

Jan Kara <jack@suse.cz> writes:

> There can be arbitrary number of listeners (potentially from different
> namespaces if I understand it correctly) listening to broadcasts. So I
> think we should pass some universal identifier rather than try to find out
> who is listening etc. I think such identifiers would be useful for other
> things too, won't they?

So internal to the kernel we have such a universal identifier.
struct user.

There are two practical questions.

- 1) How do we present that information to user space?
- 2) How does user space want to process this information?

If we only want user space to be able to look up a user and send him a message. It probably makes sense to do the struct user to uid conversion in the proper context in the kernel because we have that information.

If this is a general feature that happens to allow us to look up the user given the filesystems view of what is going on would be easier in the kernel, and not require translation. But it means that we can't support 9p and nfs for now. But since we don't support quotas on the client end anyway that doesn't sound like a big deal.

The problem with the filesystem view is that there will be occasions where we simply can not map a user into it, because the filesystem won't have a concept of that particular user.

So we could run into the situation where alice owns the file. Bob writes to the file and pushes it over quota. But the filesystem has no concept of who bob is. So we won't be able to report that it was bob that pushed things over the edge.

> BTW: Do you have some idea, when would be the infrastructure clearer?

So the plan is to get to the point where are uid comparisons in the kernel are (user namespace, uid) comparisons. Or possibly struct user comparisons (depending on the context. And struct mount will contain the user namespace of whoever mounted the filesystem.

Adding infrastructure to netlink to allow us to do conversions as the packets are enqueued for a specific user is something I would rather avoid, but that is a path we can go down if we have

to.

> Whether it makes sense to currently proceed with UIDs and later change it
> to something generic or whether I should wait before you sort it out :).

A good question. I think things are clear enough that it at least makes sense to sketch a solution to the problem even if we don't implement it at this point.

I have been hoping Cedric or Serge would jump in because I think those are the guys who have been working on the implementation.

Eric

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