
Subject: Re: [patch -mm 08/17] nsproxy: add hashtable
Posted by [Dave Hansen](#) on Mon, 11 Dec 2006 22:53:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 2006-12-11 at 16:23 +0100, Cedric Le Goater wrote:
> > Even letting the concept of nsproxy escape to user space sounds wrong.
> > nsproxy is an internal space optimization. It's not struct container
> > and I don't think we want it to become that.
>
> i don't agree here. we need that, so does openvz, vserver, people working
> on resource management.

I think what those projects need is some way to group tasks. I'm not sure they actually need nsproxies.

Two tasks in the same container could very well have different nsproxies. The nsproxy defines how the pid namespace, and pid<->task mappings happen for a given task. The init process for a container is special and might actually appear in more than one pid namespace, while its children might only appear in one. That means that this init process's nsproxy can and should actually be different from its children's. This is despite the fact that they are in the same container.

If we really need this 'container' grouping, it can easily be something pointed to by the nsproxy, but it shouldn't be the nsproxy.

-- Dave

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