360RSW05-WKB-TA04

Cisco 360 CCIE R&S Exercise Workbook Introduction

The Cisco 360 CCIE® R&S Exercise Workbook contains 20 challenging scenarios at the Cisco CCIE level that can be used for rigorous self-paced practice.

Each lab provides an extensive answer key, Mentor Guide support, and verification tables and is designed to maximize learning by providing practical experience. Also, self-paced learning resources such as the Cisco 360 CCIE R&S Reference Library and Cisco 360 CCIE R&S lessons supplement the Exercise Workbook scenarios.

Cisco 360 CCIE R&S Exercise Workbook Lab 4 Troubleshooting Section Answer Key

Table of Contents

Cisco 360 CCIE R&S Exercise Workbook Lab 4 Troubleshooting Section Answer Key...... 2
Answer Key Structure 4
Section One 4
Section Two 4
Exercise Workbook Lab 4 Troubleshooting Section Answer Key 5
Grading and Duration 5
Difficulty Level 5
Restrictions and Goals 5
Explanation of Each of the Restrictions and Goals 7
1. Switched Network Troubleshooting Section 8
1.1. Symptom: All VLANs are allowed over the trunk ports connecting SW1 and SW2. 8
1.2. Symptom: SW 4 does not have a blocked port in VLAN 25. 10
2. IPv4 OSPF Troubleshooting Section 13
2.1. Symptom: R1 has neighbors R3 and R4, but routes are not placed into the table. 13
2.2. Symptom: R4 and SW2 are not OSPF neighbors. 17
3. EIGRP Troubleshooting Section 20
4. IPv4 RIP Troubleshooting Section 22
4.1. Symptom: R1 is receiving no RIP routes from R3. 22
5. IPv4 Redistribution Troubleshooting Section 23
6. BGP Troubleshooting Section 27
7. IPv6 Troubleshooting Section 31
8. IP Multicast Troubleshooting 36

Answer Key Structure

Section One

The answer key PDF document is downloadable from the web portal.

Section Two

To obtain a comprehensive view of the configuration for a specific section, access the Mentor Guide engine in the web portal.

Exercise Workbook Lab 4 Troubleshooting Section Answer Key

Note	Regardless of any configuration you perform in this lab, it is very important that you conform to the general guidelines that are provided in the "Restrictions and Goals" section. If you do not conform to the guidelines, you could have a significant deduction of points in your final score.

Grading and Duration

- Troubleshooting lab duration: 2 hours
- Troubleshooting lab maximum score: 24 points

Note	You can assess your progress on the self-paced labs in this workbook by adding up the points that are assigned to sections and tasks. Consider taking the full Assessment Labs to assess your readiness level.

Difficulty Level

- Difficulty: Intermediate

Restrictions and Goals

Note	Read this section carefully.

- To receive any credit for a subsection, you must fully complete the subsection as per requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets displayed in the Lab IPv4 IGP diagram are /24 networks of 192.168.0.0, except as noted otherwise. *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Advertise loopback interfaces with their original masks.
- All IP addresses involved in this scenario must be reachable unless explicitly specified otherwise.
- Unless explicitly specified otherwise, addresses and networks that are advertised in the BGP section need to be reachable by all BGP routers but do not have to be reachable by routers that use only IGP.
- Use conventional routing algorithms only, unless specified otherwise.
- Do not create new interfaces to fulfill IGP requirements, and do not summarize unless explicitly asked to do so.

- *Do not* modify the hostname, console, or vty configuration unless you are specifically asked to do so.
- *Do not* modify the initial interface or IP address numbering.

Explanation of Each of the Restrictions and Goals

IPv4 subnets displayed in the IPv4 IGP diagram belong to network 192.168.0.0./16

All IP addresses in this lab belong to the 192.168.0.0/16 address space, with the exception of prefixes that are explicitly specified as being part of a different IP space.

Advertise loopback interfaces with their original masks.

The original mask is the mask that is configured on the loopback interface. Open Shortest Path First (OSPF), by default, treats loopback interfaces as host routes and advertises them as /32 prefixes. The requirement to advertise loopback interfaces with their original masks precludes using the default OSPF network type for the loopback interfaces. You need to provide a solution such as changing the OSPF network type or summarizations.

All IP addresses involved in this scenario must be reachable.

This goal a key goal to observe. It requires that all your IGPs and all your routing policy tasks must be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using the **distribute-lists**, **route-maps**, and **distance** commands. A key point to remember about this lab is that the term "redistribution" is not explicitly used. However, you must perform redistribution to assure that all IP addresses are reachable without the use of static routes or 0.0.0.0/0 routes.

Addresses and networks advertised in the BGP section need to be reachable by all BGP routers but do not have to be reachable by IGP-only routers.

This statement relaxes the requirement that all IP addresses must be reachable. The BGP prefixes need only be reachable among the routers specified in the BGP section. They can be used in other unicast tables. However, BGP routers need to have the prefixes in the routing tables as well as be able to forward traffic to the addresses known via BGP.

Use conventional routing algorithms.

This restriction prevents you from solving any problems by configuring policy routing. At the heart of this restriction is the interpretation of the phrase "conventional routing algorithms." Although this phrase can be interpreted in several different ways, this interpretation is applied in this workbook:

> Conventional routing algorithms are routing algorithms that apply destination-based prefix lookups in a routing table. Conventional routing algorithms do not use any type of information other than the destination address to make a packet-forwarding decision.

1. Switched Network Troubleshooting Section

Here is a Layer 2 diagram showing the required VLAN topology for this lab:

VLAN Propagation Diagram

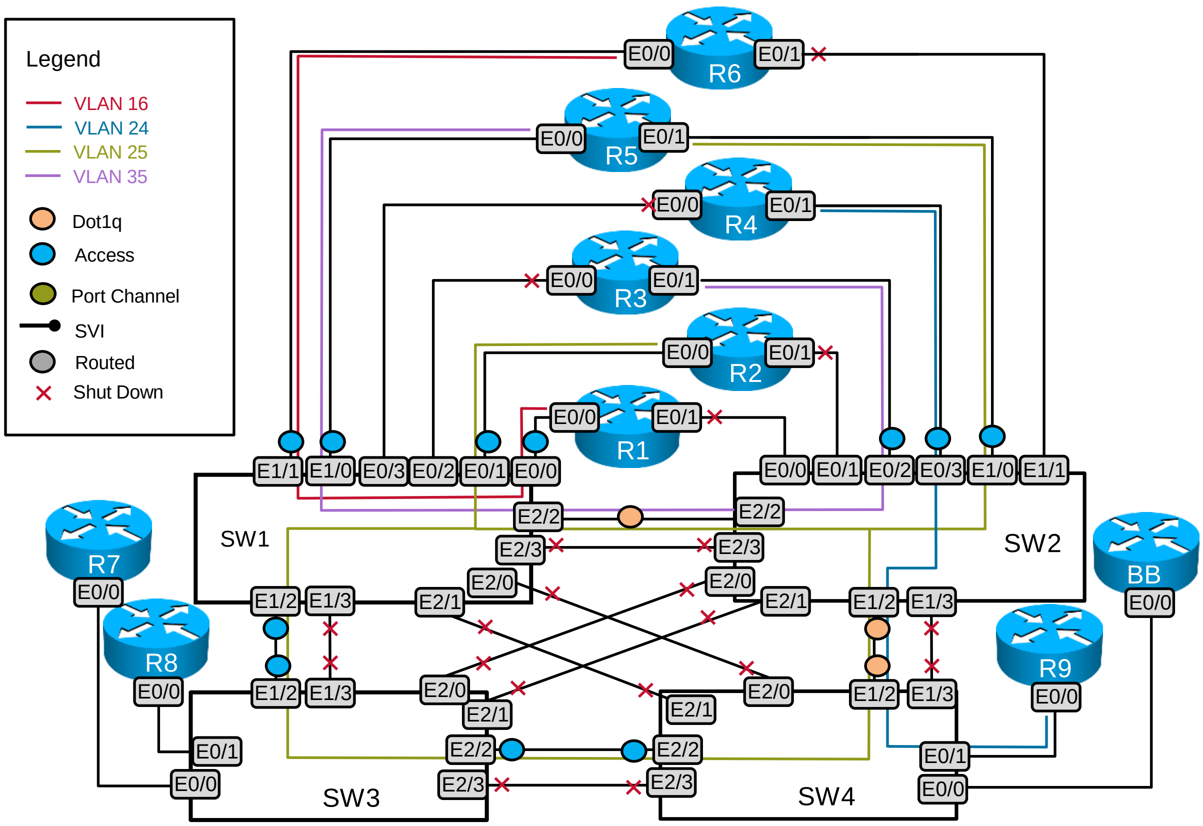

1.1. Symptom: All VLANs are allowed over the trunk ports connecting SW1 and SW2.

Analysis and testing:

According to the requirements of the scenario, the trunk connection between SW1 and SW2 should allow only the necessary VLANs across the trunk.

