


**NETMASTERCLASS**  
**ROUTING AND SWITCHING CCIE® TRACK**

# DOIT-200v6

# VOLUME II



## SAMPLE LAB ANSWER KEY

FOR

## CCIE® CANDIDATES

## Disclaimer

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

## Cisco Non-Disclosure Agreement Compliance

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

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


## DOiT-V6 SAMPLE Scenario: Spot the Issue Answer Key

### Table of Contents

S.1	Serial Interface Configuration .....	2
S.2	Catalyst 3550 Configuration.....	2
S.3	OSPF.....	2
S.4	RIP .....	2
S.5	EIGRP .....	2
S.6	ODR .....	2
S.7	BGP .....	2
S.8	IPv6 Routing.....	2
S.9	QOS .....	2
S.10	Address Administration.....	2
S.11	Gateway Redundancy .....	2
S.12	NTP Configuration .....	2
S.13	Multicast .....	2
S.14	IOS Features .....	2



**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 displayed in the Scenario diagram belong to network 172.16.0.0/16.
- Do not rely on dynamic Frame-Relay inverse ARP.
- Do not use any static routes.
- Make sure all IPv4 and IPv6 loopback interfaces are advertised with their original mask, unless noted otherwise.
- Make sure all IPv4 interfaces in the diagram are **reachable** within this internetwork. **DO NOT FORGET THIS!**
- IP subnet 10.1.1.0/24 is excluded from the previous requirement.
- Use conventional routing algorithms.

### Explanation of Each of the Goals and Restrictions:

#### IP subnets in the Scenario diagram belong to network 172.16.0.0/16

The third and fourth octets of the IP addresses displayed on the diagram belong to 172.16.0.0/16.

#### Do not rely on dynamic Frame-Relay Inverse ARP.

This requirement forces you to fulfill your Frame-Relay inverse ARP requirements with Frame-Relay map statements. Think of a Frame-Relay map statement as the equivalent of a static inverse ARP entry.

#### 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 and IPv6 loopback interfaces are advertised with their original mask, unless noted otherwise.

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

#### Make sure all IP interfaces in the diagram are **reachable** within this internetwork. **DO NOT FORGET THIS!**

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

**IP subnet 10.1.1.0/24 is excluded from the previous requirement.**

The subnet 10.1.1.0/24 is used for the QOS section only.

**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 phrase "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 INFORMATION 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

## S.1 Serial Interface Configuration



### **HIDDEN ISSUES TO SPOT WITH THE SERIAL INTERFACE CONFIGURATION**

**Issue:** *The Frame Switch is configured for a full mesh. How do you use only the specified PVCs for user traffic? Configure the specified interface types on the routers.*

**Solution:**

The Frame Relay switch is pre-configured for a full mesh of PVC's. You are instructed to use "a minimum number of DLCI's to provide Layer 3 connectivity". When examining the SAMPLE LAB diagram, you see that the Layer 3 connections over the NBMA network reflect a hub and spoke topology for the subnet 172.16.123.0/24. The interface on the hub router R1 should be a logical interface per the scenario specification, therefore it must be multipoint interface since two PVCs are assigned to this interface. The interfaces on R2 and R3 must be physical. The Frame Relay link between R1 and R4 must be terminated between the physical interface as well. To fulfill these requirements, perform the following tasks:

- Disable inverse arp on all interfaces physical and logical.
- Provide static frame-relay mappings on each of the Frame-Relay attached routers. Make sure that one spoke of the Frame-Relay topology can ping the other spoke. In order to fulfill this requirement, make sure that routers R2 and R3 not only possess a Frame-Relay map statement to router R1, each of these spoke routers must also possess map statements to one another.

**Configuration and verification:**

For the IP subnet 172.16.123.0/24 between R1, R2 and R3

R1:

```
interface Serial0/0.123 multipoint
ip address 172.16.123.1 255.255.255.0
frame-relay map ip 172.16.123.1 102 broadcast
frame-relay map ip 172.16.123.2 102 broadcast
frame-relay map ip 172.16.123.3 103 broadcast
no frame-relay inverse-arp
```

R2:

```
interface Serial0/0
ip address 172.16.123.2 255.255.255.0
encapsulation frame-relay
frame-relay map ip 172.16.123.1 201 broadcast
frame-relay map ip 172.16.123.2 201 broadcast
frame-relay map ip 172.16.123.3 201 broadcast
no frame-relay inverse-arp
```

R3:

```
interface Serial0/0
 ip address 172.16.123.3 255.255.255.0
 encapsulation frame-relay
 frame-relay map ip 172.16.123.1 301 broadcast
 frame-relay map ip 172.16.123.2 301 broadcast
 frame-relay map ip 172.16.123.3 301 broadcast
 no frame-relay inverse-arp
```

For the IP subnet 172.16.123.0/24 between R1, R2 and R3

R1:

```
interface Serial0/0
 ip address 172.16.14.1 255.255.255.0
 encapsulation frame-relay
 no fair-queue
 cdp enable
 frame-relay map ip 172.16.14.1 104 broadcast
 frame-relay map ip 172.16.14.4 104 broadcast
 no frame-relay inverse-arp
 frame-relay lmi-type ansi
```

R4:

```
interface Serial0/0
 ip address 172.16.14.4 255.255.255.0
 encapsulation frame-relay
 cdp enable
 frame-relay map ip 172.16.14.1 401 broadcast
 frame-relay map ip 172.16.14.4 401 broadcast
 no frame-relay inverse-arp
end
```

**Note:** The broadcast keyword is optional for the mapping to local IP address and to the other spoke. The keyword broadcast is used on all the mapping entries for consistency.

**Issue:** *Configure a 64000 Bit/sec serial link between R2 and R5.*

**Solution:**

HDLC is the default encapsulation on CISCO router synchronous serial interfaces. No special configuration is required to specify this encapsulation type. The speed needs to be set up with the "clock rate 64000" interface command on the DCE side of the link. To determine the DCE side, use the "show controllers s N" command where N is the interface number.

**Configuration and Verification:**

R5:



```
interface Serial1/0
 ip address 172.16.25.5 255.255.255.0
 clockrate 64000
```

```
R5#show controllers serial 1/0
Interface Serial1/0
Hardware is Quicc 68360
DCE V.35, clock rate 64000
```



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”.

## S.2 Catalyst 3550 Configuration



### HIDDEN ISSUES TO SPOT WITH THE CATALYST 3550 CONFIGURATION

Like any Catalyst 3550 configuration, you must address the following basic configuration requirements: setting the VTP mode, configuring trunk ports, 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:** Do not use any dynamic VLAN advertisement techniques. Configure VLANs consistently on both switches

**Solution:**

The VLAN Trunking Protocol (VTP) performs VLAN advertisement. A Catalyst switch can operate in one of three VTP modes: Server, Client and Transparent. Of these three modes, the VTP “transparent” mode does not relay on the exchange of VTP message to discover or advertise specific VLAN's. In order to fulfill the configuration requirements of “not using any dynamic VLAN advertisement techniques”, configure VTP transparent mode. VTP Mode Transparent does not advertise any VLAN 's that are locally created.


**Verification:**

```
CAT1#sh vtp stat | inc Operating Mode
VTP Operating Mode           : Transparent
```

```
CAT2#sh vtp stat | inc Operating Mode
VTP Operating Mode           : Transparent
```

Configure VLANs on both switches. Please look at the following diagram for the VLAN connectivity.

### VLAN Distribution Diagram



Verification:

```
CAT1#show vlan brie
```

VLAN Name	Status	Ports
1 default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4 Fa0/5, Fa0/6, Fa0/7, Fa0/9 Fa0/10, Fa0/11, Fa0/12, Fa0/15 Fa0/16, Fa0/17, Fa0/18, Fa0/19 Fa0/20, Fa0/21, Fa0/22, Fa0/23 Gi0/1, Gi0/2
17 ENG	active	
88 DEVELOP	active	
100 TEST	active	
360 PROD	active	Fa0/8
1002 fddi-default	act/unsup	
1003 token-ring-default	act/unsup	

```

1004 fddinet-default          act/unsup
1005 trnet-default           act/unsup
CAT1#

```

```
CAT2#show vlan brie
```

```

VLAN Name                Status      Ports
-----
-
1    default              active     Fa0/2, Fa0/4, Fa0/5, Fa0/6
                                           Fa0/8, Fa0/9, Fa0/10, Fa0/11
                                           Fa0/12, Fa0/15, Fa0/16, Fa0/17
                                           Fa0/18, Fa0/19, Fa0/20, Fa0/21
                                           Fa0/22, Fa0/23, Fa0/24, Gi0/1
                                           Gi0/2
17   ENG                  active     Fa0/7
88   DEVELOP              active
100  TEST                  active
360  PROD                  active     Fa0/3
1002 fddi-default         act/unsup
1003 token-ring-default   act/unsup
1004 fddinet-default      act/unsup
1005 trnet-default       act/unsup
CAT2#

```

**Issue: Do not use a CISCO proprietary tagging method. Allow necessary VLAN's to traverse both FA0/13 and Fa0/14 interfaces between the Catalyst switches. Do not allow VLAN's on Fa0/13 and Fa0/14 explicitly**

**Solution:**

The key word to pay very close attention to in this task is the very last word of this task, the word "explicitly". The word "explicitly" is included within the sentence "Do not allow VLAN's on Fa0/13 and Fa0/14 explicitly". The requirements of this task can be fulfilled by creating a port-channel interface and applying all trunk related configuration commands to the port-channel. Once the port-channel is created and configured, include ports fa0/13 and fa0/14 to be part of the port-channel. By performing these steps, the necessary trunk commands will NOT be "explicitly" configured on ports fa0/13 and fa0/14. Instead, the trunk configuration commands that end up under ports fa0/13 and fa014 in the final configuration script will be "inherited" from the port-channel configuration. By having ports fa0/13 and fa0/14 "inherit" the trunking configuration commands from the port-channel, the desired result of NOT "explicitly" configuring trunking on these ports has been fulfilled.

