SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

The MVS System Trace

Tribal Knowledge Transfer (TKT) Episode I

Ed Jaffe, Phoenix Software International

Session 24249, 7 March 2023

SHARE Atlanta

Why Am I Discussing this Topic?

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Since the pandemic, we at PSI have observed a dramatic decrease in the competency of z/OS system programmers working at our customer sites.
- Discussions with my friends working for other ISVs tells me we're not alone.
- People have retired and new people took their place without adequate tribal knowledge transfer from the retirees.
- What shall we as a community do about this? There are many options…

- Sharing information with our peers is SHARE's "raison d'être."
- It's incumbent up all of us in the industry to share what we know with our "tribe."
- Therefore, I am utilizing this avenue in an effort to help preserve and pass long important mainframe skills that are in danger of being lost.

Recommended Prerequisites

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Familiarity with z/Architecture programing basics
  - PSW & Registers
  - Addressing modes (24-bit, 31-bit, 64-bit)
  - The notion of problem vs supervisor state
  - Familiarity with assembler language
  - Some dump reading skills (IPCS, SYSABEND, SYSUDUMP, SNAP)
- Patience! Dump and trace analysis is not about instant gratification!

- These slides are (or will be) available on the SHARE event site
- You will have plenty of time to study them in detail at a later time

Today's Agenda

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Understanding system trace
- Using system trace to diagnose various problem scenarios
- Expanding system trace to capture additional events
- Adding your own entries to system trace
- Configuring system trace to best assist first-failure problem diagnosis

What is System Trace?

- System trace is a continuous, ongoing record of hardware and software events occurring during system initialization and operation.
- There is a separate trace table for each processor (one per CP and zIIP thread) and the buffers are written to the TRACE address space.
- It appears in machine-readable SVC dumps, stand-alone dumps, transaction dumps (IEATDUMP), and SYSMDUMPs as well as formatted SYSUDUMP, SYSABEND, and SNAP dumps (w/ SDATA=TRT).
- System trace is activated automatically during system initialization.
- Subsequently, you can use the TRACE command on a console with master authority (or via COMMNDxx) to customize system tracing.
- System trace is formatted in a machine-readable dump using the IPCS SYSTRACE command (wide screen recommended).

Which Events are Recorded?

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

TTE	Description
ACR	Alternate CPU recovery
AINT	Adapter interruption
ALTR	Alteration of trace option
BR	BAKR, BALR, BASR, or BASSM
BSG	Branch in subspace group
CSCH	Clear subchannel operation
DSP	Task dispatch
EXT	CALL, CLKC, EMS, SS, TIMR, WTI
HSCH	Halt subchannel operation
I/O	Input/output interruption
MCH	Machine check interruption
MOBR	AMODE change & branch
MODE	Change of addressing mode
MSCH	Modify subchannel operation
PC	Program Call instruction

TTE	Description
PCIL	PCI load instruction
PCIS	PCI store instruction
PDMX	PCIE adapter de-mux event
PGM	Program interruption
PR	Program Return instruction
PT	Program Transfer instruction
RCVY	Recovery event
RSCH	Resume subchannel operation
RST	Restart interruption
SIGA	Signal adapter operation
SPER	SLIP program event recording
SPIN	System spin start/middle/stop
SPR2	SLIP event w/STDATA specified
SRB	Initial SRB dispatch
SSAR	Set Secondary ASC mode

TTE	Description
SSIR	SSAR instruction w/instance
SSCH	Start subchannel operation
SSRB	Suspended SRB dispatch
SSRV	Internal system service call
SUSP	Lock suspension
SVC	Supervisor call interruption
SVCE	SVC error
SVCR	SVC return
SYNS	Synchronous I/O start
SYNE	Synchronous I/O end
TIME	TOD clock dynamic adjustment
USRn	User event
WAIT	Wait task dispatch
XSCH	Cancel subchannel operation
?EXPL	Unidentified trace entry

How Do Trace Entries Get Created?

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- The operating system allocates fixed, virtual storage for the 4K TBUFs and loads and maintains control register 12 (CR12) as needed.
- CR12 contains:
  - the real address of the next-to-be filled trace entry
  - bits to enable/disable the three implicit trace types (branch, mode, ASN)
  - a bit to enable/disable explicit trace (TRACE/TRACG instructions)
- When a trace entry is created, the CPU advances the address in CR12.
- If the entry would extend past the end of the current 4K TBUF, a trace exception occurs, whereby the interrupt handler:
  - points CR12 to the real frame backing the next-to-be-used 4K TBUF
  - retries the instruction that failed due to the trace exception
- For obvious reasons, trace exceptions are not recorded in the trace

Implicit Trace Entries

- The CPU automatically creates the implicit trace entries shown when the listed instruction is executed and the appropriate trace bit is enabled
- The operating system has no control over the content of these trace entries
- **Note:** This table, from the z16 PoOps, has a change bar for BSM. The z15 version had "–" (none) in the **ASN and Mode** column.

	Implicit Tracing Enabled						
	Branch	ASN	Mode	Branch and ASN	Branch and Mode	ASN and Mode	All
Instruction	Trace Entries Made						
BAKR	B	–	–	B	B	–	B
BALR	B	–	–	B	B	–	B
BASR	B	–	–	B	B	–	B
BASSM	B	–	MS	B	B I MSB	MS	B I MSB
BSA	B	–	–	B	B	–	B
BSG	B	BSG	–	BSG	B	BSG	BSG
BSM	–	–	MS	–	MS	MS	MS
PC-20	–	PC	MS	PC	MS	PC & MS	PC & MS
PC-32	–	PC	MS	PC	MS	PC	PC
PR-b	–	–	MS	–	MS	MS	MS
PR-pc	–	PR	MS	PR	MS	PR	PR
PT or PTI	–	PT	–	PT	–	PT	PT
RP	B	–	MS	B	B I MSB	MS	B I MSB
SSAR or SSAIR	–	SSAR	–	SSAR	–	SSAR	SSAR
SAM24/31/64	–	–	MS	–	MS	MS	MS
TRAP2/4	B	–	–	B	B	–	B

Explanation:

–	None.
-20	The case when PROGRAM CALL uses a 20-bit PC number.
-32	The case when PROGRAM CALL uses a 32-bit PC number.
-b	The case when PROGRAM RETURN unstacks a branch state entry.
-pc	The case when PROGRAM RETURN unstacks a program-call state entry.
I	OR.
&	AND.
B	Branch trace entry. Made only if the branch is taken and a mode-switching-branch trace entry is not made.
MS	Mode-switch trace entry. Made only if PSW bit 31 is changed.
MSB	Mode-switching-branch trace entry. Made only if PSW bit 31 is changed (which can occur only if the branch is taken.

Figure 4-6. Summary of Implicit Tracing

Explicit Trace Entries

- The remaining events are traced when the operating system issues:

TRACE $R_1,R_3,D_2(B_2)$ [RS-a]

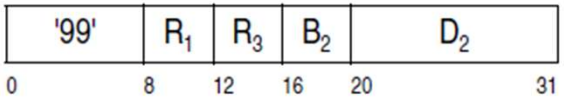

TRACG $R_1,R_3,D_2(B_2)$ [RSY-a]

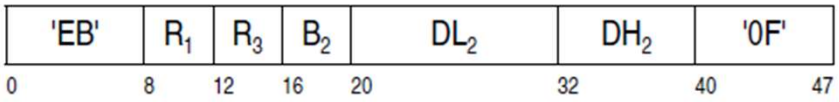

When explicit tracing is on (bit 63 of control register 12 is one), the second operand, which is a 32-bit word in storage, is fetched, and bit 0 of the word is examined. If bit 0 of the second operand is zero, a trace entry is formed at the real-storage location designated by control register 12.

If explicit tracing is off (bit 63 of control register 12 is zero), or if bit 0 of the second operand is one, no trace entry is formed, and no trace exceptions are recognized.

The displacement for TRACE is treated as a 12-bit unsigned binary integer. The displacement for TRACG is treated as a 20-bit signed binary integer.

The trace entry is composed of an entry-type identifier, a count of the number of general registers whose partial or entire contents are placed in the entry, a field whose contents indicate whether the entry was formed by TRACE (TRACE) or TRACE (TRACG), selected bits of the TOD clock, the second operand, and the partial or entire contents of a range of general registers. For TRACE (TRACE), bits 16-63 of the TOD clock and bits 32-63 of the general registers are placed in the trace entry. For TRACE (TRACG), bits 1-7 of the epoch index, bits 0-79 of the clock, and bits 0-63 of the registers are placed in the entry. See "Trace Entries" on page 4-15 for further details.

General System Trace Format

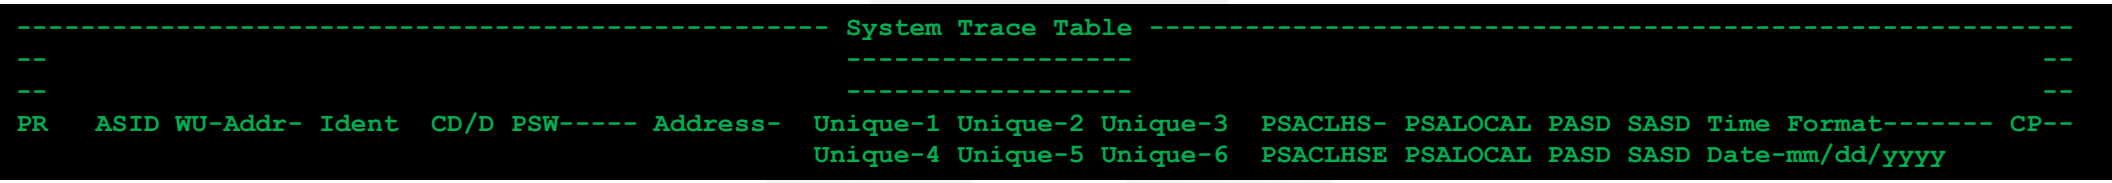

- PR – logical processor id.
- ASID – address space id.
- WU-Addr – 24-bit TCB address or 31-bit WEB address for SRB.
- Ident – TTE identifier.
- CD/D – sub-identifier, device, mask, etc.
- PSW Address – 128-bit PSW split across two lines with right half on first line.
- Unique1 – Unique6 – varies greatly by TTE type but often contain GPRs 15-1.
- PSACLHS[E] – lock status from PSA. Descriptive text for some entries.
- PSALOCAL – locally locked address space indicator from PSA.
- PASD – primary ASID. Descriptive text for some entries.
- SASD – secondary ASID. Descriptive text for some entries.
- Time Format – time of day local or hex.
- CP – four hex digits of model-dependent info to help identify the physical CP.

