



Session 8926

z/OS Debugging: a review of System Trace



SHARE
Technology • Connections • Results

MVS Core Technologies Project – March 3rd 2011

Evan Haruta Jerry Ng

IBM Poughkeepsie

haruta@us.ibm.com jerryng@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

- MVS
- OS/390®
- z/Architecture®
- z/OS®

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- **What is the System Trace Table?**
- **How do you format it in a dump?**
- **How do you read the system trace entries?**
- **A recent enhancement**

What is the MVS System Trace Table?

- Consists of trace buffers (one per CPU) residing in the Trace Address Space (ASID 4)
- System trace entries are inserted continuously into the trace buffers
- Contains detailed system activity
- Default size of 1 Meg per CPU
- Entries from trace buffers are merged in a dump

Why look at the System Trace?

- It contains the sequence of system events leading up to the problem
- It contains historical information that can be very useful for debugging
- It helps you to identify the source of the problem or confirm a failing scenario

How are system trace entries generated?

- **By hardware, when certain instructions are issued:**
 - Cross Memory instructions such as PC, PT, PR and SSAR
 - A branch is taken via a Branch and Link instruction such as BALR, BASR, BASSM, BAKR.... (if Branch Tracing is active)
 - Instructions causing switch of AMODE between 24/31 and 64 such as SAM (if Mode Tracing is active)

- **By the operating system, when the TRACE/TRACG instruction is issued by system component modules**

How can the System Trace options be changed?

- **Via the TRACE operator command:**
 - Stop or start system trace
 - Turn on Branch Tracing or Mode Tracing (or both)
 - Change the size of the trace table

- **Note: System Trace is started during IPL with Branch and Mode Tracing off**

How do I find the System Trace Table when debugging?

- The System Trace Table is dumped in SVC and Standalone dumps
- Use the IPCS subcommand SYSTRACE to format the System Trace Table in a dump
- SYSTRACE is documented in:
 - z/OS MVS IPCS Commands

Frequently used SYSTRACE parameters

■ Current

- Select entries of current address space (this is the default)

■ All

- Select all entries

■ ASID(x,y....)

- Select entries of specific address space(s) by ASID

■ Jobname(name1,name2.....)

- Select entries of specific address space(s) by jobname

■ TCB tcb_address)

- Select entries of a TCB in an address space

More SYSTRACE parameters...

■ CPU

- Select entries of a specific CPU

■ Time(z)

- z = hex, local or GMT (default is hex)

■ Exclude(x)

- x = BR or MODE or BR,MODE
- Exclude branch or mode trace entries, or both

Examples of using SYSTRACE

- **SYSTRACE**
- **SYSTRACE ALL TIME(LOCAL)**
- **SYSTRACE ASID(X'5C')**
- **SYSTRACE ALL EX(BR,MODE) TIME(LOCAL)**
- **SYSTRACE ASID (X'1F') TCB(X'9D2380') TIME(LOCAL)**
- **SYSTRACE CPU(6) ALL**

Reading the System Trace

```

IPCS OUTPUT STREAM -----
Command ==>
***** TOP OF DATA *****

----- SYSTEM TRACE TABLE -----
--
--
PR ASID WU-ADDR- IDENT  CD/D PSW----- ADDRESS-  UNIQUE-1 UNIQUE-2 UNIQUE-3  PSACLHS- PSALOCAL PASD SASD TIME
                UNIQUE-4 UNIQUE-5 UNIQUE-6  PSACLHSE

00 0021 009FE030  DSP      070C1000 814C9DDE 00000000 814C9DDE 00FB3788 00000000 00000000 0021 0021 23:0
00 0021 009FE030  SVC      4F 070C1000 814CA396 009FDCB0 00000011 FFFFFFFEF Status Start SRBs only 23:0
00 0021 009FE030  SVCR     4F 070C1000 814CA396 00000000 00000000 05616F00                23:0
02 0165 009FFB00  DSP      070C4000 8137DA1C 00800000 00000001 05882000 00000000 00000000 0049 0165 23:0
02 0165 009FFB00  PR      ... 0      09570126 29A0530E                0165
02 0165 009FFB00  PC      ... 0      095701CE                01F01
02 0165 009FFB00  PR      ... 0      095701CE 29A06B5A                0165
00 0010 009CEAD0  SSRV    132      00000000 3750E5C0 00004000 7E0EC000 Storage Obtain 23:0
                00100000
  
