360RSW05-WKB-CA05

Cisco 360 CCIE R&S Exercise Workbook Introduction

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

https://t.me/learningnets

Cisco 360 CCIE R&S Exercise Workbook Lab 5 Configuration Section Answer Key

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

https://t.me/learningnets

Table of Contents

https://t.me/learningnets

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.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Exercise Workbook Lab 5 Configuration 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

- Configuration lab duration: 6 hours
- Configuration maximum score: 76 points

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

Difficulty Level

- Difficulty: Intermediate

Restrictions and Goals

Note	Read this section carefully.

- To receive any credit for a subsection, you must fully complete the subsection as per the requirements. You will not receive partial credit for partially completed sections.
- IP subnets displayed in the scenario diagram belong to network 172.16.0.0/16.
- Do not configure any static routes. Do not modify any preconfigured static routes.
- Make sure that all IP version 4 (IPv4) and IP version 6 (IPv6) loopback interfaces are advertised with their original mask unless otherwise specified. The DVTI loopback interfaces are excluded from this requirent.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**).
- Make sure all IPv4 interfaces in the diagram are reachable within this internetwork. The 1.1.0.0/16 range is excluded from this requirement.
- IP subnets in the CustomerA VPN do not have to be reachable from outside the VPN.

https://t.me/learningnets

Explanation of Each of the Restrictions and Goals

IPv4 subnets that are displayed in the scenario diagram belong to network 172.16.0.0/16.

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

Do not configure any static routes.

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

Advertise loopback interfaces with their original masks.

The original mask is the mask that is configured on the loopback interface. Open Shortest Path First (OSPF), by default, treats loopback interfaces as host routes and advertises them as /32 prefixes. You need to provide a solution to represent the original mask of addresses assigned to loopback interfaces. A possible option is changing the OSPF network type or summarizations.

All the IP addresses that are involved in this scenario must be reachable, unless explicitly specified otherwise.

This goal is a key goal to observe. It requires that all of your IGPs and your routing policy tasks be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using the **distribute-list**, **route-map**, 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.

IP subnets in the CustomerA VPN do not have to be reachable from outside the VPN.

R7, R8, and R9 are in the CustomerA VPN. They are not required to reach IP addresses outside the VPN.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

1. Switch Configuration Section

General Tasks

As with any switch configuration, you must address the following basic configuration requirements:

- Set the VLAN Trunking Protocol (VTP) mode.
- Configure the VLANs and the VLAN names.
- Configure the trunk ports.
- Statically assign the ports of the switches to the VLANs.

Note	For a good reference on mastering basic Cisco Catalyst 3560 Switch configuration tasks, access the full set of Catalyst video-on-demand, or VoD, sessions within the "Link Layer" lesson in the Cisco 360 learning portal. These self-paced sessions provide more than 7 hours of instruction on a range of basic Catalyst switch configuration tasks. Some of the Cisco Catalyst 3560 Switch configuration commands are not available on the virtual instances of the switches.

Note that not all Cisco Catalyst 3560 Switch configuration features are supported on the virtual Cisco IOS Software on UNIX.

Use the "VLANs" tables, the "Switch-to-Router Connections" table, and the "Switch-to-Switch Connections" table for reference. It is recommended that you draw the VLAN distribution diagram for better visualization.

Carefully review the entire scenario. Closely examine the supplied diagram and any associated tables. Determine how you need to configure VTP, how to configure ports that are assigned as trunks, and how to configure ports that are assigned as simply static VLAN ports. For any ports that are statically assigned to a VLAN, it is recommended that you statically assign the **switchport mode access** command.

See the following diagram for the VLAN layout.

https://t.me/learningnets

VLAN Distribution

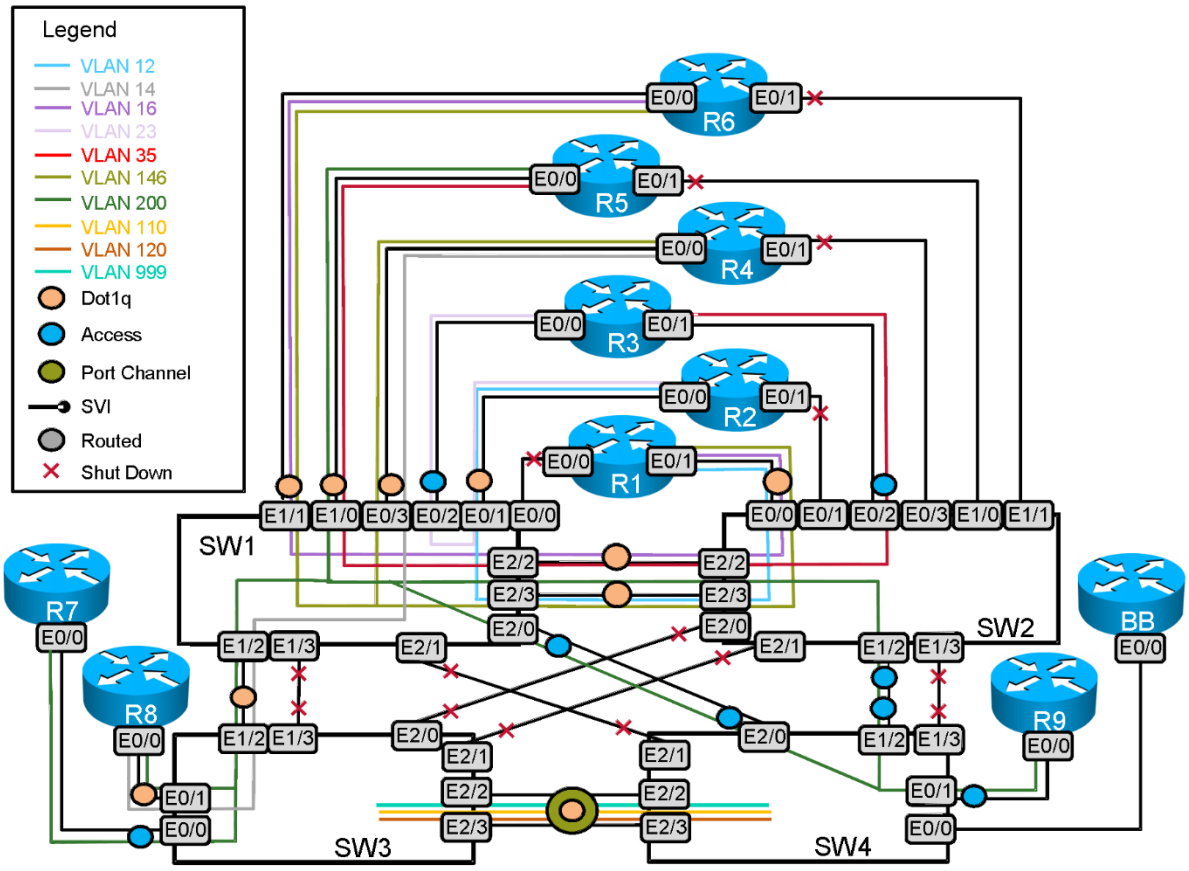

Issue: Configure the required VLANs on all switches according to the VLAN tables that are provided in the lab. VTP mode is transparent on all switches**.**

Solution:

You are told to use VTP transparent mode on all of your switches. Switches that are configured in VTP transparent mode will not advertise any VLANs that are created on them.
Configure the **vtp mode transparent** command and all necessary VLANs on SW1 and SW2 according to the scenario requirements and the "VLANs" table. The following example shows the configuration on SW1 and SW2:

```
SW1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#vtp mode transparent
Device mode already VTP Transparent for VLANS.
SW1(config)#vlan 12,14,16,23,35,146,200
SW1(config-vlan)#end
SW1#

SW2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#vtp mode transparent
Device mode already VTP Transparent for VLANS.
SW2(config)#vlan 12,14,16,23,35,146,200
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
SW2(config-vlan)#end
SW2#
```

Verify the VTP status on SW1 and SW2:

```
SW1#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 1
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.0700
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Feature VLAN:
--------------
VTP Operating Mode              : Transparent
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 12
Configuration Revision            : 0
MD5 digest                        : 0x82 0x06 0xC3 0xB4 0x18 0xC2 0x0F 0x0C
                                    0x47 0xE6 0xC7 0xD3 0x81 0x7E 0x38 0x1B
SW1#

SW2#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 1
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.0800
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Feature VLAN:
--------------
VTP Operating Mode              : Transparent
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 12
Configuration Revision            : 0
MD5 digest                        : 0x82 0x06 0xC3 0xB4 0x18 0xC2 0x0F 0x0C
                                    0x47 0xE6 0xC7 0xD3 0x81 0x7E 0x38 0x1B
SW2#
```

