
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>
