0/>



Abstract



Every installation depends on their SMF data, and the collection of that data is controlled by the SMFPRMxx member of SYS1.PARMLIB. Most people set the parameters for SMFPRMxx once and never look back. But times change, and some of the old values may actually be hurting you now. And there are many parameters that you probably don't even know exist, but they can significantly increase your performance. Cheryl and Frank will provide their recommendations for best practices when setting these parameters.

Welcome



- Hi, thank you for coming.
- Who are we and what we do.
 - Three-person company; started in 1987
 - Frank Kyne joined us in 2014 from ITSO
 - Quarterly subscription-based newsletter –
Cheryl Watson's Tuning Letter
Cheryl Watson's System z CPU Chart
 - Classes on z/OS new features, WLM, performance, software pricing, chargeback, Parallel Sysplex, and high availability
 - Consulting on all the above topics
 - Software products – GoalTender and BoxScore

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

© Copyright Watson & Walker 2015



Welcome



- Education
 - z/OS Software Pricing Strategies
 - October 12-14, Orlando, FL
 - November 11-13, Hamburg, Germany
 - Exploiting New Features of z/OS To Minimize Costs
 - November 18-20, Hamburg, Germany
 - SYS1.PARMLIB: Enhancements and New Features in z/OS 1.13, 2.1, and 2.2
 - November 16, Hamburg, Germany

- All classes are available on-site

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Agenda

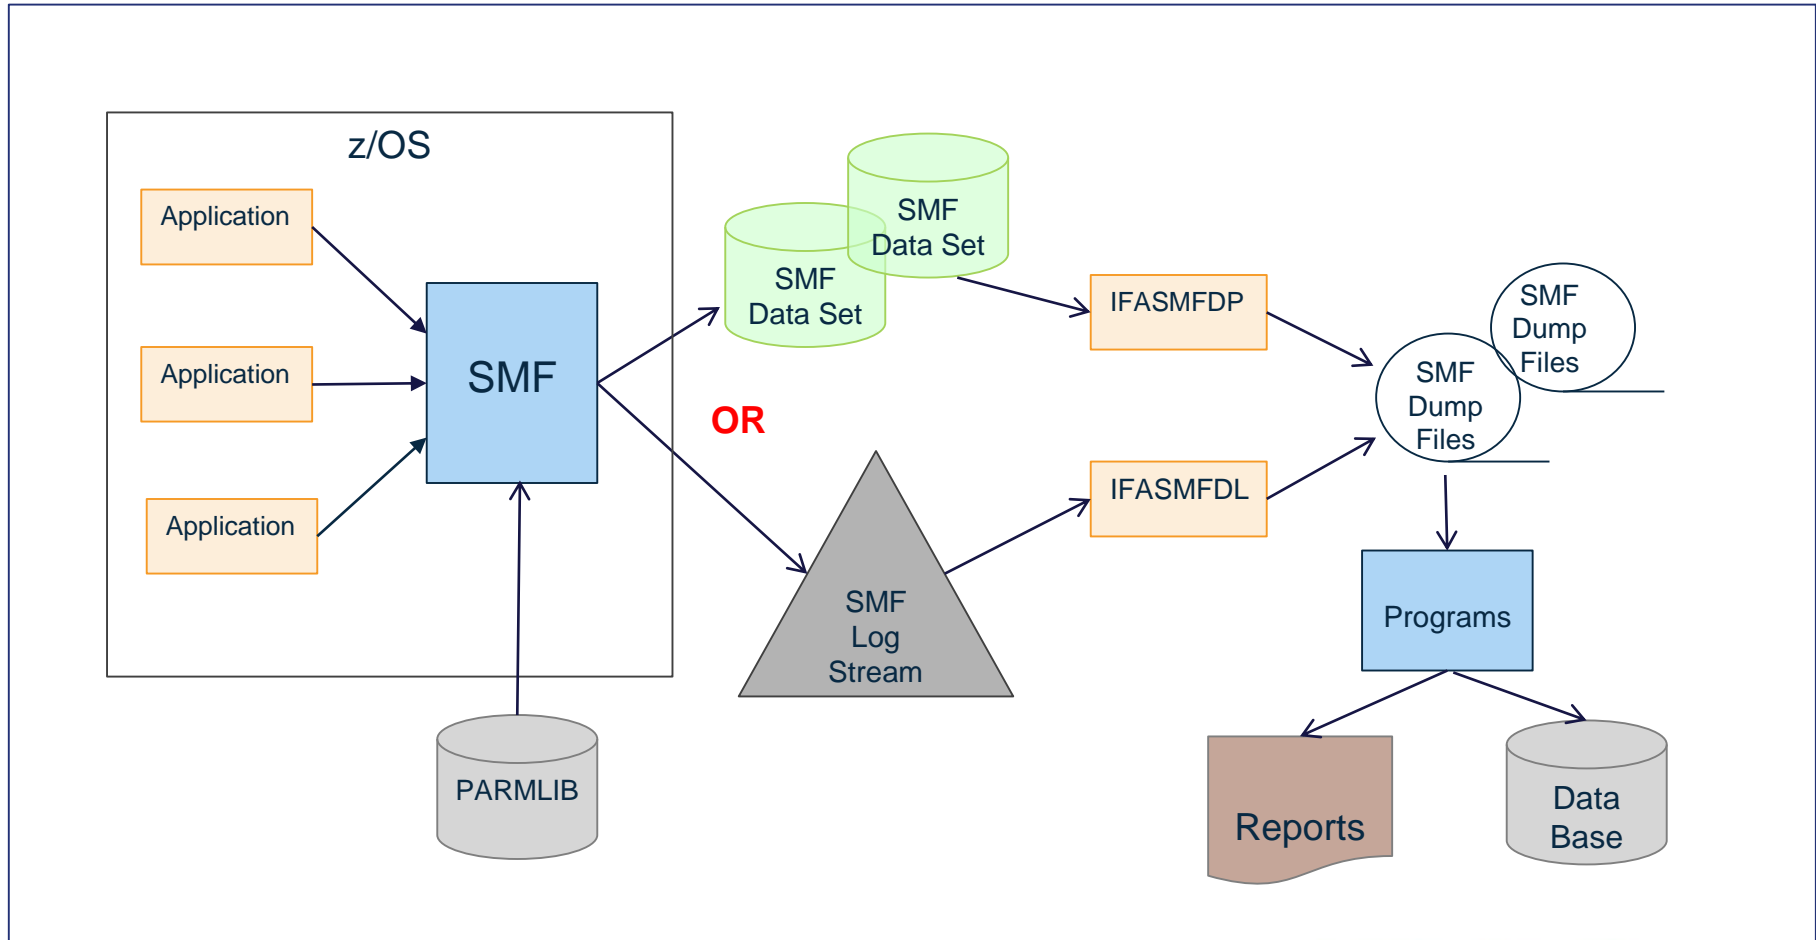
- **What is SMF?** ←
- SMF Data Sets
- SMF Logger
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

What is SMF?

- Primary manual is *z/OS MVS System Management Facilities (SMF)* – [SA22-7630](#) (z/OS 1.13) or [SA38-0667](#) (z/OS 2.1/2.2)
- Major elements:
 - Started task, SMF, created during IPL
 - Parameters specified in PARMLIB member SMFPRMxx
 - Macros – used by applications to pass data to SMF and to interrogate parameters
 - Exits – User-written exits to interrogate, change, or delete records as they are passed to SMF
 - MVS commands – used to control the execution of SMF and to dynamically modify parameters

What is SMF? Logic Flow



SMFPRMxx

- Purpose: This Parmlib member defines which record types are written to an SMF VSAM data set or SMF log stream. It also defines which exits are called to modify or look at SMF records. A global interval and synchronization are defined for use by both SMF and other subsystems.
- Most SMF parameter defaults are great and should be used. Our recommended changes to override defaults **are shown in red** throughout this presentation.
- Use system symbols to create a single SMFPRMxx that is used for all systems.

