
Subject: Re: [PATCH 11/15] Signal semantics
Posted by [dev](#) on Thu, 02 Aug 2007 08:35:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

Serge E. Hallyn wrote:

> Quoting Pavel Emelyanov (xemul@openvz.org):
>
>>[snip]
>>
>>
>>>>| Maybe it's worth disabling cross-namespaces ptracing...
>>>>
>>>>I think so too. Its probably not a serious limitation ?
>>>>
>>>>Several people think we will implement 'namespace entering' through a
>>>>ptrace hack, where maybe the admin ptraces the init in a child pidns,
>>
>>Why not implement namespace entering w/o any hacks? :)
>
>
> I did, as a patch on top of the nsproxy container subsystem. The
> response was that that is a hack, and ptrace is cleaner :)
>
> So the current options for namespace entering would be:
>
> * using Cedric's bind_ns() functionality, which assigns an
> integer global id to a namespace, and allows a process to
> enter a namespace by that global id

looks more or less good and what OVZ actually does.
So I would prefer this one.

> * using my nsproxy container subsystem patch, which lets
> a process enter another namespace using
> echo pid > /container/some/cont/directory/tasks
> and eventually might allow construction of custom
> namespaces, i.e.
> mkdir /container/c1/c2
> ln -s /container/c1/c1/network /container/c1/c2/network
> echo \$\$ > /container/c1/c2/tasks

Sound ok and logical as well.

> * using ptrace to coerce a process in the target namespace
> into forking and executing the desired program.

you'll need to change ptrace interface in this case imho...
doesn't sound ok at all... at least for me. So I agree with Pavel.

>>> makes it fork, and makes the child execute what it wants (i.e. ps -ef).
>>>
>>> You're talking about killing that functionality?
>>
>> No. We're talking about disabling the things that are not supposed
>> to work at all.
>
>
> Uh, well in the abstract that sounds like a sound policy...

Pavel simply meant that no one plans to disable functionality in question.

Thanks,
Kirill

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