
Subject: Re: [RFC][PATCH 5/7] VPIDs: vpid/pid conversion in VPID enabled case
Posted by [Alexey Kuznetsov](#) on Fri, 03 Feb 2006 14:02:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello!

> And how bad would it be for openvz to move away from the vpid approach ? do
> you have any plan for it ?

Frankly speaking, using pair (container, pid) was the first thing, which we did (year ago), so that from viewpoint of core the switch is not a big deal. :-) However, it was rejected by several reasons:

1. Replacing all the references to pid with pair (container, pid) is quite expensive. F.e. it is possible that a task has a pid from one container, but it is in process group and/or session of another container, and its controlling terminal owner by another container. Grr..
So, the structures are bloated, the functions get additional arguments. And all this is for no real purpose, the functionality comparing with VPID is even reduced.

2. It is very inconvenient not to see processes inside VPS from host system. To do ps, strace, gdb etc. we have to move inside VPS. With VPID approach I can gdb even "init" process of VPS in a way invisible to VPS, see?
Well, and main problem is that gui administration and monitoring tools, which were existing for ages stop to work and require a major rewrite. Does it answer to question about plans for moving away?

To summarize: (container, pid) approach looks clean and consistent. At first sight I loved it, even thought it will solve some of problems with inter-container access control. But the devil is in details, I have to learn this again and again: access control must be separate of real engine, otherwise you get something which does not satisfy anyone.

> I've also seen that openvz introduces a 'vps_info_t' object, which looks
> like a some virtualization backend. I'm not sure to have well understood
> this framework. What the idea behind it ? is it to handle different
> implementation of the virtualization ?

openvz _is_ a complete implementation of virtualization (<http://openvz.org>), dealing with... essentially, everything, including even things like iptables. The VPID patch is just a small part of the whole openvz, and vps_info_t is just a part of large structure containing data essential for this problem.

Alexey
