
Subject: Re: [patch 5/9] unprivileged mounts: allow unprivileged bind mounts
Posted by [Dave Hansen](#) on Tue, 08 Jan 2008 18:26:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 2008-01-08 at 12:35 +0100, Miklos Szeredi wrote:

```
> @@ -510,10 +533,16 @@ static struct vfsmount *clone_mnt(struct
>                               int flag)
> {
>     struct super_block *sb = old->mnt_sb;
> -     struct vfsmount *mnt = alloc_vfsmnt(old->mnt_devname);
> +     struct vfsmount *mnt;
>
> +     if (flag & CL_SETUSER) {
> +         int err = reserve_user_mount();
> +         if (err)
> +             return ERR_PTR(err);
> +     }
> +     mnt = alloc_vfsmnt(old->mnt_devname);
>     if (!mnt)
> -         return ERR_PTR(-ENOMEM);
> +         goto alloc_failed;
>
>     mnt->mnt_flags = old->mnt_flags;
>     atomic_inc(&sb->s_active);
```

I think there's a little race here. We could have several users racing to get to this point when `nr_user_mounts==max_user_mounts-1`. One user wins the race and gets their mount reserved. The others get the error out of `reserve_user_mount()`, and return.

But, the winner goes on to error out on some condition further down in `clone_mnt()` and never actually instantiates the mount.

Do you think this is a problem?

I think just about the one solution is to block new mounters on a sleepable lock until the race winner actually finishes their mount operation.

-- Dave

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