Introduction to SMF Performance Data Collection

Session 12916

February 4, 2013

Mary Astley

Advanced Technical Skills
IBM Corporation

Trademarks

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

CICS*	Parallel Sysplex*	OMEGAMON*
DB2*	PR/SM	Tivoli*
IBM*	RACF*	VTAM*
IMS	Rational*	WebSphere*
Lotus Domino*	Redbooks*	z/OS*
MQSeries*	RMF	z/VM*
MVS	System z*	zEnterprise*

* Registered trademarks of IBM Corporation

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.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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.

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.

Agenda

- A plan for collecting performance data
- Systems Management Facilities (SMF)
  - Overview of SMF data
  - SMF data collection
  - SMF data management
    - Data sets or logstreams
    - SMF dump program
- Resource Measurement Facility (RMF)
  - Parameters for SMF records produced by RMF
  - Post processor

Why Plan?

- In my experience
  - No planning often results in no data
  - Collecting all the measurement data does not guarantee one will have data required, it only guarantees use of DASD space
  - The data not collected will be data most needed
- Initial planning – identify data needed
  - Who needs the measurement data and why
  - How the data will be used

    System resource usage reports

    Capacity planning, performance analysis

    Reporting on service level objectives

    Historical trends

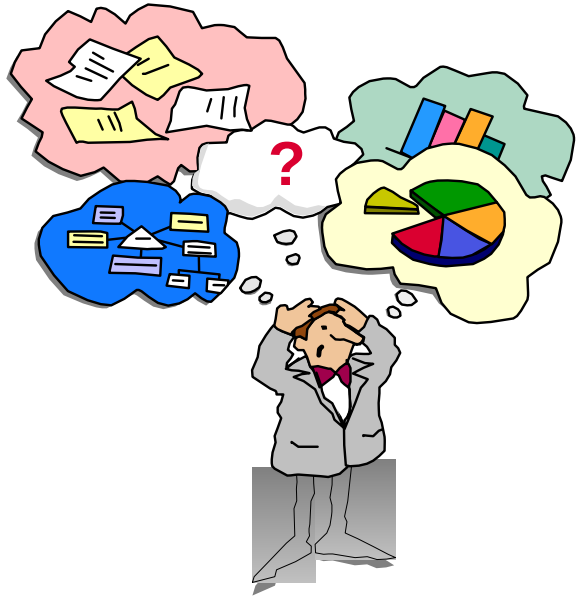

Measurement Data

- Identify sources of measurement data
  - Data collected from subsystems and program products
    - CICS, DB2, MQ, WebSphere Application Server
  - Data supplied by performance monitors
    - RMF, OMEGAMON
- Learn about the measurement data
  - What performance data is produced by each source
  - How often a data record is produced
  - Where the performance data is stored
- Examine use of the data
  - Who uses the performance data from each source
  - Possibility of data being collected, that is not used

Collecting SMF Data

Examine options to tailor the SMF data collected

- Which records are written by SMF
- Where the SMF records are written
- When SMF records are created
  Interval or frequency records produced
- Parameters used by RMF to create SMF records

Specify options in SYS1.PARMLIB members

- SMF options – SMFPRMxx members
- RMF options – ERBRMFxx members

SMF Records

- Produced by elements, features, subsystems, and program products in a z/OS system
- Each record is assigned a record type
  - Types 00-127 are reserved for IBM products
  - Types 128-255 are available for user records
- Within a record type there may be subtypes
- Each record contains an 18 or 24 byte header
  - Record type and optionally subtype
  - Date and time the record was written to SMF

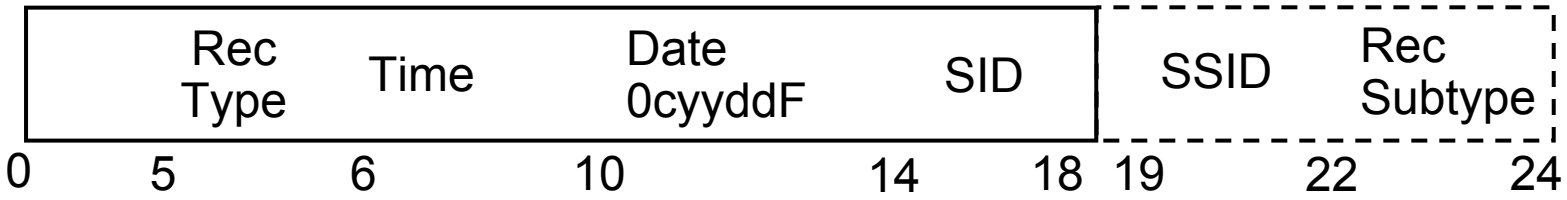

Some SMF Record Types

- CPU usage, paging, I/O activity
  - Type 30 for job and job step processor resource usage
  - Types 70-79 RMF records
- Data set activity
  - Type 42 contains DFSMS statistics
  - Type 92 for HFS activity
- System resource manager decisions
  - Type 99 written by SRM when in goal mode
- Subsystem statistics
  - Types 100-102 for DB2 statistics
  - Type 110 for CICS statistics

