Configure EIGRP to Influence Path Selection

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Background Information

Scenarios

- Network Diagram
- Initial Configurations
- Scenario 1: Influence Path Selection by Modifying the Delay Metric
- Scenario 2: Influence Path Selection with the Use an Offset-List
- Scenario 3: Influence Path Selection with Summarization
- Scenario 4: Influence Path Selection with the Use of Leak Maps
- Scenario 5: Influence Path Selection by Modifying the Administrative Distance (AD) of a Prefix
- Scenario 6: Influence Path Selection with Route Filtering

Related Information

Introduction

This document describes the process of creating a preferred path by influencing different Enhanced Interior Gateway Routing Protocol (EIGRP) features.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Knowledge of Basic IP Routing
- Knowledge of EIGRP Protocol
- Knowledge of Cisco IOS® IOS Command Line Interface (CLI)

Components Used

This document is not restricted to specific software and hardware versions, however, the information in this

document is based on these software and hardware versions:

- Router ASR 1000
- Router ISR 4000
- Cisco IOS 17.9.x

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

The EIGRP path selection can be influenced by manipulating various metrics that the protocol uses to determine the best path to a destination. EIGRP calculates the best path to a destination based on different metrics, and the path selection process involves evaluating these metrics to determine the optimal route. The EIGRP metrics include bandwidth, delay, load, reliability, and Maximum Transmission Unit (MTU).

Understanding these metrics and their significance helps network administrators to modify EIGRP path selection based on specific requirements or network conditions. By default, from the different metric values, EIGRP only uses the minimum bandwidth on the path to a destination network and the total delay to compute routing metrics. Furthermore, the bandwidth and delay metrics are determined from static values configured on the interfaces from devices along the path towards the destination, in other words these two parameters are not measured dynamically.

Aside from the metric manipulation, route filtering can also be used to influence the path selection in EIGRP. Route filtering involves controlling the information that is allowed or denied to enter or exit a router routing table. Filtering routes can be done for various reasons, including optimizing routing tables, or managing network traffic.

Some of the key features related to route filtering in EIGRP include distribute lists, prefix lists, route maps and leak maps. These mechanisms offer a powerful and flexible way for controlling routing information that can be used by network administrators to tailor EIGRP routing tables to meet specific criteria and improve network efficiency.

Scenarios

In the dynamic landscape of routing protocols, administrators often find themselves faced with the need to tailor routing decisions to align with specific network requirements and optimize traffic flow. This involves leveraging various techniques and configurations to influence how routers make path selection decisions.

These examples provide different alternatives where administrators can employ strategic configurations to manipulate EIGRP path selection:

1. Influence path selection by modifying the delay metric. Adjusting the delay metric on a router interface allows administrators to influence routing decisions by affecting this particular parameter on a link. This subtle manipulation can guide traffic to take preferred paths based on the altered delay values.

2. Influence path selection with the use an offset-list. Employing an offset-list enables the selective modification of metrics for specific prefixes, providing a targeted approach to influencing path selection over a particular interface. This mechanism is used for increasing incoming and outgoing metrics to routes learned via EIGRP and to selectively prefer some prefixes over a particular path.

3. Influence path selection with summarization. Introducing summary routes allows administrators to influence the longest match preference for a prefix. Route summarization can impact the granularity of routing decisions, optimizing routing tables and improving overall network efficiency.

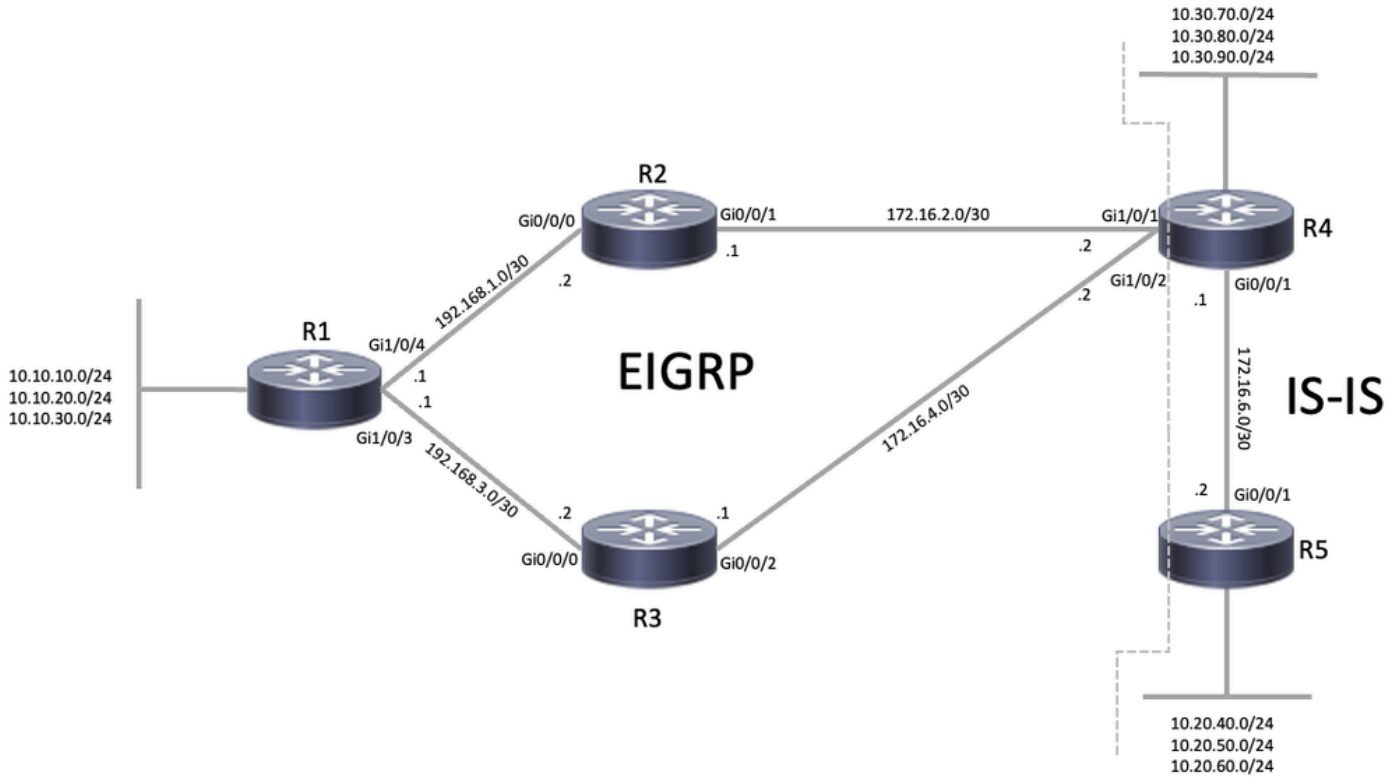
4. Influence path selection with the use of leak maps. Leveraging leak-maps during the advertisement of summary routes provides a mechanism for announcing more specific routes selectively. This approach ensures that summarized information is announced strategically, maintaining routing flexibility and influencing path selection.

5. Influence path selection by modifying the Administrative Distance (AD) of a prefix . Altering the administrative distance of a prefix is a helpful technique to control the source of routing information. This can be particularly useful in scenarios where routes from certain sources need to be excluded from the Routing Information Base (RIB).

