

---

Subject: Re: [ckrm-tech] containers development plans  
Posted by [Paul Menage](#) on Tue, 10 Jul 2007 07:53:20 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 7/10/07, Paul Jackson <pj@sgi.com> wrote:

> Kirill, Serge, et al,

>

> Is it fair to say then that Paul Menage's containers are primarily  
> for the purposes of managing resources, while namespaces are for the  
> purposes of managing identifiers?

Sort of - but one thing that we're trying to figure out how to do nicely is integrate namespaces into the container filesystem (this was the purpose of the `post_clone()` container API callback) so that we can both get a filesystem view of task namespaces, and combine namespaces with other process container subsystems.

>

> We've got some resources, like cpu cycles, memory bytes, network  
> bandwidth, that we want to allocate and account for differentially  
> by groups of tasks -- that's Menage's containers.

Plus things that aren't necessarily resource controllers, such as the container freezer, or permissions on network ports, or userspace OOM handlers. I don't think that lumping all of these in together as "resource containers" is the right thing to do.

> virtualization efforts, of which my `cpusets` is the granddaddy example,  
> being generalized by Paul Menage with his container patches. The other  
> work is, as Serge actually termed it in the body of his post, better  
> called 'namespaces'.

Purely within the kernel, yes. The more general encompassing effort to have a combined kernel/userspace solution for virtual servers is also referred to as "containers". (And to be fair that term was already in use when I started using the term "process containers" to refer to the specific framework in the kernel that handles process tracking).

Paul

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---