NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II

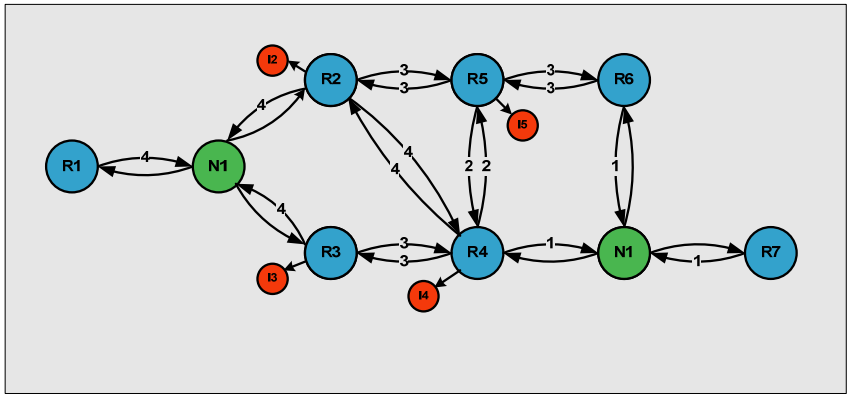

Scenario 3
ANSWER KEY

FOR

CCIE® CANDIDATES

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171
The material in this lab workbook remains the intellectual property of NetMasterClass., LLC

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.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

DOiT-V6 Scenario 3: 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.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

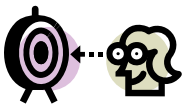

Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16.
- Do not use any static routes.
- Advertise Loopback interfaces involved in IGP configurations with their original mask. Do not change the mask.
- Do not use 0.0.0.0 anywhere in this scenario.
- Do not use "ip default-network".
- All IP addresses involved in this scenario must be reachable, unless specified otherwise.
- Networks advertised in the BGP section must be reachable only in the BGP domain.
- Use conventional routing algorithms.

Explanation of Each of the Goals and Restrictions

IP subnets on the diagram belong to network 172.16.0.0/16 unless specified otherwise.

All IP addresses in this Exam belong to the 172.16.0.0/16 address space with the exception of a set of prefixes used in the BGP section.

Do not use any static routes.

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

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

An exception is made for interfaces advertised by BGP.

Do not use 0.0.0.0 anywhere in this scenario.

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

Do not use "ip default-network".

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

This is a key goal to observe. This requires that all of your IGP's are configured properly. Also, all of your routing policy tasks must be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using distribute-lists, route-maps and the distance command. A key point to remember about this exam is: the term "redistribution" is never explicitly used in

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

this exam. However, you must perform redistribution in order to assure that all ip addresses are reachable without the use of static routes or 0.0.0.0/0 routes.

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

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

Use conventional routing algorithms.

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

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

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

The following IOS versions were used on the devices:

Device	IOS version
R1	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
R2	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
R3	IOS (tm) C2600 Software (C2600-J1S3-M), Version 12.3(15a)
R4	IOS (tm) C2600 Software (C2600-J1S3-M), Version 12.3(15a)
R5	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
R6	IOS (tm) 3600 Software (C3640-JK9O3S-M), Version 12.4(3)
FRS	IOS (tm) 2500 Software (C2500-JS-L), Version 12.2(27)
CAT1	IOS (tm) C3550 Software (C3550-I5Q3L2-M), Version 12.2(25)SEA
CAT2	IOS (tm) C3550 Software (C3550-I5Q3L2-M), Version 12.2(25)SEA

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.1 Frame-Relay

HIDDEN ISSUES TO SPOT WITH THE FRAME-RELAY CONFIGURATION

Issue: Configure Frame-Relay connections between routers R1, R3, R4 and R5.

Solution:

The challenging part of this Frame-Relay configuration is coming up with the proper subinterface configuration on R1, in accord with the OSPF requirement that R2 and R4 have different next-hop addresses on R1. The solution involves configuring addresses from the same subnet on two subinterfaces. All other addresses on Frame-Relay networks in this lab go on physical interfaces.

Frame-Relay Interfaces and Addresses on R1

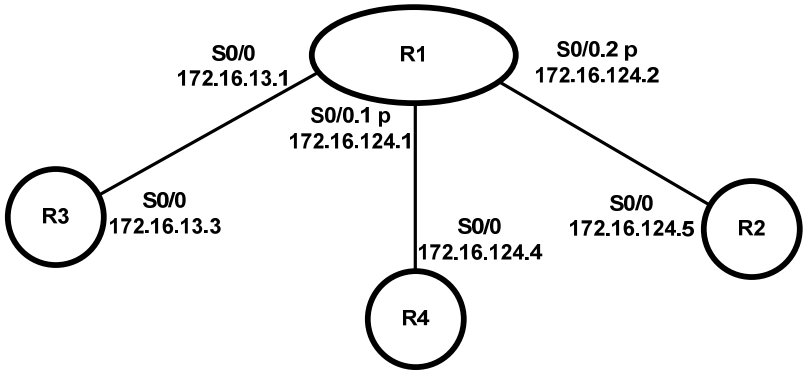

As shown above, you can create two point-to-point subinterfaces on router R1's Frame-Relay interface and assign the IP addresses 172.16.124.1/24 and 172.16.124.2/24 to each. When you attempt to do this on two separate physical interfaces, you encounter the error message: "172.16.124.0 overlaps with Serial0/0". The IOS does allow you assign the same network prefix to two different subinterfaces under the same physical interface.

Since the Frame Relay switch is pre-configured for a full mesh of PVCs between routers R1, R2, R3 and R4, and you are instructed to use only the PVCs in the diagram, you must disable inverse arp and do static Frame-Relay map statements wherever necessary. The OSPF network type will be point-to-multipoint on the 172.16.124.0 subnet, so only Frame-Relay maps to the hub are necessary. See the OSPF section for further discussion of this point.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Issue: Configure a back-to-back Frame-Relay connection between routers R2 and R5.

Solution:

The Frame-Relay connection between routers R2 and R5 is a back-to-back Frame-Relay connection with no Frame-Relay switch in between. There are two basic methods of configuring back-to-back Frame-Relay:

1. A basic method that relies on just map statements and does not use LMI, discussed at this link: http://www.cisco.com/warp/public/125/frbacktoback.html
2. One called hybrid switching, used in this lab, which can be found at this link: http://www.cisco.com/warp/public/125/frbacktoback_hybrid.htm

In the hybrid method, one router acts as a Frame-Relay switch (DCE) and the other router acts as the Frame-Relay DTE device. On the router that is configured as the Frame-Relay switch, R5 in our case, we enter the global configuration command **frame-relay switching** and the interface commands **encap Frame-Relay** and **frame-relay intf-type dce**.

When configuring a back-to-back Frame-Relay configuration you will not use any **frame-relay route** commands, because you are not switching Frame-Relay traffic from one Frame-Relay switched interface to another. Instead, use the **frame-relay interface-dlci XXX** command to advertise a DLCI from the router acting as the Frame-Relay switch to the router acting as the Frame-Relay DTE device. To assure that inverse-arp requirements are satisfied, enter a frame-relay map command on both routers referencing the DLCI announced by the router acting as the Frame-Relay switch.

This Scenario requires that router R2 use the DLCI 205 and R5 use DLCI 502. This can be accomplished using the hybrid technique described above, with one exception. The router acting as the Frame-Relay DCE needs to configure a frame-relay map statement referencing the DLCI that was assigned to it in the diagram. It will apply this DLCI to any packets it sends to the other end of the Frame-Relay connection. The remote side will de-encapsulate the IP packet and ignore the unrecognized DLCI number. On R5, we create DLCI 205 for the DTE side and map 502.

In a back-to-back serial connection, you must also set the clock rate on the router connected to the DCE side of the serial cable. The DCE side can be determined by issuing **show controllers serial N** where N is the serial interface number. Remember the space between the key word serial and the serial interface number. In the **show controllers** display, you will see the cable type and which interface is the DCE and DTE connection. Here is partial output for our R5:

```
R5#show controllers serial 1/0
CD2430 Slot 1, Port 0, Controller 0, Channel 0, Revision 18
Channel mode is synchronous serial
idb 0x63581818, buffer size 1524, V.35 DCE cable, clockrate 72000
```

Verification:

At this stage you should be able to ping across the Frame-Relay links R1 to R3 and R2 to R5. You may not be able to ping reliably. You will probably not be able to ping between R1 to R2, R1 to R4, or R2 to R4 until the OSPF section is complete.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Below, you see that R1 has three Active, Local PVCs, one each to R2, R3 and R4. There are two static mappings, one to our own address and one to 172.16.13.3.

```
R1#sh frame pvc

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

                Active      Inactive      Deleted       Static
  Local            3             0            0            0
  Switched         0             0            0            0
  Unused           0             0            0            0

R1#sh frame map
Serial0/0 (up): ip 172.16.13.1 dlci 103(0x67,0x1870), static,
              IETF, status defined, active
Serial0/0 (up): ip 172.16.13.3 dlci 103(0x67,0x1870), static,
              broadcast,
              CISCO, status defined, active
              RTP Header Compression (enabled), connections: 256
Serial0/0.1 (up): point-to-point dlci, dlci 104(0x68,0x1880), broadcast
          status defined, active
Serial0/0.2 (up): point-to-point dlci, dlci 102(0x66,0x1860), broadcast
          status defined, active
```

