

Routing

Do not confuse routing tables with iptables. Routing tables specify how to deliver a packet, whereas iptables specify whether to deliver it at all. They are completely different and unrelated.

Routing tables (ip route)

There isn't 'one' routing table in Linux. Instead, there are multiple routing tables — and a set of rules that tell the kernel how to choose the right table for each packet.

What you see when you run `ip route` without specifying a table is the contents of one particular table, `main`. Tables are identified by integer numbers (from 1 to 232–1) but can be also given textual names, which are listed in the file **/etc/iproute2/rt_tables**. The default one will look something like this:

```
cat /etc/iproute2/rt_tables
-----
#
# reserved values
#
255 local
254 main
253 default
0    unspec
#
# local
#
#1  inr.ruhep
```

You can view the contents of any table like this:

```
# ip route list table local
# ip route list table 255
```

Routing policies (ip rule)

So how does the kernel know which routing table to apply? It uses the “routing policy database”, which is managed by the `ip rule` command. In particular, `ip rule` without any arguments will print all existing rules. These are mine:

```
# ip rule
-----
0:  from all lookup local
32764: from all lookup main suppress_prefixlength 0
32765: not from all fwmark 0xca6c lookup 51820
```

```
32766: from all lookup main
32767: from all lookup default
```

The numbers you see on the left (0, 32764, ...) are rule priorities: the lower the number, the higher the priority. Rules with lower numbers are processed first. Apart from the priority, each rule has also a selector and an action. The selector tells us whether the rule applies to the packet at hand. If it does, the action is executed. The most common action is to consult a particular routing table (see the previous section). If that routing table contained a route for our packet, then we're done; otherwise, we proceed to the next rule.

The rules with priorities 0, 32766 and 32767 above are created automatically by the kernel. To quote the ip-rule(8) man page:

- Priority: 0, Selector: match anything, Action: lookup routing table local (ID 255). The local table is a special routing table containing high priority control routes for local and broadcast addresses.
- Priority: 32766, Selector: match anything, Action: lookup routing table main (ID 254). The main table is the normal routing table containing all non-policy routes. This rule may be deleted and/or overridden with other ones by the administrator.
- Priority: 32767, Selector: match anything, Action: lookup routing table default (ID 253). The default table is empty. It is reserved for some post-processing if no previous default rules selected the packet. This rule may also be deleted

Show routing

The following commands can be used to show the routing tables. They have exact the same output:

```
# ip route show
# ip route list
# ip route list table main
```

```
# ip route show
-----
default via 192.168.178.1 dev enp1s0 proto dhcp src 192.168.178.243 metric 100
169.254.0.0/16 dev enp1s0 scope link metric 1000
192.168.178.0/24 dev enp1s0 proto kernel scope link src 192.168.178.243
metric 100
```

Each entry is nothing but an entry in the routing table (Linux kernel routing table). For example, the following line represents the route for the local network. All network packets to a system in the same network are sent directly through the device enp1s0:

```
192.168.178.0/24 dev enp1s0 proto kernel scope link src 192.168.178.243
metric 100
```

Our default route is set via enp1s0 interface i.e. all network packets that cannot be sent according to the previous entries of the routing table are sent through the gateway defined in this entry i.e 192.168.178.1 is our default gateway.

Show network interfaces

Here is how to list all network interfaces on your Linux machine:

```
# ip link show
-----
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
DEFAULT group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: enp1s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state
UP mode DEFAULT group default qlen 1000
    link/ether 30:1a:11:9a:a3:64 brd ff:ff:ff:ff:ff:ff
3: wlp2s0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state
DOWN mode DORMANT group default qlen 1000
    link/ether 8e:57:51:ea:b6:70 brd ff:ff:ff:ff:ff:ff permaddr
30:35:a6:3b:78:0b
```

Add a route

An route can be added by:

```
# ip route add <network>/<netmask> via <gateway> dev <interface>
```

Where ip route add takes the following options:

- **<network>** is the network that you want to add a route to.
- **<netmask>** is the netmask for the network.
- **<gateway>** is the static Linux gateway for the network defined by the **<network>/<netmask>**.
- **<interface>** is the Linux interface that will be used to send packets to the network.

Let us see some examples. Open the terminal app under Linux and type the following command to sent all packets to the local network 192.168.1.0 directly through the device eth0:, enter:

```
# ip route add 192.168.1.0/24 dev eth0
```

In this example route traffic via 192.168.2.254 gateway for 192.168.2.0/24 network:

```
# ip route add 192.168.2.0/24 via 192.168.2.254 dev eth0
```

In other words, the above command will add a static route to the network 192.168.2.0/24. The route will use the gateway 192.168.2.254 and the Linux network interface eth0.

Set default route

All network packets that cannot be sent according to the previous entries of the routing table are sent

through the following default gateway:

```
# ip route add default via 192.168.1.254
```

Delete route

Type the following command to delete route:

```
# ip route delete 192.168.1.0/24 dev eth0
```

Let us delete default route too:

```
# ip route delete default
```

Check Route

You can also use ip command to find the route to the IP address. The following command will show the interface, metric, and gateway that is used to reach the IP address named 1.1.1.1 and 10.83.200.242:

```
# ip route get 10.83.200.242
-----
10.83.200.242 via 192.168.178.1 dev enp1s0 src 192.168.178.243 uid 0
cache
```

This output shows that the interface enp1s0 used to reach the IP address 10.83.200.242, the metric is 1000, and the gateway IP is 192.168.178.1. You can verify this using the following command assuming that netmask is /24 for 10.83.200.0 network:

From:
<https://wiki.oscardegroot.nl/> - **HomeWiki**



Permanent link:
<https://wiki.oscardegroot.nl/doku.php?id=networking:routing>

Last update: **2023/09/16 11:13**