配置EIGRP以影响路径选择

目录

简介

本文档介绍通过影响不同的增强型内部网关路由协议(EIGRP)功能来创建首选路径的过程。

先决条件

要求

Cisco 建议您了解以下主题：

- 基本IP路由的相关知识
- EIGRP协议知识
- Cisco IOS®XE命令行界面(CLI)知识

使用的组件

本文档不限于特定的软件和硬件版本，但本文档中的信息基于以下软件和硬件版本：

- 路由器ASR 1000
- 路由器ISR 4000
- 思科IOS XE 17.9.x

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原

始（默认）配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

背景信息

EIGRP路径选择会受到协议用于确定到达目的地的最佳路径的各种度量的影响。EIGRP根据不同的度量计算到达目的地的最佳路径，路径选择过程包括评估这些度量以确定最佳路由。EIGRP度量包括带宽、延迟、负载、可靠性和最大传输单位(MTU)。了解这些度量及其重要性有助于网络管理员根据特定要求或网络条件修改EIGRP路径选择。默认情况下，EIGRP仅使用通往目的网络的路径上的最低带宽和总延迟来计算路由度量，而度量值则有所不同。此外，带宽和延迟度量是根据接口上配置的静态值来确定的，这些静态值来自沿路径到达目的地的设备，换句话说，这两个参数不是动态测得的。

除了度量控制，路由过滤还可用于影响EIGRP中的路径选择。路由过滤包括控制允许或拒绝进入或退出路由器路由表的信息。由于各种原因（包括优化路由表或管理网络流量），可以过滤路由。与EIGRP中的路由过滤相关的一些关键功能包括：分发列表、前缀列表、路由映射和泄漏映射。这些机制为控制路由信息提供了一种强大而灵活的方式，网络管理员可使用这些方式定制EIGRP路由表，以满足特定标准并提高网络效率。

场景

在路由协议不断变化的环境中，管理员经常发现自己需要定制路由决策以符合特定网络要求并优化流量。这涉及到利用各种技术和配置来影响路由器做出路径选择决策的方式。接下来的示例提供了不同的备选方案，管理员可以利用策略配置来控制EIGRP路径选择。

1.通过修改延迟度量来影响路径选择

通过调整路由器接口的延迟度量，管理员可以通过影响链路上的此特定参数来影响路由决策。这种精细的操作可以引导流量根据更改后的延迟值选择首选路径。

2. 使用offset-list影响路径选择

使用offset-list可以有选择地修改特定前缀的度量，从而提供在特定接口上影响路径选择的针对性方法。此机制用于增加通过EIGRP获知的路由的传入和传出度量，并有选择地首选特定路径上的某些前缀。