Note that the initial configuration allows all VLANs across the trunk:

```
SW1#show interfaces trunk

Port        Mode         Encapsulation  Status        Native vlan
Et2/2       on           802.1q         trunking      1

Port        Vlans allowed on trunk
Et2/2       1-4094

Port        Vlans allowed and active in management domain
Et2/2       1,16,25,35
```



```
Port                Vlans in spanning tree forwarding state and not pruned
Et2/2               1,16,25,35
SW1#
SW2#show interfaces trunk

Port                Mode         Encapsulation  Status        Native vlan
Et1/2               on           802.1q         trunking      1
Et2/2               desirable    n-802.1q       trunking      1

Port                Vlans allowed on trunk
Et1/2               24-25
Et2/2               1-4094

Port                Vlans allowed and active in management domain
Et1/2               24-25
Et2/2               1,24-25,35

Port                Vlans in spanning tree forwarding state and not pruned
Et1/2               24-25
Et2/2               1,24-25,35
SW2#
```

Likely cause: No VLAN pruning commands are configured on the trunk ports connecting SW1 and SW2.

This can be clearly demonstrated in two ways with two useful Cisco IOS **show** commands:

```
SW1#show running-config int e2/2
Building configuration...

Current configuration : 103 bytes
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
end

SW2#show running-config int e2/2
Building configuration...

Current configuration : 42 bytes
!
interface Ethernet2/2
 duplex auto
end

SW2#
```

With a configuration of this type, all VLANs will be allowed over the trunk port between SW1 and SW2.

Resolution: Configure* switchport trunk allow vlans *on SW1 and SW2.

It is clear that VLAN 35 must cross this trunk, since one device on this VLAN is connected to SW1, and the other device is connected to SW2. As the diagram above shows, VLAN 25 must also be allowed across this trunk to support the spanning-tree optimization task. Manually prune this trunk using the commands below:

```
SW1#conf t
SW1(config)#int e2/2
SW1(config-if)# switchport trunk allowed vlan 25,35
SW1(config-if)#end
```

```
SW2#conf t
SW2(config)#int e2/2
SW2(config-if)# switchport trunk allowed vlan 25,35
SW2(config-if)#end
```

Now, verify this configuration with the following **show** command:

```
SW1#sh interfaces trunk

Port        Mode         Encapsulation  Status        Native vlan
Et2/2       on           802.1q         trunking      1

Port        Vlans allowed on trunk
Et2/2       25,35

Port        Vlans allowed and active in management domain
Et2/2       25,35

Port        Vlans in spanning tree forwarding state and not pruned
Et2/2       25,35

SW2#sho int trunk

Port                Mode         Encapsulation  Status        Native vlan
Et1/2               on           802.1q         trunking      1
Et2/2               desirable    n-802.1q       trunking      1

Port                Vlans allowed on trunk
Et1/2               24-25
Et2/2               25,35

Port                Vlans allowed and active in management domain
Et1/2               24-25
Et2/2               25,35

Port                Vlans in spanning tree forwarding state and not pruned
Et1/2               24-25
Et2/2               25,35
SW2#
```

You see that these trunk ports are now configured to meet the specific requirements of this scenario: to allow only the necessary VLANs over the trunk port between SW1 and SW2.

1.2. Symptom: SW4 does not have a blocked port in VLAN 25.

Analysis and testing:

Here is the initial spanning-tree state on SW4:

```
SW4#show spanning-tree vlan 25

VLAN0025
  Spanning tree enabled protocol ieee
  Root ID    Priority    8217
             Address     aabb.cc00.0800
             Cost        100
             Port        35 (Ethernet1/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    16409  (priority 16384 sys-id-ext 25)
             Address     aabb.cc00.0a00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```



```
             Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -------------------------------
-
Et1/2               Root FWD 100       128.35   Shr
Et2/2               Desg FWD 100       128.67   Shr

SW4#
```

Note the bridge priority for VLAN 25 is 16409, the root port is Et1/2, and the potential alternate root port is Et2/2, which is designated, not blocked. In the following output, you see the initial spanning-tree state for VLAN 25 on SW1, SW2, and SW3. Note that the root bridge for VLAN 25 is SW2, with a bridge priority of 8217.

```
SW1#show span vlan 25

VLAN0025
  Spanning tree enabled protocol ieee
  Root ID    Priority    8217
             Address     aabb.cc00.0800
             Cost        100
             Port        67 (Ethernet2/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32793  (priority 32768 sys-id-ext 25)
             Address     aabb.cc00.0700
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -------------------------------
-
Et0/1               Desg FWD 100       128.2    Shr
Et1/2               Desg FWD 100       128.35   Shr
Et2/2               Root FWD 100       128.67   Shr

SW1#

SW2#show span vlan 25

VLAN0025
  Spanning tree enabled protocol ieee
  Root ID    Priority    8217
             Address     aabb.cc00.0800
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8217   (priority 8192 sys-id-ext 25)
             Address     aabb.cc00.0800
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -------------------------------
-
Et1/0               Desg FWD 100       128.33   Shr
Et1/2               Desg FWD 100       128.35   Shr
Et2/2               Desg FWD 100       128.67   Shr

SW2#

SW3#show span vlan 25

VLAN0025
  Spanning tree enabled protocol ieee
  Root ID    Priority    8217
```

```
                Address     aabb.cc00.0800
                Cost        200
                Port        67 (Ethernet2/2)
                Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32793  (priority 32768 sys-id-ext 25)
             Address     aabb.cc00.0900
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et1/2               Altn BLK 100       128.35   Shr
Et2/2               Root FWD 100       128.67   Shr

SW3#
```

Likely cause: The blocked port is on SW3, not SW4.

With the root bridge at SW2 and equal cost paths, the blocked port will be on SW3, not SW4.

Resolution: Move the blocked port by making SW1 the root bridge.

The task requires that you use only bridge priorities to achieve the desired result. You can move the blocked port to SW4 by making SW1 the root bridge. One way to do this would be to lower the bridge priority for VLAN 25 on SW1 to 0 with the command **spanning-tree vlan 25 priority 0**. Here is the resulting spanning-tree state on SW4:

```
SW1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#spanning-tree vlan 25 priority 0
SW1(config)#end
SW1#


SW4#show span vlan 25

VLAN0025
  Spanning tree enabled protocol ieee
  Root ID    Priority    25
             Address     aabb.cc00.0700
             Cost        200
             Port        35 (Ethernet1/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    16409  (priority 16384 sys-id-ext 25)
             Address     aabb.cc00.0a00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 15

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et1/2               Root FWD 100       128.35   Shr
Et2/2               Altn BLK 100       128.67   Shr

SW4#
```


The root port is still Et1/2, but now Et2/2 is an alternate path to the root and is blocked. Note that the advertised root bridge priority is the sum of the configured priority and the VLAN ID; this is a consequence of the MAC address reduction feature.

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter over 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

---

2. IPv4 OSPF Troubleshooting Section

2.1. Symptom: R1 has neighbors R3 and R4, but routes are not placed into the table.

Analysis and testing:

R1 can ping the IP addresses on the Ethernet subnet that it shares with R3 and R4, and it can form a neighbor relationship with both R3 and R4. However, it is not receiving any routes from R3 and R4.

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.4           0   FULL/  -        00:01:37    192.168.134.4   Ethernet0/2
1.1.1.13          0   FULL/  -        00:01:45    192.168.134.3   Ethernet0/2
1.1.1.6           1   FULL/DR         00:00:36    192.168.1.1     Ethernet0/0
R1#
R1#sh ip route ospf | inc via.*, Ethernet0/2
R1#
```

