
Subject: Re: [PATCH 4/4] net: Make the loopback device per network namespace
Posted by [ebiederm](#) on Thu, 27 Sep 2007 16:48:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Denis V. Lunev" <den@sw.ru> writes:

> Eric W. Biederman wrote:
>> This patch makes loopback_dev per network namespace. Adding
>> code to create a different loopback device for each network
>> namespace and adding the code to free a loopback device
>> when a network namespace exits.
>>
>> This patch modifies all users the loopback_dev so they
>> access it as init_net.loopback_dev, keeping all of the
>> code compiling and working. A later pass will be needed to
>> update the users to use something other than the initial network
>> namespace.
>
> A pity that an important bit of explanation is missed. The
> initialization of loopback_dev is moved from a chain of devices
> (init_module) to a subsystem initialization to keep proper order, i.e.
> we must be sure that the initialization order is correct.

That didn't happen in the patch you mentioned. That happened
when we started dynamically allocating the loopback device.
That was the patch Daniel sent out a bit ago.

There are certainly some ordering issues and it may have helped
to talk about them. But they are because things assume the
loopback device is present. We have various bits of code that
is around such as the dst_ifdown case that assumes if another
network device is present the loopback device is present. To
fulfill that assumption I guess that means we have both an
initialization order dependency and a destruction order dependency.

The fact we were using module_init before actually appears to
me to have been racy, but we got away with it because the actual
data structure was statically allocated.

Since it appears that for a dynamically allocated loopback
registering it first and unregistering it last is necessary
for routing. It is likely worth looking at this a little
more closely and making a guarantee. So we can make it easier
for networking layers like ipv6, that want to memorize which
device is the loopback device.

Eric

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