Subject: Re: [patch 5/9] unprivileged mounts: allow unprivileged bind mounts
Posted by Szabolcs Szakacsits on Tue, 08 Jan 2008 20:44:20 GMT
View Forum Message <> Reply to Message

On Tue, 8 Jan 2008, Miklos Szeredi wrote:
> > On Tue, 2008-01-08 at 12:35 +0100, Miklos Szeredi wrote:
> > > +static int reserve_user_mount(void)
> > > +{
> > > +	int err = 0;
> > > +
> > > +	spin_lock(&vfsmount_lock);
> > > +	if (nr_user_mounts >= max_user_mounts && !capable(CAP_SYS_ADMIN))
> > > +		err = -EPERM;
> > > +	else
> > > +		nr_user_mounts++;
> > > +	spin_unlock(&vfsmount_lock);
> > > +	return err;
> > > +}
> >
> > Would -ENOSPC or -ENOMEM be a more descriptive error here?
>
> The logic behind EPERM, is that this failure is only for unprivileged
> callers. ENOMEM is too specifically about OOM. It could be changed
> to ENOSPC, ENFILE, EMFILE, or it could remain EPERM. What do others
> think?

I think it would be important to log the non-trivial errors. Several mount(8) hints to check for the reason by dmesg since it's already too challanging to figure out what's exactly the problem by the errno value. This could also prevent to mislead troubleshooters with the mount/sysctl race.

Szaka

--
NTFS-3G: http://ntfs-3g.org

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