


NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOIT-200v6

VOLUME II



DOIT-200v6 Lab 8

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

8.1 Frame Relay	8
8.2 Catalyst Configuration	8
8.3 Frame Relay QoS	9
8.4 OSPF	9
8.5 RIP	9
8.6 EIGRP	9
8.7 BGP	10
8.8 VPN	11
8.9 Router Maintenance	11
8.10 IP Features	11
8.11 IPv6 IGP	11
8.12 IPv6 BGP	12
8.13 QoS	12
8.14 IOS Specialties	12
8.15 Catalyst Specialties	12
8.16 Address Administration	13
8.17 Multicast	13
Appendices	14
Hints	16
8.1 Frame Relay	16
8.2 Catalyst configuration	16
8.3 Frame Relay QoS.....	16
8.4 OSPF.....	16
8.5 RIP.....	16
8.6 EIGRP.....	16
8.7 BGP	16
8.8 VPN	16
8.9 Router Maintenance	16
8.10 IP Features	16
8.11 IPv6 IGP.....	16
8.12 IPv6 BGP	17
8.13 QoS.....	17
8.14 IOS Specialties	17
8.15 Catalyst specialties	17
8.16 Address administration	17
8.17 Multicast.....	17




Goals and Restrictions

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

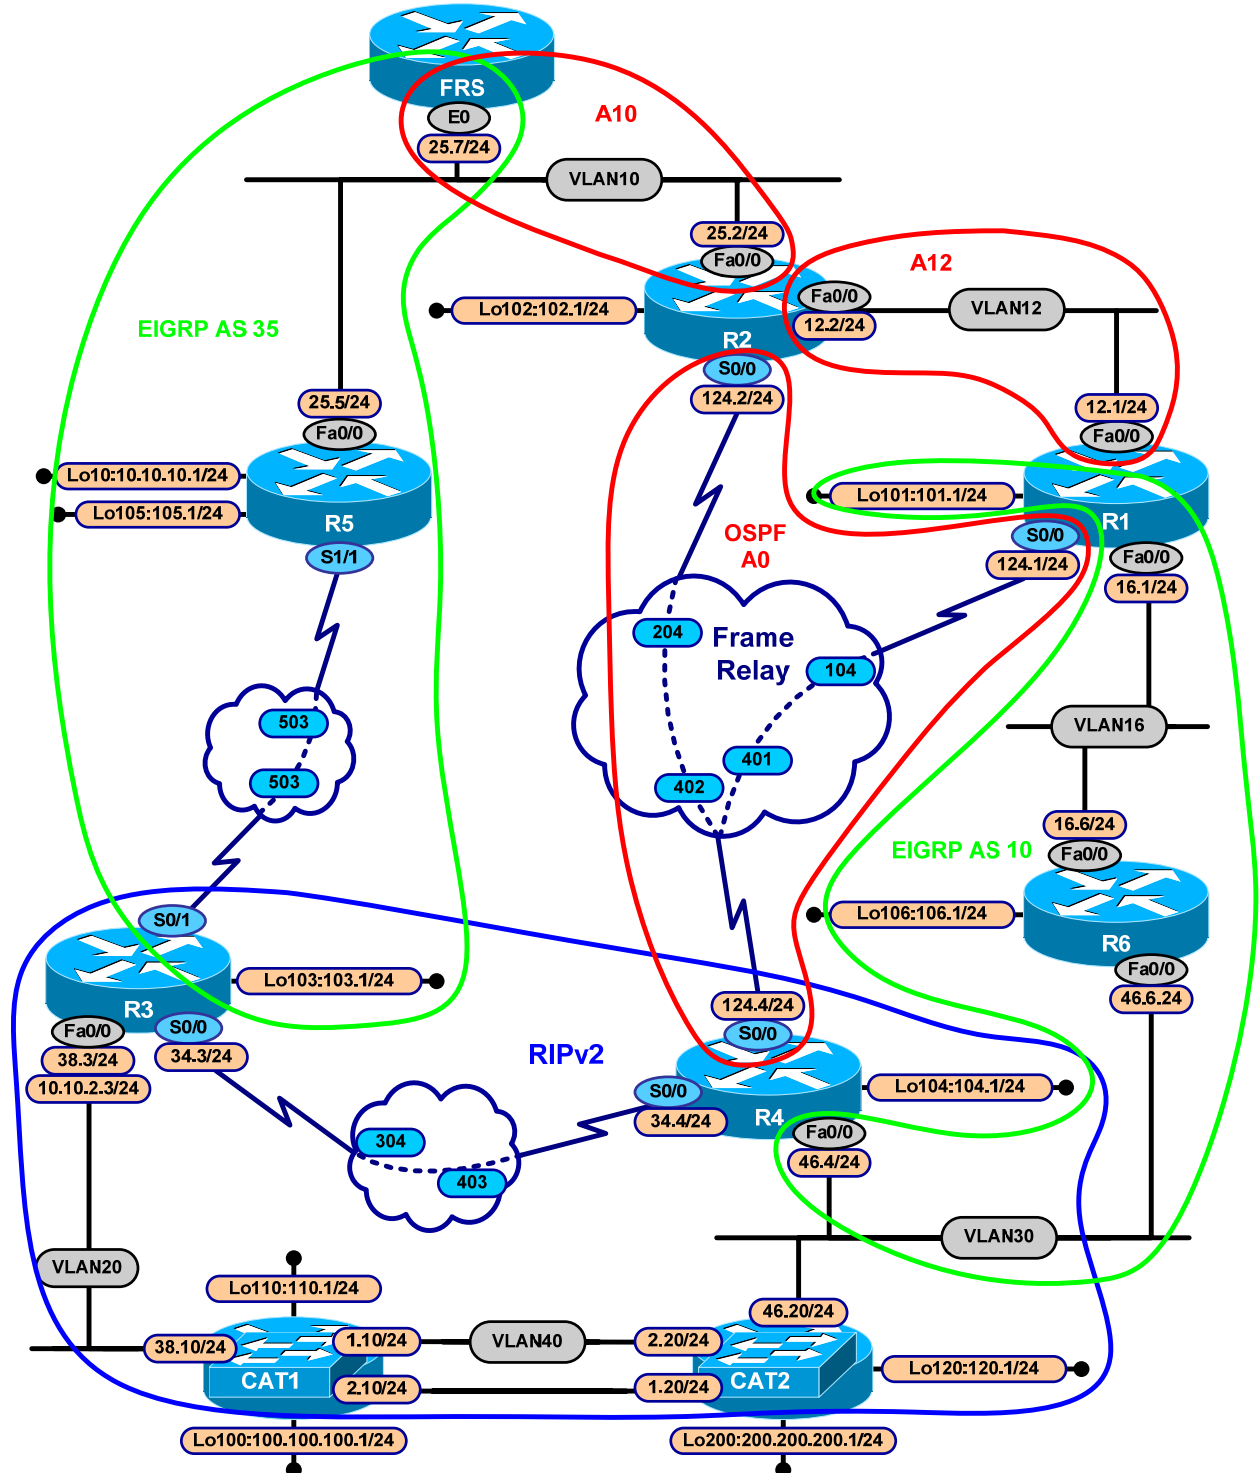


**FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM
"NETMASTERCLASS 200V6 POD LAYOUT".**

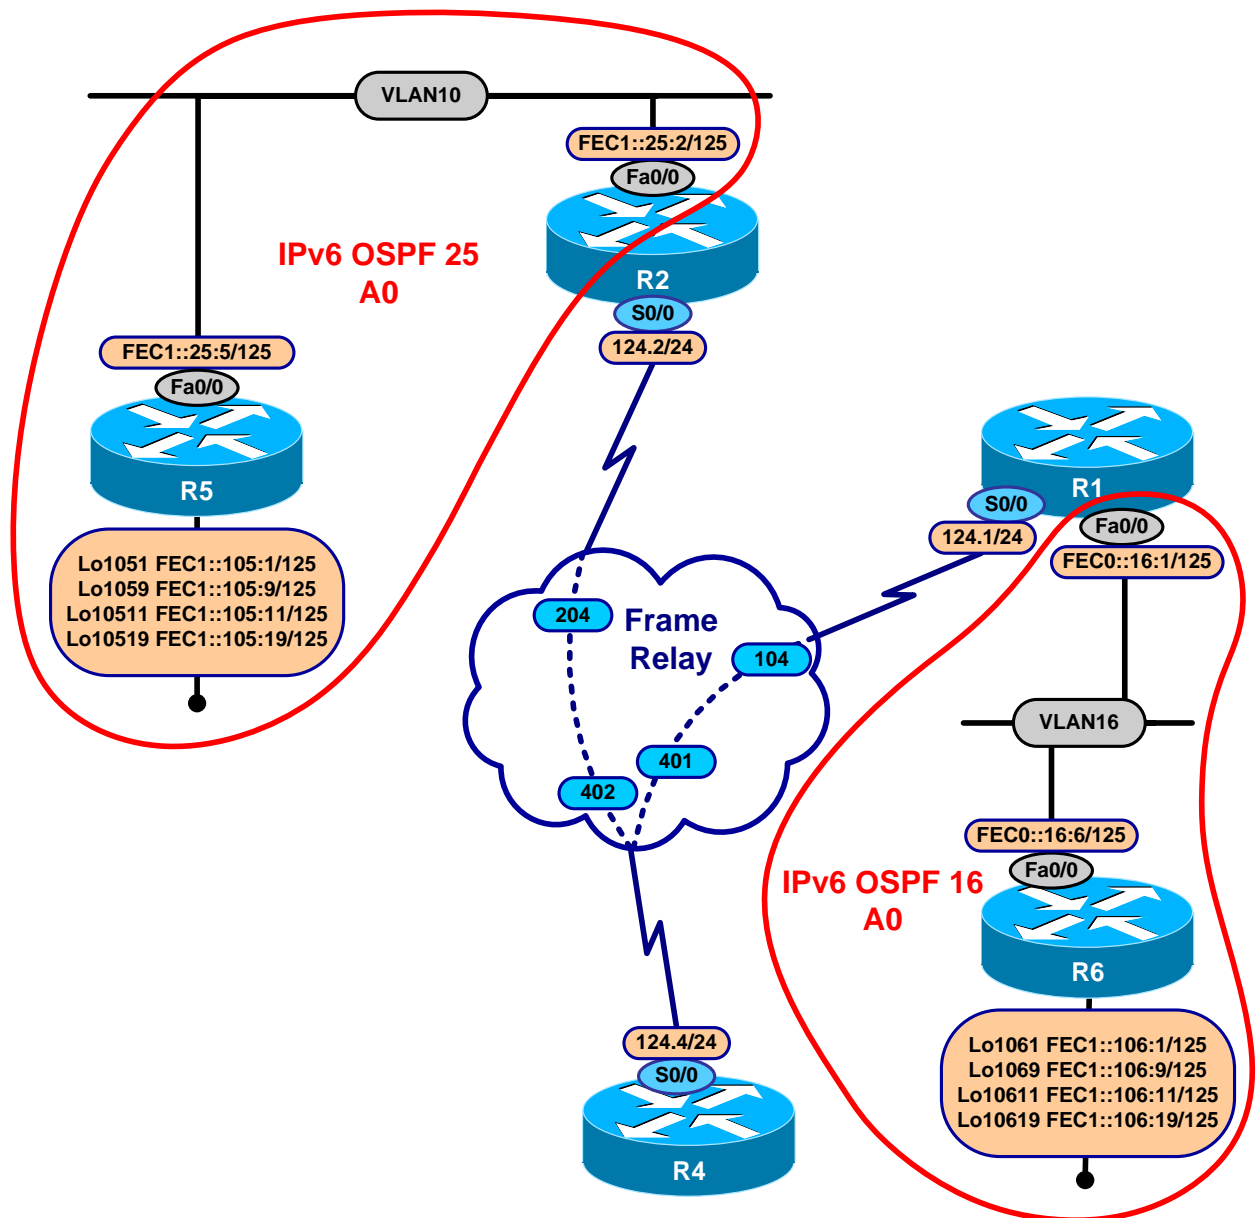
Netmasterclass 200v6 POD Layout



Netmasterclass DOiT-200v6 Scenario 8 IPv4 IGP diagram



Netmasterclass DOIT-200v6 Scenario 8 IPv6 IGP diagram



8.1 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.

- 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



Shutdown interfaces FA0/15 on both switches CAT1 and CAT2

- 8.2.1 Configure VLANs according to the diagram and the "VLAN configuration" table.

VLAN Table

Router	Interface	VLAN
R1	FastEthernet0/0	VLAN12 VLAN16
R2	FastEthernet0/0	VLAN10 VLAN12
R3	FastEthernet0/1	
R3	FastEthernet0/0	VLAN20
