


NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOIT-200v6

VOLUME II



DOIT-200v6 Lab 1

Multi-Topic CCIE-Level Scenario

For CCIE® Candidates

Disclaimer


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

Cisco Non-Disclosure Agreement Compliance

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

Table of Contents

1.1 Frame Relay	8
1.2 Catalyst Configuration	8
1.3 OSPF	9
1.4 RIP	9
1.5 EIGRP	10
1.6 BGP	10
1.7 IPv6	11
1.8 Traffic Management	12
1.9 Address Administration	12
1.10 Security	13
1.11 IOS Features	13
1.12 QOS	13
1.13 Catalyst Specialties	13
1.14 Gateway Redundancy	13
1.15 Multicast	14
1.16 NTP	14
Appendices	15
Hints	17
1.1 Frame-Relay.....	17
1.2 Catalyst 3550.....	17
1.3 OSPF.....	17
1.4 RIP.....	17
1.5 EIGRP.....	17
1.6 BGP.....	17
1.7 IPV6.....	18
1.8 Traffic Management.....	18
1.9 Address Administration.....	18
1.10 Security.....	18
1.11 IOS Features.....	18
1.12 QoS.....	18
1.13 Catalyst Specialties.....	18
1.14 Gateway Redundancy.....	18
1.15 Multicast.....	18
1.16 NTP.....	19




Goals and Restrictions

- IPv4 subnets displayed in the Scenario diagram belong to network 172.16.0.0/16.
- IPv6 networks used in the Scenario will use a FEC0::/125 network unless specified otherwise.
- Do not use any static routes.
- Advertise Loopback interfaces with their original masks.
- 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 any new IP addresses between R2 and CAT2.
- All IP addresses involved in this scenario must be reachable.
- Use conventional routing algorithms.

