9805: Recent z/OS Debugging Enhancements The Dark Arts

A Presentation by Nicole M. Fagen

August 2011

Agenda

SLIPs

IPCS Active

COPYDDIR

Dump Health Scan

GRS

Summary Format

System Trace

DOCPU

Dark Art

Art that is broadly defined as disturbing or horrific in nature

Roots are in horror

Synonyms: gothic, horror, nightmarish and disturbing

It is actually a culmination of techniques and styles

After the horror of a defect was encountered and the long torturous path to resolution suffered by many highly passionate skill technical professionals IPCS and SLIP were created. Both SLIP and IPCS have been built upon ever sense.

Dynamic System Name with REMOTE

Horror: The names of the REMOTE systems to be included in a dump must be “hardcoded” and declared when the slip was set

New Technique: Direct or indirect addresses can be used to look up system names when the trap matches

–System names assumed to be 8 characters

z/OS 1.9

Dynamic System Name Example

Problem Situation: A system in the sysplex is going in and out of the state of being status update missing, IXC426D is issued on the system detecting the temporary hang. The stall is temporary and the impacted system varies. A timely dump on the stalled system is desired.

- **IXC426D SYSTEM system IS SENDING XCF SIGNALS BUT NOT UPDATING STATUS. REPLY SYSNAME= system TO REMOVE THE SYSTEM OR R TO RETRY**

```
SLIP SET,MSGID=IXC426D,A=SVCD,
JOBLIST=(XCFAS),DSPNAME=('XCFAS'.*),
SDATA=(XESDATA,COUPLE,ALLNUC,CSA,PSA,LPA,LSQA,NUC,RGN,SQA,
SUM,SWA,TRT),
REMOTE=(UNCOND,SYSLIST=((3R?+F)),DSPNAME,JOBLIST,SDATA),END
```

This example works if and only if the system name is 8 characters!

z/OS 1.9

Gotch-ya

If only a dynamic system name is requested then the display of the slip does not indicate there is a SYSLIST.

```
RESPONSE=SD0
 IEE735I 16.38.48 SLIP DISPLAY 615
 ID=0011,NONPER,ENABLED
 ACTION=SVCD,SET BY CONS NFAGEN,RBLEVEL=ERROR,MATCHLIM=1,0
 MSGID=IXC426D,DSPNAME='XCFAS'.*
 SDATA=PSA,NUC,SQA,LSQA,RGN,LPA,TRT,CSA,SWA,SUMDUMP,ALLNUC,COUPLE,
       XESDATA
 REMOTE=(
 UNCOND
 JOBLIST,SDATA,DSPNAME
        )
 JOBLIST=XCFAS
```

If a hardcode system name is coded along with a dynamic system name then both appear in the display

```
RESPONSE=SD0
 IEE735I 16.39.40 SLIP DISPLAY 890
 ID=0012,NONPER,ENABLED
 ACTION=SVCD,SET BY CONS NFAGEN,RBLEVEL=ERROR,MATCHLIM=1,0
 MSGID=IXC426D,DSPNAME='XCFAS'.*
 SDATA=PSA,NUC,SQA,LSQA,RGN,LPA,TRT,CSA,SWA,SUMDUMP,ALLNUC,COUPLE,
       XESDATA
 REMOTE=(
 UNCOND
 SYSLIST=SB0,CURRENT.3R?+F
 JOBLIST,SDATA,DSPNAME
        )
 JOBLIST=XCFAS
```

z/OS 1.9

SDUMP Started Message

Horror: Dump is taking a long time but message log does not indicate when the dump started

New Technique: IEA045I issued to indicate SDUMP started

IEA045I AN SVC DUMP HAS STARTED AT TIME=*hh:mm:ss*
DATE=*mm/dd/yyyy* **FOR ASIDS(***xx***[,***xx,...,xx***])**
ERRORID=SEQ*yyyyyy* **CPU***zz* **ASID***asid* **TIME***hh:mm:ss.t*
QUIESCE=YES|NO

z/OS 1.11

IPCS ACTIVE

Horror: IPCS Active only lets users browse and analyze storage visible to key 8 applications.

New Technique: Access to more data!!

- Dataspaces owned by the ASID and visible to key 8 applications are now supported for IPCS users with no special authority
- Users authorized to BLSACTV ADDRSPAC can browse and analyze all storage in the ASID and its data spaces
- Users authorized to facility class BLSACTV SYSTEM can browse and analyze all systems ASIDs and data spaces, as well as, absolute storage
- Keep in mind the view is UNSERIALIZED. The system will not be stopped to execute an IPCS command.

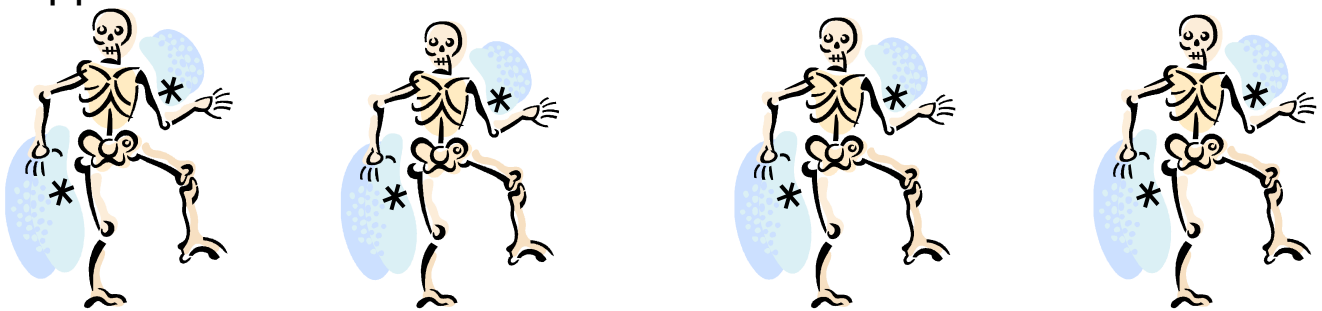

z/OS 1.9

COPYDDIR EXPORT / IMPORT

Horror: Having to engage additional assistance to pursue root cause on a huge standalone dump which requires the new person to initialize the huge standalone dump

New Technique: COPYDDIR EXPORT / IMPORT

EXPORT the directory information to a common dataset

IMPORT the directory to the new dump directory

z/OS 1.10

```
IPCS INVENTORY - NFAGEN.ZOS1B5.DIRECTRY ------------------------------------
Command ===>

AC Dump Source                                                        Status
__ DSNAME('ONTOP.GS000.P03548.C672.DUMP.PP1A.PB9AIRLM') . . . . . . . CLOSED
   Title=ABEND=S026,REASON=08118001,CONNECTOR HANG: CONNAME=DXRPJ0A$$PJ9A009,JO
   Psym=RIDS/IEANUC01#L RIDS/IXLM1TMR VALU/C$PJ9A009 VALU/H0003 PIDS/5752SCIXL
__ DSNAME('ONTOP.GS075.P61541.C724.DUMP2') . . . . . . . . . . . . . . CLOSED
   Title=SLIP DUMP ID=CJP1
   Trap=SLIP SET,ENABLE,ID=CJP1,COMP=0C4,JOBNAME=USACHPZ2,ACTION=SVCD,JOBLIST=(
__ DSNAME('ONTOP.GS100.P43129.C838.LOGGER.DUMP') . . . . . . . . . . . CLOSED
   Title=CQE1MSGL STRUCTURE DUMP
   No symptoms
__ DSNAME('ONTOP.GS113.P43516.C848.D205225') . . . . . . . . . . . . . CLOSED
   Title=ABEND=S026,REASON=08118001,CONNECTOR HANG: CONNAME=SIGPATH_030014DF,JO
   Psym=RIDS/IEANUC01#L RIDS/IXLM1TMR VALU/C030014DF VALU/H0008 PIDS/5752SCIXL
__ DSNAME('ONTOP.GS113.P43516.C848.D205436') . . . . . . . . . . . . . CLOSED
   Title=ABEND=S026,REASON=08118001,CONNECTOR HANG: CONNAME=CSQEQT00QT0202,JOBN
   Psym=RIDS/IEANUC01#L RIDS/IXLM1TMR VALU/C00QT0202 VALU/H0011 PIDS/5752SCIXL
__ DSNAME('ONTOP.GS113.P43516.C848.D205736') . . . . . . . . . . . . . CLOSED
   Title=END OF MEMORY RESOURCE MANAGER HANG DETECTED: TCB = 007EC0A8, NAME = I
   Psym=RIDS/IEAVTMMW#L RIDS/IEAVTMMW VALU/C###SCSDS PIDS/5752SCRTM AB/S030D RI
__ DSNAME('ONTOP.GS113.P43516.C848.D210927') . . . . . . . . . . . . . CLOSED
   Title=IXC102A NOT ANSWERED
   No symptoms
XP DSNAME('ONTOP.GS113.P44137.C848.SADMP') . . . . . . . . . . . . . . OPEN
 F1=HELP   F2=SPLIT   F3=END   F4=RETURN   F5=RFIND   F6=MORE   F7=UP   F8=DOWN
```

