
Subject: Re: [patch -mm] update mq_notify to use a struct pid
Posted by [Oleg Nesterov](#) on Mon, 11 Sep 2006 15:48:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 09/11, Cedric Le Goater wrote:

>
> Eric W. Biederman wrote:
> > Cedric Le Goater <clg@fr.ibm.com> writes:
> >
> >> message queues can signal a process waiting for a message.
> >>
> >> this patch replaces the pid_t value with a struct pid to avoid pid wrap
> >> around problems.
> >>
> >> Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>
> >> Cc: Eric Biederman <ebiederm@xmission.com>
> >> Cc: Andrew Morton <akpm@osdl.org>
> >> Cc: containers@lists.osdl.org
> >
> > Signed-off-by: Eric Biederman <ebiederm@xmission.com>
> >
> > I was just about to send out this patch in a couple more hours.
>
> Well, you did the same with the usb/devio.c and friends :)
>
> > So expect the fact we wrote the same code is a good sign :)
>
> How does oleg feel about it ? I've seen some long thread on possible race
> conditions with put_pid() and solutions with rcu. I didn't quite get all of
> it ... it will need another run for me.

I assume you are talking about this patch:

<http://marc.theaimsgroup.com/?l=linux-mm-commits&m=115773820415171>

I think it's ok, info->notify_owner is always used under info->lock.

This is simple. If, for example, mqqueue_read_file() didn't take info->lock, then we have a problem: pid_nr() may read a freed memory in case when __do_notify()->put_pid() happens at the same time.

In this context info->notify_owner is a usual refcounted object, no special attention is needed.

Oleg.

Containers mailing list
Containers@lists.osdl.org

