NetMasterClass, LLC

DOiT

Routing and Switching CCIE Track

Revision: 3.0

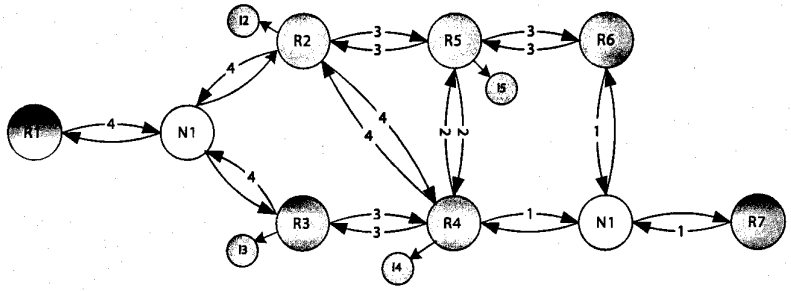

Scenario Lab BOOK For CCIE Candidates

Release 3.0

- *From The Authors of the Best Selling CCIE Preparation Guide, Cisco Certification Bridges, Routers and Switches for CCIE's*
- *8-hour Scenarios With Minimum of 15 Major Internetworking Topics In Each Scenario*
- *Scenarios Packed With "Hidden Issues"*
- *All Scenarios Use Same Equipment to Reduce Equipment Cost and Cabling Time*
- *Designed to Prepare The CCIE Candidate For "Test Day"*

Authored by Andrew Bruce Caslow and Val Pavlichenko

NetMasterClass

Revision: 3.0 (7/11/2003)

NetMasterClass
Routing and Switching CCIE Track

DOiT

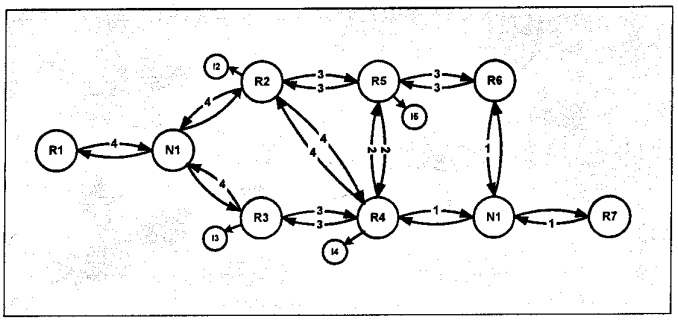

Scenario Lab BOOK

For

CCIE Candidates

Release 3.0

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

Disclaimer

NetMasterClass, LLC is an independent training and consulting company based in Herndon, Virginia. The terms "Cisco", "Cisco Systems" and "CCIE" are the trademarks of Cisco Systems, Inc. NetMasterClass, LLC is in no way affiliated with Cisco Systems, Inc.

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 contents of the NetMasterClass CCIE preparation materials is based upon the NetMasterClass "issue spotting and analysis" internetwork training methods.

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

Contents

Introduction

Thank you for purchasing the DOiT Scenario Lab Book from NetMasterClass. This lab book is the latest of a range of offerings from NetMasterClass to help CCIE candidates prepare for the CCIE lab. It is a goal of NetMasterClass to provide the CCIE candidate with a complete "end to end" suite of products and services to help CCIE candidates prepare for the CCIE lab.

The NetMasterClass slogan of "LEARNiT™, DOiT™, CHECKiT™ and DISCUSSiT™" reflects our commitment to providing an end to end suite of CCIE lab prep products and services. Listed below is a brief description of each of the NetMasterClass offerings:

LEARNiT™: Publicly Scheduled NetMasterClass Courses

Since 1997, the NetMasterClass staff has been delivering CCIE preparation courses. No other company possesses this level of experience. The NetMasterClass training staff was the first in the market to deliver a rigorous CCIE preparation class. The NetMasterClass training staff were the original innovators in developing the "issue spotting and analysis" teaching technique. This technique is applied in all of its education products and services. The current course offerings of NetMasterClass include:

Routing and Switching NetMasterClass One (RS-NMC-1)

This is a five day intensive hands-on class exposes students to a range of CCIE topics in the NetMasterClass "issue spotting and analysis" format. The course is intense. Be prepared to work late into the night during RS-NMC-1.

Routing and Switching NetMasterClass Two (RS-NMC-2)

This five day class is the sequel to RS-NMC-1. It is strongly recommended that you take RS-NMC-1 before attending RS-NMC-2. This class focuses on performing several 8-10 hour long issue spotting and analysis exams. The intensity of this class is even greater than RS-NMC-1. Again, be prepared to work late into the night during RS-NMC-2.

See www.netmasterclass.net for more details on these classes.

DOiT™: The CCIE Scenarios Lab Book

A recommendation we have for all of our CCIE lab candidates is "Perform as many well thought out issue spotting and analysis scenarios as possible". The DOiT Lab book provides 22 8-10 hour long scenarios to help you fulfill this recommendation. Each of the DOit labs possesses the following characteristics:

- Each lab is based upon a reasonable number of routers and two 3550 switches
- Each lab uses the same underlying cable scheme so that you do not need to perform extensive recabling for each lab.
- Each lab is packed full of hidden issues and challenges to test your knowledge level and analysis skills of internetworking topics.
- If a lab task becomes too challenging, you can access a "Helpful Hints" appendix

- Each lab comes with a detailed answer key that is made up of two sections: the first section provides the configuration scripts with annotations to commands that related to specific tasks; the second section is composed of explanations of hidden issues you needed to spot in each Scenario.

On-Line Rack Rentals

As a follow-on to the debut of the DoIt lab book, NetMasterClass will be offering rack rental services. The racks are designed to support the DoIT scenarios.

Continue to check www.netmasterclass.net for more DOiT scenario offerings.

CHECKiT™: The Automated Lab Grading Tool

CHECKiT is a NetMasterClass exclusive. It applies an automated grading engine to a scenario to provide the CCIE candidate with a detailed multi-page report of what the candidate did correctly in a given scenario and what the candidate missed. The multi-page report contains multiple sections focusing on individual internetworking topics such as OSPF, BGP, EIGRP, etc. As a result, CHECKiT provides the CCIE candidate with an objective assessment of where he or she stands in the preparation effort to becoming a CCIE. NetMasterClass plans to apply the CHECKiT technology to multiple scenarios.

Check www.netmasterclass.net for more details on CHECKiT

DISCUSSiT™: Enhanced Web Services for Our Students

NetMasterClass maintains an e-mail based discussion forum that is reserved only for its students and purchasers of DOiT lab book. It provides a forum for former NMC students to continue to work together towards the common goal of CCIE certification.

Also found on the NetMasterClass web-site is a collection of technical articles that you may find of interest.

General Guidelines for Performing Each DOiT Exam

Each DOiT exam is an 8+ hour long scenario consisting of at least sixteen inter-related internetworking tasks that are presented in an "issue spotting and analysis" format. What is meant by the phrase "issue spotting and analysis" is that you must CAREFULLY read the entire exam and spot the underlying hidden issues embedded within the stated exam tasks. Once you have carefully read the tasks in a specific exam, you must analyze and identify the range of possible solutions for each given task. After evaluating the different configuration options you have identified for a specific issue, you must select and implement the most appropriate configuration option that best matches the stated exam requirements. In summary, if you cannot spot the underlying hidden issues contained within the listed tasks of an exam, you cannot successfully perform a given exam.

Prerequisites for successful completion of any of the DOiT exams are:

1. A thorough understanding of all possible methods of configuring a specific type of internetworking technology. You must know what all of your options are to configure a specific internetworking technology. In order to attain this level of understanding, you must also possess a clear awareness of how the technology operates.

2. Not only must you know how to successfully configure a given technology, you must also be aware of that same technology's vulnerabilities and points of failure.

A Sample Issue Spotting and Analysis Task

A sample of an issue spotting and analysis task is presented below:

Description of exam topology relevant to the task:

You are provided with a single IP subnet configured on a hub and spoke Frame-Relay topology composed of three routers. Router R1 is the hub router and routers R2 and R3 are the spokes. All routers have their base Frame-Relay and IP configurations on a physical Serial interface. No Frame-Relay subinterfaces are used in this task.

Now, here is the issue spotting and analysis task:

"Configure OSPF over the Frame-Relay topology described above using the default OSPF network types."

From this single sentence task, your issue spotting and analysis activities should have included the following elements:

- Special attention must given to the last phrase of the task sentence "using the default OSPF network type".

- Since you possess a thorough understanding of the operation of OSPF, you have committed to memory all OSPF network types. You have also committed to memory the default OSPF network types for the type of Data-Link layer interfaces you can encounter on the CCIE lab. For example, you know that the default OSPF network type for Ethernet is broadcast and the default OSPF network type for a physical NBMA interface – Frame-Relay or ATM – is non-broadcast. If you wanted to make sure of the default OSPF network type for a specific interface, you would enter the following IOS show command: "show ip ospf

interface". Therefore, you know you must configure the OSPF non-broadcast network type on the supplied hub and spoke Frame-Relay topology.

- Now, you must spot the two underlying "hidden" issues related to this sample task:

- Hidden issue #1: the OSPF non-broadcast network type does not advertise HELLO's using the ALLSPFROUTER's multicast address 224.0.0.5. It uses a destination unicast address that must be manually supplied with a neighbor statement for each OSPF neighbor. The neighbor command is configured under the OSPF routing process.

- Hidden issue #2: We have a DR/BDR election problem on this hub and spoke topology. Here's an explanation: The OSPF non-broadcast network type performs a DR/BDR election for the segment it is configured on. All routers on a non-broadcast segment can be classified into the following three categories that relate to DR/BDR elections: a router can be the one and only DR for a segment, a router can be the one and only BDR for a segment; if a router is neither a DR or BDR, it a DROTHER. In this sample configuration task, it appears that one router will be the DR, one will be the BDR and one will be the DROTHER. However, if OSPF is to function properly, this cannot be. In conformance with the operation of OSPF, one router will be the DR; however, no router will be the BDR. In order for OSPF to function properly on a segment configured as a non-broadcast network type, all DROTHER routers must form an adjacency with both the DR and BDR routers. Since all OSPF packets have a TTL = 1, one spoke router will never form an adjacency with another spoke router. Therefore, no spoke can be either a DR or BDR router. The remedy to this hidden issue is to configure the following interface configuration command on the Frame-Relay interfaces of the spoke routers: "ip ospf priority 0". An OSPF router with a configured priority of 0 is a non-candidate DR/BDR router. Stated another way, an OSPF router configured with the interface command "ip ospf priority 0" cannot be elected DR or BDR. If we force the spoke routers to always be DROTHER routers, R1 will always be the DR. No BDR is elected on this segment.

There is more to be said about this configuration task. A more in depth discussion of the topic will be supplied in the answer key of one of your exams. You will encounter this very scenario in one of your exams. Be on the lookout for it. This sample analysis shows you how a single configuration task can contain multiple hidden issues.

It is said of many issue spotting and analysis exams, "You must pay close attention to what is said in such an exam as well as what is let unsaid".

In the sample task above, you would have had great difficulty in identifying and solving the hidden issues related to configuring the OSPF non-broadcast network type over a hub and spoke Frame-Relay network if you did not have a clear and solid understanding of the operation of OSPF as well as configuration limitations of OSPF.

Now Let's Get Ready for the Challenge

You now have ample and thorough warning. The following scenarios are full of hidden issues. It is estimated that each exam will take at least eight hours to complete. When reading each exam, keep in mind the following two points:

1. Many of the exam tasks are worded in a vague and apparently confusing manner. They are deliberately worded in such a manner to help strengthen your issue spotting and analysis skills.

2. Many of the exam tasks run counter to commonly recognized internetworking design principles. These exams are not designed to re-enforce commonly recognized internetworking design principles. They are designed to see how internetworking professionals think under complex and difficult situations. They are designed to test an internetworking professional's analysis skills when they are applied to complex and counterintuitive situations.

THEREFORE, BE PREPARED FOR STRANGELY WORDED TASKS THAT DO NOT ADHERE TO COMMONLY ACCEPTED INTERNETWORKING DESIGN PRACTICES.

When performing these exams, resist the temptation to go immediately to the answer key to learn the suggested answer for a given task. Try to spot the issues and perform the tasks to the best of your abilities.

Some other general guidelines to remember when performing this test are:

- Keep track of your time while you are performing each exam. Attempt to simulate the actual exam experience as much as possible.
- Carefully read the entire exam before actually beginning any router or switch configuration.
- Remember the "Goals and Restrictions" section provided at the beginning of each exam. These stated "Goals and Restrictions" apply to all tasks in the exam unless stated otherwise.
- Be on the lookout for words and key phrases such as: "only" and "do not" in any of the tasks.
- Prioritize your topics. Determine which topics you are going to perform first.
- Consider all of the possible options for solving a given task THEN pick the best solution among the options considered.

How to Interpret the Goals and Restrictions Section

Add interpretations to the most commonly listed Goals and Restrictions listed at the beginning of each exam such as:

1. Do not use any static routes.

Static routes can be used to provide a non-robust "stopgap" solution to a range of reachability problems. However, in the DOiTexams, use of static routes will be highly restricted. For the majority of your reachability requirements, you must rely on skillful configuration of your unicast routing protocols.

2. Network 0.0.0.0/0 should not appear in any routing table (show ip route)

A 0.0.0.0/0 entry can be used to solve a range of reachability problems. In particular, a 0.0.0.0/0 entry can be used to solve VLSM/FLSM problems that may be encountered when RIP version 1 is included within an exam. It can also be used in some type of OSPF stub area or by an ISIS Level-1 router to provide a default route. In many exams, the use any 0.0.0.0/0 entries will be restricted. When you see the restriction stated above at the beginning of a Scenario – *"Network 0.0.0.0/0 should not appear in any routing table (show ip route)"* – this means that the 0.0.0.0/0 entry cannot appear in any router or switch IP forwarding table. The absence of the

0.0.0.0/0 entry can be verified with the "show ip route" command. This restriction does not mean that you cannot enter configurations that automatically generate a 0.0.0.0/0 entry. For example, when you configure variations of OSPF stub areas, a 0.0.0.0/0 entry is created. When you configure an ISIS Level-1/Level-2 router, a 0.0.0.0/0 entry is created. These configurations will be permissible, however, if you see the restriction above, ***"network 0.0.0.0/0 should not appear in any routing table",*** make sure the 0.0.0.0/0 entry does not show up in any of your device's routing tables. This can be accomplished with the following "distance" command that is enter within the routing protocol configuration mode: "distance 255 0.0.0.0 255.255.255.255 1" This command is associated to the following standard access-list: access-list 1 permit 0.0.0.0. NOTE: If you wanted to use the 0.0.0.0/0 entry to solve a VLSM/FLSM issue or an incomplete routing information issue, some suggested alternatives to using the 0.0.0.0/0 route to solve these problems is route summarization or configuring the ip default-network command.

2. All IP addresses involved in a given scenario must be reachable.

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 includes route redistribution and the controlling of routing updates using distribute-lists, route-maps and the distance command. A key point to remember about many of these exams is: the term "redistribution" may not be explicitly stated. However, if the General Goals and Restrictions section of a given exam state that "all IP addresses involved in a given scenario must be reachable", you will more than likely need to perform redistribution in order to assure that all ip addresses are reachable.

3. 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 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.

ISDN SPIDs and Dial Plan

The NetMasterClass ISDN switch is BASIC-NI1. ISDN speeds and calling numbers are listed in Table2 Appendix B. The highlighted number is the string number used in dialer maps and dialer strings. The example of BRI configuration is as following:

R1:

```
interface BRI0/0
 ip address 172.70.12.1 255.255.255.0
 dialer map ip 172.70.12.2 name R2 broadcast 5552026
 isdn spid1 21255520240101 5552024
 isdn spid2 21255520250101 5552025
```

R2:

```
interface BRI0
 ip address 172.70.12.2 255.255.255.0
 dialer map ip 172.70.12.1 name R1 broadcast 5552024
 isdn spid1 21255520260101 5552026
 isdn spid2 21255520270101 5552027
```

Terminal Server Configuration

The terminal server is preconfigured for all exercises. See the diagram "Terminal Server Line Layout" on Figure 1.

Terminal Server Line Layout

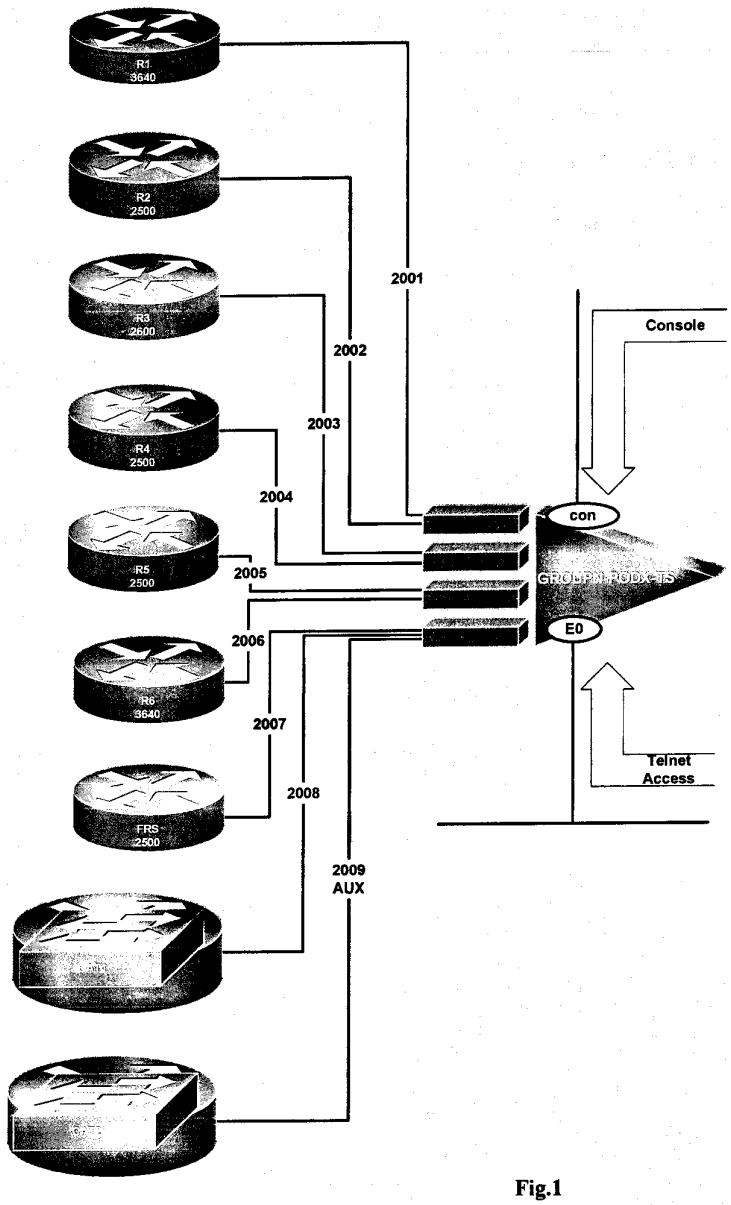

Fig.1

Example of Terminal Server configuration

The following Terminal Server configuration is configured on a Cisco 2509 router. Port 2009 is connected to the AUX port of the 2509. Privilege Level 15 is set under all line configuration modes so whenever an exec session is open, the user is immediately placed in privileged mode.

```
no ip finger
ip telnet quiet
no ip domain-lookup
ip host CAT1 2008 1.1.1.1
ip host CAT2 2009 1.1.1.1
ip host FRS 2007 1.1.1.1
ip host R6 2006 1.1.1.1
ip host R5 2005 1.1.1.1
ip host R4 2004 1.1.1.1
ip host R3 2003 1.1.1.1
ip host R2 2002 1.1.1.1
ip host R1 2001 1.1.1.1

interface Loopback0
 ip address 1.1.1.1 255.255.255.0

line con 0
 exec-timeout 0 0
privilege level 15
line 1 8
 no exec
transport input all
line aux 0
 exec-timeout 0 0
 privilege level 15
 transport input all
line vty 0 4
 exec-timeout 0 0
 privilege level 15
line vty 5 10
 exec-timeout 0 0
 privilege level 15
```

Description of NetMasterClass Student POD

Table 1 displayed below lists all of the interfaces of all the routers that comprise the rack of Cisco equipment that is used for every scenario in this workbook. When you examine the pod cable scheme in Figure 2, use this table as a reference.

If you purchase all of the routers (with the exception of the Backbone routers) in Figure 2 and cable up the routers according to this diagram, you will be able to perform all of the scenarios in this workbook without ever needing to re-cable. If you do not have this pod of equipment and you want to use a pod that is cabled to the specification of Figure 2, consider renting a rack of equipment from NetMasterClass. See Appendix B and www.netmasterclass.net for more details.

Table1. Pod Hardware Specifications

Router/Switch	Model	Interfaces
Terminal Server	cisco AS2509-RJ	1 Ethernet/IEEE 802.3 interface(s) 1 Serial network interface(s) 8 terminal line(s)
R1	cisco 3640	1 FastEthernet/IEEE 802.3 interface(s) 1 Serial network interface(s) 1 ISDN Basic Rate interface(s) 1 ATM network interface(s) 2 Voice FXS interface(s)
R2	cisco 2500	1 Ethernet/IEEE 802.3 interface(s) 2 Serial network interface(s) 1 ISDN Basic Rate interface(s)
R3	cisco 2621	2 FastEthernet/IEEE 802.3 interface(s) 2 Serial network interface(s) 2 Voice FXS interface(s)
R4	cisco 2514	2 Ethernet/IEEE 802.3 interface(s) 2 Serial network interface(s)
R5	cisco 2514	2 Ethernet/IEEE 802.3 interface(s) 2 Serial network interface(s)
R6	cisco 3640	1 FastEthernet/IEEE 802.3 interface(s) 1 Serial network interface(s) 1 ATM network interface(s)
FRS	cisco 2520	1 Ethernet/IEEE 802.3 interface(s) 2 Serial network interface(s) 2 Low-speed serial(sync/async) network interface(s) 1 ISDN Basic Rate interface(s)
CAT1	cisco WS-C3550-24	24 FastEthernet/IEEE 802.3 interface(s) 2 Gigabit Ethernet/IEEE 802.3 interface(s)
CAT2	cisco WS-C3550-24	24 FastEthernet/IEEE 802.3 interface(s) 2 Gigabit Ethernet/IEEE 802.3 interface(s)

NMC DOiT POD Layout

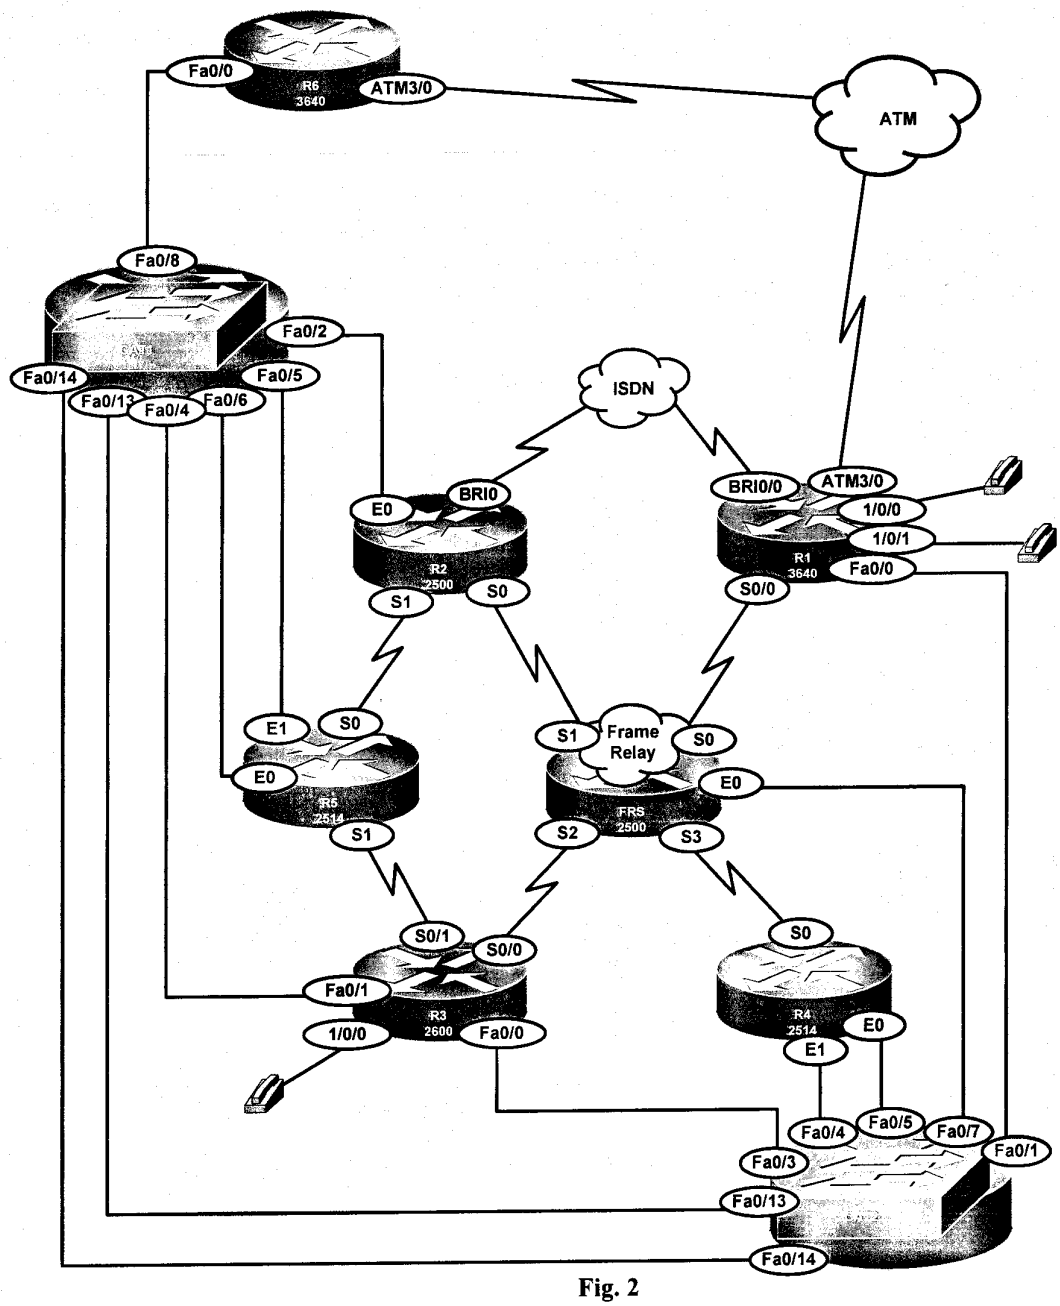

Fig. 2

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

Frame Relay Switch Configuration

The Frame Relay Switch should be preconfigured as a full mesh for all exercises. See the diagram for the DLCI numbering and corresponding interfaces. All Frame Relay Switch interfaces are connected to the DCE cable connector.

Frame Relay Switch Configuration (Full Mesh)

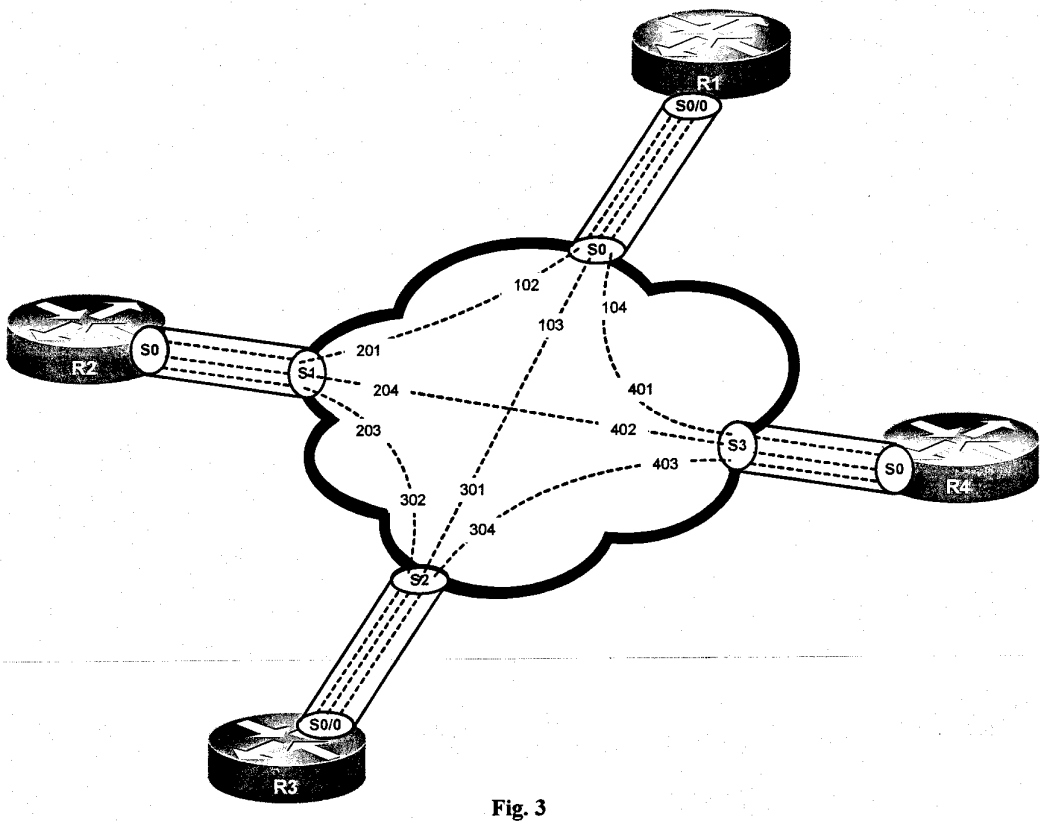

Fig. 3

Provided below is the Frame-Relay switch configuration. If you enter this configuration into a router acting as a Frame-Switch and you cable up your pod in the manner displayed in Figure 2, you will have configured a full-mesh Frame-Relay topology.

Frame Relay Switch Configuration Example

```
frame-relay switching
!
interface Serial0
 no ip address
 encapsulation frame-relay
 clockrate 64000
 frame-relay lmi-type ansi
 frame-relay intf-type dce
 frame-relay route 102 interface Serial1 201
 frame-relay route 103 interface Serial2 301
 frame-relay route 104 interface Serial3 401
!
interface Serial1
 no ip address
 encapsulation frame-relay
 clockrate 64000
 frame-relay lmi-type ansi
 frame-relay intf-type dce
 frame-relay route 201 interface Serial0 102
 frame-relay route 203 interface Serial2 302
 frame-relay route 204 interface Serial3 402
!
interface Serial2
 no ip address
 encapsulation frame-relay
 clockrate 64000
 frame-relay lmi-type ansi
 frame-relay intf-type dce
 frame-relay route 301 interface Serial0 103
 frame-relay route 302 interface Serial1 203
 frame-relay route 304 interface Serial3 403
!
interface Serial3
 no ip address
 encapsulation frame-relay
 no ip route-cache
 clockrate 64000
 frame-relay lmi-type ansi
 frame-relay intf-type dce
 frame-relay route 401 interface Serial0 104
 frame-relay route 402 interface Serial1 204
 frame-relay route 403 interface Serial2 304
```

ATM Switch Configuration

ATM Configuration Diagram

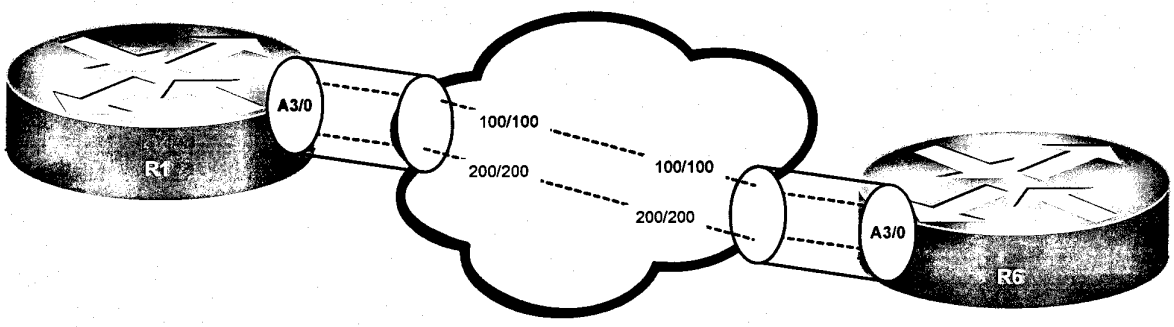

Prefix
47.0091.8100.0000.0004.4d1f.9c01

Fig. 4

```
atm address 47.0091.8100.0000.0004.4d1f.9c01.0004.4d1f.9c01.00
atm router pnni
 node 1 level 56 lowest
  redistribute atm-static
!
interface ATM4/0/0
 description PVCs
 atm pvc 100 100  interface  ATM0/0/0 100 100
 atm pvc 200 200  interface  ATM0/0/0 200 200
```

All DOiT scenarios involve two ATM interfaces within a given pod. Displayed above is a LightStream 1010 configuration that supports both SVC's and PVC's for two routers in a pod.

Scenario 1. Multi-topic CCIE level exam

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 with their original mask
- Network 0.0.0.0/0 should not appear in any routing table (show ip route)
- Do not use the "ip default-network" command
- Do not use backup interface or dialer watch in the ISDN section
- Do not introduce any new IP addresses between R2 and CAT2
- All IP addresses involved in this scenario must be reachable
- Use conventional routing algorithms

VLAN Configuration Table

Router	Interface	VLAN
R1	-	-
R2	Ethernet0	VLAN10
R3	FastEthernet0/1	VLAN30
R3	FastEthernet0/0	VLAN20
R4	Ethernet0	VLAN30
R5	Ethernet0	VLAN20
R6	FastEthernet0/0	VLAN3000
FRS	Ethernet0	VLAN3000
CAT1	-	VLAN30
CAT2	-	VLAN10

Attention For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15

Scenario 1

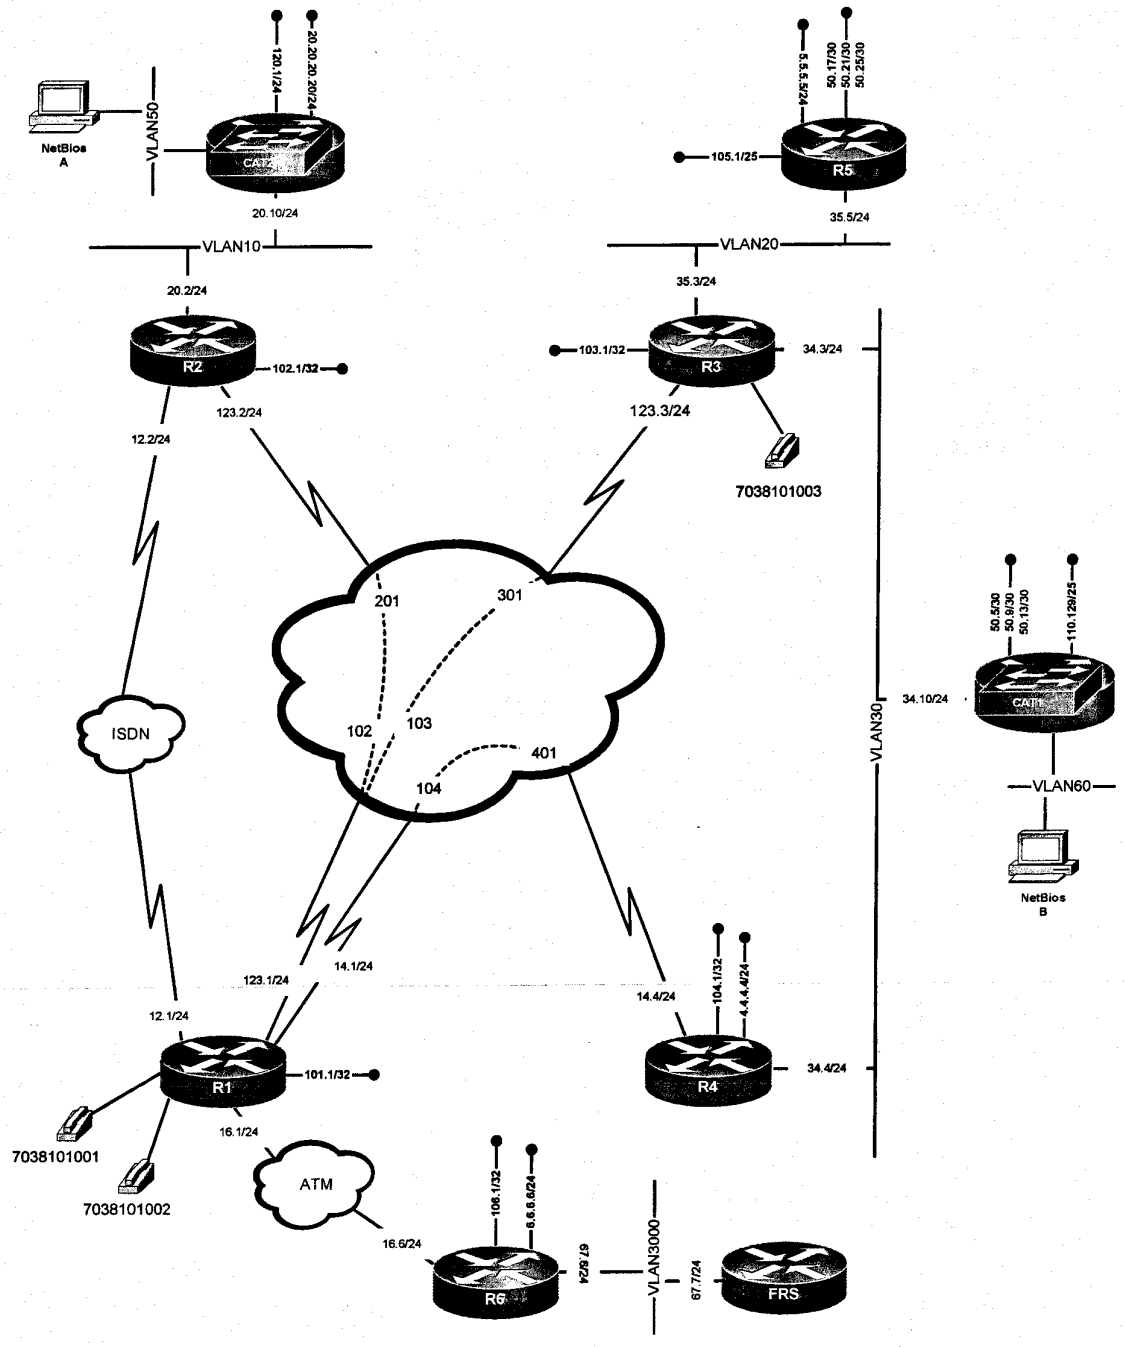

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

1.1 Frame Relay

1.1.1 Use only the PVCs displayed in the diagram.
1.1.2 Make sure you have full connectivity within Frame Relay subnets
1.1.3 Use physical interfaces wherever possible, otherwise use point-to-point logical interfaces

1.2 Catalyst Configuration

1.2.1 Configure VLANs according to the "VLAN Configuration Table"
1.2.2 Configure an ISL trunk on the link between the Fa0/13 ports of CAT1 and CAT2
1.2.3 Configure a dot1q trunk on the link between the Fa0/14 ports of CAT1 and CAT2
1.2.4 Allow only VLAN10 on the ISL trunk
1.2.5 Allow all other VLANs on the dot1q trunk
1.2.6 Configure the VTP mode suitable for all other tasks in this exam

1.3 ATM

1.3.1 Configure PVC 100/100 between R1 and R6
1.3.2 Use a service class where ATM switches make no guarantee of cell delivery, but guarantee a minimum bit rate and assure that cell loss is kept as low as possible by using a feedback mechanism
1.3.3 Set the Maximum cell rate at which the source can transmit to 128 Kbit/sec
1.3.4 Guarantee 64Kbit/sec in case of congestion

1.4 OSPF

1.4.1 Add all interfaces on the 172.16.123.0/24 subnet to OSPF area 0
1.4.2 In area 0, use the OSPF network type which forces a DR/BDR election but uses unicast packets for Hello and Database Exchange services
1.4.3 Add all interfaces on the 172.16.14.0/24 subnet to OSPF area 14. Use the best matching OSPF network type for this particular link
1.4.4 Advertise loopback 172.16.103.1/32 in OSPF without adding it in any area. It should appear in the routing tables as type E2.
1.4.5 Advertise loopback 172.16.102.1/32 in OSPF area 20
1.4.6 Advertise loopback 172.16.101.1/32 in OSPF area 10

1.5 RIP

1.5.1 Configure RIP version 1 between R2 and CAT2
1.5.2 Configure RIP version 2 between R3, R4 and CAT1
1.5.3 Advertise the following loopback interfaces on CAT1:

- 172.16.50.5/30
- 172.16.50.9/30
- 172.16.50.13/30

1.5.4 Advertise network 4.4.4.4/24 from R4

1.6 EIGRP

1.6.1 Configure EIGRP AS 100 between R3 and R5
1.6.2 Advertise the following loopback interfaces on R5:

- 172.16.50.17/30
- 172.16.50.21/30
- 172.16.50.25/30
- 172.16.105.1/25

1.6.3 Detect the loss of an EIGRP neighbor in a time period that is twice as fast as the default

1.7 ISIS

1.7.1 Configure ISIS between R1, R6 and FRS
1.7.2 Assign area 49.0001 and System ID 1111.1111.1111 to router R1
1.7.3 Assign area 49.0002 and System ID 6666.6666.6666 to router R6
1.7.4 Assign area 49.0002 and System ID 7777.7777.7777 to router FRS
1.7.5 Make sure all ISIS routers form minimally required adjacency levels on a per link basis
1.7.6 Advertise 172.16.106.1/32 on R6
1.7.7 R6 should not forward ISIS transit traffic for a time period of 20 minutes after R6 has been rebooted.

1.8 BGP

1.8.1 Configure AS600 on R6, AS700 on FRS and peer between them
1.8.2 Configure AS200 on CAT2 and AS500 on R5
1.8.3 Originate the following prefixes in BGP only within their respective ASes. Do not originate them in an IGP. These prefixes must be exchanged between all AS's:

- 20.20.20.0/24 (CAT2)
- 5.5.5.0/24 (R5)
- 6.6.6.0/24 (R6)

1.8.4 Configure AS100; do not disable synchronization within AS100. Use the minimum number of BGP speakers and peer relationships to accomplish this task
1.8.5 Provide transit paths between the above advertised prefixes through AS100

1.9 DLSW

1.9.1 Workstations A and B are connected to CAT2 VLAN50 FA0/18 and CAT1 VLAN60 FA0/18 respectively, run a NetBIOS application requiring communications between these two workstations.
1.9.2 Use a TCP transport path between R2 and R3 as a primary connection
1.9.3 Use a TCP transport path between R2 and R4 as a backup. Do not use the backup peer command
1.9.4 Use the minimal number of remote peer statements and configure them on a single router

1.10 ISDN

1.10.1 Configure ISDN to backup the frame relay link on R2.

1.10.2 Place the ISDN link in OSPF area 12.
1.10.3 Use CHAP authentication, password "nmc"
1.10.4 Only R2 should call R1
1.10.5 Make sure ISDN does not stay up indefinitely

1.11 Address Administration

1.11.1 Configure a DHCP server on R3 to supply the following settings to the workstations connected to ports Fa0/20 and Fa0/21 of CAT1 on VLAN30:

- Domain: xyz.com
- Dns server: 172.16.34.10
- Gateway

Use the gateway IP address most suitable for other tasks in this scenario.

1.11.2 Some users of the workstations connected to Fa0/20 and Fa0/21 complain about not getting the DHCP settings. Provide a solution.