Mappings are done automatically for point-to-point interfaces: it is assumed there is only one device on the subnet reachable out the interface. The configuration on R1 is unusual in that there are two point-to-point subinterfaces in the same subnet. If you turn on **debug IP packet** and try to ping 172.16.124, for example, you may see that only every other ping gets a response. As you see below, the software believes it has two interfaces it can use to reach 172.16.124.4, and it load balances between them! We will fix this later by using the point-to-multipoint OSPF network type on this link.

```
R1#ping 172.16.124.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.124.4, timeout is 2 seconds:
!
1d21h: IP: s=172.16.124.1 (local), d=172.16.124.4 (Serial0/0.1), len 100, sending
1d21h: IP: s=172.16.124.2 (local), d=172.16.124.4 (Serial0/0.2), len 100, sending.
1d21h: IP: s=172.16.124.2 (local), d=172.16.124.4 (Serial0/0.2), len 100, sending.!
1d21h: IP: s=172.16.124.1 (local), d=172.16.124.4 (Serial0/0.1), len 100, sending
1d21h: IP: s=172.16.124.2 (local), d=172.16.124.4 (Serial0/0.2), len 100, sending.
Success rate is 40 percent (2/5), round-trip min/avg/max = 56/56/56 ms
```

When you enter **show frame pvc** on R3 and R4 you should see 1 Local and 2 Unused on each, with static mappings to R1. On R2 you should see 1 Local and 2 Unused associated with the hub-and-spoke link. The back-to-back link will show 1 Local and 1 Unused. The unused PVC is DLCI 502, on which it is receiving traffic from 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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.2 Catalyst Configuration

HIDDEN ISSUES TO SPOT WITH THE CATALYST 3550 CONFIGURATION

Basic Catalyst 3550 configuration includes setting the VTP mode, configuring trunk ports, and statically assigning ports to VLAN's. For a good reference on basic Catalyst 3550 configuration tasks, download the following Tech-Note from the Technical Library on the NetMasterClass web-site: "Performing Basic Configuration tasks on the Catalyst 3550"

Issue: Configure VLANs according to the VLAN Configuration Table and Diagram

Solution:

The VLAN configuration here is fairly simple, but it might be helpful to get in the habit of creating VLAN tables like the following, especially if you have to manually prune your VLANs.

VLANs Configured on Cat1

VLAN 10	VLAN 20	VLAN 30	VLAN 60
F0/8	Not needed	F0/4	F0/8

VLANs Configured on Cat2

VLAN 10	VLAN 20	VLAN 30	VLAN 60
F0/3	F0/1, F0/7	F0/5	F0/1

Issues: On the trunks, allow only the VLANs used in this scenario. Use the dot1q trunking protocol.

Solution:

From the tables above, it is apparent that only traffic for VLANs 10, 30 and 60 needs to cross the trunks between Cat 1 and Cat2. All of the ports assigned to VLAN 20 are on Cat2 only. Cat1 port F0/8 supports VLANs 10 and 60. Cat2 port F0/1 requires VLANs 20 and 60. The process of statically configuring the allowed VLANs on trunks is sometimes referred to as manual pruning (as opposed to VTP pruning). The following command sequence, issued on both Catalysts will change the trunk encapsulation type to dot1q from the default ISL and permit traffic from only the required VLANs:

```
(config)# interface range f0/13 –14
(config-range)# switchport encapsulation dot1q
(config-range)# switchport trunk allowed VLAN 10,30,60
```

When you list multiple VLANs with the **switchport trunk allowed vlan** command, separate your listed VLANs with commas and hyphens (hyphens for ranges of VLAN's). Do not include any spaces.

Verification:

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

In the output below we can see that ports f0/13 and f0/14 are trunking with encapsulation dot1q. Only VLANs 10 and 30 are allowed across the trunks. The last line indicates that Fa0/14 is in spanning-tree blocking mode for all permitted VLANs.

```
CAT2#sh int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/1       on           802.1q         trunking      1
Fa0/13      on           802.1q         trunking      1
Fa0/14      on           802.1q         trunking      1

Port      Vlans allowed on trunk
Fa0/1       20,60
Fa0/13      10,30,60
Fa0/14      10,30,60

Port        Vlans allowed and active in management domain
Fa0/1       20,60
Fa0/13      10,30,60
Fa0/14      10,30,60

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1       20,60
Fa0/13      10,30,60
Fa0/14      none
```

Issue: Create a Layer 3 port on CAT1

Solution:

The diagram shows an Ethernet connection between R2 and CAT1 using network 172.16.25.32/27. Since no VLAN is shown, the IP address on the Catalyst must be assigned directly to the port. You turn the Catalyst port into a Layer 3, routed port by issuing the command **no switchport**. You can then configure an IP address on the port.

Verification:

The commands below are a couple of my favorites.

```
CAT1#show interfaces status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/1                        notconnect   1            auto   auto 10/100BaseTX
Fa0/2     R2  E0             connected    routed     a-full  a-100 10/100BaseTX
Fa0/3                        notconnect   1            auto   auto 10/100BaseTX
Fa0/4     R3  Fa0/1          connected    30         a-full  a-100 10/100BaseTX
…

CAT1#show ip interface brief
Interface                  IP-Address      OK? Method Status                Protocol
Vlan1                      unassigned      YES manual administratively down down
FastEthernet0/1            unassigned      YES unset  down                  down
FastEthernet0/2            172.16.25.34    YES manual up                    up
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Issue: Routing on CAT1 should be configured as appropriate for other tasks.

Solution:

By default the IP routing process is disabled on a Catalyst 3550, which is to say that the default routing configuration on a Catalyst 3550 is **no ip routing**. Later, we will run OSPF on CAT1, and IP routing will have to be enabled.

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".

3.3 OSPF

HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

Issue: Use the most suitable OSPF network type to accomplish the previous two tasks.

Solution:

The two previous tasks specified two separate requirements: (1) assign the 172.16.124.0/24 subnet to OSPF area 0 and (2) R2 and R4 must possess unique next-hop addresses for any traffic forwarded to R1. The second requirement was discussed in the Frame-Relay issue spotting section earlier in this answer key. It required that two separate subinterfaces be configured on router R1 using the same subnet prefix 172.16.124.0/24. The most suitable OSPF network type for this configuration is the point-to-multipoint OSPF network type, because it will generate host routes for each configured point-to-multipoint subinterface. Any other OSPF network-type will generate a prefix for the 172.16.124.0/24 subnet. This may cause confusion because there are two 172.16.124.0/24 subnets configured. The OSPF point-to-multipoint network type will avoid generating any such confusion and this network type will fulfill the "unique next-hop" requirement.

Verification:

Here is some of the routing detail from R2. Note that 172.16.124.4 is shown as a host route. Since the next hop to the other spoke is always the hub, no Frame-Relay maps between spokes are required.

```
R2#sh ip route 172.16.124.4
Routing entry for 172.16.124.4/32
  Known via "ospf 1", distance 110, metric 128, type intra area
  Last update from 172.16.124.2 on Serial0/0, 1d05h ago
  Routing Descriptor Blocks:
  * 172.16.124.2, from 172.16.104.1, 1d05h ago, via Serial0/0
      Route metric is 128, traffic share count is 1
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Recall from the Frame-Relay section that half of our pings failed when we tried to ping a spoke from R1. The router was load balancing across the two subinterfaces in the 172.16.124.0/24 network. With the host routes provided by the OSPF point-to-multipoint network type the confusion is gone. In order to sort out the confusion at layer 2, we had to move up to layer 3.

```
R1#sh ip route | include /32
O       172.16.124.4/32 [110/64] via 172.16.124.4, 00:12:25, Serial0/0.1
O       172.16.124.5/32 [110/64] via 172.16.124.5, 00:12:25, Serial0/0.2
```

Issue: R1 and R3 should automatically discover each other; Elect R1 as the DR. Make sure R1 is still the DR after you reboot.

Solution:

These tasks need to be read together, as they provide direction on how to configure the 172.16.13.0/24 subnet. The only OSPF network type that both allows R1 and R3 to automatically discover each other and also elects a DR is the "broadcast" OSPF network type. To assure that R1 is always the DR configure R3 with OSPF priority 0. That way it will not be elected a DR or BDR, even if R1 reloads.

Verification:

```
R1#sh ip ospf interface s0/0
Serial0/0 is up, line protocol is up
  Internet Address 172.16.13.1/24, Area 0
  Process ID 1, Router ID 172.16.101.1, Network Type BROADCAST, Cost: 64
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.101.1, Interface address 172.16.13.1
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 5, maximum is 5
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.103.1
  Suppress hello for 0 neighbor(s)
```

Issue: Configure Area 25 between R2 and R5. Use a network type with a default Hello time of 10 seconds. Do not elect a DR/BDR.

Solution:

Here are the network types available on serial links and their default Hello times:

Type	Hello	DR/BDR
Broadcast	10 seconds	Elects DR/BDR
Nonbroadcast	30 seconds	Elects DR/BDR
Point-to-multipoint	30 seconds	No DR/BDR

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

P-M non-broadcast	30 seconds	No DR/BDR
Point-to-point	10 seconds	No DR/BDR

Both Broadcast and Point-to-point network types have default Hello times of 10 seconds, but the Broadcast type elects a DR/BDR. So only the Point-to-point network type will suffice here.

