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

Posted by [Mark Lord](#) on Sat, 01 Dec 2007 23:51:28 GMT

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

Eric W. Biederman wrote:

> Mark Lord <lkml@rtr.ca> writes:

>

>> 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(),

>

> network drivers should be calling sk_alloc less then they should

> call dev_get_by_XXXX. Only protocols call sk_alloc.

..

I think I saw some bridge code that calls it, too.

Regardless, it's EXPORT_SYMBOL(), but now unusable
due to the GPL-only symbol "init_net".

Meanwhile, you are avoiding answering the question:

>>>> And what should code be passing in when "# CONFIG_NET_NS is not set" ?

..

>> and many want to do register_netdev().

>

> I haven't even touched register_netdev.

..

Bull-pucky. Somebody did:

> @@ -3361,6 +3595,8 @@ int register_netdevice(struct net_device *dev)

>

> /* When net_device's are persistent, this will be fatal. */

> BUG_ON(dev->reg_state != NETREG_UNINITIALIZED);

```
> +   BUG_ON(!dev->nd_net);
> +   net = dev->nd_net;
>
>   spin_lock_init(&dev->queue_lock);
>   spin_lock_init(&dev->_xmit_lock);
..
```

That new BUG_ON() line complains if nd_net has not been initialized, and the only thing I see drivers putting there is GPL-only "&init_net".

>> So what should they be using there ?

>
> What are you having problems with?

..

Avoiding the question again there, too.

I personally am not having problems with anything.
But I noticed the API change, and would like somebody to fix it.
Any of the three ways (or a fourth, if you've got one) might do.

Cheers

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