Networking
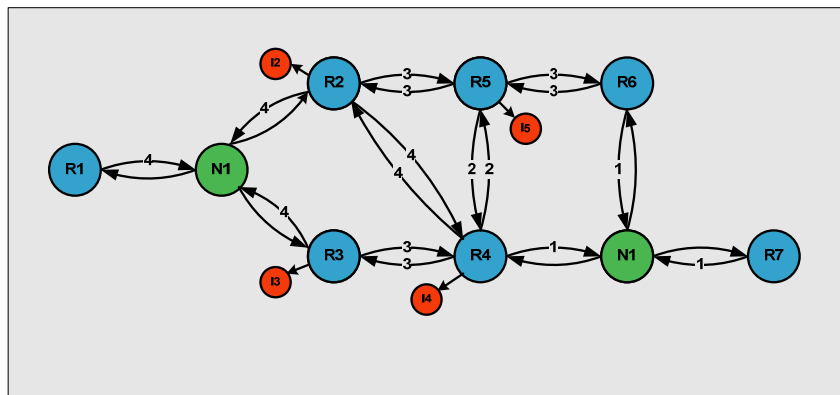Consulting
Training
www.netmasterclass.com

Cisco
Networking

NETMASTERCLASS
ROUTING AND SWITCHING CCIE® TRACK

DOiT-200v6
VOLUME II

DOiT-200v6 Lab 2
Multi-Topic CCIE-Level Scenario
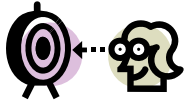
For CCIE® Candidates

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

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 with their original mask.
- Network 0.0.0.0/0 should not appear in any routing table (show ip route).
- Do not use the "ip default-network" command.
- 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.
- Use conventional routing algorithms.

