

---

Subject: Re: containers (was Re: -mm merge plans for 2.6.23)

Posted by [Ingo Molnar](#) on Wed, 11 Jul 2007 09:04:23 GMT

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

---

\* Andrew Morton <[akpm@linux-foundation.org](mailto:akpm@linux-foundation.org)> wrote:

> > c. Enable group scheduling bits only in -mm for now (2.6.23-mmX), using  
> > Paul's container patches. I can send you a short patch that hooks up  
> > cfs group scheduler with Paul's container infrastructure.

> >

> > If a. is not possible, I would prefer c.

> >

> > Let me know your thoughts ..

>

> I'm inclined to take the cautious route here - I don't think people

> will be dying for the CFS thingy (which I didn't even know about?) in

> .23, and it's rather a lot of infrastructure to add for a CPU

> scheduler configurator gadget (what does it do, anyway?)

>

> We have plenty of stuff for 2.6.23 already ;)

>

> Is this liveable with??

another option would be to trivially hook up CONFIG\_FAIR\_GROUP\_SCHED with cpuset, and to offer CONFIG\_FAIR\_GROUP\_SCHED in the Kconfig, dependent on CPUSETS and defaulting to off. That would give it a chance to be tested, benchmarked, etc.

Ingo

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---