Agenda

- What is SMF?
- **SMF Data Sets** ←
- SMF Logger
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

SMF Data Sets

- These parameters apply to both VSAM data sets and logstreams

BUFUSEWARN(25|nn)

Issue warning message (IEE986E) when buffers reach 25% full

NOBUFFS(MSG|HALT)

Action to be taken when buffers fill up

MAXDORM(3000|0100|1000)

Number of minutes to let records sit in buffer before being forced to be written.

We recommend 10 minutes for VSAM data sets and 1 minute for logstreams.

SMF VSAM Data Sets

- Traditional recording method
- VSAM data sets that are pre-formatted
- SMF writes ALL requested records to ONE of the VSAM data sets
- When that data set fills up, it calls an exit and switches to another formatted VSAM data set
- The IEFU29 exit usually issues a message to the operator and submits a job (usually IFASMFDP) to dump the records to a sequential data set and clear the VSAM data set
- The VSAM data set names are specified in SMFPRMxx (referred to as 'MANx' data sets)
- Maximum buffer size is 1 GB

SMF VSAM Data Sets

- There are several problems with SMF VSAM data sets:
 - Data sets fill up and aren't cleared, so data is lost!
 - The data sets can be easily overlaid, so data is lost!
 - Large volumes can overflow the buffers and/or the data sets, so data is lost!
 - Many sites have turned off some record types in order to reduce the volume
 - But a runaway or looping application can produce hundreds of thousands of records in a very short amount of time. This can cause overflows despite a robust configuration
 - Lost data can translate to lost revenue
 - It can take a long time to format a new SMF data set

SMF VSAM Data Sets

- File-related parameters (VSAM data sets):

RECORDING(DATASET/LOGSTREAM)

Indicates whether SMF writes records to VSAM data sets or logstreams

LISTDSN/NOLISTDSN

Determines whether SMF data sets and status are listed on the console at startup and SET SMF time

BUFSIZMAX(xxxM/1G/128M)

Specifies the max amount of buffer space to be used in the SMF STC when recording to VSAM data sets; **most sites need 1G**

SMF VSAM Data Sets

- File-related parameters (VSAM data sets):

LASTDS(MSG/HALT)

Action if no more data sets are available

DSNAME(SYS1.MANX,SYS1.MANY/datasetnames)

Defines default data set names

Recommendations:

Use system symbols, such as SYS1. &SYSID..MAN1;

Use more than two MANx data sets

DUMPABND(RETRY/NORETRY)

Defines action if SMF dump program abends.

Agenda

- What is SMF?
- SMF Data Sets
- **SMF Logger** ←
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

SMF Logger

- SMF Logger was introduced in z/OS 1.9
- SMF writes records to its dataspace buffer and also to a 'logstream' that resides in a Coupling Facility (CF) structure or a DASDONLY logstream
- Multiple LPARs can record to the same structure or logstream if using a CF
- You can record different record types to different logstreams if you wish:
 - This allows a higher write rate
 - Can reduce the post-processing time
- Initially there were some problems with managing the records, but we believe those are all resolved

SMF Logger

- SMF Logger can avoid the problems indicated before
 - Buffers are in 2 GB dataspace (one per logstream rather than one in total), so that buffers don't overflow
 - Each dataspace is managed by its own task, so that the write rate is increased
 - Logstreams are offloaded to DASD files as needed, so that you can't run out of space
 - The ability to separate record types onto different structures and logstreams can increase the write rate, as well as reducing post-processing time. This also reduces the problem of one runaway application causing other applications to lose data
 - The typical midnight SMF offload time can be decreased by dumping multiple logstreams simultaneously.

- Additional benefits:
 - The SMF post-processors that use data from multiple systems are much easier to run. When using a CF logstream containing SMF records from multiple systems, the output DASD data set can contain data from all of the systems — there is only one data source as input.
 - SMF records can be written to multiple log streams, so that there is less opportunity to lose the data. This does, however, involve some amount of overhead.
 - While it can take several minutes to CLEAR a MANx data set, this is never needed with SMF logger.

- Logstream-related parameters:

RECORDING(LOGSTREAM)

Indicates whether SMF writes records to VSAM data sets or logstreams

**DEFAULTLSNAME(logstreamname),NOBUFFS(MSG/HALT),
BUFUSEWARN(pct/25),DSPSIZMAX(nnnnM/nG/2G),
COMPRESS(PERMFIX(nnnnM)/NOPERMFIX)**

Defines the logstream that is used for any SMF record that is not sent to some other logstream.

**LSNAME(logstreamname,NOTYPE/TYPE(xx,xx),
NOBUFFS(MSG/HALT),BUFUSEWARN(pct/25),
DSPSIZMAX(nnnnM/nG/2G),
COMPRESS(PERMFIX(nnnnM)/NOPERMFIX)**

Indicates record types and options for each logstream

SMF Logger

- Logstream-related parameters:
[DSPSIZMAX\(nnnnM/nG/2G\)](#)
 - Minimum is 12M; default is 2G;
 - Because logger buffers are in data spaces in disabled reference (DREF) storage, which is NOT pageable, the use of 2G for many logstreams could cause storage shortages.
 - At initialization, SMF gets 100 buffers for each log stream (default system logger buffersize is 64KB; high activity can cause buffers to expand up to max specified in DSPSIZMAX)
- Use default initially, but monitor the space usage with the SMF Type 23 data: SMF23LFA (buffer size), SMF23LFT (buffer storage used), SMF23LFH (max buffer allocation), and SMF23LFM (DSPSIZMAX)

Recommendation

- Migrate from VSAM data sets to SMF logstreams now
 - All future enhancements (such as zEDC and SMF record signing (z/OS 2.2)) will ONLY be to logstreams
 - Start with defaults then modify after you get some experience
 - Use multiple logstreams to separate SMF records based on how you use them

Agenda

- What is SMF?
- SMF Data Sets
- SMF Logger
- **Records, records, records!** ←
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

Records, Records, Records

- Address Space
 - 4, 5, 20, 34, 35, 50 - obsolete
 - **6 – Print spool**
 - 6.0 – External writer
 - 6.2 – JES2 output writer
 - 6.5 – JES3 output writer
 - 6.7 – Print Services Facility (PSF) output
 - 6.9 – IP Printway output
 - 24 – JES2 spool offload
 - 25 – JES3 device alloc.
 - **26 – Job purge**
- Address Space (more)
 - **30 – Job processing**
 - 30.1 – Job initiation
 - 30.2 – Step interval ends
 - 30.3 – Last step interval ends
 - 30.4 – Step ends
 - 30.5 – Job ends
 - 30.6 – System STC ends
 - 32 - TSO/E commands
 - 33 – APPC/MVS activity
 - 43-49, 52-58 – JES activity
 - 97 – Foreign Enclaves

Records, Records, Records

- Volume Configuration
 - 8 – Configuration
 - 9 – Vary device online
 - 10 – Allocation recovery
 - 11 – Vary device offline
 - **19 – Direct access volume**
 - 21 – Error statistics by (tape) volume
 - 22 - Configuration
- Volume Activity
 - **74 – RMF Device Activity**
 - 74.1 – Device activity
 - 74.2 – XCF activity
 - 74.4 – CF Activity
 - 74.5 – Cache subsystem activity
 - 74.7 – FICON director statistics
 - 74.8 – Enterprise disk system statistics
 - 78.3 – RMF I/O Queuing
 - 79 – RMF Monitor II Activity
 - 79.7 – Enqueue contention
 - 79.9 – Device activity
 - 79.12 – Channel path activity
 - 79.14 – I/O queuing
 - 94 – IBM Tape Library Dataserver Statistics

Records, Records, Records

- VSAM Activity
 - 36 – ICF catalog export
 - 41 – DIV object & VLF statistics
 - 60 – VSAM Volume Data Set updated
 - 61 – Data set added or updated in ICF catalog
 - 62 – VSAM component or cluster opened
- VSAM Activity (cont.)
 - **64 – VSAM component or cluster closed**
 - 65 – Data set deleted from ICF catalog
 - 66 – IDCAMS ALTER request is made
- Operator & Events
 - **90 – Almost 35 subtypes**
 - **113 – Hardware measurements**

