Subject: Re: semantics for namespace naming
Posted by ebiederm on Fri, 15 Dec 2006 20:47:56 GMT
View Forum Message <> Reply to Message

Dave Hansen <haveblue@us.ibm.com> writes:

> On Thu, 2006-12-14 at 14:56 -0700, Eric W. Biederman wrote:
>> Because that model fundamentally keeps every process in it's own
>> container and never allows it to leave, nor does it allow things
>> from one container to cross into another container in an uncontrolled
>> fashion this feels to me like a very safe model.
>
> This is like saying that brain surgery is safe and controlled because
> the surgeon never actually goes into the patient's brain! :)

Can you think how dangerous brain surgery would be if the surgeon
actually physically went into the patients brain.

> I think of ptrace as a pretty wide-open interface. While ptrace itself
> has well-defined semantics, I could hardly consider using it in
> production, nor would I want to be the one to write the userspace apps
> to do the syscall futzing for each of Linux's architectures.

Well that isn't exactly what I am proposing. What I am proposing is
that we compare any interface to what you could do with ptrace.   If
it allows for something ptrace doesn't allow you likely have a
problem.

So I think the concept of mapping the semantics of a new interface
to the semantics of ptrace is a very interesting review exercise.

Plus thinking about ptrace changes the question from what new
interface do we add to get the semantics we want, to how do we
optimize what we can do with ptrace, so it doesn't suck to use.

Eric

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers