
Subject: Re: [RFC 1/4] memcg: provide root figures from system totals

Posted by [Michal Hocko](#) on Tue, 02 Oct 2012 09:34:33 GMT

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

On Tue 02-10-12 13:15:43, Glauber Costa wrote:

> On 10/01/2012 09:00 PM, Michal Hocko wrote:

> > On Tue 25-09-12 12:52:50, Glauber Costa wrote:

> > > For the root memcg, there is no need to rely on the res_counters.

> > This is true only if there are no children groups but once there is at

> > least one we have to move global statistics into root res_counter and

> > start using it since then. This is a tricky part because it has to be

> > done atomically so that we do not miss anything.

> >

> Why can't we shortcut it all the time?

Because it has its own tasks and we are still not at use_hierarchy := 1

> It makes a lot of sense to use the root cgroup as the sum of everything,

> IOW, global counters. Otherwise you are left in a situation where you

> had global statistics, and all of a sudden, when a group is created, you

> start having just a subset of that, excluding the tasks in root.

Yes because if there are no other tasks then, well, global == root. Once you have more groups (with tasks of course) then it depends on our favorite use_hierarchy buddy.

> If we can always assume root will have the sum of *all* tasks, including

> the ones in root, we should never need to rely on root res_counters.

but we are not there yet.

--

Michal Hocko
SUSE Labs
