CCIE® Professional Development

Troubleshooting IP Routing Protocols

The comprehensive, hands-on guide for resolving IP routing problems

Zaheer Aziz, CCIE No. 4127
Johnson Liu, CCIE No. 2637
Abe Martey, CCIE No. 2373
Faraz Shamim, CCIE No. 4131

ciscopress.com

Troubleshooting IP Routing Protocols

Zaheer Aziz, CCIE #4127
Johnson Liu, CCIE #2637
Abe Martey, CCIE #2373
Faraz Shamim, CCIE #4131

Cisco Press

Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

In Figure 7-6, Router A and Router B have a primary address in the 10.1.1.0/24 network range, while Router C has an address range of 50.1.1.0/24 configured. When Router A or Router B sends out the EIGRP hello packet, the source of the hello packet will be either 10.1.1.1 or 10.1.1.2, depending on which router sends out the hello. When Router C receives the hello packet from Router A or Router B, it notices that the source is from the 10.1.1.0 network. Because Router C has an IP address of 50.1.1.3 configured on the interface, Router C will not process the hello packet from Router A or Router B because they are from a different network. Therefore, no neighbor relationship is formed from Router C to either Router A or Router B.

The solution for this example is to match all the IP addresses on the segment to the primary address space. For the network in Figure 7-6, you need to configure Router C to be in the primary address space of 10.1.1.0/24.

Switch or Hub Between EIGRP Neighbor Connection Is Misconfigured or Is Leaking Multicast Packets to Other Ports

If the IP address configuration is correct on the interface between EIGRP neighbors, you might want to check the configuration on the switch or the hub that connects the EIGRP neighbors. If a single LAN hub connects the EIGRP neighbors for different LAN segment, the hub passes broadcast and multicast packets to other ports between two logical LAN segments. So, the multicast EIGRP hello from LAN segment 1 will be seen on the neighbor located in LAN segment 2 if a single hub connects all the LAN devices on different LAN segments. The solution is to break up the broadcast domain by using a separate hub for each LAN segment or simply configuring **no eigrp log-neighbor-warnings** under EIGRP configuration to stop seeing the error message.

If a LAN switch connects the LAN devices, you might want to check the configuration of the switch. Make sure that the switch is not configured so that different LAN segments reside within the same VLAN. Make sure that the switch is configured so that each LAN segment has its own broadcast domain and does not share its broadcast domain with other LAN segments.

EIGRP Neighbor Problem—Cause: Mismatched Masks

Sometimes, a simple misconfiguration on the interface subnet mask causes an EIGRP neighbor problem. Figure 7-7 illustrates a network diagram for such a scenario.

Figure 7-7 *Network Topology Vulnerable to EIGRP Neighbor Problems Because of Mismatched Masks*

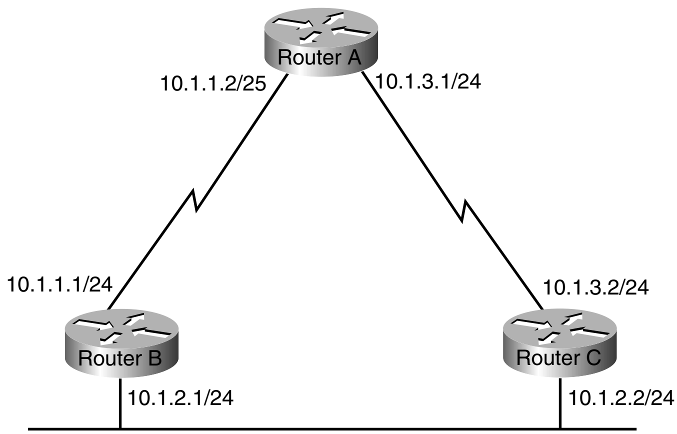

Example 7-3 shows the configuration for Routers A, B, and C.

