

---

Subject: Re: [RFC] [PATCH 2/2] namespace enter: introduce sys\_hijack (v3)

Posted by [serue](#) on Tue, 04 Sep 2007 12:50:19 GMT

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

---

Quoting Paul Menage (menage@google.com):

> On 8/29/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

> > > From: aec05999084bf3a94add66e98462652ed9408f86 Mon Sep 17 00:00:00 2001

> > From: sergeh@us.ibm.com <sergeh@us.ibm.com>

> > Date: Wed, 22 Aug 2007 15:03:57 -0700

> > Subject: [RFC] [PATCH 2/2] namespace enter: introduce sys\_hijack (v3)

> >

> > Introduce sys\_hijack (for x86 only). It is like clone, but in

> > place of a stack pointer (which is assumed null) it accepts a

> > pid. The process identified by that pid is the one which is

> > actually cloned. Some state - include the file table, the signals

> > and sighand (and hence tty), and the ->parent are taken from the

> > calling process.

>

> What do you do if there are no processes in a particular container?

The nsproxy will have been released so you couldn't enter it anyway.

> I prefer your suggestion of tying this to the nsproxy subsystem - that

> would allow you to spawn a child with a given set of namespaces, even

> if there were no appropriate process to hijack.

I can resend my original ns\_container entering patchset (maybe next week when everyone is back from summits) and we can discuss whether or not it is safe, or how to improve it if it is not.

thanks,

-serge

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---