Configure the **vtp mode transparent** command and all necessary VLANs on SW3 and SW4 according to the scenario requirements and the “VLANs” table. The following example shows the configuration on SW3 and SW4:

```
SW3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW3(config)#vtp mode transparent
Device mode already VTP Transparent for VLANS.
SW3(config-vlan)#vlan 14,110,120,200,999
SW3(config-vlan)#end
SW3#

SW4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#vtp mode transparent
Device mode already VTP Transparent for VLANS.
SW4(config)#vlan 110,120,200,999
SW4(config-vlan)#end
SW4#
```

Verify the VTP status on SW3 and SW4:

https://t.me/learningnets

```
SW3#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 1
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.0900
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Feature VLAN:
--------------
VTP Operating Mode              : Transparent
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 9
Configuration Revision            : 0
MD5 digest                        : 0x10 0xC2 0x40 0x4F 0xD2 0x7D 0x5F 0x25
                                    0x66 0x44 0x5C 0xEA 0x74 0xF0 0x89 0x97
SW3#

SW4#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 1
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.0a00
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Feature VLAN:
--------------
VTP Operating Mode              : Transparent
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 9
Configuration Revision            : 0
MD5 digest                        : 0x10 0xC2 0x40 0x4F 0xD2 0x7D 0x5F 0x25
                                    0x66 0x44 0x5C 0xEA 0x74 0xF0 0x89 0x97
SW4#
```

After you complete the VLAN configuration on all switches, verify the VLANs on all switches. Your output should resemble the following example on SW1, SW2, SW3, and SW4:

```
SW1#show vlan brief | exclude ^100[2345]|^1* +

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
12   VLAN0012                         active
14   VLAN0014                         active
16   VLAN0016                         active
23   VLAN0023                         active
35   VLAN0035                         active
146  VLAN0146                         active
200  VLAN0200                         active

SW1#

SW2# show vlan brief | exclude ^100[2345]|^1* +

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
12   VLAN0012                         active
14   VLAN0014                         active
16   VLAN0016                         active
23   VLAN0023                         active
35   VLAN0035                         active
146  VLAN0146                         active
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
200  VLAN0200                         active

SW2#

SW3# show vlan brief | exclude ^100[2345]|^1* +

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
14   VLAN0014                         active
110  VLAN0110                         active
120  VLAN0120                         active
200  VLAN0200                         active
999  VLAN0999                         active

SW3#

SW4# show vlan brief | exclude ^100[2345]|^1* +

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
110  VLAN0110                         active
120  VLAN0120                         active
200  VLAN0200                         active
999  VLAN0999                         active

SW4#
```

Issue: Use the IEEE 802.1Q protocol for trunking. Allow only the necessary VLANs on the switch-to-router trunks.

Solution:

First, identify the ports on SW1, SW2, SW3, and SW4 that are connected to the routers that have multiple VLANs associated with them. See the tables in the lab scenario. Also, examine the VLAN distribution diagram. Interfaces with multiple VLANs must be configured as trunks. Consequently, the interfaces with only one VLAN will be configured as access VLAN ports.

Configure the switch-to-router connections according to the "Switch-to-Router Connections" table.

SW1:

```
interface Ethernet0/1
 description R2 port 0/0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 12,23
 switchport mode trunk
 duplex auto
!
interface Ethernet0/2
 description R3 port 0/0
 switchport access vlan 23
 switchport mode access
 duplex auto
!
interface Ethernet0/3
 description R4 port 0/0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 14,146
 switchport mode trunk
 duplex auto
!
interface Ethernet1/0
 description R5 port 0/0
 switchport trunk encapsulation dot1q
```


https://t.me/learningnets

```
 switchport trunk allowed vlan 35,200
 switchport mode trunk
 duplex auto
!
interface Ethernet1/1
 description R6 port 0/0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 16,146
 switchport mode trunk
 duplex auto
!
```

SW2:

```
interface Ethernet0/0
 description R1 port 0/1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 12,16,146
 switchport mode trunk
 duplex auto
!
interface Ethernet0/2
 description R3 port 0/1
 switchport access vlan 35
 switchport mode access
 duplex auto
!
```

SW3:

```
interface Ethernet0/0
 description R7 port 0/0
 switchport access vlan 200
 switchport mode access
 duplex auto
!
interface Ethernet0/1
 description R8 port 0/0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 14,200
 switchport mode trunk
 duplex auto
!
```

SW4:

```
interface Ethernet0/1
 description R9 port 0/0
 switchport access vlan 200
 switchport mode access
 duplex auto
!
```

Issue: Configure switch-to-switch links according to the table in the lab.

Read the lab from beginning to end. One of the later tasks requires the configuration of the EtherChannel on the links between SW3 and SW4.

Configure the switch-to-switch connections according to the “Switch-to-Switch Connections” table.

SW1:

```
interface Ethernet1/2
 description SW3 port 1/2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
!
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
interface Ethernet2/0
 description SW4 port 2/0
 switchport access vlan 200
 switchport mode access
 duplex auto
!
interface Ethernet2/2
 description SW2 port 2/2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
!
interface Ethernet2/3
 description SW2 port 2/3
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
!
```

Verify that the following interfaces are administratively shut down:

```
SW1#show ip interface brief | inc Et.*admin
…
Ethernet0/0            unassigned      YES unset  administratively down down
Ethernet1/3            unassigned      YES unset  administratively down down
Ethernet2/1            unassigned      YES unset  administratively down down
…
SW1#
```

SW2:

```
!
interface Ethernet1/2
 description SW4 port 1/2
 switchport access vlan 200
 switchport mode access
 duplex auto
!
interface Ethernet2/2
 description SW1 port 2/2
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
!
interface Ethernet2/3
 description SW1 port 2/3
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
!
```

Verify that the following interfaces are administratively shut down:

```
SW2#show ip interface brief | inc Et.*admin
…
Ethernet1/3            unassigned      YES unset  administratively down down
Ethernet2/0            unassigned      YES unset  administratively down down
Ethernet2/1            unassigned      YES unset  administratively down down
…
SW2#
```

SW3:

```
!
interface Ethernet1/2
 description SW1 port 1/2
 switchport trunk encapsulation dot1q
```


https://t.me/learningnets

```
 switchport mode trunk
 duplex auto
!
```

Verify that the following interfaces are administratively shut down:

```
SW3#show ip interface brief | inc Et.*admin
…
Ethernet1/3            unassigned      YES unset  administratively down down
Ethernet2/0            unassigned      YES unset  administratively down down
Ethernet2/1            unassigned      YES unset  administratively down down
…
SW3#
```

SW4:

```
interface Ethernet1/2
 description SW2 port 1/2
 switchport access vlan 200
 switchport mode access
 duplex auto
!
interface Ethernet2/0
 description SW1 port 2/0
 switchport access vlan 200
 switchport mode access
 duplex auto
!
```

Verify that the following interfaces are administratively shut down:

```
SW3#show ip interface brief | inc Et.*admin
…
Ethernet1/3            unassigned      YES unset  administratively down down
Ethernet2/1            unassigned      YES unset  administratively down down
…
SW3#
```

Issue: Configure VLAN 200 bridge priority 4096 on SW4, configure bridge priority 8192 on SW2, and leave the VLAN 200 bridge priority as the default on SW1. Make sure that forwarding on VLAN 200 is done between SW1 and SW2 and between SW1 and SW4.

