360RSW05-LAB-CA12

Cisco Expert-Level Training for CCIE Routing and Switching Assessment Lab 12 Configuration Section

This assessment lab measures your ability to manage Cisco CCIE® Routing and Switching (R&S) problems when under time pressure. The tasks that this lab presents require careful evaluation and analysis of issues and options. Furthermore, you must rigorously verify all the tasks that you perform. You must complete all the tasks within a finite period.

After completing the lab, you will have access to a detailed answer key and a Mentor Guide repository that represents the completed state of the pod at the end of the lab. Each assessment lab user can compare his or her personal Mentor Guide output with the master Mentor Guide output.

Cisco Expert-Level Training for CCIE Routing and Switching Assessment Lab 12 Answer Key Configuration Section

Table of Contents

Answer Key Structure

Section 1

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

Section 2

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

Cisco Expert-Level Training for CCIE Routing and Switching Assessment Lab 12 Answer Key Configuration Section

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

Grading and Duration

- Configuration lab duration: 6 hours
- Configuration lab maximum score: 70 points
- Minimum passing score: 56 points

Restrictions and Goals

Note Read this section carefully.

- To receive any credit for a subsection, you must complete the subsection. Partial credit is not given for partially completed subsections.
- Do not introduce any new IP addresses; use only IP addresses that are specified in the scenario.
- Static routes are not allowed in this exercise unless specified otherwise.
- Advertise IPv4 and IPv6 loopback interfaces with their original masks, unless specified otherwise.
- 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.

1. Layer 2 Technologies Section

Configure the VLANs and the VLAN names according to the scenario specifications and assign the ports of the switches to these VLANs. Ensure that the VLAN names are spelled correctly and that the cases of letters match.

To help their understanding of the Layer 2 topology, many candidates find it helpful to create a VLAN propagation diagram such as the diagram that follows. To create a diagram, study the VLANs table, the Switch-to-Router Connections table, the Switch-to-Switch Links table, the interior gateway protocol (IGP) diagrams, and the other section requirements, and then carefully document each connection on a copy of the physical layer diagram. Most candidates find that, with practice, such a diagram can be created quickly and provides a valuable tool.

Headquarters VLAN Distribution

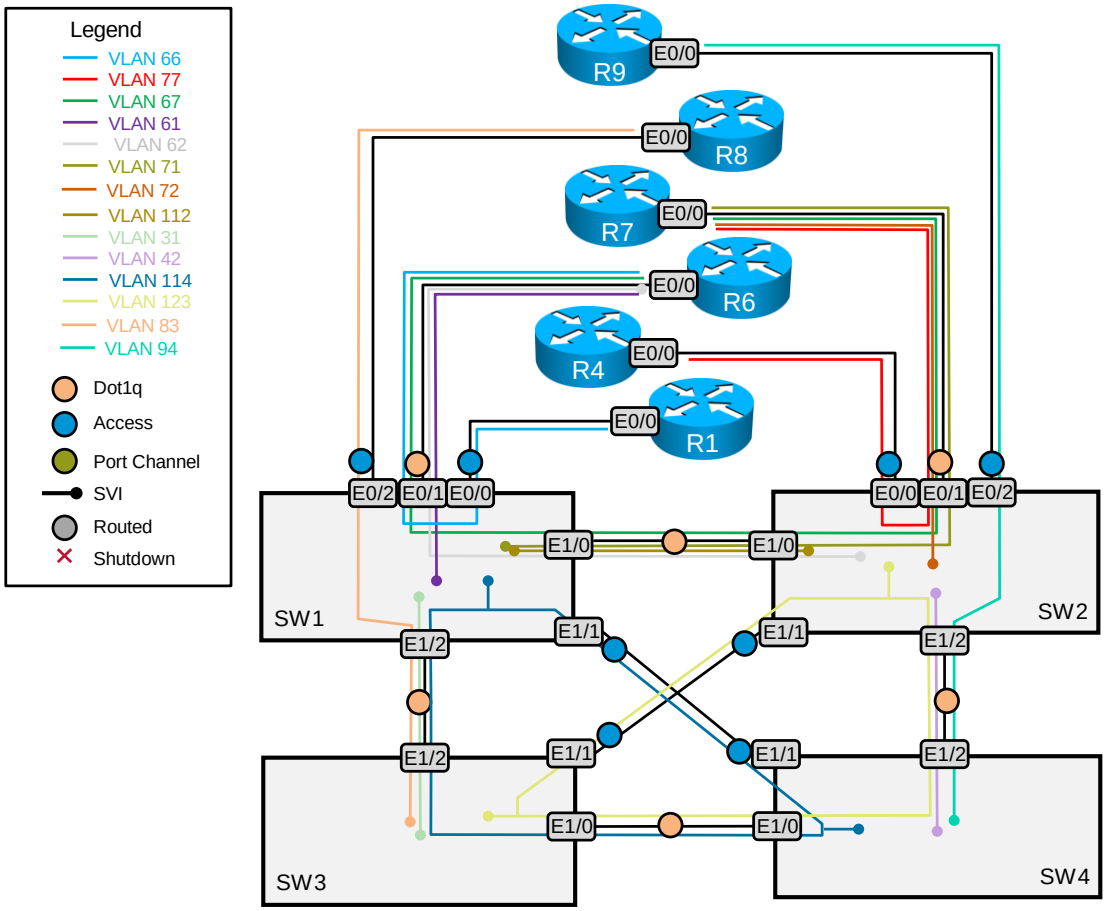

Issue: Configure the VLAN Trunking Protocol (VTP), VLANs, Switch-to-Router connection, and Switch-to-Switch links in the Headquarters, according to the lab requirements.

Solution:

You are told explicitly to use VTP version 2. Also you are not allowed to manually configure, create, delete or modify VLANs on SW2, SW3, and SW4. Therefore, SW2, SW3, and SW4 will be configured as VTP clients and will learn VLANs from the VTP server SW1. To configure VTP version 2 clients and server mode, issue the **vtp version 2, vtp mode server, and vtp mode client**

commands. Manually set the encapsulation type to dot1q by using the **switchport trunk encapsulation dot1q** command on all necessary trunk ports. Use the **vtp domain CCIE** command to set the VTP domain name according to the lab specifications.

Verification:

Verify the VTP status and password on the switches. Here are examples from SW1 and SW2:

```
SW1#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 2
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1800
Configuration last modified by 12.0.0.110 at 12-1-14 22:02:32
Local updater ID is 12.1.31.110 on interface Vl31 (lowest numbered VLAN interface found)

Feature VLAN:
--------------
VTP Operating Mode              : Server
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 19
Configuration Revision            : 2
MD5 digest                        : 0xD3 0x1E 0x65 0x35 0x9C 0xC5 0xB8 0x7C
                                    0xA6 0xAE 0xAF 0xAC 0xA4 0x17 0xD1 0xF7
SW1#

SW2#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 2
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1900
Configuration last modified by 12.0.0.110 at 12-1-14 22:02:32

Feature VLAN:
--------------
VTP Operating Mode              : Client
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 19
Configuration Revision            : 2
MD5 digest                        : 0xD3 0x1E 0x65 0x35 0x9C 0xC5 0xB8 0x7C
                                    0xA6 0xAE 0xAF 0xAC 0xA4 0x17 0xD1 0xF7
SW2#

SW3#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 2
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1a00
Configuration last modified by 12.0.0.110 at 12-1-14 22:02:32

Feature VLAN:
--------------
VTP Operating Mode              : Client
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 19
Configuration Revision            : 2
MD5 digest                        : 0xD3 0x1E 0x65 0x35 0x9C 0xC5 0xB8 0x7C
                                    0xA6 0xAE 0xAF 0xAC 0xA4 0x17 0xD1 0xF7
SW3#

SW4#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 2
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1b00
Configuration last modified by 12.0.0.110 at 12-1-14 22:02:32
```

```
Feature VLAN:
--------------
VTP Operating Mode                : Client
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 19
Configuration Revision            : 2
MD5 digest                        : 0xD3 0x1E 0x65 0x35 0x9C 0xC5 0xB8 0x7C
                                    0xA6 0xAE 0xAF 0xAC 0xA4 0x17 0xD1 0xF7
SW4#
```

Verify VLANs on the switches. Here are examples from SW1 and SW2:

```
SW1#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/3, Et1/3, Et2/0, Et2/1
                                                Et2/2, Et2/3
31   VLAN0031                         active
42   VLAN0042                         active
61   VLAN0061                         active
62   VLAN0062                         active
66   VLAN0066                         active    Et0/0
67   VLAN0067                         active
71   VLAN0071                         active
72   VLAN0072                         active
77   VLAN0077                         active
83   VLAN0083                         active    Et0/2
94   VLAN0094                         active
112  VLAN0112                         active
114  VLAN0114                         active    Et1/1
123  VLAN0123                         active
1002 fddi-default                     act/unsup
1003 trcrf-default                    act/unsup
1004 fddinet-default                  act/unsup
1005 trbrf-default                    act/unsup
SW1#

SW2#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/3, Et1/3, Et2/0, Et2/1
                                                Et2/2, Et2/3
31   VLAN0031                         active
42   VLAN0042                         active
61   VLAN0061                         active
62   VLAN0062                         active
66   VLAN0066                         active
67   VLAN0067                         active
71   VLAN0071                         active
72   VLAN0072                         active
77   VLAN0077                         active    Et0/0
83   VLAN0083                         active
94   VLAN0094                         active    Et0/2
112  VLAN0112                         active
114  VLAN0114                         active
123  VLAN0123                         active    Et1/1
1002 fddi-default                     act/unsup
1003 trcrf-default                    act/unsup
1004 fddinet-default                  act/unsup
1005 trbrf-default                    act/unsup
SW2#
```

Verify the status of the trunks on the switches. Here are examples from SW1, SW2, SW3 and SW4:

```
SW1#show interfaces trunk

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



```
Port        Vlans allowed on trunk
Et0/1       61-62,66-67
Et1/0       62,67,71,112
Et1/2       31,83,114

Port        Vlans allowed and active in management domain
Et0/1       61-62,66-67
Et1/0       62,67,71,112
Et1/2       31,83,114

Port        Vlans in spanning tree forwarding state and not pruned
Et0/1       61-62,66-67
Et1/0       62,67,71,112
Et1/2       31,83
SW1#

SW2#show interfaces trunk

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

Port        Vlans allowed on trunk
Et0/1       67,71-72,77
Et1/0       62,67,71,112
Et1/2       42,94,123

Port        Vlans allowed and active in management domain
Et0/1       67,71-72,77
Et1/0       62,67,71,112
Et1/2       42,94,123

Port        Vlans in spanning tree forwarding state and not pruned
Et0/1       67,71-72,77
Et1/0       62,67,71,112
Et1/2       42,94,123
SW2#
```

Note that only VLAN 62, VLAN 67, VLAN 71, and VLAN 112 are allowed on the port Eth1/0 between SW1 and SW2.

```
SW3#show interfaces trunk

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

Port        Vlans allowed on trunk
Et1/0       114,123
Et1/2       31,83,114

Port        Vlans allowed and active in management domain
Et1/0       114,123
Et1/2       31,83,114

Port        Vlans in spanning tree forwarding state and not pruned
Et1/0       114
Et1/2       31,83,114
SW3#
SW4#show interfaces trunk

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

Port        Vlans allowed on trunk
Et1/0       114,123
Et1/2       42,94,123

Port        Vlans allowed and active in management domain
Et1/0       114,123
Et1/2       42,94,123
```

```
Port        Vlans in spanning tree forwarding state and not pruned
Et1/0       114,123
Et1/2       42,94,123
SW4#
```

Note that only VLAN 31, VLAN 83, and VLAN114 are allowed on the ports Eth1/2 between SW1 and SW3.

Only VLAN 42, VLAN 94, and VLAN 123 are allowed on the ports Eth1/2 between SW2 and SW4.

Only VLAN 114 and VLAN 123 are allowed on the ports Eth1/0 between SW3 and SW4.

According to the lab requirements, the Ethernet1/1 interfaces are configured for access vlan mode on SW1, SW2, SW3, and SW4. The VLAN's that these access ports will be associated with can be found in two places: (1) on the IPv4 Topology Diagram and (2) in the Headquarters Switch to Switch Connection table.

```
SW1#show interfaces status | include Et1/1
Et1/1                     connected    114          auto   auto unknown
SW1#

SW2#show interfaces status | include Et1/1
Et1/1                     connected    123          auto   auto unknown
SW2#

SW3#show interfaces status | include Et1/1
Et1/1                     connected    123          auto   auto unknown
SW3#

SW4#show interfaces status | include Et1/1
Et1/1                     connected    114          auto   auto unknown
SW4#
```


Regional Main Office VLAN Distribution

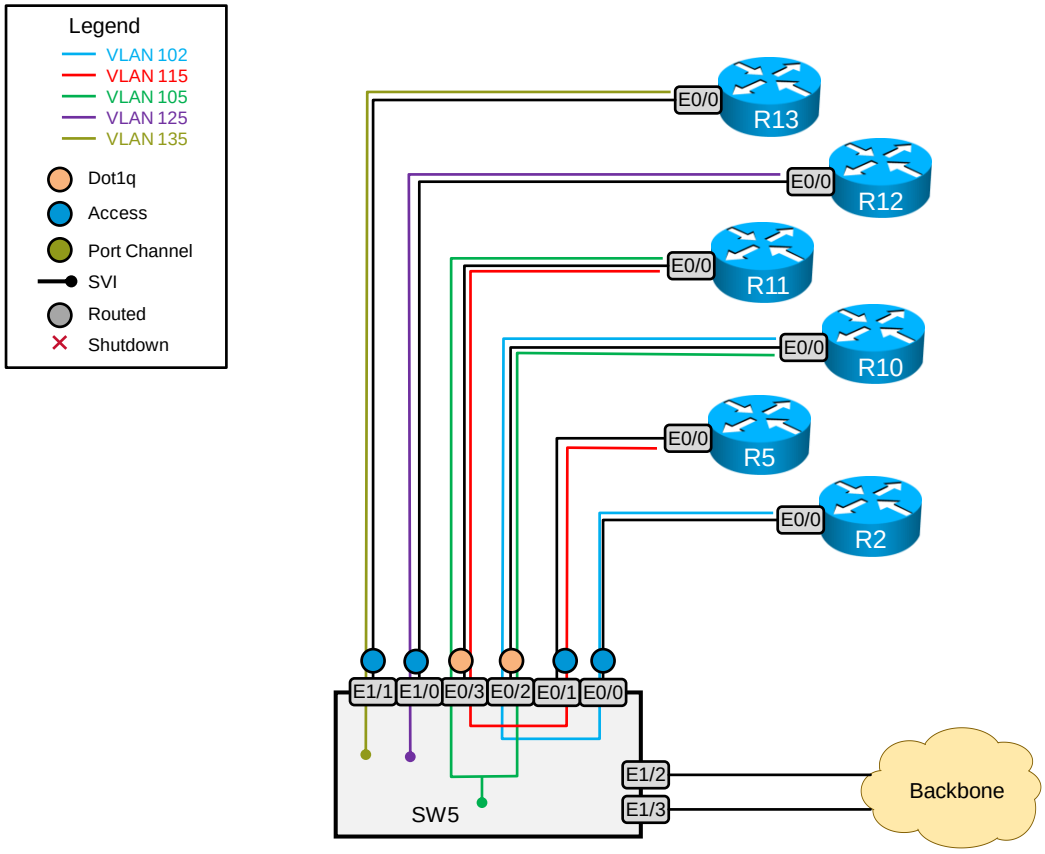

Issue: Configure the VLAN Trunking Protocol (VTP), VLANs, Switch-to-Router connections, and Switch-to-Switch links in the Regional Main Office according to the lab requirements.

Solution:

You are not allowed to create any VLANs on SW5. Note that SW5 is connected to a backbone switch. Configure SW5 as a VTP client and learn the VLANs from the backbone switch. To configure VTP version 2 client mode, issue the **vtp version 2** and **vtp mode client** commands. Manually set the encapsulation type to dot1q by using the **switchport trunk encapsulation dot1q** command on all necessary trunk ports. Use the **vtp domain CISCO** and **vtp password CISCO** commands to set the VTP domain name and password according to the lab specifications.

Verification:

Verify the VTP status and the VTP password on SW5:

```
SW5#
SW5#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 2
VTP Domain Name                 : CISCO
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1c00
Configuration last modified by 0.0.0.0 at 12-2-14 00:22:01

Feature VLAN:
--------------
```

```
VTP Operating Mode                : Client
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 10
Configuration Revision            : 5
MD5 digest                        : 0x23 0x59 0xBA 0x2D 0xB2 0x33 0x28 0xC4
                                    0x83 0x8B 0x82 0xA9 0x14 0xEA 0x3E 0xEA
SW5#

SW5#show vtp password
VTP Password: CISCO
SW5#
```

Note that SW5 shows 10 VLANs: Five default VLANs, and five learned VLANs from the backbone. The configuration revision of the backbone VTP server is 5.

Verify VLANs on SW5:

```
SW5#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active
102  VLAN0102                         active    Et0/0
105  VLAN0105                         active
115  VLAN0115                         active    Et0/1
125  VLAN0125                         active    Et1/0
135  VLAN0135                         active    Et1/1
1002 fddi-default                     act/unsup
1003 trcrf-default                    act/unsup
1004 fddinet-default                  act/unsup
1005 trbrf-default                    act/unsup
SW5#
```

Verify the status of the trunks on SW5:

```
SW5#show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/2       on               802.1q         trunking      1
Et0/3       on               802.1q         trunking      1
Et1/2       desirable        n-802.1q       trunking      1
Et1/3       desirable        n-802.1q       trunking      1

Port        Vlans allowed on trunk
Et0/2       1-4094
Et0/3       1-4094
Et1/2       1-4094
Et1/3       1-4094

Port        Vlans allowed and active in management domain
Et0/2       1,102,105,115,125,135
Et0/3       1,102,105,115,125,135
Et1/2       1,102,105,115,125,135
Et1/3       1,102,105,115,125,135

