

This document about Cisco CCNP® Routing & Switching certification is designed exclusively for the Cisco Learning Network (CLN). The goal is to provide one example of troubleshooting the Enhanced Interior Gateway Routing Protocol (EIGRP)—particularly, how easily overlooked requirements for EIGRP neighbor establishment can create difficult troubleshooting exercises.

Background

In terms of EIGRP neighbor discovery and establishment, dependencies on the media (LAN, slow-speed WAN, and high-speed WAN) are used. This document discusses only high-speed (higher than T1) WAN and LAN connections. It also does not address EIGRP authentication.

EIGRP neighbor discovery is accomplished using a periodic (every five seconds) hello packet that is sent to the multicast address 224.0.0.10; these hello packets are *not* acknowledged on multi-access and high-speed WAN networks. Rather, an EIGRP peer responds to an initial hello packet from a potentially new neighbor by sending unicast update packets (with an initialization bit set to 1), with its EIGRP topology table (the information needed to build a set of distances and vectors to each reachable network); these update packets *must* be acknowledged. The remote reacts similarly, sending unicast update packets with its EIGRP topology table. When it is acknowledged, the two devices have formed an adjacency.

At times you may see messages that say “EIGRP: Neighbor not yet found.” This indicates that a device has received an update packet without the initialization bit set, and this device has not yet received a hello packet or initialized this neighbor. This update packet will be dropped. Similarly, if a device receives an update packet with the initialization bit set from a neighbor that is currently active, it will reset the neighbor and log a “peer restarted” message. In either case, the explanation for the unexpected update packet is unknown; it could be due to a variety of reasons—to include software bugs, queuing (either input or output processing) anomalies, and so on.

It is possible to build a “one-way” EIGRP neighbor relationship if a link is passing traffic in only one direction. If, for some reason, device A could not transmit packets but device B could, then device A would show device B as a neighbor intermittently. Syslog or console messages would show **retry limit exceeded** errors, and **show ip eigrp neighbor** would show high Q counts (the number of EIGRP packets that the software is waiting to send).

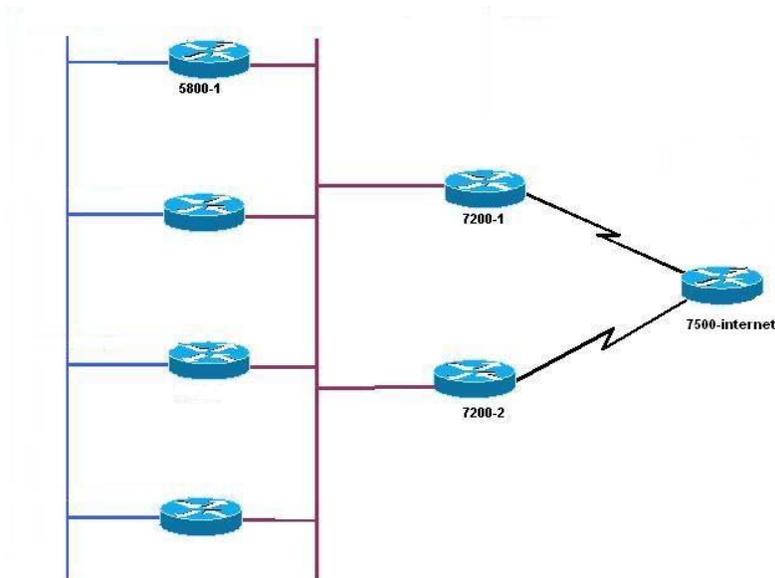
The hello packet contains information about the device that generated the packet:

- **Hold time:** By default, this is three times the hello interval and is the time that the receiving device should expect to wait before declaring the neighbor unreachable. EIGRP peers can form an adjacency even when the received hello time and hold time do not match the local settings (this unidirectional independence is facilitated by including the hold time in the hello packets).
- **K metric values:** EIGRP does *not* form an adjacency if the K values are mismatched between peers.
- **Autonomous system (AS) number:** EIGRP does *not* form an adjacency if the AS numbers are mismatched between peers.

