
Subject: Re: [PATCH, v3 2/2] cgroups: introduce timer slack subsystem

Posted by [jacob.jun.pan](#) on Wed, 02 Feb 2011 22:56:05 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2 Feb 2011 22:47:36 +0200

"Kirill A. Shutemov" <kirill@shutemov.name> wrote:

> From: Kirill A. Shutemov <kirill@shutemov.name>

>

> Provides a way of tasks grouping by timer slack value. Introduces per
> cgroup max and min timer slack value. When a task attaches to a
> cgroup, its timer slack value adjusts (if needed) to fit min-max
> range.

>

> It also provides a way to set timer slack value for all tasks in the
> cgroup at once.

>

> This functionality is useful in mobile devices where certain
> background apps are attached to a cgroup and minimum wakeups are
> desired.

>

> Signed-off-by: Kirill A. Shutemov <kirill@shutemov.name>

> Idea-by: Jacob Pan <jacob.jun.pan@linux.intel.com>

> ---

> include/linux/cgroup_subsys.h | 6 +

> include/linux/init_task.h | 4 +-

> init/Kconfig | 10 ++

> kernel/Makefile | 1 +

> kernel/cgroup_timer_slack.c | 242

> ++++++

> +

> +static int cgroup_timer_slack_check(struct task_struct *task,

> + unsigned long slack_ns)

> +{

> + struct cgroup_subsys_state *css;

> + struct timer_slack_cgroup *tslack_cgroup;

> +

use rcu_read_lock()/unlock?

lockdep complains unprotected rcu dereference check.

> + css = task_subsys_state(task, timer_slack_subsys.subsys_id);

> + tslack_cgroup = container_of(css, struct timer_slack_cgroup,

> css); +

> + if (slack_ns < tslack_cgroup->min_slack_ns)

> + return -EPERM;

> + if (slack_ns > tslack_cgroup->max_slack_ns)

```
> + return -EPERM;
> + return 0;
> +}
> +

> +
> +static struct cftype files[] = {
> + {
> + .name = "set_slack_ns",
> + .write_u64 = tslack_write_set_slack_ns,
> + },
> + },
should we also allow reading of the current slack_ns?
```

```
> + {
> + .name = "min_slack_ns",
> + .private = TIMER_SLACK_MIN,
> + .read_u64 = tslack_read_range,
> + .write_u64 = tslack_write_range,
> + },
> + {
> + .name = "max_slack_ns",
> + .private = TIMER_SLACK_MAX,
> + .read_u64 = tslack_read_range,
> + .write_u64 = tslack_write_range,
> + },
> +};
> +
```

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