Configuración de EIGRP para Influir en la Selección de Trayectoria

Contenido

Introducción

En este documento se describe el proceso de creación de una ruta preferida influyendo en las diferentes funciones del protocolo de routing de gateway interior mejorado (EIGRP).

Prerequisites

Requirements

Cisco recomienda que tenga conocimiento sobre estos temas:

- Conocimiento del ruteo IP básico
- Conocimiento del protocolo EIGRP
- Conocimiento de la interfaz de línea de comandos (CLI) de Cisco IOS®

Componentes Utilizados

Este documento no se limita a versiones específicas de software y hardware; sin embargo, la información de este documento se basa en estas versiones de software y hardware:

- Router ASR 1000

- Router ISR 4000
- Cisco IOS 17.9.x

La información que contiene este documento se creó a partir de los dispositivos en un ambiente de laboratorio específico. Todos los dispositivos que se utilizan en este documento se pusieron en funcionamiento con una configuración verificada (predeterminada). Si tiene una red en vivo, asegúrese de entender el posible impacto de cualquier comando.

Antecedentes

La selección de la trayectoria EIGRP puede verse influenciada por la manipulación de varias métricas que el protocolo utiliza para determinar la mejor trayectoria a un destino. EIGRP calcula la mejor trayectoria a un destino basándose en diferentes métricas, y el proceso de selección de trayectoria implica la evaluación de estas métricas para determinar la ruta óptima. Las métricas de EIGRP incluyen ancho de banda, retraso, carga, fiabilidad y unidad de transmisión máxima (MTU).

La comprensión de estas métricas y su importancia ayuda a los administradores de red a modificar la selección de rutas EIGRP en función de requisitos o condiciones de red específicas. De forma predeterminada, a partir de los diferentes valores de métrica, EIGRP solo utiliza el ancho de banda mínimo en la ruta a una red de destino y el retraso total para calcular las métricas de ruteo. Además, las métricas de demora y ancho de banda se determinan a partir de valores estáticos configurados en las interfaces de los dispositivos a lo largo de la trayectoria hacia el destino, en otras palabras, estos dos parámetros no se miden dinámicamente.

Aparte de la manipulación de métricas, el filtrado de rutas también se puede utilizar para influir en la selección de rutas en EIGRP. El filtrado de rutas implica controlar la información que se permite o se deniega para entrar o salir de una tabla de enrutamiento del router. El filtrado de rutas se puede realizar por varias razones, incluida la optimización de tablas de routing o la administración del tráfico de red.

Algunas de las características clave relacionadas con el filtrado de rutas en EIGRP incluyen listas de distribución, listas de prefijos, mapas de rutas y mapas de fugas. Estos mecanismos ofrecen una forma potente y flexible de controlar la información de routing que pueden utilizar los administradores de red para personalizar las tablas de routing EIGRP con el fin de cumplir criterios específicos y mejorar la eficacia de la red.

Escenarios

En el panorama dinámico de los protocolos de routing, los administradores a menudo se enfrentan a la necesidad de personalizar las decisiones de routing para ajustarlas a los requisitos específicos de la red y optimizar el flujo de tráfico. Esto implica aprovechar diversas técnicas y configuraciones para influir en la forma en que los routers toman decisiones sobre la selección de rutas.

Estos ejemplos proporcionan diferentes alternativas en las que los administradores pueden emplear configuraciones estratégicas para manipular la selección de rutas EIGRP:

1. Influya en la selección de trayectoria modificando la métrica de retraso.

El ajuste de la métrica de retraso en una interfaz de router permite a los administradores influir en las decisiones de ruteo al afectar este parámetro en particular en un link. Esta sutil manipulación puede guiar al tráfico a tomar las trayectorias preferidas en función de los valores de demora alterados.

2. Influya en la selección de trayectoria con el uso de una lista de desplazamiento.

El uso de una lista de desvío permite la modificación selectiva de métricas para prefijos específicos, proporcionando un enfoque dirigido para influir en la selección de rutas en una interfaz determinada. Este mecanismo se utiliza para aumentar las métricas de entrada y salida a las rutas aprendidas a través de EIGRP y para preferir selectivamente algunos prefijos sobre un trayecto determinado.

3. Influya en la selección del trayecto con el resumen.

La introducción de rutas de resumen permite a los administradores influir en la preferencia de coincidencia más larga para un prefijo. El resumen de rutas puede afectar a la granularidad de las decisiones de routing, optimizando las tablas de routing y mejorando la eficacia general de la red.

4. Influya en la selección de la trayectoria con el uso de mapas de fugas.

Aprovechar los mapas de fuga durante el anuncio de rutas de resumen proporciona un mecanismo para anunciar rutas más específicas de manera selectiva. Este enfoque garantiza que la información resumida se anuncie estratégicamente, manteniendo la flexibilidad de routing e influyendo en la selección de rutas.

5. Influya en la selección de la trayectoria modificando la distancia administrativa (AD) de un prefijo

Modificar la distancia administrativa de un prefijo es una técnica útil para controlar el origen de la información de ruteo. Esto puede ser particularmente útil en escenarios donde las rutas de ciertas fuentes deben ser excluidas de la Base de información de ruteo (RIB).

6. Influya en la selección de rutas con filtrado de rutas

El filtrado de rutas es un método poderoso que se utiliza para controlar el anuncio o la aceptación de rutas específicas dentro o fuera de un protocolo de ruteo. Se suele utilizar para filtrar la información de routing según criterios especificados, lo que impide que se anuncien o aprendan determinadas rutas.