**Verification:**

```

CAT1#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  s - suspended
        H - Hot-standby (LACP only)
        R - Layer3       S - Layer2
        u - unsuitable for bundling
        U - in use       f - failed to allocate aggregator

```

d - default port

```
Number of channel-groups in use: 1
Number of aggregators:          1
```

Group	Port-channel	Protocol	Ports
1	Po1(SU)	-	Fa0/13(P) Fa0/14(P)

CAT2#sh etherchannel summary

```
Flags: D - down          P - in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3       S - Layer2
       u - unsuitable for bundling
       U - in use       f - failed to allocate aggregator
       d - default port
```

```
Number of channel-groups in use: 1
Number of aggregators:          1
```

Group	Port-channel	Protocol	Ports
1	Po1(SU)	-	Fa0/13(P) Fa0/14(P)

Configure 802.1Q trunking. 802.1Q is an IEEE standard. It is an alternative trunking protocol to ISL– the CISCO proprietary protocol.

**Verification:**

CAT1#show interfaces trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/24	desirable	n-isl	trunking	1
Po1	on	802.1q	trunking	1

CAT2#show interfaces trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/1	on	802.1q	trunking	1
Po1	on	802.1q	trunking	1

**Issue:** Restrict your port configurations only to vlans displayed on the diagram “DOIT Sample Scenario”.

**Solution:**

Restrict the VLAN's that are allowed on any trunks configured in this exam only to the VLAN's that are created in this exam. In order to accomplish this goal, use the following Catalyst interface configuration command: **switchport trunk allowed vlan XXX**. Where XXX is the set of VLANs you want to allow over the trunk port. When specifying your VLAN's, make sure that you do not add any spaces between the listed VLAN's. When you list VLAN's with the **switchport trunk allowed vlan** command, use commas to separate the listed VLAN's. A hyphen can also be used between two listed VLAN numbers to represent a

contiguous range of VLAN's. You can add and delete VLAN numbers from your original list of allowed VLANs with the following: < add, remove, all, except>.

```
CAT1#show int trunk
```

Port	Mode	Encapsulation	Status	Native vlan
Fa0/24	desirable	n-isl	trunking	1
Pol	on	802.1q	trunking	1

```
Port Vlans allowed on trunk
```

```
Fa0/24 1-4094
Pol 88,100,360
```

```
Port Vlans allowed and active in management domain
```

```
Fa0/24 1,17,88,100,360
Pol 88,100,360
```

```
Port Vlans in spanning tree forwarding state and not pruned
```

```
Fa0/24 1,17,88,100,360
Pol 88,100,360
```

```
CAT1#
```

```
CAT2#show int trunk
```

