
Subject: Re: [PATCH] Virtual ethernet tunnel

Posted by [Daniel Lezcano](#) on Thu, 07 Jun 2007 15:44:37 GMT

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

Pavel Emelianov wrote:

> Daniel Lezcano wrote:

>> Pavel Emelianov wrote:

>>>> I did this at the very first version, but Alexey showed me that this
>>>> would be wrong. Look. When we create the second device it must be in
>>>> the other namespace as it is useless to have them in one namespace.
>>>> But if we have the device in the other namespace the RTNL_NEWLINK
>>>> message from kernel would come into this namespace thus confusing ip
>>>> utility in the init namespace. Creating the device in the init ns and
>>>> moving it into the new one is rather a complex task.

>>>>

>>>> Pavel,

>>>>

>>>> moving the netdevice to another namespace is not a complex task. Eric

>>>> Biederman did it in its patchset (cf. <http://lxc.sf.net/network>)

>>>>

>>>> By saying complex I didn't mean that this is difficult to implement,

>>>> but that it consists (must consist) of many stages. I.e. composite.

>>>> Making the device right in the namespace is liter.

>>>>

>>>>

>>>> When the pair device is created, both extremities are into the init

>>>> namespace and you can choose to which namespace to move one extremity.

>>>>

>>>> I do not mind that.

>>>>

>>>> When the network namespace dies, the netdev is moved back to the init

>>>> namespace.

>>>> That facilitate network device management.

>>>>

>>>> Concerning netlink events, this is automatically generated when the

>>>> network device is moved through namespaces.

>>>>

>>>> IMHO, we should have the network device movement between namespaces in

>>>> order to be able to move a physical network device too (eg. you have 4

>>>> NIC and you want to create 3 containers and assign 3 NIC to each of

>>>> them)

>>>>

>>>> Agree. Moving the devices is a must-have functionality.

>>>>

>>>> I do not mind making the pair in the init namespace and move the second

>>>> one into the desired namespace. But if we **always** will have two ends in

>>>> different namespaces what to complicate things for?

>>>>

>> Just to provide a netdev sufficiently generic to be used by people who
>> don't want namespaces but just want to do some network testing, like Ben
>> Greear does. He mentioned in a previous email, he will be happy to stop
>> redirecting people to out of tree patch.

>

> This patch creates booth devices in the init namespace. That's what
> you want, isn't it? When we have the namespaces we will be able to
> create the pair with booth ends in the init namespace - just do not
> specify the namespace id to create the 2nd end in and the driver will
> leave it in the init one.

Ok, fine.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
