

# Cisco 360 CCIE R&S Exercise Workbook Introduction

---

The Cisco 360 CCIE® R&S Exercise Workbook contains 20 challenging scenarios at the 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 2

### Troubleshooting Section

# Answer Key

---

---

COPYRIGHT 2013, CISCO SYSTEMS, INC. ALL RIGHTS RESERVED. ALL CONTENT AND MATERIALS, INCLUDING WITHOUT LIMITATION, RECORDINGS, COURSE MATERIALS, HANDOUTS AND PRESENTATIONS AVAILABLE ON THIS PAGE, ARE PROTECTED BY COPYRIGHT LAWS. THESE MATERIALS ARE LICENSED EXCLUSIVELY TO REGISTERED STUDENTS FOR THEIR INDIVIDUAL PARTICIPATION IN THE SUBJECT COURSE. DOWNLOADING THESE MATERIALS SIGNIFIES YOUR AGREEMENT TO THE FOLLOWING: (1) YOU ARE PERMITTED TO PRINT THESE MATERIALS ONLY ONCE, AND OTHERWISE MAY NOT REPRODUCE THESE MATERIALS IN ANY FORM, OR BY ANY MEANS, WITHOUT PRIOR WRITTEN PERMISSION FROM CISCO; AND (2) YOU ARE NOT PERMITTED TO SAVE ON ANY SYSTEM, MODIFY, DISTRIBUTE, REBROADCAST, PUBLISH, TRANSMIT, SHARE OR CREATE DERIVATIVE WORKS ANY OF THESE MATERIALS. IF YOU ARE NOT A REGISTERED STUDENT THAT HAS ACCEPTED THESE AND OTHER TERMS OUTLINED IN THE STUDENT AGREEMENT OR OTHERWISE AUTHORIZED BY CISCO, YOU ARE NOT AUTHORIZED TO ACCESS THESE MATERIALS.

---

# Table of Contents

## Cisco 360 CCIE R&S Exercise Workbook Lab 02 Troubleshooting Section Answer Key.....2

Table of Contents .....	3
Answer Key Structure.....	4
Section One .....	4
Section Two .....	4

## Exercise Workbook Lab 02 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. DMVPN Troubleshooting Section .....	8
1.1. Symptom: R2 cannot ping R4. ....	8
2. Switched Network Troubleshooting Section .....	10
2.1. Symptom: The routers connected to SW1 are isolated from the routers connected to SW2.....	12
2.2. Symptom: R6 cannot ping R1 at address 172.16.16.1. ....	15
2.3. Symptom: R6 cannot ping R7 at address 172.16.36.3. ....	16
3. IPv4 OSPF Troubleshooting Section .....	18
3.1. Symptom: No neighbors have been detected and no adjacencies have formed in Area 0. ....	18
3.2. Symptom: R1 and R6 are not adjacent. ....	21
3.3. Symptom: R6 and R7 are not adjacent. ....	22
3.4. Symptom: R7 received multiple OSPF interarea routes in addition to the default. ....	23
3.5. Symptom: The loopback 104 network is advertised as a host route. ....	24
4. EIGRP Troubleshooting Section .....	25
4.1. Symptom: R1 has no EIGRP neighbor relationship over subnet 172.16.12.0/24. ....	25
5. IPv4 RIP Troubleshooting Section.....	26
5.1. Symptom: R4 is learning routes directly from R3.....	26
5.2. Symptom: R3 and R4 have no RIP routes. ....	27
6. IPv4 Redistribution Troubleshooting Section .....	28
6.1. Symptom: Network 172.16.105.0/24 is not in the routing tables of R3 and R4. ....	28
6.2. Symptom: R2 has not learned any routes from R1. ....	29
6.3. Symptom: Connected networks were not redistributed into EIGRP AS12 on R1. ....	30
7. BGP Troubleshooting Section .....	33
7.1. Symptom: R4 is not forming BGP neighbor relationships with any peers.....	33
7.2. Symptom: R5 possesses a BGP prefix that does not contain <i>only</i> the AS path 65001. ....	35
8. IPv6 Troubleshooting Section .....	37
8.1. Symptom: R1 has no OSPFv3 neighbors. ....	37
8.2. Symptom: R1 has only one route to 2001:105::/64. ....	38
8.3. Symptom: R5 does not have complete routes. ....	39
9. IP Quality of Service Troubleshooting .....	40
9.1. Symptom: All DSCP AF11 traffic is being dropped. ....	40
10. IP SLA Troubleshooting Section.....	45
10.1. Symptom: IP SLAs monitor traffic is not collecting statistics.....	45

# 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 2

## 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 credit for a subsection, you must fully complete the subsection per requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the scenario diagram belong to network 135.15.0.0/16.
- *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**).
- Do not use the **ip default-gateway** or **ip default-network** commands.
- Do not introduce any new IP addresses.
- All IP addresses that are 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 must be reachable by all BGP routers but do not need 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 that are displayed in the scenario “Lab IPv4 IGP” diagram belong to network 172.16.0.0/16.**

All IP addresses in this lab belong to the 172.16.0.0/16 address space, except for prefixes that are explicitly specified as being part of a different IP space.

## **Do not use any static routes.**

Static routes can be used to solve a range of reachability problems. However, you cannot use them in this lab. You must rely on skillful configuration of all your unicast routing protocols.

## **Advertise loopback interfaces with their original masks.**

The original mask is the mask configured on the loopback interface. OSPF treats loopback interfaces as host routes by default 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 interface. You need to provide a solution such as changing the OSPF network type or summarizations.

**Network 0.0.0.0/0 should not appear in any routing table (*show ip route*) unless explicitly specified otherwise.**

A 0.0.0.0/0 entry can be used to solve a range of reachability problems. In particular, a 0.0.0.0/0 entry can be used to set up the gateway of last resort. However, in this exercise, you cannot use any 0.0.0.0/0 entries. Route summarization is an alternative to using the 0.0.0.0/0 route to solve the reachability problem is route summarization.

## **Do not use the *ip default-gateway* or *ip default-network* command.**

These commands can be used to solve reachability issues by setting the gateway of last resort. They generate a 0.0.0.0/0 route in the Routing Information Protocol (RIP) environment. You cannot use them in this scenario.

## **All IP addresses that are involved in this scenario must be reachable.**

*This is a key goal to observe.* It requires that all your IGPs and your routing policy tasks 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 ensure that all IP addresses are reachable without the use of static routes or 0.0.0.0/0 routes.

**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 IGP-only routers.**

This statement relaxes the requirement that all IP addresses must be reachable. The BGP prefixes need only be reachable only 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 and to be able to forward traffic to the addresses that are 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 “conventional routing algorithms.” Although this phrase can be interpreted in 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 other type of information other than the destination address to make a packet-forwarding decision.

# 1. DMVPN Troubleshooting Section

## 1.1. Symptom: R2 cannot ping R4.

### *Analysis and Testing:*

Ping the DMVPN IP addresses 172.16.134.3 and 172.16.134.4 from R1:

```
R1#ping 172.16.134.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#ping 172.16.134.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#
```

Note that the DMVPN hub R1 cannot ping the DMVPN spokes R3 and R4.

### *Likely Cause: DMVPN or NHRP is misconfigured.*

Use the **show ip nhrp** command to verify the NHRP registrations on R1:

```
R1#show ip nhrp
172.16.134.3/32
  Tunnel134 created 00:01:58, expire 00:01:06
  Type: incomplete, Flags: negative
  Cache hits: 7
172.16.134.4/32
  Tunnel134 created 00:02:05, expire 00:00:59
  Type: incomplete, Flags: negative
  Cache hits: 7
R1#
```

Note that R3 and R4 are not successfully registered with R1.

Use the **show running-config** command to verify the Tunnel134 configuration on R1:

```
R1#show running-config int tu134
Building configuration...

Current configuration : 209 bytes
!
interface Tunnel134
 ip address 172.16.134.1 255.255.255.0
 no ip redirects
 ip nhrp network-id 134
 ip ospf network non-broadcast
 tunnel source Ethernet0/2
 tunnel mode gre multipoint
 tunnel key 134
end

R1#
```

Note that R1 is configured correctly as according to the lab specifications.

Use the **show ip nhrp** command to verify the NHRP NHS configuration on R3:

```
R3#show ip nhrp nhs
Legend: E=Expecting replies, R=Responding, W=Waiting
Tunnel134:
192.168.1.1    E priority = 0 cluster = 0
R3#
```

Note that the NHRP NHS is configured with the physical (or nonbroadcast multi-access, or NMBA) IP address of the E0/2 interface of R1. NHS must be configured with the virtual IP address of the R1 Tunnel134 interfaces, which is 172.16.134.1.

**Resolution: Fix the NHS configuration and clear the NHRP cache on R3.**

```
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#int tu134
R3(config-if)#no ip nhrp nhs 192.168.1.1
R3(config-if)#ip nhrp nhs 172.16.134.1
R3(config-if)#end
R3#
R3#clear ip nhrp
R3#
```

In a few moments, use the **show ip nhrp** command to verify the NHRP registrations on R1:

```
R1#show ip nhrp
172.16.134.3/32 via 172.16.134.3
  Tunnel134 created 00:04:42, expire 01:57:23
  Type: dynamic, Flags: unique registered used
  NBMA address: 192.168.1.3
172.16.134.4/32
  Tunnel134 created 00:02:40, expire 00:00:24
  Type: incomplete, Flags: negative
  Cache hits: 7
R1#
```

Note that R3 is now properly registered. R4 is still not registered.

Use the **show ip nhrp** command to verify the NHRP NHS configuration on R4:

```
R4#show ip nhrp nhs
Legend: E=Expecting replies, R=Responding, W=Waiting
Tunnel134:
192.168.1.1    E priority = 0 cluster = 0
R4#
```

Note that R4 also is configured with the wrong NHS IP address. Fix the NHS configuration on R4 and clear the NHRP cache:

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int tu134
R4(config-if)#no ip nhrp nhs 192.168.1.1
R4(config-if)#ip nhrp nhs 172.16.134.1
R4(config-if)#end
R4#
R4#clear ip nhrp
R4#
```

Use the **show ip nhrp** command to verify the NHRP registrations on R1:

```
R1#show ip nhrp
172.16.134.3/32 via 172.16.134.3
```

```
Tunnel134 created 00:18:40, expire 01:43:24
Type: dynamic, Flags: unique registered used
NBMA address: 192.168.1.3
172.16.134.4/32 via 172.16.134.4
Tunnel134 created 00:03:25, expire 01:59:39
Type: dynamic, Flags: unique registered used
NBMA address: 192.168.1.4
R1#
```

Ping the DMVPN IP addresses 172.16.134.3 and 172.16.134.4 from R1 again:

```
R1#ping 172.16.134.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R1#ping 172.16.134.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.4, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R1#
```

Verify the spoke to spoke connectivity from R3:

```
R3#ping 172.16.134.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.134.4, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R3#
```

Not that the ping is now successful.

---

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

---

## 2. Switched Network Troubleshooting Section

Your initial testing of Ethernet same-subnet reachability shows the following:

- There is no reachability between R1 and R2 on VLAN 12.
- There is no reachability between R1 and R6 on VLAN 16.