Port        Vlans in spanning tree forwarding state and not pruned
Et0/2       1,102,105,115,125,135
Et0/3       1,102,105,115,125,135
Et1/2       1,102,105,115,125,135
Et1/3       1,102,105,115,125,135
SW5#
```

Note that the backbone interfaces Eth1/2 and Eth1/3 are in trunking mode. It is required to receive the VTP updates from the VTP server that is located in the backbone network.

Regional Sales and Marketing Offices VLAN Distribution

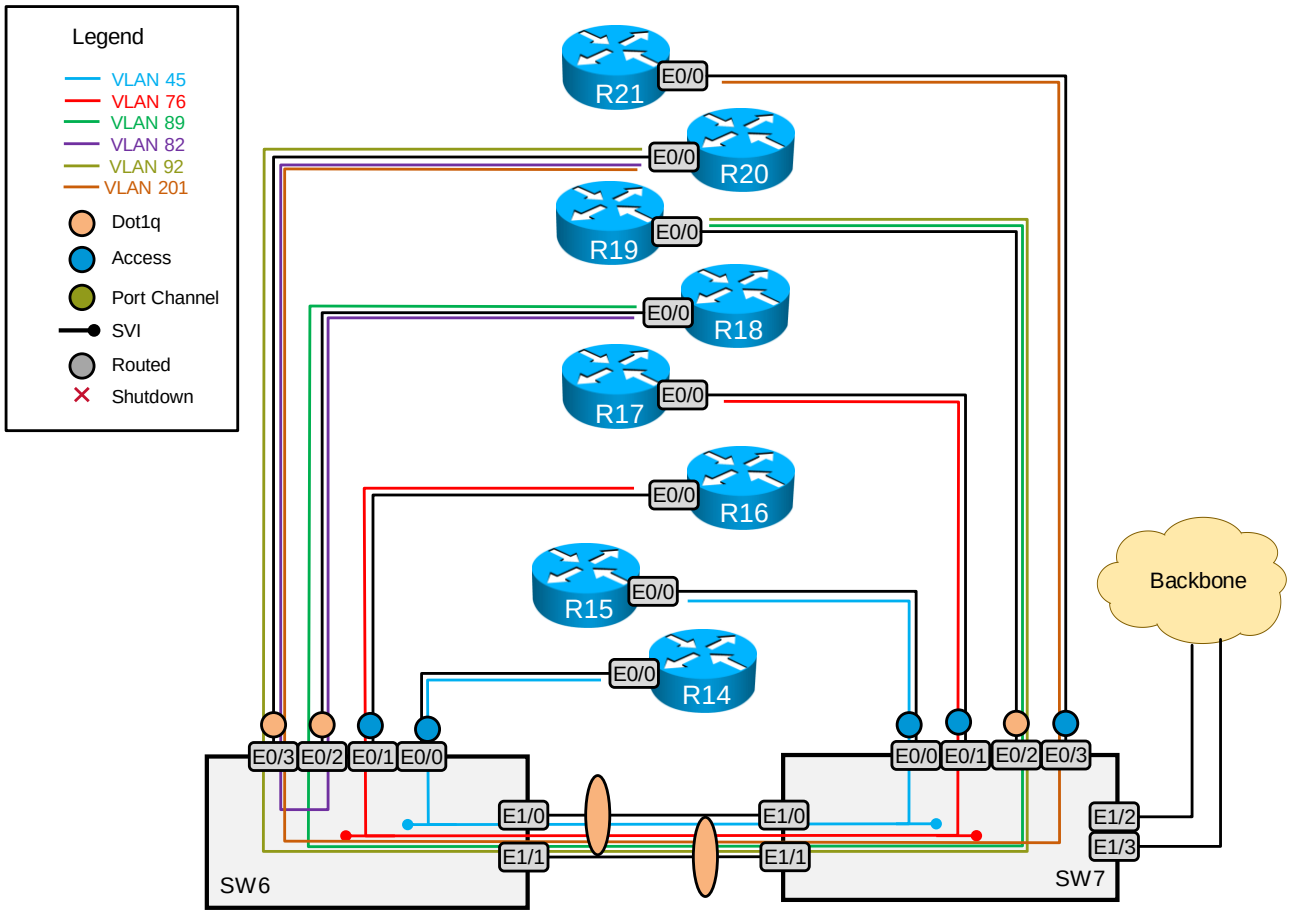

Issue: Configure the VLAN Trunking Protocol (VTP), VLANs, Switch-to-Router connections, and Switch-to-Switch links in the Regional Sales and Marketing Offices, according to the lab requirements.

All switches in the Regional Sales and Marketing Offices should not be able to advertise or synchronize their VLAN configuration based on received advertisements, but should be able to forward VTP advertisements that they receive out their trunk ports in VTP version 2

Solution:

Configure the VTP transparent mode. To configure VTP in transparent mode, issue the **vtp mode transparent** command. Issue the **vtp version 2** command. Manually set the encapsulation type to dot1q by using the **switchport trunk encapsulation dot1q** command on all necessary ports and manually configure VLANs on SW6 and SW7 according to the lab requirements.

Verify the VTP status on SW6 and SW7:

```
SW6#show vtp status
VTP Version capable             : 1 to 3
```

```
VTP version running             : 2
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1d00
Configuration last modified by 12.3.45.160 at 0-0-00 00:00:00

Feature VLAN:
--------------
VTP Operating Mode              : Transparent
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 11
Configuration Revision            : 0
MD5 digest                        : 0x84 0xBB 0x51 0xE6 0x2F 0xB1 0xAF 0x58
                                    0x30 0xE9 0x10 0xBD 0x32 0xDC 0xE0 0xFF
SW6#

SW7#show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 2
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : aabb.cc00.1e00
Configuration last modified by 12.3.45.170 at 0-0-00 00:00:00

Feature VLAN:
--------------
VTP Operating Mode              : Transparent
Maximum VLANs supported locally   : 1005
Number of existing VLANs          : 11
Configuration Revision            : 0
MD5 digest                        : 0x27 0xE6 0x1E 0xB1 0x9D 0x9F 0x33 0x22
                                    0x7E 0xD3 0xCD 0x7F 0xA8 0xC7 0xE0 0xEE
SW7#
```

Verify VLANs on SW6 and SW7:

```
SW6#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et1/2, Et1/3
45   VLAN0045                         active    Et0/0
76   VLAN0076                         active    Et0/1
82   VLAN0082                         active
89   VLAN0089                         active
92   VLAN0092                         active
201  VLAN0201                         active
1002 fddi-default                     act/unsup
1003 trcrf-default                    act/unsup
1004 fddinet-default                  act/unsup
1005 trbrf-default                    act/unsup
SW6#

SW7#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active
45   VLAN0045                         active    Et0/0
76   VLAN0076                         active    Et0/1
82   VLAN0082                         active
89   VLAN0089                         active
92   VLAN0092                         active
201  VLAN0201                         active    Et0/3
1002 fddi-default                     act/unsup
1003 trcrf-default                    act/unsup
1004 fddinet-default                  act/unsup
1005 trbrf-default                    act/unsup
SW7#
```

Verify the status of the trunks on SW6 and SW7:


```
SW6#show interfaces trunk

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

Port        Vlans allowed on trunk
Et0/2       1-4094
Et0/3       1-4094
Et1/0       1-4094
Et1/1       1-4094

Port        Vlans allowed and active in management domain
Et0/2       1,45,76,82,89,92,201
Et0/3       1,45,76,82,89,92,201
Et1/0       1,45,76,82,89,92,201
Et1/1       1,45,76,82,89,92,201

Port        Vlans in spanning tree forwarding state and not pruned
Et0/2       1,45,76,82,89,92,201
Et0/3       1,45,76,82,89,92,201
Et1/0       1,45,76,82,89,92,201
Et1/1       1,45,76,82,89,92,201
SW6#
SW7#show interfaces trunk

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

Port        Vlans allowed on trunk
Et0/2       1-4094
Et1/0       1-4094
Et1/1       1-4094
Et1/2       1-4094
Et1/3       1-4094

Port        Vlans allowed and active in management domain
Et0/2       1,45,76,82,89,92,201
Et1/0       1,45,76,82,89,92,201
Et1/1       1,45,76,82,89,92,201
Et1/2       1,45,76,82,89,92,201
Et1/3       1,45,76,82,89,92,201

Port        Vlans in spanning tree forwarding state and not pruned
Et0/2       1,45,76,82,89,92,201
Et1/0       1,76,82,89,92,201
Et1/1       45
Et1/2       1,76
Et1/3       1,76
SW7#
```

Note that the backbone interfaces Eth1/2 and Eth1/3 on SW7 are in trunking mode. The whole range of possible VLANs 1-4096 is allowed on all trunks on SW6 and SW7.

Issue:

Configure SW3 to be the root bridge for VLAN 114.

Configure SW1 so that the port E1/2 is blocking for VLAN 114.

Solution:

Here are some suggested configuration and verification examples on SW3 and SW1. These solutions will fulfill the requirements of this task. However, please note that the IOS provides other valid solutions to these tasks.

SW3:

```
SW3#show run | inc spanning-tree vlan 114
spanning-tree vlan 114 priority 24576
SW3#
```

Note that you can use the **spanning-tree vlan 114 root primary** command to configure SW3 as the VLAN 114 root bridge.

Note that there is a topological loop for VLAN 114 between SW1, SW3 and SW4. See the Headquarters VLAN Distribution diagram. Use the spanning tree cost configuration on SW1 to set the Et1/2 interface in a blocking state. Here is an example from SW1.

By increasing the cost on this port, SW1 has placed this port into a Spanning Tree blocking state.

SW1:

```
interface Ethernet1/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 31,83,114
 switchport mode trunk
 duplex auto
 spanning-tree vlan 114 cost 201
!
```

Verify the spanning tree of the VLAN 114 on SW1, SW3 and SW4:

```
SW1#show spanning-tree vlan 114

VLAN0114
  Spanning tree enabled protocol ieee
  Root ID    Priority    24690
             Address     aabb.cc00.1a00
             Cost        200
             Port        6 (Ethernet1/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32882  (priority 32768 sys-id-ext 114)
             Address     aabb.cc00.1800
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Et1/1               Root FWD 100       128.6    Shr
Et1/2               Altn BLK 201       128.7    Shr


SW1#

SW3#show spanning-tree vlan 114

VLAN0114
  Spanning tree enabled protocol ieee
  Root ID    Priority    24690
             Address     aabb.cc00.1a00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24690  (priority 24576 sys-id-ext 114)
             Address     aabb.cc00.1a00
             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


SW3#

SW4#show spanning-tree vlan 114
```



```
VLAN0114
  Spanning tree enabled protocol ieee
  Root ID    Priority    24690
             Address     aabb.cc00.1a00
             Cost        100
             Port        5 (Ethernet1/0)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32882  (priority 32768 sys-id-ext 114)
             Address     aabb.cc00.1b00
             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               Root FWD 100       128.5    Shr
Et1/1               Desg FWD 100       128.6    Shr


SW4#
```

Issue:

Configure SW2 to be the root bridge for VLAN 123.

The port E1/0 of SW3 must be blocking for VLAN 123. Do not manipulate the VLAN 123 spanning tree cost and port priority on any switch to accomplish this task.

Solution:

Note that there is a topological loop for VLAN 123 between SW2, SW3 and SW4. See the Headquarters VLAN Distribution diagram.

Here are the configuration and verification examples on SW2 and SW1.

SW2:

```
SW2#show run | inc spanning-tree vlan 123
spanning-tree vlan 123 priority 24576
SW2#
```

Note that you can use the **spanning-tree vlan 123 root primary** command to configure SW2 as the VLAN 123 root bridge.

Due to the explicit restriction stated in this exam that specifies that you are not allowed to configure the spanning tree cost or priority in this task, configure SW4 as the backup root bridge for VLAN 123. Here is an example from SW4.

SW4:

```
SW4#show run | inc spanning-tree vlan 123
spanning-tree vlan 123 priority 28672
SW4#
```

Note that you can use the **spanning-tree vlan 123 root secondary** command to configure SW4 as the VLAN 123 root bridge.

Verify the spanning tree of the VLAN 123 on SW2, SW3 and SW4:

```
SW2#show spanning-tree vlan 123

VLAN0123
  Spanning tree enabled protocol ieee
  Root ID    Priority    24699
             Address     aabb.cc00.1900
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24699  (priority 24576 sys-id-ext 123)
             Address     aabb.cc00.1900
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec
```

```
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Et1/1               Desg FWD 100       128.6    Shr
Et1/2               Desg FWD 100       128.7    Shr


SW2#
SW3#show spanning-tree vlan 123

VLAN0123
  Spanning tree enabled protocol ieee
  Root ID    Priority    24699
             Address     aabb.cc00.1900
             Cost        100
             Port        6 (Ethernet1/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32891  (priority 32768 sys-id-ext 123)
             Address     aabb.cc00.1a00
             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               Altn BLK 100       128.5    Shr
Et1/1               Root FWD 100       128.6    Shr


SW3#
SW4#show spanning-tree vlan 123

VLAN0123
  Spanning tree enabled protocol ieee
  Root ID    Priority    24699
             Address     aabb.cc00.1900
             Cost        100
             Port        7 (Ethernet1/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    28795  (priority 28672 sys-id-ext 123)
             Address     aabb.cc00.1b00
             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               Root FWD 100       128.7    Shr


SW4#
```

Issue:

Configure SW6 as the root switch for VLANs 45 and 76 using the **spanning-tree vlan** *N* **root primary** command.

On SW6, configure the E1/1 port so that the port E1/0 on SW7 is blocking for VLAN 45.

Ensure that all Ethernet interfaces on SW6 and SW7 are up and the whole range of possible VLANs is allowed on all configured and preconfigured trunks on SW6 and SW7.

Solution:

You are explicitly told to configure SW6 as the root bridge using the **spanning-tree vlan** *N* **root primary** command. The **spanning-tree vlan** *N* **root primary** command sets the bridge priority to 24576. You find that SW6 does not become the root bridge for VLANs 45 and 76 because there is another switch in the backbone that has the bridge priority 0.

Configure **spanning-tree guard root** command on the backbone Eth1/2 and Eth1/3 interfaces on SW7.

Manipulate the spanning tree port priority on the Eth1/1 interface of SW6 to place the port 1/0 of SW7 in the blocking state:

```
SW6#show running-config interface e1/1
Building configuration...

Current configuration : 143 bytes
!
interface Ethernet1/1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 duplex auto
 spanning-tree vlan 45 port-priority 64
end

SW6#
```

Verify the spanning-tree outputs on SW6 and SW7:

```
SW6#show spanning-tree vlan 45

VLAN0045
  Spanning tree enabled protocol ieee
  Root ID    Priority    24621
             Address     aabb.cc00.1d00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24621  (priority 24576 sys-id-ext 45)
             Address     aabb.cc00.1d00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Et0/0               Desg FWD 100       128.1    Shr
Et0/2               Desg FWD 100       128.3    Shr
Et0/3               Desg FWD 100       128.4    Shr
Et1/0               Desg FWD 100       128.5    Shr
Et1/1               Desg FWD 100        64.6    Shr


SW6#

SW7#show spanning-tree vlan 45

VLAN0045
  Spanning tree enabled protocol ieee
  Root ID    Priority    24621
             Address     aabb.cc00.1d00
             Cost        100
             Port        6 (Ethernet1/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32813  (priority 32768 sys-id-ext 45)
             Address     aabb.cc00.1e00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Et0/0               Desg FWD 100       128.1    Shr
Et0/2               Desg FWD 100       128.3    Shr
Et1/0               Altn BLK 100       128.5    Shr
Et1/1               Root FWD 100       128.6    Shr
Et1/2               Desg BKN*100       128.7    Shr *ROOT Inc
Et1/3               Desg BKN*100       128.8    Shr *ROOT_Inc


SW7#
```

Issue:

Configure WAN link on R6 according to the lab requirements.

Solution:

Here are PPP EAP protocol configuration and verification examples on R6. Please note that in a production environment, EAP is usually configured to communicate with a RADIUS server. Since there will be no RADIUS server in the RS-CCIEv5 lab, none is required here.

```
interface Serial1/0
 ip address 123.1.6.2 255.255.255.252
 encapsulation ppp
 ipv6 address 2000:CC1E:CAB:6::6/64
 ppp eap identity HQ
 ppp eap password 0 CISCOCCIE
 serial restart-delay 0
!

R6#show ppp all
Interface/ID OPEN+ Nego* Fail-     Stage    Peer Address    Peer Name
------------ -------------------- -------- --------------- --------------------
Se1/0        LCP+ IPCP+ IPV6CP+ C> LocalT   123.1.6.1       AS-30000
R6#


R6#show ppp interface serial1/0
PPP Serial Context Info
-------------------
Interface         : Se1/0
PPP Serial Handle: 0xD9000001
PPP Handle        : 0x6B000001
SSS Handle        : 0x26000001
AAA ID            : 12
Access IE         : 0xAF000001
SHDB Handle       : 0x0
State             : Up
Last State        : Binding
Last Event        : LocalTerm

PPP Session Info
----------------
Interface         : Se1/0
PPP ID            : 0x6B000001
Phase             : UP
Stage             : Local Termination
Peer Name         : AS-30000
Peer Address      : 123.1.6.1
Control Protocols: LCP[Open] IPCP[Open] IPV6CP[Open] CDPCP[Open]
Session ID        : 1
AAA Unique ID     : 12
SSS Manager ID    : 0x26000001
SIP ID            : 0xD9000001
PPP_IN_USE        : 0x11

Se1/0 LCP: [Open]
Our Negotiated Options
Se1/0 LCP:    MagicNumber 0xBBCC1EE0 (0x0506BBCC1EE0)
Peer's Negotiated Options
Se1/0 LCP:    AuthProto EAP (0x0304C227)
Se1/0 LCP:    MagicNumber 0xBBCC1E74 (0x0506BBCC1E74)

Se1/0 IPCP: [Open]
Our Negotiated Options
Se1/0 IPCP:    Address 123.1.6.2 (0x03067B010602)
Peer's Negotiated Options
Se1/0 IPCP:    Address 123.1.6.1 (0x03067B010601)

Se1/0 IPV6CP: [Open]
Our Negotiated Options
Se1/0 IPV6CP:    Interface-Id A8BB:CCFF:FE00:0600 (0x010AA8BBCCFFFE000600)
Peer's Negotiated Options
Se1/0 IPV6CP:    Interface-Id A8BB:CCFF:FE00:1F00 (0x010AA8BBCCFFFE001F00)

Se1/0 CDPCP: [Open]
```



```
Our Negotiated Options
  NONE
Peer's Negotiated Options
  NONE
R6#
```

In this lab, there are only two point-to-point serial links that both maintain very important connections that needed to fulfill many other tasks in this lab. This section addressed the first of these point-to-point links on R6. As can be seen, it requires the user to configure PPP with EAP. The second point-to-point serial link in this lab is on R10. It is pre-configured with HDLC.

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

2. Layer 3 Technologies Section

Issue:

Configure IPv4 EIGRP AS 100 on all links between R1, R4, R6, and R7 using the EIGRP named mode. The EIGRP AS 100 instance name is **HQ-CE.**

Configure IPv4 EIGRP AS 1000 on all links between R6, R7, R8, R9, SW1, SW2, SW3, and SW4 using the EIGRP named mode. The EIGRP AS 1000 instance name is **HQ-MAIN**.

Solution:

The following are configuration and verification examples on R1, R4, R6, and R7:

R1:

```
router eigrp HQ-CE
 !
 address-family ipv4 unicast vrf SITE1 autonomous-system 100
  !
  topology base
  exit-af-topology
  network 12.1.66.0 0.0.0.255
 exit-address-family
!
```

Note that the EIGRP configuration on R1 depends on the MP-BGP VPN requirements for VRF SITE1. This is due to the fact that R1 is a PE router in the upcoming MPLS Layer 3 VPN section. This is one of the challenges of this exam: reading the entire lab from end to end so that you can configure technologies only once to accommodate for inter-dependent tasks that arise between different sections of this exam.

R4:

