NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II

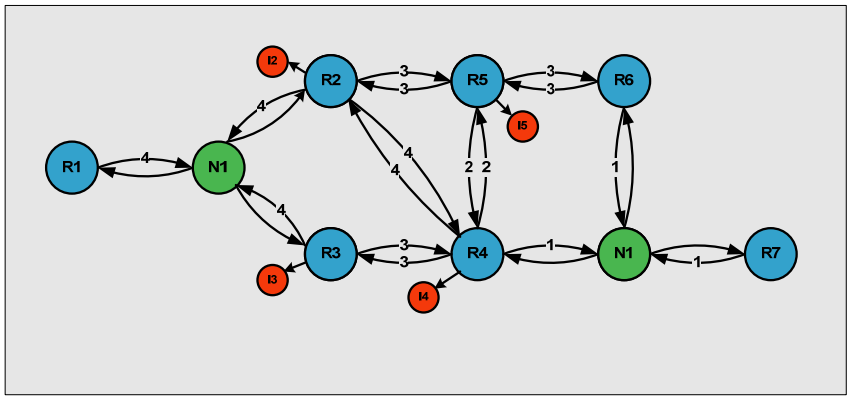

Scenario 17
ANSWER KEY

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 17: 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.

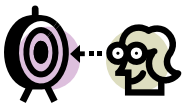

Goals and Restrictions

- IP subnets on the diagram belong to network 170.18.0.0/16.
- Do not use any static routes.
- Advertise Loopback interfaces with their original masks.
- Do not use a static default route 0.0.0.0/0.
- All IP addresses involved in this scenario must be reachable, unless specified otherwise.

Explanation of Each of the Goals and Restrictions:

IP subnets in the Scenario diagram belong to network 170.18.0.0/16

The third and forth octets of the IP addresses displayed on the diagram belong to 170.18.0.0/16.

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.

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

This requirement is primarily for the OSPF advertised loopbacks.

Do not use a static default route 0.0.0.0/0.

Other types of defaults may be allowed. The routing table should not have “S” in front of a default prefix.

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

This is a key goal to observe. This requires that all of your IGP’s are configured properly. Also, all of your routing policy tasks must be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using distribute-lists, route-maps and the distance command.

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

17.1 Serial interfaces

HIDDEN ISSUES TO SPOT WITH THE SERIAL INTERFACES CONFIGURATION

Issue: R1, R2 and R3 should be in the same subnet 170.18.123.0/24. R1, R4 and R3 should be in the same subnet 170.18.134.0/24. Use only point-to-point logical interfaces wherever possible and use physical interfaces otherwise. Use only the minimum number of DLCI's to fulfill this configuration.

Solution:

Keep in mind that you begin with an underlying full-mesh of Frame-Relay PVC's and you must end with using only the MINIMUM number of PVC's to fulfill the configuration requirement. The configuration tasks instruct you to create two subnets over the Frame-Relay cloud: the 170.18.123.0/24 subnet between routers R1, R2 and R3 and the 170.18.134.0/24 subnet between routers R1, R3 and R4. From this information, you determine that the PVC between R2 and R4 is not needed, so you can exclude this PVC from any further configuration.

The two subnets will create two hub-and-spoke topologies. You must determine which router will be the hub of each subnet. The answer to this question is found in the IGP configuration section, in particular, the OSPF configuration section. You are told to make R3 the DR of the 170.18.123.0/24 subnet. If R3 is the DR, it must be the hub. Once R3 is made the DR, the only router that can be the hub of the second subnet is R4. Therefore, R3 is the hub of the 170.18.123.0/24 subnet and R4 is the hub of the 170.18.134.0/24 subnet. R4 will be configured with a single physical interface. R3 will be configured with a physical interface for DLCI's 301 and 302 and a point-to-point subinterface for DLCI 304. R1 can be configured in one of two ways: (1) 2 point-to-point subinterfaces or (2) one point-to-point subinterface and one physical interface. R2 will be configured in one of two ways: (1) one point-to-point subinterface and (2) one physical interface.

Implementation:

- Use two point-to-point subinterfaces on R1 for each subnet. Use DLCI's 103 and 104.
- Use one point-to-point subinterface on R2. Use DLCI 203.
- Use physical interface on R3 for subnet 123.0 and point-to-point subinterface for subnet 134.0. Use DLCI's 301 and 302 for subnet 123.0, and DLCI 304 for subnet 134.0.
- Use physical interface on R4. Use DLCI's 403 and 401.

Verification:

Issue the show frame-relay map command on routers R1, R2, R3 and R4. Make sure there are no dynamic maps or maps to "0.0.0.0."

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

17.2 Catalyst Configuration

HIDDEN ISSUES TO SPOT WITH THE CATALYST CONFIGURATION

Issue: Configure the VLAN's referenced in the diagram and in the VLAN configuration table.

Solution:

Carefully review the entire Scenario. Closely examine the supplied diagram and any associated tables. Determine how you need to configure VTP, how to configure ports that are assigned as trunks and how to configure ports that are assigned as simply static VLAN ports. For any ports that are statically assigned to a VLAN, it is recommended to statically assign the following command to these types of ports: **switchport mode access**. The diagram below shows the required VLAN configuration.