Which SMF Records Are Written

SMFPRMxx parameters control records written by SMF

- SYS(record types)
  - SMF record types and subtypes to be written
  - Values used when no SUBSYS is coded and when an option is omitted on SUBSYS parameter

  SYS(NOTYPE(32,99))
  - Write all record types except type 32 and 99
- SUBSYS(subsystem,record types)
  - Specify data recording options for a subsystem
  - STC, JES2, JES3, ASCH, TSO, OMVS

  SUBSYS(STC,TYPE(30(1:4),70:78))
  - For started tasks write type 30 subtypes 1-4 and 70-78

Where SMF Data is Written

- SMF data may be written to VSAM data sets or system logger managed logstreams
- Specified in SMF parameter

  RECORDING(DATASET | LOGSTREAM)
- Operator command to switch between two modes
  - SETSMF RECORDING(DATASET | LOGSTREAM)
  - SET SMF=xx

    where xx is SMFPRMxx suffix

SMF Data Sets

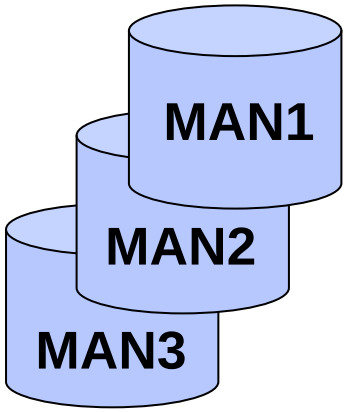

VSAM data sets

Specified in SMFPRMxx PARMLIB member

DSNAME(SYS1.MAN1,SYS1.MAN2,SYS1.MAN3)

First data set is primary, others are secondary

Data set SMF is using active data set
Empty data sets are alternates

Display SMF data sets and status: D SMF command

```
D SMF
IEE974I 13.07.20 SMF DATA SETS
    NAME              VOLSER SIZE(BLKS) %FULL   STATUS
  P-SYSA.MAN1         ZSMF01     33000   100    DUMP REQUIRED
  S-SYSA.MAN2         ZSMF01     33000    76    ACTIVE
  S-SYSA.MAN3         ZSMF01     33000     0    ALTERNATE
```

SMF Logstreams

- Defined with IXCMIAPU utility
  - RETPD and AUTODELETE LOGR parameters manage how long data is kept in logstream
- May be coupling facility or DASD only logstreams
- Specified in SMFPRMxx PARMLIB member
  - Logstream name and records to be written to it
  - Default logstream for any remaining records

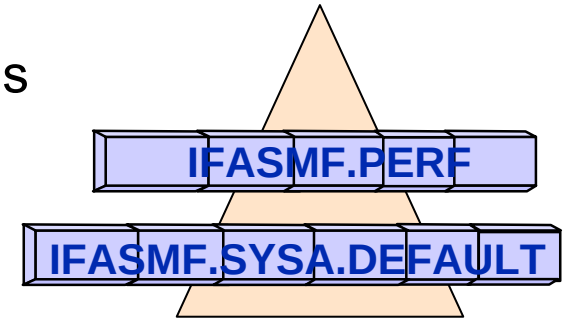

LSNAME(IFASMF.PERF,TYPE(30,70:79,100:120))
DEFAULTLSNAME(IFASMF.SYSA.DEFAULT)

Display SMF logstreams and status: D SMF command

```
D SMF
IFA714I 16.57.31 SMF STATUS 919
        LOGSTREAM NAME                BUFFERS        STATUS
      A-IFASMF.SYSA.DEFAULT             7828        CONNECTED
      A-IFASMF.PERF                     3364        CONNECTED
```

