
Subject: Re: [PATCH v3 13/28] slub: create duplicate cache
Posted by [Glauber Costa](#) on Tue, 29 May 2012 17:27:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 05/29/2012 09:25 PM, Christoph Lameter wrote:

>> I will try to at least have the page accounting done in a consistent way. How

>> > about that?

> Ok. What do you mean by "consistent"? Since objects and pages can be used
> in a shared way and since accounting in many areas of the kernel is
> intentional fuzzy to avoid performance issues there is not that much of a
> requirement for accuracy. We have problems even just defining what the
> exact set of pages belonging to a task/process is.

>

I mean consistent between allocators, done in a shared place.

Note that what we do here *is* still fuzzy to some degree. We can have a
level of sharing, and always account to the first one to touch, we don't
go worrying about this because that would be hell. We also don't
carry a process' objects around when we send it to a different cgroup.
Those are all already simplifications for performance and simplicity sake.

But we really need a page to be filled with objects from the same
cgroup, and the non-shared objects to be accounted to the right place.

Otherwise, I don't think we can meet even the lighter of isolation
guarantees.
