



Systems & Technology Group


# Utilizing z/OS Logger Support for SMF

Share in Boston, August 2010

Glenn Anderson  
IBM Technical Training  
grand@us.ibm.com



© 2010 IBM Corporation

IBM Training 

## Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by © are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\*, AS/400®, e business (logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®


The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries. Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office. IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* All other products may be trademarks or registered trademarks of their respective companies.

**Notes:**  
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.  
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.  
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.  
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.  
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.  
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.  
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

© 2010 IBM Corporation


IBM Training 

## What's SMF worth to you?

**SMF records are critical to parts of your business**

- Can you afford to have SMF records lost?
- Are you avoiding creating certain likely-useful SMF records because of the concern that doing so might result in losing other records?
  - Writing SMF data to logstreams will help


3 | © 2010 IBM Corporation

IBM Training 

## Overview

- **z/OS images are getting larger and faster**
  - More CPUs, more real storage, more jobs run
  - Resulting in more SMF data to record!
- **SMF data set recording has not kept pace**
  - Still uses *modern* technology from 20 years ago:
    - Individual, small writes (CI-sized)
    - No striping
    - VSAM ICIP (requiring no secondary extents, and explicitly cleared MANx data sets)


4 | © 2010 IBM Corporation

IBM Training 

## Overview

- **SMF recording limitations can cause data to be lost when:**
  - I/O to active dataset is held up
  - All MANx datasets full (i.e. dump required)
  - Across MANx dataset “switch” processing
    - Heavy recording can overrun buffers while processing switch
- **SMF Dump Program must read every record ...**
  1. To dump records from MANx data set to QSAM data
  2. To copy and build GDGs
  3. When filtering


5 | © 2010 IBM Corporation

IBM Training 

## Overview

- **Solution:**
  - Enhance SMF to record to System Logger log streams
    - Data buffered in **multiple data spaces** (1 per log stream),
      - Increases buffering capacity
      - Prevents data going to one log stream from interfering with data going to another
    - Data written to log streams by **multiple tasks** (1 per log stream),
      - Increases write rate
      - Eliminates single point of failure (writer SRB)
    - Allow data to be partitioned to log streams *as it is created*
      - Records can be recorded to **multiple log streams**
  - Create new dump program to read log streams
    - Provide new OUTDD parameters that reduce the need for reprocessing data


6 | © 2010 IBM Corporation

IBM Training 

## Overview

- **Benefits:**
  - System Logger captures data faster than MANx datasets
  - No I/O during switch processing to cause data overrun
  - SMF Dump program reads data in specified logstream(s) only
    - Not all data at once

7 | © 2010 IBM Corporation

IBM Training 

## Scalability

- **SMF data written to log streams will enhance Scalability in two ways:**
  1. Utilize System Logger to improve the write rate and increase the volume of data that can be recorded
    - System Logger utilizes modern technology (Coupling Facility, Media Manager) to write more data at much higher rates than SMF's "MANx" dataset allows
  2. Reduce processing time as less filtering needed by dump program
    - Provide better management of the data by separating different record types into a number of different logstreams
    - Providing keywords on the OUTDD keyword of dump program that allow data to be "read once, written many"
  - Use of log streams for SMF data is **optional**. Existing "MANx" function continues to exist for customers satisfied with this functionality

8 | © 2010 IBM Corporation

IBM Training

## Overview: Current SMF Data Flow

The diagram illustrates the SMF Address Space. A yellow box labeled 'Program' sends a 'Record' to a 'Buffer' within the 'SMF Address Space'. The space contains several 'Buffer' and 'Ctrl Info' blocks arranged in a circular pattern. A central cylinder represents the 'SYS1.MANx' dataset. Arrows show data flow from the buffers to the dataset. A 'Record' is shown being written to a 'Buffer' and then to the dataset.

