Subject: Re: [patch] unprivileged mounts update
Posted by Miklos Szeredi on Wed, 25 Apr 2007 15:18:12 GMT
View Forum Message <> Reply to Message

> From: Miklos Szeredi <mszeredi@suse.cz>
>
> - refine adding "nosuid" and "nodev" flags for unprivileged mounts:
>   o add "nosuid", only if mounter doesn't have CAP_SETUID capability
>   o add "nodev", only if mounter doesn't have CAP_MKNOD capability
>
> - allow unprivileged forced unmount, but only for FS_SAFE filesystems
>
> - allow mounting over special files, but not symlinks
>
> - for mounting and umounting check "fsuid" instead of "ruid"

Andrew, please skip this patch, for now.

Serge found a problem with the fsuid approach: setfsuid(nonzero) will remove filesystem related capabilities. So even if root is trying to set the "user=UID" flag on a mount, access to the target (and in case of bind, the source) is checked with user privileges.

Root should be able to set this flag on any mountpoint, _regardless_ of permissions.

It is possible to restore filesystem capabilities after setting fsuid, but the interfaces are rather horrible at all levels. mount(8) can probably live with these, but I'm not sure that using "fsuid" over "ruid" has enough advantages to force this.

Why did we want to use fsuid, exactly?

Thanks,
Miklos

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