
Subject: FS 'namespace'

Posted by [Herbert Poetzl](#) on Thu, 07 Dec 2006 19:47:29 GMT

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

just a question: why do we keep the fs (struct_fs)
outside of nsproxy?

TIA,
Herbert

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>

Subject: Re: FS 'namespace'

Posted by [serue](#) on Fri, 08 Dec 2006 14:40:59 GMT

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

Quoting Herbert Poetzl (herbert@13thfloor.at):

>

> just a question: why do we keep the fs (struct_fs)

> outside of nsproxy?

Good question. So we have a mounts namespace, and you would consider the per-process fs root to be an fs namespace? Practically, it would mean that chroot and pivot_mount would create a new nsproxy, but i guess that's not a real problem.

It might force us to stop our current lazy checks for 'current->nsproxy==&init_nsproxy', since the pivot_mount in early boot would make that not true.

-serge

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>

Subject: Re: FS 'namespace'

Posted by [Herbert Poetzl](#) on Fri, 08 Dec 2006 16:41:15 GMT

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

On Fri, Dec 08, 2006 at 08:40:59AM -0600, Serge E. Hallyn wrote:

> Quoting Herbert Poetzl (herbert@13thfloor.at):
> >
> > just a question: why do we keep the fs (struct_fs)
> > outside of nsproxy?
>
> Good question. So we have a mounts namespace, and you
> would consider the per-process fs root to be an fs
> namespace? Practically, it would mean that chroot
> and pivot_mount would create a new nsproxy, but i guess
> that's not a real problem.
>
> It might force us to stop our current lazy checks for
> 'current->nsproxy==&init_nsproxy', since the pivot_mount
> in early boot would make that not true.

well, IMHO those are broken anyway, I can imagine
a number of applications using private namespaces
(the old ones) without running in 'containers'

best,
Herbert

> -serge

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

Subject: Re: FS 'namespace'
Posted by [serue](#) on Fri, 08 Dec 2006 17:19:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Herbert Poetzl (herbert@13thfloor.at):
> On Fri, Dec 08, 2006 at 08:40:59AM -0600, Serge E. Hallyn wrote:
> > Quoting Herbert Poetzl (herbert@13thfloor.at):
> > >
> > > just a question: why do we keep the fs (struct_fs)
> > > outside of nsproxy?
> >
> > Good question. So we have a mounts namespace, and you
> > would consider the per-process fs root to be an fs
> > namespace? Practically, it would mean that chroot
> > and pivot_mount would create a new nsproxy, but i guess
> > that's not a real problem.
> >
> > It might force us to stop our current lazy checks for
> > 'current->nsproxy==&init_nsproxy', since the pivot_mount

> > in early boot would make that not true.
>
> well, IMHO those are broken anyway, I can imagine

Yeah I wasn't defending them by calling them lazy :)

> a number of applications using private namespaces
> (the old ones) without running in 'containers'

Do you have a patch to move the fs_struct into nsproxy? I'd be interested in running some benchmarks with and without such a patch to see the effect of dereferencing the nsproxy so frequently.

-serge

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