
Subject: Re: [PATCH v3 19/28] slab: per-memcg accounting of slab caches
Posted by [Glauber Costa](#) on Tue, 29 May 2012 16:13:28 GMT

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

On 05/29/2012 08:07 PM, Glauber Costa wrote:

> On 05/29/2012 06:52 PM, Christoph Lameter wrote:

>> On Fri, 25 May 2012, Glauber Costa wrote:

>>

>>> > This patch charges allocation of a slab object to a particular

>>> > memcg.

>> Ok so a requirement is to support tracking of individual slab

>> objects to cgroups? That is going to be quite expensive since it will

>> touch the hotpaths.

>>

>

> No, we track pages. But all the objects in the page belong to the same

> cgroup.

>

Also, please note the following:

The code that relays us to the right cache, is wrapped inside a static branch. Whoever is not using more than the root cgroup, will not suffer a single bit.

If you are, but your process is in the right cgroup, you will unfortunately pay function call penalty(*), but the code will make an effort to detect that as early as possible and resume.

(*) Not even then if you fall in the following categories, that are resolved inline:

```
+   if (!current->mm)
+       return cachep;
+   if (in_interrupt())
+       return cachep;
+   if (gfp & __GFP_NOFAIL)
+       return cachep;
```
