
Subject: Re: [PATCH v2 02/13] memcg: Kernel memory accounting infrastructure.
Posted by [KAMEZAWA Hiroyuki](#) on Tue, 13 Mar 2012 06:24:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sun, 11 Mar 2012 12:12:04 +0400

Glauber Costa <glommer@parallels.com> wrote:

```
> On 03/10/2012 12:39 AM, Suleiman Souhlal wrote:  
> > Enabled with CONFIG_CGROUP_MEM_RES_CTRLR_KMEM.  
> >  
> > Adds the following files:  
> > - memory.kmem.independent_kmem_limit  
> > - memory.kmem.usage_in_bytes  
> > - memory.kmem.limit_in_bytes  
> >  
> > Signed-off-by: Suleiman Souhlal<suleiman@google.com>  
> > ---  
> > mm/memcontrol.c | 136  
+++++  
> > 1 files changed, 135 insertions(+), 1 deletions(-)  
> >  
> > diff --git a/mm/memcontrol.c b/mm/memcontrol.c  
> > index 37ad2cb..e6fd558 100644  
> > --- a/mm/memcontrol.c  
> > +++ b/mm/memcontrol.c  
> > @@ -220,6 +220,10 @@ enum memcg_flags {  
> >     /*  
> >     MEMCG_MEMSW_IS_MINIMUM, /* Set when res.limit == memsw.limit */  
> >     MEMCG_OOM_KILL_DISABLE, /* OOM-Killer disable */  
> > + MEMCG_INDEPENDENT_KMEM_LIMIT, /*  
> > +     * kernel memory is not counted in  
> > +     * memory.usage_in_bytes  
> > +     */  
> > };
```

After looking codes, I think we need to think
whether independent_kmem_limit is good or not....

How about adding MEMCG_KMEM_ACCOUNT flag instead of this and use only
memcg->res/memcg->memsw rather than adding a new counter, memcg->kmem ?

if MEMCG_KMEM_ACCOUNT is set -> slab is accounted to mem->res/memsw.
if MEMCG_KMEM_ACCOUNT is not set -> slab is never accounted.

(I think On/Off switch is required..)

Thanks,

-Kame