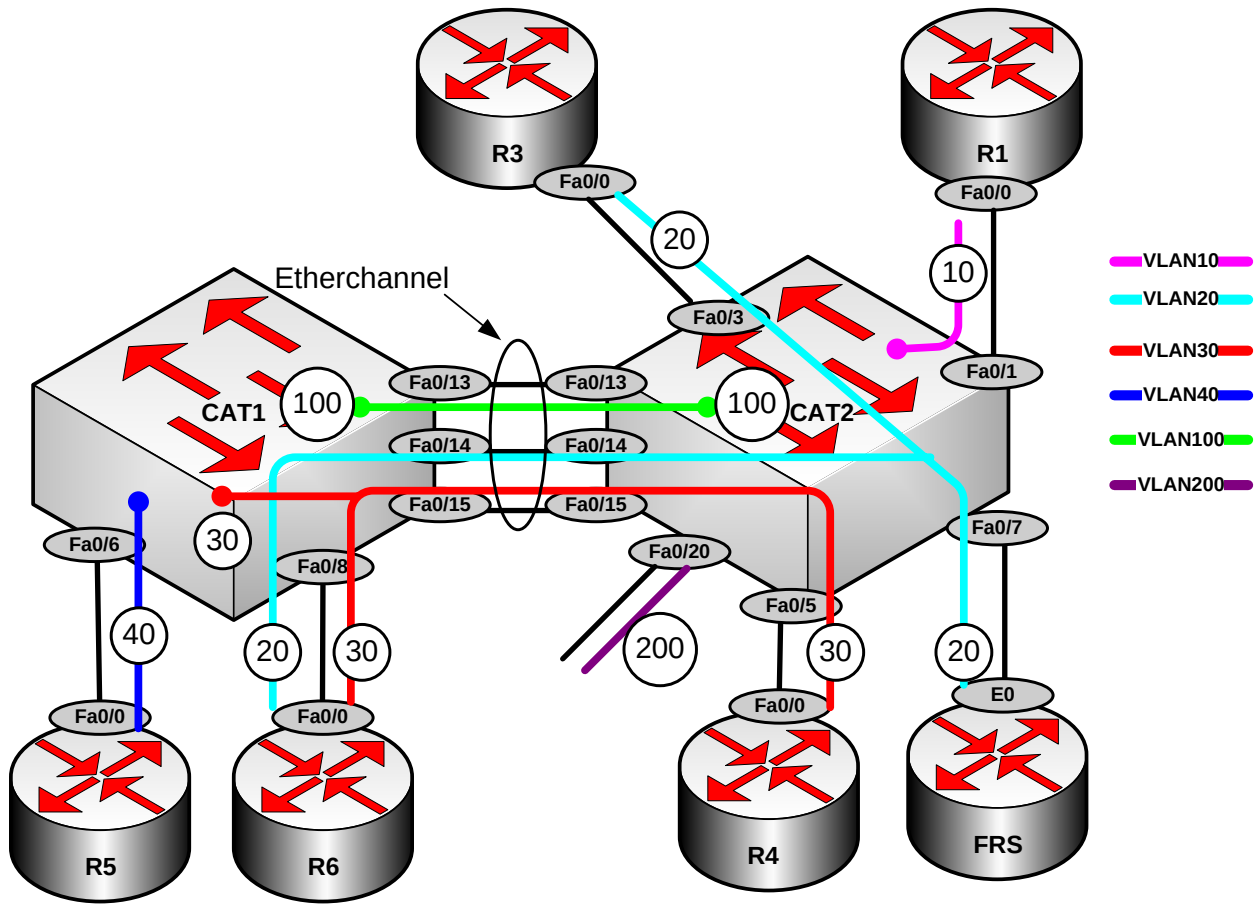

Implementation:

- Configure both switches CAT1 and CAT2 in VTP transparent mode by issuing the **vtp mode transparent** command.
- Create all the necessary vlans for this exercise: 10, 20, 30, 40 and 100 by issuing the **vlan 10,20,30,40,100** command.
- Create vlan 200 on CAT2 by issuing the **vlan 200** command.
- Create management interfaces for VLAN10, 30, 40 and 100 on related switches by issuing the **interface VLAN N** command and suppling an IP address according to the diagram.
- Configure access switch ports as follows:

```
interface FastEthernetX/X
       switchport mode access
       switchport access vlan N
```

- Configure trunk ports, allowing only the necessary vlans to be on a trunk, as follows:

On switch CAT1:

```
interface FastEthernet0/8
       switch trunk encapsulation isl
       switch mode trunk
       switch trunk allowed vlan 20,30
```

Verification:

To verify VTP mode issue the **show vtp status command**. Verify access port assignment with the command **show VLAN brief**, and verify trunk status with **show interface trunk**.

Issue: Make sure that you see the following output from the command show etherchannel summary.

This task requires you to bundle parallel links F0/13, F0/14 and F0/15 into an etherchannel using the Cisco aggregation protocol PAGP.

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".

17.3 OSPF

HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

Issue: Configure OSPF area 0 between R1, R2 and R3. Use an OSPF network type that elects a DR but does not require a neighbor statement. Make sure R3 is the DR.

Solution:

The OSPF network type that elects a DR but does require a neighbor statement is the OSPF broadcast network type. Therefore, configure the OSPF network type "broadcast" on the 170.18.123.0/24 subnet. Remember that this subnet maintains a hub and spoke topology. Therefore, set the spokes to an OSPF priority of 0 to prevent them from being elected as either a DR or BDR.

Verification:

Issue the **show ip ospf interface** command to verify the OSPF network type on the link, neighbors discovered and neighbors with adjacencies.

Issue: Configure OSPF area 0 on the subnet between R1, R3 and R4. Use an OSPF network type that elects a DR and requires neighbor statements.

Solution:

The OSPF network type that elects a DR and does require a neighbor statement is the OSPF non-broadcast network type. Therefore, configure the OSPF network type "non-broadcast" on the 170.18.134.0/24 subnet. Remember that this subnet maintains a hub and spoke topology. Therefore, set the spokes to an OSPF priority of 0 to prevent them from being elected as either a DR or BDR. Configure neighbor statements under router OSPF process on router R4 referencing spokes R1 and R3 as following:

```
neighbor 170.18.134.1
neighbor 170.18.134.3
```

Issue: For the subnet shared by R1, R3 and R4, make sure that the loss of a neighbor relationship is detected twice as fast as the default.

Solution:

To fulfill this requirement, configure one of the two following OSPF interface configuration commands; (1) **ip ospf hello-interval 15** and (2) **ip ospf dead-interval 60.** The values are set to 15 for the new HELLO interval and 60 for the new DEAD interval since the default HELLO and DEAD intervals for an OSPF non-broadcast network type are 30 and 120 respectively.

