360RSW05-WKB-TA08

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 8 Troubleshooting Section Answer Key

Table of Contents

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 8 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 10.0.0.0/8.
- *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.
- Prefixes displayed in routing tables must have mask lengths between /8 and /24, inclusive.
- 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 from within the same virtual routing and forwarding (VRF) instance. This lab does not require reachability between VPNA IP addresses and the remainder of the pod.

- Use conventional routing algorithms only, unless the instructions specify otherwise.
- Do not create new interfaces to fulfill IGP requirements; do not summarize unless you are 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 IPv4 IGP diagram must be used.

All IPv4 IP addresses in this exam use a second octet that is equal to the default administrative distance value of the IGP routing protocol that the address is assigned to.

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.

Prefixes displayed in routing tables must have mask lengths between /8 and /24, inclusive.

In general, networks should be advertised with their original masks. This requirement permits the use of summaries to provide reachability, but rules out the 0.0.0.0/0 default route and host routes.

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

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 involved in this scenario must be reachable from within the same VRF.

This is a key goal to observe. It requires that all your IGPs and all your routing policy tasks must be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using the **distribute-lists**, **route-maps**, and **distance** commands. Even when redistribution is not specifically mentioned, it may be implied by this requirement. In this lab, IP addresses in VRF VPNA do not need to be reachable from outside the VPN.

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. Switched Network Troubleshooting Section

Begin with a review of a Layer 2 diagram for this specific lab:

VLAN Propagation Diagram

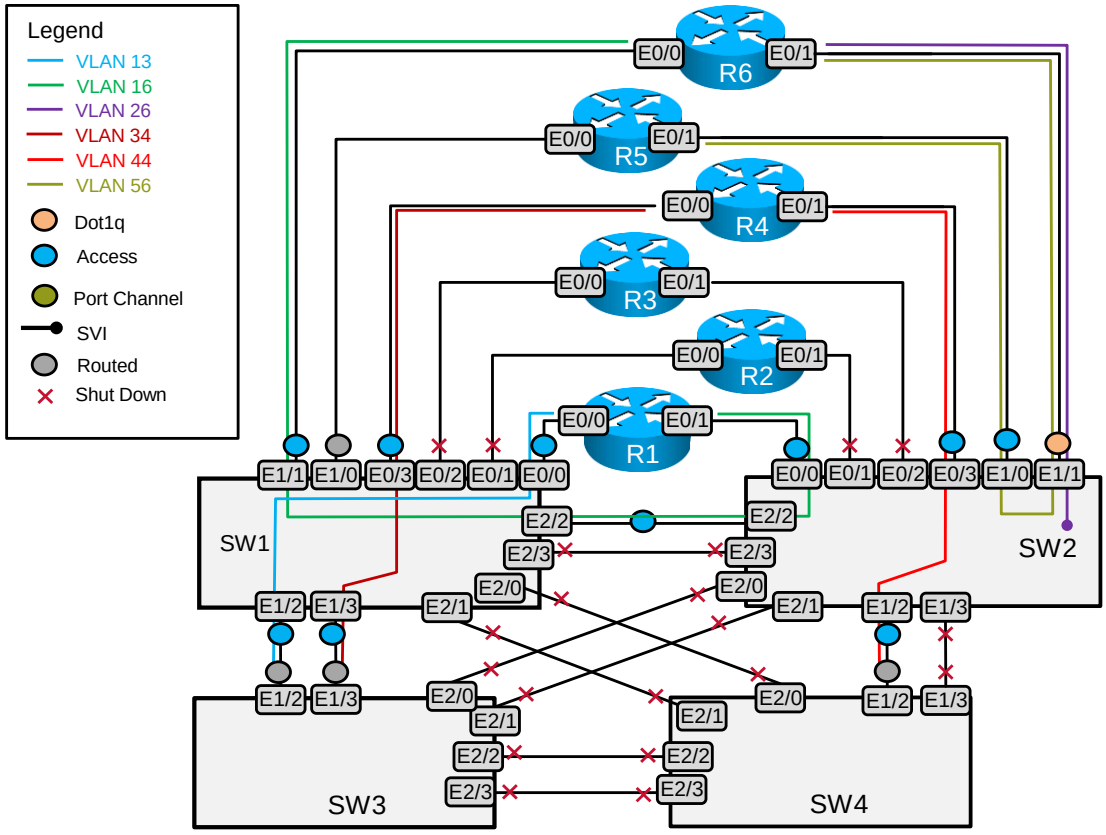

Use this diagram and the Layer 3 diagram from the scenario as guides to test Ethernet same-subnet reachability and correspondence with scenario requirements.

1.1. Symptom: R4 cannot reach SW3.

Analysis and testing:

Pings fail between IP address 10.1.34.3 on SW3 and 10.1.34.4 on R4.

```
SW3#ping 10.1.34.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
SW3#
```

The “Ethernet Topology” diagram shows these addresses connected over VLAN 34. Tracing the path of this VLAN across the switches, you see that SW1 interfaces Et0/3 and Et1/3 should be access ports in VLAN 34. Interface E1/3 on SW3 should be a routed interface. However, this is the configuration found on SW1 Et1/3:


```
SW1#sh run int Et1/3
Building configuration...

Current configuration : 64 bytes
!
interface Ethernet1/3
 no switchport
 no ip address
```

Likely cause: SW1 interface Et1/3 is configured as a routed port.

Close examination shows a misconfiguration on the Ethernet path between R4 and SW3.

Resolution: Configure SW1 Et1/3 as a switched access port in VLAN 34.