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

Netmasterclass 200v6 POD Layout
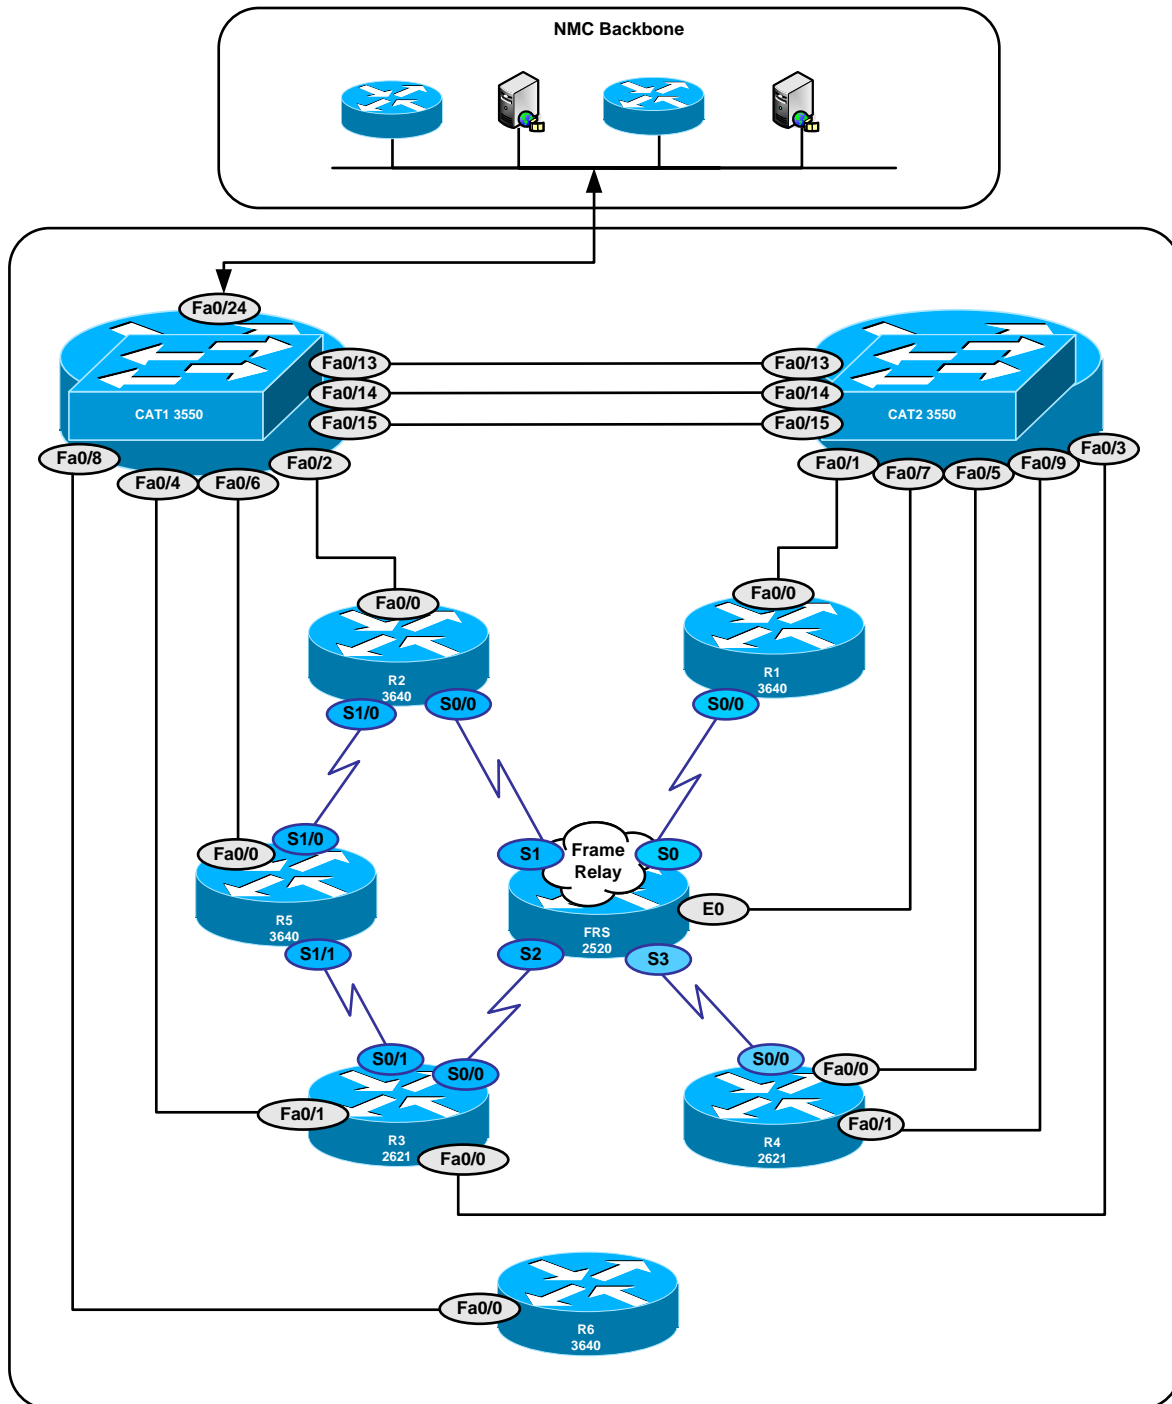
Netmasterclass DOiT-200v6 Scenario 2 IPv4 IGP diagram