```
R1#ping 172.16.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.12.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#ping 172.16.16.6
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.16.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#
```

- R4 cannot reach R3 or R5 on VLAN 34, but R3 and R5 can communicate.

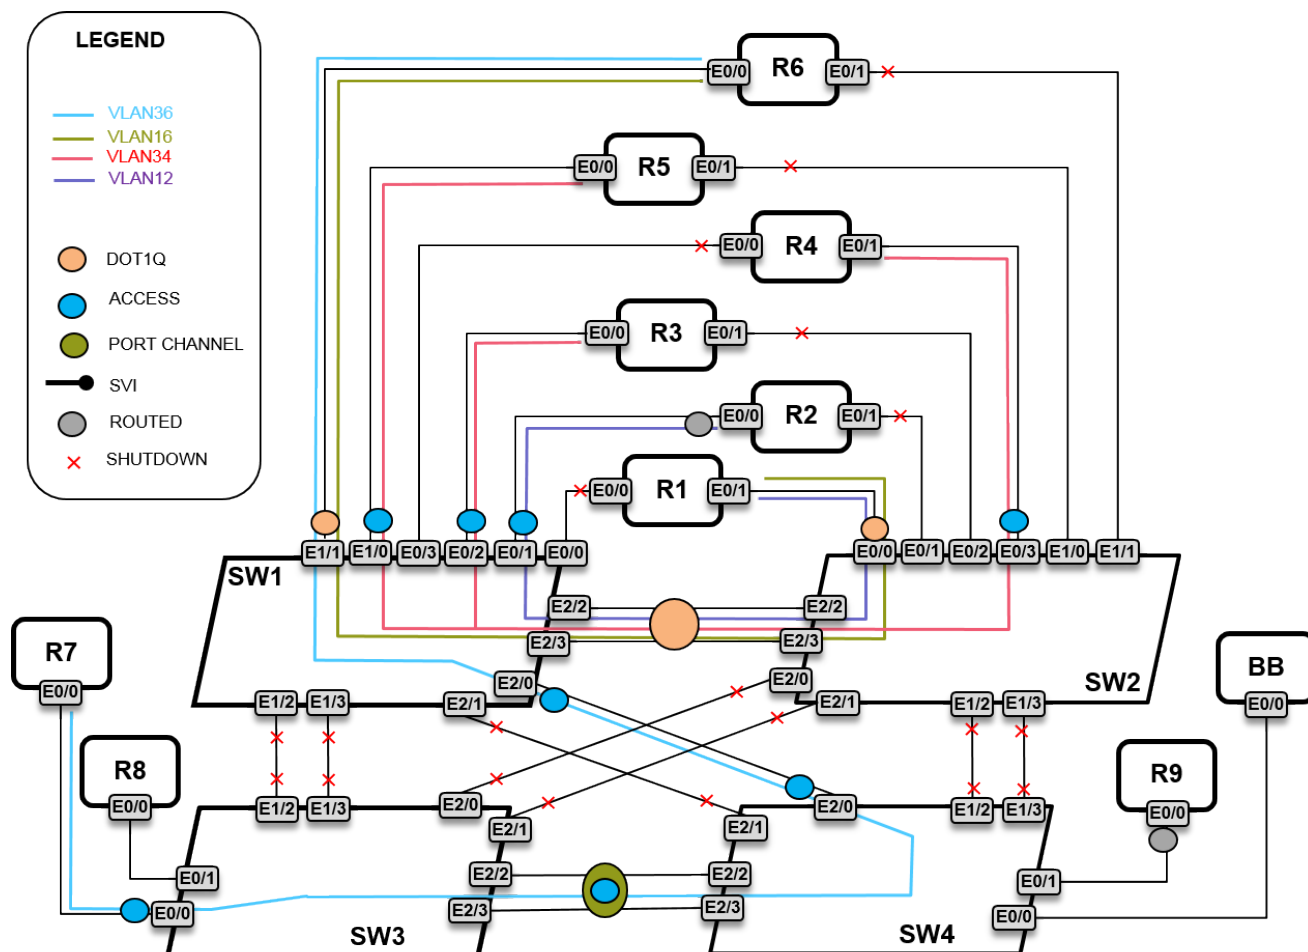
```
R4#ping 172.16.34.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.34.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R4#ping 172.16.34.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.34.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R4#
R3#ping 172.16.34.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.34.5, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R3#
```

- R6 cannot reach R7 on VLAN 36.

```
R6#ping 172.16.36.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.36.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R6#
```

The “VLAN Propagation” diagram is reproduced here for easy reference. Notice that VLANs 12, 16, and 34 all cross between SW1 and SW2 on links E2/2 and E2/3. If these links were broken, we would know why R1 and R4 are isolated from the other Ethernet devices.

## VLAN Propagation Diagram



### 2.1. Symptom: The routers connected to SW1 are isolated from the routers connected to SW2.

#### Analysis and testing:

The “VLAN Propagation” diagram indicates that E2/2 and E2/3 between SW1 and SW2 must be dot1q trunks. The command **show interfaces trunk** can help you verify that this requirement has been met.

```
SW1#show interfaces trunk

Port      Mode      Encapsulation  Status        Native vlan
Et1/1     on        802.1q         trunking      1

Port      Vlans allowed on trunk
Et1/1     16,36

Port      Vlans allowed and active in management domain
Et1/1     16,36

Port      Vlans in spanning tree forwarding state and not pruned
Et1/1     16,36
SW1#
```

Use the **show interfaces** command to verify the configuration of the E2/2 and E2/3 interfaces on SW1:

```
SW1#show interfaces e2/2 switchport
Name: Et2/2
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Operational Dot1q Ethertype: 0x8100
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Operational Native VLAN tagging: disabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: 12,16,34
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
```

```
SW1#show interfaces e2/3 switchport
Name: Et2/3
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Operational Dot1q Ethertype: 0x8100
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Operational Native VLAN tagging: disabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: 12,16,34
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
```

SW1#

Note that the E2/2 and E2/3 interfaces are configured with the dot1q trunk encapsulation, but both interfaces are statically placed in the access VLAN configuration.

**Likely cause: Interfaces E2/2 and E2/3 are incorrectly configured on SW1.**

Use the **show running-config** command to verify the configuration of these ports on SW1.

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

Current configuration : 144 bytes
!
```

```

interface Ethernet2/2
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 12,16,34
  switchport mode access
  duplex auto
end

```

```

SW1#show running-config int e2/3
Building configuration...

```

```

Current configuration : 144 bytes
!
interface Ethernet2/3
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 12,16,34
  switchport mode access
  duplex auto
end

```

```

SW1#

```

Here, you learn that the **switchport mode access** command is configured instead of the **switchport mode trunk** command.

**Resolution: Configure the switchport mode trunk command on the E2/2 and E2/3 interfaces on SW1:**

```

SW1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)#int range e2/2 - 3
SW1(config-if-range)#switchport mode trunk
SW1(config-if-range)#end
SW1#

```

Use the **show interfaces** command to verify trunks on SW1 again:

```

SW1#show interfaces trunk

```

Port	Mode	Encapsulation	Status	Native vlan
Et1/1	on	802.1q	trunking	1
Et2/2	on	802.1q	trunking	1
Et2/3	on	802.1q	trunking	1

```

Port
Et1/1
Et2/2
Et2/3

```

Port	Vlans allowed on trunk
Et1/1	16,36
Et2/2	12,16,34
Et2/3	12,16,34

```

Port
Et1/1
Et2/2
Et2/3

```

Port	Vlans allowed and active in management domain
Et1/1	16,36
Et2/2	12,16,34
Et2/3	12,16,34

```

Port
Et1/1
Et2/2
Et2/3
SW1#

```

If this is the entire cause of connectivity problems on VLANs 12, 16, and 36, then R1 should be able to ping R2 and R6, and R4 should be able to ping R3 and R4:

```

R1#ping 172.16.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.12.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

```

```
R1#ping 172.16.16.6
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.16.6, timeout is 2 seconds:
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

```
R4#ping 172.16.34.3
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.34.3, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

```
R4#ping 172.16.34.5
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.34.5, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
R4#
```

Connectivity on VLANs 16 and 34 is complete, but R1 still cannot ping R6.

## 2.2. Symptom: R6 cannot ping R1 at address 172.16.16.1.

### *Analysis and testing:*

Refer to the “VLAN Propagation” diagram to verify the configuration at each device and interface between R6 and R1, starting at R6. The “IPv4 IGP” diagram shows IP address 172.16.16.6 configured on interface E0/0 of R6. Here is the physical interface configuration:

```
R6#show running-config int e0/0
```

```
Building configuration...
```

```
Current configuration : 99 bytes
```

```
!
```

```
interface Ethernet0/0
```

```
 ip address 172.16.16.6 255.255.255.0
```

```
 ip ospf network point-to-point
```

```
end
```

```
R6#
```

This interface appears to be configured correctly. But how will these frames become associated with VLAN 16? Because this is a physical interface, the frames will be sent as untagged Ethernet frames.

Examine the next hop in the path, SW1 interface E1/1. Here is the E1/1 trunking information:

```
SW1#show interfaces e1/1 trunk
```

Port	Mode	Encapsulation	Status	Native vlan
Et1/1	on	802.1q	trunking	1

Port	Vlans allowed on trunk
Et1/1	16,36

Port	Vlans allowed and active in management domain
Et1/1	16,36

Port	Vlans in spanning tree forwarding state and not pruned
Et1/1	16,36

```
SW1#
```

The interface is a dot1q trunk, as required, and it does allow VLAN 16. But the native VLAN is still the default native VLAN, not VLAN 16, as would be required. Frames that are intended for VLAN 16 are instead being sent into VLAN 1.

**Likely cause:** *The native VLAN on SW1 interface E1/1 is incorrect.*

If the incorrect native VLAN is the only cause of the connectivity problem between R1 and R6, then changing the native VLAN on this port to VLAN 16 will fix the problem.

