Subject: Re: [patch 0/8] mount ownership and unprivileged mount syscall (v4)
Posted by Miklos Szeredi on Wed, 25 Apr 2007 07:18:28 GMT
View Forum Message <> Reply to Message

> > The following extra security measures are taken for unprivileged
> > mounts:
> >
> > - usermounts are limited by a sysctl tunable
> > - force "nosuid,nodev" mount options on the created mount
>
> The original userspace "user=" solution also implies the "noexec"
> option by default (you can override the default by "exec" option).

Unlike "nosuid" and "nodev", I don't think "noexec" has real security
benefits.

> It means the kernel based solution is not fully compatible ;-(

Oh, I don't think that matters. For traditional /etc/fstab based user
mounts, mount(8) will have to remain suid-root, the kernel can't
replace the fstab check.

In fact the latest patches don't even support these "legacy" user
mounts too well: setting the owner of a mount gives not only umount
privilege, but the ability to submount. This is not necessarily a
good thing for these kinds of user mounts.

We could add a new "nosubmount" or similar flag, to prevent
submounting, but that again would go against the simplicity of the
current approach, so I'm not sure it's worth it.

Miklos