Issue: Make sure network 172.16.108.0/24 is in the Area 25 OSPF topology table, but not in R5's routing table

Solution:

One way to keep prefixes that are in the Area OSPF topology table from going into the routing table is to use a distribute-list under the OSPF process, **distribute-list prefix 1 in**. As of 12.2(15)T you can also reference a route-map matching on any feature of the route. For more information see this link:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s24/routmap.htm

Verification:

We can verify that the inter-area route 172.16.108.0/24 is in the R5 topology table as follows:

```
R5#sh ip ospf database

            OSPF Router with ID (172.16.105.1) (Process ID 1)

                Router Link States (Area 25)

Link ID         ADV Router      Age         Seq#       Checksum Link count
172.16.102.1    172.16.102.1    1283        0x8000002E 0x0003D5 2
172.16.105.1    172.16.105.1    1283        0x8000002C 0x005E40 3

                Summary Net Link States (Area 25)

Link ID         ADV Router      Age         Seq#       Checksum
172.16.13.0     172.16.102.1    785         0x80000027 0x00C7E0
172.16.25.32    172.16.102.1    785         0x8000002A 0x0046D1
172.16.102.0    172.16.102.1    785         0x80000028 0x00F3D9
172.16.104.0    172.16.102.1    785         0x80000025 0x00E865
172.16.108.1    172.16.102.1    785         0x80000028 0x00B114
172.16.124.1    172.16.102.1    785         0x80000025 0x007505
172.16.124.2    172.16.102.1    785         0x80000025 0x006B0E
172.16.124.4    172.16.102.1    785         0x80000025 0x00D95D
172.16.124.5    172.16.102.1    785         0x80000028 0x00C4EE
```

Yet, it is not in R5's routing table:

```
R5#sh ip route 172.16.108.0
% Subnet not in table
```

Issue: Configure Area 28 between R2 and CAT1. Make sure that CAT1 receives all of the routes except for network 172.16.105.0/24 and its subnets. These prefixes should not be in CAT1's routing table or its OSPF database.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Solution:

As of 12.2(4)T you can use ABR Type 3 LSA Filtering to control the flooding of Summary LSAs into or out of an area. We entered the command **area 28 filter-list prefix 1 in** under the OSPF process on R5. The prefix-list denies 172.16.105.0/24 and all of its subnets, and then permits all others.

```
ip prefix-list 1 seq 5 deny 172.16.105.0/24 le 32
ip prefix-list 1 seq 10 permit 0.0.0.0/0 le 32
```

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".

3.4 RIP

HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

Issue: Configure RIP v.2 between R3, R4 and CAT2. Allow updates on VLAN 30 only.

Solution:

Configure **passive-interface default** under the RIP process of each specified router. Then configure **no passive-interface** for the Ethernet interfaces of each router attached to VLAN 30. This is a good general practice, even if not required, as it makes any debugs a lot easier to decipher.

Issue: Multiple Redistribution Points between OSPF and RIP v.2.

Two redistribution points exist between RIP v.2 and OSPF on routers R3 and R4. All RIP routes originating from CAT2 will be assigned the administrative distance of 120 when routers R3 and R4 receive them. These same routes will get redistributed into OSPF at both routers R3 and R4. When they are flooded through the OSPF domain, they will make it back to R3 and R4 as external OSPF routes. Since OSPF has a lower administrative distance than RIP, routers R3 and R4 will select an OSPF path for a RIP path even though they are directly connected to the RIP domain. This is a complex hidden issue. It is well described and documented in the following Tech-Note in the NetMasterClass web-site technical library: "A Scenario with Multiple Redistribution Points".

Solution:

One solution, documented in the Tech-Note"A Scenario with Multiple Redistribution Points" in the NetMasterClass web-site technical library, is to use the **distance** command under the RIP processes on R3 and R4 to lower the administrative distances for internal RIP prefixes. In essence, this makes RIP act lie EIGRP, with a low admin distance for internal routes, and a high (120) admin distance for external prefixes.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

The solution implemented here is to make OSPF look like EIGRP by using the command **distance ospf external 180.** We can verify it easily with the commands below. The general principle is that native routes should be routed by the native protocol. Here is the result on R3:

```
R3#sh ip route rip
     172.16.0.0/16 is variably subnetted, 21 subnets, 4 masks
R       172.16.124.0/24 [120/1] via 172.16.34.4, 00:00:01, FastEthernet0/1
R       172.16.120.0/24 [120/1] via 172.16.34.10, 00:00:03, FastEthernet0/1
     10.0.0.0/24 is subnetted, 1 subnets
R       10.10.10.0 [120/1] via 172.16.34.4, 00:00:01, FastEthernet0/1
```

And here is the result on R4:

```
R4#sh ip route rip
     70.0.0.0/24 is subnetted, 1 subnets
R       70.70.70.0 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
     172.16.0.0/16 is variably subnetted, 20 subnets, 4 masks
R       172.16.36.0/22 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
R       172.16.16.0/24 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
R       172.16.17.0/24 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
R       172.16.120.0/24 [120/1] via 172.16.34.10, 00:00:01, FastEthernet0/0
R       172.16.106.0/24 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
R       172.16.107.0/24 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
R       172.16.101.0/24 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
R       172.16.103.0/24 [120/1] via 172.16.34.3, 00:00:22, FastEthernet0/0
```

In the output above we see routes from the EIGRP domain, like 172.16.107.0, as RIP routes. It makes sense in terms of optimal paths to go across three 100 Mbps links rather than one low speed serial link. Note below that the prefixes native to the OSPF domain are seen as OSPF routes on R4, and none of the RIP domain routes are seen as OSPF routes.

```
R4#sh ip route ospf
     172.16.0.0/16 is variably subnetted, 20 subnets, 4 masks
O IA    172.16.25.32/27 [110/129] via 172.16.124.1, 00:16:09, Serial0/0
O IA    172.16.25.0/27 [110/909] via 172.16.124.1, 00:16:09, Serial0/0
O       172.16.13.0/24 [110/128] via 172.16.124.1, 00:16:09, Serial0/0
O       172.16.124.1/32 [110/64] via 172.16.124.1, 00:16:09, Serial0/0
O       172.16.124.2/32 [110/64] via 172.16.124.1, 00:16:09, Serial0/0
O       172.16.124.5/32 [110/128] via 172.16.124.1, 00:16:09, Serial0/0
O IA    172.16.108.0/24 [110/130] via 172.16.124.1, 00:16:09, Serial0/0
O IA    172.16.105.0/24 [110/910] via 172.16.124.1, 00:16:09, Serial0/0
O       172.16.102.0/24 [110/129] via 172.16.124.1, 00:16:09, Serial0/0
```

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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.5 EIGRP

HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

Issue: Do not form an EIGRP adjacency on the 172.16.13.0/24 subnet.

Solution:

Use the EIGRP network/mask router configuration command to limit the number of interfaces that are participating in EIGRP on routers R1 and R3. If you configure EIGRP with the classful network command – **network 172.16.0.0** – all interfaces assigned with 172.16.0.0 will be included within EIGRP, including the 172.16.13.0/24 subnet, which is a violation of task requirements. As a general practice, use the mask option when configuring a network command under the EIGRP process to gain greater control over what interfaces are participating in EIGRP.

Issue: Authenticate all EIGRP adjacencies.

Solution:

Configure a key-chain in global configuration mode, then configure EIGRP authentication on the appropriate interfaces on FRS, R1, R6 and R3.

Issue: When redistributing OSPF into EIGRP on R1, permit only routes with OSPF metrics between 0 and 1000, inclusive.

Solution:

This task is best satisfied using the EIGRP Route-Map Support feature, new in 12.3(8)T. We simply added the command **match metric 500 +-500** to the permit entries in route-map OSPF-EIGRP.

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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION

The lab presents a fairly complex redistribution challenge. There are multiple physical loops that cross routing domains, and there are multiple potential transit domains. In addition, methods of sending default routes have been largely ruled out, and subsequent tasks appear to require specific routing paths. There are many valid approaches to this problem. You are directed to mutually redistribute between the IGPs on routers R1, R3 and R4. In addition, you are directed to make all redundant paths available. Since physical loops cross all three routing domains, all three will be transit (they will advertise non-native routes), and we could run into trouble with route feedback.

BGP can provide us a model for how to implement loop-free paths with multiple transit domains. In essence, BGP tags each advertised prefix with its own AS number. BGP avoids route-feedback by dropping updates that have its own AS number in the AS path attribute.

