

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

| | |
|---|-----------|
| 17.1 Serial interfaces | 8 |
| 17.2 Catalyst Configuration | 8 |
| 17.3 OSPF | 9 |
| 17.4 RIP | 9 |
| 17.5 EIGRP | 9 |
| 17.6 Enhance Routing Stability | 10 |
| 17.7 Redistribution..... | 10 |
| 17.8 BGP | 10 |
| 17.9 Address Administration | 11 |
| 17.10 IPv6 Routing over IPv4 | 11 |
| 17.11 QOS | 11 |
| 17.12 Catalyst Specialties | 12 |
| 17.13 Gateway Redundancy..... | 12 |
| 17.14 Multicast | 12 |
| Appendices..... | 13 |
| Hints | 15 |
| 17.1 Serial Interfaces | 15 |
| 17.2 Catalyst Configuration..... | 15 |
| 17.3 OSPF | 15 |
| 17.4 RIP | 15 |
| 17.5 EIGRP | 15 |
| 17.6 Redistribution | 15 |
| 17.7 Enhance Routing Stability..... | 15 |
| 17.8 BGP..... | 15 |
| 17.9 Address Administration | 15 |
| 17.10 IPv6 Routing over IPv4 | 16 |
| 17.11 QoS..... | 16 |
| 17.12 Catalyst Specialties..... | 16 |
| 17.13 Gateway Redundancy..... | 16 |
| 17.14 Multicast..... | 16 |




Goals and Restrictions

- IP subnets on the diagram belong to network 170.18.0.0/16.
- Do not use any static routes.
- Advertise Loopback interfaces with their original mask.
- Do not use a static default route 0.0.0.0/0.
- All IP addresses involved in this scenario must be reachable, unless specified otherwise.

