
Subject: Re: [ckrm-tech] containers development plans
Posted by [Paul Menage](#) on Tue, 10 Jul 2007 06:19:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 7/9/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:
> > - splitting the memory and cpu isolation parts of cpusets into two
> > separate subsystems (still backwards-compatible)
>
> I see memory isolation using cpusets as very topology dependent
> and I am not sure if the model would work for memory controllers.

I wasn't suggesting making any changes to the page-based memory controllers as part of this.

Currently in the mainline kernel, the cpumask and nodemask portions of cpusets are essentially two mostly-independent modules that happen to be coupled together in the same file and use the same process tracking system (cpusets). Once we have generic process containers, splitting this into a "cpusets" subsystem that handles all the cpumask portions of the existing cpusets, and a "memsets" subsystem that handles all the nodemask and memory migration portions would remove that coupling and give more flexibility.

Paul
