

---

Subject: Re: namespace support requires network modules to say "GPL"

Posted by [Stephen Hemminger](#) on Sat, 01 Dec 2007 19:17:36 GMT

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

---

On Sat, 01 Dec 2007 08:10:17 -0500

Mark Lord <lkml@rtr.ca> wrote:

> > Now that we have network namespace support merged it is time to  
> > revisit the sysfs support so we can remove the dependency on !SYSFS.  
> ...  
>  
> Now that the namespace updates are part of 2.6.24,  
> there is a major inconsistency in network EXPORT\_SYMBOLs.  
>  
> It used to be that an external network module could get away without  
> having to add a MODULE\_LICENSE("GPL\*") line to the source.  
>  
> In support of that, common networking functions (still) use EXPORT\_SYMBOL()  
> rather than the more restrictive EXPORT\_SYMBOL\_GPL().  
>  
> Eg. register\_netdev(), sk\_alloc(), \_\_dev\_get\_by\_name().  
>  
> But now, none of those three are actually usable by default,  
> because they all require "init\_net", which is EXPORT\_SYMBOL\_GPL().  
>

Then init\_net needs to be not GPL limited. Sorry, we need to allow non GPL network drivers. There is a fine line between keeping the binary seething masses from accessing random kernel functions, and allowing reasonable (but still non GPL) things like ndiswrapper to use network device interface.

--

Stephen Hemminger <shemminger@linux-foundation.org>

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"

Posted by [Alan Cox](#) on Sat, 01 Dec 2007 19:23:41 GMT

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

---

> Then init\_net needs to be not GPL limited. Sorry, we need to allow  
> non GPL network drivers. There is a fine line between keeping the

Why - they aren't exactly likely to be permissible by law

> binary seething masses from accessing random kernel functions, and allowing  
> reasonable (but still non GPL) things like ndiswrapper to use network  
> device interface.

Its up to the ndiswrapper authors how the licence their code, but they should respect how we licence ours.

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"

Posted by [Stephen Hemminger](#) on Sat, 01 Dec 2007 19:38:15 GMT

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

---

On Sat, 1 Dec 2007 19:23:41 +0000

Alan Cox <[alan@lxorguk.ukuu.org.uk](mailto:alan@lxorguk.ukuu.org.uk)> wrote:

> > Then init\_net needs to be not GPL limited. Sorry, we need to allow  
> > non GPL network drivers. There is a fine line between keeping the  
>  
> Why - they aren't exactly likely to be permissible by law

Matter of debate in which there are several opinions.

I don't like binary modules either, but don't feel qualified to render a legal opinion.

>  
> > binary seething masses from accessing random kernel functions, and allowing  
> > reasonable (but still non GPL) things like ndiswrapper to use network  
> > device interface.  
>  
> Its up to the ndiswrapper authors how the licence their code, but they  
> should respect how we licence ours.

Then change the license, explicitly and get it approved, forcing license changes by technically subversive means is bad policy. It is like Euro bureaucrats sneaking in software patents in regulations. If you want to have the debate and can get it resolved, then I support you.

Actually, the whole mess would go away if the api for dev\_get\_by\_XXXX hadn't been changed in the namespace transition. IMHO the interface to dev\_get\_by\_name() should not have added a namespace parameter, of the callers in the tree, only two use a different namespace. So it would have been better to introduce dev\_get\_by\_name\_ns() with the extra parameter.

Can we get this resolved before 2.6.24 is released? Going back and forth on API's is just needless frottage.

--

Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)>

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Alan Cox](#) on Sat, 01 Dec 2007 19:45:24 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> Then change the license, explicitly and get it approved, forcing license  
> changes by technically subversive means is bad policy. It is like Euro bureaucrats

I don't need to - the licence has been the same since about 0.12

Alan

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [ebiederm](#) on Sat, 01 Dec 2007 19:54:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> writes:

> On Sat, 01 Dec 2007 08:10:17 -0500  
> Mark Lord <[lkml@rtr.ca](mailto:lkml@rtr.ca)> wrote:  
>  
>> > Now that we have network namespace support merged it is time to  
>> > revisit the sysfs support so we can remove the dependency on !SYSFS.  
>> ...  
>>  
>> Now that the namespace updates are part of 2.6.24,  
>> there is a major inconsistency in network EXPORT\_SYMBOLs.  
>>  
>> It used to be that an external network module could get away without  
>> having to add a MODULE\_LICENSE("GPL\*") line to the source.  
>>  
>> In support of that, common networking functions (still) use EXPORT\_SYMBOL()

>> rather than the more restrictive EXPORT\_SYMBOL\_GPL().  
>>  
>> Eg. register\_netdev(), sk\_alloc(), \_\_dev\_get\_by\_name().  
>>  
>> But now, none of those three are actually usable by default,  
>> because they all require "init\_net", which is EXPORT\_SYMBOL\_GPL().

