
Subject: Re: [PATCH -mm] Hook up group scheduler with control groups
Posted by [akpm](#) on Thu, 27 Sep 2007 23:42:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 28 Sep 2007 01:05:12 +0530
Dhaval Giani <dhaval@linux.vnet.ibm.com> wrote:

```
> On Thu, Sep 27, 2007 at 12:00:33PM -0700, Randy Dunlap wrote:
> > On Thu, 27 Sep 2007 23:34:15 +0530 Dhaval Giani wrote:
> > >
> > >
> > > +config RESOURCE_COUNTERS
> > > + bool "Resource counters"
> > > + help
> > > + This option enables controller independent resource accounting
> >
> > Above line is tab + 2 spaces (i.e., correct).
> >
> > > +      infrastructure that works with cgroups.
> >
> > Above line indent is 10 spaces (i.e., not correct).
> >
>
> Ah! Thanks for the explanation. Corrected patch follows.
>
> Signed-off-by : Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com>
> Signed-off-by : Dhaval Giani <dhaval@linux.vnet.ibm.com>
>
> ...
>
```

```
> @@ -219,6 +225,9 @@ static inline struct task_grp *task_grp(
>
> #ifdef CONFIG_FAIR_USER_SCHED
>   tg = p->user->tg;
> +#elif CONFIG_FAIR_CGROUP_SCHED
> + tg = container_of(task_subsys_state(p, cpu_cgroup_subsys_id),
> +   struct task_grp, css);
> #else
>   tg = &init_task_grp;
> #endif
```

that's a bit funny-looking. Are CONFIG_FAIR_CGROUP_SCHED and CONFIG_FAIR_USER_SCHED mutually exclusive? Doesn't seem that way. if they're both defined then CONFIG_FAIR_USER_SCHED "wins".

Anyway, please confirm that this is correct?

I'll switch that to ``#elif defined(CONFIG_FAIR_CGROUP_SCHED)'`. We can get gcc warnings with ``#if CONFIG_FOO'`, and people should be using ``#ifdef CONFIG_FOO'`, so I assume the same applies to `#elif`.

Containers mailing list

Containers@lists.linux-foundation.org

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