


NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6

VOLUME II



DOiT-200v6 Lab 14

Multi-Topic CCIE-Level Scenario

For CCIE® Candidates

DOiT volume II lab14 SC

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

14.1 Frame Relay	8
14.2 Catalyst Configuration	8
14.3 OSPF	8
14.4 RIP	9
14.5 EIGRP	9
14.6 BGP	9
14.7 Router Maintenance	10
14.8 Security	10
14.9 IPv6	10
14.10 Catalyst Specialties	11
14.11 Address Administration	11
14.12 Multicast	11
14.13 Quality of Service	12
14.14 IOS Features	12
Appendices	13
Hints	15
14.1 Frame-Relay	15
14.2 Catalyst Configuration.....	15
14.3 OSPF	15
14.4 RIP	15
14.5 EIGRP	15
14.6 BGP.....	15
14.7 Router Maintenance.....	15
14.8 Security	15
14.9 IPv6.....	15
14.10 Catalyst Specialties.....	15
14.11 Address Administration	16
14.12 Multicast.....	16
14.13 Quality of Service.....	16
14.14 IOS Features.....	16

DOiT volume II lab14 SC




Goals and Restrictions

- IP subnets on the diagram belong to network 172.16.0.0/16.
- Do not rely on dynamic Frame-Relay Inverse ARP.
- Do not introduce any new IP addresses.
- Do not use any static routes.
- Make sure all IPv4 and IPv6 Loopback interfaces are advertised with their original mask, unless noted otherwise.
- Make sure all IP interfaces in the diagram are reachable within this internetwork.
- Use conventional routing algorithms.
- 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.*.* as well as the 172.16.1.2 address on FRS.



