

---

Subject: Re: [RFC] Virtualization steps  
Posted by [ebiederm](#) on Fri, 24 Mar 2006 18:36:18 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Kirill Korotaev <dev@sw.ru> writes:

> Eric, Herbert,  
>  
> I think it is quite clear, that without some agreement on all these  
> virtualization issues, we won't be able to commit anything good to  
> mainstream. My idea is to gather our efforts to get consensus on most clean  
> parts of code first and commit them one by one.  
>  
> The proposal is quite simple. We have 4 parties in this conversation (maybe  
> more?): IBM guys, OpenVZ, VServer and Eric Biederman. We discuss the areas which  
> should be considered step by step. Send patches for each area, discuss, come to  
> some agreement and all 4 parties Sign-Off the patch. After that it goes to  
> Andrew/Linus. Worth trying?

Yes, this sounds like a path forward that has a reasonable chance of making progress.

> So far, (correct me if I'm wrong) we concluded that some people don't want  
> containers as a whole, but want some subsystem namespaces. I suppose for people  
> who care about containers only it doesn't matter, so we can proceed with  
> namespaces, yeah?

Yes, I think at one point I have seen all of the major parties receptive to the concept.

> So the most easy namespaces to discuss I see:  
> - utsname  
> - sys IPC  
> - network virtualization  
> - netfilter virtualization

The networking is hard simply because there is so very much of it, and it is being actively developed :)

> all these were discussed already somehow and looks like there is no fundamental  
> differences in our approaches (at least OpenVZ and Eric, for sure).

Yes. I think we agree on what the semantics should be for these parts. Which should avoid the problem with having the pid namespace.

> Right now, I suggest to concentrate on first 2 namespaces - utsname and  
> sysvipc. They are small enough and easy. Let's consider them without sysctl/proc  
> issues, as those can be resolved later. I sent the patches for these 2

> namespaces to all of you. I really hope for some \_good\_ critics, so we could  
> work it out quickly.

Sounds like a plan.

Eric

---