Enter the commands **switchport** and **switchport access VLAN 34** on SW1 interface Et1/3.

```
SW1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#int e1/3
SW1(config-if)#switchport
SW1(config-if)#switchport access vlan 34
SW1(config-if)#switchport mode access
SW1(config-if)#end
SW1#

SW3#ping 10.1.34.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.4, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
SW3#
```

1.2. Symptom: SW2 cannot reach R6.

Analysis and testing:

SW2 cannot ping IP address 10.50.26.6 on R6.

```
SW2#ping 10.50.26.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.50.26.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
SW2#
```

The "Ethernet Topology" diagram shows that R6 Et0/1 physically connects to SW2 interface Et1/1. The link connecting them should be a dot1q trunk that permits VLANs 26 and 56. Here is the configuration of this trunk link on SW2:

```
SW2#show interface e1/1 trunk

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

Port                Vlans allowed on trunk
Et1/1               26,56

Port                Vlans allowed and active in management domain
Et1/1               56

Port                Vlans in spanning tree forwarding state and not pruned
Et1/1               56
```

```
SW2#
```

Everything here looks fine, until you notice that VLAN 26 is not active on this port. When you check the status of interface VLAN 26, you see that the status is "up/down":

```
SW2#show ip int brief | inc Vlan26
Vlan26                 10.50.26.2      YES NVRAM  up                    down
```

Enter the **show vlan** command.

```
SW2#show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/1, Et0/2, Et1/3, Et2/0
                                                Et2/1, Et2/3
16   VLAN0016                         active    Et0/0, Et2/2
44   VLAN0044                         active    Et0/3, Et1/2
56   VLAN0056                         active    Et1/0
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1    enet  100001     1500  -      -      -        -    -        0      0
16   enet  100016     1500  -      -      -        -    -        0      0
44   enet  100044     1500  -      -      -        -    -        0      0
56   enet  100056     1500  -      -      -        -    -        0      0
1002 fddi  101002     1500  -      -      -        -    -        0      0
1003 tr    101003     1500  -      -      -        -    -        0      0
1004 fdnet 101004     1500  -      -      -        ieee -        0      0
1005 trnet 101005     1500  -      -      -        ibm  -        0      0

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------

SW2#
```

Likely cause: VLAN 26 was not created on SW2.

VLAN 26 was not created on SW2, so interface SW2 E1/1 dropped any traffic tagged with this VLAN ID from R6, and interface VLAN 26 was not active.

Resolution:

Configure VLAN 26 on SW2 and retest.

```
SW2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#vlan 26
SW2(config-vlan)#end
% Applying VLAN changes may take few minutes.  Please wait...

SW2#

SW2#ping 10.50.26.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.50.26.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
SW2#
```


1.3. Symptom: SW3 is not in VTP transparent mode.

Analysis and testing:

The lab requirements state that all switches be in VTP transparent mode. Verify this on each switch with the command **show vtp status**.

Likely cause: SW3 is using the default VTP mode, server.

By default, the switches are in VTP server mode.

```
SW3#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 5
VTP Operating Mode              : Server
VTP Domain Name                 :
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00
Local updater ID is 10.1.13.3 on interface Et1/2 (first layer3 interface found)
VTP version running             : 1
SW3#
```

***Resolution: On SW3, enter the command** vtp mode transparent.*

```
SW3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW3(config)#vtp mode transparent
Setting device to VTP Transparent mode for VLANS.
SW3(config)#end
SW3#
```

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**. 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.

2. IPv4 EIGRP Troubleshooting Section

2.1. Symptom: R2 and R3 lack reachability.

Analysis and testing:

R2 and R3 have only one learned route, as you see below for R2. How will these routers reach each other’s loopback subnets and IP addresses outside the EIGRP domain?

```
R2#sh ip route
[output removed for brevity]

     10.0.0.0/24 is subnetted, 3 subnets
D       10.10.101.0 [90/2297856] via 10.10.123.1, 00:00:15, Serial0/0/0
C       10.10.102.0 is directly connected, Loopback102
C       10.10.123.0 is directly connected, Serial0/0/0
```

The scenario rules out many potential reachability solutions including:

- static routes
- a default route
- a default network
- external routes
- disabling EIGRP split horizon on the hub router
- configuring an EIGRP neighbor relationship between R2 and R3

A 10.0.0.0/8 summary is configured on R1 interface E0/2. However, it is not being advertised and EIGRP has not created a discard route to null 0 for the summary. If you can get this summary to work, it should provide full reachability.

Likely cause: The configured 10/8 summary is not being advertised to R2 and R3 by R1.

Here is the configured summary:

```
R1#show run int e0/2
Building configuration...

Current configuration : 134 bytes
!
interface Ethernet0/2
 ip address 10.10.123.1 255.255.255.0
 ip pim sparse-mode
 ip summary-address eigrp 1 10.0.0.0 255.0.0.0
end

R1#
R1#show ip route eigrp
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

      10.0.0.0/8 is variably subnetted, 18 subnets, 2 masks
D        10.10.102.0/24 [90/409600] via 10.10.123.2, 00:13:39, Ethernet0/2
D        10.10.103.0/24 [90/409600] via 10.10.123.3, 00:13:35, Ethernet0/2
R1#
```

On close inspection, you can see that it references EIGRP process 1, when it should reference the existing EIGRP process 10.

Resolution: Correct the EIGRP AS 10 summary on R1 interface E0/2.