Solution:

Configure the VLAN 200 spanning-tree bridge priorities according to the requirements of the lab and verify the status of the interfaces on SW1:

```
SW2#show running-config | inc spanning-tree vlan 200
spanning-tree vlan 200 priority 8192
SW2#

SW4#show running-config | inc spanning-tree vlan 200
spanning-tree vlan 200 priority 4096
SW4#

SW1#show spanning-tree vlan 200

VLAN0200
  Spanning tree enabled protocol ieee
  Root ID    Priority    4296
             Address     aabb.cc00.0a00
             Cost        100
             Port        9 (Ethernet2/0)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32968  (priority 32768 sys-id-ext 200)
             Address     aabb.cc00.0700
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et1/0               Desg FWD 100       128.5    Shr
Et1/2               Desg FWD 100       128.7    Shr
Et2/0               Root FWD 100       128.9    Shr
Et2/2               Altn BLK 100       128.11   Shr
Et2/3               Altn BLK 100       128.12   Shr


SW1#
```

When you configure all specified bridge priorities, the spanning tree will block ports 2/2 and 2/3 on SW1 because the SW1 priority for VLAN 200 is the least desirable due to its higher value:

Spanning-Tree Configuration Diagram Initial State

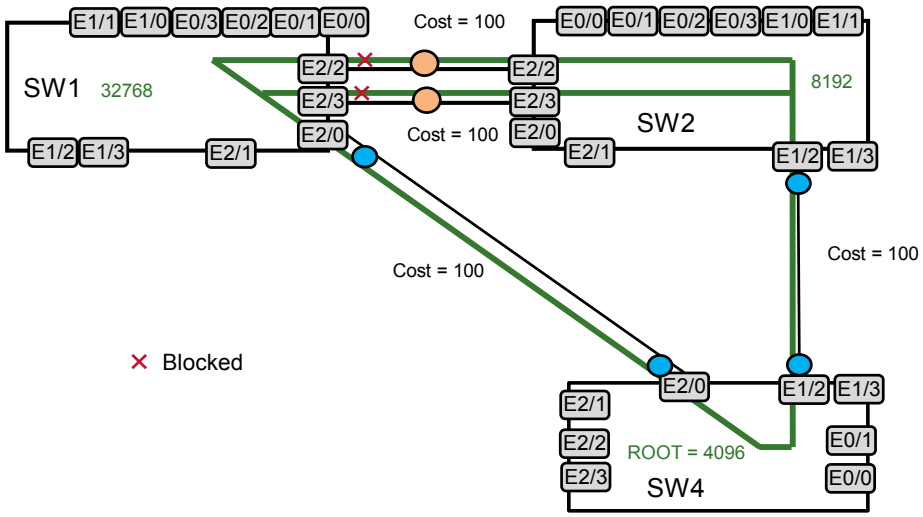

The default spanning-tree calculation will block forwarding between SW1 and SW2, which violates the requirement of the scenario. We need to move the blocking interfaces to the link between SW2 and SW4. Spanning-tree cost manipulation can help us achieve this goal. Cost is evaluated before the bridge priority in the spanning-tree calculation. Therefore, if we set the cost of the link between SW2 and SW4 to a value higher than the accumulative cost from SW2 to SW4 via SW1, we can block the undesired direct link between SW2 and SW4 (cost 201 will be applied to the 1/2 port of SW2).

https://t.me/learningnets

Spanning-Tree Configuration Diagram Final State

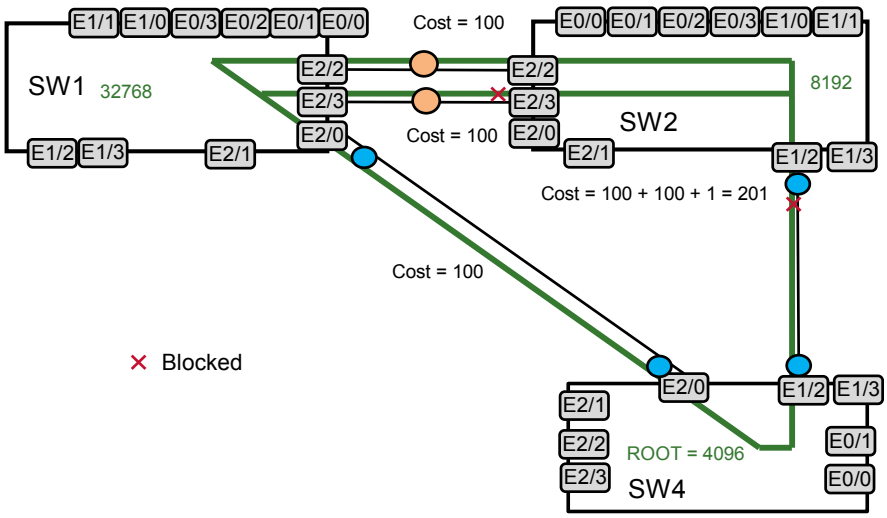

Configure the spanning-tree cost 201 on the Ethernet1/2 interface of SW2:

```
interface Ethernet1/2
  spanning-tree vlan 200 cost 201
!
```

Verify the VLAN 200 spanning tree on SW2:

```
SW2#show spanning-tree vlan 200

VLAN0200
  Spanning tree enabled protocol ieee
  Root ID     Priority    4296
              Address     aabb.cc00.0a00
              Cost        200
              Port        11 (Ethernet2/2)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    8392   (priority 8192 sys-id-ext 200)
              Address     aabb.cc00.0800
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et1/2               Altn BLK 201       128.7    Shr
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
Et2/2               Root FWD 100       128.11   Shr
Et2/3               Altn BLK 100       128.12   Shr


SW2#
```

Note that the Ethernet1/2 interface is blocked. The VLAN 200 traffic will not be forwarded on the link between SW2 and SW4. Also the spanning tree blocked the Ethernet2/3 interface on SW2 to prevent a loop on the trunk links between SW1 and SW2.

Issue: Configure an aggregated trunk link with a bandwidth of 20 Mb/s between SW3 and SW4. Use Link Aggregation Control Protocol, or LACP, and SW3 should initiate EtherChannel initialization procedures. Configure the dot1q trunk on the EtherChannel link. Allow only VLAN 110, VLAN 120, and VLAN 999 on the trunk between SW3 and SW4.

Solution:

EtherChannel can be configured to always begin negotiations using the Port Aggregation Protocol, or PAgP, and LACP. LACP operates in active and passive modes, where active mode always initiates LACP negotiations and passive mode waits for an active device on the other end.

SW3 must be configured with active mode and SW4 with passive mode.

Create an EtherChannel interface by assigning a channel group to interfaces 2/2 through 2/3:

SW3:

```
interface range Ethernet2/2-3
channel-group 10 mode active
!
```

SW4:

```
interface range Ethernet2/2-3
channel-group 10 mode passive
!
```

Configure the dot1q trunk on the EtherChannel interface:

SW3:

```
interface Port-channel10
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 110,120,999
 switchport mode trunk
!
```

SW4:

```
interface Port-channel10
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 110,120,999
 switchport mode trunk
!
```

These commands will be copied to ports 2/2 through 2/3 automatically.

https://t.me/learningnets

Verify that EtherChannel is up and trunking:

```
SW3#show etherchannel 10 port-channel
                Port-channels in the group:
                ---------------------------

Port-channel: Po10    (Primary Aggregator)

------------

Age of the Port-channel   = 0d:01h:01m:25s
Logical slot/port   = 16/0          Number of ports = 2
HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            =   LACP
Port security       = Disabled

Ports in the Port-channel:

Index   Load   Port     EC state        No of bits
------+------+------+------------------+-----------
  0     00     Et2/2    Active             0
  0     00     Et2/3    Active             0

Time since last port bundled:    0d:00h:56m:02s    Et2/3
Time since last port Un-bundled: 0d:00h:56m:11s    Et2/3

SW3#

SW3#show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/1       on               802.1q         trunking      1
Et1/2       on               802.1q         trunking      1
Po10        on               802.1q         trunking      1

Port        Vlans allowed on trunk
Et0/1       14,200
Et1/2       14,200
Po10        110,120,999

Port        Vlans allowed and active in management domain
Et0/1       200
Et1/2       1,110,120,200,999
Po10        110,120,999

Port        Vlans in spanning tree forwarding state and not pruned
Et0/1       200
Et1/2       1,110,120,200,999
Po10        110,120,999
SW3#
```

Issue: Verify that the packets coming from the same IP source to multiple IP destinations are distributed across the ports of the EtherChannel group, but the packets from different IP sources to the same destination are forwarded via the same switch port in the group.

Solution:

With destination IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the destination IP address of the incoming packet. Therefore, to provide load balancing, packets from the same IP source address sent to different IP destination addresses could be sent on different ports in the channel. But packets sent from different source IP addresses to the same destination IP address are always sent on the same port in the channel.
Configure **port-channel load-balance dst-ip** on both switches SW3 and SW4.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

SW3:

```
port-channel load-balance dst-ip
!
```

SW4:

```
port-channel load-balance dst-ip
!
```

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

https://t.me/learningnets

2. Secure DVTI Communications Section

Verify the IP connectivity between the Serial1/0 interfaces of R1, R2, and R4. The following example shows the connectivity test from R1:

```
R1#ping 1.1.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/16/17 ms
R1#ping 1.1.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/16/17 ms
R1#
```

Note that the pings between R1 and R2, and between R1 and R4 are successful.

