
Subject: [PATCH v7 10/10] Disable task moving when using kernel memory accounting

Posted by [Glauber Costa](#) on Tue, 29 Nov 2011 23:57:01 GMT

[View Forum Message](#) <> [Reply to Message](#)

Since this code is still experimental, we are leaving the exact details of how to move tasks between cgroups when kernel memory accounting is used as future work.

For now, we simply disallow movement if there are any pending accounted memory.

Signed-off-by: Glauber Costa <glommer@parallels.com>
CC: Hiroyuki Kamezawa <kamezawa.hiroyu@jp.fujitsu.com>

mm/memcontrol.c | 23 ++++++-----
1 files changed, 22 insertions(+), 1 deletions(-)

```
diff --git a/mm/memcontrol.c b/mm/memcontrol.c
index a31a278..dd9a6d9 100644
--- a/mm/memcontrol.c
+++ b/mm/memcontrol.c
@@ -5453,10 +5453,19 @@ static int mem_cgroup_can_attach(struct cgroup_subsys *ss,
{
    int ret = 0;
    struct mem_cgroup *memcg = mem_cgroup_from_cont(cgroup);
+   struct mem_cgroup *from = mem_cgroup_from_task(p);
+
+   #if defined(CONFIG_CGROUP_MEM_RES_CTLR_KMEM) && defined(CONFIG_INET)
+   if (from != memcg && !mem_cgroup_is_root(from) &&
+       res_counter_read_u64(&from->tcp_mem.tcp_memory_allocated, RES_USAGE)) {
+       printk(KERN_WARNING "Can't move tasks between cgroups: "
+             "Kernel memory held.\n");
+   return 1;
+ }
+ #endif

    if (memcg->move_charge_at_immigrate) {
        struct mm_struct *mm;
-       struct mem_cgroup *from = mem_cgroup_from_task(p);

        VM_BUG_ON(from == memcg);

@@ -5624,6 +5633,18 @@ static int mem_cgroup_can_attach(struct cgroup_subsys *ss,
    struct cgroup *cgroup,
    struct task_struct *p)
{
+   struct mem_cgroup *memcg = mem_cgroup_from_cont(cgroup);
```

```
+ struct mem_cgroup *from = mem_cgroup_from_task(p);
+
+[#if defined(CONFIG_CGROUP_MEM_RES_CTLR_KMEM) && defined(CONFIG_INET)
+ if (from != memcg && !mem_cgroup_is_root(from) &&
+     res_counter_read_u64(&from->tcp_mem.tcp_memory_allocated, RES_USAGE)) {
+     printk(KERN_WARNING "Can't move tasks between cgroups: "
+           "Kernel memory held.\n");
+     return 1;
+ }
+[#endif
+
 return 0;
}
static void mem_cgroup_cancel_attach(struct cgroup_subsys *ss,
--
```

1.7.6.4
