
Subject: Re: Re: [PATCH 1/1] Revert "[PATCH] identifier to nsproxy"
Posted by [Mishin Dmitry](#) on Tue, 12 Dec 2006 09:11:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tuesday 12 December 2006 00:02, Eric W. Biederman wrote:

> "Serge E. Hallyn" <serue@us.ibm.com> writes:

>

> > Quoting Eric W. Biederman (ebiederm@xmission.com):

> >

> > Now now, it's not being silently added, it was a very clearly commented

> > part of a proposed patchset sent to all interested parties for review,

> > and now being argued over. Sounds kosher to me.

>

> Yes. I guess the part that was moderately silent was the fact that it

> was intended to be exported to user space. If you couldn't see the implication

> that part was not explicit. But I do agree that I missed this patch

> in the first round of review, and my apologies for that.

>

> > I think the problem is that some people want to see an answer to the

> > namespace entering problem right now, but the alternate solution used on

> > using pids as implicit identifiers can't be used until the pidspaces are

> > fully implemented.

>

> I agree with that and that is a worthy discussion.

>

> One of the reasons I'm not too concerned is that sys_ptrace completely

> solves that problem today. The syscall interface completely sucks for

> handling that case but it works.

>

> The one instance where we clearly need a way to talk about namespaces

> besides enter is for moving network interfaces between network

> namespaces and I haven't looked close yet but I don't think either

> Dmitry or Daniel in their network namespace patches was using this id.

I didn't implement network devices move between I2 network namespaces yet,

but this code exists in OpenVZ. And there we use per-container id for this.

--

Thanks,
Dmitry.