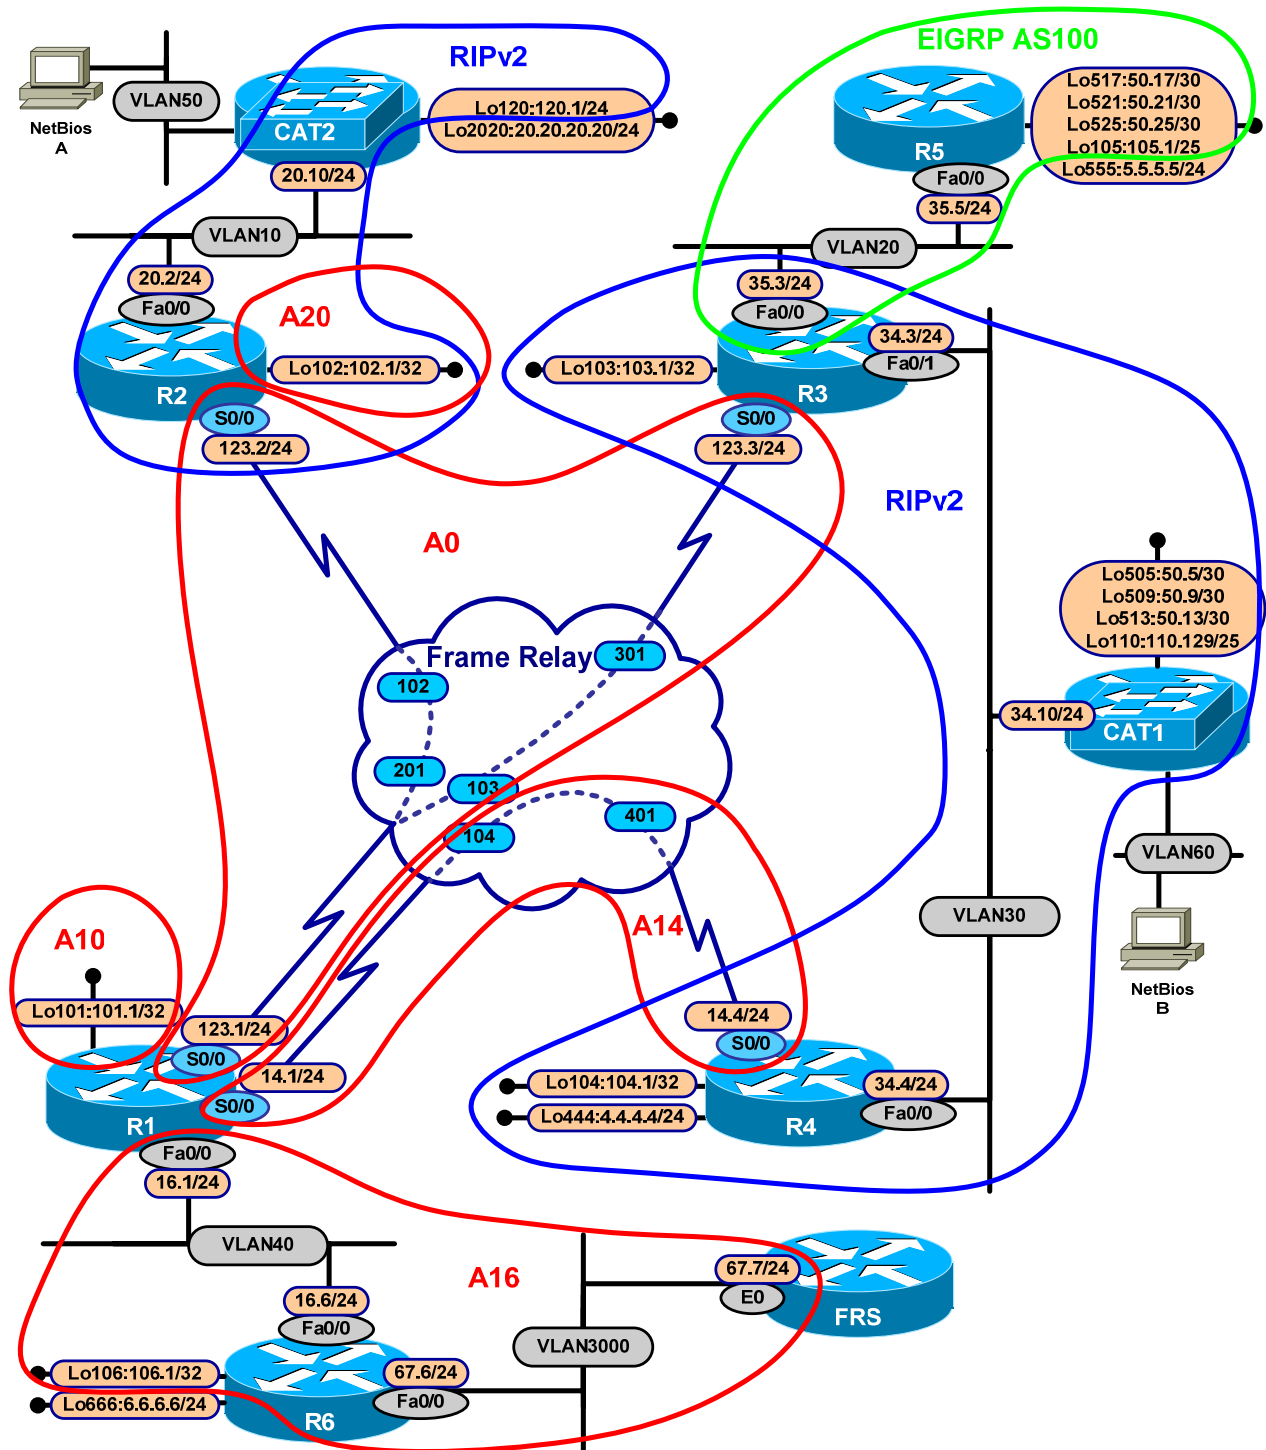


**FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM “NMC POD LAYOUT”
DISPARED IN THE APPENDIX A AT THE END OF THIS DOCUMENT.**

