
Subject: Re: [PATCH] namespaces: introduce sys_hijack (v4)

Posted by [serue](#) on Tue, 16 Oct 2007 19:12:00 GMT

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

Quoting Paul Menage (menage@google.com):

> On 10/16/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

> >

> > Currently every pid namespace's pid==1 must stick around as long as the
> > pid namespace does. If you kill the pid==1, all processes in the
> > container are killed.

>

> What about people who aren't using pid namespaces?

Not really isolated? :)

> > > Anyway, I can go ahead and add 'int which' to the parameter list now,
> > > and leave the details of how to specify a cgroup for later. That way at
> > > least the api won't fundamentally change again.

> > >

> > > Great, thanks.

> >

> > Since the goal here is to get the API right, do you know how we expect
> > to send the cgroup in? A string?

>

> My thought was to use an fd on an open cgroup directory - that can be
> trivially translated into a cgroup.

Oh good, so I can just pass in a single arg id, so

```
asmlinkage long sys_hijack(unsigned long clone_flags, int which,  
                          unsigned long id);
```

?

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
