360RSW05-WKB-TA01

Cisco 360 CCIE R&S Exercise Workbook Introduction

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

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

Cisco 360 CCIE R&S Exercise Workbook Lab 1 Troubleshooting Section Answer Key

Table of Contents

Answer Key Structure

Section One

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

Section Two

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

Exercise Workbook Lab 1 Troubleshooting Section Answer Key

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

Grading and Duration

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

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

Difficulty Level

- Difficulty: Basic to Intermediate

Restrictions and Goals

Note	Read this section carefully.

- To receive any credit for a subsection, you must fully complete the subsection as per the requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the "IPv4 IGP" diagram belong to network 172.10.0.0/16. *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**).
- Do not use the **ip default-gateway** or **ip default-network** commands.
- Do not introduce any new IP addresses.
- All the IP addresses that are involved in this scenario must be reachable, unless explicitly specified otherwise.
- Unless explicitly specified otherwise, addresses and networks that are advertised in the "BGP" section need to be reachable by all BGP routers but do not have to be reachable by routers that use only IGP.

- Use conventional routing algorithms only unless specified otherwise.
- Do not create new interfaces to fulfill IGP requirements, and do not summarize unless you are explicitly asked to do so.
- Do not modify the hostname, console, or vty configuration unless you are specifically asked to do so.
- Do not modify the initial interface or IP address numbering.

Explanation of Each of the Restrictions and Goals

IPv4 subnets that are displayed in the scenario diagram belong to network 172.10.0.0/16.

All the IP addresses in this lab belong to the 172.10.0.0/16 address space, except for prefixes that are explicitly specified as being part of a different IP space.

Do not use any static routes.

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

Advertise loopback interfaces with their original masks.

The original mask is the mask that is configured on the loopback interface. Open Shortest Path First (OSPF) treats loopback interfaces as host routes by default and advertises them as /32 prefixes. The requirement to advertise loopback interfaces with their original masks precludes using the default OSPF network type for the loopback interfaces. You need to provide a solution such as changing the OSPF network type or summarizations. Remember that this rule applies to both IPv4 and IPv6 networks.

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

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

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

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

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

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

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

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

Use conventional routing algorithms unless specified otherwise.

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

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

Because of this restrictive interpretation, no form of policy routing can be applied. Whenever you see this restriction, you will need to use dynamic routing protocols to fulfill all packet-forwarding requirements.

1. DMVPN Troubleshooting

1.1. Symptom: R2 cannot ping R4.

Analysis and Testing:

Ping the DMVPN IP addresses 172.10.124.1 and 172.10.124.4 from R2:

```
R2#ping 172.10.124.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.124.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R2#

R2#ping 172.10.124.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.124.4, timeout is 2 seconds:

*Apr   23   23:55:01.483:   %NHRP-3-PAKERROR:   Received   Error   Indication   from
172.10.124.1, code: protocol generic error(7), (trigger src: 172.10.124.2 (nbma:
10.10.1.2) dst: 172.10.124.4), offset: 0, data: 00 01 08 00 00 00 00 00 00 FF 00 48
68 02 00 34 ...

<skipped for brevity>

R2#
```

Note that the DMVPN spoke R2 can ping the DMVPN hub R1, but R2 cannot ping the other DMVPN spoke, R4.

Likely Cause: There is a DMVPN or NHRP misconfiguration on R1 or R4.

Verify the NHRP registrations on R1**:**

```
R1#show ip nhrp
172.10.124.2/32 via 172.10.124.2
   Tunnel124 created 00:10:18, expire 01:49:41
   Type: dynamic, Flags: unique registered used
   NBMA address: 10.10.1.2
172.10.124.4/32
   Tunnel124 created 00:00:27, expire 00:02:37
   Type: incomplete, Flags: negative
   Cache hits: 6
R1#
```

Note that R4 is not registered with R1.

Verify the NHRP next-hop server (NHS) configuration on R4.

```
R4#show ip nhrp nhs
Legend: E=Expecting replies, R=Responding, W=Waiting
Tunnel124:
172.10.124.1    E priority = 0 cluster = 0

R4#
```

Note that the NHRP NHS is configured correctly on R4.

Verify the NHS physical (or NMBA) IP address:

```
R4#show ip nhrp detail
172.10.124.1/32 via 172.10.124.1
   Tunnel124 created 00:15:56, never expire
   Type: static, Flags: used
   NBMA address: 10.1.1.1

R4#
```

Note that the NBMA IP address is not correct. It should be 10.10.1.1.

Resolution: Fix the mapping for the NBMA IP address in the NHS configuration on R4.

```
R4(config)#interface tu124
R4(config-if)# no ip nhrp map 172.10.124.1 10.1.1.1
R4(config-if)# ip nhrp map 172.10.124.1 10.10.1.1
```

Clear the NHRP cache on R4 by using the **clear ip nhrp** command and in a few moments verify the NHRP registrations on R1 again:

```
R1#show ip nhrp
172.10.124.2/32 via 172.10.124.2
   Tunnel124 created 00:21:09, expire 01:38:50
   Type: dynamic, Flags: unique registered used
   NBMA address: 10.10.1.2
172.10.124.4/32 via 172.10.124.4
   Tunnel124 created 00:03:18, expire 01:59:43
   Type: dynamic, Flags: unique registered used
   NBMA address: 10.10.1.4
R1#
```

Try to ping 172.10.124.4 from R2 again:

```
R2#ping 172.10.124.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.124.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R2#
```

The ping is now successful.