Records, Records, Records

- Specialty Data Sets (HFS, zFS, OAM)
 - 42 – DFSMS Statistics
 - 42.1 / 42.5 – Storage class
 - 42.2 – Cache control unit
 - 42.3 – SMS configuration change
 - 42.4 – System Data Mover
 - **42.6 – Data set I/O statistics**
 - 42.7 / 42.8 – NFS file activity
 - 42.9 – Out of space ABEND (Sx37)
 - 42.10 – Allocation volume selection failure
 - 42.11 – XRC interval statistics
 - 42.14 – ADSM accounting
 - 42.15 – 42.19 – VSAM RLS
 - 42.20, 42.21, 42.24, 42.25 – PDS/PDSE activity
 - 42.22 / 42.23 – DFSMSrmm
- 74.6 – RMF HFS
- 77 – RMF Enqueues
- 85 – OAM Transaction Performance
- 92 – File System Activity for UNIX files
 - **Turn on only for special studies:**
 - 92.5, 92.6, 92.7 – Unmounts, remounts, moves
 - 92.11 – File close

Records, Records, Records

- Non-VSAM Data Set Activity
 - **14 – Input data set close**
 - **15 – Output data set close**
 - 17 – Data set deleted
 - 18 – Data set renamed
- Workload Manager
 - 99 – WLM activity
 - 99.6 – Service class period summary (most useful)
 - 99.10 – Processor changes
- Security Records
 - 80 – Security product processing – identifies unauthorized attempts
 - 81 – RACF initialization
 - 82 – Crypto
 - 82.2 – Cryptographic Unit Support Program (CUSP)
 - 82.1/3-29 – Integrated Cryptographic Service Facility (ICSF)
 - 82.1 – PCF
 - 83 – RACF audit record for data sets

Records, Records, Records

- RMF Records
 - **70.1 – CPU Activity**
 - 70.2 – Crypto Activity
 - 71 – Paging Activity
 - **72.3 – Workloads (Service classes)**
 - 72.4 – Storage
 - 72.5 – Serialization Delay
 - 73 – Channel Path Activity
 - 74.1 – Device Activity
 - 74.2 – XCF Activity
 - 74.3 – OMVS Kernel Activity
 - 74.5 – Cache Subsystem
 - 74.6 – HFS Statistics
 - 74.7 – FICON Director Stats
 - 74.8 – Enterprise Disk System Stats
 - 75 – Page Data Sets
 - 76 – RMF Trace Activity
 - 77 – Enqueue
 - 78.2 – Virtual Storage
 - 78.3 - I/O Queuing
 - 79 – RMF Monitor II records (12 subtypes)

Records, Records, Records

- Subsystems
 - **110 – CICS**
 - 111 – CICS Transaction Gateway
 - **100 – DB2 Statistics**
 - **101 – DB2 Accounting – can account for 75% of all SMF**
 - 102 – DB2
 - 108 – Domino Server
 - 109 – TCP/IP
 - **118** / 119 – FTP; use 119 instead of 118
- WebSphere, MQ, MB
 - 103 – HTTP Server
 - 115 – MQSeries Stats
 - **116 – MQ Accounting**
 - 117 – Message Broker
 - **120 - WAS**

All of these are high volume SMF records!

Records, Records, Records

- What you should learn first
 - 6, 26, 30 – Address spaces
 - 14, 15, 42.6, 64 – Data sets
 - 70 – CPU activity
 - 72 – Service class activity
 - 90 – Operator events
 - 113 – Hardware Measurements
 - Subsystems that represent critical applications (e.g. CICS, DB2)

Records, Records, Records

- Consider record types when creating logstreams
 - **LSNAME(logstreamname,NOTYPE/TYPE(xx,xx) . . .**
 - Consider separate logstreams for each of the following (based on how the data will be used and who will be accessing it):
 - All RMF records, device records, volume records, data set records
 - CICS records
 - DB2 records
 - MQ records
 - WebSphere records
 - Security records
 - Long-term records, such as type 89, 70
 - Vendor/ISV records (128:255)
 - Everything else (DEFAULTLSNAME)

Agenda

- What is SMF?
- SMF Data Sets
- SMF Logger
- Records, records, records!
- **General Parameters** ←
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

General Parameters

- General parameters

 - ACTIVE/NOACTIVE

 - Activate record collection

 - INTVAL(mm/30)

 - Defines the default global interval for any application that does interval recording

 - SYNCVAL(mm/00)

 - Specifies which minute in the hour the global interval will start on; Used for types 23, 30, 32, 42, 70-79, 89, 108, 113, 119, CICS PA, DB2, etc;

 - Default is 00, but some sites still use SYNCVAL(59)

- Why we like these defaults:

 - Most sites use them
 - Length of interval has proven very usable
 - Smaller interval incurs more overhead
 - Longer interval smooths out spikes of activity
 - You REALLY should use the same value on every system in your sysplex.

General Parameters

- General parameters (cont.):

DUMPABND(RETRY/NORETRY)

Action to take if IFASMFDP abends (does not apply to logstream mode)

JWT(hhmm/0010/0030) (also SWT and TWT)

Maximum time that job (or STC or TSO user) can be in a wait

JWT includes tape mounts and HSM recalls, so most sites needed to use 30 minutes, which is too long for TSO. See z/OS 2.1 recommendation later.

SID(yyyy)/(yyyy,SYSDNAME(sysname))/(yyyy,ser#,ser#)/ (yyyy,COMBIN(ser#,ser#))

Indicates the 4-character system id that will be placed in every record

Better to do filtering in IEASYMxx and create SID value there

Recommendation: SID(&SYSDNAME(1:4)) (or &SID as defined in IEASYM)

SUBPARM(name(parm))

Passes the 'parm' value to the named subsystem

(Does anyone use this??)

General Parameters

- General parameters (cont.):

PROMPT(IPLR/LIST/ALL)/NOPROMPT

Operator options when SMF starts; IPLR=issue WTOR for IPL reason;
LIST=List all the SMF options; ALL=IPLR+LIST

PROMPT forces operator intervention during IPL, which most sites don't appreciate;

NOPROMPT or PROMPT(IPLR) won't allow the operator to use the SETSMF command for the life of the IPL.

AUTHSETSMF/NOAUTHSETSMF

AUTHSETSMF says operator *can* use SETSMF *regardless* of PROMPT setting.

NOAUTHSETSMF says operator can NOT use SETSMF regardless of PROMPT setting

Note: If you want SMF parameters displayed on the system log at IPL, simply add a display command of 'D SMF,O,S' in the COMMNDxx member

General Parameters

- General parameters (cont.):

MEMLIMIT(xxM/xxG/xxT/xxP/2G)/NOMEMLIMIT

Defines maximum amount of storage above 2GB that can be used by an address space if MEMLIMIT was not specified on the JCL

SMFDLEXIT(USER1(exit,exit)/NOUSER1),
(USER2(exit...)/NOUSER2),(USER3(exit...)/NOUSER3)

SMFDPEXIT(USER1(exit,exit)/NOUSER1),
(USER2(exit...)/NOUSER2),(USER3(exit...)/NOUSER3)

These control which IFASMFD_P and IFASMFD_L exits can be used, and the names of the exit load modules. Defaults are:

SMFDLEXIT(USER2(IRRADU00),USER3(IRRADU86))

SMFDPEXIT(USER2(IRRADU00),USER3(IRRADU86))

We **STRONGLY** recommend looking at the IFASMFD_x exits provided in the additional material (Appendix E) for IBM SMF Redbook [SG24-7919](#). See slide 65.

Agenda

- What is SMF?
- SMF Data Sets
- SMF Logger
- Records, records, records!
- General Parameters
- **Record-Specific Parameters** ←
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

Record-Specific Parameters

- SYS

SYS(TYPE/NOTYPE(xx,xx:xx,xx(xx),0:255),

Defines which records are to be written for all types of work (if not overridden)

INTERVAL(hhmmss/SMF,SYNC/NOSYNC)/NOINTERVAL,

Defines the interval at which interval records should be created for all types of work (if not overridden) for the Type 30 (job-related) and Type 32 (TSO) records.

