
Subject: Re: [ckrm-tech] [patch00/05]: Containers(V2)- Introduction
Posted by [Paul Menage](#) on Thu, 21 Sep 2006 00:42:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 9/20/06, Paul Jackson <pj@sgi.com> wrote:

> Chandra wrote:

> > AFAICS, That doesn't help me in over committing resources.

>

> I agree - I don't think cpuset's plus fake numa ... handles over commit.

> You might could hack up a cheap substitute, but it wouldn't do the job.

I have some patches locally that basically let you give out a small set of nodes initially to a cpuset, and if memory pressure in `try_to_free_pages()` passes a specified threshold, automatically allocate one of the parent cpuset's unused memory nodes to the child cpuset, up to specified limit. It's a bit ugly, but lets you trade of performance vs memory footprint on a per-job basis (when combined with fake numa to give lots of small nodes).

Paul
