
Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by [Paul Menage](#) on Wed, 20 Sep 2006 17:37:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 9/20/06, Rohit Seth <rohitseth@google.com> wrote:

>
> cpusets provides cpu and memory NODES binding to tasks. And I think it
> works great for NUMA machines where you have different nodes with its
> own set of CPUs and memory. The number of those nodes on a commodity HW
> is still 1. And they can have 8-16 CPUs and in access of 100G of
> memory. You may start using fake nodes (untested territory) to

I've been experimenting with this, and it's looking promising.

>
> Containers also provide a mechanism to move files to containers. Any
> further references to this file come from the same container rather than
> the container which is bringing in a new page.

This could be done with memory nodes too - a vma can specify its
memory policy, so binding individual files to nodes shouldn't be hard.

>
> In future there will be more handlers like CPU and disk that can be
> easily embeded into this container infrastructure.

I think that at least the userspace API for adding more handlers would
need to be defined before actually committing a container patch, even
if the kernel code isn't yet extensible. The CKRM/RG interface is a
good start towards this.

Paul
