


**NETMASTERCLASS**  
**ROUTING AND SWITCHING CCIE® TRACK**

# DOiT-200v6

# VOLUME II



## DOiT-200v6 Lab 19

## 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

<b>19.1 Frame Relay</b> .....	<b>8</b>
<b>19.2 Catalyst Configuration</b> .....	<b>8</b>
<b>19.3 OSPF</b> .....	<b>9</b>
<b>19.4 RIP</b> .....	<b>9</b>
<b>19.5 EIGRP</b> .....	<b>9</b>
<b>19.6 Redistribution</b> .....	<b>10</b>
<b>19.7 BGP</b> .....	<b>10</b>
<b>19.8 Traffic Optimization part 1</b> .....	<b>10</b>
<b>19.9 Traffic Optimization part 2</b> .....	<b>10</b>
<b>19.10 IPv6 BGP</b> .....	<b>11</b>
<b>19.11 QOS</b> .....	<b>11</b>
<b>19.12 Catalyst Specialties</b> .....	<b>12</b>
<b>19.13 Gateway Redundancy</b> .....	<b>12</b>
<b>19.14 Multicast</b> .....	<b>12</b>
<b>Appendices</b> .....	<b>13</b>
<b>Hints</b> .....	<b>15</b>
19.1 Frame-Relay .....	15
19.2 Catalyst Configuration.....	15
19.3 OSPF .....	15
19.4 RIP .....	15
19.5 EIGRP .....	15
19.6 Redistribution .....	15
19.7 BGP.....	15
19.8 Traffic Optimization part 1.....	16
19.9 Traffic Optimization part 2.....	16
19.10 IPv6 BGP .....	16
19.11 QoS.....	16
19.12 Catalyst Specialties.....	16
19.13 Gateway Redundancy.....	16
19.14 Multicast.....	16




## Goals and Restrictions

- IP subnets on the diagram belong to network 172.10.0.0/16.
- Do not use any static routes except those used for backup purposes.
- Advertise Loopback interfaces with their original mask.
- Do not change any prefix masks.
- Do not use policy-routing.
- 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.