Issue: Place VLAN10 and Loopback 120 into OSPF area 1. Make sure OSPF sends the minimum information to CAT2.

Solution:

Since dynamic default routes are not ruled out by the scenario ground rules, the most straight-forward way to accomplish this task would be to make Area 1 totally stubby. Issue the command **area 1 stub** on CAT2 and the command **area 1 stub no-summary** on R1. Note that all routers in a stub area must have the stub keyword, but only the ABR must have the no-summary keyword. Making the area stub keeps out external prefixes (no type-5 LSAs). Adding the **no-summary** keyword keeps the ABR from sending Type-3 LSAs into the area.

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".

17.4 RIP

HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

Issue: Configure RIP Version 2 between R2 and R5. Send the minimum required information from R2 to R5

Solution:

The minimum required information would be a 170.18.0.0/16 summary, since all addresses that must be reached are from this subnet.

Issue: Allow only the following networks to be advertised to router R2 using the minimum number of access list statements.

- 192.80.2.0
- 192.80.3.0
- 192.88.2.0
- 192.88.3.0
- 170.18.102.0

Solution:

This configuration requirement can be fulfilled with the following access-list:

```
access-list 17 permit 192.80.2.0 0.8.1.0.
access-list 17 permit 170.18.102.0
```

This list allows the “8” bit in the second octet and the 1 bit in the third octet to vary, resulting in the required match. Since the task says “advertised to” the access-list should be applied to an outbound distribute-list under the RIP routing process of R5. The outbound distribute-list should explicitly reference the Serial 1/0 interface of 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”.

17.5 EIGRP

HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

Issue: Configure EIGRP AS 100 between R4, R3 and R6.

Solution:

EIGRP is being configured on two routers that are also configured with OSPF – routers R3 and R4. Since this configuration establishes two redistribution points between the same two routing protocols, there is a potential for routing loops to form or suboptimal paths to be selected.

However, EIGRP has a useful feature that avoids these negative effects of multiple redistribution points between the same two routing protocols. EIGRP assigns an administrative distance of 90 to all internally learned EIGRP prefixes and assigns an administrative distance of 170 to all prefixes that are redistributed into EIGRP.

By setting a higher administrative distance to all prefixes that are redistributed into EIGRP, routers running EIGRP and some other routing protocol will by default select the other routing protocol as the preferred routing source over an EIGRP routing process possessing the same prefix as an external route. Therefore, very little needs to be done here to avoid routing loops or suboptimal path selection. .

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”.

17.6 Redistribution

HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION

In this scenario the core protocols are OSPF and EIGRP. RIP is an edge protocol; it will not provide transit services. R3 and R4 are redistribution points between OSPF and EIGRP. Mutual redistribution at two or more points should always raise a red flag for route feedback. But, as discussed in the EIGRP section, we have little worry when one of the routing domains is EIGRP.

The worst-case scenario would be if network 170.18.105.0, for example, was redistributed from OSPF into EIGRP on, say, R3, and was then redistributed back into OSPF on R4. This could cause R2 to see R4 as an ASBR for a route in the RIP domain! We want to make sure that R2 always sees the RIP domain routes as coming from R5. One way to protect ourselves is to raise the administrative distance on R1 of external OSPF routes.

To make sure that you can reach all configured IP addresses, try using the TCL script shown on the following page. In privileged mode of each router, enter the command **tclsh**, then paste in this script. Enter **tclq** to leave this mode. You can quit a failing ping by holding down CTL and SHIFT, then hitting 6 twice.

```
foreach addr {
170.18.134.1
170.18.10.1
170.18.123.1
170.18.101.1

170.18.25.2
170.18.123.2
170.18.102.1

170.18.134.3
170.18.255.3
170.18.123.3
170.18.103.1

170.18.134.4
170.18.104.1
170.18.64.1
170.18.64.4

192.80.3.1
192.80.2.1
192.88.3.1
192.88.2.1
170.18.25.5
170.18.105.1

170.18.255.6
170.18.255.10
170.18.106.1
```

```
170.18.64.6

170.18.255.1

170.18.110.1
170.18.64.10

170.18.10.20
170.18.120.1
} {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".

17.7 Enhance Routing Stability

HIDDEN ISSUES TO SPOT WITH THE ROUTING STABILITY CONFIGURATION

Issue: Imagine that the F0/0 interface on R4 is flapping, causing instability throughout the network. Implement a feature on this interface that will isolate failures so that disturbances are not propagated

Solution:

We implemented the damping feature on this interface. You can find it documented at this link:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ftipevdp.htm

Verification:

After implementing event dampening, we went to the connected CAT2 interface f0/5 and shut/no shut it a few times. In order to get a response we lowered the thresholds and increased the half-life. Note in the output shown below that when the interface is dampened the connected interface is not in the routing table. It returns when the penalty expires.

```
R4#sh interface dampening
FastEthernet0/0 VLAN30
  Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
      5     259    TRUE       43      30     100     200     200   10158       0

R4#sh ip route connected
     170.18.0.0/16 is variably subnetted, 12 subnets, 3 masks
C       170.18.134.0/24 is directly connected, Serial0/0
C       170.18.253.0/24 is directly connected, Loopback100
C       170.18.44.0/29 is directly connected, Loopback0

FastEthernet0/0 VLAN30
```

```
  Flaps Penalty    Supp ReuseTm   HalfL  ReuseV   SuppV  MaxSTm    MaxP Restart
      5      55   FALSE       0      30     100     200     200   10158       0
R4#

