Subject: Re: Re: [patch 05/10] add "permit user mounts in new namespace" clone flag
Posted by Miklos Szeredi on Tue, 17 Apr 2007 18:58:57 GMT
View Forum Message <> Reply to Message

> > mount --make-rshared /
> > mkdir -p /mnt/ns/$USER
> > mount --rbind / /mnt/ns/$USER
> > mount --make-rslave /mnt/ns/$USER
>
> This was my main point - that the tree in which users can mount will be
> a slave of /, so that propagating the "are user mounts allowed" flag
> among peers is safe and intuitive.

I think it's equally intuitive if flags are not propagated. After
all, I may want to set the mount flag _only_ on this particular mount,
and not anything else.

This is something I wouln't be sure about either way without
consulting the man.

> True. But the kernel functionality you provide enables both ways so no
> problem in either case :)
>
> > Another point: user mounts under /proc and /sys shouldn't be allowed.
> > There are files there (at least in /proc) that are seemingly writable
> > by the user, but they are still not writable in the sense, that
> > "normal" files are.
>
> Good point.
>
> > Anyway, there are lots of userspace policy issues, but those don't
> > impact the kernel part.
>
> Though it might make sense to enforce /proc and /sys not allowing user
> mounts under them in the kernel.

I think not, it would just be another policy decision having to be
made in the kernel on an fs by fs basis, and getting it wrong would be
much worse than getting it wrong in userspace.

Miklos