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

Posted by [Ben Greear](#) on Tue, 04 Dec 2007 18:57:01 GMT

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

Eric W. Biederman wrote:

> Ben Greear <greearb@candelatech.com> writes:

>

>

>>> Regardless of infringement it is incompatible with a complete network
>>> namespace implementation. Further it sounds like the module you are
>>> describing defines a kernel ABI without being merged and hopes that
>>> ABI will still be supportable in the future. Honestly I think doing so
>>> is horrible code maintenance policy.

>>>

>>>

>> I don't mind if the ABI changes, so long as I can still use something similar.

>>

>

> It has occurred to me that I am seeing an implication here that may in fact not
> exist.

>

> My impression of dev_get_by_xxxx is that the function is only able to be used
> sanely when being part of the definition of a kernel/userspace interface. With
> the further assumption on my part that you need to define a new instance of
> dev_get_by_xxxx

>

> It has just occurred to me that it is possible to reuse the SIOCBRADDIF
> and SIOCBRDELIF for that same purpose without defining a new kernel/userspace
> interface.

>

> What and how are you using dev_get_by_xxx?

>

I have a module that has a collection of 2-port bridges. These bridges
are used for emulation
purposes (somewhat similar to netem's feature set). Each bridge is
logically independent
of the others. To set up a bridge, I do something like:

```
echo add_my_bridge my_br1 eth0 eth1 > /proc/net/foo/config
```

Inside the module, it reads "eth0" and "eth1" and needs to find those
devices (ie, dev_get_by_name). It then registers to receive all pkts from
eth1 and transmit them on eth0, and vice versa.

If it would not require GPL symbols, I have no problem changing my API
to be something
like:

```
echo add_my_bridge my_br1 eth0 namespaceX eth1 namespaceY >  
/proc/net/foo/config
```

I am using procs so that I don't have to define any new 'official' kernel ABI, as that would more likely be a derivative work, and is a pain to keep up to date with changing kernels anyway...

Personally, it seems useful for my module to be able to have eth0 in one namespace and eth1 in another, but I won't complain if they both have to be in the same namespace or even just in the default namespace due to GPL symbol issues.

Thanks,
Ben

--

Ben Greear <greearb@candelatech.com>
Candela Technologies Inc <http://www.candelatech.com>

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