On each redistribution, we apply a route-map to (1) deny routes that were sourced from the destination protocol, (2) permit routes that are neither from the source protocol nor the target protocol, and (3) tag the remaining routes, which will be from the source protocol. Here is an example from R3:

```
route-map EIGRP-OSPF deny 10
 match tag 110            <denies ospf routes back into ospf>
!
route-map EIGRP-OSPF permit 20
 match tag 120            <permits rip routes>
!
route-map EIGRP-OSPF permit 30
 set tag 90         <sets tag for eigrp (remaining) routes>
```

Verification:

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 through 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, and paste the script below. When finished, be sure to issue the command **tclq** to kill the process. CTRL-Z will exit the TCL shell interface, but will leave the process running in the background.

```
foreach address {
172.16.101.1
172.16.17.1
172.16.13.1
172.16.124.1
172.16.102.1
172.16.124.2
172.16.16.1
172.16.25.33
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
172.16.124.5
172.16.25.2
172.16.103.1
172.16.36.3
172.16.13.3
172.16.34.3
172.16.104.1
172.16.34.4
172.16.124.4
172.16.105.1
172.16.25.5
172.16.106.1
172.16.36.6
172.16.16.6
70.70.70.70
172.16.107.1
172.16.17.7
172.16.25.34
172.16.108.1
172.16.120.1
10.10.10.10
172.16.34.10
} {ping $address}
```

We also need to make sure that our solution is a stable one. If we have split-horizon or other route feedback problems routes may continually be inserted and removed from our routing tables. We can test stability by observing the output of **debug IP routing**. Finally, we need to make sure that our routes are optimal: that native prefixes are routed by native protocols, that we are using the shortest paths, and that redundant paths are available. This requires close examination of each routing table. To test redundancy shut various key interfaces and make sure you still have full reachability.

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".

3.6 BGP

HIDDEN ISSUES TO SPOT WITH THE BGP CONFIGURATION

They say a picture is worth a thousand words, so you might want to focus on the diagram of the required BGP topology shown on the next page. We create four autonomous systems. AS 700 contains just FRS. It is peered with R1 in AS 100 and sources the prefix 70.70.70.0/24. AS100 contains R6, R1 and R2, with R1 as a route reflector. AS 100 peers with both of the routers in AS 300, R3 and R4. Cat2 is alone in AS1000, sources the prefix 10.10.10.0/24 and peers with both of the routers in AS300. The basic challenge is to engineer the preferred path CAT2>R4>R3>R6>R1>FRS by using local preference. In addition, we need to meet the synchronization requirement in AS 300.

Issue: Basic Internal and External BGP Peering:

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Solution:

The issues to spot here are the route-reflector requirement in AS100 and the multihop requirement on the peering between R4 and R2. Without **ebgp-multihop**, EBGP packets have a TTL of 1 and will not be forwarded from one spoke to another in a hub and spoke topology, because the hub must decrement the TTL. Therefore, both R2 and R4 need to be configured with **ebgp-multihop**.

Verification:

Issue the command **show IP bgp summary** on each router. Partial output for R4 is shown below. It shows three good neighbors with zero prefixes learned as yet. An "active" state is misconfigured.

```
R4#sh ip bgp summary
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.34.3     4   300      51      51       40    0    0 00:41:20        0
172.16.34.10    4  1000      45      44       38    0    0 00:35:33        0
172.16.124.5    4   100      50      51       38    0    0 00:41:03        0
```

BGP Peerings and Required Traffic Path

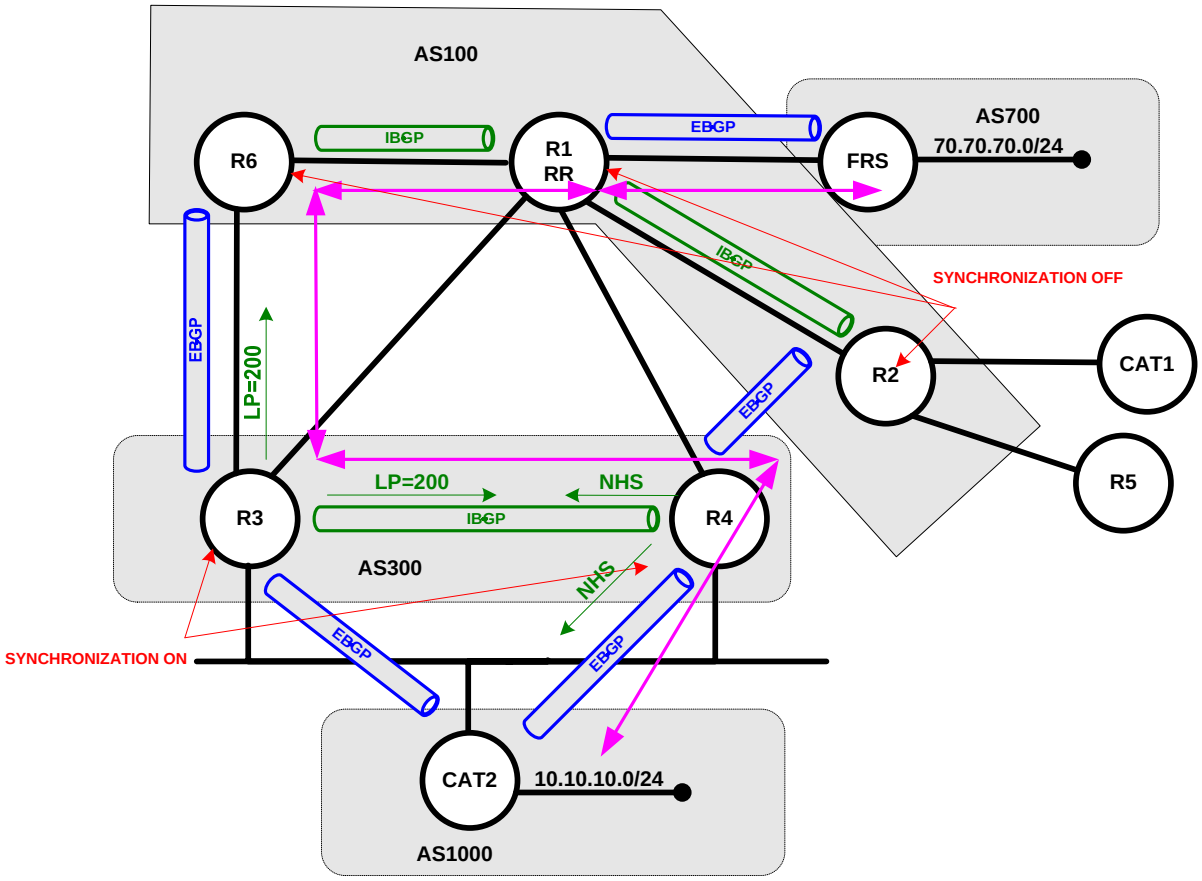

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Issue: Advertise the 10.10.10.0/24 network as IGP and 70.70.70.0/24 network as incomplete.

Solution:

The 10.10.10.0/24 prefix will have an origin code of "IGP" if we advertise into BGP with a **network** statement. Remember that you must use the **mask** keyword if advertising other than a classful network. To advertise the 70.70.70.0/24 prefix with an origin code of "incomplete" we must redistribute it into BGP. We use a route-map to specify the single connected prefix required.

Issue: Use the synchronization method in AS300.

The rule of synchronization says that a router will not advertise externally or use an IBGP-learned route, unless that route is in the forwarding table from a source other than BGP. This feature is intended to make sure that you do not advertise a route that you cannot in fact reach. It is based on a model that assumes that not all routers in the AS will run BGP and that you will redistribute BGP into your IGP. In practice, this model is rare, and synchronization is turned off in most production networks. But you should be prepared to meet this requirement in the CCIE lab. Note that synchronization was on by default until 12.2(8)T, when the default became **no synchronization**. You can verify whether it is enabled or disabled with **show IP protocols**. In this exercise, you need only synchronize routes that are used in the preferred path.
Part 1: Synchronize network 70.70.70.0/24 on R4 and advertise to CAT2.

As depicted on the diagram above, the preferred path requires R4 to advertise the network 70.70.70.0/24, learned from R3, to CAT2. So this IBGP-learned route must be in the routing table of R4 from a source other than BGP. We accomplish this by redistributing BGP into RIP on R3 so that RIP will advertise network 70.70.70.0/24 to R4. The prefix will still not be in the R4's routing table from RIP, because the route is also learned via external BGP from R2, and EBGP has an administrative distance of just 20. So we have raised the administrative distance to 200 for all bgp-learned routes. An alternative here would have been the **backdoor** command. This will allow the RIP-learned route for prefix 70.70.70.0/24 to be put into the forwarding table, satisfying the synchronization requirement.

At this point we will find that R4 prefers to advertise the external route to 70.70.70.0, learned from R2. To fix this we increase the local preference to 200 on R3 for all prefixes learned from R6. R4 now prefers the route learned from R3, advertises it to CAT2, and does not send the external route to CAT2 or R3. One final issue to clear up in Part 1 is demonstrated in the output below. Notice that CAT2, at this point, has two routes for 70.70.70.0. It has learned one from R3 and one from R4, but both have a next-hop of R3!

```
CAT2#sh ip bgp
BGP table version is 3, local router ID is 172.16.120.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.10.0/24    0.0.0.0                  0         32768 i
*  70.70.70.0/24    172.16.34.3                            0 300 100 700 ?
*>                  172.16.34.3                            0 300 100 700 ?
CAT2#sh ip bgp 70.70.70.0
BGP routing table entry for 70.70.70.0/24, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  172.16.34.3
  300 100 700
    172.16.34.3 from 172.16.34.3 (172.16.140.22)
      Origin incomplete, localpref 100, valid, external
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
300 100 700
  172.16.34.3 from 172.16.34.4 (172.16.104.1)
    Origin incomplete, localpref 100, valid, external, best
```

When three or more BGP speakers are configured on a single multi-access network, rules for setting the next-hop attribute are unique. When R4 advertises the 70.70.70.0/24 prefix to CAT2 over an EBGP neighbor relationship, it does not set its own update source as the new next-hop address for the prefix. The 70.70.70.0/24 prefix retains the next-hop attribute of 172.16.34.3. Since we require that traffic from CAT2 to 70.70.70.0 pass through R4, we need to issue the command **neighbor 172.16.34.10 next-hop-self** on R4. This will override BGP's attempt to shortcut across the Ethernet directly to R3. Finally, just to make sure that CAT2 prefers the path through R4 to the path through R3, we have raised the **weight** on the R4 route.