Recommend using the same value on all systems in the sysplex if possible.

EXITS(. . .)/NOEXITS,

Defines the interval at which interval records should be created for all types of work (if not overridden) for the Type 30 (job-related) and Type 32 (TSO) records.

DETAIL/NODETAIL)

For Type 30 records, defines whether the EXCP sections are created in interval records for started tasks. For Type 32 records, indicates whether resource usage (CPU, EXCP, etc.) are recorded for each TSO command (if no SUBSYS override).

Record-Specific Parameters

- SUBSYS

SUBSYS(xxx,TYPE/NOTYPE(xx,xx:xx,xx(xx),0:255),

Defines parameters that override SYS statement for a subset of workloads; **xxx can be: TSO, STC, JES2, JES3, ASCH, OMVS**

INTERVAL(hhmmss/SMF/SMF,SYNC/NOSYNC)/NOINTERVAL,

EXITS(. . .)/NOEXITS,

DETAIL/NODETAIL)

Record-Specific Parameters

EMPTYEXCPSEC(NOSUPPRESS/SUPPRESS)

Indicates whether empty EXCP sections in the Type 30 record are recorded. Empty EXCP sections are generated for each non-allocated SMS candidate volume per DD statement. They are also generated as a result of non-dataset allocations, such as for DD DUMMY or spool file allocations. You get smaller records with no loss of data if SUPPRESS is specified.

MULCFUNC/NOMULCFUNC

Determines whether subsystems, such as CICS and DB2, record resource usage information in the Type 89 records (for Measured Usage License Charge). MULC is not supported with sub-capacity pricing, so for most customers, there is no value in collecting these records. Only collect if you use MULC.

REC(ALL/PERM)

Used for Type 17 records (delete of a data set) to determine whether temporary data set deletions are recorded

Record-Specific Parameters

STATUS(hhmmss/010000/SMF,SYNC/NOSYNC)/NOSTATUS

Defines frequency to create Type 23 (SMF status) records. Using a value not larger than 30-minutes provides more valuable information for tuning SMF buffers.

MAXEVENTINTRECS(00/nn/60)

Indicates the maximum number of processor change events (in an SMF interval) that will cause new Type 30 and Type 89 records to be written. The default is 0, meaning that records will only be written at the original intervals, thus invalidating the value of the CPU times in these records if there was a processor change during that interval.

Record-Specific Parameters

DDCONS(YES/NO)

Consolidate duplicate Type 30 EXCP sections before writing the interval/step end/job end records; YES results in smaller records, but potentially takes more CPU and elapsed time at interval/step/job termination.

Interesting note: In 1991, customers complained when using SYNC on interval recording. Problem was that CICS and DB2 were creating very large records due to the high number of EXCP sections and consolidating them was using extra CPU time; IBM recommended setting DDCONS(NO) and problem was resolved. So performance recommendation was to set DDCONS(NO) for the last 20+ years. BUT in 2009, Cheryl and I ran tests while working on the Mean Time to Recovery Best Practices Redbook ([SG24-7816](#)) to see if DDCONS(NO) would shorten the DB2 shutdown process, and it had no effect. The question still remains as to whether it impacts other workloads during consolidation. One recent customer experience said that they improved CICS response time during interval recording by specifying NO.

Record-Specific Parameters

- SMF Record Flood Control:
 - Many customers experience problems when some SMF record creator intermittently generates massive numbers of records, completely filling the SMF buffers. When the buffers fill, subsequent records will be discarded. You have no way to know which records will arrive next, so you don't know which ones will be discarded – they might be unimportant, or they might be vital.
 - To give you more control in such situations, SMF introduced a function known as SMF record flood control.
 - It lets you specify a policy indicating (at the record type level) an arrival rate above which those records can be discarded, thereby protecting the buffers for your more critical records.

Record-Specific Parameters

- SMF Flood Control parameters:

FLOOD(OFF/ON)

Indicates whether SMF should activate SMF flooding control

FLOODPOL(TYPE(aa),RECTHRESH(bbbbb),
INTVLTIME(ccccc in tenths-of-seconds),
MAXHIGHINTS(dddd),
ENDINTVL(eeeee in tenths-of-seconds),
ACTION(MSG/DROP))

Defines a flood policy for specific types of records, their allowable rates, and actions when they exceed their rates; MSG issued is IFA780I.

A policy reads like this: If record type aa creates more than bbbbb records per cccc interval and lasts for dddd intervals, then either issue a message or drop the record until there are fewer than bbbbb records per cccc interval for a period of eeeee.

Record-Specific Parameters

- **RECOMMENDATION:**
 - SMF record flooding is more likely to be an issue if you are still using MANx data sets.
 - Add the FLDSTATS(nnnn) parameter to IFASMFDP/L statements to get a report showing the arrival rate for each record type.
 - For more info, see page 642 of [SG24-7853](#) z/OS 1.12 Implementation Guide.
 - Regardless, everyone should set up an SMF flood policy, identifying the records that you are willing to sacrifice if the buffers start filling.
 - Customer comment – “We are using SMF logger files and SMF message flooding and we have not lost one SMF record since the conversion.” Do you know how many records you have lost?

SMFPRMxx

- **Our recommendations (only code changes to defaults):**

```
SID(&SYSNAME(1:4)),  
NOPROMPT,AUTHSETSMF,  
DDCONS(NO),EMPTYEXCPSEC(SUPPRESS),  
MAXEVENTINTRECS(60),  
SWT(0010),TWT(0010), JWT(30),          /* zOS 2.1 only */  
STATUS(SMF,SYNC),  
MAXDORM(0100),  
RECORDING(LOGSTREAM),  
DEFAULTLSNAME(IFASMF.&SYSPLEX..DEFAULT),  
LSNAME(IFASMF.&SYSPLEX..DB2,TYPE(100:102)),  
LSNAME(IFASMF.&SYSPLEX..OLTP,TYPE(110:111,108:120)),  
LSNAME(IFASMF.&SYSPLEX..ISV,TYPE(128:255)),  
SYS(NOTYPE(4,5,19,20,34,35,40,92,118),INTERVAL(SMF,SYNC)),  
FLOOD(ON),  
FLOODPOL(TYPE(00:255),RECTHRESH(1000),  
    INTVLTIME(10),MAXHIGHINTS(20),ENDINTVL(300),ACTION(MSG)),  
FLOODPOL(TYPE(aaa),RECTHRESH(1000),  
    INTVLTIME(5),MAXHIGHINTS(30),ENDINTVL(300),ACTION(DROP))  
FLOODPOL(TYPE(bbb),RECTHRESH(2000),  
    INTVLTIME(5),MAXHIGHINTS(30),ENDINTVL(300),ACTION(DROP))
```