1.12 Security

1.12.1 Make R3 an http server. Configure R1 to prevent http access from the R2 serial interface to R3 during the working day from 8:00am to 5:00pm
1.12.2 Test it with the telnet command

1.13 VOICE

1.13.1 The telephone number 703-810-1001 is assigned to port 1/0/0 on router R1
1.13.2 The telephone number 703-810-1002 is assigned to port 1/0/1 on router R1
1.13.3 The telephone number 703-810-1003 is assigned to port 1/0/0 on router R3
1.13.4 Configure Voice over IP to allow calls from R3 to R1 and vice versa
1.13.5 Use only the last three digits to make calls. Do not use any kind of regular expressions

1.14 QOS

1.14.1 Allocate 20% of bandwidth to ftp traffic, packet size 1500 bytes
1.14.2 Allocate 30% of bandwidth to http traffic, packet size 600 bytes
1.14.3 Allocate 40% of bandwidth to udp port 5001 traffic, packet size 300 bytes
1.14.4 Allocate 10% of bandwidth to other types of traffic, packet size 1500 bytes
1.14.5 The traffic listed above is sent out the interface s0 of R4
1.14.6 Use Custom Queuing to accomplish this section

1.15 Catalyst Specialties

1.15.1 Change the MTU size on VLAN 3000 to twice higher than default
1.15.2 Ports Fa0/11 and Fa/012 of CAT1 are connected to jacks located in the visitor room. Only two vlans, 700 and 800, are allowed in the visitor room. Visitors can be either on VLAN800 or VLAN700 depending on their security clearance.
1.15.3 The visitors' configuration file is on the switch with the ip address 172.16.200.1

1.16 Gateway Redundancy

1.16.1 Configure an HSRP standby group 10 between R3 and R4.
1.16.2 Make R4 a primary and R3 a backup gateway for the workstations connected to VLAN30
1.16.3 The virtual gateway IP address is 172.16.34.1

1.17 Multicast

1.17.1 Source the IP multicast stream destined to 225.5.5.5 from router R5
1.17.2 Build a Shortest Path Tree (SPT) only. Routers R1, R2, R3, R4 and R6 will be the members of this tree.
1.17.3 Join one of the loopback interfaces on each router within the SPT to multicast group 225.5.5.5
1.17.4 Make sure you receive replies from all of the routers listed above on R5

1.18 NTP

1.18.1 Configure R1 as NTP master stratum 3
1.18.2 Configure R4 as an NTP client
1.18.3 Configure a peer association between R3 and R4

Scenario 2. Multi-topic CCIE level exam

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 with their original mask
- Network 0.0.0.0/0 should not appear in any routing table (show ip route)
- Do not use the "ip default-network" command
- Do not introduce new links to R4
- Do not change the original mask unless it is explicitly specified otherwise
- Do not use backup interface or dialer watch in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10 VLAN20 VLAN30
R2	Ethernet0	VLAN40
R3	-	-
R3	FastEthernet0/0	VLAN50
R4	-	-
R5	Ethernet0	VLAN10
R6	FastEthernet0/0	VLAN20
FRS	Ethernet0	VLAN50
CAT1	-	VLAN40
CAT2	-	VLAN30

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 2

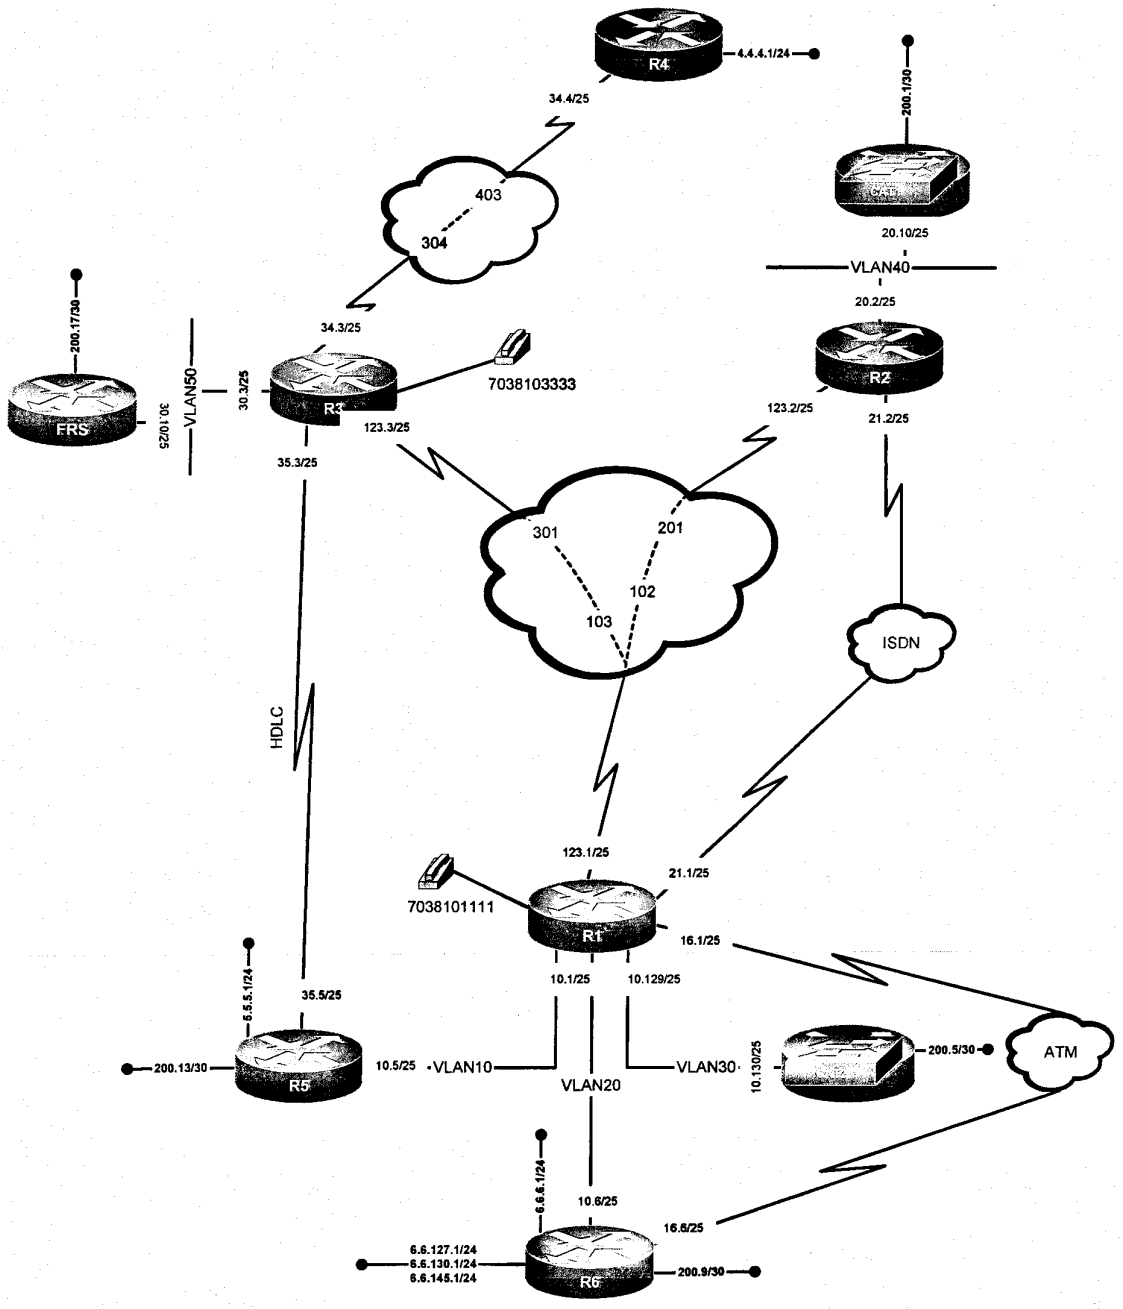

2.1 Frame Relay

2.1.1 Use only the PVC's displayed in the diagram
2.1.2 Use logical interfaces on R3 and a physical interface on R4 on the Frame Relay subnets
2.1.3 Use a logical interface on R1 and a physical interface on R2 on the Frame Relay subnets

2.2 Catalyst Configuration

2.2.1 Configure VLAN's according to the VLAN configuration table and the diagram
2.2.2 Advertise VLAN's from CAT1 to CAT2
2.2.3 Make sure ports Fa0/13 and Fa0/14 are not in a blocking state for ***any*** VLAN
2.2.4 Use the dot1q protocol for all trunks involved in this scenario

2.3 ATM

2.3.1 All ATM devices share the prefix 47.0091.8100.0000.0004.4d1f.9c01
2.3.2 Configure the End System Identifier 0001.0001.0001 on R1
2.3.3 Configure the End System Identifier 0006.0006.0006 on R6
2.3.4 Use logical interfaces in this section
2.3.5 Do not use dynamic methods of Layer3/Layer2 mapping in the configuration

2.4 OSPF

2.4.1 Configure OSPF area 3 on VLAN40
2.4.2 The devices connected to VLAN40 should become fully adjacent only if they pass non-clear text authentication using the string "nmc
2.4.3 Advertise loopback interface 172.16.200.1/30 into OSPF

2.5 RIP

2.5.1 Configure RIP version 1 between routers R3 and R4
2.5.2 Configure RIP between R1, R2 and R3. Unicast RIP updates on the Frame Relay network
2.5.3 Exchange RIP updates on Frame Relay links and ISDN link only

2.6 EIGRP

2.6.1 Configure EIGRP AS100 on the link between R3 and R5
2.6.2 Configure EIGRP AS200 between routers R1, R5, R6 and CAT2. Use a unicast EIGRP packet exchange between these routers
2.6.3 Configure EIGRP AS200 on the ATM link
2.6.4 Advertise loopback networks on R6 to AS200 without using a network statement:

- 6.6.127.0/24
- 6.6.130.0/24
- 6.6.145.0/24

2.6.5 Summarize these subnets with the most optimal mask

2.6.6 Advertise the following Loopback networks into EIGRP AS 200:

- 172.16.200.5/30
- 172.16.200.9/30
- 172.16.200.13/30

2.7 ISIS

2.7.1 Configure Network Entity Title 49.0001.3333.3333.3333.00 on R3
2.7.2 Configure Network Entity Title 49.0002.1010.1010.1010.00 on FRS
2.7.3 Do not use the default level configuration
2.7.4 R3 and FRS should become adjacent on VLAN50 only if they pass authentication using the string "nmc". The Authentication string must be carried in the ISIS HELLO packets
2.7.5 Advertise Loopback 172.16.200.17/30 from FRS

2.8 BGP

2.8.1 Configure AS 400 on R4
2.8.2 Configure AS 100 on R3 and R1
2.8.3 Configure AS 500 on R5
2.8.4 Configure AS 600 on R6
2.8.5 Peer AS 400 with AS100 between R4 and R3
2.8.6 Peer AS 500 and AS 100 between R5 and R1
2.8.7 Peer AS 600 and AS 100 between R1 and R6
2.8.8 Advertise Loopback network 4.4.4.0/24 from R4, Loopback network 5.5.5.0/24 from R5 and Loopback network 6.6.6.0/24 from R6
2.8.9 AS 100 should provide transit service to network 6.6.6.0/24 from 4.4.4.0/24 only if AS100 loses reachability to network 5.5.5.0
2.8.10 Turn synchronization off in this section

2.9 DLSW

2.9.1 Configure DLSW peer relationships from R3 to R5 and from R3 to R6
2.9.2 NetBIOS workstations are located on VLAN50 and VLAN20
2.9.3 The DLSW peer relationship between R3 and R6 should be used to provide a TCP transport between NetBIOS workstations
2.9.4 In case R6 fails, the DLSW peer relationship between R3 and R5 should be used
2.9.5 Tear down the DLSW peer relationship between R3 and R5 5 minutes after the DLSW peer relationship between R3 and R6 is up again.

2.10 ISDN

2.10.1 Use CHAP authentication, password "nmc"
2.10.2 Change the host names in the challenges to "R1R1" and "R2R2"
2.10.3 Only R1 should call R2
2.10.4 Make sure ISDN does not stay up indefinitely

2.11 Router Maintenance

2.11.1 There are diskless SUN workstations on VLAN 40
2.11.2 The SUN Server, where the workstations can boot their OS from, is located on VLAN 50. Its IP address is 172.16.30.100
2.11.3 Provide communications between WS1(IP address 172.16.20.100, MAC 8:0:20:ac:24:b8) , WS2 ((IP address 172.16.20.101, MAC 8:0:20:ac:24:b9) and the SUN server.
2.11.4 Provide the solution which does not require a DHCP server

2.12 Security

2.12.1 Allow Finger to R4 from R1, R2 and R5. Disallow finger connections from R4 to R1, R2 and R5. Do not use the "established" keyword.

2.13 VOICE

2.13.1 The telephone number 703-810-1111 is assigned to port 1/0/0 on router R1
2.13.2 The telephone number 703-810-3333 is assigned to port 1/0/0 on router R3
2.13.3 Configure the connection between these two phones using the Real Time Protocol transport. When the telephone attached to R3 is taken off-hook, automatically ring the telephone on router R1.

2.14 QOS

2.14.1 Assign DSCP value 45 to frames received by CAT1 on ports Fa0/11 and Fa0/12

2.15 Catalyst Specialties

2.15.1 CAT1 ports Fa0/11 and Fa0/12 are patched to the Conference Room
2.15.2 Only the two workstations described in the Router Maintenance section can be used in the Conference Room.
2.15.3 Also, two Cisco 7960 IP Phones with the MAC addresses of 00-07-85-95-D1-A7 and 00-07-85-95-D2-B7 are connected to the CAT1 ports Fa0/11and Fa0/12
2.15.4 Configure Priority Tagged Frames to carry voice traffic on the Native VLAN
2.15.5 Override the Class of Service value in the packets generated by the SUN workstations with a value of 4

2.16 Gateway Redundancy

2.16.1 Assign 172.16.10.126 to a virtual gateway and make sure the MAC address associated with the virtual gateway is set to 0000.0c07.ac1a
2.16.2 Authenticate HSRP on the 172.16.10.0/25 subnet (password nmc). Exchange Hello packets 3 times faster than by default.
2.16.3 Select R1 as the preferred gateway by using priority 150.
2.16.4 Select R5 as the preferred gateway if the Frame Relay connection on R1 goes down
2.16.5 R1 should take responsibility back after the Frame Relay connection comes back up

2.17 Multicast

2.17.1 Simulate the multicast stream (172.16.34.4, 232.2.2.2) from R4
2.17.2 Statically configure a Shared Tree between R3, R5, R1, R6 and R2 rooted from R5
2.17.3 Join group 232.2.2.2 on every member of the Shared Tree by using one of the router interfaces
2.17.4 ISDN and ATM are not part of the multicast network
2.17.5 Make sure you get a reply from all joined routers on R4

Scenario 3. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16
- Do not use any static routes, unless required for backup purposes
- Advertise Loopback interfaces involved in IGP's with their original mask. Do not change the mask.
- Network 0.0.0.0/0 should not appear in any routing table (show ip route)
- Do not use "ip default-network"
- Do not use backup interface or dialer watch in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN20
R2	Ethernet0	-
R3	FastEthernet0/0	VLAN10
R3	FastEthernet0/1	VLAN30
R4	Ethernet0	VLAN30
R5	Ethernet0	-
R6	FastEthernet0/0	VLAN10
FRS	Ethernet0	VLAN20
CAT1	-	-
CAT2	-	VLAN30

Attention **For physical connectivity, check the diagram "NMC Pod Layout" on Figure 2 on page 15**

Scenario 3

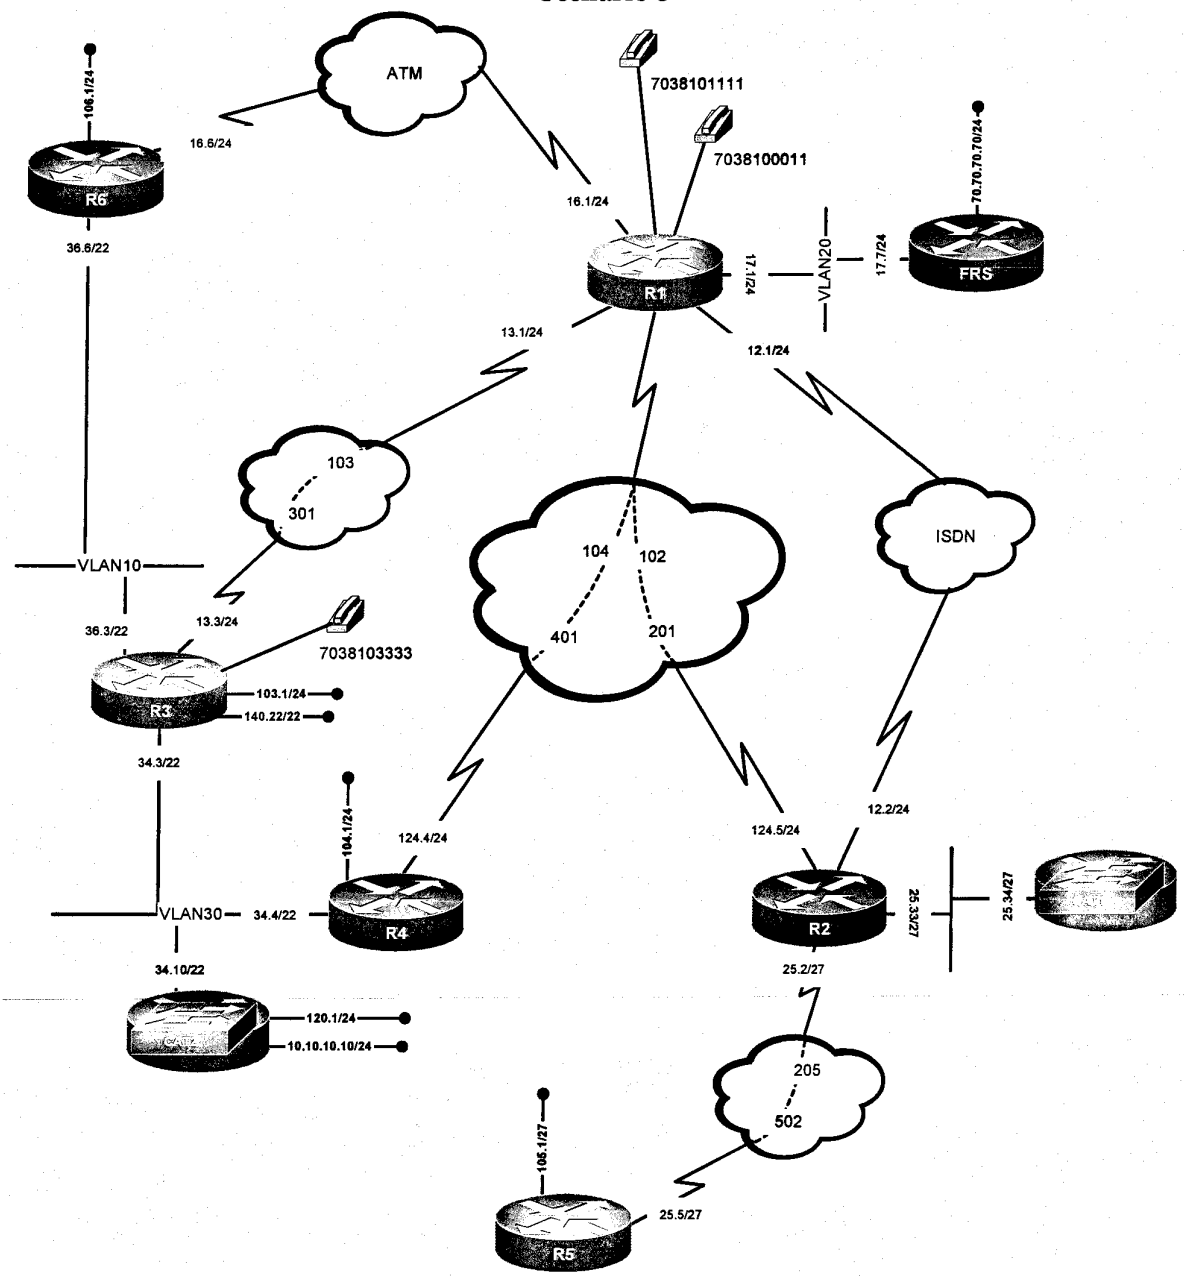

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

3.1 Frame Relay

3.1.1 Use only the PVC's displayed in the diagram
3.1.2 Use physical interfaces on subnet 172.16.13.0/24, use physical interfaces for the 172.16.124.0/24 subnet on routers R2 and R4 only.
3.1.3 Configure physical interfaces for subnet 172.16.25.0/27
3.1.4 Check the diagram for IP addressing

3.2 Catalyst Configuration

3.2.1 Configure VLAN's according to the "VLAN Configuration Table" and diagram
3.2.2 The VTP mode is the default setting
3.2.3 On the trunks, allow only VLAN's involved in this scenario
3.2.4 Routing on CAT1 should be configured as its default setting.
3.2.5 The trunking protocol dot1q should be configured for this scenario

3.3 ATM

3.3.1 The ATM switch advertises prefix 47.0091.8100.0000.0004.4d1f.9c01
3.3.2 Configure End System Identifier 0001.0001.0001 on R1
3.3.3 Configure End System Identifier 0006.0006.0006 on R6
3.3.4 Use physical interfaces on the ATM link
3.3.5 Build the Virtual Circuit when there is traffic. Dynamically map Layer3 IP addresses to Layer 2

3.4 OSPF

3.4.1 Configure OSPF area 0 on the subnet 172.16.124.0/24
3.4.2 R2 and R4 must possess unique next hop addresses for traffic forwarded to R1
3.4.3 Configure OSPF area 13 between R1 and R3. Elect R1 as DR. Make sure R1 is still DR after you reboot it.
3.4.4 R1 and R3 should automatically discover each other
3.4.5 Place the ISDN link in OSPF area 12
3.4.6 On R3, place network 172.16.103.0/24 into the OSPF process as external Type 1
3.4.7 On R3, advertise network 172.16.140.0/22 as OSPF area 10

3.5 RIP

3.5.1 Configure RIP version 1 between routers R1 and FRS
3.5.2 Configure RIP version 2 between R3, R4 and CAT2. Allow RIP updates only on VLAN30

3.6 EIGRP

3.6.1 Configure EIGRP AS1 between R1,R3 and R6
3.6.2 Do not form the EIGRP adjacency on subnet 172.16.13.0/24
3.6.3 Advertise network 172.16.106.0/24 from R6
3.6.4 Authenticate all EIGRP adjacencies with password "nmc"

3.7 ISIS

3.7.1 Configure Network Entity Title on R2 to 49.0002.2222.2222.2222.00
3.7.2 Configure Network Entity Title on R5 to 49.0005.5555.5555.5555.00
3.7.3 Use the default "is-type" for this configuration
3.7.4 Restrict the adjacency only to one type
3.7.5 Advertise the network 172.16.105.1/27 with the initial metric of 0.

3.8 BGP

3.8.1 Configure AS 700 on FRS
3.8.2 Configure AS 100 on R1, R2 and R6. Do not use the default synchronization method. Do not form a BGP peer relationship between R2 and R6
3.8.3 Peer AS700 and AS100 between FRS and R1
3.8.4 Peer AS 100 and AS 300 between routers R6 and R3 as well as R4 and R2.
3.8.5 Configure AS 300 on R3 and R4. Use the default synchronization method in AS300
3.8.6 Configure AS1000 on CAT2 and peer it with both routers of AS300
3.8.7 Advertise network 70.70.70.0/24 on AS700 as incomplete
3.8.8 Advertise network 10.10.10.0/24 on AS1000 as IGP
3.8.9 The transit path from network 10.10.10.0 to network 70.70.70.0 should be comprised of CAT2, R4, R3, R6, R1 and FRS in this particular order
3.8.10 The transit path from network 70.70.70.0 to network 10.10.10.0 should be comprised of FRS, R1, R6, R3, R4 and CAT2 in this particular order
3.8.11 Accomplish the previous two tasks with LOCAL PREF manipulation

3.9 DLSW

3.9.1 Configure a DLSW peer relationship between R1 and R2 using the TCP protocol
3.9.2 Make sure the DLSW peer is not established over Frame-Relay

3.10 ISDN

3.10.1 Use clear text authentication, use password "nmcr1" on R1 and "nmcr2" on R2
3.10.2 R1 should trigger an ISDN call to R2 based on the absence of the 172.16.105.0/27 prefix
3.10.3 Only R1 should call R2
3.10.4 Make sure ISDN does not stay up indefinitely

3.11 Router Maintenance

3.11.1 The Network Administrator (SALLY) would like to access R1 via Internet Explorer at any time from only the subnet assigned to VLAN30
3.11.2 Sally would like to access router R1 using a non-standard port and with her user name SALLY
3.11.3 Her user name is defined on router R1

3.12 Security

3.12.1 None of the devices on VLAN30 should be able to telnet straight to R5 unless SALLY authenticates herself on R1
3.12.2 The users on VLAN30 should be able to maintain a telnet session for no longer than 10 minutes after SALLY logs onto router R1
3.12.3 Limit your configuration so that other traffic involved in this scenario is not disrupted

3.13 VOICE

3.13.1 The telephone number 703-810-1111 is assigned to port 1/0/0 on router R1
3.13.2 The telephone number 703-810-0011 is assigned to port 1/0/1 on router R1
3.13.3 The telephone number 703-810-3333 is assigned to port 1/0/0 on router R3
3.13.4 When the telephone attached to R1 port 1/0/0 is taken off-hook, ring the telephone connected to port 1/0/1 on R1
3.13.5 Configure one dial peer from R3 to R1 for both 703-810-1111 and 703-810-0011
3.13.6 Use the minimum packet overhead technique to consume less bandwidth for voice connections

3.14 QOS

3.14.1 Configure QOS on the GigabitEthernet0/1 interface of CAT2 using the most simplified method of configuration and allowing the discovery of IP phones.
3.14.2 Two IP phones are connected to ports Fa0/10 and Fa0/11 of CAT1. Configure QOS on these ports using the same technique applied in the previous task

3.15 Catalyst Specialties

3.15.1 Enable message logging on CAT1
3.15.2 Write messages to the file nmc.log stored on the local flash
3.15.3 Do not allow the logging process to exceed a file size that is twice as large as the default size.
3.15.4 Error messages about software or hardware malfunctions only should be logged

3.16 Address Administration

3.16.1 Besides two devices connected to VLAN10, the rest of the address space of VLAN10 is taken for the workstations
3.16.2 The network administrator would like to control the traffic generated by workstations by forwarding it to both R3 and R6, with the traffic distribution being approximately half and half
3.16.3 R3 and R6 should be configured to provide the TCP/IP information to the workstations, such as domain name nmc.com, IP address of DNS server 172.16.17.2
3.16.4 All workstations are Microsoft clients using a hybrid NetBIOS node type
3.16.5 In case of failure of either Frame Relay or ATM links on R3 or R6 respectively, the traffic should be sent through the remaining link
3.16.6 When the failed link comes back on, return to the previous forwarding scheme

3.17 Multicast

3.17.1 Configure PIM neighbor relationships between R4 and R1, R1 and R2 as well as R2 and R5
3.17.2 The PIM network is based on the Flood and Prune concept
3.17.3 ISDN is not involved in this section
3.17.4 Generate traffic to multicast group 233.3.3.3 from CAT2
3.17.5 Join one of the interfaces on routers R4, R1, R2 and R5 to group 233.3.3.3
3.17.6 Make sure CAT2 receives ping replies from all multicast routers

Scenario 4. Multi-topic CCIE level exam

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 with their original mask
- Network 0.0.0.0/0 should not appear in any routing table (show ip route)
- Do not use dialer watch in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	Ethernet0	VLAN10
R3	FastEthernet0/0	VLAN10
R3	FastEthernet0/1	-
R4	Ethernet0	VLAN10
R4	Ethernet1	VLAN50
R5	Ethernet0	VLAN20
R5	Ethernet1	VLAN50
R6	FastEthernet0/0	VLAN10 VLAN30
FRS	Ethernet0	VLAN40
CAT1	-	VLAN20
CAT2	-	VLAN30 VLAN40

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 4

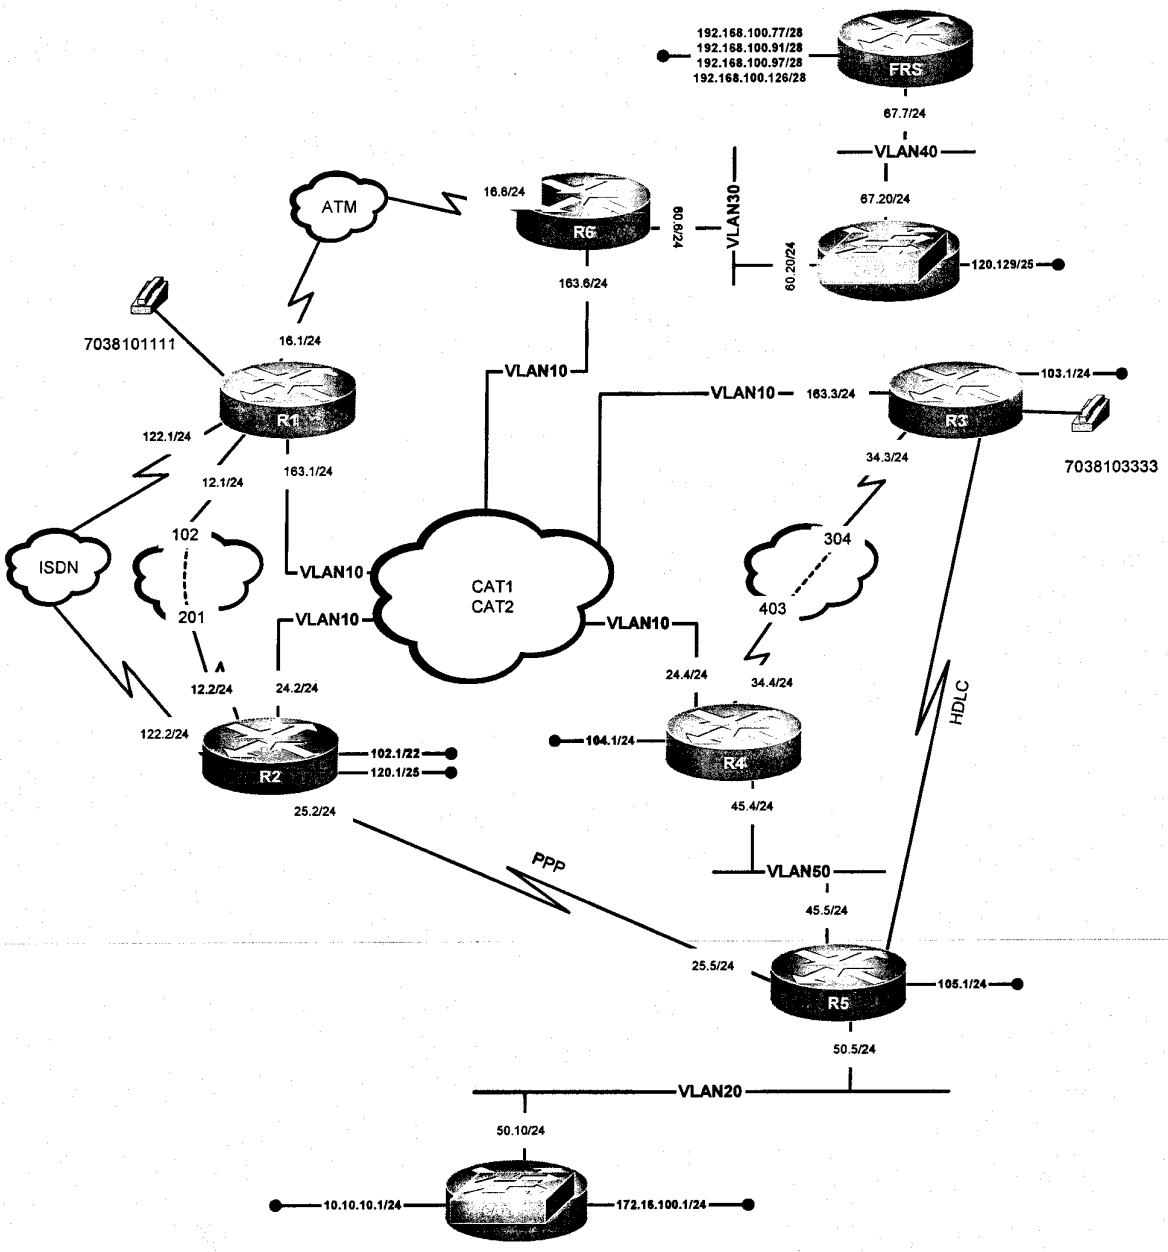

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

4.1 Frame Relay

4.1.1 Use only the PVC's displayed in the diagram
4.1.2 Use a physical interface on R3 and a point-to-point interface on R4 on subnet 172.16.34.0/24
4.1.3 Use logical interfaces on subnet 172.16.12.0/24

4.2 Catalyst Configuration

4.2.1 Configure VLAN10 for subnet 172.16.163.0/24 between R1, R6 and R3
4.2.2 Configure VLAN10 for subnet 172.16.24.0/24 between R2 and R4
4.2.3 Networks 172.16.163.0/24 and 172.16.24.0/24 must represent separate broadcast domains. Do not use filtering to accomplish this task.
4.2.4 Configure other VLAN's according to the "VLAN configuration" table
4.2.5 Use dot1q trunking protocol in this scenario

4.3 ATM

4.3.1 Configure PVC 100/100 between R1 and R6
4.3.2 Do not use any static and dynamic mapping between Layer 3 and Layer 2
4.3.3 Assign IP addresses given on the diagram

4.4 OSPF

4.4.1 Configure OSPF area 0 on subnet 172.16.163.0/24 between R1 and R6
4.4.2 Configure OSPF area 16 on subnet 172.16.16.0/24 between R1 and R6
4.4.3 Configure OSPF area 12 on subnets 172.16.12.0/24 and 172.16.122.0/24 between R1 and R2
4.4.4 Advertise the Loopback interface 172.16.102.0/22 from R2 as area 20
4.4.5 Advertise the Loopback interface 172.16.120.0/25 from R2 without including it in any areas
4.4.6 Configure OSPF area 60 on subnet 172.16.60/24.
4.4.7 Do not allow OSPF external routing information into area 60. Do not use any filtering technique to accomplish this task

4.5 RIP

4.5.1 Configure RIP version 1 between R5 and CAT1
4.5.2 Configure RIP version 2 between CAT2 and FRS
4.5.3 Configure RIP version 2 between R2, R4 and R5
4.5.4 Configure mutual redistribution between RIP and other routing protocols on R2 and R4
4.5.5 RIP must exchange updates only on the links between R2 and R5, R4 and R5, R5 and CAT1 and on the Ethernet link between CAT2 and FRS only
4.5.6 Router R5 should send traffic to R4 as a preferred next hop.

4.6 EIGRP

4.6.1 Configure EIGRP AS 10 only on the subnet 172.16.24.0/24
4.6.2 R2 and R4 should consider each other as valid neighbors for a period that is twice longer than the default

4.7 ISIS

4.7.1 Configure area 49.0001 and System ID 6666.6666.6666 on router R6
4.7.2 Configure area 49.0001 and System ID 4444.4444.4444 on router R4
4.7.3 Configure area 49.0001 and System ID 3333.3333.3333 on router R3
4.7.4 Restrict the adjacency to only one type
4.7.5 Make sure traffic from R1 destined to 172.16.34.0/24 is sent directly to R3

4.8 BGP

4.8.1 Configure AS 100 on routers R5 and CAT1
4.8.2 Configure AS 200 on routers R2, R1 and R6
4.8.3 Configure AS 300 on routers R3 and R4
4.8.4 Configure AS 400 on routers CAT2 and FRS
4.8.5 Peer AS400 and AS200 between CAT2 and R6
4.8.6 Peer AS400 and AS300 between CAT2 and R3
4.8.7 Peer AS 200 and AS100 between R2 and R5
4.8.8 Peer AS 300 and AS100 between R4 and R5
4.8.9 Routers R6, R1 and R2 should be in different AS's
4.8.10 Advertise the loopback interfaces from FRS:

- 192.168.100.77/28
- 192.168.101.91/28
- 192.168.102.97/28
- 192.168.103.126/28

4.8.11 All other BGP routers must possess only one prefix best describing these loopbacks.
4.8.12 In AS 100, prefer R4 as a next hop for networks advertised by AS400
4.8.13 The Administrator of AS100 decided to use the default synchronization configuration
4.8.14 Advertise network 10.10.10.0/8 from CAT1. This network should be seen in AS400 with the AS PATH three times longer via R3

4.9 DLSW

4.9.1 Configure a DLSW peer relationship between R5 and R3 using the minimal packet overhead
4.9.2 Configure a DLSW peer relationship between R5 and R2. Only R2 will possess the DLSW configuration type "conf".

4.10 ISDN

4.10.1 Bring the ISDN link up only in case R1's DLCI goes to an INACTIVE state
4.10.2 Make sure OSPF does not keep the ISDN link up indefinitely

4.11 Network Monitoring

4.11.1 The system Administrator would like to monitor the network performance, primarily response time, on router R1.
4.11.2 The specific list of parameters to monitor is the following:

- RTT taken to connect and access data from an HTTP server on CAT2 based on GET requests
- Time taken to establish a TCP session to R5 using port 80

4.12 Security

4.12.1 R5 must accept packets sourced from AS400 networks advertised in a previous task from R4 only
4.12.2 Do not use "ip access-group" to accomplish this task
4.12.3 In case of link failure between R4 and R5, R5 should accept this traffic from R2

4.13 VOICE

4.13.1 Provide a VOIP connection between two PBXs connected to port 1/0/0 of R1 and port 1/0/0 of R3
4.13.2 The number 7038101111 is assigned to port 1/0/0 on router R1
4.13.3 The number 7038103333 is assigned to port 1/0/0 on router R3

4.14 QOS

4.14.1 Prioritize traffic sent via R2 based on the following regulations:

- DNS traffic is very critical
- NTP traffic is very critical
- HTTP traffic is critical
- SMTP traffic is normal
- FTP traffic is not critical

4.14.2 Prioritize traffic coming from VLAN20 and forwarded to R4 as medium
4.14.3 Apply this configuration on R5

4.15 Catalyst Specialties

4.15.1 Activate HTTP server on CAT1
4.15.2 Configure both switches to store configuration files in CAT1.conf and CAT2.conf on the flash
4.15.3 Limit configuration files to a size of 200000 bytes

4.16 Address Administration

4.16.1 Loopback network 172.16.100.0/24 should be in the routing table of only two routers: CAT1 and R5
4.16.2 Make sure you can ping the rest of the network from 172.16.100.1

4.17 Multicast

4.17.1 Source traffic (172.16.50.10, 230.4.4.4) from CAT1
4.17.2 Configure PIM Dense Mode multicast routing between R5, R4 and R2
4.17.3 Configure Loopback interfaces as members of the destination multicast group on routers R5, R4 and R2
4.17.4 An IPTV server is connected to VLAN20. The server is streaming 64 audio programs
4.17.5 Configure links between R2 and R5 for minimal bandwidth consumption by the audio streams

Scenario 5. Multi-topic CCIE level exam

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 with their original mask
- Do not use "ip default-network"
- Do not use backup interface or dialer watch in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN 146
R2	Ethernet0	VLAN23
R3	FastEthernet0/0	VLAN23
R3	FastEthernet0/1	-
R4	Ethernet0	VLAN146
R5	Ethernet0	VLAN200
R6	FastEthernet0/0	VLAN146
FRS	Ethernet0	VLAN200
CAT1	-	VLAN146, VLAN200
CAT2	-	VLAN200

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 5

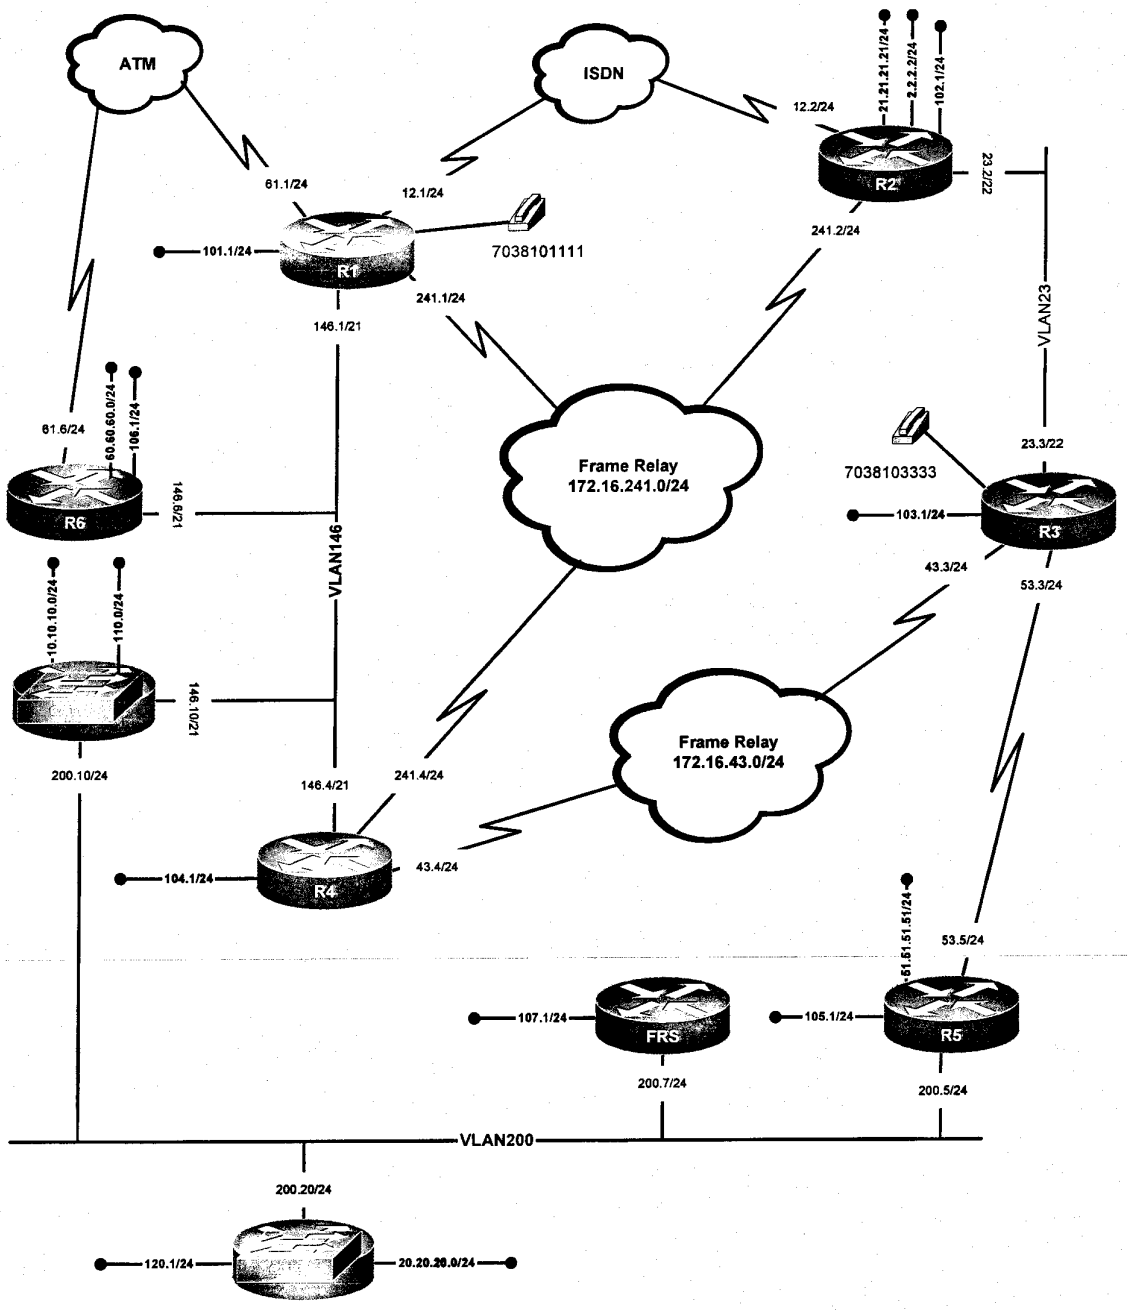

5.1 Serial Interfaces

5.1.1 Use the minimum number of PVC's to comprise the subnets 172.16.241.0/24 and 172.16.43.0/24
5.1.2 Configure logical interfaces on all Frame Relay serial interfaces

5.2 Catalyst Configuration

5.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table
5.2.2 Provide a solution ensuring link integrity on the interfaces Fa0/13 and Fa0/14

5.3 ATM

5.3.1 Configure the PVC 100/100 between R1 and R6
5.3.2 Use the minimum overhead encapsulation for IP connectivity
5.3.3 Assign IP addresses given in the diagram

5.4 OSPF

5.4.1 Configure OSPF area 0 on the subnet 172.16.43.0/24 and on the subnet 172.16.53.0/24
5.4.2 Elect a DR on R3 for all subnets in area 0. Use the default network type on the subnet 172.16.43.0/24
5.4.3 Configure OSPF area 23 on the subnet 172.16.23.0/22
5.4.4 Do not elect a DR on the subnet 172.16.23.0/22
5.4.5 Do not allow external routing information of any type into area 23.
5.4.6 Advertise network 172.16.103.0/24 in OSPF area 103

5.5 RIP

5.5.1 Configure only RIP version1 between R5 and CAT2
5.5.2 Configure only RIP version2 between CAT1 and FRS
5.5.3 In case of a serial link failure between R3 and R5, R5 should forward locally sourced traffic to FRS. DLSW traffic is excluded from this requirement

5.6 EIGRP

5.6.1 Configure EIGRP AS 100 on the subnet between R1, R6, CAT1 and R4
5.6.2 Configure EIGRP AS 10 on the ATM links between R1 and R6
5.6.3 Advertise the Loopback interfaces 172.16.101.0/24 and 172.16.106.0/24 into EIGRP10.

5.7 ISIS

5.7.1 Configure area 49.0001 and System ID 1111.1111.1111 on router R1
5.7.2 Configure area 49.0001 and System ID 2222.2222.2222 on router R2
5.7.3 Configure area 49.0001 and System ID 4444.4444.4444 on router R4
5.7.4 Restrict ISIS adjacencies to minimally required level over the Frame Rely media
5.7.5 Advertise networks 172.16.102.0/24 and 172.16.104.0/24 assigned to Loopback interfaces into ISIS

5.8 BGP

5.8.1 Configure BGP AS 600 on router R6
5.8.2 Configure BGP AS 10 on router CAT1
5.8.3 Configure BGP AS 140 on routers R1 and R4
5.8.4 Configure AS 2000 on router R2 and AS1000 on R3, R5 and FRS
5.8.5 Routers R2, R3, R5 and FRS belong to AS100
5.8.6 Configure AS 20 on CAT2
5.8.7 Peer AS 600 with both routers of AS140
5.8.8 Peer AS 10 with both routers of AS140
5.8.9 Peer AS 140 and AS100 between R1 and R2 as well as R4 and FRS
5.8.10 Do not allow a full mesh to be configured within AS 1000
5.8.11 Peer AS20 with FRS and R5
5.8.12 Advertise networks 10.10.10.0/24, 20.20.20.0/24 and 60.60.60.0/24 from their respective AS's. See the diagram
5.8.13 Make sure AS 600, AS 10 and AS20 load balance traffic between the above specified networks
5.8.14 You may use "no synchronization" in this exam

5.9 DLSW

5.9.1 Configure DLSW peer relationships between routers R2 and R3, R3 and R5 as well as R5 and FRS
5.9.2 Provide the solution allowing direct communications between NetBIOS workstations on R2 and FRS sites.
5.9.3 No additional "configured" DLSW peer relationships are allowed
5.9.4 Routers R2, R3, R5 and FRS must be involved in the forwarding of DLSW traffic.

5.10 ISDN

5.10.1 Voice traffic should be forwarded across the Frame-Relay if the link is up.
5.10.2 If the Frame relay link is down, forward traffic to 172.16.12.1
5.10.3 Do not run any routing protocol on the ISDN link

5.11 VPN

5.11.1 Configure an IP VPN between customer networks 21.21.21.0/24 on R2 and 51.51.51.0/24 on R5
5.11.2 Do not advertise networks 21.21.21.0/24 on R2 and 51.51.51.0/24 into any IGP used in this scenario. Rather, create a separate OSPF area 0 for the VPN between R2 and R5 and form an OSPF adjacency between routers R2 and R5
5.11.3 Advertise subnet 21.21.21.0/24 in OSPF area 20 and subnet 51.51.51.0/24 in area 50
5.11.4 Drop corrupted and out-of-order VPN packets
5.11.5 All tunnel packets should carry the ID key 10

5.12 Security

5.12.1 Authenticate EIGRP AS100 adjacencies with the string "nmceigrp". This key should be used from midnight Jan 1 2003 to midnight Jan 1 2006

5.12.2 Authenticate RIP version 2 with the string "nmcrip". This key should be used for 72 hours.

5.13 VOICE

5.13.1 The phone with the number 7038101111 is connected to port 1/0/0 on router R1
5.13.2 The phone with the number 7038103333 is connected to port 1/0/0 on router R3
5.13.3 Configure VOIP between R3 and R1 so that only R3 can call R1. Use the 172.16.101.1 IP address for this peer relationship
5.13.4 Forward voice traffic to R2 only. Mark it as immediate. Do not mark traffic within dial-peer configuration mode.

5.14 QOS

5.14.1 Overwrite ip precedence setting for voice traffic from immediate to critical on R2
5.14.2 Guarantee 50,000 bps with additional bursts of 10,000 bps for the DLSW communications between NetBIOS workstations on R2 and NetBIOS workstations on FRS
5.14.3 Make sure you see the reservations made dynamically along the DLSW forwarding path.

