Subject: Re: [PATCH v3 07/16] memcg: skip memcg kmem allocations in specified code regions

Posted by Glauber Costa on Mon, 24 Sep 2012 09:09:23 GMT

View Forum Message <> Reply to Message

```
On 09/21/2012 11:59 PM, Tejun Heo wrote:
> Hello.
>
> On Tue, Sep 18, 2012 at 06:12:01PM +0400, Glauber Costa wrote:
>> +static void memcg stop kmem account(void)
>> +{
>> + if (!current->mm)
>> + return;
>> +
>> + current->memcg_kmem_skip_account++;
>> +}
>> +
>> +static void memcg_resume_kmem_account(void)
>> +{
>> + if (!current->mm)
>> + return;
>> +
>> + current->memcg_kmem_skip_account--;
>> +}
> I can't say I'm a big fan of this approach. If there are enough
> users, maybe but can't we just annotate the affected allocations
> explicitly? Is this gonna have many more users?
>
```

What exactly do you mean by annotating the affected allocations?

There are currently two users of this. In both places, we are interested in disallowing recursion, because cache creation will trigger new cache allocations that will bring us back here.

We can't rely on unsetting the GFP flag we're using for this, because that affects only the page allocation, not the metadata allocation for the cache.

- Also, in general, can we please add some comments? I know memcg codeis dearth of comments but let's please not keep it that way.
- All right here.