

---

Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [serue](#) on Wed, 07 Mar 2007 21:59:19 GMT

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

---

Quoting Paul Menage (menage@google.com):

> On 3/7/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

> >

> >All that being said, if it were going to save space without overly

> >complicating things I'm actually not opposed to using nsproxy, but it

>

> If space-saving is the main issue, then the latest version of my

Space saving was the only reason for nsproxy to exist.

Now of course it also provides the teensiest reduction in # instructions since every clone results in just one reference count inc for the nsproxy rather than one for each namespace.

> containers patches uses just a single pointer in the task\_struct, and

> all tasks in the same set of containers (across all hierarchies) will

> share a single container\_group object, which holds the actual pointers

> to container state.

Yes, that's why this consolidation doesn't make sense to me.

Especially considering again that we will now have nsproxies pointing to containers pointing to... nsproxies.

> Effectively, container\_group is to container as nsproxy is to namespace.

>

> Paul

---

Containers mailing list

Containers@lists.osdl.org

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

---