Task 1.1

SW1:

```
monitor session 1 source vlan 1011 rx
monitor session 1 destination interface fastEthernet 0/12 ingress vlan 5
```

Task 1.1 Verification

```
Rack1SW1#show monitor session 1
Session 1
---------
Type              : Local Session
Source VLANs      :
    RX Only       : 1011
Destination Ports : Fa0/12
    Encapsulation : Native
          Ingress : Enabled, default VLAN = 5
    Ingress encap : Untagged
```

Task 1.2

SW1:

```
Rack1SW1#mkdir archive
Create directory filename [archive]?
Created dir flash:archive
Rack1SW1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1SW1(config)#alias exec backup copy running-config
flash:/archive/backup.config
Rack1SW1(config)#boot config-file flash:/archive/backup.config
```

Task 1.2 Verification

```
Rack1SW1#dir flash:
Directory of flash:/

    2  -rwx     7963136   Jan 1 1970 02:44:50 +00:00  c3560-advipservicesk9-
mz.122-25.SEE2.bin
    3  -rwx        1197   Mar 1 1993 00:05:09 +00:00  config.old
    4  -rwx         856   Mar 1 1993 00:02:01 +00:00  vlan.dat
    5  -rwx        1914   Mar 1 1993 00:02:05 +00:00  config.text
    7  -rwx         831   Mar 1 1993 23:54:15 +00:00  log.txt
    8  drwx          64   Mar 1 1993 00:45:57 +00:00  archive
   10  -rwx          24   Mar 1 1993 00:45:57 +00:00  private-config.text

32514048 bytes total (24540672 bytes free)

Rack1SW1#show aliases | include backup
  backup               copy running-config flash:/archive/backup.config
```

```
Rack1SW1#show boot
BOOT path-list      : flash:c3560-advipservicesk9-mz.122-25.SEE2.bin
Config file         : flash:/archive/backup.config
Private Config file : flash:/private-config.text
Enable Break        : no
Manual Boot         : no
HELPER path-list    :
Auto upgrade        : yes
```

Task 1.3

R5:

```
interface FastEthernet0/0
 mac-address 0000.0c12.3456
```

SW1:

```
interface FastEthernet0/5
 switchport mode access
 switchport port-security
 switchport port-security mac-address sticky
```

Task 1.3 Verification

```
Rack1SW1(config)#interface f0/5
Rack1SW1(config-if)#shutdown
%LINK-5-CHANGED: Interface FastEthernet0/5, changed state to administratively
      down
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state
      to down
Rack1SW1(config-if)#switchport port-security
Rack1SW1(config-if)#switchport port-security mac-address sticky

Rack1R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R5(config)#interface e0/0
Rack1R5(config-if)#mac-address 0000.0c12.3456
Rack1R5(config-if)#

Rack1SW1(config-if)#no shutdown
%LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5, changed state
      to up
Rack1SW1(config-if)#do show run interface fa0/5
Building configuration...

Current configuration : 231 bytes
!
interface FastEthernet0/5
 switchport access vlan 5
 switchport mode access
 switchport port-security
 switchport port-security mac-address sticky
 switchport port-security mac-address sticky 0000.0c12.3456
```

```
 no ip address
end
```

Task 1.4

SW3 and SW4:
```
no spanning-tree vlan 1363
```

SW4:
```
interface FastEthernet0/20
 switchport backup interface Fa0/21
```

Task 1.4 Verification

```
Rack1SW4#show interface f0/20 switchport backup

Switch Backup Interface Pairs:

Active Interface        Backup Interface       State
------------------------------------------------------------------------
FastEthernet0/20        FastEthernet0/21       Active Up/Backup Standby

Rack1SW4#show spanning-tree vlan 1363

Spanning tree instance(s) for vlan 1363 does not exist.

Rack1SW4#
```

🕮 Further Reading

Configuring Flex Links

Task 1.5

SW3:

```
interface range fa0/14, 15
 no shutdown
 switchport mode access
 switchport access vlan 1363
```

SW1:

```
interface fastEthernet 0/1
 switchport access vlan 4
!
interface fastEthernet 0/17
 switchport access vlan 4
 no shut
 no cdp enable
!
interface fastEthernet 0/3
 switchport access vlan 42
!
interface fastEthernet 0/18
 switchport access vlan 42
 no shutdown
 no cdp enable
```

Task1.5 Verification

For testing, you can ping from R1 to R3 and look at the output of **show interface counters** on SW3 while the ping traffic is passing. If the traffic is passing through SW3, you should see activity that corresponds. The section states that VLANs can't be added, but does not prohibit you from using existing VLANs. In this example, CDP is disabled so that you don't get error messages about the VLANs not matching on both sides.

```
Rack1R1#ping 204.12.1.3 repeat 100000

Rack1SW3#show int counters | i Port|0/1(4|5)
Port            InOctets   InUcastPkts   InMcastPkts   InBcastPkts
Fa0/14          12527415        105838           608             4
Fa0/15          12514046        105970           176             0
Port           OutOctets  OutUcastPkts  OutMcastPkts  OutBcastPkts
Fa0/14          12573479        105841          1075            79
Fa0/15          12619973        106018          1507            83
Rack1SW3#
```

Task 1.6

> **Note**
>
> The EIGRP requirement is to have separate passwords for two different neighbors on the same subnet. By using logical interfaces, this can be achieved. If you just configure using the physical interfaces, you may need to reconfigure later.

R1:

```
interface Virtual-Template13
ip address 167.1.135.1 255.255.255.0
!
interface Serial0/0
 encapsulation frame-relay
 frame-relay interface-dlci 103 ppp Virtual-Template13
```

R3:

```
interface Virtual-Template13
ip address 167.1.135.3 255.255.255.0
!
interface Virtual-Template35
ip address 167.1.135.3 255.255.255.0
!
interface Serial1/0
 encapsulation frame-relay
 frame-relay interface-dlci 301 ppp Virtual-Template13
 frame-relay interface-dlci 305 ppp Virtual-Template35
```

R5:

```
interface Virtual-Template35
ip address 167.1.135.5 255.255.255.0
!
interface Serial0/0
 encapsulation frame-relay
 frame-relay interface-dlci 503 ppp Virtual-Template35
```

Task 1.6 Verification

```
Rack1R3#show frame-relay pvc 301

PVC Statistics for interface Serial1/0 (Frame Relay DTE)

DLCI = 301, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial1/0

  input pkts 30            output pkts 19           in bytes 6188
  out bytes 334            dropped pkts 0           in pkts dropped 0
  out pkts dropped 0                out bytes dropped 0
  in FECN pkts 0           in BECN pkts 0           out FECN pkts 0
  out BECN pkts 0          in DE pkts 0             out DE pkts 0
  out bcast pkts 0         out bcast bytes 0
  5 minute input rate 0 bits/sec, 0 packets/sec
```

```
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 01:44:02, last time pvc status changed 00:48:48
  Bound to Virtual-Access1 (up, cloned from Virtual-Template13)
```

```
Rack1R3#show frame-relay pvc 305

PVC Statistics for interface Serial1/0 (Frame Relay DTE)

DLCI = 305, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial1/0

  input pkts 33            output pkts 48           in bytes 8124
  out bytes 8370           dropped pkts 0           in pkts dropped 0
  out pkts dropped 0                out bytes dropped 0
  in FECN pkts 0           in BECN pkts 0           out FECN pkts 0
  out BECN pkts 0          in DE pkts 0             out DE pkts 0
  out bcast pkts 24        out bcast bytes 7968
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 01:44:05, last time pvc status changed 01:28:31
  Bound to Virtual-Access2 (up, cloned from Virtual-Template35)

Rack1R3#ping 167.1.135.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 167.1.135.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/36 ms

Rack1R3#ping 167.1.135.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 167.1.135.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/60/64 ms
```

