Subject: Re: [patch 5/8] allow unprivileged bind mounts
Posted by Miklos Szeredi on Sun, 22 Apr 2007 07:19:45 GMT
View Forum Message <> Reply to Message

> > From: Miklos Szeredi <mszeredi@suse.cz>
> >
> > Allow bind mounts to unprivileged users if the following conditions
> > are met:
> >
> > - mountpoint is not a symlink or special file
>
> Why? This sounds like a left over from when we were checking permissions.

Hmm, yes. Don't know. Maybe only the symlink check.

Bind mounts of directory over non-directy, and vica versa are already
excluded, even for root.

> > - parent mount is owned by the user
> > - the number of user mounts is below the maximum
> >
> > Unprivileged mounts imply MS_SETUSER, and will also have the "nosuid"
> > and "nodev" mount flags set.
>
> So in principle I agree, but in detail I disagree.
>
> capable(CAP_SETUID) should be required to leave MNT_NOSUID clear.
> capable(CAP_MKNOD) should be required to leave MNT_NODEV clear.
>
> I.e. We should not special case this as a user mount but rather
> simply check to see if the user performing the mount has the appropriate
> capabilities to allow the flags.

Sounds sane. Will fix.

Miklos

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