Task 1.1

```
SW1:
mac-address-table static 0030.1369.87a0 vlan 17 drop

errdisable recovery cause psecure-violation
errdisable recovery interval 60
!
interface FastEthernet0/7
 switchport mode access
 switchport port-security maximum 2
 switchport port-security
!
interface FastEthernet0/8
 switchport mode access
 switchport port-security maximum 2
 switchport port-security
```

Task 1.1 Breakdown

In addition to being used to restrict access to a specific MAC address, port-security can be used to limit the amount of MAC addresses that are allowed to send traffic into a port. This can be used on shared segments of the network in order to limit the amount of hosts that are allowed to access the network through a single port. As the default violation mode is shutdown, when the number of MAC addresses exceeds two, the interface is put into err-disabled state.

For the MAC restriction, the immediate reaction to this task is typically to use an extended MAC address access-list to deny traffic from this MAC address from entering interfaces Fa0/7 or Fa0/8. However, MAC address access-lists only affect non-IP traffic. Therefore, assuming that hosts on VLAN 17 are running IP (a fair assumption), using a MAC access-list to filter this host will have no effect.

As an alternative, traffic from this host has been effectively black holed by creating a static MAC address table (CAM table) entry for its MAC address. Much like static IP routing, a static MAC entry in the CAM table takes precedence over any dynamically learned reachability information.

Task 1.1 Verification

```
Rack1SW1#show port-security interface fa0/7
Port Security              : Enabled
Port Status                : Secure-down
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 0000.0000.0000:0
Security Violation Count   : 0

Rack1SW1#show port-security interface fa0/8
Port Security              : Enabled
Port Status                : Secure-down
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 0000.0000.0000:0
Security Violation Count   : 0
```

An additional MAC address is heard on the port and a violation occurs
↓ ↓ ↓

```
Rack1SW1#
%PM-4-ERR_DISABLE: psecure-violation error detected on Fa0/7, putting
Fa0/7 in err-disable state
Rack1SW1#
%PORT_SECURITY-2-PSECURE_VIOLATION: Security violation occurred, caused
by MAC address 00d0.586e.b930 on port FastEthernet0/7.
Rack1SW1#
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/7,
changed state to down
Rack1SW1#

Rack1SW1#show port-security interface fa0/7
Port Security              : Enabled
Port Status                : Secure-shutdown ← port disabled
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address        : 00d0.586e.b930
Security Violation Count   : 1
```

```
Rack1SW1#show interface status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/7                        err-disabled 17           auto   auto
10/100BaseTX
                             ↑     ↑     ↑
                          err-disabled state

Rack1SW1#show errdisable recovery
ErrDisable Reason    Timer Status
-----------------    --------------
udld                 Disabled
bpduguard            Disabled
security-violatio    Disabled
channel-misconfig    Disabled
vmps                 Disabled
pagp-flap            Disabled
dtp-flap             Disabled
link-flap            Disabled
l2ptguard            Disabled
psecure-violation    Enabled
gbic-invalid         Disabled
dhcp-rate-limit      Disabled
unicast-flood        Disabled
storm-control        Disabled
arp-inspection       Disabled
loopback             Disabled

Timer interval: 60 seconds

Interfaces that will be enabled at the next timeout:

Rack1SW1#show mac-address-table vlan 17 | inc Drop|Vlan|--
-------------------------------------------
Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  17    0030.1369.87a0    STATIC      Drop
```

Task 1.2

```
SW2:
interface FastEthernet0/2
 storm-control unicast level 3.00
```

Task 1.2 Breakdown

Storm control limits the amount of unicast, multicast, or broadcast traffic that is received in a layer 2 switchport. When the threshold of unicast or broadcast traffic is exceeded, traffic in excess of the threshold is dropped. When the multicast threshold is exceeded, all unicast, multicast, or broadcast traffic above the threshold is dropped. To configure storm-control, issue the **storm-control**

[unicast | broadcast | multicast] level [level] interface level command.

Task 1.2 Verification

```
Rack1SW1#show storm-control unicast
Interface  Filter State   Level    Current ← shows real-time level
---------  -------------  -------  -------
Fa0/1      inactive       100.00%  N/A
Fa0/2      Forwarding     3.00%    0.00%
Fa0/3      inactive       100.00%  N/A
```

☠ Pitfall

The storm-control command takes the level argument as a percentage of interface bandwidth. If you are asked to suppress traffic based on an absolute bandwidth level, such as 2Mbps, ensure to take into account whether the interface is running in 10Mbps or 100Mbps mode.

Task 1.3

```
SW1:
interface FastEthernet0/7
 switchport protected
!
interface FastEthernet0/8
 switchport protected
```

Task 1.3 Breakdown

Port protection prevents hosts that are in the same broadcast domain from directly communicating with each other at layer 2. This feature is especially useful when devices are placed in the same VLAN that would not normally be communicating with each other, such as web servers in a DMZ. Since there is typically not a valid case in which one server would initiate a connection to another server, this feature is very useful.

Task 1.3 Verification

```
Rack1SW1#show interfaces fastEthernet 0/7 switchport  | include
Protected
Protected: true

Rack1SW1#show interfaces fastEthernet 0/8 switchport  | include
Protected
Protected: true
```

Task 1.4

```
R4:
interface Serial0/0/0.54 point-to-point
 frame-relay interface-dlci 405
  class EEK
!
map-class frame-relay EEK
 frame-relay end-to-end keepalive mode bidirectional
 frame-relay end-to-end keepalive timer send 15

R5:
interface Serial0/0/0.54 point-to-point
 frame-relay interface-dlci 504
  class EEK
!
map-class frame-relay EEK
 frame-relay end-to-end keepalive mode bidirectional
 frame-relay end-to-end keepalive timer send 15
```

Task 1.4 Breakdown

When problems occur in the provider cloud, the end devices of the Frame Relay cloud may not detect a problem, as LMI communication with the local Frame Relay switch continues without interruption. For this reason, the DLCI may appear to be *active*, however, in reality no user traffic can be sent across the PVC. Frame Relay end-to-end keepalives can be used to detect this problem.

By participating in active request/response polling, Frame Relay end-to-end keepalives behave much like the hello packets in IGP. If a response is not heard back within the configured timer, the DLCI is brought to inactive state.

Task 1.4 Verification

```
Rack1R5#show frame-relay map
Serial0/0/0.54 (up): point-to-point dlci, dlci 504(0x1F8,0x7C80),
broadcast
          status defined, active

Rack1R5#ping 129.1.54.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 129.1.54.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

```
Rack1R5#show frame-relay end-to-end keepalive

End-to-end Keepalive Statistics for Interface Serial0/0/0 (Frame Relay
DTE)

DLCI = 504, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 20,       Receive Sequence Number: 21
Configured Event Window: 3,     Configured Error Threshold: 2
Total Observed Events: 23,      Total Observed Errors: 0
Monitored Events: 3,            Monitored Errors: 0
Successive Successes: 3,        End-to-end VC Status: UP

RECEIVE SIDE STATISTICS

Send Sequence Number: 20,       Receive Sequence Number: 19
Configured Event Window: 3,     Configured Error Threshold: 2
Total Observed Events: 22,      Total Observed Errors: 0
Monitored Events: 3,            Monitored Errors: 0
Successive Successes: 3,        End-to-end VC Status: UP
```

Task 2.1

SW3 and SW4:
```
interface Port-channel34
 ip ospf network point-to-point
```

Task 2.1 Breakdown

With an OSPF network type of broadcast, you will see both net link states and summary net link states for the area. Since a network type of point-to-point treats the local network slightly different, it will not have a net link entry for the area. Alternatively, you could also use the network type of point-to-multipoint.

Task 2.2

R1:
```
router bgp 200
 neighbor 129.1.17.7 route-reflector-client
```

R3:
```
router bgp 200
 neighbor 129.1.23.2 route-reflector-client
```

R4:
```
router bgp 100
 neighbor 129.1.46.6 route-reflector-client
```

R5:
```
router bgp 100
```

```
 neighbor 129.1.58.8 route-reflector-client
```

Task 2.2 Verification

```
Rack1R1#show ip bgp quote-regexp ^254 | begin Netw
   Network          Next Hop            Metric LocPrf Weight Path
*>i205.90.31.0      129.1.23.2               0    100      0 254 ?
*>i220.20.3.0       129.1.23.2               0    100      0 254 ?
*>i222.22.2.0       129.1.23.2               0    100      0 254 ?

Rack1R1#show ip bgp quote-regexp ^100 | begin Netw
   Network          Next Hop        Metric LocPrf Weight Path
*> 28.119.16.0/24   129.1.124.4                            0 100 54 i
* i                 129.1.17.7           0    100      0 100 54 i
*> 28.119.17.0/24   129.1.124.4                            0 100 54 i
* i                 129.1.17.7           0    100      0 100 54 i
*> 112.0.0.0        129.1.124.4                            0 100 54 50 60 i
* i                 129.1.17.7           0    100      0 100 54 50 60 i
*> 113.0.0.0        129.1.124.4                            0 100 54 50 60 i
* i                 129.1.17.7           0    100      0 100 54 50 60 i
*> 114.0.0.0        129.1.124.4                            0 100 54 i
* i                 129.1.17.7           0    100      0 100 54 i
*> 115.0.0.0        129.1.124.4                            0 100 54 i
* i                 129.1.17.7           0    100      0 100 54 i
<output omitted>

Rack1R5#show ip bgp quote-regexp ^54 | begin Netw
   Network          Next Hop            Metric LocPrf Weight Path
*>i28.119.16.0/24   129.1.58.8               0    100      0 54 i
* i                 129.1.46.6               0    100      0 54 i
*>i28.119.17.0/24   129.1.58.8               0    100      0 54 i
* i                 129.1.46.6               0    100      0 54 i
*>i112.0.0.0        129.1.58.8               0    100      0 54 50 60 i
* i                 129.1.46.6               0    100      0 54 50 60 i
*>i113.0.0.0        129.1.58.8               0    100      0 54 50 60 i
* i                 129.1.46.6               0    100      0 54 50 60 i
*>i114.0.0.0        129.1.58.8               0    100      0 54 i
* i                 129.1.46.6               0    100      0 54 i
*>i115.0.0.0        129.1.58.8               0    100      0 54 i
* i                 129.1.46.6               0    100      0 54 i
<output omitted>

Rack1R4#show ip bgp quote-regexp ^200 | beg Netw
   Network          Next Hop            Metric LocPrf Weight Path
* i205.90.31.0      129.1.58.8               0    100      0 200 254 ?
*                   129.1.124.1                            0 200 254 ?
*>                  129.1.124.2                            0 200 254 ?
* i220.20.3.0       129.1.58.8               0    100      0 200 254 ?
*                   129.1.124.1                            0 200 254 ?
*>                  129.1.124.2                            0 200 254 ?
* i222.22.2.0       129.1.58.8               0    100      0 200 254 ?
*                   129.1.124.1                            0 200 254 ?
*>                  129.1.124.2                            0 200 254 ?
```

Task 2.3

```
R1:
router bgp 200
 network 129.1.17.0 mask 255.255.255.0

R3:
router bgp 200
 network 129.1.3.0 mask 255.255.255.128
 network 129.1.3.128 mask 255.255.255.128

R4:
router bgp 100
 network 129.1.45.0 mask 255.255.255.248
 network 129.1.46.0 mask 255.255.255.0

SW2:
router bgp 100
 network 129.1.58.0 mask 255.255.255.0
```

Task 2.3 Verification

Verify BGP prefix origination

```
Rack1SW2#show ip bgp quote-regexp ^$
BGP table version is 21, 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
r>i129.1.45.0/29    150.1.4.4                0    100      0 i
r>i129.1.46.0/24    150.1.4.4                0    100      0 i

Rack1SW1#show ip bgp quote-regexp ^$
BGP table version is 25, local router ID is 150.1.7.7
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
r>i129.1.3.0/25     129.1.13.3               0    100      0 i
r>i129.1.3.128/25   129.1.13.3               0    100      0 i
r>i129.1.17.0/24    129.1.17.1               0    100      0 i
```

These devices show RIB failure (r), which is not something to be worried about in this case. Here, it just means that even though the route made it through the best path selection process for BGP, the route was not installed in the routing table. Here, it is due to a better route. In earlier IOS versions, networks with a RIB failure were not advertised to BGP peers, but that is no longer the case. Other items that could cause a RIB failure include memory issues or restrictions on the number of routes.

Task 2.4

```
R1:
router bgp 200
 neighbor 129.1.124.4 route-map BGP_OUT_TO_R4 out
!
ip prefix-list VLAN_3 seq 5 permit 129.1.3.0/25
!
ip prefix-list VLAN_33 seq 5 permit 129.1.3.128/25
!
route-map BGP_OUT_TO_R4 permit 10
 match ip address prefix-list VLAN_3
 set metric 20
!
route-map BGP_OUT_TO_R4 permit 20
 match ip address prefix-list VLAN_33
 set metric 10
!
route-map BGP_OUT_TO_R4 permit 1000

R2:
router bgp 200
 neighbor 129.1.124.4 route-map BGP_OUT_TO_R4 out
!
ip prefix-list VLANs_3_&_33 seq 5 permit 129.1.3.0/24 ge 25 le 25
!
route-map BGP_OUT_TO_R4 deny 10
 match ip address prefix-list VLANs_3_&_33
!
route-map BGP_OUT_TO_R4 permit 1000

SW1:
router bgp 200
 neighbor 129.1.78.8 route-map BGP_OUT_TO_SW2 out
!
ip prefix-list VLAN_3 seq 5 permit 129.1.3.0/25
!
ip prefix-list VLAN_33 seq 5 permit 129.1.3.128/25
!
route-map BGP_OUT_TO_SW2 permit 10
 match ip address prefix-list VLAN_3
 set metric 10
!
route-map BGP_OUT_TO_SW2 permit 20
 match ip address prefix-list VLAN_33
 set metric 20
!
route-map BGP_OUT_TO_SW2 permit 1000
```

Task 2.4 Breakdown

Recall how to influence the BGP best path selection process:

Attribute	Direction Applied	Traffic Flow Affected
Weight	Inbound	Outbound
Local-Preference	Inbound	Outbound
AS-Path	Outbound	Inbound
MED	Outbound	Inbound

In the above task, traffic engineering is applied on traffic destined for VLANs 3 and 33. AS 200 wants to affect how traffic is entering its AS that is destined for these VLANs. In order to effect an inbound traffic flow, either the MED or AS-Path attributes should be modified on outbound BGP updates. In the above solutions, MED has been used to influence the selection path. However, AS-Path could have been used in the same manner.

Traffic for VLAN 3 is preferred to come in the link between SW1 and SW2. This has been accomplished by advertising VLAN 3 with a more preferable (lower) MED value to SW2 than that which has been advertised to R4.

Additionally, traffic for VLAN 33 has a preferred entry point of the link between R1 and R4. This has been similarly accomplished by advertising VLAN 33 with a more preferable (lower) MED value to R4 than that which has been advertised to SW2.

Lastly, this requirement states that the link between R2 and R4 can not be used by AS 100 to get to VLAN 3 or VLAN 33. This is simply accomplished by filtering the advertisement of these networks from R2 to R4. Specifically, this has been configured by creating a prefix-list which matches both VLAN 3 and 33. Next, a route-map is configured that will be applied outbound from R2 to R4. The first sequence of the route-map is a deny sequence in which the previously created prefix-list is matched. This effectively stops the advertisement of VLANs 3 and 33 to R4.

☠ Pitfall

When changing BGP attributes through a route-map, don't forget to add an explicit permit sequence of the route-map at the end. If you leave the explicit permit out, all other prefixes not matched in the route-map will be denied.

```
Rack1R4#show ip bgp
BGP table version is 19, 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
```

```
*>i28.119.16.0/24   129.1.46.6                    100      0 54 i
*>i28.119.17.0/24   129.1.46.6                    100      0 54 i
*>i112.0.0.0        129.1.46.6               0    100      0 54 50 60 i
*>i113.0.0.0        129.1.46.6               0    100      0 54 50 60 i
*>i114.0.0.0        129.1.46.6               0    100      0 54 i
*>i115.0.0.0        129.1.46.6               0    100      0 54 i
*>i116.0.0.0        129.1.46.6               0    100      0 54 i
*>i117.0.0.0        129.1.46.6               0    100      0 54 i
*>i118.0.0.0        129.1.46.6               0    100      0 54 i
*>i119.0.0.0        129.1.46.6               0    100      0 54 i

The > denotes the best path                              1. weight both 0
 ↓                                                            ↓
*>i129.1.3.0/25     129.1.58.8              10    100      0 200 i
*                   129.1.124.1             20             0 200 i
<snip>

Rack1R4#show ip bgp 129.1.3.0 255.255.255.128
BGP routing table entry for 129.1.3.0/25, version 19
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  129.1.46.6 129.1.124.1 129.1.124.2
  200 ← 3. AS-Path both 1 AS long
    129.1.58.8 (metric 74) from 150.1.5.5 (150.1.5.5)
4. Origin both IGP  5. MED is tiebreaker  ↙ 2. local-preference both
100
      Origin IGP, metric 10, localpref 100, valid, internal, best
      Originator: 150.1.8.8, Cluster list: 150.1.5.5
  200 ← 3. AS-Path both 1 AS long
```

Task 2.5

R1:

```
ip as-path access-list 1 permit ^254$
!
route-map BGP_OUT_TO_R4 deny 30
 match as-path 1
```

SW1:

```
ip as-path access-list 1 permit ^254$
!
route-map BGP_OUT_TO_SW2 deny 30
 match as-path 1
```

Task 2.5 Breakdown

By filtering the advertisement of prefixes learned from AS 254 to AS 100, AS 100 is forced to use the path between R2 and R4 to reach these prefixes. This has been accomplished by creating an AS-Path access-list which matches prefixes that are from AS 254. Next, this AS-Path access-list is added to a new deny sequence of the route-map previously defined on R1 and SW1.

Task 2.5 Verification

```
Rack1R4#show ip bgp quote-regexp _254_ | begin Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 205.90.31.0      129.1.124.2                            0 200 254 ?
*> 220.20.3.0       129.1.124.2                            0 200 254 ?
*> 222.22.2.0       129.1.124.2                            0 200 254
```

Task 2.6

```
R4:
router bgp 100
 neighbor 129.1.124.1 default-originate
 neighbor 129.1.124.2 default-originate

SW2:
router bgp 100
 neighbor 129.1.78.7 default-originate
```

Task 2.6 Verification

```
Rack1SW1#show ip bgp 0.0.0.0
BGP routing table entry for 0.0.0.0/0, version 27
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Flag: 0x1860
  Advertised to update-groups:
     2
  100
    129.1.78.8 from 129.1.78.8 (150.1.8.8)
      Origin IGP, localpref 100, valid, external, best
  100
    129.1.17.1 from 129.1.17.1 (150.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal
```

Task 2.7

```
SW1:
router bgp 200
 neighbor 129.1.78.8 route-map BGP_IN_FROM_SW2 in
!
ip prefix-list DEFAULT seq 5 permit 0.0.0.0/0
!
route-map BGP_IN_FROM_SW2 permit 10
 match ip address prefix-list DEFAULT
 set local-preference 200
```

Task 2.7 Breakdown

In the above task, it is asked that SW1 be configured as the most preferable default exit point from AS 200. Since it is also stated that this configuration must be done on SW1, either local-preference or weight are candidates to affect the BGP best path selection. However, as weight is only locally significant, it is not a valid attribute to impact how the entire AS chooses the best path. Therefore, local-preference must be used to affect the selection.

In the above configuration, an IP prefix-list has been created which matches a default route. Next, a route-map is created that matches this prefix-list and sets the local-preference. As the default local-preference value is 100, any value above 100 would accomplish the desired goal.

Task 2.7 Verification

```
Rack1R1#show ip bgp
BGP table version is 75, local router ID is 150.1.1.1
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
*>i0.0.0.0          129.1.17.7               0    200      0 100 i
*                   129.1.124.4              0             0 100 i
<output omitted>

Rack1R1#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 200", distance 200, metric 0, candidate default path
  Tag 100, type internal
  Last update from 129.1.17.7 00:02:20 ago
  Routing Descriptor Blocks:
  * 129.1.17.7, from 129.1.17.7, 00:02:20 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
```

Shutdown the link to SW2 and verify the default routing again:

```
Rack1R1#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 200", distance 20, metric 0, candidate default path
  Tag 100, type external
  Last update from 129.1.124.4 00:00:36 ago
  Routing Descriptor Blocks:
  * 129.1.124.4, from 129.1.124.4, 00:00:36 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
```

Task 2.8

```
R2:
ip as-path access-list 1 permit ^100(_[0-9]+)?$
!
router bgp 200
 neighbor 129.1.124.4 filter-list 1 in
```

Task 2.8 Breakdown

Recall the special characters used in regular expressions:

Character	Meaning
^	Start of string
$	End of string
[]	Range of characters
-	Used to specify range (i.e. [0-9])
( )	Logical grouping
.	Any single character
*	Zero or more instances
+	One or more instance
?	Zero or one instance
_ (underscore)	Comma, open or close brace, open or close parentheses, start or end of string, or space

The above task requires that R2 only accept prefixes that have been originated in its directly connected provider's AS, as well as the provider's directly connected customers. This is a common view of the BGP table to take, since it is usually a safe assumption that your provider will have the best path to a destination if they are directly peering with that destination's AS.

The easiest way to create a regular expression is to think logically about what you are first try to match, and to write out all possibilities of these matches. For example, R2's directly connected AS is AS 100. Therefore, we can assume that there may be paths that have been originated inside AS 100. This is the first possibility we must match:

^100$

The ^ means that the path begins, the 100 matches AS 100, and the $ means that the path ends.

Next, be must also match the condition in which prefixes are originated from AS 100's directly connected ASs. However, we do not know which explicit AS numbers these are. Therefore, for the time being we will use the placeholder X. The second possibility is therefore as follows:

^100_X$

The ^ means that the path begins, the 100 matches AS 100, the _ matches a space, the X is our place holder for any single AS, and the $ means that the path ends.

Next let's reason out what X can represent. Since X is only one single AS, there will be no spaces, commas, parentheses, or any other special type characters. In other words, X must be a combination of integers. However, since we don't know what the exact path is, we must take into account that X may be more than one integer (i.e. 10 is two integers, 123 is three integers). The character used to match one or more instances is the plus sign. Therefore our second path is now:

^100_X+$

Where X is any single integer. Next we should define X. Again since we do not know what specific number or combination of numbers X will be, we can reason that it can be any combination of any number from zero to nine. This can be denoted as a the range from 0 to 9 by using brackets. Therefore our second choice is now:

^100_[0-9]+$

This will match all of AS 100's directly connected customers. Now we can stop where we are, and list both of these combinations in an as-path access-list, or we can try to combine them into one single line. To combine them, first let us compare what is different between them.

^100$
^100_[0-9]+$

From looking at the expressions, it is evident that the sequence _[0-9]+ is the difference. For the time being let us represent this sequence with the variable A. In the first case, A does not exist in the expression. In the second case, A does exist in the expression. In other words, A is either true or false. True or false (0 or 1) is represented by the character ?

Therefore we can reduce our expression to:

^100A?$

However, if we simply write the expression as ^100_[0-9]+?$, the question mark will apply to the plus sign. Instead, we want the question mark to apply to the string _[0-9]+ as a whole. Therefore, this string can be grouped together using parentheses. Parentheses are used in regular expressions as simply a logical grouping. Therefore, our final expression reduces to:

^100(_[0-9]+)?$

In order to meet the requirement of still being eligible as a default exit point, make sure to verify that the policy does not block the default 0.0.0.0 route from R4.

Note

To match a question mark in IOS, the escape sequence CTRL-V or ESC-Q must be entered first.

Task 2.8 Verification

```
Rack1R2#show ip bgp neighbors 129.1.124.4 routes
BGP table version is 106, 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
*  0.0.0.0          129.1.124.4              0             0 100 i
*> 28.119.16.0/24   129.1.124.4                            0 100 54 i
*> 28.119.17.0/24   129.1.124.4                            0 100 54 i
*> 114.0.0.0        129.1.124.4                            0 100 54 i
*> 115.0.0.0        129.1.124.4                            0 100 54 i
*> 116.0.0.0        129.1.124.4                            0 100 54 i
*> 117.0.0.0        129.1.124.4                            0 100 54 i
*> 118.0.0.0        129.1.124.4                            0 100 54 i
*> 119.0.0.0        129.1.124.4                            0 100 54 i
*> 129.1.45.0/29    129.1.124.4              0             0 100 i
*> 129.1.46.0/24    129.1.124.4              0             0 100 i
*> 129.1.58.0/24    129.1.124.4                            0 100 i
```

Verify paths for non-direct customers of AS100:

```
Rack1R2#show ip bgp quote-regexp ^100_[0-9]+(_[0-9]+)+$
BGP table version is 106, 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
*>i112.0.0.0        129.1.13.1           0    100      0 100 54 50 60 i
*>i113.0.0.0        129.1.13.1           0    100      0 100 54 50 60 i
```

Task 2.9

```
R1:
ip prefix-list DEFAULT seq 5 permit 0.0.0.0/0
!
route-map BGP_IN_FROM_R4 permit 10
 match ip address prefix-list DEFAULT
 set local-preference 50
!
route-map BGP_IN_FROM_R4 permit 1000
!
router bgp 200
  neighbor 129.1.124.4 route-map BGP_IN_FROM_R4 in
```

Task 2.9 Breakdown

Similar to task 6.17, the local-preference of the default route learned from AS 100 has been modified in order to affect how traffic leaves AS 200. In this case, R1 is configured as the least preferred exit point by setting the local-preference lower than the other two values of 100 and 200.

Task 2.9 Verification

Verify the default routing in AS200. Look for the most preferred default route when all links to AS100 are up:

```
Rack1R3#show ip bgp 0.0.0.0
BGP routing table entry for 0.0.0.0/0, version 132
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  100
    129.1.17.7 (metric 20514560) from 129.1.13.1 (150.1.1.1)
      Origin IGP, metric 0, localpref 200, valid, internal, best
      Originator: 150.1.7.7, Cluster list: 150.1.1.1
```

Next, shutdown the link between SW1 and SW2. Then, verify the BGP default route again:

```
Rack1R3#show ip bgp 0.0.0.0
BGP routing table entry for 0.0.0.0/0, version 134
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x840
  Advertised to update-groups:
```

```
     1
  100, (Received from a RR-client)
    129.1.23.2 from 129.1.23.2 (150.1.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Finally, shut down the serial interface on R2 and verify the BGP routes again:

```
Rack1R3#show ip bgp 0.0.0.0
BGP routing table entry for 0.0.0.0/0, version 160
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Advertised to update-groups:
     2
  100
    129.1.13.1 from 129.1.13.1 (150.1.1.1)
      Origin IGP, metric 0, localpref 50, valid, internal, best
```

Task 2.10

R2:
```
router bgp 200
 aggregate-address 129.1.0.0 255.255.0.0
 aggregate-address 150.1.0.0 255.255.240.0
 neighbor 129.1.23.3 route-map BGP_OUT_TO_R3 out
!
ip prefix-list AGGREGATE seq 5 permit 129.1.0.0/16
ip prefix-list AGGREGATE seq 10 permit 150.1.0.0/20
!
route-map BGP_OUT_TO_R4 deny 20
 match ip address prefix-list AGGREGATE
!
route-map BGP_OUT_TO_R3 deny 10
 match ip address prefix-list AGGREGATE
!
route-map BGP_OUT_TO_R3 permit 1000
```

R6:
```
router bgp 100
 aggregate-address 129.1.0.0 255.255.0.0
 aggregate-address 150.1.0.0 255.255.240.0
 neighbor 129.1.46.4 route-map BGP_OUT_TO_R4 out
!
ip prefix-list AGGREGATE seq 5 permit 129.1.0.0/16
ip prefix-list AGGREGATE seq 10 permit 150.1.0.0/20
!
route-map BGP_OUT_TO_R4 deny 10
 match ip address prefix-list AGGREGATE
!
route-map BGP_OUT_TO_R4 permit 1000
```

SW2:
```
router bgp 100
 aggregate-address 129.1.0.0 255.255.0.0
 aggregate-address 150.1.0.0 255.255.240.0
```

```
 neighbor 129.1.78.7 route-map BGP_OUT out
 neighbor 129.1.58.5 route-map BGP_OUT out
!
ip prefix-list AGGREGATE seq 5 permit 129.1.0.0/16
ip prefix-list AGGREGATE seq 10 permit 150.1.0.0/20
!
route-map BGP_OUT deny 10
 match ip address prefix-list AGGREGATE
!
route-map BGP_OUT permit 1000
```

Task 2.10 Breakdown

The above task illustrates a straightforward aggregation configuration, in which the border routers of the network are advertising an aggregate block of the internal address space to the backbones. In addition to this, the aggregate block is denied from being advertised to the internal routers by matching it in a prefix-list, and denying it in a route-map applied to the iBGP neighbors.

Task 2.10 Verification

Verify the summary prefix generation. For example on SW2:

```
Rack1SW2#show ip bgp 129.1.0.0
BGP routing table entry for 129.1.0.0/16, version 59
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  Local, (aggregated by 100 150.1.8.8)
    0.0.0.0 from 0.0.0.0 (150.1.8.8)
      Origin IGP, localpref 100, weight 32768, valid, aggregated,
local, atomic-aggregate, best
```

Confirm that SW2 does not send summary to internal routers:

```
Rack1SW2#show ip bgp neigh 129.1.58.5 advertised-routes | inc 129.1.0.0
Rack1SW2#

Rack1SW2#show ip bgp neigh 129.1.78.7 advertised-routes | inc 129.1.0.0
Rack1SW2#
```

Task 3.1

R1, R2, R3, R4 and R6:

```
ipv6 unicast-routing
```

R1:

```
interface FastEthernet0/0
 ipv6 address 2001:CC1E:1:1::1/64
```

R2:

```
interface Serial0/1
 ipv6 address 2001:CC1E:1:23::2/64
```

R3:
```
interface FastEthernet0/0
 ipv6 address 2001:CC1E:1:3::3/64
!
interface Serial1/3
 ipv6 address 2001:CC1E:1:23::3/64
```

R4:
```
interface FastEthernet0/1
 ipv6 address 2001:CC1E:1:46::4/64
```

R6:
```
interface FastEthernet0/0
 ipv6 address 2001:CC1E:1:46::6/64
```

Task 3.2

R1:
```
interface Serial0/0
 ipv6 address 2001:CC1E:1:124::1/64
 ipv6 address FE80::1 link-local
 frame-relay map ipv6 FE80::2 104
 frame-relay map ipv6 FE80::4 104 broadcast
 frame-relay map ipv6 2001:CC1E:1:124::2 104
 frame-relay map ipv6 2001:CC1E:1:124::4 104
```

R2:
```
interface Serial0/0
 ipv6 address 2001:CC1E:1:124::2/64
 ipv6 address FE80::2 link-local
 frame-relay map ipv6 FE80::4 204 broadcast
 frame-relay map ipv6 2001:CC1E:1:124::1 204
 frame-relay map ipv6 2001:CC1E:1:124::4 204
 frame-relay map ipv6 FE80::1 204
```

R4:
```
interface Serial0/0/0.124 multipoint
 ipv6 address 2001:CC1E:1:124::4/64
 ipv6 address FE80::4 link-local
 frame-relay map ipv6 FE80::2 402 broadcast
 frame-relay map ipv6 2001:CC1E:1:124::1 401
 frame-relay map ipv6 2001:CC1E:1:124::2 402
 frame-relay map ipv6 FE80::1 401 broadcast
```

Task 3.2 Verification

```
Rack1R4#show frame-relay map
Serial0/0/0.124 (up): ipv6 FE80::2 dlci 402(0x192,0x6420), static,
              broadcast,
              CISCO, status defined, active
Serial0/0/0.124 (up): ip 129.1.124.2 dlci 402(0x192,0x6420), static,
              broadcast,
```

```
              CISCO, status defined, active
Serial0/0/0.124 (up): ipv6 2001:CC1E:1:124::1 dlci 401(0x191,0x6410),
static,
              CISCO, status defined, active
Serial0/0/0.124 (up): ipv6 2001:CC1E:1:124::2 dlci 402(0x192,0x6420),
static,
              CISCO, status defined, active
Serial0/0/0.124 (up): ipv6 FE80::1 dlci 401(0x191,0x6410), static,
              broadcast,
              CISCO, status defined, active
Serial0/0/0.124 (up): ip 129.1.124.1 dlci 401(0x191,0x6410), static,
              broadcast,
              CISCO, status defined, active
Serial0/0/0.54 (up): point-to-point dlci, dlci 405(0x195,0x6450),
broadcast
          status defined, active

Rack1R2#show frame-relay map
Serial0/0 (up): ipv6 FE80::4 dlci 204(0xCC,0x30C0), static,
              broadcast,
              CISCO, status defined, active
Serial0/0 (up): ip 129.1.124.4 dlci 204(0xCC,0x30C0), static,
              broadcast,
              CISCO, status defined, active
Serial0/0 (up): ipv6 2001:CC1E:1:124::1 dlci 204(0xCC,0x30C0), static,
              CISCO, status defined, active
Serial0/0 (up): ipv6 2001:CC1E:1:124::4 dlci 204(0xCC,0x30C0), static,
              CISCO, status defined, active
Serial0/0 (up): ipv6 FE80::1 dlci 204(0xCC,0x30C0), static,
              CISCO, status defined, active
Serial0/0 (up): ip 129.1.124.1 dlci 204(0xCC,0x30C0), static,
              CISCO, status defined, active

Rack1R1#show frame-relay map
Serial0/0 (up): ipv6 FE80::2 dlci 104(0x68,0x1880), static,
              CISCO, status defined, active
Serial0/0 (up): ip 129.1.124.2 dlci 104(0x68,0x1880), static,
              CISCO, status defined, active
Serial0/0 (up): ipv6 FE80::4 dlci 104(0x68,0x1880), static,
              broadcast,
              CISCO, status defined, active
Serial0/0 (up): ip 129.1.124.4 dlci 104(0x68,0x1880), static,
              broadcast,
              CISCO, status defined, active
Serial0/0 (up): ipv6 2001:CC1E:1:124::2 dlci 104(0x68,0x1880), static,
              CISCO, status defined, active
Serial0/0 (up): ipv6 2001:CC1E:1:124::4 dlci 104(0x68,0x1880), static,
              CISCO, status defined, active
```

Test basic connectivity:

```
Rack1R1#ping 2001:CC1E:1:124::2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:CC1E:1:124::2, timeout is 2
seconds:
```

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/112
ms

Rack1R1#ping 2001:CC1E:1:124::4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:CC1E:1:124::4, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/32 ms

Rack1R4#ping ipv6 2001:CC1E:1:46::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:CC1E:1:46::6, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms

Rack1R2#ping 2001:CC1E:1:23::3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:CC1E:1:23::3, timeout is 2
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

Task 3.3

R4:

```
ipv6 router eigrp 46
 no shut
!
interface fastEtherent 0/1
 ipv6 eigrp 46

ipv6 prefix-list TEST permit 0::0/0 le 64
```

R6:

```
interface fastEthernet 0/0
 ipv6 eigrp 46
!
interface loopback601
 ipv6 address 2001:205:90:31::1/48
 ipv6 eigrp 46
!
interface loopback602
 ipv6 address 2001:220:20:3::1/64
 ipv6 eigrp 46
!
interface loopback603
 ipv6 address 2001:222:22:2::1/80
 ipv6 eigrp 46
```

```
!
ipv6 router eigrp 46
 no shut
```

Task 3.3 Verification

```
Rack1R4#show ipv6 route eigrp
IPv6 Routing Table - Default - 8 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   2001:205:90::/48 [90/156160]
     via FE80::219:56FF:FED4:F878, FastEthernet0/1
D   2001:220:20:3::/64 [90/156160]
     via FE80::219:56FF:FED4:F878, FastEthernet0/1
D   2001:222:22:2::/80 [90/156160]
     via FE80::219:56FF:FED4:F878, FastEthernet0/1
Rack1R4#
```

For now, we will just configure the prefix list since there is not currently any advertisements going to R2 or R3. Note: Some IOS versions may be missing part of the context sensitive help for the command. Try typing in the entire command.

Task 3.4

```
R4:
interface serial 0/0/0.124
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-multipoint

R2:
interface Serial0/1
ipv6 ospf 1 area 0
ipv6 ospf network point-to-point

interface Serial0/0
ipv6 ospf 1 area 0
ipv6 ospf network point-to-multipoint

R1:
interface Serial0/0
ipv6 ospf 1 area 0
ipv6 ospf network point-to-multipoint

interface FastEthernet0/0
ipv6 ospf 1 area 0

R3:
```

```
interface Serial1/3
 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point

interface FastEthernet0/0
 ipv6 ospf 1 area 0
```

Task 3.4 Verification

Verify OSPFv3 neighbors and routes:

```
Rack1R4#show ipv6 ospf neigh

Neighbor ID     Pri   State           Dead Time   Interface ID
Interface
150.1.1.1         1   FULL/  -        00:01:34    5
Serial0/0/0.124
150.1.2.2         1   FULL/  -        00:01:46    5
Serial0/0/0.124
Rack1R4#

Rack1R4#show ipv6 route ospf
IPv6 Routing Table - Default - 12 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
O   2001:CC1E:1:1::/64 [110/65]
     via FE80::1, Serial0/0/0.124
O   2001:CC1E:1:3::/64 [110/129]
     via FE80::2, Serial0/0/0.124
O   2001:CC1E:1:23::/64 [110/128]
     via FE80::2, Serial0/0/0.124
O   2001:CC1E:1:124::1/128 [110/64]
     via FE80::1, Serial0/0/0.124
O   2001:CC1E:1:124::2/128 [110/64]
     via FE80::2, Serial0/0/0.124
Rack1R4#
```

Task 3.5

R4:

```
ipv6 router eigrp 46
 redistribute ospf 1
 redistribute connected
 default-metric 10000 10 255 1 1500
!
ipv6 router ospf 1
 redist eigrp 46 route-map NO65
 redist conn


route-map NO65
 match ipv6 address prefix TEST

interface FastEthernet0/1
 ipv6 summary-address eigrp 46 2001:222:22:2::/64
```

Task 3.5 Verification

Make sure to verify by looking at your routing tables on R6 and R3, and verify that both show all the networks. To restrict to prefixes with a mask of 64 bits or less, you can add the prefix list configured earlier to a route map with the redistribution. In order to still have reachability to the loopback on R6, a summary needs to be configured with a mask length less than 64 bits.

```
Rack1R6#show ipv6 route
IPv6 Routing Table - Default - 16 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
C   2001:205:90::/48 [0/0]
     via Loopback601, directly connected
L   2001:205:90:31::1/128 [0/0]
     via Loopback601, receive
C   2001:220:20:3::/64 [0/0]
     via Loopback602, directly connected
L   2001:220:20:3::1/128 [0/0]
     via Loopback602, receive
D   2001:222:22:2::/64 [90/158720]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
C   2001:222:22:2::/80 [0/0]
     via Loopback603, directly connected
L   2001:222:22:2::1/128 [0/0]
     via Loopback603, receive
EX  2001:CC1E:1:1::/64 [170/261120]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
EX  2001:CC1E:1:3::/64 [170/261120]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
EX  2001:CC1E:1:23::/64 [170/261120]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
C   2001:CC1E:1:46::/64 [0/0]
     via FastEthernet0/0, directly connected
L   2001:CC1E:1:46::6/128 [0/0]
     via FastEthernet0/0, receive
EX  2001:CC1E:1:124::/64 [170/261120]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
EX  2001:CC1E:1:124::1/128 [170/261120]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
EX  2001:CC1E:1:124::2/128 [170/261120]
     via FE80::207:EFF:FE7A:1125, FastEthernet0/0
L   FF00::/8 [0/0]
     via Null0, receive
```

```
Rack1R3#show ipv6 route
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
OE2  2001:205:90::/48 [110/20]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
OE2  2001:220:20:3::/64 [110/20]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
OE2  2001:222:22:2::/64 [110/20]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
O   2001:CC1E:1:1::/64 [110/910]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
C   2001:CC1E:1:3::/64 [0/0]
     via ::, FastEthernet0/0
L   2001:CC1E:1:3::3/128 [0/0]
     via ::, FastEthernet0/0
C   2001:CC1E:1:23::/64 [0/0]
     via ::, Serial1/3
L   2001:CC1E:1:23::3/128 [0/0]
     via ::, Serial1/3
OE2  2001:CC1E:1:46::/64 [110/20]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
O   2001:CC1E:1:124::1/128 [110/909]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
O   2001:CC1E:1:124::2/128 [110/781]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
O   2001:CC1E:1:124::4/128 [110/845]
     via FE80::211:BBFF:FEA2:6C00, Serial1/3
L   FE80::/10 [0/0]
     via ::, Null0
L   FF00::/8 [0/0]
     via ::, Null0
Rack1R3#
```

Task 4.1

SW4:

```
sdm prefer extended-match

ip vrf TEST
 rd 44:44

interface FastEthernet0/6
 ip vrf forwarding TEST
 no switchport
 ip address 10.0.0.10 255.255.255.0


router ospf 129 vrf TEST
 network 10.0.0.10 0.0.0.0 area 0
```

Task 4.1 Breakdown

Configuring a VRF on the switch may require a change to the SDM profile for 3550 switches.

Task 4.1 Verification

```
Rack1SW4#show ip vrf TEST
  Name                             Default RD          Interfaces
  TEST                             44:44               Fa0/6
Rack1SW4#

Rack1R6#ping vrf VPNB 10.0.0.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.0.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Rack1R6#
```

Task 4.2

```
R6:
interface FastEthernet0/0
 mpls ip

R4:
Interface Serial0/0/0.54
mpls ip
!
Interface FastEthernet0/0
 mpls ip
!
interface FastEthernet0/1
 mpls ip

R5:
interface Serial0/0/0.54
 mpls ip
!
interface FastEthernet0/1
 mpls ip
```


Task 4.2 Breakdown

LDP is the default label protocol, so all that is needed is to enable MPLS on the interfaces.

Task 4.2 Verification

Verify that the neighbor adjacencies form, and check the output of ***show mpls ldp neighbor*** *and* ***show mpls ldp discovery****.*

```
Rack1R4#show mpls ldp neigh
    Peer LDP Ident: 150.1.5.5:0; Local LDP Ident 150.1.4.4:0
        TCP connection: 150.1.5.5.22578 - 150.1.4.4.646
        State: Oper; Msgs sent/rcvd: 16/15; Downstream
        Up time: 00:00:47
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 129.1.45.5
          Serial0/0/0.54, Src IP addr: 129.1.54.5
        Addresses bound to peer LDP Ident:
          129.1.58.5      129.1.45.5      129.1.54.5      150.1.5.5
          150.1.0.255
    Peer LDP Ident: 150.1.6.6:0; Local LDP Ident 150.1.4.4:0
        TCP connection: 150.1.6.6.65364 - 150.1.4.4.646
        State: Oper; Msgs sent/rcvd: 16/16; Downstream
        Up time: 00:00:47
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 129.1.46.6
        Addresses bound to peer LDP Ident:
          129.1.46.6      54.1.1.6        150.1.6.6

Rack1R4#show mpls ldp discovery
 Local LDP Identifier:
    150.1.4.4:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 150.1.5.5:0
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 150.1.6.6:0
        Serial0/0/0.54 (ldp): xmit/recv
            LDP Id: 150.1.5.5:0
Rack1R4#

Rack1R5#show mpls ldp neigh
    Peer LDP Ident: 150.1.4.4:0; Local LDP Ident 150.1.5.5:0
        TCP connection: 150.1.4.4.646 - 150.1.5.5.22578
        State: Oper; Msgs sent/rcvd: 18/19; Downstream
        Up time: 00:02:32
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 129.1.45.4
          Serial0/0/0.54, Src IP addr: 129.1.54.4
        Addresses bound to peer LDP Ident:
```

```
129.1.45.4      129.1.46.4      129.1.54.4      129.1.124.4
150.1.4.4       150.1.0.255     129.1.45.6
```

```
Rack1R5#show mpls ldp disc
 Local LDP Identifier:
    150.1.5.5:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 150.1.4.4:0
        Serial0/0/0.54 (ldp): xmit/recv
            LDP Id: 150.1.4.4:0

Rack1R6#show mpls ldp discovery
 Local LDP Identifier:
    150.1.6.6:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 150.1.4.4:0

Rack1R6#show mpls ldp neigh
    Peer LDP Ident: 150.1.4.4:0; Local LDP Ident 150.1.6.6:0
        TCP connection: 150.1.4.4.646 - 150.1.6.6.65364
        State: Oper; Msgs sent/rcvd: 19/20; Downstream
        Up time: 00:03:25
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 129.1.46.4
        Addresses bound to peer LDP Ident:
          129.1.45.4      129.1.46.4      129.1.54.4      129.1.124.4
          150.1.4.4       150.1.0.255     129.1.45.6
Rack1R6#
```

For testing, you can also ping from R6 to R5, and verify that you see the counters increment in the output of show mpls forwarding.

```
Rack1R4#show mpls forw 150.1.5.5
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next
Hop
Label  Label or VC   or Tunnel Id      Switched      interface
18     No Label      150.1.5.5/32      570           Fa0/0
129.1.45.5
Rack1R4#
```

Task 4.3

```
R5:
router bgp 100
  no bgp default ipv4-unicast
  neighbor 150.1.6.6 remote-as 100
  neighbor 150.1.6.6 update-source lo0
 address-family vpnv4 uni
  neighbor 150.1.6.6 activate
  address-family ipv4 vrf VPNA
  redistribute connected
```

```
R6:
router bgp 100
no bgp default ipv4-unicast
 neighbor 150.1.5.5 remote-as 100
 neighbor 150.1.5.5 upd lo0
address-family vpnv4 uni
 neighbor 150.1.5.5 activate
address-family ipv4 vrf VPNB
 redistribute connected

router ospf 12 vrf VPNB
 redist bgp 100 subnets
```

Task 4.3 Breakdown

Here, we have the neighbors added to BGP for the address families, in addition to redistribution for the VRFs. When redistributing into BGP for the VRFs on the endpoints, normally you would redistribute based on the VRF routing protocols. Since R6 only has the connected network in OSPF, redistribute connected is sufficient for the reachability for this section.

Task 4.3 Verification

Verify that R5 and R6 show the routes.

```
Rack1R5#show ip bgp vpnv4 all
BGP table version is 7, 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
Route Distinguisher: 4:4 (default for vrf VPNA)
*>i10.0.0.0/24      150.1.6.6                0    100      0 ?
*> 50.0.0.1/32      0.0.0.0                  0         32768 ?
*> 51.0.0.1/32      0.0.0.0                  0         32768 ?
Route Distinguisher: 6:6
*>i10.0.0.0/24      150.1.6.6                0    100      0 ?

Rack1R6#show ip bgp vpnv4 all
BGP table version is 7, 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
Route Distinguisher: 4:4
*>i50.0.0.1/32      150.1.5.5                0    100      0 ?
*>i51.0.0.1/32      150.1.5.5                0    100      0 ?
Route Distinguisher: 6:6 (default for vrf VPNB)
```

```
*> 10.0.0.0/24      0.0.0.0                  0         32768 ?
*>i50.0.0.1/32      150.1.5.5                0    100      0 ?
*>i51.0.0.1/32      150.1.5.5                0    100      0 ?
Rack1R6#
```

```
Rack1R5#show ip route vrf VPNA | beg Gate
Gateway of last resort is not set

     51.0.0.0/32 is subnetted, 1 subnets
C       51.0.0.1 is directly connected, Loopback51
     50.0.0.0/32 is subnetted, 1 subnets
C       50.0.0.1 is directly connected, Loopback50
     10.0.0.0/24 is subnetted, 1 subnets
B       10.0.0.0 [200/0] via 150.1.6.6, 00:06:50
Rack1R5#
```

Looking at the mpls forwarding table on R4, you can see entries for R5 and R6's loopbacks with a pop tag.

```
Rack1R4#show mpls forw 150.1.5.5
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next
Hop
Label  Label or VC   or Tunnel Id      Switched      interface
18     Pop Label     150.1.5.5/32      4293          Fa0/0
129.1.45.5
Rack1R4#show mpls forw 150.1.6.6
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next
Hop
Label  Label or VC   or Tunnel Id      Switched      interface
19     Pop Label     150.1.6.6/32      6227          Fa0/1
129.1.46.6
Rack1R4#
```

Next, take a look at the traffic flow. Starting on R5, look at the CEF entry for the destination network.

```
Rack1R5#show ip cef vrf VPNA 10.0.0.0/24
10.0.0.0/24
  nexthop 129.1.45.4 FastEthernet0/1 label 19 26
```

The CEF table gives us the label information, which can be traced through R5 to R6.

```
Rack1R5#show mpls forw label 19
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next
Hop
Label  Label or VC   or Tunnel Id      Switched      interface
19     19            150.1.6.6/32      0             Fa0/1
129.1.45.4
```

```
Rack1R4#show mpls forw label 19
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next
Hop
Label  Label or VC   or Tunnel Id      Switched      interface
```

```
19     Pop Label     150.1.6.6/32      6805          Fa0/1
129.1.46.6
Rack1R4#

Rack1R6#show mpls forwarding label 26
Local  Outgoing      Prefix            Bytes Label   Outgoing   Next
Hop
Label  Label or VC   or Tunnel Id      Switched      interface
26     No Label      10.0.0.0/24[V]    1140          aggregate/VPNB
Rack1R6#
```

Task 5.1

R3:

```
interface Serial1/2
 ip multicast helper-map 225.25.25.25 129.1.23.255 111
!
interface Serial1/3
 ip directed-broadcast
!
access-list 111 permit udp any any eq 31337
!
ip forward-protocol udp 31337
```

R2:

```
interface Serial0/1
 ip multicast helper-map broadcast 225.25.25.25 111
!
access-list 111 permit udp any any eq 31337
!
ip forward-protocol udp 31337
```

Task 5.1 Verification

In order to test the above configuration, a router configured with the IP SLA monitor feature in VLAN 17 will be designated as the multicast server, while another router in VLAN 22 will be the multicast client:

SW1:

```
rtr 1
 type udpEcho  dest-ipaddr 225.25.25.25 dest-port 31337 source-ipaddr
129.1.17.7 source-port 31337 control disable
 timeout 1
 frequency 5
rtr schedule 1 start-time now
!
ip multicast-routing distributed
!
interface Vlan 17
 ip pim dense-mode
! Make sure to remove the PIM mode when done testing!
```

```
R1:
Rack1R1(config)#interface fastethernet 0/0
Rack1R1(config-if)#no ip mroute-cache
                   ↑       ↑       ↑
            multicast fast switching disabled on
              the incoming interface so debug
                  output can be seen

Rack1R1#show ip mroute
<snip>

(*, 225.25.25.25), 00:08:28/stopped, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/1, Forward/Dense, 00:08:28/00:00:00

(129.1.17.7, 225.25.25.25), 00:08:28/00:02:50, flags: T
  Incoming interface: FastEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/1, Forward/Dense, 00:08:28/00:00:00
                              ↑      ↑      ↑
Indicates a multicast feed destined for 225.25.25.25
is being received from 129.1.17.7 in interface
FastEthernet0/0, and is forwarded out interface Serial0/1

Rack1R1#debug ip mpacket
IP multicast packets debugging is on
Rack1R1#
IP(0): s=129.1.17.7 (FastEthernet0/0) d=225.25.25.25 (Serial0/1) id=0,
prot=17, len=44(44), mforward
Rack1R1#
IP(0): s=129.1.17.7 (FastEthernet0/0) d=225.25.25.25 (Serial0/1) id=0,
prot=17, len=44(44), mforward
Rack1R1#
IP(0): s=129.1.17.7 (FastEthernet0/0) d=225.25.25.25 (Serial0/1) id=0,
prot=17, len=44(44), mforward
↑       ↑       ↑
packets generated by SLA are received by R1
in the Ethernet interface connecting
to VLAN 17 and are forwarded out
interface Serial 0/1 to R3

Rack1R3#show ip mroute
<snip>

(*, 225.25.25.25), 00:18:53/stopped, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/2, Forward/Dense, 00:18:53/00:00:00

(129.1.17.7, 225.25.25.25), 00:12:52/00:02:57, flags: PLTX
  Incoming interface: Serial1/2, RPF nbr 129.1.13.1
  Outgoing interface list: Null
  ↑             ↑             ↑
  Feed is received in Serial1/2
```

```
  but it is not forwarded anywhere

Rack1R2#debug ip packet detail 111
IP packet debugging is on (detailed) for access list 111
↑               ↑                ↑
Previously defined access-list 111
used to filter debug output

Rack1R2#
IP: s=129.1.17.7 (Serial0/1), d=255.255.255.255, len 44, rcvd 2
    UDP src=31337, dst=31337
Rack1R2#
IP: s=129.1.17.7 (Serial0/1), d=255.255.255.255, len 44, rcvd 2
    UDP src=31337, dst=31337
↑               ↑                ↑
R2 received the feed as an IP broadcast

Rack1R2#show access-lists
Extended IP access list 111
    10 permit udp any any eq 31337 (319 matches)
      ↑               ↑                ↑
      Broadcast feed hits the helper-map and
      is translated back into a multicast feed

SW4: (Testing only, remove when done)
ip multicast-routing
ip mroute 129.1.17.7 255.255.255.255 192.10.1.2
!
interface vlan 22
ip address 192.10.1.10 255.255.255.0
ip pim dense


Client#
IP(0): s=129.1.17.7 (FastEthernet0/0) d=225.25.25.25 id=0, prot=17,
len=60(44), mroute olist null
Rack1SW4#
IP(0): s=129.1.17.7 (FastEthernet0/0) d=225.25.25.25 id=0, prot=17,
len=60(44), mroute olist null
                    ↑               ↑                ↑
              Client receives transmission as a multicast
           Broadcast conversion is transparent to the client


Rack1SW4#show ip mroute | beg \(129
(129.1.17.7, 225.25.25.25), 00:01:43/00:02:56, flags: PT
  Incoming interface: Vlan22, RPF nbr 192.10.1.2, Mroute
  Outgoing interface list: Null

Rack1R2#show ip mroute | beg \(129
(129.1.17.7, 225.25.25.25), 00:08:34/00:02:57, flags: T
  Incoming interface: Serial0/1, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Dense, 00:01:26/00:00:00
```

A few notes on testing:

Without a PIM neighbor on R2's FastEthernet segment, you may see "Null" for the outgoing interface list in the output of show ip mroute.

In the testing / verification shown, SW1 and SW4 had PIM modes configured. Our section did explicitly state to not add PIM on additional interfaces for the traffic to pass, so make sure that you remove the PIM statements from the interfaces on SW1 and SW4.

Task 5.2

R4 and R5:
```
ip multicast-routing
!
interface Loopback1
 ip address 150.1.0.255 255.255.255.255
 ip pim sparse-mode
!
interface FastEthernet0/0
 ip pim sparse-mode
!
interface FastEthernet0/1
 ip pim sparse-mode
!
router ospf 1
 network 150.1.0.255 0.0.0.0 area 0
!
ip pim rp-address 150.1.0.255
```

R4:
```
ip msdp peer 150.1.5.5 connect-source Loopback0
```

R5:
```
ip msdp peer 150.1.4.4 connect-source Loopback0
```

R6:
```
ip multicast-routing
!
ip pim rp-address 150.1.0.255
!
interface FastEthernet0/0
 ip pim sparse-mode
```

SW2:
```
ip multicast-routing distributed
!
ip pim rp-address 150.1.0.255
!
interface Vlan58
 ip pim sparse-mode
```


Further Reading

Anycast RP

Task 5.3 Verification

```
Rack1R6#show ip pim rp map
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 150.1.0.255 (?)

Rack1R4#show ip msdp peer
MSDP Peer 150.1.5.5 (?), AS 100
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback0 (150.1.4.4)
    Uptime(Downtime): 00:00:40, Messages sent/received: 3/3
    Output messages discarded: 0
    Connection and counters cleared 00:01:40 ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 2
  Input queue size: 0, Output queue size: 0
Rack1R4#

Rack1R5#show ip msdp peer
MSDP Peer 150.1.4.4 (?), AS 100
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback0 (150.1.5.5)
    Uptime(Downtime): 00:00:58, Messages sent/received: 3/4
    Output messages discarded: 0
    Connection and counters cleared 00:01:46 ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 2
  Input queue size: 0, Output queue size: 0
Rack1R5#
```

For testing purposes, we will have R6's Loopback0 join multicast group 226.26.26.26

```
R6:
interface Loopback0
 ip address 150.1.6.6 255.255.255.0
 ip igmp join-group 226.26.26.26
 ip pim sparse-mode

Rack1SW2#ping 226.26.26.26

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 226.26.26.26, timeout is 2 seconds:

Reply to request 0 from 129.1.46.6, 9 ms
Rack1SW2#

Rack1R4#show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
(129.1.58.8, 226.26.26.26), RP 150.1.0.255, BGP/AS 0,
00:00:12/00:05:47, Peer 150.1.5.5
Learned from peer 150.1.5.5, RPF peer 150.1.5.5,
SAs received: 1, Encapsulated data received: 1
Rack1R4#
```

