通过改变 EIGRP 度量值设置优先路由

目录

简介

本文档介绍了如何通过影响增强型内部网关路由协议 (EIGRP) 度量来创建首选路由。基于网络图中显示的拓扑，本文档介绍了几种方法用于影响从客户端传输到服务器的 IP 数据流，其中首选路径为 R1>R2>R3。目标是使路径 R1>R2>R4 成为仅在 R3 上出现故障时使用的备用路径。

先决条件

要求

本文档要求对 IP 路由和 EIGRP 路由有基本的了解。有关 IP 路由和 EIGRP 的详细信息，请参阅以下文档：

- 路由基本知识
- 增强的内部网关路由协议

使用的组件

本文档中的信息基于以下软件和硬件版本。

- Cisco IOS® 软件版本 9.21 及更高版本支持 EIGRP。本文档中的信息基于 Cisco IOS 软件版本

12.3(3)。

- 在所有的路由器（例如 Cisco 2500 系列和 Cisco 2600 系列）以及所有的第 3 层交换机上都可以配置 EIGRP。

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

规则

关于文件规则的更多信息，参考。

背景信息

有多种方法可用于通过影响 EIGRP 度量来设置首选路由。本文档介绍了这些方法，并详述了各自的优缺点。本文档还讨论了修改带宽所产生的影响，不过在本例中修改带宽并不是改变路径的可行手段。

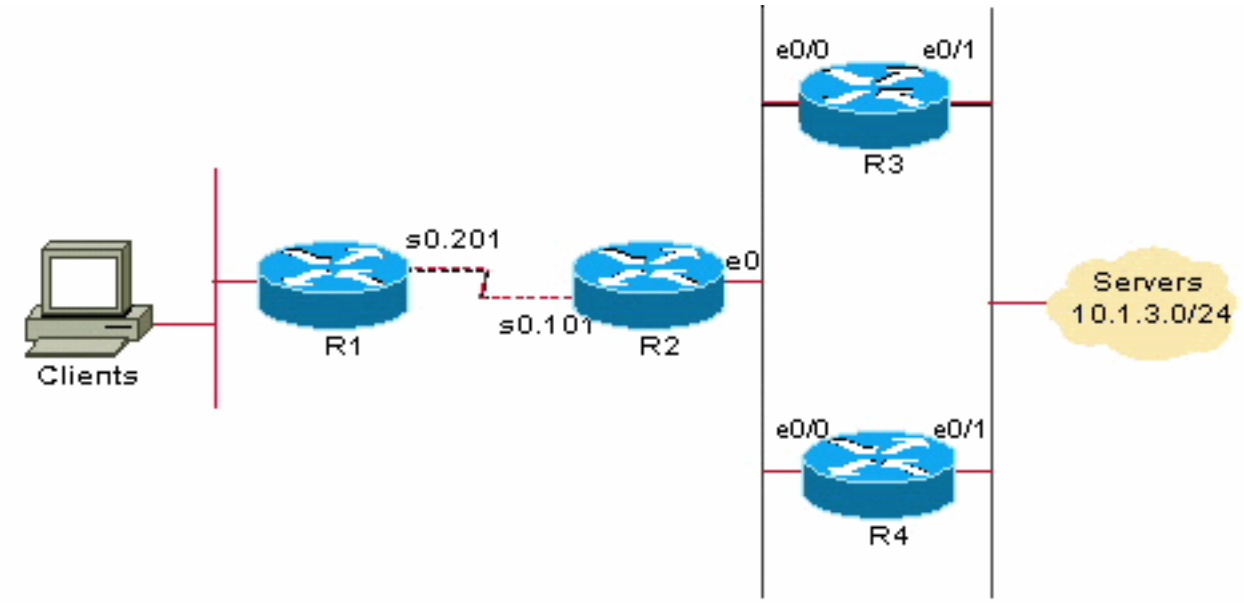

单击网络图可在单独的浏览器窗口中进行显示，以便在本文档稍后部分用作参考。

本文档中用于验证 EIGRP 行为的两条命令为 show ip eigrp topology 和 show ip eigrp topology network-ip subnet-mask。