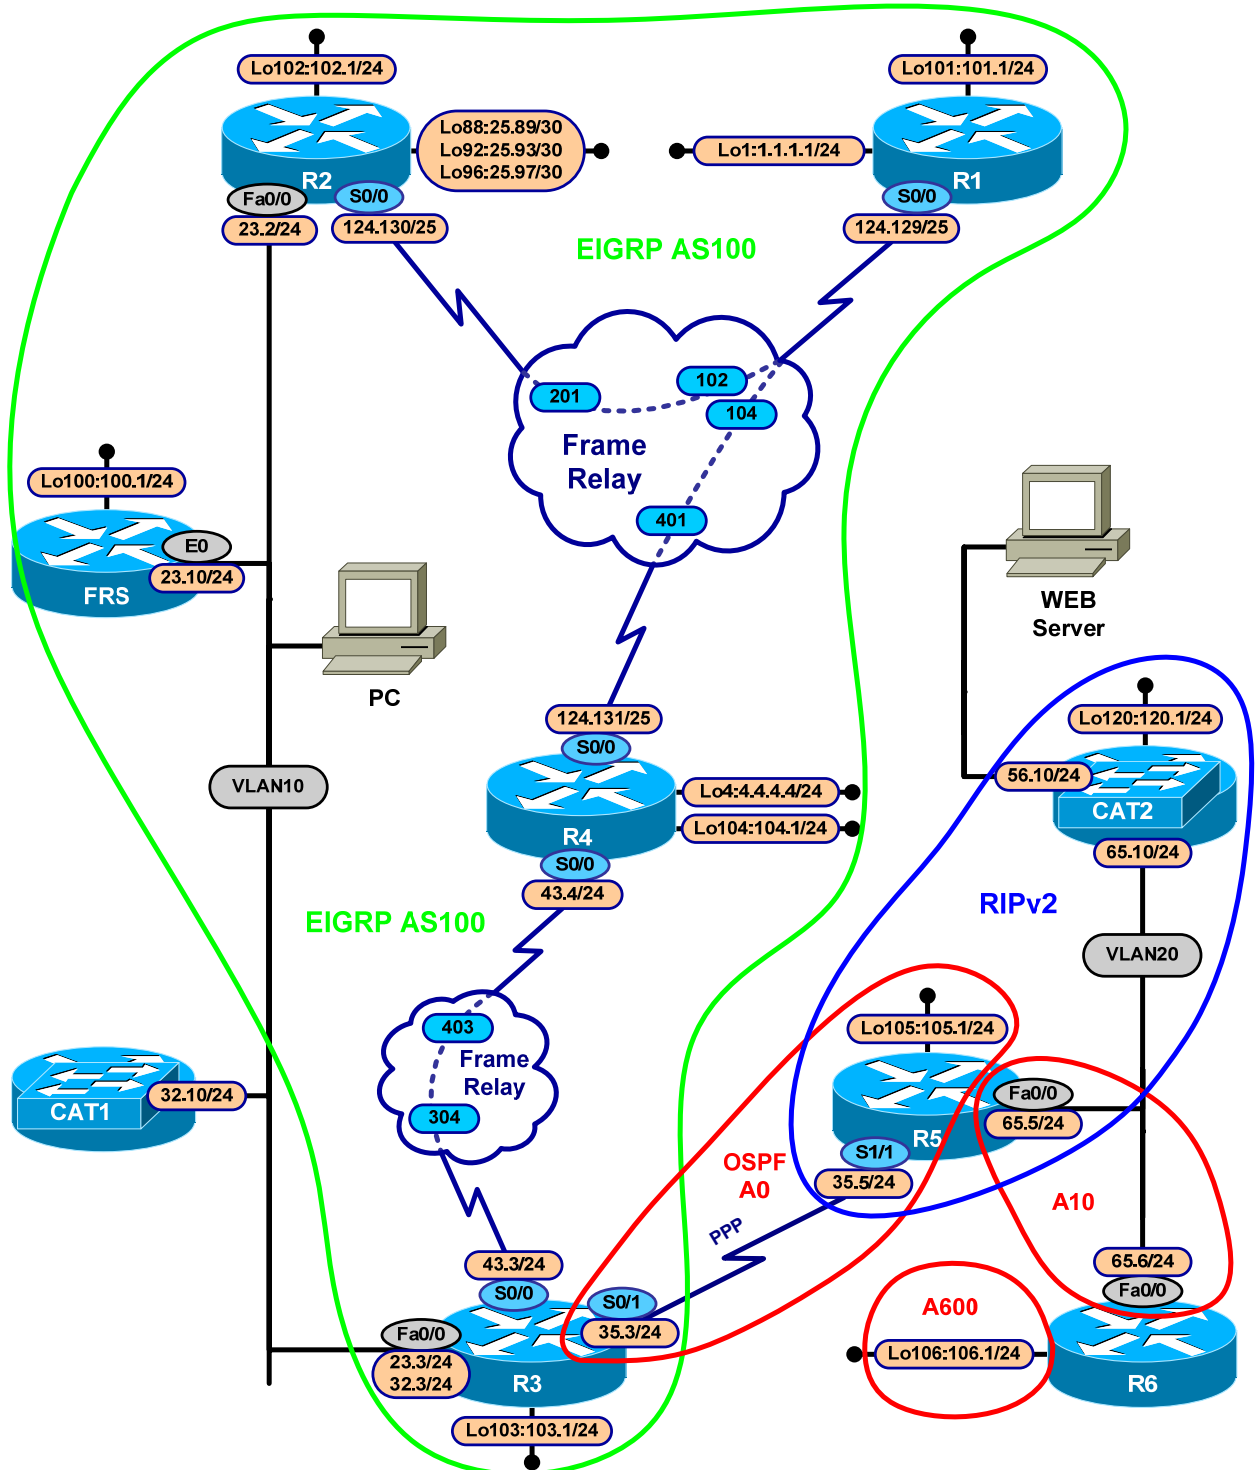


**FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM  
"NETMASTERCLASS 200V6 POD LAYOUT".**

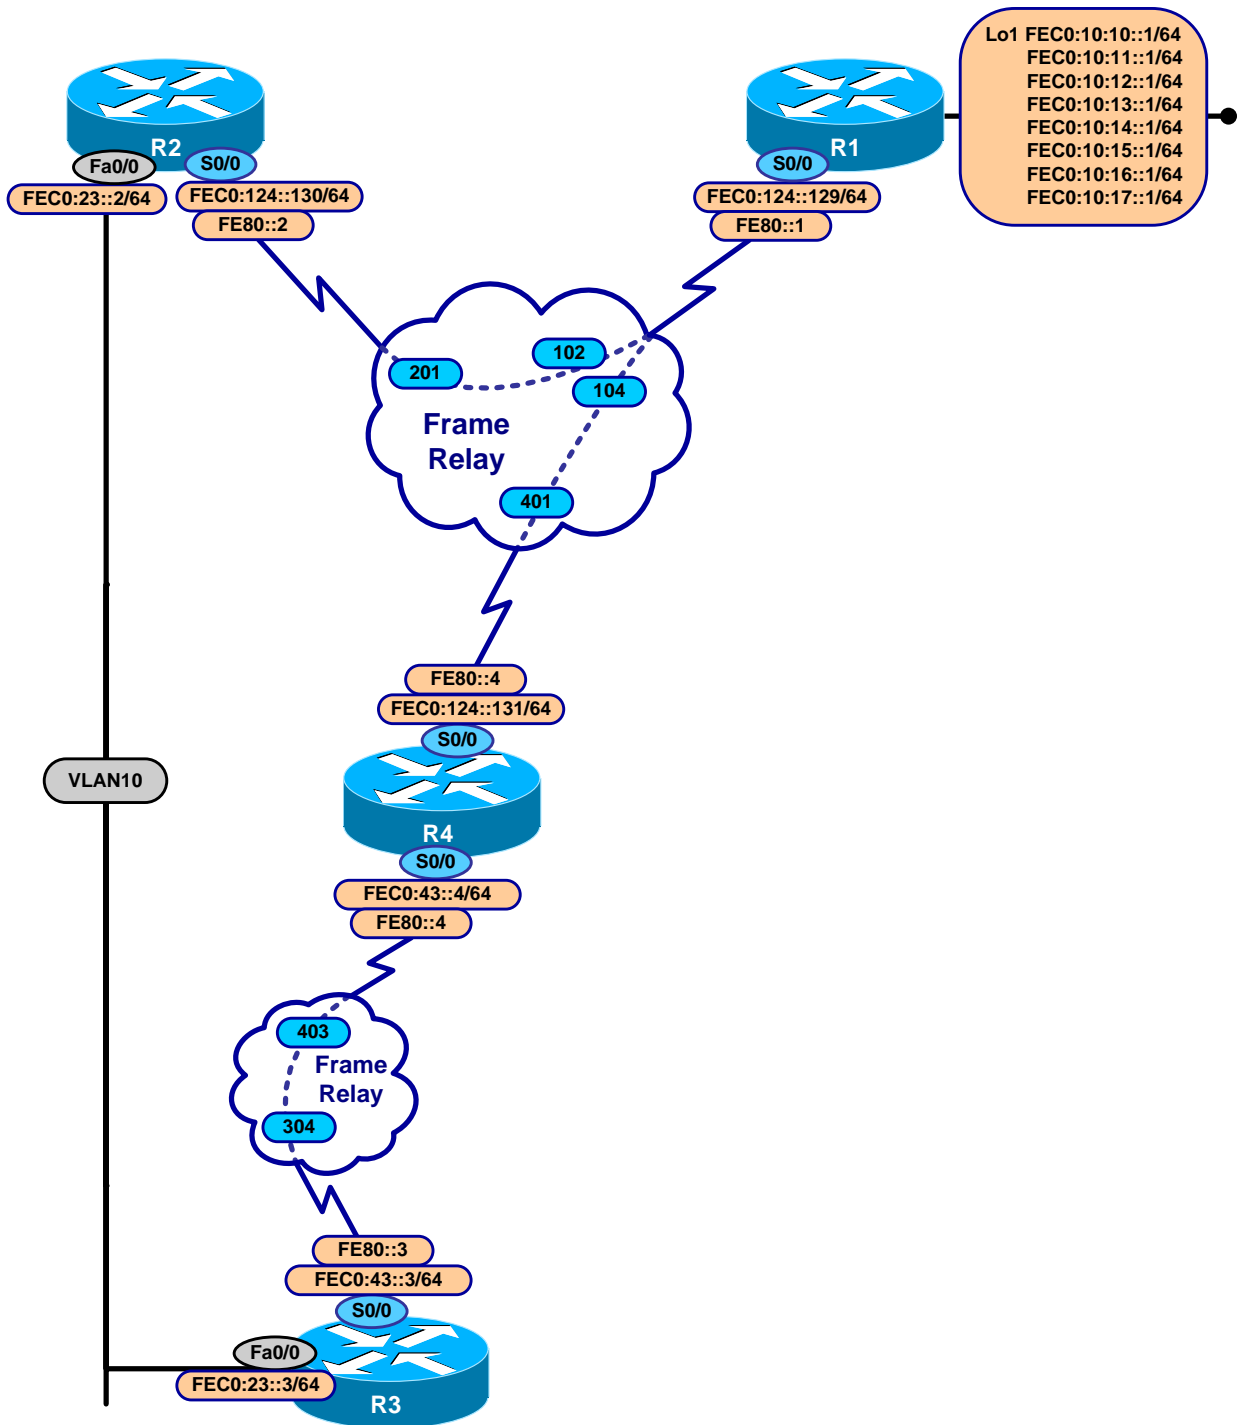
### Netmasterclass 200v6 POD Layout



Netmasterclass DOIT-200v6 Scenario 19 IPv4 IGP diagram



### Netmasterclass DOiT-200v6 Scenario 19 IPv6 IGP diagram



## 19.1 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.**

- 19.1.1 Use only the PVCs displayed on the diagram to fulfill this configuration.
- 19.1.2 Use physical interfaces on subnet 172.10.124.128/25. Use physical and point-to-point interfaces on the subnet 172.10.43.0/24.

