NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II

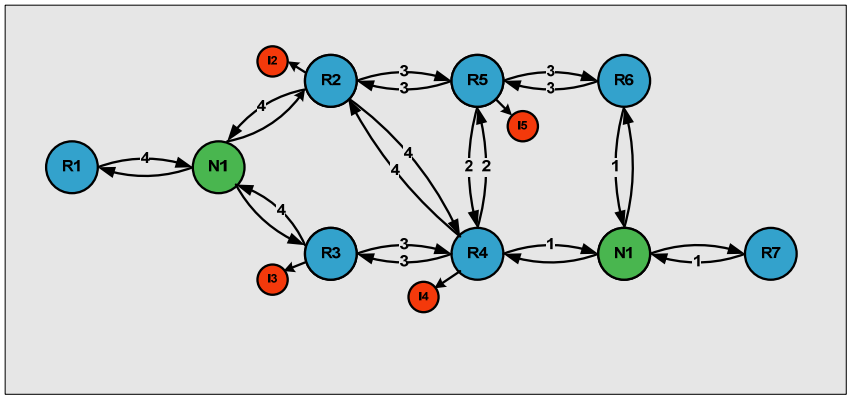

Scenario 19
ANSWER KEY

FOR

CCIE® CANDIDATES

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

Disclaimer

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

Cisco Non-Disclosure Agreement Compliance

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

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

DOiT-V6 Scenario 19: Spot the Issue Answer Key

Table of Contents

19.1 Frame-Relay 6
19.2 Catalyst Configuration 7
19.3 OSPF 10
19.4 RIP 12
19.5 EIGRP 13
19.6 Redistribution 14
19.7 BGP 16
19.8 Traffic Optimization part 1 18
19.9 Traffic Optimization part 2 19
19.10 IPv6 BGP 19
19.11 QOS 24
19.12 Catalyst Specialties 26
19.13 Gateway Redundancy 27
19.14 Multicast 29

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

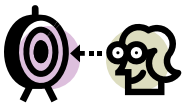

Goals and Restrictions

- · IP subnets on the diagram belong to network 172.10.0.0/16
- · Do not use any static routes except those used for backup purposes
- · Advertise Loopback interfaces with their original mask
- · Do not change any prefix masks
- · Do not use policy-routing
- · 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 on the diagram belong to network 172.10.0.0/16.

This is specified to make sense of the two octets supplied in the diagram.

Do not use any static routes except those used for backup purposes.

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

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

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

Do not use policy routing.

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

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

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

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

This is a key goal to observe. This requires that all of your 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. However, 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 relaxes the requirement for full-reachability.

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

19.1 Frame-Relay

HIDDEN ISSUES TO SPOT WITH THE FRAME-RELAY CONFIGURATION

Issue: Use only the PVC's displayed on the diagram to fulfill this configuration.

Solution:

When you examine the diagram, you notice that the Scenario is based upon a hub and spoke Frame-Relay topology, with router R1 as the hub and routers R2 and R4 as the spokes. In order to fulfill the requirements of the task stated above disable Frame-Relay inverse arp on all Frame-Relay physical interfaces and statically map only the DLCIs referenced in the diagram.

Implementation:

- Disable inverse arp on physical interfaces of R1, R2, R3 and R4 by issuing the no frame-relay inverse-arp command.
- Map DLCIs 102 and 104 to physical Frame-Relay interface Serial0/0 on router R1.
- Map DLCI 201 to physical Frame-Relay interface Serial0/0 on router R2.
- Map DLCI 401 to physical Frame-Relay interface Serial0/0 on router R4.

Verification:

To verify DLCI-to-interface mapping issue the **show frame-relay map** command:

```
R1#sh frame-relay map
Serial0/0 (up): ip 172.10.124.129 dlci 102(0x66,0x1860), static,
              CISCO, status defined, active
Serial0/0 (up): ip 172.10.124.130 dlci 102(0x66,0x1860), static,
              broadcast,
              CISCO, status defined, active
Serial0/0 (up): ip 172.10.124.131 dlci 104(0x68,0x1880), static,
              broadcast,
              CISCO, status defined, active
R1#
```

Issue: Use physical interfaces on subnet 172.10.124.128/25. Use physical and point-to-point interfaces on the subnet 172.10.43/0/24.

Solution:

For both of these subnets, carefully read ahead to the IGP section to see what IGPs are running over these Frame-Relay links. You will find that EIGRP is running over the 172.10.124.128/25 subnet and RIP is running over the 172.10.43.0/24 subnet. Since both of these routing protocols are Distance Vector routing protocols, you must consider the issue of split-horizon. For EIGRP, split-horizon is enabled on all Frame-Relay interfaces. You will want to disable split-horizon on R1, the hub of the partial mesh subnet, 172.10.124.128/25. For RIP, split horizon is disabled by default on physical Frame-Relay interfaces.

Manually enable split-horizon on the link that is configured with the physical interface on the 172.10.43.0/24 subnet.

Implementation:

To disable split-horizon for EIGRP on interface Serial0/0 of router R1 issue the command **no ip split-horizon eigrp 124**. To enable split-horizon for RIP on interface Serial0/0 of router R3 issue the command **ip split-horizon**.

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".

19.2 Catalyst Configuration

HIDDEN ISSUES TO SPOT WITH THE CATALYST CONFIGURATION

Issue: Do not advertise VLAN's in this scenario.

Solution:

If you are told not to advertise VLAN's in a scenario, this is directing you on how to configure VTP. If you are instructed to NOT advertise VLAN's, configure VTP transparent mode on both of your Catalyst switches. Catalyst switches configured in VTP transparent mode will not advertise any VLAN's that are created on them. To configure VTP in transparent mode issue the **vtp mode transparent** command on CAT1 and CAT2.

Verification:

```
CAT1#sh vtp status
VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 15
       VTP Operating Mode               : Transparent
CAT1#
```

Issue: Use the ISL protocol for trunking.

Solution:

If you are told to use the ISL protocol for trunking, you know that you will not need to configure any of the extensive 802.1Q features such as QOS or 802.1Q tunneling. To configure a switch port as an ISL trunk, issue the **switchport trunk encapsulation isl** command under interface configuration for that port.

