
Subject: Re: [PATCH v2 04/11] kmem accounting basic infrastructure
Posted by [Glauber Costa](#) on Wed, 15 Aug 2012 19:22:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 08/15/2012 10:25 PM, Christoph Lameter wrote:

> On Wed, 15 Aug 2012, Ying Han wrote:

>
>>> How can you figure out which objects belong to which memcg? The ownerships
>>> of dentries and inodes is a dubious concept already.

>>
>> I figured it out based on the kernel slab accounting.
>> obj->page->kmem_cache->memcg

>
> Well that is only the memcg which allocated it. It may be in use heavily
> by other processes.
>

Yes, but a lot of the use cases for cgroups/containers are pretty local.
That is why we have been able to get away with a first-touch mechanism
even in user pages memcg. In those cases - which we expect to be the
majority of them - this will perform well.

Now, this is not of course representative of the whole range of possible
use cases, and others are valid. There are people like Greg
and Ying Han herself that want a more fine grained control on which
memcg gets the accounting. That is one of the topics for the summit.

But even then: regardless of what mechanism is in place, one cgroup is
to be accounted (or not accounted at all, meaning it belongs to a
non-accounted cgroup). And then we can just grab whichever memcg it was
allocated from and shrink it.
