NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II

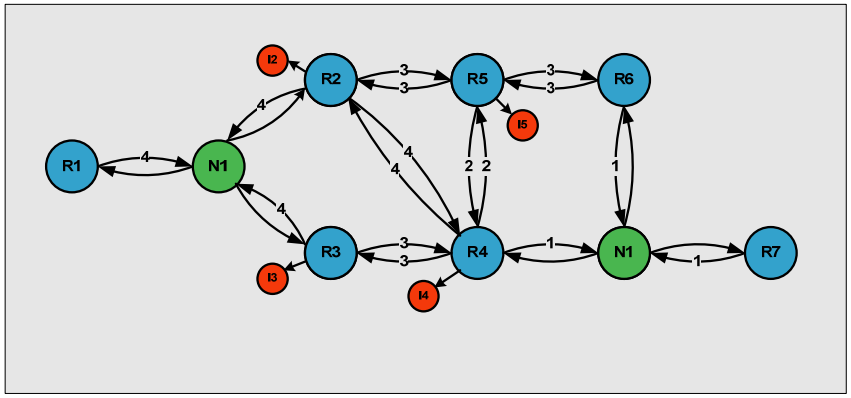

Scenario 15
ANSWER KEY

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.

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

DOiT-V6 Scenario 15: 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.

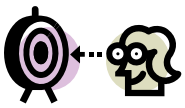

Goals and Restrictions

- IP subnets on the diagram belong to network 160.20.0.0/16.
- Do not use any static routes.
- Advertise Loopback interfaces with their original masks.
- Do not use ip default-network.
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks 192.50.*.* are excluded from the previous requirement.
- Networks advertised in the BGP section must be reachable only in the BGP domain.

Explanation of Each of the Goals and Restrictions:

IP subnets in the Scenario diagram belong to network 160.20.0.0/16

The third and forth octets of the IP addresses displayed on the diagram belong to 160.20.0.0/16. This is specified to make sense of the two octets supplied in the diagram.

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 with their original masks.

By having OSPF routers advertise their loopbacks as host routes you can solve certain VLSM/FLSM problems, since all RIP versions redistribute /32 routes. This solution is not permitted.

Do not use ip default-network

The ip default-network command can be used to solve a range of reachability problems. In particular, the default-network command can be used to solve redistribution problems, however in this exercise, you cannot use it. A suggested alternative is route summarization.

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

This restriction relaxes the previous restriction. The BGP originated networks need only be reachable from the BGP speaking routers.

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

15.1 Frame Relay

HIDDEN ISSUES TO SPOT WITH THE FRAME RELAY CONFIGURATION

Issue: Use only the PVCs displayed on the diagram.

Solution:

1. Disable inverse arp so that no undesirable dynamic inverse-arp entries are found on any of the routers. Unneeded inverse-arp entries could cause you to use more PVCs than are displayed in the diagram.
2. Provide static frame-relay mappings on each of the Frame-Relay attached routers. Make sure the frame-relay map statements reflect the topology in the scenario diagram.

Issue: Configure logical interfaces on the subnet 160.20.123.0/24, use point-to-point interfaces on the spokes.

Solution:

The 160.20.123.0/24 subnet is made up of a hub and spoke NBMA topology with R1 as the hub. To fulfill the logical interfaces requirement, you need to configure a multipoint subinterface on router R1 and two point-to-point subinterfaces on routers R2 and R3.

Verification:

Use the command **show Frame-Relay pvc** to verify that you are only using the permitted DLCIs. Note in the output below for R3 that only DLCI 301 has status Active Local. This should be the result when you turn off inverse-arp, statically map and then no-shut the interface. Some IOS versions will require a reboot.

```
R3#sh frame-relay pvc

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

                Active      Inactive      Deleted       Static
  Local            1             0            0            0
  Switched         0             0            0            0
  Unused           2             0            0            0

DLCI = 301, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0/0.123
R3#
```

Issue: Enable Frame-Relay between R2 and R5. Use physical interfaces on this link.

Solution:

First step in configuration of back-to-back connection is to configure clock on the physical interface. You will need to determine which end is the DCE connection from the synchronous serial interface perspective

so that you can configure a clock rate on it. You can do this with the command **show controllers serial N** where N is the serial interface number. Remember a space is required between the key word serial and the serial interface number. In the **show controllers** display, you will see which interface is the DCE connection and which is the DTE connection. Then configure **clock *rate*** command on the DCE side. Note that supported clock rate values are HW dependent; IOS will list the possible rates in response to a "?" when command is being entered.

The Frame-Relay connection between routers R2 and R5 is a back-to-back Frame-Relay connection with no Frame-Relay switch in between. There are 2 ways to configure back-to-back frame relay.

Less common method is to disable LMI on both routers.

http://www.cisco.com/warp/public/125/frbacktoback.html

More commonly used method which will use is to configure one of the routers to act as Frame Relay DCE. For consistency's sake, it is recommended to make both physical and Frame Relay DCE to be configured on the same router.

http://www.cisco.com/warp/public/125/frbacktoback_hybrid.html

You will configure one of the routers to act as the Frame-Relay switch and the other router to act as the Frame-Relay DCE device or a 'a hybrid FR switch'. On the router that is configured as the Frame-Relay switch, you will enter the global configuration command **frame-relay switching** and the interface command **frame-relay intf-type dce** as well as your usual interface commands relevant to Frame Relay configuration such as **encapsulation frame-relay**.