Scenario

In this document, you will work with the network topology shown in the following figure. All of the routers in this network are running EIGRP in AS 212. The 7500-internet router is sending a default network (20.0.0.0) for access to the Internet. The network topology is designed to incorporate load balancing between routers 7200-1 and 7200-2, as well as for redundancy and backup. However, while network access is not impaired, there is concern that load balancing is not being accomplished with the current configurations. The goal is to identify and correct the reason that load balancing might be failing.

In this document, you will work with only one downstream router: 5800-1.



Configuration

The relevant configurations for the three routers are as follows:

7500-internet

```
router eigrp 212
 redistribute static metric 9600 10 200 200 1500
 network 20.0.0.0 0.0.0.255
 network 192.168.1.0
 no auto-summary
 no eigrp log-neighbor-changes
!
ip default-network 20.0.0.0
```

7500-internet#**show ip interfaces brief**

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	20.0.0.1	YES	NVRAM	up	up
Serial1/0	192.168.1.1	YES	manual	up	up
Serial2/0	192.168.1.10	YES	manual	up	up

7200-1

```
router eigrp 212
 network 192.168.1.0
 network 212.50.185.96 0.0.0.31
 no auto-summary
 no eigrp log-neighbor-changes
```

7200-1#**show ip interfaces brief**

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	212.50.185.125		YES NVRAM	up	up
Serial1/0	192.168.1.9	YES	manual	up	up

7200-2

```
router eigrp 212
 network 192.168.1.0
 network 212.50.185.96 0.0.0.31
 no auto-summary
 no eigrp log-neighbor-changes
```

7200-2#**show ip interfaces brief**

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	212.50.185.126		YES NVRAM	up	up
Serial1/0	192.168.1.2	YES	manual	up	up

Verification

The first clear indication that load balancing across the 7200s is not occurring is from the **show ip route** output from 5800-1.

5800-1#**show ip route**

```
Codes: 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
 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
 o - ODR, P - periodic downloaded static route
```

Gateway of last resort is 212.50.185.126 to network 20.0.0.0

```
D    212.50.167.0/24 [90/160000] via 212.50.185.82, 00:32:16, Ethernet1/0
    212.50.166.0/24 is variably subnetted, 4 subnets, 2 masks
D    212.50.166.0/24 is a summary, 00:32:19, Null0
C    212.50.166.1/32 is directly connected, Loopback1
C    212.50.166.2/32 is directly connected, Loopback2
C    212.50.166.20/32 is directly connected, Loopback20
    20.0.0.0/24 is subnetted, 1 subnets
D*   20.0.0.0 [90/2200576] via 212.50.185.126, 00:00:04, Ethernet0/0
    212.50.185.0/27 is subnetted, 3 subnets
C    212.50.185.64 is directly connected, Ethernet1/0
C    212.50.185.96 is directly connected, Ethernet0/0
C    212.50.185.32 is directly connected, Ethernet2/0
    192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks
D    192.168.1.8/30 [90/2174976] via 212.50.185.126, 00:21:17, Ethernet0/0
D    192.168.1.0/29 [90/2174976] via 212.50.185.125, 00:00:06, Ethernet0/0
D    192.168.1.4/30 [90/2686976] via 212.50.185.126, 00:00:05, Ethernet0/0
D*EX 0.0.0.0/0 [170/2177536] via 212.50.185.126, 00:00:06, Ethernet0/0
```

```
5800-1#show ip route 20.0.0.0
```

```
Routing entry for 20.0.0.0/24, 1 known subnets
  Redistributing via eigrp 212
```

```
D*    20.0.0.0 [90/2200576] via 212.50.185.126, 00:14:35, Ethernet0/0
```

Ideally, the 5800 should have two routes to the default-network: 7200-1 (212.50.126.125) and 7200-2 (212.50.126.126).

