Subject: Re: Re: [patch 05/10] add "permit user mounts in new namespace" clone flag
Posted by Miklos Szeredi on Tue, 17 Apr 2007 11:09:25 GMT
View Forum Message <> Reply to Message

> Interesting....
>
> So far even today these things can happen, however they are sufficiently
> unlikely the tools don't account for them.
>
> Once a hostile user can cause them things are more of a problem.
>
> > (Unless you want to tackle each problem legacy tool one at a time to
> > remove problems - i.e. deluser should umount everything under
> > /home/hallyn before deleting, backup should be spawned from it's own
> > namespace cloned right after boot or just back up on one filesystem,
> > etc.)
>
> I don't see a way that backup and deluser won't need to be modified
> to work properly in a system where non-priveleged mounts are allowed,
> at least they will need to account for /share.
>
> That said it is clearly a hazard if we enable this functionality by
> default.
>
> If we setup a pam module that triggers on login and perhaps when
> cron and at jobs run to setup an additional mount namespace I think
> keeping applications locked away in their own mount namespace is
> sufficient to avoid hostile users from doing unexpected things to
> the initial mount namespace. So unless I am mistake it should be
> relatively simple to prevent user space from encountering problems.
>
> That still leaves the question of how we handle systems with an old
> user space that is insufficiently robust to deal with mounts occurring
> at unexpected locations.
>
>
>   I think a simple sysctl to enable/disable of non-priveleged mounts
>   defaulting to disabled is enough.
>
> Am I correct or will it be more difficult than just a little pam
> module to ensure non-trusted users never run in the initial mount
> namespace?

I'm still not sure, what your problem is.

With the v3 of the usermounts patchset, by default, user mounts are
disabled, because the "allow unpriv submounts" flag is cleared on all

mounts.

There are several options available to sysadmins and distro builders to enable user mounts in a secure way:

- pam module, which creates a private namespace, and sets "allow unpriv submounts" on the mounts within the namespace

- pam module, which rbinds / onto /mnt/ns/$USER, and chroots into /mnt/ns/$USER, then sets the "allow unpriv submounts" on the mounts under /mnt/ns/$USER.

- sysadmin creates /mnt/usermounts writable to all users, with sticky bit (same as /tmp), does "mount --bind /mnt/usermounts /mnt/usermounts" and sets the "allow unpriv submounts" on /mnt/usermounts.

All of these are perfectly safe wrt userdel and backup (assuming it doesn't try back up /mnt).

Miklos

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