NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II

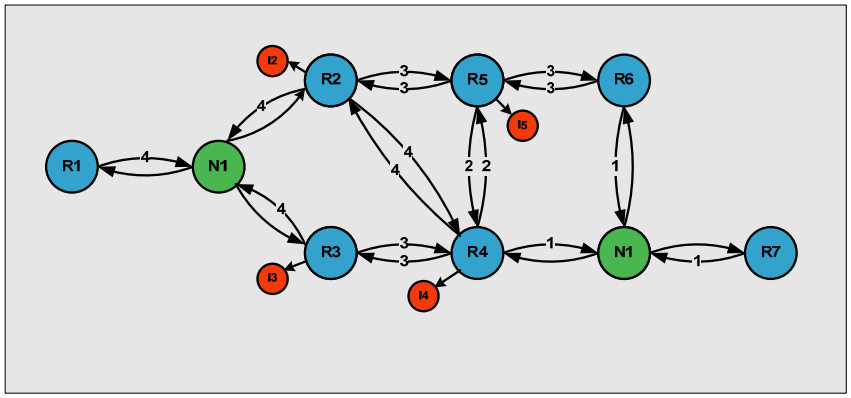

Scenario 4
ANSWER KEY

FOR

CCIE® CANDIDATES

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

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-200v6 Scenario 4: Spot the Issue Answer Key

Table of Contents

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.

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

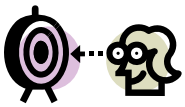

Goals and Restrictions

- IP subnets on the diagram belong to network 151.90.0.0/16.
- Do not use any static routes.
- Advertise Loopback interfaces involved in IGP configurations with their original mask. Do not change the mask.
- Network 0.0.0.0/0 should not appear in any routing table (show ip route).
- All IP addresses involved in this scenario must be reachable, unless specified otherwise.
- Networks advertised in the BGP section must be reachable only in the BGP domain.
- Use conventional routing algorithms.

Explanation of Each of the Goals and Restrictions

IP subnets on the diagram belong to network 151.90.0.0/16 unless specified otherwise.

All IP addresses in this Exam belong to the 151.90.0.0/16 address space with the exception of a set of prefixes used in the BGP section.

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.

Advertise Loopback interfaces involved in IGP configurations with their original masks. Do not change the masks.

An exception is made for interfaces advertised by BGP.

Do not use 0.0.0.0 anywhere in this scenario.

A 0.0.0.0/0 entry may NOT be used to solve reachability problems.

All IP addresses involved in this scenario must be reachable unless specified otherwise.

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 or 0.0.0.0/0 routes.

Networks advertised in the BGP section must be reachable only in the BGP domain.

This restriction relaxes Restriction #3 above. The loopbacks configured for the BGP section need to be reachable only by BGP speakers. They do not have to be reachable from non-BGP speakers, but the routes may be found in the forwarding tables of some non-BGP speakers.

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 INFORMAITON 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

4.1 Frame Relay

HIDDEN ISSUES TO SPOT WITH THE FRAME-RELAY CONFIGURATION

The Frame Relay switch is pre-configured for a full mesh of PVC's. You are instructed to use "PVCs displayed in the diagram **only**". When examining the Scenario 4 diagram, you see two Frame-Relay clouds. Both clouds have only two connections on them. The construction of these two clouds is discussed below.

Issue: You are told to configure logical Frame-Relay interfaces on the 151.90.12.0/24 subnet.

Solution:

You can limit the PVCs used on routers R1 and R2 by configuring a single point-to-point subinterface. All unused DLCIs will remain on the physical interface.

Implementation:

```
R1:
interface Serial0/0.12 point-to-point
 ip address 151.90.12.1 255.255.255.0
frame-relay interface-dlci 102

R2:
interface Serial0/0.12 point-to-point
 ip address 151.90.12.2 255.255.255.0
 frame-relay interface-dlci 201
```

Verification:

Issue the **show frame-relay map** command to verify DLCI mapping:

```
R1#sh frame-relay map
Serial0/0.12 (up): point-to-point dlci, dlci 102(0x66,0x1860), broadcast
          status defined, active
```

Issue: The second Frame-Relay cloud also has just two connections on it. However, we are told to use the physical Frame-Relay interface on R3 and the point-to-point subinterface on R4.

Solution:

On the physical interface on R3 it is recommended to disable frame-relay inverse-arp and assign a Frame-Relay map statement for router R4.

Implementation:

Configure the physical interface on router R3:

```
interface Serial0/0
 ip address 151.90.34.3 255.255.255.0
 encapsulation frame-relay
 no fair-queue
 frame-relay map ip 151.90.34.3 304
 frame-relay map ip 151.90.34.4 304 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
```

Configure the point-to-point subinterface on R4:

```
interface Serial0/0.34 point-to-point
 ip address 151.90.34.4 255.255.255.0
 frame-relay interface-dlci 403
```

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".

4.2 Catalyst Configuration

HIDDEN ISSUES TO SPOT WITH THE CATALYST 3550 SWITCH CONFIGURATION

Like any Catalyst 3550 configuration, you must address the following basic configuration requirements: setting the VTP mode, configuring trunk ports, and statically assigning ports to VLAN's. For a good reference on basic Catalyst 3550 configuration tasks, download the following Tech-Note from the Technical Library on the NetMasterClass web-site: "Performing Basic Configuration Tasks on the Catalyst 3550" Below is a diagram depicting the VLAN configuration for this labs. Most students find it helpful to create a similar drawing or table.

The diagram makes it clear which ports are associated with which VLANs, which ports must be trunk ports, and which VLANs must be allowed on the trunk connecting the switches.