- Una lista de distribución es una de las herramientas principales que se utilizan para filtrar prefijos en EIGRP y puede funcionar junto con una lista de acceso (ACL), una lista de prefijos o un mapa de rutas.
- El uso de una lista de prefijos facilita el filtrado granular de prefijos de vecinos específicos. Este nivel de control es esencial para administrar las actualizaciones de ruteo para modificar la preferencia de trayectoria.

Diagrama de la red

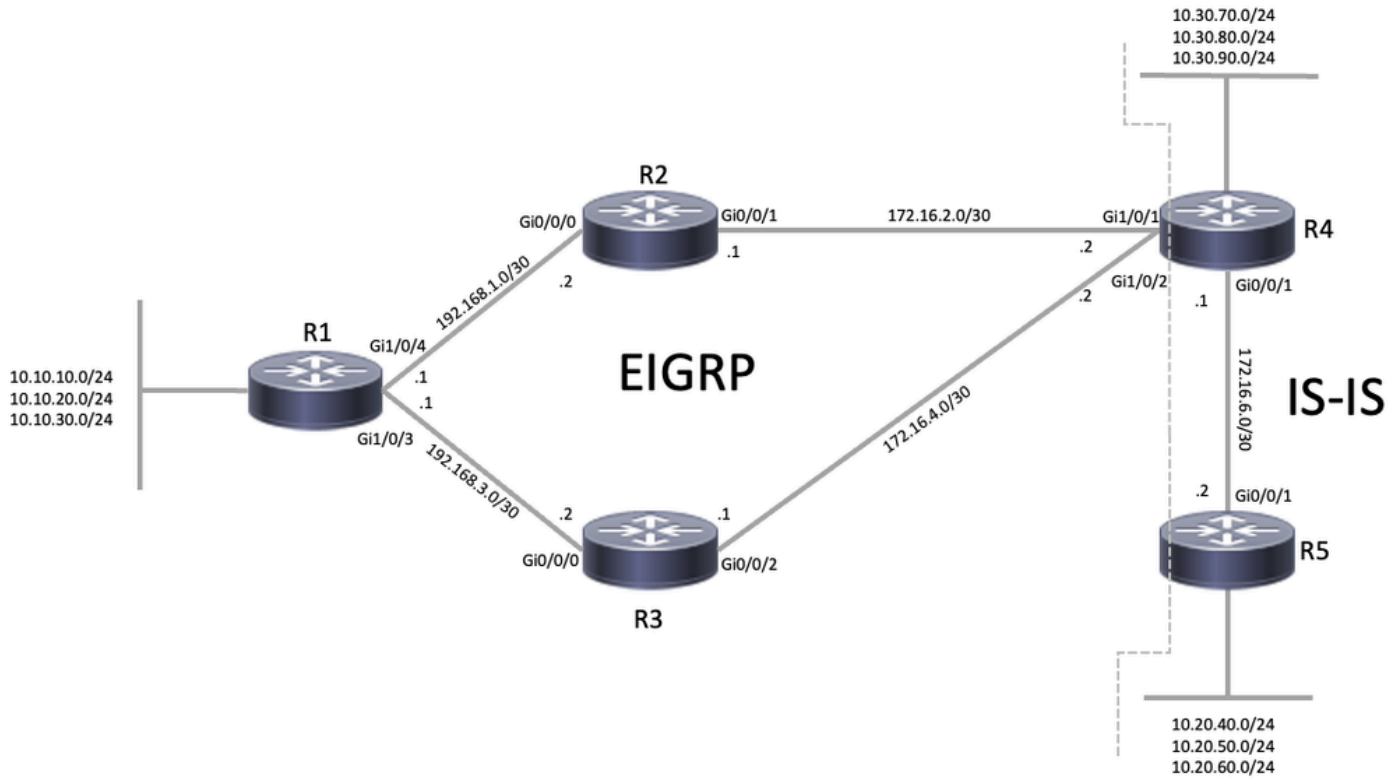

Topología EIGRP

Configuraciones iniciales

Antes de modificar cualquier configuración, es importante revisar la configuración inicial y el estado de los dispositivos (la configuración inicial es la misma en cada escenario). Según el diagrama de red, R1, R2, R3 y R4 son vecinos EIGRP (cada router tiene dos adyacencias) con R4 también formando parte del dominio Sistema intermedio a sistema intermedio (IS-IS) y realizando la redistribución mutua entre IS-IS y EIGRP. Es importante observar que R1 tiene dos trayectorias en la tabla de ruteo (a través de la interfaz Gi1/0/3 y Gi1/0/4) a las subredes 10.20.x.x y 10.30.x.x a través de EIGRP, y las subredes 10.10.x.x están conectadas directamente.

R1	
Configuraciones	Estado
<#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	<#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

```
  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
```

```
       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,
       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
C
```

10.10.10.0/24 is directly connected, Loopback10

```
L        10.10.10.10/32 is directly connected, Loopba
C
```

