Subject: Re: [patch 0/8] unprivileged mount syscall
Posted by serue on Mon, 09 Apr 2007 14:38:02 GMT
View Forum Message <> Reply to Message

Quoting Miklos Szeredi (miklos@szeredi.hu):
> > > This patchset adds support for keeping mount ownership information in
> > > the kernel, and allow unprivileged mount(2) and umount(2) in certain
> > > cases.
> >
> > No replies, huh?
>
> All we need is a comment from Andrew, and the replies come flooding in ;)
>
> > My knowledge of the code which you're touching is not strong, and my spare
> > reviewing capacity is not high. And this work does need close review by
> > people who are familar with the code which you're changing.
> >
> > So could I suggest that you go for a dig through the git history, identify
> > some individuals who look like they know this code, then do a resend,
> > cc'ing those people? Please also cc linux-kernel on that resend.
>
> OK.
>
> > > One thing that is missing from this series is the ability to restrict
> > > user mounts to private namespaces. The reason is that private
> > > namespaces have still not gained the momentum and support needed for
> > > painless user experience. So such a feature would not yet get enough
> > > attention and testing. However adding such an optional restriction
> > > can be done with minimal changes in the future, once private
> > > namespaces have matured.
> >
> > I suspect the people who developed and maintain nsproxy would disagree ;)
>
> Well, they better show me some working and simple-to-use userspace
> code, because I've not seen anything like that related to mount
> namespaces.

If you mean to test/exploit them, see
http://lxc.sourceforge.net/patches/2.6.20/2.6.20-lxc8/broken-out/tests/

Compile the ns_exec.c program and do

 ns_exec -m /bin/sh

to get a shell in a new mounts namespace.

> pam_namespace.so is one example of a non-working, but probably-not-too-
> hard-to-fix one.

Non-working? I sure hope the one used for LSPP certification is working... As is the ugly version I wrote 18 mounts ago and use on my laptop.

> I'm just saying this is not yet something that Joe Blow would just
> enable by ticking a box in their desktop setup wizard, and it would
> all work flawlessly thereafter. There's still a _long_ way towards
> that, and mostly in userspace.

I'm not sure there's a that long a way to go, but clearly we need to be showing users what they can do, or they'll never work their way towards there.

For instance, as you say, a user admin gui with a checkmark and text boxes saying 'enter new namespace on login', 'create private /tmp', and 'create private dmcrypted /home' would be trivial right now.

-serge

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