Configure the static VTI Tunnel241 interfaces on spokes R1 and R4 according to the "Configure VTI Tunnel Interfaces" section requirements and the "IPv4 IGP" diagram:
R1:

```
interface Tunnel241
 ip unnumbered Loopback241
 tunnel source Serial1/0
 tunnel mode ipsec ipv4
 tunnel destination 1.1.2.2
!
```

R4:

```
interface Tunnel241
 ip unnumbered Loopback241
 tunnel source Serial1/0
 tunnel mode ipsec ipv4
 tunnel destination 1.1.2.2
!
```

Configure the dynamic VTI Virtual-Template124 interface on hub R2 according to the "Configure VTI Tunnel Interfaces" section requirements and the "IPv4 IGP" diagram:
R2:

```
interface Virtual-Template124 type tunnel
 ip unnumbered Loopback241
 tunnel mode ipsec ipv4
```

Configure IPsec security on the DVTI subnet 172.16.241.0/24 according to the "Configure IPsec DVTI Communications Between R1, R2, and R4" section requirements.
Configure IKE Phase 1 and Phase 2 and the IPsec crypto profile configuration on R1, R2, and R4.

R1:

```
crypto isakmp policy 10
 encr 3des
 hash md5
 authentication pre-share
crypto isakmp key sharedpswd address 0.0.0.0
crypto isakmp diagnose error
!
!
crypto ipsec transform-set vti_transform esp-3des esp-md5-hmac
 mode tunnel
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
!
!
crypto ipsec profile vti_profile
 set transform-set vti_transform
!
```

R2:

```
crypto isakmp policy 10
 encr 3des
 hash md5
 authentication pre-share
crypto isakmp key sharedpswd address 0.0.0.0
crypto isakmp profile vti_isakmp_profile
   keyring default
   match identity address 0.0.0.0
   virtual-template 124
!
!
crypto ipsec transform-set vti_transform esp-3des esp-md5-hmac
 mode tunnel
!
!
crypto ipsec profile vti_profile
 set transform-set vti_transform
!
```

R4:

```
crypto isakmp policy 10
 encr 3des
 hash md5
 authentication pre-share
crypto isakmp key sharedpswd address 0.0.0.0
crypto isakmp diagnose error
!
!
crypto ipsec transform-set vti_transform esp-3des esp-md5-hmac
 mode tunnel
!
!
crypto ipsec profile vti_profile
 set transform-set vti_transform
!
```

Apply the IPsec profile on the Virtual-Tunnel124 interface on R2, and on the Tunnel241 interfaces on R1 and R4:

R2:

```
interface Virtual-Template124 type tunnel
tunnel protection ipsec profile vti_profile
!
```

R1:

```
interface Tunnel241
tunnel protection ipsec profile vti_profile
!
```

R4:

```
interface Tunnel241
tunnel protection ipsec profile vti_profile
!
```

Verify the interfaces on R2:

```
R2#show ip interface brief | inc Loop|Vi
Loopback102                172.16.102.1    YES manual up                    up
Loopback241                172.16.241.2    YES manual up                    up
Virtual-Access1            172.16.241.2    YES unset  up                    up
Virtual-Access2            172.16.241.2    YES unset  up                    up
```


https://t.me/learningnets

```
Virtual-Template124                    172.16.241.2          YES   unset      up
down
R2#
```

Note that two virtual access interfaces Virtual-Access1 and Virtual-Access2 are cloned from the Virtual-Template124 interface and assigned with the IP address 172.16.241.2, which is the same as the one on the Loopback241 interface that was used for the unnumbered address on the Virtual-Template124.

Verify the ISAKMP SAs on R2:

```
R2#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
1.1.2.2         1.1.4.4         QM_IDLE           1002 ACTIVE
1.1.2.2         1.1.1.1         QM_IDLE           1001 ACTIVE

IPv6 Crypto ISAKMP SA

R2#
```

Note that R2 shows active ISAKMP SAs with R1 and R4.

Verify the IPsec SAs on the Virtual-Access1 on R2:

```
R2#show crypto ipsec sa interface vi1 | inc
interface|encrypt|decrypt|Status|peer
interface: Virtual-Access1
   current_peer 1.1.1.1 port 500
    #pkts encaps: 0, #pkts encrypt: 0, #pkts digest: 0
    #pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0
        Status: ACTIVE(ACTIVE)
        Status: ACTIVE(ACTIVE)
R2#
```

Note that R2 shows active IPsec SAs with the peer 1.1.1.1 (R1). The traffic counters are currently displaying 0 because there is no interesting IPsec traffic between R1 and R2 at this moment.
Verify the IPsec SAs on the Virtual-Access2 on R2:

```
R2#show crypto ipsec sa interface vi2 | inc
interface|encrypt|decrypt|Status|peer
interface: Virtual-Access2
   current_peer 1.1.4.4 port 500
    #pkts encaps: 0, #pkts encrypt: 0, #pkts digest: 0
    #pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0
        Status: ACTIVE(ACTIVE)
        Status: ACTIVE(ACTIVE)
R2#
```

Note that R2 shows active IPsec SAs with the peer 1.1.4.4 (R4). The traffic counters are currently displaying 0 because there is no interesting IPsec traffic between R2 and R4 at this moment.

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

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

3. IPv4 OSPF Section

Note	All OSPF routers must be configured with only one OSPF process ID, or PID. Use your IGP diagram to help guide your configuration.

Issue: Elect a designated router, or DR, on R3 for all subnets in Area 0. Configure the network type that uses multicast to transmit the OSPF packets.

Solution:

The default network type for a serial link is point-to-point. If R3 is going to be the DR, configure the OSPF network type broadcast on the serial interfaces that are connected to the 172.16.43.0/24 subnet that is in the OSPF Area 0. The broadcast network type uses the 224.0.0.5 and 224.0.0.6 multicast groups for the OSPF packet destination.

Issue the **ip ospf priority 0** command under the interface configuration for Serial1/1 on router R4 to ensure that R3 is always elected as the DR for the 172.16.43.0/24 subnet.

Issue: Do not elect a DR on the 172.16.23.0/24 subnet.

Solution:

If you do not want to elect a DR on a subnet, make sure the OSPF network type is neither broadcast nor nonbroadcast. Since the 172.16.23.0/24 subnet is Ethernet, the default OSPF network type is broadcast. It must be changed to one of the following: point-to-point, point-to-multipoint, or point-to-multipoint nonbroadcast. The point-to-multipoint network type is chosen in this answer key.

Issue the **ip ospf network point-to-multipoint** command for interfaces that are connected to VLAN 23

Issue: Do not allow external routing information into Areas 23 and 25 on the redistribution from EIGRP to OSPF. However, allow that information to transit through Areas 23 and 25 when such external information is originated by the routers within the areas.

Solution:

Areas 23 and 25 must be configured as not-so-stubby-areas (NSSAs) to support external OSPF information to transit through the areas. The redistribution section requires R4 to perform the route redistribution from EIGRP to OSPF. You have to configure the **no-redistribution** keyword in the NSSA configuration on R4 to prevent the EIGRP routes from being redistributed into OSPF.

Issue: Make sure that R2 can still forward packets to all routers. Use the R4 172.16.241.4 address as the next hop of last resort.

Solution:

Configure the **default-information-originate** keyword in the NSSA configuration on R4.

Configure OSPF on R1, R2, R3, R4, and R5 according to the scenario requirements:

https://t.me/learningnets

R1:

```
router ospf 100
 area 25 nssa
 network 172.16.12.0 0.0.0.255 area 25
 network 172.16.241.0 0.0.0.255 area 25
```

R2:

```
!
interface Ethernet0/0.23
 encapsulation dot1Q 23
 ip address 172.16.23.2 255.255.252.0
 ip ospf network point-to-multipoint
!
router ospf 100
 area 23 nssa
 area 25 nssa
 network 172.16.12.0 0.0.0.255 area 25
 network 172.16.20.0 0.0.3.255 area 23
 network 172.16.102.0 0.0.0.255 area 25
 network 172.16.241.0 0.0.0.255 area 25
!
```

R3:

```
interface Loopback103
 ip address 172.16.103.1 255.255.255.0
 ip ospf network point-to-point
!
interface Serial1/0
 ip address 172.16.43.3 255.255.255.0
 ip ospf network broadcast
!
interface Ethernet0/0
 ip address 172.16.23.3 255.255.252.0
 ip ospf network point-to-multipoint
!
router ospf 100
 area 23 nssa
 network 172.16.20.0 0.0.3.255 area 23
 network 172.16.43.0 0.0.0.255 area 0
 network 172.16.53.0 0.0.0.255 area 0
 network 172.16.103.0 0.0.0.255 area 103
!
!
```