To investigate, start with the **show ip ospf interface** command:

```
R1#show ip ospf interface e0/2
Ethernet0/2 is up, line protocol is up
  Internet Address 192.168.134.1/28, Area 0, Attached via Interface Enable
  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_MULTIPOINT, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:13
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 1.1.1.4
    Adjacent with neighbor 1.1.1.13
  Suppress hello for 0 neighbor(s)
R1#
```

```
R3#show ip ospf interface e0/2
Ethernet0/2 is up, line protocol is up
  Internet Address 192.168.134.3/28, Area 0, Attached via Interface Enable
  Process ID 1, Router ID 1.1.1.13, Network Type NON_BROADCAST, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
```

```
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.13, Interface address 192.168.134.3
  Backup Designated router (ID) 1.1.1.1, Interface address 192.168.134.1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:28
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
R3#

R4#show ip ospf interface e0/2
Ethernet0/2 is up, line protocol is up
  Internet Address 192.168.134.4/28, Area 0, Attached via Interface Enable
  Process ID 1, Router ID 1.1.1.4, Network Type NON_BROADCAST, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.4, Interface address 192.168.134.4
  Backup Designated router (ID) 1.1.1.1, Interface address 192.168.134.1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:14
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
R4#
```

As you can see, all three OSPF speakers see each other as OSPF neighbors. However, R1 is configured with an OSPF point-to-multipoint network type, and both R3 and R4 are configured with an OSPF nonbroadcast network type. This configuration results in no routes learned on R1 from either R3 or R4. This can be clearly seen from the **show ip route** display below:

```
R1#show ip route ospf | beg Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 3 subnets
O        1.1.1.2 [110/21] via 192.168.1.1, 00:32:34, Ethernet0/0
O        1.1.1.6 [110/11] via 192.168.1.1, 00:32:59, Ethernet0/0
      192.168.1.0/24 is variably subnetted, 6 subnets, 2 masks
O        192.168.1.2/31 [110/20] via 192.168.1.1, 00:32:59, Ethernet0/0
O        192.168.1.4/31 [110/30] via 192.168.1.1, 00:32:34, Ethernet0/0
R1#
```

Likely cause: An OSPF network type mismatch has occurred.

Because R1 is configured as with an OSPF point-to-multipoint network type and R3 and R4 are configured with an OSPF nonbroadcast network type, a problem will arise when the OSPF shortest path first (SPF) tree is constructed. These two network types model the shared link differently: transit versus point-to-point. The OSPF process identifies this inconsistency in its database and marks the router link-state advertisements (LSAs) learned from the neighbor as not reachable. Here you see router LSAs from R1 and R4:

```
R1#show ip ospf database router 1.1.1.4

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Router Link States (Area 0)

  Adv Router is not-reachable in topology Base with MTID 0
  LS age: 69
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.4
  Advertising Router: 1.1.1.4
  LS Seq Number: 80000005
  Checksum: 0xCFF6
  Length: 72
  Area Border Router
  Number of Links: 4

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 1.1.1.13
     (Link Data) Router Interface address: 192.168.234.4
      Number of MTID metrics: 0
       TOS 0 Metrics: 10

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 1.1.1.12
     (Link Data) Router Interface address: 192.168.234.4
      Number of MTID metrics: 0
       TOS 0 Metrics: 10

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 192.168.234.4
     (Link Data) Network Mask: 255.255.255.255
      Number of MTID metrics: 0
       TOS 0 Metrics: 0

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 192.168.134.4
     (Link Data) Router Interface address: 192.168.134.4
      Number of MTID metrics: 0
       TOS 0 Metrics: 10


R1#

R4#show ip ospf database router 1.1.1.1

            OSPF Router with ID (1.1.1.4) (Process ID 1)

                Router Link States (Area 0)

  Adv Router is not-reachable in topology Base with MTID 0
  LS age: 153
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000004
  Checksum: 0x23E7
  Length: 60
```

```
  Area Border Router
  AS Boundary Router
  Number of Links: 3

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 1.1.1.4
     (Link Data) Router Interface address: 192.168.134.1
      Number of MTID metrics: 0
       TOS 0 Metrics: 10

    Link connected to: another Router (point-to-point)
     (Link ID) Neighboring Router ID: 1.1.1.13
     (Link Data) Router Interface address: 192.168.134.1
      Number of MTID metrics: 0
       TOS 0 Metrics: 10

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 192.168.134.1
     (Link Data) Network Mask: 255.255.255.255
      Number of MTID metrics: 0
       TOS 0 Metrics: 0


R4#
```

Resolution: Change the OSPF network type to point-to-multipoint on R3 and R4.

The resolution to this issue is to change the OSPF network type to point-to-multipoint on R3 and R4. You must do so because the requirements of the scenario state the following:

"OSPF should advertise IPv4 addresses 192.168.134.0/28 and 192.168.234.0/24 subnets as host routes."

As soon as you change these network types, the routers will reform the neighbor relationships and the proper OSPF routes will be placed into the routing table.

```
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#int e0/2
R3(config-if)#ip ospf network point-to-multipoint
R3(config-if)#end
R3#
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int e0/2
R4(config-if)#ip ospf network point-to-multipoint
R4(config-if)#end
R4#
R1#sh ip route ospf | inc via.*, Ethernet0/2
O IA     1.1.1.4 [110/11] via 192.168.134.4, 00:00:35, Ethernet0/2
O E2     1.1.1.12 [110/20] via 192.168.134.4, 00:00:35, Ethernet0/2
                  [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O E2     1.1.1.13 [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O E2     1.1.1.15 [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O E2  5.0.0.0/8 [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O IA     192.168.1.0/28 [110/20] via 192.168.134.4, 00:00:35, Ethernet0/2
O E2     192.168.1.8/31 [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O E2     192.168.1.10/31 [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O        192.168.134.3/32 [110/10] via 192.168.134.3, 00:00:57, Ethernet0/2
O        192.168.134.4/32 [110/10] via 192.168.134.4, 00:00:35, Ethernet0/2
O        192.168.234.2 [110/20] via 192.168.134.4, 00:00:35, Ethernet0/2
                       [110/20] via 192.168.134.3, 00:00:57, Ethernet0/2
O        192.168.234.3 [110/10] via 192.168.134.3, 00:00:57, Ethernet0/2
O        192.168.234.4 [110/10] via 192.168.134.4, 00:00:35, Ethernet0/2
R1#
```


2.2. Symptom: R4 and R9 are not OSPF neighbors.

Analysis and testing:

R4 and R9 are able to ping each other on VLAN 24; however, they cannot form an OSPF neighbor relationship.

```
R4#ping 192.168.1.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R4#
```

A suggested opening verification command to troubleshoot any problem of OSPF neighbor relationship formation is the **show ip ospf interface** command. This command verifies that OSPF is configured on a specific interface. If OSPF is not configured on the interface, then no other OSPF operations on that interface will work. Therefore, begin your troubleshooting process with this command on R4:

```
R4#show ip ospf interface e0/1
Ethernet0/1 is up, line protocol is up
  Internet Address 192.168.1.4/28, Area 2, Attached via Interface Enable
  Process ID 1, Router ID 1.1.1.4, Network Type BROADCAST, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.4, Interface address 192.168.1.4
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:04
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
R4#
```

Next, run the same command on R9:

```
R9#show ip ospf interface e0/0
Ethernet0/0 is up, line protocol is up
  Internet Address 192.168.1.5/31, Area 2, Attached via Network Statement
  Process ID 1, Router ID 1.1.1.2, Network Type BROADCAST, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.2, Interface address 192.168.1.5
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:05
  Supports Link-local Signaling (LLS)
```

```
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 3/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
R9#
```

You see that OSPF is enabled on these interfaces but that no neighbors are discovered.

Now you know two things:

- You can ping the IP addresses of R4 and R9.
- OSPF is enabled on both interfaces.