You will not need any Frame Relay switching commands (i.e. **frame-relay route**) commands, since you are not switching Frame-Relay traffic from one Frame-Relay switched interface to another. Instead, use the **frame-relay interface-dlci XXX** command to advertise a DLCI from the router acting as the Frame-Relay switch to the router acting as the Frame-Relay DTE device. Optionally **frame-relay map** can also be used. With this type of configuration both sides possess the same DLCI. Usual Frame Relay configuration rules apply with such configuration, i.e. inverse mapping, traffic shaping, etc.

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".

15.2 Catalyst Configuration

HIDDEN ISSUES TO SPOT WITH THE CATALYST 3550 CONFIGURATION

Issue: Basic VLAN and Trunk creation

Solution:

In addition to the table provided in the scenario, and the provided diagrams, it might help to create a drawing similar to the following one. It makes it very clear which switch ports are access, which are trunks, and which VLANs must be permitted across which trunks.

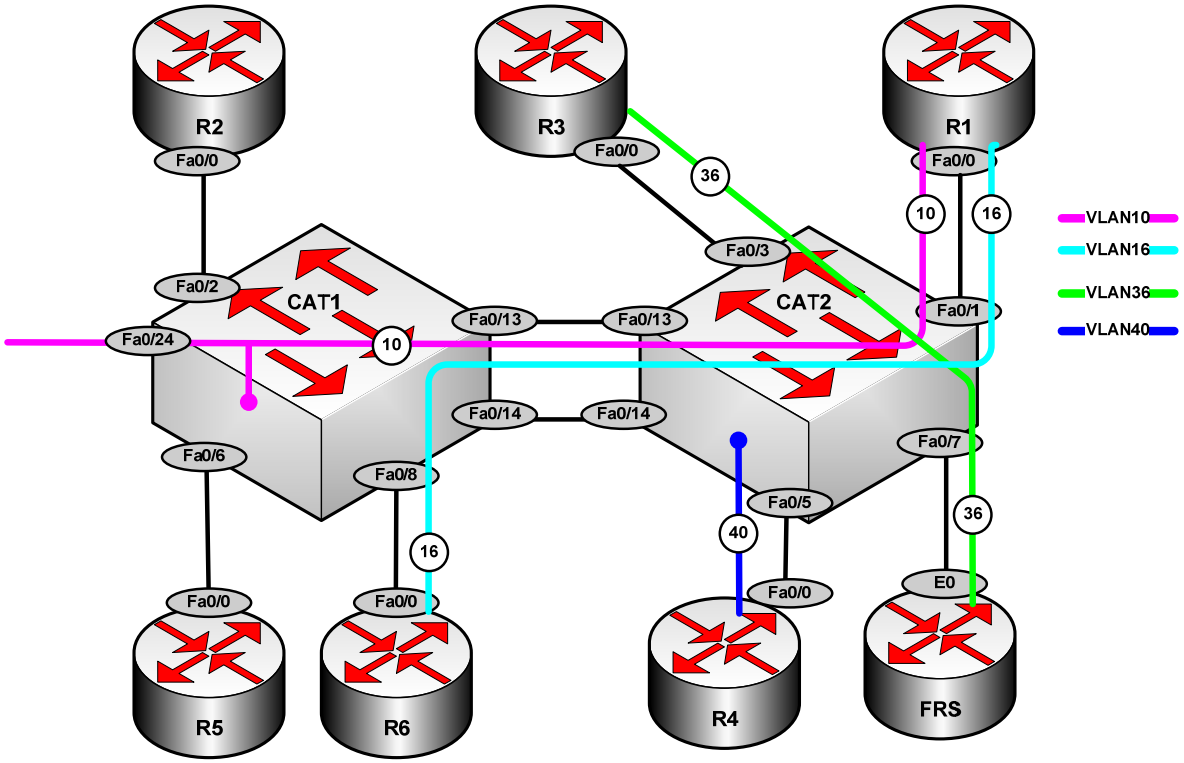

Issue: Use a trunk protocol that does not support native VLANs

Solution:

Frames on the native VLAN are not tagged with a VLAN ID. The dot1q trunking protocol supports a native VLAN, but the ISL protocol tags all frames with a VLAN ID. This requirement points to the use of ISL for your trunking requirements in this Scenario. ISL is the default trunking protocol on the Catalyst 3550.

Issue: Make sure only VLANs used in this scenario are permitted on the configured trunks.

Solution:

You control the VLANs permitted on trunks with the command **switchport trunk allowed vlans XXX**, where XXX is a list or of VLAN's to permit. It also supports the keywords add, all, remove, none and except. By default, all VLANs are allowed. Only ports F0/13 on CAT1 and CAT2 form a trunk.

Verification:

Below is output from the command **show interfaces trunk** on CAT1. Note that the encapsulation protocol in use is ISL. This output helps us verify that only the VLANs 10 and 16 are allowed across trunk ports F0/13 and F0/2.

```
CAT2#sh int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/1       on           isl            trunking      1
Fa0/13      desirable    isl            trunking      1

Port      Vlans allowed on trunk
Fa0/1       10,16
Fa0/13      10,16
CAT2#
```

Issue: Use interfaces F0/14 on subnet 160.20.17.0/24. Do not create a VLAN on this subnet Issue:

Solution:

Catalyst physical ports are Layer 2, switched ports by default. However each of them can become a routed, layer 3 port by simply entering the command **no switchport**.

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".