Task 1.7

R4:

```
interface Loopback45
 ip address 167.1.45.4 255.255.255.255
!
Interface Multilink 1
 Ip unnumbered Loopback45

interface Serial0/1/0
 encap ppp
 ppp multilink
 ppp multilink group 1
 encapsulation ppp
```

R5:

```
interface Loopback45
 ip address 167.1.45.5 255.255.255.255
!
Interface Multilink 1
 Ip unnumbered Loopback45
interface Serial0/1/0
 encapsulation ppp
 ppp multilink
 ppp multilink group 1
```

```
clockrate 64000
```

Task 1.7 Verification

Verify the PPP peer-neighbor route:

```
Rack1R4#show ip route
<snip>

C    192.10.1.0/24 is directly connected, FastEthernet0/1
     167.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
C       167.1.34.0/24 is directly connected, Serial0/0
C       167.1.45.5/32 is directly connected, Multilink1
C       167.1.45.4/32 is directly connected, Loopback45
C       167.1.4.0/24 is directly connected, FastEthernet0/0
     150.1.0.0/24 is subnetted, 1 subnets
C       150.1.4.0 is directly connected, Loopback0
```

Verify connectivity:

```
Rack1R4#ping 167.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 167.1.45.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

Task 2.1

R4:

```
key chain RIP
 key 1
  key-string CISCO
!
interface FastEthernet0/1
 ip rip authentication mode md5
 ip rip authentication key-chain RIP
 ip rip v2-broadcast
!
router rip
 version 2
 no auto-summary
 network 192.10.1.0
```

Task 2.1 Breakdown

> **Note**
>
> RIPv2 updates are typically sent to the multicast address 224.0.0.9. However, these packets can be sent to the all subnet broadcast address of 255.255.255.255 by issuing the **`ip rip v2-broadcast`** interface level command.

Task 2.1 Verification

```
Rack1R4#show ip protocols
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 10 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/1           2     2                    RIP
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    192.10.1.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.10.1.254         120      00:00:06
  Distance: (default is 120)
```

Verify the RIP updates:

```
Rack1R4#debug ip rip
RIP protocol debugging is on
RIP: sending v2 update to 255.255.255.255 via FastEthernet0/1 (192.10.1.4)
RIP: build update entries - suppressing null update
RIP: received packet with MD5 authentication
RIP: received v2 update from 192.10.1.254 on FastEthernet0/1
      205.90.31.0/24 via 0.0.0.0 in 7 hops
      220.20.3.0/24 via 0.0.0.0 in 7 hops
      222.22.2.0/24 via 0.0.0.0 in 7 hops
```

Task 2.2

R2, R5, SW1, and SW2:

```
router ospf 1
 auto-cost reference-bandwidth 1000
```

Previous Reference

OSPF Reference Bandwidth: Lab 3

Task 2.2 Verification

```
Rack1SW2#show ip ospf interface port-channel 1
Port-channel1 is up, line protocol is up (connected)
  Internet Address 167.1.78.8/24, Area 2578
  Process ID 1, Router ID 150.1.8.8, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 150.1.8.8, Interface address 167.1.78.8
  Backup Designated router (ID) 150.1.7.7, Interface address 167.1.78.7
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:04
  Supports Link-local Signaling (LLS)
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.7.7  (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

Task 2.3

SW3:

```
no ip igmp snooping vlan 1363
!
interface FastEthernet0/24
 ip access-group DENY_EIGRP in
!
ip access-list extended DENY_EIGRP
 deny   eigrp any any
 permit ip any any
!
mac-address-table static 0100.5e00.000a vlan 1363 int fa0/14 fa0/15 fa0/20
fa0/21
```

Task 2.3 Breakdown

Configuring the static MAC entry will limit the destinations for the multicast MAC to the ports where the other EIGRP neighbors are reachable, and will prevent the hello traffic being sent out Fa0/24 to BB3. Configuring the access list will block the traffic inbound from BB3. The OUI for multicast mapping is 01:00:5E, and the last 23bits of the address correspond to the last 23 bits of the Ethernet address. EIGRP's hello traffic uses the multicast address 224.0.0.10, the last three octets are 0.0.10, which converts to 00:00:0a for the last 23 bits in hex. Your configuration for this section may vary, depending on how you configured section 1.5

Task 2.3 Verification

Check EIGRP neighbors:

```
Rack1R6#show ip eigrp neighbors
IP-EIGRP neighbors for process 10
H   Address          Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                     (sec)         (ms)       Cnt Num
1   204.12.1.3       Fa0/0             14 00:00:24  339  2034  0  26
0   204.12.1.1       Fa0/0             13 00:00:24  384  2304  0  13
```

Check EIGRP routes:

```
Rack1R3#show ip route eigrp
     167.1.0.0/16 is variably subnetted, 7 subnets, 2 masks
D       167.1.45.5/32 [90/21024000] via 167.1.34.4, 00:14:25, Serial1/1.34
D       167.1.45.4/32 [90/20640000] via 167.1.34.4, 00:14:25, Serial1/1.34
     150.1.0.0/24 is subnetted, 4 subnets
D       150.1.6.0 [90/409600] via 204.12.1.6, 00:14:12, FastEthernet0/0
D       150.1.4.0 [90/20640000] via 167.1.34.4, 00:14:25, Serial1/1.34
D       150.1.1.0 [90/409600] via 204.12.1.1, 00:14:25, FastEthernet0/0
```

The following verification could not be reproduced in the rental racks or in the real lab, as it configures the backbone router. However, it is presented here to illustrate how the task could potentially be verified.

BB3:

```
router eigrp 10
 network 204.12.1.0
!
access-list 100 permit eigrp 204.12.1.0 0.0.0.255 any

BB3#debug ip packet detail 100

IP: s=204.12.1.6 (Ethernet0), d=224.0.0.10, len 60, rcvd 2, proto=88
IP: s=204.12.1.3 (Ethernet0), d=224.0.0.10, len 60, rcvd 2, proto=88
IP: s=204.12.1.1 (Ethernet0), d=224.0.0.10, len 60, rcvd 2, proto=88
IP: s=204.12.1.254 (local), d=224.0.0.10 (Ethernet0), len 60, sending
broad/multicast, proto=88
IP: s=204.12.1.6 (Ethernet0), d=224.0.0.10, len 60, rcvd 2, proto=88
IP: s=204.12.1.3 (Ethernet0), d=224.0.0.10, len 60, rcvd 2, proto=88
```

```
BB3#show ip eigrp neighbors
IP-EIGRP neighbors for process 10
H   Address      Interface       Hold Uptime   SRTT   RTO  Q  Seq Type
                                 (sec)         (ms)       Cnt Num
2   204.12.1.1    Et0               3 00:03:12 1439  5000  0  15
1   204.12.1.3    Et0               1 00:03:12   24   200  0  28
0   204.12.1.6    Et0               1 00:03:12   19   200  0  10
```

Enable filtering and check debugging output again:

```
BB3#
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 10: Neighbor 204.12.1.6 (Ethernet0) is down:
holding time expired
destroy peer: 204.12.1.6
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 10: Neighbor 204.12.1.3 (Ethernet0) is down:
holding time expired
destroy peer: 204.12.1.3
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 10: Neighbor 204.12.1.1 (Ethernet0) is down:
holding time expired
destroy peer: 204.12.1.1

BB3#debug ip packet detail 100
IP packet debugging is on (detailed) for access list 100
BB3#
IP: s=204.12.1.254 (local), d=224.0.0.10 (Ethernet0), len 60, sending
broad/multicast, proto=88
IP: s=204.12.1.254 (local), d=224.0.0.10 (Ethernet0), len 60, sending
broad/multicast, proto=88
IP: s=204.12.1.254 (local), d=224.0.0.10 (Ethernet0), len 60, sending
broad/multicast, proto=88
IP: s=204.12.1.254 (local), d=224.0.0.10 (Ethernet0), len 60, sending
broad/multicast, proto=88
```