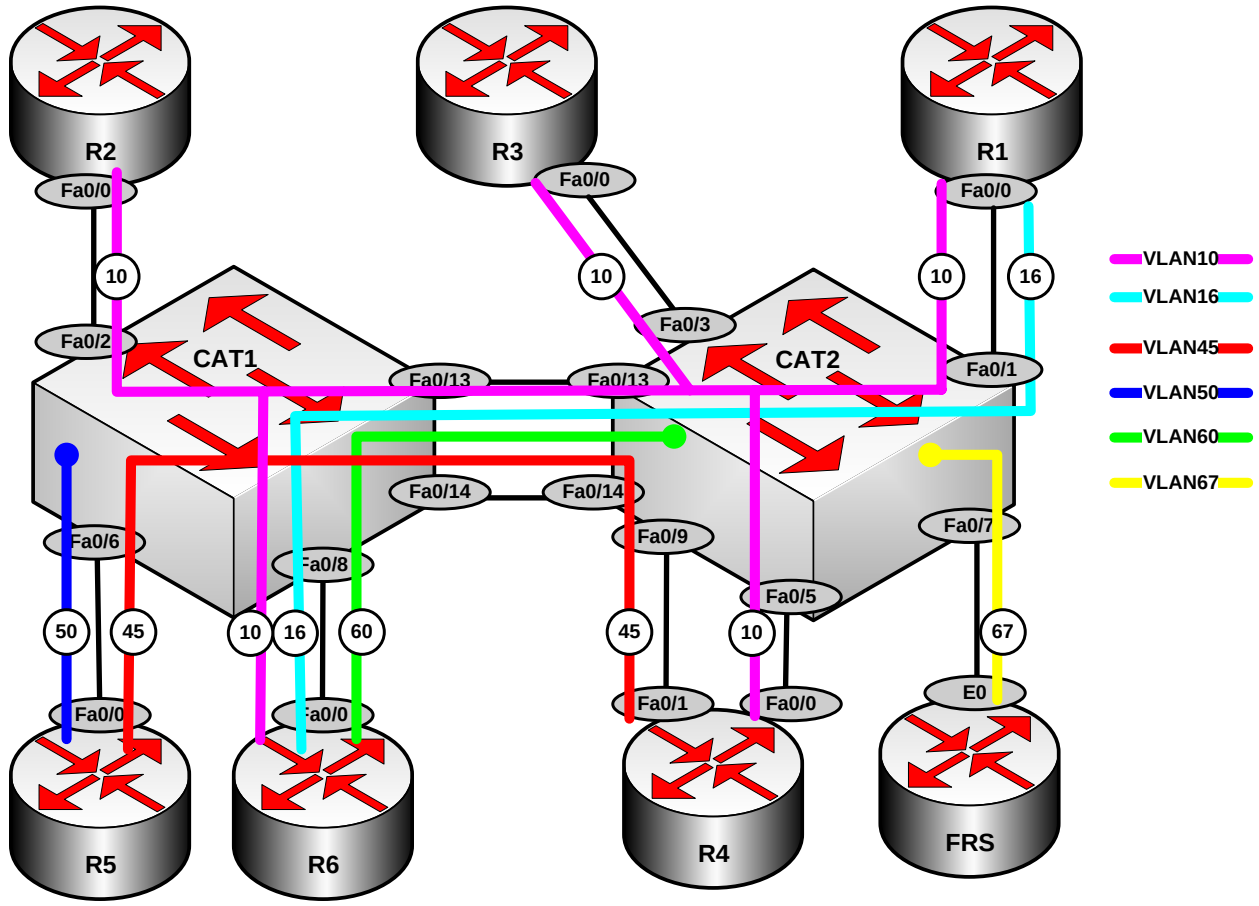

Issue: Configure two different subnets on VLAN 10

Solution:

Since the Scenario has placed multiple subnets on the same VLAN, this implies that there are multiple broadcast domains on the same VLAN. A solution for this type of scenario is deploying a DOT1Q-tunnel. Since R6 has only one FastEthernet interface and since the Scenario diagram requires that R6 have a connection to VLAN 10 as well as VLANs 16 and 60, VLAN 60 must be configured as a native VLAN on router R6 to avoid any type of configuration conflict with the DOT1Q tunnel.

The configuration requirements for this Task can be divided into three categories:

1). Fulfilling the VLAN 10 configuration requirements for the 151.90.163.0/24 subnet
2). Fulfilling the VLAN 10 configuration requirements for the 151.90.24.0/24 subnet
3). Fulfilling the VLAN 60 configuration requirements for the 151.90.60.0/24 subnet

Each of these configuration tasks will be addressed separately. We will begin with the most complex task fulfilling the configuration requirements for the 151.90.163.0/24 subnet. It is the configuration of this

subnet that drives all other configuration requirements in the Section of the scenario. It is this configuration that involves the vast majority of the 802.1Q tunneling requirements.

Implementation:

STEP 1: Configuring VLAN 10 for Subnet 151.90.163.0/24

R1 & R3 (R1 & R3 have identical configurations. The differences are marked by an "X".)

```
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 151.90.163.X 255.255.255.0

(X = host address of R1 and R3)
```

CAT2 Connection to R1 and R3

```
interface FastEthernet0/X
 switchport access vlan 60
 switchport mode dot1q-tunnel
 spanning-tree bpdufilter enable

(X = CAT2 port number attached to routers R1 and R3)
```

R6

```
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 151.90.163.6 255.255.255.0
!
interface FastEthernet0/0.60
 encapsulation dot1Q 60 native
 ip address 151.90.60.6 255.255.255.0
```

CAT1 Connection to R6

```
interface FastEthernet0/8
 switchport access vlan 60
 switchport mode dot1q-tunnel
 spanning-tree bpdufilter enable
```

From the perspective of routers R1, R3 and R6, the 802.1Q tunnel configuration is completely transparent. The routers act as if they are attached to a standard 802.1Q trunk port. In this configuration, they send out their Ethernet frames with an 802.1Q tag associated with VLAN 10.

(Note that R6 possesses FastEthernet subinterfaces associated with VLANs 16 and 60. For the moment, disregard this. It will be explained shortly.)

