

---

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

Posted by [Daisuke Nishimura](#) on Thu, 05 Jun 2008 05:23:20 GMT

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

---

Hi.

I think this would be usefull.

On Thu, 05 Jun 2008 13:43:06 +0900, KOSAKI Motohiro <kosaki.motohiro@jp.fujitsu.com> wrote:

> Index: b/kernel/cgroup\_task.c

> =====

> --- /dev/null

> +++ b/kernel/cgroup\_task.c

> @@ -0,0 +1,185 @@

> +/\* cgroup\_task.c - #task control group

> + \*

> + \* Copyright: KOSAKI Motohiro <kosaki.motohiro@jp.fujitsu.com>

> + \*

> + \* This program is free software; you can redistribute it and/or modify

> + \* it under the terms of the GNU General Public License as published by

> + \* the Free Software Foundation; either version 2 of the License, or

> + \* (at your option) any later version.

> + \*

> + \* This program is distributed in the hope that it will be useful,

> + \* but WITHOUT ANY WARRANTY; without even the implied warranty of

> + \* MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the

> + \* GNU General Public License for more details.

> + \*/

> +

> +#include <linux/res\_counter.h>

I don't think it's needed.

Or, are you planning to implement this feature by using res\_counter?

> +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;

> +

> + taskcg = task\_cgroup\_from\_cgrp(cgrp);

> +

> + spin\_lock(&taskcg->lock);

> + if (max\_tasks < taskcg->nr\_tasks)

```
> + return -EBUSY;  
need spin_unlock().
```

Thanks,  
Daisuke Nishimura.

---

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

---