Subject: Re: [PATCH v2 06/11] memcg: kmem controller infrastructure
Posted by Michal Hocko on Wed, 15 Aug 2012 14:23:38 GMT
View Forum Message <> Reply to Message

On Wed 15-08-12 18:01:51, Glauber Costa wrote:
> On 08/15/2012 05:09 PM, Michal Hocko wrote:
> > On Wed 15-08-12 13:42:24, Glauber Costa wrote:
> > [...]
> >>>> +
> >>>> + ret = 0;
> >>>> +
> >>>> + if (!memcg)
> >>>> + return ret;
> >>>> +
> >>>> + _memcg = memcg;
> >>>> + ret = __mem_cgroup_try_charge(NULL, gfp, delta / PAGE_SIZE,
> >>>> +   &_memcg, may_oom);
> >>>
> >>> This is really dangerous because atomic allocation which seem to be
> >>> possible could result in deadlocks because of the reclaim.
> >>
> >> Can you elaborate on how this would happen?
> >
> > Say you have an atomic allocation and we hit the limit so we get either
> > to reclaim which can sleep or to oom which can sleep as well (depending
> > on the oom_control).
> >
>
> I see now, you seem to be right.

No I am not because it seems that I am really blind these days...
We were doing this in mem_cgroup_do_charge for ages:
 if (!(gfp_mask & __GFP_WAIT))
  return CHARGE_WOULDBLOCK;

/me goes to hide and get with further feedback with a clean head.

Sorry about that.
--
Michal Hocko
SUSE Labs