6. Influence path selection with route filtering . Route filtering is a powerful method used to control the advertisement or acceptance of specific routes in or out of a routing protocol. It is commonly used to filter routing information based on specified criteria, preventing certain routes from being advertised or learned.

- A distribute-list is one of the main tools used to filter prefixes in EIGRP and it can work in conjunction with an access-list (ACL), prefix-list or route-map.
- Employing a prefix-list facilitates the granular filtering of prefixes from specific neighbors. This level of control is essential to manage routing updates to modify the path preference.

Network Diagram



EIGRP Topology

Initial Configurations

Before modifying any configuration, it is important to review the initial configuration and status of the devices (initial configuration is the same in each scenario). Based on the network diagram, R1, R2, R3 and R4 are EIGRP neighbors (each router has two adjacencies) with R4 also being part of the Intermediate System-to-Intermediate System (IS-IS) domain, and doing mutual redistribution between IS-IS and EIGRP. It is important to note that R1 has two paths in the routing table (through interface Gi1/0/3 and Gi1/0/4) to subnets 10.20.x.x and 10.30.x.x via EIGRP, and that subnets 10.10.x.x are directly connected.

Configurations	
<#root> R1# **show run | section router eigrp** router eigrp LAB ! address-family ipv4 unicast autonomous-system 100 ! topology base exit-af-topology network 10.10.10.0 0.0.0.255 network 10.10.20.0 0.0.0.255 network 10.10.30.0 0.0.0.255 network 192.168.1.0 0.0.0.3	<#root> R1# **show ip route eigrp** Codes: L - local, C - connected, S - static, R - RIP, D - EIGRP, EX - EIGRP external, O - OSPF, IA - N1 - OSPF NSSA external type 1, N2 - OSPF NSSA E1 - OSPF external type 1, E2 - OSPF external n - NAT, Ni - NAT inside, No - NAT outside, No i - IS-IS, su - IS-IS summary, L1 - IS-IS leve ia - IS-IS inter area, * - candidate default, H - NHRP, G - NHRP registered, g - NHRP regist o - ODR, P - periodic downloaded static route,

```
  network 192.168.3.0 0.0.0.3
 exit-address-family

R1#

show run interface GigabitEthernet1/0/3

Building configuration...

Current configuration : 93 bytes
!
interface GigabitEthernet1/0/3
 no switchport
 ip address 192.168.3.1 255.255.255.252
end

R1#

show run interface GigabitEthernet1/0/4

Building configuration...

Current configuration : 93 bytes
!
interface GigabitEthernet1/0/4
 no switchport
 ip address 192.168.1.1 255.255.255.252
end
```

```
       a - application route
       + - replicated route, % - next hop override, p
       & - replicated local route overrides by connec

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2
D EX     10.20.40.0/24
           [170/66560] via 192.168.3.2, 00:31:39, Gig
           [170/66560] via 192.168.1.2, 00:31:39, Gig
D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 00:31:39, Gig
           [170/66560] via 192.168.1.2, 00:31:39, Gig
D EX     10.20.60.0/24
           [170/66560] via 192.168.3.2, 00:31:39, Gig
           [170/66560] via 192.168.1.2, 00:31:39, Gig
D        10.30.70.0/24
           [90/16000] via 192.168.3.2, 00:29:39, Giga
           [90/16000] via 192.168.1.2, 00:29:39, Giga
D        10.30.80.0/24
           [90/16000] via 192.168.3.2, 00:29:39, Giga
           [90/16000] via 192.168.1.2, 00:29:39, Giga
D        10.30.90.0/24
           [90/16000] via 192.168.3.2, 00:29:38, Giga
           [90/16000] via 192.168.1.2, 00:29:38, Giga
      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 6d21h
D        172.16.4.0 [90/15360] via 192.168.3.2, 6d21h

R1#

show ip route connected

 Codes: L - local, C - connected, S - static, R - RIP

10.10.10.0/24 is directly connected, Loopback10

 L 10.10.10.10/32 is directly connected, Loopback10 C

10.10.20.0/24 is directly connected, Loopback20

 L 10.10.20.20/32 is directly connected, Loopback20 C

10.10.30.0/24 is directly connected, Loopback30

 L 10.10.30.30/32 is directly connected, Loopback30 1

R1#

show interfaces GigabitEthernet1/0/3

 GigabitEthernet1/0/3 is up, line protocol is up (con

MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec, rel

 Encapsulation ARPA, loopback not set Keepalive set (

show interfaces GigabitEthernet1/0/4

 GigabitEthernet1/0/4 is up, line protocol is up (con

MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec, rel

 Encapsulation ARPA, loopback not set Keepalive set (

show ip eigrp neighbors
```

	EIGRP-IPv4 VR(LAB) Address-Family Neighbors for AS(1

In the case of R2 and R3, all the prefixes 10.10.x.x, 10.20.x.x and 10.30.x.x are being learned through EIGRP.

R2	
Configurations	**Status**
<#root> R2# **show run | section router eigrp** router eigrp LAB ! address-family ipv4 unicast autonomous-system 100 ! topology base exit-af-topology network 172.16.2.0 0.0.0.3 network 192.168.1.0 0.0.0.3 exit-address-family R2# **show run interface GigabitEthernet 0/0/0** Building configuration... Current configuration : 96 bytes ! interface GigabitEthernet0/0/0 ip address 192.168.1.2 255.255.255.252 negotiation auto end R2# **show run interface GigabitEthernet 0/0/1** Building configuration... Current configuration : 95 bytes ! interface GigabitEthernet0/0/1 ip address 172.16.2.1 255.255.255.252 negotiation auto end	<#root> R2# **show ip route eigrp** Codes: L - local, C - connected, S - static, R - RIP, D - EIGRP, EX - EIGRP external, O - OSPF, IA - N1 - OSPF NSSA external type 1, N2 - OSPF NSSA E1 - OSPF external type 1, E2 - OSPF external i - IS-IS, su - IS-IS summary, L1 - IS-IS leve ia - IS-IS inter area, * - candidate default, o - ODR, P - periodic downloaded static route, a - application route + - replicated route, % - next hop override, p Gateway of last resort is not set 10.0.0.0/24 is subnetted, 9 subnets D 10.10.10.0 [90/10880] via 192.168.1.1, 6d22h D 10.10.20.0 [90/10880] via 192.168.1.1, 6d22h D 10.10.30.0 [90/10880] via 192.168.1.1, 6d22h D EX 10.20.40.0 [170/61440] via 172.16.2.2, 01:32 D EX 10.20.50.0 [170/61440] via 172.16.2.2, 01:32 D EX 10.20.60.0 [170/61440] via 172.16.2.2, 01:32 D 10.30.70.0 [90/10880] via 172.16.2.2, 01:30: D 10.30.80.0 [90/10880] via 172.16.2.2, 01:30: D 10.30.90.0 [90/10880] via 172.16.2.2, 01:30: 172.16.0.0/16 is variably subnetted, 3 subnets, D 172.16.4.0/30 [90/15360] via 172.16.2.2, 6d2 192.168.3.0/30 is subnetted, 1 subnets D 192.168.3.0 [90/15360] via 192.168.1.1, 6d22 R2# **show interfaces GigabitEthernet0/0/0** GigabitEthernet0/0/0 is up, line protocol is up Hardware is BUILT-IN-2T+6X1GE, address is 0062.ec8a Internet address is 192.168.1.2/30 **MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec, rel** Encapsulation ARPA, loopback not set Keepalive not supported Full Duplex, 1000Mbps, link type is auto, media typ output flow-control is on, input flow-control is on ARP type: ARPA, ARP Timeout 04:00:00 Last input 00:00:01, output 00:03:30, output hang n Last clearing of "show interface" counters never Input queue: 0/375/0/0 (size/max/drops/flushes); To

```
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     208297 packets input, 18918243 bytes, 0 no buffe
     Received 718 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ign
     0 watchdog, 145070 multicast, 0 pause input
     134239 packets output, 10474478 bytes, 0 underru
     0 output errors, 0 collisions, 4 interface reset
     11577 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapp

R2#

show interfaces GigabitEthernet0/0/1

GigabitEthernet0/0/1 is up, line protocol is up
  Hardware is BUILT-IN-2T+6X1GE, address is 0062.ec8a
  Internet address is 172.16.2.1/30


 MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec, re

  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 1000Mbps, link type is auto, media typ
  output flow-control is on, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:05, output 00:03:35, output hang n
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); To
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     145790 packets input, 15086179 bytes, 0 no buffe
     Received 2 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     1 input errors, 0 CRC, 0 frame, 0 overrun, 0 ign
     0 watchdog, 145679 multicast, 0 pause input
     134227 packets output, 10473816 bytes, 0 underru
     0 output errors, 0 collisions, 4 interface reset
     11575 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapp

R2#

show ip eigrp neighbors


EIGRP-IPv4 VR(LAB) Address-Family Neighbors for AS(10
H   Address                 Interface              Ho
                                                   (s
1   172.16.2.2              Gi0/0/1
0   192.168.1.1             Gi0/0/0
```

