Subject: Re: [patch 7/9] unprivileged mounts: allow unprivileged fuse mounts
Posted by Pavel Machek on Wed, 09 Jan 2008 13:35:06 GMT
View Forum Message <> Reply to Message

Hi!

> > ...this will break with FUSE enabled, right? (Minor security hole by
> > allowing users to stop c-a-delete, where none existed before?)
>
> Yup (or I don't know, I'm sure there was or is some problem with
> ptrace, that could be used to create unkillable processes).
>
> Fuse could actually be fixed to exit reliably for 'killall5 -9' (it
> used to), but that has other problems, and it doesn't seem very
> important to me. But this can be discussed.

I think it is better to fix fuse than to rewrite all the shutdown scripts.

> What cannot be fixed is if one process is inside an fs operation
> (e.g. unlink), holding a VFS lock (i_mutex) and another process goes
> to uninterruptible sleep on that lock. There's no way (other than
> rewriting the VFS) in which that second process could be killed unless
> you kill the first one or the fuse server.

I believe VFS should be rewritten here. Perhaps new "TASK_KILLABLE"
state can help?

> > I really believe FUSE vs. signals needs fixing. Either that, or
> > updating all the manpages
> >
> > man 1 kill:
> > -   KILL   9   exit   this signal may not be blocked
> > +   KILL   9   exit   this signal may not be blocked, except by FUSE user mount
>
> Heh, there are all very interesting, but most of these issues are not
> even on my todo list (which has grown into quite a big pile over the
> years), which means, that they don't seem to matter to people in
> practice.
>
> You seem to be implying that fuse is worthless if these issues are not
> fixed, but that is very far from the truth, I think.

I'm not saying fuse is worthless. It is a nice toy for single-user
systems. But I do not think we should be merging "allow ordinary users
to mount their own fuse's" before issues above are fixed.
Pavel
--
(english) http://www.livejournal.com/~pavelmachek

(cesky, pictures) http://atrey.karlin.mff.cuni.cz/~pavel/picture/horses/blog.html