After correcting this summary, you see the 10.0.0.0/8 summary in the routing tables of both R2 and R3. They can now ping each other's loopback IP addresses and all connected routes on R1.

```
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#interface Ethernet0/2
R1(config-if)#no  ip summary-address eigrp 1 10.0.0.0 255.0.0.0
R1(config-if)# ip summary-address eigrp 10 10.0.0.0 255.0.0.0
R1(config-if)#end
```



```
R1#
R1#
R1#show ip route eigrp
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

      10.0.0.0/8 is variably subnetted, 19 subnets, 3 masks
D        10.0.0.0/8 is a summary, 00:00:13, Null0
D        10.10.102.0/24 [90/409600] via 10.10.123.2, 00:15:07, Ethernet0/2
D        10.10.103.0/24 [90/409600] via 10.10.123.3, 00:15:03, Ethernet0/2
R1#
```

Verify the EIGRP table on R2:

```
R2#sh ip route eigrp
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

      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
D        10.0.0.0/8 [90/409600] via 10.10.123.1, 00:00:29, Ethernet0/2
R2#
```

Try to ping R3 from R2:

```
R2#ping 10.10.103.1 source 10.10.102.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.103.1, timeout is 2 seconds:
Packet sent with a source address of 10.10.102.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
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.

3. IPv4 RIP Troubleshooting Section

3.1. Symptom: SW4 is not learning any RIP routes from R4.

Analysis and testing:

When you check the routing table on SW4, you find only connected routes.

```
SW4#show ip route
```

```
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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.40.44.0/24 is directly connected, Ethernet1/2
L        10.40.44.40/32 is directly connected, Ethernet1/2
C        10.40.140.0/24 is directly connected, Loopback140
L        10.40.140.1/32 is directly connected, Loopback140
SW4#
```

R4 interface Et0/1 is configured to send a 10.0.0.0/8 summary to SW4. You enable the **debug ip rip** command on R4 and find that R4 is receiving prefix 10.40.140.0/24 from SW4, but it is not advertising any routes.

R4 is receiving RIP routes, but is not sending them, and there appears to be no filter configured on these devices. Check RIP operation on R4:

```
R4#show ip protocols | section rip
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 24 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 1, receive any version
    Interface             Send  Recv  Triggered RIP  Key-chain
    Ethernet0/1           1     1 2
  Automatic network summarization is in effect
  Address Summarization:
    10.0.0.0/8 for Ethernet0/1
  Maximum path: 4
  Routing for Networks:
     10.0.0.0
    10.0.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
    Serial1/2
  Passive Interface(s):
    Serial1/3
    Loopback104
    RG-AR-IF-INPUT1
    Tunnel0
    Tunnel1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.40.44.40           120      00:00:22
  Distance: (default is 120)
  Redistributing: connected, rip
R4#
```

This output confirms that no filters have been applied and that Et0/1 is not configured as a passive interface. The “version control” section of this output indicates default operation; Et0/1 is configured to send RIP version 1 updates and to receive either version 1 or version 2.

The task requires only RIP version 2 updates be exchanged on this link, but you are not permitted to make any configuration changes under the global RIP process.

Likely cause: RIP version 1 does not support the advertisement of locally created summary routes.

RIP version 1 does not support the advertisement of locally created summary routes, but you must send this summary to meet the lab reachability requirements within the given restrictions.

Resolution: Configure R4 interface Et0/1 to send RIP version 2 advertisements.

Since you cannot change the global RIP process to version 2, you will configure the sending of RIP version 2 updates at the interface level. On R4 interface Et0/1, configure the command **ip rip send version 2**. In a moment, you will see a routing table similar to the following on SW4, and SW4 will have reachability to the required pod addresses.

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int e0/1
R4(config-if)#ip rip send version 2
R4(config-if)#end
R4#

SW4#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
R        10.0.0.0/8 [120/1] via 10.40.44.4, 00:00:03, Ethernet1/2
SW4#
```

4. BGP Troubleshooting Section

4.1. Symptom: The BGP peering between R4 and R5 is down.

Analysis and testing:

To verify basic BGP configuration, run the command **show ip bgp summary** on R1, R4, R5, R6, and SW3. The expected results are found, except on R4 and R5:

```
R4#show ip bgp summary
[output removed for brevity]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.34.3       4     3     855     872      134    0    0 13:12:08        10
10.1.45.5       4    56     850     878        0    0    0 00:01:03 Active

R5#show ip bgp summary
[output removed for brevity]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.54.4       4     4       0       0        0    0    0 never    Idle
10.56.106.1     4    56      81      75       19    0    0 00:01:19        14
```

On R4, the status is shown as "Active," indicating that attempts have been made to form a peering, but they have not yet succeeded. The status of this peering on R5 is "Idle." This status usually indicates that no route can be found to the configured peer.

Likely cause: R5's peering statement with R4 uses an incorrect peering address.

When you examine this output more closely, you see that the peering has been configured to non-existent IP address 10.1.54.4, instead of 10.1.45.4, as required.

Resolution: Reconfigure R5's peering statement to use IP address 10.1.45.4.

Correct the BGP configuration on R5:

```
R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R5(config)#router bgp 56
R5(config-router)#neighbor 10.1.45.4 remote-as 4
R5(config-router)#no neighbor 10.1.54.4 remote-as 4
R5(config-router)#end
R5#
```

After correcting this error, the peering between R4 and R5 comes up and routes are exchanged.

```
R5#show ip bgp summary
BGP router identifier 10.56.105.1, local AS number 56
BGP table version is 21, main routing table version 21
14 network entries using 1960 bytes of memory
19 path entries using 1444 bytes of memory
9/7 BGP path/bestpath attribute entries using 1260 bytes of memory
5 BGP AS-PATH entries using 120 bytes of memory
1 BGP extended community entries using 40 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 4824 total bytes of memory
BGP activity 16/0 prefixes, 28/7 paths, scan interval 60 secs

Neighbor           V               AS  MsgRcvd  MsgSent     TblVer   InQ  OutQ  Up/Down
State/PfxRcd
10.1.45.4        4              4       14       14        21     0     0 00:00:38        10
10.56.106.1      4             56       41       35        21     0     0 00:25:42         9
R5#
```

