Subject: Re: [patch 0/8] unprivileged mount syscall
Posted by Karel Zak on Fri, 13 Apr 2007 20:07:20 GMT
View Forum Message <> Reply to Message

On Fri, Apr 13, 2007 at 01:58:59PM +0200, Miklos Szeredi wrote:
> > On Wed, 2007-04-11 at 12:44 +0200, Miklos Szeredi wrote:
> > > > 1. clone the master namespace.
> > > >
> > > > 2. in the new namespace
> > > >
> > > > move the tree under /share/$me to /
> > > > for each ($user, $what, $how) {
> > > > move /share/$user/$what to /$what
> > > > if ($how == slave) {
> > > > make the mount tree under /$what as slave
> > > > }
> > > > }
> > > >
> > > > 3. in the new namespace make the tree under
> > > > /share as private and unmount /share
> > >
> > > Thanks. I get the basic idea now: the namespace itself need not be
> > > shared between the sessions, it is enough if "share" propagation is
> > > set up between the different namespaces of a user.
> > >
> > > I don't yet see either in your or Viro's description how the trees
> > > under /share/$USER are initialized. I guess they are recursively
> > > bound from /, and are made slaves.
> >
> > yes. I suppose, when a userid is created one of the steps would be
> >
> > mount --rbind / /share/$USER
> > mount --make-rslave /share/$USER
> > mount --make-rshared /share/$USER
>
> Thinking a bit more about this, I'm quite sure most users wouldn't
> even want private namespaces. It would be enough to
>
>   chroot /share/$USER
>
> and be done with it.

I don't think so. How to you want to implement non-shared /tmp
directories? The chroot is overkill in this case. See:

http://www.coker.com.au/selinux/talks/sage-2006/PolyInstantiatedDirectories.html
http://danwalsh.livejournal.com/

> Private namespaces are only good for keeping a bunch of mounts
> referenced by a group of processes. But my guess is, that the natural
> behavior for users is to see a persistent set of mounts.
>
> If for example they mount something on a remote machine, then log out
> from the ssh session and later log back in, they would want to see
> their previous mount still there.

They can mount to /mnt where the directory is shared ("mount --make-shared /mnt") and visible and all namespaces.

I think /share/$USER is an extreme example. You can found more situations when private namespaces are nice solution.

Karel

--
Karel Zak <kzak@redhat.com>

_____

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