
Subject: [PATCH 0/4] user namespaces: introduction
Posted by [serue](#) on Mon, 28 Jan 2008 19:08:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Here is a small patchset I've been sitting on for awhile to make signaling mostly subject to user namespaces. In particular,

1. store user_namespace in user struct
2. introduce CAP_NS_OVERRIDE
3. require CAP_NS_OVERRIDE to signal another user namespace

The first step should have been done all along. Else wouldn't a hash collision on (ns1, uid) and (ns2, uid), however unlikely, give us wrong results at uid_hash_find()?

The main remaining signaling+usersns issue is of course the siginfo. Tacking a usersns onto siginfo is a pain due to lifetime mgmt issues. I haven't decided whether to just catch all the callers and fake uid=0 if user namespaces aren't the same, introduce some unique non-refcounted id to represent (user,user_ns), or find some other way to deal with it.

thanks,
-serge

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