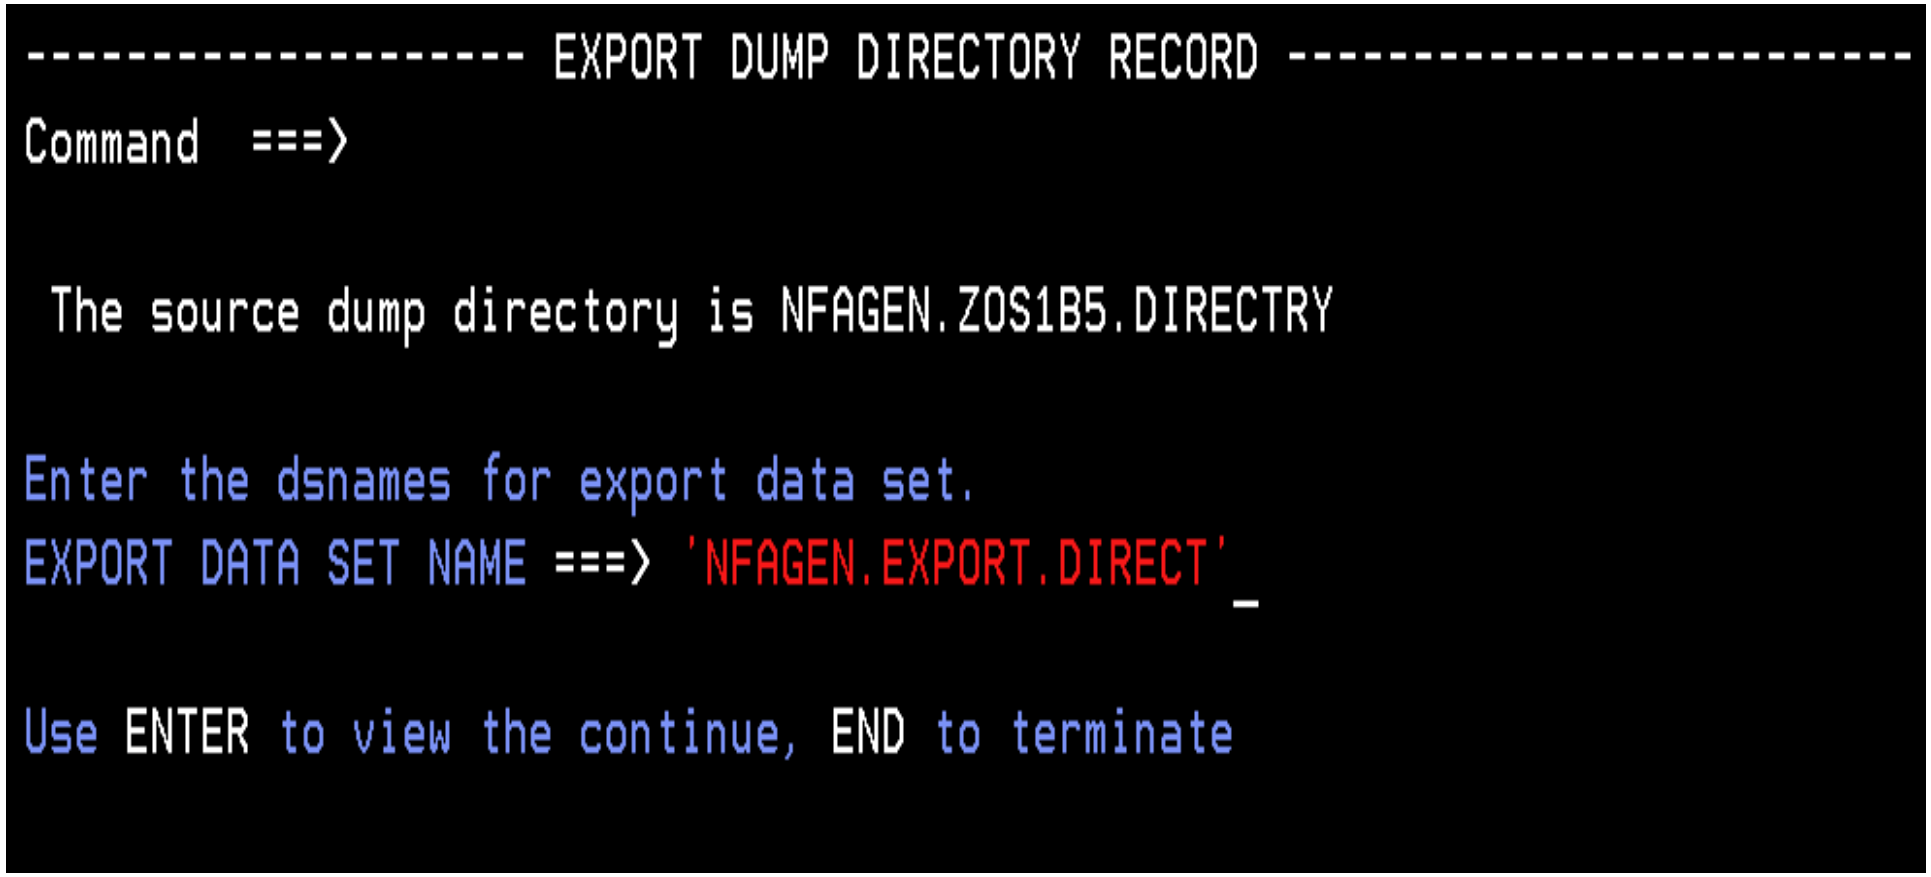

Place data in a dataset the new user can access

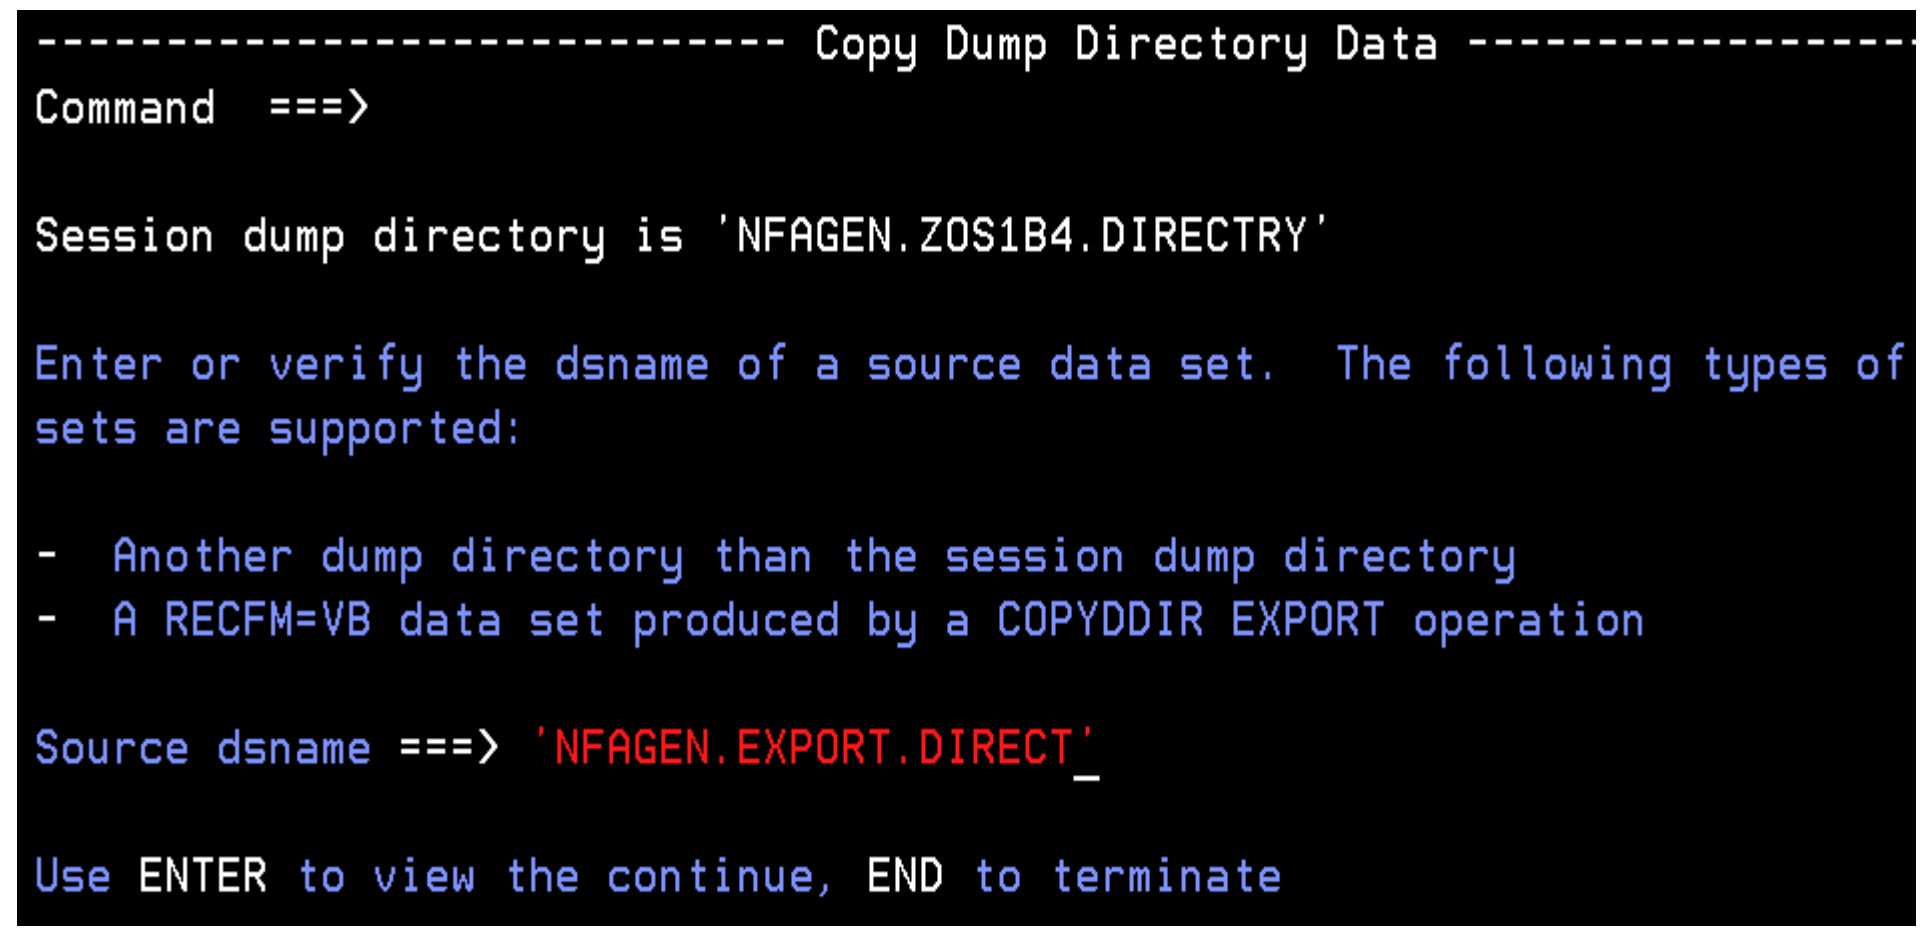

Copy into the current directory from the dataset exported into in previous step. Access from Utilities -> COPYDDIR selections of IPCS.

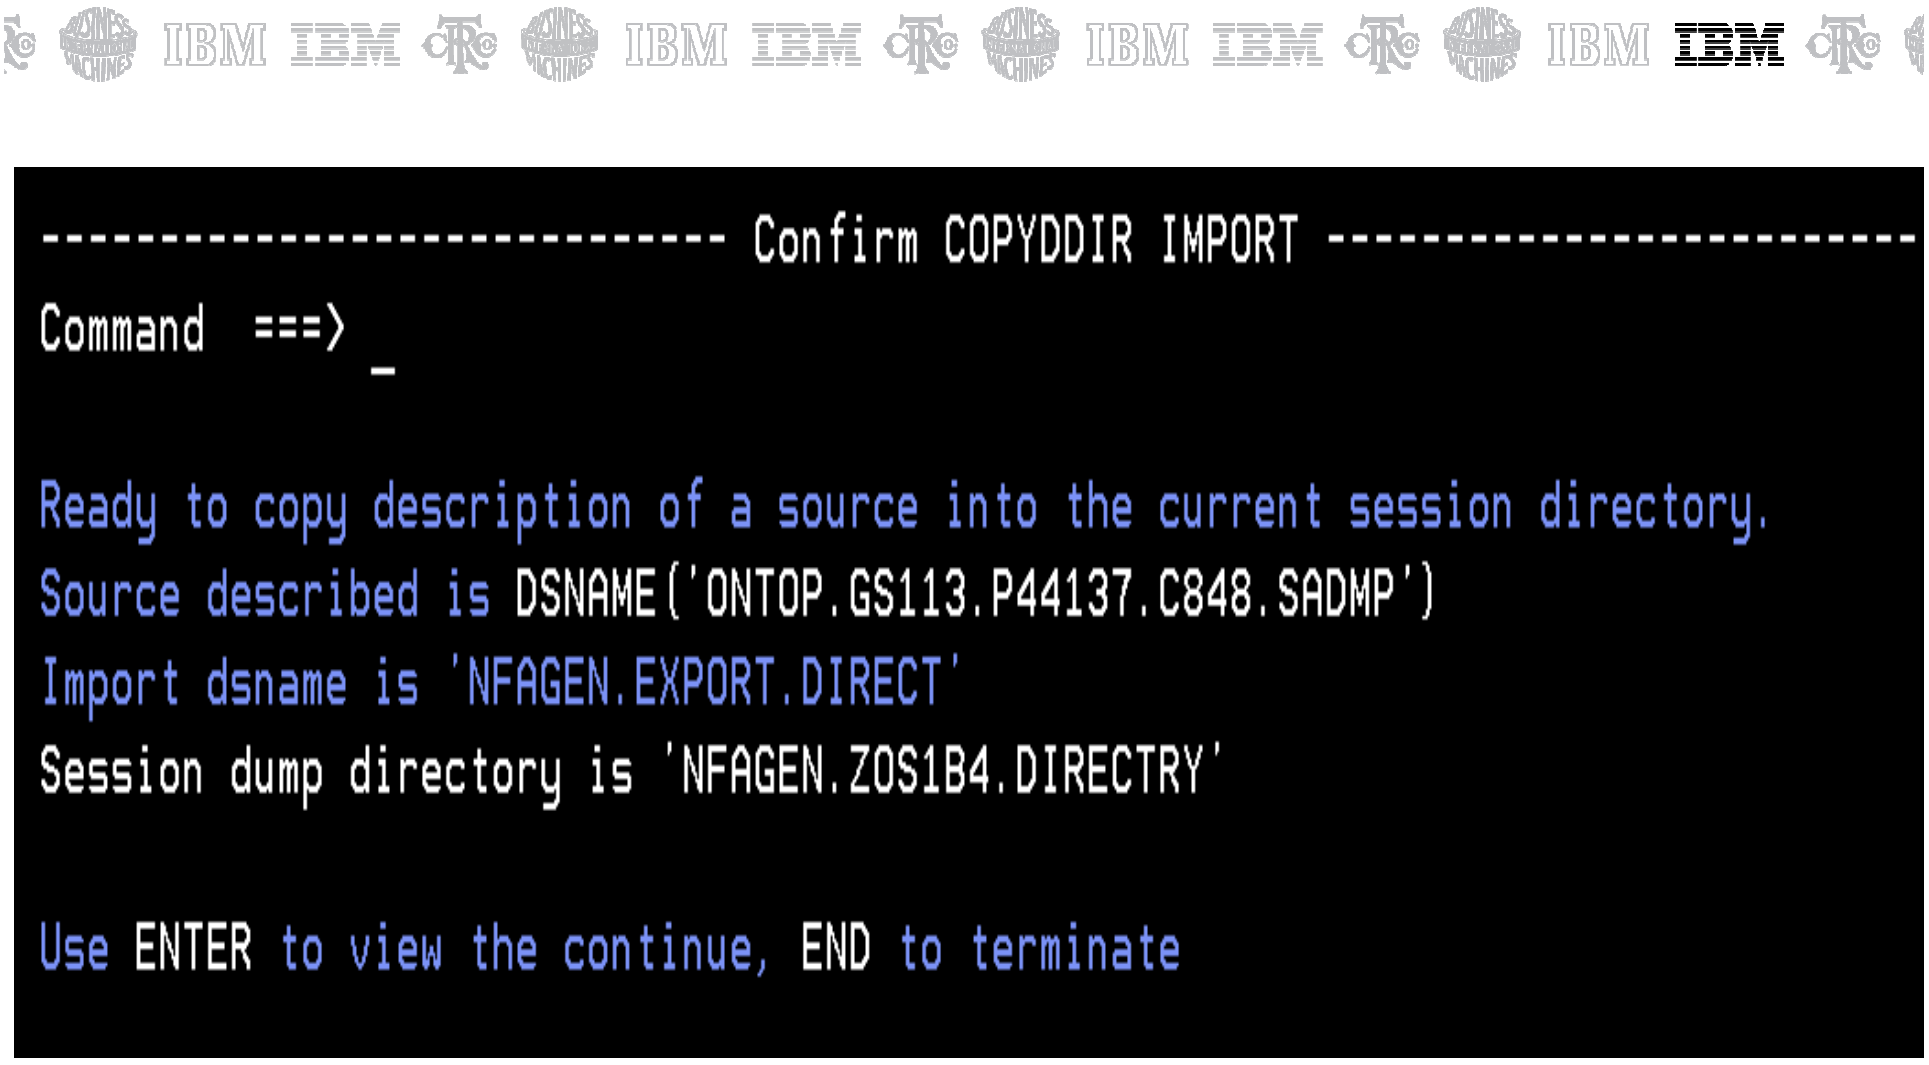

Go ahead and start using the dump

Dump Health Scan

Horror: Large SAD must be FTPd to IBM, delays in the customer's network, IBM's network, etc. results in delays in problem resolution

New Technique: Get started with initial problem diagnosis faster using the dump health scan. Key system level IPCS command outputs routed to a PDS.

z/OS 1.12

Dump Health Scan, How it Works

Physically allocate the PDS

ALLOC the PDS as IPCSPDS

- –ALLOC DDNAME(IPCSPDS) DSNAME('PMRxxxxx.yyy.zzz') SHR

Go into IPCS and issue

- –IP SETDEF PDS
- –IP PROFILE PAGESIZE(2147483647)

Issue IPCS commands one at a time

- –PRINT NOTERM is highly recommended

z/OS 1.12

PDS Allocation

```
                          Allocate New Data Set
Command ===>
                                                                More:
Data Set Name  . . . : NFAGEN.PDS.TEST2

Management class . . . MIGONLY     (Blank for default management class)
Storage class  . . . . STANDARD    (Blank for default storage class)
 Volume serial . . . . SL629D      (Blank for system default volume) **
 Device type . . . . . _______     (Generic unit or device address) **
Data class . . . . . . _______     (Blank for default data class)
 Space units . . . . . CYLINDER    (BLKS, TRKS, CYLS, KB, MB, BYTES
                                    or RECORDS)
 Average record unit   _           (M, K, or U)
 Primary quantity  . . 1           (In above units)
 Secondary quantity    20          (In above units)
 Directory blocks  . . 20          (Zero for sequential data set) *
 Record format . . . . VBA
 Record length . . . . 255
 Block size  . . . . . 32760
 Data set name type    PDS         (LIBRARY, HFS, PDS, LARGE, BASIC, *
                                    EXTREQ, EXTPREF or blank)
```

DUMP Health Scan PDS Contents

IPCS subcommand	PDS member
ANALYZE RESOURCE	ANALYZE
ASMK	ASMK
COMK	COMK
COPYCAPD	COPYCAPD
IOSK ALL VALIDATE	IOSK
STATUS CPU WORKSHEET	STATUS
SYSTRACE ALL TIME(LOCAL)	SYSTRACE
SYSTRACE TTCH(LIST)	SYSTRACE
VERBX MTRACE	VERBX
VERBX SADMPMSG	VERBX

GRS

Horror: Large dump, lots of resources, takes forever and a day to do finds in GRS report outputs

New Techniques

- Code restructured to yield performance improvements
- New filtering options to reduce data presented to user
- Panel driven interface
- More attributes relayed about resources of interest
  - Event TOD and ENQ history (request, contention, ownership)
  - Altered by RNLs, Exits, 3rd party managed
  - Directed ENQ details, requesting TCB / Target TCB

