
Subject: Re: [PATCH 13/16] net: Make the device list and device lookups per namespace.

Posted by [davem](#) on Wed, 12 Sep 2007 11:39:47 GMT

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

From: ebiederm@xmission.com (Eric W. Biederman)

Date: Sat, 08 Sep 2007 15:35:46 -0600

>
> This patch makes most of the generic device layer network
> namespace safe. This patch makes dev_base_head a
> network namespace variable, and then it picks up
> a few associated variables. The functions:
> dev_getbyhwaddr
> dev_getfirsthwbytype
> dev_get_by_flags
> dev_get_by_name
> __dev_get_by_name
> dev_get_by_index
> __dev_get_by_index
> dev_ioctl
> dev_ethtool
> dev_load
> wireless_process_ioctl
>
> were modified to take a network namespace argument, and
> deal with it.
>
> vlan_ioctl_set and briocli_set were modified so their
> hooks will receive a network namespace argument.
>
> So basically anything in the core of the network stack that was
> affected to by the change of dev_base was modified to handle
> multiple network namespaces. The rest of the network stack was
> simply modified to explicitly use &init_net the initial network
> namespace. This can be fixed when those components of the network
> stack are modified to handle multiple network namespaces.
>
> For now the ifindex generator is left global.
>
> Fundametally ifindex numbers are per namespace, or else
> we will have corner case problems with migration when
> we get that far.
>
> At the same time there are assumptions in the network stack
> that the ifindex of a network device won't change. Making
> the ifindex number global seems a good compromise until
> the network stack can cope with ifindex changes when

> you change namespaces, and the like.

>

> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

Applied to net-2.6.24, thanks.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
