
Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by [Paul Menage](#) on Wed, 20 Sep 2006 20:17:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 9/20/06, Paul Jackson <pj@sgi.com> wrote:

> Paul M. wrote:

> > Rather than adding a new process container abstraction, wouldn't it
> > make more sense to change cpuset to make it more extensible (more
> > separation between resource controllers), possibly rename it to
> > "containers",
>

> Without commenting one way or the other on the overall advisability
> of this (for lack of sufficient clues), if we did this and renamed
> "cpusets" to "containers", we would still want to export the /dev/cpuset
> interface to just the CPU/Memory controllers. Perhaps the "container"
> pseudo-filesystem could optionally be mounted with a "cpuset" option,
> that just exposed the cpuset relevant interface, or some such thing.

Absolutely - I was thinking that as a first cut, any subsystem (e.g. cpusets, res_groups, etc) that wanted to use per-task containers could declare what files it wanted a container dir populated with, so you could have it looking just like cpusets if you wanted to, and mount it on /dev/cpuset and use it exactly as before. If you then added the res_group patch to your kernel, you would also get the appropriate resource group files appearing in each directory, but the cpuset support would work as before.

Longer term we'd probably want to figure out a better naming partitioning scheme, or maybe just a convention that each directory entry was prefixed with the subsystem name. Also, maybe have a convention that control files and subcontainer names be in different namespaces (e.g. all control files start with ".", all subcontainer names start with something else).

Paul