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

```
Port Vlans allowed on trunk
```

```
Fa0/1 10,17,88,100
Pol 88,100,360
```

```
Port Vlans allowed and active in management domain
```

```
Fa0/1 17,88,100
Pol 88,100,360
```

```
Port Vlans in spanning tree forwarding state and not pruned
```

```
Fa0/1 17,88,100
Pol 88,100,360
```

```
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".**

## ***IPv4 Routing Protocols Redistribution***



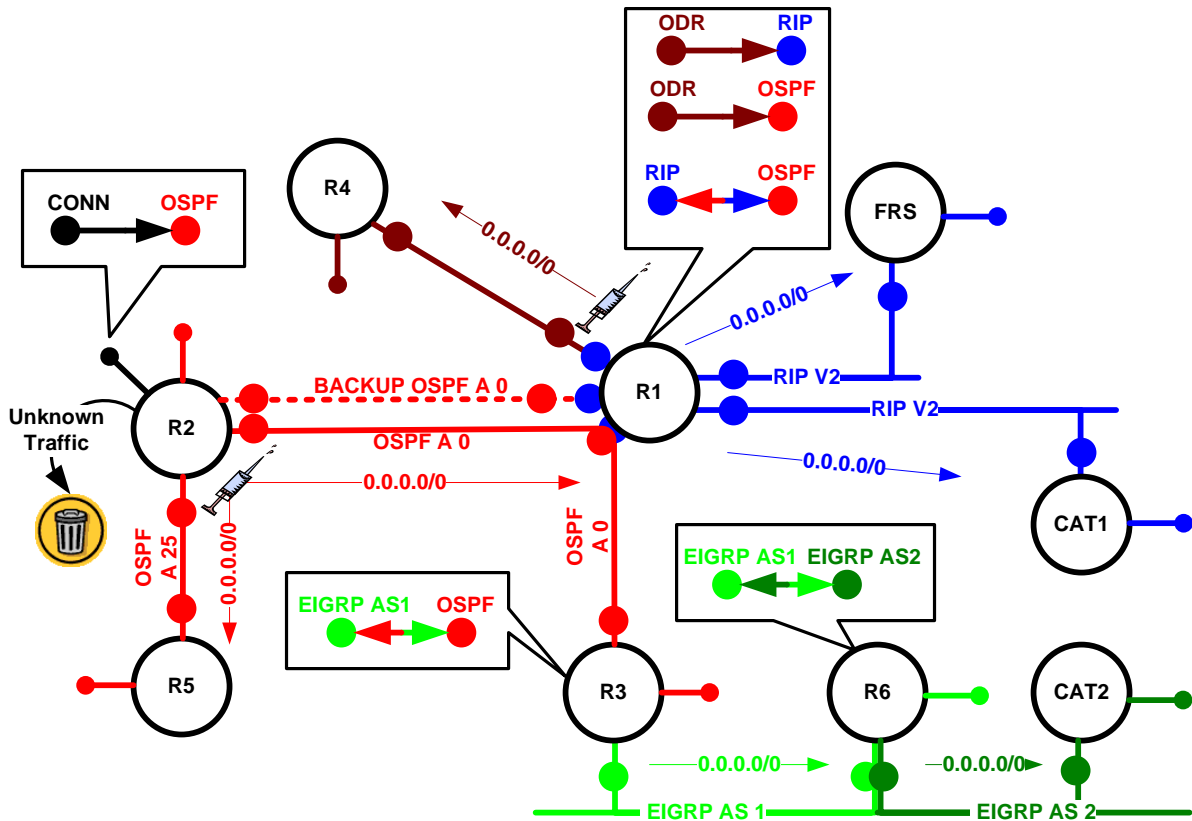
### ***HIDDEN ISSUES TO SPOT WITH ROUTE REDISTRIBUTION***

Before examining the specific issues related to configuring each of the IGP's involved in this Scenario, let's survey the entire topology and determine how all of the different IGP's will interoperate. Performing such a survey will force us to consider the issues related to route redistribution.

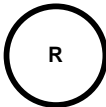




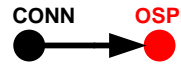




When evaluating a single internetwork topology that contains multiple routing protocols, a good starting point of analysis is to determine whether there is more than one direct or indirect connecting point between two routing protocols. If there is only one connecting point between two routing protocols, providing connectivity between them is relatively simple. If there are two or more connecting points, then providing connectivity between the two routing protocols can be complex. When two or more connecting points exist, you can use them to provide redundancy as well as load balancing and optimum path selection. However, when two or more connecting points exist, you must also assure, at the very least, that no routing loops exist and, whenever possible, no suboptimal paths are selected.

When evaluating this Scenario's internetwork topology and how the routing protocols have been assigned to it, there is only one connecting point between any two routing protocols. When the topology is examined more closely, you notice that there are only two routing protocols that provide transit services to other routing protocols. These two routing protocols are OSPF and EIGRP AS 1. EIGRP AS 1 provides transit routing services in a very limited capacity. It provides transit services only between EIGRP AS 2 and OSPF. Taken as a whole, EIGRP (both EIGRP AS 1 and AS 2) comprise a non-transit routing domain to the rest of the Scenario topology.

OSPF is the primary transit routing protocol for this Scenario. OSPF provides transit services to EIGRP, RIP and ODR. All of the other routing protocols are acting as edge/non-transit routing protocols. The Scenario topology is represented in the following diagram:



### Legend

	Router		Loopback
	RIP		Mutual redistribution, eg. EIGRP and OSPF
	EIGRP		One way redistribution, eg. CONNECTED into OSPF
	OSPF		Prefix injection
	ODR		Trash can

See the scenario master diagram and VLAN table for data link details!

**NOTE:** The colors used in this diagram greatly add to the understanding of redistribution applied to this Scenario. If you print this document, attempt to print this page with a color printer.

As the diagram reflects, OSPF, represented in the color red, resides in the center of the Scenario topology. All other routing protocols must traverse OSPF to reach each other. Closely examine the diagram above. It identifies where route redistribution occurs and whether the redistribution performed is one-way or two-way. A brief description of the redistribution operations performed in this Scenario is provided below:

### **Locations Where Two-Way Redistribution is Performed**

Two-way redistribution is performed between OSPF and the following routing protocols: RIP v2 and EIGRP AS 1. Two-way redistribution is performed to provide reachability between all routers in this topology. As the diagram reflects, the redistribution operations are performed on routers R1 and R3. All prefixes from RIP v2 and EIGRP AS 1 (as well as EIGRP AS 2 via EIGRP AS 1) are injected into OSPF. At the very least, the only prefix that OSPF needs to inject into these routing protocols is a 0.0.0.0/0 prefix. By injecting a 0.0.0.0/0 prefix into the edge non-transit routing protocols, universal connectivity will be fulfilled. OSPF has the ability to inject a 0.0.0.0/0 prefix into these routing protocols since the OSPF router R2 is originating a 0.0.0.0/0 prefix due to the fact that R2 is configured with a “default-information origination always” command. To learn why router R2 is configured with “default-information origination always”, read the OSPF Spot the Issues section for more details.

Two-way redistribution is also performed on router R4 between EIGRP AS 1 and EIGRP AS 2. See the EIGRP Spot the Issues section for more details.

### **Locations Where One-Way Redistribution is Performed**

One-way redistribution is performed in two places: (1) from the “connected” route source into OSPF on router R2 and (2) from ODR into OSPF on router R1. Each of these one-way redistribution operations warrants a brief discussion.

Regarding the one-way redistribution of a connected route into OSPF on R2, remember to configure a route-map or a distribute-list to assure that only the desired connected interface(s) is injected into OSPF. Without a route-map or distribute-list, all connected interfaces on R2 will be injected into OSPF as Type 5 LSA's.



### Some Thoughts to Consider with This Redistribution Plan

You might consider configuring the following highly restricted one-way redistribution operation on router R1: redistribute only a 0.0.0.0/0 from OSPF into RIP v2; and not redistribute any routing protocol into OSPF. You might consider performing this operation for two reasons: (1) Since R1 is running ODR and RIP v2, it already possesses complete routing information for each of these routing protocols and (2) since the internal OSPF router R2 is originating a 0.0.0.0/0 route that is in turn learned by R1, router R1 possesses a 0.0.0.0/0 route that it can redistribute into RIP v2. Following this logic, router R1 merely needs to redistribute the 0.0.0.0/0 prefix it learns from R2 into RIP v2 to provide universal connectivity between these routing protocols. Due to the 0.0.0.0/0 prefix generated by R2 and necessarily traversing R1, all traffic that matches the 0.0.0.0/0 prefix will be forwarded to R1.

This logic seems to be valid until you consider the forwarding information possessed by two sets of routers: Set #1, OSPF internal routers R2 and R5 and Set #2, internal EIGRP speaking routers R6 and CAT2.

Regarding the first set of routers – OSPF internal routers R2 and R5 – the RIP v2 routes must be redistributed into OSPF to fulfill their universal connectivity requirements. Since router R2 originated the 0.0.0.0/0 route, it will black hole all traffic generated by R5 destined to the RIP v2 routing domain. Also, R2 will possess no 0.0.0.0/0 entry in its routing table since it is the originator of the default route. As a result, it possesses no routing information to reach the RIP v2 networks. Therefore, R2 must receive the prefixes from the RIP v2 routing domains. Consequently, RIP v2 must be redistributed into OSPF to fulfill the connectivity requirements of the internal OSPF routers R2 and R5.

For the second set of routers, the EIGRP internal routers, their prefixes must be redistributed into OSPF to provide routing information to all internal OSPF routers as well as the routers residing in ODR and RIP v2.

### A Comment on the OSPF “Default-Information Originate” Command

A comment needs to be added about the effect of originating a 0.0.0.0/0 network with commands such as **default-information originate always** with OSPF. By configuring this command, routers will attract all traffic that is not represented in any other routers’ routing table. All routers in this Scenario will ultimately forward all “unknown traffic” to router R2. When router R2 receives this traffic, it will drop the traffic since it does not possess a default route itself. When you examine the supplied redistribution diagram, you will see a trash can next to router R2. This represents the behavior of R2 when it receives traffic that it has no routing table entry for. It will drop the traffic into the bit bucket or “trash can”.

## Redistribution Table

The following table provides a useful summary of which prefixes were imported into a given routing protocol. Pay special attention to the color-coding of the table. The colors exactly match the colors used in the diagram. Whenever a permit column for a given routing protocol is completely empty, it reflects that no prefixes were redistributed into the routing protocol. This represents that the routing protocol is involved in one-way redistribution.

Redist Point	Into RIP		Into OSPF		Into EIGRP	
	PERMIT	DENY	PERMIT	DENY	PERMIT	DENY
R1	All ODR All OSPF Inc 0.0.0.0		All ODR All RIP			
R3			All EIGRP		All OSPF Inc 0.0.0.0	
R4					All AS1 to AS2 All AS2 to AS1	

*NOTE: The colors used in this table greatly add to the understanding of redistribution applied to this Scenario. If you print this document, attempt to print this page with a color printer.*

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".

## S.3 OSPF



### HIDDEN ISSUES TO SPOT WITH THE OSPF CONFIGURATION

**Issue:** Let backbone OSPF speakers automatically discover each other and elect a DR.

**Solution:**

The OSPF network type "broadcast" is the correct answer. OSPF speakers will "automatically" discover each other via the multicast address 224.0.0.5 during the initial HELLO exchange. The OSPF speakers will also elect at least a DR and possibly a BDR (A hidden issue is associated with electing a BDR. See the comments below.) Therefore, both conditions of the task are fulfilled.

An issue to be aware of: placement of the DR. The DR must be on the hub of the Frame-Relay cloud. There should not be any BDR elected in this topology. All DROTHER's must form an adjacency with **BOTH** the DR and the BDR. If the hub router is elected as the DR, all spokes can form an adjacency with the DR. Since this is a hub and spoke topology and since the hub router is the DR, the only candidate for a BDR is a spoke router. Any other spoke router **CANNOT** form an OSPF adjacency with another spoke router. Therefore, a BDR cannot be designated in a hub and spoke topology.

To assure that the Frame-Relay hub router is elected the DR and to assure that a spoke router is not elected a BDR, configure the **ip ospf priority** command on the spoke routers and set its value to 0. An OSPF router that is assigned a priority of 0 will never be elected DR or BDR. It is a statically configured DROTHER router.

**Configuration and Verification:**

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.20.5	0	FULL/DROTHER	00:00:33	172.16.123.2	Serial0/0.123
172.16.30.3	0	FULL/DROTHER	00:00:30	172.16.123.3	Serial0/0.123

R1#

Note: Field "Pri" shows the priority configured on the adjacent router. Both R2 and R3 are in the DROTHER state.

**Issue: Add Loopback 172.16.20.5/30 on R2 in OSPF area 20 and summarize it to 172.16.20.0/24**

**Solution:**

Use the network statement under the OSPF process to add the interface to OSPF area 20. Since this prefix is introduced to OSPF as a summary LSA (inter-area) use the **area 20 range 172.16.20.0 255.255.255.0** to summarize the give /30 prefix.

**Issue: Add a Loopback interface on R2 into the OSPF routing process without use of the "network" command.**

**Solution:**

Add the IP address assigned to the loopback interface you created into OSPF using "redistribute connected". Make sure you filter the redistribution process so that only the IP address on the loopback interface you created - and no other IP address on a "connected" network - is injected into OSPF. Applying either a route-map or a distribute-list to the redistribution of the connected networks can fulfill this filtering requirement.

**Issue: Make sure that the network 2.0.0.0 and its subnets do not appear in the routing tables of any router except of R2.**

**Solution:**

The challenge in this task is to make the 2.2.2.0/24 network reachable throughout the pod without announcing it to any other router. Configuring **default-information originate always** on router R2 provides the solution. By originating a 0.0.0.0/0 on router R2, it will attract traffic to destinations that are not represented in any other router's forwarding table for this Scenario. Since no other router in this Scenario has a routing table entry for the 2.2.2.0/24 network, all traffic destined to this address will be forwarded to router R2 since it originated the 0.0.0.0/0 network.

**Issue: Configure area 25 between routers R2 and R5. Add a loopback interface on R5 into OSPF area 50.**

**Solution:**

R5 possesses a connection to R2 via area 25. R5 has no direct connection to Area 0. However, R5 also has a loopback interface assigned to area "50". Since R5 does not possess a direct connection to Area 0, it requires a virtual-link for area 50. Area 25 will be the transit area for the virtual-link.


**Verification:**

```
R2#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 172.16.50.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 25, via interface Serial1/0, Cost of using 781
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
  Adjacency State FULL (Hello suppressed)
  Index 1/2, retransmission queue length 0, number of retransmission 1
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec
R2#

R2#show ip ospf neighbor | inc VL0
172.16.50.1      0    FULL/  -           -           172.16.25.5    OSPF_VL0
R2#

R2#show ip route ospf | inc 172.16.50
O IA 172.16.50.0/24 [110/782] via 172.16.25.5, 00:58:13, Serial1/0
R2#
```

You see a similar virtual link output on 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".

## S.4 RIP



### HIDDEN ISSUES TO SPOT WITH THE RIP CONFIGURATION

**Issue: Restrict the advertisement of RIP updates to only VLAN17 and VLAN88 interfaces.**

**Solution:**

Configure routers R1 and FRS with **passive interface default**. Then, disable passive interface on the VLAN 17 and VLAN 88 interfaces.

**Issue: Configure RIP version 2 only between devices connected to VLAN's 17 and 88.**

**Solution:**

Under the desired interface, specify the appropriate RIP version for sending and receiving updates. The syntax for these interface commands are: **ip rip send version <1 2>** and **ip rip receive version <1 2 >**. Also, you can restrict the version for the rip routing process under RIP process using the **version 2** command, this method was chosen in this answer key

**Verification:**

Output on R1, the other RIP speakers have similar configuration:

```

R1#show ip protocols | begin Routing Protocol is "rip"
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 16 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: ospf 100, rip
  Default version control: send version 2, receive version 2
  Interface          Send  Recv  Triggered RIP  Key-chain
  FastEthernet0/0.17  2     2
  FastEthernet0/0.88  2     2
Automatic network summarization is in effect
Maximum path: 4
Routing for Networks:
  172.16.0.0
Passive Interface(s):
  FastEthernet0/0
  FastEthernet0/0.100
  Serial0/0
  Serial0/0.123
  BRI0/0
  BRI0/0:1
  BRI0/0:2
  ATM3/0
Passive Interface(s):
  Loopback10
  Loopback100
  VoIP-Null0
Routing Information Sources:
  Gateway          Distance    Last Update
  172.16.17.7      120        00:00:25
  172.16.88.2     120        00:00:15
Distance: (default is 120)

```

**Issue:** Provide reachability to 2.2.2.2 from FRS and CAT1.

**Solution:**

Redistributing OSPF into RIP will redistribute a 0.0.0.0/0 network into RIP. The default network will be propagated to FRS and CAT1. FRS and CAT1 will set the gateway of last resort pointing to R1. When the gateway of last resort is set on FRS and CAT1 towards R1, these routers will be able to reach the 2.2.2.2 address.

**Verification:**