- Program requests to write a SMF record
- Locates appropriate buffer in SMF A.S. to write the record
- When ready to write, writes full buffers to the SMF dataset

**Note:** Each buffer is numbered to correspond to a particular record in the SMF dataset

9 | © 2010 IBM Corporation

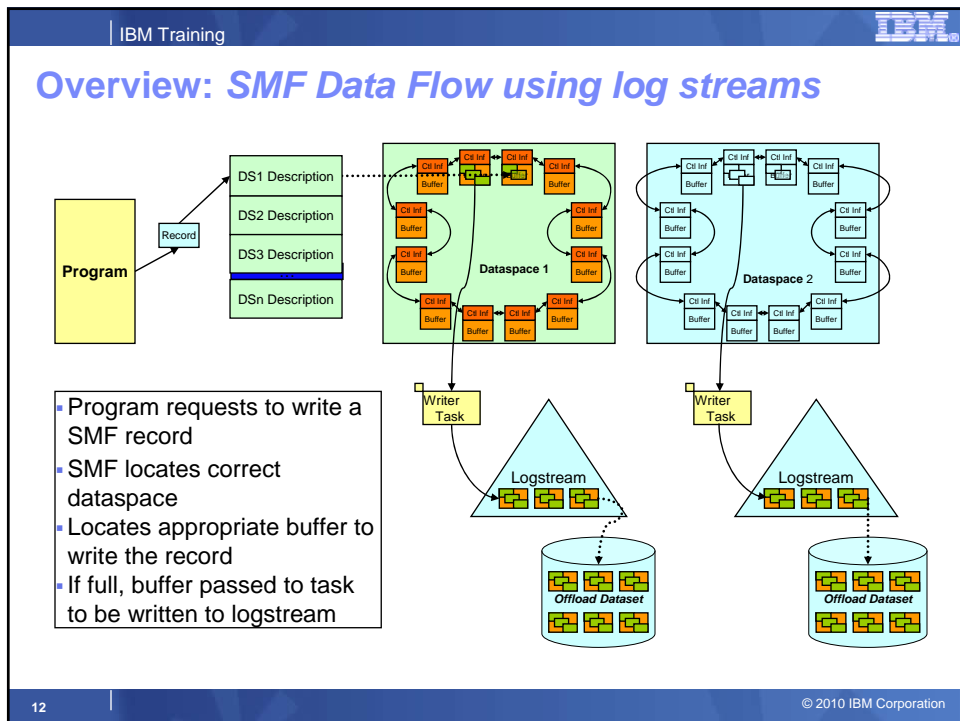
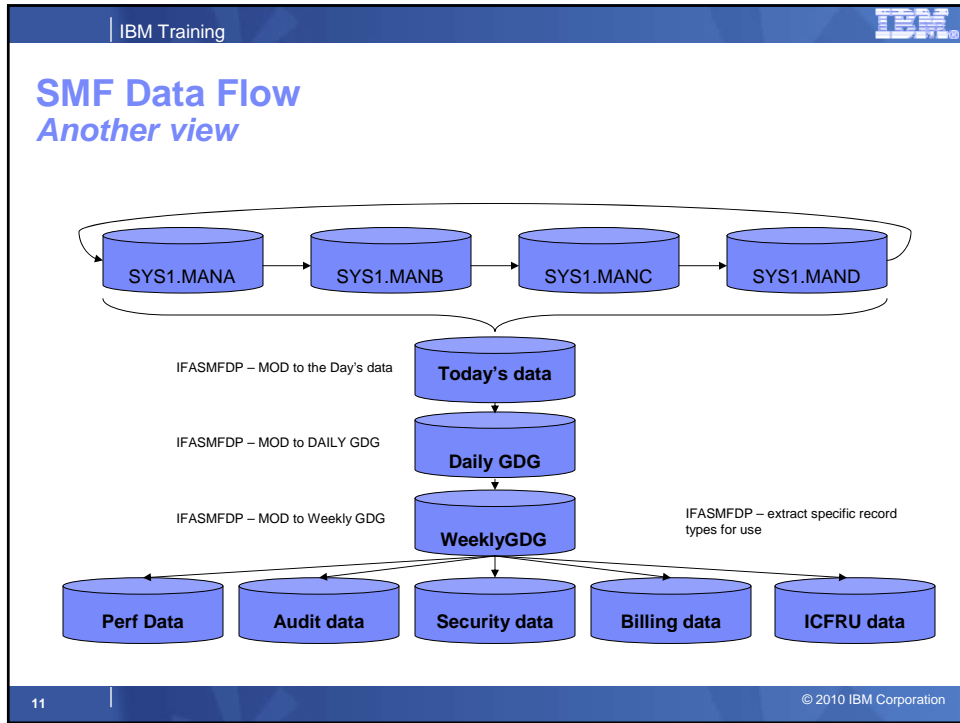
IBM Training


## Current SMF data flow

The diagram shows the 'SMF Dump Routine' as a central box. It has arrows pointing to two 'GDG - Archive' datasets (represented as cylinders with horizontal bars) and one 'SYS1.MANx' dataset (represented as a cylinder with vertical bars). The routine is shown copying data from the 'SYS1.MANx' dataset to the 'GDG - Archive' datasets.

