
Subject: Re: [patch] unprivileged mounts update
Posted by [Jan Engelhardt](#) on Fri, 27 Apr 2007 07:01:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Apr 26 2007 22:27, Miklos Szeredi wrote:
>> On Apr 25 2007 11:21, Eric W. Biederman wrote:
>> >>
>> >> Why did we want to use fsuid, exactly?
>> >
>> >- Because ruid is completely the wrong thing we want mounts owned
>> > by whomever's permissions we are using to perform the mount.
>>
>> Think nfs. I access some nfs file as an unprivileged user. knfsd, by
>> nature, would run as euid=0, uid=0, but it needs fsuid=jengelh for
>> most permission logic to work as expected.
>
>I don't think knfsd will ever want to call mount(2).

I was actually out at something different...

```
/* Make sure a caller can chown. */  
if ((ia_valid & ATTR_UID) &&  
    (current->fsuid != inode->i_uid ||  
     attr->ia_uid != inode->i_uid) && !capable(CAP_CHOWN))  
    goto error;
```

for example. Using current->[e]uid would not make sense here.

>But yeah, I've been convinced, that using fsuid is the right thing to
>do.

Jan
--

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