

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

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

Permanent link:
<https://wiki.oscardegroot.nl/doku.php?id=networking:routing&rev=1694856845>

Last update: **2023/09/16 09:34**