R4:

```
interface Loopback104
 ip address 172.16.104.1 255.255.255.0
 ip ospf network point-to-point
!
interface Serial1/1
 ip address 172.16.43.4 255.255.255.0
 ip ospf network broadcast
 ip ospf priority 0
!
router ospf 100
 area 25 nssa no-redistribution default-information-originate
 network 172.16.43.0 0.0.0.255 area 0
 network 172.16.104.0 0.0.0.255 area 0
 network 172.16.241.0 0.0.0.255 area 25
!
!
```

R5:

```
interface Loopback105
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
 ip address 172.16.105.1 255.255.255.0
 ip ospf network point-to-point
!
router ospf 100
 network 172.16.53.0 0.0.0.255 area 0
 network 172.16.105.0 0.0.0.255 area 0
!
```

Verify the OSPF neighbors on R2:

```
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.103.1      0   FULL/  -        00:01:44    172.16.23.3     Ethernet0/0.23
172.16.241.1      0   FULL/  -        00:00:38    172.16.241.1    Virtual-Access2
172.16.241.4      0   FULL/  -        00:00:39    172.16.241.4    Virtual-Access1
172.16.241.1      1   FULL/BDR        00:00:35    172.16.12.1     Ethernet0/0.12
R2#
```

Note that R2 forms all required OSPF neighbor relationships with R1, R3, and R4.

Verify the OSPF neighbors on R4:

```
R4#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.103.1      1   FULL/DR         00:00:36    172.16.43.3     Serial1/1
172.16.241.2      0   FULL/  -        00:00:33    172.16.241.2    Tunnel241
R4#
```

Note that R4 shows R3 as the OSPF DR on the subnet 172.16.43.0/24 that is connected to the Serial1/1 interface.

Verify the OSPF routing table on R3:

```
R3#show ip route ospf | inc ^O
O IA     172.16.12.0/24 [110/1074] via 172.16.43.4, 01:04:32, Serial1/0
O        172.16.23.2/32 [110/10] via 172.16.23.2, 01:04:42, Ethernet0/0
O IA     172.16.102.0/24 [110/1065] via 172.16.43.4, 01:04:32, Serial1/0
O        172.16.104.0/24 [110/65] via 172.16.43.4, 01:04:32, Serial1/0
O        172.16.105.0/24 [110/11] via 172.16.53.5, 00:39:40, Ethernet0/1
O IA     172.16.241.1/32 [110/1075] via 172.16.43.4, 01:04:31, Serial1/0
O IA     172.16.241.2/32 [110/1065] via 172.16.43.4, 01:04:32, Serial1/0
O IA     172.16.241.4/32 [110/65] via 172.16.43.4, 01:04:32, Serial1/0
R3#
```

Verify the OSPF routing table on R4:

```
R4#show ip route ospf | inc ^O
O        172.16.12.0/24 [110/1010] via 172.16.241.2, 01:06:42, Tunnel241
O IA     172.16.23.2/32 [110/74] via 172.16.43.3, 01:06:12, Serial1/1
O IA     172.16.23.3/32 [110/64] via 172.16.43.3, 01:06:12, Serial1/1
O        172.16.53.0/24 [110/74] via 172.16.43.3, 00:41:16, Serial1/1
O        172.16.102.0/24 [110/1001] via 172.16.241.2, 01:06:42, Tunnel241
O IA     172.16.103.0/24 [110/65] via 172.16.43.3, 01:06:12, Serial1/1
O        172.16.105.0/24 [110/75] via 172.16.43.3, 00:41:06, Serial1/1
O        172.16.241.1/32 [110/1011] via 172.16.241.2, 01:06:07, Tunnel241
O        172.16.241.2/32 [110/1001] via 172.16.241.2, 01:06:42, Tunnel241
R4#
```

Note that the OSPF intra-area routes are learned via the Tunnel241 interface that is connected to the secure DVTI network 172.16.241.0/24.

https://t.me/learningnets

Verify the OSPF connectivity. Here is an example from R4:

```
R4#ping 172.16.241.1 source loopback241
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.241.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.241.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 21/21/22 ms
R4#ping 172.16.102.1 source loopback241
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.102.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.241.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/21/22 ms
R4#ping 172.16.103.1 source loopback241
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.103.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.241.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/8/9 ms
R4#ping 172.16.105.1 source loopback241
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.105.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.241.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
R4#
```

Note that R4 can successfully ping interfaces on R1, R2, R3, and R5.

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

4. IPv4 EIGRP Section

Issue: Configure EIGRP AS 100 on the subnet between R1, R6, and R4. Configure EIGRP AS 10 on the VLAN 16 links between R1 and R6. Advertise the loopback interfaces 172.16.101.0/24 and 172.16.106.0/24 into EIGRP AS 10.

Solution:

By default, two EIGRP processes running on the same router using different AS numbers will not automatically exchange routes. Therefore, manual redistribution must be performed on either R1 or R6 (or both routers). The recommended procedure is to perform redistribution between the two EIGRP autonomous systems on R6.

Issue: Authenticate EIGRP adjacencies for a period from January 1, 2013, to January 1, 2030.

Solution:

Configure EIGRP Message Digest 5, or MD5, authentication on the routers in EIGRP AS 100. Apply the authentication configuration to the Ethernet interfaces that are attached to VLAN 146. In order to apply the authentication for the period of time specified, enter the specified starting and ending times and dates with the **accept-lifetime** and **send-lifetime** commands.

Configure EIGRP on R1, R4, and R6 according to the scenario requirements:

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

R1:

```
key chain eigrp-100
 key 1
  key-string cisco
  accept-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2030
  send-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2030
!
interface Ethernet0/1.146
 encapsulation dot1Q 146
 ip address 172.16.146.1 255.255.248.0
  ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-100
!

router eigrp 100
 network 172.16.144.0 0.0.7.255
 auto-summary
!
!
router eigrp 10
 network 172.16.61.0 0.0.0.255
 network 172.16.101.0 0.0.0.255
!
```

R4:

```
!
key chain eigrp-100
 key 1
  key-string cisco
  accept-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2030
  send-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2030
!
interface Ethernet0/0.146
 encapsulation dot1Q 146
 ip address 172.16.146.4 255.255.248.0
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-100
!
router eigrp 100
 network 172.16.144.0 0.0.7.255
!
```

R6:

```
key chain eigrp-100
 key 1
  key-string cisco
  accept-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2030
  send-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2030
!
interface Ethernet0/0.146
 encapsulation dot1Q 146
 ip address 172.16.146.6 255.255.248.0
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-100
 !
router eigrp 100
 network 172.16.144.0 0.0.7.255
 redistribute eigrp 10
!
!
router eigrp 10
 network 172.16.61.0 0.0.0.255
 network 172.16.106.0 0.0.0.255
!
```


https://t.me/learningnets

Verify the EIGRP neighbors on R6:

```
R6#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   172.16.146.4            Et0/0.146                11 02:22:39  529  3174  0  9
0   172.16.146.1            Et0/0.146                10 02:22:39  528  3168  0  5
EIGRP-IPv4 Neighbors for AS(10)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   172.16.61.1             Et0/0.16                 10 02:21:47   11   100  0  3
R6#
```

Note that R6 forms the EIGRP neighbor relationships with R1 and R4 in AS 100.
Verify the EIGRP routing table on R4:

```
R4#show ip route eigrp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 21 subnets, 4 masks
D EX     172.16.61.0/24
           [170/307200] via 172.16.146.6, 13:18:36, Ethernet0/0.146
D EX     172.16.101.0/24
           [170/435200] via 172.16.146.6, 13:18:36, Ethernet0/0.146
D EX     172.16.106.0/24
           [170/409600] via 172.16.146.6, 13:18:35, Ethernet0/0.146
R4#
```

Note that R4 learns the EIGRP AS 100 prefixes as external.

Verify the EIGRP connectivity. Here is an example from R4:

```
R4#ping 172.16.101.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.101.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R4#ping 172.16.106.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.106.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R4#ping 172.16.61.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.61.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R4#ping 172.16.61.6
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.61.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R4#
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Note that R4 can successfully ping interfaces on R1 and R6.

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

5. Redistribution Section

Issue: Mutually redistribute EIGRP and OSPF on R4.

Solution:

Configure the route redistribution on R4 according to the scenario requirements.

R4:

```
router eigrp 100
 redistribute ospf 100 metric 10000 100 255 1 1500
!
router ospf 100
 redistribute eigrp 100 metric 1 subnets
```

You can use the following Tcl script to test universal reachability. To use the script, enter the command **tclsh** in privileged mode, and paste in the script. To kill failing pings, hold down **Ctrl-Shift** and press the **6** key twice. When you are finished, enter **tclquit** to leave Tcl mode.

Note	Tcl connectivity verification scripts for each router are available via the Verification link in the CIERSWB service tab on the web portal.

