360RSW05-LAB-TA02-AK

Cisco Expert-Level Training for CCIE Routing and Switching Assessment Lab 2 (CIERS1) Troubleshooting Section

Cisco Expert-Level Training for CCIE® Routing and Switching (R&S) Advanced Workshop 1 introduces a blended learning approach for CCIE preparation.

Advanced Workshop 1 includes an assessment lab for the instructor and student to assess student readiness. This two-hour online troubleshooting assessment lab assesses your grasp of the technology and configuration skills needed for the core topics. You must master the core technologies before moving on to multicast, quality of service (QoS), and security.

The assessment lab tests interpretation skills using the hidden issues format. The in-depth reports provide not only the score but also ample feedback to master those technologies where skills are weak. You can compare performed work with the Master Answer Key and receive detailed explanations in the answer key for any lost points. You will know where to place study emphasis before continuing studies. This assessment saves many hours of study time by making sure that what you think you know is what you actually know.

Cisco Expert-Level Training for CCIE Routing and Switching Assessment Lab 2 (CIERS1) 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.

Cisco Expert-Level Training for CCIE Routing and Switching CCIE R&S Assessment Lab 2 (CIERS1) Troubleshooting Section Answer Key

Note	Regardless of any configuration that you perform in this lab, you must conform to the general guidelines provided. If you do not conform, you can expect a significant deduction of points in your final exam score.

Grading and Duration

- Troubleshooting lab duration: 2 hours
- Troubleshooting lab maximum score: 24 points
- Troubleshooting lab pass score: 19 points

1. Incident 1: Layer 3 — DMVPN

Troubleshooting strategy:

Review the lab diagrams and verify Dynamic Multipoint Virtual Private Network (DMVPN) connectivity, DMVPN mapping, and permanent virtual circuit (PVC) and Local Management Interface (LMI) status.

Note	Refer to the ”Troubleshooting” lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

You verified the connectivity from R12 on the DMVPN subnet 10.20.192.0/24 and found that you cannot ping DMVPN interfaces between R9 and R12.

Troubleshooting:

The ping from R9 to R12 fails, but a ping from R9 to R11 is successful:

```
R9#ping 10.20.192.12

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.20.192.12, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R9#


R9#ping 10.20.192.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.20.192.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/40/44 ms
R9#
```

After examining the lab IPv4 Addressing and IGP Topology diagram, you realize that R9, R11, and R12 are connected to the DMVPN network. You need to find the DMVPN interfaces and verify their status. R9 and R11 can ping each other.

Note that the service provider transport network for the DMVPN is based on the fully operational Frame Relay. You do not need to troubleshoot or change the configuration of service provider Frame Relay network.

Examine the DMVPN interface status on R12:

```
R12#show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
Ethernet0/0            unassigned      YES NVRAM  administratively down down
Ethernet0/1            unassigned      YES NVRAM  administratively down down
Ethernet0/2            unassigned      YES NVRAM  administratively down down
Ethernet0/3            unassigned      YES NVRAM  administratively down down
Serial1/0              1.1.192.12      YES NVRAM  up                    up
Serial1/1              unassigned      YES NVRAM  administratively down down
Serial1/2              unassigned      YES NVRAM  administratively down down
Serial1/3              unassigned      YES NVRAM  administratively down down
```



```
Loopback0              10.1.1.12       YES NVRAM  up                    up
Tunnel129              10.20.192.12    YES NVRAM  up                    up
R12#
```

Note that the Tunnel129 interface that is connected to the DMVPN subnet 10.20.192.0/24 is up.

Verify the output of the s**how ip nhrp** commands on R9, R11, and R12.

```
R9#show ip nhrp
10.20.192.11/32 via 10.20.192.11
  Tunnel129 created 00:01:08, expire 01:58:51
  Type: dynamic, Flags: unique registered used
  NBMA address: 1.1.192.11
10.20.192.12/32
  Tunnel129 created 00:00:18, expire 00:02:46
  Type: incomplete, Flags: negative
  Cache hits: 5
R9#

R11#show ip nhrp
10.20.192.9/32 via 10.20.192.9
  Tunnel129 created 00:02:46, never expire
  Type: static, Flags: used
  NBMA address: 1.1.192.9
R11#

R12#show ip nhrp
10.20.192.9/32 via 10.20.192.9
  Tunnel129 created 00:02:58, never expire
  Type: static, Flags:
  NBMA address: 1.1.192.9
R12#
```

Note that the Tunnel129 interfaces on R11 and R12 are statically mapped to the nonbroadcast multi-access (NBMA) address 1.1.1.192.9 of R9. Therefore R9 is the hub and R11 and R12 are the spokes on DMVPN subnet 10.20.192.0/24. R11 is successfully registered with the DMVPN hub R9, but R12 is not registered with R9.

Verify the DMVPN next-hop server configuration on the DMVPN spokes, R11 and R12:

```
R11#show ip nhrp nhs
Legend: E=Expecting replies, R=Responding, W=Waiting
Tunnel129:
10.20.192.9  RE priority = 0 cluster = 0

R11#
R12#show ip nhrp nhs
R12#
```

Note that R9 10.20.192.9 is configured as a next-hop server on R11, but the next-hop server is missing on R12.

Verify the Tunnel129 interface configuration on R12:

```
R12#show run interface Tunnel129
Building configuration...

Current configuration : 291 bytes
!
```

```
interface Tunnel129
 ip address 10.20.192.12 255.255.255.0
 no ip redirects
 ip nhrp map 10.20.192.9 1.1.192.9
 ip nhrp map multicast 1.1.192.9
 ip nhrp network-id 1
 ip ospf network non-broadcast
 ip ospf priority 0
 tunnel source Serial1/0
 tunnel mode gre multipoint
 tunnel key 1
end

R12#
```

Note that the NHRP next-hop server configuration is missing on R12.

Note	Refer to the "VPN" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Configure the Next Hop Resolution Protocol (NHRP) next-hop server on the Tunnel129 interface of R12. You can run the **debug nhrp** command and observe the NHRP messages:

```
R12#debug nhrp
NHRP protocol debugging is on
R12#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R12(config)#interface Tunnel129
R12(config-if)#ip nhrp nhs 10.20.192.9
R12(config-if)#
*Sep  8 16:12:53.453: NHRP: NHS 10.20.192.9 Tunnel129 vrf 0 Cluster 0 Priority 0 Transitioned to 'E' from ' '

*Sep  8 16:12:53.453: NHRP: Attempting to send packet through interface Tunnel129 via DEST  dst 10.20.192.9
*Sep  8 16:12:53.453: NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 1.1.192.9
*Sep  8 16:12:53.453: NHRP: Send Registration Request via Tunnel129 vrf 0, packet size: 92
*Sep  8 16:12:53.453:      src: 10.20.192.12, dst: 10.20.192.9
*Sep  8 16:12:53.454: NHRP: 120 bytes out Tunnel129
*Sep  8 16:12:53.471: NHRP: Receive Registration Reply via Tunnel129 vrf 0, packet size: 112
*Sep  8 16:12:53.471: NHRP:
R12(config-if)#NHS 10.20.192.9 Tunnel129 vrf 0 Cluster 0 Priority 0 Transitioned to 'RE' from 'E'

*Sep  8 16:12:53.471: NHRP: NHS-UP: 10.20.192.9
R12(config-if)#end
R12#
```

Check the output of the **show ip nhrp nhs** command on R12 again:

```
R12#show ip nhrp nhs
Legend: E=Expecting replies, R=Responding, W=Waiting
Tunnel129:
10.20.192.9  RE priority = 0 cluster = 0

R12#
```

Note that the NHRP next-hop server is configured on R12.

Verify the NHRP registrations on R9:

```
R9#show ip nhrp
10.20.192.11/32 via 10.20.192.11
  Tunnel129 created 00:26:31, expire 01:33:28
  Type: dynamic, Flags: unique registered used
  NBMA address: 1.1.192.11
10.20.192.12/32 via 10.20.192.12
  Tunnel129 created 00:06:26, expire 01:55:19
  Type: dynamic, Flags: unique registered used
  NBMA address: 1.1.192.12
R9#
```

Note that R12 is now successfully registered with R9.

Verify by pinging from R9 to R12 again:

```
R9#ping 10.20.192.12

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.20.192.12, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 40/44/48 ms
R9#
```

Note that the ping is working now.

Note	To obtain a configuration view of the tasks in this and following sections, access the Mentor Guide engine. You can retrieve the available commands by querying the Mentor Guide engine via “Command Line" field.

2. Incident 2: Layer 2 — PPP Multilink

Troubleshooting strategy:

Review the lab diagrams and verify connectivity status on PPP serial links; also, check other requirements of the PPP configuration.

Note	Refer to the ”Troubleshooting” lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

Users who are connected to R14 complain about the bandwidth on the link represented by the 10.30.134.0/24 subnet. The bandwidth must be comprised by the combined bandwidth of four physical links, according to their contract.

Note that the users do not complain about connectivity, just about the bandwidth. Nonetheless, verify the connectivity to start your troubleshooting, and ensure that you do not see any packet drops (“.”) in the ping output:

```
R14#ping 10.30.134.13 rep 1000
```

```
Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 10.30.134.13, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

Note that there are no packet drops in the ping output, therefore the bandwidth degradation is caused by something else.

Troubleshooting:

Examine the lab IPv4 Addressing and IGP Topology diagram and verify what interface the subnet 10.30.134.0/24 is connected to (also, you can use the **show ppp multilink** command to retrieve the information about all multilink interfaces that are configured on the device):

```
R14#show ip route connected
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, 37 subnets, 2 masks
C        10.1.1.14/32 is directly connected, Loopback0
C        10.30.134.0/24 is directly connected, Multilink1314
C        10.30.134.13/32 is directly connected, Multilink1314
L        10.30.134.14/32 is directly connected, Multilink1314
R14#
```

Note that the interface associated with 10.30.134.0/24 is the Multilink PPP (MLP) interface Multilink1314.

Verify the status of the PPP Multilink bundles on the Multilink1314 interface on R13:

```
R13#show ppp multilink interface Multilink1314

Multilink1314
  Bundle name: R14
  Remote Endpoint Discriminator: [1] R14
  Local Endpoint Discriminator: [1] R13
  Bundle up for 1d01h, total bandwidth 4632, load 1/255
  Receive buffer limit 36000 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 649 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x5E0F received sequence, 0x5E1A sent sequence
  Member links: 3 active, 1 inactive (max 255, min not set)
    Se1/3, since 1d01h
    Se1/1, since 1d01h
    Se1/0, since 1d01h
    Se1/2 (inactive)
R13#
```

Note that the interface S1/2 is in the MLP bundle, but it is not active.

Also if you enabled the logging to console on R13 and/or R14, you would see periodic messages similar to the following:

```
R13#
 %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/2, changed state to up
R13#
…%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/2, changed state to down
…: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/2, changed state to up
R13#
…%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/2, changed state to down
… %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/2, changed state to up
R13#
```

The inactive PPP link connected to the S1/2 interface on R13 and R14 in the Multilink1314 bundle may be causing the bandwidth degradation.

Verify the configuration of the interface S1/2 on R13:

```
R13#show run int s1/2
Building configuration...