```
FRS#show ip route | inc 0.0.0.0
Gateway of last resort is 172.16.17.1 to network 0.0.0.0
R* 0.0.0.0/0 [120/1] via 172.16.17.1, 00:00:27, Ethernet0
FRS#
```

```
CAT1#show ip route | inc 0.0.0.0
Gateway of last resort is 172.16.88.1 to network 0.0.0.0
R* 0.0.0.0/0 [120/1] via 172.16.88.1, 00:00:14, Vlan88
CAT1#
```



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

## S.5 EIGRP



### HIDDEN ISSUES TO SPOT WITH THE EIGRP CONFIGURATION

**Issue: Restrict the advertisement of EIGRP updates to only VLAN360 interfaces.**

**Solution:**

Configure routers R3, R6 and CAT2 with an EIGRP “network” statement that possesses wildcard mask. When you configure an EIGRP “network” statement without a wildcard mask, you must supply a classful prefix. By configuring an EIGRP “network” statement with a classful prefix, you inject all prefixes on that router that share the specified classful prefix into EIGRP. Therefore, when you want to restrict a limited set of prefixes to be injected into EIGRP, use the “network” statement with the wildcard mask to define only the prefixes you want to include in EIGRP. The other method is to use the passive-interface default along with the classful network statement and do no passive interface for the interface connected to VLAN360. The network with the wildcard method is used in this answer key.

**Issue: Redistribution between two EIGRP ASes is not automatic.**

**Solution:**

Mutual redistribution must be done on R6 between EIGRP AS1 and AS2.

**Issue: Solving the reachability issue on routers R4 and CAT2.**

**Solution:**

The scenario requires only one prefix to be advertised from R3 to R6 and therefore to CAT2. R6 and CAT2 must be able to ping the rest of the network. Solution: In the OSPF section, the use of the **default-information originate always** command on router R2 was explained. R2 injects a 0.0.0.0/0 into the OSPF network. The 0.0.0.0/0 route can be used to solve the reachability issue facing both R6 and CAT2. R3 should redistribute OSPF into EIGRP. Therefore, 0.0.0.0/0 will be redistributed into EIGRP. All other prefixes advertised from R3 to R6 should be filtered. Only 0.0.0.0/0 should be advertised from R3 to R6. Based upon the receipt of the 0.0.0.0/0 network, R6 will set up the gateway of last resort pointing to R3. Likewise, CAT2 will receive 0.0.0.0/0 from R6, if the redistribution between EIGRP AS1 and AS2 is done. Consequently, CAT2 will set up the gateway of last resort to R6.

**Verification:**

```
R6#show ip route | inc 0.0.0.0
Gateway of last resort is 172.16.36.3 to network 0.0.0.0
D*EX 0.0.0.0/0 [170/1734656] via 172.16.36.3, 01:10:15, FastEthernet0/0
R6#
```

```
CAT2#show ip route | inc 0.0.0.0
Gateway of last resort is 172.16.36.6 to network 0.0.0.0
D*EX 0.0.0.0/0 [170/1734912] via 172.16.36.6, 01:10:41, Vlan360
```

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".

## S.6 ODR



### HIDDEN ISSUES TO SPOT WITH THE ODR CONFIGURATION

**Issue: Configure on demand routing between R1 and R6.**

**Solution:**

ODR is a feature that provides IP routing for stub sites, with minimum overhead. The overhead of a general, dynamic routing protocol is avoided without incurring the configuration and management overhead of static routing. R4 will be a stub router and R1 will play a role of the hub router.

ODR uses the Cisco Discovery Protocol (CDP) to carry minimal routing information between the hub and stub routers. The stub routers send IP prefixes to the hub router via CDP. ODR supports VLSM ( Variable Subnet Length Mask).

Do not forget to redistribute ODR routing protocol into OSPF and RIP on R1 to propagate R4's network to RIP and OSPF domains..

**Issue: CDP packets exchange is not enabled by default on the physical frame relay serial interface.**

**Solution:**

Enable the CDP on the respective interfaces:

**Configuration and Verification:**

On R1:

```
router odr
network 172.16.0.0
```

```
R1#show run int s0/0 | inc cdp
  cdp enable
```

```
R1#show cdp neig | inc R4
R4          Ser 0/0          136          R          2621          Ser 0/0
R1#
```

```
R4#show run int s0/0 | inc cdp
  cdp enable
```

```
R4#show cdp neighbor | inc R1
R1          Ser 0/0          147          R S I          3640          Ser 0/0
```



R4#

R4#show ip ro

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2  
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  
ia - IS-IS inter area, \* - candidate default, U - per-user static route  
o - ODR, P - periodic downloaded static route

Gateway of last resort is 172.16.14.1 to network 0.0.0.0

172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks

C 172.16.40.0/25 is directly connected, Loopback40

C 172.16.14.0/24 is directly connected, Serial0/0

O\* 0.0.0.0/0 [160/1] via 172.16.14.1, 00:00:45, Serial0/0

R4#



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

## S.7 BGP



### HIDDEN ISSUES TO SPOT WITH THE BGP CONFIGURATION

**Issue:** Advertise the following prefixes from FRS with the origin code "incomplete".

**Solution:**

Do not use the "network command" to originate these prefixes in BGP. Use redistribution to originate the prefixes. When prefixes are originated in BGP via redistribution, their origin code is set to "incomplete".

**Issue:** Do not form a BGP peer relationship between R2 and R4. Use the AS numbers given in the exam.

**Solution:**

Routers R2, R3 and R4 are I-BGP speakers within the same Autonomous System. By default, a full mesh of I-BGP neighbor relationships must be formed between I-BGP speakers. However, since you are instructed not to form a BGP peer relationship between R2 and R4, a full mesh cannot be formed. The remedy for this non full-mesh requirement is to configure a route reflector on router R3.

**Issue:** All BGP speakers in AS 64600 have networks 192.168.104.0/24, 192.168.105.0/24 and a summary for the other 192.\*.\* networks in both their BGP and IP routing tables...Do not change the AS-PATH configuration.

**Solution:**

Configure the BGP aggregate command with the following parameters:

```
aggregate-address 192.168.100.0 255.255.252.0 as-set summary-only
```

The mask of the aggregate covers only 192.168.100.0 – 192.168.103.0. Therefore, the specified additional subnets of 192.168.104.0/24 and 192.168.105.0/24 can also be advertised. By default, the longer matching subnets of an aggregate are advertised with the aggregate. This behavior is suppressed with the **summary-only** keyword. The second requirement – “Do not change the AS-PATH configuration” – can be fulfilled by including the **as-set** option in the aggregate statement.

**Issue:** Use the synchronization method on R2 and R3.

**Solution:**

Redistribute EBGP-learned updates on router R2 into OSPF. This will fulfill the synchronization requirements for all IBGP learned routes received by routers R2 and R3. Synchronization must be disabled on router R4. Router R4 will never receive 192.168.\*.\* prefixes via its IGP because of restrictions in the EIGRP section. Therefore, the only way to assign the best path to the networks 192.168.\*.\* on R4 is to disable synchronization.

**Issue:** All BGP speakers should have only a classful prefix of the IP address assigned to R3's loopback 10 interface in their BGP tables. Do not redistribute connected.

**Solution:**

Originate the 3.0.0.0/8 network on router R3 into BGP using a network statement WITHOUT a mask. This will generate/originate a classful prefix in BGP.

**Issue:** Network 3.0.0.0/8 does not show up in the bgp table of R1.

**Solution:**

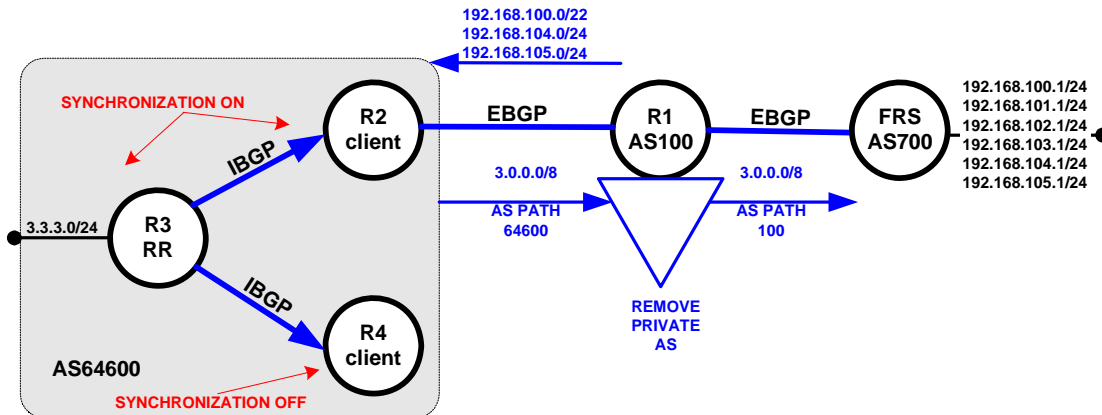
This is because the prefix is not synchronized on R2. You advertise 3.3.3.0/24 via OSPF and 3.0.0.0/8 via BGP. They do not match on R2. Solution: summarize 3.3.3.0/24 to 3.0.0.0/0 on R3 by using OSPF area range command.

**Issue:** On FRS, this major network should be shown as originated from AS 100.

**Solution:**

Apply the **remove private-as** command to the neighbor relationship between R1 and FRS. As the following diagram illustrates, the 3.0.0.0/8 prefix will be received by R1 in AS 100 with an AS path of “64600”. This “64600” AS will be stripped off when R1 advertises the 3.0.0.0/8 prefix to FRS. FRS will

receive the 3.0.0.0/8 prefix as if it was originated by AS 100. Therefore, AS 64600 is completely hidden from FRS.



**Verification:**

Verify BGP peer relationship and the BGP table, for more show commands check the SHOWit engine:

```
FRS#sh ip bgp summary | begin State/PfxRcd
Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.17.1   4    100  5134    5133    8       0     0  3d13h   1
FRS#
```

```
FRS#sh ip bgp | begin Weight Path
Network      Next Hop      Metric LocPrf Weight Path
*> 3.0.0.0    172.16.17.1  0      100    i
*> 192.168.100.0  0.0.0.0      0      32768 ?
*> 192.168.101.0  0.0.0.0      0      32768 ?
*> 192.168.102.0  0.0.0.0      0      32768 ?
*> 192.168.103.0  0.0.0.0      0      32768 ?
*> 192.168.104.0  0.0.0.0      0      32768 ?
*> 192.168.105.0  0.0.0.0      0      32768 ?
FRS#
```

```
R1#sh ip bgp summary | begin State/PfxRcd
Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.17.7   4    700  5134    5135    13      0     0  3d13h   6
172.16.123.2  4   64600  5134    5135    13      0     0  3d13h   1
FE80::3333:3333 4  64600    99      99      0     0     0  01:33:09 (NoNeg)
R1#
```

```
R1#sh ip bgp | begin Weight Path
Network      Next Hop      Metric LocPrf Weight Path
*> 3.0.0.0    172.16.123.3  0      64600 i
s> 192.168.100.0  172.16.17.7  0      700 ?
*> 192.168.100.0/22 0.0.0.0      100    32768 700 ?
s> 192.168.101.0  172.16.17.7  0      700 ?
s> 192.168.102.0  172.16.17.7  0      700 ?
```

```
s> 192.168.103.0 172.16.17.7 0 0 700 ?
*> 192.168.104.0 172.16.17.7 0 0 700 ?
*> 192.168.105.0 172.16.17.7 0 0 700 ?
```

R1#

```
R2#sh ip bgp summary | begin State/PfxRcd
Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.123.1  4      100   5138   5137    6      0    0 3d13h    3
172.16.123.3  4 64600   5138   5140    6      0    0 3d13h    1
```

R2#

```
R2#sh ip bgp | begin Weight Path
Network      Next Hop      Metric LocPrf Weight Path
r>i3.0.0.0   172.16.123.3  0      100    0  i
*> 192.168.100.0/22 172.16.123.1  0      0      0 100 700 ?
*> 192.168.104.0   172.16.123.1  0      0      0 100 700 ?
*> 192.168.105.0   172.16.123.1  0      0      0 100 700 ?
```

R2#

```
R3#sh ip bgp summary | begin State/PfxRcd
Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.34.4   4 64600   5132   5135    8      0    0 3d13h    0
172.16.123.2  4 64600   5136   5135    8      0    0 3d13h    3
```

R3#

```
R3#sh ip bgp | begin Weight Path
Network      Next Hop      Metric LocPrf Weight Path
*> 3.0.0.0     0.0.0.0      0      0      32768 i
r>i192.168.100.0/22 172.16.123.1  0      100    0 100 700 ?
r>i192.168.104.0   172.16.123.1  0      100    0 100 700 ?
r>i192.168.105.0   172.16.123.1  0      100    0 100 700 ?
```

R3#

```
R6#sh ip bgp summary | begin State/PfxRcd
Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.36.3   4 64600   5135   5132    5      0    0 3d13h    4
```

R6#

```
R6#sh ip bgp | begin Weight Path
Network      Next Hop      Metric LocPrf Weight Path
*>i3.0.0.0     172.16.36.3  0      100    0  i
*>i192.168.100.0/22 172.16.123.1  0      100    0 100 700 ?
*>i192.168.104.0   172.16.123.1  0      100    0 100 700 ?
*>i192.168.105.0   172.16.123.1  0      100    0 100 700 ?
```

R6#



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

## S.8 IPv6 Routing



### HIDDEN ISSUES TO SPOT WITH THE IPv6 ROUTING CONFIGURATION

**Issue:** Assign all the Site Local and Link Local IPv6 addresses according to the Sample Scenario IPv6 IGP diagram. Make sure all the IPV6 addresses are reachable within the subnet, including the locally assigned addresses.

**Solution:**

1. Enable the ipv6 unicast-routing in global configuration mode.
2. Under the specified on the diagram interfaces configure IPv6 addresses. The following shows the configuration on the Frame Relay interface on R2 as an example. Configure all other interfaces similarly:

```
interface Serial0/0
  ipv6 address FEC0::BBBB:2/120
  ipv6 address FE80::2222:2222 link-local
```

3. Make sure you have a within the subnet reachability. This issue is very similar to IPV4 subnet reachability over the Frame Relay. Statically map the IPv6 addresses including the local ones:

```
interface Serial0/0
  ipv6 address FEC0::BBBB:2/120
  ipv6 address FE80::2222:2222 link-local
  frame-relay map ipv6 FE80::1111:1111 201 broadcast
  frame-relay map ipv6 FE80::2222:2222 201 broadcast
  frame-relay map ipv6 FE80::3333:3333 201 broadcast
  frame-relay map ipv6 FEC0::BBBB:1 201 broadcast
  frame-relay map ipv6 FEC0::BBBB:2 201 broadcast
  frame-relay map ipv6 FEC0::BBBB:3 201 broadcast
```

### OSPF IPv6

**Issue:** Configure OSPF IPv6 area 0 on the Frame Relay link between R1 and R2. Use the default OSPF network type. R1 should initiate the HELLO packet. R2 should never initiate the HELLO packets.

**Solution:**

The default OSPF IPv6 network type for the physical and logical multipoint interfaces and the encapsulation Frame Relay is non-broadcast. The "non-broadcast" network type instructs the OSPF process to elect the DR and BDR on the link. Configure the OSPF IPv6 priority 0 on the serial frame relay interface of R2 to stop the HELLO initiation on this interface. R1 will initiate the HELLO packets with the neighbor command. Make sure you use the Link Local IPv6 address for the neighbor command.

### Configuration and Verification:

On R1:

```
int s0/0.123
ipv6 ospf neighbor FE80::2222:2222
ipv6 ospf 1 area 0
!
ipv6 router ospf 1
```

Note! You must use the link local addresses for the OSPF neighbor statements.

```
R1#show ipv6 ospf int s0/0.123
Serial0/0.123 is up, line protocol is up
  Link Local Address FE80::1111:1111, Interface ID 17
  Area 0, Process ID 1, Instance ID 0, Router ID 172.16.10.1
  Network Type NON_BROADCAST, Cost: 64
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.10.1, local address FE80::1111:1111
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    Hello due in 00:00:15
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 7, maximum is 7
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.20.5
  Suppress hello for 0 neighbor(s)
R1#
```

```
R1#show ipv6 ospf neighbor detail
Neighbor 172.16.20.5
  In the area 0 via interface Serial0/0.123
  Neighbor: interface-id 4, link-local address FE80::2222:2222
  Neighbor priority is 0 (configured 0), State is FULL, 10 state changes
  DR is 172.16.10.1 BDR is 0.0.0.0
  Poll interval 120
  Options is 0x65EEBFA5
  Dead timer due in 00:01:46
  Neighbor is up for 01:43:27
  Index 1/1/1, retransmission queue length 0, number of retransmission 1
  First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
  Last retransmission scan length is 2, maximum is 2
  Last retransmission scan time is 0 msec, maximum is 0 msec
R1#
```

On R2:

```
Interface serial0/0
ipv6 ospf priority 0
ipv6 ospf 1 area 0
```

## RIP IPv6

**Issue:** Configure RIP IPv6 instance SAMPLE on the link between the routers R2 and R5. Advertise the loopback 500 network into RIP IPV6 process. This RIP IPv6 process should be identified as RIPIPV6. Configure RIP IPv6 instance SAMPLE on the link between the routers R3 and R6. Advertise the loopback 600 on the router R6 without enabling RIP on the interface.

### Solution:

When you configure RIP IPv6, make sure you create two instances and label them as SAMPLE and RIPIPV6 on R5.

### Configuration and Verification:

Router R5 is used as an example. R2, R3 and R6 should be configured similarly for the IPv6 RIP instance SAMPLE. Use the SHOWit engine to view all related configurations and IOS "show" command output.

#### R5:

```

interface Loopback500
  no ip address
  ipv6 address FEC0::5555:1/125
  ipv6 rip RIPIPV6 enable
!
interface Serial1/0
  ip address 172.16.25.5 255.255.255.0
  ipv6 address FEC0::AAAA:5/120
  ipv6 address FE80::2525:5555 link-local
  ipv6 rip SAMPLE enable
!
ipv6 router rip SAMPLE
  redistribute connected metric 1
  redistribute rip RIPIPV6 metric 1
!
ipv6 router rip RIPIPV6
!
R5#show ipv6 rip database
RIP process "SAMPLE", local RIB
  FEC0::1111:0/125, metric 2, installed
    Serial1/0/FE80::2525:2222, expires in 161 secs
  FEC0::6666:0/125, metric 2, installed
    Serial1/0/FE80::2525:2222, expires in 161 secs
  FEC0::AAAA:0/120, metric 2
    Serial1/0/FE80::2525:2222, expires in 161 secs
  FEC0::BBBB:0/120, metric 2, installed
    Serial1/0/FE80::2525:2222, expires in 161 secs
  FEC0::D3D3:0/120, metric 2, installed
    Serial1/0/FE80::2525:2222, expires in 161 secs
RIP process "RIPIPV6", local RIB

```

```

R5#show ipv6 route rip
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, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   FEC0::1111:0/125 [120/2]
    via FE80::2525:2222, Serial1/0
R   FEC0::6666:0/125 [120/2]
    via FE80::2525:2222, Serial1/0
R   FEC0::BBBB:0/120 [120/2]
    via FE80::2525:2222, Serial1/0
R   FEC0::D3D3:0/120 [120/2]
    via FE80::2525:2222, Serial1/0
R5#

```

## BGP IPv6

**Issue:** *Configure IPv6 BGP peer relationship between router R3 AS 64600 and R1 AS 100. You must use the link local IPv6 address to form this peer relationship.*

### Solution:

When peering between the link local IPv6 address make sure you update the source and the next hop attribute

### Configuration and Verification:

Router R1 is used as an example. R3 should be configured similarly. Check the SHOWIT engine for more details.

