

Linux Routing

Presented by
Kishor Panth

What are we going to do ???

- **route**

- route

- route to display the routing cache

- route add

- route add default

- route del

- **ip route**

- ip route show

- ip route show cache

- ip route add to populate a routing table

- default route with ip route add default

- NAT with ip route add nat

- Removeing routes with ip route del

- ting routes with ip route change

- ip route get

- ip route flush

- ip route flush cache

ip rule

ip rule show

Displaying the RPDB with ip rule show

Adding a rule to the RPDB with ip rule add

ip rule add nat

ip rule del

Route -n

```
[root@classroom root]# route -n
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
203.78.162.96	0.0.0.0	255.255.255.248	U	0	0	0	eth0
192.168.100.0	192.168.100.1	255.255.255.0	UG	0	0	0	eth1
192.168.100.0	0.0.0.0	255.255.255.0	U	0	0	0	eth1
172.16.100.0	192.168.100.1	255.255.255.0	UG	0	0	0	eth1
192.168.10.0	192.168.100.1	255.255.255.0	UG	0	0	0	eth1
169.254.0.0	0.0.0.0	255.255.0.0	U	0	0	0	eth1
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	203.78.162.97	0.0.0.0	UG	0	0	0	eth0

```
.
```

route

```
[root@classroom root]# route
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
203.78.162.96	*	255.255.255.248	U	0	0	0	eth0
192.168.100.0	host1.kcm.edu.n	255.255.255.0	UG	0	0	0	eth1
192.168.100.0	*	255.255.255.0	U	0	0	0	eth1
172.16.100.0	host1.kcm.edu.n	255.255.255.0	UG	0	0	0	eth1
192.168.10.0	host1.kcm.edu.n	255.255.255.0	UG	0	0	0	eth1
169.254.0.0	*	255.255.0.0	U	0	0	0	eth1
127.0.0.0	*	255.0.0.0	U	0	0	0	lo
default	leased-cust2-97	0.0.0.0	UG	0	0	0	eth0
-	-	-	-	-	-	-	-

Viewing the routing cache with route

```
[root@classroom root]# route -C
```

```
Kernel IP routing cache
```

Source	Destination	Gateway	Flags	Metric	Ref	Use	Iface
host168.kcm.edu	192.168.100.255	192.168.100.255	b1	0	0	43	lo
host170.kcm.edu	192.168.100.255	192.168.100.255	b1	0	0	6	lo
host166.kcm.edu	192.168.100.255	192.168.100.255	b1	0	0	0	lo
host1.kcm.edu.n	classroom.kcm.e	classroom.kcm.e	1	0	0	0	lo
203.116.51.18	classroom.kcm.e	classroom.kcm.e	1	0	0	2	lo
classroom.kcm.e	www44.thny.bbc.	leased-cust2-97		0	0	0	eth0
classroom.kcm.e	www.my-etrust.c	leased-cust2-97		0	0	2	eth0
classroom.kcm.e	messenger.hotma	leased-cust2-97		0	0	0	eth0
baym-gw35.msgr.	classroom.kcm.e	classroom.kcm.e	1	0	0	24	lo
host130.kcm.edu	192.168.100.255	192.168.100.255	b1	0	0	20	lo
host158.kcm.edu	192.168.100.255	192.168.100.255	b1	0	0	69	lo
host150.kcm.edu	192.168.100.255	192.168.100.255	b1	0	0	38	lo
classroom.kcm.e	www44.thny.bbc.	leased-cust2-97		0	0	0	eth0
classroom.kcm.e	203.116.51.18	leased-cust2-97		0	0	2	eth0
classroom.kcm.e	www.my-etrust.c	leased-cust2-97		0	0	2	eth0
classroom.kcm.e	messenger.hotma	leased-cust2-97		0	0	0	eth0
cisco.intranet.	cisco.intranet.	cisco.intranet.	1	0	0	812	lo

```
■
```

Adding a static entry with route add

```
[root@classroom root]# route add -net 202.79.55.0 netmask 255.255.255.0 gateway 203.78.162.97
SIOCADDRT: File exists
[root@classroom root]# route add - host 202.79.55.14 netmask 255.255.255.0 gw 203.78.162.97
-: Unknown server error
[root@classroom root]#
```

Creating a default route

```
[root@classroom root]# route add default gw 203.78.162.97
```

alternate method

```
[root@classroom root]# route add -net 0.0.0.0 netmask 0.0.0.0 gw 203.78.162.97
```

