Subject: Re: [patch 07/10] unprivileged mounts: add sysctl tunable for "safe" property
Posted by serue on Wed, 23 Jan 2008 00:00:40 GMT

View Forum Message <> Reply to Message

Quoting Miklos Szeredi (miklos@szeredi.hu):

```
> > > > What do you think about doing this only if FS_SAFE is also set,
> > > > so for instance at first only FUSE would allow itself to be
> > > > made user-mountable?
> > > >
> > > > A safe thing to do, or overly intrusive?
> > >
> > > It goes somewhat against the "no policy in kernel" policy ;).  I think
> > > the warning in the documentation should be enough to make sysadmins
> > > think twice before doing anything foolish:
> >
> > Warning in which documentation?  A sysadmin considering setting fs_safe
> > for ext2 or xfs isn't going to be looking at fuse docs, which I think is
> > what you're talking about.  Are you going to add a file under
> > Documentation/filesystems?
>
> Yes, I meant documentation of the new sysctl tunable in
> Documentation/filesystems/proc.txt:
```

Argh, sorry.

```
> > Index: linux/Documentation/filesystems/proc.txt
> > ===================================================================
> > --- linux.orig/Documentation/filesystems/proc.txt 2008-01-16 13:25:07.000000000 +0100
> > +++ linux/Documentation/filesystems/proc.txt 2008-01-16 13:25:09.000000000 +0100
> > @@ -43,6 +43,7 @@ Table of Contents
> >    2.13 /proc/<pid>/oom_score - Display current oom-killer score
> >    2.14 /proc/<pid>/io - Display the IO accounting fields
> >    2.15 /proc/<pid>/coredump_filter - Core dump filtering settings
> > +  2.16 /proc/sys/fs/types - File system type specific parameters
> >
> >  ------------------------------------------------------------------------------
> >  Preface
> > @@ -2283,4 +2284,21 @@ For example:
> >    $ echo 0x7 > /proc/self/coredump_filter
> >    $ ./some_program
> >
> > +2.16 /proc/sys/fs/types/ - File system type specific parameters
> > +---------------------------------------------------------------
> > +
> > +There's a separate directory /proc/sys/fs/types/<type>/ for each
> > +filesystem type, containing the following files:
> > +
```

> > +usermount_safe
> > +--------------
> > +
> > +Setting this to non-zero will allow filesystems of this type to be
> > +mounted by unprivileged users (note, that there are other
> > +prerequisites as well).
> > +
> > +Care should be taken when enabling this, since most
> > +filesystems haven't been designed with unprivileged mounting
> > +in mind.
> > +
> > ---
> >
>
> Do you think this is enough? Or do we need something more, to prevent
> sysadmin inadvertently setting this for an unsafe filesystem?

I would think something more would be good. First explaining
that fuse should be safe modulo warnings in the fuse documentation,
procfs and sysfs may be safe, while other filesystems are not known safe
at all.

Then explaining the dangers with not-known-safe filesystems and what is
needed to make them safe. Clearly making sure input validation is
properly done so for instance getsb() doesn't turn into a buffer
overflow, etc.

Such a checklist also would be useful for holding a meaningful discussion
about the other filesystems and maybe turning some people loose on
an audit of other filesystems.

thanks,
-serge