

Re: Traffic NOT moving through the correct network interface?

## Re: Traffic NOT moving through the correct network interface?

---

*Source:* <http://linux.derkeiler.com/Newsgroups/alt.os.linux/2008-01/msg00220.html>

---

- *From:* [ibuprofin@xxxxxxxxxxxxxxxxxxxxxxxxxxxx](mailto:ibuprofin@xxxxxxxxxxxxxxxxxxxxxxxxxxxx) (Moe Trin)
  - *Date:* Fri, 11 Jan 2008 18:30:43 -0600
- 

On Fri, 11 Jan 2008, in the Usenet newsgroup alt.os.linux, in article <[BPahj.53359\\$uV6.49818@pd7urf1no](mailto:BPahj.53359$uV6.49818@pd7urf1no)>, Calab wrote:

```
ROOT:~#route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
10.0.11.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
68.145.52.0 0.0.0.0 255.255.252.0 U 0 0 0 eth1
68.145.56.0 0.0.0.0 255.255.252.0 U 0 0 0 eth2
239.0.0.0 0.0.0.0 255.0.0.0 U 0 0 0 eth0
```

Comment: 239.0.0.0/8 is part of Multicast, and only 239.192.0.0/10 is in current use – Site Local stuff under RFC2365. You're also missing the loopback interface, but that's not part of this problem.

```
0.0.0.0 68.145.52.1 0.0.0.0 UG 0 0 0 eth1
0.0.0.0 68.145.56.1 0.0.0.0 UG 0 0 0 eth2
0.0.0.0 10.0.11.1 0.0.0.0 UG 0 0 0 eth0
```

And there's your problem. Neither the interface, or the application know anything about routing. Everything is done by the kernel. If the packet is destined for 10.0.11.0 to 10.0.11.255, this table says that the remote host is directly attached to the eth0 interface. 68.145.52.0 to 68.145.55.255 is directly attached to eth1, and 68.145.56.0 to 68.145.59.255 is on eth2. 239.0.0.0 to 239.255.255.255 is also directly attached to the eth0 interface. That far – no problem. So how is the kernel going to send something that is going to 64.59.135.176? Well, that's not directly attached to anything here, so we have to look at using a gateway... lookie here – we got three of them, and they all lead to 0.0.0.0/0 ("everywhere"). Hmmmm, they all have the same metric too... so which do we use? The assumption the kernel makes is that you changed your mind – you first thought 68.145.52.1 was the way to the world, and then you decided (later) that "no, the world is reachable via 68.145.56.1" so we'll ignore 68.145.52.1. Then, you changed your mind

Re: Traffic NOT moving through the correct network interface?

## Re: Traffic NOT moving through the correct network interface?

again, and said "the world is reachable via 10.0.11.1" – right, we'll ignore the other two.

And why should the kernel use this in preference to eth0?

Because the connections on eth1 and eth2 are INBOUND.

kernel doesn't care about inbound. It looks at that routing table, and says how do I get "there".

eth0 is a gigabit ethernet connection to my LAN. It is used for file access within the LAN and any incidental (unplanned) outbound internet traffic such as package retrieval, etc.

eth2 is USB to cable modem #1. This is the inbound connection used for services accessible from the internet – web, ssh, ftp, etc.

As said – inbound means nothing as to how the stuff goes back out. When you have multiple ways to the world, you need one of three modes of operation:

Bonding – use is combined, and packets are shoved out more or less at random, but *\_equally\_* among the interfaces.

Fall-over – Use one interface, unless the network goes down, in which case, use a second as a backup (this needs a routing daemon on your end, and a gateway protocol between you and your ISP).

Policy-based – route web traffic through interface 1, ftp through 2, ssh through 3 and so on (this requires *iproute2*).

Bonding and fall-over require cooperation of your upstreams, and can be a real bear to get running cleanly. Policy based routing requires you to set up a policy daemon – this is lightly covered in that HOWTO I told you to look at. The *iproute2* package also has a lot of documentation.

The other problem is how you are to be reached from the Internet. I know 68.144.0.0/13 is Shaw, though I thought that was all residential services. Grabbing a hostname at random... oh, wonderful – shaw is calling these *\_really\_* generic looking, that's going to prevent you from using them for mail service. Anyway, let's say you've got one address called foo.bar.example.com. Which interface do you expect that to come in on? Well, obviously, you'll need three real addresses – one that reaches each of these NICs. So people out here on the Internet are going to need to know which of those addresses leads to the correct hose to find the application they are looking to connect to. If they choose the wrong hostname, they don't connect. That's going to make things messy – whoever is running the DNS is going to love it.

Re: Traffic NOT moving through the correct network interface?

## Re: Traffic NOT moving through the correct network interface?

All three interfaces are on separate subnets... Shouldn't make a difference if the IP's are actually real.

You identified part of the problem – same subnets. But that only screws up trying to reach hosts on that subnet. The other problem is those three default routes – and that's where your problem lies.

I assumed that outbound connections would pass through the same interface as the inbound connection that spawned them.

Never has worked that way. Same concept as trying to ping your own interface from the same computer. Let's assume you have addresses 10.0.11.11, 68.145.52.133 and 68.145.57.214, and from this computer, you try to ping each one. Which interface will be used? None of them. Why? Because your kernel is the one running the networking, and it knows the addresses that it "owns". So when it's trying to talk to itself, it uses the loopback. You could take a blowtorch and burn the NICs to a cinder, and as long as you don't damage the computer pinging each address works just fine. The kernel is doing the job, not the NIC or the application.

In your Fri, 11 Jan 2008 05:59:14 GMT post back to Unruh, you mention that the default route has shifted to eth1. You're running DHCP, aren't you. That reset the default route. DHCP is lousy if you are trying to run servers. You'll probably want to speak to Shaw about getting static IP addresses for this.

You also ask

1) Is it possible to route based on port number or service? Have all outbound traffic on port 12345 through eth1 and all traffic on port 54321 through eth2, for example?

```
[compton ~]$ whatis ip tc
ip (8) – show / manipulate routing, devices, policy routing and tunnels
tc (8) – show / manipulate traffic control settings
[compton ~]$
```

You're in for hours and hours of fun – good luck.

2) Can I permanently specify that eth1 be the default gateway? I had asked about specifying metrics for interfaces back in September in another group and was basically told that I didn't need to.

Re: Traffic NOT moving through the correct network interface?

Re: Traffic NOT moving through the correct network interface?

Yes – only specify that one as the default, and not the other two. If you are using DHCP, that's going to be messy, and this will be much more complicated if you are using policy routing.

Old guy

.