
Subject: Re: [PATCH v3 06/13] memcg: kmem controller infrastructure
Posted by [Glauber Costa](#) on Fri, 28 Sep 2012 11:34:19 GMT

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

On 09/27/2012 05:44 PM, Michal Hocko wrote:

>> > the reference count aquired by mem_cgroup_get will still prevent the
>> > memcg from going away, no?
> Yes but you are outside of the rcu now and we usually do css_get before
> we rcu_unlock. mem_cgroup_get just makes sure the group doesn't get
> deallocated but it could be gone before you call it. Or I am just
> confused - these 2 levels of ref counting is really not nice.
>
> Anyway, I have just noticed that __mem_cgroup_try_charge does
> VM_BUG_ON(css_is_removed(&memcg->css)) on a given memcg so you should
> keep css ref count up as well.
>

IIRC, css_get will prevent the cgroup directory from being removed.
Because some allocations are expected to outlive the cgroup, we
specifically don't want that.