15.3 OSPF

HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

Issue: Configure OSPF over Frame-Relay between routers R1, R2 and R3. Use the OSPF non-broadcast network type. Make the Frame-Relay network the OSPF backbone area.

Solution:

The OSPF non-broadcast network type models the hub and spoke Frame-Relay topology as a multi-access network. A DR is elected to source the network LSA and to reduce flooding of LSAs. Each router on the link must form an adjacency with the DR. Since all OSPF packets have a TTL of 1, spoke routers will not be able to form adjacencies with a DR that is on another spoke. With a hub and spoke Frame-Relay topology, therefore, the DR needs to be at the hub, and spoke routers should not be allowed to become either a DR or BDR. To assure that this never happens, set the OPSF priority to 0 at the interface level on the spoke routers. The OSPF non-broadcast network type requires neighbor statements so that OSPF packets can be unicast rather than multicast. At Frame-Relay hub router R1, enter two neighbor statements, one for R2 and the second for R3.

Issue: On R3, create three loopback interfaces with the following subnets and place them into OSPF Area 4. Summarize the three entries with the most efficient mask

- 160.20.163.0/24
- 160.20.169.0/24
- 160.20.174.0/24

Solution:

With OSPF you possess two summarization tools: the area range command and the summary address command. The area range command is used to summarize between OSPF areas and makes that router an Area Border Router (ABR). The summary address command is used to summarize routes from outside of OSPF, and configuring it makes that router an Autonomous System Border Router (ASBR). Since the prefixes to be summarized originate from an OSPF area, the area range command is the appropriate summarization tool. The most efficient mask is the one that summarizes the required addresses and as few others as possible. When determining what the most efficient mask length is, carefully examine the bit boundaries that the addresses fall within. In this case, the most efficient match for the three subnets supplied above is configured as:

```
area 4 range 160.20.160.0 255.255.240.0
```

Issue: Configure OSPF area 10 on the link between R1 and R6. Make sure that the routers in this OSPF area possess the minimum amount of routing information to reach all destinations within your pod.

Solution:

To minimize the amount of routing information in a non-area 0 OPSF area, configure the OSPF stub area feature. An OSPF stub area restricts External routes from entering the designated area. An OSPF totally stubby area restricts both external and inter-area OSPF routes from entering an area. When you configure a non-zero OSPF area as a stub area, you need to configure all routers in the stub area with the stub area keyword as well. If you want to configure an area as a Totally Stubby area, you must still configure all routers within the stub area with the **area X stub command**. On the router that is the ABR for the Totally Stubby area, enter the command **area X stub no-summary** under the OSPF routing process.

Issue: Configure OSPF area 25 on the link between R2 and R5. Advertise the loopback 160.20.105.1/24 in OSPF area 105.

Solution:

Since router R5 has no direct connection to Area 0, you will need to configure a virtual-link through area 25 to allow Area 105 to be accessible to the rest of the OSPF network.

Verification:

To verify that your virtual link is truly operational, look for a neighbor over the virtual link, and look for the string “adjacency state full” in the output of **show ip ospf virtual-link**, as shown below:

```
R2#sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address         Interface
160.20.105.1      0   FULL/  -          -         160.20.25.5     OSPF_VL3
160.20.101.1      1   FULL/DR         00:01:48    160.20.123.1    Serial0/0.123
160.20.105.1      0   FULL/  -        00:00:34    160.20.25.5     Serial1/0
R2#

R2#sh ip ospf virtual-links
Virtual Link OSPF_VL3 to router 160.20.105.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 25, via interface Serial1/0, Cost of using 781
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
    Adjacency State FULL (Hello suppressed)
    Index 2/3, 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
R2#
```

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”.

15.4 RIP

HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

Issue: Configure RIP version 2 update exchange over the Frame-Relay connection between R1 and R4 and between R4 and CAT2. Make sure that RIP advertises only over the specified links.

Solution:

Whenever you configure RIP over Frame-Relay, check to see if there is a split-horizon issue by observing the debug of a few updates. Remember that split-horizon is disabled by default on physical Frame-Relay interfaces, like R4's S0/0 interface. There is no reason to have split-horizon disabled on the physical interface on R4, and there is a potential for feedback, between routers R1 and R4. It is recommended that you enable split-horizon on the Frame-Relay interface on R4. As a general practice, take a few minutes to run **debug ip rip**, and make sure the updates make sense!

Issue: Make sure that RIP advertises only over the specified links.

Solution:

RIP version 2 is enabled under the RIP routing process. You can control what interfaces that RIP updates get advertised out of with the passive-interface command. A recommended general practice to apply to a standard RIP configuration is to enter the **passive-interface default command** under the router rip configuration mode and then enter **no passive-interface Xy** where Xy is the specific interface you want RIP updates to be advertised on. You can verify both the RIP version and passive status with the **show ip protocols** command.

Issue: On CAT2 configure the following addresses on one loopback interface and advertise them into RIP without using a network statement

Solution:

This task calls for redistribution of connected routes. Note that a route-map referencing an interface also includes secondary addresses on the interface.

Issue: Use a single filtering statement to permit only the following RIP updates learned by R4 from CAT2:

- 192.50.152.0
- 192.50.153.0
- 192.50.154.0
- 192.50.155.0

Solution:

This filtering task requires you to determine the bit boundaries that this range of addresses falls between, which in this case is the /22 bit boundary. The range of addresses for this bit boundary in the third octet is: 192.50.152.0 -192.50.155.0, which is precisely the range supplied above. Therefore, to permit only this range of prefixes create the access-list **access-list 1 permit 192.50.152.0 0.0.3.0** and apply it with the command **distribute-list 1 in F0/0** under the RIP process on R4.