z/OS 1.10

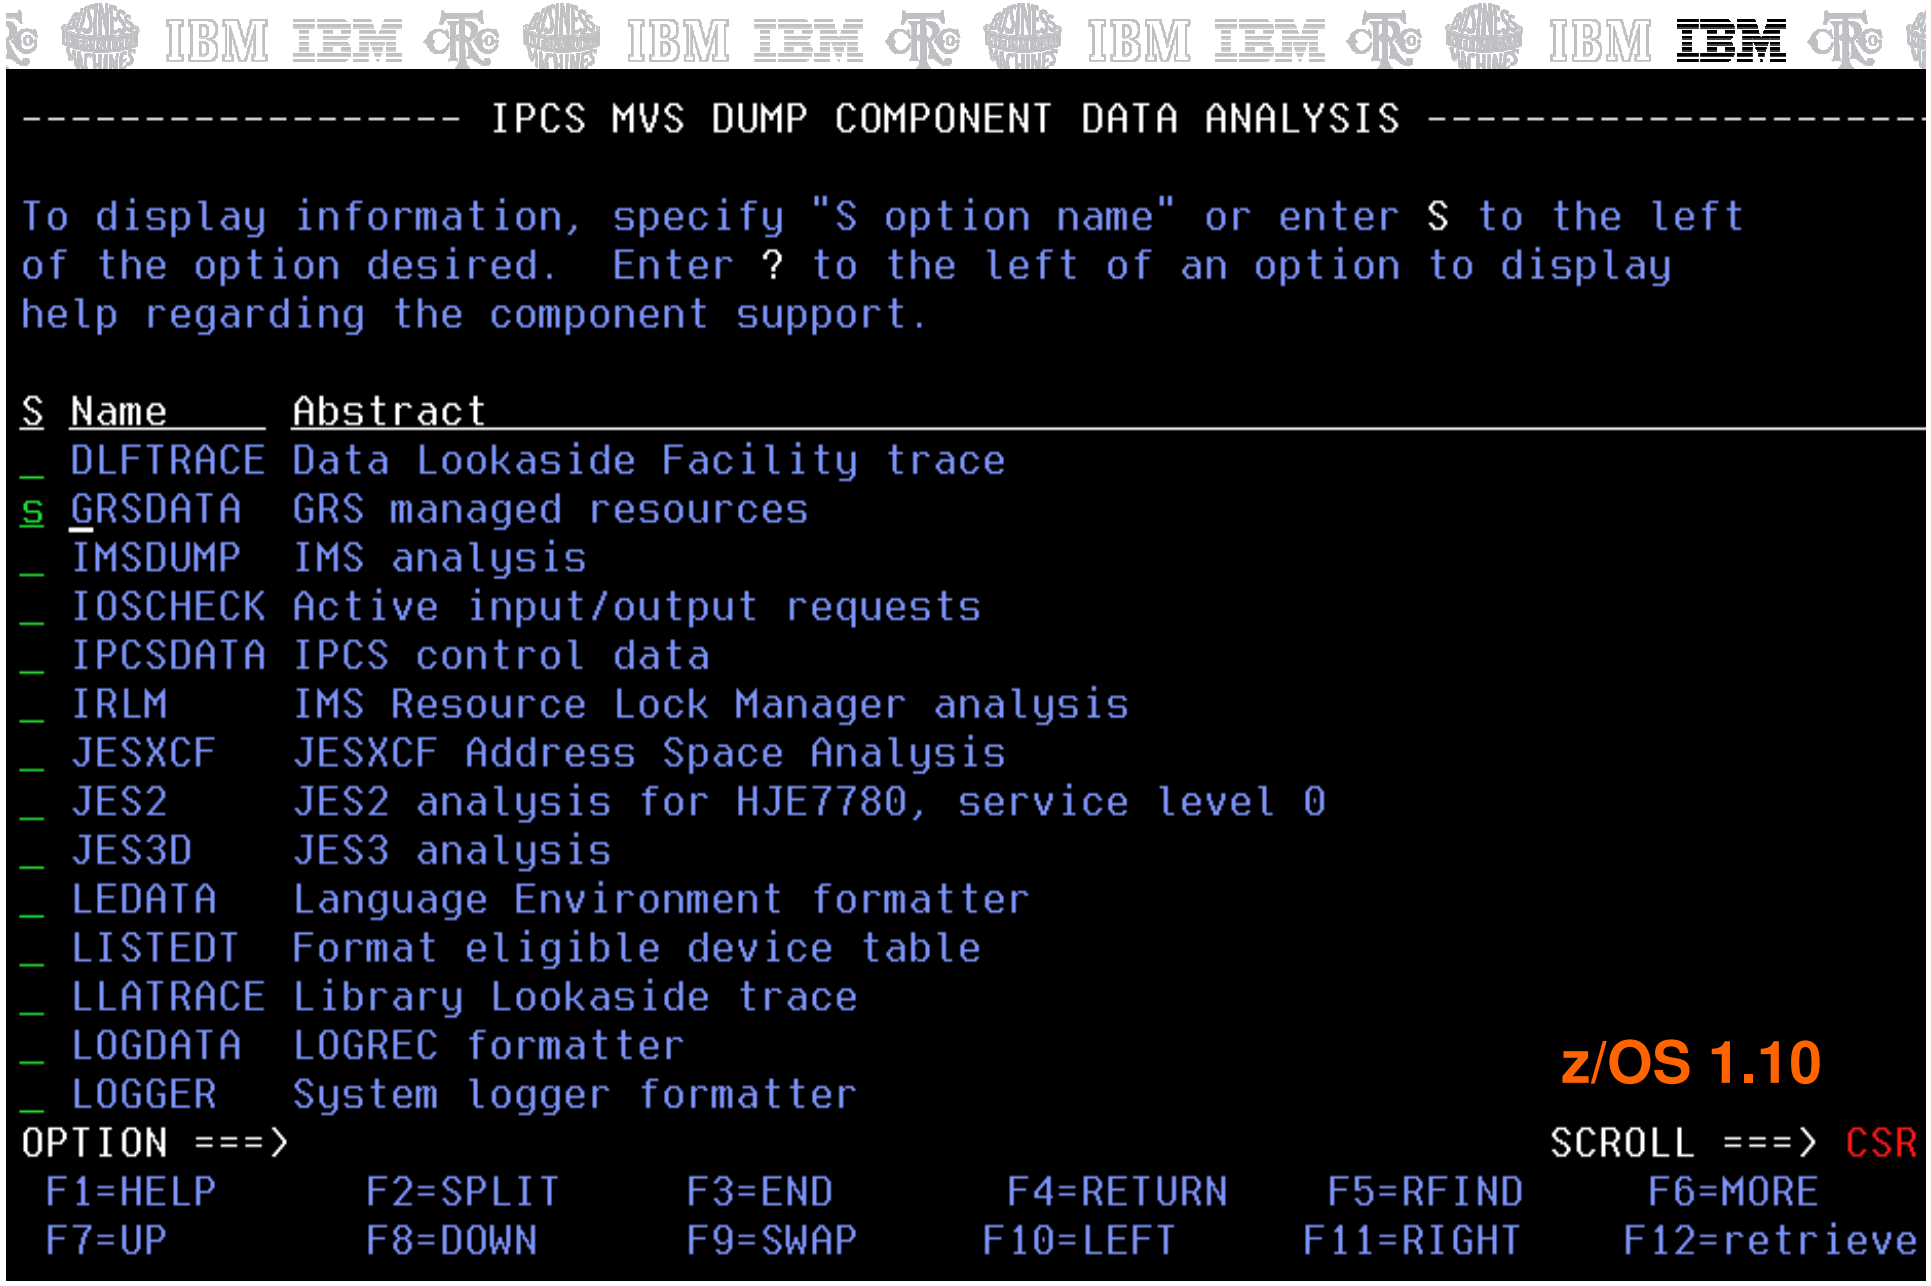

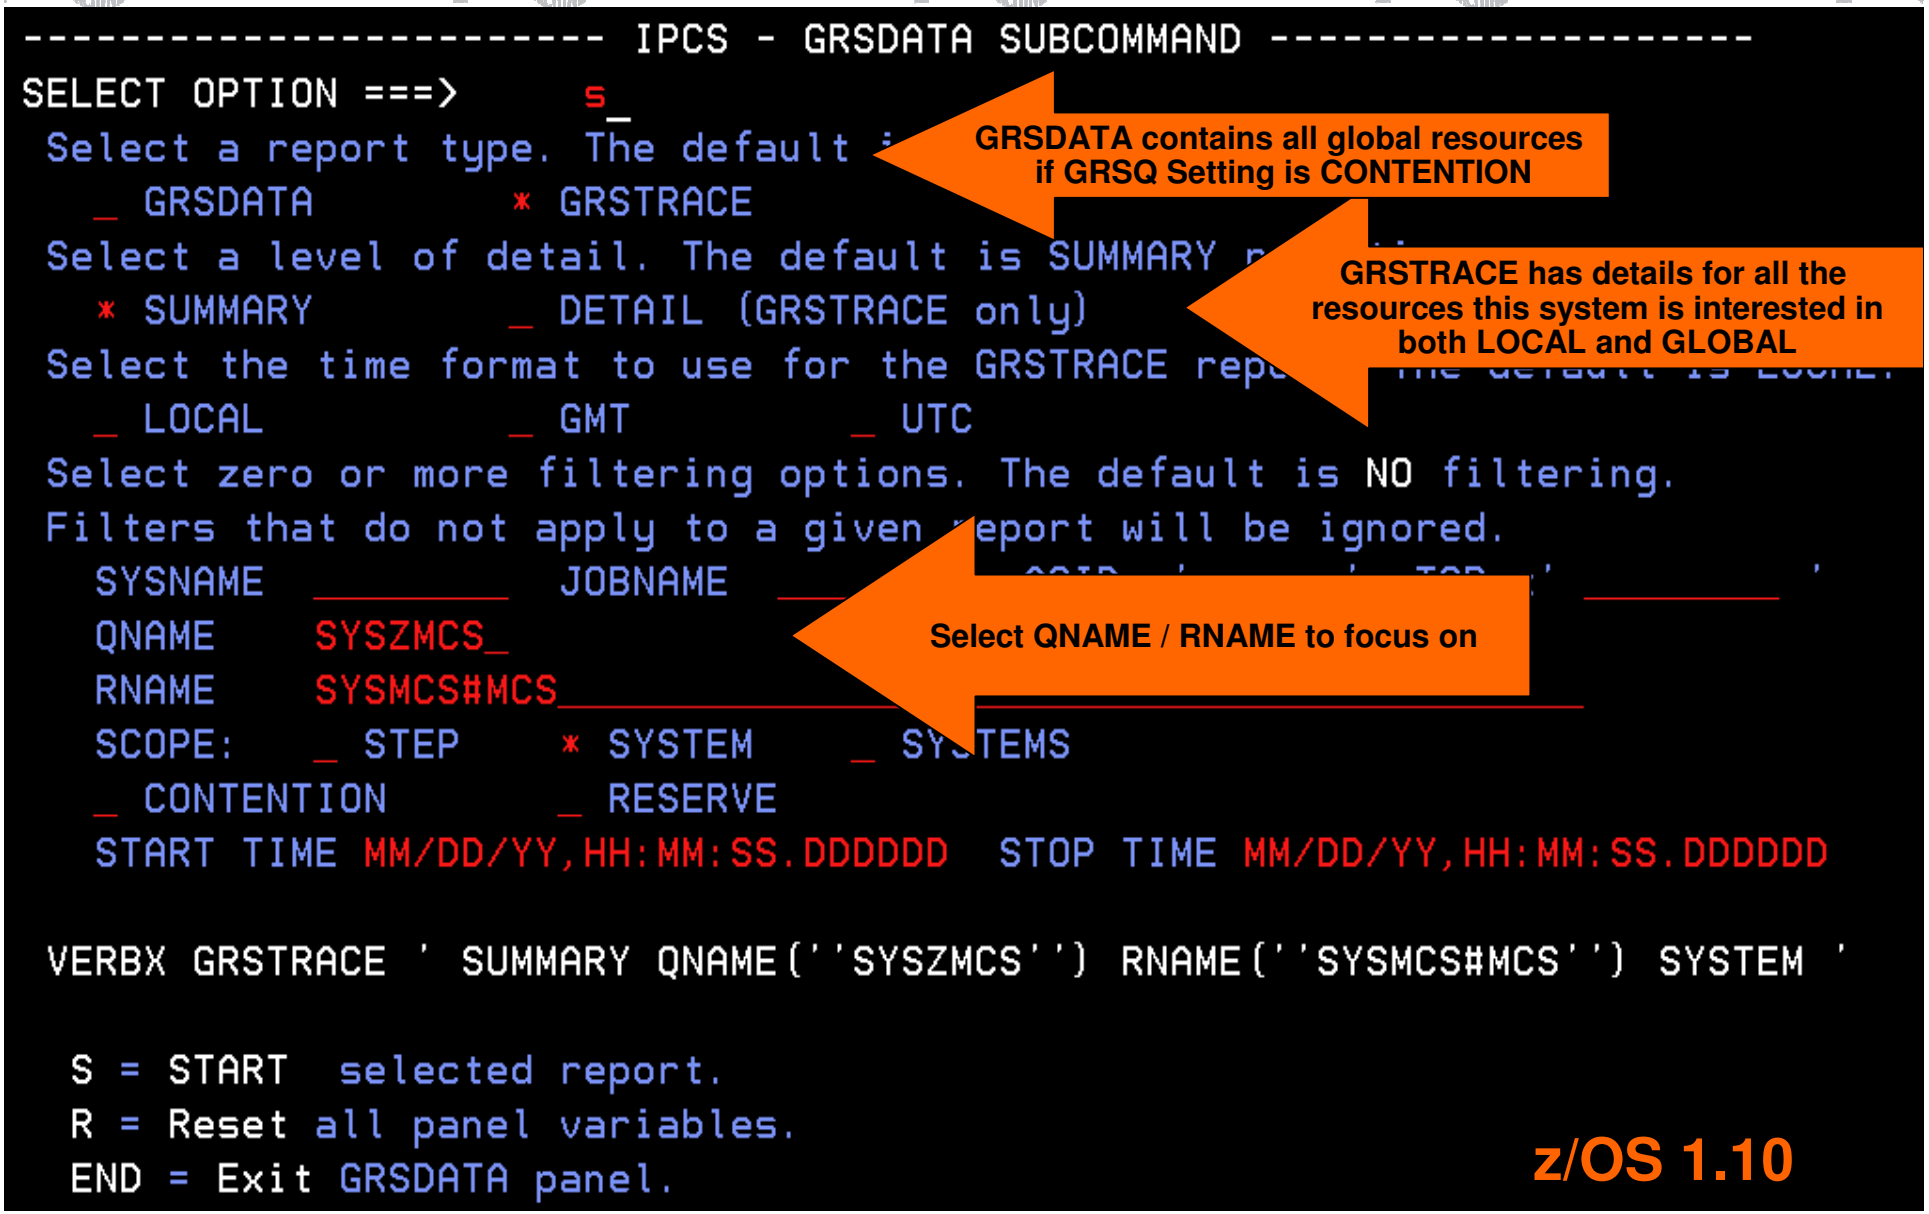

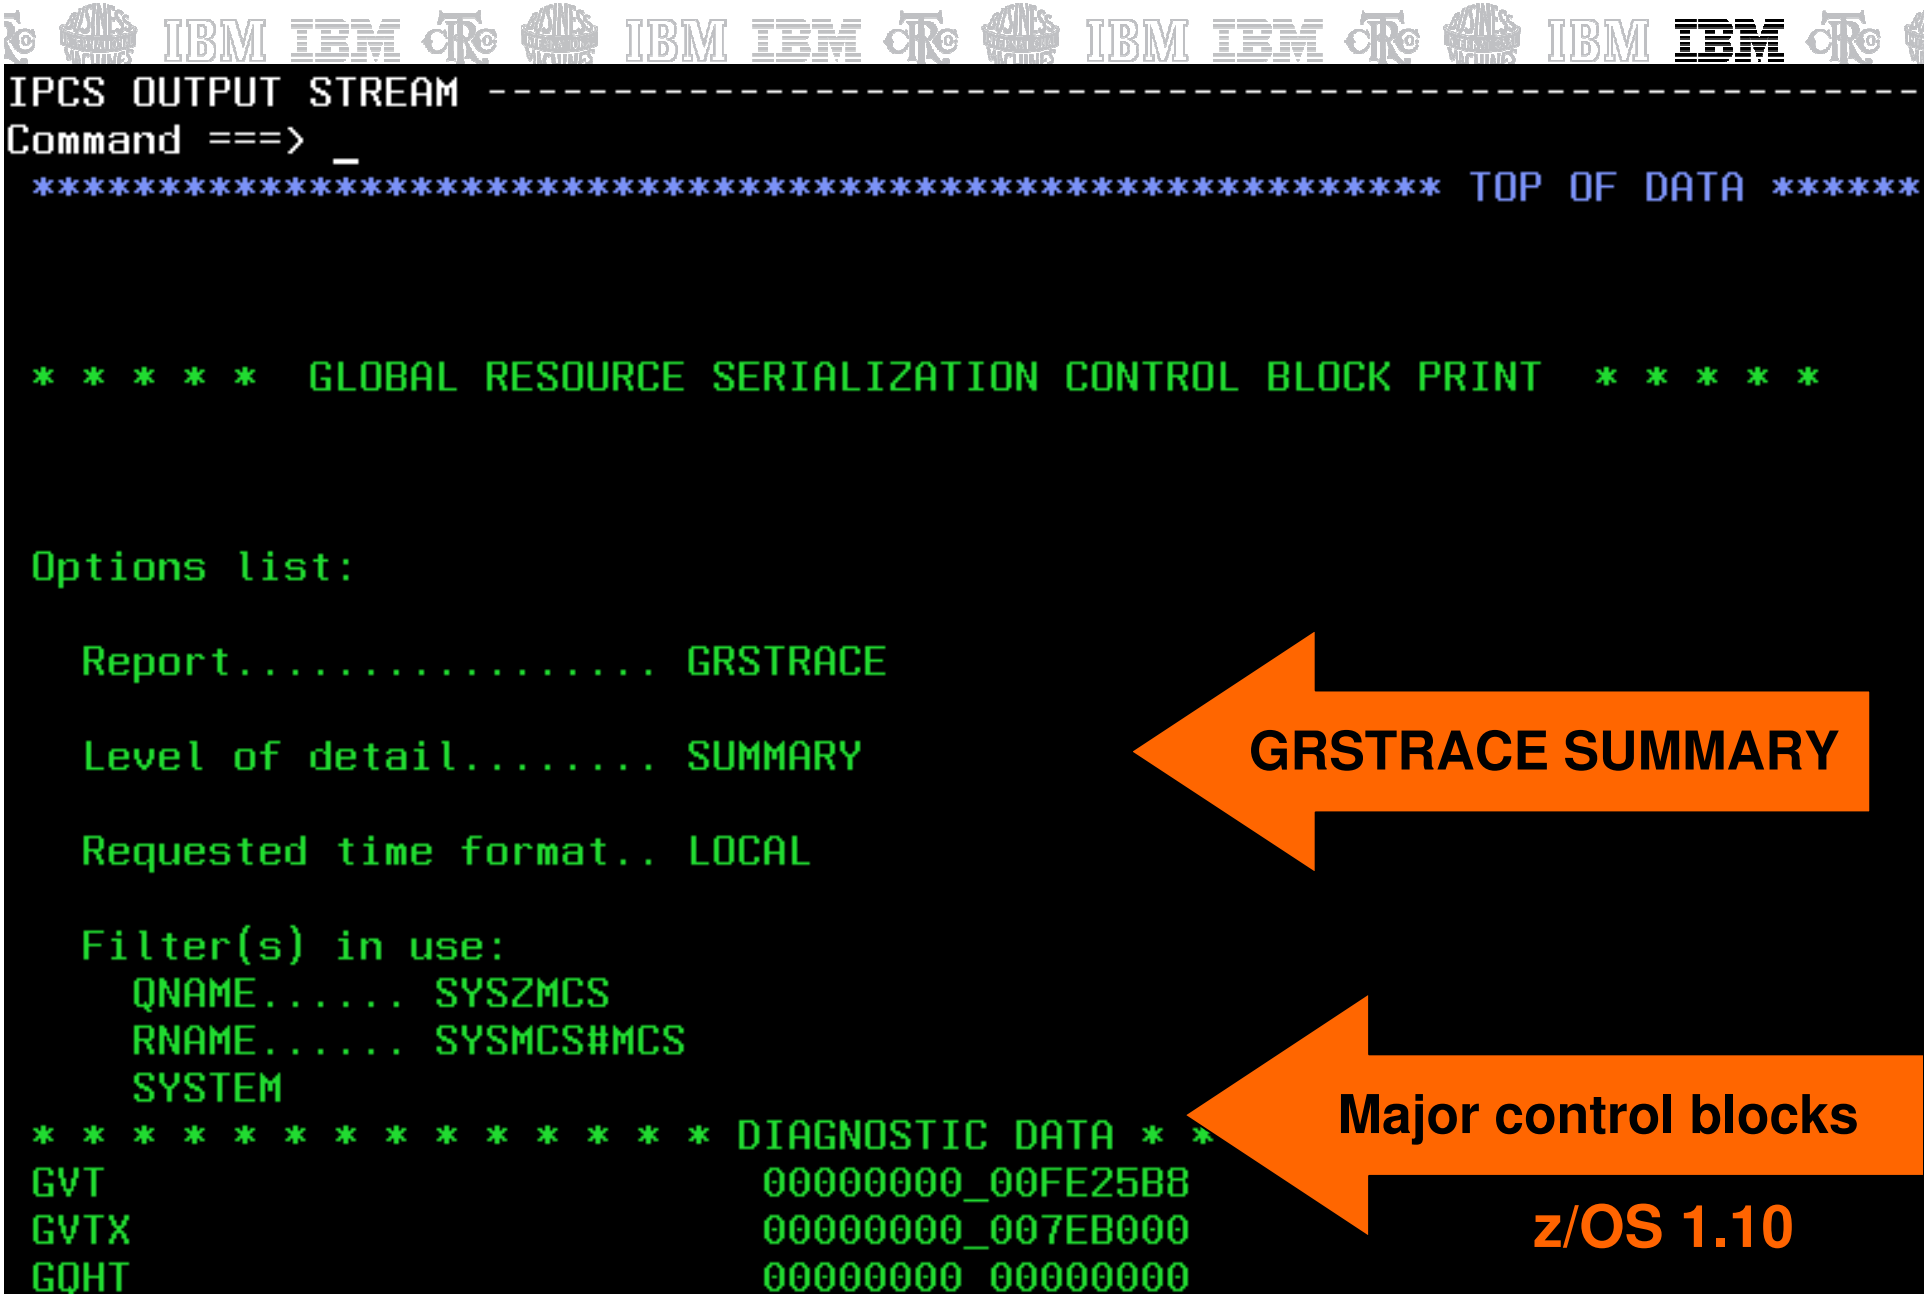

```
IPCS OUTPUT STREAM ----------------------------------------------------------------
Command ===> _
 SGHT                          00000028_F8F00000
 LQHT                          00000028_F8E00000
 STHT                          00000000_7F58A000
 RPT                           00000000_7FFB7EE0

 * * * * * * * * * * * * * * * CONFIGURATION INFORMATION * * * * * * * * *
 GRS Mode                                  STAR
 Synchres Setting                           Yes
 Active Exits
      ISGNQXITFAST
 GRSQ Setting                         Contention
 CNS                                       R22
 CTRACE Buffer size (K)                    128
 The EQDQ monitor is                       Off
 ENQMAXA                                250000
 ENQMAXU                                 16384


                              ********************************************************
                              ********************************************************
                              *****                                            *****
                              *****   STEP QUEUE (STHT) CONTROL BLOCK PRINT    *****
                              *****                                            *****
                              ********************************************************
```