5. IPv4 Redistribution Troubleshooting Section

5.1. Symptom: Thorough testing shows that IP address 10.56.105.1 is not reachable from outside AS 56.

Analysis and testing:

After completing redistribution, it is a good practice to thoroughly test for required reachability. Ping tests can be automated on routers using a Tool Command Language (Tcl) script like the following:

```
tclsh
foreach address {

10.1.13.1
10.1.16.1
10.10.101.1
10.10.123.1

10.10.102.1
10.10.123.2

10.10.103.1
10.10.123.3
```



```
10.40.44.4
10.1.45.4
10.1.34.4
10.40.104.1

10.56.56.5
10.1.45.5
10.56.105.1

10.56.56.6
10.1.16.6
10.56.106.1

10.1.13.3
10.1.34.3
10.30.130.1

10.40.44.40
10.40.140.1

} {ping $address}
```

Note that this script does not include the IP addresses in VPNA. When you run this script from R1, R2, R3, and R4, you find that IP address 10.56.105.1 is unreachable, though it is reachable form R5 and R6. The pattern of reachability suggests that connected routes may not have been redistributed into BGP on R5.

Likely cause: Connected routes have not been redistributed into BGP on R5.

On R5, check the BGP configuration:

```
R5#show run | section bgp
 redistribute bgp 56 subnets
router bgp 56
 neighbor 10.1.45.4 remote-as 4
 neighbor 10.56.106.1 remote-as 56
 neighbor 10.56.106.1 update-source Loopback105
 no auto-summary
 !
 address-family vpnv4
  neighbor 10.56.106.1 activate
  neighbor 10.56.106.1 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf VPNA
  redistribute ospf 50 vrf VPNA
  no synchronization
 exit-address-family
```

Resolution: Add the** redistribute connected **command under the BGP process on R5.

Under the BGP process on R5, add the command **redistribute connected**. Repeat your reachability tests. You will find that all devices outside VPNA can now reach all non-VPNA IP addresses.

```
R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R5(config)#router bgp 56
R5(config-router)#redistribute connected
R5(config-router)#end
R5#
```

Verify the Tcl script again on R1:

```
R1#tclsh
R1(tcl)#foreach address {
```

```
+>(tcl)#
+>(tcl)#10.1.13.1
+>(tcl)#10.1.16.1
+>(tcl)#10.10.101.1
+>(tcl)#10.10.123.1
+>(tcl)#
+>(tcl)#10.10.102.1
+>(tcl)#10.10.123.2
+>(tcl)#
+>(tcl)#10.10.103.1
+>(tcl)#10.10.123.3
+>(tcl)#
+>(tcl)#10.40.44.4
+>(tcl)#10.1.45.4
+>(tcl)#10.1.34.4
+>(tcl)#10.40.104.1
+>(tcl)#
+>(tcl)#10.56.56.5
+>(tcl)#10.1.45.5
+>(tcl)#10.56.105.1
+>(tcl)#
+>(tcl)#10.56.56.6
+>(tcl)#10.1.16.6
+>(tcl)#10.56.106.1
+>(tcl)#
+>(tcl)#10.1.13.3
+>(tcl)#10.1.34.3
+>(tcl)#10.30.130.1
+>(tcl)#
+>(tcl)#10.40.44.40
+>(tcl)#10.40.140.1
+>(tcl)#
+>(tcl)#} {ping $address}
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.16.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.101.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.123.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.102.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.123.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.103.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.123.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.40.44.4, timeout is 2 seconds:
!!!!!
```



```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.40.104.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.56.56.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.56.105.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.56.56.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.16.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.56.106.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.30.130.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.40.44.40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.40.140.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
R1(tcl)#exit
R1
```

6. MPLS Layer 3 VPN Troubleshooting Section

6.1. Symptom: OSPF routes from the far VPNA site are not available in local routing tables.

Analysis and testing:

The basic operation of OSPF processes 0, 50, and 60 appears to be correct: R5 sees R6 and SW1 as OSPF neighbors, and R6 sees R5 and SW2 as OSPF neighbors. R5 has an OSPF route to R6 loopback 106, and R6 has an OSPF route to R5 loopback 105. As shown below, the provider edge (PE) routers are learning OSPF routes from the attached customer edge (CE) routers, SW1 and SW2.

```
R5#show ip route vrf VPNA
[output removed for brevity]

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C        10.50.15.0/24 is directly connected, Ethernet0/0
L        10.50.15.5/32 is directly connected, Ethernet0/0
O       10.50.110.0 [110/2] via 10.50.15.1, 01:49:35, Ethernet0/0

R6#show ip route vrf VPNA

Routing Table: VPNA
[output removed for brevity]

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C        10.50.26.0/24 is directly connected, Ethernet0/1.26
L        10.50.26.6/32 is directly connected, Ethernet0/1.26
O       10.50.120.0 [110/2] via 10.50.26.2, 01:49:02, Ethernet0/1.26
```

However, both SW1 and SW2 have only connected routes in their routing tables; they are not learning OSPF routes from the far VPN site.

