
Subject: Re: [ckrm-tech] [RFC] [PATCH 0/3] Add group fairness to CFS
Posted by [William Lee Irwin III](#) on Wed, 30 May 2007 20:13:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sat, May 26, 2007 at 08:41:12AM -0700, William Lee Irwin III wrote:
>> The smpnice affair is better phrased in terms of task weighting. It's
>> simple to honor nice in such an arrangement. First unravel the
>> grouping hierarchy, then weight by nice. This looks like
[...]
>> conveniently collapse to 1). Now that the hierarchy is flattened,
>> nice numbers can be factored in for t_1's final weight being
>> $0.7*0.36/(0.7*0.36+0.3*0.24+0.7*0.24+0.3*0.16) = 0.252/0.54 = 0.467..$
>> and the others being 0.133.. (t_2), 0.311.. (t_3), and 0.0889.. (t_4).

On Wed, May 30, 2007 at 10:44:05PM +0530, Srivatsa Vaddagiri wrote:
> Hmm ..so do you think this weight decomposition can be used to flatten
> the tree all the way to a single level in case of cfs? That would mean we can
> achieve group fairness with single level scheduling in cfs ..I am
> somewhat skeptical that we can achieve group fairness with a single
> level rb-tree (and w/o substantial changes to pick_next_task logic in cfs
> that is), but if it can be accomplished would definitely be a great win.

Yes, the hierarchy can be flattened completely and global task weights
computed and used to achieve group fairness. The changes aren't to
pick_next_task() but rather to the ->fair_key computations. In fact, I
went a step beyond that.

On Sat, May 26, 2007 at 08:41:12AM -0700, William Lee Irwin III wrote:
>> In such a manner nice numbers obey the principle of least surprise.

The step beyond was to show how nice numbers can be done with all that
hierarchical task grouping so they have global effects instead of
effects limited to the scope of the narrowest grouping hierarchy
containing the task. I had actually assumed the weighting and
flattening bits were already in your plans from some other post you
made and was building upon that.

-- wli

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