Subject: Re: + user-ns-implement-user-ns-unshare-remove-config_user_ns.patch added to -mm tree
Posted by Cedric Le Goater on Fri, 26 Jan 2007 14:40:31 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> "Serge E. Hallyn" <serue@us.ibm.com> writes:
>
>> Allow me to re-ask a fundamental question: do we want the uid namespace
>> to stick to turning uid checks into (uid,ns) checks? or do we want the
>> uid namespaces to try to protect against root in other namespaces?
>
> I am fairly certain we want to at least make the checks (uid, ns) checks.
> That gives a minimal level of protection against root in other namespaces,
> as the lesser root does not match the (uid, ns) check for the system root.

I agree that we need the (uid, ns) checks. I would say that these are the next steps to complete the user namespace feature and remove its experimental status. But I'm glad that the kernel supports the initial framework, it should make it easier to fill in the gap. IMO.

Note that we did the opposite with the pid namespace, clean up the kernel before doing the framework, but the framework is also much more complex to get right.

C.

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