

Re: Multiple NICS on same subnet

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.networking/2006-03/msg00440.html>

- *From:* floyd@xxxxxxxxxx (Floyd L. Davidson)
 - *Date:* Wed, 15 Mar 2006 06:28:57 -0900
-

andrew.bell.ia@xxxxxxxxxx wrote:

Hi,

I'm trying to understand the behavior of a configuration where I have multiple NICs on the same subnet. (I know, its wrong, I'm just trying to understand.)

Here's the configuration:

```
eth0 : 172.16.11.4
eth1: 129.143.4.5
eth2: 129.143.4.6
```

```
$ netstat -rn
Kernel IP routing table
Destination Gateway Genmask Flags MSS Window
  irrt Iface
129.143.4.0 0.0.0.0 255.255.255.0 U 0 0
0 eth1
```

The fact that this entry for eth1 comes before the next entry, for eth2, means that (since the routing is identical), all packets to the 129.143.4.x subnet are routed via eth1, and none are routed to eth2.

```
129.143.4.0 0.0.0.0 255.255.255.0 U 0 0
0 eth2
```

No packets are **ever** route to this interface.

```
172.16.11.0 0.0.0.0 255.255.255.0 U 0 0
0 eth0
```

Re: Multiple NICS on same subnet

All packets to the 172.16.11.x subnet are routed to this interface

```
0.0.0.0 129.143.4.254 0.0.0.0 UG 0 0
0 eth1
```

This is a "default" route, which sends anything not matched by the above routes to IP address 129.143.4.254 via the eth1 interface.

eth2 is not cabled. eth0 and eth1 are cabled.

That will be fine, given the above route table.

When you ping the address of eth2 (129.143.4.6), you get a response.

If sent on the same machine, the packet actually goes to the lo device, not the eth2 device! The internals of the TCP/IP stack are of course aware of any address assigned to an interface on this machine, and any packet routed to an IP on the machine goes to the lo device. (Do and "ifconfig lo" command, note the counters, do a ping to any local IP address, and then check the counters again with ifconfig. You'll see that they increment nicely...)

If you look at the arp entry on the machine that originated the ping for 129.143.4.6, you get the MAC of eth0, which isn't even on the same logical network. No proxy ARP is turned on.

Can someone explain what is happening inside the stack to make this happen?

Sounds as if you have IP forwarding enabled, and the machine from which you sent the ping is gatewayed to 129.143.4.6 via eth0 (172.16.11.0).

Ping that address again, and check to lo device counts as described above. Interesting...

—
Floyd L. Davidson <http://www.apaflo.com/floyd_davidson>
Ukpeagvik (Barrow, Alaska) floyd@xxxxxxxxxx