5.15 Catalyst Specialties

5.15.1 Three workstations using a multicast application tuned for traffic destined to the group 225.5.5.5 are connected to ports Fa0/15 – 17 on CAT2
5.15.2 Allow the switch to remove the port from the multicast distribution as soon as the application is terminated
5.15.3 The multicast application uses IGMP version 2

5.16 Address Administration

5.16.1 Loopback 2.2.2.2/24 is not advertised by any dynamic routing protocol
5.16.2 Ping devices on subnet 172.16.146.0/24 from 2.2.2.2. The packets should appear on VLAN 146 as if they were sourced from the subnet on VLAN 23
5.16.3 Ping the devices on subnet 172.16.200.0/24 from 2.2.2.2. The packets should appear on VLAN 200 as if they were sourced from the subnet 172.16..241.0/24

5.17 Multicast

5.17.1 Configure PIM Dense Mode on all interfaces connected to VLAN200
5.17.2 Configure PIM Dense Mode on all interfaces connected to VLAN146
5.17.3 Configure PIM Dense Mode between R3, R5 and R4
5.17.4 Configure PIM Dense Mode on the subnet 172.16.241.0/24 between R1 and R4
5.17.5 Join any Loopback interface on each router participating in PIM Dense Mode to group 225.5.5.5
5.17.6 Generate traffic (172.16.102.1, 225.5.5.5)
5.17.7 Multicast traffic to VLAN 200 should be forwarded by R5
5.17.8 Multicast traffic to VLAN 146 should be forwarded by R1
5.17.9 R2 should receive the replies from all joined interfaces

Scenario 6. Multi-topic CCIE level exam

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 with their original mask
- Network 0.0.0.0/0 should not appear in any routing table (show ip route)
- Do not create new interfaces to fulfill IGP requirements, do not summarize
- Do not use backup interface in the ISDN section
- 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. The Voice Section is excluded from this restriction.

VLAN Configuration Table

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

Attention For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15

Scenario 6

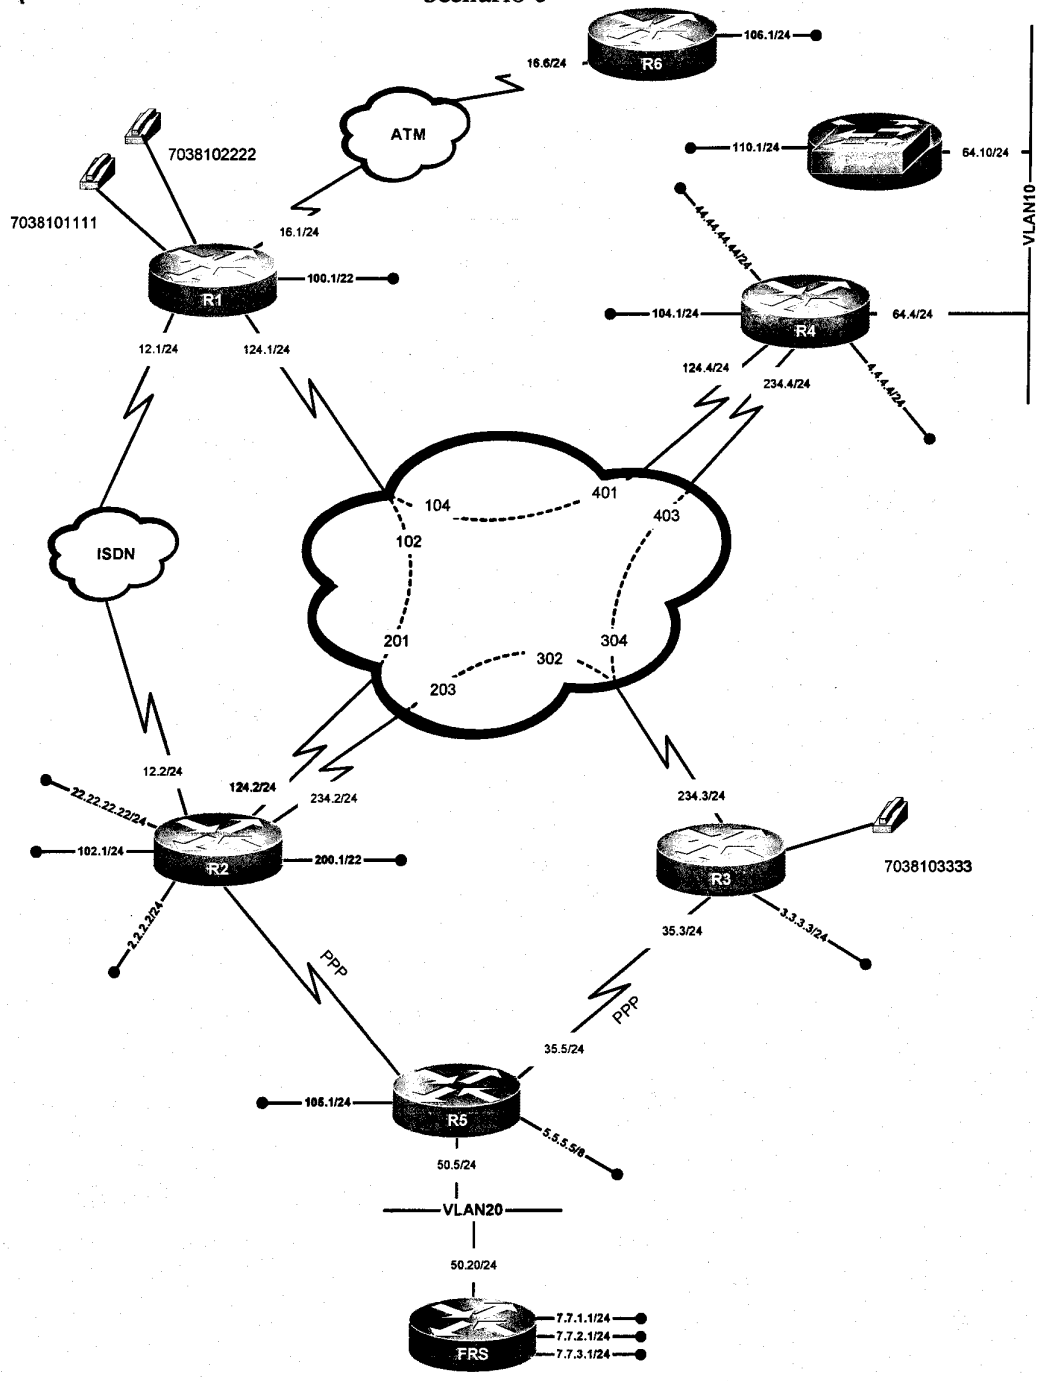

6.1 Serial Interfaces

6.1.1 Configure point-to-point logical Frame-Relay interfaces whenever possible
6.1.2 Otherwise configure Frame-Relay multipoint interfaces
6.1.3 Configure PPP encapsulation on the serial links between R2 and R5 as well as R3 and R5
6.1.4 Do not assign any explicit IP addresses to the link between R2 and R5
6.1.5 Change the MTU size only on R5 on the link to R3 to 2500. Do not change the MTU size on R3.

6.2 Catalyst Configuration

6.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table

6.3 ATM

6.3.1 Configure PVC 100/100 between R1 and R6
6.3.2 Make sure you can send the cells at the Peak Cell Rate of 256 Kbits/sec at any time, for any duration
6.3.3 Assign the IP addresses given on the diagram

6.4 OSPF

6.4.1 Configure OSPF area 0 on the links between R2 and R5
6.4.2 Configure OSPF area 10 on the subnet 172.16.50.0/24
6.4.3 Advertise network 7.7.1.0/24 as external to OSPF, network 7.7.2.0/24 in area 10 and 7.7.3.0/24 as any area you would like
6.4.4 Do not allow OSPF external networks in area 10 from R2
6.4.5 Advertise 172.16.105.0/24 in OSPF area 105 on router R5

6.5 RIP

6.5.1 Configure only RIP version 2 between routers R2, R3 and R4
6.5.2 Advertise networks 2.2.2.0/24 and 4.4.4.0/24 in RIP
6.5.3 R3 should not advertise networks learned from its Frame Relay interface back out the same interface
6.5.4 Make sure R4 and R2 still can exchange RIP routing information
6.5.5 Configure RIP version 1 only between R4 and CAT1

6.6 EIGRP

6.6.1 Configure EIGRP AS 100 between routers R1, R2, R4 and R6
6.6.2 Advertise networks 22.22.22.0/24 and 44.44.44.0/24 into EIGRP AS 100, as well as 172.16.100.0/22 and 172.16.106.0/24
6.6.3 R1 should not advertise networks learned from its Frame Relay interface back out the same interface
6.6.4 Make sure R2 and R4 still can exchange EIGRP routing information
6.6.5 Configure EIGRP AS 200 on the ISDN link

6.7 ISIS

6.7.1 Configure NET 49.0003.3333.3333.3333.00 on R3
6.7.2 Configure NET 49.0005.5555.5555.5555.00 on R5
6.7.3 Advertise networks 3.3.3.0/24 and 5.5.5.0/8 as External networks

6.8 BGP

6.8.1 Configure AS 500 on router R5
6.8.2 Configure AS 100 on routers R1, R2, R3,R4 and CAT1
6.8.3 Originate prefix 22.22.22.0/24 into BGP with the IGP origin code on R2
6.8.4 Advertise it to other peers of AS 100 with the local preference of 300 and tag it with 100:500
6.8.5 Use the most efficient configuration technique within AS100. R4 should be selected for this purpose.
6.8.6 A full mesh and new AS numbers are not allowed in AS 100
6.8.7 Configure AS 600 on router R6
6.8.8 Configure BGP peer relationships between R6 and R1, R2 and R5 as well as R3 and R5
6.8.9 Advertise only network 7.7.1.0/24 on R5 into BGP. Do not use any filtering techniques. Do not use the network statement
6.8.10 Network 7.7.1.0/24 should show up as "D EX"" in the routing table of routers R1, R4 and R6
6.8.11 Network 7.7.1.0/24 should show up as "B" in the routing table of all other routers of AS 100
6.8.12 You can disable synchronization only on one router in this scenario

6.9 IOS Features

6.8.9 Limit community "nmc-exam6" to read the following MIB object with OID 1.3.6.1.2.1.15 on R6

6.10 ISDN

6.10.1 Configure the ISDN link as a backup to the Frame Rely link on 172.16.124.0/24
6.10.2 Make a call from R2 when the Frame Relay PVC goes down
6.10.3 Use PPP CHAP authentication with password "nmc"
6.10.4 The ISDN link should not stay up when Frame Relay connection is active

6.11 Router Maintenance

6.11.1 For security purposes, disable the router feature which allows you to generate a stream of ASCII characters via a telnet session to the chargen port on R6

6.12 Security

6.12.1 Do not allow R6 to initiate any IP traffic to the rest of the network with the exception of control IP traffic such as routing updates and ICMP packets
6.12.2 On R1, block all java applets except those originating from 172.16.106.0/24

6.13 VOICE

6.13.1 The number 7038101111 is associated with port 1/0/0 on router R1
6.13.2 The number 7038102222 is associated with port 1/0/1 on router R1
6.13.3 The number 7038103333 is associated with port 1/0/0 on router R3
6.13.4 R3 should be able to call 7038101111 and 7038102222 with a single dial-peer configuration
6.13.5 Mark all voice packets with IP precedence 5.
6.13.6 Make sure that voice traffic between R1 and R3 passes through only R2.

6.14 QOS

6.14.1 Allocate 60 Kbps for voice traffic between R3 and R1
6.14.2 Restrict the voice conversation to 24 Kbps
6.14.3 Allow the voice conversation to proceed only if the previously specified allocation succeed.

6.15 Catalyst Specialties

6.15.1 Accept all unicast traffic on the VLAN 20
6.15.2 Accept broadcast traffic up to 25% of the total amount of traffic on VLAN 20
6.15.3 Do not accept any multicast traffic

6.16 Address Administration

6.16.1 Configure Network Address Translation on R4 with the outside interface on VLAN10
6.16.2 When you ping from the 172.16.110.1 interface of CAT1 to the rest of the network on the other side of R4, the IP address 172.16.110.1 must be translated to 172.16.104.10

6.17 Multicast

6.17.1 Configure PIM Dense Mode multicast routing between R5 and R3, R3 and R2, R3 and R4, R4 and R1 as well as R2 and R1
6.17.2 ISDN is not part of the multicast distribution
6.17.3 Source multicast traffic to destination 226.6.6.6 from different interfaces on router R5
6.17.4 Join one of the Loopback interfaces of routers R3, R2, R4 and R1 to multicast group 226.6.6.6
6.17.5 Make sure you receive the replies on R5 from all joined interfaces
6.17.6 Make sure both serial links are utilized for the multicast distribution from R3 to R1

Scenario 7. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets in the diagram belong to network 172.16.0.0/16
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- Do not use "ip default-network"
- Do not use backup interface in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	Ethernet0	VLAN20
R3	FastEthernet0/1	VLAN20
R3	FastEthernet0/0	VLAN10 VLAN30
R4	Ethernet0	VLAN40
R5	-	-
R6	FastEthernet0/0	VLAN40
FRS	-	-
CAT1	-	VLAN40
CAT2	-	VLAN20

Attention **For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15**

Scenario 7

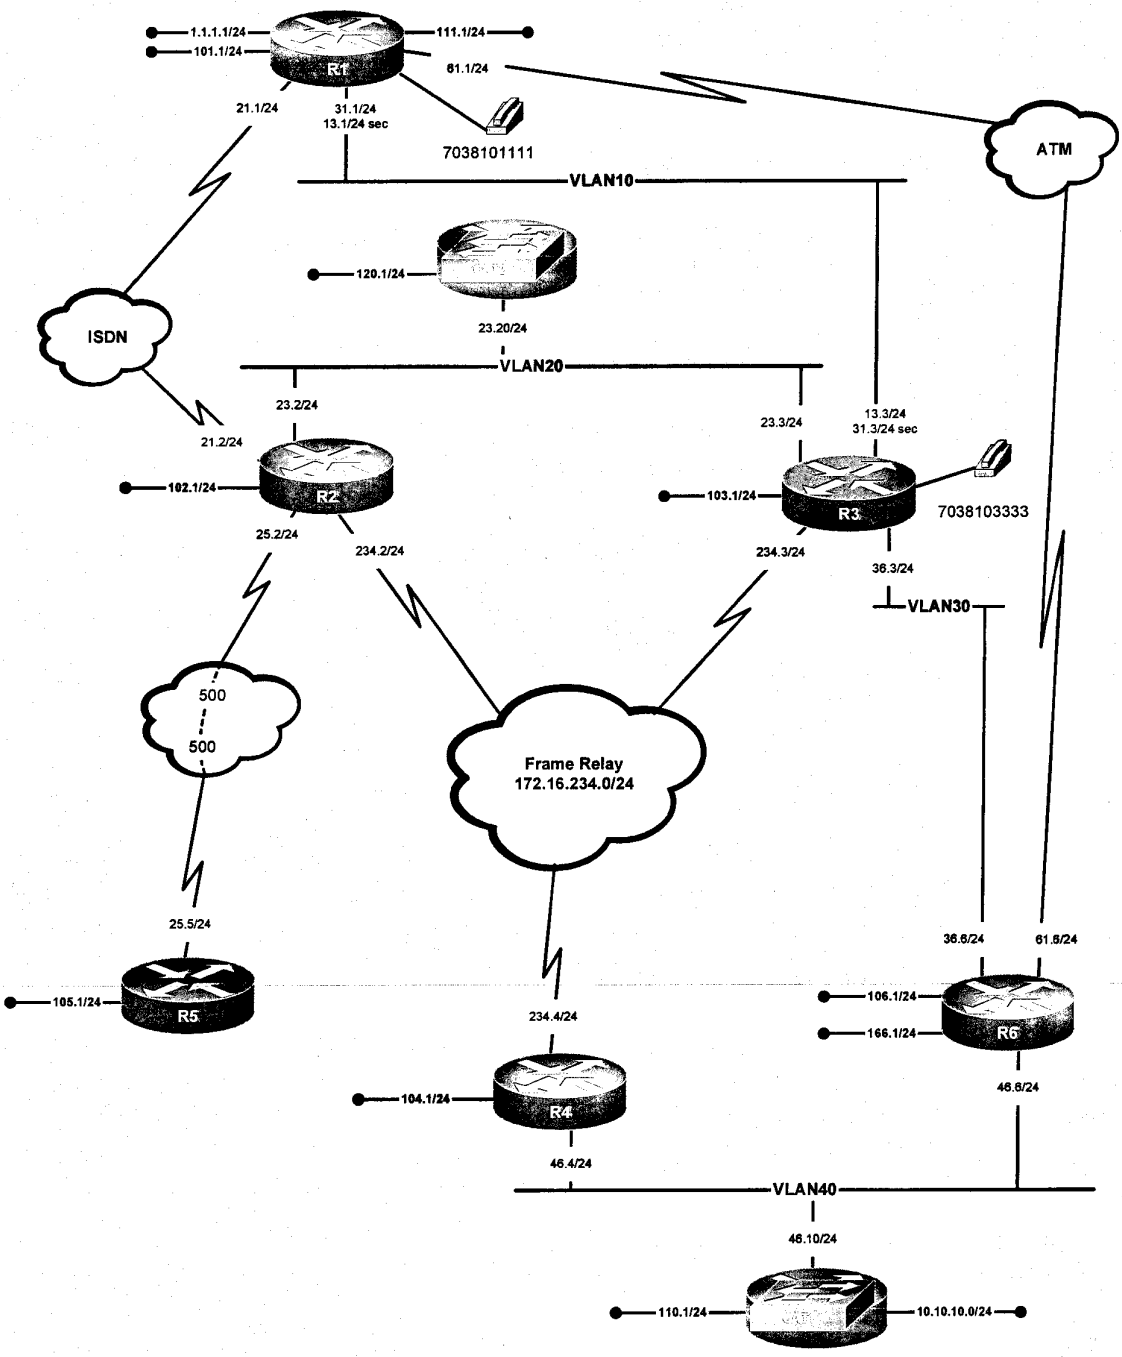

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

7.1 Serial Interfaces

7.1.1 Configure the minimal number of PVC's for the subnet 172.16.234.0/24
7.1.2 Configure logical interfaces on the subnet 172.16.234.0/24
7.1.3 Configure Frame Relay physical interfaces on the subnet 172.16.25.0/24

7.2 Catalyst Configuration

7.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table

7.3 ATM

7.3.1 Configure End System Identifier 1111.1111.1111 on router R1
7.3.2 Configure End System Identifier 6666.6666.6666 on router R6
7.3.3 Configure a logical interface on R1 and a physical interface on R6 for the subnet 172.16.61.0/24
7.3.4 Configure R6 to resolve its IP address to the appropriate data link identifier.

7.4 OSPF

7.4.1 Configure OSPF area 0 on subnet 172.16.25.0/24
7.4.2 Configure OSPF area 20 on all interfaces connected to subnet 172.16.23.0/24
7.4.3 Add the Loopback interface 172.16.120.1/24 of CAT2 to area 120
7.4.4 Add the Loopback interface 172.16.105.1/24 of R5 to area 105

7.5 RIP

7.5.1 Configure RIP version 2 between R1 and R2
7.5.2 Advertise RIP routing updates on the ISDN link only
7.5.3 Accept RIP updates only if they pass clear password authentication using the string "exam7"

7.6 EIGRP

7.6.1 Configure EIGRP AS 30 on the interfaces connected to VLAN30 only
7.6.2 Advertise the Loopback interface 172.16.103.1/24 in EIGRP AS 103 on R3
7.6.3 Advertise the Loopback interface 172.16.106.1/24 in EIGRP AS 106 on R6

7.7 ISIS

7.2.1 The following table displays the Area numbers and System IDs of the routers involved in this section:

Router	Area	System ID
R2	234	2.2.2
R3	234	3.3.3
R4	234	4.4.4
R6	46	6.6.6
R6	61	6.6.6
R1	61	1.1.1

7.7.2 Configure ISIS between R2, R4 and R3. Use the minimal router level configuration.
7.7.3 Configure ISIS between R4 and R6 on VLAN 40. Make it a backbone area between 234 and 46
7.7.4 Configure ISIS between R6 and R1 on the ATM link. Form a level-1 adjacency only.
7.7.5 Advertise networks 172.16.111.0/24 and 172.16.166.0/24 on R1 and R6 respectively.

7.8 BGP

7.8.1 Configure AS 100 on R1 and AS 300 on R3. Peer them over VLAN 10 without using any secondary IP addresses.
7.8.2 Advertise VLAN10 in both AS 100 and AS 300
7.8.3 Advertise the ISDN subnet 172.16.21.0/24 as well the ISDN end point addresses
7.8.4 Make sure the AS 100 does not destabilize the BGP domain if ISDN link goes up and down
7.8.5 Add routers R2 and R5 to AS 300
7.8.6 Peer all routers within AS300. Make sure they exchange NLRI with the local preference of 200. Do not use a route map to accomplish this task.
7.8.7 Configure AS 400 on router R4 and AS 600 on router R6
7.8.8 Configure AS 3000 on R5 as well and peer AS3000 and AS 400 between R5 and R4
7.8.9 Peer AS 600 and AS 300 between R6 and R3
7.8.10 Configure AS 1000 on CAT1
7.8.11 Peer AS 1000 with both AS 400 and AS 600
7.8.12 Advertise networks 1.1.1.0/24 and 10.10.10.0/24 from AS 100 and AS 1000 respectively. Provide transit service between these networks.
7.8.13 Do not originate any other networks into BGP.

7.9 IOS Features

7.9.1 Configure R5 so that you can reboot it with an snmp "set" operation from a network management station.

7.10 ISDN

7.10.1 Configure R1 and R2 to call each other
7.10.2 Use PPP encapsulation and the CHAP protocol with the password "exam7"
7.10.3 Use logical interfaces on ISDN link
7.10.4 Make sure ISDN does not stay up indefinitely

7.11 Router Maintenance

7.11.1 Configure at least 40 small buffers on R1
7.11.2 Configure not more than 170 big buffers on R1

7.12 Security

7.12.1 Router R5 will be connected to the Internet in the future.
7.12.2 Develop and apply the security protections against:

- ICMP redirect messages
- Broadcast packets

- Multicast packets
- Packets destined to the default network
- Packets sourced from the reserved Loopback addresses
- Packets from private non routed IP addresses

7.12.3 Do not use a standard or extended numbered access-list to accomplish this task

7.13 VOICE

7.13.1 The number 7038101111 is assigned to port 1/0/0 on router R1
7.13.2 The number 7038103333 is assigned to port 1/0/0 on router R3
7.13.3 R3 and R1 should be able to call each other
7.13.4 Select the digitizing technology that allows VOIP communication between the above mentioned phones to consume 6300 bits per second for the voice payload.
7.13.5 The network administrator decided to reduce the voice bandwidth by increasing the payload of each RTP packet to a value twice higher than default

7.14 QOS

7.14.1 Configure WFQ on the Frame Relay links between R2 and R3 with the following parameters:

- congestive discard threshold of 128 messages
- 512 dynamic queues
- 10 RSVP queues for future RSVP deployments

7.15 Catalyst Specialties

7.15.1 Configure VTP version 2 on CAT2
7.15.2 CAT1 will be a server in VTP domain "NMC" with the password "NMC" as well
7.15.3 Restrict the amount of flooded traffic on the trunks configured in this Scenario

7.16 Address Administration

7.16.1 CAT1 should prefer R4 as its primary gateway based on a CISCO developed protocol
7.16.2 In case router R4 fails, CAT1 should prefer R6
7.16.3 Be able to ping the rest of the network from 172.16.110.1 from CAT1

7.17 Multicast

7.17.1 Configure a PIM Sparse mode network. Do not use the CISCO proprietary method of distributing RP information through the multicast network
7.17.2 CAT2 is the root of the shared tree for the multicast groups 228.8.8.8 and 227.7.7.7
7.17.3 R2, R3, R4, R5 and R6 are members of the shared tree
7.17.4 ISDN and ATM links are not part of this exercise
7.17.5 Join one of the loopback interfaces on each member router to the multicast group 227.7.7.7
7.17.6 Configure each member router to process only Join and Prune messages destined for 172.16.120.1 and 227.7.7.7 group only
7.17.7 Generate (172.16.31.1, 227.7.7.7) traffic, make sure you receive responses from each multicast router

Scenario 8. Multi-topic CCIE level exam

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 with their original mask
- Do not use "ip default-network"
- Do not use backup interface in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	-	-
R2	Ethernet0	VLAN10
R3	FastEthernet0/1	-
R3	FastEthernet0/0	VLAN20
R4	Ethernet0	VLAN30
R5	Ethernet0	VLAN10
R6	FastEthernet0/0	VLAN30
FRS	Ethernet0	VLAN10
CAT1	-	VLAN20 VLAN40
CAT2	-	VLAN30 VLAN40

Attention For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15

Scenario 8

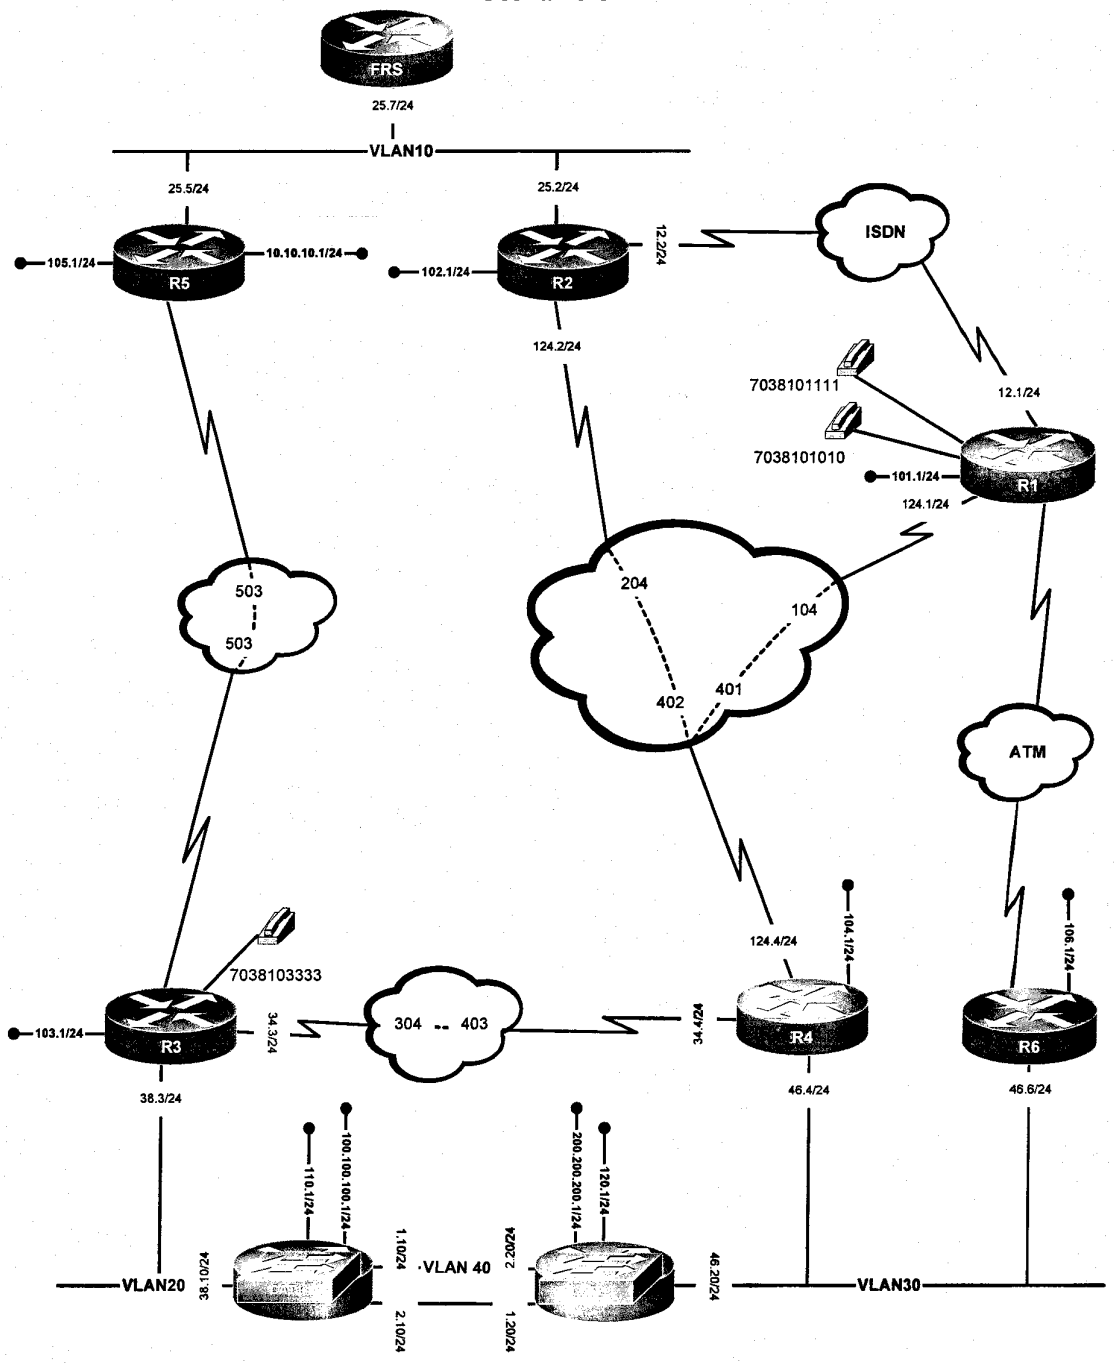

8.1 Serial Interfaces

8.1.1 Configure the Frame Relay links using only the PVC's displayed in the diagram
8.1.2 Configure logical interfaces on the Frame Relay link between R5 and R3. This link does not have an explicit IP subnet assigned to it.
8.1.3 None of the devices on the Frame Relay link between R3 and R5 should perform frame relay switching
8.1.4 Configure logical interfaces on the Frame Relay link between R3 and R4. Use CHAP authentication with the password "nmc"

8.2 Catalyst Configuration

8.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table
8.2.2 Use ISL for all trunking needs.
8.2.3 Make sure you allow only VLAN's configured in this Scenario to pass through the trunks
8.2.4 Assign 172.16.2.10/24 on the Fa0/14 interface of CAT1 and 172.16.1.10/24 on the Fa0/14 interface of CAT2

8.3 ATM

8.3.1 PVC 100/100 should be configured in compliance with IETF LLC specifications
8.3.2 The IP traffic should be flowing through this PVC only if CHAP authentication is passed with the password "nmc"

8.4 OSPF

8.4.1 Configure area 0 between R1, R2 and R4
8.4.2 Configure Area 10 between R2 and FRS. Use the OSPF "non-broadcast" network type
8.4.3 Configure redistribution between ISIS and OSPF on FRS
8.4.4 Make sure packets destined to the network 10.10.10.0/24 are forwarded to R5 from R2
8.4.5 Make sure the prefix 10.10.10.0/24 is in the routing tables of R1 and R4 with the next hop pointing to R2

8.5 RIP

8.5.1 Configure RIP version 2 between R3, R4, CAT1 and CAT2
8.5.2 Do not use secondary addresses in the RIP section
8.5.3 CAT1 and CAT2 should forward traffic directed to 172.16.103.1 via R4
8.5.4 R4 should forward traffic directed to 172.16.103.1 across the subnet 172.16.34.0/24
8.5.5 CAT1 and CAT2 should load balance the traffic destined to 172.16.110.0/24 and 172.16.120.0/24 respectively.
8.5.6 Do not use policy routing to accomplish the any of the tasks in this Section

8.6 EIGRP

8.6.1 Configure EIGRP AS 10 on the link between R1 and R6. Do not use the major network under the EIGRP routing process
8.6.2 Advertise the subnets 172.16.106.0/24 and 172.16.101.0/24 respectively
8.6.3 Configure EIGRP AS 10 on VLAN30 between R4 and R6. CAT2 should not be in EIGRP domain

8.7 ISIS

8.7.1 Configure ISIS between R3 and R5
8.7.2 Configure ISIS between R5 and FRS
8.7.3 Advertise the Loopback subnet 10.10.10.1/24 into ISIS

8.8 BGP

8.8.1 Configure AS 1000 on CAT1 and AS 2000 on CAT2
8.8.2 Advertise networks 100.100.100.0/24 and 200.200.200.0/24 respectively in the AS's 1000 and 2000. Do not advertise these networks into any IGP
8.8.3 Configure BGP AS 100 on routers R2, R3, R4, R5 and FRS
8.8.4 In AS 100, restrict your peering relationships to the following: R3 and R5, R4 and R2, FRS and R5, FRS and R2.
8.8.5 Peer AS 100 and AS 1000 between R3 and CAT1
8.8.6 Peer AS 100 and AS 2000 between R4 and CAT2
8.8.7 R2 must not accept any NLRI transiting through R5. R5 must not accept any NLRI transiting through R2. Do not use any filtering techniques.
8.8.8 Make sure you provide transit connectivity between 100.100.100.0/24 and 200.200.200/24 across AS 100 and each BGP speaker in AS 100 possesses both prefixes even if one of the given IBGP peers is lost (Test it with the neighbor shutdown command)
8.8.9 Configure the minimum number of IBGP peers to fulfill the criteria above. Do not introduce new AS numbers

8.9 VPN

8.9.1 Configure the IP address 10.10.1.2/24 on the Ethernet interface of router R2. Preserve the existing IP addressing provided in the Scenario diagram
8.9.2 Configure the IP address 10.10.2.3/24 on the interface of router R3 connected to VLAN 20. Preserve the existing IP addressing provided in the Scenario diagram
8.9.3 Configure 10.10.11.0/24 between the routers R2 and R3. Use RIP version 2 within this VPN. Make sure that only the three networks specified above are routed within the VPN
8.9.4 Do not use any generic encapsulation. Use IP into IP encapsulation

8.10 ISDN

8.10.1 Use a form of encapsulation which provides no authentication capabilities
8.10.2 Configure OSPF area 12 over the ISDN link
8.10.3 The ISDN link should not stay up indefinitely

8.11 Router Maintenance

8.11.1 Configure the following access characteristics for the aux line on all routers and the VTY line number 5 for the Catalyst 3550 switches:

- The device should supply the enable prompt whenever it is accessed
- The device should not prompt for a username and password
- The device should not terminate a session after the expiration of some period of time

8.11.2 Also, the router administrator wants to wait only for the minimal amount of time whenever a TELNET session with a mistyped (incorrect) IP address is initiated from one of the Catalysts

8.12 Security

8.12.1 A web server farm is located on VLAN20. The servers are using non-standard ports for the HTTP services, namely:

- 8080
- 8081
- 8082
- 8083

8.12.2 Also the servers with the IP addresses of 172.16.38.16 through 172.16.38.19 run a SMTP service on port 2525

8.12.3 Allow access to these services only through R3, but not CAT1

8.12.4 Also, configure security based on application layer criteria for the service mentioned above

8.13 VOICE

8.13.1 The number 7038101111 is assigned to port 1/0/0 on router R1

8.13.2 The number 7038101010 is assigned to port 1/0/1 on router R1

8.13.3 The number 7038103333 is assigned to port 1/0/0 on router R3

8.13.4 R3 should be able to call both phones on R1

8.13.5 Configure a prioritized list of CODECS:

- g723ar63
- g726r24
- g728
- g729r8

8.13.6 CODECS G729r8 and G728 should generate frames of 130 bytes. The others will operate at their default settings

8.13.7 All VOIP dial-peers should use this list of CODECS

8.14 QOS

8.14.1 Provide a QOS solution to the VPN users allowing traffic between R2 and R3 to be shaped to 30Kbit per second

8.14.2 Let the selected algorithm calculate the sustained bits per interval and the excess bits permitted in the first interval

8.15 Catalyst Specialties

8.15.1 Two workstations are connected to the following ports assigned to VLAN 30 on CAT2: Fa0/10 and Fa0/11

8.15.2 The network administrator should not allow any traffic between these two workstations.

8.15.3 However, these workstations should be able to freely communicate with the rest of the network
8.15.4 Do not use any filtering access lists

8.16 Address Administration

8.16.1 Allow TFTP, DNS, BOOTP and Time broadcasts to be propagated on the links between R3, R4, CAT1 and CAT2
8.16.2 The solution should protect against broadcast storms

8.17 Multicast

8.17.1 Configure Multicast PIM Sparse Mode routing between R3, R4, R5, R1 and R2
8.17.2 Use the Auto-RP discovery protocol
8.17.3 Configure R3 to be both the RP and Mapping agent
8.17.4 Join one of the loopback interfaces of routers R3, R4, R5, R1 and R2 to group 230.17.17.17
8.17.5 Ping the group 230.17.17.17 from CAT1
8.17.6 Make sure you get responses only from routers R3, R5 and R4

Scenario 9. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16
- Advertise Loopback interfaces with their original mask
- Do not use backup interface in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	-	-
R3	FastEthernet0/1	VLAN40
R3	FastEthernet0/0	VLAN20
R4	Ethernet0	VLAN40
R5	Ethernet0	VLAN30
R6	FastEthernet0/0	VLAN20
FRS	Ethernet0	VLAN40
CAT1	-	VLAN30
CAT2	-	VLAN10

Attention **For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15**

Scenario 9

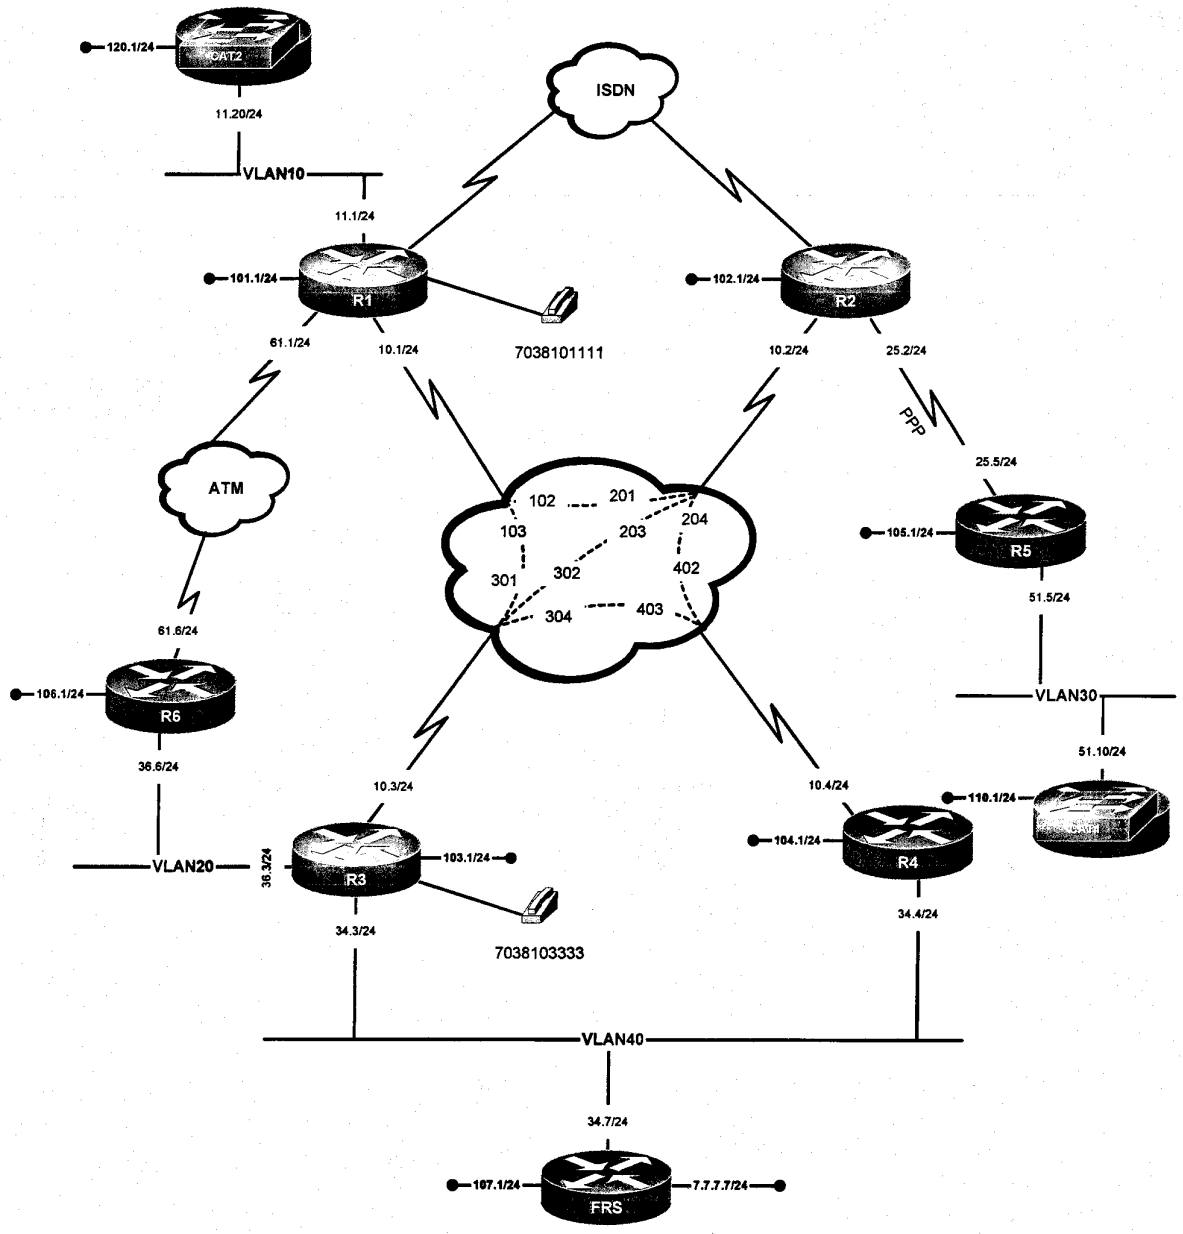

9.1 Serial Interfaces

9.1.1 Configure physical interfaces on the Frame Relay connections between R1, R2, R3 and R4
9.1.2 Configure PPP encapsulation on the Serial connection between R2 and R5

9.2 Catalyst Configuration

9.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table

9.3 ATM

9.3.1 Configure the following two PVC's , 100/100 and 200/200, between R1 and R6
9.3.2 These two PVC's must represent one IP subnet and load balance traffic. Do not use any bridging technique to accomplish this requirement
9.3.3 ATM connectivity must be established only if CHAP authentication is passed with the password "linkR1R6"

9.4 OSPF

9.4.1 Configure OSPF area 0 between routers R1, R2, R3 and R4 on the 172.16.10.0/24 subnet. Assure that the routers which can properly fulfill flooding requirements get elected to perform this operation. OSPF neighbors on this subnet should discover one another automatically.
9.4.2 Advertise Loopback subnet 172.16.102.0/24 into area 102 on R2
9.4.3 Configure OSPF between R2 and R5 on the subnet 172.16.25.0/24
9.4.4 Advertise Loopback subnet 172.16.105.0/24 into area 105 on R5
9.4.5 Configure OSPF area 51 on the subnet 172.16.51.0/24 between R5 and CAT1
9.4.6 Advertise Loopback subnet 172.16.110.0/24 into area 110 on CAT1
9.4.7 The ISDN link between R1 and R2 should be placed into area 12
9.4.8 Do not create any new interfaces on R2

9.5 RIP

9.5.1 Configure RIP version 1 on the routers R1 and CAT2

9.6 EIGRP

9.6.1 Configure EIGRP AS 1 on the ATM link between routers R1 and R6 as well as on the link 172.16.36.0/24 between R6 and R3.
9.6.2 Advertise the Loopback subnets 172.16.103.0/24 and 172.16.106.0/24 into EIGRP AS1

9.7 ISIS

9.7.1 Configure ISIS between routers R3, R4 and FRS. Use the following NET's:

- R3 49.7777.3333.3333.3333.00
- R4 49.7777.4444.4444.4444.00
- FRS 49.7777.7777.7777.7777.00

9.7.1 Form only a Level-2 adjacency between R3 and R4
9.7.2 Form only a Level-1 adjacency between R4 and FRS
9.7.3 Elect R3 and FRS to be DR routers
9.7.4 Advertise 172.16.107.0/24 in ISIS on FRS

9.8 BGP

9.8.1 Configure AS 700 on router FRS and advertise network 7.7.7.0/24
9.8.2 Configure AS 100 on routers R1, R2, R3 and R4
9.8.3 Configure R3 and R4 to be in one confederation
9.8.4 Configure R1 and R2 in the other confederation
9.8.5 Peer these two confederations between R1, R3 and R2, R4
9.8.6 Peer AS 700 and 100 between FRS and R3, R4
9.8.7 On R3 and R4, create a bit bucket for all destinations not matching more specific prefixes
9.8.8 Forward all unknown traffic to FRS from AS100, if the network 172.16.107.0/24 is up
9.8.9 Discard all unknown traffic on R3 and R4 if one of the following two conditions is met:

- AS 700 is unreachable
- 172.16.107.0 is down

9.9 IOS Features

9.9.1 Configure CAT2 to be a TFTP server only for CAT1, FRS and R5 routers
9.9.2 The IOS image residing on CAT2 should be available as "catos.bin" to the TFTP catalyst clients.