10.10.20.0/24 is directly connected, Loopback20

```
L        10.10.20.20/32 is directly connected, Loopba
C
```

```
10.10.30.0/24 is directly connected, Loopback30

L        10.10.30.30/32 is directly connected, Loopba
      192.168.1.0/24 is variably subnetted, 2 subnets
C        192.168.1.0/30 is directly connected, Gigabi
L        192.168.1.1/32 is directly connected, Gigabi
      192.168.3.0/24 is variably subnetted, 2 subnets
C        192.168.3.0/30 is directly connected, Gigabi
L        192.168.3.1/32 is directly connected, Gigabi

R1#

show interfaces GigabitEthernet1/0/3

GigabitEthernet1/0/3 is up, line protocol is up (conn
  Hardware is Gigabit Ethernet, address is dc77.4c0d.
  Internet address is 192.168.3.1/30


MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255

  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, media type is 10/100/1000Bas
  input flow-control is on, output flow-control is un
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, 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
     133448 packets input, 10412767 bytes, 0 no buffe
     Received 133325 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ign
     0 watchdog, 133323 multicast, 0 pause input
     0 input packets with dribble condition detected
     207232 packets output, 18832310 bytes, 0 underru
     Output 738 broadcasts (0 IP multicasts)
     0 output errors, 0 collisions, 4 interface reset
     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 swapp

R1#

show interfaces GigabitEthernet1/0/4

GigabitEthernet1/0/4 is up, line protocol is up (conn
  Hardware is Gigabit Ethernet, address is dc77.4c0d.
  Internet address is 192.168.1.1/30


MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255

  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, media type is 10/100/1000Bas
  input flow-control is on, output flow-control is un
```

```
  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
     133435 packets input, 10411748 bytes, 0 no buffe
     Received 133318 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ign
     0 watchdog, 133317 multicast, 0 pause input
     0 input packets with dribble condition detected
     207061 packets output, 18806457 bytes, 0 underru
     Output 714 broadcasts (0 IP multicasts)
     0 output errors, 0 collisions, 4 interface reset
     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 swapp

R1#

show ip eigrp neighbors


EIGRP-IPv4 VR(LAB) Address-Family Neighbors for AS(10
H   Address                 Interface              Ho
                                                   (s
1   192.168.1.2             Gi1/0/4
0   192.168.3.2             Gi1/0/3
```

En el caso de R2 y R3, todos los prefijos 10.10.x.x, 10.20.x.x y 10.30.x.x se aprenden a través de EIGRP.

R2	
Configuraciones	Estado

<#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
	<#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

```
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
```

```
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,
reliability 255/255, txload 1/255, rxload 1/255

```
  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,
reliability 255/255, txload 1/255, rxload 1/255

```
  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

Configuraciones

```
<#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
```

Estado

```
<#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
```

```
 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
```

```
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,
     reliability 255/255, txload 1/255, rxload 1/255

  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,
     reliability 255/255, txload 1/255, rxload 1/255

  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 igr
     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

Configuraciones

```
<#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
```

Estado

```
<#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
```

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

```
      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
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
     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 00:05:38, 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,
     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 00:08:36, 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 ur
     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
```

Escenario 1: Influya en la Selección de Trayectorias Modificando la Métrica de Demora

En este ejemplo, el valor Delay se utiliza para influir en EIGRP para que prefiera la trayectoria a través de R3. Antes de hacer cualquier cambio, puede confirmar que EIGRP está balanceando la carga entre las interfaces Gi1/0/3 y Gi1/04 ya que ambas interfaces tienen el mismo valor Delay

de 10 microsegundos.

<#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

,

Ahora, modifique y aumente la demora para la interfaz GigabitEthernet1/0/4. Al cambiar el valor de demora a 100 (decenas de microsegundos), el RIB instala la trayectoria a través de la interfaz Gi1/0/3 solamente.

Al observar la tabla de topología EIGRP, puede confirmar que la interfaz Gi1/0/4 sigue mostrándose como un sucesor factible para todos los prefijos y tiene un retardo total más alto.

<#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 - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.3.0/30, 1 successors, FD is 1310720
        via Connected, GigabitEthernet1/0/3
P 10.30.70.0/24, 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 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
        via 192.168.1.2 (66846720/1310720), GigabitEthernet1/0/4
P 10.10.30.0/24, 1 successors, FD is 163840
        via Connected, Loopback30
```

```
P 10.20.60.0/24, 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 192.168.1.0/30, 1 successors, FD is 66191360
        via Connected, GigabitEthernet1/0/4
        via 192.168.3.2 (3276800/2621440), GigabitEthernet1/0/3
P 10.20.40.0/24, 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 10.10.20.0/24, 1 successors, FD is 163840
        via Connected, Loopback20
P 10.30.90.0/24, 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 172.16.4.0/30, 1 successors, FD is 1966080
        via 192.168.3.2 (1966080/1310720), GigabitEthernet1/0/3
P 10.10.10.0/24, 1 successors, FD is 163840
        via Connected, Loopback10

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, Query origin flag is 1, 1 Successor(s), FD is 8519680, RIB is 66560
  Descriptor Blocks:
  192.168.3.2 (GigabitEthernet1/0/3), from 192.168.3.2, Send flag is 0x0
      Composite metric is (8519680/7864320), route is External
      Vector metric:
        Minimum bandwidth is 1000000 Kbit


Total delay is 120000000 picoseconds

        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 172.16.6.1
      External data:
        AS number of route is 0
        External protocol is IS-IS, external metric is 20
        Administrator tag is 0 (0x00000000)
  192.168.1.2 (GigabitEthernet1/0/4), from 192.168.1.2, Send flag is 0x0
      Composite metric is (73400320/7864320), route is External
      Vector metric:
        Minimum bandwidth is 1000000 Kbit


Total delay is 1110000000 picoseconds

        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 172.16.6.1
```

```
      External data:
        AS number of route is 0
        External protocol is IS-IS, external metric is 20
        Administrator tag is 0 (0x00000000)

R1#

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/id)
  1 192.168.3.2 1 msec 0 msec 0 msec
  2 172.16.4.2 0 msec 0 msec 1 msec
  3 172.16.6.2 1 msec 1 msec *

R1#

show ip cef 10.20.40.1

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

La modificación del retraso puede ser una herramienta útil para controlar el flujo de tráfico y cambiar el comportamiento general de la red. Retraso es un valor acumulado que crece en función del retraso de cada segmento dentro del trayecto. También es importante tener en cuenta que dado que el ancho de banda puede ser utilizado por otros cálculos de protocolos, los cambios en el parámetro de demora de la interfaz es un método preferido. Sin embargo, los cambios en la demora sólo son útiles en escenarios donde se prefiere una trayectoria sobre otra para todas las rutas que se reciben.

