Subject: Re: [PATCH 3/5] Switch caches notification dynamically Posted by Pavel Emelianov on Tue, 25 Sep 2007 07:46:44 GMT

View Forum Message <> Reply to Message

Christoph Lameter wrote:

> On Fri, 21 Sep 2007, Pavel Emelyanov wrote:

>

- >> The /sys/slab/<name>/cache_notify attribute controls
- >> whether the cache <name> is to be accounted or not.

>>

- >> For the reasons described before kmalloc caches cannot
- >> be turned on.

>

- > It looks like the patch is forbidding to turn the notification off? On is
- > okay even if there are already objects present?

Yup. If we "on" the notifications we will just receive new events.

- > Full slabs may exist that cannot be accounted for now. I guess that is
- > okay because we only want to track new allocations?

If we turn notifications off with full slabs in memory we won't be able to (in case of kmem accounting) release the page's metadata that the slub_newpage event attached to them.

We can make this by tracking all the full slabs, and at the moment of "off" walk them and generate slab_free and slab_release_page events, but does it worth doing it ever?

Thanks, Pavel