Subject: [PATCH] namespaces: fix race at task exit
Posted by serue on Thu, 25 Jan 2007 15:05:42 GMT
View Forum Message <> Reply to Message

In do_exit(), the exit_task_namespaces() was placed after
exit_notify() because exit_notify ends up using the pid
namespace both to access the reaper, and for detaching the
pid. However, this placement allows an nfs server to reap
the task before exit_task_namespaces() completes.

This patch moves the exit_task_namespaces() into release_task,
below release_thread() which puts the pids(), and just above
the call_rcu(delayed_put_task_struct). I believe this should
solve both problems.

Signed-off-by: Serge E. Hallyn <serue@us.ibm.com>

---

kernel/exit.c | 2 +-
1 files changed, 1 insertions(+), 1 deletions(-)

765277a4170d7bbd1c4613de661ec6ac64d5580a
diff --git a/kernel/exit.c b/kernel/exit.c
index 3540172..ab9ae30 100644
--- a/kernel/exit.c
+++ b/kernel/exit.c
@@ -174,6 +174,7 @@ repeat:
write_unlock_irq(&tasklist_lock);
proc_flush_task(p);
release_thread(p);
+ exit_task_namespaces(p);
call_rcu(&p->rcu, delayed_put_task_struct);

p = leader;
@@ -939,7 +940,6 @@ fastcall NORET_TYPE void do_exit(long co
tsk->exit_code = code;
proc_exit_connector(tsk);
exit_notify(tsk);
- exit_task_namespaces(tsk);
#ifdef CONFIG_NUMA
mpol_free(tsk->mempolicy);
tsk->mempolicy = NULL;
--
1.1.6

_____
Containers mailing list
Containers@lists.osdl.org