**Resolution:** *Configure VLAN 16 as the native VLAN on SW1 interface E1/1.*

```
SW1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)#interface e1/1
SW1(config-if)#switchport trunk native vlan 16
```

Now try a ping to R1 address 172.16.16.1:

```
R6#ping 172.16.16.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.16.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
R6#
```

This change probably had no effect on the inability of R6 to reach R7 address 172.16.36.3, but try again:

```
R6#ping 172.16.36.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.36.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Connectivity on VLANs 16 and 34 is complete, but R1 still cannot ping R7.

### 2.3. Symptom: R6 cannot ping R7 at address 172.16.36.3.

**Analysis and testing:**

Perform an analysis between the Layer 3 interfaces involved while also checking the Layer 2 configuration. Here is the configuration of E0/0.3 on R6:

```
interface Ethernet0/0.3
 encapsulation dot1q 36
 ip address 172.16.36.6 255.255.255.0
```

This configuration looks correct. It is encapsulated for dot1q VLAN 36, and it appears to have the IP address that is specified on the “Layer 3” diagram.

Enter the **show ip interface brief** command to verify that the interface is active:

```
R6#show ip interface brief
Interface                               IP-Address      OK?  Method  Status
Protocol
Ethernet0/0                             172.16.16.6     YES NVRAM  up
Ethernet0/0.3                           172.16.36.6     YES NVRAM  up
[output removed for brevity]
```

This output looks correct; the status of both interfaces is “up/up.” The “VLAN Propagation” diagram indicates that E1/1 on SW1 is the next interface in the path toward R7. You verified earlier that this interface was a dot1q trunk and that it allowed VLAN 36.

Use the **show vlan** command to verify that VLAN 36 is active on SW1 and that it is configured on access link E2/0, as the “VLAN Propagation” diagram shows.

```
SW1#show vlan | include 36
36    VLAN0036                active    Et2/0
```

The output of the **show vlan** command on SW1 clearly indicates that SW1 is correctly configured to support traffic between R6 and R7 on VLAN 36. According to the diagram, the next device in the path is SW4, and it should have two interfaces configured as access links in VLAN 36: E2/0 and a port channel interface.

Here is the output from the **show vlan** command for SW4:

```
SW4#show vlan | include 36
36 VLAN0036 active Et2/0, Et2/2, Et2/3
```

This output is not exactly what you were looking for. You should see a port channel interface here, not the physical ports. Therefore, we must verify that there is an EtherChannel on SW4:

```
SW4#show etherchannel summary
Flags: D - down P - bundled in port-channel
       I - stand-alone s - suspended
       H - Hot-standby (LACP only)
       R - Layer3 S - Layer2
       U - in use N - not in use, no aggregation
       f - failed to allocate aggregator

       M - not in use, no aggregation due to minimum links not met
       m - not in use, port not aggregated due to minimum links not met
       u - unsuitable for bundling
       d - default port

       w - waiting to be aggregated
Number of channel-groups in use: 1
Number of aggregators: 1

Group  Port-channel  Protocol    Ports
-----+-----+-----+-----
1      Po1 (SD)        -           Et2/2 (D)  Et2/3 (D)
```

SW4#

You may be close to the cause. SW4 is configured for an EtherChannel, but the ports are in a down state.

Here is the output from the **show interface status** command for these ports on SW4:

```
SW4#show interface status

Port      Name           Status      Vlan    Duplex  Speed  Type
Et0/0     Et0/0          connected   1       auto    auto   unknown
Et0/1     Et0/1          connected   1       auto    auto   unknown
Et0/2     Et0/2          connected   1       auto    auto   unknown
Et0/3     Et0/3          connected   1       auto    auto   unknown
Et1/0     Et1/0          connected   1       auto    auto   unknown
Et1/1     Et1/1          disabled    1       auto    auto   unknown
Et1/2     Et1/2          disabled    1       auto    auto   unknown
Et1/3     Et1/3          disabled    1       auto    auto   unknown
Et2/0     Et2/0          connected   36      auto    auto   unknown
Et2/1     Et2/1          disabled    1       auto    auto   unknown
Et2/2     Et2/2          disabled    36      auto    auto   unknown
Et2/3     Et2/3          disabled    36      auto    auto   unknown
Po1       Po1            disabled    36      auto    auto
```

SW4#

E2/2 and E2/3 have been disabled on SW4.

**Likely cause: R6 cannot reach R7 on VLAN 36 because ports of the EtherChannel are disabled .**

```
SW4#show ip interface brief | inc Port
Port-channell1 unassigned YES unset administratively down down
SW4#
```

You can re-enable the ports by issuing the command **no shutdown**.

```
SW4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW4(config)#int pol
SW4(config-if)#no shut
SW4(config-if)#
*Apr 25 21:24:25.655: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
SW4(config-if)#en
*Apr 25 21:24:25.659: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-
channel1, changed state to up
SW4(config-if)#end
SW4#
*Apr 25 21:24:27.183: %SYS-5-CONFIG_I: Configured from console by console
(cierswbv5-te-lab02-sc, SJ)
SW4#
*Apr 25 21:24:27.639: %LINK-3-UPDOWN: Interface Ethernet2/2, changed state to up
*Apr 25 21:24:27.639: %LINK-3-UPDOWN: Interface Ethernet2/3, changed state to up
*Apr 25 21:24:28.643: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet2/2,
changed state to up
*Apr 25 21:24:28.643: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet2/3,
changed state to up
SW4#
```

The interface is “up/up,” try to ping R7 from R6 again.

```
R6#ping 172.16.36.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.36.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R6#
```

Same-subnet IPv4 unicast reachability appears to be complete for the Ethernet. All Layer 3 interfaces can ping the IPv4 IP addresses that are on the same subnet. Next, verify the performance of the routing protocols, starting with OSPF.

- 
- Note** The Mentor Guide engine in the web portal can help you use Cisco IOS Software commands to see a comprehensive view of the configuration for a specific section. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands as well as a collection of proprietary commands such as **show all**.
- Note** To learn more about Cisco Catalyst switch troubleshooting methods and techniques, download and watch the VoD sessions from the Cisco 360 “Troubleshooting” lesson module. This lesson module contains more than 8 hours of video content that is dedicated to the subject of troubleshooting.
- 

### 3. IPv4 OSPF Troubleshooting Section

#### 3.1. Symptom: No neighbors have been detected and no adjacencies have formed in Area 0.

##### *Analysis and testing:*

A suggested general approach to troubleshooting OSPF adjacencies begins by verifying that the correct interfaces have been associated with the correct areas. A command such as **show ip ospf interface** can provide a quick answer.

Once the OSPF interfaces are verified, verify support for the selected OSPF network types. In this case, all three of the interfaces on subnet 172.16.134.0/24 have been associated with area 0 with the appropriate network statements. Here is partial output from the command **show ip ospf interface tu134** on each router:

```

R1#show ip ospf interface tu134
Tunnel134 is up, line protocol is up
  Internet Address 172.16.134.1/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.16.101.1, Network Type NON_BROADCAST, Cost: 1000
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                1000        no            no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.101.1, Interface address 172.16.134.1
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:10
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/3, 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)
R1#
R3#show ip ospf interface t134
Tunnel134 is up, line protocol is up
  Internet Address 172.16.134.3/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.16.103.1, Network Type NON_BROADCAST, Cost: 1000
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                1000        no            no            Base
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  No designated router on this network
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:07
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2, 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)
R3#
R4#show ip ospf interface t134
Tunnel134 is up, line protocol is up
  Internet Address 172.16.134.4/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.16.104.1, Network Type NON_BROADCAST, Cost: 1000
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                1000        no            no            Base
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  No designated router on this network
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:25
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2, 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#
```

This output shows that all interfaces are using the OSPF nonbroadcast network type. Furthermore, it shows that R1 is the designated router (DR), that the two other routers are in a DROTHER state, and that the hello and dead times are all equal, as is as expected. Because the routers have not detected any neighbors, a possible issue is that hello packets are not being exchanged.

The OSPF nonbroadcast network type unicasts all OSPF packets, including hellos. Before these unicasts can be formed, OSPF requires neighbor statements on a router that is eligible to become a DR. Therefore, R1 must be checked to see if it has been configured with OSPF neighbor statements.

**Likely cause: R1 has not been configured with OSPF neighbor statements.**

A debug of OSPF hello packets will indicate whether hello packets are being sent and received. But take the direct approach here and simply search the configuration of R1 for a neighbor statement:

```
R1#show run | section router ospf
router ospf 1
 log-adjacency-changes
 area 16 virtual-link 172.16.106.1
 redistribute connected subnets route-map CONNECTED
 redistribute eigrp 12 subnets
 network 172.16.134.1 0.0.0.0 area 0
 network 172.16.16.1 0.0.0.0 area 16
```

Missing configuration statements can be difficult to identify, but in this case you know exactly what you are looking for and you do not find it. IPv4 OSPF network statements must be configured under the process, and evidently they are not there.

**Resolution: Under the OSPF process on R1, configure OSPF neighbor statements to R3 and R4.**

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#router ospf 1
R1(config-router)#neighbor 172.16.134.3
R1(config-router)#neighbor 172.16.134.4
R1(config-router)#end
```

In a few moments, use the **show ip ospf neighbor** command to verify the OSPF neighbors on R1:

```
R1#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address         Interface
172.16.103.1     0     FULL/DROTHER    00:01:48   172.16.134.3   Tunnel134
172.16.104.1     0     FULL/DROTHER    00:01:53   172.16.134.4   Tunnel134
R1#
```

After the neighbor statements are configured, adjacencies form with R3 and R4. You now see OSPF routes in the routing table of R1:

```
R1#show ip route ospf
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

```

172.16.0.0/16 is variably subnetted, 14 subnets, 2 masks
O E2    172.16.34.0/24 [110/20] via 172.16.134.4, 00:01:05, Tunnel134
        [110/20] via 172.16.134.3, 00:01:05, Tunnel134
O IA    172.16.103.0/24 [110/1001] via 172.16.134.3, 00:01:05, Tunnel134
O IA    172.16.104.0/24 [110/1001] via 172.16.134.4, 00:01:05, Tunnel134
O E2    172.16.105.0/24 [110/20] via 172.16.134.4, 00:01:05, Tunnel134
        [110/20] via 172.16.134.3, 00:01:05, Tunnel134
R1#

```

R1 has OSPF neighbors in area 0, but it does not see R6 as an OSPF neighbor on its interface E0/1.6.

### 3.2. Symptom: R1 and R6 are not adjacent.

#### Analysis and testing:

Here is partial output from the command **debug ip ospf hello** on R1:

```

*
*Apr 25 22:54:12.069: OSPF-1 HELLO Et0/1.6: Rcv hello from 172.16.106.1 area 16
172.16.16.6
*Apr 25 22:54:12.069: OSPF-1 HELLO Et0/1.6: Mismatched hello parameters from
172.16.16.6
*Apr 25 22:54:12.069: OSPF-1 HELLO Et0/1.6: Dead R 40 C 120, Hello R 10 C 30

```

This output clearly indicates that there is a timer mismatch between R1 and R6. This mismatch is most commonly caused by mismatching OSPF network types, so verify the operation of OSPF on the two interfaces:

```

R1#show ip ospf interface E0/1.6
Ethernet0/1.6 is up, line protocol is up
 Internet Address 172.16.16.1/24, Area 16, Attached via Network Statement
 Process ID 1, Router ID 172.16.101.1, Network Type POINT_TO_MULTIPOINT, Cost: 10
 Topology-MTID    Cost    Disabled    Shutdown    Topology Name
      0            10         no          no          Base
 Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT
 Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
 [output removed for brevity]

```

```

R6#show ip ospf interface E0/0
Ethernet0/0 is up, line protocol is up
 Internet Address 172.16.16.6/24, Area 16, Attached via Network Statement
 Process ID 1, Router ID 172.16.106.1, Network Type POINT_TO_POINT, Cost: 10
 Topology-MTID    Cost    Disabled    Shutdown    Topology Name
      0            10         no          no          Base
 Transmit Delay is 1 sec, State POINT TO POINT
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 [output removed for brevity]

```

#### Likely cause: OSPF network type mismatch.

You could easily fix this mismatch by changing the OSPF network type, but this option has been forbidden by the lab restrictions. Because both of these network types model the link as point-to-point, their use will not result in the kind of OSPF database inconsistency that would result from mixing, for example, point-to-point and broadcast network types.

These network types are compatible, so all you need to do is make sure that the timers match. Change the hello time to 10 seconds on R1. The dead time will automatically adjust to four times the hello interval.

**Resolution: Change the OSPF hello time to 10 seconds on R1 E0/1.6.**

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#int E0/1.6
R1(config-subif)#ip ospf hello 10
R1(config-subif)#end

R1#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
172.16.103.1     0    FULL/DROTHER    00:01:55   172.16.134.3   Tunnel134
172.16.104.1     0    FULL/DROTHER    00:01:55   172.16.134.4   Tunnel134
172.16.106.1     0    FULL/ -         00:00:27   172.16.16.6    OSPF VL0
172.16.106.1     0    FULL/ -         00:00:36   172.16.16.6    Ethernet0/1.6
R1#
```

The only adjacency left to test is the one between R6 and R7 in OSPF area 36. You have verified that the required interfaces are associated with the area, and you know that they have unicast reachability, yet R7 shows blank output when you check for OSPF neighbors.