Current configuration : 126 bytes
!
interface Serial1/2
 no ip address
 encapsulation ppp
 serial restart-delay 0
 ppp multilink
 ppp multilink group 1314
end

R13#
```

Verify the status of the PPP Multilink bundles on the Multilink1314 interface on R14:

```
R14#show ppp multilink interface Multilink1314

Multilink1314
  Bundle name: R13
  Remote Endpoint Discriminator: [1] R13
  Local Endpoint Discriminator: [1] R14
  Bundle up for 03:23:26, total bandwidth 4632, load 1/255
  Receive buffer limit 36000 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 434 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x14E0 received sequence, 0x14CB sent sequence
  Member links: 3 active, 0 inactive (max not set, min not set)
    Se1/3, since 03:23:26
    Se1/1, since 03:23:26
    Se1/0, since 03:23:26
R14#
```

Note that the interface S1/2 is not listed in the MLP bundle. Also, the total bandwidth is 4632 (4632 = 1544 * 3, where 1544 is the T1 bandwidth derived from the interfaces S1/0 through S1/3).

Verify the configuration of the interface S1/2 on R14:

```
R14#show run int s1/2
Building configuration...

Current configuration : 126 bytes
!
```

```
interface Serial1/2
 no ip address
 encapsulation ppp
 serial restart-delay 0
 ppp multilink
 ppp multilink group 1413
end

R14#
```

Note that the multilink group number is not correct on the interface S1/2; it must be 1314 to match the number of the multilink interface.

Configure the correct multilink number on the interface S1/2 on R14:

```
R14#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R14(config)#interface Serial1/2
R14(config-if)# no ppp multilink group 1413
R14(config-if)# ppp multilink group 1314
R14(config-if)#end
R14#
```

Verify the status of the PPP Multilink bundles on the MLP interface on R13 and R14 again:

```
R13#show ppp multilink interface Multilink1314

Multilink1314
  Bundle name: R14
  Remote Endpoint Discriminator: [1] R14
  Local Endpoint Discriminator: [1] R13
  Bundle up for 03:29:37, total bandwidth 6176, load 1/255
  Receive buffer limit 48000 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 389 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x1529 received sequence, 0x153E sent sequence
  Member links: 4 active, 0 inactive (max not set, min not set)
    Se1/3, since 03:29:37
    Se1/1, since 03:29:37
    Se1/0, since 03:29:37
    Se1/2, since 00:01:04
R13#
```

Note that the total bandwidth is changed from 4632 to 6176 (6176 = 1544 * 4, where 1544 is the T1 bandwidth).

```
R14#show ppp multilink interface Multilink1314

Multilink1314
  Bundle name: R13
  Remote Endpoint Discriminator: [1] R13
  Local Endpoint Discriminator: [1] R14
  Bundle up for 03:28:53, total bandwidth 6176, load 1/255
  Receive buffer limit 48000 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 434 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x1531 received sequence, 0x151C sent sequence
  Member links: 4 active, 0 inactive (max not set, min not set)
    Se1/3, since 03:28:53
    Se1/1, since 03:28:53
    Se1/0, since 03:28:53
```



```
    Se1/2, since 00:00:20
R14#
```

3. Incident 3: Layer 3 — MPLS Core

Troubleshooting strategy:

After reviewing the lab "IPv4 Addressing and IGP Topology" and the "MPLS VPN Core and BGP" diagrams, verify connectivity between attached virtual routing and forwarding instances (VRFs), verify VRF configuration, and verify Multiprotocol Label Switching (MPLS) label distribution and forwarding.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

Users are complaining about connectivity via the MPLS VPN core to the Enhanced Interior Gateway Routing Protocol (EIGRP) autonomous system (AS) 1000 domain.

Troubleshooting:

Verify connectivity to EIGRP AS 1000; for example, ping the Loopback0 interface of R15 from the Loopback0 interface of R27:

```
R27#ping 10.1.1.15 source 10.1.1.27

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.15, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.27
.....
Success rate is 0 percent (0/5)
R27#
```

Check the routing tables on R27 and R15:

```
R27#show ip route 10.1.1.15
Routing entry for 10.1.1.15/32
  Known via "rip", distance 120, metric 4
  Redistributing via rip
  Last update from 10.10.26.26 on Ethernet0/0, 00:00:06 ago
  Routing Descriptor Blocks:
  * 10.10.26.26, from 10.10.26.26, 00:00:06 ago, via Ethernet0/0
      Route metric is 4, traffic share count is 1

R27#

R15#show ip route 10.1.1.27
Routing entry for 10.1.1.27/32
  Known via "eigrp 1000", distance 170, metric 332800, type external
  Redistributing via eigrp 1000
  Last update from 10.30.135.13 on Ethernet0/0, 04:03:19 ago
  Routing Descriptor Blocks:
  * 10.30.135.13, from 10.30.135.13, 04:03:19 ago, via Ethernet0/0
      Route metric is 332800, traffic share count is 1
      Total delay is 3000 microseconds, minimum bandwidth is 10000 Kbit
```

```
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2

R15#
```

Note that the Routing Information Protocol (RIP) and EIGRP AS 1000 routes are exchanged via MPLS and the VPN core. Verify the Multiprotocol Border Gateway Protocol (MP-BGP) tables on provider edge (PE) routers R4 and R3:

```
R4#show bgp vpnv4 unicast all | inc 10.1.1.15|27/
*>i10.1.1.15/32    10.1.30.1        435200   100     0 ?
*> 10.1.1.27/32    10.1.26.26            2        32768 ?
R4#

R3#show bgp vpnv4 unicast all | inc 10.1.1.15|27/
*> 10.1.1.15/32    10.1.133.13      435200       32768 ?
*>i10.1.1.27/32    10.1.40.1             2    100     0 ?
R3#
```

Note that VRF VPN prefixes are exchanged between R3 and R4.

When you access R3, you notice the following message:

```
*Sep   1 12:37:11.855: %BGP-4-VPNV4NH_MASK: Nexthop 10.1.30.1 may not be reachable from neighbor 10.1.10.1 - not
/32 mask
```

This message is also logged in the R3 logging buffer:

```
R3#show logging | inc VPN
*Sep   4 21:59:36.951: %BGP-4-VPNV4NH_MASK: Nexthop 10.1.30.1 may not be reachable from neighbor 10.1.10.1 - not
/32 mask
R3#
```

This error message will be explained later in this section.

Verify the VRF configuration and forwarding path across the MPLS VPN core:

```
R4#show ip vrf
  Name                    Default RD        Interfaces
  VPN                     1:1               Et0/3
                                            Se1/0
R4#

R4#show ip route vrf VPN 10.1.1.15

Routing Table: VPN
Routing entry for 10.1.1.15/32
  Known via "bgp 100", distance 200, metric 435200, type internal
  Redistributing via ospf 10, rip
  Advertised by ospf 10 subnets
            rip metric 3
  Last update from 10.1.30.1 03:56:09 ago
  Routing Descriptor Blocks:
  * 10.1.30.1 (default), from 10.1.30.1, 03:56:09 ago
     Route metric is 435200, traffic share count is 1
     AS Hops 0
     MPLS label: 32
     MPLS Flags: MPLS Required
R4#
```


Note that next hop 10.1.30.1 is listed in the default routing table represented by **show ip route**:

```
R4#show ip route 10.1.30.1
Routing entry for 10.1.30.1/32
  Known via "ospf 1", distance 110, metric 21, type intra area
  Last update from 10.1.45.5 on Ethernet0/1, 04:14:45 ago
  Routing Descriptor Blocks:
  * 10.1.45.5, from 10.1.30.1, 04:14:45 ago, via Ethernet0/1
     Route metric is 21, traffic share count is 1

R4#
```

Note that the traffic is forwarded via R5, as it should be according to the lab "IPv4 Addressing and IGP Topology" diagram. R5 is a pure label switching router and it is not configured as a BGP speaker.

Verify the MPLS forwarding table on R5:

```
R5#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
20         No Label   10.1.30.1/32     162741        Et0/3       10.1.35.3
21         Pop Label  10.1.1.4/32      0             Et0/0       10.1.45.4
22         Pop Label  10.1.1.3/32      0             Et0/3       10.1.35.3
23         Pop Label  10.1.1.2/32      0             Et0/2       10.1.25.2
24         Pop Label  10.1.1.1/32      0             Et0/1       10.1.15.1
25         Pop Label  10.1.40.0/24     254148        Et0/0       10.1.45.4
26         Pop Label  10.1.34.0/24     0             Et0/3       10.1.35.3
27         Pop Label  10.1.32.0/24     0             Et0/2       10.1.25.2
28         Pop Label  10.1.20.0/24     254649        Et0/2       10.1.25.2
29         Pop Label  10.1.14.0/24     0             Et0/0       10.1.45.4
30         Pop Label  10.1.12.0/24     0             Et0/2       10.1.25.2
31         Pop Label  10.1.10.0/24     254908        Et0/1       10.1.15.1
R5#
```

Note that the prefix 10.1.30.1/32, which is the address of the Label Distribution Protocol (LDP) peer, is untagged ("No Label"). This would cause packets to be dropped on the outgoing interface E0/3 of R5 connected to R3.

The message that you saw earlier on R3—" *…: %BGP-4-VPNV4NH_MASK: Nexthop 10.1.30.1 may not be reachable from neighbor 10.1.10.1 - not /32 mask"—indicates that:

Excerpt from the Cisco documentation system error messages

%BGP-4-VPNV4NH_MASK—A VPNv4 route is being sent to the IBGP neighbor. The address of the next hop is a loopback interface that does not have a /32 mask defined. OSPF is being used on this loopback interface, and the OSPF network type of this interface is LOOPBACK. OSPF advertises this IP address as a host route (with mask /32), regardless of what mask is configured. This advertising conflicts with LDP, which uses configured masks, so the LDP neighbors may not receive a tag for the route indicated in this error message. This condition could break connectivity between sites that belong to the same VPN.

Verify MPLS bindings on R5:

```
R5#show mpls ip binding | begin ^ +10.1.30
  10.1.30.0/24
      out label:    imp-null  lsr: 10.1.30.1:0
  10.1.30.1/32
      in label:     20
```

```
    out label:    24       lsr: 10.1.20.1:0
    out label:    21       lsr: 10.1.10.1:0
    out label:    21       lsr: 10.1.40.1:0
 10.1.32.0/24
    in label:     27
    out label:    imp-null  lsr: 10.1.20.1:0     inuse
    out label:    27       lsr: 10.1.10.1:0
    out label:    25       lsr: 10.1.30.1:0
    out label:    26       lsr: 10.1.40.1:0
 10.1.34.0/24
    in label:     26
    out label:    29       lsr: 10.1.20.1:0
    out label:    26       lsr: 10.1.10.1:0
    out label:    imp-null  lsr: 10.1.30.1:0     inuse
    out label:    25       lsr: 10.1.40.1:0
 10.1.35.0/24
    in label:     imp-null
    out label:    23       lsr: 10.1.20.1:0
    out label:    20       lsr: 10.1.10.1:0

R5#
```

Note that the 10.1.30.0/24 advertised by LDP from R3 has no matching route in the routing table of R5. R3 sent an implied null for its connected network 10.1.30.0/24, but R5 does not have this network in its routing table:

```
R5#show ip route 10.1.30.0
% Subnet not in table
R5#
```

Verify the Loopback30 IP configuration on R3:

```
R3#show ip int Loopback30
Loopback30 is up, line protocol is up
  Internet address is 10.1.30.1/24
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1514 bytes
  <skipped>
R3#
```

The Loopback30 interface is configured with a /24 mask on R3, but Open Shortest Path First (OSPF) by default advertises it as a host /32 entry:

```
R3#show ip ospf int lo30
Loopback30 is up, line protocol is up
  Internet Address 10.1.30.1/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 10.1.30.1, 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
R3#

R4#show ip route | inc 10.1.30.1
O       10.1.30.1/32 [110/21] via 10.1.45.5, 04:28:23, Ethernet0/1
R4#
R5#show ip route | inc 10.1.30.1
O       10.1.30.1/32 [110/11] via 10.1.35.3, 00:07:33, Ethernet0/3
R5#
```


You cannot freely change the IP addresses and masks according to the lab restrictions, therefore you need to configure the **ip ospf network point-to-point** type on the loopback 30 of R3 to change the default OSPF loopback advertisement and advertise Looback30 as /24 route.

Try to ping from R27 to R15 again:

```
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#int lo30
R3(config-if)#ip ospf network point-to-point
R3(config-if)#end
R3#

R27#ping 10.1.1.15 source 10.1.1.27

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.15, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.27
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/25/36 ms
R27#
#
```

4. Incident 4: Layer 3 — OSPF

After reviewing the lab "IPv4 Addressing and IGP Topology" diagram, verify OSPF neighbor relationships and connectivity.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

After examining the OSPF domain, you realize that the OSPF neighbor relationships are not formed correctly in Area 1 on the subnet 10.20.192.0/24. You may see periodic OSPF messages on R9, for example:

```
R9#
*Sep  9 13:59:18.342: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.1.12 on Tunnel129 from 2WAY to DOWN, Neighbor Down:
Dead timer expired
R9#
*Sep  9 14:03:18.375: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.1.12 on Tunnel129 from 2WAY to DOWN, Neighbor Down:
Dead timer expired
R9#
```

Verify the OSPF neighbor relationships on R9. R9 should form OSPF neighbor relationships with R11 and R12:

```
R9#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.7          1   FULL/DR         00:00:32    10.20.79.7      Ethernet0/0
10.1.1.6          1   FULL/DR         00:00:30    10.20.69.6      Ethernet0/1
10.1.1.11         1   FULL/DR         00:01:51    10.20.192.11    Tunnel129
N/A               0   ATTEMPT/DROTHER     -       10.20.192.12    Tunnel129
R9#
```

Note that R9 unsuccessfully attempts to form a neighbor relationship with R12 and switches between the states ATTEMPT/DROTHER and 2WAY/DROTHER. Also R9 forms the FULL neighbor relationship with R11, but R9 sees R11 as the designated router (DR). You need to investigate this issue.

R9 is the DMVPN hub router in this lab. According to the lab specifications, OSPF control packets (HELLO, Database Exchange, ACKs, and so on) must be exchanged in unicast IP packets and the DR must be on R9, the hub. This combination of requirements suggests that the OSPF network type must be configured as **ip ospf network non-broadcast** on the subnet 10.20.192.0/24. Verify OSPF neighbors on R11:

```
R11#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.9          0   FULL/  -        00:01:43    10.20.192.9     Tunnel129
R11#
```

Note that R11 forms the point-to-point or point-to-multipoint type of adjacency “FULL/ -”, with no DR/BDR/DROTHER references. Although it is FULL, it is not correct, because there is an OPSF network type mismatch between R9 and R11.

Verify the OSPF interfaces on R9 and R11:

```
R9#show ip ospf interface tu129
Tunnel129 is up, line protocol is up
  Internet Address 10.20.192.9/24, Area 1, Attached via Network Statement
  Process ID 1, Router ID 10.1.1.9, 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 1
  Designated Router (ID) 10.1.1.11, Interface address 10.20.192.11
  Backup Designated router (ID) 10.1.1.11, Interface address 10.20.192.11
  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 1/4, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 11, maximum is 11
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.11  (Designated Router)
  Suppress hello for 0 neighbor(s)
R9#

R11#show ip ospf interface tu129
Tunnel129 is up, line protocol is up
  Internet Address 10.20.192.11/24, Area 1, Attached via Network Statement
  Process ID 1, Router ID 10.1.1.11, Network Type POINT_TO_POINT, Cost: 1000
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0          1000     no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:12
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
```



```
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.9
  Suppress hello for 0 neighbor(s)
R11#
```

Note that the OSPF network types are different.

The adjacency is formed as FULL because, during the initialization of the lab, the OSPF timers were configured to match between R9 and R11.

```
R11#show run int tu129
Building configuration...

Current configuration : 296 bytes
!
interface Tunnel129
 ip address 10.20.192.11 255.255.255.0
 no ip redirects
 ip nhrp map 10.20.192.9 1.1.192.9
 ip nhrp map multicast 1.1.192.9
 ip nhrp network-id 1
 ip nhrp nhs 10.20.192.9
 ip ospf hello-interval 30
 tunnel source Serial1/0.192
 tunnel mode gre multipoint
 tunnel key 1
end

R11#
```

Although the OSPF neighbor relationship is FULL, it is not a healthy OSPF adjacency.

Configure the correct OSPF network type on the interface of R11, and set the interface OSPF priority to 0 to ensure that R9 is always chosen as the DR:

```
R11#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R11(config)#int tu129
R11(config-subif)#ip ospf network non-broadcast
R11(config-subif)#ip ospf priority 0
R11(config-subif)#end
R11#
```

Verify the OSPF neighbor relationships on R11 and R9 again:

```
R11#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.9          1   FULL/DR         00:01:56    10.20.192.9     Tunnel129
R11#

R9#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.7          1   FULL/DR         00:00:33    10.20.79.7      Ethernet0/0
10.1.1.6          1   FULL/DR         00:00:39    10.20.69.6      Ethernet0/1
```

```
10.1.1.11       0  FULL/DROTHER  00:01:50  10.20.192.11  Tunnel129
10.1.1.12       0  FULL/DROTHER  00:01:55  10.20.192.12  Tunnel129
R9#
```

5. Incident 5: Layer 3 — EIGRP

Review the lab "IPv4 Addressing and IGP Topology" diagram, and verify EIGRP configuration and connectivity.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

You tried to ping the IP addresses of R18 from R15, for example 10.1.1.18, and you found that the ping was not successful:

```
R15#ping 10.1.1.18

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.18, timeout is 2 seconds:
....
Success rate is 0 percent (0/5)
R15#
```

Troubleshooting:

Some routers have been initialized with the **no logging console** command to hide the periodic error and informational messages. You should enable **logging console** before you attempt the troubleshooting.

On R15, you notice messages similar to the following:

```
R15#
*Sep  9 14:17:56.856: %DUAL-6-NBRINFO: EIGRP-IPv4 1000: Neighbor 10.30.156.15 (Ethernet0/2) is blocked: not on common subnet (10.30.157.15/24)
R15#
*Sep  9 14:18:08.749: %DUAL-6-NBRINFO: EIGRP-IPv4 1000: Neighbor 10.30.156.16 (Ethernet0/2) is blocked: not on common subnet (10.30.157.15/24)
R15#
```

Verify the IP configuration on the E0/1 interface of R15 and the E0/1 interface of R16:

```
R15#show ip int e0/1 | inc address is
  Internet address is 10.30.156.15/24
  Broadcast address is 255.255.255.255
  Helper address is not set
R15#

R16#show ip int e0/1 | inc address is
  Internet address is 10.30.156.16/24
  Broadcast address is 255.255.255.255
  Helper address is not set
R16#
```


Note that the IP addresses on the 10.30.156.0/24 subnet between R15 and R16 are configured correctly.

Verify the IP configuration on the E0/2 interface of R15 and the E0/1 interface of R17:

```
R15#show ip int e0/2 | inc address is
  Internet address is 10.30.157.15/24
  Broadcast address is 255.255.255.255
  Helper address is not set
R15#

R17#show ip int e0/1 | inc address is
  Internet address is 10.30.157.17/24
  Broadcast address is 255.255.255.255
  Helper address is not set
R17#
```

Note that the IP addresses on the 10.30.157.0/24 subnet between R15 and R17 are configured correctly.

Why does R15 complain about R16 and R17 not being on the common subnet?

Try to ping the IP addresses 10.30.156.16 and 10.30.157.17 from R15:

```
R15#ping 10.30.156.16

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.30.156.16, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R15#ping 10.30.157.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.30.157.17, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R15#
```

Note that R15 cannot ping R17.

According to the "Ethernet Switched Cabling Topology" diagram, R15, R16, and R17 are connected via SW1 and SW2. R15 and R16 should share VLAN 10. R15 and R17 should share VLAN20.

Verify VLAN configuration on SW1 and SW2:

```
SW1#show vlan brie

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Et1/0, Et1/1, Et1/2, Et1/3
                                                 Et2/0, Et2/1
10   MARKET                            active    Et0/0
30   ENG                               active    Et0/1, Et0/3
40   GUEST                             active    Et0/2
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup
SW1#
SW2#show vlan brief

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
```

```
1    default                 active    Et1/0, Et1/1, Et1/2, Et1/3
                                       Et2/0, Et2/1
10   MARKET                  active    Et0/0, Et0/1
20   DEVELOP                 active    Et0/2
40   GUEST                   active    Et0/3
1002 fddi-default            act/unsup
1003 trcrf-default           act/unsup
1004 fddinet-default         act/unsup
1005 trbrf-default           act/unsup
SW2#
```

Note that the E0/0 interface of SW2 is placed in VLAN 10, but it should be in VLAN 20, according to the "Ethernet Switched Cabling Topology" diagram.

Fix the VLAN 20 configuration on SW2:

```
SW2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#int e0/0
SW2(config-if)#switchport access vlan 20
SW2(config-if)#end
SW2#

SW2#show vlan brie

VLAN Name                        Status    Ports
---- ----------------------------- --------- -------------------------------
1    default                     active    Et1/0, Et1/1, Et1/2, Et1/3
                                           Et2/0, Et2/1
10   MARKET                      active    Et0/1
20   DEVELOP                     active    Et0/0, Et0/2
40   GUEST                       active    Et0/3
1002 fddi-default                act/unsup
1003 trcrf-default               act/unsup
1004 fddinet-default             act/unsup
1005 trbrf-default               act/unsup
SW2#
```

Verify the EIGRP neighbor relationships on R15:

```
R15#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1000)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
2   10.30.157.17            Et0/2            13 02:11:43   20   200  0  974
0   10.30.156.16            Et0/1            14 15:10:05   15   200  0  11472
1   10.30.135.13            Et0/0            13 16:03:30   15   200  0  22994
R15#
```

Note that R15 forms the EIGRP neighbor relationships with R16 and R17 now, and R15 no longer displays the periodic message about common subnet.

On R15, verify the EIGRP prefixes that are advertised to R15 from R16 and R17:

```
R15#show ip route eigrp | inc via 10.30.15[67]
D       10.1.1.16/32 [90/409600] via 10.30.156.16, 15:11:15, Ethernet0/1
D       10.1.1.17/32 [90/409600] via 10.30.157.17, 02:12:52, Ethernet0/2
D       10.30.178.0/24 [90/307200] via 10.30.157.17, 02:12:52, Ethernet0/2
D       10.30.168.0/24 [90/307200] via 10.30.156.16, 15:11:15, Ethernet0/1
R15#
```


Note that R15 learns only the EIGRP prefixes that are advertised from R16 and R17, but not from R18.

Verify the EIGRP neighbor relationships on R16 and R17:

```
R16#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1000)
H   Address             Interface       Hold Uptime   SRTT   RTO  Q  Seq
                            (sec)       (ms)      Cnt Num
0   10.30.156.15        Et0/1            14 15:16:39   15   200  0  29891
R16#
R17#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1000)
H   Address             Interface       Hold Uptime   SRTT   RTO  Q  Seq
                            (sec)       (ms)      Cnt Num
0   10.30.157.15        Et0/1            14 02:18:38   19   200  0  29904
R17#
```

Note that R16 and R17 form the EIGRP neighbor relationship with R15, but not with R18.

Verify IP connectivity between R16, R17, and R18:

```
R18#ping 255.255.255.255

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 255.255.255.255, timeout is 2 seconds:
.
Reply to request 1 from 10.30.168.16, 1 ms
Reply to request 2 from 10.30.168.16, 1 ms
Reply to request 3 from 10.30.168.16, 4 ms
Reply to request 4 from 10.30.168.16, 4 ms
R18#
```

Note that R18 can ping R16, but not R17.

According to the "Ethernet Switched Cabling Topology" diagram, R17 and R18 are connected via SW1 and SW2. R17 and R18 should share VLAN 40, which is trunked between SW1 and SW2 across the dot1q trunks.

Verify the trunk configuration on SW1 and SW2:

```
SW1#show interfaces trunk

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

Port        Vlans allowed on trunk
Et2/2       10
Et2/3       10

Port        Vlans allowed and active in management domain
Et2/2       10
Et2/3       10

Port        Vlans in spanning tree forwarding state and not pruned
Et2/2       10
Et2/3       10
SW1#
```

```
SW2#show interfaces trunk

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

Port          Vlans allowed on trunk
Et2/2          1-4094
Et2/3          1-4094

Port          Vlans allowed and active in management domain
Et2/2          1,10,20,40
Et2/3          1,10,20,40

Port          Vlans in spanning tree forwarding state and not pruned
Et2/2          1,10,20,40
Et2/3          1,20,40
SW2#
```

Note that only VLAN 10 is allowed on the trunks between SW1 and SW2 on SW1.

Verify the E2/2 and E2/3 interfaces configuration on SW1:

```
SW1#show run int e2/2
Building configuration...

Current configuration : 137 bytes
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10
 switchport mode trunk
 duplex auto
end

SW1#show run int e2/3
Building configuration...

Current configuration : 137 bytes
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10
 switchport mode trunk
 duplex auto
end

SW1#
```

You can add VLAN 40 to the list of the allowed VLANs, or, because there is no specific requirement for manually pruned VLANs in this lab, you can remove the **switchport trunk allowed vlan 10** command from the E2/2 and E2/3 interfaces of SW1 to allow all VLANs through the trunks. Allow VLAN 40 on the trunks on SW1:

```
SW1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#int e2/2
SW1(config-if)#no  switchport trunk allowed vlan 10
SW1(config-if)#int e2/3
SW1(config-if)#no  switchport trunk allowed vlan 10
SW1(config-if)#end
SW1#
```



```
SW1#show interfaces trunk

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

Port        Vlans allowed on trunk
Et2/2       1-4094
Et2/3       1-4094

Port        Vlans allowed and active in management domain
Et2/2       1,10,30,40
Et2/3       1,10,30,40

Port        Vlans in spanning tree forwarding state and not pruned
Et2/2       1,10,30,40
Et2/3       1,10,30,40
SW1#
```

Verify IP connectivity between R18 and R16 and R17 again:

```
R18#ping 255.255.255.255

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 255.255.255.255, timeout is 2 seconds:

Reply to request 0 from 10.30.168.16, 4 ms
Reply to request 1 from 10.30.178.17, 1 ms
Reply to request 1 from 10.30.168.16, 1 ms
Reply to request 2 from 10.30.168.16, 1 ms
Reply to request 2 from 10.30.178.17, 1 ms
Reply to request 3 from 10.30.178.17, 4 ms
Reply to request 3 from 10.30.168.16, 4 ms
Reply to request 4 from 10.30.168.16, 1 ms
Reply to request 4 from 10.30.178.17, 4 ms
R18#
```

Note that R18 can successfully ping R16 and R17.

Issue:

You found only connected routes on R18 and no EIGRP neighbor relationship between R18 and R16 and R17. The logging console is turned off during the lab initialization on R16, R17, and R18.

Troubleshooting:

You verified EIGRP neighbor relationships and the routing table on R16, R17 and R18, and noticed that R18 does not form EIGRP neighbor relationships with R16 and R17. Here is an example of routing table and EIGRP neighbor relationships on R18:

```
R18#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
        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, 2 masks
C        10.1.1.18/32 is directly connected, Loopback0
C        10.30.168.0/24 is directly connected, Ethernet0/0
L        10.30.168.18/32 is directly connected, Ethernet0/0
C        10.30.178.0/24 is directly connected, Ethernet0/1
L        10.30.178.18/32 is directly connected, Ethernet0/1
R18#
R18#show ip eigrp nei
EIGRP-IPv4 Neighbors for AS(1000)
R18#
```

You tried to debug, but you do not see any console messages, because the routers R18, R16, and R17 were initialized with the **no logging console** command. After you configure the **logging console** command, R18 displays periodic messages about the K-value mismatch, if error message logging is configured to output to a router console:

```
R18#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R18(config)#logging con
R18(config)#end
R18#
*Sep   1 18:36:34.594: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1000: Neighbor 10.30.168.16 (Ethernet0/0) is down:
K-value mismatch
*Sep   1 18:36:34.606: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1000: Neighbor 10.30.168.16 (Ethernet0/0) is down:
Interface Goodbye received
R18#
```

EIGRP forms a neighbor relationship only if the K constant values (or EIGRP metric weights) are the same for both neighbors.

Verify **show ip protocols** output on R18, R16, and R17:

```
R18#show ip protocols | inc K
  EIGRP metric weight K1=1, K2=1, K3=1, K4=1, K5=0
R18#

R16#show ip protocols | inc K
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
R16#

R17#show ip protocols | inc K
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
R17#
```

Note that R18 is configured with different EIGRP weights.

Remove the **metric weights** command from R18:

```
R18#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R18(config)#router eigrp 1000
R18(config-router)#no metric weights 0 1 1 1 1 0
R18(config-router)#end
```

Verify the EIGRP neighbor relationships and learned prefixes on R18 again:


```
R18#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1000)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.30.178.17            Et0/1            12 00:01:44  662  3972  0  53
0   10.30.168.16            Et0/0            12 00:01:45   14   200  0  53
R18#

R18#show ip route summary | section eigrp
eigrp 1000      0        49        6768       6272
R18#
```

Note that the actual number of the EIGRP prefixes can vary, but it must be more than zero.

Verify IP connectivity between R15 and R18:

```
R15#ping 10.1.1.18

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.18, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R15#

R15#trace 10.1.1.18

Type escape sequence to abort.
Tracing the route to 10.1.1.18

  1 10.30.157.17 4 msec
    10.30.156.16 12 msec
    10.30.157.17 0 msec
  2 10.30.168.18 0 msec
    10.30.178.18 4 msec *
R15#
```

Note that R15 can ping networks of R18 and the **traceroute** command shows two redundant paths via 10.30.157.17 on R17 and 10.30.156.16 on R16.

6. Incident 6: Layer 3 — IPv6 OSPF Routing

Troubleshooting Strategy:

After reviewing the lab IP version 6 (IPv6) routing diagram, verify IPv6 OSPF neighbors, prefixes, configuration, and connectivity.

Note	Refer to the ”Troubleshooting” lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

You found that you cannot ping between the routers in IPv6 OSPF Area 1 and Area 2.

Troubleshooting:

For example, verify connectivity between R21 and R24:

```
R21#ping 8888:234::24
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8888:234::24, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
R21#
```

Verify the routing table on R21:

```
R21#show ipv6 route
IPv6 Routing Table - default - 4 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
OI  8888:122::/64 [110/1010]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
C   8888:221::/64 [0/0]
     via Ethernet0/0, directly connected
L   8888:221::21/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
R21#
```

Note that the prefix 8888:234::/64 in not in the routing table of R21. Verify whether this prefix is in the routing table of R20:

```
R20#show ipv6 route
IPv6 Routing Table - default - 5 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
C   8888:122::/64 [0/0]
     via Tunnel903, directly connected
L   8888:122::20/128 [0/0]
     via Tunnel903, receive
C   8888:221::/64 [0/0]
     via Ethernet0/0, directly connected
L   8888:221::20/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
R20#
```

R20 is missing this prefix as well. Check the IPv6 OSPF neighbor relationships:

```
R20#show ipv6 ospf neighbor
```



```
Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
10.1.1.19       0   2WAY/DROTHER    00:00:36    12              Tunnel903
10.1.1.21       1   FULL/DR         00:00:35    2               Ethernet0/0
10.1.1.22       1   FULL/BDR        00:00:36    2               Ethernet0/0
R20#
```

Note that the type of the adjacency with R19 is not correct. After examining the DMVPN configuration on R19, R20, and R23, you realize that R19 is the DMVPN hub router. Therefore, it must be the DR for the IPv6 subnet to reliably maintain the network link-state advertisement (LSA). Remember, the IPv6 OSPF network type on the DMVPN subnet is broadcast, according to lab IPv6 routing diagram.

Verify IPv6 OSPF properties on all DMVPN interfaces that are attached to Area 0:

```
R20#show ipv6 ospf interface tu903
Tunnel903 is up, line protocol is up
  Link Local Address FE80:122::20, Interface ID 15
  Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.20
  Network Type BROADCAST, Cost: 1000
  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 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/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)
R20#
R19#show ipv6 ospf interface tu903
Tunnel903 is up, line protocol is up
  Link Local Address FE80:122::19, Interface ID 15
  Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.19
  Network Type BROADCAST, Cost: 1000
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  Designated Router (ID) 10.1.1.23, local address FE80:122::23
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 2, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 2, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.23  (Designated Router)
  Suppress hello for 0 neighbor(s)
R19#
R23#show ipv6 ospf interface tu903
Tunnel903 is up, line protocol is up
  Link Local Address FE80:122::23, Interface ID 15
  Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.23
  Network Type BROADCAST, Cost: 1000
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.1.1.23, local address FE80:122::23
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
```

```
 Graceful restart helper support enabled
 Index 1/1/1, flood queue length 0
 Next 0x0(0)/0x0(0)/0x0(0)
 Last flood scan length is 2, maximum is 4
 Last flood scan time is 0 msec, maximum is 0 msec
 Neighbor Count is 1, Adjacent neighbor count is 1
   Adjacent with neighbor 10.1.1.19
 Suppress hello for 0 neighbor(s)
R23#
```

Note that the spoke R23 became a DR for the DMVPN subnet, but R23 does not have direct connectivity to all OSPF speakers on the Area 0 subnet. Therefore, the network LSA is not complete, and it affects the OSPF routes installation procedures:

