Subject: Re: + user-ns-implement-user-ns-unshare-remove-config_user_ns.patch added to -mm tree
Posted by ebiederm on Thu, 25 Jan 2007 19:38:31 GMT
View Forum Message <> Reply to Message

"Serge E. Hallyn" <serue@us.ibm.com> writes:

>> As it sits right now using the user namespace instead of being an
>> enhancement of security as it should feels like security loophole
>> 101.
>
> That's a bit of a callous exaggeration, don't you think? It takes what
> used to be one big pool and partitions it, offering isolation between
> partitions with one clearly defined exception. Given that there used to
> be no partitions at all, this still nets added isolation over the
> original, no loopholes.

Except from what I could tell from my quick review the partition is far
from complete. That concerns me. Especially the way this is expected
to be used it to setup a user who appears to be root in his partition.

The fact that for any non-filesystem based permission check he appears
to be root to the rest of the system disturbs me.

Now maybe I'm blind but that is my current perception of the situation
and until the partition is complete I don't want people thinking it is.

The little minor details about your exception are important but not
really what concerned me at this point.

>> For the stable namespaces I'm fine with removing the config options
>> but the user namespace is not at all complete.
>
> I'm fine with that, it certainly is still very new.
>
> I was liking Cedric's patch because with the CONFIG_USER_NS option,
> we just end up with all the more corner cases to test, so Cedric's
> patch, considering how little ends up actually CONFIG'd out, actually
> seems an improvement.

As long as we retain the ability to compile out the ability to actually
create a second user namespace I am happy. I have no problem with
running all of the code with just the initial user namespace present.

My apologies for not being able to have the security conversation yet
and not having done a more thorough review.

The truth is that coming anywhere near the user namespace I find scary because of all it's security implications. Almost by definition any bug there is a security issue.

If we limit ourselves to exactly allowing overlapping uid and gids by making the comparisons for equality include their corresponding namespace, and we look carefully to ensure we get every such comparison. I feel fairly comfortable in saying we have simply added a slightly more sophisticated form of naming users but the current security model remains. If we deviate from that one iota I figure we need to completely think through the new rules very carefully as we have changed the rules and sometimes innocuous changes combined to allow unexpected loop holes.

I think it is healthy to be paranoid about security issues, isn't it?

So in summary my only real complaint with removing CONFIG_USER_NS is that it appears to me that the code is incomplete and has not been closely scrutinized. As such making it available to end users without even a warning when that is the case appears irresponsible. Especially as much of the code that is sitting in Andrews tree is merged into the production kernel, when the window opens.

Eric

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