When a Catalyst 3550 port configured with 802.1Q tunneling receives the 802.1Q frame from the router it takes the received 802.1Q frame and places it into another 802.1Q frame. You now have an 802.1Q frame inside of another 802.1Q frame. Whatever VLAN the switchport is associated with locally – in this case it is VLAN 60 for the configurations above – the frame is switched on that VLAN. Therefore, in this scenario, frames are received from routers R1, R3 and R6 with a VLAN tag of 10 and are then encapsulated in an Ethernet frame associated with VLAN 60. When the frames get switched to their final

destination the 802.1Q tunnel frame associated with VLAN 60 is stripped off and the original 802.1Q frame associated with VLAN 10 is sent to the destination router. Again, the routers in this Scenario do not know that 802.1Q tunneling is configured. You do not need to configure the routers in any special way for 802.1Q tunneling.

On the Catalyst switch side, you need to configure the following two commands at the switchport level:

```
Interface fa 0/X
 switchport access vlan XX
 switchport mode dot1q-tunnel
```

The **switchport mode dot1q-tunnel** command enables 802.1Q tunneling on the switchport. The **switchport access vlan XX** command establishes the VLAN you want the 802.1Q tunnel traffic to be transported on. A point of confusion for many students is: Why have I not configured 802.1Q trunking on this interface? The router connected to this interface is configured as an 802.1Q trunk. ANSWER: You have configured 802.1Q on this interface to communicate with the directly connected 802.1Q trunk interface on the router; however, you have not configured 802.1Q trunking on the Catalyst interface; you have configured 802.1Q tunneling on the interface. In summary, configuring 802.1Q trunking and 802.1Q tunneling on the same Catalyst port cannot be done. They are mutually exclusive. Now let's examine the configurations of the other VLAN 10 subnet 151.90.24.0/24.

STEP 2: Configuring VLAN 10 for Subnet 151.90.24.0/24.

R2 and R4

```
interface Ethernet0
 ip address 151.90.24.X 255.255.255.0
```

CAT1 Connections to R2 and R4

```
interface FastEthernet0/X (X=2 for CAT1 and 5 for CAT2)
 switchport access vlan 10
 switchport mode access
```

As you can see, it is much simpler than the previous VLAN 10 configuration that involved the 802.1Q tunnel. This VLAN 10 configuration is configured with no involvement of 802.1Q at all. It is configured as a standard static VLAN. However, due to the 802.1Q tunnel configuration supplied for the 151.90.163.0/24 subnet, the two IP subnets in this Scenario – 151.90.163.0/24 and 151.90.24.0/24 - which are both configured for VLAN 10, are completely separate. Finally, let's examine VLAN 60 running between R6 and CAT2. The configuration related to this VLAN is supplied below.

STEP 3: Configuring VLAN 60 for Subnet 151.90.60.0/24.

R6

```
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 151.90.163.6 255.255.255.0
!
interface FastEthernet0/0.60
 encapsulation dot1Q 60 native
 ip address 151.90.60.6 255.255.255.0
```

CAT1 Connection to R6

```
interface FastEthernet0/8
 switchport access vlan 60
 switchport mode dot1q-tunnel
```

CAT2 Switched Virtual Interface VLAN 60

```
interface Vlan60
 ip address 151.90.60.20 255.255.255.0
 ip mtu 1500
```

Router R6 has only one physical FastEthernet interface that is connected to VLAN 10, VLAN 16 and VLAN 60 via 802.1Q trunking. R6 receives all VLAN 10 traffic from CAT2 through the 802.1Q trunk. These frames are delivered to the CAT1 fa0/8 port (the port that attaches to the fa0/0 interface of R6) over VLAN 60. When they are transmitted out of this port the 802.1Q tunnel frame associated to VLAN 60 is stripped off and the original 802.1Q frame, tagged for VLAN 10, is delivered to R6.

However, when R6 wants to send a frame out to VLAN 60 on the 151.90.60.0/24 subnet, it does not want that frame to go into the 802.1Q tunnel. Therefore, R6 sets VLAN 60 as the "native" VLAN. Frames sent out the native VLAN do not get tagged as 802.1Q frames. These frames get sent directly onto VLAN 60 with no involvement of the 802.1Q tunnel. These frames will get delivered to the CAT2 VLAN 60 Switched Virtual Interface just as they would on a standard VLAN.

When you configure 802.1Q tunneling, it will add four additional bytes to an Ethernet frame. This adjustment will have an effect on the MTU size of packets traversing the VLAN's involved with 802.1Q tunneling. This change in MTU size will have an effect on forming adjacencies with certain routing protocols, namely OSPF and Integrated IS-IS.

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".

4.3 OSPF

HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

Issue: Make sure R1 sees two next-hops to R6 for external prefixes learned from R6

Solution:

R1 has two paths to external prefixes advertised by R6, one across F0/0.10 and one across f0/0.16. Since R1, R6 and R3 are all on a shared, /24 subnet, by default R6 will advertise routes learned from R3 via EIGRP into OSPF in Type 5 LSAs with R3 as the forwarding address, and R1 will see just this single best

path. For example, R1 will have a next-hop to network 151.90.34.0 of R3's interface address, rather than R6's address. Here is the LSA on R1:

```
R1#sh ip ospf database ext 151.90.34.0

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 266
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 151.90.34.0 (External Network Number )
  Advertising Router: 6.6.6.6
  LS Seq Number: 80000009
  Checksum: 0x5029
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 151.90.163.3
        External Route Tag: 90
```

How can we get R1 to ignore the fact that they are on a shared network? How can we get R6 to send External LSAs with 0.0.0.0 (itself) as the forwarding address? With some IOS you could fool OSPF by using a 0.0.0.0 network mask. But this appears not to work in 12.4. One way would be to have R1 and R6 treat the Ethernet as a point-to-point network. Here is the same LSA after changing the OSPF network type on the subnet 151.90.163.0 interfaces to point-to-point and clearing the ospf process on R6:

```
R1#sh ip ospf database ext 151.90.34.0

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 14
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 151.90.34.0 (External Network Number )
  Advertising Router: 6.6.6.6
  LS Seq Number: 80000002
  Checksum: 0x25F3
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 90
```

The 0.0.0.0 forwarding address indicates the router sourcing the External LSA should be the destination for traffic to the advertised prefix. R1 will then have two equal-cost paths to network 151.90.34.0, both with R6 as the next hop:

```
O E2    151.90.34.0/24 [110/20] via 151.90.163.6, 00:00:46, FastEthernet0/0.10
                       [110/20] via 151.90.16.6, 00:00:46, FastEthernet0/0.16
```

Issue: Make Area 60 a stub area. CAT2 and FRS may have one Inter-Area route and may not have any default route. Make sure CAT2 and FRS can still reach other prefixes in the 151.90.0.0/16 range.

