

# *Bit Bucket X'1C'*

R.P. Shannon, bob.shannon@rocketsoftware.com  
Carl Youngren, cyoungre@hhsdc.ca.gov

Brian Peterson, brian.d.peterson@bluecrossmn.com  
Mike Loos, mikeloos@us.ibm.com  
Skip Robinson, robinsjo@sce.com  
Deb Carnes, deb.carnes@boeing.com

**SHARE 103**  
**Session 2817**  
**New York, NY**  
**August 20, 2004**



# Trademarks

CICS  
RACF

GDPS  
RMF

Lotus Notes  
WebSphere

OS/390  
z/OS

All of the 950+ trademarks of IBM found at <http://www.ibm.com/legal/copytrade/shtml> are incorporated by reference (including "surf shack")

*Omegaview*

Any omissions are completely unintentional

## MPFCMD



Used an old MPF exit (MPFXTALL) off the CBT Tape

It uses the message ID to identify a PDS member, then it submits the member

Unfortunately started to behave poorly. Stopped working altogether in z/OS 1.5

Wrote a new exit (MPFCMD) using same design but based on IEFPRMLB service

Member IEE391A:

```
* S SMF DUMP  
S SMFDUMP
```

# MPFCMD

```
<snip>
HOUSKEEP DS      0H
USING WORKAREA,R13          map workarea
MVC  SAVEAREA+4,=C'FMT1'    indicate stack in use
MVC  MEMBER,=CL8'  '        init member to blanks
MVC  CALLER,=CL8'  '        init caller to blanks
MVC  DDNAME,=CL8'  '        init ddname to blanks
```

```
*      Address the CTXT and Locate the Message
USING CTXT,R10              map ctxt
CLC  CTXTACRN,=C'CTXT'      verify it's the ctxt
BNE  FREESTOR               bif not ctxt
L    R9,CTXTTXPJ            load addr msg attributes
USING CTXTATTR,R9          map attributes
```

```
*      Find Message Id - It will be the first 8 bytes
*      of the message, or fewer than 8 if a blank is
*      encountered
```

```
LA   R1,CTXTTMSG           load message addr
XR   R2,R2                 initialize kount
LA   R15,8                 set loop limit
```

# MPFCMD

```
MSGLOOP  DS      0H
          CR      R2,R15      have we processed 8 chars?
          BNL     MSGLOOPE    bif yes
          CLI     0(R1),C' '   blank?
          BE      MSGLOOPE    bif yes
          LA      R2,1(,R2)    bump kount
          LA      R1,1(,R1)    advance to next character
          B       MSGLOOP     iterate

MSGLOOPE DS      0H
          BCTR    R2,0         decrement kount for execute
          EX      R2,MOVEMSG   move message id
          B       PRMSETUP     branch

MOVEMSG  MVC      MEMBER(0),CTXTMSG  executed instruction

PRMSETUP DS      0H
          MVC     CALLER,=CL8'MPFCMD' set caller id
          L       R0,=AL4(WALEN)     load workarea size
          LA      R1,BUFFER-WORKAREA(0,0) calc size not used for buffer
          SR      R0,R1              calc buffer size
          ST      R0,BUFFSIZE        save buffer size
          AR      R0,R13             calc buffer end addr
          ST      R0,BUFFEND@        and save it
          LA      R8,BUFFER          load buffer addr
          ST      R8,BUFFER@         and save it
```

# MPFCMD

```
USING PRM_READ_BUFFER,R8          map buffer
MVC  PRM_READ_BUFF_SIZE,BUFFSIZE  set buffer size

*      Allocate Parmlib and Read the Member
IEFPRMLB REQUEST=ALLOCATE,ALLOCDNAME=DDNAME,BLANK72=YES,      x
      CALLERNAME=CALLER,MEMNAME=MEMBER,MOUNT=NO,              x
      READ=YES,READBUF=(8),RETMMSG=NO,S99RB=NO,WAITDSN=NO,    x
      RETCODE=RC,RSNCODE=RSN,PLISTVER=MAX,MF=(E,DYNPRMLB)

      LA      R8,PRM_RECORD          load addr 1st record
      LTR     R15,R15                work ok?
      BZ      READREC                bif yes

*      Issue an Error Message Indicating the Parmlib
*      Read Routine Failed
<snip>

READREC DS      0H
      L       R7,BUFFEND@           load buffer end addr
      MVC     DYNMGCRE(MGCRELEN),MGCRE copy static mgcre to dynamic
```

# MPFCMD

```
READREC DS 0H
        L R7,BUFFEND@      load buffer end addr
        MVC DYNMGCRE(MGCRELEN),MGCRE  copy static mgcre to dynamic

RECLOOP DS 0H
        CR R8,R7           end of buffer?
        BNL FREPRMLB      bif yes
        CLI 0(R8),C'*'    comment?
        BE RECLOOPE       bif yes
        CLC 0(80,R8),=CL80' ' blank line?
        BE FREPRMLB      bif yes - all done
        CLC 0(80,R8),=XL80'00' null line?
        BE FREPRMLB      bif yes - all done
        XC CMD(4),CMD      clear flags
        MVI CMDLEN,X'20'   set cmd length to 32 bytes
        MVC CMDTEXT(32),0(R8) move record to command text
        XR R0,R0          reg0 = 0
        LA R2,CMD         reg2 ---> command

MGCRE TEXT=(2),CONSID=CONSID,MF=(E,DYNMGCRE) issue command

RECLOOPE DS 0H
        LA R8,80(,R8)     bump to next record
        B RECLOOP        iterate
```

# MPFCMD

```

        IEFPRMLB REQUEST=FREE,DDNAME=DDNAME,CALLERNAME=CALLER,          x
                RETCODE=RC,RSNCODE=RSN,PLISTVER=MAX,MF=(E,DYNPRMLB)
        B      FREESTOR          branch
        <snip>
SAVEAREA DS      18F          savearea
RC        DS      F          iefprmlb return code
RSN       DS      F          iefprmlb reason code
CALLER    DS      CL8        iefprmlb caller
MEMBER    DS      CL8        iefprmlb member name
DDNAME    DS      CL8        iefprmlb ddname
BUFFEND@  DS      F          end addr of parm buffer
BUFFSIZE  DS      F          buffer size
BUFFER@   DS      F          buffer addr
CONSID    DS      F          X'00000000' (MASTER)
CMD        DS      0D        command
CMDFLAG1  DS      X          1st flag byte (x'00000000')
CMDLEN    DS      X          cmd length (32 bytes max)
CMDTEXT   DS      CL80       command text
WTOLEN    DS      XL2        wto length
WTOFLAGS  DS      XL2        wto flags
WTOTEXT   DS      CL80       wto text
DYNPRMLB  DS      0D,CL(PRMLBLEN) dynamic iefprmlb
DYNMGCRE  DS      0D,CL(MGCRELEN) dynamic mgcre
BUFFER    DS      0D,CL(8192-(*-WORKAREA)) iefprmlb buffer
WALEN     EQU    *-WORKAREA
        IEZVX100          message dsect
        IEFZPMAP          iefprmlb buffer mapping
        IEFZPRC          iefprmlb completion codes

```

# IEAVG700

This code fragment is in TSSO:

```
HRDCONS  DS      0H                                     29OCT92
          STH     R15,MASTCONS                          AND STORE IT...
          ST      R15,CURRCONS                          INDICATE START FROM MSTR?? 30OCT92
          DROP   R3,R14
          L       R3,CVTCUCB                            GET UCM PREFIX
          DROP   R7
          USING  UCM,R3
          LH     R1,UCMBRDST
          LTR    R1,R1
          BM     BRDPROB
          LA     R1,1(R1)
          STH    R1,UCMBRDST
          B      ENDSTP15
```

```
UCMBRDST DC      H'0' -                               COUNT OF REQUESTS TO HAVE WTOS BROADCAST
                                                    TO ALL SUBSYSTEMS
```

Since I had no idea what it did, I asked Scott Fagen. He didn't tell me to RTFM but could have .....

