Subject: Re: [PATCH] namespaces: fix race at task exit
Posted by Cedric Le Goater on Thu, 25 Jan 2007 15:20:23 GMT
View Forum Message <> Reply to Message

Serge E. Hallyn wrote:
> 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>

I've run some tests on x86 and x86_64: mounted a NFS share after
having unshare(CLONE_NEWNS) and I didn't reproduce the bug Daniel
had found.

it looks safe.

C.