Next, enable the **debug ip ospf hello** utility to see if you can discover any hidden issues with this Cisco IOS tool. Here is an example on R9:

```
R9#debug ip ospf hello
OSPF hello debugging is on
R9#
*May 26 18:25:37.501: OSPF-1 HELLO Et0/1: Rcv hello from 1.1.1.6 area 2
192.168.1.2
R9#
*May 26 18:25:41.190: OSPF-1 HELLO Et0/0: Rcv hello from 1.1.1.4 area 2
192.168.1.4
*May 26 18:25:41.190: OSPF-1 HELLO Et0/0: Mismatched hello parameters from
192.168.1.4
*May 26 18:25:41.190: OSPF-1 HELLO Et0/0: Dead R 40 C 40, Hello R 10 C 10 Mask
R 255.255.255.240 C 255.255.255.254
*May 26 18:25:42.007: OSPF-1 HELLO Et0/1: Send hello to 224.0.0.5 area 2 from
192.168.1.3
R9#u
*May 26 18:25:42.549: OSPF-1 HELLO Et0/0: Send hello to 224.0.0.5 area 2 from
192.168.1.5
R9#u all
All possible debugging has been turned off
R9#
```

This debug utility tells you that there are mismatched hello parameters from 192.168.1.4, specifically from R4.

Likely cause: The hello parameters between router R4 and R9 are mismatched.

Between any two OSPFv2 neighbors, the following parameters must match in order for an adjacency to form:

- Hello and dead timers
- Prefix and prefix mask length
- Maximum transmission unit (MTU) size

On R4, the **debug ip ospf hello** output states the following:

```
*May 26 18:25:41.190: OSPF-1 HELLO Et0/0: Dead R 40 C 40, Hello R 10 C 10 Mask
R 255.255.255.240 C 255.255.255.254
```


The IP address mask length of the local router—R9, in this case—is /31, as indicated in the scenario diagram. However, according to this debug output, the IP address mask length of the remote router—in this case, R4—is /28. Check this with a **show run** command:

```
R4#show running-config int e0/1
Building configuration...

Current configuration : 137 bytes
!
interface Ethernet0/1
 ip address 192.168.1.4 255.255.255.240
 ip ospf 1 area 2
 ipv6 address 2001:2:24::4/64
 ipv6 ospf 1 area 2
end

R4#

R9#show running-config int e0/0
Building configuration...

Current configuration : 119 bytes
!
interface Ethernet0/0
 ip address 192.168.1.5 255.255.255.254
 ipv6 address 2001:2:24::2/64
 ipv6 ospf 1 area 2
 ipv6 ospf 1 area 2
end

R9#
```

You can see that there is an IP address mask length mismatch between R4 and R9.

Resolution: Change the R4 E0/1 IP address mask length from /28 to /31.

When this change is made, you can predict that the OSPF adjacency will form successfully. This troubleshooting ticket is successfully resolved.

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int e0/1
R4(config-if)#ip address 192.168.1.4 255.255.255.254
% Warning: use /31 mask on non point-to-point interface cautiously
R4(config-if)#end
R4#
*May 26 18:31:33.083: %SYS-5-CONFIG_I: Configured from console by console
(cierswbv5-te-lab04-sc, SJ)
R4#
*May 26 18:31:35.388: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.2 on Ethernet0/1
from LOADING to FULL, Loading Done
R4#
```

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter over 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

3. EIGRP Troubleshooting Section

3.1. Symptom: The neighbor relationship between R2 and R5 is flapping.

Analysis and testing:

Enhanced Interior Gateway Routing Protocol (EIGRP)-speaking routers R2 and R5 are able to ping each other on VLAN 25.

```
R2#ping 192.168.1.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.11, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
R2#
```

However, they cannot form an EIGRP neighbor relationship with each other. The log shows the EIGRP neighbor relationship going up and down every 80 seconds.

```
R2#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1)
R2#
```

A suggested opening verification command to troubleshoot any EIGRP neighbor relationship formation problem is the **show ip eigrp interface** command. Here, you see output for R2 and R5:

```
R2#show ip eigrp interface detail e0/0
EIGRP-IPv4 Interfaces for AS(1)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et0/0                    0        0/0       0/0           0       0/0
0           0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 0/0
  Hello's sent/expedited: 760/1
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is not set
R2#

R5#show ip eigrp interface detail e0/1
EIGRP-IPv4 Interfaces for AS(1)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et0/1                    0        0/0       0/0           0       0/0
0           0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 0/0
  Hello's sent/expedited: 765/1
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
```



```
  Authentication mode is not set
R5#
```

Note that the EIGRP interfaces are configured correctly, but the EIGRP neighbor relationship is not formed between R2 and R5.

Enable the **debug eigrp neighbor** command. Here is what you find:

```
R2#deb ip eigrp neighbor
R2#
R2#
R2#
*May 26 18:40:06.335: EIGRP: Neighbor blocked on Ethernet0/0; insufficient MTU
R2#
*May 26 18:40:11.000: EIGRP: Neighbor blocked on Ethernet0/0; insufficient MTU
R2#u all
All possible debugging has been turned off
R2
```

Likely cause: R2 E0/0 is configured with insufficient MTU size.

Verify the Etherent0/0 interface configuration:

```
R2#show running-config int e0/0
Building configuration...

Current configuration : 127 bytes
!
interface Ethernet0/0
 ip address 192.168.1.10 255.255.255.254
 ip mtu 150
 ipv6 address 2001:116::2/64
 ipv6 eigrp 111
end

R2#
```

Resolution: Change the MTU size on the Ethernet0/0 interface on R2.

Correct this problem by reconfiguring the MTU size on the E0/0 interface of R2 to the default value of 1500. Once this is completed, R2 will accept and reply to the reliable packets from R5, and the EIGRP neighbor relationship will form.

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#int e0/0
R2(config-if)#no  ip mtu 150
R2(config-if)#end
R2#
*May 26 18:45:53.806: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.1.11
(Ethernet0/0) is up: new adjacency
R2#
*May 26 18:45:55.033: %SYS-5-CONFIG_I: Configured from console by console
(cierswbv5-te-lab04-sc, SJ)
R2#
```

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter over 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

---

4. IPv4 RIP Troubleshooting Section

4.1. Symptom: R1 is receiving no RIP routes from R3.

Analysis and testing:

Verify the RIP routing table on R1:

```
R1#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       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, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

R1#
```

To verify that RIP is advertising only over the link between R1 and R3, enable **debug ip rip** on both routers. After a few minutes, you will find that R1 is sending updates, but is not receiving any from R3. To confirm, enter **show ip protocols** on R1. You will see that R3 is not listed as the RIP gateway.

Likely cause: All interfaces on R3 are passive for RIP.

Investigating on R3, you enter the **show ip protocols** command, and see output similar to the following:

```
R3#show ip protocols | section rip
    rip, includes subnets in redistribution
  Redistributing: ospf 1, rip
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 16 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: ospf 1 (internal, external 1 & 2, nssa-external 1 & 2)

  Redistributing: eigrp 1, rip
  Default version control: send version 2, receive version 2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    192.168.1.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/1
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
    Serial1/2
    Serial1/3
    Loopback103
    RG-AR-IF-INPUT1
    Tunnel0
     VoIP-Null0
    VoIP-Null0
  Routing Information Sources:
```



```
    Gateway         Distance      Last Update
    192.168.1.16         120      00:00:12
  Distance: (default is 120)
R3#
```

Resolution: Disable** passive-interface **on S1/0 of R3.

