

Setting a Preferred Route by Influencing EIGRP Metrics

Document ID: 13673

Interactive: This document offers customized analysis of your Cisco device.

Introduction

Prerequisites

Requirements

Components Used

Conventions

Background Information

Background – EIGRP Metrics Basics

Possible Configurations

Default Load-sharing Configuration

Change the Interface Delay Parameter at R4

Use an offset-list on R4 to Modify the Composite Metric at R2

Change the Administrative Distance at R2

Change the Bandwidth at R2

NetPro Discussion Forums – Featured Conversations

Related Information

Introduction

This document describes how to create a preferred route by influencing Enhanced Interior Gateway Routing Protocol (EIGRP) metrics. Given the topology shown in the network diagram, this document describes several ways to influence IP traffic from the clients to the servers so that the path R1>R2>R3 is preferred. The goal is to make the path R1>R2>R4 into a backup that is only used in the event of a failure at R3.

Prerequisites

Requirements

This document requires a basic understanding of IP routing and EIGRP routing. To learn more about IP routing and EIGRP, please refer to these documents:

- Routing Basics
- Enhanced Interior Gateway Routing Protocol

Components Used

The information in this document is based on these software and hardware versions.

- EIGRP is supported in Cisco IOS® Software Release 9.21 and later. The information in this document is based on Cisco IOS Software Release 12.3(3).
- EIGRP can be configured on all routers (such as the Cisco 2500 series and the Cisco 2600 series) and on all Layer 3 switches.


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

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Background Information

There are several methods for setting a preferred route by influencing EIGRP metrics. This document describes these methods and details their advantages and disadvantages. This document also discusses the effect of modifying bandwidth, even though doing so is not a viable means of altering the path in this example.



Click on the network diagram to display it in a separate browser window for use as a reference later in this document.

Two of the commands used throughout this document to verify of EIGRP behavior are the **show ip eigrp topology** and the **show ip eigrp topology network-ip subnet-mask** commands.

Note: For additional information about the commands used in this document, refer to the Command Lookup Tool [↗](#) (registered customers only) .

If you have the output of a **show ip eigrp topology** command or a **show ip eigrp topology network-ip subnet-mask** command from your Cisco device, you can use Output Interpreter [↗](#) (registered customers only) to display potential issues and fixes. To use Output Interpreter [↗](#) (registered customers only) , you must be logged in and have JavaScript enabled in your web browser.

Background – EIGRP Metrics Basics

EIGRP updates contain five metrics: minimum bandwidth, delay, load, reliability, and maximum transmission unit (MTU). Of these five metrics, by default, only minimum bandwidth and delay are used to compute best path. Unlike most metrics, minimum bandwidth is set to the minimum bandwidth of the entire path, and it does not reflect how many hops or low bandwidth links are in the path. Delay is a cumulative value which increases by the delay value of each segment in the path. For more information on EIGRP metrics refer to the Enhanced Interior Gateway Routing Protocol White Paper.

Possible Configurations

These configurations can be used to set a preferred route.

Default Load-sharing Configuration

R1

```
R1# show run

Current configuration: 640 bytes
!
version 12.3
!
hostname R1
!
interface Serial0
  no ip address
  encapsulation frame-relay

  !--- Enables Frame Relay encapsulation.

!
interface Serial0.201 point-to-point

  !--- Enables a point-to-point link on the sub-interface.

  ip address 10.1.1.1 255.255.255.0
  frame-relay interface-dlci 201

  !--- Assigns a data-link connection identifier (DLCI)
  !--- to a Frame Relay sub-interface.

!
router eigrp 1
  network 10.0.0.0
!
end
```