Mode, CNS system, ENQMAXA and ENQMAXU

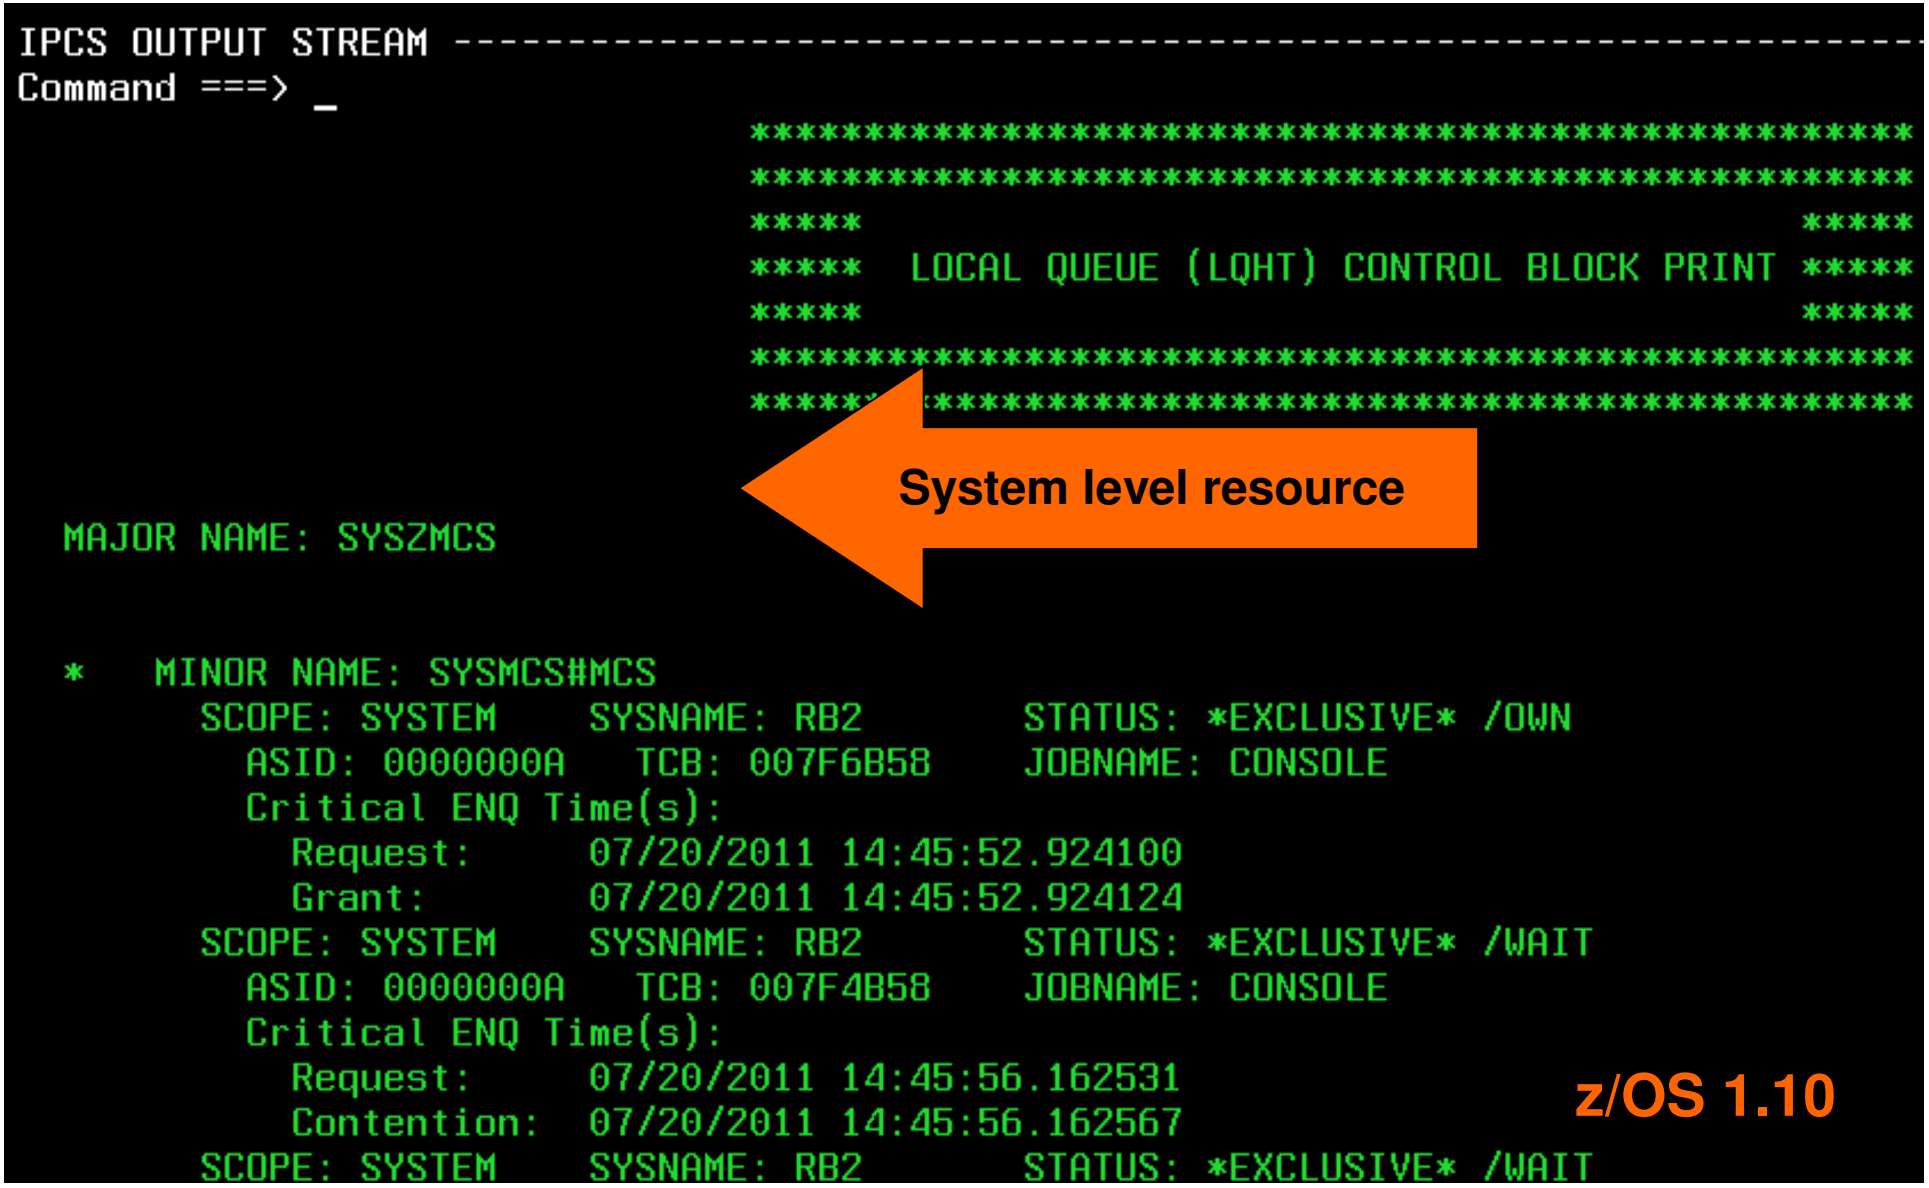

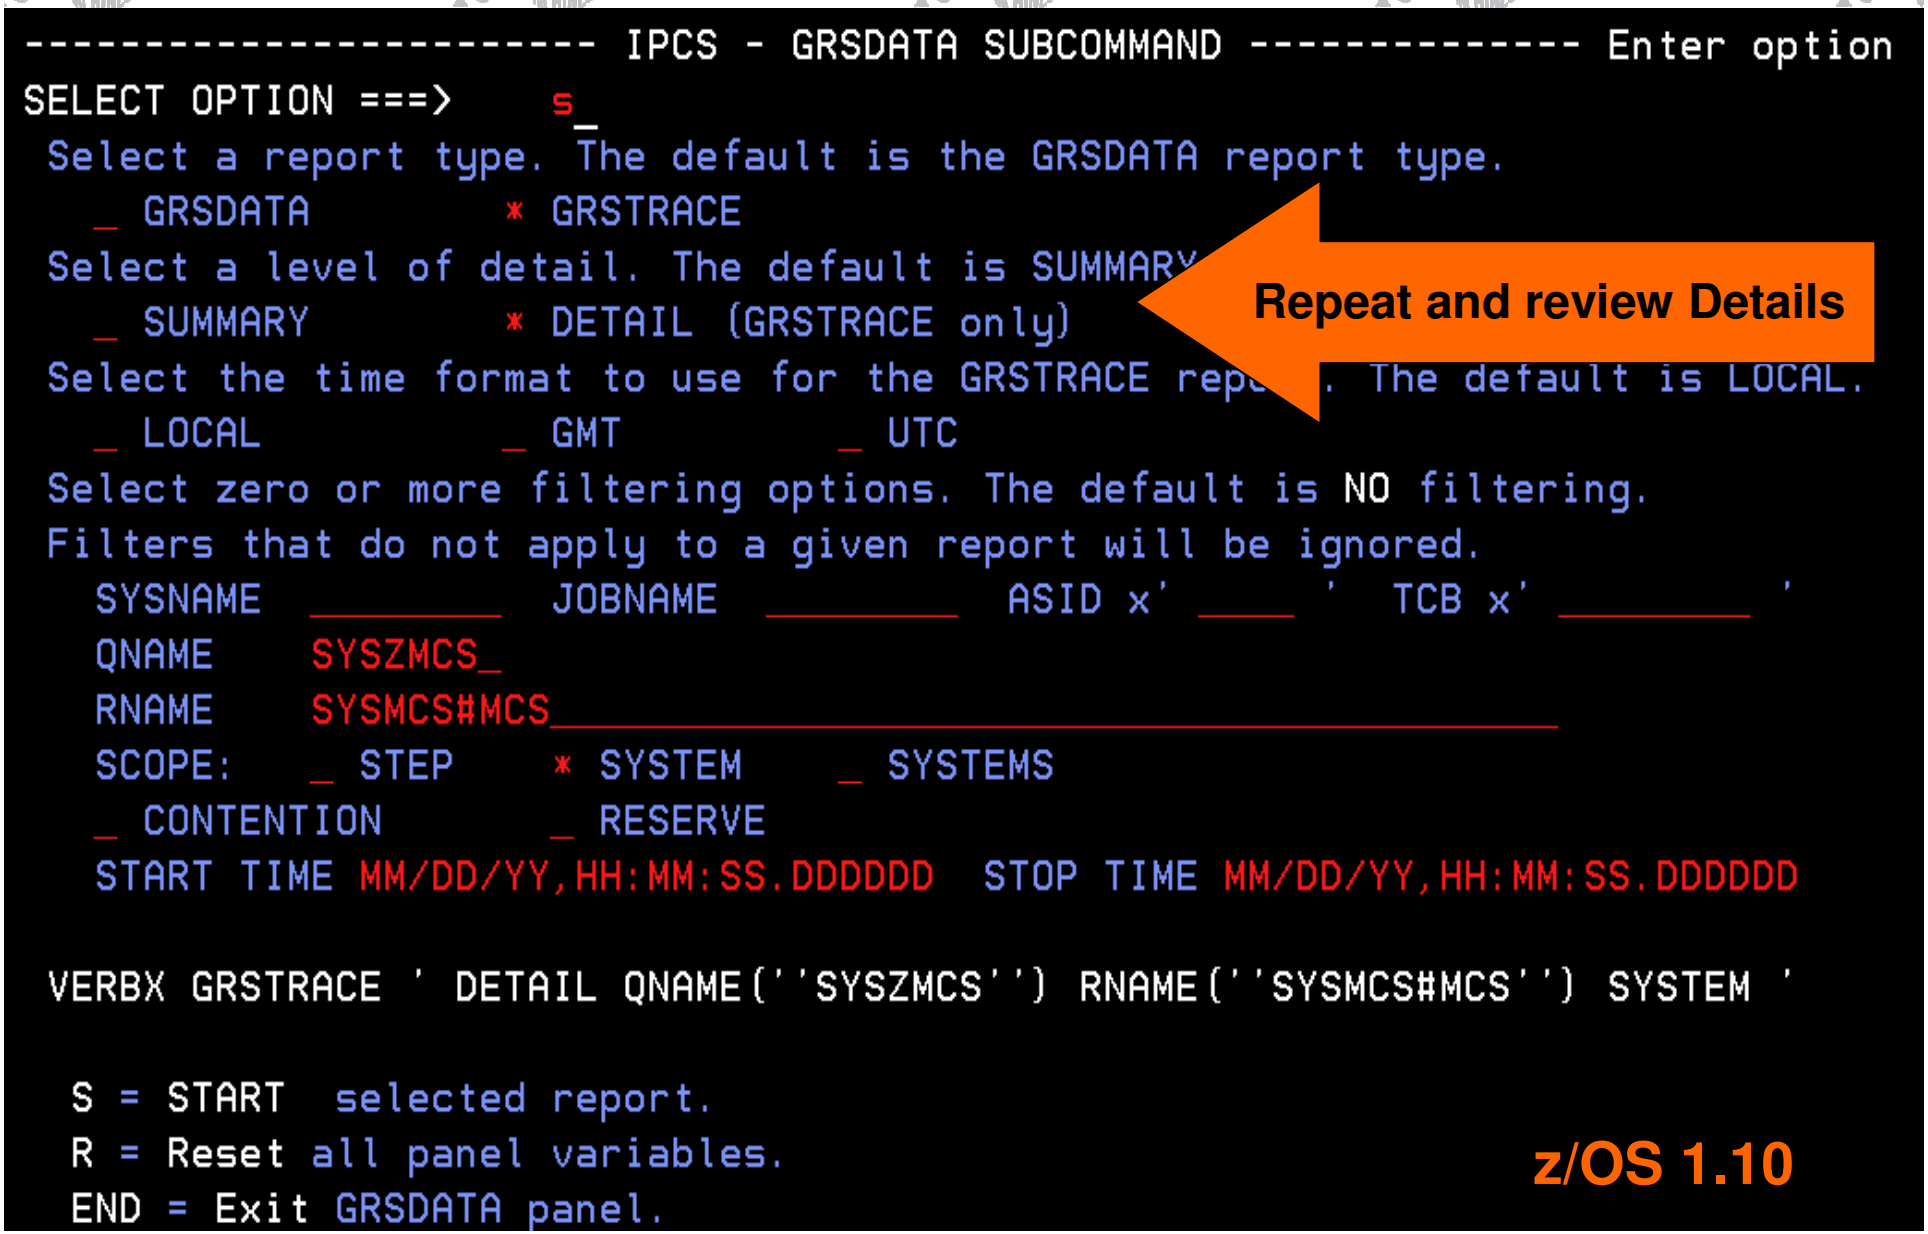

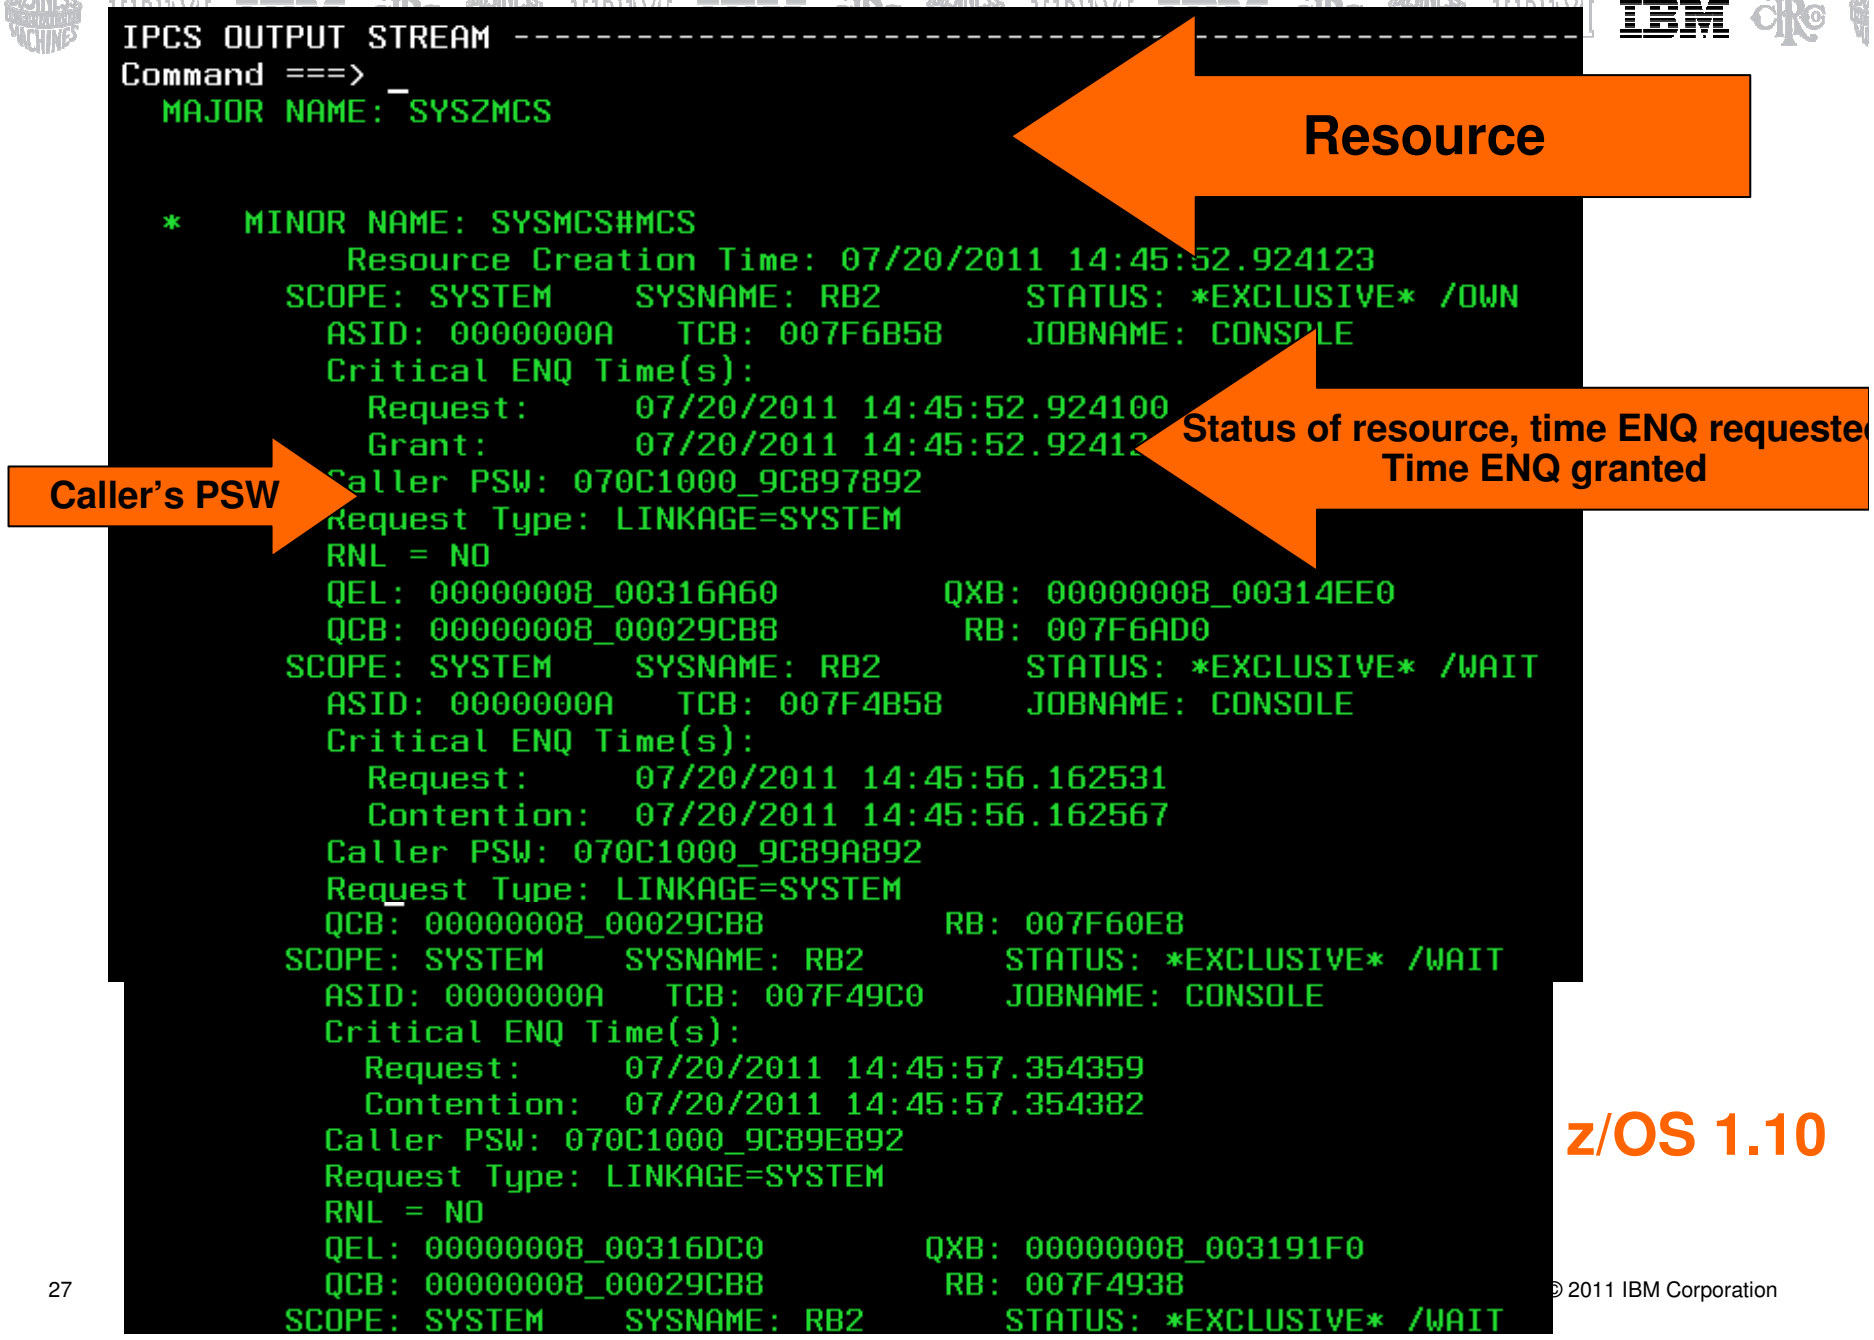

z/OS 1.10

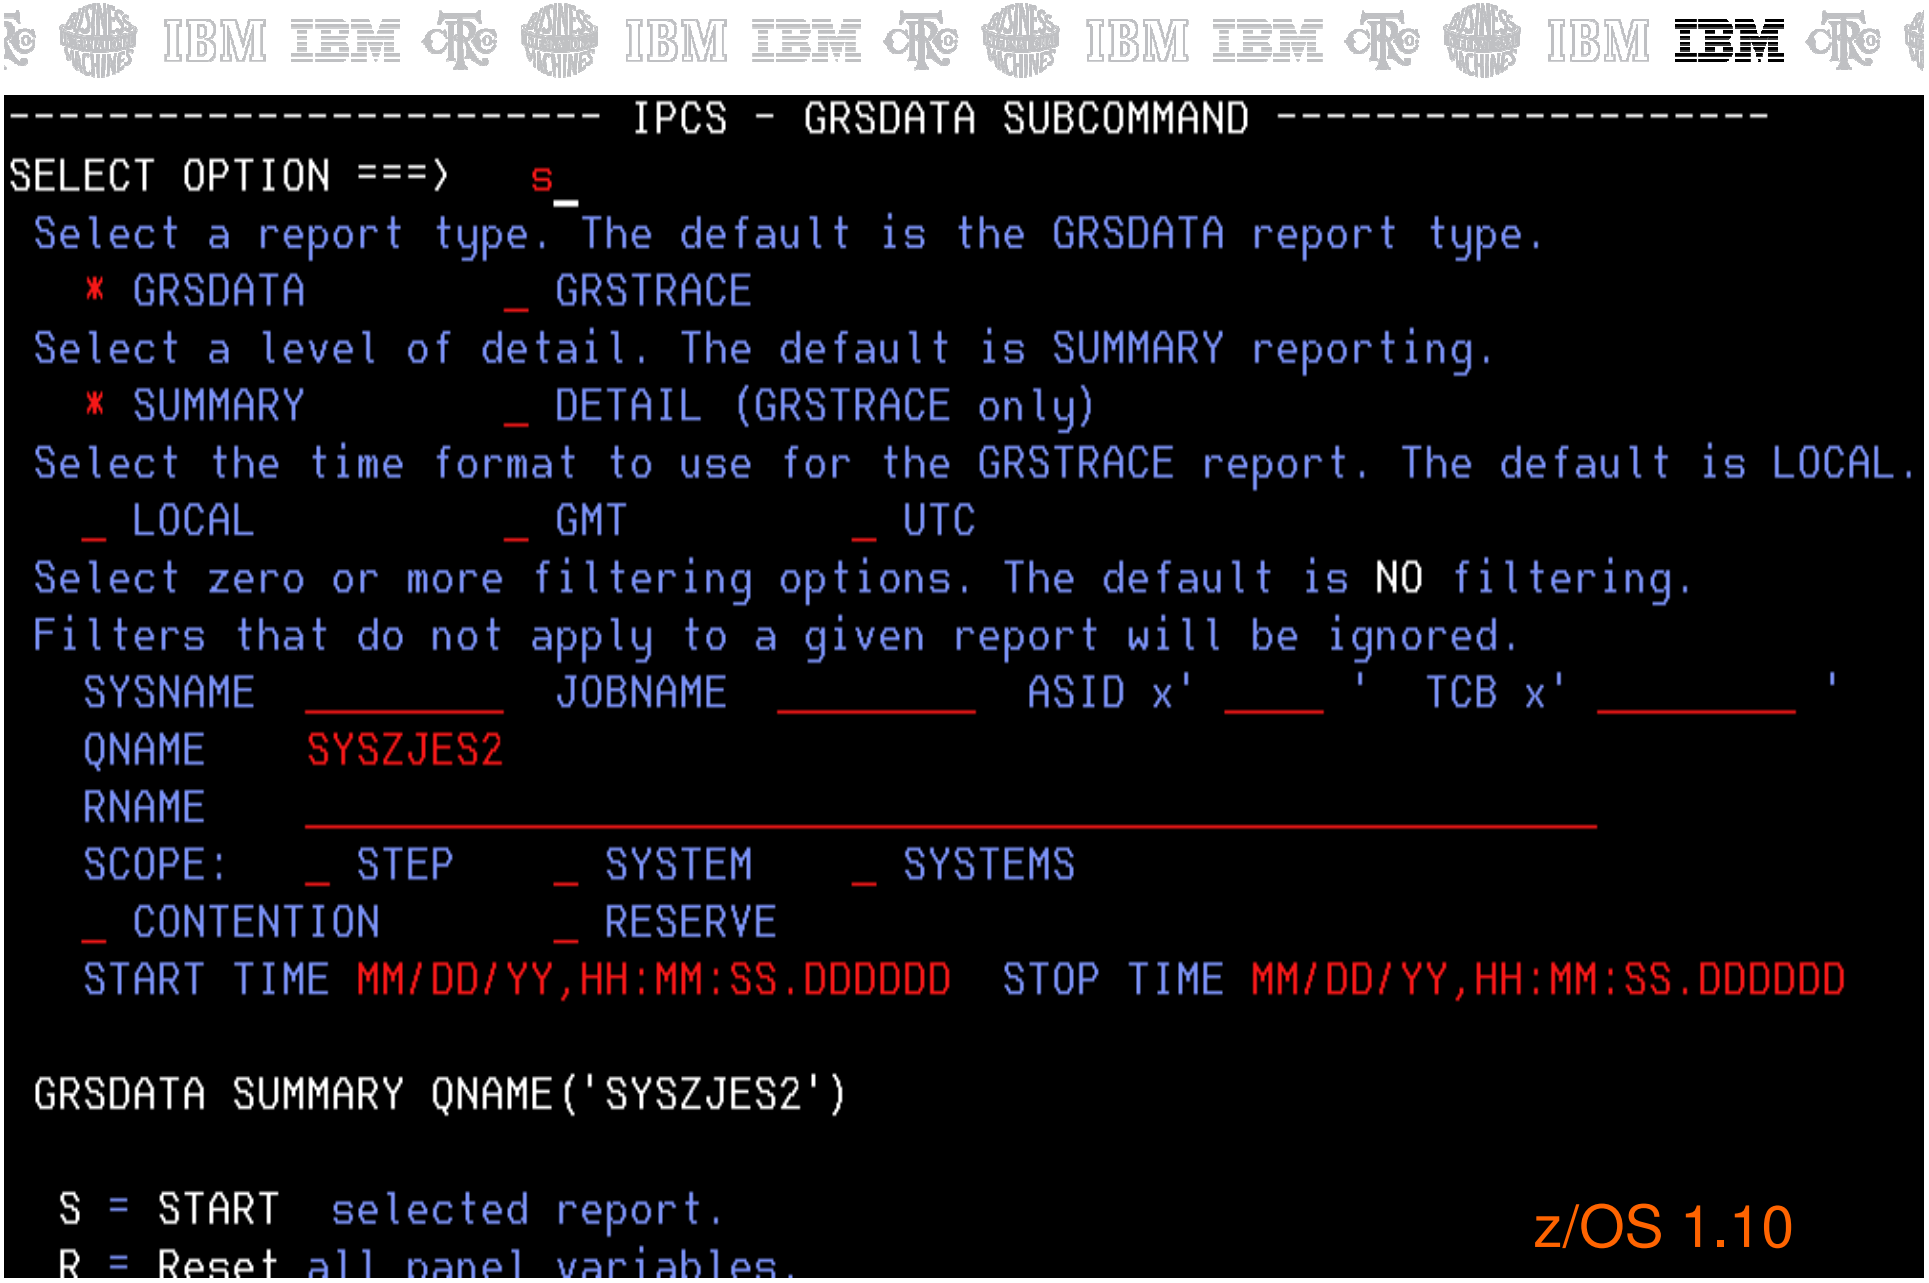

z/OS 1.10

GRSDATA Output for SYSZJES2

```
Global systems resources

  Major.. CL8'SYSZJES2'

    Minor.. CL050'PP1JS1SYS1.JESCKPT1'

      SCOPE. SYSTEMS  SYSNAME. PP1B      JOBNAME. JES2      RESERVE
      ASID.. 002C     TCB..... 008FF370  STATUS.. EXCLUSIVE

      SCOPE. SYSTEMS  SYSNAME. PP1D      JOBNAME. JES2      RESERVE
      ASID.. 0022     TCB..... 008FF370  STATUS.. WAITEXC   ECB..... 2A5BA170

      SCOPE. SYSTEMS  SYSNAME. PP1C      JOBNAME. JES2      RESERVE
      ASID.. 002C     TCB..... 008FF370  STATUS.. WAITEXC   ECB..... 2A5BA170

                 Latch Statistics
```

Not a lot of detail

SUMMARY Subcommand

Horror: Output of IP SUMM FO can be very very very long

New Technique: New options to enable IPCS users to focus on a sole TCB or EXCLUDE(GLOBAL,JPQ,LOADLIST). Panel driven or command driven.

