
Subject: Re: [patch 0/8] unprivileged mount syscall
Posted by [Miklos Szeredi](#) on Sun, 15 Apr 2007 20:21:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

> > 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?

```
mount --bind /.tmp/$USER /share/$USER/tmp
```

or whatever else this polyunsaturated thingy does within the cloned namespace.

> The chroot is overkill in this case.

What do you mean it's an overkill? clone(CLONE_NS) duplicates all the mounts, just as mount --rbind does.

> > 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.

Private to a single login session? I'd like to hear examples.

Thanks,
Miklos

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