
Subject: [RFC] [PATCH 0/8] user namespaces: add ns to user_struct

Posted by [serue](#) on Fri, 07 Dec 2007 19:12:59 GMT

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

I'm working toward fixing up some of the remaining uid==0 and uid1==uid2 checks, and beginning to restrict capabilities within namespaces.

This patchset starts to do that by

1. improving per-ns user_struct storing
2. introducing CAP_NS_OVERRIDE
3. requiring CAP_NS_OVERRIDE to signal another user namespace
4. remove a few uid==0 checks

Especially the last 3 patches are a definite security improvement in the face of user namespaces.

The next steps would be

- * add user_ns to siginfo
- * signals delivered to another userns (like sigchld) send uid 0.
- * fix up more uid and gid checks (sigh)
- * convert struct key_user?
- * introduce uid aliases
- * per-process keyring
- * stores (user_ns,uid) keys
- * allows process which is really (user_ns1, uid1) to act as though it were (user_ns2, uid2) on objects in user_ns2
- * convert struct kstat (may have serious lifetime issues)

That should leave us in a reasonable shape to start considering how to really handle file access.

I still have a set of patches which tag struct inode with user_ns and patch ext2+ext3. But it's at the end of my patch set for now.

Comments welcome, on these patches, on the outlined next steps, or on anything I'm forgetting.

(Against 2.6.24-rc3-mm2)

thanks,
-serge

Containers mailing list