Task 2.4

R4:

```
interface Serial0/1/0
 bandwidth 1536
 ip bandwidth-percent eigrp 10 25
!
```

R1:

```
key chain EIGRP
 key 13
  key-string CISCO13
!
interface Virtual-Template13
 ip authentication mode eigrp 10 md5
 ip authentication key-chain eigrp 10 EIGRP
!
router eigrp 10
 network 167.1.135.1 0.0.0.0
```

R3:

```
key chain EIGRP13
  key 13
   key-string CISCO13
!
key chain EIGRP35
  key 35
    key-string CISCO35
!
interface Virtual-Template13
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 EIGRP13
!
interface Virtual-Template35
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 EIGRP35
!
router eigrp 10
  network 167.1.135.3 0.0.0.0
```

R5:

```
key chain EIGRP
 key 35
  key-string CISCO35
!
interface Virtual-Template35
 ip authentication mode eigrp 10 md5
 ip authentication key-chain eigrp 10 EIGRP
!
interface Serial0/1/0
 bandwidth 1536
 ip bandwidth-percent eigrp 10 25

router eigrp 10
 network 167.1.135.5 0.0.0.0
```

Task 2.4 Verification

Verify EIGRP authentication:

```
Rack1R3#show ip eigrp interfaces detail virtual-access 1
IP-EIGRP interfaces for process 10

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface Peers    Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Vi1         1        0/0         72       0/10          314           0
  Hello interval is 5 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 1/4
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 1  Out-of-sequence rcvd: 0
  Authentication mode is md5,  key-chain is "EIGRP13"

Rack1R3#show ip eigrp interfaces detail virtual-access 2
IP-EIGRP interfaces for process 10

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface Peers    Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Vi2         1        0/0        1320       0/10         6538          0
  Hello interval is 5 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 2/3
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 1
  Retransmissions sent: 0  Out-of-sequence rcvd: 1
  Authentication mode is md5,  key-chain is "EIGRP35"
```

Verify the EIGRP neighbors:

```
Rack1R3#show ip eigrp neighbors
IP-EIGRP neighbors for process 10
H   Address         Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                    (sec)         (ms)       Cnt Num
5   167.1.135.5     Vi2               11 00:01:35 1320  5000  0  17
4   167.1.135.1     Vi1               12 00:01:39   72   432  0  23
3   204.12.1.6      Fa0/0             14 00:21:59    1   200  0  14
2   204.12.1.1      Fa0/0             12 00:22:08  277  1662  0  25
1   167.1.34.4      Se1/1.34         158 00:30:47  203  1218  0  28
0   167.1.13.1      Se1/2             14 00:31:26   24  1140  0  24
```

Task 2.5

R1:

```
router eigrp 10
 eigrp stub connected
 timers active-time 1
```

Task 2.5 Verification

```
        Accessed by p_saffari@yahoo.com from 94.182.214.117 at 05:26:47 Oct 21, 2009
Rack1R3#show ip eigrp neighbors detail | include CONNECTED
```



```
Stub Peer Advertising ( CONNECTED ) Routes
```

Task 2.6

```
R4:
interface FastEthernet0/1
 ip summary-address rip 167.1.0.0 255.255.0.0
 ip summary-address rip 150.1.0.0 255.255.240.0
!
router eigrp 10
 redistribute rip metric 10000 10 255 1 1500
!
router rip
 redistribute connected metric 1
 redistribute eigrp 10 metric 1

R5:
interface Virtual-Template35
 ip summary-address eigrp 10 0.0.0.0 0.0.0.0 5
!
interface Multilink1
 ip summary-address eigrp 10 0.0.0.0 0.0.0.0 5
!
router ospf 1
 default-information originate always
```

Task 2.6 Verification

Check for the default route:

```
Rack1R4#show ip route | begin Gate
Gateway of last resort is 167.1.45.5 to network 0.0.0.0

R    222.22.2.0/24 [120/7] via 192.10.1.254, 00:00:14, FastEthernet0/1
D    204.12.1.0/24 [90/2195456] via 167.1.34.3, 00:32:58, Serial0/0/0
R    220.20.3.0/24 [120/7] via 192.10.1.254, 00:00:14, FastEthernet0/1
C    192.10.1.0/24 is directly connected, FastEthernet0/1
     167.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D       167.1.135.1/32 [90/4729856] via 167.1.34.3, 00:07:02, Serial0/0/0
D       167.1.135.0/24 [90/4729856] via 167.1.34.3, 00:07:02, Serial0/0/0
D       167.1.135.5/32 [90/4729856] via 167.1.34.3, 00:16:39, Serial0/0/0
C       167.1.34.0/24 is directly connected, Serial0/0/0
C       167.1.45.5/32 is directly connected, Multilink1
C       167.1.45.4/32 is directly connected, Loopback45
C       167.1.4.0/24 is directly connected, FastEthernet0/0
D       167.1.13.0/24 [90/21024000] via 167.1.34.3, 00:41:31, Serial0/0/0
     150.1.0.0/24 is subnetted, 4 subnets
D       150.1.6.0 [90/2323456] via 167.1.34.3, 00:32:40, Serial0/0/0
C       150.1.4.0 is directly connected, Loopback0
D       150.1.3.0 [90/2297856] via 167.1.34.3, 00:07:03, Serial0/0/0
D       150.1.1.0 [90/2323456] via 167.1.34.3, 00:09:57, Serial0/0/0
R    205.90.31.0/24 [120/7] via 192.10.1.254, 00:00:15, FastEthernet0/1
D*   0.0.0.0/0 [90/2306560] via 167.1.45.5, 00:07:03, Multilink1
```

```
Rack1SW1#show ip route ospf
     167.1.0.0/24 is subnetted, 3 subnets
O       167.1.58.0 [110/310] via 167.1.78.8, 00:07:18, Port-channel1
     150.1.0.0/24 is subnetted, 3 subnets
O       150.1.5.0 [110/311] via 167.1.78.8, 00:07:18, Port-channel1
O       150.1.2.0 [110/31] via 167.1.27.2, 00:07:18, FastEthernet0/2
O*E2 0.0.0.0/0 [110/1] via 167.1.78.8, 00:07:18, Port-channel1
```

Finally, test full connectivity with the following Tcl script:

```
foreach i {
167.1.135.1
150.1.1.1
167.1.13.1
204.12.1.1
150.1.2.2
167.1.27.2
167.1.135.3
167.1.34.3
150.1.3.3
167.1.13.3
204.12.1.3
167.1.34.4
167.1.45.4
150.1.4.4
192.10.1.4
167.1.135.5
167.1.45.5
150.1.5.5
167.1.58.5
150.1.6.6
204.12.1.6
150.1.7.7
167.1.27.7
167.1.78.7
167.1.58.8
150.1.8.8
167.1.78.8
222.22.2.1
} {ping $i}
```

Note that VLAN4, VLAN5, and Serial link from R6 to BB1 are excluded from connectivity test. Also, SW3 will not have reachability until later in the lab.

