NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II

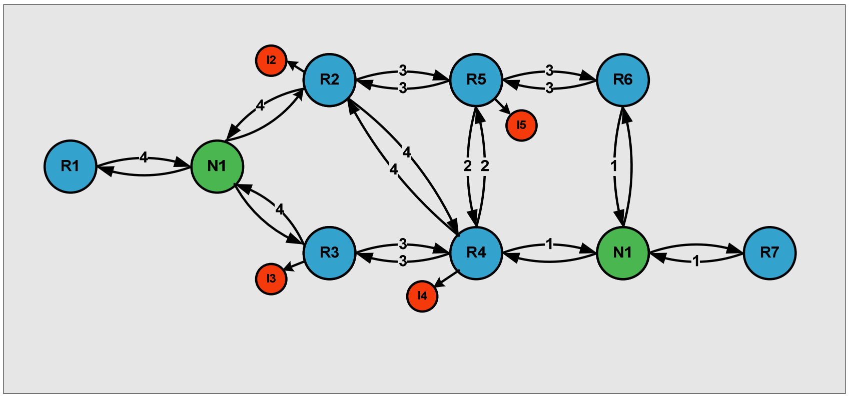

Scenario 16
ANSWER KEY

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-200V6 Scenario 16: 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 160.16.0.0/16.
- Do not use any static routes.
- Advertise Loopback interfaces with their original masks.
- Do not use the default route 0.0.0.0/0 or default-information originate, unless specified otherwise.
- 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.

Explanation of Each of the Goals and Restrictions:

IP subnets in the Scenario diagram belong to network 160.16.0.0/16

The third and forth octets of the IP addresses displayed on the diagram belong to 160.16.0.0/16.

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 mask.

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.

Do not use the default route 0.0.0.0/0 or default-information originate

Unless specified otherwise (see the ODR task).

Make sure all IP interfaces in the diagram are *reachable* within this internetwork.

This is a key goal to observe. This requires that all of your IGPs 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. You must perform redistribution in order to assure that all IP addresses are reachable without the use of static routes.

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

16.1 Frame Relay

HIDDEN ISSUES TO SPOT WITH THE FRAME RELAY CONFIGURATION

Issue: Configure a physical interface on router R3 and logical interfaces on all other Frame-Relay interfaces. Use point-to-point logical interfaces wherever possible.

Solution:

The Scenario diagram shows a hub and spoke Frame-Relay topology, with R1 as the hub. R1 is using three DLCI’s while all other routers are only using one DLCI. Of the three DLCI’s R1 is using, two are associated with one subnet and the third one is associated with a second subnet. If R1 is to be configured with only logical Frame-Relay interfaces, at least one of the interfaces must be multipoint, since one of the Frame-Relay subnets that R1 is attached to requires the configuration of two DLCI’s on R1. Since you are required to use point-to-point logical interfaces wherever possible, the second logical interface on R1 must be point-to-point. Spokes R2 and R4 will be configured with point-to-point subinterfaces, and R3 will be configured on the physical interface, as instructed.

As result of this configuration requirement, the 160.16.123.0/24 subnet will have three routers attached to it with three different Frame-Relay interface types. R1 will be configured with a multipoint Frame-Relay subinterface, R2 will be configured with a point-to-point Frame-Relay subinterface, and R3 will be configured with a physical Frame-Relay interface. This diversity of interface types will create configuration challenges for Link-State routing protocols such as OSPF. Please read ahead in this Scenario to determine which IGP is running on the Frame-Relay cloud.

Issue: R1, R2 and R3 should be in the same subnet. R1 and R4 should be in the same subnet. Use only the minimum number of DLCI’s to fulfill this configuration.

Solution:

While the Scenario specifies a hub and spoke topology, the underlying Frame-Relay PVC configuration is a full-mesh. In order to fulfill the configuration requirement of “use only the minimum number of DLCIs” in this Scenario, it is recommended that you disable frame-relay inverse arp on all routers and configure frame-relay map statements for only the DLCIs referenced in the Scenario diagram. This will assure that only the “minimum” number of DLCIs are being used for this configuration.

Implementation:

- Assign IP addresses to Frame-Relay interfaces according to the IP connectivity diagram.
- Map DLCIs 102 and 103 to point-to-multipoint Frame-Relay subinterface on R1 by issuing the frame-relay map ip 160.16.123.2 102 broadcast and frame-relay map ip 160.16.123.3 103 broadcast commands.
- Map DLCI 301 to physical Frame-Relay interface on R3 by issuing the frame-relay map ip 160.16.123.1 301 broadcast and frame-relay map ip 160.16.123.2 301 commands.
- Assign DLCI 201 to point-to-point Frame-Relay subinterface of R2 by issuing the frame-relay interface-dlci 201 command.

- Assign DLCI 401 to point-to-point Frame-Relay subinterface of R4 by issuing the frame-relay interface-dlci 401 command.

Verification:

Issue the **show frame-relay pvc** command on each router and verify the number of DLCIs marked as Local, Switched and Unused. Router R1 should indicate 3 Local Active PVCs. Routers R2, R3 and R4 should indicate 1 Local Active PVC and 2 unused PVCs.

In order to fulfill the configuration requirement of “use only the minimum number of DLCI’s” in this Scenario, it is recommended that you disable frame-relay inverse arp on all routers and configure frame-relay map statements for only the DLCI’s referenced in the Scenario diagram

Verify number of DLCI’s in use by issuing the **show frame-relay map** command. There should be no maps to “0.0.0.0” and no entries marked “dynamic.”

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”.

16.2 Catalyst Configuration

HIDDEN ISSUES TO SPOT WITH THE CATALYST CONFIGURATION

Issue: Do not use a CISCO proprietary trunking protocol on router R1.

Solution:

Two trunking protocols are available on a Cisco 3550: ISL and 802.1Q. ISL is a Cisco proprietary trunking protocol. If you are explicitly instructed not to use a Cisco proprietary trunk protocol, interpret that to mean that you cannot configure ISL trunking. As an alternative to using ISL, configure 802.1Q trunking.

Verification:

Issue the command **show interface trunk** on CAT2.

Issue: Create the VLAN’s referenced in the diagram and in the VLAN table. When creating these VLAN’s, allow the VLAN’s to be advertised from CAT2 to CAT1.