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

2. Switched Network Troubleshooting

VLAN Distribution Diagram

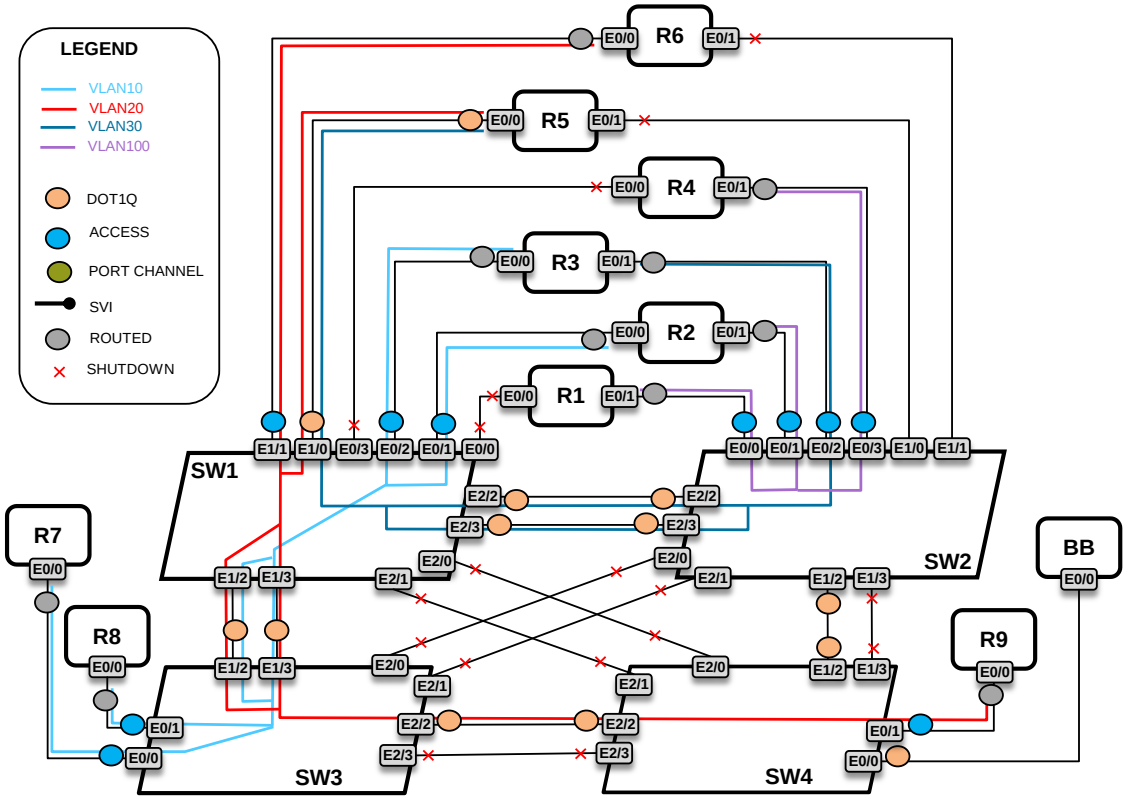

2.1. Symptom: There is a Multiple Spanning Tree (MST) configuration.

Analysis and Testing:

The "Switched Network Troubleshooting" section indicates that only one VLAN should be allowed between SW1 and SW2, but you do not know which VLAN it should be.

An MST configuration requires that VLAN 10 and VLAN 20 be joined together within the same spanning tree, and VLAN 30 must be alone in a single spanning tree.

Therefore the only VLAN that is allowed on the trunk between SW1 and SW2 should be VLAN 30.

Verify the trunk configuration between SW1 and SW2:

```
SW1#show interfaces trunk
```

```
Port                Mode         Encapsulation  Status        Native vlan
Et1/0               on           802.1q         trunking      1
Et1/2               on           802.1q         trunking      1
Et1/3               on           802.1q         trunking      1
Et2/2               on           802.1q         trunking      1
Et2/3               on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et1/0               1-4094
Et1/2               1-4094
Et1/3               1-4094
Et2/2               1-4094
Et2/3               1-4094

Port                Vlans allowed and active in management domain
Et1/0               1,10,20,30
Et1/2               1,10,20,30
Et1/3               1,10,20,30
Et2/2               1,10,20,30
Et2/3               1,10,20,30

Port                Vlans in spanning tree forwarding state and not pruned

Port                Vlans in spanning tree forwarding state and not pruned
Et1/0               1,10,20,30
Et1/2               1,10,20,30
Et1/3               1,10,20,30
Et2/2               1,10,20,30
Et2/3               1,10,20
SW1#
```

Likely Cause: All VLANs are allowed on the trunk link between SW1 and SW2.

We need to add a configuration line on both interfaces E2/2 and E2/3, on both SW1 and SW2, in order to reduce the number of allowed VLANS.

Resolution: Allow only VLAN 30.

```
SW1(config)# interface range e2/2 - 3
SW1(config-if-range)#switchport trunk allowed vlan 30

SW2(config)# interface range e2/2 - 3
SW2(config-if-range)#switchport trunk allowed vlan 30
```

Verify your configuration on SW1 again:

```
SW1#show interface trunk

Port                Mode         Encapsulation  Status        Native vlan
Et1/0               on           802.1q         trunking      1
Et1/2               on           802.1q         trunking      1
Et1/3               on           802.1q         trunking      1
Et2/2               on           802.1q         trunking      1
Et2/3               on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et1/0               1-4094
Et1/2               1-4094
Et1/3               1-4094
Et2/2               30
Et2/3               30

Port                Vlans allowed and active in management domain
```



```
Et1/0               1,10,20,30
Et1/2               1,10,20,30
Et1/3               1,10,20,30
Et2/2               30
Et2/3               30

Port                Vlans in spanning tree forwarding state and not pruned

Port                Vlans in spanning tree forwarding state and not pruned
Et1/0               1,10,20,30
Et1/2               1,10,20,30
Et1/3               1,10,20,30
Et2/2               30
Et2/3               none
SW1#
```

Verify the configuration on SW2:

```
SW2#show interfaces trunk

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

Port                Vlans allowed on trunk
Et1/2               1-4094
Et2/2               30
Et2/3               30

Port                Vlans allowed and active in management domain
Et1/2               1,10,20,30,100
Et2/2               30
Et2/3               30

Port                Vlans in spanning tree forwarding state and not pruned
Et1/2               1,10,20,30,100
Et2/2               30
Et2/3               30
SW2#
```

2.2. Symptom: Traffic for VLAN 10 between SW1 and SW3 crosses port E1/2.

Analysis and Testing:

Verify how the VLAN 10 traffic is forwarded between SW1 and SW3. According to the lab requirements, this traffic should cross the link between the E1/3 interface of SW1 and SW3.

SW1 should be the root bridge for VLAN 10, which is mapping to the MST instance 1. Verify the spanning-tree configuration on SW1:

```
SW1#show spanning-tree mst 1

##### MST1    vlans mapped:   10,20
Bridge        address aabb.cc00.0700  priority      32769 (32768 sysid 1)
Root          this switch for MST1

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Et0/1            Desg FWD 2000000   128.2    Shr
Et0/2            Desg FWD 2000000   128.3    Shr
Et1/0            Desg FWD 2000000   128.33   Shr
Et1/1            Desg FWD 2000000   128.34   Shr
Et1/2            Desg FWD 2000000   128.35   Shr
```

```
Et1/3            Desg FWD 2000000   128.36   Shr

SW1#
```

Note that SW1 is the root bridge for VLAN 10.