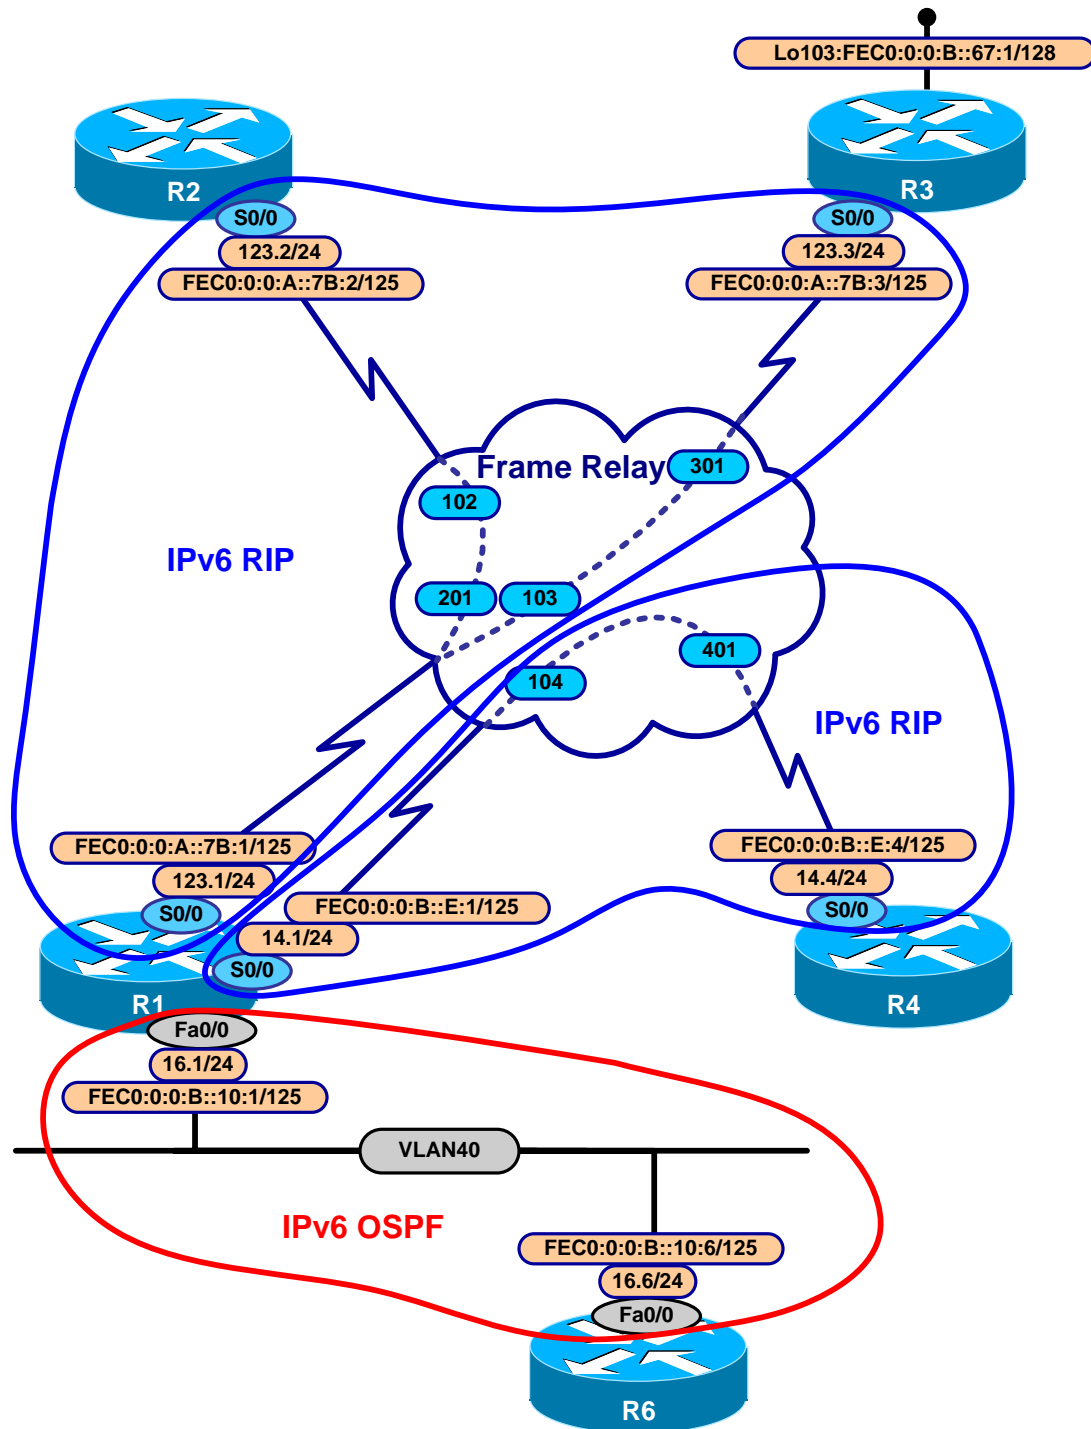
Netmasterclass 200v6 POD Layout



Netmasterclass DOiT-200v6 Scenario 1 IPv4 IGP diagram



Netmasterclass DOiT-200v6 Scenario 1 IPv6 IGP diagram



1.1 Frame Relay



The Frame-Relay Switch configuration is provided in the Appendix. Do not modify the frame relay route statements on the frame relay switch FRS.

- 1.1.1 Use only the PVCs displayed in the “Netmasterclass DOIT-200v6 Scenario 1 IPv4 IGP diagram”. No dynamic entries are allowed in the Frame Relay map tables.
- 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



Shutdown interfaces FA0/15 on both switches CAT1 and CAT2

- 1.2.1 Configure the VLANs for all Ethernet segments involved in this scenario. See the diagram “Netmasterclass DOIT-200v6 Scenario 1 IPv4 IGP diagram” and the tables provided below.

VLAN Tables

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN40
R2	FastEthernet0/0	VLAN10
R3	FastEthernet0/1	VLAN30
R3	FastEthernet0/0	VLAN20
R4	FastEthernet0/0	VLAN30
R5	FastEthernet0/0	VLAN20
R6	FastEthernet0/0	VLAN40 VLAN3000
FRS	Ethernet0	VLAN3000
CAT1	-	VLAN30
CAT2	-	VLAN10

VLAN	VLAN NAME
VLAN10	NET-10
VLAN20	NET-20
VLAN30	NET-30
VLAN40	NET-40
VLAN50	NET-50

VLAN60	NET-60
VLAN700	NET-700
VLAN800	NET-800
VLAN3000	NET-3000

- 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, as well as on the other trunks involved in this scenario.
- 1.2.4 Allow only VLAN10 on the ISL trunk.
- 1.2.5 Allow all other VLANs on the dot1q trunk over Fa0/14 interfaces of CAT1 and CAT2.
- 1.2.6 Configure the VTP mode suitable for all other tasks in this exam.

