1. Bridging and Switching

Task 1.1

SW1:
```
interface range FastEthernet0/13 - 15
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
 no shutdown
!
interface FastEthernet0/16
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
 no shutdown
```

SW2:
```
interface range FastEthernet0/13 - 15
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
 no shutdown
!
interface FastEthernet0/16
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
 no shutdown
```

SW3:
```
interface FastEthernet0/13
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
 no shutdown
!
interface FastEthernet0/16
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
 no shutdown
```

Task 1.1 Verification

```
Rack1SW1#show interfaces trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/13      desirable    802.1q         trunking      1
Fa0/14      desirable    802.1q         trunking      1
Fa0/15      desirable    802.1q         trunking      1
Fa0/16      desirable    802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/13      1-4094
Fa0/14      1-4094
Fa0/15      1-4094
Fa0/16      1-4094

Port        Vlans allowed and active in management domain
```

Copyright © 2009 Internetwork Expert www.INE.com

```
Fa0/13      1,3-5,7,17,23,28,38,56
Fa0/14      1,3-5,7,17,23,28,38,56
Fa0/15      1,3-5,7,17,23,28,38,56
Fa0/16      1,3-5,7,17,23,28,38,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/13      none
Fa0/14      none
Fa0/15      none
Fa0/16      1,3-5,7,17,23,28,38,56

Rack1SW2#show interfaces trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/13      desirable    802.1q         trunking      1
Fa0/14      desirable    802.1q         trunking      1
Fa0/15      desirable    802.1q         trunking      1
Fa0/16      desirable    802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/13      1-4094
Fa0/14      1-4094
Fa0/15      1-4094
Fa0/16      1-4094

Port        Vlans allowed and active in management domain
Fa0/13      1,3-5,7,17,23,28,38,56
Fa0/14      1,3-5,7,17,23,28,38,56
Fa0/15      1,3-5,7,17,23,28,38,56
Fa0/16      1,3-5,7,17,23,28,38,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/13      1,3-5,7,17,23,28,38,56
Fa0/14      1,3-5,7,17,23,28,38,56
Fa0/15      1,3-5,7,17,23,28,38,56
Fa0/16      1,3-5,7,17,23,28,38,56

Rack1SW3#show interfaces trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/13      desirable    802.1q         trunking      1
Fa0/16      desirable    802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/13      1-4094
Fa0/16      1-4094

Port        Vlans allowed and active in management domain
Fa0/13      1,3-5,7,17,23,28,38,56
Fa0/16      1,3-5,7,17,23,28,38,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/13      1,3-5,7,17,23,28,38,56
Fa0/16      1,3-5,7,17,23,28,38,56
Rack1SW3#
```

Copyright © 2009 Internetwork Expert

Task 1.2

SW1:
```
spanning-tree vlan 4,28,38,56 priority 4096
spanning-tree vlan 1,3,5,7,17,23 priority 61440
```

Task 1.2 Breakdown

Spanning-tree root bridge election is determined by the lowest bridge-ID. Bridge-ID is made up of two portions, the bridge priority and a MAC address. The bridge priority defaults to 32768, half of the maximum value 65535. Since each bridge-ID must be unique, and since each VLAN (by default) runs its own instance of spanning-tree, there must be some way to distinguish bridge-IDs between difference spanning-tree instances.

In older platforms, this was accomplished by assigning a single MAC address per VLAN. This solution results in a waste of MAC addresses, since each VLAN requires its own simply for identification. New Cisco switch platforms use the system-id extension to deal with this problem. The bridge-ID for a specific spanning-tree VLAN instance will be the configured priority plus the system-id extension. The system-id extension is effectively the VLAN number. Therefore, in order to ensure that SW1 is the root for VLANs 4, 28, 38, and 56 (even VLANs), and that SW2 is the root for VLANs 3, 5, 7, 17, and 23 (odd VLANs), the priority must be adjusted accordingly on SW1. Since a lower priority value is better, SW1 has been set with the lowest priority value, zero, for even VLANs.

For odd VLANs, SW1's priority has been set to the configurable maximum value of 61440. These values are arbitrary as long as SW1 priority for the even VLANs is less than SW2's default priority (32768) plus the system-id extension (VLAN number). Furthermore, SW1 can use any arbitrary number to force SW2 to be the root for the odd VLANs, as long as it is greater than SW2's priority plus the system-id extension.

> **Note**
>
> SW3's spanning-tree priority is set to 61440 in the initial configuration. This should have been noticed before starting the lab.

Copyright © 2009 Internetwork Expert

Task 1.2 Verification

```
Rack1SW1#show spanning-tree vlan 1 | include ID|Address
  Root ID     Priority    32769
              Address     0016.9d31.8380
  Bridge ID   Priority    61441  (priority 61440 sys-id-ext 1)
              Address     0019.55e6.6580

Rack1SW1#show spanning-tree vlan 3 | include ID|Address
  Root ID     Priority    32771
              Address     0016.9d31.8380
  Bridge ID   Priority    61443  (priority 61440 sys-id-ext 3)
              Address     0019.55e6.6580

Rack1SW1#show spanning-tree vlan 4 | include ID|Address
  Root ID     Priority    24580
              Address     0019.55e6.6580
  Bridge ID   Priority    24580  (priority 24576 sys-id-ext 4)
              Address     0019.55e6.6580

Rack1SW1#show spanning-tree vlan 28 | include ID|Address
  Root ID     Priority    24604
              Address     0019.55e6.6580
  Bridge ID   Priority    24604  (priority 24576 sys-id-ext 28)
              Address     0019.55e6.6580

Rack1SW2#show spanning-tree vlan 1 | include ID|Address
  Root ID     Priority    32769
              Address     0016.9d31.8380
  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     0016.9d31.8380

Rack1SW2#show spanning-tree vlan 3 | include ID|Address
  Root ID     Priority    32771
              Address     0016.9d31.8380
  Bridge ID   Priority    32771  (priority 32768 sys-id-ext 3)
              Address     0016.9d31.8380

Rack1SW2#show spanning-tree vlan 4 | include ID|Address
  Root ID     Priority    24580
              Address     0019.55e6.6580
  Bridge ID   Priority    32772  (priority 32768 sys-id-ext 4)
              Address     0016.9d31.8380

Rack1SW2#show spanning-tree vlan 28 | include ID|Address
  Root ID     Priority    24604
              Address     0019.55e6.6580
  Bridge ID   Priority    32796  (priority 32768 sys-id-ext 28)
              Address     0016.9d31.8380


Rack1SW1#show span vlan 1-4094 | i root|VLAN
VLAN0001
VLAN0003
VLAN0004
              This bridge is the root
```

Copyright © 2009 Internetwork Expert

```
VLAN0005
VLAN0007
VLAN0017
VLAN0023
VLAN0028
             This bridge is the root
VLAN0038
             This bridge is the root
VLAN0056
             This bridge is the root
Rack1SW1#
Rack1SW2#show span vlan 1-4094 | i VLAN|root
VLAN0001
             This bridge is the root
VLAN0003
             This bridge is the root
VLAN0004
VLAN0005
             This bridge is the root
VLAN0007
             This bridge is the root
VLAN0017
             This bridge is the root
VLAN0023
             This bridge is the root
VLAN0028
VLAN0038
VLAN0056
Rack1SW2#
Rack1SW3#show span vlan 1-4094 | i VLAN|root
VLAN0001
VLAN0003
VLAN0004
VLAN0005
VLAN0007
VLAN0017
VLAN0023
VLAN0028
VLAN0038
VLAN0056
Rack1SW3#
```

Copyright © 2009 Internetwork Expert

Task 1.3

SW1:

```
interface FastEthernet0/14
 spanning-tree vlan 4,28,38,56 port-priority 16
!
interface FastEthernet0/15
 spanning-tree vlan 4,28,38,56 port-priority 32
```

Previous Reference

Spanning-tree port-priority: Lab 3

Task 1.3 Verification

Verify the spanning-tree root ports for even numbered VLANs on SW2:

```
Rack1SW2#show spanning-tree vlan 4,28,38,56 | include VLAN|Interface|Fa
VLAN0004
             Port        16 (FastEthernet0/14)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/4            Desg FWD 100       128.6    Shr
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Root FWD 19        128.16   P2p
Fa0/15           Altn BLK 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0028
             Port        16 (FastEthernet0/14)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Root FWD 19        128.16   P2p
Fa0/15           Altn BLK 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
Fa0/24           Desg FWD 100       128.26   Shr
VLAN0038
             Port        16 (FastEthernet0/14)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Root FWD 19        128.16   P2p
Fa0/15           Altn BLK 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0056
             Port        16 (FastEthernet0/14)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/6            Desg FWD 19        128.8    P2p
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Root FWD 19        128.16   P2p
Fa0/15           Altn BLK 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
```

Copyright © 2009 Internetwork Expert

Shutdown Fa0/14 on SW1 and view the spanning-tree information:

```
Rack1SW2#show spanning-tree vlan 4,28,38,56 | include VLAN|Interface|Fa
VLAN0004
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/4            Desg FWD 100       128.6    Shr
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/15           Root FWD 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0028
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/15           Root FWD 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
Fa0/24           Desg FWD 100       128.26   Shr
VLAN0038
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/15           Root FWD 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0056
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/6            Desg FWD 19        128.8    P2p
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/15           Root FWD 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
Rack1SW2#
```

Task 1.4

SW1:

```
interface FastEthernet0/15
 spanning-tree vlan 3,5,7,17,23 cost 1
```

Task 1.4 Breakdown

By default, all three of these interfaces will have a tie in port cost at 19 (FastEthernet). By adjusting the cost of interface Fa0/15 to less than 19, it will be preferred for these VLANs. Once Fa0/15 is down, the choice will be between port Fa0/13 and Fa0/14, both with a cost of 19. Since cost is a tie, and since the priority has not been adjusted on SW2, the tie breaker will be the lowest port ID. As 13 is lower than 14, port Fa0/13 will be chosen without any further configuration.

