Subject: Re: Re: [patch 05/10] add "permit user mounts in new namespace" clone flag
Posted by serue on Tue, 17 Apr 2007 17:07:37 GMT
View Forum Message <> Reply to Message

Quoting Miklos Szeredi (miklos@szeredi.hu):
> > > > > > Also for bind-mount and remount operations the flag has to be propagated
> > > > > > down its propagation tree. Otherwise a unpriviledged mount in a shared
> > > > > > mount wont get reflected in its peers and slaves, leading to unidentical
> > > > > > shared-subtrees.
> > > > > >
> > > > > > That's an interesting question. Do we want shared mounts to be
> > > > > > totally identical, including mnt_flags? It doesn't look as if
> > > > > > do_remount() guarantees that currently.
> > > > >
> > > > > Depends on the semantics of each of the flags. Some flags like of the
> > > > > read/write flag, would not interfere with the propagation semantics
> > > > > AFAICT. But this one certainly seems to interfere.
> > > >
> > > > That depends. Current patches check the "unprivileged submounts
> > > > allowed under this mount" flag only on the requested mount and not on
> > > > the propagated mounts. Do you see a problem with this?
> > >
> > > Don't see a problem if the flag is propagated to all peers and slave
> > > mounts.
> > >
> > > If not, I see a problem. What if the propagated mount has its flag set
> > > to not do un-priviledged mounts, whereas the requested mount has it
> > > allowed?
> >
> > Then the mount is allowed.
> >
> > It is up to the sysadmin/distro to design set up the propagations in a
> > way that this is not a problem.
> >
> > I think it would be much less clear conceptually, if unprivileged
> > mounting would have to check propagations as well.
> >
> > Miklos

I'm a bit lost about what is currently done and who advocates for what.

It seems to me the MNT_ALLOWUSERMNT (or whatever :) flag should be propagated. In the /share rbind+chroot example, I assume the admin would start by doing

```
mount --bind /share /share
mount --make-slave /share
```

mount --bind -o allow_user_mounts /share (or whatever)
mount --make-shared /share

then on login, pam does

chroot /share/$USER

or some sort of

mount --bind /share /home/$USER/root
chroot /home/$USER/root

or whatever. In any case, the user cannot make user mounts except under /share, and any cloned namespaces will still allow user mounts.

Or are you guys talking about something else?

-serge