- EXCLUDE(GLOBAL) causes global SRB formatting to be omitted.
- EXCLUDE(JPQ) causes job pack queue formatting to be omitted.
- EXCLUDE(LOADLIST) cause load list formatting to be omitted.

IP SUMM FO ASID(6) TCB(xxxxx)

IP SUMM FO ASID(6) EXCLUDE(GLOBAL,JPQ,LOADLIST)

z/OS 1.10

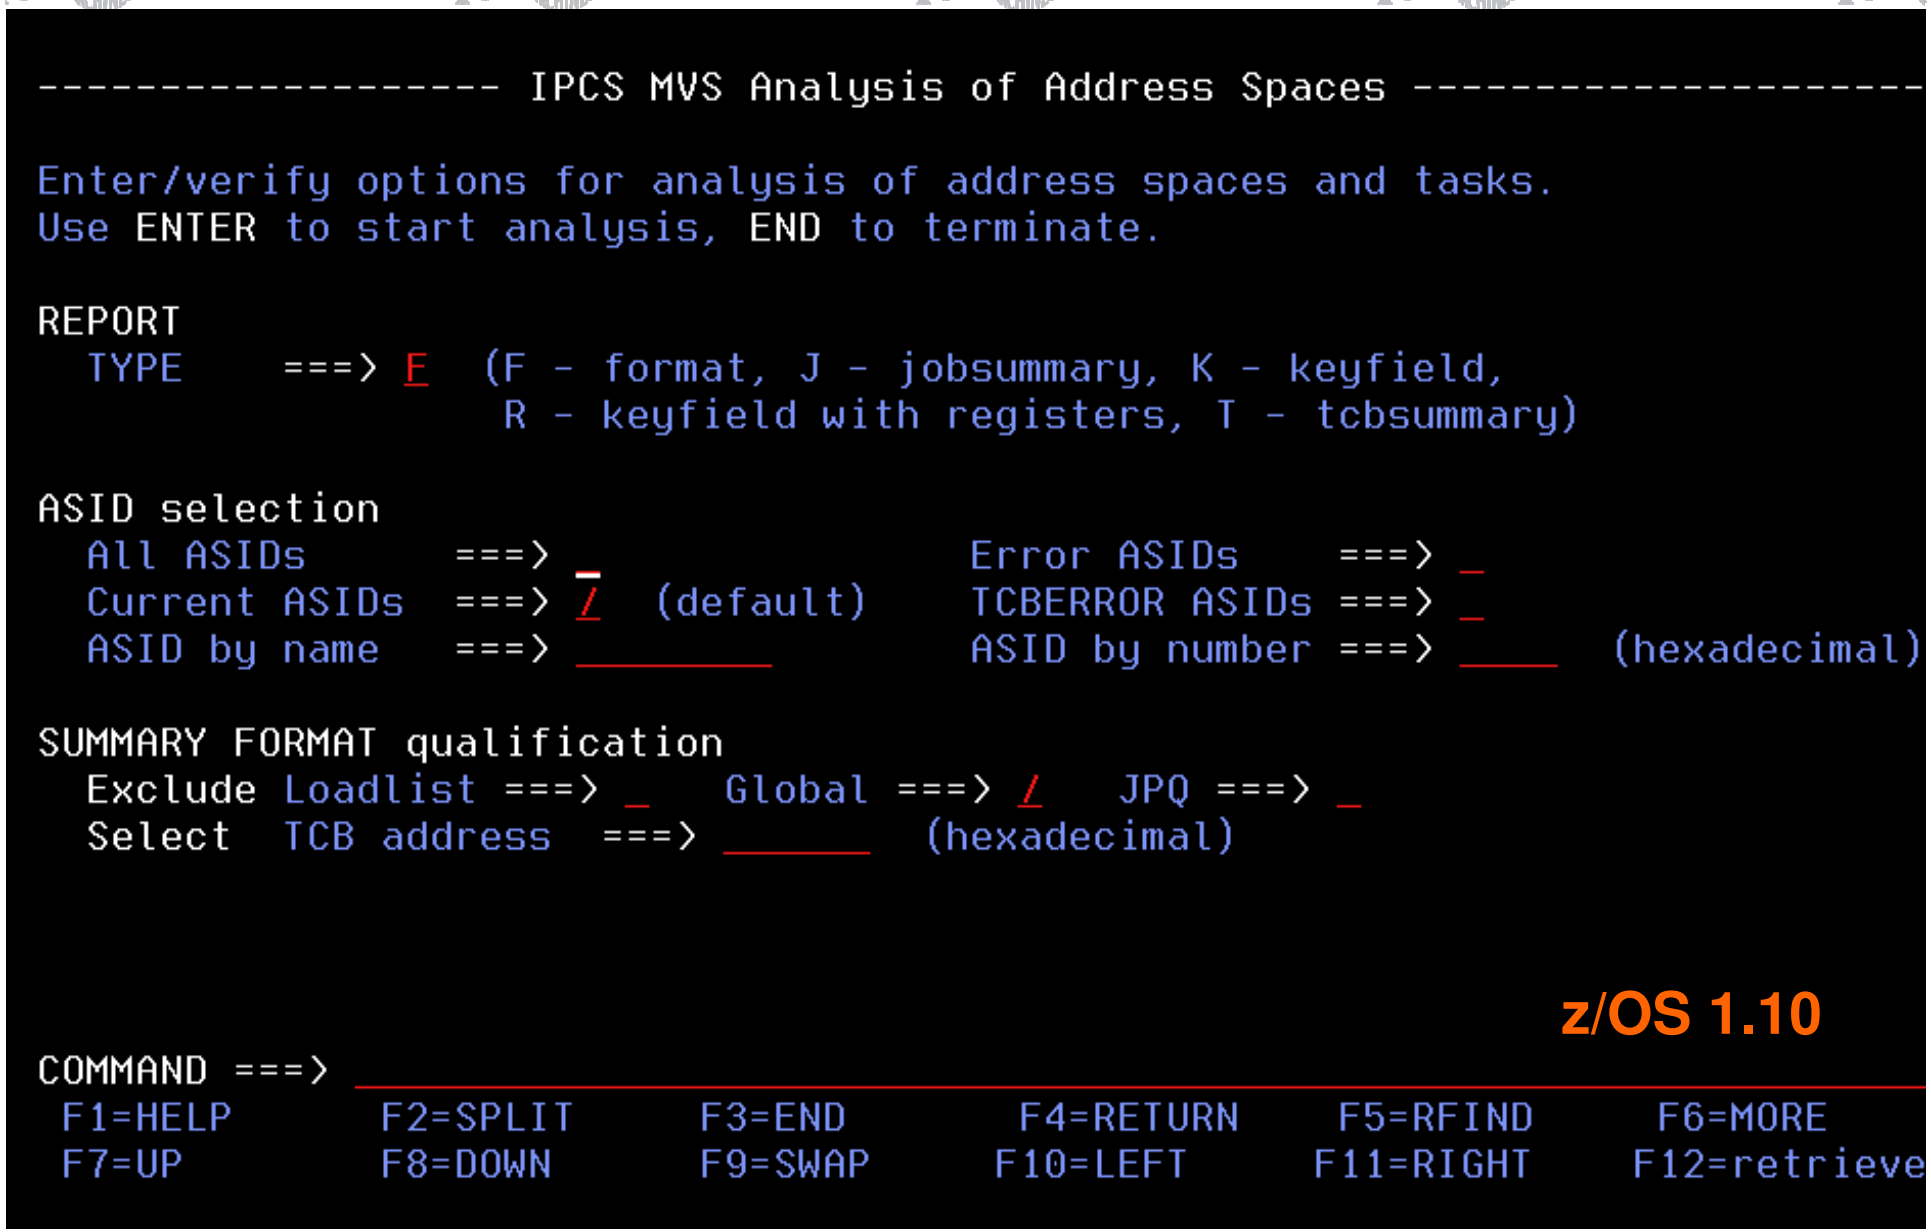

IP SUMM FO ASID(6) TCBADDR(x'7d1178')

SRBs

Work Unit Queues

ASCB

ASSB

ASXB

PC tables

Formatted TCB x’7D1178’

*Does not have all of the “other” TCBs

z/OS 1.10

IP SUMM FO ASID(6) EXCLUDE(GLOBAL)

Work Unit Queues

ASCB

ASSB

ASXB

PC tables

All TCBs in ASID 6 formatted

SRBs not formatted

z/OS 1.10

System Trace

Horror: History is not long enough to see the problem in the system trace table. Faster processors make this even worse.

New Technique: System trace buffers can be made larger and new “M” for megabyte notation accepted.

TR ST,{nnnK,nnnM}

Horror Remains: Buffers are below the line and that’s a lot of storage!

z/OS 1.9

System Trace

Horror: We have the ability to specify HUGE system trace buffers and they reside below the bar!

New Technique: System trace buffers moved above the bar,

- –Ability to specify in “G” gigabytes
- –New default per processor is 1M
- –1M is the minimum per processor.

z/OS 1.10

System Trace

Horror: Turning on branch tracing also turns on mode tracing

New Technique: Branch tracing and mode tracing are controlled individually

TRACE ST,BR={ON|OFF},MODE={ON|OFF}

z/OS 1.9

System Trace

TRACE ST,{nnnK|nnnM|nG}

- –Each CPU will be allotted the amount of storage declared
- –If “K” is specified 1M is allotted

TRACE ST,BUFSIZE={nnnnnK|nnnnnM|nnnG}

- –All CPUs will share the storage allotment declared
- –Much wiser and easier for complete system planning

Caution: Always take into consideration the amount of real storage on the system. Larger trace buffers means less pages available for other system work. More storage for trace buffers may cause a system level performance problem.

z/OS 1.10

5 Words Captured on SPER

Horror: there is some key data that needs to be captured immediately when SPER is trigger

New Technique: Ability to capture 5 words of data in SPER system trace entry when a slip hits.

z/OS 1.9

5 Words Captured on SPER

Data to be captured

- displaces Unique Data Fields
  - Instruction length code
  - translation error address
  - PSACLHS/E
  - trap ID
- Specify by direct or indirect addressing
- Declare a start and end range
- Processing will always rounded up to word boundary
- Data to be capture truncated at 5 words

SLIP SET,IF,…,A=STRACE,STDATA=((2R?,+8),(10000,+4)),END

z/OS 1.9

SYSTRACE SORTCPU

Horror: System trace sorts entries in time order. Some problems occur on a single CPU and analysis takes a long time when the traces are fully merged.

New Technique: Enable users to sort system trace entries by CPU

IP SYSTRACE SORTCPU(mm/dd/yy,hh:mm:ss:dddddd,n)

- "n" is the number of entries before and after the specified time to show in the output

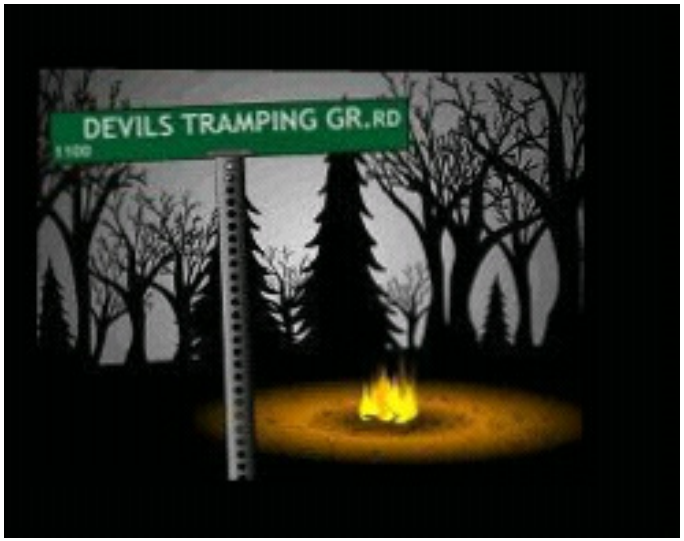

z/OS 1.12

IP SYSTRACE SORTCPU(07/05/11,13:29:59.200299,5)

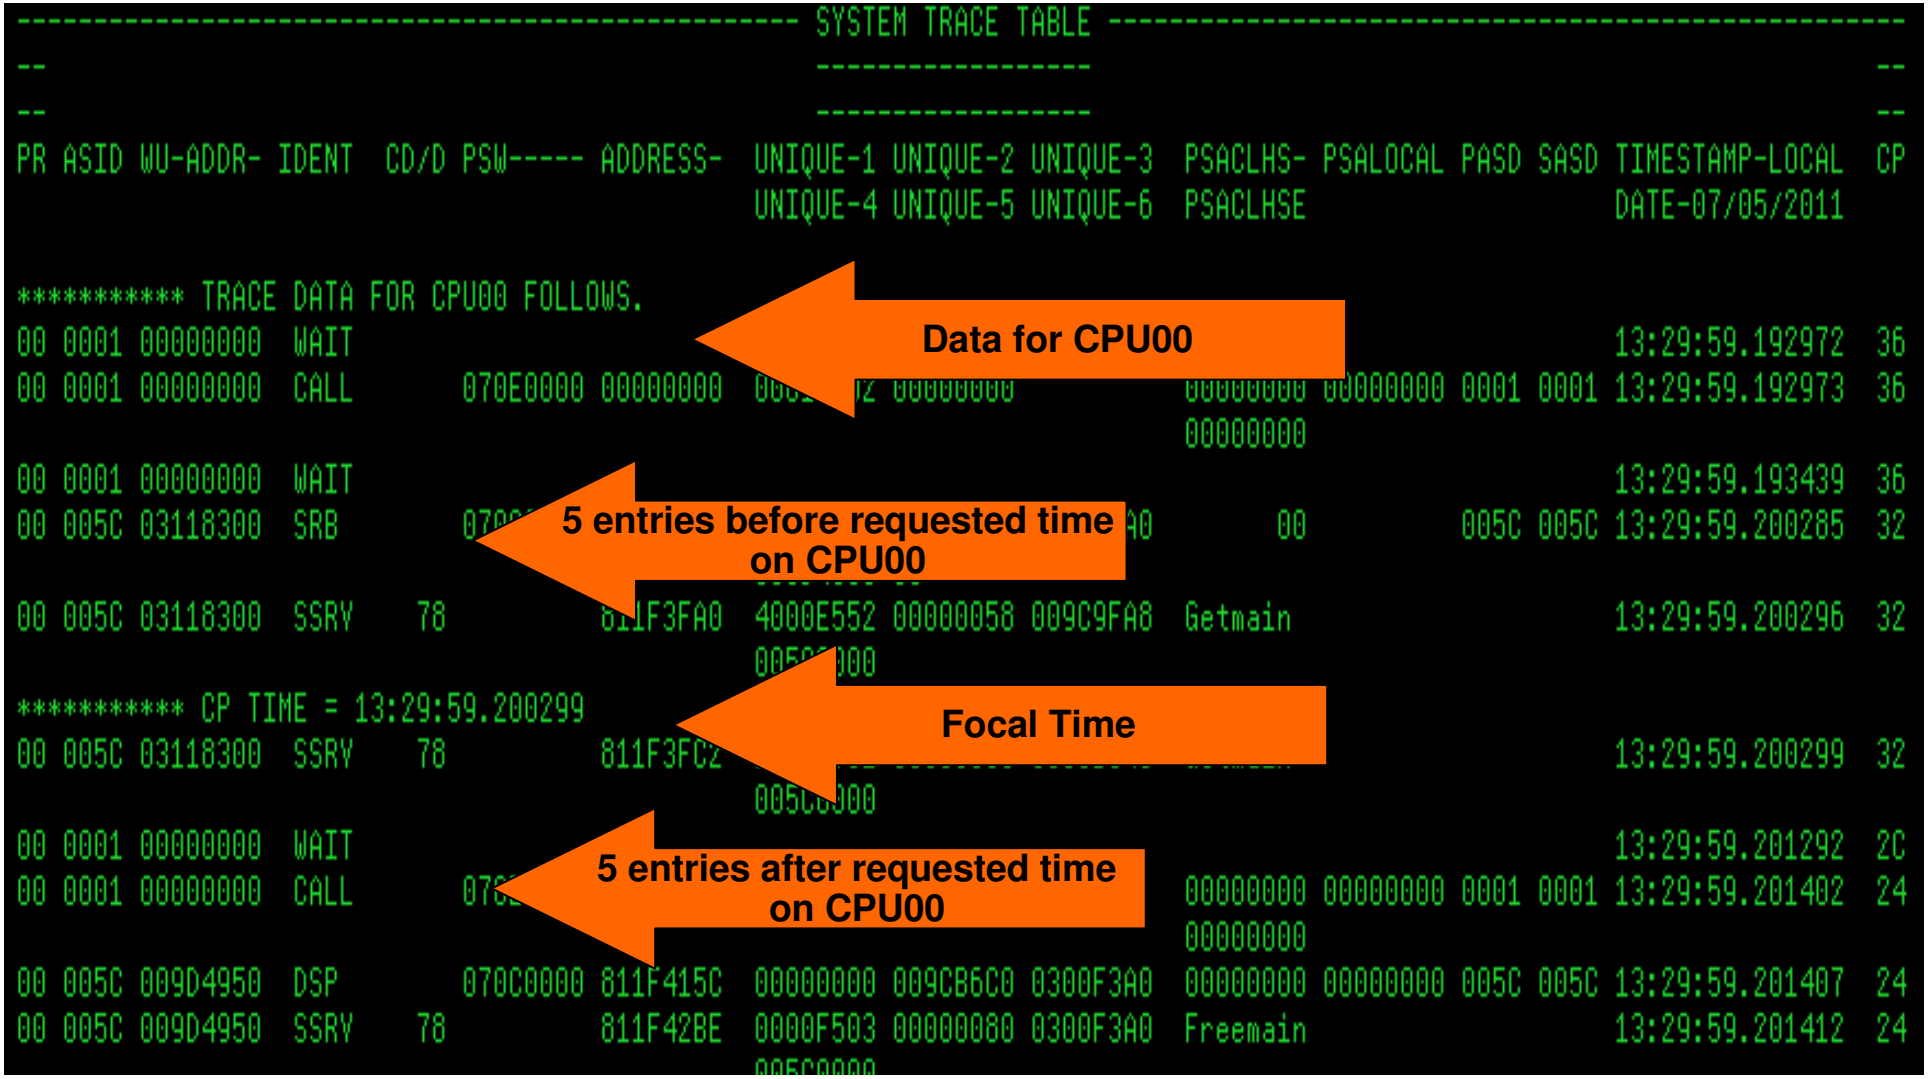

Data for each CPU follows in the same format

System Trace Performance Analysis

Horror: Nightmare to do performance analysis by hand in a dump.

