
Subject: Re: [patch 0/8] unprivileged mount syscall
Posted by [Miklos Szeredi](#) on Fri, 13 Apr 2007 14:05:16 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.
> >
> > 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.
> >
> > Miklos
>
> Agreed on desired behavior, but not on chroot sufficing. It actually
> sounds like you want exactly what was outlined in the OLS paper.
>
> Users still need to be in a different mounts namespace from the admin
> user so long as we consider the deluser and backup problems

I don't think it matters, because /share/\$USER duplicates a part or
the whole of the user's namespace.

So backup would have to be taught about /share anyway, and deluser
operates on /home/\$USER and not on /share/*, so there shouldn't be any
problem.

There's actually very little difference between rbind+chroot, and
CLONE_NEWNS. In a private namespace:

- 1) when no more processes reference the namespace, the tree will be
disbanded
- 2) the mount tree won't be accessible from outside the namespace

Wanting a persistent namespace contradicts 1).

Wanting a per-user (as opposed to per-session) namespace contradicts
2). The namespace `_has_` to be accessible from outside, so that a new
session can access/copy it.

So both requirements point to the rbind/chroot solution.

Miklos

Containers mailing list

Containers@lists.linux-foundation.org

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