
Subject: Re: [patch 0/8] unprivileged mount syscall
Posted by [Ram Pai](#) on Mon, 16 Apr 2007 07:59:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 2007-04-13 at 13:58 +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.

>

> 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 will continue see their previous mount tree.

Even if all the namespaces belonging to the different sessions of the user get dismantled when all the sessions exit, the a mirror of those mount trees continue to exist under /share/\$USER in the original namespace. So I don't think we have a issue.

NOTE: when I say 'original namespace' I mean the admin namespace; the first namespace that gets created when the machine boots.

RP

>

> Miklos

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