Agenda

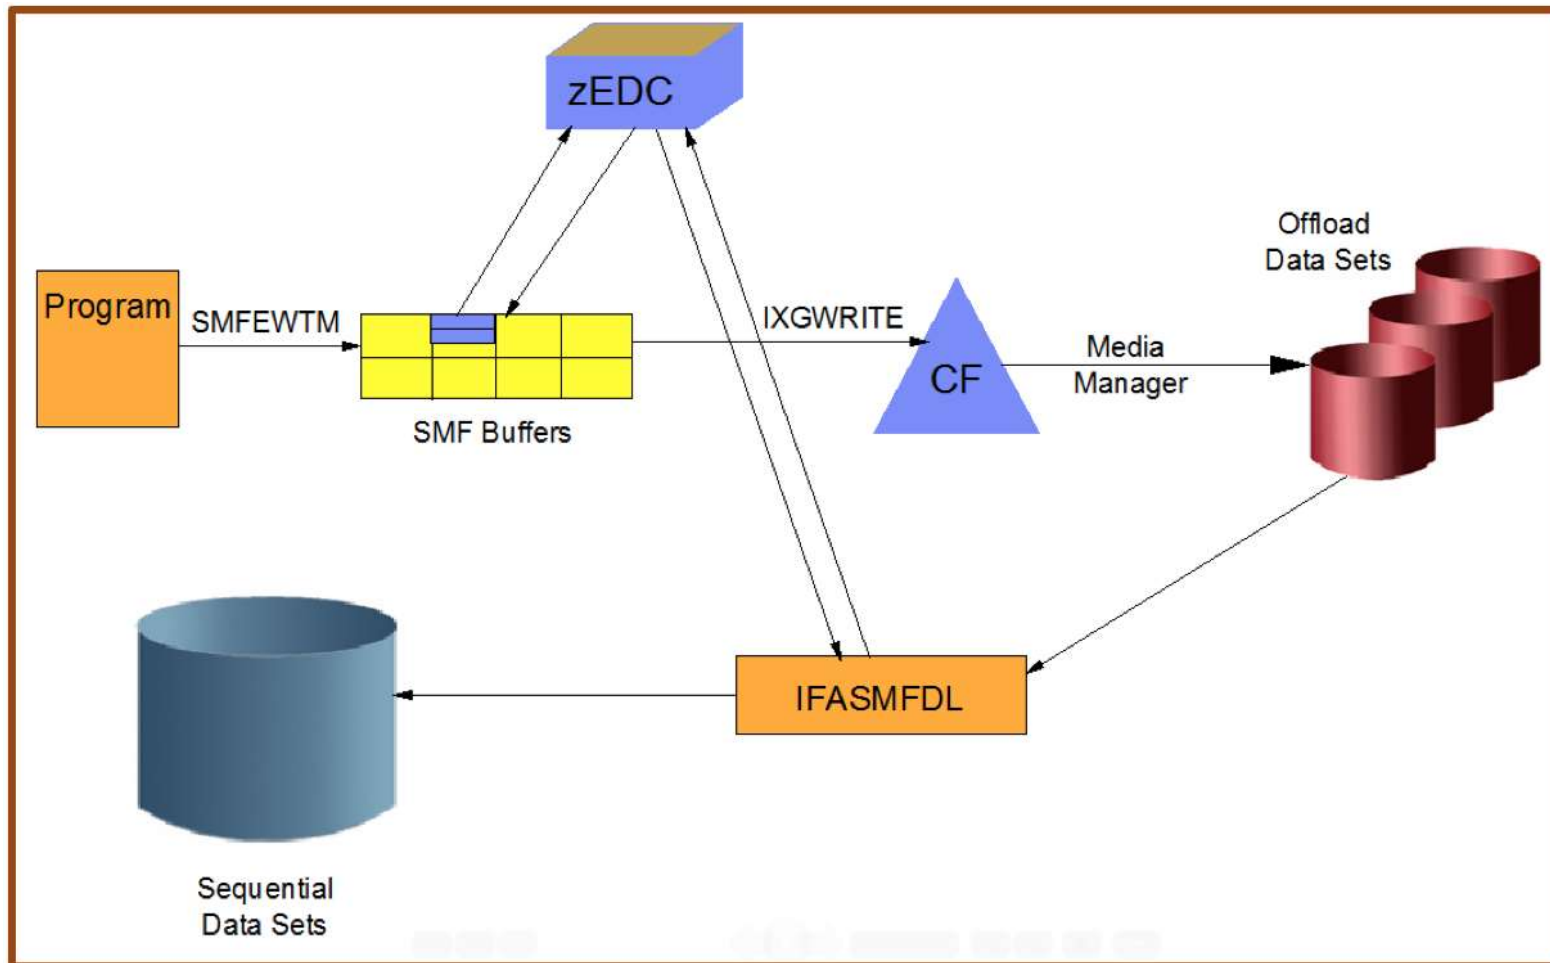
- What is SMF?
- SMF Data Sets
- SMF Logger
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- **z/OS 2.1 Enhancements** ←
- z/OS 2.2 Enhancements
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

z/OS 2.1 - zEDC

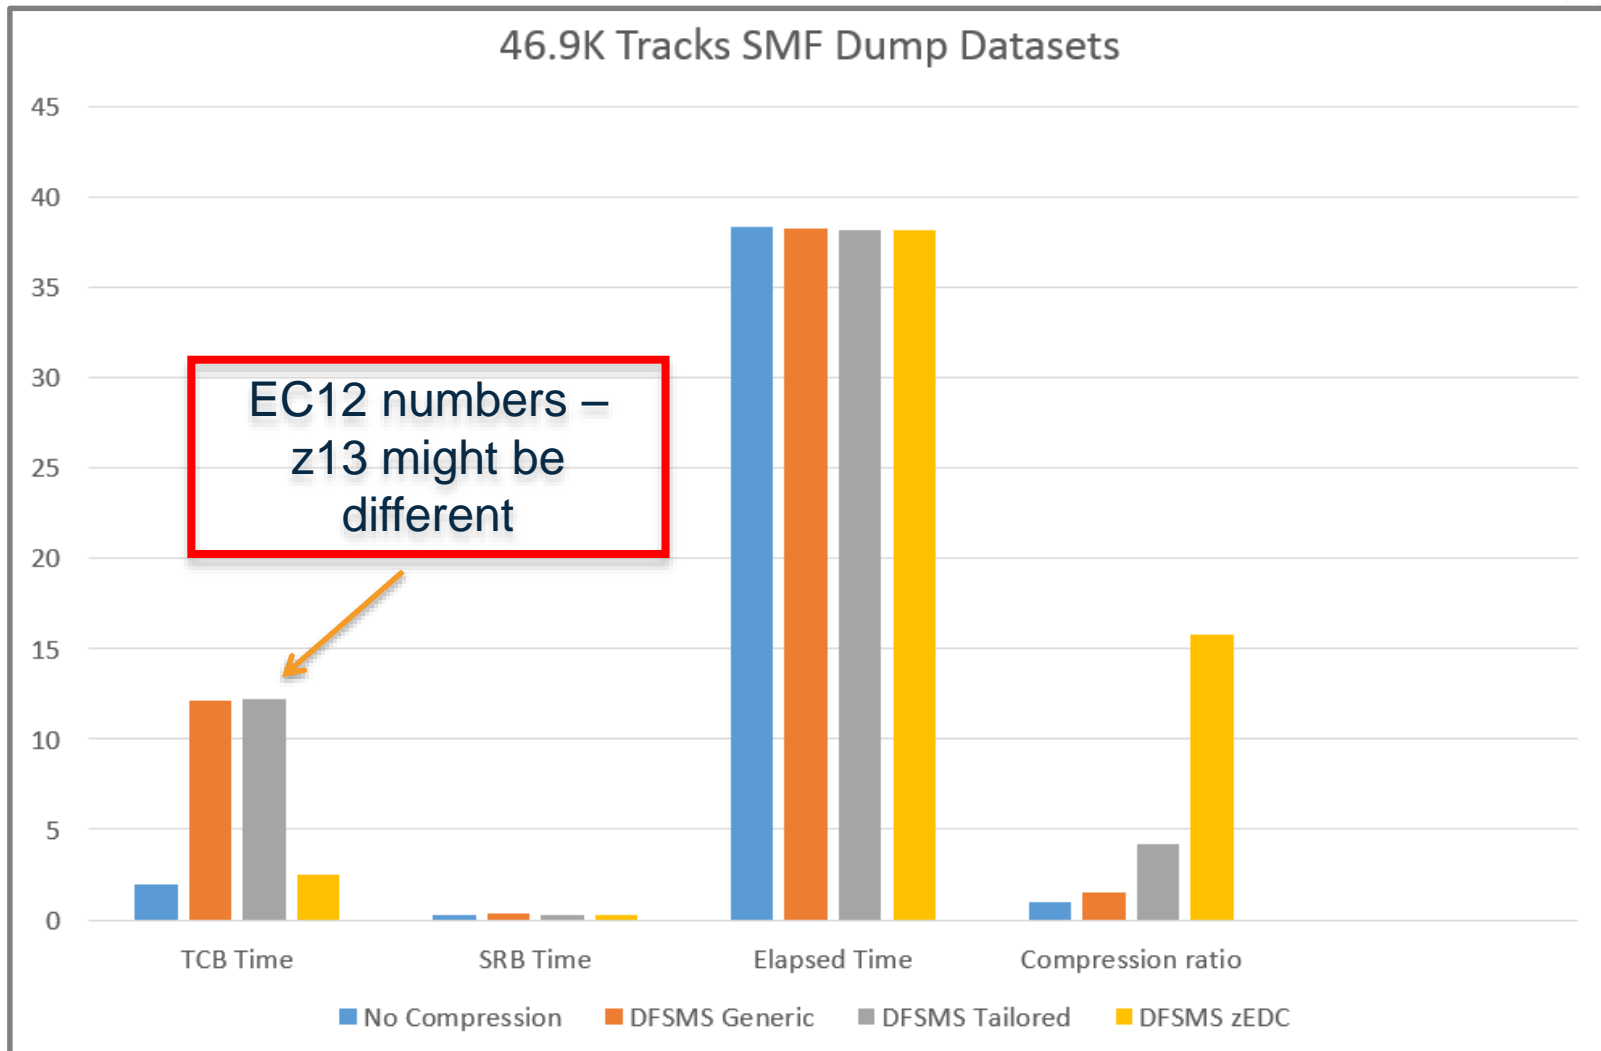
- z/OS 2.1 added compression of SMF records using zEDC
 - Applies to SMF log streams only
 - LSNAME/DEFAULTLSNAME(...**COMPRESS(NOPERMFIX / PERMFIX(nnnnnM))...**)
 - **PERMFIX(nnnnnM / 2M) | NOPERMFIX**
 - Global PERMFIX indicates the default amount of storage that can be fixed for SMF zEDC use for *each* logstream.

