

Hi

```
> > v1 -> v2
> > o implement can_attach() and attach() method.
> > o remove can_fork member from cgroup_subsys.
> > instead, copy_process() call task_cgroup_can_fork() directly.
> > o fixed dead lock bug.
> > o add Documentation/controller/task.txt
> >
>
> I think Documentation/kernel-parameters.txt should also be updated.
>
> cgroup_disable= [KNL] Disable a particular controller
> Format: {name of the controller(s) to disable}
> {Currently supported controllers - "memory"}
```

Yup, You are right.
I'll fix it.

```
> > +config CGROUP_TASK
> > + bool "Simple number of task accounting cgroup subsystem"
> > + depends on CGROUPS && EXPERIMENTAL
> > + default n
> > + help
> > +     Provides a simple number of task accounting cgroup subsystem for
>
> should use TAB
```

Ah, Thanks.

```
> > +static void task_cgroup_destroy(struct cgroup_subsys *ss, struct cgroup *cgrp)
> > +{
> > + kfree(cgrp->subsys[task_cgroup_subsys_id]);
>
> You should kfree task_cgroup* instead of cgroup_subsys_state*.
```

There point to the same address.
but, your opinion has better code clearness, of course.

Thanks.

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