Part 2: Synchronize network 10.10.10.0 on R3. Make R4 the next hop.

Our preferred path requires that R3 learn 10.10.10.0/24 from R4 and advertise it to R6. Since synchronization is enabled, that same prefix will have to be in R3's routing table from a source other than BGP. In brief, our solution is as follows:

1. Advertise the 10.10.10.0/24 prefix via BGP from R4 to R3 with **next-hop-self**.
2. On R3, raise the local preference to 200 for routes learned from R4
3. On R4, redistribute BGP into RIP.
4. Add the command **IP summary-address 10.10.10.0 255.255.255.0** to R4 interface F0/0.
5. On R3 raise the administrative distance of EBGP routes above 120.

Steps 1 and 2 advertise the prefix from R4 to R3 with a local preference of 200, so that R3 will prefer the IBGP path through R4, rather than the EBGP path to CAT2. Steps 3 and 4 allow us to get the prefix 10.10.10.0/24 into the routing table on R3 as a RIP route with next-hop of R4. If we simply advertised it into RIP from CAT2, then the next-hop from R3 would be CAT2, which is unacceptable. If we only redistributed 10.10.10.0/24 from BGP into RIP on R4, then the next-hop on R3 would again be CAT2, because of the way BGP preserves next-hop on multi-access links. By using the **IP summary-address** command on R4, we achieve a "re-sourcing" of the route on R4, making R4 the next-hop for R3. Step 5 uses the **distance bgp** command to make sure the RIP route is installed even if the EBGP route is learned first.

Verification:

We can verify the required path with the traceroute command, from Cat2 and FRS:

```
CAT2#traceroute 70.70.70.70

Type escape sequence to abort.
Tracing the route to 70.70.70.70

  1 172.16.34.4 0 msec 4 msec 0 msec
  2 172.16.34.3 4 msec 0 msec 4 msec
  3 172.16.36.6 20 msec 20 msec 20 msec
  4 172.16.16.1 24 msec 20 msec 20 msec
  5 172.16.17.7 24 msec *  20 msec
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
FRS#traceroute 10.10.10.10

Type escape sequence to abort.
Tracing the route to 10.10.10.10

  1 172.16.17.1 4 msec 4 msec 4 msec
  2 172.16.16.6 8 msec 4 msec 4 msec
  3 172.16.36.3 20 msec 20 msec 20 msec
  4 172.16.34.4 20 msec 24 msec 24 msec
  5 172.16.34.10 24 msec *  20 msec
```

The following output verifies route synchronization in AS300 for the required prefixes on R3 and R4.

```
R3#sh ip bgp 10.10.10.0
BGP routing table entry for 10.10.10.0/24, version 2
Paths: (2 available, best #2)
  Advertised to non peer-group peers:
  172.16.34.10 172.16.36.6
  1000
    172.16.34.10 from 172.16.34.10 (172.16.120.1)
      Origin IGP, metric 0, localpref 100, valid, external
  1000
    172.16.34.4 from 172.16.34.4 (172.16.104.1)
      Origin IGP, metric 0, localpref 200, valid, internal, synchronized, best
```

Here is the same output for R4.

```
R4#sh ip bgp 70.70.70.0
BGP routing table entry for 70.70.70.0/24, version 12
Paths: (2 available, best #1)
  Advertised to non peer-group peers:
  172.16.34.10 172.16.124.5
  100 700
    172.16.36.6 (metric 1) from 172.16.34.3 (172.16.140.22)
      Origin incomplete, localpref 200, valid, internal, synchronized, best
  100 700
    172.16.124.2 (metric 64) from 172.16.124.5 (172.16.124.5)
      Origin incomplete, localpref 100, valid, external
```

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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.7 Router Maintenance

HIDDEN ISSUES TO SPOT WITH ROUTER MAINTENANCE CONFIGURATION

Issue: A network administrator named Sally would like to access R1 via Internet Explorer at any time from the subnet on VLAN 30 only.

Solution:

Enable the global configuration **ip http server** on router R1. This will allow router R1 to be accessed via a web browser. To restrict the access to just users on VLAN 30, create an access-list permitting only 172.16.10.0/24. Apply this access-list to http with the command **ip http access-class 1** in global configuration mode.

Issue: Sally would like to access R1 using a non-standard http port and with a username/password level of security with her user name as Sally.

Solution:

The standard http port assignment is 80. You can change the http port for Cisco router access with the command **ip http port XX** where XX is an arbitrarily selected TCP port. You can assign username/password security to the http access of a Cisco router with the command: **ip http authentication local**. To complete this configuration you must create a username/password entry on R1 for Sally.

In the 12.2 doc CD, you can find a discussion on http access and customization in the Cisco IOS Configuration Fundamentals Configuration Guide> Cisco IOS User Interfaces> Using the Cisco Web Browser User Interface. You might note that starting in 12.2(15)T there is a **IP http secure-server** option that permits SSL-encrypted communication with the router. This feature is discussed in the 12.3 Configuration Guide.

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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.8 Security

HIDDEN ISSUES TO SPOT WITH THE SECURITY CONFIGURATION

Issue: None of the devices on VLAN 30 should be able to telnet straight to router R5 unless Sally authenticates herself on router R1. Once Sally has authenticated herself on router R1, the users on VLAN 30 should be able to maintain a telnet session for no longer than 10 minutes.

Solution:

This task refers to two characteristics of Lock-and-Key Security: (1) access to a specific service based upon a user being authenticated and (2) a time limit to the access. This feature is well documented in the 12.2 Cisco IOS Security Configuration Guide on the doc CD, under Traffic Filtering and Firewalls.

The first step is to create and apply an access-list on R1 that denies VLAN 30 traffic to R5, permits required traffic, and denies everything else. We created access-list 120 and applied it inbound to S0/0 and S0/0.1.

1. access-list 120 remark ---------- Router Security ----------
2. access-list 120 dynamic VLAN30 timeout 10 permit tcp any host 172.16.25.5 eq telnet
3. access-list 120 deny tcp 172.16.32.0 0.0.3.255 host 172.16.25.5 eq telnet
4. access-list 120 permit ospf any any
5. access-list 120 permit eigrp any any
6. access-list 120 permit pim any any
7. access-list 120 permit icmp any any
8. access-list 120 permit igmp any any
9. access-list 120 permit tcp any host 172.16.124.1 eq telnet
10. access-list 120 permit tcp any host 172.16.13.1 eq telnet
11. access-list 120 permit tcp 172.16.32.0 0.0.3.255 host 172.16.124.1 eq 8080
12. access-list 120 permit tcp any any eq bgp
13. access-list 120 permit tcp any eq bgp any
14. access-list 120 deny ip any any log

Line 3, above, denies telnet traffic from VLAN 30 to R5. Line 2 is the dynamic access-list statement that will turn into a “permit” when SALLY authenticates. Notice the timeout of 10 (minutes) is placed in the access-list statement. Many of the other entries are self-explanatory. Lines 9 and 10 are necessary to allow SALLY to telnet to the router and trigger the dynamic entry. The final line is there to help us determine if we forgot to permit any required traffic.

When the **access-enable** command is run the dynamic access-list entry is created on the interface used by the vty session. You can demonstrate the link between the interface used by the vty session and the access-list grouping as follows: telnet into R2 from Cat 1 and enter access-enable. Then telnet into R2 from R5 and do the same thing. You should get output like that below complaining there is no access-list associated with the incoming interface.

```
CAT1#telnet 172.16.124.5
R2#access-enable
No input access group defined for Serial0/0.
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
R2#

R5#telnet 172.16.125.5
R2#access-enable
No input access group defined for Serial1/0.
```

Verification:

Before testing the Lock-and-Key operation, make sure you can telnet from R1 to R5. Remember, too, that you will have to enter **login local** on the vty lines of R1 so that Sally can authenticate to the local database. Below you see a sequence of telnets to verify the configuration.

```
1.  CAT2#telnet 172.16.25.5
2.  Trying 172.16.25.5 ...
3.  % Destination unreachable; gateway or host down

4.  CAT2#telnet 172.16.124.1
5.  Trying 172.16.124.1 ... Open

6.  User Access Verification

7.  Username: sally
8.  Password:
9.  [Connection to 172.16.124.1 closed by foreign host]

10. CAT2#telnet 172.16.25.5
11. Trying 172.16.25.5 ... Open

12. User Access Verification

13. Password:
14. R5#
15. R5#

16. R1#sh access-list 120
17. Extended IP access list 120
18. 10 Dynamic VLAN30 permit tcp any host 172.16.25.5 eq telnet
    permit tcp any host 172.16.25.5 eq telnet (33 matches)
19. 20 deny tcp 172.16.32.0 0.0.3.255 host 172.16.25.5 eq telnet (4 matches)
20. 30 permit ospf any any (437 matches)
21. [remainder not shown]
```

At line 1, you see a failed attempt to telnet from CAT2 to R5. This is as expected. We then telnet to R1 and authenticate as Sally. Note that the session is then automatically dropped at line 9. This is also the expected result, but it can be disconcerting to unsophisticated users. However, Sally's successful telnet triggered the dynamic access-list entry, so when we again attempt to telnet to R5, at line 10, it is permitted. At line 18 you see the dynamic access-list entry, and below it you see the temporary entry it created in response to the access-enable command issued on the vty session.

