
Subject: Re: [RFC | PATCH 0/9] CPU controller over process container
Posted by [Srivatsa Vaddagiri](#) on Thu, 12 Apr 2007 18:19:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Apr 12, 2007 at 07:56:47PM +0200, Herbert Poetzl wrote:

> > - Each task-group gets its own runqueue on every cpu.
>
> how does that scale for, let's say 200-300 guests on a
> 'typical' dual CPU machine?

Scheduling complexity is still $O(1)$ and hence I would say CPU-wise, it should be very scalable. Memory-wise, I agree that this can consume more memory if number of guests are large ..But this I feel is a memory vs cpu tradeoff ..If you had only one queue in which tasks from all groups were present, then it increases the schedule() complexity ?

If there are specific tests you had in mind to test this scalability aspect, I would be happy to run them.

> > - In addition, there is an active and expired array of
> > task-groups themselves. Task-groups that have expired their
> > quota are put into expired array.
>
> how much overhead does that add to the scheduler, cpu
> and memory wise?

cpu-wise, it should add very little overhead (since $O(1)$ behavior is retained). memory-wise, same points as above.

> > - Scheduling the next task involves picking highest priority task-group
> > from active array first and then picking highest-priority task
> > within it. Both steps are $O(1)$.
>
> how does that affect interactivity?

Note that I define task-group priority = highest priority tasks the group has, which IMO should give decent if not good interactivity ..But is (good) interactivity a big requirement here? As we know, that's a hard thing to achieve even in today's $O(1)$ scheduler ..

> > - SMP load-balancing is accomplished on the lines of smpnice.
>
> what about strict CPU limits (i.e. 20% regardless of
> the idle state of the machine)

Not supported in these patches. Any idea how/where that would be usefull?

--

Regards,
vatsa

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