Subject: Re: [RFC] ns containers (v2): namespace entering
Posted by serue on Wed, 21 Feb 2007 21:04:01 GMT
View Forum Message <> Reply to Message

Quoting Eric W. Biederman (ebiederm@xmission.com):
> "Serge E. Hallyn" <serue@us.ibm.com> writes:
>
> > Quoting Eric W. Biederman (ebiederm@xmission.com):
> >>
> >> You miss an issue here. One of the dangers of enter is leaking
> >> capabilities into a contained set of processes. Once you show up in
> >
> > Good point. As wrong as it feels to me to use ptrace for this, the
> > advantage is that none of my task attributes leak into the target
> > namespace, and that's a very good thing.
> >
> > I do wonder how you specify what the forced clone should run.
> > Presumably you want to run something not in the target container.
> > I suppose we can pass the fd over a socket or something.
>
> Yes. At least in the case without a network namespace I can setup
> a unix domain socket and pass file descriptors around. I think my solution
> to the network namespace case was to just setup a unix domain socket in
> the parent namespace and leave it open in init. Not a real solution :(

How about we solve both this and the general ugliness of using ptrace
with a new

 hijack_and_clone(struct task_struct *tsk, int fd)

Which takes tsk, clones it, and execs the contents of fd?

-serge

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