Subject: Re: [patch 2/8] allow unprivileged umount
Posted by ebiederm on Sat, 21 Apr 2007 13:29:51 GMT
View Forum Message <> Reply to Message

Miklos Szeredi <miklos@szeredi.hu> writes:

> From: Miklos Szeredi <mszeredi@suse.cz>
>
> The owner doesn't need sysadmin capabilities to call umount().
>
> Similar behavior as umount(8) on mounts having "user=UID" option in
> /etc/mtab. The difference is that umount also checks /etc/fstab,
> presumably to exclude another mount on the same mountpoint.
>
> Signed-off-by: Miklos Szeredi <mszeredi@suse.cz>
> ---
>
> Index: linux/fs/namespace.c
> ===
> --- linux.orig/fs/namespace.c 2007-04-20 11:55:05.000000000 +0200
> +++ linux/fs/namespace.c 2007-04-20 11:55:06.000000000 +0200
> @@ -659,6 +659,25 @@ static int do_umount(struct vfsmount *mn
> }
>
> /*
> + * umount is permitted for
> + * - sysadmin
> + * - mount owner, if not forced umount
> + */
> +static bool permit_umount(struct vfsmount *mnt, int flags)
> +{
> + if (capable(CAP_SYS_ADMIN))
> + return true;
> +
> + if (!(mnt->mnt_flags & MNT_USER))
> + return false;
> +
> + if (flags & MNT_FORCE)
> + return false;
> +
> + return mnt->mnt_uid == current->uid;
> +}

I think this should be:

static bool permit_umount(struct vfsmount *mnt, int flags)
{
 if ((mnt->mnt_uid != current->fsuid) && !capable(CAP_SETUID))

```
        return false;

 if ((flags & MNT_FORCE) && !capable(CAP_SYS_ADMIN))
  return false;

 return true;
}
```

I.e.
  MNT_USER gone.
  compare against fsuid.
  Only require setuid for unmounts unless force is specified.

  I suspect we can allow MNT_FORCE for non-privileged users
  as well if we can trust the filesystem.

```
> +/*
>   * Now umount can handle mount points as well as block devices.
>   * This is important for filesystems which use unnamed block devices.
>   *
> @@ -681,7 +700,7 @@ asmlinkage long sys_umount(char __user *
>    goto dput_and_out;
>
>   retval = -EPERM;
> - if (!capable(CAP_SYS_ADMIN))
> + if (!permit_umount(nd.mnt, flags))
>    goto dput_and_out;
>
>   retval = do_umount(nd.mnt, flags);
>
> --
```