
Subject: Re: [PATCH 11/11] sysfs: user namespaces: fix bug with
clone(CLONE_NEWUSER) with fairsched

Posted by [serue](#) on Thu, 26 Jun 2008 13:07:52 GMT

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

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

> "Serge E. Hallyn" <serue@us.ibm.com> writes:

>

> > Except the sysfs mount holds no refcount on the userns. So as long as we

> > do the ida tagging as you suggested in your response to patch 6, there

> > should be no reference to the user_ns left in sysfs code.

> >

> > The extra reference in patch #9 is for a light ref on the network

> > namespace. I'm still not sure that needs to be there, since if the

> > network namespace goes away, it will properly unregister its sysfs

> > mounts. Eric, Benjamin, I really don't see any use for the hold_net()

> > from sysfs. What is it doing?

>

> Mostly just being a sanity check. We can remove that if it easier.

>

> Eric

In itself it seems an ok check (temporarily) to make sure that the
net_ns hook to unset the tag.netns at netns release is properly working,
but given that it's only checking for net_ns coding errors, and Tejun
wants a whole generic hooking infrastructure for net_ns to register with
so as to keep mention of net_ns out of sysfs code, yeah I think the sane
thing is to just remove it.

But the first thing (after Daniel is finished with the patch 5 fallout)
is to address the patch 6 comments about using ida. At first I didn't
like it (seemed like too much bookkeeping) but I think it'll actually
work out very nicely.

-serge

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
