Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Paul Menage on Tue, 13 Mar 2007 02:25:48 GMT
View Forum Message <> Reply to Message

On 3/12/07, Herbert Poetzl <herbert@13thfloor.at> wrote:
>
> why? you simply enter that specific space and
> use the existing mechanisms (netlink, proc, whatever)
> to retrieve the information with _existing_ tools,

That's assuming that you're using network namespace virtualization, with each group of tasks in a separate namespace. What if you don't want the virtualization overhead, just the accounting?

Paul

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

Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Herbert Poetzl on Tue, 13 Mar 2007 15:57:06 GMT
View Forum Message <> Reply to Message

On Mon, Mar 12, 2007 at 07:25:48PM -0700, Paul Menage wrote:
> On 3/12/07, Herbert Poetzl <herbert@13thfloor.at> wrote:
> >
> > why? you simply enter that specific space and
> > use the existing mechanisms (netlink, proc, whatever)
> > to retrieve the information with _existing_ tools,
>
> That's assuming that you're using network namespace virtualization,

or isolation :)

> with each group of tasks in a separate namespace.

correct ...

> What if you don't want the virtualization overhead, just the
> accounting?

there should be no 'virtualization' overhead, and what
do you want to account for, if not by a group of tasks?

maybe I'm missing the grouping condition here, but I
assume you assign tasks to the accounting containers

note: network isolation is not supposed to add overhead
compared to the host system (at least not measureable
overhead)

best,
Herbert

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

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

Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Paul Menage on Sat, 17 Mar 2007 01:40:23 GMT
View Forum Message <> Reply to Message

On 3/13/07, Herbert Poetzl <herbert@13thfloor.at> wrote:
>
> note: network isolation is not supposed to add overhead
> compared to the host system (at least not measureable
> overhead)

Does "network isolation" involve giving each network namespace a
separate IP address, etc? If so, that's something that some people
might not want (and specifically, that we don't want). If not, what
does it add over existing Linux traffic control (via tc)?

Thanks,

Paul

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

---