SMF Logstream Recording

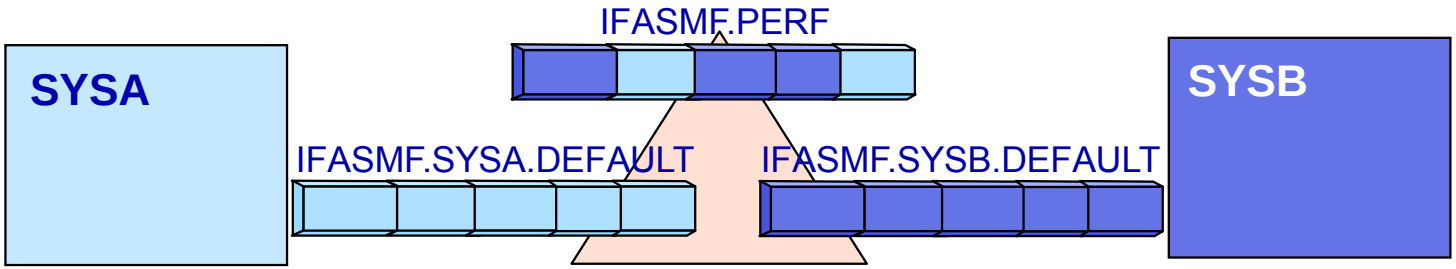

- Decide on number of logstreams to be defined
  - Multiple logstreams may be used for SMF data
  - One logstream or multiple logstreams per system
  - Separate record types into different logstreams
  - Merge records from multiple systems into a CF logstream
- Determine a naming convention for the logstreams
  - Name must begin with IFASMF.
    and can be up to 26 characters in length
  - Name could include system name or type of data

    IFASMF.SYSA IFASMF.PERF

When SMF Records Are Created

- z/OS components, subsystems, and program products determine when records are created
- Records are created for events, such as
  - Job or work unit start, job step end, and job end
  - Data set close or end of volume processing
  - File system is mounted or unmounted
  - RACF protected resource auditing
- Some records may be created at defined intervals
  - SMF Global recording interval
  - Intervals defined in subsystem and product
    - May provide parameters to tailor interval, record content
    - For example: DB2 zparms, RMF parameters

SMF Interval Recording

- SMF writes records at a specific time interval
  - Interval records written each recording interval
  - SMF record types 30 (job), 32 (TSO)
  - Interval starts when unit of work starts
- Minimize loss of data for long running jobs
- Interval parameter to enable and provide length

  SYS(INTERVAL(hhmmss))

  SUBSYS(INTERVAL(hhmmss)

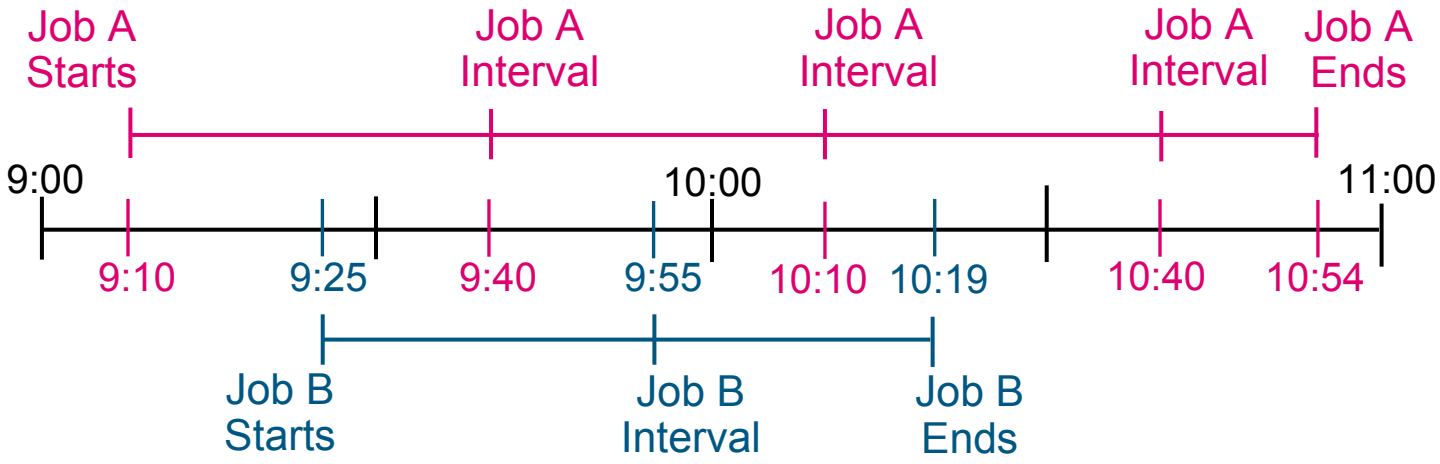

Synchronizing Data

- Synchronize the reporting interval with a minute in the hour
- Reasons to synchronize measurement data
  - Report data by time interval, e.g. by hour
  - Combine measurement data from different sources
  - Compare data for same time window across systems and program products
- Same synchronization value on all systems for sysplex data