General System Trace Format

- The oldest trace entries appear at the beginning and the newest appear at the end. The trace is read top-down.
- Trace entries with a dash (-) between the logical processor and ASID columns appear outside the range where trace data from all processors is available. Avoid using such entries!
- Trace entries with an asterisk (*) before the identifier represent an unusual condition, for example an unhandled program check or an I/O error.
- Embedded messages alert you when trace data from all processors is or is not available and when there are significant time gaps on a given processor. Time gaps have become more common since the advent of HiperDispatch.

System Trace Tells a Remarkable Story

SHARE
EDUCATE • NETWORK • INFLUENCE

- After an abend or during a loop, the system trace relates a story so valuable that many problems are solved by trace inspection alone and some problems can't be solved without it.
- It is the first thing I look at right after inspecting PSW and registers.
- Generally, customer trace tables are sized to hold no less than one full second of trace history. If they are ¼ sec or less, we inform the customer and politely suggest their configuration may be sub-optimal.
- The following, from a random customer dump, is exemplary…

```
******** Trace data is not available from all processors before this time.
0000 0064 009F8130  I/O  094F2 00000000_3174F2C6  00104007 3FFCE440 0C000100  00000080 00000000 0064 0064 18:11:24.968806982 0077
                           07842000 80000000           026E3D90 00400001  00000000
~
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~
0001 0045 009D4E88  EXT    EMS 00000000_010DBC9A  00011201 20000000 27A2A398  04000001 00F9E300 0004 0045 18:11:27.556463174 0077
                           05041000 80000000  8146C0E8                    00000000
0001 0045 009D4E88  PC     ...   0       0146C202             00503       SysTrace Processor Snap
******** Trace data is not available from all processors after this time.
```

~2.6 sec!

Make Good Use of the Documentation

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Practice makes perfect; the more you traces you examine, the more you will recognize and interpret without assistance.
- But make no mistake, even people that analyze many traces refer to the manual on a regular basis. Mine is on “speed dial!”
  - https://www.ibm.com/docs/en/zos/2.5.0?topic=trace-reading-system-output

Reading system trace output

Last Updated: 2021-10-28

The following topics describe system trace table entries (TTE) as they appear in a dump formatted with the IPCS SYSTRACE subcommand.

- Example of a system trace in a dump
- Summary of system trace entry identifiers shows a table of the system trace identifiers for each system trace entry in a dump and shows where you can find the format of the entry in this section. If you are looking for a particular entry start with this table, because many of the entries are similar and are grouped together.

A Little More Background

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- The MVS Recovery Termination Manager (RTM) creates some of the trace entries we will be looking at today. There are two components:
  - RTM1:
    - The part of RTM that runs in the unit of work that is terminating.
    - Generally given control by the Second-Level Interrupt Handler (SLIH) e.g., to handle a program check or by code that issues the CALLRTM macro.
    - In many cases, after doing preliminary analysis, RTM1 uses RTM2 to finish the job.
  - RTM2:
    - The part of RTM that runs under an SVRB in a user TCB. It handles:
      - Normal end initiated by issuing SVC 3 (Exit).
      - Abnormal end initiated by issuing SVC 13 (ABEND).
    - Drives recovery routines and resource managers, schedules dumps, removes RBs from the RB chain, cleans up TCBs, cleans up address spaces, etc.

Common TTEs: SVC/SVCR PC/PR SSRV PGM

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

```
0002 00A8 00AA4040  SVC     7B 00000000_05382FA4  00000000 00000000 01474E10  PurgeDq                        15:09:25.255927355 0002
                               47043000 80000000
0002 00A8 00AA4040  SVC     4F 00000000_00FEE6BC  00FBC880 0000000D 0000000D  Status   Stop  SRBs/TCBs     15:09:25.255927597 0002
                               47040000 80000000
0002 00A8 00AA4040  SVCR    4F 00000000_00FEE6BC  00000000 00FBC880 7FFFAE98                                 15:09:25.255951981 0002
                               47040000 80000000
0002 00A8 00AA4040  PGM    030 00000000_0113C3D6  00040030 00000000           80000001 00000000 00A8 00A8 15:09:25.255964205 0002
                               44042000 80000000           7F645401           00000000
0002 00A8 00AA4040  PC     ...   0      0113C3DA             0030B            Storage  Obtain
0002 00A8 00AA4040  SSRV   132         00000000  0000EF02 00000318 0276B6D0  Storage  Obtain                15:09:25.255984901 0002
                                                  00A80000
0002 00A8 00AA4040  PR     ...   0      0113C3DA 014EE42C                                    00A8
0002 00A8 00AA4040  PGM    031 00000000_0113C3F0  00040031 00000000           80000001 00000000 00A8 00A8 15:09:25.255993106 0002
                               44042000 80000000           7F645401           00000000
0002 00A8 00AA4040  PGM    030 00000000_0113D5DA  00040030 00000000           80000001 00000000 00A8 00A8 15:09:25.256050219 0002
                               44041000 80000000           7F645401           00000000
0002 00A8 00AA4040  PC     ...   0      0113D5DE             00311            Storage  Release
0002 00A8 00AA4040  SSRV   133         00000000  0000EF03 00000318 0276B6D0  Storage  Release               15:09:25.256062518 0002
                                                  00A80000
0002 00A8 00AA4040  PR     ...   0      0113D5DE 014EE42C                                    00A8
0002 00A8 00AA4040  PGM    031 00000000_0113D5DE  00020031 00000000           80000001 00000000 00A8 00A8 15:09:25.256066867 0002
                               44041000 80000000           7F645401           00000000
0002 00A8 00AA4040  SVC     4F 00000000_00FEE9A4  00FBC880 0000000D FFFFFFF3  Status   Start SRBs/TCBs     15:09:25.256099591 0002
                               47041000 80000000
0002 00A8 00AA4040  SVCR    4F 00000000_00FEE9A4  00000000 00000000 03D0DE00                                 15:09:25.256111901 0002
                               47041000 80000000
0002 00A8 00AA4040  SVCR    7B 00000000_05382FA4  00000000 00000000 03D0DE00                                 15:09:25.256115856 0002
                               47043000 80000000
```

Common TTEs: SVC/SVCR PC/PR SSRV PGM

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

```
0002 00A8 00AA4040  SVC     7B 00000000_05382FA4  00000000 00000000 01474E10  PurgeDq                          15:09:25.255927355 0002
                               47043000 80000000
0002 00A8 00AA4040  SVC     4F 00000000_00FEE6BC  00FBC880 0000000D 0000000D  Status   Stop  SRBs/TCBs         15:09:25.255927597 0002
                               47040000 80000000
0002 00A8 00AA4040  SVCR    4F 00000000_00FEE6BC  00000000 00FBC880 7FFFAE98                                   15:09:25.255951981 0002
                               47040000 80000000
0002 00A8 00AA4040  PGM    030 00000000_0113C3D6  00040030 00000000           80000001 00000000 00A8 00A8 15:09:25.255964205 0002
                               44042000 80000000           7F645401           00000000
0002 00A8 00AA4040  PC     ...    0      0113C3DA             0030B           Storage  Obtain
0002 00A8 00AA4040  SSRV   132            00000000  0000EF02 00000318 0276B6D0 Storage  Obtain                15:09:25.255984901 0002
                                                     00A80000
0002 00A8 00AA4040  PR     ...    0      0113C3DA 014EE42C                                  00A8
0002 00A8 00AA4040  PGM    031 00000000_0113C3F0  00040031 00000000           80000001 00000000 00A8 00A8 15:09:25.255993106 0002
                               44042000 80000000           7F645401           00000000
0002 00A8 00AA4040  PGM    030 00000000_0113D5DA  00040030 00000000           80000001 00000000 00A8 00A8 15:09:25.256050219 0002
                               44041000 80000000           7F645401           00000000
0002 00A8 00AA4040  PC     ...    0      0113D5DE             00311           Storage  Release
0002 00A8 00AA4040  SSRV   133            00000000  0000EF03 00000318 0276B6D0 Storage  Release               15:09:25.256062518 0002
                                                     00A80000
0002 00A8 00AA4040  PR     ...    0      0113D5DE 014EE42C                                  00A8
0002 00A8 00AA4040  PGM    031 00000000_0113D5DE  00020031 00000000           80000001 00000000 00A8 00A8 15:09:25.256066867 0002
                               44041000 80000000           7F645401           00000000
0002 00A8 00AA4040  SVC     4F 00000000_00FEE9A4  00FBC880 0000000D FFFFFFF3  Status   Start SRBs/TCBs      15:09:25.256099591 0002
                               47041000 80000000
0002 00A8 00AA4040  SVCR    4F 00000000_00FEE9A4  00000000 00000000 03D0DE00                                   15:09:25.256111901 0002
                               47041000 80000000
0002 00A8 00AA4040  SVCR    7B 00000000_05382FA4  00000000 00000000 03D0DE00                                   15:09:25.256115856 0002
                               47043000 80000000
```