9.10 ISDN

9.10.1 Do not configure any explicit IP addresses on the ISDN link
9.10.2 Use CHAP authentication with the password "exam9"
9.10.3 Do not let the ISDN link flap

9.11 NTP

9.11.1 Configure NTP on all routers. All routers should get their time from R1
9.11.2 Configure US Eastern Time and daylight saving time on the routers

9.12 Security

9.12.1 Configure an Access List on R3 which does explicitly the following:

- Only permit SMTP and POP3 e-mail to R1
- Only permit R6 to ping R2
- Only permit UDP ports 6000-7000 to R5
- Only allow WWW traffic to FRS
- Do not allow SNMP traffic to R4
- Do not allow Tacacs+ traffic to CAT1

9.13 VOICE

9.13.1 The number 7038101111 is assigned to port 1/0/0 on router R1
9.13.2 The number 7038103333 is assigned to the port 1/0/0 on router R3
9.13.3 Configure a VOIP relationship between routers R1 and R3. Make sure the voice traffic does not flow over the Frame Relay connection.
9.13.4 R1's voice port 1/0/1 is connected to the local office exchange reachable by 1(571)333
9.13.5 Only the last 4 digits prepended by 9 with one second delay are used to call within the office exchange

9.14 QOS

9.14.1 Guarantee 30 Kbit/sec for voice conversations and minimal delays.

9.15 Catalyst Specialties

9.15.1 The workstation running a packet Sniffer application is connected to port Fa0/23 on CAT2
9.15.2 Monitor the Ethernet interface of R2 on the workstation running the Sniffer application

9.16 Address Administration

9.16.1 Servers providing TCP services are located on VLAN 10 with IP addresses ranging from 172.16.11.100 through 172.16.11.110
9.16.2 Provide a solution allowing TCP/IP requests coming from R1 to be load balanced between the web servers

9.17 Multicast

9.17.1 Configure PIM Sparse Mode multicast routing between R1, R2, R3, R4, R5, R6, CAT1 and FRS
9.17.2 Use the BSR protocol for RP announcements
9.17.3 Configure R6 to be an RP and BSR candidate
9.17.4 Join one of the loopback interfaces of routers R1, R2, R3, R4, R5, R6, CAT1 and FRS to group 229.17.17.17
9.17.5 Ping the group 229.17.17.17 from CAT2
9.17.6 Make sure you get responses only from routers R1, R2, R3, R4, R5,R6 and FRS but not from CAT1. Configure R5 to accomplish this task. Do not use any solution based on access-lists
9.17.7 The ISDN link is not involved in this exercise

Scenario 10. Multi-topic CCIE level exam

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 with their original mask
- Do not use "ip default-network"
- Do not use backup interface in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN40
R2	Ethernet0	VLAN40
R3	FastEthernet0/1	VLAN30
R3	FastEthernet0/0	VLAN10 VLAN20
R4	Ethernet0	VLAN30
R5	Ethernet0	VLAN50
R6	-	VLAN10 VLAN20
FRS	-	VLAN60
CAT1	-	VLAN50 VLAN60
CAT2	-	VLAN20 VLAN60

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

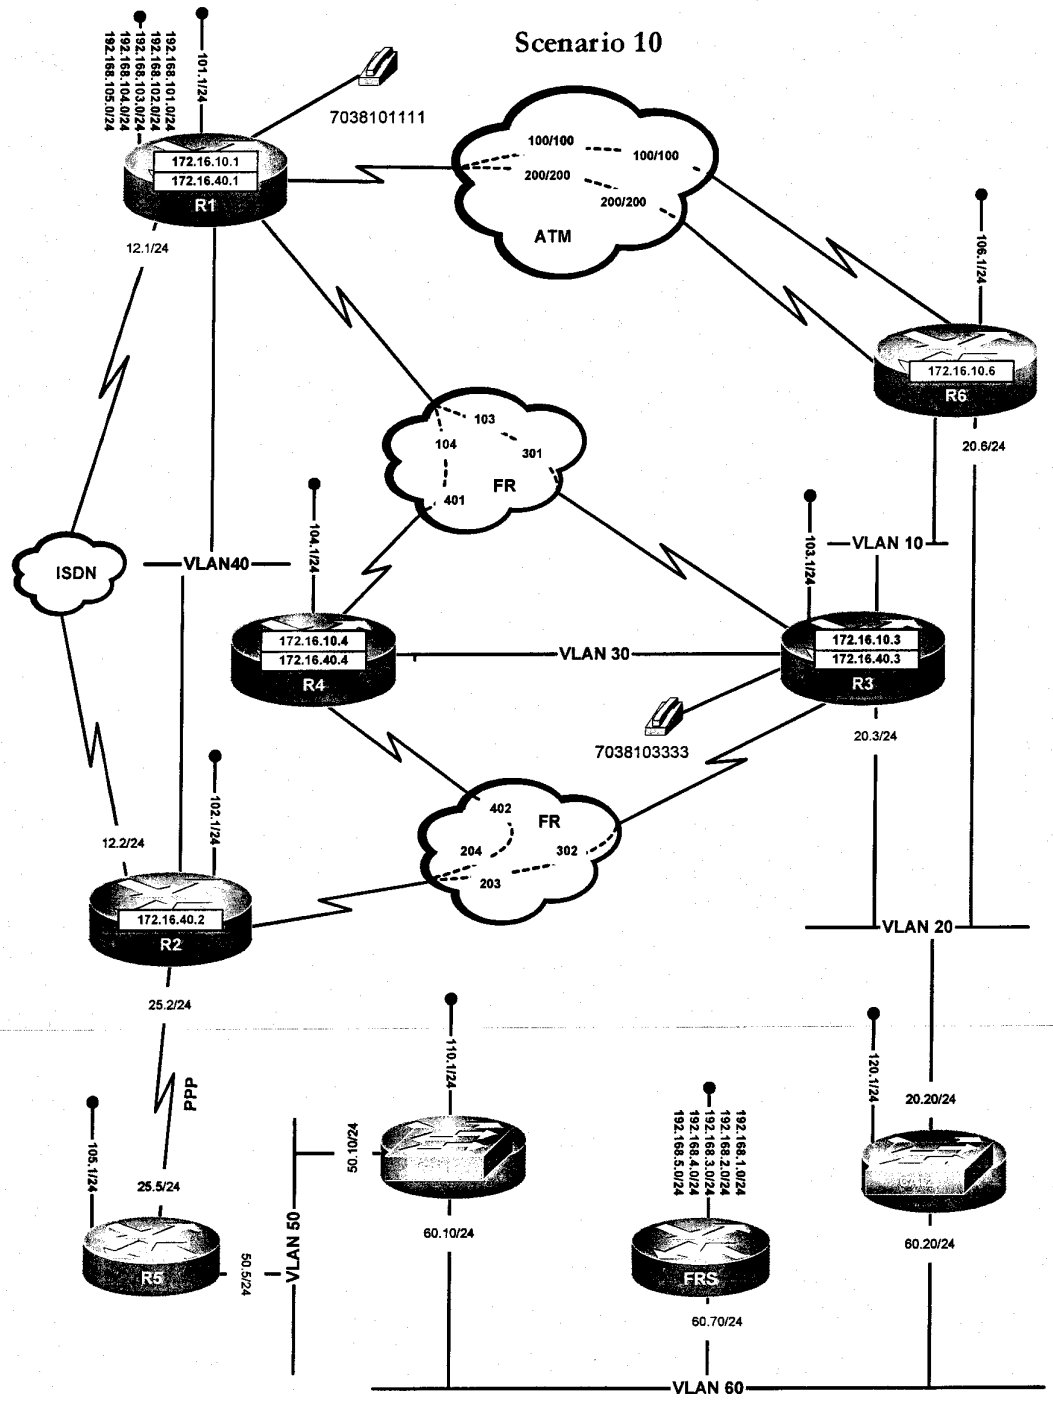

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

10.1 Serial Interfaces

10.1.1 Configure multipoint interfaces for all Frame Relay links displayed on the diagram
10.1.2 Make sure that only PVC's shown on the diagram are used in this scenario
10.1.3 Configure PPP on the link between R2 and R5
10.1.4 All serial interfaces (physical and logical) involved in this scenario must be configured for the logical bandwidth of a T1 line.

10.2 Catalyst Configuration

10.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table
10.2.2 Configure CAT1 as the root bridge for VLAN 10
10.2.3 Configure R1 as the root bridge for VLAN 40. Configure R2 as the root bridge if R1 is not available

10.3 ATM

10.3.1 Configure a physical ATM interface on R1 and multipoint interfaces on R6
10.3.2 Check the diagram for the PVC numbers
10.3.3 All ATM interfaces (physical and logical) involved in this scenario must be configured for the logical bandwidth of a T1 line

10.4 Address Administration

10.4.1 Configure subnet 172.16.10.0/24 between routers R1, R6, R3 and R4. This IP subnet is composed of the following data-links:

ATM	PVC 100/100 PVC 200/200
FRAME RELAY	PVC 103/301 PVC 104/401
ETHERNET	VLAN10

10.4.2 Configure subnet 172.16.40.0/24 between routers R1, R2, R3 and R4. This IP subnet is composed of the following data-links:

FRAME RELAY	PVC 204/402 PVC 203/302
ETHERNET	VLAN 40 VLAN 30

10.5 OSPF

10.5.1 Configure OSPF area 0 on subnet 172.16.40.0/24
10.5.2 Configure OSPF are 10 on subnet 172.16.20.0/24
10.5.3 Configure OSPF area 12 on the ISDN link.
10.5.4 Reduce refreshing and flooding of LSA's in the OSPF domain

10.5.5 Advertise Loopback 172.16.106.0/24 in area 60

10.6 RIP

10.6.1 Configure RIP version 2 between R1, R4, R3 and R6
10.6.2 Allow RIP updates to be exchanged only on the 172.16.10.0/24 subnet

10.7 EIGRP

10.7.1 Configure EIGRP AS 100 between CAT1 and CAT2
10.7.2 Configure EIGRP AS 200 between CAT1 and R5
10.7.3 Advertise network 172.16.120.0/24 in EIGRP AS 100
10.7.4 Advertise network 172.16.110.0/24 in EIGRP AS 200

10.8 ISIS

10.8.1 Configure ISIS area 49.9999 between R2 and R5
10.8.2 Advertise the Loopback networks 172.16.102.0/24 and 172.16.105.0/24 from R2 and R4 respectively into ISIS

10.9 BGP

10.9.1 Configure BGP AS 700 on CAT1, CAT2 and FRS
10.9.2 All other routers should be in BGP AS 100
10.9.3 Configure BGP AS 1 on routers R2, R3 and R5. Configure BGP AS 2 on routers R1, R4 and R6
10.9.4 A full-mesh is not allowed within either AS1 or AS2
10.9.5 Configure BGP peer relationships between AS 700 and AS 1 using peers CAT1 and R5 as well as CAT2 and R3
10.9.6 Configure BGP peer relationships between AS 1 and AS 2 using peers R2 and R4 as well as R3 and R6
10.9.7 Advertise the following networks from FRS :

- 192.168.1.0/24
- 192.168.2.0/24
- 192.168.3.0/24
- 192.168.4.0/24
- 192.168.5.0/24

10.9.8 AS 2 should prefer R2 as the next-hop for the above networks.
10.9.9 Advertise the following networks from R1:

- 192.168.101.0/24
- 192.168.102.0/24
- 192.168.103.0/24
- 192.168.104.0/24
- 192.168.105.0/24

10.9.10 Prefer incoming traffic to AS2 networks via R2. Do not use communities, local preference manipulation and filtering

10.9.11 Notify the console operator monitoring the peer relationship between CAT1 and R5 if R5 receives more then 3 prefixes from CAT1

10.10 IOS Features

10.10.1 Trap SNMP messages on the host 172.16.50.100 from router R5
10.10.2 Configure the community string "public" with Read/Write rights on R5
10.10.3 Configure remote monitoring on R5 with traps sent to previously defined SNMP community
10.10.4 The events should be trapped with the description "R5_trapped_events" and the owner is joedoe

10.11 ISDN

10.11.1 Configure PPP on the ISDN link between R1 and R2
10.11.2 Use password "nmc?doit" as a password for CHAP authentication
10.11.3 Use the OSPF demand circuit extension for the ISDN adjacency
10.11.4 Make sure the ISDN link remains stable.

10.12 NTP

10.12.1 Configure R6 to be a master clock with its stratum set to 5
10.12.2 Configure CAT2 and R3 in stratum 6, get the clock from R6, do not use "ntp server" and "ntp peer" commands
10.12.3 Configure R1 and R4 in stratum 7, get the clock from R6 and R3. R6 should be preferred. Do not use "ntp peer" command

10.13 Security

10.13.1 All IP communications are allowed on VLAN60 except telnet
10.13.2 Make sure that only the three devices connected to VLAN 60 (see the diagram) are able to communicate and pass transit traffic.
10.13.3 All non-IP traffic must be permitted on VLAN 60

10.14 VOICE

10.14.1 The number 7038101111 is assigned to port 1/0/0 on router R1
10.14.2 The number 7038103333 is assigned to port 1/0/0 on router R3
10.14.3 Configure Voice over Frame relay on the PVC 103/301 between these two phones
10.14.4 Make sure it does not disturb other traffic involved in this scenario

10.16 QOS

10.16.1 On the port of the Catalyst connected to R5, configure 4 queues according to the following classification:

- Queue1 for Class of Service 1 and 3
- Queue2 for Class of Service 0 and 2
- Queue3 for Class of Service 6 and 7
- Queue4 for Class of Service 4 and 5

10.16.2 Make the switch empty one of the queues first
10.16.3 Allocate the following percentage of switch attention to the queues:

- Queue 1 – 30%
- Queue 2 – 10%
- Queue 3 – 40%
- Queue 4 – 20%

10.17 Catalyst Specialties

10.17.1 The Network Administrator wants to allow only authorized Windows XP workstations on CAT2. The workstations are connected to ports Fa0/10 and Fa0/11 on CAT2.
10.17.2 The Windows XP workstations are authenticated with the imaginary RADIUS server (172.16.20.100) using the Extensible Authentication Protocol (EAP) extension. Use the authentication key "nmc" between the switch and RADIUS server.
10.17.3 Also, some workstations share a wireless access point connected to Fa0/12. They must be authenticated as well.

10.18 Multicast

10.18.1 VLAN 20 is comprised of imaginary Catalyst switches equipped with old supervisory modules with limited multicast processing capabilities
10.18.2 Select the device on VLAN 20 to control the multicast distribution to the workstations
10.18.3 All devices on VLAN 20 support IGMP version 2

Scenario 11. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16
- Do not introduce any new IP addresses
- Do not use any static routes
- Advertise Loopback interfaces with their original mask and do not change this mask
- Do not use 0.0.0.0/0 in the RIP domain of this network
- Do not use backup interface in the ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN20
R2	-	-
R3	FastEthernet0/1	VLAN30
R3	FastEthernet0/0	VLAN20
R4	Ethernet0	VLAN40
R5	Ethernet0	VLAN10
R6	FastEthernet0/0	VLAN10
FRS	Ethernet0	VLAN20
CAT1	-	VLAN30 VLAN40
CAT2	-	VLAN30 VLAN40

Attention For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15

Scenario 11

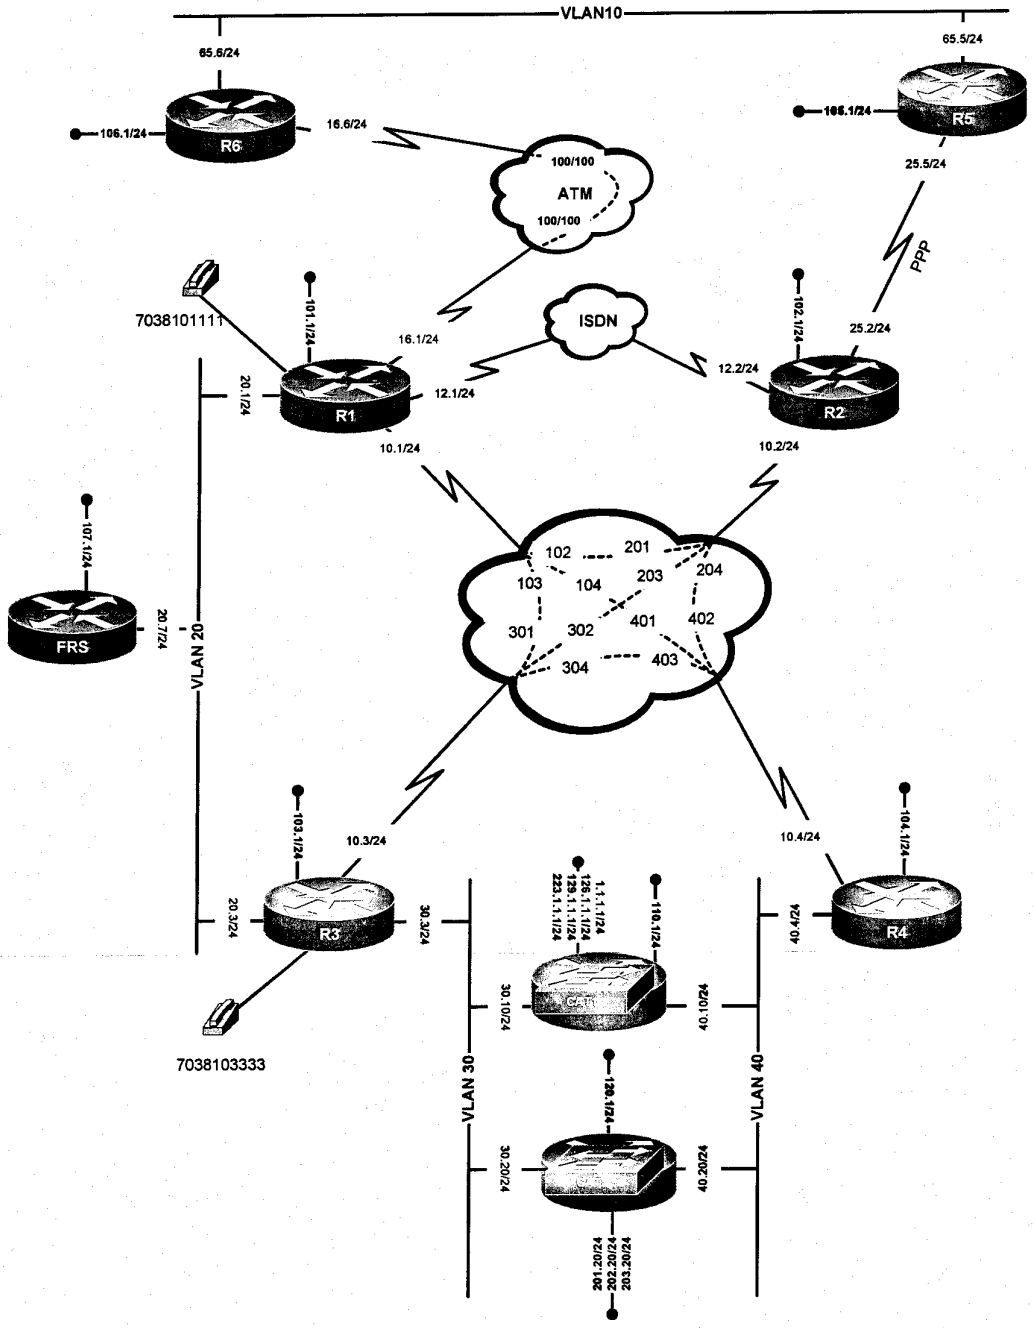

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

11.1 Serial Interfaces

11.1.1 Configure multipoint interfaces for all Frame Relay links displayed on the diagram
11.1.2 Use the Inverse ARP configuration on the Frame Relay subnet
11.1.3 Make sure there is full connectivity within the 172.16.10.0/24 network
11.1.4 Configure PPP on the link between R2 and R5

11.2 Catalyst Configuration

11.1.1 Configure VLAN's according to the diagram and the "VLAN configuration" table
11.1.2 Use transparent VTP mode and the ISL trunking protocol

11.3 ATM

11.3.1 Configure PVC 100/100 between routers R1 and R6
11.3.2 Select the service category which is most suitable for voice and video, but allows for bursts
11.3.3 Use the following parameters:

- Peak Cell Rate 128 Kbit/sec
- Sustained Cell Rate 128 Kbit/sec
- Burst Size 1 Kbyte

11.4 Address Administration

11.4.1 Workstations are connected to VLAN 40 via ports Fa0/18 and Fa0/19 of CAT1 and ports Fa0/18 of Fa0/19 of CAT2
11.4.2 The DHCP server is 172.16.65.100
11.4.3 The Network Administrator would like to assign IP addresses from different address scopes to the workstations attached to VLAN 40 based on the port numbers they are connected to.

11.5 OSPF

11.5.1 Configure OSPF area 0 on the Frame Relay cloud between R1, R2, R3 and R4
11.5.2 Make sure all routers attain a FULL adjacency state on the Frame Relay subnet
11.5.3 If one or two PVC-s fail, retain the adjacencies between all routers connected to Frame Relay subnet
11.5.4 If one or two PVC-s fail, retain full connectivity between all routers connected to Frame Relay subnet
11.5.5 Configure OSPF area 16 on the ATM link between routers R1 and R6. Use the default network type. Originate the hello packets from R6
11.5.6 Advertise Loopback networks 172.16.101.0/24 into OSPF area 101 and 172.16.103/24 into OSPF area 103 respectively

11.6 RIP

11.6.1 Configure RIP version 1 between R1, R3 and FRS

11.7 EIGRP

11.7.1 Configure EIGRP AS 30 on VLAN 30 between routers R3, CAT1 and CAT2
11.7.2 Configure EIGRP AS 40 on VLAN 40 between routers R4, CAT1 and CAT2
11.7.3 Advertise the following networks from CAT2 into EIGRP:

- 172.16.201.0/24
- 172.16.202.0/24
- 172.16.203.0/24

11.7.4 Router R3 and R4 should receive the most efficient summary for the above mentioned networks. Accomplish this by configuring CAT1

11.8 ISIS

11.8.1 Configure ISIS between R6 (NET 49.1000.6666.6666.6666.00) and R5 (NET 49.1000.5555.5555.5555.00) on VLAN 10.
11.8.2 Configure ISIS between R5 and R2 (NET 49.1000.2222.2222.2222.00)
11.8.3 Configure authentication between ISIS speakers. Use a method that transports the authentication string in the ISIS LSP. Use a different ISIS adjacency level on each segment.
11.8.4 Configure a mutual redistribution between ISIS and OSPF on R6 and R2. R5 should send packets via R6 for non-ISIS originated networks
11.8.5 Advertise the Loopback network 172.16.102.1/24 into ISIS. R1 should use ATM link to reach this network

11.9 BGP

11.9.1 Configure the BGP AS numbers according to the following table:

AS Number	Routers
AS 1000	CAT1 CAT2
AS 300	R3
AS 400	R4
AS100	R1 R2 R5 R6

11.9.2 Peer the AS's using the following scheme:

- AS 1000 and AS 300 between CAT1 and R3
- AS 1000 and AS 400 between CAT1 and R4
- AS 300 and AS 100 between R3 and R1
- AS 400 and AS 100 between R4 and R2

11.9.3 Do not use route reflectors and confederations within AS 100. Use the "no sync" method.
11.9.4 Provide the solution that allows the load sharing of outbound traffic from AS 100 to networks advertised from AS 1000:

- 1.1.1.0/24
- 126.1.1.0/24
- 129.1.1.0/24
- 223.1.1.0/24

11.9.5 Generalize the solution for the full IP address space

11.10 IOS Features

11.10.1 An administrator is concerned that the NVRAM size will not suffice for future configuration files on R6.
11.10.2 When you telnet from R3, the telnet client should not output any ip addresses
11.10.3 Suppress all telnet messages when you telnet from R3 to CAT1

11.11 ISDN

11.11.1 Use CHAP authentication on the ISDN link. The password to use is "^R^Snmc", where:

- ^R is Control-R
- ^S is Control-S

11.11.2 Bring the ISDN line up based on the absence of 172.16.102.0/24
11.11.3 R1 is a calling party

11.12 NTP

11.12.1 Configure R3 as a master with the stratum 4
11.12.2 Configure CAT1 as a client
11.12.3 Make sure NTP messages are exchanged between 172.16.103.1 and 172.16.110.1

11.13 Security

11.13.1 ICMP packets between networks 172.16.101.0/24 and 172.16.102.0/24 should be encrypted
11.13.2 The payload of encrypted packets should contain minimal IP overhead information
11.13.3 Use a pre-shared method of authentication using the password "nmc"

11.14 Multicast

11.14.1 Join the interfaces attached to the following VLAN to the multicast group 224.11.11.11 :

- VLAN 20

11.14.2 Ping group 224.11.11.11 from the Ethernet interface of R4 and make sure R4 receives replies from all joined interfaces
11.14.3 Multicast routing is not allowed between the VLAN 20 and the source

11.15 QOS

11.15.1 Configure different thresholds for two types of traffic. One type (threshold 2) is represented by DSCP values of 20 through 29 and the other type (threshold 1) is represented by all remaining DSCP values
11.15.2 The following table displays the threshold assignment to the queues:

Queue 1	threshold 1	threshold 2
1	10%	100%
2	20%	100%
3	30%	100%
4	40%	100%

11.15.3 Configure WRED on the interface Gig0/1 of CAT1 based on the information given above. Do not use the tail drop algorithm.

11.16 Catalyst Specialties

11.16.1 Configure a Port Aggregation Protocol between CAT1 and CAT2 on the ports Fa0/13 and Fa0/14
11.16.2 Load balance traffic between CAT1 and CAT2 based on the destination MAC address

Scenario 12. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 151.10.0.0/16
- Do not introduce any new IP addresses
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- Do not use "ip default-network"
- Do not use dialer watch in The ISDN section
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	Ethernet0	-
R3	FastEthernet0/1	-
R3	FastEthernet0/0	-
R4	Ethernet0	VLAN10
R5	Ethernet0	-
R6	FastEthernet0/0	-
FRS	Ethernet 0	VLAN10
CAT1	-	-
CAT2	-	-

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 12

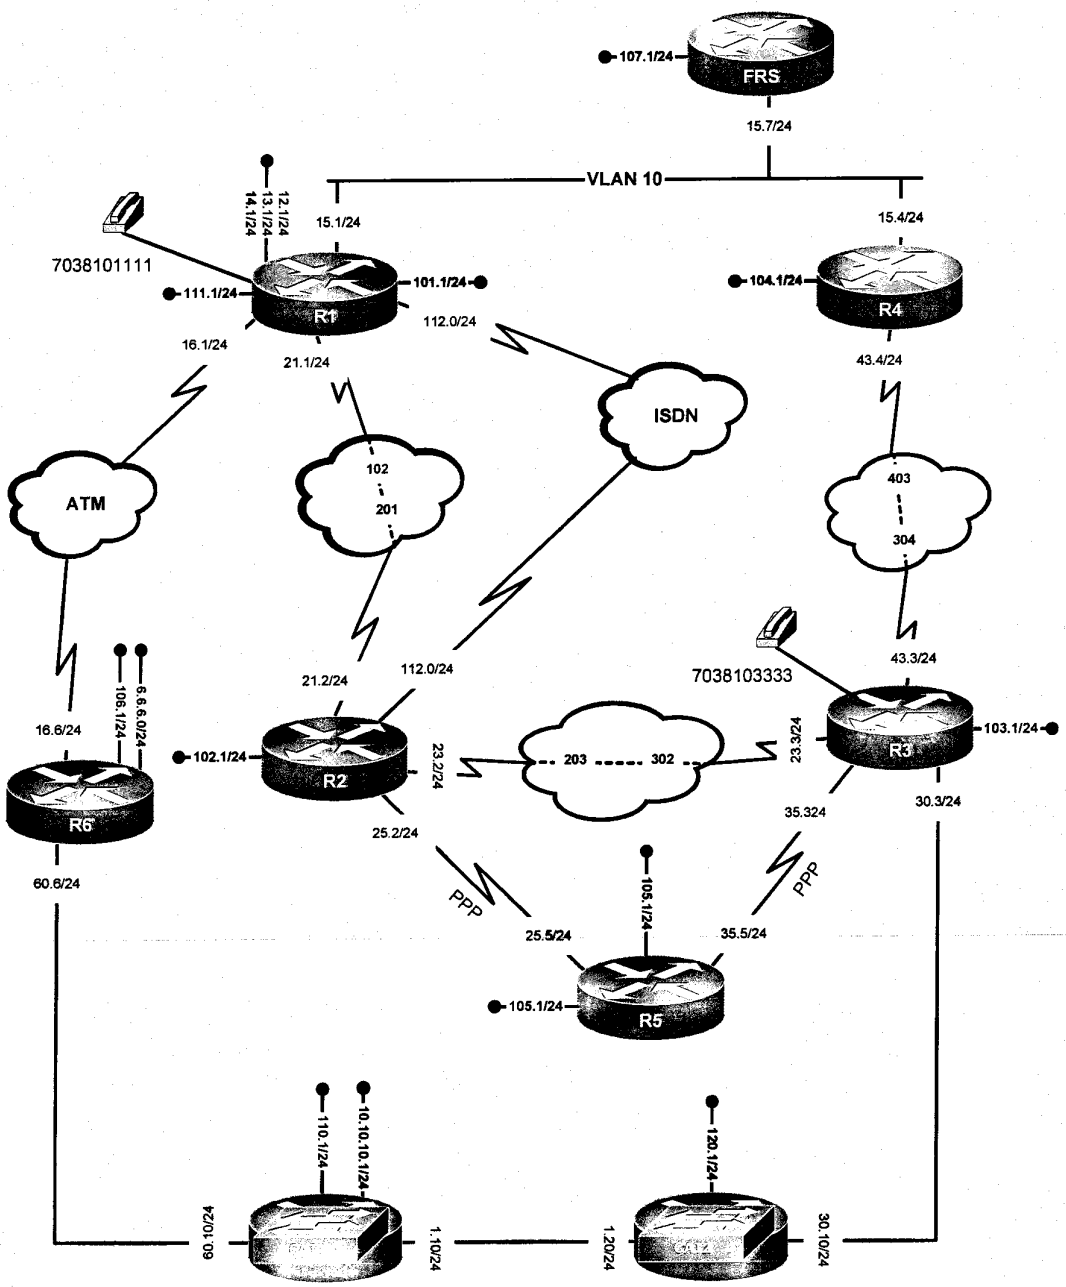

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

12.1 Serial Interfaces

12.1.1 Configure physical and logical point-to-point Frame Relay interfaces on R3
12.1.2 Configure a multipoint Frame Relay interface on R4
12.1.3 Configure two multipoint Frame Relay interfaces on R2
12.1.4 Configure a point to point Frame Relay interface on R1
12.1.5 Both links from R5 to R2 as well as R5 to R3 are configured as PPP links

12.2 Catalyst Configuration

12.2.1 Configure VLAN's according to the diagram and the "VLAN configuration" table
12.2.2 You may use ISL protocol for all trunks involved in this scenario.
12.2.3 Use interface FastEthernet 0/0 on R3 for this scenario

12.3 ATM

12.3.1 Configure ESI 1111.1111.1111.00 on R1
12.3.2 Configure ESI 6666.6666.6666.00 on R6
12.3.3 Configure variable bit rate non-real-time service on the ATM link between R1 and R6 with the following parameters:

- 64000 bps Peak Cell Rate
- 64000 bps Sustainable Cell Rate
- 1000 cells Minimum Burst Size
- 64000 bps Input Peak Cell Rate
- 64000 bps Input Sustainable Cell Rate
- 1000 cells Input Maximum Burst Size

12.4 Address Administration

12.4.1 The Network Administrator wants to distribute mail and HTTP services to different machines giving the appearance that the services are all reachable at one address
12.4.2 Configure the test bed using FRS, R1 and R4. Use the 151.10.15.100 IP address as virtual server IP address.
12.4.3 Originate packets from 151.10.103.1 to the SMTP port using the virtual server address. Make sure the packets are delivered to the following R1 IP address: 151.10.15.1.
12.4.4 Originate packets from 151.10.103.1 to the HTTP port using the virtual server address. Make sure the packets are delivered to following FRS IP address: 151.10.15.7.

12.5 OSPF

12.5.1 Configure OSPF area 0 on the ATM link, network 151.10.16.0/24
12.5.2 Configure OSPF area 21 on the interfaces connected to 151.10.21.0/24
12.5.3 Configure OSPF area 12 on the interfaces connected to 151.10.112.0/24
12.5.4 Advertise Loopback network 151.10.101.1/24 in OSPF area 101
12.5.5 Configure OSPF area 15 on the only R1 and FRS interfaces connected to 151.10.15.0/24

12.5.6 Advertise the following Loopback networks from R1 in area 15:

- 151.10.12.0/24
- 151.10.13.0/24
- 151.10.14.0/24

Summarize these networks into OSPF with the most efficient mask

12.5.7 Advertise Loopback network 151.10.106.0/24 in area 106
12.5.8 Advertise Loopback network 151.10.107.0/24 in area 107
12.5.9 Advertise Loopback network 151.10.111.0/24 in area 0

12.6 RIP

12.6.1 Configure RIP version 2 on R3, CAT1, CAT2 and R6
12.6.2 RIP updates must be exchanged only on the segments between these routers
12.6.3 Mutually redistribute RIP into the other routing protocols on R3 only
12.6.4 R6 should not receive prefixes from the other RIP speakers. Do not use passive interface and any route filtering technique based on prefixes

12.7 EIGRP

12.7.1 Configure EIGRP AS 100 on the Fame Relay link 151.10.23.0/24 between R3 and R2 as well as on the PPP links between R2 and R5; and R3 and R5.
12.7.2 Advertise network 151.10.105.0/24 from R5 in EIGRP
12.7.3 When redistributing routing information into EIGRP, use the default metrics:

- Bandwidth 1000
- Delay 100
- Reliability 255
- Load 1
- MTU 1500

12.7.4 Authenticate the EIGRP adjacency over Frame Relay using key "rs?nmc"
12.7.5 R5 should load balance traffic to EIGRP external networks using a 6:1 ratio favoring R2. Influence the routing decision by configuring R5 only.

12.8 ISIS

12.8.1 Configure ISIS on the following routers:

Router	NET
FRS	10.7777.7777.7777.00
R4	10.4444.4444.4444.00
R3	10.3333.3333.3333.00

12.8.2 Build a Level-2 network only. Do not allow a Level-1 adjacency anywhere

12.8.3 Advertise Loopback network 151.10.104.0/24 in ISIS
12.8.4 Mutually redistribute OSPF and ISIS on FRS
12.8.5 Mutually redistribute EIGRP AS 100 and ISIS on R3

Attention

- R2 should prefer R1 to reach networks 151.10.104.0/24 and 151.10.43.0/24
- R1 should forward packets destined to networks 151.10.104.0/24 and 151.10.43.0/24 via 151.10.15.4
- In the case where R2 loses connected networks to R1, R2 should forward packets destined to 151.10.104.0/24 and 151.10.43.0/24 via R3
- Do not use any filtering of routing information.

12.9 BGP

12.9.1 Configure AS 600 on router R6
12.9.2 Configure AS 1000 on CAT1 and CAT2
12.9.3 Peer AS 600 with the loopback interface 151.10.110.1 of CAT1
12.9.4 Advertise networks 6.6.6.0/24 and 10.10.10.0/24 from AS 600 and AS 1000 respectively

12.10 IOS Features

12.10.1 Configure the switching method based on IP flows on the Fa0/0 interface of R6
12.10.2 Send the flow statistics to the management workstation on 151.10.60.100
12.10.3 Generate BGP autonomous system-to-autonomous system traffic flow statistics as well.Send the data to 151.10.60.200

12.11 ISDN

12.11.1 Use CHAP authentication on the ISDN link. Password is "rsnmc"
12.11.2 Backup the Frame-Relay interface with ISDN. Change of the PVC state should trigger the ISDN call
12.11.3 ISDN should remain in standby when the PVC is Active
12.11.4 Do not rely on the OSPF protocol extensions
12.11.5 Make sure at any given time, IP traffic between 151.10.112.1 and 151.10.112.2 port 23 can traverse the ISDN link.

12.12 Policy Routing

12.12.1 All packets originating from 151.10.111.1 and destined to 151.10.105.1 must be forwarded to R6
12.12.2 On R6 all traffic coming from 151.10.111.1 must be forwarded to CAT1 with the TOS field set to Network Control. Do not use a numeric value.
12.12.3 Forward traffic to CAT1 only if it is in the CDP table

12.13 Security

12.13.1 On port Fa0/7 of CAT2, allow only FRS's IP address

12.14 QOS

12.14.1 Traffic with the IP precedence 7 arriving from R6 should be marked with the DSCP 25
12.14.2 Traffic arriving from CAT1 to CAT2 should be marked with DSCP 20
12.14.3 Traffic marked as DSCP 20 should be marked as IP precedence critical on R3

12.15 Catalyst Specialties

12.10.4 Configure two physical links between CAT1 and CAT2 in one IP subnet. Do not use bridging for this task.
12.10.5 Configure the virtual IP address 151.10.1.1 as a gateway on the subnet between CAT1 and CAT2

12.16 Multicast

12.16.1 Configure R2 as an RP. Use 151.10.102.1
12.16.2 Use the static method of RP configuration in this task
12.16.3 Configure PIM on the link between CAT2 R3, links between R3 and R2, the Frame-Relay link between R2 and R1 and on all interfaces on VLAN10
12.16.4 Join one of the loopback interfaces of each multicast router to group 225.12.12.12
12.16.5 Source multicast traffic from CAT1.
12.16.6 Make sure all joined interfaces reply.

Scenario 13. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16
- Do not introduce any new IP addresses
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- Do not use "ip default-network"
- Do not use dialer watch, backup interface in the ISDN section
- 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
- Networks 192.168.*.* are excluded from the reachability requirements
- Networks 1.1.1.0/24, 2.2.2.0/24 and 3.3.3.0/24 are not part of IGP and do not have to be pingable

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10 VLAN20
R2	Ethernet0	VLAN40
R3	FastEthernet0/1	VLAN50
R3	FastEthernet0/0	-
R4	Ethernet0	VLAN30
R5	Ethernet0	VLAN50
R6	FastEthernet0/0	VLAN40
FRS	Ethernet0	VLAN30
CAT1	-	VLAN20
CAT2	-	VLAN10

Attention **For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15**

Scenario 13

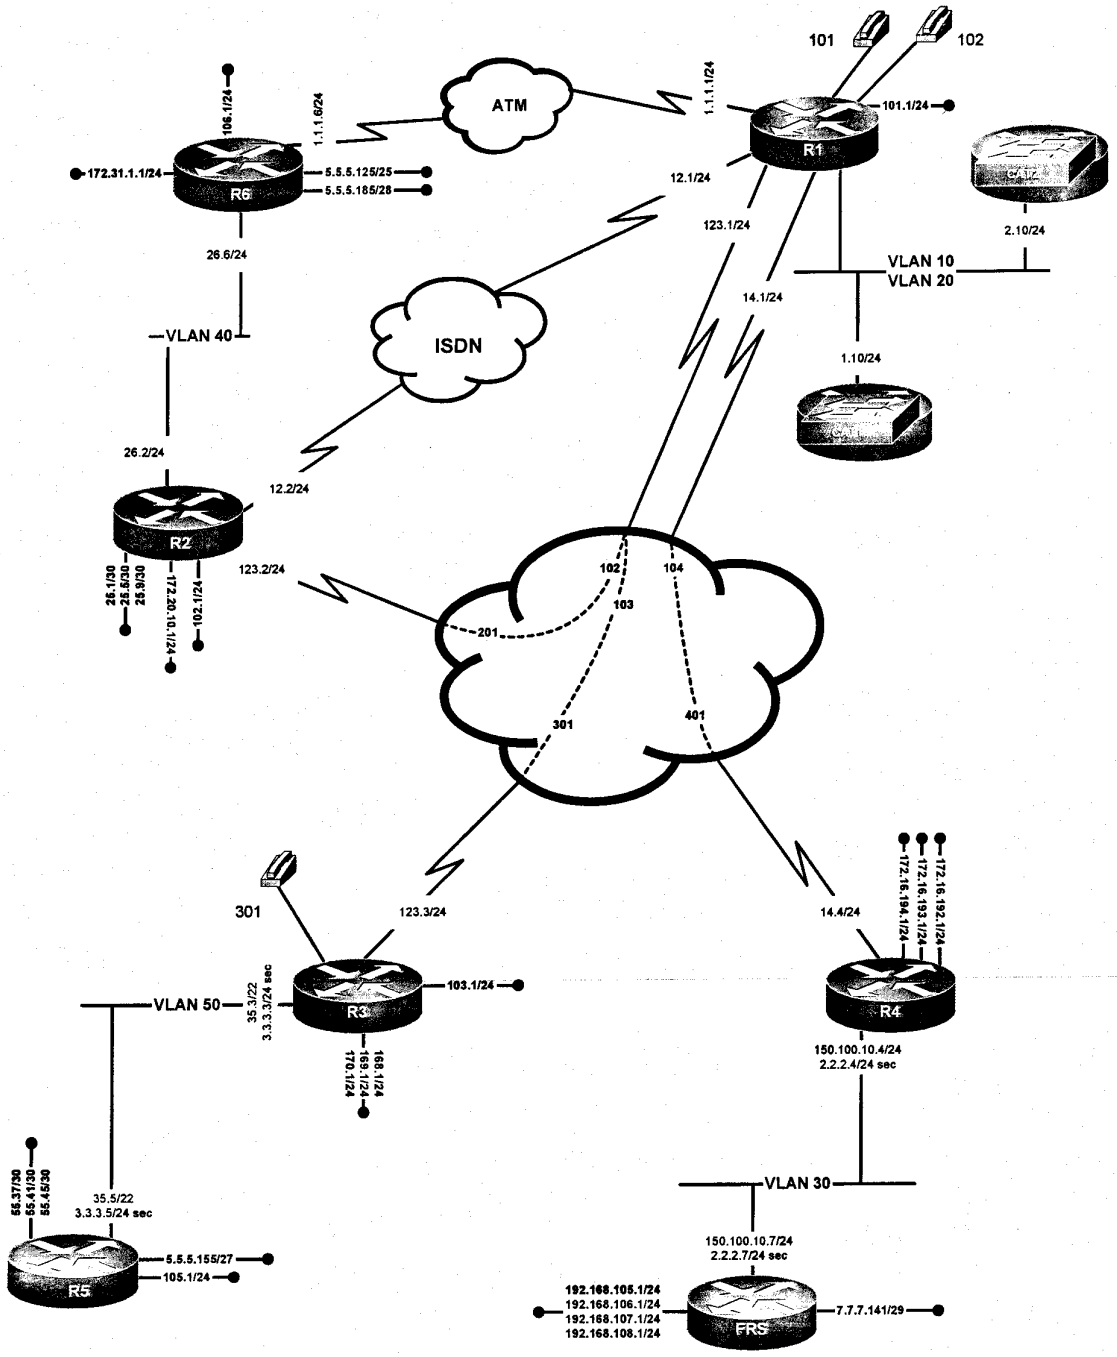

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

13.1 Frame Relay

13.1.1 Configure the Frame Switch according to DLCI assignments listed on the diagram.
13.1.2 Make sure only PVCs listed on the diagram are used for user traffic.
13.1.3 One and only one router must be configured with two and only two logical Frame-Relay interfaces. **The three remaining routers with Frame-Relay interfaces must use physical interfaces only. The logical interfaces must be of a different type.**
13.1.4 Make sure that all routers can ping all attached Frame-Relay interfaces, including their own local interface.

13.2 Catalyst Configuration

13.2.1 Configure the CAT1 management interface with an IP address of 172.16.1.10/24. Assign the management interface to VLAN 20
13.2.2 Create the VLAN's referenced in the VLAN configuration table. When creating these VLAN's do not use any type of dynamic VLAN advertisement protocol.
13.2.3 Enable trunking on the necessary ports to fulfill connectivity requirements.
13.2.4 Supply IP addresses for all specified router interfaces in the diagram
13.2.5 Configure a "router on a stick" configuration on R1 for VLAN's 10 and 20. Make sure that all interfaces on VLAN 20 can ping the R1 router interface. Make sure that all interfaces on VLAN 10 can ping the R1 router interface. Also devices on the VLAN10 and VLAN 20 must be reachable from the rest of the network. CAT1 and CAT2 must be configured for default routing.
13.2.6 Configure a link between CAT1 and CAT2 with the aggregate bandwidth 200 Mbit/second. Use ports Fa0/13 and Fa0/14 to this accomplish this task.

13.3 ATM

13.3.1 Configure PVC 100/100 between R1 and R6
13.3.2 Configure the IP addresses according to the diagram.
13.3.3 Do not include the subnet 1.1.1.0/24 in any IGP

13.4 OSPF

13.4.1 Configure OSPF over Frame-Relay between routers R1, R2 and R3. Do not use an OSPF network type that uses a DR/BDR. Make the Frame-Relay network the OSPF backbone area.
13.4.2 On R2, create the following three loopback interfaces with 30-bit address masks and place them into OSPF area 2:

- 172.16.25.1/30
- 172.16.25.5/30
- 172.16.25.9/30

Summarize the subnets with a 27-bit mask.

13.4.3 Configure OSPF area 3 on the interface fa0/1 of R3
13.4.4 Configure VLAN 40 in OSPF area 26. Create a loopback on R6 with the address of 172.31.1.1/24 and place it into OSPF area 5.
13.4.5 Assign the IP address 172.20.10.1/24 to a Loopback interface on R2 and add it into the OSPF routing process as an external major network.

13.4.6 Authenticate the OSPF backbone area. Do not use a clear text password. Use password "nmc"
13.4.7 Advertise the following Loopback subnets in OSPF area 101:

- 172.16.101.0/24
- 172.16.102.0/24
- 172.16.103.0/24
- 172.16.106.0/24

13.4.8 Configure OSPF area 12 on the network 172.16.12.0/24 between R1 and R2

13.5 RIP

13.5.1 Configure RIP version 1 over the Frame-Relay connection between R1 and R4.
13.5.2 Make sure that RIP advertises only over the specified interfaces.
13.5.3 Advertise the following networks from FRS

- 192.168.105.0/24
- 192.168.106.0/24
- 192.168.107.0/24
- 192.168.108.0/24

13.5.4 Allow only the following networks from FRS running RIP version2 into R4:

- 192.168.105.0/24
- 192.168.106.0/24
- 192.168.107.0/24

Use a minimal number of filtering statements to accomplish this task.

13.5.5 Do not send the unnecessary RIP updates on the subnets 172.16.14.0/24 and 150.100.10.0/24
13.5.6 Create the following three loopback interfaces on R4 with the addresses displayed below:

- 172.16.192.0/24
- 172.16.193.0/24
- 172.16.194.0/24

Summarize these routes when they are redistributed into OSPF.

13.6 EIGRP

13.6.1 Configure EIGRP AS 35 on the VLAN between R3 and R5.
13.6.2 Make sure that EIGRP AS 35 advertises only over the specified interfaces.
13.6.3 Create the following three loopback interfaces on R5 with the addresses displayed below:

- 172.16.55.37/30
- 172.16.55.41/30
- 172.16.55.45/30

Make sure that R3 has only a summary of these routes in its routing table.

13.6.4 Advertise Loopback subnet 172.16.105.0/24 in EIGRP AS35
13.6.5 Advertise the minimal number of prefixes to router R5; however, make sure that R5 is able to reach all IP addresses within your pod. Do not use a 0.0.0.0/0 route.
13.6.6 Restrict the bandwidth utilization to half of the default value for EIGRP traffic on VLAN 50.

13.7 BGP

13.7.1 Configure AS 500, AS 600 and AS 700 on routers R5, R6 and FRS respectively
13.7.2 Routers R1, R2, R3 and R4 are part of AS 100
13.7.3 Configure AS 62222 on routers R1, R2 and R3. Create a full mesh of BGP peer relationships
13.7.4 Configure AS 61111 on router R4
13.7.5 AS 62222 and AS 61111 should be peering between R3 and R4 frame relay interfaces only
13.7.6 Peer AS 600 and AS 100 between routers R6 and R1 over the ATM link, 1.1.1.0/24
13.7.7 Peer AS 500 and AS 100 between R5 and R3 over the secondary network 3.3.3.0/24
13.7.8 Peer AS 700 and AS 100 between FRS and R4 over the secondary network 2.2.2.0/24
13.7.9 Networks 1.1.1.0/24, 2.2.2.0/24 and 3.3.3.0/24 must not be included in any IGP
13.7.10 You may use "no synchronization" in this scenario

13.8 BGP Advertising and Filtering

13.8.1 Advertise the following prefixes from the corresponding AS's:

5.5.5.125/25 5.5.5.185/28	AS 600
5.5.5.155/27	AS 500
7.7.7.141/29	AS 700

13.8.2 R1 must receive only /25 from R6. Do not use filtering techniques based on prefixes and communities.
13.8.3 Make sure you can ping between the following IP addresses through AS 100:

- 5.5.5.125/25
- 5.5.5.185/28
- 5.5.5.155/27
- 7.7.7.141/29

13.9 DLSW

13.9.1 Configure a DLSW peer relationship between R1 and R5. Both sides of the peer relationship should have type "conf"

13.9.2 Use only one command on each side to establish a peer relationship, although you should see 4 connected peers in the output of "show dlsw peer"

13.10 ISDN

13.10.1 Configure ISDN physical interfaces to accomplish this task

13.10.2 The ISDN link should be used to backup the Frame Relay connection between R1 and R2. CHAP authentication will be used for this link.
13.10.3 Either side can call, but if R1 calls R2, both sides must challenge each other. If R2 calls R1, R2 will send the challenge only. R1 will never challenge R2
13.10.4 Use password "exam13"

13.11 Router Maintenance

13.11.1 The router R3 crashes for an unknown reason. The technical support requested crash dump data. The TFTP server is set on the station with the ip address 172.16.35.100
13.11.2 Send the crash information
13.11.3 Sometimes, while switching between reverse telnet sessions, users enter the X28 inline editor. Apply the solution to disable this editor on router R1.

