

Module 8 – Policy Based Routing

Objective: Using interesting lab exercises, the student will implement some of the fundamental techniques of Policy Routing, as applied in Cisco IOS.

Prerequisite: Module 1 and the Policy Routing Presentation

REFERENCES

Cisco Systems Documentation CD.

INTRODUCTION

The workshop lab exercises which accompany this module appear after the following discussion and documentation about Policy Routing. Please read this text first, referring to the provided CD Documentation in the event of any questions.

WHAT IS POLICY ROUTING?

Policy routing is a more flexible mechanism for routing packets than destination routing. It is a process whereby the router puts packets through a route map before routing them. The route map determines which packets are routed to which router next. You might enable policy routing if you want certain packets to be routed some way other than the obvious shortest path. Some possible applications for policy routing are to provide equal access, protocol-sensitive routing, source-sensitive routing, routing based on interactive versus batch traffic, or routing based on dedicated links.

To enable policy routing, you must identify which route map to use for policy routing and create the route map. The route map itself specifies the match criteria and the resulting action if all of the match clauses are met. These steps are described in the following three task tables.

To enable policy routing on an interface, indicate which route map the router should use by performing the following task in interface configuration mode. All packets arriving on the specified interface will be subject to policy routing. This command disables fast switching of all packets arriving on this interface.

1. Identify the route map to use for policy routing.

```
ip policy route-map map-tag
```

2. Define a route map to control where packets are output.

Tuesday, February 10, 2004

```
route-map map-tag[permit | deny] [sequence-number]
```

3. Define the criteria by which packets are policy routed. The next step is to define the criteria by which packets are examined to see if they will be policy routed. No match clause in the route map indicates all packets. Perform one or more of the following tasks in route-map configuration mode:

- Match the Level 3 length of the packet:

```
match length min max
```

- Match the destination IP address that is permitted by one or more standard or extended access lists:

```
match ip address {access-list-number | name} [...access-list-number |  
name]
```

4. **Pick Destination Port.** The last step is to specify where the packets that pass the match criteria are output. To do so, perform one or more of the following tasks in route-map configuration mode:

- Specify the next hop to which to route the packet Next-hop must be adjacent, early documentation is NOT correct on this point.

```
set ip next-hop ip-address [...ip-address]
```

- Specify the output interface for the packet.

```
set interface type number [... type number]
```

- Specify the next hop to which to route the packet, if there is no explicit route for this destination.

```
set ip default next-hop ip-address [... ip-address]
```

- Specify the output interface for the packet, if there is no explicit route for this destination.

```
set default interface type number [...type number]
```

The set commands can be used in conjunction with each other. They are evaluated in the order shown in the previous task table. A usable next hop implies an interface. Once the local router finds a next hop and a usable interface, it routes the packet.

Enable Local Policy Routing

Packets that are generated by the router are not normally policy-routed. To enable local policy routing for such packets, indicate which route map the router should use by performing the following task in global configuration mode. All packets originating on the router will then be subject to local policy routing.

Identify the route map to use for local policy routing.

```
ip local policy route-map map-tag
```

Use the *show ip local policy* command to display the route map used for local policy routing, if one exists.

Caveats

- Minimum version is 11.0(1) – policy routed traffic is process switched, so can be a processor intensive operation for high data rates.
- Make sure "ip policy route-map" is applied on the INPUT interface.
- Make sure that IP routing or something similar is used to get the packets back to the source. Policy routing is a **static** route to the next hop – the next hop **must** know how to route the packet onwards.
- Policy routed traffic can be fast-switched (see below) as from 11.3 and 12.0 software releases. Note that *ip route-cache policy* is off by default and needs to be applied to the input interface.
- No IP options work with policy routing, e.g. IP record route ...

Fast-Switched Policy Routing

IP policy routing can now be fast-switched. Prior to this feature, policy routing could only be process switched, which meant that on most platforms, the switching rate was approximately 1,000 to 10,000 packets per second. This was not fast enough for many applications or the size of trunks being used. Users who need policy routing to occur at faster speeds can now implement policy routing without slowing down the router.