3. 通过汇总影响路径选择

引入总结路由后，管理员可以影响前缀的最长匹配优先级。路由总结会影响路由决策的精细度，从而优化路由表并提高整体网络效率。

4. 使用泄漏映射影响路径选择

在通告总结路由的过程中利用泄漏映射提供一种机制，以便有选择地通告更具体的路由。此方法可确保从战略角度通告汇总信息，从而保持路由灵活性并影响路径选择。

5. 通过修改前缀的管理距离(AD)影响路径选择

更改前缀的管理距离是控制路由信息来源的有用技术。在需要从路由信息库(RIB)中排除来自特定源

的路由的情况下，此功能尤其有用。

6. 通过路由过滤影响路径选择

路由过滤是一种用于控制路由协议内外的特定路由通告或接受的有效方法。它通常用于根据指定条件过滤路由信息，防止通告或获知某些路由。

distribute-list是用于过滤EIGRP中前缀的主要工具之一，它可以与访问列表(ACL)、前缀列表或路由映射配合使用。

使用前缀列表有助于从特定邻居精细过滤前缀。此控制级别对于管理路由更新以修改路径首选项至关重要。

网络图

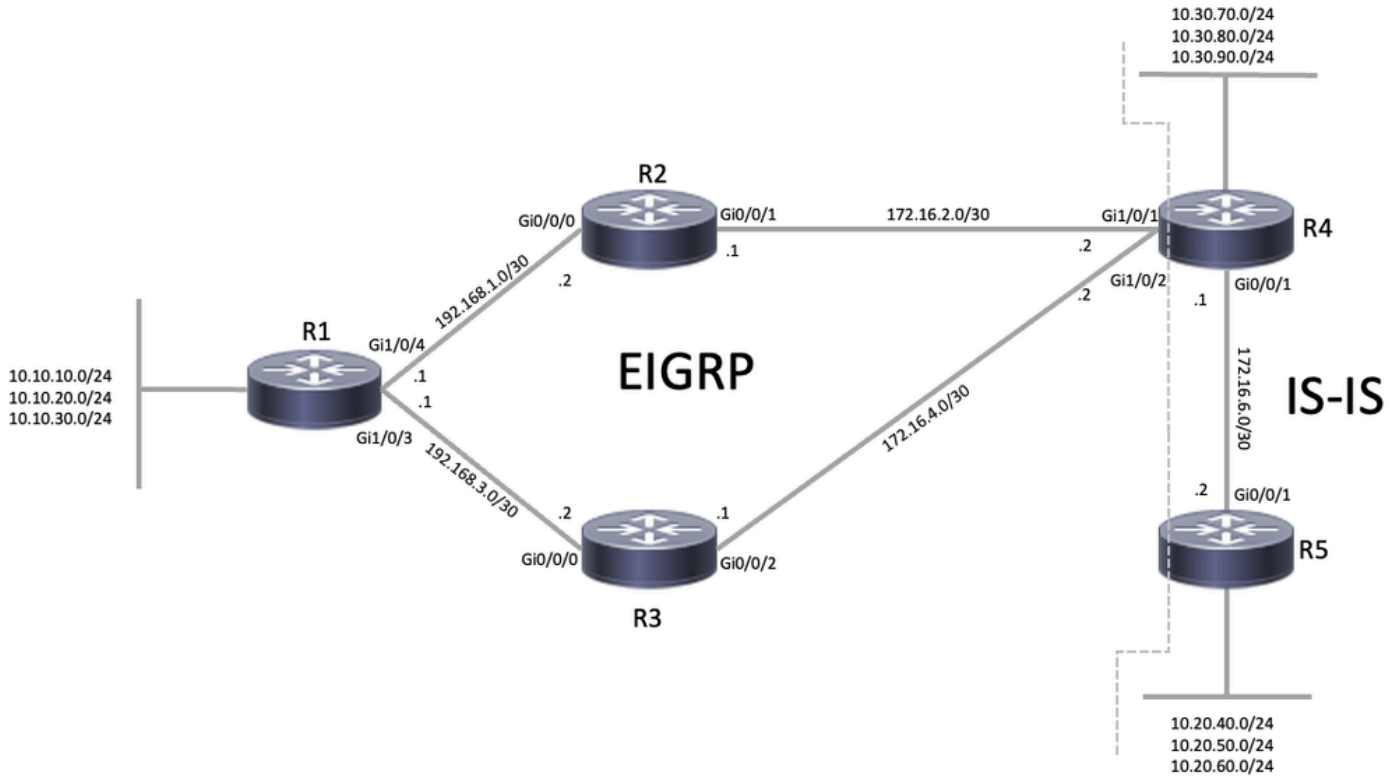

EIGRP拓扑

初始配置

修改任何配置之前，务必查看设备的初始配置和状态（每种场景中的初始配置相同）。根据网络图，R1、R2、R3和R4是EIGRP邻居（每台路由器都有两个邻接关系），R4也是中间系统到中间系统(IS-IS)域的一部分，并且在IS-IS和EIGRP之间执行相互重分配。请注意，R1的路由表中有两条通过EIGRP到10.20.x.x和10.30.x.x子网的路径（通过接口Gi1/0/3和Gi1/0/4），子网10.10.x.x是直接连接的。

R1	
配置	
<#root>	<#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
  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
```

```
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,
       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

对于R2和R3，所有前缀10.10.x.x、10.20.x.x和10.30.x.x都是通过EIGRP获知的。

R2	
配置	状态
<#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 !	<#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**

```
interface GigabitEthernet0/0/1
 ip address 172.16.2.1 255.255.255.252
 negotiation auto
end
```

```
  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

配置