Solution:

References to advertising VLAN’s between Catalyst switches are indirectly guiding you on how to configure VTP. The primary purpose of VTP is to advertise VLAN’s from one switch to another. There are three VTP modes: (1) VTP server, (2) VTP client and (3) VTP transparent mode. A Catalyst switch configured in VTP server mode and assigned to a VTP domain, will advertise all VLAN's that are created on it to other Catalyst switches in the same VTP domain. In order to fulfill the stated configuration requirement, configure CAT2 as a VTP server. CAT1 can be configured as either a VTP server or client. At the very least, CAT1 must be in the same VTP domain as CAT2.

Verification:

To verify the VTP information, issue the **show vtp status** command. To verify the VLAN information issue the **show vlan brief** command.

Issue: Enable trunking only on the necessary ports. Permit only the required VLANs across configured trunks. Create trunk links of Catalyst ports F0/13 and F0/14. Do not permit VLAN 100 across any trunk link between the Catalysts

Solution:

A diagram like the following can be helpful:

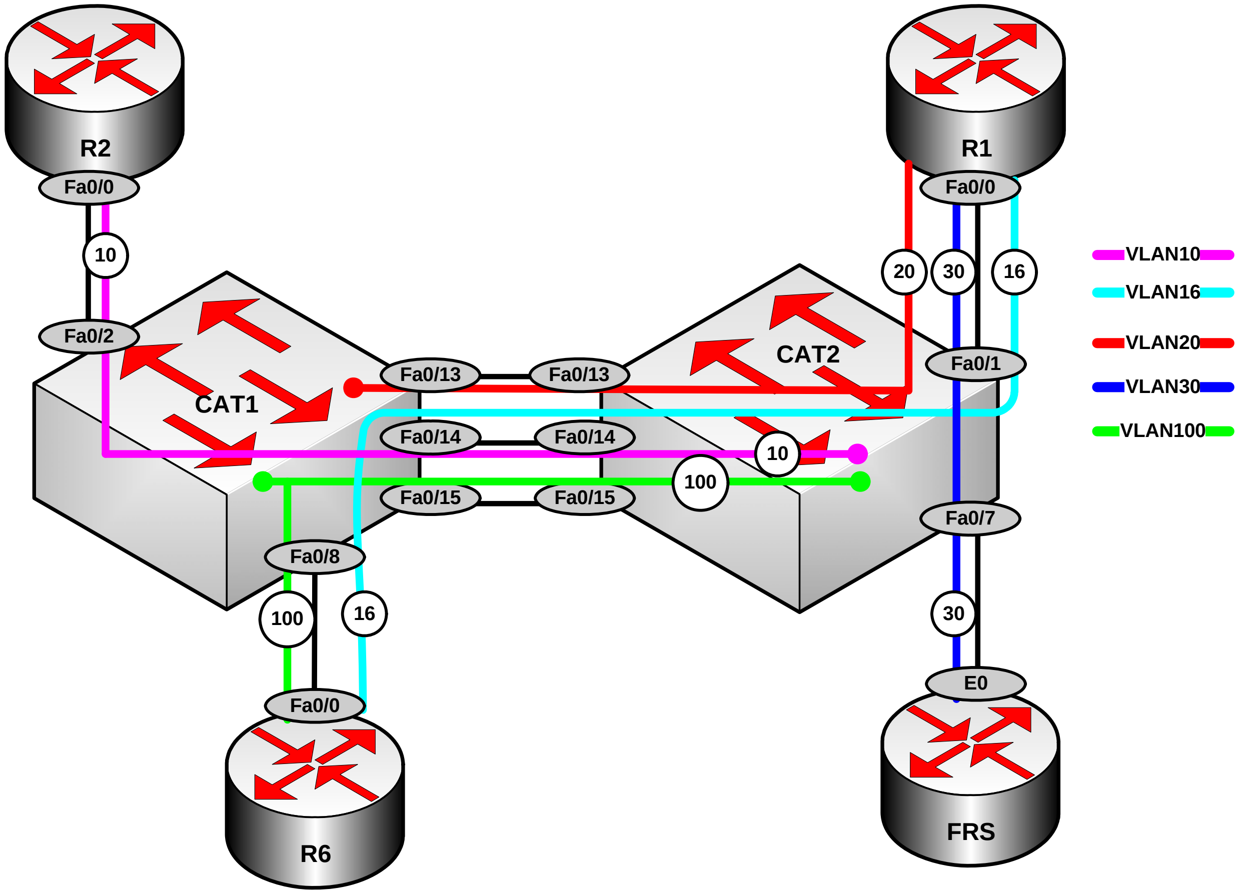

Configure the following Catalyst 3550 interface configuration command to restrict the number of VLAN's that pass over a trunk port: **switchport trunk allowed vlans XX** where XX is the range of VLAN's to allow over the trunk port. VLAN 100 must cross between the Catalysts. Since it cannot go on a trunk link, make the F0/15 ports access links in VLAN 100.

Verification:

To verify trunk information issue the command **show interfaces trunk**.

Issue: Configure VLAN 20 between CAT1 and the fa0/0 interface on R1 as well as VLAN 30 between FRS and the fa0/0 interface on R1.

Solution:

Since there are two VLANs on the same R1 interface, router R1 will need to be configured with FastEthernet subinterfaces, with each subinterface referencing a unique VLAN. Use the command **encapsulation dot1q XX** where XX is the VLAN number on each subinterface.

In order to determine IP addressing on these FastEthernet interfaces you must read ahead to the IGP configuration section. The OSPF section requires us to "Advertise FRS and the Catalyst VLAN 20 interface as one 160.16.11.0/24 OSPF network." A solution to the seemingly irreconcilable configuration requirements is to configure IRB on router R1. Place the VLAN 20 and VLAN 30 FastEthernet subinterfaces in the same bridge-group, enable IRB, create a BVI interface on router R1 and place it in the 160.16.11.0/24 subnet. Enabling IRB on router R1 will require the following commands:

- bridge irb (to activate Integrated Routing and Bridging)
- bridge 1 route ip (to send IP packets to the BVI)
- bridge 1 protocol ieee (to create the bridge and assign a spanning-tree protocol)
- bridge-group 1 (to associate the subinterfaces with the bridge)

Verification:

The command **show interface IRB** will show which interfaces are bridging and routing for which protocols.

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".

16.3 OSPF

HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

Issue: Configure the Frame-Relay network 160.16.123.0/24 as the OSPF backbone area. Do not elect a DR/BDR in the OSPF area.