Check the configuration on SW3 to see which port has become the root port:

```
SW3#show spanning-tree mst 1

##### MST1    vlans mapped:   10,20
Bridge        address aabb.cc00.0900  priority      32769 (32768 sysid 1)
Root          address aabb.cc00.0700  priority      32769 (32768 sysid 1)
            port    Et1/2           cost      2000000              rem hops 19

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Et0/0            Desg FWD 2000000   128.1    Shr
Et0/1            Desg FWD 2000000   128.2    Shr
Et1/2            Root FWD 2000000   128.35   Shr
Et1/3            Altn BLK 2000000   128.36   Shr
Et2/2            Desg FWD 2000000   128.67   Shr

SW3#
```

Likely Cause: The cost or port-priority statement is missing.

When two parallel links are established between two devices, there are two ways to influence the selection of the root port: we may either modify the cost or modify the port priority.

The port priority should be modified on the designated side, whereas the cost should be modified on the nondesignated side.

According to the section requirements, you are not allowed to use the same method to influence the spanning-tree calculation for VLAN 10 and VLAN 30. Also you know that you cannot configure SW1. Since SW2 is the root bridge for VLAN 30, you may want to leave the port-priority manipulation option for VLAN 30 if you find any issues with the VLAN 30 forwarding between SW1 and SW2.

Resolution: Decrease the cost of E1/3 on SW3.

Modify the cost on the nondesignated side.

Since SW1 is the root switch, the cost has to be modified on SW3.

You can either increase the cost of E1/2, or you can decrease the cost of E1/3.

Both interfaces have a cost of 2,000,000.

Decrease the cost of E1/3:

```
SW3(config)#interface E1/3
SW3(config-if)#spanning-tree mst 1 cost 1999999
```

Now verify that E1/3 has become the root port on SW3:

```
SW3#show spanning-tree mst 1

##### MST1    vlans mapped:   10,20
```



```
Bridge        address aabb.cc00.0900  priority      32769 (32768 sysid 1)
Root          address aabb.cc00.0700  priority      32769 (32768 sysid 1)
            port    Et1/3           cost      1999999              rem hops 19

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Et0/0            Desg BLK 2000000   128.1    Shr
Et0/1            Desg BLK 2000000   128.2    Shr
Et1/2            Altn BLK 2000000   128.35   Shr
Et1/3            Root FWD 1999999   128.36   Shr
Et2/2            Desg BLK 2000000   128.67   Shr

SW3#
```

2.3. Symptom: Traffic for VLAN 30 between SW1 and SW2 crosses port E2/2.

Analysis and Testing:

Again, by default the switch selects the port with the lowest port number to become the root port.

SW2 should be the root bridge for VLAN 30, which is mapping to the MST instance 2. Verify the spanning-tree configuration on SW2:

```
SW2#show spanning-tree mst 2

##### MST2     vlans mapped:   30
Bridge        address aabb.cc00.0800  priority      24578 (24576 sysid 2)
Root          this switch for MST2

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Et0/2            Desg FWD 2000000   128.3    Shr
Et1/2            Desg FWD 2000000   128.35   Shr
Et2/2            Desg FWD 2000000   128.67   Shr
Et2/3            Desg FWD 2000000   128.68   Shr

SW2#
```

Note that SW2 is the root bridge for VLAN 30.

Check the configuration on SW1 to see which port has become the root port:

```
SW1#show spanning-tree mst 2

##### MST2     vlans mapped:   30
Bridge        address aabb.cc00.0700  priority      32770 (32768 sysid 2)
Root          address aabb.cc00.0800  priority      24578 (24576 sysid 2)
            port    Et2/2           cost      2000000              rem hops 19

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Et1/0            Desg FWD 2000000   128.33   Shr
Et1/2            Desg FWD 2000000   128.35   Shr
Et1/3            Desg FWD 2000000   128.36   Shr
Et2/2            Root FWD 2000000   128.67   Shr
Et2/3            Altn BLK 2000000   128.68   Shr

SW1#
```

We can deduct that traffic for VLAN 30 crosses interface E2/2, instead of E2/3 as required by the lab.

Likely Cause: The** cost **or** port-priority **statement is missing.

The only method you can use on the root bridge is changing the port priority.

You may either increase the port priority of E2/2, or you can decrease the port priority of E2/3.

By default the port priority is equal to 128.

Resolution: Decrease the port priority of E2/3 on SW2.

The port priority may be modified in increments of 64:

```
SW2(config)#interface e2/3
SW2(config-if)#spanning-tree mst 2 port-priority 64
```

Verify that E2/3 has become the root port on SW1:

```
SW1#show spanning-tree mst 2

##### MST2     vlans mapped:   30
Bridge        address aabb.cc00.0700  priority      32770 (32768 sysid 2)
Root          address aabb.cc00.0800  priority      24578 (24576 sysid 2)
            port    Et2/3           cost      2000000              rem hops 19

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Et1/0            Desg FWD 2000000   128.33   Shr
Et1/2            Desg FWD 2000000   128.35   Shr
Et1/3            Desg FWD 2000000   128.36   Shr
Et2/2            Altn BLK 2000000   128.67   Shr
Et2/3            Root FWD 2000000   128.68   Shr

SW1#
```

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

3. IPv4 OSPF Troubleshooting

3.1. Symptom: R1, R2, and R3 do not form the OSPF neighbor relationships over the DMVPN.

Analysis and Testing:

Verify the neighbor relationship on R1 and R2:

```
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.10.23.10      0   2WAY/DROTHER    00:00:33    172.10.23.10    FastEthernet0/0
172.10.23.30      0   2WAY/DROTHER    00:00:34    172.10.23.30    FastEthernet0/0
172.10.103.1      1   FULL/DR         00:00:37    172.10.23.3     FastEthernet0/0
172.10.101.1      0   DOWN/DROTHER      -         172.10.124.1    Serial0/0/0.124

R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.10.102.1      1   2WAY/DROTHER    00:01:59    172.10.124.2    Tunnel124
172.10.104.1      1   FULL/DR         00:01:35    172.10.124.4    Tunnel124
```



```
R1#
```

Note that R1 sees R4 as a designated router (DR).