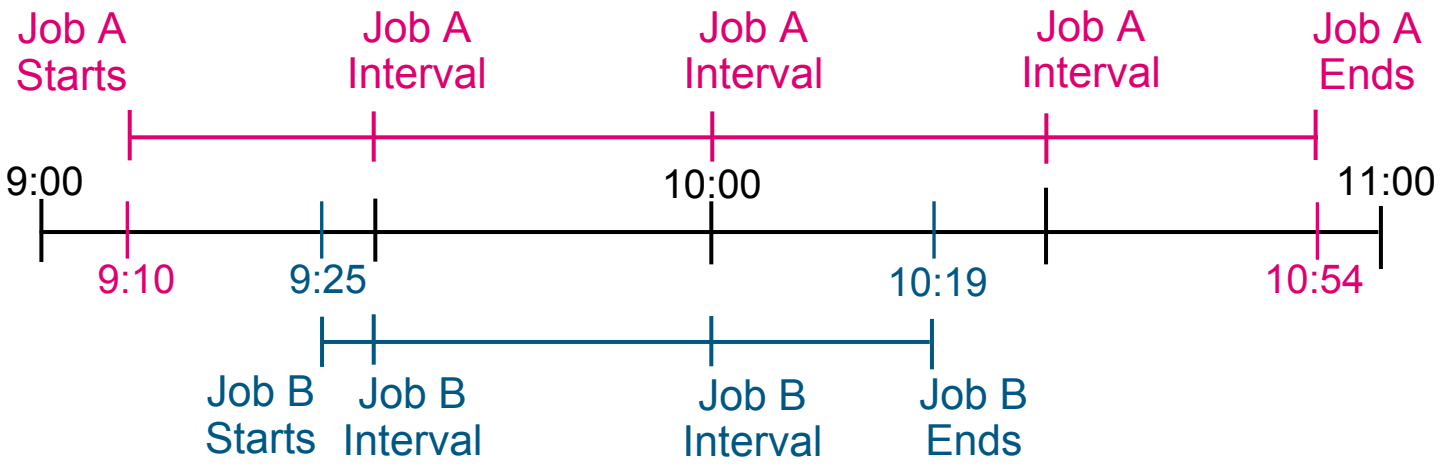

Global Recording Interval

- SMF recording interval available to programs
- Interval length is from 1 to 60 minutes
  - Default is 30
  - Specified with INTVAL(mm) SMF parameter
- Always synchronized with some part of hour
  - Value specified in minutes past the hour
  - Default is 00
  - Specified with SYNCVAL(mm) SMF parameter
- Synchronize SMF type 30, 32 interval records and RMF records with global recording interval
  - SMF: SYS(INTERVAL(SMF,SYNC))
    SUBSYS(INTERVAL(SMF,SYNC))
  - RMF: SYNC(SMF) in RMF monitor I parameters

SMFPRMxx PARMLIB Member

```
ACTIVE                                          /* SMF is active   */
SID(SYSA)                                       /* SMF system id */
RECORDING(DATASET | LOGSTREAM)   /*  Where written  */
DSNAME(SYS1.MAN1,SYS1.MAN2,SYS1.MAN3)
LSNAME(IFASMF.PERF,TYPE(30,70:79))
DEFAULTLSNAME(IFASMF.SYSA.DEFAULT)
INTVAL(30)                              /* Global recording interval  */
SYNCVAL(00)                             /* Global SYNC value          */
SYS(NOTYPE(32,99),INTERVAL(SMF,SYNC))   /* No type 32, 99 */
SUBSYS(STC,INTERVAL(SMF,SYNC),
            TYPE(0,30,70:79,100:102,110))       /* STC records   */
```

SMF Data Summary

- Each record has a record type
- SMFPRMxx options control
  - record types written to data sets or logstreams
  - Interval recording and global recording interval

z/OS, IBM products - Types 0-127
  RMF – Types 70-79
  DB2 – Types 100-102
  CICS – Type 110
  HTTP Server – Type 103
  WebSphere Appl. Server – Type 120
User records – Types 128-255