```

- Entries are presented chronologically and intermixed
- Oldest entry at the top
- Newest entry at the bottom

Reading the System Trace: *who did it?*

```

IPCS OUTPUT STREAM -----
Command ==>
***** TOP OF DATA *****

----- SYSTEM TRACE TABLE -----
--
--
PR ASID WU-ADDR- IDENT  CD/D PSW----- ADDRESS-  UNIQUE-1 UNIQUE-2 UNIQUE-3  PSACLHS- PSALOCAL PASD SASD TIME
                               UNIQUE-4 UNIQUE-5 UNIQUE-6  PSACLHSE

00 0021 009FE030  DSP           070C1000 814C9DDE 00000000 814C9DDE 00FB3788 00000000 00000000 0021 0021 23:0
00 0021 009FE030  SVC           4F 070C1000 814CA396 009FDCB0 00000011 FFFFFFFEF Status Start SRBs only 23:0
00 0021 009FE030  SVCR          4F 070C1000 814CA396 00000000 00000000 05616F00                23:0
02 0165 009FFB00  DSP           070C4000 8137DA1C 00800000 00000001 05882000 00000000 00000000 0049 0165 23:0
02 0165 009FFB00  PR            ... 0         09570126 29A0530E                0165
02 0165 009FFB00  PC            ... 0         095701CE                01F01
02 0165 009FFB00  PR            ... 0         095701CE 29A06B5A                0165
00 0010 009CEAD0  SSRV          132         00000000 3750E5C0 00004000 7E0EC000 Storage Obtain                23:0
                               00100000
  
```

- PR (processor id)
- ASID
- WU-ADDR (TCB mode: TCB address / SRB mode: WEB address or zero)
- PSW

Reading the System Trace: *what was done?*

```

IPCS OUTPUT STREAM -----
Command ==>
***** TOP OF DATA *****

----- SYSTEM TRACE TABLE -----
--
--
PR ASID WU-ADDR- IDENT  CD/D PSW----- ADDRESS-  UNIQUE-1  UNIQUE-2  UNIQUE-3  PSACLHS-  PSALOCAL  PASD  SASD  TIME
                                UNIQUE-4  UNIQUE-5  UNIQUE-6  PSACLHSE

00 0021 009FE030  DSP           070C1000 814C9DDE 00000000 814C9DDE 00FB3788 00000000 00000000 0021 0021 23:0
00 0021 009FE030  SVC           4F 070C1000 814CA396 009FDCB0 00000011 FFFFFFFEF Status Start SRBs only 23:0
00 0021 009FE030  SVCR          4F 070C1000 814CA396 00000000 00000000 05616F00 23:0
02 0165 009FFB00  DSP           070C4000 8137DA1C 00800000 00000001 05882000 00000000 00000000 0049 0165 23:0
02 0165 009FFB00  PR            ... 0          09570126 29A0530E 01F01 0165
02 0165 009FFB00  PC            ... 0          095701CE 29A06B5A 0165
02 0165 009FFB00  PR            ... 0          095701CE 29A06B5A 0165
00 0010 009CEAD0  SSRV          132          00000000 3750E5C0 00004000 7E0EC000 Storage Obtain 23:0
                                00100000
  
```

- IDENT - trace table entry identifier
- CD/D – a number describing the entry
- SVC descriptors

Reading the System Trace: *how was it done?*

```

IPCS OUTPUT STREAM -----
Command ==>
***** TOP OF DATA *****
----- SYSTEM TRACE TABLE -----
--
--
PR ASID WU-ADDR- IDENT  CD/D PSW----- ADDRESS-  UNIQUE-1 UNIQUE-2 UNIQUE-3 PSACLHS- PSALOCAL PASD SASD TIME
                UNIQUE-4 UNIQUE-5 UNIQUE-6 PSACLHSE
00 0021 009FE030  DSP      070C1000 814C9DDE 00000000 814C9DDE 00FB3788 00000000 00000000 0021 0021 23:0
00 0021 009FE030  SVC      4F 070C1000 814CA396 009FDCB0 00000011 FFFFFFFEF Status Start SRBs only 23:0
00 0021 009FE030  SVCR     4F 070C1000 814CA396 00000000 00000000 05616F00 00000000 00000000 23:0
02 0165 009FFB00  DSP      070C4000 8137DA1C 00800000 00000001 05882000 00000000 00000000 0049 0165 23:0
02 0165 009FFB00  PR      ... 0 09570126 29A0530E 01F01 0165
02 0165 009FFB00  PC      ... 0 095701CE 29A06B5A 0165
02 0165 009FFB00  PR      ... 0 095701CE 29A06B5A 0165
00 0010 009CEAD0  SSRV    132 00000000 3750E5C0 00004000 7E0EC000 Storage Obtain 23:0
                00100000

