
Subject: Re: [PATCH] introduce task cgroup v2
Posted by [Balbir Singh](#) on Sat, 21 Jun 2008 08:56:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On Sun, Jun 15, 2008 at 6:32 PM, KOSAKI Motohiro
> <kosaki.motohiro@jp.fujitsu.com> wrote:
>> honestly, I used res_counter on early version.
>> but I got bad performance.
>
> Bad performance on the charge/uncharge?
>
> The only difference I can see is that res_counter uses
> spin_lock_irqsave()/spin_unlock_irqrestore(), and you're using plain
> spin_lock()/spin_unlock().
>
> Is the overhead of a pushf/cli/popf really going to matter compared
> with the overhead of forking/exiting a task?
>
> Or approaching this from the other side, does res_counter really need
> irq-safe locking, or is it just being cautious?

We really need irq-safe locking. We can end up uncharging from reclaim context (called under zone->lru_lock and mem->zone->lru_lock - held with interrupts disabled)

I am going to convert the spin lock to a reader writers lock, so that reads from user space do not cause contention. I'll experiment and look at the overhead.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

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