
Subject: Re: [RFC] [PATCH 0/3] Add group fairness to CFS
Posted by [William Lee Irwin III](#) on Wed, 23 May 2007 19:46:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

* William Lee Irwin III <wli@holomorphy.com> wrote:

>> [...] As an interface it may be poor and worse yet poorly specified,
>> [...]

On Wed, May 23, 2007 at 09:26:54PM +0200, Ingo Molnar wrote:

> that's the only thing that matters to fundamental design questions like
> this.

I'm not sure where it comes in as a design question. Tuning whatever's done for sched_yield() is just one of the drudgery tasks of rounding out the implementation as I see it. It can be painful depending on what's being changed, though it's not clear to me precisely how painful tuning it will be for this change.

If it turns out to be a large enough issue, I can enforce reasonable semantics for it with no impact on the core scheduling algorithm anyway.

* William Lee Irwin III <wli@holomorphy.com> wrote:

>> The content of my comment was that the patch does something to
>> sched_yield() semantics, so it raises the question of what will happen
>> in benchmarks and other performance affairs that are sensitive to
>> sched_yield() semantics changes.

On Wed, May 23, 2007 at 09:26:54PM +0200, Ingo Molnar wrote:

> the correct approach to the "sys_sched_yield() is an API that sucks"
> problem is to simply _not use it_. User-space is figuring that out now,
> fortunately.

We do need API's to displace it for that to really take off. I think the yield_to() API's will help there. Sadly, I expect it to be a very long time for the transition to really happen.

-- wli

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