SVC/SVCR Trace Entry Breakdown

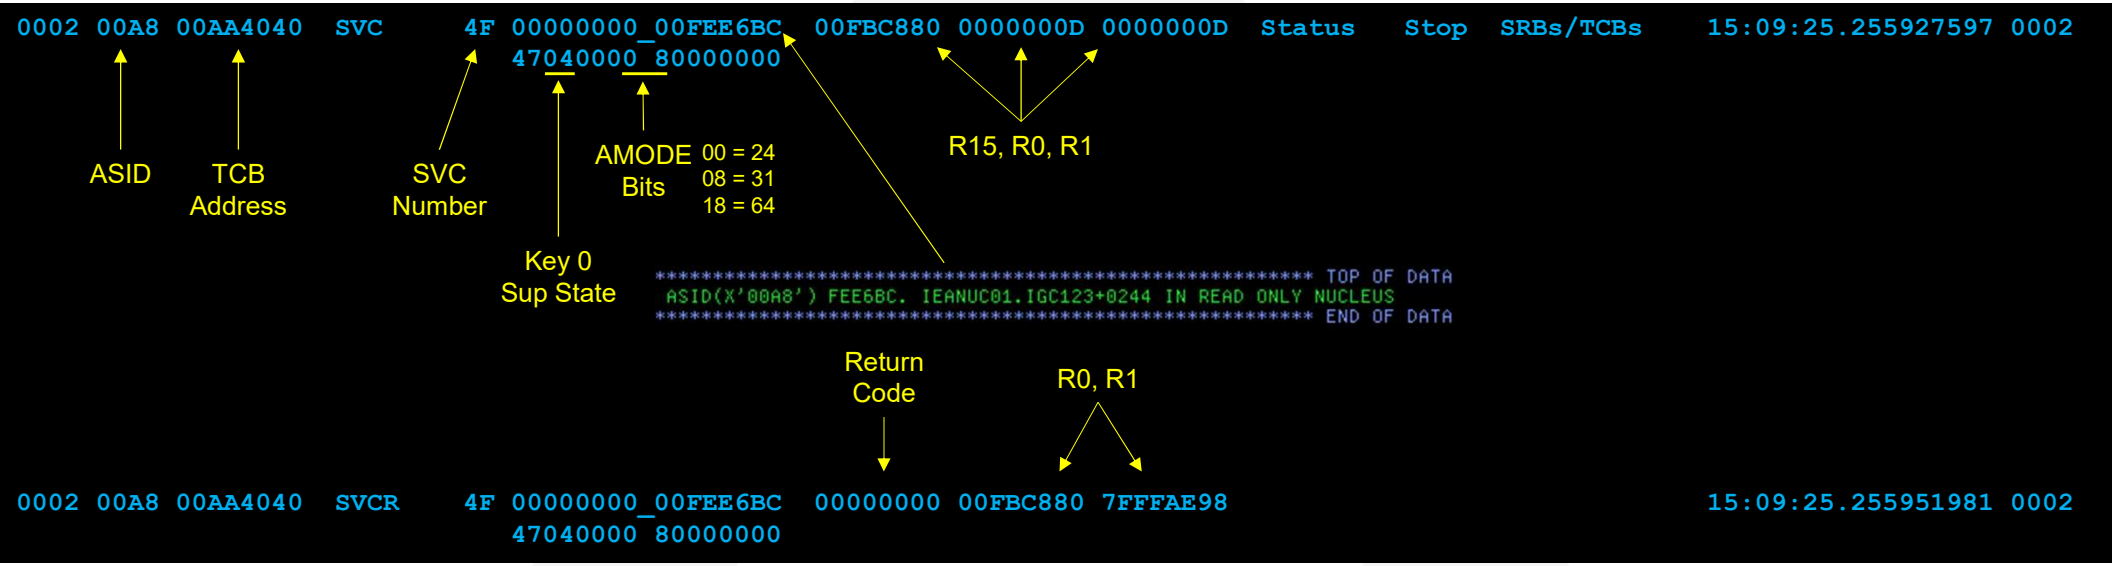

- If high-order byte of work unit address is zeros, you know it's a TCB.
- Non-zero would indicate a Work Element Block (WEB) for an SRB
- Execution location obtained via IP W 00FEE6BC command

PGM Trace Entry Breakdown

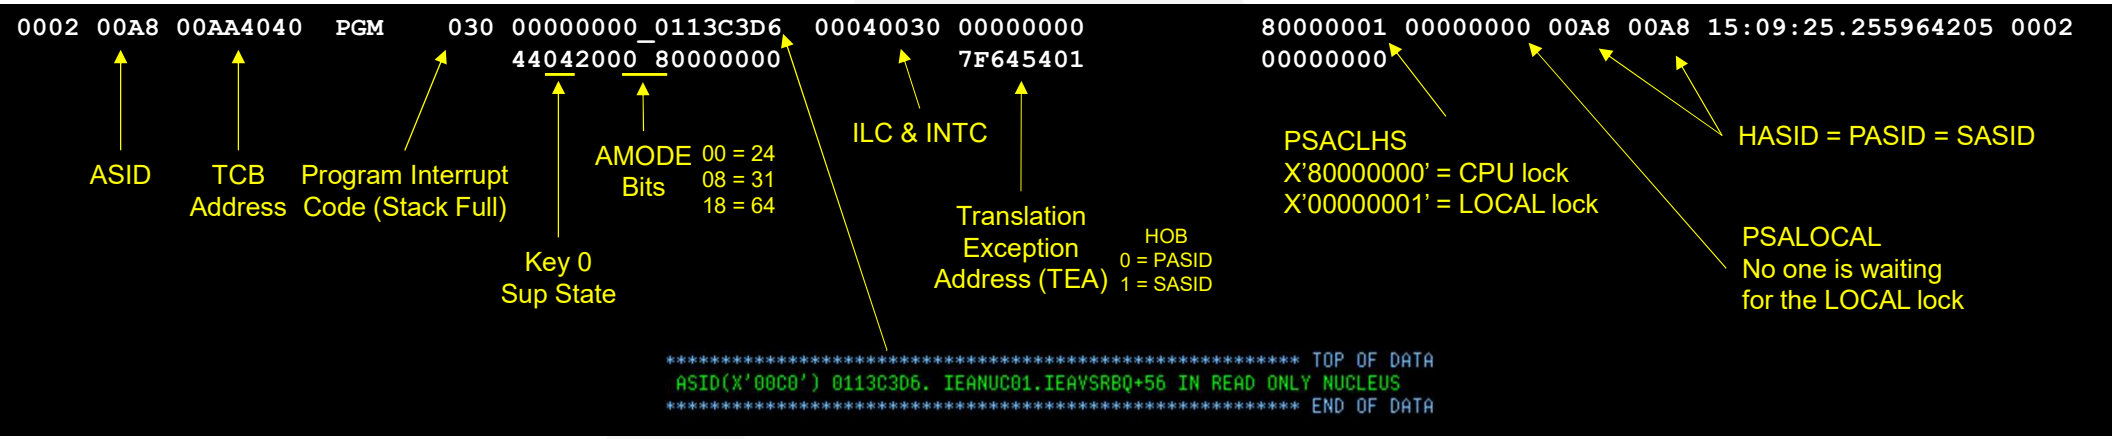

- Execution location obtained via IP W 0113C3D6 command
- Always ignore the last three nybbles of the TEA (in this case 7F645000)
  - The TEA has no relevance unless the program check is a translation exception
- SYS1.MACLIB(IHAPSA) has the definitions of the bits in PSACLHS and PSACLHSE
- PSALOCAL is a pointer to a FIFO list describing all other units of work waiting for you to release the local lock