Verification:

```
CAT1#show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/13      on           isl            trunking      1
Fa0/14      on           isl            trunking      1
Fa0/24      desirable    n-isl          trunking      1
CAT1#
```

Issue: Configure the required VLANs as shown on the diagram and in the following table:

Router	Interface	VLAN
R1	FastEthernet0/0	-
R2	Ethernet0	VLAN10
R3	-	-
R3	FastEthernet0/0	VLAN10
R4	-	-
R5	Ethernet0	VLAN20
R6	FastEthernet0/0	VLAN20
FRS	Ethernet0	VLAN10
CAT1	-	VLAN10
CAT2	-	VLAN20

Solution:

Use the physical layer diagram to determine which router ports are connected to which switch ports. Associate access ports with the required VLANs with the command **switchport access vlan x**. If the VLAN does not yet exist, this command will create it. It is a good practice to nail down the switch mode on access ports with the command **switchport mode access**. Ports that are associated with multiple VLANs are trunk ports; the port connected to R6 F0/0, for example. Note that when we indicate that VLAN 10 is associated with CAT1, for example, we mean that it has an Interface VLAN 10.

Verification:

You can verify access port assignment with the command **show vlan brief**, as you see here for CAT2:

```
CAT2#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/4, Fa0/5
                                                Fa0/6, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12, Fa0/16, Fa0/17
                                                Fa0/18, Fa0/19, Fa0/20, Fa0/21
                                                Fa0/22, Fa0/23, Fa0/24, Gi0/1
                                                Gi0/2
10   VLAN0010                         active    Fa0/3, Fa0/7
20   VLAN0020                         active
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
CAT2#
```

Note that by default, all access ports are in VLAN 1. Only ports F0/3 and F0/7 are configured as access ports on CAT2.

Issue: Configure the FRS Ethernet interface with an IP address of 172.10.23.10/24.

Solution:

This configuration step is setting the stage for upcoming IGP configuration requirements. See the EIGRP discussion, below for more details.

Issue: Make CAT1 the root bridge for VLAN 10. Put the Fa0/13 interface on CAT2 into the Spanning Tree blocking state. Do not use spanning tree path-cost manipulation to accomplish this task.

Solution:

CAT1 and CAT2 are connected on ports fa0/13 and fa0/14. With these two connections, a looped topology exists between the two Catalysts. This is a necessary condition for placing a Spanning Tree port into a blocking state. You must have a looped topology. If you have no looped topology in your Spanning Tree, you will have no blocked ports. If CAT1 is made the root bridge of the Spanning Tree running on VLAN 10, none of CAT1's ports can be placed in a blocking state. A root bridge will never have any of its ports in a blocked state (unless it is looped to itself with a cross-over cable or hub).

The fa0/13 port on CAT2 could be placed in a blocking state by increasing the fa0/13 port path-cost on CAT2 to be higher that the fa0/14 port path-cost on CAT2 as well. However, this method of placing a port into a blocking state is restricted by the task. If you can't manipulate a port's path-cost, the next parameter to determine whether a port goes into a blocking state is the Bridge-ID. This parameter will be of no use because the Bridge-ID's to be compared by CAT2 are the same on each link (CAT1's Bridge ID).

After comparing Catalyst switch Bridge-ID's the next parameter to compare is the port priority. A higher port priority advertised by the Designated Bridge on a segment will cause the remote end to go into a blocking state. Since CAT1 is the root bridge of VLAN, it is the designated bridge for all segments it is attached to. The default port priority is 128. Therefore, configure a port priority of 129 or higher on the fa0/13 port of CAT1. This will be advertised to CAT2 and CAT2 will place its fa0/13 port into a blocking state. Issue the **command spanning-tree vlan 10 port-priority 144** on CAT1 F0/13.

Verification:

```
CAT1#show spanning-tree vlan 10 detail

 VLAN0010 is executing the ieee compatible Spanning Tree protocol
  Bridge Identifier has priority 24576, sysid 10, address 000a.b7f7.7900
  Configured hello time 2, max age 20, forward delay 15
  We are the root of the spanning tree
  Topology change flag not set, detected flag not set
```

To verify that port FastEthernet0/13 on CAT2 is in blocking state, issue the **show spanning-tree vlan 10** command:

```
R5#sh spanning-tree vlan 10
[skipped]

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/3            Desg FWD 19        128.3    P2p
Fa0/7            Desg FWD 100       128.7    Shr
Fa0/13           Altn BLK 19        128.13   P2p
Fa0/14           Root FWD 19        128.14   P2p
R5#
```

Issue: A PC with the NIC MAC address 00-07-85-92-D0-E7 is connected to port Fa0/10 of CAT2 on the default VLAN. Make sure only that PC is allowed to access port Fa0/10.

Solution:

To restrict access to only the data-link address listed, configure the following two interface configuration commands on the fa0/10 port of Cat2: **switchport port-security mac-address 0007.8592.D0E7** followed by **switchport port-security**.

Verification:

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

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".

19.3 OSPF

HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

Issue: Place the PPP link between R3 and R5 in area 0.

Solution:

When you configure the PPP link between R3 and R5, host routes appear on both routers listing the IP address of the other end of the PPP connection as a "Connected" routing table entry. This is called the PPP "peer neighbor route". If you want to eliminate this route, enter the command **no peer neighbor-route** under the PPP interface.

Issue: Configure OSPF area 10 on VLAN 20 between R5 and R6. Do not elect a DR/BDR on this subnet. Make sure OSPF packets are exchanged without use of a multicast address due to security reasons.

Solution:

If you can't elect a DR or BDR on VLAN 20, then you cannot use the OSPF network types broadcast or non-broadcast. This leaves you with the following OSPF network types: point-to-point, point-to-multipoint and point-to-multipoint non-broadcast. Of these three, both point-to-point and point-to-multipoint use the 224.0.0.5 multicast for advertising HELLO messages. The OSPF network type point-to-multipoint non-broadcast does not use the 224.0.0.5 multicast at all. Therefore, configure VLAN 20 using the OSPF network type point-to-multipoint non-broadcast. Remember to configure neighbor statements for point-to-multipoint non-broadcast to identify the unicast destination of OSPF packets.

Verification:

```
R5#show ip ospf neighbor detail
[skipped]
Neighbor 172.10.106.1, interface address 172.10.65.6
    In the area 10 via interface FastEthernet0/0
    Neighbor priority is 0 (configured 0), State is FULL, 13 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:01:59
    Neighbor is up for 05:03:07
    Index 1/2, retransmission queue length 0, number of retransmission 2
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

Issue: Create Loopback 106 on R6 and place it in area 600.

Solution:

Since R6 possesses no direct link to OSPF Area 0, a virtual-link must be configured over Area 10, allowing the Area 600 prefix to be learned by all OSPF routers. Remember to include the virtual-link in the Area 0 authentication configuration.

Verification:

To verify that virtual link is up issue the **show ip ospf virtual-links** command. The "up" indication on the first line of the output can be deceiving; look for the "Adjacency State Full".