```
<#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
```

状态

```
<#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

配置	状

<#root>

R4#

```
show run | section router 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
```

<#root>

R4#

```
show ip route eigrp


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, 0
i L2     10.20.50.0/24 [115/20] via 172.16.6.2, 0
```

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

```
i L2     10.20.60.0/24 [115/20] via 172.16.6.2, 0

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

      10.0.0.0/8 is variably subnetted, 12 subnet
C        10.30.70.0/24 is directly connected, Loo
L        10.30.70.70/32 is directly connected, Lo
C        10.30.80.0/24 is directly connected, Loo
L        10.30.80.80/32 is directly connected, Lo
C        10.30.90.0/24 is directly connected, Loo
L        10.30.90.90/32 is directly connected, Lo
      172.16.0.0/16 is variably subnetted, 6 subn
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
```

场景1：通过修改延迟度量来影响路径选择

在本例中，延迟值用于影响EIGRP首选通过R3的路径。 执行任何更改之前，您可以确认EIGRP在接口Gi1/0/3和Gi1/04之间执行负载均衡，因为两个接口具有相同的10微秒延迟值。

<#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

,

现在，我们修改并增加接口GigabitEthernet1/0/4的延迟。通过将延迟值更改为100（十微秒），RIB将仅安装通过接口Gi1/0/3的路径。

通过查看EIGRP拓扑表，您可以确认接口Gi1/0/4仍显示为所有前缀的可行后继路由，且具有更高的总延迟。

<#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

修改延迟是控制流量和更改整体网络行为的有用工具。延迟是基于路径中每个网段的延迟而增长的累积值。同样需要注意的是，由于带宽可供其他协议计算使用，因此更改接口延迟参数是首选方法。但是，只有在所有路由优先于另一条路径的情况下，对延迟的更改才有用。

