On 1/11/07, Balbir Singh <balbir@in.ibm.com> wrote:
>
> I tried something similar, I added an activated field, which is set
> to true when the ->create() callback is invoked. That did not help
> either, the machine still panic'ed.

I think that marking it active when create() is called may be too soon.

Is this with my unchanged cpuacct subsystem, or with the version that
you've extended to track load over defined periods? I don't see it
when I test under VMware (with two processors in the VM), but I
suspect that's not going to be quite as parallel as a real SMP system.

>
> I see the need for it, but I wonder if we should start with that
> right away. I understand that people might want to group cpusets
> differently from their grouping of let's say the cpu resource
> manager. I would still prefer to start with one hierarchy and then
> move to multiple hierarchies. I am concerned that adding complexity
> upfront might turn off people from using the infrastructure.

That's what I had originally and people objected to the lack of flexibility :-)

The presence or absence of multiple hierarchies is pretty much exposed
to userspace, and presenting the right interface to userspace is a
fairly important thing to get right from the start.

Paul