Task 2.7

R1:

```
router bgp 100
 neighbor 150.1.3.3 remote-as 100
 neighbor 150.1.3.3 update-source Loopback0
```

R3:
```
router bgp 100
 neighbor iBGP peer-group
 neighbor iBGP remote-as 100
 neighbor iBGP update-source Loopback0
 neighbor iBGP route-reflector-client
 neighbor iBGP send-community
 neighbor 150.1.1.1 peer-group iBGP
 neighbor 150.1.4.4 peer-group iBGP
 neighbor 167.1.135.5 peer-group iBGP
 neighbor 150.1.6.6 peer-group iBGP
 neighbor 150.1.9.9 peer-group iBGP
 neighbor 150.1.9.9 shutdown
 neighbor 150.1.10.10 peer-group iBGP
 neighbor 150.1.10.10 shutdown
```

R4:
```
router bgp 100
 neighbor 150.1.3.3 remote-as 100
 neighbor 150.1.3.3 update-source Loopback0
```

R5:
```
router bgp 100
 neighbor 150.1.3.3 remote-as 100
```

R6:
```
router bgp 100
 neighbor 150.1.3.3 remote-as 100
 neighbor 150.1.3.3 update-source Loopback0
 neighbor 150.1.3.3 next-hop-self
```

Task 2.7 Breakdown

BGP peer groups are a way to minimize redundant configuration between neighbors that share common attributes. For example, R3 is peering with R1, R4, R5, R6, and two additional devices. These devices are all in AS 100 and are route-reflector clients of R3. Instead of specifying two neighbor statements applying the **remote-as** and **route-reflector-client** options, a peer group has been defined that has these options applied. Then, instead of applying the options directly on the neighbor, the neighbor is simply specified as part of the predefined peer-group.

The **shutdown** option of the BGP neighbor command is typically used for the case that is described in this task. For example, a new circuit may be on order that involves a BGP peering session. Instead of waiting until the circuit is installed and up, the BGP configuration can be applied beforehand, and the neighbor disabled with the **neighbor [*address*] shutdown** option. Therefore, the only configuration that is required once the new circuit is up is to issue a **no** statement for the command with the shutdown applied.

Task 2.7 Verification

Verify the BGP neighbors:

```
Rack1R3#show ip bgp summary | begin Neighbor
Neighbor     V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
150.1.1.1    4 100   8       9       12    0    0 00:00:57       10
150.1.4.4    4 100   4       9       12    0    0 00:00:41        0
150.1.6.6    4 100   9       9       12    0    0 00:00:02       11
150.1.9.9    4 100   0       0        0    0    0 never    Idle (Admin)
150.1.10.10  4 100   0       0        0    0    0 never    Idle (Admin)
167.1.135.5  4 100   4       9       12    0    0 00:00:11        0
204.12.1.254 4 54    22     18       12    0    0 00:13:11       10
```

Task 2.8

R4:

```
router bgp 100
 neighbor 192.10.1.254 remote-as 254
 neighbor 192.10.1.254 local-as 200
 neighbor 192.10.1.254 password CISCO
```

Task 2.8 Verification

```
Rack1R4#show ip bgp summary | begin Neighbor
Neighbor     V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down  State/PfxRcd
150.1.3.3    4 100  15       9       15    0    0 00:04:23       11
192.10.1.254 4 254   5       8       12    0    0 00:00:27        3
```

Check local-AS configuration:

```
Rack1R4#show ip bgp  neighbors 192.10.1.254
BGP neighbor is 192.10.1.254,  remote AS 254,  local AS 200, external link
  BGP version 4, remote router ID 222.22.2.1
  BGP state = Established, up for 00:01:03
  Last read 00:00:02, last write 00:00:02, hold time is 180, keepalive interval
is 60 seconds
<output omitted>
```

Check for any prepended AS:

```
Rack1R4#show ip bgp quote-regexp _254$
BGP table version is 15, local router ID is 150.1.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 205.90.31.0      192.10.1.254             0             0 200 254 ?
*> 220.20.3.0       192.10.1.254             0             0 200 254 ?
*> 222.22.2.0       192.10.1.254             0             0 200 254 ?
```

Task 2.9

R4:

```
router bgp 100
 neighbor 192.10.1.254 local-as 200 no-prepend
```

Previous Reference

BGP Local AS Feature: Lab 2

Task 2.9 Verification

Confirm that AS 200 is not prepended:

```
Rack1R4#show ip bgp quote-regexp _254$
BGP table version is 21, local router ID is 150.1.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 205.90.31.0      192.10.1.254             0             0 254 ?
*> 220.20.3.0       192.10.1.254             0             0 254 ?
*> 222.22.2.0       192.10.1.254             0             0 254 ?
```

Task 2.10

R1:

```
router bgp 100
 neighbor 204.12.1.254 route-map TO_BB3 out
!
ip prefix-list VLAN4_AND_VLAN5 seq 5 permit 167.1.4.0/23  ge 24 le 24
!
route-map TO_BB3 permit 10
 match ip address prefix-list VLAN4_AND_VLAN5
 set as-path prepend 100 100
!
route-map TO_BB3 permit 1000
```

R3:

```
router bgp 100
 neighbor 204.12.1.254 route-map TO_BB3 out
!
ip prefix-list VLAN4_AND_VLAN5 seq 5 permit 167.1.4.0/23 ge 24 le 24
!
route-map TO_BB3 permit 10
 match ip address prefix-list VLAN4_AND_VLAN5
set as-path prepend 100 100
```

```
!
route-map TO_BB3 permit 1000
```

R4:
```
router bgp 100
 network 167.1.4.0 mask 255.255.255.0
```

R5:
```
router bgp 100
 network 167.1.5.0 mask 255.255.255.0
```

R6:
```
router bgp 100
 neighbor 204.12.1.254 route-map TO_BB3 out
!
ip prefix-list VLAN4_AND_VLAN5 seq 5 permit 167.1.4.0/23 ge 24 le 24
!
route-map TO_BB3 permit 10
 match ip address prefix-list VLAN4_AND_VLAN5
 set as-path prepend 100 100
!
route-map TO_BB3 permit 1000
```

Task 2.10 Verification

Verify the BGP tables of BB1 and BB3:

```
BB1 >show ip bgp quote-regexp _100$
BGP table version is 987, local router ID is 212.18.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 167.1.4.0/24     54.1.1.6                               0 100 i
*> 167.1.5.0/24     54.1.1.6                               0 100 i

BB3>show ip bgp quote-regexp _100$
BGP table version is 35, local router ID is 31.3.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight Path
*>i167.1.4.0/24     172.16.4.1           0    100      0 100 i
*                   204.12.1.3                        0 100 100 100 i
*                   204.12.1.3                        0 100 100 100 i
*                   204.12.1.3                        0 100 100 100 i
*>i167.1.5.0/24     172.16.4.1           0    100      0 100 i
*                   204.12.1.3                        0 100 100 100 i
*                   204.12.1.3                        0 100 100 100 i
*                   204.12.1.3                        0 100 100 100 i
```

Task 2.11

SW1:
```
router bgp 65078
 network 150.1.7.0 mask 255.255.255.0
```

SW2:
```
router bgp 65078
 network 150.1.8.0 mask 255.255.255.0
 aggregate-address 150.1.0.0 255.255.240.0 summary-only
```

Task 2.11 Verification

Check for the summary received from SW2:

```
Rack1R5#show ip bgp neighbors 167.1.58.8 routes
BGP table version is 31, local router ID is 150.1.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 150.1.0.0/20     167.1.58.8               0             0 65078 i

Total number of prefixes 1
```

Task 2.12

R1, R3, and R6:
```
router bgp 100
 neighbor 204.12.1.254 remove-private-as
```

R4:
```
router bgp 100
 neighbor 192.10.1.254 remove-private-as
```

R6:
```
router bgp 100
neighbor 54.1.1.254 remove-private-as
```

Task 2.12 Verification

Check AS-path for aggregated prefix on BB1:

```
BB1>show ip bgp 150.1.0.0
BGP routing table entry for 150.1.0.0/20, version 990
Paths: (2 available, best #2, table Default-IP-Routing-Table)
Flag: 0x840
  Advertised to non peer-group peers:
  172.16.4.3
  100, (aggregated by 65078 150.1.8.8)
    172.16.4.3 from 172.16.4.3 (31.3.0.1)
      Origin IGP, metric 0, localpref 100, valid, internal, atomic-aggregate
  100, (aggregated by 65078 150.1.8.8)
    54.1.1.6 from 54.1.1.6 (150.1.6.6)
      Origin IGP, localpref 100, valid, external, atomic-aggregate, best
```

> **Caution**
>
> As mentioned in previous labs, you will not have access to the BB routers to execute commands on during the real lab.

