

---

Subject: Re: [patch 0/8] unprivileged mount syscall  
Posted by [akpm](#) on Fri, 06 Apr 2007 23:02:38 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 04 Apr 2007 20:30:12 +0200 Miklos Szeredi <[miklos@szeredi.hu](mailto:miklos@szeredi.hu)> wrote:

> This patchset adds support for keeping mount ownership information in  
> the kernel, and allow unprivileged mount(2) and umount(2) in certain  
> cases.

No replies, huh?

My knowledge of the code which you're touching is not strong, and my spare reviewing capacity is not high. And this work does need close review by people who are familiar with the code which you're changing.

So could I suggest that you go for a dig through the git history, identify some individuals who look like they know this code, then do a resend, cc'ing those people? Please also cc linux-kernel on that resend.

> This can be useful for the following reasons:

>

> - mount(8) can store ownership ("user=XY" option) in the kernel  
> instead, or in addition to storing it in /etc/mstab. For example if  
> private namespaces are used with mount propagations /etc/mstab  
> becomes unworkable, but using /proc/mounts works fine

>

> - fuse won't need a special suid-root mount/umount utility. Plain  
> umount(8) can easily be made to work with unprivileged fuse mounts

>

> - users can use bind mounts without having to pre-configure them in  
> /etc/fstab

>

> All this is done in a secure way, and unprivileged bind and fuse  
> mounts are disabled by default and can be enabled through sysctl or  
> /proc/sys.

>

> One thing that is missing from this series is the ability to restrict  
> user mounts to private namespaces. The reason is that private  
> namespaces have still not gained the momentum and support needed for  
> painless user experience. So such a feature would not yet get enough  
> attention and testing. However adding such an optional restriction  
> can be done with minimal changes in the future, once private  
> namespaces have matured.

I suspect the people who developed and maintain nsproxy would disagree ;)

Please also cc [containers@lists.osdl.org](mailto:containers@lists.osdl.org).

> An earlier version of these patches have been discussed here:  
>  
> <http://lkml.org/lkml/2005/5/3/64>  
>  
> --

---

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

---