Note: The Frame Relay switch is hidden in the network diagram.

```
R1# show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, L1 - ISIS level-1, L2 - ISIS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

  10.0.0.0/24 is subnetted, 3 subnets
D       10.1.3.0 [90/2221056] via 10.1.1.2, 00:07:08, Serial0.201
D       10.1.2.0 [90/2195456] via 10.1.1.2, 00:07:08, Serial0.201
C       10.1.1.0 is directly connected, Serial0.201

R1# show ip eigrp topology 10.1.3.0 255.255.255.0

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2221056
  Routing Descriptor Blocks:
    10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
```

```
Composite metric is (2221056/307200), Route is Internal
Vector metric:
  Minimum bandwidth is 1544 Kbit
  Total delay is 22000 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 2
```

R2

R2# **show run**

```
Current configuration: 618 bytes
!
version 12.3
!
hostname R2
!
interface Ethernet0
 ip address 10.1.2.2 255.255.255.0
 no ip directed-broadcast
!
!
interface Serial0
 no ip address
 encapsulation frame-relay
!
interface Serial0.101 point-to-point
 ip address 10.1.1.2 255.255.255.0
 frame-relay interface-dlci 101
!
router eigrp 1
 network 10.0.0.0
!
end
```

R2# **show ip route**

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, * - candidate default
       U - per-user static route, o - ODR
```

Gateway of last resort is not set

```
10.0.0.0/24 is subnetted, 3 subnets
D    10.1.3.0 [90/307200] via 10.1.2.4, 00:03:47, Ethernet0
     [90/307200] via 10.1.2.3, 00:03:48, Ethernet0
C    10.1.2.0 is directly connected, Ethernet0
C    10.1.1.0 is directly connected, Serial0.101
```

Note: R2 has two equal-cost paths to 10.1.3.0/24 through R3 (10.1.2.3) and R4 (10.1.2.4).

R2# **show ip eigrp topology 10.1.3.0 255.255.255.0**

```
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 307200
  Routing Descriptor Blocks:
    10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
      Composite metric is (307200/281600), Route is Internal
      Vector metric:
        Minimum bandwidth is 10000 Kbit
```

```

    Total delay is 2000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
  Composite metric is (307200/281600), Route is Internal
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 2000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1

```

Note: Both paths have the same Composite Metric (Distance/Reported Distance). The Feasible Distance (FD) to R1 is advertised, which then becomes the Reported Distance for R1.

R3

R3# **show run**

```

Current configuration: 556 bytes
!
version 12.3
!
hostname R3
!
interface Ethernet0/0
 ip address 10.1.2.3 255.255.255.0
 no ip directed-broadcast
!
interface Ethernet0/1
 ip address 10.1.3.3 255.255.255.0
 no ip directed-broadcast
!
router eigrp 1
 network 10.0.0.0
!
end

```

R3# **show ip route**

```

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, ia - ISIS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

```

Gateway of last resort is not set

```

    10.0.0.0/24 is subnetted, 3 subnets
C       10.1.3.0 is directly connected, Ethernet0/1
C       10.1.2.0 is directly connected, Ethernet0/0
D       10.1.1.0 [90/20537600] via 10.1.2.2, 00:16:14, Ethernet0/0

```

R3# **show ip eigrp topology 10.1.3.0 255.255.255.0**

```

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
  Routing Descriptor Blocks:
    0.0.0.0 (Ethernet0/1), from Connected, Send flag is 0x0
      Composite metric is (281600/0), Route is Internal

```

```
Vector metric:
  Minimum bandwidth is 10000 Kbit
  Total delay is 1000 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 0
10.1.2.4 (Ethernet0/0), from 10.1.2.4, Send flag is 0x0
  Composite metric is (307200/281600), Route is Internal
Vector metric:
  Minimum bandwidth is 10000 Kbit
  Total delay is 2000 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 1
```

R3# **show interface ethernet0/1**

```
Ethernet0/1 is up, line protocol is up
  Hardware is AmdP2, address is 0050.7329.52e1 (bia 0050.7329.52e1)
  Internet address is 10.1.3.3/24
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:02, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  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
    291 packets input, 28402 bytes, 0 no buffer
  Received 283 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 input packets with dribble condition detected
  500 packets output, 50876 bytes, 0 underruns
  0 output errors, 0 collisions, 2 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier
  0 output buffer failures, 0 output buffers swapped out
```

R4

R4# **show run**

