ibm.com

e-business

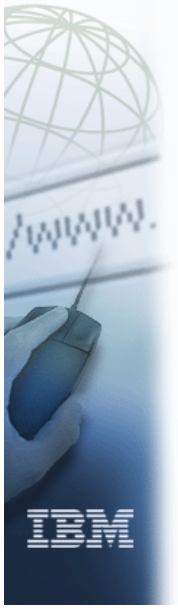

Message Flood Automation

Redbooks

International Technical Support Organization

Trademarks

eNetwork	DFSMS/MVS	IMS	RMF
geoManager	DFSMSdfp	IMS/ESA	RS/6000
AD/Cycle	DFSMSdss	IP PrintWay	S/390
ADSTAR	DFSMShsm	IPDS	S/390 Parallel Enterprise Server
AFP	DFSMSrmm	Language Environment	SecureWay
APL2	DFSORT	Multiprise	StorWatch
APPN	Enterprise System 3090	MQSeries	Sysplex Timer
BookManger	Enterprise System 4381	MVS/ESA	System/390
BookMaster	Enterprise System 9000	Network Station	System REXX
C/370	ES/3090	NetSpool	SystemView
CallPath	ES/4381	OfficeVision/MVS	SOM
CICS	ES/9000	Open Class	SOMobjects
CICS/ESA	ESA/390	OpenEdition	SP
CICS/MVS	ESCON	OS/2	VisualAge
CICSPlex	First Failure Support Technology	OS/390	VisualGen
COBOL/370	FLowMark	Parallel Sysplex	VisualLift
DataPropagator	FFST	Print Services Facility	VTAM
DisplayWrite	GDDM	PrintWay	WebSphere
DB2	ImagePlus	ProductPac	3090
DB2 Universal Database	Intelligent Miner	PR/SM	3890/XP
DFSMS	IBM	QMFr	z/OS
	IBM System z	RACF	z/OS.e

Domino (Lotus Development Corporation)
DFS (Transarc Corporation)
Java (Sun Microsystems, Inc.)
Lotus (Lotus Development Corporation)
Tivoli (Tivoli Systems Inc.)
Tivoli Management Framework (Tivoli Systems Inc.)
Tivoli Manger (Tivoli Systems Inc.)
UNIX (X/Open Company Limited)
Windows (Microsoft Corporation)
Windows NT (Microsoft Corporation)

Console Message Flow Problems

- Large numbers of messages to the z/OS consoles can obscure important messages and delay them from being acted on
- Large numbers of messages to the automation system can delay the processing of normal messages
- Messages can use excessive CPU and storage resources
  - Buffering excessive message traffic may use large amounts of virtual and real storage and it can cause SQA to overflow into CSA
    - This can cause jobs, subsystems, and even complete systems to be delayed or even to fail

Message Flood Implementation

- Use of Message Flood Automation is optional
- You must take overt action
  - To install Message Flood Automation
  - To activate Message Flood Automation

Based on the Message Flood Automation function that has been distributed with the Geographically Dispersed Parallel Sysplex (GDPS) high availability product since 2003

