
Subject: Re: [PATCH] introduce task cgroup v2
Posted by [Paul Menage](#) on Sat, 21 Jun 2008 15:48:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sat, Jun 21, 2008 at 2:10 AM, KOSAKI Motohiro
<kosaki.motohiro@jp.fujitsu.com> wrote:

>
> I am going to convert spinlock in task limit cgroup to atomic_t.
> task limit cgroup has following caractatics.
> - many write (fork, exit)
> - few read
> - fork() is performance sensitive syscall.

This is true, but I don't see how it can be more performance-sensitive than the overhead of allocating/freeing a page.

What kinds of performance regressions did you see?

> if increase fork overhead, system total performance cause degression.

What kind of overhead were you seeing? How about if you delay doing any task accounting until the task_limit subsystem is bound to a hierarchy? That way there's no noticeable overhead for people who aren't using your subsystem.

Paul

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