13.12 Security

13.12.1 The Network Administrator is planning to secure the network by implementing the following steps:

- Stop sending and receiving ICMP redirects on Ethernet networks
- Lower the risk of being an amplifier network for any smurf attacks
- Prevent a quick port scan (Port scan is used by the attacker to find out what services are available on your network)

13.12.2 Develop a sample configuration, apply it on the FRS router

13.13 VOICE

13.13.1 Using a 3 digit dial-plan, configure two phones connected to R1 using numbers 101 for port 1/0/0 and 102 for port 1/0/1. Configure the phone connected to port 1/0/0 of router R3 using number 301
13.13.2 Configure VOIP between R1 and R3 to provide voice communications. Use IP addresses 172.16.123.1 and 172.16.123.3 for this task
13.13.3 Make sure you do not hear the dial tone if the Frame Relay interface is DOWN on either side.

13.14 QOS

13.14.1 For ingress traffic coming to CAT1 from R3, mark all packets with CoS=4. All packets with CoS=4 should have an internal DSCP value of 45.
13.14.2 CAT2 is a QoS boundary switch for a different domain. DSCP values should be adjusted according to the following rule: for incoming packets carrying DSCP=45 set the DSCP=35 on the CAT1-CAT2 trunk.
13.14.3 On the switch port connected to FRS, DSCP=35 should be converted to COS=2. Make sure that only COS=2 is mapped to output queue 2 for that port

13.15 Catalyst Specialties

13.15.1 Ports Fa0/16 through Fa0/18 of CAT1 are patched to the visitor room to provide connectivity to different visitors
13.15.2 Provide the NMC monitoring server located at 66.170.103.40 with the information about MAC addresses learned on the interfaces Fa0/16 through Fa0/18

13.15.3 Use the community string "visitor-room". Report to the monitoring server changes in the MAC table every 20 seconds if there are any
13.15.4 Keep the last 15 records of MAC address changes in the switch notification history table.

13.16 Multicast

13.16.1 Configure PIM Dense-Mode multicast routing on routers R1, R3, R4 and FRS.
13.16.2 Form PIM neighbor relationships between R1 and R3, R1 and R4 as well as R4 and FRS
13.16.3 Generate multicast traffic destined to 229.13.13.13 from the Frame-Relay interface of R2
13.16.4 Simulate a member of multicast group 229.13.13.13 on each router R1, R3, R4, FRS. Use one of the loopback interfaces
13.16.5 Simulate a member of the 229.13.13.13 multicast group on R5 using the Ethernet interface of R5
13.16.6 Simulate a member of the 229.0.0.13 multicast group on R5 using the Ethernet interface of R5
13.16.7 Configure static multicast router R3 on CAT2
13.16.8 Allow only (*,229.13.13.13) to be created on R3. Configure CAT1 to accomplish this task
13.16.9 Make sure you receive the responses on R2 from all members of 229.13.13.13

Scenario 14. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16
- Do not introduce any new IP addresses
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- In this exercise, R5 and FRS are used for backbone router simulation
- In this exercise, it's not required to ping R5 and FRS originated networks 192.*.*.* and 140.*.*.*

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10 VLAN20
R2	Ethernet0	VLAN40
R3	FastEthernet0/1	VLAN50 VLAN30
R3	FastEthernet0/0	-
R4	Ethernet0	VLAN40
R5	Ethernet0	VLAN50
R6	FastEthernet0/0	VLAN40
FRS	Ethernet0	VLAN20
CAT1	-	VLAN10
CAT2	-	-

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 14

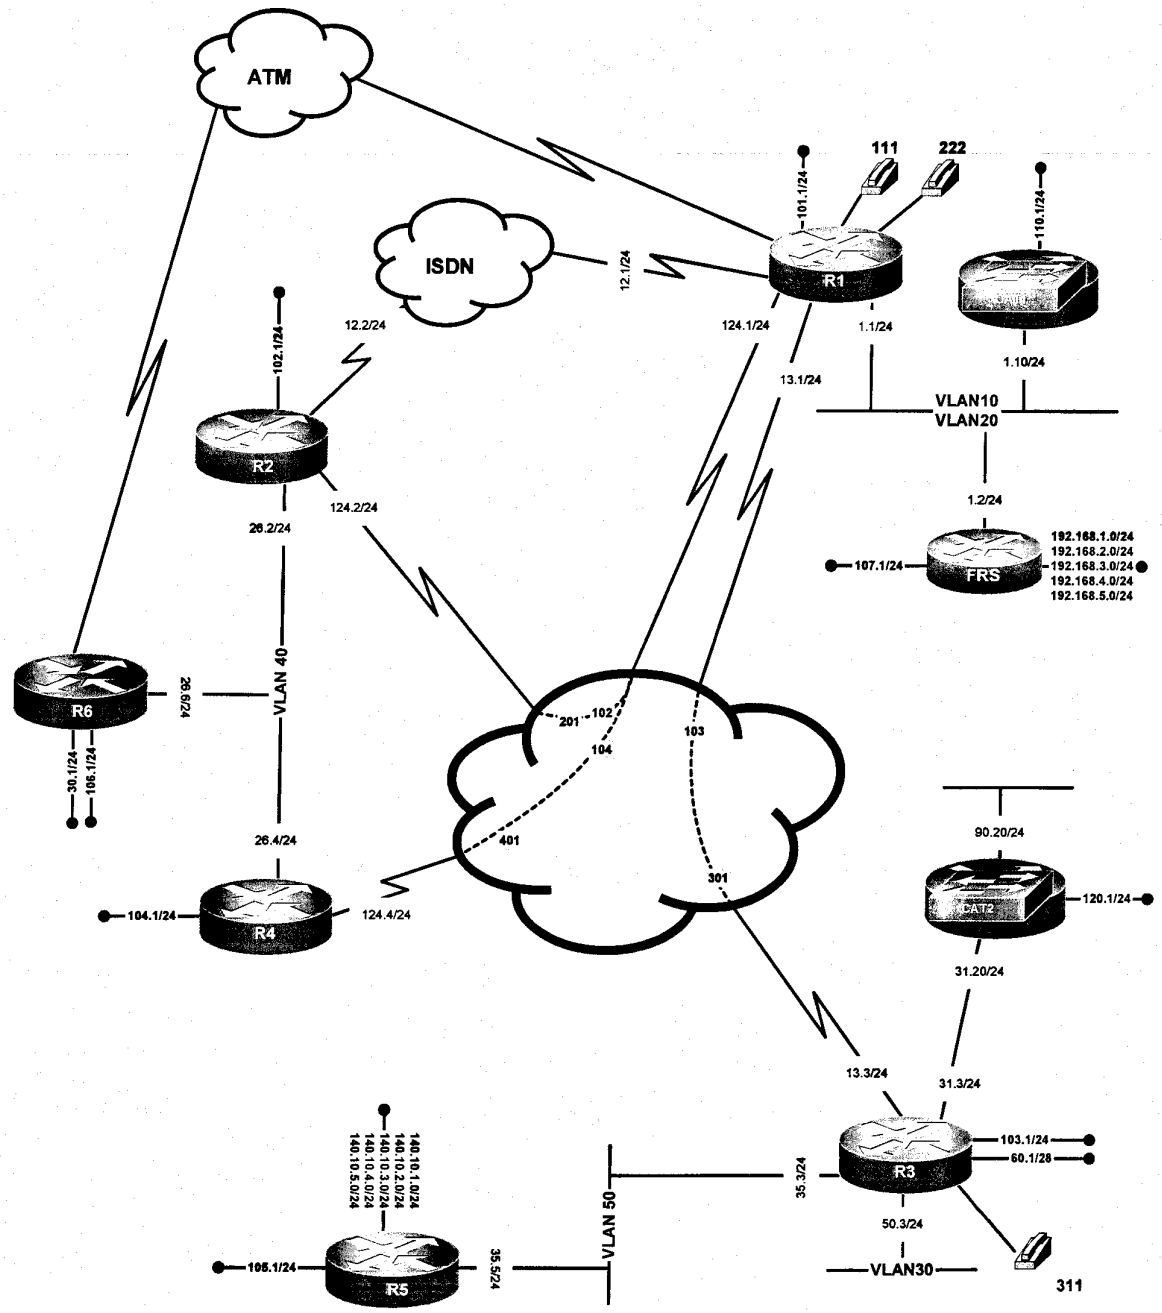

14.1 Frame Relay

14.1.1 Configure Frame-Relay according to DLCI assignments listed on the diagram.
14.1.2 Make sure only PVCs listed on the diagram are used for user traffic.

14.2 Catalyst Configuration

14.2.1 Configure VLAN 10 between R1 and FRS. Configure VLAN 20 between R1 and the CAT1 VLAN 20 switched virtual interface. Assign IP addresses from the 172.16.1.0/24 subnet to R1, FRS and the CAT1 switched virtual interface.

14.3 ATM

14.3.1 Configure ATM PVC 200/200 between R6 and R1 using physical interfaces
14.3.2 PVC 200/200 will be used for voice traffic only. You do not have to configure IP addresses on the ATM link.

14.4 OSPF

14.4.1 Configure the Frame Relay network as the OSPF area 0. Use physical interface on R3, and logical interfaces on all other routers. R1, R2, and R4 should be in the same subnet and R1 and R3 should be in the same subnet. **Do not include the R1/R3 subnet in the backbone.**
14.4.2 Configure OSPF on the link between R1 and R3 for OSPF area 33. On R3, assign a loopback interface with the IP address 172.16.60.1/28 to OSPF area 44. Also, make R3 the Designated router for the R1/R3 link
14.4.3 Use the OSPF network type non-broadcast for all Frame Relay connections
14.4.4 On R3, have the VLAN30 subnet advertised via OSPF without including it as one of your OSPF networks. Make sure that it is viewed as a type 1 route by OSPF and the network can be reached from everywhere. Also, ensure that there is only one entry for the loopback /28 subnet and that the loopback subnet can be seen by all other routers. Do not include the VLAN between R3 and R5 within any IGP.

14.5 RIP

14.5.1 On R6, configure RIP v.1. Advertise the loopbacks on R6 to R2 and R4 via RIP, do not broadcast/multicast RIP updates on VLAN 40 for security reasons.
14.5.2 On R2 and R4, configure mutual redistribution between OSPF and RIP. Use the minimum non-zero metric in redistribution.
14.5.3 Configure RIP between R3 FastEthernet 0/0 and CAT2. Send updates to 224.0.0.9 only on the link between R3 and CAT2

14.6 EIGRP

14.6.1 On R1, configure EIGRP AS 70 and only allow networks 192.168.2.0 and 192.168.3.0 to be accepted into R1 from the router FRS. Also do not have R1 send EIGRP updates.
14.6.2 Make sure that EIGRP routes allowed on R1 are seen by all other routers.
14.6.3 Configure EIGRP 20 between CAT1 and FRS and advertise network 172.16.110.0/24 from CAT1

14.7 ISIS

14.7.1 Configure ISIS between R1 and FRS with the NET 49.0001.1111.1111.1111.00 on R1 and 49.0007.7777.7777.7777.00 on FRS
14.7.2 Configure domain authentication using password "R1FRS"
14.7.3 Advertise network 172.16.107.0/24 in ISIS

14.8 BGP

14.8.1 On R3, configure BGP AS 100 (your AS) to speak with AS500 (backbone AS) configured on R5. AS500 will be advertising routes from 140.10.1.0/24 - 140.10.5.0/24 Allow only networks 140.10.2.0/24 - 140.10.5.0/24 into your AS.
14.8.2 Configure IBGP between R3 and R1. Make sure that R1 has networks from AS500 installed in its routing table.
14.8.3 Modify your BGP configuration so that R4 can see the networks from 140.10.*.* range in its local routing table. Do not use any type of redistribution and do not use a full mesh.
14.8.4 Configure BGP so that you are sending an aggregate for 172.0.0.0/8 to AS500 and supressing all other routes. Also, make sure R4 sees only the following AS 500 originated BGP update, 140.10.2.0/24, in its routing table.

14.9 DLSW

14.9.1 Configure DLSw+ between R1 and R5 using TCP.
14.9.2 Configure DLSw+ between R4 and R5 using FST.
14.9.3 Modify your DLSw+ configuration to allow any to any connectivity between R1 and R4 without forming a full mesh.
14.9.4 Configure an access list on R4 that will allow cluster controllers located on R4's Ethernet to access a FEP (3333.2222.1111) and allow NetBIOS clients to access two servers named NMCVASERVER and NMCMDSERVER. (The FEP and NetBIOS servers are located on R1's Ethernet.)

14.10 ISDN

14.10.1 Configure dial backup so that R2 calls R1 when routes learned over the Frame-Relay interface disappear. Do not use the backup interface command. Make sure the ISDN link does not stay up indefinitely. (Include the backup link in OSPF area 12).
14.10.2 Make sure all traffic originating from network 172.16.30.0/24 uses the ISDN link instead of the Frame Relay network.

14.11 Router Maintenance

14.11.1 Configure R4 to supply configuration information to a new router which will be connected to VLAN 40 in the future. The new router should receive its configuration from the TFTP server 172.16.50.100 located on the VLAN 30.
14.11.2 The new router will have IP address 172.16.26.100/24 and MAC address 0010.7be8.131d
14.11.3 The new router should send the request for the configuration R100.cfg via R2

14.12 Security

14.12.1 The network administrator needs to secure the following part of the network:

- The Administrator does not want any packet to be routed in his network based on the routing path carried in the IP packet
- Do not allow BOOTP services
- Disable autoconfiguration

14.12.2 Develop a sample configuration and apply it on R5

14.13 VOICE over IP

14.13.1 Use the three digit dial plan 311 on the phone connected to port 1/0/0 on router R3.
14.13.2 Use the three digit dial plan 111 on the phone connected to port 1/0/0 on router R1.
14.13.3 Configure two VoIP peers between R1 and R3 using two separate dial-peers. Make the VOIP dial-peer with the higher IP address the more preferred VoIP path from R1 to R3.
14.13.4 Make sure that all dial-peers do not suppress any silence during a call.
14.13.5 Make sure that each call uses a minimum of IP packet overhead.

14.14 VOICE over ATM

14.14.1 Configure voice communications between R6 and R1 over the ATM link, so that R6 can call either phone on R1. For the phone connected to the R1 voice-port 1/0/1, the phone number is 222. For the phone connected to the R1 voice-port 1/0/0, the phone number is 111.
14.14.2 Allow a Peak value of 128 Kbps, average value of 64 Kbps on the ATM PVC used for voice.
14.14.3 Restrict fax transmissions to 9600 bps to number 111 and do not allow fax transfers to number 222

14.15 Catalyst Specialties

14.15.1 CAT2 will be connected to the test environment network 172.16.90.0/24 in the future. Interface FastEthernet0/7 will be used for this purpose
14.15.2 Provide a solution to allow traffic sourced on the network 172.16.90.0/24 only from the selected hosts (.1, .3, .5,.7) to get into your network. Apply the solution on CAT2. Do not use filtering techniques based on Layer 2 filtering. Use minimal number of statements for this task
14.15.3 Only telnet and ICMP must be allowed as user data traffic on VLAN40. Configure 3550 switches to accommodate this restriction. Do not use VLAN map based filtering

14.16 Address Administration

14.16.1 Configure the 1.1.1.3/24 address on R3's FastEthernet0/0 interface without changing any pre-existing IP addresses. Use the 1.1.1.0 private address space and use a portion of the legal address space of the R3 FastEthernet subnet.
14.16.2 Make sure that CAT1 can ping all interfaces in your network. However, allow TELNET access only from the R3's management loopback interface.

14.17 Multicast

14.17.1 Configure Multicast Routing on R1, R2, R3 and R4 using R1 as the shared root.
14.17.2 Announce the shared root without use of any dense groups or static configurations.
14.17.3 Join management loopback interfaces of R1, R2, R3 and R4 to group 239.10.10.10. Join the FastEthernet interface of R6 to group 239.10.10.10
14.17.4 Join interfaces on VLAN 40 to group 239.10.10.10. Make sure traffic to 239.10.10.10 is flooded out of the appropriate VLAN 40 ports only.

Scenario 15. Multi-topic CCIE level exam

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 with their original mask
- Do not use "ip default-network"
- Do not use dialer watch in the ISDN section
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks 192.50.*.* are excluded from the previous requirement
- Networks advertised in the BGP section must be reachable only in the BGP domain

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	Ethernet0	VLAN20
R3	-	-
R3	FastEthernet0/0	VLAN30
R4	Ethernet0	VLAN40
R5	-	-
R6	-	-
FRS	Ethernet0	VLAN30
CAT1	-	VLAN10
CAT2	-	VLAN40

Attention For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15

Scenario 15

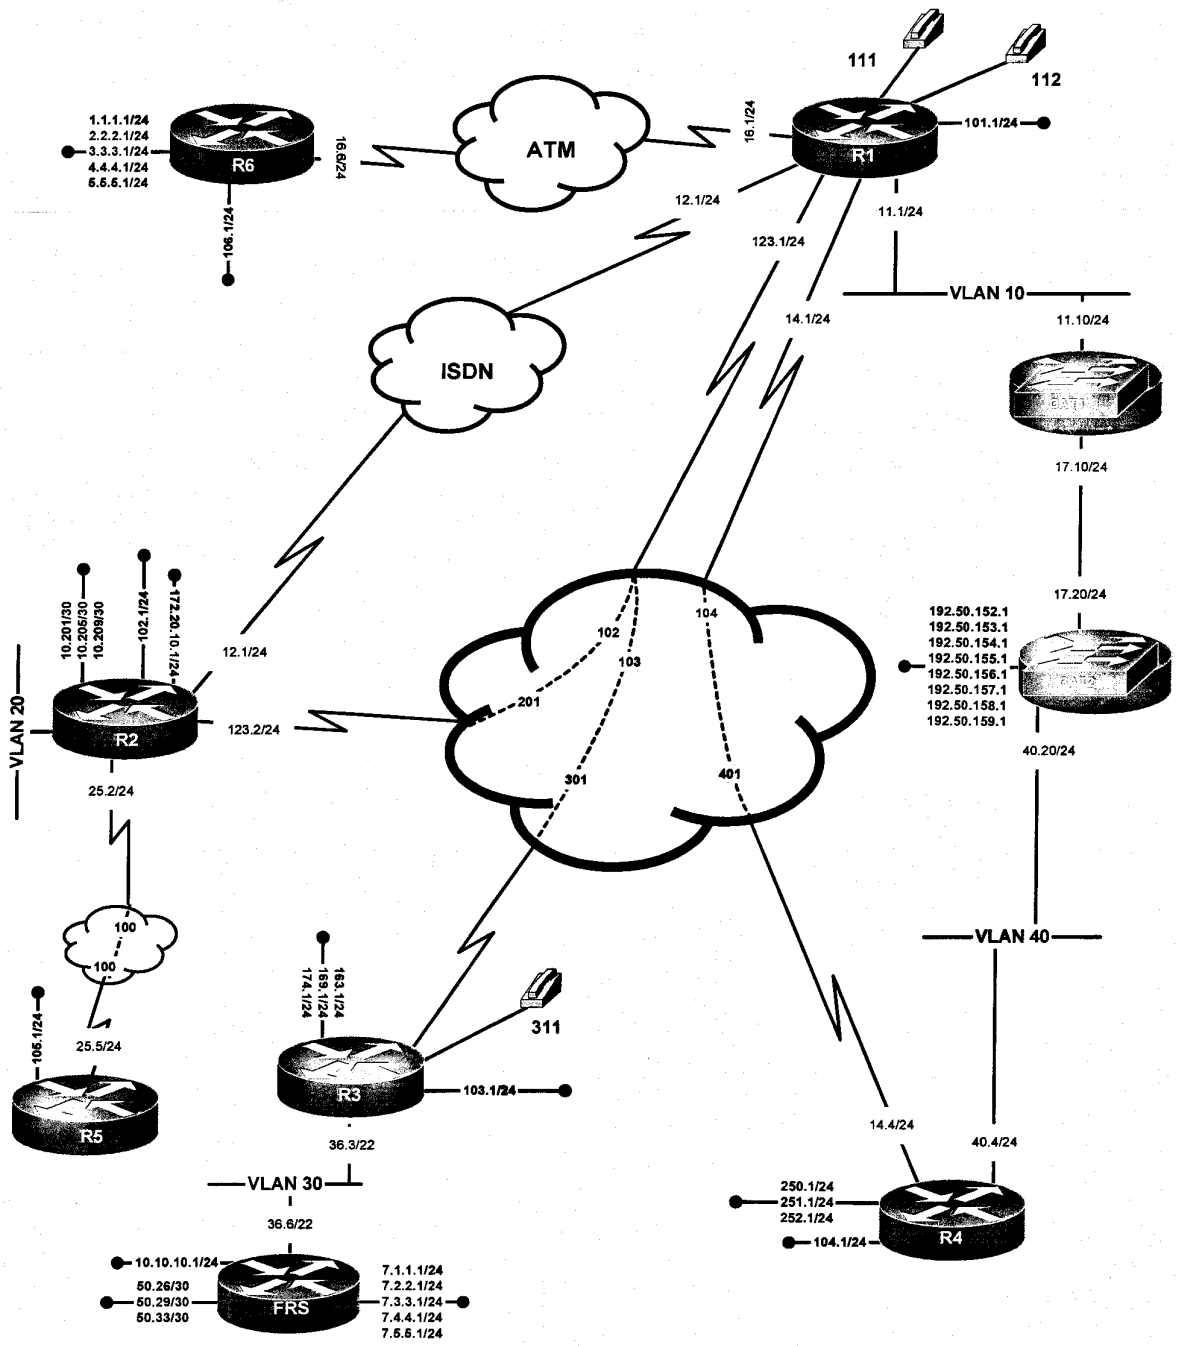

15.1 Frame Relay

15.1.1 Configure the Frame links according to DLCI assignments listed on the diagram. Also, configure a back-to-back Frame-Relay configuration between routers R2 and R5. Refer to diagram.
15.1.2 Enable Frame-Relay interfaces on the R2 and R5 routers. Use physical interfaces on this link
15.1.3 Configure logical interfaces on the subnet 172.16.123.0/24, use point-to-point interfaces on the spokes
15.1.4 Configure physical interfaces on the link 172.16.14.0/24
15.1.5 Use only the PVCs displayed on the diagram

15.2 Catalyst Configuration

15.2.1 Configure VLAN's according to the diagram and the VLAN table
15.2.2 Use a trunk protocol which does not support native vlans
15.2.3 Make sure only VLAN's used for this scenario are permitted on the configured trunks

15.3 ATM

15.3.1 Configure PVC 100/100 between R1 and R6 according to the diagram
15.3.2 Configure logical point-to-point interfaces to accomplish this task
15.3.3 Make sure that your ATM connection does not exceed a transfer rate of 512Kbps.

15.4 OSPF

15.4.1 Configure OSPF over Frame-Relay between routers R1, R2 and R3. Use the OSPF non-broadcast network type. Make the Frame-Relay network the OSPF backbone area.
15.4.2 On R2, create the following three loopback interfaces with 30-bit address masks and place them into OSPF area 2. Summarize the subnets with the most efficient mask

- 172.16.10.201/30
- 172.16.10.205/30
- 172.16.10.209/30

15.4.3 On R3, create three loopback interfaces with the following subnets and place them into OSPF area 4:

- 172.16.163.0/24
- 172.16.169.0/24
- 172.16.174.0/24

Summarize the three entries with the most efficient mask

15.4.4 Advertise the loopback 172.16.103.1/24 on R3 via OSPF without including it in any areas.
15.4.5 Configure OSPF area 10 on the ATM link between R1 and R6. Make sure that the routers in this OSPF area possess the minimum amount of routing information to reach all destinations within your pod. Advertise the loopback network 172.16.106.0/24 in area 10 from R6
15.4.6 Configure OSPF area 25 on the link between R2 and R5. Advertise the loopback 172.16.105.1/24 in OSPF area 105

15.5 RIP

15.5.1 Configure RIP version 2 update exchange over the Frame-Relay connection between R1 and R4.
15.5.2 Make sure that RIP version 2 advertises only over the specified link
15.5.3 R4 is connected to CAT2 via VLAN 40. CAT2 is a RIP version 2 speaking router as well. It advertises the following networks on VLAN40 only:

- 192.50.152.0
- 192.50.153.0
- 192.50.154.0
- 192.50.155.0
- 192.50.156.0
- 192.50.157.0
- 192.50.158.0
- 192.50.159.0

Use one loopback interface to advertise all these networks from CAT2

15.5.4 Make sure R4 learns only the following networks from CAT2:

- 192.50.152.0
- 192.50.153.0
- 192.50.154.0
- 192.50.155.0

Use a single filtering statement to accomplish this task

15.5.5 Create three loopback interfaces on R4 with the addresses displayed below:

- 172.16.250.0 /24
- 172.16.251.0/24
- 172.16.252.0/24

15.5.6 Summarize them using the most efficient mask length.

15.6 EIGRP

15.6.1 Configure EIGRP AS 10 on the subnets 172.16.11.0/24 and 172.16.17.0/24 between R1, CAT1 and CAT2
15.6.2 CAT2 should learn network 10.0.0.0/8 from CAT1. Do not use "summary-address" and do not originate network 10.0.0.0/8 and its subnets on any loopback interfaces of R1, CAT1 and CAT2
15.6.3 CAT1 should accept only odd networks from the following range from CAT2:

- 192.50.156.0
- 192.50.157.0
- 192.50.158.0
- 192.50.159.0

Attention

- Do not perform the redistribution between RIP and EIGRP on CAT2

15.7 ISIS

15.7.1 Configure ISIS level2 routing between R3 and FRS on VLAN 30. Configure NET 49.0003.3333.3333.3333.00 on R3 and NET 40.0007.7777.7777.7777.00 on FRS.
15.7.2 Advertise network 10.10.10.0/24 in ISIS from FRS
15.7.3 Advertise the following networks in ISIS from FRS:

- 172.16.50.26/30
- 172.16.50.29/30
- 172.16.50.33/30

Summarize these prefixes using the most efficient mask length.

15.8 BGP

15.8.1 Make R1, R2, R3, R4, CAT1 and CAT2 BGP speakers within AS 100. Do not allow a full mesh of IBGP speakers within AS 10.
15.8.2 Provide redundant NLRI exchange using R2 and R3 routers.
15.8.3 Configure AS 600 on R6 and AS 700 on FRS. Peer them to AS 100 between R6 and R1 as well as FRS and R3
15.8.4 Advertise the following networks from AS 600:

- 1.1.1.1/24
- 2.2.2.1/24
- 3.3.3.1/24
- 4.4.4.1/24
- 5.5.5.1/24

15.8.5 Advertise the following networks from AS 700:

- 7.1.1.1/24
- 7.2.2.1/24
- 7.3.3.1/24
- 7.4.4.1/24
- 7.5.5.1/24

15.8.6 On each EBGP peer, accept only prefixes with a third octet of only 4 and 5 and originated from a locally connected AS into AS 100.

15.9 DLSW

15.9.1 Configure DLSw+ on routers R1, R2 and R4. Provide DLSw+ connectivity between VLAN 20 on R2 and VLAN 40 on R4. Allow only remote peer statements to be configured from R2 to R1 and R4 to R1. **Do not configure any remote peer commands from r2 to r4 or vice versa.** Configure R1 with a single DLSw+ statement.

15.9.2 NetBIOS hosts NMC1 and NMC2 are connected to VLAN 40 of router R4.

15.9.3 Permit VLAN 20 workstations to communicate with the host NMC1 only .

15.9.4 Cache the MAC address of 0200.61F8.E8E8 originating on VLAN 20 in the DLSw+ reachability cache tables of R1 and R4.

15.10 ISDN

15.10.1 Activate the ISDN interfaces on routers R1 and R2

15.10.2 Use a form of authentication that does not pass a password over the network.

15.10.3 Configure OSPF on the ISDN link. Assign the ISDN interfaces to OSPF area 100. Using OSPF configuration commands only, configure the ISDN link as a backup to the Frame Relay connection between R2 and R1.

15.10.4 Make sure that no routing traffic such as periodic updates or HELLO's keep the interface up indefinitely.

15.11 Address Administration

15.11.1 Configure the R3 VLAN30 interface and the FRS Ethernet interface with an IP address from the 11.1.1.0/24 subnetwork. Configure these new addresses without removing the currently configured IP addresses. Using a source IP address from the 11.1.1.0/24 subnet, ping R1 from both R3 and FRS. When performing these pings, do not let the packets originating from R3 and R6 reach R1 with the IP address 11.1.1.0/24

15.11.2 Ping the FRS Ethernet interface belonging to 11.1.1.0/24 subnet using a pre-selected destination address extracted from the 22-bit subnet assigned to VLAN 30.

15.12 Security

15.12.1 Configure an IP Access-List that will allow only a user named Bill to activate an access-list that allows TELNET services from router R2 into router R4 using the R4 ip address 172.16.14.4 for a period of five minutes. Apply the access-list on the multipoint subinterface of R1. Make sure that ***only*** traffic generated within this Scenario is allowed by this access-list. Also, make sure that pings and traceroute can still be successfully performed from any part of the network. Use a minimum number of access-list statements.

15.13 VOICE

15.13.1 Two telephones are attached to the two FXS ports on router R1. See the diagram. Use the following three digit dial plan: 111 for voice-port 1/0/0 and 112 for voice-port 0/0/1. Ring the phone attached to R1 voice-port 1/0/0 when the phone attached to R1 voice-port 0/0/1 is placed off-hook.

15.13.2 One telephone is attached to the FXS port 1/0/0 on R3. See the diagram. Use the following three digit dial plan: 311. When the phone on R3 is placed off-hook, make the phone attached to port 1/0/1 on router R1 ring. Configure a single dial peer on R3 that would forward digits to either of the two voice ports on R1.
15.13.3 Adjust your existing VoIP configuration so that a minimum of per packet overhead is consumed by voice traffic.

15.14 QOS

15.14.1 A Traffic Generator on VLAN10 is connected to port Fastethernet 0/24 of CAT1. It is generating 5 UDP packets per second. Packet size is 1024 Bytes. The UDP stream is destined to 172.16.11.1 port 5111. See the diagram
15.14.2 Limit incoming UDP traffic destined to port 5111 to a rate of 8000 bit/sec on the interface Fa0/0 of router R1. Configure the minimal syntax values for burst size and extended burst size.
15.14.3 Drop excessive traffic.

15.15 Catalyst Specialties

15.15.1 Set port fa0/18 of CAT2 to bypass the learning and listening states of Spanning Tree.
15.15.2 On CAT2, monitor both transmitted and received traffic on port fa0/19 from a SNIFFER attached on port fa0/20.
15.15.3 Configure the Catalyst so that learned MAC addresses in VLAN 40 are retained by the Catalyst for a period that is one and one half times as long as the default.
15.15.4 Send all Catalyst error messages to the syslog server at 172.16.1.150.
The table displayed below lists and briefly describes the severity levels supported by system message logs:

Severity Level	Keyword	Description
0	emergencies	System unusable
1	alerts	Immediate action required
2	critical	Critical condition
3	errors	Error conditions
4	warning	Warning conditions
5	notifications	Normal but significant event
6	informational	Informational messages
7	debugging	Debugging messages

15.16 Multicast

15.16.1 Enable multicast routing between routers R1, R2, R3, R5 and R6. Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree.
15.16.2 Make R1 the root of the shared tree. Do not use any dynamic methods to discover or advertise the root of the shared tree.
15.16.3 Configure R1, R2, R3, R5 and R6 to join the multicast group 227.7.7.7 only. Associate this multicast group with a loopback interface on each router.
15.16.4 Ping the multicast group 227.7.7.7 from R4.

Scenario 16. Multi-topic CCIE level exam

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 with their original mask
- Do not use the default route 0.0.0.0/0 or default-information originate
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN20 VLAN30
R2	Ethernet0	VLAN10
R3	-	-
R3	-	-
R4	-	-
R5	-	-
R6	FastEthernet0/0	VLAN100
FRS	Ethernet0	VLAN30
CAT1	-	VLAN20 VLAN100
CAT2	-	VLAN10 VLAN100

Attention For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15

Scenario 16

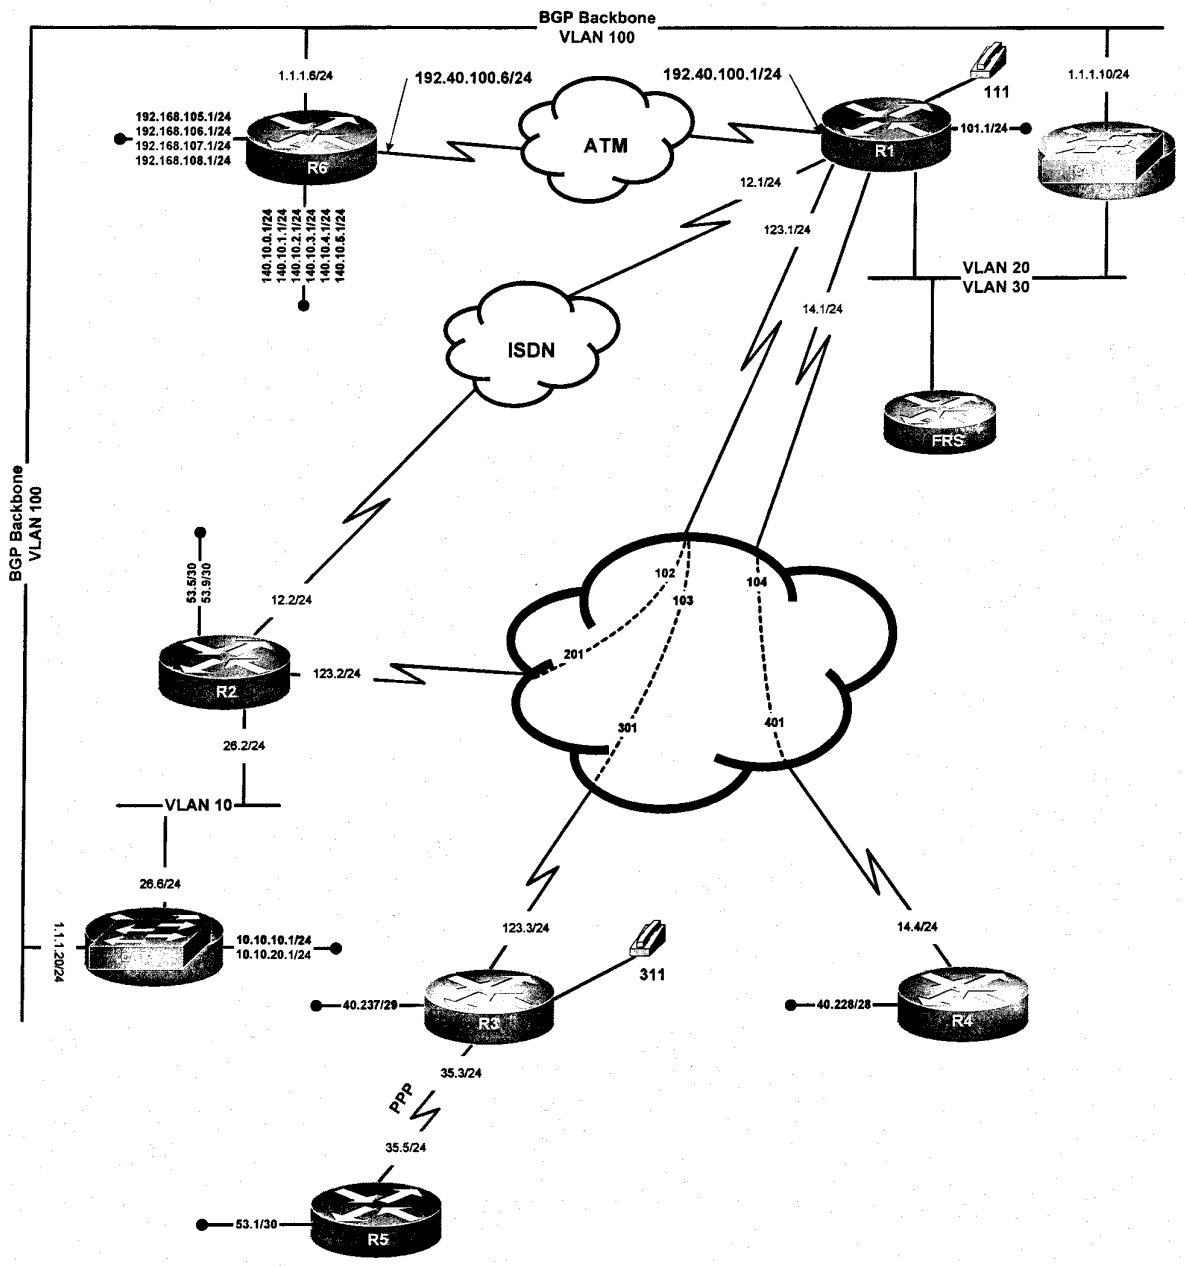

16.1 Frame Relay

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

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

16.2 Catalyst Configuration

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

16.2.2 Configure the CAT1 VLAN20 interface with an IP address of 172.16.10.11/24.

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

16.2.4 Enable trunking only on the necessary ports. Trunk only VLAN's used in this scenario

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

16.3 ATM

16.3.1 Configure an ATM connection between R1 and R6 routers. Use a configuration that will involve the NSAP addresses of R1 and R6 . The ESI of R6's NSAP address is:

100000100000.00

16.3.2 The R1 ESI address is 1111.1111.1111.00. All ATM devices are using same prefix.

16.3.3 Assign the IP address of 192.40.100.1/24 to the R1 physical ATM interface. The IP address assigned to the logical ATM interface on R6 is 192.40.100.6

16.3.4 Do not rely on Inverse ARP in binding Layer3 to Layer2

16.4 OSPF

16.4.1 Configure the Frame-Relay network 172.16.123.0/24 as the OSPF backbone area. Do not elect a DR/BDR in the OSPF area.

16.4.2 Configure OSPF on the link between R1 and R4. Place this link in OSPF area 41.

16.4.3 Make R4 the designated router for the R1/R4 link.

16.4.4 On R4, configure a loopback network 172.16.40.228/28 and assign it to OSPF area 4.

16.4.5 Configure OSPF MD5 authentication for area 0.

16.4.6 On R3, advertise a loopback 172.16.40.237/29 without including it in any routing protocol. Make sure it is viewed as a Type 1 OSPF External route.

16.4.7 Advertise the following loopbacks from R2 as area 22:

- 172.16.53.5/30
- 172.16.53.9/30

16.4.8 Advertise FRS and the Catalyst VLAN 30 interface as one 172.16.11.0/24 OSPF network. Assign this network to OSPF area 1.

16.4.9 Advertise Loopback 172.16.101.0/24 in OSPF area 101

16.5 RIP

16.5.1 Configure RIP version 1 over the ATM connection between R1 and R6.
16.5.2 Do not broadcast RIP updates on the ATM connection.
16.5.3 Advertise the following networks from R6:

- 192.168.105.0/24
- 192.168.106.0/24
- 192.168.107.0/24
- 192.168.108.0/24

16.5.4 Receive only the 192.168.106.0/24 and 192.168.107.0/24 subnets on R1 from the R6.

16.6 EIGRP

16.6.1 Configure EIGRP AS 1 between R2 and CAT2
16.6.2 Create one and only one loopback on router CAT2 and assign the following two addresses to the single loopback interface: 10.10.10.1/24 and 10.10.20.1/24.
16.6.3 Advertise these addresses via EIGRP AS1.
16.6.4 Make sure that EIGRP advertises only over VLAN10 interfaces.
16.6.5 Restrict the bandwidth utilization to half of the default value for EIGRP traffic on VLAN 10.

16.7 ISIS

16.7.1 Configure Integrated IS-IS on the point-to-point link between R3 and R5. Configure R3 in area 49.0009. Configure R5 in area 49.0099.
16.7.2 Create a loopback interface on R5. Assign an IP address 172.16.53.1/30 to the loopback. Advertise this IP address into Integrated IS-IS.
16.7.3 Make sure that the ISIS configured routers generate the minimum number of IS-IS packet types.

16.8 BGP

Attention

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

16.8.1 Configure AS 1581 on CAT1, AS 1771 on R6 and AS 1776 on CAT2. Configure VLAN100 on R6, CAT1 and CAT2 and assign an IP address from the 1.1.1.0/24 subnet to each of these routers that possess an connection to VLAN 100. See the Scenario diagram
16.8.2 Peer AS 1581 and AS 1771 over 1.1.1.0/24. Peer AS 1771 and AS 1776 over 1.1.1.0/24
16.8.3 Advertise the following networks from AS 1771:

- 140.10.0.0/24
- 140.10.1.0/24

- 140.10.2.0/24
- 140.10.3.0/24
- 140.10.4.0/24
- 140.10.5.0/24

16.8.4 Configure the BGP router-id 172.16.110.1 and 172.16.120.1 on CAT1 and CAT2 respectively
16.8.5 Make R1, R2, R3, R4 ,R5 and FRS BGP speakers within AS 100.
16.8.6 Do not allow a full mesh of IBGP speakers within AS100. R3 and R1 should be involved in NLRI exchange within AS 100. Also peer R1 and CAT1 as well as R3 and CAT2
16.8.7 Configure R1,R4 and FRS in AS 65001.
16.8.8 Configure R2, R3 and R5 in AS 65000.
16.8.9 Only one peering relationship can exist between private AS's 65000 and 65001. It must be between R3 and R4.
16.8.10 Allow only the prefix range of 140.10.2.0/24 – 140.10.5.0/24 into AS100.
16.8.11 Forward traffic from AS 100 destined to the 140.10.2.0/24 – 140.10.5.0/24 subnets via CAT1.

16.9 DLSW

16.9.1 Configure DLSw+ on routers R1, R4 and R5.
16.9.2 Maintain the primary connection between R1 and R5. In case this peer fails, allow R1 to connect to R4. If R5 regains connectivity, have the failed over connection from R1 to R4 to be redirected back to the primary DLSw+ peer, R5, after five minutes.
16.9.3 On R4, do not allow any NETBIOS traffic. Advertise this restriction to other peers.

16.10 ISDN

16.10.1 Activate the ISDN interfaces on routers R1 and R2.
16.10.2 Use a form of authentication that does not use an MD5 hash. Make sure that the passwords used in the authentication are different.
16.10.3 Configure OSPF area 12 on the ISDN link. If the 172.16.101.0/24 sub-network disappears from R2's routing table, have R2 activate the ISDN link.
16.10.4 Make sure that no routing traffic such as periodic updates or HELLO's keep the interface up indefinitely.

16.11 Address Administration

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

16.12 NTP

16.12.1 Make R1 the primary time source. Set R1 as a stratum 5.
16.12.2 Configure router R2 so that it obtains its time from R1.
16.12.3 Configure CAT2 to peer with R2. Make sure R2 and CAT2 are synchronized only if they pass authentication. Make sure that CAT2 is synchronized with R2 only.
16.12.4 Configure CAT1 as an NTP client. Enable Daylight Savings Time.

16.13 VOICE

16.13.1 Connect a phone to the FXS port 1/0/0 on router R3. Use the following three digit dial plan: 311.
16.13.2 Connect the phone to the FXS port 1/0/0 on router R1. Use the following three digit dial plan: 111.
16.13.3 Configure two VOIP peers between R1 and R3 using two separate dial-peers. Make the VOIP dial-peer with the higher IP address the more preferred VOIP path from R1 to R3.
16.13.4 For all other three digit destination patterns, forward the digits to 172.30.4.1
16.13.5 Make sure that all dial-peers do not suppress any silence during a call.

16.14 QOS

16.14.1 A traffic generator on VLAN100 is connected to port Fastethernet 0/24 of CAT1. It's is generating 5 UDP packets per second. Packet size is 1024 Bytes. The UDP stream is destined to the VLAN30 interface of CAT1 port 5011.
16.14.2 Police this traffic to 8000 bit/sec on the interface FastEthernet 0/24 of CAT1. Allow a minimal burst size.
16.14.3 Excessive traffic should be dropped.

16.15 Catalyst specialties

16.15.1 Allow only a workstation with a Data-Link address of 0050.04fd.9f73 to use port fa0/20 of CAT2.
16.15.2 Allow TELNET access to CAT1 from the loopback interface of R3 only.
16.15.3 Configure a greeting message "Welcome to RS-NMC-2!" on CAT1.

16.16 Multicast

16.16.1 Enable multicast routing between routers R1, R2, R3 and R5. Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a flood and prune protocol.
16.16.2 Configure all of the above listed routers to join the multicast group 229.9.9.9. Associate this multicast group with a loopback interface on each router.
16.16.3 Join ethernet interfaces of VLAN 10 to the group 229.9.9.9
16.16.4 Statically configure a MAC table entry for ports in VLAN 10 for 229.9.9.9.
16.16.5 Ping the multicast group 229.9.9.9 from R4 to all other multicast routers.