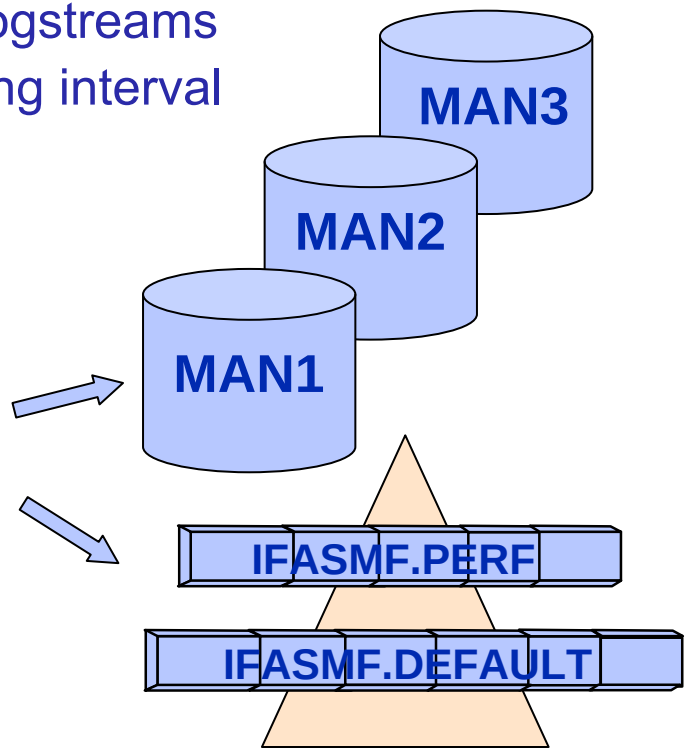

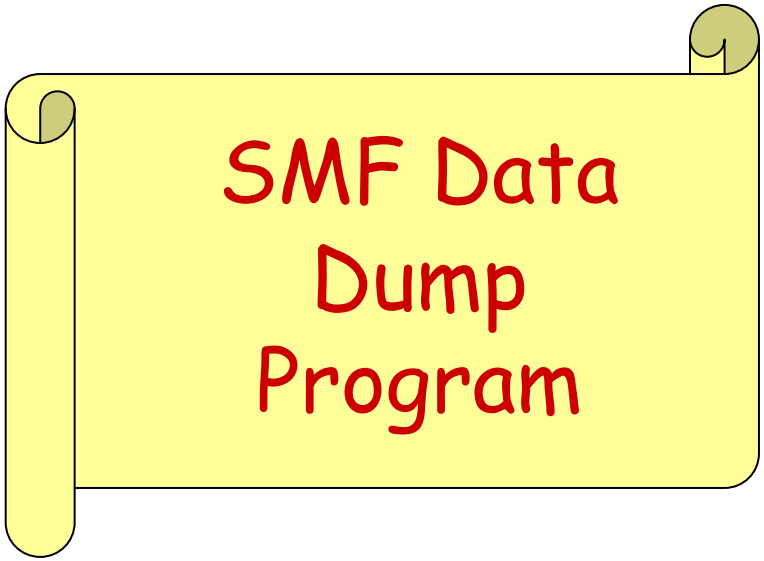

SMF Data Set Dump Program

IFASMFDP program

- Copies the input SMF data to output data sets
- Input is SMF records
  - SMF data set
  - Output data set from previous IFASMFDP execution
- Output is a sequential data set
- May have multiple input and output data sets
- Used to dump and clear SMF data sets
- Specify processing options through parameters
  - Input and Output DD names
  - Start and end dates and times for the records written
  - Record types and subtypes to be written to output data sets

SMF Logstream Dump Program

IFASMFDL program

- Copies the SMF data from logstreams to data sets
  - Input is SMF logstreams
  - May contain data for multiple systems
  - Probably will have data for more than one day
- Output is a sequential data set or data sets
- Specify record types and subtypes to be written to the output data sets
- Specify processing options through parameters
  - Input LSNAME and Output DD names
  - Start and end dates and times for the records written
  - System identifier

Dump & Clear SMF DS

Sample IFASMFDP JCL to dump and clear SMF data set

```
//STEP1    EXEC PGM=IFASMFDP
//SMFIN     DD  DSN=SYS1.MANX,DISP=SHR
//SMFOUT    DD  DSN=SYS1.SMFDATA,DISP=MOD
//SYSPRINT DD  SYSOUT=A
//SYSIN     DD  *
      INDD(SMFIN,OPTIONS(ALL))
      OUTDD(SMFOUT,TYPE(0:255))
/*
```

OPTIONS(ALL)
DUMP and CLEAR
SMFIN DD data set

TYPE(0:255)
Dump record types 0-255
to SMFOUT DD

IFASMFDP Example

Copy records from a SMF data set created by IFASMFDP

```
//STEP1    EXEC PGM=IFASMFDP
//SMFIN     DD  DSN=SMFDATA.D08044,DISP=SHR
//SMFOUT1   DD  DSN=SMFDATA.T30,DISP=SHR
//SMFOUT2   DD  DSN=SMFDATA.RMF,DISP=SHR
//SYSPRINT  DD  SYSOUT=A
//SYSIN     DD  *
      INDD(SMFIN,OPTIONS(DUMP))
      OUTDD(SMFOUT2,TYPE(70:78))
      OUTDD(SMFOUT1,TYPE(30(2:3)))
      DATE(2012215,2012215)
      START(0800)
      END(1700)
/*
```

For this day
Copy first shift data

IFASMFDL Example

Dump records from a SMF logstream to data sets

```
//STEP1    EXEC PGM=IFASMFDL
//SMFOUT1   DD  DSN=SMFDATA.T30,DISP=SHR
//SMFOUT2   DD  DSN=SMFDATA.RMF,DISP=SHR
//SYSPRINT  DD  SYSOUT=A
//SYSIN     DD  *
  LSNAME(IFASMF.SYSA.DEFAULT)
  OUTDD(SMFOUT2,TYPE(70:78))
  OUTDD(SMFOUT1,TYPE(30(2:3)))
  DATE(2012215,2012215)
  START(0800)
  END(1700)
/*
```

