
Subject: Re: [PATCH 0/7] containers (V7): Generic Process Containers
Posted by [ebiederm](#) on Tue, 20 Feb 2007 17:34:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

> On 2/12/07, Sam Vilain <sam@vilain.net> wrote:
>>
>> I know I'm a bit out of touch, but AUI the NSProxy **is** the container.
>> We decided a long time ago that a container was basically just a set of
>> namespaces, which includes all of the subsystems you mention.
>
> You may have done that, but the CKRM/ResGroups independently decided a
> long time ago that the fundamental unit was the resource class, and
> the OpenVZ folks decided that the fundamental unit was the
> BeanCounter, and the CPUSet folks decided that the fundamental unit
> was the CPUSet, etc ... :-)

Using the container name is bad and it led to this stupid argument.

The fundamental unit of what we have merged into the kernel is the namespace. The aggregate of all namespaces and everything is the container.

Please, please pick a different name so people don't take one look at your stuff and get confused like Sam did.

> But there's a lot of common ground between these different approaches,
> and potential for synergy, so the point of this patch set is to
> provide a unification point for all of them, and a stepping stone for
> other new resource controllers and process control modules.

For the case of namespaces I don't see how your code makes things better. I do not see a real problem that you are solving.

I do agree that a common interface to the code and a common set of infrastructure may help.

Personally until the pid namespace (which includes within it process groupings) is complete I'm not certain the problem will be concrete enough to solve. Hopefully we can have a working version of that shortly.

Eric

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