

Basic System Problem Analysis



Bill Cadier
Hewlett-Packard Co.
TCSD MPE/iX Lab
25 April 2003

roduction



commonly used DAT macros

some OS structures and their types

PA-RISC Registers

short vs. long pointers

overview of the procedure calling convention

case studies

Commonly Used Macros



`sys_abort`

`om_ptree`

`om_family`

`om_errors`

`om_fplib`

`mi_showjob`

`mi_cihistory`

`mi_showvar`

`is_open_files`

`is_file`

`is_find_gufd_entry`

`lcx`

- `io_ios_diag_log`
- `rm_format_sirs`
- `rm_semaphore`
- `rm_sem_deadlock`
- `process_dispatcher`
- `process_wait`
- `vs_page_info`
- `mm_page_info`
- `mm_active_io`
- `mm_completed_io`
- `tbl_info`

Process Management Structures



P_{IB}: process information block, type "pib_type"

P_{IBX}: process information block extension, type "pibx_type"

P_{CB}: process control block (CM), type "pcb_type"

P_{CBX}: process control block extension (CM), type "pcbx_type"

Job/Session Management Structures



JMAT: job master table, type "jmat_entry_type"

JIT: job information table, type "jit_entry_type"

JDT: job directory table, type "jdt_header_type"

PLFD: process local file descriptor (file handle), type "plfd_type"

GDPD: global data pointer descriptor (file pointers), type "gdpd_t"

GUFD: global unique file descriptor, type "gufd_t"

FLAB: file label, type "flab_t"

Virtual Space Management Structures



VSOD: virtual space object descriptor, type
"vs_od_type"

Cache entry: type "cache_entry_type"

B-tree's:

extent b-tree, type "b_tree_root_type"

extent A/R (variable access rights) b-tree, same type

Memory Management Structures



HPDIR: hashed page directory, type "hpdir_rec"

IPDIR: indexed page directory, known system object (KSO) 3, type "ipdir_rec"

MIB: memory management information block, type "mib_type"

Memory Management Globals, known system object (KSO) 4, type "mm_global_info_rec"

Dispatcher Structures



dispatcher globals, KS 0 127, type "disp_globals_type"

TCB, task control block, type "tcb_type"

Table Management



used extensively in the OS, table header type `tbl_hdr`

characteristic of a "table management" table is that the first two words of the table point to itself

various types of tables are used, FIFO, LIFO, monotonic etc.

`TBL_INFO` macro is the easiest way to view a table management header

System Globals



System globals is ALWAYS found at address `$a.c0000000`

The type is "SYSTEM_GLOBALS_TYPE"

The macro `SYSGLOB` will return a field within the object.

Short pointer can be created with `ZDEPI 3, 1, 2, rx`

AA-RISC General Registers

AA-RISC uses 32 general registers. The procedure calling convention defines:

R30 is "SP" or the stack pointer

R27 is "DP" or the data pointer (global variables)

R2 is "RP" or the procedure return pointer

R28 and R29 are function return (ret0 and ret1)

R26, R25, R24 and R23 can contain the first four arguments passed to procedures (arg0..arg3)

R31 is used as the "millicode RP"

R0 is a read only register whose value is zero

A-RISC Space Registers



There are 8 space registers, SR0 to SR7

SR0 saves space ID for external branches

SR1 to SR3 loaded by software as needed

SR4, SR5, SR6 and SR7 are defined by the calling convention

SR4 is code, typically the space ID of your program

SR5 is data, the space ID of a process STACK

SR6 is always \$b (#11), OS structures and short mapped files

SR7 is always \$a (#10), OS structures, NL.PUB.SYS and short mapped files



Short vs. Long Pointers

Load and Store instructions that specify a space register of zero intend that the hardware will derive the space register using the first 2 bits of the offset portion of the address and add 4 to that giving the SR number to use.

LDW -296(0, 30), 22

SR30 contains 418432f0 the '4' is 0100 in binary. The first 2 bits, are 01 + 4 = 5. So SR5 will be used to complete the pointer.

Short vs. Long Pointers

Short pointer addressing the high order 2 bits of an offset are used to denote the space register therefore they are NOT used as part of the address. This means that using short pointers limits addressability to $2^{(30)}-1$ or 1 GB.

Long pointer reference would specify a space register from 0-7, for example:

```
W 272(sr1, r19), r21
```

Procedure Calling Convention



Stack Frames are built by **non-leaf** procedures so that when they call other procedures registers can be spilled into the frame and restored from there on return.

