
Subject: Re: Network Namespace status
Posted by [ebiederm](#) on Sun, 16 Sep 2007 23:47:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Miller <davem@davemloft.net> writes:

```
> From: ebiederm@xmission.com (Eric W. Biederman)
> Date: Thu, 13 Sep 2007 13:12:08 -0600
>
>> The final blocker to having multiple useful instances of network
>> namespaces is the loopback device. We recognize the network namespace
>> of incoming packets by looking at dev->nd_net. Which means for
>> packets to properly loopback within a network namespace we need a
>> loopback device per network namespace. There were some concerns
>> expressed when we posted the cleanup part of the patches that allowed
>> for multiple loopback devices a few weeks ago so resolving this one
>> may be tricky.
>
> There was a change posted recently to dynamically allocate the
> loopback device. I like that (sorry I don't have a reference
> to the patch handy), and you can build on top of that to get
> the namespace local loopback objects you want.
>
> static struct net_device *loopback_dev(struct net_namespace *net)
> {
> ...
> }
>
> You get the idea.
```

Sure. Thanks.

Since the change got dropped I figured it for a rejection, and that I would have to rework that patch.

On a similar note. It recently occurred to me that I can make creating multiple network namespaces depend on !CONFIG_SYSFS. Which will allow most of the rest of the patches I am sure of to be merged now. And give me just a little more time to work with Tejun and finish up the sysfs support.

Eric

Containers mailing list
Containers@lists.linux-foundation.org

