
Subject: Re: [PATCH] Mark timer_stats as incompatible with multiple pid namespaces

Posted by [Ingo Molnar](#) on Thu, 13 Dec 2007 21:37:11 GMT

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

* Eric W. Biederman <ebiederm@xmission.com> wrote:

> >> Well struct pid * works in that case if you grab the reference to
> >> it.
> >
> > but the display of the stats might happen much later. The point of
> > this API is to save pid+comm, which gives users a good idea about
> > what caused the events in the past - without having to pin any
> > resource of that task.
>
> Likewise struct pid is designed not to be a problem if pinned. It is a
> little heavier then it used to be with the addition of pid namespace
> support but not much. And if it is to heavy struct pid needs to be
> fixed.
>
> Holding the struct pid very much does not pin the task struct, and it
> shouldn't pin any other resources. I agree 64bytes or so is a bit
> more to pin then 4 bytes but it really isn't a lot.

yeah, and i have no conceptual objections - i just wanted to outline the thinking behind /proc/timer_stats.

Ingo

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
