Subject: Re: [PATCH] Send quota messages via netlink
Posted by serge on Thu, 30 Aug 2007 19:10:10 GMT
View Forum Message <> Reply to Message

Quoting Jan Kara (jack@suse.cz):
> On Wed 29-08-07 15:06:43, Eric W. Biederman wrote:
> > Jan Kara <jack@suse.cz> writes:
> > >> However I'm still confused about the use of current->user. If that
> > >> is what we really want and not the user who's quota will be charged
> > >> it gets to be a really trick business, because potentially the uid
> > >> we want to deliver varies depending on who opened the netlink socket.
> > > I see it's a complicated matter :). What I need to somehow pass to
> > > userspace is something (and I don't really care whether it will be number,
> > > string or whatever) that userspace can read and e.g. find a terminal
> > > window or desktop the affected user has open and also translate the
> > > identity to some user-understandable name (average user Joe has to
> > > understand that he should quickly cleanup his home directory ;).
> > > Thinking more about it, we could probably pass a string to userspace in
> > > the format:
> > > <namespace type>:<user identification>
> > >
> > > So for example we can have something like:
> > > unix:1000 (traditional unix UIDs)
> > > nfs4:joe@machine
> > >
> > > The problem is: Are we able to find out in which "namespace type" we are
> > > and send enough identifying information from a context of unpriviledged
> > > user?
> >
> > Ok. This provides enough context to understand what you are trying to do.
> > You do want the unix user id, not the filesystem notion. Because you
> > are looking for the user.
> >
> > So we have to figure out how to do the hard thing which is look at
> > who opened our netlink broadcast see if they are in the same user
> > namespace as current->user. Which is a pain and we don't currently
> > have the infrastructure for.
> There can be arbitrary number of listeners (potentially from different
> namespaces if I understand it correctly) listening to broadcasts. So I

Currently that is true, but i think isolating netlink sockets is going
to have to be done pretty soon.

On the one hand cloning a new netlink socket ns when you unshare
CLONE_NEWNET may seem 'obvious', but I think doing so when you unshare
CLONE_NEWUSER make much more sense considering netlink's use for audit
and now for quota.

> think we should pass some universal identifier rather than try to find out

Even with isolating netlink we still may want to send out an identifier.
However, just as with mounts extensions we're printing out the memory
address of vfsmounts, we might just want to print out the memory address
of the userns. It's not universal, but should be good enough.

-serge

> who is listening etc. I think such identifiers would be useful for other
> things too, won't they?
>   BTW: Do you have some idea, when would be the infrastructure clearer?
> Whether it makes sence to currently proceed with UIDs and later change it
> to something generic or whether I should wait before you sort it out :).
>
>   Honza
> --
> Jan Kara <jack@suse.cz>
> SuSE CR Labs
> -
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
> Please read the FAQ at http://www.tux.org/lkml/

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