Scenario 17. Multi-topic CCIE level exam

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 with their original mask
- Do not use a static default route 0.0.0.0/0
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks advertised in the BGP section must be reachable only in the BGP domain

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	-	-
R3	-	-
R3	FastEthernet0/0	VLAN20
R4	Ethernet0	VLAN30
R5	Ethernet0	VLAN40
R6	-	VLAN20 VLAN30
FRS	Ethernet0	VLAN20
CAT1	-	VLAN30 VLAN100 VLAN40
CAT2	-	VLAN10 VLAN100

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 17

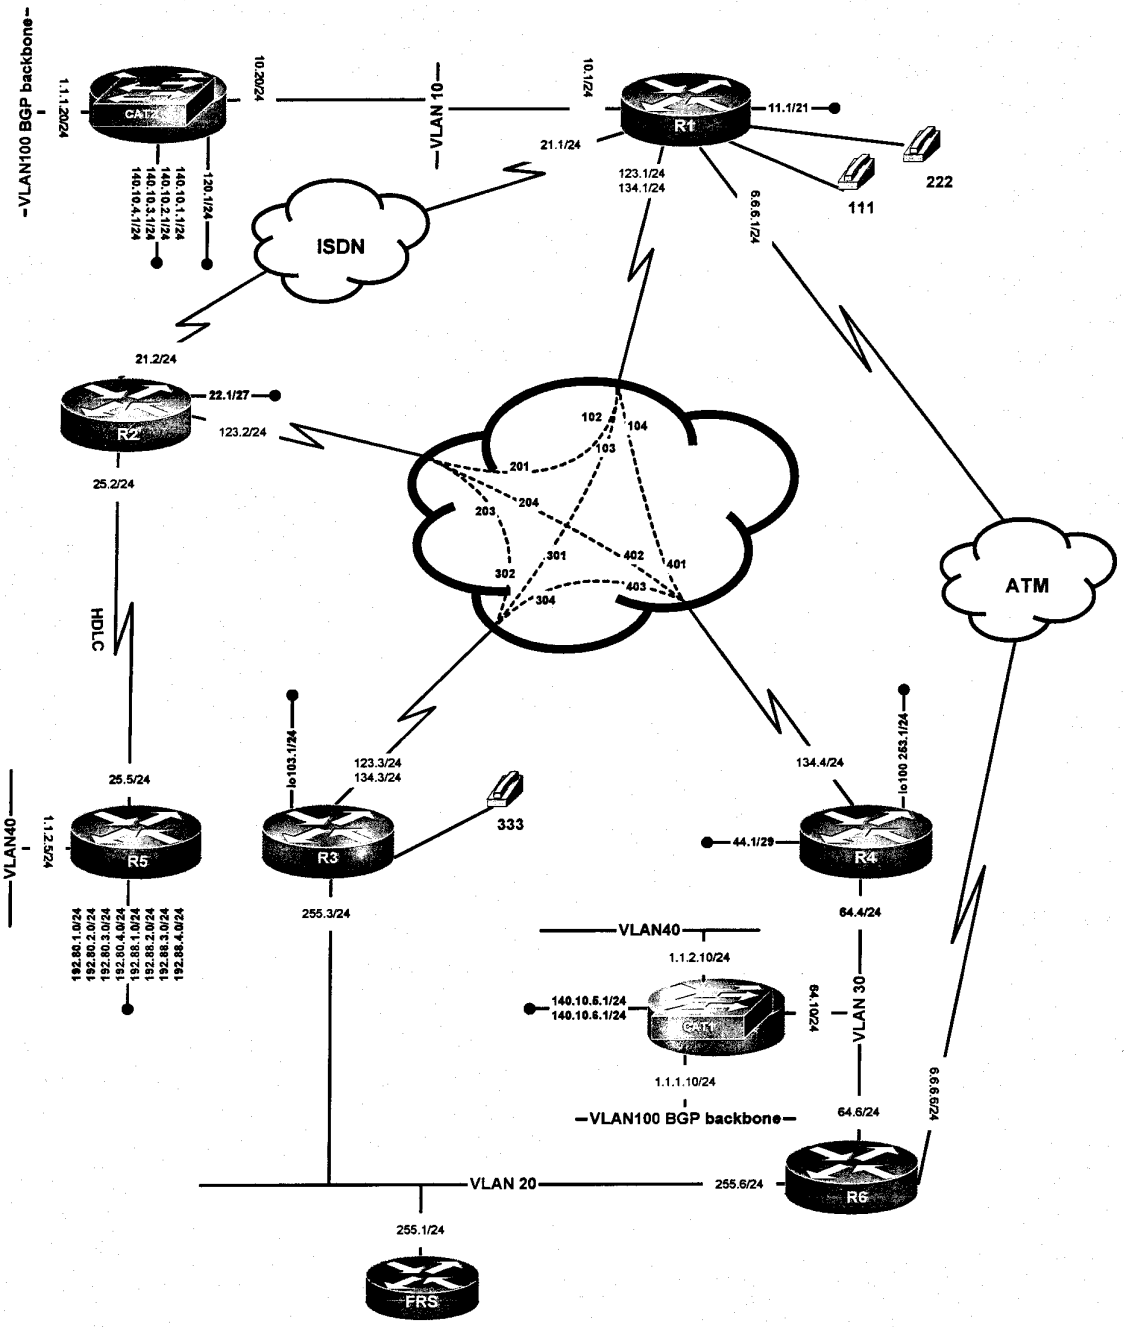

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

17.1 Serial interfaces

17.1.1 Configure Frame Relay encapsulation between the routers R1, R2, R3 and R4. See the diagram
17.1.2 R1, R2 and R3 should be in the same subnet 172.16.123.0/24. R1, R4 and R3 should be in the same subnet 172.16.134.0/24
17.1.3 Use only point-to-point logical interfaces wherever possible and use physical interfaces otherwise.
17.1.4 **Use only the minimum number of DLCI's to fulfill this configuration.**
17.1.5 Configure the HDLC protocol on the serial link between R2 and R5

17.2 Catalyst Configuration

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

17.3 ATM

17.3.1 Configure ATM network 6.6.6.0/24 between R1 and R6. Interfaces should be configured as multipoint interfaces. Use PVC 100/100 for this network. Network 6.6.6.0 will be used for telephone traffic only. It should not be included in any routing protocols, therefore does not have to be reachable from the rest of the network.

17.4 OSPF

17.4.1 Configure OSPF area 0 between R1, R2 and R3. Use an OSPF network type that elects a DR but does not require a neighbor statement. Make sure R3 is the DR.
17.4.2 Configure OSPF area 90 on the subnet between R1, R3 and R4. Use an OSPF network type that elects a DR and requires neighbor statements
17.4.3 For the subnet shared by R1, R3 and R4, make sure that the loss of a neighbor relationship is detected twice as fast as the default.
17.4.4 On the following routers, advertise the following loopback addresses:

On R1: loopback interface :	172.16.11.1/21	Assign to Area 11
On R2: loopback interface :	172.16.22.1/27	Assign to Area 22
On R4: loopback interface :	172.16.44.1/29	Assign to Area 44

Whenever possible, make sure that other routers learn these loopback addresses with their original mask length. Do not use OSPF summarization commands.

17.5 RIP version 1

17.5.1 Configure RIP version 1 between R2 and R5.
17.5.2 Advertise the following networks on R5 into RIP:

- 192.80.1.0/24
- 192.80.2.0/24
- 192.80.3.0/24
- 192.80.4.0/24
- 192.88.1.0/24
- 192.88.2.0/24

- 192.88.3.0/24
- 192.88.4.0/24

17.5.3 Allow only the following networks to be advertised to router R2 with a single access list statement:

- 192.80.2.0
- 192.80.3.0
- 192.88.2.0
- 192.88.3.0

17.6 EIGRP

17.6.1 Configure EIGRP AS 100 between R4, R3 and R6.
17.6.2 Do not form a neighbor relationship over the PVC between R3 and R4 .
17.6.3 On R4, create Loopback 100 and assign the 172.16.253.1/24 IP address to it. Include the interface within EIGRP AS 100.

17.7 RIP version 2

17.7.1 Configure RIP version 2 under the RIP routing process between R1 and CAT2.
17.7.2 Restrict RIP updates to VLAN10 only
17.7.3 Make R1 and CAT2 exchange RIP version 2 updates only.

17.8 BGP

Attention

- VLAN 100 is used for BGP backend connectivity only. It is not part of any IGP routing protocol and does not have to be reachable
- VLAN 40 is used for the BGP connectivity only. It is not part of any IGP routing protocol and does not have to be reachable

17.8.1 Configure BGP AS 1581 on CAT1 and AS 1771 on CAT2
17.8.2 Peer AS 1581 and AS1771 over 1.1.1.0/24. See the diagram.
17.8.3 Advertise the following networks from AS 1771:

- 140.10.1.0/24
- 140.10.2.0/24
- 140.10.3.0/2
- 140.10.4.0/24

17.8.4 Advertise the following networks from AS 1581:

- 140.10.5.0/24

- 140.10.6.0/24

17.8.5 Configure BGP AS 100 on R1, R2 and R5. Configure BGP AS 10 on R3, R4 and R6.
17.8.6 Do not use a full mesh in AS 10.
17.8.7 Peer AS 100 and AS 10 between the following two pairs of EBGP peers: R1 and R4; R2 and R3.
17.8.8 Peer AS 100 and AS 1581 between CAT1 and R5 over VLAN 40
17.8.9 Peer AS100 and AS 1771 between CAT2 and R1 over VLAN 10
17.8.10 Prefer R5 as an exit point to networks 140.10.*.0/24 advertised earlier in this section
17.8.11 In AS 10, set the local-pref for all prefixes originating from AS 1771 and traversing AS 1581 to a value of 300. Also, set all prefixes originating from AS 1581 to a value of 200 in AS 10. Do not use the AS-PATH or the IP address prefix as match criteria to set the local-pref in AS 10. Allow only prefixes originating from AS 1771 into AS 100.

17.9 DLSW

17.9.1 Configure DLSw+ between R5 and loopback 100 on R4 . Make sure DLSW packets are sent through ISDN.
17.9.2 Configure DLSw+ between FRS and loopback 100 on R4.
17.9.3 On R4, do not allow any SNA traffic except type 4 . Advertise this restriction to other peers.

17.10 ISDN

17.10.1 Activate the ISDN interfaces on routers R1 and R2.
17.10.2 Configure RIP on the ISDN link. Make sure the ISDN link does not stay up indefinitely.
17.10.3 Use ppp chap authentication, password "netmaster".
17.10.4 Use the ISDN link only in absence of OSPF routes.

17.11 Address Administration

17.11.1 Configure FRS to act as a host with the ip address of 172.16.255.1/24.
17.11.2 Prefer R6 as a gateway for FRS. Do not use HSRP. Do not use any static configuration.
17.11.3 All packets originating from FRS, should have the source IP address changed at the first hop router. The source IP address must be in 172.16.0.0/16 range.

17.12 Voice over IP over ATM

17.12.1 Configure voice/data integration over the ATM link 6.6.6.0/24 between R6 and R1 using IP connectivity. Use the following three digit dial plan: 222 for R1 voice port 1/0/1
17.12.2 Configure R6 to be able call the phone 222 of R1 using IP address 6.6.6.1
17.12.3 Test the call with the "csim start 222" command executed from R6

17.13 VOICE over Frame Relay

17.13.1 Configure voice/data integration over the Frame-Relay cloud without using IP connectivity. Make the phone on R1 ring the phone on R3. Make sure that both voice and data can be transported over the DLCI. Use the following three digit dial plan: 111 for R1 voice port 1/0/0 and 333 for R3 voice port 1/0/0

17.14 QOS

17.14.1 Restrict TTCP traffic from CAT2 destined to 172.16.11.1 port 5001 to 1000000 Bit/sec on R1. Allow burst traffic up to 512K.
17.14.2 If traffic burst exceeds the additional 512K , drop it.
17.14.3 See the diagram for ip addressing

17.15 Catalyst Specialties

17.15.1 Create VLAN 200 and assign port fa0/20 to it on CAT2. Do not allow BPDU traffic on this VLAN.
17.15.2 Configure R4 as the preferred gateway on CAT1 for networks originating from R1, R2 and R5.
17.15.3 Configure R6 as a backup gateway on CAT1. Do not use HSRP or IRDP protocols.
17.15.4 Configure CAT1 so that it can be managed by a network management service that uses UDP port 161.
17.15.5 Set read-only access with the string of RS-NMC. Set read-write access with the string NMC.

17.16 Gateway Redundancy

17.16.1 Configure HSRP between R4 and R6. Make R4 the preferred gateway. Switch over to R6 if the frame relay connection on R4 goes down.
17.16.2 The virtual gateway IP address is 172.10.64.1

17.17 Multicast

17.17.1 Enable multicast routing between routers R1, R2, and R3. Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree.
17.17.2 Configure all of the above listed routers to join the multicast group 229.10.10.10. Associate this multicast group with a loopback interface on each router.
17.17.3 Use the 224.0.1.39 PIM dense group for this configuration. Make R3 the root of the shared tree. Accomplish this task by configuring only R3.
17.17.4 Ping the multicast group 229.10.10.10 from R4 to all other multicast routers.

Scenario 18. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 151.10.0.0/16
- Do not use any static routes except those used for backup purposes
- Advertise Loopback interfaces with their original mask
- Do not use default-information originate and ip default-network commands
- Do not change any prefix masks
- 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

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN50, Determine others
R2	Ethernet0	VLAN10
R3	FastEthernet0/1	VLAN40
R3	FastEthernet0/0	Determine
R4	Ethernet0	VLAN40
R5	Ethernet0	Determine
R6	FastEthernet0/0	VLAN10
FRS	Ethernet0	VLAN40
CAT1	-	Determine
CAT2	-	VLAN50

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 18

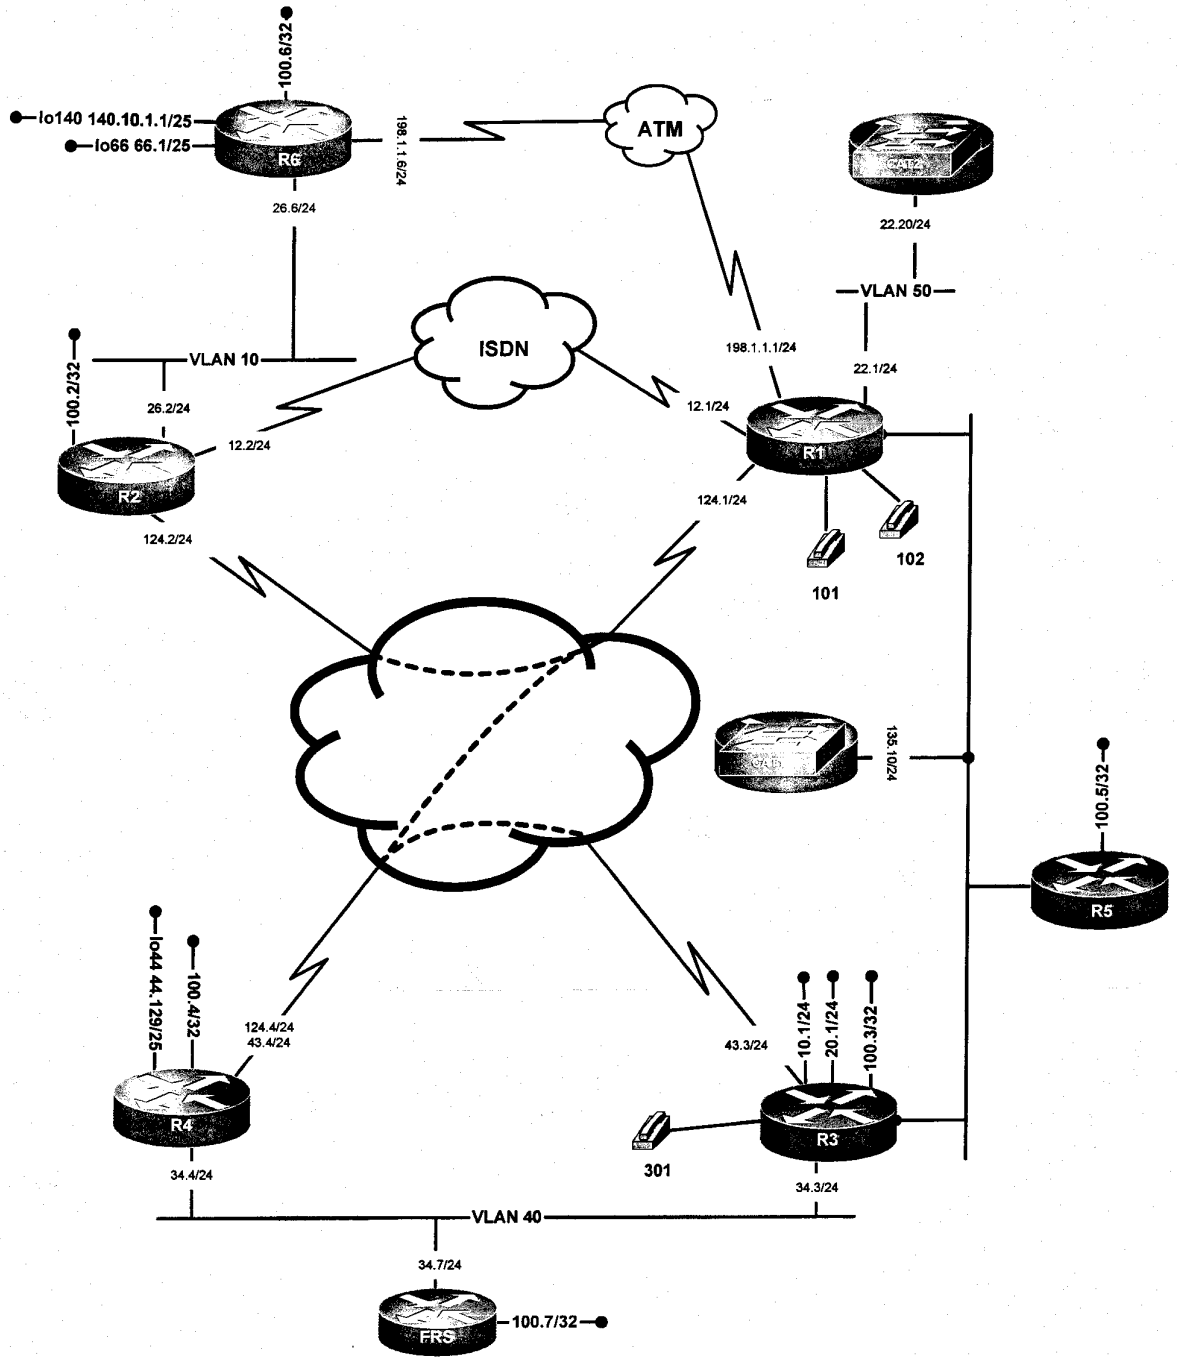

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

18.1 Frame Relay

18.1.1 Configure Frame Relay encapsulation between R1, R2, R3 and R4. See the diagram.
18.1.2 All entries in the "show frame-relay map" display must be dynamic.
18.1.3 **Use only PVCs shown in the diagram to fulfill this configuration.**

18.2 Catalyst Configuration

18.2.1 Do not advertise VLAN's in this scenario. Use the ISL protocol for trunking.
18.2.2 Configure the CAT1 VLAN 20 interface with an IP address of 151.10.135.10/24.
18.2.3 Configure VLAN 20 between CAT1, R1 and R3. Create VLAN 30 between R5, R1 and R3. Place CAT1, R1,R3 and R5 on the same network 151.10.135.0/24. Use the router number for the host part of the IP address.

18.3 ATM

18.3.1 Configure an ATM connection between R1 and R6 routers.
18.3.2 The R1 ESI address is 001100110011.00. The R6 ESI address is 006600660066.00. All ATM devices are using the same prefix.
18.3.3 See the diagram for the IP address assignments.

18.4 OSPF

18.4.1 Place all Frame-Relay subnets in area 0
18.4.2 Choose the OSPF network type which will provide full reachability on subnet 151.10.124.0/24
18.4.3 Choose the default OSPF network type on the Frame-Relay subnet between R3 and R4
18.4.4 Configure OSPF area 10 on the subnet between R4, R3 and FRS. Make FRS the most preferred candidate for DR and R4 the most preferred candidate for BDR.
18.4.5 Configure loopback 151.10.20.1 on router R3 and assign it to area 20
18.4.6 Configure loopback 151.10.10.1 on router R3. Add it in OSPF without using a network statement or redistribute connected.
18.4.7 Configure loopback 44 on R4 and add it in OSPF area 44. Advertise it throughout the network with its original mask.
18.4.8 Make sure all OSPF routers can reach networks 151.10.10.0 and 151.10.20.0 using the shortest possible OSPF path.

18.5 RIP

18.5.1 Configure RIP version 1 between R1, R3 and R5.
18.5.2 Make RIP send updates 3 times more often than by default.

18.6 EIGRP

18.6.1 Configure EIGRP AS 100 on the link VLAN50 between R1 and CAT2
18.6.2 Allow only one prefix redistributed from ISIS into EIGRP. Filter all other prefixes. Do not redistribute any other routing protocol into EIGRP AS 100
18.6.3 Do not use any summarization in the EIGRP AS 100 domain

18.6.4 Prevent EIGRP queries from being sent to router R1.
18.6.5 CAT2 should be able to ping the rest of the network

18.7 ISIS

18.7.1 Configure R2 in the ISIS area 49.0002. Peer R2 with R6. Make sure routers maintain minimum routing information.
18.7.2 Configure R6 and R1 in the same area 49.0001 and make sure routers maintain minimum routing information
18.7.3 Use the default ISIS Level-type configuration on R6.
18.7.4 R1 should have its gateway set to R6. R2 should not have the gateway of last resort set to R6
18.7.5 Advertise loopback 66 without an ISIS interface command or redistribute command.
18.7.6 Advertise loopback 140 from R6. Advertise it throughout the network with its original mask.
18.7.7 Make sure all OSPF speakers receive all ISIS networks

18.8 BGP

18.8.1 Use synchronization method in this section
18.8.2 Configure BGP AS 20 on R2 and add network 140.10.1.0/25 in the R2 bgp routing process.
18.8.3 Configure BGP AS 134 on routers R1 and R3
18.8.4 Configure BGP AS 50 on router R5
18.8.5 Peer AS20 and AS134 between routers R1 and R2
18.8.6 Peer AS134 and AS50 between routers R3 and R5
18.8.7 Make sure that network 140.10.1.0/25 resides in the bgp table on R5

18.9 DLSW

18.9.1 A channel gateway with the MAC address 0000.0c4e.54a3 is connected to VLAN 40
18.9.2 Configure a DLSW peer relationship between R1 and R3 using the minimal-overhead option for transport.
18.9.3 Configure a DLSW peer relationship between R1 and R4 using a reliable IP transport
18.9.4 R1 should send DLSW traffic to the channel gateway via R3. Accomplish this without modifying the R1 DLSW configuration.

18.10 ISDN

18.10.1 Activate the ISDN interfaces on routers R1 and R2.
18.10.2 Use CHAP as the authentication method.
18.10.3 Configure ISDN to backup the serial link on R2. Activate ISDN if R2 loses the VPN network 151.10.55.0/24. Do not use a backup interface or dialer watch. Do not run any routing protocol over ISDN.
18.10.4 Make sure R1 initiates a call to R2 only after R2 has called R1.

18.11 VPN

18.11.1 Configure a VPN between routers R2 and R5, use RIP verion 2 as the VPN routing protocol. Make R2 and R5 adjacent VPN neighbors. Use network 151.10.50.0/24 between R2 and R5. Make sure that all VPN traffic is terminated between below mentioned loopbacks.
18.11.2 Configure the VPN network 151.10.52.0/24 on a loopback interface of R2. Terminate the endpoints of the VPN network 151.10.55.0.0/24 using the same loopback interfaces specified in the Security Section below.

18.11.3 Make sure the VPN does not introduce any suboptimal routing paths.

18.12 Security

18.12.1 Use 151.10.100.2/32 and 151.10.100.5/32 loopbacks for IPSEC traffic. Use Encapsulation Security Payload with the 56-bit DES encryption algorithm. Configure a Pre-Shared Key and use Message Digest for ISAKMP.
18.12.2 Make sure networks 151.10.52.0/24 and 151.10.55.0.0/24 communicate in an encrypted manner.
18.12.3 The payload of encrypted packets should not contain minimum IP overhead information

18.13 VOICE

18.13.1 Using a 3 digit dial-plan, configure two phones connected to R1 using numbers 101 for port 1/0/0 and 102 for port 1/0/1.Configure the phone connected to port 1/0/0 of router R3 using number 301
18.13.2 Configure VOIP between R1 and R3 to provide voice communications.
18.13.3 When you dial the number "5" digit by itself on the phone assigned to #101, the phone #102 will ring.
18.13.4 When you dial the number "5" digit by itself on the phone assigned to #301, the phone #101 will ring.

18.14 QOS

18.14.1 Frame size of protocol X is 500 bytes
18.14.2 Frame size of protocol Y is 300 bytes
18.14.3 Frame size of protocol Z is 200 bytes
18.14.4 Allocate 10% of the bandwith for protocol X
18.14.5 Allocate 65% of the bandwith for protocol Y
18.14.6 Allocate 25% of the bandwith for protocol Z

In this task you should calculate the byte count for each protocol. The configuration does not have to be applied to any router.

18.15 Catalyst specialties

18.15.1 On CAT1, ports Fa0/15, Fa0/16, Fa0/17 and Fa0/20 should be configured according to the following specifications:

- Describe port as: Testing Fast Ethernet? Why not.
- During periods of congestion, the port should be able to receive "pause link operation" frames

18.15.2 On CAT1, ports Gig0/1, Gig0/2 should be configured according to the following specifications:

- Describe port as: Testing Gig Ethernet? Why not.
- During periods of congestion, the port should be able to send and receive "pause link operation" frames

18.15.3 Apply the configuration to both groups of interfaces without explicitly referring to the interfaces

18.16 Address Administration

18.16.1 From R4, HTTP to the IP address 151.10.10.1 and get the router prompt from R5

18.17 Multicast

18.17.1 Enable multicast routing between routers R1and R4, R4 and R3 as well as R3 and R5. Enable a multicast routing protocol that will use PIM V2 messaging only to build the shared tree.
18.17.2 Configure routers R3, R4 and R5 to join the multicast group 229.20.20.20. Associate this multicast group with a loopback interface on each router.
18.17.3 Make R3 the root of the shared tree. Accomplish this task by configuring only R3.
18.17.4 Ping the multicast group 229.20.20.20 from R1.

Scenario 19. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 172.10.0.0/16
- Do not use any static routes except those used for backup purposes
- Advertise Loopback interfaces with their original mask
- Do not change any prefix masks
- Do not use policy-routing
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks advertised in the BGP section must be reachable only in the BGP domain

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN30
R2	Ethernet0	VLAN10
R3	-	-
R3	FastEthernet0/0	VLAN10
R4	-	-
R5	Ethernet0	VLAN20
R6	FastEthernet0/0	VLAN20 VLAN30
FRS	Ethernet0	VLAN10
CAT1	-	VLAN10
CAT2	-	VLAN20

Attention **For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15**

Scenario 19

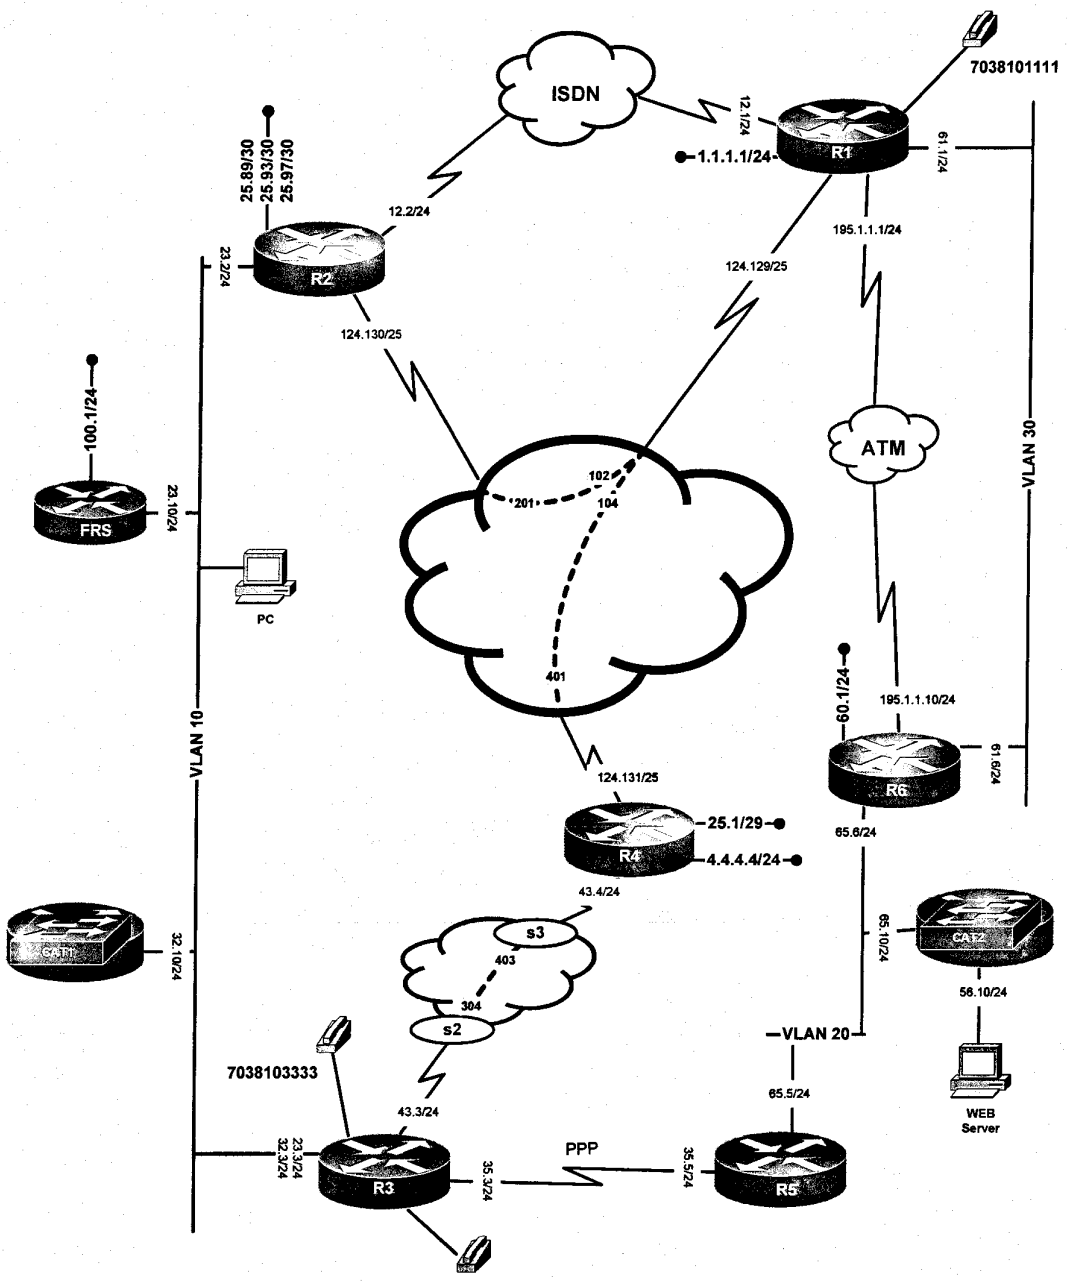

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

19.1 Frame Relay

19.1.1 Use ***only the PVC's displayed on the diagram*** to fulfill this configuration.
19.1.2 Use physical interfaces on subnet 172.10.124.128/25. Use physical and point-to-point interfaces on the subnet 172.10.43.0/24.

19.2 Catalyst Configuration

19.2.1 Do not advertise VLAN's in this scenario. Use the ISL protocol for trunking.
19.2.2 Configure the FRS Ethernet interface with an IP address of 172.10.23.10/24.
19.2.3 Make CAT1 the root bridge for VLAN 10. Put the Fa0/13 interface on CAT2 in the Spanning Tree blocking state. Do not use spanning tree path-cost manipulation to accomplish this task.
19.2.4 A PC with the NIC MAC address 00-07-85-92-D0-E7 is connected to port Fa0/10 of CAT2 on the default VLAN. Make sure only that PC is allowed to access port Fa0/10.

19.3 ATM

19.3.1 Configure an ATM PVC 100/100 connection between R1 and R6.
19.3.2 See the diagram for the ip addresses assignments.

19.4 OSPF

19.4.1 Place PPP link between R3 and R5 in area 0.
19.4.2 Configure OSPF area 10 on VLAN 20 between R5 and R6. Do not elect a DR/BDR on this subnet. Make sure OSPF packets are exchanged without use of a multicast address due to security reasons.
19.4.3 Assign the following IP address,172.10.60.1/24, to a loopback interface on R6 and place it in area 600
19.4.4 Use clear text authentication on area 0. Password "nmc"
19.4.5 Use md5 authentication on area 10. Password "rsnmc"

19.5 RIP

19.5.1 Configure RIP version 2 between R3 and R4.
19.5.2 RIP speakers should not multicast the updates.
19.5.3 Configure RIP version 1 between only R5 and CAT2. Set the gateway of last resort on CAT2 only if 172.10.124.128/25 is in the R5 routing table.

19.6 EIGRP

19.6.1 Configure EIGRP AS 124 between routers R1, R2 and R4.
19.6.2 Do not allow multicast EIGRP traffic on the subnet between routers R1, R2 and R4
19.6.3 Advertise Loopbacks into EIGRP AS124 on R2, see the diagram:

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

Summarize these loopbacks with the most optimal mask.

19.6.4 Make sure you have a summary for these loopbacks on R1 and R4 only.
19.6.5 Advertise the loopback 172.10.25.1 mask 255.255.255.248 on R4 as an external EIGRP prefix:

- ip address 172.10.25.1 255.255.255.248

19.6.6 Change the administrative distance for this prefix to 1 higher than default on router R2
19.6.7 Advertise the following loopback on R4 as internal EIGRP prefix:

- ip address 4.4.4.4 255.255.255.0

19.6.8 Advertise the following loopback on R1 as internal EIGRP prefix:

- ip address 1.1.1.1 255.255.255.0

19.6.9 Configure EIGRP AS124 on the VLAN 10 subnet 172.10.23.0/24 between R2, R3 and FRS.
19.6.10 EIGRP AS124 on FRS must be configured with the "network 172.10.0.0" statement. Make sure FRS does not advertise loopback 172.10.100.0/24. Do not use any route filtering techniques.
19.6.11 Ping 172.10.100.1 from the rest of the network using address 172.10.23.100
19.6.12 Configure subnet 172.10.32.0/24 on VLAN 10 between R3 and CAT1. CAT1 should not run the "ip routing" process. CAT1 should be reachable from the rest of the network

19.7 ISIS

19.7.1 Configure NET 49.0001.1111.1111.1111.00 on R1
19.7.2 Form an adjacency between R1 and R6 over the ATM link. The NET of R6 is 49.0006.6666.6666.6666.00
19.7.3 Authenticate ISIS adjacencies on the links between R1 and R6 using password "nmc".
19.7.4 Perform redistribution between OSPF and ISIS on R6
19.7.5 Do not redistribute between EIGRP and ISIS

19.8 BGP

19.8.1 Configure AS 23 between R2 and R3. Configure AS 4 on R4 .
19.8.2 Peer AS 23 and AS 4 between R2 and R4 as well as R3 and R4. Do not use loopback interfaces for peering.
19.8.3 Advertise network 4.4.4.0/24 in AS4.
19.8.4 Advertise network 172.10.23.0/24 in AS23.
19.8.5 Outbound traffic from a PC connected to the 172.10.23.0/24 subnet destined to the 4.0.0.0/24 network should flow through R2.
19.8.6 Incoming traffic from the 4.0.0.0/24 network to a PC connected to the 172.10.23.0/24 subnet should flow through R2.
19.8.7 If the frame-Relay link on R2 goes down, the aforementioned traffic should pass through R3.
19.8.8 Return the traffic pattern through R2 when frame relay link on R2 is back up.
19.8.9 Use the minimal number of bgp decision steps to accomplish this task
19.8.10 Configure AS 1 on R1 and peer with AS4 using loopbacks 1.1.1.1 and 4.4.4.4
19.8.11 Advertise networks 1.1.1.0/24 and 4.4.4.0/24 at AS1 and AS4 respectively

19.9 DLSW

19.9.1 Configure a DLSW peer relationship between R5 and R4.
19.9.2 Guarantee the average bit rate for end-to-end DLSW traffic to 20Kbps allowing bursts not higher than 10 Kbps

19.10 ISDN

19.10.1 Activate the ISDN interfaces on routers R1 and R2.
19.10.2 Use CHAP as the authentication method. Use a CHAP method which does not require challenges sent from both sides. Use password "nmc100". Do not use the default host names. R2 should initiate the call.
19.10.3 Configure ISDN to backup the serial link on R2. Activate ISDN and provide reachability to 172.10.60.0/24 in the event that the R2 Frame-Relay and Ethernet links fail. Do not use backup interface or dialer watch. Do not run any routing protocol over ISDN.

19.11 Traffic Optimization part 1

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

19.12 Traffic Optimization part 2

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

19.13 VOICE

19.13.1 The dial plan is as follows:

- Assign #7038101111 to the phone on port 1/0/0 on R1
- Assign #7038103333 to the phone on port 1/0/0 on R3

19.13.2 Call from R3 to R1 without dialing the entire number, use only three last digits to accomplish this task

19.14 QOS

19.14.1 Reserve bandwidth for telnet application traffic between R4 and R5. The telnet traffic is sourced from 172.10.43.4 port 5000 and destined to 172.10.35.5 port 23.

19.14.2 Send PATH messages from R4 to R5 and make sure a single reservation for a Guaranteed Bit Rate of 5 kbps allowing bursts up to 2 Kbytes is established on R3.

19.15 Catalyst Specialties

19.15.1 Prohibit all traffic of ethertype 8042 from entering VLAN30
19.15.2 Prohibit all SNA traffic from entering VLAN10
19.15.3 Do not apply any filtering configurations to Catalyst switchports to accomplish this task.

19.16 Gateway Redundancy

19.16.1 Assign the IP address of 172.10.23.1 to the virtual gateway and make sure the mac-address associated with the virtual gateway is set to 0000.0c07.ac14
19.16.2 Authenticate HSRP on the 172.10.23.0/24 subnet (password nmc). Make sure Hello packets are exchanged 3 times faster than by default.
19.16.3 Select the prefered gateway that is most suitable for other tasks of this exam by using priority 150.

19.17 Multicast

19.17.1 Configure dense mode cast routing between R4, R3 and R5
19.17.2 Join dense group 229.50.50.50 on interfaces S1 and E0 of R5.
19.17.3 Make sure you can ping 229.50.50.50 from R1.

Scenario 20. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 112.10.0.0/16
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- Do not use default-information originate and ip default-network commands in RIP, OSPF and BGP domains
- Do not change any prefix masks, unless specified otherwise
- Do not use policy-routing
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks advertised in the BGP section must be reachable only in the BGP domain

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10 VLAN11
R2	Ethernet0	VLAN100
R3	FastEthernet0/1	VLAN10
R3	FastEthernet0/0	VLAN10
R4	Ethernet0	VLAN101
R5	-	-
R6	FastEthernet0/0	VLAN10 VLAN111
FRS	Ethernet0	VLAN101
CAT1	-	VLAN60
CAT2	-	VLAN11 VLAN111

Attention **For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15**

Scenario 20

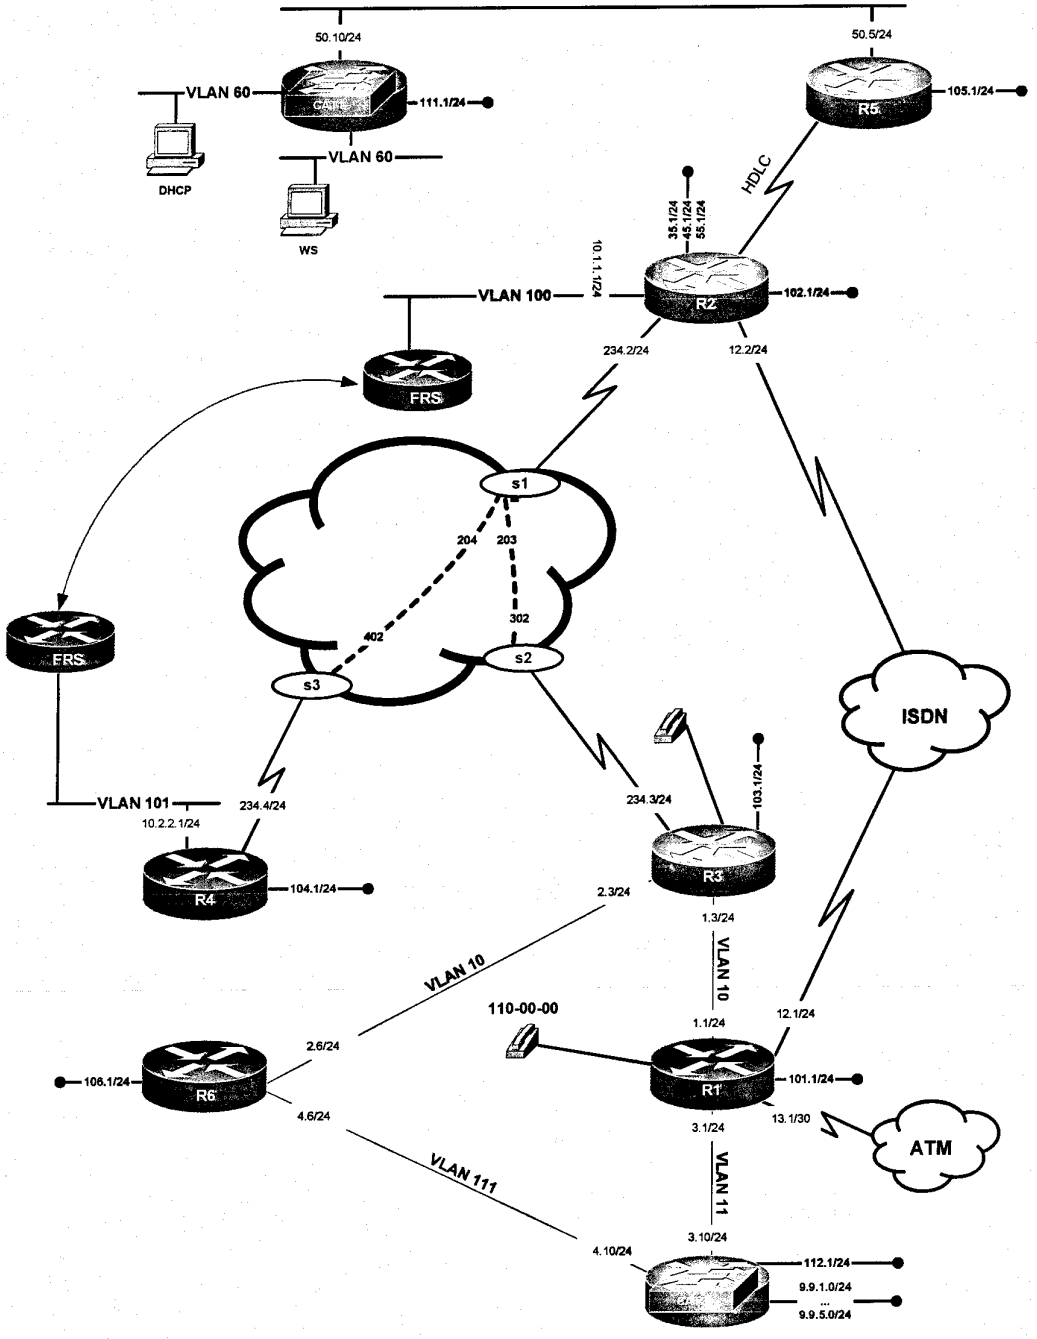

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

20.1 Frame Relay

20.1.1 Use ***only the PVC's displayed on the diagram*** to fulfill this configuration.
20.1.2 Use physical interfaces on R4 and R3 on the subnet 112.10.234.0/24. Use a logical interface on R2 on the subnet 112.10.234.0/24..

20.2 Catalyst Configuration

20.2.1 Use dot1q trunking protocol for all trunks involved in this scenario, see the diagram.
20.2.2 Configure VLANs according to diagram
20.2.3 Assign 112.10.1.0/24 subnet on VLAN10 between R3 and R1
20.2.4 Assign 112.10.2.0/24 subnet on VLAN10 between R3 and R6
20.2.5 Packets generated on the 112.10.1.0/24 and 112.10.2.0/24 subnets should be isolated within the VLANs representing these subnets.
20.2.6 Configure VLANs and subnets between R1, R6 and CAT2 according to diagram.
20.2.7 Configure a link between R5 and CAT1 without configuring any VLAN.

20.3 ATM

20.3.1 Configure the IP subnet 112.10.13.0/30 on a point-to-point ATM link using PVC 200/200.
20.3.2 Set the Peak Cell Rate to 128 Kbps, Sustainable Cell Rate to 64 Kbps , allow bursts to up to 16 cells on the PVC 200/200

20.4 OSPF

20.4.1 Place 112.10.1.0/24 and 112.10.2.0/24 in area 0.
20.4.2 Configure area 26 on the network 112.10.6.0/24 between routers R2 and R6.
20.4.3 Place the HDLC link between R2 and R5 in area 0.
20.4.4 Configure area 10 on the link between R5 and CAT1.

20.5 RIP

20.5.1 Configure RIP version 2 between R1, R2, R3 and R4
20.5.2 Summarize networks:

- 172.16.35.0/24
- 172.16.45.0/24
- 172.16.55.0/24

Use the most efficient mask

20.5.3 Do not allow RIP updates on the 112.10..2.0/24 link.
20.5.4 Configure RIP version 2 on the ISDN link between R1 and R2

20.6 EIGRP

20.6.1 Configure EIGRP AS 10 between routers R1, R6 and CAT2.
20.6.2 Adjust EIGRP to make sure that CAT2 sends 75% of outbound traffic to R1 and 25% of outbound traffic to R6.

20.7 Catalyst Specialties

20.7.1 The DHCP server and workstations are connected to ports Fa0/20 and Fa0/21 on VLAN 60. Workstation users complain about not always getting an IP address and the associated workstation configuration from the DHCP server. Provide a solution.
20.7.2 Make sure the provided solution prevents a network meltdown if other switches are accidentally added to VLAN 60 and create a spanning tree loop.

20.8 BGP

20.8.1 Use the most efficient method of configuration to reduce the load on router system resources.
20.8.2 Configure AS 100 on CAT2 and advertise networks 9.9.1.0/24 through 9.9.5.0/24
20.8.3 Configure AS 1000 on R1, AS 600 on R6, AS 300 on R3 and AS 200 on R2
20.8.4 Form all BGP peer relationships between the management loopbacks. Peer AS 100 with AS 1000, peer AS 1000 with all other ASes.
20.8.5 Allow all networks coming from CAT2 to R1.
20.8.6 Allow only odd networks to all other BGP speakers. Do not use any filters to accomplish this task.
20.8.7 Set the MED to 35 for all prefixes that are allowed.

20.9 DLSW

20.9.1 Set up DLSW peers between R2 and R3 as well as R4 and R3. Use loopback interfaces to form the peer relationships.
20.9.2 Make sure that traffic coming from peers between R2 and R3 is bridged to the R3 Fa0/0 interface.
20.9.3 Make sure that traffic coming from peers between R4 and R3 is bridged to the R3 Fa0/1 interface.

20.10 ISDN

20.10.1 Activate the ISDN interfaces on routers R1 and R2.
20.10.2 Use CHAP as the authentication method. Use password "nmclass". R2 should initiate the call.
20.10.3 Configure ISDN to backup the PVC between R2 and R3. Do not use a backup interface and a dialer watch.
20.10.4 Start using both B channels as soon as possible
20.10.5 Make sure ISDN does not stay up indefinitely.

20.11 Address Administration

20.11.1 Simulate a host on FRS by disabling IP routing (no ip routing in global configuration mode)
20.11.2 Assign the IP address 112.10..5.100 to the FRS Ethernet 0 interface. Ping R2 from FRS
20.11.3 Move FRS from VLAN101 to VLAN100 and ping R2 from FRS without changing FRS's IP address.

20.12 Security

20.12.1 SMTP packet exchange is occuring via R2's Ethernet interface. Configure R2 to check SMTP packets for illegal commands. Drop the packet if it contains an illegal SMTP command. Hang the session and time it out in 300 seconds.

20.12.2 The network adminnistrator is concerned about TCP SYN-flooding attacks to the TCP servers located on VLAN 11. Provide a solution to prevent this type of attack.

