
Subject: Re: [PATCH v3 05/13] Add a __GFP_KMEMCG flag
Posted by [Mel Gorman](#) on Thu, 27 Sep 2012 13:34:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Sep 18, 2012 at 06:04:02PM +0400, Glauber Costa wrote:

> This flag is used to indicate to the callees that this allocation is a
> kernel allocation in process context, and should be accounted to
> current's memcg. It takes numerical place of the of the recently removed
> __GFP_NO_KSWAPD.
>
> Signed-off-by: Glauber Costa <glommer@parallels.com>
> CC: Christoph Lameter <cl@linux.com>
> CC: Pekka Enberg <penberg@cs.helsinki.fi>
> CC: Michal Hocko <mhocko@suse.cz>
> CC: Johannes Weiner <hannes@cmpxchg.org>
> CC: Suleiman Souhlal <suleiman@google.com>
> CC: Rik van Riel <riel@redhat.com>
> CC: Mel Gorman <mel@csn.ul.ie>
> Acked-by: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

I agree with Christophs recommendation that this flag always exist instead of being 0 in the !MEMCG_KMEM case. If __GFP_KMEMCG ever is used in another part of the VM (which would be unexpected but still) then the behaviour might differ too much between MEMCG_KMEM and !MEMCG_KMEM cases. As unlikely as the case is, it's not impossible.

For tracing __GFP_KMEMCG should have an entry in
include/trace/events/gfpflags.h

Get rid of the CONFIG_MEMCG_KMEM check and update
include/trace/events/gfpflags.h and then feel free to stick my Acked-by
on it.

--
Mel Gorman
SUSE Labs