R4	FastEthernet0/0	VLAN30
R5	FastEthernet0/0	VLAN10
R6	FastEthernet0/0	VLAN16 VLAN30
FRS	Ethernet0	VLAN10
CAT1	-	VLAN20 VLAN40
CAT2	-	VLAN30 VLAN40

- 8.2.2 Use ISL for all trunking needs.
- 8.2.3 Make sure you allow only VLANs 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 172.16.1.20/24 on the Fa0/14 interface of CAT2.

8.3 Frame Relay QoS

- 8.3.1 On router R4 guarantee that future IPv6 transit traffic generated by IPv6 BGP task through router R4 will have guaranteed 25% of overall bandwidth outbound for PVCs that it is going to use. Apply only on necessary PVCs.

8.4 OSPF

- 8.4.1 Configure area 0 between R1, R2 and R4.
- 8.4.2 Configure area 10 between R2 and FRS.
- 8.4.3 Configure area 12 on VLAN12 between R1 and R2.
- 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 domain.
- 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.
- 8.5.7 Configure OSPF to not to establish extra adjacencies on VLAN10 with unnecessary neighbors.

8.6 EIGRP

- 8.6.1 Configure EIGRP AS 10 on the VLAN16 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 the EIGRP domain.
- 8.6.4 Configure EIGRP AS 35 between FRS and R5 on VLAN10. R2 is not part of this AS. Extend this AS to the link between R5 and R3.
- 8.6.5 Advertise 10.10.10.0/24 on R5 into EIGRP as native.
- 8.6.6 Do not allow EIGRP routing updates to go beyond necessary devices on VLAN10.



TO OBTAIN UNIVERSAL CONNECTIVITY, PERFORM A MUTUAL REDISTRIBUTION OF DYNAMIC INTERIOR GATEWAY PROTOCOLS: BETWEEN EIGRP 35 AND OSPF ON FRS, BETWEEN RIP, OSPF AND EIGRP 10 ON R4. DO NOT PERFORM ANY OTHER REDISTRIBUTION IN THIS SCENARIO. PERFORM REDISTRIBUTE CONNECTED WHERE REQUIRED AND NOT RESTRICTED BY THE SCENARIO.

