

Policy based routing

By default networks packets are routed based on their destination address.

Linux supports routing policies, and this allows you to have multiple routing tables and make routing decisions based on other variables.



Kernel must be compiled with CONFIG_IP_ADVANCED_ROUTER (Networking/IP: Advanced Router) and CONFIG_IP_MULTIPLE_TABLES (Networking/IP: policy routing).

If you will use this in combination with iptables you will also need CONFIG_NETFILTER_XT_MARK & IP_NF_MANGLE from Network packet filtering framework (Netfilter).

Routing tables

By default you should have 3 routing tables: **local**, **main** & **default**. You can check what tables are currently existing by checking the content of [/etc/iproute2/rt_tables](#).

By default the file should look something like this:

[/etc/iproute2/rt_tables](#)

```
#
# reserved values
#
255    local
254    main
253    default
0     unspec
#
# local
#
#1     inr.ruhep
```

You can check the routing table(s) with

```
ip route show table [Table name || Table ID]
```

Source based routing example

In order to start making routing based on source address you'll need to do just two simple things.

First you'll need to create a new **routing table** instance. Let's name it **srcroute**.

```
echo 200 srcroute >> /etc/iproute2/rt_tables
```

Then you need to create the policy routing rule. In the current case we'll match all traffic with srcip 10.10.10.1 and route it based on the **srcroute** routing table entries.

```
ip rule add from 10.10.10.1 lookup srcroute
```



You will probably need to add a default gateway to the new routing table(s) you are creating.

ip rule

Full list of the possible parameters you can get from the [man page](#).

You can combine ip-rule with fwmark that you can set with iptables, so can create an ip rule with every match iptables is capable of.

If you notice any packets getting lost, make sure to enable martian packages logging

```
net.ipv4.conf.default.log_martians=1  
net.ipv4.conf.all.log_martians=1
```

If you notice any martian packets getting logged/dropped, you can disable the reverse path filter of the kernel

```
net.ipv4.conf.default.rp_filter=0
```

iptables mark

Simply use `-j MARK --set-mark <MARK>`. The target MARK only works in mangle.

- For incoming packages use `-t mangle -A PREROUTING`
- For outgoing packages `-t mangle -A OUTPUT`.

Sources

* Originally written by [lamerix](#)

[howtos](#), [network](#), [routing](#), [author lamerix](#)

From:

<https://docs.slackware.com/> - **SlackDocs**

Permanent link:

https://docs.slackware.com/howtos:misc:network_policy_based_routing

Last update: **2018/03/09 22:24 (UTC)**