Message Processing before V1R11

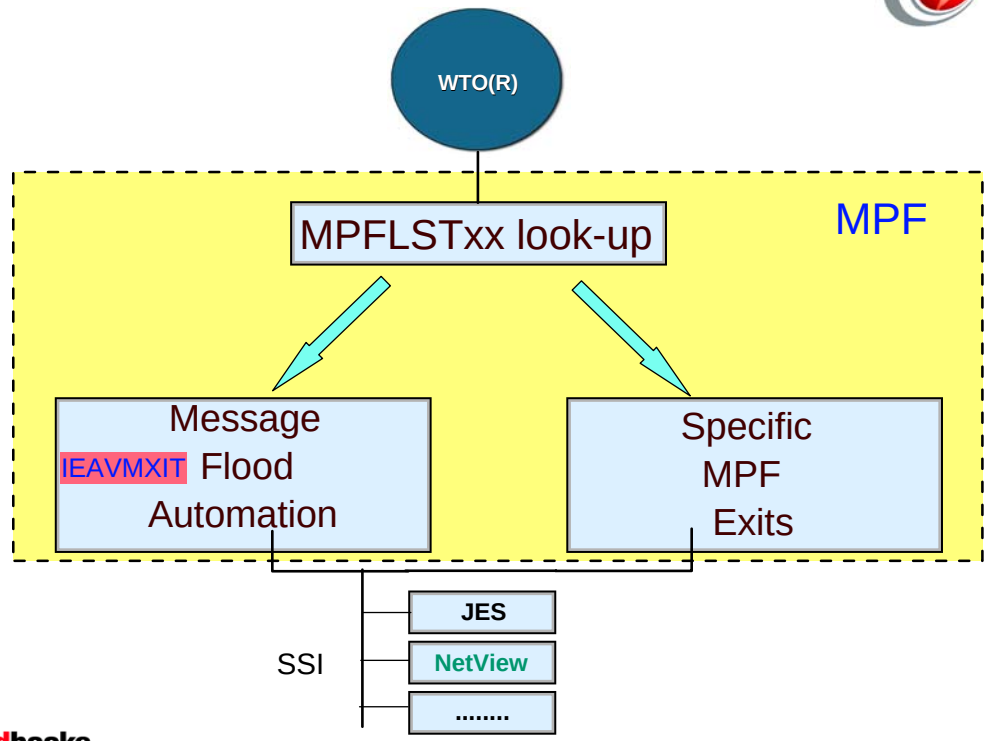

Redbooks

Message Processing with z/OS V1R11

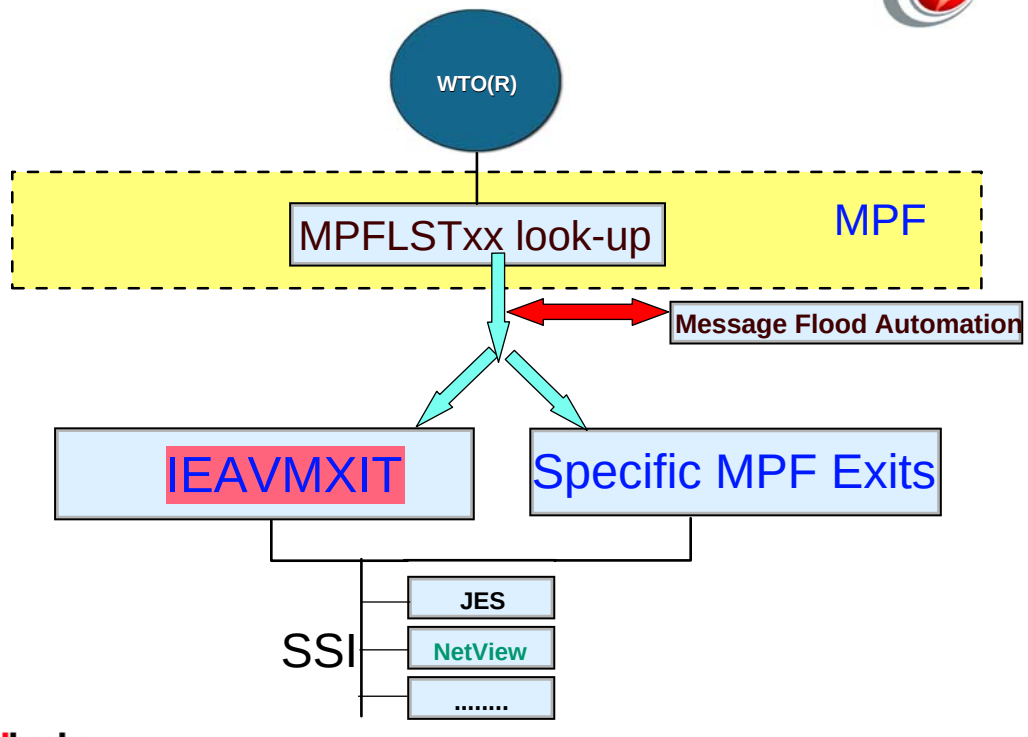

Redbooks

Message Flood Automation - z/OS V1R12

- When it became possible to record SMF records in log streams using System Logger
  - It also became more difficult to prevent system outages due to an unexpected increase of SMF record arrival rates
- When using SMF data set recording
  - The SYS1.MANxx data sets filled up more quickly
- Indicators from the Logger subsystem may happen too late to allow time for correcting the issue
  - With z/OS V1R12, you have the ability to setup rules for matching a flood of SMF data records and either issue a warning message or begin dropping that record type

