
Subject: Re: [ckrm-tech] [RFC] [PATCH 0/3] Add group fairness to CFS
Posted by [Srivatsa Vaddagiri](#) on Thu, 31 May 2007 08:56:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, May 31, 2007 at 02:03:53PM +0530, Srivatsa Vaddagiri wrote:

> Its ->wait_runtime will drop less significantly, which lets it be
> inserted in rb-tree much to the left of those 1000 tasks (and which indirectly
> lets it gain back its fair share during subsequent schedule cycles).
>
> Hmm ..is that the theory?

My only concern is the time needed to converge to this fair distribution, especially in face of fluctuating workloads. For ex: a container who does a fork bomb can have a very adverse impact on other container's fair share under this scheme compared to other schemes which dedicate separate rb-trees for different containers (and which also support two level hierarchical scheduling inside the core scheduler).

I am inclined to have the core scheduler support atleast two levels of hierarchy (to better isolate each container) and resort to the flattening trick for higher levels.

--

Regards,
vatsa

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