Likely Cause: The hub is not the DR.

We can see that R4 is the DR on the subnet 172.10.124.0/24. Therefore OSPF updates from R4 cannot reach R2.

In a hub-and-spoke DMVPN topology, the hub R1 should become the DR.

Resolution: Set OSPF priorities to 0 on the spokes.

```
R2(config)#interface tu124
R2(config-subif)#ip ospf priority 0

R4(config)#interface tu124
R4(config-subif)#ip ospf priority 0
```

Reload the OSPF process on R1:

```
R1#clear ip ospf process
Reset ALL OSPF processes? [no]: yes
```

Now verify that R1 has become the DR:

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
N/A               0   ATTEMPT/DROTHER 00:00:52    172.10.124.2    Tunnel124
N/A               0   ATTEMPT/DROTHER 00:00:52    172.10.124.4    Tunnel124
R1#
```

It looks like R1 is still not able to become the DR.

Verify the interface configuration of R1:

```
R1#show running-config int tu124
Building configuration...

Current configuration : 227 bytes
!
interface Tunnel124
 ip address 172.10.124.1 255.255.255.0
 no ip redirects
 ip nhrp network-id 10
 ip ospf network non-broadcast
 ip ospf priority 0
 tunnel source Ethernet0/1
 tunnel mode gre multipoint
 tunnel key 10
end

R1#
```

You need to remove the priority of 0 to allow R1 to become the DR:

```
R1(config)#interface tu124
R1(config-subif)#no ip ospf priority 0
```

R1 is now the DR on the subnet 172.10.124.0/24:

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.10.102.1      0   FULL/DROTHER    00:01:51    172.10.124.2    Tunnel124
172.10.104.1      0   FULL/DROTHER    00:01:54    172.10.124.4    Tunnel124
R1#
```

You have not fully corrected the reachability issues to the subnet 172.10.124.0/24 yet.

Let us investigate further.

3.2. Symptom: R5 is also missing the OSPF route to 172.10.124.0.

Analysis and Testing:

We corrected the situation among the devices in area 124, making sure that R1 becomes the DR.

But R5, R7, and R8 are still missing the subnet 172.10.124.0/24 in their OSPF routing tables.

The subnet 172.10.124.0/24 belongs to area 124, which is not connected to area 0.

Likely Cause: There is a misconfigured virtual link.

You should have a virtual link between R2 and R3.

```
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.10.23.10      0   2WAY/DROTHER    00:00:36    172.10.23.10    Ethernet0/0
172.10.23.30      0   2WAY/DROTHER    00:00:37    172.10.23.30    Ethernet0/0
172.10.103.1      1   FULL/DR         00:00:38    172.10.23.3     Ethernet0/0
172.10.101.1      1   FULL/DR         00:01:57    172.10.124.1    Tunnel124
R2#
```

There is no interface virtual link on R2.

Verify the virtual configuration on R2:

```
R2#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 172.10.103.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 23, via interface Ethernet0/0
 Topology-MTID    Cost    Disabled     Shutdown      Topology Name
        0           10        no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:09
  Message digest authentication enabled
    Youngest key id is 1
R2#
```

Note that the virtual link is configured but it is not adjacent with R3.

Run the **debug ip ospf adj** command to examine the adjacency formation process:

```
R2#
```



```
R2#deb ip ospf adj
OSPF adjacency debugging is on
R2#
*Apr 24 01:21:13.019: OSPF-1 ADJ   NV0: Send with youngest Key 0
*Apr 24 01:21:13.907: OSPF-1 ADJ   VL0: Send with youngest Key 1
R2#
*Apr 24 01:21:14.225: OSPF-1 ADJ     VL0: Rcv pkt from 172.10.23.3 : Mismatched
Authentication Key - Message Digest Key 1
R2#
```

Carefully examine the virtual configuration on R2 and R3:

```
R2#show running-config | section router ospf
router ospf 1
 area 23 authentication message-digest
 area 23 virtual-link 172.10.103.1 authentication message-digest
 area 23 virtual-link 172.10.103.1 message-digest-key 1 md5 CISCO
 area 124 filter-list prefix No23 out
 redistribute connected subnets route-map C2O
 network 172.10.23.2 0.0.0.0 area 23
 network 172.10.124.2 0.0.0.0 area 124
 distance ospf intra-area 121 inter-area 121 external 171
R2#


R3#show running-config | section router ospf
router ospf 1
 area 23 authentication message-digest
 area 23 virtual-link 172.10.102.1 authentication message-digest
 area 23 virtual-link 172.10.102.1 message-digest-key 1 md5 CISC0
 redistribute rip subnets route-map R2O
 network 172.10.23.3 0.0.0.0 area 23
 network 172.10.35.3 0.0.0.0 area 0
 distance ospf intra-area 121 inter-area 121 external 171
ipv6 router ospf 2
ipv6 router ospf 1
 router-id 233.233.233.233
R3#
```

Note that the MD5 string CISCO is misconfigured on R3.

Resolution: Fix the virtual link configuration on R3.

```
R3(config)#router ospf 1
R3(config-router)#no area 23 virtual-link 172.10.102.1 message-digest-key 1 md5
CISC0
R3(config-router)#area 23 virtual-link 172.10.102.1 message-digest-key 1 md5 CISCO
```

You now have the interface virtual link on R2:

```
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.10.103.1      0   FULL/  -        00:00:10    172.10.23.3     OSPF_VL0
172.10.23.10      0   2WAY/DROTHER    00:00:33    172.10.23.10    Ethernet0/0
172.10.23.30      0   2WAY/DROTHER    00:00:39    172.10.23.30    Ethernet0/0
172.10.103.1      1   FULL/DR         00:00:35    172.10.23.3     Ethernet0/0
172.10.101.1      1   FULL/DR         00:01:35    172.10.124.1    Tunnel124
R2#
```

3.3. Symptom: R5, R7, and R8 are still missing an OSPF route to 172.10.124.0/24.

Analysis and Testing:

The 172.10.124.0/24 subnet is missing in the routing tables of R5, R7, and R8:

```
R5#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      172.10.0.0/16 is variably subnetted, 12 subnets, 2 masks
O IA     172.10.23.0/24 [121/20] via 172.10.35.3, 00:43:09, Ethernet0/0.30
O E2     172.10.43.0/24 [171/20] via 172.10.35.3, 00:43:09, Ethernet0/0.30
O E2     172.10.100.0/24 [171/20] via 172.10.35.3, 00:29:25, Ethernet0/0.30
O E2     172.10.102.0/24 [171/20] via 172.10.35.3, 00:29:25, Ethernet0/0.30
O E2     172.10.103.0/24 [171/20] via 172.10.35.3, 00:43:09, Ethernet0/0.30
O E2     172.10.104.0/24 [171/20] via 172.10.35.3, 00:43:09, Ethernet0/0.30
R5#
```

R2 should advertise this subnet.