- After full dataset switched with empty dataset:
  - Run dump routine to copy data for archiving to GDGs
  - Ready SYS1.MANx for reuse
  - Installation can choose to dump different records to different datasets

10 | © 2010 IBM Corporation



IBM Training 


## Usage & Invocation

- **The support is invoked by:**
  - Define new logstreams in system logger
    - See "Setting up a Sysplex" for documentation
  - Defining new keywords in SMFPRMxx:
 

```
LSNAME(IFASMF.q1.q2,TYPE(xx:yy))
DEFAULTLSNAME(IFASMF.q1.q2)
RECORDING(DATASET|LOGSTREAM)
```

    - SETSMF operator command can be used to toggle recording settings (for fallback, as an example)
  - Creating new JCL to use IFASMF DL with new logstreams
  - Update processes to use data from logstreams, if necessary
    - Ex. Automate periodic "Switch SMF" commands to drive new SMF Dump program
  - Activate PARMLIB changes via IPL or SET SMF=xx command


13 | © 2010 IBM Corporation

IBM Training 

## Usage & Invocation

- **Called by:**
  - Any function recording SMF data, no change to programs writing SMF data
- **New/Changed External Output**
  - New messages: IFA7xx messages, instead of IEE360, etc.
  - Notable messages:
    - IFA710I LOGSTREAM PARAMETERS WILL NOT BE USED DUE TO ERROR.
    - IFA711I LOGSTREAM PARAMETERS ARE IN EFFECT
    - IFA718E – LOGSTREAM IFASMF.aaaa UNAVAILABLE, RC=xxx,RSN=yyy DATA BEING BUFFERED, TIME=time

14 | © 2010 IBM Corporation

IBM Training 


## Usage & Invocation (examples)

- **Using logstreams, you can:**
  - Write data to a DASDONLY logstream, simply replacing SMF SYS1.MANx datasets
    - Use DEFAULTLSNAME(IFASMF.xxx) or LSNAME(IFASMF.xxx,TYPE(0:255)) to specify logstream
    - Run new SMF Dump Program to archive data

**Value**

1. Simplest approach to using logstreams
2. Better performance using logstream vs. SMF data sets

15 | © 2010 IBM Corporation

IBM Training 

## Usage & Invocation (examples)

- **A more sophisticated approach is also possible...**
  - Write data to a logstream on a 'task oriented' basis
    - Record Types (30,70:72,99) to one log stream (eg. IFASMF.PERF.DATA)
    - Record Types 30,80:81,83 to another log stream (eg. IFASMF.AUDIT.DATA)
    - Record DB2 data (Type 101) to a third stream (eg. IFASMF.DB2.DATA)
    - And use the DEFAULTLSNAME keyword to record all other record types.

16 | © 2010 IBM Corporation





## SMF Dump Program for Logstream - IFASMF DL


- **Differences between IFASMF DP and IFASMF DL**
  - IFASMF DL does not support CLEAR or any DELETE function
  - IFASMF DL requires SYSIN DD
  - IFASMF DL requires DATE parameter within the SYSIN DD to be effective
    - You don't want to read the whole log stream each and every time!