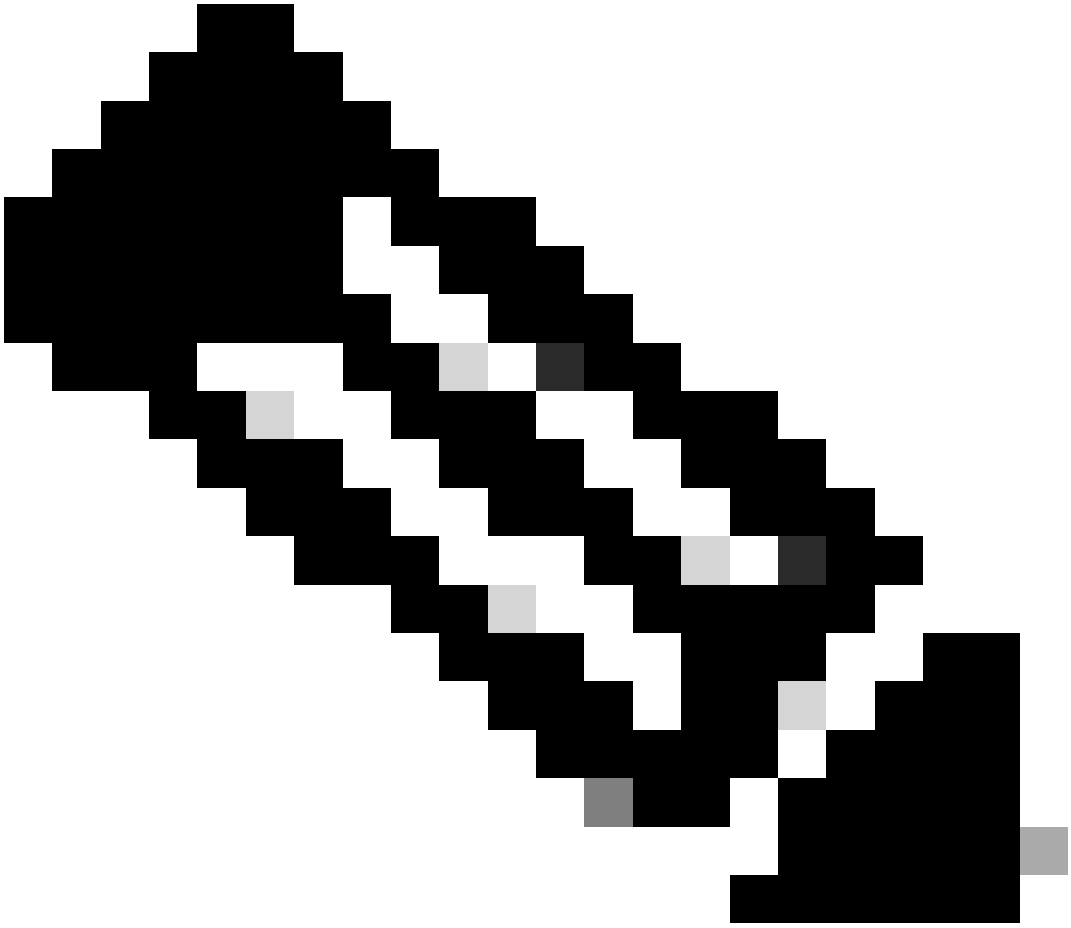

注意：当选择新的延迟值时，请务必谨慎，因为您不希望将延迟增加到EIGRP不再将这些路由视为可行后继路由的点。

场景2：使用偏移列表影响路径选择

在此场景中，使用ACL选择需要处理的相关流量或前缀。ACL用于匹配这些前缀，并且对于本示例，添加下一个配置以控制发往子网10.20.60.0/24和10.30.90.0/24的流量。

<#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
```

目标是修改特定前缀的度量，但不会影响所有其他EIGRP流量。本示例使用offset-list在R1的入站方向上将偏移量添加到所选前缀(10.20.60.0/24和10.30.90.0/24)的度量中。

其思想是在到达子网10.20.60.0/24（来自R1）时首选通过接口Gi1/0/4的R2路径，在到达子网10.30.90.0/24（来自R1）时首选通过接口Gi1/0/3的R3路径。

该配置使用了命令offset-list {ACL name|ACL number} {in|out} <offset> <interface>，如下所示：

```
<#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
```

通过检查RIB、转发信息库(FIB)和EIGRP拓扑表，可以验证配置结果。在接下来的输出中，可以看到，应用到接口Gi1/0/3的偏移量影响了该特定前缀的度量，也就是说，使此路径不再适合：

```
<#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

对前缀10.30.90.0/24完成了类似的过程，现在添加了offset-list以首选R3通过接口Gi1/0/3的路径（但将偏移应用于Gi1/0/4）。同样，通过查看RIB、FIB和EIGRP拓扑，可以看到所选前缀的首选路径是通过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
```

查看show ip route eigrp命令，您可以确认配置是否成功，以及只有特定前缀受到影响而所有其他路由保持不变。此外，运行traceroute可确认流量所选择的路径为：

<#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 *
```

场景3：通过汇总影响路径选择

在本场景中，路由汇总用于首选一条路径。EIGRP可以灵活地为每个接口配置总结路由，在本例中，在R4上配置总结路由来总结10.30.x.x前缀，并为10.20.x.x前缀配置另一个前缀。其思路是，R4在接口GigabitEthernet1/0/1上通告汇总路由10.30.0.0/16，在接口GigabitEthernet1/0/2上通告汇总路由10.20.0.0/16，借助此配置，流量受最长匹配首选项的影响。这会导致从R1发往10.30.x.x子网的流量选择通过R3的路径，而发往子网10.20.x.x的流量选择通过R2的路径。配置如下所示：

```
<#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#
```

现在，通过检查R1的路由表，可以验证存在一条通过接口GigabitEthernet1/0/3（连接到R3）获知的10.20.0.0/16总结路由和一条通过GigabitEthernet1/0/4（连接到R2）获知的总结路由10.30.0.0/16。此配置的结果是，目的地址为10.20.60.1的流量通过R2路由，目的地址为10.30.90.1的流量通过R3路由。原因是R1首选仍通过其他接口获知的最长匹配前缀，并可通过FIB和traceroute输出进行确认：

```
<#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 *
```