Example 7-3 *Router A, B, and C Configurations for the Network in Figure 7-7*

```
Router A#interface  serial 0
ip address 10.1.1.2 255.255.255.128
interface  serial 1
ip address 10.1.3.1 255.255.255.0
```

```
Router B#interface  serial 0
ip address 10.1.1.1 255.255.255.0
interface ethernet 0
ip address 10.1.2.1 255.255.255.0
```

```
Router C#interface ethernet 0
ip address 10.1.2.2 255.255.255.0
interface  serial 0
ip address 10.1.3.2 255.255.255.0
```

Notice the mismatched mask on the serial interface of Router A and Router B. Router A has a mask of 255.255.255.128, while Router B has a mask of 255.255.255.0 on Serial 0. Initially, EIGRP has no problem forming the neighbor between Router A and Router B because 10.1.1.1 and 10.1.1.2 are in the same subnet. The problem occurs when a neighbor relationship is established and Router A and Router B begin to exchange EIGRP topology tables and install routes based on the EIGRP topology table, as demonstrated in Example 7-4.

Example 7-4 *Routing Tables from Router B and Router C*

```
Router B#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 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set
C 10.1.1.0/24  Serial 0
D 10.1.1.0/25  10.1.2.2
```

```
Router c#show ip route eigrp
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 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set
D 10.1.1.0/24  10.1.2.1
D 10.1.1.0/25  10.1.3.1
```

When Router B sends Router A an EIGRP update, Router A responds to the update with an EIGRP acknowledgement packet with a destination address of 10.1.1.1 to Router B. When Router B receives the packet, it forwards the ACK packet to Router C instead of processing it because Router B has a more specific route from Router C. Router B has a more specific route of 10.1.1.0/25 with the next hop to 10.1.2.2. This /25 route overrides the /24 route because /25 is more specific than /24. When Router C receives the ACK packet from Router B, it looks at its routing table for the 10.1.1.1 entry, and the routing table points to Router A. Router C then forwards the ACK packet back to Router A. This creates a routing loop. The packet to 10.1.1.1 loops from Router A to Router B, from Router B to Router C, and back from Router C to Router A. As a result, Router B won't process the ACK packet from Router A; Router B will think that Router A never ACK'ed the update packet, and Router B will reset the neighbor after 16 retries.

The solution for this problem: Configure the right subnet mask on Router A's Serial 0 interface to 255.255.255.0.

EIGRP Neighbor Problem—Cause: Mismatched K Values

For EIGRP to establish its neighbors, the K constant value to manipulate the EIGRP metric must be the same. Refer to Chapter 6 for an explanation of the K values. In EIGRP's metric calculation, the default for the K value is set so that only the bandwidth and the delay of the interface are used to calculate the EIGRP metric. Many times, the network administrator might want other interface factors, such as load and reliability, to determine the EIGRP metric. Therefore, the K values are changed. Because only bandwidth and delay are used in calculations, the remaining K values are set to a value of 0 by default. However, the K values must be the same for all the routers, or EIGRP won't establish a neighbor relationship. Figure 7-8 shows an example of this case.

Figure 7-8 *Network Vulnerable to EIGRP Neighbor Problems Because of Mismatched K Values*

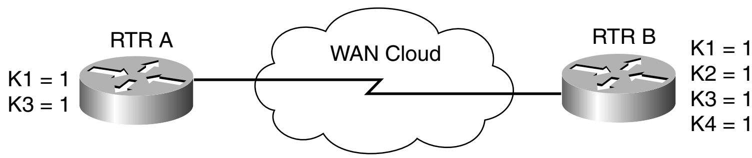

For the network in Figure 7-8, K1 is bandwidth and K3 is delay. The network administrator changed the K values of RTR B to all 1s from K1 to K4, while RTR A retains the default value of K1 and K3 to be 1. In this example, RTR A and RTR B will not form EIGRP neighbor relationship because the K values don't match. Example 7-5 shows the configuration for RTR B.

Example 7-5 *Configuration for RTR B in Figure 7-8*

```
RTR B#router eigrp 1
network xxxx
metric weights 0 1 1 1 1 0
```

RTR B's configuration includes the extra **metric weights** command. The first number is the type of service (ToS) number, which, because it's not supported, gets a value of 0. The five numbers after the ToS are the K1 through K5 values.

Troubleshooting this problem requires careful scrutiny of the router's configuration. The solution for this problem is to change all the K values to be the same on all the neighboring routers. In this example, in Router A, changing the K values to match the K value of Router B will solve the problem, as demonstrated in Example 7-6.