```

- **UNIQUE** – up to 6 unique fields containing information related to the entry

Reading the System Trace: *under what conditions?*

```

IPCS OUTPUT STREAM -----
Command ==>
***** TOP OF DATA *****
----- SYSTEM TRACE TABLE -----
--
--
PR ASID WU-ADDR- IDENT  CD/D PSW----- ADDRESS-  UNIQUE-1  UNIQUE-2  UNIQUE-3  PSACLHS-  PSALOCAL  PASD  SASD  TIME
                                UNIQUE-4  UNIQUE-5  UNIQUE-6  PSACLHSE
00 0021 009FE030  DSP      070C1000 814C9DDE 00000000 814C9DDE 00FB3788 00000000 00000000 0021 0021 23:0
00 0021 009FE030  SVC      4F 070C1000 814CA396 009FDCB0 00000011 FFFFFFFEF Status   Start SRBs only 23:0
00 0021 009FE030  SVCR     4F 070C1000 814CA396 00000000 00000000 05616F00                23:0
02 0165 009FFB00  DSP      070C4000 8137DA1C 00800000 00000001 05882000 00000000 00000000 0049 0165 23:0
02 0165 009FFB00  PR      ...  0      09570126 29A0530E                0165
02 0165 009FFB00  PC      ...  0      095701CE                01F01
02 0165 009FFB00  PR      ...  0      095701CE 29A06B5A                0165
00 0010 009CEAD0  SSRV    132      00000000 3750E5C0 00004000 7E0EC000 Storage Obtain                23:0
                                00100000
  
```

- PSACLHS / PSACLHSE / PSALOCAL – local lock information
- PASD / SASD - Primary and Secondary ASID
- TIME – timestamp

Note: the rightmost column is CP (not shown above). It is the physical CPU number which is usually not useful for common debugging purposes

Common System Trace entries

■ Dispatch

- DSP (TCB dispatch)
- SRB (SRB dispatch)
- SSRB (Suspended SRB dispatch)
- WAIT (no-work wait dispatch)

■ Interrupts

- SVC, SVCR, SVCE (SVC issuance, SVC completion, SVC error)
- I/O
- CLKC, EMS, EXT, CALL, SS (external interrupts such as clock comparator, emergency signal, general external, external call and service signal)
- PGM (program check)
- MCH (machine check)
- RST (restart interrupt)

- See **z/OS MVS Diagnosis: Tools and Service Aids** manual, chapter on System Trace for detailed description of system trace entries

Common System Trace entries...

- **I/O operations**
 - SSCH (start subchannel)
 - MSCH (modify subchannel)
 - HSCH (halt subchannel)
 - RSCH (resume subchannel)

- **Cross Memory instructions**
 - PC (program call)
 - PT (program transfer)
 - PR (program return)
 - SSAR (set secondary address space number)

- **Recovery**
 - RCVY (recovery termination manager processing)
 - ACR (alternate CPU recovery)

Common System Trace entries...

■ Other entries

- SSRV (system service entered through branch entry or PC)
- SUSP (suspended for a lock)
- BR (branch trace entry)
- MODE (mode trace entry)
- BSG (branch to subspace group)