Under the RIP process on R3, enter the command **no passive-interface s1/0**.

```
R3#conf t
R3(config-router)#end
R3#
```

Verify the RIP routing table on R1 again:

```
R1#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       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, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      192.168.234.0/24 is variably subnetted, 4 subnets, 2 masks
R        192.168.234.0/28 [120/3] via 192.168.1.17, 00:00:06, Serial1/0
R1#
```

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter over 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

5. IPv4 Redistribution Troubleshooting Section

5.1. Symptom: R2 uses the OSPF E2 route instead of the EIGRP direct route for 5.0.0.0/8.

Analysis and testing:

The scenario requires that R2 and R3 use the fastest path to subnet 5.0.0.0/8 on R5. This will be the EIGRP routes across the Ethernet links. This route will be problematic in this scenario because its administrative distance will be 170 by default. This value will make it less desirable than any other routing source. Therefore, you can expect that R2 and R3 will not select the EIGRP path for the 5.0.0.0/8 network.

```
R2#show ip route 5.0.0.0
Routing entry for 5.0.0.0/8
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 10
  Redistributing via eigrp 1
  Advertised by eigrp 1 metric 10000 100 255 1 1500
```

```
  Last update from 192.168.234.3 on Ethernet0/3, 01:23:52 ago
  Routing Descriptor Blocks:
  * 192.168.234.3, from 1.1.1.13, 01:23:52 ago, via Ethernet0/3
      Route metric is 20, traffic share count is 1
R2#
R3#show ip route 5.0.0.0
Routing entry for 5.0.0.0/8
  Known via "rip", distance 120, metric 3
  Redistributing via ospf 1, eigrp 1, rip
  Advertised by ospf 1 subnets
                eigrp 1 metric 10000 100 255 1 1500
  Last update from 192.168.1.16 on Serial1/0, 00:00:08 ago
  Routing Descriptor Blocks:
  * 192.168.1.16, from 192.168.1.16, 00:00:08 ago, via Serial1/0
      Route metric is 3, traffic share count is 1
R3#
```

R2, the problematic router in this specific task, has a preconfigured solution in its initial configuration. However, this solution is not working:

```
router ospf 1
 router-id 1.1.1.12
 redistribute eigrp 1 subnets
 distance 171 1.1.1.13 0.0.0.0 5
!
access-list 5 permit 5.5.5.0 0.0.0.255
```

R2 knows that it will learn the 5.0.0.0/8 route from EIGRP with an administrative distance of 170. R2 also knows that it will learn the same route from OSPF via R3. R2 has set the administrative distance for this route to 171, yet it is still selecting OSPF as the preferred route.

Likely cause: An access list is misconfigured on R2.

Although the **distance** command that was entered on R2 is correct, a closer look at the access list reflects a problem. The access list is written as follows:

```
access-list 5 permit 5.5.5.0 0.0.0.255
```

Yet the route is actually as shown here:

```
R2#sh ip ospf database | b Ex
                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
…
5.0.0.0         1.1.1.13        61          0x80000052 0x00A397 0
…
```

The access list is matching on 5.5.5.0, but the route is only 5.0.0.0.

Resolution: Rewrite the access list to reflect the correct prefix.

To match the correct prefix, the access list on router R2 must be rewritten:

```
access-list 5 permit 5.0.0.0
```

Make this correction and apply it to the following configuration on router R2:

```
!
router ospf 1
```



```
 router-id 1.1.1.12
 log-adjacency-changes
 redistribute eigrp 1 subnets
 distance 171 1.1.1.13 0.0.0.0 5
!
access-list 5 permit 5.0.0.0
```

```
R2#clear ip route *
R2#show ip route 5.0.0.0
Routing entry for 5.0.0.0/8
  Known via "eigrp 1", distance 170, metric 409600, type external
  Redistributing via ospf 1, eigrp 1
  Advertised by ospf 1 subnets
  Last update from 192.168.1.11 on Ethernet0/0, 00:00:03 ago
  Routing Descriptor Blocks:
  * 192.168.1.11, from 192.168.1.11, 00:00:03 ago, via Ethernet0/0
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
R2#
```

6.2 Symptom: The IP address 5.5.5.5 is unreachable from router R1.

Analysis and testing:

A ping from R1 to IP address 5.5.5.5 fails, and a traceroute to the 5.5.5.5 address provides output similar to the following:

```
R1#trace 5.5.5.5
Type escape sequence to abort.
Tracing the route to 5.5.5.5
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.134.3 1 msec 0 msec 0 msec
  2 192.168.1.16 4 msec 4 msec 5 msec
  3 192.168.134.3 5 msec 4 msec 5 msec
  4 192.168.1.16 8 msec 8 msec 9 msec
  5 192.168.134.3 8 msec 9 msec 6 msec
  6 192.168.1.16 13 msec 13 msec 13 msec
  7 192.168.134.3 13 msec 12 msec 13 msec
  8 192.168.1.16 16 msec 16 msec 17 msec
  9 192.168.134.3 17 msec 17 msec 16 msec
 10 192.168.1.16 21 msec 23 msec 21 msec
 11 192.168.134.3 21 msec 21 msec 21 msec
 12 192.168.1.16 25 msec 26 msec 25 msec
 13 192.168.134.3 24 msec 25 msec 25 msec
 14 192.168.1.16 30 msec 28 msec 31 msec
 15 192.168.134.3 29 msec 28 msec 30 msec
 16 192.168.1.16 33 msec 34 msec 33 msec
 17 192.168.134.3 34 msec 34 msec 32 msec
 18 192.168.1.16 40 msec 38 msec 39 msec
 19 192.168.134.3 38 msec 36 msec 38 msec
 20 192.168.1.16 40 msec 41 msec 49 msec
 21 192.168.134.3 44 msec 43 msec 44 msec
 22 192.168.1.16 47 msec 44 msec 47 msec
 23 192.168.134.3 45 msec 51 msec 44 msec
 24 192.168.1.16 56 msec 49 msec 50 msec
 25 192.168.134.3 50 msec 50 msec 50 msec
 26 192.168.1.16 58 msec 54 msec 55 msec
 27 192.168.134.3 54 msec 51 msec 53 msec
 28 192.168.1.16 59 msec 58 msec 58 msec
 29 192.168.134.3 58 msec 55 msec 59 msec
 30 192.168.1.16 60 msec 62 msec 64 msec
R1#
```

On further inspection of the path between R1 and R5, you find the following route on R3:

```
R3#show ip route 5.5.5.5
Routing entry for 5.0.0.0/8
  Known via "rip", distance 120, metric 3
  Redistributing via ospf 1, eigrp 1, rip
  Advertised by ospf 1 subnets
                eigrp 1 metric 10000 100 255 1 1500
  Last update from 192.168.1.16 on Serial0/0/1, 00:00:08 ago
  Routing Descriptor Blocks:
  * 192.168.1.16, from 192.168.1.16, 00:00:08 ago, via Serial0/0/1
      Route metric is 3, traffic share count is 1
```

R3 has a RIP route to the 5.0.0.0/8 network that refers to R1 as the next hop. This is the opposite of the direction it needs to take to deliver the packet to its ultimate destination. Instead, R3 should use the EIGRP route to R5.

Likely cause: The administrative distance for RIP learned routes is lower than the administrative distance for EIGRP external routes.

R3 learns this route with administrative distance 120 from R1 and administrative distance 170 from R5.

Resolution: Raise the administrative distance of RIP routes to 172.

Raise the administrative distance of RIP routes on R3 to greater than 172. This is two more than the competing EIGRP external administrative distance and one more than the OSPF administrative distance of prefix 5.0.0.0/8. Recall that RIP routes are to be last resort. Raise the administrative distance of all RIP routes on R3 with the command **distance 172** under the RIP process.

```
R3(config)#router rip
R3(config-router)#
*May 26 19:10:16.249: %BGP-5-ADJCHANGE: neighbor 1.1.1.1 Up
R3(config-router)#distance 172
R3(config-router)#end
R3#
R3#clear ip route *
R3#show ip route 5.5.5.5
Routing entry for 5.0.0.0/8
  Known via "eigrp 1", distance 170, metric 409600, type external
  Redistributing via ospf 1, eigrp 1, rip
  Advertised by ospf 1 subnets
                rip metric 3
  Last update from 192.168.1.9 on Ethernet0/1, 00:00:02 ago
  Routing Descriptor Blocks:
  * 192.168.1.9, from 192.168.1.9, 00:00:02 ago, via Ethernet0/1
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
R3#
```

To verify the solution, clear IP routes on R3 and retry the trace from R1 to IP address 5.5.5.5.

```
R1#trace 5.5.5.5
Type escape sequence to abort.
Tracing the route to 5.5.5.5
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.134.3 0 msec 0 msec 0 msec
  2 192.168.1.9 1 msec *  1 msec
```



```
R1#
```

