

Lab — Linux routing, hands-on

Lesson 02 · ~12 min · Single host, uses ip netns

This lab modifies kernel routing. Run on a throwaway VM if you can. All namespaces and veth pairs vanish on reboot.

Step 1 — Read your current routing table

```
ip route show
ip route show table all      # all tables
ip rule show
```

Step 2 — Ask the kernel a "what would you do?" question

```
ip route get 8.8.8.8
# Example output: 8.8.8.8 via 192.168.1.1 dev eth0 src 192.168.1.50 ...
```

Step 3 — Add a bogus route, watch the answer change

```
sudo ip route add 10.20.0.0/16 via 192.168.1.254 dev eth0
ip route get 10.20.5.5      # should now resolve via the new route
sudo ip route del 10.20.0.0/16
```

Step 4 — Build a 2-namespace topology

```
# Create two namespaces
sudo ip netns add red
sudo ip netns add blue
```

```
# Connect them with a veth pair
sudo ip link add veth-red type veth peer name veth-blue
sudo ip link set veth-red netns red
sudo ip link set veth-blue netns blue

# Assign IPs and bring up
sudo ip -n red addr add 10.10.0.1/24 dev veth-red
sudo ip -n blue addr add 10.10.0.2/24 dev veth-blue
sudo ip -n red link set veth-red up
sudo ip -n blue link set veth-blue up
sudo ip -n red link set lo up
sudo ip -n blue link set lo up

# Ping across!
sudo ip netns exec red ping -c 3 10.10.0.2
```

Step 5 — Add a policy rule, observe table selection

```
# Create a separate routing table (just memory – no /etc edits needed)
sudo ip route add default via 10.10.0.2 dev veth-red table 100
sudo ip rule add from 10.10.0.0/24 lookup 100

# Verify
sudo ip rule show
sudo ip route show table 100
```

Cleanup

```
sudo ip netns del red
sudo ip netns del blue
sudo ip rule del from 10.10.0.0/24 lookup 100
sudo ip route flush table 100
```

Self-check

- ✓ You can read the routing table and a policy rule list.
- ✓ You have used `ip route get` to ask the kernel.
- ✓ You have built and pinged across a two-namespace topology.

- ✓ You have added an `ip rule` entry that selects a non-main routing table.