```
router eigrp HQ-CE
 !
 address-family ipv4 unicast vrf SITE2 autonomous-system 100
  !
  topology base
  exit-af-topology
  network 12.1.77.0 0.0.0.255
 exit-address-family
!
```

Note that the EIGRP configuration on R4 depends on the MP-BGP VPN requirements for VRF SITE2. Like R1, R4 is also a PE in the upcoming MPLS Layer 3 VPN section.

R6:

```
router eigrp HQ-CE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  topology base
   default-metric 1500 100 3 255 1500
   redistribute eigrp 1000
  exit-af-topology

  network 12.1.66.0 0.0.0.255
  network 192.168.123.0
 exit-address-family
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
   default-metric 1500 100 3 255 1500
   redistribute eigrp 100
  exit-af-topology
  network 12.1.0.6 0.0.0.0
  network 12.1.61.0 0.0.0.255
  network 12.1.62.0 0.0.0.255
  network 12.1.67.0 0.0.0.255
 exit-address-family
 !
```

R7:

```
!
router eigrp HQ-CE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  topology base
   default-metric 1500 100 3 255 1500
   redistribute eigrp 1000
  exit-af-topology
  network 12.1.77.0 0.0.0.255
 exit-address-family
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
   default-metric 1500 100 3 255 1500
   redistribute eigrp 100
  exit-af-topology
  network 12.1.0.7 0.0.0.0
  network 12.1.67.0 0.0.0.255
  network 12.1.71.0 0.0.0.255
  network 12.1.72.0 0.0.0.255
 exit-address-family
 !
```

Note that the redistribution between EIGRP AS 100 and EIGRP AS 1000 is performed on R6 and R7 to provide redundant IPv4 EIGRP connectivity in the Headquarters network.

The following are configuration and verification examples on R8, R9, SW1, SW2, SW3, and SW4:

R8:

```
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  network 12.1.0.0 0.0.255.255
  network 192.168.8.0
  eigrp stub connected
 exit-address-family
```



```
 !
```

R9:

```
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  network 12.1.0.0 0.0.255.255
  network 192.168.9.0
  eigrp stub connected
 exit-address-family
 !
```

Note that the **eigrp stub connected** command is configured on R8 and R9 to permit advertising only connected routes.

SW1:

```
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  network 12.1.0.0 0.0.255.255
 exit-address-family
 !
```

SW2:

```
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  network 12.1.0.0 0.0.255.255
 exit-address-family
 !
```

SW3:

```
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  network 12.1.0.0 0.0.255.255
 exit-address-family
 !
```

SW4:

```
router eigrp HQ-MAIN
 !
 address-family ipv4 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  network 12.1.0.0 0.0.255.255
 exit-address-family
 !
```

EIGRP neighbors verification examples on R6 and R7:

```
R6#show ip eigrp neighbors
EIGRP-IPv4 VR(HQ-CE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   12.1.66.1               Et0/0.66                 12 1w1d       20   120  0  23
EIGRP-IPv4 VR(HQ-MAIN) Address-Family Neighbors for AS(1000)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   12.1.62.120             Et0/0.62                 13 1w1d        1   100  0  98
1   12.1.61.110             Et0/0.61                 13 1w1d        2   100  0  73
0   12.1.67.7               Et0/0.67                 14 1w1d        2   100  0  57
R6#

R7#show ip eigrp neighbors
EIGRP-IPv4 VR(HQ-CE) Address-Family Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   12.1.77.4               Et0/0.77                 12 1w1d       45   270  0  18
EIGRP-IPv4 VR(HQ-MAIN) Address-Family Neighbors for AS(1000)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   12.1.72.120             Et0/0.72                 13 1w1d       56   336  0  101
1   12.1.71.110             Et0/0.71                 13 1w1d        1   100  0  75
0   12.1.67.6               Et0/0.67                 10 1w1d       46   276  0  72
R7#
```

IPv4 reachability verification:

One way to test that your redistribution satisfies the goal of universal connectivity is to run a Tool Command Language (Tcl) script, such as the script that follows, on each router. Tcl scripting support is available in the Cisco IOS Software versions that are used here on the routers. The following simple script lists all the Loopback0 IP addresses in your EIGRP AS 100 and EIGRP AS 100 in Headquarters. It can be built once in Notepad and then pasted into each router to automate pings. A paper on Tcl scripting is available in the Network Library section of your portal.

```
tclsh
foreach address {
12.1.0.6
12.1.0.7
12.1.0.8
12.1.0.9
12.1.0.110
12.1.0.120
12.1.0.130
12.1.0.140
} {
ping $address source lo0
}
```

Run **tclsh** in privileged mode, paste the following script, and then issue the command **exit** or **tclquit.**

```
R8#tclsh
R8(tcl)#foreach address {
+>(tcl)#12.1.0.6
+>(tcl)#12.1.0.7
+>(tcl)#12.1.0.8
+>(tcl)#12.1.0.9
+>(tcl)#12.1.0.110
+>(tcl)#12.1.0.120
+>(tcl)#12.1.0.130
+>(tcl)#12.1.0.140
+>(tcl)#} {
+>(tcl)#ping $address source lo0
+>(tcl)#}
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.0.6, timeout is 2 seconds:
Packet sent with a source address of 12.1.0.8
!!!!!
```

```
<skipped for brevity>

R8(tcl)#exit
R8#
```

You also need to make sure that your solution is stable. If you have split horizon or other route feedback problems, routes may continually be inserted and removed from your routing tables. You can test stability by observing the output of the **debug IP routing** command. Finally, you need to make sure that your routes are optimal—that native prefixes are routed by native protocols and that you are using the shortest paths. This requires close examination of each routing table.

Issue:

Configure IPv4 EIGRP AS 100 on all links between R2, R5, R10, and R11 using the EIGRP classic mode

Configure IPv4 EIGRP AS 1000 on all links between R10, R11, R12, R13, and SW5 using the EIGRP classic mode.

Solution:

The following are configuration and verification examples on R2, R5, R10, and R11:

R2:

```
router eigrp 100
 !
 address-family ipv4 vrf SITE1
  redistribute bgp 10001 metric 1000 100 255 3 1500
  network 12.2.102.0 0.0.0.255
  autonomous-system 100
 exit-address-family
!
```

Note that the EIGRP configuration on R2 depends on the MP-BGP VPN requirements for VRF SITE1. This is due to the fact that R2 is a PE router in the upcoming MPLS Layer 3 VPN section.

R5:

```
router eigrp 100
 !
 address-family ipv4 vrf SITE2
  redistribute bgp 10001 metric 1000 100 255 3 1500
  network 12.2.115.0 0.0.0.255
  autonomous-system 100
 exit-address-family
!
```

Note that the EIGRP configuration on R5 depends on the MP-BGP VPN requirements for VRF SITE2. This is due to the fact that R5 is a PE router in the upcoming MPLS Layer 3 VPN section.

R10:

```
router eigrp 100
 default-metric 1000 100 3 255 1500
 network 12.2.102.0 0.0.0.255
 network 192.168.123.0
 redistribute eigrp 1000
router eigrp 1000
 default-metric 1000 100 3 255 1500
 network 12.2.0.10 0.0.0.0
 network 12.2.105.0 0.0.0.255
 redistribute eigrp 100
```

```
!
```

R11:

```
router eigrp 100
 default-metric 1000 100 3 255 1500
 network 12.2.115.0 0.0.0.255
 router eigrp 1000
 default-metric 1000 100 3 255 1500
 network 12.2.0.11 0.0.0.0
 network 12.2.105.0 0.0.0.255
 redistribute eigrp 100
```

Note that the redistribution between EIGRP AS 100 and EIGRP AS 1000 is performed on R10 and R11 to provide redundant IPv4 EIGRP connectivity in the Main Regional Office.

The following are configuration and verification examples on R12, R13, and SW5:

R12:

```
router eigrp 1000
 network 12.2.0.0 0.0.255.255
 network 192.168.12.0
 eigrp stub connected
```

R13:

```
router eigrp 1000
 network 12.2.0.0 0.0.255.255
 network 192.168.13.0
 eigrp stub connected
```

Note that the **eigrp stub connected** command is configured on R12 and R13 to permit advertising only connected routes.

SW5:

```
router eigrp 1000
 network 12.2.0.0 0.0.255.255
```

EIGRP neighbors verification examples on R10 and R11:

```
R10#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   12.2.102.2              Et0/0.102                12 1w1d        1   100  0  20
EIGRP-IPv4 Neighbors for AS(1000)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   12.2.105.150            Et0/0.105                10 1w1d        1   100  0  65
0   12.2.105.11             Et0/0.105                12 1w1d        5   100  0  28
R10#

R11#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   12.2.115.5              Et0/0.115                14 1w1d        1   100  0  17
EIGRP-IPv4 Neighbors for AS(1000)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   12.2.105.10             Et0/0.105                11 1w1d        1   100  0  26
0   12.2.105.150            Et0/0.105                12 1w1d        1   100  0  65
R11#
```

IPv4 reachability verification:

The following simple scripts can be used to test connectivity between the Loopback0 interfaces of all devices in the Main Regional Office:


```
tclsh
foreach address {
12.2.0.10
12.2.0.11
12.2.0.12
12.2.0.13
12.2.0.150
} {
ping $address source lo0
}
```

Issue:

Configure IPv4 OSPF in the Regional Sales Office on R14, R15, R16, R17, SW6, and SW7. These devices should support only the IPv4 protocol. Use the interface level commands to accomplish this task.

Solution:

Given the explicit exam specification that the OSPF devices must "support only the IPv4 protocol", the only option is to configure OSPFv2. Since OSPFv3 supports both IPv4 and IPv6, it does not meet this requirement.

The following are configuration examples on R14, R15, R16, R17, SW6, and SW7.

R14:

```
interface Loopback0
 ip address 12.3.0.14 255.255.255.255
 ip ospf 1 area 0
!
interface Ethernet0/0
 ip address 12.3.45.14 255.255.255.0
 ip ospf 1 area 0
!
router ospf 1
```

R15:

```
interface Loopback0
 ip address 12.3.0.15 255.255.255.255
 ip ospf 1 area 0
!
interface Ethernet0/0
 ip address 12.3.45.15 255.255.255.0
 ip ospf 1 area 0
!
router ospf 1
```

R16:

```
interface Loopback0
 ip address 12.3.0.16 255.255.255.255
 ip ospf 1 area 1.6.7.1
!
interface Ethernet0/0
 ip address 12.3.76.16 255.255.255.0
 ip ospf 1 area 1.6.7.1
!
interface Ethernet0/1
 ip address 192.168.16.1 255.255.255.0
 ip ospf 1 area 1.6.7.1
!
!
router ospf 1
```

R17:

```
interface Loopback0
 ip address 12.3.0.17 255.255.255.255
```

```
 ip ospf 1 area 1.6.7.1
!
interface Ethernet0/0
 ip address 12.3.76.17 255.255.255.0
 ip ospf 1 area 1.6.7.1
!
interface Ethernet0/1
 ip address 192.168.17.1 255.255.255.0
 ip ospf 1 area 1.6.7.1
!
!
router ospf 1
```

SW6:

```
interface Loopback0
 ip address 12.3.0.170 255.255.255.255
 ip ospf 1 area 0
!
interface Vlan45
 ip address 12.3.45.160 255.255.255.0
 ip ospf 1 area 0
!
interface Vlan76
 ip address 12.3.76.160 255.255.255.0
 ip ospf 1 area 1.6.7.1
!
!
router ospf 1
```

SW7:

```
interface Vlan45
 ip address 12.3.45.170 255.255.255.0
 ip ospf 1 area 0
!
interface Vlan76
 ip address 12.3.76.170 255.255.255.0
 ip ospf 1 area 1.6.7.1
!
router ospf 1
```

Note that the **ip ospf** *PID* **area** *N* interface OSPF command is used in this answer key.

The following are OSPF interface and neighbor verification examples on SW6 and SW7:

```
SW6#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               12.3.0.160/32      1     LOOP  0/0
Vl45         1     0               12.3.45.160/24     1     BDR   3/3
Vl76         1     1.6.7.1         12.3.76.160/24     1     BDR   3/3
SW6#
SW6#
SW6#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
12.3.0.14         1   FULL/DROTHER    00:00:35    12.3.45.14      Vlan45
12.3.0.15         1   FULL/DROTHER    00:00:38    12.3.45.15      Vlan45
12.3.0.170        1   FULL/DR         00:00:34    12.3.45.170     Vlan45
12.3.0.16         1   FULL/DROTHER    00:00:30    12.3.76.16      Vlan76
12.3.0.17         1   FULL/DROTHER    00:00:32    12.3.76.17      Vlan76
12.3.0.170        1   FULL/DR         00:00:36    12.3.76.170     Vlan76
SW6#

SW7#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               12.3.0.170/32      1     LOOP  0/0
Vl45         1     0               12.3.45.170/24     1     DR    3/3
Vl76         1     1.6.7.1         12.3.76.170/24     1     DR    3/3
SW7#
SW7#
SW7#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
```



```
12.3.0.14         1   FULL/DROTHER    00:00:35    12.3.45.14      Vlan45
12.3.0.15         1   FULL/DROTHER    00:00:38    12.3.45.15      Vlan45
12.3.0.160        1   FULL/BDR        00:00:35    12.3.45.160     Vlan45
12.3.0.16         1   FULL/DROTHER    00:00:30    12.3.76.16      Vlan76
12.3.0.17         1   FULL/DROTHER    00:00:32    12.3.76.17      Vlan76
12.3.0.160        1   FULL/BDR        00:00:31    12.3.76.160     Vlan76
SW7#
```

IPv4 reachability verification:

The following simple scripts can be used to connectivity between the Loopback0 interfaces of all devices in the Regional Sales Office:

```
tclsh
foreach address {
12.3.0.14
12.3.0.15
12.3.0.16
12.3.0.17
12.3.0.160
12.3.0.170
} {
ping $address source lo0
}
```

Issue:

Configure IPv4 OSPF in Regional Marketing Office on R18, R19, R20, and R21.

Redistribute the Ethernet0/1 interface into the IPv4 OSPF Area 201 on R21.

The OSPF external type “E1” or “E2” networks are not allowed in the IPv4 OSPF Area 201.

Solution:

Given the following two conditions:

(1) OSPF external type “E1” or “E2” networks are not allowed in the IPv4 OSPF Area 201 and,

(2) R21 is redistributing a connected route into Area 201,

The only option to fulfill these requirements would be to configure Area 201 as an NSSA area.

See the configuration example on R18, R19, R20, and R21:

R18:

```
!
ipv6 unicast-routing
interface Loopback0
 ip address 12.4.0.18 255.255.255.255
 ipv6 address 12:4::18/128
 ospfv3 1 ipv4 area 0
!
interface Ethernet0/0.82
 encapsulation dot1Q 82
 ip address 12.4.82.18 255.255.255.0
 ipv6 address 12:4:82::18/64
 ospfv3 1 ipv4 area 0
 !
interface Ethernet0/0.89
 encapsulation dot1Q 89
 ip address 12.4.89.18 255.255.255.0
 ipv6 address 12:4:89::18/64
 ospfv3 1 ipv4 area 0
!
router ospfv3 1
 address-family ipv4 unicast
 exit-address-family
 !
```

Note that because of the restriction “The same OSPF process must be used for the IPv6 and IPv4 routing” in the “IPv6 OSPF in Regional Marketing Office”configuration section, the **ospfv3** *PID* **ipv4 area** *N* interface OSPFv3 command is used to provide the configuration compatibility with the IPv6 OSPFv3 configuration.

R19:

```
!
ipv6 unicast-routing
interface Loopback0
ip address 12.4.0.19 255.255.255.255
 ipv6 address 12:4::19/128
 ospfv3 1 ipv4 area 0
!
interface Ethernet0/0.89
 encapsulation dot1Q 89
 ip address 12.4.89.19 255.255.255.0
 ipv6 address 12:4:89::19/64
 ospfv3 1 ipv4 area 0
!
interface Ethernet0/0.92
 encapsulation dot1Q 92
 ip address 12.4.92.19 255.255.255.0
 ipv6 address 12:4:92::19/64
 ospfv3 1 ipv4 area 0
!
router ospfv3 1
 address-family ipv4 unicast
 exit-address-family
 !
```

R20:

```
!
ipv6 unicast-routing
!
interface Loopback0
 ip address 12.4.0.20 255.255.255.255
 ipv6 address 12:4::20/128
 ospfv3 1 ipv4 area 0
!
!
interface Ethernet0/0.82
 encapsulation dot1Q 82
 ip address 12.4.82.20 255.255.255.0
 ipv6 address 12:4:82::20/64
 ospfv3 1 ipv4 area 0
!
interface Ethernet0/0.92
 encapsulation dot1Q 92
 ip address 12.4.92.20 255.255.255.0
 ip accounting output-packets
 ipv6 address 12:4:92::20/64
 ospfv3 1 ipv4 area 0
!
interface Ethernet0/0.201
 encapsulation dot1Q 201
 ip address 12.4.201.20 255.255.255.0
 ipv6 address 12:4:201::20/64
 ospfv3 1 ipv4 area 201
!
!
router ospfv3 1
!
address-family ipv4 unicast
  area 201 nssa default-information-originate
 exit-address-family
 !
```

Note that the OSPFv3 NSSA area 201 is configured to disallow The OSPF external type “E1” or “E2” networks in the IPv4 OSPF area 201. The keyword **default-information-originate** is configured to provide IPv4 connectivity between R21 and the rest of the network.

R21:

```
!
ipv6 unicast-routing
!
interface Loopback0
 ip address 12.4.0.21 255.255.255.255
 ipv6 address 12:4::21/128
 ospfv3 1 ipv4 area 201
 ospfv3 1 ipv6 area 201
!
interface Ethernet0/0
 ip address 12.4.201.21 255.255.255.0
 ipv6 address 12:4:201::21/64
 ipv6 enable
 ospfv3 1 ipv4 area 201
 ospfv3 1 ipv6 area 201
!
interface Ethernet0/1
 ip address 192.168.21.1 255.255.255.0
 ipv6 address 192:168:21::1/64
!
!
router ospfv3 1
 !
 address-family ipv4 unicast
  redistribute connected
  area 201 nssa
 exit-address-family
 !
```

The following displays OSPF interface and neighbor verification examples on R20 and R21:

```
R20#show ospfv3 ipv4 interface brief
Interface    PID   Area            AF          Cost  State Nbrs F/C
Lo0          1     0               ipv4        1     LOOP  0/0
Et0/0.92     1     0               ipv4        10    DR    1/1
Et0/0.82     1     0               ipv4        10    DR    1/1
Et0/0.201    1     201             ipv4        10    BDR   1/1
R20#

