
Subject: Re: [RFC] slub: show dead memcg caches in a separate file

Posted by [Suleiman Souhlal](#) on Mon, 07 May 2012 22:04:39 GMT

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

On Thu, May 3, 2012 at 11:47 AM, Glauber Costa <glommer@parallels.com> wrote:

> One of the very few things that still unsettles me in the kmem
> controller for memcg, is how badly we mess up with the
> /proc/slabinfo file.
>
> It is alright to have the cgroup caches listed in slabinfo, but once
> they die, I think they should be removed right away. A box full
> of containers that come and go will rapidly turn that file into
> a supreme mess. However, we currently leave them there so we can
> determine where our used memory currently is.
>
> This patch attempts to clean this up by creating a separate proc file
> only to handle the dead slabs. Among other advantages, we need a lot
> less information in a dead cache: only its current size in memory
> matters to us.
>
> So besides avoiding pollution of the slabinfo files, we can access
> dead cache information itself in a cleaner way.
>
> I implemented this as a proof of concept while finishing up
> my last round for submission. But I am sending this separately
> to collect opinions from all of you. I can either implement
> a version of this for the slab, or follow any other route.

I don't really understand why the "dead" slabs are considered as
polluting slabinfo.

They still have objects in them, and I think that hiding them would
not be the right thing to do (even if they are available in a separate
file): They will incorrectly not be seen by programs like slabtop.

-- Suleiman