### 3.3. Symptom: R6 and R7 are not adjacent.

#### *Analysis and testing:*

The lack of any type of neighbor entry indicates that R7 is not receiving acceptable hello packets from R6. Here is partial output from the **debug ip ospf hello** command on R7:

```
R7#
*Apr 25 23:14:32.669: OSPF-1 HELLO Et0/0: Rcv hello from 172.16.106.1 area 36
172.16.36.6
*Apr 25 23:14:32.669: OSPF-1 HELLO Et0/0: Hello from 172.16.36.6 with mismatched
Stub/Transit area option bit
R7#
```

Apparently the hello packets from R6 are rejected because of a mismatched stub flag.

**Likely cause: Area 36 stub flags differ on R6 and R7.**

Look at the area configuration:

```
R7#show run | include area
network 172.16.36.3 0.0.0.0 area 36

R6#show run | include area
area 16 virtual-link 172.16.111.1
area 36 stub
network 172.16.36.6 0.0.0.0 area 36
network 172.16.16.6 0.0.0.0 area 16
```

The stub feature is intended to minimize the OSPF use of router resources by filtering unnecessary routes. This configuration appears to be an attempt to meet the requirement that R6 advertise only a single route to R7. However, all routers must have the OSPF stub bit set, and every router in an area must have the same stub type; you cannot mix the not-so-stubby area (NSSA) and the stub. Therefore, set the stub bit on R7 by entering the **area 36 stub** command.

**Resolution: On R7, enter the area 36 stub command.**

```
R7#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R7(config)#router ospf 1
```

```
R7(config-router)#area 36 stub
R7(config-router)#end
```

```
R7#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.106.1	1	FULL/DR	00:00:37	172.16.36.6	Ethernet0/0

```
R7#
```

```
R7#
```

```
R7#show ip route ospf
```

```
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 172.16.36.6 to network 0.0.0.0
```

```
O*IA 0.0.0.0/0 [110/11] via 172.16.36.6, 00:00:40, Ethernet0/0
      172.16.0.0/16 is variably subnetted, 7 subnets, 2 masks
O IA   172.16.16.0/24 [110/20] via 172.16.36.6, 00:00:40, Ethernet0/0
O IA   172.16.16.1/32 [110/20] via 172.16.36.6, 00:00:40, Ethernet0/0
O IA   172.16.103.0/24 [110/1021] via 172.16.36.6, 00:00:40, Ethernet0/0
O IA   172.16.104.0/24 [110/1021] via 172.16.36.6, 00:00:40, Ethernet0/0
O IA   172.16.134.0/24 [110/1020] via 172.16.36.6, 00:00:40, Ethernet0/0
R7#
```

R7 now has a full adjacency with R6. However, R7 has learned multiple routes from R6: the interarea default and the type 7 external route, 172.16.106.0/24. The **no-summary** keyword on R6 caused R6 to send the default route to R7. Network policies require that R6 advertise only the default route.

### 3.4. Symptom: R7 received multiple OSPF interarea routes in addition to the default.

#### *Analysis and testing:*

Why is R6 advertising multiple routes to R7, and how can this advertising be stopped? Network policies restrict the use of filters on R6, and creating a filter on R7 would not keep R6 from sending the information. You know that the OSPF stub feature filters type 5 link-state advertisements (LSAs), so R6 would not advertise any of its E1 or E2 routes into area 36. However, the stub feature alone does permit the advertisement of OSPF interarea routes.

**Likely cause:** *By default, the OSPF stub feature does not filter interarea routes.*

OSPF interarea routes can be filtered by adding the **no-summary** keyword to the OSPF stub configuration command **area 36 stub** on the Area Border Router (ABR) R6.

**Resolution:** *On R6, add the no-summary keyword to the area 36 stub command.*

```
R6#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R6(config)#router ospf 1
R6(config-router)#area 36 stub no-summary
```

The following output shows the result on R7:

```
R7#show ip route ospf
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 172.16.36.6 to network 0.0.0.0

```
O*IA 0.0.0.0/0 [110/11] via 172.16.36.6, 00:00:10, Ethernet0/0
```

R7#

As required by network policies, R7 will now have full reachability to the remainder of the network using just a single, learned route.

All your routers now have the required adjacencies. Here are the current OSPF routes on R1:

```
R1#show ip route ospf
```

```

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

```

172.16.0.0/16 is variably subnetted, 15 subnets, 2 masks
O E2 172.16.34.0/24 [110/20] via 172.16.134.4, 00:37:41, Tunnel134
      [110/20] via 172.16.134.3, 00:37:41, Tunnel134
O IA 172.16.36.0/24 [110/20] via 172.16.16.6, 00:24:51, Ethernet0/1.6
O IA 172.16.103.0/24 [110/1001] via 172.16.134.3, 00:37:41, Tunnel134
O IA 172.16.104.1/32 [110/1001] via 172.16.134.4, 00:00:07, Tunnel134
O E2 172.16.105.0/24 [110/20] via 172.16.134.4, 00:37:41, Tunnel134
      [110/20] via 172.16.134.3, 00:37:41, Tunnel134
O E2 172.16.106.0/24 [110/20] via 172.16.16.6, 00:25:01, Ethernet0/1.6
R1#

```

When you focus on just the internal OSPF routes, you notice that the loopback 104 subnet is not advertised with its original mask, as required by network policies. It is advertised as a host route instead of as a /24 network.

### 3.5. Symptom: The loopback 104 network is advertised as a host route.

#### *Analysis and testing:*

You recall that OSPF uses a special network type of loopback for subnets that are configured on loopback interfaces, which causes OSPF to advertise the subnet as a host route rather than with its original mask.

#### *Likely cause: The OSPF network type remains loopback.*

The typical practice is to avoid this symptom by changing the OSPF network type to point-to-point. Check whether this change has been made on R4:

```

interface Loopback104
 ip address 172.16.104.1 255.255.255.0
 ipv6 address 2001:104::1/64
 ipv6 ospf 1 area 104

```

```
Loopback104 is up, line protocol is up
```

```
Internet Address 172.16.104.1/24, Area 104, Attached via Network Statement
Process ID 1, Router ID 172.16.104.1, Network Type LOOPBACK, Cost: 1
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
   0                1          no            no            Base
```

Loopback interface is treated as a stub Host

Network policies do not permit you to change the existing OSPF network type, so how can you fix this situation within the rules? Notice that loopback 104 is in its own area 104, which gives you the option of summarizing this network to the required /24 prefix length.

**Resolution: On R4, create an interarea summary to 172.16.104.0/24.**

```
R4(config)#router ospf 1
R4(config-router)#area 104 range 172.16.104.0 255.255.255.0
```

```
R1#show ip route 172.16.104.1
Routing entry for 172.16.104.0/24
  Known via "ospf 1", distance 110, metric 65, type inter area
```

R1 now shows the loopback network with the required /24 mask. Now that OSPF appears to be operating as required, verify the Enhanced Interior Gateway Routing Protocol (EIGRP).

---

**Note** To learn more about OSPF troubleshooting methods and techniques, download and watch the VoD sessions from the Cisco 360 "Troubleshooting" lesson module. This lesson module contains more than 8 hours of video content that is dedicated to the subject of troubleshooting.

---

## 4. EIGRP Troubleshooting Section

### 4.1. Symptom: R1 has no EIGRP neighbor relationship over subnet 172.16.12.0/24.

#### *Analysis and testing:*

R1 has no EIGRP neighbor relationship with R2. Start your analysis with subnet 172.16.12.0/24. Network policies require that all EIGRP traffic be unicast on this link.

Check the configuration on both R1 and R2 to verify correct neighbor statements:

```
R1#sh run | section eigrp
router eigrp 12
 redistribute connected route-map CONNECTED
 redistribute ospf 1
 network 172.16.12.0 0.0.0.255
 auto-summary
 neighbor 172.16.12.2 Ethernet0/1.2
```

```
R2#show run | section eigrp
router eigrp 12
 network 172.16.2.0 0.0.0.255
 network 172.16.102.0 0.0.0.255
 auto-summary
```

Here, you see that R1 has the required neighbor statement but R2 does not.

**Likely cause: R2 does not have the required neighbor statement.**

The following output shows the result when you add the neighbor statement on R2. It has no effect.

```
R2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#router eigrp 12
R2(config-router)#neighbor 172.16.12.1 E0/0
R2(config-router)#end
```

```
R2#sh run | se eigrp
router eigrp 12
network 172.16.2.0 0.0.0.255
network 172.16.102.0 0.0.0.255
auto-summary
```

If you look more closely, you see that the root of the issue is not the missing neighbor statement; it is that EIGRP has not been activated at all on the required interface. Note that E0/0 is missing from the output:

```
R2#show ip eigrp interface
EIGRP-IPv4 Interfaces for AS(12)
      Xmit Queue   Mean   Pacing Time   Multicast   Pending
Interface   Peers  Un/Reliable  SRTT  Un/Reliable  Flow Timer  Routes
Lo102             0      0/0         0      0/1         0           0
```

The root cause of the problem appears to be an incorrect network statement. Note in the configuration shown that there is no network statement for subnet 172.16.12.0, but there is one for the nonexistent network 172.16.2.0.

**Resolution:** On R2, configure the correct network and neighbor statements.