```
R23#show ipv6 ospf database network

        OSPFv3 Router with ID (10.1.1.23) (Process ID 1)

            Net Link States (Area 0)

  LS age: 1443
  Options: (V6-Bit E-Bit R-bit DC-Bit)
  LS Type: Network Links
  Link State ID: 12 (Interface ID of Designated Router)
  Advertising Router: 10.1.1.23
  LS Seq Number: 80000001
  Checksum: 0x5542
  Length: 32
      Attached Router: 10.1.1.23
      Attached Router: 10.1.1.19


            Net Link States (Area 2)

  LS age: 919
  Options: (V6-Bit E-Bit R-bit DC-Bit)
  LS Type: Network Links
  Link State ID: 2 (Interface ID of Designated Router)
  Advertising Router: 10.1.1.25
  LS Seq Number: 8000000E
  Checksum: 0x4D17
  Length: 36
      Attached Router: 10.1.1.25
      Attached Router: 10.1.1.24
      Attached Router: 10.1.1.23

R23#
```

Note that R20 is not in the list of the devices of the network LSA for the Area 0 DMVPN subnet.

To ensure that R19 is elected as the DR, change the IPv6 OSPF priority on R19 to greater than zero, and set the priority to 0 on R23:

```
R23#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R23(config)#int tu903
R23(config-subif)#ipv6 ospf priority 0
R23(config-subif)#end
```



```
R19#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R19(config)#int tu903
R19(config-subif)#ipv6 ospf priority 1
R19(config-subif)#end
```

You may have to clear OSPFv3 processes on R19, R20, and R23, or reload these devices.

Verify the IPv6 routing tables on R21, R20, R19, R23, and R24:

```
R21#show ipv6 route | inc 8888:221|234
C   8888:221::/64 [0/0]
L   8888:221::21/128 [0/0]
OI  8888:234::/64 [110/84]
R21#
R20#show ipv6 route | inc 8888:221|234
C   8888:221::/64 [0/0]
L   8888:221::20/128 [0/0]
OI  8888:234::/64 [110/74]
R20#

R19#show ipv6 route | inc 8888:221|234
OI  8888:221::/64 [110/74]
OI  8888:234::/64 [110/74]
R19#

R23#show ipv6 route | inc 8888:221|234
OI  8888:221::/64 [110/74]
C   8888:234::/64 [0/0]
L   8888:234::23/128 [0/0]
R23#
```

Now that the routing path is established, try to ping R24 from R21 again:

```
R21#ping 8888:234::24
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8888:234::24, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 30/37/55 ms
R21#
```

7. Incident 7: Layer 3 — BGP

Troubleshooting Strategy:

After reviewing the lab "IPv4 Addressing and IGP Topology" diagram, verify BGP neighbors, prefixes, configuration, and connectivity.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

The NOC reports a connectivity issue between R15 and networks in the customer BGP AS 1000 and AS 10.

Troubleshooting:

Verify BGP prefixes on R11 and R12:

```
R11#show ip bgp | inc 10.10|30|40
R11#

R12#show ip bgp | inc 10.10|30|40
R12#
```

Check BGP on the upstream neighbor R9:

```
R9#show ip bgp | inc 10.10|30|40
* i10.10.26.0/24    10.1.27.2              0    100      0 100 ?
* i10.10.29.0/24    10.1.27.2              0    100      0 100 ?
* i10.10.127.0/24   10.1.27.2              0    100      0 100 ?
* i10.10.128.0/24   10.1.27.2              0    100      0 100 ?
* i10.30.134.0/24   10.1.27.2              0    100      0 100 ?
* i10.30.134.14/32  10.1.27.2              0    100      0 100 ?
* i10.30.135.0/24   10.1.27.2              0    100      0 100 ?
* i10.30.156.0/24   10.1.27.2              0    100      0 100 ?
* i10.30.157.0/24   10.1.27.2              0    100      0 100 ?
* i10.30.168.0/24   10.1.27.2              0    100      0 100 ?
* i10.30.178.0/24   10.1.27.2              0    100      0 100 ?
* i10.40.122.19/32  10.1.27.2              0    100      0 100 ?
* i10.40.122.20/32  10.1.27.2              0    100      0 100 ?
* i10.40.122.23/32  10.1.27.2              0    100      0 100 ?
* i10.40.221.0/24   10.1.27.2              0    100      0 100 ?
* i10.40.234.0/24   10.1.27.2              0    100      0 100 ?
R9#
```

Note that the BGP networks are not marked as “>” (best).

Verify the interior gateway protocol (IGP) reachability to the next hop 10.1.27.2:

```
R9#show ip route  10.1.27.2
% Subnet not in table
R9#
```

The issue with the next-hop reachability must be fixed in order for R8 and R9 to advertise BGP prefixes to R10, R11, and R12, and in order to provide connectivity from R10, R11, and R12 to the networks that are beyond PEs R3 and R4.

Because you are not allowed to change the preconfigured BGP peer relationships and BGP and OSPF redistributions, you can solve the unreachable next hop on R9 and R8 by configuring the **next-hop-self** keyword on the R7 and R6 BGP neighbor relationships with R9 and R8. Here is an example configuration on R7; a similar configuration must be done on R6:

```
R7#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R7(config)#router bgp 1000
R7(config-router)# neighbor 10.1.1.9 next-hop-self
R7(config-router)# neighbor 10.1.1.8 next-hop-self
R7(config-router)#end

R7#clear ip bgp *
R7#
```


In a few moments, check the tables on R9, R11, and R12:

```
R9#show ip bgp | inc 10.10|30|40
*>i10.10.26.0/24    10.1.1.7             0    100     0 100 ?
*>i10.10.29.0/24    10.1.1.7             0    100     0 100 ?
*>i10.10.127.0/24   10.1.1.7             0    100     0 100 ?
*>i10.10.128.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.134.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.134.14/32  10.1.1.7             0    100     0 100 ?
*>i10.30.135.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.156.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.157.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.168.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.178.0/24   10.1.1.7             0    100     0 100 ?
*>i10.40.122.19/32  10.1.1.7             0    100     0 100 ?
*>i10.40.122.20/32  10.1.1.7             0    100     0 100 ?
*>i10.40.122.23/32  10.1.1.7             0    100     0 100 ?
*>i10.40.221.0/24   10.1.1.7             0    100     0 100 ?
*>i10.40.234.0/24   10.1.1.7             0    100     0 100 ?
R9#

R11#sh ip bgp | inc 10.10|30|40
*>i10.10.26.0/24    10.1.1.7             0    100     0 100 ?
*>i10.10.29.0/24    10.1.1.7             0    100     0 100 ?
*>i10.10.127.0/24   10.1.1.7             0    100     0 100 ?
*>i10.10.128.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.134.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.134.14/32  10.1.1.7             0    100     0 100 ?
*>i10.30.135.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.156.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.157.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.168.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.178.0/24   10.1.1.7             0    100     0 100 ?
*>i10.40.122.19/32  10.1.1.7             0    100     0 100 ?
*>i10.40.122.20/32  10.1.1.7             0    100     0 100 ?
*>i10.40.122.23/32  10.1.1.7             0    100     0 100 ?
*>i10.40.221.0/24   10.1.1.7             0    100     0 100 ?
*>i10.40.234.0/24   10.1.1.7             0    100     0 100 ?
R11#

R12#show ip bgp | inc 10.10|30|40
*>i10.10.26.0/24    10.1.1.7             0    100     0 100 ?
*>i10.10.29.0/24    10.1.1.7             0    100     0 100 ?
*>i10.10.127.0/24   10.1.1.7             0    100     0 100 ?
*>i10.10.128.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.134.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.134.14/32  10.1.1.7             0    100     0 100 ?
*>i10.30.135.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.156.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.157.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.168.0/24   10.1.1.7             0    100     0 100 ?
*>i10.30.178.0/24   10.1.1.7             0    100     0 100 ?
*>i10.40.122.19/32  10.1.1.7             0    100     0 100 ?
*>i10.40.122.20/32  10.1.1.7             0    100     0 100 ?
*>i10.40.122.23/32  10.1.1.7             0    100     0 100 ?
*>i10.40.221.0/24   10.1.1.7             0    100     0 100 ?
*>i10.40.234.0/24   10.1.1.7             0    100     0 100 ?
R12#
```

Try to ping from R11 or R12 again:

```
R12#ping 10.30.135.15

Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.30.135.15, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 40/46/48 ms
R12#
```

Issue:

You found that the External BGP (EBGP) peer relationship is not formed between R8 in AS 1000 and R10 in AS 10.

Troubleshooting:

You verified all BGP peer relationships and found that the BGP peer relationship between R8 and R10 is not established:

```
R8#show ip bgp summary
BGP router identifier 10.1.1.8, local AS number 1000
BGP table version is 444, main routing table version 444
51 network entries using 6120 bytes of memory
99 path entries using 5148 bytes of memory
7/5 BGP path/bestpath attribute entries using 868 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
Bitfield cache entries: current 1 (at peak 2) using 28 bytes of memory
BGP using 12188 total bytes of memory
BGP activity 109/58 prefixes, 337/238 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.1.6        4  1000     573     510      444    0    0 08:27:29       48
10.1.1.7        4  1000     604     517      444    0    0 00:09:05       51
10.1.1.10       4    10       0       0        0    0    0 never    Idle
R8#

R10#sh ip bgp summary
BGP router identifier 10.1.1.10, local AS number 10
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.1.8        4  1000       0       0        0    0    0 never    Idle
R10#
```

Note that both neighbors are in an active state (not established).

Verify **show ip bgp neighbors** output, for example on R10:

```
R10#show ip bgp neighbors 10.1.1.8
BGP neighbor is 10.1.1.8,  remote AS 1000, external link
 BGP version 4, remote router ID 0.0.0.0
 BGP state = Idle
 Last read 00:00:00, last write 00:00:00, hold time is 180, keepalive interval is 60 seconds
 Message statistics:
  InQ depth is 0
  OutQ depth is 0
                 Sent       Rcvd
  Opens:            5          5
  Notifications:    2          3
  Updates:          0         38
  Keepalives:      29         24
  Route Refresh:    0          0
  Total:           36         70
 Default minimum time between advertisement runs is 30 seconds
```



```
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 0/0
  Output queue size: 0
  Index 1, Offset 0, Mask 0x2
  1 update-group member
                        Sent       Rcvd
  Prefix activity:            ----       ----
    Prefixes Current:           0          0
    Prefixes Total:             0          0
    Implicit Withdraw:          0          0
    Explicit Withdraw:          0          0
    Used as bestpath:         n/a          0
    Used as multipath:        n/a          0

                          Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Total:                      0          0
  Number of NLRIs in the update sent: max 0, min 0

  Connections established 5; dropped 5
  Last reset 00:05:11, due to BGP Notification sent, hold time expired
  External BGP neighbor not directly connected.
  No active TCP connection
R10#
```

Note that the EBGP neighbor is not directly connected. The EBGP Time To Live (TTL) value is 1 by default. R8 and R10 form their EBGP peer relationship between the Loopback 0 interfaces. R8 and R10 will decrement the TTL, making EBGP peer-relationship setup impossible.

Increment the TTL value by configuring the **neighbor** *Lo0_IP_Addresss* **ebgp-multihop** command on R8. Perform a similar configuration on R10.