Task 2.13

R3:

```
router bgp 100
 bgp inject-map ORIGINATE exist-map LEARNED_PATH
 neighbor iBGP next-hop-self
 neighbor iBGP route-map TO_IBGP_PEERS out
!
ip prefix-list ORIGINATED_ROUTES seq 10 permit 150.1.8.0/24
ip prefix-list ROUTE seq 5 permit 150.1.0.0/20
ip prefix-list ROUTE_SOURCE seq 5 permit 167.1.135.5/32
ip prefix-list SPECIFIC_ROUTES seq 10 permit 150.1.8.0/24
!
route-map LEARNED_PATH permit 10
 match ip address prefix-list ROUTE
 match ip route-source prefix-list ROUTE_SOURCE
!
route-map ORIGINATE permit 10
 set ip address prefix-list ORIGINATED_ROUTES
!
route-map TO_IBGP_PEERS deny 10
 match ip address prefix-list SPECIFIC_ROUTES
!
route-map TO_IBGP_PEERS permit 1000
```

R6:

```
router bgp 100
 bgp inject-map ORIGINATE exist-map LEARNED_PATH
 neighbor 150.1.3.3 route-map TO_R3 out
 neighbor 204.12.1.254 route-map TO_R3 out

!
ip prefix-list ORIGINATED_ROUTES seq 10 permit 150.1.7.0/24
ip prefix-list ROUTE seq 5 permit 150.1.0.0/20
ip prefix-list ROUTE_SOURCE seq 5 permit 150.1.3.3/32
ip prefix-list SPECIFIC_ROUTES seq 5 permit 150.1.7.0/24
!
route-map LEARNED_PATH permit 10
 match ip address prefix-list ROUTE
 match ip route-source prefix-list ROUTE_SOURCE
!
route-map TO_R3 deny 10
 match ip address prefix-list SPECIFIC_ROUTES
!
route-map TO_R3 permit 1000
!
route-map ORIGINATE permit 10
 set ip address prefix-list ORIGINATED_ROUTES
!
route-map TO_BB3 deny 5
 match ip address prefix-list SPECIFIC_ROUTES
```

Task 2.13 Breakdown

The BGP conditional route injection feature allows a router to originate an arbitrary network block based on the existence of a prefix in the BGP table. This feature is designed to be used in the case that is described in this task.

In the above task, AS 100 is learning the aggregate block 150.1.0.0/20 from AS 65078. Since AS 100 has multiple exit points to AS 54, it may be desirable for AS 100 to create a traffic engineering policy based on longer matches. By re-injecting subnets that make up the aggregate, AS 100 can force it's upstream peers (AS 54 in this case) to follow a forwarding policy based on the longer match to the destination.

The BGP conditional route injection feature relies on two parts, the inject-map and the exist-map. When the prefix and route-source matched in the exist-map exist in the BGP table, the prefix or prefixes set in the inject-map are injected into the BGP table.

Task 2.13 Verification

Verify the BGP prefix injection:

```
Rack1R6#show ip bgp injected-paths
BGP table version is 18, local router ID is 150.1.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i150.1.7.0/24     167.1.58.8                             0 ?

Rack1R3#show ip bgp injected-paths
BGP table version is 32, local router ID is 150.1.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i150.1.8.0/24     167.1.58.8                             0 ?
```

Verify the specific prefix advertisements:

```
Rack1R3#show ip bgp neighbors 204.12.1.254 advertised-routes
BGP table version is 32, local router ID is 150.1.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i150.1.0.0/20     167.1.58.8               0    100      0 65078 i
*>i150.1.8.0/24     167.1.58.8                            0 ?
*>i167.1.4.0/24     150.1.4.4                0    100      0 i
*>i167.1.5.0/24     167.1.135.5              0    100      0 i
*>i205.90.31.0      192.10.1.254             0    100      0 254 ?
*>i220.20.3.0       192.10.1.254             0    100      0 254 ?
*>i222.22.2.0       192.10.1.254             0    100      0 254 ?

Rack1R6#show ip bgp neighbors 204.12.1.254 advertised-routes
BGP table version is 18, local router ID is 150.1.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i150.1.0.0/20     167.1.58.8               0    100      0 65078 i
*>i167.1.4.0/24     150.1.4.4                0    100      0 i
*>i167.1.5.0/24     167.1.135.5              0    100      0 i
*>i205.90.31.0      192.10.1.254             0    100      0 254 ?
*>i220.20.3.0       192.10.1.254             0    100      0 254 ?
*>i222.22.2.0       192.10.1.254             0    100      0 254 ?

Total number of prefixes 6

Rack1R6#show ip bgp neigh 54.1.1.254 advertised-routes
BGP table version is 18, local router ID is 150.1.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 28.119.16.0/24   204.12.1.254             0             0 54 i
*> 28.119.17.0/24   204.12.1.254             0             0 54 i
*> 112.0.0.0        204.12.1.254                           0 54 50 60 i
*> 113.0.0.0        204.12.1.254                           0 54 50 60 i
*> 114.0.0.0        204.12.1.254                           0 54 i
*> 115.0.0.0        204.12.1.254                           0 54 i
*> 116.0.0.0        204.12.1.254                           0 54 i
*> 117.0.0.0        204.12.1.254                           0 54 i
*> 118.0.0.0        204.12.1.254                           0 54 i
*> 119.0.0.0        204.12.1.254                           0 54 i
*>i150.1.0.0/20     167.1.58.8               0    100      0 65078 i
*>i150.1.7.0/24     167.1.58.8                            0 ?
<snip>
```

Task 3.1

```
R6:
ipv6 unicast-routing
!
interface Loopback100
 ipv6 address 2001:150:1:26::6/64
 ipv6 eigrp 101
!
interface Loopback101
 ipv6 address 2001:150:1:2E::6/64
 ipv6 eigrp 101

ipv6 router eigrp 101
 no shut
```

Task 3.1 Verification

```
Rack1R6#show ipv6 interface brief | beg Loop
Loopback100                [up/up]
    FE80::213:C4FF:FEA6:9420
    2001:150:1:26::6
Loopback101                [up/up]
    FE80::213:C4FF:FEA6:9420
    2001:150:1:2E::6

Rack1R6#show ipv6 eigrp 101 interface detail
IPv6-EIGRP interfaces for process 101

                      Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo100              0        0/0         0       0/1            0           0
  Hello interval is 5 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Authentication mode is not set
  Use multicast
Lo101              0        0/0         0       0/1            0           0
  Hello interval is 5 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Authentication mode is not set
  Use multicast
Rack1R6#
```

Task 3.2

R4:
```
ipv6 unicast-routing
!
interface Tunnel46
 ipv6 address 2001:167:1:46::4/64
 tunnel source Loopback0
 tunnel destination 150.1.6.6
 tunnel mode ipv6ip
```

R6:
```
interface Tunnel46
 ipv6 address 2001:167:1:46::6/64
 tunnel source Loopback0
 tunnel destination 150.1.4.4
 tunnel mode ipv6ip
```

Task 3.2 Verification

```
Rack1R6#show interfaces tunnel 46
Tunnel46 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 100 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 150.1.6.6 (Loopback0), destination 150.1.4.4
  Tunnel protocol/transport IPv6/IP
<output omitted>

Rack1R6#ping 2001:167:1:46::4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:167:1:46::4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 68/68/72 ms
```

Task 3.3

R4:
```
interface FastEthernet0/0
 ipv6 address 2001:167:1:4::4/64
 ipv6 eigrp 101
!
interface Tunnel46
 ipv6 eigrp 101

ipv6 router eigrp 101
 no shut
```

R6:
```
interface Tunnel46
 ipv6 eigrp 101
```

```
ipv6 summary-address eigrp 101 0::0/0 7
ipv6 summary-address eigrp 101 2001:150:1:20::0/60 7
```