```
R2(config)#router eigrp 12
R2(config-router)#no network 172.16.2.0 0.0.0.255
R2(config-router)#network 172.16.12.0 0.0.0.255
```

Now go back to R1 and check for an EIGRP neighbor across this subnet and for learned routes:

```
R1#show ip eigrp neighbor
EIGRP-IPv4 Neighbors for AS(12)
H   Address                Interface      Hold Uptime    SRTT   RTO  Q  Seq
                               (sec)          (ms)          Cnt  Num
0   172.16.12.2              E0/1.2        11 00:03:47    1    200  0  1423
```

```
R1#show ip route eigrp
172.16.0.0/24 is subnetted, 13 subnets
D    172.16.102.0 [90/156160] via 172.16.12.2, 00:03:53, Ethernet0/1.2
```

R1 now has a neighbor relationship across subnet 172.16.12.0, and it has learned a route to the loopback network on R2.

---

**Note** To learn more about EIGRP troubleshooting methods and techniques, download and watch the VoD sessions from the Cisco 360 “Troubleshooting” lesson module. This lesson module contains more than 8 hours of video content that is dedicated to the subject of troubleshooting.

---

## 5. IPv4 RIP Troubleshooting Section

### 5.1. Symptom: R4 is learning routes directly from R3.

#### *Analysis and testing:*

Network policies require that all RIP updates on subnet 172.16.34.0/24 be unicast through R5. However, the partial output of the **show ip protocols** command on R4 indicates that it has received updates from R3:

```
R4#show ip protocols
[output removed for brevity]
Routing Information Sources:
  Gateway         Distance      Last Update
  172.16.34.5     120           00:00:19
  172.16.34.3     120           00:00:01
```

The log of the **debug ip rip** command on R3 indicates that it is sending multicast updates in addition to the required unicast updates:

```
R3#show log | include sending
*Mar  9 16:37:43.790: RIP:  sending v2 update to 224.0.0.9 via Ethernet0/0
(172.16.34.3)
*Mar  9 16:37:43.790: RIP:  sending v2 update to 172.16.34.5 via Ethernet0/0
(172.16.34.3)
```

**Likely cause: R3 is sending multicast updates onto the subnet.**

RIPv2 can be kept from sending multicast updates in different ways, but the simplest and most direct way is to make the interface passive. To determine if this option has been applied, review the configuration. But the direct answer as to which RIP protocols are passive can be found in the output of the **show ip protocols** command:

```
R3#show ip protocols | begin Passive
Passive Interface(s):
 Ethernet0/1
 Loopback103
 VoIP-Null0
```

Notice that E0/0 is not included on this list.

**Resolution: On R3, add E0/0 to the list of RIP passive interfaces.**

```
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#router rip
R3(config-router)#passive-interface E0/0
R3(config-router)#do show ip protocols | begin Passive
Passive Interface(s):
 Ethernet0/0
 Ethernet0/1
 Loopback103
```

Now that RIP is sending only unicast updates, check that they are being propagated through R5 as required. When you check the routing tables, you see the following pattern: R5 has learned routes from both R3 and R4, but neither R3 nor R4 has any RIP routes in its routing tables.

## 5.2. Symptom: R3 and R4 have no RIP routes.

**Analysis and testing:**

RIP and OSPF are mutually redistributed on both R3 and R4. To isolate the performance of RIP itself, temporarily shut the Ethernet0/2 interfaces on R3 and R4. Here are the routing tables on R3 and R4 after the shutdown:

```
R3#show ip route rip
172.16.0.0/24 is subnetted, 3 subnets
R    172.16.105.0 [120/1] via 172.16.34.5, 00:00:11, Ethernet0/

R4#show ip route rip
172.16.0.0/24 is subnetted, 3 subnets
R    172.16.105.0 [120/1] via 172.16.34.5, 00:00:26, Ethernet0/1

R5#show ip route rip
172.16.0.0/24 is subnetted, 4 subnets
R    172.16.104.0 [120/1] via 172.16.34.4, 00:00:22, Ethernet0/0
R    172.16.103.0 [120/1] via 172.16.34.3, 00:00:18, Ethernet0/0
```

R5 has learned network 172.16.103.0/24 from R3 but does not advertise it to R4. R5 has learned network 172.16.104.0/24 from R4 but does not advertise it to R3.

R5 apparently is not advertising routes that it learns on its interface E0/0 back out E0/0 because of the operation of split horizon. Is split horizon enabled on R5 interface E0/0?

**Likely cause: IP split horizon has not been disabled on R5 interface E0/0.**

```
R5#show ip interface E0/0 | inc horizon
Split horizon is enabled
```

The answer is that split horizon has not been disabled.

**Resolution: On R5 interface E0/0, disable IP split horizon.**

```
R5(config)#int E0/0
R5(config-if)#no ip split-horizon
```

The impact of this change on the routing tables of R3 and R4 validates your analysis:

```
R3#show ip route rip
172.16.0.0/24 is subnetted, 4 subnets
R    172.16.104.0 [120/2] via 172.16.34.4, 00:00:22, Ethernet0/0
R    172.16.105.0 [120/1] via 172.16.34.5, 00:00:22, Ethernet0/0
```

```
R4#show ip route rip
172.16.0.0/24 is subnetted, 4 subnets
R    172.16.105.0 [120/1] via 172.16.34.5, 00:00:05, Ethernet0/1
R    172.16.103.0 [120/2] via 172.16.34.3, 00:00:05, Ethernet0/1
```

When you reactivate the Ethernet0/2 interfaces on R3 and R4, you find that all these RIP routes disappear. Read the “IPv4 Redistribution Troubleshooting Section” to troubleshoot the IPv4 redistribution.

## 6. IPv4 Redistribution Troubleshooting Section

### 6.1. Symptom: Network 172.16.105.0/24 is not in the routing tables of R3 and R4.

#### *Analysis and testing:*

You are required to mutually redistribute between OSPF and RIP on R3 and R4. Optimal path requirements mean that the networks 172.16.105.0/24, 172.16.103.0/24, and 172.16.104.0/24 must be represented in the routing tables of R3 and R4 as RIP routes.

What happened to these routes? Here is partial output from the **debug ip routing** command from R3:

```
*Mar  9 17:23:48.193: RT: add 172.16.105.0/24 via 172.16.34.5, rip metric [120/1]
*Mar  9 17:23:48.193: RT: NET-RED 172.16.105.0/24
*Mar  9 17:23:48.241: RT: closer admin distance for 172.16.105.0, flushing 1 routes
*Mar  9 17:23:48.241: RT: NET-RED 172.16.105.0/24
*Mar  9 17:23:48.241: RT: SET_LAST_RDB for 172.16.105.0/24
NEW rdb: via 172.16.134.4
*Mar  9 17:23:48.241: RT: add 172.16.105.0/24 via 172.16.134.4, ospf metric [110/20]
```

This cycle continues every few seconds; the route is added by RIP and then replaced by OSPF. You attempted to address this instability by lowering the RIP administrative distance to 109 for these routes on both R3 and R4.

The following output shows the **distance** command and access list as configured on R3. The same configuration was entered on R4; however, it appears not to be working.

```
R3#show run | include distance
distance 109 0.0.0.0 0.0.0.0 10
```

```
R3#show access-list 10
```

```
Standard IP access list 10
 10 permit 172.16.104.0
 20 permit 172.16.105.0
 30 permit 172.16.103.0
```

**Likely cause:** *The effort to lower the administrative distance of selected networks is ineffective.*

What could be wrong with this situation? The **distance** command specifies the source as “any” and references the access list. The access list looks correct. Do you see the problem? The only source acceptable to RIP routes for R3 and R4 is R5 IP address 172.16.34.5, so rewrite the **distance** command on both R3 and R4.

**Resolution:** *Rewrite the distance command to reference the R5 interface address.*

```
R3(config-router)#no distance 109 0.0.0.0 0.0.0.0 10
R3(config-router)#distance 109 172.16.34.5 0.0.0.0 10

R3(config-router)#do sh ip route rip
 172.16.0.0/16 is variably subnetted, 11 subnets, 2 masks
R    172.16.104.0/24 [109/2] via 172.16.34.4, 00:00:12, Ethernet0/0
R    172.16.105.0/24 [109/1] via 172.16.34.5, 00:00:12, Ethernet0/0
```

This configuration appears to work. What was the problem? The “any” expression for the route source should have been “0.0.0.0 255.255.255.255.”

Have you met the optimal path requirement for R1? It should now have two paths to both RIP networks 172.16.104.0/24 and 172.16.105.0/24.

Here is the current routing table on R1:

```
R1#show ip route
[removed for brevity]

Gateway of last resort is not set

 172.16.0.0/24 is subnetted, 14 subnets
C    172.16.16.0 is directly connected, Ethernet0/1.6
C    172.16.134.0 is directly connected, Tunnel134
O IA  172.16.36.0 [110/2] via 172.16.16.6, 00:36:23, Ethernet0/1.6
O E2  172.16.34.0 [110/20] via 172.16.134.4, 00:36:23, Tunnel134
      [110/20] via 172.16.134.3, 00:36:23, Tunnel134
C    172.16.12.0 is directly connected, Ethernet0/1.2
C    172.16.111.0 is directly connected, Loopback111
O IA  172.16.104.0 [110/65] via 172.16.134.4, 00:36:24, Tunnel134
O E2  172.16.105.0 [110/20] via 172.16.134.4, 00:03:32, Tunnel134
      [110/20] via 172.16.134.3, 00:03:32, Tunnel134
O E2  172.16.106.0 [110/20] via 172.16.16.6, 00:36:24, Ethernet0/1.6
C    172.16.101.0 is directly connected, Loopback101
D    172.16.102.0 [90/156160] via 172.16.12.2, 01:33:14, Ethernet0/1.2
O IA  172.16.103.0 [110/65] via 172.16.134.3, 00:36:24, Tunnel134
```

Use the **show ip route** command to verify the routing table on R2:

```
R2#show ip route
[removed for brevity]

Gateway of last resort is not set

 172.16.0.0/24 is subnetted, 3 subnets
C    172.16.12.0 is directly connected, Ethernet0/0
C    172.16.102.0 is directly connected, Loopback102
```

Note that R2 has no learned routes.

## 6.2. Symptom: R2 has not learned any routes from R1.

### *Analysis and testing:*

Earlier, you verified EIGRP neighbors between these routers, so the problem is most likely related to redistribution. Here is the configuration of EIGRP and the EIGRP topology table on R1:

```
router eigrp 12
  redistribute connected route-map CONNECTED
  redistribute ospf 1
  network 172.16.12.0 0.0.0.255
  auto-summary
  neighbor 172.16.12.2 Ethernet0/1.2
  redistribute eigrp 12 subnets

R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(12)/ID(172.16.111.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.12.0/24, 1 successors, FD is 28160
   via Connected, Ethernet0/1.2
P 172.16.102.0/24, 1 successors, FD is 156160
   via 172.16.12.2 (156160/128256), Ethernet0/1.2
```

Despite the **redistribution** commands under the EIGRP AS12 process, no external routes have found their way into the topology table.

### *Likely cause: Failure of redistribution into EIGRP on R1.*

Now that you have the configuration and see its results, you may realize that no metric has been identified for the redistribution. Add a **default-metric** command:

### *Resolution: Add the default-metric command under EIGRP AS12 on R1.*

This addition certainly seems to have helped; you now have many more routes in the EIGRP topology table:

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#router eigrp 12
R1(config-router)#default-metric 1 1 1 1 1
R1(config-router)#do show ip eigrp topology
IP-EIGRP Topology Table for AS(12)/ID(172.16.111.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.16.36.0/24, 1 successors, FD is 2560000256
   via Redistributed (2560000256/0)
P 172.16.34.0/24, 1 successors, FD is 2560000256
   via Redistributed (2560000256/0)
P 172.16.12.0/24, 1 successors, FD is 28160
   via Connected, Ethernet0/1.2
P 172.16.104.0/24, 1 successors, FD is 2560000256
   via Redistributed (2560000256/0)
P 172.16.105.0/24, 1 successors, FD is 2560000256
   via Redistributed (2560000256/0)
P 172.16.106.0/24, 1 successors, FD is 2560000256
   via Redistributed (2560000256/0)
P 172.16.102.0/24, 1 successors, FD is 156160
   via 172.16.12.2 (156160/128256), Ethernet0/1.2
P 172.16.103.0/24, 1 successors, FD is 2560000256
   via Redistributed (2560000256/0)
```

However, you do not see the connected networks 172.16.101.0/24, 172.16.16.0/24, or 172.16.134.0/24.