# IEAVG700

In order for messages to be broadcast, UCMBRDST must be > 0

IEAVG700 is the service for this:

- Only need to run once per IPL
- JES usually takes care of it
- Requires supervisor state key 0

```

        USING  SCSRPLST,SCSRLIST
        LA     R1,SCSRPLST
        ST     R1,SCSR@
        XC     SCSRPLST(SCSPLEN),SCSRPLST
        MVC    SCSACRO,=CL4'SCSR'
        MVI    SCSVER,SCSVERSN
        OI     SCSFUNC1,SCSBRDON
        LA     R1,SCSR@
        SPKA   0
        LINK   EP=IEAVG700
        LTR    R15,R15
        BZ     STEP15X

SCSR@    DS    F
SCSRLIST DS    0D,CL(SCSPLEN)
IEZVG100

        map parmlist
        load parmlist addr
        and save it
        clear parmlist
        set acronym value
        set version level
        ind broadcast wto/wtors
        r1 ---> parmlist
        switch to key 0
        call subsys cns1 service
        work ok?
        bif yes

        ---> scsr parmlist
        subsys cns1 srv parmlist
```

# ASASYMBM



Very easy to use service to resolve symbolic parameters:

- Key 8
- Problem state

Not new, but I never needed to use it before

Provide an input (symbolic) and length, and target field

The most difficult thing is when appending a symbolic to a static symbol and worrying about length and position

# ASASYMBM

```
CALLSYMB DS      0H
*           We have a symbolic parameter located in PATTERN
*           Now we populate the symb parameter list in order
*           to call asasymbm
XC      SYMBPWRK(SYMBP_LEN),SYMBP  initialize parmlist
LA      R0,PATTERN                  load pattern addr
ST      R0,SYMBPPATTERN@           store in parmlist
ST      R14,SYMBPPATTERNLENGTH     store length in parmlist
LA      R0,SYMOUT                  load target addr
ST      R0,SYMBPTARGET@           store in parmlist
MVC     SYMOUTLN,=F'8'             set target length
LA      R0,SYMOUTLN                load target length addr
ST      R0,SYMBPTARGETLENGTH@     store in parmlist
LA      R0,RC                      load return code addr
ST      R0,SYMBPRETURNCODE@       store in parmlist
ST      R1,SAVEREG1                save reg 1
LINKX   EP=ASASYMBM,MF=(E,SYMBPWRK),SF=(E,DYNLINKX)

LTR     R15,R15                    work ok?
BNZ     SYMERROR                   bif no
```

# ASASYMBP

SYMBPPATTERN@	DS 1AL4	Address of input pattern containing symbols
SYMBPPATTERNLENGTH	DS 1FL4	Length of input pattern
SYMBPTARGET@	DS 1AL4	Address of output target area. The target, for the length specified in the word pointed to by SymbpTargetLength@, is initialized to blanks before substitution processing begins
SYMBPTARGETLENGTH@	DS 1AL4	Address of input output fullword field containing length of target area. On input, it contains the length provided. On output, it contains the actual length of the target resulting from the symbolic substitution.
SYMBPSYMBOLTABLE@	DS 1AL4	Address of symbol table mapped by SYMBT DSECT. If only the default set of symbols are wanted, and none of the functions indicated by the flags in SymbtFlags are needed, then this field should be 0.
SYMBPTIMESTAMP@	DS 1AL4	Address of 8-character area containing the time stamp to use. If this field is hexadecimal zeroes, the system will obtain the current time stamp for use in evaluating symbolics related to time or date.
SYMBPRETURNCODE@	DS 1FL4	Address of fullword which is to contain the return code

## RACF Dynamic Templates - z/OS 1.5



Will automatically determine whether RACF templates are at the correct level:

- If correct, nothing happens
- If incorrect, will dynamically build the templates from IRRTEMP2 and issue a message:

```
ICH579E RACF TEMPLATES ON DATABASE ARE DOWNLEVEL: 468  
      HRF7707 00000000.00000000; USING TEMPLATES AT LEVEL  
      OA03853 00000010.00000010 FROM IRRTEMP2.  
      RUN IRRMIN00 PARM=UPDATE.
```

IBM still recommends running IRRMIN00

But the good news is that if you forgot to update the templates, a second IPL is no longer required!

# DIAG Traps

Contains all of the traps that can be set in DIAGxx  
Should only be used under IBM's direction

IGVINITCPOOL	Initialize Cpool storage
IGVUNCOND	Make all Freemains and STORAGE RELEASEs unconditional
IGVINITGETMAIN	Initialize GETMAINed storage
IGVNEWPAGE24	Use new page for 24 bit Getmains
IGVNEWPAGE31	Use new page for 31 bit Getmains
IGVDIAGXXABEND	Issue abend for some DIAGxx processing errors
IGVNOUSERKEYCSA	Abend requestors of user key CSA
IGVCPOOLGETV	CPOOL GET validity
IEANOSUSPSYSTRC	Do not suspend system trace when WAITtask is dispatched
IEASCHEDULEV	SCHEDULE validity checking
IEASPINLOCKV	Spin lock validity checking
IEAINITARSRB	Initialize access registers for SRB dispatch. Also does G64H
IEACMSETV	CMSET validity checking
IEASCHEDULETRACE	SCHEDULE tracing
IEARISGNLTRACE	RISGNL tracing
IEARPSGNLTRACE	RPSGNL tracing
IEANOSDWA	(E)STAE(X) and ARR routines get no SDWA
IEANOSUSPSYSTRC	Do not suspend system trace when WAIT task is dispatched
IEASCHEDULEV	SCHEDULE validity checking