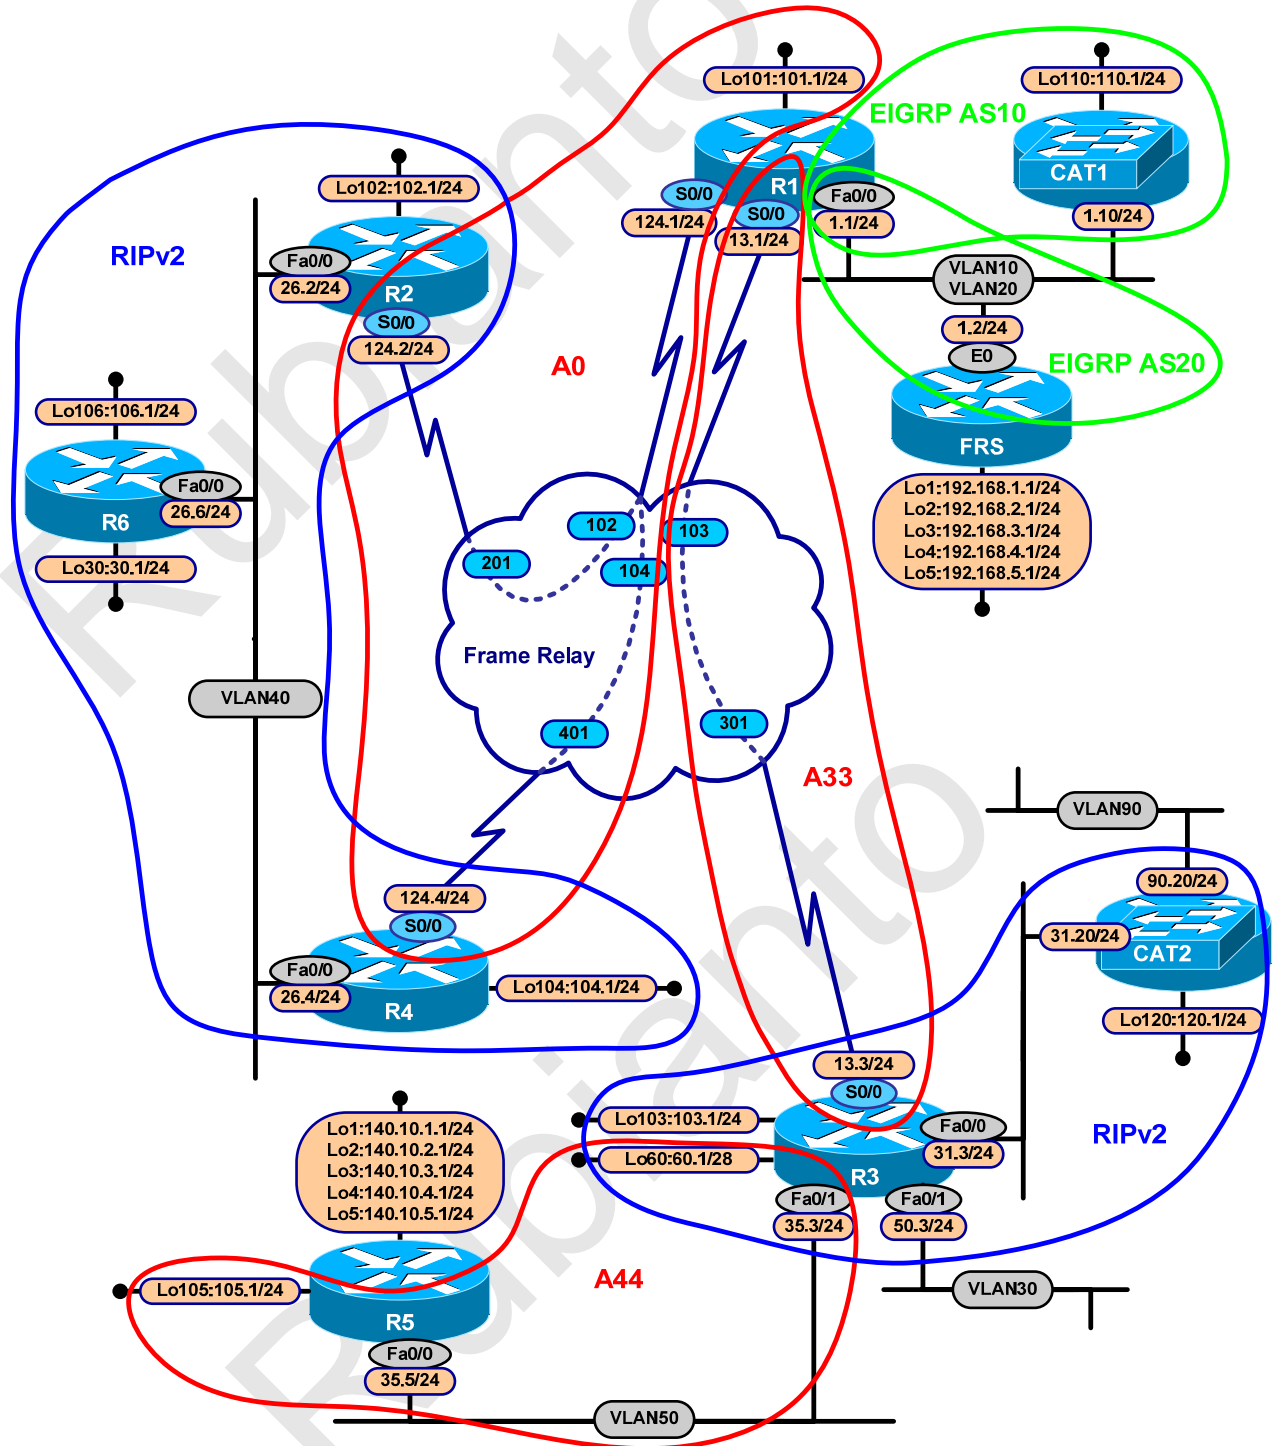
FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM "NMC POD LAYOUT" DISPLAYED IN THE APPENDIX A AT THE END OF THIS DOCUMENT.

