

Subject: Re: [RFC][PATCH v2 7/7] Do not recompute msgmni anymore if explicitly set by user

Posted by [Nadia Derby](#) on Tue, 05 Feb 2008 14:55:10 GMT

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

Yasunori Goto wrote:

> Thanks Nadia-san.

>

> I tested this patch set on my box. It works well.

> I have only one comment.

>

>

>

>>---

>> ipc/ipc_sysctl.c | 43 ++++++

>> 1 file changed, 41 insertions(+), 2 deletions(-)

>>

>>Index: linux-2.6.24/ipc/ipc_sysctl.c

>>=====

>>--- linux-2.6.24.orig/ipc/ipc_sysctl.c 2008-01-29 16:55:04.000000000 +0100

>>+++ linux-2.6.24/ipc/ipc_sysctl.c 2008-01-31 13:13:14.000000000 +0100

>>@@ -34,6 +34,24 @@ static int proc_ipc_dointvec(ctl_table *

>> return proc_dointvec(&ipc_table, write, filp, buffer, lenp, ppos);

>> }

>>

>>+static int proc_ipc_callback_dointvec(ctl_table *table, int write,

>>+ struct file *filp, void __user *buffer, size_t *lenp, loff_t *ppos)

>>+{

>>+ size_t lenp_bef = *lenp;

>>+ int rc;

>>+

>>+ rc = proc_ipc_dointvec(table, write, filp, buffer, lenp, ppos);

>>+

>>+ if (write && !rc && lenp_bef == *lenp)

>>+ /*

>>+ * Tunable has successfully been changed from userland:

>>+ * disable its automatic recomputing.

>>+ */

>>+ unregister_ipcns_notifier(current->nsproxy->ipc_ns);

>>+

>>+ return rc;

>>+}

>>+

>

>

>

> Hmm. I suppose this may be side effect which user does not wish.

>

- > I would like to recommend there should be a switch which can turn on/off
- > automatic recomputing.
- > If user would like to change this value, it should be turned off.
- > Otherwise, his request will be rejected with some messages.
- >
- > Probably, user can understand easier than this side effect.
- >

Hi Yasunori,

Hope you're feeling better!

Well the idea behind this was the following: if msgmni is changed say via procs it is usually to increase it, in order for applications that need more msg queues to be able to run.

So even if a new ipc namespace is created or memory is removed, the application that has set that new value doesn't care: it wants msgmni to be unchanged. I agree with you that unconditionally turning the recomputing off may appear coarse, but I'm afraid that adding the functionality of turning that recomputing on/off will make things more complicated:

- 1) manage the tunable recomputing state: it shouldn't be turned on twice
- 2) adds more things to do at the user level.

I'll try to think about it more deeply and may be come up with an intermediate solution.

Regards,
Nadia

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