


NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II



DOiT-200v6 Lab 16

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

16.1 Frame Relay	8
16.2 Catalyst Configuration	8
16.3 OSPF	9
16.4 RIP	9
16.5 EIGRP	9
16.6 On-Demand Routing	9
16.7 Redistribution	10
16.8 BGP	10
16.9 Address Administration	11
16.10 NTP	11
16.11 IPv6	11
16.12 QOS	12
16.13 Catalyst Specialties	12
16.14 Multicast	12
Appendices	13
Hints	15
16.1 Frame-Relay	15
16.2 Catalyst Configuration.....	15
16.3 OSPF	15
16.4 RIP	15
16.5 EIGRP	15
16.6 ODR	15
16.7 Redistribution	15
16.8 BGP.....	15
16.9 Address Administration	15
16.10 NTP	15
16.11 IPV6	16
16.12 QoS	16
16.13 Catalyst Specialties.....	16
16.14 Multicast.....	16




Goals and Restrictions

- IP subnets on the diagram belong to network 160.16.0.0/16.
- 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, unless specified otherwise.
- 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.

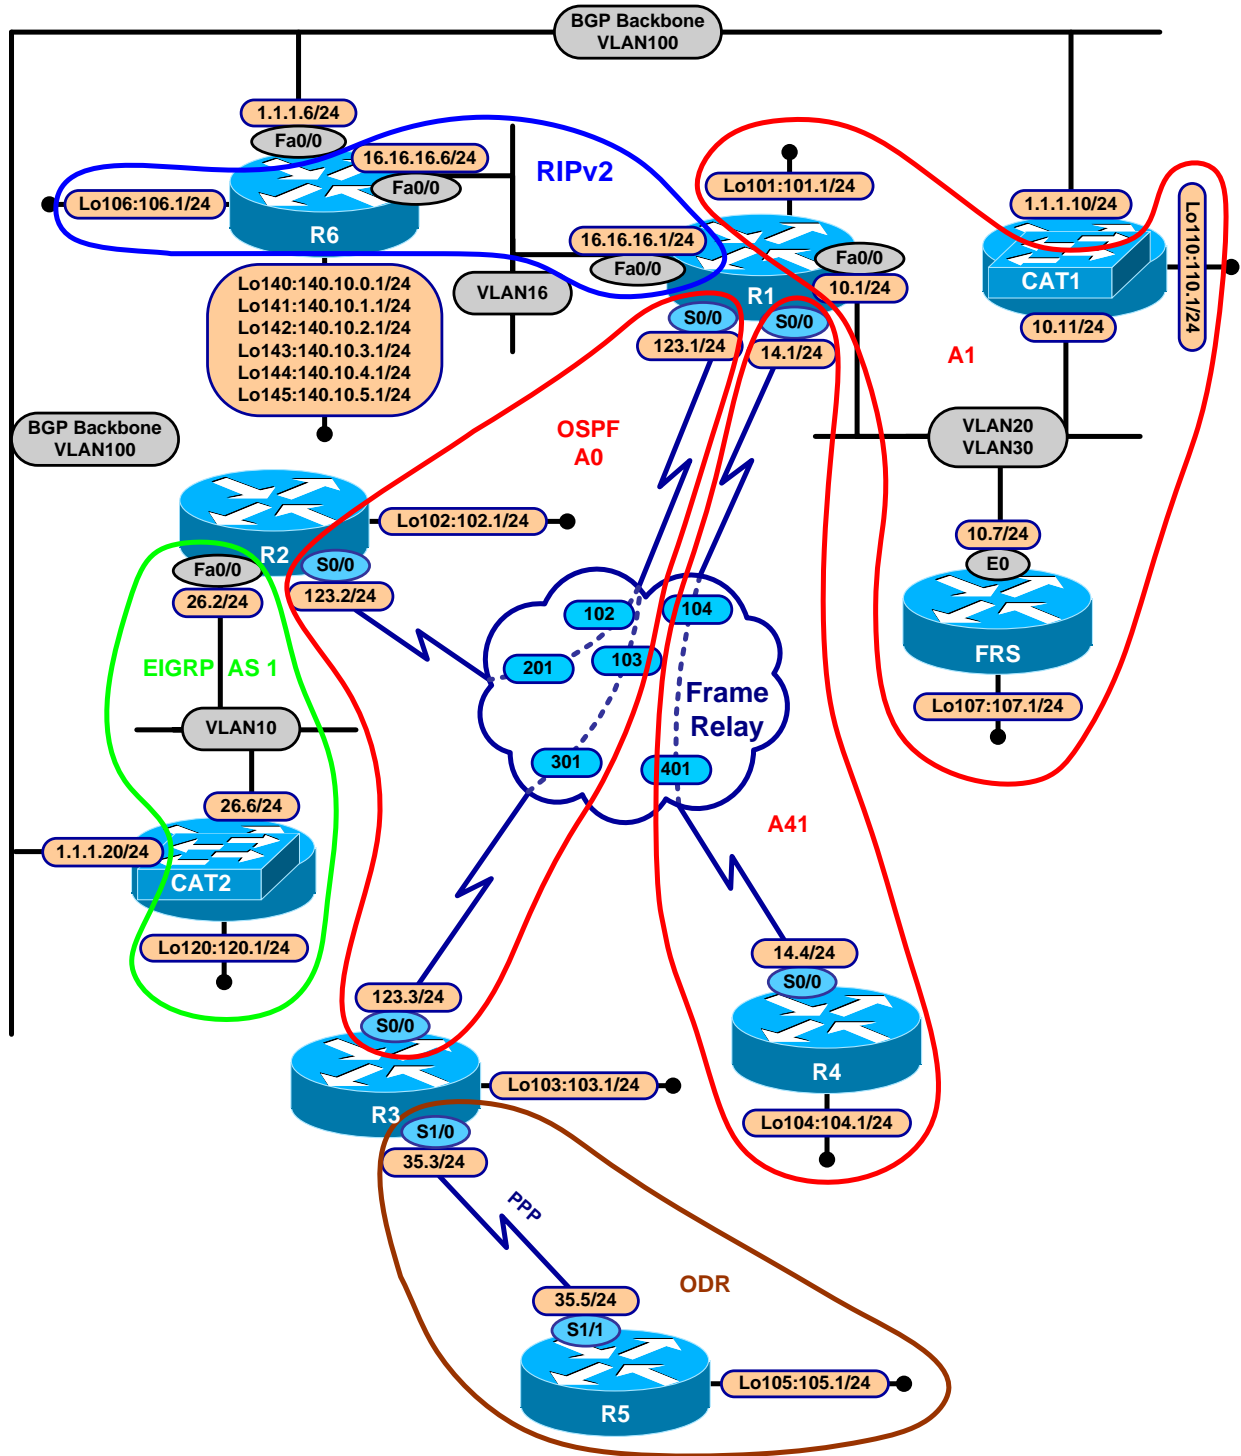


**FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM
"NETMASTERCLASS 200V6 POD LAYOUT".**

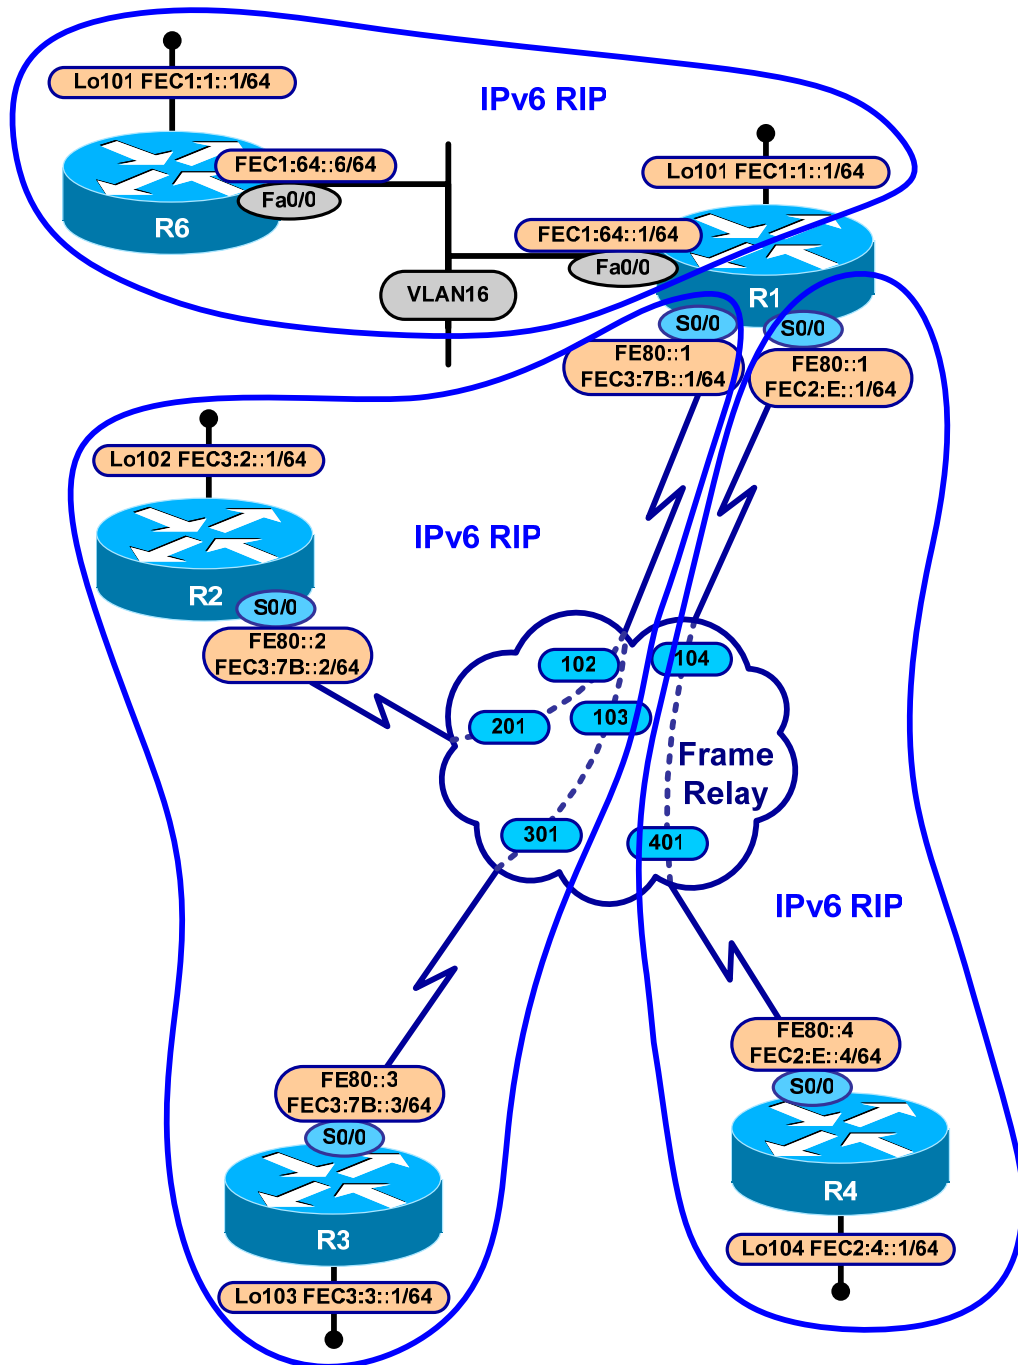
Netmasterclass 200v6 POD Layout



Netmasterclass DOiT-200v6 Scenario 16 IPv4 IGP diagram



Netmasterclass DOiT-200v6 Scenario 16 IPv6 IGP diagram



16.1 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.

