
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>