如果有输出－**show ip eigrp topology命令**或－**show ip eigrp topology network-ip subnet-mask命令**从您的Cisco设备，您能使用Cisco CLI分析器 (仅限注册用户)显示潜在问题和修正。要使用Cisco CLI分析器 (仅限注册用户)，您必须登陆和有在您的Web浏览器启用的Javascript。

背景 - EIGRP 度量基本知识

EIGRP 更新包含五个度量：最低带宽、延迟、负载、可靠性和最大传输单元 (MTU)。默认情况下，这五个度量中只有最低带宽和延迟用于计算最佳路径。与多数度量都不相同，最低带宽设置为整个路径的最低带宽，并不反映路径中的跳数或低带宽链路数。延迟是一种累计值，会按照路径中各网段延迟值的增加而增加。有关 EIGRP 度量的详细信息，请参阅增强型内部网关路由协议白皮书。

可能的配置

这些配置可用于设置首选路由。

默认负载分担配置

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
```

注意： 帧中继交换机在网络图中已隐藏。

```
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
```

注意： R2 有两个指向 10.1.3.0/24 的等价路径，分别经过 R3 (10.1.2.3) 和 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
```

注意： 两个路径的复合度量（距离/报告距离）相同。通往 R1 的可行距离 (FD) 会进行通告，随后就成为 R1 的报告距离。

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
```

更改 R4 的接口延迟参数

由于对延迟度量所做的更改会传播至所有的下行路由器，因此对于以下两种情况，更改接口延迟参数是影响路径选择的首选方法：

- 以太网段 10.1.3.0/24 仅包含服务器，且 10.1.3.0/24 子网后无其他子网。（此配置是服务器群的理想设置。）
- 您希望针对所有通过 10.1.3.0/24 网段上的 EIGRP 邻居获知的路由影响路径选择。

1. 在进行任何更改前，请检查接口上的延迟。其当前设置与 R3 相同，如此处所示。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. 更改 10.1.3.0/24 网段上的延迟值。选择新的延迟时，一定要谨慎。您肯定不希望延迟的增加导致 R2 不再将该路由视为可行后继路由。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. 确认此接口的延迟已更改为 1200 微秒。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
4. 确认 R2 只有一个“最佳”路由可以到达 10.1.3.0，也就是经过 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 **show ip eigrp topology** 命令显示 R4 通告的延迟度量已按预期方式增加了 200（增至 2200 微秒）。这种延迟的增加使得两个路由具有不同的成本，并使 R2 无法执行负载均衡。**注意：** 由于 R4 通告的距离 (286720) 小于 R2 通告的距离（可行距离 307200），因此该路径被视为无环路路径。由于 R4 通告的路径被视为无环路路径，因此是一种可行后继路由，可在 R3 停止通告指向 10.1.3.0/24 的路由时立即安装。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 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
```

在 R4 上使用 offset-list 修改 R2 的复合度量

可通过在 R4 路由器上使用 offset-list 来修改 R2 的复合度量。如果 R4 上的 offset-list 值为 20，则 R2 上 R2-R4 路径的复合度量将增加 20。于是，R2-R4 路径成为相对于 R2-R3 的备用路径。对于以下情况，使用 offset-list 是首选方法：

- 您只希望影响正在通告的特定路径。
- 有其他一些路由器连接至 10.1.3.0/24 子网，而您不希望影响由这些路由器发起的路径。

1. 对 R4 上的 offset-list 进行配置，使所有以 **10.1.3.x** 开头的路由的延迟都增加 (20)。`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. 您在offset-list在R4的EIGRP拓扑表里不更改什么的此输出中能看到。只有在通告了路由的情况下，度量才会发生更改。`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. 在 R2 上，确认经过 R3 (10.1.2.3) 的路由还是唯一的最佳路径。`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` EIGRP 拓扑表反映了 R4 (10.1.2.4) 中延迟的增加情况。R4 可行距离 (281600) + R4 offset-list (20) = R4 报告距离 (281620)。**注意：** Cisco IOS 软件版本 12.0(7) 的表面缺陷使得增加的延迟无法精确反映在此处所示输出内容的 Total Delay 部分中。`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`