## 19.2 Catalyst Configuration



**Shutdown interfaces FA0/15 on both switches CAT1 and CAT2**

- 19.2.1 Do not advertise VLAN's in this scenario. Use the ISL protocol for trunking.

**VLAN Configuration Table**

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

- 19.2.2 Configure the FRS Ethernet interface with an IP address of 172.10.23.10/24.
- 19.2.3 Make CAT1 the root bridge for VLAN 10. Put the Fa0/13 interface on CAT2 in the Spanning Tree blocking state. Do not use spanning tree path-cost manipulation to accomplish this task.
- 19.2.4 A PC with the NIC MAC address 00-07-85-92-D0-E7 is connected to port Fa0/10 of CAT2 on the default VLAN. Make sure only that PC is allowed to access port Fa0/10.



## 19.3 OSPF

- 19.3.1 Place PPP link between R3 and R5 in area 0.
- 19.3.2 Configure OSPF area 10 on VLAN 20 between R5 and R6. Do not elect a DR/BDR on this subnet. Make sure OSPF packets are exchanged without use of a multicast address due to security reasons.
- 19.3.3 Create Loopback 106 on R6 and place it in area 600.
- 19.3.4 Use clear text authentication on area 0. Password "nmc"
- 19.3.5 Use md5 authentication on area 10. Password "rsnmc"

## 19.4 RIP

- 19.4.1 Configure RIP version 2 between only R5 and CAT2.
- 19.4.2 Set the gateway of last resort on CAT2 only if 172.10.124.128/25 is in the R5 routing table.

