Subject: [patch] unprivileged mounts update
Posted by Miklos Szeredi on Wed, 25 Apr 2007 07:45:31 GMT
View Forum Message <> Reply to Message

From: Miklos Szeredi <mszeredi@suse.cz>

- refine adding "nosuid" and "nodev" flags for unprivileged mounts:
  o add "nosuid", only if mounter doesn't have CAP_SETUID capability
  o add "nodev", only if mounter doesn't have CAP_MKNOD capability

- allow unprivileged forced unmount, but only for FS_SAFE filesystems

- allow mounting over special files, but not symlinks

- for mounting and umounting check "fsuid" instead of "ruid"

Thanks to everyone for the comments, with special thanks to Serge
Hallyn and Eric Biederman.

For testing the new functionality provided by this patchset a simple
tool similar in syntax to mount(8) is available from:

  http://www.kernel.org/pub/linux/kernel/people/mszeredi/mmount

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

Index: linux/fs/namespace.c
===
--- linux.orig/fs/namespace.c 2007-04-22 17:48:18.000000000 +0200
+++ linux/fs/namespace.c 2007-04-22 18:19:51.000000000 +0200
@@ -252,10 +252,12 @@ static int reserve_user_mount(void)
 static void __set_mnt_user(struct vfsmount *mnt)
 {
  BUG_ON(mnt->mnt_flags & MNT_USER);
- mnt->mnt_uid = current->uid;
+ mnt->mnt_uid = current->fsuid;
  mnt->mnt_flags |= MNT_USER;
- if (!capable(CAP_SYS_ADMIN))
- mnt->mnt_flags |= MNT_NOSUID | MNT_NODEV;
+ if (!capable(CAP_SETUID))
+ mnt->mnt_flags |= MNT_NOSUID;
+ if (!capable(CAP_MKNOD))
+ mnt->mnt_flags |= MNT_NODEV;
 }

 static void set_mnt_user(struct vfsmount *mnt)
@@ -725,10 +727,10 @@ static bool permit_umount(struct vfsmoun

```
  if (!(mnt->mnt_flags & MNT_USER))
    return false;

- if (flags & MNT_FORCE)
+ if ((flags & MNT_FORCE) && !(mnt->mnt_sb->s_type->fs_flags & FS_SAFE))
    return false;

- return mnt->mnt_uid == current->uid;
+ return mnt->mnt_uid == current->fsuid;
 }

 /*
@@ -792,13 +794,13 @@ static bool permit_mount(struct nameidat
  if (type && !(type->fs_flags & FS_SAFE))
    return false;

- if (!S_ISDIR(inode->i_mode) && !S_ISREG(inode->i_mode))
+ if (S_ISLNK(inode->i_mode))
    return false;

  if (!(nd->mnt->mnt_flags & MNT_USER))
    return false;

- if (nd->mnt->mnt_uid != current->uid)
+ if (nd->mnt->mnt_uid != current->fsuid)
    return false;

  *flags |= MS_SETUSER;
```

_____

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