Are the PE routers R5 and R6 advertising VPN routes between the sites? Here are the VPNv4 BGP tables on R5 and R6:

```
R5#show bgp vpnv4 unicast vrf VPNA
BGP table version is 5, local router ID is 10.56.105.1
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
Route Distinguisher: 56:1 (default for vrf VPNA)
*> 10.50.15.0/24    0.0.0.0                  0         32768 ?
*> 10.50.110.0/24   10.50.15.1               2         32768 ?

R6#show bgp vpnv4 unicast vrf VPNA
BGP table version is 5, local router ID is 10.56.106.1
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
Route Distinguisher: 56:1 (default for vrf VPNA)
*> 10.50.26.0/24    0.0.0.0                  0         32768 ?
*> 10.50.120.0/24   10.50.26.2               2         32768 ?
```

Each router has only its own local-site routes. Is there a good VPNv4 peering between R5 and R6?

```
R5#show bgp vpnv4 unicast all summary
[table legend removed for brevity]
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.56.106.1     4    56      19      19        5    0    0 00:07:50        0
```

This output shows that there is a good VPNv4 peering between the PE routers, but no VPN prefixes are being sent. The result suggests that some type of filter is being applied, but none is apparent in the configuration.

Going back to the basics of MPLS Layer 3 VPN operation, you recall that VPN route advertisement is determined by VRF import route targets. If a PE router does not require routes for a particular VPN, then it does not request them from its peers.

Likely cause: PE routers are not requesting VPN routes from other sites, because no import route targets are configured.

When you examine the VRF configurations on R5 and R6, you see that no import route targets have been applied.

R5:

```
!
ip vrf VPNA
 rd 56:1
 route-target export 56:6
!
```

R6:

```
!
ip vrf VPNA
 rd 56:1
 route-target export 56:5
!
```

Resolution: Configure the appropriate import route targets on R5 and R6.

VRF VPNA uses export route target 56:5 on R5, so R6 should import this route target. R6 uses export route target 56:6, so R5 should import this route target.

```
R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R5(config)#ip vrf VPNA
R5(config-vrf)# route-target import 56:6
R5(config-vrf)#end
R5#
R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)#ip vrf VPNA
R6(config-vrf)#route-target import 56:5
R6(config-vrf)#end
R6#
```

After these configurations changes are made, VPN routes are advertised between the PE routers:

```
R5#show bgp vpnv4 unicast vrf VPNA
[table legend removed for brevity]

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 56:1 (default for vrf VPNA)
*> 10.50.15.0/24    0.0.0.0                  0         32768 ?
*>i10.50.26.0/24    10.56.106.1              0    100      0 ?
*> 10.50.110.0/24   10.50.15.1               2         32768 ?
*>i10.50.120.0/24   10.56.106.1              2    100      0 ?

R6#show bgp vpnv4 unicast vrf VPNA
[table legend removed for brevity]

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 56:1 (default for vrf VPNA)
*>i10.50.15.0/24    10.56.105.1              0    100      0 ?
*> 10.50.26.0/24    0.0.0.0                  0         32768 ?
*>i10.50.110.0/24   10.56.105.1              2    100      0 ?
*> 10.50.120.0/24   10.50.26.2               2         32768 ?
```

6.2. Symptom: OSPF routes from the far VPN site are OE2.

Analysis and testing:

You now have reachability between sites in VPNA, but the routes learned from the far site appear as OSPF external routes. Here are the VPN OSPF routes as they appear on SW1 and SW2:

```
SW1#show ip route ospf
     10.0.0.0/8 is subnetted, 4 subnets
O E2    10.50.26.0/24 [110/1] via 10.50.15.5, 00:08:11, Ethernet1/0
O E2    10.50.120.0/24 [110/2] via 10.50.15.5, 00:08:11, Ethernet1/0

SW2#show ip route ospf
     10.0.0.0/8 is subnetted, 4 subnets
O E2    10.50.15.0/24 [110/1] via 10.50.26.6, 00:11:48, Vlan26
O E2    10.50.110.0/24 [110/2] via 10.50.26.6, 00:11:48, Vlan26
```

Recall that OSPF routes from distant VPN sites will be shown as internal routes when domain IDs in each site are the same, and will be shown as external routes when the domain IDs differ. Domain IDs can be seen in the details of the local process.

```
R5#show ip ospf 50
 Routing Process "ospf 50" with ID 10.50.15.5
   Domain ID type 0x0005, value 0.0.0.50
[Remaining output removed for brevity]

R6#show ip ospf 60
 Routing Process "ospf 60" with ID 10.50.26.6
   Domain ID type 0x0005, value 0.0.0.60
[Remaining output removed for brevity]
```

By default, the domain ID is equal to the local OSPF process ID, and this value is attached to each OSPF route that is advertised by the PE router. Here you see domain ID 60 represented in hexadecimal and associated with prefix 10.50.26.0/24 in R5's VPNA RIB:

```
R5#show bgp vpnv4 unicast vrf VPNA 10.50.26.0
BGP routing table entry for 56:1:10.50.26.0/24, version 12
Paths: (1 available, best #1, table VPNA)
  Not advertised to any peer
  Local
    10.56.106.1 (metric 2) from 10.56.106.1 (10.56.106.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:56:6 OSPF DOMAIN ID:0x0005:0x0000003C0200
        OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.50.26.6:0
      mpls labels in/out nolabel/16
```

Likely cause: The OSPF domain IDs differ in VPNA site 5 and site 6.

By default, the OSPF domain ID is equal to the OSPF process ID configured on the PE router. If the domain IDs do not match between sites, then OSPF routes from the far site will be injected into OSPF as external routes.

Resolution: Configure matching domain IDs for OSPF processes 50 and 60.

