Subject: Re: + user-ns-implement-user-ns-unshare-remove-config_user_ns.patch added to -mm tree
Posted by serue on Fri, 26 Jan 2007 16:37:56 GMT
View Forum Message <> Reply to Message

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

Well we do that, the question is whether we want to stick to just that.

> That gives a minimal level of protection against root in other namespaces,

Not really.

> as the lesser root does not match the (uid, ns) check for the system root.

But it will pass capable(CAP_DAC_OVERRIDE) and such checks.

> Exactly how capabilities play into this I'm not quite certain, but something
> important to understand. Especially for suid root executables.
>
> > 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.
>
> What is the hole you see with root in other namespaces that needs an
> LSM, the only hole I know of currently is the incomplete state of the
> (uid/gid, ns) checks.

The hole is that most permission checks are of the form

```
 if (uid1 == uid2 || capable(CAP_DAC_OVERRIDE))
  allow permission;
```

The root user in a vserver needs CAP_DAC_OVERRIDE within his own userns,
so we can't just take that away.

My patch is one way to handle that for files. Another, arguably
cleaner approach, would be to not handle the problem with user
namespaces, but do so with security modules. But that does feel
like a leak across user namespaces.

> Not that I don't think an LSM couldn't improve the situation.
> Although if I have to deal with the LSM insanity much more I'm going
> to lobby for changing the concept it to an interapplication firewall,

A whatzit?

> and get all of the stupid code into the kernel.

It sounds like you're having an interesting problem - would love to
see it explained on the lsm or selinux list.

-serge

_____

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