
Subject: Re: [PATCH v2 08/11] memcg: disable kmem code when not in use.
Posted by [Michal Hocko](#) on Fri, 17 Aug 2012 08:04:04 GMT

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

On Fri 17-08-12 11:01:06, Glauber Costa wrote:

> On 08/17/2012 11:02 AM, Michal Hocko wrote:

> > On Thu 09-08-12 17:01:16, Glauber Costa wrote:

> > > We can use jump labels to patch the code in or out when not used.

> > >

> > > Because the assignment: memcg->kmem_accounted = true is done after the

> > > jump labels increment, we guarantee that the root memcg will always be

> > > selected until all call sites are patched (see memcg_kmem_enabled).

> > >

> > > Not that it would be really important because kmem_accounted goes away

And just found out it doesn't go away completely, it just transforms
from bool to unsigned long (with flags). The rest still holds...

> > in a subsequent patch but I think the wording is a bit misleading here.

> > First of all there is no guarantee that kmem_accounted=true is seen

> > before atomic_inc(&key->enabled) because there is no memory barrier and

> > the lock serves just a leave barrier. But I do not think this is

> > important at all because key->enabled is what matters here. Even if

> > memcg_kmem_enabled is true we do not consider it if the key is disabled,

> > right?

> >

>

> Right.

--

Michal Hocko
SUSE Labs
