
Subject: Re: Re: namespace and nsproxy syscalls
Posted by [Cedric Le Goater](#) on Wed, 27 Sep 2006 13:47:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert Poetzl wrote:

[...]

> that's just a waste of resources ... IMHO it is
> a little weird to actually consider having an init
> process 'just' to have a reference for a bunch of
> namespaces, given that you might want to access
> them individually, am I missing something?
>
> for me this suggestion sounds like making a dog
> mandatory for each household, so that when you
> want to get the younger son on the phone you
> can refer to him as 'the younger son of the family
> with the dog charly' :) ...

there is still the child reaping issue.

we need sometask to collect the SIGCHLD of a pidspace because we
can't redirect them the real init. This would create pid collisions
and the real init might well respawn a process that has never died.

C.

Subject: Re: Re: namespace and nsproxy syscalls
Posted by [Herbert Poetzl](#) on Wed, 27 Sep 2006 14:47:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Sep 27, 2006 at 03:47:08PM +0200, Cedric Le Goater wrote:

> Herbert Poetzl wrote:
>
> [...]
>
> > that's just a waste of resources ... IMHO it is
> > a little weird to actually consider having an init
> > process 'just' to have a reference for a bunch of
> > namespaces, given that you might want to access
> > them individually, am I missing something?
> >
> > for me this suggestion sounds like making a dog
> > mandatory for each household, so that when you
> > want to get the younger son on the phone you
> > can refer to him as 'the younger son of the family

> > with the dog charly' :) ...
>
> there is still the child reaping issue.
>
> we need sometask to collect the SIGCHLD of a pidspace because
> we can't redirect them the real init.

why not simply `_reserve_` the `pid=1` in such light-weight
guests as we do now, and have the host init take this
place in all those contexts

> This would create pid collisions and the real init might well
> respawn a process that has never died.

how so? either the pid is present inside the context
(because the process was not reaped) or it isn't, because
it was already reaped (by parent or fallback host init)

btw, IIRC, the reaping does not involve the pid except
for the returned pid, which could be made special
(e.g. 0) without any issues ...

it would also be fine with some magic auto reaping
done by a 'shared' kernel thread, so it would not have
to be the host init process at all

best,
Herbert

> C.
> _____
> 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>
