Subject: [PATCH 3/12] Use find_task_by_vpid in taskstats
Posted by Pavel Emelianov on Tue, 29 Jan 2008 13:46:15 GMT
View Forum Message <> Reply to Message

The pid to lookup a task by is passed inside taskstats code
via genetlink message.

Since netlink packets are now processed in the context of the
sending task, this is correct to lookup the task with
find_task_by_vpid() here.

Besides, I fix the call to fill_pid() from taskstats_exit(),
since the tsk->pid is not required in fill_pid() in this case,
and the pid field on task_struct is going to be deprecated
as well.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

---
 kernel/taskstats.c | 6 +++---
 1 files changed, 3 insertions(+), 3 deletions(-)

diff --git a/kernel/taskstats.c b/kernel/taskstats.c
index 07e86a8..4a23517 100644
--- a/kernel/taskstats.c
+++ b/kernel/taskstats.c
@@ -183,7 +183,7 @@ static int fill_pid(pid_t pid, struct task_struct *tsk,

  if (!tsk) {
   rcu_read_lock();
- tsk = find_task_by_pid(pid);
+ tsk = find_task_by_vpid(pid);
   if (tsk)
    get_task_struct(tsk);
   rcu_read_unlock();
@@ -230,7 +230,7 @@ static int fill_tgid(pid_t tgid, struct task_struct *first,
   */
  rcu_read_lock();
  if (!first)
- first = find_task_by_pid(tgid);
+ first = find_task_by_vpid(tgid);

  if (!first || !lock_task_sighand(first, &flags))
   goto out;
@@ -547,7 +547,7 @@ void taskstats_exit(struct task_struct *tsk, int group_dead)
  if (!stats)
   goto err;

```
- rc = fill_pid(tsk->pid, tsk, stats);
+ rc = fill_pid(-1, tsk, stats);
  if (rc < 0)
   goto err;

--
1.5.3.4
```