Solution:

ABRs automatically generate a 0.0.0.0 into stub areas. Since this is an inter-area route, we can use the Area Filter-List feature to deny the default and permit a 151.90.0.0/16 summary.

***Issue:* MTU mismatch issue between R6 and CAT2 on VLAN 60.**

Solution:

On some IOS, the 802.1Q tunneling configuration changes the system mtu on the catalyst, causing a mismatch between router R6 and CAT2. You can check this with the command **show system mtu**. You can remedy this problem by manually setting the MTU size on the Catalyst 2 switched virtual interface with the command: "**ip mtu 1500**". Alternatively, you have OSPF ignore the mismatch, by using the command **ip ospf mtu-ignore**.

Implementation: Configure SVI Vlan60 with MTU=1500:

```
interface Vlan60
 ip address 151.90.60.20 255.255.255.0
 ip mtu 1500
end
```

Verification:

Check that OSPF adjacency is established between CAT2 and R6:

```
CAT2#sh ip ospf nei
Neighbor ID     Pri   State           Dead Time   Address         Interface
6.6.6.6           1   FULL/DROTHER    00:00:39    151.90.60.6     Vlan60
```

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".

4.4 RIP

HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

Issue: RIP must exchange updates only on the following links: R2 -R5, R4-R5 and R5-CAT1.

Solution:

Configure **passive-interface default** under the RIP process of each specified router. Then configure **no passive-interface Xy** (X = interface type; y = interface number) for each of the interfaces connected to the link with the specified peer.

Verification:

Issue the show ip protocols | be Passive command and verify that necessary interfaces are not in the list of passive interfaces for RIP.

Issue: Router R5 should send traffic to R4 as a preferred next-hop.

Solution:

Configure the offset-list command under the "router rip" configuration mode on router R5. With the offset-list command, you can manually increase the hop count of RIP routes. Here is a suggested offset-list configuration to enter under the "router rip" configuration mode on R5: "**offset-list 0 in 10 Serial1/0**".

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".

4.5 EIGRP

HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

Issue: R2 and R4 should consider each other as a valid neighbor for a period that is twice longer than the default.

Solution:

This is accomplished by adjusting the EIGRP holdtime. The default EIGRP holdtime is 15 seconds or three times the default EIGRP hello time on a LAN interface. To fulfill the requirement, set the EIGRP holdtime to 30 seconds on the Ethernet interfaces on routers R2 and R4.

Implementation:

Router R2:

```
interface FastEthernet0/0
 ip address 151.90.24.2 255.255.255.0
 ip hold-time eigrp 10 30
```

Router R4:

```
interface FastEthernet0/0
 ip address 151.90.24.4 255.255.255.0
 ip hold-time eigrp 10 30
```

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".

4.6 Redistribution

HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION

Since we are required to do mutual redistribution on all border routers, and alternate physical paths must b e available, all three routing domains will be transit. That is, each domain will import external routes and advertise them to other domains. One way to sort out the potential route feedback, looping and suboptimal path issues is to tag each native prefix as it leaves the domain and then deny these prefixes as they re-enter the domain. On each redistribution point we added a route-map that 1) denies our own prefixes back in, 2) permits prefixes from other domains, 3) and then tags our prefixes as they leave. Seems like a lot of route-maps, and they may not all be strictly necessary. But the route-maps can be created once and reused on other routers. Given the number of loops crossing routing domains, we don't want to take any chances.

Notice the command **ospf distance external 130** configured under the OSPF process on R2. Its purpose is to avoid suboptimal paths to RIP domain routes. Without this command, R2 may well prefer the path through the OSPF and EIGRP domains to reach subnet 151.90.105.0, for example. Now OSPF will act more like EIGRP: very low administrative distance for internal prefixes (110) and very high for external prefixes (130). With external OSPF prefixes rated less trustworthy than RIP prefixes, R2 will take the direct path to 151.19.105.0, yet the longer path will be available if the direct route fails.

Here is diagram showing the routing domains and some of the routing issues:

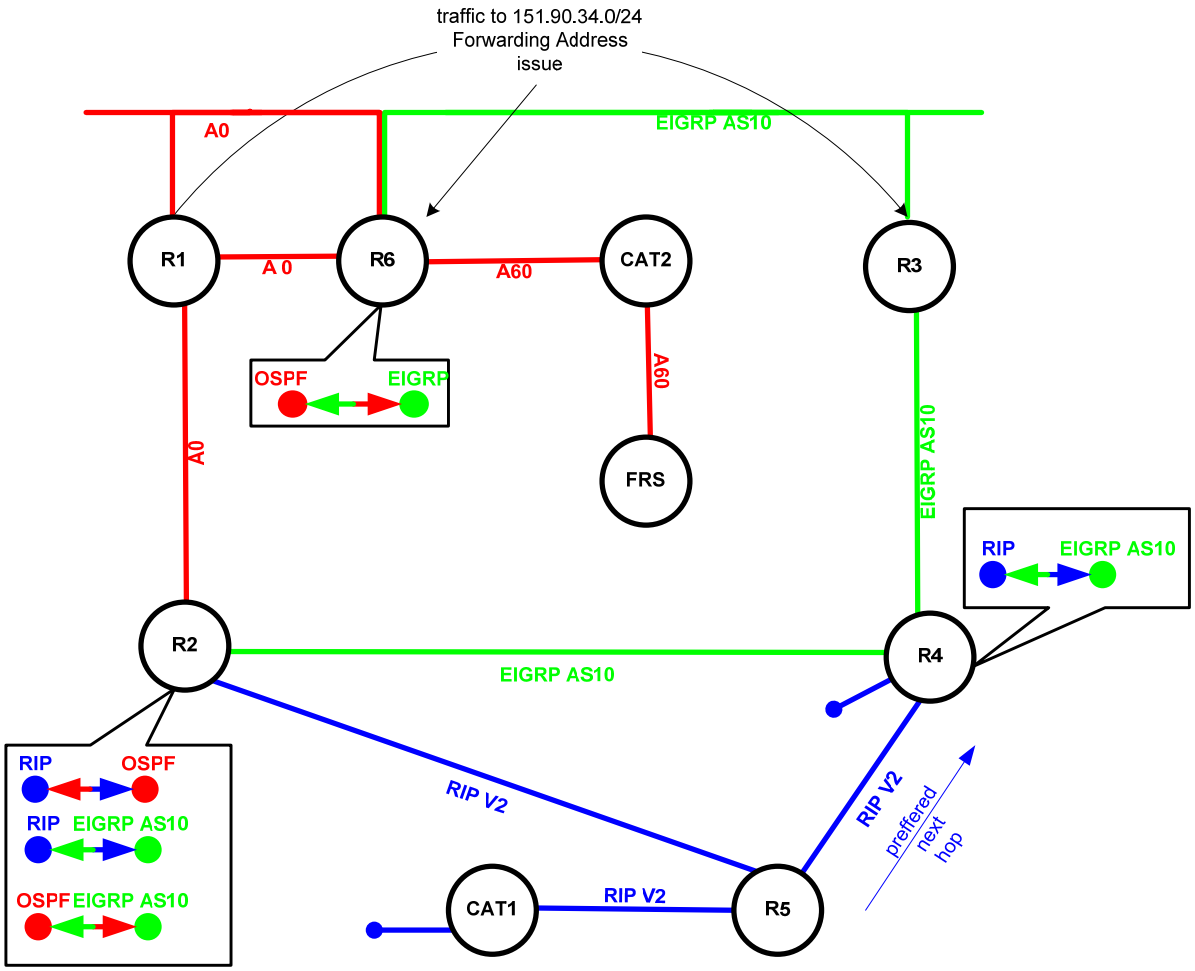

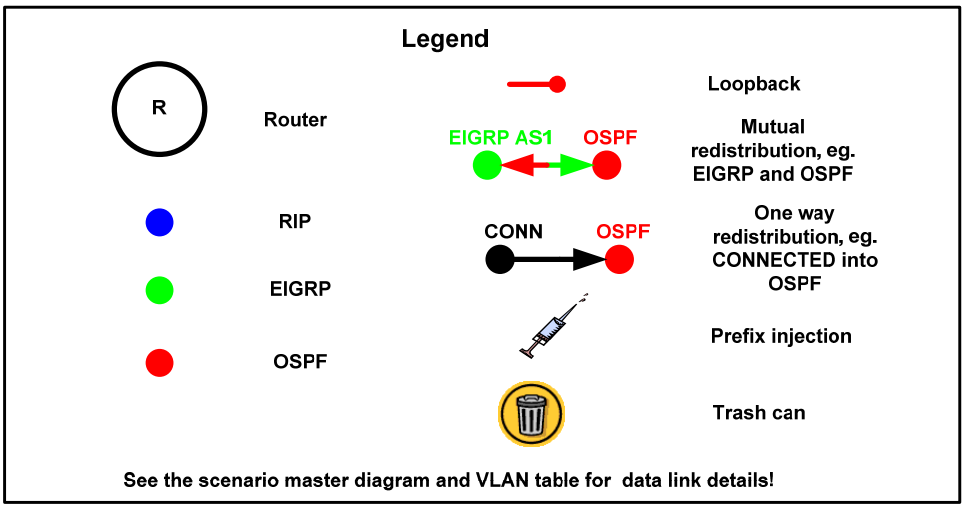

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

This is just one solution to the redistribution challenge that could be correct here. All correct solutions will have at least these three attributes:

1. Provides for universal reachability. Test with a TCL script.
2. Provides stable routes. No route flapping. Test with **debug ip routing**.
3. Provides optimal paths. To some degree, "optimal" is a value judgment. For our purposes, make sure that all paths meet the stated lab requirements, that the routes generally take the shortest paths, and that native prefixes are routed by the native protocols.

Here is a TCL script you can use for this lab. There is a paper on TCL scripting available in the READiT section of the Netmasterclass website. Some addresses are used in later tasks and may not be reachable at this point. Run **tclsh** in privileged mode, paste the script below, and then issue the command **tclq.** All routes may not be reachable at this point. If you do have failures, make sure you understand why they happened. Are the prefixes part of a BGP task? Are they required to be reachable?

```
foreach addr {
151.90.163.1
151.90.101.1
151.90.12.1
151.90.16.1

151.90.102.1
151.90.12.2
151.90.25.2
151.90.24.2

151.90.163.3
151.90.103.1
151.90.34.3

151.90.104.1
151.90.24.4
151.90.45.4
151.90.34.4

151.90.105.1
151.90.25.5
151.90.45.5
151.90.50.5

151.90.163.6
151.90.16.6
151.90.106.1
151.90.60.6

192.168.100.77
192.168.100.91
192.168.100.97
192.168.100.126
151.90.67.7
151.90.107.1
```

```
10.10.10.1
151.90.100.1
151.90.50.10

151.90.67.20
151.90.60.20} {ping $addr}
```

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".

4.7 BGP

HIDDEN ISSUES TO SPOT WITH THE BGP CONFIGURATION

Issue: All routers other than FRS must possess only one prefix best describing the 192.168.X.0/24 networks originating from router FRS.

Solution:

Summarize the set of BGP prefixes on router FRS with an aggregate address of 192.168.100.0/24. Also, use the **summary-only** option with the aggregate command to avoid leaking out the longer matching subnetworks.