20.12.3 Allow TCP sessions to be continued without any traffic passing through router R2 for not more than 200 seconds.

20.13 VOICE

20.13.1 Configure Voice over IP on the ATM link. Use PVC 200/200 to configure IP subnet 112.10.13.0/30. A remote phone is connected to a device on the other side of ATM link. The local phone number is 110-00-00. The remote phone number is 110-00-01.

20.13.2 Make sure you guarantee 20Kbps for voice and use the rest for the other types of traffic. LLQ should be used to accomplish this task

20.13.3 Make sure the voice channel consumes the lowest amount of bandwidth

20.14 QOS

20.14.1 Configure ip addresses on the Ethernet interfaces of R2 and R4 and connect them to the corresponding VLANs. See the diagram.

20.14.2 A traffic generator is connected to VLAN 101. It generates 8 UDP streams with precedence settings ranging from 0 through 7. The destination IP address of these streams is 10.1.1.1.

20.14.3 Allocate 80% of bandwidth for the UDP flows with immediate and critical ip precedence.

20.14.4 Allocate 20% of bandwidth for the other types of traffic.

20.15 Router Administration

20.15.1 Log only alerts and emergencies from CAT1 to the server with the IP address 112.10.50.100. Use the facility local6 to identify the syslog messages.

20.15.2 Create a message of the day banner on router R3. Display the hostname of the router and a line number in the banner when you login onto the router.

20.16 Multicast

20.16.1 Configure a multicast domain based on flood and prune technology between routers R1, R2, R3, R4 and R5 for the group 225.5.5.5. ISDN is not involved in this configuration.

20.16.2 Hosts residing on both VLAN 11 and the LAN segment between CAT1 and R5 are running an application that generates broadcast traffic destined to UDP port 5000.

20.16.3 Using the application mentioned above, hosts residing on the segment between CAT1 and R5 communicate with hosts on VLAN11 by broadcasting UDP datagrams. Broadcast forwarding is not allowed between above specified LANs. Provide a configuration that would allow the above described application to successfully operate given the broadcast forwarding restriction.

Scenario 21. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 141.37.0.0/16
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- Network 0.0.0.0/0 should not appear in any routing table (show ip route)
- Do not change any prefix masks, unless specified otherwise
- Do not use policy-routing
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks advertised in the BGP section must be reachable only in the BGP domain

VLAN Configuration Table

Router	Interface	VLAN
R1	-	-
R2	-	-
R3	-	-
R3	-	-
R4	Ethernet0	VLAN10
R5	Ethernet0	VLAN20
R6	FastEthernet0/0	VLAN10
FRS	Ethernet0	VLAN30
CAT1	-	VLAN20 VLAN30
CAT2	-	VLAN20 VLAN30

Attention For physical connectivity, check the diagram "NMC Pod Layout", Figure 2 on page 15

Scenario 21

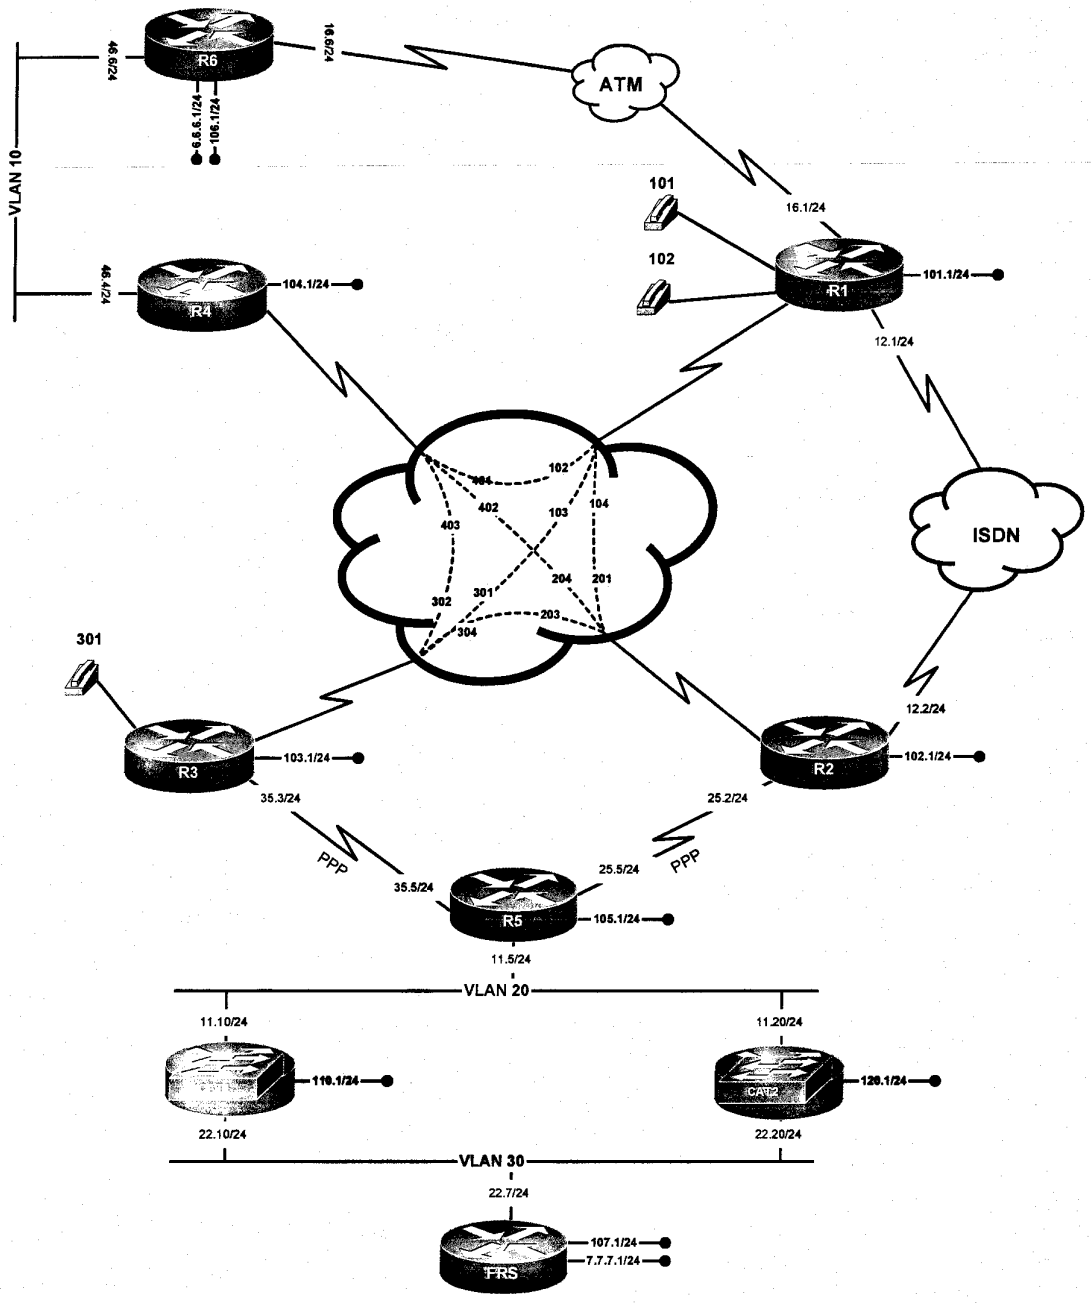

21.1 Serial interfaces

21.1.1 Configure the Frame Relay subnet 141.37.2.0/24 between routers R1, R2, R3 and R4. Assign the router hostname number for the host part of the IP address for interfaces attached to the 141.37.2.0/24 subnet. Configure the Frame Relay subnet 141.37.1.0/24 between routers R1, R2 and R4. Assign the router hostname number for the host part of the IP address for interfaces attached to the 141.37.1.0/24 subnet.
21.1.2 R3 must have only one physical interface
21.1.3 All other routers will have two interfaces. One them must be a logical point-to-point interface.
21.1.4 Use a full-mesh of PVCs on the 141.37.1.0/24 subnet. Static Frame-Relay mapping and Frame-Relay interface-dlci commands are not allowed on the 141.37.1.0/24 subnet.
21.1.5 Configure PPP encapsulation on the serial links between R2and R5 as well as R3 and R5.

21.2 Catalyst Configuration

21.2.1 Configure VLAN's according to the diagram and VLAN table

21.3 ATM

21.3.1 Configure ATM PVC 100/100 between R6 and R1. Use physical interfaces on this link

21.4 OSPF

21.4.1 Configure OSPF area 0 on VLAN 10 between routers R4 and R6.
21.4.2 Configure OSPF area 10 on the ATM link between R6 and R1
21.4.3 Configure OSPF area 100 on the subnet 141.37.1.0/24
21.4.4 Advertise Loopback networks 141.37.104.0/24 and 141.37.102.0/24 in OSPF area 100 on routers R2 and R4 respectively
21.4.5 Form a full mesh of OSPF adjacencies between R1, R4 and R2 on the subnet 141.37.1.0/24. R4 should not advertise OSPF routes to R2 and vice versa. (To test it, shutdown R1's frame relay interface and clear the ospf process on routers R2 and R4. Check the OSPF database content)
21.4.6 Configure OSPF area 100 on the subnet 141.37.2.0/24. Use the OSPF non-broadcast network type.
21.4.7 Advertise Loopback networks 141.37.106.0/24 and 141.37.101.0/24 in OSPF area 106 and 101 on routers R6 and R1 respectively

21.5 RIP

21.5.1 Configure RIP version 2 on routers R3, R2 and R5.
21.5.2 Perform the mutual redistribution between OSPF and RIP on routers R2 and R3 with the same metric value.
21.5.3 Send RIP updates only on the PPP links
21.5.4 Accept the RIP updates only if they are authenticated with the clear text key "whynmc?"
21.5.5 Have router R5 prefer R2 for the destination network 141.37.106.0/24 and prefer R3 for the destination network 141.37.101.0/24
21.5.6 The configuration should be applied on router R5 to accomplish the previous task
21.5.7 Configure RIP version 2 between CAT1, CAT2 and FRS. Exchange the updates on VLAN30 only
21.5.8 Advertise 141.37.107.0/24 from FRS into RIP
21.5.9 Traffic destined to the network 141.37.107.0/24 should transit via CAT2
21.5.10 When network 141.37.120.0/24 is down, traffic destined to the network 141.37.107.0/24 should transit via CAT1

21.5.11 When network 141.37.120.0/24 is back up, the traffic flow should return to CAT2
21.5.12 The traffic flow control solution for the previous tasks should not be based on any metric manipulation and filtering.

21.6 EIGRP

21.6.1 Configure EIGRP AS 20 on the interfaces of VLAN 20 only
21.6.2 Advertise the Loopback networks 141.37.105.0/24 and 141.37.120.0/24 in EIGRP AS 20

21.7 ISIS

21.7.1 Configure ISIS on the ISDN link between routers R1 and R2
21.7.2 Configure NET 49.1111.1111.1111.00 on R1 and 49.2222.2222.2222.00 on R2
21.7.3 Form Level-2 adjacency only, do not use the default ISIS router type.
21.7.4 Advertise loopback networks 141.37.101.0/24 and 141.37.102.0/24 in ISIS on routers R1 and R2 respectively

21.8 BGP

21.8.1 Configure AS 700 on routers FRS, CAT1 and CAT2
21.8.2 Allow IBGP learned prefixes to be forwarded to another IBGP peer on router FRS
21.8.3 Routers R1, R2, R3, R4 and R5 comprise AS 1000
21.8.4 Configure AS 400 on router R4
21.8.5 Configure AS 100 on router R1
21.8.6 Configure AS 235 on routers R2, R3 and R5
21.8.7 Prefixes must be exchanged between all routers in AS 1000 via R5
21.8.8 Configure AS 65000 on router R6
21.8.9 Peer AS 65000 and AS 1000 between routers R6 and R4 as well as R6 and R1
21.8.10 Peer AS 700 and AS 1000 between routers CAT1 and R3 as well as CAT2 and R2
21.8.11 Advertise Loopback networks 7.7.7.0/24 and 6.6.6.0/24 from FRS and R6 respectively
21.8.12 Prefix 6.6.6.0/24 should show up in the BGP tables of AS 700 as if it was originated from AS 1000
21.8.13 R6 should receive the 7.7.7.0/24 prefix from R4 only. Do not use filtering techniques based on AS PATH and prefixes. The configuration must be applied on the peer between R5 and R1. All routers in AS1000 should possess 7.7.7.0/24 in their BGP tables as well as in their local routing tables as a "B" entry.
21.8.14 Make sure transit service for the networks 6.6.6.0/24 and 7.7.7.0/24 is provided. Verify it with a ping.

21.9 IP Accounting

21.9.1 Configure IP accounting on router R5 to generate the following output: the number of packets and bytes based on IP precedence.

21.9.2 To test your configuration, clear counters on R5, send 10 ping packets destined to the network 141.37.101.0/24 and 5 ping packets destined to the network 141.37.106.0/24 from the FRS
21.9.3 Make sure R5 displays accurate statistics based on IP precedence

21.10 ISDN

21.10.1 The ISDN link should be activated if R1 loses the prefix 141.37.107.0/24

21.10.2 R1 is the only router that can call. Make sure R1 sends the PPP CHAP challenge only.
21.10.3 Use password "exam21"

21.11 Router Maintenance

21.11.1 On R1, configure user admin with the password "cisco"
21.11.2 When user admin logins via telnet to the router R1, the following options should appear on the screen:

1 Display Routes
2 Exit

21.11.3 When number 1 is entered, the content of the routing table should appear on the screen
21.11.4 When either number 2 or the CR (Enter) is typed, router R1 should give an exec prompt

21.12 Security

18.13.1 On R3, configure authentication for telnet users. Allow access to router R3 to only the users that pass the authentication. Do not use the command "login local", Do not configure passwords anywhere, except with the username command.
18.13.2 Use the following username and password for this section:

Username: doit
Password: cisco

18.13.3 Make sure user "doit" can access the IOS enable mode

21.13 VOICE

21.13.1 Using a3 digit dial-plan, configure two phones connected to R1 using numbers 101 for port 1/0/0 and 102 for port 1/0/1. Configure the phone connected to port 1/0/0 of router R3 using number 301
21.13.2 Configure VOIP between R1 and R3 to provide voice communications
21.13.3 Configure R3 so that a call from R1 port 1/0/0 to R3 uses a g729br8 CODEC. All other calls should use a g728 CODEC.
21.13.4 R3 should forward all calls to numbers starting with the digit "1" to R1. Use only one dial-peer to accomplish this task.

21.14 QOS

21.14.1 For ping packets sourced from FRS and destined to the following destinations, set the IP precedence:

141.37.101.0/24	precedence	2
141.37.106.0/24	precedence	5

21.14.2 IP precedence must be set on the Catalyst switches.

21.15 Catalyst Specialties

21.15.1 Use the dot1q protocol for all trunks

21.15.2 Use VLAN 40 to send untagged Ethernet frames on all trunk interfaces
21.15.3 Configure CAT2 to be the server of VTP domain "doit" and CAT1 to be the client. Use password cisco to exchange VTP messages
21.15.4 MAC addresses should be stored in the tables for no longer than 20 seconds

21.16 Gateway Redundancy

21.16.1 Configure HSRP group 20 between CAT1 and CAT2
21.16.2 Make CAT2 the active gateway. CAT1 and CAT2 should be both suitable for the ACTIVE/STANDBY state.
21.16.3 The IP address 141.37.22.1 should be used for the virtual gateway
21.16.4 Reserve the name "doit" for this HSRP peer relationship

21.17 Multicast

21.17.1 Configure IP multicast based on push technology using the unicast routing information between R6, R1, R2, R3, R4 and R5. ISDN is not part of this exercise.
21.17.2 Configure R6, R1 and R3 to deliver data using the reliable multicast transport protocol for hosts running applications requiring reliable delivery
21.17.3 Use a virtual interface to exchange packets from the reliable multicast transport protocol
21.17.4 Join the following loopback interfaces to the group 225.21.21.21:

- 141.37.101.0/24
- 141.37.102.0/24
- 141.37.103.0/24
- 141.37.104.0/24
- 141.37.105.0/24
- 141.37.106.0/24

21.17.5 Ping 225.21.21.21 from CAT1

Scenario 22. Multi-topic CCIE level exam

Goals and Restrictions

- IP subnets on the diagram belong to network 135.15.0.0/16
- Do not use any static routes
- Advertise Loopback interfaces with their original mask
- Do not use default-information originate and ip default-network commands in RIP, OSPF and BGP domains
- Do not change any prefix masks, unless specified otherwise
- Do not use backup interface, dial watch, or snapshot routing
- All IP addresses involved in this scenario must be reachable, unless specified otherwise
- Networks advertised in the BGP section must be reachable only in the BGP domain

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN10
R2	Ethernet0	VLAN20
R3	FastEthernet0/1	-
R3	FastEthernet0/0	VLAN10
R4	Ethernet0	VLAN30
R5	Ethernet0	VLAN30
R6	FastEthernet0/0	VLAN20
FRS	Ethernet0	VLAN30
CAT1	-	VLAN30
CAT2	-	VLAN30

Attention For physical connectivity, check the diagram "NMC Pod Layout" , Figure 2 on page 15

Scenario 22

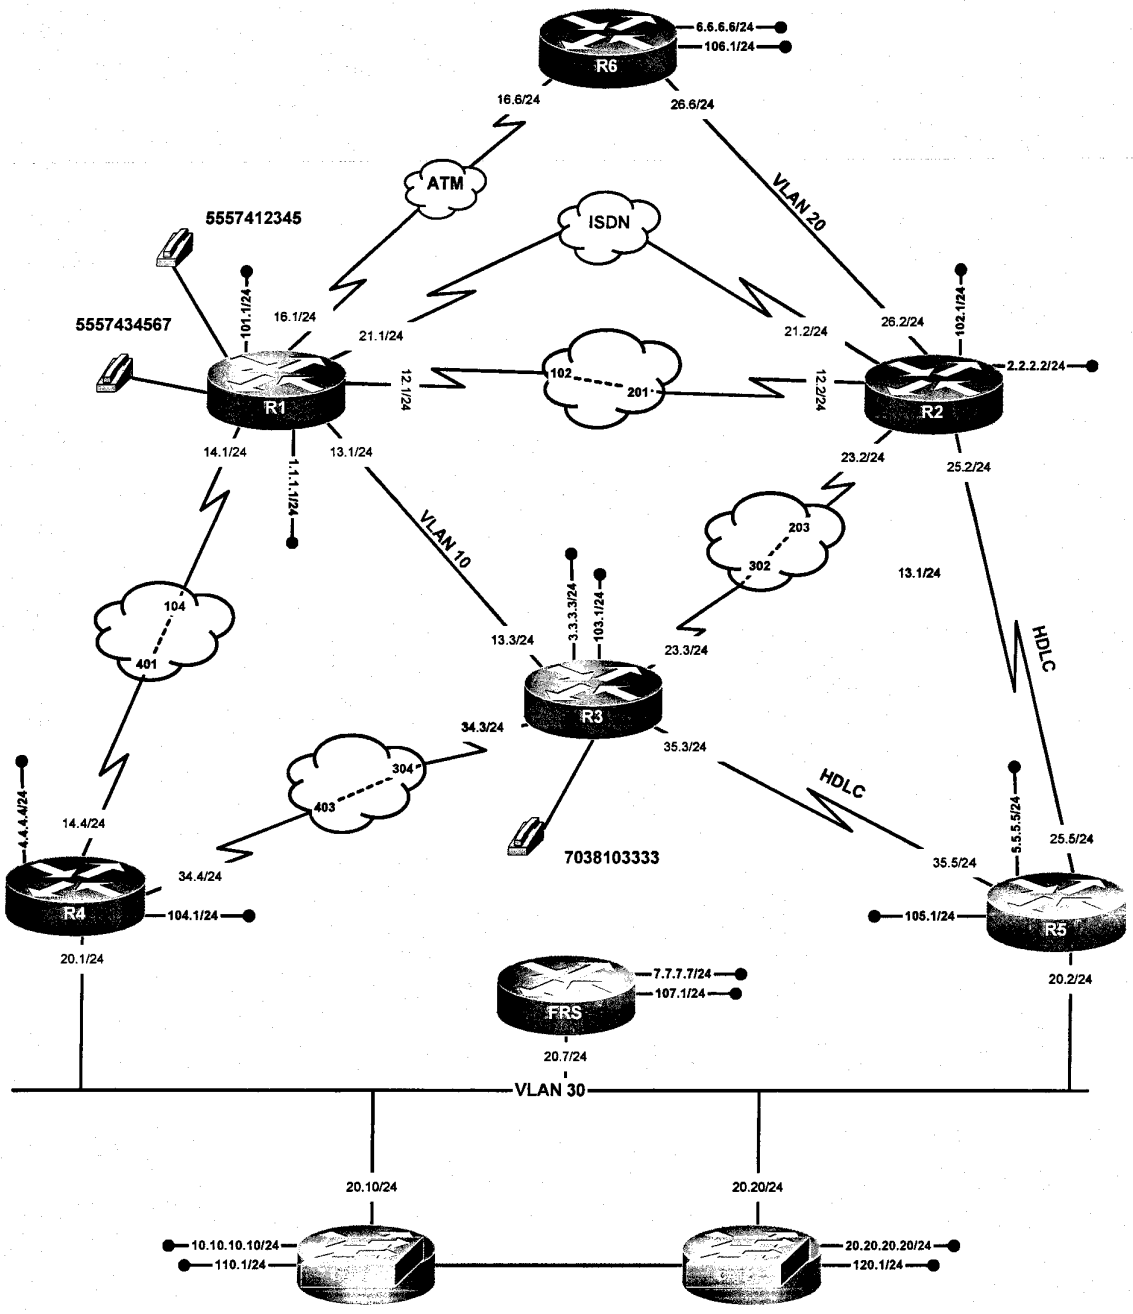

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

22.1 Serial Interfaces

22.1.1 Configure Frame Relay subnets according to the diagram
22.1.2 Configure physical Frame Relay interfaces on the subnet 135.15.23.0/24
22.1.3 Configure physical and point-to point logical interfaces on the subnet 135.15.34.0/24
22.1.4 Configure point-to-point and multipoint interfaces on the subnet 135.15.14.0/24; the point-to-point interface should be placed on R1
22.1.5 Configure point-to-point interfaces on the subnet 135.15.12.0/24
22.1.6 Only PVC's displayed on the diagram should be used in this scenario
22.1.7 Configure HDLC encapsulation on the serial links between R5 and R2 as well as R5and R3. Configure the clock to 1Mbit/sec for these HDLC links.

22.2 Catalyst Configuration

22.2.1 Configure VLAN's according to the diagram and VLAN table
22.2.2 Use the ISL protocol for all trunks you create

22.3 ATM

22.3.1 Configure a PVC between R6 and R1. Use a physical interface on R1 and point-to-multipoint interface on R6
22.3.2 Use the following VPI/VCI pairs:

- R1 100/100
- R6 100/100

22.3.3 Configure Unspecified Rate with the QOS parameters provided in the table below:

PCR	256 Kbits
MCR	128 Kbps

22.4 OSPF

22.4.1 Assign a new network that will be shared by R4, R5 and FRS. The network number is 135.15.10.0/24. Use the following host numbers for the devices:

- R4 .4
- R5 .5
- FRS .7

22.4.2 The solution should not use bridging techniques.
22.4.3 Configure OSPF area 0 between R4, FRS and R5 on the network 135.15.10.0/24
22.4.4 The network type that elects a DR/BDR and does not use the unicast packet exchange should be used. R5 should be configured with the higher chance to become the DR and R4 the BDR.
22.4.5 Configure OSPF area 10 on the links 135.15.13.0/24, 135.15.14.0/24 and 135.15.34.0/24. OSPF network type point-to-point should be used on 135.15.13.0/24 and 135.15.14.0/24 and OSPF network type broadcast should be used on 135.15.34.0/24 with the R4 as the DR.

22.4.6 Configure OSPF area 20 on the link 135.15.12.0/24, use the non-broadcast network type, originate OSPF traffic from R1 only on this subnet

22.4.7 The adjacency on the 135.15.14.0/24 link should be authenticated with a cleartext password. The adjacency on the 135.15.34.0/24 link should be authenticated with MD5

22.5 RIP

22.5.1 Configure RIP version 2 between routers R2, R3 and R5

22.5.2 Advertise updates only on the HDLC and Frame-Relay links.

22.5.3 Configure RIP version 2 on the 135.15.21.0/24 link

22.6 EIGRP

22.6.1 Configure EIGRP AS 30 on all interfaces on VLAN 30

22.7 ISIS

22.7.1 Configure ISIS between R1, R6 and R2. Use the IS-type Level-1 on all routers.

22.7.2 Configure the following NET's on the routers:

- R1 01.1350.1510.1001.00
- R2 01.1350.1510.2001.00
- R6 01.1350.1510.6001.00

22.7.3 R6 should have R1 and R2 as redundant exits to the rest of the network. However, prefer R1.

22.8 BGP

22.8.1 Configure AS 1000 on CAT1; AS 235 on routers CAT2, R3 and R5; AS 47 on the routers R4 and FRS; AS 126 on the routers R1, R2 and R6

22.8.2 You may disable synchronization in this scenario

22.8.3 Do not configure a full mesh in AS 126 and AS 235. Use R6 and R5 to exchange reachability information.

22.8.4 Peer AS 1000 and AS 235 between CAT1 and CAT2

22.8.5 Peer AS 1000 and AS 47 between CAT1 and R4

22.8.6 Peer AS 47 and AS 235 between FRS and R3

22.8.7 Peer AS 47 and AS 126 between R4 and R1

22.8.8 Advertise the following loopback networks:

- 1.1.1.0/24 from R1
- 2.2.2.0/24 from R2
- 3.3.3.0/24 from R3
- 4.4.4.0/24 from R4
- 5.5.5.0/24 from R5
- 6.6.6.0/24 from R6
- 7.7.7.0/24 from FRS
- 10.10.10.0/24 from CAT1
- 20.20.20.0/24 from CAT1

22.9 Traffic control

22.9.1 If traffic is originated from the 135.15.103.0/24 subnet and destined to the 135.15.106.0/24 subnet, it must be forwarded to R2 from R3

22.9.2 This traffic flow must traverse router R2 twice and R1 once in one direction. The term "traverse" is defined as follows: enter the route on one interface and exit the router on a different interface.

22.9.3 This traffic flow must be symmetric.

22.9.4 Do not make forwarding decisions based on the precise source and destination match on R1

22.10 ISDN

22.10.1 Configure ISDN clear text authentication. Use the router hostnames and password "doit"

22.10.2 ISDN should be used only for the ICMP traffic between 135.15.103.0/24, 135.15.106.0/24, 135.15.101.0/24, 135.15.102.0/24 and routing control traffic. Do not use policy based routing for the communication between 135.15.101.0/24 and 135.15.102.0/24.

22.10.3 If there is no other communication between these routers, R1 and R2 must exchange full routing tables over the ISDN link.

22.10.4 R1 and R2 can call each other, make sure ISDN link does not stay up indefinitely.

22.11 Router Access Management

22.11.1 The network administrator would like to have his telnet session to R6's VTY line 6 redirected to R2's VTY line 7.

22.11.2 Configure R6 to perform this task. Redirection should occur only if user admin with the password doit is authenticated on both routers R6 and R2

22.11.3 Allow a telnet session to R6's VTY line 6 only from the 135.35.101.0/24 network

22.12 Security

22.12.1 Configure a user "administrator" in the group "administratorgrp" to be able to read and write to view "doit", if it passes authentication. Use MD5 and password nmc

22.12.2 Configure user "operator" in the group "operatorgrp' to be able to read the view "doit",. No authentication is required for the user operator

22.12.3 Configure security model "V3" for this task

22.13 VOICE

22.13.1 Configure two phones connected to R1 using numbers 5557412345 for port 1/0/0 and 5557434567 for port 1/0/1. Configure the phone connected to port 1/0/0 of router R3 using number 7038103333

22.13.2 Configure VOIP between R1 and R3 to provide voice communications

22.13.3 Configure only one VOIP dial peer on router R3 to be able to ring only the two telephone numbers assigned R1 provided in this scenario.

22.14 QOS

22.14.1 Traffic between the networks 135.15.103 to 135.15.106 should be marked with the IP precedence 2
22.14.2 Accomplish this without use of route maps

22.15 IOS Features

22.15.1 User NMCADMIN should be able to see the R5's output of show run and show version from R4
22.15.2 User NMCOPERATOR should be able to see R5's output of show version but not show run from R4
22.15.3 Telnet and reverse telnet are not allowed

22.16 Network Management

22.16.1 On R1, create an SNMP view "doit" of all objects in the system group except for system 7 objects. Also, add all objects of Cisco private MIB
22.16.2 Configure read only access for "operator" from VLAN20 only.
22.16.3 Configure read/write access for "administrator" from VLAN 30 only

22.17 Multicast

22.17.1 Configure PIM Sparse Mode routing on all routers involved in this scenario
22.17.2 Add all subnets in multicast routing except for the ISDN link.
22.17.3 Join multicast group 225.22.22.22 on the following loopback interfaces:

- 1.1.1.0/24 from R1
- 2.2.2.0/24 from R2
- 3.3.3.0/24 from R3
- 4.4.4.0/24 from R4
- 5.5.5.0/24 from R5
- 6.6.6.0/24 from R6
- 7.7.7.0/24 from FRS
- 10.10.10.0/24 from CAT1
- 20.20.20.0/24 from CAT1

22.17.4 Configure R1, R3, R4 as redundant RP's in the network. Use 111.1.1.1/32 as the RP address and use a static configuration on all routers and switches that need to access this RP address.
22.17.5 Each router and switch should receive a reply from every router and switch when the address 225.22.22.22 is pinged.

Appendix A.

Issues to Consider When Performing DOiT Labs (Some Helpful Hints)

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 1

1.1 Frame-Relay

In order to assure that you use only the PVC's displayed in the diagram, you may need to explicitly disable a Frame-Relay feature that is enabled by default. The feature that you may want to disable simplifies a Frame-Relay configuration by providing automatic Layer 3 to Layer 2 mappings on Frame-Relay physical interfaces as well as Frame-Relay multipoint subinterfaces. Remember the Frame-Relay topology you are supplied with is a hub and spoke topology. You are explicitly told to use only the PVC's displayed in the exam diagram.

1.2 Catalyst 3550

When choosing a VTP mode, make sure you select one that allows you to advertise ALL VLAN's mentioned in this exercise. Read the entire exercise carefully. List all of the VLAN's mentioned. Pay particular attention to VLAN's with high numbers. Some VTP modes may be limited in how they can propagate certain VLAN numbers.

1.3 ATM

Know all of the different ATM classes of service such as CBR, VBR-rt, UBR etc. Only one of these service classes uses a feedback mechanism and it isn't CBR or UBR.

1.4 OSPF

Know all of the different OSPF network types. Of the five configurable OSPF network types, only two of them perform a DR/BDR election. The two network types that elect a DR/BDR have nothing to do with any variation of point-to-point or point-to-multipoint designs.

If you are going to add a network to OSPF without assigning it to an area, would that make the network an external OSPF entry (a Type-5 LSA)? Consider this thought.

1.5 RIP

When configuring RIP in an exam, ask yourself whether the configuration specifically states RIP version 1. If it does, be on the look out for the VLSM/FLSM issue. Remember RIP v.1 does not advertise mask length information with its updates.

When configuring any version of RIP in an exam, ask yourself whether the configuration specifies two or more redistribution points. If it does, be on the look out for routing loops or suboptimal path selection. Remember that RIP does not make any administrative distance distinction between internal and external RIP routes. This can cause problems on the router performing the redistribution, especially since RIP has the highest administrative distance of all IGP's. You might find that the redistributing routers are selecting a suboptimal path for native RIP routes. Check for this behavior carefully on the redistributing routers.

1.7 ISIS

By default, ISIS routers are configured as Level-1-Level-2 routers. This can be changed under within the "router isis" configuration mode. You may need to change this to fulfill all configuration requirements.

1.8 BGP

When synchronization is enabled for BGP, all routers in a given AS do not need to be configured with BGP. It is obvious that you must configure all routers that are acting as both IBGP and EBGP speakers. However, you may overlook the fact that routers within an AS that possess nothing but IBGP neighbor relationships may not need BGP configured on them at all. Identify all routers within AS 100 that would possess nothing but IBGP neighbor relationships. Determine which of these must necessarily be configured with BGP. This will help you determine what the minimal number of BGP speakers should be.

1.9 DLSw+

You can make one dlsw peer more preferred over another with the backup peer command. If the backup peer command is not available, consider using the dlsw cost command. Assigning a lower cost to a peer makes that peer more desirable. The default dlsw cost is 3.

If you are told to use the minimal number of DLSw+ peer commands, your configuration will probably involve a "promiscuous" configuration on some dlsw local-peer command.

1.11 Address Administration

Make note that in this DHCP configuration requirement, no gateway information is supplied for the specified DHCP pool. Perhaps you need to read ahead in the exam to see if this missing information is supplied somewhere else. In particular, look for any other configuration tasks involving VLAN 30.

1.13 Voice

For any given voice configuration task, if you are given a complete telephone number and you are instructed to dial that number using a subset of the number, consider using the num-exp command or a translation-rule. Translation-rule provide you with more flexibility. They allow you to apply a form of regular-expressions.

1.14 QOS

When performing this custom queuing configuration task, it is recommended that you take into account the differing packet sizes for each protocol assigned to a given queue before you adjust an individual queue's byte-count.

1.15 Catalyst Specialties

Dynamic VLAN's allow users that connect to a specific port to dynamically join a VLAN using VMPS. A Catalyst 3550 cannot be configured as a VMPS server. It can be a VMPS client. In order to configure a Catalyst 3550 as a VMPS client, you need the IP address of the VMPS server.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 2

2.2 Catalyst Configuration

When considering how to prevent any of the fa0/13 and fa0/14 ports from attaining a Spanning Tree blocking state, remember that these two ports interconnect CAT1 to CAT2. Therefore, they form a loop between CAT1 and CAT2. Think about how to make these two ports appear to both Catalysts as one single port. Think of techniques to group individual Catalyst ports so that they act as one aggregated port.

2.4 OSPF

There seems to be no OSPF Area 0 in this exercise. Can you configure OSPF without an Area 0? Yes, provided that there are no other areas.

OSPF offers two authentication methods: one based on passwords passed in clear-text and the second based on the MD5 hashing algorithm.

If you attempt to inject a prefix into an OSPF without associating the prefix to an OSPF area, you need to make that prefix an external OSPF LSA.

2.5 RIP

A good general practice to perform with RIP is to configure passive-interface default under the RIP routing process. Then apply the no passive-interface command to specify the interfaces you want RIP updates to go out on. This will assure that RIP updates get exchanged on the interfaces you desire.

2.6 EIGRP

The EIGRP domain involves three different VLAN's (VLAN's 10, 20 and 30) are configured with the same IP address (172.16.10.0/24). What options are available to make three separate VLAN's appear as a single subnet? Perhaps you should consider some type of bridging technique.

When selecting the most optimal mask for a set of addresses to summarize, pay very close attention to the bit boundaries the supplied prefixes traverse.

2.7 ISIS

The default ISIS router level is Level-1/Level-2. Remember that ISIS Level-1 routers only communicate with other Level-1 routers in their own area. Level-2 routers will communicate with any other Level-2 router regardless of what area it is in.

2.8 BGP

When you configure BGP to advertise a specific set of prefixes only when some other specified prefix in no longer received, consider configuring a BGP conditional advertisement.

2.9 DLSw+

DLSw+ has a backup peer option on the dlsw remote-peer configuration command. It also has an anti-flapping parameter called "linger" for the above mentioned backup option.

2.10 ISDN

Snapshot routing is designed to prevent the periodic updates generated by the traditional Distance-Vector routing protocols such as RIP from keeping up an ISDN link indefinitely.

2.11 Router Maintenance

A Cisco router can be configured as a RARP server to service workstation requests for an IP address.

2.12 Security

If you cannot use the "established" parameter in an extended access-list to fulfill the stated requirement, consider using a reflexive access-list. The Finger service can be activated on a Cisco router. You can test it by telneting to the Finger TCP port.

2.14 QOS

Read carefully!!! Do not just read this section carefully. Read all sections carefully. In particular, read Task 2.15.5 in Section 2.15 very carefully.

2.15 Catalyst Specialty Commands

Do some research on the "switchport voice vlan" Catalyst 3550 interface configuration command. Remember, to enable mls qos on the 3550 whenever you configure a QOS option.

2.16 Gateway Configuration

Did you know that the last two digits of the MAC address used by HSRP are determined by the standby group number you configure. In order to understand this statement, you must be able to convert a decimal number (the standby group number) to a hex number a MAC address.

2.17 Multicasting

Whenever you see a reference to a "shared root", you know it is a PIM Sparse Mode problem. With a Sparse Mode problem, you will need to configure a Rendezvous Point.
There are three ways to configure Rendezvous Point: one static and two dynamic. You must determine what is the one method of "statically" configuring a Rendezvous Point.

As with any multicast scenario, be on the lookout for any RPF lookup problems.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 3

3.1 Frame-Relay

The selection of Frame-Relay interfaces on router R1 is driven by the OSPF configuration requirements in Section 3.1. In order to determine precisely what Frame-Relay interfaces to select for router R1, carefully read Task 3.4.2: "For traffic forwarded to router R1, routers R2 and R4 must possess unique next-hop addresses". Also, carefully examine the scenario diagram. When you examine the Scenario diagram, make note of the fact that routers R2 and R4 have IP addresses assigned to their Serial interfaces but R1 does not. The IP addresses assigned to the serial interfaces of R2 and R4 are from the same subnet 172.16.124.0/24. Therefore, whatever interface or interfaces are assigned to the serial interface of router R1 must be assigned to the 172.16.124.0/24 subnet. For example, when R2 forwards a packet to router R1, it will send the packet to the next-hop address of 172.16.124.1 – an IP address assigned to router R1. When R4 forwards a packet to router R1, it will send the packet to some 172.16.124.0/24 address other than 172.16.124.1 since that is the next-hop address used by router R2. Remember, the next-hop addresses must be unique. An option to consider when fulfilling this requirement is to configure two point-to-point subinterfaces on router R1 and assign them both to the 172.16.124.0/24 subnet. You can do this with point-to-point subinterfaces, you can't do this with physical interfaces. See the OSPF section to complete the configuration requirements for providing IP connectivity over Frame-Relay. In order to fulfill this requirement, you must carefully select the correct OSPF network type.

3.2 Catalyst Configuration

By default, all VLAN's are passed over a trunk port. Review whatever configuration commands are available to limit the range of VLAN's that are passed over a trunk port.

3.3 ATM

Classical IP is designed to dynamically map IP addresses to NSAP addresses.

3.4 OSPF

As mentioned in the Frame-Relay section, the OSPF and Frame-Relay configuration for this scenario are tightly intertwined. In order to assure reachability of the 172.16.124.0/24 subnet from all other routers in this scenario, it cannot be advertised as a single /24 subnet. Instead, advertise all devices that are connected to it as individual /32 host routes. This can be performed by selecting the correct OSPF network type.

3.7 ISIS

You can include a loopback interface within an ISIS process without using the "ip router isis" command. Just configure the interface as passive under the router isis configuration mode. This configuration technique may also result in surprising metric assignments to isis configured interfaces.

3.8 BGP

A method to consider in fulfilling the synchronization requirement of task 3.8.5 is to include the 10.10.10.0/24 prefix in the RIP process running on CAT2. When it gets advertised to another router, perhaps you can use the "ip summary-address rip" command to generate a /24 entry on the routers that need to fulfill the synchronization requirement.

3.9 DLSw+

Select you IP addresses carefully for the local-peer statements. By selecting the correct IP addresses, you can avoid the DLSw+ traffic from going over the Frame-Relay connection.

3.10 ISDN

Read the Goals and Restrictions at the beginning of this Scenario very carefully. It may direct your choice of how to configure ISDN.

3.11 Router Maintenance

Remember that a router an be configured as an HTTP server.

3.12 Security

Whenever a specific username is referenced to a task that is related to permitting or denying a specific network resource, consider configuring a dynamic access-list.

3.13 Voice

An option to consider when reducing packet overhead for voice traffic on slow speed WAN links is "ip rtp header-compression".

3.14 QOS

An option to consider as the most simplified method of configuring QOS is an "auto qos" configurations available on the Catalyst 3550. There is an interface configuration command called "auto qos voip trust" that you might consider as an option.

3.16 Gateway Redundancy

This section makes references to distributing traffic as well as redundancy for workstations connected to VLAN 10. This can be accomplished by configuring multiple DHCP servers as well as MHSRP on VLAN 10.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 4

4.2 Catalyst Configuration

In order to configure, two different subnets on VLAN 10, you might consider configuring 802.1Q Tunneling.

4.3 ATM

The ATM configuration involves an ATM PVC. Since you cannot use any static or dynamic mappings, consider configuring the ATM interface as a point-to-point subinterface. Point-to-point subinterfaces require no static or dynamic inverse arp entries.

4.4 OSPF

An option to consider to when excluding external routing information from Area 60 is a stub area configuration. However, remember that Area 60 also possesses an ASBR. Therefore, consider making Area 60 a Not So Stubby Area.

4.5 RIP

When attempting to fulfill the RIP configuration requirements, carefully review the Goals and Restrictions Section at the beginning of the Scenario. A number of techniques are available to solve the VLSM/FLSM issue including "default information-originate", "rip version 2" and "ip default-network". Carefully review which one of these is excluded from this Scenario. Also, remember that you can configure a specific version of RIP at the interface level.

4.7 ISIS

Remember that ISIS has no equivalent to the OSPF interface command "ip ospf network". Therefore, you cannot change ISIS network types like you can with OSPF. ISIS has only two network types point-to-point and LAN. The default ISIS network type for a point-to-point NBMA subinterface is "point-to-point". The default ISIS network type for a physical NBMA interface or multipoint subinterface is "LAN". Two NBMA interfaces configured on the same link that have different ISIS network types will not form an adjacency.

4.8 BGP

Remember that you use local-preference to influence "outbound" routing decisions. Also, remember that you can use the AS prepend command within a route-map configuration to increase the AS-Path length for a given prefix.

4.9 DLSw+

The minimum overhead for a DLSw+ peer relationship would involve a DLSw direct encapsulation. If only one DLSw+ peer is possesses the "conf" peer type, the other DLSw+ peeer must be configured as promiscuous.

4.11 Network Monitoring

To perform the type of network monitoring described in this section, consider configuring Cisco's Service Assurance Agent (SAA).

4.12 Security

If you cannot filter packets using an IP access-list, consider using the following interface configuration command: " ip verify unicast reverse-path". This allows you to filter packets based upon an examination of the source address of packets received on a specific interface.

4.13 Voice

In order to fulfill this requirement, consider using the "connection trunk" command.

4.14 QOS

When you see in a QOS scenario, the language of dividing traffic into categories of "critical", "medium" and "normal", consider configuring priority queuing.

4.17 Multicasting

By default, RTP header compression only supports 32 compression connections. You can adjust this number if desired.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 5

5.1 Serial interfaces

You are instructed to configure logical interfaces on all Frame-Relay interfaces. Read ahead and determine what IGP will be running over these Frame-Relay interfaces. If it is ISIS, you must configure a full mesh of PVC's for ISIS. If there are three routers on one of the Frame-Relay cloud used by ISIS, you will need to configure a logical interface that can have more than one DLCI assigned to it.

5.2 Catalyst Configuration

On the Catalyst 3550, Cisco has implemented configuration tools that allow you to determine whether traffic is being transmitted and received on a given port. A technology related to this is UniDirectional Link Detection (UDLD).

5.3 ATM

Out of all the encapsulation types available for ATM PVC's, aal5mux is the most limited. It may also be the one with the least overhead.

5.4 OSPF

If you cannot elect a DR on a segment, you cannot configure that segment with the non-broadcast or broadcast OSPF network types.

5.5 RIP

When attempting to fulfill the RIP configuration requirements, carefully review the Goals and Restrictions Section at the beginning of the Scenario. A number of techniques are available to solve the VLSM/FLSM issue including "default information-originate", "rip version 2" and "ip default-network". Carefully review which one of these is excluded from this Scenario.

5.8 BGP

Remember that BGP will select one and only one best path for a specific prefix. If you get a task requirement that instructs you to select more than one best path, consider using the "maximum-paths" command.

5.9 DLSw+

If there are more than two DLSw+ peers and you are instructed to allow all peers to communicate directly with each other but you must use a limited number of remote-peer commands, consider configuring dlsw border peers.

5.10 ISDN

When you are told to use the ISDN link only when another link is not available and at the same time, you are told not to configure a routing protocol over the link, consider using a floating static route or policy routing. Before you do, check the Scenario wide Goals and Restrictions section at the beginning of the test.

5.11 VPN

Remember that a VPN does not always involve IPSEC. The most basic VPN could be a GRE tunnel.

5.13 Voice

To assure that traffic that a router sends to specific destination, consider a static route or using some form of policy routing. When considering policy routing, account for the differences between standard policy routing and local policy routing. Also, remember that policy routing configurations involve route maps. You can set packet precedence bits with a route-map. Before configuring either of these options, check the Scenario wide Goals and Restrictions section at the beginning of the test.

5.14 QOS

If QOS parameters are supposed to be dynamically reserved along a path of a given flow of traffic, then consider configuring RSVP.

5.16 Address Administration

When a task asks you to make a packet appear "as if" it was originated from a source other than its actual source, consider configuring NAT.

5.17 Multicast

When two PIM multicast routers are able to forward the same multicast traffic from the same source onto the same subnet, only one will perform the forwarding. The multicast router that will perform the forwarding will be selected in a PIM Assert election. This type of election may be important for you to successfully fulfill the configuration requirements.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 6

6.1 Serial Interfaces

When you select what Frame-Relay interfaces will be point-to-point and what interfaces will be multipoint, read ahead to the IGP sections. The IGP sections may influence your interface selections. If you cannot supply an explicit IP address to a point-to-point Serial interface, consider configuring "ip unnumbered".

6.3 ATM

For any ATM configuration task that requires you to set a cell rate of any type, review the different types of ATM quality of service classes such as CBR, VBR, ABR, etc.

6.4 OSPF

Remember that you cannot configure a virtual-link across a NSSA area.

6.5 RIP

You are configuring RIP on a hub and spoke NBMA topology where R3 is the hub. Split-horizon is enabled by default on R3's multipoint subinterface. Task 6.5.3 indicates that you cannot disable split-horizon. If you can't disable split-horizon, remember that RIP updates have a TTL = 2. Consider converting the default RIP broadcasts to unicast updates. See what the effect of this change will be. Remember when you use RIP unicast updates you must configure the related interfaces as passive.

6.6 EIGRP

You are configuring EIGRP on a hub and spoke NBMA topology where R1 is the hub. Split-horizon is enabled by default on R1's multipoint subinterface. Task 6.6.3 indicates that you cannot disable split-horizon. If you can't disable split-horizon, remember that EIGRP updates have a TTL = 2. Consider converting the default EIGRP multicasts to unicast updates. See what the effect of this change will be. Remember when you use EIGRP unicast updates you do not need to configure the related interfaces as passive.

6.7 ISIS

For ISIS to operate properly, make sure that interface MTU's are the same. Read the entire Scenario and determine whether any MTU sizes were changed.