```
R8#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R8(config)#router bgp 1000
R8(config-router)# neighbor 10.1.1.10 ebgp-multihop
R8(config-router)#end
R8#
R10(config)#router bgp 10
R10(config-router)# neighbor 10.1.1.8 ebgp-multihop
R10(config-router)#end
```

Note that the backdoor configuration on R8 is optional because R8 learns network 10.1.1.10/32 via EBGP and OSPF. R8 prefers 10.1.1.10/32 via OSPF. The backdoor configuration on R8 is provided in the answer key for consistency.

In a few seconds, verify the peer relationship:

```
R10#show ip bgp summary
BGP router identifier 10.1.1.10, local AS number 10
BGP table version is 158, main routing table version 158
51 network entries using 6120 bytes of memory
51 path entries using 2652 bytes of memory
3/2 BGP path/bestpath attribute entries using 372 bytes of memory
2 BGP AS-PATH entries using 48 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 9192 total bytes of memory
BGP activity 102/51 prefixes, 102/51 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.1.8        4  1000      21      11      158    0    0 00:00:07       51
R10#show ip bgp
BGP table version is 158, local router ID is 10.1.1.10
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
*> 10.1.1.6/32      10.1.1.8                 0 1000 ?
*> 10.1.1.7/32      10.1.1.8                 0 1000 ?
*> 10.1.1.8/32      10.1.1.8                 0 1000 ?
*> 10.1.1.9/32      10.1.1.8                 0 1000 ?

<skipped>
```

Issue:

The EBGP peer relationship between R8 and R10 is flapping.

Troubleshooting:

You see messages on R8 and R10 similar to the following:

```
R10#
*Sep 13 21:40:52.484: %BGP-5-NBR_RESET: Neighbor 10.1.1.8 reset (Peer closed the session)
*Sep 13 21:40:52.484: %BGP-3-NOTIFICATION: received from neighbor 10.1.1.8 4/0 (hold time expired) 0 bytes
R10#
*Sep 13 21:40:52.484: %BGP-5-ADJCHANGE: neighbor 10.1.1.8 Down Peer closed the session
*Sep 13 21:40:52.484: %BGP_SESSION-5-ADJCHANGE: neighbor 10.1.1.8 IPv4 Unicast topology base removed from session  Peer closed the session
R10#
*Sep 13 21:41:04.315: %BGP-5-ADJCHANGE: neighbor 10.1.1.8 Up
R10#
```

This flapping condition most likely is caused by a routing instability. OSPF prefixes are redistributed into BGP on R6 and R7 during the lab initialization, therefore 10.1.1.8/32 is in both OSPF and BGP routing tables of R6 and R7. R8 rejects 10.1.1.8/32 from R6 and R7. If you ran the **debug ip bgp update** command on R8, you would see a message similar to the following:

```
*Sep 13 16:38:57.715: BGP(0): 10.1.1.6 rcv UPDATE about 10.1.1.8/32 -- DENIED due to: NEXTHOP is our own address;
```

Because you configured the BGP next-hop-self commands on R6 and R7, R8 will learn 10.1.1.8/32 with the following next hops and will not deny it from R6 and R7:

```
R8#show ip bgp 10.1.1.8
BGP routing table entry for 10.1.1.8/32, version 65
Paths: (2 available, best #1, table default, RIB-failure(17))
  Advertised to update-groups:
     7
  Refresh Epoch 1
  Local
    10.1.1.6 (metric 11) from 10.1.1.6 (10.1.1.6)
      Origin incomplete, metric 11, localpref 100, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  Local
    10.1.1.7 (metric 11) from 10.1.1.7 (10.1.1.7)
      Origin incomplete, metric 11, localpref 100, valid, internal
      rx pathid: 0, tx pathid: 0
R8#
```


When the EBGP peer relationship is established between R10 and R8, R10 learns 10.1.1.8/32 via EBGP because of its better administrative distance (AD):

```
R10#show ip bgp | inc 10.1.1.8/32
*> 10.1.1.8/32      10.1.1.8                 0 1000 ?
R10#
```

Note that the EBGP neighbor 10.1.1.8 is currently reachable via the BGP route on R10. The BGP neighbor must not be reachable via the BGP route; it must be reachable via IGP. Therefore, in this lab, the EBGP neighbors 10.1.1.8 and 10.1.1.10 must be reachable via OSPF on R8 and 10.

When EBGP peering is established between R8 and R10, the OSPF route on R10 to 10.1.1.8 (AD = 110) is replaced with the EBGP learned route to 10.1.1.8/32 (AD = 20).

```
R10#show ip route 10.1.1.8
Routing entry for 10.1.1.8/32
  Known via "bgp 10", distance 20, metric 0
  Tag 1000, type external
  Last update from 10.1.1.8 00:00:47 ago
  Routing Descriptor Blocks:
  * 10.1.1.8, from 10.1.1.8, 00:00:47 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 1000
      MPLS label: none

R10#
```

It creates a recursion error condition (the BGP neighbor must not be reachable via the BGP prefix). Cisco IOS Software detects this condition and shuts down the EBGP peer relationship. When EBGP is shut down, the OSPF prefix 10.1.1.8/32 is installed in the routing table:

```
R10#show ip route 10.1.1.8
Routing entry for 10.1.1.8/32
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 10.20.80.8 on Serial1/0, 00:00:17 ago
  Routing Descriptor Blocks:
  * 10.20.80.8, from 10.1.1.8, 00:00:17 ago, via Serial1/0
      Route metric is 65, traffic share count is 1

R10#
```

Cisco IOS Software attempts to form the EBGP peer relationship again and succeeds because the BGP neighbor 10.1.1.8 is reachable via OSPF. Then R10 learns 10.1.1.8/32 via EBGP, and the cycle starts again.

You can also see this recursion issue in the output of the **debug ip routing** command:

```
R10#deb ip routing
IP routing debugging is on
R10#
*Sep  2 17:21:00.027: RT: recursion error routing 10.1.1.8 - probable routing loop
R10#
```

You can filter the BGP advertisements of 10.1.1.8/32 and 10.1.1.10/32 between R8 and R10, or you can configure BGP backdoor networks to change the administrative distance of 10.1.1.8 and 10.1.1.10 from AD = 20 to AD = 200 on R10 and optionally on R8 since R8 is the IBGP peer:

```
R10#conf t
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
R10(config)#router bgp 10
R10(config-router)#network 10.1.1.8 mask 255.255.255.255 backdoor
R10(config-router)#end

R10#show ip bgp | inc 10.1.1.8/32
r> 10.1.1.8/32     10.1.1.8                 0 1000 ?
R10#

R8#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R8(config)#router bgp 1000
R8(config-router)#network 10.1.1.10 mask 255.255.255.255 backdoor
R8(config-router)#end

R8#show ip bgp | inc 10.1.1.10/32
r>i10.1.1.10/32    10.1.1.7            75    100     0 ?
R8#
```

8. Incident 8: Layer 3 — Multicast

Troubleshooting Strategy:

Review the lab "IPv4 Multicast" diagram, verify Protocol Independent Multicast (PIM) neighbor relationships, Internet Group Management Protocol (IGMP) memberships, and multicast connectivity. Ensure that Reverse Path Forwarding (RPF) lookup does not fail for the source of multicast traffic and for the rendezvous point (RP). Verify the PIM mode of operation (dense, sparse, or sparse-dense) and RP configuration such as static, Auto-RP, or bootstrap router (BSR). Verify multicast connectivity.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

You verified all PIM neighbor relationships, and they are configured as specified in the lab, but when you ping 225.5.5.5 from R21, for example, R21 receives replies only from R20.

```
R21#ping 225.5.5.5 rep 5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 225.5.5.5, timeout is 2 seconds:
.
Reply to request 1 from 10.1.1.20, 1 ms
Reply to request 2 from 10.1.1.20, 1 ms
Reply to request 3 from 10.1.1.20, 1 ms
Reply to request 4 from 10.1.1.20, 1 ms
R21#
```


Troubleshooting:

Try running the **debug ip pim** and **clear ip mroute *** commands on R20 and ping again from R21.

```
R20#deb ip pim
PIM debugging is on
R20#clear ip mroute *
R20#


R20#
*Sep  9 15:02:32.913: PIM(0): Received RP-Reachable on Tunnel903 from 10.1.1.19
*Sep  9 15:02:32.913: PIM(0): Received RP-Reachable on Tunnel903 from 10.1.1.19
*Sep  9 15:02:32.913:      for group 225.5.5.5
*Sep  9 15:02:32.913: PIM(0): RP not known, or mismatch
*Sep  9 15:02:32.913: PIM(0): Received RP-Reachable on Tunnel903 from 10.1.1.19
*Sep  9 15:02:32.914: PIM(0): Received RP-Reachable on Tunnel903 from 10.1.1.19
*Sep  9 15:02:32.914:      for group 224.0.1.40
*Sep  9 15:02:32.914: PIM(0): RP not known, or mismatch
R20#
*Sep  9 15:03:13.382: PIM(0): Building Periodic (*,G) Join / (S,G,RP-bit) Prune message for 225.5.5.5
*Sep  9 15:03:14.181: PIM(0): Building Periodic (*,G) Join / (S,G,RP-bit) Prune message for 224.0.1.40
R20#
```

Check the RP configuration on the router R20:

```
R20#show ip pim rp map
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
   RP: 100.1.1.19 (?)
R20#
```

There is a typo, a bug that was introduced during the lab initialization. Correct the RP configuration on R20, and ping from R21:

```
R20#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R20(config)#no ip pim rp-address 100.1.1.19
R20(config)# ip pim rp-address 10.1.1.19
R20(config)#end
R20#
R21#ping 225.5.5.5 rep 5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 225.5.5.5, timeout is 2 seconds:

Reply to request 0 from 10.1.1.20, 4 ms
Reply to request 0 from 10.1.1.23, 84 ms
Reply to request 1 from 10.1.1.20, 1 ms
Reply to request 1 from 10.1.1.19, 44 ms
Reply to request 2 from 10.1.1.20, 1 ms
Reply to request 2 from 10.1.1.19, 44 ms
Reply to request 3 from 10.1.1.20, 1 ms
Reply to request 3 from 10.1.1.19, 44 ms
Reply to request 4 from 10.1.1.20, 1 ms
Reply to request 4 from 10.1.1.19, 44 ms
R21#
```

Note that R21 now starts receiving responses from R20 (10.40.221.20) and from R19 (10.40.122.19). Also, at the beginning, R21 received only a couple of replies from R23. Most likely, these replies from R23 were generated when the RP (R19) was sending the multicast packets encapsulated in the unicast packets during the registration process.

When R19 switches to a multicast delivery, the multipoint DMVPN interface becomes an issue. This interface is incoming for the multicast traffic arriving to R19 from R20 and cannot be outgoing at the same time for the traffic from R19 to R23—see the diagram below. Configure **ip pim nbma-mode** on the DMVPN interface of R19:

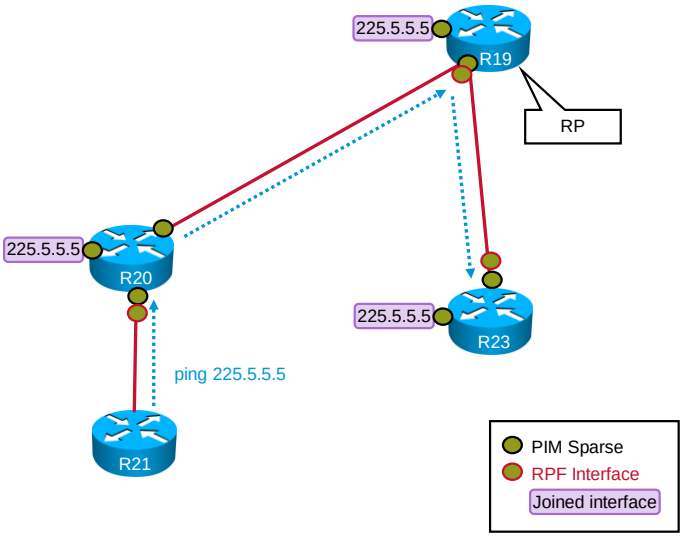

The following output represents the multicast routing table before the **ip pim nbma-mode** command is configured:

```
R19#show ip mroute 225.5.5.5
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.5.5.5), 01:16:56/00:03:20, RP 10.1.1.19, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel903, Forward/Sparse, 00:11:34/00:03:20
```



```
    Loopback0, Forward/Sparse, 01:16:55/00:02:12

(10.40.221.21, 225.5.5.5), 00:01:08/00:01:55, flags: LT
  Incoming interface: Tunnel903, RPF nbr 10.40.122.20
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:01:08/00:02:12

R19#
```

The **ip pim nbma-mode** command should be configured on R19:

```
R19#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R19(config)#int Tu903
R19(config-subif)#ip pim nbma-mode
R19(config-subif)#end
```

The following output represents the multicast routing table after the **ip pim nbma-mode** command is configured:

```
R19#show ip mroute 225.5.5.5
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.5.5.5), 01:19:24/00:03:12, RP 10.1.1.19, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel903, 10.40.122.20, Forward/Sparse, 00:00:40/00:02:48
    Tunnel903, 10.40.122.23, Forward/Sparse, 00:01:15/00:03:12
    Loopback0, Forward/Sparse, 01:19:23/00:02:38

(10.40.221.21, 225.5.5.5), 00:00:29/00:03:04, flags: LT
  Incoming interface: Tunnel903, RPF nbr 10.40.122.20
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:00:29/00:02:38
    Tunnel903, 10.40.122.23, Forward/Sparse, 00:00:29/00:03:12

R19#
R21#ping 225.5.5.5 rep 5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 225.5.5.5, timeout is 2 seconds:

Reply to request 0 from 10.1.1.20, 1 ms
Reply to request 0 from 10.1.1.23, 72 ms
Reply to request 1 from 10.1.1.20, 1 ms
Reply to request 1 from 10.1.1.23, 80 ms
Reply to request 1 from 10.1.1.19, 40 ms
Reply to request 2 from 10.1.1.20, 1 ms
Reply to request 2 from 10.1.1.23, 80 ms
```

```
Reply to request 2 from 10.1.1.19, 44 ms
Reply to request 3 from 10.1.1.20, 1 ms
Reply to request 3 from 10.1.1.23, 72 ms
Reply to request 3 from 10.1.1.19, 36 ms
Reply to request 4 from 10.1.1.20, 4 ms
Reply to request 4 from 10.1.1.23, 88 ms
Reply to request 4 from 10.1.1.19, 44 ms
R21#
```

Now you receive replies from all three routers—R20, R19, and R23.

9. Incident 9: Application — QoS

After reviewing the lab "IPv4 Addressing and IGP Topology" diagram, verify configuration and connectivity.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

The service policy is not applied on the egress direction. Also, the IP differentiated services code point (DSCP) is not set to the value that is required in the lab.

Troubleshooting:

Verify the policy map on R8:

```
R8#show policy-map interface
 Serial1/0

  Service-policy input: Finger-Policy

    Class-map: Finger (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: protocol finger
      QoS Set
        dscp ef
          Packets marked 0

    Class-map: class-default (match-any)
      708 packets, 54608 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

Note that the DSCP value EF in the QoS set action is not correct, according to the lab requirements. Also, the direction of the service policy is not correct; it must be output instead of input.

Configure the correct DSCP value 56 and change the direction for the service policy:


```
R8#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R8(config)#policy-map Finger-Policy
R8(config-pmap)# class Finger
R8(config-pmap-c)#  set ip dscp 56
R8(config-pmap-c)#end
R8#
```

Note that the Cisco IOS Software will translate the DSCP value “56” to the value “cs7” and will change it in the router configuration:

```
policy-map Finger-Policy
 class Finger
  set ip dscp cs7
```

```
R8(config)#int s1/0
R8(config-if)#no service-policy input Finger-Policy
R8(config-if)#service-policy output Finger-Policy
R8(config-if)#end
R8#
```

Verify that the classification and marking are working according to the lab requirements. (This verification is optional.)

Run the **telnet 10.20.80.10 finger** command from R8:

```
R8#telnet 10.20.80.10 finger
Trying 10.20.80.10, 79 ... Open

-------------------------------------------------
Cisco 360 R&S Assessment, Troubleshooting
Product, POD location:
Device:              R10
-------------------------------------------------

    Line      User      Host(s)          Idle      Location
   0 con 0              idle             00:01:31   CIERS1-TU-LAB02, VIRT
*  2 vty 0              idle             00:00:00 10.20.80.8

  Interface    User           Mode       Idle     Peer Address

[Connection to 10.20.80.10 closed by foreign host]
R8#
```

Verify the output of the **show policy-map interface** command on R8:

```
R8#show policy-map interface
 Serial1/0

  Service-policy output: Finger-Policy

    Class-map: Finger (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name FINGER
      QoS Set
        dscp cs7
          Packets marked 0

    Class-map: class-default (match-any)
      18798 packets, 1404822 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
```

```
    Match: any
R8#
```

Note that although the QoS service policy is fixed, you do not see any packets marked with DSCP 56 (cs7).

Verify the match criteria:

```
R8#show access-lists
Extended IP access list FINGER
   10 permit tcp any any eq 78
R8#
```

Note that the TCP port 78 is incorrect. The TCP port assigned to the FINGER protocol is 79, you can see it from the following output:

```
R8#telnet 10.20.80.10 finger
Trying 10.20.80.10, 79 ... Open
```

Fix the port number in the access list FINGER:

```
R8#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R8(config)#ip access-list extended FINGER
R8(config-ext-nacl)#no  permit tcp any any eq 78
R8(config-ext-nacl)# permit tcp any any eq 79
R8(config-ext-nacl)#end
R8#
```

Run the **telnet 10.20.80.10 finger** command from R8 again and check the output of the **show policy-map interface** command:

```
R8#telnet 10.20.80.10 finger
Trying 10.20.80.10, 79 ... Open

-------------------------------------------------
Cisco 360 R&S Assessment, Troubleshooting
Product, POD location:
Device:              R10
-------------------------------------------------

   Line      User      Host(s)          Idle      Location
  0 con 0              idle             00:24:37   CIERS1-TU-LAB02, VIRT
*  2 vty 0             idle             00:00:00 10.20.80.8

  Interface    User           Mode       Idle     Peer Address

[Connection to 10.20.80.10 closed by foreign host]

R8#show policy-map interface
 Serial1/0

  Service-policy output: Finger-Policy

    Class-map: Finger (match-all)
      5 packets, 224 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name FINGER
      QoS Set
        dscp cs7
          Packets marked 5

    Class-map: class-default (match-any)
      18944 packets, 1415486 bytes
```



```
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: any
R8#
```

Note that the TCP port 79 packets are matched and marked with DSCP 56 (cs7).

10. Incident 10: Services — NTP

After reviewing the lab "IPv4 Addressing and IGP Topology" diagram, verify configuration and connectivity.

Note	Refer to the "Troubleshooting" lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

The Network Time Protocol (NTP) association on R25 is insane and invalid.

Troubleshooting:

Verify the NTP association on R23, R24, and R25:

```
R23#show ntp associations detail
127.127.1.1 configured, ipv4, our_master, sane, valid, stratum 14
ref ID .LOCL., time D5E6BA69.A872B1F0 (12:13:29.658 PST Fri Sep 20 2013)
our mode active, peer mode passive, our poll intvl 16, peer poll intvl 16
root delay 0.00 msec, root disp 0.00, reach 377, sync dist 2.36
delay 0.00 msec, offset 0.0000 msec, dispersion 1.20, jitter 0.97 msec
precision 2**10, version 4
assoc id 44519, assoc name 127.127.1.1
assoc in packets 41, assoc out packets 41, assoc error packets 0
org time D5E6BA69.A872B1F0 (12:13:29.658 PST Fri Sep 20 2013)
rec time 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
xmt time D5E6BA69.A872B1F0 (12:13:29.658 PST Fri Sep 20 2013)
filtdelay =     0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00
filtoffset =    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00
filterror =     0.97    1.21    1.45    1.69    1.93    2.17    2.41    2.65
minpoll = 4, maxpoll = 4

R23#
R24#sho ntp associations detail
10.40.234.23 configured, ipv4, our_master, sane, valid, stratum 15
ref ID 127.127.1.1    , time D5E6B929.A872B1F0 (12:08:09.658 PST Fri Sep 20 2013)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 2.19, reach 37, sync dist 6.94
delay 0.00 msec, offset 0.0000 msec, dispersion 2.87, jitter 0.97 msec
precision 2**10, version 4
assoc id 24860, assoc name 10.40.234.23
assoc in packets 10, assoc out packets 11, assoc error packets 0
org time 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
rec time D5E6B92C.974BC848 (12:08:12.591 PST Fri Sep 20 2013)
xmt time D5E6B92C.974BC848 (12:08:12.591 PST Fri Sep 20 2013)
```

```
filtdelay =    0.00   0.00   0.00   0.00   1.00   1.00   1.00   1.00
filtoffset =   0.00   0.00   0.00   0.00   0.50   0.50   0.50   0.50
filterror =    1.95   2.95   3.99   5.02   5.88   5.91   5.94   5.97
minpoll = 6, maxpoll = 10

R24#

R25#show ntp associations detail
10.40.234.24 configured, ipv4, insane, invalid, stratum 16
ref ID 10.40.234.23   , time D5E6B9B7.974BC848 (12:10:31.591 PST Fri Sep 20 2013)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 6.25, reach 0, sync dist 15945.58
delay 0.00 msec, offset 0.0000 msec, dispersion 15937.50, jitter 0.00 msec
precision 2**10, version 4
assoc id 49592, assoc name 10.40.234.24
assoc in packets 15, assoc out packets 16, assoc error packets 15
org time 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
rec time D5E6B9BA.2D4FDFB8 (12:10:34.177 PST Fri Sep 20 2013)
xmt time D5E6B9BA.2D4FDFB8 (12:10:34.177 PST Fri Sep 20 2013)
filtdelay =    0.00   0.00   0.00   0.00   0.00   0.00   0.00   0.00
filtoffset =   0.00   0.00   0.00   0.00   0.00   0.00   0.00   0.00
filterror =  16000.0 16000.0 16000.0 16000.0 16000.0 16000.0 16000.0 16000.0
minpoll = 6, maxpoll = 10

R25#
```

Note that the NTP client on R24 is synchronized with the NTP master on R23, but R25 is not synchronized with R24.
Note that the NTP master advertises stratum 15 to R24. This means that the NTP client on R25 will have stratum 16, which is an invalid stratum.

