
Subject: Re: [PATCH v4 05/14] Add a __GFP_KMEMCG flag
Posted by [Michal Hocko](#) on Tue, 09 Oct 2012 15:04:03 GMT

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

On Mon 08-10-12 14:06:11, 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.
>
> [v4: make flag unconditional, also declare it in trace code]
>
> 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: Suleiman Souhlal <suleiman@google.com>
> Acked-by: Johannes Weiner <jannes@cmpxchg.org>
> Acked-by: Rik van Riel <riel@redhat.com>
> Acked-by: Mel Gorman <mel@csn.ul.ie>
> Acked-by: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

Acked-by: Michal Hocko <mhocko@suse.cz>

> ---
> include/linux/gfp.h | 3 +-+
> include/trace/events/gfpflags.h | 1 +
> 2 files changed, 3 insertions(+), 1 deletion(-)
>
> diff --git a/include/linux/gfp.h b/include/linux/gfp.h
> index 02c1c97..9289d46 100644
> --- a/include/linux/gfp.h
> +++ b/include/linux/gfp.h
> @@ -31,6 +31,7 @@ struct vm_area_struct;
> #define __GFP_THISNODE 0x40000u
> #define __GFP_RECLAMABLE 0x80000u
> #define __GFP_NOTRACK 0x200000u
> +#define __GFP_KMEMCG 0x400000u
> #define __GFP_OTHER_NODE 0x800000u
> #define __GFP_WRITE 0x1000000u
>
> @@ -87,7 +88,7 @@ struct vm_area_struct;
>
> #define __GFP_OTHER_NODE ((__force gfp_t)__GFP_OTHER_NODE) /* On behalf of other
node */
> #define __GFP_WRITE ((__force gfp_t)__GFP_WRITE) /* Allocator intends to dirty page */
> -
> +#define __GFP_KMEMCG ((__force gfp_t)__GFP_KMEMCG) /* Allocation comes from a

```
memcg-accounted resource */  
> /*  
> * This may seem redundant, but it's a way of annotating false positives vs.  
> * allocations that simply cannot be supported (e.g. page tables).  
> diff --git a/include/trace/events/gfpflags.h b/include/trace/events/gfpflags.h  
> index 9391706..730df12 100644  
> --- a/include/trace/events/gfpflags.h  
> +++ b/include/trace/events/gfpflags.h  
> @@ -36,6 +36,7 @@  
> {{(unsigned long)__GFP_RECLAMABLE, "GFP_RECLAMABLE"}, \  
> {{(unsigned long)__GFP_MOVABLE, "GFP_MOVABLE"}, \  
> {{(unsigned long)__GFP_NOTRACK, "GFP_NOTRACK"}, \  
> + {{(unsigned long)__GFP_KMEMCG, "GFP_KMEMCG"}, \  
> {{(unsigned long)__GFP_OTHER_NODE, "GFP_OTHER_NODE"} \  
> ) : "GFP_NOWAIT"  
>  
>--  
> 1.7.11.4  
>  
>--  
> To unsubscribe from this list: send the line "unsubscribe cgroups" in  
> the body of a message to majordomo@vger.kernel.org  
> More majordomo info at http://vger.kernel.org/majordomo-info.html
```

--
Michal Hocko
SUSE Labs
