
Subject: Re: [-mm] CPU controller statistics (v5)
Posted by [Balbir Singh](#) on Wed, 11 Jun 2008 08:40:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On Tue, Jun 3, 2008 at 1:05 PM, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

```
>> diff --git a/include/linux/cgroup.h b/include/linux/cgroup.h
>> index e155aa7..60a25cb 100644
>> --- a/include/linux/cgroup.h
>> +++ b/include/linux/cgroup.h
>> @@ -293,6 +293,7 @@ int cgroup_is_descendant(const struct cgroup *cgrp);
>> struct cgroup_subsys {
>>     struct cgroup_subsys_state *(*create)(struct cgroup_subsys *ss,
>>                                           struct cgroup *cgrp);
>> + void (*initialize)(int early);
>>     void (*pre_destroy)(struct cgroup_subsys *ss, struct cgroup *cgrp);
>>     void (*destroy)(struct cgroup_subsys *ss, struct cgroup *cgrp);
>>     int (*can_attach)(struct cgroup_subsys *ss,
>> diff --git a/kernel/cgroup.c b/kernel/cgroup.c
>> index 15ac0e1..77569d7 100644
>> --- a/kernel/cgroup.c
>> +++ b/kernel/cgroup.c
>> @@ -2553,6 +2553,9 @@ int __init cgroup_init_early(void)
>>
>>     if (ss->early_init)
>>         cgroup_init_subsys(ss);
>> +
>> +     if (ss->initialize)
>> +         ss->initialize(1);
>> }
>> return 0;
>> }
>> @@ -2577,6 +2580,9 @@ int __init cgroup_init(void)
>>     struct cgroup_subsys *ss = subsys[i];
>>     if (!ss->early_init)
>>         cgroup_init_subsys(ss);
>> +
>> +     if (ss->initialize)
>> +         ss->initialize(0);
>> }
```

>

> This seems a little weird - even if the subsystem didn't want early
> initialization, we call its initialize() during early setup?

>

> I assume the idea is to move away from the current model where the
> subsystem is expected to initialize itself during the first call to
> its create() method, when it gets passed a cgroup with a NULL parent?

> I agree that was a bit icky. How about we call ss->initialize() from

> cgroup_init_subsys()? Then we wouldn't need the "early" parameter,
> since it would be implicit based on whether the subsystem wanted early
> initialization or not.
>

The motivation was to ensure that kmalloc* calls are available at the time of
ss_initialize

> Also, if you're adding a new subsystem method, you should document it
> in Documentation/cgroups.txt

Thanks for the comments. There were some more things that are wrong with this
patch and Andrew pointed them out. We'll work on a newer vesion.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

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