For this day
Copy first shift data

Input logstream

Summary Activity Report

- Provided by the IFASMFDP and IFASMFDL programs
- For SMF records read from the input file
  - START DATE-TIME date and time of earliest record read
  - END DATE-TIME date and time of latest record read
  - RECORD TYPE SMF record number
  - RECORDS READ number of each SMF record type read
- For SMF records written to the output file
  - RECORDS WRITTEN number of records by record type

Summary Activity Report

SHARE
Technology · Connections · Results

SUMMARY ACTIVITY REPORT

START DATE-TIME 08/01/2012-00:05:00 END DATE-TIME 08/03/2012-15:24:00

RECORD TYPE	RECORDS READ	PERCENT OF TOTAL	AVG. RECORD LENGTH	MIN. RECORD LENGTH	MAX. RECORD LENGTH	RECORDS WRITTEN
0	1	.00 %	64.00	64	64	0
2	1	.00 %	18.00	18	18	2
3	1	.00 %	18.00	18	18	2
14	29	.02 %	397.65	372	516	0
15	9	.01 %	372.00	372	372	0
20	2	.00 %	96.00	96	96	0
26	1	.00 %	447.00	447	447	0
30	42,204	27.67 %	1,663.17	400	32,752	4,522
34	1	.00 %	215.00	215	215	0
35	1	.00 %	150.00	150	150	0
40	71	.05 %	218.67	74	514	0
43	1	.00 %	32.00	32	32	0
45	1	.00 %	28.00	28	28	0
70	2,274	1.49 %	12,619.02	1,188	28,424	324
71	758	.50 %	2,012.00	2,012	2,012	108
72	54,574	35.78 %	1,474.33	1,132	20,316	7,776
73	758	.50 %	20,915.24	19,892	21,008	108
74	15,168	9.94 %	10,909.75	364	32,632	2,160
75	4,548	2.98 %	264.00	264	264	648
77	758	.50 %	321.05	320	640	108
78	1,516	.99 %	4,152.00	1,888	6,416	216
82	2	.00 %	154.00	40	268	0
88	9,096	5.96 %	234.50	161	308	0
89	1,516	.99 %	1,863.01	418	3,182	0
90	11	.01 %	260.72	72	632	0
100	9,925	6.51 %	1,396.29	306	3,526	0
101	4	.00 %	1,482.00	734	2,288	0
102	9,302	6.10 %	995.80	194	2,850	0
TOTAL	152,533	100 %	2,609.69	18	32,752	15,974

