


**NETMASTERCLASS**  
**ROUTING AND SWITCHING CCIE® TRACK**

# DOIT-200v6

# VOLUME II



## Scenario 14 ANSWER KEY

FOR

CCIE® CANDIDATES

## Disclaimer

NetMasterClass, LLC is an independent training and consulting company based in Herndon, Virginia. The terms "Cisco", "Cisco Systems" and "CCIE" are the trademarks of Cisco Systems, Inc. NetMasterClass, LLC is Cisco Learning Partner.

## Cisco Non-Disclosure Agreement Compliance

All products and services offered by NetMasterClass, LLC are in full compliance with the Cisco CCIE Lab non-disclosure agreement. The content of the NetMasterClass CCIE preparation materials is based upon the NetMasterClass "issue spotting and analysis" internetwork training methods.

***NOTE: To use this document to its maximum effectiveness, access the SHOWiT engine while reviewing each section of this Answer Key.***


## DOiT-V6 Scenario 14: Spot the Issue Answer Key

### Table of Contents

14.1	Frame Relay .....	6
14.2	Catalyst Configuration .....	7
14.3	OSPF .....	8
14.4	RIP .....	11
14.5	EIGRP.....	13
14.6	BGP .....	17
14.7	Router Maintenance .....	20
14.8	Security.....	21
14.9	IPv6.....	21
14.10	Catalyst Specialties .....	28
14.11	Address Administration.....	29
14.12	Multicast .....	31
14.13	Quality of Service .....	34
14.14	IOS Features .....	40



**REGARDLESS OF ANY CONFIGURATION YOU PERFORM IN THIS EXAM, IT IS VERY IMPORTANT TO CONFORM TO THE GENERAL GUIDELINES PROVIDED BELOW. IF YOU DO NOT CONFORM TO THEM, THIS CAN RESULT IN A SIGNIFICANT DEDUCTION OF POINTS IN YOUR FINAL EXAM SCORE.**



## Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16.
- Do not rely on dynamic Frame-Relay Inverse ARP.
- Do not introduce any new IP addresses.
- Do not use any static routes.
- Make sure all IPv4 and IPv6 Loopback interfaces are advertised with their original mask, unless noted otherwise.
- Make sure all IP interfaces in the diagram are reachable within this internetwork.
- Use conventional routing algorithms.
- In this exercise, R5 and FRS are used for backbone router simulation.
- In this exercise, it's not required to ping R5 and FRS originated networks 192.\*.\* and 140.\*.\* as well as the 172.16.1.2 address on FRS.

### ***Explanation of Each of the Goals and Restrictions:***

#### **IP subnets in the Scenario diagram belong to network 172.16.0.0/16**

The third and fourth octets of the IP addresses displayed on the diagram belong to 172.16.0.0/16.

#### **Do not rely on dynamic Frame-Relay Inverse ARP.**

This requirement forces you to fulfill your Frame-Relay inverse arp requirements with Frame-Relay map statements. Think of a Frame-Relay map statement as the equivalent of a static inverse arp entry.

#### **Do not use any static routes.**

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

#### **Make sure all IPv4 and IPv6 Loopback interfaces are advertised with their original mask, unless noted otherwise.**

This requirement is primarily for the OSPF advertised loopbacks. Use "ip ospf network point-to-point" under the loopback interface. Otherwise, the loopback will be advertised as a /32 host entry by default.

#### **Make sure all IP interfaces in the diagram are *reachable* within this internetwork. DO NOT FORGET THIS!**

This is a key goal to observe. This requires that all of your IGP's are configured properly. Also, all of your routing policy tasks must be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using distribute-lists, route-maps and the distance command. A key point to remember about this exam is: the term "redistribution" is never explicitly used in

this exam. However, you must perform redistribution in order to assure that all IP addresses are reachable without the use of static routes.

### Use conventional routing algorithms.

This restriction prevents you from solving any problems by configuring policy routing. At the heart of this restriction is the interpretation of the “conventional routing algorithms”. Although this phrase can be interpreted in a number of different ways, the interpretation applied in this workbook is:

CONVENTIONAL ROUTING ALGORITHMS ARE ROUTING ALGORITHMS THAT APPLY DESTINATION BASED PREFIX LOOKUPS IN A ROUTING TABLE. CONVENTIONAL ROUTING ALGORITHMS DO NOT USE ANY OTHER TYPE OF INFORMATION OTHER THAN THE DESTINATION ADDRESS TO MAKE A PACKET FORWARDING DECISION.

Due to this restrictive interpretation, no form of policy routing can be applied. Whenever you see this restriction, you will need to use dynamic routing protocols to fulfill all packet forwarding requirements

### The following IOS versions were used on the devices:

Device	IOS version
R1	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
R2	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
R3	IOS (tm) C2600 Software (C2600-J1S3-M), Version 12.3(15a)
R4	IOS (tm) C2600 Software (C2600-J1S3-M), Version 12.3(15a)
R5	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
R6	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
FRS	IOS (tm) 2500 Software (C2500-JS-L), Version 12.2(27)
CAT1	IOS (tm) C3550 Software (C3550-I5Q3L2-M), Version 12.2(25)SEA
CAT2	IOS (tm) C3550 Software (C3550-I5Q3L2-M), Version 12.2(25)SEA

## 14.1 Frame Relay



### HIDDEN ISSUES TO SPOT WITH FRAME RELAY

**Issue:** Make sure only PVCs listed on the diagram are used for user traffic.

**Solution:**

When you examine the diagram, you notice that the Scenario is based upon a hub and spoke Frame-Relay topology. In order to fulfill the requirements of the task stated above – “Make sure only PVC’s listed on the diagram are used for user traffic – disable Frame-Relay inverse arp and statically map only the DLCI’s referenced in the diagram.

**Verification:**

Use **show frame-relay map** command to verify attached PVCs:

```
R1# sh frame map
Serial0/0.124 (up): ipv6 FE80::2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ipv6 FE80::4 dlci 104(0x68,0x1880), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ipv6 FEC0::124:2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ipv6 FEC0::124:4 dlci 104(0x68,0x1880), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ip 172.16.124.1 dlci 102(0x66,0x1860), static,
CISCO, status defined, active
Serial0/0.124 (up): ip 172.16.124.2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ip 172.16.124.4 dlci 104(0x68,0x1880), static,
broadcast,
CISCO, status defined, active
Serial0/0.13 (up): point-to-point dlci, dlci 103(0x67,0x1870), broadcast, CISCO
status defined, active
```

**Note:** The table above is rather large because it contains both IPv4 and IPv6 entries.

```
R2# sh frame map
Serial0/0.124 (up): ipv6 FE80::4 dlci 201(0xC9,0x3090), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ipv6 FEC0::124:1 dlci 201(0xC9,0x3090), static,
broadcast,
CISCO, status defined, active
Serial0/0.124 (up): ipv6 FEC0::124:4 dlci 201(0xC9,0x3090), static,
CISCO, status defined, active
Serial0/0.124 (up): ip 172.16.124.1 dlci 201(0xC9,0x3090), static,
broadcast,
```

```

CISCO, status defined, active
Serial0/0.124 (up): ip 172.16.124.2 dlci 201(0xC9,0x3090), static,
CISCO, status defined, active
Serial0/0.124 (up): ip 172.16.124.4 dlci 201(0xC9,0x3090), static,
CISCO, status defined, active
Serial0/0.124 (up): ipv6 FE80::1 dlci 201(0xC9,0x3090), static,
broadcast,
CISCO, status defined, active

```



**To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as “show all”.**

## 14.2 Catalyst Configuration



### HIDDEN ISSUES TO SPOT WITH CATALYST CONFIGURATION

**Issue:** Assign IP addresses from the 172.16.1.0/24 subnet to R1, FRS and the CAT1 switched virtual interface.

**Solution:**

The core issue involved with this configuration requirement is: how to make devices attached to two separate VLAN's appear on the same IP subnet. R1 and FRS are assigned to VLAN 10 while R1 and a CAT1 Switched Virtual Interface are assigned to VLAN 20. This requirement can be fulfilled by configuring IRB on router R1. Assign both of the FastEthernet subinterfaces on R1 to a bridge-group. Then enable IRB and create a BVI on R1 using the 172.16.1.0/24 subnet. Do not assign an IP address to the FastEthernet subinterfaces. Finally, do not forget to enter the global configuration command “bridge 1 route ip” on R1 to allow bridged IP packets to get forwarded to the BVI.

**Verification:**

Verify that bridge is active on R1 and is forwarding by using **show bridge** command:

```

R1# sh bridge

Total of 300 station blocks, 298 free
Codes: P - permanent, S - self

Bridge Group 1:

    Address          Action  Interface    Age  RX count  TX count
0009.e8f4.0e00     forward Fa0/0.10      0    182307    6
0050.5480.66e1     forward Fa0/0.20      0    462751   40

```

Use show spanning-tree command to see additional parameters of spanning tree:

R1# sh spanning-tree

```
Bridge group 1 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, address 00d0.5895.c8a1
Configured hello time 2, max age 20, forward delay 15
We are the root of the spanning tree
Topology change flag not set, detected flag not set
Number of topology changes 1 last change occurred 4d21h ago
    from FastEthernet0/0.10
Times: hold 1, topology change 35, notification 2
    hello 2, max age 20, forward delay 15
Timers: hello 1, topology change 0, notification 0, aging 300
```

Port 11 (FastEthernet0/0.10) of Bridge group 1 is forwarding

```
Port path cost 19, Port priority 128, Port Identifier 128.11.
Designated root has priority 32768, address 00d0.5895.c8a1
Designated bridge has priority 32768, address 00d0.5895.c8a1
Designated port id is 128.11, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 2
BPDU: sent 210817, received 5
```

Port 12 (FastEthernet0/0.20) of Bridge group 1 is forwarding

```
Port path cost 19, Port priority 128, Port Identifier 128.12.
Designated root has priority 32768, address 00d0.5895.c8a1
Designated bridge has priority 32768, address 00d0.5895.c8a1
Designated port id is 128.12, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 2
BPDU: sent 210815, received 3
```



To obtain a comprehensive view of the configuration tasks in this section, access the **SHOWIT engine**. With the **SHOWIT engine**, you can enter in over 1000 IOS commands as well a collection of **NMC proprietary commands** such as “show all”.

## 14.3 OSPF



### HIDDEN ISSUES TO SPOT WITH OSPF

**Issue:** On R3, assign a loopback interface with the IP address 172.16.60.1/28 to OSPF area 44.

**Solution:**

R3 does not possess a direct connection to area 0. It possesses a connection to Area 33. If you assign an interface on R3 to an area other than 33, you will need to configure a virtual-link with Area 33 as the transit area.

**Verification:**

Issue **show ip ospf virtual-links** command:



```
R3# show ip ospf virtual
Virtual Link OSPF_VL0 to router 172.16.101.1 is up
Run as demand circuit
DoNotAge LSA allowed.
Transit area 33, via interface Serial0/0, Cost of using 64
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Index 1/2, 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
```

**Issue: Use the OSPF network type non-broadcast for all Frame Relay connections**

**Solution:**

Remember that the 172.16.124.0/24 subnet is configured on a Frame-Relay hub and spoke topology. Since all OSPF packets have a TTL =1, OSPF spoke routers will never communicate with other spoke routers. Therefore, no spoke routers can become either a DR or BDR. To assure that this never happens, set their OSPF priority to 0 on the spoke routers R2 and R4 at the Frame-Relay interface level. At hub router R1, enter in two neighbor statements, one for R2 and the second for R4.

**Verification:**

```
R1# show ip ospf inte s0/0.124
Serial0/0.124 is up, line protocol is up
Internet Address 172.16.124.1/24, Area 0
Process ID 100, Router ID 172.16.101.1, Network Type NON_BROADCAST, Cost: 64
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 172.16.101.1, Interface address 172.16.124.1
No backup designated router on this network
Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
oob-resync timeout 120
Hello due in 00:00:00
Index 4/6, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 4, maximum is 8
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 2, Adjacent neighbor count is 2
Adjacent with neighbor 172.16.102.1
Adjacent with neighbor 172.16.104.1
Suppress hello for 0 neighbor(s)
```

```
R1# show ip ospf interface s0/0.13
Serial0/0.13 is up, line protocol is up
Internet Address 172.16.13.1/24, Area 33
Process ID 100, Router ID 172.16.101.1, Network Type NON_BROADCAST, Cost: 64
Transmit Delay is 1 sec, State DROTHER, Priority 0
Designated Router (ID) 172.16.103.1, Interface address 172.16.13.3
No backup designated router on this network
Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
oob-resync timeout 120
Hello due in 00:00:08
Index 1/4, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 3, maximum is 11
```

```
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 172.16.103.1 (Designated Router)
Suppress hello for 0 neighbor(s)
```

**Issue: On R3, have the VLAN30 subnet advertised via OSPF without including it as one of your OSPF networks.**

**Solution:**

Configure redistribute connected on R3 to inject the VLAN30 subnet into OSPF without using an OSPF network command. Whenever you configure “redistribute connected”, remember to apply either a distribute-list or route-map to inject only the connected prefix (es) you intended to redistribute. Without such a filter, you will redistribute ALL connected networks.

**Verification:**

Issue **show ip ospf database | be Ext** command:

```
R3# show ip ospf database | be Ext
Type-5 AS External Link States

Link ID        ADV Router    Age           Seq#           Checksum Tag
172.16.1.0     172.16.101.1 1702          0x800000D4    0x00519A 0
172.16.26.0    172.16.102.1 198           0x800000D3    0x00796C 0
172.16.26.0    172.16.104.1 170           0x800000D3    0x006B78 0
172.16.30.0    172.16.102.1 199           0x800000D3    0x004D94 0
172.16.30.0    172.16.104.1 171           0x800000D3    0x003FA0 0
172.16.31.0    172.16.103.1 1789          0x800000D2    0x00FBD1 0
172.16.50.0    172.16.103.1 1789          0x800000D2    0x00A694 0
172.16.102.0   172.16.102.1 199           0x800000D3    0x003267 0
172.16.102.0   172.16.104.1 171           0x800000D3    0x002473 0
172.16.103.0   172.16.103.1 1789          0x800000D2    0x00E0A4 0
172.16.104.0   172.16.102.1 199           0x800000D3    0x001C7B 0
172.16.104.0   172.16.104.1 171           0x800000D3    0x000E87 0
172.16.106.0   172.16.102.1 199           0x800000D3    0x00068F 0
172.16.106.0   172.16.104.1 171           0x800000D3    0x00F79B 0
172.16.110.0   172.16.101.1 1703          0x800000D2    0x00A1DE 0
172.16.120.0   172.16.103.1 1789          0x800000D2    0x00254F 0
192.168.2.0    172.16.101.1 1703          0x800000D2    0x001E22 0
192.168.3.0    172.16.101.1 1704          0x800000D2    0x00132C 0
```

As you see, 35.3/24 is not in this list, because it is filtered by redistribution command.

**Issue: Ensure that there is only one entry for the loopback /28 subnet and that the loopback subnet can be seen by all other routers**

**Solution:**

When you configure redistribute connected, you might inject into a routing process many more routes than you intended. Therefore, it is a recommended general practice, to always configure either a distribute-list or a route-map to explicitly limit what you want to inject with redistribute connected. Emphasis in this issue

is placed on the /28 address assigned to Area 44 on router R3. This subnet is assigned to a loopback interface. By default, loopback interfaces are advertised by OSPF as /32 routes. If you performed redistribute connected on R3 and did not filter out this /28 loopback, it is possible that you would end up with a 32/ IA route and a /28 E1 route for the same prefix. The task stated above explicitly states that only ONE entry should be seen. You can assure that only one entry is seen by all other routers by either configuring a complementary distribute-list or route-map for the redistribute connected subnet command under OSPF. Use either the distribute-list or route-map to filter out the /28 subnet from being injected into OSPF as an External route. An alternative solution is to enter the following command under the loopback interface where the /28 address is assigned : "ip ospf network point-to-point". When this command is entered, the interface no longer is advertised as a /32. It will be advertised with its native mask. In this case, that is /28. If OSPF learns the same prefix from two different sources, it will select an intra-area route over an inter-area route and an inter-area route over an external route. Therefore, OSPF routers will be the IA /28 route over the E1 /28 route. This would fulfill the "single entry" requirement for the /28 prefix.



To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".

## 14.4 RIP



### HIDDEN ISSUES TO SPOT WITH RIP

**Issue:** Do not broadcast/multicast RIP updates on VLAN 40 for security reasons.

**Solution:**

If you are instructed neither broadcast nor multicast RIP updates, configure RIP to unicast its updates. This is accomplished by making the interface passive and then configuring neighbor statements for every device that needs to receive the RIP updates. In this Scenario, it is recommended to configure two neighbor statements on each VLAN 40 RIP speaking router.

**Verification:**

Enable debugging for ip packets on R2 (for example), you will see that RIP sends unicast updates on VLAN 40:

```
00:45:18: IP: s=172.16.26.2 (local), d=172.16.26.4 (FastEthernet0/0), len 312, sending
00:45:18:   UDP src=520, dst=520
00:45:18: IP: s=172.16.26.2 (local), d=172.16.26.6 (FastEthernet0/0), len 312, sending
00:45:18:   UDP src=520, dst=520
```

**Issue: On R2 and R4, configure mutual redistribution between OSPF and RIP.**

**Solution:**

This configuration task is loaded with danger and instability! You are instructed to perform mutual redistribution between RIP and OSPF at two redistribution points. Such a configuration requirement is pregnant with routing instabilities.

Redistribution between RIP and OSPF will be performed under route-maps RIP→OSPF and OSPF→RIP. Also, RIP native prefixes will be protected with administrative distance 109, which is 1 less than default administrative distance of OSPF. This will make sure that OSPF will not push RIP routes out of routing table.

**Issue: Send updates to 224.0.0.9 only on the link between R3 and CAT2**

**Solution:**

If you are going to send RIP updates using the multicast address of 224.0.0.9, you need to configure RIP version 2. Therefore, configure RIP version 2 between R3 and CAT2.

```
interface FastEthernet0/3
description R3 Fa0/0
no switchport
ip address 172.16.31.20 255.255.255.0
ip access-group 100 out
ip rip send version 2
```

**Verification:**

Enable debugging of IP packets on CAT2 and make sure that RIP sends updates on required multicast address:

```
01:20:58: datagramsize=390, IP 0: s=172.16.31.3 (FastEthernet0/3), d=224.0.0.9, tohlen
372, fragment 0, fo 0, rcvd 2
01:20:58: UDP src=520, dst=520
```



**To obtain a comprehensive view of the configuration tasks in this section, access the SHOWiT engine. With the SHOWiT engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as “show all”.**

## 14.5 EIGRP



### HIDDEN ISSUES TO SPOT WITH EIGRP

**Issue:** Allow networks 192.168.2.0 and 192.168.3.0 to be accepted into R1 from the router FRS.

**Solution:**

On R1, configure an access-list allowing only the two listed prefixes to be accepted by EIGRP. Apply the access-list to a “distribute-list in” command referencing the BVI interface.

```
router eigrp 20
 network 172.16.1.0 0.0.0.255
 distribute-list 10 in BVI1
 no auto-summary
 eigrp stub receive-only

access-list 10 permit 192.168.2.0 0.0.1.0
```

**Verification:**

Issue **show ip route eigrp** command on R1:

```
R1#sh ip route eigrp
D   192.168.2.0/24 [90/281600] via 172.16.1.2, 01:26:48, BVI1
D   192.168.3.0/24 [90/281600] via 172.16.1.2, 01:26:48, BVI1
```

**Issue:** Do not have R1 send EIGRP updates.

**Solution:**

At first, this seems like a passive interface configuration requirement. However, if you make an EIGRP interface passive, it will not transmit any EIGRP HELLO's. If no HELLO packets are generated, the EIGRP speaker will never form an adjacency with another EIGRP speaker. If no adjacency is formed, then the EIGRP speaker will not receive any routing updates. R1 must receive updates from FRS. Therefore, configuring the passive-interface command is unacceptable. The solution to this problem is to configure a distribute-list that denies all routes and applying the distribute-list to the BVI interface.

Another option to fulfill this configuration requirement is to configure “eigrp stub receive-only” under the EIGRP routing process on router R1. With this command, R1 can be configured to silently listen to FRS without advertising any updates to FRS.

**Issue:** Configure EIGRP 10 between R1 and CAT1.

Since the R1 EIGRP AS 20 routing process is configured as “a stub receive-only” router, CAT1 will never receive updates for the rest of the topology if it is also in EIGRP AS 10. Therefore, a second EIGRP process is configured on R1 – EIGRP 20 – and this is used to advertise routing information to CAT1. What

is interesting about this configuration is: the 172.16.1.0/24 prefix assigned to BVI1 on R1 resides in both EIGRP AS 10 and EIGRP AS 20.



*To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".*



### **HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION**

In this scenario the core protocol is OSPF. The edge protocols are RIP and EIGRP. Therefore, each of these edge routing protocols shall be redistributed into OSPF and propagated to other routing domains through OSPF.

The redistribution points are R2, R4, R1, FRS and R3.

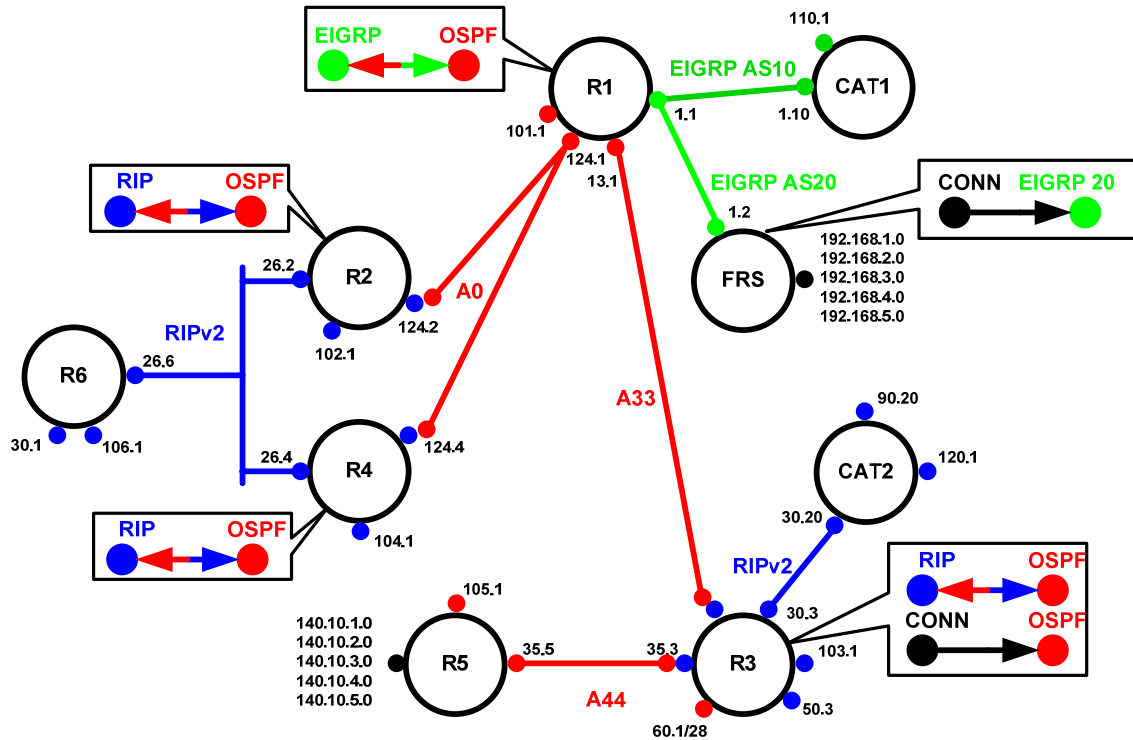
There is a two-point mutual redistribution between RIP and OSPF in this scenario. This particular topology represents an academic redistribution problem, when loops and prefix loss will occur unless special precautions are taken.

For RIP and OSPF to redistribute flawlessly and seamlessly, RIP must not redistribute it is own prefixes back from OSPF, and also RIP shall set the administrative distance for it's own routes to 109 (or any value lower than OSPF distance 110).

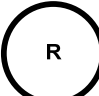





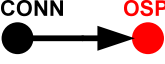


This will make sure that RIP will not lose it is own prefixes and the redistributing router will not install OSPF external prefixes instead of RIP native prefixes in its local routing table.

There is also a loop avoidance technique applied in this Scenario. If you shut down loopback 30.1 on router R6, RIP must not accept this network from OSPF and feed it back to OSPF on the ASBR. Since RIP is configured to not redistribute it is native prefixes from OSPF, this will not happen. Otherwise, assuming that R4 is the ASBR, R1 will point to R4, R4 will point to R2, R2 will point to R1 when 30.1 goes down on R6. This behavior must be prevented.

R1 redistributes all routes between OSPF and EIGRP.



**Legend**

	Router		
	RIP		Loopback
	EIGRP		Mutual redistribution, eg. EIGRP and OSPF
	OSPF		One way redistribution, eg. CONNECTED into OSPF
			Prefix injection
			Trash can

See the scenario master diagram and VLAN table for data link details!

R3 performs mutual redistribution between RIP and OSPF. Connected networks are redistributed into OSPF. One of the connected networks is redistributed as E1, others as E2. Therefore, a route map will have two entries. The first entry will **set metric-type type-1**.

### Redistribution Table

Redist point	Into RIP		Into OSPF		Into EIGRP	
	PERMIT	DENY	PERMIT	DENY	PERMIT	DENY
R1			All routes from EIGRP, All connected networks			
R2	All routes from OSPF	All native RIP routes	Only native RIP routes			
R3	All routes from OSPF		All routes from RIP, Connected networks			
R4	All routes from OSPF	All native RIP routes	Only native RIP routes			
FRS					Connected networks	

The table above provides a useful summary of which prefixes were imported into a given routing protocol. Pay special attention to the color coding of the table. The colors exactly match the colors used in the diagram. Whenever a permit column for a given routing protocol is completely empty, it reflects that no prefixes were redistributed into the routing protocol. This represents that the routing protocol is involved in one-way redistribution.



## 14.6 BGP



### HIDDEN ISSUES TO SPOT WITH BGP

**Issue:** Allow only networks 140.10.2.0/24 -140.10.5.0/24 into AS 100

**Solution:**

Configure either an access-list or prefix-list to allow only the four listed prefixes into AS 100. Since the range of addresses crosses a bit boundary, it will require at least to access-list or prefix-list statements.

In this example, distribute-list in is used as a variant and access-list 100 controls which prefixes are allowed in:

```
neighbor 172.16.35.5 distribute-list 100 in

access-list 100 permit ip 140.10.2.0 0.0.1.0 host 255.255.255.0
access-list 100 permit ip 140.10.4.0 0.0.1.0 host 255.255.255.0
```

**Verification:**

Issue **show ip bgp** command on R3:

```
R3#sh ip bgp
BGP table version is 5, local router ID is 172.16.103.1
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
*> 140.10.2.0/24    172.16.35.5      0         0   500  i
*> 140.10.3.0/24    172.16.35.5      0         0   500  i
*> 140.10.4.0/24    172.16.35.5      0         0   500  i
*> 140.10.5.0/24    172.16.35.5      0         0   500  i
```

**Issue:** Make sure that R1 has networks from AS500 installed in its routing table.

**Solution:**

When R1 receives the BGP updates from R3, they are marked as “I”BGP updates. The two most commonly encountered configuration requirements that need to be addressed when installing an IBGP learned update into a local routing table are: (1) the next-hop reachability of the IBGP learned update and (2) the issue of synchronization. There is no next-hop reachability issue; however, you must address a synchronization issue. The synchronization issue is based on the rule of synchronization, an anti-blackholing provision. The rule of synchronization states: “A BGP speaker cannot advertise an IBGP learned update to another BGP speaker until a matching entry for the IBGP learned update is in the local routing table.” The rule of synchronization is commonly addressed in one of two ways: (1) redistribute the BGP update into the IGP and the edge of the AS or (2) disable synchronization altogether. In Task 14.8.3,

it instructs you to not use any type of redistribution. Therefore, the best solution for addressing synchronization in this scenario is to disable it.

**Verification:**

Issue **sh ip route bgp** command on R1:

```
R1#sh ip route bgp
    140.10.0.0/24 is subnetted, 4 subnets
B       140.10.4.0 [200/0] via 172.16.35.5, 01:47:20
B       140.10.5.0 [200/0] via 172.16.35.5, 01:47:20
B       140.10.2.0 [200/0] via 172.16.35.5, 01:47:20
B       140.10.3.0 [200/0] via 172.16.35.5, 01:47:20
```

**Issue: Modify your BGP configuration so that R4 can see the networks from 140.10.\*.\* range in its local routing table.... Do not use a full mesh.**

**Solution:**

R4 will be made the third IBGP speaker within AS100. In order for R4 to see the 140.10.\*.\* prefixes in its local routing table, synchronization will need to be disabled on R4 as well. To fulfill the “do not use a full mesh” requirement, configure a route-reflector within AS 100. If you do not configure a route-reflector or confederation, you will need to configure a full-mesh of IBGP speakers.

In this example, R1 is a route reflector:

```
neighbor 172.16.13.3 route-reflector-client
neighbor 172.16.124.4 route-reflector-client
```

**Verification:**

Issue **sh ip bgp** and **sh ip route bgp** on R4. Only 140.10.2.0 will be installed, other prefixes are filtered by distance command and access-list:

```
distance 255 0.0.0.0 255.255.255.255 BGP-distance
ip access-list standard BGP-distance
    deny 140.10.2.0
    permit any

R4#sh ip bgp
BGP table version is 8, local router ID is 172.16.104.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

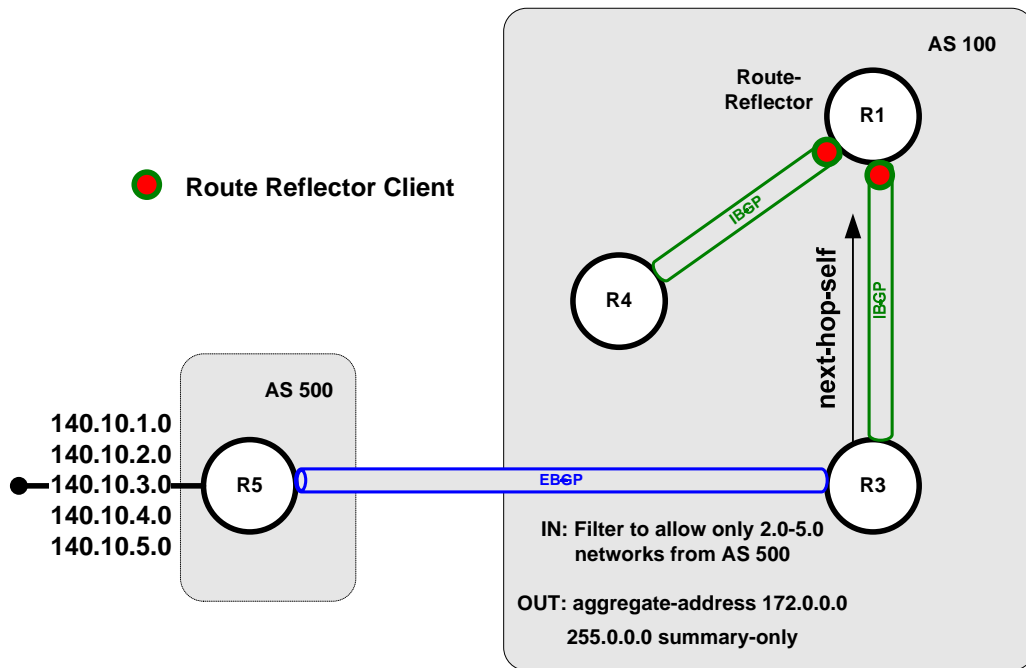
   Network          Next Hop          Metric LocPrf Weight Path
*>i140.10.2.0/24    172.16.13.3         0     100    0 500 i
r>i140.10.3.0/24    172.16.13.3         0     100    0 500 i
r>i140.10.4.0/24    172.16.13.3         0     100    0 500 i
r>i140.10.5.0/24    172.16.13.3         0     100    0 500 i
R4#sh ip route bgp
    140.10.0.0/24 is subnetted, 1 subnets
B       140.10.2.0 [200/0] via 172.16.13.3, 01:53:55
```