## Usage - Example 1

- **Run new SMF Dump program to dump performance data with out reading other data**

```
//AHMAD1 JOB MSGLEVEL=(1,1),MSGCLASS=A
/**
/* Function: Dump SMF records from a log stream
/**
//DUMPMAN EXEC PGM=IFASMF DL
//OUTDD1 DD DSN=AHMAD.JUSTONE.DATA,DISP=(NEW,PASS,DELETE),
// UNIT=SYSDA,VOL=SER=DEPT32,
// SPACE=(CYL,(5,1),RLSE),DCB=(LRECL=32760,RECFM=VBS,BLKSIZE=32756)
//SYSPRINT DD SYSOUT=*
//DUMPOUT DD DUMMY
//SYSIN DD *
LSNAME (IFASMF.PERF.DATA)
OUTDD (OUTDD1,TYPE(72),START(0000),END(2400))
DATE(2007011,2007017)
```

IBM Training 


## Usage - Example 2

- Run new SMF Dump program to dump from multiple logstreams

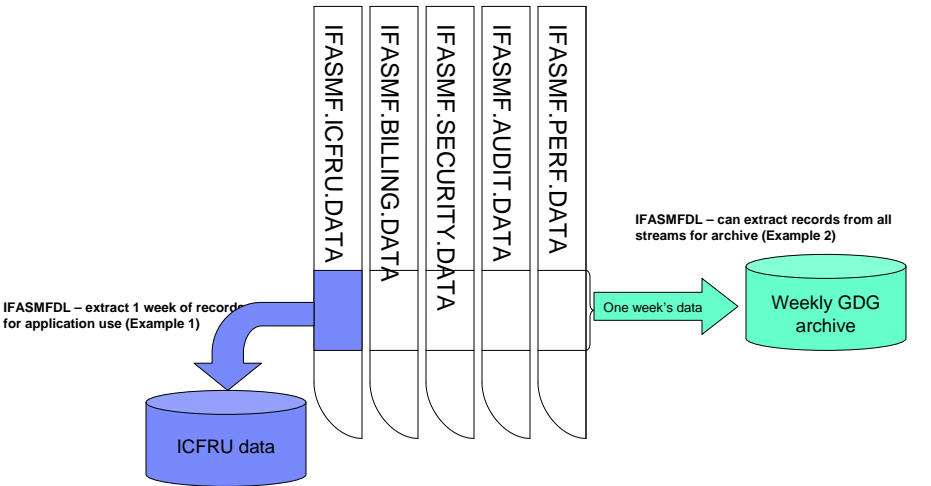
```

//AHMAD2 JOB MSGLEVEL=(1,1),MSGCLASS=A,NOTIFY=AHMAD
//****
//* Function: Dump SMF records from 4 Log streams
//**
//DUMPMAN EXEC PGM=IFASMF DL
//OUTDD1 DD DSN=AHMAD.ALLDATA,DISP=(NEW,PASS,DELETE),
// UNIT=SYSDA,VOL=SER=DEPT32,
// SPACE=(CYL,(5,1),RLSE),DCB=(LRECL=32760,RECFM=VBS,BLKSIZE=32756)
//SYSPRINT DD SYSOUT=*
//DUMPOUT DD DUMMY
//SYSIN DD *
LSNAME (IFASMF.AUDIT.DATA)
OUTDD (OUTDD1,START(0000),END(2400))
LSNAME (IFASMF.PERF.DATA)
OUTDD (OUTDD1,START(0000),END(2400))
    
```

19 | © 2010 IBM Corporation

IBM Training 

## SMF Data Flow – cont'd



IFASMFDL – extract 1 week of records for application use (Example 1)

IFASMFDL – can extract records from all streams for archive (Example 2)

ICFRU data

Weekly GDG archive

20 | © 2010 IBM Corporation



## Usage & Invocation (examples)

- **Tune logstream parameters via logger interfaces**
  - Update the storage class specified to alter attributes of offload and staging data sets.
  - Set Retention period (RETPD) for Performance data to 90 days, and auto-delete after that time
  - Set Retention period for Audit data to 7 years (2562 days) with auto-delete set to Yes.
  - Alter the LOWOFFLOAD and HIGHOFFLOAD values to ensure that Logger will offload data from the Logger CF structure in a timely manner.