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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.9 IPv6

HIDDEN ISSUES TO SPOT WITH THE IPv6 CONFIGURATION

IPv6 Frame-Relay

Issue: Configure IPv6 subnet FEC0::12:0/125 between R1 and R2.

Solution:

Frame-Relay is already configured between R1 and R2, but IPv6 will not let you configure the same subnet on different point-to-point subinterfaces of the same interface. Therefore the IP addresses between R1 and R2, and between R1 and R3 are in different IPv6 subnets. As usual with Frame-Relay, the point-to-point subinterface is the easiest part of the configuration, because all you need is to add IPv6 addresses to it and it will work. The other end (R2) is a physical interface, however, and both routable and link-level IPv6 addresses will have to be mapped to the correct PVCs.

R1

```
interface Serial0/0.2 point-to-point
 ip address 172.16.124.2 255.255.255.0
 ip pim dense-mode
 ip ospf network point-to-multipoint
 ip igmp join-group 233.3.3.3
 ipv6 address FEC0::12:1/125
 ipv6 ospf network broadcast
 ipv6 ospf 100 area 0
 frame-relay interface-dlci 102
```

R2

```
interface Serial0/0
 ip address 172.16.124.5 255.255.255.0
 ip pim dense-mode
 encapsulation frame-relay IETF
 ip ospf network point-to-multipoint
 ipv6 address FEC0::12:2/125
 ipv6 ospf network broadcast
 ipv6 ospf 100 area 0
 no fair-queue
 frame-relay map ipv6 FEC0::12:1 201 broadcast
 frame-relay map ipv6 FE80::2D0:58FF:FE95:C8A1 201 broadcast
 frame-relay map ip 172.16.124.2 201 broadcast
 frame-relay map ip 172.16.124.5 201
 no frame-relay inverse-arp
 frame-relay lmi-type cisco
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Verification:

Verify that frame-relay maps are set up correctly between IPv6 addresses and PVCs:

```
R2#sh frame-relay map
Serial0/0 (up): ipv6 FEC0::12:1 dlci 201(0xC9,0x3090), static,
              broadcast,
              IETF, status defined, active
Serial0/0 (up): ipv6 FE80::2D0:58FF:FE95:C8A1 dlci 201(0xC9,0x3090), static,
              broadcast,
              IETF, status defined, active
```

Issues: Configure the required IPv6 subnets.

Solution:

Configuration for IPv6 subnets between R1, R2, R3 and R5 is similar to the previously described task, and following the same rules will result in a working configuration.

IPv6 RIP

Issue: Configure IPv6 RIP between R1, R3 and R4.

***Solution*:**

RIP for IPv6 is configured on the interface level. In this task, the process is given the identifier "RIPV6." This same identifier is used across the network, even though it has only local significance on each router. Under each RIP interface enter the command **ipv6 rip RPV6 enable**.

***Verification*:**

To verify that RIP is working, make sure that R3 and R4 have learned the IPv6 prefix on the other side of R1.

```
R4#sh ipv6 route fec0::13:0/125
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
R   FEC0::13:0/125 [120/2]
     via FE80::2D0:58FF:FE95:C8A1, Serial0/0
```

IPv6 OSPF

Issue: Configure OSPF area 0 on Frame-Relay link between R1 and R2, and on the link between R1 and R6. Configure OSPF area 25 on the Frame-Relay link between R2 and R5. OSPF neighbors must elect a DR on every link. OSPF neighbors must find each other automatically on all links.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

***Solution*:**

IPv6 OSPF areas are configured on the interface level using the **ipv6 ospf *process-id* area** command. For example, on R1 you enter the command **ipv6 ospf 100 area 0** under interface S0/0.2. To change the OSPF network type enter the command **ipv6 ospf network broadcast**. Note that a virtual link will need to be configured across area 25 and that the syntax is done under the OSPFv3 process. Also notice that the IPv6 OSPFv3 RID is a 32-bit number, by default chosen just as the OSPFv2 RID is chosen (highest IPv4 loopback).

***Verification*:**

To verify that OSPF is adjacent everywhere, make sure that the OSPF database on each router has full information about all OSPF routers and networks. Issuing the **show ipv6 ospf database** command can do this. Alternatively, go to each router and enter **show ipv6 ospf neighbor**.

IPv6 Redistribution

Issue: Redistribute RIP and OSPF on R1.
***Solution*:**

Redistribution between IPv6 RIP and IPv6 OSPF is done on R1 in router configuration mode:

R1

```
ipv6 router ospf 100
 log-adjacency-changes
 redistribute connected metric 1
 redistribute rip RIPV6 metric 1
!
ipv6 router rip RIPV6
 redistribute connected metric 1
 redistribute ospf 100 metric 1
```

Remember that IPv6 routing protocols will not redistribute connected networks into another routing protocol. You must redistribute connected networks manually.

***Verification*:**

At this point, you should have reachability to each IPv6 address from every IPv6 router. The only exception is the 2001::105:1 address used in the BGP section. One way to verify this is by running the following TCL script on each router.

```
foreach address {
#-----------R1---------------
FEC0::12:1
FEC0::13:1
FEC0::14:1
FEC0::16:1
#-----------R2---------------
FEC0::12:2
FEC0::25:2
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
#------------R3----------------
FEC0::13:3
#------------R4----------------
FEC0::14:4
#------------R5----------------
FEC0::25:5
FECO::105:1
2001::105:1
#------------R6----------------
FEC0::16:6
} {ping $address}
```

IPv6 BGP

Issue: Configure BGP AS 500 on router R5. Advertise prefix 2001::105:0/125 into BGP as internal. Make sure that only R2 and R6 have this prefix in their IPv6 BGP tables. Do not change AS numbers. Leave IPv4 BGP configuration intact. Automatic configuration changes and reordering by IOS are exempt from this requirement.

***Solution* for basic IPv6 BGP peering:**

First, configure the BGP routing process on R5 and the neighbor statement to IBGP peer R6:

```
router bgp 500
 no synchronization
 bgp log-neighbor-changes
 neighbor FEC0::25:2 remote-as 100
 no auto-summary
```

Note that as soon as the configuration above is entered, IOS automatically adds the address-family entries as shown below. You must activate the IPv6 BGP neighbor under the IPv6 address-family.

```
 address-family ipv4 multicast
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv6
 neighbor FEC0::25:2 activate
 network 2001::105:0/125
 exit-address-family
 !
 address-family ipv4
 no neighbor FEC0::25:2 activate
 no auto-summary
 no synchronization
 exit-address-family
```

On R2, configure the IPv6 EBGP peer relationship with R5. R2 has an existing BGP AS 100 configuration. Add the command **neighbor fec0::25:2 remote-as 500** directly under the BGP 100 process, then enter the following commands to activate the neighbor:

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
address-family ipv6
 neighbor fec0::25:5 activate
```

You can verify that the neighbor relationship has formed as follows:

```
R2#show bgp ipv6 unicast summary
BGP router identifier 172.16.124.5, local AS number 100
[output removed for brevity]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
FEC0::25:5      4   500      10       9        2    0    0 00:05:01        1
```