**Issue:** Send an aggregate for 172.0.0.0/8 to AS500 and suppress all other routes.

**Solution:**

Configure an aggregate for 172.0.0.0/8 on R3 with the summary-only option. The hidden issue here is that R3 must have a 172.X.X.X entry in its BGP if it is going to advertise the aggregate. Therefore, you must originate a 172.X.X.X entry into the BGP table of R3. This can be accomplished with the following command: "network 172.16.50.0 mask 255.255.255.0".

It was possible to do the same by typing network 172.16.0.0, but latest IOS doesn't advertise network into BGP unless there is an entry of the same network and mask in routing table.



**Verification:**

Issue **show ip bgp** command on R5:

```

R5#sh ip bgp
BGP table version is 7, local router ID is 172.16.105.1
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
*> 140.10.1.0/24  0.0.0.0          0         0   32768  i
*> 140.10.2.0/24  0.0.0.0          0         0   32768  i
*> 140.10.3.0/24  0.0.0.0          0         0   32768  i
*> 140.10.4.0/24  0.0.0.0          0         0   32768  i
*> 140.10.5.0/24  0.0.0.0          0         0   32768  i
*> 172.0.0.0/8    172.16.35.3      0         0     100  i
  
```

## 14.7 Router Maintenance



### **HIDDEN ISSUES TO SPOT WITH ROUTER MAINTENANCE**

**Issue: Configure R4 to supply configuration information to a new router which will be connected to VLAN 40 in the future.**

**Solution:**

You need to read the entire Router Maintenance section to figure out how to use all supplied information. It is the Autoinstall over Ethernet feature that allows you to automate the CISCO router configuration. R4 should be configured as a DHCP server supplying all necessary information to the autoinstall client (a new router). The client is on the same subnet with the DHCP server.

**Issue: The new router should receive its configuration from the TFTP server 172.16.50.100 located on the VLAN 30.**

**Solution:**

You can specify the address of TFTP server by configuring the following command under the dhcp-pool configuration mode: "option 150 ip X.X.X.X", where X.X.X.X is a given TFTP server's IP address.

**Issue: The new router will have IP address 172.16.26.100/24 and MAC address 0010.7be8.131d.**

**Solution:**

Configure the binding between a specified IP address and a MAC address using "hardware-address" command under the ip dhcp pool configuration mode.

**Issue: The new router should send the request for the configuration R100.cfg via R2.**

**Solution:**

This can be done by configuring the following command "bootfile <filename>" under the ip dhcp pool configuration mode.



**To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".**

## Security



### HIDDEN ISSUES TO SPOT WITH SECURITY

**Issue:** *The Administrator does not want any packet to be routed in his network based on the routing path carried in the IP packet.*

**Solution:**

Disable source routing with the “no ip source-route” global configuration command. On some IOS’s, it might be disabled by default.

**Issue:** *Do not allow BOOTP services.*

**Solution:**

Disable BOOTP services with the “no boot network” global configuration command.

**Issue:** *Disable autoconfiguration.*

**Solution:**

Disable autoconfiguration with the “no service config” global configuration command.



*To obtain a comprehensive view of the configuration tasks in this section, access the SHOWiT engine. With the SHOWiT engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as “show all”.*

## 14.8 IPv6



### HIDDEN ISSUES TO SPOT WITH IPv6

**Issue:** *Configure subnet FEC0::124:0/125 on Frame-Relay links between R1, R2 and R4.*

**Issue:** *Configure subnet FEC0::13:0/125 on Frame-Relay link between R1 and R3.*

**Issue:** *Configure subnet FEC0::26:0/125 on VLAN40.*

**Issue:** *Configure subnet FEC0::35:0/125 on VLAN50.*

**Solution:**

IPv6 configuration on interface includes setting ipv6 address using **ipv6 address** command. Additional configuration must be done on NBMA networks, such as assigning map statements and possibly link-local

addresses for easy administration. For example, on R1:

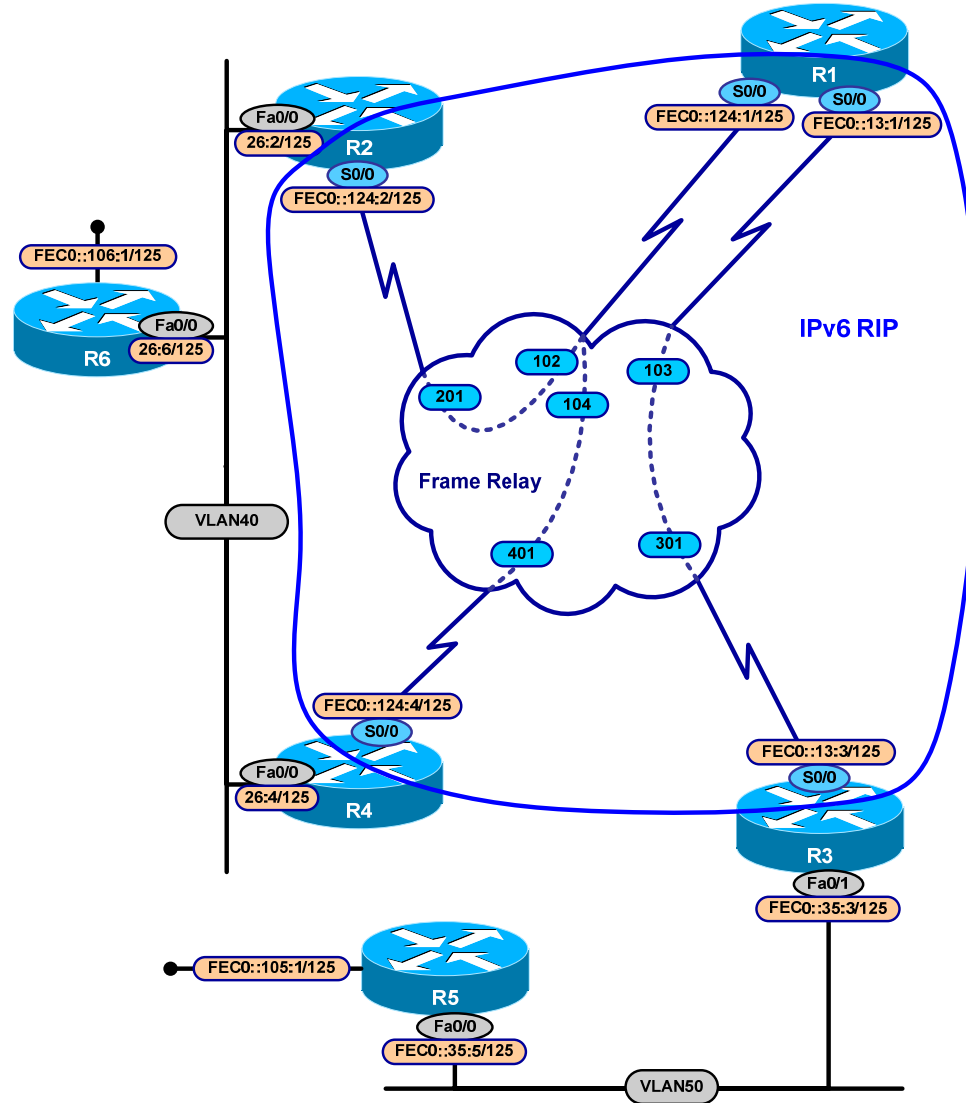
```
interface Serial0/0.124 multipoint
 ip address 172.16.124.1 255.255.255.0
 ip pim sparse-mode
 ip ospf network non-broadcast
 ipv6 address FEC0::124:1/125
 ipv6 address FE80::1 link-local
 ipv6 rip IBGP enable
 frame-relay map ipv6 FE80::2 102 broadcast
 frame-relay map ipv6 FE80::4 104 broadcast
 frame-relay map ipv6 FEC0::124:2 102 broadcast
 frame-relay map ipv6 FEC0::124:4 104 broadcast
 frame-relay map ip 172.16.124.1 102
 frame-relay map ip 172.16.124.2 102 broadcast
 frame-relay map ip 172.16.124.4 104 broadcast
 no frame-relay inverse-arp

interface Serial0/0.13 point-to-point
 ip address 172.16.13.1 255.255.255.0
 ip pim sparse-mode
 ip ospf network non-broadcast
 ip ospf priority 0
 ipv6 address FEC0::13:1/125
 ipv6 address FE80::1 link-local
 ipv6 rip IBGP enable
 frame-relay interface-dlci 103 CISCO
```

**Issue:** Configure loopbacks *FEC0::105:1/125* on R5 and *FEC0::106:1/125* on R6.

**Solution:**

Use **ipv6 address** command on these loopbacks.



**Issue:** Configure RIP for IPv6 between R1, R2, R3 and R4. Only Frame-Relay networks must be included in RIP process.

**Solution:**

RIP is configured on interface for IPv6. Use command **ipv6 rip IBGP enable** on R1, R2, R3 and R4. IBGP – is an identifier, and can be anything.

For example:

```
interface Serial0/0.13 point-to-point
  ipv6 rip IBGP enable
```

**Verification:**

Make sure that R3 knows about 124:0/125 and vice versa:

Use **sh ipv6 route rip** command On R3:

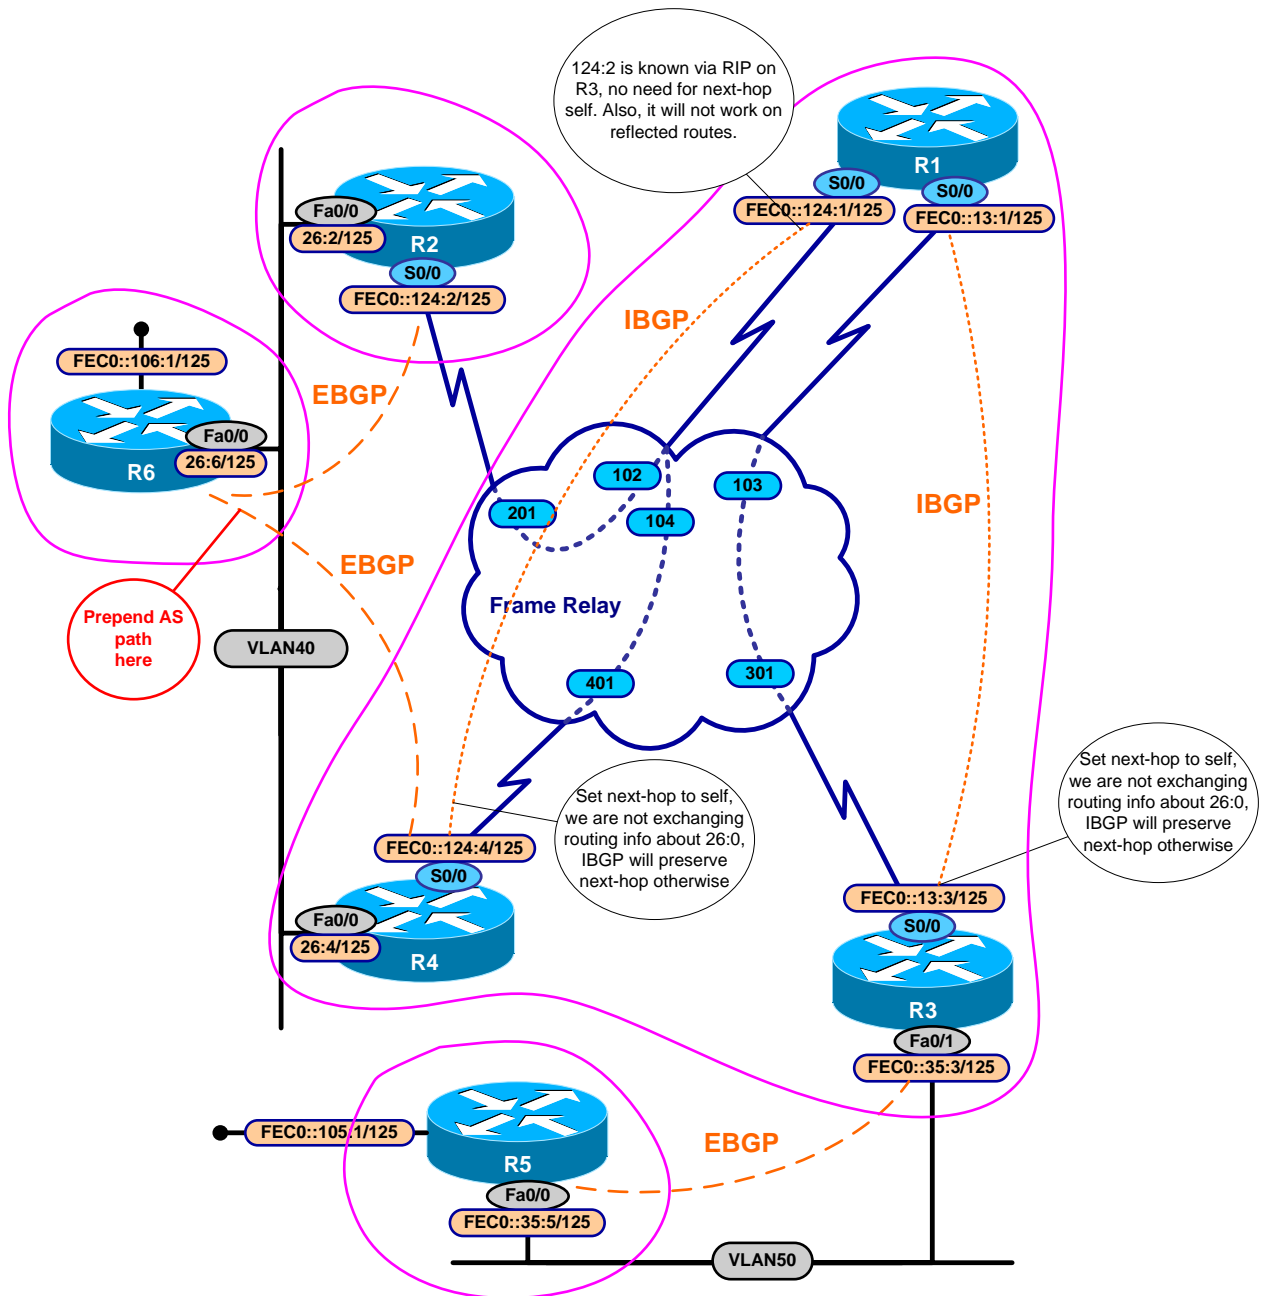
```
R3#sh ipv6 route rip
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
R   FEC0::124:0/125 [120/2]
    via FE80::1, Serial0/0
```



**To obtain a comprehensive view of the configuration tasks in this section, access the SHOWiT engine. With the SHOWiT engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".**



### IPv6 BGP diagram



**Issue: Configure BGP AS 100 for IPv6 family on R1, R3 and R4.**

**Issue: Configure BGP AS 200 for IPv6 family on R2.**

**Issue: Configure BGP AS 600 for IPv6 family on R6.**

**Issue: Configure BGP AS 500 for IPv6 family on R5.**

**Solution:**

AS 100 is already configured for IPv4 and BGP process exists. For others, BGP process doesn't exist. To configure BGP for IPv6 use **address-family ipv6** command under **router bgp AS#** command. Configuration of BGP for IPv6 is similar to BGP configuration for IPv4, except that addresses and networks are in v6 format. For example (R6), areas marked yellow appear automatically:

```
router bgp 600
  no synchronization
  bgp log-neighbor-changes
  neighbor FEC0::26:2 remote-as 200
  neighbor FEC0::26:4 remote-as 100
  no auto-summary
  !
  address-family ipv4 multicast
  no auto-summary
  no synchronization
  exit-address-family
  !
  address-family ipv6
  neighbor FEC0::26:2 activate
  neighbor FEC0::26:4 activate
  neighbor FEC0::26:4 route-map ADD-AS out
  network FEC0::106:0/125
  exit-address-family
  !
  address-family ipv4
  no neighbor FEC0::26:2 activate
  no neighbor FEC0::26:4 activate
  no auto-summary
  no synchronization
  exit-address-family
```

**Issue:** Connect AS 600 to AS 200 between R2 and R6, connect AS 600 and AS 100 between R6 and R4, connect AS 100 and AS 500 between R3 and R5.

**Solution:**

Use **neighbor** statement under **address-family ipv6** to establish neighbor relationship between AS's.

**Verification:**

Use **sh bgp ipv6 summary** command:

```
R6#sh bgp ipv6 summary
[skipped]

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ Up/Down  State/PfxRcd
FEC0::26:2   4    200    182    183     6    0    0 02:55:25  1
FEC0::26:4   4    100    182    181     6    0    0 02:54:54  1
```

**Issue: AS 100 IBGP connections must remain the same with IPv4 BGP task.**

**Solution:**

You must keep route-reflector on R1 for IPv6 part too.

**Issue: Advertise FEC0::105:0/125 and FEC0::106:0/125 into BGP as internal.**

**Solution:**

Use **network** command on R5 and R6. The point is that once you have these prefixes advertised, since networks others than 105 and 106 are not known to the remote systems, AS 100 must replace next hop on R4 and R3 to provide reachability. **Next-hop-self** must be configured on R3 and R4 towards R1. Reachability between 124:0/125 and 13:0/125 is provided with RIP.

**Verification:**

Make sure you have prefixes installed on R5 and R6:

```
R5#sh ipv6 route bgp
IPv6 Routing Table - 7 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
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
B   FEC0::106:0/125 [20/0]
    via FE80::250:54FF:FE7C:A561, FastEthernet0/0
```

```
R6#sh ipv6 route bgp
IPv6 Routing Table - 7 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
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
B   FEC0::105:0/125 [20/0]
    via FE80::204:C1FF:FE8E:BC0, FastEthernet0/0
```

**Issue: Make sure traffic from FEC0::105.1 to FEC0::106:1 goes thru R3, R1, R2 to R6, and return traffic goes thru R4, R1, R3 to R5. Apply configuration on R6. Use one well-known mandatory attribute. Do not change origin.**

**Solution:**

Traffic going towards R6 from R5 must go thru R2. Default path will go thru R4. You are allowed to use only one well-known mandatory attribute. Well-known mandatory attributes are Origin, Next-hop and AS-path. You can't change origin. Changing next hop will not help in this task. Therefore you need to prepend AS-path on R6 towards R4 with multiple AS 600 entries so it will be longer and R2 will be more preferred.

```
neighbor FEC0::26:4 route-map ADD-AS out
```

```
route-map ADD-AS permit 10
  set as-path prepend 600 600
```

**Verification:**

Verify that R1 prefers R2 for packets going towards 106:0/125:

```
R1#sh ipv6 route bgp
IPv6 Routing Table - 8 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
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
B   FEC0::105:0/125 [200/0]
    via FEC0::13:3, Null
B   FEC0::106:0/125 [20/0]
    via FE80::2, Serial0/0.124
```



*To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as “show all”.*

## 14.9 Catalyst Specialties



### HIDDEN ISSUES TO SPOT WITH CATALYST SPECIALTIES

**Issue:** CAT2 will be connected to the test environment network 172.16.90.0/24 in the future. Interface FastEthernet0/7 will be used for this purpose.

**Solution:**

This configuration requirement sets the stage for the other tasks that need to be performed in this section. You need to determine what interface is chosen to apply the remaining configuration tasks of this section.

**Issue:** Provide a solution to allow traffic sourced on the network 172.16.90.0/24 only from the selected hosts (.1, .3, .5, .7) to get into your network. Apply the solution on CAT2. Do not use filtering techniques based on Layer 2 filtering. Use minimal number of statements for this task.