Issue: Peer AS400 and AS300 between CAT2 and R3.

Solution:

Since these routers do not share a common subnet, **ebgp-multihop** must be configured on both routers in the **neighbor** statement.

Issue: R5 shoud prefer router R4 as a next-hop to networks originated by AS400.

Solution:

One solution to fulfill this configuration requirement is increasing the administrative weight on router R5's neighbor statement to R4. Administrative weight is a Cisco proprietary feature and (unlike local-preference) influences the path selection process of only the router it is configured on.

Implementation:

The only prefixes in this scenario that R5 could learn from R4 are those sourced from AS 400. We have used the minimal configuration steps by simply adding the following command to R5's BGP configuration:

neighbor 151.90.45.4 weight 40000. Alternatively, we could have created an as-path access-list that specifies prefixes learned from AS400 (_400$), matched this list in a route-map, set the weight, and applied the route-map to the neighbor statement, as shown here:

```
ip as-path access-list 10 permit _400$
!
route-map Set-Weight permit 10
 match as-path 10
 set weight 40000
!
route-map Set-Weight permit 20
```

Issue: Advertise network 10.10.10.0/8 from CAT1. This network should be seen in AS400 with the AS PATH three times longer via R3 than via R6..

Solution:

On router R3, configure a route-map that will perform AS-Path prepending for the 10.0.0.0/8 prefix when it is advertised to CAT2.

The following diagram represents the BGP configuration for this Scenario:

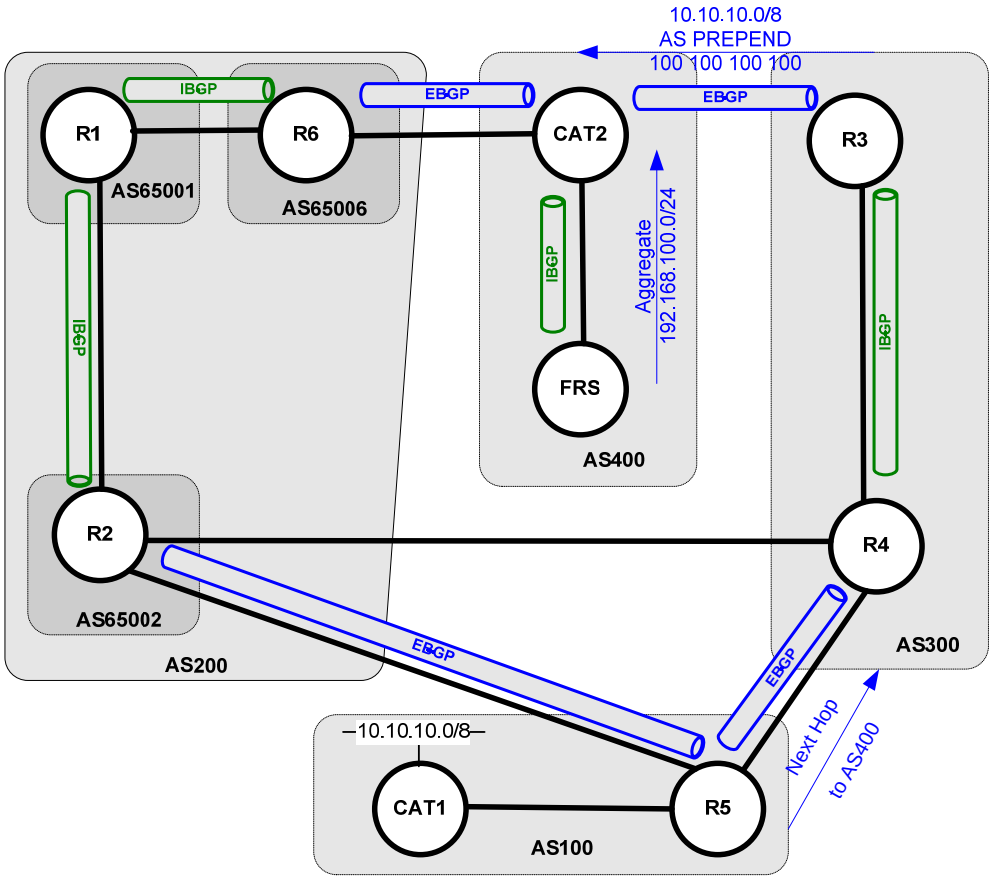

4.8 Network Monitoring

HIDDEN ISSUES TO SPOT WITH THE NETWORK MONITORING CONFIGURATION

Issue: The system Administrator would like to monitor network performance. Specifically, he wants monitor echos to 151.90.102.1 from R1. The collection interval should be every thirty seconds. Collection should continue indefinitely.

Solution:

This task can be completed using the IP SLA feature. First, create a monitor:

```
ip sla monitor 10
 type echo protocol ipIcmpEcho 151.90.102.1
 frequency 30
```

Then, schedule it with a lifetime of "forever" and a start-time of "now:"

```
ip sla monitor schedule 10 life forever start-time now
```

Verification:

The IP SLA feature provides many verification commands, for example:

```
R1#sh ip sla monitor statistics detail
Round trip time (RTT)   Index 10
        Latest RTT: 35 ms
Latest operation start time: *17:39:34.619 UTC Wed Mar 6 2002
Latest operation return code: OK
Over thresholds occurred: FALSE
Number of successes: 39
Number of failures: 1
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never

R1#sh ip sla monitor collection-statistics
Entry number: 10
Start Time Index: *17:20:04.659 UTC Wed Mar 6 2002
Number of successful operations: 40
Number of operations over threshold: 0
Number of failed operations due to a Disconnect: 0
Number of failed operations due to a Timeout: 1
Number of failed operations due to a Busy: 0
Number of failed operations due to a No Connection: 0
Number of failed operations due to an Internal Error: 0
Number of failed operations due to a Sequence Error: 0
Number of failed operations due to a Verify Error: 0
RTT Values:
RTTAvg: 35      RTTMin: 33      RTTMax: 40
NumOfRTT: 40    RTTSum: 1405    RTTSum2: 49381
```

Issue: If the test times out, then R1 should send an SNMP trap to an imaginary management station at 151.90.106.2. .

