
Subject: Re: [PATCH v3 07/13] mm: Allocate kernel pages to the right memcg
Posted by [Mel Gorman](#) on Fri, 28 Sep 2012 13:28:39 GMT

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

On Fri, Sep 28, 2012 at 01:43:47PM +0400, Glauber Costa wrote:

> On 09/27/2012 05:50 PM, Mel Gorman wrote:

> >> +void __free_accounted_pages(struct page *page, unsigned int order)

> >> > +{

> >> > + memcg_kmem_uncharge_page(page, order);

> >> > + __free_pages(page, order);

> >> > +}

> >> > +

> >> > +void free_accounted_pages(unsigned long addr, unsigned int order)

> >> > +{

> >> > + if (addr != 0) {

> >> > + VM_BUG_ON(!virt_addr_valid((void *)addr));

> > This is probably overkill. If it's invalid, the next line is likely to

> > blow up anyway. It's no biggie.

> >

>

> So this is here because it is in free_pages() as well. If it blows, at

> least we know precisely why (if debugging), and VM_BUG_ON() is only

> compiled in when CONFIG_DEBUG_VM.

>

Ah, I see.

> But I'm fine with either.

> Should it stay or should it go ?

>

It can stay. It makes sense that it look similar to free_pages() and as
you say, it makes debugging marginally easier.

--

Mel Gorman

SUSE Labs