Removing a static route

```
[~/Quick Connect]om root]# route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
203.78.162.96    0.0.0.0         255.255.255.248 U        0      0      0 eth0
192.168.100.0    192.168.100.1   255.255.255.0   UG        0      0      0 eth1
192.168.100.0    0.0.0.0         255.255.255.0   U        0      0      0 eth1
172.16.100.0     192.168.100.1   255.255.255.0   UG        0      0      0 eth1
202.79.55.0      203.78.162.97   255.255.255.0   UG        0      0      0 eth0
192.168.10.0     192.168.100.1   255.255.255.0   UG        0      0      0 eth1
169.254.0.0      0.0.0.0         255.255.0.0     U        0      0      0 eth1
127.0.0.0        0.0.0.0         255.0.0.0       U        0      0      0 lo
0.0.0.0          203.78.162.97   0.0.0.0         UG        0      0      0 eth0
[root@classroom root]# route del -net 202.79.55.0 netmask 255.255.255.0 gw 203.78.162.97
[root@classroom root]# route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
203.78.162.96    0.0.0.0         255.255.255.248 U        0      0      0 eth0
192.168.100.0    192.168.100.1   255.255.255.0   UG        0      0      0 eth1
192.168.100.0    0.0.0.0         255.255.255.0   U        0      0      0 eth1
172.16.100.0     192.168.100.1   255.255.255.0   UG        0      0      0 eth1
192.168.10.0     192.168.100.1   255.255.255.0   UG        0      0      0 eth1
169.254.0.0      0.0.0.0         255.255.0.0     U        0      0      0 eth1
127.0.0.0        0.0.0.0         255.0.0.0       U        0      0      0 lo
0.0.0.0          203.78.162.97   0.0.0.0         UG        0      0      0 eth0
[root@classroom root]#
```

ip route

- **ip route** provides management tools for manipulating any of the routing tables
- using the **ip route** is that you can operate on any of the 255 routing tables with this command
- the **iproute2** suite of tools does not rely on DNS for any operation so, the ubiquitous -n switch in previous examples will not be required in any example
- **ip route** utility when used in conjunction with the **ip rule** utility can create stateless NAT tables

Operation includes

- displaying routes
- routing cache
- adding routes
- deleting routes
- modifying existing routes
- fetching a route
- modifying existing routes

Displaying a routing table with ip route show

```
You have new mail in /var/spool/mail/root
[root@classroom root]# ip route show
203.78.162.96/29 dev eth0  scope link
192.168.100.0/24 via 192.168.100.1 dev eth1
192.168.100.0/24 dev eth1  scope link
172.16.100.0/24 via 192.168.100.1 dev eth1
192.168.10.0/24 via 192.168.100.1 dev eth1
169.254.0.0/16 dev eth1  scope link
127.0.0.0/8 dev lo  scope link
default via 203.78.162.97 dev eth0
[root@classroom root]#
```

local routing table with ip route show table local

```
[root@classroom root]# ip route show table local
broadcast 192.168.100.0 dev eth1 proto kernel scope link src 192.168.100.100
local 203.78.162.100 dev eth0 proto kernel scope host src 203.78.162.100
broadcast 127.255.255.255 dev lo proto kernel scope link src 127.0.0.1
broadcast 203.78.162.103 dev eth0 proto kernel scope link src 203.78.162.100
broadcast 203.78.162.96 dev eth0 proto kernel scope link src 203.78.162.100
local 192.168.100.100 dev eth1 proto kernel scope host src 192.168.100.100
broadcast 192.168.100.255 dev eth1 proto kernel scope link src 192.168.100.100
broadcast 127.0.0.0 dev lo proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo proto kernel scope host src 127.0.0.1
[root@classroom root]#
```

ip route show table local

- This gives a good deal of information about the IP networks to which the machine is directly connected, and look into the way that the routing tables treat special addresses like broadcast addresses and locally configured addresses.
- The first field in this output tells whether the route is for a broadcast address or an IP address or range locally hosted on this machine. Subsequent fields inform us through which device the destination is reachable, and notably (in this table) that the kernel has added these routes as part of bringing up the IP layer interfaces.
- For each IP hosted on the machine, it makes sense that the machine should restrict accessibility to that IP or IP range to itself only. This explains why, in 192.168.100.100 has a host scope. Because classroom hosts this IP, there's no reason for the packet to be routed off the box. Similarly, a destination of localhost (127.0.0.1) does not need to be forwarded off this machine. In each of these cases, the scope has been set to host.

