

# 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 9

### Troubleshooting Section

# Answer Key

---

---

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

---

# Table of Contents

## **Cisco 360 CCIE R&S Exercise Workbook Lab 9 Troubleshooting Section Answer Key..... 2**

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

## **Exercise Workbook Lab 9 Troubleshooting Section Answer Key ..... 5**

|   |    |
|---|----|
| Grading and Duration.....   | 5  |
| Difficulty Level.....   | 5  |
| Restrictions and Goals.....   | 5  |
| Explanation of Each of the Restrictions and Goals .....               | 7  |
| 1. Switched Network Troubleshooting Section.....                      | 8  |
| 1.1. Symptom: R2 is not able to ping R6 on subnet 148.49.26.0/24..... | 8  |
| 2. IPv4 OSPF Troubleshooting Section .....                            | 23 |
| 3. IPv4 EIGRP Troubleshooting Section.....                            | 29 |
| 4. IPv4 RIP Troubleshooting Section.....                              | 32 |
| 5. IPv4 Redistribution Troubleshooting Section .....                  | 36 |
| 6. Security Troubleshooting Section .....                             | 42 |
| 7. IPv6 Troubleshooting Section .....                                 | 45 |
| 8. QoS Troubleshooting Section.....                                   | 46 |
| 9. IP Services Troubleshooting Section .....                          | 49 |

# 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 9

## Troubleshooting Section

### Answer Key

---

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

---

## Grading and Duration

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

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

---

## Difficulty Level

- Difficulty: Basic to Intermediate

## Restrictions and Goals

---

**Note** Read this section carefully.

---

- To receive credit for a subsection, you must fully complete the subsection per requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the scenario diagram belong to network 148.49.0.0/16.
- *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.
- Unless explicitly specified otherwise, addresses and networks that are advertised in the “Border Gateway Protocol” (BGP) section need to be reachable by all BGP routers but do not have to be reachable by interior gateway protocol (IGP)-only routers.

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

# Explanation of Each of the Restrictions and Goals

**IPv4 subnets that are displayed in the scenario “IPv4 IGP” diagram belong to network 148.49.0.0/16.**

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

**Do not use any static routes.**

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

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

The original mask is the mask configured on the loopback interface. Open Shortest Path First (OSPF) by default treats loopback interfaces as host routes and advertises them as /32 prefixes. The requirement to advertise loopback interfaces with their original masks precludes using the default OSPF network type for the loopback interfaces. You need to provide a solution such as changing the OSPF network type or summarizations. 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) except on SW1.**

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

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

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

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

**Use conventional routing algorithms.**

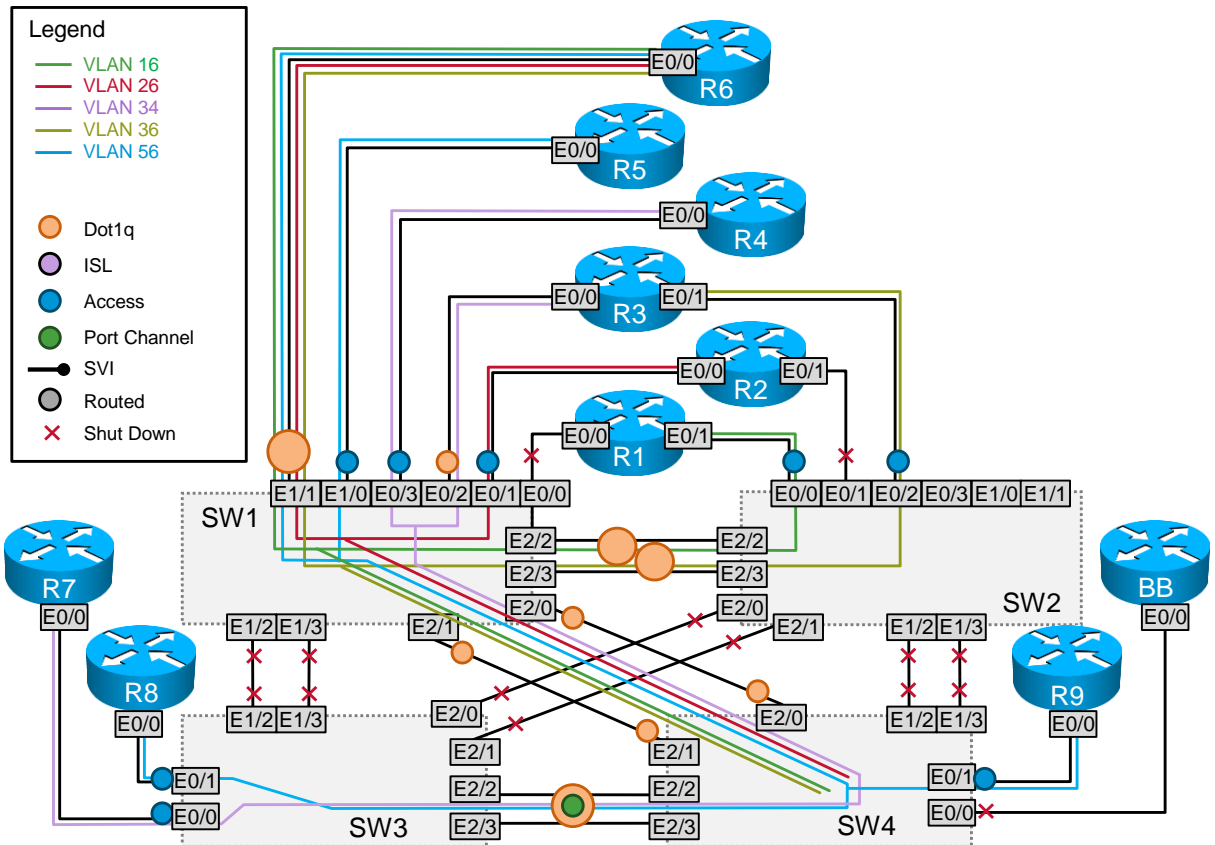
This restriction prevents you from solving any problems by configuring policy routing. At the heart of this restriction is the interpretation of “conventional routing algorithms.” Although this phrase can be interpreted in different ways, this interpretation is applied in this workbook:

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

# 1. Switched Network Troubleshooting Section

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

## VLAN Propagation



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

### 1.1. Symptom: R2 is not able to ping R6 on subnet 148.49.26.0/24.

#### Analysis and Testing:

Interfaces are up, but there is no connectivity between R2 and R6 on subnet 148.49.26.0/24.

```
R2#ping 148.49.26.6
```

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

```
Sending 5, 100-byte ICMP Echos to 148.49.26.6, timeout is 2 seconds:
```

```
.....
```

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

```
R2#show ip interface brief
```

| Interface      | IP-Address  | OK? | Method | Status |
|----------------|-------------|-----|--------|--------|
| Ethernet0/0    | unassigned  | YES | NVRAM  | up     |
| Ethernet0/0.26 | 148.49.26.2 | YES | NVRAM  | up     |

```

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            148.49.21.2    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
Loopback102         148.49.102.1   YES NVRAM   up          up
R2#
R6#show ip interface brief
Interface                IP-Address      OK?  Method  Status
Protocol
Ethernet0/0              unassigned      YES NVRAM   up          up
Ethernet0/0.16           148.49.16.6    YES NVRAM   up          up
Ethernet0/0.26           unassigned      YES DHCP    up          up
Ethernet0/0.36           148.49.36.6    YES NVRAM   up          up
Ethernet0/0.56           148.49.56.6    YES NVRAM   up          up
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                unassigned      YES NVRAM   administratively down down
Serial1/1                unassigned      YES NVRAM   administratively down down
Serial1/2                unassigned      YES NVRAM   administratively down down
Serial1/3                unassigned      YES NVRAM   administratively down down
Loopback106              148.49.106.1   YES NVRAM   up          up
R6#

```

No IP address has been configured on R6 interface Et0/0.26. When you read ahead, you will see that this IP address should be assigned dynamically by a DHCP server configured on R2.

First, skip this requirement, type an IP address manually, and verify the ping on VLAN 26 again:

```

R6(config)#interface Ethernet 0/0.26
R6(config-subif)#ip address 148.49.26.6 255.255.255.0
R6(config-subif)#end
R6#
R6#ping 148.49.26.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 148.49.26.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R6#

```

The ping between R2 and R6 is still not working, even after manually configuring an IP address on R6 interface Et0/0.26.

Verify the interface configuration on R2:

```

R2#show running-config interface Ethernet 0/0.26
Building configuration...

Current configuration : 147 bytes
!
interface Ethernet0/0.26
 encapsulation dot1Q 26
 ip address 148.49.26.2 255.255.255.0

```

The lab indicates that no IP address may be configured on main interface Et0/0 on R2. Therefore, the IP address 148.49.26.2 has been configured on subinterface Et0/0.26 with dot1q encapsulation.

Let us verify how SW1 Et0/1 has been configured:

```

SW1#show running-config interface Ethernet 0/1
Building configuration...

Current configuration : 107 bytes

```

```

!
interface Ethernet0/1
  switchport access vlan 26
  switchport mode access
end

```

**Likely cause: Interface Et0/1 is not trunking on SW1.**

Even though only one VLAN is configured on the link between R2 Et0/0 and SW1 Et0/1, we have to enable trunking on Et0/1 on SW1, since R2 is trunking on Et0/0. This is because no IP address was allowed on main interface Et0/0 on R2.

We see on SW1 that interface Et0/2 is in access mode, in VLAN 26:

```
SW1#show vlan brief
```

| VLAN | Name               | Status    | Ports               |
|------|--------------------|-----------|---------------------|
| 1    | default            | active    | Et0/0, Et1/2, Et1/3 |
| 16   | VLAN0016           | active    |                     |
| 26   | VLAN0026           | active    | Et0/1               |
| 34   | VLAN0034           | active    | Et0/2, Et0/3        |
| 36   | VLAN0036           | active    |                     |
| 56   | VLAN0056           | active    | Et1/0               |
| 1002 | fddi-default       | act/unsup |                     |
| 1003 | token-ring-default | act/unsup |                     |
| 1004 | fddinet-default    | act/unsup |                     |
| 1005 | trnet-default      | act/unsup |                     |

**Resolution: Enable trunking on SW1 Et0/1.**

We simply enable dot1q on SW1 Et0/1:

```

SW1(config)#interface Ethernet0/1
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport mode trunk
SW1(config-if)# switchport trunk allowed vlan 26
SW1(config-if)#end
SW1#

```

And now we can verify on SW1:

```
SW1#show interfaces trunk
```

| Port  | Mode | Encapsulation | Status   | Native vlan |
|-------|------|---------------|----------|-------------|
| Et0/1 | on   | 802.1q        | trunking | 1           |
| Et1/1 | on   | 802.1q        | trunking | 1           |
| Et2/0 | on   | 802.1q        | trunking | 1           |
| Et2/1 | on   | 802.1q        | trunking | 1           |
| Et2/2 | on   | 802.1q        | trunking | 1           |
| Et2/3 | on   | 802.1q        | trunking | 1           |

```
Port Vlan allowed on trunk
```

|       |             |
|-------|-------------|
| Et0/1 | 26          |
| Et1/1 | 16,26,36,56 |
| Et2/0 | 56          |
| Et2/1 | 56          |
| Et2/2 | 16,36       |
| Et2/3 | 16,36       |

```
Port Vlan allowed and active in management domain
```

|       |             |
|-------|-------------|
| Et0/1 | 26          |
| Et1/1 | 16,26,36,56 |
| Et2/0 | 56          |

```

Et2/1          56
Et2/2          16,36

Port           Vlans allowed and active in management domain
Et2/3          16,36

Port           Vlans in spanning tree forwarding state and not pruned
Et0/1          26
Et1/1          16,26,36,56
Et2/0          56
Et2/1          none
Et2/2          16,36
Et2/3          16,36
SW1#

```

Verify if R2 can ping R6:

```

R2#ping 148.49.26.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 148.49.26.6, timeout is 2 seconds:
!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

```

Connectivity on VLAN 26 is complete, but in the IP Services troubleshooting ticket, we will need to remove the manually configured IP address on R6 Et0/0.26 to let R6 obtain its IP address by the DHCP server that is configured on R2.

