
Subject: Re: [RFC PATCH] namespaces: document unshare security implications
Posted by [serue](#) on Thu, 15 Nov 2007 17:01:17 GMT

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

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

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

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

> > So I think CAP_SYS_ADMIN is a good starting place. It is trivial verifiable
> > that it is safe. So starting there allows us to work on other aspects
> > of the problem for now.

>

> It was a good starting place, but at this point I have two concerns with
> sticking with CAP_SYS_ADMIN:

>

> 1. now that file capabilities are upstream, people may want to
> add just the requisite capability in fP for an unsharing helper
> program. Cedric had mentioned wanting to do that.
> If we are going to switch to unprivileged unshares, then doing
> so later is ok. But if we're going to switch to a custom
> capability later, then that could be seen as an API change
> since users will have to switch the capability on all the
> unsharing programs.

>

> 2. As I pointed out a few times, we can cleanly separate
> unsharing namespace and actually manipulating the resources.
> By requiring CAP_SYS_ADMIN for both unsharing a mounts namespace
> and for performing privileged mounts, any program given the
> authority to unshare is automatically given the authority to
> also completely manipulate the mounts, both in the new private
> namespace and the original namespace (by just not unsharing).

>

> It's even worse with the net namespace, since the privilege
> needed to unshare the namespace authorizes you to update
> *other* namespaces in the system, but *not* network devices!
> But like you say let's stick with established namespaces.

Ok I'm being inconsistent (waffling between talking about not needing capabilities and using a separate capability), imprecise, and overly verbose.

Point 2 above is my key motivating factor.

So to attempt to state a clear, precise goal:

If limited unprivileged updates to a namespace are possible,
then the privilege needed to unshare the namespace should
be as isolated from other privileges as possible.

If limited unprivileged updates are not possible, or if

unsharing implicitly equals updating (*1) then the privilege needed to unshare should equal that to update the namespace, not another namespace (*2).

*1: as may be the case with NETNS since the new network namespace is created empty

*2: i.e. not CAP_SYS_ADMIN to unshare(NETNS) and CAP_NET_ADMIN to update.

-serge

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