## 19.5 EIGRP

- 19.5.1 Configure EIGRP AS 100 between routers R1, R2 and R4. Advertise all prefixes on these routers as EIGRP internal routes.
- 19.5.2 Do not allow multicast EIGRP traffic on the subnet between routers R1, R2 and R4
- 19.5.3 Summarize the networks below with the most optimal mask. Make sure you have a summary for these loopbacks on R1 and R4 only.
  - o IP address 172.10.25.89 255.255.255.252
  - o IP address 172.10.25.93 255.255.255.252
  - o IP address 172.10.25.97 255.255.255.252
- 19.5.4 Configure EIGRP AS100 on the VLAN 10 subnet 172.10.23.0/24 between R2, R3 and FRS.
- 19.5.5 EIGRP AS100 on FRS must be configured with the "network 172.10.0.0" statement. Make sure FRS does not advertise loopback 172.10.100.0/24. Do not use any route filtering techniques.
- 19.5.6 Ping 172.10.100.1 from the rest of the network using address 172.10.23.100.
- 19.5.7 Configure subnet 172.10.32.0/24 on VLAN 10 between R3 and CAT1. CAT1 should not run the "ip routing" process. CAT1 should be reachable from the rest of the network.

## 19.6 Redistribution

- 19.6.1 Mutually redistribute EIGRP and OSPF on R3. Redistribute RIP into OSPF on R5. Do redistribute OSPF into RIP.

## 19.7 BGP

- 19.7.1 Configure AS 23 between R2 and R3. Configure AS 4 on R4.
- 19.7.2 Peer AS 23 and AS 4 between R2 and R4 as well as R3 and R4. Do not use loopback interfaces for peering.
- 19.7.3 Advertise network 4.4.4.0/24 in AS4.
- 19.7.4 Advertise network 172.10.23.0/24 in AS23.
- 19.7.5 Outbound traffic from a PC connected to the 172.10.23.0/24 subnet destined to the 4.0.0.0/24 network should flow through R2.
- 19.7.6 Incoming traffic from the 4.0.0.0/24 network to a PC connected to the 172.10.23.0/24 subnet should flow through R2.
- 19.7.7 If the frame-Relay link on R2 goes down, the aforementioned traffic should pass through R3. Return the traffic pattern through R2 when the frame relay link on R2 is back up. Use the minimal number of BGP decision steps to accomplish this task.
- 19.7.8 Configure AS 1 on R1 and peer it with AS4 using loopbacks 1.1.1.1 and 4.4.4.4.
- 19.7.9 Advertise networks 1.1.1.0/24 and 4.4.4.0/24 at AS1 and AS4 respectively.

## 19.8 Traffic Optimization part 1

- 19.8.1 A Distributed Director is connected to VLAN 20. The IP address of the director is 172.10.65.1
- 19.8.2 Configure R5 and R6 to supply the Distributed Director with BGP and IGP metrics for efficient traffic distribution.
- 19.8.3 R5 and R6 must supply the routing metrics to only the Distributed Director specified above.

## 19.9 Traffic Optimization part 2

- 19.9.1 A web server is connected to port Fa0/15 of CAT2. Users should not configure their browser for any web proxy. Configure CAT2 to offload HTTP requests from the Web server. Check the diagram for IP address requirements.

## 19.10 IPv6 BGP

19.10.1 Assign IPv6 addresses as follows:

Router	Interface	Link-local IPv6 addresses	Site-local IPv6 address
R1	S0/0	FE80::1	FEC0:124::129/64
R2	F0/0	Default	FEC0:23::2/64
R2	S0/0	FE80::2	FEC0:124::130/64
R3	F0/0	Default	FEC0:23::3/64
R3	Serial to R4	FE80::3	FEC0::43:3/64
R4	Serial to R3	FE80::4	FEC0::43:4/64
R4	S0/0	FE80::4	FEC0::124:131/64

19.10.2 Create the following IPv6 BGP peers using directly connected, site-local addresses:

- o Peer R1 and R4 across the Frame-Relay link between them.
- o Peer R2 and R3 across VLAN10.
- o Peer R2 and R4 across the Frame-Relay link between them.
- o Peer R3 and R4 across the Frame-Relay link between them.

19.10.3 Advertise all connected IPv6 addresses into BGP using network statements. All IPv6 addresses should be reachable within the IPv6 BGP domain.

19.10.4 Add the following prefixes to Loopback 0 on R1. Advertise them into BGP with a single statement. AS23 should see only an aggregate for the highest four of these addresses. R1 should not see this aggregate.

