Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by Paul Menage on Wed, 20 Sep 2006 19:25:15 GMT
View Forum Message <> Reply to Message

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
};
```

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

Paul