

---

Subject: Re: [PATCH v2 06/11] memcg: kmem controller infrastructure  
Posted by [Glauber Costa](#) on Tue, 14 Aug 2012 11:00:53 GMT

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

---

On 08/10/2012 09:27 PM, Kamezawa Hiroyuki wrote:

```
>> +bool __memcg_kmem_new_page(gfp_t gfp, void *_handle, int order)
>> > +{
>> > + struct mem_cgroup *memcg;
>> > + struct mem_cgroup **handle = (struct mem_cgroup **)_handle;
>> > + bool ret = true;
>> > + size_t size;
>> > + struct task_struct *p;
>> > +
>> > + *handle = NULL;
>> > + rcu_read_lock();
>> > + p = rcu_dereference(current->mm->owner);
>> > + memcg = mem_cgroup_from_task(p);
>> > + if (!memcg_kmem_enabled(memcg))
>> > + goto out;
>> > +
>> > + mem_cgroup_get(memcg);
>> > +
> This mem_cgroup_get() will be a potential performance problem.
> Don't you have good idea to avoid accessing atomic counter here ?
> I think some kind of percpu counter or a feature to disable "move task"
> will be a help.
>
>
```

I have just sent out a proposal to deal with this. I tried the trick of marking only the first charge and last uncharge, and it works quite alright at the cost of a bit test on most calls to memcg\_kmem\_charge.

Please let me know what you think.

---