Solution:

This task requires an IP SLA reaction-configuration. It also will require you to configure snmp and allow snmp traps.

```
ip sla monitor reaction-configuration 10 timeout-enable action-type trapOnly

snmp-server community nmc RO
snmp-server community ccie RW
snmp-server enable traps snmp
snmp-server host 151.90.106.2 password
```

Verification:

When you shut the loopback102 on R2, the IP SLA monitor should detect the timeout and trigger a trap. You can test whether the trap is sent by turning on **debug snmp**.

```
R1#sh ip sla mon statis det
Round trip time (RTT)   Index 10
        Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: *17:24:34.619 UTC Wed Mar 6 2002
Latest operation return code: Timeout
Over thresholds occurred: FALSE
Number of successes: 9
Number of failures: 1
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never


R1#
00:31:50: SNMP: Queuing packet to 151.90.106.2
00:31:50: SNMP: V1 Trap, ent rttMonNotificationsPrefix, addr 151.90.163.1, gentrap 6, spectrap 2
 rttMonCtrlAdminTag.10 =
 rttMonHistoryCollectionAddress.10 = 97 5A  66 01
 rttMonCtrlOperTimeoutOccurred.10 =
```

When connectivity to 151.90.102.1 is restored, another trap will be sent.

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".

4.9 Security

HIDDEN ISSUES TO SPOT WITH THE SECURITY CONFIGURATION

Issue: **Under normal network conditions R5 must accept packets sourced from AS400 networks advertised in a previous task only if the packets arrive via R4. Do not use the command "ip access-group" to accomplish this task.**

Solution:

This task can be fulfilled using the **ip verify unicast reverse-path** interface configuration command. By configuring this command, you are applying a process known as "uRPF". This interface configuration command compares the source address of incoming packets to information listed in the local routing table to determine if the packet should have been received on the interface it was actually received on.

uRPF is based upon the principle that an interface that is used to forward a packet to a given destination address should be the same interface used when a packet is received from the previously mentioned destination address.

uRPF check is successful (i.e. packet is forwarded) if the packet is received on the interface where the local routing table points to for the source ip address in the packet.

If a packet fails unicast RPF check an optional ACL is checked to see if the packet should be dropped (using a deny statement in the ACL) or forwarded (using a permit statement in the ACL). If no ACL is specified in the uRPF command, the router drops the packet immediately.

The ACL can be configured in the form of :

- standard ACL – the source IP address in the packet is matched against ACL entries, or
- extended ACL – the packet is matched as usual, but ACL entries must specify protocol 'ip' and destination of 'any' (never versions of IOS may remove this restriction).

Logging can be used with the ACL.

uRPF requires the activation of CEF processing (issue the **ip cef** command).

Under normal network conditions R5 would have a route for AS400 pointing to R4:

```
------------------------------R5-------------------------------
R5#show ip route | i 192.168.100.0
B    192.168.100.0/24 [20/0] via 151.90.45.4, 11:25:52
```

Configuration of uRPF on R5↔R2 interface

```
interface Serial1/0
 ip address 151.90.25.5 255.255.255.0
 ip verify unicast reverse-path 107
 !
access-list 107 permit ip 192.168.100.64 0.0.0.63 any log
access-list 107 permit ip any any
```

will instruct the router to log packets from AS400 networks if they arrive via R2. Note that the ACL is

constructed to permit all packets not sourced from AS400 even if they fail uRPF check.

Issue: **In case of network failure resulting in R5 not receiving BGP routes for AS400 networks from R4, R5 should accept packets sourced from AS400 networks only if they arrive via R2. Do not use "ip access-group" to accomplish this task.**

Solution:

Solution to this requirement is similar to the one described above.

As the R5 route after such failure would point to R2, configure uRPF on R5↔R4 interface.

```
interface FastEthernet0/0.50
 encapsulation dot1Q 50
 ip address 151.90.45.5 255.255.255.0
 ip verify unicast reverse-path 107
!
access-list 107 deny   ip 192.168.100.64 0.0.0.63 any
access-list 107 permit ip any any
```

Note that both solutions work very nicely together. For example the uRPF configured to satisfy requirements of 4.12.2 does not prevent traffic forwarding in normal network conditions, as in such case uRPF check would not fail. This behavior of uRPF check that it follows the local routing table is the essence of the feature.

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".

4.10 IPv6

Issue: **Configure IPv6 on VLAN 10 between R1, R5 and R3; on VLAN 10 between R2 and R4; on Frame-Relay links between R1 and R2, R3 and R4 and over Serial link between R2 and R5.**

Solution:

IPv6 configuration on routers is enabled by the **ipv6 unicast** global configuration command. After that, use the interface-level command **ipv6 address X:X:X::X/X** will configure IPv6 address on that interface:

```
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 151.90.163.1 255.255.255.0
 ipv6 address FEC0::163:1/125
```

Interface S0/0 on R3 is a physical Frame-Relay interface, and will require static maps to both the FEC0::34:4 global and FE80::4 link-local addresses. We have you hard-code the link-local address on

R4's S0/0.34 interface to make this mapping easier. Note that just this one interface will need a static map; all others are logical point-to-point interfaces.

Issue: **Configure RIP for IPv6 process on all IPv6 links using id "RIPV6". Make sure you can ping all IPv6 interfaces from all IPv6 routers.**

Solution:

RIPv6 configuration is done on the interface level using the **"ipv6 rip *ProcessID* enable"** command. Here is an example on R1. Note that the interface-level command automatically creates the global process.

```
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 151.90.163.1 255.255.255.0
 ipv6 address FEC0::163:1/125
 ipv6 rip RIPv6 enable
```

You can verify IPv6 connectivity with a TCL script like the following:

```
foreach addr {
fec0::163:1
fec0::12:1
fec0::24:2
fec0::12:2
fec0::163:3
fec0::34:3
fec0::34:4
fec0::24:4
fec0::25:5
fec0::163:6
} {ping $addr}
```

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".