R3

Configurations	Status
<#root> R3# **show run | section router eigrp** router eigrp LAB ! address-family ipv4 unicast autonomous-system 100 ! topology base exit-af-topology network 172.16.4.0 0.0.0.3 network 192.168.3.0 0.0.0.3 exit-address-family R3# **show run interface GigabitEthernet 0/0/0** Building configuration... Current configuration : 96 bytes ! interface GigabitEthernet0/0/0 ip address 192.168.3.2 255.255.255.252 negotiation auto end R3# **show run interface GigabitEthernet 0/0/2** Building configuration... Current configuration : 95 bytes ! interface GigabitEthernet0/0/2 ip address 172.16.4.1 255.255.255.252 negotiation auto end	<#root> R3# **show ip route eigrp** Codes: L - local, C - connected, S - static, R - RIP, D - EIGRP, EX - EIGRP external, O - OSPF, IA - N1 - OSPF NSSA external type 1, N2 - OSPF NSSA E1 - OSPF external type 1, E2 - OSPF external i - IS-IS, su - IS-IS summary, L1 - IS-IS leve ia - IS-IS inter area, * - candidate default, o - ODR, P - periodic downloaded static route, a - application route + - replicated route, % - next hop override, p Gateway of last resort is not set 10.0.0.0/24 is subnetted, 9 subnets D 10.10.10.0 [90/10880] via 192.168.3.1, 6d22h D 10.10.20.0 [90/10880] via 192.168.3.1, 6d22h D 10.10.30.0 [90/10880] via 192.168.3.1, 6d22h D EX 10.20.40.0 [170/61440] via 172.16.4.2, 01:46 D EX 10.20.50.0 [170/61440] via 172.16.4.2, 01:46 D EX 10.20.60.0 [170/61440] via 172.16.4.2, 01:46 D 10.30.70.0 [90/10880] via 172.16.4.2, 01:44: D 10.30.80.0 [90/10880] via 172.16.4.2, 01:44: D 10.30.90.0 [90/10880] via 172.16.4.2, 01:44: 172.16.0.0/16 is variably subnetted, 3 subnets, D 172.16.2.0/30 [90/15360] via 172.16.4.2, 6d2 192.168.1.0/30 is subnetted, 1 subnets D 192.168.1.0 [90/15360] via 192.168.3.1, 6d22 R3# **show interfaces GigabitEthernet0/0/0** GigabitEthernet0/0/0 is up, line protocol is up Hardware is BUILT-IN-2T+6X1GE, address is 0062.ec8a Internet address is 192.168.3.2/30 **MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec, rel** Encapsulation ARPA, loopback not set Keepalive not supported Full Duplex, 1000Mbps, link type is auto, media typ output flow-control is on, input flow-control is on ARP type: ARPA, ARP Timeout 04:00:00 Last input 00:00:01, output 00:00:01, output hang n Last clearing of "show interface" counters never Input queue: 0/375/0/0 (size/max/drops/flushes); To Queueing strategy: fifo Output queue: 0/40 (size/max) 5 minute input rate 0 bits/sec, 0 packets/sec 5 minute output rate 0 bits/sec, 0 packets/sec 208616 packets input, 18949840 bytes, 0 no buffe Received 726 broadcasts (0 IP multicasts) 0 runts, 0 giants, 0 throttles 2 input errors, 0 CRC, 0 frame, 0 overrun, 0 ign 0 watchdog, 145285 multicast, 0 pause input 134420 packets output, 10488621 bytes, 0 underru

```
     0 output errors, 0 collisions, 5 interface reset
     11597 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapp
     10 carrier transitions

R3#

show interfaces GigabitEthernet0/0/2

GigabitEthernet0/0/2 is up, line protocol is up
  Hardware is BUILT-IN-2T+6X1GE, address is 0062.ec8a
  Internet address is 172.16.4.1/30


 MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec, re

  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 1000Mbps, link type is auto, media typ
  output flow-control is on, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:01, output 00:00:01, output hang n
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); To
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     145895 packets input, 15083732 bytes, 0 no buffe
     Received 1 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     1 input errors, 0 CRC, 0 frame, 0 overrun, 0 ign
     0 watchdog, 145785 multicast, 0 pause input
     134433 packets output, 10489999 bytes, 0 underru
     0 output errors, 0 collisions, 5 interface reset
     11543 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapp
     6 carrier transitions

R3#

show ip eigrp neighbors


EIGRP-IPv4 VR(LAB) Address-Family Neighbors for AS(10
H   Address                 Interface              Ho
                                                   (s
1   172.16.4.2              Gi0/0/2
0   192.168.3.1             Gi0/0/0
```

