

[OpenVMS] Discussion Of Unusual MUTEX Wait State

Last Technical Review: 6-APR-2000

Copyright (c) 1995, 2000. Compaq Computer Corporation. All rights reserved.

PRODUCT: Compaq OpenVMS VAX, All Versions  
Compaq OpenVMS Alpha, All Versions

COMPONENT: Scheduler

SOURCE: Compaq Computer Corporation

SYMPTOM:

The DCL command \$ SHOW SYSTEM shows a process in a MUTEX wait state. However, the event flag wait mask (in the SDA> SHOW PROCESS display) points to the Job Information Block (JIB), instead of to a mutex.

DISCUSSION:

This is a brief discussion of the problem; a detailed analysis and an example follow the SOLUTION section.

A process usually goes into a MUTEX wait state when the mutex it has requested is not available. However, in OpenVMS V5.n, a MUTEX wait state can occur if a process depletes a pooled quota.

To see if this is the case, examine the process with the System Dump Analyzer (SDA), with the command SDA> SHOW PROCESS. When a process is in MUTEX wait state, the "event flag wait mask" usually points to the address of the mutex the process is waiting for. However, in this case, the address is that of the JIB.

You can also tell from this display if there are remaining "timer entries", and if there is sufficient "BUFIO byte count/limit" for at least one buffered QIO.

You can also format the JIB to examine the counts and limits.

You can find more information on general performance tuning in the OpenVMS database article entitled:

[OpenVMS] Cookbook of Performance Slowdown, VAX and Alpha, V6.0 and Up

More information on investigating MUTEX wait states is in the OpenVMS database article entitled:

[OPENVMS] How to Troubleshoot a Process in MUTEX State

SOLUTION 1:

If DECamds is installed execute an "Adjust Process Quota Limit" fix from the Data Analyzer node on the target process. Adjust the depleted quota to allow the process to continue.

Before the process is restarted, adjust the appropriate User Authorization File (UAF) quota for the process.

SOLUTION 2:

You must reboot to clear the process.

Before you run the process again, raise the depleted limit in the UAF account for this process.

If the process logs in prior to the change, it must log out and back in to receive the new UAF limits.

SOLUTION 3:

Contact the Customer Support Center. As part of the Custom Projects Program, we can supply you with a kernel-mode program to add either TQELM or BYTLM to a process. The nominal cost is 2 hours at our standard time and materials rate.

Separate programs are available in this manner for VAX or Alpha systems. We supply both the .EXE and the .OBJ files so that the program may be relinked after OpenVMS upgrades.

ANALYSIS:

Starting in OpenVMS V5.n, a process can go into a MUTEX wait state without requesting a mutex.

When a process requests a pooled resource, OpenVMS checks the resource limit in the JIB. If not enough of the resource remains to satisfy the request, the process is put in a MUTEX wait state, and the event flag wait mask is set to the address of the process's JIB.

Currently, only the timer queue count and buffered I/O count return a MUTEX wait state when depleted. The byte field JIB\$B\_FLAGS indicates which of the two quotas has been exhausted:

Note:

On OpenVMS Alpha, this field is JIB\$L\_FLAGS. Only the bottom byte of the longword is used, and usage is identical.

bit set	value	UAF quota
0	1	BYTLM
1	2	TQELM

In the example below, JIB\$B\_FLAGS = 1, indicating that bit 0 is set, thus confirming the lack of BYTCNT; the UAF quota BYTLM should be raised.

