

---

Subject: Re: [RFC][PATCH] VPIDs: Virtualization of PIDs (OpenVZ approach)  
Posted by [Herbert Poetzl](#) on Fri, 03 Feb 2006 03:01:43 GMT

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

---

On Thu, Feb 02, 2006 at 06:54:21PM +0300, Kirill Korotaev wrote:

> [RFC][PATCH] VPIDs: Virtualization of PIDs (OpenVZ approach)  
>  
> OpenVZ project would like to present another approach for VPIDs  
> developed by Alexey Kuznetsov. These patches were heavily tested and  
> used for half a year in OpenVZ already.  
>  
> OpenVZ VPIDs serves the same purpose of VPS (container) checkpointing  
> and restore. OpenVZ VPS checkpointing code will be released soon.  
>  
> These VPIDs patches look less intrusive than IBM's, so I would ask to  
> consider it for inclusion. Other projects can benefit from this code  
> as well.

well, IMHO this approach lacks a few things which would be very useful in a mainline pid virtualization, which pretty much explains why it is relatively small

- hierarchical structure (it's flat, only one level)
- a proper administration scheme
- a 'view' into the child pid spaces
- handling of inter context signalling

and, more important, it does not deal with the existing issues and error cases, where references to pids, tasks, task groups and sessions aren't handled properly ...

I think that in real world virtualization scenarios with hundreds of namespaces those 'imprecisions' will occasionally lead to very strange and random behaviour which in many cases will go completely unnoticed.

so I really prefer to cleanup the existing pid handling first, to avoid big surprises later ...

best,  
Herbert

> Patch set consists of 7 patches. The main patches are: -  
> diff-vpids-macro, which introduces pid access interface -  
> diff-vpids-core, which introduces pid translations where required  
> using new interfaces - diff-vpids-virt, which introduces virtual pids  
> handling in pid.c and pid mappings for VPSs  
>

> Patches are against latest 2.6.16-rc1-git6  
>  
> Kirill Korotaev,  
> OpenVZ team  
>  
> -  
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
> the body of a message to majordomo@vger.kernel.org  
> More majordomo info at <http://vger.kernel.org/majordomo-info.html>  
> Please read the FAQ at <http://www.tux.org/lkml/>

---