## 1.2. Symptom: STP is blocking ports between SW3 and SW4.

### Analysis and Testing:

Verify the Spanning Tree Protocol (STP) port state on SW3:

```

SW3#show spanning-tree int e2/2

```

| Vlan     | Role | Sts | Cost | Prio.Nbr | Type |
|----------|------|-----|------|----------|------|
| VLAN0034 | Root | FWD | 100  | 128.67   | Shr  |
| VLAN0056 | Root | FWD | 100  | 128.67   | Shr  |

```

SW3#show spanning-tree int e2/3

```

| Vlan     | Role | Sts | Cost | Prio.Nbr | Type |
|----------|------|-----|------|----------|------|
| VLAN0034 | Altn | BLK | 100  | 128.68   | Shr  |
| VLAN0056 | Altn | BLK | 100  | 128.68   | Shr  |

```

SW3#

```

We can see that interface E2/3 is blocking.

The reasons for this are as follows:

- SW4 is the root switch.
- All ports on SW4 are designated ports.
- Interface E2/2 on SW3 is the root port, because the bridge protocol data unit (BPDU) received on E2/2 is better than the BPDU received on E2/3. To compare the BPDU received on E2/2 and the BPDU received on E2/3, we first compare the root path cost—it is the same (100). Because the root path cost is the same, we compare the bridge ID—it is also the same (SW4). Because the bridge ID is the same, we compare the port ID. The port ID indicated on the BPDU received on E2/2 is 128.67. The port ID indicated on the BPDU received on E2/3 is 128.68.

To make sure that there is no blocking interface between two switches, configure an EtherChannel between them.

**Likely cause: No EtherChannel has been configured.**

Verify if a port channel has been configured on SW3 and SW4:

```
SW3#sh ip interface brief | include Port
Port-channell1          unassigned      YES unset  down      down

SW4#sh ip interface brief | include Port
Port-channell1          unassigned      YES unset  down      down
```

Port channel 1 has been configured.

Let us see if the encapsulation configured on these port channels is coherent between SW3 and SW4:

```
SW3#sh running-config interface port-channel 1
Building configuration...
```

```
Current configuration : 126 bytes
!
interface Port-channell1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34,56
 switchport mode trunk
end
```

```
SW4#sh running-config interface port-channel 1
Building configuration...
```

```
Current configuration : 126 bytes
!
interface Port-channell1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34,56
 switchport mode trunk
end
```

Both port channels have been configured as trunking with dot1q, and both allow only VLAN 34 and VLAN 56.

Let us see why the state of the port channel is **down**:

```
SW3# show etherchannel
      Channel-group listing:
-----
```

```
Group: 1
-----
Group state = L2
Ports: 0      Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol:    -
Minimum Links: 0
```

```
SW3#
```

```
SW4#show etherchannel
      Channel-group listing:
-----
```

```
Group: 1
-----
Group state = L2
Ports: 0      Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol:    -
Minimum Links: 0
```

SW4#

SW3#show etherchannel summary

```
Flags: D - down          P - in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3       S - Layer2
       U - in use       f - failed to allocate aggregator
       u - unsuitable for bundling
       w - waiting to be aggregated
       d - default port
```

Number of channel-groups in use: 1

Number of aggregators: 1

| Group | Port-channel | Protocol | Ports |
|-------|--------------|----------|-------|
| 1     | Po1(S)       | -        |       |

SW4#show etherchannel summary

```
Flags: D - down          P - in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3       S - Layer2
       U - in use       f - failed to allocate aggregator
       u - unsuitable for bundling
       w - waiting to be aggregated
       d - default port
```

Number of channel-groups in use: 1

Number of aggregators: 1

| Group | Port-channel | Protocol | Ports |
|-------|--------------|----------|-------|
| 1     | Po1(S)       | -        |       |

We can see that **no interface** has been grouped under the port channel interface.

### **Resolution: Group interfaces E2/2 and E2/3 under port channel 1, on SW3 and SW4.**

The lab requires the use of a standard protocol.

When enabling an EtherChannel, we may use a Cisco proprietary protocol (Port Aggregation Protocol, or **PAgP**) or a standard protocol (Link Aggregation Control Protocol, or **LACP**). We will configure the EtherChannel using LACP:

```
SW3(config)#interface range Ethernet2/2 - 3
SW3(config-if-range)#channel-group 1 mode ?
  active      Enable LACP unconditionally
  auto        Enable PAgP only if a PAgP device is detected
  desirable   Enable PAgP unconditionally
  on          Enable Etherchannel only
  passive     Enable LACP only if a LACP device is detected

SW3(config-if-range)#channel-group 1 mode active
SW3(config-if-range)#end
SW3#

SW4(config)# interface range Ethernet2/2 - 3
SW4(config-if-range)#channel-group 1 mode active
SW4(config-if-range)#end
SW4#
```

Let us verify that the EtherChannel is working properly on SW4:

```
SW4#show etherchannel
      Channel-group listing:
      -----
```

```
Group: 1
-----
Group state = L2
Ports: 2      Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol:    -
```

```
SW4#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3       S - Layer2
        U - in use       f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

```
Number of channel-groups in use: 1
Number of aggregators:          1
```

| Group | Port-channel | Protocol | Ports               |
|-------|--------------|----------|---------------------|
| 1     | Po1(SU)      | -        | Et2/2 (P) Et2/3 (P) |

Verify that the EtherChannel is working properly on SW3:

```
SW3#show etherchannel
      Channel-group listing:
      -----
```

```
Group: 1
-----
Group state = L2
Ports: 2      Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol:    -
```

```
SW3#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3       S - Layer2
        U - in use       f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
```

```
Number of channel-groups in use: 1
Number of aggregators:          1
```

| Group | Port-channel | Protocol | Ports               |
|-------|--------------|----------|---------------------|
| 1     | Po1(SU)      | -        | Et2/2 (P) Et2/3 (P) |

Now verify that there is no blocking port on SW3:

```
SW3#show spanning-tree int e2/2
```

| Vlan     | Role | Sts | Cost | Prio.Nbr | Type |
|----------|------|-----|------|----------|------|
| VLAN0034 | Root | FWD | 56   | 128.514  | Shr  |

```

VLAN0056          Root FWD 56          128.514 Shr
SW3#show spanning-tree int e2/3

Vlan              Role Sts Cost          Prio.Nbr Type
-----
VLAN0034          Root FWD 56          128.514 Shr
VLAN0056          Root FWD 56          128.514 Shr
SW3#
SW3#show spanning-tree vlan 56

VLAN0056
Spanning tree enabled protocol ieee
Root ID            Priority    4152
Address            aabb.cc00.0a00
Cost               56
Port              514 (Port-channel1)
Hello Time        2 sec    Max Age 20 sec    Forward Delay 15 sec

Bridge ID         Priority    32824 (priority 32768 sys-id-ext 56)
Address            aabb.cc00.0900
Hello Time        2 sec    Max Age 20 sec    Forward Delay 15 sec
Aging Time        300

Interface         Role Sts Cost          Prio.Nbr Type
-----
Et0/1             Desg FWD 100          128.2    Shr
Po1               Root FWD 56          128.514 Shr

```

We can see that the STP considers interfaces E2/2 and E2/3 as one logical port, **port channel 1**, which is therefore in **forwarding** mode.

### 1.3. Symptom: SW4 is not the root switch for all VLANs.

#### Analysis and Testing:

You verified the root bridge configuration on SW1, SW2, SW3, and SW4 and found that SW4 is not the root switch for all VLANs:

```

SW1#show spanning-tree root

Vlan              Root ID            Root Cost  Hello Time  Max Age  Fwd Dly  Root Port
-----
VLAN0001          32769 aabb.cc00.0700    0         2         20      15
VLAN0016          32784 aabb.cc00.0700    0         2         20      15
VLAN0026          32794 aabb.cc00.0700    0         2         20      15
VLAN0034          32802 aabb.cc00.0700    0         2         20      15
VLAN0036          32804 aabb.cc00.0700    0         2         20      15
VLAN0056          4152  aabb.cc00.0a00    100        2         20      15  Et2/0
SW1#

SW2#show spanning-tree root

Vlan              Root ID            Root Cost  Hello Time  Max Age  Fwd Dly  Root Port
-----
VLAN0001          32769 aabb.cc00.0800    0         2         20      15
VLAN0016          32784 aabb.cc00.0700    100        2         20      15  Et2/2
VLAN0036          32804 aabb.cc00.0700    100        2         20      15  Et2/2

```

We can see that for **VLAN 56**, SW4 is considered the root switch by SW1, SW3, and SW4.

But for **VLANs 16 and 36**, SW1 is considered the root switch by SW1 and SW2.

Finally, for VLANs 26 and 34, SW1 is considered the root switch by itself.

To make sure SW4 is the root bridge, let us verify the bridge priorities of each switch.

SW1:

```
SW1#show spanning-tree bridge
```

| Vlan     | Bridge ID         | Hello Time     | Max Age | Fwd Dly | Protocol |      |
|----------|-------------------|----------------|---------|---------|----------|------|
| VLAN0001 | 32769 (32768, 1)  | aabb.cc00.0700 | 2       | 20      | 15       | ieee |
| VLAN0016 | 32784 (32768, 16) | aabb.cc00.0700 | 2       | 20      | 15       | ieee |
| VLAN0026 | 32794 (32768, 26) | aabb.cc00.0700 | 2       | 20      | 15       | ieee |
| VLAN0034 | 32802 (32768, 34) | aabb.cc00.0700 | 2       | 20      | 15       | ieee |
| VLAN0036 | 32804 (32768, 36) | aabb.cc00.0700 | 2       | 20      | 15       | ieee |
| VLAN0056 | 32824 (32768, 56) | aabb.cc00.0700 | 2       | 20      | 15       | ieee |

```
SW1#
```

SW2:

```
SW2#show spanning-tree bridge
```

| Vlan     | Bridge ID         | Hello Time     | Max Age | Fwd Dly | Protocol |      |
|----------|-------------------|----------------|---------|---------|----------|------|
| VLAN0001 | 32769 (32768, 1)  | aabb.cc00.0800 | 2       | 20      | 15       | ieee |
| VLAN0016 | 32784 (32768, 16) | aabb.cc00.0800 | 2       | 20      | 15       | ieee |
| VLAN0036 | 32804 (32768, 36) | aabb.cc00.0800 | 2       | 20      | 15       | ieee |

```
SW2#
```

SW3:

```
SW3#show spanning-tree bridge
```

| Vlan     | Bridge ID         | Hello Time     | Max Age | Fwd Dly | Protocol |      |
|----------|-------------------|----------------|---------|---------|----------|------|
| VLAN0001 | 32769 (32768, 1)  | aabb.cc00.0900 | 2       | 20      | 15       | ieee |
| VLAN0034 | 32802 (32768, 34) | aabb.cc00.0900 | 2       | 20      | 15       | ieee |
| VLAN0056 | 32824 (32768, 56) | aabb.cc00.0900 | 2       | 20      | 15       | ieee |

```
SW3#
```

SW4:

```
SW4#show spanning-tree bridge
```

| Vlan     | Bridge ID       | Hello Time     | Max Age | Fwd Dly | Protocol |      |
|----------|-----------------|----------------|---------|---------|----------|------|
| VLAN0001 | 4097 (4096, 1)  | aabb.cc00.0a00 | 2       | 20      | 15       | ieee |
| VLAN0034 | 4130 (4096, 34) | aabb.cc00.0a00 | 2       | 20      | 15       | ieee |
| VLAN0056 | 4152 (4096, 56) | aabb.cc00.0a00 | 2       | 20      | 15       | ieee |

```
SW4#
```

We can see that SW4 has a lower priority (4096) than SW1, SW2, or SW3 (32768), for VLANs. Therefore, we should expect SW4 to become the root bridge for the listed VLANs. Note that SW4 does not show the spanning tree details for VLAN 16, VLAN 26, and VLAN 36,

Verify the VLANs on SW4:

```
SW4#show vlan brie
```

| VLAN Name   | Status | Ports  |
|-------------|--------|--|
| 1 default   | active | Et0/0, Et0/1, Et0/2, Et0/3<br>Et1/0, Et1/1, Et1/2, Et1/3 |
| 16 VLAN0016 | active |  |

```

26 VLAN0026 active
34 VLAN0034 active
36 VLAN0036 active
56 VLAN0056 active
1002 fddi-default act/unsup
1003 token-ring-default act/unsup
1004 fddinet-default act/unsup
1005 trnet-default act/unsup
SW4#

```

Note that all necessary VLANs are configured on SW4. But they are not listed in the output of the spanning tree because they are not associated with any ports on SW4.

**Likely cause: VLANs 16, 26, 34, and 36 are not allowed on trunk links between SW4 and SW1.**

If a trunk link does not allow a specific VLAN, and if this VLAN is configured on both sides of this trunk, then the root switch will not be consistent on both sides of the trunk.

Let us verify which VLANs are allowed between SW1 and SW4:

```
SW1#show interfaces trunk
```

```

Port                Mode           Encapsulation  Status        Native vlan
Et0/1                on             802.1q         trunking      1
Et1/1                on             802.1q         trunking      1
Et2/0                on             802.1q         trunking      1
Et2/1                on             802.1q         trunking      1
Et2/2                on             802.1q         trunking      1
Et2/3                on             802.1q         trunking      1

Port                Vlans allowed on trunk
Et0/1                26
Et1/1                16,26,36,56
Et2/0                56
Et2/1                56
Et2/2                16,36
Et2/3                16,36

Port                Vlans allowed and active in management domain
Et0/1                26
Et1/1                16,26,36,56
Et2/0                56
Et2/1                56
Et2/2                16,36

Port                Vlans allowed and active in management domain
Et2/3                16,36

Port                Vlans in spanning tree forwarding state and not pruned
Et0/1                26
Et1/1                16,26,36,56
Et2/0                56
Et2/1                none
Et2/2                16,36
Et2/3                16,36
SW1#

```

```
SW4#show interfaces trunk
```

```

Port                Mode           Encapsulation  Status        Native vlan
Et2/0                on             802.1q         trunking      1
Et2/1                on             802.1q         trunking      1
Po1                  on             802.1q         trunking      1

```

```

Port                Vlans allowed on trunk
Et2/0                56
Et2/1                56
Po1                  34,56

Port                Vlans allowed and active in management domain
Et2/0                56
Et2/1                56
Po1                  34,56

Port                Vlans in spanning tree forwarding state and not pruned
Et2/0                56
Et2/1                56
Po1                  34,56
SW4#

```

**Resolution: Allow VLANs 16, 26, 34, and 36 on trunk links between SW1 and SW4.**

```

SW1(config)#interface range Ethernet2/0 - 1
SW1(config-if-range)#switchport trunk allowed vlan add 16,26,34,36

SW4(config)#interface range Ethernet2/0 - 1
SW4(config-if-range)#switchport trunk allowed vlan add 16,26,34,36

```

Let us first verify that these four VLANs are now added on trunk links:

```

SW1#show interfaces trunk

Port                Mode                Encapsulation  Status        Native vlan
Et0/1                on                  802.1q         trunking      1
Et1/1                on                  802.1q         trunking      1
Et2/0                on                  802.1q         trunking      1
Et2/1                on                  802.1q         trunking      1
Et2/2                on                  802.1q         trunking      1
Et2/3                on                  802.1q         trunking      1

Port                Vlans allowed on trunk
Et0/1                26
Et1/1                16,26,36,56
Et2/0                16,26,34,36,56
Et2/1                16,26,34,36,56
Et2/2                16,36
Et2/3                16,36

Port                Vlans allowed and active in management domain
Et0/1                26
Et1/1                16,26,36,56
Et2/0                16,26,34,36,56
Et2/1                16,26,34,36,56
Et2/2                16,36

Port                Vlans allowed and active in management domain
Et2/3                16,36

Port                Vlans in spanning tree forwarding state and not pruned
Et0/1                26
Et1/1                16,26,36,56
Et2/0                34,56
Et2/1                none
Et2/2                16,36
Et2/3                none
SW1#
SW4#show interfaces trunk

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

```

```

Po1                on                802.1q                trunking                1

Port                Vlans allowed on trunk
Et2/0              16,26,34,36,56
Et2/1              16,26,34,36,56
Po1                34,56

Port                Vlans allowed and active in management domain
Et2/0              16,26,34,36,56
Et2/1              16,26,34,36,56
Po1                34,56

Port                Vlans in spanning tree forwarding state and not pruned
Et2/0              16,26,34,36,56
Et2/1              none
Po1                34,56
SW4#

```

Let us verify that SW1, SW2, and SW3 now recognize SW4 as the root switch for all VLANs available in their respective spanning tree instances:

```
SW1#show spanning-tree root
```

| Vlan     | Root ID              | Root Cost | Hello Time | Max Age | Fwd Dly | Root Port |
|----------|----------------------|-----------|------------|---------|---------|-----------|
| VLAN0001 | 32769 aabb.cc00.0700 | 0         | 2          | 20      | 15      |           |
| VLAN0016 | 4112 aabb.cc00.0a00  | 100       | 2          | 20      | 15      | Et2/0     |
| VLAN0026 | 4122 aabb.cc00.0a00  | 100       | 2          | 20      | 15      | Et2/0     |
| VLAN0034 | 4130 aabb.cc00.0a00  | 100       | 2          | 20      | 15      | Et2/0     |
| VLAN0036 | 4132 aabb.cc00.0a00  | 100       | 2          | 20      | 15      | Et2/0     |
| VLAN0056 | 4152 aabb.cc00.0a00  | 100       | 2          | 20      | 15      | Et2/0     |

```
SW2#show spanning-tree root
```

| Vlan     | Root ID              | Root Cost | Hello Time | Max Age | Fwd Dly | Root Port |
|----------|----------------------|-----------|------------|---------|---------|-----------|
| VLAN0001 | 32769 aabb.cc00.0800 | 0         | 2          | 20      | 15      |           |
| VLAN0016 | 4112 aabb.cc00.0a00  | 200       | 2          | 20      | 15      | Et2/2     |
| VLAN0036 | 4132 aabb.cc00.0a00  | 200       | 2          | 20      | 15      | Et2/2     |

```
SW3#show spanning-tree root
```

| Vlan     | Root ID              | Root Cost | Hello Time | Max Age | Fwd Dly | Root Port |
|----------|----------------------|-----------|------------|---------|---------|-----------|
| VLAN0001 | 32769 aabb.cc00.0900 | 0         | 2          | 20      | 15      |           |
| VLAN0034 | 4130 aabb.cc00.0a00  | 56        | 2          | 20      | 15      | Po1       |
| VLAN0056 | 4152 aabb.cc00.0a00  | 56        | 2          | 20      | 15      | Po1       |

```
SW4#show spanning-tree root
```

| Vlan     | Root ID             | Root Cost | Hello Time | Max Age | Fwd Dly | Root Port |
|----------|---------------------|-----------|------------|---------|---------|-----------|
| VLAN0001 | 4097 aabb.cc00.0a00 | 0         | 2          | 20      | 15      |           |
| VLAN0016 | 4112 aabb.cc00.0a00 | 0         | 2          | 20      | 15      |           |
| VLAN0026 | 4122 aabb.cc00.0a00 | 0         | 2          | 20      | 15      |           |
| VLAN0034 | 4130 aabb.cc00.0a00 | 0         | 2          | 20      | 15      |           |
| VLAN0036 | 4132 aabb.cc00.0a00 | 0         | 2          | 20      | 15      |           |
| VLAN0056 | 4152 aabb.cc00.0a00 | 0         | 2          | 20      | 15      |           |

#### 1.4. Symptom: Broadcast traffic crosses all trunk links.

##### *Analysis and Testing:*

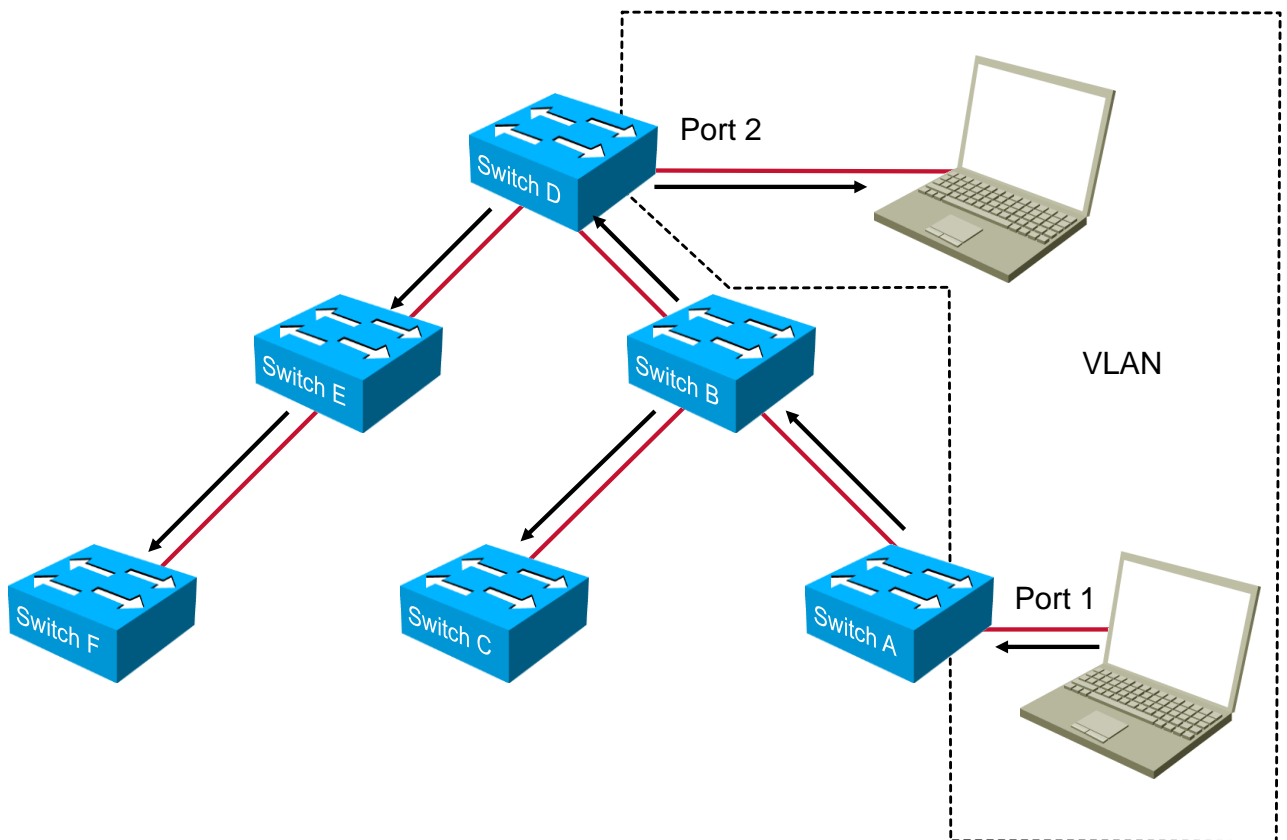
When you configure VLAN Trunking Protocol (VTP), by default VTP pruning is not enabled.

As per Cisco documentation: “VTP pruning increases network available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to reach the destination devices.

Without VTP pruning, a switch floods broadcast, multicast, and unknown unicast traffic across all trunk links within a VTP domain even though receiving switches might discard them.”

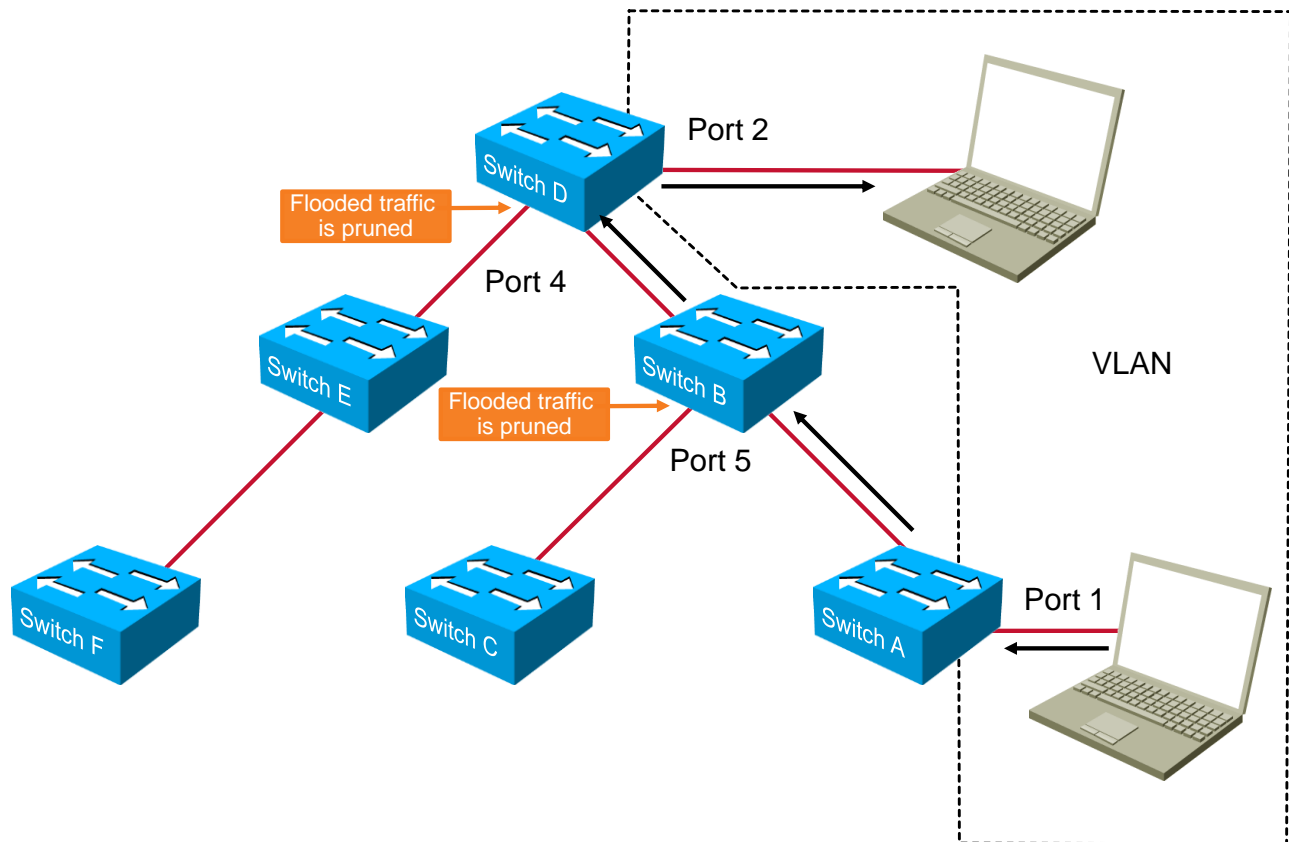
The figure below shows a switched network without VTP pruning enabled. Port 1 on Switch A and Port 2 on Switch D are assigned to the red VLAN. If a broadcast is sent from the host connected to Switch A, Switch A floods the broadcast, and every switch in the network receives it, even though Switches C, E, and F have no ports in the red VLAN.

#### Flooding Traffic Without VTP Pruning



The next figure shows a switched network with VTP pruning enabled. The broadcast traffic from Switch A is not forwarded to Switches C, E, and F because traffic for the red VLAN has been pruned on the links shown (Port 5 on Switch B and Port 4 on Switch D).

## Optimized Flooded Traffic with VTP Pruning



**Likely cause: VTP pruning has not been enabled.**

VTP pruning may be enabled only on the VTP server.

Since the lab indicates that “if SW1, SW2, or SW3 reboot, they should rely on SW4 to rebuild their VLAN configuration,” this means that SW4 is the VTP server and SW1, SW2, and SW3 are VTP clients.

Let us verify the status of VTP pruning on SW4:

```
SW4#show vtp status
VTP Version                : 3 (capable)
Configuration Revision      : 5
Maximum VLANs supported locally : 1005
Number of existing VLANs    : 10
VTP Operating Mode         : Server
VTP Domain Name            : TROUBLESHOOTING
VTP Pruning Mode           : Disabled (Operationally Disabled)
VTP V2 Mode                : Disabled
VTP Traps Generation       : Disabled
MD5 digest                 : 0x6B 0xC8 0xB7 0x5D 0x91 0x3C 0x9E 0x66
Configuration last modified by 0.0.0.0 at 8-14-13 13:23:29
```

```
Local updater ID is 0.0.0.0 (no valid interface found)
VTP version running      : 1
SW4#
```

**Resolution: Enable VTP pruning.**

```
SW4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW4(config)#vtp pruning
Pruning switched on
SW4(config)#
*Aug 14 16:23:41.859: %SW_VLAN-6-VTP_PRUNING_CHANGE: VTP Operational Pruning
Enabled.
SW4(config)#end
SW4#
```

Let us verify that pruning has been enabled:

```
SW4#show vtp status
VTP Version          : 3 (capable)
Configuration Revision : 6
Maximum VLANs supported locally : 1005
Number of existing VLANs : 10
VTP Operating Mode   : Server
VTP Domain Name     : TROUBLESHOOTING
VTP Pruning Mode    : Enabled (Operationally Enabled)
VTP V2 Mode         : Disabled
VTP Traps Generation : Disabled
MD5 digest          : 0xDD 0xDF 0x34 0x2B 0x2E 0xAF 0xB7 0x29
Configuration last modified by 0.0.0.0 at 8-14-13 16:23:41
Local updater ID is 0.0.0.0 (no valid interface found)
VTP version running  : 1
SW4#
```

You can check on any VTP client that VTP pruning has also been automatically enabled.

Verify SW1, for example:

```
SW1#show vtp status
VTP Version          : 3 (capable)
Configuration Revision : 6
Maximum VLANs supported locally : 1005
Number of existing VLANs : 10
VTP Operating Mode   : Client
VTP Domain Name     : TROUBLESHOOTING
VTP Pruning Mode    : Enabled (Operationally Enabled)
VTP V2 Mode         : Disabled
VTP Traps Generation : Disabled
MD5 digest          : 0xDD 0xDF 0x34 0x2B 0x2E 0xAF 0xB7 0x29
Configuration last modified by 0.0.0.0 at 8-14-13 16:23:41
VTP version running  : 1
SW1#
```

---

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

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

---

## 2. IPv4 OSPF Troubleshooting Section

### 2.1. Symptom: The area is incorrect, and virtual links are down.

#### Analysis and Testing:

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

```
R1#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
Se1/0      1    0     148.49.21.1/24  64    P2P   1/1
VL1        1    0     0.0.0.0/0        65535 DOWN  0/0
VL0        1    0     0.0.0.0/0        65535 DOWN  0/0
Et0/2     1   134   148.49.134.1/24  10    DR    0/0
R1#
```

```
R2#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
Lo102     1    0     148.49.102.1/24  1     P2P   0/0
Se1/0     1    0     148.49.21.2/24   64    DR    1/1
R2#
```

```
R3#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
VL0        1    0     0.0.0.0/0        65535 DOWN  0/0
Lo103     1    34    148.49.103.1/24  1     P2P   0/0
Et0/2     1    34    148.49.134.3/24  10    DROTH 0/0
Et0/0     1    34    148.49.34.3/24   10    P2MP  2/2
R3#
```

```
R4#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
VL0        1    0     0.0.0.0/0        65535 DOWN  0/0
Lo104     1    34    148.49.100.1/24  1     P2P   0/0
Et0/0     1    34    148.49.34.4/24   10    P2MP  2/2
Et0/2     1   134   148.49.134.4/24  10    DROTH 0/0
R4#
```

```
R7#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
Et0/0     1    34    148.49.34.7/24   10    P2MP  2/2
R7#
```

From this, we see that the Ethernet0/2 interface on R3 has been assigned to Area 34 instead of Area 134. Furthermore, all virtual links are down.

Verify the OSPF virtual links on R1, R3, and R4:

```
R1#show ip ospf virtual-links
Virtual Link OSPF_VL1 to router 148.49.103.1 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 134
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                65535      no            no            Base
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Virtual Link OSPF_VL0 to router 148.49.104.1 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 134
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                65535      no            no            Base
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

R1#

```
R3#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 148.49.111.1 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 134
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                65535      no            no            Base
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
R3#
```

```
R4#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 148.49.111.1 is down
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 134
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                65535      no            no            Base
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
R4#
```

When configuring a virtual link, the address to be used is the router ID of the distant OSPF device.

**Likely cause: The OSPF router IDs are incorrect.**

Verify the OSPF router IDs of R1, R3, and R4:

```
R1#show ip ospf | include ID
Routing Process "ospf 1" with ID 148.49.101.1

R3#show ip ospf | include ID
Routing Process "ospf 1" with ID 148.49.103.1

R4#show ip ospf | include ID
Routing Process "ospf 1" with ID 148.49.100.1
```

Note that on R1, the router ID that is used to configure the virtual link toward R4 is incorrect. The router ID IP address should be 148.49.104.1 instead of 148.49.100.1.

Also, remember that lab requires the router ID of R1 to be 148.49.111.1. Therefore, the router IDs that are used on R3 and R4 to configure the virtual link toward R1 are correct, but the router ID of R1 itself is not conforming to lab requirements.

**Resolution: Correct the router IDs.**

First, we need to configure the Ethernet0/2 interface on R3 in Area 134 instead of Area 34.

```
R3(config)#router ospf 1
R3(config-router)#network 148.49.134.0 0.0.00.255 area 134
```

We now see that the Ethernet0/2 interface on R3 has been assigned to Area 134.

```
R3#show ip ospf interface brief
Interface  PID  Area      IP Address/Mask      Cost  State Nbrs F/C
VL0        1    0         0.0.0.0/0            65535 DOWN  0/0
Lo103     1    34        148.49.103.1/24      1     P2P   0/0
Et0/0     1    34        148.49.34.3/24       10    P2MP  2/2
Et0/2     1    134       148.49.134.3/24     10    DROTH 0/1
R3#
```

Then, we need to remove the virtual link to the IP address 148.49.104.1 on R1 and replace it with a virtual link to the router ID of R4.

```
R1(config-router)#router ospf 1
R1(config-router)#no area 134 virtual-link 148.49.104.1
R1(config-router)#area 134 virtual-link 148.49.100.1
```

Finally, we need to modify the router ID of R1 to make sure it becomes 148.49.111.1.

But remember that the lab requirement prohibits the use of the command **router-id**. So the only way we have to influence the calculation of the router ID is to make sure that R1 has a loopback interface with IP address 148.49.111.1.

We may not modify the IP address of the Loopback101 interface on R1; otherwise, subnet 148.49.101.1/24 will not be present anymore in our network.

Therefore, we will create a new loopback interface:

```
R1(config)#interface loopback 111
R1(config-if)#ip address 148.49.111.1 255.255.255.0
```

Remember to reload R1 to have OSPF take into account the new router ID. To avoid having to reload the router, you may manually configure any OSPF router ID, reload the OSPF process, remove the router ID, and again reload the OSPF process. This way, 148.49.111.1 will become the router ID of R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1#clear ip ospf 1 process
Reset OSPF process? [no]: yes

R1(config)#router ospf 1
R1(config-router)#no router-id 1.1.1.1

R1#clear ip ospf 1 process
Reset OSPF process? [no]: yes

R1#show ip ospf | include ID
Routing Process "ospf 1" with ID 148.49.111.1
```

Let us verify that all virtual links are now up:

```
R1#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
Se1/0      1    0     148.49.21.1/24   64    P2P   1/1
VL1        1    0     148.49.134.1/24  10    P2P   1/1
VL0        1    0     148.49.134.1/24  10    P2P   1/1
Et0/2     1    134   148.49.134.1/24  10    DR    2/2
R1#
```

```
R3#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
VL0        1    0     148.49.134.3/24  10    P2P   1/1
Lo103     1    34    148.49.103.1/24  1     P2P   0/0
Et0/0     1    34    148.49.34.3/24   10    P2MP  2/2
Et0/2     1    134   148.49.134.3/24  10    DROTH 1/1
R3#
```

```
R4#show ip ospf interface brief
Interface  PID  Area  IP Address/Mask  Cost  State Nbrs F/C
VL0        1    0     148.49.134.4/24  10    P2P   1/1
Lo104     1    34    148.49.100.1/24  1     P2P   0/0
```

```

Et0/0      1      34      148.49.34.4/24      10      P2MP  2/2
Et0/2      1      134     148.49.134.4/24    10      DROTH 1/1
R4#

```

## 2.2. Symptom: The adjacency between R1 and R2 flaps.

### *Analysis and Testing:*

A good starting point when OSPF adjacency flaps is to analyze the OSPF network type.

The lab requires that routers in Area 0 should elect a designated router (DR) and backup designated router (BDR). The OSPF network type may either be broadcast or nonbroadcast.

Furthermore, the lab requires that in Area 0, the hello interval should be 10 seconds, and the dead interval should be 40 seconds. But you are prohibited from using the commands **ip ospf hello-interval** and **ip ospf dead-interval**.

Therefore, the OSPF network type in Area 0 may only be broadcast.

Let us verify if this is the case:

```

R1#show ip ospf interface s1/0
Serial1/0 is up, line protocol is up
  Internet Address 148.49.21.1/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 148.49.111.1, Network Type POINT_TO_POINT, Cost: 64
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                  64        no            no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 3/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 8
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 148.49.102.1
  Suppress hello for 0 neighbor(s)
R1#

```

```

R2#show ip ospf interface s1/0
Serial1/0 is up, line protocol is up
  Internet Address 148.49.21.2/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 148.49.102.1, Network Type BROADCAST, Cost: 64
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
  0                  64        no            no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 148.49.102.1, Interface address 148.49.21.2
  Backup Designated router (ID) 148.49.111.1, Interface address 148.49.21.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:08
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 7
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 148.49.111.1 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
R2#

```

**Likely cause: The OSPF network type is incorrect.**

We should modify the network type on R1 interface S1/0.