Networking
Consulting
Training
www.netmasterclass.com
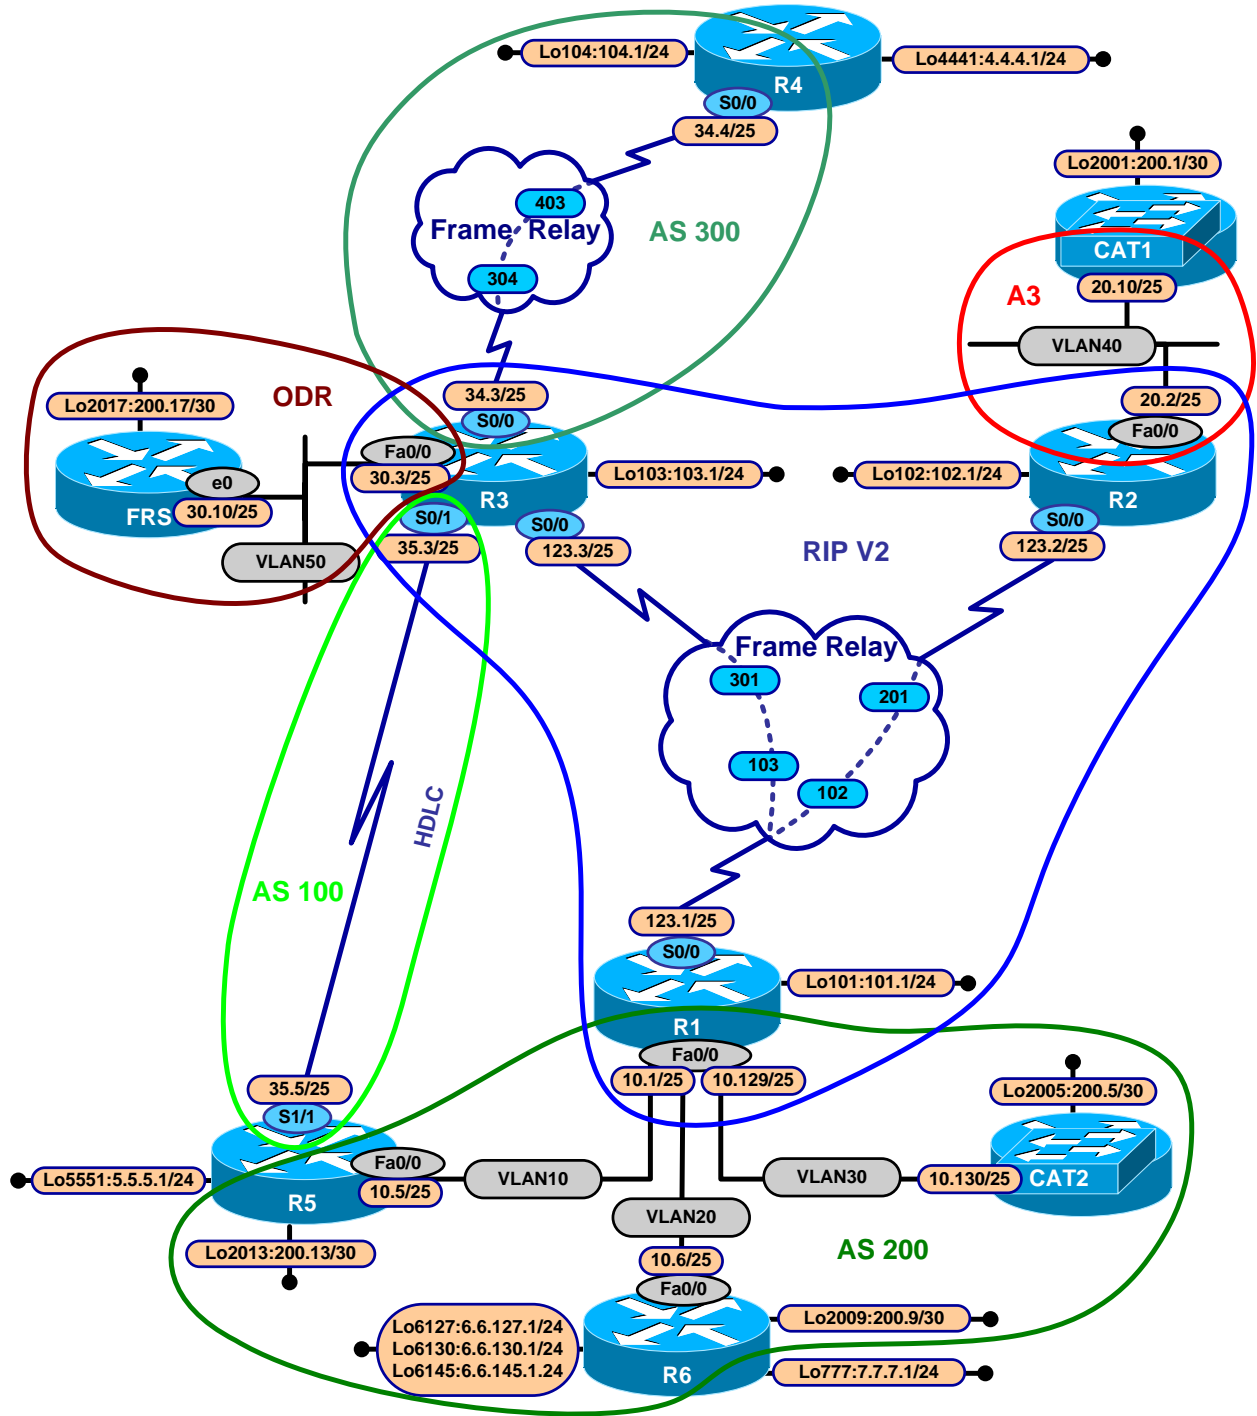
Netmasterclass DOiT-200v6 Scenario 2 IPv6 IGP diagram

