
Subject: Re: [patch] unprivileged mounts update
Posted by [serue](#) on Thu, 26 Apr 2007 14:57:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Miklos Szeredi (miklos@szeredi.hu):

> > Right, I figure if the normal action is to always do
> > mnt->user = current->fsuid, then for the special case we
> > pass a uid in someplace. Of course... do we not have a
> > place to do that? Would it be a no-no to use 'data' for
> > a non-fs-specific arg?
>
> I guess it would be OK for bind, but not for new- and remounts, where
> 'data' is already used.
>
> Maybe it's best to stay with fsuid after all, and live with having to
> restore capabilities. It's not so bad after all, this seems to do the
> trick:
>
> cap_t cap = cap_get_proc();
> setfsuid(uid);
> cap_set_proc(cap);
>
> Unfortunately these functions are not in libc, but in a separate
> "libcap" library. Ugh.

Ok, are you still planning to nix the MS_SETUSER flag, though, as Eric suggested? I think it's cleanest - always set the mnt->user field to current->fsuid, and require CAP_SYS_ADMIN if the mountpoint->mnt->user != current->fsuid.

-serge

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