场景4：使用泄漏映射影响路径选择

在通告总结路由期间使用泄漏映射提供了一种灵活的机制，可以选择性地通告更具体的路由，然后利用最长匹配来首选期望的路径。

在本示例中，从R4在两个接口（Gi1/0/1和Gi1/0/2）上通告汇总路由10.0.0.0/8。让我们来看看配置：

```
<#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
```

之前的配置反映在R1的路由表中（如下所示），但是，这仍然对来自R1的两条路径上的流量进行负载均衡：

```
<#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
```

但是，从R1到子网10.20.60.0/24和10.30.70.0/24的流量需要优先于GigabitEthernet1/0/4（连接到R2）。为实现此结果，可以在R4上配置一个泄漏映射，以泄漏更具体的前缀，但保留汇总不变。

```
<#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
```

在应用了之前的配置后，R1开始看到10.20.60.0/24和10.30.70.0/24的更具体的条目，这些条目现在通过接口GigabitEthernet1/0/4获知，如下所示：

```
<#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
```

场景5：通过修改前缀的管理距离(AD)影响路径选择

本示例的目的是修改前缀10.30.90.0/24的AD，因此，发往它的流量可以通过R3路由。

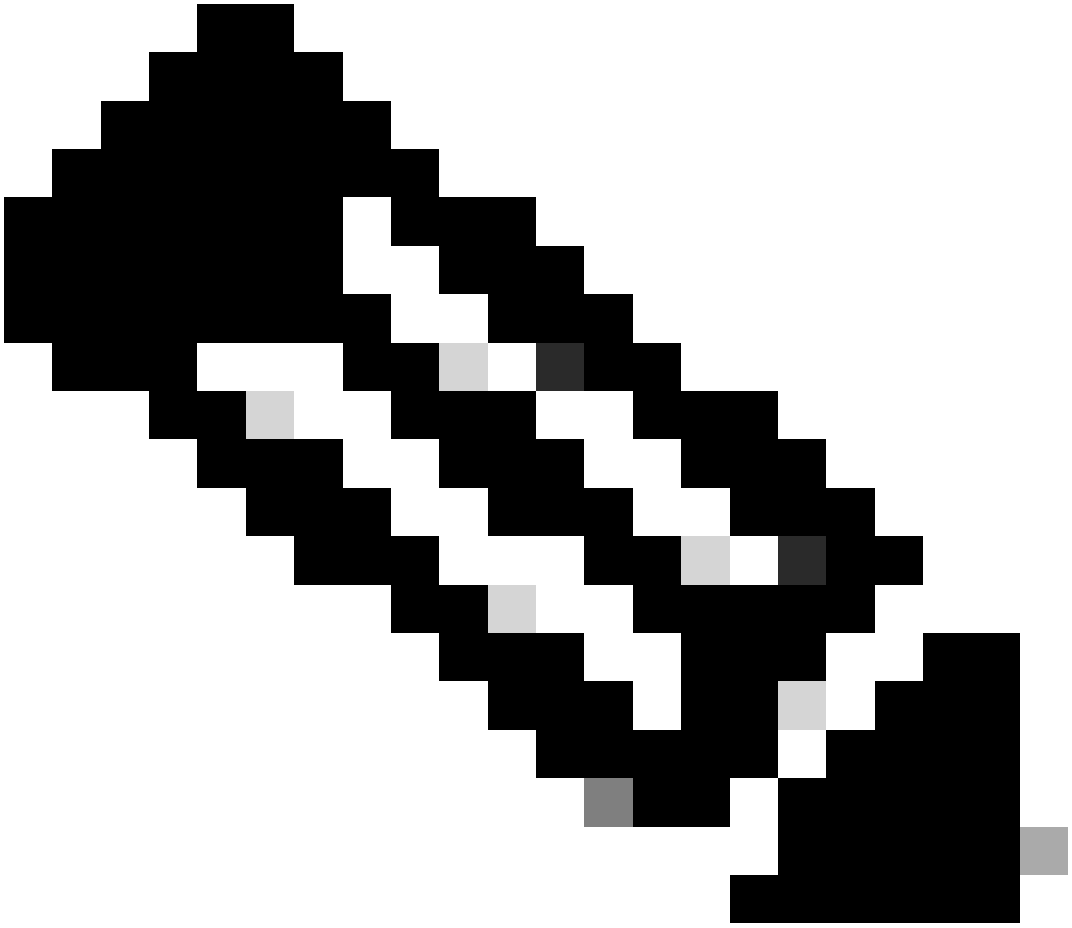

注意：此方法也是一种影响EIGRP的资源，但是，它不像Offset-List那样优先。如果在同一设备上使用多个路由协议，请务必小心，因为此方法也会影响它们。

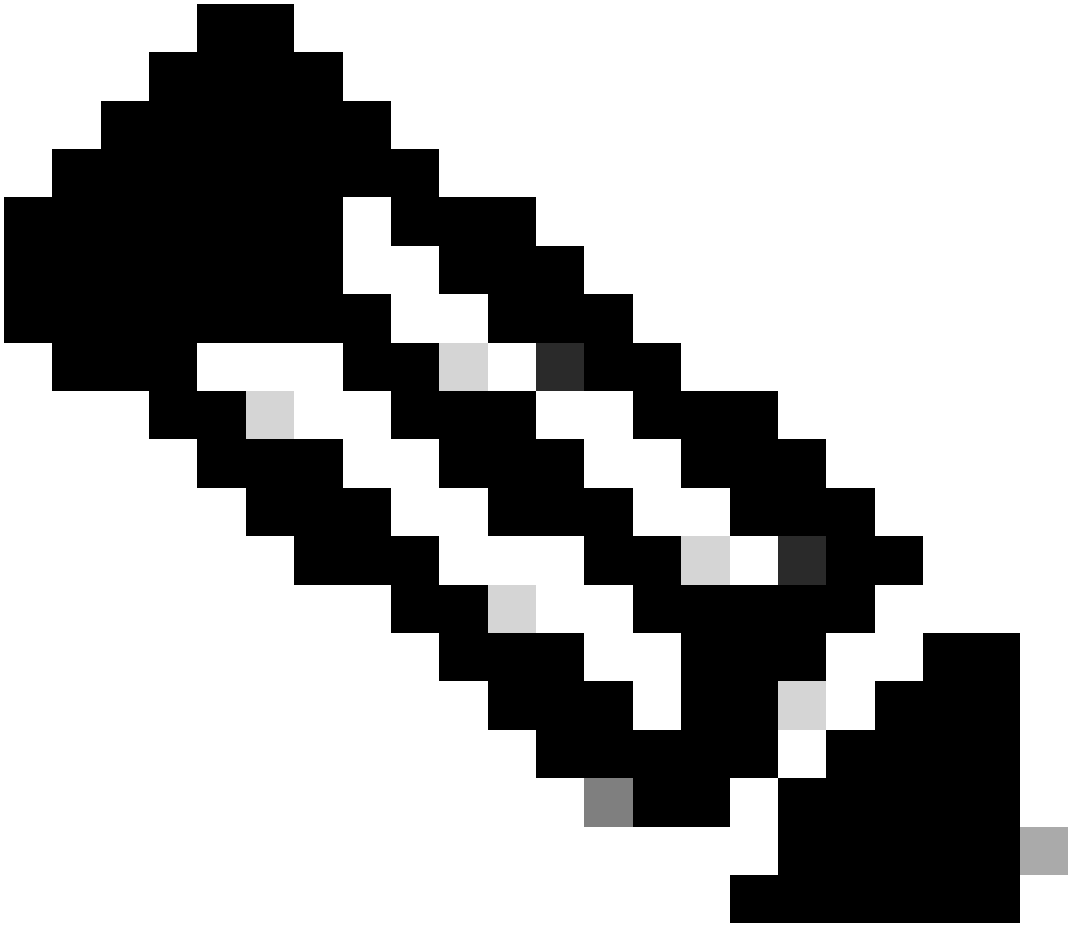

注意：此方法只影响内部EIGRP路由，配置不会修改外部EIGRP路由的AD。

请注意R1正在通过具有相同度量的路由R2 (192.168.1.2)和R3 (192.168.3.2)获知路由10.30.90.0/24：

<#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
```

为了实现此更改，需要配置用于匹配所需子网的ACL，然后可以通过使用命令distance <route AD> <IP Source address> <Wildcard bits> <ACL>指定通告邻居来修改前缀的AD。

在本示例中，为了首选来自R3的通告，使用较低的AD值(85)，使用通配符0.0.0.0添加R3 EIGRP邻居(192.168.3.2)的IP地址，然后添加匹配前缀的ACL：

<#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
```