Solution:

If you cannot elect a DR/BDR on the 160.16.123.0/24 subnet, you cannot use the OSPF Broadcast or Non-Broadcast network types. This leaves you to choose between the following three OSPF network types: (1) point-to-point, (2) point-to-multipoint and (3) point-to-multipoint non-broadcast. You cannot use the OSPF point-to-point network type because you can only discover one neighbor with this OSPF network type. R1 is going to discover 2 neighbors. Therefore, you need to configure either the OSPF network type point-to-multipoint or point-to-multipoint non-broadcast. The difference between these two network types is: the point-to-multipoint OSPF network type uses the 244.0.0.0.5 multicast for HELLO advertisement, while the point-to-multipoint non-broadcast OSPF network type uses unicast addresses for HELLO advertisements. Since the point-to-multipoint non-broadcast OSPF network type unicasts its HELLO packets, it requires the configuration of neighbor statements to specify the unicast addresses to be used by the HELLO packets. Again, it must be emphasized that neither point-to-multipoint non-broadcast nor point-to-multipoint OSPF network types employ a DR/BDR election. Issue the command **ip ospf network point-to-multipoint** on R1 under multipoint Frame-Relay subinterface Serial0/0.123. On R2 and R3 enter this command under the point-to-point Frame-Relay subinterfaces Serial0/0.123.

Verification:

Issue the command **show ip ospf interface** on each router to verify the OSPF network types. Verify that each router has a full adjacency with the command **show ip ospf neighbor**.

Issue: Create the necessary OSPF topology on R1 using a single OSPF process. You are permitted only the following network statement under this process:

```
network 160.60.123.0 0.0.0.255 area 0
```

Solution:

IOS 12.3T permits specification of OSPFv2 areas on interfaces. Here is an example for Loopback 101:

```
interface Loopback101
 ip address 160.60.101.1 255.255.255.0
 ip ospf network point-to-point
 ip ospf 1 area 1
```

This feature could be helpful for certain scenarios involving unnumbered links.

Issue: Make R4 the designated router for the R1/R4 link.

Solution:

If R4 is going to be the designated router for the R1/R4 link, you must make sure that both ends of the link are configured as either the OSPF network type broadcast or non-broadcast. By default, both ends of the link will be set to the OSPF network type point-to-point since both ends are configured with Frame-Relay point-to-point subinterfaces. Whether you use the broadcast or non-broadcast OSPF network type to fulfill the DR/BDR configuration requirement set the OSPF priority to 0 on the R1 end of the link. This will assure that R4 becomes the DR for this link.

Verification:

Issue the **show ip ospf neighbor Serial0/0.14** command on router R1. The output below verifies that the neighbor with OSPF router ID 160.16.40.228 is the DR for the link.

```
R1#sh ip ospf nei s0/0.14

Neighbor ID     Pri   State           Dead Time   Address         Interface
160.60.104.1      1   FULL/DR         00:01:47    160.60.14.4     Serial0/0.14
R1#
```

Issue: Configure OSPF MD5 authentication for area 0.

Solution:

When you configure MD5 authentication for OSPF Area 0, you must remember that a virtual-link extends OSPF Area 0. To fulfill the configuration requirements of this task, you must configure MD5 authentication on routers R1, R2 and R. Under the OSPF routing process of each of these routers, enter the following command: **area 0 authentication message-digest**. Apply the key to each interface in area 0.

Verification:

To verify that authentication is enabled on the link, issue the **show ip ospf interface** command, as show here for R1:

```
R1#sh ip ospf interface s0/0.123
Serial0/0.123 is up, line protocol is up
  Internet Address 160.60.123.1/24, Area 0
  Process ID 1, Router ID 160.60.101.1, Network Type POINT_TO_MULTIPOINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT,
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:17
  Supports Link-local Signaling (LLS)
  Index 1/4, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, 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 160.60.103.1
    Adjacent with neighbor 160.60.102.1
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
R1#
```

Issue: Advertise FRS and the Catalyst VLAN 20 interface as one 160.16.10.0/24 OSPF network. Assign this network to OSPF area 1.

Solution:

This task provided with some configuration clues that lead you to configuring IRB on router R1. This task must be read in conjunction with the tasks under the Catalyst Configuration Section.

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".

16.4 RIP

HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

Issue: Configure RIP over the connection between R1 and R6. Do not broadcast or multicast the RIP updates.

Solution:

In order to fulfill this requirement, you must unicast RIP updates over the F0/0.16 connection. This is accomplished by configuring the VLAN 16 interfaces as RIP passive-interfaces and configuring neighbor statements referencing the unicast IP addresses of the remote end.

Verification:

Issue the command **debug ip rip**, watch the updates, and make sure that you only see unicast RIP packets.

Issue: R6 should receive the minimal number of necessary prefixes from R1 via RIP. It may not be a 0.0.0.0/0 route, and you may not create summaries on any interface.

Solution:

Notice that the network covered by RIP is Class A, 16.0.0.0/8, and all the routes it learns via OSPF are subnets from the Class B 160.16.0.0/16 network. Enabling auto-summary on RIP will make it summarize the learned prefix to their classful boundary. R6 will receive just the 160.16.0.0/16 subnet from R1, and it comprises all the subnets that must be reachable.

Issue: Make sure R1 has the RIP prefix 160.60.106.0/24 in its routing table.

Solution:

R6 has only its loopback prefix, 160.16.106.0/24 to send to R1. Disable auto-summary on R6 to make sure it sends the /24 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".

16.5 EIGRP

HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

Issue: Configure EIGRP AS 1 between R2 and CAT2. Make sure that EIGRP advertises only over VLAN10 interfaces.

Solution:

The recommended general practice to make sure that EIGRP advertises only over the specified set of interfaces or VLANS is to use the mask option with the EIGRP network command. An alternative method is to use the passive-interface command. However, when you use the passive-interface command, you are only restricting that interface from advertising EIGRP control traffic out of the interface. The interface is still part of the EIGRP process and any IP address associated with the interface will be advertised out to other EIGRP speakers. If you want to restrict an interface from generating any EIGRP control traffic in addition to restricting the same interface from being advertised as an EIGRP prefix, use the EIGRP network command with the mask option to limit precisely which prefixes you want to participate in EIGRP.

Verification:

Use the **show ip eigrp interface** command to see which interfaces are involved in EIGRP process.

