
Subject: Re: [ckrm-tech] [PATCH 4/7] UBC: syscalls (user interface)
Posted by [Magnus Damm](#) on Mon, 21 Aug 2006 08:42:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 2006-08-21 at 09:48 +0200, Andi Kleen wrote:

> > You may be looking for the NUMA emulation patches posted here:
> >
> > <http://marc.theaimsgroup.com/?l=linux-mm&m=112806587501884&w=2>
> >
> > There is a slightly updated x86_64 version here too:
> >
> > <http://marc.theaimsgroup.com/?l=linux-mm&m=113161386520342&w=2>
>
> Hmm, I must have missed that version. Seems like a improvement. Best you
> resubmit it, although I'll probably only take it after the .19 merge

No problem. The second URL pointed to a x86_64 version where I tried to break out code to make some kind of generic NUMA emulation layer. At that time no one seemed interested in that strategy as a simple resource control solution so I gave that up.

For x86_64 I think it's only worth mucking around with the code if people believe that it is the right way to go for in-kernel resource control.

The x86_64 patches above include code to divide each real NUMA node into several smaller emulated nodes, but that is kind of pointless if people only use it for non-resource control purposes, ie just to play with CPUSETS and NUMA on non-NUMA hardware. For simple purposes like that I think the existing NUMA emulation code for x86_64 works perfectly well.

I still think that i386 users would benefit from NUMA emulation though. If you want me to up-port the i386-specific code just let me know.

Thanks,

/ magnus