Example: DSP - TCB Dispatch

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNIQUE-3	PSACLHS-	PSALOCAL	PASD	SASD
							UNIQUE-4	UNIQUE-5	UNIQUE-6	PSACLHSE			
00	0021	009FE030	DSP		070C1000	814C9DDE	00000000	814C9DDE	00FB3788	00000000	00000000	0021	0021
00	0021	009FE030	SVC	4F	070C1000	814CA396	009FDCB0	00000011	FFFFFFFF	Status	Start	SRBs	only
00	0021	009FE030	SVCR	4F	070C1000	814CA396	00000000	00000000	05616F00				
00	0021	009FE030	I/O	02E76	070C1000	814C97A2	00C04007	7E94C6F0	0C000001	00000080	00000000	0021	0021
								022B8850	00400002	00000000			
00	012A	0096ECD8	DSP		070C0000	874C2C12	00000000	00000001	3B5EF6BC	00000000	00000000	012A	012A
00	012A	0096ECD8	PC	...	0	3B5158E8		00107		PC	Deq		

- WU ADDR = TCB address
- UNIQUE-2 = R0
- UNIQUE-3 = R1
- TCB continues to run until a new unit of work is dispatched on the same CPU

Example: SVC/SVCR/SVCE

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNIQUE-3	PSACLHS-	PSALOCAL	PASD	SASD
							UNIQUE-4	UNIQUE-5	UNIQUE-6	PSACLHSE			
00	0021	009FE030	DSP		070C1000	814C9DDE	00000000	814C9DDE	00FB3788	00000000	00000000	0021	0021
00	0021	009FE030	SVC	4F	070C1000	814CA396	009FDCB0	00000011	FFFFFFEF	Status	Start SRBs only		
00	0021	009FE030	SVCR	4F	070C1000	814CA396	00000000	00000000	05616F00				
00	0021	009FE030	I/O	02E76	070C1000	814C97A2	00C04007	7E94C6F0	0C000001	00000080	00000000	0021	0021
								022B8850	00400002	00000000			
00	012A	0096ECD8	DSP		070C0000	874C2C12	00000000	00000001	3B5EF6BC	00000000	00000000	012A	012A
00	012A	0096ECD8	PC	...	0	3B5158E8		00107		PC Deq			

- SVC entry indicates an SVC issued by a TCB
- SVCR entry indicates completion of SVC
- SVCE entry indicates an error condition
- UNIQUE-1 = R15 UNIQUE-2 = R0 UNIQUE-3 = R1
- See SVC description after the UNIQUE fields
- See [z/OS MVS Diagnosis: Reference](#) for SVC summary and registers usage

Example: SRB/SSRB

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNIQUE-3	PSACLHS-	PSALOCAL	PASD	SASD
							UNIQUE-4	UNIQUE-5	UNIQUE-6	PSACLHSE			
05	0032	071671F0	SRB			470C0000 814EBF98	00000032	0689C500	0689C500	00		0032	0032
							008EDE88	20					
05	0032	00000000	SSRV	78		80FEC8BA	4080E552	00000058	008E6FA8	Getmain			
							00320000						
05	0032	008EDE88	SUSP			80067964	00000000	LOCL	00000000	00000000	00000000		
							09F247E0			00000000			
.....													
..... <i>Entries omitted here</i>													
...													
01	0032	071671F0	SSRB			470C3000 80067964	00000032		09F247E0		01	00000000	0032 0032
							008EDE88						

- SRB entry indicates initial SRB dispatch
- SSRB entry indicates suspended SRB re-dispatch
- WU-ADDR has WEB (Work Element Block) address
- UNIQUE-2 = R0 UNIQUE-3 = R1
- While SRB/SSRB is running, WU-ADDR will contain zero or associated TCB address (PurgeDeq TCB)

Example: Program Check

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNI	
							UNIQUE-4	UNIQUE-5	UNI	
01	0017	008FABD0	PGM	011	071C3000	81DA506E	00040011	008E2EA4		
01	0017	008FABD0	SVC	78	071C2000	81DA5346	0000E512	0000027A	000	
01	0017	008FABD0	SVCR	78	071C2000	81DA5346	00000000	00000280	008	
01	0017	008FABD0	SVC	77	071C2000	81DA5588	00000000	FFF00000	000	
01	0017	008FABD0	SVCR	77	071C2000	81DA5588	00000000	FFF00000	800	
01	0017	008FABD0	PGM	004	071C3000	81DA5AAA	00040004	008E2EA4		
01	0017	008FABD0	*RCVY PROG					940C4000	00000004	000

- A program interrupt can be resolvable
- A PGM entry indicates an error condition only if it is followed by RCVY
- RCVY entry indicates RTM (Recovery Termination Manager) in control
- UNIQUE-1 of PGM entry contains instruction length and interrupt code