### 6.3. Symptom: Connected networks were not redistributed into EIGRP AS 12 on R1.

#### *Analysis and testing:*

The redistribution of connected routes was filtered using the route map CONNECTED so that network 172.16.101.0/24 would not be advertised. Here is that route map:

```
R1#show route-map CONNECTED
route-map CONNECTED not found
```

#### *Likely cause: Route map CONNECTED was not created or was misspelled.*

What is the effect of applying a nonexistent route map during redistribution? It is a “deny any.”

#### *Resolution: Create route map CONNECTED on R1.*

```
R1(config)#route-map CONNECTED deny 10
R1(config-route-map)# match interface Loopback111
R1(config-route-map)#!
R1(config-route-map)#route-map CONNECTED permit 20
```

Now use the **show ip route** command to verify the EIGRP routes on R2:

```
R2#show ip route eigrp
172.16.0.0/24 is subnetted, 13 subnets
D EX 172.16.16.0 [170/30720] via 172.16.12.1, 00:01:06, Ethernet0/0
D EX 172.16.134.0 [170/2172416] via 172.16.12.1, 00:01:06, Ethernet0/0
D EX 172.16.36.0
[170/2560002816] via 172.16.12.1, 00:11:14, Ethernet0/0
D EX 172.16.34.0
[170/2560002816] via 172.16.12.1, 00:11:14, Ethernet0/0
D EX 172.16.104.0
[170/2560002816] via 172.16.12.1, 00:11:14, Ethernet0/0
D EX 172.16.105.0
[170/2560002816] via 172.16.12.1, 00:11:14, Ethernet0/0
D EX 172.16.106.0
[170/2560002816] via 172.16.12.1, 00:11:14, Ethernet0/0
D EX 172.16.101.0 [170/156160] via 172.16.12.1, 00:01:06, Ethernet0/0
D EX 172.16.103.0
[170/2560002816] via 172.16.12.1, 00:11:14, Ethernet0/0
```

R2 now appears to have all the routes that it should have.

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

---

**Note** The QoS section includes a policy that drops ICMP traffic.

---

```
tclsh
foreach address {
172.16.16.1
172.16.134.1
172.16.12.1
172.16.101.1
```

172.16.12.2  
172.16.102.1

172.16.134.3  
172.16.34.3  
172.16.103.1

172.16.134.4  
172.16.34.4  
172.16.104.1

172.16.34.5  
172.16.105.1

172.16.16.6  
172.16.36.6  
172.16.106.1

172.16.36.3

} {ping \$address}

## 7. BGP Troubleshooting Section

### 7.1. Symptom: R4 is not forming BGP neighbor relationships with any peers.

#### Analysis and testing:

Begin your analysis by entering the following BGP **show** command on routers R1, R4, and R5:

```
R1#show ip bgp summary
BGP router identifier 172.16.111.1, local AS number 65001
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
1 BGP AS-PATH entries using 24 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 441 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
172.16.134.4  4 65002    65     64      0    0    0 00:13:49 Active
172.16.16.6   4   600     9      9      2    0    0 00:04:03      1

R4#sh ip bgp summary
BGP router identifier 172.16.104.1, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.34.5   4 65002     8      8      0    0    0 never Active
172.16.134.1  4 65001    11     11      0    0    0 never Active

R5#show ip bgp summary
BGP router identifier 172.16.105.1, local AS number 65002
BGP table version is 1, main routing table version 1

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.34.4   4 65002    27     27      0    0    0 00:04:46 Active
```

The **show ip bgp summary** command is a good “opening moves” troubleshooting tool for BGP neighbor relationship issues. As this command displays, all BGP neighbor relationships are in an “active” state, which is not good. These neighbor relationships must transition to an “established” state.

Review the BGP configuration of each of these three routers:

```
R1#sh run | section bgp
router bgp 65001
no synchronization
bgp log-neighbor-changes
bgp confederation identifier 100
bgp confederation peers 65002
neighbor 172.16.134.4 remote-as 65002
neighbor 172.16.16.6 remote-as 600
no auto-summary

R4#sh run | section bgp
router bgp 65001
no synchronization
bgp log-neighbor-changes
neighbor 172.16.34.5 remote-as 65002
neighbor 172.16.34.5 filter-list 1 out
neighbor 172.16.134.1 remote-as 65001
```

```
R5#sh run | s bgp
router bgp 65002
no synchronization
bgp log-neighbor-changes
neighbor 172.16.34.4 remote-as 65002
no auto-summary
```

Although R1 has a complete confederation configuration, R4 and R5 do not. Furthermore, R4 has a misconfigured **neighbor remote-as** command pointing to R1.

**Likely cause: R4 and R5 are missing the BGP confederation configuration commands and R4 has a misconfigured neighbor remote-as command.**

Examination of the configuration of both R4 and R5 shows that both contain absolutely no BGP confederation configuration commands. The two BGP confederation configuration commands are:

- **bgp confederation identifier**
- **bgp confederation peers**

Neither of these commands appears on either R4 or R5, which is the likely cause of why no BGP confederation peer neighbor relationships have been formed.

The **bgp confederation peers** command needs to be configured on a BGP confederation peer that is forming a neighbor relationship with another BGP confederation peer on another subconfederation AS. Therefore, the **bgp confederation peers** command needs to be configured on both R1 and R4 in this scenario. It is already configured on R1; it needs to be configured on R4.

The **bgp confederation identifier** command needs to be configured on a router that is forming a neighbor relationship with a BGP speaker in a different AS—both a different public AS and a different subconfederation AS. Therefore, the **bgp confederation identifier** command needs to be configured on both R1 and R4 in this scenario. It is already configured on R1; it needs to be configured on R4.

The Internal BGP (IBGP) confederation router R5 is adequately configured:

```
R5#sh run | s bgp
router bgp 65002
bgp log-neighbor-changes
neighbor 172.16.34.4 remote-as 65002
no auto-summary
```

An IBGP confederation router that maintains no EBGP neighbor relationships of any kind does not need any configuration commands on it.

**Resolution: Configure the appropriate BGP confederation commands on R4:**

```
R4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R4(config-router)#no router bgp 65001
R4(config)#router bgp 65002
R4(config-router)#neighbor 172.16.134.1 remote-as 65001
R4(config-router)#bgp confederation peers 65001
R4(config-router)#bgp confederation identifier 100
```

With these commands entered, check the status of the BGP neighbor relationships in your BGP confederation:

```
R1#show ip bgp summary
BGP router identifier 172.16.111.1, local AS number 65001
```

```

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
1 BGP AS-PATH entries using 24 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 441 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
172.16.134.4	4	65002	126	125	2	0	0	00:02:13	0
172.16.16.6	4	600	29	29	2	0	0	00:24:59	1

#### R4#sh ip bgp summary

```

BGP router identifier 172.16.104.1, local AS number 65002
BGP table version is 3, main routing table version 3
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
1 BGP AS-PATH entries using 24 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 441 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
172.16.34.5	4	65002	4	5	3	0	0	00:00:22	0
172.16.134.1	4	65001	9	9	3	0	0	00:04:42	1

#### R5#show ip bgp summary

```

BGP router identifier 172.16.105.1, local AS number 65002
BGP table version is 3, main routing table version 3
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
1 BGP AS-PATH entries using 24 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 441 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
172.16.34.4	4	65002	122	121	3	0	0	00:00:38	1

As expected, all BGP confederation peer neighbor relationships are fully established. Only R4 needed additional BGP configuration commands. On R5, the IBGP confederation peer does not need to be configured with either of the two BGP confederation configuration commands.

## 7.2. Symptom: R5 has a BGP prefix that does not contain *only* the AS path 65001.

### Analysis and testing:

In accordance with the expected behavior and network policies stated in the BGP section of this scenario, R5 must learn prefixes that include only 65001 in their AS path. Examination of the R5 BGP table reveals that it contains the following prefixes:

```

R5#sh ip bgp
BGP table version is 9, local router ID is 172.16.105.1
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

```

```
*>i172.16.101.0/24 172.16.134.1 0 100 0 (65001) i
r>i172.16.106.0/24 172.16.16.6 0 100 0 (65001) 600 i
```

The 172.16.106.0/24 prefix contains AS600 in its AS path, which violates the requirements of this scenario.

The requirements of this scenario were enforced by the following AS path access list that was configured on R4 and applied on a BGP **neighbor** command pointing to R5:

```
R4#sh run | begin bgp
router bgp 65002
 no synchronization
 bgp log-neighbor-changes
 bgp confederation identifier 100
 bgp confederation peers 65001
 neighbor 172.16.34.5 remote-as 65002
 neighbor 172.16.34.5 filter-list 1 out
 neighbor 172.16.134.1 remote-as 65001
 no auto-summary
!
ip as-path access-list 1 permit (65001)
```

The AS path access list should permit only a BGP update that includes 65001 in its AS path list. Both of the BGP prefixes in the R5 BGP table include 65001 in its AS path.

**Likely cause: The AS path regular expression on R4 must be revised.**

The AS path regular expression on R4 must be revised to allow only BGP updates with 65001 and only 65001 in the AS path. This revision will involve the use of the regular expression special characters ^ and \$.

This revision can be tested with the **show ip bgp regexp** command:

```
R4#show ip bgp regexp ^(65001)$
```

```
R4#
```

Surprisingly, this regular expression does not work because there is no matching AS path. The reason is that the parentheses symbols are regular expression special characters themselves. To be interpreted literally, they must be preceded by a backslash.

Re-enter this regular expression using backslash symbols:

```
R4#show ip bgp regexp ^\(65001\)
BGP table version is 4, local router ID is 172.16.104.1
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
*> 172.16.101.0/24 172.16.134.1      0      100      0 (65001) i
R4#
```

The regular expression now matches exactly what is requested. You can now change your configuration accordingly. The **show ip bgp regexp** command is useful because it allows for the testing of a regular expression without changing a configuration. Only when the desired regular expression is formulated and tested with the **show ip bgp regexp** command does a configuration change need to be made.

**Resolution: Revise the AS path regular expression on R4 to ^\(65001\) \$**

First, make the desired configuration change on R4:

```
R4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R4(config)#no ip as-path access-list 1
R4(config)#ip as-path access-list 1 permit ^\(65001\) $
R4(config)#end
```

```
R4#clea ip bgp * soft out
```

Next, verify that the desired result is achieved. Use two commands. One command is entered on R4 to see what R4 is advertising to R5:

```
R4#show ip bgp neighbors 172.16.34.5 advertised-routes
BGP table version is 4, local router ID is 172.16.104.1
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
*> 172.16.101.0/24	172.16.134.1	0	100	0	(65001) i

```
Total number of prefixes 1
```

The desired result is achieved. Only one prefix with 65001 in its AS path is advertised.