Netmasterclass 200v6 POD Layout



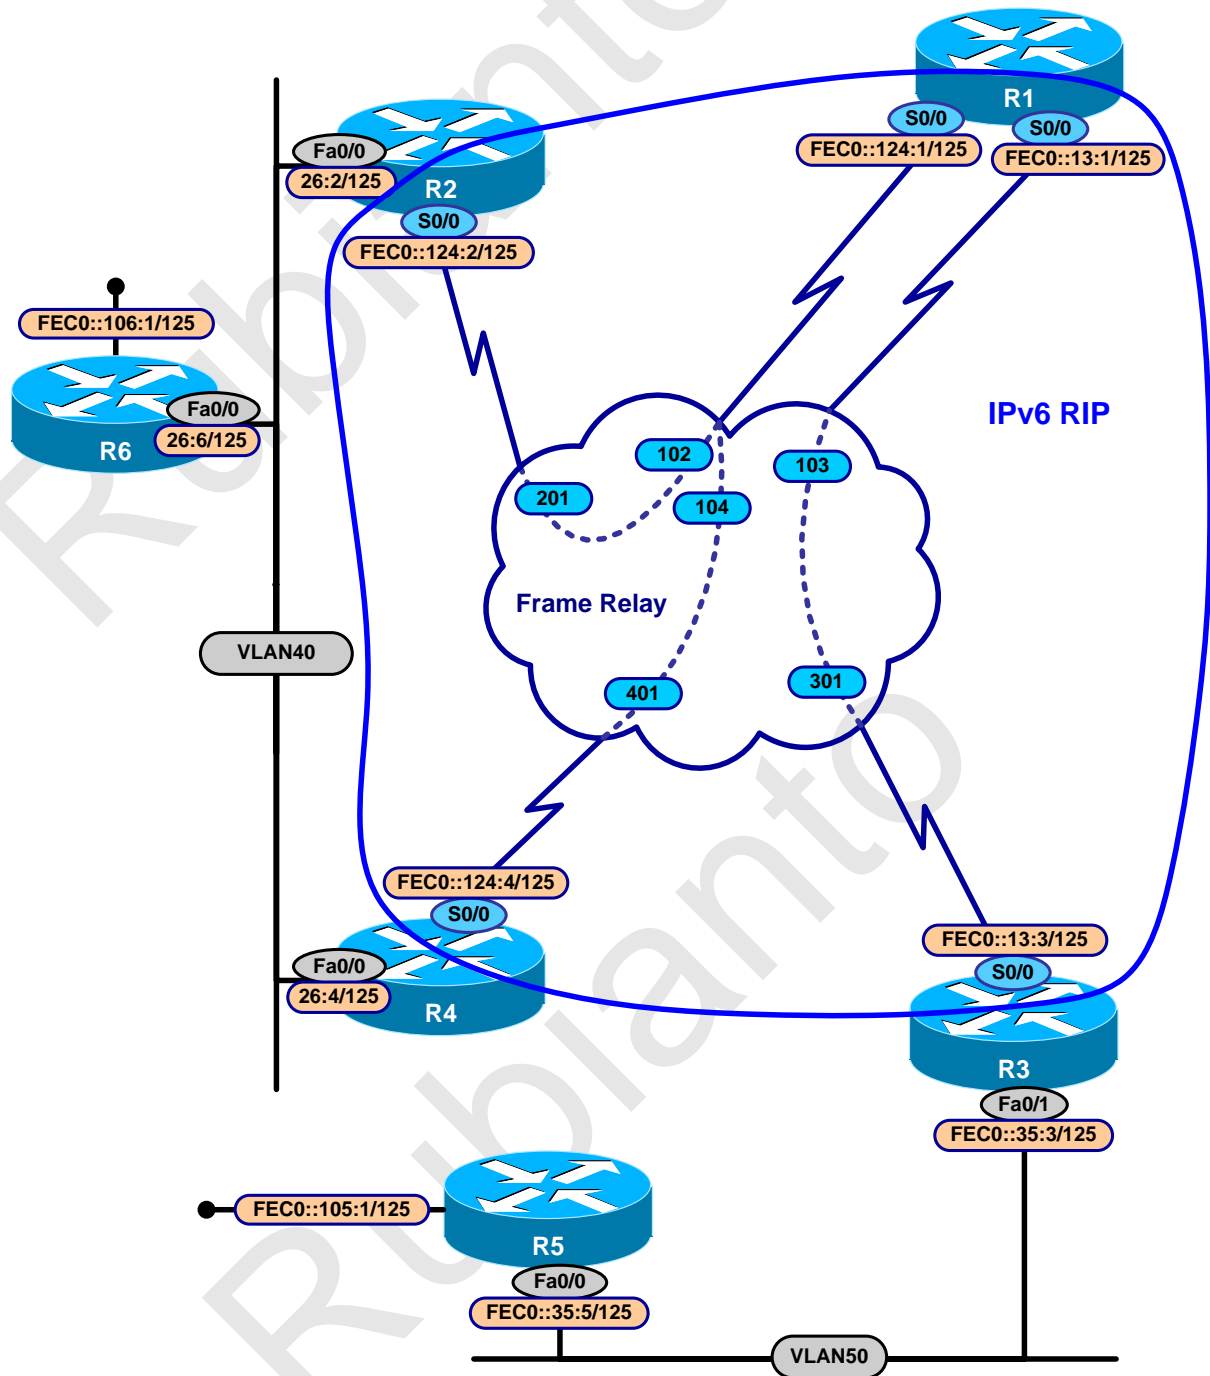
DoIT volume II lab14 SC

Netmasterclass DOiT-200v6 Scenario 14 IPv4 IGP diagram



DOiT volume II lab14 SC

Netmasterclass DOiT-200v6 Scenario 14 IPv6 IGP diagram



DOiT volume II lab14 SC

14.1 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.

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



Shutdown interfaces FA0/15 on both switches CAT1 and CAT2

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

VLAN Configuration Table

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

14.3 OSPF

14.3.1 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.

14.3.2 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 as well as VLAN 35 and to OSPF

area 44. Assign the 172.16.105.0/24 subnet on R5 to OSPF area 44. Also, make R3 the Designated router for the R1/R3 link.

14.3.3 Use the OSPF network type non-broadcast for all Frame Relay connections.

14.3.4 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.

14.4 RIP

14.4.1 On R6, configure RIP v.2. Advertise the loopbacks on R6 to R2 and R4 via RIP, do not broadcast/multicast RIP updates on VLAN 40 for security reasons.

14.4.2 On R2 and R4, configure mutual redistribution between OSPF and RIP. Use the minimum non-zero metric in redistribution.

14.4.3 Configure RIP between R3 FastEthernet 0/0 and CAT2. Send updates to 224.0.0.9 only on the link between R3 and CAT2.

14.5 EIGRP

14.5.1 On R1 and FRS, configure EIGRP AS 20 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.

14.5.2 Make sure that EIGRP routes allowed on R1 are seen by all other routers.

14.5.3 Configure EIGRP 10 between CAT1 and R1 and advertise network 172.16.110.0/24 from CAT1.

14.6 BGP

14.6.1 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.

14.6.2 Configure IBGP between R3 and R1. Make sure that R1 has networks from AS500 installed in its routing table.

14.6.3 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.

- 14.6.4 Configure BGP so that you are sending an aggregate for 172.0.0.0/8 to AS500 and suppressing all other routes.
- 14.6.5 Make sure R4 sees only the following AS 500 originated BGP update, 140.10.2.0/24, in its IPv4 routing table; do not introduce modifications in R4 BGP table to meet this requirement.