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".

15.5 EIGRP

HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

Issue: CAT1 should accept only networks 192.50.157.0 and 192.50.159.0 from CAT2. Implement this requirement using a single-line access-list.

Solution:

You can fulfill this requirement by applying a distribute-list under the EIGRP process on CAT1. The challenge is to match the required networks with an access list that has the fewest possible number of statements. The following single-line access-list should work:

```
access-list 1 permit 192.50.157.0 0.0.2.0
```

This can be seen more clearly when the two network addresses are compared in binary form. The only difference is the value of the "2" place in the third octet.

```
192.50.157.0       11000000.00110010.10011101.00000000
192.50.159.0       11000000.00110010.10011111.00000000
Mask (XOR)         00000000.00000000.00000010.00000000
```

As with an access-list, the number of bits you set to one in the wildcard mask will determine the number of possible matching combinations. In the access-list above only one bit is set to one: the seventh bit in the third octet. All other bits are set to zero. It will in fact only match 192.50.157.0 and 192.50.159.0.

Verification:

Is CAT1 really learning only these two prefixes from CAT2? To verify you can run **show ip route** with a "pipe" that matches on CAT2 as a next hop:

```
CAT1#sh ip route | include 160.20.17.20
D EX 192.50.157.0/24 [170/156160] via 160.20.17.20, 04:23:52, FastEthernet0/14
D EX 192.50.159.0/24 [170/156160] via 160.20.17.20, 04:23:52, FastEthernet0/14
CAT1#
```

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".

15.6 Redistribution

HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION

Redistribution can become complex to the extent that physical loops cross routing domains and multiple paths are required. There is a physical loop that crosses the RIP and EIGRP domains, but we are explicitly directed not to do redistribution between them on CAT2. So in essence we have a topology with OSPF as a core and RIP and EIGRP as stub domains.

We are not permitted to redistribute into the RIP domains, and they are not allowed to have a default route. Our solution is to send 160.20.0.0/16 summaries from R1 and from R3. This works here because all of the networks that MUST be reachable are from this major network.

Does this solution provide full reachability, stability and optimality? We can test reachability using the TCL script shown at the end of this section. TCL scripting is supported on R1, R2, R3, R4, R5 and R6. We can test stability by entering **debug ip routing** on each of our routers. If route feedback is causing routes to be periodically installed and withdrawn, the output of this debug will reveal it. To test optimality, we can use the commands **sh ip route <protocol>** and **sh ip route | include <interface>**.

Here is a suggested TCL script to test reachability. Note that this list does not include the prefixes involved in the BGP task.

```
foreach address {
160.20.14.1
160.20.123.1
160.20.16.1
160.20.101.1
160.20.10.1
160.20.123.2
160.20.25.2
160.20.102.1
160.20.36.3
160.20.123.3
160.20.103.1
160.20.163.1
160.20.169.1
160.20.174.1
160.20.40.4
```

```
160.20.14.4
160.20.104.1
160.20.25.5
160.20.105.1
160.20.106.1
160.20.36.6
160.20.107.1
160.20.10.10
160.20.17.10
160.20.40.20
160.20.17.20
192.50.152.1
} {ping $address}
```

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".

15.7 BGP

HIDDEN ISSUES TO SPOT WITH THE BGP CONFIGURATION

Issue: Make R1, R2, R3, R4, CAT1 and CAT2 BGP speakers within AS 100. Do not allow a full mesh of IBGP speakers within AS 10. Provide redundant NLRI exchange using R2 and R3 routers.

Solution:

You have three options when configuring IBGP neighbor relationships: (1) Configure a full-mesh of IBGP speakers (2) configure a route-reflector or a collection of route-reflectors within an AS and (3) configure a confederation of private AS's. Route-reflectors can be combined with confederations. In this configuration, you are instructed not to allow a full mesh of IBGP speakers within AS 10. This requirement can be fulfilled with either a route-reflector or confederation. However, the requirement to provide redundant NLRI exchange suggests route-reflectors as the solution. A redundant route-reflector can provide the redundant NLRI exchange.

Both R2 and R3 will be configured as route-reflectors with the remaining routers in AS 10 being route-reflector clients to both R2 and R3. Both routers R2 and R3 will be configured with an identical BGP cluster-id for loop avoidance purposes. Remember you must enter the BGP cluster-id before you enter the route-reflector-client commands.

