

---

Subject: Re: Re: [RFC] network namespaces  
Posted by [Daniel Lezcano](#) on Mon, 11 Sep 2006 15:04:38 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Herbert Poetzl wrote:

> On Mon, Sep 11, 2006 at 04:40:59PM +0200, Daniel Lezcano wrote:  
>

>>I am currently working on this and I am finishing a prototype bringing  
>>isolation at the ip layer. The prototype code is very closed to  
>>Andrey's patches at TCP/UDP level. So the next step is to merge the  
>>prototype code with the existing network namespace layer 2 isolation.

>

>

> you might want to take a look at the current Linux-VServer  
> implementation for the network isolation too, should be  
> quite similar to Andrey's approach, but maybe you can  
> gather some additional information from there

ok, thanks. I will.

>>IHMO, the solution of splitting CONFIG\_NET\_NS into CONFIG\_L2\_NET\_NS  
>>and CONFIG\_L3\_NET\_NS is for me not acceptable because you will need  
>>to recompile the kernel. The proper way is certainly to have a  
>>specific flag for the unshare, something like CLONE\_NEW\_L2\_NET and  
>>CLONE\_NEW\_L3\_NET for example.

>

>

> I completely agree here, we need a separate namespace  
> for that, so that we can combine isolation and virtualization  
> as needed, unless the bind restrictions can be completely  
> expressed with an additional mangle or filter table (as  
> was suggested)

What is the bind restriction ? Do you want to force binding to a  
specific source address ?

-- Daniel

---

Containers mailing list  
Containers@lists.osdl.org  
<https://lists.osdl.org/mailman/listinfo/containers>

---