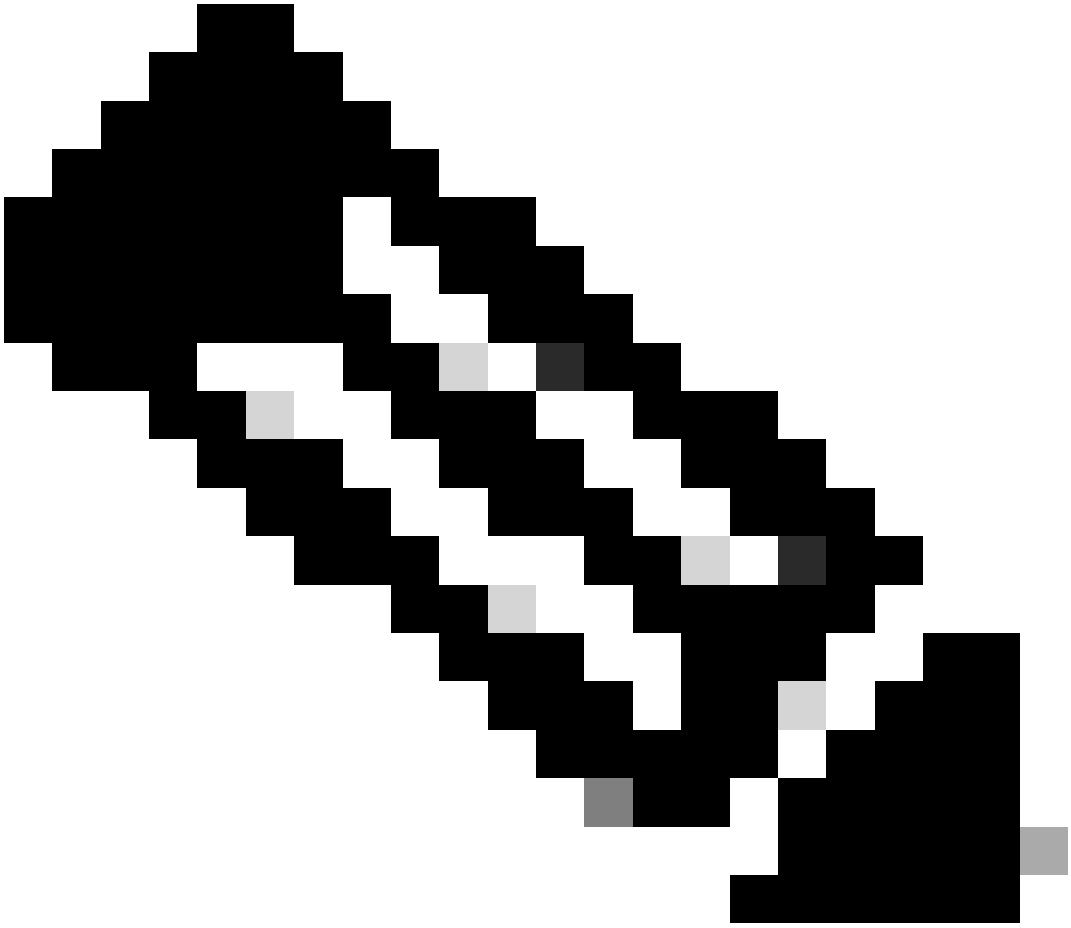

Nota: Tenga cuidado cuando seleccione el nuevo valor de retraso, no desea aumentar el retraso a un punto en el que EIGRP ya no vea esas rutas como un sucesor factible.

Escenario 2: Influya en la selección de rutas con Use an Offset-List

En este escenario, el tráfico interesante o el prefijo que debe manipularse se selecciona con el uso de una ACL. Una ACL se utiliza para hacer coincidir estos prefijos y, para este ejemplo, se agrega la siguiente configuración para manipular el tráfico destinado a las subredes 10.20.60.0/24 y 10.30.90.0/24.

<#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
```

El objetivo es modificar la métrica de los prefijos específicos pero sin afectar al resto del tráfico EIGRP. Este ejemplo utiliza una lista de desplazamiento para agregar un desplazamiento a la métrica de los prefijos seleccionados (10.20.60.0/24 y 10.30.90.0/24) en la dirección entrante de R1.

La idea es preferir el trayecto a través de R2 a través de la interfaz Gi1/0/4 cuando se llega a la subred 10.20.60.0/24 (desde R1) y preferir el trayecto a través de R3 a través de la interfaz Gi1/0/3 cuando se llega a la subred 10.30.90.0/24 (desde R1).

La configuración utiliza el comando offset-list {ACL name|ACL number} {in|out} <offset> <interface> como se muestra a continuación:

<#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
```

Los resultados de la configuración se pueden verificar mediante la verificación del RIB, la Base de información de reenvío (FIB) y la tabla de topología EIGRP. En los siguientes resultados, se puede ver que el desplazamiento aplicado a la interfaz Gi1/0/3 afectó la métrica de este prefijo específico, en otras palabras, haciendo que este trayecto sea menos deseable:

<#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 (0), type external
  Redistributing via eigrp 100
  Last update from 192.168.1.2 on GigabitEthernet1/0/4, 00:01:31 ago
  Routing Descriptor Blocks:
  * 192.168.1.2, from 192.168.1.2, 00:01:31 ago,
```

via GigabitEthernet1/0/4