- o FEC0:10:10::1/64
- o FEC0:10:11::1/64
- o FEC0:10:12::1/64
- o FEC0:10:13::1/64
- o FEC0:10:14::1/64
- o FEC0:10:15::1/64
- o FEC0:10:16::1/64
- o FEC0:10:17::1/64

19.10.5 Traffic leaving AS 23 for prefixes FEC0:10:10::/64 and FEC0:10:11::/64 should have a next-hop of FEC0:124::129. Traffic leaving AS23 for prefixes FEC0:10:12::/64 and FEC0:10:13::/64 should have a next-hop of FEC0:43::4.

## 19.11 QOS

19.11.1 Allocate a reservable bandwidth of 60 Kbps on the interfaces involved in this section. Send a PATH message from R4 to R5 requesting bandwidth reservation for telnet sourced from 172.10.43.4 port 5000 on R4 and destined to 172.10.35.5 port 23 on R5.

Make sure you have a single reservation for a Guaranteed Bit Rate of 5 kbps allowing bursts up to 2 Kbytes. Verify the reservation setup with the command `show ip rsvp reservation`.

## 19.12 Catalyst Specialties

19.12.1 Prohibit all traffic of ethertype 8042 from entering VLAN20.

19.12.2 Prohibit all SNA traffic from entering VLAN10.

19.12.3 Do not apply any filtering configurations to Catalyst switchports to accomplish this task.

## 19.13 Gateway Redundancy

19.13.1 Assign the IP address of 172.10.23.1 to the virtual gateway and make sure the mac-address associated with the virtual gateway is set to 0000.0c07.ac14

19.13.2 Authenticate HSRP on the 172.10.23.0/24 subnet (password nmc). Make sure Hello packets are exchanged 3 times faster than by default.

19.13.3 Select the preferred gateway that is most suitable for other tasks of this exam by using priority 150.

## 19.14 Multicast

19.14.1 Configure dense mode multicast routing between R4, R3 and R5.

19.14.2 Join dense group 229.50.50.50 on Loopback 105 or R5.


19.14.3 Make sure you can ping 229.50.50.50 from R1.

## 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

### 19.1 Frame-Relay

No hint.

### 19.2 Catalyst Configuration

A port is placed in a spanning tree blocking state based upon the following criteria in the order presented: (1) what port has the lower total path cost to the root bridge (2) what bridge has a lower bridge id (3) what port has a lower port priority.

### 19.3 OSPF

Both of the OSPF network types point-to-multipoint and point-to-multipoint non-broadcast do not perform DR/BDR elections. However, the point-to-multipoint non-broadcast network type requires neighbor statements and the point-to-multipoint network type uses the 224.0.0.5 multicast.

### 19.4 RIP

No hint.

### 19.5 EIGRP

EIGRP is a Distance Vector routing protocol. Split-horizon is applied to EIGRP updates and is enabled by default for all EIGRP configured interfaces.

When you create an EIGRP summary address, it inserts the prefix into the local forwarding table referencing a Null0 interface. It then advertises this prefix to other EIGRP routers.

### 19.6 Redistribution

Make sure you thoroughly test for universal reachability.

### 19.7 BGP

The first comparison performed during the BGP path selection process is the BGP administrative weight.

When determining how to configure many of the BGP tasks listed in this Scenario, remember to carefully read tasks in other sections of this scenario. Other tasks in this Scenario will constrain your options for configuring BGP.

Remember that EBGP learned updates have an administrative distance of 20. When you advertise EBGP prefixes that are also used to form the underlying EBGP neighbor relationship between two peers, your BGP peer stability will be adversely affected.

## 19.8 Traffic Optimization part 1

No hint.

## 19.9 Traffic Optimization part 2

No hint.

## 19.10 IPv6 BGP

Note that these peerings are the same as the IPv4 BGP peerings. The prefixes involved in the aggregation and preferred exit tasks fall on very neat bit boundaries. Remember that there are sixteen bits between colons in an IPv6 address.

## 19.11 QoS

Research the commands: `ip rsvp sender-host` and `ip rsvp reservation host`.

## 19.12 Catalyst Specialties

VLAN filter.

## 19.13 Gateway Redundancy

No hint.

## 19.14 Multicast

Always look for RPF issues.