

## Re: Getting rid of SHMMAX/SHMALL ?

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-08/1095.html>

---

**From:** Matti Aarnio ([matti.aarnio\\_at\\_zmailer.org](mailto:matti.aarnio_at_zmailer.org))

**Date:** 08/04/05

Date: Thu, 4 Aug 2005 17:20:40 +0300  
To: Andi Kleen <[ak@suse.de](mailto:ak@suse.de)>

On Thu, Aug 04, 2005 at 03:23:38PM +0200, Andi Kleen wrote:

> On Thu, Aug 04, 2005 at 02:19:21PM +0100, Hugh Dickins wrote:

>> On Thu, 4 Aug 2005, Andi Kleen wrote:

>>

>>> I noticed that even 64bit architectures have a ridiculously low

>>> max limit on shared memory segments by default:

>>>

>>> #define SHMMAX 0x2000000 /\* max shared seg size (bytes) \*/

>>> #define SHMMNI 4096 /\* max num of segs system wide \*/

>>> #define SHMALL (SHMMAX/PAGE\_SIZE\*(SHMMNI/16)) /\* max shm system wide (pages) \*/

>>>

>>> Even on 32bit architectures it is far too small and doesn't

>>> make much sense. Does anybody remember why we even have this limit?

>>

>> To be like the UNIXes.

>

> Ok, no other more fundamental reason ? :)

> I cannot think of any at least.

Those supply DEFAULT values for bootup time, and they can be adjusted with sysctl. Existence of the limits is good.

Their easy tunability (even easier than at Solaris, where you tune them only with a reboot) is even better.

SHM resources are non-swappable, thus I would not by default let user programs go and allocate very much SHM spaces at all. Such is usually spelled as: "denial-of-service-attack" For that reason I would not raise builtin defaults either.

...

>

> I think we should just get rid of the per process limit and keep  
> the global limit, but make it auto tuning based on available memory.

Err... No thanks! I would prefer to have even finer grained control of how much SHM somebody can allocate. For normal user the value

## Linux-Kernel: Re: Getting rid of SHMMAX/SHMALL ?

might be zero, but for users in a group "SHM1" there could be a level of N MB, etc. (Except that such mechanisms are rather complex...)

For dedicated servers there is no problem of letting there be single global limit and its default value being in highish realms, but pick any machine with multiple user