R4#sh ip route connected
     170.18.0.0/16 is variably subnetted, 12 subnets, 3 masks
C       170.18.134.0/24 is directly connected, Serial0/0
C       170.18.253.0/24 is directly connected, Loopback100
C       170.18.44.0/29 is directly connected, Loopback0
C       170.18.64.0/24 is directly connected, FastEthernet0/0.
```

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

17.8 BGP

HIDDEN ISSUES TO SPOT WITH THE BGP CONFIGURATION

VLAN 100 IS USED FOR BGP BACKEND CONNECTIVITY ONLY. IT IS NOT PART OF ANY IGP ROUTING PROTOCOL AND DOES NOT HAVE TO BE REACHABLE. VLAN 40 IS USED FOR THE BGP CONNECTIVITY ONLY. IT IS NOT PART OF ANY IGP ROUTING PROTOCOL AND DOES NOT HAVE TO BE REACHABLE.

Issue: There are a number of tasks directing you to create complex BGP peerings among your devices.

Solution:

Read each task carefully, and try to draw out the topology. The target topology is shown in the diagram below:

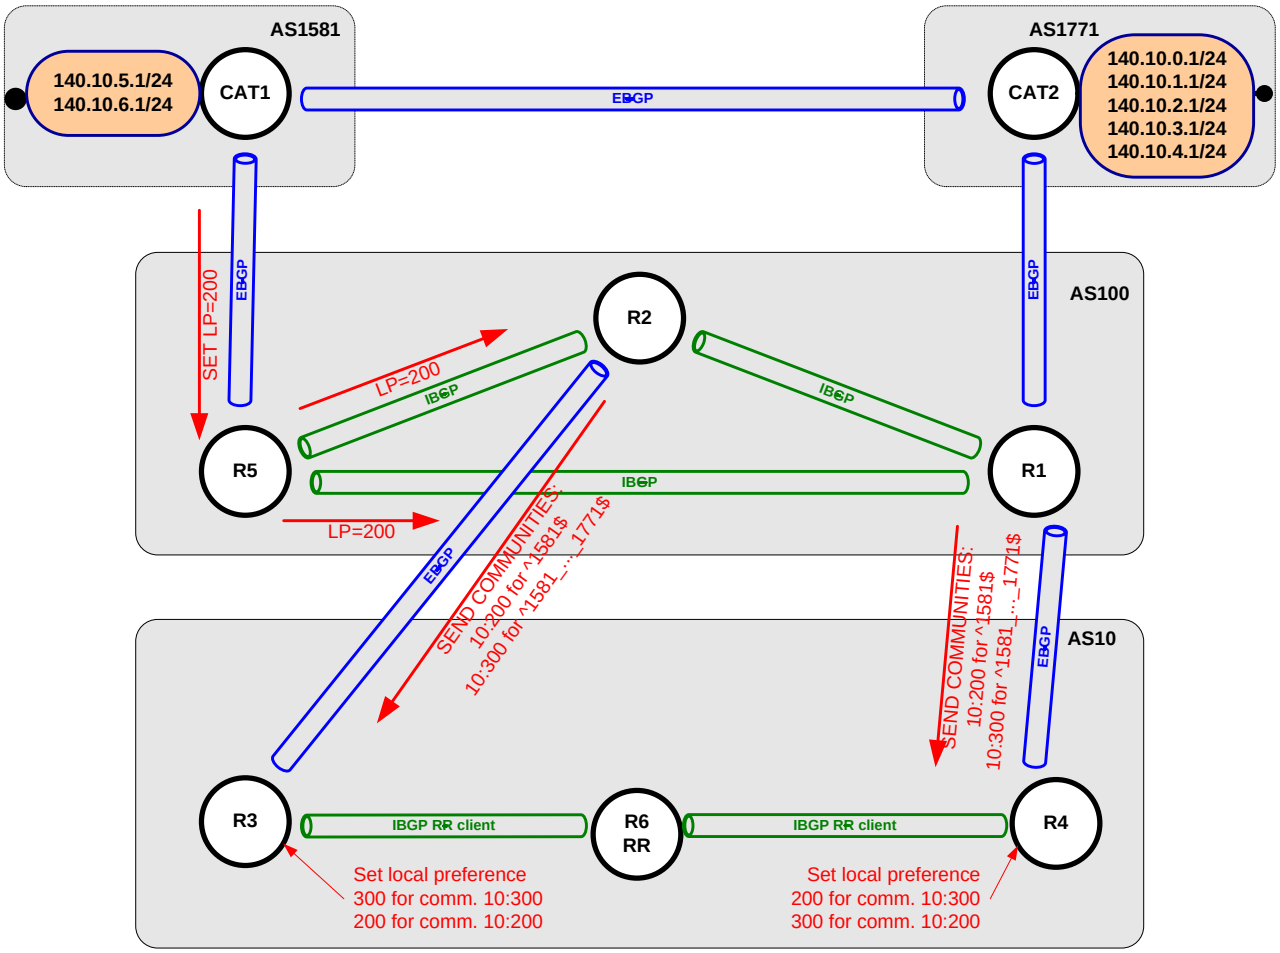

Issue: Configure BGP AS 10 on R3, R4 and R6. Do not use a full mesh in AS 10.

Solution:

In order to avoid configuring a full-mesh in AS 10, you can configure a route-reflector or confederation. If you are not directed any farther to configure a confederation, configure a route-reflector since it is simpler to configure. Configure R6 to be a route-reflector by configuring its neighbors as follows:

- neighbor 170.18.64.4 remote-as 10
- neighbor 170.18.64.4 route-reflector-client
- neighbor 170.18.255.3 remote-as 10
- neighbor 170.18.255.3 route-reflector-client

Issue: Prefer R5 as an exit point to networks 140.10.*.0/24 advertised earlier in this section.

Solution:

To make R5 the preferred exit point for the 140.10.*.0/24 networks, assign a higher local preference for these prefixes on router R5 and retain the default local-preference setting on all other routers. The default local-preference is 100 and a higher local-preference is more preferred. By setting all R5 learned 140.10.*.0/24 prefixes to a higher local-preference, R5 will act as the exit point for these prefixes in AS 100.

Verification:

Issue the **show ip bgp** command on router R1, and verify that R5 is a preferred exit point:

```
R1#sh ip bgp
BGP table version is 7, local router ID is 170.18.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  140.10.1.0/24    170.18.10.20             0             0 1771 i
*>i                 170.18.25.5              0    200      0 1581 1771 i
*  140.10.2.0/24    170.18.10.20             0             0 1771 i
*>i                 170.18.25.5              0    200      0 1581 1771 i
*  140.10.3.0/24    170.18.10.20             0             0 1771 i
*>i                 170.18.25.5              0    200      0 1581 1771 i
*  140.10.4.0/24    170.18.10.20             0             0 1771 i
*>i                 170.18.25.5              0    200      0 1581 1771 i
*  140.10.5.0/24    170.18.10.20                           0 1771 1581 i
*>i                 170.18.25.5              0    200      0 1581 i
*  140.10.6.0/24    170.18.10.20                           0 1771 1581 i
*>i                 170.18.25.5              0    200      0 1581 I
R1#
```

Issue: In AS 10, set the local-preference for all prefixes originating from AS 1771 and traversing AS 1581 to prefer R3 as a next hop. Also, set the local-preference for all prefixes originating from AS 1581 to prefer R4 as a next hop. Use values 200 and 300 to accomplish this task. Do not use the AS-PATH or the IP address prefix as match criteria to set the local-preference in AS 10.

Solution:

If you cannot set the local-preference in AS 10 by matching on AS-PATH strings or prefixes, you can use either the AS-PATH or address prefix as a matching criteria in AS 100 and set a unique community string to the desired prefixes in AS 100. The restriction stated above – "Do not use the AS-PATH or the IP address prefix as match criteria to set the local-pref in AS 10." – applies only to AS 10 and NOT to AS 100. Once the community is set in AS 100, it will be propagated to AS 10. Once it reaches AS 10, you can set the local-preferences to the specified values based upon community strings. Remember that for all BGP speakers that need to advertise a prefix with a community attached to it enter the following BGP neighbor command: **neighbor X.X.X.X send-community**. Check the diagram for the community and local preference values sent between the autonomous systems.

CAT1 and CAT2 are backbone simulators in this scenario; they are both external AS's to the rest of the network. CAT1 is AS 1581, CAT2 is AS 1771. Both ASs originate prefixes into your network, 4 prefixes from AS 1771 and 2 prefixes from AS 1581.

When these prefixes enter AS 100, R5 is selected as a preferred exit point in AS 100. Setting local preference 200 to all prefixes received on R5 from CAT1 does this. R5 will advertise the new local preference value to the other AS 100 peers.

In AS 10, the outbound traffic to the networks originated from AS1581 should be forwarded via R3, and traffic towards networks originated in 1771 and traversing 1581 should prefer R4 as an exit point. This will require us to set local preference to 300 for prefixes originated in 1771 and passing thru 1581, and 200 for others. On R4, local preference 300 should be assigned for prefixes originated in AS 1581, and local preference 200 should be assigned to others.

The restriction of not using as-path or ip address filters in AS 10 does not prohibit using them in AS 100. In AS 100 these prefixes will be classified using as-path access-lists and assigned communities 10:200 for prefixes originated in 1581 and 10:300 for prefixes originated in 1771 and traversing 1581. AS 10 will set local preferences according to these community numbers.

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".

17.9 Address Administration

HIDDEN ISSUES TO SPOT WITH THE ADDRESS ADMINISTRATION CONFIGURATION

Issue: Configure FRS to act as a host with the ip address of 170.18.255.1/24. Prefer R6 as a gateway for FRS. Do not use HSRP or VRRP. Do not use any static configuration.

Solution:

If FRS is going to act as a host device, you will then be able to disable ip routing in global configuration mode and rely on a gateway discovery protocol like the ICMP Router Discovery Protocol (IRDP). Configure IRDP on the VLAN 20 assigned interfaces of routers R3 and R6, and configure the following command in global configuration mode on FRS: **ip gdp irdp**. To make R6 the more preferred gateway on VLAN 20 for FRS, set R6 with a LOWER IRDP preference than R3. Use the interface configuration command **ip irdp preference X** on R6.

Implementation:

1. Configure R3 and R6 as IRDP information sources on VLAN20:

Router R3:

```
interface FastEthernet0/0
 description VLAN20
 ip address 170.18.255.3 255.255.255.0
 no ip proxy-arp
 ip irdp
 ip irdp maxadvertinterval 30
 ip irdp minadvertinterval 10
 ip irdp holdtime 90
 ip irdp preference 200
 ip pim sparse-dense-mode
 ip igmp join-group 229.10.10.10
 duplex auto
 speed auto
end
```

Router R6:

```
interface FastEthernet0/0.20
 encapsulation isl 20
 ip address 170.18.255.6 255.255.255.0
 no ip redirects
 no ip proxy-arp
 ip nat inside
 ip irdp
 ip irdp maxadvertinterval 30
 ip irdp minadvertinterval 10
 ip irdp holdtime 90
 ip irdp preference 100
```

2. Configure FRS as an IRDP client: Issue the **ip gdp irdp** global configuration command.