PC/SSRV/PR Trace Entry Breakdown

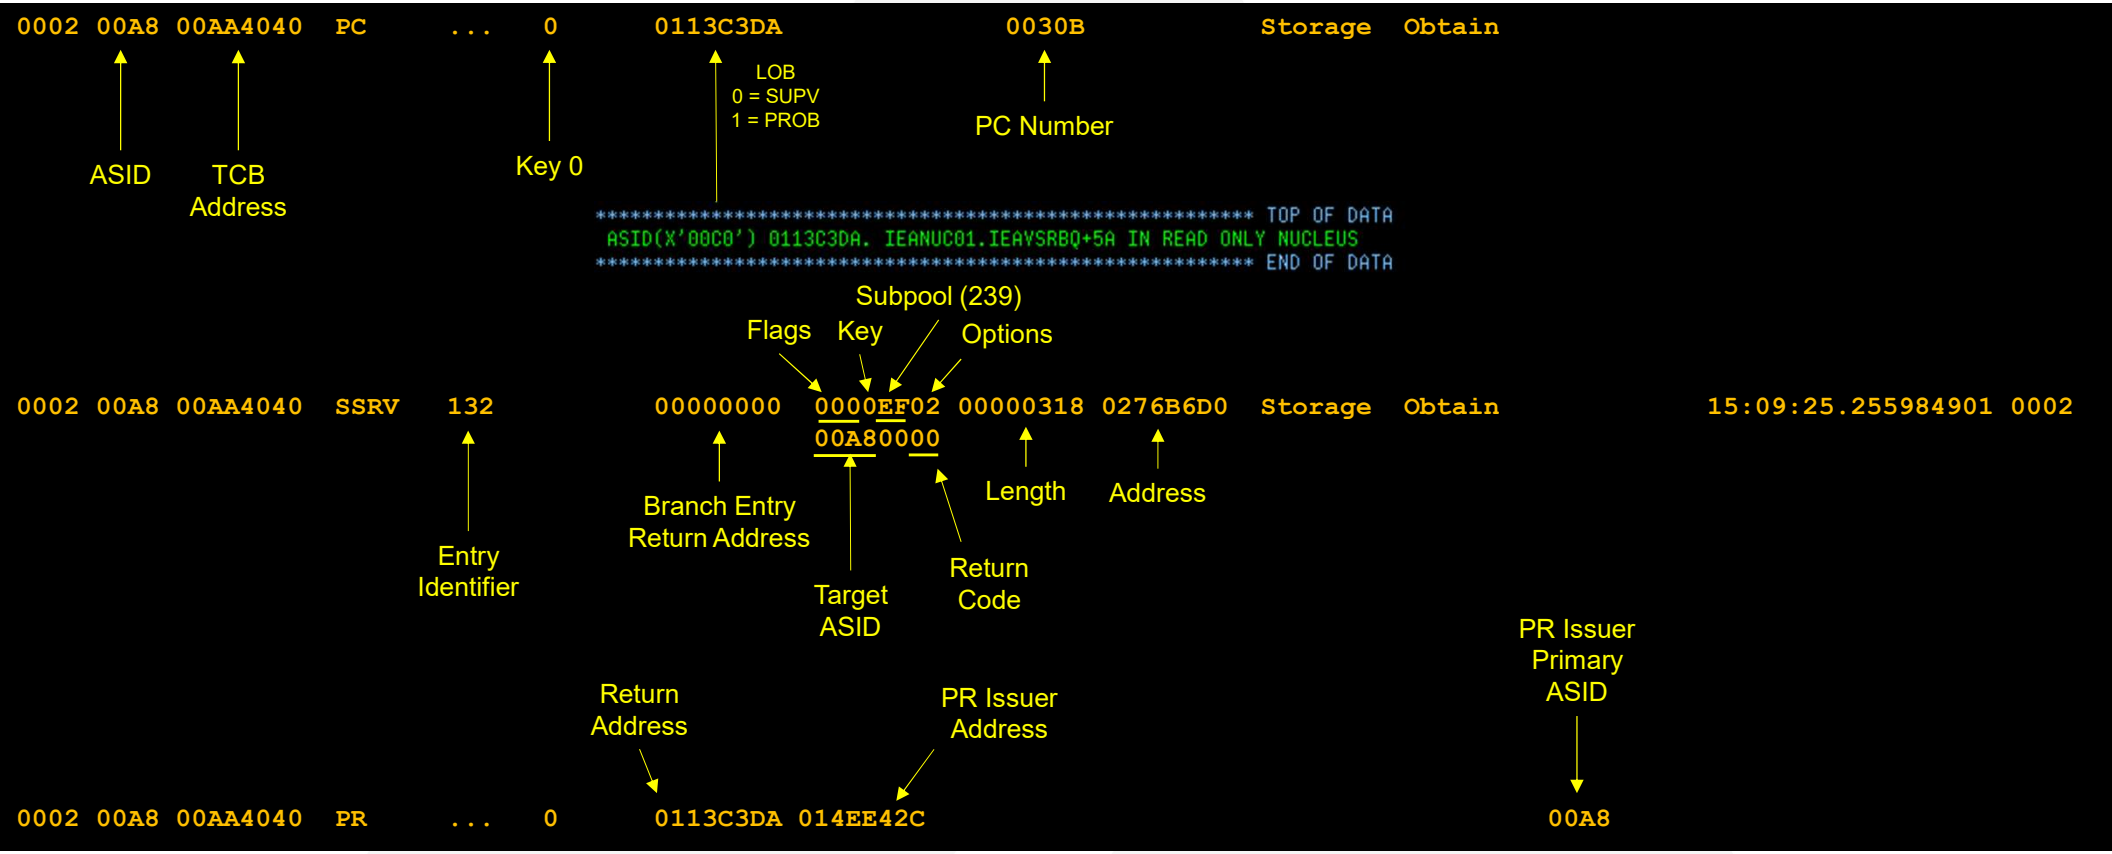

- There are 61 different SSRV IDs and each has unique parameters

My Favorite Ultra-Simple Test Programs

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- // EXEC PGM=ABEND001 to abend with U0001

```
ABEND001 CSECT ,
         ABEND 1,DUMP
         END   ,
```

- // EXEC PGM=ABEND0C1 to abend with S0C1

```
ABEND001 CSECT ,
         DC    D'0'
         END   ,
```

- // EXEC PGM=ABEND806 to abend with S806 ☺

Unhandled TCB Abend (U0001)

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

```
0000 00C0 00ABBE88  SVCR  FF00 00000000_00023000  00023000 00000064 00022000                              15:02:07.903563172 0001
                              47850000 00000000
0000 00C0 00ABBE88 *SVC      D 00000000_00023010  00023000 80000000 80000001                              15:02:07.903563346 0001
                              47851000 00000000
0000 00C0 00ABBE88  SSRV    78          8595642E  1000FF50 00000248 00ABB2C0  Getmain                     15:02:07.903571458 0001
                                                  00C00000
0000 00C0 00ABBE88  SSRV    78          85956484  1000FF70 00001310 7F5FDCF0  Getmain                     15:02:07.903608276 0001
                                                  00C00000
0000 00C0 00ABBE88  SSRV    78          81485066  0000E540 00000150 7F5F3EB0  Getmain                     15:02:07.904103985 0001
                                                  00C00000
0000 00C0 00ABBE88  PC      ...   0     0148524A            00506          SysTrace Snapshot
0000 00C0 00ABBE88  SSRV    78          90B05874  1000FF70 00000358 7FF70CA8  Getmain                     15:02:07.904157117 0001
                                                  00040000
0000 00C0 00ABBE88  EXT   TIMR 00000000_00FDFA8C  00001005                    00000001 0285C080 0004 00C0 15:02:07.904919110 0001
                              47040000 80000000                               00000000
0000 00C0 00ABBE88  SSRV    78          90B05EB2  1000E574 0005E000 7F5AB000  Getmain                     15:02:07.904927266 0001
                                                  00040000
0000 00C0 00ABBE88  PC      ...   0     10B05FEE            0090E          IarV64
0000 00C0 00ABBE88  PC      ...   0           00_01CF9EF8   0030A          LocAscb
0000 00C0 00ABBE88  PR      ...   0           00_01CF9EF8 0148036C                    0004
0000 00C0 00ABBE88  SSRV    14B         01A06000  00000000 00000000 00000000  IarV64   GetStor            15:02:07.905119914 0001
                                                  00000050_01E00000
                                                  00000000_0000001E
                                                  7F5AB000_00000000
0000 00C0 00ABBE88  PR      ...   0     10B05FEE 01D0CDAE                           0004
0000 00C0 00ABBE88  PC      ...   0     10B060F6            0090E          IarV64
0000 00C0 00ABBE88  EXT   CLKC 00000000_00FDFB68  00001004 00FD5AD8     0000  00000001 0285C080 0004 0004 15:02:07.907167971 0001
                              47040000 80000000                               00000000
```

Unhandled TCB U0001 Breakdown

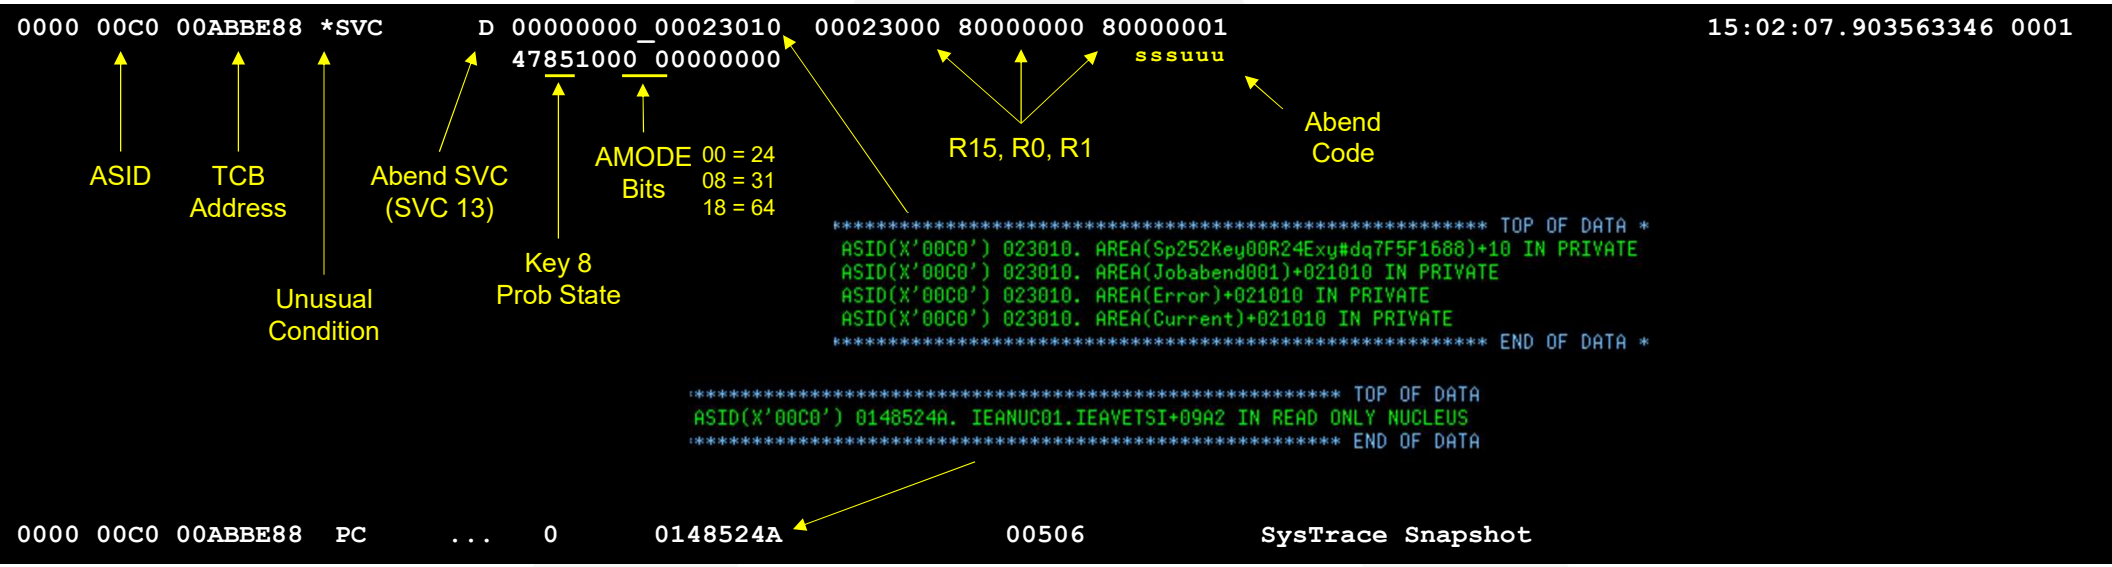

- SysTrace Snapshot occurs when SNAPTRC macro is issued by RTM.
- This is a dead giveaway that a dump is being taken.
- If too many SNAPTRC macros are issued in close succession, a so-called “mini-snapshot” (only 64K per processor) is taken to avoid affecting overall system performance.