**Solution:**

Notice that all hosts have odd numbers and are in the range from 1 to 7. In the binary form, all hosts can be represented as follows:

- o 00000001 - .1
- o 00000011 - .3
- o 00000101 - .5
- o 00000111 - .7

Notice that the first 5 digits from the left must be 0 and last digit must be 1. We do not care about the second and third digits from the right. This logic can be represented by the following base and wildcard:  
172.16.90.1 0.0.0.6

This fulfills the requirement of matching the desired range of addresses with the minimum number of statements.

**Issue: Only telnet and ICMP must be allowed as user data traffic on VLAN40. Configure 3550 switches to accommodate this restriction. Do not use VLAN map based filtering.**

**Solution:**

Since VLAN map based filtering is not allowed, you need to identify the interfaces of the switches attached to VLAN40 and configure access-lists that allow ICMP and telnet traffic as well as routing protocols used, in our case that would be RIP UDP port 520.



To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".

## 14.10 Address Administration



### HIDDEN ISSUES TO SPOT WITH ADDRESS ADMINISTRATION

**Issue: Configure the 1.1.1.3/24 address on R3's FastEthernet0/0 interface without changing any pre-existing IP addresses. Use the 1.1.1.0 private address space and use a portion of the legal address space of the R3 FastEthernet subnet.**

**Solution:**

This task is a NAT configuration requirement without ever explicitly mentioning NAT. The 1.1.1.3 address assigned to R3's FastEthernet interface will be assigned as a secondary IP address. This 1.1.1.0 will be the NAT inside address. The primary IP address of the FastEthernet interface will be the NAT outside address.

Configure Fa0/0 as NAT inside interface:

```
interface FastEthernet0/0
 ip address 1.1.1.3 255.255.255.0 secondary
 ip address 172.16.31.3 255.255.255.0
 ip nat inside
 ip rip send version 2
 duplex auto
 speed auto
```

Configure Serial0/0, FastEthernet0/1.30 and FastEthernet0/1.50 as NAT outside interface:

```
interface Serial0/0
  ip nat outside

interface FastEthernet0/1.30
  encapsulation isl 30
  ip nat outside

interface FastEthernet0/1.50
  encapsulation isl 50
  ip nat outside
```

Configure NAT translation and list of source addresses:

```
access-list 50 permit 1.1.1.0 0.0.0.255
ip nat inside source list 50 interface FastEthernet0/0 overload
```

**Verification:**

Since this is a dynamic NAT, you will not see translations until packets start flowing. You can still check that configuration is accepted by using **sh ip nat stat** command:

```
R3#sh ip nat stat
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Outside interfaces:
  Serial0/0, FastEthernet0/1.30, FastEthernet0/1.50
Inside interfaces:
  FastEthernet0/0
Hits: 0 Misses: 0
Expired translations: 0
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 50 interface FastEthernet0/0 refcount 0
```

**Issue: Make sure that CAT1 can ping all interfaces in your network. However, allow TELNET access only from the R3's management loopback interface.**

**Solution:**

The methods to apply to configure CAT1 so that it can ping all interfaces in the Scenario network are straightforward. The method to apply to allow TELNET access only from R3's management loopback interface is not as clear. In order to accomplish this task, create an access-list permitting only the R3 loopback interface. Then, apply the access-list under the "line vty" mode with the following command: "access-class 1 in".

```
line vty 0 4
  access-class 1 in
  privilege level 15
  password cisco
  login
line vty 5 15
```

```
access-class 1 in
privilege level 15
password cisco
login
```

### Verification:

Telnet to CAT1 from R3 using loopback0 as source interface and otherwise:

```
R3#telnet 172.16.1.10 /source-interface Loopback0
Trying 172.16.1.10 ... Open

User Access Verification

Password:
CAT1#exit

[Connection to 172.16.1.10 closed by foreign host]
R3#telnet 172.16.1.10
Trying 172.16.1.10 ...
% Connection refused by remote host
```



To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as “show all”.

## 14.11 Multicast



### HIDDEN ISSUES TO SPOT WITH MULTICAST

**Issue:** Announce the shared root without use of any dense groups or static configurations.

#### Solution:

Since a shared root is involved, you know this configuration requirement is a sparse mode configuration. When you configure sparse mode, a rendezvous point is involved. The challenge with sparse mode is how to advertise to all sparse mode routers the location of the rendezvous point. Three methods of advertising the rendezvous point are: static configuration, auto-rp using pim sparse-dense mode, the bootstrap routing protocol. The configuration requirement above prohibits the use of “static configurations or dense groups.” The static configuration restriction eliminates the static method of PIM Sparse Mode Rendezvous Point advertisement. The “dense group” restriction eliminates the auto-rp method since it requires pim sparse/dense mode. Therefore, only one Rendezvous Point advertisement method remains: the bootstrap routing protocol. Therefore, to fulfill this configuration requirement, you need to configure PIM Sparse Mode with the Bootstrap Routing Protocol. Since it is specified that R1 is to be configured as the shared root, configure R1 as the candidate RP for the 239.10.10.10 multicast group as well as the bootstrap router.

**Verification:**

```
R1#sh ip pim bsr-router
PIMv2 Bootstrap information
This system is the Bootstrap Router (BSR)
  BSR address: 172.16.101.1 (?)
  Uptime:      13:19:32, BSR Priority: 0, Hash mask length: 0
  Next bootstrap message in 00:00:49
  Candidate RP: 172.16.101.1(Loopback0)
  Advertisement interval 60 seconds
  Next advertisement in 00:00:27
  Group acl: 80

R1#sh access-li 80
Standard IP access list 80
  10 permit 239.10.10.10
```

And on R2:

```
R2#sh ip pim bsr-r
PIMv2 Bootstrap information
  BSR address: 172.16.101.1 (?)
  Uptime:      13:18:25, BSR Priority: 0, Hash mask length: 0
  Expires:     00:02:25
```

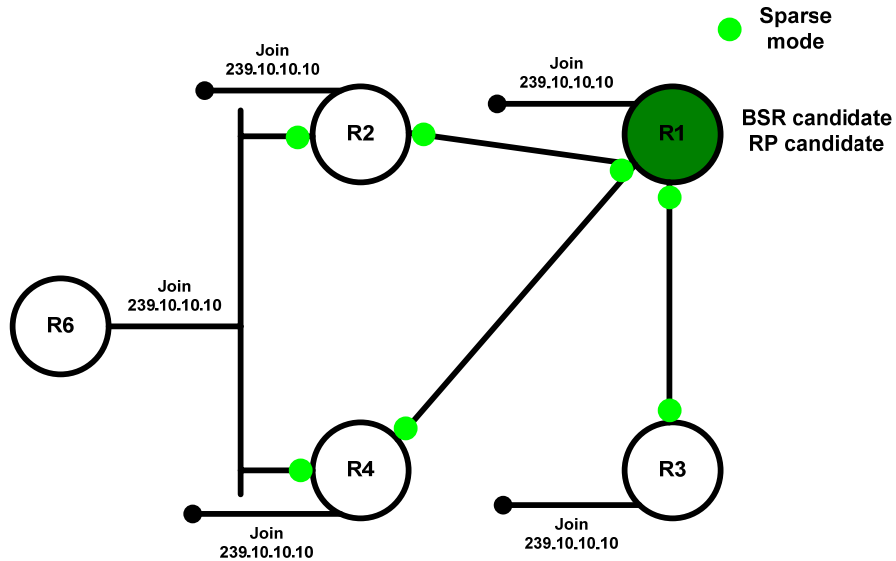
**Issue:** Make sure traffic to 239.10.10.10 is flooded out of the appropriate VLAN 40 ports only.

**Solution:**

Configure the following global configuration command on the Catalyst 3550: “mac-address-table static 0100.5e0a.0a0a vlan 40 interface Xy” where Xy is the interface name and number. Do not forget to correctly translate the IP multicast address to the Ethernet multicast address. An IP multicast address translated to an Ethernet multicast address always has the first six hex-digits (or first 24 bits) set to 01-00-5E. The 25<sup>th</sup> bit of an IP multicast address mapped to an Ethernet multicast address is always 0. The remaining 23 bits of the Ethernet multicast address are directly derived from the low order 23 bits of the IP multicast address.

**Implementation:**





To fulfill the requirement 14.17.4 on vlan 40 static mac-address-table entries shall be created mapping mac address 01-00-5e-0a-0a-0a to all ports on vlan 40. IGMP snooping on vlan 40 must be disabled before static mac address can be entered.

**CAT1:**

```
no ip igmp snooping
mac-address-table static 0100.5e0a.0a0a vlan 40 interface FastEthernet0/8
```

**CAT2:**

```
no ip igmp snooping
mac-address-table static 0100.5e0a.0a0a vlan 40 interface FastEth0/13 FastEth0/14
```

**Verification:**

Make sure that configuration is accepted:

```
CAT1#sh mac-address-table static | inc 5e0a
 40 0100.5e0a.0a0a  STATIC Fa0/8
```



To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".

## 14.12 Quality of Service

**Issue:** *Traffic originating from the 172.16.35.5 address of R5 with an IP precedence setting of either 3 or 4 must not consume more than 50% of the bandwidth of the Frame-Relay interface on R3.*

**Solution:**

This configuration requirement is limiting the bandwidth consumption of a specific classification of traffic. Two possible options can be considered to limit the bandwidth consumption by a specific class of traffic.

Option #1: A policing mechanism.

Option #2: A shaping mechanism.

Regarding limiting bandwidth consumption to 50% of the bandwidth of the Frame-Relay interface on R3, it is important to reference the Frame-Relay switch configuration supplied at the end of the DOiT lab and determine the clock rate of this Frame-Relay connection. This configuration displays a Frame-Relay link clock rate of 64,000 bps. Therefore, the specified class of traffic should not consume more than 32000 bps of bandwidth.

An option to consider is to set this threshold is the MQC. It is recommended that you always consider using the MQC when fulfilling a QoS configuration requirement since the MQC is the most versatile and state of the art QoS tool in the IOS. The MQC allows you to set either shaping or policing bandwidth consumption thresholds with either a fixed value or a percentage of an interface's IOS reference bandwidth.

Since there is a configuration requirement in this specific Scenario stating that you must "not use any "percentage" keyword when fulfilling this configuration requirements of this section", any configuration options involving the "percentage" keyword are excluded. Therefore, you must specify the specific bandwidth of 32000 bps.

