

---

Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction  
Posted by [Chandra Seetharaman](#) on Wed, 20 Sep 2006 19:35:23 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2006-09-20 at 12:25 -0700, Paul Menage wrote:

> On 9/20/06, Chandra Seetharaman <sekharan@us.ibm.com> wrote:

> >

> > We had this discussion more than 18 months back and concluded that it is  
> > not the right thing to do. Here is the link to the thread:

>

> Even if the resource control portions aren't totally compatible,  
> having two separate process container abstractions in the kernel is  
> sub-optimal, both in terms of efficiency and userspace management. How  
> about splitting out the container portions of cpuset from the actual  
> resource control, so that CKRM/RG can hang off of it too? Creation of  
> a cpuset or a resource group would be driven by creation of a  
> container; at fork time, a task inherits its parent's container, and  
> hence its cpuset and/or resource groups.

>

> At its most crude, this could be something like:

>

```
> struct container {  
> #ifdef CONFIG_CPUSETS  
>   struct cpuset cs;  
> #endif  
> #ifdef CONFIG_RES_GROUPS  
>   struct resource_group rg;  
> #endif  
> };
```

Won't it restrict the user to choose one of these, and not both.

It will also prevent the possibility of having resource groups within a cpuset.

>

> but at least it would be sharing some of the abstractions.

>

> Paul

>

> -----

> Take Surveys. Earn Cash. Influence the Future of IT

> Join SourceForge.net's Techsay panel and you'll get the chance to share your

> opinions on IT & business topics through brief surveys -- and earn cash

> [http://www.techsay.com/default.php?page=join.php&p=sourc eforge&CID=DEVDEV](http://www.techsay.com/default.php?page=join.php&p=sourc%20eforge&CID=DEVDEV)

>

> \_\_\_\_\_  
> ckrm-tech mailing list

> <https://lists.sourceforge.net/lists/listinfo/ckrm-tech>

--

-----  
Chandra Seetharaman | Be careful what you choose....  
- sekharan@us.ibm.com | .....you may get it.  
-----

---

Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction  
Posted by [Paul Menage](#) on Wed, 20 Sep 2006 19:57:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 9/20/06, Chandra Seetharaman <sekharan@us.ibm.com> wrote:

> > At its most crude, this could be something like:

> >

> > struct container {

> > #ifdef CONFIG\_CPUSETS

> > struct cpuset cs;

> > #endif

> > #ifdef CONFIG\_RES\_GROUPS

> > struct resource\_group rg;

> > #endif

> > };

>

> Won't it restrict the user to choose one of these, and not both.

Not necessarily - you could have both compiled in, and each would only worry about the resource management that they cared about - e.g. you could use the memory node isolation portion of cpusets (in conjunction with fake numa nodes/zones) for memory containment, but give every cpuset access to all CPUs and control CPU usage via the resource groups CPU controller.

The generic code would take care of details like container creation/destruction (with appropriate callbacks into cpuset and/or res\_group code, tracking task membership of containers, etc.

Paul

---