Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by ebiederm on Wed, 07 Mar 2007 22:32:48 GMT
View Forum Message <> Reply to Message

"Paul Menage" <menage@google.com> writes:

> On 3/7/07, Serge E. Hallyn <serue@us.ibm.com> wrote:
>>
>> All that being said, if it were going to save space without overly
>> complicating things I'm actually not opposed to using nsproxy, but it
>
> If space-saving is the main issue, then the latest version of my
> containers patches uses just a single pointer in the task_struct, and
> all tasks in the same set of containers (across all hierarchies) will
> share a single container_group object, which holds the actual pointers
> to container state.

Yes.

However:
> Effectively, container_group is to container as nsproxy is to namespace.

The statement above nicely summarizes the confusion in terminology.

In the namespace world when we say container we mean roughly at the level
of nsproxy and container_group. Although it is expected to be a user space
concept like an application, not a concept implemented directly in the kernel.
i.e. User space is expected to combine separate resource controls and namespaces
and run processes inside that combination.

You are calling something that is on par with a namespace a container. Which
seriously muddies the waters. About as much as calling as referring to your
shoe as your whole outfit.

Without fixing the terminology it is going to be very hard to
successfully communicate.

Eric

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