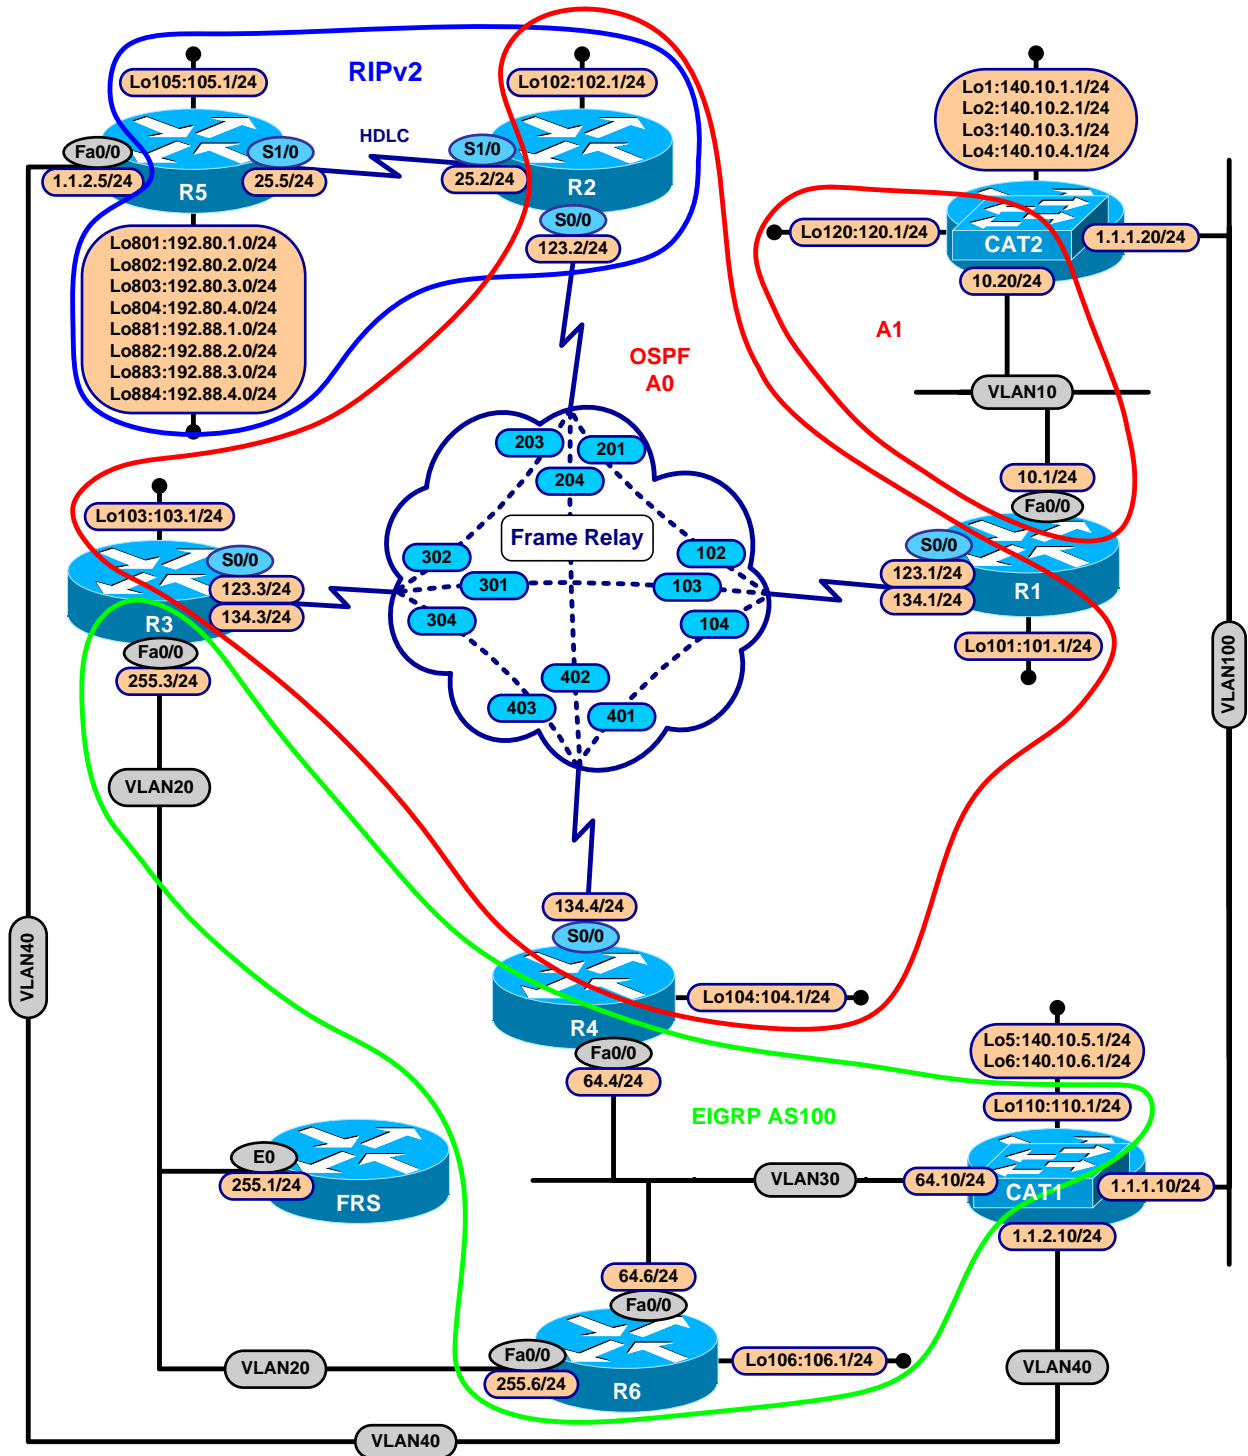


**FOR PHYSICAL CONNECTIVITY, CHECK THE DIAGRAM
"NETMASTERCLASS 200V6 POD LAYOUT".**

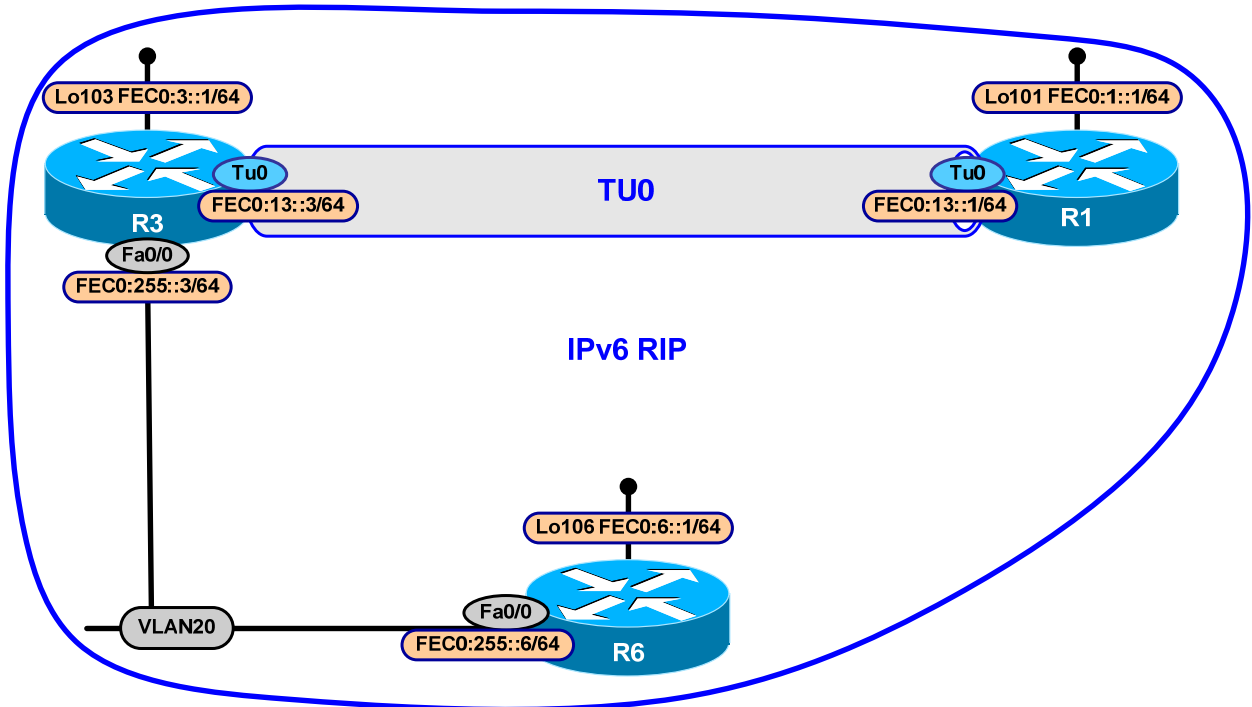
Netmasterclass 200v6 POD Layout



Netmasterclass DOiT-200v6 Scenario 17 IPv4 IGP diagram



Netmasterclass DOiT-200v6 Scenario 17 IPv6 IGP diagram



17.1 Serial interfaces



The Frame-Relay Switch configuration is provided in the Appendix. Do not modify the frame relay route statements on the frame relay switch FRS.

- 17.1.1 Configure Frame Relay encapsulation between the routers R1, R2, R3 and R4. See the diagram.
- 17.1.2 R1, R2 and R3 should be in the same subnet 170.18.123.0/24. R1, R4 and R3 should be in the same subnet 170.18.134.0/24.
- 17.1.3 Use only point-to-point logical interfaces wherever possible and use physical interfaces otherwise.
- 17.1.4 Use only the minimum number of DLCI's to fulfill this configuration.
- 17.1.5 Configure the HDLC protocol on the serial link between R2 and R5.

17.2 Catalyst Configuration

- 17.2.1 Configure the VLAN's referenced in the diagram and in the VLAN configuration table.