NUMBER OF RECORDS IN ERROR 0

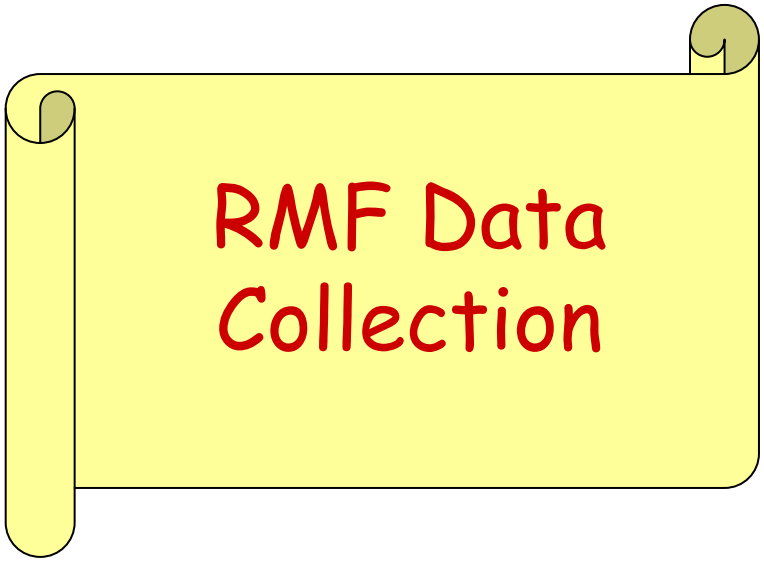

RMF Performance Data

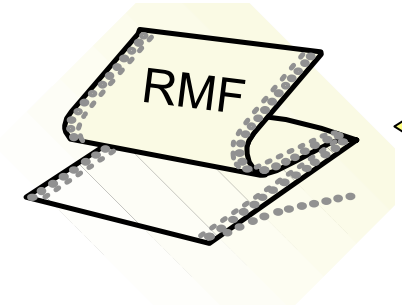

Monitor I Data

Interval reports

SMF record types 70-78

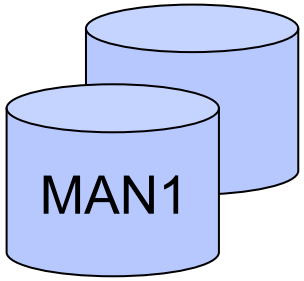

Monitor II Data

SMF record type 79

Interactively displayed

Monitor III Data

SMF record types 72, 74

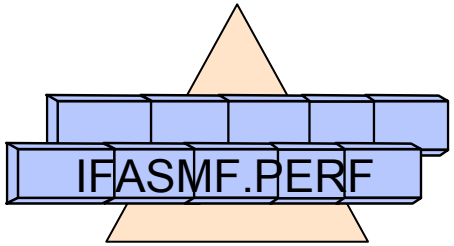

Interactively displayed with RMF Monitor III session

Written to VSAM data sets

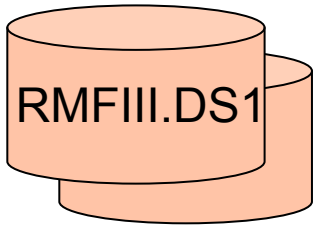

Synchronizing Data

- Data Scope
  - Single system
  - Sysplex
  - Other (e.g. cache data)
- Synchronize intervals with TOD clock
  - Sysplex reports need same SYNC value on all systems
  - RMF intervals same time frame as SMF intervals
- Specified with SYNC parameter
  - NOSYNC — no synchronization
  - SYNC(RMF,mm) — sync RMF with clock
  - SYNC(SMF) — use SMF Global Interval

Collecting RMF Data

- Parameters for performance data to be collected
  - Resource usage - CPU, CF, paging, storage, I/O
  - Workload information - tran. rates, response times
- Parameters controlling sampling, reports, SMF records
  - CYCLE(nnnn) Sample frequency in ms
  - INTERVAL(mm) Reporting interval
  - SYNC() | NOSYNC Synchronize intervals
  - NOSTOP | STOP(mmm) When monitoring stops
  - NORECORD | RECORD Write SMF records
- Parameters specified in ERBRMFxx PARMLIB members
- SMF parameter to write the RMF records
  - SYS(TYPE(70:79))

RMF Monitor I - Sample ERBRMFxx

```
/*********************************************************************/
/*  MEASUREMENT DATA TO BE COLLECTED                                 */
/*********************************************************************/
  CACHE                    /* CACHE STATISTICS                      */
  CHAN                     /* CHANNEL STATISTICS                    */
  CPU                      /* CPU STATISTICS                        */
  DEVICE(DASD)             /* DIRECT ACCESS DEVICES MEASURED        */
  DEVICE(NOTAPE)           /* NO TAPE DEVICES MEASURED              */
  ENQ(SUMMARY)             /* ENQUEUE SUMMARY                       */
  IOQ(DASD)                /* DASD I/O QUEUEING MEASURED            */
  PAGESP                   /* PAGE/SWAP DATASET STATISTICS          */
  PAGING                   /* PAGING DATA                           */
  NOTRACE                  /* NO TRACE REPORT                       */
  VSTOR(S)                 /* VIRTUAL STORAGE SUMMARY DATA          */
  WKLD(PERIOD)             /* WORKLOAD MANAGER DATA                 */
/*********************************************************************/
/*   SAMPLING AND REPORTING DATA                                     */
/*********************************************************************/
  CYCLE(1000)              /* SAMPLE EVERY SECOND (1000 MSEC)       */
  NOSTOP                   /* ACTIVE UNTIL OPERATOR ISSUES STOP     */
  SYNC(SMF)                /* USE INTVAL/SYNCVAL FROM SMFPRMXX      */
  NOOPTIONS                /* OPTIONS NOT DISPLAYED, NO REPLY       */
  RECORD                   /* WRITE SMF RECORDS EVERY INTERVAL      */
  NOREPORT                 /* NO WRITTEN REPORTS TO SYSOUT          */
  SYSOUT(A)                /* REPORTS TO CLASS A, IF REPORT         */
/*********************************************************************/
```

RMF Post Processor

- Reports system utilization and performance by "post processing " RMF and SMF data
  - SMF records written by RMF (Types 70-78)
  - SMF type 103 records from HTTP Server
  - SMF type 108 records from Lotus Domino Server
- Reports include
  - Workload Activity (sysplex)
  - Coupling Facility Activity (sysplex)
  - CPU Activity (system)
  - Device Activity (system)
  - Cache Subsystem Activity

Sample RMF Postprocessor JCL

```
//RMFPP EXEC PGM=ERBRMFPP,REGION=0M
//MFPINPUT DD DISP=SHR,DSN=RMFDATA.SYSPLEX
//MFPMSGDS DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSIN  DD *
 SYSRPTS(WLMGL(SCPER))  /* Workload Activity */
 SYSRPTS(CF)            /* Coupling Facility */
 REPORTS(CPU)           /* CPU Activity      */
 REPORTS(DEVICE(DASD))  /* DASD Activity     */
 SYSRPTS(WLMGL(SCPER))  /* Workload Activity */
 SYSRPTS(CF)            /* Coupling Facility */
 REPORTS(CACHE(SUBSYS)) /* Cache Subsystem   */
 RTOD(0800,1700)        /* HHMM to HHMM      */
/*
```