```
R1(config)#int s1/0
R1(config-subif)#ip ospf network broadcast
```

Verify the OSPF configuration and adjacency on R1:

```
R1#show ip ospf interface s1/0
Serial1/0 is up, line protocol is up
Internet Address 148.49.21.1/24, Area 0, Attached via Network Statement
Process ID 1, Router ID 148.49.111.1, Network Type BROADCAST, Cost: 64
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
0                  64         no            no            Base
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 148.49.102.1, Interface address 148.49.21.2
Backup Designated router (ID) 148.49.111.1, Interface address 148.49.21.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:06
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Index 3/3, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 8
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 148.49.102.1 (Designated Router)
Suppress hello for 0 neighbor(s)
R1#
```

**2.3. Symptom: The metric for subnet 148.49.26.0 is not the same on R1 and R7.**

**Analysis and Testing:**

The lab requires that the OSPF metric for subnet 148.49.26.0 should be the same on R1 and R7.

Let us verify this:

```
R1#show ip route 148.49.26.0
Routing entry for 148.49.26.0/24
  Known via "ospf 1", distance 110, metric 90, type extern 1
  Redistributing via rip
  Last update from 148.49.21.2 on Serial1/0, 00:10:15 ago
  Routing Descriptor Blocks:
  * 148.49.21.2, from 148.49.102.1, 00:10:15 ago, via Serial1/0
    Route metric is 90, traffic share count is 1
R1#
```

```
R7# show ip route 148.49.26.0
Routing entry for 148.49.26.0/24
  Known via "ospf 1", distance 110, metric 110, type extern 1
  Last update from 148.49.34.3 on Ethernet0/0, 00:11:04 ago
  Routing Descriptor Blocks:
  * 148.49.34.4, from 148.49.102.1, 00:11:04 ago, via Ethernet0/0
    Route metric is 110, traffic share count is 1
  148.49.34.3, from 148.49.102.1, 00:11:04 ago, via Ethernet0/0
    Route metric is 110, traffic share count is 1
R7#
```

When a route is redistributed into OSPF, it may either retain the same metric within all areas or see its OPSF metric increase when propagated away from the router redistributing this route into OSPF.

- In the first case, the metric type would be set to 2, and the route will appear with the prefix O **E2** in the routing table.
- In the second case, the metric type would be set to 1, and the route will appear with the prefix O **E1** in the routing table.

By default, when a subnet is redistributed into OSPF, the metric type is 2.

R2 is redistributing subnet 148.49.26.0/24 into OSPF.

***Likely cause: The metric type on R2 is incorrect.***

Verify how subnet 148.49.26.0/24 is redistributed on R2 in the running configuration:

```
router ospf 1
 redistribute connected subnets route-map Ospf26
 !
 !
route-map Ospf26 permit 10
 match interface Ethernet0/0.26
 set metric 26
 set metric-type type-1
```

***Resolution: On R2, modify the metric type of route map OSPF 26.***

```
R2(config)#route-map Ospf26
R2(config-route-map)#set metric-type type-2
```

Let us verify that the OSPF metric is the same on R1 and R7:

```
R1#show ip route 148.49.26.0
Routing entry for 148.49.26.0/24
  Known via "ospf 1", distance 110, metric 26, type extern 2, forward metric 64
  Redistributing via rip
  Last update from 148.49.21.2 on Serial1/0, 00:00:10 ago
  Routing Descriptor Blocks:
  * 148.49.21.2, from 148.49.102.1, 00:00:10 ago, via Serial1/0
    Route metric is 26, traffic share count is 1
R1#
R7# show ip route 148.49.26.0
Routing entry for 148.49.26.0/24
  Known via "ospf 1", distance 110, metric 26, type extern 2, forward metric 84
  Last update from 148.49.34.3 on Ethernet0/0, 00:00:28 ago
  Routing Descriptor Blocks:
  * 148.49.34.4, from 148.49.102.1, 00:00:28 ago, via Ethernet0/0
    Route metric is 26, traffic share count is 1
  148.49.34.3, from 148.49.102.1, 00:00:28 ago, via Ethernet0/0
    Route metric is 26, traffic share count is 1
R7#
```

---

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

---

### 3. IPv4 EIGRP Troubleshooting Section

#### 3.1. Symptom: R5 has no EIGRP neighbor relationship over subnet 172.16.56.0/24.

##### *Analysis and Testing:*

There are four devices in Enhanced Interior Gateway Routing Protocol (EIGRP) autonomous system (AS) 56. Therefore, each device should have three EIGRP neighbor relationships. However, R5 has no EIGRP neighbor relationships at all.

```
R5#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
R5#
```

```
R6#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface                Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)              (ms)              Cnt  Num
1   148.49.56.8             Et0/0.56                10 00:00:41    13    100  0  3
0   148.49.56.9             Et0/0.56                12 00:02:09    12    100  0  3
R6#
```

```
R8#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface                Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)              (ms)              Cnt  Num
1   148.49.56.9             Et0/0                    13 00:00:38    266   1596  0  5
0   148.49.56.6             Et0/0                    12 00:02:20     5    100  0  5
R8#
```

```
R9#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface                Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)              (ms)              Cnt  Num
1   148.49.56.8             Et0/0                    11 00:01:07     1    100  0  6
0   148.49.56.6             Et0/0                    11 00:04:13     8    100  0  3
R9#
```

***Likely cause: R5 does not have the required neighbor statements.***

The lab requires that EIGRP does not send any multicast frames on VLAN 56. Check the configuration on each device to verify the correct neighbor statements.

```
R5#sh run | section eigrp
router eigrp 56
 network 148.49.56.0 0.0.0.255
 network 148.49.105.0 0.0.0.255
 auto-summary
```

```
R6#sh run | section eigrp
router eigrp 56
 redistribute connected route-map Eigrp106
 redistribute rip metric 1 1 1 1 1
 network 148.49.36.0 0.0.0.255
 network 148.49.56.0 0.0.0.255
 auto-summary
 neighbor 148.49.56.5 Ethernet0/0.56
 neighbor 148.49.56.9 Ethernet0/0.56
 neighbor 148.49.56.8 Ethernet0/0.56
 redistribute eigrp 56 metric 14
```

```
R8#sh run | begin router eigrp
router eigrp 56
```

```

network 148.49.56.0 0.0.0.255
auto-summary
neighbor 148.49.56.9 Ethernet0/0
neighbor 148.49.56.6 Ethernet0/0
neighbor 148.49.56.5 Ethernet0/0
!

R9#sh run | begin router eigrp
router eigrp 56
network 148.49.56.0 0.0.0.255
redistribute connected
auto-summary
neighbor 148.49.56.6 Ethernet0/0
neighbor 148.49.56.5 Ethernet0/0
neighbor 148.49.56.8 Ethernet0/0
!
!

```

Here, we see that R6, R8, and R9 have the required neighbor statement, but R5 does not.

**Resolution: On R5, configure the correct neighbor statements.**

```

R5(config)#router eigrp 56
R5(config-router)# neighbor 148.49.56.8 Ethernet0/0
R5(config-router)# neighbor 148.49.56.9 Ethernet0/0
R5(config-router)# neighbor 148.49.56.6 Ethernet0/0

```

```

R5#sh run | section eigrp
router eigrp 56
network 148.49.56.0 0.0.0.255
network 148.49.105.0 0.0.0.255
auto-summary
neighbor 148.49.56.8 Ethernet0/0
neighbor 148.49.56.9 Ethernet0/0
neighbor 148.49.56.6 Ethernet0/0
R5#

```

Check the EIGRP neighbor relationships on VLAN 56 again:

```

R5#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface          Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)          (ms)          Cnt  Num
2   148.49.56.6             Et0/0             14 00:01:18    12   100  0  8
1   148.49.56.8             Et0/0             10 00:01:18    10   100  0  9
0   148.49.56.9             Et0/0             14 00:01:18     8   100  0  8
R5#
R6#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface          Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)          (ms)          Cnt  Num
2   148.49.56.5             Et0/0.56          11 00:01:24     8   100  0  7
1   148.49.56.8             Et0/0.56          13 00:10:47    13   100  0  9
0   148.49.56.9             Et0/0.56          13 00:12:15    10   100  0  8
R6#
R8#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface          Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)          (ms)          Cnt  Num
2   148.49.56.5             Et0/0             14 00:01:35     7   100  0  7
1   148.49.56.9             Et0/0             11 00:09:15   213  1278  0  8
0   148.49.56.6             Et0/0             11 00:10:58     5   100  0  8
R8#
R9#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(56)
H   Address                Interface          Hold Uptime    SRTT    RTO  Q  Seq

```

|   |             |       | (sec)       | (ms) | Cnt | Num |
|---|-------------|-------|-------------|------|-----|-----|
| 2 | 148.49.56.5 | Et0/0 | 11 00:01:43 | 5    | 100 | 0 7 |
| 1 | 148.49.56.8 | Et0/0 | 13 00:09:28 | 1    | 100 | 0 9 |
| 0 | 148.49.56.6 | Et0/0 | 12 00:12:34 | 6    | 100 | 0 8 |

R9#

Note that there are three neighbors for each device in EIGRP AS 56.

### 3.2. Symptom: The delay configuration for redistribution of subnet 148.49.106.0/24 on R6 is incorrect.

#### *Analysis and Testing:*

The lab requires that Loopback106 be redistributed into EIGRP AS 56 with specific metric values. Delay should be set to 100,000 microseconds.

```
R6#show ip eigrp topology 148.49.106.0/24
EIGRP-IPv4 Topology Entry for AS(56)/ID(148.49.106.1) for 148.49.106.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 25600512
Descriptor Blocks:
0.0.0.0, from Rconnected, Send flag is 0x0
Composite metric is (25600512/0), route is External
Vector metric:
  Minimum bandwidth is 4000000 Kbit
  Total delay is 1000000 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1514
  Hop count is 0
  Originating router is 148.49.106.1
External data:
  AS number of route is 0
  External protocol is Connected, external metric is 0
  Administrator tag is 106 (0x00000006A)
R6#
```

Note that the delay is 1,000,000 microseconds instead of 100,000 microseconds.

Verify the route map that is used on R6:

```
router eigrp 56
 redistribute connected route-map Eigrp106
!
route-map Eigrp106 permit 10
 match interface Loopback106
 set metric 4000000 100000 255 1 1514
 set tag 106
```

#### *Likely cause: The delay is configured with an incorrect unit of measurement.*

Verify the unit of measurement in which the metric should be configured:

```
R6(config)#route-map Eigrp106
R6(config-route-map)#set metric 4000000 ?
+/-<delay>      Add or subtract delay
<1-4294967295>  IGRP delay metric, in 10 microsecond units
<cr>
```

We used the question mark (?) key to get help from Cisco IOS Software, and the output indicated that the delay should be expressed in tens of microseconds.

**Resolution: On R6, modify the metric for route map EIGRP 106.**

```
R6(config)#route-map Eigrp106
R6(config-route-map)#set metric 4000000 10000 255 1 1514
```

Let us now verify that the correct metric is used in the topology database:

```
R6#show ip eigrp topology 148.49.106.0/24
EIGRP-IPv4 Topology Entry for AS(56)/ID(148.49.106.1) for 148.49.106.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2560512
Descriptor Blocks:
 0.0.0.0, from Rconnected, Send flag is 0x0
  Composite metric is (2560512/0), route is External
  Vector metric:
    Minimum bandwidth is 4000000 Kbit
    Total delay is 100000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1514
    Hop count is 0
    Originating router is 148.49.106.1
  External data:
    AS number of route is 0
    External protocol is Connected, external metric is 0
    Administrator tag is 106 (0x0000006A)
R6#
```

---

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

---

## 4. IPv4 RIP Troubleshooting Section

### 4.1. Symptom: R2 is not listed as a RIP information source.

**Analysis and Testing:**

A good starting point is to check RIP configuration on each device running RIP to verify interfaces where RIP has been activated and is sending updates.

We need to make sure only VLAN 16 and VLAN 26 see multicast RIP updates, as indicated in the lab.

R6 should send RIP updates to E0/0.16 and E0/0.26 and receive updates from R1 and R2:

```
R6#show ip protocols | begin "rip"
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 25 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: eigrp 56, rip
  Default version control: send version 2, receive version 2
  Interface          Send Recv  Triggered RIP  Key-chain
  Ethernet0/0.16     2      2
  Ethernet0/0.26     2      2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    148.49.0.0
  Passive Interface(s):
```

```

Ethernet0/0
Ethernet0/0.36
Ethernet0/0.56
Ethernet0/1
Ethernet0/2
Ethernet0/3
Serial1/0
Serial1/1
Passive Interface(s):
Serial1/2
Serial1/3
Loopback106
RG-AR-IF-INPUT1
VoIP-Null0
Routing Information Sources:
Gateway          Distance      Last Update
148.49.16.1      120           00:00:25
Distance: (default is 120)

```

R6#

Note that R6 sends the RIP version 2 multicast updates on the Ethernet0/0.16 and Ethernet0/0.26 interfaces and receives updates from R1. R6 does not show R2 in the list of the RIP routing information sources.

R1 should send RIP updates to E0/1 and receive RIP updates from R6:

```

R1#show ip protocols | begin "rip"
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: ospf 1 (internal, external 1 & 2, nssa-external 1 & 2)

  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface          Send Recv Triggered RIP Key-chain
    Ethernet0/1         2        2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    148.49.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
    Serial1/2
    Serial1/3
  Passive Interface(s):
    Loopback101
    Loopback111
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway          Distance      Last Update
    148.49.16.6      120           00:00:04
  Distance: (default is 120)

```

R1#

Note that R1 sends the RIP version 2 multicast updates on the Ethernet0/1 interface and receives updates from R6.

R2 should send RIP updates to E0/0.26 and receive updates from R6:

```
R2#show ip protocols | begin "rip"
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Outgoing routes will have 2 added to metric if on list Rip_metric_2
  Sending updates every 30 seconds, next due in 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: ospf 1 (internal, external 1 & 2, nssa-external 1 & 2)

  Redistributing: rip
  Default version control: send version 2, receive version 2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    148.49.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/0.26
    Ethernet0/1
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
    Serial1/2
  Passive Interface(s):
    Serial1/3
    Loopback102
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    148.49.26.6      120          00:00:20
  Distance: (default is 120)
```

R2#

Note that R2 does not send the RIP version 2 multicast updates on the Ethernet0/0.26 interface. But R2 receives updates from R6.

**Likely cause: R2 is not sending updates to E0/0.26 because the interface is passive.**

Verify the RIP configuration on R2:

```
R2#show run | section rip
router rip
  version 2
  redistribute ospf 1 route-map O2R
  passive-interface default
  offset-list Rip_metric_2 out 2
  network 148.49.0.0
R2#
```

Note that all interfaces are passive on R2.

**Resolution: On R2, remove interface E0/0.26 from the list of RIP passive interfaces.**

```
R2(config)#router rip
R2(config-router)#no passive-interface Ethernet 0/0.26
```

Verify that R2 is now sending RIP updates out of E0/0.26.

```
R2#show ip protocols | begin "rip"
Routing Protocol is "rip"
```

```
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Outgoing routes will have 2 added to metric if on list Rip_metric_2
Sending updates every 30 seconds, next due in 19 seconds
Invalid after 180 seconds, hold down 180, flushed after 240
Redistributing: ospf 1 (internal, external 1 & 2, nssa-external 1 & 2)
```

```
Redistributing: rip
Default version control: send version 2, receive version 2
Interface          Send Recv Triggered RIP Key-chain
Ethernet0/0.26     2      2
Automatic network summarization is in effect
Maximum path: 4
Routing for Networks:
 148.49.0.0
Passive Interface(s):
 Ethernet0/0
 Ethernet0/1
 Ethernet0/2
 Ethernet0/3
 Serial1/0
 Serial1/1
Passive Interface(s):
 Serial1/2
 Serial1/3
 Loopback102
 RG-AR-IF-INPUT1
 VoIP-Null0
Routing Information Sources:
 Gateway          Distance      Last Update
 148.49.26.6      120           00:00:08
Distance: (default is 120)
```

R2#

### Verify the sources of RIP routing information on R6 again:

```
R6#show ip protocols | begin "rip"
Routing Protocol is "rip"
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Sending updates every 30 seconds, next due in 25 seconds
Invalid after 180 seconds, hold down 180, flushed after 240
Redistributing: eigrp 56, rip
Default version control: send version 2, receive version 2
Interface          Send Recv Triggered RIP Key-chain
Ethernet0/0.16     2      2
Ethernet0/0.26     2      2
Automatic network summarization is in effect
Maximum path: 4
Routing for Networks:
 148.49.0.0
Passive Interface(s):
 Ethernet0/0
 Ethernet0/0.36
 Ethernet0/0.56
 Ethernet0/1
 Ethernet0/2
 Ethernet0/3
 Serial1/0
 Serial1/1
Passive Interface(s):
 Serial1/2
 Serial1/3
 Loopback106
 RG-AR-IF-INPUT1
```

```

VoIP-Null0
Routing Information Sources:
  Gateway         Distance      Last Update
  148.49.16.1     120           00:00:13
  148.49.26.2     120           00:00:17
Distance: (default is 120)

```

R6#

## 5. IPv4 Redistribution Troubleshooting Section

### 5.1. Symptom: Network 148.49.100.0/24 is not in the routing tables of R5, R6, R8, and R9.

#### *Analysis and Testing:*

You found that 148.49.100.0/24 is not in the routing tables of R5, R6, R8, and R9. Here is an example from R6:

```

R6#show ip route 148.49.100.0
% Subnet not in table
R6#

```

You are required to mutually redistribute between OSPF and RIP on R1 and R2. Network 172.16.100.0 should be redistributed from OSPF to RIP, and then from RIP to EIGRP.

Verify which subnets are redistributed from OSPF to RIP on R1:

```

R1#show running-config | section rip
  redistribute rip subnets
router rip
  version 2
  redistribute ospf 1 route-map O2R
  passive-interface default
  no passive-interface Ethernet0/1
  network 148.49.0.0
R1#
R1#show route-map O2R
route-map O2R, permit, sequence 10
  Match clauses:
    ip address (access-lists): O2R
  Set clauses:
    metric 1
  Policy routing matches: 0 packets, 0 bytes

R1#show ip access-lists O2R
Standard IP access list O2R
 10 permit 148.49.103.0, wildcard bits 0.0.0.255 (2 matches)
 20 permit 148.49.107.0, wildcard bits 0.0.0.255 (2 matches)
 30 permit 148.49.34.0, wildcard bits 0.0.0.255 (5 matches)

```

Note that subnet 148.49.100.0 is not permitted by access list O2R.

Now verify how R2 is redistributing OSPF to RIP:

```

R2#show run | section rip
router rip
  version 2
  redistribute ospf 1 route-map O2R
  passive-interface default
  no passive-interface Ethernet0/0.26
  offset-list Rip_metric_2 out 2
  network 148.49.0.0
R2#
R2#show route-map O2R
route-map O2R, permit, sequence 10
  Match clauses:

```

```
ip address (access-lists): O2R
Set clauses:
metric 1
Policy routing matches: 0 packets, 0 bytes
```

```
R2#show ip access-lists O2R
Standard IP access list O2R
10 permit 148.49.103.0, wildcard bits 0.0.0.255 (4 matches)
20 permit 148.49.107.0, wildcard bits 0.0.0.255 (2 matches)
30 permit 148.49.34.0, wildcard bits 0.0.0.255 (18 matches)
```

Note that subnet 148.49.100.0/24 is never redistributed from OSPF to RIP.

**Likely cause: There is a missing statement in access list O2R.**

You could say that you just need to add a fourth line on access list O2R, to permit subnet 148.49.100.0/24.

But remember that the lab says, “On R1 and R2, you may not create any additional route maps or access lists. You may modify the current access lists and current route maps, but each access list may not have more than three statements.”

Rewrite the access list O2R to meet the following requirements:

- There are no more than three statements.
- Include subnet 148.49.100.0/24 with the other three subnets already permitted.

**Resolution: Rewrite the access list O2R.**

How can you do this? You need to permit two subnets using a single statement on R1 and R2:

```
R2(config)#ip access-list standard O2R
R2(config-std-nacl)#no 10
R2(config-std-nacl)#10 permit 148.49.96.0 0.0.7.255
```

Now check the routing table on R6:

```
R6#show ip route 148.49.100.0
Routing entry for 148.49.100.0/24
Known via "rip", distance 120, metric 1
Redistributing via eigrp 56, rip
Advertised by eigrp 56 metric 1 1 1 1 1
Last update from 148.49.16.1 on Ethernet0/0.16, 00:00:05 ago
Routing Descriptor Blocks:
* 148.49.26.2, from 148.49.26.2, 00:00:05 ago, via Ethernet0/0.26
  Route metric is 1, traffic share count is 1
  148.49.16.1, from 148.49.16.1, 00:00:05 ago, via Ethernet0/0.16
    Route metric is 1, traffic share count is 1
R6#
```

R6 may reach subnet 148.49.100.0/24.

**5.2. Symptom: Networks 148.49.16.0/24 and 148.49.26.0/24 are not in the routing tables of R5, R8, and R9.**

Example from R5:

```
R5#show ip route 148.49.16.0
% Subnet not in table
R5#show ip route 148.49.26.0
% Subnet not in table
R5#
```

## ***Analysis and Testing:***

Networks 148.49.16.0 and 148.49.26.0 are connected to R6.

R6, R5, R8, and R9 belong to EIGRP AS 56.

Verify if R6 is redistributing connected routes into EIGRP:

```
R6#sh run | section eigrp
router eigrp 56
network 148.49.36.0 0.0.0.255
network 148.49.56.0 0.0.0.255
redistribute connected route-map Eigrp106
redistribute rip metric 1 1 1 1
auto-summary
neighbor 148.49.56.9 Ethernet0/0.56
neighbor 148.49.56.5 Ethernet0/0.56
neighbor 148.49.56.8 Ethernet0/0.56
redistribute eigrp 56 metric 14
R6#

R6#sh route-map Eigrp106
route-map Eigrp106, permit, sequence 10
Match clauses:
  interface Loopback106
Set clauses:
  metric 4000000 100000 255 1 1514
  tag 106
Policy routing matches: 0 packets, 0 bytes
```

You can see that interfaces E0/0.16 and E0/0.26 are not redistributed into EIGRP.

## ***Likely cause: There is a missing statement in route map EIGRP 106.***

You need to add a new line to the route map to allow redistribution of subnets 148.49.16.0/24 and 148.49.26.0/24.

## ***Resolution: Add a statement to route map EIGRP 106 to allow missing subnets.***

Make sure that you allow these two additional subnets using an additional route map permit statement:

```
R6(config)#route-map Eigrp106 20
R6(config-route-map)#match interface e0/0.16 e0/0.26
```

Verify that EIGRP neighbors may now reach subnets 148.49.16.0 and 148.49.26.0:

```
R5#show ip route 148.49.16.0
Routing entry for 148.49.16.0/24
  Known via "eigrp 56", distance 170, metric 307200, type external
  Redistributing via eigrp 56
  Last update from 148.49.56.6 on Ethernet0/0, 00:01:13 ago
  Routing Descriptor Blocks:
  * 148.49.56.6, from 148.49.56.6, 00:01:13 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
R5#show ip route 148.49.26.0
Routing entry for 148.49.26.0/24
  Known via "eigrp 56", distance 170, metric 307200, type external
  Redistributing via eigrp 56
  Last update from 148.49.56.6 on Ethernet0/0, 00:01:31 ago
  Routing Descriptor Blocks:
  * 148.49.56.6, from 148.49.56.6, 00:01:31 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
```

R5#

### 5.3. Symptom: R2 is sending traffic for 148.49.56.0 to R1 instead of R6.

#### *Analysis and Testing:*

The lab indicates that R2 should prefer R6 as the next hop for all traffic destined to VLAN 56.

Verify forwarding on R2:

```
R2#show ip route 148.49.56.0
Routing entry for 148.49.56.0/24
  Known via "ospf 1", distance 119, metric 20, type extern 2, forward metric 64
  Redistributing via rip
  Last update from 148.49.21.1 on Serial1/0, 01:50:01 ago
  Routing Descriptor Blocks:
  * 148.49.21.1, from 148.49.111.1, 01:50:01 ago, via Serial1/0
    Route metric is 20, traffic share count is 1
```

R2#

You can see that R2 is sending traffic for 148.49.56.0 to R1 via the Serial1/0 interface.

You also note that the administrative distance for this route is 119 instead of the regular 110 value attributed to OSPF by default.

Find out where this unusual administrative distance comes from.

