Subject: Re: [patch 1/8] add user mounts to the kernel
Posted by Miklos Szeredi on Sun, 22 Apr 2007 07:02:08 GMT
View Forum Message <> Reply to Message

> > The MNT_USER flag is not copied on any kind of mount cloning:
> > namespace creation, binding or propagation.
>
> I half agree, and as an initial approximation this works.
> Ultimately we should be at the point that for mount propagation
> that we copy the owner of the from the owner of our parent mount
> at the propagation destination.

Yes, that sounds the most sane.

Ram, what do you think?

> > + if (mnt->mnt_flags & MNT_USER)
> > + seq_printf(m, ",user=%i", mnt->mnt_uid);
> How about making the test "if (mnt->mnt_user != &root_user)"

We don't want to treat root_user special. That's what capabilities
were invented for.

> > Index: linux/include/linux/fs.h
> > ===
> > --- linux.orig/include/linux/fs.h 2007-04-20 11:55:02.000000000 +0200
> > +++ linux/include/linux/fs.h 2007-04-20 11:55:05.000000000 +0200
> > @@ -123,6 +123,7 @@ extern int dir_notify_enable;
> > #define MS_SLAVE (1<<19) /* change to slave */
> > #define MS_SHARED (1<<20) /* change to shared */
> > #define MS_RELATIME (1<<21) /* Update atime relative to mtime/ctime. */
> > +#define MS_SETUSER (1<<22) /* set mnt_uid to current user */
>
> If we unconditionally use the fsuid I think we can get away without
> this flag.

That coudl work if we wouldn't have to worry about breaking the user
interface. As it is, we cannot be sure, that existing callers of
mount(2) don't have fsuid set to some random value.

> > #define MNT_SHRINKABLE 0x100
> > +#define MNT_USER 0x200
>
> If we assign a user to all mount points and root gets to own the
> initial set of mounts then we don't need the internal MNT_USER
> flag.

I think we do want to treat "owned" mounts special, rather than

treating user=0 mounts special.

> > +
> > + uid_t mnt_uid; /* owner of the mount */
>
> Can we please make this a user struct. That requires a bit of
> reference counting but it has uid namespace benefits as well
> as making it easy to implement per user mount rlimits.

OK, can you ellaborate, what the uid namespace benifits are?

Miklos

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