

---

Subject: Re: [RFC][PATCH 5/7] VPIDs: vpid/pid conversion in VPID enabled case  
Posted by [Alexey Kuznetsov](#) on Mon, 06 Feb 2006 11:24:35 GMT

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

---

Hello!

> In the very, very rare cases where we can't do that (like a fork()  
> boundary), we do change the APIs to take both task and container.

I promise you much more of boundary cases, unless  
you make some windowsish `sys_fork_exec_deal_with_all_the_rest`(100 arguments)

Look how this works in `openvz`. It uses pure traditional unixish api.  
Fork is not a boundary at all. To enter to a container you  
do all the work in steps:

1. change accounting space (`sys_setuid()`)
2. plain `fork()`
2. tune communication (pipes, ptys etc)
3. `chroot()`
- ...
- N. enter container (at the moment it is `ioctl` on a special device, could be  
`syscall`).

You can omit any step, if you need. You can add entering any subsystem,  
which you invent in future. Simple and stupid. And, nevertheless, universal.

Alexey

---