- For broadcast addresses, which are intended for any listeners who happen to share the IP network, the destination only makes sense as for a scope of devices connected to the same link layer
- The final characteristic available to us in each line of the local routing table output is the src keyword. This is treated as a hint to the kernel about what IP address to select for a source address on outgoing packets on this interface. Naturally, this is most commonly used (and abused) on multi-homed hosts, although almost every machine out there uses this hint for connections to localhost

routing cache with ip route show cache

```
cache mtu 1500 advmss 1460
local 203.78.162.100 from 203.78.160.225 dev lo src 203.78.162.100
cache <local> iif eth0
local 203.78.162.100 from 65.54.194.117 dev lo src 203.78.162.100
cache <local> iif eth0
65.54.246.250 via 203.78.162.97 dev eth0 src 203.78.162.100
cache mtu 1500 advmss 1460
broadcast 192.168.100.255 from 192.168.100.150 dev lo src 192.168.100.100
cache <local,brd> iif eth1
local 203.78.162.100 from 66.150.2.101 dev lo src 203.78.162.100
cache <local> iif eth0
broadcast 192.168.100.255 from 192.168.100.158 dev lo src 192.168.100.100
cache <local,brd> iif eth1
broadcast 192.168.100.255 from 192.168.100.156 dev lo src 192.168.100.100
cache <local,brd> iif eth1
broadcast 192.168.100.255 from 192.168.100.130 dev lo src 192.168.100.100
cache <local,brd> iif eth1
65.54.208.222 from 203.78.162.100 via 203.78.162.97 dev eth0
cache mtu 1500 rtt 296ms rttvar 167ms cwnd 2 advmss 1460
207.46.104.20 via 203.78.162.97 dev eth0 src 203.78.162.100
cache mtu 1500 advmss 1460
local 203.78.162.100 from 207.46.110.33 dev lo src 203.78.162.100
cache <local> iif eth0
broadcast 192.168.100.255 from 192.168.100.178 dev lo src 192.168.100.100
cache <local,brd> iif eth1
local 203.78.162.100 from 66.150.2.71 dev lo src 203.78.162.100
cache <local> iif eth0
216.109.118.73 from 203.78.162.100 via 203.78.162.97 dev eth0
cache mtu 1500 rtt 55ms rttvar 57ms cwnd 2 advmss 1460
65.54.246.250 from 203.78.162.100 via 203.78.162.97 dev eth0
```

route add to populate a routing table

```
[root@classroom root]# ip route add 202.79.55.0/24 via 192.168.100.1  
[root@classroom root]#
```

prohibit route with route add

```
[root@classroom root]# ip route add prohibit 202.79.55.56
```

Using from in a routing command with route add

```
[root@classroom root]# ip route add prohibit 202.79.55.56 from 192.168.100.117
```

default route with ip route add default

```
[root@classroom root]# ip route add default via 203.78.162.97
```

NAT with ip route

```
[root@classroom root]# ip route add nat 0.0.0.0 via 203.78.162.100
RTNETLINK answers: File exists
[root@classroom root]# ip route show table local | grep ^nat
nat 0.0.0.0 via 203.78.162.100 scope host
[root@classroom root]#
```

ip route del

```
[root@classroom root]# ip route show
prohibit 202.79.55.56
prohibit 209.10.26.51
203.78.162.96/29 dev eth0 scope link
192.168.100.0/24 via 192.168.100.1 dev eth1
192.168.100.0/24 dev eth1 scope link
172.16.100.0/24 via 192.168.100.1 dev eth1
202.79.55.0/24 via 192.168.100.1 dev eth1
192.168.10.0/24 via 192.168.100.1 dev eth1
10.38.0.0/16 via 192.168.100.1 dev eth1
169.254.0.0/16 dev eth1 scope link
127.0.0.0/8 dev lo scope link
default via 203.78.162.97 dev eth0
You have new mail in /var/spool/mail/root
[root@classroom root]# ip route del 202.79.55.0/24 via 192.168.100.1
[root@classroom root]# █
```

```
[root@classroom root]# ip route show
prohibit 202.79.55.56
prohibit 209.10.26.51
203.78.162.96/29 dev eth0 scope link
192.168.100.0/24 via 192.168.100.1 dev eth1
192.168.100.0/24 dev eth1 scope link
172.16.100.0/24 via 192.168.100.1 dev eth1
192.168.10.0/24 via 192.168.100.1 dev eth1
10.38.0.0/16 via 192.168.100.1 dev eth1
169.254.0.0/16 dev eth1 scope link
127.0.0.0/8 dev lo scope link
default via 203.78.162.97 dev eth0
[root@classroom root]#
```