Similarly, 7200-1 indicates the following routing information:

```
7200-1#show ip route
```

```
Codes: 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
        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
        o - ODR, P - periodic downloaded static route
```

Gateway of last resort is 212.50.185.126 to network 20.0.0.0.

```
D    212.50.167.0/24 [90/158720] via 212.50.185.98, 10:46:30, Ethernet0/0
      [90/158720] via 212.50.185.97, 10:46:30, Ethernet0/0
D    212.50.166.0/24 [90/156160] via 212.50.185.101, 10:46:30, Ethernet0/0
      20.0.0.0/24 is subnetted, 1 subnets
D*   20.0.0.0 [90/2198016] via 212.50.185.126, 22:21:08, Ethernet0/0
      212.50.185.0/27 is subnetted, 3 subnets
D    212.50.185.64 [90/30720] via 212.50.185.98, 10:46:30, Ethernet0/0
      [90/30720] via 212.50.185.97, 10:46:30, Ethernet0/0
C    212.50.185.96 is directly connected, Ethernet0/0
D    212.50.185.32 [90/30720] via 212.50.185.98, 10:46:30, Ethernet0/0
      [90/30720] via 212.50.185.97, 10:46:30, Ethernet0/0
      192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks
D    192.168.1.8/30 [90/2172416] via 212.50.185.126, 22:21:09, Ethernet0/0
C    192.168.1.0/29 is directly connected, Serial1/0
D    192.168.1.4/30 [90/2684416] via 212.50.185.126, 22:21:09, Ethernet0/0
D*EX 0.0.0.0/0 [170/2174976] via 212.50.185.126, 22:21:10, Ethernet0/0
```

Again the route to the 20.0.0.0 default network is by way of the LAN rather than the WAN (which is the expectation; see Figure 1):

```
7200-1#show ip route 20.0.0.0
```

```
Routing entry for 20.0.0.0/24, 1 known subnets
  Redistributing via eigrp 212
```

```
D*    20.0.0.0 [90/2198016] via 212.50.185.126, 00:13:49, Ethernet0/0
```

While this data is sufficient to indicate that the serial interface between 7200-1 and 7500-internet is not being utilized for load-balancing purposes, a final data point would be the 7200-1 serial interface counters:

```
7200-1#show interfaces serial 1/0
```

```
Serial1/0 is up, line protocol is up
  Hardware is M4T
  Internet address is 192.168.1.5/29
  MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 11/255
  Encapsulation HDLC, crc 16, loopback not set
  Keepalive set (10 sec)
  Restart-Delay is 0 secs
  Last input 00:00:00, 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 70000 bits/sec, 84 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    204852 packets input, 21168726 bytes, 0 no buffer
    Received 683 broadcasts, 0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    3061 packets output, 178155 bytes, 0 underruns
    0 output errors, 0 collisions, 2 interface resets
```

```
1 unknown protocol drops
0 output buffer failures, 0 output buffers swapped out
3 carrier transitions      DCD=up  DSR=up  DTR=up  RTS=up  CTS=up
```

The problem—a routing issue that is preventing load balancing from occurring—has now been confirmed. In particular, the interface statistics (in [blue in the preceding list](#)) suggest that the problem is unidirectional; there are inbound packets to 7200-1 from 7500-internet, but 7200-1 is not sending comparable packet amounts out the same interface.

It is time to troubleshoot the load-balancing failure, with the suspicion that given the unexpected routing table entries and dissimilar packet counts on the 7200-1 serial interface, the problem most likely lies at the 7200-1.

Troubleshooting the Load-Balancing Failure

First, on the 7200-1 you will confirm that the EIGRP topology reflects what you have already seen in the IP routing table. Next, you will check to be sure that the EIGRP neighbors are as you would expect.