You are not permitted to change the OSPF process IDs. Instead, you must statically configure equal domain IDs under OSPF process 50 on R5 and OSPF process 60 on R6.

```
R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R5(config)#router ospf 50 vrf VPNA
R5(config-router)#domain-id 0.0.0.56
R5(config-router)#end
```



```
R5#

R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)#router ospf 60 vrf VPNA
R6(config-router)#domain-id 0.0.0.56
R6(config-router)#end
R6#
```

Here is the resulting routing table on SW1 and SW2:

```
SW1#show ip route ospf
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA     10.50.26.0/24 [110/11] via 10.50.15.5, 05:31:28, Ethernet1/0
O IA     10.50.120.0/24 [110/21] via 10.50.15.5, 05:31:28, Ethernet1/0
SW1#

SW2#show ip route ospf
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA     10.50.15.0/24 [110/2] via 10.50.26.6, 05:33:14, Vlan26
O IA     10.50.110.0/24 [110/12] via 10.50.26.6, 05:33:14, Vlan26
SW2#
```

7. Router Monitoring Troubleshooting Section

7.1. Symptom: When R1 Loopback 106 is shut down, the BGP, SW1 can no longer reach SW2.

Analysis and testing:

R6 is configured with an Embedded Event Manager (EEM) applet. It is designed to monitor the syslog for an indication that Loopback 106 protocol status has gone to a down state, and then configure a replacement loopback interface. Here is the applet as it is currently configured on R6:

```
event manager applet LoopbackWatch
 event syslog pattern ".*UPDOWN.*Loopback106.*down"
 action 1.0 syslog priority warning msg "EEM says Loopback 106 is down"
 action 1.1 cli command "enable"
 action 1.2 cli command "configure terminal"
 action 1.3 cli command "interface Loopback166"
 action 1.4 cli command "ip address 10.56.106.1 255.255.255.0"
 action 1.5 cli command "end"
```

When you test the applet by shutting Loopback 106, you see the following output to the console:

```
R6(config)#int loop106
R6(config-if)#shut
R6(config-if)#
*Feb 28 10:21:54.659: %TDP-5-INFO: Default-IP-Routing-Table: TDP ID removed
*Feb 28 10:21:54.659: %LDP-5-NBRCHG: LDP Neighbor 10.56.105.1:0 (1) is DOWN (LDP
Router ID changed)
*Feb 28 10:21:56.659: %LINK-5-CHANGED: Interface Loopback106, changed state to
administratively down
*Feb 28 10:21:57.659: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback106,
changed state to down
*Feb 28 10:21:57.663: %HA_EM-4-LOG: LoopbackWatch: EEM says Loopback 106 is down
*Feb 28 10:21:57.755: %SYS-5-CONFIG_I: Configured from console by  on vty0
(EEM:LoopbackWatch)
*Feb 28 10:21:58.707: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback166,
changed state to up
R6(config-if)#
*Feb 28 10:21:59.763: %LDP-5-NBRCHG: LDP Neighbor 10.56.105.1:0 (1) is UP
```

Label Distribution Protocol (LDP) is using Loopback 106 as its router ID, so the LDP neighbor goes down, EEM sources a syslog message, interface Loopback 166 comes up, and the LDP neighbor relationship reforms. Even though BGP used Loopback 106 as its update source with the peer on R5, the BGP neighbor relationship did not go down. Everything seems to be back to normal. However, pings from SW1 to IP address 10.50.26.2 on SW2 now fail. In fact, there is no longer any reachability between VPNA sites 5 and 6.

First, note that SW1 has a route to 10.50.26.2, and SW2 has a route back. Enable **debug ip icmp** on SW1 and SW2 and then ping from one switch to the other. When you ping from SW1 to SW2, you see no response on SW2, but when you ping from SW2 to SW1, you see that SW1 receives the echo request and sends an echo reply. The path from SW1 to SW2 is broken, but the path from SW2 to SW1 appears to be working.

Here is the BGP routing entry for subnet IP address 10.50.26.2 on R5:

```
R5#sho bgp vpnv4 unicast all 10.50.26.0/24
BGP routing table entry for 56:1:10.50.26.0/24, version 6
Paths: (1 available, best #1, table VPNA)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.56.106.1 (metric 11) from 10.56.106.1 (10.56.106.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:56:6 OSPF DOMAIN ID:0x0005:0x000000380200
        OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:6.6.6.6:0
      mpls labels in/out nolabel/16
      rx pathid: 0, tx pathid: 0x0
R5#
```

The next hop is the BGP next hop on R6, which is IP address 10.56.106.1. To send VPN traffic to this address, R5 must find an LDP label for the 10.56.106.0/24 subnet:

```
R5#show mpls ip binding 10.56.106.1 24

R5#
```

But there is no label for this subnet in the binding table. Instead, there is a binding for subnet 10.56.106.1/32:

```
R5#show mpls ip binding
  10.1.16.0/24
        in label:     19
        out label:    imp-null  lsr: 10.56.56.6:0     inuse
  10.1.45.0/24
        in label:     imp-null
        out label:    19        lsr: 10.56.56.6:0
  10.56.56.0/24
        in label:     imp-null
        out label:    imp-null  lsr: 10.56.56.6:0
  10.56.105.0/24
        in label:     imp-null
        out label:    18        lsr: 10.56.56.6:0
  10.56.106.0/24 (no route)
        out label:    imp-null  lsr: 10.56.56.6:0
  10.56.106.1/32
        in label:     20
R5#
```

LDP must find a label that matches a subnet that is connected to an LDP peer. Via BGP, R5 knows that the subnet for IP address 10.56.106.0 is /24, but the only label it has is for a /32 subnet.

Likely cause: R5 does not have a label to the next-hop subnet.

There was reachability between the sites when Loopback 106 was up, but no reachability when Loopback 166 tried to replace it. What changed? If you closely examine the configuration of these two loopback interfaces on R6, you will see that Loopback 166 lacks the command **ip ospf network point-to-point**.

Resolution: Add the required OSPF network type to the EEM commands.

