

---

Subject: Re: [PATCH 1/15] Move exit\_task\_namespaces()  
Posted by [Pavel Emelianov](#) on Mon, 06 Aug 2007 13:36:51 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Oleg Nesterov wrote:

> On 08/06, Pavel Emelianov wrote:

>> Oleg Nesterov wrote:

>>> On 08/06, Pavel Emelianov wrote:

>>>> Oleg Nesterov wrote:

>>>>> On 08/06, Pavel Emelianov wrote:

>>>>>> Oleg Nesterov wrote:

>>>>>>> On 07/26, Pavel Emelianov wrote:

>>>>>>>> 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.

>>>>>>>> I almost forgot about this one...

>>>>>>>

>>> I guess I missed something stupid and simple...

>> In other words. Let task X live in init\_pid\_ns, task Y is his child and

>> lives

>> in another namespace. task X and task Y both die. This will happen:

>>

>> 1. Task X call exit\_task\_namespaces()

>> and sets its nsproxy to NULL

>

> Ah, got it, thanks. So the problem is not namespace itself (parent's or

> child's), there are still valid (even if different but related).

>

> We just can't get ->parent->nsproxy. I was greatly confused by the "parent

> can be from different namespace" above. We have exactly same problem if

> namespaces are not differ.

>

> IOW, the problem is: we can't clear ->nsproxy (exit\_task\_namespaces) until

> we get rid of ->children. This has nothing to do with different namespace.

No. If the parent is always in the same namespace we do not need to  
get its nsproxy :) Problem is exactly in that the parent's namespace  
is to be known.

> Oleg.

>

>

---