R4

Configurations	Sta
<#root> R4# **show run | section router eigrp**	<#root> R4# **show ip route eigrp**

```
router eigrp LAB
 !
 address-family ipv4 unicast autonomous-system 100
  !
  topology base


redistribute isis level-2 metric 1000000 10 255 1 1500

  exit-af-topology
  network 10.30.70.0 0.0.0.255
  network 10.30.80.0 0.0.0.255
  network 10.30.90.0 0.0.0.255
  network 172.16.2.0 0.0.0.3
  network 172.16.4.0 0.0.0.3
 exit-address-family

R4#
show run | section ^router isis

router isis
 net 49.0001.0000.0000.0004.00
 is-type level-2-only
 metric-style wide


redistribute eigrp 100


R4#
show run interface GigabitEthernet1/0/1

Building configuration...

Current configuration : 95 bytes
!
interface GigabitEthernet1/0/1
 ip address 172.16.2.2 255.255.255.252
 negotiation auto
end

R4#
show run interface GigabitEthernet1/0/2

Building configuration...

Current configuration : 95 bytes
!
interface GigabitEthernet1/0/2
 ip address 172.16.4.2 255.255.255.252
 negotiation auto
end

R4#
show run interface GigabitEthernet0/0/1

Building configuration...

Current configuration : 112 bytes
!
interface GigabitEthernet0/0/1
```

```
Codes: L - local, C - connected, S - static, R -
       D - EIGRP, EX - EIGRP external, O - OSPF,
       N1 - OSPF NSSA external type 1, N2 - OSPF
       E1 - OSPF external type 1, E2 - OSPF exter
       i - IS-IS, su - IS-IS summary, L1 - IS-IS
       ia - IS-IS inter area, * - candidate defau
       o - ODR, P - periodic downloaded static ro
       a - application route
       + - replicated route, % - next hop overrid

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnet
D        10.10.10.0/24 [90/16000] via 172.16.4.1,
                       [90/16000] via 172.16.2.1,
D        10.10.20.0/24 [90/16000] via 172.16.4.1,
                       [90/16000] via 172.16.2.1,
D        10.10.30.0/24 [90/16000] via 172.16.4.1,
                       [90/16000] via 172.16.2.1,
      192.168.1.0/30 is subnetted, 1 subnets
D        192.168.1.0 [90/15360] via 172.16.2.1, 6
      192.168.3.0/30 is subnetted, 1 subnets
D        192.168.3.0 [90/15360] via 172.16.4.1, 6

R4#
show ip route isis


Codes: L - local, C - connected, S - static, R -
       D - EIGRP, EX - EIGRP external, O - OSPF,
       N1 - OSPF NSSA external type 1, N2 - OSPF
       E1 - OSPF external type 1, E2 - OSPF exter
       i - IS-IS, su - IS-IS summary, L1 - IS-IS
       ia - IS-IS inter area, * - candidate defau
       o - ODR, P - periodic downloaded static ro
       a - application route
       + - replicated route, % - next hop overrid

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnet
i L2     10.20.40.0/24 [115/20] via 172.16.6.2, C
i L2     10.20.50.0/24 [115/20] via 172.16.6.2, C
i L2     10.20.60.0/24 [115/20] via 172.16.6.2, C

R4#
show ip route connected

Codes: L - local, C - connected, S - static, R -
       D - EIGRP, EX - EIGRP external, O - OSPF,
       N1 - OSPF NSSA external type 1, N2 - OSPF
       E1 - OSPF external type 1, E2 - OSPF exter
       i - IS-IS, su - IS-IS summary, L1 - IS-IS
       ia - IS-IS inter area, * - candidate defau
       o - ODR, P - periodic downloaded static ro
       a - application route
       + - replicated route, % - next hop overrid

Gateway of last resort is not set
```

```
 ip address 172.16.6.1 255.255.255.252
 ip router isis
 negotiation auto
end
```

```
      10.0.0.0/8 is variably subnetted, 12 subnet
C        10.30.70.0/24 is directly connected, Loc
L        10.30.70.70/32 is directly connected, Lo
C        10.30.80.0/24 is directly connected, Loc
L        10.30.80.80/32 is directly connected, Lo
C        10.30.90.0/24 is directly connected, Loc
L        10.30.90.90/32 is directly connected, Lo
      172.16.0.0/16 is variably subnetted, 6 subr
C        172.16.2.0/30 is directly connected, Gig
L        172.16.2.2/32 is directly connected, Gig
C        172.16.4.0/30 is directly connected, Gig
L        172.16.4.2/32 is directly connected, Gig
C        172.16.6.0/30 is directly connected, Gig
L        172.16.6.1/32 is directly connected, Gig

R4#

show interfaces GigabitEthernet1/0/1

GigabitEthernet1/0/1 is up, line protocol is up
  Hardware is SM-X-4X1G-1X10G, address is 0027.90
  Internet address is 172.16.2.2/30


 MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec

  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 1000Mbps, link type is auto, media
  output flow-control is on, input flow-control i
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:05:38, output 00:00:30, output ha
  Last clearing of "show interface" counters neve
  Input queue: 0/375/0/0 (size/max/drops/flushes)
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     134612 packets input, 9965393 bytes, 0 no bu
     Received 5 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0
     0 watchdog, 134482 multicast, 0 pause input
     146207 packets output, 14544461 bytes, 0 und
     0 output errors, 0 collisions, 1 interface r
     0 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers s

R4#

show interfaces GigabitEthernet1/0/2

GigabitEthernet1/0/2 is up, line protocol is up
  Hardware is SM-X-4X1G-1X10G, address is 0027.90
  Internet address is 172.16.4.2/30


MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,

  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 1000Mbps, link type is auto, media
```

```
  output flow-control is on, input flow-control i
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:08:36, output 00:00:01, output ha
  Last clearing of "show interface" counters neve
  Input queue: 0/375/0/0 (size/max/drops/flushes)
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     134654 packets input, 9968624 bytes, 0 no bu
     Received 2 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0
     0 watchdog, 134535 multicast, 0 pause input
     146139 packets output, 14525699 bytes, 0 und
     0 output errors, 0 collisions, 1 interface r
     0 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers s

R4#

show interfaces GigabitEthernet0/0/1

GigabitEthernet0/0/1 is up, line protocol is up
  Hardware is ISR4331-3x1GE, address is 0027.9064
  Internet address is 172.16.6.1/30
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 use
     reliability 255/255, txload 1/255, rxload 1/
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 1000Mbps, link type is auto, media
  output flow-control is on, input flow-control i
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:01, output 00:00:03, output ha
  Last clearing of "show interface" counters neve
  Input queue: 0/375/0/0 (size/max/drops/flushes)
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     576123 packets input, 655123623 bytes, 0 no
     Received 2 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0
     0 watchdog, 576069 multicast, 0 pause input
     154335 packets output, 216885838 bytes, 0 un
     0 output errors, 0 collisions, 1 interface r
     0 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers s

R4#

show ip eigrp neighbors


EIGRP-IPv4 VR(LAB) Address-Family Neighbors for A
H   Address                 Interface

1   172.16.4.1              Gi1/0/2
```

```
0   172.16.2.1              Gi1/0/1

R4#

show isis neighbors



System Id       Type Interface     IP Address
R5              L2   Gi0/0/1       172.16.6.2
```

Scenario 1: Influence Path Selection by Modifying the Delay Metric

In this example, the Delay value is used to influence EIGRP to prefer the path through R3. Before doing any change, you can confirm that EIGRP is load balancing between interfaces Gi1/0/3 and Gi1/04 as both interfaces have the same Delay value of 10 microseconds.