Example 7-6 *Configuring the K Values on Router A to Match Router B*

```
RTR A#router eigrp 1
network xxxx
metric weights 0 1 1 1 1 0
```

EIGRP Neighbor Problem—Cause: Mismatched AS Number

EIGRP won't form any neighbor relationships with neighbors in different autonomous systems. If the AS numbers are mismatched, no adjacency is formed. This problem is usually caused by misconfiguration on the routers. Figure 7-9 illustrates such a problem.

Figure 7-9 *Network Experiencing an EIGRP Neighbor Problem Because Mismatched AS Numbers*

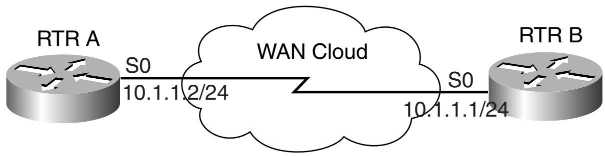

In the network shown in Figure 7-9, RTR A and RTR B are in the EIGRP AS number of 1 and the proper network numbers have been configured; however, no EIGRP neighbor relationship is formed between RTR A and RTR B. Begin by checking the configuration of RTR A and RTR B in Example 7-7.

Example 7-7 *Configurations for RTR A and RTR B in Figure 7-9*

```
RTR B#show running-config
interface serial 0
IP address 10.1.1.1 255.255.255.0
router eigrp 11
network 10.0.0.0
```

```
RTR A#show running-config
Interface serial 0
IP address 10.1.1.2 255.255.255.0
router eigrp 1
network 10.0.0.0
```

You should notice the misconfiguration immediately. RTR B's Serial 0 interface is configured to be in EIGRP AS number 11, while RTR A's Serial 0 is configured to be in EIGRP AS number 1. Because the AS numbers don't match across the link, no EIGRP neighbor relationship will be formed. To resolve this problem, simply configure both routers with the same EIGRP AS number, as shown in Example 7-8. In this example, both routers will be configured to be in EIGRP AS 1.

Example 7-8 *Configuring Both Routers with the Same EIGRP AS Numbers*

```
RTR A#router eigrp 1
network 10.0.0.0
```

```
RTR B#router eigrp 1
network 10.0.0.0
```

EIGRP Neighbor Problem—Cause: Stuck in Active

Sometimes, EIGRP resets the neighbor relationship because of a "stuck in active" condition. The error message is

```
%DUAL-3-SIA: Route network mask stuck-in-active state in IP-EIGRP AS. Cleaning up
```

This section discusses the method of troubleshooting the EIGRP stuck in active error.

Reviewing the EIGRP DUAL Process

To resolve an EIGRP stuck in active error, you need to understand the DUAL process in EIGRP. Refer to Chapter 6 for thorough coverage of the DUAL process, although it is reviewed here as well.

EIGRP is an advanced distance-vector protocol; it doesn't have LSA flooding, like OSPF, or a link-state protocol to tell the protocol the overall view of the network. EIGRP relies only on its neighbors for information on network reachability and availability. EIGRP keeps a list of backup routes called *feasible successors*. When the primary route is not available, EIGRP immediately uses the feasible successor as the backup route. This shortens convergence time. Now, if the primary route is gone and no feasible successor is available, the route is in active state. The only way for EIGRP to converge quickly is to query its neighbors about the unavailable route. If the neighbor doesn't know the status of the route, the neighbor asks its neighbors, and so on, until the edge of the network is reached. The query stops if one of the following occurs:

- All queries are answered from all the neighbors.
- The end of network is reached.
- The lost route is unknown to the neighbors.

The problem is that, if there are no query boundaries, EIGRP potentially can ask every router in the network for a lost route. When EIGRP first queries its neighbor, a stuck in active timer starts. By default, the timer is three minutes. If, in three minutes, EIGRP doesn't receive the query response from all its neighbors, EIGRP declares that the route is stuck in active state and resets the neighbor that has not responded to the query. Figure 7-10 illustrates the query process of EIGRP when a route is lost.