After you correct the OSPF priorities on the subnet 172.10.124.0/24 and add the missing virtual link, you expect to see the missing subnet in the OSPF routing tables of R5, R7, and R8.

Why is this route not advertised from area 124 to area 23?

First, analyze the OSPF configuration on R2:

```
R2#show running-config | section ospf
 ip ospf message-digest-key 1 md5 CISCO
 ip ospf priority 0
 ipv6 ospf network point-to-multipoint non-broadcast
 ipv6 ospf 2 area 0
 ip ospf priority 0
router ospf 1
 log-adjacency-changes
 area 23 authentication message-digest
 area 23 virtual-link 172.10.103.1 authentication message-digest
 area 23 virtual-link 172.10.103.1 message-digest-key 1 md5 CISCO
 area 124 filter-list prefix No23 out
 redistribute connected subnets route-map C2O
 network 172.10.23.2 0.0.0.0 area 23
 network 172.10.124.2 0.0.0.0 area 124
 distance ospf intra-area 121 inter-area 121 external 171


R2#show ip prefix-list
ip prefix-list No23: 2 entries
   seq 5 deny 172.16.23.0/24
   seq 10 permit 0.0.0.0/24
```

Likely Cause: There is an incorrect filter list on R2.

The incorrect filter list that is applied on R2 allows only subnets 0.0.0.0/24 to be advertised out of area 124.

The "IPv4 OSPF Troubleshooting" section requires that the subnet 172.10.23.0/24 not be advertised into area 124.

Therefore the filter list should be applied as inbound, not outbound.

Moreover, to allow all subnets, the second line should read 0.0.0.0/0 le 32 instead of 0.0.0.0/24.

And the first line should read 172.10.23.0 instead of 172.16.23.0.

Resolution: Apply the filter as inbound instead of outbound.

```
R2(config-router)#router ospf 1
R2(config-router)#area 124 filter-list prefix No23 in
R2(config-router)#no area 124 filter-list prefix No23 out
```

You need to correct the first and second lines of the prefix list:

```
R2(config-router)#no ip prefix-list No23
R2(config)#ip prefix-list No23 seq 5 deny 172.10.23.0/24
R2(config)#ip prefix-list No23 seq 10 permit 0.0.0.0/0 le 32
```

Verify that R5 and R7 now have the OSPF route to the subnet 172.10.124.0/24:

```
R5#show ip route ospf | inc ^O
O IA     172.10.23.0/24 [121/20] via 172.10.35.3, 00:51:42, Ethernet0/0.30
O E2     172.10.43.0/24 [171/20] via 172.10.35.3, 00:51:42, Ethernet0/0.30
O E2     172.10.100.0/24 [171/20] via 172.10.35.3, 00:37:58, Ethernet0/0.30
O IA     172.10.101.1/32 [121/1021] via 172.10.35.3, 00:01:20, Ethernet0/0.30
O E2     172.10.102.0/24 [171/20] via 172.10.35.3, 00:37:58, Ethernet0/0.30
O E2     172.10.103.0/24 [171/20] via 172.10.35.3, 00:51:42, Ethernet0/0.30
O E2     172.10.104.0/24 [171/20] via 172.10.35.3, 00:51:42, Ethernet0/0.30
O IA     172.10.124.0/24 [121/1020] via 172.10.35.3, 00:01:20, Ethernet0/0.30
R5#

R7#show ip route ospf | inc ^O
O IA     172.10.35.0/24 [121/20] via 172.10.23.3, 00:52:37, Ethernet0/0
O E2     172.10.43.0/24 [171/20] via 172.10.23.3, 00:52:37, Ethernet0/0
O E2     172.10.65.0/24 [171/20] via 172.10.23.3, 00:52:37, Ethernet0/0
O E2     172.10.100.0/24 [171/20] via 172.10.23.2, 00:38:56, Ethernet0/0
O IA     172.10.101.1/32 [121/1011] via 172.10.23.2, 00:02:13, Ethernet0/0
O E2     172.10.102.0/24 [171/20] via 172.10.23.2, 00:38:56, Ethernet0/0
O E2     172.10.103.0/24 [171/20] via 172.10.23.3, 00:52:37, Ethernet0/0
O E2     172.10.104.0/24 [171/20] via 172.10.23.3, 00:52:37, Ethernet0/0
O IA     172.10.105.0/24 [121/21] via 172.10.23.3, 00:52:37, Ethernet0/0
O IA     172.10.124.0/24 [121/1010] via 172.10.23.2, 00:02:13, Ethernet0/0
R7#
R8#show ip route | inc 124
O IA     172.10.124.0/24 [121/1010] via 172.10.23.2, 00:03:06, Ethernet0/0
R8#
```

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

4. IPv4 EIGRP Troubleshooting

4.1. Symptom: R6 does not have any EIGRP routes in its routing table.

Analysis and Testing:

Verify the neighbors on R6:

```
R6#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
R6#
```

There are no neighbors.

Verify the EIGRP interfaces on R6:

```
R6#show ip eigrp interfaces detail
EIGRP-IPv4 Interfaces for AS(100)
                              Xmit  Queue     PeerQ              Mean     Pacing  Time
Multicast    Pending
Interface               Peers  Un/Reliable  Un/Reliable  SRTT    Un/Reliable    Flow
Timer   Routes
Et0/0                    0        0/0       0/0           0       0/2            50
0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 5/1
  Hello's sent/expedited: 6809/4
  Un/reliable mcasts: 0/4  Un/reliable ucasts: 6/5
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 2  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is not set
R6#
```

Likely Cause: Authentication is not correctly configured on R6.