6.8 BGP

A use of the term "IGP" when used within the context of a BGP configuration might reference a BGP origin code. When you are injecting the 7.7.1.0/24 prefix into BGP on R5, remember that it is a specific subset of an OSPF route. Consider redistributing this prefix into BGP while at the same time matching on the prefix's unique characteristics.

6.10 ISDN

Before configuring ISDN, carefully read the IGP section of this Scenario and determine if there is an IGP running on the ISDN link. If the IGP is OSPF, consider configuring ospf demand-circuit; if it is RIP consider configuring snapshot routing; if it is EIGRP, consider configuring dialer watch.

6.11 Router Maintenance

When you are told to disable a specific router feature, review what is enabled and disabled by the global configuration "service" commands.

6.14 QOS

A possible method of fulfilling the requirements of this section might involve RSVP.

6.15 Catalyst Specialties

The Catalyst 3550 has an interface command that allows you to limit broadcasts and multicasts. It is designed to control heavy "storms" of broadcast and multicast traffic.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 7

7.3 ATM

In Task 7.3.4, it states to configure R6 to resolve IP addresses to a data-link identifier. Other ATM configuration tasks in this section are referencing NSAP addresses. Consider configuring Classical IP on the ATM link with R6 as the CLIP server. If you make R6 the CLIP server, it will resolve an IP address to an NSAP address and then an NSAP address to the ATM data-link identifier: the VPI/VCI pair.

7.7 ISIS

While OSPF has Area 0 as its backbone area, there is no explicit backbone area in ISIS. The equivalent to OSPF Area 0 in ISIS is a contiguous collection of ISIS Level 2 routers.

7.8 BGP

Closely examine the IP addresses assigned to the interfaces on VLAN 10. Two IP addresses are assigned to the FastEthernet interfaces of R1 and R3 on VLAN 10. The address prefixes are flip-flopped on each of the two routers. Since the IP addresses are not the same, this will create a complication when you attempt to form an EBGP peer on these interfaces. You must treat this EBGP neighbor relationship as if it was NOT on a shared subnet. The IOS has a specific BGP neighbor configuration command designed to allow EBGP peers to form neighbor relationships using addresses that do not share a common subnet between the two EBGP speakers.

If you are concerned that a set of any prefixes advertised by any routing protocol may destabilize a link, consider summarizing the set of prefixes into a single summary. Then, if one of the prefixes within the set is flapping, the summary will conceal the flapping condition. The same logic applies to BGP updates. You can conceal destabilizing link conditions cause by BGP route flaps by configuring the BGP aggregate command.

7.10 ISDN

When logical interfaces are associated with an ISDN interface, consider configuring dialer profiles. Dialer profile configurations involve the creation of logical dialer interfaces.

7.12 Security

If you can't use a standard or extended "numbered" access-list to protect a router from a specific set of IP traffic such as ICMP redirects, broadcast packets, etc; consider configuring a "named" access-list.

7.13 Voice

When determining the sampling rate applied to a given voice call, you need to identify the CODEC configured under the relevant dial-peer. Each CODEC has its own sampling rate. When you manually select a CODEC, the IOS will list an associated bandwidth measurement with each CODEC. On the same line, you can specify how many bytes a single voice sample can produce.

7.15 Catalyst Specialties

A technique to restrict the amount of flooded traffic on trunk links is to limit the number of VLAN's traversing the trunk links. If you want to limit the number of VLAN's traversing a trunk link, configure the "switchport trunk allowed vlan" port configuration command.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 8

8.1 Serial Interfaces

If you want to configure a Frame-Relay interface so that the link does not possess an explicit IP subnet but yet it still forwards IP traffic, consider configuring point-to-point subinterfaces with ip unnumbered.

If you configure a back to back Frame-Relay configuration as is done between routers R3 and R5 and you are instructed to not enable Frame-Relay switching, consider disabling the interface keepalive feature.

If you are instructed to configure CHAP over a Frame-Relay connection, you will need to configure PPP over that same connection and you will also need to create a virtual-template.

8.3 ATM

If you are instructed to configure CHAP over an ATM connection, you will need to configure PPP over that same connection and you will also need to create a virtual-template.

8.5 OSPF

If packets destined to network 10.10.10.0/24 are not being forwarded to router R5 from R2, you may have an issue with the redistribution between OSPF and ISIS. This issue may involve the forwarding address setting for the 10.10.10.0/24 External OSPF prefix.

8.8 BGP

The challenge of this section is centered on how you configure IBGP within AS 100. The last requirement instructs you to use a minimum number of IBGP peers. This means that you should not form a full mesh within AS 100. Therefore, you will configure either a router-reflector or a set of route-reflectors or a confederation. Furthermore, you are instructed to prevent router R2 from accepting any NLRI transiting through R5 and vice versa. You are told to fulfill this restriction without using any filtering techniques. If you can't use a filtering technique, consider configuring a route-reflector with the same cluster-id on routers R2 and R5. This may prevent R2 from accepting routes from R5 and vice versa.

8.9 VPN

If you are told to use "IP into IP" encapsulation for a VPN scenario, you are going to configure an "IP in IP" tunnel between two routers.

8.12 Security

If you want to apply security techniques that are based on application layer inspection of traffic, use the "ip inspect" interface command. This involves the configuration of CBAC.

8.15 Catalyst Specialties

If you want to prevent traffic from flowing between two workstations on the same switch and you cannot use an access-list, consider configuring the following Catalyst 3550 port configuration command: "switchport protected".

8.16 Address Administration

When you are attempting to fulfill this task, keep in mind that you are supposed to forward the broadcasts for the specific list of application traffic. You are not supposed to convert the broadcasts into unicast packets.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 9

9.4 OSPF

When you select routers to perform flooding of OSPF LSA's, you are probably identifying the routers that will be elected a Designated Router and Backup Designated Router.

9.8 BGP

A "bit bucket" for all destinations not matching more specific prefixes would be a 0.0.0.0/0 entry to Null0.

A method of forwarding all unknown traffic to a specific router is to have that router advertise a 0.0.0.0/0 address to down stream routers. BGP can accomplish this with the default-originate neighbor command.

If you want BGP to advertise one prefix as long as it possesses some other prefix, consider configuring a conditional advertisement.

9.10 ISDN

When you are told not to configure an explicit IP address on any point to point link such as an ISDN link, consider using ip unnumbered.

9.14 QOS

One technique of guaranteeing bandwidth with minimal delays is to configure RSVP.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 10

10.4 Address Administration

When the same IP address spans over multiple links, consider configuring Transparent Bridging. Whenever configure Transparent Bridging, pay special attention to how the spanning tree gets formed to prevent bridging loops. When examining how the spanning tree gets formed, determine which device becomes the root bridge, what ports become the designated ports, what ports become the root ports and what ports become the blocked ports. You will have a port go into a blocking state only if there is a loop in the bridge topology. Therefore, for this scenario, carefully mark which interfaces belong to which IP subnet. This will help you determine which interfaces are in a specific transparent bridge group, which will then help you determine which interfaces are participating in a specific spanning tree. Once you have identified the interfaces in this manner, determine which device is the root of a given spanning tree, then construct the path of the spanning tree. This should lead you to which ports are in a blocking state.

When determining how to configure transparent bridging in this exercise, pay special attention to the IP addresses superimposed on the following routers: R1, R3, R4 and R6.

10.5 OSPF

When you want to reduce the refreshing and flooding of OSPF LSA's, review the available OSPF interface configuration commands.

10.9 BGP

When a full mesh is not allowed among IBGP peers, consider configuring either a route-reflector or confederation.

When you need to influence incoming traffic for a given autonomous system, consider configuring a MED or AS PATH prepend. The MED would be sent from the AS that will be receiving the traffic (in this case that would be AS 2). The AS Path prepend would be performed by the AS advertising the relevant prefixes.

10.11 ISDN

The "?" character is a special character when entered with the IOS command line interface.

10.13 Security

When you want to limit what devices can communicate with each other on a given VLAN, consider configuring Catalyst 3550 access-lists or vlan maps.

10.16 QOS

Remember that Catalyst 3550 ports are configured to support Weighted Round Robin queuing.

10.18 Multicast

How do you reduce multicast flood on the older Catalyst switches. Cisco developed a proprietary technique for solving this problem of multicast flooding on a given VLAN.

******ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 11**

11.3 ATM

Know all of the different ATM classes of service such as CBR, VBR-rt, UBR etc. For any ATM configuration task that requires you to set a cell rate of any type, review the different types of ATM quality of service classes such as CBR, VBR, ABR, etc.

11.4 Address Administration

When you encounter a configuration that requires you to assign IP addresses on a per port basis on a Catalyst 3550 switch consider configuring dhcp with a specific option related to this very task. Consult UNIVERCD and search for configuring DHCP on the Catalyst 3550. Closely review any references to specific DHCP options.

11.5 OSPF

In this scenario, you are presented with a full mesh of Frame-Relay PVC's. You are asked to configure OSPF over the Frame-Relay full-mesh topology. Then, you are presented with the following challenge: If one or two (but no more than two) PVC's fail, retain all adjacencies between the OSPF routers connected to the Frame-Relay network. In order for OSPF routers to be adjacent on a Frame-Relay topology, they must share a direct PVC connection. All OSPF packets possess a TTL of 1. Therefore, they cannot be forwarded through a transit "hub" router on a hub and spoke Frame-Relay topology. If one or two Frame-Relay PVC's fail in this scenario, a hub and spoke Frame-Relay topology will be created. What options do you have to overcome this apparent dilemma of retaining OSPF adjacencies among all routers connected to a full-mesh Frame-Relay network when one or two of the PVC's fail? An option to consider is to not form the OSPF adjacencies directly on the Frame-Relay interfaces themselves. Consider using a set of tunnel interfaces over the Frame-Relay topology and associate the IP address to the tunnel interfaces with ip unnumbered.

11.9 BGP

Load sharing is defined as preferring one half of the IP address space on one path and the other half of IP address space on a second path.

11.16 Catalyst Specialties

Load balancing options are available with EtherChannel.

******ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 12**

12.5 OSPF

When performing this OSPF task, carefully read and review all other IGP configuration requirements. Analyze how other IGP routing protocols interact with OSPF. Examine the characteristics of all OSPF External routes. In particular, examine the forwarding address attribute of OSPF External routes.

12.6 RIP

When a RIP update attains a metric of 15 hops, it is unreachable. You can manually influence the RIP metric with the offset-list command under the RIP routing process.

12.7 EIGRP

When you want to load balance traffic with EIGRP to a specified ratio, consider using the "variance" command under the eigrp routing process.

12.8 ISIS

When performing this ISIS task, carefully read and review all other IGP configuration requirements. Analyze how other IGP routing protocols interact with OSPF.

12.11 ISDN

If an ISDN backup configuration involves placing the ISDN interface in a "standby" mode, then you are configuring the backup interface command. If you are supposed to pass traffic over an ISDN interface while it is in "standby", this cannot be done with a legacy DDR configuration. Consider using an alternative method of DDR configuration.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 13

13.3 ATM

If you are given incomplete configuration information for ATM, consider using ATM configuration discovery tools such as ATM ILMI PVC discovery.

13.4 OSPF

Only two OSPF network types elect a DR/BDR: the OSPF network types broadcast and non-broadcast.

Remember that an OSPF virtual-link extends Area 0.

13.5 RIP

RIP versions can be specified on a per interface basis.

RIP Version 1 does not pass prefix mask length information. RIP version 2 passes prefix mask length information.

13.6 EIGRP

EIGRP is able to mark a prefix as a candidate default with the "ip default-network" command.

13.7 BGP

You need to form EBGP neighbor relationships on a common shared subnet shared by both EBGP speakers. If you do not form EBGP neighbor relationships on a common shared subnet, you must apply the ebgp multihop configuration command. When you configure EBGP neighbor relationships using secondary addresses, the address is not considered to be on a common shared subnet shared by two or more EBGP speakers.

EBGP speakers between private AS's in a confederation must also share a common subnet. If they do not, you must configure ebgp multihop.

13.11 Router Maintenance

Router crash information is contained within a "core dump". The IOS possesses commands to direct core dumps to files on tftp or ftp servers.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 14

14.2 Catalyst Configuration

What options are available to make multiple separate VLAN's appear as a single subnet? Perhaps you should consider some type of bridging technique.

14.4 – 14.5 OSPF and RIP

Pay special attention to the mutual redistribution performed between OSPF and RIP. Make sure that RIP native routes are reachable RIP and not reachable via OSPF External routes on routers R2 and R4. Make note that there are two redistribution points between OSPF and RIP. Closely examine the administrative distance of the RIP and OSPF routing table entries.

14.8 BGP

When configuring BGP, pay attention to what IGP's are or ARE NOT running over the links where BGP neighbor relationships are formed.

14.9 DLSw+

When configuring DLSw+ MAC address filters, remember that DLSw+ processes all MAC addresses in non-canonical address format and Ethernet uses a canonical address format.

14.13 Voice

Under the dial-peer configuration mode, the preference command is used to make on dial-peer more preferred over another.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 15

15.2 Catalyst Configuration

The 802.1Q trunking protocol implements the concept of a native VLAN.

15.8 BGP

You can configure multiple route-reflectors within the same AS. Multiple route-reflectors can be configured in a manner that provides BGP update redundancy.

15.9 DLSw+

If there are more than two DLSw+ peers and you are instructed to allow all peers to communicate directly with each other but you must use a limited number of remote-peer commands, consider configuring dlsw border peers.

When configuring any DLSw+ command that references MAC addresses, remember that DLSw+ processes all MAC addresses in non-canonical address format and Ethernet uses a canonical address format.

15.10 ISDN

PPP PAP passes a password between ISDN peers.

15.11 Address Administration

If you need to access a specific private address from an outside untrusted side, consider using a ip nat inside source static configuration.

15.12 Security

Whenever a specific username is referenced to a task that is related to permitting or denying a specific network resource, consider configuring a dynamic access-list.

15.13 Voice

VOIP packets are transported with a 20 byte IP header, 12 byte UDP header and a 12 byte RTP header. The default VOIP CODEC is G729a. A single G729a voice sample generates a 20 byte voice payload. Therefore, a default VOIP packet using a G729a CODEC possesses 40 bytes of overhead and 20 bytes of voice payload. This results in a 2:1 ratio of packet overhead to actual voice payload. Is there some way to make this ratio more favorable?

15.14 QOS

Traffic shaping buffers traffic that is out of profile. Traffic policing drops or marks traffic that is out of profile. You can deploy traffic policing with CAR or the Modular QOS CLI.

15.16 Multicasting

Only PIM Sparse Mode constructs a shared tree. The root of the shared tree is the Rendezvous Point.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 16

16.1 Frame-Relay

Inverse ARP is activated by default on Frame-Relay PVC's Inverse ARP may cause you to use more PVC's than you need to use. Remember you are supposed to use a minimum number of DLCI's in this Scenario.

16.4 OSPF

Remember that an OSPF virtual-link extends Area 0.

16.6 EIGRP

The recommended method for controlling what interfaces are participating in EIGRP is to configure the network command with a mask.

16.7 ISIS

By default, ISIS routers are designated as Level-1/Level-2 routers. Level-1/Level-2 routers generate both Level-1 and Level-2 hellos out every interface. If only one of type ISIS router is on the same interface, generating the unnecessary packets is undesirable.

16.8 BGP

EBGP speakers between private AS's in a confederation must also share a common subnet. If they do not, you must configure ebgp multihop.

16.9 DLSw+

DLSw+ has a backup peer option on the dlsw remote-peer configuration command. It also has an anti-flapping parameter called "linger" for the above mentioned backup option.

Whenever you see a reference to "advertising" or "caching" a NETBIOS host name, MAC address or LSAP number, consider deploying DLSw+ capabilities exchange.

16.14 QOS

Traffic shaping buffers traffic that is out of profile. Traffic policing drops or marks traffic that is out of profile. You can deploy traffic policing with CAR or the Modular QOS CLI.

16.16 Multicast

PIM Dense Mode uses the flood and prune protocol.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 17

17.1 Serial Interfaces

Before you begin configuring your Frame-Relay interfaces, carefully read the IGP configuration tasks in this Scenario.

Remember that you can configure the Frame-Relay physical interface even when you have a subinterface configured under the interface.

17.4 OSPF

You cannot place a secondary IP address in a different OSPF area than that assigned to the primary IP address.

17.5 RIP

Pay special attention to the mutual redistribution performed between OSPF and RIP. Make sure that RIP native routes are reachable RIP and not reachable via OSPF External routes on routers R2 and R4. Make note that there are two redistribution points between OSPF and RIP. Closely examine the administrative distance of the RIP and OSPF routing table entries.

17.8 BGP

If you can't match on prefix and or AS-Path at AS 10, perhaps you can match by prefix or AS-Path in another AS. When you match by prefix or AS-Path in another AS, perhaps you can tag the desired prefixes with a tag such as a BGP community.

17.9 DLSw+

When you must make a forwarding decision based upon criteria other than a destination IP address, consider configuring policy routing.

Whenever you see a reference to "advertising" or "caching" a NETBIOS host name, MAC address or LSAP number, consider deploying DLSw+ capabilities exchange.

When configuring DLSw+, carefully read the entire command. Pay particular attention to any process that might change IP addresses between DLSw+ peers.

17.10 ISDN

Snapshot routing is designed to prevent the periodic updates generated by the traditional Distance-Vector routing protocols such as RIP from keeping up an ISDN link indefinitely.

17.11 Address Administration

HSRP makes two or more routers appear as one router to end systems on a subnet. IRDP allows an end system to select a preferred router as a default gateway.

17.14 QOS

Traffic shaping buffers traffic that is out of profile. Traffic policing drops or marks traffic that is out of profile. You can deploy traffic policing with CAR or the Modular QOS CLI.

17.17 Multicast

224.0.1.39 and 224.0.1.40 are reserved multicast addresses used by PIM Auto-RP.

******ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 18**

18.1 Frame-Relay

Interpret "dynamic mappings" as inverse arp. Therefore, you must use inverse arp to fulfill the task of this section.

However, you must use only the PVC's in the supplied diagram. Therefore, you cannot have inverse arp occur over all Frame-Relay PVC's. You must control inverse arp without disabling inverse arp. An option to consider is to isolate unneeded inverse arp entries on interfaces that have no IP addresses.

18.2 Catalyst Configuration

What options are available to make multiple separate VLAN's appear as a single subnet? Perhaps you should consider some type of bridging technique.

18.4 OSPF

Since you could only use inverse arp on the Frame-Relay connections and since you have a hub and spoke NBMA topology on the 151.10.124.0/24 subnet, router R2 cannot ping the Frame-Relay interface of router R4. If you can configure an OSPF network type that generates a host route for each interface, this will solve the reachability problem between R2 and R4.

If you can't redistribute, inject a prefix into OSPF with a network statement, or redistribute connected, consider redistributing the prefix into OSPF from another routing protocol.

OSPF paths through Area 0 are always preferred over non-Area 0 paths. Virtual-Links extend Area 0 through non-Area 0 areas.

18.7 ISIS

You can inject a prefix into ISIS with the "ip router isis" command or via redistribution. The passive-interface command can also allow you to inject a prefix into ISIS.

18.8 BGP

When synchronization is enabled and OSPF is the IGP configured in the Autonomous System that will fulfill the synchronization requirement, make sure that the OSPF ASBR RID matches the advertising IBGP speaker RID for all IBGP advertised updates.

18.9 DLSw+

The DLSw+ peer relationship that uses minimal overhead does not use IP at all. Determine whether there is a direct PVC between the routers R1 and R3. You can use this PVC even if it has not IP address configured on it.

You can make one peer more preferred over another peer by manipulating the DLSw+ cost parameter. You can see the configured as well as the default DLSw+ cost command with the "show dlsw capabilities" command.

18.10 ISDN

When configuring ISDN, carefully reread the Goals and Restrictions at the beginning of this scenario. After rereading the Goals and Restrictions, consider what options you have available when you use ISDN to provide connectivity when a particular prefix is removed from the routing table and you must replace it without using an IGP.

Consider a range of options to make sure that R1 allows calls only from R2's ISDN number. One option requires that the ISDN switch announces R2's telephone number to R1. Another option – DDR Callback – provides similar functionality independently of how the ISDN switch is configured.

18.11 VPN

When you configure RIP on a router, all interfaces on the router that share the classful prefix configured with the network command under the RIP process. Since you are running RIP over the VPN tunnel, make sure you do not allow too many prefixes to be advertised over the tunnel.

18.14 QOS

When you need to allocate a percentage of bandwidth to a particular protocol, consider using custom queuing or class-based weighted fair queuing.

18.16 Address Administration

You can translate both IP addresses as well as TCP ports with NAT.

18.17 Multicast

Only one PIM Sparse mode dynamic Rendezvous Advertisement method requires PIM version 2. It is not Auto-RP.

******ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 19**

19.2 Catalyst Configuration

A port is placed in a spanning tree blocking state based upon the following criteria in the order presented: (1) what port has the lower total path cost to the root bridge (2) what bridge has a lower bridge id (3) what port has a lower port priority.

19.4 OSPF

Both of the OSPF network types point-to-multipoint and point-to-multipoint non-broadcast do not perform DR/BDR elections. However, the point-to-multipoint non-broadcast network type requires neighbor statements and the point-to-multipoint network type uses the 224.0.0.5 multicast.

19.6 EIGRP

EIGRP is a Distance Vector routing protocol. Split-horizon is applied to EIGRP updates. Split-horizon is enabled by default for all EIGRP configured interfaces.

When you create an EIGRP summary address, it inserts the prefix into the local forwarding table referencing a Null0 interface. It then advertises this prefix to other EIGRP routers.

19.7 ISIS

ISIS does not use IP for transport. ATM will not forward the non-IP ISIS traffic by default.

19.8 BGP

The first comparison performed during the BGP path selection process is the BGP administrative weight.

When determining how to configure many of the BGP tasks listed in this Scenario, remember to carefully read tasks in other sections of this scenario. Other tasks in this Scenario will constrain your options for configuring BGP.

Remember that EBGP learned updates have an administrative distance of 20. When you advertise EBGP prefixes that are also used to form the underlying EBGP neighbor relationship between two peers, your BGP peer stability will be adversely affected.

19.9 DLSw+

You can guarantee a DLSw+ peer a specific of amount of end to end bandwidth using RSVP.

19.10 ISDN

When configuring ISDN, carefully reread the Goals and Restrictions at the beginning of this scenario. After rereading the Goals and Restrictions, consider what options you have available when you use ISDN to provide connectivity when a particular prefix is removed from the routing table and you must replace it without using an IGP.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 20

20.2 Catalyst Configuration

In order to configure, two different subnets on VLAN 10, you might consider configuring 802.1Q Tunneling.

20.3 ATM

Know all of the different ATM classes of service such as CBR, VBR-rt, UBR etc. For any ATM configuration task that requires you to set a cell rate of any type, review the different types of ATM quality of service classes such as CBR, VBR, ABR, etc.

20.4 OSPF

If the task directs you to form an adjacency with an OSPF peer that does not possess a physical connection with another peer, consider creating a virtual common subnet.

20.8 BGP

When you form BGP neighbor relationships using loopback interfaces, remember to consider hidden issues related to forming EBGP neighbor relationship with IP addresses that reside on an interface other than one that is on a common subnet shared by both EBGP peers.

20.10 ISDN

When you want to use both B Channels for a single ISDN call you are inverse multiplex multiple B channels. This technique is sometimes referred to as "multilink".

20.12 Security

The IOS possesses a global configuration command designed to prevent Denial of Service attacks due to half open TCP sessions.

20.13 Voice

LLQ requires the configuration of the Modular QOS CLI (MQC).

20.14 QOS

When you need to allocate a percentage of bandwidth to a particular protocol, consider using custom queuing or class-based weighted fair queuing.

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 21

21.1 Serial Interfaces

Read the tasks of this section carefully. In particular, pay attention to the fact that R3 must have one physical interface and then the next task states all OTHER (excluding) R3 are allow to have only one logical point-to-point interface.

Even when you have configured a subinterface under a physical Frame-Relay interface, you can still configure an IP address and assign DLCI's to the physical Frame-Relay interface.

21.4 OSPF

Carefully examine the options associated with the OSPF neighbor command. In particular, search for options related to filtering traffic.

21.5 RIP

In order to perform the RIP configuration, you must carefully read the entire exam. In order to fulfill the last tasks in the RIP section, FRS must advertise to one and only one neighbor but yet involve both routers on VLAN 30. Perhaps FRS will advertise to a virtual IP address shared by CAT1 and CAT2. Again, read the entire exam carefully.

21.8 BGP

Review the well known communities.

21.10 ISDN

What IP protocol number is used by ISIS for control traffic? Does ISIS use IP at all for control traffic? This is an "interesting" question.

21.12 Security

In order to fulfill the tasks in this section, read the remainder of the exam very carefully.

21.17 Multicast

Carefully review the multicast configuration options on the Catalyst 3550. Pay close attention to any configuration options that reference the word "reliable".

****ISSUES TO CONSIDER WHEN PERFORMING SCENARIO 22

22.4 OSPF

Not only is there a point-to-point tunnel interface, there is also a multipoint tunnel interface.

22.7 ISIS

R6 will not view prefixes learned from R1 and R2 as equal since one link is ATM and the other is FastEthernet. An ATM interface has a different metric than a FastEthernet interface.

22.9 Traffic Control

You can determine any type of forwarding path with policy routing.

22.10 ISDN

Snapshot routing is not the only way to supporess RIP updates. Another method is configuring RIP triggered updates.

22.12 SNMP

There are three versions of SNMP. Each version provides a different level of SNMP security.

22.14 QOS

You have four options to mark traffic: MQC, Route-maps, CAR and QPPB

Appendix B.

On line instructions for the POD rentals

The Terminal server naming convention and ip addresses for the on-online access:

Table1. Terminal Server names and IP addresses

Terminal Server Host Name	IP Address
GROUP1	
GROUP1-POD1-TS	66.170.103.11
GROUP1-POD2-TS	66.170.103.12
GROUP1-POD3-TS	66.170.103.13
GROUP1-POD4-TS	66.170.103.14
GROUP1-POD5-TS	66.170.103.15
GROUP1-POD6-TS	66.170.103.16
GROUP1-POD7-TS	66.170.103.17
GROUP2	
GROUP2-POD1-TS	66.170.103.21
GROUP2-POD2-TS	66.170.103.22
GROUP2-POD3-TS	66.170.103.23
GROUP2-POD4-TS	66.170.103.24
GROUP2-POD5-TS	66.170.103.25
GROUP2-POD6-TS	66.170.103.26
GROUP2-POD7-TS	66.170.103.27

For scheduling check www.netmasterclass.net

For the security restrictions check www.netmasterclass.net

For pod access information, such as user name, password, IP address, etc. check www.netmasterclass.net

Procedure for on-line access to NetMasterClass, LLC pod:

1. Telnet to the Terminal Server GROUPN-PODX-TS
2. Provide username and password
3. To open reverse telnet sessions type "telnet hostname", for example:

   telnet R1

Table2. ISDN SPIDs and Calling Numbers

GROUP1		GROUP2	
POD1		POD1	
R1	R2	R1	R2
SPID1:21255520000101 SPID2:21255520010101	SPID1:21255520020101 SPID2:21255520030101	SPID1:21255520000101 SPID2:21255520010101	SPID1:21255520020101 SPID2:21255520030101
POD2		POD2	
R1	R2	R1	R2
SPID1:21255520040101 SPID2:21255520050101	SPID1:21255520060101 SPID2:21255520070101	SPID1:21255520040101 SPID2:21255520050101	SPID1:21255520060101 SPID2:21255520070101
POD3		POD3	
R1	R2	R1	R2
SPID1:21255520080101 SPID2:21255520090101	SPID1:21255520100101 SPID2:21255520110101	SPID1:21255520080101 SPID2:21255520090101	SPID1:21255520100101 SPID2:21255520110101
POD4		POD4	
R1	R2	R1	R2
SPID1:21255520120101 SPID2:21255520130101	SPID1:21255520140101 SPID2:21255520150101	SPID1:21255520120101 SPID2:21255520130101	SPID1:21255520140101 SPID2:21255520150101
POD5		POD5	
R1	R2	R1	R2
SPID1:21255520160101 SPID2:21255520170101	SPID1:21255520180101 SPID2:21255520190101	SPID1:21255520160101 SPID2:21255520170101	SPID1:21255520180101 SPID2:21255520190101
POD6		POD6	
R1	R2	R1	R2
SPID1:21255520200101 SPID2:21255520210101	SPID1:21255520220101 SPID2:21255520230101	SPID1:21255520200101 SPID2:21255520210101	SPID1:21255520220101 SPID2:21255520230101
POD7		POD7	
R1	R2	R1	R2
SPID1:21255520240101 SPID2:21255520250101	SPID1:21255520260101 SPID2:21255520270101	SPID1:21255520240101 SPID2:21255520250101	SPID1:21255520260101 SPID2:21255520270101

Appendix C.

Recommended Reading List for Routing and Switching CCIE Candidates

Cisco Certification: Bridges, Routers and Switches for CCIE's 2nd Edition , Caslow and Pavlichenko, Prentice-Hall

TCP/IP Routing Volume I, Jeff Doyle, Cisco Press

TCP/IP Routing Volume II, Jeff Doyle, Cisco Press

Internet Routing Architectures 2nd Edition, Bassam Halabi, Cisco Press

Cisco BGP-4 Command and Configuration Handbook, William Parkhurst, Cisco Press

Cisco OSPF Command and Configuration Handbook, William Parkhurst, Cisco Press

Developing IP Multicast Networks Volume I, Beau Williamson, Cisco Press

IP Quality of Service, Srinivas Vegesna, Cisco Press

CCIE Practical Studies Volume I, Karl Solie, Cisco Press

Cisco Access-List Field Guide, Kent Hundley & Gil Held, Mc Graw-Hill

Hardening Cisco Routers, Thomas Akin, O'Reilly

Appendix D.

OSPF Network Type Table

Family	OSPF Network Type	Hello Addressing	Database Exchange Addressing	DR/BDR Election	Issues	Link Advertisement	Next Hop (see next page)
Broadcast algorithm	Broadcast	multicast 224.0.0.5	multicast 224.0.0.5 224.0.0.6	Yes	Partial mesh, NBMA topology hub and spoke Placement of the DR (hub) Placement of the BDR (none)	Link is advertised as prefix	Next Hop is the advertising router
	non–broadcast	unicast	unicast	Yes	Partial mesh, NBMA topology hub and spoke Placement of the DR (hub) Placement of the BDR (none) neighbor statements	Link is advertised as prefix	Next Hop is the advertising router
Point to Point algorithm	point to point	multicast 224.0.0.5	multicast 224.0.0.5	No	Flood Administration overhead related to the subnetting	Link is advertised as prefix	Next Hop is the other end of the point to point link
	point to multipoint	multicast 224.0.0.5	multicast 224.0.0.5	No	Flood host entries /32 for the ends of the mulipoint collection	Link is advertised as a number of host entires	Next Hop is the other end of the point to point link
	point to multipoint non–broadcast	unicast	unicast	No	Flood host entries /32 for the ends of the mulipoint collection	Link is advertised as a number of host entires	Next Hop is the other end of the point to point link

The Next-Hop Selection Methods of OSPF Network Types

To illustrate the relationship between the OSPF network types, consider the following two sets of diagrams. In the first set of diagrams, we compare the next-hop address selection behavior of the broadcast and non-broadcast OSPF network types. In the second set of diagrams, we compare the next-hop address selection behavior of the point-to-point, point-to-multipoint and point-to-multipoint non-broadcast OSPF network types. By reviewing these two sets of diagrams and their accompanying explanations, a clearer understanding of the differences between the two different families of OSPF network types (the family using the broadcast algorithm and the family using the point-to-point algorithm) will be obtained.

Next-Hop Selection Methods Based Upon a Broadcast Algorithm

In the diagram below, prefix "X" is advertised by R2 to R1 and R3 on an Ethernet segment. When R1 and R3 receive the "X" prefix from R2, the next-hop address associated with the prefix will be "IP2", the IP address of the R2 Ethernet interface. Since Ethernet is a shared medium, it is both logical and efficient to preserve the next-hop IP address of whichever router originally advertised the prefix on the shared segment. To change the next-hop IP address on a shared medium would result in forwarding packets to superfluous and unnecessary next-hops. By default, when an Ethernet segment is configured for OSPF, it is assigned the OSPF "broadcast" network type. The OSPF "broadcast" network type preserves the next-hop address of the router that originally advertised the prefix onto the shared segment.

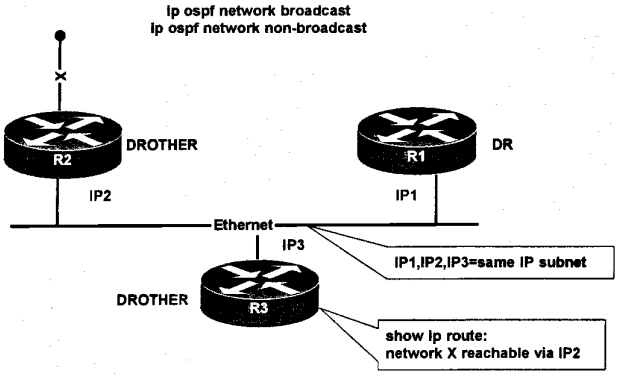

Once you understand how the next-hop is set on a segment assigned with the OSPF "broadcast" network type, it is easy to understand the method of setting the next-hop address on an OSPF "non-broadcast" network type.

In the diagram below, prefix "X" is advertised by R2 to R1 and R3 over an non-broadcast multi-access network such as ATM or Frame-Relay. The setting of the next-hop IP address for this advertised prefix is identical with the behavior observed on the Ethernet segment above. R1 and R3 will receive the "X" prefix with the next-hop set to R2. This behavior underscores the close relationship between the OSPF broadcast and non-broadcast network types. When this method of setting the next-hop is performed over a hub-and spoke NBMA network, complications with

forwarding packets and even failures can occur. Because of this, it is recommended to configure the OSPF non-broadcast network on a fully meshed NBMA network.

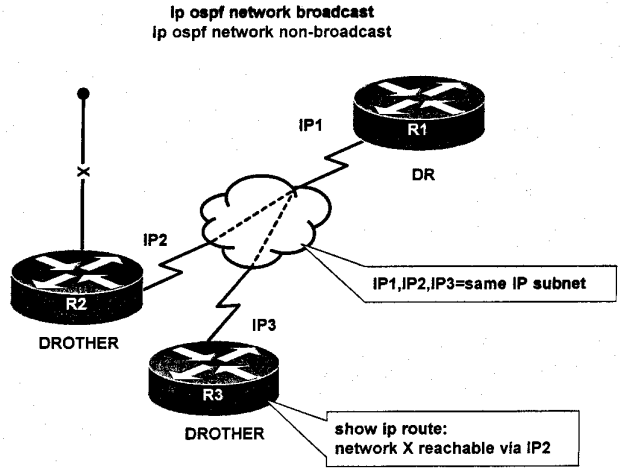

Next-Hop Selection Methods Based Upon a Point-to-Point Algorithm

In the diagram below, two separate physical point-to-point links are configured between the three routers. R1 is terminating both of the point-to-point links on separate interfaces. The resulting topology is a hub and spoke topology composed of two separate physical interfaces. When R2 advertises prefix X over a collection of point-to-point links, the next-hop is set to the advertising end of the point-to-point link. In the diagram below, the next-hop for prefix X for R1 will be R2, and the next-hop for prefix X for R2 will be R1. This type of behavior is what is applied when the OPSF "point-to-point" network type is used. The OSPF "point-to-point" network type is the default configuration for point-to-point links.

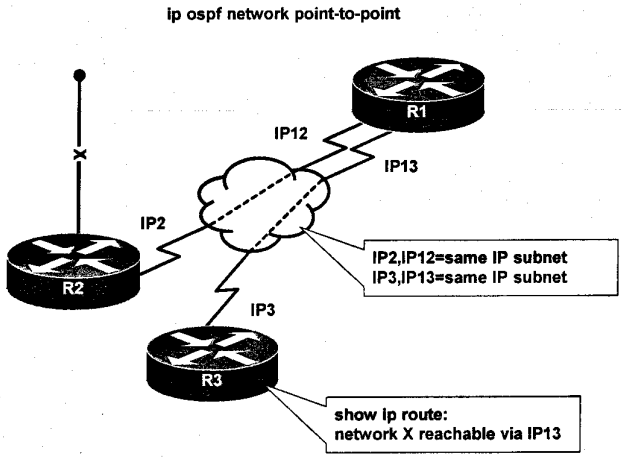

In the diagram below, a logical hub and spoke topology is presented. Unlike the topology above where R1 is maintaining two physical interfaces, R1 is now maintaining only one physical interface with two logical connections to R2 and R3. This is a common configuration in an NBMA environment. In order for this topology to apply the point-to-point next-hop selection method, these routers must be configured with the OSPF point-to-multipoint and point-to-multipoint non-broadcast networks types. When these network types are configured, the next-hop address selection process is identical with the point-to-point OSPF network type displayed above. The next-hop IP address is set to the IP address of the remote end of the point-to-point connection.

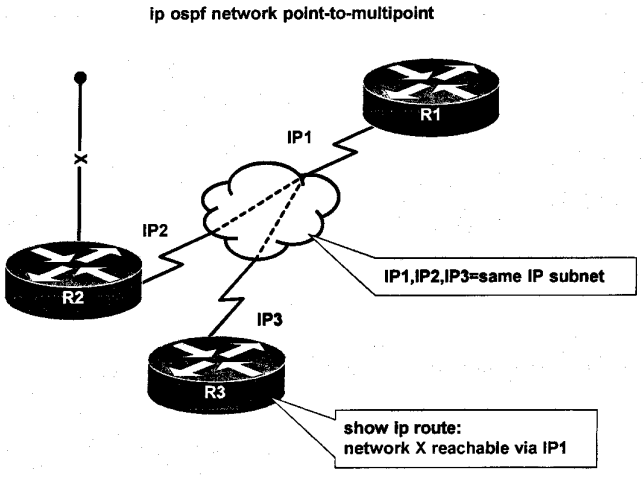

Conclusion

Since the OSPF broadcast and non-broadcast network types set the next-hop address for prefixes advertised in the exact same manner, they belong to the same family of OSPF network types. This family is called the "broadcast algorithm" family because all network types within this group treat the segment as a single "shared" or "broadcast" segment.

Since the OSPF point-to-point, point-to-multipoint and point-to-multipoint non-broadcast network types set the next-hop address for prefixes advertised in the exact same manner, they belong to the same famil t of OSPF network types. This family is called the "point-to-point algorithm" family because all network types within this group treat the segment as a collection of point-to-point links.

Appendix E.

OSPF Default Network Type Table

Interface	Encapsulation	Default OSPF network type
Serial (physical)	PPP	Point to point
Serial (physical)	HDLC	Point to point
Serial (physical) Serial (point to point) Serial (multipoint)	Frame Relay	Non broadcast Point to point Non broadcast
ATM (physical) ATM (point to point) ATM (multipoint)	All types	Non broadcast Point to point Non broadcast
Ethernet (physical) Ethernet (logical)	Ethernet ISL, dot1q	Broadcast Broadcast
Token Ring	Token Ring	Broadcast
FDDI	FDDI	Broadcast
BRI (physical)	PPP, HDLC	Point to point
Dialer profile	PPP, HDLC	Point to point
Loopback	Loopback	Loopback
Tunnel	All types	Point to Point

Appendix F.

IP Subnet Line (check www.netmasterclass.com for details)

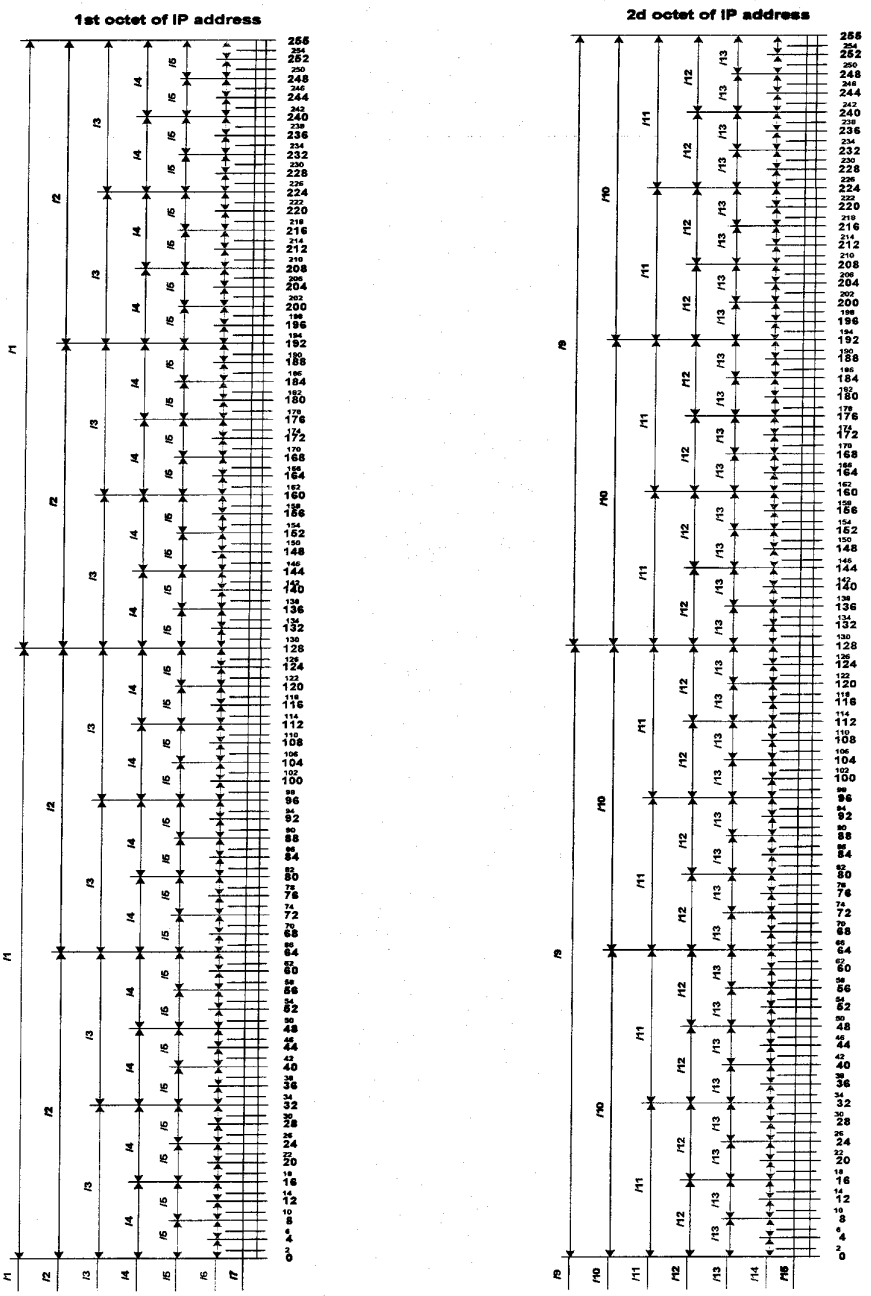

☎ 1-888-677-2669 http://www.netmasterclass.net 13530 Dulles Technology Drive, Suite #150, Herndon, VA 20171

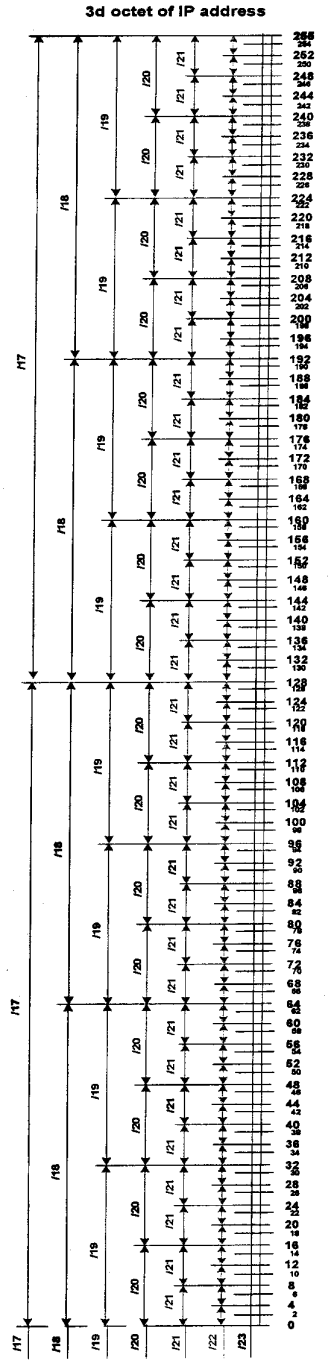

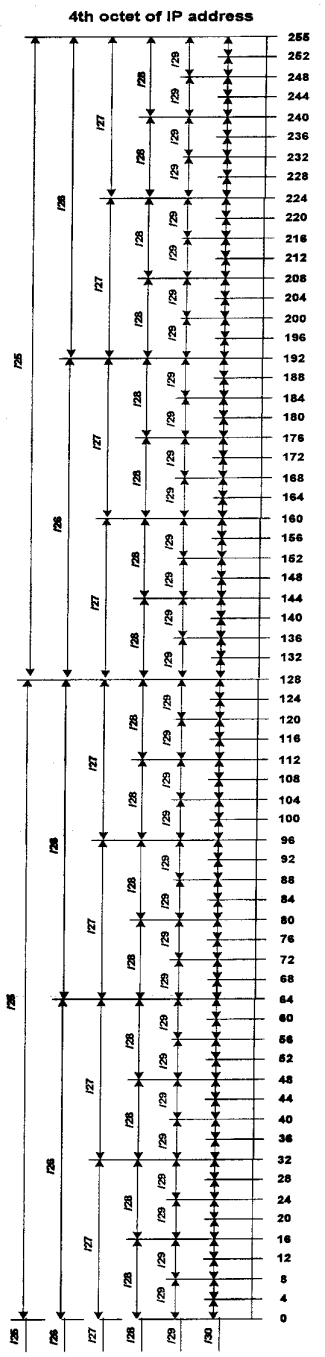

DOiT, Routing and Switching CCIE Track
Revision: 3.0

We at NetMasterClass want to thank you for selecting the NetMasterClass Lab Scenarios to help you as you continue your studies to prepare for the CCIE exam.

We based the development of the labs on our classroom training experience and our real world experience working with complex, multilayered networks for large companies. Each of us has over 10 years experience in training, consulting, and publishing technical material. As the CCIE exam changes and becomes increasingly difficult, we know you need practice labs that are challenging and reflect the level of difficulty that you will experience in the real CCIE lab exam. To this end, we believe the labs contained in this workbook are among the most challenging available in the marketplace today.

Based on our experience, we believe nothing will ever replace the classroom-training experience; therefore, we continue to enhance our RS-NMC-1 and RS-NMC-2 classes. In the classroom, we can cover additional topics such as time management, building checklists, and evaluating your strengths and weaknesses. If you have not taken one of our courses, we would love to see you in a future class.

Good luck on your studies!
We trust you will enjoy the lab scenarios.

Sincerely,

Bruce Caslow
Val Pavlichenko

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