<#root>

R1#

show ip route eigrp

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX     10.20.40.0/24
           [170/66560] via 192.168.3.2, 5d22h, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 5d22h, GigabitEthernet1/0/4
D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 5d22h, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 5d22h, GigabitEthernet1/0/4
D EX     10.20.60.0/24
           [170/66560] via 192.168.3.2, 5d22h, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 5d22h, GigabitEthernet1/0/4
D        10.30.70.0/24 [90/16000] via 192.168.3.2, 5d22h, GigabitEthernet1/0/3
                       [90/16000] via 192.168.1.2, 5d22h, GigabitEthernet1/0/4
D        10.30.80.0/24 [90/16000] via 192.168.3.2, 5d22h, GigabitEthernet1/0/3
                       [90/16000] via 192.168.1.2, 5d22h, GigabitEthernet1/0/4
D        10.30.90.0/24 [90/16000] via 192.168.3.2, 5d22h, GigabitEthernet1/0/3
                       [90/16000] via 192.168.1.2, 5d22h, GigabitEthernet1/0/4
      172.16.0.0/30 is subnetted, 2 subnets
```

```
D        172.16.2.0 [90/15360] via 192.168.1.2, 1w5d, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 1w5d, GigabitEthernet1/0/3


R1#
```

show interface GigabitEthernet1/0/3 | i DLY

```
  MTU 1500 bytes, BW 1000000 Kbit/sec,
```

DLY 10 usec

```
,
R1#
```

show interface GigabitEthernet1/0/4 | i DLY

```
  MTU 1500 bytes, BW 1000000 Kbit/sec,
```

DLY 10 usec

```
,
```

Now, modify and increase the delay to interface GigabitEthernet1/0/4. By changing the delay value to 100 (tens of microseconds) the RIB installs the path through interface Gi1/0/3 only.

By looking at the EIGRP topology table, you can confirm that interface Gi1/0/4 is still showing as a feasible successor for the all the prefixes and has a higher total delay.

<#root>

R1#

configure terminal

```
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#
```

interface GigabitEthernet1/0/4

```
R1(config-if)#
```

delay 100

```
R1(config-if)#
```

end

```
R1#

show ip route eigrp


Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX     10.20.40.0/24
           [170/66560] via 192.168.3.2, 00:05:52,

GigabitEthernet1/0/3


D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 00:05:52,

GigabitEthernet1/0/3


D EX     10.20.60.0/24
           [170/66560] via 192.168.3.2, 00:05:52,

GigabitEthernet1/0/3


D        10.30.70.0/24
           [90/16000] via 192.168.3.2, 00:05:52,

GigabitEthernet1/0/3


D        10.30.80.0/24
           [90/16000] via 192.168.3.2, 00:05:52,

GigabitEthernet1/0/3


D        10.30.90.0/24
           [90/16000] via 192.168.3.2, 00:05:52,

GigabitEthernet1/0/3


      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/20480] via 192.168.3.2, 00:05:52, GigabitEthernet1/0/3
D        172.16.4.0 [90/15360] via 192.168.3.2, 00:05:52, GigabitEthernet1/0/3
```

R1#

show interface GigabitEthernet1/0/4 | i DLY

MTU 1500 bytes, BW 1000000 Kbit/sec,

DLY 1000 usec

,

R1#

show ip eigrp topology

EIGRP-IPv4 VR(LAB) Topology Table for AS(100)/ID(192.168.3.1) Codes: P - Passive, A - Active, U - Upda

via 192.168.1.2 (66928640/1392640), GigabitEthernet1/0/4

P 10.20.50.0/24, 1 successors, FD is 8519680 via 192.168.3.2 (8519680/7864320), GigabitEthernet1/0/3

via 192.168.1.2 (73400320/7864320), GigabitEthernet1/0/4

P 10.30.80.0/24, 1 successors, FD is 2048000 via 192.168.3.2 (2048000/1392640), GigabitEthernet1/0/3

via 192.168.1.2 (66928640/1392640), GigabitEthernet1/0/4

P 172.16.2.0/30, 1 successors, FD is 2621440 via 192.168.3.2 (2621440/1966080), GigabitEthernet1/0/3 v

via 192.168.1.2 (73400320/7864320), GigabitEthernet1/0/4

P 192.168.1.0/30, 1 successors, FD is 66191360 via Connected, GigabitEthernet1/0/4 via 192.168.3.2 (32

via 192.168.1.2 (73400320/7864320), GigabitEthernet1/0/4

P 10.10.20.0/24, 1 successors, FD is 163840 via Connected, Loopback20 P 10.30.90.0/24, 1 successors, F

via 192.168.1.2 (66928640/1392640), GigabitEthernet1/0/4

P 172.16.4.0/30, 1 successors, FD is 1966080 via 192.168.3.2 (1966080/1310720), GigabitEthernet1/0/3 P

R1#

show ip eigrp topology 10.20.40.0/24

EIGRP-IPv4 VR(LAB) Topology Entry for AS(100)/ID(192.168.3.1) for 10.20.40.0/24 State is Passive, Quer

Total delay is 120000000 picoseconds

Reliability is 255/255 Load is 1/255 Minimum MTU is 1500 Hop count is 2 Originating router is 172.16.6

Total delay is 1110000000 picoseconds

```
 Reliability is 255/255 Load is 1/255 Minimum MTU is 1500 Hop count is 2 Originating router is 172.16.6
```

```
traceroute 10.20.40.1 source loopback10
```

```
 Type escape sequence to abort. Tracing the route to 10.20.40.1 VRF info: (vrf in name/id, vrf out name
```

```
R1#
```

```
show ip cef 10.20.40.1
```

```
 10.20.40.0/24 nexthop 192.168.3.2 GigabitEthernet1/0/3
```

Modifying the delay can be a useful tool to control the traffic flow and change the overall network behavior. Delay is a cumulative value that grows based on the delay of each segment within the path. It is also important to note that given that bandwidth can be used by other protocols calculations, changes to the interface delay parameter is a preferred method. The changes to delay, however, are useful only in scenarios where one path is preferred over another for all the routes being received.



Note: Be careful when you select the new delay value, you do not want to increase the delay to a point where EIGRP no longer views those routes as a feasible successor.

Scenario 2: Influence Path Selection with the Use an Offset-List

In this scenario, the interesting traffic or prefix that needs to be manipulated is selected with the use of an ACL. An ACL is used to match these prefixes and, for this example, next configuration is added to manipulate the traffic destined to subnets 10.20.60.0/24 and 10.30.90.0/24.



Note: As a reminder, when specifying subnets such as 10.20.60.0/24 and 10.30.90.0/24 in the ACL, all more specific prefixes (for example, subnets with longer prefix lengths that fall within these ranges, such as 10.20.60.128/25 or 10.30.90.192/26) are also matched and affected by the same configuration. This behavior is inherent to how ACL wildcards operate. Be sure to consider this when defining ACL entries to avoid the unintentional inclusion of more specific routes.

<#root>

```
R1#
```

```
configure terminal
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#
```

```
access-list 20 permit 10.20.60.0 0.0.0.255
```

```
R1(config)#
```

```
access-list 30 permit 10.30.90.0 0.0.0.255
```

```
!
R1#
```

```
show access-lists 20
```

```
Standard IP access list 20
    10 permit 10.20.60.0, wildcard bits 0.0.0.255
R1#
```