Task 3.3 Verification

```
Rack1R4#show ipv6 route eigrp
IPv6 Routing Table - Default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, M - MIPv6, R - RIP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
D   ::/0 [90/27008000]
     via FE80::9601:606, Tunnel46
D   2001:150:1:20::/60 [90/27008000]
     via FE80::9601:606, Tunnel46
Rack1R4#
```

Task 5.1

R3, R4 and R5:

```
ip pim rp-address 150.1.4.4 override
```

Task 5.1 Verification

Verify the PIM RP to group mapping:

```
Rack1R4#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static-Override
    RP: 150.1.4.4 (?)

Rack1R3#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static-Override
    RP: 150.1.4.4 (?)

Rack1R5#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static-Override
    RP: 150.1.4.4 (?)
```

Task 5.2

R3:

```
interface Serial1/1.34
 ip dvmrp unicast-routing
!
interface Virtual-Template35
 ip pim sparse
 ip dvmrp unicast-routing
```

R4:
```
interface Tunnel0
 ip unnumbered FastEthernet0/0
 ip pim sparse-mode
 tunnel source FastEthernet0/1
 tunnel destination 220.20.3.192
 tunnel mode dvmrp
!
interface Multilink1
 ip pim sparse
 ip dvmrp unicast-routing
!
interface Serial0/0
 ip dvmrp unicast-routing
```

R5:
```
interface Virtual-Template35
 ip pim sparse
 ip dvmrp unicast-routing
!
interface Multilink1
 ip pim sparse
 ip dvmrp unicast-routing
```

Task 5.2 Verification

Verify the DVMRP routes:

```
Rack1R4#show ip dvmrp route
DVMRP Routing Table - 7 entries
150.1.0.0/16 [0/2] uptime 00:09:21, expires 00:02:09
    via 167.1.34.3, Serial0/0
167.1.5.0/24 [0/2] uptime 00:09:21, expires 00:02:39
    via 167.1.45.5, Multilink1
167.1.45.4/32 [0/3] uptime 00:09:21, expires 00:02:09
    via 167.1.34.3, Serial0/0
167.1.135.0/24 [0/2] uptime 00:09:21, expires 00:02:09
    via 167.1.34.3, Serial0/0
167.1.135.3/32 [0/2] uptime 00:09:21, expires 00:02:39
    via 167.1.45.5, Multilink1
167.1.135.5/32 [0/2] uptime 00:08:51, expires 00:02:09
    via 167.1.34.3, Serial0/0
204.12.1.0/24 [0/2] uptime 00:08:51, expires 00:02:09
    via 167.1.34.3, Serial0/0
Rack1R4#
```

```
Rack1R5#show ip dvmrp route
DVMRP Routing Table - 8 entries
150.1.0.0/16 [0/2] uptime 00:09:28, expires 00:02:44
    via 167.1.45.4, Multilink1
167.1.4.0/24 [0/2] uptime 00:09:28, expires 00:02:44
    via 167.1.45.4, Multilink1
167.1.34.0/24 [0/2] uptime 00:09:28, expires 00:02:44
    via 167.1.45.4, Multilink1
167.1.45.5/32 [0/3] uptime 00:09:28, expires 00:02:32
    via 167.1.135.3, Virtual-Access2
167.1.135.0/24 [0/3] uptime 00:09:28, expires 00:02:44
    via 167.1.45.4, Multilink1
167.1.135.3/32 [0/4] uptime 00:08:28, expires 00:02:32
    via 167.1.135.3, Virtual-Access2
167.1.135.5/32 [0/3] uptime 00:09:28, expires 00:02:44
    via 167.1.45.4, Multilink1
204.12.1.0/24 [0/2] uptime 00:09:28, expires 00:02:32
    via 167.1.135.3, Virtual-Access2
Rack1R5#

Rack1R3#show ip dvmrp route
DVMRP Routing Table - 9 entries
150.1.0.0/16 [0/32] uptime 00:11:21, expires 00:02:57
    via 167.1.135.5, Virtual-Access3
167.1.4.0/24 [0/2] uptime 00:11:21, expires 00:02:39
    via 167.1.34.4, Serial1/1.34
167.1.5.0/24 [0/2] uptime 00:11:03, expires 00:02:57
    via 167.1.135.5, Virtual-Access3
167.1.34.0/24 [0/3] uptime 00:11:21, expires 00:02:57
    via 167.1.135.5, Virtual-Access3
167.1.45.4/32 [0/2] uptime 00:11:03, expires 00:02:57
    via 167.1.135.5, Virtual-Access3
167.1.45.5/32 [0/2] uptime 00:11:21, expires 00:02:39
    via 167.1.34.4, Serial1/1.34
167.1.135.0/24 [0/2] uptime 00:11:03, expires 00:02:57
    via 167.1.135.5, Virtual-Access3
167.1.135.3/32 [0/3] uptime 00:11:03, expires 00:02:39
    via 167.1.34.4, Serial1/1.34
167.1.135.5/32 [0/4] uptime 00:10:03, expires 00:02:57
    via 167.1.135.5, Virtual-Access3
Rack1R3#
```

Task 5.2 Breakdown

Configuring the tunnel and adding DMVRP routing to the interfaces is fairly straightforward. Make sure that you are careful about the interfaces themselves. When using logical interfaces, such as the Virtual Template or Multilink interface, then the multicast settings need to be applied on the logical interface. Section 5.1 stated that PIM was enabled on the Serial interfaces. Creating logical interfaces does not mean that the logical interfaces will inherit these settings, so you need to configure the logical interfaces for sparse mode. You can remove the sparse mode from the physical interfaces if you wish, and just have it configured on the logical interfaces. You should see PIM neighbors, and should see DVMRP routes from both interfaces that connect these devices.

Task 5.3

R4:
```
interface Tunnel0
 ip dvmrp metric 1 list VLAN4_AND_VLAN5
 ip dvmrp summary-address 167.1.4.0 255.255.254.0
 no ip dvmrp auto-summary
 ip dvmrp metric 0 list ALL
!
ip access-list standard VLAN4_AND_VLAN5
 permit 167.1.4.0 0.0.0.255
 permit 167.1.5.0 0.0.0.255
ip access-list standard ALL
 permit 0.0.0.0 255.255.255.255
```

Task 5.3 Verification

Verify the summary generation. Apply the configuration below to R4 Serial0/0 temporarily:

R4:
```
interface Serial0/0
 ip dvmrp metric 1 list VLAN4_AND_VLAN5
 ip dvmrp summary-address 167.1.4.0 255.255.254.0
 no ip dvmrp auto-summary
```

Verify the DVMRP routes on R3:

```
Rack1R3#show ip dvmrp route interface s1/1.34
DVMRP Routing Table - 9 entries
167.1.4.0/23 [0/2] uptime 00:01:45, expires 00:02:14
    via 167.1.34.4, Serial1/1.34
167.1.135.0/24 [0/3] uptime 00:07:50, expires 00:02:10
    via 167.1.34.4, Serial1/1.34
167.1.135.3/32 [0/3] uptime 00:07:50, expires 00:02:10
    via 167.1.34.4, Serial1/1.34
```

Task 6.1

R4:
```
ip tcp intercept list 100
ip tcp intercept watch-timeout 30
ip tcp intercept max-incomplete low 500 high 1000
!
access-list 100 permit tcp any host 167.1.4.119
```

Task 6.1 Verification

For testing, you can temporarily change the address of SW4 on the VLAN. Telnet to protected servers from R3, and R5:

```
Rack1R3#telnet 167.1.4.119 80
Trying 167.1.4.119, 80 ... Open

Rack1R5#telnet 167.1.4.119 80
Trying 167.1.4.119, 80 ... Open

Rack1R4#show tcp intercept connections
Incomplete:
Client             Server                State    Create   Timeout  Mode
167.1.34.3:60228   167.1.4.119:80        SYNSENT  00:00:32 00:00:01 I
167.1.45.5:51445   167.1.4.119:80        SYNSENT  00:00:06 00:00:03 I

Established:
Client             Server                State    Create   Timeout  Mode
```

