

# SUMMARY: Solaris 10, Oracle 10G, Shared Memory Segments

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/SunManagers/2007-07/msg00092.html>

---

- *From:* Bryan Pepin <bpepin@xxxxxxx>
  - *Date:* Mon, 23 Jul 2007 09:57:09 -0400
- 

Hello All,

Our DBA's were able to uncover the mystery for us. It seems that Oracle, running on Solaris 10, has purposely broken up the SGA into multiple segments for performance improvements. See Oracle Metalink Note: 399261.1 (10G SGA is split in multiple shared memory segments). To disable this "feature",

`_enable_NUMA_optimization=FALSE` in the Oracle parameter file...

We are going to test this change ourselves, to see if there is an actual improvement running multiple segments.

For those without MetaLink access, here is the important information from the note...

...snip....

Applies to:

Oracle Server – Enterprise Edition – Version: 10.1.0.0 to 10.2.0.3

10g SGA is split in multiple shared memory segments

In 10g NUMA optimization is enabled by default while in 9i it is not because of this we see multiple segments in 10g. NUMA optimization is set by parameter `_enable_numa_optimization=true` in parameter file.

This use of multiple shared memory segments is expected in 10g for performance reasons. Where we create 1 per process group (lgrp to use sun terminology) and 1 that stripes across all lgrps and one that is a really small bootstrap segment. Performance should be better with multiple segments.

NUMA optimization is an internal optimization on the way the data structures are laid out and how the buffer cache is laid out such that we reduce the total number of remote cache misses on a large system.

If you want the database to start with a single segment, then set `_enable_NUMA_optimization=FALSE`.

....snip....

## SUMMARY: Solaris 10, Oracle 10G, Shared Memory Segments

Thanks.

–Bryan

Bryan Pepin wrote:

Hello,

Has anyone else noticed the following behavior regarding Oracle's use of shared memory, specifically around the shmmax (Solaris 8) and shm-max-memory (Solaris 10). Before we upgraded our servers from Solaris 8, our Oracle 10G databases were grabbing just 1 large shared memory segment for the database. For example, this is what ipcs said while running the DB on Solaris 8:

```
T ID KEY MODE OWNER GROUP CREATOR CGROUP
NATTCH SEGSZ CPID LPID ATIME DTIME CTIME ISMATTCH
Shared Memory:
m 19458 0xcf1bb170 --rw-r----- oracle dba oracle dba 790
31138578432 21097 22653 10:47:09 10:47:09 7:52:43 790
```

After "Live Upgrading" to Solaris 10, and converting all our old /etc/system settings to the new /etc/project format, we are noticing that now Oracle is grabbing the same "total" size shared mem segments, but instead of just 1, it is being split up into multiple ISM segments?

```
T ID KEY MODE OWNER GROUP CREATOR CGROUP
NATTCH SEGSZ CPID LPID ATIME DTIME CTIME ISMATTCH
PROJECT
Shared Memory:
m 112 0x8c07a648 --rw-r----- oracle dba oracle dba 1043
65536 23260 15444 14:25:28 14:25:28 16:21:59 1043 Oracle
m 111 0 --rw-r----- oracle dba oracle dba 1043
1811939328 23260 15444 14:25:28 14:25:28 16:21:55 1043
Oracle
m 109 0 --rw-r----- oracle dba oracle dba 1043
1811939328 23260 15444 14:25:28 14:25:28 16:21:51 1043
Oracle
m 107 0 --rw-r----- oracle dba oracle dba 1043
1795162112 23260 15444 14:25:28 14:25:28 16:21:47 1043
Oracle
m 106 0 --rw-r----- oracle dba oracle dba 1043
1795162112 23260 15444 14:25:28 14:25:28 16:21:43 1043
Oracle
m 105 0 --rw-r----- oracle dba oracle dba 1043
1795162112 23260 15444 14:25:28 14:25:28 16:21:40 1043
Oracle
m 104 0 --rw-r----- oracle dba oracle dba 1043
1795162112 23260 15444 14:25:28 14:25:28 16:21:36 1043
Oracle
m 103 0 --rw-r----- oracle dba oracle dba 1043
```

SUMMARY: Solaris 10, Oracle 10G, Shared Memory Segments

```
1795162112 23260 15444 14:25:28 14:25:28 16:21:33 1043
Oracle
m 102 0 --rw-r----- oracle dba oracle dba 1043
1795162112 23260 15444 14:25:28 14:25:28 16:21:29 1043
Oracle
m 101 0 --rw-r----- oracle dba oracle dba 1043
1795162112 23260 15444 14:25:28 14:25:28 16:21:26 1043
Oracle
m 100 0 --rw-r----- oracle dba oracle dba 1043
15015608320 23260 15444 14:25:28 14:25:28 16:20:51 1043
Oracle
```

Has anyone run into this? We are not sure if this is a Solaris 10 problem, or an Oracle issue? Sun Support is saying that have not seen this before.

Any solutions out there to get Oracle back to using just 1 large shared memory segment?

My last thing to test will be to put the old shmmax variable into /etc/system, and reboot to see if that magically fixes the situation?

Thanks.

-Bryan

---

\*\*\*\*\*

Bryan Pepin  
Unix Enterprise Systems

EMC Corporation  
4400 Computer Drive  
Westboro, MA 01580  
508-898-4776bpepin@xxxxxxx

---

sunmanagers mailing list  
sunmanagers@xxxxxxxxxxxxxxxxx  
<http://www.sunmanagers.org/mailman/listinfo/sunmanagers>