

---

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

Posted by [ebiederm](#) on Mon, 03 Dec 2007 17:35:21 GMT

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

---

Patrick McHardy <kaber@trash.net> writes:

> Ben Greear wrote:

>> I have a binary module that uses dev\_get\_by\_name...it's sort of a bridge-like

>> thing and

>> needs user-space to tell it which device to listen for packets on...

>>

>> This code doesn't need or care about name-spaces, so I don't see how it could

>> really

>> be infringing on the author's code (any worse than loading a binary driver

>> into the kernel

>> ever does).

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 would certainly prefer to not have to patch around any problems with calling

>> dev\_get\_by\_name

>> from a non-gpl module, but if required, I can probably figure something out...

>

>

> For all I care binary modules can break, but frankly I don't see

> how encapsulating a couple of structures and pointers in a new

> structure and adding a new argument to existing functions shifts

> the decision about how a function should be usable to the namespace

> guys. IMO all functions should continue to be usable as before,

> as decided by whoever actually wrote them. The only exception

> might be stuff where an existing EXPORT\_SYMBOL is clearly wrong,

> but that would be a separate discussion.

I don't think we have actually shifted the decision.

Further from a namespace perspective if I had to support out of tree modules and the current in kernel API the implementation would be impossible short of loading kernel modules multiple times once for each namespace. I totally refuse to give out of tree modules that power whatever their license.

Right now the network namespace code that has been merged isn't that interesting as it does not include ipv4 and ipv6 support which everyone uses.

One of the tests for completion of the network namespace work is grepping for `&init_net` and making certain we have cleanly removed all references to except in a handful of cases like the boot code.

Once things are largely complete it makes sense to argue with out of tree module authors that because they don't have network namespace support in their modules, their modules are broken.

Right now I suspect to many developers even of in-tree modules I have just shifted code around in an annoying looking way. I can completely see other developers not getting the point.

Eric

---

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

---