RIP FRAME-34

S0/0
FE80::4
FECO::34:4/125
R4

Lo604 FECO::104:1/125

RIP FRAME-34

403

Frame Relay

304

FE80::3
FECO::34:3/125
S0/0

Lo603 FECO::103:1/125

RIP FRAME-34

R3

RIP FRAME-123

Lo602 FECO::102:1/125

S0/0
FECO::123:3/125
FE80::3

RIP FRAME-123

R2

S0/0
FECO::123:2/125
FE80::2

Frame Relay

301

201

103

102

FE80::1
FECO::123:1/125
S0/0

IPv6 OSPF
A0

Lo601 FECO::101:1/125

IPv6 OSPF
A0

Fa0/0
FECO::122:1/125
VLAN30

FECO::15:1/122

R1

Fa0/0
10.1/25

FECO::16:1/122

IPv6 OSPF
A0

TUNNEL

FECO::15:5/122

TUNNEL

IPv6 OSPF
A56

FECO::16:6/122

Fa0/0
10.5/25
FECO::156:5/125

VLAN10

R5

Lo605 FECO::105:1/125

IPv6 OSPF
A0

IPv6 OSPF
A56

VLAN20

Fa0/0
10.6/25

R6

Lo610 2001:0:0:6::/64
Lo611 2001:0:0:7::/64
Lo612 2001:0:0:8::/64
Lo613 2001:0:0:9::/64
Lo614 2001:0:0:10::/64

Lo606 FECO::106:1/125

2.1 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.

2.1.1 Use only the PVCs displayed in the "Netmasterclass DOiT-200v6 Scenario 2 Pv4 IGP diagram". No dynamic entries are allowed in the Frame Relay map tables.

2.1.2 Use different type logical interfaces on R3 and a physical interface on R4 on the Frame Relay subnets.

2.1.3 Use a logical interface on R1 and a physical interface on R2 on the Frame Relay subnets.

2.2 Catalyst Configuration

Shutdown interfaces FA0/15 on both switches CAT1 and CAT2

2.2.1 Configure the VLANs for all Ethernet segments involved in this scenario. See the diagram "Netmasterclass DOiT-200v6 Scenario 2 Pv4 IGP diagram" and the tables provided below.

VLAN Tables

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

VLAN	VLAN NAME
VLAN10	LANA
VLAN20	LANB
VLAN30	LANC
VLAN40	LAND
VLAN50	LANE

2.2.2 Advertise VLAN's from CAT1 to CAT2. Make sure CAT2 does not advertise any VLANs.

2.2.3 Make sure ports Fa0/13 and Fa0/14 are not in a blocking state for any VLAN.

2.2.4 Use the dot1q protocol for all trunks involved in this scenario.

2.2.5 Configure domain name as doitlab2.

2.3 OSPF

2.3.1 Configure OSPF area 3 on VLAN40 between R2 and CAT1.

2.3.2 The devices connected to VLAN40 should become fully adjacent only if they pass non-clear text authentication using the string "nmc" on a per-interface basis.

2.3.3 Advertise loopback interface 172.16.200.1/30 into OSPF without any reference to an OSPF area.

2.4 RIP

2.4.1 Configure RIP version 2 between R1, R2 and R3. Unicast RIP updates on the Frame Relay network only between the router R1 and R3 and the routers R1 and R2.

2.4.2 Exchange RIP updates on Frame Relay link only.

2.5 EIGRP

2.5.1 Configure EIGRP AS100 on the HDLC link between R3 and R5.

2.5.2 Configure EIGRP AS200 between routers R1, R5, R6 and R1, CAT2. Use a unicast EIGRP packet exchange between these routers.

2.5.3 Advertise the following loopback networks on R6 to AS200 without using a network statement:

- o 6.6.127.0/24
- o 6.6.130.0/24
- o 6.6.145.0/24

2.5.4 Summarize these subnets on R6 with the most optimal mask.

2.5.5 Advertise the following loopback networks into EIGRP AS 200 as internal EIGRP prefixes:
- o 172.16.200.5/30
- o 172.16.200.9/30
- o 172.16.200.13/30

2.5.6   Configure EIGRP AS 300 between the routers R3 and R4 and advertise the loopback 104 in AS 300.

2.6 ODR

2.6.1   Configure on demand routing between R3 and FRS.

2.6.2   FRS should be configured as "no ip routing" device and should be able to ping all IP addresses from the 172.16.0.0/16 address space.

