
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 <= 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