```
R5#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 172.10.106.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 10, via interface FastEthernet0/0, Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
    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
  Simple password authentication enabled
```

Issue: Use clear text authentication on area 0. Password "nmc".

Solution:

This authentication configuration is applied to all interfaces assigned to Area 0, including all virtual-links configured in this scenario

Issue: Use md5 authentication on area 10. Use password "rsnmc".

Solution:

This authentication configuration is applied to all interfaces assigned to Area 10, but NOT to any virtual-links configured in this scenario that may transit OSPF Area 10.

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".

19.4 RIP

HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

Issue: Configure RIP version 2 between only R5 and CAT2. Set the gateway of last resort on CAT2 only if 172.10.124.128/25 is in the R5 routing table

Solution:

To fulfill this configuration requirement, enter the following command under the router rip configuration mode: **default-information originate route-map rip-default**. The route-map "rip-default" will match on an access-list permitting only the 172.10.124.128/25 prefix. The effect of this configuration will be to allow router R5 to advertise to CAT2 a 0.0.0.0/0 route only if R5 possesses the 172.10.124.128/25 prefix in its local routing table. The 172.10.124.128/25 prefix is an EIGRP prefix that will be learned by R5 through redistribution.

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".

19.5 EIGRP

HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

Issue: Configure EIGRP AS 100 between routers R1, R2 and R4. Do not allow multicast EIGRP traffic on the subnet between routers R1, R2 and R4.

Solution:

As mentioned in the Frame-Relay section, make sure that split-horizon is disabled on the Frame-Relay physical interface of R1. R1 is the hub of a hub and spoke topology on the 172.10.124.128/25 subnet. Spilt-horizon is enabled by default on all interface types for EIGRP. Therefore, you must manually disable split-horizon on router R1. To prevent EIGRP traffic from being multicasted on the 172.10.124.128/25 subnet, configure neighbor statements on routers R1, R2 and R4. Unlike RIP, do not put the EIGRP interfaces that will unicast traffic into a passive state.

Issue: Summarize the networks below with the most optimal mask. Make sure you have a summary for these loopbacks on R1 and R4 only, not on R2:

- ip address 172.10.25.89 255.255.255.252
- ip address 172.10.25.93 255.255.255.252
- ip address 172.10.25.97 255.255.255.252

Solution:

EIGRP summarization is performed at the interface level. The key challenge to this task is determining precisely which interface to place the summarization command on. The three prefixes to be summarized reside on router R2. If you place the summary command on the Frame-Relay interface of R2, you will end up with a summary of the three loopbacks on R2, referencing the NULL0 interface, as well as in the forwarding tables of R1 and R4. To prevent the summary from appearing in the R2 routing table, allow the three prefixes to be advertised from R2 to R1 and then perform the summarization on router R1. R1 will then have the summary referencing the NULL0 interface, and R4 will learn the summary from R1. By placing the summary on router R1, the configuration requirements are fulfilled. You can then add a distribute-list inbound on R2 to block the summary created on R1.

Issue: Configure EIGRP AS100 on the VLAN 10 subnet 172.10.23.0/24 between R2, R3 and FRS. EIGRP AS100 on FRS must be configured with the "network 172.10.0.0" statement. Make sure FRS does not advertise loopback 172.10.100.0/24. Do not use any route filtering techniques.

Solution:

Configure **eigrp stub receive-only** under the EIGRP routing process on FRS.

Issue: Ping 172.10.100.1 from the rest of the network using address 172.10.23.100.

Solution:

This creates a NAT configuration requirement on FRS.

- Configure interface Loopback0 on FRS as the ip nat inside interface.
- Configure interface Ethernet0 on FRS as the ip nat outside interface.
- Configure NAT translation with the command **ip nat inside source static 172.10.100.1 172.10.23.100**.

Verification:

Issue the **show ip nat translations** command:

```
FRS#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
--- 172.10.23.100      172.10.100.1       ---                ---
FRS#
```

Issue: Configure subnet 172.10.32.0/24 on VLAN 10 between R3 and CAT1. CAT1 should not run the "ip routing" process. CAT1 should be reachable from the rest of the network.

Solution:

Configure the 172.10.32.0/24 subnet as a secondary address. In order to advertise this prefix to R2, you will need to disable split-horizon for EIGRP 124 on Interface F0/0. You could configure **ip default-gateway** on CAT1. You may find that this is unnecessary, because with ip routing disabled the CAT will arp for all destinations, and proxy-arp on R3 may take care of routing its packets.

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".

19.6 Redistribution

HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION

In Scenario 19 the core protocol is EIGRP. It spans almost the whole topology. RIP is a stub area, and OSPF provides transit between the RIP and EIGRP domains. Router R5 is a redistribution point between the OSPF and RIP domains.

RIP only sends two networks into OSPF, 172.10.56.0/24 and 172.10.120.0/24. Notice that the redistribution of RIP into OSPF on R5 does not result in an E2 route on R3 for 172.10.64.0/24. The OSPF network type point-to-multipoint models that link a collection of point-to-point links, and it advertise it as a collection of /32, host routes, not as a /24. Because OSPF does not accept RIP's version of the link as a /24, CAT2's address 172.20.65.10 becomes unreachable from off the link. One remedy is to do an inter-area summary to 172.10.65.0/24.

Redistribution from OSPF into RIP is not required for full reachability. Instead, RIP generates a conditional default route (0.0.0.0/0) into the RIP domain.

One way to test that your redistribution satisfies the goal of universal connectivity is to run a TCL script like the one below on each router. TCL scripting support is available in the IOS versions used here on routers R1, R2, R3, R4, R5 and R6. The simple script below lists all of the IP addresses in our pod. It can be built once in notepad, and then pasted into each router to automate pings. There is a paper on TCL scripting available in the READiT section of the Netmasterclass website. Some addresses are used in later tasks and may not be reachable at this point. Run **tclsh** in privileged mode, paste the script below, and then issue the command **tclq.**

Simple TCL script to test reachability:

```
foreach addr {
1.1.1.1
172.10.124.129
172.10.101.1
172.10.124.130
172.10.23.1
172.10.23.2
172.10.102.1
172.10.25.97
172.10.25.93
172.10.25.89
172.10.32.3
172.10.35.3
172.10.43.3
172.10.23.3
172.10.103.1
4.4.4.4
172.10.124.131
172.10.43.4
172.10.104.1
172.10.35.5
172.10.105.1
172.10.65.5
172.10.106.1
172.10.65.6
172.10.23.10
172.10.100.1
172.10.23.100
172.10.32.10
172.10.56.10
172.10.120.1
172.10.65.10} {ping $addr}
```

19.7 BGP

HIDDEN ISSUES TO SPOT WITH THE BGP CONFIGURATION

Issue: Configure AS 23 between R2 and R3. Configure AS 4 on R4. Peer AS 23 and AS 4 between R2 and R4 as well as R3 and R4. Do not use loopback interfaces for peering. Advertise network 4.4.4.0/24 in AS4. Advertise network 172.10.23.0/24 in AS23.

Solution:

These are fairly straightforward peering and advertising directions that are setting the stage for the remaining BGP tasks. For configuration details, see the online SHOWiT engine.

Note that R2-R4 eBGP peering is performed over IP subnet `172.10.124.128/25` over Frame Relay where R2 and R4 are the spokes and R1 is the hub. As IP packets forwarded between spokes hub decreases TTL by 1, so even that R1 and R2 are on the same subnet eBGP multihop needs to be set to 2 for BGP session to be successfully established.

Verification:

Verify peering with the command **show ip bgp summary**. Here is the result on R4. The number of prefixes learned, at the end of each line, indicates a good peering. Indications of “Active” or “Idle” would indicate peering problems.

```
R4#show ip bgp summary
BGP router identifier 172.10.25.1, local AS number 4
[output removed for brevity]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
1.1.1.1         4     1     156     157        5    0    0 02:31:36        1
172.10.43.3     4    23     159     159        5    0    0 02:32:11        1
172.10.124.130  4    23     157     158        5    0    0 02:32:08        1
R4#
```

Issue: Outbound traffic from a PC connected to the 172.10.23.0/24 subnet destined to the 4.0.0.0/24 network should flow through R2.

Solution:

This task will influence the HSRP configuration that is specified later in this Scenario.

Issue: Incoming traffic from the 4.0.0.0/24 network to a PC connected to the 172.10.23.0/24 subnet should flow through R2. If the frame-Relay link on R2 goes down, the aforementioned traffic should pass through R3. Return the traffic pattern through R2 when frame relay link on R2 is back up.

Solution:

This task is also related to the HSRP configuration that is discussed later in this Scenario. See the diagram below for more detail on the BGP topology.

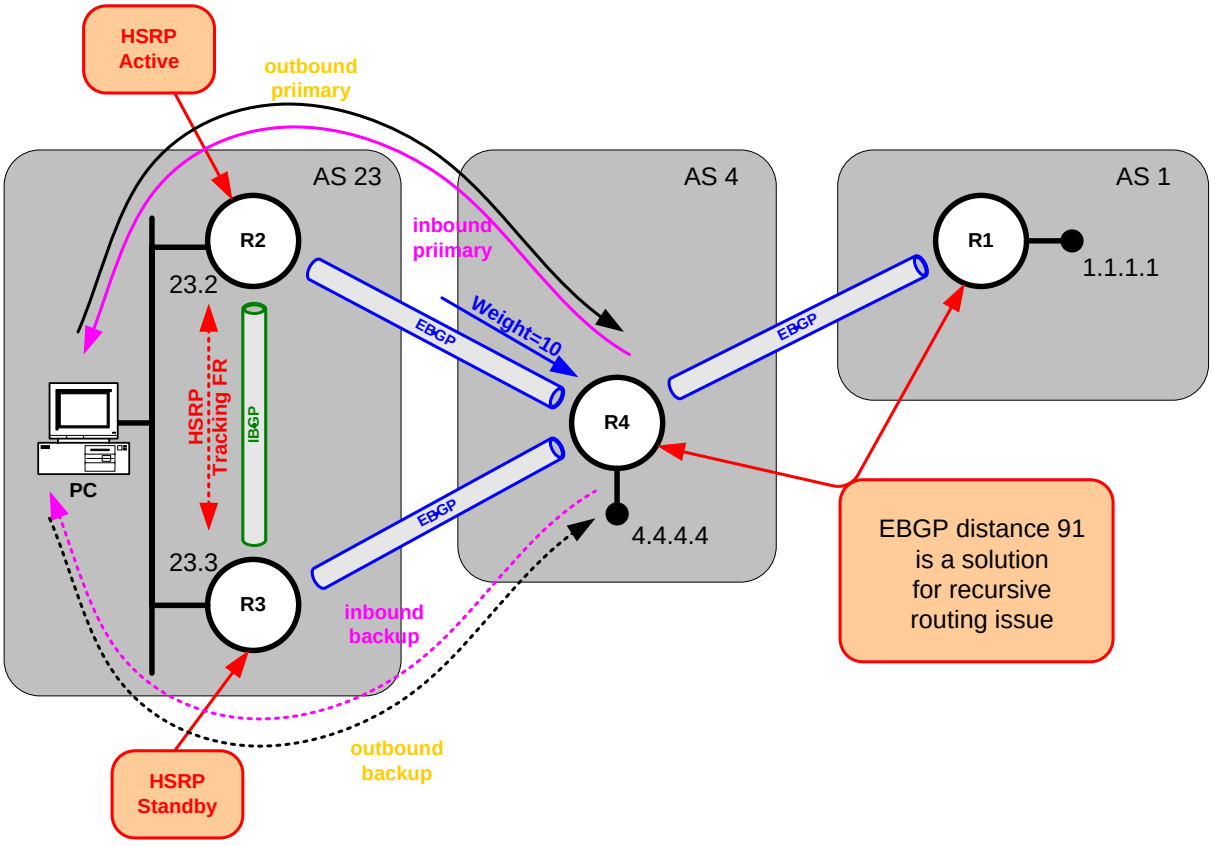

Issue: Use the minimal number of bgp decision steps to accomplish this task.

Solution:

Use the BGP administrative weight since it is the first attribute to be compared in the BGP path selection process. Since the BGP administrative weight is the first attribute compared between two candidate BGP paths, it fulfills the configuration requirement of using "the minimal number of bgp decision steps" to accomplish the stated task. Configure weight on R4 to the prefixes received from R2 by issuing the command n**eighbor 172.10.124.130 weight 10**.

Verification:

Issue the **show ip bgp** command:

```
R4#show ip bgp

BGP table version is 25, local router ID is 172.10.104.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
r> 1.1.1.0/24       1.1.1.1                  0             0 1 i
*> 4.4.4.0/24       0.0.0.0                  0         32768 i
r> 172.10.23.0/24   172.10.124.130           0            10 23 i
r                   172.10.43.3              0             0 23 i
R4#
```

Issue: Configure AS 1 on R1 and peer it with AS4 using loopbacks 1.1.1.1 and 4.4.4.4. Advertise networks 1.1.1.0/24 and 4.4.4.0/24 at AS1 and AS4 respectively.

Solution:

This configuration creates a recursive routing situation. Since the 1.1.1.0/24 and the 4.4.4.0/24 prefixes are advertised via EBGP speakers, these prefixes administrative distance is set to 20. Since these same prefixes are used to form the EBGP neighbor relationship between R1 and R4, they need to be learned via an IGP like OPSF or EIGRP. (NOTE: These prefixes have already been assigned to the EIGRP routing process.) To eliminate the problem, set the administrative distance for these prefixes to 91 under the BGP routing process of routers R1 and R4 so that EIGRP will be the more preferred routing source.

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".

19.8 Traffic Optimization part 1

HIDDEN ISSUES TO SPOT WITH THE TRAFFIC OPTIMIZATION CONFIGURATION

Issue: A Distributed Director is connected to VLAN 20. The IP address of the director is 172.10.65.1. Configure R5 and R6 to supply the Distributed Director with BGP and IGP metrics for efficient traffic distribution. R5 and R6 must supply the routing metrics to only the Distributed Director specified above.

Solution:

Configure the following two commands on both R5 and R6 in global configuration mode: **ip drp server** and **ip drp access-group 1** where access-group 1 matches the complete IP address of the Distributed Director: 172.10.65.1.

19.9 Traffic Optimization part 2

HIDDEN ISSUES TO SPOT WITH THE TRAFFIC OPTIMIZATION CONFIGURATION

Issue: A web server is connected to port Fa0/15 of CAT2. Users should not configure their browser for any web proxy. Configure CAT2 to offload HTTP requests from the Web server. Check the diagram for IP address requirements.

Solution:

Configure the Web Cache Coordination Protocol (WCCP) on CAT2. Enter the following command under the VLAN 20 switched virtual interface on CAT2, **ip wccp web-cache redirect in**. Enable WCCP with the **ip wccp web-cache** command.

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

19.10 IPv6 BGP

HIDDEN ISSUES TO SPOT WITH THE IPv6 BGP CONFIGURATION

Issue: Assign IPv6 addresses.

Solution:

This task can be broken down into three steps:

- Entering the command **IPv6 unicast-routing** in global configuration mode.
- Assigning the required addresses to interfaces.
- Mapping far-side addresses on Frame-Relay multipoint interfaces.

These tasks are shown below for R2:

```
R2(config)#ipv6 unicast-routing

interface FastEthernet0/0
ipv6 address FEC0:23::2/64

interface Serial0/0
encapsulation frame-relay
 ipv6 address FEC0:124::130/64
 ipv6 address FE80::2 link-local
frame-relay map ipv6 FE80::4 201 broadcast
```

```
 frame-relay map ipv6 FEC0:124::129 201 broadcast
 frame-relay map ipv6 FEC0:124::131 201 broadcast
frame-relay map ipv6 FE80::1 201 broadcast
 no frame-relay inverse-arp
```

Note that both link-local and site-local addresses need to be mapped. As a general practice we hard-code link-local addresses whenever they have to be mapped. Unlike with IPv4 addresses, connected IPv6 addresses do not have to be mapped in order to be pinged.

Verification:

Make sure you can ping within the same subnet before moving forward. Can R2 ping all of the addresses on the connected frame link?

```
R2#ping fe80::2
Output Interface: Serial0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R2#ping fe80::1
Output Interface: Serial0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

R2#ping fe80::4
Output Interface: Serial0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/88/88 ms
Success rate is 0 percent (0/4)

R2#ping fec0:124::130
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:124::130, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#ping fec0:124::129
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:124::129, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms

R2#ping fec0:124::131
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:124::131, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/87/88 ms
```

Issue: Create the following IPv6 BGP peers using directly connected, site-local addresses.

Solution:

The routers are peered just as they were for the IPv4 BGP exercise, as shown in the diagram below.

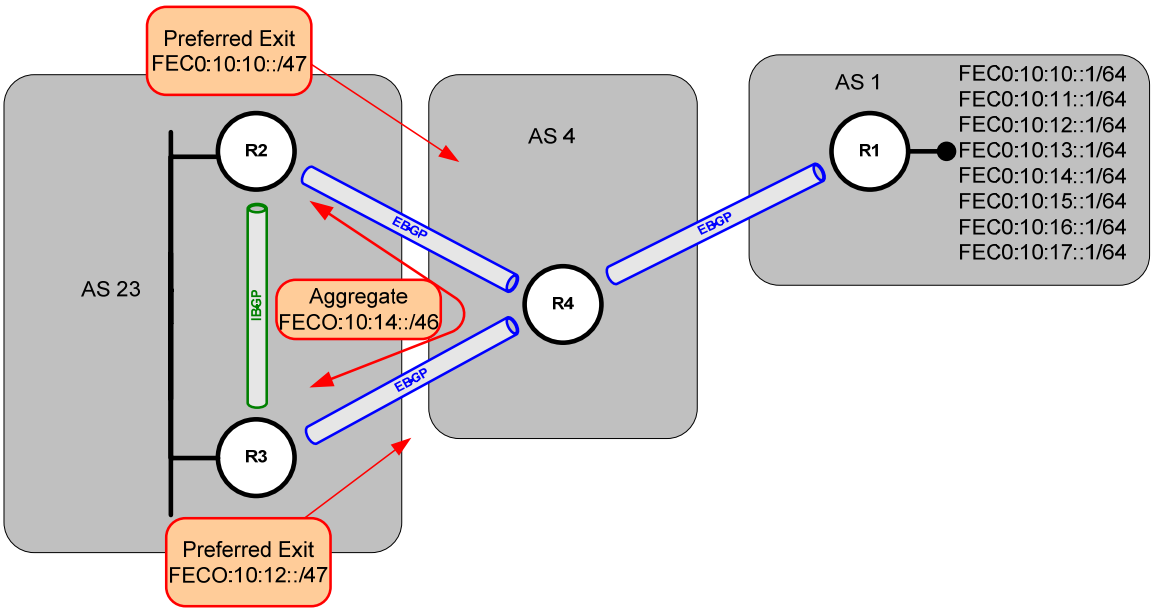

Under the primary BGP process, configure the neighbor and remote-as. Then, under address-family IPv6, activate the neighbor. Here is the relevant part of the configuration on R2:

```
router bgp 23
neighbor FEC0:23::3 remote-as 23
 neighbor FEC0:124::131 remote-as 4
 no auto-summary
 !
address-family ipv6
 neighbor FEC0:23::3 activate
 neighbor FEC0:124::131 activate
exit-address-family
```

Verification:

The command show BGP IPv6 summary can be used to verify the required peering. The output below shows 5 prefixes learned from R3 and 7 learned from R4, as it would at the end of the exercise. You might see 0 at this stage. Entries of “active” or “idle” under “PfxRcd” would indicate failed peering.

