Subject: Re: [PATCH 2/4] Add a __GFP_SLABMEMCG flag
Posted by Christoph Lameter on Mon, 11 Jun 2012 14:24:37 GMT
View Forum Message <> Reply to Message

On Sat, 9 Jun 2012, James Bottomley wrote:

> On Fri, 2012-06-08 at 14:31 -0500, Christoph Lameter wrote:
> > On Fri, 8 Jun 2012, Glauber Costa wrote:
> >
> > > */
> > > #define __GFP_NOTRACK_FALSE_POSITIVE (__GFP_NOTRACK)
> > >
> > > -#define __GFP_BITS_SHIFT 25 /* Room for N __GFP_FOO bits */
> > > +#define __GFP_BITS_SHIFT 26 /* Room for N __GFP_FOO bits */
> > > #define __GFP_BITS_MASK ((__force gfp_t)((1 << __GFP_BITS_SHIFT) - 1))
> >
> > Please make this conditional on CONFIG_MEMCG or so. The bit can be useful
> > in particular on 32 bit architectures.
>
> I really don't think that's at all a good idea. It's asking for trouble
> when we don't spot we have a flag overlap. It also means that we're
> trusting the reuser to know that their use case can never clash with
> CONFIG_MEMGC and I can't think of any configuration where this is
> possible currently.

Flag overlap can be avoided using the same method as we have done with the
page flags (which uses an enum). There are other uses of N bits after
GFP_BITS_SHIFT. On first look this looks like its 4 right now so we cannot
go above 28 on 32 bit platforms. It would also be useful to have that
limit in there somehow so that someone modifying the GFP_BITS sees the
danger.

> I think making the flag define of __GFP_SLABMEMCG conditional might be a
> reasonable idea so we get a compile failure if anyone tries to use it
> when !CONFIG_MEMCG.

Ok that is another reason to do so.