1.3 OSPF

- 1.3.1 Add all interfaces on the 172.16.123.0/24 subnet to OSPF area 0.
- 1.3.2 In area 0, use the OSPF network type which forces a DR/BDR election but uses unicast packets for Hello and Database Exchange services.
- 1.3.3 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. Advertise the Loopback network 172.16.104.1/32 in OSPF area 14.
- 1.3.4 Advertise loopback 172.16.103.1/32 in OSPF from R3 without adding it in any area and without redistribution connected prefixes. It should appear in the routing tables of the OSPF speakers R1 and R2 as type E2.
- 1.3.5 Advertise loopback 172.16.102.1/32 in OSPF area 20.
- 1.3.6 Advertise loopback 172.16.101.1/32 in OSPF area 10.
- 1.3.7 Configure OSPF area 16 and area 67 on the 172.16.16.0/24 and 172.16.67.0/24 respectively. Advertise the Loopback network 172.16.107.1/32 as E2 with the initial metric 100 and tag 200.
- 1.3.8 Set the dead interval to 20 seconds for the OSPF adjacency over the 172.16.16.0/24 link. Do not configure **ip ospf dead-interval** interface command to accomplish this task.

1.4 RIP

- 1.4.1 Configure only RIP version 2 between R2 and CAT2.

- 1.4.2 Configure only RIP version 2 between R3, R4 and CAT1.
- 1.4.3 Advertise the following loopback interfaces on CAT1:
 - 172.16.50.5/30
 - 172.16.50.9/30
 - 172.16.50.13/30
 - 172.16.110.129/25
- 1.4.4 Advertise interfaces 4.4.4.4/24 and 172.16.104.1/32 from R4 into RIPv2.
- 1.4.5 Make sure that router CAT2 receives only /24 and /32 prefixes. Do not use any filtering to accomplish this task.

