
Subject: Re: [PATCH] Hookup group-scheduler with task container infrastructure
Posted by [Srivatsa Vaddagiri](#) on Tue, 18 Sep 2007 10:36:28 GMT

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

On Tue, Sep 18, 2007 at 05:19:45PM +0900, KAMEZAWA Hiroyuki wrote:

> Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com> wrote:

> > + tg->cfs_rq = kzalloc(sizeof(cfs_rq) * num_possible_cpus(), GFP_KERNEL);

> > + if (!tg->cfs_rq)

> > + goto err;

> > + tg->se = kzalloc(sizeof(se) * num_possible_cpus(), GFP_KERNEL);

> > + if (!tg->se)

> > + goto err;

> Sorry for very lazy response..

>

> num_possible_cpus() just returns # of possible cpus. Then it will not return

> Max-cpu-id to be used. I think just use NR_CPUS here is an easy way.

Good catch! Yes, you are right, the pointer array needs to be NR_CPUS size.

On closer examination, I think I can simply use alloc_percpu() here.
Will send a patch after some testing.

--

Regards,
vatsa

Containers mailing list

Containers@lists.linux-foundation.org

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