

---

Subject: Re: [PATCH 0/7] containers (V7): Generic Process Containers

Posted by [serue](#) on Mon, 12 Feb 2007 22:47:17 GMT

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

---

Quoting Sam Vilain (sam@vilain.net):

> menage@google.com wrote:

> > Generic Process Containers

> > -----

> >

> > There have recently been various proposals floating around for  
> > resource management/accounting and other task grouping subsystems in  
> > the kernel, including ResGroups, User BeanCounters, NSProxy  
> > containers, and others. These all need the basic abstraction of being  
> > able to group together multiple processes in an aggregate, in order to  
> > track/limit the resources permitted to those processes, or control  
> > other behaviour of the processes, and all implement this grouping in  
> > different ways.

> >

>

> I know I'm a bit out of touch, but AUI the NSProxy *is* the container.

The nsproxy is an implementation detail, actually, and might go away at any time. But yes we had been talking for quite awhile about sets of namespaces as containers, so that a vserver would be a container, as would a lightweight checkpoint/restart job.

> We decided a long time ago that a container was basically just a set of  
> namespaces, which includes all of the subsystems you mention.

>

> This would suggesting re-write this patchset, part 2 as a "CPUSet

Well it's an unfortunate conflict, but I don't see where we have any standing to make Paul change his terminology :)

Lately I've been talking about "Paul's containers", our "namespaces", and the "namespace container subsystem" in patch 7.

> namespace", part 4 as a "CPU scheduling namespace", parts 5 and 6 as  
> "Resource Limits Namespace" (drop this "BeanCounter" brand), and of  
> course part 7 falls away.

Part 7 does not fall away, it uses Paul's patchset to start to provide a management interface for namespaces.

-serge

---

Containers mailing list

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