## Another Example

- **A feature in the SMF Dump Program for log streams allow you to partition output data based on date, time, SMF id (SID)**

```
//DUMPMULT JOB MSGLEVEL=(1,1),MSGCLASS=A
//*****
//* Function: Dump SMF records from a Logstream
//*****
//DUMPMAN EXEC PGM=IFASMPDL
//OUTDD1 DD DSN=SYS1.SMFDATA.REC4.SY1(+1),DISP=(NEW,PASS,DELETE),
// UNIT=3390,VOL=SER=D72CT3,SPACE=(CYL,(5,1),RLSE),
// DCB=(LRECL=32760,RECFM=VBS,BLKSIZE=4096)
//OUTDD2 DD DSN=SYS1.SMFDATA.REC4.SY2(+1),DISP=(NEW,PASS,DELETE),
// UNIT=3390,VOL=SER=D72CT3,SPACE=(CYL,(5,1),RLSE),
// DCB=(LRECL=32760,RECFM=VBS,BLKSIZE=4096)
//SYSPRINT DD SYSOUT=A
//DUMPOUT DD DUMMY
//SYSIN DD *
LSNAME(IFASMF.MULTSYS.STREAM1)
OUTDD(OUTDD1,TYPE(4),START(0000),END(2400),SID(SY1))
OUTDD(OUTDD2,TYPE(4),START(0000),END(2400),SID(SY2))
DATE(2007011,2007011)
```

IBM Training

## Usage & Invocation (examples)

- Value:
  1. Segment data according to purpose
  2. Reduce reprocessing
  3. Record more data than previously

23 | © 2010 IBM Corporation

## SMF Processing

- Relative date processing in IFASMF DL intended to mirror typical GDG processing
- New RELATIVEDATE keyword
  - Specify DAILY, WEEKLY, or MONTHLY range and number of units
- IFASMF DL LSNAME OPTIONS to dump and/or delete data from logstream (vs. waiting for retention period to expire)
- SMFPRMxx MAXDORM to apply to SMF log streams (in addition to dataset recording)

24

## SMF Processing

- This supports allows for:
  1. A new RELATIVEDATE option for selecting a date range of records from the logstream with IFASMF DL
  2. A new ARCHIVE and DELETE option in IFASMF DL
  3. Allow MAXDORM value in SMFPRMxx to be applied to logstream recording so that buffered data can be moved to the logstream at regular intervals.
- Value:
  1. Allows the user to remove SMF data from the logstream
  2. Allows for grouping SMF logstream data by generic date masks (daily, weekly, monthly), eliminating the need for secondary post processing handling of the data (perhaps by propagating it into GDGs).
  3. Prevents SMF records from stagnating in the buffer

25

## Usage and Invocation

- The support for ARCHIVE, DELETE and RELATIVEDATE is invoked by the IFASMF DL program. The support for MAXDORM is invoked by updating your SMFPRMxx.
- RELATIVEDATE Parameter
  - Used to specify a date range based on the current day, week or month
    - **RELATIVEDATE(u, x, y)**
      - u – BYDAY, BYWEEK or BYMONTH
      - x – Number of units to move back
      - y – Number of units to gather
- DELETE/ARCHIVE Option
  - LSNAME(IFASMF.LS1,OPTIONS(**ARCHIVE**))
  - LSNAME(IFASMF.LS1,OPTIONS(**DELETE**))

26

## Usage and Invocation


- ARCHIVE and DELETE both operate on logger block boundaries.
- Entire logstream is read through only once.
- An IXGDelete is not done until all output datasets have been successfully closed.
  - ▶ There is no case where a “partial delete” of data can occur
- No “holes” in the logstream can exist. Blocks must be deleted in a contiguous subset starting at the beginning of the logstream.
- ARCHIVE and DELETE always begin at the start of the logstream regardless of the start date specified.
- Every record in the time range needs to match an OUTDD for ARCHIVE
- Entire logstream can not be ARCHIVED or DELETED. Only a subset of the logstream can be processed.

27

## Dependencies

- **No Hardware Dependencies**
- **Software Dependencies**
  - Major emphasis on leaving SMF data set (MANx) path intact
  - Also emphasis on maintaining end user data view
    - Record data from IFASMF DL should be indistinguishable from data from IFASMF DP
  - Expected hits to ISV code that uses “live” SMF data (i.e. programs that read MAN datasets rather than using output from SMF Dump program)
    - Check with ISVs that use the SMF data sets in this way


28

IBM Training 

## Migration & Coexistence Considerations

- **SMF MAN datasets can still be defined in SMFPRMxx**
- **Recording can work to logstream or dataset, but not both!**
  - SETSMF command can toggle environment for new exploitation or for fallback, but buffered data may be lost if it could not be written
- **A robust exploitation of new function may require business process updates!**
  - e.g, will Billing Department find required data in same datasets?
- **Coexistence :**
  - Logstreams can be single system (DASDONLY) or CF-based
  - When CF-based, be sure each system has a unique SMF “SID”

