Subject: Re: [PATCH 2/2] hijack: update task_alloc_security
Posted by serue on Wed, 28 Nov 2007 14:57:33 GMT
View Forum Message <> Reply to Message

Quoting Crispin Cowan (crispin@crispincowan.com):
> Serge E. Hallyn wrote:
> > Quoting Casey Schaufler (casey@schaufler-ca.com):
> >
> >> Could y'all bring me up to speed on what this is intended to
> >> accomplish so that I can understand the Smack implications?
> >>
> > It's basically like ptracing a process, forcing it to fork, then having
> > the child execute a file and continue as your child. It takes part of
> > its state from the current process (stack etc), some from the hijacked
> > process (namespaces, keys?), and an lsm can decide for itself whose ->security
> > should be used for the child process.
> >
> That just doesn't gob smack me with the obvious abstract intention of
> this API :)
>
> So it is like I want to run a process inside a name space, but I am not
> inside that name space, so I highjack one that is in there, force it to
> fork, and then give me its child. Ugh.

Well that's how it started, but in reality you are the one who forks,
and you only copy namespace and some related data.

So what you are objecting to is the process you'd have to go through if
we don't have hijack.

> Couldn't we just implement "put me in that namespace over there?" AFAIK

There are two patchsets out there to do that, and there are standing
objections to them. Really doing it at clone is the safest thing,
and with pid namespaces the only sane thing to do.

> namespaces don't actually have names, making it hard to implement "put
> me in namespace Foo", but I view that as a defect of namespaces that
> should be fixed, rather than hacked around.

You can name them using the ns cgroup.

thanks,
-serge