```
      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 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, Query origin flag is 1, 1 Successor(s), FD is 8519680, RIB is 66560
  Descriptor Blocks:
  192.168.1.2 (GigabitEthernet1/0/4), from 192.168.1.2, Send flag is 0x0
      Composite metric is (8519680/7864320), route is External
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
        Total delay is 120000000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 172.16.6.1
      External data:
        AS number of route is 0
        External protocol is IS-IS, external metric is 20
        Administrator tag is 0 (0x00000000)
  192.168.3.2 (
```

GigabitEthernet1/0/3

```
), from 192.168.3.2, Send flag is 0x0
```

```
      Composite metric is (8519880/7864520), route is External
      Vector metric:
        Minimum bandwidth is 1000000 Kbit

Total delay is 120003052 picoseconds        <---

        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 172.16.6.1
      External data:
        AS number of route is 0
        External protocol is IS-IS, external metric is 20
        Administrator tag is 0 (0x00000000)
```

Un proceso similar se completa para el prefijo 10.30.90.0/24, la lista de desplazamiento se agrega ahora para preferir el trayecto R3 a través de la interfaz Gi1/0/3 (pero aplicando el desplazamiento a Gi1/0/4). De manera similar, al revisar la topología RIB, FIB y EIGRP, se puede ver que la trayectoria preferida para el prefijo seleccionado es a través de 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, Query origin flag is 1, 1 Successor(s), FD is 2048000, RIB is 16000
  Descriptor Blocks:
  192.168.3.2 (GigabitEthernet1/0/3), from 192.168.3.2, Send flag is 0x0
      Composite metric is (2048000/1392640), route is Internal
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
        Total delay is 21250000 picoseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 172.16.6.1
  192.168.1.2 (GigabitEthernet1/0/4), from 192.168.1.2, Send flag is 0x0
      Composite metric is (2048300/1392940), route is Internal
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
```

Total delay is 21254578 picoseconds   <---

```
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 172.16.6.1
```

Si observa el comando show ip route eigrp, puede confirmar que la configuración es exitosa y que sólo los prefijos específicos se vieron afectados y que todas las demás rutas permanecieron intactas. La ejecución también de un traceroute confirma que el tráfico está tomando la trayectoria deseada:

<#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 *
```

Escenario 3: Influya en la selección de rutas con resumen

En este escenario, el resumen de ruta se utiliza para preferir una trayectoria sobre la otra. EIGRP tiene la flexibilidad de configurar una ruta de resumen por interfaz, y en este ejemplo se configura una ruta de resumen en R4 para resumir los prefijos 10.30.x.x y otra para los prefijos 10.20.x.x. La idea es que R4 anuncie la ruta de resumen 10.30.0.0/16 sobre la interfaz GigabitEthernet1/0/1 y la ruta de resumen 10.20.0.0/16 sobre la interfaz GigabitEthernet1/0/2, y con esta configuración el tráfico esté influenciado por la preferencia de coincidencia más larga. Esto hace que el origen de tráfico de R1 y destinado a las subredes 10.30.x.x seleccione la trayectoria a través de R3 y el tráfico destinado a las subredes 10.20.x.x seleccione la trayectoria a través de R2. La configuración se muestra a continuación:

<#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#
```

Ahora, al verificar la tabla de ruteo desde R1, se puede verificar que hay una ruta de resumen para 10.20.0.0/16 aprendida a través de la interfaz GigabitEthernet1/0/3 (conectada a R3) y una ruta de resumen 10.30.0.0/16 aprendida a través de GigabitEthernet1/0/4 (conectada a R2). El resultado de esta configuración es que el tráfico con un destino de 10.20.60.1 se rutea a través de R2 y el tráfico con el destino de 10.30.90.1 se rutea a través de R3. La razón es que R1 prefiere los prefijos de coincidencia más largos que todavía se aprenden a través de la otra interfaz, y se pueden confirmar a través de las salidas de FIB y traceroute:

<#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/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#

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/id)
```

```
1 192.168.3.2 1 msec 0 msec 1 msec          <--- R3
  2 172.16.4.2 0 msec 1 msec *
```

Escenario 4: Influya en la Selección de Trayectorias con el Uso de Leaks Maps

El uso de mapas de fuga durante el anuncio de rutas de resumen proporciona un mecanismo flexible para anunciar rutas más específicas de manera selectiva y luego aprovechar la coincidencia más larga para preferir una trayectoria deseada.

En este ejemplo, una ruta de resumen 10.0.0.0/8 se anuncia desde R4 en ambas interfaces (Gi1/0/1 y Gi1/0/2). Eche un vistazo a la configuración:

```
<#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
```

La configuración anterior se refleja en la tabla de ruteo R1 como se muestra a continuación; sin embargo, esto sigue equilibrando la carga del tráfico a través de las dos trayectorias desde 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, 00:04:16, GigabitEthernet1/0/4

      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
```

Sin embargo, el tráfico de R1 a la subred 10.20.60.0/24 y 10.30.70.0/24 debe ser preferido sobre GigabitEthernet1/0/4 (conectado a R2). Para lograr este resultado, se puede configurar un mapa de fugas en R4 para filtrar los prefijos más específicos pero manteniendo el resumen en su lugar.