> **TO OBTAIN UNIVERSAL CONNECTIVITY, PERFORM A MUTUAL REDISTRIBUTION OF DYNAMIC INTERIOR GATEWAY PROTCOLS: BETWEEN RIP AND OSPF ON ROUTER R2, BETWEEN EIGRP AS 100, EIGRP AS 300, RIP AND ODR ON R3, BETWEEN RIP AND EIGRP ON R3, BETWEEN RIP AND EIGRP 200 ON R1. DO NOT PERFORM ANY OTHER REDISTRIBUTION IN THIS SCENARIO. PERFORM REDISTRIBUTE CONNECTED WHERE REQUIRED AND NOT RESRICTED BY THE SCENARIO.**

2.7 BGP

2.7.1   Configure AS 400 on R4.

2.7.2   Configure AS 100 on R3 and R1.

2.7.3   Configure AS 500 on R5.

2.7.4   Configure AS 600 on R6.

2.7.5   Peer AS 400 with AS100 between R4 and R3.

2.7.6   Peer AS 500 and AS 100 between R5 and R1.

2.7.7   Peer AS 600 and AS 100 between R1 and R6.

2.7.8   Advertise Loopback network 4.4.4.0/24 from R4, Loopback network 5.5.5.0/24 from R5 and Loopback network 7.7.7.0/24 from R6.

2.7.9   AS 100 should provide transit service to network 7.7.7.0/24 from 4.4.4.0/24 only if AS100 loses reachability to network 5.5.5.0

2.7.10 Turn synchronization off in this section.

2.8 IPv6 Addressing

2.8.1   Configure subnet FEC0::123:0/125 on Frame-Relay links between R1, R2 and R3. Use same interfaces that you have chosen for IPv4 task. Do not use any extra DLCIs above those used for IPv4 configuration.

2.8.2   Configure subnet FEC0::34:0/125 on Frame-Relay link between R3 and R4. Use same interfaces that you have chosen for IPv4 task. Do not use any extra DLCIs about those used for IPv4 configuration.

2.8.3   Configure link-local IPv6 addresses on Frame Relay links `FE80::R`, where `R` is the router number.

2.8.4   Configure IPv6 address FEC0::122:1/125 on R1 interface that belongs to Vlan 30.

2.8.5   Configure interfaces Loopback60X with IPv6 address FEC0::10X:1/126, where X is number of the router (1 for R1, 2 for R2, …, 6 for R6). Do not configure loopbacks on FRS and CAT switches.

2.8.6   On R6 configure interfaces Loopback610 – Loopback614 using IPv6 addresses 2001:0:0:6::/64 - 2001:0:0:10::/64, using an EUI-64 interface ID in the low-order 64 bits of the address.

2.9 IPv6 Routing

2.9.1   Configure IPv6 RIP FRAME-34 on subnet FEC0::34:0/125. Add loopback 603 and 604 IPv6 subnets to IPv6 RIP FRAME-34.

2.9.2   Configure IPv6 RIP FRAME-123 on subnet FEC0::123:0/125. Add loopback 602 IPv6 subnet to IPv6 RIP FRAME-123.

2.9.3   Configure IPv6 OSPF area 0 on FastEthernet subinterface of R1 that belongs to Vlan 30, on R1s Loopback601, on R5s Loopback605

2.9.4   Configure IPv6 OSPF area 0 on the link between R1 and R5. Use IPv6 addresses FEC0::15:1/112 and FEC0::15:5/112 between R1 and R5; encapsulate IPv6 packets within IPv4 packets for a transmission across the IPv4 link between R1 and R5.

2.9.5   Configure IPv6 OSPF area 56 on link between R1 and R6. Use IPv6 addresses FEC0::16:1/112 and FEC0::16:6/112 between R1 and R6; encapsulate IPv6 packets within IPv4 GRE packets for a transmission across the IPv4 link between R1 and R6. Do not allow Type 5 LSA's into area 56.

2.9.6   Configure IPv6 OSPF area 56 on Fa0/0 interface of R5.

2.9.7   Advertise Loopbacks from R6 into OSPF as non-OSPF prefixes. Make sure that every IPv6 router except R6 has only the best summary for globally aggregatable addresses from R6.