U0001 Breakdown

- Although IP WHERE was not super helpful identifying which program issued the ABEND macro, IP SUMMARY FORMAT CURRENT REGS has:

```
PRB: 00AFB058
   -0020  XSB...... 7FFFDC10  FLAGS2... 00       RTPSW1... 00000000  00000000          RTPSW2... 00000000  00000000
   -0008  FLAGS1... 00000000  WLIC..... 0002000D
   +0000  RSV...... 00000000  00000000           SZSTAB... 00110082  CDE...... 00AFE970  OPSW..... 478D1000  00023010
   +0018  SQE...... 00000000  LINK..... 00ABBE88

          64-Bit GPRs from the RB/XSB
   Left halves of all registers contain zeros
    0-3   00000064  00022000  00000040  00AD9D64
    4-7   00AD9D40  00AFB1A0  00AC1FC8  00FA2E80
    8-11  00AFD608  00AFB4A8  01DB0200  08000003
   12-15  05404098  00022008  00AFD654  00AFD638
   +0060  RSV...... C1C2C5D5  C4F0F0F1

   EP....... ABEND001
   ENTPT.... 00000000  00023000           RRBP..... 00AFB058  USE...... 0001      SP....... FC
   Reenterable. Reusable.
   APF library.
   NRFAC.... 00000001  MSBAD.... 00000000  00023000  LNTH..... 00000010
   NAMEL.... 0008      ASID..... 00C0      PROVIDI.. 00000002  PROVIDD.. 80000000  000EF900  7FF73BE4  7F5F4E60
   EPTOKEN.. 0000005B  00C00041
```

Interrupt Code

TCB Address

Old PSW Address (64-bit)

Unhandled TCB Program Check (0C1)

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

```
0000 00C0 00ABBE88  SVCR  FF00 00000000_00031000  00031000 00000064 00030000                                21:39:15.458348656 0001
                              47850000 00000000
0000 00C0 00ABBE88 *PGM    001 00000000_00031002  00020001 00000000          00000000 00000000 00C0 00C0 21:39:15.458361011 0001
                              47850000 00000000           7F5B2800          00000000
0000 00C0 00ABBE88 *RCVY  PROG                    940C1000 00000001 00000000 00000000 00000000 00C0 00C0 21:39:15.458594509 0001
                                                                             00000000
0000 00C0 00ABBE88  SSRV   12D          8120B1A0  00ABBE88 000C8000 FF3A0000 Status   NDETcb Reset       21:39:15.458684652 0001
                                                  00000000
0000 00C0 00ABBE88  SSRV   12D          8120B1C0  00ABBE88 000B8000 00000000 Status   SDETcb Reset       21:39:15.458692758 0001
                                                  00000000
0000 00C0 00ABBE88  DSP          00000000_0120DA68  00000000 00000064 00030000  00000000 00000000 00C0 00C0 21:39:15.458709212 0001
                              47850000 80000000
0000 00C0 00ABBE88 *SVC      D 00000000_0120DA6A  00031000 00000064 00030000                                21:39:15.458713126 0001
                              47850000 80000000
0000 00C0 00ABBE88  SSRV    78          8595642E  1000FF50 00000248 00ABB2C0 Getmain                        21:39:15.458721473 0001
                                                  00C00000
0000 00C0 00ABBE88  SSRV    78          85956484  1000FF70 00001310 7F5AACF0 Getmain                        21:39:15.458790979 0001
                                                  00C00000
0000 00C0 00ABBE88  SSRV    78          81485066  0000E540 00000150 7F61EEB0 Getmain                        21:39:15.459257919 0001
                                                  00C00000
0000 00C0 00ABBE88  PC      ...   0     0148524A            00506          SysTrace Snapshot
0000 00C0 00ABBE88  SSRV    78          90B05874  1000FF70 00000358 7FF70CA8 Getmain                        21:39:15.459303220 0001
                                                  00040000
0000 00C0 00ABBE88  EXT   CLKC 00000000_00FDFA8C  00001004 02246080     001D 00000001 0285C080 0004 00C0 21:39:15.460085477 0001
                              47040000 80000000                              00000000
0000 00C0 00ABBE88  SSRV   110          8106B2A6  022460AC 00F9CE80 8106BB1A Schedule                       21:39:15.460093552 0001
                                                  00000000
0000 00C0 00ABBE88  SSRV   110          8106B2A6  03EC6144 00F9CE80 8106BB1A Schedule                       21:39:15.460097716 0001
                                                  00000000
```

