


NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOIT-200v6

VOLUME II



DOIT-200v6 Lab 3

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

3.1 Frame Relay	8
3.2 Catalyst Configuration	8
3.3 OSPF	9
3.4 RIP	9
3.5 EIGRP	9
3.6 BGP	10
3.7 Router Maintenance	10
3.8 Security	10
3.9 IPv6	11
3.10 QOS	12
3.11 Catalyst Specialties	12
3.12 Address Administration	12
3.13 Multicast	13
Appendices	14
Hints	16
3.1 Frame-Relay.....	16
3.2 Catalyst Configuration	16
3.3 OSPF	16
3.4 RIP.....	16
3.5 EIGRP.....	16
3.6 BGP	16
3.7 Router Maintenance	17
3.8 Security.....	17
3.9 IPv6.....	17
3.10 QoS.....	17
3.11 Catalyst Specialties.....	17
3.12 Address Administration	17
3.13 Multicast.....	17




Goals and Restrictions

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

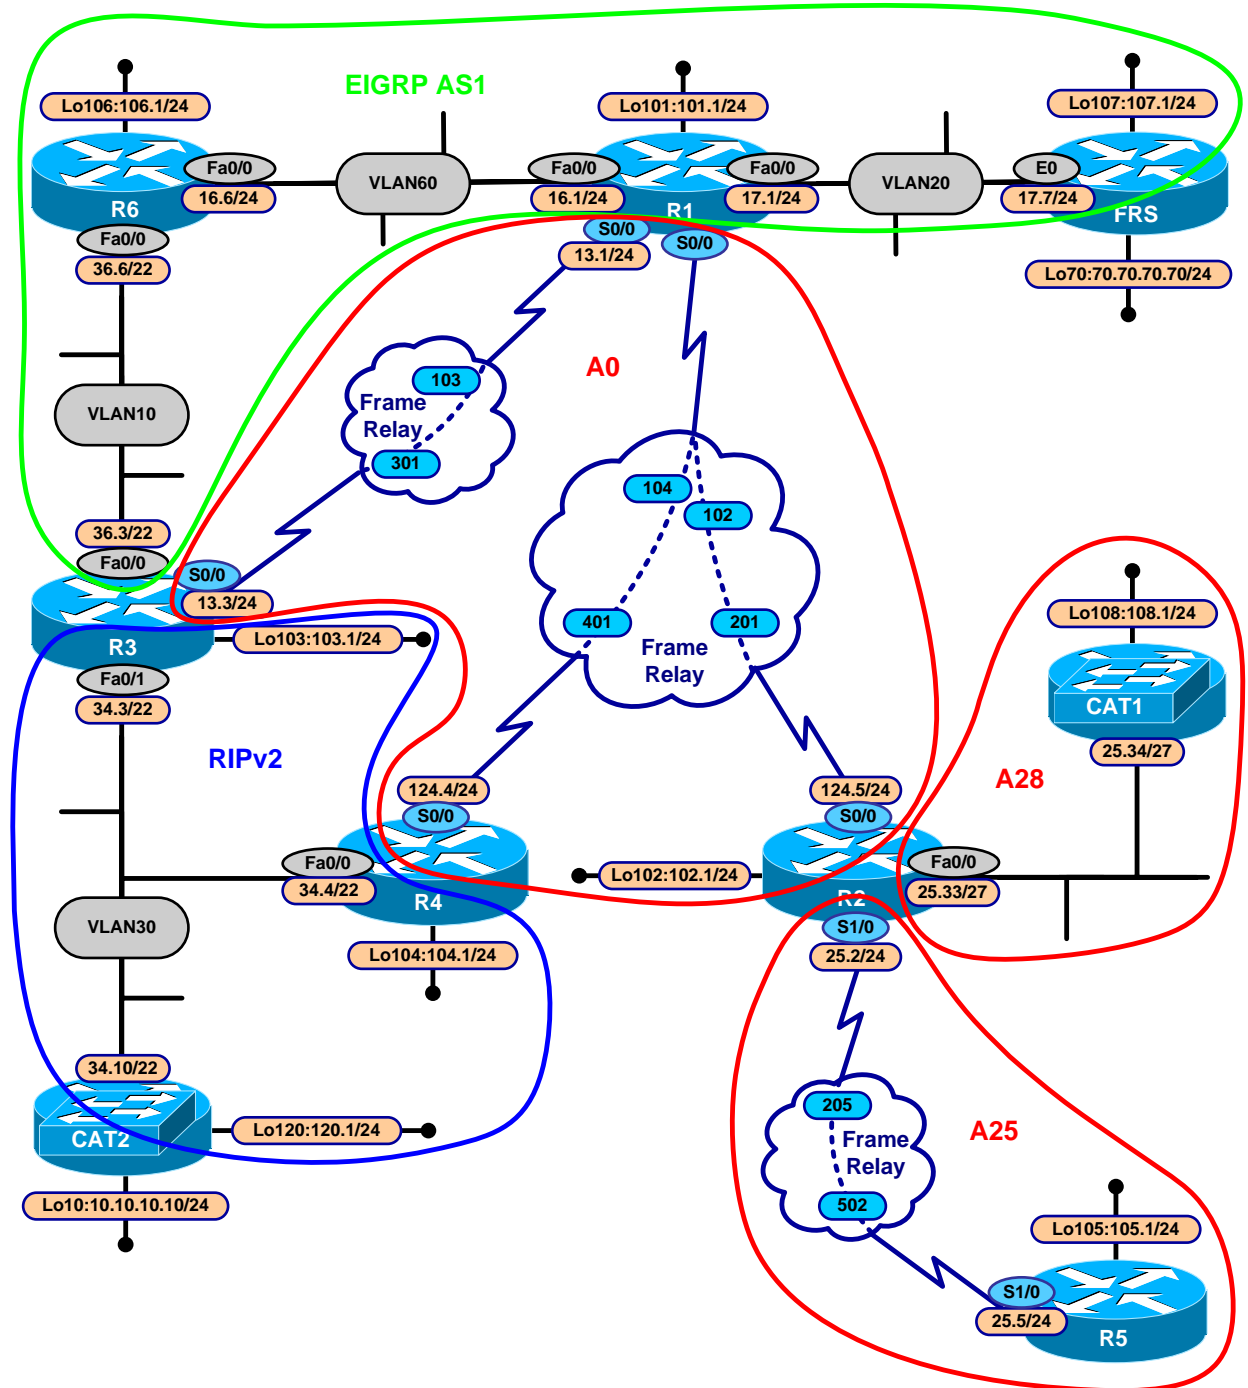


FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM "NMC POD LAYOUT" DISPLAYED IN THE APPENDIX A AT THE END OF THIS DOCUMENT.


Netmasterclass 200v6 POD Layout




Netmasterclass DOIT-200v6 Scenario 3 IPv4 IGP diagram



Netmasterclass DOiT-200v6 Scenario 3 IPv6 IGP diagram




3.1 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.

- 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. Also 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



Shutdown interfaces FA0/15 on both switches CAT1 and CAT2

- 3.2.1 Configure VLAN's according to the "VLAN Configuration Table" and diagram.

VLAN Configuration Table

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

- 3.2.2 Use a VTP mode that does not propagate VLAN information over the trunks.
- 3.2.3 On the trunks allow only the VLAN's involved in this scenario.
- 3.2.4 Routing on CAT1 should be configured as appropriate for other tasks.