You are not allowed to change the preconfigured NTP client server relationships between R23, R24, and R25.

Verify the NTP master configuration on R23, and set the stratum on R23 to a value less than 15—for example, 14, as configured in this answer key.

```
R23#show run | inc ntp
ntp master 15
R23#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R23(config)#ntp master 14
R23(config)#end
R23#
```

Wait for a few minutes before checking the NTP association again:

```
R24#sho ntp associations detail
10.40.234.23 configured, ipv4, our_master, sane, valid, stratum 14
ref ID 127.127.1.1    , time D5E6BD4F.A872B1F0 (12:25:51.658 PST Fri Sep 20 2013)
our mode client, peer mode server, our poll intvl 128, peer poll intvl 128
root delay 0.00 msec, root disp 2.34, reach 377, sync dist 9.55
delay 0.00 msec, offset 0.0000 msec, dispersion 5.56, jitter 0.97 msec
precision 2**10, version 4
assoc id 24860, assoc name 10.40.234.23
assoc in packets 27, assoc out packets 28, assoc error packets 0
org time 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
rec time D5E6BD5C.974BC848 (12:26:04.591 PST Fri Sep 20 2013)
xmt time D5E6BD5C.974BC848 (12:26:04.591 PST Fri Sep 20 2013)
filtdelay =    1.00   1.00   1.00   1.00   1.00   1.00   0.00   1.00
```



```
filtoffset =   0.50   0.50   0.50   0.50  -0.50   0.50   0.00   0.50
filterror =    1.95   1.98   3.93   3.96   5.98   6.01   8.01   8.04
minpoll = 6, maxpoll = 10

R24#

R25#show ntp associations detail
10.40.234.24 configured, ipv4, our_master, sane, valid, stratum 15
ref ID 10.40.234.23   , time D5E6BBC8.974BC848 (12:19:20.591 PST Fri Sep 20 2013)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 7951.11, reach 377, sync dist 7956.51
delay 1.00 msec, offset 0.5000 msec, dispersion 2.92, jitter 0.97 msec
precision 2**10, version 4
assoc id 49592, assoc name 10.40.234.24
assoc in packets 31, assoc out packets 32, assoc error packets 23
org time 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
rec time D5E6BDBE.2D4FDFB8 (12:27:42.177 PST Fri Sep 20 2013)
xmt time D5E6BDBE.2D4FDFB8 (12:27:42.177 PST Fri Sep 20 2013)
filtdelay =    1.00   1.00   0.00   0.00   1.00   1.00   1.00   1.00
filtoffset =   0.50   0.50   0.00   0.00   0.50  -0.50   0.50   0.50
filterror =    1.95   2.97   3.99   4.99   6.00   7.02   8.04   9.00
minpoll = 6, maxpoll = 10

R25#
```

11. Incident 11: Security — MPPE Ecryption

After reviewing the lab “IPv4 Addressing and IGP Topology” diagram, verify configuration and connectivity.

Note	Refer to the ”Troubleshooting” lesson available at the Cisco Expert-Level Training Program for the CCIE Routing and Switching portal.

Issue:

Packets between R28 and R29 must be exchanged in the Microsoft Point-to-Point Encryption (MPPE) encrypted format.

Troubleshooting:

Verify the interfaces S1/1 on R29 and S1/0 on R28:

Verify the status of MPPE on the interfaces S1/1 of R29 and S1/0 of R28:

```
R29#show ppp mppe s1/1
```

```
MPPE not configured on Se1/1

R29#s


R28#show ppp mppe s1/0
MPPE not configured on Se1/0

R28#
```

Note that the MPPE encryption is not configured on the serial link between R28 and R29.

Verify the S1/0 interface configuration on R28 and S1/1 interface configuration on R29:

```
R28#show run int s1/0
Building configuration...

Current configuration : 139 bytes
!
interface Serial1/0
 ip address 10.10.128.28 255.255.255.0
 encapsulation ppp
 serial restart-delay 0
 ppp encrypt mppe 40 required
end

R28#
R29#show run int s1/1
Building configuration...

Current configuration : 139 bytes
!
interface Serial1/1
 ip address 10.10.128.29 255.255.255.0
 encapsulation ppp
 serial restart-delay 0
 ppp encrypt mppe 40 required
end

R29#
```

Note that the MPPE is configured on the interfaces, but it is not operational. You have to configure PPP authentication to enable the MPPE encryption; for example, you can use a preconfigured **username R28 password 0 cisco** command on R29:

```
R29#sh run | inc username
username R28 password 0 cisco

R29#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R29(config)#int s1/1
R29(config-if)#ppp authentication ms-chap
R29(config-if)#end
R29#
```

Note that the authentication can be configured bidirectionally, but it is not required in this lab.

```
R28#deb ppp mpp events
MPPE Events debugging is on
R28#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R28(config)#int s1/0
R28(config-if)#ppp chap password cisco
```



```
R28(config-if)#end
R28#
*Sep 20 11:58:19.243: Se1/0 MPPE: Generate keys using local database
*Sep 20 11:58:19.243: Se1/0 MPPE: Initialize keys
*Sep 20 11:58:19.243: Se1/0 MPPE: [40 bit encryption]  [stateless mode]

R28#
*Sep   1 23:15:03.271: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/0, changed state to up
R28#
```

Check the status of MPPE encryption:

```
R29#show int s1/1 | inc Encap|Listen|Open
  Encapsulation PPP, LCP Open
  Open: IPCP, CCP, CDPCP, crc 16, loopback not set
R29#


R29#show ppp mppe s1/1
Interface Serial1/1 (current connection)
  Software encryption, 40 bit encryption, Stateless mode
  packets encrypted = 20      packets decrypted  = 19
  sent CCP resets   = 0       receive CCP resets = 0
  next tx coherency = 20      next rx coherency  = 19
  tx key changes    = 20      rx key changes     = 19
  rx pkt dropped    = 0       rx out of order pkt= 0
  rx missed packets = 0

R29#

R28#show int s1/0 | inc Encap|Listen|Open
  Encapsulation PPP, LCP Open
  Open: IPCP, CCP, CDPCP, crc 16, loopback not set
R28#

R28#show ppp mppe s1/0
Interface Serial1/0 (current connection)
  Software encryption, 40 bit encryption, Stateless mode
  packets encrypted = 19      packets decrypted  = 20
  sent CCP resets   = 0       receive CCP resets = 0
  next tx coherency = 19      next rx coherency  = 20
  tx key changes    = 19      rx key changes     = 20
  rx pkt dropped    = 0       rx out of order pkt= 0
  rx missed packets = 0

R28#
```

Connectivity Verification Script

Troubleshooting strategy:

Review the lab “IPv4 Addressing and IGP Topology” diagram and verify connectivity between the IP addresses.

A Tool Command Language (Tcl) script can be used to test universal reachability. To use the script, enter the **tclsh** command in privileged mode and paste in this script. To kill failing pings,

hold down the **Ctrl-Shift** keys while pressing the **6** key twice. After you are done, enter the **tclquit** command to leave Tcl mode.

```
foreach address {
10.1.1.6
10.1.1.7
10.1.1.8
10.1.1.9
10.1.1.10
10.1.1.11
10.1.1.12
10.1.1.13
10.1.1.14
10.1.1.15
10.1.1.16
10.1.1.17
10.1.1.18
10.1.1.19
10.1.1.20
10.1.1.21
10.1.1.22
10.1.1.23
10.1.1.24
10.1.1.25
10.1.1.26
10.1.1.27
10.1.1.28
10.1.1.29} {ping $address}
```

You must ensure that the solution is stable. If there are split-horizon or other route-feedback problems, routes may continually be inserted and removed from the routing tables. To test stability, observe the output of the **debug ip routing** command.

Troubleshooting Issues Summary

	Issue: Brief Description	Solution: Brief Description
Incident 1: Layer 3 – DMVPN		
1	Interface Tu129 of R12 is configured with NHRP NHS.	Configure the **ip nhrp nhs** command on the tunnel interface on R12.
Incident 2: Layer 2 – PPP Multilink		
2	The PPP Multilink bundle over multiple serial links between R13 and R14 is not used at full capacity. The PPP Multilink group number on interface S1/2 of R14 is not correct	Correct the PPP Multilink group configuration on R14.
Incident 3: Layer 3 – MPLS core		
3	There is a connectivity issue between RIP, EIGRP, and OSPF VPN sites. The LDP endpoint is not advertised into OSPF correctly on R3.	Configure **ip ospf network point-to-point** on the loopback interfaces of R3 and R4 involved in the LDP neighbor relationships.
Incident 4: Layer 3 – OSPF		
4	OSPF adjacency between R9 and R11 is misleading. It is FULL but is between incompatible OSPF network types.	Configure **ip ospf network non-broadcast** on the interface of R11 connected to 10.20.192.0/24
Incident 5: Layer 3 – EIGRP		
5	An EIGRP K1, K2, K3, K4, and K5 value mismatch exists between R18 and its neighbors.	Remove the **metric weights** command from R18.
Incident 6: Layer 3 – IPv6 Routing		
6	The IPv6 OSPF DR was configured on the wrong router, R23.	Move the IPv6 OSPF DR from R23 to R19.
Incident 7: Layer 3 – BGP		
7	No BGP peer relationship between R1 and R2; EBGP TTL issue. Routing recursion.	Configure the BGP **ebgp-multihop** command on R1 and R2. Configure a BGP backdoor network on R8 and R10.
8	There is a connectivity issue between the networks connected to PEs R3 and R4 and networks connected to R7 in BGP AS 1000.	Configure **next-hop-self** on the BGP neighbor statement from R7 to R9.
Incident 8: Layer 3 – Multicast		
9	The multicast PIM RP IP address is misconfigured on R20: **ip pim rp-address 100.1.1.19**	Correct the RP IP address: **ip pim rp-address 10.1.1.19**
10	R23 does not respond to a multicast ping from R21, the command **ip pim nbma-mode** is missing.	Configure the **ip pim nbma-mode** command on the R19 DMVPN multipoint interface.
Incident 9: Application – QoS		

	Issue: Brief Description	Solution: Brief Description
11	The service policy is applied on ingress direction on R8 instead of egress.	Change the keyword **input** to **output** on the service policy command on R8.
12	The DSCP was set to a wrong value EF, TCP port 78 is wrong.	Change it to **ip dscp 56.** Change the TCP port to 79.
Incident 10: Services – NTP		
13	NTP master stratum is set to 15 on R23. It makes the NTP client stratum set to 16, which is an invalid stratum value.	Change the stratum value on R23 to a value lower than 15—for example, 14.
Incident 11: Security – MPPE encryption		
14	MPPE is not negotiated, packets are not encrypted.	Fix **ppp ms-chap** authentication to enable the MPPE negotiation.