BGP Topology

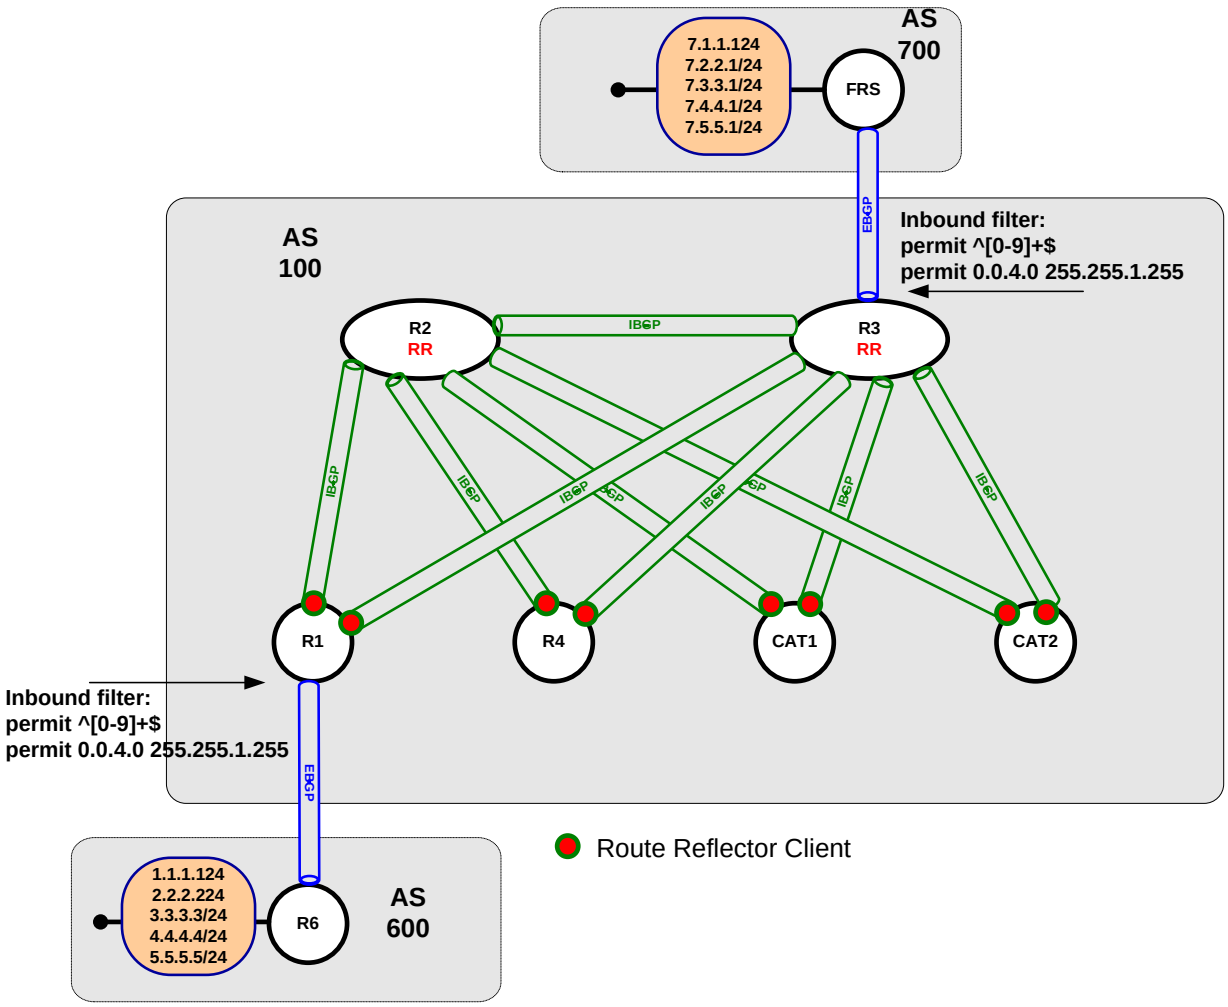

Issue: On each EBGP peer, accept only prefixes with a third octet of 4 or 5 and originated from a locally connected AS into AS 100.

Solution:

This filtering requirement possesses two components: one based upon an IP prefix and one based upon AS-path characteristics. However, in order for a BGP update to be accepted both criteria must match, resulting in a logical AND match requirement. A logical AND match requirement is best fulfilled by configuring a route-map. Configure a route-map with the two match statements under a single route-map stanza: 1 match criterion for a prefix match and a second match criterion based upon an AS-path match. The route-map is applied to inbound EBGP neighbor statement on R1 and R3. The supporting access-lists for the route-map could be written as shown below.
For the prefix match: access-list 1 permit 0.0.4.0 255.255.1.255

This combination of base and mask says that the first, second and fourth octets can be anything, but the third octet must be 0000 0100 or 0000 0101 (4 or 5). The regular expression below says that the path must contain only one or more (+) numerals. This expression does not permit a blank path or a path with any spaces. Since there must be at least one numeral and there can be no spaces, the path must consist of a single AS number, which would be true only of prefixes originated from neighboring autonomous systems. If these filters are working as designed, then we should see only the prefixes 4.4.4.0, 5.5.5.0, 7.4.4.0 and 7.5.5.0 in AS 100.

For the as-path match: ip as-path access-list 1 permit ^[0-9]+$.

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".

15.8 Address Administration

HIDDEN ISSUES TO SPOT WITH THE ADDRESS ADMINISTRATION CONFIGURATION

Issue: Using a source IP address from the 11.1.1.0/24 subnet, ping R1 from both R3 and FRS. When performing these pings, do not let the packets originating from R3 and R6 reach R1 with the IP address 11.1.1.0/24. Ping the FRS Ethernet interface belonging to 11.1.1.0/24 subnet using a pre-selected destination address extracted from the 22-bit subnet assigned to VLAN 36.

Solution:

This is a Network Address Translation (NAT) problem that does not explicitly mention NAT. The router performing NAT – in this scenario it is router R3 – will re-write the source IP address of packets that originate from the 11.1.1.0/24 subnet to a source IP address that is known to router R1. Router R1 will never know that the packets originated from the 11.1.1.0/24 subnet.

