Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction Posted by Paul Menage on Thu, 21 Sep 2006 01:52:30 GMT

View Forum Message <> Reply to Message

On 9/20/06, Chandra Seetharaman < sekharan@us.ibm.com> wrote:

>

- > Interesting. So you could set up the fake node with "guarantee" and let > it grow till "limit" ?
- Sure that works great. (Theoretically you could do this all in

Sure - that works great. (Theoretically you could do this all in userspace - start by assigning "guarantee" nodes to a container/cpuset and when it gets close to its memory limit assign more nodes to it. But in practice userspace can't keep up with rapid memory allocators.

>

- > BTW, can you do these with fake nodes:
- > dynamic creation
- > dynamic removal
- > dynamic change of size

The current fake numa support requires you to choose your node layout at boot time - I've been working with 64 fake nodes of 128M each, which gives a reasonable granularity for dividing a machine between multiple different sized jobs.

>

- > Also, How could we account when a process moves from one node to > another?
- If you want to do that (the systems I'm working on don't really) you could probably do it with the migrate_pages() syscall. It might not be that efficient though.

Paul