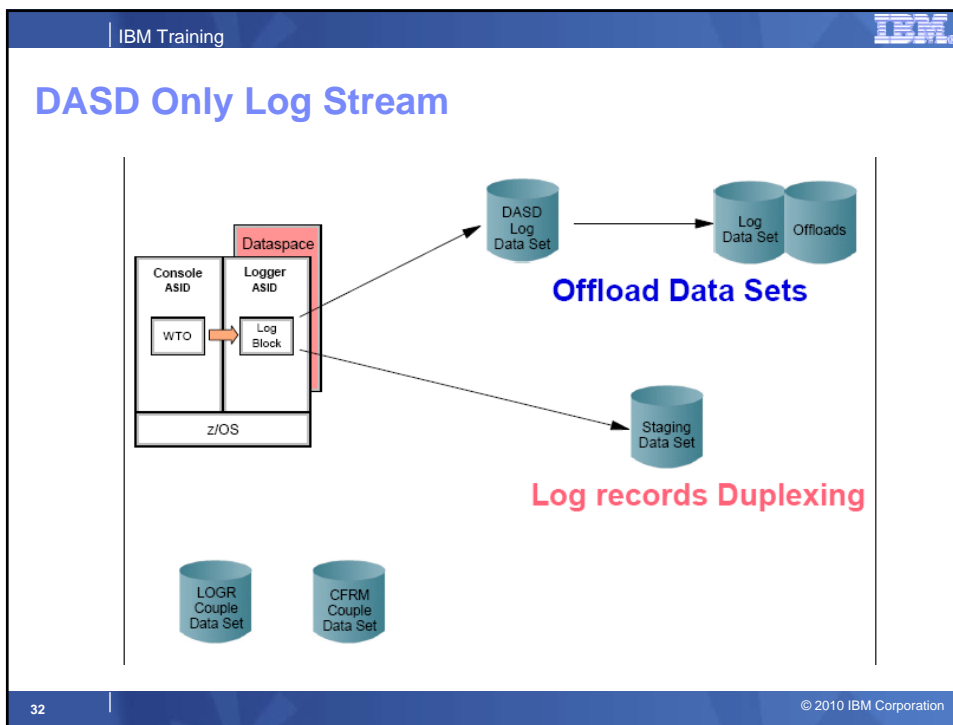
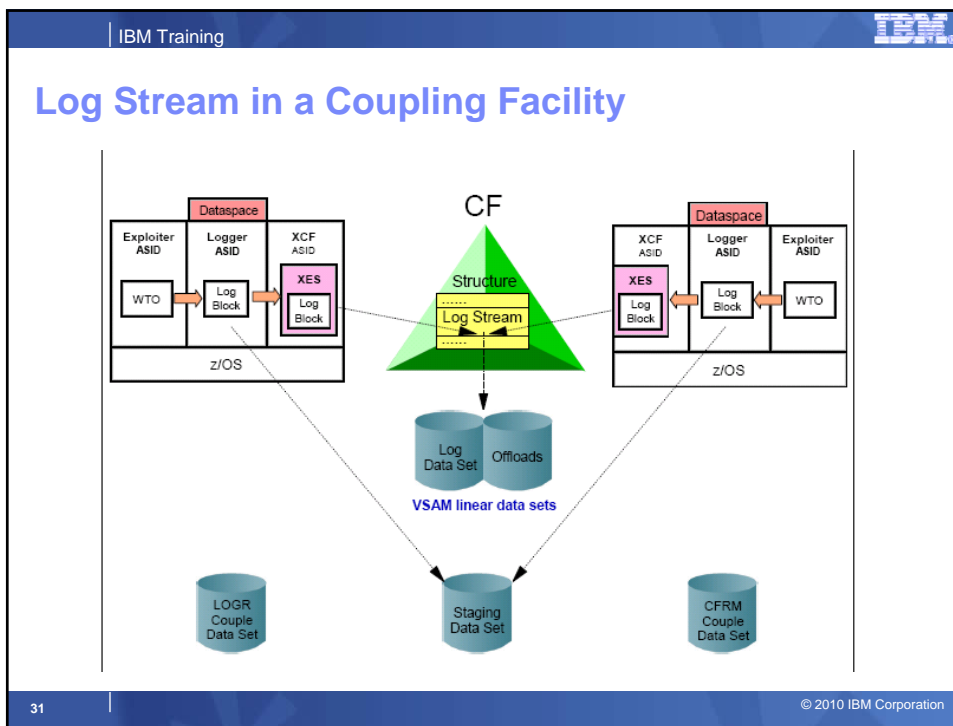
29 | © 2010 IBM Corporation