The second command is entered on R5 to check what the next-hop router has received on R5:

```
R5#show ip bgp
BGP table version is 10, local router ID is 172.16.105.1
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
*>i172.16.101.0/24	172.16.134.1	0	100	0	(65001) i

Once again, the desired result is achieved. Only one prefix with 65001 in its AS path is advertised. To restore reachability to the 172.16.106.0/24 network on R5, the **network backdoor** command was used under BGP on R1 for this subnet; the network is then redistributed into EIGRP on R1 from OSPF.

## 8. IPv6 Troubleshooting Section

### 8.1. Symptom: R1 has no OSPFv3 neighbors.

#### *Analysis and testing:*

For OSPF version 3 (OSPFv3), as with OSPF for IPv4, start on R1 by verifying neighbor relationships across area 0:

```
R1#sh ipv ospf nei
R1#show ipv os interface brief
Interface  PID  Area          Intf ID  Cost  State  Nbrs  F/C
Et0/2     1    0             6        64   DR     0/0
Lo101     1    0             12       1    P2P    0/0
```

```
R1#show ip ospf interface E0/2
Ethernet0/2 is up, line protocol is up
Internet Address 172.16.134.1/24, Area 0
Process ID 1, Router ID 172.16.111.1, Network Type NON_BROADCAST, Cost: 64
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 172.16.111.1, Interface address 172.16.134.1
```

Just as you saw earlier with OSPF for IPv4, OSPFv3 is configured for interface E0/2, and the network type is the default nonbroadcast. As with OSPF for IPv4, neighbor statements must be supplied to unicast IP version 6 (IPv6) hellos.

OSPFv3 neighbor addresses are link-local addresses and must be configured on the interface. Add them on R1, and see if the addition makes any difference:

**Likely cause: Missing OSPFv3 neighbor statements on R1 interface E0/2.**

```
R1(config)#int E0/2
R1(config-if)#ipv6 ospf neighbor fe80::3
R1(config-if)#ipv6 ospf neighbor fe80::4
R1(config-if)#end
```

```
R1#show ipv6 os neighbor
Neighbor ID      Pri   State           Dead Time   Interface ID  Interface
172.16.104.1     1    FULL/DROTHER    00:01:53   6             Ethernet0/2
172.16.103.1     0    FULL/DROTHER    00:01:53   6             Ethernet0/2
```

Here is the resulting IPv6 routing table on R1:

```
R1#show ipv6 route ospf
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OE2  2001:34::/64 [110/20]
      via FE80::3, Ethernet0/2
      via FE80::4, Ethernet0/2
OI   2001:103::/64 [110/65]
      via FE80::3, Ethernet0/2
OI   2001:104::1/128 [110/64]
      via FE80::4, Ethernet0/2
OE2  2001:105::/64 [110/20]
      via FE80::4, Ethernet0/2
```

As required, you have mutually redistributed between OSPFv3 and IPv6 EIGRP on R3 and R5. However, R1 has only one path to the 2001:105::/64 network.

## 8.2. Symptom: R1 has only one route to 2001:105::/64.

### *Analysis and testing:*

At the moment, R1 has a path to prefix 2001:105::/64 only through R4. Check the routing table on R3 for this route:

```
R3#show ipv route 2001:105::/64
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OE2  2001:105::/64 [110/20]
      via FE80::4, Ethernet0/2
```

R3 has this route, but it is in the routing table with a source of OSPF. This situation will not permit the EIGRP source of the route to be redistributed into OSPF on this router.

***Likely cause: The route 2001:105::/64 is not in the routing table of R3 with a source of EIGRP.***

You can address this problem by either lowering the administrative distance of this EIGRP route or by raising the administrative distance of the OSPF source for this route.

***Resolution: On R3 and R4, raise the administrative distance of external OSPFv3 routes to 171.***

```
R3(config)#ipv6 router ospf 1
R3(config-rtr)#distance ospf external 171
R4(config)#ipv6 router ospf 1
R4(config-rtr)#distance ospf external 171
```

If you have identified the correct cause and properly implemented the solution, then R1 should now have two next hops for prefix 2001:105::/64.

```
R1#show ipv6 route ospf
...
```

```
OE2 2001:105::/64 [110/20]
    via FE80::3, Ethernet0/2
    via FE80::4, Ethernet0/2
```

Use the **show ipv6 route** command to verify the IPv6 routes on R5:

```
R5#show ipv6 route eigrp
...
EX 2001:101::/64 [120/3]
    via FE80::222:90FF:FEB0:DB68, Ethernet0/0
    via FE80::222:90FF:FEB0:DDD9, Ethernet0/0
EX 2001:103::/64 [120/3]
    via FE80::222:90FF:FEB0:DB68, Ethernet0/0
    via FE80::222:90FF:FEB0:DDD9, Ethernet0/0
EX 2001:104::1/128 [120/3]
    via FE80::222:90FF:FEB0:DB68, Ethernet0/0
EX 2001:134::/64 [120/3]
    via FE80::222:90FF:FEB0:DB68, Ethernet0/0
```

### 8.3. Symptom: R5 does not have complete routes.

#### *Analysis and testing:*

It seems odd that R5 has two equal-cost routes to 2001:103::/64 but only one route to 2001:134::/64. Simplify this issue by shutting the E0/0 interface on R3 and observing the result:

```
R5#show ipv6 route eigrp
IPv6 Routing Table - 8 entries
[legend removed for brevity]
EX 2001:101::/64 [120/3]
    via FE80::222:90FF:FEB0:DDD9, Ethernet0/0
EX 2001:103::/64 [120/3]
    via FE80::222:90FF:FEB0:DDD9, Ethernet0/0
```

When you take away the R3 route source from R5, you see only two routes. The missing routes are the routes that are connected to R4. By simplifying the issue, you see that R4 is not redistributing connected routes into EIGRP.

#### *Likely cause: R4 is not redistributing connected routes into EIGRP.*

Here is the global EIGRP configuration on both R3 and R4:

```
R3#sh run | section ipv6 router eigrp
ipv6 router eigrp 1
 redistribute ospf 1 metric 1500 100 255 3 1500 include-connected
R3#
```

```
# R4#sh run | section ipv6 router eigrp
ipv6 router eigrp 1
 redistribute ospf 1 metric 1000 100 255 3 1500
R4#
```

Note the absence of the **include-connected** keyword on R4. Add that keyword to the configuration, reactivate R3 interface E0/0, and recheck the routes on R5.

#### *Resolution: Add the include-connected keyword to the redistribution on R4.*

```
R4(config)# ipv6 router eigrp 1
R4(config-rtr)# redistribute ospf 1 metric 1500 100 255 3 1500 include-connected
```

R5 now has two next hops to the 2001:134::/64 prefix, as it should. However, there is one issue left to address before you check for full IPv6 reachability: Why is the host route 2001:104::1/128 in the routing table? Configure an area range so that the Loopback104 interface on R4 is advertised with its original mask.

R4:

```
ipv6 router ospf 1
area 104 range 2001:104::/64
distance ospf external 171
redistribute eigrp 1 include-connected

!
```

```
R5#show ipv6 route eigrp
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDR - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
EX 2001:101::/64 [170/1757696]
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/0
   via FE80::A8BB:CCFF:FE00:300, Ethernet0/0
EX 2001:103::/64 [170/1757696]
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/0
   via FE80::A8BB:CCFF:FE00:300, Ethernet0/0
EX 2001:104::/64 [170/1757696]
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/0
   via FE80::A8BB:CCFF:FE00:300, Ethernet0/0
EX 2001:134::/64 [170/1757696]
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/0
   via FE80::A8BB:CCFF:FE00:300, Ethernet0/0
R5#
```

As with IPv4, you can verify full reachability with a Tcl script. Here is a Tcl script that should now indicate full reachability from each of the IPv6 routers:

```
tclsh
foreach address {
    2001:134::1
    2001:101::1

    2001:34::3
    2001:134::3
    2001:103::1

    2001:34::4
    2001:134::4
    2001:104::1

} {ping $address}
```

## 9. IP Quality of Service Troubleshooting

### 9.1. Symptom: All DSCP AF11 traffic is being dropped.

#### *Analysis and testing:*

Begin your analysis with a standard ping from R6:

```
R6#ping 172.16.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.12.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

R6#

Note that R6 sends the ICMP packets with the IP precedence 0 by the router defaults.

Continue with an extended ping specifying the exact type of service (ToS) for AF11 and pinging a destination address that transits through the interface on R1 that is assigned to VLAN 16 on the E0/1.6 interface of R1. In this case, the 172.16.12.2 address will be pinged again. Remember that you must translate AF11 into a ToS value. In this case, AF11 is translated into differentiated services code point (DSCP) value 10, and ToS value 40:

```
R6#ping
Protocol [ip]:
Target IP address: 172.16.12.2
Repeat count [5]:
Datagram size [100]: 1000
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]: 40
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 1000-byte ICMP Echos to 172.16.12.11, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

All AF11 ping packets are dropped.

A simple test to determine whether all the ping packets are being dropped due to the MQC configuration is to remove the MQC configuration from the interface and attempt the ping again.

First, remove the **service-policy** command from the E0/1.6 interface of R1:

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#inte E0/1.6
R1(config-subif)#no service-policy input ccie
R1(config-subif)#end
```

Now, rerun your ping test:

```
R1#ping
Protocol [ip]:
Target IP address: 172.16.12.2
Repeat count [5]:
Datagram size [100]: 1000
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]: 40
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 1000-byte ICMP Echos to 172.16.12.2, timeout is 2 seconds:
!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

The pings are successful. Therefore, you can deduce that the ping packets are being dropped by the MQC configuration.

Reapply the MQC configuration to the E0/1.6 interface on R1, and then perform more analysis:

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#inte E0/1.6
R1(config-subif)#service-policy input ccie
R1(config-subif)#end
```

If all traffic is being dropped due to an MQC configuration, a likely cause is a misordering of the entry of the class maps in the policy map applied to an interface.

