

---

Subject: Re: [PATCH v2 10/11] memcg: allow a memcg with kmem charges to be destructed.

Posted by [Glauber Costa](#) on Wed, 22 Aug 2012 08:36:17 GMT

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

---

On 08/21/2012 12:22 PM, Michal Hocko wrote:

> On Thu 09-08-12 17:01:18, Glauber Costa wrote:

>> Because the ultimate goal of the kmem tracking in memcg is to track slab  
>> pages as well, we can't guarantee that we'll always be able to point a  
>> page to a particular process, and migrate the charges along with it -  
>> since in the common case, a page will contain data belonging to multiple  
>> processes.

>>

>> Because of that, when we destroy a memcg, we only make sure the  
>> destruction will succeed by discounting the kmem charges from the user  
>> charges when we try to empty the cgroup.

>

> This changes the semantic of memory.force\_empty file because the usage  
> should be 0 on success but it will show kmem usage in fact now. I guess  
> it is inevitable with u+k accounting so you should be explicit about  
> that and also update the documentation.

aaand, it's done.

---