
Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [Paul Menage](#) on Thu, 08 Mar 2007 00:42:05 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 3/7/07, Sam Vilain <sam@vilain.net> wrote:

> Paul Menage wrote:

> >> In the namespace world when we say container we mean roughly at the level
> >> of nsproxy and container_group.

> >>

> > So you're saying that a task can only be in a single system-wide container.

> >

>

> Nope, we didn't make the mistake of nailing down what a "container" was

> too far before it is implemented. We talked before about

> containers-within-containers because, inevitably if you provide a

> feature you'll end up having to deal with virtualising systems that in

> turn use that feature.

Sure, my aproach allows containers hierarchically as children of other containers too.

>

> > My patch provides multiple potentially-independent ways of dividing up

> > the tasks on the system - if the "container" is the set of all

> > divisions that the process is in, what's an appropriate term for the

> > sub-units?

> >

>

> namespace, since 2.4.x

>

> > That assumes the viewpoint that your terminology is "correct" and

> > other people's needs "fixing". :-)

> >

>

> Absolutely. Please respect the semantics established so far; changing

> them adds nothing at the cost of much confusion.

But "namespace" has well-established historical semantics too - a way of changing the mappings of local names to global objects. This doesn't describe things liek resource controllers, cpusets, resource monitoring, etc.

Trying to extend the well-known term namespace to refer to things that aren't namespaces isn't a useful approach, IMO.

Paul

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