Which alternative kernel does the above comment apply to?

> Then init\_net needs to be not GPL limited. Sorry, we need to allow  
> non GPL network drivers.

For the record network drivers should not be affected. As a practical measure that just gets unmanageable and it is unnecessary.

There are specific exceptions where network drivers mess with the userspace interfaces where I do have some impact. However if you are messing with our userspace interface especially with network namespaces in place I don't see how it is possible for you to be anything other than a derivative work, and something we need in tree to keep maintenance a manageable thing.

It should just be the core of the network stack that struct net has some effect on.

> There is a fine line between keeping the  
> binary seething masses from accessing random kernel functions, and allowing  
> reasonable (but still non GPL) things like ndiswrapper to use network  
> device interface.

Does ndiswrapper break? If so what dubious and unsupportable thing is it doing?

Eric

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [ebiederm](#) on Sat, 01 Dec 2007 20:13:46 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> writes:

> Actually, the whole mess would go away if the api for dev\_get\_by\_XXXX hadn't  
> been changed in the namespace transition. IMHO the interface to  
> dev\_get\_by\_name()

> should not have added a namespace parameter, of the callers in the tree, only  
> two use a different namespace. So it would have been better to to introduce  
> dev\_get\_by\_name\_ns() with the extra parameter.

As a general rule if you are calling dev\_get\_by\_name and taking an &init\_net parameter that means you code has not yet been converted to actually support network namespaces.

Not everything can be safely changed at once so we take it by steps. When the code fully supports network namespaces practically nothing will take an &init\_net parameter. The network namespace parameter will come in some form from userspace. Either from current or from the network socket.

Except for boot time initialization I don't know of any cases using dev\_get\_by\_XXXX that won't need to be modified before the network namespace work is complete.

I believe I mentioned that this getting the fully network namespace support was going to take a while and a bunch of patches at the outset.

> Can we get this resolved before 2.6.24 is released? Going back and forth  
> on API's is just needless frottage.

Sure. We keep the updated dev\_get\_by\_XXXX that takes a network namespace parameter.

Or is their some legitimate usage of it by out of tree code that I'm not aware of?

Eric

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Mark Lord](#) on Sat, 01 Dec 2007 20:21:12 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Eric W. Biederman wrote:

> Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> writes:

>

>> Actually, the whole mess would go away if the api for dev\_get\_by\_XXXX hadn't  
>> been changed in the namespace transition. IMHO the interface to  
>> dev\_get\_by\_name()

>> should not have added a namespace parameter, of the callers in the tree, only  
>> two use a different namespace. So it would have been better to introduce  
>> dev\_get\_by\_name\_ns() with the extra parameter.  
>  
> As a general rule if you are calling dev\_get\_by\_name and taking an &init\_net  
> parameter that means your code has not yet been converted to actually support  
> network namespaces.  
>  
> Not everything can be safely changed at once so we take it by steps. When  
> the code fully supports network namespaces practically nothing will take  
> an &init\_net parameter. The network namespace parameter will come in  
> some form from userspace. Either from current or from the network  
> socket.  
>  
> Except for boot time initialization I don't know of any cases using  
> dev\_get\_by\_XXXX that won't need to be modified before the network  
> namespace work is complete.  
>  
> I believe I mentioned that this getting the fully network namespace  
> support was going to take a while and a bunch of patches at the  
> outset.  
>  
>> Can we get this resolved before 2.6.24 is released? Going back and forth  
>> on API's is just needless frottage.  
>  
> Sure. We keep the updated dev\_get\_by\_XXXX that takes a network  
> namespace parameter.  
..

And what should code be passing in when "# CONFIG\_NET\_NS is not set" ?

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Mark Lord](#) on Sat, 01 Dec 2007 22:13:13 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Arjan van de Ven wrote:  
> On Sat, 01 Dec 2007 15:21:12 -0500  
> Mark Lord <lkml@rtr.ca> wrote:  
>  
>> Eric W. Biederman wrote:  
>>> Stephen Hemminger <shemminger@linux-foundation.org> writes:  
>>> Sure. We keep the updated dev\_get\_by\_XXXX that takes a network  
>>> namespace parameter.