New Technique: Introduction of SYSTRACE PERFDATA

PERFDATA ([SHOWTRC] [DOWHERE] [SIGCPU(sss.dddddd)])

- SHOWTRC requests the system trace table be displayed in the output
- DOWHERE requests WHERE on the PSWs within CLKC and SRB analysis sections of PERFDATA
- SIGCPU requests CLKC analysis and WHERE analysis for SRB events be bypassed for usage less than the specified time

Always always always keep in mind this is a very tiny focused picture. It is relevant for the time covered but may not reflect they overall system view.

z/OS 1.12

Summary of Contents of PERFDATA Report

CPU usage summary

CPU breakdown by ASID

SRB breakdown by ASID

- With WHERE info

TCB breakdown by ASID

CLKC events

- With WHERE info

Lock information

SSCH to I/O times

z/OS 1.12

IP SYSTRACE PERFDATA

```
No address spaces with the CURRENT attribute were found

Note: Only SYSTRACE records available for ALL PROCESSORS are considered.

System: MCEYS1 SP7.1.2 HBB7770


PERFDATA Analysis:

 CPU#    Went from          To          Seconds     SRB Time     TCB Time    Idle Time   CPU Overhead
 ---- -------------- -------------- ----------- ----------- ----------- ----------- -----------
  01  17:29:58.257288 17:30:02.782979    4.525691    0.224277    0.327805    3.868902    0.128643
  00  17:29:58.258513 17:30:02.785919    4.527406    0.396144    0.783116    3.141108    0.240969
  02  17:29:58.280770 17:30:02.785971    4.505201    0.001450    0.001013    4.501923    0.001918
                                       ----------- ----------- ----------- ----------- -----------
                                          8299    0.621872    1.111       11.511933    0.371531


 SRB time      :     0.621872
 TCB time      :     1.111934
 Idle time     :    11.511933
 CPU Overhead  :     0.371531
                  -----------
         Total :    13.558299
```

Operates on the Intersection of all trace buffer

Summary for each processor

Summary for all process
SRB Time, TCB time, Id
CPU Overhead

Lots of Idle Time
Caution: zIIPs, zAAPs processe
The same as GP

```
Found 104 address spaces in SYSTRACE.
Found 230 SRB and SSRB PSWs in SYSTRACE.

CPU breakdown by ASID:

ASID Jobname    SRB Time      TCB Time    Tota
---- -------- ------------ ------------ -------
0067 RMF           0.006995     0.005196     0.01219
0079 ZESCM         0.000095     0.002359     0.002454
0024 JES2MON       0.003534     0.057579     0.061114
0001 *MASTER*      0.006579     0.002225     0.008805
006B NET           0.135695     0.018589     0.154284
0006 XCFAS         0.131358     0.367775     0.499133
0071 RCS           0.003644     0.001970     0.005614
0083 EKM           0.007061     0.001013     0.008075
0061 TCPIP         0.079679     0.019590     0.099269
0075 TN3270        0.045144     0.027801     0.072946
0085 ZT17MSTR      0.012100     0.005796     0.017897
0027 NFSC          0.000470     0.000895     0.001366
0056 ZCICSZJ       0.001101     0.000760     0.001862
005C SMSVSAM       0.022346     0.055299     0.077646
000B WLM           0.004230     0.117218     0.121448
0029 JES2S001      0.004554     0.004564     0.009119
0026 JES2          0.002373     0.040293     0.042666
```

... lines omitted ...

```
00A7 LDAPSRV       0.000194     0.000000     0.000194
0051 BPXOINIT      0.000993     0.003040     0.004033
                ------------ ------------ ------------
                   0.621872     1.111934     1.733807 (No.  f  PUs in Systrace:  3)
```

Total address spaces observed
Running in the intersect of trace buffer

Breakdown by asid / jobn
SRB Time, TCB Time
Total Time
Caution: Not sorted!

Caution: Includes work ru
On zIIP and zAAPs
If the problem is on a GP
Data may not be applic