R20#show ospfv3 ipv4 neighbor

          OSPFv3 1 address-family ipv4 (router-id 12.4.0.20)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
12.4.0.19         1   FULL/BDR        00:00:33    16              Ethernet0/0.92
12.4.0.18         1   FULL/BDR        00:00:32    15              Ethernet0/0.82
12.4.0.21         1   FULL/DR         00:00:31    3               Ethernet0/0.201
R20#

R21#show ospfv3 ipv4 database nssa-external self-originate

          OSPFv3 1 address-family ipv4 (router-id 12.4.0.21)

                Type-7 AS External Link States (Area 201)

  LS age: 761
  LS Type: AS External Link
  Link State ID: 0
  Advertising Router: 12.4.0.21
  LS Seq Number: 8000017F
  Checksum: 0xB2F1
  Length: 48
  Prefix Address: 192.168.21.0
  Prefix Length: 24, Options: P
  Metric Type: 2 (Larger than any link state path)
  Metric: 20
  Forward Address: 12.4.0.21

R21#show ip route ospfv3
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 12.4.201.20 to network 0.0.0.0

O*N2  0.0.0.0/0 [110/1] via 12.4.201.20, 1w1d, Ethernet0/0
      12.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
O IA     12.4.0.18/32 [110/20] via 12.4.201.20, 1w1d, Ethernet0/0
O IA     12.4.0.19/32 [110/20] via 12.4.201.20, 1w1d, Ethernet0/0
O IA     12.4.0.20/32 [110/10] via 12.4.201.20, 1w1d, Ethernet0/0
O IA     12.4.82.0/24 [110/20] via 12.4.201.20, 1w1d, Ethernet0/0
O IA     12.4.89.0/24 [110/30] via 12.4.201.20, 1w1d, Ethernet0/0
O IA     12.4.92.0/24 [110/20] via 12.4.201.20, 1w1d, Ethernet0/0
R21#
```

IPv4 reachability verification:

The following simple scripts can be used to test connectivity between the Loopback0 interfaces of all devices in Regional Sales Office:

```
tclsh
foreach address {
12.4.0.18
12.4.0.19
12.4.0.20
12.4.0.21
} {
ping $address source lo0
}
```

Issue:

Configure IPv4 BGP AS 12000 according to the lab requirements.

Solution:

BGP configuration example on R6:

```
!
router bgp 12000
 bgp log-neighbor-changes
 bgp default local-preference 101
 neighbor 12.2.0.10 remote-as 12000
 neighbor 12.2.0.10 update-source Loopback0
neighbor 123.1.6.1 remote-as 30000
 !
 address-family ipv4
  network 12.1.0.6 mask 255.255.255.255
  aggregate-address 12.1.0.0 255.255.0.0 summary-only
  neighbor 12.2.0.10 activate
  neighbor 12.2.0.10 next-hop-self
  neighbor 123.1.6.1 activate
 exit-address-family
 !
```

Note that because of the lab restriction “Do not use prefix lists and route maps to accomplish this task”, the **bgp default local-preference** command is configured on R6 to prefer the Internet prefixes via R6 in the BGP AS 12000. Specifically these prefixes are:

- Prefix 123.1.0.0/16 that is originated from the Internet BGP AS 30000 and is used to provide connectivity between the IP addresses of the sources of the DMVPN tunnels.
- Summary prefixes 12.3.0.0/16 and 12.4.0.0/16 that represent Regional Sales and Marketing offices respectively.

Also note that after you configure the BGP peer relationship between R6 and the Internet BGP AS 30000 you notice that R6 does not learn the default route 0.0.0.0/0 from the Internet BGP AS 30000.

Only R10 learns the default route 0.0.0.0/0 from the Internet BGP AS 30000 in this lab. R6 should learn the default route 0.0.0.0/0 from R10 as displayed in the following output:

```
R6#show ip bgp regexp _30000_
BGP table version is 10, local router ID is 12.1.0.6
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
 r>i 0.0.0.0          12.2.0.10                0    100      0 30000 i
 *>  12.3.0.0/16     123.1.6.1              200             0 30000 12300 i
 *>  12.4.0.0/16     123.1.6.1              200             0 30000 12400 i
 *>  123.1.0.0/16    123.1.6.1              200             0 30000 i
R6#
```

The following is a BGP configuration example on R10:

```
!
router bgp 12000
 bgp log-neighbor-changes
 network 12.2.0.10 mask 255.255.255.255
 aggregate-address 12.2.0.0 255.255.0.0 summary-only
 neighbor 12.1.0.6 remote-as 12000
 neighbor 12.1.0.6 update-source Loopback0
 neighbor 12.1.0.6 next-hop-self
 neighbor 123.1.10.1 remote-as 30000
!
```

Note that the BGP peer relationship between R6 and R10 will be established only after you configure MP-BGP VPN.

Verify the BGP configuration on R6 and R10:

```
R6#show bgp ipv4 unicast summary
BGP router identifier 12.1.0.6, local AS number 12000
BGP table version is 12, main routing table version 12
7 network entries using 980 bytes of memory
7 path entries using 532 bytes of memory
7/7 BGP path/bestpath attribute entries using 952 bytes of memory
3 BGP AS-PATH entries using 72 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 2536 total bytes of memory
BGP activity 13/0 prefixes, 13/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.2.0.10       4        12000   14136   14147       12    0    0 1w1d            2
123.1.6.1       4        30000   14145   14133       12    0    0 1w1d            3
R6#
R6#show bgp ipv4 unicast regexp ^$
BGP table version is 12, local router ID is 12.1.0.6
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
 *>  12.1.0.0/16      0.0.0.0                            32768 i
 s>  12.1.0.6/32      0.0.0.0                  0         32768 i
 *>i 12.2.0.0/16      12.2.0.10                0    100      0 i
R6#
```

Note that only one prefix 12.1.0.0/16 is originated on R6, more specific prefixes are suppressed.

```
R10#show bgp ipv4 unicast summary
BGP router identifier 12.2.0.10, local AS number 12000
BGP table version is 11, main routing table version 11
7 network entries using 980 bytes of memory
10 path entries using 760 bytes of memory
10/7 BGP path/bestpath attribute entries using 1360 bytes of memory
3 BGP AS-PATH entries using 72 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 3172 total bytes of memory
BGP activity 7/0 prefixes, 10/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.1.0.6        4        12000   14147   14136       11    0    0 1w1d            4
123.1.10.1      4        30000   14139   14150       11    0    0 1w1d            4
R10#show bgp ipv4 unicast regexp ^$
BGP table version is 11, local router ID is 12.2.0.10
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
 *>i 12.1.0.0/16      12.1.0.6                 0    101      0 i
 *>  12.2.0.0/16      0.0.0.0                                32768 i
 s>  12.2.0.10/32     0.0.0.0                  0             32768 i
R10#
```

Note that only one prefix 12.2.0.0/16 is originated on R10, more specific prefixes are suppressed.

Issue:

Configure IPv4 BGP AS 12300 and AS 12400 according to the lab requirements.

Solution:

The following are configuration examples on R14 and R15:

R14:

```
router ospf 1
 redistribute bgp 12300 subnets
!
router bgp 12300
 bgp log-neighbor-changes
 aggregate-address 12.3.0.0 255.255.0.0 summary-only
 redistribute ospf 1
 neighbor 12.3.0.15 remote-as 12300
 neighbor 12.3.0.15 update-source Loopback0
 neighbor 12.3.0.15 next-hop-self
 neighbor 123.1.14.1 remote-as 30000
!
```

R15:

```
router bgp 12300
 bgp log-neighbor-changes
 aggregate-address 12.3.0.0 255.255.0.0 summary-only
 redistribute ospf 1
 neighbor 12.3.0.14 remote-as 12300
 neighbor 12.3.0.14 update-source Loopback0
 neighbor 12.3.0.14 next-hop-self
 neighbor 123.1.15.1 remote-as 30000
!
```

Verify the BGP configuration on R14 and R15:

```
R14#show bgp ipv4 unicast summary
BGP router identifier 12.3.0.14, local AS number 12300
BGP table version is 30, main routing table version 30
16 network entries using 2240 bytes of memory
24 path entries using 1824 bytes of memory
17/10 BGP path/bestpath attribute entries using 2312 bytes of memory
3 BGP AS-PATH entries using 72 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 6448 total bytes of memory
BGP activity 16/0 prefixes, 25/1 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.3.0.15       4        12300   14282   14285       30    0    0 1w2d            8
123.1.14.1      4        30000   14291   14267       30    0    0 1w2d            5
R14#
R14#show bgp ipv4 unicast regexp ^$
BGP table version is 30, local router ID is 12.3.0.14
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
 r i 12.3.0.0/16      12.3.0.15                0    100      0 i
 r>                   0.0.0.0                                32768 i
 s>  12.3.0.14/32     0.0.0.0                  0             32768 ?
 s>  12.3.0.15/32     12.3.45.15              11             32768 ?
 s>  12.3.0.16/32     12.3.45.160             12             32768 ?
 s>  12.3.0.17/32     12.3.45.160             12             32768 ?
 s>  12.3.0.160/32    12.3.45.160             11             32768 ?
 s>  12.3.0.170/32    12.3.45.170             11             32768 ?
 s>  12.3.45.0/24     0.0.0.0                  0             32768 ?
 s>  12.3.76.0/24     12.3.45.160             11             32768 ?
 * i 192.168.16.0     12.3.0.15               21    100      0 ?
 *>                   12.3.45.160             21             32768 ?
 * i 192.168.17.0     12.3.0.15               21    100      0 ?
 *>                   12.3.45.160             21             32768 ?
R14#
```

Note that the prefix 12.3.0.0/16 is originated on R14 and more specific prefixes are suppressed. Also, the Ethernet0/1 networks of R16 and R17 are advertised too.

The following are configuration examples on R18 and R19:

```
R15#show bgp ipv4 unicast summary
BGP router identifier 12.3.0.15, local AS number 12300
BGP table version is 31, main routing table version 31
16 network entries using 2240 bytes of memory
24 path entries using 1824 bytes of memory
17/10 BGP path/bestpath attribute entries using 2312 bytes of memory
3 BGP AS-PATH entries using 72 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 6448 total bytes of memory
BGP activity 16/0 prefixes, 25/1 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.3.0.14       4        12300   14288   14285       31    0    0 1w2d            8
123.1.15.1      4        30000   14285   14272       31    0    0 1w2d            5
R15#
R15#show bgp ipv4 unicast regexp ^$
BGP table version is 31, local router ID is 12.3.0.15
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
 *>  12.3.0.0/16      0.0.0.0                                32768 i
 * i                   12.3.0.14                0    100      0 i
 s>  12.3.0.14/32     12.3.45.14              11             32768 ?
 s>  12.3.0.15/32     0.0.0.0                  0             32768 ?
 s>  12.3.0.16/32     12.3.45.160             12             32768 ?
 s>  12.3.0.17/32     12.3.45.160             12             32768 ?
 s>  12.3.0.160/32    12.3.45.160             11             32768 ?
 s>  12.3.0.170/32    12.3.45.170             11             32768 ?
 s>  12.3.45.0/24     0.0.0.0                  0             32768 ?
```

```
 s>  12.3.76.0/24     12.3.45.160              11         32768 ?
 * i 192.168.16.0     12.3.0.14                21    100      0 ?
 *>                   12.3.45.160              21         32768 ?
 * i 192.168.17.0     12.3.0.14                21    100      0 ?
 *>                   12.3.45.160              21         32768 ?
R15#
```

Note that the prefix 12.3.0.0/16 is originated on R14 and more specific prefixes are suppressed. Also the Ethernet0/1 networks of R16 and R17 are advertised too.

Configuration examples on R18 and R19:

R18:

```
router bgp 12400
 bgp log-neighbor-changes
 neighbor 12.4.0.19 remote-as 12400
 neighbor 12.4.0.19 update-source Loopback0
 neighbor 123.1.18.1 remote-as 30000
 !
 address-family ipv4
  aggregate-address 12.4.0.0 255.255.0.0 summary-only
  redistribute ospfv3 1
  neighbor 12.4.0.19 activate
  neighbor 12.4.0.19 next-hop-self
  neighbor 123.1.18.1 activate
 exit-address-family
 !
router ospfv3 1
 !
 address-family ipv4 unicast
  redistribute bgp 12400
 exit-address-family
 !
```

R19:

```
router bgp 12400
 bgp log-neighbor-changes
 neighbor 12.4.0.18 remote-as 12400
 neighbor 12.4.0.18 update-source Loopback0
 neighbor 123.1.19.1 remote-as 30000
 !
 address-family ipv4
  aggregate-address 12.4.0.0 255.255.0.0 summary-only
  redistribute ospfv3 1
  neighbor 12.4.0.18 activate
  neighbor 12.4.0.18 next-hop-self
  neighbor 123.1.19.1 activate
 exit-address-family
 !
router ospfv3 1
 !
 address-family ipv4 unicast
  redistribute bgp 12400
 exit-address-family
 !
```

Verify the BGP configuration on R18 and R19:

```
R18#show bgp ipv4 unicast summary
BGP router identifier 12.4.0.18, local AS number 12400
BGP table version is 70, main routing table version 70
16 network entries using 2240 bytes of memory
25 path entries using 1900 bytes of memory
17/9 BGP path/bestpath attribute entries using 2312 bytes of memory
3 BGP AS-PATH entries using 72 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 6524 total bytes of memory
BGP activity 21/0 prefixes, 48/14 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.4.0.19       4        12400   14309   14323       70    0    0 1w2d            8
123.1.18.1      4        30000   14298   14291       70    0    0 1w2d            6
```

```
R18#
R18#show bgp ipv4 unicast regexp ^$
BGP table version is 70, local router ID is 12.4.0.18
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
 *>  12.4.0.0/16      12.4.89.19               1         32768 ?
 *                    0.0.0.0                            32768 i
 * i                  12.4.0.19                0    100      0 i
 s>  12.4.0.18/32     0.0.0.0                  0         32768 ?
 s>  12.4.0.19/32     12.4.89.19              10         32768 ?
 s>  12.4.0.20/32     12.4.82.20              10         32768 ?
 s>  12.4.0.21/32     12.4.82.20              20         32768 ?
 s>  12.4.82.0/24     0.0.0.0                  0         32768 ?
 s>  12.4.89.0/24     0.0.0.0                  0         32768 ?
 s>  12.4.92.0/24     12.4.82.20              20         32768 ?
 s>  12.4.201.0/24    12.4.82.20              20         32768 ?
 * i 192.168.21.0     12.4.0.19                20    100      0 ?
 *>                   12.4.82.20               20         32768 ?
R18#

R19#show bgp ipv4 unicast summary
BGP router identifier 12.4.0.19, local AS number 12400
BGP table version is 67, main routing table version 67
16 network entries using 2240 bytes of memory
24 path entries using 1824 bytes of memory
16/9 BGP path/bestpath attribute entries using 2176 bytes of memory
3 BGP AS-PATH entries using 72 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 6312 total bytes of memory
BGP activity 21/0 prefixes, 47/14 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.4.0.18       4        12400   14323   14309       67    0    0 1w2d            8
123.1.19.1      4        30000   14275   14278       67    0    0 1w2d            6
R19#
R19#show bgp ipv4 unicast regexp ^$
BGP table version is 67, local router ID is 12.4.0.19
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
 * i 12.4.0.0/16      12.4.0.18                1    100      0 ?
 *>                   0.0.0.0                            32768 i
 s>  12.4.0.18/32     12.4.89.18              10         32768 ?
 s>  12.4.0.19/32     0.0.0.0                  0         32768 ?
 s>  12.4.0.20/32     12.4.92.20              10         32768 ?
 s>  12.4.0.21/32     12.4.92.20              20         32768 ?
 s>  12.4.82.0/24     12.4.89.18              20         32768 ?
 s>  12.4.89.0/24     0.0.0.0                  0         32768 ?
 s>  12.4.92.0/24     0.0.0.0                  0         32768 ?
 s>  12.4.201.0/24    12.4.92.20              20         32768 ?
 * i 192.168.21.0     12.4.0.18                20    100      0 ?
 *>                   12.4.92.20               20         32768 ?
R19#
```

Note that the prefix 12.4.0.0/16 is originated on R18 and more specific prefixes are suppressed. Also the Ethernet0/1 network of R21 is advertised too.

```
R19#show bgp ipv4 unicast summary
BGP router identifier 12.4.0.19, local AS number 12400
BGP table version is 67, main routing table version 67
16 network entries using 2240 bytes of memory
24 path entries using 1824 bytes of memory
16/9 BGP path/bestpath attribute entries using 2176 bytes of memory
```

```
3 BGP AS-PATH entries using 72 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 6312 total bytes of memory
BGP activity 21/0 prefixes, 47/14 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.4.0.18       4        12400   14323   14309       67    0    0 1w2d            8
123.1.19.1      4        30000   14275   14278       67    0    0 1w2d            6
R19#
R19#show bgp ipv4 unicast regexp ^$
BGP table version is 67, local router ID is 12.4.0.19
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
 * i 12.4.0.0/16      12.4.0.18                1    100      0 ?
 *>                   0.0.0.0                               32768 i
 s>  12.4.0.18/32     12.4.89.18              10             32768 ?
 s>  12.4.0.19/32     0.0.0.0                  0             32768 ?
 s>  12.4.0.20/32     12.4.92.20              10             32768 ?
 s>  12.4.0.21/32     12.4.92.20              20             32768 ?
 s>  12.4.82.0/24     12.4.89.18              20             32768 ?
 s>  12.4.89.0/24     0.0.0.0                  0             32768 ?
 s>  12.4.92.0/24     0.0.0.0                  0             32768 ?
 s>  12.4.201.0/24    12.4.92.20              20             32768 ?
 * i 192.168.21.0     12.4.0.18               20    100      0 ?
 *>                   12.4.92.20              20             32768 ?
R19#
```

Note that the prefix 12.4.0.0/16 is originated on R19 and more specific prefixes are suppressed. Also the Ethernet0/1 network of R21 is advertised too.

Issue:

Configure IPv4 BGP AS 65022 and AS 65023 according to the lab requirements.

Solution:

The following are configuration examples on R22 and R23:

R22:

```
!
router bgp 65022
 bgp log-neighbor-changes
 neighbor 123.1.22.1 remote-as 30000
!
```

R23:

```
!
router bgp 65023
 bgp log-neighbor-changes
 neighbor 123.1.23.1 remote-as 30000
!
```

BGP verification examples on R22 and R23:

```
R22#show bgp ipv4 unicast summary
BGP router identifier 192.168.123.22, local AS number 65022
BGP table version is 11, main routing table version 11
8 network entries using 1120 bytes of memory
8 path entries using 608 bytes of memory
7/7 BGP path/bestpath attribute entries using 952 bytes of memory
4 BGP AS-PATH entries using 96 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 2776 total bytes of memory
BGP activity 8/0 prefixes, 8/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
123.1.22.1      4        30000   14297   14291       11    0    0 1w2d            8
R22#
R22#show bgp ipv4 unicast regexp ^$
R22#
```

Note that R22 does not advertise any BGP prefixes.

