

---

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

Posted by [ebiederm](#) on Thu, 08 Mar 2007 01:32:10 GMT

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

---

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

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

>> But "namespace" has well-established historical semantics too - a way  
>> of changing the mappings of local \* to global objects. This  
>> accurately describes things like resource controllers, cpusets, resource  
>> monitoring, etc.

>

> Sorry, I think this statement is wrong, by the generally established  
> meaning of the term namespace in computer science.

>

>>

>> Trying to extend the well-known term namespace to refer to things that  
>> are semantically equivalent namespaces is a useful approach, IMHO.

>>

>

> Yes, that would be true. But the kinds of groupings that we're talking  
> about are supersets of namespaces, not semantically equivalent to  
> them. To use Eric's "shoe" analogy from earlier, it's like insisting  
> that we use the term "sneaker" to refer to all footwear, including ski  
> boots and birkenstocks ...

Pretty much. For most of the other cases I think we are safe referring to them as resource controls or resource limits. I know that roughly covers what cpusets and beancounters and ckrm currently do.

The real trick is that I believe these groupings are designed to be something you can setup on login and then not be able to switch out of. Which means we can't use sessions and process groups as the grouping entities as those have different semantics.

Eric

---

Containers mailing list

[Containers@lists.osdl.org](mailto:Containers@lists.osdl.org)

<https://lists.osdl.org/mailman/listinfo/containers>

---