Subject: Re: [PATCH v3 04/13] kmem accounting basic infrastructure
Posted by Michal Hocko on Wed, 26 Sep 2012 16:01:27 GMT
View Forum Message <> Reply to Message

On Wed 26-09-12 18:33:10, Glauber Costa wrote:
> On 09/26/2012 06:03 PM, Michal Hocko wrote:
> > On Tue 18-09-12 18:04:01, Glauber Costa wrote:
[...]
> >> @@ -4961,6 +5015,12 @@ mem_cgroup_create(struct cgroup *cont)
> >> int cpu;
> >> enable_swap_cgroup();
> >> parent = NULL;
> >> +
> >> +#ifdef CONFIG_MEMCG_KMEM
> >> + WARN_ON(cgroup_add_cftypes(&mem_cgroup_subsys,
> >> + kmem_cgroup_files));
> >> +#endif
> >> +
> >> if (mem_cgroup_soft_limit_tree_init())
> >> goto free_out;
> >> root_mem_cgroup = memcg;
> >> @@ -4979,6 +5039,7 @@ mem_cgroup_create(struct cgroup *cont)
> >> if (parent && parent->use_hierarchy) {
> >> res_counter_init(&memcg->res, &parent->res);
> >> res_counter_init(&memcg->memsw, &parent->memsw);
> >> + res_counter_init(&memcg->kmem, &parent->kmem);
> > >
> > Haven't we already discussed that a new memcg should inherit kmem_accounted
> > from its parent for use_hierarchy?
> > Say we have
> > root
> > |
> > A (kmem_accounted = 1, use_hierachy = 1)
> > \
> > B (kmem_accounted = 0)
> > \
> > C (kmem_accounted = 1)
> > >
> > B find's itself in an awkward situation becuase it doesn't want to
> > account u+k but it ends up doing so becuase C.
> > >
> >
> Ok, I haven't updated it here. But that should be taken care of in the
> lifecycle patch.

I am not sure which patch you are thinking about but I would prefer to
have it here because it is safe wrt. races and it is more obvious as
well.

--
Michal Hocko
SUSE Labs