

[HPADM] RE: SHMMAX modification

Source: <http://unix.derkeiler.com/Mailing-Lists/HP-UX-Admin/2005-05/0020.html>

From: Vidal, Ignacio – (Arg) (ignacio.vidal_at_telmex.com)

Date: 05/09/05

Date: Mon, 9 May 2005 09:31:50 -0300

To: <hpx-admin@dutchworks.nl>

Hello:

Here are some answers to my original question.

A lot of thanks to Ilic, Daniel, Bill (Ryan), Bill (Hassel), Martin and Pal.

The problem has to do with 32-bit Oracle.

Thanks again!

Ignacio

This is a VERY common problem and has nothing to do with RAM!!

The problem is that you are using a 32bit version of Oracle so getting access to anything more than one quadrant (which is 1Gb) requires special relinking (called SHMEM_MAGIC). This will allow access to as much as 1750 megs of shared memory.

Now comes the icky part: All shared memory for all programs must fit into a single shared memory map, and while there may lots of shared memory pieces, none of them are large enough when Oracle requests 1000megs or 1500megs. You need to get a copy of shminfo ASAP to see how badly fragmented your shared memory area is. The `ipcs -bmop` command only shows successful segments, but cannot show fragmented free areas. These fragments are often caused by killing processes with `-9` (never a good idea) and not performing proper `ipcs` cleanup afterwards. Get shminfo from:

<ftp://contrib.9unsupp8@hprc.external.hp.com/sysadmin/programs/shminfo/>

and be sure to read the text file about SHMEM_MAGIC, SHARED_MAGIC and EXEC_MAGIC.

The solution is memory windows. What this does is to allow 32bit programs to run within their own window for shared memory. Since it is separate, all possible memory is available to your applications. NOTE: EVERY process that uses this shared memory area must be started in the same window. You may have several databases running at the same

time, each with their own window.

A 32bit program is severely limited on memory mapping. Without special linking and other considerations, using more than 1Gb for shared memory will be difficult. And depending on shared memory fragmentation, you might not even get 1Gb.

The absolute best solution is to convert to 64bit Oracle. All memory limitations go away.

```
--
Bill Hassell
#####
This tells me that you're using a 32-bit version of Oracle. Even on a
64-bit machine, 32-bit processes follow 32-bit rules, and it looks to me
like you're bumping up against them. Here's what I think is happening:
For every process, the kernel maintains a map relating the virtual
addresses
of memory pages allocated to the process, to the physical addresses the
kernel uses. Prior to the Itanium version of 11.23, HP-UX broke the
process
address space into four quadrants, each with an assigned function, for
performance reasons. This is still the default mode in 11.23, but a
flat
address map is available under some circumstances. In the 32-bit maps,
each
quadrant contains 1 Gb, for a total of 4 Gb of addressable space. In
64-bit
maps the quadrant size is 4 Tb and they're numbered differently, but
basic
principles are the same. In the 32-bit map, Q3 and Q4 contain the
shared
objects by default, and the same maps for those quadrants appear in all
32-bit process, except for a few exceptions not germane to this
discussion.
Which means that all the 32-bit process on the system combined can only
address 2 Gb of shared memory. But, 256 Mb of Q4 is reserved for I/O
addresses, so it's really only 1.75 Gb! And then there are the shared
libraries, which take the form of memory-mapped files. If you're using
SSH
and have UsePrivilegeSeparation set to "on," each connection will take
another 2.5 Mb bite out of that same scarce resource, also in the form
of
memory-mapped files. Telnet is much thriftier, but still needs to
create a
memory-mapped file. Your reading of the mmap man page is correct:
shared-mode memory-mapped files and shared memory segments compete for
the
same map space.
There are three basic ways to get around this. The first is to move
what
you can into the 64-bit space. Since that's a completely different set
of
maps and they're much, much bigger, you'll run out of global memory
before
you hit issues like this. The second is that it's possible to link
executables as "EXEC_MAGIC" and then use chatr(1) to change it to
"SHMEM_MAGIC," which will use Q2 in that process's map for shared
objects
also. The problem with that approach is that data in that shared
```

HP-UX-Admin: [HPADM] RE: SHMMAX modification

quadrant
can only be shared with other SHMEM_MAGIC programs. The third way is
"memory windows," which creates a new Q3 which is not the global Q3.
There
too, only processes assigned to use that specific Q3 can share data.
Bottom line: If you're hitting this limit with 32-bit Oracle but you
need
to assign more space and create more memory-mapped files, it's time to
consider 64-bit Oracle.
I hope this helps!
- Martin