Now, let's set up the IBGP peering between R2 and R6. On R2, add the following commands:

```
router bgp 100
 neighbor fec0::16:6 remote-as 100

address-family ipv6
 neighbor fec0::16:6 activate
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

Netmasterclass DOIT-Lab-3 IPv6 BGP Diagram

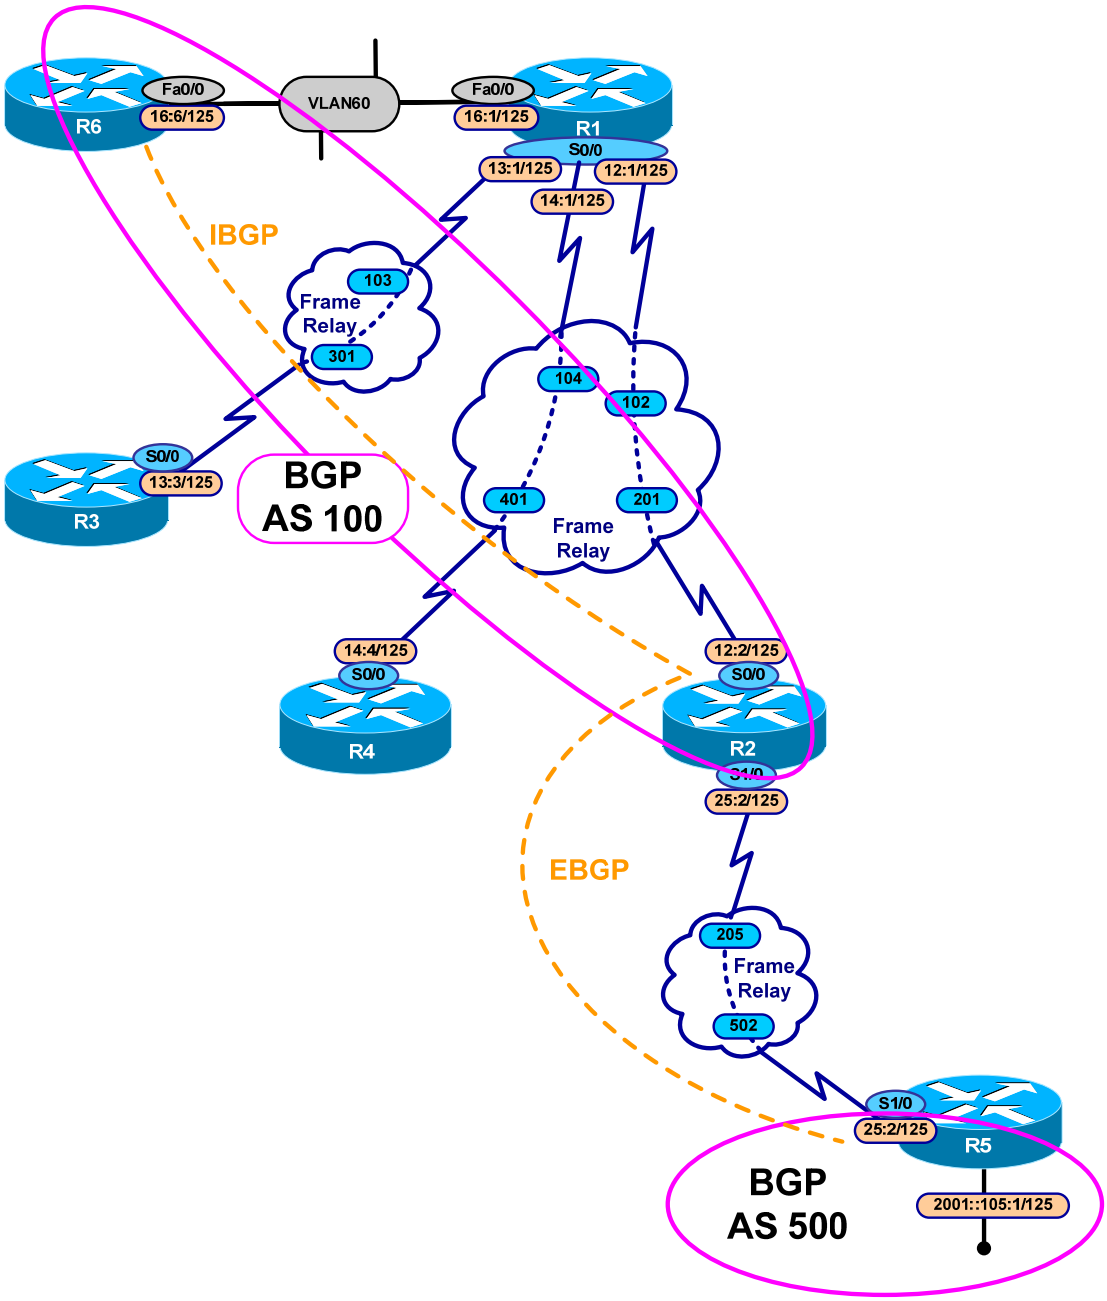

To complete the basic peering, enter the following commands to the existing BGP 100 configuration on R6:

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
router bgp 100
 neighbor fec0::12:2 remote-as 100

address-family ipv6
 neighbor fec0::12:2 activate
```

Verify that your peering has been successfully by issuing the **show bgp ipv6 unicast summary** command on R2:

```
R2#sh bgp ipv6 unicast summary
BGP router identifier 172.16.124.5, local AS number 100
[output removed for brevity]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
FEC0::16:6      4   100       4       5        2    0    0 00:00:42        0
FEC0::25:5      4   500      22      21        2    0    0 00:17:17        1
```

Solution: Advertise prefix 2001::105:0/125 into BGP as internal. Make sure that only R2 and R6 have this prefix in their IPv6 BGP tables.

All we need to do is enter the command **network 2001::105:0/125** under **address-family ipv6** on R5. To verify the advertisement, enter the command **show bgp IPv6 unicast** on R5, R2 and R6

```
R5#sh bgp ipv6 unicast
[output removed for brevity]

   Network          Next Hop            Metric LocPrf Weight Path
*> 2001::105:0/125  ::                       0         32768 I


R2# sh bgp ipv6 unicast
[output removed for brevity]

   Network          Next Hop            Metric LocPrf Weight Path
*> 2001::105:0/125  FEC0::25:5               0             0 500 I


R6# sh bgp ipv6 unicast
[output removed for brevity]

   Network          Next Hop            Metric LocPrf Weight Path
*>i2001::105:0/125  FEC0::25:5               0    100      0 500 i
```

Solution: Make sure that only R2 and R6 have this prefix in their IPv6 BGP tables.

Even though R6 has the 2001::105:0/125 prefix in its table, it cannot ping it, because R1 is in the path between R6 and the R2, and it is not running IPv6 BGP. In order to provide reachability within the requirements we have redistributed BGP into OSPFv3 on R2.

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

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.10 QoS

HIDDEN ISSUES TO SPOT WITH THE QOS CONFIGURATION

Issue: Configure QOS on the GigabitEthernet 0/1 interface on CAT2 using the most simplified method of configuration

Solution:

The **auto qos voip** command has the effect of entering a number of other commands, both global and interface level, which are appropriate for voice quality of service. It has two options, **trust** and **cisco-phone**. The former will trust existing CoS marking unconditionally. The latter trusts existing CoS only if CDP detects a Cisco phone connected to the port. The **trust** option would be most appropriate for a Gigabit Ethernet uplink port, and the **cisco-phone** option would be most appropriate for an access port. The command **show auto qos** prints out the commands entered by the macro. Each command is described in the table below.

Command	Effect
Gobal Commands:	
mls qos map cos-dscp 0 8 16 26 32 46 48 56	maps CoS 3 to DSCP 26 (AF31), CoS 5 to DSCP 46 (EF) - preserves original DSCP in voice control and voice data packets
mls qos min-reserve 5 170	"reserve level 5" defined as depth of 170 packets
mls qos min-reserve 6 10	"reserve level 6" defined as depth of 10 packets
mls qos min-reserve 7 65	"reserve level 7" defined as depth of 65 packets
mls qos min-reserve 8 26	"reserve level 8" defined as depth of 26 packets
mls qos	enables qos – ports default to untrusted
Interface Gig0/1:	
mls qos trust cos	maintain existing CoS markings
wrr-queue bandwidth 20 1 80 1	about 20% bw to Q1, 80% bw to Q3
wrr-queue min-reserve 1 5	assign queue depth of 170 to Q1
wrr-queue min-reserve 2 6	assign queue depth of 10 to Q2
wrr-queue min-reserve 3 7	assign queue depth of 65 to Q3
wrr-queue min-reserve 4 8	assign queue depth of 26 to Q4
no wrr-queue cos-map	clear existing CoS to queue map
wrr-queue cos-map 1 0 1 2 4	map CoS 0,1,2, and 4 to Q1
wrr-queue cos-map 3 3 6 7	map CoS 3, 6 and 7 to Q3
wrr-queue cos-map 4 5	map CoS 5 to Q4
priority-queue out	makes Q4 the Priority Queue

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

The global configuration mode commands made by the macro do the following: customize the default cos-dscp map to maintain original voice control and data DSCP markings, define four queue depth levels to be applied on the interface, and enable mls qos. The interface-level commands enable trust for existing CoS, assign about 20% of bandwidth to queue 1 and 80% to queue 3, assign non-default packet depths to the queues, map CoS to queues, and make queue 4 a priority queue. The bottom line result in terms of voice QoS is that voice data gets assigned to a priority queue with a low queue depth. Note that only queue 4 can be a priority queue.

Issue: Two IP phones are connected to ports Fa0/10 and Fa0/11 of CAT1. Configure QOS on these ports using the same technique as in the previous task.

Solution:

Configure the command **auto qos voip cisco-phone** under the FastEthernet ports 0/10 and 0/11 on CAT1. These commands make nearly the same changes to CAT1 as the previous command issued on CAT2. The only difference is the commands **mls qos trust device cisco-phone** issued on each interface. It implements CoS trust based on CDP discovery of a Cisco phone.

Since we have our phones on CAT1 and the uplink on CAT2, we need to make changes to the links connecting the switches. With **mls qos** enabled, existing CoS and DSCP markings will be zeroed-out if we do not explicitly trust. Enter the command **mls qos trust cos** on ports f0/13 and f0/14 of both CAT1 and CAT2.

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".

3.11 Catalyst Specialties

HIDDEN ISSUES TO SPOT WITH CATALYST SPECIALTIES

Issue: Enable message logging on CAT1 writing the messages to a file on local flash called "nmc.log". The log file should be twice as large as its default size. Only error messages about hardware and software should be logged.