```

interface Loopback100
  no ip address
  ipv6 address FEC0::1111:1/125
!

router bgp 100
  bgp log-neighbor-changes
  neighbor 172.16.17.7 remote-as 700
  neighbor 172.16.123.2 remote-as 64600
  neighbor FE80::3333:3333 remote-as 64600
  neighbor FE80::3333:3333 ebgp-multihop 255
  neighbor FE80::3333:3333 update-source Serial0/0.123
  !
  !
  address-family ipv6
  neighbor FE80::3333:3333 activate
  neighbor FE80::3333:3333 route-map NH out
  network FEC0::1111:0/125
  redistribute ospf 1 match internal external 1 external 2
  no synchronization
  exit-address-family

```



```
!
!
route-map NH permit 10
set ipv6 next-hop FEC0::BBBB:1
!
```

Verify the peer relationship and updates:

```
R1#show bgp ipv6 unicast summary
BGP router identifier 172.16.10.1, local AS number 100
BGP table version is 7, main routing table version 7
6 network entries using 894 bytes of memory
6 path entries using 456 bytes of memory
8/4 BGP path/bestpath attribute entries using 992 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2390 total bytes of memory
BGP activity 14/0 prefixes, 14/0 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
FE80::3333:3333	4	64600	143	143	7	0	0	02:17:17	3

```
R1#show bgp ipv6 unicast
BGP table version is 7, local router ID is 172.16.10.1
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
*> FEC0::1111:0/125	::	0		32768	i
*> FEC0::5555:0/125	::	20		32768	?
*> FEC0::6666:0/125	FEC0::BBBB:3	2		0	64600 ?
*> FEC0::AAAA:0/120	::	20		32768	?
*> FEC0::BBBB:0/120	FEC0::BBBB:3	0		0	64600 ?
*> FEC0::D3D3:0/120	FEC0::BBBB:3	0		0	64600 ?

R1#

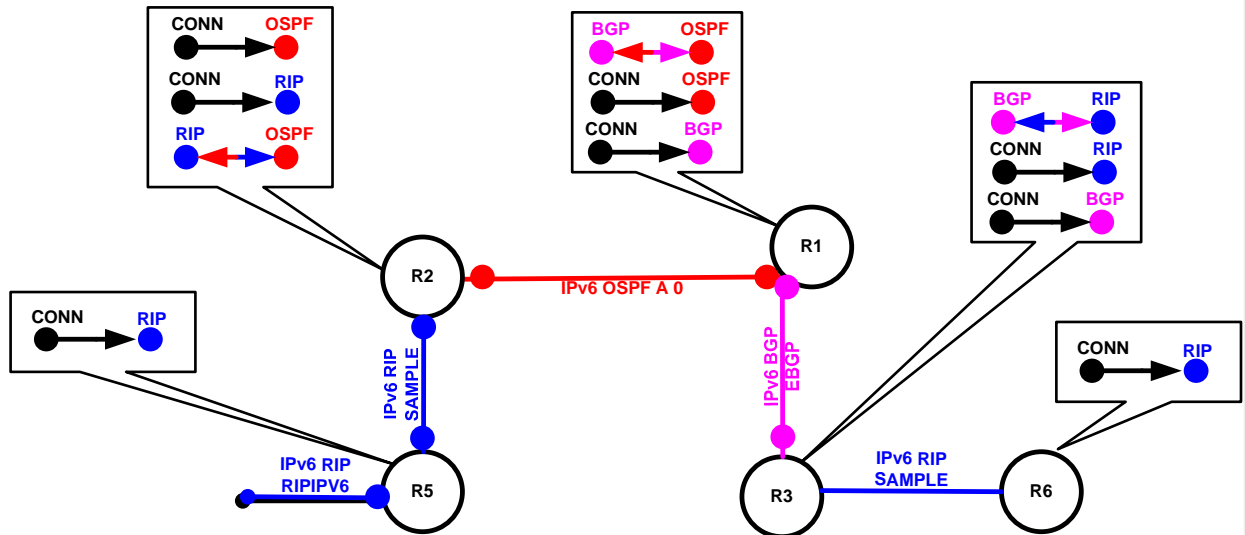
**Issue: IPv6 Redistribution**



**TO OBTAIN UNIVERSAL CONNECTIVITY FOR IPV6 GLOBAL ADDRESSES, PERFORM A MUTUAL REDISTRIBUTION OF DYNAMIC IPV6 ROUTING PROTOCOLS: BETWEEN RIP AND OSPF ON ROUTER R2, BETWEEN OSPF AND BGP ON ROUTER R1, BETWEEN BGP AND RIP ON ROUTER R3. DO NOT PERFORM ANY OTHER REDISTRIBUTION IN THIS SCENARIO. PERFORM REDISTRIBUTE CONNECTED WHERE REQUIRED AND NOT RESTRICTED BY THE SCENARIO.**

**Solution:**

Perform mutual redistribution between OSPF IPv6 and RIP IPv6 on router R2. Perform mutual redistribution between OSPF IPv6 and BGP IPv6 on router R1. Perform mutual redistribution between RIP IPv6 and BGP IPv6 on R3. Make sure you redistribute the connected networks into the dynamic routing protocols as well. The IPv6 routing protocols will not take connected networks from the dynamic IPv6 routing protocols. Please see the following diagram for the redistribution strategy:



**Configuration and Verification:**

Configuration on R1 is shown as an example. Check the SHOWIT for the configuration on the other redistributing routers:

R1:

```

address-family ipv6
 neighbor FE80::3333:3333 activate
 neighbor FE80::3333:3333 route-map NH out
 network FEC0::1111:0/125
 redistribute connected
 redistribute ospf 1 match internal external 1 external 2
 no synchronization
 exit-address-family
!
ipv6 router ospf 1
 log-adjacency-changes
 redistribute connected
 redistribute bgp 100
!

```

## Reachability Verification

One way to test that your redistribution configuration satisfies the goal of universal connectivity is to run a TCL script like the one displayed below on each router. TCL scripting support is available in the IOS versions used on routers R1, R2, R3, R4, R5 and R6 (the 3600 and 2600 models.) The simple script below lists all of the IP addresses in a pod configured for this Scenario. It can be built once in notepad, and then pasted into each router to automate “ping” testing process. There is a paper on TCL scripting available in the READiT section of the Netmasterclass website. Some addresses are used in later tasks and may not be reachable at this point. Run **tclsh** in privileged mode, paste the script below, and then issue the command **tclq**.

### IPv4 and IPv6 script

```
foreach address {  
172.16.10.1  
172.16.17.1  
172.16.88.1  
172.16.123.1  
172.16.14.1  
FEC0::BBBB:1  
FEC0::1111:1  
172.16.20.5  
172.16.123.2  
172.16.2.1  
172.16.25.2  
2.2.2.2  
FEC0::BBBB:2  
FEC0::AAAA:2  
172.16.123.3  
172.16.36.3  
172.16.30.3  
3.3.3.3  
FEC0::BBBB:3  
FEC0::D3D3:3  
172.16.14.4  
172.16.40.1  
172.16.25.5  
172.16.50.1  
FEC0::AAAA:5  
FEC0::5555:1  
172.16.60.1  
172.16.36.6  
FEC0::6666:1  
FEC0::D3D3:6  
172.16.17.7  
172.16.77.7  
192.168.100.1  
192.168.101.1
```








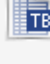





```

192.168.102.1
192.168.103.1
192.168.104.1
192.168.105.1
172.16.80.1
172.16.88.2
172.16.36.20
172.16.120.1
} {
ping $address
}

```

You can generate the TCL ping script for each lab from the “Connectivity check results” Table. Please look at the DOIT Lab3 example.

**DOIT vol.II Lab-3 info**

<p><b>Level of Difficulty: Moderate</b>  <b>Total Number of Tasks/Sections in this Scenario: 13</b>  <b>Key Features of this Scenario:</b>            Frame-Relay interface addressing challenge!            Catalyst trunking, routed ports, QoS            IGP (OSPF, RIP v2, EIGRP, redistribution)            Challenging BGP with synchronization and preferred hops            Securing HTTP access            IPv6 (OSPF, RIP, BGP)            MHSRP            Multicast</p> <p><b>Some Observations to Remember Before Getting to the “Golden Moment”:</b> Yes, you can put addresses from the same subnet on different interfaces!  <b>Number of loops in the topology:</b> loops within loops!</p>	<p> DOIT vol.II Lab-3 Scenario</p> <p> DOIT vol.II Lab-3 Answer Key</p> <p> Initial Configuration Script</p> <p> SHOWIT II</p> <p> Forum II</p>	<p> IP protocol table</p> <p> BGP peer Relationship table</p> <p> BGP prefix table</p> <p> Gateway of last resort</p> <p> EIGRP neighbor table</p> <p> CDP neighbor table</p> <p> Suggested list of commands</p> <p> Connectivity check results</p>
---	---	--

Click on the Connectivity check results, the browser will open the Connectivity Report:



N/A

### Connectivity Report for DOIT Lab 3 generated on NMC POD 4

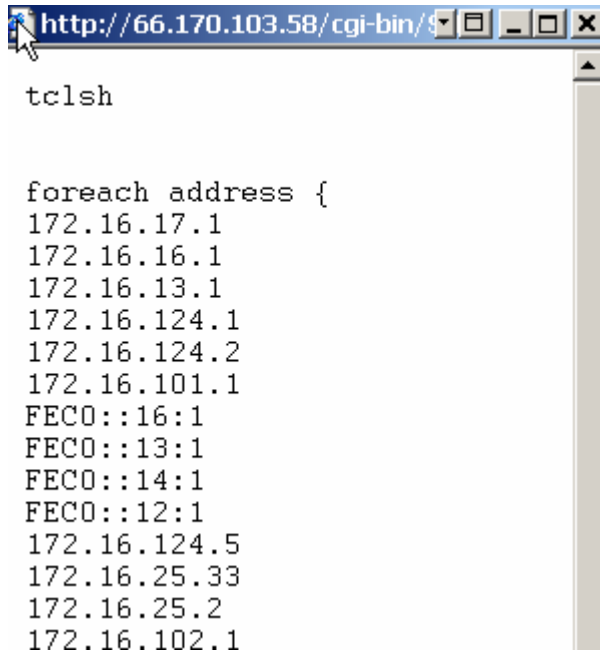
Based on file generated:  
10/18/05 12:27

To see the connectivity status for your POD you may use MultiPing TCL Script.

[TCL Script](#)

Destination		Source								
Device	IP address	R1	R2	R3	R4	R5	R6	FRS	CAT1	CAT2
R1	172.16.17.1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!
R1	172.16.16.1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!
R1	172.16.13.1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!
R1	172.16.124.1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!
R1	172.16.124.2	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!
R1	172.16.101.1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!
R1	FEC0::16:1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!			
R1	FEC0::13:1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!			
R1	FEC0::14:1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!			
R1	FEC0::12:1	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!			
R2	172.16.124.5	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!	!!!!

Click on the "TCL Script" button, the browser will open the text window with the TCL script in it:



```

tclsh

foreach address {
172.16.17.1
172.16.16.1
172.16.13.1
172.16.124.1
172.16.124.2
172.16.101.1
FEC0::16:1
FEC0::13:1
FEC0::14:1
FEC0::12:1
172.16.124.5
172.16.25.33
172.16.25.2
172.16.102.1

```

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**. Also, for more details on TCL and RSH, see the Netmasterclass article "**CISCO IOS TCL and RCMD testing and troubleshooting scripting**" available at <http://www.netmasterclass.net/site/lib.php>



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".

## S.9 QOS



### HIDDEN ISSUES TO SPOT WITH THE QOS CONFIGURATION

**Issue:** Limit only incoming UDP traffic destined to port 512X to an 8000 bit/sec rate on the FastEthernet interface of router R1. Configure the minimal values for burst size and extended burst size. Drop excessive traffic. Use the MQC.

### **Solution:**

Policing can be done on IOS routers in two ways: (1) configure a committed access-rate (CAR) command on the R1 FastEthernet interface or (2) perform a Modular Quality of Service CLI (MQC) configuration with a "police" statement on router R1 and apply it to its FastEthernet interface with an inbound service-policy command. Both solutions will perform the inbound policing function that is outlined in the configuration task. The scenario explicitly specifies the MQC method. It must be a policing function since the task is directing you to drop "excessive traffic" that exceeds the specified 8000 bit/sec rate. This cannot be interpreted as a traffic-shaping task for two reasons: (1) traffic shaping is performed on outbound traffic only. The traffic specified in this task is inbound; (2) you are directed to drop excessive traffic. Traffic shaping will attempt to queue excess traffic.

### **Configuration and Verification:**

On R1:

1. Classification of the traffic. You needed to create an extended access-list to match on the specified UDP traffic.

```
access-list 101 permit udp host 10.1.1.1 host 10.1.1.124 eq 5124
```

2. Create a class-map to match the traffic specified in the access-list.

```
class-map match-all UDP-STREAM  
match access-group 101
```

3. Create a policy-map to police the traffic according to scenario specifications

```
policy-map UDP-TRAFFIC  
class UDP-STREAM  
police cir 8000 bc 1500 be 1500  
conform-action transmit  
exceed-action drop
```

4. Apply the policy on the interface connected to VLAN 100 where the traffic generator is.

```
interface FastEthernet0/0.100  
encapsulation dot1q 100  
ip address 10.1.1.124 255.255.255.0  
service-policy input UDP-TRAFFIC
```

5. Verify the actions of your policing

```
R1#show policy-map int fa 0/0.100  
FastEthernet0/0.100
```

```
Service-policy input: UDP-TRAFFIC
```

```
Class-map: UDP-STREAM (match-all)  
708674 packets, 758281180 bytes  
5 minute offered rate 42000 bps, drop rate 34000 bps
```

```
Match: access-group 101
police:
  cir 8000 bps, bc 1500 bytes
  conformed 132466 packets, 141738620 bytes; actions:
  transmit
  exceeded 576208 packets, 616542560 bytes; actions:
  drop
  conformed 8000 bps, exceed 34000 bps
```

```
Class-map: class-default (match-any)
  708696 packets, 758295312 bytes
  5 minute offered rate 42000 bps, drop rate 0 bps
```

Match: any



*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".*



## S.10 Address Administration



### **HIDDEN ISSUES TO SPOT WITH THE ADDRESS ADMINISTRATION CONFIGURATION**

**Issue:** *The lowest 10 ip addresses will be used for routers, servers and printers.*

**Solution:**

Configure an **ip dhcp excluded-address** command for these addresses.

**Issue:** *Specific workstations with supplied MAC addresses should always receive the same IP address.*

**Solution:**

Create separate DHCP pools for each supplied MAC address and configure the corresponding IP address to each separate pool.

**Issue:** *Supply the appropriate gateway ip address.*

**Solution:**

This task is tied to the HSRP configuration in the next section. In order to fulfill this task, you must read ahead and determine the virtual IP address used by HSRP.



**REMEMBER, ALWAYS READ A CCIE LAB EXAM END-TO-END AND CAREFULLY LOOK FOR HIDDEN ISSUES THAT MIGHT INVOLVE MULTIPLE TASKS.**



**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”.**

## S.11 Gateway Redundancy



### HIDDEN ISSUES TO SPOT WITH THE GATEWAY REDUNDANCY CONFIGURATION

**Issue:** *Prefer R3 when the Frame-Relay connection becomes active.*

**Solution:**

Configure the HSRP **preempt** command to allow R3 to regain the active standby status when its Frame-Relay link becomes active again. The HSRP virtual IP address will be used in DHCP pools that were mentioned in the Address Administration section.

**Issue:** *If the Frame Relay connection fails, prefer R6.*

**Solution:**

Configure HSRP with the “track” option so that when the Frame-Relay interface fails, R3 will no longer be the preferred gateway. The “track” option will allow you to decrement the locally assigned HSRP priority value so that a given HSRP peer will have a lower priority when the “tracked” interface fails. In this Scenario, we want the assigned priority to R3 to be decremented when the Frame-Relay connection fails. To explicitly assign a value to be decremented from an HSRP peer’s assigned priority, specify the amount to be decremented in the following command: **standby 1 track Serial0/0 XXX** where XXX is the value to decrement from the locally assigned priority. By default, the track option will decrement the locally assigned HSRP priority by 10.

**Issue:** *Assign the lowest IP address on VLAN34 to the virtual gateway.*

**Solution:**

From the DHCP configuration in this Scenario’s Address Administration section, select the lowest IP address from the DHCP excluded-address pool.



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”.

## S.12 NTP Configuration



### HIDDEN ISSUES TO SPOT WITH THE NTP CONFIGURATION

**Issue:** *Make R1 the NTP master with stratum 5.*

**Solution:**

Configure the NTP “master” command on router R1 setting the stratum level to 5.

**Issue: Configure a server association between R3 and R1.****Solution:**

To fulfill this requirement, use the NTP “server” command on router R3.

**Issue: Configure a peer association between R3 and R6.****Solution:**

To fulfill this requirement, use the NTP “peer” command on router R6 referencing router R3. No “ntp peer” command needs to be added on router R3. By configuring NTP in this manner, R6 will become an NTP “symmetric active” peer and R3 will become an NTP “symmetric passive” peer. For more details on the difference between NTP symmetric active and passive peers, see RFC 1305, Section 3.3 (Modes of Operation).



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”.

## S.13 Multicast



### HIDDEN ISSUES TO SPOT WITH THE MULTICAST CONFIGURATION

**Issue: What multicast routing protocol is based on “Flood and Prune” technology?****Solution:**

The PIM Dense Mode multicast routing protocol uses “flood and prune” technology. With PIM Dense Mode, the multicast distribution tree is built by the multicast traffic itself. The source of the tree is the source of the IP Multicast traffic. Each router receiving multicast datagrams from an upstream node will check if the datagram arrived via the appropriate incoming interface. This operation is called the Reverse Path Forwarding Check or “RPF” check. If the received traffic passes the RPF check, the PIM Dense Mode enabled router will forward the datagrams out to the downstream nodes. This process is known as the “flooding” process. The nodes not interested in receiving the multicast traffic will send a PRUNE message to upstream, multicast routers to stop the multicast traffic from being sent downstream via that interface. The multicast tree will end up having only branches that possess active receivers at the leaf nodes of the tree.

**Issue: Configure multicast routing between R1, R2 and R3.**

**Solution:**

Simply enable **ip multicast-routing** on routers R1, R2 and R3 in global configuration mode and configure **-mode** on the interfaces connected to the 172.16.123.0/24 network.

**Issue. Configure R2 to start processing multicast packets arriving from R5.**

**Solution:**

You need to configure **ip pim dense-mode** on the Serial 1/0 interface of R2 as well.

**Issue: Simulate an active receiver of traffic destined to 230.30.30.30 using the loopback interfaces.**

**Solution:**

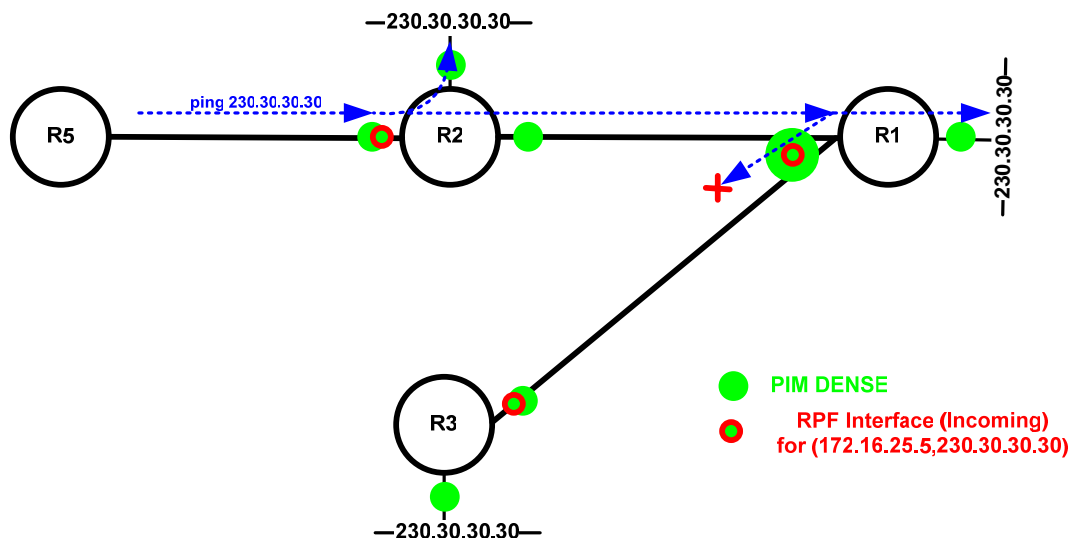
Configure the command **ip igmp join-group 230.30.30.30** under the loopback interfaces. Also, it is recommended to configure **ip pim dense** under the loopback interface as well.

**Issue I'm getting replies from R1 but not from R3.**

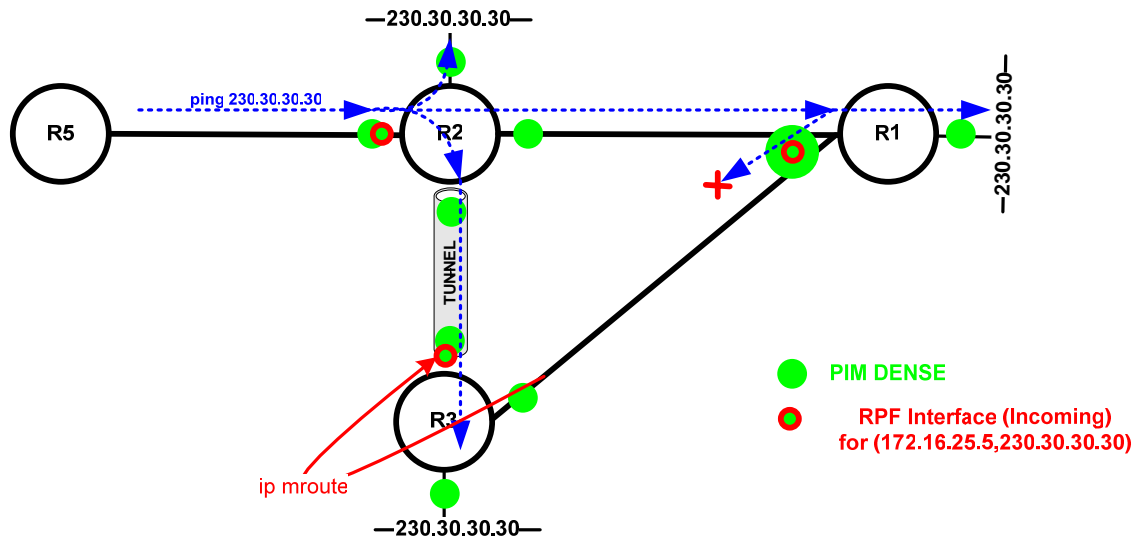
**Solution:**

There are two reasons for this behavior:

**Reason #1)** R1 is receiving traffic to 230.30.30.30 via Serial0/0.123. R1 will not forward traffic back out the same interface to router R3. For any given stream of multicast traffic, the incoming interface cannot also be an outgoing interface at the same time. See the following diagram:



**Solution to Reason #1):** Build a tunnel between R2 and R3 so that R2 will be able to send multicast traffic to R3 by performing encapsulation of multicast datagrams into unicast datagrams, which will be delivered to R3 via a unicast routing path.



**Reason #2):** R3's RPF check against the source of the traffic (172.16.25.5) needs to be adjusted to prefer the tunnel interface as the selected incoming interface.

**Solution to Reason #2):** Configure an **ip mroute** command to force the RPF lookup process to select the tunnel interface as the incoming interface for the specified multicast traffic.

**Verification:**

Check the PIM neighbor relationship:

