
Subject: Re: [RFC][PATCH] Do not set /proc inode->pid for non-pid-related inodes
Posted by [Herbert Poetzl](#) on Fri, 23 Mar 2007 01:06:27 GMT

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

On Thu, Mar 22, 2007 at 02:14:48PM +0100, Cedric Le Goater wrote:

>
> >> So I suggested to have a kthread be pid == 1 for each new pid
> >> namespace. the kthread can do the killing of all tasks if needed
> >> and will die when the refcount on the pid namespace == 0.
> >>
> >> Would such a (rough) design be acceptable for mainline ?
> >
> > The case that preserves existing semantics requires us to be able
> > to run /sbin/init in a container. Therefore pid 1 should be a user
> > space process.
>
> /sbin/init can't run without being pid == 1. hmm ? need to check. When
> we have more of the pid namespace, it should be easier.
>
> > So I don't think a design that doesn't allow us to run /sbin/init as
> > in a container would be acceptable for mainline.
>
> I agree that user space is assuming that /sbin/init has pid == 1 but
> don't you think that's a strong assumption ?

most inits around even act differently depending on
the pid, e.g. they act as telinit when pid != 1
so while it might be a wrong assumption, almost all
inits on Linux make that assumption and would need
to be changed ...

best,
Herbert

> on the kernel side we have is_init() so it shouldn't be an issue.
>
> C.
> _____
> Containers mailing list
> Containers@lists.linux-foundation.org
> <https://lists.linux-foundation.org/mailman/listinfo/containers>

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