
Subject: Re: [PATCH 4/9] pid: Generalize task_active_pid_ns
Posted by [Oleg Nesterov](#) on Thu, 13 Dec 2007 17:07:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 12/13, Eric W. Biederman wrote:

>
> Oleg Nesterov <oleg@tv-sign.ru> writes:
>
> > On 12/12, Eric W. Biederman wrote:
> >>
> >> Currently task_active_pid_ns is not safe to call after a
> >> task becomes a zombie and exit_task_namespaces is called,
> >> as nsproxy becomes NULL. By reading the pid namespace from
> >> the pid of the task we can trivially solve this problem
> >
> > Confused. If the task becomes a zombie, we can't assume it has a valid
> > ->pids[].pid. The parent can release us as soon as exit_notify() drops
> > tasklist.
>
>
> What this allows me to do (as seen later in the patchset) is to send
> to call pid_nr_ns and deliver a signal to a task group without caring
> if the element of the task group I am talking to is a zombie or not.

Yes I see, thanks Eric. I was confused by changelog, and I missed the subsequent changes which need task_active_pid_ns() of the reciever.

Oleg.

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