- Benefits:
 - Significantly better compression than other methods and there was NO compression option for the SYS1.MAN data sets.
 - Our customers are reporting capture ratios between 7.5:1 to over 10:1.
 - Significantly less CPU time than other compression methods.
 - Significantly less disk space and network transfer times.
 - It's SO EASY!

z/OS 2.1 - zEDC



z/OS 2.1 - zEDC



Complete your session evaluations online at www.SHARE.org/Orlando-Eval

z/OS 2.1 - zEDC

RMF Postprocessor XML Toolkit\PCIE.xml

File Edit View Favorites Tools Help

RMF Postprocessor Interval Report [System Z2] : PCIE Activity Report

RMF Version : z/OS V2R1 SMF Data : z/OS V2R1
 Start : 08/13/2013-05.45.00 End : 08/13/2013-06.00.01 Interval : 15:00:000 minutes

▼ General PCIE Activity

Function ID	Function PCHID	Function Name	Function Type	Function Status	Owner Job Name	Owner Address Space ID	Function Allocation Time	PCI Load Operations Rate	PCI Store Operations Rate	PCI Store Block Operations Rate	Refresh PCI Translations Operations Rate	DMA Address Space Count	DMA Read Data Transfer Rate	DMA Write Data Transfer Rate
0001	0380	Hardware Accelerator	1014044B	Allocated	FPGHWAM	0014	900	0	0.091	0	2.91	1	62.0	0
0011	05C4	Hardware Accelerator	1014044B	Allocated	FPGHWAM	0014	900	0	0.091	0	2.92	1	14.3	19.1
0020	038C	10GbE RoCE	15B31003	Allocated	VTAM	00DE	900	0.889	0	0	0	1	42.9	0

▼ Hardware Accelerator Activity

Function ID	Time Busy %	Request Execution Time	Std Dev for Request Execution Time	Request Queue Time	Std Dev for Request Queue Time	Request Size	Transfer Rate Total
0001	0.005	31.4	4.88	545	68.0	75.2	0.110
0011	0.004	30.7	5.35	541	93.3	74.4	0.109

▼ Hardware Accelerator Compression Activity

Function ID	Compression Request Rate	Compression Throughput	Compression Ratio	Decompression Request Rate	Decompression Throughput	Decompression Ratio	Buffer Pool Size	Buffer Pool Utilization
0001	1.46	0.088	4.11	0	0		64	0
0011	1.46	0.087	3.94	0	0		64	0

100%

- **RECOMMENDATION** on PERMFIX:
 - Monitor buffer usage as described in DSPSIZMAX above
 - Monitor SMF Type 23 fields: SMF23PFT (total storage SMF is using for zEDC), SMF23PFM (PERMFIX for this log stream), SMF23PFH (maximum storage used for zEDC), SMF23CWN (number of blocks written containing compressed records), SMF23NCN (number of blocks written containing non-compressed records)
 - Adjust if fixed storage is a problem

- How to charge for use of zEDC? If zEDC is not free, but it moves processing off your CPs, so how do you recover your costs?
- APAR [OA45767](#) delivers support for SMF Type 30 accounting of zEDC use.
 - If you do chargeback and are interested in zEDC, that APAR is a must-have.
 - APAR [OA48268](#) (still open) corrects contents of some fields.

- z/OS 2.2, and rolled back to 2.1 via APAR [OA47917](#) (still open as of 8/12/15), added 2 new fields to the Type 23 records: SMF23BBC and SMF23BAC.
 - These new fields report the number of bytes before and after zEDC compression for that log stream. Allows you to determine the benefit you are getting from zEDC for each logstream.
 - Because of the compression algorithm used by zEDC, IBM expect that the compression ratio for SMF logstreams will be less than that for sequential data sets containing SMF records – use these fields to see what compression YOU are getting.

- **RECOMMENDATION:**

- If you have the zEDC feature on your CPC, we definitely recommend that you exploit it with your SMF log streams. If you are already exploiting the SMF log stream support, adding zEDC support is only a few minutes work.
- If you are not currently using the SMF log stream support and your CPC has zEDC, this is definitely a good reason to revisit that decision. Moving SMF from data set to Logstream mode is not a trivial exercise, but adding zEDC support to an SMF that is already running in Logstream mode IS trivial.
- After exploiting zEDC for SMF, monitor results with RMF PCIE report.
- Apply (available) PTF for HIPER APAR [OA46498](#)

z/OS 2.1 - SWT/TWT

- z/OS 2.1 adds new wait time limits **SWT** and **TWT**
 - Previously, and as a default, the JWT (Job Wait Time Limit) provided wait time limits for all types of work
 - Value is specified as **JWT=hhmm|0010**; if wait limit is exceeded and TIME=1440 is not coded in the JCL, SMF exit IEFUTL is called (which can cancel job or let it continue). The majority of sites use JWT(0030).
 - New parameters SWT (Started Task Wait Time) and TWT (TSO/E Wait Time) lets you override JWT for STC and TSO users. If not specified, uses JWT value.
 - See BPXPRMxx parameter of PWT for Unix tasks
 - APAR [OA38478](#) provides toleration for z/OS 1.13

- **RECOMMENDATION:**
 - Wait times include TSO user not entering at keyboard, HSM recalling data set, started task waiting for an operator reply, job waiting for a tape mount.
 - Possible options:
 - JWT(0030)** (To be consistent with most sites' settings)
 - TWT(0010)** (For TSO security)
 - SWT(0010)** (Probably only WTOR)
- References:
 - Cheryl Watson's *Tuning Letter 2014 No 1*
 - Cheryl Watson's *Tuning Letter 2008 No 3*

z/OS 2.1 - SMF30COUNT

- z/OS 2.1 adds new Type 30 instruction counts
 - Specify **SMF30COUNT/NOSMF30COUNT**
 - If SMF30COUNT is specified, you get new fields containing the number of instructions executed broken out by TCB vs SRB, General Purpose vs specialty, and enabled for specialty but run on GP.
 - Because CPU time is so variable, you might be able to identify major changes if CPU time changes but instruction counts don't change; calculate CPU/instruction.
 - To collect this data, you MUST enable Basic counters for the LPAR in the HMC and start the HIS address space on z/OS.
 - There are differing opinions in IBM about the usefulness of this data.

z/OS 2.1 - SMF30COUNT

- **RECOMMENDATION:**
 - You should already be running HIS, so we would turn this on (at least long enough to collect the data and analyze it)
- **References:**
 - Cheryl Watson's Tuning Letter 2014 No 1 – see reasons for variability in this issue on page 40.
 - IBM Redpaper [REDP-4727](#) - *Setting Up and Using the IBM System z CPU Measurement Facility with z/OS.*
 - IBM manual [SA38-0666](#) – *z/OS MVS System Commands* - see the sections titled 'Setting up hardware event data collection' and 'Appendix A HIS Map Format'.