As anticipated, 20.0.0.0 points to the next-hop addresses out of the LAN interface, with no route learned from serial 1/0:

```
7200-1#show ip eigrp topology 20.0.0.0/24
IP-EIGRP (AS 212): Topology entry for 20.0.0.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2198016
  Routing Descriptor Blocks:
    212.50.185.126 (Ethernet0/0), from 212.50.185.126, Send flag is 0x0
      Composite metric is (2198016/2195456), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 21100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
      Exterior flag is set
```

The interfaces look OK from the perspective of EIGRP, although it is important to keep in mind one odd bit of data. The mean smooth round-trip time (SRTT) is the number of milliseconds required for an EIGRP packet to be sent to a neighbor and for the local router to receive an acknowledgment of that packet. For Se1/0 the SRTT is 0, which suggests a potential connectivity issue:

```
7200-1#show ip eigrp interface
IP-EIGRP interfaces for process 212
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Se1/0	1	0/0	0	0/15	50	0
Et0/0	4	0/0	12	0/1	52	0

You can see 192.168.1.6 (7500-internet) as a neighbor:

```

7200-1#show ip eigrp neighbor serial 1/0
IP-EIGRP neighbors for process 212
H   Address          Interface Hold Uptime      SRTT    RTO   Q     Seq
   (sec)    (ms)  Cnt   Num
4   192.168.1.6   Se1/0      14   00:01:04    1     5000  2     0

```

There are two unusual points in the preceding example. The first is the uptime for 192.168.1.6 (7500-internet); recall that the uptime is the length in time since the 7200-1 first heard from the 7500-internet neighbor; you would expect this value to be higher, and that 7200-1 and 7500-internet had established an adjacency longer than 1 minute ago.

Next, recall that the Q count is the number of EIGRP packets that the 7200-1 is waiting to send. Normally this would be 0, because the SRTT would be small enough that the odds of the IOS parser catching the brief moment when EIGRP packets are queued is quite small. However, with an SRTT of five seconds (5000ms) the likelihood of EIGRP packet congestion increases. Again, this suggests connectivity issues between 7200-1 and 7500-internet.

At this point, it is helpful to enable the logging of EIGRP neighbor changes on 7200-1 with the EIGRP command **eigrp log-neighbor-changes** (this command is enabled by default). After you do this, you see the following messages generated to the syslog and console on 7200-1:

```

*Nov  5 21:57:51.865: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 212: Neighbor 192.168.1.6
(Serial1/0) is up: new adjacency
*Nov  5 21:59:11.593: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 212: Neighbor 192.168.1.6
(Serial1/0) is down: retry limit exceeded

```

Recall from the "Background" section at the beginning of this document that the sequence for two neighbors to establish an adjacency is for one to send a multicast hello packet, which causes the receiver to react with a unicast update packet (which must be acknowledged) with the initialization bit set.

The new adjacency message is an indication that the 7200-1 has received a multicast hello packet from 192.168.1.6 (7500-internet). The "retry limit exceeded" message (80 seconds later) indicates that either the acknowledgments from 7500-internet for the update packets sent by 7200-1 have not been received, or the update packets never made it to the 7500-internet router in the first place.

But the important point is to differentiate between what works and what does not work: multicast traffic is OK (as evidenced by the new adjacency message), while unicast traffic is problematic (as evidenced by the "retry limit exceeded" message).

The next logical test is to ping 192.168.1.6 (the 7500-internet) from the 7200-1:

```

7200-1#ping 192.168.1.6

```

```

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/24/44 ms

```

Earlier you noticed that all of the routes on the 7200-1 were pointing out of the LAN interface; now you need to confirm that the route to 192.168.1.6 is indeed connected, versus learned from the LAN segment (in other words, confirm that the ICMP packets actually went out the serial interface):

