Subject: [patch 00/10] mount ownership and unprivileged mount syscall (v3)
Posted by Miklos Szeredi on Mon, 16 Apr 2007 11:03:08 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

The following security measures are taken for unprivileged mounts:

 - only allow submounting under mounts which have a special mount flag set
 - only allow mounting on files/directories writable by the user
 - limit the number of user mounts
 - force "nosuid,nodev" mount options

Changes from the previous submissions:

 - add mount flags to set/clear mnt_flags individually
 - add "usermnt" mount flag. If it is set, then allow unprivileged
   submounts under this mount
 - make max number of user mounts default to 1024, since now the
   usermnt flag will prevent user mounts by default

--