



How to connect two network interfaces on the same subnet?

🔒 SOLUTION VERIFIED - Updated April 28 2016 at 2:05 AM - English ▾()

Environment

- Red Hat Enterprise Linux 6
- Red Hat Enterprise Linux 5
- Multiple network interfaces, each with an IP address in the same subnet

Issue

- How to connect two network interfaces on the same subnet?
- In our environment, there are three bonding devices connected with the same segment.
- We have captured packet and found that packet should be transmitted from bond0 was actually transmitted from bond1
- Also, it was confirmed that the transmit port staggered even if there was no bonding setting.

```

+-----+
|           Linux           |
|   .168       .169       |
+-----+
|           |           |
+-----+
|           Switch          |
+-----+
|           |           |
+-----+
|           .1             |
|           Gateway        |
+-----+

```

Resolution

Add routing tables and rules binding source IP address for each route, and add those as default gateway for each network interface.

Assuming this networking environment :

```

+-----+
|           Linux           |
|   eth0           eth1     |
| 10.64.208.180    10.64.208.208 |
+-----+

```

```

# ip addr show
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
   inet 10.64.208.180/24 brd 10.65.211.255 scope global eth0
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
   inet 10.64.208.208/24 brd 10.65.211.255 scope global eth1

```

- Add new routing tables in `/etc/iproute2/route`

```

# cat /etc/iproute2/route

100 t1
101 t2

```

- Add routes to those routing tables

```

# ip route add 10.64.208.0/24 dev eth0 src 10.64.208.180 table t1
# ip route add table t1 default via 10.64.208.254 dev eth0
# ip route show table t1
10.64.208.0 dev eth0 scope link src 10.64.208.180
default via 10.64.208.254 dev eth0

```



```
# ip route add 10.64.208.0/24 dev eth1 src 10.64.208.208 table t2
# ip route add table t2 default via 10.64.208.254 dev eth1
# ip route show table t2
10.64.208.0 dev eth1 scope link src 10.64.208.208
default via 10.64.208.254 dev eth0
```

- Add rules to apply traffic to the routing tables

```
# ip rule add table t1 from 10.64.208.180
# ip rule add table t2 from 10.64.208.208
# ip route show
10.64.208.0/24 dev eth0 proto kernel scope link src 10.64.208.180
10.64.208.0/24 dev eth1 proto kernel scope link src 10.64.208.208
169.254.0.0/16 dev eth1 scope link
default via 10.64.208.254 dev eth0
```

- Set interfaces ready for receiving ARP replies

```
# sysctl net.ipv4.conf.default.arp_filter=1
```

- Checking ping with -I IPADDR

```
# ping -I 10.64.208.180 DSTADDR
```

- To make this routes persistent following configuration files have to be changed

For network addresses and routes:

```
# cat /etc/sysconfig/network-scripts/ifcfg-eth*

# ifcfg-eth0
DEVICE=eth0
BOOTPROTO=none
ONBOOT=yes
NETMASK=255.0.0.0
IPADDR=10.64.208.180
GATEWAY=10.64.208.254
TYPE=Ethernet

# ifcfg-eth1
DEVICE=eth1
BOOTPROTO=none
ONBOOT=yes
NETMASK=255.0.0.0
IPADDR=10.64.208.208
GATEWAY=10.64.208.254
TYPE=Ethernet
```

```
# cat /etc/sysconfig/network-scripts/route-eth*
# route-eth0
10.0.0.0/8 dev eth0 src 10.64.208.180 table t1
default via 10.64.208.254 dev eth0 table t1

# route-eth1
10.0.0.0/8 dev eth1 src 10.64.208.208 table t2
default via 10.64.208.254 dev eth1 table t2
```

For routing rules:

```
# cat /etc/sysconfig/network-scripts/rule-eth*
# rule-eth0
table t1 from 10.64.208.180

# rule-eth1
table t2 from 10.64.208.208
```

For receiving ARP replies:

```
# grep arp_filter /etc/sysctl.conf
net.ipv4.conf.all.arp_filter = 1
net.ipv4.conf.default.arp_filter = 1
```

For sending ARP:

```
# grep /etc/sysctl.conf
net.ipv4.conf.all.arp_announce = 2
net.ipv4.conf.default.arp_announce = 2
```



Refer to </usr/share/doc/kernel-doc-/Documentation/networking/ip-sysctl.txt> for more information about these settings.

Root Cause

- When there are 2 interfaces on the same subnet there is no assurance as to which interface will be used to transmit traffic and the machine will accept traffic for either IP on either interface.
- This is because in Linux the IP address belongs to the host and is not associated with the interface.
- If you ping with `-I DEV`, attempting to use a given interface, there is no guarantee the reply packet (if there even is one) will come back to the same interface, so pings done with `-I DEV` may not work.

Diagnostic Steps

- Setup system with 2 interfaces on the same subnet.
- Ping a target and capture packets with `tcpdump`.

Product(s) [Red Hat Enterprise Linux \(/taxonomy/products/red-hat-enterprise-linux\)](/taxonomy/products/red-hat-enterprise-linux) **Component** [initscripts \(/components/initscripts\)](/components/initscripts) [kernel \(/components/kernel\)](/components/kernel)

Category [Learn more \(/category/learn-more\)](/category/learn-more)

Tags [network \(/tags/network\)](/tags/network) [networking \(/tags/networking\)](/tags/networking) [network_stack \(/tags/network_stack\)](/tags/network_stack) [route \(/tags/route\)](/tags/route) [routing \(/tags/routing\)](/tags/routing)

This solution is part of Red Hat's fast-track publication program, providing a huge library of solutions that Red Hat engineers have created while supporting our customers. To give you the knowledge you need the instant it becomes available, these articles may be presented in a raw and unedited form.

[Privacy Policy \(http://www.redhat.com/en/about/privacy-policy\)](http://www.redhat.com/en/about/privacy-policy)

[Customer Portal Terms of Use \(https://access.redhat.com/help/terms/\)](https://access.redhat.com/help/terms/)

[All Policies and Guidelines \(http://www.redhat.com/en/about/all-policies-guidelines\)](http://www.redhat.com/en/about/all-policies-guidelines)

Copyright © 2017 Red Hat, Inc.