```
R2#sh run | section ospf 1
router ospf 1
 redistribute connected subnets route-map Ospf26
 network 148.49.21.0 0.0.0.255 area 0
 network 148.49.102.0 0.0.0.255 area 0
 distance 119 0.0.0.0 255.255.255.255 Vlan56
 redistribute ospf 1 route-map O2R
```

R2#

```
R2#sh ip access-lists Vlan56
Standard IP access list Vlan56
 10 permit 148.49.56.0, wildcard bits 0.0.0.255 (10 matches)
```

#### *Likely cause: The administrative distance value is incorrect.*

In order for R2 to use R6 instead of R1 as a next hop, R2 must prefer a RIP route to an OSPF route. Therefore, you need to make sure the RIP route has a lower administrative distance than the OSPF route.

In this lab case:

- The RIP administrative distance was left at 120.
- The OSPF administrative distance was configured to 119.

#### *Resolution: Increase the OSPF administrative distance to 121 for subnet 148.49.56.0/24.*

```
R2(config)#router ospf 1
R2(config-router)#distance 121 0.0.0.0 255.255.255.255 Vlan56
```

Now verify the routing table on R2:

```
R2#show ip route 148.49.56.0
Routing entry for 148.49.56.0/24
  Known via "rip", distance 120, metric 1
```

```

Redistributing via rip
Last update from 148.49.26.6 on Ethernet0/0.26, 00:00:15 ago
Routing Descriptor Blocks:
* 148.49.26.6, from 148.49.26.6, 00:00:15 ago, via Ethernet0/0.26
  Route metric is 1, traffic share count is 1

```

## 5.4. Symptom: R6 load-balances between R1 and R2.

### *Analysis and Testing:*

The lab indicates that R6 should prefer R1 as the next hop for all traffic destined to all subnets advertised by R1, with the following exception: R6 should prefer R2 as the next hop for all traffic destined to subnet 148.49.21.0/24.

However, we can see that R6 load-balances between R1 and R2:

```

R6#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

148.49.0.0/16 is variably subnetted, 23 subnets, 2 masks
R       148.49.21.0/24 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.34.3/32 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.34.4/32 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.34.7/32 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.100.0/24 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.101.0/24 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.102.0/24 [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.103.0/24 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.107.0/24 [120/1] via 148.49.26.2, 00:00:20, Ethernet0/0.26
        [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.111.0/24 [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R       148.49.134.0/24 [120/1] via 148.49.16.1, 00:00:15, Ethernet0/0.16
R6#

```

The section indicates that this behavior on R6 is only a preference; this means that if the preferred link ever goes down, a backup link should be available to send traffic. The only way to meet this requirement is to modify metrics for these subnets.

But the lab prohibits any configuration on R6. Therefore, the only way is to modify these metrics on R1 or R2, or both.

### *Likely cause: The metric values are incorrect.*

If you want to make sure that R6 uses R1 as a next hop, you can increase the metric sent by R2 to R6:

```

R2(config)#route-map O2R
R2(config-route-map)#set metric 2

```

Check the routing table of R6 again:

```
R6# show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override
```

Gateway of last resort is not set

```
148.49.0.0/16 is variably subnetted, 23 subnets, 2 masks
R      148.49.21.0/24 [120/1] via 148.49.26.2, 00:00:13, Ethernet0/0.26
      [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.34.3/32 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.34.4/32 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.34.7/32 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.100.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.101.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.102.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.103.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.107.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.111.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R      148.49.134.0/24 [120/1] via 148.49.16.1, 00:00:14, Ethernet0/0.16
R6#
```

R6 now sends all traffic to R1, but it still continues load-balancing between R1 and R2 for traffic destined to subnet 148.49.21.0/24.

So you have not yet completely fulfilled the requirement. On R6, you need to increase the metric for subnet 148.49.21.0 via 148.49.16.1. But again, you are not allowed to make any configuration changes on R6. Therefore, you will need to make the change on R1.

**Resolution: Increase the metric sent by R1 to R6 for subnet 148.49.21.0/24.**

```
R1(config)#ip access-list standard Rip_metric_2
R1(config-std-nacl)# permit 148.49.21.0 0.0.0.255

R1(config)#router rip
R1(config-router)# offset-list Rip_metric_2 out 2
```

Verify the routing table on R6:

```
R6# show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

148.49.0.0/16 is variably subnetted, 23 subnets, 2 masks
R      148.49.21.0/24 [120/1] via 148.49.26.2, 00:00:12, Ethernet0/0.26
R      148.49.34.3/32 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
```

```

R      148.49.34.4/32 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.34.7/32 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.100.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.101.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.102.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.103.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.107.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.111.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R      148.49.134.0/24 [120/1] via 148.49.16.1, 00:00:11, Ethernet0/0.16
R6#

```

You can see that the next hop is always 148.49.16.1 except for subnet 148.49.21.0.

Now that you seem to have addressed all of the IPv4 unicast issues, you will test reachability using this simple Tool Command Language (Tcl) script. You enter the command **tclsh** and paste in this script. When it is complete, you will have a record of successful and unsuccessful pings.

You enter the command **tclquit** to exit the command interpreter.

```

tclsh
foreach address {
148.49.101.1
148.49.16.1
148.49.134.1
148.49.21.1
148.49.21.2
148.49.102.1
148.49.26.2
148.49.36.3
148.49.134.3
148.49.103.1
148.49.34.3
148.49.100.1
148.49.134.4
148.49.34.4
148.49.56.5
148.49.105.1
148.49.106.1
148.49.36.6
148.49.56.6
148.49.16.6
148.49.26.6
148.49.107.1
148.49.34.7
148.49.56.9
148.49.109.1
148.49.56.8} {ping $address}

```

## 6. Security Troubleshooting Section

### 6.1. Symptom: On SW2, more than two MAC addresses are allowed on E1/0 to E1/1.

#### *Analysis and Testing:*

Interfaces E1/0 to E1/1 should allow only two MAC addresses. Verify if port security has been configured:

```

SW2#show port-security
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
              (Count)           (Count)           (Count)
-----
Total Addresses in System (excluding one mac per port)  : 0
Max Addresses limit in System (excluding one mac per port) : 4096

```

It looks as if port security has not been configured.

**Likely cause: Port security is not configured.**

Check the configuration of interface E1/0, for example:

```
SW2#show run int e1/0
Building configuration...

Current configuration : 147 bytes
!
interface Ethernet1/0
  switchport mode access
  switchport port-security maximum 2
  switchport port-security mac-address sticky
  duplex auto
end

SW2#
```

Port security has been configured but not activated.

**Resolution: Activate port security.**

```
SW2(config)#interface range Ethernet1/0 - 1
SW2(config-if-range)#switchport port-security
```

```
SW2#show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
          (Count)          (Count)          (Count)
-----
Et1/0           2                0                0                Shutdown
Et1/1           2                0                0                Shutdown
-----
Total Addresses in System (excluding one mac per port)      : 0
Max Addresses limit in System (excluding one mac per port) : 4096

SW2#
```

**6.2. Symptom: SW2 PortFast interfaces do not go into the error-disabled state when a BPDU is received.**

**Analysis and Testing:**

A PortFast interface goes into the error-disabled state if BPDU guard is enabled.

You may enable BPDU guard either per interface or globally on all PortFast interfaces.

**Likely cause: BPDU guard is not enabled for PortFast interfaces.**

```
SW2#show spanning-tree summary
Switch is in pvst mode
Root bridge for: VLAN0001
EtherChannel misconfig guard          is enabled
Extended system ID                    is enabled
Portfast Default                      is disabled
Portfast Edge BPDU Guard Default      is disabled
Portfast Edge BPDU Filter Default     is disabled
Loopguard Default                    is disabled
Platform PVST Simulation              is enabled
PVST Simulation Default               is enabled but inactive in pvst mode
Bridge Assurance                      is enabled but inactive in pvst mode
```

```

Pathcost method used          is short
UplinkFast                   is disabled
BackboneFast                  is disabled

```

| Name     | Blocking | Listening | Learning | Forwarding | STP Active |
|----------|----------|-----------|----------|------------|------------|
| VLAN0001 | 0        | 0         | 0        | 2          | 2          |
| VLAN0016 | 1        | 0         | 0        | 2          | 3          |
| VLAN0036 | 1        | 0         | 0        | 2          | 3          |
| 3 vlans  | 2        | 0         | 0        | 6          | 8          |

SW2#

**Resolution: Configure PortFast BPDU guard.**

```

SW2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW2(config)#spanning-tree portfast bpduguard default
SW2(config)#end
SW2#

```

```

SW2#show spanning-tree summary
Switch is in pvst mode
Root bridge for: VLAN0001
EtherChannel misconfig guard          is enabled
Extended system ID                   is enabled
Portfast Default                      is disabled
Portfast Edge BPDU Guard Default     is enabled
Portfast Edge BPDU Filter Default    is disabled
Loopguard Default                    is disabled
Platform PVST Simulation             is enabled
PVST Simulation Default              is enabled but inactive in pvst mode
Bridge Assurance                     is enabled but inactive in pvst mode
Pathcost method used                 is short
UplinkFast                           is disabled
BackboneFast                         is disabled

```

| Name     | Blocking | Listening | Learning | Forwarding | STP Active |
|----------|----------|-----------|----------|------------|------------|
| VLAN0001 | 0        | 0         | 0        | 2          | 2          |
| VLAN0016 | 1        | 0         | 0        | 2          | 3          |
| VLAN0036 | 1        | 0         | 0        | 2          | 3          |
| 3 vlans  | 2        | 0         | 0        | 6          | 8          |

SW2#

**6.3. Symptom: Any switch behind port E0/3 on SW2 may become a root switch.**

**Analysis and Testing:**

To mitigate STP manipulation and enforce root bridge placement in the network, we use root guard.

Let us see whether root guard has been configured on E0/3:

```

SW2#sh run int e0/3
Building configuration...

