
Subject: Re: [ckrm-tech] [PATCH 4/7] UBC: syscalls (user interface)

Posted by [dev](#) on Mon, 21 Aug 2006 13:33:19 GMT

[View Forum Message](#) <> [Reply to Message](#)

Andrew Morton wrote:

> I have this mad idea that you can divide a 128GB machine up into 256 fake
> NUMA nodes, then you use each "node" as a 512MB unit of memory allocation.
> So that 4.5GB job would be placed within an exclusive cpuset which has nine
> "mems" (what are these called?) and voila: the job has a hard 4.5GB limit,
> no kernel changes needed.

this doesn't allow memory overcommitment, does it?

Kirill
