Subject: Re: Re: [patch 05/10] add "permit user mounts in new namespace" clone flag
Posted by Miklos Szeredi on Thu, 19 Apr 2007 09:02:10 GMT
View Forum Message <> Reply to Message

> Checking the permissions on the mountpoint to allow unmounting is
>
> - rather inelegant: user can't see those permissions, can only
>   determine if umount is allowed by trial and error
>
> - may be a security hole, e.g.:
>
>   sysadmin:
>
>   mkdir -m 777 /mnt/disk
>   mount /dev/hda2 /mnt/disk
>
> Of course the user doesn't have the right to delete the contents of
> the mount, yet the permissions on the mountpoint would imply that s/he
> has permission to umount the disk.

It is becoming increasingly apparent, that mount/umount permission
based on file permissions is inherently broken:

1) there are user-writable files under /proc/$PID/, which definitely
   shouldn't be allowed to be overmounted

2) if user mounts an fs read-only, then wants to create a submount of
   this, it will fail with the current patchset

Solving 2) should be trivial: submounting a mount owned by the user
should be always allowed regardless of the file permissions.

Maybe this could be generaized to say, that a mount can be submounted
by an unprivileged user IFF parent mount is owned by said user.

This would get rid of some of the complications in the current
patchset, namely the functionality of MNT_ALLOWUSERMNT and MNT_USER
flags would be merged, and the permission checking would be removed.

For example on login, the user could get a private namespace set up
some that the home directory is owned by the user, and hence can be
freely submounted:

```
clone(CLONE_NEWNS)
mount --bind /home/$USER /home/$USER
mount --remount -ouser=$USER /home/$USER
```

This is of course more limiting than allowing mounts based on file permissions, but it's also a lot cleaner.

Hmm?

Miklos

_____

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