SMF Message Flood Automation

- Using logstream recording it was difficult to determine the write rate of records
  - With data set recording it was obvious that write rate increased when MANx data sets would cycle quickly
  - Indicators from the Logger subsystem may happen too late and not allow time for correcting the issue
- SMF logstream recording did not offer NOBUFFS or BUFUSEWARN support
  - SMF logstream support shipped without the ability to warn when buffer usage began growing for a given logstream dataspace
  - There was no ability to halt the system in the event that the buffers for a given logstream where exhausted

Message Processing Changes

- OA25602 (z/OS V1R9 & V1R10)
  - Implements the V1R11 changes on previous releases
  - Message restrictions have been removed since IEAVMXIT is no longer used
  - All messages can now be acted upon
  - The limit of 30 SPECIFIC message ids has been increased to 50

SMF Record to SMF Address Space

Pre-z/OS V1R9

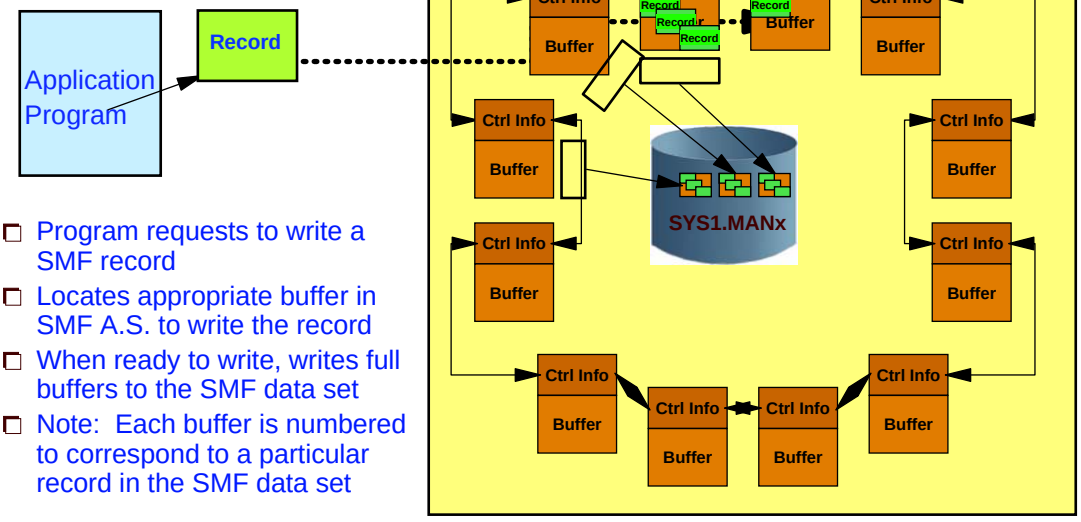

- Program requests to write a SMF record
- Locates appropriate buffer in SMF A.S. to write the record
- When ready to write, writes full buffers to the SMF data set
- Note: Each buffer is numbered to correspond to a particular record in the SMF data set