```
R23#show bgp ipv4 unicast summary
BGP router identifier 192.168.123.23, local AS number 65023
BGP table version is 11, main routing table version 11
8 network entries using 1120 bytes of memory
8 path entries using 608 bytes of memory
7/7 BGP path/bestpath attribute entries using 952 bytes of memory
4 BGP AS-PATH entries using 96 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 2776 total bytes of memory
BGP activity 8/0 prefixes, 8/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
123.1.23.1      4        30000   14297   14291       11    0    0 1w2d            8
R23#
R23#show bgp ipv4 unicast regexp ^$
R23#
```

Note that R23 does not advertise any BGP prefixes.

Issue:

Configure IPv6 EIGRP named instance **HQ-MAIN** in Headquarters according to the lab requirements.

Solution:

When fulfilling the IPv4 uncast connectivity requirements of this lab in an earlier section, EIGRP named mode is already configured on these routers. All that is needed is to add an IPv6 address-family and the following commands.

The following are configuration examples on R6, R7, R8, R9, SW1, SW2, SW3, and SW4:

R6:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !

 !
 address-family ipv6 unicast autonomous-system 1000
  !
  af-interface Serial1/0
   shutdown
  exit-af-interface
  !
  topology base
   redistribute bgp 12000 metric 1000 100 255 3 1500
  exit-af-topology
 exit-address-family
!
```

Note that the redistribution of IPv6 BGP into IPv6 EIGRP is configured to provide IPv6 connectivity between the Headquarters and the Regional Marketing Office.

Note that the Serial1/0 interface is excluded from the IPv6 EIGRP 1000 on R6.

R7:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
 exit-address-family
```

R8:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  eigrp stub connected
 exit-address-family
```

R9:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
  eigrp stub connected
 exit-address-family
```

SW1:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
 exit-address-family
```

SW2:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
 exit-address-family
```

SW3:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
 exit-address-family
```



```
!
```

SW4:

```
ipv6 unicast-routing
!
router eigrp HQ-MAIN
 !
address-family ipv6 unicast autonomous-system 1000
  !
  topology base
  exit-af-topology
 exit-address-family
!
```

IPv6 EIGRP verification examples on R6 and R7:

```
R6#show ipv6 eigrp interfaces
EIGRP-IPv6 VR(HQ-MAIN) Address-Family Interfaces for AS(1000)
                              Xmit Queue   PeerQ        Mean   Pacing Time   Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Et0/0.61                 1        0/0       0/0        1023       0/2          5116          0
Et0/0.62                 1        0/0       0/0        1021       0/2          5108          0
Et0/0.67                 1        0/0       0/0         820       0/2          4092          0
Lo0                      0        0/0       0/0           0       0/0             0          0
R6#
R6#show ipv6 eigrp neighbors
EIGRP-IPv6 VR(HQ-MAIN) Address-Family Neighbors for AS(1000)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   Link-local address:     Et0/0.62               13 1w2d     1021  5000  0  25
    FE80::A8BB:CCFF:FE80:1900
1   Link-local address:     Et0/0.61               12 1w2d     1023  5000  0  22
    FE80::A8BB:CCFF:FE80:1800
0   Link-local address:     Et0/0.67               11 1w2d      820  4920  0  27
    FE80::A8BB:CCFF:FE00:700
R6#

R7#show ipv6 eigrp interfaces
EIGRP-IPv6 VR(HQ-MAIN) Address-Family Interfaces for AS(1000)
                              Xmit Queue   PeerQ        Mean   Pacing Time   Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Et0/0.67                 1        0/0       0/0           2       0/2            50          0
Et0/0.71                 1        0/0       0/0           2       0/2            50          0
Et0/0.72                 1        0/0       0/0        1024       0/2          5116          0
Lo0                      0        0/0       0/0           0       0/0             0          0
R7#
R7#show ipv6 eigrp neighbors
EIGRP-IPv6 VR(HQ-MAIN) Address-Family Neighbors for AS(1000)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   Link-local address:     Et0/0.72               13 1w2d     1024  5000  0  23
    FE80::A8BB:CCFF:FE80:1900
1   Link-local address:     Et0/0.71               10 1w2d        2   100  0  21
    FE80::A8BB:CCFF:FE80:1800
0   Link-local address:     Et0/0.67               14 1w2d        2   100  0  14
    FE80::A8BB:CCFF:FE00:600
R7#
```

The following simple script lists all the Loopback0 IPv6 EIGRP addresses in Headquarters:

```
tclsh
foreach address {
12:1::6
12:1::7
12:1::8
12:1::9
12:1::110
12:1::120
12:1::130
12:1::140
} {
```

```
ping $address source lo0
}
```

Issue:

Configure IPv6 OSPF in Regional Marketing Office according to the lab requirements.

Solution:

The following configuration examples on R18, R19, R20, and R21:

R18:

```
ipv6 unicast-routing
!
interface Loopback0
 ip address 12.4.0.18 255.255.255.255
 ipv6 address 12:4::18/128
 ospfv3 1 ipv6 area 0
!
interface Ethernet0/0
 no ip address
!
interface Ethernet0/0.82
 encapsulation dot1Q 82
 ip address 12.4.82.18 255.255.255.0
 ipv6 address 12:4:82::18/64
 ospfv3 1 ipv6 area 0
!
interface Ethernet0/0.89
 encapsulation dot1Q 89
 ip address 12.4.89.18 255.255.255.0
 ipv6 address 12:4:89::18/64
 ospfv3 1 ipv6 area 0
!
router ospfv3 1
 !
!
 address-family ipv6 unicast
  redistribute bgp 12400
 exit-address-family
!
```

Note that the redistribution of IPv6 BGP into IPv6 OSPF is configured to provide IPv6 connectivity between the Headquarters and the Regional Marketing Office.

R19:

```
ipv6 unicast-routing
!
interface Loopback0
 ip address 12.4.0.19 255.255.255.255
 ipv6 address 12:4::19/128
 ospfv3 1 ipv6 area 0
!
!
interface Ethernet0/0.89
 encapsulation dot1Q 89
 ip address 12.4.89.19 255.255.255.0
 ipv6 address 12:4:89::19/64
 ospfv3 1 ipv6 area 0
!
interface Ethernet0/0.92
 encapsulation dot1Q 92
 ip address 12.4.92.19 255.255.255.0
 ipv6 address 12:4:92::19/64
 ospfv3 1 ipv6 area 0
!
router ospfv3 1
 !
 !
 address-family ipv6 unicast
  redistribute bgp 12400
 exit-address-family
!
```


Note that the redistribution of IPv6 BGP into IPv6 OSPF is configured to provide IPv6 connectivity between the Headquarters and the Regional Marketing Office.

R20:

```
ipv6 unicast-routing
!
interface Loopback0
 ip address 12.4.0.20 255.255.255.255
 ipv6 address 12:4::20/128
 ospfv3 1 ipv6 area 0
!
!
interface Ethernet0/0.82
 encapsulation dot1Q 82
 ip address 12.4.82.20 255.255.255.0
 ip accounting output-packets
 ipv6 address 12:4:82::20/64
 ospfv3 1 ipv6 area 0
!
interface Ethernet0/0.92
 encapsulation dot1Q 92
 ip address 12.4.92.20 255.255.255.0
 ip accounting output-packets
 ipv6 address 12:4:92::20/64
 ospfv3 1 ipv6 area 0
!
interface Ethernet0/0.201
 encapsulation dot1Q 201
 ip address 12.4.201.20 255.255.255.0
 ipv6 address 12:4:201::20/64
 ospfv3 1 ipv6 area 201
!
router ospfv3 1
 !
 address-family ipv6 unicast
  area 201 nssa default-information-originate
 exit-address-family
!
```

Note that the OSPFv3 NSSA area 201 is configured to disallow the OSPF external type “E1” or “E2” networks in IPv6 OSPF area 201. The keyword **default-information-originate** is configured to provide IPv6 connectivity between R21 and the rest of the network.

R21:

```
ipv6 unicast-routing
!
interface Loopback0
 ip address 12.4.0.21 255.255.255.255
 ipv6 address 12:4::21/128
 ospfv3 1 ipv6 area 201
!
interface Ethernet0/0
 ip address 12.4.201.21 255.255.255.0
 ipv6 address 12:4:201::21/64
 ospfv3 1 ipv6 area 201
!
interface Ethernet0/1
 ip address 192.168.21.1 255.255.255.0
 ipv6 address 192:168:21::1/64
!
router ospfv3 1
!
 address-family ipv6 unicast
  redistribute connected
  area 201 nssa
 exit-address-family
!
```

The following are OSPF interface and neighbor verification examples on R20 and R21:

```
R20#show ospfv3 ipv6 interface brief
Interface    PID   Area            AF          Cost  State Nbrs F/C
Lo0          1     0               ipv6        1     LOOP  0/0
Et0/0.92     1     0               ipv6        10    DR    1/1
Et0/0.82     1     0               ipv6        10    DR    1/1
Et0/0.201    1     201             ipv6        10    BDR   1/1
R20#
R20#show ospfv3 ipv6 neighbor

          OSPFv3 1 address-family ipv6 (router-id 12.4.0.20)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
12.4.0.19         1   FULL/BDR        00:00:33    16              Ethernet0/0.92
12.4.0.18         1   FULL/BDR        00:00:39    15              Ethernet0/0.82
12.4.0.21         1   FULL/DR         00:00:34    3               Ethernet0/0.201
R20#

R21#show ospfv3 ipv6 database nssa-external self-originate

          OSPFv3 1 address-family ipv6 (router-id 12.4.0.21)

               Type-7 AS External Link States (Area 201)

  LS age: 995
  LS Type: AS External Link
  Link State ID: 0
  Advertising Router: 12.4.0.21
  LS Seq Number: 8000018A
  Checksum: 0xF1C9
  Length: 52
  Prefix Address: 192:168:21::
  Prefix Length: 64, Options: P
  Metric Type: 2 (Larger than any link state path)
  Metric: 20
  Forward Address: 12:4::21

R21#
R21#show ipv6 route ospf
IPv6 Routing Table - default - 13 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
ON2 ::/0 [110/1]
     via 12:4:201::20, Ethernet0/0
OI  12:4::18/128 [110/20]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
OI  12:4::19/128 [110/20]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
OI  12:4::20/128 [110/10]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
OI  12:4:82::/64 [110/20]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
OI  12:4:89::/64 [110/30]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
OI  12:4:92::/64 [110/20]
     via FE80::A8BB:CCFF:FE00:1400, Ethernet0/0
R21#
```

The following simple script lists all the Loopback0 IPv6 OSPF addresses in the Regional Marketing Office:

```
tclsh
foreach address {
12:4::18
12:4::19
12:4::20
```



```
12:4::21
} {
ping $address source lo0
}
```

Issue:

Configure IPv6 BGP according to the lab requirements.

Solution:

The following are configuration examples on R6, R18, and R19:

R6:

```
router bgp 12000
 bgp log-neighbor-changes
 neighbor 2000:CC1E:CAB:6::1 remote-as 30000
 !
!
 address-family ipv6
  network 12:1::6/128
  network 192:168:8::/64
  aggregate-address 192:168:8::/47 summary-only
  aggregate-address 12:1::/32 summary-only
  neighbor 2000:CC1E:CAB:6::1 activate
 exit-address-family
!
!
```

R18:

```
router bgp 12400
 bgp log-neighbor-changes
 neighbor 12:4::19 remote-as 12400
 neighbor 12:4::19 update-source Loopback0
 neighbor 2000:CC1E:CAB:18::1 remote-as 30000
 !
address-family ipv6
  network 12:4::18/128
  network 192:168:21::/64
  aggregate-address 12:4::/32 summary-only
  neighbor 12:4::19 activate
  neighbor 12:4::19 next-hop-self
  neighbor 2000:CC1E:CAB:18::1 activate
 exit-address-family
!
```

R19:

```
router bgp 12400
 bgp log-neighbor-changes
 neighbor 12:4::18 remote-as 12400
 neighbor 12:4::18 update-source Loopback0
 neighbor 2000:CC1E:CAB:19::1 remote-as 30000
 !
!
 address-family ipv6
  network 12:4::19/128
  network 192:168:21::/64
  aggregate-address 12:4::/32 summary-only
  neighbor 12:4::18 activate
  neighbor 12:4::18 next-hop-self
  neighbor 2000:CC1E:CAB:19::1 activate
 exit-address-family
!
```

IPv6 BGP verification examples on R6, R18 and R19:

```
R6#show bgp ipv6 unicast summary
BGP router identifier 12.1.0.6, local AS number 12000
BGP table version is 9, main routing table version 9
6 network entries using 984 bytes of memory
6 path entries using 600 bytes of memory
```

```
5/5 BGP path/bestpath attribute entries using 680 bytes of memory
3 BGP AS-PATH entries using 72 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 2336 total bytes of memory
BGP activity 13/0 prefixes, 13/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2000:CC1E:CAB:6::1
                4        30000   14538   14543        9    0    0 1w2d            2
R6#show bgp ipv6 unicast neighbors 2000:CC1E:CAB:6::1 advertised-routes
BGP table version is 9, local router ID is 12.1.0.6
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
 *>  12:1::/32        ::                                   32768 i
 *>  192:168:8::/47   ::                                   32768 i

Total number of prefixes 2
R6#
```

Note that R6 advertises the IPv6 summary prefixes for the Loopback0 and Ethernet0/1 interfaces of Headquarters.

```
R18#show bgp ipv6 unicast summary
BGP router identifier 12.4.0.18, local AS number 12400
BGP table version is 7, main routing table version 7
5 network entries using 820 bytes of memory
9 path entries using 900 bytes of memory
7/4 BGP path/bestpath attribute entries using 952 bytes of memory
3 BGP AS-PATH entries using 72 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 2744 total bytes of memory
BGP activity 21/0 prefixes, 48/14 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12:4::19        4        12400   14534   14547        7    0    0 1w2d            4
2000:CC1E:CAB:18::1
                4        30000   14541   14540        7    0    0 1w2d            2
R18#
R18#show bgp ipv6 unicast neighbors 2000:CC1E:CAB:18::1 advertised-routes
BGP table version is 7, local router ID is 12.4.0.18
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
 r>  12:4::/32        ::                                   32768 i
 *>  192:168:21::/64  FE80::A8BB:CCFF:FE00:1400
                                          20         32768 i

Total number of prefixes 2
R18#
```

Note that R6 advertises the IPv6 summary prefixes for the Loopback0 interfaces and the Ethernet0/1 interface of the Regional Marketing office.

The following simple scripts list all the IPv6 addresses in Headquarters and Regional Marketing Office that are required for a connectivity test in accordance with the lab requirements:

```
tclsh
foreach address {
12:1::6
12:1::7
12:1::8
```



```
12:1::9
12:1::110
12:1::120
12:1::130
12:1::140
12:4::18
12:4::19
12:4::20
12:4::21
} {
ping $address source lo0
}

tclsh
foreach address {
192:168:8::1
192:168:9::1
192:168:21::1
} {
ping $address source e0/1
}
```

Issue:

Configure the multicast routing according to the lab requirements.

Solution:

This multicast task required the configuration of bidirectional PIM for two reasons:

(1) the lab requirements stated that "no multicast routers should display (S,G) state in the respective multicast routing tables in this lab" and,

(2) bidirectional PIM is explicitly listed in the lower left corner of the supplied multicast diagram.

The following are configuration examples on R6, R22, and R23:

R6:

```
ip multicast-routing
!
interface Tunnel192
 bandwidth 100000
 ip address 192.168.123.6 255.255.255.0
 ip pim nbma-mode
 ip pim sparse-mode
 ip nhrp map multicast dynamic
!
ip pim bidir-enable
ip pim rp-address 192.168.123.6 MACL bidir
!
ip access-list standard MACL
 permit 232.12.12.12
```

R22:

```
ip multicast-routing
!
interface Tunnel192
 bandwidth 100000
 ip address 192.168.123.6 255.255.255.0
 ip pim nbma-mode
 ip pim sparse-mode
 ip nhrp map multicast 123.1.6.2
!
interface Ethernet0/1
 ip address 192.168.22.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 232.12.12.12
```

```
!
ip pim bidir-enable
ip pim rp-address 192.168.123.6 MACL bidir
!
ip access-list standard MACL
 permit 232.12.12.12
```

R23:

```
ip multicast-routing
!
interface Tunnel192
 bandwidth 100000
 ip address 192.168.123.6 255.255.255.0
 ip pim nbma-mode
 ip pim sparse-mode
 ip nhrp map multicast 123.1.6.2
!
!
interface Ethernet0/1
 ip address 192.168.23.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 232.12.12.12
!
ip pim bidir-enable
ip pim rp-address 192.168.123.6 MACL bidir
!
ip access-list standard MACL
 permit 232.12.12.12

!
```

IPv4 multicast verification examples on R6, R22, and R23:

R6:

```
R6#show ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.123.6    Tunnel192                v2/S   2      30     1      192.168.123.23
R6#
R6#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
192.168.123.23    Tunnel192                23:05:26/00:01:41 v2    1 / DR B S P G
192.168.123.22    Tunnel192                1w2d/00:01:22     v2    1 / B S P G
R6#
R6#
R6#show ip pim rp map
PIM Group-to-RP Mappings

Acl: MACL, Static, Bidir Mode
    RP: 192.168.123.6 (?)
R6#
R6#show ip mroute bidirectional
(*,232.12.12.12), 1w2d/-, RP 192.168.123.6, flags: B
  Bidir-Upstream: Tunnel192, RPF nbr: 0.0.0.0
  Incoming interface list:
    Tunnel192, Accepting/Sparse
  Outgoing interface list:
    Tunnel192, 192.168.123.22, Forward/Sparse, 1w2d/00:03:17
    Tunnel192, 192.168.123.23, Forward/Sparse, 1w2d/00:03:15

(*, 232.12.12.12), 1w2d/00:03:17, RP 192.168.123.6, flags: B
  Bidir-Upstream: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel192, 192.168.123.22, Forward/Sparse, 1w2d/00:03:17
    Tunnel192, 192.168.123.23, Forward/Sparse, 1w2d/00:03:15

R6#
```


