
Subject: Re: [RFC] cgroup basic counting
Posted by [Paul Menage](#) on Fri, 16 Dec 2011 16:35:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, Dec 16, 2011 at 4:29 AM, Glauber Costa <glommer@parallels.com> wrote

```
> diff --git a/kernel/cgroup.c b/kernel/cgroup.c
> index 1fd7867..e894a4f 100644
> --- a/kernel/cgroup.c
> +++ b/kernel/cgroup.c
> @@ -1211,9 +1211,9 @@ static int parse_cgroupfs_options(char *data, struct cgroup_sb_opts
> *opts)
>         set_bit(i, &opts->subsys_bits);
>         one_ss = true;
>
> -         break;
> +         continue;
>     }
> -     if (i == CGROUP_SUBSYS_COUNT)
> +     if (opts->subsys_bits == 0)
>         return -ENOENT;
```

This is broken - it will silently ignore unknown/misspelled subsystems that are specified after a valid subsystem. Replacing the break with a continue is harmless but doesn't make sense - if the token already matched a subsystem name then it won't match any other subsystem.

What change are you actually trying to effect with this patch?

Paul