Altering existing routes with ip route change

```
[root@classroom root]# ip route show
prohibit 202.79.55.56
prohibit 209.10.26.51
203.78.162.96/29 dev eth0  scope link
192.168.100.0/24 via 192.168.100.1 dev eth1
192.168.100.0/24 dev eth1  scope link
172.16.100.0/24 via 192.168.100.1 dev eth1
192.168.10.0/24 via 192.168.100.1 dev eth1
10.38.0.0/16 via 192.168.100.1 dev eth1
169.254.0.0/16 dev eth1  scope link
127.0.0.0/8 dev lo  scope link
default via 203.78.162.97 dev eth0
[root@classroom root]# ip route change -net 10.38.0.0/16 via 192.168.100.100
Error: an inet prefix is expected rather than "-net".
[root@classroom root]# ip route change 10.38.0.0/16 via 192.168.100.100
[root@classroom root]#
```

fetching route information with ip route get

```
[root@classroom root]# ip route get 0.0.0.0
```

```
local 127.0.0.1 dev lo src 127.0.0.1
```

```
cache <local> mtu 16436 advmss 16396
```

```
[root@classroom root]# ip route get 172.16.100.1
```

```
172.16.100.1 via 192.168.100.1 dev eth1 src 192.168.100.100
```

```
cache mtu 1500 advmss 1460
```

```
[root@classroom root]#
```

Clearing routing tables with ip route flush

```
[root@classroom root]# ip route show
prohibit 209.10.26.51
203.78.162.96/29 dev eth0 scope link
192.168.100.0/24 via 192.168.100.1 dev eth1
192.168.100.0/24 dev eth1 scope link
172.16.100.0/24 via 192.168.100.1 dev eth1
192.168.10.0/24 via 192.168.100.1 dev eth1
10.38.0.0/16 via 192.168.100.100 dev eth1
169.254.0.0/16 dev eth1 scope link
127.0.0.0/8 dev lo scope link
default via 203.78.162.97 dev eth0
[root@classroom root]# ip route flush 209.10.26.51
[root@classroom root]#
```