```
Current configuration: 549 bytes
!
version 12.3
!
hostname R4
!
interface Ethernet0/0
 ip address 10.1.2.4 255.255.255.0
 no ip directed-broadcast
!
interface Ethernet0/1
 ip address 10.1.3.4 255.255.255.0
 no ip directed-broadcast
!
router eigrp 1
 network 10.0.0.0
!
end
```

```
R4# show ip route
```

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, IA - ISIS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 3 subnets
C    10.1.3.0 is directly connected, Ethernet0/1
C    10.1.2.0 is directly connected, Ethernet0/0
D    10.1.1.0 [90/20537600] via 10.1.2.2, 00:17:08, Ethernet0/0
```

```
R4# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
Routing Descriptor Blocks:
0.0.0.0 (Ethernet0/1), from Connected, Send flag is 0x0
Composite metric is (281600/0), Route is Internal
Vector metric:
Minimum bandwidth is 10000 Kbit
Total delay is 1000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 0
10.1.2.3 (Ethernet0/0), from 10.1.2.3, Send flag is 0x0
Composite metric is (307200/281600), Route is Internal
Vector metric:
Minimum bandwidth is 10000 Kbit
Total delay is 2000 microseconds
Reliability is 255/255
Load is 1/255
Minimum MTU is 1500
Hop count is 1
```

Change the Interface Delay Parameter at R4

Because changes to the delay metric are propagated to all downstream routers, changing the interface delay parameter is the preferred method of influencing path selection for these two scenarios:

- The Ethernet segment 10.1.3.0/24 contains only servers and there are no other subnets behind the 10.1.3.0/24 subnet. (This configuration is an ideal setup for a server farm.)
- You wish to influence path selection for all routes learned via EIGRP neighbors on the 10.1.3.0/24 segment.

1. Check the delay on the interface before you make any changes.

It is currently set to the same as R3, as shown here.

```
R4# show interface ethernet0/1
```

```
Ethernet0/1 is up, line protocol is up
Hardware is AmdP2, address is 0050.7329.5321 (bia 0050.7329.5321)
Internet address is 10.1.3.4/24
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
```

```
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:02, output 00:00:02, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
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
  284 packets input, 27914 bytes, 0 no buffer
  Received 276 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 input packets with dribble condition detected
  482 packets output, 49151 bytes, 0 underruns
  0 output errors, 0 collisions, 2 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier
  0 output buffer failures, 0 output buffers swapped out
```

2. Change the delay value on the 10.1.3.0/24 segment.

You must be careful when you select the new delay. You do not want to increase the delay to a point where R2 no longer views the route as a feasible successor.

```
R4# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R4(config)# interface ethernet0/1
```

```
R4(config-if)# delay 120
```

```
!--- Delay is entered in tens of microseconds.
```

```
R4(config-if)# end
```

```
R4#
```

3. Confirm that the delay has changed to 1200 microseconds for this interface.

```
R4# show interface ethernet0/1
```

```
Hardware is AmdP2, address is 0050.7329.5321 (bia 0050.7329.5321)
Internet address is 10.1.3.4/24
MTU 1500 bytes, BW 10000 Kbit, DLY 1200 usec,
  reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:03, output 00:00:00, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
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
  345 packets input, 33508 bytes, 0 no buffer
  Received 333 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 input packets with dribble condition detected
  575 packets output, 57863 bytes, 0 underruns
  0 output errors, 0 collisions, 2 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier
  0 output buffer failures, 0 output buffers swapped out
```

Confirm that R2 only has one 'best' route to 10.1.3.0 via R3.

```
R2# show ip route
```

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, * - candidate default
       U - per-user static route, o - ODR
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 3 subnets
D    10.1.3.0 [90/307200] via 10.1.2.3, 00:02:43, Ethernet0
C    10.1.2.0 is directly connected, Ethernet0
C    10.1.1.0 is directly connected, Serial0.101
```

