
Subject: Re: [PATCH 1/15] Move exit_task_namespaces()
Posted by [Pavel Emelianov](#) on Mon, 06 Aug 2007 13:38:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Oleg Nesterov wrote:

> On 08/06, Pavel Emelianov wrote:

>> If task X is exiting and has already exit_task_namespaces()-ed task
>> Y will OOPs during its exit in determining parent's namespace. I agree
>> that in that case this is not important what namespace X belongs to,
>> but we need to handle the race with changing the nsproxy from not-NULL
>> to NULL. This is OK to make this under task_lock() but what to add
>> extra locking for if we can avoid it?

>

> No, we can't take task_lock() under write_lock(tasklist).

I meant that we could save the namespace earlier (w/o tasklist),
carry it up to the place we need and release it later. But all
this is too complex and it's easier to get rid of children and
then release the namespaces.

> Oleg.

>

>