GR3 through GR18 are "callee save" registers, they are spilled, if necessary by the procedure that is being called.

GR19 through GR22 are "caller save" registers, saved by the procedure making the call.



Procedure Calling Convention: Registers

```
n = FREAD ( MPE_fd, &buf, -32767 );
```

```
=00000000 40100480 013dc097 41845630 R4 =d66e8018 d66ea018 00000000 00000000  
=00000000 00000000 00000000 00000000 R12=00000000 00000000 00000000 00000000  
5=00000000 00000000 00000000 00000000 R20=0000000a 013dc08c c01075a0 000002d6  
4=41845abc d44b1400 0000000a c0202008 R28=00000020 00000000 4184db30 0000008b
```

```
($2c5) nmdebug > dv sp-60,10
```

```
RT $2d6. 4184dad0 $ 00000000 4184568c 0000004d 4184567c  
RT $2d6. 4184dae0 $ 00000029 00000000 00000000 00000000  
RT $2d6. 4184daf0 $ 00000000 00000000 00000000 ffff8001 <- sp-34  
RT $2d6. 4184db00 $ 40bbe000 00000000 d44b1400 4164671c
```

The number is \$a or 10. The 'buffer' parameter to FREAD is a long pointer. As a result it must be aligned in registers and R23 and R24 will contain the value 0x41845abc and R25 is skipped. We only use R26..R23 so the "length" parameter is saved in the stack at SP-\$34.

Procedure Calling Convention: Stack Frame



3dc0c8	FREAD	6bc23fd9	STW	r2, -20(sr0, r30)
3dc0cc	FREAD+\$4	6fc30200	STWM	r3, 256(sr0, r30)
3dc0d0	FREAD+\$8	6bc43e09	STW	r4, -252(sr0, r30)
3dc0d4	FREAD+\$c	6bc53e11	STW	r5, -248(sr0, r30)
3dc0d8	FREAD+\$10	6bd73da1	STW	r23, -304(sr0, r30)
3dc0dc	FREAD+\$14	6bd83da9	STW	r24, -300(sr0, r30)
3dc0e0	FREAD+\$18	08000240	OR	r0, r0, r0
.
3dc160	FREAD+\$98	d35a1ff0	EXTRS	r26, 31, 16, r26

Procedure Calling Convention: SP & PSP



Since the STWM (or IDO) instruction is executed to build a new stack frame, all references to a procedure's parameters become "PSP" (previous stack pointer) relative.

GR26 to GR23 **may** be spilled to PSP-\$24 to PSP-\$30 respectively.

You cannot count on that occurring! There may be no need to save a register to memory.

Case Study: SA663



=a.0019fe78 system_abort

- 0) SP=418562e0 RP=a.00a51bc8 sm_quarantine_gufd+\$1fc
- 1) SP=418562e0 RP=a.00ee5a5c tm_close_common.tm_unlink_plfd_and_gdpd+\$18
- 2) SP=418558e0 RP=a.00ee75cc tm_close_common+\$1a98
- 3) SP=41855860 RP=a.0158a8e4 tm_ord_fix_buf_disc+\$1e4
- 4) **SP=418548a0 RP=a.01164370 fclose_nm+\$5d4**
- 5) SP=418547e0 RP=a.01163d68 ?fclose_nm+\$8
export stub: a.013d22a8 FCLOSE+\$b8
- 6) SP=41854560 RP=a.013d21bc ?FCLOSE+\$8
export stub: 298.00279b68 cr_fclose+\$1c
- 7) SP=418544a0 RP=298.00272350 COB_CLOSE+\$17c
- 8) SP=41854468 RP=298.0026e804 ?COB_CLOSE+\$8
export stub: 97c.0000e1d4
- 9) SP=418543f0 RP=97c.00000000
(end of NM stack)

Case Study: SA663



```
a ($70) nmdat > lev 5
```

```
b ($70) nmdat > dv psp-60, 10
```

```
T $866.41854500 $ d6ef0a94 41854418 05650003 4f4bbec1  
T $866.41854510 $ ca12a970 0300000a 84000000 013d21bc  
T $866.41854520 $ 06020000 00000000 00000003 00000000  
T $866.41854530 $ 01030000 00000000 41850000 4185000d
```

`close_nm` has spilled the file number `$d` to the stack.
Remember that the file number is a 16 bit value (see the
`CLOSE` intrinsic definition).