1.5 EIGRP

- 1.5.1 Configure EIGRP AS 100 between R3 and R5. Configure only a single **network 172.16.35.0 0.0.0.255** statement under the EIGRP routing process
- 1.5.2 Advertise the following loopback interfaces on R5 as internal EIGRP networks:
 - 172.16.50.17/30
 - 172.16.50.21/30
 - 172.16.50.25/30
 - 172.16.105.1/25
- 1.5.3 Detect the loss of an EIGRP neighbor in a time period that is twice as fast as the default.



TO OBTAIN UNIVERSAL CONNECTIVITY, PERFORM A MUTUAL REDISTRIBUTION OF DYNAMIC INTERIOR GATEWAY PROTOCOLS: BETWEEN RIP AND OSPF ON ROUTERS R3 AND R4, BETWEEN OSPF AND EIGRP ON R3, BETWEEN RIP AND EIGRP ON R3, BETWEEN OSPF AND RIP ON R2. DO NOT PERFORM ANY OTHER REDISTRIBUTION IN THIS SCENARIO. PERFORM REDISTRIBUTE CONNECTED WHERE REQUIRED AND NOT RESTRICTED BY THE SCENARIO.

1.6 BGP

- 1.6.1 Configure AS600 on R6, AS700 on FRS and peer between them.
- 1.6.2 Configure AS200 on CAT2 and AS500 on R5.
- 1.6.3 Originate the following prefixes in BGP only within their respective ASes. These prefixes must be exchanged between all AS's:

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

Do not introduce the prefixes into IGP within AS200, AS500, AS600, AS700.

- 1.6.4 Configure AS100; do not disable synchronization within AS100. Peer AS100 and AS200 between R2 and CAT2. Peer AS500 and AS 100 between R3 and R5. Peer AS 600 and AS 100 between R6 and R1.
- 1.6.5 Provide transit paths between the above-advertised prefixes through AS100.
- 1.6.6 Use the minimum number of BGP speakers and peer relationships to accomplish this section.

1.7 IPv6

- 1.7.1 Configure a site-local subnet 7B between R1, R2 and R3. Use SLA number A for this part of your network. Use minimal number of PVCs.
- 1.7.2 Configure a site-local subnet E between R1 and R4. Use SLA number B here and in the rest of the network. Use point-to-point sub-interface on R1 and physical interface on R4.
- 1.7.3 Configure link local IPv6 addresses for all frame relay serial interfaces with corresponding IPv4 addresses according to the following table:

Router	IPv4 address	IPv6 link local address
R1	172.16.123.1	FE80::7B:1
R1	172.16.14.1	FE80::7B:14
R1	172.16.16.1	FE80::16:1
R2	172.16.123.2	FE80::7B:2
R3	172.16.123.3	FE80::7B:3
R4	172.16.14.4	FE80::7B:4
R6	172.16.16.6	FE80::16:6

- 1.7.4 Configure a site-local subnet 10 between R1 and R6.
- 1.7.5 Configure RIP on Frame Relay links, name the RIP routing process as FRAME-123 between the routers R1, R2 and R3, name the RIP routing process as FRAME-14 between the routers R1 and R4.
- 1.7.6 Assign the IPv6 address of FEC0::B:0:0:67:1/128 to Loopback100 of R3 and advertise it into the existing IPv6 RIP routing process.

- 1.7.7 Configure OSPF backbone area on the link between R1 and R6. Make sure R1 is always the DR and R6 is the DROTHER on the subnet. Do not multicast OSPF packets on the subnet between R1 and R6
- 1.7.8 Configure your network to provide reachability between R6 FastEthernet interface and R3 Fa0/1. Make sure you can ping Fa0/1 on R3 using IPv6 address. Do not assign IPV6 address to R3's Fa0/1. Apply configuration only on R4 to accomplish this task.



TO OBTAIN UNIVERSAL IPv6 CONNECTIVITY, PERFORM A MUTUAL REDISTRIBUTION OF DYNAMIC INTERIOR GATEWAY PROTOCOLS: BETWEEN IPv6 RIP AND IPv6 OSPF ON ROUTER R1. PERFORM REDISTRIBUTE CONNECTED WHERE REQUIRED AND NOT RESTRICTED BY THE SCENARIO.