- 3.2.5 The trunking protocol dot1q should be configured for this scenario.

3.3 OSPF

- 3.3.1 Configure OSPF area 0 on the subnet 172.16.124.0/24.
- 3.3.2 R2 and R4 must possess unique next hop addresses for traffic forwarded to R1.
- 3.3.3 Configure OSPF area 0 between R1 and R3. Elect R1 as DR. Make sure R1 is still DR after you reboot it. R1 and R3 should automatically discover each other.
- 3.3.4 Configure Area 25 between R2 and R5. Use a network type with a default Hello time of 10 seconds. Do not elect a DR/BDR. Make sure network 172.16.108.0/24 is in the Area 25 OSPF topology table, but not in R5's routing table. This network does not have to be reachable from R5.
- 3.3.5 Configure Area 28 between R2 and CAT1. Make sure that CAT1 receives all of the routes except for network 172.16.105.0/24 and its subnets. These prefixes should not be in CAT1's routing table or its OSPF database. This network does not have to be reachable from CAT1.

3.4 RIP

- 3.4.1 Configure RIP version 2 between R3, R4 and CAT2. Allow RIP updates only on VLAN30.

3.5 EIGRP

- 3.5.1 Configure EIGRP AS1 between FRS, R1, R6 and R3.
- 3.5.2 Do not form an EIGRP adjacency on subnet 172.16.13.0/24.
- 3.5.3 Advertise network 172.16.106.0/24 from R6 and network 172.16.107.0/24 from FRS.
- 3.5.4 Authenticate all EIGRP adjacencies with password "nmc"
- 3.5.5 When redistributing OSPF into EIGRP on R1, permit only routes with OSPF metrics between 0 and 1000, inclusive.



TO OBTAIN UNIVERSAL CONNECTIVITY, PERFORM MUTUAL REDISTRIBUTION BETWEEN IGPS ON R1, R3 AND R4. MAKE SURE REDUNDANT PHYSICAL PATHS COULD BE USED IN CASE OF LINK FAILURE. DO NOT PERFORM ANY OTHER REDISTRIBUTION IN THIS SCENARIO. PERFORM REDISTRIBUTE CONNECTED WHERE REQUIRED AND NOT RESRICTED BY THE SCENARIO.