```

R2#sh ip pim neighbor
PIM Neighbor Table
Neighbor      Interface      Uptime/Expires  Ver  DR
Address
172.16.123.1  Serial0/0      2d16h/00:01:39  v2   1 / S
172.16.30.3   Tunnel203      2d16h/00:01:31  v2   1 / S
R2#
  
```

```

R1#sh ip pim neighbor
PIM Neighbor Table
Neighbor      Interface      Uptime/Expires  Ver  DR
Address
172.16.123.3  Serial0/0.123  2d16h/00:01:34  v2   1 / DR S
172.16.123.2  Serial0/0.123  2d16h/00:01:19  v2   1 / S
R1#
  
```

```

R3#sh ip pim neighbor
PIM Neighbor Table
  
```

Neighbor Address	Interface	Uptime/Expires	Ver	DR Prio/Mode
172.16.123.1	Serial0/0	2d16h/00:01:26	v2	1 / S
172.16.25.2	Tunnel302	2d16h/00:01:38	v2	1 / S

R3#

Send the traffic to multicast group 230.30.30.30 form R5:

R5#ping 230.30.30.30

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 230.30.30.30, timeout is 2 seconds:

```
Reply to request 0 from 172.16.25.2, 32 ms
Reply to request 0 from 172.16.123.1, 220 ms
Reply to request 0 from 172.16.123.1, 208 ms
Reply to request 0 from 172.16.123.1, 192 ms
Reply to request 0 from 172.16.30.3, 180 ms
Reply to request 0 from 172.16.123.1, 108 ms
Reply to request 0 from 172.16.123.1, 96 ms
Reply to request 0 from 172.16.123.1, 80 ms
R5#
```

Check the (\*,G) and (S,G) for the 230.30.30.30 group in the mroute table:

```
R3#sh ip mroute 230.30.30.30
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

(*, 230.30.30.30), 2d17h/stopped, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel302, Forward/Dense, 2d17h/00:00:00
    Serial0/0, Forward/Dense, 2d17h/00:00:00
    Loopback31, Forward/Dense, 2d17h/00:00:00

(172.16.25.5, 230.30.30.30), 00:01:01/00:02:07, flags: LT
  Incoming interface: Tunnel302, RPF nbr 172.16.25.2, Mroute
  Outgoing interface list:
    Loopback31, Forward/Dense, 00:01:02/00:00:00
    Serial0/0, Forward/Dense, 00:01:02/00:00:00
```

## S.14 IOS Features



## **HIDDEN ISSUES TO SPOT WITH THE IOS FEATURES CONFIGURATION**

**Issue: Configure R6 to allow the router configuration and monitoring from the WEB browser**

**Solution:**

An HTTP WEB server can be activated on a CISCO router. For security reasons, it is generally not recommended for a production environment, but the feature does exist on Cisco routers. Use the command **ip http server** in global configuration mode to activate the HTTP server on R6.

**Issue: A WEB browser should be able to open a session with the router using the following url:**  
<http://R6-IP-ADDRESS:8090>.

**Solution:**

The default HTTP port is port 80. You can change the default port to port 8090 using the command **ip http port 8090**

**Issue: Restrict web access to router R6 to only the administrators residing on VLAN360 and limited to 2 sessions at a time .**

**Solution:**

Use the following combination of commands:

```
access-list 34 permit 172.16.34.0 0.0.0.255
ip http access-class 34
ip http max-connections 2
```

**Issue: Administrators will use username "admin" and password "adminnmc". This will be locally configured on router R6.**

**Solution:**

First, configure a user name on router R6. Second, specify the http authentication method to use the "local" username configuration on router R6. This can be performed with the following commands:

```
username admin password adminnmc
username admin privilege 15
ip http authentication local
```

**Issue: On the router R5 provide a solution allowing administrators to track any configuration change made to the IOS software running configuration and identify the user that made that change.**

**Solution:**

The Configuration Change Notification and Logging (Configuration Logging) feature was introduced in 12.3(4)T IOS. It allows the tracking of configuration changes entered on a per-session and per-user basis by implementing a configuration log. The configuration log will track each configuration command that is applied, who applied the command, the parser return code for that command, and the time that the command was applied.

**Issue: Configure maximum 300 entries retained in the configuration file**

**Solution:**

Configure logging size. See the configuration and verification section below.

**Issue: All the password information must be suppressed in the configuration log files**

**Solution:**

Configure hidekey. See the configuration and verification section below.

**Configuration and verification**

On R5 configure the following:

```

archive
 log config
  logging enable
  logging size 300
 hidekeys
  
```

Here are the verification steps:

1. Clear log and make a configuration change, for example configure a banner:

```

R5#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R5(config)#archive
R5(config-archive)#log config
R5(config-archive-log-cfg)#logging size 1
R5(config-archive-log-cfg)#logging size 300
R5(config-archive-log-cfg)#exit
R5(config-archive)#exit
R5(config)#banner motd "We are testing this new feature"
R5(config)#end
  
```

2. Verify the log:

```

R5#show archive log config all
idx  sess          user@line      Logged command
 69   53           console@console | logging size 1
  
```



```
70 53 console@console logging size 300
71 53 console@console exit
72 53 console@console exit
73 53 console@console banner motd "We are testing this new feature"
```

R5#

As you see the changes are logged for the user console@console. Do not forget to undo the banner it is not part of this scenario ☺ .

Please read more about this new feature at:

[http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t\\_4/gtconlog.htm](http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_4/gtconlog.htm)



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