

---

Subject: Re: [patch -mm 08/17] nsproxy: add hashtable

Posted by [serue](#) on Mon, 11 Dec 2006 15:56:22 GMT

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

---

Quoting Serge E. Hallyn ([serue@us.ibm.com](mailto:serue@us.ibm.com)):

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

> > Herbert Poetzl <[herbert@13thfloor.at](mailto:herbert@13thfloor.at)> writes:

> > > Beyond that yes it seems to make sense to let user space

> > > maintain any mapping of containers to ids.

> > >

> > > I agree with that, but we need something to move

> > > around between the various spaces ...

> >

> > If you have CAP\_SYS\_PTRACE or you have a child process

> > in a container you can create another with ptrace.

> >

> > Now I don't mind optimizing that case, with something like

> > the proposed bind\_ns syscall. But we need to be darn certain

> > why it is safe, and does not change the security model that

> > we currently have.

>

> Sigh, and that's going to have to be a discussion per namespace.

Well, assuming that we're using pids as identifiers, that means we can only enter decendent namespaces, which means 'we' must have created them. So anything we could do by entering the ns, we could have done by creating it as well, right?

-serge

---

Containers mailing list

[Containers@lists.osdl.org](mailto:Containers@lists.osdl.org)

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

---