
Subject: [PATCH 24/29] memory controller task migration v7
Posted by [Paul Menage](#) on Tue, 11 Sep 2007 19:53:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Balbir Singh <balbir@linux.vnet.ibm.com>

Allow tasks to migrate from one cgroup to the other. We migrate mm_struct's mem_cgroup only when the thread group id migrates.

Signed-off-by: Balbir Singh <balbir@linux.vnet.ibm.com>
Cc: Pavel Emelianov <xemul@openvz.org>
Cc: Paul Menage <menage@google.com>
Cc: Peter Zijlstra <a.p.zijlstra@chello.nl>
Cc: "Eric W. Biederman" <ebiederm@xmission.com>
Cc: Nick Piggin <nickpiggin@yahoo.com.au>
Cc: Kirill Korotaev <dev@sw.ru>
Cc: Herbert Poetzl <herbert@13thfloor.at>
Cc: David Rientjes <rientjes@google.com>
Cc: Vaidyanathan Srinivasan <svaidy@linux.vnet.ibm.com>
Signed-off-by: Andrew Morton <akpm@linux-foundation.org>

mm/memcontrol.c | 35 ++++++
1 files changed, 35 insertions(+)

diff -puN mm/memcontrol.c~memory-controller-task-migration-v7 mm/memcontrol.c
--- a/mm/memcontrol.c~memory-controller-task-migration-v7

+++ a/mm/memcontrol.c
@@ -326,11 +326,46 @@ static int mem_cgroup_populate(struct
ARRAY_SIZE(mem_cgroup_files));
}

+static void mem_cgroup_move_task(struct cgroup_subsys *ss,
+ struct cgroup *cont,
+ struct cgroup *old_cont,
+ struct task_struct *p)
+{
+ struct mm_struct *mm;
+ struct mem_cgroup *mem, *old_mem;
+
+ mm = get_task_mm(p);
+ if (mm == NULL)
+ return;
+
+ mem = mem_cgroup_from_cont(cont);
+ old_mem = mem_cgroup_from_cont(old_cont);
+
+ if (mem == old_mem)

```
+ goto out;
+
+ /*
+ * Only thread group leaders are allowed to migrate, the mm_struct is
+ * in effect owned by the leader
+ */
+ if (p->tgid != p->pid)
+ goto out;
+
+ css_get(&mem->css);
+ rcu_assign_pointer(mm->mem_cgroup, mem);
+ css_put(&old_mem->css);
+
+out:
+ mmput(mm);
+ return;
+}
+
+struct cgroup_subsys mem_cgroup_subsys = {
+ .name = "memory",
+ .subsys_id = mem_cgroup_subsys_id,
+ .create = mem_cgroup_create,
+ .destroy = mem_cgroup_destroy,
+ .populate = mem_cgroup_populate,
+ .attach = mem_cgroup_move_task,
+ .early_init = 1,
+};
```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>