

Purpose

- To provide an overview of the use of semaphore debug output used to assist Lotus® Domino® performance investigations.

Introduction

- A Domino server is a multithread application that uses semaphores to protect access to shared resources. This includes databases, shared memory and code paths.
- With debug enabled a Domino server will provide semaphore information that can be used to assist in investigating issues such as a server hang.

What is a semaphore and a semaphore timeout?

- **Semaphore** - A data structure which restricts thread and process access to shared resources. This includes databases, shared memory and code paths.
- **Semaphore timeout** – Occurs when a thread attempts to access a semaphore-protected resource for more than a set time limit (30 seconds).

What is Lock Manager?

- Lock Manager (LkMgr) is the device which protects access to Notes databases (NSF files) when Transactional Logging is enabled for Domino.
- When enabled, Lock Manager replaces the database semaphore (0x0244) as the device which controls thread access to Notes databases.
- Lock Manager is used for all Notes databases when Transactional logging is enabled, even if a individual database is not enabled for transactional logging.

Why do semaphore timeouts occur?

- A heavy load on Domino or the operating system can cause a thread to delay releasing a semaphore for other threads to access.
- A process or thread crashing while it is holding a semaphore.
- A semaphore contention where two threads are each holding a semaphore that the other needs.
- Semaphore timeouts may occur under other circumstances.

How to enable semaphore debug

- Set the following parameters in the notes.ini of the target server and then restart the Domino server.

- ▶ debug_capture_timeout=1
- ▶ debug_show_timeout=1
- ▶ debug_show_blockingthreadcallstack=1

NOTE for Domino 7.0 (and above). Debug_show_blockingthreadcallstack=1 shows the function call stack of the thread holding a semaphore if another thread times out on that semaphore.

- Lock Manager debug is not required as it posts semaphore timeouts to the Domino server console by default.

Semaphore debug output files


- With `debug_capture_timeout=1` enabled, for every 30 seconds a thread waits to access a semaphore it posts a timeout entry in the file `semdebug.txt`.
- By default the `semdebug.txt` file is located in the `IBM_TECHNICAL_SUPPORT` directory in the Domino server's data directory.
- Current `semdebug` log is named `semdebug.txt` until the server is restarted. It is renamed when the server starts to...
`semdebug_<hostname>_<Date>@<Time>.txt`
- With `debug_capture_timeout=1` enabled, semaphore timeouts are written to the server console and the `console.log` (if console logging is enabled).

Semaphore debug output

- `ti="000280A0-802573EB" sq="00210BA3" THREAD [573940:00002-00001] WAITING FOR READ LOCK ON FRWSEM 0x030B Collection semaphore (@8C94B1A4) (R=0,W=1,WRITER=258242:00001,1STREADER=00000:00000) FOR 30000 ms`
- Output must be annotated to convert the 'ti' field to date and time values
- `02/10/2008 12:27:20 AM GMT sq="00210BA3" THREAD [573940:00002-00001] WAITING FOR READ LOCK ON FRWSEM 0x030B Collection semaphore (@8C94B1A4) (R=0,W=1,WRITER=258242:00001,1STREADER=00000:00000) FOR 30000 ms`

Annotate semdebug.txt

- The Lotus Notes® Diagnostic Tool can be used to annotate semdebug files.
- <http://www.lotus.com/idd/sandbox.nsf/ByDate/c0c64aa07fc9abae85257356004ac7b3>
- After installing Lotus Notes Diagnostic Tool, alternate click on the semdebug file and select “Convert Semdebug Time Innards.”



Annotated semdebug output

- Annotated output – SEM

<date/time> <log sequence number> THREAD [PID:VID-PTID] WAITING FOR SEM <SEM TYPE>
<Description> (@<address>) (**OWNER**=PID:PTID) FOR 30000 ms

- Annotated output – RWSEM and FRWSEM

<date/time> <log sequence number> THREAD [PID:VID-PTID] WAITING FOR <RWSEM/FRWSEM>
<SEM TYPE> <Description> (@<address>) (R=#readers, W=#writers, **WRITER**=PID:PTID,
1STREADER=PID:PTID) FOR 30000 ms

- Key fields are OWNER (SEM) and WRITER (RWSEM and FRWSEM)
 - ▶ They indicate the process and thread ID that is holding the semaphore.

