
Subject: Re: [RFC][PATCH] introduce task cgroup (#task restrictioon for prevent fork bomb by cgroup)

Posted by [KAMEZAWA Hiroyuki](#) on Thu, 05 Jun 2008 05:09:46 GMT

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

On Thu, 05 Jun 2008 13:43:06 +0900

KOSAKI Motohiro <kosaki.motohiro@jp.fujitsu.com> wrote:

```
> +static int task_cgroup_max_tasks_write(struct cgroup *cgrp,
> +      struct cftype *cftype,
> +      s64 max_tasks)
> +{
> +    struct task_cgroup *taskcg;
> +
> +    if ((max_tasks > INT_MAX) ||
> +        (max_tasks < INT_MIN))
> +    return -EINVAL;
```

should be (max_tasks > INT_MAX) || (max_tasks < -1)) ?

```
> +
> + taskcg = task_cgroup_from_cgrp(cgrp);
> +
> + spin_lock(&taskcg->lock);
> + if (max_tasks < taskcg->nr_tasks)
> + return -EBUSY;
> + taskcg->max_tasks = max_tasks;
> + spin_unlock(&taskcg->lock);
```

This will cause dead lock.

And it seems this doesn't handle "attach failure".

It will be helpful Documentation somewhere.

Thanks,
-Kame

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