```
R2# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 307200
Routing Descriptor Blocks:
10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
  Composite metric is (307200/281600), Route is Internal
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 2000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
  Composite metric is (312320/286720), Route is Internal
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 2200 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
```

The **show ip eigrp topology** command shows that the delay metric advertised by R4 has been increased by 200 (to 2200 microseconds), as expected. This increase makes the two routes have different costs and prevents R2 from load balancing.

Note: Because the distance advertised by R4 (286720) is less than the distance advertised by R2 (the feasible distance, 307200), the path is considered loop-free. Because the path advertised by R4 is considered loop-free, it is a feasible successor and is installed immediately if R3 stops advertising a route to 10.1.3.0/24.

```
R1# show ip route
```

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, * - candidate default
       U - per-user static route, o - ODR
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 3 subnets
D    10.1.3.0 [90/2221056] via 10.1.1.2, 00:25:27, Serial0.201
```

```
4.      D      10.1.2.0 [90/2195456] via 10.1.1.2, 00:25:27, Serial0.201
      C      10.1.1.0 is directly connected, Serial0.201
```

```
R1# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2221056
Routing Descriptor Blocks:
 10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
   Composite metric is (2221056/307200), Route is Internal
   Vector metric:
     Minimum bandwidth is 1544 Kbit
     Total delay is 22000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 2
```

Use an offset-list on R4 to Modify the Composite Metric at R2

The Composite Metric on R2 can be modified using an offset-list on router R4. An offset-list value of 20 on R4 increases the Composite Metric for the path R2–R4 by 20 on R2. Therefore, the path R2–R4 becomes a backup path to R2–R3. An offset-list is the preferred method if:

- You only want to influence a particular path being advertised.
- Additional routers are connected to the 10.1.3.0/24 subnet and you do not want to influence paths originated by the routers.

1. Configure an offset-list on R4 that increases (by 20) the delay for any route starting with **10.1.3.x**.

```
R4# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R4(config)# access-list 99 permit 10.1.3.0 0.0.0.255
```

```
R4(config)# router eigrp 1
```

```
R4(config-router)# offset-list 99 out 20 e0/0
```

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

```
R4#
```

2. You can in this output that the offset-list does not change anything in the EIGRP topology table on R4.

The metric changes only when the route is advertised.

```
R4# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
Routing Descriptor Blocks:
 0.0.0.0 (Ethernet0/1), from Connected, Send flag is 0x0
   Composite metric is (281600/0), Route is Internal
   Vector metric:
     Minimum bandwidth is 10000 Kbit
     Total delay is 1000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 0
```

```

10.1.2.3 (Ethernet0/0), from 10.1.2.3, Send flag is 0x0
Composite metric is (307200/281600), Route is Internal
Vector metric:
  Minimum bandwidth is 10000 Kbit
  Total delay is 2000 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 1

```

3. On R2, confirm that the route through R3 (10.1.2.3) is once again the only best path.

```
R2# show ip route
```

```

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, * - candidate default
       U - per-user static route, o - ODR

```

```
Gateway of last resort is not set
```

```

      10.0.0.0/24 is subnetted, 3 subnets
D       10.1.3.0 [90/307200] via 10.1.2.3, 00:00:20, Ethernet0
C       10.1.2.0 is directly connected, Ethernet0
C       10.1.1.0 is directly connected, Serial0.101

```

The EIGRP topology table reflects the increase in delay from R4 (10.1.2.4).

R4 feasible distance (281600) + R4 offset-list (20) = R4 reported distance (281620).

Note: A cosmetic defect in Cisco IOS Software Release 12.0(7) prevents the increased delay from being accurately reflected in the Total Delay section of the output shown here.