14.7 Router Maintenance

- 14.7.1 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.
- 14.7.2 The new router will have IP address 172.16.26.100/24 and MAC address 0010.7be8.131d.
- 14.7.3 The new router should send the request for the configuration R100.cfg via R2.

14.8 Security

- 14.8.1 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
- 14.8.2 Develop a sample configuration that achieves these requirements and apply it on R5

14.9 IPv6

- 14.9.1 Configure subnet FEC0::124:0/125 on Frame-Relay links between R1, R2 and R4.
- 14.9.2 Configure subnet FEC0::13:0/125 on Frame-Relay link between R1 and R3.
- 14.9.3 Configure subnet FEC0::26:0/125 on VLAN40.
- 14.9.4 Configure subnet FEC0::35:0/125 on VLAN50.
- 14.9.5 Configure loopbacks FEC0::105:1/125 on R5 and FEC0::106:1/125 on R6.
- 14.9.6 Configure RIP for IPv6 between R1, R2, R3 and R4. Only Frame-Relay networks must be included in RIP process.
- 14.9.7 Configure BGP AS 100 for IPv6 family on R1, R3 and R4.

- 14.9.8 Configure BGP AS 200 for IPv6 family on R2.
- 14.9.9 Configure BGP AS 600 for IPv6 family on R6.
- 14.9.10 Configure BGP AS 500 for IPv6 family on R5.
- 14.9.11 Connect AS 600 to AS 200 between R2 and R6, connect AS 600 and AS 100 between R6 and R4, connect AS 100 and AS 500 between R3 and R5.
- 14.9.12 AS 100 IBGP connections must remain the same with IPv4 BGP task.
- 14.9.13 Advertise FEC0::105:0/125 and FEC0::106:0/125 into BGP as internal.
- 14.9.14 Make sure traffic from FEC0::105:1 to FEC0::106:1 goes thru R3, R1, R2 to R6, and return traffic goes thru R4, R1, R3 to R5. Apply configuration on R6. Use one well-known mandatory attribute. Do not change origin.

14.10 Catalyst Specialties

- 14.10.1 CAT2 will be connected to the test environment network 172.16.90.0/24 in the future. Interface FastEthernet0/17 will be used for this purpose.
- 14.10.2 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.
- 14.10.3 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.

14.11 Address Administration

- 14.11.1 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.
- 14.11.2 Make sure that CAT1 can ping all interfaces in your network. However, allow TELNET access only from the R3's management loopback interface.

14.12 Multicast

- 14.12.1 Configure Multicast Routing on R1, R2, R3 and R4 using R1 as the shared root.
- 14.12.2 Announce the shared root without use of any dense groups or static configurations.

14.12.3 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

14.12.4 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.

14.13 Quality of Service

14.13.1 Traffic originating from the 172.16.35.5 address of R5 with an IP precedence setting of either 3 or 4 must not consume more than 50% of the bandwidth of the Frame-Relay interface on R3.

14.13.2 If traffic that meets the classification stated above attempts to exceed this bandwidth consumption threshold, buffer the traffic. The traffic with an IP precedence of 4 should get twice as much bandwidth as traffic marked with IP precedence 3. The maximum bandwidth allocation for a single class of traffic must be 16000 bps.

14.13.3 Use a technique that provides the ability to consume additional bandwidth above the stated threshold only after periods of inactivity.

14.13.4 As a bonus for periods of inactivity, allow this classification of traffic to transmit an additional 12,000 bits. Transmit an equal value each and every interval as well.

14.13.5 Apply the QoS mechanism directly to the Frame-Relay interface on R3. Apply a QoS mechanism that can be applied to any type of physical interface. Also, do not use any "percentage" keyword when fulfilling these configuration requirements of this section.

14.14 IOS Features

14.14.1 Simulate traffic every 2 minutes by sending 100 test packets and monitor the jitter of this simulated traffic from R5 to R1. Provide a minimal configuration on R1.

