

---

Subject: Re: [RFC][PATCH 5/7] VPIDs: vpid/pid conversion in VPID enabled case  
Posted by [serue](#) on Mon, 06 Feb 2006 14:51:04 GMT

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

---

Quoting Alexey Kuznetsov (kuznet@ms2.inr.ac.ru):

> > > 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?  
> >  
> > that's another container model issue again. your init process of a VPS  
> > could be the real init. why do you need a fake one ? just trying to  
> > understand all the issues you had to solve and I'm sure they are valid.  
>  
> It is not a fake init, it is a real init. Main goal of openvz is to allow  
> to start even f.e. the whole instance of stock redhat inside container  
> including standard init. It is not a strict architectural requirement,  
> but this option must be present.  
>  
> BTW the question sounds strange. I would think that in (container, pid)  
> approach among another limitations you get requirement that you must  
> have a child reaper inside container. With VPIDs this does not matter,  
> wait() made by parent inside host always returns global pid of child.  
> But with (container, pid) children can be reaped only inside container, right?

Nah, we can just special-case the lookup for pid==1 to always return the  
single real init, or insert the single real init into the pidhash for  
every container.

If we're going to provide the option which you need to have your own  
init, then I guess the best behavior is to add a flag to whatever  
mechanism creates a container specifying to either hash the real init  
into the pidhash, or make the first exec()d process the container's  
init? Does that seem reasonable?

thanks,  
-serge

---