
Subject: Re: [patch 3/8] account user mounts

Posted by [Miklos Szeredi](#) on Sun, 22 Apr 2007 08:08:49 GMT

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

> > > From: Miklos Szeredi <mszeredi@suse.cz>
> > >
> > > Add sysctl variables for accounting and limiting the number of user
> > > mounts.
> > >
> > > The maximum number of user mounts is set to 1024 by default. This
> > > won't in itself enable user mounts, setting a mount to be owned by a
> > > user is first needed
> > >
> > > Since each mount has a user can we just make this a per user rlimit?
> > >
> > > If we are going to implement a sysctl at this point I think it should
> > > be a global limit that doesn't care if who you are. Even root can
> > > have recursive mounts that attempt to get out of control.
> > >
> > > Recursive bind mounts are done carefully enough, so they don't get out
> > > of control.
> > >
> > > Recursive mount propagations can get out of control. But root can
> > > shoot itself in the foot any number of ways, and it's not for the
> > > kernel to police that.
> > >
> > > Yes. It is.
> > >
> > > This is mostly about removing special cases.
> > >
> > > We routinely have limits on resources that are global and apply
> > > to root along with every one else. Root can change them but
> > > they still apply to root. Things like the number of inodes
> > > in the system or the total number of files.

There's no max_inodes any more. As for max_files:

get_empty_filp():

```
/*  
 * Privileged users can go above max_files  
 */  
if (get_nr_files() >= files_stat.max_files && !capable(CAP_SYS_ADMIN)) {
```

Miklos

Containers mailing list

Containers@lists.linux-foundation.org