Previous Reference

Spanning-tree port cost: Lab 4

Copyright © 2009 Internetwork Expert

Task 1.4 Verification

Verify the spanning-tree root ports for odd numbered VLANs:

```
Rack1SW1#show spanning-tree vlan 3,5,7,17,23 | inc VLAN|Interface|Fa
VLAN0003
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/3            Desg FWD 100       128.5    Shr
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/15           Root FWD 1         128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0005
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/5            Desg FWD 100       128.7    Shr
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/15           Root FWD 1         128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0007
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/15           Root FWD 1         128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0017
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/15           Root FWD 1         128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0023
             Port        17 (FastEthernet0/15)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/15           Root FWD 1         128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
```

Shutdown Fa0/15 on SW2 and view the spanning-tree information:

```
Rack1SW1#show spanning-tree vlan 3,5,7,17,23 | inc VLAN|Interface|Fa
VLAN0003
             Port        15 (FastEthernet0/13)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/3            Desg FWD 100       128.5    Shr
Fa0/13           Root FWD 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/16           Desg FWD 19        128.18   P2p
```

Copyright © 2009 Internetwork Expert

```
VLAN0005
             Port        15 (FastEthernet0/13)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/5            Desg FWD 100       128.7    Shr
Fa0/13           Root FWD 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0007
             Port        15 (FastEthernet0/13)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Root FWD 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0017
             Port        15 (FastEthernet0/13)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/13           Root FWD 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/16           Desg FWD 19        128.18   P2p
VLAN0023
             Port        15 (FastEthernet0/13)
Interface        Role Sts Cost      Prio.Nbr Type
Fa0/13           Root FWD 19        128.15   P2p
Fa0/14           Altn BLK 19        128.16   P2p
Fa0/16           Desg FWD 19        128.18   P2p
```

Task 1.5

SW2:

```
interface FastEthernet0/24
 snmp trap mac-notification added
!
snmp-server enable traps MAC-Notification
snmp-server host 187.1.3.100 CISCOTRAP  MAC-Notification
mac-address-table notification
```

Task 1.5 Breakdown

To enable SNMP trapping when a MAC address is added or removed from the CAM table, issue the global configuration commands **mac-address-table notification** and **snmp-server enable traps MAC-Notification**. Then, these traps are selectively enabled on a per-interface basis by issuing the **snmp trap mac-notifications** interface level command. These traps are then forwarded to the NMS station located at 187.1.3.100, using the community string CISCOTRAP.

Task 1.5 Verification

Verify SNMP MAC Address logging configuration:

```
Rack1SW2#clear mac-address-table dynamic interface fa0/24
```

Copyright © 2009 Internetwork Expert

```
Rack1SW2#show mac-address-table notification
MAC Notification Feature is Enabled on the switch
Interval between Notification Traps : 1 secs
Number of MAC Addresses Added : 1
Number of MAC Addresses Removed : 0
Number of Notifications sent to NMS : 1
Maximum Number of entries configured in History Table : 1
Current History Table Length : 1
MAC Notification Traps are Enabled
History Table contents
----------------------
History Index 0, Entry Timestamp 348747, Despatch Timestamp 348747
MAC Changed Message :
Operation: Added   Vlan: 28  MAC Addr: 0060.7015.ac7a Dot1dBasePort: 24
```

Task 1.6

SW1, SW2 and SW3:

```
ip access-list extended IPONLY
 permit ip any any
!
mac access-list extended IP_ARP
 permit any any 0x806 0x0
!
mac access-list extended PVST_PLUS
 permit any any 0x010B 0x0
!
mac access-list extended PVST
 permit any any lsap 0x4242 0x0
 permit any any lsap 0xaaaa 0x0
!
vlan access-map IPONLY 10
 action forward
 match ip address IPONLY
!
vlan access-map IPONLY 20
 action forward
 match mac address IP_ARP
!
vlan access-map IPONLY 30
 action forward
 match mac address PVST_PLUS
!
vlan access-map IPONLY 40
 action forward
 match mac address PVST
!
vlan access-map IPONLY 100
 action drop
!
vlan filter IPONLY vlan-list 56
```

Copyright © 2009 Internetwork Expert

Task 1.6 Breakdown

The above task describes a seemingly straightforward scenario in which only IP traffic is allowed to transit VLAN 56. This is accomplished by creating a VLAN access-list (VACL) which permits IP traffic, and denies all other. However, when this access-map is applied, other behind the scenes protocols stop working. These protocols include IP ARP and STP (PVST+ in our case). PVST+ BPDUs are transported in Ethernet frames using 802.3 LLC SNAP encapsulation over 802.1q trunks, having PID (Protocol ID) of 0x010B. Additionally, some PVST+ BPDUs are encapsulated into Ethernet 802.3 LLC frames, having SSAP/DSAP 0x42 to interoperate with classic IEEE STP.

In addition to permitting IP, these above protocols must be permitted. Although IP uses the ethertype 0x800, IP ARP uses its own ethertype value of 0x806. This value must also be permitted, otherwise ARP cannot work. Note that even though PVST+ uses LLC SNAP encapsulation, you can match the PID value using the “ethertype” keyword in MAC access-lists.

Previous Reference

VLAN Access-Lists: Lab 5

Task 1.6 Verification

To verify the filtering, you can simulate a simple IPX network between R5 and R6, assuming the IOS versions support it.

R5:

```
ipx routing
!
interface Fa 0/1
 ipx encapsulation sap
 ipx network 56
```

R6:

```
ipx routing
!
interface FastEthernet0/0
 ipx encapsulation sap
 ipx network 56
```

With the VLAN filter applied, try to IPX ping R6 from R5:

```
Rack1R6#show ipx interface FastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  IPX address is 56.0015.62d0.4830, SNAP [up]
  Delay of this IPX network, in ticks is 1
  IPXWAN processing not enabled on this interface.
```

Copyright © 2009 Internetwork Expert www.INE.com

```
  IPX SAP update interval is 60 seconds
  IPX type 20 propagation packet forwarding is disabled
<output omitted>

Rack1R5#ping 56.0015.62d0.4830

Translating "56.0015.62d0.4830"

Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to 56.0015.62d0.4830, timeout is
2 seconds:
.....
Success rate is 0 percent (0/5)
```

Ensure that IP/ARP works fine:

```
Rack1R5#ping 187.1.56.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 187.1.56.6, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
```

Verify the spanning-tree status. You should see a root port on SW2:

```
Rack1SW2#show spanning-tree vlan 56

VLAN0056
  Spanning tree enabled protocol rstp
  Root ID    Priority    24632
             Address     000f.8fe0.3500
             Cost        19
             Port        13 (FastEthernet0/13)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32824  (priority 32768 sys-id-ext 56)
             Address     000f.8fb2.e800
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
------
Fa0/6            Desg FWD 19        128.8    P2p
Fa0/13           Altn BLK 19        128.15   P2p
Fa0/14           Root FWD 19        128.16   P2p
Fa0/15           Altn BLK 19        128.17   P2p
Fa0/16           Desg FWD 19        128.18   P2p
```

Remove VLAN filter:

```
Rack1SW1(config)#no vlan filter IPONLY vlan-list 56
Rack1SW2(config)#no vlan filter IPONLY vlan-list 56
Rack1SW3(config)#no vlan filter IPONLY vlan-list 56

Rack1R5#ping 56.0015.62d0.4830
```

Copyright © 2009 Internetwork Expert

```
Translating "56.0015.62d0.4830"

Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to 56.0015.62d0.4830, timeout is
2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Make sure that you turn the vlan filter back on, or you will lose the points for the section.

Task 1.7

```
SW1:
mls qos
!
interface FastEthernet0/7
 switchport access vlan 17
 switchport voice vlan 7
 mls qos trust cos
!
interface FastEthernet0/8
 switchport access vlan 17
 switchport voice vlan 7
 mls qos trust cos
!
define interface-range VPORTS FastEthernet 0/7 - 8
```

Task 1.7 Breakdown

The first step in configuring the 3560 to communicate with Cisco IP phones is to define how VoIP traffic will be carried. This task states that data traffic will be encapsulated in VLAN 7, and VoIP traffic will be encapsulated in VLAN 17. As the default port state of the 3560 is dynamic, a dot1q trunk will automatically be negotiated with the Cisco IP phone. The only configuration required to communicate with the phone is to apply both the access and voice VLAN to the port. Ensure that these VLANs are defined in the VLAN database.

Quality of Service processing is disabled on the 3560 by default. To enable QoS processing, issue the **mls qos** global configuration command. Next, the command **mls qos trust cos** has been issued on the interfaces connected to the IP phones. This instructs the switch to maintain the CoS value that is received on the interface.

Lastly, an interface range macro has been defined named VPORTS. This macro can be used in the future to reference ports Fa0/7 and Fa0/8 together. These macros can be used to reduce the administrative overhead of keeping track of which interfaces contain the same configuration. For example, if a certain range of interfaces are configured in an EtherChannel bundle, a macro could be

Copyright © 2009 Internetwork Expert

created to manage all the member interfaces. This way, the member interfaces could be referenced by the macro, and it would be ensured that all member interfaces receive the same configuration.

Task 1.7 Verification

Verify MLS QoS configuration:

```
Rack1SW1#show mls qos interface fa0/7
FastEthernet0/7
trust state: trust cos
trust mode: trust cos
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none

Rack1SW1#show mls qos interface fa0/8
FastEthernet0/8
trust state: trust cos
trust mode: trust cos
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
```

Verify Voice VLAN and appliance trust:

```
Rack1SW1#show interfaces fa0/7 switchport | inc Voice|Appl
Voice VLAN: 7 (VLAN0007)
Appliance trust: none

Rack1SW1#show interfaces fa0/8 switchport | inc Voice|Appl
Voice VLAN: 7 (VLAN0007)
Appliance trust: none
```

Copyright © 2009 Internetwork Expert

Task 1.8

SW1-SW3:
```
spanning-tree backbonefast
```

SW1:
```
spanning-tree vlan 4,28,38,56 forward-time 10
```

SW2:
```
spanning-tree vlan 1,3,5,7,17,23 forward-time 10
```

Task 1.8 Verification

```
Rack1SW1#show spanning-tree vlan 4 | include Forward
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 10 sec
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 10 sec
...

