Subject: Re: [PATCH v2 02/13] memcg: Kernel memory accounting infrastructure. Posted by Glauber Costa on Mon, 12 Mar 2012 12:38:26 GMT View Forum Message <> Reply to Message

On 03/10/2012 12:39 AM, Suleiman Souhlal wrote: > +#ifdef CONFIG_CGROUP_MEM_RES_CTLR_KMEM > + int> +memcg_charge_kmem(struct mem_cgroup *memcg, gfp_t gfp, long long delta) > +{ > + struct res counter *fail res; > + struct mem_cgroup *_memcg; > + int may oom, ret; > + > + may_oom = (gfp& __GFP_WAIT)&& (gfp& __GFP_FS)&& !(gfp& __GFP_NORETRY); > + > + > + ret = 0;> + > + memcg = memcg; > + if (memcg&& !mem_cgroup_test_flag(memcg, MEMCG INDEPENDENT KMEM LIMIT)) { > + > + ret = mem cgroup try charge(NULL, gfp, delta / PAGE SIZE, > + & memcq, may oom); > + if (ret == -ENOMEM) > + return ret; > + } > + > + if (memcg&& memcg == memcg) > + ret = res counter charge(&memcg->kmem, delta,&fail res); > + > + return ret; > +} > + > +voidOk.

So I've spent most of the day today trying to come up with a way not to kill the whole performance we gain from consume_stock() by this res_counter_charge() to kmem afterwards...

You mentioned you want to still be able to bill to memcg->kmem mostly for debugging/display purposes. So we're surely not using all of the res_counter infrastructure (limiting, soft limits, etc)

I was thinking: Can't we have a percpu_counter that we use for this purpose when !kmem_independent ?

we may not even need to bloat the struct, since we can fold it into a

union with struct res_counter kmem (which is bigger than a percpu counter anyway).

We just need to be a bit more careful not to allow kmem_independent to change when we already have charges to any of them (but we need to do it anyway)