14.14.2 Provide two separate sets of simulated traffic from R5 to R1 with one set of traffic marked as precedence 4 and the second set of traffic marked as precedence 3.

14.14.3 For any type of IP Service, that needs to be configured for the task above, make sure the R1 is the server of any supporting service.

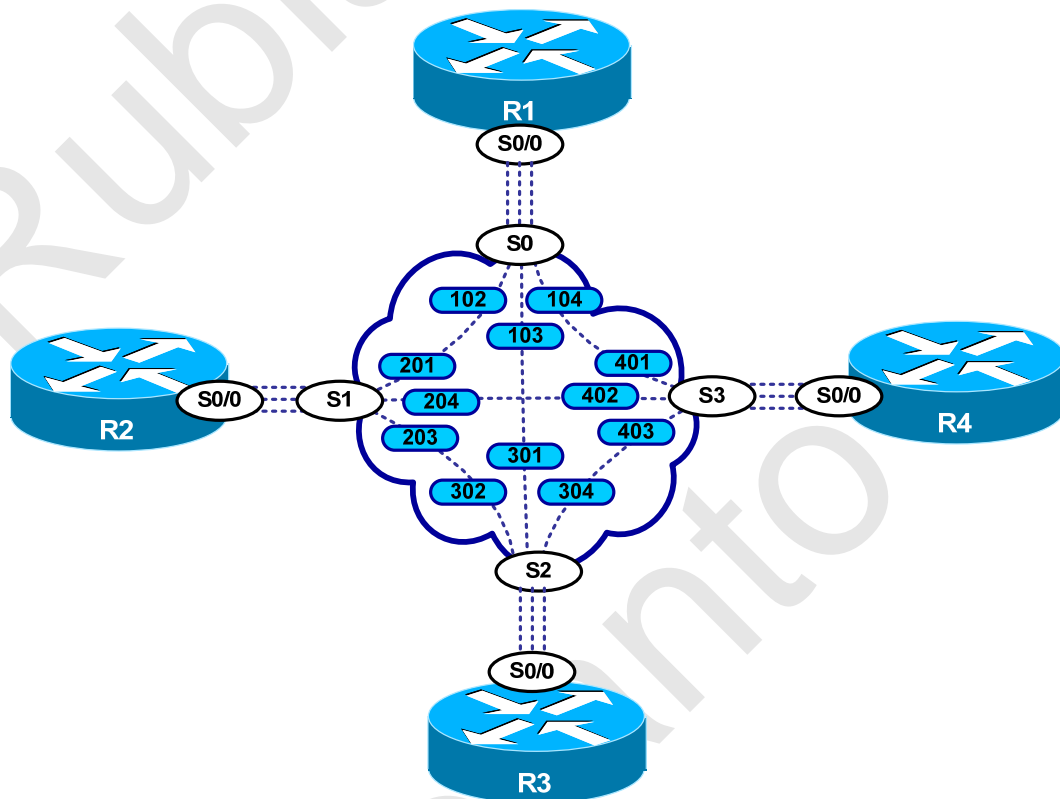
14.14.4 Provide a mechanism that can be used on R1 to count the number of packets received by the router on its point-to-point Frame-Relay subinterface. Packet counts should be classified by the received packets' precedence setting. Do not use any access-lists to fulfill this requirement.

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
  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
```

Hints

14.1 Frame-Relay

No hint.

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.

14.3 OSPF

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.

14.4 RIP

See OSPF section.

14.5 EIGRP

No hint.

14.6 BGP

When configuring BGP, pay attention to what IGP's are or ARE NOT running over the links where BGP neighbor relationships are formed.

14.7 Router Maintenance

No hint.

14.8 Security

No hint.

14.9 IPv6

No hint.

14.10 Catalyst Specialties

No hint.

14.11 Address Administration

No hint.

14.12 Multicast

No hint.

14.13 Quality of Service

No hint.

14.14 IOS Features

No hint.

DOIT volume II lab14 SC

Rubianto