结果可在R1的RIB和FIB输出中看到，其中10.30.90.0/24的路由条目将AD更改为85，首选EIGRP邻居为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
```

场景6：通过路由过滤影响路径选择

在本例中，我们通过过滤进入R1的一些路由或前缀来选择性影响路径选择。

当目的地为下一子网10.30.70.0/24、10.30.80.0/24和10.20.40.0/24时，R1必须首选R2路径。当目标是子网10.30.90.0/24、10.20.50.0/24和10.20.60.0/24时，R1必须首选R3路径。

为此，使用前缀列表匹配所需的路由，并在EIGRP进程下配置分发列表以在入站方向应用路由过滤器，如下所示：

```
<#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
```

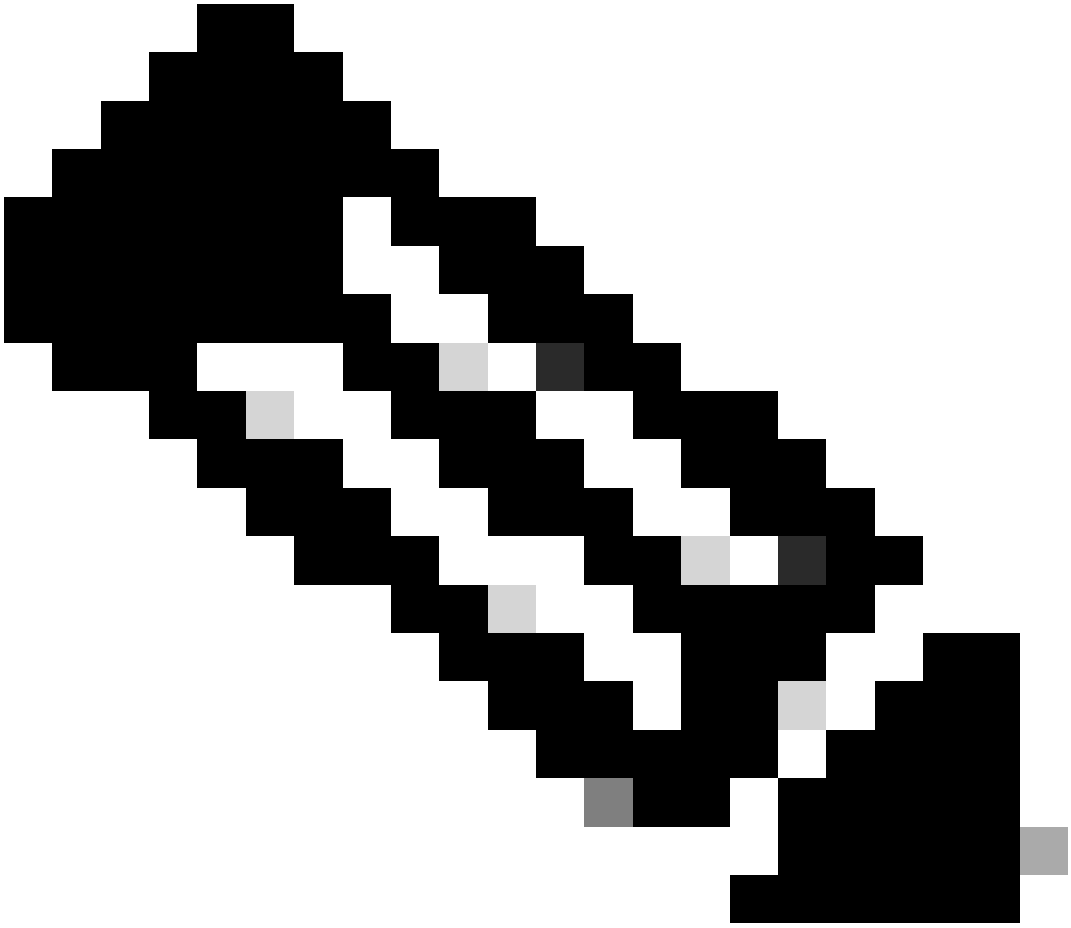

注意：请注意，当使用ip prefix-list匹配所需的路由时，应用distribute-list时需要“prefix”选项

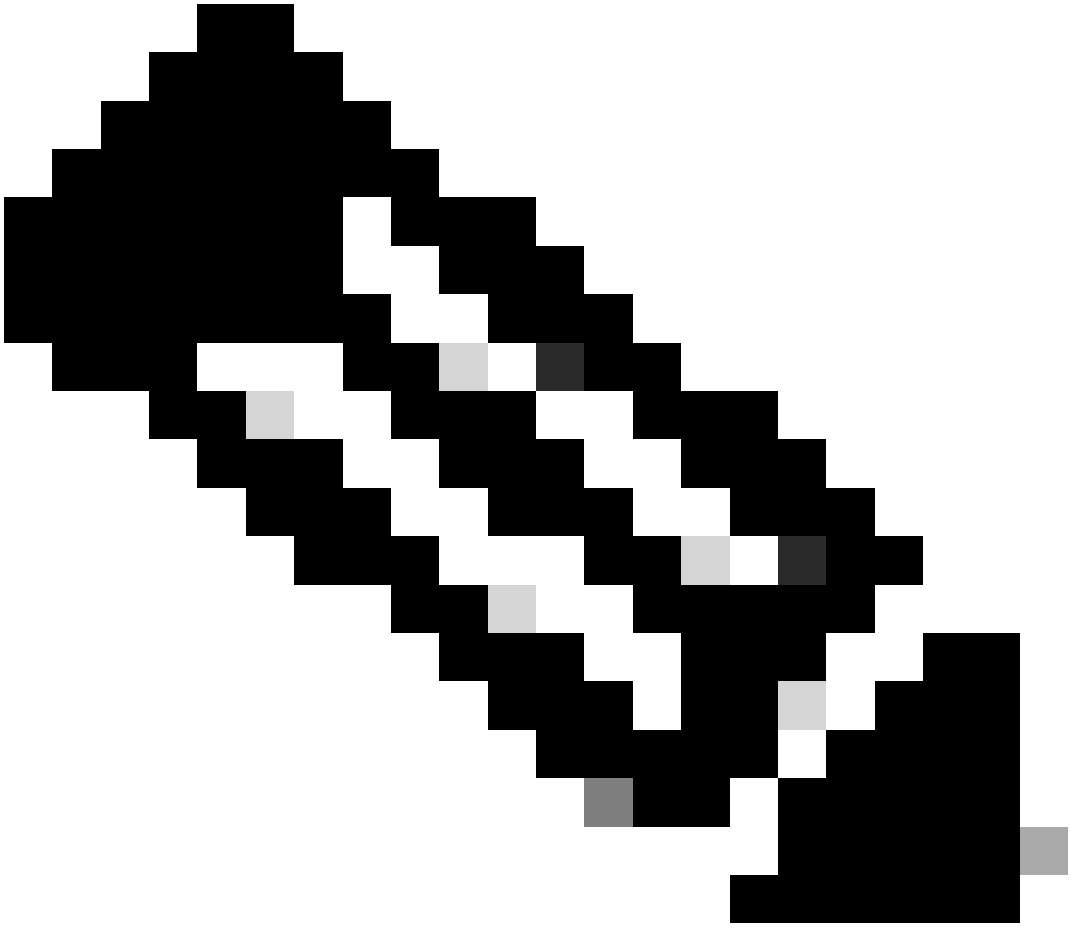

注意：使用offset-list等方法之间的主要区别之一是distribute-list阻止将不允许的前缀插入RIB和EIGRP拓扑表。

结果是R1路由表显示了所需的路径选择：

<#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
```

相关信息

- 了解和使用增强型内部网关路由协议
- EIGRP 介绍
- IP路由配置指南，Cisco IOS XE 17.x

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。