
Subject: Re: [PATCH 1/2] namespaces: introduce sys_hijack (v10)

Posted by [Stephen Smalley](#) on Tue, 27 Nov 2007 18:09:24 GMT

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

On Tue, 2007-11-27 at 10:11 -0600, Serge E. Hallyn wrote:

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

> > Just the name "sys_hijack" makes me concerned.

> >

> > This post describes a bunch of "what", but doesn't tell us about "why"

> > we would want this. What is it for?

>

> Please see my response to Casey's email.

>

> > And I second Casey's concern about careful management of the privilege

> > required to "hijack" a process.

>

> Absolutely. We're definately still in RFC territory.

>

> Note that there are currently several proposed (but no upstream) ways to

> accomplish entering a namespace:

>

> 1. bind_ns() is a new pair of syscalls proposed by Cedric. An

> nsproxy is given an integer id. The id can be used to enter

> an nsproxy, basically a straight current->nsproxy = target_nsproxy;

>

> 2. I had previously posted a patchset on top of the nsproxy

> cgroup which allowed entering a nsproxy through the ns cgroup

> interface.

>

> There are objections to both those patchsets because simply switching a

> task's nsproxy using a syscall or file write in the middle of running a

> binary is quite unsafe. Eric Biederman had suggested using ptrace or

> something like it to accomplish the goal.

>

> Just using ptrace is however not safe either. You are inheriting *all*

> of the target's context, so it shouldn't be difficult for a nefarious

> container/vserver admin to trick the host admin into running something

> which gives the container/vserver admin full access to the host.

I don't follow the above - with ptrace, you are controlling a process already within the container (hence in theory already limited to its container), and it continues to execute within that container. What's the issue there?

> That's where the hijack idea came from. Yes, I called it hijack to make

> sure alarm bells went off :) bc it's definately still worrisome. But at

> this point I believe it is the safest solution suggested so far.

--

Stephen Smalley
National Security Agency

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