- 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



Shutdown interfaces FA0/15 on both switches CAT1 and CAT2

- 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 160.16.10.11/24.
- 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.

VLAN Configuration Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN 16 VLAN20 VLAN30
R2	FastEthernet0/0	VLAN10
R3	FastEthernet0/1	-
R3	FastEthernet0/0	-
R4	FastEthernet0/0	-
R5	FastEthernet0/0	-
R6	FastEthernet0/0	VLAN 16 VLAN100
FRS	Ethernet0	VLAN30
CAT1	-	VLAN 20 VLAN100
CAT2	-	VLAN 10 VLAN100

- 16.2.4 Enable trunking only on the necessary ports. Permit only the required VLANs across configured trunks. Create trunk links of Catalyst ports F0/13 and F0/14. Do not permit VLAN 100 across any trunk link between the Catalysts.
- 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 OSPF

- 16.3.1 Configure the Frame-Relay network 160.16.123.0/24 as the OSPF backbone area. Do not elect a DR/BDR in the OSPF area.
- 16.3.2 Create the necessary OSPF topology on R1 using a single OSPF process. You are permitted only the following network statement under this process:

```
network 160.60.123.0 0.0.0.255 area 0
```
- 16.3.3 Configure OSPF on the link between R1 and R4. Place this link in OSPF area 41. Make R4 the designated router for the R1/R4 link.
- 16.3.4 Configure OSPF MD5 authentication for area 0.
- 16.3.5 Advertise FRS and the Catalyst VLAN 20 interface as one 160.16.10.0/24 OSPF network. Assign this network to OSPF area 1.
- 16.3.6 Advertise Loopback 160.16.101.0/24 as part of OSPF area 1. Remember, you are only allowed one network statement under the OSPF process on R1!

16.4 RIP

- 16.4.1 Configure RIP over the connection between R1 and R6. Do not broadcast or multicast the RIP updates.
- 16.4.2 R6 should receive the minimal number of necessary prefixes from R1 via RIP. It may not be a 0.0.0.0/0 route, and you may not create summaries on any interface.
- 16.4.3 Make sure R1 has the RIP prefix 160.60.106.0/24 in its routing table.

16.5 EIGRP

- 16.5.1 Configure EIGRP AS 1 between R2 and CAT2. Make sure that EIGRP advertises only over VLAN10 interfaces.
- 16.5.2 Send a summary for the entire Class B range from R2 to CAT2 only.

16.6 On-Demand Routing

- 16.6.1 Configure ODR between R3 and R5. R5 is permitted 0.0.0.0/0 route.

16.7 Redistribution

- 16.7.1 Mutually redistribute between RIP and OSPF on R1.
- 16.7.2 Mutually Redistribute between EIGRP and OSPF on R2.
- 16.7.3 Redistribute ODR into OSPF on R3.
- 16.7.4 On R3, redistribute 160.16.103.0 into OSPF as a Type 1 External route. This should be the only OSPF Type 1 prefix in the network.



NOTE: VLAN 100 IS USED FOR THE BGP BACKEND CONNECTIVITY ONLY. ITS NETWORK 1.1.1.0/24 IS NOT PART OF ANY IGP AND DOES NOT HAVE TO BE REACHABLE.

16.8 BGP

- 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 a 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:
 - o 140.10.0.0/24
 - o 140.10.1.0/24
 - o 140.10.2.0/24
 - o 140.10.3.0/24
 - o 140.10.4.0/24
 - o 140.10.5.0/24
- 16.8.4 Configure the BGP router-id 160.16.110.1 and 160.16.120.1 on CAT1 and CAT2 respectively.
- 16.8.5 Make R1, R3, R4 and FRS BGP speakers within AS 100. Do not peer R4 and FRS.
- 16.8.6 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 R3 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. Note the mask length in this requirement!

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 Address Administration

16.9.1 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. Its IP address is 1.1.1.1.