Remember that “order of entry” does matter when MQC class maps are applied to a policy map.

First, review the class maps that have been configured on R1:

```
R1#show class-map
Class Map match-all af11 (id 1)
  Match  dscp af11 (10)

Class Map match-all icmp-prec0 (id 3)
  Match ip  precedence 0
  Match protocol  icmp

Class Map match-all icmp (id 2)
  Match protocol  icmp

Class Map match-any class-default (id 0)
  Match any

R1#
```

Three class maps have been configured. The fourth class map, **class-default**, is the MQC default class map; it includes a **match-any** setting, which is why an MQC policy map is said to end with an “implicit permit” action.

Now, review the policy map that was configured on R1 and applied to the E0/1.6 interface that is attached to VLAN 16:

```
R1#show policy-map interface E0/1.6
Service-policy input: ccie

Class-map: icmp (match-all)
  5 packets, 5070 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: protocol icmp
  drop

Class-map: af11 (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: dscp ef (46)
  police:
    cir 8000 bps, bc 2000 bytes
    conformed 0 packets, 0 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0 bps, exceed 0 bps

Class-map: class-default (match-any)
  1 packets, 60 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any
```

```

R1#show policy-map interface E0/1.6
Ethernet0/1.6

Service-policy input: ccie

Class-map: icmp-prec0 (match-all)
  5 packets, 590 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: ip precedence 0
  Match: protocol icmp
  QoS Set
  precedence 2
  Packets marked 5

Class-map: icmp (match-all)
  5 packets, 590 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: protocol icmp
  drop

Class-map: af11 (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: dscp af11 (10)
  police:
    cir 8000 bps, bc 2000 bytes
    conformed 0 packets, 0 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps

Class-map: class-default (match-any)
  55 packets, 4656 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any
R1#

```

This **show** command output shows that the broader class map “icmp” is coupled with an MQC “drop” action. Following this broader MQC classification is the narrower class map “af,” which matches only DSCP Expedited Forwarding (EF) traffic.

**Likely cause:** *The MQC class maps are entered in the policy map in the wrong order.*

To determine which class maps have been entered in the wrong order, run the following **show** command:

```

R1#sh run policy-map
Building configuration...

Current configuration : 174 bytes
!
policy-map ccie
 class icmp-prec0
   set ip precedence 2
 class icmp
   drop
 class af11
   police cir 8000 bc 2000
     conform-action transmit
     exceed-action drop
!
end
R1#

```

As expected, the broader class map of “icmp” with the MQC “drop” action is listed before the more specific class map “af11.”

**Resolution: Reorder the entry of the class maps in the MQC policy map.**

Begin by removing the current policy map configuration on R1:

```
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#no policy-map ccie
```

Once the policy map is removed, reconfigure the policy map and re-enter the class maps in the correct order:

```
R1(config)#policy-map ccie
R1(config-pmap)# class icmp-prec0
R1(config-pmap-c)# set ip precedence 2
R1(config-pmap-c)# class af11
R1(config-pmap-c)# police cir 8000 bc 2000
R1(config-pmap-c-police)# conform-action transmit
R1(config-pmap-c-police)# exceed-action drop
R1(config-pmap-c-police)# class icmp
R1(config-pmap-c)# drop
R1(config-pmap-c)#interface Ethernet0/1.6
R1(config-subif)#service-policy input ccie
R1(config-subif)#
R1(config-pmap-c)#end
R1#
```

Use the **show policy-map** command to verify that the order is correct:

```
R1#show policy-map
Policy Map ccie
  Class icmp-prec0
    set ip precedence 2
  Class af11
    police cir 8000 bc 2000
      conform-action transmit
      exceed-action drop
  Class icmp
    drop
R1#
```

The more specific “af11” class map is now listed before the broader “icmp” class map. Remember this rule when entering class maps into an MQC policy map: Order matters. As with configuring an access list, a class map must be entered in the proper order. As with an access list, the more specific match criteria should be entered first.

Now, with this minimal configuration entered, attempt your ping test from R6:

Rerun your extended ping:

```
R6#ping
Protocol [ip]:
Target IP address: 172.16.12.2
Repeat count [5]:
Datagram size [100]: 1000
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]: 40
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
```

```
Sending 5, 1000-byte ICMP Echos to 172.16.12.11, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5)
```

Run the **show policy-map** command:

```
R1#show policy-map interface E0/1.6
Ethernet0/1.6
```

```
Service-policy input: ccie
```

```
Class-map: icmp-prec0 (match-all)
  5 packets, 590 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: ip precedence 0
  Match: protocol icmp
  QoS Set
    precedence 2
    Packets marked 5
```

```
Class-map: af11 (match-all)
  5 packets, 590 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: dscp af11 (10)
  police:
    cir 8000 bps, bc 2000 bytes
    conformed 5 packets, 590 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps
```

```
Class-map: icmp (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: protocol icmp
  drop
```

```
Class-map: class-default (match-any)
  11 packets, 963 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any
```

```
R1#
```

The **show policy-map interface** counters are incrementing as expected.

---

<b>Note</b>	The Mentor Guide engine in the web portal can help you use Cisco IOS Software commands to see a comprehensive view of the configuration for a specific section. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands as well as a collection of proprietary commands such as <b>show all</b> .
-------------	--

---

## 10. IP SLA Troubleshooting Section

### 10.1. Symptom: IP SLAs monitor traffic is not collecting statistics

#### *Analysis and testing:*

The IP service level agreements (SLAs) process is not collecting statistics. However, R2 is receiving the following ICMP error messages:

```
R2#debug ip icmp
ICMP packet debugging is on
```

```
R2#ICMP: dst (172.16.12.2) port unreachable rcv from 172.16.16.6
```

```
R2#ICMP: dst (172.16.12.2) port unreachable rcv from 172.16.16.6
```

Here is the baseline IP SLAs configuration on R2:

```
R2#sh run | section sla
ip sla monitor 1
  type jitter dest-ipaddr 172.16.106.1 dest-port 16387
  frequency 300
ip sla monitor schedule 1 life forever start-time now
```

```
R2#sh run | section sla
ip sla 1
  udp-jitter 172.16.106.1 16387
  frequency 300
ip sla schedule 1 life forever start-time now
R2#
```

The baseline configuration on R2 looks fine. Run the following IP SLAs **show** command:

```
R2#show ip sla statistics
IPSLAs Latest Operation Statistics

IPSLA operation id: 1
Type of operation: udp-jitter
  Latest RTT: 1 milliseconds
Latest operation start time: 17:15:38 PST Thu Apr 25 2013
Latest operation return code: No connection
RTT Values:
  Number Of RTT: 10                      RTT Min/Avg/Max: 1/1/1 milliseconds
Latency one-way time:
  Number of Latency one-way Samples: 5
  Source to Destination Latency one way Min/Avg/Max: 0/0/1 milliseconds
  Destination to Source Latency one way Min/Avg/Max: 0/0/1 milliseconds
Jitter Time:
  Number of SD Jitter Samples: 9
  Number of DS Jitter Samples: 9
  Source to Destination Jitter Min/Avg/Max: 0/1/1 milliseconds
  Destination to Source Jitter Min/Avg/Max: 0/1/1 milliseconds
Over Threshold:
  Number Of RTT Over Threshold: 0 (0%)
Packet Loss Values:
  Loss Source to Destination: 0
  Source to Destination Loss Periods Number: 0
  Source to Destination Loss Period Length Min/Max: 0/0
  Source to Destination Inter Loss Period Length Min/Max: 0/0
  Loss Destination to Source: 0
  Destination to Source Loss Periods Number: 0
  Destination to Source Loss Period Length Min/Max: 0/0
  Destination to Source Inter Loss Period Length Min/Max: 0/0
  Out Of Sequence: 0      Tail Drop: 0
  Packet Late Arrival: 0  Packet Skipped: 0
Voice Score Values:
  Calculated Planning Impairment Factor (ICPIF): 0
  Mean Opinion Score (MOS): 0
Number of successes: 11
Number of failures: 0
Operation time to live: Forever
Number of successes: 0
Number of failures: 21
Operation time to live: Forever
```

Notice how all the IP SLAs monitor timers are zeroed out. Notice that the latest operation code is “no connection,” and the number of IP SLAs successes is “0.” The configuration appears to be incomplete.

**Likely cause: The IP SLAs responder is not enabled on R6.**

Before the IP SLAs monitor jitter service can be used, the command **ip sla responder** must be configured on the target router. In this case, it would be on R6:

```
R2#sh run | section sla
R2#
```

Only the router prompt is returned. Therefore, there is no **ip sla responder** command on R6.

**Resolution: Enable the IP SLAs responder on R6.**

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

Once the IP SLAs monitor responder is configured on R6, the **show ip sla statistics** command generates counters that show that the IP SLAs monitor is operational:

```
R2#show ip sla statistics
IPSLAs Latest Operation Statistics

IPSLA operation id: 1
Type of operation: udp-jitter
  Latest RTT: 1 milliseconds
Latest operation start time: 17:15:38 PST Thu Apr 25 2013
Latest operation return code: OK
RTT Values:
  Number Of RTT: 10                RTT Min/Avg/Max: 1/1/1 milliseconds
Latency one-way time:
  Number of Latency one-way Samples: 5
  Source to Destination Latency one way Min/Avg/Max: 0/0/1 milliseconds
  Destination to Source Latency one way Min/Avg/Max: 0/0/1 milliseconds
Jitter Time:
  Number of SD Jitter Samples: 9
  Number of DS Jitter Samples: 9
  Source to Destination Jitter Min/Avg/Max: 0/1/1 milliseconds
  Destination to Source Jitter Min/Avg/Max: 0/1/1 milliseconds
Over Threshold:
  Number Of RTT Over Threshold: 0 (0%)
Packet Loss Values:
  Loss Source to Destination: 0
  Source to Destination Loss Periods Number: 0
  Source to Destination Loss Period Length Min/Max: 0/0
  Source to Destination Inter Loss Period Length Min/Max: 0/0
  Loss Destination to Source: 0
  Destination to Source Loss Periods Number: 0
  Destination to Source Loss Period Length Min/Max: 0/0
  Destination to Source Inter Loss Period Length Min/Max: 0/0
  Out Of Sequence: 0      Tail Drop: 0
  Packet Late Arrival: 0  Packet Skipped: 0
Voice Score Values:
  Calculated Planning Impairment Factor (ICPIF): 0
  Mean Opinion Score (MOS): 0
Number of successes: 11
Number of failures: 0
Operation time to live: Forever
```

```
R2#
```

The “Latest operation return code” previously listed “no connection,” but it now lists “OK.” Also, the “Number of successes” counter, which previously listed 0, now lists 11. The baseline IP SLAs monitor is now operational.