The NAT inside interface is the FastEthernet interface attached to VLAN 36 on R3, and the NAT outside interface is the R3 Frame-Relay interface. Line 1 below defines the inside source addresses that can be translated. Line 2 creates a pool of global addresses called **mypool**. This pool consists of the upper half of the connected 160.20.36.0/24 network, which is advertised via OSPF to the rest of the pod. Line 3 creates the translation from permitted inside addresses to the global pool. Though not strictly necessary given this topology, we use the overload keyword because we are permitting up to 254 inside addresses to be translated to just half that number of global addresses. Finally, line 4 makes our FRS E0 secondary address reachable through the global address 160.20.36.130.

```
1.     access-list 11 permit 11.1.1.0 0.0.0.255
2.     ip nat pool mypool 160.20.36.129 160.20.36.254 netmask 255.255.255.128
3.     ip nat inside source list 11 pool mypool overload
4.     ip nat inside source static 11.1.1.7 160.20.36.130.
```

15.9 Network Monitoring

HIDDEN ISSUES TO SPOT WITH THE NETWORK MONITORING CONFIGURATION

On R6, create loopback166 and leave it shutdown. If Loopback 106 goes down, the router should bring up Loopback 166 and attempt to email bob@netmasterclass.net. Use the following environmental variables. Set the domain name on R6 to netmasterclass.net

Solution:

The Embedded Event Manager 2.1 feature, available in IOS 12.3T, enables the router to monitor itself, take corrective action, and send informational messages. Included with the IOS on R6 are two sample scripts. The EEM allows you to build your own applications and scripts, but we can use a built-in sample script for our needs. The one we need is sl_intf_down.tcl.

Step one: make sure the sample script is available:

```
R6#sh event manager policy available
No.  Type    Time Created                 Name
1    system  Thu Feb 7  06:28:15 2036     sl_intf_down.tcl
2    system  Thu Feb 7  06:28:15 2036     tm_cli_cmd.tcl
```

Step two, set the environmental variables:

```
R6(config)#event manager environment _email_server 160.20.101.2
R6(config)#event manager environment _email_to bob@netmasterclass.net
R6(config)#event manager environment _email_from candidate@netmasterclass.net
R6(config)#event manager environment _email_cc instructor@netmasterclass.net
R6(config)#event manager environment _cron_entry 0-59/2 0-23/1 * * 0-7
R6(config)#$er environment _show_cmd show event manager policy registered
R6(config)#event manager environment _syslog_pattern .*UPDOWN.*Loopback106
R6(config)#event manager environment _config_cmd1 interface Loopback166
R6(config)#event manager environment _config_cmd2 no shut

R6#sh event manager environment
No.  Name                        Value
1    _email_server               160.20.101.2
2    _email_to                   bob@netmasterclass.net
3    _email_from                 candidate@netmasterclass.net
4    _email_cc                   instructor@netmasterclass.net
5    _cron_entry                 0-59/2 0-23/1 * * 0-7
6    _show_cmd                   show event manager policy regi
7    _syslog_pattern             .*UPDOWN.*Loopback106
8    _config_cmd1                interface Loopback166
9    _config_cmd2                no shut
```

Step three: Register the sample script:

```
R6(config)# event manager policy sl_intf_down.tcl

R6#sh event manager policy registered
No.  Class   Type    Event Type          Trap  Time Registered           Name
1    script  system  syslog              Off   Thu Mar 7 04:34:52 2002   sl_intf_down.tcl
```

```
        occurs 1 pattern {.*UPDOWN.*Loopback106}
        nice 0 queue-priority normal maxrun 90.000
```

Verification:

To test the configuration, simply shut Loopback 106. Here is what we got:

```
R6(config)#int loop106
R6(config-if)#shut
R6(config-if)#
05:34:22: %LINK-5-CHANGED: Interface Loopback106, changed state to administratively down
05:34:23: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback106, changed state to down
R6(config-if)#
05:34:25: %SYS-5-CONFIG_I: Configured from console by  on vty0
R6(config-if)#s
05:34:26: %LINK-3-UPDOWN: Interface Loopback166, changed state to up
05:34:27: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback166, changed state to up
R6(config-if)#
05:40:59: IP: s=160.20.16.6 (local), d=160.20.101.2 (FastEthernet0/0), len 44, sending
05:40:59:     TCP src=19121, dst=25, seq=1703828083, ack=0, win=4128 SYN
```

The output above is the result of a debug ip packet detail to the address of the imaginary email server. Note the packet to destination TCP port 25. The event manager sent syslog messages complaining about the failed email, but at least it did try!

Our use of the EEM feature is quite elementary. Check out this link for more information:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_14/gteem21.htm

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".

15.10 IPv6

HIDDEN ISSUES TO SPOT WITH THE IPv6 CONFIGURATION

Issue: Configure Link-local and Site-local IPv6 addresses in accordance with the IPv6 Diagram and the table.

Solution:

Remember to map both the link-local and site-local addresses on the Frame-Relay interfaces. Verify IPv6 connectivity between R2 and R5 with a ping before moving forward.

Issue: Configure IPv6 OSPF areas 6 and 25 as shown in the IPv6 diagram. Connect them via IPv6 area 0. You may add a single IPv6 network, as required. Do not change the default OSPF network

types. Do not configure IPv6 addresses on the Frame Relay hub-and-spoke or on the VLAN16 interfaces.