```
ip gdp irdp
```

Verification:

Issue the **show ip route** command on FRS:

```
FRS#sh ip route
Gateway         Using  Interval  Priority   Interface
170.18.255.3    IRDP        5       200     Ethernet0
170.18.255.6    IRDP       15       100     Ethernet0

Default gateway is 170.18.255.6

Host              Gateway           Last Use    Total Uses  Interface
170.18.253.1      170.18.255.6          0:00            24  Ethernet0
```

Issue: All packets originating from FRS should have the source IP address changed at the first hop router. The source IP address must be in 170.180.0/16 range

Solution:

This is NAT problem without ever explicitly mentioning NAT. From a basic NAT configuration perspective, the R6 VLAN 20 interface will be the NAT inside interface and the R6 VLAN 30 interface will be the NAT outside interface.

Implementation:

Configure NAT on router R6:

1. Issue the **ip nat inside** command under interface FastEthernet0/0.20
2. Issue the **ip nat outside** command under interface FastEthernet0/0.30
3. Issue the **ip nat inside source static 170.18.255.1 170.18.255.10** global configuration command.

Implement same configuration on R3 to ensure that this requirement is satisfied in the situation when FRS looses its primary gateway and falls back to the secondary gateway (R3), and to support asymmetric traffic flows (i.e. communication between FRS and R2 – direct path FRS→R6…→R2, return path R2→R3→FRS). Note that since it is a static address translation scenario there is no need to use any of new 12.2T / 12.3T NAT features. To ensure that traffic forwarded via R3 (i.e. Frame Relay traffic) use `no-alias` option with `ip nat` command.

Router R3:

```
interface FastEthernet0/0
 ip nat inside
!
interface Serial0/0
 ip nat outside
interface Serial0/0.134 point-to-point
 ip nat outside
!
ip nat inside source static 170.18.255.1 170.18.255.10 no-alias
```

Verification:

Issue the **show ip nat translations** command on router R6:

```
R6#sh ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
--- 170.18.255.10      170.18.255.1       ---                ---
```

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".

17.10 IPv6 Routing Over IPv4

HIDDEN ISSUES TO SPOT WITH THE IPv6 ROUTING OVER IPv4 CONFIGURATION

Issue: Remove the IPv6 address from R1 Loopback 0. Assign the following IPv6 addresses in addition to those specified in the previous task.

Solution:

Since we are not allowed to put IPv6 addresses on the Frame-Relay links, we tunnel across the IPV4 domain.

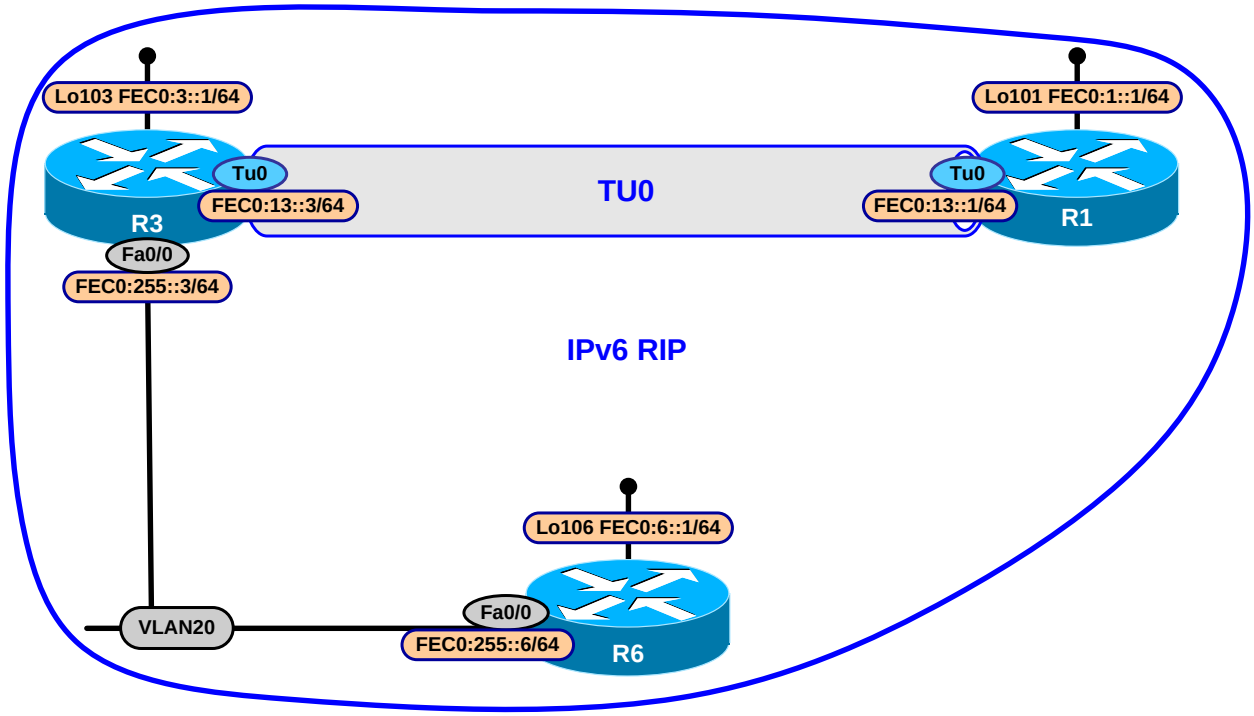

Here we add IPv6 addresses for the path between R1, R6 and R3. The IPv6 loopback on R1 is no longer needed.

We have configured the simplest kind of manual tunnel here. Other types are documented in an excellent text by Regis Desmeules:
Cisco Self-Study: Implementing Cisco IPv6 Networks (IPV6). ISBN: 1-58705-086-2

To implement the IPv6 RIP routing requirement enter the command **IPv6 rip RIP enable** on each of the interfaces with IPv6 addresses.

Verification:

Here is a test ping from R to Loopback 106 on R6:

```
R1#ping fec0:6::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0:6::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 76/78/84 ms
```

And here is the IPv6 rip routing table on R1:

```
R1#sh ipv6 route
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, 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
L   FE80::/10 [0/0]
     via ::, Null0
C   FEC0:1::/64 [0/0]
     via ::, Loopback101
L   FEC0:1::1/128 [0/0]
     via ::, Loopback101
R   FEC0:3::/64 [120/2]
     via FE80::AA12:8603, Tunnel0
R   FEC0:6::/64 [120/3]
     via FE80::AA12:8603, Tunnel0
C   FEC0:13::/64 [0/0]
     via ::, Tunnel0
L   FEC0:13::1/128 [0/0]
     via ::, Tunnel0
R   FEC0:255::/64 [120/2]
     via FE80::AA12:8603, Tunnel0
L   FF00::/8 [0/0]
     via ::, Null0
R1#
```

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".

17.11 QOS

HIDDEN ISSUES TO SPOT WITH THE QOS CONFIGURATION

Issue: Restrict TTCP traffic from CAT2 destined to 170.18.11.1 port 5001 to 1000000 Bit/sec on R1. Allow burst traffic up to 512,000 bytes. Traffic which exceeds the above specified condition shall be dropped.

Solution:

Configure traffic policing on the appropriate inbound R1 FastEthernet subinterface. This can be performed using either the rate-limit command or the modular QOS Command-line interface (MQC). If you use the MQC, you need to perform four steps (1) create an access-list to match the appropriate TCP traffic, (2) associate the access-list with a Class-Map (3) associate the Class-Map to a Policy-Map, (4) enter a police command under the policy-map and (5) apply the policy-map to the R1 appropriate inbound FastEthernet subinterface. If you use the rate-limit command, you need to perform two steps (1) create an access-list to match the appropriate TCP traffic and (2) apply the access-list to an inbound rate-limit command on the appropriate outbound FastEthernet subinterface on R1.

Implementation:

On router R1, create an access-list to match the traffic pattern:

```
access-list 107 permit tcp host 170.18.10.20 host 170.18.11.1 eq 5001
```

Under interface FastEthernet0/0 configure CAR by issuing the command:

```
rate-limit input access-group 107 1000000 512000 512000 conform-action transmit
exceed-action drop
```

The rate-limit command has syntax "rate-limit <rate> <bc> <bc+be>".

Verification:

Issue the **show interface FastEthernet0/0 rate-limit** command:

```
R1#show interface FastEthernet0/0 rate-limit
FastEthernet0/0 VLAN10
  Input
    matches: access-group 107
      params:  1000000 bps, 512000 limit, 512000 extended limit
      conformed 0 packets, 0 bytes; action: transmit
      exceeded 0 packets, 0 bytes; action: drop
      last packet: 198642380ms ago, current burst: 0 bytes
      last cleared 2d07h ago, conformed 0 bps, exceeded 0 bps
R1#
```

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".

17.12 Catalyst Specialties

HIDDEN ISSUES TO SPOT WITH THE CATALYST SPECIALTIES CONFIGURATION

Issue: Create VLAN 200 and assign port fa0/20 to it on CAT2. Do not allow BPDU traffic on this VLAN

Solution:

One way to accomplish this configuration requirement is to disable Spanning Tree on VLAN200 by issuing the command **no spanning-tree vlan 200**.

Issue: Configure R4 as the preferred gateway on CAT1 for networks originating from R1, R2 and R5. Configure R6 as a backup gateway on CAT1. Do not use HSRP, VRRP or IRDP protocols.

Solution:

Configure EIGRP on CAT1. It will prefer R4 for networks originating from R1, R2 and R5 since prefixes from these routers will possess a lower EIGRP composite metric when they are advertised from R4. If the routes are not available this way, CAT1 will then learn them from R6 via EIGRP. Remember to activate the ip routing process before you enable EIGRP on CAT1.

Issue: Configure CAT1 so that it can be managed by a network management service that uses UDP port 161. Set read-only access with the string of RS-NMC. Set read-write access with the string NMC.

Solution:

Configure SNMP on CAT1 using the READ-ONLY and READ-WRITE community strings specified above. SNMP uses UDP port 161.

```
snmp-server community RS-NMC RO
snmp-server community NMC RW
```

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".

17.13 Gateway Redundancy

HIDDEN ISSUES TO SPOT WITH THE GATEWAY REDUNDANCY CONFIGURATION

Issue: Configure HSRP between R4 and R6. Make R4 the preferred gateway. Switch over to R6 if the frame relay connection on R4 goes down.

Solution:

To fulfill this requirement, configure HSRP between R4 and R6 using the virtual gateway IP address of 172.10.64.1. Set R4 with a higher standby-group priority than R6. Since you are told to switch the preferred gateway from R4 to R6 when the Frame-Relay connection of R4 goes down, configure the "standby..track" option on R4 so that R4 will decrease its standby priority so that R6 will become more preferred. Here are the commands to configure HSRP on R4 and R6:

Router R4:

```
interface FastEthernet0/0
 description VLAN30
 ip address 170.18.64.4 255.255.255.0
 no ip redirects
 duplex auto
 speed auto
 standby ip 170.18.64.1
 standby priority 110
 standby preempt
 standby track Serial0/0 20
```

Router R6:

```
interface FastEthernet0/0.30
 encapsulation isl 30
 ip address 170.18.64.6 255.255.255.0
 no ip redirects
 ip nat outside
 standby ip 170.18.64.1
 standby preempt
```

Verification:

Issue the **show standby** command:

```
R6#show standby
FastEthernet0/0.30 - Group 0
  State is Standby
    12 state changes, last state change 00:32:33
  Virtual IP address is 170.18.64.1
  Active virtual MAC address is 0000.0c07.ac00
    Local virtual MAC address is 0000.0c07.ac00 (default)
  Hello time 3 sec, hold time 10 sec
```

