
Subject: Re: Re: namespace and nsproxy syscalls
Posted by [Cedric Le Goater](#) on Sat, 07 Oct 2006 21:40:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert Poetzl wrote:

>> * but we also said that a pid namespace can not survive the death of its
>> pid 1.
>
> which makes it unusable for our lightweight guest
> purpose if it requires a separate init process

the pid 1 process in a namespace can be the same for multiple namespaces,
which makes it a SPOF one would say, but we need a child reaper different
from the "real" init process to avoid pid value collisions.

>> yes, i'm testing such a patch as discussed on the list. I have good
>> results for a full nsproxy but i'm having trouble with the mnt namespace
>> (used to be called namespace) which is stored in nsproxy and the
>> fs_struct which is stored in the task_struct.
>
> what's the problem with handing out *space handles to userspace, which
> can be later used to reach a specific namespace and/or manipulate
> specific settings?

no problem. that's fine.

I'm being cautious with the mnt namespace.

cheers,

C.
