Subject: Re: [RFC][patch 3/11][CFQ-cgroup] Introduce cgroup subsystem
Posted by Li Zefan on Thu, 03 Apr 2008 02:39:09 GMT
View Forum Message <> Reply to Message

Satoshi UCHIDA wrote:
> Thank you for reply.
>
>>> +
>>> +static struct cgroup_subsys_state *
>>> +cfq_cgroup_create(struct cgroup_subsys *ss, struct cgroup *cont)
>>> +{
>>> + struct cfq_cgroup *cfqc;
>>> +
>>> + if (!capable(CAP_SYS_ADMIN))
>>> + return ERR_PTR(-EPERM);
>>> +
>>> + if (!cgroup_is_descendant(cont))
>>> + return ERR_PTR(-EPERM);
>> What are these checks for? Cgroups already provides filesystem
>> permissions to control directory creation, and the "descendant" check
>> looks like it may have been cut/pasted from the nsproxy subsystem.
>>
>
> This code was referred one of io-throttle.
> Is it not necessary these checks?
> IF not necessary, remove this code.
>
>>> /* */
>>> +
>>> +#ifdef CONFIG_CGROUP_CFQ
>>> +SUBSYS(cfq_cgroup)
>>> +#endif
>>> +
>>> +/* */
>> To fit with the convention for other subsystems, simply "cfq" would be
>> a better name than "cfq_cgroup". (Clearly it's a cgroup subsystem from
>> context).
>>
>
> Ok, I change name.
> I hesitated whether using "_cgroup".
> The cpuset and the cpuacct does not use it,
> but cpu and memory uses it(cpu_cgroup and mem_cgroup).
> In this patchset, I select the latter case.
>

+struct cgroup_subsys cfq_cgroup_subsys = {
+ .name = "cfq_cgroup",

+ ...
+};

but memory controller has the name 'memory', similar for cgroup sched.

So we do this:
 mount -t cgroup -omemory xxx /dev/memcg
but not:
 mount -t cgroup -omemory_cgroup xxx /dev/memcg

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