Verify the EIGRP configuration on the E0/0 interface on R6:

```
R6#show running-config interface e0/0
Building configuration...

Current configuration : 151 bytes
!
interface Ethernet0/0
 ip address 172.10.65.6 255.255.255.0
 ip authentication mode eigrp 10 md5
 ip authentication key-chain eigrp 10 KeyEigrp
end

R6#
```

As you can see, AS 10 has been used instead of AS 100 for authentication command lines.

Resolution: Modify the autonomous system for authentication command lines.

```
R6(config-if)#interface e0/0
R6(config-if)# no ip authentication mode eigrp 10 md5
R6(config-if)# no ip authentication key-chain eigrp 10 KeyEigrp
R6(config-if)# ip authentication mode eigrp 100 md5
R6(config-if)# ip authentication key-chain eigrp 100 KeyEigrp
```


Verify the EIGRP neighbors on R6:

```
R6#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.10.65.5             Et0/0                    13 00:00:35    9   100  0  11
0   172.10.65.20            Et0/0                    10 00:00:35    9   100  0  11
R6#
```

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

5. IPv4 RIP Troubleshooting

5.1. Symptom: The debug indicates the use of the multicast address 224.0.0.9.

Analysis and Testing:

The "IPv4 RIP Troubleshooting" section requires that all RIP updates be sent to the IP destination address 255.255.255.255, whereas the following debug output indicates that this requirement is not met:

```
R4#debug ip rip
RIP protocol debugging is on

[output omitted for brevity]
 RIP: sending v2 update to 224.0.0.9 via Serial1/0 (172.10.43.4)
```

R3 should send only RIPv1 updates. Therefore, there is no possibility for R3 to break the requirement of using only a broadcast address in RIP updates.

However, R4 is sending RIPv2 updates.

Likely Cause: RIPv2 has been configured.

The "IPv4 RIP Troubleshooting" section requires that R4 send RIPv2 updates.

How can you configure RIPv2 and force routers to broadcast updates?

Cisco IOS Software offers a command that forces this type of behavior among RIPv2 routers.

Resolution: Add the** ip rip v2-broadcast **interface configuration command line.

```
R4(config)#interface serial1/0
R4(config-if)#ip rip v2-broadcast
```

Verify the IP destination address of the RIP updates now:

```
R4#debug ip rip
RIP protocol debugging is on

[output omitted for brevity]
RIP: sending v2 update to 255.255.255.255 via Serial1/0 (172.10.43.4)
```

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

6. IPv4 Redistribution Troubleshooting

6.1. Symptom: R6 and SW2 cannot reach the subnet 172.10.35.0/24.

Analysis and Testing:

The subnet 172.10.35.0/24 is not present in the IP routing tables of R6 and R9. Therefore these devices cannot ping 172.10.35.3 and 172.10.35.5:

```
R9#ping 172.10.35.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R9#ping 172.10.35.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R9#


R6#ping 172.10.35.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R6#ping 172.10.35.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R6#
```

Verify the routing tables to make sure the problem is in sending the echo packets, not in receiving the echo-reply packets:

```
R6#show ip route 172.10.35.0
% Subnet not in table


R9#show ip route 172.10.35.0
% Subnet not in table
```

These subnets should have been redistributed on R5 into EIGRP AS 100.

Likely Cause: There is an incorrect redistribution of connected networks on R5.

Remember that the EIGRP section requires that the subnet 172.10.105.0/24 be redistributed on R5 into EIGRP.

Therefore a route map had been configured to allow this subnet only:

```
R5#show run | section eigrp
```



```
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 KeyEigrp
router eigrp 100
 redistribute connected route-map E2C
 redistribute ospf 1 metric 1 1 1 1 1
 network 172.10.65.5 0.0.0.0
 distance eigrp 123 169
 auto-summary
 redistribute eigrp 100 subnets


R5#show route-map
route-map E2C, permit, sequence 10
  Match clauses:
    interface Loopback105
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
R5#
```

This route map is preventing all other connected networks from being redistributed into EIGRP.

Resolution: Match the interface E0/0.30 in the route map E2C.

```
R5(config)#route-map E2C
R5(config-route-map)#match interface e0/0.30
```

Verify now that the subnet 172.10.35.0/24 is now present in the routing tables of R6 and R9:

```
R6#show ip route 172.10.35.0
Routing entry for 172.10.35.0/24
  Known via "eigrp 100", distance 169, metric 307200, type external
  Redistributing via eigrp 100
  Last update from 172.10.65.5 on Ethernet0/0, 00:00:32 ago
  Routing Descriptor Blocks:
  * 172.10.65.5, from 172.10.65.5, 00:00:32 ago, via Ethernet0/0
      Route metric is 307200, traffic share count is 1
      Total delay is 2000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
R6#

R9#show ip route 172.10.35.0
Routing entry for 172.10.35.0/24
  Known via "eigrp 100", distance 169, metric 307200, type external
  Redistributing via eigrp 100
  Last update from 172.10.65.5 on Ethernet0/0, 00:00:50 ago
  Routing Descriptor Blocks:
  * 172.10.65.5, from 172.10.65.5, 00:00:50 ago, via Ethernet0/0
      Route metric is 307200, traffic share count is 1
      Total delay is 2000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
R9#
```

Also verify pings on R6 and R9:

```
R9#ping 172.10.35.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R9#ping 172.10.35.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
R9#

R6#ping 172.10.35.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R6#ping 172.10.35.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.35.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R6#
```

6.2. Symptom: R1 cannot ping the subnet 172.10.23.0/24.

Analysis and Testing:

The subnet 172.10.23.0/24 is not present in the routing table of R1:

```
R1#show ip route | include 23
R1#
```

At first, you could have expected OSPF alone to add a route for the subnet 172.10.23.0/24 for R1.

But the OSPF section required you to prevent this route from entering area 124.

Therefore, the only way to inform R1 is by redistributing this route from RIP into OSPF on R4.

Likely Cause: There is an incorrect redistribution from RIP into OSPF on R4.