z/OS V1R9 SMF Records to Log Streams

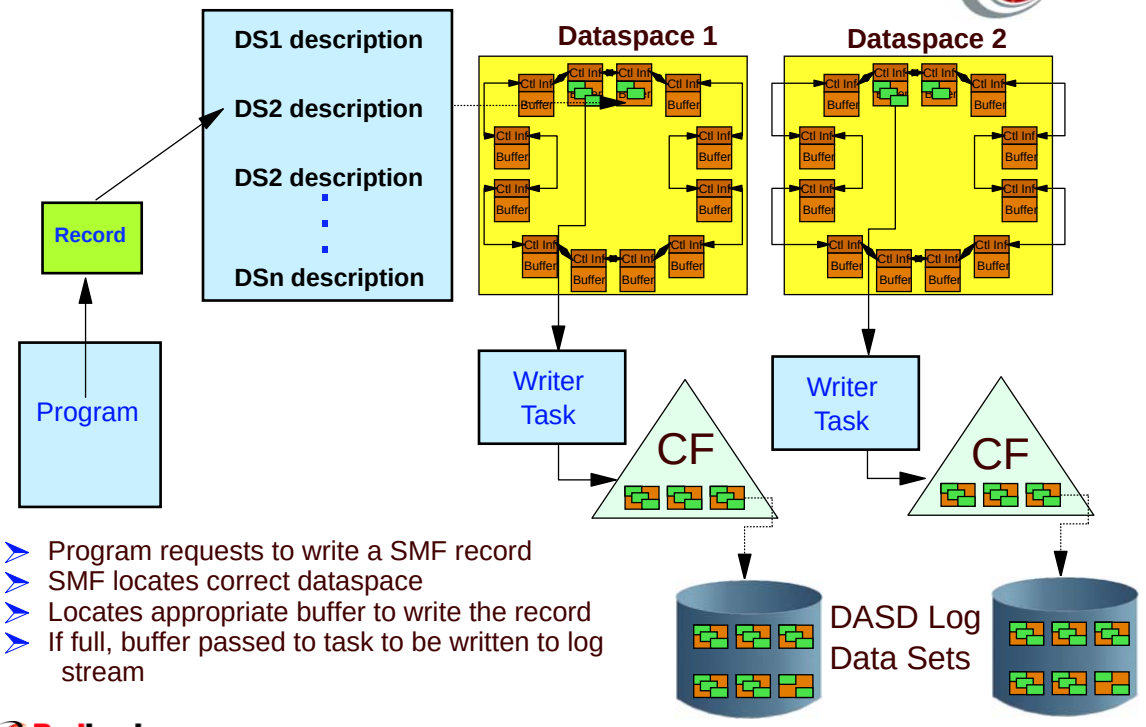

- Program requests to write a SMF record
- SMF locates correct dataspace
- Locates appropriate buffer to write the record
- If full, buffer passed to task to be written to log stream

Redbooks

ibm.com

e-business

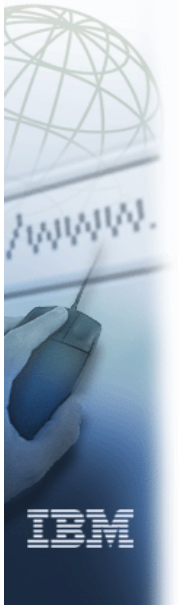

z/OS V1R12 Enhancements

Redbooks

International Technical Support Organization

z/OS V1R12 Enhancements

- Provide the ability to setup rules for matching a flood of SMF data and either issue a message or begin dropping that record type
- Provide support for the NOBUFFS and BUFUSEWARN options for SMF logstream buffers•
- Benefit:
  - These two functions will help provide a more robust environment for SMF recording to logstream

z/OS V1R12 Enhancements

- Get statistical information for existing and new SMF data
  - Keeps historical SMF data to aid in defining flood policies
- Define flood policies to either issue a message or drop the records until the rate returns to normal
  - Detect abnormal SMF behaviour and allow for automated actions to be taken
- Setup individual logstream buffers with different BUFUSEWARN and NOBUFFS settings
  - Each logstream can be configured with it's own BUFUSEWARN and NOBUFFS policy
    - This allows for logstream with important data to be managed more aggressively

SMPPRMxx Parmlib Member

- Choose SMF recording type - RECORDING keyword
  - SYS1.MANx or log stream
- Define log stream names - LSNAME keyword
  - Specify record types by log stream
- Activate parmlib via IPL or SET SMF=xx command
- SETSMF RECORDING can be used to change recording settings (for fallback)

```
RECORDING(DATASET | LOGSTREAM)
DEFAULTLSNAME(logstreamname)
LSNAME(logstreamname,TYPE{aa,bb}|({aa,bb:zz}))
```


Defining SMF System Logging

- If you specify the same record type on two or more different LSNAME parameters
  - System writes the record to all specified log streams
- Could create a SMFPRMxx parmlib member with the following contents:

```
DEFAULTLSNAME(IFASMF.DEFAULT)
  LSNAME(IFASMF.PERF,TYPE(30,89))
  LSNAME(IFASMF.JOB,TYPE(30))
  RECORDING(LOGSTREAM)
```

Buffer Management Setup

- Each logstream can now be targeted with a unique buffer management setup
  - Using the options on the LSNAME or DEFAULTLSNAME keyword will override the global values
- With z/OS V1R12, a new feature allows management of logstream buffers, in the same way that it was possible with SMF data sets - SMFPRMxx parmlib member
  - Now, you can use the NOBUFFS and BUFUSEWARN options for managing SMF log stream buffers