Issue: Send a summary for the entire Class B range only from R2 to CAT2.

Solution:

Configure the following command on R2's interface f0/0:

```
ip summary-address eigrp 1 128.0.0.0 192.0.0.0
```

To make sure that this summary only goes to CAT2, add a route-map to your redistribution that denies this summary into OSPF.

16.6 On-Demand Routing

HIDDEN ISSUES TO SPOT WITH THE ODR CONFIGURATION

Issue: Configure ODR between R3 and R5. R5 is permitted 0.0.0.0/0 route.

Solution:

On-demand routing uses CDP to report prefixes from stub routers to hub routers, and to send a default route to the stub router. The stub router must not be running any dynamic routing protocol. Simply issue the command **router odr** on the hub router.

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".

16.7 Redistribution

HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION

Since there are no physical loops that cross routing domains in this lab, redistribution can be done without fear of route-feedback, vastly simplifying the task. Our design sends all routes into the core protocol, OSPF, and sends the minimum amount of information allowed to the edge protocols.

Issue: On R3, redistribute 160.16.103.0 into OSPF as a Type 1 External route. This should be the only OSPF Type 1 prefix in the network

Solution:

The redistribution of the loopback address on R3 as a connected prefix will require a route-map. There are two connected interfaces that need to be redistributed into OSPF, Loopback 103 and S0/1, and both must be accounted for by the route-map.

Below is a TCL script you can use to test universal reachability. To use the script, enter the command **tclsh** in privileged mode, and paste in this script. To kill failing pings, hold down CTL and SHIFT and hit the 6 key twice. When you are done enter **tclq** to leave the tcl mode. This list excludes the 140.10.0.0/16 and 1.1.1.0/24 addresses that are part of the BGP task.

```
foreach addr {
16.16.16.1
160.60.14.1
160.60.10.1
160.60.123.1
160.60.101.1
160.60.26.2
160.60.123.2
160.60.102.1

160.60.35.3
160.60.123.3
160.60.103.1

160.60.14.4
160.60.104.1

160.60.35.5
160.60.105.1

16.16.16.6
160.60.106.1

160.60.10.7
160.60.107.1
160.60.10.11
160.60.110.1
160.60.26.6
160.60.120.1
} {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”.

16.8 BGP

HIDDEN ISSUES TO SPOT WITHT THE BGP CONFIGURATION

Issue: Configure AS 1581 on CAT1, AS 1771 on R6 and AS 1776 on CAT2. Configure VLAN100 on R6, CAT1 and CAT2, and assign an IP address from the 1.1.1.0/24 subnet to each of these routers that possess a connection to VLAN 100. See the Scenario diagram.

Solution:

VLAN 100 is used for BGP backend connectivity only. It is not part of any IGP and does not have to be reachable.

Verification:

Issue the ping command from one of the routers on subnet 1.1.1.0/24. Make sure all others respond. Issue the **show ip bgp summary** command on all routers; make sure the necessary adjacencies are formed over subnet 1.1.1.0/24

Scenario BGP Diagram

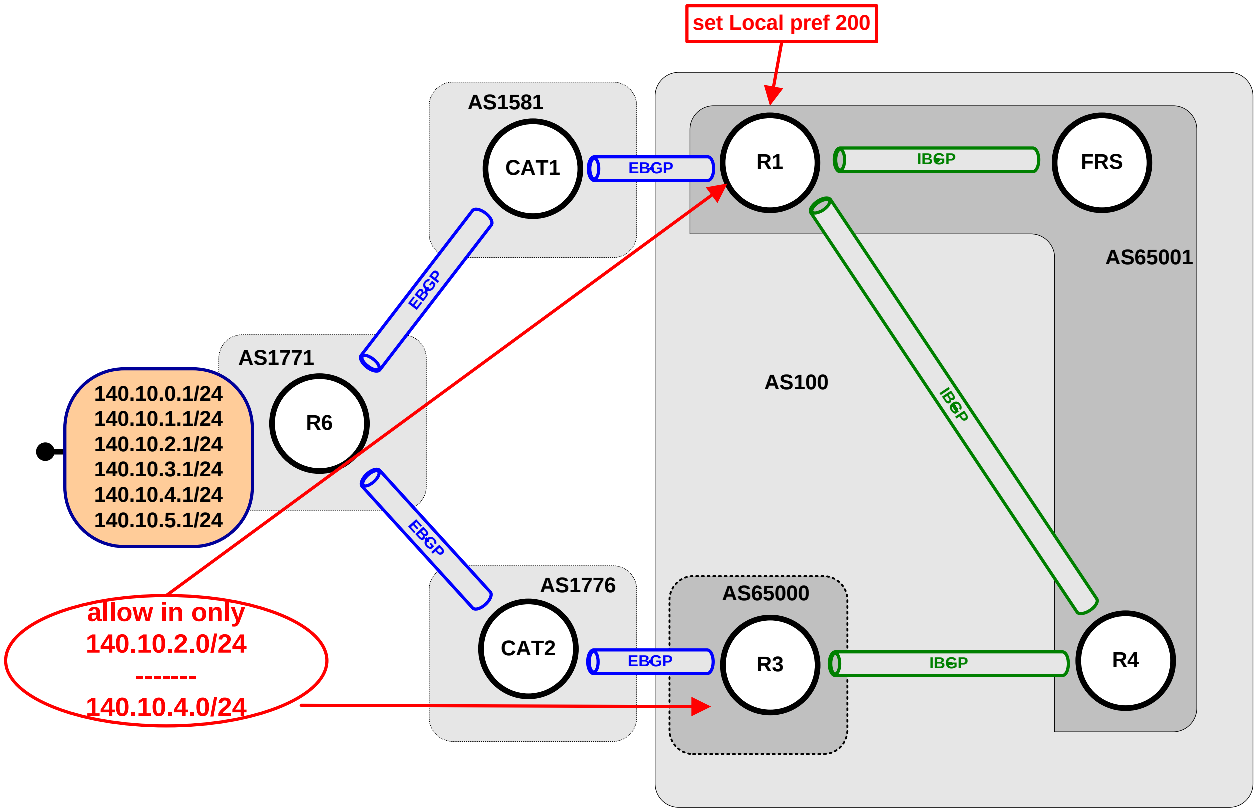

Issue: Configure the BGP router-id 160.16.110.1 and 160.16.120.1 on CAT1 and CAT2 respectively.

Solution:

When you can the router-id on a BGP speaker, beware of what it might do to your synchronization requirements. In this case CAT1 and CAT2 are in their own AS's with no other IBGP speakers, so they have no synchronization issues. Synchronization is an issue for only IBGP learned updates. A second issue associated with BGP router-ids, is that the router-id acts as a tiebreaker for the BGP path selection algorithm.

Verification:

Issue the **show run | begin router bgp** command on CAT1 and CAT2; make sure you see the router-id set up in the configuration.

Issue: Make R1, R3, R4 and FRS BGP speakers within AS 100. Do not peer R4 and FRS.

Solution:

The configuration requirements listed above are spread out over multiple configuration tasks in this Scenario's BGP Configuration Section. Taken together, these tasks suggest that you configure a confederation among the routers listed above using two private AS's, 65000 and 65001.

Implementation:

Configure R1, R4 and FRS with BGP AS number 65001. Set up a confederation id 100 and list other confederation sub-Autonomous Systems. Issue the **router bgp 65001** command on routers R1, R4 and FRS. Issue **bgp confederation identifier 100** command on routers R1, R4 and FRS. Issue the **bgp confederation peers 65000** command on routers R1, R4 and FRS.

Configure R3 with BGP AS number 65000. Set up a confederation id 100 and list other confederation sub-Autonomous Systems. Issue the **router bgp 65000** command on router R3. Issue **bgp confederation identifier 100** on router R3. Issue **bgp confederation peers 65001** on router R3. Issue the **show ip protocols** command on BGP speaking routers to verify confederation members.

Since FRS and R4 may not peer within AS 65001, R3 will be configured as a route-reflector. R3 and CAT2 need to configure **ebgp-multihop** in order to form their EBGP neighbor relationship.

Verification:

Issue the **show ip bgp summary** command on BGP speaking routers to verify local AS number and configured neighbors.

Issue: Configure R1, R4 and FRS in AS 65001. Configure R3 in AS 65000.

Solution:

Within each private AS, you will configure a route-reflector. R1 will be the route-reflector for AS 65001. R3 will be the route-reflector for AS 65000. R1 is able to form an EBGP neighbor relationship on a commonly shared subnet with Cat1; however R3 cannot do the same with CAT2. R3 and CAT2 need to configure **ebgp-multihop** in order to form their EBGP neighbor relationship.

Issue: Only one peering relationship can exist between private AS's 65000 and 65001. It must be between R3 and R4.

Solution:

R3 and R4 are not directly connected. Since R3 and R4 are not directly connected, they will need to be configured with EBGP multihop. Be on the lookout for any routing loops in this configuration since the path the R3-R4 neighbor relationship is taking is the same path used for forwarding IP end user packets.

Verification:

Issue the **show ip bgp summary** command on router R3 and make sure that neighbor relationship to R4 exists and is UP.

Issue: Allow only the prefix range of 140.10.2.0/24 – 140.10.5.0/24 into AS100.

Solution:

To fulfill this configuration requirement, you must configure a minimum of two lines in your access-list or prefix-list, since the range you are supplied with crosses the number “4” bit boundary. The access-list to configure is:

```
        access-list 1 permit  140.10.2.0 0.0.1.0
        access-list 1 permit  140.10.4.0 0.0.1.0
```

The prefix-list is:

```
        ip prefix-list test permit 140.10.2.0/23 ge 24 le 24
        ip prefix-list test permit 140.10.4.0/23 ge 24 le 24
```

Apply either one of these filters to an inbound neighbor statement on routers R1 and R3 for each EBGP neighbor relationships these routers maintain. That is, the R1-CAT1 EBGP neighbor relationship and the R3-CAT2 EBGP neighbor relationship.

Verification:

On routers R1 and R3 issue the **show ip bgp** command to verify received routes.

Issue: Forward traffic from AS 100 destined to the 140.10.2.0/24 – 140.10.5.0/24 subnets via CAT1.

Solution:

For the prefixes specified, configure a route-map to match on the specified prefixes and set the local-preference on R1 to 200, the only AS 100 EBGP peer with CAT1 (AS 1581). By setting the local-preference on R1 to 200 and leaving the default local-preference setting on all other IBGP speakers within AS 100 (the default local-preference setting is 100), R1 will be the preferred exit path and CAT1 will consequently be the preferred router to forward traffic for the specified prefixes. All IBGP speakers regardless of whether or not they are in different sub-Autonomous Systems share the local-preference attribute. Local-preference is the most commonly used BGP attribute for influencing how traffic exits an AS.

Verification:

On all bgp speakers issue the command **show ip bgp** and make sure that the path traversing AS 1581 is preferred.

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”.

16.9 Address Administration

HIDDEN ISSUES TO SPOT WITH THE ADDRESS ADMINISTRATION CONFIGURATION

Issue: Configure R2 to assign IP addresses from the VLAN 10 IP address range to clients. Make sure the clients use CAT2 as a gateway. All clients are a part of netmasterclass.com. The IP address is negotiated for 10 minutes. The DNS server is ns.siteprotect.com; the ip address is 1.1.1.1

Solution:

This configuration requirement is a DHCP configuration task without ever explicitly mentioning DHCP. The DHCP configuration will be performed on router R2. It is recommended to begin the DHCP configuration by entering the following global configuration command for each IP address you want to exclude from being used by the DHCP server process: **ip dhcp excluded-address X.X.X.X** where X.X.X.X is the unicast IP address you want to exclude.

Once you have entered the IP addresses to exclude from DHCP, begin configuring your DHCP pool with the following global configuration command: **ip dhcp pool <NAME>**. After you enter this command, you will be placed in the “dhcp-config” mode. In this mode, you can configure a wide range of DHCP options including: the address range to use for dhcp (example: network 160.16.26.0 255.255.255.0), default gateway to be advertised by DHCP (example: default-router 160.16.26.2) and the DHCP lease time in DAYS HOURS MINUTES (configuration for a 10 minute lease time: lease 0 0 10).

Verification:

Issue the **show ip dhcp pool** command:

```
R2#sh ip dhcp pool

Pool NMC :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 254
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 160.16.26.1          160.16.26.1      - 160.16.26.254     0
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”.

16.10 NTP

HIDDEN ISSUES TO SPOT WITH THE NTP CONFIGURATION

Issue: Make R1 the primary time source. Set R1 as a stratum 5.