Now that you have addressed all of the IPv4 unicast issues, you will test reachability using this simple Tcl script. Enter the command **tclsh** and paste in this script. When it is complete, you will have a record of successful and unsuccessful pings. Enter the command **tclquit** to exit the command interpreter.

```
tclsh
foreach address {
1.1.1.1
192.168.1.0
192.168.1.16
192.168.134.1
1.1.1.12
192.168.1.10
192.168.234.2
1.1.1.13
192.168.1.8
192.168.1.17
192.168.234.3
192.168.134.3
1.1.1.4
192.168.1.4
192.168.234.4
192.168.134.4
5.5.5.5
1.1.1.15
192.168.1.9
192.168.1.11
1.1.1.6
192.168.1.1
192.168.1.2
1.1.1.2
192.168.1.3
192.168.1.5
} {ping $address}
```

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter more than 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

6. BGP Troubleshooting Section

6.1. Symptom: The BGP neighbor relationship between R1 and R3 goes up and down.

Analysis and testing:

The log shows that the peering relationship between R1 and R3 goes up and down. You will see messages similar to the following.

```
01:54:41.593: %BGP-3-NOTIFICATION: received from neighbor 1.1.1.1 4/0 (hold
time expired) 0 bytes
R3#
01:54:41.593: %BGP-5-ADJCHANGE: neighbor 1.1.1.1 Down BGP Notification received
```

When you examine the peering relationship between External Border Gateway Protocol (EBGP) speakers R1 and R3, you may see that a neighbor relationship has not been formed:

```
R1#sh ip bgp summ
BGP router identifier 1.1.1.1, local AS number 16
BGP table version is 2, main routing table version 2
1 network entries using 117 bytes of memory
1 path entries using 52 bytes of memory
2/1 BGP path/bestpath attribute entries using 248 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 417 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
1.1.1.13        4     1       0       0        0    0    0 never    Active
```

Begin your analysis of BGP by reviewing the supplied BGP topology diagram:

IPv4 BGP Diagram

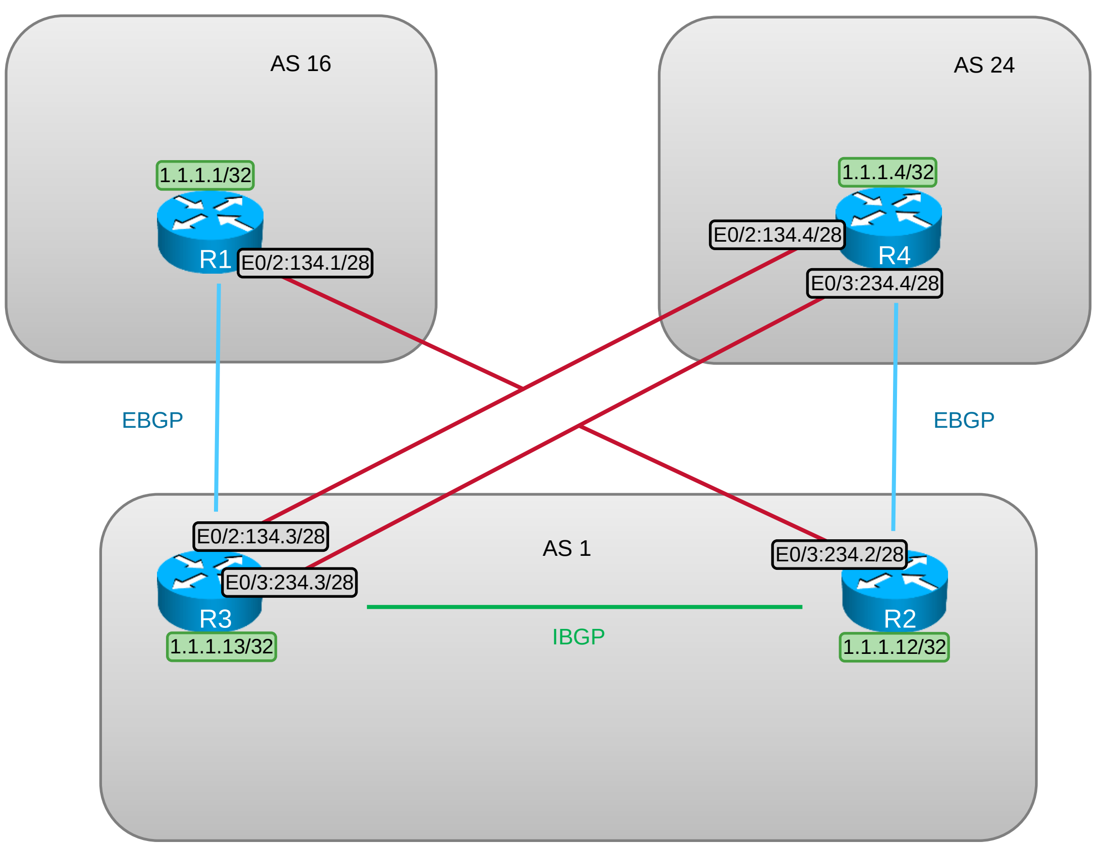

This cycling suggests route instability. You enter the **debug ip routing** command on R1 and R3 and see output similar to this:


```
01:50:29.469: RT: recursion error routing 1.1.1.1 - probable routing loop
```

Likely cause: There is a recursive routing error in the EBGP neighbor relationship formed between R1 and R .

Notice in the diagram that R1 and R3 are configured as EBGP peers with their respective loopback interfaces as the terminating points for their peering relationship. For this peering relationship to be formed, R1 and R3 need to advertise their respective loopback interfaces to each other using an IGP. However, according to the requirements of the scenario, R1 must also advertise its loopback interface via BGP. This requirement means that R1 will advertise its loopback interface to R3 as an EBGP route with an administrative distance of 20.

Resolution: Add the appropriate subnet mask to the BGP* network backdoor *commands on R3 and R4.

When the BGP **network** command is configured with the **backdoor** option, it sets the administrative distance for the route that it matches to 200. This configuration will solve the problem in this situation, because the administrative distance of the peering route will be higher for the EBGP source than the OSPF administrative distance of 110. The initialization script includes the command **network 1.1.1.1 backdoor** on R3 and R4. However, it does not work. Like any BGP network statement, the **network backdoor** command requires a precise match in the routing table. You will need to add the appropriate mask to this command to make it effective.

On R3 and R4, remove the command **network 1.1.1.1 backdoor** and replace it with **network 1.1.1.1 mask 255.255.255.255 backdoor**.

```
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#router bgp 1
R3(config-router)#no  network 1.1.1.1 backdoor
R3(config-router)#network 1.1.1.1 mask 255.255.255.255 backdoor
R3(config-router)#end
R3#
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#router bgp 24
R4(config-router)#no  network 1.1.1.1 backdoor
R4(config-router)#network 1.1.1.1 mask 255.255.255.255 backdoor
R4(config-router)#end
R4#
```

When you enter the **network backdoor** command, the 1.1.1.1/32 route is inserted in the R3 routing table as an OSPF route. The recursive routing problem is solved.

```
R3#sh ip ro 1.1.1.1
Routing entry for 1.1.1.1/32
  Known via "ospf 1", distance 110, metric 65, type inter area
  Redistributing via eigrp 1, rip
  Advertised by eigrp 1 metric 10000 100 255 1 1500
                rip metric 3
  Last update from 192.168.134.1 on Serial0/0/0.1, 1d02h ago
  Routing Descriptor Blocks:
  * 192.168.134.1, from 1.1.1.1, 1d02h ago, via Serial0/0/0.1
```

The route appears in the BGP table as the best route, but with a Routing Information Base (RIB) failure:

```
R3#sh ip bgp
```

```
BGP table version is 3, local router ID is 1.1.1.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
r> 1.1.1.1/32       1.1.1.1                  0             0 16 i
```

As shown here, the **rib-failure** output indicates that the route was inserted into the local IP routing table because of the lower administrative distance of another route source—in this case, OSPF:

```
R3#sh ip bgp rib-failure
Network            Next Hop                      RIB-failure   RIB-NH Matches
1.1.1.1/32         1.1.1.1             Higher admin distance              n/a
```

6.2. Symptom: R2 is not advertising the 1.1.1.1/32 route to R4.

Analysis and testing:

The scenario requires that R2 advertise prefix 1.1.1.1/32 to R4 only if this prefix is in the local routing table from OSPF. This is a description of the BGP synchronization requirement, and synchronization has been enabled on R2. Here, you see the condition of this route on R2:

```
R2#show ip b 1.1.1.1
BGP routing table entry for 1.1.1.1/32, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  16
    1.1.1.1 (metric 68) from 1.1.1.13 (1.1.1.13)
      Origin IGP, metric 0, localpref 100, valid, internal, not synchronized
```

Likely cause: The OSPF router ID and BGP router ID are mismatched.