# DIAG Traps

IEASPINLOCKV	Spin lock validity checking
IEAINITARSRB	Initialize access registers for SRB dispatch. Also does G64H
IEACMSETV	CMSET validity checking
IEASCHEDULETRACE	SCHEDULE tracing
IEARISGNLTRACE	RISGNL tracing
IEARPSGNLTRACE	RPSGNL tracing
DGNB_IEANOSDWA	(E)STAE(X) and ARR routines get no SDWA
IXCRECSTRALLOC	Do symrec recording for structure alloc
IEAINITREGSTASK	Initialize ARs and G64H for task dispatch
IGVINITFREEMAIN	Initialize FREEMAINED storage
IGVCPPOOLFREEQ	Check for already freed CPOOL cell
CNZTRON	Early SETCON TR=ON
CNZTRONWITHABEND	Early SETCON TR=OnWithAbend
IXLDUPOUTOFSYNCH	Initiate CFCC diag cmd for duplex out of synch conditions
IOSPROTCAPTUCB	Protect captured views of UCBs
CSVRENTSP252	Put all private RENT modules in SP252
CSVRENTPROTECT	Page Protect full pages of RENT modules
IXLBREAKDUPLEX	Initiate SVC dump for break duplex condition
CSVSP252ROUNDUP	Round extent sizes of SP 252 modules up to a page multiple
CSVSP228ROUNDUP	Round extent sizes of SP 228 modules up to a page multiple
CSVSP241ROUNDUP	Round extent sizes of SP 241 modules up to a page multiple
IGVCPPOOLFREEQPXT	Use HEXDATA as a list of PXT addresses for filtering
IARSERIALIZEPIN	Serialize RSMPIN processing

## RMF PM



Free extension to RMF

Java-based, browser interface

Collects and summarizes performance information across Sysplex

Also has collection agent for Linux for zSeries

Explained in "RMF User's Guide"

About 200 possible displays ranging from nearly useless to quite valuable

Did we mention it's free?

# RMF PM

## RMF DDS Browser Interface

Overview

My View

Explore

RMF

Home

,RSPLEX01,SYSPLEX  
% processor utilization by MVS image

Local Time: 08/12/2004 09:48:20

RS22	16	
RS25	10	
RS23	9	
RS21	4	

,RSPLEX01,SYSPLEX  
performance index by important WLM service class period

Local Time: 08/12/2004 09:48:20

DB2.1	26.3	
STCHI.1	18.9	
T80.1	1.15	
STCMD.1	1.11	
BATCH.1	0.7	

,RSPLEX01,SYSPLEX  
i/o intensity by volume

Local Time: 08/12/2004 09:48:20

RS22.TSP110	261	
RS22.MGP11E	73	
RS25.PGP115	43	
RS25.TSP105	19	
RS22.TSP10D	17	
RS22.RZ152A	17	
RS22.S1P106	12	
RS25.S1P106	11	
RS22.JSP100	10	
RS25.PXP102	9	

,RSPLEX01,SYSPLEX  
% CSA utilization by MVS image

Local Time: 08/12/2004 09:48:20

RS22	87	
RS25	25	
RS23	22	
RS21	9	

Automatic refresh in 24 seconds ...

Internet



9:49 AM

## How Long Is Your IPL?

The "IPL Timing Statistics" feature was introduced in OS/390 R3 (and examined in Bit Bucket X'14', 7/2000 in Boston)

System tracks the time to complete IPL/NIP/MSI

Data stored in IPL Statistics control block (IPST)

IPST allocated in SP245 (SQA) by IPL and relocated to SP241 (ECSA) during NIP

The IPST is basically undocumented

## How Long Is Your IPL?

CVT + x'8C'	-	ECVT pointer
ECVT + x'1B0'	-	IPST pointer

### IPST Specifics

+0	Eyecatcher (IPST)
+4	Address of the end of the IPST +1
+8	Address of next available entry
+C	Start of variable length items

### Variable Length Item Specifics

+0	Item_Key - 01 = IPL, 02 = NIP, 03 = MSI
+1	Item_Length - Length of following data
+2	Item_Data - Component specific data

## How Long Is Your IPL?

### Item Data Specifics

#### 01 Entry (IPL)

- +x'00' - 2 character suffix of module
- +x'02' - 4 byte Elapsed Time (in Timer Units)

#### 02 Entry (NIP)

- +x'00' - 2 character suffix of module
- +x'02' - 4 byte Elapsed Time (in Timer Units)

#### 03 Entry (MSI)

- +x'00' - 8 character module name
- +x'08' - 4 byte Elapsed Time (in Timer Units)

## How Long Is Your IPL?

The IPST is NOT an Intended programming interface and is subject to change at any time!

Most of the modules involved are OCO

Elapsed time values are relative to the work involved

Intended to be an IBM diagnostic aid for excessive IPL time problems

Originally, only IBM Level 2 had tools to display data (based on a dump from your system)

Today, IPCS has a function to display the information

## How Long Is Your IPL?

### IPLDATA Function (new with z/OS 1.3)

- Instructs IPCS to display information related to the IPL process

### IPLDATA INFORMATION/STATUS

ACTIVE/DATASET('dsn')/FILE('ddn')/PATH('path')  
FLAG(INFORMATIONAL/WARNING/ERROR/  
SERIOUS/SEVERE/TERMINATING)  
PRINT/NOPRINT TERMINAL/NOTERMINAL  
TEST/NOTEST

“IPLDATA INFORMATION” looks like “D IPLINFO”  
console command

“IPLDATA STATUS” is the display of IPST

# How Long Is Your IPL?

## \*\*\* IPL Statistics \*\*\*

```
IEAIPL10  00:00:00.000  ISNIRIM - Read SCPINFO
IEAIPL20  00:00:05.955  Test Block storage to 2G
IEAIPL11  00:00:00.023  Fast FIND service
IEAIPL31  00:00:00.018  LOAD service
IEAIPL30  00:00:00.020  IPLWTO service
IEAIPL46  00:00:00.318  Read SCHIBs into IPL workspace
IEAIPL49  00:00:00.000  Process Load and Default parameters
IEAIPL50  00:00:00.804  IPL parmlib - process LOADxx and NUCLSTxx
  <snip of 19 lines>
Total IPL Time:  00:00:19.697
```

## \*\*\* NIP Statistics \*\*\*

```
IEAVNIP0  00:00:00.058  NIP Base
IEAVNIPM  00:00:00.305  Invoke NIP RIMs
IEAVNPE6  00:00:00.069  Service Processor Interface
IEAVNPFF  00:00:00.027  Loadwait/Restart
IEAVNPA6  00:00:00.011  RTM - RTCT and recording buffer
IEAVNPC6  00:00:00.021  WTO
IEAVNPC3  00:00:00.010  Issue messages from IPL message queue
IEAVNP24  00:00:00.034  SMS Open/Mount
  <snip of 64 lines>
Total NIP Time:  00:02:02.252
```

# How Long Is Your IPL?

## \*\*\* IEEVIPL Statistics \*\*\*

```
IEETRACE  00:00:00.001  Master trace
ISNMSI    00:00:01.032  SPI
UCMPECBM  00:00:01.835  CONSOLE address space
ENFPC005  00:00:00.000  CONSOLE ready ENF
IEFSCHIN  00:00:00.624  IEFSCHAS address space
IEFJSINT  00:00:00.009  Subsystem interface
IEFSJLOD  00:00:00.054  JESCT
  <snip of 4 lines>
IEEVIPL   00:00:04.136          Uncaptured time:  00:00:00.000
```