DDTS	Description
CSCdp36097 (registered customers only)	EIGRP: offset-list adds incorrect delay value

```
R2# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 307200
  Routing Descriptor Blocks:
  10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
    Composite metric is (307200/281600), Route is Internal
    Vector metric:
      Minimum bandwidth is 10000 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
  10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
    Composite metric is (307220/281620), Route is Internal
    Vector metric:
      Minimum bandwidth is 10000 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1

```

Change the Administrative Distance at R2

It is also possible to alter the path selection process by changing the administrative distance, on R2, of the route learned from R4. This method is less ideal than the others. It can increase the potential for routing loops for the reasons listed here:

- Administrative distance is generally used to determine the method by which a route was learned. If set incorrectly, the individual router cannot choose a redistributed route instead of the actual best path.
- Administrative distance is not propagated to other routers. Routing protocols rely on the fact that all routers choose the same path given the same set of parameters. Altering parameters on a single router can lead to routing loops.

1. Change R2's configuration so that when a routing update from R4 (10.1.2.4) for subnet 10.1.3.0/24 is heard, the administrative distance is increased to 91.

91 is chosen because it is 1 greater than the default EIGRP administrative distance for internals (which is 90). The default administrative distance for EIGRP externals (routes redistributed into EIGRP) is 170. Refer to the Default Distance Value Table in the What Is Administrative Distance? document for the default values of all routing protocols.

```
R2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# access-list 99 permit 10.1.3.0 0.0.0.255
R2(config)# router eigrp 1
R2(config-router)# distance 91 10.1.2.4 0.0.0.0 99
R2(config-router)# end
R2#
```

2. At this point, you may need to issue the **clear ip route** command for the changes to take effect.

Note: There is now only one path to 10.1.3.0/24 through R3 (10.1.2.3).

```
R2# show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - ISIS, L1 - ISIS level-1, L2 - ISIS level-2, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

      10.0.0.0/24 is subnetted, 3 subnets
D       10.1.3.0 [90/307200] via 10.1.2.3, 00:05:28, Ethernet0
C       10.1.2.0 is directly connected, Ethernet0
C       10.1.1.0 is directly connected, Serial0.101
```

Note: Nothing in the EIGRP topology table has changed.

```
R2# show ip eigrp topology 10.1.3.0 255.255.255.0
IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 307200
  Routing Descriptor Blocks:
```

```

10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
  Composite metric is (307200/281600), Route is Internal
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 2000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
  Composite metric is (307200/281600), Route is Internal
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 2000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1

```

Potential Problems

To illustrate a potential problem that this method can cause if not used carefully, imagine that R1 and R2 are running Open Shortest Path First (OSPF), with an administrative distance 110, for the 11.0.0.0/8 network. Imagine also that R4 has a static route for 11.1.1.0/24 that points at R2 (10.1.2.2). R4 is redistributing static routes into EIGRP so that some new routers on 10.1.3.0/24 can get to 11.1.1.0/24.

Normally, R2 receives the EIGRP external route for 11.1.1.0/24 from R4 with an administrative distance of 170. Because this distance is greater than that of the OSPF route (110), it is not installed.

This output is an example of the distance command used above when it is improperly configured.

```

R2# configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.

R2(config)# access-list 99 permit 11.1.1.0 0.0.0.255

R2(config)# router eigrp 1

R2(config-router)# distance 91 10.1.2.4 0.0.0.0 99

R2(config-router)# end

R2#

```

This configuration creates a routing loop between R2 and R4 for the subnet 11.1.1.0/24. R2 now prefers the 11.1.1.0/24 route advertised by R4. This is because the administrative distance (91) is lower than the administrative distance for the OSPF route (110).

Change the Bandwidth at R2

Using the bandwidth to influence EIGRP paths is discouraged for two reasons:

- Changing the bandwidth can have impact beyond affecting the EIGRP metrics. For example, quality of service (QoS) also looks at the bandwidth on an interface.
- EIGRP throttles to use 50 percent of the configured bandwidth. Lowering the bandwidth can cause problems like staving EIGRP neighbors from getting hello packets because of the throttling back.

Changing the delay does not impact other protocols nor does it cause EIGRP to throttle back.