Rack1SW1#show spanning-tree vlan 1 | include Forward
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 10 sec
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
...

Rack1SW1#show spanning-tree backbonefast
BackboneFast is enabled

BackboneFast statistics
-----------------------
Number of transition via backboneFast (all VLANs)           : 0
Number of inferior BPDUs received (all VLANs)               : 0
Number of RLQ request PDUs received (all VLANs)             : 0
Number of RLQ response PDUs received (all VLANs)            : 0
Number of RLQ request PDUs sent (all VLANs)                 : 0
Number of RLQ response PDUs sent (all VLANs)                : 0
...
```

Task 1.9

R4:
```
username Rack1R5 password 0 C1SC0?2000
!
interface Serial0/1
 encapsulation ppp
 ppp authentication chap
```

R5:
```
interface Serial0/1
 encapsulation ppp
 clockrate 64000
 ppp chap password 0 C1SC0?2000
```

Task 1.9 Breakdown

Copyright © 2009 Internetwork Expert

Note that the escape sequence CTRL-V or ESC-Q must be used in order to enter a question mark in the password field. This username/password pair must also be configured in R4's local username database in order to authenticate R5.

The **username** and **ppp chap** commands with the "0" option after the password is telling the router that the password to come is in plain text format (i.e. unencrypted). This is also the default option when entering a password so the commands below will achieve the same result:

```
username Rack1R5 password 0 C1SC0?2000
username Rack1R5 password C1SC0?2000
```

If the commands are used with the "7" option after the password, the router will be expecting the password to come to be in encrypted form. Commonly this is used when a configuration is being copied from one router that has the **service password-encryption** command applied to another router. Below is the output of the command with the password in encrypted form:

```
username Rack1R5 password 7 123A5424312453567A7B74
```

Copyright © 2009 Internetwork Expert

Task 1.9 Verification

Verify PPP authentication:

```
Rack1R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R5(config)#interface Serial 0/1
Rack1R5(config-if)#do debug ppp authentication
PPP authentication debugging is on
Rack1R5(config-if)#shutdown
Rack1R5(config-if)#
%LINK-5-CHANGED: Interface Serial0/1, changed state to administratively
down
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/1, changed
state to down
Rack1R5(config-if)#no shutdown
%LINK-3-UPDOWN: Interface Serial0/1, changed state to up
Se0/1 PPP: Using default call direction
Se0/1 PPP: Treating connection as a dedicated line
Se0/1 PPP: Session handle[1A000004] Session id[3]
Se0/1 PPP: Authorization required
Se0/1 PPP: No authorization without authentication
Se0/1 CHAP: I CHALLENGE id 2 len 28 from "Rack1R4"
Se0/1 CHAP: Using hostname from unknown source
Se0/1 CHAP: Using password from interface CHAP
Se0/1 CHAP: O RESPONSE id 2 len 28 from "Rack1R5"
Se0/1 CHAP: I SUCCESS id 2 len 4
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/1, changed
state to up
```

2. Interior Gateway Protocol

Task 2.1

SW2:

```
ip routing
!
key chain RIP
 key 1
  key-string CISCO
!
interface Vlan28
 ip rip authentication mode md5
 ip rip authentication key-chain RIP
!
router rip
 version 2
 network 192.10.1.0
 no auto-summary
```

Copyright © 2009 Internetwork Expert

Task 2.1 Verification

Verify that SW2 receives authenticated RIP updates:

```
Rack1SW2#debug ip rip
RIP protocol debugging is on
 RIP: received packet with MD5 authentication
 RIP: received v2 update from 192.10.1.254 on Vlan28
      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

SW2:

```
router rip
 redistribute connected route-map CONNECTED->RIP metric 1
!
route-map CONNECTED->RIP permit 10
 match interface Loopback0
```

Task 2.2 Verification

Verify that the Loopback0 interface is being advertised:

```
Rack1SW2#show ip rip database
150.1.0.0/16    auto-summary
150.1.8.0/24    redistributed
    [1] via 0.0.0.0,
187.1.0.0/16 is possibly down
187.1.38.0/24 is possibly down
192.10.1.0/24    auto-summary
192.10.1.0/24    directly connected, Vlan28
205.90.31.0/24    auto-summary
205.90.31.0/24
    [7] via 192.10.1.254, 00:00:06, Vlan28
220.20.3.0/24    auto-summary
220.20.3.0/24
    [7] via 192.10.1.254, 00:00:06, Vlan28
222.22.2.0/24    auto-summary
222.22.2.0/24
    [7] via 192.10.1.254, 00:00:06, Vlan28
```

Copyright © 2009 Internetwork Expert

Task 2.3

R1:
```
router ospf 1
 router-id 150.1.1.1
 network 187.1.17.1 0.0.0.0 area 0
```

R3:
```
router ospf 1
 router-id 150.1.3.3
 network 187.1.3.3 0.0.0.0 area 0
 network 187.1.38.3 0.0.0.0 area 38
```

R4:
```
router ospf 1
 router-id 150.1.4.4
 network 187.1.4.4 0.0.0.0 area 0
 network 187.1.45.4 0.0.0.0 area 45
```

R5:
```
router ospf 1
 router-id 150.1.5.5
 network 187.1.45.5 0.0.0.0 area 45
```

SW1:
```
ip routing
!
router ospf 1
 router-id 150.1.7.7
 network 187.1.7.7 0.0.0.0 area 7
 network 187.1.13.7 0.0.0.0 area 7
 network 187.1.17.7 0.0.0.0 area 0
```

SW2:
```
ip routing
!
router ospf 1
 router-id 150.1.8.8
 network 187.1.38.8 0.0.0.0 area 38
```

Copyright © 2009 Internetwork Expert

Task 2.4

R1:
```
interface Serial0/0.134 multipoint
 ip ospf network point-to-multipoint
!
router ospf 1
 area 134 range 187.1.134.0 255.255.255.0
 area 134 virtual-link 150.1.3.3
 network 187.1.134.1 0.0.0.0 area 134
```

R3:
```
interface Serial1/0
 ip ospf network point-to-multipoint
!
router ospf 1
 area 134 range 187.1.134.0 255.255.255.0
 area 134 virtual-link 150.1.1.1
 area 134 virtual-link 150.1.4.4
 network 187.1.134.3 0.0.0.0 area 134
```

R4:
```
interface Serial0/0.134 multipoint
 ip ospf network point-to-multipoint
!
router ospf 1
 area 134 range 187.1.134.0 255.255.255.0
 area 134 virtual-link 150.1.3.3
 network 187.1.134.4 0.0.0.0 area 134
```

Tasks 2.3 – 2.4 Verification

Verify the OSPF neighbors:

```
Rack1R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.3.3         0   FULL/  -        -           187.1.134.3       OSPF_VL0
150.1.7.7         1   FULL/BDR        00:00:38    187.1.17.7 FastEthernet0/0
150.1.3.3             0   FULL/  - 00:01:57    187.1.134.3  Serial0/0.134
```

```
Rack1R3#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.4.4         0   FULL/  -        -           187.1.134.4       OSPF_VL1
150.1.1.1         0   FULL/  -        -           187.1.134.1       OSPF_VL0
150.1.8.8         1   FULL/BDR        00:00:30    187.1.38.8        Ethernet0/1
150.1.4.4         0   FULL/  -        00:01:39    187.1.134.4       Serial1/0
150.1.1.1         0   FULL/  -        00:01:36    187.1.134.1       Serial1/0
```

Copyright © 2009 Internetwork Expert

```
Rack1R4#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.3.3         0   FULL/  -        -           187.1.134.3     OSPF_VL0
150.1.5.5         0   FULL/  -        00:00:34    187.1.45.5      Serial0/1
150.1.3.3         0   FULL/  -        00:01:57    187.1.134.3     Serial0/0.134
```

Verify the OSPF network type on Frame Relay segment between R1, R3, and R4:

```
Rack1R3#show ip ospf interface s1/0
Serial1/0 is up, line protocol is up
  Internet Address 187.1.134.3/24, Area 134
  Process ID 1, Router ID 150.1.3.3, Network Type POINT_TO_MULTIPOINT,
Cost: 781
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT,
  Timer intervals configured, Hello 30, Dead 120, Wait 120,Retransmit 5
<output omitted>
```

Task 2.5

R1:

```
router ospf 1
 network 150.1.1.1 0.0.0.0 area 0
```

R3:

```
router ospf 1
 network 150.1.3.3 0.0.0.0 area 0
```

R4:

```
router ospf 1
 network 150.1.4.4 0.0.0.0 area 0
```

R5:

```
router ospf 1
 redistribute connected subnets route-map CONNECTED->OSPF
!
route-map CONNECTED->OSPF
 set metric 20
 set metric-type type-2
 match interface Loopback0
```