Here is a rule to remember when synchronization is enabled and the underlying IGP is OSPF:

When the rule of synchronization is enabled and the underlying IGP is OSPF, the router ID of the OSPF Autonomous System Boundary Router (ASBR) must match the router ID of the advertising IBGP speaker. If you re-examine the following two Cisco IOS **show** command outputs, you will see that this is not the case:

```
R2#sh ip b 1.1.1.1
BGP routing table entry for 1.1.1.1/32, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  16
    1.1.1.1 (metric 68) from 1.1.1.13 (1.1.1.13)
      Origin IGP, metric 0, localpref 100, valid, internal, not synchronized
R2#sh ip route 1.1.1.1
Routing entry for 1.1.1.1/32
  Known via "ospf 1", distance 110, metric 21, type inter area
  Redistributing via eigrp 1
  Advertised by eigrp 1 metric 10000 100 255 1 1500
  Last update from 192.168.234.3 on Ethernet0/3, 00:04:19 ago
  Routing Descriptor Blocks:
  * 192.168.234.4, from 1.1.1.1, 00:04:19 ago, via Ethernet0/3
      Route metric is 21, traffic share count is 1
    192.168.234.3, from 1.1.1.1, 00:04:19 ago, via Ethernet0/3
      Route metric is 21, traffic share count is 1
R2#
```


Resolution: Disable IGP synchronization on R2.

Verify the BGP configuration on R2:

```
router bgp 1
 synchronization
 bgp log-neighbor-changes
 neighbor 1.1.1.4 remote-as 24
 neighbor 1.1.1.4 ebgp-multihop 255
 neighbor 1.1.1.4 update-source Loopback102
 neighbor 1.1.1.13 remote-as 1
 neighbor 1.1.1.13 update-source Loopback102
!
```

A solution to this problem is to configure the **no synchronization** command on R2:

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#router bgp 1
R2(config-router)#no synchronization
R2(config-router)#end
R2#
```

R2 will then display the BGP learned route as best, and will advertise it to R4.

```
R2#show ip bgp 1.1.1.1
BGP routing table entry for 1.1.1.1/32, version 2
Paths: (1 available, best #1, table default, RIB-failure(17))
  Advertised to update-groups:
     2
  Refresh Epoch 1
  16
    1.1.1.1 (metric 21) from 1.1.1.13 (1.1.1.13)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
R2#
```

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter over 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

7. IPv6 Troubleshooting Section

7.1. Symptom: There are no OSPFv3 adjacencies across Area 0.

Analysis and testing:

Begin by reviewing the IPv6 diagram for this scenario:

IPv6 IGP Diagram

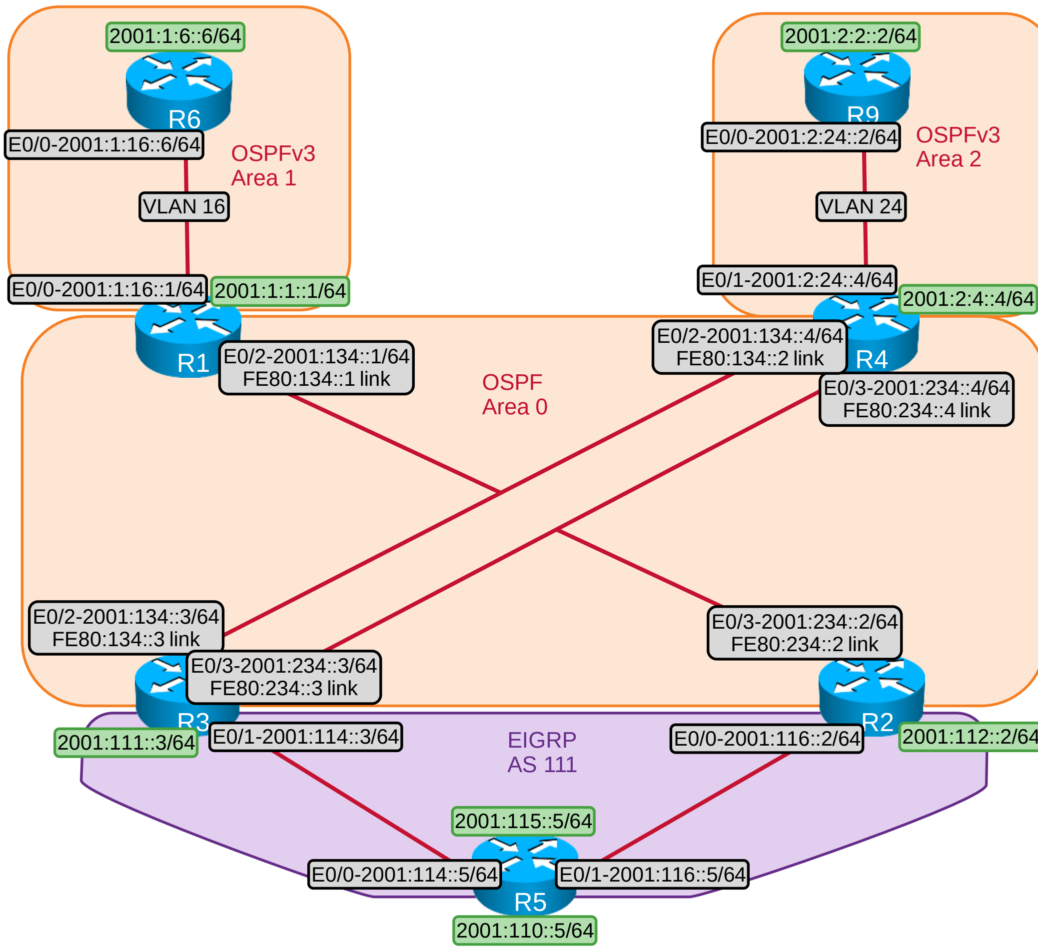

Likely cause: The OSPFv3 network type point-to-multipoint no-broadcast requires neighbor statements. The neighbor statements are missing in the configuration.

OSPFv3 network type point-to-multipoint nonbroadcast uses a destination unicast address for all packet transmissions. The IPv6 link-local addresses are used for the neighbor statements configuration.

Verify the OSPFv3 configuration on the Ethernet0/2 interface of R1:

```
R1#show running-config int e0/2
Building configuration...

Current configuration : 285 bytes
!
interface Ethernet0/2
 ip address 192.168.134.1 255.255.255.240
 ip pim sparse-mode
 ip ospf network point-to-multipoint
 ip ospf 1 area 0
 ipv6 address FE80:134::1 link-local
 ipv6 address 2001:134::1/64
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-multipoint non-broadcast
end

R1#
```


Note that the OSPFv3 neighbors are not configured on R1.

Verify the OSPFv3 configuration on the Ethernet0/2 interface of R3:

```
R3#show running-config int e0/2
Building configuration...

Current configuration : 285 bytes
!
interface Ethernet0/2
 ip address 192.168.134.3 255.255.255.240
 ip pim sparse-mode
 ip ospf network point-to-multipoint
 ip ospf 1 area 0
 ipv6 address FE80:134::3 link-local
 ipv6 address 2001:134::3/64
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-multipoint non-broadcast
end

R3#
```

Note that the OSPFv3 neighbors are not configured on R3.

Verify the OSPFv3 configuration on the Ethernet0/2 interface of R4:

```
R4#show running-config int e0/2
Building configuration...

