
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [serge](#) on Tue, 03 Apr 2007 15:54:14 GMT

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

Quoting Paul Menage (menage@google.com):

> On 4/3/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

> >

> >But frankly I don't know where we stand right now wrt the containers
> >patches. Do most people want to go with Vatsa's latest version moving
> >containers into nsproxy? Has any other development been going on?
> >Paul, have you made any updates?

>

> I've not made major changes since the last patch post, just some small
> optimizations and fixes - I've been too tied up with other stuff.

>

> Whilst I've got no objection in general to using nsproxy rather than
> the container_group object that I introduced in my latest patches, I

Hmm, my largest objection had been that the nsproxy as a container structure would end up pointing to nsproxy as a namespace proxy.

But if we do as Eric suggests and have one subsystem per namespace type, rather than one subsystem for all namespaces, I guess that is no longer a problem.

That still leaves yours.

> think that Vatsa's approach of losing the general container object is
> flawed, since it loses any kind of per-group generic state (e.g. "this
> container is being deleted") and last time I saw it, I think it would
> tend to lose processes so that they didn't show up in any directory in
> the container fs.

>

> Paul

> -

> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at <http://vger.kernel.org/majordomo-info.html>
> Please read the FAQ at <http://www.tux.org/lkml/>