Unhandled TCB 0C1 Breakdown

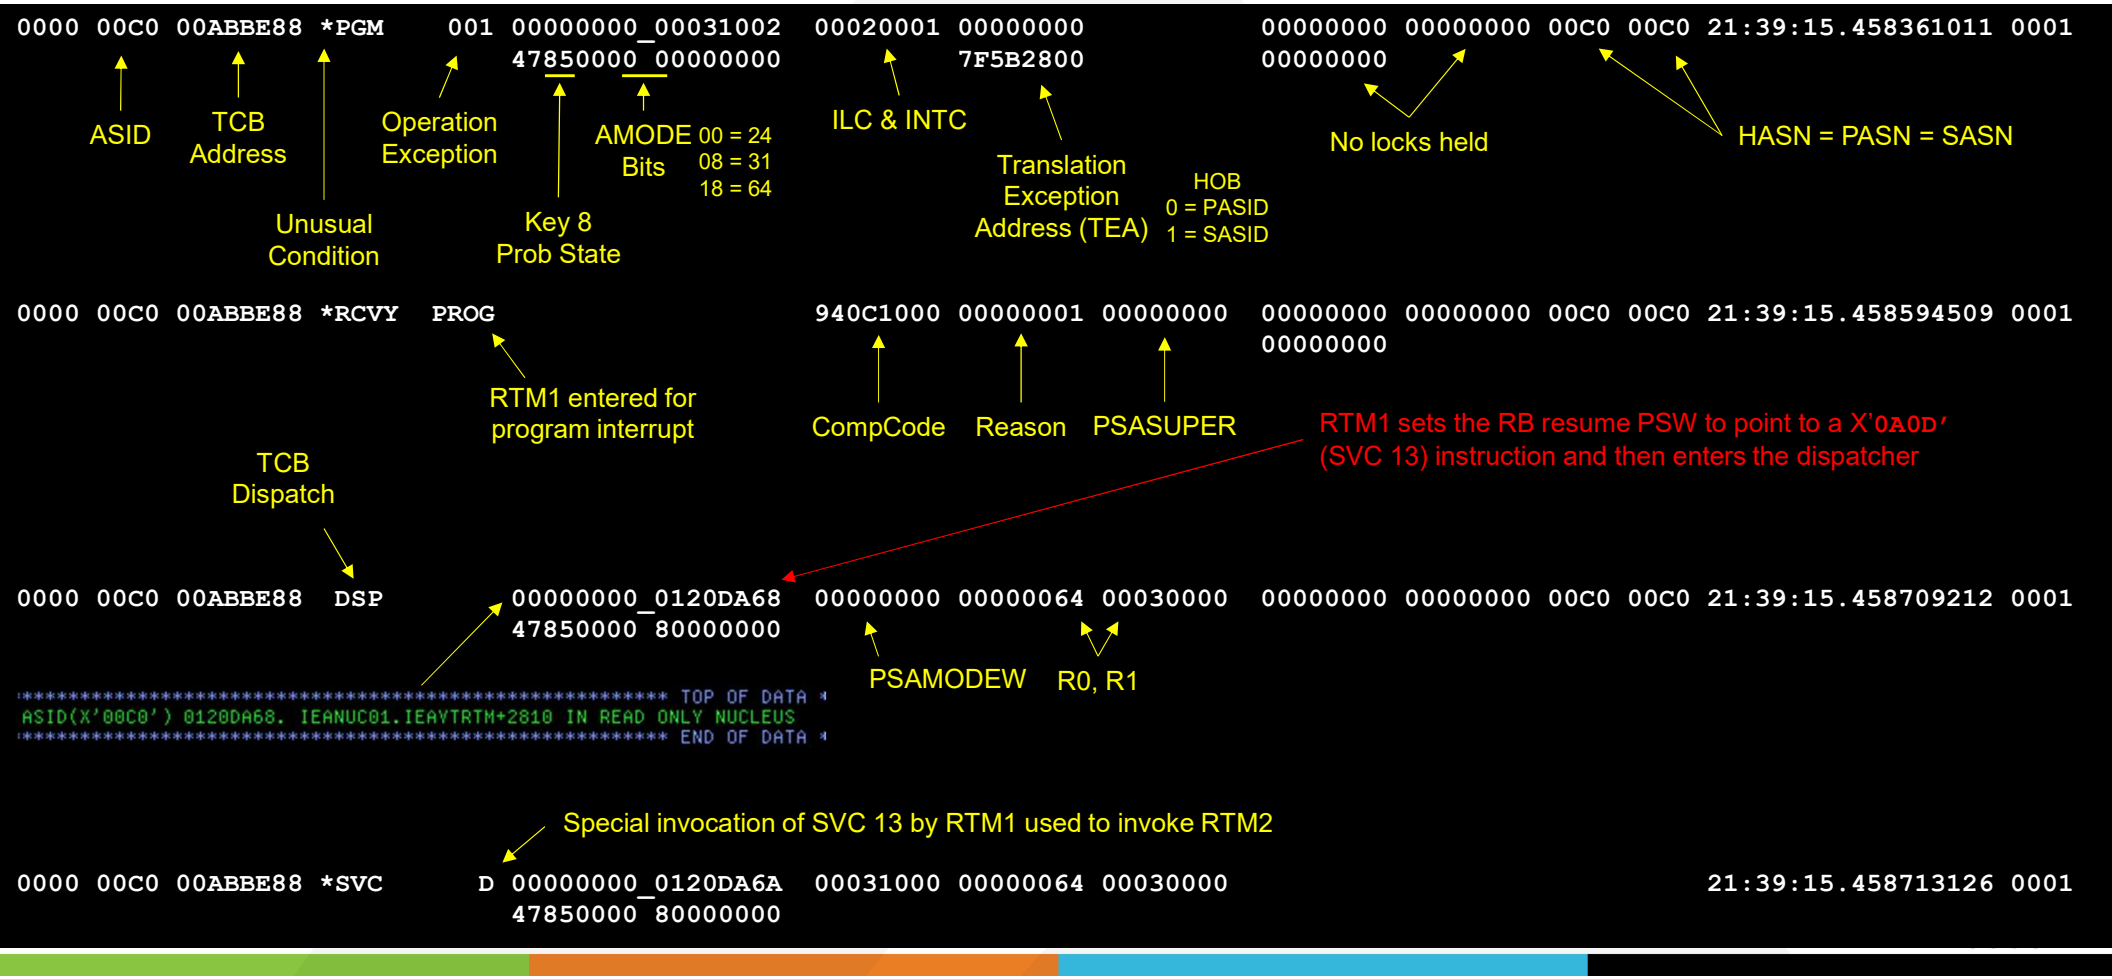

Types of RCVY Trace Entries

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

CD/D	Description
ABRT	Serious RTM failure
ABT	CALLRTM TYPE=ABTERM macro issued
ABTR	Rescheduling of a CALLRTM TYPE=ABTERM request
ESTA	Preparing to invoke ESTAE recovery routine
ESTR	Retry requested by ESTAE recovery routine
FRR	RTM1 invoking a function recovery routine (FRR)
ITRM	The system requested RTM1 to end an interrupted task
ITRR	ITRM reentry to fulfill ITRM request
MCH	RTM1 entered for a machine check interruption
MEM	CALLRTM TYPE=MEMTERM macro issued
MEMR	Abnormal memory end following a MEM event
PERC	Percolation from RTM1 to RTM2 to continue recovery
PROG	RTM1 was entered for a program check interruption

CD/D	Description
RCML	Special processing for task in a failing address space holding the local lock of another address space
RCMR	RCML reentry to process RMGR abend
RESM	Resume from an FRR after a RESTART request (RSRT)
RSRT	RTM entered for a RESTART request from the operator
RTRY	Retry from an FRR
SABN	The system requested RTM1 to abend the current UOW
SKFE	RTM bypassed FESTAE exit because its address is zero
SRB	Percolation from service request block (SRB) recovery
SRBT	CALLRTM TYPE=SRBTERM macro issued
STRM	The system requested RTM1 to abend suspended task
STRR	STRM reentry to abend suspended task

ESTAE-Protected Abend (U4000)

- Unimportant trace entries removed (Getmain, Storage, CsvQuery, etc.)
- For this slide, you've now graduated to proper "green screen" mode ☺

```
0000 00A8 00AE8298 *SVC      D 00000000_10D735D8  00000000 80000000 80000FA0                        15:09:59.822357917 0001
                               47141000 80000000
0000 00A8 00AE8298 *RCVY  ESTA          90C8F34C  7F59B6D8 00000000 7F598000  00000000 00000000 00A8 00A8 15:09:59.822825080 0001
                                                  00000000 00AFF3E0           00000000
0000 00A8 00AE8298  SVC      C 00000000_0595EBBA  0595F401 90C8F34C 7F59B6D8  Synch                      15:09:59.822829029 0001
                               47041000 80000000
0000 00A8 00AE8298  SVCR  FF00 00000000_0595F400  0595F400 90C8F34C 7F59B6D8                           15:09:59.822845536 0001
                               47040000 80000000
0000 00A8 00AE8298  DSP        00000000_0595F400  00000000 90C8F34C 7F59B6D8  00000000 00000000 00A8 00A8 15:09:59.822853673 0001
                               47040000 80000000
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
0002 00A8 00AE8298  SVC     33 00000000_10D8AEC2  00000002 7F36FDB8 7F36FBE0  Snap/Sdump                 15:09:59.823900805 0000
                               47041000 80000000
0002 00A8 00AE8298  SVCR    33 00000000_10D8AEC2  00000000 FFF00000 FF36FBE0                           15:09:59.824592603 0000
                               47041000 80000000
0002 00A8 00AE8298  SVC      1 00000000_10CCDB74  90C88000 00000001 7F36FCE0  Wait                       15:09:59.824654344 0000
                               47140000 80000000
0002 00A8 00AE8298  SVCR     1 00000000_10CCDB74  80AE3F98 00000001 7F36FCE0                           15:09:59.824654502 0000
                               47140000 80000000
******** Trace data is not available from all processors after this time.
```

Recovery exit address

Recovery exit parm

SYNCH macro creates new PRB to run exit in program's state/key

ECB address

Wait for dump capture phase to complete before proceeding

ESTAE-Protected Abend With Retry (0C4)

SHARE
EDUCATE • NETWORK • INFLUENCE

```
0000 00A5 00AB7E78  PGM    004 00000000_10D44852  00040004 00000000           00000000 00000000 00A5 00A5 11:04:17.552208707 0000
                               47141000 80000000           00000480           00000000
0000 00A5 00AB7E78 *RCVY  PROG                    940C4000 00000004 00000000  00000000 00000000 00A5 00A5 11:04:17.552515855 0000
                                                                              00000000
0000 00A5 00AB7E78  DSP        00000000_0120DD68  00000000 7F2DE8C0 7F50C800  00000000 00000000 00A5 00A5 11:04:17.552657064 0000
                               47141000 80000000
0000 00A5 00AB7E78  SVC      D 00000000_0120DD6A  7F2DE1E0 7F2DE8C0 7F50C800                             11:04:17.552657190 0000
                               47141000 80000000                               Recovery exit parm
0000 00A5 00AB7E78 *RCVY  ESTA          90CDA34C  7F4D46D8 00000000_7F4EE000  00000000 00000000 00A5 00A5 11:04:17.552995105 0000
                 Recovery exit address            00000000 00AFF2E0           00000000
0000 00A5 00AB7E78  SVC      C 00000000_0595EBBA  0595F401 90CDA34C 7F4D46D8  Synch                      11:04:17.552995191 0000
                               47041000 80000000                               SYNCH macro  creates new PRB to
0000 00A5 00AB7E78  SVCR  FF00 00000000_0595F400  0595F400 90CDA34C 7F4D46D8   run exit in program's state/key 11:04:17.553007505 0000
                               47040000 80000000
0000 00A5 00AB7E78  DSP        00000000_0595F400  00000000 90CDA34C 7F4D46D8  00000000 00000000 00A5 00A5 11:04:17.553015586 0000
                               47040000 80000000
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
0000 00A5 00AB7E78  SVC      3 00000000_00FD7C32  7F4D4970 00000000 7F4D46D8  Exit                       11:04:17.553617774 0000
                               47041000 80000000
0000 00A5 00AB7E78  SVCR     C 00000000_0595EBBA  7F4D4970 00000000 7F4D46D8                             11:04:17.553625899 0000
                               47041000 80000000
0000 00A5 00AB7E78  DSP        00000000_0595EBBA  00000000 00000000 7F4D46D8  00000000 00000000 00A5 00A5 11:04:17.553630064 0000
                               47041000 80000000           Local lock held
0000 00A5 00AB7E78 *RCVY  ESTR 00000000_90CDBBE6  90CDA34C 00AFF2E0           00000001 00000000 00A5 00A5 11:04:17.553765103 0000
                  Retry address                      Recovery exit address    00000000
0000 00A5 00AB7E78  SVCR     4 00000000_10CDBBE6  7F4D4B68 00000000 90D4487A                             11:04:17.553859406 0000
                               47141000 80000000
0000 00A5 00AB7E78  DSP        00000000_10CDBBE6  00000000 00000000 90D4487A  00000000 00000000 00A5 00A5 11:04:17.553863645 0000
                               47141000 80000000
```

Some ugliness here. This is the return from the SVRB created for RTM2. It appears as 'SVCR 4' because RTM handled a PGM 004. ☹

