
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 cpuset (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
