Subject: Re: [PATCH v5 07/14] mm: Allocate kernel pages to the right memcg
Posted by Glauber Costa on Thu, 18 Oct 2012 09:24:47 GMT
View Forum Message <> Reply to Message

On 10/18/2012 02:12 AM, Andrew Morton wrote:
> On Tue, 16 Oct 2012 14:16:44 +0400
> Glauber Costa <glommer@parallels.com> wrote:
>
>> When a process tries to allocate a page with the __GFP_KMEMCG flag, the
>> page allocator will call the corresponding memcg functions to validate
>> the allocation. Tasks in the root memcg can always proceed.
>>
>> To avoid adding markers to the page - and a kmem flag that would
>> necessarily follow, as much as doing page_cgroup lookups for no reason,
>> whoever is marking its allocations with __GFP_KMEMCG flag is responsible
>> for telling the page allocator that this is such an allocation at
>> free_pages() time.
>
> Well, why? Was that the correct decision?
>

I don't fully understand your question. Is this the same question you
posed in patch 0, about marking some versus marking all? If so, I
believe I should have answered it there.

If not, please explain.

>> This is done by the invocation of
>> __free_accounted_pages() and free_accounted_pages().
>
> These are very general-sounding names. I'd expect the identifiers to
> contain "memcg" and/or "kmem", to identify what's going on.
>

Changed.