BLDL SVC Showing SSCH, WAIT, I/O and SRB

- GETMAIN and FREEMAIN SVC/SVCR pairs removed

```
0000 00A5 00AB7E78  SVC     12 00000000_05E0E40A  FF52FFA0 FF52FFA8 000244F0  Bldl                          11:04:17.457907286 0001
                               47041000 80000000
0000 00A5 00AB7E78  SVC      0 00000000_01262ED8  00FE30C8 01000800 00AACE98  Excp                          11:04:17.457935915 0001
                               47540000 80000000
0000 00A5 00AB7E78  SSCH 07381 00  02   01FCEC2C  02645E38 52C2F081 3FF08C60                                11:04:17.457980865 0001
                                                              00AE1E38
0000 00A5 00AB7E78  SVCR     0 00000000_01262ED8  00000000 02645E38 00AACE98                                11:04:17.457980963 0001
                               47540000 80000000
0000 00A5 00AB7E78  SVC      1 00000000_01262EEA  00000000 00000001 00AACE94  Wait                          11:04:17.457984830 0001
                               47540000 80000000
0000 00A5 00AB7E78  SVCR     1 00000000_01262EEA  80AFF800 00000001 00AACE94                                11:04:17.457984955 0001
                               47540000 80000000
0000 0001 00000000  WAIT                                                                                     11:04:17.457993043 0001
0000 0001 00000000  I/O  07381 00000000_00000000  50C04007 7E892A70 0C000000  00000080 00000000 0001 0001 11:04:17.458389583 0001
                               07060000 00000000            02645E38 00100002  00000000
0000 0001 00000000  SSRV   112          810C5302  01FCEC00 00F91B80 8129AF18  Schedule                      11:04:17.458389814 0001
                                                  00800000
0000 00A5 01F6CB00  SRB         00000000_0129AF18  000000A5 01FCEC00 01FCEC2C        00          00A5 00A5 11:04:17.458390316 0001
                               47040000 80000000  00AB7E78 80
0000 00A5 01F6CB00  SSRV     2          80FF55DA  00AACE94 7F000000 00000000  Post                          11:04:17.458402617 0001
                                                  00000000
0000 00A5 00AB7E78  DSP         00000000_01262EEA  00000000 00000001 00AACE94  00000000 00000000 00A5 00A5 11:04:17.458410891 0001
                               47540000 80000000
0000 00A5 00AB7E78  SVCR    12 00000000_05E0E40A  00000004 00000000 00AACC98                                11:04:17.458423270 0001
                               47041000 80000000
```

Use the Trace to Help Find Root Cause

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Sometimes a dump is produced only after two or more back-to-back abends. The trace can help you find root cause.
- Don't focus on the last abend in the dump, focus on the first one.
- Find the following in the trace:

```
******** Trace data is not available from all processors before this time.
```

- Starting from this point, search for unusual conditions represented by an asterisk (e.g., *PGM, *RCVY, etc.)
- Generally you can ignore I/O errors, they are most-likely not yours.
- In most cases, the first and only problem you find will be the abend being debugged. But, every once in a while, you discover something unexpected (perhaps an earlier abend already recovered from), that is responsible in some way for the current problem.

Another Favorite Ultra-Simple Test Program

- // EXEC PGM=ABEND322 to abend with S322 (aka IEFBR15)

```
ABEND322 CSECT ,
         BR    15
         END   ,
```

```
0000 00A5 00ABBE88  DSP        00000000_0001C000  00000000 00000064 0001B000  00000000 00000000 00A5 00A5 22:13:06.477408232 0000
                               47850000 00000000
0000 00A5 00ABBE88  EXT    WTI 00000000_0001C000  00001007                    00000000 00000000 00A5 00A5 22:13:06.477609064 0000
                               47850000 00000000                              00000000
0000 00A5 00ABBE88  DSP        00000000_0001C000  00000000 00000064 0001B000  00000000 00000000 00A5 00A5 22:13:06.477615942 0001
                               47850000 00000000
0000 00A5 00ABBE88  EXT   TIMR 00000000_0001C000  00001005                    00000000 00000000 00A5 00A5 22:13:06.478884364 0001
                               47850000 00000000                              00000000
0000 00A5 00ABBE88  EXT   CLKC 00000000_0001C000  00001004 03EBB230     0072  00000000 00000000 00A5 00A5 22:13:06.478998950 0001
                               47850000 00000000                              00000000
0000 00A5 00ABBE88  SSRV   110          8106A606  03EBB25C 00FBAB80 8106AE7A  Schedule                    22:13:06.479007091 0001
                                                  000000C2
0002 00A5 00ABBE88  DSP        00000000_0001C000  00000000 00000064 0001B000  00000000 00000000 00A5 00A5 22:13:06.479017945 0002
                               47850000 00000000
0002 00A5 00ABBE88  EXT   TIMR 00000000_0001C000  00001005                    00000000 00000000 00A5 00A5 22:13:06.480291691 0002
                               47850000 00000000                              00000000
0000 00A5 00ABBE88  DSP        00000000_0001C000  00000000 00000064 0001B000  00000000 00000000 00A5 00A5 22:13:06.480324578 0000
                               47850000 00000000
0000 00A5 00ABBE88  I/O  0F962 00000000_0001C000  6080C0C9 3C5F52E8 00800007  00000080 00000000 00A5 00A5 22:13:06.480586823 0000
                               47850000 00000000           026D9D08 00000000  00000000
```

Debugging a CPU Loop

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- All you see are external and I/O interrupts. Nothing of any real value?
- My example shows the tightest possible loop (all PSW NSIs are 0001C000), but most real-world loops are wider in scope.
- Look through a couple/few pages of these interrupts and jot down the PSW NSIs. You'll usually begin to see a pattern emerge.
- Issue IP WHERE for each unique value. If you don't get a name you recognize from IP WHERE, browse the location in storage and back up looking for an eyecatcher (often a copyright statement).
- Addresses in Not-My-Code represent other product or component code that gets control as a result of the application loop. Helpful clues!
- Addresses in application code are the most valuable. A sufficiently large sample of such addresses, coupled with the rest of the information in the dump as well as a product listing, is usually enough to solve the issue.

Branch and/or Mode Tracing

- Extraordinarily difficult or perplexing problems will sometimes require a recreate with branch or mode tracing enabled.
  - Full-disclosure, I have never solved any issue through the use of mode tracing, but branch tracing has been my savior many times over the years!
- These options are enabled by issuing the TRACE command.

```
D TRACE
IEE843I 08.39.15  TRACE DISPLAY 378
        SYSTEM STATUS INFORMATION
 ST=(ON,0006M,00030M) AS=ON   BR=OFF EX=ON   MO=OFF MT=(ON,064K)

TRACE ST,BR=ON
IEE839I ST=(ON,0006M,00030M) AS=ON   BR=ON   EX=ON   MO=OFF MT=(ON,064K) 416
```

- There is a performance penalty (~10%?) so it's a good idea to disable branch/mode tracing when no longer needed.
- **NOTE:** It's fully supported when z/OS runs as a guest under z/VM.

Branch Tracing Example

- BR trace line contains ten entries for 24- and 31-bit mode programs
- Entries are longer when the code is 64-bit (e.g., 00_12345678)
- Each entry represents the target of a BALR, BASR, BASSM or BAKR

```
0000 008D 00AB7408  SVCR    78 00000000_083A7F78  00000000 00002380 0008F000                          13:57:44.520616198 0002
                               47852000 80000000
0000 008D 00AB7408  BR         08366678 083656F8 083656F8 08366398 083A2688 083A1EC0 083A24E8 08366678 083656F8 083656F8
                    BR         08366398 083A2688 083A1EC0 083A24E8 08366678 083656F8 083656F8 08366398 083A2688 083A1EC0
                    BR         083A24E8 0838BA78 083656F8 083DDCE0 083DF978 083E1DA0 083991A0 0839D6E8 083A1EC0 083A24E8
                    BR         083A2688 083A1EC0 083A24E8 0839BC18 0839BC18 08369CF0 083656F8 08366238 0839D6E8 083A1EC0
                    BR         083A24E8 083656F8 083656F8 08369CF0 083656F8 08366238 0839D6E8 083A1EC0 083A24E8 083656F8
                    BR         083656F8 08366678 083656F8 083656F8 08366238 0839D6E8 083A1EC0 083A24E8 083656F8 08366398
                    BR         083A2688 083A1EC0 083A24E8 083A6808 08369CF0 083656F8 08366238 0839D6E8 083A1EC0 083A24E8
                    BR         083656F8 083656F8 08369CF0 083656F8 08366238 0839D6E8 083A1EC0 083A24E8 083656F8 083656F8
                    BR         08392928 083A7EC8
0000 008D 00AB7408  SVC     78 00000000_083A7F78  00000003 0000002F 10C933C8  Freemain                 13:57:44.520894730 0002
                               47852000 80000000
0000 008D 00AB7408  SVCR    78 00000000_083A7F78  00000000 0000002F 10C933C8                          13:57:44.520902886 0002
                               47852000 80000000
0000 008D 00AB7408  SVC     78 00000000_083DDCC2  00000003 00000060 10C94000  Freemain                 13:57:44.520911172 0002
                               47852000 80000000
0000 008D 00AB7408  BR           FEA008
0000 008D 00AB7408  SVCR    78 00000000_083DDCC2  00000000 00000060 10C94000                          13:57:44.520931533 0002
                               47852000 80000000
```

Mode Tracing Example

- The MODE trace displays only one mode change per line
- All entries describe a change to or from 64-bit mode only – no 24/31

```
0006 008D 03AAF700  MODE    ... 24 OR 31       00_014B0832
0006 008D 03AAF700  MODE    ... 64       014B085C
0006 008D 03AAF700  MODE    ... 24 OR 31       00_014B06B6
0006 008D 03AAF700  MODE    ... 64       014B06B0
0006 008D 03AAF700  MODE    ... 24 OR 31       00_014B0832
0006 008D 03AAF700  MODE    ... 64       014B085C
0006 008D 03AAF700  MODE    ... 24 OR 31       00_014B06B6
0006 008D 03AAF700  MODE    ... 64       014B06B0
0006 008D 03AAF700  MODE    ... 24 OR 31       00_014B0832
0006 008D 03AAF700  MODE    ... 64       014B085C
0006 008D 03AAF700  MODE    ... 24 OR 31       00_014B06B6
0006 008D 03AAF700  MODE    ... 64       01AA353E
0006 008D 03AAF700  MODE    ... 24 OR 31       00_01AA3588
0006 008D 03AAF700  MODE    ... 64       01AA35B6
0006 008D 03AAF700  MODE    ... 24 OR 31       00_01AA364A
0006 008D 03AAF700  MODE    ... 64       01472134
0006 008D 03AAF700  MODE    ... 24 OR 31       00_01AA38AE
0006 008D 03AAF700  MODE    ... 64       01AA38B6
0006 008D 03AAF700  MODE    ... 24 OR 31       00_01AA3AF2
0006 008D 03AAF700  MODE    ... 64       01472134
0006 008D 03AAF700  MODE    ... 24 OR 31       00_01AA3B46
0006 008D 03AAF700  MODE    ... 64       01AA3B50
0006 008D 03AAF700  MODE    ... 24 OR 31       00_01AA3C0C
```