## \*\*\* IEEMB860 Statistics \*\*\*

```
ILRTMRLG  00:00:01.148  ASM
IECVIOSI  00:01:19.337  IOS dynamic pathing
ATBINSYS  00:00:00.030  APPC
IKJEFXSR  00:00:00.184  TSO
IXGBLF00  00:00:00.038  Logger
ILMINITM  00:00:00.044  License manager
COMMNDXX  00:00:00.190  COMMANDxx processing
  <snip of 6 lines>
IEEMB860  00:01:29.642          Uncaptured time:  00:00:00.140
Total Time:  00:03:55.728
```

## How Long Is Your IPL?

On 6/25/2004, one contributor to IBM-Main posted  
"*...keep an eye on OA07335..., ...it is shaving 6-15 minutes  
off the IPL time..*"

OA07335 (Base 1.6, PTFs for 1.4, 1.4.1, and 1.5)  
Pre-Reqs OA03030 (GDPS Hyperswap)

New function APAR to "parallelize" IOS initialization  
processing (IECVIOSI) done during MSI processing

- Before: one I/O for each UCB path at a time
- After: multiple I/Os for "a bunch" of UCB paths are done at one time

Question: How do you save 6-15 minutes on something  
that only takes 3:55?

## How Long Is Your IPL?

Non-scientific IECVIOSI (before OA07335) ROTs:

- 10,000 UCBs - 1-2 minutes
- 40,000 UCBs - 7-8 minutes

For our installation (~10,000 UCBs):

- Before OA07335, IECVIOSI took 01:19.337
- After OA07335, IECVIOSI takes 00:57.647

The IBM-Main contributor is now getting a consistent IECVIOSI time of 1 minute

- The 6-15 minutes was frequent, but occasionally much less

The speculation is the excessive time is a hardware "feature"

"Parallelizing" the I/O covers up the paralyzing misbehavior

## How Long Is Your IPL? (Bit Bucket Freebie)

A few “thrifty” installations have only one path to their tape drives

When a Hot I/O occurs on one of those drives, the default is to go through Channel Recovery

- Issue WTOR IOS062E to the operator
- Operator has to respond (correctly) to the complex question
- Wait too long and a WAIT062 occurs

OW51963 (Closed FIN 10/2001) requested the default for IECIOSxx “BOX\_LP=(device\_classes) be changed from “None” to “All”

- Box the device offline rather than doing a channel recovery

In z/OS 1.4, the default was changed to “All” and is actually documented

## RACF Application Identity Mapping (AIM)

Lots of high performance CICS regions, some with extensive RACF controls

One day, CICS regions started blinking red on Omegaview

Icons would go back to green before we could investigate

An unusual number of Help Desk Tickets were opened reporting poor response scattered across the state

Operations reported an intermittent enqueue on **SYSZRAC2**

Fired up RMFMON III to capture whatever we could

# RACF Application Identity Mapping (AIM)

```

RMF V1R2      ENQ Resource Delays
Command ==>
Line 1 of 109
Scroll ==> CSR

Samples: 100      System: S1S1  Date: 07/08/04  Time: 16.10.00  Range: 100  Sec

----- Resource Name -----
Major/Minor      (Scope)
SYSZRAC2         (SYS)
  SYS1.RACF1

----- Delayed -----
% Name          STAT
48 EDMREYWB     SW
48 RH70         SW
48 DMSAR        SW
48 Z33TCLI      SW
48 HWDIRECT     EW
48 EMTDDBM1     SW
47 JES2         SW
47 OMDB2        SW
46 B4IFCR3      SW
45 HDPDWICP     SW
45 HWCIDOR2     SW
43 HWCISTAR     SW
42 BPABNKB      SW
42 B300007      SW
41 CSJIM        SW
40 EU4IR        SW

----- Holding -----
% Name/SYS      STAT
48 MHKXK1       SO
48 MHKXK1       SO
48 MHKXK1       SO
48 MHKXK1       SO
48 MHKXK1       SO
48 MHKXK1       SO
47 MHKXK1       SO
47 MHKXK1       SO
46 MHKXK1       SO
45 MHKXK1       SO
45 MHKXK1       SO
43 MHKXK1       SO
42 MHKXK1       SO
42 MHKXK1       SO
41 MHKXK1       SO
40 MHKXK1       SO

```

## RACF Application Identity Mapping (AIM)

One of our neighbors heard the commotion about MHKXK and reported it was a newly defined USS user

We have had USS for a long time

- Standard MVS stuff (TCP/IP, SMTP, HTTP Server, etc.)
- Been this way from OS/390 R6, through OS/390 R10, and now on z/OS 1.4
- Very little USS work

There was a direct correlation between the enqueues and MHKXK's activities

At this point, the light bulb went on

## RACF Application Identity Mapping (AIM)

Some USS functions (get\_UMAP and get\_GMAP) have a UID/GID and want to know the RACF Userid / Group name

Three generations of RACF behavior

- Search each User/Group profile (default)
- UNIXMAP Profiles (MVS 5.1+)
- AIM (OS/390 R10+)

The original RACF design did not index UID/GID, the only way was to search each profile looking for the first match

- Okay if few IDs and little activity

## RACF Application Identity Mapping (AIM)

UNIXMAP was invented to provide an index

- REXX Execs provided (some actually worked) to create profiles
- Profiles took a lot of space
- Profiles could be manipulated (RALT, Permit, etc) = corruption
- Solves the performance problem

AIM introduced with OS/390 2.10 and is the default for newly created RACF databases

- The documentation takes almost two whole pages in the RACF SPG
- Implied it was for Lotus Notes and Novell uses

AIM uses an "Alias Record" which is a new structure in RACF

Compared to UNIXMAP it is:

- Faster and uses less space
- Not user modifiable
- Not visible

## RACF Application Identity Mapping (AIM)

Use PGM=IRRIRA00 to migrate to AIM

- Transition from Stage 0 (no AIM), to Stage 1, then Stage 2, and then Stage 3

Of course, you have committed the SPG to memory

Study II12972 - Index Reorganization latest information

- No more than 130 logonids with the same UID (OA06979-DOC)
- Run IRRUT200 before to insure there is free space and no structure problems (and create a backup)
- Allow adequate CPU time for execution

OA02721 - re-enables the use of VLF for IRRUMAP and IRRGMAP

## RACF Application Identity Mapping (AIM)

Our conversion (Stage 0 to 1) of 600,000+ logonids (less than 100 have a UID) took 1:08 minutes elapsed (2.67 sec. CPU)

Subsequent Stages took much less time

All the performance problems immediately disappeared!

An informal survey of other shops, shows we weren't the only ones unaware of AIM

Lack of AIM (or UNIXMAP) may not appear to be a problem, but it is!

# Create Your Own InfoCenter



## Eclipse-based

- Product: Platform, Component: Help

## Download from:

- <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ102912>
  - Contains a Pointer to the download site...Roughly 31 MB

Install per included instructions

## installing\_and\_managing.htm

To install the help system, extract the archive containing this document to any location on your system. Ensure that you preserve the folder structure in the archive (I.E.if using WinZip(R), enable "Use folder names option"...)