```
tclsh
foreach address {

172.16.146.1
172.16.241.1
172.16.61.1
172.16.12.1
172.16.101.1

172.16.241.2
172.16.23.2
172.16.12.2
172.16.102.1

172.16.53.3
172.16.43.3
172.16.23.3
172.16.103.1

172.16.146.4
172.16.241.4
172.16.43.4
172.16.104.1

172.16.53.5
```


https://t.me/learningnets

```
172.16.105.1

172.16.146.6
172.16.61.6
172.16.106.1


} {ping $address}
tclquit
```

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

6. BGP Section

Issue: Read the Border Gateway Protocol (BGP) peering requirements and create a diagram of peer relationships similar to the diagram below:

BGP Diagram

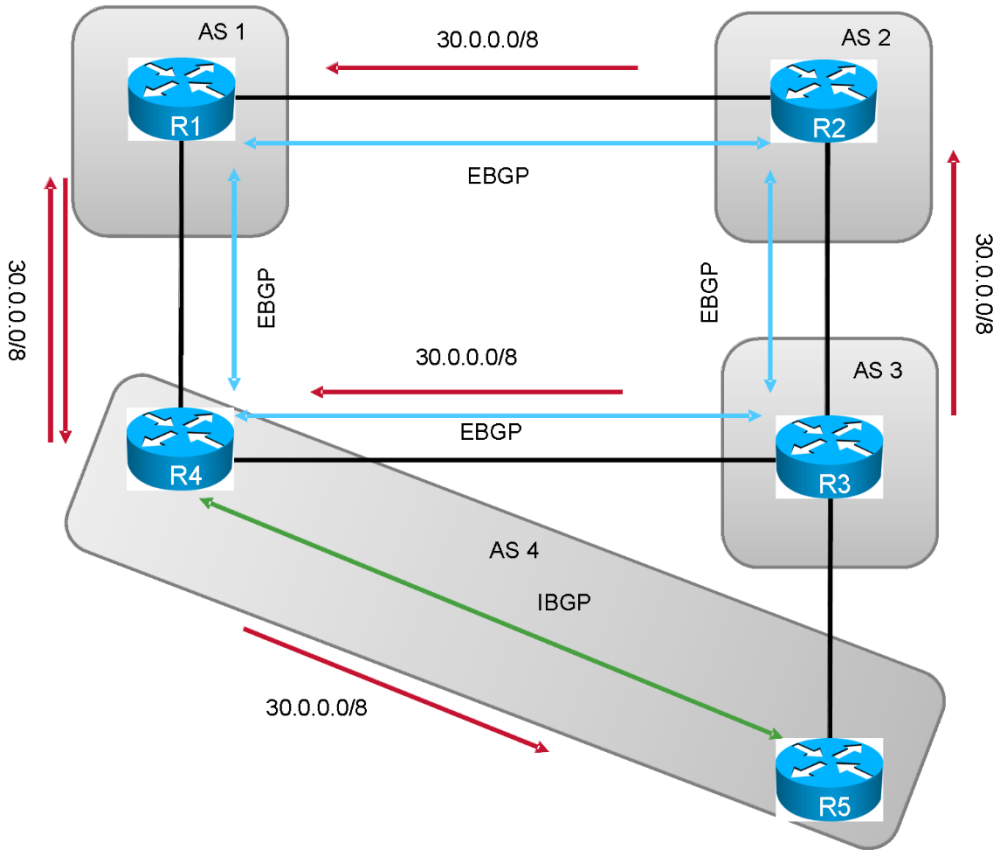

Solution:

These fairly straightforward peering and advertising directions set the stage for the remaining BGP tasks. For configuration details, see the online Mentor Guide engine.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Verify peering with the **show ip bgp summary** command. The number of prefixes learned, at the end of each line, indicates a good peering. Entries such as "Active" or "Idle" would indicate peering problems.

Issue: On R3, advertise the Loopback 30 subnet into BGP as a classful prefix. Do not use aggregation.

Solution:

To add the prefix 30.0.0.0/8 into BGP without using aggregation or a static route, enable auto summarization under BGP and enter the classful network statement on R3.

Issue: On router R2, configure BGP to track next-hop changes and report these changes directly to BGP. Increase the default delay for next-hop changes by 1 second.

Solution:

The BGP Support for Next-Hop Address Tracking feature is enabled by default when a supporting Cisco IOS Software image is installed. BGP next-hop address tracking is event-driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to the BGP routing process as they are updated in the Routing Information Base (RIB). This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes that are installed in the RIB. When a best-path calculation is run between BGP scanner cycles, only next-hop changes are tracked and processed.

The trigger is enabled by default, so adjust the delay (the default is 5 seconds):

```
R2(config-router)#bgp nexthop trigger delay 6
```

Configure BGP on R1, R2, R3, R4, and R5 according to the scenario requirements.

R1:

```
router bgp 1
 bgp log-neighbor-changes
 neighbor 172.16.12.2 remote-as 2
 neighbor 172.16.146.4 remote-as 4
!
```

R2:

```
router bgp 2
 bgp log-neighbor-changes
 bgp nexthop trigger delay 6
 neighbor 172.16.12.1 remote-as 1
 neighbor 172.16.23.3 remote-as 3
!
```

R3:

```
router bgp 3
 bgp log-neighbor-changes
 network 30.0.0.0
 neighbor 172.16.23.2 remote-as 2
 neighbor 172.16.43.4 remote-as 4
 auto-summary
!
```


https://t.me/learningnets

R4:

```
router bgp 4
 bgp log-neighbor-changes
 neighbor 172.16.43.3 remote-as 3
 neighbor 172.16.105.1 remote-as 4
 neighbor 172.16.105.1 update-source Loopback104
 neighbor 172.16.146.1 remote-as 1
!
!
```

R5:

```
router bgp 4
 bgp log-neighbor-changes
 neighbor 172.16.104.1 remote-as 4
 neighbor 172.16.104.1 update-source Loopback105
!
!
```

Verify the BGP peer relationships on the routers. Here is an example on R4:

```
R4#show ip bgp summary
BGP router identifier 172.16.241.4, local AS number 4
BGP table version is 2, main routing table version 2
1 network entries using 140 bytes of memory
1 path entries using 76 bytes of memory
1/1 BGP path/bestpath attribute entries using 140 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 428 total bytes of memory
BGP activity 7/0 prefixes, 9/2 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.43.3     4            3      94      95        2    0    0 01:22:40        1
172.16.105.1    4            4      68      71        2    0    0 00:56:50        0
172.16.146.1    4            1      95      93        2    0    0 01:22:21        0
R4#
```

Note that R4 forms BGP peer relationships with R1, R3, and R5.
Verify the BGP prefixes on R4:

```
R4#show ip bgp
BGP table version is 2, local router ID is 172.16.241.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  30.0.0.0         172.16.43.3              0             0 3 i
R4#
```

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

7. MPLS Layer 3 VPN Section

Issue: Configure a VRF called CustomerA on R4 and R5. Use the route distinguisher value 4:100. Assign interface E0/0.14 on R4 and interface E0/0.200 on R5 to this VRF.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Solution:

Import and export route targets are not specified, so you will have to choose values that meet the overall requirements. A simple and common approach is to use the route distinguisher value for both route targets. Here is our configuration for R4 and R5:

```
ip vrf CustomerA
 rd 4:100
 route-target export 4:100
 route-target import 4:100
```

You then associate the required interfaces with the VRF using the command **ip vrf forwarding** *CustomerA*.

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#int e0/0.14
R4(config-subif)#ip vrf forwarding CustomerA
%  Interface  Ethernet0/0.14  IPv4  disabled  and  address(es)  removed  due  to
disabling VRF CustomerA
R4(config-subif)# ip address 172.16.14.4 255.255.255.0
R4(config-subif)#end
R4#
R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R5(config)#int e0/0.200
R5(config-subif)# ip vrf forwarding CustomerA
%  Interface  Ethernet0/0.200  IPv4  disabled  and  address(es)  removed  due  to
disabling VRF CustomerA
R5(config-subif)# ip address 172.16.25.5 255.255.255.0
R5(config-subif)#end
R5#
```

Note that adding the interface to the VRF removes it from the default VRF; it will no longer be listed when you enter **show ip route connected**. Remember to add back the original IP address.

