
Subject: Re: [PATCH 2/2] hijack: update task_alloc_security

Posted by [serue](#) on Mon, 03 Dec 2007 14:50:12 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Crispin Cowan (crispin@crispincowan.com):

> Serge E. Hallyn wrote:

> > Quoting Crispin Cowan (crispin@crispincowan.com):

> >

> >> I find that ptrace, specifically CAP_SYS_PTRACE, is overloaded. AppArmor

> >> is having problems because we have to choose between granting

> >> cap_sys_ptrace, or not allowing the process to read /proc/pid/self &

> >> such like. So there, the problem is that we have to grant too much power

> >> to a process to just let it read some /proc stuff about itself.

> >>

> >> Here the problem appears to be the other way. cap_sys_ptrace is powerful

> >> enough to mess with other user's processes on the system, but if ptrace

> >> gives you hijack, then that seems to give you the power to control

> >> processes in someone else's namespace.

> >>

> > The user namespace patchset I'm working on right now to start having

> > signals respect user namespaces introduces CAP_NS_OVERRIDE. Once that

> > is in, then hijack would require CAP_NS_OVERRIDE|CAP_SYS_PTRACE.

> >

> > Of course, since we're considering only allowing HIJACK_NS which is

> > only allowed into a different namespace, hijack would then always

> > require CAP_NS_OVERRIDE...

> >

> > Does that suffice?

> >

> I think that CAP_NS_OVERRIDE|CAP_SYS_PTRACE is a problem because of the

Oops, yeah I meant &.

> | making ptrace more powerful than it is now. If you make it

> CAP_NS_OVERRIDE only, then the problem goes away.

I was seeing CAP_NS_OVERRIDE as more of a modifier to other capabilities. So (CAP_NS_OVERRIDE&CAP_KILL) means you can signal processes in another namespace.

(CAP_NS_OVERRIDE&CAP_DAC_OVERRIDE) means you can override (currently nonexistent) DAC file access checks in another user namespace.

(CAP_NS_OVERRIDE&CAP_MKNOD) means you can create devices which your container isn't allowed to create.

Yup, this is all somewhat looking ahead...

-serge

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