```
show access-lists 30
```

```
Standard IP access list 30
    10 permit 10.30.90.0, wildcard bits 0.0.0.255
```

The objective is to modify the metric of the specific prefixes but without affecting all other EIGRP traffic. This example uses an offset-list to add an offset to the metric of the selected prefixes (10.20.60.0/24 and 10.30.90.0/24) in the inbound direction of R1.

The idea is to prefer the path through R2 via interface Gi1/0/4 when reaching the subnet 10.20.60.0/24 (from R1) and prefer the path through R3 via interface Gi1/0/3 when reaching the subnet 10.30.90.0/24 (from R1).

The configuration make use of the command **offset-list {ACL name|ACL number} {in|out} <offset> <interface>** as shown next:

<#root>

```
R1#
```

```
configure terminal
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#
```

```
router eigrp LAB
```

```
R1(config-router)#
```

```
address-family ipv4 unicast autonomous-system 100
```

```
R1(config-router-af)#
```

```
topology base
```

```
R1(config-router-af-topology)#
```

```
offset-list 20 in 200 GigabitEthernet1/0/3

R1(config-router-af-topology)#

end
```

The results of the configuration can be verified by checking the RIB, the Forwarding Information Base (FIB) and EIGRP Topology table. In the next outputs, it can be seen that the offset applied to the interface Gi1/0/3 affected the metric of this specific prefix, in other words, making this path less desirable:

<#root>

```
R1#

show ip route 10.20.60.0

 Routing entry for 10.20.60.0/24 Known via "eigrp 100", distance 170, metric 66560, precedence routine

via GigabitEthernet1/0/4

 Route metric is 66560, traffic share count is 1 Total delay is 120 microseconds, minimum bandwidth is

R1#

show ip cef 10.20.60.0

 10.20.60.0/24

nexthop 192.168.1.2 GigabitEthernet1/0/4

 R1#

show ip eigrp topology 10.20.60.0/24

 EIGRP-IPv4 VR(LAB) Topology Entry for AS(100)/ID(192.168.3.1) for 10.20.60.0/24 State is Passive, Quer

GigabitEthernet1/0/3

), from 192.168.3.2, Send flag is 0x0 Composite metric is (8519880/7864520), route is External Vector m

Total delay is 120003052 picoseconds <---

 Reliability is 255/255 Load is 1/255 Minimum MTU is 1500 Hop count is 2 Originating router is 172.16.6
```

A similar process is completed for prefix 10.30.90.0/24, the offset-list is added now to prefer R3 path through interface Gi1/0/3 (but applying the offset to Gi1/0/4). In a similar way, by reviewing the RIB, FIB and EIGRP topology, it can be seen that the preferred path for the selected prefix is through R3:

<#root>

R1#

configure terminal

Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#

router eigrp LAB

R1(config-router)#

address-family ipv4 unicast autonomous-system 100

R1(config-router-af)#

topology base

R1(config-router-af-topology)#

offset-list 30 in 300 gigabitEthernet 1/0/4

R1(config-router-af-topology)#

end

R1#

show ip route 10.30.90.0

Routing entry for 10.30.90.0/24
  Known via "eigrp 100", distance 90, metric 16000, precedence routine (0), type internal
  Redistributing via eigrp 100
  Last update from 192.168.3.2 on

GigabitEthernet1/0/3

, 00:00:25 ago
  Routing Descriptor Blocks:
  * 192.168.3.2, from 192.168.3.2, 00:00:25 ago, via GigabitEthernet1/0/3
      Route metric is 16000, traffic share count is 1
      Total delay is 21 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

R1#

show ip cef 10.30.90.0

 10.30.90.0/24

```
nexthop 192.168.3.2 GigabitEthernet1/0/3
```

```
R1#
```

```
show ip eigrp topology 10.30.90.0/24
```

```
 EIGRP-IPv4 VR(LAB) Topology Entry for AS(100)/ID(192.168.3.1) for 10.30.90.0/24 State is Passive, Quer
```

```
Total delay is 21254578 picoseconds <---
```

```
 Reliability is 255/255 Load is 1/255 Minimum MTU is 1500 Hop count is 2 Originating router is 172.16.6
```

Looking at **show ip route eigrp** command you can confirm that the configuration is successful and only the specific prefixes were affected and all other routes remained intact. Also running a **traceroute** confirms that the traffic is taking the desired path:

<#root>

```
R1#
```

```
show ip route eigrp
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX     10.20.40.0/24
           [170/66560] via 192.168.3.2, 00:22:32, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:22:32, GigabitEthernet1/0/4
D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 00:22:32, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:22:32, GigabitEthernet1/0/4

D EX 10.20.60.0/24 [170/66560] via 192.168.1.2, 00:16:54, GigabitEthernet1/0/4

D        10.30.70.0/24
           [90/16000] via 192.168.3.2, 00:22:32, GigabitEthernet1/0/3
           [90/16000] via 192.168.1.2, 00:22:32, GigabitEthernet1/0/4
D        10.30.80.0/24
```

```
           [90/16000] via 192.168.3.2, 00:22:32, GigabitEthernet1/0/3
           [90/16000] via 192.168.1.2, 00:22:32, GigabitEthernet1/0/4

D 10.30.90.0/24 [90/16000] via 192.168.3.2, 00:04:56, GigabitEthernet1/0/3

      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 00:22:32, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 00:22:32, GigabitEthernet1/0/3

R1#

traceroute 10.20.60.1 source loop10

Type escape sequence to abort.
Tracing the route to 10.20.60.1
VRF info: (vrf in name/id, vrf out name/id)


1 192.168.1.2 1 msec 1 msec 0 msec <--- R2

  2 172.16.2.2 1 msec 1 msec 0 msec
  3 172.16.6.2 1 msec 1 msec *

R1#

traceroute 10.30.90.1 source loop10

Type escape sequence to abort.
Tracing the route to 10.30.90.1
VRF info: (vrf in name/id, vrf out name/id)


 1 192.168.3.2 0 msec 1 msec 0 msec <--- R3

  2 172.16.4.2 1 msec 1 msec *
```

Scenario 3: Influence Path Selection with Summarization

In this scenario route summarization is used to prefer one path over the other. EIGRP has the flexibility to configure a summary route per interface, and in this example a summary route is configured on R4 to summarize the 10.30.x.x prefixes and another one for the 10.20.x.x prefixes. The idea is that R4 announces the summary route 10.30.0.0/16 over interface GigabitEthernet1/0/1 and summary route 10.20.0.0/16 over interface GigabitEthernet1/0/2, and with this configuration traffic is influenced by the longest match preference. This causes that traffic source from R1 and destined to 10.30.x.x subnets selects the path through R3 and traffic destined to subnets 10.20.x.x selects the path through R2. The configuration is shown next:

<#root>

R4#

```
configure terminal
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#
router eigrp LAB
R4(config-router)#
address-family ipv4 unicast autonomous-system 100
R4(config-router-af)#
af-interface gigabitEthernet 1/0/1
R4(config-router-af-interface)#
summary-address 10.30.0.0/16

R4(config-router-af-interface)#
exit
R4(config-router-af)#
af-interface gigabitEthernet 1/0/2
R4(config-router-af-interface)#
summary-address 10.20.0.0/16
R4(config-router-af-interface)#
end
R4#
```

Now, by checking the routing table from R1, it can be verified that there is a summary route for 10.20.0.0/16 being learned through interface GigabitEthernet1/0/3 (connected to R3) and a summary route 10.30.0.0/16 learned over GigabitEthernet1/0/4 (connected to R2). The result of this configuration is that traffic with a destination of 10.20.60.1 is routed trough R2 and traffic with destination of 10.30.90.1 is routed though R3. The reason is that R1 prefers the longest match prefixes that are still learned through the other interface, and can be confirmed via FIB and traceroute outputs:

```
<#root>
R1#
show ip route eigrp
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 14 subnets, 3 masks

D 10.20.0.0/16 [90/66560] via 192.168.3.2, 00:00:16, GigabitEthernet1/0/3


D EX     10.20.40.0/24
           [170/66560] via 192.168.1.2, 00:00:16, GigabitEthernet1/0/4
D EX     10.20.50.0/24
           [170/66560] via 192.168.1.2, 00:00:16, GigabitEthernet1/0/4
D EX     10.20.60.0/24
           [170/66560] via 192.168.1.2, 00:00:16, GigabitEthernet1/0/4

D 10.30.0.0/16 [90/16000] via 192.168.1.2, 00:00:44, GigabitEthernet1/0/4


D        10.30.70.0/24
           [90/16000] via 192.168.3.2, 00:00:44, GigabitEthernet1/0/3
D        10.30.80.0/24
           [90/16000] via 192.168.3.2, 00:00:44, GigabitEthernet1/0/3
D        10.30.90.0/24
           [90/16000] via 192.168.3.2, 00:00:44, GigabitEthernet1/0/3
      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 02:42:44, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 02:42:44, GigabitEthernet1/0/3

R1#
```

show ip route 10.20.0.0

```
Routing entry for 10.20.0.0/16
  Known via "eigrp 100", distance 90, metric 66560, precedence routine (0), type internal
  Redistributing via eigrp 100


Last update from 192.168.3.2 on GigabitEthernet1/0/3, 00:12:07 ago


  Routing Descriptor Blocks:
  * 192.168.3.2, from 192.168.3.2, 00:12:07 ago, via GigabitEthernet1/0/3
      Route metric is 66560, traffic share count is 1
      Total delay is 120 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

R1#
```

show ip route 10.30.0.0

```
Routing entry for 10.30.0.0/16
  Known via "eigrp 100", distance 90, metric 16000, precedence routine (0), type internal
  Redistributing via eigrp 100
```

Last update from 192.168.1.2 on GigabitEthernet1/0/4, 00:12:50 ago

```
  Routing Descriptor Blocks:
  * 192.168.1.2, from 192.168.1.2, 00:12:50 ago, via GigabitEthernet1/0/4
      Route metric is 16000, traffic share count is 1
      Total delay is 21 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

R1#

show ip cef exact-route 10.10.10.1 10.20.60.1

```
 10.10.10.1 -> 10.20.60.1 =>IP adj out of GigabitEthernet1/0/4, addr 192.168.1.2
```

R1#

traceroute 10.20.60.1 source loop10

```
 Type escape sequence to abort. Tracing the route to 10.20.60.1 VRF info: (vrf in name/id, vrf out name
```

1 192.168.1.2 1 msec 1 msec 0 msec <--- R2

```
 2 172.16.2.2 1 msec 1 msec 0 msec 3 172.16.6.2 1 msec 1 msec *    R1#
```

show ip cef exact-route 10.10.10.1 10.30.90.1

```
 10.10.10.1 -> 10.30.90.1 =>IP adj out of GigabitEthernet1/0/3, addr 192.168.3.2 R1#
```

traceroute 10.30.90.1 source loop10

```
 Type escape sequence to abort. Tracing the route to 10.30.90.1 VRF info: (vrf in name/id, vrf out name
```

1 192.168.3.2 1 msec 0 msec 1 msec <--- R3

```
 2 172.16.4.2 0 msec 1 msec *
```

Scenario 4: Influence Path Selection with the Use of Leak Maps

Using leak-maps during the advertisement of summary routes provides a flexible mechanism to announce more specific routes selectively and then take advantage from the longest match to prefer a desired path.

In this example, a summary route 10.0.0.0/8 is advertised from R4 on both interfaces (Gi1/0/1 and Gi1/0/2). Take a look at the configuration:

```
<#root>

R4#

configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#

router eigrp LAB

R4(config-router)#

address-family ipv4 unicast autonomous-system 100

R4(config-router-af)#

af-interface GigabitEthernet1/0/1

R4(config-router-af-interface)#

summary-address 10.0.0.0 255.0.0.0

R4(config-router-af-interface)#

exit

R4(config-router-af)#

af-interface GigabitEthernet1/0/2

R4(config-router-af-interface)#

summary-address 10.0.0.0 255.0.0.0

R4(config-router-af-interface)#

end
```

The previous configuration is reflected in R1 routing table as shown next, this however, is still load balancing the traffic across the two paths from R1:

```
<#root>

R1#
```

```
show ip route eigrp


Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks

D 10.0.0.0/8 [90/16000] via 192.168.3.2, 00:04:16, GigabitEthernet1/0/3 [90/16000] via 192.168.1.2, 00:0


      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 03:50:08, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 03:50:08, GigabitEthernet1/0/3
```

However, traffic from R1 to subnet 10.20.60.0/24 and 10.30.70.0/24 needs to be preferred over GigabitEthernet1/0/4 (connected to R2). To accomplish this result, a leak-map can be configured on R4 to leak the more specific prefixes but keeping the summarization in place.

<#root>

```
R4#

configure terminal


Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#

ip prefix-list LEAKED-PREFIXES permit 10.20.60.0/24

R4(config)#

ip prefix-list LEAKED-PREFIXES permit 10.30.70.0/24

R4(config)#

route-map LEAKED-PREFIXES

 R4(config-route-map)#

match ip address prefix-list LEAKED-PREFIXES
```

```
 R4(config-route-map)#
exit

R4(config)#
router eigrp LAB

R4(config-router)#
address-family ipv4 unicast autonomous-system 100

R4(config-router-af)#
af-interface GigabitEthernet1/0/1

R4(config-router-af-interface)#
summary-address 10.0.0.0 255.0.0.0 leak-map LEAKED-PREFIXES

R4(config-router-af-interface)#
end
```

After applying the previous configuration, R1 starts seeing a more specific entry for 10.20.60.0/24 and 10.30.70.0/24 that are now being learned through interface GigabitEthernet1/0/4, as shown next:

```
<#root>
R1#
show ip route eigrp

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
```

```
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        10.0.0.0/8 [90/16000] via 192.168.3.2, 01:26:41, GigabitEthernet1/0/3
                    [90/16000] via 192.168.1.2, 01:26:41, GigabitEthernet1/0/4


D EX 10.20.60.0/24 [170/66560] via 192.168.1.2, 00:01:29, GigabitEthernet1/0/4 D 10.30.70.0/24 [90/16000


      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 05:12:33, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 05:12:33, GigabitEthernet1/0/3

R1#
```

```
show ip cef exact-route 10.10.10.1 10.20.60.1
```

```
10.10.10.1 -> 10.20.60.1 =>IP adj out of GigabitEthernet1/0/4, addr 192.168.1.2

R1#
```

```
show ip cef exact-route 10.10.10.1 10.30.70.1
```

```
10.10.10.1 -> 10.30.70.1 =>IP adj out of GigabitEthernet1/0/4, addr 192.168.1.2
```

Scenario 5: Influence Path Selection by Modifying the Administrative Distance (AD) of a Prefix

The idea of this example is to modify the AD for the prefix 10.30.90.0/24, therefore, the traffic destined to it can be routed via R3.



