
Subject: Re: [PATCH 4/4] get rid of populate for memcg
Posted by [Tejun Heo](#) on Tue, 20 Mar 2012 18:31:49 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello, Glauber.

On Tue, Mar 20, 2012 at 08:50:56PM +0400, Glauber Costa wrote:

```
> @@ -4929,7 +4929,9 @@ mem_cgroup_create(struct cgroup *cont)
>     atomic_set(&memcg->refcnt, 1);
>     memcg->move_charge_at_immigrate = 0;
>     mutex_init(&memcg->thresholds_lock);
> - return &memcg->css;
> +
> + if (!register_kmem_files(memcg, &mem_cgroup_subsys))
> + return &memcg->css;
```

After the change, I think register_kmem_files() is a quite misleading name.

```
> @@ -2484,6 +2484,11 @@ int proto_register(struct proto *prot, int alloc_slab)
>     }
>   }
>
> +#ifdef CONFIG_CGROUP_MEM_RES_CTLR_KMEM
> + if (prot->init_cgroup)
> +   prot->init_cgroup(NULL, NULL);
> +#endif
```

So, init_cgroup() is overloaded to do two things - one load time init and per-cgroup init, depending on the args.

```
> @@ -37,7 +37,6 @@ static struct cftype tcp_files[] = {
>   },
>   { } /* terminate */
> };
> -CGROUP_SUBSYS_CFTYPES(mem_cgroup_subsys, tcp_files);
```

What I don't get is why you can't just keep this. Is it because the files might appear before the protocol is registered? Wouldn't it be much better to add ipv4_tcp_init_cgroup() or whatever call to inet_init() instead of overloading init_cgroup() with mostly unrelated stuff?

Thanks.

--
tejun
