
Subject: Re: Processes with multiple pid_t values
Posted by [Herbert Poetzl](#) on Wed, 13 Dec 2006 08:10:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Dec 12, 2006 at 02:34:20PM -0700, Eric W. Biederman wrote:
> "Serge E. Hallyn" <serue@us.ibm.com> writes:
>
> > Sigh, here we go back again to the question of what to do in the case
> > of a lightweight container which doesn't have a /sbin/init. Let's say
> > I do
> > spawn_container(fork_and_exit(sleep 30m));
> >
> > so pid 1 forks, pid 2 sleeps 30 minutes, but pid 1 exits right after
> > the fork. What do we do? Create an idle pid 1? Tack a struct pid
> > to the parent namespaces' pid=1 making it pid=1 for the child namespace?

it would be perfectly fine, as long as the following is true:

- only the pid remains, not the associated task
- a 'fallback' reaper takes care of reaping the children (in the lightweight case)
- the task creating the container does not need to 'hang' around outside the container

> I was thinking whoever called spawn container would be pid 1.
> But I agree that is only a partial solution.

> It definitely makes sense to worry about this question later when all
> of the rest of the details for system containers are sorted out.

okay, as long as 'we' do not forget about this
on the way ...

best,
Herbert

> Eric
> _____
> Containers mailing list
> Containers@lists.osdl.org
> <https://lists.osdl.org/mailman/listinfo/containers>

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