Solution:

Enable IPv6 OSPF on each specified interface with the command **ipv6 ospf 1 area X**. The default OSPF network type on the Frame-Relay interfaces connecting routers R2 and R5 is NBMA, so you will need to configure neighbor statements on these interfaces. A manual ipv6v4 tunnel between R2 and R6 is one way to meet the Area 0 requirement. Here is the tunnel configuration for R2:

```
interface Tunnel0
 no ip address
 ipv6 address FEC0:26::2/64
 ipv6 address FE80::262 link-local
 ipv6 traffic-filter NOEF in
 ipv6 ospf 1 area 0
 tunnel source Serial0/0.123
 tunnel destination 160.20.16.6
 tunnel mode ipv6ip
```

And here is the configuration on R6:

```
interface Tunnel0
 no ip address
 ipv6 address FEC0:26::6/64
 ipv6 address FE80::266 link-local
 ipv6 ospf 1 area 0
 tunnel source F0/0
 tunnel destination 160.20.123.2
 tunnel mode ipv6ip
```

Issue: Make sure that routers R2 and R5 see only a summary route representing the IPv6 loopback interfaces on R6

Solution:

This task can be accomplished by configuring an OSPF inter-area summary on R6. The challenge lies in recognizing that the addresses are given in hexadecimal and finding the appropriate mask length. The original addresses were configured with /124 masks. The last 32 bits of each address can be analyzed as follows:

```
::250:1       ::0000 0010 0101 0000:0000:0000:0000:0001

::251:1       ::0000 0010 0101 0001:0000:0000:0000:0001

::252:1       ::0000 0010 0110 0010:0000:0000:0000:0001
```

With our /124 mask, the host bits are those marked in green above. The most efficient summary would be a /110, which is the number of bits in common, to the left of those marked in yellow above. We issued the command **area 6 range FEC0::250:0/110** under the ipv6 router ospf process on R6. Here is the resulting routing table on R2:

```
R2#sh ipv6 route
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
L   FE80::/10 [0/0]
     via ::, Null0
OI  FEC0::250:0/110 [110/11111]
     via FE80::266, Tunnel0
C   FEC0:25::/64 [0/0]
     via ::, Serial1/0
L   FEC0:25::2/128 [0/0]
     via ::, Serial1/0
C   FEC0:26::/64 [0/0]
     via ::, Tunnel0
L   FEC0:26::2/128 [0/0]
     via ::, Tunnel0
O   FEC0:105::1/128 [110/781]
     via FE80::255, Serial1/0
L   FF00::/8 [0/0]
     via ::, Null0
```

Issue: All traffic that leaves R6 destined to IPv6 addresses should be tagged with DSCP value EF. Routers R2 and R5 should not see these DSCP values.

Solution:

The simplest solution is to take advantage of the 12.2(8)T Tunnel TOS feature. Entering the command **tunnel tos 184** on R6 sets the TOS byte to binary 10111000 in the outer (IPv4) header, regardless of the value of the TOS byte in the encapsulated IPv6 packet. Remember that DSCP is represented by the first six bits of the TOS byte, so 101110 is decimal 46 and PHB EF, but the command sets the full 8 bits, and must be padded with two zeros on the right, resulting in the decimal 184. This feature is documented at the link below.

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s17/12s_tos.htm

Verification:

To verify the operation we created access lists that matched on dscp EF and applied them inbound on the F0/0.16 subinterface of R1, the Tunnel interface of R2 and the Frame Relay interface on R5. When we generated pings from R6 to FEC0:105::1, we got hits on R1 (IPv4), but not on R2 or R5 (IPv6).

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".

15.11 QOS

HIDDEN ISSUES TO SPOT WITH THE QOS CONFIGURATION

Issue: Limit incoming UDP traffic destined to port 5111 to a rate of 8000 bit/sec on the interface Fa0/0 of router R1. Configure the minimal syntax values for burst size and extended burst size. Drop excessive traffic.

Solution:

Even though the term "traffic policing" is never explicitly stated, this task possesses all of the characteristics of a traffic-policing requirement. An inbound rate is specified for a defined set of traffic. If that inbound rate is exceeded, the traffic is to be dropped. Traffic policing will never buffer traffic. A traffic-policing configuration will drop traffic, pass it through, or pass it through and mark the packets using tools such as IP precedence bits or DSCP bits. This configuration requirement can be fulfilled using the rate-limit command (CAR) or the Modular QOS CLI (MQC). You can verify your configuration with the commands **show interface rate-limit** or **show policy-map interface**.

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".

15.12 Catalyst Specialties

HIDDEN ISSUES TO SPOT WITH THE CATALYST SPECIALTIES CONFIGURATION

Issue: Set port Fa0/18 of CAT2 to bypass the learning and listening states of Spanning Tree.

Solution:

In order to bypass the learning and listening states of Spanning Tree, configure the following port configuration command on the Catalyst 3550: spanning-tree portfast.

Issue: On CAT2, monitor both transmitted and received traffic on port fa0/19 from a SNIFFER attached on port Fa0/20.

Solution:
To fulfill this configuration, enter the following two commands in global configuration mode on CAT2. Note that monitoring both transmitted and received traffic is the default.

```
monitor session 1 source interface Fa0/19
monitor session 1 destination interface Fa0/20.
```

Issue: Configure the Catalyst so that learned MAC addresses in VLAN 40 are retained by the Catalyst for a period that is one and one half times as long as the default.