SW1:

```
router ospf 1
 network 150.1.7.7 0.0.0.0 area 0
```

SW2:

```
router ospf 1
 network 150.1.8.8 0.0.0.0 area 38
```

Copyright © 2009 Internetwork Expert

Task 2.5 Verification

Verify the OSPF networks origination:

```
Rack1SW1#show ip route ospf
     187.1.0.0/24 is subnetted, 7 subnets
O IA    187.1.134.0 [110/1] via 187.1.17.1, 00:01:05, Vlan17
O IA    187.1.45.0 [110/910] via 187.1.17.1, 00:01:05, Vlan17
O IA    187.1.38.0 [110/75] via 187.1.17.1, 00:01:05, Vlan17
O       187.1.3.0 [110/75] via 187.1.17.1, 00:01:05, Vlan17
O       187.1.4.0 [110/856] via 187.1.17.1, 00:01:05, Vlan17
     150.1.0.0/16 is variably subnetted, 6 subnets, 2 masks
O E2    150.1.5.0/24 [110/20] via 187.1.17.1, 00:00:34, Vlan17
O IA    150.1.8.8/32 [110/76] via 187.1.17.1, 00:00:39, Vlan17
O       150.1.4.4/32 [110/847] via 187.1.17.1, 00:01:06, Vlan17
O       150.1.3.3/32 [110/66] via 187.1.17.1, 00:01:06, Vlan17
O       150.1.1.1/32 [110/2] via 187.1.17.1, 00:01:06, Vlan17
```

Task 2.6

R1:

```
interface FastEthernet0/0
ip ospf authentication null
!
router ospf 1
area 134 virtual-link 150.1.3.3 authentication authentication-key CISCO
```

R3:

```
router ospf 1
 area 134 virtual-link 150.1.1.1 authentication authentication-key
CISCO
 area 134 virtual-link 150.1.4.4 authentication message-digest
 area 134 virtual-link 150.1.4.4 message-digest-key 1 md5 CISCO
```

R4:

```
router ospf 1
 area 134 virtual-link 150.1.3.3 authentication message-digest
 area 134 virtual-link 150.1.3.3 message-digest-key 1 md5 CISCO
```

SW1:

```
interface Vlan17
 ip ospf authentication null
```

Copyright © 2009 Internetwork Expert www.INE.com

Task 2.6 Verification

Verify the OSPF virtual-link authentication:

```
Rack1R3#show ip ospf virtual-links
Virtual Link OSPF_VL1 to router 150.1.4.4 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 134, via interface Serial1/0, Cost of using 781
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
    Adjacency State FULL (Hello suppressed)
    Index 2/5, retransmission queue length 0,number of retransmission 1
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
  Message digest authentication enabled
    Youngest key id is 1
Virtual Link OSPF_VL0 to router 150.1.1.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 134, via interface Serial1/0, Cost of using 781
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
    Adjacency State FULL (Hello suppressed)
    Index 1/4, retransmission queue length 0,number of retransmission 1
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
  Simple password authentication enabled
```

Confirm that no authentication is enabled on area0 interfaces on R1 and SW1:

```
Rack1R1#show ip ospf interface fa0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 187.1.17.1/24, Area 0
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 150.1.1.1, Interface address 187.1.17.1
  Backup Designated router (ID) 150.1.7.7, Interface address 187.1.17.7
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  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)
```

Copyright © 2009 Internetwork Expert

```
Rack1SW1#show ip ospf interface vl17
Vlan17 is up, line protocol is up
  Internet Address 187.1.17.7/24, Area 0
  Process ID 1, Router ID 150.1.7.7, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.1.1, Interface address 187.1.17.1
  Backup Designated router (ID) 150.1.7.7, Interface address 187.1.17.7
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  Supports Link-local Signaling (LLS)
  Index 1/2, 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 150.1.1.1  (Designated Router)
  Suppress hello for 0 neighbor(s)
```

Task 2.7

R2:

```
interface Serial0/0.235 multipoint
 no ip split-horizon eigrp 10
!
router eigrp 10
 network 150.1.2.2 0.0.0.0
 network 187.1.235.2 0.0.0.0
 no auto-summary
 eigrp router-id 150.1.2.2
```

R3:

```
interface Serial1/1.235 multipoint
 no ip split-horizon eigrp 10
!
router eigrp 10
 network 187.1.235.3 0.0.0.0
 no auto-summary
 eigrp router-id 150.1.3.3
```

R5:

```
interface Serial0/0
 no ip split-horizon eigrp 10
!
router eigrp 10
 network 187.1.5.5 0.0.0.0
 network 187.1.56.5 0.0.0.0
 network 187.1.235.5 0.0.0.0
 no auto-summary
 eigrp router-id 150.1.5.5
```

Copyright © 2009 Internetwork Expert

R6:

```
router eigrp 10
 redistribute connected metric 10000 10 255 1 1500 route-map CONNECTED-
>EIGRP
 network 187.1.56.6 0.0.0.0
 no auto-summary
 eigrp router-id 150.1.6.6
!
route-map CONNECTED->EIGRP permit 10
 match interface Loopback0
```

Task 2.7 Verification

Verify the EIGRP neighbors:

```
Rack1R5#show ip eigrp neighbors
IP-EIGRP neighbors for process 10
H   Address          Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                       (sec)         (ms)
Cnt Num
2   187.1.235.2      Se0/0            138 00:03:34   48   288  0  4
1   187.1.56.6       Et0/1             12 00:03:44  135   810  0  7
0   187.1.235.3      Se0/0            130 00:04:05  824  4944  0  7
```

Verify the EIGRP routes:

```
Rack1R2#show ip route eigrp
     187.1.0.0/24 is subnetted, 3 subnets
D       187.1.56.0 [90/2195456] via 187.1.235.5, 00:09:44,
Serial0/0.235
D       187.1.5.0 [90/2195456] via 187.1.235.5, 00:09:44, Serial0/0.235
     150.1.0.0/24 is subnetted, 2 subnets
D EX    150.1.6.0 [170/2198016] via 187.1.235.5, 00:09:44,
Serial0/0.235
```

Task 2.8

R2:

```
router eigrp 10
 eigrp stub connected summary
```

Task 2.8 Verification

```
Rack1R5#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 10
H   Address          Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                       (sec)         (ms)
Cnt Num
2   187.1.235.2      Se0/0            169 00:00:14   32   200  0  5
   Version 12.2/1.2, Retrans: 1, Retries: 0, Prefixes: 2
   Stub Peer Advertising ( CONNECTED SUMMARY ) Routes
1   187.1.56.6       Et0/1             12 00:14:42   54   324  0  12
   Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 1
0   187.1.235.3      Se0/0            170 00:15:03  296  1776  0  14
   Version 12.3/1.2, Retrans: 0, Retries: 0, Prefixes: 5
```

Copyright © 2009 Internetwork Expert

Task 2.9

R3 and R5:

```
ip access-list standard EVEN
 permit 0.0.0.0 254.255.255.255
!
route-map EIGRP_TO_OSPF deny 5
 match tag 110
!
route-map EIGRP_TO_OSPF permit 10
 match ip address EVEN
 set metric-type type-1
 set tag 90
!
route-map EIGRP_TO_OSPF permit 20
 set metric 100
 set tag 90
!
route-map OSPF_TO_EIGRP deny 5
 match tag 90
!
route-map OSPF_TO_EIGRP permit 10
 set tag 110
```

R5:

```
router eigrp 10
 redistribute connected route-map CONNECTED_TO_EIGRP
 redistribute ospf 1 metric 1500 10 255 1 1500 route-map OSPF_TO_EIGRP
!
router ospf 1
 redistribute eigrp 10 subnets route-map EIGRP_TO_OSPF
 distance 171 0.0.0.0 255.255.255.255 R3_R6_LOOPBACKS
!
! R5 should see the below Loopbacks via EIGRP
!
ip access-list standard R3_R6_LOOPBACKS
 permit 150.1.6.0
 permit 150.1.3.0
!
route-map CONNECTED_TO_EIGRP permit 10
 match interface Loopback0
!
route-map CONNECTED_TO_EIGRP permit 20
 match interface Serial0/1
!
```

R3:

```
router eigrp 10
 redistribute ospf 1 metric 1500 10 255 1 1500 route-map OSPF_TO_EIGRP
!
router ospf 1
 redistribute eigrp 10 subnets route-map EIGRP_TO_OSPF
 distance 171 0.0.0.0 255.255.255.255 R6_LOOPBACK
!
ip access-list standard R6_LOOPBACK
 permit 150.1.6.0
```

Copyright © 2009 Internetwork Expert

SW2:

```
interface Vlan28
 ip summary-address rip 187.1.0.0 255.255.0.0
!
router ospf 1
 redistribute rip subnets route-map RIP_TO_OSPF
 redistribute connected subnets
!
router rip
 redistribute ospf 1 metric 1
!
access-list 1 permit 0.0.0.0 254.255.255.255
!
route-map RIP_TO_OSPF permit 10
 match ip address 1
 set metric-type type-1
!
route-map RIP_TO_OSPF permit 20
 set metric 100
 set metric-type type-2
```

Task 2.9 Breakdown

Task 3.2 states that the Loopback 0 interface of SW2 should be advertised into the RIP domain without using the network statement. This is accomplished by redistributing connected. However, an additional stipulation on this task is that no other interfaces should be advertised into RIP while this configuration is performed. Therefore, a route-map is configured on SW2 that matches only the Loopback 0 interface, and is used to filter networks that are redistributed into RIP as connected. This configuration presents a problem with reachability from R3 to BB2.

