

---

Subject: [PATCH 1/15] Move exit\_task\_namespaces()  
Posted by [Pavel Emelianov](#) on Thu, 26 Jul 2007 14:46:28 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Make task release its namespaces after it has reparented all his children to child\_reaper, but before it notifies its parent about its death.

The reason to release namespaces after reparenting is that when task exits it may send a signal to its parent (SIGCHLD), but if the parent has already exited its namespaces there will be no way to decide what pid to deliver to him - parent can be from different namespace.

The reason to release namespace before notifying the parent is that when task sends a SIGCHLD to parent it can call wait() on this task and release it. But releasing the mnt namespace implies dropping of all the mounts in the mnt namespace and NFS expects the task to have valid sighand pointer.

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

---

exit.c | 5 ++++-

1 files changed, 4 insertions(+), 1 deletion(-)

diff -upr linux-2.6.23-rc1-mm1.orig/kernel/exit.c linux-2.6.23-rc1-mm1-7/kernel/exit.c

--- linux-2.6.23-rc1-mm1.orig/kernel/exit.c 2007-07-26 16:34:45.000000000 +0400

+++ linux-2.6.23-rc1-mm1-7/kernel/exit.c 2007-07-26 16:36:37.000000000 +0400

@ @ -788,6 +804,10 @ @ static void exit\_notify(struct task\_stru

BUG\_ON(!list\_empty(&tsk->children));

BUG\_ON(!list\_empty(&tsk->ptrace\_children));

+ write\_unlock\_irq(&tasklist\_lock);

+ exit\_task\_namespaces(tsk);

+ write\_lock\_irq(&tasklist\_lock);

+

/\*

\* Check to see if any process groups have become orphaned

\* as a result of our exiting, and if they have any stopped

@ @ -999,7 +1021,6 @ @ fastcall NORET\_TYPE void do\_exit(long co

tsk->exit\_code = code;

proc\_exit\_connector(tsk);

- exit\_task\_namespaces(tsk);

exit\_notify(tsk);

#ifdef CONFIG\_NUMA

mpol\_free(tsk->mempolicy);

---