3.6 BGP

- 3.6.1 Configure AS 700 on FRS.
- 3.6.2 Configure AS 100 on R1, R2 and R6. Do not use the synchronization method. Do not form a BGP peer relationship between R2 and R6.
- 3.6.3 Peer AS700 and AS100 between FRS and R1.
- 3.6.4 Peer AS 100 and AS 300 between routers R6 and R3 as well as R4 and R2.
- 3.6.5 Configure AS 300 on R3 and R4. Use the synchronization method in AS300.
- 3.6.6 Configure AS1000 on CAT2 and peer it with both routers of AS300.
- 3.6.7 Advertise network 70.70.70.0/24 on AS700 as incomplete.
- 3.6.8 Advertise network 10.10.10.0/24 on AS1000 as IGP.
- 3.6.9 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.
- 3.6.10 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.
- 3.6.11 Accomplish the previous two tasks with LOCAL PREF manipulation.

3.7 Router Maintenance

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

3.8 Security

- 3.8.1 None of the devices on VLAN30 should be able to telnet straight to R5 unless SALLY authenticates herself on R1.
- 3.8.2 The users on VLAN30 should be able to maintain a telnet session for no longer than 10 minutes after SALLY logs onto router R1.

3.8.3 Limit your configuration so that other traffic involved in this scenario is not disrupted.

3.9 IPv6

Addressing: Assign the following addresses to the indicated interfaces. Link-local addresses are provided on the Frame-Relay links to simplify mapping.

Router	Interface	IPv6 address
R1	F0/0.60	FEC0::16:1/125
R1	S0/0	FEC0::13:1/125 FE80::1 link-local
R1	S0/0.1	FEC0::14:1/125 FE80::1 link-local
R1	S0/0.2	FEC0::12:1/125 FE80::1 link-local
R2	S0/0	FEC0::12:2/125 FE80::2 link-local
R2	S1/0	FEC0::25:2/125 FE80::2 link-local
R3	S0/0	FEC0::13:3/125 FE80::3 link-local
R4	S0/0	FEC0::14:4/125 FE80::4 link-local
R5	S1/0	FEC0::25:5/125 FE80::5 link-local
R5	Loopback 105	FEC0::105:1/125
R5	Loopback 2001	2001::105:1/125
R6	F0/0.60	FEC0::16.6/125

IPv6 Routing

- 3.9.1 Configure IPv6 RIP between R1, R3 and R4.
- 3.9.2 Configure OSPF area 0 on Frame-Relay link between R1 and R2, and on the link between R1 and R6.
- 3.9.3 Configure OSPF area 25 on Frame-Relay link between R2 and R5.
- 3.9.4 OSPF neighbors must elect a DR on every link. OSPF neighbors must find each other automatically on all links

IPv6 Redistribution

- 3.9.5 Redistribute as necessary in order to provide full IPv6 unicast reachability. This requirement excludes 2001::105:1 on R5.

IPv6 BGP

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

3.10 QOS

- 3.10.1 Configure QOS on the GigabitEthernet0/1 interface of CAT2 using the most simplified method of configuration and allowing the discovery of IP phones.
- 3.10.2 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.

3.11 Catalyst Specialties

- 3.11.1 Enable message logging on CAT1.
- 3.11.2 Write messages to the file nmc.log stored on the local flash.
- 3.11.3 Do not allow the logging process to exceed a file size that is twice as large as the default size.
- 3.11.4 Only error messages about software or hardware malfunctions should be logged.

3.12 Address Administration

- 3.12.1 Besides two devices connected to VLAN10, the rest of the address space of VLAN10 is taken for workstations.
- 3.12.2 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.
- 3.12.3 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.
- 3.12.4 All workstations are Microsoft clients using a hybrid NetBIOS node type.

- 3.12.5 In case of failure of either Frame Relay link on R3 or the F0/0.60 link on R6, the traffic should be sent through the remaining link.
- 3.12.6 When the failed link comes back on, return to the previous forwarding scheme.

3.13 Multicast


- 3.13.1 Configure PIM neighbor relationships between R4 and R1, R1 and R2 as well as R2 and R5.
- 3.13.2 The PIM network is based on the Flood and Prune concept.
- 3.13.3 Generate traffic to multicast group 233.3.3.3 from CAT2.
- 3.13.4 Join one of the interfaces on routers R4, R1, R2 and R5 to group 233.3.3.3.
- 3.13.5 Make sure CAT2 receives ping replies from all multicast routers

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

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 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 to 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. Selecting the correct OSPF network type can perform this.

3.4 RIP

No hint

3.5 EIGRP

No hint

3.6 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.

3.7 Router Maintenance

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

3.8 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.

3.9 IPv6

No hint.

3.10 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.

3.11 Catalyst Specialties

No hint.

3.12 Address Administration

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.

3.13 Multicast

No hint.