Task 6.1

```
R6:
access-list 100 permit tcp host 129.1.46.100 any eq telnet
access-list 100 deny   tcp any any eq telnet log
!
line vty 0 4
 access-class 100 in
```

Task 6.1 Verification

```
Rack1R6#telnet 150.1.6.6
Trying 150.1.6.6 ...
% Connection refused by remote host

Rack1R6#
%SEC-6-IPACCESSLOGP: list 100 denied tcp 150.1.6.6(14768) ->
0.0.0.0(23), 1 packet
```

Task 6.2

```
R2:
access-list 22 permit 129.1.0.0 0.0.255.255
!
login block-for 300 attempts 10 within 60
login quiet-mode access-class 22

username cisco password cisco
login on-failure
line vty 0 181
```

```
login local
```

Task 6.2 Verification

```
Rack1R2#show login
     A default login delay of 1 seconds is applied.
     Quiet-Mode access list 22 is applied.

     Router enabled to watch for login Attacks.
     If more than 9 login failures occur in 60 seconds or less,
     logins will be disabled for 300 seconds.

     Router presently in Normal-Mode.
     Current Watch Window
         Time remaining: 54 seconds.
         Login failures for current window: 0.
     Total login failures: 0.
```

For testing, start with lower values for attempts:

```
Rack1R2(config)#login block 300 attempts 2 within 600
Rack1R2(config)#end
Rack1R2#telnet 150.1.2.2 /source lo0
Trying 150.1.2.2 ... Open


User Access Verification

Username: c
Password:
% Login invalid

%SEC_LOGIN-4-LOGIN_FAILED: Login failed [user: c] [Source: 150.1.2.2]
[localport: 23] [Reason: Login Authentication Failed - BadUser] at
23:51:09 PST Sat Mar 2 2002

Username: c
Password:
% Login invalid

%SEC_LOGIN-4-LOGIN_FAILED: Login failed [user: c] [Source: 150.1.2.2]
[localport: 23] [Reason: Login Authentication Failed - BadUser] at
23:51:19 PST Sat Mar 2 2002
%SEC_LOGIN-1-QUIET_MODE_ON: Still timeleft for watching failures is 574
secs, [user: c] [Source: 150.1.2.2] [localport: 23] [Reason: Login
Authentication Failed - BadUser] [ACL: 22] at 23:51:19 PST Sat Mar 2
2002
[Connection to 150.1.2.2 closed by foreign host]
Rack1R2#
```

Adjust to the values specified in the section:

```
Rack1R2#show login
```

```
A default login delay of 1 seconds is applied.
Quiet-Mode access list 22 is applied.

Router enabled to watch for login Attacks.
If more than 9 login failures occur in 60 seconds or less,
logins will be disabled for 300 seconds.

Router presently in Normal-Mode.
Current Watch Window
    Time remaining: 54 seconds.
    Login failures for current window: 0.
Total login failures: 0.
```

Task 6.2 Breakdown

Security enhancements allow conditional blocking to prevent the router from being impacted by a denial of service or brute force attack. The login block-for command allows you to set a threshold time period, such that if a certain number of failed attempts are received, access will be blocked. The quiet-mode ACL allows you to specify which hosts are allowed to access the device, even if the block threshold is exceeded. The login on-failure command, although not mandated by the section, will allow you to see the failed attempts logged locally. If you are just using a password on the line, it will not trigger the feature, so username and password are configured, along with login local on the VTY lines.

Further Reading

[Cisco IOS Login Enhancements]()

Task 7.1

```
R6:
logging host ipv6 2001:CC1E:1:1::100
!
ip access-list log-update threshold 10
```

Task 4.1 Breakdown

This task is very straightforward. Configure the logging destination and adjust the threshold. Make sure that your logging level is informational or debugging in order to get hits for ACL entries. By default, the logging severity level is high enough, but it is possible that a lower level could have been set in the initial configuration.

```
Rack1R6#show logging | beg Trap
    Trap logging: level informational, 95 message lines logged
        Logging to 2001:CC1E:1:1::100  (udp port 514,  audit disabled,
              authentication disabled, encryption disabled, link up),
              4 message lines logged,
```


Task 7.2

```
R1, R2, SW1:
ntp server 150.1.3.3

R3, R6:
ntp master

R4, R5, SW2, SW3, SW4:
ntp server 150.1.6.6

R1, R2, R3, SW1:
clock timezone PST -8
clock summer-time PDT recurring

R4, R5, R6, SW2:
clock timezone CST -6
clock summer-time CDT recurring

SW3 and SW4:
ntp server 150.1.6.6
```

> ☞ **Quick Note**
> The actual NTP server that SW3 and SW4 point it is irrelevant for this task

Task 7.2 Verification

Verify that the clocks are synchronized. For instance on R1:

```
Rack1R1#show ntp status
Clock is synchronized, stratum 9, reference is 150.1.3.3
nominal freq is 249.5901 Hz, actual freq is 249.5902 Hz, precision is
2**18
reference time is CCF5C2A7.03975C21 (06:50:15.014 UTC Fri Dec 19 2008)
clock offset is -1.2667 msec, root delay is 25.18 msec
root dispersion is 1.74 msec, peer dispersion is 0.43 msec
```

R6 is in Chicago (UTC -6), while R2 is in Reno (UTC -8):

```
Rack1R6#show clock
00:55:36.888 CST Fri Dec 19 2009

Rack1R6#show ntp status
Clock is synchronized, stratum 8, reference is 127.127.7.1
nominal freq is 249.5901 Hz, actual freq is 249.5901 Hz, precision is
2**18
reference time is CCF5C3E7.59B407B2 (00:55:35.350 CST Fri Dec 19 2009)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.02 msec, peer dispersion is 0.02 msec

Rack1R2#show clock
22:56:45.523 PST Thu Dec 18 2009
```

```
Rack1R2#show clock
.23:02:54.691 PST Thu Dec 18 2009

Rack1R2#show ntp status
Clock is unsynchronized, stratum 16, no reference clock
nominal freq is 249.5901 Hz, actual freq is 249.5901 Hz, precision is
2**18
reference time is CCF5C583.0522C1A8 (23:02:27.020 PST Thu Dec 18 2009)
clock offset is -774.5739 msec, root delay is 24.67 msec
root dispersion is 8649.80 msec, peer dispersion is 16000.00 msec


Rack1SW3#show version | include started
System restarted at 01:09:16 UTC Sun Jan 15 2010
Rack1SW3#
```

Note

When NTP is configured, the device will also timestamp the last configuration change and the last time the configuration was saved to NVRAM in the configuration itself.

```
Rack1SW3#show running-config | include Last|NVRAM
! Last configuration change at 08:00:33 UTC Sun Jan 15 2010
! NVRAM config last updated at 08:06:55 UTC Sun Jan 15 2010
```

Task 7.2 Breakdown

NTP advertisements are always sent in Coordinated Universal Time (UTC), also commonly known as Greenwich Mean Time (GMT). In order to avoid log inconsistencies due to devices being located in different time zones, it is common practice to leave the local time in UTC. However, the time zone of the router's local clock can be adjusted by issuing the **clock timezone [*timezone*] [*offset*]** global configuration command. Additionally, daylight savings time can be configured with the **clock summer-time [*daylight timezone*] recurring** command. Time zone configuration is always locally significant, and is never propagated via NTP.

Task 7.3

```
R1, R2, SW1:
ip domain-lookup
ip name-server 150.1.3.3

R3:
ip dns server
ip domain-lookup
!
ip host Rack1R1 150.1.1.1
ip host Rack1R2 150.1.2.2
ip host Rack1R3 150.1.3.3
ip host Rack1SW1 150.1.7.7
```

Task 7.3 Verification

Verify the new domain server:

```
Rack1R1#ping Rack1R2
Translating "Rack1R2"...domain server (150.1.3.3)

Translating "Rack1R2"...domain server (150.1.3.3) [OK]

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

Rack1R1#ping Rack1SW1
Translating "Rack1SW1"...domain server (150.1.3.3)

Translating "Rack1SW1"...domain server (150.1.3.3) [OK]

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 7.4

```
R4:
interface FastEthernet0/0
 glbp 1 ip 129.1.45.6
 glbp 1 preempt
 glbp 1 weighting 30
 glbp 1 load-balancing weighted

R5:
interface FastEthernet0/1
 glbp 1 ip 129.1.45.6
 glbp 1 priority 50
 glbp 1 preempt
 glbp 1 weighting 70
 glbp 1 load-balancing weighted
```

Task 7.4 Verification

```
Rack1R4#show glbp
FastEthernet0/0 - Group 1
  State is Active
    2 state changes, last state change 00:00:57
  Virtual IP address is 129.1.45.6
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.367 secs
  Redirect time 600 sec, forwarder time-out 14400 sec
  Preemption enabled, min delay 0 sec
  Active is local
  Standby is 129.1.45.5, priority 50 (expires in 8.361 sec)
  Priority 100 (default)
  Weighting 30 (configured 30), thresholds: lower 1, upper 30
  Load balancing: weighted
  Group members:
    000f.90fa.ed60 (129.1.45.4) local
    000f.90fb.0a21 (129.1.45.5)
  There are 2 forwarders (1 active)
  Forwarder 1
    State is Active
      1 state change, last state change 00:00:47
    MAC address is 0007.b400.0101 (default)
    Owner ID is 000f.90fa.ed60
    Redirection enabled
    Preemption enabled, min delay 30 sec
    Active is local, weighting 30
  Forwarder 2
    State is Listen
    MAC address is 0007.b400.0102 (learnt)
    Owner ID is 000f.90fb.0a21
    Redirection enabled, 597.572 sec remaining (maximum 600 sec)
    Time to live: 14397.572 sec (maximum 14400 sec)
    Preemption enabled, min delay 30 sec
```