Solution:

The default aging time for a MAC address residing in a mac-address table on a Catalyst 3550 is 300 seconds. To fulfill this requirement, enter the following command in global configuration mode on a Catalyst 3550: **mac-address-table aging-time 450 vlan 40**.

Issue: Send all Catalyst error messages to the syslog server at 160.20.1.150.

Solution:

To configure a Catalyst 3550 to forward messages to a syslog server, you perform the same two steps that you would perform on a Cisco router.

Step 1: Enter the following global configuration command: **logging 160.20.1.150**.
Step 2: Enter the following global configuration command to send all syslog ERROR messages as well as lower level severity level messages to the syslog server: **logging trap errors**.

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".

15.13 Multicast

HIDDEN ISSUES TO SPOT WITH THE MULTICAST CONFIGURATION

Issue: Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree. Do not use any dynamic methods to discover or advertise the root of the shared tree.

Solution:

The multicast routing protocol that meets the requirements stated above is PIM Sparse Mode. If you cannot use any dynamic methods to discover or advertise the root of the shared tree, you must configure each and every multicast router with a manual configuration that statically identifies the Rendezvous Point. This can be accomplished with the following global configuration command: **ip pim rp-address X.X.X.X** where X.X.X.X is the IP address of the Rendezvous Point.

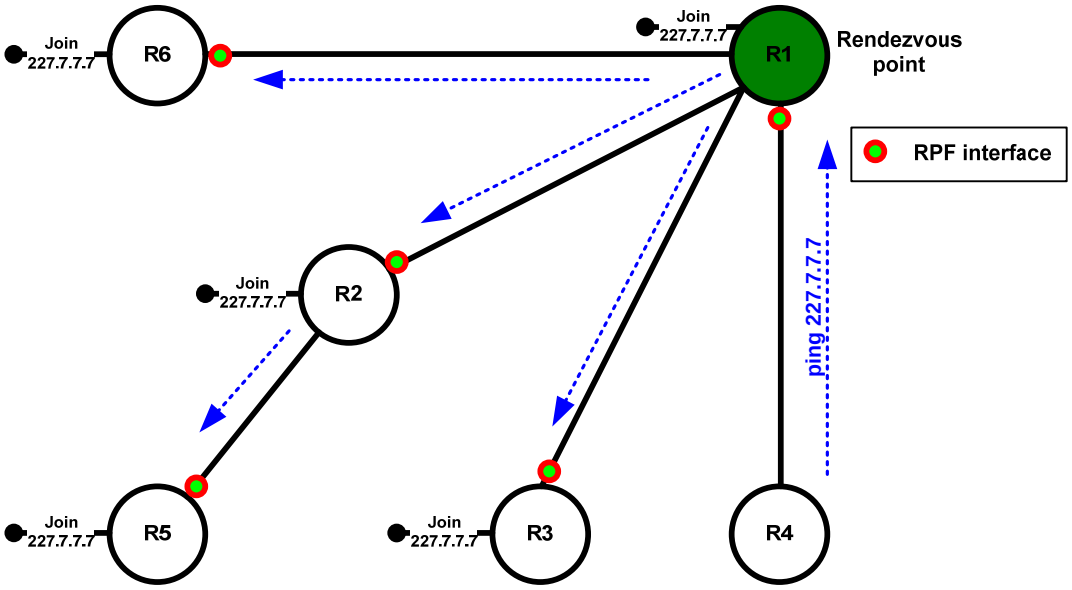

Issue: Configure R1, R2, R3 R5 and R6 to join the multicast group 227.7.7.7 only. Associate this group with a loopback interface on each router.

Solution:

In this lab we are using the loopback interfaces to simulate subnets with multicast clients attached. We can simulate a multicast client by configuring a multicast routing protocol and an **ip igmp join-group** command on the interface. To restrict the groups that clients attached to a router interface may join, we can define permitted groups in an access list and apply it to the interface with the **ip igmp access-group** command.

Verification:

You should get replies from your pings similar to the following:

```
R4#ping 227.7.7.7 rep 100

Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 227.7.7.7, timeout is 2 seconds:

Reply to request 0 from 160.20.14.1, 44 ms
Reply to request 0 from 160.20.25.5, 120 ms
Reply to request 0 from 160.20.123.2, 92 ms
Reply to request 0 from 160.20.123.3, 84 ms
Reply to request 0 from 160.20.16.6, 68 ms
```

For a more detailed look at the multicast performance, examine the output of **show ip mroute**, as shown below for R2. Make sure you have an active ping going from R4, because all of these states will timeout in a few minutes.

```
(*, 227.7.7.7), 05:56:31/00:02:34, RP 160.20.101.1, flags: SJCL
  Incoming interface: Serial0/0.123, RPF nbr 160.20.123.1
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 05:54:48/00:02:34
    Loopback102, Forward/Sparse, 05:56:31/00:02:16
```

The (*,G) entry above has an “S” flag, indicating that this group is being distributed in sparse mode. This entry describes the RPF interface and Outgoing Interface List for the SHARED tree, which is rooted at the Rendezvous Point. The (S,G) entry below shows the RPF interface and Outgoing Interface List for the SOURCE tree, rooted at the source of our ping.

```
(160.20.14.4, 227.7.7.7), 00:04:29/00:03:17, flags: LT
  Incoming interface: Serial0/0.123, RPF nbr 160.20.123.1
  Outgoing interface list:
    Loopback102, Forward/Sparse, 00:04:30/00:02:14
    Serial1/0, Forward/Sparse, 00:04:30/00:02:33
```

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”.