
Subject: Re: [patch 3/8] account user mounts
Posted by [ebiederm](#) on Sun, 22 Apr 2007 07:49:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Miklos Szeredi <miklos@szeredi.hu> writes:

```
>> > 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.

Since it is perfectly possible to do a per user rlimit at this stage in the design. I contend that either:

- We implement a per user rlimit of mounts.
- We implement a global limit on mounts.

No other case makes sense. The previous objections were at least in part because the limit only applied to user mounts but the name of the limit did not apply to user mounts.

```
>> Also currently you are not checking the max_users. It looks like
```

>> you do this in a later patch but still it is a little strange to
>> allow user own mounts and have accounting but to not check the
>> limit at this state.

>

> Yeah, but at this stage user mounts are not yet allowed, so this is
> safe.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

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