
Subject: Re: [PATCH v2 07/11] mm: Allocate kernel pages to the right memcg
Posted by [Mel Gorman](#) on Mon, 13 Aug 2012 08:57:49 GMT

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

On Mon, Aug 13, 2012 at 12:03:38PM +0400, Glauber Costa wrote:

> 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>

> >> CC: Christoph Lameter <cl@linux.com>

> >> CC: Pekka Enberg <penberg@cs.helsinki.fi>

> >> CC: Michal Hocko <mhocko@suse.cz>

> >> CC: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

> >> CC: Johannes Weiner <hannes@cmpxchg.org>

> >> CC: Suleiman Souhlal <suleiman@google.com>

> >

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

> >

> > Acked-by: KAMEZAWA Hiroyuki <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?

>

It'll take me a few days but I'll get around to it.

--

Mel Gorman
SUSE Labs
