

---

Subject: Re: [patch 7/8] allow unprivileged mounts  
Posted by [ebiederm](#) on Sat, 21 Apr 2007 21:33:25 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Andi Kleen <[andi@firstfloor.org](mailto:andi@firstfloor.org)> writes:

> Andrew Morton <[akpm@linux-foundation.org](mailto:akpm@linux-foundation.org)> writes:

>  
>> On Fri, 20 Apr 2007 12:25:39 +0200 Miklos Szeredi <[miklos@szeredi.hu](mailto:miklos@szeredi.hu)> wrote:

>>  
>> > Define a new fs flag FS\_SAFE, which denotes, that unprivileged  
>> > mounting of this filesystem may not constitute a security problem.

>> >  
>> > Since most filesystems haven't been designed with unprivileged  
>> > mounting in mind, a thorough audit is needed before setting this flag.

>>  
>> Practically speaking, is there any realistic likelihood that any filesystem  
>> apart from FUSE will ever use this?

>  
> If it worked for mount --bind for any fs I could see uses of this. I haven't  
> thought  
> through the security implications though, so it might not work.

Binding a directory that you have access to in other was is essentially the same thing as a symlink. So there are no real security implications there. The only problem case I can think of is removal media that you want to remove but someone has made a bind mount to. But that is essentially the same case as opening a file so there are no new real issues. Although our diagnostic tools will likely fall behind for a bit.

We handle the security implications by assigning an owner to all mounts and only allowing you to add additional mounts on top of a mount you already own.

If you have the right capabilities you can create a mount owned by another user.

For a new mount if you don't have the appropriate capabilities nodev and nosuid will be forced.

Initial super block creation is a lot more delicate so we need the FS\_SAFE flag, to know that the kernel is prepared to deal with the crazy things that a hostile user space is prepared to do.

Eric

---

Containers mailing list