Solution:

To be configured as the primary time source, R1 must be configured as an NTP master. It is the NTP master that is the primary time source for all NTP sessions. As the configuration requirements state, set the NTP master to a stratum 5. Configure R1 as NTP master by issuing the **ntp master 5** command.

Verification:

Issue the show ntp status command on R1:

```
R1#sh ntp status
Clock is synchronized, stratum 5, reference is 127.127.7.1
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**18
reference time is AF415FAA.ACCF937C (05:12:10.675 UTC Fri Mar 5 1993)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.02 msec, peer dispersion is 0.02 msec
R1#
```

Issue: Configure router R2 so that it obtains its time from R1.

Solution:

Configure R2 to be an NTP client of R1. NTP client obtains its time from an NTP server. An NTP client does not exchange synchronization messages with an NTP server. In order to configure R2 as an NTP client, configure the **ntp server** command on router R2 referencing a reachable IP address on R1.

Verification:

Issue the **show ntp status** command on R2:

```
R2#sh ntp status
Clock is synchronized, stratum 6, reference is 160.16.123.1
nominal freq is 250.0000 Hz, actual freq is 249.9997 Hz, precision is 2**18
reference time is AF415F33.470C44A0 (05:10:11.277 UTC Fri Mar 5 1993)
clock offset is -0.0531 msec, root delay is 46.66 msec
root dispersion is 0.23 msec, peer dispersion is 0.14 msec
R2#
```

Issue: Configure CAT2 to peer with R2. Make sure R2 and CAT2 are synchronized only if they pass authentication. Make sure that CAT2 is synchronized with R2 only.

Solution:

The word “synchronization” possesses special meaning with NTP. The two most basic types of NTP relationships are: (1) client/server mode and (2) peer mode. In client/server mode, the NTP client merely obtains time from the NTP server. In NTP peer mode, two NTP peers exchange NTP messages with each other to keep their clocks synchronized and to avoid clock “drift”. This configuration task requires that CAT2 and R2 become NTP peers. In addition, both R2 and CAT2 must pass authentication. NTP uses the MD5 hashing algorithm for authentication. Whenever you use MD5, you must make sure that the authentication strings are identical on all participating devices because MD5 performs a hash comparison.

Configure R2 to peer with CAT2, enable authentication and configure a key:

```
ntp authentication-key 1 md5 nmc
ntp authenticate
ntp trusted-key 1
ntp clock-period 17179887
ntp peer 160.16.26.6 key 1
```

Configure access list on CAT2 to allow it to synchronize to R2 only:

```
access-list 17 permit 160.16.26.2
```

Configure CAT2 to peer with R2, enable authentication and configure a key:

```
ntp authentication-key 1 md5 nmc
ntp authenticate
ntp trusted-key 1
ntp clock-period 17180503
ntp access-group peer 17
ntp peer 160.16.26.2 key 1
```

Verification: Issue the show NTP status command on CAT2:

```
CAT2#sh ntp status
Clock is synchronized, stratum 7, reference is 160.16.26.2
nominal freq is 250.0000 Hz, actual freq is 249.9907 Hz, precision is 2**18
reference time is AF4165F0.3C27516A (05:38:56.234 UTC Fri Mar 5 1993)
clock offset is -1.3079 msec, root delay is 48.81 msec
root dispersion is 7877.08 msec, peer dispersion is 7875.56 msec
```

Issue the show NTP associations detail command on CAT2:

```
CAT2#sh ntp associations det
160.16.26.2 configured, authenticated, our_master, sane, valid, stratum 6
ref ID 160.16.123.1, time AF416333.475828F5 (05:27:15.278 UTC Fri Mar 5 1993)
our mode active, peer mode active, our poll intvl 64, peer poll intvl 64
root delay 46.81 msec, root disp 0.20, reach 6, sync dist 7900.162
delay 2.00 msec, offset -1.3079 msec, dispersion 7875.56
precision 2**18, version 3
org time AF4165F0.3B8FB454 (05:38:56.232 UTC Fri Mar 5 1993)
rcv time AF4165F0.3C27516A (05:38:56.234 UTC Fri Mar 5 1993)
xmt time AF416618.43127756 (05:39:36.262 UTC Fri Mar 5 1993)
filtdelay =     2.00    2.20    0.00    0.00    0.00    0.00    0.00    0.00
filtoffset =   -1.31   -0.92    0.00    0.00    0.00    0.00    0.00    0.00
filterror =     0.38    1.36 16000.0 16000.0 16000.0 16000.0 16000.0 16000.0.
```

16.11 IPv6

HIDDEN ISSUES TO SPOT WITH THE IPv6 CONFIGURATION

Issue: Assign the following IPv6 addresses. Map as necessary, and verify same-subnet reachability

Solution:

The fact that the link-local addresses on R1's external interfaces are all the same emphasizes that they are, in fact, local to the link. We often hard-code these external addresses, rather than accepting the default EUI-64 address, so that the same configurations can be loaded on different pods.

The challenging part may be the required mapping of both the link-local and site-local addresses on the multipoint Frame-Relay interfaces. This is no different in concept than IPv4 mapping requirements. Unlike IPv4, it is not necessary to map local IPv6 addresses on these interfaces in order to ping them.

Issue: Place all of the above interfaces into IPv6 RIP with tag "RIP". Make sure all IPv6 routers can ping all IPv6 Site-Local addresses.

Solution:

The basic requirement can be met by issuing the command **ipv6 unicast-routing** in global configuration mode and the command **ipv6 rip RIP enable** on each of the specified interfaces. Thorough testing will reveal that the loopback network FEC3:3::/64 is missing on R2 and the loopback network FEC3:2:/64 is missing on R3. This is the result of split-horizon protection enabled by default on R1. In IPv4, RIP split-horizon is disabled on the interface, but in IPv6 RIP it is disabled under the process.