Note: This approach is another resource to influence EIGRP, however, this is less preferred than the use of an Offset-List. Be careful if you are using multiple routing protocols on the same device as this method can affect them as well.



Note: This method only affects Internal EIGRP routes, the configuration does not modify the AD of External EIGRP routes.

Notice that R1 is learning the route 10.30.90.0/24 through R2 (192.168.1.2) and R3 (192.168.3.2) with the same metric:

<#root>

R1#

show ip route eigrp

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX     10.20.40.0/24
           [170/66560] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:00:26, GigabitEthernet1/0/4
D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:00:26, GigabitEthernet1/0/4
D EX     10.20.60.0/24
           [170/66560] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:00:26, GigabitEthernet1/0/4
D        10.30.70.0/24
           [90/16000] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3
           [90/16000] via 192.168.1.2, 00:00:26, GigabitEthernet1/0/4
D        10.30.80.0/24
           [90/16000] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3
           [90/16000] via 192.168.1.2, 00:00:26, GigabitEthernet1/0/4
```

D 10.30.90.0/24 [90/16000] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3 [90/16000] via 192.168.1.2, 0

```
      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 00:00:26, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 00:00:26, GigabitEthernet1/0/3
```

To accomplish the change, an ACL needs to be configured that is used to match the desired subnet, afterwards the AD of the prefix can be modified by also specifying the advertising neighbor with the use of the command **distance <route AD> <IP Source address> <Wildcard bits> <ACL>**.

In this example, in order to prefer the advertisement from R3, a lower AD value is used (85), the IP address of R3 EIGRP neighbor (192.168.3.2) is added with a wild card of 0.0.0.0 and then the ACL to match the prefix is added:

<#root>

R1#

configure terminal

```
 Enter configuration commands, one per line. End with CNTL/Z. R1(config)#
access-list 30 permit 10.30.90.0 0.0.0.255
 R1(config)#
router eigrp LAB
 R1(config-router)#
address-family ipv4 unicast autonomous-system 100
 R1(config-router-af)#
topology base
 R1(config-router-af-topology)#
distance 85 192.168.3.2 0.0.0.0 30
 R1(config-router-af-topology)#
end
```

The result can be seen in the RIB and FIB output from R1, where the routing entry for 10.30.90.0/24 has its AD changed to 85 and the preferred EIGRP neighbor is R3 (192.168.3.2):

<#root>

R1#

```
show ip route eigrp
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
```

```
D EX     10.20.40.0/24
           [170/66560] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:00:14, GigabitEthernet1/0/4
D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:00:14, GigabitEthernet1/0/4
D EX     10.20.60.0/24
           [170/66560] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3
           [170/66560] via 192.168.1.2, 00:00:14, GigabitEthernet1/0/4
D        10.30.70.0/24
           [90/16000] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3
           [90/16000] via 192.168.1.2, 00:00:14, GigabitEthernet1/0/4
D        10.30.80.0/24
           [90/16000] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3
           [90/16000] via 192.168.1.2, 00:00:14, GigabitEthernet1/0/4
```

D 10.30.90.0/24 [85/16000] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3

```
      172.16.0.0/30 is subnetted, 2 subnets
D        172.16.2.0 [90/15360] via 192.168.1.2, 00:00:14, GigabitEthernet1/0/4
D        172.16.4.0 [90/15360] via 192.168.3.2, 00:00:14, GigabitEthernet1/0/3

R1#
```

show ip route 10.30.90.0

```
Routing entry for 10.30.90.0/24
  Known via "eigrp 100", distance 85, metric 16000, precedence routine (0), type internal
  Redistributing via eigrp 100
  Last update from 192.168.3.2 on GigabitEthernet1/0/3, 00:00:31 ago
  Routing Descriptor Blocks:
```

*** 192.168.3.2, from 192.168.3.2, 00:00:31 ago, via GigabitEthernet1/0/3**

```
      Route metric is 16000, traffic share count is 1
      Total delay is 21 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

R1#
```

show ip cef 10.30.90.0

```
10.30.90.0/24
```

nexthop 192.168.3.2 GigabitEthernet1/0/3

Scenario 6: Influence Path Selection with Route Filtering

In this example, the idea is to selectively influence the path selection by filtering some routes or prefixes coming into R1.

R1 must prefer R2 path when the destination is any of the next subnets 10.30.70.0/24, 10.30.80.0/24 and 10.20.40.0/24. When the destination is subnet 10.30.90.0/24, 10.20.50.0/24 and 10.20.60.0/24 R1 must prefer R3 path.

To accomplish this, a prefix-list is used to match the desired routes and a distribute-list is configured under the EIGRP process to apply the route filter in an inbound direction, as shown next:

```
<#root>
R1#

configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#

ip prefix-list R2-Preferred permit 10.30.70.0/24

R1(config)#

ip prefix-list R2-Preferred permit 10.30.80.0/24


R1(config)#

ip prefix-list R2-Preferred permit 10.20.40.0/24

R1(config)#
R1(config)#

ip prefix-list R3-Preferred permit 10.30.90.0/24

R1(config)#

ip prefix-list R3-Preferred permit 10.20.50.0/24

R1(config)#

ip prefix-list R3-Preferred permit 10.20.60.0/24


R1(config)#

router eigrp LAB

R1(config-router)#
```

```
address-family ipv4 unicast autonomous-system 100

R1(config-router-af)#

topology base

R1(config-router-af-topology)#

distribute-list prefix R2-Preferred in GigabitEthernet1/0/4


R1(config-router-af-topology)#

distribute-list prefix R3-Preferred in GigabitEthernet1/0/3

R1(config-router-af-topology)#

end
```



Note: Notice that the "prefix" option is needed when applying the distribute-list as an **ip prefix-list** is being used to match the desired routes



Note: One of the main differences between methods such as the use of an offset-list, is that the distribute-list prevents the not permitted prefixes from being inserted into the RIB and the EIGRP Topology table.

The result is that R1 routing table shows the desired path selection:

<#root>

R1#

```
show ip route eigrp

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
```

```
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX     10.20.40.0/24
           [170/66560] via 192.168.1.2, 00:00:12,
```

GigabitEthernet1/0/4 <--- R2

```
D EX     10.20.50.0/24
           [170/66560] via 192.168.3.2, 00:00:24,
```

GigabitEthernet1/0/3 <--- R3

```
D EX     10.20.60.0/24
           [170/66560] via 192.168.3.2, 00:00:24,
```

GigabitEthernet1/0/3

```
D        10.30.70.0/24
           [90/16000] via 192.168.1.2, 00:00:12,
```

GigabitEthernet1/0/4

```
D        10.30.80.0/24
           [90/16000] via 192.168.1.2, 00:00:12,
```

GigabitEthernet1/0/4

```
D        10.30.90.0/24
           [90/16000] via 192.168.3.2, 00:00:24,
```

GigabitEthernet1/0/3

Related Information

- Understand and Use the Enhanced Interior Gateway Routing Protocol
- Introduction to EIGRP
- IP Routing Configuration Guide, Cisco IOS 17.x