R22:

```
R22#show ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.123.22   Tunnel192                v2/S   1      30     1      192.168.123.22
192.168.22.1     Ethernet0/1              v2/S   0      30     1      192.168.22.1
R22#
R22#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
192.168.123.6     Tunnel192                1w2d/00:01:44     v2    1 / B S P G
R22#
R22#show ip pim rp map
PIM Group-to-RP Mappings

Acl: MACL, Static, Bidir Mode
    RP: 192.168.123.6 (?)
R22#
R22#show ip mroute bidirectional
(*,232.12.12.12), 1w2d/-, RP 192.168.123.6, flags: BCL
  Bidir-Upstream: Tunnel192, RPF nbr: 192.168.123.6
  Incoming interface list:
    Ethernet0/1, Accepting/Sparse
    Tunnel192, Accepting/Sparse
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse, 1w2d/00:02:05
    Tunnel192, Bidir-Upstream/Sparse, 1w2d/stopped

(*, 232.12.12.12), 1w2d/00:02:05, RP 192.168.123.6, flags: BCL
  Bidir-Upstream: Tunnel192, RPF nbr 192.168.123.6
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse, 1w2d/00:02:05
    Tunnel192, Bidir-Upstream/Sparse, 1w2d/stopped

R22#
R22#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter   Group Accounted
232.12.12.12     Ethernet0/1              1w2d      00:02:15  192.168.22.1
224.0.1.40       Tunnel192                1w2d      00:02:33  192.168.123.6
224.0.1.40       Ethernet0/1              1w2d      00:02:14  192.168.22.1
R22#
```

R23:

```
R23#show ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.123.23   Tunnel192                v2/S   1      30     1      192.168.123.23
192.168.23.1     Ethernet0/1              v2/S   0      30     1      192.168.23.1
R23#
R23#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
192.168.123.6     Tunnel192                23:52:09/00:01:23 v2    1 / B S P G
R23#
R23#show ip pim rp map
PIM Group-to-RP Mappings

Acl: MACL, Static, Bidir Mode
    RP: 192.168.123.6 (?)
R23#
R23#show ip mroute bidirectional
(*,232.12.12.12), 1w2d/-, RP 192.168.123.6, flags: BCL
```

```
  Bidir-Upstream: Tunnel192, RPF nbr: 192.168.123.6
  Incoming interface list:
    Ethernet0/1, Accepting/Sparse
    Tunnel192, Accepting/Sparse
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse, 23:52:28/00:02:22
    Tunnel192, Bidir-Upstream/Sparse, 23:52:28/stopped

(*, 232.12.12.12), 1w2d/00:02:22, RP 192.168.123.6, flags: BCL
  Bidir-Upstream: Tunnel192, RPF nbr 192.168.123.6
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse, 23:52:28/00:02:22
    Tunnel192, Bidir-Upstream/Sparse, 23:52:28/stopped

R23#
R23#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter   Group Accounted
232.12.12.12     Ethernet0/1              1w2d      00:02:13  192.168.23.1
224.0.1.40       Tunnel192                1w2d      00:02:32  192.168.123.6
224.0.1.40       Ethernet0/1              1w2d      00:02:10  192.168.23.1
R23#
```

Verify multicast connectivity according to the lab requirements:

```
R10#ping 232.12.12.12
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 232.12.12.12, timeout is 2 seconds:

Reply to request 0 from 192.168.23.1, 24 ms
Reply to request 0 from 192.168.22.1, 24 ms
R10#
```

3. VPN Technologies Section

Issue:

Configure MPLS according to the lab requirements. See the MP-BGP MPLS VPN and DMVPN Topology Diagram.

Solution:

The following are configuration examples on R1, R2, R3, and R4:

R1:

```
interface Ethernet0/1
 ip address 10.1.13.1 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
mpls ldp router-id Loopback0
!
```

R2:

```
!
interface Ethernet0/1
 ip address 10.1.23.2 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
mpls ldp router-id Loopback0
!
```


R3:

```
!
interface Ethernet0/0
 ip address 10.1.13.3 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
interface Ethernet0/1
 ip address 10.1.23.3 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
interface Ethernet0/2
 ip address 10.1.34.3 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
interface Ethernet0/3
 ip address 10.1.35.3 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
mpls ldp router-id Loopback0
!
!
```

R4:

```
!
interface Ethernet0/1
 ip address 10.1.34.4 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
mpls ldp router-id Loopback0
!
!
```

R5:

```
!
interface Ethernet0/1
 ip address 10.1.35.5 255.255.255.0
 ip ospf 1 area 0
 mpls ip
!
!
mpls ldp router-id Loopback0
!
!
```

MPLS verification example on R3:

```
R3#show mpls interfaces
Interface              IP            Tunnel   BGP Static Operational
Ethernet0/0            Yes (ldp)     No       No  No     Yes
Ethernet0/1            Yes (ldp)     No       No  No     Yes
Ethernet0/2            Yes (ldp)     No       No  No     Yes
Ethernet0/3            Yes (ldp)     No       No  No     Yes
R3#
R3#show mpls ldp neighbor
    Peer LDP Ident: 10.0.0.4:0; Local LDP Ident 10.0.0.3:0
        TCP connection: 10.0.0.4.44561 - 10.0.0.3.646
        State: Oper; Msgs sent/rcvd: 15927/15940; Downstream
        Up time: 1w2d
        LDP discovery sources:
          Ethernet0/2, Src IP addr: 10.1.34.4
        Addresses bound to peer LDP Ident:
          10.1.34.4       10.0.0.4
    Peer LDP Ident: 10.0.0.5:0; Local LDP Ident 10.0.0.3:0
        TCP connection: 10.0.0.5.46930 - 10.0.0.3.646
        State: Oper; Msgs sent/rcvd: 15931/15923; Downstream
```

```
        Up time: 1w2d
        LDP discovery sources:
          Ethernet0/3, Src IP addr: 10.1.35.5
        Addresses bound to peer LDP Ident:
          10.1.35.5       10.0.0.5
    Peer LDP Ident: 10.0.0.2:0; Local LDP Ident 10.0.0.3:0
        TCP connection: 10.0.0.2.646 - 10.0.0.3.24610
        State: Oper; Msgs sent/rcvd: 15914/15914; Downstream
        Up time: 1w2d
        LDP discovery sources:
          Ethernet0/1, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.0.0.2
    Peer LDP Ident: 10.0.0.1:0; Local LDP Ident 10.0.0.3:0
        TCP connection: 10.0.0.1.646 - 10.0.0.3.23664
        State: Oper; Msgs sent/rcvd: 15945/15918; Downstream
        Up time: 1w2d
        LDP discovery sources:
          Ethernet0/0, Src IP addr: 10.1.13.1
        Addresses bound to peer LDP Ident:
          10.1.13.1       10.0.0.1
R3#
```

Issue:

Configure MP-BGP VPN according to the lab requirements. See the MP-BGP MPLS VPN and DMVPN Topology Diagram.

Solution:

The following are configuration verification examples on R1, R2, R3, R4 and R5:

R1 (a PE router for the VRF SITE-1):

```
!
!
ip vrf SITE1
 rd 10001:1
 route-target export 1:1
 route-target import 1:1
!
interface Ethernet0/0
 ip vrf forwarding SITE1
 ip address 12.1.66.1 255.255.255.0
!
!
router bgp 10001
 bgp router-id 10.0.0.1
 bgp log-neighbor-changes
 neighbor 10.0.0.3 remote-as 10001
 neighbor 10.0.0.3 update-source Loopback0
 neighbor 10.0.0.3 next-hop-self
 !
 address-family vpnv4
  neighbor 10.0.0.3 activate
  neighbor 10.0.0.3 send-community both
 exit-address-family
 !
 address-family ipv4 vrf SITE1
  redistribute eigrp 100 route-map HQ-LO0ET10
 exit-address-family
!
!
ip prefix-list HQ-LO0ET10 seq 5 permit 12.1.0.0/24 ge 32
ip prefix-list HQ-LO0ET10 seq 10 permit 192.168.8.0/23 ge 24 le 24
ip prefix-list HQ-LO0ET10 seq 15 permit 123.1.0.0/16
!
route-map HQ-LO0ET10 permit 10
 match ip address prefix-list HQ-LO0ET10
!
!
```

Note that only the Loopback0 and Ethernet0/1 interfaces of the Headquarters and the summary network 123.1.0.0/16 that is received from the Internet are allowed from R1 to R2 in the VRF SITE1.

R2 (a PE router for the VRF SITE-1):

```
!
!ip vrf SITE1
 rd 10001:1
 route-target export 1:1
 route-target import 1:1
!
!
interface Ethernet0/0
 ip vrf forwarding SITE1
 ip address 12.2.102.2 255.255.255.0
!
!
router bgp 10001
 bgp router-id 10.0.0.2
 bgp log-neighbor-changes
 neighbor 10.0.0.3 remote-as 10001
 neighbor 10.0.0.3 update-source Loopback0
 neighbor 10.0.0.3 next-hop-self
 neighbor 10.0.0.3 default-originate
 !
 address-family vpnv4
  neighbor 10.0.0.3 activate
  neighbor 10.0.0.3 send-community both
exit-address-family
 !
 address-family ipv4 vrf SITE1
  redistribute eigrp 100 route-map MAIN-LO0ET10
  default-information originate
 exit-address-family
!
!
ip prefix-list MAIN-LO0ET10 seq 5 permit 12.2.0.0/24 ge 32
ip prefix-list MAIN-LO0ET10 seq 10 permit 192.168.12.0/23 ge 24 le 24
ip prefix-list MAIN-LO0ET10 seq 15 permit 0.0.0.0/0
!
route-map MAIN-LO0ET10 permit 10
 match ip address prefix-list MAIN-LO0ET10
!
!
```

Note that only the Loopback0 and Ethernet0/1 interfaces of the Regional Main Office and the default network 0.0.0.0/0 that is received from the Internet are allowed from R2 to R1 in the VRF SITE1.

R3 (a P router for all PE's in this scenario):

```
!
!
router bgp 10001
 bgp router-id 10.0.0.3
 bgp log-neighbor-changes
 neighbor reflect peer-group
 neighbor reflect remote-as 10001
 neighbor reflect update-source Loopback0
 neighbor 10.0.0.1 peer-group reflect
 neighbor 10.0.0.2 peer-group reflect
 neighbor 10.0.0.4 peer-group reflect
 neighbor 10.0.0.5 peer-group reflect
 !
 address-family ipv4
  neighbor reflect route-reflector-client
  neighbor 10.0.0.1 activate
  neighbor 10.0.0.2 activate
  neighbor 10.0.0.4 activate
  neighbor 10.0.0.5 activate
 exit-address-family
 !
 address-family vpnv4
  neighbor reflect send-community both
  neighbor reflect route-reflector-client
  neighbor 10.0.0.1 activate
  neighbor 10.0.0.2 activate
```

```
  neighbor 10.0.0.4 activate
  neighbor 10.0.0.5 activate
 exit-address-family
!
```

Note that the same update policies are grouped into a peer group on R3. Also, R3 is configured as a route-reflector for the PE routers. While R3 is not a PE itself, it must maintain a BGP "address-family vpnv4" to provide route-reflector services to the PE routers.

R4 (a PE router for the VRF SITE-2):

```
!
ip vrf SITE2
 rd 10001:2
 route-target export 2:2
 route-target import 2:2
!
interface Ethernet0/0
 ip vrf forwarding SITE2
 ip address 12.1.77.4 255.255.255.0
!
router bgp 10001
 bgp router-id 10.0.0.4
 bgp log-neighbor-changes
 neighbor 10.0.0.3 remote-as 10001
 neighbor 10.0.0.3 update-source Loopback0
 neighbor 10.0.0.3 next-hop-self
 !
 address-family vpnv4
  neighbor 10.0.0.3 activate
  neighbor 10.0.0.3 send-community both
 exit-address-family
 !
 address-family ipv4 vrf SITE2
  redistribute eigrp 100 route-map HQ-LO0ET10
 exit-address-family
!
!
ip prefix-list HQ-LO0ET10 seq 5 permit 12.1.0.0/24 ge 32
ip prefix-list HQ-LO0ET10 seq 10 permit 192.168.8.0/23 ge 24 le 24
!
route-map HQ-LO0ET10 permit 10
 match ip address prefix-list HQ-LO0ET10
!
```

Note that only the Loopback0 and Ethernet0/1 interfaces of the Headquarters are allowed from R4 to R5 in the VRF SITE2.

R5 (a PE router for the VRF SITE-2):

```
!
!
ip vrf SITE2
 rd 10001:2
 route-target export 2:2
 route-target import 2:2
!
!
interface Ethernet0/0
 ip vrf forwarding SITE2
 ip address 12.2.115.5 255.255.255.0
!
router bgp 10001`
 bgp router-id 10.0.0.5
 bgp log-neighbor-changes
 neighbor 10.0.0.3 remote-as 10001
 neighbor 10.0.0.3 update-source Loopback0
 neighbor 10.0.0.3 next-hop-self
 !
 address-family vpnv4
  neighbor 10.0.0.3 activate
  neighbor 10.0.0.3 send-community both
```



```
 exit-address-family
 !
 address-family ipv4 vrf SITE2
  redistribute eigrp 100 route-map MAIN-LO0ET10
 exit-address-family
!
!
ip prefix-list MAIN-LO0ET10 seq 5 permit 12.2.0.0/24 ge 32
ip prefix-list MAIN-LO0ET10 seq 10 permit 192.168.12.0/23 ge 24 le 24
!
route-map MAIN-LO0ET10 permit 10
 match ip address prefix-list MAIN-LO0ET10
!
```

Note that only the Loopback0 and Ethernet0/1 interfaces of the Regional Main Office are allowed from R5 to R4 in the VRF SITE2.

The following are VRF and VPN MP-BGP verification examples on R1, R2, R4 and R5:

R1 (a PE router for the VRF SITE-1):

```
R1#show ip vrf
  Name                              Default RD          Interfaces
  SITE1                             10001:1             Et0/0
R1#
R1#show bgp vpnv4 unicast all summary
BGP router identifier 10.0.0.1, local AS number 10001
BGP table version is 38, main routing table version 38
19 network entries using 2888 bytes of memory
19 path entries using 1444 bytes of memory
12/12 BGP path/bestpath attribute entries using 1728 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
12 BGP extended community entries using 3000 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 9084 total bytes of memory
BGP activity 20/0 prefixes, 22/2 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.3        4        10001   15373   15369       38    0    0 1w2d            8
R1#
R1#show bgp vpnv4 unicast all neighbors 10.0.0.3 advertised-routes
BGP table version is 38, local router ID is 10.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 10001:1 (default for vrf SITE1)
 *>  12.1.0.6/32      12.1.66.6           4437333         32768 ?
 *>  12.1.0.7/32      12.1.66.6           1537024         32768 ?
 *>  12.1.0.8/32      12.1.66.6           1551360         32768 ?
 *>  12.1.0.9/32      12.1.66.6           1551360         32768 ?
 *>  12.1.0.110/32    12.1.66.6           4096000         32768 ?
 *>  12.1.0.120/32    12.1.66.6           4096000         32768 ?
 *>  12.1.0.130/32    12.1.66.6           4101120         32768 ?
 *>  12.1.0.140/32    12.1.66.6           4101120         32768 ?
 *>  123.1.0.0/16     12.1.66.6           4437333         32768 ?
 *>  192.168.8.0      12.1.66.6           2058240         32768 ?
 *>  192.168.9.0      12.1.66.6           2058240         32768 ?

Total number of prefixes 11
R1#
```

R2 (a PE router for the VRF SITE-1):

```
R2#show ip vrf
  Name                              Default RD          Interfaces
  SITE1                             10001:1             Et0/0
```

```
R2#
R2#show bgp vpnv4 unicast all summary
BGP router identifier 10.0.0.2, local AS number 10001
BGP table version is 38, main routing table version 38
19 network entries using 2888 bytes of memory
19 path entries using 1444 bytes of memory
12/12 BGP path/bestpath attribute entries using 1728 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
12 BGP extended community entries using 3000 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 9084 total bytes of memory
BGP activity 22/2 prefixes, 22/2 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.3        4        10001   15390   15352       38    0    0 1w2d           11
R2#
R2#show bgp vpnv4 unicast all neighbors 10.0.0.3 advertised-routes
BGP table version is 38, local router ID is 10.0.0.2
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: 10001:1 (default for vrf SITE1)
 *>  0.0.0.0          12.2.102.10        2611200         32768 ?
 *>  12.2.0.10/32     12.2.102.10        2611200         32768 ?
 *>  12.2.0.11/32     12.2.102.10         435200         32768 ?
 *>  12.2.0.12/32     12.2.102.10         435456         32768 ?
 *>  12.2.0.13/32     12.2.102.10         435456         32768 ?
 *>  12.2.0.150/32    12.2.102.10         435200         32768 ?
 *>  192.168.12.0     12.2.102.10         333056         32768 ?
 *>  192.168.13.0     12.2.102.10         333056         32768 ?

Total number of prefixes 8
R2#
```

R4 (a PE router for the VRF SITE-2):

```
R4#   show ip vrf
  Name                              Default RD          Interfaces
  SITE2                             10001:2             Et0/0
R4#
R4#show bgp vpnv4 unicast all summary
BGP router identifier 10.0.0.4, local AS number 10001
BGP table version is 29, main routing table version 29
17 network entries using 2584 bytes of memory
17 path entries using 1292 bytes of memory
10/10 BGP path/bestpath attribute entries using 1440 bytes of memory
2 BGP rrinfo entries using 48 bytes of memory
10 BGP extended community entries using 2500 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 7864 total bytes of memory
BGP activity 18/0 prefixes, 20/2 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.3        4        10001   15392   15367       29    0    0 1w2d            7
R4#
R4#show bgp vpnv4 unicast all neighbors 10.0.0.3 advertised-routes
BGP table version is 29, local router ID is 10.0.0.4
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: 10001:2 (default for vrf SITE2)
 *>  12.1.0.6/32      12.1.77.7          1537024         32768 ?
 *>  12.1.0.7/32      12.1.77.7          1024640         32768 ?
 *>  12.1.0.8/32      12.1.77.7          1551360         32768 ?
 *>  12.1.0.9/32      12.1.77.7          1551360         32768 ?
 *>  12.1.0.110/32    12.1.77.7          4096000         32768 ?
```



```
 *>  12.1.0.120/32    12.1.77.7          4096000         32768 ?
 *>  12.1.0.130/32    12.1.77.7          4101120         32768 ?
 *>  12.1.0.140/32    12.1.77.7          4101120         32768 ?
 *>  192.168.8.0      12.1.77.7          2058240         32768 ?
 *>  192.168.9.0      12.1.77.7          2058240         32768 ?

Total number of prefixes 10
R4#
```

R5 (a PE router for the VRF SITE-2):

```
R5#show ip vrf
  Name                              Default RD          Interfaces
  SITE2                             10001:2             Et0/0
R5#
R5#show bgp vpnv4 unicast all summary
BGP router identifier 10.0.0.5, local AS number 10001
BGP table version is 34, main routing table version 34
17 network entries using 2584 bytes of memory
17 path entries using 1292 bytes of memory
10/10 BGP path/bestpath attribute entries using 1440 bytes of memory
2 BGP rrinfo entries using 48 bytes of memory
10 BGP extended community entries using 2500 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 7864 total bytes of memory
BGP activity 18/0 prefixes, 20/2 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.3        4        10001   15390   15375       34    0    0 1w2d           10
R5#
R5#show bgp vpnv4 unicast all neighbors 10.0.0.3 advertised-routes
BGP table version is 34, local router ID is 10.0.0.5
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: 10001:2 (default for vrf SITE2)
 *>  12.2.0.10/32    12.2.115.11         435200         32768 ?
 *>  12.2.0.11/32    12.2.115.11         409600         32768 ?
 *>  12.2.0.12/32    12.2.115.11         435456         32768 ?
 *>  12.2.0.13/32    12.2.115.11         435456         32768 ?
 *>  12.2.0.150/32   12.2.115.11         435200         32768 ?
 *>  192.168.12.0    12.2.115.11         333056         32768 ?
 *>  192.168.13.0    12.2.115.11         333056         32768 ?

Total number of prefixes 7
R5#
```