>> ..  
>>  
>> And what should code be passing in when "# CONFIG\_NET\_NS is not set" ?  
>  
> network drivers probably really really don't want to call  
> dev\_get\_by\_XXX...  
..

Fine. But all of them want to call sk\_alloc(),  
and many want to do register\_netdev().

So what should they be using there ?

And please STOP trimming the CC list.

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Stephen Hemminger](#) on Sun, 02 Dec 2007 00:30:35 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Sat, 1 Dec 2007 11:17:36 -0800  
Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> wrote:

> Then init\_net needs to be not GPL limited. Sorry, we need to allow  
> non GPL network drivers. There is a fine line between keeping the  
> binary seething masses from accessing random kernel functions, and allowing  
> reasonable (but still non GPL) things like ndiswrapper to use network  
> device interface.

>  
I spoke too soon earlier, ndiswrapper builds and loads against current  
2.6.24-rc3. Vmware and proprietary VPN software probably do not. Once again I don't  
give a damn, but the enterprise distro vendors certainly care.

--  
Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)>

--  
Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)>

---

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

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [ebiederm](#) on Sun, 02 Dec 2007 02:02:37 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> writes:

> On Sat, 1 Dec 2007 11:17:36 -0800  
> Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> wrote:  
>  
>> Then init\_net needs to be not GPL limited. Sorry, we need to allow  
>> non GPL network drivers. There is a fine line between keeping the  
>> binary seething masses from accessing random kernel functions, and allowing  
>> reasonable (but still non GPL) things like ndiswrapper to use network  
>> device interface.  
>>  
> I spoke too soon earlier, ndiswrapper builds and loads against current  
> 2.6.24-rc3. Vmware and proprietary VPN software probably do not. Once again I  
> don't  
> give a damn, but the enterprise distro vendors certainly care.

It looks like someone found a work around for vmware.

As for proprietary VPN software, the one case that came of David Miller  
said we had sufficient alternatives in the kernel that he didn't care.

We have to make a lot of changes to get the network namespaces  
complete. I don't have pity on any external code that breaks because  
they won't let me at their code.

If it turns out the code that was broken was an unacknowledged  
derivative work and it can't be fixed I have even less pity on  
them.

Eric

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Mark Lord](#) on Sun, 02 Dec 2007 03:34:09 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Stephen Hemminger wrote:

> On Sat, 1 Dec 2007 11:17:36 -0800  
> Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> wrote:  
>



>> Then init\_net needs to be not GPL limited. Sorry, we need to allow  
>> non GPL network drivers. There is a fine line between keeping the  
>> binary seething masses from accessing random kernel functions, and allowing  
>> reasonable (but still non GPL) things like ndiswrapper to use network  
>> device interface.  
>>  
> I spoke too soon earlier, ndiswrapper builds and loads against current  
> 2.6.24-rc3. Vmware and proprietary VPN software probably do not. Once again I don't  
> give a damn, but the enterprise distro vendors certainly care.  
...

Naw, enterprise (or any other) distro vendors shouldn't have any issues here,  
since they can just patch their kernels around any issues.

But it looks like Eric has this one thought out well enough.

Cheers

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Alan Cox](#) on Sun, 02 Dec 2007 13:51:04 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Sat, 1 Dec 2007 16:30:35 -0800  
Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> wrote:

> On Sat, 1 Dec 2007 11:17:36 -0800  
> Stephen Hemminger <[shemminger@linux-foundation.org](mailto:shemminger@linux-foundation.org)> wrote:  
>  
> > Then init\_net needs to be not GPL limited. Sorry, we need to allow  
> > non GPL network drivers. There is a fine line between keeping the  
> > binary seething masses from accessing random kernel functions, and allowing  
> > reasonable (but still non GPL) things like ndiswrapper to use network  
> > device interface.  
> >  
> I spoke too soon earlier, ndiswrapper builds and loads against current  
> 2.6.24-rc3. Vmware and proprietary VPN software probably do not. Once again I don't  
> give a damn, but the enterprise distro vendors certainly care.

Enterprise distro vendors ship kernels from the 2.6.19 era, so I don't  
see why they care.

Alan

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Valdis.Kletnieks](#) on Sun, 02 Dec 2007 19:56:42 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Sun, 02 Dec 2007 13:51:04 GMT, Alan Cox said:

> On Sat, 1 Dec 2007 16:30:35 -0800  
> > I spoke too soon earlier, ndiswrapper builds and loads against current  
> > 2.6.24-rc3. Vmware and proprietary VPN software probably do not. Once again I don't  
> > give a damn, but the enterprise distro vendors certainly care.  
>  
> Enterprise distro vendors ship kernels from the 2.6.19 era, so I don't  
> see why they care.

