

---

Subject: Re: [patch 0/8] unprivileged mount syscall  
Posted by [Miklos Szeredi](#) on Mon, 16 Apr 2007 15:55:39 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> >> Arn't there ways to escape chroot jails? Serge had pointed me to a URL  
> >> which showed chroots can be escaped. And if that is true than having all  
> >> user's private mount tree in the same namespace can be a security issue?  
> >  
> > No. In fact chrooting the user into /share/\$USER will actually  
> > \_grant\_ a privilege to the user, instead of taking it away. It allows  
> > the user to modify it's root namespace, which it wouldn't be able to  
> > in the initial namespace.  
> >  
> > So even if the user could escape from the chroot (which I doubt), s/he  
> > would not be able to do any harm, since unprivileged mounting would be  
> > restricted to /share. Also /share/\$USER should only have read/search  
> > permission for \$USER or no permissions at all, which would mean, that  
> > other users' namespaces would be safe from tampering as well.  
>  
> A couple of points.  
> - chroot can be escaped, it is just a chdir for the root directory  
> it is not a security feature. The only security is that you have to  
> be root to call chdir. A carefully done namespace setup won't have  
> that issue.  
>  
> - While it may not violate security as far as what a user is allowed  
> to modify it may violate security as far as what a user is allowed  
> to see.

I think that's just up to the permissions in the global namespace. In this example if you 'chmod 0 /share' there won't be anything for the user to see.

> There are interesting per login cases as well such as allowing a  
> user to replicate their mount tree from another machine when they  
> log in. When /home is on a network filesystem this can be very  
> practical and can allow propagation of mounts across machines not  
> just across a single login session.

Yeah, sounds interesting, but I think it's better to get the basics working first, and then we can start to think about the extras.

Btw, there's nothing that prevents cloning the namespace `_after_` chrooting into the per-user tree. That would still be simpler than doing it the other way round: first creating per-session namespaces and then setting up mount propagation between them.

Miklos

---

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

---