```
  Active is 129.1.45.5 (primary), weighting 70 (expires in 7.568 sec)
```

Task 8.1

```
R2:
interface Serial0/0
 frame-relay traffic-shaping
 frame-relay class DLCI_204
!
map-class frame-relay DLCI_204
 frame-relay cir 512000
 frame-relay bc 5120
 frame-relay be 0
 frame-relay fragment 640

R4:
interface Serial0/0/0
 frame-relay traffic-shaping
!
interface Serial0/0/0.124 multipoint
 frame-relay interface-dlci 401
  class DLCI_401
 frame-relay interface-dlci 402
  class DLCI_402
!
interface Serial0/0/0.54 point-to-point
 frame-relay interface-dlci 405
  class EEK
!
map-class frame-relay EEK
 frame-relay cir 512000
 frame-relay bc 5120
 frame-relay be 0
 frame-relay fragment 640
!
map-class frame-relay DLCI_401
 frame-relay cir 512000
 frame-relay bc 5120
 frame-relay be 0
 frame-relay fragment 640
!
map-class frame-relay DLCI_402
 frame-relay cir 512000
 frame-relay bc 5120
 frame-relay be 0
 frame-relay fragment 640
```

> ☞ **Quick Note**
> Previously applied.

> ☞ **Quick Note**
> Previously applied.

Task 8.1 Breakdown

The smaller the Frame Relay Traffic Shaping interval (Tc), the less time traffic is delayed in the output queue as it is waiting to exit to the transmit ring. This in turn equates to less delay, and better performance, for low bandwidth delay sensitive traffic such as VoIP. However, lowering the shaping interval does not accomplish anything when the MTU of a packet exceeds the Bc value.

Suppose that the MTU of the interface is 1500 bytes, and that in each Tc the FRTS algorithm has allotted 5120 bits of committed burst. This means that it will take a minimum of three intervals (30ms in this case) in order to clock this packet onto the interface. Depending on the serialization delay of the interface (dependent on the hardware clocking speed), this delay in sending the packet can result in unacceptable delay for real time traffic, even if it is prioritized. This is due to the fact that even if a packet is in the low latency queue, it must wait for whatever packet is on the transmit ring to exit the interface.

In order to further reduce the delay of real time traffic as it exits the output queue, Frame Relay fragmentation can be used to reduce the MTU of packets transmitted out the interface. By reducing the maximum fragment size to Bc (in bytes), a real time packet such as VoIP is guaranteed that the worst case scenario delay that will be incurred in the output queue is one single Tc (10ms in this case).

Task 8.1 Verification

```
Rack1R4#show frame-relay pvc 402

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

DLCI = 402, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial0/0/0.124

  input pkts 716           output pkts 758          in bytes 133624
  out bytes 128601         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 303       out bcast bytes 97464
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 05:13:08, last time pvc status changed 01:17:53
  Queueing strategy: weighted fair
  Current fair queue configuration:
   Discard     Dynamic      Reserved
   threshold   queue count  queue count
    64          16           0
  Output queue size 0/max total 600/drops 0
  fragment type end-to-end fragment size 640
  cir 512000    bc   5120      be 0         limit 640    interval 10
  mincir 256000    byte increment 640   BECN response no  IF_CONG no
  frags 5         bytes 653       frags delayed 0       bytes delayed 0
  shaping inactive
```

```
  traffic shaping drops 0

Rack1R2#show frame-relay pvc 204

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

DLCI = 204, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial0/0

  input pkts 644           output pkts 600          in bytes 94568
  out bytes 96298          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 196       out bcast bytes 69702
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 03:16:45, last time pvc status changed 01:18:42
  Queueing strategy: weighted fair
  Current fair queue configuration:
   Discard     Dynamic      Reserved
   threshold   queue count  queue count
    64          16           0
  Output queue size 0/max total 600/drops 0
  fragment type end-to-end fragment size 640
  cir 512000    bc   5120      be 0         limit 640    interval 10
  mincir 256000    byte increment 640   BECN response no  IF_CONG no
  frags 16        bytes 2152      frags delayed 0       bytes delayed 0
  shaping inactive
  traffic shaping drops 0
```

Task 8.2

```
R2:
class-map match-all VoIP
  match access-group name VoIP
!
 policy-map LLQ
  class VoIP
   priority 192
!
map-class frame-relay DLCI_204
 service-policy output LLQ
!
ip access-list extended VoIP
 permit udp any 129.1.46.0 0.0.0.255 range 16384 32767
```

```
R4:
class-map match-all VoIP
  match access-group name VoIP
!
 policy-map LLQ
  class VoIP
   priority 192
!
map-class frame-relay DLCI_402
 service-policy output LLQ
!
ip access-list extended VOIP
 permit udp 129.1.46.0 0.0.0.255 any range 16384 32767
```

Task 8.2 Breakdown

By putting VoIP traffic in the low latency queue by using the **priority** keyword under the MQC policy-map, VoIP traffic is always guaranteed to be dequeued first on the Frame Relay circuit between R2 and R4 up to 192Kbps. When VoIP traffic exceeds 192Kbps of the output queue, it is not guaranteed low latency, but may be transmitted. When VoIP traffic exceeds 192Kbps of the output queue, and there is congestion in the queue, VoIP in excess of 192Kbps will be dropped.

Task 8.2 Verification

```
Rack1R4#show frame-relay pvc 402

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

DLCI = 402, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial0/0/0.124

  input pkts 731           output pkts 769          in bytes 135652
  out bytes 130340         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 306       out bcast bytes 98574
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 05:15:48, last time pvc status changed 01:20:34
  service policy LLQ
 Serial0/0/0.124: DLCI 402 -

  Service-policy output: LLQ

    Class-map: VoIP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name VoIP
      Queueing
        Strict Priority
        Output Queue: Conversation 40
        Bandwidth 192 (kbps) Burst 4800 (Bytes)
        (pkts matched/bytes matched) 0/0
        (total drops/bytes drops) 0/0

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
  Output queue size 0/max total 600/drops 0
  fragment type end-to-end fragment size 640
  cir 512000    bc   5120      be 0         limit 640    interval 10
  mincir 256000    byte increment 640   BECN response no  IF_CONG no
  frags 16        bytes 2392      frags delayed 0       bytes delayed 0
  shaping inactive
  traffic shaping drops 0

Rack1R2#show frame-relay pvc 204

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

DLCI = 204, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE =
Serial0/0

  input pkts 658           output pkts 618          in bytes 96546
  out bytes 98834          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 200      out bcast bytes 71306
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  pvc create time 03:20:06, last time pvc status changed 01:22:03
  service policy LLQ
 Serial0/0: DLCI 204 -

  Service-policy output: LLQ

    Class-map: VoIP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name VoIP
      Queueing
        Strict Priority
        Output Queue: Conversation 40
        Bandwidth 192 (kbps) Burst 4800 (Bytes)
        (pkts matched/bytes matched) 0/0
        (total drops/bytes drops) 0/0

    Class-map: class-default (match-any)
      13 packets, 1860 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
  Output queue size 0/max total 600/drops 0
  fragment type end-to-end fragment size 640
  cir 512000    bc   5120      be 0         limit 640    interval 10
  mincir 256000    byte increment 640   BECN response no  IF_CONG no
  frags 34        bytes 4688      frags delayed 0       bytes delayed 0
  shaping inactive
  traffic shaping drops 0
```

Task 8.3

SW3:

```
mls qos
!
ip access-list extended HTTP_REPLIES
 permit tcp any eq 80 any
!
ip access-list extended SMTP_REPLIES
 permit tcp any eq 25 any
!
class-map HTTP_REPLIES
 match access-group name HTTP_REPLIES
!
class-map SMTP_REPLIES
 match access-group name SMTP_REPLIES
!
mls qos aggregate-policer POLICE_2M 2000000 128000 exceed-action drop
!
policy-map MARK_AND_POLICE
 class HTTP_REPLIES
  set dscp af21
```



```
  police aggregate POLICE_2M
 class SMTP_REPLIES
  set dscp af23
  police aggregate POLICE_2M
!
interface FastEthernet 0/5
 service-policy input MARK_AND_POLICE
```

Task 8.3 Verification

```
Rack1SW3#show policy-map interface fastEthernet 0/5
 FastEthernet0/5

  Service-policy input: MARK_AND_POLICE

    Class-map: HTTP_REPLIES (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name HTTP_REPLIES

    Class-map: SMTP_REPLIES (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name SMTP_REPLIES

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

Rack1SW3#show mls qos aggregate-policer
aggregate-policer POLICE_2M 2000000 128000 exceed-action drop
Used by policy map MARK_AND_POLICE

Rack1SW3#show mls qos interface FastEthernet 0/5 ?
  buffers     Show buffer information
  policers    Show policers information
  queueing    Show queueing information
  statistics  Show statistics
  |           Output modifiers
  <cr>
```

```
Rack1SW3#show mls qos interface FastEthernet 0/5 policers
FastEthernet0/5
policymap=MARK_AND_POLICE
type=Shared, id=0 name=POLICE_2M

Rack1SW3#show mls qos interface FastEthernet 0/5 statistics
FastEthernet0/5
Ingress
  dscp: incoming   no_change  classified policed    dropped (in bytes)
Others: 1165       0          1165       0          0
Egress
  dscp: incoming   no_change  classified policed    dropped (in bytes)
Others: 2436          n/a       n/a      0          0
```