
Subject: PostgreSQL can't allocate shared memory
Posted by [Gregor Mosheh](#) on Mon, 16 Apr 2007 05:42:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi, all.

First: OpenVZ; nice. A slick alternative to VMWare, less overhead. Nice.
But, a hitch in my first experiments...

I have PostgreSQL 8.1 in a VE (Fedora Core 6, for what it's worth), and I cannot set `shared_buffers` higher than 3700 (about 30 MB) and have it start. Anything higher and it complains that it can't allocate the shared memory.

In the UBC, I verified that I set `shmpages` to 32000, which should be 128 MB at 4k apiece (or 256 MB if it's 8k pages; I'm not sure which it is for shared memory pages) and therefore more than enough for PgSQL's request. I have tried raising `shmpages`, but it doesn't change anything. The `privvmpages` is 262144 (1 GB) so this should also be plenty.

The UBC is showing 0 failcount across the board.

And the HN's `kernel.shmmax` is set to 1 GB. (Is that relevant in a VE?)

I'd be very appreciative of any help. And thanks for the software!

--

HostGIS
Cartographic development and hosting services
707-822-9355
<http://www.HostGIS.com/>

"Remember that no one cares if you can back up, only if you can restore."
- AMANDA

Subject: Re: PostgreSQL can't allocate shared memory
Posted by [dev](#) on Mon, 16 Apr 2007 07:37:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

Gregor Mosheh wrote:

> Hi, all.
> First: OpenVZ; nice. A slick alternative to VMWare, less overhead. Nice.
> But, a hitch in my first experiments...
>
> I have PostgreSQL 8.1 in a VE (Fedora Core 6, for what it's worth), and I
> cannot set `shared_buffers` higher than 3700 (about 30 MB) and have it
> start. Anything higher and it complains that it can't allocate the shared
> memory.

>
> In the UBC, I verified that I set shmpages to 32000, which should be 128
> MB at 4k apiece (or 256 MB if it's 8k pages; I'm not sure which it is for
> shared memory pages) and therefore more than enough for PostgreSQL's request. I
> have tried raising shmpages, but it doesn't change anything. The
> privvmpages is 262144 (1 GB) so this should also be plenty.
>
> The UBC is showing 0 failcount across the board.
>
> And the HN's kernel.shmmax is set to 1 GB. (Is that relevant in a VE?)
check the same in VE please.
AFAICS, by default kernel sets:
#define SHMMAX 0x2000000
which is 32Mb.
So you have to increase it in VE /etc/sysctl.conf file or in /proc.

> I'd be very appreciative of any help. And thanks for the software!
You are welcome!

Thanks,
Kirill

Subject: Re: PostgreSQL can't allocate shared memory
Posted by [Gregor Mosheh](#) on Mon, 16 Apr 2007 13:57:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

>> And the HN's kernel.shmmax is set to 1 GB. (Is that relevant in a VE?)
> check the same in VE please.
> AFAICS, by default kernel sets:
> #define SHMMAX 0x2000000
> which is 32Mb.
> So you have to increase it in VE /etc/sysctl.conf file or in /proc.

Ahhh, very tricky -- in that it's exactly what one should expect! Thanks,
Kirill; I can't believe that this escaped me. I had thought that the VE
would have shmmax set by the UBC; evidently it's the case that each VE
gets their own set of sysctls for perfectly ordinary use.

So, let me summarize the shared memory situation for VEs:

- * The HN's kernel.shmmax sets an absolute limit on shared memory,
for both the HN and any VEs.
- * A VE's shmpages sets the limit on shared memory for that VE.
Therefore it must be <= kernel.shmmax
- * Then the VE uses sysctl kernel.shmmax to set their allowable limit,
which of course must be <= the VE's shmpages.

Sound right?

--

HostGIS
Cartographic development and hosting services
707-822-9355
<http://www.HostGIS.com/>

"Remember that no one cares if you can back up, only if you can restore."
- AMANDA

Subject: Re: PostgreSQL can't allocate shared memory
Posted by [dev](#) on Mon, 16 Apr 2007 17:32:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

Gregor Mosheh wrote:

>>>And the HN's kernel.shmmax is set to 1 GB. (Is that relevant in a VE?)
>>

>>check the same in VE please.

>>AFAICS, by default kernel sets:

>>#define SHMMAX 0x2000000

>>which is 32Mb.

>>So you have to increase it in VE /etc/sysctl.conf file or in /proc.

>

>

> Ahhh, very tricky -- in that it's exactly what one should expect! Thanks,

> Kirill; I can't believe that this escaped me. I had thought that the VE

> would have shmmax set by the UBC; evidently it's the case that each VE

> gets their own set of sysctls for perfectly ordinary use.

>

> So, let me summarize the shared memory situation for VEs:

>

> * The HN's kernel.shmmax sets an absolute limit on shared memory,

> for both the HN and any VEs.

HN's kernel.shmmax controls the maximum single shmem segment size in VE0 (HN) only.

One can allocate no more than shmni segments in total

and no more than shmall bytes in total.

Same parameters inside VE mean the same things, but for VE.

> * A VE's shmpages sets the limit on shared memory for that VE.

> Therefore it must be <= kernel.shmmax

UBC shmpages parameter controls all the shared memory which possible to allocate
either via IPC or shmem (e.g. tmpfs).

sysctls above are only for IPC SYSv5 shared memory.

> * Then the VE uses sysctl kernel.shmmax to set their allowable limit,
> which of course must be \leq the VE's shmpages.

right, taking into account that UBC shmpages accounts non-IPC shmem also.

> Sound right?

right. Can you please describe this on wiki?

Thanks,
Kirill

Subject: Re: PostgreSQL can't allocate shared memory
Posted by [Gregor Mosheh](#) on Mon, 16 Apr 2007 17:40:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

> right. Can you please describe this on wiki?

Sure thing; will do this evening. Thanks for the clarifications; and
thanks for the correction that the HN's shmmax is really for the HN and
not for the VEs.
