Subject: Re: [PATCH v5 09/14] memcg: kmem accounting lifecycle management
Posted by Michal Hocko on Thu, 18 Oct 2012 06:14:51 GMT

On Wed 17-10-12 16:28:38, David Rientjes wrote:
> On Tue, 16 Oct 2012, Glauber Costa wrote:
[...]
> > +
> > +static void memcg_kmem_mark_dead(struct mem_cgroup *memcg)
> > +{
> > + if (test_bit(KMEM_ACCOUNTED_ACTIVE, &memcg->kmem_accounted))
> > + set_bit(KMEM_ACCOUNTED_DEAD, &memcg->kmem_accounted);
> > +}
>
> The set_bit() doesn't happen atomically with the test_bit(), what
> synchronization is required for this?

The group has to be active in order to become dead so the ordering is
natural and you do not need to test&set atomicaly. Also once a group
becomes active it is always marked that way until it goes away.
--
Michal Hocko
SUSE Labs