Subject: Re: [PATCH 1/6] user namespace : add the framework
Posted by serue on Mon, 16 Jul 2007 14:38:48 GMT
View Forum Message <> Reply to Message

Quoting Serge E. Hallyn (serue@us.ibm.com):
> Quoting Andrew Morton (akpm@linux-foundation.org):
> > On Mon, 4 Jun 2007 14:40:24 -0500 "Serge E. Hallyn" <serue@us.ibm.com> wrote:
> >
> > > Add the user namespace struct and framework
> > >
> > > Basically, it will allow a process to unshare its user_struct table, resetting
> > > at the same time its own user_struct and all the associated accounting.
> > >
> > > A new root user (uid == 0) is added to the user namespace upon creation. Such
> > > root users have full privileges and it seems that theses privileges should be
> > > controlled through some means (process capabilities ?)
> >
> > The whole magical-uid-0-user thing in this patch seem just wrong to
> > me.
> >
> > I'll merge it anyway, mainly because I want to merge _something_ (why oh
> > why do the git-tree guys leave everything to the last minute?) but it strikes
> > me that there's something fundamentally wrong whenever the kernel starts
> > "knowing" about the significance of UIDs in this fashion.
>
> $(&(%
>
> I thought I disagreed, but now I'm pretty sure I completely agree.
>
> 'root_user' exists in the kernel right now, but the root_user checks
> which exist (in fork.c and sys.c) shouldn't actually be applied for root
> in a container, since the container may not be trusted.

By the way, I don't think these two uses of 'user == &root_user' are
legitimate. CAP_SYS_RESOURCE should probably be checked for both. Hah,
it already is in fork.c, in addition to root_user check.

But I guess I should take that up elsewhere.

thanks,
-serge

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