Should we disable split-horizon protection and just move on? Isn't split-horizon an important protection against route feedback and instability? If you disable split-horizon on R1 and then debug ipv6 rip on the spoke routers you will see that they are learning their own loopback networks from R1. This is unlikely to cause instability here, because these loopbacks are directly connected. As a general practice, however, we recommend inbound filters on each spoke to preserve split-horizon protection. An example is shown below for R2. We create an IPv6 access list that denies our routes and permits all others, and apply it in a distribute-list statement.

```
ipv6 router rip RIP
  distribute-list prefix-list NOTMINE in Serial0/0.1
!
ipv6 prefix-list NOTMINE seq 5 deny FEC3:2::/64
ipv6 prefix-list NOTMINE seq 10 permit ::/0 le 128
```

Verification

Just as with IPv4, a good way to thoroughly test universal reachability with IPv6 is to create and run a TCL script. An example suitable for this lab is shown below. On each router, enter **tclsh**, paste this script, and stand back! To leave TCL mode enter **tclq**.

```
foreach address {
FEC3:7B::1
FEC2:E::1
FEC1:64::1
FEC1:1::1
FEC3:7B::2
FEC3:2::1
FEC3:7B::3
FEC3:3::1
FEC2:E::4
FEC2:4::1
FEC1:64::6
FEC1:6::1
} {ping $address}
```

Issue: Configure R1 to send only a default route to R4.

Solution:

There are two interface configuration mode commands that can be used in IPv6 RIP to send a default route: **ipv6 rip RIP default-information originate** and **ipv6 rip RIP default-information only**. The former sends the network ::/0 in addition to the more-specific prefixes. The latter sends ::/0 and suppresses other routes, so we enter this command under interface S0/0.14 on R1.

Verification:

Here is the relevant part of the resulting table on R4:

```
R4#show ipv6 route rip
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
R   ::/0 [120/2]
     via FE80::1, Serial0/0.1
R4#
```

Issue: Configure R1 to send a single, longest summary to R6 representing the networks connected to R2, R3 and R4.

Solution:

The networks connected to R2, R3 and R4 begin with the first 16 bits of FEC2: and FEC3. As shown below, the first 15 bits in all these addresses are the same, and are different from the first 15 bits of FEC1:

FEC1:	1111	1110	1100	0001
FEC2:	1111	1110	1100	0010
FEC3	1111	1110	1100	0011

The longest summary would therefore be FEC2::/15. This can be configured on the F0/0.16 interface of R1 with the command **ipv6 rip RIP summary-address FEC2::/15**.

Verification:

Here is the relevant part of the IPv6 routing table on R6:

```
R6#sh ipv6 route rip
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, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   FEC1:1::/64 [120/2]
     via FE80::2B0:64FF:FEEF:A001, FastEthernet0/0.16
R   FEC2::/15 [120/2]
     via FE80::2B0:64FF:FEEF:A001, FastEthernet0/0.16
R6#
```

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".

16.12 QOS

HIDDEN ISSUES TO SPOT WITH THE QOS CONFIGURATION

Issue: Police this traffic to 8000 bit/sec to port Fa0/24 of CAT1. Configure the minimal syntax values for burst size and extended burst size. Drop excessive traffic.

Solution:

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 a mark the packets using tools such as IP precedence bits or DSCP bits. Since the configuration is on a Catalyst 3550, this configuration requirement must be fulfilled using the Modular QOS CLI (MQC). You will need to configure an access-list and associate it with a CLASS-MAP. Once the CLASS-MAP is configured, you will need to associate it with a policy-map. Once the class-map has been associated with a policy map, you will need to enter a **police** command under the policy map. Finally, you must apply the policy map to Fa0/24 of CAT1 with the **service-policy** interface configuration command.

Implementation:

Enable MLS QoS by issuing the **mls qos** command.
Configure class-map to match UDP traffic:

```
class-map match-all udp-5011
  match access-group 101
```

Configure policy-map to specify the rate limit parameters:

```
policy-map filter-udp-5011
  class udp-5011
    police 8000 8000 exceed-action drop
```

Attach policy-map to interface:

```
interface FastEthernet0/24
 service-policy input filter-udp-5011
```

Verification:

Issue the **show mls qos interface Fa0/24 policers** command:

```
CAT1#sh mls qos interface fa0/24 policers
FastEthernet0/24
policymap=filter-udp-5011
type=Single, id=0 rate=8000, qlimit=8000, drop=1
```

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".

16.13 Catalyst Specialties

HIDDEN ISSUES TO SPOT WITH THE CATALYST SPECIALTIES CONFIGURATION

Issue: Allow only a workstation with a Data-Link address of 0050.04fd.9f73 to use port fa0/20 of CAT2.

Solution:

To restrict access to only the data-link address listed, configure the following two interface configuration commands: **switchport port-security mac-address 0050.04fd.9f73** followed by **switchport port-security**.

Verification:

You can verify the status of switchport port-security with the command **show port-security Xy** where Xy is the interface configured with port security.

Issue: Allow TELNET access to CAT1 from the loopback interface of R3 only.

Solution:

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 command **access-class 1 in**.

Verification:

Issue the **show line vty 0 15** command and verify that access-list is applied to all vty lines:

```
CAT1#sh line vty 0 15
   Tty Typ     Tx/Rx    A Modem  Roty AccO AccI   Uses   Noise  Overruns   Int
     1 VTY              -    -      -    -    1      0       0     0/0       -
     2 VTY              -    -      -    -    1      0       0     0/0       -
     3 VTY              -    -      -    -    1      0       0     0/0       -
     4 VTY              -    -      -    -    1      0       0     0/0       -
     5 VTY              -    -      -    -    1      0       0     0/0       -
     6 VTY              -    -      -    -    1      0       0     0/0       -
     7 VTY              -    -      -    -    1      0       0     0/0       -
     8 VTY              -    -      -    -    1      0       0     0/0       -
     9 VTY              -    -      -    -    1      0       0     0/0       -
    10 VTY              -    -      -    -    1      0       0     0/0       -
    11 VTY              -    -      -    -    1      0       0     0/0       -
    12 VTY              -    -      -    -    1      0       0     0/0       -
    13 VTY              -    -      -    -    1      0       0     0/0       -
    14 VTY              -    -      -    -    1      0       0     0/0       -
    15 VTY              -    -      -    -    1      0       0     0/0       -
    16 VTY              -    -      -    -    1      0       0     0/0       -
```

Issue: “Welcome to RS-NMC-2!” on CAT1.

Solution:

Configure the Catalyst 3550 global configuration command **banner motd # Welcome to RS-NMC-2! #** where the symbol “#” represent a delimiter to the message “Welcome to RS-NMC-2!”. The “#” will be replaced with “^C” in the running configuration and will not appear in the message. Configuring this command on a Catalyst 3550 is the same as configuring the command on a router.