8.7 BGP

- 8.7.1 Configure AS 1000 on CAT1 and AS 2000 on CAT2.
- 8.7.2 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.
- 8.7.3 Configure BGP AS 100 on routers R2, R3, R4, R5 and FRS.
- 8.7.4 In AS 100, restrict your peering relationships to the following: R3 and R5, R4 and R2, FRS and R5, FRS and R2.
- 8.7.5 Peer AS 100 and AS 1000 between R3 and CAT1.
- 8.7.6 Peer AS 100 and AS 2000 between R4 and CAT2.
- 8.7.7 R2 must not accept any NLRI transiting through R5. R5 must not accept any NLRI transiting through R2. Do not use any filtering techniques.
- 8.7.8 Make sure you provide transit connectivity between 100.100.100.0/24 and 200.200.200/24 across AS 100. Make sure that 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.
- 8.7.9 Configure the minimum number of IBGP peer relationships to fulfill the criteria above. Do not introduce new AS numbers.

8.8 VPN

- 8.8.1 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.
- 8.8.2 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.
- 8.8.3 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.
- 8.8.4 Do not use any generic encapsulation. Use IP into IP encapsulation.

8.9 Router Maintenance

- 8.9.1 Configure the following access characteristics for the aux line on all routers and the VTY line number 5 for the Catalyst 3550 switches:
 - o The device should supply the enable prompt whenever it is accessed
 - o The device should not prompt for a username and password
 - o The device should not terminate a session after the expiration of some period of time
- 8.9.2 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.

8.10 IP Features

- 8.10.1 On router R3 configure solution recommended to improve TCP performance in network paths with large bandwidth, long-delay characteristics that are called Long Fat Networks (LFNs).

8.11 IPv6 IGP

- 8.11.1 Configure two separated IPv6 networks, first between R2 and R5, second between R1 and R6.
- 8.11.2 Create loopback interfaces according to the IPv6 diagram.
- 8.11.3 Configure OSPFv3 area 25 on the network between R2 and R5. Make it area 0.
- 8.11.4 Configure OSPFv3 area 16 on the network between R1 and R6. Make it area 0.

8.12 IPv6 BGP

- 8.12.1 Configure IPv6 BGP AS 100 on R2.
- 8.12.2 Configure IPv6 BGP AS 16 on R1.
- 8.12.3 Advertise loopbacks and connected OSPF networks into these BGP processes.
- 8.12.4 Connect BGP AS over link level address. Do not configure IPv6 on any adapter configured with IPv4 address for this interconnection. Do not use EBGP multihop. Solution must utilize Frame Relay physical layer only.
- 8.12.5 Summarize both AS links to /96 prefix and advertise only this prefix to the other peer. Use BGP configuration to perform summarization.
- 8.12.6 Make sure all IPv6 routers can reach any subnet in the other IPv6 network.

8.13 QOS

- 8.13.1 Provide a QOS solution to the VPN users allowing traffic between R2 and R3 to be shaped to 30Kbit per second.
- 8.13.2 Let the selected algorithm calculate the sustained bits per interval and the excess bits permitted in the first interval.

8.14 IOS Specialties

- 8.14.1 Configure a user "user" on R1, use password "cisco".
- 8.14.2 Configure router R1 to block login for 120 seconds after 5 failed attempts to login within 120 seconds.
- 8.14.3 Report each successful and unsuccessful login to log.
- 8.14.4 Allow telnet connection from Lo106 of R6 to pass at all times.

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 Shared Tree 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.

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
  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

8.1 Frame Relay

No hint.

8.2 Catalyst configuration

No hint.

8.3 Frame Relay QoS

Enable per-PVC traffic shaping. Read IPv6 BGP part ahead.

8.4 OSPF

No hint.

8.5 RIP

No hint.

8.6 EIGRP

No hint.

8.7 BGP

No hint.

8.8 VPN

No hint.

8.9 Router Maintenance

No hint.

8.10 IP Features

No hint.

8.11 IPv6 IGP

No hint.

8.12 IPv6 BGP

Check for feedbacks.

8.13 QoS

No hint.

8.14 IOS Specialties

Check 12.4 features.

8.15 Catalyst specialties

No hint.

8.16 Address administration

No hint.

8.17 Multicast

No hint.