更改 R2 的管理距离

也可以在 R2 上更改从 R4 所获知路由的管理距离，以此更改路径选择过程。相比其他方法，这种方法不太理想。由于下述原因，这种方法会增加出现路由环路的可能性：

- 管理距离通常用于确定获知路由的方法。如果设置有误，单个路由器将无法选择重分配的路由以代替实际最佳路径。
- 管理距离不会传播至其他路由器。路由协议会取决于这样一种事实：在采用同一组参数时，所

有路由器都会选择同一种路径。更改单个路由器上的参数会导致路由环路。

1. 更改 R2 的配置，使得在接收到 R4 (10.1.2.4) 发出的子网 10.1.3.0/24 路由更新时，管理距离将增至 91。选择 91 是因为它比内部路由的默认 EIGRP 管理距离 (90) 大 1。EIGRP 外部路由（重分配到 EIGRP 中的路由）的默认管理距离为 170。参考在默认值的本文所有路由协议。 `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. 此时，您可能需要发出 clear ip route 命令，以便让所做的更改生效。**注意：** 现在只有一个路径可以到达 10.1.3.0/24，也就是经过 R3 (10.1.2.3) 的路径。`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` **注意：** EIGRP 拓扑表的内容没有任何更改。`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`

潜在问题

如果使用不当，此方法会导致一种潜在问题；为了说明这一点，请设想这种情形：在 11.0.0.0/8 网络中，R1 和 R2 正在运行开放最短路径优先 (OSPF) 协议，管理距离为 110。再设想对于 11.1.1.0/24，R4 有一个静态路由指向 R2 (10.1.2.2)。R4 正在将静态路由重分配到 EIGRP 中，以使 10.1.3.0/24 上的一些新路由器能到达 11.1.1.0/24。

通常，R2 从 R4 接收 11.1.1.0/24 的 EIGRP 外部路由，管理距离为 170。由于此距离大于 OSPF 路由距离 (110)，因此不予安装。

此输出是上面所用的 distance 命令的错误配置示例。

`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#`

这种配置为子网 11.1.1.0/24 在 R2 和 R4 之间创建了一个路由环路。现在，R4 通告的 11.1.1.0/24 路由是 R2 的首选路由。这是因为此管理距离 (91) 小于 OSPF 路由的管理距离 (110)。

更改 R2 的带宽

不鼓励使用带宽来影响 EIGRP 路径，原因有两个：

- 除影响 EIGRP 度量外，更改带宽还会造成其他影响。例如，服务质量 (QoS) 也和接口带宽有关。
- EIGRP 降至使用一半的配置带宽。降低带宽会引起一些问题，例如由于降低了流量导致 EIGRP 邻居无法获得 hello 数据包。

更改延迟不会影响其他协议，也不会导致 EIGRP 流量降低。

1. 在进行任何更改前，请检查 R1 的 EIGRP 拓扑表。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. 检查 R2 上 ethernet0 接口的起始值。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. 降低带宽，以观察对 R1 产生的影响。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. 确认更改。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. 确认 EIGRP 拓扑表中也有变化。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. 查看对 R1 中的 EIGRP 拓扑表产生的影响。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 无变化，这是因为 R1 和 R2 之间的帧中继连接仍是最低速的链路。只有将 R2 的 ethernet0 接口带宽降至 1544 以下，您才能看到变化。
7. 将 R2 的 ethernet0 接口带宽降至 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. 查看对 R1 中的 EIGRP 拓扑表产生的影响。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
```

相关信息

- EIGRP（增强型内部网关路由协议）支持页
- 技术支持 - Cisco Systems