
Subject: [patch 00/10] (resend) mount ownership and unprivileged mount syscall
Posted by [Miklos Szeredi](#) on Thu, 12 Apr 2007 16:45:41 GMT

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

This patchset adds support for keeping mount ownership information in the kernel, and allow unprivileged mount(2) and umount(2) in certain cases.

This can be useful for the following reasons:

- mount(8) can store ownership ("user=XY" option) in the kernel instead, or in addition to storing it in /etc/mtab. For example if private namespaces are used with mount propagations /etc/mtab becomes unworkable, but using /proc/mounts works fine
- fuse won't need a special suid-root mount/umount utility. Plain umount(8) can easily be made to work with unprivileged fuse mounts
- users can use bind mounts without having to pre-configure them in /etc/fstab

Unprivileged mounts are restricted to private namespaces created with a special clone flag.

Changes from the previous submission:

- add namespace flag for allowing user mounts
- add clone flag to set above namespace flag
- make max number of user mounts default to 1024, since now the namespace flag will prevent user mounts by default

--

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