--
Martin Lodahl of Auburn, CA

Oooooohhhh a classic one...
Your oracle is probably 32 bit.
And a 32 bit process can use at most 1.75 GB shared object (default).
Shared object is loaded (mmaped) shared libraries or a shared
memory segment. Oracle SGA is a shared memory segment, and the
error messages from telnet are due to the filled shared memory
region. This 1.75 region is shared among all 32bit processes!
You can use this:
file <oracle executable>
if the output is " PA-RISC1.1 (or PA-RISC2.0) shared executable"
then this is a 32bit application.
For really 64bit PA-RISC executable file command gives:
ELF-64 executable object file - PA-RISC 2.0 (LP64)
You could try to relink oracle as SHMEM_MAGIC.
I have never tried and never used this and
I do _not_ know whether Oracle supports this.
So you should (must...) use 64bit oracle.
Anyway raising shmmax above 1GB can not help on 32bit processes.
The largest continous shmem segment is 1GB for this.
In this case Oracle will use at least two shmmem segemnts.
Regards,
Pal Baranyai

Hello,
I didn't have exactly the same problem, but some problem with my
aplication.
To resolve this problem, I've set SHMMAX to max of RAM in server.
e.g. RAM(16GB) => shmmax = 0x400000000
Best regards
Daniel

I believe you will find that it is not enough swap space.
Run the command swapinfo -atm
You add memory , how about swap?
TOTAL swap avail must be equal to or greater than memory (preferably
1.5-2x of
memeory).
From your output I don't see how much memory you have:
Apparently somebody (probably Oracle process) made a call to mmap to map
a file
into memory, and there wasn't
enough memory..or is better to say not able to map space in swap..
Please be aware adding memory without extending swap make no real use of
that.
Answer it depends of your OS version? 11.00 has different memory
management than
11.i (and up). There is nice explanation on HP web site. Simply there
is no

HP-UX-Admin: [HPADM] RE: SHMMAX modification

```
enough contiguous amount of memory ( swap) to mmap process..
> -----Mensaje original-----
> De: hpux-admin-owner@DutchWorks.nl
> [mailto:hpux-admin-owner@DutchWorks.nl] En nombre de Vidal,
> Ignacio - (Arg)
> Enviado el: Viernes, 06 de Mayo de 2005 13:18
> Para: hpux-admin@dutchworks.nl
> Asunto: [HPADM] SHMMAX modification
>
>
> Good afternoon:
> Last night we changed the value for SHMMAX in one of our
> N4000 servers.
> This server has 10GB of RAM (we upgraded from 5GB some weeks ago).
> The previous value for SHMMAX was 1GB (0x40000000).
> We intended to raise SHMMAX to 2GB (0x80000000) because our dba wanted
> to enlarge Oracle's shared_pool_size's from 1300MB to 2000MB (due to a
> specific requirement for an application).
>
> After the change for SHMMAX (to 2GB), our dba could only change
> shared_pool_size's to 1350MB, and today I saw that some users get the
> following error executing a telnet to the server:
>
> # telnet localhost (<<<< this is an example, I can use SSH
> successfully)
> /usr/lib/dld.sl: Call to mmap() failed - TEXT /usr/lib/libsis.sl
> /usr/lib/dld.sl: Not enough space
> Abort(coredump)
>
> As I could see in mmap's man page, there seems to be a relationship
> between this function and shared memory, but I don't realize how to
> solve this issue.
>
> Does anyone have a similar problem?
> I'll appreciate very much your help
>
> Best regards
>
> Ignacio
>
> --
>           ---> Please post QUESTIONS and SUMMARIES only!! <---
>           To subscribe/unsubscribe to this list, contact
> majordomo@dutchworks.nl
>           Name: hpux-admin@dutchworks.nl      Owner:
> owner-hpux-admin@dutchworks.nl
>
> Archives:  ftp.dutchworks.nl:/pub/digests/hpux-admin
> (FTP, browse only)
>           http://www.dutchworks.nl/htbin/hpsysadmin   (Web,
> browse & search)
>
>
>
>           ---> Please post QUESTIONS and SUMMARIES only!! <---
>           To subscribe/unsubscribe to this list, contact majordomo@dutchworks.nl
> Name: hpux-admin@dutchworks.nl      Owner: owner-hpux-admin@dutchworks.nl
>
> Archives:  ftp.dutchworks.nl:/pub/digests/hpux-admin      (FTP, browse only)
>           http://www.dutchworks.nl/htbin/hpsysadmin   (Web, browse & search)
```