**Issue:** *If traffic that meets the classification stated above attempts to exceed this bandwidth consumption threshold, buffer the traffic.*

**Solution:**

This statement limits your configuration options to a traffic shaping option since there is a requirement to buffer any traffic that exceeds the bandwidth consumption threshold. A policing mechanism never buffers traffic. Now, that you have narrowed the configuration requirement down to a traffic shaping requirement, you must now determine which traffic shaping tool to use. Cisco provides many options including:

Frame-Relay Traffic Shaping

Class-Based Traffic Shaping using the "shape peak" option

Class-Based Traffic Shaping using the "shape average" option

The Frame-Relay traffic shaping option is eliminated by the language of the last configuration requirement provided at the end of the QoS section. There is a statement at the end of this section that directs you to select "a QoS mechanism that can be applied to any type of physical interface". This should be inferred to mean to select a QoS mechanism that could be applied to an Ethernet interface as well as a Frame-Relay

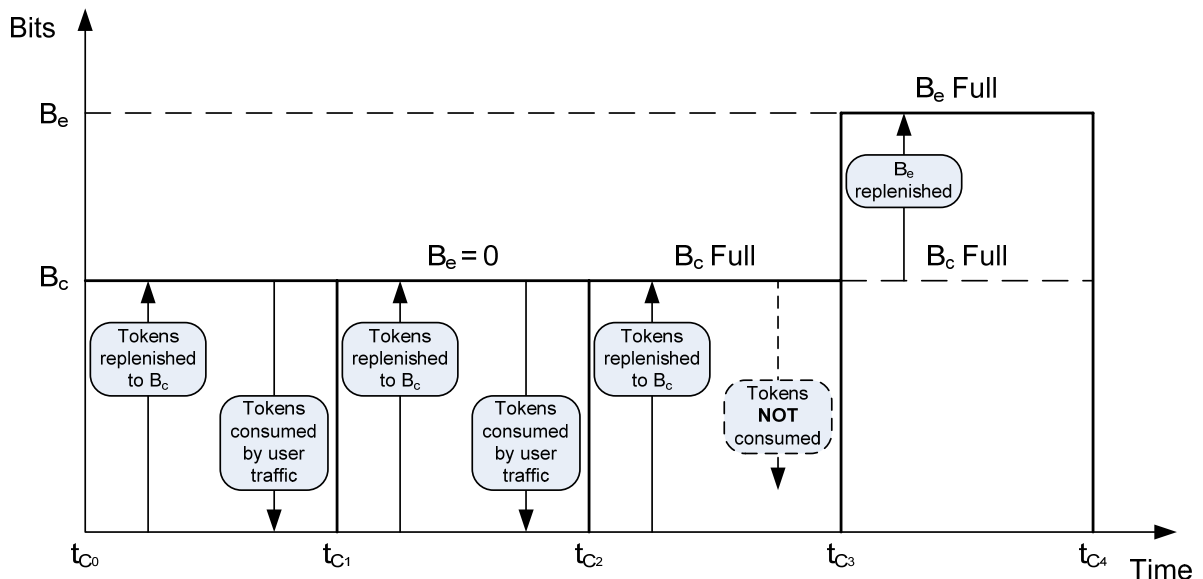
interface. This requirement eliminates the Frame-Relay Traffic Shaping option since this can only be applied to Frame-Relay interfaces. Now, our shaping options are limited down to the following two choices:

Class-Based Traffic Shaping using the “shape peak” option

Class-Based Traffic Shaping using the “shape average” option

Let’s examine two of these requirements in particular. The MQC “shape peak” configuration option transmits traffic that amounts to both  $B_c$  (committed burst) and  $B_e$  (excessive burst) every interval. The MQC “shape average” configuration option transmits traffic that amount to only  $B_c$  every interval; however, it will also transmit traffic that amount to  $B_e$  after periods of inactivity. This acts as a “bonus round” for low usage of your interface by the “shape average” mechanism. Here is a diagram that explains this feature of the “shape average” command. It involves an understanding of the Token-Bucket algorithm and how it operates with the “shape average” command. Please make careful note of the following two separate operations with the “shape average” command:

(1) the replenishment of the token-bucket and (2) the consumption of tokens due to transmitting traffic.



The key thing to note in the diagram above is that the “shape average” mechanism only replenishes the  $B_e$  value after an interval of no use of  $B_c$  tokens. Stated another way,  $B_e$  tokens are replenished only after an interval when no  $B_c$  tokens have been consumed AND the  $B_c$  token bucket completely full. Therefore, with the “shape average” mechanism, you only receive  $B_e$  tokens when the  $B_c$  token bucket is full. It is possible that you will never receive the ability to burst at the  $B_e$  rate. This possibility is fulfilled when traffic is constantly being transmitted through a “shape average” mechanism and the  $B_c$  token bucket is never able to fill. This behavior of the “shape average” mechanism is markedly different than the “shape peak” mechanism.

Now, we must apply what we have learned to the specific language of this Scenario. We have narrowed our possible configuration options down to two choices:

## Class-Based Traffic Shaping using the “shape peak” option Class-Based Traffic Shaping using the “shape average” option

Since this Scenario includes a configuration requirement that explicitly states: “Use a technique that provides the ability to consume additional bandwidth above the configured threshold only after periods of inactivity”, this directs us to configuring the “shape average” mechanism. Remember, the “shape peak” mechanism transmits both Bc and Be amounts of data for each and every interval. The “shape average” mechanism transmits Be only after periods of inactivity.

Now, that it is determined which MQC shaping mechanism to use (shape average), we need to determine which values to apply to the shape average configuration. The configuration task specified that the bandwidth consumption threshold should be 50% of the interface bandwidth. Therefore, the CIR must be configured with the value of 32000 bps. Furthermore, the section included the following configuration requirement:

“As a bonus for periods of inactivity, allow this classification of traffic to transmit an additional 12,000 bits. Transmit an equal value each and every interval as well.”

This is directing you to configure both the Bc and Be to a value of 12000 bits each.

***Issue: If traffic that meets the classification stated above attempts to exceed this bandwidth consumption threshold, buffer the traffic. The traffic with an IP precedence of 4 should get twice as much bandwidth as traffic marked with IP precedence 3. The maximum bandwidth allocation for a single class of traffic must be 16000 bps.***

### **Solution:**

As already stated, this QoS configuration requirement must apply a traffic shaping mechanism since the task specifies to buffer traffic that exceeds a specified bandwidth consumption threshold. This presents a second issue to address: how to queue the excess traffic. In the configuration task presented above, you are directed to provide twice as much bandwidth for traffic marked with an IP precedence of 4 versus traffic marked with an IP precedence of 3. Since two different classes of traffic must be queued in different ways, a second MQC policy-map must be configured and nested inside of the MQC policy-map configured with the “shape average” command. Therefore, this QoS section requires that a nested MQC policy-map be configured. As a result, two policy-maps must be created.

Since the maximum bandwidth allocation for a single class must be 16000 bps and since there is only 32000 bps of second of available bandwidth allocated to the shaping mechanism, 16000 bps of bandwidth must be allocated to shaped traffic marked with IP precedence 4 and 8000 bps must be allocated to shaped traffic marked with IP precedence 3 (half the amount of 16000 bps).

### **Configuration**

Since this QoS section involves the configuration of two nested MQC policy-maps, two sets of the following IOS commands must be found in the configuration of R3:

Two class-maps  
Two policy-maps

## Two service-policy commands

First, an explanation of the policy-map for queuing traffic that exceeds the shaping mechanism's threshold will be applied. Once this secondary policy-map is presented, the primary policy-map will be presented. The primary policy-map contains the traffic shaping configuration commands as well as a service-policy statement to apply the secondary policy-map.

### Configuring the MQC for Nested Queuing within a Shaper

First, the class-maps must be configured. Only two class-maps need to be configured:

```
class-map match-all prec4
  match ip precedence 4
class-map match-all prec3
  match ip precedence 3
```

Second, the class-maps must be associated with a policy-map and bandwidth must be assigned to the queues:

```
policy-map q-shaper
  class prec4
    bandwidth 16
  class prec3
    bandwidth 8
```

Please note bandwidth allocation for the queues is in “kilobits” per second.

Finally, the policy-map must be applied. In this scenario it is applied inside of another policy-map:

```
policy-map shaper
  class shaper
    shape average 32000 12000 12000
    service-policy q-shaper
```

### Configuring the MQC for Class-Based Traffic Shaping

On again, first, the class-maps must be configured. Only two class-maps need to be configured:

```
class-map match-all shaper
  match access-group 101
```

This particular class-map references an access-list. Here is the two line access-list that meets the requirements of this section:

```
access-list 101 permit ip host 172.16.35.5 host 172.16.124.1 precedence flash (3)
access-list 101 permit ip host 172.16.35.5 host 172.16.124.1 precedence flash-override (4)
```

Second, the class-map must be associated with a policy-map:

```
policy-map shaper
class shaper
  shape average 32000 12000 12000
  service-policy q-shaper
```

The first command applies the shape average command. It is followed by the value of 32000 for the CIR in bits-per-second, the Committed Burst (Bc) 12000 in bits and the Excess Burst (Be) 12000 in bits.

Finally, the policy-map must be applied.

This policy-map is applied to the physical Frame-Relay interface on R3 as directed by the instructions of the Scenario:

```
interface Serial0/0
ip address 172.16.13.3 255.255.255.0
service-policy output shaper
end
```

## Verification

An extremely useful and detailed verification IOS command for any MQC configuration is the “show policy interface” command. A sample display of this command is provided below for this Scenario’s QoS configuration requirements:

```
R3#show policy-map interface serial 0/0

Serial0/0
  Service-policy output: shaper

  Class-map: shaper (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group 101
    Traffic Shaping
      Target/Average   Byte   Sustain   Excess   Interval   Increment
      Rate             Limit  bits/int  bits/int  (ms)       (bytes)
      32000/32000     3000   12000    12000    375        1500

      Adapt Queue   Packets  Bytes   Packets  Bytes  Shaping
      Active Depth
      -      0       0       0       0       0      no

  Service-policy : q-shaper

  Class-map: prec4 (match-all)
    0 packets, 0 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: ip precedence 4
    Queueing
      Output Queue: Conversation 25
      Bandwidth 16 (kbps) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 0/0
      (depth/total drops/no-buffer drops) 0/0/0

  Class-map: prec3 (match-all)
    0 packets, 0 bytes
```

```

5 minute offered rate 0 bps, drop rate 0 bps
Match: ip precedence 3
Queueing:
  Output Queue: Conversation 26
  Bandwidth 8 (kbps) Max Threshold 64 (packets)
  (pkts matched/bytes matched) 0/0
  (depth/total drops/no-buffer drops) 0/0/0

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

Class-map: class-default (match-any)
  456 packets, 23452 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any

```

The output of this very useful show command has been subdivided into four sections. Each section is worthy of note. The first section (in yellow) at the top provides traffic shaping statistics. If the “queue depth” field in this first section is greater than 1, then the second section statistics will increment. The second section is dedicated to the nested queuing configuration provided in this Scenario. The last two sections are the “class-default” sections of this configuration. The first class-default section (red) applies to the nested queuing policy-map. The second class-default section (yellow) applies to the primary shaper policy-map. Notice in this particular display, the yellow class-default section has processed 456 packets while the red class-default section has processed 0 packets. In fact, if you look at the packet counters in the queuing section of this display, they are also set to 0. This is because no packets exceeded the bandwidth consumption threshold and have needed to be queued.