VLAN Configuration Table

| Router | Interface | VLAN |
|--------|-----------------|-----------------------|
| R1 | FastEthernet0/0 | VLAN10 |
| R2 | FastEthernet0/0 | - |
| R3 | FastEthernet0/1 | - |
| R3 | FastEthernet0/0 | VLAN20 |
| R4 | FastEthernet0/0 | VLAN30 |
| R5 | FastEthernet0/0 | VLAN40 |
| R6 | FastEthernet0/0 | VLAN20 VLAN30 |
| FRS | Ethernet0 | VLAN20 |
| CAT1 | - | VLAN30 VLAN100 VLAN40 |
| CAT2 | - | VLAN10 VLAN100 |

- 17.2.2 Make sure that you see the following output from the command **show etherchannel summary** :

```

Number of channel-groups in use: 1
Number of aggregators:          1

Group  Port-channel  Protocol    Ports
-----+-----+-----+-----
1      Po1(SU)        PAgP        Fa0/13(P) Fa0/14(P) Fa0/15(P)
  
```


17.3 OSPF

- 17.3.1 Configure OSPF area 0 between R1, R2 and R3. Use an OSPF network type that elects a DR but does not require a neighbor statement. Make sure R3 is the DR.
- 17.3.2 Configure OSPF area 0 on the subnet between R1, R3 and R4. Use an OSPF network type that elects a DR and requires neighbor statements.
- 17.3.3 For the subnet shared by R1, R3 and R4, make sure that the loss of a neighbor relationship is detected twice as fast as the default.
- 17.3.4 Place VLAN10 and Loopback 120 into OSPF area 1. Make sure OSPF sends the minimum information to CAT2.

17.4 RIP

- 17.4.1 Configure RIP Version 2 between R2 and R5. Send the minimum required information from R2 to R5.
- 17.4.2 Advertise the following networks on R5 into RIP:
 - o 192.80.1.0/24
 - o 192.80.2.0/24
 - o 192.80.3.0/24
 - o 192.80.4.0/24
 - o 192.88.1.0/24
 - o 192.88.2.0/24
 - o 192.88.3.0/24
 - o 192.88.4.0/24
- 17.4.3 Allow only the following networks to be advertised to router R2 using the minimum number of access list statements.
 - o 192.80.2.0
 - o 192.80.3.0
 - o 192.88.2.0
 - o 192.88.3.0
 - o 170.18.102.0

17.5 EIGRP

- 17.5.1 Configure EIGRP AS 100 between R4, R3 and R6.
- 17.5.2 Do not form a neighbor relationship over the PVC between R3 and R4.

17.6 Enhance Routing Stability

- 17.6.1 Imagine that the F0/0 interface on R4 is flapping, causing instability throughout the network. Implement a feature on this interface that will isolate failures so that disturbances are not propagated.

17.7 Redistribution

- 17.7.1 Redistribute RIP into OSPF on R2
- 17.7.2 Mutually redistribute OSPF and EIGRP on R3 and R4.
- 17.7.3 Note that VLANs 100 and 40 are only used for BGP backend connectivity. These networks do not have to be reachable from routers not directly connected.

17.8 BGP

- 17.8.1 Configure BGP AS 1581 on CAT1 and AS 1771 on CAT2
- 17.8.2 Peer AS 1581 and AS1771 over 1.1.1.0/24. See the diagram.
- 17.8.3 Advertise the following networks from AS 1771:
- o 140.10.1.0/24
 - o 140.10.2.0/24
 - o 140.10.3.0/24
 - o 140.10.4.0/24
- 17.8.4 Advertise the following networks from AS 1581:
- o 140.10.5.0/24
 - o 140.10.6.0/24
- 17.8.5 Configure BGP AS 100 on R1, R2 and R5. Configure BGP AS 10 on R3, R4 and R6.
- 17.8.6 Do not use a full mesh in AS 10.
- 17.8.7 Peer AS 100 and AS 10 between the following two pairs of EBGP peers: R1 and R4; R2 and R3.
- 17.8.8 Peer AS 100 and AS 1581 between CAT1 and R5 over VLAN 40.
- 17.8.9 Peer AS100 and AS 1771 between CAT2 and R1 over VLAN 10.

17.8.10 Prefer R5 as an exit point to networks 140.10.*.0/24 advertised earlier in this section.

17.8.11 In AS 10, set the local-pref for all prefixes originating from AS 1771 and traversing AS 1581 to prefer R3 as a next hop. Also, set the local-pref for all prefixes originating from AS 1581 to prefer R4 as a next hop. Use values 200 and 300 to accomplish this task. Do not use the AS-PATH or the IP address prefix as match criteria to set the local-pref in AS 10.

