
Subject: Re: Re: [patch 05/10] add "permit user mounts in new namespace" clone flag

Posted by [Christoph Hellwig](#) on Wed, 18 Apr 2007 15:01:10 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, Apr 18, 2007 at 10:26:29AM -0400, Trond Myklebust wrote:

> > That doesn't really change my argument though. _If_ the flag is per
> > mount, then it makes sense to be able to change it on a master and not
> > on a slave. If mount flags are propagated, this is not possible.

>

> Read-only isn't the only issue. On something like NFS, there are flags
> to set the security flavour, turn on/off encryption etc.

>

> If I mount your home directory using no encryption in my namespace, for
> instance, then neither you nor the administrator will be able to turn it
> on afterwards in yours without first unmounting it from mine so that the
> superblock is destroyed.

I suspect the right answer here is to make nfs mount handling smarter.
The way mounting works the filesystem is allowed to choose whether it
can re-used a superblock or needs a new one. In the NFS case we probably
want to allow multiple superblocks for the same export if important
parameters like the security model mismatch.

This is however only the technical part of the answer. There's a more
important one, and that's the user experience - we need a much better
way to document which flags are per-superblock and which are per-mount.

Due to this and some other issues we'll probably need a new mount API
in some time.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
