Subject: Re: + user-ns-implement-user-ns-unshare-remove-config_user_ns.patch added to -mm tree
Posted by serue on Thu, 25 Jan 2007 20:46:35 GMT
View Forum Message <> Reply to Message

Quoting Eric W. Biederman (ebiederm@xmission.com):
> "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.

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?

If we go with the first, we can always enforce protection against root
in other namespaces using LSMs. SELinux users have what they need, and
others can use a trivial new LSM.

Eric? Herbert? Kirill? Cedric? Anyone with an opinion?

-serge

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