```
7200-1#show ip route 192.168.1.6
Routing entry for 192.168.1.4/30
  Known via "eigrp 212", distance 90, metric 2684416, type internal
  Redistributing via eigrp 212
  Last update from 212.50.185.126 on Ethernet0/0, 1d01h ago
  Routing Descriptor Blocks:
  * 212.50.185.126, from 212.50.185.126, 1d01h ago, via Ethernet0/0
    Route metric is 2684416, traffic share count is 1
    Total delay is 40100 microseconds, minimum bandwidth is 1544 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2
```

You can see that 192.168.1.4/30 is being unexpectedly learned from the LAN interface. Notice the routing information for the entire 192.168.1.0 network:

```
7200-1#show ip route 192.168.1.0
Routing entry for 192.168.1.0/24, 3 known subnets
  Attached (1 connections)
  Variably subnetted with 2 masks
  Redistributing via eigrp 212

D       192.168.1.8/30 [90/2172416] via 212.50.185.126, 1d01h, Ethernet0/0
C       192.168.1.0/29 is directly connected, Serial1/0
D       192.168.1.4/30 [90/2684416] via 212.50.185.126, 1d01h, Ethernet0/0
```

The problem is now obvious: there is a mask mismatch between 7200-1 (/29) and 7500-internet (/30).

Explanation

Remember that a router chooses the longest match for a destination address when making a route determination. First, from a routing perspective, the following situation specific to the network 192.168.1.4/30 occurs:

1. The 7500-internet is directly connected to the network 192.168.1.4/30 by way of serial 2/0. 192.168.1.4/30 will not be advertised out serial 2/0 (due to split horizon), but will be advertised out serial 1/0 (to 7200-2) via EIGRP.
2. The 7200-2, after adding 192.168.1.4/30 to its routing table, advertises the route to 7200-1 by way of the LAN.
3. The 7200-1 adds the unique route 192.168.1.4/30 to its routing table; the connected route 192.168.1.0/29 is also in the routing table. The two routes coexist in the routing table of 7200-1, since the masks are unique.

Now from the EIGRP perspective, here is what happens (one of two scenarios):

1. The 7200-1 sends a multicast hello packet out serial 1/0, which is received by 7500-internet.
2. The 7500-internet responds with a unicast update packet to 7200-1 (192.168.1.5).
3. The 7200-1 receives the update packet from 7500-internet and generates an acknowledgment to send back to 7500-internet (192.168.1.6). Note that, as EIGRP neighbors are required to be directly connected, the acknowledgment has a hop count (the IP Time to Live field, or TTL) of 1.
4. The 7200-1 routing table causes the acknowledgment to be sent to 7200-2, because the longest route match for the destination address 192.168.1.6 is 192.168.1.4/30.
5. The 7200-2 receives the acknowledgment for 7500-internet, but before forwarding it, decrements the hop count (TTL) by 1. The TTL is now 0, which means that 7200-2 discards the packet and sends an ICMP dispose ip.hopcount message back to 7200-1.

The second scenario is similar; the only difference is that the 7500-internet sends a multicast hello to 7200-1, and the 7200-1 unicast update packet is forwarded to 7200-2 (for the same reasons explained earlier), where it is dropped.

Conclusion

Here are a few tips for troubleshooting:

- Proceed methodically, from the simple to the complex. In the example presented in this document, you started at the farthest device (5800-1) and worked toward the problem, using simple **show** command output.
- Try to make the investigation as specific as possible; for example, **show ip route 20.0.0.0** is more specific and can be more useful than **show ip route**.
- Use debug commands only when required. Often, the expertise of troubleshooters is inversely proportional to the number of debug commands that they enable to solve the problem. Debugs may have negative consequences on a device performance (CPU in particular), and so should be used only when necessary.
- Do not commit yourself to early conclusions. At one point during troubleshooting, it was tempting to conclude there is a physical layer problem between 7500-internet and 7200-1. Pursuing that conclusion may have wasted valuable time testing the WAN or working with the carrier. Since there were no errors or drops on the physical interface (as seen during the verification stage), WAN errors would not have been a logical issue to suspect.