4.11 QOS

HIDDEN ISSUES TO SPOT WITH THE QUALITY OF SERVICE CONFIGURATION

Issue: **Prioritize traffic sent via R2 to router R5 based upon a specified listing of traffic types (DNS, HTTP, etc.).**

Solution:

When examining the queueing requirements listed in the task, notice that there are four general categories of traffic (1) very critical (2) critical (3) medium and (4) normal. These 4 classifications correspond to the 4 categories provided by priority queuing. What is listed as "very critical" place in the high priority queue. What is listed as "critical" place in the medium priority queue. What is listed as "medium" in place in the normal priority queue. What is listed as "normal" place in the low priority queue.

Implementation:

Configure priority-queueing on R5. Group 2 is for protocol-specific queueing, and Group 4 is for everything coming from FastEthernet0/0:

```
priority-list 2 protocol ip high udp domain
priority-list 2 protocol ip high tcp domain
priority-list 2 protocol ip high udp ntp
priority-list 2 protocol ip medium tcp www
priority-list 2 protocol ip normal tcp smtp
priority-list 2 protocol ip low tcp ftp
priority-list 2 protocol ip low tcp ftp-data

priority-list 4 interface FastEthernet0/0 medium
```

Apply these groups to the interfaces:

```
interface FastEthernet0/0
priority-group 4
!
interface Serial1/0
priority-group 2
```

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".

4.12 Address Administration

HIDDEN ISSUES TO SPOT WITH THE ADDRESS ADMINISTRATION CONFIGURATION

Issue: **The 151.90.100.0/24 network should be in the routing table of ONLY routers CAT1 and R5. Make sure you can ping the rest of the network from 151.90.100.1.**

Solution:

Configuring NAT on router R5 can fulfill these requirements. From the wording of the tasks, it should be clear that the 151.90.100.0/24 prefix should only be known by CAT1 and R5. Distribute list is applied on R5 towards the rest of the network to block 151.90.100.0/24. Therefore the 151.90.100.0/24 prefix will be advertised to router R5 and no other. On R5, NAT will be configured with the FastEthernet interface as the NAT inside interface and the two interfaces on R5 going towards R2 and R4 as the NAT outside interfaces. Since the address to be translated is an inside address that will not be known to any other router outside of the RIP domain (consisting of only CAT1 and R5), a standard NAT configuration should be configured on R5 using an **ip nat pool** statement and a **ip nat inside source list** statement. Since there is only one IP address to translate, the nat pool created can consist of only a single address. As with any NAT configuration, make sure the address you select to use in your NAT pool is reachable by outside routers.

Implementation:

Configure the NAT pool, and NAT translations:

```
ip nat pool NAT 151.90.50.11 151.90.50.11 netmask 255.255.255.0
ip nat inside source list NAT pool NAT overload
!
ip access-list standard NAT
 permit 151.90.100.1
```

Configure interfaces as nat inside and nat outside:

```
interface FastEthernet0/0
 ip nat inside

interface FastEthernet0/0.50
 ip nat outside
```

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".

4.13 Multicast

HIDDEN ISSUES TO SPOT WITH THE MULTICASTING CONFIGURATION

Issue: **An IPTV server is connected on VLAN 50. The server is streaming 64 programs. Configure the link between routers R2 and R5 for the minimal bandwidth consumption for audio streams.**

Solution:

You are directed to configure rtp header compression on the PPP link between routers R2 and R5. However, when rtp header-compression is configured with its default settings, it supports only 32 compression connections. In order to support the number of connections specified, you must enter the following interface configuration command: **ip rtp compression-connections XXX** (where XXX is the number of connections you want to support). Set the number of rtp connections to 64, as follows:

```
ip tcp header-compression iphc-format
ip rtp header-compression iphc-format
ip rtp compression-connections 64
```

The following diagram represents the multicast configuration for this Scenario.

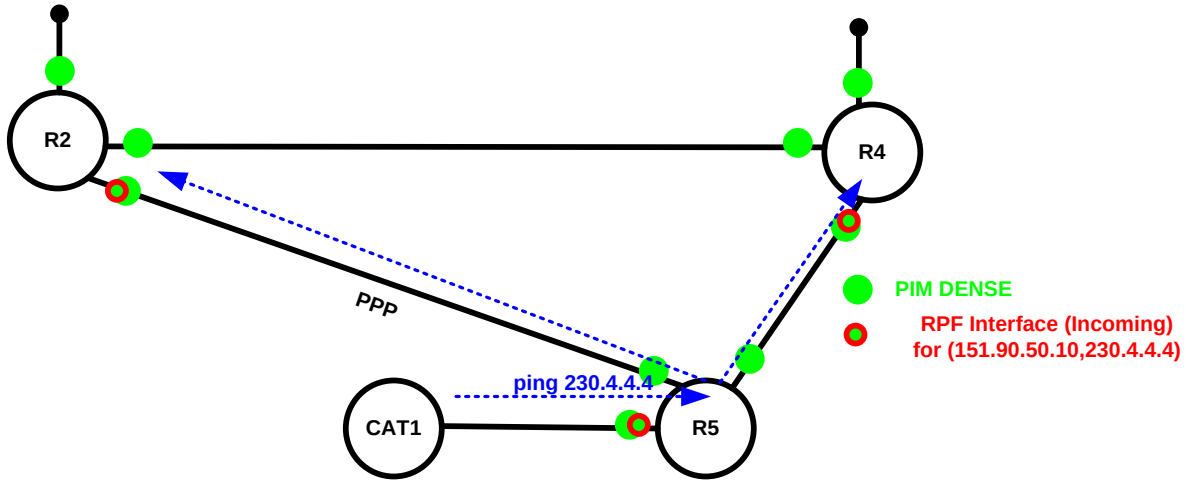

To obtain a comprehensive view of the configuration tasks in this section, access the SHOWiT engine. With the SHOWiT engine, you can enter in over 1000 IOS commands as well a collection of NMC proprietary commands such as "show all".