
Subject: Re: [patch 1/8] add user mounts to the kernel

Posted by [akpm](#) on Sat, 21 Apr 2007 07:55:03 GMT

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

On Fri, 20 Apr 2007 12:25:33 +0200 Miklos Szeredi <miklos@szeredi.hu> wrote:

> From: Miklos Szeredi <mszeredi@suse.cz>

>

> Add ownership information to mounts.

>

> A new mount flag, MS_SETUSER is used to make a mount owned by a user.

> If this flag is specified, then the owner will be set to the current

> real user id and the mount will be marked with the MNT_USER flag. On

> remount don't preserve previous owner, and treat MS_SETUSER as for a

> new mount. The MS_SETUSER flag is ignored on mount move.

So is a modified mount(8) needed? If so, is there some convenient way in which testers can get hold of it?

> The MNT_USER flag is not copied on any kind of mount cloning:

> namespace creation, binding or propagation. For bind mounts the

> cloned mount(s) are set to MNT_USER depending on the MS_SETUSER mount

> flag. In all the other cases MNT_USER is always cleared.

>

> For MNT_USER mounts a "user=UID" option is added to /proc/PID/mounts.

> This is compatible with how mount ownership is stored in /etc/mtab.

>

> Signed-off-by: Miklos Szeredi <mszeredi@suse.cz>

> ---

>

> Index: linux/fs/namespace.c

> =====

> --- linux.orig/fs/namespace.c 2007-04-20 11:55:02.000000000 +0200

> +++ linux/fs/namespace.c 2007-04-20 11:55:05.000000000 +0200

> @@ -227,6 +227,13 @@ static struct vfsmount *skip_mnt_tree(st

> return p;

> }

>

> +static void set_mnt_user(struct vfsmount *mnt)

> +{

> + BUG_ON(mnt->mnt_flags & MNT_USER);

> + mnt->mnt_uid = current->uid;

> + mnt->mnt_flags |= MNT_USER;

> +}

I'm a bit surprised to see this. Using uids in-kernel is all rather old-fashioned and restricted. I'd have expected

```
mnt->user = get_uid(current->user);
```

Containers mailing list

Containers@lists.linux-foundation.org

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