On 08/06, Pavel Emelyanov wrote:
>
> Oleg Nesterov wrote:
> >On 08/06, Pavel Emelyanov wrote:
> >>Oleg Nesterov wrote:
> >>>On 08/06, Pavel Emelyanov wrote:
> >>>>Oleg Nesterov wrote:
> >>>>>On 07/26, Pavel Emelyanov 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 dever 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
> int 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 have nothing to do with different namespace.

Oleg.