When the Loopback 0 network of SW2 is redistributed into RIP, all other networks are implicitly denied. As the VLAN 38 interface of SW2 is directly connected, this network will not be advertised into RIP. This presents the problem that R3 no longer has IP reachability to SW2, however other devices in the routing domain will have reachability due to the redistribution of OSPF into RIP on SW2. In order to maintain reachability while staying within the requirements, a manual summary has been configured to BB2.

By adding the **`ip summary-address rip 187.1.0.0 255.255.0.0`** on the VLAN 28 interface, the entire major network 187.1.0.0/16 will be advertised on to BB2, and will therefore resolve the issue of connectivity between R3 and BB2.

Copyright © 2009 Internetwork Expert

Task 2.9 Verification

Verify the external routes redistributed into OSPF:

```
Rack1R4#show ip route ospf
     187.1.0.0/16 is variably subnetted, 15 subnets, 3 masks
O       187.1.134.1/32 [110/845] via 187.1.134.3, 00:34:59,
Serial0/0.134
O       187.1.134.3/32 [110/64] via 187.1.134.3, 00:34:59,
Serial0/0.134
O E2    187.1.235.0/24 [110/100] via 187.1.134.3, 00:19:58,
Serial0/0.134
O E2    187.1.56.0/24 [110/100] via 187.1.134.3, 00:13:19,
Serial0/0.134
O IA    187.1.38.0/24 [110/74] via 187.1.134.3, 00:34:34, Serial0/0.134
O       187.1.17.0/24 [110/846] via 187.1.134.3, 00:34:34,
Serial0/0.134
O       187.1.3.0/24 [110/74] via 187.1.134.3, 00:34:34, Serial0/0.134
O IA    187.1.7.0/24 [110/847] via 187.1.134.3, 00:34:34, Serial0/0.134
O E1 222.22.2.0/24 [110/94] via 187.1.134.3, 00:34:24, Serial0/0.134
O E1 220.20.3.0/24 [110/94] via 187.1.134.3, 00:34:24, Serial0/0.134
O E2 192.10.1.0/24 [110/20] via 187.1.134.3, 00:34:24, Serial0/0.134
     150.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
O E1    150.1.6.0/24 [110/84] via 187.1.134.3, 00:13:19, Serial0/0.134
                     [110/84] via 187.1.45.5, 00:13:19, Serial0/1
O E2    150.1.5.0/24 [110/20] via 187.1.45.5, 00:20:33, Serial0/1
O       150.1.3.0/24 [110/65] via 187.1.134.3, 00:34:35, Serial0/0.134
O E1    150.1.2.0/24 [110/84] via 187.1.134.3, 00:13:23, Serial0/0.134
                     [110/84] via 187.1.45.5, 00:13:23, Serial0/1
O IA    150.1.8.8/32 [110/75] via 187.1.134.3, 00:34:35, Serial0/0.134
O       150.1.7.7/32 [110/847] via 187.1.134.3, 00:34:35, Serial0/0.134
O       150.1.1.1/32 [110/846] via 187.1.134.3, 00:34:35, Serial0/0.134
```

Verify the summary route generation on SW2:

```
Rack1SW2#debug ip rip
RIP protocol debugging is on
 RIP: sending v2 update to 224.0.0.9 via Vlan28 (192.10.1.8)
 RIP: build update entries
       150.1.1.1/32 via 0.0.0.0, metric 1, tag 0
       150.1.2.0/24 via 0.0.0.0, metric 1, tag 0
       150.1.3.3/32 via 0.0.0.0, metric 1, tag 0
       150.1.4.4/32 via 0.0.0.0, metric 1, tag 0
       150.1.5.0/24 via 0.0.0.0, metric 1, tag 0
       150.1.6.0/24 via 0.0.0.0, metric 1, tag 0
       150.1.7.7/32 via 0.0.0.0, metric 1, tag 0
       150.1.8.0/24 via 0.0.0.0, metric 1, tag 0
       187.1.0.0/16 via 0.0.0.0, metric 2, tag 0
```

Copyright © 2009 Internetwork Expert

Test full connectivity between all internal networks with the following TCL script:

```
foreach i {
187.1.134.1
150.1.1.1
187.1.17.1
187.1.235.2
150.1.2.2
187.1.134.3
187.1.235.3
150.1.3.3
187.1.38.3
187.1.134.4
187.1.45.4
150.1.4.4
187.1.4.4
187.1.235.5
187.1.56.5
187.1.45.5
150.1.5.5
187.1.5.5
187.1.56.6
150.1.6.6
150.1.7.7
187.1.17.7
187.1.7.7
187.1.13.7
187.1.13.9
187.1.38.8
150.1.8.8
192.10.1.8
} { ping $i }
```

Note that VLAN23 and the Frame Relay link between R6 and BB1 are excluded from any IGP and thus are not verified.

Task 2.10

R1:

```
ipv6 unicast-routing
!
interface Tunnel14
 ipv6 address 2001:187:1:14::1/64
 tunnel source 150.1.1.1
 tunnel destination 150.1.4.4
!
interface Tunnel16
 ipv6 address 2001:187:1:16::1/64
 tunnel source 150.1.1.1
 tunnel destination 150.1.6.6
!
interface FastEthernet0/0
 ipv6 address 2001:187:1:17::1/64
```

R4:

Copyright © 2009 Internetwork Expert

```
ipv6 unicast-routing
!
interface Tunnel14
 ipv6 address 2001:187:1:14::4/64
 tunnel source 150.1.4.4
 tunnel destination 150.1.1.1
!
interface Tunnel46
 ipv6 address 2001:187:1:46::4/64
 tunnel source 150.1.4.4
 tunnel destination 150.1.6.6
!
interface FastEthernet0/0
 ipv6 address 2001:187:1:4::4/64
```

R6:

```
ipv6 unicast-routing
!
interface Tunnel16
 ipv6 address 2001:187:1:16::6/64
 tunnel source 150.1.6.6
 tunnel destination 150.1.1.1
!
interface Tunnel46
 ipv6 address 2001:187:1:46::6/64
 tunnel source 150.1.6.6
 tunnel destination 150.1.4.4
!
interface FastEthernet 0/0
 ipv6 address 2001:187:1:56::6/64
```

SW1:

```
sdm prefer dual-ipv4-and-ipv6 routing
!
! Reboot SW1
!
ipv6 unicast-routing
!
interface Vlan17
 ipv6 address 2001:187:1:17::7/64
```

Copyright © 2009 Internetwork Expert

Task 2.10 Verification

Verify basic connectivity:

```
Rack1R1#ping 2001:187:1:14::4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:187:1:14::4, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 140/141/144
ms

Rack1R1#ping 2001:187:1:16::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:187:1:16::6, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 156/157/160
ms

Rack1R1#ping 2001:187:1:17::7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:187:1:17::7, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms

Rack1R4#ping 2001:187:1:46::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:187:1:46::6, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/88/89 ms
```

Copyright © 2009 Internetwork Expert

Task 2.11

R1:
```
interface Tunnel14
 ipv6 rip RIPng enable
!
interface Tunnel16
 ipv6 rip RIPng enable
!
interface FastEthernet0/0
 ipv6 rip RIPng enable
 ipv6 rip RIPng summary-address 2001:187:1::/57
```

R4:
```
interface Tunnel14
 ipv6 rip RIPng enable
!
interface Tunnel46
 ipv6 rip RIPng enable
!
interface FastEthernet0/0
 ipv6 rip RIPng enable
```

R6:
```
interface Tunnel16
 ipv6 rip RIPng enable
!
interface Tunnel46
 ipv6 rip RIPng enable
!
interface FastEthernet 0/0
 ipv6 rip RIPng enable
```

SW1:
```
interface Vlan 17
 ipv6 rip RIPng enable
```

Copyright © 2009 Internetwork Expert

Task 2.11 Verification

Verify the RIPng routes on R1 and R4:

```
Rack1R6#show ipv6 route rip
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS
summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF
ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   2001:187:1:4::/64 [120/2]
     via FE80::211:92FF:FE08:2D80, Tunnel46
R   2001:187:1:14::/64 [120/2]
     via FE80::20F:23FF:FED5:5220, Tunnel16
     via FE80::211:92FF:FE08:2D80, Tunnel46
R   2001:187:1:17::/64 [120/2]
     via FE80::20F:23FF:FED5:5220, Tunnel16

Rack1R1#show ipv6 route rip
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS
summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF
ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   2001:187:1:4::/64 [120/2]
     via FE80::211:92FF:FE08:2D80, Tunnel14
R   2001:187:1:46::/64 [120/2]
     via FE80::20F:23FF:FEF4:E640, Tunnel16
     via FE80::211:92FF:FE08:2D80, Tunnel14
R   2001:187:1:56::/64 [120/2]
     via FE80::20F:23FF:FEF4:E640, Tunnel16

Rack1R4#show ipv6 route rip
IPv6 Routing Table - 14 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS
summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF
ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   2001:187:1:16::/64 [120/2]
     via FE80::20F:23FF:FEF4:E640, Tunnel46
     via FE80::20F:23FF:FED5:5220, Tunnel14
R   2001:187:1:17::/64 [120/2]
     via FE80::20F:23FF:FED5:5220, Tunnel14
R   2001:187:1:56::/64 [120/2]
     via FE80::20F:23FF:FEF4:E640, Tunnel46

Rack1SW1#show ipv6 route rip
```

```
IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS
summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF
ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   2001:187:1::/57 [120/2]
     via FE80::20F:23FF:FED5:5220, Vlan17
```

Task 2.12

R1:

```
ipv6 router rip RIPng
  distribute-list prefix-list NONE in Tunnel16
!
ipv6 prefix-list NONE seq 5 deny ::/0 le 128
```