1. Check the EIGRP topology table for R1 before you make any changes.

```
R1# show ip eigrp topology 10.1.3.0 255.255.255.0

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2221056
  Routing Descriptor Blocks:
    10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
      Composite metric is (2221056/307200), Route is Internal
  Vector metric:
    Minimum bandwidth is 1544 Kbit
    Total delay is 22000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2
```

2. Check the starting values for the ethernet0 interface on R2.

```
R2# show interface ethernet0

Ethernet0 is up, line protocol is up
  Hardware is Lance, address is 0010.7b3c.6786 (bia 0010.7b3c.6786)
  Internet address is 10.1.2.2/24
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:01, output 00:00:02, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  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
    1938 packets input, 165094 bytes, 0 no buffer
    Received 1919 broadcasts, 0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    0 input packets with dribble condition detected
    1482 packets output, 124222 bytes, 0 underruns
    0 output errors, 0 collisions, 18 interface resets
    0 babbles, 0 late collision, 0 deferred
    0 lost carrier, 0 no carrier
    0 output buffer failures, 0 output buffers swapped out
```

3. Decrease the bandwidth to see the impact on R1.

```
R2# configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.

R2(config)# interface ethernet0

R2(config-if)# bandwidth 5000

R2(config-if)# end

R2#
```

4. Confirm the change.

```
R2# show interface ethernet0

Ethernet0 is up, line protocol is up
  Hardware is Lance, address is 0010.7b3c.6786 (bia 0010.7b3c.6786)
  Internet address is 10.1.2.2/24
  MTU 1500 bytes, BW 5000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
```

```

Last input 00:00:02, output 00:00:01, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
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
 1995 packets input, 169919 bytes, 0 no buffer
  Received 1969 broadcasts, 0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  0 input packets with dribble condition detected
 1525 packets output, 127831 bytes, 0 underruns
  0 output errors, 0 collisions, 18 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier
  0 output buffer failures, 0 output buffers swapped out

```

5. Confirm that it also changed in the EIGRP topology table.

```
R2# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 2 Successor(s), FD is 563200
Routing Descriptor Blocks:
 10.1.2.4 (Ethernet0), from 10.1.2.4, Send flag is 0x0
   Composite metric is (563200/281600), Route is Internal
   Vector metric:
     Minimum bandwidth is 5000 Kbit
     Total delay is 2000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 1
 10.1.2.3 (Ethernet0), from 10.1.2.3, Send flag is 0x0
   Composite metric is (563200/281600), Route is Internal
   Vector metric:
     Minimum bandwidth is 5000 Kbit
     Total delay is 2000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 1

```

6. Check the impact on the EIGRP topology table for R1.

```
R1# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```

IP-EIGRP (AS 1): topology entry for 10.1.3.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2221056
Routing Descriptor Blocks:
 10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
   Composite metric is (2221056/563200), Route is Internal
   Vector metric:
     Minimum bandwidth is 1544 Kbit
     Total delay is 22000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 2

```

There is no change, as the Frame Relay connection between R1 and R2 is still the lowest speed link. You see a change only if you decrease that bandwidth on the ethernet0 interface for R2 to a value less than 1544.

7. Decrease the bandwidth on the ethernet0 interface for R2 to 1000.

```
R2# configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
R2(config)# interface ethernet 0
```

```
R2(config-if)# bandwidth 1000
```

```
R2(config-if)# end
```

```
R2#
```

8. Check the impact on the EIGRP topology table for R1.

```
R1# show ip eigrp topology 10.1.3.0 255.255.255.0
```

```
IP-EIGRP (AS 1): Topology entry for 10.1.3.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 312320
  Routing Descriptor Blocks:
  10.1.1.2 (Serial0.201), from 10.1.1.2, Send flag is 0x0
    Composite metric is (3123200/2611200), Route is Internal
  Vector metric:
    Minimum bandwidth is 1000 Kbit
    Total delay is 22000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2
```

NetPro Discussion Forums – Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

NetPro Discussion Forums – Featured Conversations for RP
Service Providers: MPLS
Virtual Private Networks: Services
Virtual Private Networks: Security

Related Information

- [EIGRP \(Enhanced Interior Gateway Routing Protocol\) Support Page](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 13673