
Subject: Re: [PATCH 1/15] Move exit_task_namespaces()
Posted by [Oleg Nesterov](#) on Mon, 06 Aug 2007 09:54:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

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 deliver to him - parent can be from different namespace.

> >

> > I almost forgot about this one...

> >

> > After reading the whole series, I can't understand the above explanation
> > any longer. The parent can't be from different namespace, either we have
> > another sub-thread, or we reparent the child to /sbin/init which should
> > be from the same namespace.

>

> If the child that is a new namespace's init is exiting its parent is from the
> different namespace.

In that case it doesn't have childs. They were SIGKILL'ed before exit_notify().

> Moreover, we will probably want to implement "entering"
> the pid namespace, so having tasks with parents from another namespace will
> be OK.

Well. I saw this word "entering", but I don't know the meaning. Just curious,
could you explain?

And, if an exiting task has a child which is already from another namespace,
why can't we release our namespace before re-parenting? I guess I need to
know what "entering" means to understand this...

Oleg.