<#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
```

Después de aplicar la configuración anterior, R1 comienza a ver una entrada más específica para 10.20.60.0/24 y 10.30.70.0/24 que ahora se aprenden a través de la interfaz GigabitEthernet1/0/4, como se muestra a continuación:

```
<#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] via 192.168.1.2, 00:01:29, GigabitEthernet1/0/4

      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
```

Escenario 5: Influir en la selección de rutas modificando la distancia administrativa (AD) de un prefijo

La idea de este ejemplo es modificar el AD para el prefijo 10.30.90.0/24, por lo tanto, el tráfico destinado a él se puede rutear a través de R3.

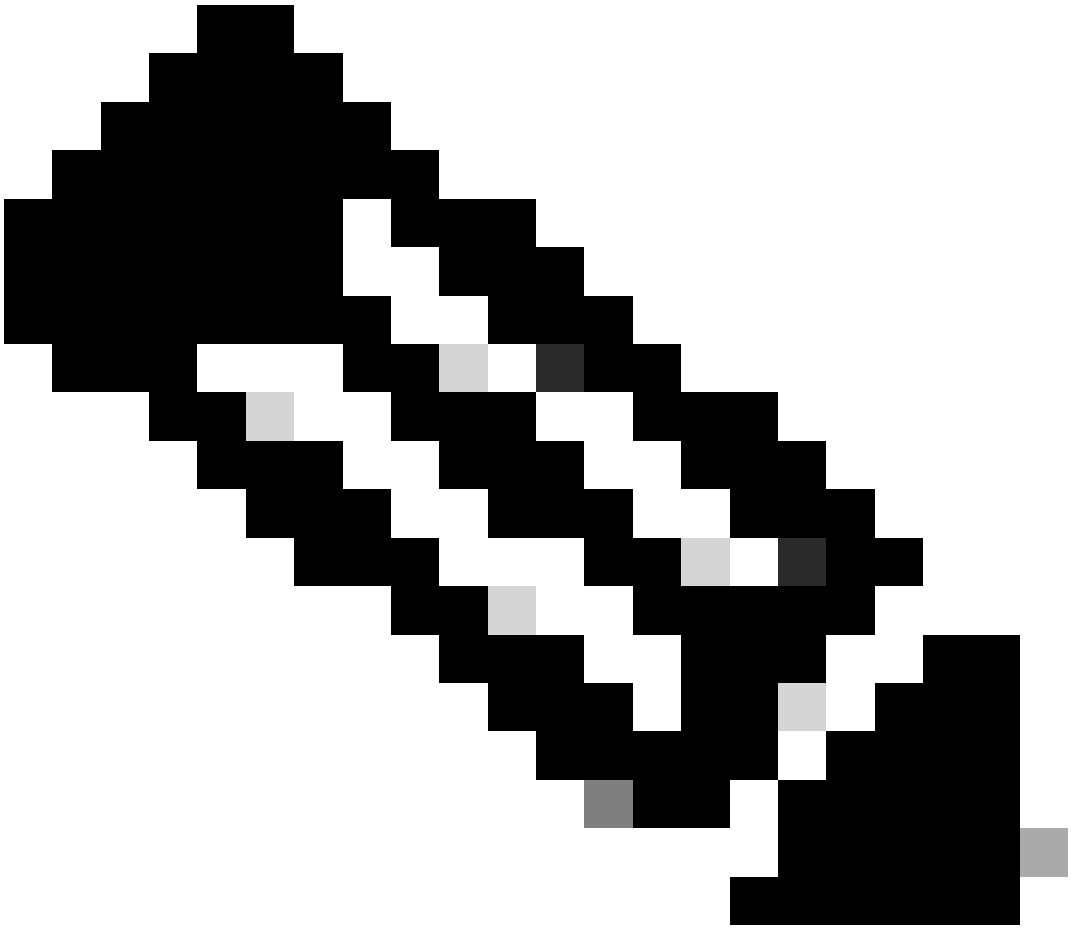

Nota: Este enfoque es otro recurso para influir en EIGRP; sin embargo, es menos preferible que el uso de una lista de compensación. Tenga cuidado si utiliza varios protocolos de ruteo en el mismo dispositivo, ya que este método también puede afectarlos.

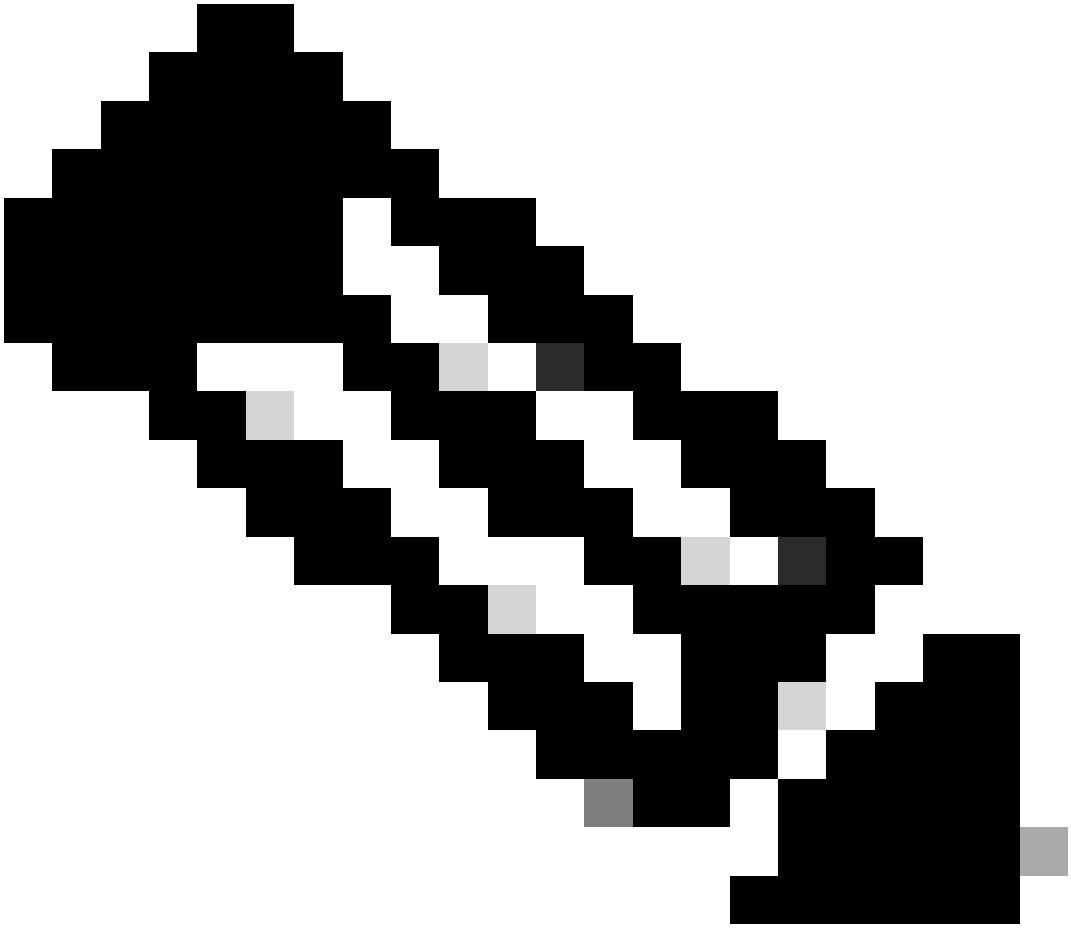

Nota: Este método sólo afecta las rutas EIGRP internas, la configuración no modifica el AD de las rutas EIGRP externas.

Observe que R1 está aprendiendo la ruta 10.30.90.0/24 a través de R2 (192.168.1.2) y R3 (192.168.3.2) con la misma métrica:

<#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, 00:00:26, GigabitEthernet1/0/4

      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
```