**Note:** It is oftentimes useful to reset the counters of the “show policy-map interface” table with the “clear counters interface” command. Here is an example of using this command:

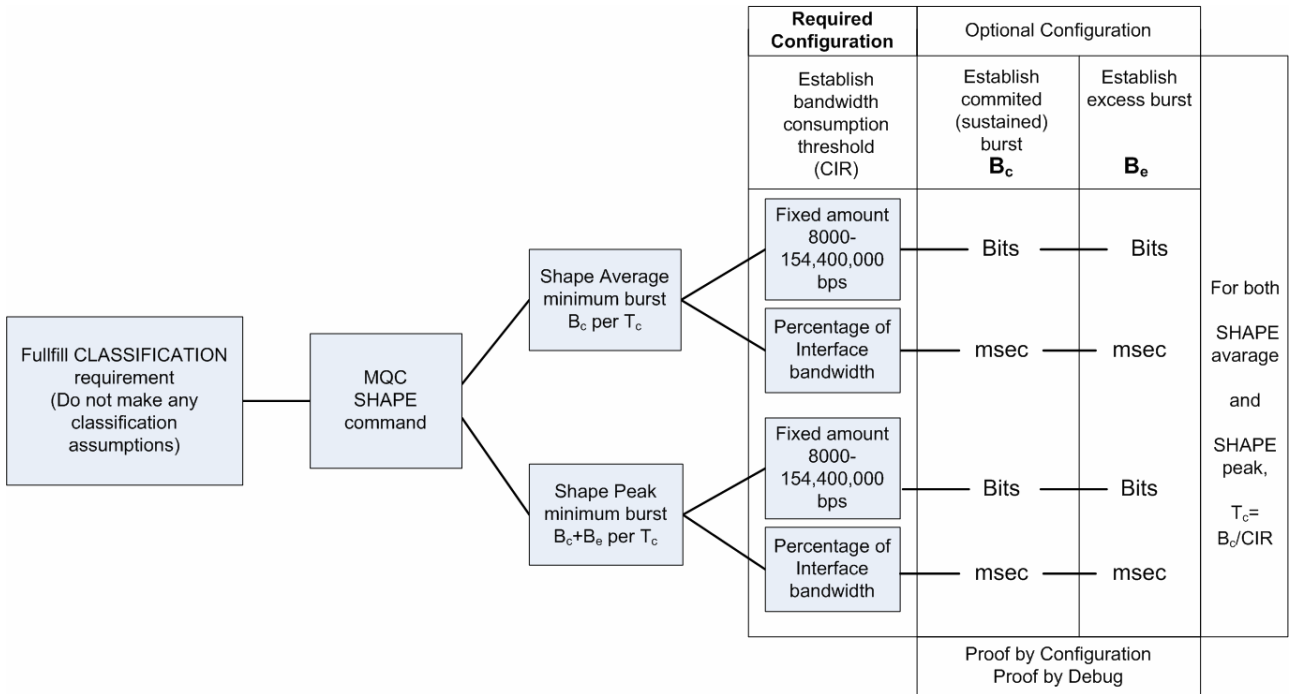
```
R5#clear counters serial 1/0
```

Clear "show interface" counters on this interface [confirm]

Always specify the physical interface with this command.

## Conclusion

The general recommendation of approaching any CCIE lab configuration requirement is to know all of your IOS configuration options. In an effort to provide a visual representation of your MQC traffic shaping options, consider the following decision diagram:



Consider building a similar decision diagram for as many technologies that you might encounter in the CCIE lab. You will find similar diagrams in the NetMasterClass Technical Library.



To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".

### 14.13 IOS Features

**Issue:** Simulate traffic every 2 minutes by sending 100 test packets and monitor the jitter of this simulated traffic from R5 to R1. Provide a minimal configuration on R1.

**Solution:**

This is an SAA (IP SLA in IOS 12.4) configuration requirement. You can configure SAA/IP SLA to simulate traffic on a periodic basis and measure delay variation between packets. It involves entering a collection of "rtr" or "ip sla" commands on the router that will be generating the traffic. In this specific configuration, these commands will be entered on R5. Only one global configuration command – "rtr responder" or "ip sla monitor responder" needs to be configured on R1. This fulfills the "minimal configuration on R1" requirement stated above.



**Issue:** Provide two separate sets of simulated traffic from R5 to R1 with one set of traffic marked as precedence 4 and the second set of traffic marked as precedence 3.

**Solution:**

A new feature of SAA is that you can specify specific TOS settings for a configured set of test packets. It is very important to remember that you configure this feature with TOS values and not IP Precedence or DSCP values. Therefore, you must be very careful to map the TOS value to the correct IP Precedence or DSCP value. Provided below is a table to help perform this translation to meet the specific requirements of this Scenario.

Baseline TOS octet Decimal Value	128	64	32	16	8	4	2	1
IP Precedence Value	4	2	1					

The scenario directs you to mark a set of test packets with the IP Precedence setting of 4. Using the table above, an IP Precedence value of 4 maps to a TOS value of 128.

The scenario also directs you to mark a separate set of test packets with the IP Precedence setting of 3. Using the table above, an IP Precedence value of 3 maps to a TOS value of 96.

**Configuration for SAA/IP SLA**

**R5**

```
R5#sh run | begin ip sla
ip sla monitor 1
  type jitter dest-ipaddr 172.16.101.1 dest-port 16387 num-packets 100
  tos 128
  frequency 120
ip sla monitor schedule 1 life forever start-time now
ip sla monitor 2
  type jitter dest-ipaddr 172.16.101.1 dest-port 16388 num-packets 100
  tos 96
  frequency 120
ip sla monitor schedule 2 life forever start-time now
```

The commands highlighted in yellow place you into SAA/IP SLA configuration mode. The commands highlighted in green activate SAA/IP SLA.

**R1**

```
R1#sh run | include ip sla
ip sla monitor responder
```

Only one command is all you need on the target side of an SAA/IP SLA configuration when collecting jitter statistics.

## Verification for SAA/IP SLA

```
R1#show interfaces s0/0.13 precedence
Serial0/0.13
  Input
    Precedence 3: 101 packets, 6484 bytes
    Precedence 4: 101 packets, 6484 bytes
```

Notice how only 101 only packets are received. SAA/IP SLA is configured to send exactly 100 packets. Consider the first packet as an SAA/IP SLA control plane packet. This is how you end up with the amount of 101 packets.

**Issue:** For any type of IP Service, that needs to be configured for the task above, make sure the R1 is the server of any supporting service.

### Solution:

This configuration task directs your NTP configuration for SAA. In order to collect and process SAA jitter values, NTP must be running between the two test devices. The two SAA test devices, the generator and the responder, must be synchronized with NTP. While the configuration requirements in this Scenario are silent about this, you should know this if you have developed an expertise in configuring the jitter testing feature of SAA. Now, that you know you need to configure NTP with the SAA configuration, you must determine whether to configure an NTP peer relationship or client/server relationship. Since the configuration task above states that you must make R1 the “server for any supporting service”, this directive applies to directly to your NTP configuration. Therefore, configure R1 as the NTP server and R5 as the NTP client.

## Configuration for NTP

### R1

```
R1#sh run | begin ntp
ntp master
```

### R5

```
R5#sh run | begin ntp
ntp server 172.16.101.1
```

## Verification for NTP

```
R5#sh ntp associations detail
172.16.101.1 configured, our_master, sane, valid, stratum 8
ref ID 127.127.7.1, time C6F7EB66.4EE230C6 (20:07:02.308 UTC Wed Oct 12 2005)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 125.03, reach 3, sync dist 8014.206
delay 28.18 msec, offset -0.6516 msec, dispersion 7875.09
precision 2**24, version 3
org time C6F7EB9C.84ECC0C8 (20:07:56.519 UTC Wed Oct 12 2005)
rcv time C6F7EB9C.88B34887 (20:07:56.533 UTC Wed Oct 12 2005)
xmt time C6F7EB9C.81781B63 (20:07:56.505 UTC Wed Oct 12 2005)
filtdelay = 28.18 28.27 0.00 0.00 0.00 0.00 0.00 0.00
filtoffset = -0.65 -0.50 0.00 0.00 0.00 0.00 0.00 0.00
filterror = 0.02 0.99 16000.0 16000.0 16000.0 16000.0 16000.0 16000.0
R5#
```

```
R5#sh clock detail
20:07:17.605 UTC Wed Oct 12 2005
Time source is NTP
```

**Note:** NTP commands appear at the end of the configuration script.

**Note:** Always remember that NTP advertises its time in UTC format.

**Issue:** Provide a mechanism that can be used on R1 to count the number of packets received by the router on its point-to-point Frame-Relay subinterface. Packet counts should be classified by the received packets' precedence setting. Do not use any access-lists to fulfill this requirement.

One method to count the number packets received on an interface on a per-IP precedence basis, is the interface configuration command "ip accounting".

### Configuration for IP Accounting

R1#

```
interface Serial0/0.13 point-to-point
ip address 172.16.13.1 255.255.255.0
ip accounting precedence input
end
```

### Verification for IP Accounting

```
R1#show interfaces s0/0.13 precedence
Serial0/0.13
  Input
    Precedence 6:  2 packets, 107 bytes
```

**Note:** It is oftentimes useful to reset the counters of the IP Accounting table with the "clear counters interface" command. Here is an example of using this command:

```
R5#clear counters serial 1/0
```

Clear "show interface" counters on this interface [confirm]



To obtain a comprehensive view of the configuration tasks in this section, access the SHOWit engine. With the SHOWit engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".