Task 6.2

```
R6:
interface Serial0/0/0
 ip access-group FROM_BB1 in
!
ip access-list extended FROM_BB1
 deny   ip any any option any-options
 permit ip any any
```

Task 6.2 Verification

To verify, issue ping with ip options enabled from BB1:

```
BB1>ping
Protocol [ip]:
Target IP address: 54.1.1.6
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]: T
Number of timestamps [ 9 ]:
Loose, Strict, Record, Timestamp, Verbose[TV]:
Sweep range of sizes [n]:
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 54.1.1.6, timeout is 2 seconds:
Packet has IP options:  Total option bytes= 40, padded length=40
 Timestamp: Type 0.  Overflows: 0 length 40, ptr 5
  >>Current pointer<<
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)

Unreachable from 54.1.1.6.  Received packet has options
 Total option bytes= 40, padded length=40
 Timestamp: Type 0.  Overflows: 0 length 40, ptr 5
  >>Current pointer<<
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)

Unreachable from 54.1.1.6.  Received packet has options
 Total option bytes= 40, padded length=40
 Timestamp: Type 0.  Overflows: 0 length 40, ptr 5
  >>Current pointer<<
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
  Time= 16:00:00.000 PST (00000000)
<output omitted>
```

Alternatively, there is also a global command “ip options drop” that could be used to drop packets with ip options. Using that command on R6 would have slightly different output on BB1, since the traffic is dropped and unreachables are not sent back to BB1. Output shown below:

```
Packet has IP options:  Total option bytes= 40, padded length=40
 Timestamp: Type 0.  Overflows: 0 length 40, ptr 5
  >>Current pointer<<
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)
  Time= 17:00:00.000 PDT (00000000)

Request 0 timed out
Request 1 timed out
Request 2 timed out
Request 3 timed out
Request 4 timed out
Success rate is 0 percent (0/5)
BB1>ping 54.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 54.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/33/36 ms
```

Task 6.3

> **Note**
>
> In order to configure views, you need to use the global command “enable view” and enter the enable password first.

```
Rack1R4#enable view
Password:

Rack1R4#
```

R4:

```
aaa new-model
username OPERATOR password CISCO
username ADMIN privilege 15 password CISCO
enable secret CISCO
!
! Local authentication is needed to configure the roles.
!
aaa authentication login default none
aaa authentication login VTY local
aaa authorization exec VTY local
!
parser view HTTP
```

```
 secret CISCO
 commands configure include-exclusive all ip http
 commands exec include show run
 commands exec include config t
!
username OPERATOR view HTTP

line vty 0 4
 login authentication VTY
 authorization exec VTY
```

Task 6.3 Breakdown

Role-based CLI allows you to assign commands to views. Views can then be grouped with other views, creating a superview, or can be assigned to users. When configuring Role based CLI, you need to enable AAA, and login using a username and password. The command "**enable view**" will allow you to configure the views with the **parser view** command. Similarly to the functionality seen when moving privilege levels, the command **show run** is a special case and will only allow the user to see items that they have the correct privilege to configure.

Task 6.3 Verification

```
Rack1R4#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

Username: OPERATOR
Password: CISCO

Rack1R4>show parser view
Current view is 'HTTP'
Rack1R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R4(config)>?
Configure commands:
  do    To run exec commands in config mode
  exit  Exit from configure mode
  ip    Global IP configuration subcommands

Rack1R4(config)#ip ?
Global IP configuration subcommands:
  http  HTTP server configuration

Rack1R4(config)#ip http ?
  access-class                 Restrict http server access by access-class
  active-session-modules       Set up active http server session modules
  authentication               Set http server authentication method
  client                       Set http client parameters
  max-connections              Set maximum number of concurrent http server
                               connections
```

```
  path                          Set base path for HTML
  port                          Set http port
  secure-active-session-modules Set up active http secure server session
                                modules

<snip>
```

```
Rack1R4>show run
Building configuration...

Current configuration : 60 bytes
!
!
!
!
no ip http server
no ip http secure-server
!
!
end

Rack1R4>
```

Task 6.4

R6:

```
ip traffic-export profile BB1
 interface FastEthernet0/0
 bidirectional
 mac-address 0060.0060.0060
 incoming sample one-in-every 50
 outgoing sample one-in-every 20

interface Serial 0/0/0
 ip traffic-export apply BB1
```

SW3:

```
mac-address 60.60.60 vlan 1363 drop
```

Task 6.4 Verification

For testing, enable debug ip traffic-export events and generate some traffic on the interface. (PING traffic from R6 to BB1 is sufficient) You can also look at the output of **show ip traffic-export.**

```
RITE: exported output packet # 60
RITE: exported input packet # 30

Rack1R6#show ip traffic-export
Router IP Traffic Export Parameters
Monitored Interface             Serial0/0/0
        Export Interface                FastEthernet0/0
        Destination MAC address 0060.0060.0060
        bi-directional traffic export is on
Output IP Traffic Export Information    Packets/Bytes Exported    112/11080
        Packets Dropped           2132
        Sampling Rate             one-in-every 20 packets
        No Access List configured
Input IP Traffic Export Information     Packets/Bytes Exported    61/5764
        Packets Dropped           3031
        Sampling Rate             one-in-every 50 packets
        No Access List configured
```

```
        Profile BB1 is Active


Rack1R6#
```

Task 7.1

R6:

```
username NOC privilege 15 password 0 CISCO
username NOC autocommand menu NOC
!
menu NOC title #
Menu for Level 1 NOC users
#
menu NOC text 1. View Current Configuration
menu NOC command 1. show running-config
menu NOC text 2. Backup Current Configuration
menu NOC command 2. copy running-config
https://NOC:CISCO@167.1.5.115:8080/CONFIGS/R6_CONFIG.txt
menu NOC text 3. Exit
menu NOC command 3. exit
!
line vty 0 4
login local
```

Task 4.1 Verification

Verify the menu:

```
Rack1R6#telnet 150.1.6.6
Trying 150.1.6.6 ... Open


User Access Verification

Username: NOC
Password: <CISCO>
Menu for Level 1 NOC users


    1.         View Current Configuration

    2.         Backup Current Configuration

    3.         Exit

<2>

Address or name of remote host [167.1.5.115]?
Destination filename [CONFIGS/R6_CONFIG.txt]?
%Error writing https://NOC:CISCO@167.1.5.115:8080/CONFIGS/R6_CONFIG.txt (I/O
error)
```

Task 4.1 Breakdown

In order to view the entire configuration, the user will need to be at privilege level 15. Make sure that you test your menu after configuring. In general, it is recommended to test from a telnet session, since forgetting the 'exit' option can prevent you from being able to exit.