Restrictions

Fast-switched policy routing supports all of the **match** commands and most of the **set** commands, except for the following restrictions:

- The **set ip default** command is not supported.
- The **set interface** command is supported only over point-to-point links, unless a route-cache entry exists using the same interface specified in the **set interface** command in the route map. Also, at process level, the routing table is consulted to determine if the interface is on a reasonable path to the destination. During fast switching, the software does not make this check. Instead, if the packet matches, the software blindly forwards the packet to the specified interface.

Tuesday, February 10, 2004

Platforms

This feature is supported on these platforms:

- Cisco 2500 series
- Cisco 2600 series
- Cisco 3600 series
- Cisco 4x00 series
- Cisco 7x00 series

Prerequisites

Policy routing must be configured before you configure fast-switched policy routing. See the section “Enable Policy Routing” in the chapter “Configuring IP Routing Protocols” in the Cisco IOS Release 11.2 *Network Protocols Configuration Guide, Part 1*.

Configuration Task

Fast switching of policy routing is disabled by default. To have policy routing be fast switched, perform the following task in interface configuration mode:

Configuration Example

The following example enables policy routing to be fast switched on an ethernet interface.

```
interface ethernet 0/0
 ip route-cache policy
```

LIST OF RELATED COMMANDS

- ip policy route-map
- ip local policy
- match ip address
- match length
- route-map
- set default interface
- set interface
- set ip default next-hop
- set ip next-hop
- set ip precedence (undocumented as of 11.2(7)P)
- set ip tos (undocumented as of 11.2(7)P)

- ip route-cache policy
- show ip cache policy

POLICY-BASED ROUTING CASE STUDIES

Case Study1: Example from beta training

```
input-network ----- async1 Router e0 ----- 6.6.6.6
                                           7.7.7.7   two neighbours
2 users                               one access-point   towards 2 networks
```

```
int async 1
ip policy route-map equal-access  ! equal-access is an arbitrary name

route-map equal-access permit 30
  set default interface Null0  ! this is the bit bucket (drop unknown guys)
!
route-map equal-access permit 20
  match ip address 2
  set ip default next-hop 7.7.7.7
!
route-map equal-access permit 10
  match ip address 1
  set ip default next-hop 6.6.6.6
!
! ..... the evaluation sequence number is 10,20,30
!

access-list 1 permit 123.123.123.123
access-list 2 permit 124.124.124.124
! ..... differentiate your 2 allowed "users"
!       this example shows simple access-lists but extended IP access lists
!       are supported too.
```

Case Study2: telnet and ftp traffic separation

```

      |               +-----+ S1               +-----+
      |               |               |               | e0
+-----+ left  |               | right  +-----+
      |               +-----+
193.104.93.0    +-----+ BRI0               +-----+

```

Tuesday, February 10, 2004

The config on left is:

```
username right password blah

interface Ethernet0
 ip address 193.104.93.10 255.255.255.0
 ip policy route-map foo
!
interface Serial1
 ip address 22.0.0.2 255.0.0.0
 encapsulation x25
 x25 address 2
 x25 idle 1
 x25 map ip 22.0.0.1 1

interface BRI0
 ip address 10.0.0.2 255.0.0.0
 encapsulation ppp
 dialer idle-timeout 60
 dialer map ip 10.0.0.1 name right 12345678
 dialer-group 1
 ppp compression predictor
 ppp authentication chap

access-list 101 permit tcp 193.104.93.0 0.0.0.255 any eq telnet
access-list 101 permit icmp any any
access-list 103 permit tcp 193.104.93.0 0.0.0.255 any eq ftp
access-list 103 permit udp 193.104.93.0 0.0.0.255 any eq tftp
!
route-map foo permit 12
 set default interface Null0
!
route-map foo permit 11
 match ip address 103
 set ip next-hop 10.0.0.1
!
route-map foo permit 10
 match ip address 101
 set ip next-hop 22.0.0.1
```

Case Study 3 - IP Precedence and TOS bits.

