
Subject: [patch 3/3][netns] net: hide master/linked interface from netlink
Posted by [Daniel Lezcano](#) on Mon, 24 Sep 2007 12:21:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

From: Daniel Lezcano <dlezcana@fr.ibm.com>

Actually when a network device is linked to another, the name appears to be @<link>. For example, if a macvlan0 is created on top of eth0, the ip link show is:

```
6: macvlan0@eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop
    link/ether 6a:d4:10:0d:a8:55 brd ff:ff:ff:ff:ff:ff
```

But if we move macvlan0 to a network namespace, eth0 does no longer exist inside it and the result will be:

```
6: macvlan0@if2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop
    link/ether 6a:d4:10:0d:a8:55 brd ff:ff:ff:ff:ff:ff
```

if2 is, I guess, some random value. That can do invalid memory access or inconsistent data showing.

The patchset will avoid such case, it checks if the linked device exist into the current network namespace and if it doesn't the result will be:

```
6: macvlan0@NONE: <BROADCAST,MULTICAST> mtu 1500 qdisc noop
    link/ether 6a:d4:10:0d:a8:55 brd ff:ff:ff:ff:ff:ff
```

Signed-off-by: Daniel Lezcano <dlezcana@fr.ibm.com>

```
net/core/rtnetlink.c | 14 ++++++++
1 file changed, 11 insertions(+), 3 deletions(-)
```

Index: linux-2.6-netns/net/core/rtnetlink.c

```
=====
--- linux-2.6-netns.orig/net/core/rtnetlink.c
+++ linux-2.6-netns/net/core/rtnetlink.c
@@ -636,6 +636,8 @@ static int rtnl_fill_ifinfo(struct sk_bu
{
    struct ifinfomsg *ifm;
    struct nlmsghdr *nlh;
+   int ifindex = 0;
+   struct net_device *d;

    nlh = nlmsg_put(skb, pid, seq, type, sizeof(*ifm), flags);
    if (nlh == NULL)
@@ -656,11 +658,17 @@ static int rtnl_fill_ifinfo(struct sk_bu
```

```
NLA_PUT_U8(skb, IFLA_LINKMODE, dev->link_mode);
NLA_PUT_U32(skb, IFLA_MTU, dev->mtu);

- if (dev->ifindex != dev->iflink)
- NLA_PUT_U32(skb, IFLA_LINK, dev->iflink);
+ if (dev->ifindex != dev->iflink) {
+ d = dev_get_by_index(dev->nd_net, dev->iflink);
+ ifindex = d?dev->iflink:0;
+ NLA_PUT_U32(skb, IFLA_LINK, ifindex);
+ }

- if (dev->master)
+ if (dev->master) {
+ d = dev->master;
+ ifindex = dev->nd_net == d->nd_net?dev->master->ifindex:0;
 NLA_PUT_U32(skb, IFLA_MASTER, dev->master->ifindex);
+ }

if (dev->qdisc_sleeping)
 NLA_PUT_STRING(skb, IFLA_QDISC, dev->qdisc_sleeping->ops->id);
```

--

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