

---

Subject: Re: [RFC PATCH 0/4] [RESEND] Change default MSGMNI tunable to scale with lowmem

Posted by [Nadia Derby](#) on Tue, 15 Jan 2008 15:30:19 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

Resending with the right address for the containers mailing list.  
So sorry...

Nadia Derby wrote:

> Hi,  
>  
> Yesterday, I posted to lkml a series of patches that make the ipc  
> tunable msgmni scale with lowmem (see thread  
> <http://lkml.org/lkml/2008/1/14/196>).  
>  
> Since these patches watch for memory hotplug notifications, I thought  
> that comments from linux-mm people would be interesting.  
>  
> Also, since these patches change the ipc\_namespace structure, comments  
> from people subscribed at the containers mailing list would be welcome too.  
>  
> Notes:  
> 1) please, Cc me since I'm not subscribed to linux-mm.  
> 2) sorry for lkml subscribers who are receiving this mail for the 2nd  
> time, but I wanted all the comments to be shared.  
>  
>  
> Here is patch 0, the complete series can be found in the thread  
> mentioned above.  
>  
> -----  
>  
> On large systems we'd like to allow a larger number of message queues.  
> In some cases up to 32K. However simply setting MSGMNI to a larger value  
> may cause problems for smaller systems.  
>  
> The first patch of this series introduces a default maximum number of  
> message queue ids that scales with the amount of lowmem.  
>  
> Since msgmni is per namespace and there is no amount of memory dedicated  
> to each namespace so far, the second patch of this series scales msgmni  
> to the number of ipc namespaces.  
>  
> In the last patch, a notifier block is added to the ipc namespace  
> structure to manage memory hotplug. The callback routine is activated  
> upon memory add/remove and it recomputes msgmni. One callback routine is  
> added to the memory notifier chain each time an ipc namespace is

> allocated. It is removed when the coresponding ipc namespace is freed.  
>  
> I still have 1 issue that I'll try to solve next:  
> . use the notification mechanism to recompute all the msg\_ctlmni each  
> time an ipc namespace is created / removed.  
>  
>  
> These patches should be applied to 2.6.24-rc7, in the following order:  
>  
> [PATCH 1/4]: ipc\_scale\_msgmni\_with\_lowmem.patch  
> [PATCH 2/4]: ipc\_scale\_msgmni\_with\_namespaces.patch  
> [PATCH 3/4]: ipc\_slab\_memory\_callback\_prio\_to\_const.patch  
> [PATCH 4/4]: ipc\_recompute\_msgmni\_on\_memory\_hotplug.patch  
>  
>  
> Regards,  
> Nadia  
>

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---