1.8 Traffic Management

- 1.8.1 Workstations A and B are connected to CAT2 VLAN50 FA0/18 and CAT1 VLAN60 FA0/18 respectively. They run a NetBIOS application requiring communications between these two workstations.
- 1.8.2 Configure Catalyst switches to enable such communications.

1.9 Address Administration

- 1.9.1 Configure a DHCP server on R1 to supply the following settings to the workstations connected to ports Fa0/20 and Fa0/21 of CAT1 on VLAN40:
 - o Domain: xyz.com
 - o DNS server: 172.16.16.10
 - o Gateway
- 1.9.2 Use the gateway IP address most suitable for other tasks in this scenario.
- 1.9.3 First 10 IP addresses of VLAN 30 are reserved for the routers, printers, wireless access points.
- 1.9.4 Some users of the workstations connected to Fa0/20 and Fa0/21 of CAT1 complain about not getting the DHCP settings. Provide a solution.
- 1.9.5 Wireless workstations also obtain the IP address from this DHCP server. The administrator is concerned about IP address spoofing. Provide a solution allowing the DHCP server to explicitly know when a user logs out. Only authorized users should be able to respond to the ARP request.

1.10 Security

- 1.10.1 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.
- 1.10.2 Test it with the telnet command.

1.11 IOS Features

- 1.11.1 On router R4 create user "operator" with password "nmc". In "show run" output password must be encrypted with MD5 hash. Make sure user "operator" goes directly to level 15 upon login.

1.12 QOS

- 1.12.1 Allocate 20% of bandwidth to ftp traffic, packet size 1500 bytes.
- 1.12.2 Allocate 30% of bandwidth to http traffic, packet size 600 bytes.
- 1.12.3 Allocate 40% of bandwidth to udp port 5001 traffic, packet size 300 bytes.
- 1.12.4 Allocate 10% of bandwidth to other types of traffic, packet size 1500 bytes.
- 1.12.5 The traffic listed above is sent out the interface S0/0 of R4.
- 1.12.6 Use Custom Queuing to accomplish tasks in this section.

1.13 Catalyst Specialties

- 1.13.1 Change the MTU size on VLAN 3000 to twice higher than default.
- 1.13.2 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.
- 1.13.3 The visitors' configuration file is on the switch with the ip address 172.16.200.1.

1.14 Gateway Redundancy

- 1.14.1 Configure an HSRP standby group 3000 between R1 and R6. The HSRP group must be MD5 authenticated with the password "doitlab1"
- 1.14.2 Make R1 a primary and R6 a backup gateway for the workstations connected to VLAN40.

1.14.3 The virtual gateway IP address is 172.16.16.254.

1.15 Multicast

1.15.1 Build a Shortest Path Tree (SPT) only. Routers R1, R2, R3, R4 and R6 will be the members of this tree.

1.15.2 Join one of the loopback interfaces on each router within the SPT to multicast group 225.5.5.5.

1.15.3 Make sure you receive replies from all of the routers listed above on R5.

1.16 NTP

1.16.1 Configure R1 as NTP master stratum 3.

1.16.2 Configure R4 as an NTP client.


1.16.3 Configure a peer association between R3 and R4.

Appendices

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)



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 on this page, 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
  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
```


Hints

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 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.

1.4 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.

1.5 EIGRP

No hint

1.6 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.

1.7 IPV6

Remember to map the IPV6 addresses on the multipoint Frame-Relay links. First step in RIP configuration for IPV6 is enabling it on the interface using **ipv6 rip id enable** and mapping link-local addresses.

1.8 Traffic Management

Can you bridge the NON IP traffic between two VLANs on the Catalyst switches?

1.9 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.

1.10 Security

No hint.

1.11 IOS Features

No hint.

1.12 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.

1.13 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.

1.14 Gateway Redundancy

No hint.

1.15 Multicast

Do not forget about the RPF lookup issue. Make sure you understand what incoming interface on each multicast router is.

1.16 NTP

Recall the difference between NTP server and peer association.