Verify your configuration with the command **show ip vrf**:

```
R4#show ip vrf
  Name                              Default RD          Interfaces
  CustomerA                         4:100               Et0/0.14
R4#

R5#show ip vrf
  Name                              Default RD          Interfaces
  CustomerA                         4:100               Et0/0.200
R5#
```

Issue: Enable the exchange of vpnv4 addresses between R4 and R5. Enable the Label Distribution Protocol (LDP) between R3 and R4 and between R3 and R5.

Solution:

R4 and R5 created an IBGP peering in AS 4 in an earlier section. Enter the **neighbor activate** command under the VPNv4 address family on both R4 and R5. The neighbor **send-community extended** command will be added automatically. Verify the peering as you see here:

Enter the command **mpls ip** on the interfaces that are connecting the specified routers. Verify LDP peering with the command **show mpls ldp neighbor**: R3 should see both R4 and R5 as LDP neighbors. Remember that the LDP router ID must be a reachable IP address in the default VRF.

https://t.me/learningnets

Issue: Configure R7 and R8 as peers in BGP AS 100. Peer R8 with R4. Configure R9 in BGP AS 100 and peer it with R5.

Solution:

Remember to activate the customer edge neighbor under the IPv4 address family for the CustomerA VRF on both R4 and R5.

When customers use the same AS number at different sites, the BGP loop prevention mechanism will keep the provider edge router from advertising routes to the customer edge router. Overcome this by configuring the **as-override** feature. It will replace each leading instance of the duplicate AS number with the core AS number.

Configure the MPLS Layer 3 VPN on R3, R4, R5, R7, R8, and R9 according to the scenario requirements.

R3:

```
interface Ethernet0/1
 ip address 172.16.53.3 255.255.255.0
 mpls ip
!
!
interface Serial1/0
 ip address 172.16.43.3 255.255.255.0
 mpls ip
!
```

R4:

```
interface Ethernet0/0.14
 encapsulation dot1Q 14
 ip vrf forwarding CustomerA
 ip address 172.16.14.4 255.255.255.0
!
interface Serial1/1
 ip address 172.16.43.4 255.255.255.0
 mpls ip
!
router bgp 4
 bgp log-neighbor-changes
 neighbor 172.16.43.3 remote-as 3
 neighbor 172.16.105.1 remote-as 4
 neighbor 172.16.105.1 update-source Loopback104
 neighbor 172.16.146.1 remote-as 1
 !
 address-family vpnv4
  neighbor 172.16.105.1 activate
  neighbor 172.16.105.1 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CustomerA
  neighbor 172.16.14.10 remote-as 100
  neighbor 172.16.14.10 activate
  neighbor 172.16.14.10 as-override
 exit-address-family
!
```

R5:

```
interface Ethernet0/0.35
 encapsulation dot1Q 35
 ip address 172.16.53.5 255.255.255.0
 mpls ip
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
!
interface Ethernet0/0.200
 encapsulation dot1Q 200
 ip vrf forwarding CustomerA
 ip address 172.16.25.5 255.255.255.0
!
!
router bgp 4
 bgp log-neighbor-changes
 neighbor 172.16.104.1 remote-as 4
 neighbor 172.16.104.1 update-source Loopback105
 !
 address-family vpnv4
  neighbor 172.16.104.1 activate
  neighbor 172.16.104.1 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CustomerA
  neighbor 172.16.25.20 remote-as 100
  neighbor 172.16.25.20 activate
  neighbor 172.16.25.20 as-override
 exit-address-family
!
```

R7:

```
router bgp 100
 bgp log-neighbor-changes
 redistribute connected
 neighbor 172.16.200.10 remote-as 100
!
```

R8:

```
router bgp 100
 bgp log-neighbor-changes
 redistribute connected
 neighbor 172.16.14.4 remote-as 4
 neighbor 172.16.200.7 remote-as 100
!
```

R9:

```
router bgp 100
 bgp log-neighbor-changes
 redistribute connected
 neighbor 172.16.25.5 remote-as 4
!
```

Verify the VPN BGP peer relationships on the routers. Here is an example on R4:

```
R4#show bgp vpnv4 unicast all summary
BGP router identifier 172.16.241.4, local AS number 4
BGP table version is 9, main routing table version 9
6 network entries using 912 bytes of memory
6 path entries using 456 bytes of memory
5/3 BGP path/bestpath attribute entries using 740 bytes of memory
3 BGP AS-PATH entries using 72 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2204 total bytes of memory
BGP activity 7/0 prefixes, 8/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.14.10    4          100      41      41        9    0    0 00:34:10        4
172.16.105.1    4            4      61      61        9    0    0 00:50:02        2
R4#
```

Note that R4 forms VPN BGP peer relationships with R5 and R8.

https://t.me/learningnets

Verify the VPN BGP prefixes on R4:

```
R4#show bgp vpnv4 unicast all
BGP table version is 9, local router ID is 172.16.241.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 4:100 (default for vrf CustomerA)
 r>  172.16.14.0/24   172.16.14.10             0             0 100 ?
 *>i 172.16.25.0/24   172.16.105.1             0    100      0 100 ?
 *>  172.16.107.0/24  172.16.14.10                           0 100 ?
 *>  172.16.110.0/24  172.16.14.10             0             0 100 ?
 *>i 172.16.120.0/24  172.16.105.1             0    100      0 100 ?
 *>  172.16.200.0/24  172.16.14.10             0             0 100 ?
R4#
```

Note that R4 learns the VPN subnet 172.16.107.0/24 from R8 and 172.16.120.0/24 from R5.

Verify connectivity between 172.16.107.0/24 and 172.16.120.0/24 within AS 100. Here is an example from R7:

```
R7#ping 172.16.120.1 source 172.16.107.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.120.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.107.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/9/11 ms
R7#
```

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

8. IPv6 Routing Section

Issue: Assign IPv6 addresses.

Solution:

Configure IPv6 addresses on R1, R2, R3, R4, R5 and R6 according to the scenario requirements.

Assign the required addresses to the interfaces on the following devices:

R1:

```
interface Ethernet0/1.12
 ipv6 address 2001::12:1/125
!
interface Ethernet0/1.16
  ipv6 address 2001::16:1/125
!
```

R2:

```
interface Ethernet0/0.12
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
  ipv6 address 2001::12:2/125
!
interface Ethernet0/0.23
ipv6 address 2001::23:2/125
```

R3:

```
interface Ethernet0/0
ipv6 address 2001::23:3/125
!
interface Ethernet0/1
 ipv6 address 2001::53:3/125

interface Serial1/0
ipv6 address 2001::43:3/125
!
```

R4:

```
interface Serial1/1
ipv6 address 2001::43:4/125
```

R5:

```
interface Ethernet0/0.35
ipv6 address 2001::53:5/125
!
```

R6:

```
interface Ethernet0/0.16
ipv6 address 2001::16:6/125
!
```

Make sure that you can ping within the same subnet before moving forward. Can R3 ping all of the addresses on the connected links?

```
R3#ping  2001::23:2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::23:2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/11 ms
R3#ping  2001::43:4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::43:4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
R3#ping  2001::53:5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::53:5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/17 ms
R3#
```

Issue: Configure OSPF Area 0 on VLAN 16 between R1 and R6. Configure OSPF Area 12 on VLAN 12. Do not send link-state advertisement (LSA) type 5 into Area 12; however, R1 and R6 must have external prefixes received from Area 12.

Solution:

Enter the command **ipv6 unicast-routing** in global configuration mode.

https://t.me/learningnets

Configure Area 12 as an NSSA. This configuration is applied under the IPv6 router OSPF configuration on R1 and R2.

Configure IPv6 OSPF on R1, R2, and R6:
R1:

```
ipv6 unicast-routing
interface Ethernet0/1.12
 ipv6 address 2001::12:1/125
!
interface Ethernet0/1.16
  ipv6 address 2001::16:1/125
!
ipv6 router ospf 1
 area 12 nssa
!
```

R2:

```
ipv6 unicast-routing
interface Ethernet0/0.12
ipv6 address 2001::12:2/125
ipv6 ospf 1 area 12
!
ipv6 router ospf 1
 area 12 nssa
!
```

R6:

```
ipv6 unicast-routing
interface Ethernet0/0.16
ipv6 address 2001::16:6/125
ipv6 ospf 1 area 0
!
```

Verify IPv6 OSPF neighbor relationships on R1:

```
R1#show ipv6 ospf neighbor

            OSPFv3 Router with ID (172.16.241.1) (Process ID 1)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
172.16.106.1      1   FULL/BDR        00:00:36    15              Ethernet0/1.16
172.16.241.2      1   FULL/BDR        00:00:38    16              Ethernet0/1.12
R1#
```

Note that R1 forms the IPv6 OSPF neighbor relationships with R2 and R6.

