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

Posted by [Paul Jackson](#) on Fri, 18 Aug 2006 21:19:45 GMT

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

Chris wrote:

> Hypothetically, if you can guarantee that those threads get a specified
> amount of time, but may possibly get *more* cpu time and thus finish
> faster, what's the problem?

Each thread is already getting ~100% of one dedicated CPU. It can't get any *more* than that ;). 256 active threads, 256 CPUs, *very* tightly coupled, in this example, on say a 512 CPU system running unrelated stuff on the other half.

If anything else interferes with the memory bandwidth of even one thread to its node local memory, or takes any of that node local memory for unrelated uses, it's a big problem. A few percent variation in job runtime, for a two day job, means hours difference in the finish time. People notice, when that's your 'big money' app.

This is not your fathers PC ;).

--

I won't rest till it's the best ...
Programmer, Linux Scalability
Paul Jackson <pj@sgi.com> 1.925.600.0401