2.9.8 Mutually redistribute IPV6 RIP processes on R3.

2.9.9 Mutually redistribute IPv6 RIP and IPv6OSPF on R1.

2.10 Router Maintenance

2.10.1 There are two diskless SUN workstations WS1 and WS2 on VLAN 40

2.10.2 The SUN Server, where the workstations can boot their OS from, is located on VLAN 50. Its IP address is 172.16.30.100.

2.10.3 Provide communications between WS1 (IP address 172.16.20.100, MAC 8:0:20:ac:24:b8), WS2 (IP address 172.16.20.101, MAC 8:0:20:ac:24:b9) connected to CAT1 ports FA0/11 and FA0/12 respectively and the SUN server.

2.10.4 Provide the solution allowing the workstations to boot their OS which does not require a DHCP server. Apply the configuration on R2

2.11 Security

2.11.1 Allow the Finger application to R4 from R1, R2 and R5. Disallow Finger connections from R4 to routers R1, R2 and R5. Do not use the "established" keyword.

2.12 QOS

2.12.1 Assign DSCP value 45 to frames received by CAT1 on ports Fa0/11 and Fa0/12.

2.13 Catalyst Specialties

2.13.1 CAT1 ports Fa0/11 and Fa0/12 are patched to the Conference Room.

2.13.2 Only the two workstations described in the Router Maintenance section can be used in the Conference Room.

2.13.3 Also, two Cisco 7960 IP Phones with the MAC addresses of 00-07-85-95-D1-A7 and 00-07-85-95-D2-B7 are connected to the CAT1 ports Fa0/11and Fa0/12 respectively.

2.13.4 Configure Priority Tagged Frames to carry voice traffic on the Native VLAN.

2.13.5 Override the Class of Service value in the packets generated by the SUN workstations with a value of 4.

2.14 Gateway Redundancy

2.14.1 Assign 172.16.10.126 to a virtual gateway and make sure the MAC address associated with the virtual gateway is set to 0000.0c07.ac1a.

2.14.2 Authenticate HSRP on the 172.16.10.0/25 subnet (password nmc). Exchange Hello packets 3 times faster than by default.

2.14.3 Select R1 as the preferred gateway by using priority 150.

2.14.4 Select R5 as the preferred gateway if the Frame Relay connection on R1 goes down.

2.14.5 R1 should take responsibility back after the Frame Relay connection comes back up.

2.15 Multicast

2.15.1 Simulate the multicast stream (172.16.34.4, 232.2.2.2) from R4.

2.15.2 Statically configure a Shared Tree between R3, R5, R1, R6 and R2 rooted from R5.

2.15.3 Join group 232.2.2.2 on every member of the Shared Tree by using one of the router interfaces.

2.15.4 Make sure you get a reply from all joined routers on 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)
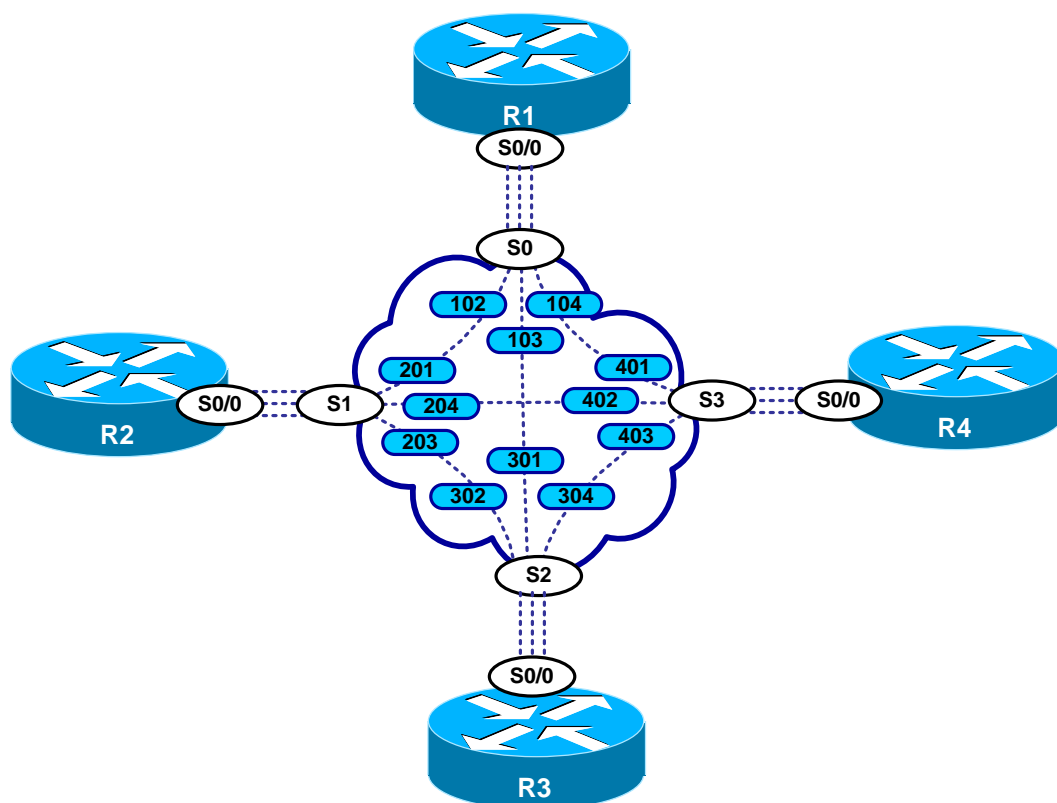
Provided below is the Frame-Relay switch configuration. If you enter this configuration into a router acting as a Frame-Switch and you cable up your pod in the manner displayed on this page, you will have configured a full-mesh Frame-Relay topology.

