
Subject: Re: [RFC] [PATCH 2/2] namespace enter: introduce sys_hijack (v3)

Posted by [Dave Hansen](#) on Tue, 04 Sep 2007 19:06:18 GMT

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

On Tue, 2007-09-04 at 07:50 -0500, Serge E. Hallyn wrote:

> > 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.

Yeah, we'd need some kind of other object to keep the nsproxy around and hold a reference to it.

But, it also begs other questions about how we define the namespace boundaries vs. containers. What if we have a normal container with chroot'd process inside of it? Two such processes will not share an nsproxy because the chroot'd one has switched filesystem namespaces.

Who is to say that the "container" is represented by one process's nsproxy more than another?

-- Dave

Containers mailing list

Containers@lists.linux-foundation.org

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