Policy Routing

Version: 3511

Table of Contents

Policy Routing

Overview

Policy Routing is used for advanced control over network traffic.

No additional packages are required for use of Policy Routing

Configuration

Source IP Configuration

To configure Policy Routing, we will use the sample configuration below. In this example, we will have two ISPs called A and B. Each ISP is routing us a block of IP addresses, and these IP addresses need to leave out their respective ISPs:

```
!
interface Ethernet0
description ISP-A
ip address 192.168.5.2 255.255.255.252
!
interface Ethernet1
description ISP-B
ip address 172.16.6.2 255.255.255.252
!
interface Ethernet2
description Internal Network
#ISP-A Netblock
ip address 1.2.3.1 255.255.255.0
#ISP-B Netblock
ip address 10.20.30.1 255.255.255.0
!
```

We will route the 1.2.3.0/24 network through ISP-A and the 10.20.30.0/24 network through ISP-B.

```
!
#Rules for ISP-A
ip rule add from 192.168.5.0/30 table 100
ip rule add from 1.2.3.0/24 table 100
ip route add 192.168.5.0/30 dev eth0 table 100
ip route add 1.2.3.0/24 dev eth2 table 100
ip route add default via 192.168.5.1 table 100
#Rules for ISP-B
ip rule add from 172.16.6.0/30 table 200
ip rule add from 10.20.30.0/24 table 200
ip route add 172.16.6.0/30 dev eth1 table 200
ip route add 10.20.30.0/24 dev eth2 table 200
ip route add default via 172.16.6.1 table 200
#Router's default route for primary routing table
ip route add default via 192.168.5.1
```

```
!
```

In the above configuration, it would not be possible for the two internal networks to communicate. If you want them to be able to route between each other, you'll need to add an interface route to each of the tables. Those rules would look like this:

```
ip route add 10.20.30.0/24 dev eth2 table 100
ip route add 1.2.3.0/24 dev eth2 table 200
```

Interface Based Configuration

With the above example, can also specify policy routing based on incoming interface as well.

```
!
 ip rule add iif eth0 table 100
 ip rule add iif eth1 table 200
!
```

With the above examples, all the traffic coming in the specified interfaces will use the tables listed. We can write rules in these tables to choose where the traffic will be routed to.

Iptables Fwmark Based Configuration

The Iptables configuration allows greater flexibility in directing traffic into a routing table. Any iptables match can be used to mark packets with an iptables fwmark. Once a packet is marked with a fwmark it can be directed to any routing table with a ip rule.

The major disadvantage of this method is that it requires an iptables rule(located under the firewall configuration) and an "ip rule" in the network interface configuration. In the following example use an iptables rule to match all tcp port 80 traffic and then route it using an example routing table.

Iptables configuration

Note, while it's possible use iptables to fwmark packets in the FORWARD and OUTPUT chains, if you wish to use fwmark to select a routing table you will need to mark the packets in the PREROUTING chain.

```
iptables -t mangle -A PREROUTING -p tcp --dport 80 -j MARK --set-mark 0x01
```

Network Interface configuration

The network interface configuration example below

```
!
 # Match fwmark 0x01 and us routing table 100 for these packets
 ip rule add fwmark 0x01 table 100
 # Set a default route via Serial0 in table 100
 ip route add default dev Serial0 table 100
!
 # Normal default route
 ip route add default via 173.35.24.1 dev eth0
!
```

Complete Example: External Web proxy Redirection

Policy routing is commonly used to direct traffic to an external Web Proxy. It has the advantage of not modifying the traffic before it arrives at the proxy server. What follows below is complete example of how to setup this common application. Specific implementation details may change based on your configuration.

In this example we have a Proxy server connected to Ethernet1 @ ip address 34.25.33.10. The Router has a static default gateway connected to Ethernet0 via 12.137.51.1. The customers are located on Ethernet2 in the 12.139.68.0/24 network. The Proxy server is set to use the routers local ip address on Ethernet1 as the default gateway.

Network Interface configuration(wan.conf)

The network interface configuration defines how we route traffic on the router and the ip addresses configured on the interfaces.

```
!
interface Ethernet0
 description Upstream internet connection
 ip address 12.137.51.2 255.255.255.252
!
interface Ethernet1
 description Connection to Proxy server
 ip address 34.25.33.1 255.255.255.0
!
interface Ethernet2
 description Customer interface
 ip address 12.139.68.1 255.255.255.0
!
 # Policy routing rule to match traffic marked by iptables
 ip rule add fwmark 0x01 table 100
 # Add network routes to table 100
 ip route add 34.25.33.0/24 dev eth1 table 100
 ip route add 12.139.68.0/24 dev eth2 table 100
 # Add a default route for traffic in table 100 to the Proxy server
 ip route add default via 34.25.33.10 table 100
!
 # Add the normal static default route
 ip route add default via 12.137.51.1
!
```

This configuration sets the ip addresses on the interfaces, sets a normal default router and creates a routing table 100 with associated routes added to the table.

The key item to note about the Network interface configuration shown above is that we have to add network routes to table 100. The normal kernel routing table adds network routes automatically. But when you define a new routing table these routes are not added. Therefore in the example above we add network routes for 34.25.33.0/24 and 12.139.68.0/24 into table 100 before we add the default route pointing at the proxy server.

Firewall Configuration

This shows the firewall rules needed to implement the redirection to the proxy server. Other iptables rules may effect these rules.

```
iptables -t mangle -A PREROUTING -i eth2 -s 12.139.68.0/24 -p tcp --dport 80 -j MARK --set-mark 0x01
```

This rule sets an FWmark on packets coming in from the Customer facing interface(Ethernet2) that are going to port 80. This FWmark is then used by the "ip rule" in the network interface configuration to put these packets into routing table 100. Since we have added a different default route in table 100 that points to the proxy server, these packets will then be routed to the proxy server.

It is very important that we only match packets coming from customers and not packets coming from the Proxy server itself. This is why the iptables rule above is very specific.