The following is a MP-BGP verification example on R3:

```
R3#show bgp vpnv4 unicast all
BGP table version is 45, local router ID is 10.0.0.3
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: 10001:1
 *>i 0.0.0.0          10.0.0.2           2611200    100      0 ?
 *>i 12.1.0.6/32      10.0.0.1           4437333    100      0 ?
 *>i 12.1.0.7/32      10.0.0.1           1537024    100      0 ?
 *>i 12.1.0.8/32      10.0.0.1           1551360    100      0 ?
 *>i 12.1.0.9/32      10.0.0.1           1551360    100      0 ?
 *>i 12.1.0.110/32    10.0.0.1           4096000    100      0 ?
 *>i 12.1.0.120/32    10.0.0.1           4096000    100      0 ?
 *>i 12.1.0.130/32    10.0.0.1           4101120    100      0 ?
 *>i 12.1.0.140/32    10.0.0.1           4101120    100      0 ?
 *>i 12.2.0.10/32     10.0.0.2           2611200    100      0 ?
```

```
 *>i 12.2.0.11/32     10.0.0.2            435200    100      0 ?
 *>i 12.2.0.12/32     10.0.0.2            435456    100      0 ?
 *>i 12.2.0.13/32     10.0.0.2            435456    100      0 ?
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 12.2.0.150/32    10.0.0.2            435200    100      0 ?
 *>i 123.1.0.0/16     10.0.0.1           4437333    100      0 ?
 *>i 192.168.8.0      10.0.0.1           2058240    100      0 ?
 *>i 192.168.9.0      10.0.0.1           2058240    100      0 ?
 *>i 192.168.12.0     10.0.0.2            333056    100      0 ?
 *>i 192.168.13.0     10.0.0.2            333056    100      0 ?
Route Distinguisher: 10001:2
 *>i 12.1.0.6/32      10.0.0.4           1537024    100      0 ?
 *>i 12.1.0.7/32      10.0.0.4           1024640    100      0 ?
 *>i 12.1.0.8/32      10.0.0.4           1551360    100      0 ?
 *>i 12.1.0.9/32      10.0.0.4           1551360    100      0 ?
 *>i 12.1.0.110/32    10.0.0.4           4096000    100      0 ?
 *>i 12.1.0.120/32    10.0.0.4           4096000    100      0 ?
 *>i 12.1.0.130/32    10.0.0.4           4101120    100      0 ?
 *>i 12.1.0.140/32    10.0.0.4           4101120    100      0 ?
 *>i 12.2.0.10/32     10.0.0.5            435200    100      0 ?
 *>i 12.2.0.11/32     10.0.0.5            409600    100      0 ?
 *>i 12.2.0.12/32     10.0.0.5            435456    100      0 ?
 *>i 12.2.0.13/32     10.0.0.5            435456    100      0 ?
 *>i 12.2.0.150/32    10.0.0.5            435200    100      0 ?
 *>i 192.168.8.0      10.0.0.4           2058240    100      0 ?
 *>i 192.168.9.0      10.0.0.4           2058240    100      0 ?
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 192.168.12.0     10.0.0.5            333056    100      0 ?
 *>i 192.168.13.0     10.0.0.5            333056    100      0 ?
R3#
```

Issue:

Configure DMVPN and DMVPN routing according to the lab requirements. See the MP-BGP MPLS VPN and DMVPN Topology Diagram.

Solution:

The following are configuration examples on R6, R10, R22, and R23:

R6 (the DMVPN hub router):

R6:

```
interface Tunnel192
 bandwidth 100000
 ip address 192.168.123.6 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip pim nbma-mode
 ip pim sparse-mode
 ip nhrp authentication nhrpkey
 ip nhrp map multicast dynamic
 ip nhrp network-id 123
 ip nhrp redirect
 ip tcp adjust-mss 1360
 delay 100
 tunnel source Serial1/0
 tunnel mode gre multipoint
!
router eigrp HQ-CE
 !
 address-family ipv4 unicast autonomous-system 100
  !
  af-interface Tunnel192
   no split-horizon
  exit-af-interface
  !
  topology base
   default-metric 1500 100 3 255 1500
   distribute-list prefix NET-192 out Tunnel192
   redistribute eigrp 1000
   redistribute bgp 12000
  exit-af-topology
```

```
  neighbor 192.168.123.10 Tunnel192
  neighbor 192.168.123.22 Tunnel192
  neighbor 192.168.123.23 Tunnel192
  network 12.1.66.0 0.0.0.255
  network 192.168.123.0
 exit-address-family
!
ip prefix-list NET-192 seq 5 permit 192.168.8.0/23 ge 24 le 24
ip prefix-list NET-192 seq 10 permit 192.168.12.0/23 ge 24 le 24
ip prefix-list NET-192 seq 15 permit 192.168.22.0/23 ge 24 le 24
```

R10 (a DMVPN spoke router):

```
interface Tunnel192
 bandwidth 100000
 ip address 192.168.123.10 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip nhrp authentication nhrpkey
 ip nhrp map 192.168.123.6 123.1.6.2
 ip nhrp map multicast 123.1.6.2
 ip nhrp network-id 123
 ip nhrp nhs 192.168.123.6
 ip nhrp shortcut
 ip tcp adjust-mss 1360
 delay 100
 tunnel source Serial1/0
 tunnel mode gre multipoint
!
router eigrp 100
 distribute-list prefix NET-192 out Tunnel192
 default-metric 1000 100 3 255 1500
 network 12.2.102.0 0.0.0.255
 network 192.168.123.0
 redistribute eigrp 1000
 redistribute bgp 12000
 neighbor 192.168.123.6 Tunnel192
!
ip prefix-list NET-192 seq 5 permit 192.168.12.0/23 ge 24 le 24
```

R22 (a DMVPN spoke router):

```
!
interface Tunnel192
 ip address 192.168.123.22 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip pim sparse-mode
 ip nhrp authentication nhrpkey
 ip nhrp map 192.168.123.6 123.1.6.2
 ip nhrp map multicast 123.1.6.2
 ip nhrp network-id 123
 ip nhrp nhs 192.168.123.6
 ip nhrp shortcut
 ip tcp adjust-mss 1360
 tunnel source Ethernet0/2
 tunnel mode gre multipoint
!
router eigrp 100
 network 192.168.22.0
 network 192.168.123.0
 neighbor 192.168.123.6 Tunnel192
!
```

R23 (a DMVPN spoke router):

```
interface Tunnel192
 ip address 192.168.123.23 255.255.255.0
 no ip redirects
 ip mtu 1400
 ip pim sparse-mode
 ip nhrp authentication nhrpkey
 ip nhrp map 192.168.123.6 123.1.6.2
 ip nhrp map multicast 123.1.6.2
```

```
 ip nhrp network-id 123
 ip nhrp nhs 192.168.123.6
 ip nhrp shortcut
 ip tcp adjust-mss 1360
 tunnel source Ethernet0/2
 tunnel mode gre multipoint
!
router eigrp 100
 network 192.168.23.0
 network 192.168.123.0
 neighbor 192.168.123.6 Tunnel192
!
```

Here are IPv4 DMVPN verification examples on R6, R10, R22, and R23:

```
R6#show dmvpn ipv4 interface tunnel192
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel192, IPv4 NHRP Details
Type:Hub, NHRP Peers:3,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 123.1.10.2       192.168.123.10    UP     1w2d     D
     1 123.1.22.2       192.168.123.22    UP     1w2d     D
     1 123.1.23.2       192.168.123.23    UP    1d11h     D

R6#
R10#show dmvpn ipv4 interface tunnel192
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel192, IPv4 NHRP Details
Type:Spoke, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 123.1.6.2         192.168.123.6    UP     1w2d     S

R10#

R22#show dmvpn ipv4 interface tunnel192
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel192, IPv4 NHRP Details
Type:Spoke, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 123.1.6.2         192.168.123.6    UP     1w2d     S

R22#

R23#show dmvpn ipv4 interface tunnel192
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel192, IPv4 NHRP Details
```

```
Type:Spoke, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1 123.1.6.2         192.168.123.6    UP    1d11h     S

R23#
```

Verify DMVPN connectivity from the DMVPN hub R6:

```
R6#ping 192.168.123.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.123.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 14/14/15 ms
R6#ping 192.168.123.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.123.22, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/13/14 ms
R6#ping 192.168.123.23
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.123.23, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 10/12/14 ms
R6#
```

Verify reachability between the Ethernet0/1 interfaces of Headguarters, Branch Office 1 and Branch Office 2:

```
R8#ping 192.168.22.1 source ethernet0/1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.8.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 10/10/11 ms
R8#ping 192.168.23.1 source ethernet0/1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.23.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.8.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 10/10/11 ms
R8#

R9#ping 192.168.22.1 source ethernet0/1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.22.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.9.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 10/10/11 ms
R9#
R9#ping 192.168.23.1 source ethernet0/1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.23.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.9.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 10/10/11 ms
R9#
```

Verify reachability between the Ethernet0/1 interfaces of Headquarters and Main Regional Office:

```
R9#traceroute 192.168.13.1 source ethernet0/1
Type escape sequence to abort.
Tracing the route to 192.168.13.1
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.94.140 1 msec 1 msec 0 msec
  2 12.1.42.120 1 msec 0 msec 0 msec
  3 12.1.62.6 1 msec 2 msec 1 msec
  4 192.168.123.10 11 msec 11 msec 11 msec
  5 12.2.105.150 11 msec 11 msec 11 msec
  6 12.2.135.13 12 msec *  12 msec
R9#
```

Note that the reachability between Headquarters and the Main Regional Office is established over the DMVPN. This fulfills the requirements of this exam.

Issue:

Configure Internet Routing according to the lab requirements. See the MP-BGP MPLS VPN and DMVPN Topology Diagram.

Solution:

The following are configuration examples on R2:

```
!
router bgp 10001
!
 address-family ipv4 vrf SITE1
  redistribute eigrp 100 route-map MAIN-LO0ET10
  default-information originate
 exit-address-family
!
```

Verify the 0.0.0.0/0 network propagation in the Main Regional Office and in the Headquarters via MPLS VPN:

R10:

```
R10#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 12000", distance 20, metric 0, candidate default path
  Tag 30000, type external
  Redistributing via eigrp 100, eigrp 1000
  Advertised by eigrp 100
                eigrp 1000
  Last update from 123.1.10.1 1w2d ago
  Routing Descriptor Blocks:
  * 123.1.10.1, from 123.1.10.1, 1w2d ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 30000
      MPLS label: none
R10#
```

R12:

```
R12#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "eigrp 1000", distance 170, metric 2611456, candidate default path
  Tag 30000, type external
  Redistributing via eigrp 1000
  Last update from 12.2.125.150 on Ethernet0/0, 1w2d ago
  Routing Descriptor Blocks:
  * 12.2.125.150, from 12.2.125.150, 1w2d ago, via Ethernet0/0
      Route metric is 2611456, traffic share count is 1
      Total delay is 2010 microseconds, minimum bandwidth is 1000 Kbit
      Reliability 3/255, minimum MTU 1500 bytes
      Loading 255/255, Hops 2
      Route tag 30000
R12#
```

R13:

```
R13#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "eigrp 1000", distance 170, metric 2611456, candidate default path
  Tag 30000, type external
  Redistributing via eigrp 1000
```



```
  Last update from 12.2.135.150 on Ethernet0/0, 1w2d ago
  Routing Descriptor Blocks:
  * 12.2.135.150, from 12.2.135.150, 1w2d ago, via Ethernet0/0
      Route metric is 2611456, traffic share count is 1
      Total delay is 2010 microseconds, minimum bandwidth is 1000 Kbit
      Reliability 3/255, minimum MTU 1500 bytes
      Loading 255/255, Hops 2
      Route tag 30000
R13#
```

R1:

```
R1#show ip route vrf SITE1 0.0.0.0

Routing Table: SITE1
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 10001", distance 200, metric 2611200, candidate default path, type internal
  Redistributing via eigrp 100
  Advertised by eigrp 100 metric 1000 100 255 3 1500
  Last update from 10.0.0.2 1w2d ago
  Routing Descriptor Blocks:
  * 10.0.0.2 (default), from 10.0.0.3, 1w2d ago
      Route metric is 2611200, traffic share count is 1
      AS Hops 0
      MPLS label: 31
      MPLS Flags: MPLS Required
R1#
```

R8:

```
R8#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "eigrp 1000", distance 170, metric 7178240, candidate default path, type external
  Redistributing via eigrp 1000
  Last update from 12.1.83.130 on Ethernet0/0, 1w2d ago
  Routing Descriptor Blocks:
  * 12.1.83.130, from 12.1.83.130, 1w2d ago, via Ethernet0/0
      Route metric is 7178240, traffic share count is 1
      Total delay is 4020 microseconds, minimum bandwidth is 1000 Kbit
      Reliability 3/255, minimum MTU 1500 bytes
      Loading 255/255, Hops 5
R8#
```

R9:

```
R9#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "eigrp 1000", distance 170, metric 7178240, candidate default path, type external
  Redistributing via eigrp 1000
  Last update from 12.1.94.140 on Ethernet0/0, 1d12h ago
  Routing Descriptor Blocks:
  * 12.1.94.140, from 12.1.94.140, 1d12h ago, via Ethernet0/0
      Route metric is 7178240, traffic share count is 1
      Total delay is 4020 microseconds, minimum bandwidth is 1000 Kbit
      Reliability 3/255, minimum MTU 1500 bytes
      Loading 255/255, Hops 5
R9#
```

Verify connectivity to the Internet from R8, R9, R12, and R13 according to the lab requirements:

```
R8#ping 33.33.33.33 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.33, timeout is 2 seconds:
Packet sent with a source address of 12.1.0.8
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 10/10/11 ms
R8#

R9#ping 33.33.33.33 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.33, timeout is 2 seconds:
Packet sent with a source address of 12.1.0.9
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/10/11 ms
R9#

R12#ping 33.33.33.33 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.33, timeout is 2 seconds:
Packet sent with a source address of 12.2.0.12
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/9/10 ms
R12#

R13#ping 33.33.33.33 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.33.33.33, timeout is 2 seconds:
Packet sent with a source address of 12.2.0.13
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/9/10 ms
R13#
```

4. Infrastructure Security Section

Issue:

Configure DMVPN Security according to the lab requirements. See the MP-BGP MPLS VPN and DMVPN Topology Diagram.

Solution:

Here are configuration and verification examples of configuring ISAKMP and IPSec on routers R6, R10, R22, and R23:

R6:

```
!
crypto isakmp policy 1
 encr aes
 authentication pre-share
 group 2
crypto isakmp key SHAREDKEY address 0.0.0.0
!
!
crypto ipsec transform-set dmvpn-transform esp-aes esp-sha-hmac
 mode transport
!
!
crypto ipsec profile dmvpn-profile
 set transform-set dmvpn-transform
!
!
interface Tunnel192
 tunnel protection ipsec profile dmvpn-profile
!
```

R10:

```
!
crypto isakmp policy 1
 encr aes
 authentication pre-share
 group 2
crypto isakmp key SHAREDKEY address 0.0.0.0
!
!
crypto ipsec transform-set dmvpn-transform esp-aes esp-sha-hmac
 mode transport
!
!
crypto ipsec profile dmvpn-profile
 set transform-set dmvpn-transform
!
```



```
!
interface Tunnel192
 tunnel protection ipsec profile dmvpn-profile
!
```

R22:

```
!
crypto isakmp policy 1
 encr aes
 authentication pre-share
 group 2
crypto isakmp key SHAREDKEY address 0.0.0.0
!
!
crypto ipsec transform-set dmvpn-transform esp-aes esp-sha-hmac
 mode transport
!
!
crypto ipsec profile dmvpn-profile
 set transform-set dmvpn-transform
!
!
interface Tunnel192
 tunnel protection ipsec profile dmvpn-profile
!
```

R23:

```
!
crypto isakmp policy 1
 encr aes
 authentication pre-share
 group 2
crypto isakmp key SHAREDKEY address 0.0.0.0
!
!
crypto ipsec transform-set dmvpn-transform esp-aes esp-sha-hmac
 mode transport
!
!
crypto ipsec profile dmvpn-profile
 set transform-set dmvpn-transform
!
!
interface Tunnel192
 tunnel protection ipsec profile dmvpn-profile
!
```