Agenda

- What is SMF?
- SMF Data Sets
- SMF Logger
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- **z/OS 2.2 Enhancements** ←
- SMFLIMxx
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

z/OS 2.2 – RECSIGN

- z/OS 2.2 adds new facility to protect SMF records that are written to a logstream.
 - NORECSIGN/RECSIGN(HASH(SHA1|SHA256|SHA384|SHA512),
SIGNATURE(RSA|ECDSA),
TOKENNAME(name))
 - Can be globally set, or defined on DEFAULTLSNAME parameter, or defined on LSNAME parameter
 - Provides additional level of security via ICSF
 - See SHARE session 17792, Fri 8:30

Agenda

- What is SMF?
- SMF Data Sets
- SMF Logger
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- **SMFLIMxx** ←
- Additional Slides
 - SMF Reports
 - MVS Commands
 - References

z/OS 2.2 – SMFLIMxx

- New Parmlib member SMFLIMxx lets customer set storage limits using Parmlib member instead of having to code an IEFUSI exit.
 - Not in base zOS 2.2, but added with New Function APAR [OA47062](#)
 - Documentation will be found at <http://www-01.ibm.com/support/docview.wss?crawler=1&uid=isg1OA47062> (but not until about August 20, 2015)
 - Keywords in SMFLIMxx member such as JOBCLASS, USER, JOBNAME, etc. can be specified with limits for REGIONABOVE, REGIONBELOW, EXECUTE (with a value of CANCEL), SYSRESVABOVE, SYSRESVBELOW
 - New JCL keyword of REGIONX=(below,above)

Agenda

- What is SMF?
- SMF Data Sets
- SMF Logger
- Records, records, records!
- General Parameters
- Record-Specific Parameters
- z/OS 2.1 Enhancements
- z/OS 2.2 Enhancements
- SMFLIMxx
- **Additional Slides** ←
 - SMF Reports
 - MVS Commands
 - References

SMF Reports

- Three ways to get listings of record types written:
 - IFASMFDP – Used for offloading and clearing VSAM data sets, as well as selecting records from sequential files. See first figure.
 - IFASMF DL – Used for offloading SMF log streams. Creates same report.
 - Download information with utility program provided in appendix from SMF Logger Redbook to get CSV file. See second figure.
 - Use user exits to get listing of subtypes (also in appendix from SMF Logger Redbook). See third figure.

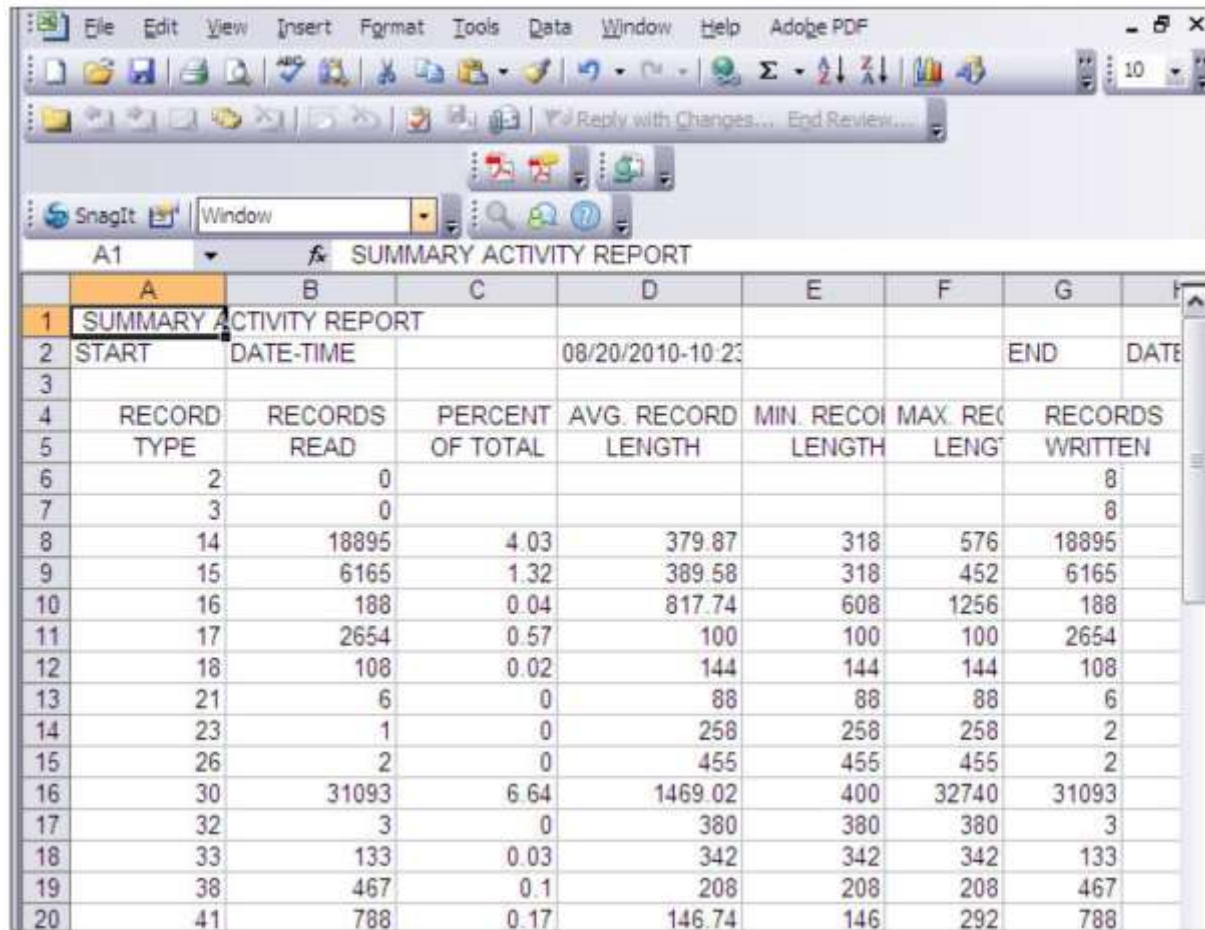
SMF Reports

– IFASMFDP/DL

START DATE-TIME 07/29/2010-15:04:18				SUMMARY ACTIVITY REPORT				END DATE-TIME 07/29/2010-17:22:42			
RECORD TYPE	RECORDS READ	PERCENT OF TOTAL	AVG. RECORD LENGTH	MIN. RECORD LENGTH	MAX. RECORD LENGTH	RECORDS WRITTEN					
2	0					1					
3	0					1					
14	35	.15 %	354.40	344	608	35					
15	37	.16 %	344.75	344	372	37					
17	1	.00 %	100.00	100	100	1					
23	13	.06 %	258.00	258	258	13					
30	448	1.93 %	1,300.87	400	3,925	448					
32	3	.01 %	232.00	224	248	3					
41	9	.04 %	412.00	412	412	9					
42	323	1.39 %	493.06	176	7,948	323					
60	1,724	7.44 %	546.71	338	547	1,724					
61	5	.02 %	464.40	304	958	5					
62	35	.15 %	188.00	188	188	35					
64	1,713	7.40 %	462.00	462	462	1,713					
65	1	.00 %	304.00	304	304	1					
70	18	.08 %	11,472.00	852	22,092	18					
71	9	.04 %	1,920.00	1,920	1,920	9					
72	576	2.49 %	1,389.43	1,108	13,260	576					
73	9	.04 %	21,728.00	21,728	21,728	9					