16.10 NTP

16.10.1 Make R1 the primary time source. Set R1 as a stratum 5.

16.10.2 Configure router R2 so that it obtains its time from R1.

16.10.3 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.

16.10.4 Configure CAT1 as an NTP client. Enable Daylight Savings Time.

16.11 IPv6

16.11.1 Assign the following IPv6 addresses. Map as necessary, and verify same-subnet reachability.

Router	Interface	Link-local IPv6 addresses	Site-local IPv6 address
R1	Serial 0/0.123	FE80::1	FEC3:7B::1/64
R1	Serial 0/0.14	FE80::1	FEC2:E::1/64
R1	F0/0.16	Default	FEC1:64::1/64
R1	Loopback 0	Default	FEC1:1::1/64
R2	S0/0.123	FE80::2	FEC3:7B::2/64
R2	Loopback 0	Default	FEC3:2::1/64
R3	Serial 0/0	FE80::3	FEC3:7B::3/64
R3	Loopback 0	Default	FEC3:3::1/64
R4	Serial 0/0.14	FE80::4	FEC2:E::4/64
R4	Loopback0	Default	FEC2:4::1/64
R6	F0/0.16	Default	FEC1:64::6/64
R6	Loopback 0	Default	FEC1:6::1/64

16.11.2 Place all of the above interfaces into IPv6 RIP with tag "RIP". Make sure all IPv6 routers can ping all IPv6 Site-Local addresses.

16.11.3 Configure R1 to send only a default route to R4. Configure R1 to send a single, longest summary to R6 representing the networks connected to R2, R3 and R4.

16.12 QOS

16.12.1 A traffic generator on VLAN100 is connected to port FastEthernet 0/24 of CAT1. It's generating 5 UDP packets per second. Packet size is 1024 Bytes. The UDP stream is destined to the VLAN20 interface of CAT1 with destination port 5011.

16.12.2 Police this traffic to 8000 bit/sec on the interface FastEthernet 0/24 of CAT1. Allow a minimal burst size.

16.12.3 Excess traffic should be dropped.

16.13 Catalyst Specialties

16.13.1 Allow only a workstation with a Data-Link address of 0050.04fd.9f73 to use port fa0/20 of CAT2.

16.13.2 Allow TELNET access to CAT1 from the loopback interface of R3 only.

16.13.3 Configure a greeting message "Welcome to RS-NMC-2!" on CAT1.

16.14 Multicast

16.14.1 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.

16.14.2 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.

16.14.3 Join Ethernet interfaces of VLAN 20 and 30 to the group 229.9.9.9.

16.14.4 Statically configure a MAC table entry for ports in VLAN 20 and VLAN 30 for 229.9.9.9.


16.14.5 Ping the multicast group 229.9.9.9 from R4 and CAT1.

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 72000
  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
  no shutdown
!
interface Serial1
  no ip address
  encapsulation frame-relay
  clockrate 72000
  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
  no shutdown
!
interface Serial2
  no ip address
  encapsulation frame-relay
  clockrate 72000
  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
  no shutdown
!
interface Serial3
  no ip address
  encapsulation frame-relay
  no ip route-cache
  clockrate 72000
  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
  no shutdown
```

Hints

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.

16.2 Catalyst Configuration

Map out the VLANs. If a VLAN cannot be on trunk, then it must be on an access link!

16.3 OSPF

Check out the new interface-level features of 12.3T.

16.4 RIP

Consider RIP's auto-summary feature.

16.5 EIGRP

The recommended method for controlling which interfaces are participating in EIGRP is to configure network commands with the appropriate masks.

16.6 ODR

No hint.

16.7 Redistribution

Don't forget prefix 160.60.35.0!

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 Address Administration

No hint.

16.10 NTP

No hint.

16.11 IPV6

Beware of Split Horizon issues.

16.12 QoS

No hint.

16.13 Catalyst Specialties

No hint.

16.14 Multicast

PIM Dense Mode uses the flood and prune protocol. Turn off IGMP snooping and statically map.