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

Posted by [Dhaval Giani](#) on Thu, 05 Jun 2008 09:27:45 GMT

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

```
> +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))
```

It should be < -1 I think.

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

Shouldn't you drop the lock here?

```
> +    taskcg->max_tasks = max_tasks;
> +    spin_unlock(&taskcg->lock);
> +
> +    return 0;
> +}
```

How does this controller affect performance? Do you have some numbers?

Thanks,

--

regards,
Dhaval

Containers mailing list

Containers@lists.linux-foundation.org

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