Lock Manager output

- [0680:00AE-0A7C] LkMgr BEGIN Long Held Lock Dump -----
 - [0680:00AE-0A7C] Lock(Mode=SIX* LockID(DB DB=/mail/testuser1.nsf)) Waiters
countNonIntentLocks = 1 countIntentLocks = 0, queueLength = 3
 - [0680:00AE-0A7C] Req(**Status=Granted** Mode=SIX Class=Manual Nest=-1 Cnt=1
Tran=0 Func=N/A [0968:0002-0964])
 - [0680:00AE-0A7C] Req(Status=Waiting Mode=S Class=Manual Nest=-1 Cnt=1
Tran=0 Func=N/A [0C88:0002-0C84] Delay=22min)
 - [0680:00AE-0A7C] Req(Status=Waiting Mode=S Class=Manual Nest=-1 Cnt=1
Tran=0 Func=N/A [0904:002D-0EB0] Delay=20min)
 - [0680:00AE-0A7C] Req(Status=Waiting Mode=S Class=Manual Nest=-1 Cnt=1
Tran=0 Func=N/A [0680:0469-088C] Delay=19min)
 - [0680:00AE-0A7C] LkMgr END Long Held Lock Dump -----
-
- Key field is **Status=Granted**, which indicates the tread that is holding the lock.

Lock Manager output explanations

- LockID - database involved in the lock
- Status – status of lock
 - ▶ Granted :: locked held
 - ▶ Waiting :: waiting for lock
- Mode – level of lock
 - ▶ Six :: write
 - ▶ S :: read
- [PID : VTID PTIN] process ID, virtual thread ID and physical ID of each thread waiting on this database thread

Review of output

- You should review timeouts on semaphores and Lock Manager concurrently.
- Review semaphore and Lock Manager timeouts from the time of a occurrence (hang or performance) backwards. The larger a gap between semaphore timeouts the less likely the timeouts are related.
- Suspect threads taken from semaphore timeouts are used to investigate other diagnostic files. For example, the Notes System Diagnostic (NSD) log or the console log.

Using timeout data to investigate NSD files

- One or more manual NSDs are typically gathered during a performance or hang issue.
- For more information on manual NSDs on UNIX® platforms, see the following technote:

How to run NSD manually on a Domino server for UNIX platforms (#1214298)

<http://www.ibm.com/support/docview.wss?rs=899&uid=swg21214298>

Example - Timeout data to investigate NSD files

- Review of a semdebug.xt from the time of a server hang back reveals a suspect thread. The thread is suspect because it was the first in a chain of consistent timeouts occurring.

```
02/10/2008 12:27:20 AM GMT sq="00210BA3" THREAD [573940:00002-00001] WAITING FOR READ LOCK ON  
FRWSEM 0x030B Collection semaphore (@8C94B1A4)  
(R=0,W=1,WRITER=258242:00002,1STREADER=00000:00000) FOR 30000 ms
```

- A search on the process ID 258242 in a manual NSD taken during the issue helps locate the function call stack for thread 2 for process 258242 (fixup)
- The call stack was then found for the first thread (00001)

```
[1] 258242: /opt/lotus702/lotus/notes/latest/ibmpow/fixup -F mail  
#####  
##### thread 2/3 :: fixup, pid=2582420, tid=2183653 #####  
#####  
[1] 0xd03b12e8 select(??, ??, ??, ??, ??) + ??  
[2] 0xd1a926fc unix_usleep(??) + 0x58  
[3] 0xd1a845e8 OSDelayThread(??) + 0x10  
[4] 0xd1a75338 StaticHangEnable() + 0x78  
[5] 0xd198fcf8 FileRead(??, ??, ??, ??) + 0xb0
```


References

Semaphores and Semaphore Timeouts (Technote #1094630)

<http://www.ibm.com/support/docview.wss?rs=899&uid=swg21094630>

Notes/Domino Best Practices: Master Checklist (Technote #7008523)

<http://www.ibm.com/support/docview.wss?rs=203&uid=swg27008523>

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

Current Domino IBM Lotus Lotus Notes

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.