```
R2#show bgp ipv6 unicast summary
[output removed fro brevity]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
FEC0:23::3      4    23     427     432       17    0    0 00:44:18        5
FEC0:124::131   4     4     421     429       17    0    0 06:49:35        7
```

Issue: Advertise all connected IPv6 addresses into BGP using network statements. All IPv6 addresses should be reachable within the IPv6 BGP domain.

Solution:

Under the address-family IPv6 on each IPv6 router, issue a network statement for each connected, site-local prefix. We are using BGP to provide IPv6 reachability throughout the pod. Since we are peering to directly connected addresses, this should not cause recursive routing or peering address issues. Here is the relevant configuration from R2:

```
router bgp 23
!
 address-family ipv6
 network FEC0:23::/64
 network FEC0:124::/64
 exit-address-family
```

Issue: Add the following prefixes to Loopback 0 on R1. Advertise them into BGP with a single statement. AS23 should see only an aggregate for the highest four of these addresses. R1 should not see this aggregate.

Solution:

The simplest way to advertise these 8 addresses into BGP is to issue the command **redistribute connected** under the address-family IPv6 on R1. To avoid seeing the aggregate on R1, we create it on R4 with the **as-set** keyword. This keyword preserves the AS path attribute, causing to R1 to drop the update. If you are not used to seeing IP addresses in Hex, it may not be apparent that these four addresses fall on a very neat bit boundary. Here are the first 48 bits of each address:

```
FEC0:10:14 = 1111 1110 1100 0000 : 0000 0000 0001 0000 : 0000 0000 0001 0100 :
FEC0:10:15 = 1111 1110 1100 0000 : 0000 0000 0001 0000 : 0000 0000 0001 0101 :
FEC0:10:16 = 1111 1110 1100 0000 : 0000 0000 0001 0000 : 0000 0000 0001 0110 :
FEC0:10:17 = 1111 1110 1100 0000 : 0000 0000 0001 0000 : 0000 0000 0001 0111 :
                   16 bits             16 bits              16 bits
```

Of the first 48 bits in each address, only the last two vary. Since the first 46 bits are identical, we can summarize them as FEC0:10:14::/46.

Here is a TCL script you can use to test for universal IPv6 reachability.

```
foreach address {
FEC0:10:10::1
FEC0:10:11::1
FEC0:10:12::1
FEC0:10:13::1
FEC0:10:14::1
FEC0:10:15::1
FEC0:10:16::1
FEC0:10:17::1
FEC0:124::129
FEC0:23::2
```

```
FEC0:124::130
FEC0:23::3
FEC0:43::3
FEC0:124::131
FEC0:43::4
} {ping $address}
```

Issue: Traffic leaving AS 23 for prefixes FEC0:10:10::/64 and FEC0:10:11::/64 should have a next-hop of FECO:124::129. Traffic leaving AS23 for prefixes FEC0:10:12::/64 and FEC0:10:13::/64 should have a next-hop of FEC0:43::4.

Solution:

Local preference is commonly used within a dual-homed AS to indicate a preferred exit. We prefer R2 as the exit for the prefixes that start FECO:10:10::/47, so we raise the local preference on these prefixes as they arrive at R2. We raise the local preference on the prefixes that start FECO:10:12::/47 as they arrive at R3. Here is the relevant configuration for R2. Remember to reset your peers when you change policy. Some prefer the command **clear ip BGP ***. Others prefer to add the **soft** keyword.

```
ipv6 prefix-list LOCALPREF seq 5 permit FEC0:10:10::/47 ge 64 le 64

route-map LOCALPREF permit 10
 match ipv6 address prefix-list LOCALPREF
 set local-preference 200

address-family ipv6
neighbor FEC0:124::131 route-map LOCALPREF in
```

Verification:

In the partial BGP table below you will see that R2 prefers the EBGP paths to the first two prefixes and the IBGP paths to the other two, based on the local preference attributes. Note that the next hop for FEC0:10:10::/64 is R1's address, even though R2 is not peering with R1. BGP is smart enough to use a forwarding address when peers are on a shared network.

```
R2#show bgp IPv6 unicast
BGP table version is 30, local router ID is 172.10.25.97
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> FEC0:10:10::/64  FEC0:124::129                200      0 4 1 ?
*> FEC0:10:11::/64  FEC0:124::129                200      0 4 1 ?
*  FEC0:10:12::/64  FEC0:124::129                         0 4 1 ?
*>i                 FEC0:43::4               0    200      0 4 1 ?
*  FEC0:10:13::/64  FEC0:124::129                         0 4 1 ?
*>i                 FEC0:43::4               0    200      0 4 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".

19.11 QOS

HIDDEN ISSUES TO SPOT WITH THE QOS CONFIGURATION

Issue: Allocate a reservable bandwidth of 60 Kbps on the interfaces involved in this section. Send a PATH message from R4 to R5 requesting bandwidth reservation for telnet sourced from 172.10.43.4 port 5000 on R4 and destined to 172.10.35.5 port 23 on R5. Make sure you have a single reservation for a Guaranteed Bit Rate of 5 kbps allowing bursts up to 2 Kbytes. Verify the reservation setup with the command show ip rsvp reservation.

Solution:

Configure RSVP bandwidth on all interfaces that make up the path between R4 and R5. This includes two interfaces on router R3. You can send the PATH message for the specified application from R4 to R5 by configuring the **ip rsvp sender-host** command on R4. The **ip rsvp sender-host** command enables a router to simulate a host generating RSVP PATH messages. Configure the router R5 with the **ip rsvp reservation-host** command to behave as though it is continuously receiving an RSVP RESV message from the originator containing the indicated attributes. The **ip rsvp reservation-host** command enables a router to simulate a host generating RSVP RESV messages.

