
Subject: [patch -mm 14/17] user namespace: maintain user ns for priv_userns mounts to vfsmount

Posted by [Cedric Le Goater](#) on Tue, 05 Dec 2006 10:28:06 GMT

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

From: Serge E. Hallyn <serue@us.ibm.com>

If a vfsmount is marked priv_userns, then a bind mount should maintain the old->user_ns. Not doing so presents a simple way around the priv_userns security mechanism :)

Signed-off-by: Serge E. Hallyn <serue@us.ibm.com>

fs/namespace.c | 2 ++
1 file changed, 2 insertions(+)

Index: 2.6.19-rc6-mm2/fs/namespace.c

=====

--- 2.6.19-rc6-mm2.orig/fs/namespace.c
+++ 2.6.19-rc6-mm2/fs/namespace.c
@@ -240,6 +240,8 @@ static struct vfsmount *clone_mnt(struct

```
if (mnt) {  
    mnt->mnt_flags = old->mnt_flags;  
+ if (mnt->mnt_flags & MNT_PRIV_USERNS)  
+ mnt->mnt_user_ns = old->mnt_user_ns;  
    atomic_inc(&sb->s_active);  
    mnt->mnt_sb = sb;  
    mnt->mnt_root = dget(root);  
}
```

--

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