IBM Training 


## Installation

- **Prerequisites for installation**
  - Use IXCM2APU to create log streams for SMF
    - Decide on retention periods, CF vs DASDONLY, staging/offload dataset size, etc.
    - Recommend using staging data sets for early implementers
  - Use CFSIZER to plan size of CF Structures:
    - Consider data from recent IFASMFDP summaries to determine current data volume
    - Consider how long data should be retained in CF
      - If size too small, logger will have to offload frequently
    - Consider the Logger “HIGHOFFLOAD” specification
      - HIGHOFFLOAD(80) means 20% of structure space will be “white space”, intended to hold records while offloading during peak recording
    - Future re-planning exercises can use SMF Type 88 record data

30 | © 2010 IBM Corporation





IBM Training 

## IXCMIAPU – Define CF Logstream


```

//STEP1 EXEC PGM=IXCMIAPU
//STEPLIB DD DSN=SYS1.MIGLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSIN DD *
DATA TYPE(LOGR) REPORT(YES)
DEFINE STRUCTURE NAME(LOGGER_SMF) LOGSNUM(2)

DEFINE LOGSTREAM NAME(IFASMF.ALLSYS.DATA)
STRUCTNAME(LOGGER_SMF)
LOGGERDUPLEX(UNCOND)
DUPLXMODE(UNCOND)
STG_DUPLEX(YES)
STG_DATACLAS(MVSLOGR)
LS_DATACLAS(MVSLOGR)
LS_SIZE(500000)
HLQ(LOGGER)
HIGHOFFLOAD(85)
LOWOFFLOAD(0)
AUTODELETE(YES)
RETPD(2)

```

33 | © 2010 IBM Corporation

IBM Training 

## Installation

- **Prerequisites for installation**
  - Ensure sufficient SMS storage for peak recording/offload datasets
    - Consider data retention vs. available DASD pool
      - RETPD and AUTODELETE(YES) parameters will help ensure that old data is discarded when no longer needed
    - Consider using separate LS\_DATACLAS if concerned about DASD consumption
    - Having different data classes can help ensure that one application does not monopolize the DASD pool

34 | © 2010 IBM Corporation



## SMFPRMxx – SYS1.PARMLIB

- Be sure to plan for fallback to datasets in the event of problems !
- Add the LSNAME/DEFAULTLSNAME keywords

```

ACTIVE                /*ACTIVE SMF RECORDING*/
BUFSIZMAX(0800M)     /* MAXIMUM BUFFER SIZE */
DSNAME(&SYSNAME..MAN1,&SYSNAME..MAN2,&SYSNAME..MAN3,&SYSNAME..MAN4)
LISTDSN              /* LIST DATA SET STATUS AT IPL*/
NOPROMPT             /*DON'T PROMPT THE OPERATOR */

DEFAULTLSNAME(IFASMF.ALLSYS.DEFAULT)
LSNAME(IFASMF.ALLSYS.DATA,TYPE(100:255))

RECORDING(LOGSTREAM)

INTVAL(05)           /* SMF GLOBAL RECORDING INTERVAL */
MEMLIMIT(50G)        /* LIMIT ABOVE THE BAR      */
SYNCVAL(45)          /* GLOBAL SYNC VALUE        */
REC(PERM)            /*TYPE 17 PERM RECORDS ONLY*/
MAXDORM(3000)        /* WRITE AN IDLE BUFFER AFTER 30 MIN*/
STATUS(010000)       /* WRITE SMF STATS AFTER 1 HOUR*/
JWT(0030)            /* 522 AFTER 30 MINUTES*/
SID(&SYSNAME(1:4))    /* USE SYSNAME AS SID      */
SYS(NOTYPE(32,99),
  EXITS(IEFACTRT,IEFUTL,IEFUS1,IEFU83,IEFU84,IEFU29),
  INTERVAL(SMF,SYNC),NODETAIL)
/* WRITE ALL RECORDS EXCEPT TYPE 32 (TSO RECORDS), TAKE THE

