
Subject: [PATCH 2/3] user_ns: bugfix and cleanup
Posted by [serue](#) on Wed, 06 Dec 2006 23:19:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Serge E. Hallyn <serue@us.ibm.com>
Subject: [PATCH 2/3] user_ns: bugfix and cleanup

The first two updates are used in the next patchset:

1. change get_user_ns to return the namespace
2. change get_user_ns and put_user_ns to handle NULL input

This fixes a bug in the lxc patchset:

3. when clone_mnt uses the user_ns from the original vfsmnt (for MNT_PRIV_USERNS mounts), it needs to put the user_ns which was gotten in alloc_mnt, and get the new one.

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

```
fs/namespace.c          |  9 ++++++  
include/linux/user_namespace.h | 11 +++++++----  
2 files changed, 12 insertions(+), 8 deletions(-)
```

```
diff --git a/fs/namespace.c b/fs/namespace.c  
index 55fcaa7..fa52e24 100644  
--- a/fs/namespace.c  
+++ b/fs/namespace.c  
@@ -57,8 +57,7 @@ struct vfsmount *alloc_vfsmnt(const char  
    struct vfsmount *mnt = kmalloc_cache_alloc(mnt_cache, GFP_KERNEL);  
    if (mnt) {  
        memset(mnt, 0, sizeof(struct vfsmount));  
-       mnt->mnt_user_ns = current->nsproxy->user_ns;  
-       get_user_ns(mnt->mnt_user_ns);  
+       mnt->mnt_user_ns = get_user_ns(current->nsproxy->user_ns);  
        atomic_set(&mnt->mnt_count, 1);  
        atomic_set(&mnt->mnt_writers, 0);  
        INIT_LIST_HEAD(&mnt->mnt_hash);  
@@ -261,8 +260,10 @@ 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;  
+       if (mnt->mnt_flags & MNT_PRIV_USERNS) {  
+           put_user_ns(mnt->mnt_user_ns);  
+           mnt->mnt_user_ns = get_user_ns(old->mnt_user_ns);  
+       }
```

```

atomic_inc(&sb->s_active);
mnt->mnt_sb = sb;
mnt->mnt_root = dget(root);
diff --git a/include/linux/user_namespace.h b/include/linux/user_namespace.h
index f25c54c..f9d76de 100644
--- a/include/linux/user_namespace.h
+++ b/include/linux/user_namespace.h
@@ -17,9 +17,11 @@ extern struct user_namespace init_user_n

#endif CONFIG_USER_NS

-static inline void get_user_ns(struct user_namespace *ns)
+static inline struct user_namespace *get_user_ns(struct user_namespace *ns)
{
- kref_get(&ns->kref);
+ if (ns)
+ kref_get(&ns->kref);
+ return ns;
}

extern int unshare_user_ns(unsigned long unshare_flags,
@@ -29,12 +31,13 @@ extern void free_user_ns(struct kref *kr

static inline void put_user_ns(struct user_namespace *ns)
{
- kref_put(&ns->kref, free_user_ns);
+ if (ns)
+ kref_put(&ns->kref, free_user_ns);
}

#else

-static inline void get_user_ns(struct user_namespace *ns)
+static inline struct user_namespace *get_user_ns(struct user_namespace *ns)
{
}

--
```

1.4.1

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