Case Study: SA663



```
193 ($70) nmdat > fs_file(,d)
```

```
Filename: TESTFILE.PUB.AP
```

```
Relative Mode file
```

```
Access options: APPEND, NOMR, LOCK, SHR, BUF, NOMULTI, WAIT, NOCOPY
```

```
Access method: $0
```

```
Last error number: $0
```

```
...
```

```
File options: SYS, BINARY, FORMAL, F, NOCCTL, DEQ, STD, NOLABEL
```

```
File code: $9c5
```

```
Record size: $100
```

```
Block size: $100
```

```
Record limit: $ffffeff00
```

```
...
```

```
dev: $76
```

Case Study: SA663



9d (\$70) nmdat > lev 2

9e (\$70) nmdat > dc pc

S \$a. ee5a5c

ee5a5c tm_close_common. tm_unlin*+\$184 2000008f ** Stmt 143

9f (\$70) nmdat > dcx pc-184, 188/4

Case Study: SA663



```
. . .
_close_common. tm_unlink*+$118      4bda3eb1  LDW      -168(sr0, r30), r26
_close_common. tm_unlink*+$11c      287fefff  ADDI L   $ffffff000, r3, 1
_close_common. tm_unlink*+$120      343900c8  LD0      100(r1), r25
_close_common. tm_unlink*+$124      4bd83ea9  LDW      -172(sr0, r30), r24
_close_common. tm_unlink*+$128      23ed0012  LDI L    $91a000, r31
_close_common. tm_unlink*+$12c      e7e02430  BLE      536(sr4, r31)
. ee5a04          *Call To:  tm_unlink_gdpd
_close_common. tm_unlink*+$130      081f0242  OR       r31, r0, r2
_close_common. tm_unlink*+$134      2000008d  ** Stmt  141
_close_common. tm_unlink*+$138      4bd63ea9  LDW      -172(sr0, r30), r22
_close_common. tm_unlink*+$13c      4ac10000  LDW      0(sr0, r22), r1
_close_common. tm_unlink*+$140      84202132  COMI BT, =, N 0, r1, tm_unlink
l fd_and_gdpd+$1e0
. . .
```


Case Study: SA663



```
1 ($70) nmdat > dv sp-#172  
T $866.41855834 $ 418545a8
```

```
2 ($70) nmdat > dv [sp-#172]  
T $866.418545a8 $ fc0e008f
```

```
3 ($70) nmdat > wl errmsg(S16(fc0e), 8f)  
e manager; unable to unlink the GDPD.
```

Case Study: SA663



```
sa6 ($70) nmdat > fv fs_gufd(fs_plfd(,d)) 'gufd_t'
```

```
CORD
```

```
. .
```

```
FILE_VIR_ADDR : 2e4.0
```

```
GDPD_PTR : 0
```

```
. .
```

```
QUARANTINE_REASON :
```

```
ALL : fc0e008f
```

```
QUARANTINE_TIME : 3b167877d4453
```

```
EOF_OFFSET : 94dd300
```

```
. .
```

```
STORE_ACTIVE : 1
```

hangs



hangs are usually difficult to diagnose.

determine the scope of the hang, what is affected

gather as much information as possible **BEFORE**
deciding to get a memory dump.

if you have to reboot the system to clear a hang you
may as well get a memory dump too, time permitting

memory dumps of hangs can be **MUCH** larger than
system abort memory dumps

Memory Dump is a Static Image



this aircraft ...

taking off or landing?

If landing, where has it been?

If taking off, where is it going?

What is its speed?

What is its current heading?



Before a Memory Dump



repeating the `SHOW PROC` command can tell if processes are using CPU time or not

`SHOWJOB` will tell what is presently running

`SHOW Q/SHOW WG` will show the present queue & workgroup settings

are disks active or idle

use debug to trace suspect processes

macros such as `pm_semaphore`, `rm_semaphore` can help

Case Study: Hang Memory Dump



```
0 ($0) nmdat > process_wai t
```

```
===== DISPATCHER INFORMATION FOR A PROCESS =====
```

IN #	State	Wait Event	Pri	Class	Blocked Reason
1	LONG_WAIT	IPC	\$7918	AS	Known Port fffffffed Progen Global Port
2	LONG_WAIT	IPC	\$38ff	BS	JUNK_WAIT
86	LONG_WAIT	Control Block	\$33ff	CS	CNTL_BLOCK_WAIT
87	LONG_WAIT	IPC	\$33ff	CS	TERMINAL_READ_WAIT
88	LONG_WAIT	IPC	\$33ff	BS	Jsmain Port ffff7fb0
89	LONG_WAIT	IPC	\$33ff	CS	CHILD_WAIT
8a	LONG_WAIT	Control Block	\$1bff	CS	CNTL_BLOCK_WAIT

