
Subject: Re: [patch 0/8] unprivileged mount syscall
Posted by [serue](#) on Wed, 11 Apr 2007 14:45:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Ian Kent (raven@themaw.net):

> On Wed, 2007-04-11 at 09:26 -0500, Serge E. Hallyn wrote:
> > Quoting Ian Kent (raven@themaw.net):
> > > On Wed, 2007-04-11 at 12:48 +0200, Miklos Szeredi wrote:
> > > > > >
> > > > > > - users can use bind mounts without having to pre-configure them in
> > > > > > /etc/fstab
> > > > > >
> > > > > >
> > > > > This is by far the biggest concern I see. I think the security
> > > > > implication of allowing anyone to do bind mounts are poorly understood.
> > > >
> > > > And especially so since there is no way for a filesystem module to veto
> > > > such requests.
> > > >
> > > > The filesystem can't veto initial mounts based on destination either.
> > > > I don't think it's up to the filesystem to police bind/move mounts in
> > > > any way.
> > > >
> > > But if a filesystem can't or the developer thinks that it shouldn't for
> > > some reason, support bind/move mounts then there should be a way for the
> >
> > Can you list some valid reasons why an fs could care where it is
> > mounted? The only thing I could think of is a stackable fs, but it
> > shouldn't care whether it is overlay-mounted or not.
>
> For my part, autofs and autofs4.

Ah, thanks.

I can see I'm going to have start using autofs to get to know the implementation, because it seems clear we'll run into it in the containers work again (beyond the struct pid conv) at some point.

> Moving or binding isn't valid.
> I tried to design that limitation out version 5 but wasn't able to.
> In time I probably can but couldn't continue to support older versions.

thanks,
-serge

> >
> > thanks,
> > -serge

> >
> > > filesystem to tell the kernel that.
> > >
> > > Surely a filesystem is in a good position to be able to decide if a
> > > mount request "for it" should be allowed to continue based on it's "own
> > > situation and capabilities".
> > >
> > > lan
> > >
> > >
> > >
> > > -
> > > To unsubscribe from this list: send the line "unsubscribe linux-fsdevel" in
> > > the body of a message to majordomo@vger.kernel.org
> > > More majordomo info at <http://vger.kernel.org/majordomo-info.html>

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