Task 2.12 Verification

```
Rack1SW1#traceroute 2001:187:1:56::6

Type escape sequence to abort.
Tracing the route to 2001:187:1:56::6

  1 2001:187:1:17::1 0 msec 0 msec 0 msec
  2 2001:187:1:14::4 117 msec 118 msec 109 msec
  3 2001:187:1:56::6 176 msec 176 msec 176 msec

Rack1R6#traceroute  2001:187:1:17::7

Type escape sequence to abort.
Tracing the route to 2001:187:1:17::7

  1 2001:187:1:16::1 120 msec 120 msec 124 msec
  2 2001:187:1:17::7 120 msec 125 msec 144 msec
```

Copyright © 2009 Internetwork Expert

3. Exterior Gateway Routing

Task 3.1

R3:
```
router bgp 200
 neighbor 187.1.235.2 remove-private-as
 neighbor 187.1.235.5 remove-private-as
```

R4:
```
router bgp 200
 neighbor 187.1.45.5 remove-private-as
```

SW1:
```
interface Loopback77
 ip address 187.1.77.7 255.255.255.0
!
router bgp 65017
 network 187.1.77.0 mask 255.255.255.0
```

SW2:
```
router bgp 200
 neighbor 192.10.1.254 remove-private-as
```

Task 3.1 Breakdown

The above task states that BGP devices outside AS 200 should see this prefix as originated in AS 200. By removing the private AS number when AS 200 passes updates upstream, the private AS configuration is transparent to the rest of the network.

Copyright © 2009 Internetwork Expert

Task 3.1 Verification

Verify Loopback77 prefix in the BGP table on R3:

```
Rack1R3#show ip bgp | include 77|Netw
   Network          Next Hop            Metric LocPrf Weight Path
*> 187.1.77.0/24    187.1.134.1                              0 65017 i
```

Verify the same prefix in AS100:

```
Rack1R5#show ip bgp 187.1.77.0
BGP routing table entry for 187.1.77.0/24, version 18
Paths: (3 available, best #3, table Default-IP-Routing-Table)
  Advertised to update-groups:
     1          2
  200
    187.1.235.3 from 187.1.235.3 (150.1.3.3)
      Origin IGP, localpref 100, valid, external
  200, (Received from a RR-client)
    187.1.235.3 from 187.1.235.2 (150.1.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    187.1.45.4 from 187.1.45.4 (150.1.4.4)
      Origin IGP, localpref 100, valid, external, best
```

Task 3.2

R2:

```
router bgp 100
 network 187.1.235.0 mask 255.255.255.0
 aggregate-address 187.1.0.0 255.255.0.0 summary-only
 neighbor 204.12.1.254 unsuppress-map UNSUPPRESS
!
ip prefix-list NETWORK_235 seq 5 permit 187.1.235.0/24
!
route-map UNSUPPRESS permit 10
 match ip address prefix-list NETWORK_235
```

R6:

```
router bgp 100
 aggregate-address 187.1.0.0 255.255.0.0 summary-only
```

Task 3.2 Breakdown

When BGP aggregation is configured, the aggregate-address (along with all subnets of the aggregate) are candidate to be advertised to the rest of the BGP domain. By adding the **summary-only** keyword, these subnets advertisements are suppressed. By configuring unsuppress map on R2, traffic from AS 54 will prefer to come in to R2. This is due to the fact that all routers throughout the network will always choose the longest match in the IP routing table. As R6 is only advertising the shorter match, this path will not be used unless the subnet information is lost from R2.

Copyright © 2009 Internetwork Expert

Task 3.2 Verification

Verify the prefixes advertised to AS54 by R6:

```
Rack1R6#show ip bgp neighbors 54.1.1.254 advertised-routes
BGP table version is 26, 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
*> 187.1.0.0        0.0.0.0                            32768 i
*>i205.90.31.0      187.1.235.3              0    100      0 200 254 ?
*>i220.20.3.0       187.1.235.3              0    100      0 200 254 ?
*>i222.22.2.0       187.1.235.3              0    100      0 200 254 ?
```

Verify the prefixes advertised to AS54 by R2:

```
Rack1R2#show ip bgp neighbors 204.12.1.254 advertised-routes
BGP table version is 41, local router ID is 150.1.2.2
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
*> 187.1.0.0        0.0.0.0                            32768 i
s> 187.1.235.0/24   0.0.0.0                  0         32768 i
*> 205.90.31.0      187.1.235.3                            0 200 254 ?
*> 220.20.3.0       187.1.235.3                            0 200 254 ?
*> 222.22.2.0       187.1.235.3                            0 200 254 ?
```

Task 3.3

SW2:

```
router bgp 200
 network 192.10.1.0
 network 205.90.31.0
 network 220.20.3.0
 network 222.22.2.0
!
router bgp 200
 distance bgp 121 200 200
```

Copyright © 2009 Internetwork Expert

Task 3.3 Breakdown

Debugging RIP will show you what routes are learned. Switch 2 is learning routes via both RIP and BGP. BGP will win for the routing table, since eBGP has an AD of 20, compared to the AD of 120 for RIP.

```
07:21:13: RIP: received v2 update from 192.10.1.254 on Vlan28
07:21:13:      205.90.31.0/24 via 0.0.0.0 in 7 hops
07:21:13:      220.20.3.0/24 via 0.0.0.0 in 7 hops
07:21:13:      222.22.2.0/24 via 0.0.0.0 in 7 hops

Since the task explicitly asks
```

Task 3.3 Verification

See if prefixes appear in BGP table:

```
Rack1SW2#show ip bgp
BGP table version is 47, local router ID is 150.1.8.8
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
<output omitted>
*> 192.10.1.0       0.0.0.0                  0         32768 i
*  205.90.31.0      192.10.1.254             0             0 254 ?
*>                  192.10.1.254             7         32768 i
*  220.20.3.0       192.10.1.254             0             0 254 ?
*>                  192.10.1.254             7         32768 i
*  222.22.2.0       192.10.1.254             0             0 254 ?
*>                  192.10.1.254             7         32768 i
```

Task 3.4

R3:

```
interface Loopback 33
 ip address 150.1.33.33 255.255.255.0
!
interface Loopback 133
 ip address 150.1.133.133 255.255.255.0
!
ip prefix-list LOOPBACK33 permit 150.1.33.0/24
ip prefix-list LOOPBACK133 permit 150.1.133.0/24
!
route-map SET_COMMUNITY permit 10
 match ip address prefix-list LOOPBACK33
 set community 100:542
!
route-map SET_COMMUNITY permit 20
 match ip address prefix-list LOOPBACK133
 set community 100:546
!
```

Copyright © 2009 Internetwork Expert

```
ip bgp-community new-format
!
router bgp 200
 network 150.1.33.0 mask 255.255.255.0 route-map SET_COMMUNITY
 network 150.1.133.0 mask 255.255.255.0 route-map SET_COMMUNITY
 neighbor 187.1.235.2 send-community
 neighbor 187.1.235.5 send-community
```

R5:

```
router bgp 100
 neighbor 187.1.56.6 send-community
 neighbor 187.1.235.2 send-community
```

R2:

```
ip bgp-community new-format
!
ip community-list standard 100:542 permit 100:542
!
route-map TO_BB3 permit 10
 match community 100:542
 set as-path prepend 100 100 100
!
route-map TO_BB3 permit 100

router bgp 100
 neighbor 187.1.235.5 send-community
 neighbor 204.12.1.254 route-map TO_BB3 out
```

R6:

```
ip bgp-community new-format
!
ip community-list standard 100:546 permit 100:546
!
route-map TO_BB1 permit 10
 match community 100:546
 set as-path prepend 100 100 100
!
route-map TO_BB1 permit 100
!
router bgp 100
 neighbor 187.1.56.5 send-community
 neighbor 54.1.1.254 route-map TO_BB1 out
```

Task 3.4 Verification

```
Rack1R6#show ip bgp 150.1.33.0
BGP routing table entry for 150.1.33.0/24, version 55
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200
    187.1.235.3 (metric 2172416) from 187.1.56.5 (150.1.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: 100:542

Rack1R6#show ip bgp 150.1.133.0
BGP routing table entry for 150.1.133.0/24, version 56
```

Copyright © 2009 Internetwork Expert

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200
    187.1.235.3 (metric 2172416) from 187.1.56.5 (150.1.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: 100:546
```

```
Rack1R2#show ip bgp 150.1.33.0
BGP routing table entry for 150.1.33.0/24, version 12
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     1         3
  200
    187.1.235.3 from 187.1.235.5 (150.1.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal
      Community: 100:542
  200
    187.1.235.3 from 187.1.235.3 (150.1.3.3)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 100:542
```

```
Rack1R2#show ip bgp 150.1.133.0
BGP routing table entry for 150.1.133.0/24, version 13
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     1         3
  200
    187.1.235.3 from 187.1.235.5 (150.1.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal
      Community: 100:546
  200
    187.1.235.3 from 187.1.235.3 (150.1.3.3)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 100:546
```

```
RS.42.1.BB1>sh ip bgp
BGP table version is 477, 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
...
*> 150.1.33.0/24    54.1.1.6                               0 100 200 i
*  150.1.133.0/24   54.1.1.6                               0 100 100
100 100 200 i
```

```
RS.42.1.BB3>show ip bgp
BGP table version is 579, 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
```

Copyright © 2009 Internetwork Expert

```
...*>i150.1.33.0/24    172.16.4.1               0    100      0 100 200
i
*                   204.12.1.2                             0 100 100
100 100 200 i
*> 150.1.133.0/24   204.12.1.2                             0 100 200 i
```

4. IP and IOS Features

Task 4.1

R6:

```
archive
 log config
  logging enable
  logging size 500
  notify syslog
!
logging 187.1.5.155
```

Task 4.1 Verification

Verify the change logging configuration:

```
Rack1R6#show archive log config all
 idx   sess           user@line      Logged command
    1     1        console@console  |  logging enable
    2     1        console@console  |  logging size 500
    3     1        console@console  |  notify syslog
    4     1        console@console  |  logging 187.1.5.155

Rack1R6#show logging
Syslog logging: enabled (11 messages dropped, 2 messages rate-limited,
                0 flushes, 0 overruns, xml disabled, filtering
disabled)
    Console logging: level debugging, 156 messages logged, xml
disabled,
                     filtering disabled
    Monitor logging: level debugging, 0 messages logged, xml disabled,
                     filtering disabled
    Buffer logging: disabled, xml disabled,
                    filtering disabled
    Logging Exception size (4096 bytes)
    Count and timestamp logging messages: disabled

No active filter modules.

