
Subject: Re: [PATCH 0/13] Pid namespaces (OpenVZ view)

Posted by [ebiederm](#) on Thu, 24 May 2007 16:00:13 GMT

[View Forum Message](#) <> [Reply to Message](#)

Pavel Emelianov <xemul@openvz.org> writes:

> That's how OpenVZ sees the pid namespaces.
>
> The main idea is that kernel keeps operating with tasks pid
> as it did before, but each task obtains one more pid for each
> pid type - the virtual pid. When putting the pid to user or
> getting the pid from it kernel operates with the virtual ones.

Just a quick reaction.

- I would very much like to see a minimum of 3 levels of pids, being supported. Otherwise it is easy to overlook some of the cases that are required to properly support nesting, which long terms seems important.
- Semantically fork is easier then unshare. Unshare can mean a lot of things, and it is easy to pick a meaning that has weird side effects. Your implementation has a serious problem in that you change the value of getpid() at runtime. Glibc does not know how to cope with the value of getpid() changing.

Eric