ip route flush cache

```
[root@classroom root]# ip route show cache | more
64.124.140.148 from 203.78.162.100 via 203.78.162.97 dev eth0
  cache mtu 1500 advmss 1460
66.102.7.99 via 203.78.162.97 dev eth0 src 203.78.162.100
  cache mtu 1500 advmss 1460
65.54.194.117 via 203.78.162.97 dev eth0 src 203.78.162.100
  cache mtu 1500 advmss 1460
203.78.160.36 via 203.78.162.97 dev eth0 src 203.78.162.100
  cache mtu 1500 advmss 1460
local 203.78.162.100 from 65.54.140.158 dev lo src 203.78.162.100
  cache <local> iif eth0
192.168.100.171 from 192.168.100.100 via 192.168.100.1 dev eth1
  cache mtu 1500 rtt 73ms rttvar 72ms cwnd 2 advmss 1460
local 203.78.162.100 from 209.73.164.91 dev lo src 203.78.162.100
  cache <local> iif eth0
203.78.162.102 dev eth0 src 203.78.162.100
  cache mtu 1500 advmss 1460
65.54.140.158 from 203.78.162.100 via 203.78.162.97 dev eth0
  cache mtu 1500 rtt 763ms rttvar 1232ms cwnd 2 advmss 1460
local 203.78.162.100 from 12.111.50.57 dev lo src 203.78.162.100
  cache <local> iif eth0
203.78.160.39 via 203.78.162.97 dev eth0 src 203.78.162.100
  cache mtu 1500 advmss 1460
192.168.100.168 from 192.168.100.100 via 192.168.100.1 dev eth1
  cache mtu 1500 rtt 62ms rttvar 70ms ssthresh 2 cwnd 3 advmss 1460
207.68.177.124 from 203.78.162.100 via 203.78.162.97 dev eth0
  cache mtu 1500 advmss 1460
129.137.2.136 via 203.78.162.97 dev eth0 src 203.78.162.100
  cache mtu 1500 advmss 1460
broadcast 192.168.100.255 from 192.168.100.123 dev lo src 192.168.100.100
  cache <local,brd> iif eth1
local 203.78.162.100 from 131.103.198.127 dev lo src 203.78.162.100
  cache <local> iif eth0
.....
```

```
[root@classroom root]# ip route flush cache
[root@classroom root]# ip route show cache | more
local 203.78.162.100 from 12.111.50.57 dev lo src 203.78.162.100
    cache <local> iif eth0
local 203.78.162.100 from 203.78.160.20 dev lo src 203.78.162.100
    cache <local> iif eth0
local 192.168.100.100 from 192.168.100.128 dev lo src 192.168.100.100
    cache <local> iif eth1
203.78.162.101 from 203.78.162.100 tos 0x10 dev eth0
    cache mtu 1500 advmss 1460
192.168.100.128 from 192.168.100.100 via 192.168.100.1 dev eth1
    cache mtu 1500 advmss 1460
local 203.78.162.100 from 207.230.101.149 dev lo src 203.78.162.100
    cache <local> iif eth0
12.111.50.57 from 203.78.162.100 via 203.78.162.97 dev eth0
    cache mtu 1500 advmss 1460
192.168.100.123 from 192.168.100.100 via 192.168.100.1 dev eth1
    cache mtu 1500 rtt 36ms rttvar 35ms cwnd 2 advmss 1460
local 203.78.162.100 from 203.78.162.101 dev lo src 203.78.162.100
    cache <local,src-direct> iif eth0
local 192.168.100.100 from 192.168.100.173 dev lo src 192.168.100.100
    cache <local> iif eth1
203.78.160.20 from 203.78.162.100 via 203.78.162.97 dev eth0
    cache ipid 0xf675 mtu 1500 advmss 1460
local 192.168.100.100 from 192.168.100.123 dev lo src 192.168.100.100
    cache <local> iif eth1
207.230.101.149 from 203.78.162.100 via 203.78.162.97 dev eth0
    cache mtu 1500 advmss 1460
[root@classroom root]#
```

Ip Rule

- Another part of the **iproute2** software package
- **ip rule** is the single tool for manipulating the routing policy database under linux (RPDB)

ip rule show

```
You have new mail in /var/spool/mail/root  
[root@classroom root]# ip rule show  
0:      from all lookup local  
32766:  from all lookup main  
32767:  from all lookup 253  
[root@classroom root]# █
```

Adding a rule to the RPDB with ip rule add

```
[root@classroom root]# ip route add 202.79.55.0 via 203.78.162.97 table 9
```

```
[root@classroom root]# ip rule add tos 0x08 table 9
```

```
[root@classroom root]# ip route flush cache
```

```
[root@classroom root]# ip rule show
0:      from all lookup local
32765:  from all tos 0x08 lookup 9
32766:  from all lookup main
32767:  from all lookup 253
[root@classroom root]#
```

- the rule we inserted was added to the next available higher priority in the RPDB because we did not specify a priority. If we wished to specify a priority, we could use **prio**.
- any packet with an IP ToS field matching 0x08 will be routed according to the instructions in table 9
- no route in table 9 applies to the matched packet (not possible, since we added a default route), the packet would be routed according to the instructions in table "main"

ip rule add nat

```
[root@classroom root]# ip rule add nat 203.78.162.100 from 192.168.100.0
[root@classroom root]# ip rule show
0:      from all lookup local
32764:  from 192.168.100.0 lookup main map-to 203.78.162.100
32765:  from all tos 0x08 lookup 9
32766:  from all lookup main
32767:  from all lookup 253
[root@classroom root]#
```

ip rule del

```
[root@classroom root]# ip rule add nat 203.78.162.100 from 192.168.100.0
[root@classroom root]# ip rule show
0:      from all lookup local
32764:  from 192.168.100.0 lookup main map-to 203.78.162.100
32765:  from all tos 0x08 lookup 9
32766:  from all lookup main
32767:  from all lookup 253
[root@classroom root]# ip rule del nat 203.78.162.100 from 192.168.100.0
You have new mail in /var/spool/mail/root
[root@classroom root]# ip rule show
0:      from all lookup local
32765:  from all tos 0x08 lookup 9
32766:  from all lookup main
32767:  from all lookup 253
[root@classroom root]# █
```