Subject: Re: [PATCH v3 12/13] execute the whole memcg freeing in rcu callback
Posted by Glauber Costa on Mon, 24 Sep 2012 08:48:01 GMT
View Forum Message <> Reply to Message

> And the above description too makes me scratch my head quite a bit. I
> can see what the patch is doing but can't understand the why.
>
> * Why was it punting the freeing to workqueue anyway? ISTR something
>   about static_keys but my memory fails. What changed? Why don't we
>   need it anymore?
>
> * As for locking context, the above description seems a bit misleading
>   to me. Synchronization constructs involved there currently doesn't
>   require softirq or irq safe context. If that needs to change,
>   that's fine but that's a completely different reason than given
>   above.
>
> Thanks.
>

I just suck at changelogs =(

The problem here is very much like the one we had with static branches.
In that case, we had the problem with the cgroup_lock() being held, in
which case the jump label lock could not be called.

In here, after the kmem patches are in, the destruction function could
be called directly from memcg_kmem_uncharge_page() when the last put is
done. But this can actually be called from the page allocator, with an
incompatible softirq context. So it is not that it could be called, they
are actually called in that context at this point.