See how RIP has been redistributed into OSPF on R4:

```
R4#sh run | section ospf
 ipv6 ospf priority 0
 ipv6 ospf 1 area 0
 ip ospf priority 0
router ospf 1
 log-adjacency-changes
 redistribute rip subnets route-map R2O
 network 172.10.124.4 0.0.0.0 area 124
 distance ospf intra-area 121 inter-area 121 external 171
 redistribute ospf 1 metric 10 route-map O2R
ipv6 router ospf 1
 router-id 244.244.244.244
 log-adjacency-changes
 redistribute rip FRAME include-connected
 redistribute ospf 1 include-connected


R4#show route-map R2O
route-map R2O, permit, sequence 10
  Match clauses:
    ip address (access-lists): RIP
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes


R4#sh ip access-lists RIP
Standard IP access list RIP
    10 permit 172.10.43.0, wildcard bits 0.0.0.255 (1 match)
    20 permit 172.10.104.0, wildcard bits 0.0.0.255 (1 match)
```



```
        30 permit 172.10.103.0, wildcard bits 0.0.0.255 (1 match)
```

Resolution: On R4, add a line for the access list RIP that allows the subnet 172.10.23.0.

```
R4(config)#ip access-list standard RIP
R4(config-std-nacl)#permit 172.10.23.0 0.0.0.255
```

Now verify the routing table on R1:

```
R1#show ip route | include 23
O E2     172.10.23.0/24 [171/20] via 172.10.124.4, 00:00:19, Tunnel124
R1#

R1#ping 172.10.23.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.23.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/100/100 ms


R1#ping 172.10.23.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.10.23.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/75/76 ms
```

Note that R1 reaches the subnet 172.10.23.0 via R4 and then via R3:

```
R1#traceroute 172.10.23.2
Type escape sequence to abort.
Tracing the route to 172.10.23.2
VRF info: (vrf in name/id, vrf out name/id)
  1 172.10.124.4 0 msec 1 msec 0 msec
  2 172.10.43.3 9 msec 9 msec 6 msec
  3 172.10.23.2 5 msec *  2 msec
R1#
```

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

```
tclsh
foreach addr {
172.10.124.1
172.10.101.1
172.10.124.2
172.10.23.2
172.10.102.1
172.10.100.1
172.10.23.3
172.10.35.3
172.10.43.3
172.10.103.1
172.10.124.4
172.10.43.4
172.10.104.1
172.10.35.5
172.10.105.1
172.10.65.5
172.10.106.1
172.10.65.6
```

```
172.10.23.10
172.10.120.1
172.10.65.20
172.10.23.30
} {ping $addr}
tclquit
```

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

7. Security Troubleshooting

7.1. Symptom: Communication between the hosts on VLAN 10 and the DHCP server needs to be verified.

Analysis and Testing:

The "Security Troubleshooting" section requires that a PC that is located in VLAN 10 should communicate with a DHCP server through R3. The IP address of the DCHP server is 172.10.35.254.

Verify the configuration of interface e0/0:

```
R3#show running-config interface e0/0
Building configuration...

Current configuration : 294 bytes
!
interface FastEthernet0/0
 ip address 172.10.23.3 255.255.255.0
 ip access-group NoSpoof in
  ip ospf message-digest-key 1 md5 CISCO
 ipv6 address FEC0:23::3/125
 ipv6 ospf 2 area 0
 ipv6 ospf neighbor FE80::2
 ipv6 ospf network point-to-multipoint non-broadcast
endend



R3#sh ip access-lists NoSpoof
Extended IP access list NoSpoof
    10 permit ip 172.10.0.0 0.0.255.255 any (14298 matches)
    20 deny ip any any log
```

Likely Cause: A permit statement is missing from the access list NoSpoof.

To allow communication between a host and a DHCP server, you need to permit hosts with the IP address 0.0.0.0 to communicate with the DHCP server at 172.10.35.254.

Resolution: Add a permit statement on R3.

```
R3(config)#no ip access-list extended NoSpoof
R3(config)#ip access-list extended NoSpoof
R3(config-ext-nacl)# permit ip 172.10.0.0 0.0.255.255 any
R3(config-ext-nacl)# permit udp host 0.0.0.0 host 172.10.35.254
```



```
R3(config-ext-nacl)# deny   ip any any log


R3#show ip access-lists NoSpoof
Extended IP access list NoSpoof
    10 permit ip 172.10.0.0 0.0.255.255 any (14381 matches)
    20 permit udp host 0.0.0.0 host 172.10.35.254
    30 deny ip any any log
```

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

8. IPv6 Troubleshooting

R2 cannot ping the IPv6 subnets that are advertised via OSPFv3 from R3:

```
R2#ping FEC0:103::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:103::1, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
R2#ping FEC0:133::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:133::1, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
R2#
```

8.1. Symptom: An OSPF neighbor is missing between R2 and R3.

Analysis and Testing:

R2 should have one OSPF neighbor, but it does not have any:

```
R2#show ipv6 ospf neighbor
R2#
```

Remember that the "IPv6 Troubleshooting" section requires that no DR be elected on VLAN 10 and that no OSPFv3 multicast packets should be seen on VLAN 10.

Therefore VLAN 10 should become the OSPF point-to-multipoint nonbroadcast network type.

First, verify this point:

```
R2#show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Et0/0        2     0               3          10    P2MP  0/0
R2#

R3#show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Se1/0        1     0               7          64    P2P   1/1
Lo103        2     0               14         1     P2P   0/0
```

```
Lo133          2     0               15          1      P2P    0/0
Et0/0          2     0               3           10     P2MP   0/0
R3#
```

For the OSPFv3 point-to-multipoint nonbroadcast network type, a **neighbor** command is required on either R2 or R3.

Likely Cause: The** neighbor **command is missing.

The **neighbor** command is not present on R2:

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

Current configuration : 262 bytes
!
interface Ethernet0/0
 ip address 172.10.23.2 255.255.255.0
 ip ospf message-digest-key 1 md5 CISCO
 ip ospf priority 0
 ipv6 address FE80::2 link-local
 ipv6 address FEC0:23::2/125
 ipv6 ospf 2 area 0
 ipv6 ospf network point-to-multipoint non-broadcast
end

R2
```

The command is also not on R3:

```
R3#show running-config interface e0/0
Building configuration...