Branch and Mode Tracing Example

- BR, MODE, and MOBR (BASSM) entries all together

```
0000 008D 00AA64F8  PC      ...   0      052A226A              0041C          Console
0000 008D 00AA64F8  BR            FF7C08 066B82E8   FF7FFA
0000 008D 00AA64F8  PR      ...   0      052A226A 10BA5166                                 008D
0000 008D 00AA64F8  BR          0148EE50 013EA5E8
0000 008D 00AA64F8  PR      ...   0      1122C492 052A0C9E                                 008D
0000 008D 00AA64F8  MODE    ... 64       111B65DE
0000 008D 00AA64F8  MOBR    ... 31       111B6000
0000 008D 00AA64F8  BR          1122F000
0000 008D 00AA64F8  MODE    ... 64       111B65DE
0000 008D 00AA64F8  BR          11181000
0000 008D 00AA64F8  MOBR    ... 31       1114D000
0000 008D 00AA64F8  MOBR    ... 64             00_11181000
0000 008D 00AA64F8  MOBR    ... 31       1117EFD0
0000 008D 00AA64F8  PC      ...   1      1117F07C              00318          Resume   SRB
0000 008D 00AA64F8  BR            FE9E5E 01483500
0000 008D 00AA64F8  SSRV   119            9117EE46  01FB3720 8003927C 04089E40  Resume                    14:42:03.500014386 0001
                                                    00000000
0000 008D 00AA64F8  BR          0113FE7A   FDE968 01125A18   FE9E5E   FE9EB2   FDEE30   FE9EB2
0000 008D 00AA64F8  PR      ...   0      1117F07C 01498D84                                 008D
0000 008D 00AA64F8  BR          11154000
0000 008D 00AA64F8  SVC       1 00000000_111A8B74  91154000 00000001 7F3F436C  Wait                      14:42:03.500021402 0001
                                47140000 80000000
0000 008D 00AA64F8  SVCR      1 00000000_111A8B74  80AA63B8 00000001 7F3F436C                             14:42:03.500021688 0001
                                47140000 80000000
0000 008D 00AA64F8  BR            FFE3E4 01483680
```

Creating Your Own Trace Entries

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Special debugging scenarios might necessitate creation of your own trace entries. The PTRACE macro is used for this purpose.
- It is available to supervisor state, key 0 programs only.
- You can be disabled; you can hold any locks; you can be in a TCB or an SRB. The only restrictions are DAT-ON and primary ASC mode and you must provide a save area (72 or 144 bytes depending on AMODE).
- A hex digit (0-F) is appended to USR to make the trace identifier.
- Either the TRACE or TRACG instruction may be generated.
- You can record the contents of up to 11 consecutive registers or up to 1024 words (4K) of storage. Each trace entry documents up to 24 bytes, so one PTRACE call can result in multiple trace entries

Example PTRACE Output


```
151678           PTRACE TYPE=USR3,REGS=(R3,R12),SAVEAREA=STANDARD
151681+          LA     0,X'3'                    .INDICATE USER ENTRY NUMBER    01-PTRACE
151682+          LA     1,(R3*X'10')+R12          .LOAD REG PAIR NUMBERS         01-PTRACE
151683+          L      15,PSATRVT-PSA(0,0)       .GET TRVT ADDRESS              01-PTRACE
151684+          O      1,TRVTRPMK-TRVT(0,15)     .INDICATE REG PAIR GIVEN       01-PTRACE
151685+          L      15,TRVTETUR-TRVT(0,15)    .GET ROUTINE EP ADDRESS        01-PTRACE
151686+          BASSM 14,15                      .CALL PTRACE ROUTINE           01-PTRACE
```

```
0006 0051 02B2F400  USR3              90E32868  6166 000 7F47700C 7F1BBDD8              0051 0051 18:03:56.945423860 0003
                                                00000007 10E328AA 7F1EAE30
0006 0051 02B2F400  USR3              90E32868  6166 014 7F2C0500 7F472500              0051 0051 18:03:56.945423887 0003
                                                7F470000 10E1D000 10E3261A
0006 0051 02B2F400  USR3              90E32868  6167 000 7F1BB192 7F1BBDD8              0051 0051 18:03:56.945425001 0003
                                                00000006 10E328AA 7F1EB080
0006 0051 02B2F400  USR3              90E32868  6167 014 7F2C0500 7F472500              0051 0051 18:03:56.945425021 0003
                                                7F470000 10E1D000 10E3261A
0006 0051 02B2F400  USR3              90E32868  6168 000 7F47700C 7F1BBDD8              0051 0051 18:03:56.945426424 0003
                                                00000005 10E328AA 7F1EB2D0
0006 0051 02B2F400  USR3              90E32868  6168 014 7F2C0500 7F472500              0051 0051 18:03:56.945426449 0003
                                                7F470000 10E1D000 10E3261A
0006 0051 02B2F400  USR3              90E32868  6169 000 7F4771AA 7F1BBDD8              0051 0051 18:03:56.945428992 0003
                                                00000004 10E328AA 7F1EB520
0006 0051 02B2F400  USR3              90E32868  6169 014 7F2C0500 7F472500              0051 0051 18:03:56.945429019 0003
                                                7F470000 10E1D000 10E3261A
0006 0051 02B2F400  USR3              90E32868  616A 000 7F4771AA 7F1BBDD8              0051 0051 18:03:56.945433489 0003
                                                00000003 10E328AA 7F1EB770
0006 0051 02B2F400  USR3              90E32868  616A 014 7F2C0500 7F472500              0051 0051 18:03:56.945433517 0003
                                                7F470000 10E1D000 10E3261A
```

Timing Duration of Events

- Many trace entries have a time stamp.
  - 18:03:56.945423887 (milliseconds, microseconds, nanoseconds)
- You can calculate the elapsed clock time of events that occur between two time stamps.
  - Be mindful that clock time is not necessarily the same as CPU time.
  - By displaying the trace for all ASIDs, you can get an idea if your CP was in control the whole time between two time stamps or if it gave up control to other MVS work or even another LPAR (if you see WTI).
- You can use PTRACE to time events that occur without any intervening time stamps.
- For example, synchronous coupling facility service calls do not generate system trace events of any kind.

Timing an IXLCACHE WRITE_DATALIST

- PTRACE USR0 and USR1 wrapped around the IXLCACHE call
- Followed by a DC H'0' to generate an PIC 001 (0C1) abend

```
0006 009A 03C28C80  USR0       00000000_1198347E  EE9B 000 00000000 00000000                    0014 009A 05:25:47.440063500 0003
                                                  00000050 2140F400 00000000
                                                  7F336628 00000000 00000000
                                                  00000000 7F070328 00000050
                                                  18301000 00000000 7F5AB000
0006 009A 03C28C80  USR0       00000000_1198347E  EE9B 038 00000000 00000000                    0014 009A 05:25:47.440063529 0003
                                                  00000050 1B302000 00000050
                                                  45E0100C 00000000 11983300
0006 009A 03C28C80  USR1       00000000_1198359E  EE9C 000 00000000 00000000                    0014 009A 05:25:47.440097105 0003
                                                  00000050 2140F400 00000000
                                                  7F336628 00000000 00000000
                                                  00000000 7F070328 00000050
                                                  18301000 00000000 7F5AB000
0006 009A 03C28C80  USR1       00000000_1198359E  EE9C 038 00000000 7F337650                    0014 009A 05:25:47.440097137 0003
                                                  00000050 1B302000 00000050
                                                  45E0100C 00000000 11983300
0006 009A 03C28C80 *PGM    001 00000000_119835A4  00020001 00000000           00000000 00000000 0014 009A 05:25:47.440097628 0003
                               47740001 80000000           7F46D800           00000000
0006 009A 03C28C80 *RCVY   PROG                   940C1000 00000001 00000000  00000000 00000000 0014 009A 05:25:47.440117985 0003
                                                                              00000000
0006 009A 03C28C80  SSRV    78          814AA5A2  4000EF50 00000A28 00F99180  Getmain                    05:25:47.440120972 0003
                                                  00010000
0006 009A 03C28C80 *RCVY   FRR  470C0000 90D6E000  940C1000 00000001 00000000  00000000 00000000 0014 009A 05:25:47.440133300 0003
                                                                    00000002  00000000
```

Subtracting these two values shows it took ~34 µsecs

Maximizing First-Failure Problem Diagnosis

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

- Size Does Matter!
  - Trace tables that are too small are an impediment to successful diagnosis
  - You may specify trace table buffer sizes between 1M and 9G per processor
  - The default buffer size is 1M per processor (the minimum allowed)
  - It's up to you to size them properly (keep in mind they occupy fixed storage)
- Content Matters Too!
  - Several DIAGxx TRAPS affects system trace processing
  - IeaRISignlTrace – enable tracing of RISIGNL requests
  - IeaRPSignlTrace – enable tracing of RPSIGNL requests
  - IeaRtm2SnapX22 – causes RTM2 to always create a system trace snapshot as it did prior to the z/OS 1.12 trace snapshot avoidance enhancement
  - IeaScheduleTrace – enable tracing of SCHEDULE and IEAMSCHD ☺
  - IeaSysTrcNoLimit – overrides limits on max size traces specified by TRACE ST command

Recommended!

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

Questions?

Your feedback is important!

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

Submit a session evaluation for each session you attend:

www.share.org/evaluation