Task 7.2

R2:

```
interface Loopback0
 ip nat inside
!
interface FastEthernet0/0
 ip address 172.16.0.2 255.255.255.0 secondary
 ip address 167.1.27.2 255.255.255.0
 ip nat outside
 ip policy route-map POLICY
!
ip nat pool INSIDE_GLOBAL 167.1.27.100 167.1.27.199 netmask 255.255.255.0
ip nat inside source list INSIDE_LOCAL pool INSIDE_GLOBAL
!
ip access-list standard INSIDE_LOCAL
 permit 172.16.0.0 0.0.0.255
!
route-map POLICY permit 10
 match ip address INSIDE_LOCAL
 set interface Loopback0
```

Task 7.2 Verification

```
Rack1R2#debug ip nat detailed
IP NAT detailed debugging is on
Rack1R2#debug ip policy
Policy routing debugging is on
```

Configure SW1 to simulate packets from the virtual host:

SW1:

```
ip local policy route-map LOCAL
!
ip access-list standard LOCAL
 permit 172.16.0.0 0.0.0.255
!
route-map LOCAL permit 10
 match ip address LOCAL
 set ip default next-hop 167.1.27.2
!
interface FastEthernet0/14
 ip address 172.16.0.8 255.255.255.0 secondary

Rack1SW1#ping 167.1.13.3 source 172.16.0.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 167.1.13.3, timeout is 2 seconds:
Packet sent with a source address of 172.16.0.8
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 80/83/84 ms
```

View R2's debugging output:

```
Rack1R2#
IP: s=172.16.0.8 (FastEthernet0/0), d=167.1.13.3, len 100, policy match
IP: route map POLICY, item 10, permit
IP: s=172.16.0.8 (FastEthernet0/0), d=167.1.13.3 (Loopback0), len 100, policy
routed
IP: FastEthernet0/0 to Loopback0 167.1.13.3
NAT: installing alias for address 167.1.27.100
NAT: i: icmp (172.16.0.8, 4) -> (167.1.13.3, 4) [20]
NAT: s=172.16.0.8->167.1.27.100, d=167.1.13.3 [20]
```

Note the return packets:

```
NAT*: o: icmp (167.1.13.3, 4) -> (167.1.27.100, 4) [21]
NAT*: s=167.1.13.3, d=167.1.27.100->172.16.0.8 [21]
IP: s=167.1.13.3 (FastEthernet0/0), d=172.16.0.8 (FastEthernet0/0), len 100,
policy rejected -- normal forwarding
```

```
Rack1R2#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
--- 167.1.27.100       172.16.0.8         ---                ---
```

```
Rack1R3#
ICMP: echo reply sent, src 167.1.13.3, dst 167.1.27.100
ICMP: echo reply sent, src 167.1.13.3, dst 167.1.27.100
ICMP: echo reply sent, src 167.1.13.3, dst 167.1.27.100
ICMP: echo reply sent, src 167.1.13.3, dst 167.1.27.100
ICMP: echo reply sent, src 167.1.13.3, dst 167.1.27.100
```

Task 7.3

R5:

```
ip icmp rate-limit unreachable 5000
```

Task 7.3 Verification

Ping the unreachable destination from R4:

```
Rack1R4#ping 167.1.8.8 repeat 10

Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 167.1.8.8, timeout is 2 seconds:
U...U...U.
Success rate is 0 percent (0/10)
```

Without rate-limit configuration you would get:

```
Rack1R4#ping 167.1.8.8 repeat 10

Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 167.1.8.8, timeout is 2 seconds:
UUUUUUUUUU
Success rate is 0 percent (0/10)
```

Task 7.3 Breakdown

The default value for the 'ICMP rate limit unreachable' command is 500, which means that an unreachable will only be sent once every 500 milliseconds.

Task 7.4

R1:
```
track 1 interface Serial0/0 line-protocol
!
interface FastEthernet0/0
 standby 1 ip 204.12.1.100
 standby 1 priority 101
 standby 1 track 1
```

R3:
```
interface FastEthernet0/0
 standby 1 ip 204.12.1.100
 standby 1 preempt
```

R6:
```
interface FastEthernet0/0
 standby 1 ip 204.12.1.100
 standby 1 preempt
 standby 1 track Serial0/0
```

Task 7.4 Breakdown

R6 will be the active router over R3 if their priorities are the same, since R6's IP address is numerically higher. If this weren't the case, R6 would require a higher HSRP priority than R3.

Task 7.4 Verification

```
Rack1R1#show standby
FastEthernet0/0 - Group 1
  State is Active
    2 state changes, last state change 01:01:32
  Virtual IP address is 204.12.1.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.742 secs
  Preemption disabled
  Active router is local
  Standby router is 204.12.1.6, priority 100 (expires in 9.439 sec)
  Priority 101 (configured 101)
    Track object 1 state Up decrement 10
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

```
Rack1R6#show standby
FastEthernet0/0 - Group 1
  State is Standby
    1 state change, last state change 00:05:34
  Virtual IP address is 204.12.1.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.720 secs
  Preemption enabled
  Active router is 204.12.1.1, priority 101 (expires in 8.028 sec)
  Standby router is local
  Priority 100 (default 100)
    Track interface Serial0/0/0 state Up decrement 10
  IP redundancy name is "hsrp-Fa0/0-1" (default)

Rack1R3#show standby
FastEthernet0/0 - Group 1
  State is Listen
    2 state changes, last state change 00:06:21
  Virtual IP address is 204.12.1.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
  Preemption enabled
  Active router is 204.12.1.1, priority 101 (expires in 9.732 sec)
  Standby router is 204.12.1.6, priority 100 (expires in 7.412 sec)
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Fa0/0-1" (default)
Rack1R3#

Rack1R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R1(config)#interface s0/0
Rack1R1(config-if)#shutdown
Rack1R1(config-if)#
%HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Active -> Speak
%LINK-5-CHANGED: Interface Serial0/0, changed state to administratively down
Rack1R1(config-if)#^Z
Rack1R1#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed state to down
%SYS-5-CONFIG_I: Configured from console by consoles
Rack1R1#show track
Track 1
  Interface Serial0/0 line-protocol
  Line protocol is Down (hw admin-down)
    2 changes, last change 00:00:08
  Tracked by:
    HSRP FastEthernet0/0 1
Rack1R1#
```

```
Rack1R6#show standby
FastEthernet0/0 - Group 1
  State is Active
    2 state changes, last state change 00:01:30
  Virtual IP address is 204.12.1.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.000 secs
  Preemption enabled
  Active router is local
  Standby router is 204.12.1.3, priority 100 (expires in 7.992 sec)
  Priority 100 (default 100)
    Track interface Serial0/0 state Up decrement 10
  IP redundancy name is "hsrp-Gi0/0-1" (default)
Rack1R6#

Rack1R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R6(config)#interface s0/0/0
Rack1R6(config-if)#shutdown
Rack1R6(config-if)#
%LINK-5-CHANGED: Interface Serial0/0, changed state to administratively down
%HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Active -> Speak
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed state to down

Rack1R3#show standby
FastEthernet0/0 - Group 1
  State is Active
    4 state changes, last state change 00:00:09
  Virtual IP address is 204.12.1.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.848 secs
  Preemption enabled
  Active router is local
  Standby router is unknown
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

Task 8.1

R4:

```
class-map VIP
 match access-group name VIP
!
policy-map LLQ
 class VIP
   priority percent 99
!
interface FastEthernet0/1
service-policy output LLQ
!
ip access-list extended VIP
 permit ip host 167.1.4.204 any
```

Task 8.1 Breakdown

In some earlier IOS versions, it was necessary to enter the command “max-reserved-bandwidth” when applying a policy with items that added up to more than 75%. In some IOS versions, you can specify a priority percentage of 100. In other versions, you may receive the error “Sum total of class bandwidths exceeds 99 percent.” Here, we have specified a percentage of 99.

Task 8.1 Verification

Verify the LLQ configuration:

```
Rack1R4#show policy-map int fa0/1
 FastEthernet0/1

  Service-policy output: LLQ

    queue stats for all priority classes:

      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

    Class-map: VIP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name VIP
      Priority: 99% (99000 kbps), burst bytes 2475000, b/w exceed drops: 0


    Class-map: class-default (match-any)
      40 packets, 4628 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any

      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 25/2714
```

```
Rack1R4#
```

Task 8.2

R6:

```
interface Serial0/0/0
 custom-queue-list 1
!
access-list 182 permit tcp host 167.1.4.119 eq www any
!
queue-list 1 protocol ip 1 list 182
queue-list 1 default 2
```

Task 8.2 Verification

Verify that there are only two queues in the custom queue configuration:

```
Rack1R6#show queueing custom
Current custom queue configuration:

List   Queue  Args
1      2      default
1      1      protocol ip          list 182
```