Case Study: Hang Memory Dump



```
51 ($0) nmdat > pin 86
```

```
52 ($86) nmdat > pm_semaphores
```

```
ADDRESS OF SEMAPHORE WAITED ON: $b.88ae19b0
```

```
54 ($86) nmdat > rm_semaphore b.88ae19b0
```

```
list of pins waiting on semaphore at $b.88ae19b0
```

```
0 $6e $76 $7e $86 $8e $92 $96 $9a $9e $a2 $a6 $aa $ae
```

```
pid $35 has an exclusive lock on shareable semaphore at $b.88ae19b0
```

Case Study: Hang Memory Dump



```
56 ($86) nmdat > pin 35
```

```
57 ($35) nmdat > tr, d, i
```

```
PC=a.0017099c enable_int+$2c
```

```
* 0) SP=41853ef0 RP=a.00786004 notify_dispatcher.block_current_process+$3
```

```
1) SP=41853ef0 RP=a.00787e44 notify_dispatcher+$268
```

```
2) SP=41853e70 RP=a.001b6034 sem_block.wait_for_resource+$1bc
```

```
3) SP=41853d70 RP=a.001b6428 sem_block+$358
```

```
4) SP=41853cb0 RP=a.00757ce8 cb_shr_lock+$240
```

```
5) SP=41853bb0 RP=a.00757a94 ?cb_shr_lock+$8
```

```
export stub: fb.011380f8 lock' set' exclusive_345+$230
```

```
6) SP=41853a70 RP=fb.0113bc78 nmdbunlock+$16f4
```

```
7) SP=41853a30 RP=fb.0109ae70 dblock+$10c
```

```
8) SP=418522b0 RP=fb.0109ad38 ?dblock+$8
```

```
export stub: 48f.000060a0
```


Case Study: Hang Memory Dump



```
58 ($35) nmdat >pm_semaphores
```

```
ADDRESS OF SEMAPHORE WAITED ON: $b.88ae18d0
```

```
59 ($35) nmdat > rm_semaphore b.88ae18d0
```

```
list of pins waiting on semaphore at $b.88ae18d0
```

```
5 $72 $7a $82 $8a
```

```
in $60 has an exclusive lock on shareable semaphore at $b.88ae18d0
```

Case Study: Hang Memory Dump



```
5a ($35) nmdat > pin 60
```

```
5b ($60) nmdat > tr, d, i
```

```
PC=a.0017099c enable_int+$2c
```

```
* 0) SP=41853ef0 RP=a.00786004 notify_dispatcher.block_current_process+$3
```

```
1) SP=41853ef0 RP=a.00787e44 notify_dispatcher+$268
```

```
2) SP=41853e70 RP=a.001b6034 sem_block.wait_for_resource+$1bc
```

```
3) SP=41853d70 RP=a.001b6428 sem_block+$358
```

```
4) SP=41853cb0 RP=a.00757ce8 cb_shr_lock+$240
```

```
5) SP=41853bb0 RP=a.00757a94 ?cb_shr_lock+$8
```

```
export stub: fb.011380f8 lock'set' exclusive_345+$230
```

```
6) SP=41853a70 RP=fb.0113bc78 nmdbunlock+$16f4
```

```
7) SP=41853a30 RP=fb.0109ae70 dblock+$10c
```

```
8) SP=418522b0 RP=fb.0109ad38 ?dblock+$8
```

```
export stub: 309.000060a0
```

Case Study: Hang Memory Dump



```
5c ($60) nmdat > pm_semaphores
```

```
ADDRESS OF SEMAPHORE WAITED ON: $b.88ae19b0
```

```
5d ($60) nmdat > rm_semaphore b.88ae19b0
```

```
list of pins waiting on semaphore at $b.88ae19b0
```

```
0 $6e $76 $7e $86 $8e $92 $96 $9a $9e $a2 $a6 $aa $ae
```

```
in $35 has an exclusive lock on shareable semaphore at $b.88ae19b0
```

Haven't we been here before?

Case Study: Hang Conclusion



The hang is due to a database locking problem

The memory dump probably was not necessary,
DBUTIL "SHOW LOCKS" would probably have helped
determine what the problem was

The TELESUP utility "UNDELOCK" might even have
been able to correct it

the End!



THANK YOU



i n v e n t