The help system works with any information packaged as an Eclipse documentation plug-in, including IBM product documentation plug-ins. IBM documentation plug-ins are easily identified because they use a common directory-naming convention (*com.ibm.xxx.doc*)

To install or upgrade documentation plug-ins, complete the following steps:

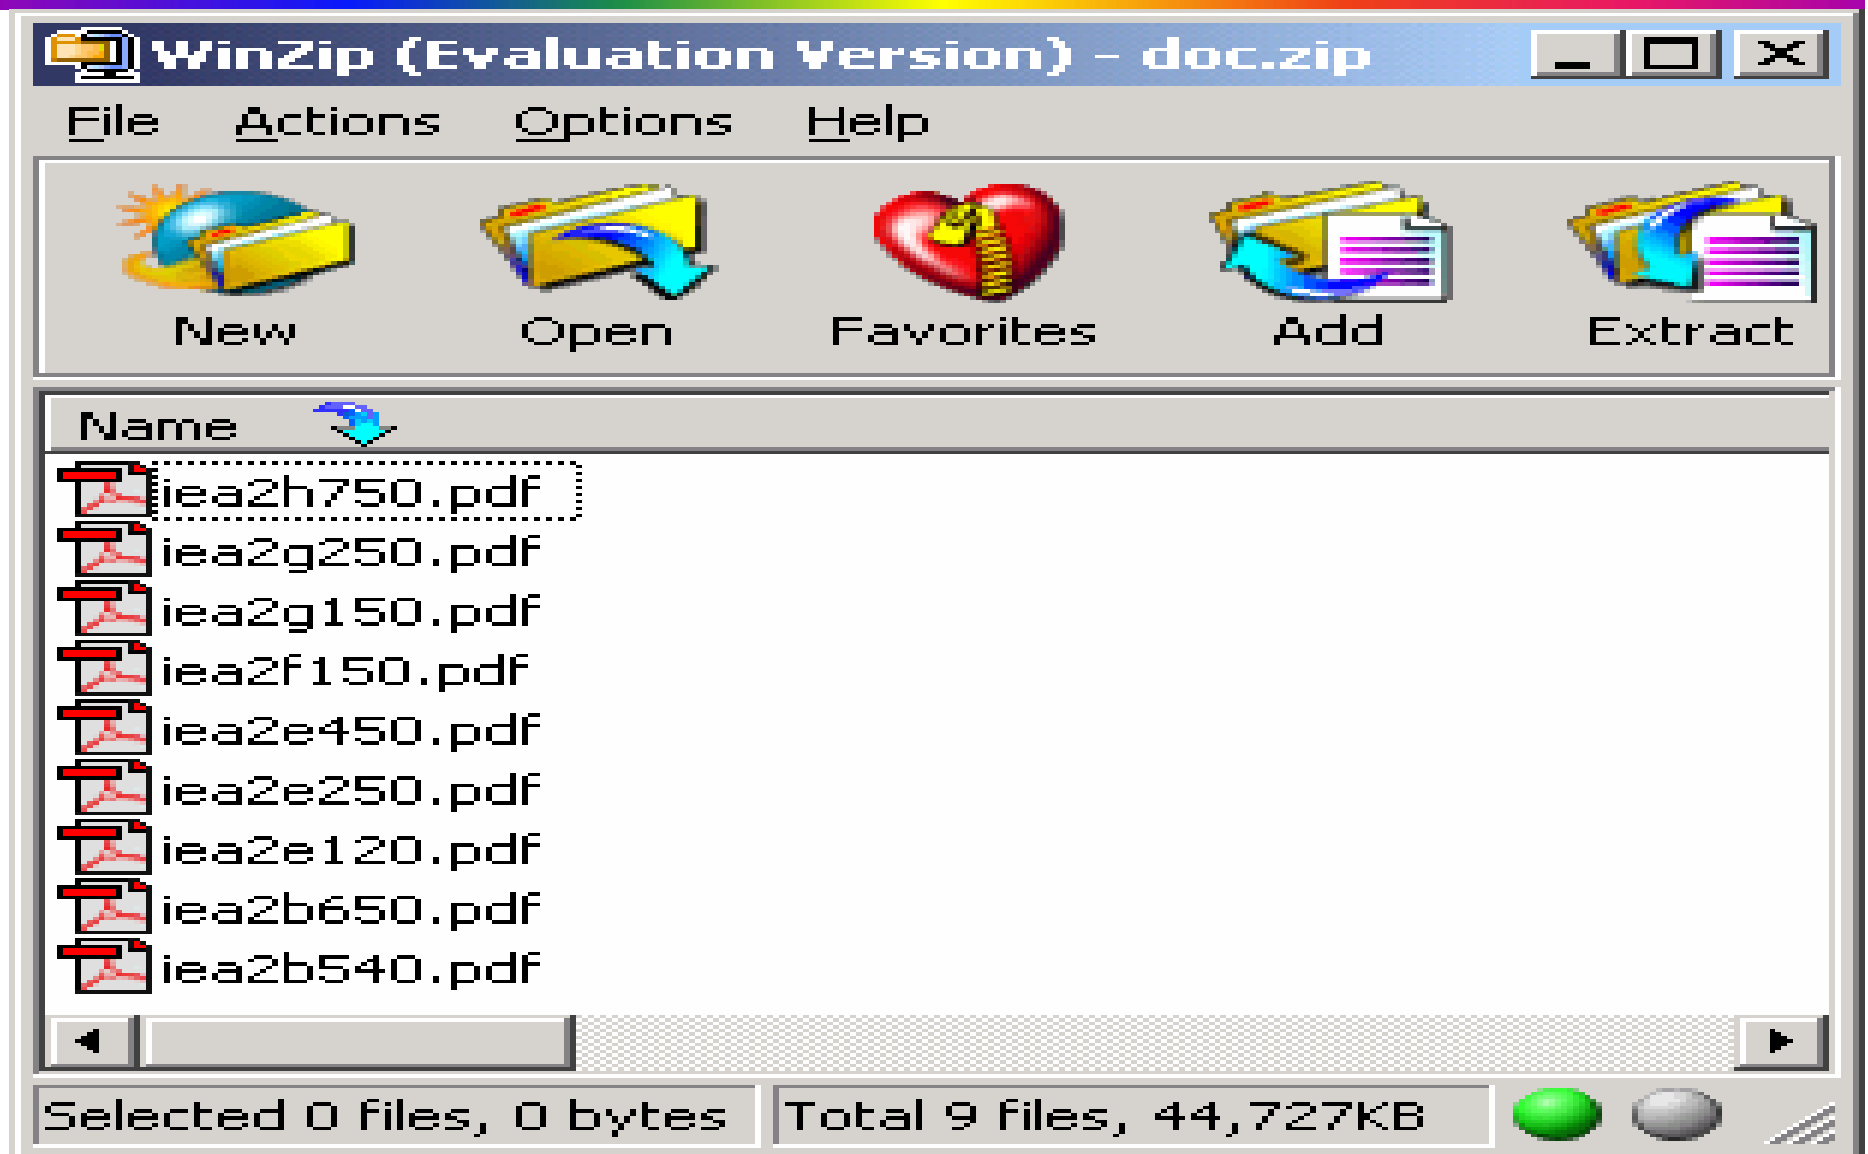
- Obtain the documentation plug-ins. For example, you might go to an IBM product web page to download documentation plug-ins. **Or, you may wish to create your own...(next page)**
- Copy the documentation plug-in folders to the eclipse\plugins folder of the directory where the help system is installed. To see the newly added documentation plug-in, start the help system (or stop and restart it if it was running)

