Subject: Re: [PATCH 09/23] kmem slab accounting basic infrastructure
Posted by Glauber Costa on Wed, 02 May 2012 15:15:43 GMT
View Forum Message <> Reply to Message

```
>> @@ -3951,8 +3966,26 @@ static int mem_cgroup_write(struct cgroup *cont, struct cftype *cft,
>>                          break;
>>                  if (type == _MEM)
>>                          ret = mem_cgroup_resize_limit(memcg, val);
>> -                else
>> +                else if (type == _MEMSWAP)
>>                          ret = mem_cgroup_resize_memsw_limit(memcg, val);
>> +#ifdef CONFIG_CGROUP_MEM_RES_CTLR_KMEM
>> +                else if (type == _KMEM) {
>> +                        ret = res_counter_set_limit(&memcg->kmem, val);
>> +                        if (ret)
>> +                                break;
>> +                        /*
>> +                         * Once enabled, can't be disabled. We could in theory
>> +                         * disable it if we haven't yet created any caches, or
>> +                         * if we can shrink them all to death.
>> +                         *
>> +                         * But it is not worth the trouble
>> +                         */
>> +                        if (!memcg->kmem_accounted&&  val != RESOURCE_MAX)
>> +                                memcg->kmem_accounted = true;
>> +                }
>> +#endif
>> +                else
>> +                        return -EINVAL;
>>                  break;
>>          case RES_SOFT_LIMIT:
>>                  ret = res_counter_memparse_write_strategy(buffer,&val);
>
> Why is RESOURCE_MAX special?
```

Because I am using the convention that setting it to any value different
than that will enable accounting.