Example: ABEND

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNI
							UNIQUE-4	UNIQUE-5	UNI
03	0014	02BB43A0	SRB		070C0000	80FE27BE	00000014	01DA612C	81D
							008DE580	20	
03	0014	00000000	SSRV	129		80FE28BE	01DA611C	00000000	000
							00000000		
03	0014	00000000	*SVCE	D	070C1000	80FE7B02	00000000	84000000	842
							00000004		
03	0014	00000000	*RCVY	FRR	070C0000	8100C544	84202000	00000000	000

- For ABEND (SVC D), R0 contains ABEND code (UNIQUE-3)
- Trace entry is SVCE because SVC was issued in an invalid environment
- Note RCVY entry following the SVCE D

Example: Loop

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNI
							UNIQUE-4	UNIQUE-5	UNI
03-0055	006D6248		CLKC		070C0000	81F79C88	00001004	006CAA80	
03-0055	006D6248		DSP		070C0000	81F79C88	00000900	00006000	000
03-0055	006D6248		CLKC		070C0000	81F79C90	00001004	006CAA80	
03-0055	006D6248		DSP		070C0000	81F79C90	00000900	00006000	000
03-0055	006D6248		I/O	968	070C0000	81F79C88	60800011	01B936D8	800
							0080FFFF	00FCFD98	
03-0055	006D6248		DSP		070C0000	81F79C88	00000900	00006000	000
03-0055	006D6248		CLKC		070C0000	81F79C94	00001004	006CAA80	

- A simple loop is indicated by trace entries with similar PSW addresses
- A complicated loop consists of groups of entries repeating and may not be obvious

Branch Tracing: is it useful?

PR	ASID	WU-ADDR-	IDENT	CD/D	PSW-----	ADDRESS-	UNIQUE-1	UNIQUE-2	UNI
							UNIQUE-4	UNIQUE-5	UNI
01	0017	008FABD0	BR		00FF5480	00FEAA00	00FF4A00		
01	0017	008FABD0	SVCR	38	075C3000	81FF774C	00000000	00000000	800
01	0017	008FABD0	BR		01FF7C20				
01	0017	008FABD0	SVC	63	070C1000	81FF850A	00000008	00000000	008
01	0017	008FABD0	BR		01DA671C	0110F8B6			
01	0017	008FABD0	BR		00FD5648	01186888	00FEA40C	011C7F28	02D
01	0017	008FABD0	SVC	78	071C2000	81DA5346	0000E512	0000027A	000
01	0017	008FABD0	BR		01144464	01133F34	011CED38	00FEA40C	011

- Branch Tracing traces target address of instructions such as BALR, BASR, BASSM, BAKR when they cause branching during their execution
- The return address is not traced
- Helpful when investigating module flow but does not give a complete picture

Missing System Trace Table: **why?**

```
BLS17541I No address spaces with the ERROR attribute were found
***** INVALID CONTROL BLOCK TTCH/01 AT 7EC1D000 = E3E3C3C8/02
***** INVALID CONTROL BLOCK TTCH/01 AT 7EC1D000 = E3E3C3C8/02
***** SYSTEM TRACE PROCESSING IS TERMINATED.
```

- Using the wrong level of IPCS (compared to level of system where dump was taken)
- ASID(4) is not dumped due to TRT not requested in SDATA of SVC Dump
- A major MVS control block is overlaid and IPCS cannot get to the system trace table

z/OS R12 SYSTRACE enhancement

- Traditionally SYSTRACE output consists of trace entries sorted by time
- New SYSTRACE parameter SORTCPU will provide output with trace entries sorted by CPU (in ascending CPU number)
- Sub-parameters of SORTCPU allow all or N entries for each CPU to be displayed before and after a specific time
- See [z/OS V1R12.0 MVS IPCS Commands](#) for details

Reference Information

■ Manuals

- [z/OS MVS IPCS Commands](#)
- [z/OS MVS IPCS Customization](#)
- [z/OS MVS IPCS User's Guide](#)
- [z/OS MVS Diagnosis: Reference](#)
- [z/OS MVS Diagnosis: Tools and Service Aids](#)
- [z/OS MVS System Codes](#)