Totals

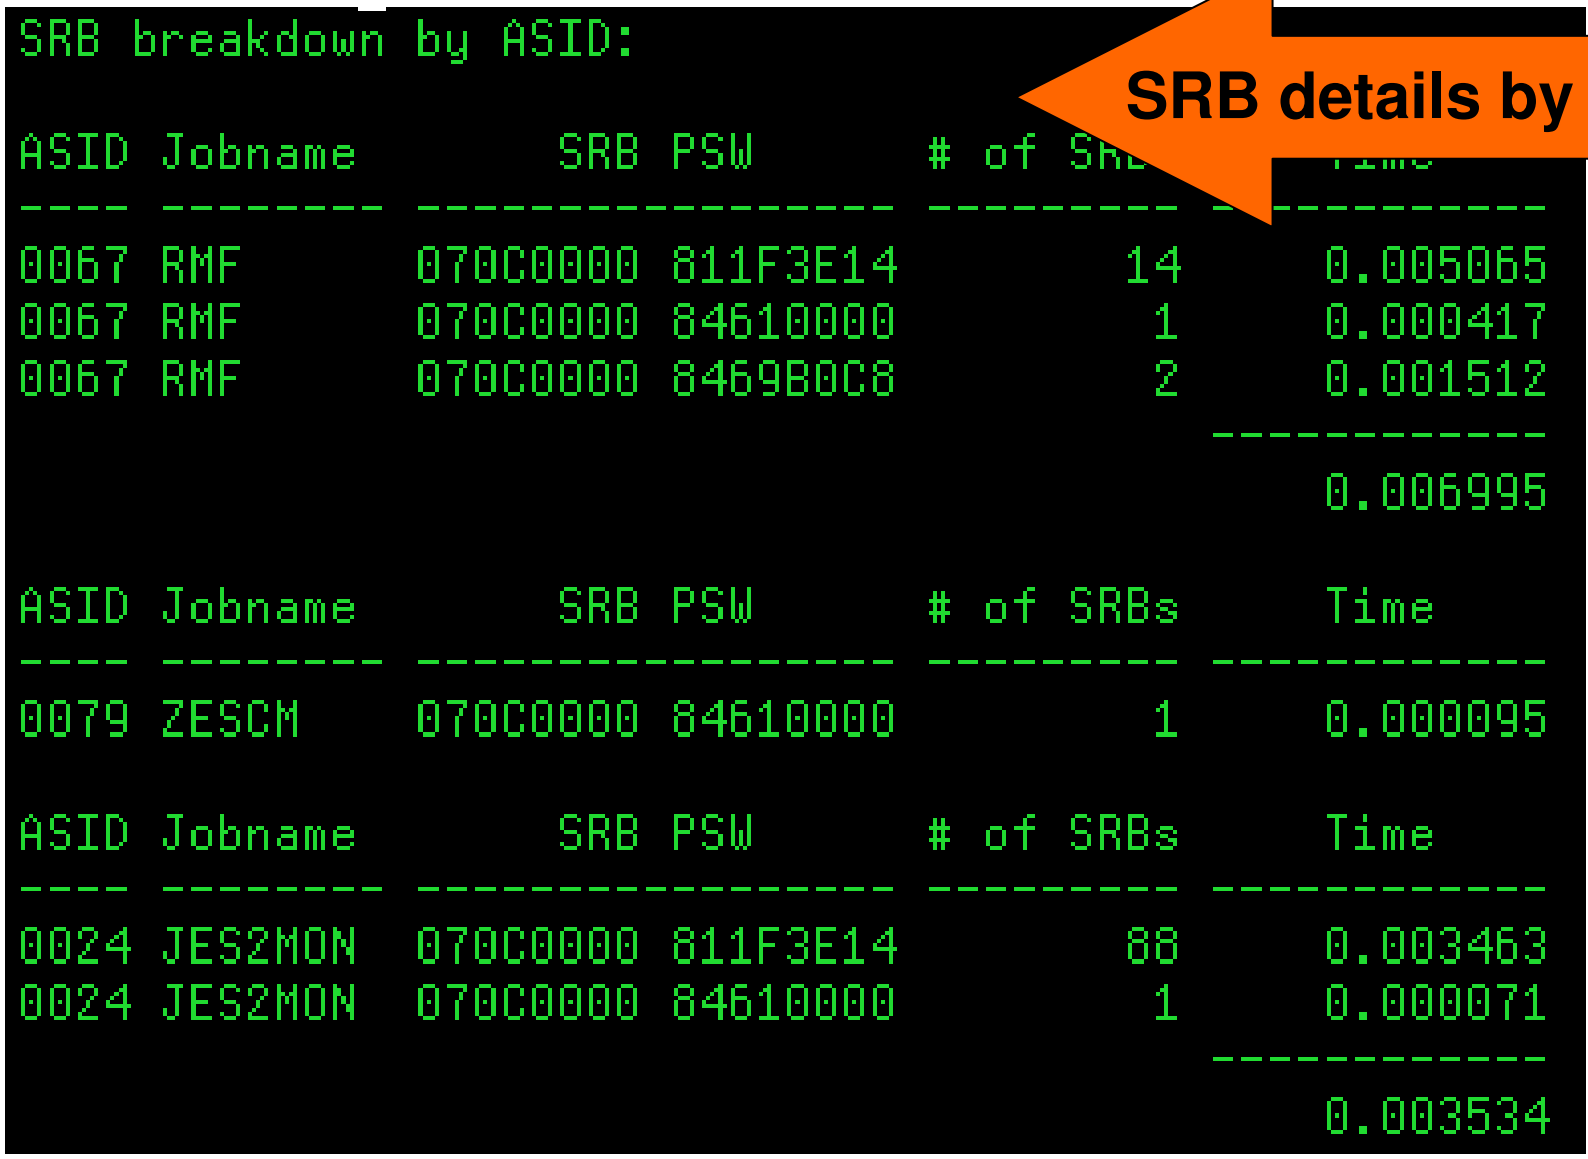

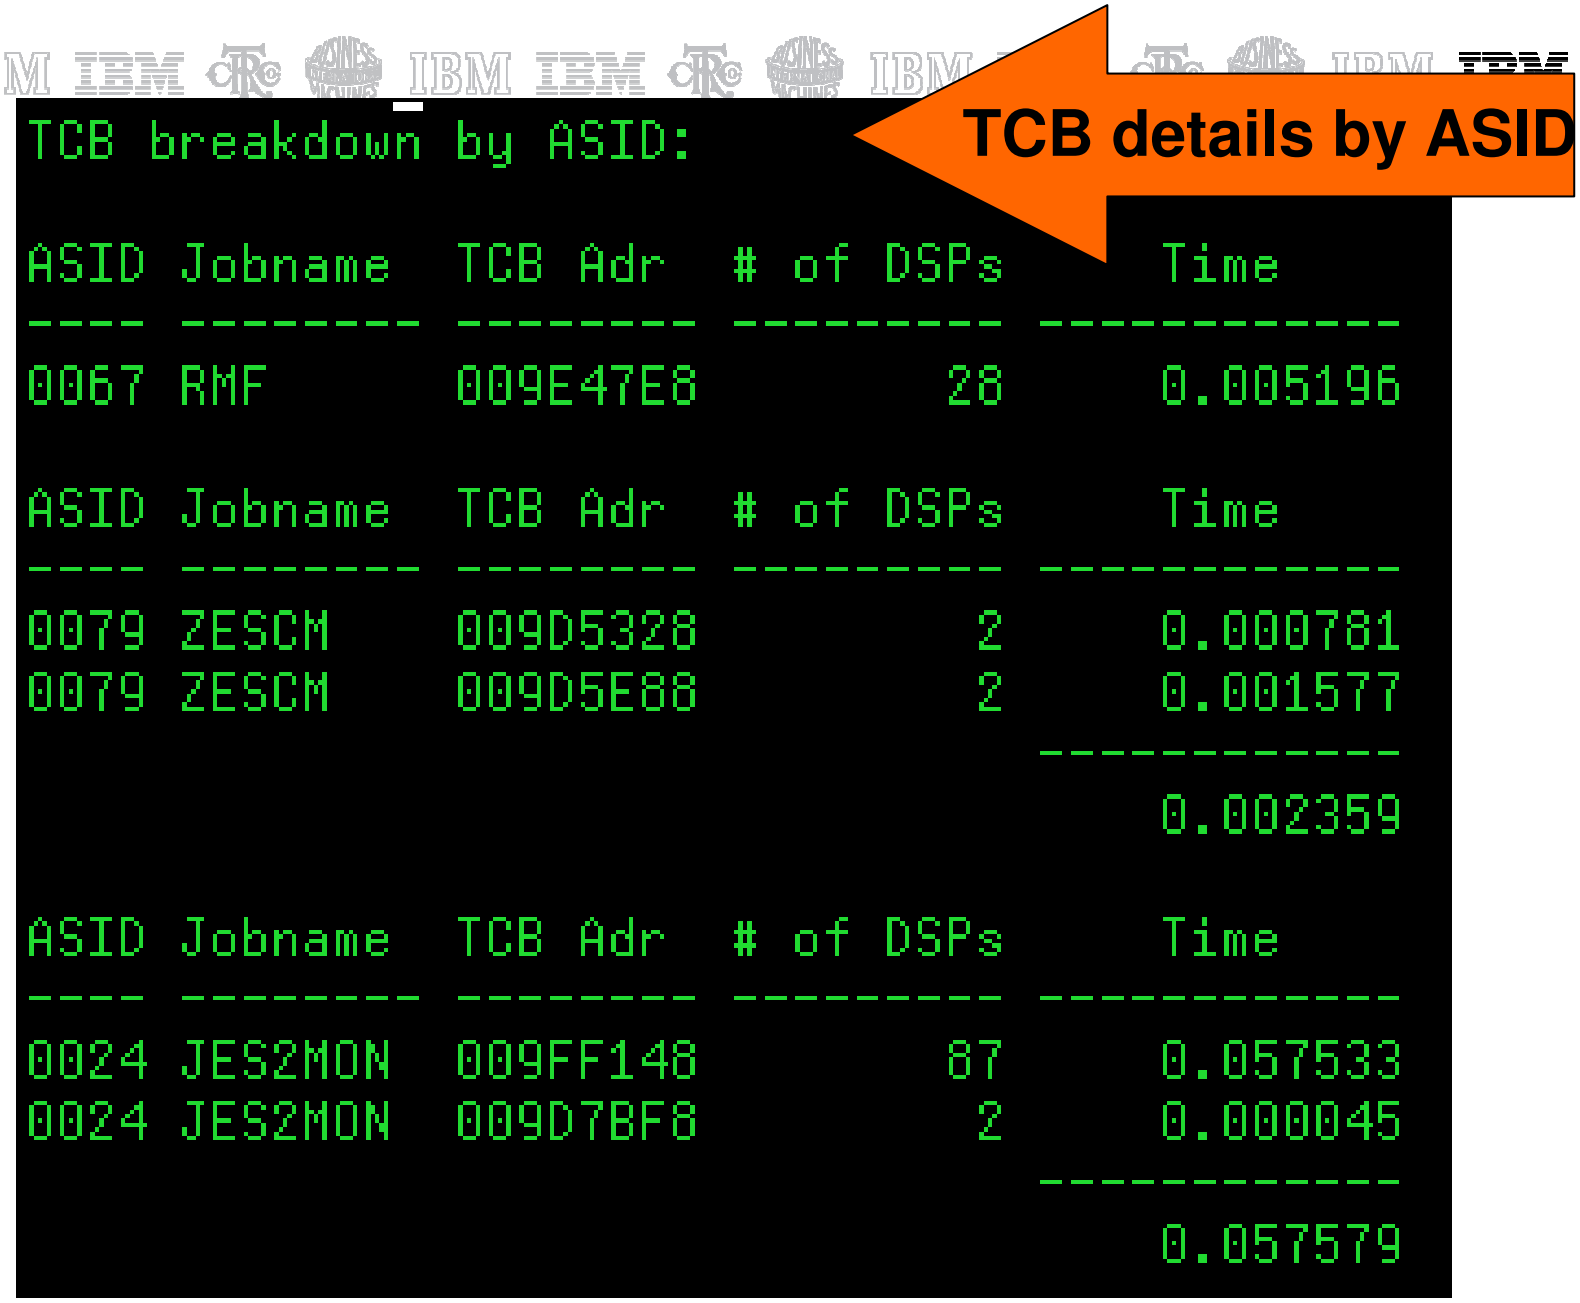

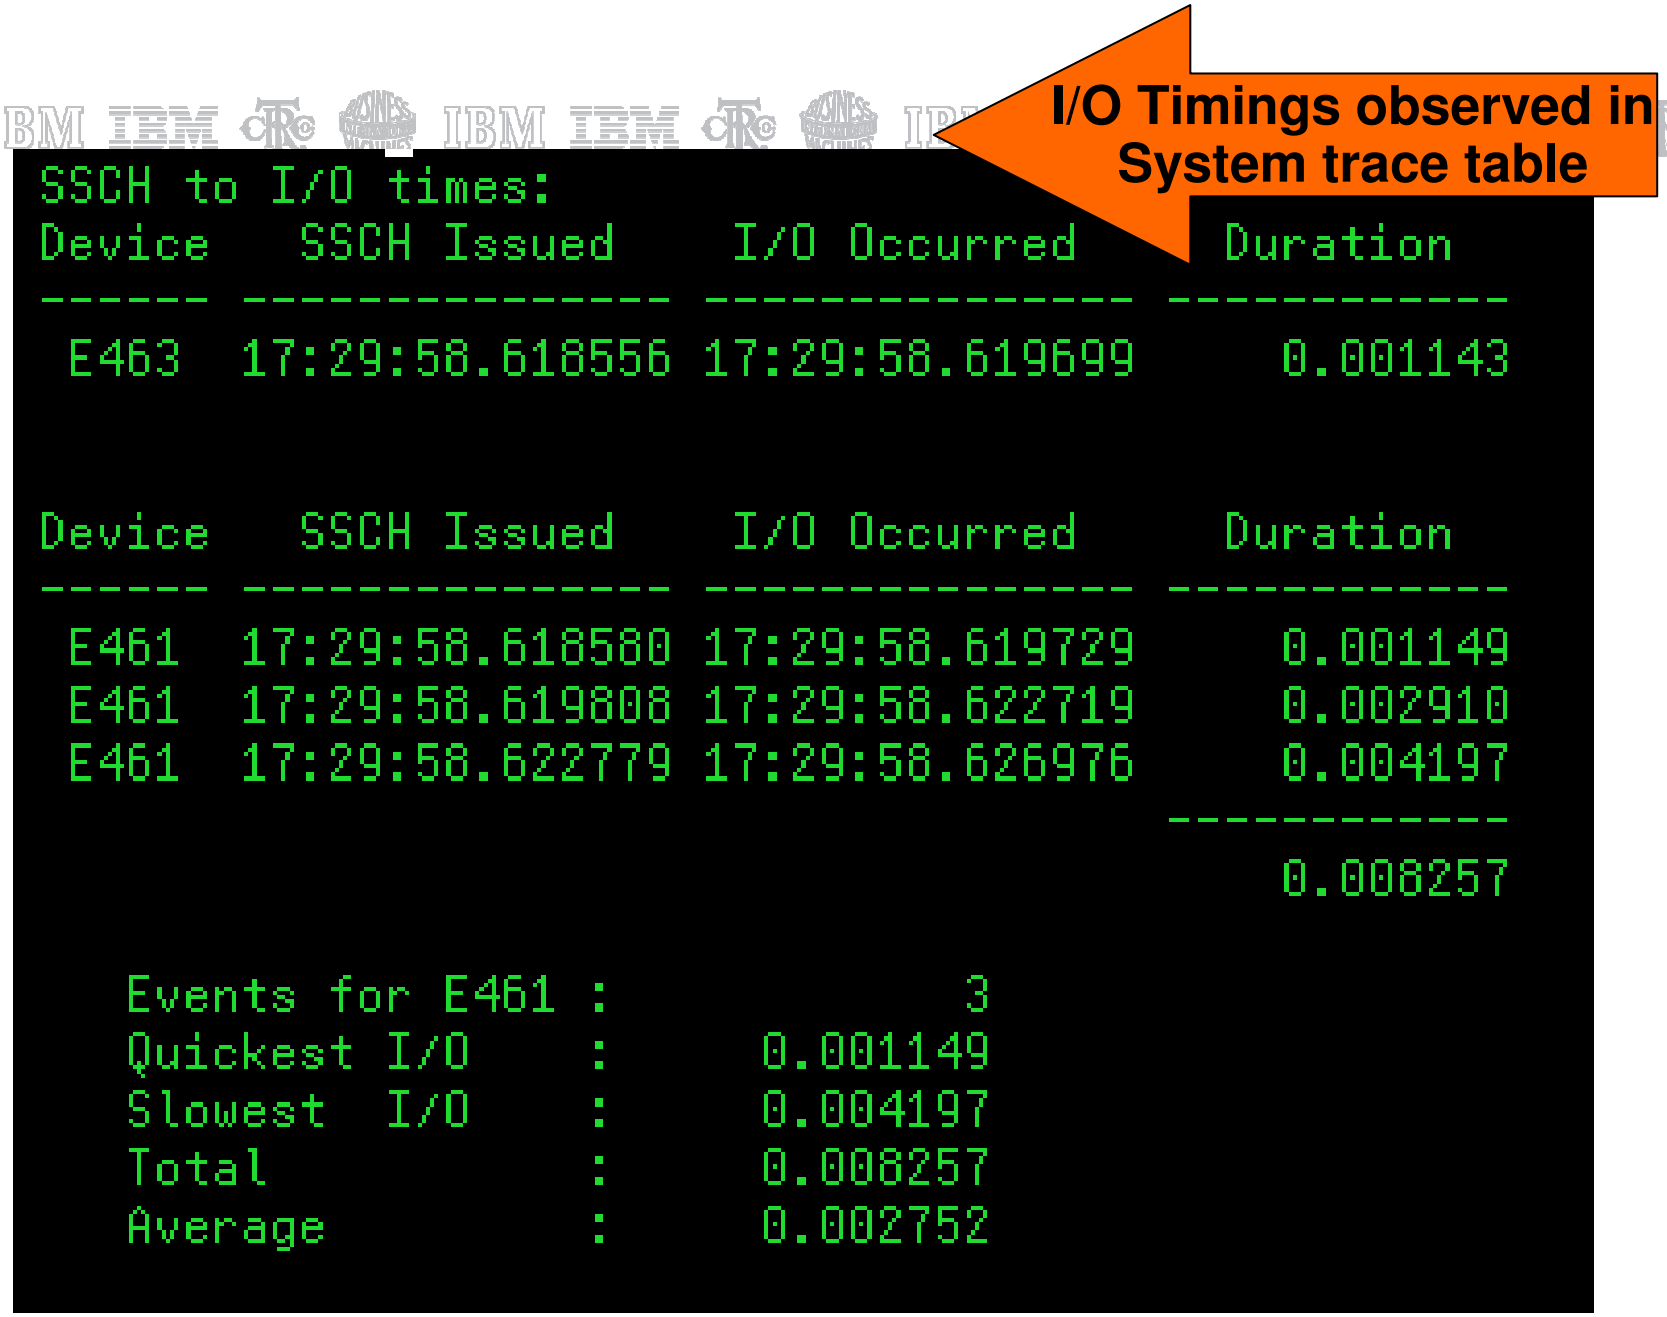

IP SYSTRACE PERFDATA(DOWHERE)

Summary sections at the top are the same

```
SRB breakdown by ASID: (WHERE command bypassed for CPU usage less than 0.100000):
```

WHERE success same ratio as WHERE success in dumps

```
ASID Jobname       SRB PSW      # of SRBs     Time                                    Where Info
---- -------- ---------------- --------- ------------ ----------------------------------------
006B NET      070C0000 811F3E14        61     0.022359
Errors detected in STRUCTURE(Xtlst) at NOCPU ASID(X'0001') FBCA2C:
Located via STRUCTURE(Cdemajor) at NOCPU ASID(X'0001') FBCA00
   Not on doubleword boundary
Errors detected in STRUCTURE(Cdemajor) at NOCPU ASID(X'0001') FBCA00:
   Locator CDXLMJP=00FBCA2C It may be damaged
006B NET      070C0000 B6A4C6D0       293     0.112009 RC=  4 for WHERE 36A4C6D0 ASID(07) command
006B NET      070C0000 8105BDC8         4     0.000187
006B NET      070C0000 816B8F00         1     0.000705
006B NET      070C0000 84610000         1     0.000432
                                          ------------
                                            0.135695
```

IP SYSTRACE PERFDATA(DOWHERE) cont.

TCB breakdown by ASID:

ASID	Jobname	TCB Adr	# of DSPs	Time
0006	XCFAS	009F8E88	206	0.336749
0006	XCFAS	009F9AA0	179	0.011334
0006	XCFAS	009FAA30	239	0.019591
0006	XCFAS	009F8CD8	4	0.000054
0006	XCFAS	009FAE88	3	0.000045
				0.367775

IP SYSTRACE PERFDATA(DOWHERE) cont.

```
Lock Events:
Lock ASID TCB/SRB  Type PSW Adr   Suspended at     Resumed at    Suspend Time
---- ---- -------- ---- -------- --------------- --------------- ------------
CEDQ 0009 009FF230 TCB  B840F2A0 17:30:00.988762 17:30:00.988777     0.000014


CLKC Events:

ASID Jobname  SRB/TCB      Clkc PSW      Where processing (CPU usage for this ASID is:     0.499133)
---- -------- -------- ---------------- ------------------------------------------------------------
0006 XCFAS    009FAA30 070C2000 80FE81CA ASID(X'0006') 00FE81CA. IEANUC01.IEAVTSFR+03CA IN READ ONLY NUCLEUS
0006 XCFAS    009F8E88 070C0000 80FF2086 ASID(X'0006') 00FF2086. IEANUC01.IEAVELK+048E IN READ ONLY NUCLEUS
0006 XCFAS    009FAA30 070C0000 80FF2086 (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA ASID(X'0006') 00FF24FA. IEANUC01.IEAVELK+0902 IN READ ONLY NUCLEUS
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
0006 XCFAS    009F8E88 070C0000 80FF24FA (Same as above)
```

System Trace – CPUMASK

Horror: More processors, more data to dig through in the system trace table and elongated problem identification times.

New Technique: CPUMASK and CPUTYPE

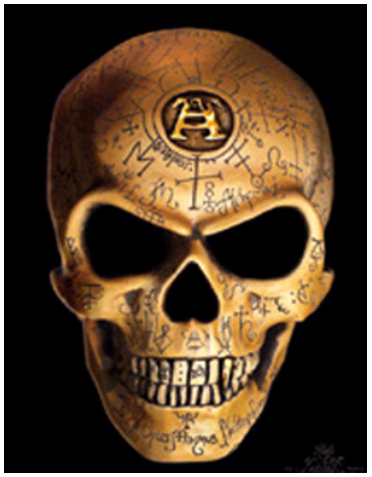

IP SYSTRACE CPUMASK(24F)

- –X'24F' is a bit mask where each bit represents a CPU
  - X'24F' represents CPUs 3,5,8,9,10 and 11
- –CPU numbers start with CPU 0 and go to maximum number of CPUs 128
- –Really helpful for very large LPARs with a high number of CPUs.

z/OS 1.13

System Trace – CPUTYPE

Horror: More types of processors, where the profiles for the processors differ greatly as the work eligible to run on the processors differs, more data to dig through in the system trace table and elongated problem identification times.

New Technique: CPUTYPE

IP SYSTRACE CPUTYPE(ZAAP,ZIIP,STANDARD)

- ZAAP is abbreviated as ZA
- ZIIP is abbreviated as ZI
- STANDARD represents the classic general purpose processors and is abbreviated as CP or S
- Blanks or commas can be used to separate the options

z/OS 1.13

System Trace – CPUMASK and CPUTYPE

CPUMASK and CPUTYPE options can be combined the output is the union of the two

z/OS 1.13

DOCPU

Horror: Obtaining diagnostic data from each processor in a standalone dump

New Technique: DOCPU

IP DOCPU

IP DOCPU CPU(list)

- CPU(0) – only processor 0
- CPU(3:5) – processors 3,4 and 5
- CPU(x'B') – only processor 11

IP DOCPU CPUTYPE(ZA,ZI,S)

IP DOCPU CPUMASK(401)

z/OS 1.13

DOCPU EXEC((ipcs command))

DOCPU was built with the ability to execute an IPCS command against each processor

IP DOCPU CPU(0,1) EXEC((L 1000 LEN(50)))

- For processors 0 and 1, display 50 bytes of data starting at address 1000

Gotchya – If you accidentally enter DOCPU on a z/OS 1.13 SVC dump the command will take and no error message is returned. No data is returned either.

z/OS 1.13

The End...

Questions??

IPCS GRSTRACE Summary Output

MAJOR NAME: xmajorname

MINOR NAME: xminorname
SCOPE: xscope SYSNAME: xsysname STATUS: xstatus
ASID: xasid TCB: xtcb JOBNAME: xjobname
MASID: xmasid MTCB: xmtcb
Reserve Device: xdevice Volser: xvolser
Critical ENQ Time(s):
Request: xdate xtime
Contention: xdate xtime
Grant: xdate xtime
Delta Time Waiting: xdeltatime
Movewaiter: xdate xtime

z/OS 1.10

Example of IP VERBX GRSTRACE 'QNAME(''TES?ENQ'') SUMMARY'

IPCS GRSTRACE Summary Example

```
MAJOR NAME: TESTENQ

*   MINOR NAME: DUMMYENQ
     SCOPE: SYSTEMS   SYSNAME: S1       STATUS: *SHARED*    /OWN
       ASID: 0000002C   TCB: 006FF020    JOBNAME: GRSTOOL
       Critical ENQ Time(s):
         Request:     06/04/2007 15:30:05.804018
         Grant:       06/04/2007 15:30:05.834250
     SCOPE: SYSTEMS   SYSNAME: S1       STATUS: *SHARED*    /OWN
       ASID: 00000028   TCB: 006FF020    JOBNAME: GRSTOOL
      Critical ENQ Time(s):
        Request:     06/04/2007 15:32:18.460284
        Contention:  06/04/2007 15:32:18.484524
        Grant:       06/04/2007 15:32:34.846436
        Delta Time Waiting: 00:00:16.361911
     SCOPE: SYSTEMS   SYSNAME: S2       STATUS: *EXCLUSIVE* /WAIT
      ASID: 0000002F   TCB: 006FF020    JOBNAME: GRSTOOL
      Critical ENQ Time(s):
        Request:     06/04/2007 15:33:18.738913
      Some ENQ information is unavailable for this remote request
```

z/OS 1.10

IPCS GRSTRACE Detail Output

```
MAJOR NAME: xmajorname
  MINOR NAME: xminorname
     Resource Creation Time: xdate xtime
     Last Movewaiter Time:   xdate xtime
   SCOPE: xscope    SYSNAME: xsysname  STATUS: xstatus
     ASID: xasid      TCB: xtcb       JOBNAME: xjobname
     MASID: xmasid    MTCB: xmtcb
     Reserve Device: xdevice Volser: xvolser
     Critical ENQ Time(s):
       Request:     xdate xtime
       Contention:  xdate xtime
       Grant:       xdate xtime
       Delta Time Waiting: xdeltatime
       Movewaiter:  xdate xtime
    Caller PSW: xpsw              Caller TCB: xcallertcb
    Request Type: xrequesttype
    RNL Processing Actions: xrnlactions
    Affected by ISGNQXIT/FAST
    Affected by ISGNQXITBATCH/CND
    Managed by an Alternate Serialization Product
    ISGENQ Userdata:
     xuserdata
    QEL: xqeladdr              QXB: xqxbaddr
    QCB: xqcbaddr              SVRB: xsvrbaddr
```

z/OS 1.10