Subject: Re: [PATCH v4 23/25] memcg: propagate kmem limiting information to children
Posted by KAMEZAWA Hiroyuki on Mon, 18 Jun 2012 12:37:01 GMT
View Forum Message <> Reply to Message

(2012/06/18 19:28), Glauber Costa wrote:
> The current memcg slab cache management fails to present satisfatory hierarchical
> behavior in the following scenario:
>
> -> /cgroups/memory/A/B/C
>
> * kmem limit set at A
> * A and B empty taskwise
> * bash in C does find /
>
> Because kmem_accounted is a boolean that was not set for C, no accounting
> would be done. This is, however, not what we expect.
>

Hmm....do we need this new routines even while we have mem_cgroup_iter() ?

Doesn't this work ?

```
 struct mem_cgroup {
  .....
  bool kmem_accounted_this;
  atomic_t kmem_accounted;
  ....
 }
```

at set limit

```
 ....set_limit(memcg) {

  if (newly accounted) {
   mem_cgroup_iter() {
    atomic_inc(&iter->kmem_accounted)
   }
  } else {
   mem_cgroup_iter() {
    atomic_dec(&iter->kmem_accounted);
   }
 }
```

hm ? Then, you can see kmem is accounted or not by atomic_read(&memcg->kmem_accounted);

Thanks,

-Kame