```

## SMFPRMxx – SYS1.PARMLIB

```


ACTIVE                /*ACTIVE SMF RECORDING*/
BUFSIZMAX(0800M)     /* MAXIMUM BUFFER SIZE */
DSNAME(&SYSNAME..MAN1,&SYSNAME..MAN2,&SYSNAME..MAN3,&SYSNAME..MAN4)
LISTDSN              /* LIST DATA SET STATUS AT IPL*/
NOPROMPT             /*DON'T PROMPT THE OPERATOR */

DEFAULTLSNAME(IFASMF.ALLSYS.DEFAULT)
LSNAME(IFASMF.ALLSYS.DATA,TYPE(100:255))

RECORDING(LOGSTREAM)

INTVAL(05)           /* SMF GLOBAL RECORDING INTERVAL */
MEMLIMIT(50G)        /* LIMIT ABOVE THE BAR      */
SYNCVAL(45)          /* GLOBAL SYNC VALUE        */
REC(PERM)            /*TYPE 17 PERM RECORDS ONLY*/
MAXDORM(3000)        /* WRITE AN IDLE BUFFER AFTER 30 MIN*/
STATUS(010000)       /* WRITE SMF STATS AFTER 1 HOUR*/
JWT(0030)            /* 522 AFTER 30 MINUTES*/
SID(&SYSNAME(1:4))    /* USE SYSNAME AS SID      */
SYS(NOTYPE(32,99),
  EXITS(IEFACTRT,IEFUTL,IEFUS1,IEFU83,IEFU84,IEFU29),
  INTERVAL(SMF,SYNC),NODETAIL)
/* WRITE ALL RECORDS EXCEPT TYPE 32 (TSO RECORDS), TAKE THE

```

IBM Training 

## Performance


|                                 | Base run with SMF MANx data sets | Using 1 log stream | Split across 3 logstreams | Mult. Logstreams and type 30 duplicated | Mult. Logstreams and type 30, 100:102 duplicated |
|---------------------------------|----------------------------------|--------------------|---------------------------|---|--|
| CPU%                            | 86.56%                           | 86.19%             | 87.05%                    | 86.34%                                  | 86.95%   |
| Tot dasd I/O rate               | 4643                             | 3622               | 3387                      | 3436                                    | 3256   |
| SMFLOGR # of req                |                                  | 82769              | 90474                     | 91879                                   | 149324   |
| SMF data logging rate (rec/sec) | 17355.19                         | 17010.23           | 17221.54                  | 17199.62                                | 34472.71   |
| SMF avg. rec length             | 298.12                           | 298.12             | 298.12                    | 298.12                                  | 298.3  |
| SMF size in MB                  | 1776.33                          | 1741.02            | 1762.65                   | 1760.40                                 | 3530.46  |

SMF Logger study using Trade6 on a z9 system with 16CPs (single image)

99.9% of the SMF records were type 102 records

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

37 | © 2010 IBM Corporation

IBM Training 

## Summary

- Using new SMF function, you can:
  - Increase the speed and capacity of SMF recording
  - Decrease the reprocessing required for SMF data using keywords available in new IFASMF DL dump program

38 | © 2010 IBM Corporation

## Appendix

### ▪ Publication References

- SA22-7630 MVS System Management Facilities (SMF)
- SA22-7592 MVS Initialization and Tuning Reference
- SA22-7625 MVS Setting Up a Sysplex
- SA22-7593 MVS Installation Exits
- SA22-7627 MVS System Commands
- SA22-7637 MVS Messages, Volume 7 (IEE messages)
- SA22-7638 MVS Messages, Volume 8 (IFA messages)
- SQ24-6898 System Logger Redbook

### ▪ IBM Washington Systems Center – White Papers

- z/OS SMF Recording with MVS Logger – WP101130
- Migrating SMF from Data Set Recording to Log Stream Logging – WP101271
  - Available at: [www.ibm.com/support/techdocs](http://www.ibm.com/support/techdocs)