First define the flows you want to change the precedence,
in my case are all the flows destined to 224.1.1.1:

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
access-list 101 permit ip any host 224.1.1.1
```

Use the route map to change the precedence:

```
route-map precedence permit 1
match ip address 101
set ip precedence priority
```

or the TOS:

```
route-map tos permit 1
match ip address 101
set ip tos max-reliability
```

Now you want to apply the route map to the (incoming) interface(s) where you want the precedence change to take effect:

```
interface Ethernet1/0
ip policy route-map precedence
```

Case Study 4 - Equal Access

The following example provides two sources with equal access to two different service providers. Packets arriving on async interface 1 from the source 1.1.1.1 are sent to the router at 6.6.6.6 if the router has no explicit route for the packet's destination. Packets arriving from the source 2.2.2.2 are sent to the router at 7.7.7.7 if the router has no explicit route for the packet's destination. All other packets for which the router has no explicit route to the destination are discarded.

```
access-list 1 permit ip 1.1.1.1
access-list 2 permit ip 2.2.2.2
!
interface async 1
ip policy route-map equal-access
!
route-map equal-access permit 10
match ip address 1
set ip default next-hop 6.6.6.6
route-map equal-access permit 20
match ip address 2
set ip default next-hop 7.7.7.7
route-map equal-access permit 30
set default interface null0
```

LABS & EXERCISES

1. **Check Physical Connectivity.** The connectivity for this workshop should be as in Figure 1. Ensure that all physical connections are complete and pingable. The addresses used for links between routers should be left the same as those chosen for Module 1.




Figure 1 – Network Configuration

2. **Clean up from past labs.** Delete any remaining BGP configuration, route maps, communities, access-lists and AS-Path access-lists.
3. **Create OSPF PID 100, area 0 on all routers.** All Router Teams will configure one OSPF domain – Area 0 – for the network. This configuration should be as was used in Module 1 checkpoint #1.
4. **Ping Test.** The system specified by the instructors (usually Router15 or the time server ntp0.workshop.net) will be used for the target of this lab. All routers should be able to ping and traceroute to the Time Server. Capture the trace to the Time Server. This will be used as a baseline.

Checkpoint #1: Call the workshop instructor, explain the OSPF configuration, and demonstrate the connectivity to the Time Server.

5. **Policy Assignment – Create a Ping/Traceroute Ring.** Using policy routing, configure interfaces on the entire network so that a traceroute from **Router6** to the Time Server would take the following path (see Figure 2):

Router6 ? Router8 ? Router10 ? Router12 ? Router14 ? Router13 ? Router11 ? Router9 ? Router7 ? Router5
 ? Router3 ? Router1 ? Router2 ? Router4 ? Router6 ? Router15/TimeSource

6. **Approaches to applying policy routing.** Policy-Based Routing is as dangerous as static routes. Hence, you need to approach the application of a Policy Route very methodically. OSPF must be configured properly before any of consideration is made to implement policy routing. The first step is to change the characteristic of the trace leaving the router. If the trace to the Time Server originally went out one port, change it to go out another. To accomplish this you use the *ip local policy route-map* command (see explanation above).

Hint: Work out which interface is the outbound interface and which is inbound, and where the policy configuration needs to be applied. Think before typing in configuration.

Example:

```
! destination address to match against
!
access-list 100 permit ip any host 192.168.1.1
!
! configuration to determine policy for packets generated on router
!
route-map redirect-local-traffic permit 10
  match ip address 100
  set interface serial 0/0
!
ip local policy route-map redirect-local-traffic
```

Tuesday, February 10, 2004