## To Create Your Own Plug-in



- Create a doc.zip file containing the actual documentation
  - .pdf, .txt, .html, .??? open within the existing window
  - .doc, .ppt, .xls, etc. open in a new window
- Create a plugin.xml file
- Create a table of contents - name\_toc.xml

# Create a doc.zip File



## Create a plugin.xml File

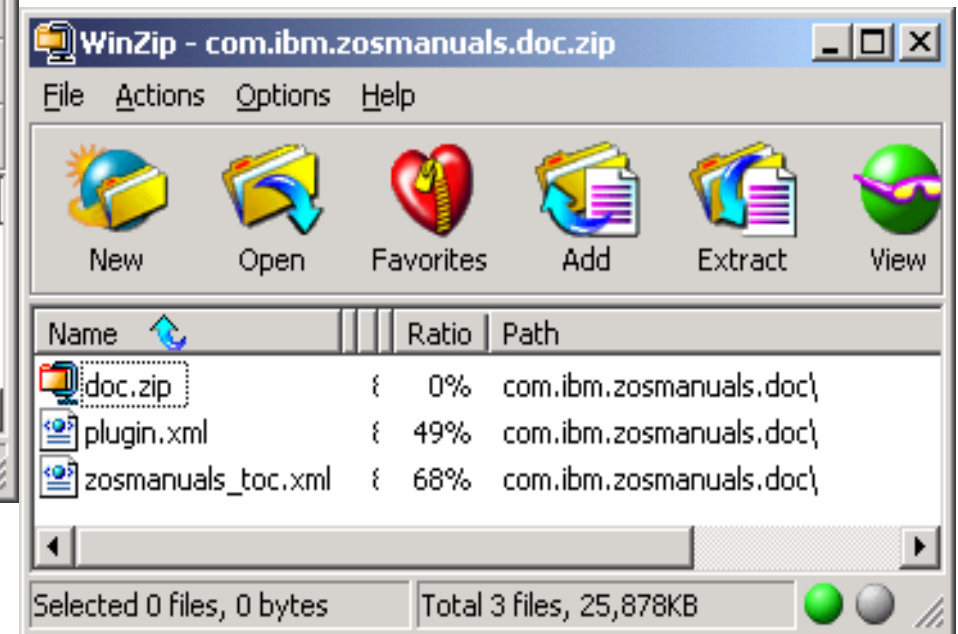
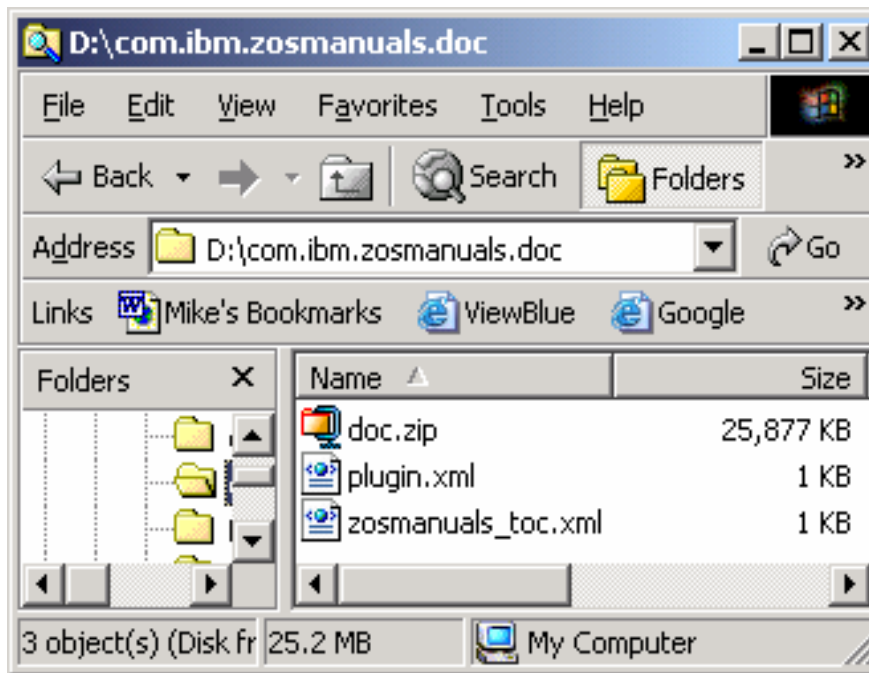
```
<?xml version="1.0" encoding="UTF-8" ?>
<?NLS TYPE="org.eclipse.help.toc"?>
<!-- ===== -->
<!-- This is the plugin for declaring the help -->
<!-- ===== -->
<plugin name = "Selected z/OS manuals"
      id = "com.ibm.zosmanuals.doc"
      version = "5"
      vendor-name = "IBM">
  <extension point="org.eclipse.help.toc">
    <toc file="zosmanuals_toc.xml" primary="true" />
  </extension>
</plugin>
```

# Create a `_toc.xml` File

```
<?xml version="1.0" encoding="UTF-8"?>
<?NLS TYPE="org.eclipse.help.toc"?>
<!--
=====
===== -->
<!-- Define the "info set" -->
<!--
=====
===== -->
<toc label="Selected z/OS manuals" >
  <topic label="MVS JCL Users Guide" href="iea2b540.pdf" />
  <topic label="MVS JCL Reference" href="iea2b650.pdf" />
  <topic label="MVS Init and Tuning Guide" href="iea2e120.pdf" />
  <topic label="MVS Init and Tuning Reference" href="iea2e250.pdf" />
  <topic label="MVS Installation Exits" href="iea2e450.pdf" />
  <topic label="MVS Setting up a Sysplex" href="iea2f150.pdf" />
  <topic label="MVS System Commands" href="iea2g150.pdf" />
  <topic label="MVS SMF" href="iea2g250.pdf" />
  <topic label="MVS System Codes" href="iea2h750.pdf" />
</toc>
```

# Zip 'Em Up and It's a Plug-In...

From a proper directory structure...



Then install the plugin by extracting it into the `\eclipse\plugins` directory, preserving the path name by using the appropriate option during the extract.

# And It Looks Like This...

The screenshot shows a Microsoft Internet Explorer browser window titled "Help - Microsoft Internet Explorer". The address bar shows "http://127.0.0.1:2505/help/index.jsp". The page content is titled "Software information center" and features a search bar with the text "Search scope: All topics".

