
Subject: [patch 02/10] unprivileged mounts: allow unprivileged umount
Posted by [Miklos Szeredi](#) on Wed, 16 Jan 2008 12:31:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

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>

Acked-by: Serge Hallyn <serue@us.ibm.com>

Index: linux/fs/namespace.c

```
=====
--- linux.orig/fs/namespace.c 2008-01-16 13:25:05.000000000 +0100
+++ linux/fs/namespace.c 2008-01-16 13:25:06.000000000 +0100
@@ -894,6 +894,27 @@ static int do_umount(struct vfsmount *mn
    return retval;
}

+static bool is_mount_owner(struct vfsmount *mnt, uid_t uid)
+{
+ return (mnt->mnt_flags & MNT_USER) && mnt->mnt_uid == uid;
+}
+
+/*
+ * 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 (flags & MNT_FORCE)
+ return false;
+
+ return is_mount_owner(mnt, current->fsuid);
+}
+
+/*
+ * Now umount can handle mount points as well as block devices.
+ * This is important for filesystems which use unnamed block devices.
```

```
@@ -917,7 +938,7 @@ asmlinkage long sys_umount(char __user *
goto dput_and_out;

retval = -EPERM;
- if (!capable(CAP_SYS_ADMIN))
+ if (!permit_umount(nd.path.mnt, flags))
goto dput_and_out;

retval = do_umount(nd.path.mnt, flags);

--
```

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