```
!  
! configuration to determine policy for incoming packets  
!  
route-map redirect-incoming-traffic permit 10  
  match ip address 100  
  set interface serial 0/0  
!  
interface serial 1/0  
  ip policy route-map redirect-incoming-traffic
```

HINT: For two routers that interconnect via serial links – the *set interface* command may be easier. *Set interface* will not work for broadcast interconnects like ethernet.

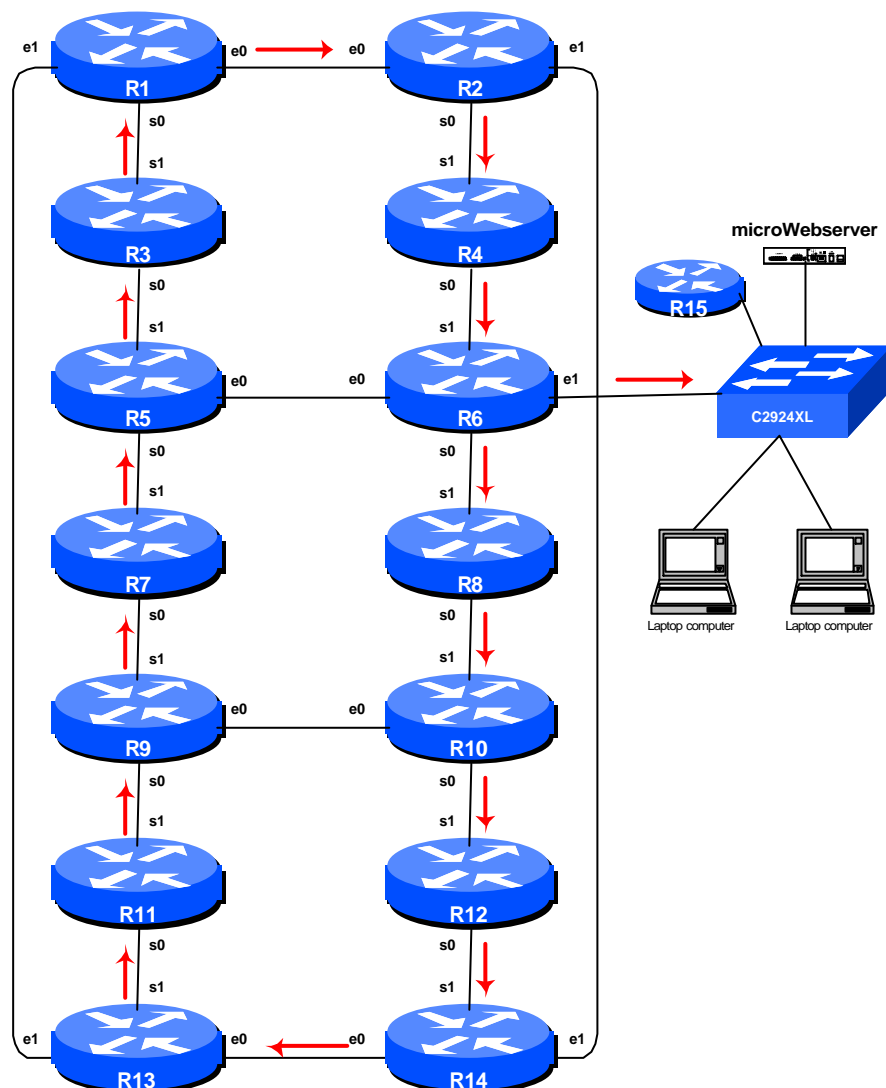


Figure 2 – Traceroute Ring from Router 6 to Time Server (Router15)

Question: What should be used in lieu of *set interface* on broadcast mediums like ethernet?

Answer: Use “*set ip next-hop <ip address>*”.

Question: What differences are there between the configuration used for Router 6 and the remaining routers in the workshop? Why?

Answer: Remember that the trace from the console has to go round the class, not direct to the Time Server...

Q. Will Policy Routing drop my packets if I set in on a router's interface, but have no *set next-hop* configured?

A. Policy Routing will first look for another *set* command first. If there is no other *set* command configured or if the *set* command doesn't make sense, then the router falls back to the normal routing (destination-based).

Checkpoint #2: The lab instructor will now demonstrate a trace from Router6 to the Time Server (or substitute device). If the policy routing has been set up successfully, a trace will follow the paths shown in Figure 2.

Q. What happens if the connection between one of the routers in the path of the policy routed connection to the Time Server is broken? Why?

Checkpoint #3: The lab instructor will now break the connection between Router4 and Router6. Observe what happens to the paths followed now.

A. If the policy destination is down, the router reverts to using the routing table for its forwarding decisions.

Tuesday, February 10, 2004

CONFIGURATION NOTES

Documentation is critical! You should record the configuration at each ***Checkpoint***, as well as the configuration at the end of the module.