SMFPRMxx Parmlib Member

- NOBUFFS - Specifies the system action when the logstream specified on the LSNAME parameter buffer area is full - NOBUFFS(HALT | MSG)
  - **MSG** - Specifies that the system is to issue a message and continue processing; SMF data is lost until buffer storage is again available
  - **HALT** - Specifies that the system is to enter a restartable wait state. HALT means that no SMF data is lost
- BUFUSEWARN({nn}) - BUFUSEWARN specifies the overall buffer warning level percentage (nn) when SMF starts to issue warning message IFA785E for the logstream specified on this LSNAME parameter

```
IFA785E SMF HAS USED nn% OF AVAILABLE BUFFER SPACE FOR LOGSTREAM lsname
Explanation: This message is issued when nn percent of a given data
space for SMF log stream name lsname has been filled with SMF buffers.
```

NOBUFFS and BFUSEWARN

```
NOBUFFS(MSG)
BUFUSEWARN(75)
LSNAME(IFASMF.PERF.LS,TYPE(..))
LSNAME(IFASMF.CRIT.LS,TYPE(..),NOBUFFS(HALT),BUFUSEWARN(25))
```

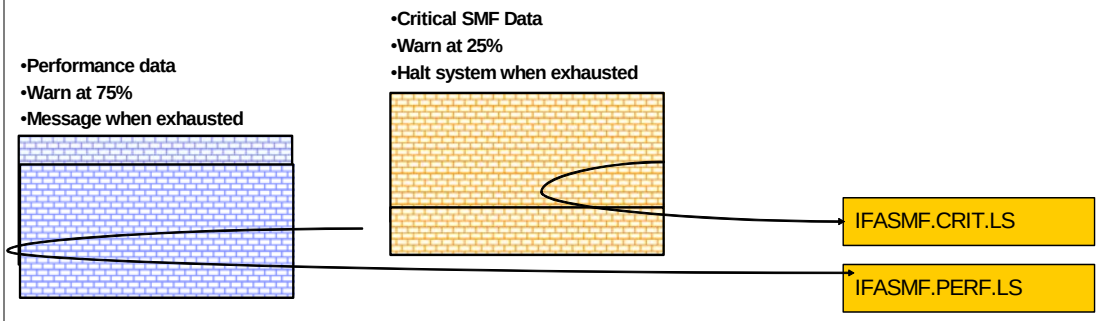

Turning on|off - SMF Flood Automation

- A new option is available for use in the SMFPRMxx parmlib member and the SETSMF command which will turn the facility on and off
  - SETSMF FLOOD(ON|OFF)
- When the facility is turned off all active counts are reset and any active flood actions are ended
  - By default the facility is turned off
- The facility applies to both SMF data set record and SMF logstream recording

SMF Flood Policy with - FLOODPOL

- You can use the FLOOD and FLOODPOL parameters to specify SMF record flood options
  - Use the FLOODPOL parameter in SMFPRMxx to specify a filter for the SMF record in flooding

```
FLOODPOL(
{TYPE({aa,bb})
{aa,bb:zz}
{aa,bb:zz,...}) },
RECTHRESH(xxxx),
INTVLTIME(xxxx),
MAXHIGHINTS(xxxx),
ENDINTVL(xxxx),
ACTION({MSG|DROP})
```


SMF Flood Policy with - FLOODPOL

- TYPE({aa,bb}|{aa,bb:zz} - Specifies the records for this filter
- INTVLTIME(ssss) - Specifies a flood interval time value, given in tenths of seconds
- MAXHIGHINTS(xxxx) - Specifies the number of intervals, defined by INTVLTIME, that must occur at or above the flood rate before action is taken
- ENDINTVL(ssss) - Specifies the amount of time, in tenths of second, that must pass before the RECTHRESH records are generated to determine that a flood state has ended
- RECTHRESH(xxxx) - Specifies the number of records in an interval for this filter
- ACTION({MSG|DROP}) - Specifies the action to be taken when a flood state is entered