Verification:

Issue the telnet 160.16.10.11 command from router R3:

```
R3#telnet 160.16.10.11 /source-interface lo0
Trying 160.16.10.11 ... Open
CCC
Welcome to RS-NMC-2!

Password required, but none set
[Connection to 160.16.10.11 closed by foreign host]
```

Issue the same command from router R4:

```
R4#telnet 160.16.10.11 /source-interface lo0
Trying 160.16.10.11 ...
% 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".

16.14 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 flood and prune protocol.

Solution:

A protocol that uses any unicast routing protocol for source address determination is PIM. Now you must determine which version of PIM: Sparse Mode or Dense Mode. The answer lies in the last sentence displayed above: "based on a flood and prune protocol". It is PIM Dense Mode that is based upon a "flood and prune" protocol. Therefore, you need to configure PIM Dense Mode to fulfill the configuration requirement. Configure PIM dense mode by issuing the **ip pim dense-mode** command under interfaces that participate in the multicast task.

Verification:

Verify with the command **show ip pim neigbors**.

Multicast Topology

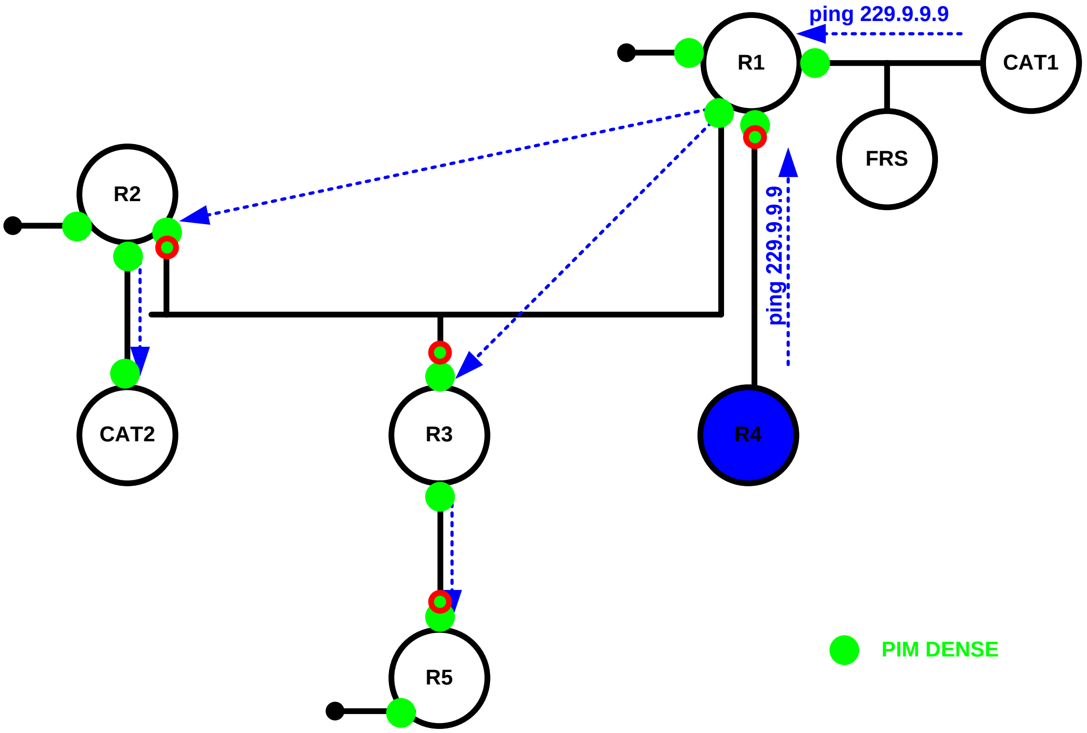

Issue: Configure all of the above listed routers to join the multicast group 229.9.9.9. Associate this multicast group with a loopback interface on each router.

Solution:

Configure the command **ip igmp join-group 229.9.9.9** under a loopback interface on each of the listed routers.

Verification:

Issue the command **show ip igmp interface** on routers R1, R2, R3 and R5:

```
R1#sh ip igmp int loop101
Loopback101 is up, line protocol is up
  Internet address is 160.60.101.1/24
  IGMP is enabled on interface
  Current IGMP host version is 2
  Current IGMP router version is 2
  IGMP query interval is 60 seconds
  IGMP querier timeout is 120 seconds
  IGMP max query response time is 10 seconds
  Last member query count is 2
  Last member query response interval is 1000 ms
  Inbound IGMP access group is not set
```

```
IGMP activity: 2 joins, 0 leaves
Multicast routing is enabled on interface
Multicast TTL threshold is 0
Multicast designated router (DR) is 160.60.101.1 (this system)
IGMP querying router is 160.60.101.1 (this system)
Multicast groups joined by this system (number of users):
    224.0.1.40(1)  229.9.9.9(1)
```

Issue: Statically configure a MAC table entry for ports in VLAN 20 and VLAN 30 for 229.9.9.9.

Solution:

Configure the following global configuration command on the Catalyst 3550: **mac-address-table static 0100.5e09.0909 interface Xy** where Xy is the interface name and number. 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 25th 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. By configuring this command, you will have fulfilled the static configuration requirement on a Catalyst 3550. To accomplish this task, IGMP snooping must be disabled on VLAN 20 and 30, and the static mac address 01.00.5e.09.09.09 must be assigned to CAT2 port FA0/1 for vlan 20 and to CAT2 port FA0/7 for vlan 30. This will fulfill the requirements of static configuration of vlans 20 and 30.

Verification:

Issue the **sh mac-address-table multicast vlan X** command, where X is a vlan number:

```
CAT2#sh mac-address-table multicast vlan 20
Vlan    Mac Address       Type       Ports
----    -----------       ----       -----
  20    0100.5e09.0909    USER        Fa0/1 Fa0/13 Fa0/14

CAT2#sh mac-address-table multicast vlan 30
Vlan    Mac Address       Type       Ports
----    -----------       ----       -----
  30    0100.5e09.0909    USER        Fa0/1 Fa0/7 Fa0/13 Fa0/14
```

Issue: Ping the multicast group 229.9.9.9 from R4 and CAT1

Solution:

When fulfilling this configuration requirement, carefully determine whether there is an RPF lookup problem on any of the routers. Since the PING is originating from router R4, the multicast packets will get forwarded to R1 and R1 will then forward them out all its interfaces. All RPF lookups should operate properly. There is not a hidden RPF lookup issue in this Scenario. You should get echo replies from R1, R2, R3, CAT2 and R5.

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".