Configuration and Verification:

1. Reserve the bandwidth along the IP forwarding path between R4 and R5, that's how PATH IP messages will be forwarded:

On R4:

```
interface Serial0/0
ip rsvp bandwidth
!
interface Serial0/0.43 point-to-point
 ip address 172.10.43.4 255.255.255.0
frame-relay interface-dlci 403
 ip rsvp bandwidth 60
!
```

On R3:

```
interface Serial0/0
 ip address 172.10.43.3 255.255.255.0
ip rsvp bandwidth 60
!
interface Serial0/1
 ip address 172.10.35.3 255.255.255.0
ip rsvp bandwidth 60
!
```

On R5:

```
interface Serial1/1
 ip address 172.10.35.5 255.255.255.0
ip rsvp bandwidth 60
!
```

2. Simulate a PATH message on the router R4 and RESV messaged on R5:

On R4:

```
ip rsvp sender-host 172.10.35.5 172.10.43.4 TCP 23 5000 5 2
```

To see the PATH message turn on the following debug:

```
deb ip rsvp dump-messages path
```

```
*Mar  1 10:22:22.789: RSVP:     version:1 flags:0000 type:Path cksum:4E13 ttl:255
reserved:0 length:136
*Mar  1 10:22:22.789:  SESSION               type 1 length 12:
*Mar  1 10:22:22.789:   Destination 172.10.35.5, Protocol_Id 6, Don't Police , DstPort 23
*Mar  1 10:22:22.789:  HOP                   type 1 length 12:
*Mar  1 10:22:22.789:   Neighbor 172.10.43.4, LIH 0x01000402
*Mar  1 10:22:22.793:  TIME_VALUES           type 1 length 8 :
*Mar  1 10:22:22.793:   Refresh period is 30000 msecs
*Mar  1 10:22:22.793:  SENDER_TEMPLATE       type 1 length 12:
*Mar  1 10:22:22.793:   Source 172.10.43.4, udp_source_port 5000
*Mar  1 10:22:22.793:  SENDER_TSPEC          type 2 length 36:
*Mar  1 10:22:22.793:   version=0, length in words=7
*Mar  1 10:22:22.793:   Token bucket fragment (service_id=1, length=6 words
*Mar  1 10:22:22.793:       parameter id=127, flags=0, parameter length=5
*Mar  1 10:22:22.793:       average rate=625 bytes/sec, burst depth=2000 bytes
*Mar  1 10:22:22.793:       peak rate   =625 bytes/sec
*Mar  1 10:22:22.793:       min unit=0 bytes, max pkt size=4294967295 bytes
*Mar  1 10:22:22.797:  ADSPEC                type 2 length 48:
*Mar  1 10:22:22.797:  version=0  length in words=10
*Mar  1 10:22:22.797:  General Parameters  break bit=0  service length=8
*Mar  1 10:22:22.797:                                      IS Hops:1
*Mar  1 10:22:22.797:             Minimum Path Bandwidth (bytes/sec):193000
*Mar  1 10:22:22.797:                    Path Latency (microseconds):0
*Mar  1 10:22:22.797:                                   Path MTU:1500
*Mar  1 10:22:22.797:  Controlled Load Service  break bit=0  service length=0
```

On R5:

```
ip rsvp reservation-host 172.10.35.5 172.10.43.4 TCP 23 5000 FF RATE 5 2
```

To see the RESV message turn on the following debug:

```
debug ip rsvp dump-messages resv
```

```
*Mar 15 07:28:41.550: RSVP:     version:1 flags:0000 type:Resv cksum:3483 ttl:255
reserved:0 length:96
*Mar 15 07:28:41.550:  SESSION               type 1 length 12:
*Mar 15 07:28:41.550:   Destination 172.10.35.5, Protocol_Id 6, Don't Police , DstPort 23
*Mar 15 07:28:41.550:  HOP                   type 1 length 12:
*Mar 15 07:28:41.550:   Neighbor 172.10.35.5, LIH 0x12000407
*Mar 15 07:28:41.550:  TIME_VALUES           type 1 length 8 :
*Mar 15 07:28:41.550:   Refresh period is 30000 msecs
*Mar 15 07:28:41.550:  STYLE                 type 1 length 8 :
*Mar 15 07:28:41.550:   Fixed-Filter (FF)
*Mar 15 07:28:41.550:  FLOWSPEC              type 2 length 36:
*Mar 15 07:28:41.550:   version = 0 length in words = 7
*Mar 15 07:28:41.550:   service id = 5, service length = 6
*Mar 15 07:28:41.550:   tspec parameter id = 127, flags = 0, length = 5
*Mar 15 07:28:41.550:   average rate = 625 bytes/sec, burst depth = 2000 bytes
*Mar 15 07:28:41.550:   peak rate    = 625 bytes/sec
```

```
*Mar 15 07:28:41.550:   min unit = 0 bytes, max pkt size = 0 bytes
*Mar 15 07:28:41.554:  FILTER_SPEC          type 1 length 12:
*Mar 15 07:28:41.554:   Source 172.10.43.4, udp_source_port 5000
```

3. Verify the reservation made on all three routers:

```
R4#show ip rsvp reservation
To              From            Pro DPort Sport Next Hop      I/F      Fi Serv BPS
172.10.35.5     172.10.43.4     TCP 23    5000  172.10.43.3   Se0/0.43 FF RATE 5K
R4#

R3#show ip rsvp reservation
To              From            Pro DPort Sport Next Hop      I/F      Fi Serv BPS
172.10.35.5     172.10.43.4     TCP 23    5000  172.10.35.5   Se0/1    FF RATE 5K
R3#

R5#show ip rsvp reservation
To              From            Pro DPort Sport Next Hop      I/F      Fi Serv BPS
172.10.35.5     172.10.43.4     TCP 23    5000  172.10.35.5             FF RATE 5K
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".

19.12 Catalyst Specialties

HIDDEN ISSUES TO SPOT WITH THE CATALYST SPECIALTIES CONFIGURATION

Issue: Prohibit all traffic of Ethertype 8042 from entering VLAN20. Prohibit all SNA traffic from entering VLAN10. Do not apply any filtering configurations to Catalyst switch ports to accomplish this task.

Solution:

The key phrase to focus on in the configuration statements supplied above is "Do not apply any filtering configurations to Catalyst SWITCHPORTS.", emphasis on the word "switchports". You know you need to filter out the Ethernet type 8042 and SNA traffic, but whatever filtering technique you use cannot be applied to a specific switchport. Therefore, you must come up with a filtering solution that allows the filtering of the listed traffic that is applied in a manner other than at the switchport level. Applying the filter at the VLAN level using VLAN access-maps and VLAN-filters can fulfill this configuration requirement.

To use VLAN access-maps and VLAN filters, perform the following three-step configuration;

Step 1: Create a named "mac access-list"

