Subject: [PATCH 1/3] don't attach a task to a dead cgroup
Posted by Glauber Costa on Thu, 19 Apr 2012 22:49:16 GMT
View Forum Message <> Reply to Message

Not all external callers of cgroup_attach_task() test to
see if the cgroup is still live - the internal callers at
cgroup.c does.

With this test in cgroup_attach_task, we can assure that
no tasks are ever moved to a cgroup that is past its
destruction point and was already marked as dead.

Signed-off-by: Glauber Costa <glommer@parallels.com>
CC: Tejun Heo <tj@kernel.org>
CC: Li Zefan <lizefan@huawei.com>
CC: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
---
 kernel/cgroup.c | 3 +++
 1 files changed, 3 insertions(+), 0 deletions(-)

diff --git a/kernel/cgroup.c b/kernel/cgroup.c
index b61b938..932c318 100644
--- a/kernel/cgroup.c
+++ b/kernel/cgroup.c
@@ -1927,6 +1927,9 @@ int cgroup_attach_task(struct cgroup *cgrp, struct task_struct *tsk)
 struct cgroup_taskset tset = { };
 struct css_set *newcg;

+ if (cgroup_is_removed(cgrp))
+ return -ENODEV;
+
 /* @tsk either already exited or can't exit until the end */
 if (tsk->flags & PF_EXITING)
  return -ESRCH;
--
1.7.7.6