Current configuration : 285 bytes
!
interface Ethernet0/2
 ip address 192.168.134.4 255.255.255.240
 ip pim sparse-mode
 ip ospf network point-to-multipoint
 ip ospf 1 area 0
 ipv6 address FE80:134::4 link-local
 ipv6 address 2001:134::4/64
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-multipoint non-broadcast
end

R4#
```

Resolution: Add neighbor statements on R1, R3, and R4.

It is worth noting that, although OSPFv2 neighbor statements are configured under routing protocol configuration mode for IPv4, OSPFv3 neighbor statements are configured under interface configuration mode using the IPv6 unicast link-local address.

Configure the OSPFv3 neighbor statements on R1, R3, and R4:

```
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#interface Ethernet0/2
R1(config-if)#ipv6 ospf nei FE80:134::3
R1(config-if)#ipv6 ospf nei FE80:134::4
R1(config-if)#end
R1#

R3(config)#int e0/2
R3(config-if)#ipv6 ospf nei FE80:134::1
```

```
R3(config-if)#ipv6 ospf nei FE80:134::4
R3(config-if)#end
R3#

R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int e0/2
R4(config-if)#ipv6 ospf nei FE80:134::1
R4(config-if)#ipv6 ospf nei FE80:134::3
R4(config-if)#end
R4#
```

Verify the OSPFv3 neighbor relationship again. Here is an example from R4:

```
R4#show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.4) (Process ID 1)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
1.1.1.13          0   FULL/  -        00:01:42    6               Ethernet0/3
1.1.1.12          0   FULL/  -        00:01:53    6               Ethernet0/3
1.1.1.1           0   FULL/  -        00:01:46    5               Ethernet0/2
1.1.1.13          0   FULL/  -        00:01:51    5               Ethernet0/2
1.1.1.2           1   FULL/BDR        00:00:32    3               Ethernet0/1
R4#
```

7.2. Symptom: OSPFv3 is not summarizing the IPv6 EIGRP domain routes correctly.

Analysis and testing:

The scenario requirements state that OSPFv3 must summarize the IPv6 EIGRP domain routes with the single best longest match. Here are the OSPFv3 external routes as seen on R1:

```
R1#show ipv6 route | inc OE2
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
OE2 2001:1::/45 [110/20]
OE2 2001:110::/30 [110/20]
OE2 2001:114::/64 [110/20]
OE2 2001:115::/64 [110/20]
OE2 2001:116::/64 [110/20]
R1#
```

Notice that there is a /30 summary of these routes in the table, but the longer matches are not suppressed.

Likely cause: The specified OSPFv3 area range address mask length is too long.

Here are the IPv6 EIGRP domain routes with bits 17 through 32 written out in binary:

2001:0000 0001 0001 0000::/64

2001:0000 0001 0001 0010::/64

2001:0000 0001 0001 0100::/64

2001:0000 0001 0001 0101::/64

2001:0000 0001 0001 0110::/64

Note that they differ only in the last 3 bits; the first 29 bits of each address are the same.

When these addresses are converted into binary, it becomes evident that a /30 will not include all these prefixes.

Resolution: Shorten the specified OSPFv3 area range address mask length to /29.

Remove the existing summary and shorten the OSPFv3 summary prefix mask length to the appropriate value. Once you have done this on R2 and R3, the IPv6 EIGRP domain routes will be represented by a single /29 route in the OSPF domain, and the longer matches will be suppressed.

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#ipv6 router ospf 1
R2(config-rtr)#no  summary-prefix 2001:110::/30
R2(config-rtr)# summary-prefix 2001:110::/29
R2(config-rtr)#end
R2#
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#ipv6 router ospf 1
R3(config-rtr)#no  summary-prefix 2001:110::/30
R3(config-rtr)# summary-prefix 2001:110::/29
R3(config-rtr)#end
R3#

R1#show ipv6 route | inc OE2
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
OE2 2001:1::/45 [110/20]
OE2 2001:110::/29 [110/20]
R1#
```

Now that you have addressed all the IPv6 unicast connectivity issues, you will test reachability using this simple Tcl script. Enter the command **tclsh** and paste in this script. When it is complete, you will have a record of successful and unsuccessful pings. Enter the command **tclquit** to exit the command interpreter.

```
tclsh
foreach address {
2001:1:16::1
2001:1:1::1
2001:116::2
2001:112::2
2001:114::3
2001:111::3
2001:2:24::4
2001:2::4
2001:114::5
2001:116::5
2001:110::5
2001:115::5
2001:1:16::6
2001:1:6::6
2001:2:24::2
2001:2:2::2
} {ping $address}
```

Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter more than 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.

8. IP Multicast Troubleshooting

1.1. Symptom: No replies are received to pings to 239.255.1.1 from R2.

Analysis and testing:

A ping to the multicast group 239.255.1.1 from R2 is not successful:

```
R2#ping 239.255.1.1 source e0/3
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.255.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.234.2
.
R2#
```

Verify the mroute table on the first-hop multicast router R3.

```
R3#sh ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.255.1.1), 00:00:48/00:02:13, RP 1.1.1.6, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 192.168.134.1
  Outgoing interface list:
    Loopback103, Forward/Sparse, 00:00:47/00:02:13

(*, 224.0.1.40), 00:00:48/00:02:13, RP 1.1.1.6, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 192.168.134.1
  Outgoing interface list:
    Loopback103, Forward/Sparse, 00:00:47/00:02:13

R3#
```

Note that there is no (S,G) entry for the traffic from R2.

Verify the PIM interfaces on R3:

```
R3#show ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
1.1.1.13         Loopback103              v2/S   0      30     1      1.1.1.13
192.168.134.3    Ethernet0/2              v2/S   1      30     1
192.168.134.3
R3#
```

Note that the Ethernet0/3 interface is not in the list of PIM interfaces. Enable PIM on the Ethernet0/3 interface on R3:

```
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#int e0/3
```

```
R3(config-if)#ip pim sparse-mode
R3(config-if)#end
R3#
```

Ping the multicast group 239.255.1.1 from R2 again:

```
R2#ping 239.255.1.1 source e0/3
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.255.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.234.2

Reply to request 0 from 1.1.1.13, 9 ms
R2#
```

Note that R2 receives the reply from R3, but not from R1 and R6.

Verify the mroute table on R1:

```
R1#sh ip mroute 239.255.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.255.1.1), 00:07:22/00:03:13, RP 1.1.1.6, flags: SJCL
  Incoming interface: Ethernet0/0, RPF nbr 192.168.1.1
  Outgoing interface list:
    Ethernet0/2, Forward/Sparse, 00:06:45/00:03:13
    Loopback101, Forward/Sparse, 00:07:21/00:02:43

(192.168.234.2, 239.255.1.1), 00:01:16/00:01:43, flags: LR
  Incoming interface: Ethernet0/0, RPF nbr 192.168.1.1
  Outgoing interface list:
    Loopback101, Forward/Sparse, 00:01:16/00:02:43

R1#
```

Note that R1 does not show the Ethernet0/2 interface in the mroute entries. Is R1 registered with the RP on R6?

```
R6#show ip mroute
IP Multicast Forwarding is not enabled.
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

R6#
```

R6 does not have any entries in its mroute table. Is multicast routing enabled on R6?

This issue is made even clearer by the following Cisco IOS **show** command output:

```
R6#show run | inc multi
ipv6 multicast rpf use-bgp
multilink bundle-name authenticated
R6#
```

Multicast routing is not enabled on R6. Enable multicast routing on R6:

```
R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)#ip multicast-routing
R6(config)#end
R6#

R6#show ip mroute 239.255.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.255.1.1), 00:02:42/00:03:22, RP 1.1.1.6, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 00:02:05/00:03:22
    Loopback106, Forward/Sparse, 00:02:42/00:02:10
```

Ping the multicast group 239.255.1.1 from R2 again:

```
R2#ping 239.255.1.1 source e0/3
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.255.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.234.2

Reply to request 0 from 1.1.1.13, 1 ms
Reply to request 0 from 1.1.1.6, 43 ms
Reply to request 0 from 1.1.1.1, 18 ms
Reply to request 0 from 1.1.1.6, 1 ms
R2#ping 239.255.1.1 source e0/3
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.255.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.234.2

Reply to request 0 from 1.1.1.13, 1 ms
Reply to request 0 from 1.1.1.6, 1 ms
Reply to request 0 from 1.1.1.1, 1 ms
R2#ping 239.255.1.1 source e0/3
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.255.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.234.2

Reply to request 0 from 1.1.1.13, 1 ms
Reply to request 0 from 1.1.1.6, 1 ms
Reply to request 0 from 1.1.1.1, 1 ms
R2#
```


Note To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter more than 1000 Cisco IOS Software commands into the engine, as well as a collection of proprietary commands such as **show all**.