

---

Subject: Re: The issues for agreeing on a virtualization/namespaces implementation.

Posted by [dev](#) on Mon, 20 Feb 2006 14:25:35 GMT

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

---

> as does Linux-VServer currently, but do you have  
> any proof that putting all the fields together in  
> one big structure actually has any (dis)advantage  
> over separate structures?

have no proof and don't mind if there are many pointers. Though this doesn't look helpful to me as well.

>>mmm, how do you plan to pass additional flags to clone()?

>>e.g. strong or weak isolation of pids?

> do you really have to pass them at clone() time?

> would shortly after be more than enough?

> what if you want to change those properties later?

I don't think it is always suitable to do configuration later.

We had races in OpenVZ on VPS create/stop against exec/enter etc. (even introduced flag is\_running). So I have some experience to believe it will be painful place.

>>this syscalls will start handling this new namespace and that's all.

>>this is not different from many syscalls approach.

> well, let's defer the 'how many syscalls' issue to

> a later time, when we know what we want to implement :)  
agreed.

Kirill

---