Current configuration : 42 bytes
!
interface Ethernet0/3
 duplex auto
end

```

**Likely cause: Root guard is not configured.**

To make sure a port shuts down if a better BPDU is received on this port, activate root guard.

**Resolution: Configure root guard.**

```
SW2(config)#int e0/3
SW2(config-if)#spanning-tree guard root
```

## 7. IPv6 Troubleshooting Section

### 7.1. Symptom: R3 and R4 do not have a route to FEC0::106:0/125.

**Analysis and Testing:**

When verifying the IPv6 prefixes, you found that FEC0::106:0/125 is missing in the routing tables of R1, R3, and R4:

```
R1#show ipv6 route FEC0::106:0
% Route not found
R1#
R3#show ipv6 route FEC0::106:0
% Route not found
R3#
R4#show ipv6 route FEC0::106:0
% Route not found
R4#
```

**Likely cause: FEC0::106:0/125 in OSPF on R6 is advertised incorrectly.**

According to the lab requirements, the Loopback106 network should be redistributed into OSPFv3 Area 0.

Verify the OSPFv3 configuration on R6:

```
R6#show ipv6 ospf database external self-originate
          OSPFv3 Router with ID (148.49.106.1) (Process ID 1)
R6#
```

Verify the IPv6 redistribution configuration on R6:

```
R6#show running-config | section ipv6 router ospf
ipv6 router ospf 1
 redistribute connected nssa-only
R6#
```

Note that the redistribution is configured with the incorrect key nssa-only. The OSPFv3 Area0 is not a not-so-stubby area (NSSA).

**Resolution: Remove the nssa-only keyword from the redistribution statement on R6.**

```
R6(config)#ipv6 router ospf 1
R6(config-rtr)# no redistribute connected nssa-only
R6(config-rtr)# redistribute connected
R6(config-rtr)# end
R6#
```

Verify the OSPFv3 database on R6 again:

```
R6#show ipv6 ospf database external self-originate
```

OSPFv3 Router with ID (148.49.106.1) (Process ID 1)

Type-5 AS External Link States

```
LS age: 8
LS Type: AS External Link
Link State ID: 1
Advertising Router: 148.49.106.1
LS Seq Number: 80000001
Checksum: 0xA6D7
Length: 44
Prefix Address: FEC0::106:0
Prefix Length: 125, Options: None
Metric Type: 2 (Larger than any link state path)
Metric: 20
```

R6#

---

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

---

## 8. QoS Troubleshooting Section

### 8.1. Symptom: Telnet traffic from the Loopback103 interface of R3 to the Loopback109 interface of R9 is policed.

#### *Analysis and Testing:*

When you connect with Telnet from the Loopback103 interface of R3 to the Loopback109 interface of R9, you notice that the traffic is rated to 8000 b/s on R9. According to the lab requirements, this type of traffic should be excluded from the control plane policy traffic classification criteria.

Clear control plane policing counters and verify the QoS policy configuration on R9:

```
R9#clear control-plane *
R9#
R9#show policy-map control-plane
Control Plane

Service-policy input: TELNET

Class-map: TELNET (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group 100
  police:
    cir 8000 bps, bc 1500 bytes
    conformed 0 packets, 0 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps

Class-map: class-default (match-any)
  49 packets, 3816 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any

R9#
```

Connect with Telnet from the Loopback103 interface of R3 to the Loopback109 interface of R9:

```
R3#telnet 148.49.109.1 /source-interface loo103
Trying 148.49.109.1 ... Open

-----
Cisco 360 R&S Exercise Workbook
Product, POD location: cierswbv5-te-lab09-sc, SJ
Device: R9
-----
R9#exit

[Connection to 148.49.109.1 closed by foreign host]
R3#
```

Verify the QoS policy configuration on R9:

```
R9#show policy-map control-plane
Control Plane

Service-policy input: TELNET

Class-map: TELNET (match-all)
  21 packets, 1180 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group 100
  police:
    cir 8000 bps, bc 1500 bytes
    conformed 21 packets, 1180 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps

Class-map: class-default (match-any)
  152 packets, 11790 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any
R9#
```

Note that the Telnet traffic from the Loopback103 interface of R3 to the Loopback109 interface of R9 is conformed and policed to 8000 b/s.

**Likely cause:** *The traffic classification criteria are incorrect.*

Verify access list 100 on R9:

```
R9#show access-lists 100
Extended IP access list 100
 10 permit tcp host 148.49.103.1 host 148.49.109.1 eq telnet (21 matches)
R9#
```

Note that the Telnet traffic from the Loopback103 interface of R3 to the Loopback109 interface of R9 is permitted to be policed.

**Resolution:** *On R9, modify the traffic classification criteria.*

```
R9#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R9(config)#no access-list 100
R9(config)# access-list 100 deny tcp host 148.49.103.1 host 148.49.109.1 eq telnet
R9(config)#access-list 100 permit tcp any any eq telnet
R9(config)#end
R9#
```

Note that access list 100 restricts the Telnet traffic from the Loopback103 interface of R3 to the Loopback109 interface of R9 from the policing requirements and permits any other Telnet traffic for policing.

Clear control plane policing counters and verify the QoS policy configuration on R9:

```
R9#clear control-plane *
R9#
R9#show policy-map control-plane
Control Plane

Service-policy input: TELNET

Class-map: TELNET (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group 100
  police:
    cir 8000 bps, bc 1500 bytes
    conformed 0 packets, 0 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps

Class-map: class-default (match-any)
  49 packets, 3816 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any
R9#
```

Connect with Telnet from the Loopback103 interface of R3 to the Loopback109 interface of R9:

```
R3#telnet 148.49.109.1 /source-interface loo103
Trying 148.49.109.1 ... Open

-----
Cisco 360 R&S Exercise Workbook
Product, POD location: cierswbv5-te-lab09-sc, SJ
Device: R9
-----
R9#exit

[Connection to 148.49.109.1 closed by foreign host]
R3#
```

Verify the QoS policy configuration on R9:

```
R9#show policy-map control-plane
Control Plane

Service-policy input: TELNET

Class-map: TELNET (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group 100
  police:
    cir 8000 bps, bc 1500 bytes
    conformed 0 packets, 0 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps
```

```
Class-map: class-default (match-any)
  47 packets, 3322 bytes
  5 minute offered rate 1000 bps, drop rate 0000 bps
  Match: any
R9#
```

Note that the Telnet traffic from the Loopback103 interface of R3 to the Loopback109 interface of R9 is not policed.

Connect with Telnet from the default interface of R3 to the Loopback109 interface of R9:

```
R3#telnet 148.49.109.1
Trying 148.49.109.1 ... Open

-----
Cisco 360 R&S Exercise Workbook
Product, POD location: cierswbv5-te-lab09-sc, SJ
Device:                R9
-----
R9#exit

[Connection to 148.49.109.1 closed by foreign host]
R3#
```

Verify the QoS policy configuration on R9:

```
R9#show policy-map control-plane
Control Plane

Service-policy input: TELNET

Class-map: TELNET (match-all)
  24 packets, 1344 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group 100
  police:
    cir 8000 bps, bc 1500 bytes
    conformed 24 packets, 1344 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceeded 0000 bps

Class-map: class-default (match-any)
  127 packets, 9650 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any
R9#
```

Note that this type of Telnet traffic is conformed and policed to 8000 b/s.

## 9. IP Services Troubleshooting Section

### 9.1. Symptom: The IP address for interface E0/0.26 on R6 is incorrect.

#### *Analysis and Testing:*

Remember that R2 should attribute an IP address to all devices located on VLAN 26.

Therefore, you need to reconfigure the interface E0/0.26 on R6 to make it receive an IP address dynamically by having R2 act as a DHCP server.

```
R6#conf t
```

```

Enter configuration commands, one per line. End with CNTL/Z.
R6(config)#interface Ethernet 0/0.26
R6(config-subif)#ip address dhcp
R6(config-subif)#end
R6#
*Aug 15 03:44:09.373: %SYS-5-CONFIG_I: Configured from console by console
(cierswbv5-te-lab09-sc, SJ)
R6#

```

Verify which IP address that R6 received from R2:

```

R6#
*Aug 15 03:44:14.507: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0.26 assigned
DHCP address 148.49.26.1, mask 255.255.255.0, hostname VLAN26_R6

R6#show ip interface brief

```

| Interface      | IP-Address   | OK? | Method | Status                |
|----------------|--------------|-----|--------|-----------------------|
| Ethernet0/0    | unassigned   | YES | NVRAM  | up                    |
| Ethernet0/0.16 | 148.49.16.6  | YES | NVRAM  | up                    |
| Ethernet0/0.26 | 148.49.26.1  | YES | DHCP   | up                    |
| Ethernet0/0.36 | 148.49.36.6  | YES | NVRAM  | up                    |
| Ethernet0/0.56 | 148.49.56.6  | YES | NVRAM  | up                    |
| Ethernet0/1    | unassigned   | YES | NVRAM  | administratively down |
| Ethernet0/2    | unassigned   | YES | NVRAM  | administratively down |
| Ethernet0/3    | unassigned   | YES | NVRAM  | administratively down |
| Serial1/0      | unassigned   | YES | NVRAM  | administratively down |
| Serial1/1      | unassigned   | YES | NVRAM  | administratively down |
| Serial1/2      | unassigned   | YES | NVRAM  | administratively down |
| Serial1/3      | unassigned   | YES | NVRAM  | administratively down |
| Loopback106    | 148.49.106.1 | YES | NVRAM  | up                    |

```

R6#

```

**Likely cause: The DHCP server is incorrectly configured.**

The Layer 3 diagram of this scenario requires R6 to get the IP address 148.49.26.6 for interface E0/0.26.

Verify the configuration of the DHCP server on R2:

```

R2#sh run | section dhcp
ip dhcp pool VLAN26
  network 148.49.26.0 255.255.255.0
  default-router 148.49.26.6
ip dhcp pool SPECIFIC
  host 148.49.26.54 255.255.255.0
  client-name SPECIFIC
ip dhcp pool VLAN26_R6
  host 148.49.26.6 255.255.255.0

```

**Resolution: Configure the client identifier on R2.**

We need to obtain the client identifier of R6 to make sure that R2 attributes only the IP address 148.49.26.6 to interface E0/0.26 of R6:

```

R2#debug ip dhcp server packet
DHCP server packet debugging is on.

```

Release the DHCP address to R6:

```

R6#release dhcp E0/0.26
R6#

```

Return to R2:

```
R2#
*Aug 15 03:47:53.842: DHCPD: client's VPN is .
*Aug 15 03:47:53.842: DHCPD: No option 125
*Aug 15 03:47:53.842: DHCPD: DHCPRELEASE message received from client
0056.4c41.4e32.365f.5236 (148.49.26.1).
*Aug 15 03:47:53.842: DHCPD: removing ARP entry (148.49.26.1 vrf default).
R2#
*Aug 15 03:47:55.108: DHCPD: client's VPN is .
*Aug 15 03:47:55.108: DHCPD: No option 125
*Aug 15 03:47:55.108: DHCPD: DHCPRELEASE message received from client
0056.4c41.4e32.365f.5236 (148.49.26.1).
*Aug 15 03:47:55.108: DHCPD: Finding a relay for client 0056.4c41.4e32.365f.5236 on
interface Ethernet0/0.26.
R2#
*Aug 15 03:47:57.112: DHCPD: client's VPN is .
*Aug 15 03:47:57.112: DHCPD: No option 125
*Aug 15 03:47:57.112: DHCPD: DHCPRELEASE message received from client
0056.4c41.4e32.365f.5236 (148.49.26.1).
*Aug 15 03:47:57.112: DHCPD: Finding a relay for client 0056.4c41.4e32.365f.5236 on
interface Ethernet0/0.26.
R2#
```

Configure R2 to recognize requests sent by R6, using its client identifier:

```
R2(config)#ip dhcp pool VLAN26_R6
R2(dhcp-config)# host 148.49.26.6 255.255.255.0
R2(dhcp-config)# client-identifier 0056.4c41.4e32.365f.5236
```

Renew the DHCP address and verify which IP address is assigned to R6 now:

```
R6#renew dhcp E0/0.26
R6#
*Aug 15 03:52:36.136: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0.26 assigned
DHCP address 148.49.26.6, mask 255.255.255.0, hostname VLAN26_R6
R6#
```

## 9.2. Symptom: The IP address is incorrect for the default router advertised by the DHCP server.

### *Analysis and Testing:*

In this scenario, there may not be a symptom that would get our attention. Remember that the lab requires that R2 should announce R6 as the default gateway for workstations connected to VLAN 26.

Let us verify the configuration of the DHCP server:

```
R2#show run | section dhcp
ip dhcp pool VLAN26
 network 148.49.26.0 255.255.255.0
 default-router 148.49.26.2
ip dhcp pool SPECIFIC
 host 148.49.26.54 255.255.255.0
 client-name SPECIFIC
ip dhcp pool VLAN26_R6
 host 148.49.26.6 255.255.255.0
 client-identifier 0056.4c41.4e32.365f.5236
```

**Likely cause: DHCP pools are configured incorrectly.**

We can see that the IP address used for the default router in pool VLAN 26 is not 148.49.26.6 as required by this lab.

Additionally, the pool named SPECIFIC is missing an IP address for the default router.

**Resolution: Modify the configuration of the IP address for the default router.**

```
R2(config)#ip dhcp pool VLAN26
R2(dhcp-config)#default-router 148.49.26.6
R2(dhcp-config)#ip dhcp pool SPECIFIC
R2(dhcp-config)#default-router 148.49.26.6
```

Verify the final configuration for both DHCP pools VLAN 26 and SPECIFIC:

```
R2#show run | section dhcp

ip dhcp pool VLAN26
network 148.49.26.0 255.255.255.0
default-router 148.49.26.6
ip dhcp pool SPECIFIC
host 148.49.26.54 255.255.255.0
client-name SPECIFIC
default-router 148.49.26.6
ip dhcp pool VLAN26_R6
host 148.49.26.6 255.255.255.0
client-identifier 0056.4c41.4e32.365f.5236
```

---

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

---