Summary

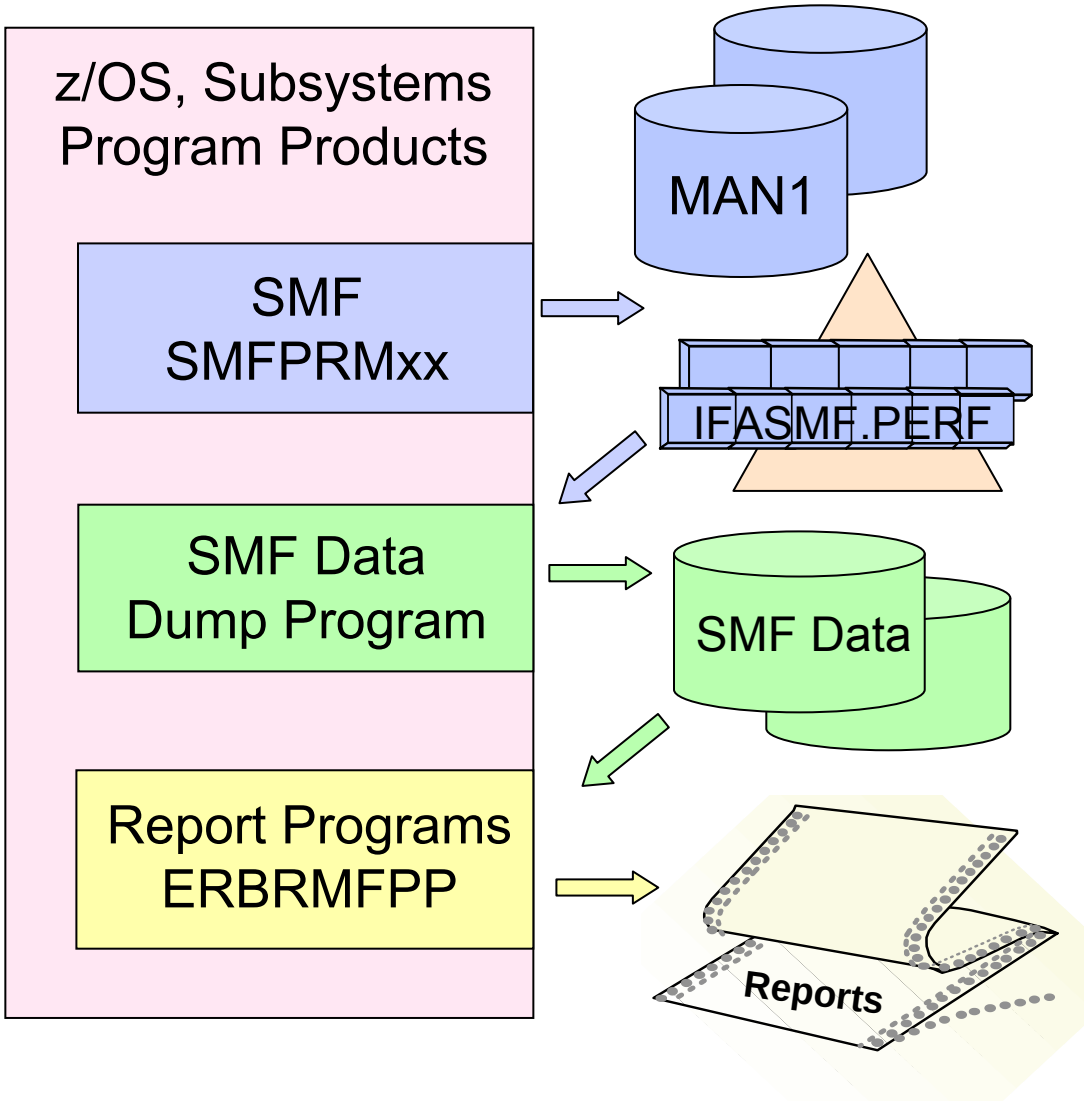

- Performance data reports usage of system resources
- Amount and granularity of SMF and RMF data is controlled by parameters, such as interval value
- SMF and RMF parameters are specified in PARMLIB
- Program to copy SMF data
  IFASMFDP from data sets
  IFASMFDL from logstreams

References

- z/OS Publications
  - MVS System Management Facilities (SMF), SA22-7630
  - MVS Initialization and Tuning Reference, SA22-7592
  - MVS Planning: Workload Management, SA22-7602
  - MVS Setting Up a Sysplex, SA22-7625
  - RMF User's Guide, SC33-7990
- Redbooks
  - SMF Logstream Mode: Optimizing the New Paradigm, SG24-7919
  - ABCs of z/OS System Programming Volume 11, SG24-6327