They don't care \*now\*. They will care when they try to rev forward from .19.

Not that they'll care a \*lot\* - it took \*me\* all of about an hour to get VMware Server 1.0.4 working under -rc3-mm2. Probably will take an enterprise distro 4-5 hours, 30 mins for the port and 4 1/2 hours for the paperwork. :)

---

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

---

---

Subject: RE: namespace support requires network modules to say "GPL"  
Posted by [David Schwartz](#) on Mon, 03 Dec 2007 00:02:17 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> > Then init\_net needs to be not GPL limited. Sorry, we need to allow  
> > non GPL network drivers. There is a fine line between keeping the

> Why - they aren't exactly likely to be permissible by law

Really? What law and/or what clause in the GPL says that derivative works have to be licensed under the GPL? Or does the kernel have some new technique to determine whether or not code has been distributed?

As I read the GPL, it only requires you to release something under the GPL if you distribute it. The kernel has no idea whether or not code has been

distributed. So if it's enforcing the GPL, it cannot prohibit anything non-distributed code can lawfully do. (Ergo, it's *\*NOT\* \*ENFORCING\** the GPL.)

> > binary seething masses from accessing random kernel functions,  
> and allowing  
> > reasonable (but still non GPL) things like ndiswrapper to use network  
> > device interface.  
>  
> Its up to the ndiswrapper authors how the licence their code, but they  
> should respect how we licence ours.

You license yours under the GPL, so they should respect the GPL.

It sounds like we're back to where we were years ago. Didn't we already agree that `EXPORT_SYMBOL_GPL` was *\*NOT\** a GPL-enforcement mechanism and had nothing to do with respecting the GPL? After all, if it s a GPL-enforcement mechanism, why is it not a "further restriction" which is prohibited by the GPL? (The GPL contains no restrictions on what code can use what symbols if that code is not distributed, but `EXPORT_SYMBOL_GPL` does.)

Are you now claiming that `EXPORT_SYMBOL_GPL` is intended to enforce the GPL?

DS

---

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

---

---

Subject: Re: namespace support requires network modules to say "GPL"  
Posted by [Romano Giannetti](#) on Mon, 03 Dec 2007 08:24:15 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Sat, 2007-12-01 at 22:34 -0500, Mark Lord wrote:

> Stephen Hemminger wrote:.  
> >  
> > I spoke too soon earlier, ndiswrapper builds and loads against current  
> > 2.6.24-rc3. Vmware and proprietary VPN software probably do not. Once again I don't  
> > give a damn, but the enterprise distro vendors certainly care.  
> ...  
>  
> Naw, enterprise (or any other) distro vendors shouldn't have any issues here,  
> since they can just patch their kernels around any issues.

Please pardon me for jumping in; I am not a kernel developer, but I try

to help with debugging whenever I can (and it's not just hand-waving, I helped to track down a couple of nasty bugs on MMC or ACPI EC, recently). And I am an engineer and IANAL, so I wouldn't speak about laws here. But I think it's not just a distribution's problem.

Unfortunately, I need VMware and ndiswrapper to get work done with my laptop. It's not the perfect world, but the only alternative is to boot in XP. So I normally stick with vendors kernels and, when I have time to "play" with new kernel, I go for it. If ndiswrapper and VMware work, perfect, I can test extensively the new kernel; if I find problems, I \*know\* I have to restart without proprietary modules, try to reproduce, report back. I did it a lot of times.

What I think is that every time VMware or (worst) ndiswrapper breaks, the kernel loose an awful lot of testers. In the span of time before Giri and the VMware team post a patch (-rc1 and -rc4, typically), my testing activity is just occasional. And I guess a lot of people is in the same situation.

These are just my 2cents. I will continue to test new kernels every time I can, and to use native solutions as often as I can (go, ath5k, go!; and LabWindows/CVI for Linux, anyone?). But maybe a bit of tolerance can help everyone...

Back in my hole,

Romano

--

--

indicado en la misma. Si Ud. no es el destinatario indicado, le informamos que cualquier forma de

This communication contains confidential information. It is for the exclusive use of the intended addressee. If you are not the intended addressee, please note that any form of distribution, copying or use of this communication or the information in it is strictly prohibited by law. If you have received this communication in error, please immediately notify the sender by reply e-mail and destroy this message. Thank you for your cooperation.

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

---