

---

Subject: Re: [PATCH] Virtual ethernet tunnel

Posted by [Pavel Emelianov](#) on Thu, 07 Jun 2007 15:33:26 GMT

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

---

>>  
>> no one is against generic code and ability to create 2 interfaces in  
>> \*one\* namespace.  
>> (Like we currently allow to do so in OpenVZ)  
>>  
>> However, believe me, moving an interface is a \*hard\* operation. Much  
>> harder then netdev  
>> register from the scratch.  
>>  
>> Because it requires to take into account many things like:  
>> - packets in flight which requires synchronize and is slow on big  
>> machines  
>> - asynchronous sysfs entries registration/deregistration from  
>> rtln\_unlock -> netdev\_run\_todo  
>> - name/ifindex collisions  
>> - shutdown/cleanup of addresses/routes/qdisc and other similar stuff  
>>  
>>  
> All of what you are describing is already implemented in the Eric's  
> patchset.

Daniel. We \*agree\* that this task \*is implementable\*. We just want  
to say that creating the device in the namespace is less comp...  
(oh, well, forget this word) faster than creating and then moving.

> You can have a look at :  
>  
> [http://lxc.sourceforge.net/patches/2.6.20/2.6.20-netns1/broken\\_out/](http://lxc.sourceforge.net/patches/2.6.20/2.6.20-netns1/broken_out/)  
>  
> And more precisly:  
>  
> for sysfs issues:  
>  
> [http://lxc.sourceforge.net/patches/2.6.20/2.6.20-netns1/broken\\_out/0065-sysfs-Shadow-directory-support.patch](http://lxc.sourceforge.net/patches/2.6.20/2.6.20-netns1/broken_out/0065-sysfs-Shadow-directory-support.patch)  
>  
>  
> for network device movement:  
>  
> [http://lxc.sourceforge.net/patches/2.6.20/2.6.20-netns1/broken\\_out/0096-net-Implement-network-de  
vice-movement-between-namesp.patch](http://lxc.sourceforge.net/patches/2.6.20/2.6.20-netns1/broken_out/0096-net-Implement-network-device-movement-between-namesp.patch)

Good job. Can you prove that this code is less buggy than the existing  
register\_netdevice() one? This requires testing, doesn't it? But on the

other hand we have the ability to create the device right in the namespace using well known (and thus well debugged) code with minimal actions from the kernel.

>  
> Thanks,  
> Daniel  
> -  
> To unsubscribe from this list: send the line "unsubscribe netdev" in  
> the body of a message to majordomo@vger.kernel.org  
> More majordomo info at <http://vger.kernel.org/majordomo-info.html>  
>

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---