SMF Reports

- IFASMFDP using spreadsheet utility



1	SUMMARY ACTIVITY REPORT						
2	START	DATE-TIME		08/20/2010-10:23		END	DATE
4	RECORD	RECORDS	PERCENT	AVG. RECORD	MIN. RECOI	MAX. RECO	RECORDS
5	TYPE	READ	OF TOTAL	LENGTH	LENGTH	LENGTH	WRITTEN
6	2	0					8
7	3	0					8
8	14	18895	4.03	379.87	318	576	18895
9	15	6165	1.32	389.58	318	452	6165
10	16	188	0.04	817.74	608	1256	188
11	17	2654	0.57	100	100	100	2654
12	18	108	0.02	144	144	144	108
13	21	6	0	88	88	88	6
14	23	1	0	258	258	258	2
15	26	2	0	455	455	455	2
16	30	31093	6.64	1469.02	400	32740	31093
17	32	3	0	380	380	380	3
18	33	133	0.03	342	342	342	133
19	38	467	0.1	208	208	208	467
20	41	788	0.17	146.74	146	292	788

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

SMF Reports

– IFASMFDP using user exits

DETAILED ACTIVITY REPORT FOR OUTPUT DDNAME: SMFOUT

RECORD TYPE	RECORD SUBTYPE	SYSTEM ID	RECORDS WRITTEN	BYTES WRITTEN	RECORD MINLEN	RECORD MAXLEN	START DATE	START TIME	END DATE	END TIME
2	N/A	#@\$A	1	18	18	18	10.210	17:23:43.50	10.210	17:23:43.50
3	N/A	#@\$A	1	18	18	18	10.210	17:23:51.17	10.210	17:23:51.17
14	N/A	#@\$A	35	12404	344	608	10.210	15:04:18.55	10.210	17:21:49.04
15	N/A	#@\$A	37	12756	344	372	10.210	15:05:29.28	10.210	17:22:09.05
17	N/A	#@\$A	1	100	100	100	10.210	15:28:10.01	10.210	15:28:10.01
23	N/A	#@\$A	13	3354	258	258	10.210	15:13:24.32	10.210	17:13:24.32
30	1	#@\$A	3	1221	400	413	10.210	15:04:18.51	10.210	17:19:23.27
30	2	#@\$A	331	450861	1119	3925	10.210	15:15:00.00	10.210	17:21:39.20
30	3	#@\$A	5	6361	1209	1435	10.210	15:06:17.14	10.210	16:35:31.19
30	4	#@\$A	5	8341	1119	2109	10.210	15:06:17.14	10.210	16:35:31.20
30	5	#@\$A	5	8393	1119	2122	10.210	15:06:17.14	10.210	16:35:31.20
30	6	#@\$A	99	107613	1087	1087	10.210	15:15:00.00	10.210	17:15:00.00
32	1	#@\$A	3	696	224	248	10.210	15:15:11.51	10.210	17:21:39.20
41	3	#@\$A	9	3708	412	412	10.210	15:13:46.31	10.210	17:13:46.33
42	1	#@\$A	6	1056	176	176	10.210	15:13:25.58	10.210	17:13:26.39
42	2	#@\$A	126	32040	212	436	10.210	15:15:00.00	10.210	17:15:00.03
42	5	#@\$A	9	57708	5980	7948	10.210	15:15:00.00	10.210	17:15:00.00
42	6	#@\$A	158	61976	308	1604	10.210	15:04:18.42	10.210	17:22:09.05
42	24	#@\$A	24	6480	270	270	10.210	15:10:52.23	10.210	17:22:09.05
60	N/A	#@\$A	1724	942538	338	547	10.210	15:04:18.41	10.210	17:22:40.48
61	N/A	#@\$A	5	2322	304	958	10.210	15:04:18.54	10.210	17:19:41.31
62	N/A	#@\$A	35	6580	188	188	10.210	15:04:18.55	10.210	17:09:50.35
64	N/A	#@\$A	1713	791406	462	462	10.210	15:04:18.41	10.210	17:22:40.49
65	N/A	#@\$A	1	304	304	304	10.210	15:28:10.01	10.210	15:28:10.01
70	1	#@\$A	9	198828	22092	22092	10.210	15:15:00.01	10.210	17:15:00.01
70	2	#@\$A	9	7668	852	852	10.210	15:15:00.03	10.210	17:15:00.03

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

MVS Commands

- Display SMF
 - ‘d smf,s’ or ‘d smf,o’
 - s – displays status and datasets
 - o – displays the parameters in effect
 - SET SMF
 - Replaces SMFPRMxx or can be used to restart SMF
 - Format is ‘t smf=xx’, where ‘xx’ is the suffix for SMFPRMxx
 - SETSMF parameter(value)
 - If NOPROMPT is NOT used at IPL or AUTHSETSMF is specified, this can be used to override any parameter of SMFPRMxx except ACTIVE, PROMPT, SID, and EXITS
 - Example: ‘ss jwt(0030)’

References

- **z/OS MVS System Management Facilities (SMF)** – [SA22-7630-26](#) (z/OS 1.13) or [SA38-0667-02](#) (z/OS 2.1)
 - Contains most record layouts; dump programs; exits; macros
- **z/OS MVS Initialization and Tuning Reference** – [SA22-07592-23](#) (z/OS 1.13) or [SA23-1379-03](#) – SMFPRMxx parameters
- [SG24-7919-00](#) – **SMF Logstream Mode: Optimizing the New Paradigm** (10Feb2011) – See appendix for sample IFASMFDP exits to show subtypes; program to eliminate dupes; sample IEFU29L exit
- [GG24-4453-00](#) – **Enhanced Auditing Using the RACF SMF Data Unload Utility** (26Oct1994)

- References (cont.):
 - SHARE 2013 in San Francisco [Session 13043](#) – **Cheryl Watson** – *SMF 101 – Everything You Should Know and More*
 - IBM Techdocs [Flash10165](#) – **Riaz Ahmad** – *Limiting Storage use above the bar in z/Architecture*. Information on MEMLIMIT.

References

- WSC Techdocs
 - [WP101271](#) – *Migrating SMF from Data Set Recording to Log Stream Logging* (29Jul2008)
 - [TD102183](#) – *New SMF Support for zAAPs and SMF Accounting* (8Jun2005)
 - [WP101130](#) – *z/OS SMF Recording with MVS Logger* (12Nov2007)
 - [FQ131484](#) – *z/OS Frequently Asked Question* (23Jul2013)

References


- Cheryl Watson's Tuning Letters
 - 2008 No 2 – SMF Update – Part 1 (buffers, intervals, logger)
 - 2008 No 3 – SMF Update – Part 2 (SMFPRMxx, logger)
 - 2008 No 5 – SMF – Part 3 (address space records)
 - 2008 No 6 – SMF – Part 4 (user survey results, volume and data set records)
 - 2009 No. 2 – Type 19 SMF Records
 - 2009 No 3 – SMF Update – Part 5 (BUFSIZMAX, dumping, logger recommendations)
 - 2009 No 4 – SMF Update – Part 6 (security, WLM, CICS, DB2 records)
 - 2011 No 3 – SMFPRMxx update, SMF Type 113s

Other sessions?

- If you found this session helpful, here are our other sessions:
 - Wednesday 11:15 – **Cheryl and Frank** – [A Consultant's View on Recent z/OS Pricing Options.](#)
 - Friday 10:00 – **Cheryl and Frank** – [The Cheryl and Frank zRoadshow](#)
- If you like SMF data (and who doesn't??!!), please see our new AND IMPROVED(!) *SMF Reference Summary* at www.watsonwalker.com/references.html
- Also, please see our new series on SMF in Enterprise Systems Media TechJournal starting Aug/Sep 2014 - <http://ourdigitalmags.com/publication/?i=222435&p=10>

Thank You

- Thank you for coming and please enjoy the rest of the conference and your time in Orlando.



- Please remember to complete an evaluation – the session number is 18034.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval