Subject: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by Glauber Costa on Thu, 26 Jul 2012 09:16:22 GMT

View Forum Message <> Reply to Message

On 07/26/2012 07:57 AM, Eric W. Biederman wrote:
> Glauber Costa <glommer@parallels.com> writes:
>
>> I just came up with the following preliminary list of sessions:
>>
>> http://wiki.linuxplumbersconf.org/2012:containers
>>
>> Since people mostly said what they wanted to talk about, but without
>> extensive descriptions, I took the liberty of coming up with a small
>> text for each in the blueprints. If you believe this is inaccurate, or
>> would like to see it extended (although I personally don't see the point
>> about going into very formal and deep details here), just let me know
>> and I will edit it.
>>
>> This is all still subject to change.
>
> Something that just came up recently and worth looking at if it hasn't
> already be resolved.
>
> The network namespace, the user namespace, and the memory control group
> are not meshing well.
>
> In particular we need some additional checks for an unprivileged user
> who can set tcp_mem. If you are the creator of a network namespace you
> should at least be able to set the values down. I don't know at all
> about increasing the amount of memory consumed by the tcp stack.

This is between the user namespace and net namespace only, right ?

To be quite honest, I haven't looked thoroughly at UNS after your last
work. How do you yourself believe this should be?

>
> The non-nesting nature of memory control groups with respect to the
> network stack also seems very bizarre.

Correction:

The non-nesting nature of memory control groups is very bizarre. No need
for modifiers. It does support nesting, though. Just that it is
selectable, and not the default. But there is work in progress to change
that.

>

> Another old issue is that unless I have missed something control groups
> are still broken for generic use in containers. Does anyone care?
> Are there any plans on fixing this issue?
>
> Eric
>