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

Posted by [serue](#) on Tue, 03 Apr 2007 15:32:20 GMT

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

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

> Srivatsa Vaddagiri <vatsa@in.ibm.com> writes:

>

> > On Mon, Apr 02, 2007 at 09:09:39AM -0500, Serge E. Hallyn wrote:

> >> Losing the directory isn't a big deal though. And both unsharing a

> >> namespace (which causes a ns_container_clone) and mounting the hierarchy

> >> are done by userspace, so if for some installation it is a big deal,

> >> the init scripts on initrd can mount the hierarchy before doing any

> >> unsharing.

> >

> > Yes I thought of that after posting this (that initrd can mount the

> > hierarchies), so this should not be an issue in practice ..

> >

> >> Can you think of a reason why losing a few directories matters?

> >

> > If we loose directories, then we don't have a way to manage the

> > task-group it represents thr' the filesystem interface, so I consider

> > that bad. As we agree, this will not be an issue if initrd

> > mounts the ns hierarchy atleast at bootup.

>

> I suspect that could be a problem if we have recursive containers.

> Even by having a separate mount namespace for isolation you really

> don't want to share the mount. If you can see all of the processes

> you do want to be able to see and control everything.

>

> I guess I want to ask before this gets to far. Why are all of the

> namespaces lumped into one group? I would think it would make much

> more sense to treat each namespace individually (at least for the

> user space interface).

You mean why is there one subsystem for all namespaces?

Convenience and no reason not to. What advantage would there be for the userspace tools if we have them split up?

There's no real reason other than memory waste not to split them up - vserver users could always just mount all namespace subsystems under one hierarchy.

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?

-serge
