

---

Subject: Re: [PATCH 00/10] memcg: Kernel Memory Accounting.  
Posted by [Suleiman Souhlal](#) on Tue, 28 Feb 2012 22:12:47 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hello,

On Tue, Feb 28, 2012 at 12:49 AM, Pekka Enberg <penberg@kernel.org> wrote:

> On Mon, 27 Feb 2012, Suleiman Souhlal wrote:

>> The main difference with Glauber's patches is here: We try to

>> track all the slab allocations, while Glauber only tracks ones

>> that are explicitly marked.

>> We feel that it's important to track everything, because there

>> are a lot of different slab allocations that may use significant

>> amounts of memory, that we may not know of ahead of time.

>> This is also the main source of complexity in the patchset.

>

> Well, what are the performance implications of your patches? Can we

> reasonably expect distributions to be able to enable this thing on

> generic kernels and leave the feature disabled by default? Can we

> accommodate your patches to support Glauber's use case?

I don't have up to date performance numbers, but we haven't found any critical performance degradations on our workloads, with our internal versions of this patchset.

There are some conditional branches added to the slab fast paths, but I think it should be possible to come with a way to get rid of those when the feature is disabled, maybe by using a static\_branch. This should hopefully make it possible to keep the feature compiled in but disabled at runtime.

I think it's definitely possible to accommodate my patches to support Glauber's use case, with a bit of work.

-- Suleiman

---