To fix the EEM applet, insert the command **ip ospf network point-to-point**, as you see here:

```
event manager applet LoopbackWatch
 event syslog pattern ".*UPDOWN.*Loopback106.*down"
 action 1.0 syslog priority warning msg "EEM says Loopback 106 is down"
 action 1.1 cli command "enable"
 action 1.2 cli command "configure terminal"
 action 1.3 cli command "interface Loopback166"
 action 1.4 cli command "ip address 10.56.106.1 255.255.255.0"
 action 1.5 cli command "ip ospf network point-to-point"
 action 1.6 cli command "end"
```

The easiest way to do this might be to copy it into a text editor, make the change, delete the applet, and add the corrected version.

7.2. Symptom: A syslog message is generated at the wrong severity level when R6 Loopback 106 goes down.

Analysis and testing:

Examine the details of syslog messages generated on R6 when you shut down the Loopback 106 interface:

```
R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)#int lo 106
R6(config-if)#shut
R6(config-if)#
*Jun 19 13:26:30.494: %LINK-5-CHANGED: Interface Loopback106, changed state to
administratively down
*Jun 19 13:26:31.498: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback106,
changed state to down
*Jun 19 13:26:31.499: %HA_EM-4-LOG: LoopbackWatch: EEM says Loopback 106 is down
R6(config-if)#
```

Note that the syslog message is generated at level 4.

The following table summarizes the syslog severity levels:

Syslog Message Severity

Severity Level	Keyword	Description
0	emergencies	System unusable
1	alerts	Immediate action required
2	critical	Critical condition
3	errors	Error condition
4	warning	Warning condition
5	notifications	Normal but significant event
6	informational	Informational message
7	debugging	Debugging message

When you closely examine the EEM applet, you see that the syslog message is configured as priority "warning."

```
event manager applet LoopbackWatch
 event syslog pattern ".*UPDOWN.*Loopback106.*down"
 action 1.0 syslog priority warning msg "EEM says Loopback 106 is down"
 action 1.1 cli command "enable"
 action 1.2 cli command "configure terminal"
 action 1.3 cli command "interface Loopback166"
 action 1.4 cli command "ip address 10.56.106.1 255.255.255.0"
 action 1.5 cli command "ip ospf network point-to-point"
 action 1.6 cli command "end"
```

Likely cause: The configured syslog priority value needs to be changed.

This lab requires R6 to generate messages with the priority "errors," which corresponds to severity level 3.

Resolution: Change the applet message priority to "errors," "critical," "alerts," or "emergencies."

Change the severity level of the message sourced by the EEM applet:

```
R6#sh run | section event
event manager applet LoopbackWatch
 event syslog pattern ".*UPDOWN.*Loopback106.*down"exit

 action 1.0 syslog priority errors msg "EEM says Loopback 106 is down"
 action 1.1 cli command "enable"
 action 1.2 cli command "configure terminal"
 action 1.3 cli command "interface Loopback166"
 action 1.4 cli command "ip address 10.56.106.1 255.255.255.0"
 action 1.5 cli command "ip ospf network point-to-point"
 action 1.6 cli command "end"
```

With this correction, you should see output similar to the following when R6 Loopback 106 goes down:

```
R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)#int lo 106
R6(config-if)#shut
R6(config-if)#
*Jun 19 13:26:30.494: %LINK-5-CHANGED: Interface Loopback106, changed state to
administratively down
*Jun 19 13:26:31.498: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback106,
changed state to down
*Jun 19 13:26:31.499: %HA_EM-3-LOG: LoopbackWatch: EEM says Loopback 106 is down
R6(config-if)#
```


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**.

8. IP Multicast Troubleshooting Section

8.1. Symptom: R4 does not respond when R2 pings IP multicast address 224.1.1.1.

Analysis and testing:

Here are the results when you ping from R2 Loopback 102 to IP multicast address 224.1.1.1:

```
R2#ping 224.1.1.1 rep 3 source loop102
Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 224.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 10.10.102.1

Reply to request 0 from 10.10.103.1, 21 ms
Reply to request 1 from 10.10.103.1, 1 ms
Reply to request 1 from 10.30.130.1, 25 ms
Reply to request 1 from 10.30.130.1, 25 ms
Reply to request 1 from 10.10.103.1, 1 ms
Reply to request 2 from 10.10.103.1, 1 ms
Reply to request 2 from 10.30.130.1, 1 ms
Reply to request 2 from 10.10.103.1, 1 ms
Reply to request 2 from 10.30.130.1, 1 ms
R2#
```

Replies are received from R3 and SW3, but not from R4. Is SW3 forwarding the traffic to R4? Here is SW3's **mroute** table for this group:

```
SW3#show ip mroute 224.1.1.1
[table legend removed for brevity]

(*, 224.1.1.1), 00:20:20/stopped, RP 10.30.130.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback130, Forward/Sparse, 00:20:19/00:02:13

(10.10.102.1, 224.1.1.1), 00:00:11/00:02:57, flags: LMT
  Incoming interface: Ethernet0/19, RPF nbr 10.1.13.1
  Outgoing interface list:
    Loopback130, Forward/Sparse, 00:00:11/00:02:48
```

Note that the outgoing interface list does not include the link that connects to R4. Since this is sparse mode, the most likely explanation is that R4 did not send a PIM join message for this group to SW3. Here is the mroute table and Source-Active (SA) cache information for this group on R4:

```
R4#sh ip mroute 224.1.1.1
[table legend removed for brevity]

(*, 224.1.1.1), 00:22:48/00:02:58, RP 10.40.104.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback104, Forward/Sparse, 00:22:48/00:02:58

R4#show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
(10.10.102.1,  224.1.1.1),  RP  10.10.101.1,  BGP/AS  1,  00:14:51/00:05:20,  Peer
10.30.130.1
```

