

---

Subject: Re: [PATCH v2 07/11] mm: Allocate kernel pages to the right memcg  
Posted by [Glauber Costa](#) on Mon, 13 Aug 2012 08:03:38 GMT

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

---

On 08/10/2012 09:33 PM, Kamezawa Hiroyuki wrote:

> (2012/08/09 22:01), Glauber Costa 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. This is done by the invocation of  
>> `__free_accounted_pages()` and `free_accounted_pages()`.

>>

>> Signed-off-by: Glauber Costa <[glommer@parallels.com](mailto:glommer@parallels.com)>

>> CC: Christoph Lameter <[cl@linux.com](mailto:cl@linux.com)>

>> CC: Pekka Enberg <[penberg@cs.helsinki.fi](mailto:penberg@cs.helsinki.fi)>

>> CC: Michal Hocko <[mhocko@suse.cz](mailto:mhocko@suse.cz)>

>> CC: Kamezawa Hiroyuki <[kamezawa.hiroyu@jp.fujitsu.com](mailto:kamezawa.hiroyu@jp.fujitsu.com)>

>> CC: Johannes Weiner <[hannes@cmpxchg.org](mailto:hannes@cmpxchg.org)>

>> CC: Suleiman Souhlal <[suleiman@google.com](mailto:suleiman@google.com)>

>

> Ah, ok. `free_accounted_page()` seems good.

>

> Acked-by: KAMEZAWA Hiroyuki <[kamezawa.hiroyu@jp.fujitsu.com](mailto:kamezawa.hiroyu@jp.fujitsu.com)>

>

> I myself is okay with this. But...

>

> Because you add a new hook to `alloc_pages()`, please get Ack from Mel  
> before requesting merge.

>

> Thanks,

> -Kame

Absolutely.

Mel, would you mind taking a look at this series and commenting on this?

Thanks in advance.

---