Para realizar el cambio, se debe configurar una ACL que se utilice para coincidir con la subred deseada. Posteriormente, se puede modificar el AD del prefijo especificando también el vecino anunciante con el uso del comando distance <route AD> <IP Source address> <Wildcard bits> <ACL>.

En este ejemplo, para preferir el anuncio de R3, se utiliza un valor AD más bajo (85), la dirección IP del vecino EIGRP R3 (192.168.3.2) se agrega con un comodín de 0.0.0.0 y luego se agrega la ACL que coincide con el prefijo:

```
<#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
```

El resultado se puede ver en la salida RIB y FIB de R1, donde la entrada de ruteo para 10.30.90.0/24 tiene su AD cambiado a 85 y el vecino EIGRP preferido es 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
```

Escenario 6: Influya en la selección de rutas con filtrado de rutas

En este ejemplo, la idea es influir selectivamente en la selección de trayectoria filtrando algunas rutas o prefijos que entran en R1.

R1 debe preferir la ruta R2 cuando el destino es cualquiera de las subredes siguientes 10.30.70.0/24, 10.30.80.0/24 y 10.20.40.0/24. Cuando el destino es la subred 10.30.90.0/24, 10.20.50.0/24 y 10.20.60.0/24 R1 debe preferir la ruta R3.

Para lograr esto, se utiliza una lista de prefijos para coincidir con las rutas deseadas y se configura una lista de distribución bajo el proceso EIGRP para aplicar el filtro de rutas en una dirección entrante, como se muestra a continuación:

```
<#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
```

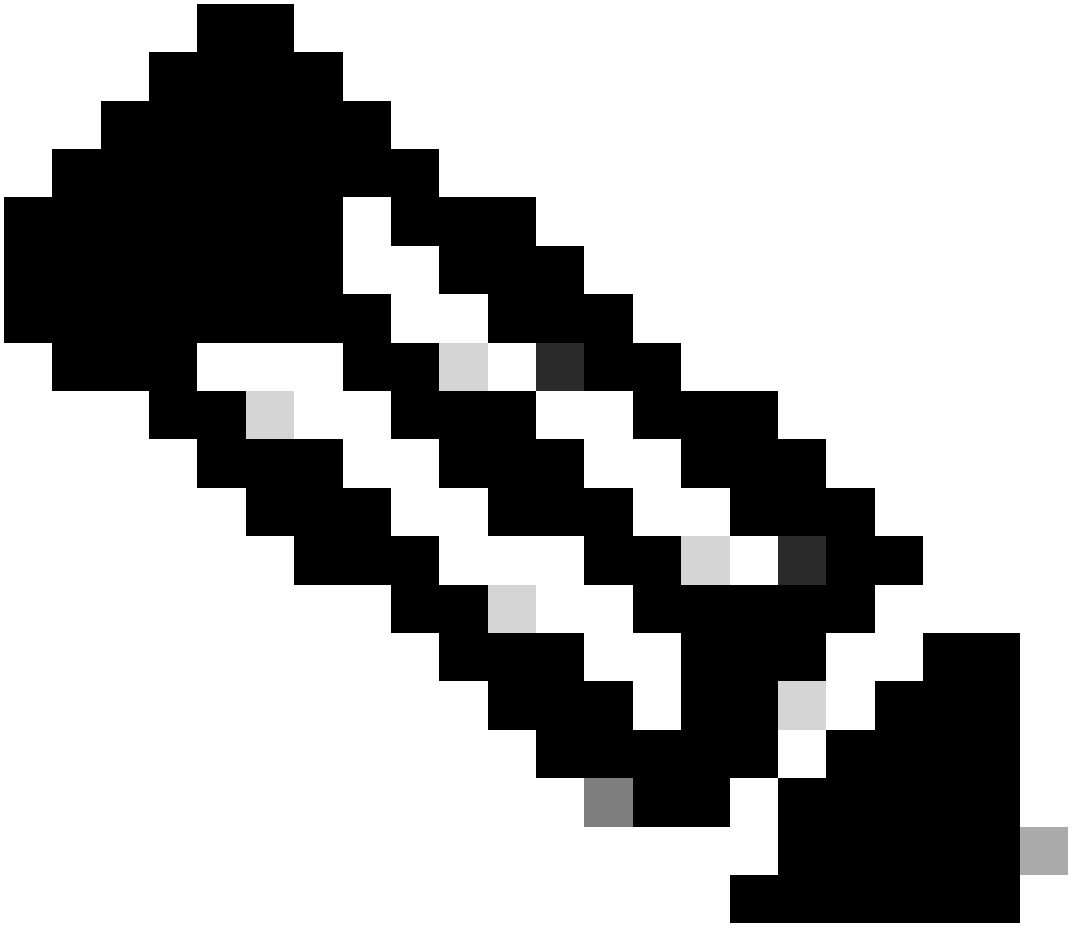

Nota: Observe que la opción "prefix" es necesaria cuando se aplica la lista de distribución como una lista de prefijos ip que se utiliza para coincidir con las rutas deseadas

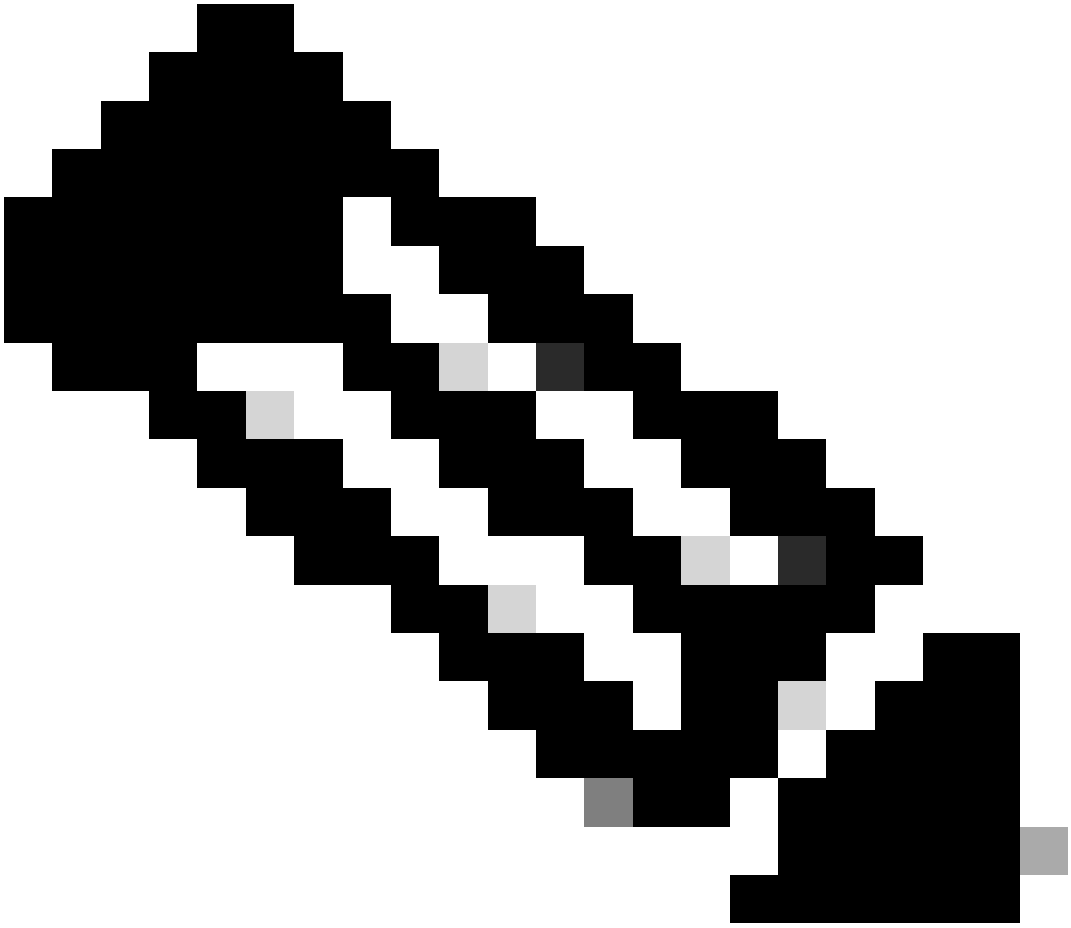

Nota: Una de las principales diferencias entre métodos, como el uso de una lista de desplazamiento, es que la lista de distribución evita que los prefijos no permitidos se inserten en el RIB y en la tabla de topología EIGRP.

El resultado es que la tabla de ruteo R1 muestra la selección de trayectoria deseada:

<#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
```

Información Relacionada

- Comprensión y uso del protocolo de enrutamiento de gateway interior mejorado
- Introducción a EIGRP
- Guía de Configuración de IP Routing, Cisco IOS 17.x

Acerca de esta traducción
Cisco ha traducido este documento combinando la traducción automática y los recursos humanos a fin de ofrecer a nuestros usuarios en todo el mundo contenido en su propio idioma.

Tenga en cuenta que incluso la mejor traducción automática podría no ser tan precisa como la proporcionada por un traductor profesional.

Cisco Systems, Inc. no asume ninguna responsabilidad por la precisión de estas traducciones y recomienda remitirse siempre al documento original escrito en inglés (insertar vínculo URL).