Notice that R4 has a client for this group configured on its Loopback 104, and it has an entry for source 10.10.102.1 for this group in its SA cache, but it does not list the (S,G) state for 224.1.1.1. Check the reverse multicast reverse path to the source address:

```
R4#show ip rpf 10.10.102.1
RPF information for ? (10.10.102.1) failed, no route exists
```

No multicast RPF path is available for this source. Check the unicast routing table:

```
R4#show ip route 10.10.102.1
Routing entry for 10.10.102.0/24
  Known via "bgp 4", distance 20, metric 0
  Tag 56, type external
  Last update from 10.1.45.5 00:26:05 ago
  Routing Descriptor Blocks:
  * 10.1.45.5, from 10.1.45.5, 00:26:05 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 56
```

The unicast routing table indicates a next hop of 10.1.45.5 for this prefix, but this path is not enabled for PIM. You need R4 to send its PIM join on the PIM-enabled path, out Et0/0 toward SW3. You could override the unicast routing table with a static **mroute**, but the task requires that you use BGP routing information for RPF lookups. A unicast BGP policy assures that R5 will always be chosen as the unicast next hop to this prefix, and you are not permitted to change it.

The BGP peerings connecting SW3 to R1 and R4 are enabled to support the IPv4 multicast address family. Routes learned via this address family will override routes in the unicast routing table for the purposes of multicast RPF lookups.

Examine R1 for comparison, since the ping sourced from 10.40.140.1 on SW4 to group 224.1.1.1 is answered by R3.

R1 has a unicast route to 10.40.140.1 via R6, which is not PIM-enabled:

```
R1#show ip route 10.40.140.1
Routing entry for 10.40.140.0/24
  Known via "bgp 1", distance 20, metric 0
  Tag 56, type external
  Last update from 10.1.16.6 00:39:58 ago
  Routing Descriptor Blocks:
  * 10.1.16.6, from 10.1.16.6, 00:39:58 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 56
```

However, for the purposes of multicast RPF lookups to this source address, R1 uses the PIM-enabled path through SW3, which it learned from multiprotocol BGP.

```
R1#show ip rpf 10.40.140.1
RPF information for ? (10.40.140.1)
  RPF interface: Ethernet0/0
  RPF neighbor: ? (10.1.13.3)
  RPF route/mask: 10.40.140.0/24
  RPF type: multicast (bgp 1)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
```

On R1, the command **show bgp ipv4 multicast** indicates this learned path, but on R4, this output shows only the locally sourced route.

```
R1#sh bgp ipv4 multicast
BGP table version is 4, local router ID is 10.10.101.1
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
*> 10.40.140.0/24   10.1.13.3                              0 3 4 i
R4#show bgp ipv4 multicast
BGP table version is 6, local router ID is 10.40.104.1
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
*> 10.40.140.0/24   10.40.44.40              1         32768 i
```

Likely cause: Subnet 10.10.102.0/24 is not advertised into MBGP on R1.

When you examine R4's configuration, you see the following partial results:

```
address-family ipv4 multicast
  neighbor 10.1.34.3 activate
  no auto-summary
  network 10.40.140.0 mask 255.255.255.0
 exit-address-family
```

Here is the corresponding configuration on R1:

```
address-family ipv4 multicast
  neighbor 10.1.13.3 activate
  no auto-summary
 exit-address-family
```

Though the IPv4 multicast address family has been activated on R1, the local route has not been advertised into it.

Resolution: On R1, add a network statement for subnet 10.10.102.0/24 under the IPv4 multicast address family of BGP.

Fix the BGP configuration on R1:

```
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#router bgp 1
R1(config-router)# address-family ipv4 multicast
R1(config-router-af)#network 10.10.102.0 mask 255.255.255.0
R1(config-router-af)#end
R1#
```

A few moments after making this change, the route should be installed on R4:

```
R4#show bgp ipv4 multicast
BGP table version is 3, local router ID is 10.40.104.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.10.102.0/24   10.1.34.3                              0 3 1 i
 *>  10.40.140.0/24   10.40.44.40              1         32768 i
R4#
```

R4 will choose its connection to SW3 as its RPF interface to source address 10.10.102.1, and R2 will receive replies from R4 when it pings multicast address 224.1.1.1.

```
R2#ping 224.1.1.1 rep 3 source loop102
Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 224.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 10.10.102.1
```

```
Reply to request 0 from 10.30.130.1, 19 ms
Reply to request 0 from 10.40.104.1, 31 ms
Reply to request 0 from 10.40.104.1, 31 ms
Reply to request 0 from 10.40.104.1, 31 ms
Reply to request 0 from 10.10.103.1, 28 ms
Reply to request 0 from 10.10.103.1, 28 ms
Reply to request 0 from 10.30.130.1, 19 ms
Reply to request 0 from 10.30.130.1, 19 ms
Reply to request 1 from 10.10.103.1, 1 ms
Reply to request 1 from 10.40.104.1, 6 ms
Reply to request 1 from 10.40.104.1, 6 ms
Reply to request 1 from 10.30.130.1, 1 ms
Reply to request 1 from 10.10.103.1, 1 ms
Reply to request 1 from 10.30.130.1, 1 ms
Reply to request 2 from 10.10.103.1, 1 ms
Reply to request 2 from 10.40.104.1, 6 ms
Reply to request 2 from 10.40.104.1, 6 ms
Reply to request 2 from 10.10.103.1, 1 ms
Reply to request 2 from 10.30.130.1, 1 ms
Reply to request 2 from 10.30.130.1, 1 ms
R2#
```