These quotas correspond to the following fields:

In the UAF	In the JIB	SDA's "SHOW PROCESS" Display
BYTLM	JIB\$L_BYTLM	BUFIO byte xxxxx/limit
	JIB\$L_BYTCNT	BUFIO byte count/xxxxx
TQELM	JIB\$W_TQLM	
	JIB\$W_TQCNT	Timer entries allowed left

Example:

```
$ SHOW SYSTEM
VAX/VMS V5.0-2 on node TUBORG 8-FEB-1989 08:50:23.33 Uptime
  Pid Process Name State Pri I/O CPU
2AA00201 SWAPPER HIB 16 0 0 00:01:17.57
:
:
2AA00356 VTXSRV HIB 6 140 0 00:00:01.06
2AA00357 VISTASRV HIB 6 119 0 00:00:01.19
2AA00358 VTXUPDSRV LEF 6 97 0 00:00:00.97
2AA00759 K_RICHARDS MUTEX 7 6453 0 00:00:37.94
:
:
```

```
$ ANALYZE/SYSTEM
VAX/VMS System analyzer
```

```
SDA> READ SYS$SYSTEM:SYSDEF ! NOTE: On Alpha it's:
! SYS$LOADABLE_IMAGES:SYSDEF
!
%SDA-I-READSYM, reading symbol table SYS$COMMON:[SYSEXE]SYSDEF.STB;1
```

```
SDA> SHOW SUMMARY
Current process summary
```

Extended PID	Indx	Process name	Username	State	Pri	PCB
2AA00201	0001	SWAPPER		HIB	16	801EB798
:	:	:	:	:	:	:
2AA00759	0159	K_RICHARDS	K_RICHARDS	MUTEX	7	8087FD20
:	:	:	:	:	:	:

```
SDA> SHOW PROCESS/INDEX=159
Process index: 0159 Name: K_RICHARDS Extended PID: 2AA00759
```

```
Process status: 02040001 RES,PHDRES
```

PCB address	8087FD20	JIB address	80E2E470
PHD address	83D73600	Swapfile disk address	00000000
Master internal PID	00030159	Subprocess count	0
Internal PID	00030159	Creator internal PID	00000000
Extended PID	2AA00759	Creator extended PID	00000000
State	MUTEX	Termination mailbox	0000
Current priority	7	AST's enabled	KESU
Base priority	4	AST's active	NONE
UIC	[00022,000026]	AST's remaining	78

```

Mutex count                0      Buffered I/O count/limit    49/65
Waiting EF cluster         0      Direct I/O count/limit     18/18
Starting wait time        1B001B18  BUFIO byte count/limit    528/16320
Event flag wait mask      80E2E470  # open files allowed left  499
Local EF cluster 0       E0000002  Timer entries allowed left  50
Local EF cluster 1       DC000000  Active page table count    0
Global cluster 2 pointer 00000000  Process WS page count     277
Global cluster 3 pointer 00000000  Global WS page count      67

```

The "Event flag wait mask" field, which would usually contain the address of a mutex, points instead to the Job Information Block (JIB). None of these limits has reached 0, but note that "BUFIO byte count/limit" is 528/16320, also seen in the JIB\$L\_BYTCNT field below as a hexadecimal 210.

These values show that this user has only 528 bytes available for buffered I/O. However, the program that was running needs 1000 bytes for a buffered I/O. Since there are not enough bytes currently available, the process goes into the MUTEX wait.

```

SDA> FORMAT 80E2E470          ; format the JIB to see all the pooled
                               ; quota.
80E2E470  JIB$L_MTLFL          80E2E470
80E2E474  JIB$L_MTLBL          80E2E470
      :
      :
80E2E488  JIB$T_ACCOUNT
80E2E488                20205049
80E2E48C                20202020
80E2E490  JIB$L_BYTCNT        00000210      = 528 decimal <--+
80E2E494  JIB$L_BYTLM         00003FC0      |
80E2E498  JIB$L_PBYTCNT       00000000      |
80E2E49C  JIB$L_PBYTLIM       00000000      pooled limits |
80E2E4A0  JIB$W_FILCNT        01F3          |
      :
      :
80E2E4D9  JIB$B_FLAGS         01          Bit 0 is set; <--+
                               so we ran out of
                               BYTLM
SDA> EXIT

```