    Trap logging: level informational, 97 message lines logged
        Logging to 187.1.38.100 (udp port 514, audit disabled, link
up), 8 message lines logged, xml disabled,
              filtering disabled
        Logging to 187.1.5.155 (udp port 514, audit disabled, link up),
4 message lines logged, xml disabled,
              filtering disabled
```

Copyright © 2009 Internetwork Expert

Task 4.2

R6:
```
service timestamps log datetime msec localtime show-timezone
!
clock timezone PST -8
clock summer-time PDT recurring
!
ntp server 150.1.1.1
```

R1:
```
ntp master 1
```

Task 4.2 Verification

Verify the logging timestamps:

```
Rack1R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R6(config)#exit
Mar 12 06:21:52.438 PST: %SYS-5-CONFIG_I: Configured from console by
console
```

Make sure to give NTP a couple minutes to synchronize before you check ntp status:

```
Rack1R6#show ntp status
Clock is synchronized, stratum 2, reference is 150.1.1.1
nominal freq is 249.5901 Hz, actual freq is 249.5901 Hz, precision is
2**18
reference time is C029973F.A5C189BE (21:59:27.647 PST Thu Feb 28 2002)
clock offset is -0.0261 msec, root delay is 102.89 msec
root dispersion is 0.24 msec
```

Task 4.3

R3:
```
ip wccp web-cache redirect-list 25
!
interface FastEthernet0/0
 ip wccp web-cache redirect in
!
access-list 25 deny 187.1.3.50
access-list 25 permit any
```

Task 4.3 Breakdown

By default, traffic from all hosts received or sent on an interface (depending on how redirection is configured) is candidate for redirection to a web cache engine. In the above scenario, all traffic except that which is sourced from 187.1.3.50 is eligible for caching.

Copyright © 2009 Internetwork Expert

Task 4.3 Verification

Verify the WCCP configuration:

```
Rack1R3#show ip wccp web-cache
Global WCCP information:
    Router information:
        Router Identifier:                   -not yet determined-
        Protocol Version:                    2.0

    Service Identifier: web-cache
        Number of Cache Engines:             0
        Number of routers:                   0
        Total Packets Redirected:            0
        Process:                             0
        Fast:                                0
        CEF:                                 0
        Redirect access-list:                25
        Total Packets Denied Redirect:       0
        Total Packets Unassigned:            0
        Group access-list:                   -none-
        Total Messages Denied to Group:      0
        Total Authentication failures:       0
        Total Bypassed Packets Received:     0

Rack1R3#show ip wccp interfaces
WCCP interface configuration:
    FastEthernet0/0
        Output services: 0
        Input services:  1
        Mcast services:  0
        Exclude In:      FALSE
```

Copyright © 2009 Internetwork Expert

Task 4.4

R5:

```
interface FastEthernet0/0
 ip helper-address 187.1.56.255
 ip directed-broadcast
!
interface FastEthernet0/1
 ip directed-broadcast
```

Task 4.4 Verification

Verify the broadcast forwarding configuration:

```
Rack1R5#show ip interface FastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  Internet address is 187.1.5.5/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is 187.1.56.255
  Directed broadcast forwarding is enabled

<output omitted>

Rack1R5#show ip interface FastEthernet0/1
FastEthernet0/1 is up, line protocol is up
  Internet address is 187.1.56.5/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is enabled

<output omitted>
```

Copyright © 2009 Internetwork Expert

5. IP Multicast

Task 5.1

R3:

```
interface Loopback0
 ip pim sparse-mode
 ip ospf network point-to-point
!
ip pim bsr-candidate Loopback0 0
```

R4:

```
ip pim rp-candidate Serial0/0.134 group-list R4_GROUP
!
ip access-list standard R4_GROUP
 permit 224.0.0.0 7.255.255.255
```

R5:

```
ip pim rp-candidate Serial0/0 group-list R5_GROUP
!
ip access-list standard R5_GROUP
 permit 232.0.0.0 7.255.255.255
!
router ospf 1
 distance 171 0.0.0.0 255.255.255.255 R3_LOOPBACK
```

Task 5.1 Verification

The AD for R3's loopback is adjusted so that R5's path via EIGRP is preferred.

Verify the RP mappings:

```
Rack1R1#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/5
  RP 187.1.134.4 (?), v2
    Info source: 150.1.3.3 (?), via bootstrap, priority 0
         Uptime: 00:43:45, expires: 00:03:20
Group(s) 232.0.0.0/5
  RP 187.1.235.5 (?), v2
    Info source: 150.1.3.3 (?), via bootstrap, priority 0
         Uptime: 00:00:32, expires: 00:03:20
```

Copyright © 2009 Internetwork Expert

Task 5.2

R1:
```
interface Tunnel14
 ip address 187.1.14.1 255.255.255.0
 ip pim sparse-mode
 tunnel source 150.1.1.1
 tunnel destination 150.1.4.4
!
ip mroute 0.0.0.0 0.0.0.0 Tunnel 14
```

R4:
```
interface Tunnel14
 ip address 187.1.14.4 255.255.255.0
 ip pim sparse-mode
 tunnel source 150.1.4.4
 tunnel destination 150.1.1.1
```

SW1:
```
interface Vlan7
 ip igmp join-group 228.34.28.100
```

Task 5.2 Verification

Try pinging multicast group from R4 before configuring the tunnel:

```
Rack1R4#ping 228.34.28.100 repeat 5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 228.34.28.100, timeout is 2 seconds:
Packet sent with a source address of 187.1.4.4
.....

Rack1R3#show ip mroute
IP Multicast Routing Table

<output omitted>

(*, 228.34.28.100), 00:01:20/stopped, RP 187.1.134.4, flags: SP
  Incoming interface: Serial1/0, RPF nbr 187.1.134.4
  Outgoing interface list: Null

(187.1.134.4, 228.34.28.100), 00:01:21/00:02:59, flags: PT
  Incoming interface: Serial1/0, RPF nbr 187.1.134.4
  Outgoing interface list: Null
```

Copyright © 2009 Internetwork Expert

```
Rack1R1#show ip mroute

<output omitted>

(*, 228.34.28.100), 00:03:09/00:03:18, RP 187.1.134.4, flags: S
  Incoming interface: Serial0/0.134, RPF nbr 187.1.134.4
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse, 00:03:09/00:03:18

Now establish the tunnel and try to ping again:

Rack1R4#ping 228.34.28.100 repeat 5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 228.34.28.100, timeout is 2 seconds:

Reply to request 0 from 187.1.17.7, 132 ms
Reply to request 1 from 187.1.17.7, 112 ms
Reply to request 2 from 187.1.17.7, 108 ms
Reply to request 3 from 187.1.17.7, 112 ms
Reply to request 4 from 187.1.17.7, 108 ms

Rack1R1#show ip mroute
IP Multicast Routing Table

<output omitted>

(187.1.14.4, 228.34.28.100), 00:00:36/00:03:02, flags: FT
  Incoming interface: Tunnel14, RPF nbr 187.1.14.4
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse, 00:00:38/00:02:52
```

Task 5.3

R4 & R5:

```
ip access-list extended R3_R4_GROUPS
 permit ip host 150.1.3.3 any
 permit ip host 150.1.4.4 any
!
ip pim accept-register list R3_R4_GROUPS
```

R3 & R4:

```
ip pim register-source Loopback0
```

R1, R3, R4, R5, and SW1:

```
ip access-list standard R4_GROUP
 permit 224.0.0.0 7.255.255.255
!
ip access-list standard R5_GROUP
 permit 232.0.0.0 7.255.255.255
!
ip pim accept-rp 150.1.4.4 R4_GROUP
ip pim accept-rp 150.1.5.5 R4_GROUP
```

Task 5.3 Verification

```
Rack1R5#ping 228.34.28.100 repeat 100

Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 228.34.28.100, timeout is 2
seconds:
.....

