
Subject: Re: [PATCH 00/11] kmem controller for memcg: stripped down version
Posted by [akpm](#) on Mon, 25 Jun 2012 23:27:45 GMT

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

On Mon, 25 Jun 2012 18:15:17 +0400

Glauber Costa <glommer@parallels.com> wrote:

> What I am proposing with this series is a stripped down version of the
> kmem controller for memcg that would allow us to merge significant parts
> of the infrastructure, while leaving out, for now, the polemic bits about
> the slab while it is being reworked by Cristoph.
>
> Me reasoning for that is that after the last change to introduce a gfp
> flag to mark kernel allocations, it became clear to me that tracking other
> resources like the stack would then follow extremely naturally. I figured
> that at some point we'd have to solve the issue pointed by David, and avoid
> testing the Slab flag in the page allocator, since it would soon be made
> more generic. I do that by having the callers to explicit mark it.
>
> So to demonstrate how it would work, I am introducing a stack tracker here,
> that is already a functionality per-se: it successfully stops fork bombs to
> happen. (Sorry for doing all your work, Frederic =p). Note that after all
> memcg infrastructure is deployed, it becomes very easy to track anything.
> The last patch of this series is extremely simple.
>
> The infrastructure is exactly the same we had in memcg, but stripped down
> of the slab parts. And because what we have after those patches is a feature
> per-se, I think it could be considered for merging.

hm. None of this new code makes the kernel smaller, faster, easier to
understand or more fun to read!

Presumably we're getting some benefit for all the downside. When the
time is appropriate, please do put some time into explaining that
benefit, so that others can agree that it is a worthwhile tradeoff.
