Subject: Re: [PATCH 04/10] memcg: skip memcg kmem allocations in specified code regions

Posted by Glauber Costa on Mon, 30 Jul 2012 14:09:23 GMT

View Forum Message <> Reply to Message

On 07/30/2012 04:50 PM, Kirill A. Shutemov wrote:

- > On Wed, Jul 25, 2012 at 06:38:15PM +0400, Glauber Costa wrote:
- >> This patch creates a mechanism that skip memcg allocations during
- >> certain pieces of our core code. It basically works in the same way
- >> as preempt_disable()/preempt_enable(): By marking a region under
- >> which all allocations will be accounted to the root memcg.

>>

- >> We need this to prevent races in early cache creation, when we
- >> allocate data using caches that are not necessarily created already.

>

> Why not a GFP_* flag?

>

The main reason for this is to prevent nested calls of kmem_cache_create(), since they could create (and in my tests, do create) funny circular dependencies with each other. So the cache creation itself would proceed without involving memcg.

At first, it is a bit weird to have cache creation itself depending on a allocation flag test.