Rack1R4#debug ip pim
PIM debugging is on

%PIM-4-INVALID_SRC_REG: Received Register from 187.1.235.5 for
(187.1.5.5, 228.34.28.100), not willing to be RP
```

Copyright © 2009 Internetwork Expert

6. QoS

Task 6.1

```
R3:
interface Serial1/0
 frame-relay traffic-shaping
 frame-relay class FRTS
!
map-class frame-relay FRTS
 frame-relay cir 192000
 frame-relay bc 19200
 frame-relay be 12800
```

Task 6.1 Breakdown

This task states that R3 should average 192Kbps on both VC 301 and 304, and that traffic bursts of up to 320Kbps should be allowed for a maximum period of 100ms. The following values can therefore be inferred from this description:

CIR = 192000bps
AR = 320000bps
Tc = 100ms

Using the formula Bc = CIR * Tc/1000:

Bc = 192000 * 100/1000
Bc = 192000 * 1/10
Bc = 19200

Using the formula Be = (AR - CIR) * Tc/1000

Be = (320000 - 192000) * 100/1000
Be = 128000 * 1/10
Be = 12800

Previous Reference
Frame Relay Traffic Shaping: Lab 1

Copyright © 2009 Internetwork Expert

Task 6.1 Verification

```
Rack1R3#show frame-relay pvc 304

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

DLCI = 304, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial1/0

  input pkts 2593          output pkts 2711         in bytes 221401
  out bytes 242072         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 971       out bcast bytes 75926
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 06:09:50, last time pvc status changed 06:09:45
  cir 192000    bc 19200     be 12800     byte limit 4000   interval
100
  mincir 96000     byte increment 2400  Adaptive Shaping none
  pkts 6         bytes 528       pkts delayed 0         bytes delayed 0
  shaping inactive
  traffic shaping drops 0
  Queueing strategy: fifo
  Output queue 0/40, 0 drop, 0 dequeued

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 2373          output pkts 2752         in bytes 202607
  out bytes 246973         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 972       out bcast bytes 75960
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 06:09:53, last time pvc status changed 06:09:28
  cir 192000    bc 19200     be 12800     byte limit 4000   interval
100
  mincir 96000     byte increment 2400  Adaptive Shaping none
  pkts 7         bytes 868       pkts delayed 0         bytes delayed 0
  shaping inactive
  traffic shaping drops 0
  Queueing strategy: fifo
  Output queue 0/40, 0 drop, 0 dequeued
```

Task 6.2

R1:

```
class-map match-any CRITICAL
 match packet length min 80 max 100
 match protocol ospf
!
class-map ANY
 match any
!
policy-map MARK
 class CRITICAL
 class ANY
  set fr-de
!
interface Serial 0/0.134
  service-policy output MARK
```

Task 6.2 Verification

```
Rack1R1#show frame-relay pvc 103

PVC Statistics for interface Serial0/0 (Frame Relay DTE)

DLCI = 103, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial0/0.134

  input pkts 30970         output pkts 27394        in bytes 2863566
  out bytes 2645571        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 21
  out bcast pkts 10381     out bcast bytes 624772
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 1d08h, last time pvc status changed 1d06h

Rack1R1#show policy-map interface serial 0/0.134

 Serial0/0.134

  Service-policy output: MARK

    Class-map: CRITICAL (match-any)
      5 packets, 420 bytes
      5 minute offered rate 0 bps
      Match: packet length min 80 max 100
        5 packets, 420 bytes
        5 minute rate 0 bps
      Match: protocol ospf
        0 packets, 0 bytes
        5 minute rate 0 bps

    Class-map: ANY (match-all)
      26 packets, 2552 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      QoS Set
        fr-de
```

Copyright © 2009 Internetwork Expert

```
          Packets marked 24

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

Task 6.3

SW2:

```
mls qos
!
interface FastEthernet0/24
 speed 10
 srr-queue bandwidth shape  0  0  10  0
 srr-queue bandwidth limit 30
```

Task 6.3 Verification

```
Rack1SW2#show mls qos maps dscp-output-q
   Dscp-outputq-threshold map:
     d1 :d2    0     1     2     3     4     5     6     7     8     9
     ------------------------------------------------------------
      0 :    02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01
      1 :    02-01 02-01 02-01 02-01 02-01 02-01 03-01 03-01 03-01 03-01
      2 :    03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01
      3 :    03-01 03-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      4 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 04-01 04-01
      5 :    04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      6 :    04-01 04-01 04-01 04-01

Rack1SW2#show mls qos interface FastEthernet 0/24 queueing
FastEthernet0/24
Egress Priority Queue : disabled
Shaped queue weights (absolute) :  0 0 10 0
Shared queue weights  :  25 25 25 25
The port bandwidth limit : 30  (Operational Bandwidth:30.44)
The port is mapped to qset : 1
```

Task 6.3 Breakdown

Setting the speed to 10M will meet the objective of the connection to BB2 being 10Mbps. The command srr-queue bandwidth limit will limit the overall egress to 30% of the physical speed, for the 3Mbps requirement. The command srr-queue bandwidth shape 0 0 10 0 will give 1/10th to the third queue. As mentioned in the QoS section of Volume 1, shaped weights still apply to the physical speed, not the bandwidth limit, when calculating queue rates.

Task 6.4

R5:

```
policy-map MARK
 class class-default
  police cir 256000 pir 512000
    conform-action set-prec-transmit 1
    exceed-action set-prec-transmit 0
```

Copyright © 2009 Internetwork Expert

```
    violate-action drop
!
interface FastEthernet 0/0
 service-policy input MARK
!
interface FastEthernet 0/1
 service-policy input MARK
```

Task 6.4 Verification

```
Rack1R5#show policy-map interface fastEthernet 0/0
 FastEthernet0/0

  Service-policy input: MARK

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      police:
          cir 256000 bps, bc 8000 bytes
          pir 512000 bps, be 16000 bytes
        conformed 0 packets, 0 bytes; actions:
          set-prec-transmit 1
        exceeded 0 packets, 0 bytes; actions:
          set-prec-transmit 0
        violated 0 packets, 0 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps, violate 0 bps

Rack1R5#show policy-map interface fastEthernet 0/1
 FastEthernet0/1

  Service-policy input: MARK

    Class-map: class-default (match-any)
      2 packets, 148 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      police:
          cir 256000 bps, bc 8000 bytes
          pir 512000 bps, be 16000 bytes
        conformed 2 packets, 148 bytes; actions:
          set-prec-transmit 1
        exceeded 0 packets, 0 bytes; actions:
          set-prec-transmit 0
        violated 0 packets, 0 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps, violate 0 bps
```

Copyright © 2009 Internetwork Expert

7. Security

Task 7.1

R2, R6, and SW2:
```
access-list 100 permit tcp any any
access-list 100 permit udp any any
access-list 100 deny 53 any any log
access-list 100 deny 55 any any log
access-list 100 deny 77 any any log
access-list 100 deny 103 any any log
access-list 100 permit ip any any
!
logging 187.1.38.100
```

R2:
```
interface FastEthernet0/0
 ip access-group 100 in
 ip access-group 100 out
```

R6:
```
interface Serial0/0
 ip access-group 100 in
 ip access-group 100 out
```

SW2:
```
interface Vlan28
 ip access-group 100 in
 ip access-group 100 out
```

Task 7.2 Breakdown

For the most part, this section is very straightforward. You are given specific items to block, and the devices are explicitly stated. There are a few additional items to keep in mind. When logging ACL entries, make sure that you have the logging set to an appropriate level. Since those messages are informational, logging will need to be at level 6 or 7. In this particular case, the default logging level is high enough, so no further configuration is needed. One other thing to watch carefully is the section that states “interest in the amount of packets that are denied by this filtering policy”. In this case, we are just logging when traffic is denied. It is possible that the section could also be alluding to using IP accounting with the “access-violations” option. This would be an example of a section where you may want to get additional clarification from the proctor whether they were just looking for general information, or tracking statistics for the denied traffic.

Task 7.2

R6:
```
ip inspect name FIREWALL http audit-trail on
ip inspect name FIREWALL ftp audit-trail on
ip inspect name FIREWALL dns
ip inspect name FIREWALL h323 router-traffic
```

Copyright © 2009 Internetwork Expert

```
!
ip access-list extended FROM_BB1
 permit tcp any eq bgp any
 permit tcp any any eq bgp
!
interface Serial0/0
 ip inspect FIREWALL out
 ip access-group FROM_BB1 in
```

Task 7.2 Verification

The access list created here takes the place of the access list from the prior section. Since you can only apply a single access list to an interface per direction, make sure that the access list integrates the requirements of both sections. In this case, the earlier section's requirements were to block some IP protocols, and to permit other traffic. This section's requirements are to only allow certain traffic inbound. By only allowing BGP traffic inbound, we are blocking the other protocols, so the requirements for both sections are met.

```
Rack1R6#show ip inspect all
Session audit trail is disabled
Session alert is enabled
one-minute (sampling period) thresholds are [400:500] connections
max-incomplete sessions thresholds are [400:500]
max-incomplete tcp connections per host is 50. Block-time 0 minute.
tcp synwait-time is 30 sec -- tcp finwait-time is 5 sec
tcp idle-time is 3600 sec -- udp idle-time is 30 sec
dns-timeout is 5 sec
Inspection Rule Configuration
 Inspection name FIREWALL
    http alert is on audit-trail is on timeout 3600
    ftp alert is on audit-trail is on timeout 3600
    dns alert is on audit-trail is off timeout 30
    sip alert is on audit-trail is off timeout 30
    h323 alert is on audit-trail is off timeout 3600
 inspection of router local traffic is enabled

Interface Configuration
 Interface Serial0/0
  Inbound inspection rule is not set
  Outgoing inspection rule is FIREWALL
    http alert is on audit-trail is on timeout 3600
    ftp alert is on audit-trail is on timeout 3600
    dns alert is on audit-trail is off timeout 30
    sip alert is on audit-trail is off timeout 30
    h323 alert is on audit-trail is off timeout 3600
 inspection of router local traffic is enabled
  Inbound access list is FROM_BB1
  Outgoing access list is 100

Rack1R6#show ip bgp summary
...

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
```

Copyright © 2009 Internetwork Expert

```
54.1.1.254      4    54     632     638       76    0    0 10:24:23       10
187.1.56.5      4   100    2030    2017       76    0    0 10:32:31
```

Copyright © 2009 Internetwork Expert