Current configuration : 270 bytes
!
interface Ethernet0/0
 ip address 172.10.23.3 255.255.255.0
 ip access-group NoSpoof in
 ip ospf message-digest-key 1 md5 CISCO
 ipv6 address FE80::3 link-local
 ipv6 address FEC0:23::3/125
 ipv6 ospf 2 area 0
 ipv6 ospf network point-to-multipoint non-broadcast
end

R3#
```

Note that the IPv6 link-local addresses are configured on R2 and R3.

Resolution: Add a** neighbor **command.

Configure the OSPFv3 **neighbor** command on R2 or R3 or on both routers. The IPv6 link-local address must be used for the OSPFv3 **neighbor** command:

```
R2(config)#interface e0/0
R2(config-if)#ipv6 ospf neighbor FE80::3
```

R2 and R3 are now neighbors:


```
R2#show ipv6 ospf neighbor

            OSPFv3 Router with ID (172.10.102.1) (Process ID 2)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
172.10.103.1      0   FULL/  -        00:01:45    3               Ethernet0/0
R2#
```

Try to ping the R3 IPv6 subnets from R2 again:

```
R2#ping FEC0:103::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:103::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R2#ping FEC0:133::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:133::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R2#
```

8.2. Symptom: R2 cannot ping the subnet FEC0:43::/125.

Analysis and Testing:

Ping FEC0:43::4 from R2:

```
R2#ping FEC0:43::4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:43::4, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
R2#
```

The subnet FEC0:43::/125 is not present in the IPv6 routing table of R2:

```
R2#show ipv6 route | inc 43
R2#
```

This subnet is present on R1:

```
R1#show ipv6 route FEC0:43::/125
Routing entry for FEC0:43::/125
  Known via "eigrp 100", distance 170, metric 1757696, type external
  Route count is 1/1, share count 0
  Routing paths:
    FE80::4, Ethernet0/1
      Last updated 00:34:25 ago

R1#
```

We can see that R1 learned this route from R4 via IPv6 EIGRP.

Why does R1 not advertise this subnet to R2?

Likely Cause: The IPv6 EIGRP split horizon prevents R1 from advertising the subnet FEC0:43::/125 to R2.

Since R1 learns this route through its interface E0/1, it refuses to announce it out to the same interface, as a loop prevention mechanism. Note that R2 and R4 do not form the direct IPv6 EIGRP neighbor relationship, and R2 and R4 are supposed to learn the IPv6 EIGRP prefixes via R1.

You need to disable the IPv6 EIGRP split horizon on R1.

Resolution: Disable the IPv6 EIGRP split horizon on R1.

```
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#int e0/1
R1(config-if)#no ipv6 split-horizon eigrp 100
R1(config-if)#en
*Apr   24   12:46:44.634:   %DUAL-5-NBRCHANGE:   EIGRP-IPv6   100:   Neighbor   FE80::4
(Ethernet0/1) is resync: split horizon changed
*Apr   24   12:46:44.635:   %DUAL-5-NBRCHANGE:   EIGRP-IPv6   100:   Neighbor   FE80::2
(Ethernet0/1) is resync: split horizon changed
R1(config-if)#end
R1#
```

Allow enough time to let R1 advertise its route to R2 and then check the routing table on R2:

```
R2#show ipv6 route eigrp
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
EX  FEC0:43::/125 [170/1783296]
     via FE80::1, Ethernet0/1
D   FEC0:124::/125 [90/307200]
     via FE80::1, Ethernet0/1
R2#
```

You can test the universal IPv6 reachability by using this Tcl script:

```
tclsh
foreach address {
FEC0:103::1
FEC0:133::1
FEC0:124::1
FEC0:124::2
FEC0:23::2
FEC0:23::3
FEC0:43::3
FEC0:124::4
FEC0:43::4
} {ping $address}

R2#tclsh
R2(tcl)#foreach address {
+>(tcl)#FEC0:103::1
+>(tcl)#FEC0:133::1
+>(tcl)#FEC0:124::1
+>(tcl)#FEC0:124::2
+>(tcl)#FEC0:23::2
+>(tcl)#FEC0:23::3
+>(tcl)#FEC0:43::3
+>(tcl)#FEC0:124::4
+>(tcl)#FEC0:43::4
```



```
+>(tcl)#} {ping $address}
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:103::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:133::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:124::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:124::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:23::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:23::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:43::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/12/26 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:124::4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/17 ms
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:43::4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R2(tcl)#exit
R2#
```

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

9. Quality of Service Troubleshooting

9.1. Symptom: The class maps need to be verified.

Analysis and Testing:

Verify the class map on R2:

```
R2#show class-map
 Class Map match-any class-default (id 0)
   Match any

 Class Map match-all CLASS_VOICE (id 1)
```

```
    Match access-group name  VOICE

 Class Map match-all CLASS4 (id 2)
    Match access-group name  QOS4

 Class Map match-all CLASS2 (id 3)
    Match access-group name  QOS2

 Class Map match-all CLASS3 (id 4)
    Match access-group name  QOS3

 Class Map match-all CLASS1 (id 5)
    Match access-group name  QOS1

R2#
```

Note that the classes are configured to match the type of traffic that is defined in the scenario.

Verify the associated access lists on R2:

```
R2#show access-lists
Extended IP access list QOS1
    10 permit tcp any any eq www
Extended IP access list QOS2
    10 permit tcp any any eq smtp
Extended IP access list QOS3
    10 permit tcp any any eq domain
Extended IP access list QOS4
    10 permit tcp any any eq telnet
Extended IP access list VOICE
    10 permit tcp any any eq 1720
    20 permit udp any any range 16384 32767
R2#
```

Likely Cause: The access list does not match the DNS packets.

Remember that DNS uses both UDP and TCP protocols.

Resolution: Add a line to match the UDP packets.

```
R2(config)#ip access-list extended QOS3
R2(config-ext-nacl)# permit udp any any eq domain


R2#show ip access-lists  QOS3
Extended IP access list QOS3
    10 permit tcp any any eq domain
    20 permit udp any any eq domain
```

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

10. IP Service Troubleshooting

10.1. Symptom: The logging size on R4 needs to be verified.

Analysis and Testing:

The "IP Service Troubleshooting" section requires setting the maximum number of entries retained in the configuration log as two times larger as the default value.

Verify the configuration on R4:

```
R4#show run | section archive
archive
 log config
  logging enable
  logging size 1000
  hidekeys
```

Likely Cause: The logging size is incorrect.

The default value is 100.

Resolution: Set the logging size to 200.

```
R4(config)#archive
R4(config-archive)# log config
R4(config-archive-log-cfg)#  logging size 200
```

Verify the configuration on R4 again:

```
R4#show run | section archive
archive
 log config
  logging enable
  logging size 200
  hidekeys
R4#
```

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