
Subject: Re: [BUG]: Crash with CONFIG_FAIR_CGROUP_SCHED=y
Posted by [Dmitry Adamushko](#) on Fri, 09 Nov 2007 10:59:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 09/11/2007, Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com> wrote:

> [...]

>

> As a solution to this problem, I moved sched_fork() call, which
> initializes scheduler related fields on a new task, before
> copy_namespaces(). I am not sure though whether moving up will
> cause other side-effects. Do you see any issue?

Should be ok (IMHO and at first glance :-)

> - The second problem exposed by this test is that task_new_fair()
> assumes that parent and child will be part of the same group (which
> needn't be as this test shows). As a result, cfs_rq->curr can be NULL
> for the child.

Would it be better, logically-wise, to use is_same_group() instead?
Although, we can't have 2 groups with cfs_rq->curr != NULL on the same
CPU... so if the child belongs to another group, it's cfs_rq->curr is
automatically NULL indeed.

--

Best regards,
Dmitry Adamushko

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