
Subject: Re: [PATCH v3 06/13] memcg: kmem controller infrastructure
Posted by [Glauber Costa](#) on Mon, 01 Oct 2012 08:28:28 GMT

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

On 09/30/2012 12:25 PM, Tejun Heo wrote:

> On Fri, Sep 28, 2012 at 03:34:19PM +0400, Glauber Costa wrote:

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

>>> 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.

>

> That synchronous ref draining is going away. Maybe we can do that

> before kmemcg? Michal, do you have some timeframe on mind?

>

Since you said yourself in other points in this thread that you are fine with some page references outliving the cgroup in the case of slab, this is a situation that comes with the code, not a situation that was incidentally there, and we're making use of.