Here is a secure DMVPN verification example on R6:

```
R6#show crypto isakmp key
Keyring      Hostname/Address                            Preshared Key

default      0.0.0.0        [0.0.0.0]                    SHAREDKEY
R6#
R6#show crypto isakmp policy

Global IKE policy
Protection suite of priority 1
        encryption algorithm:   AES - Advanced Encryption Standard (128 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #2 (1024 bit)
        lifetime:               86400 seconds, no volume limit
R6#
R6#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
123.1.6.2       123.1.22.2      QM_IDLE           1028 ACTIVE
123.1.6.2       123.1.23.2      QM_IDLE           1029 ACTIVE
```

```
123.1.10.2      123.1.6.2       QM_IDLE           1030 ACTIVE

IPv6 Crypto ISAKMP SA

R6#
R6#show crypto ipsec profile
IPSEC profile default
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={
                default:  { esp-aes esp-sha-hmac  } ,
        }

IPSEC profile dmvpn-profile
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={
                dmvpn-transform:  { esp-aes esp-sha-hmac  } ,
        }

R6#
R6#show crypto ipsec transform-set
Transform set default: { esp-aes esp-sha-hmac  }
   will negotiate = { Transport,  },

Transform set dmvpn-transform: { esp-aes esp-sha-hmac  }
   will negotiate = { Transport,  },

R6#
R6#show crypto session
Crypto session current status

Interface: Tunnel192
Session status: UP-ACTIVE
Peer: 123.1.10.2 port 500
  IKEv1 SA: local 123.1.6.2/500 remote 123.1.10.2/500 Active
  IPSEC FLOW: permit 47 host 123.1.6.2 host 123.1.10.2
        Active SAs: 2, origin: crypto map

Interface: Tunnel192
Session status: UP-ACTIVE
Peer: 123.1.22.2 port 500
  IKEv1 SA: local 123.1.6.2/500 remote 123.1.22.2/500 Active
  IPSEC FLOW: permit 47 host 123.1.6.2 host 123.1.22.2
        Active SAs: 2, origin: crypto map

Interface: Tunnel192
Session status: UP-ACTIVE
Peer: 123.1.23.2 port 500
  IKEv1 SA: local 123.1.6.2/500 remote 123.1.23.2/500 Active
  IPSEC FLOW: permit 47 host 123.1.6.2 host 123.1.23.2
        Active SAs: 2, origin: crypto map

R6#
R6#show crypto ipsec sa

interface: Tunnel192
    Crypto map tag: Tunnel192-head-0, local addr 123.1.6.2

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (123.1.6.2/255.255.255.255/47/0)
   remote ident (addr/mask/prot/port): (123.1.10.2/255.255.255.255/47/0)
   current_peer 123.1.10.2 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 252343, #pkts encrypt: 252343, #pkts digest: 252343
    #pkts decaps: 182058, #pkts decrypt: 182058, #pkts verify: 182058
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 123.1.6.2, remote crypto endpt.: 123.1.10.2
     path mtu 1500, ip mtu 1500, ip mtu idb (none)
     current outbound spi: 0xE0643A65(3764664933)
     PFS (Y/N): N, DH group: none
```



```
      inbound esp sas:
       spi: 0xF07AB705(4034574085)
         transform: esp-aes esp-sha-hmac ,
         in use settings ={Transport, }
         conn id: 471, flow_id: SW:471, sibling_flags 80000000, crypto map: Tunnel192-head-0
         sa timing: remaining key lifetime (k/sec): (4215530/3247)
         IV size: 16 bytes
         replay detection support: Y
         Status: ACTIVE(ACTIVE)

      inbound ah sas:

      inbound pcp sas:

      outbound esp sas:
       spi: 0xE0643A65(3764664933)
         transform: esp-aes esp-sha-hmac ,
         in use settings ={Transport, }
         conn id: 472, flow_id: SW:472, sibling_flags 80000000, crypto map: Tunnel192-head-0
         sa timing: remaining key lifetime (k/sec): (4215526/3247)
         IV size: 16 bytes
         replay detection support: Y
         Status: ACTIVE(ACTIVE)

      outbound ah sas:

      outbound pcp sas:

    protected vrf: (none)
    local  ident (addr/mask/prot/port): (123.1.6.2/255.255.255.255/47/0)
    remote ident (addr/mask/prot/port): (123.1.22.2/255.255.255.255/47/0)
    current_peer 123.1.22.2 port 500
      PERMIT, flags={origin_is_acl,}
     #pkts encaps: 261655, #pkts encrypt: 261655, #pkts digest: 261655
     #pkts decaps: 224685, #pkts decrypt: 224685, #pkts verify: 224685
     #pkts compressed: 0, #pkts decompressed: 0
     #pkts not compressed: 0, #pkts compr. failed: 0
     #pkts not decompressed: 0, #pkts decompress failed: 0
     #send errors 0, #recv errors 0

      local crypto endpt.: 123.1.6.2, remote crypto endpt.: 123.1.22.2
      path mtu 1500, ip mtu 1500, ip mtu idb (none)
      current outbound spi: 0xE6487C86(3863510150)
      PFS (Y/N): N, DH group: none

      inbound esp sas:
       spi: 0xBD430E95(3175288469)
         transform: esp-aes esp-sha-hmac ,
         in use settings ={Transport, }
         conn id: 469, flow_id: SW:469, sibling_flags 80000000, crypto map: Tunnel192-head-0
         sa timing: remaining key lifetime (k/sec): (4264810/2232)
         IV size: 16 bytes
         replay detection support: Y
         Status: ACTIVE(ACTIVE)

      inbound ah sas:

      inbound pcp sas:

      outbound esp sas:
       spi: 0xE6487C86(3863510150)
         transform: esp-aes esp-sha-hmac ,
         in use settings ={Transport, }
         conn id: 470, flow_id: SW:470, sibling_flags 80000000, crypto map: Tunnel192-head-0
         sa timing: remaining key lifetime (k/sec): (4264804/2232)
         IV size: 16 bytes
         replay detection support: Y
         Status: ACTIVE(ACTIVE)

      outbound ah sas:

      outbound pcp sas:

    protected vrf: (none)
    local  ident (addr/mask/prot/port): (123.1.6.2/255.255.255.255/47/0)
```

```
   remote ident (addr/mask/prot/port): (123.1.23.2/255.255.255.255/47/0)
   current_peer 123.1.23.2 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 40579, #pkts encrypt: 40579, #pkts digest: 40579
    #pkts decaps: 34845, #pkts decrypt: 34845, #pkts verify: 34845
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 123.1.6.2, remote crypto endpt.: 123.1.23.2
     path mtu 1500, ip mtu 1500, ip mtu idb (none)
     current outbound spi: 0x114012D1(289411793)
     PFS (Y/N): N, DH group: none

     inbound esp sas:
      spi: 0xF33B50EF(4080750831)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Transport, }
        conn id: 467, flow_id: SW:467, sibling_flags 80000000, crypto map: Tunnel192-head-0
        sa timing: remaining key lifetime (k/sec): (4368832/1285)
        IV size: 16 bytes
        replay detection support: Y
        Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:

     outbound esp sas:
      spi: 0x114012D1(289411793)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Transport, }
        conn id: 468, flow_id: SW:468, sibling_flags 80000000, crypto map: Tunnel192-head-0
        sa timing: remaining key lifetime (k/sec): (4368822/1285)
        IV size: 16 bytes
        replay detection support: Y
        Status: ACTIVE(ACTIVE)

     outbound ah sas:

     outbound pcp sas:
R6#
```

Issue: Configure SSH according to the lab requirements.
Solution:

SSH version 2 uses an RSA private or public key pair to secure the session. You can generate a key pair using the router hostname and domain name, or you can supply a key pair that has been independently generated. Use the first option in this exercise, because the scenario specifies the domain name. Create the user and password entry for **administrator**/**cisco**. The following configuration would meet the requirements of this section.

```
R21(config)#ip domain ccie.cisco.com
R21(config)#crypto key generate rsa
…
Choose the size of the key modulus in the range of 360 to 2048 for your
  General Purpose Keys. Choosing a key modulus greater than 512 may take
  a few minutes.

How many bits in the modulus [512]: 1024
% Generating 1024 bit RSA keys, keys will be non-exportable... [OK]

R21(config)#
```

It is important to remember that SSH version 2 requires at least 768 bits.

The generated public key is not stored in the startup or running configurations. It is stored in a private area of NVRAM and can be verified by the following command:

```
R21#show crypto key mypubkey rsa
```



```
% Key pair was generated at: 12:14:07 PST Dec 8 2014
Key name: R21.ccie.cisco.com.server
Key type: RSA KEYS
Temporary key
 Usage: Encryption Key
 Key is not exportable.
 Key Data:
  307C300D 06092A86 4886F70D 01010105 00036B00 30680261 00B40AF8 473D8847
  6B6A6886 03E6A916 248AE6B7 0A43AD8F 5F3040E3 88BA4343 A967B13D 6AE917E4
  D33FFF6A 1C3DFD9A 0C98FDD7 B13CD1AE A78C55EF 9384A704 88164941 34740B6D
  DFB2AEF7 4C0C2593 6CC1BD77 CDC15469 CB50DD10 F2CBF054 9B020301 0001
R21#
```

The following is an SSH configuration example on R21.

R21:

```
!
ip domain name ccie.cisco.com
!
username administrator password 0 cisco
!
ip ssh port 2121 rotary 1
ip ssh version 2
!
line vty 0 1
 location 360rsw05-lab-ca12, SJ
 exec-timeout 0 0
 privilege level 15
 login local
 rotary 1
 transport input ssh
line vty 2 4
 location 360rsw05-lab-ca12, SJ
 exec-timeout 0 0
 privilege level 15
 no login
 transport input all
!
```

The following are SSH verification examples on R16 and R17:

```
R16#ssh -l administrator -p 2121 -v 2 192.168.21.1
Password:

-------------------------------------------------
Cisco 360 R&S Graded Assessment Labs
Product, POD location:  360rsw05-lab-ca12, SJ
Device:                 R21
-------------------------------------------------

R21>who
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:02:31   360rsw05-lab-ca12, SJ
*  2 vty 0     administra idle                 00:00:00 12.3.76.16

  Interface    User               Mode         Idle     Peer Address

R21>exit

[Connection to 192.168.21.1 closed by foreign host]
R16#

R17#ssh -l administrator -p 2121 -v 2 192.168.21.1
Password:

-------------------------------------------------
Cisco 360 R&S Graded Assessment Labs
Product, POD location:  360rsw05-lab-ca12, SJ
Device:                 R21
-------------------------------------------------

R21>who
```

```
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:02:50   360rsw05-lab-ca12, SJ
   2 vty 0     administra idle                 00:00:18 12.3.76.16
*  3 vty 1     administra idle                 00:00:00 12.3.76.17

  Interface    User               Mode         Idle     Peer Address

R21>exit

[Connection to 192.168.21.1 closed by foreign host]
R17#
```

5. Infrastructure Services Section

Issue: Configure access to Internet Shared Service according to the lab requirements and the "IPv4 IGP" diagram.

Solution:

The following is a configuration example on R10.

R10:

```
interface Loopback0
 ip address 12.2.0.10 255.255.255.255
 ip nat enable
 ip virtual-reassembly in
!
!
interface Ethernet0/0.102
 encapsulation dot1Q 102
 ip address 12.2.102.10 255.255.255.0
 ip nat enable
!
interface Ethernet0/0.105
 encapsulation dot1Q 105
 ip address 12.2.105.10 255.255.255.0
 ip nat enable
!
interface Serial1/0
 ip address 123.1.10.2 255.255.255.252
 ip nat enable
 serial restart-delay 0
!
ip nat source route-map LOOPBACKS interface Loopback0 overload
!
!

!
route-map LOOPBACKS permit 10
 match ip address 111
!
access-list 111 permit ip 12.2.0.0 0.0.0.255 host 30.30.30.30
access-list 111 permit ip 12.1.0.0 0.0.0.255 host 30.30.30.30
!
```

Note that the keywords inside and outside are not allowed in the configuration. The **ip nat enable** command is used in this lab.

The following are network address translation verification examples on R8, R9, R10, R12, and R13.

R10:

```
R10#show ip nat nvi statistics
Total active translations: 2 (0 static, 2 dynamic; 2 extended)
NAT Enabled interfaces:
```



```
  Ethernet0/0.102, Ethernet0/0.105, Serial1/0, Loopback0
Hits: 574  Misses: 46
CEF Translated packets: 332, CEF Punted packets: 0
Expired translations: 44
Dynamic mappings:
-- Source [Id: 1] route-map LOOPBACKS interface Loopback0 refcount 2
R10#

R10#clear ip nat nvi translation *
R10#show ip nat nvi translations
R10#
```

To test this NAT configuration, open the following Telnet sessions.

R8:

```
R8#telnet 30.30.30.30 79 /source-interface Loopback0
Trying 30.30.30.30, 79 ... Open

------------------------------------------------
Cisco 360 R&S Graded Assessment Labs
Product, POD location:  360rsw05-lab-ca12, SJ
Device:                 AS-30000
------------------------------------------------

    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:04:40   360rsw05-lab-ca12, SJ
*  2 vty 0                 idle                 00:00:00 12.2.0.10

  Interface    User               Mode         Idle     Peer Address
  Se2/0        HQ                 Sync PPP     00:00:01 123.1.6.2

[Connection to 30.30.30.30 closed by foreign host]
R8#
```

R9:

```
R9#telnet 30.30.30.30 79 /source-interface Loopback0
Trying 30.30.30.30, 79 ... Open

------------------------------------------------
Cisco 360 R&S Graded Assessment Labs
Product, POD location:  360rsw05-lab-ca12, SJ
Device:                 AS-30000
------------------------------------------------

    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:04:45   360rsw05-lab-ca12, SJ
*  2 vty 0                 idle                 00:00:00 12.2.0.10

  Interface    User               Mode         Idle     Peer Address
  Se2/0        HQ                 Sync PPP     00:00:00 123.1.6.2

[Connection to 30.30.30.30 closed by foreign host]
R9#
```

R12:

```
R12#telnet 30.30.30.30 79 /source-interface Loopback0
Trying 30.30.30.30, 79 ... Open

------------------------------------------------
Cisco 360 R&S Graded Assessment Labs
Product, POD location:  360rsw05-lab-ca12, SJ
Device:                 AS-30000
------------------------------------------------

    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:00:00   360rsw05-lab-ca12, SJ
*  2 vty 0                 idle                 00:00:00 12.2.0.10

  Interface    User               Mode         Idle     Peer Address
```

```
  Se2/0        HQ                Sync PPP     00:00:01 123.1.6.2

[Connection to 30.30.30.30 closed by foreign host]
R12#
```

R13:

```
R13#telnet 30.30.30.30 79 /source-interface Loopback0
Trying 30.30.30.30, 79 ... Open

------------------------------------------------
Cisco 360 R&S Graded Assessment Labs
Product, POD location:  360rsw05-lab-ca12, SJ
Device:                 AS-30000
------------------------------------------------

    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:00:04   360rsw05-lab-ca12, SJ
*  2 vty 0                 idle                 00:00:00 12.2.0.10

  Interface    User               Mode         Idle     Peer Address
  Se2/0        HQ                 Sync PPP     00:00:00 123.1.6.2

[Connection to 30.30.30.30 closed by foreign host]
R13#
```

Verify the IP NAT translation table is correct.

R10:

```
R10#show ip nat nvi translations
Pro Source global      Source local       Destin  local      Destin  global
tcp 12.2.0.10:20700    12.1.0.8:20700     30.30.30.30:79     30.30.30.30:79
tcp 12.2.0.10:11790    12.1.0.9:11790     30.30.30.30:79     30.30.30.30:79
tcp 12.2.0.10:11714    12.2.0.12:11714    30.30.30.30:79     30.30.30.30:79
tcp 12.2.0.10:28144    12.2.0.13:28144    30.30.30.30:79     30.30.30.30:79
R10#
```

Issue:

Configure Netwrok Time Services according to the lab requirements and the “IPv4 IGP” diagram.

Solution:

Here is a NTP configuration example on R8:

```
!
ntp authentication-key 1 md5 1511021F0725 7
ntp authenticate
ntp trusted-key 1
ntp server 33.33.33.33 key 1
!
```

Note that the password is defined in the Restrictions and Goals section. The password is **cisco**.

Here is a NTP verification example on R8:

```
R8#show ntp associations detail
33.33.33.33 configured, ipv4, authenticated, our_master, sane, valid, stratum 1
ref ID .LOCL., time D83D7E7E.7A1CAD58 (08:25:34.477 PST Thu Dec 18 2014)
our mode client, peer mode server, our poll intvl 1024, peer poll intvl 1024
root delay 0.00 msec, root disp 2.21, reach 377, sync dist 22.66
delay 11.00 msec, offset 0.5000 msec, dispersion 1.98, jitter 0.97 msec
precision 2**10, version 4
assoc id 38976, assoc name 33.33.33.33
assoc in packets 6539, assoc out packets 6541, assoc error packets 0
org time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
rec time D83D7E82.34BC6B10 (08:25:38.206 PST Thu Dec 18 2014)
```



```
xmt time D83D7E82.34BC6B10 (08:25:38.206 PST Thu Dec 18 2014)
filtdelay =    11.00   12.00   11.00   11.00   11.00   11.00   11.00   11.00
filtoffset =    0.50    0.00    0.50    0.50    0.50    0.50    0.50    0.50
filterror =     1.95    1.98    2.01    2.04    2.07    2.10    2.13    2.16
minpoll = 6, maxpoll = 10

R8#
```

Issue:

Configure IP Statistics Services according to the lab requirements and the "IPv4 IGP" diagram.

Solution:

Cisco IP accounting support provides basic IP accounting functions. By enabling IP accounting, users can see the number of bytes and packets switched through the Cisco IOS Software based on source and destination IP addresses. Only transit IP traffic is measured and only on an outbound basis; traffic that is generated by the software or terminates in the software is not included in the accounting statistics. To maintain accurate accounting totals, the software maintains two accounting databases: an active database and a checkpointed database.

On R20, configure the **ip accounting** command for the outbound transit packets on the interfaces that are connected to VLAN 82 and VLAN 92:

```
interface Ethernet0/0.82
 encapsulation dot1Q 82
 ip address 12.4.82.20 255.255.255.0
 ip accounting output-packets
 ipv6 address 12:4:82::20/64
!
interface Ethernet0/0.92
 encapsulation dot1Q 92
 ip address 12.4.92.20 255.255.255.0
 ip accounting output-packets
 ipv6 address 12:4:92::20/64
!
```

From R21, ping any destination that is reachable via R20:

```
R21#  ping 192.168.16.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.16.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
R21#  ping 192.168.17.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.17.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
R21#
```

Verify the IP statistics on R20:

```
R20#show ip accounting
   Source           Destination              Packets               Bytes
 12.4.0.21        12.4.0.19                      5                 500
 12.4.0.21        12.4.0.18                      5                 500
 192.168.21.1     12.3.76.17                    78                8050
 192.168.21.1     12.3.76.16                    77                8050
 12.4.201.21      192.168.17.1                   5                 500
 12.4.201.21      192.168.16.1                   5                 500

Accounting data age is 1w2d
R20#
```

Note	Every ping sends five 100-byte IP packets.

IP accounting will not create any entries for the traffic that is originated from the interfaces of R20.

```
R20#
```

Issue:

Configure MAC Statistics Services according to the lab requirements and the "IPv4 IGP" diagram.

Solution:

The MAC address accounting function provides accounting information for IP traffic based on the source and destination MAC addresses on LAN interfaces. MAC accounting calculates the total packet and byte counts for a LAN interface that receives IP packets from or sends IP packets to a unique MAC address. It also records a timestamp for the last packet that is received or sent. For example, with IP MAC accounting, you can determine how much traffic is being sent to peers, received from peers, or both.

MAC address accounting is configured on the physical interfaces. This lab requires that you measure both inbound and outbound traffic.

On R20:

```
interface Ethernet0/0
 no ip address
 ip accounting mac-address input
 ip accounting mac-address output
```

Verify the MAC statistics on R20:

```
R20#show interfaces ethernet0/0 mac-accounting
Ethernet0/0
  Input   (509 free)
    aabb.cc00.1500(200):  185 packets, 22756 bytes, last: 3845317ms ago
    aabb.cc00.1300(206):  191 packets, 17246 bytes, last: 3854311ms ago
    aabb.cc00.1200(207):  5 packets, 590 bytes, last: 86890324ms ago
                  Total:  381 packets, 40592 bytes
  Output   (509 free)
    aabb.cc00.1500(200):  204 packets, 18648 bytes, last: 3845317ms ago
    aabb.cc00.1300(206):  83 packets, 10044 bytes, last: 13743109ms ago
    aabb.cc00.1200(207):  92 packets, 11206 bytes, last: 3853855ms ago
                  Total:  379 packets, 39898 bytes
R20#
```

