

---

Subject: Re: The issues for agreeing on a virtualization/namespaces implementation.

Posted by [serue](#) on Wed, 08 Feb 2006 21:22:12 GMT

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

---

Quoting Dave Hansen (haveblue@us.ibm.com):

> On Wed, 2006-02-08 at 12:03 -0600, Serge E. Hallyn wrote:

> > Now I believe Eric's code so far would make it so that you can only

> > refer to a namespace from it's \*creating\* context. Still restrictive,

> > but seems acceptable.

>

> The same goes for filesystem namespaces. You can't see into random

> namespaces, just the ones underneath your own. Sounds really reasonable

> to me.

Hmmm? I suspect I'm misreading what you're saying, but to be clear:

Let's say I start a screen session. In one of those shells, I clone, specify CLONE\_NEWNS, and exec a shell. now i do a bunch of mounting.

Other shells in the screen session won't see the results of those mounts, and if i ctrl-d, the shell which started the screen session can't either. Each of these is in the "parent filesystem namespace".

OTOH, shared subtrees specified in the parent shell could make it such that the parent ns, but not others, see the results. Is that what you're referring to?

thanks,

-serge

---