```
mac access-list extended Etype-8042
 permit any any etype-8042
```

```
mac access-list extended Permit-any
 permit any any
mac access-list extended SNA
 permit any any lsap 0x0 0xD0D
```

Step 2: Create a vlan access-map

```
vlan access-map VLAN10 10
 action drop
 match mac address SNA
vlan access-map VLAN10 20
 action forward
 match mac address Permit-any
vlan access-map VLAN20 10
 action drop
 match mac address Etype-8042
vlan access-map VLAN20 20
 action forward
 match mac address Permit-any
```

Step 3: Apply the vlan access-map to a vlan filter

```
vlan filter VLAN10 vlan-list 10
vlan filter VLAN20 vlan-list 20
```

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

19.13 Gateway Redundancy

HIDDEN ISSUES TO SPOT WITH THE GATEWAY REDUNDANCY CONFIGURATION

Issue: Assign the IP address of 172.10.23.1 to the virtual gateway and make sure the mac-address associated with the virtual gateway is set to 0000.0c07.ac14.

Solution:

To use the MAC address specified above, assign the standby group to the value of "20." The standby group number is in decimal. It gets translated into hex and is used as the last two hex digits of the MAC address used by HSRP. When you translate "20" from decimal to hex, you end up with 0x14.

Verification:

Issue the **show standby** command:

```
R2#show standby
FastEthernet0/0 - Group 20
  State is Active
    5 state changes, last state change 02:53:22
  Virtual IP address is 172.10.23.1
  Active virtual MAC address is 0000.0c07.ac14
    Local virtual MAC address is 0000.0c07.ac14 (default)
R2#
```

Issue: Authenticate HSRP on the 172.10.23.0/24 subnet (password nmc). Make sure Hello packets are exchanged 3 times faster than by default.

Solution:

The default HSRP HELLO time is 3 seconds. Therefore, set it to 1 second. Also, configure HSRP authentication between the HSRP peers. Configure HSRP timers by issuing the command **standby 20 timers 1 4**. Authenticate HRSP adjacencies with the command **standby 20 authentication nmc**.

Verification:

Issue the **show standby** command and verify timer settings.

Issue: Select the preferred gateway that is most suitable for other tasks of this exam by using priority 150.

Solution:

The BGP section requires that devices on the VLAN 10 link prefer R2 as an exit, unless the Frame-Relay link is down. The following configuration on R2 helps achieve that result by making it primary, unless the tracked interface goes down. The decrement value of 51 would reduce the priority to 99, which is one below the default priority of 100 on R3.

```
interface FastEthernet0/0
 description VLAN10
 ip address 172.10.23.2 255.255.255.0
 no ip redirects
 duplex auto
 speed auto
 standby 20 ip 172.10.23.1
 standby 20 timers 1 4
 standby 20 priority 150
 standby 20 preempt
 standby 20 authentication nmc
 standby 20 track Serial0/0 51
```

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".

19.14 Multicast

HIDDEN ISSUES TO SPOT WITH THE MULTICAST CONFIGURATION

Issue: Join dense group 229.50.50.50 on interfaces Loopback 105 of R5.

Solution:

Configure the command **ip igmp join-group 229.50.50.50** under the specified interfaces of the above listed routers. Configure PIM on these interfaces to receive all the replies from a ping later on.

Issue: Make sure you can ping 229.50.50.50 from R1.

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 R1, the multicast packets will get forwarded to R4 and R4 will then forward them out all its interfaces.

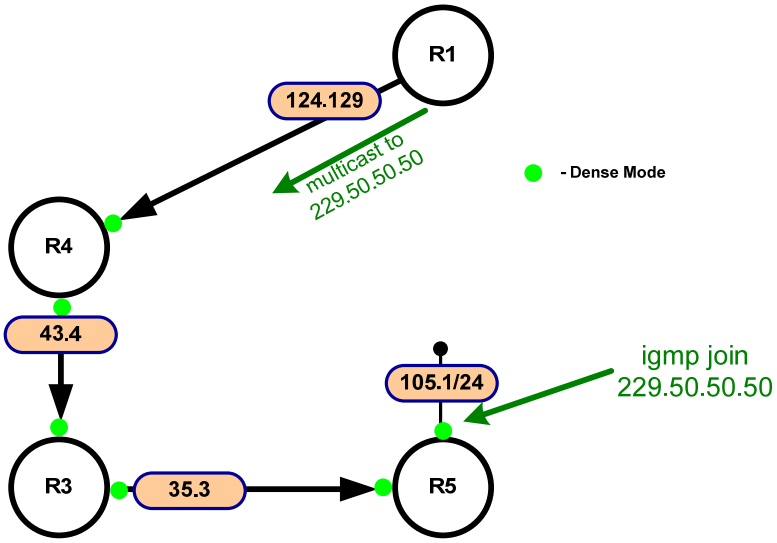

There is an RPF check failure on R3 towards the source of multicast traffic. R3 prefers the path to the source over F0/0, where traffic comes from R4, which is on the slower path to R1. To make the RPF check successful, a static mroute is added on R3, pointing to the IP address of R4's s0/0 interface as the next hop towards the source. Configure a static mroute on router R3 by issuing the command:

```
ip mroute 172.10.124.129 255.255.255.255 172.10.43.4.
```

Verification:

You can verify the static mroute with the command show ip mroute static, as shown here:

```
R3#show ip mroute static
Mroute: 172.10.124.129/32, RPF neighbor: 172.10.43.4
  Protocol: none, distance: 0, route-map: none
R3#
```

To verify its operation, enter the mtrace command to the source address:

```
R3#mtrace 172.10.124.129
Type escape sequence to abort.
Mtrace from 172.10.124.129 to 172.10.23.3 via RPF
From source (?) to destination (?)
Querying full reverse path...
 0  172.10.23.3
-1  172.10.23.3 PIM/Static  [172.10.124.129/32]
-2  172.10.43.4 PIM  [172.10.124.128/25]
-3  172.10.124.129
R3#
```

You see that the 172.10.23.3 outgoing interface is overridden by the static mroute. This can also be seen in the output of **show ip mroute** for this (S,G) entry:

```
R3#show ip mroute 229.50.50.50
IP Multicast Routing Table
[removed for brevity]

(172.10.124.129, 229.50.50.50), 00:00:31/00:02:57, flags: T
  Incoming interface: Serial0/0, RPF nbr 172.10.43.4, Mroute
  Outgoing interface list:
    Serial0/1, Forward/Dense, 00:00:31/00:00:00
R3#
```

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".