
Subject: Re: [PATCH v3 04/13] kmem accounting basic infrastructure

Posted by [Michal Hocko](#) on Thu, 27 Sep 2012 17:56:43 GMT

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

On Thu 27-09-12 10:46:05, Tejun Heo wrote:

[...]

> > > The part I nacked is enabling kmemcg on a populated cgroup and then
> > > starting accounting from then without any apparent indication that any
> > > past allocation hasn't been considered. You end up with numbers which
> > > nobody can't tell what they really mean and there's no mechanism to
> > > guarantee any kind of ordering between populating the cgroup and
> > > configuring it and there's *no* way to find out what happened
> > > afterwards neither. This is properly crazy and definitely deserves a
> > > nack.

> > >

> >

> > Mel suggestion of not allowing this to happen once the cgroup has tasks
> > takes care of this, and is something I thought of myself.

>

> You mean Michal's? It should also disallow switching if there are
> children cgroups, right?

Right.

--

Michal Hocko
SUSE Labs
