
Subject: Re: [PATCH, v3 2/2] cgroups: introduce timer slack subsystem
Posted by [Kirill A. Shutemov](#) on Mon, 07 Feb 2011 11:06:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, Feb 04, 2011 at 09:27:55AM -0800, Jacob Pan wrote:

> On Fri, 4 Feb 2011 15:34:39 +0200
> "Kirill A. Shutemov" <kirill@shutemov.name> wrote:
> > What's mean "original timer slack" if you are free to move a task
> > between a lot of cgroups and process itself free to change it anytime?
> >
>
> I need to manage tasks by a management software instead of letting the
> task change timer_slack by itself. The goal is to make management
> transparent and no modifications to the existing apps. Therefore, it is
> desirable to automatically enforce timer_slack when the apps are in the
> cgroup while automatically restore it when it is no longer under cgroup
> management.

Tasks are always under cgroup management. Root cgroup is still cgroup.

> So the "original timer slack" can be the default 50us or whatever value
> chosen by the task itself. But the app itself should not care or even be
> aware of which cgroup it is in.
>
> So here are two options i can think of
> 1. add a new variable called cg_timer_slack_ns to struct task_struct{}
> cg_timer_slack_ns will be set by cgroup timer_slack subsystem, then we
> can retain the original per task value in timer_slack_ns.
> timer code will pick max(cg_timer_slack_ns, timer_slack_ns) if
> cg_timer_slack_ns is set.
>
> 2. leave task_struct unchanged, add a current_timer_slack to the
> cgroup. timer_slack cgroup does not modify per task timer_slack_ns.
> similar to option #1, let timer code pick the timer_slack to use based
> on whether the task is in timer_slack cgroup.
>
> Any thoughts?

I think it's over-engineering.

What about configuration like this:

```
root_cgroup
|--timer_slack.min_slack_ns = 0
|--timer_slack.max_slack_ns = ULONG_MAX
|--50us
| |--timer_slack.min_slack_ns = 50000
| |--timer_slack.max_slack_ns = 50000
```

```
|--500us  
|--timer_slack.min_slack_ns = 500000  
|--timer_slack_max_slack_ns = ULONG_MAX
```

If you want a task allow to drive its timer_slack, just leave it in root_cgroup.

If you want to drive timer_slack of a task, move it between 50us and 500um based on your policy.

--

Kirill A. Shutemov

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