
Subject: [RFC][PATCH 2/2] add the user namespace to the unshare syscall [try #2]
Posted by [Cedric Le Goater](#) on Mon, 28 Aug 2006 14:56:53 GMT

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

This patch allows a process to unshare its user namespace through the unshare() syscall.

Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>
Cc: Andrew Morton <akpm@osdl.org>
Cc: Kirill Korotaev <dev@openvz.org>
Cc: Eric W. Biederman <ebiederm@xmission.com>
Cc: Herbert Poetzl <herbert@13thfloor.at>
Cc: Serge E. Hallyn <serue@us.ibm.com>
Cc: Dave Hansen <haveblue@us.ibm.com>

kernel/fork.c | 21 ++++++-----
1 file changed, 17 insertions(+), 4 deletions(-)

Index: 2.6.18-rc4-mm3/kernel/fork.c

```
=====
--- 2.6.18-rc4-mm3.orig/kernel/fork.c
+++ 2.6.18-rc4-mm3/kernel/fork.c
@@ -1603,6 +1603,7 @@ asmlinkage long sys_unshare(unsigned lon
 struct nsproxy *new_nsproxy = NULL, *old_nsproxy = NULL;
 struct uts_namespace *uts, *new_uts = NULL;
 struct ipc_namespace *ipc, *new_ipc = NULL;
+ struct user_namespace *user, *new_user = NULL;

 check_unshare_flags(&unshare_flags);

@@ -1610,7 +1611,7 @@ asmlinkage long sys_unshare(unsigned lon
 err = -EINVAL;
 if (unshare_flags & ~(CLONE_THREAD|CLONE_FS|CLONE_NEWNS|CLONE_SIGHAND|
 CLONE_VM|CLONE_FILES|CLONE_SYSVSEM|
- CLONE_NEWUTS|CLONE_NEWIPC))
+ CLONE_NEWUTS|CLONE_NEWIPC|CLONE_NEWUSER))
     goto bad_unshare_out;

 if ((err = unshare_thread(unshare_flags)))
@@ -1631,18 +1632,20 @@ asmlinkage long sys_unshare(unsigned lon
     goto bad_unshare_cleanup_semundo;
 if ((err = unshare_ipcs(unshare_flags, &new_ipc)))
     goto bad_unshare_cleanup_uts;
+ if ((err = unshare_user_ns(unshare_flags, &new_user)))
+     goto bad_unshare_cleanup_ipc;

- if (new_ns || new_uts || new_ipc) {
```

```

+ if (new_ns || new_uts || new_ipc || new_user) {
    old_nsproxy = current->nsproxy;
    new_nsproxy = dup_namespaces(old_nsproxy);
    if (!new_nsproxy) {
        err = -ENOMEM;
    - goto bad_unshare_cleanup_ipc;
    + goto bad_unshare_cleanup_user;
    }
}

if (new_fs || new_ns || new_sigh || new_mm || new_fd || new_ulist ||
- new_uts || new_ipc) {
+ new_uts || new_ipc || new_user) {

    task_lock(current);

@@ @ -1696,12 +1699,22 @@ asmlinkage long sys_unshare(unsigned long
    new_ipc = ipc;
}

+ if (new_user) {
+ user = current->nsproxy->user_ns;
+ current->nsproxy->user_ns = new_user;
+ new_user = user;
+ }
+
    task_unlock(current);
}

if (new_nsproxy)
    put_nsproxy(new_nsproxy);

+bad_unshare_cleanup_user:
+ if (new_user)
+ put_user_ns(new_user);
+
bad_unshare_cleanup_ipc:
if (new_ipc)
    put_ipc_ns(new_ipc);

```

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