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
 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

2.1   Frame-Relay

No hint

2.2   Catalyst Configuration

When considering how to prevent any of the fa0/13 and fa0/14 ports from attaining a Spanning Tree blocking state, remember that these two ports interconnect CAT1 to CAT2. Therefore, they form a loop between CAT1 and CAT2. Think about how to make these two ports appear to both Catalysts as one single port. Think of techniques to group individual Catalyst ports so that they act as one aggregated port.

2.3   OSPF

There seems to be no OSPF Area 0 in this exercise. Can you configure OSPF without an Area 0? Yes, provided that there are no other areas.

OSPF offers two authentication methods: one based on passwords passed in clear-text and the second based on the MD5 hashing algorithm.

If you attempt to inject a prefix into an OSPF without associating the prefix to an OSPF area, you need to make that prefix an external OSPF LSA.

2.4   RIP

A good general practice to perform with RIP is to configure passive-interface default under the RIP routing process. Then apply the no passive-interface command to specify the interfaces you want RIP updates to go out on. This will assure that RIP updates get exchanged on the interfaces you desire.

2.5   EIGRP

The EIGRP domain involves three different VLAN's (VLAN's 10, 20 and 30) are configured with the same IP address (172.16.10.0/24). What options are available to make three separate VLAN's appear as a single subnet? Perhaps you should consider some type of bridging technique.

When selecting the most optimal mask for a set of addresses to summarize, pay very close attention to the bit boundaries the supplied prefixes traverse.

2.6   ODR

No Hint.

2.7   BGP

When you configure BGP to advertise a specific set of prefixes only when some other specified prefix in no longer received, consider configuring a BGP conditional advertisement.

2.8   IPv6 Addressing

No hint

2.9   IPv6 Routing

No hint

2.10   Router Maintenance

Cisco router can be configured as a RARP server to service workstation requests for an IP address.

2.11   Security

If you cannot use the "established" parameter in an extended access-list to fulfill the stated requirement, consider using a reflexive access-list. The Finger service can be activated on a Cisco router. You can test it by telneting to the Finger TCP port.

2.12   QOS

Read carefully!!! Do not just read this section carefully. Read all sections carefully. In particular, read Task 2.15.5 in Section 2.15 very carefully.

2.13   Catalyst Specialties

Do some research on the "switchport voice vlan" Catalyst 3550 interface configuration command. Remember, to enable mls qos on the 3550 whenever you configure a QOS option.

2.14   Gateway Redundancy

Did you know that the last two digits of the MAC address used by HSRP are determined by the standby group number you configure? In order to understand this statement, you must be able to convert a decimal number (the standby group number) to a hex number a MAC address.

2.15   Multicasting

Whenever you see a reference to a "shared root", you know it is a PIM Sparse Mode problem. With a Sparse Mode problem, you will need to configure a Rendezvous Point.

There are three ways to configure Rendezvous Point: one static and two dynamic. You must determine what the only method of "statically" configuring a Rendezvous Point is.

As with any multicast scenario, be on the lookout for any RPF lookup problems.