Issue: Configure the EIGRP AS 555 process on R2, R3, R4, and R5. Summarize 2001::43:0/125 and 2001::53:0/125 on R3 using the best possible mask.

Solution:

Enter the command **ipv6 unicast-routing** in global configuration mode.

Configure IPv6 EIGRP AS 555 on R2, R3, R4, and R5:
R2:

```
!
interface Ethernet0/0.23
ipv6 address 2001::23:2/125
ipv6 eigrp 555
!
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
ipv6 router eigrp 555
!
```

R3:

```
ipv6 unicast-routing
interface Ethernet0/0
ipv6 address 2001::23:3/125
ipv6 eigrp 555
ipv6 summary-address eigrp 555 2001::40:0/107
!
interface Ethernet0/1
 ipv6 address 2001::53:3/125
ipv6 eigrp 555
!
interface Serial1/0
ipv6 address 2001::43:3/125
ipv6 eigrp 555
!
ipv6 router eigrp 555
!
```

R4:

```
ipv6 unicast-routing
interface Serial1/1
ipv6 address 2001::43:4/125
ipv6 eigrp 555
!
ipv6 router eigrp 555
!
```

R5:

```
ipv6 unicast-routing
interface Ethernet0/0.35
ipv6 address 2001::53:5/125
ipv6 eigrp 555
!
ipv6 router eigrp 555
!
```

Verify IPv6 EIGRP neighbor relationships on R3:

```
R3#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(555)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   Link-local address:     Se1/0                    10 00:33:25   13   100  0  7
    FE80::204:C1FF:FE8E:BC0
1   Link-local address:     Et0/1                    13 00:34:50    5   100  0  6
    FE80::A8BB:CCFF:FE00:500
0   Link-local address:     Et0/0                    11 00:35:45    9   100  0  9
    FE80::A8BB:CCFF:FE00:200
R3#
```

Note that R3 forms the IPv6 EIGRP neighbor relationships with R2, R4, and R5.

Issue: Perform mutual redistribution between IPv6 EIGRP AS 555 and IPv6 OSPF on R2. Verify IPv6 connectivity.

Solution:

Configure IPv6 route redistribution on R2:

R2:

https://t.me/learningnets

```
ipv6 router eigrp 555
 redistribute ospf 1
 redistribute connected
 default-metric 1500 1000 255 3 1500
!
ipv6 router ospf 1
 area 12 nssa
 redistribute connected
 redistribute eigrp 555
!
!
```

Verify the IPv6 routing table on R1 and R5:

```
R1#show ipv6 route ospf
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
ON2 2001::23:0/125 [110/20]
     via 2001::12:2, Ethernet0/1.12
ON2 2001::40:0/107 [110/20]
     via 2001::12:2, Ethernet0/1.12
R1#
```

Note that R1 learns prefixes that are originated in the IPv6 EIGRP AS 555 as IPv6 OSPF external NSSA prefixes. The IPv6 summary network 2001::40:0/107 represents the 2001::43:0/125 and 2001::53:0/125 IPv6 subnets.

```
R5#show ipv6 route eigrp
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
EX  2001::12:0/125 [170/2013696]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.35
EX  2001::16:0/125 [170/2013696]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.35
D   2001::23:0/125 [90/307200]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.35
D   2001::43:0/125 [90/2195456]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.35
R5#
```

Note that R5 learns prefixes that are originated in the IPv6 OSPF as IPv6 EIGRP external.

Ping R6 from R5:

```
R5#ping 2001::16:6
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::16:6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/14 ms
R5#
```

Note that R5 can successfully ping R6 across the IPv6 OSPF and EIGRP routing domains.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

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

9. Cisco IOS Software Features Section

Issue: Configure R1 to report EIGRP notifications to the Simple Network Management Protocol (SNMP) server with the address 172.16.12.10 (imaginary). Send this information using traps. Use SNMPv2c and the community string SNMPSERVER. Allow remote access to this router with the community string EIGRPR1.

Solution:

The EIGRP MIB can now be queried using SNMP.
Configure the SNMP on R1.

R1:
```
snmp-server community EIGRPR1 RO
snmp-server enable traps eigrp
snmp-server host 172.16.12.10 version 2c SNMPSERVER
!
```

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

10. QoS Section

Issue: There is a flood attack. R1 receives an excessive number of IPv6 packets with random protocol, source, and destination addresses and ports, but the length is equal to 110.
Packets are coming to R1 via an Ethernet interface, probably from R6. It has been confirmed that the packets transferred via R2 are not part of the attack. Block that traffic on R1. Make sure that packets of size 110 coming from R2 are not affected. Set the R2 MAC address as 0004.c18e.0bc0.

Solution:

Block excessive traffic on R1. Make sure that packets of size 110 coming from R2 are not affected.

The class-based weighted fair queuing (CBWFQ) QoS method provides an option to create a match criteria based on the packet length. Packets that are coming from R2 are differentiated based on the R2 MAC address. R1 will use the class “attack” to match all IPv6 length-110 packets (declared malicious) and drop them on ingress interfaces.

The interface of R2 is configured with MAC address 0004.c18e.0bc0.

Configure the QoS policy and apply it to the interfaces on R1:

R1:
```
!
class-map match-all ATTACK-Ethernet
 match protocol ipv6
```


https://t.me/learningnets

```
 match not source-address mac 0004.C18E.0BC0
 match packet length min 110 max 110
!
policy-map ATTACK-Ethernet
 class ATTACK-Ethernet
  drop
!
interface Ethernet0/1.12
 encapsulation dot1Q 12
 ip address 172.16.12.1 255.255.255.0
 service-policy input ATTACK-Ethernet
!
interface Ethernet0/1.16
 encapsulation dot1Q 16
 ip address 172.16.61.1 255.255.255.0
 service-policy input ATTACK-Ethernet
!
```

Verify the QoS policy by pinging the R2 IPv6 address from R6 using an extended **ping** command specifying a different size packet.

```
R6#ping
Protocol [ip]: ipv6
Target IPv6 address: 2001::12:2
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands? [no]:
Sweep range of sizes? [no]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::12:2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R6#
```

Note that the standard size 100 byte IPv6 ping is successful.

Try to ping using the packet size 110:

```
R6#ping
Protocol [ip]: ipv6
Target IPv6 address: 2001::12:2
Repeat count [5]:
Datagram size [100]: 110
Timeout in seconds [2]:
Extended commands? [no]:
Sweep range of sizes? [no]:
Type escape sequence to abort.
Sending 5, 110-byte ICMP Echos to 2001::12:2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R6#
```

Note that the ping is blocked.

Verify the QoS policy on R1

```
R1#show policy-map interface e0/1.16
 Ethernet0/1.16

  Service-policy input: ATTACK-Ethernet

    Class-map: ATTACK-Ethernet (match-all)
      5 packets, 640 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: protocol ipv6
      Match: not source-address mac 0004.C18E.0BC0
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
      Match: packet length min 110 max 110
      drop

    Class-map: class-default (match-any)
      46 packets, 3752 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
R1#
```

Note that R1 shows that five packets are dropped by the QoS policy.

Ping R1 from R2 using the same packet size of 110 bytes.

```
R2#ping
Protocol [ip]: ipv6
Target IPv6 address: 2001::16:1
Repeat count [5]:
Datagram size [100]: 110
Timeout in seconds [2]:
Extended commands? [no]:
Sweep range of sizes? [no]:
Type escape sequence to abort.
Sending 5, 110-byte ICMP Echos to 2001::16:1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R2#
```

Note that ping from R2 to R1 is successful.
Verify the QoS policy on R1

```
R1#show policy-map interface e0/1.12
 Ethernet0/1.12

  Service-policy input: ATTACK-Ethernet

    Class-map: ATTACK-Ethernet (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: protocol ipv6
      Match: not source-address mac 0004.C18E.0BC0
      Match: packet length min 110 max 110
      drop

    Class-map: class-default (match-any)
      117 packets, 10772 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
R1#
```

Note that R1 shows that the QoS policy does not classify the traffic from R2 to be dropped.

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

https://t.me/learningnets