The main content area is titled "Selected z/OS manuals" and contains a list of chapters. Chapter 25, "IEFUSI — Step Initiation Exit", is highlighted. The right-hand pane displays the content of Chapter 25, including a table of contents for the chapter and a list of topics that appear as follows:

- "Comparing IEFUSI with IEALIMIT when Limiting Region Size" on page 174
- "Locking the Exit in SMP/PMX" on page 174
- "Controlling the Exit Routine Through the Dynamic Exits Facility" on page 175
- "Exit Routine Environment" on page 175
- "Exit Recovery" on page 175
- "Exit Routine Processing" on page 176
- "Programming Considerations" on page 178
- "Entry Specifications" on page 178
  - Registers at entry
  - Parameter Descriptions
  - "Return Specifications" on page 180
  - "Registers at Exit"
- "Examples" on page 182
  - Example of Using IEFUSI to Limit Region Size
  - Examples of Storage Allocations Based on Values Set by IEFUSI

Below the table of contents, the text states: "IEFUSI receives control before each job step is started (prior to allocation). A return code from this exit indicates whether the job step is to be started or the job should be cancelled."

It then lists the uses of IEFUSI:

- You can use IEFUSI to:
  - Validate job step accounting information.
  - Write to a user data set.
  - For long-running jobs, create and write a user step-initiation SMF record in case of system failure.

Additional information includes:

- Set the region size and region limit for all programs that run under this job step. For more information on controlling region size and region limit, see [z/OS MVS Initialization and Tuning Guide](#).
- Set limits on the use of data spaces and hiperspaces created by application programs with storage key 8-F.
- Limit the number of pages that can be shared at one time through the use of the IAFVSEPV macro.
- Set the default size of data spaces and hiperspaces.
- Limit the use of the 16 exabyte address space above two gigabytes.
- Reduce the values of LDALJLK, LDKEYVRG, LDALIMIT, and LDAPVRG. Initially, these fields contain the maximum amount of storage available to the user. Specifically:

## Recent Adventures With Standalone Dump



Recent instability (JES2) prompted a review of SA Dump  
(JES2 instability turned out to be mostly self-inflicted)

Review of SA Dump turned out to be Very Timely!

IBM has made many improvements in SA Dump

More enhancements coming in z/OS 1.6; This experience  
based on 1.4

## Recent Adventures With Standalone Dump

Manual documents plethora of options, then recommends a subset. For my system, the following made sense:

```
//OSG          EXEC PGM=AMDSAOSG
//SYSLIB       DD    DSN=SYS1.MACLIB,DISP=SHR
//            DD    DSN=SYS1.MODGEN,DISP=SHR
//GENPRINT     DD    SYSOUT=*
//GENPARMS    DD    *
               AMDSADMP IPL=DSYSALLDA,VOLSER=SDST00, <== SADUMP IPL VOLUME X
               OUTPUT=D620C, <===== ADDRESS OF VOL=SER=SSD001 X
               DUMP=('SP(ALL) IN ASID(1,'JES2','JESXCF') X
               ALSO DATASPACE OF ASID(1,'DUMPSRV','JESXCF','APPC','SMSX
               VSAM','CONSOLE','SMSPDSE') X
               ALSO PAGETABLES OF DATASPACE'), X
               CONSOLE=SYSC,MINASID=PHYSIN, X
               MSG=ALL,REUSED=ALWAYS
               END
```

**REUSED=ALWAYS** - use SA Dump disk data sets no matter what  
**CONSOLE=SYSC** - use the HMC for the console  
**DUMP= Options** - recommended by manual  
plus **SMSPDSE**, plus **DUMPSRV**

# Recent Adventures With Standalone Dump

8:22:50 am Thursday morning - Kaboom!

New SA Dump in action

SA Dump timestamps - 1 minute 29 seconds - 9 gig LPAR

\*\*\* STAND-ALONE DUMP MESSAGE LOG

08:38:10.12 AMD083I AMDSADMP: STAND-ALONE DUMP INITIALIZED

08:38:11.29 AMD005I DUMPING OF REAL STORAGE NOW IN PROGRESS.

08:39:37.11 AMD102I FINAL DUMP OPTIONS:

08:39:37.11 CSA ALSO LSQA, SP(229:230,236:237,249) IN ASID(PHYSIN)

08:39:37.11 ALSO SP(ALL) IN ASID(1,'JES2','JESXCF')

08:39:37.11 ALSO DATASPACE OF ASID(1,'JESXCF','APPC','SMSVSAM','CONSOLE','

08:39:37.11 ALSO PAGETABLES OF DATASPACE

08:39:39.44 AMD056I DUMPING OF VIRTUAL STORAGE COMPLETED.

08:39:39.51 AMD104I DEVICE VOLUME USED DATA SET NAME

08:39:39.51 1 620C SSD001 39% SYS1.SADMP

08:39:39.51 8 6598 SSD008 41% SYS1.SADMP

# Recent Adventures With Standalone Dump

## Lessons Learned

- Read the manual - things change/improve
- Dump to disk - disk is cheap
- Review dump options
- IPL SA Dump
  - \* Load address (SA dump SYSRES)
  - \* Load parameter (SA dump options)
- Load Parameter - constant "S" followed by option "M"
  - "SM" means use defaults for console, dump device, and dump title, with no operator prompting
- View messages using HMC console (always there)

## Recent Adventures With Standalone Dump (Epilog)

Remember the time (8:22:50 AM)? The system didn't die (unfortunately) until we stopped it at 8:31 AM

SYSTRACE of course wraps frequently - on busy system, many times per minute (see OA07896, changes default table size)

System sets aside certain SYSTRACE entries for just this reason.

TTCH: System Trace Table Snapshot Copy Header

```
ipcs systrace ttch(list) time(local)
```

TTCH	ASID	TCB	TIME
71E80000	0064	009DBD90	08/05/2004 08:30:00.051179
71F42000	04D4	0095A1F0	08/05/2004 08:29:49.752802
7F436000	04CA	009A3D60	08/05/2004 08:22:51.686201
7F4F8000	0001	00000000	08/05/2004 08:22:50.868423 <--- !
7F5BA000	001C	009E3B60	08/05/2004 08:22:50.858826

The second to the last entry turned out to provide the answer (which happened to be an ISV software problem)

*Bit Bucket X'IC'*

## Enhanced Data Integrity (EDI)

New DFSMS (1H0) function introduced in z/OS 1.5

Prevents multiple tasks writing SHR to a sequential file

For example, two jobs both have -

```
// DD DISP=SHR,DSN=A.B.SEQ
```

Both jobs open file for output and this has always been allowed. However results can be unpredictable. But some applications are designed to work this way!