```
  Next hello sent in 0.372 secs
Preemption enabled
Active router is 170.18.64.4, priority 110 (expires in 9.908 sec)
Standby router is local
Priority 100 (default 100)
IP redundancy name is "hsrp-Fa0/0.30-0" (default)
```

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".

17.14 Multicast

HIDDEN ISSUES TO SPOT WITH THE MULTICAST CONFIGURATION

Issue: Enable multicast routing between routers R1, R2, and R3. Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree.

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, Dense Mode or Sparse-Dense Mode. The answer lies in the last sentence displayed above: "based on a shared tree". It is PIM Sparse Mode that is based upon a "shared tree." At the root of the Shared Tree is a given multicast-group's Rendezvous Point. Therefore, you need to configure PIM Sparse-Dense Mode to fulfill the configuration requirement. Use the **ip pim sparse-dense-mode** command on each required interface. Remember to enable multicast routing by issuing the command **ip multicast-routing** in global configuration mode before configuring PIM or IGMP.

Verification:

You can use the command **show ip pim neighbors** to verify your PIM configuration. See the diagram on the following page.

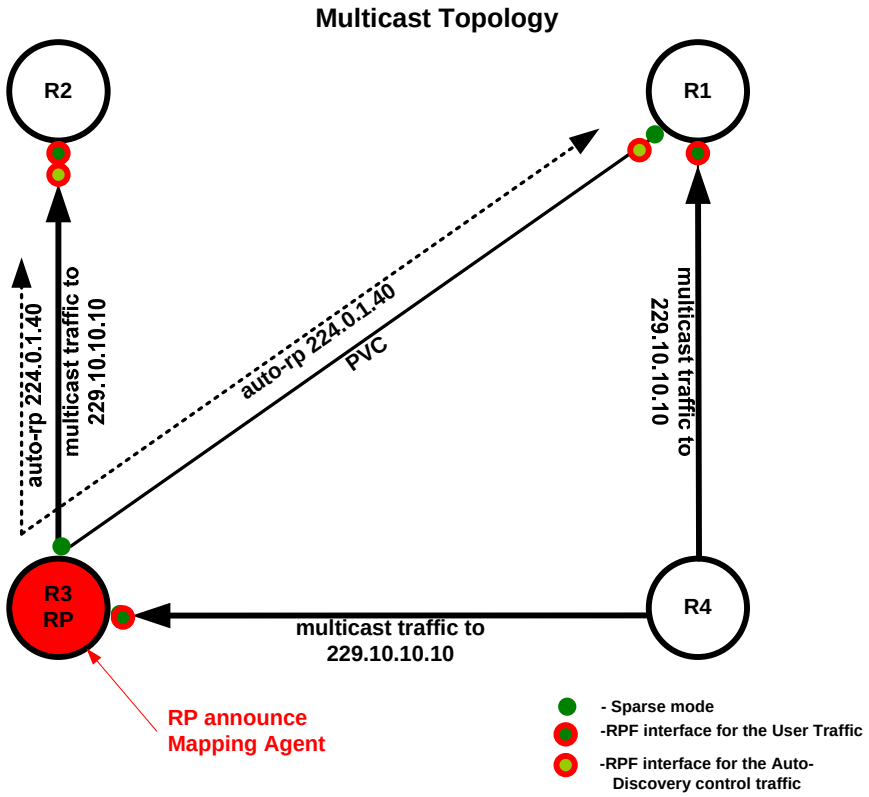

Issue: Configure all of the above listed routers to join the multicast group 229.10.10.10. Associate this multicast group with a loopback interface on each router.

Solution:

Configure the **ip igmp join-group X.X.X.X** command under a loopback interface on the listed routers. Also, remember to configure **ip pim sparse-dense mode** on the same loopbacks to make them respond to pings.

Issue: Use the 224.0.1.39 PIM dense group for this configuration. Make R3 the root of the shared tree. Accomplish this task by configuring only R3.

Solution:

Since you are told to use 224.0.1.39, you must configure Auto-RP. Auto-RP uses the reserved multicast groups 224.0.1.39 and 224.0.1.40. When you configure Auto-RP, you must not only configure a candidate RP, you must also configure a Mapping Agent. Both the RP Announcer and the Mapping Agent can be configured on the same router or different routers. Use the following command on R3:

```
ip pim send-rp-announce Serial0/0 scope 2
ip pim send-rp-discovery Serial0/0 scope 2
```

Verification:

To verify that all routers have learned the RP address, use the command **show ip pim rp [mapping].**

```
R1#sh ip pim rp mapping
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 170.18.123.3 (?), v2v1
    Info source: 170.18.123.3 (?), elected via Auto-RP
         Uptime: 02:31:00, expires: 00:02:47
```

Issue: Ping the multicast group 229.10.10.10 from R4 to all other multicast routers.

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 R3. Once R1 receives the multicast packets from R4, it will forward the packets down to router R3.

This is the central issue related to this task: R3 will receive multicast traffic for the 229.10.10.10 group on both its physical interface from R1 and on its point-to-point subinterface from R4. R3 should not accept the traffic coming from R1, because, the incoming interface for a given multicast traffic stream cannot also be an outgoing interface. R3 has the same interface connected to R2 and R1, therefore R1's multicast traffic will not be forwarded back out the same interface to R2.

You have to make sure that R3 receives the traffic directly from R4 on one interface and then forwards traffic to R2 via a different interface. A method to consider using to make sure that R3 receives the multicast traffic directly from R4 is to configure an "mroute" entry. Please refer to the final configuration scripts to determine the solution used for this Scenario.

A second potential issue to be aware for this task is: since the 170.18.123.0/24 subnet is configured as an OSPF "broadcast" network type, it is possible that R2 will have R1 listed as the next-hop for the 170.18.134.0/24 subnet. Remember the following rule: if a prefix is advertised from one spoke to another spoke on a partial mesh NBMA topology and the partial mesh is configured as either an OSPF "broadcast" or "non-broadcast" network type, the next-hop WILL NOT be changed at the hub router. Therefore, the next-hop for a prefix advertised from one spoke to another spoke will remain the advertising spoke router. This creates a problem for an RPF lookup. For an RPF lookup to be successful, the next-hop IP address of a routing entry selected by the RPF process must match a neighboring PIM router. If the next-hop is a spoke router, it will not be a neighboring PIM router. To assure that the RPF process on router R2 does not use a routing table entry for the 170.18.134.0/24 subnet (this is the where R4, the source of the multicast traffic, resides) that lists R1 as the next-hop, configure an mroute on R2 explicitly designating R3 as the upstream next-hop. See the final configuration for more details.

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".