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

Posted by [Tejun Heo](#) on Wed, 26 Sep 2012 16:36:48 GMT

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

Hello, Michal, Glauber.

On Wed, Sep 26, 2012 at 04:03:47PM +0200, Michal Hocko wrote:

```
> Haven't we already discussed that a new memcg should inherit kmem_accounted
> from its parent for use_hierarchy?
> Say we have
> root
> |
> A (kmem_accounted = 1, use_hierarchy = 1)
> \
>  B (kmem_accounted = 0)
>  \
>   C (kmem_accounted = 1)
>
> B find's itself in an awkward situation because it doesn't want to
> account u+k but it ends up doing so because C.
```

Do we really want this level of flexibility? What's wrong with a global switch at the root? I'm not even sure we want this to be optional at all. The only reason I can think of is that it might screw up some configurations in use which are carefully crafted to suit userland-only usage but for that isn't what we need a transition plan rather than another ultra flexible config option that not many really understand the implication of?

In the same vein, do we really need both `.kmem_accounted` and config option? If someone is turning on MEMCG, just include kmem accounting.

Thanks.

--

tejun
