

---

Subject: Re: [RFC] [PATCH 0/3] Add group fairness to CFS  
Posted by [William Lee Irwin III](#) on Wed, 23 May 2007 19:04:48 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

\* William Lee Irwin III <[wli@holomorphy.com](mailto:wli@holomorphy.com)> wrote:  
>> [...] sched\_yield() semantics are yet another twist.

On Wed, May 23, 2007 at 08:40:35PM +0200, Ingo Molnar wrote:  
> that's nonsense, sched\_yield() semantics are totally uninteresting. It  
> is a fundamentally broken interface.

They're not totally uninteresting. People will complain when apps break or perform poorly due to changes in its semantics. As an interface it may be poor and worse yet poorly specified, but it has non-negligible effects on performance issues that can't be ignored and that will remain the case for the foreseeable future.

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.

-- wli

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: [RFC] [PATCH 0/3] Add group fairness to CFS  
Posted by [Ingo Molnar](#) on Wed, 23 May 2007 19:26:54 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

\* William Lee Irwin III <[wli@holomorphy.com](mailto:wli@holomorphy.com)> wrote:

> \* William Lee Irwin III <[wli@holomorphy.com](mailto:wli@holomorphy.com)> wrote:  
> >> [...] sched\_yield() semantics are yet another twist.

>  
> On Wed, May 23, 2007 at 08:40:35PM +0200, Ingo Molnar wrote:  
> > that's nonsense, sched\_yield() semantics are totally uninteresting. It  
> > is a fundamentally broken interface.

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

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

> 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.

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.

Ingo

---

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

---