On Thu, Oct 18, 2012 at 12:07 AM, David Rientjes <rientjes@google.com> wrote:
> On Wed, 17 Oct 2012, Glauber Costa wrote:
>
>> Some flags are used internally by the allocators for management
>> purposes. One example of that is the CFLGS_OFF_SLAB flag that slab uses
>> to mark that the metadata for that cache is stored outside of the slab.
>>
>> No cache should ever pass those as a creation flags. We can just ignore
>> this bit if it happens to be passed (such as when duplicating a cache in
>> the kmem memcg patches).
>>
>> Because such flags can vary from allocator to allocator, we allow them
>> to make their own decisions on that, defining SLAB_AVAILABLE_FLAGS with
>> all flags that are valid at creation time. Allocators that doesn't have
>> any specific flag requirement should define that to mean all flags.
>>
>> Common code will mask out all flags not belonging to that set.
>>
>> [v2: leave the mask out decision up to the allocators]
>> [v3: define flags for all allocators]
>> [v4: move all definitions to slab.h]
>>
>> Signed-off-by: Glauber Costa <glommer@parallels.com>
>> Acked-by: Christoph Lameter <cl@linux.com>
>> CC: David Rientjes <rientjes@google.com>
>> CC: Pekka Enberg <penberg@cs.helsinki.fi>
>
> Acked-by: David Rientjes <rientjes@google.com>

Applied, thanks!