Solution:

All of these requirements can be fulfilled with a single global configuration command on CAT1, **logging file flash:nmc.log 8192 4.** The default maximum file log size is 4096. You are asked to make the log file twice the size of the default. Therefore, it should be 4096*2 = 8192. You are supposed to log only error messages related to Catalyst hardware or software. There are 8 logging error levels numbered 0-7. By setting the logging level to 4, the "system warning" level, you will log all levels from 0-4. which cover all levels that relate to system hardware and software errors. Information on this topic can be found in the 3550 Configuration Guide under Configuring System Message Logging.

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

3.12 Address Administration

HIDDEN ISSUES TO SPOT WITH THE ADDRESS ADMINISTRATION

Issue: Forward traffic generated by the workstations on VLAN 10 to both routers R3 and R6 in a 50/50 approximate distribution. R6 should be configured to provide TCP/IP connectivity information to the workstations on VLAN 10 (IP address, DNS, etc.)

Solution:

We can accomplish these tasks by configuring two HSRP groups and two DHCP servers on the subnet. Each DHCP server will give out the virtual address of one of the HSRP groups. Each group will be active on a different router. There is a 50/50 chance that any workstation connecting to VLAN 10 will obtain its DHCP-supplied IP address configuration from either router R3 or R6, so the traffic distribution on VLAN 10 will be divided between routers R3 and R6 at an approximate distribution of half and half.

Standby group 1 has a virtual address of 172.16.36.1 and has a higher priority on router R3. The DHCP server on R3 gives out the address 172.16.36.1 as the gateway. Standby group 2 has a virtual address of 172.176.36.2 and has a higher priority on R6. The DHCP server on R6 gives out the address 172.16.36.2 as the gateway. Below you see verification of the MHSRP configuration as seen on R3.

```
R3#sh standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Prio P State    Active addr     Standby addr    Group addr
Fa0/0       1    110  P Active   local           172.16.36.6     172.16.36.1
Fa0/0       2    90   P Standby  172.16.36.6     local           172.16.36.2
```

Issue: All workstations are Microsoft clients using a hybrid NETBIOS node type.

Solution:

To fulfill this requirement, configure the following command under the DHCP server or “pool” configuration on both routers R3 and R6: **netbios-node-type h-node**. This node type checks a WINS server among its NetBIOS name resolution methods.

Issue: In case of the failure of the respective links on either routers R3 or R6, traffic should be forwarded through the remaining link. When the failed link recovers, return the previous forwarding scheme.

Solution:

To configure MHSRP to adjust itself due to the failure of a local link on the preferred router, configure the HSRP track option on both routers R3 and R6. Router R3 will track its Frame-Relay connection, and R6 will track its F0/0.60 connection. If the F0/0.60 interface on R6 goes down, for example, then the HSRP priority of standby group 2 will decrease by 30, to 80. R6 will now have a lower priority than R3 for

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

standby group 2, and R3 will take over routing for virtual address 172.16.36.2. To have the configuration return to the previous forwarding scheme when these links recover, configure both HSRP groups with the "preempt" option.

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".

3.13 Multicast

HIDDEN ISSUES TO SPOT WITH THE MULTICAST CONFIGURATION

Issue: Configure PIM neighbor relationships between R4 and R1, R1 and R2 and R2 and R5. The PIM network is based on a Flood and Prune concept.

Solution:

These requirements suggest PIM Dense Mode, since it builds trees based on a "Flood and Prune" concept. You are then directed to form very specific PIM Dense Mode neighbor relationships, as shown on the diagram below.

Make sure that all of your Frame-Relay map statements have the **broadcast** parameter so that you can forward multicast traffic out of all Frame interfaces.

Issue: Join one of the interfaces on the routers R1, R2, R4 and R5 to group 233.3.3.3

Solution:

Configure **ip igmp join-group 233.3.3.3** on the specified routers. Whenever you configure the **ip igmp join-group command** on an interface, make sure the following three configuration requirements are also met: (1) an IP address is on the interface configured with the ip igmp join-group command, (2) the IP address is propagated by a unicast routing protocol and (3) multicast routing is also configured on the same interface with a command such as **ip pim dense-mode**.

Issue: Make sure you receive ping replies from all routers on CAT2

For traffic that is sourced from router CAT2 (which is represented in the (S,G) entry of 172.16.34.10,233.3.3.3), no RPF lookup or Outgoing Interface List problems should exist. This is because R1 has **two separate subinterfaces** connecting it to routers R4 and R2. In this particular multicast scenario, all traffic generated by CAT2 will flow through R4 up to R1. If R1 were configured with a single multipoint subinterface, then multicast traffic would have to come into that multipoint and then be

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

forwarded back out of the same interface. One of the rules of PIM Dense Mode is that the incoming interface cannot be on the outgoing interface list.

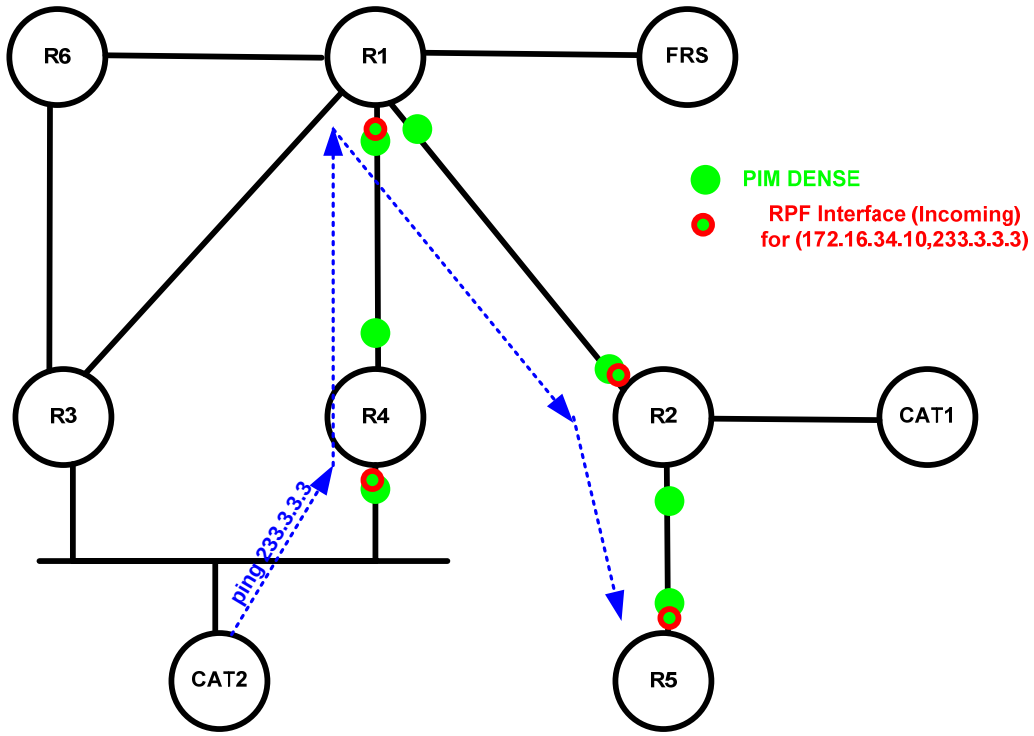

Verification:

First, lets do an extended ping from CAT2 to the multicast address to generate a flow of traffic. Note that we have set it for a large repeat count. That way we will have a long-lasting, consistent stream of multicast packets to use as a troubleshooting tool on the other routers, if necessary. Here you see we are getting returns as expected from R4 (172.16.34.4), R1 (172.16.124.1), R2 (172.16.124.5) and R5 (172.16.25.5).

```
CAT2#ping ip
Target IP address: 233.3.3.3
Repeat count [1]: 999
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 999, 100-byte ICMP Echos to 233.3.3.3, timeout is 2 seconds:

Reply to request 0 from 172.16.34.4, 1 ms
Reply to request 0 from 172.16.25.5, 232 ms
Reply to request 0 from 172.16.124.5, 204 ms
Reply to request 0 from 172.16.124.1, 132 ms
Reply to request 1 from 172.16.34.4, 1 ms
Reply to request 1 from 172.16.25.5, 228 ms
Reply to request 1 from 172.16.124.5, 204 ms
```

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook

```
[Etcetera]
```

The critical router in this scenario is R1. Here is a portion of its mroute table while the traffic is still streaming. Notice that the incoming interface is S0/0.1 and the outgoing interface is S0/0.2. As you will remember from the Frame-Relay section of this lab, these two subinterfaces are in the same subnet. PIM Dense mode has no problem allowing the stream to come in one subinterface and to be sent out another.

```
R1#sh ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report, Z - Multicast Tunnel
       Y - Joined MDT-data group, y - Sending to MDT-data group
Outgoing interface flags: H - Hardware switched
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 233.3.3.3), 1d00h/stopped, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0.1, Forward/Dense, 1d00h/00:00:00
    Serial0/0.2, Forward/Dense, 1d00h/00:00:00

(172.16.34.10, 233.3.3.3), 00:07:08/00:02:59, flags: LT
  Incoming interface: Serial0/0.1, RPF nbr 172.16.124.4
  Outgoing interface list:
    Serial0/0.2, Forward/Dense, 00:07:08/00:00:00
```

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".

© 2005 Copyright NetMasterClass, LLC - All rights reserved.
☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

The material in this lab workbook remains the intellectual property of NetMasterClass., LLC
The purchaser cannot re-distribute the materials in any format or resell this workbook