With EDI you control the situation: allow or disallow

## EDI Overview

EDI works from tables built in storage. By default, nothing happens

Two actions are required to use EDI

- Create member IFGPSEDI in PARMLIB  
No default member is supplied. The member name is fixed and is either present or not
- Start task IFGEDI to build tables from IFGPSEDI  
The PROC is supplied in SYS1.IBM.PROCLIB. It should run SUB=MSTR early (e.g. from COMMNDxx). The task runs and disappears, but can be run any number of times to rebuild tables

## EDI Overview



Three possible modes:

- WARN - messages only, no failures
- DISABLE - no action at all
- ENABLE - non-excluded violators are ABENDED

Excluded datasets can be named generically with "\*" and "%" having the usual meanings. For performance, don't have too many exclusions

# EDI PROC - IEFGEDI

```
//IFGEDI PROC
//*****
//*$MAC(IFGEDI); DFSMS OCE EXTENDED DATA INTEGRITY STARTED
//* TASK
//*PROPRIETARY V3 STATEMENT
//*LICENSED MATERIALS - PROPERTY OF IBM
//*"RESTRICTED MATERIALS OF IBM"
//*5694-A01
//*(C) COPYRIGHT 2003, 2004 IBM CORP.
//*END PROPRIETARY V3 STATEMENT
//* FUNCTION -
//* THIS PROCEDURE IS USED TO START DFSMS OCE STARTED
//* TASK IFG0DSNT TO BUILD 2 IN STORAGE TABLES: 1) DATA SET
//* NAMES CONTAINED IN PARMLIB MEMBER IFGPSEDI. 2) COMMON
//* FILTER SERVICES TABLE OF POINTERS AND LENGTH OF DATA SET
//* NAMES IN THE DSN TABLE.
//*$00=EDI,HDZ11H0,020630,SJPLJAL: OCE EXTENDED DATA INTEGRITY
//* SUPPORT
//*$P1=K1H0620,HDZ11H0,020924,SJPLGEC: ADD TIME PARM TO AVOID @P1A
//* ABEND 322(TIME OUT) @P1A
//*$01=OA06021,HDZ11H0,040130,SJPLJAL: REDESIGN $01A
//*****
//EDIST EXEC PGM=IFG0DSNT,REGION=512K,TIME=20 $01C
```

# IFGPSEDI for Warn Mode

```
MODE(WARN) /* Must be first line in this member! */
/*
/* Member controls DFSMS EDI (Enhanced Data Integrity),
/* which prevents multiple OPEN/SHR/WRITE applications from
/* clobbering each other.
/*
/* SYNTAX:
/*
/*     MODE{ (WARN|ENFORCE|DISABLE) } [ ,DSN(dsn) [ ,DSN(dsn) ] ...
/*
/* where DSN is a dataset to be excluded from control. DSN may
/* contain generic characters * and % .
/*
/* The controlling task must be started to enable EDI:
/*
/*     S IFGED1,SUB=MSTR
```

## EDI in Warn Mode

Job EDIJOB1 is already writing to a sequential file when Job EDIJOB2 begins writing to the same file. In warn mode, EDI issues a warning message:

```
IEF403I EDIJOB2 - STARTED - TIME=04.56.35
IEC813I EDI OWNER: JN=EDIJOB1 ,SYSNAME=B0,
        TCB=00AD0D40,ASID=00A1
IEC984I EDIJOB2,A812,BR3003,TED066.TESTEDI.DATA
        NOT EXCLUDED BUT HAD AN ENHANCED DATA
        INTEGRITY OUTPUT VIOLATION
```

# IFGPSEDI for Enforce Mode

```
MODE(ENFORCE) /* MUST BE FIRST LINE IN THIS MEMBER! */
/*
/* Member controls RACF EDI (Enhanced Data Integrity),
/* which prevents multiple OPEN/SHR/WRITE applications from
/* clobbering each other.
/*
/* SYNTAX:
/*
/*     MODE{(WARN|ENFORCE|DISABLE)}[,DSN(dsn)[,DSN(dsn)]...
/*
/* where DSN is a dataset to be excluded from control. DSN may
/* contain generic characters * and % .
/*
/* The controlling task must be started to enable EDI:
/*
/*     S IFGED1,SUB=MSTR
```

## EDI in Enforce Mode

Job EDIJOB1 is already writing to a sequential file when job EDIJOB2 begins writing to the same file. In Enforce mode, EDI ABENDs the second job:

```
IEF403I EDIJOB2 - STARTED - TIME=05.00.31
IEC813I EDI OWNER: JN=EDIJOB1 ,SYSNAME=B0,
        TCB=00AD0D40,ASID=0068
IEC143I 213-FD,IFG0194D,EDIJOB2,STEP1,SYSPRINT,
        A812,BR3003,TED066.TESTEDI.DATA
IEA995I SYMPTOM DUMP OUTPUT 009
SYSTEM COMPLETION CODE=213 REASON CODE=000000FD
...
IEF450I EDIJOB2 STEP1 - ABEND=S213 U0000 REASON=000000FD
```

# EDI IFGPSEDI with Dataset Exclusion(s)

```
MODE(ENFORCE), /* MUST BE FIRST LINE IN THIS MEMBER! */
DSN(TED066.TESTEDI.DATA)
/*
/* Member controls RACF EDI (Enhanced Data Integrity),
/* which prevents multiple OPEN/SHR/WRITE applications from
/* clobbering each other.
/*
/* SYNTAX:
/*
/*     MODE{ (WARN|ENFORCE|DISABLE) } [ ,DSN(dsn) [ ,DSN(dsn) ] ...
/*
/* where DSN is a dataset to be excluded from control. DSN may
/* contain generic characters * and % .
/*
/* The controlling task must be started to enable EDI:
/*
/*     S IFGED1,SUB=MSTR
```

## EDI Allowing Excluded Dataset

Job EDIJOB1 is already writing to a sequential file when EDIJOB2 begins writing to the same file. EDI has nothing at all to say:

```
IEF403I EDIJOB2 - STARTED - TIME=05.04.40  
IEF404I EDIJOB2 - ENDED - TIME=05.05.15
```

# Bibliography

APARs II12972, OA02721, OA06979, OA07335, OW51963

"Anonymous Sources"; 1-800-237-5511; IBM Corporation

## All manuals of the IBM Corporation:

GA22-7589-05 "z/OS V1R4.0 MVS Diagnosis Tools and Service Aids"

SA22-7592-06 "z/OS V1R4.0 MVS Initialization and Tuning Reference"

SA22-7681-03 "z/OS V1R4.0 Security Server RACF System Programmer's Guide"

GC26-7398-06 "z/OS V1R5.0 DFSMS Migration Guide"

SC26-7410-03 "z/OS V1R5.0 DFSMS: Using Data Sets"

SA22-7592-07 "z/OS V1R5.0 MVS Initialization and Tuning Reference"

SA22-7606-05 "z/OS V1R5.0 MVS Programming Assembler Services Reference Vol 1  
(ABEND-HSPSRV)"

SA22-7642-03 "z/OS V1R5.0 MVS Using the Subsystem Interface"

SC33-7990-07 "z/OS V1R5.0 Resource Measurement Facility (RMF) User's Guide"

SA22-7683-04 "z/OS V1R5.0 Security Server RACF Security Administrator's Guide"

## Contributors, Willingly or Unwillingly, Knowingly or Unknowingly:

Dave Danner - Summit Technical Services

Scott Fagen - IBM Corporation

Russ Hardgrove - IBM Corporation

Mike Keyes - IBM Corporation

Jerry Ng - IBM Corporation