Message Flood Automation Examples

- Consider the following SMFPRMxx parmlib option settings when a flood state is entered
  - **MSG** - Issue warning message IFA780A at the start of the flood state - Message IFA781I is issued when the flooding has stopped
  - **DROP** - Issue message IFA782A at the start of the flood state and also begin dropping records - Any attempts to write a record through the SMFEWTM or SMFWTM macro results in a return code 52 - At the end of the flood state, message IFA783I is issued that indicates the number of records that have been dropped

```
FLOOD(ON)
FLOODPOL(TYPE(4,5),RECTHRESH(1000),INTVLTIME(50),MAXHIGHINTS(15),ENDINTVL(120),ACTION(MSG))
FLOODPOL(TYPE(102),RECTHRESH(5000),INTVLTIME(10),MAXHIGHINTS(15),ENDINTVL(100),ACTION(MSG))
FLOODPOL(TYPE(102),RECTHRESH(5000),INTVLTIME(10),MAXHIGHINTS(15),ENDINTVL(50),ACTION(DROP))
```


Message Flood Automation Examples

- The first filter sets up a monitor for both record types 4, 5
- This filter will detect when 1000 records have been generated within 5 seconds, and if records continue to be generated at this rate for 15 consecutive one-second intervals - message IFA780A is issued
- The flood state ends when it takes more than 12 seconds to generate 1000 records - Message IFA781I is issued when the flood state ends

```
FLOOD(ON)
FLOODPOL(TYPE(4,5),RECTHRESH(1000),INTVLTIME(50),MAXHIGHINTS(15),ENDINTVL(120),ACTION(MSG))
FLOODPOL(TYPE(102),RECTHRESH(5000),INTVLTIME(10),MAXHIGHINTS(15),ENDINTVL(100),ACTION(MSG))
FLOODPOL(TYPE(102),RECTHRESH(5000),INTVLTIME(10),MAXHIGHINTS(15),ENDINTVL(50),ACTION(DROP))
```

```
IFA780A SMF RECORD FLOOD MSG FILTER FOR TYPE xx EXCEEDED AT TIME=hh.mm.ss
IFA781I SMF RECORD FLOOD MSG FILTER FOR TYPE xx RETURNED TO NORMAL AT TIME=hh.mm.ss

IFA782A SMF RECORD FLOOD DROP FILTER FOR TYPE xx EXCEEDED AT TIME=hh.mm.ss
IFA783I SMF RECORD FLOOD DROP FILTER FOR TYPE xx RETURNED TO NORMAL AT
TIME=hh.mm.ss, RECORDS DROPPED=yyyyyyyy
```

SMF Flood Policy with - FLOODPOL

- When the MSG filter is specified, message IFA780A is issued at the start of the flood to warn the user
- Message IFA781I is issued when the flooding has stopped
- When the DROP filter is specified, message IFA782A is issued at the start of the flood and dropping of records begins
- Any attempts to write a record through SMFEWTM or SMFWTM macro will get a return code= 52
- At the end of the flooding, message IFA783I is issued with the number of records dropped

New Message Flood Messages

```
New flood messages
IFA782A SMF RECORD FLOOD DROP FILTER FOR TYPE xx EXCEEDED AT
TIME=hh.mm.ss

Explanation: This message is issued when a flood is
triggered and the action for the flood filter is DROP
processing.

IFA783I SMF RECORD FLOOD DROP FILTER FOR TYPE xx RETURNED TO
NORMAL AT TIME=hh.mm.ss, RECORDS DROPPED=yyyyyyyy

Explanation: This message is issued when a DROP
filter has detected the end of a flood.
```

z/OS V1R12 - Dump Programs

- The IFASMFDP and IFASMFDL programs have been updated to allow the gathering of SMF flooding statistics
  - New parameter FLDSTATS(xxxx)
    - The xxxx parameter is the number of records that make a single interval
  - The new report will be generated separately for each z/OS image processed
    - The report will include statistics for each record type encountered

SMF Flood Statistics

- IFASMFDP and IFASMFDL dump programs are updated to allow gathering of SMF flooding statistics
  - FLDSTATS options in the SMF log stream dump programs IFASMFDL and IFASMFDP to display flood related statistics - This value indicates the number of records that make a single interval for the statistical calculations

```
//RUNSMFDL JOB MSGLEVEL=(1,1),NOTIFY=&SYSUID
//SMFDMP EXEC PGM=IFASMFDL
//DUMP04 DD DUMMY
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
LSNAME(IFASMF.SOME.LOGSTREAM,OPTIONS(DUMP))
OUTDD(DUMP04,TYPE(90,101))
FLDSTATS(1000)
```