17.9 Address Administration

17.9.1 Configure FRS to act as a host with an IP address of 170.18.255.1/24.

17.9.2 FRS should dynamically prefer R6 as a gateway. Do not use HSRP or VRRP. Do not use any static configuration.

17.9.3 All packets originating from FRS should have the source IP address changed at the first hop router. The source IP address must be in the 170.18.0.0/16 range.

17.10 IPv6 Routing over IPv4

17.10.1 Assign the IPv6 addresses in the table below. Assign additional IPv6 addresses as necessary.

| Router | Interface | Site-local IPv6 address |
|--------|--------------|-------------------------|
| R1 | Loopback101 | FEC0:1::1/64 |
| R3 | Loopback 103 | FEC0:3::1/64 |
| R6 | Loopback106 | FEC0:6:6/64 |
| R3 | F0/0 | FEC0:255::3/64 |
| R6 | F0/0.20 | FEC0:255::6/64 |

17.10.2 Assure full IPv6 connectivity along the path R1 <> R3 <> R6. Do not put IPv6 addresses on any Frame-Relay link. Use IPv6 RIP with tag RIP between R1 and R3 and R6.

17.11 QOS

17.11.1 Restrict TTCP traffic from CAT2 destined to 170.18.11.1 port 5001 to 1000000 Bit/sec on R1. Allow burst traffic up to 512,000 bytes.

17.11.2 Traffic which exceeds the above specified condition shall be dropped.

17.11.3 See the diagram for IP addressing

17.12 Catalyst Specialties

- 17.12.1 On CAT2 VLAN 200 and assign port fa0/20 to it. Do not allow BPDU traffic on this VLAN.
- 17.12.2 Configure R4 as the preferred gateway on CAT1 for networks originating from R1, R2 and R5.
- 17.12.3 Configure R6 as a backup gateway on CAT1. Do not use HSRP or IRDP protocols.
- 17.12.4 Configure CAT1 so that it can be managed by a network management service that uses UDP port 161. Set read-only access with the string of RS-NMC. Set read-write access with the string NMC.

17.13 Gateway Redundancy

- 17.13.1 Configure HSRP between R4 and R6. Make R4 the preferred gateway. Switch over to R6 if the frame relay connection on R4 goes down.
- 17.13.2 The virtual gateway IP address is 172.10.64.1.

17.14 Multicast


- 17.14.1 Enable multicast routing between routers R1, R2, and R3. Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree.
- 17.14.2 Configure all of the above listed routers to join the multicast group 229.10.10.10. Associate this multicast group with a loopback interface on each router.
- 17.14.3 Use the 224.0.1.39 PIM dense group for this configuration. Make R3 the root of the shared tree. Accomplish this task by configuring only R3.
- 17.14.4 Ping the multicast group 229.10.10.10 from R4 to all other multicast routers.

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

17.1 Serial Interfaces

Before you begin configuring your Frame-Relay interfaces carefully read the IGP configuration tasks in this Scenario.

Remember that you can configure the Frame-Relay physical interface even when you have a subinterface configured under the interface.

17.2 Catalyst Configuration

No hint.

17.3 OSPF

You cannot place a secondary IP address in a different OSPF area than that assigned to the primary IP address.

17.4 RIP

Only the 170.18.0.0/16 class B networks must be reachable.

17.5 EIGRP

No hint.

17.6 Redistribution

Keep It Simple

17.7 Enhance Routing Stability

A 12.2T feature usually associated with just BGP route instability.

17.8 BGP

If you can't match on prefix and or AS-Path at AS 10, perhaps you can match by prefix or AS-Path in another AS. When you match by prefix or AS-Path in another AS, perhaps you can tag the desired prefixes with a tag such as a BGP community.

17.9 Address Administration

HSRP makes two or more routers appear as one router to end systems on a subnet. IRDP allows an end system to select a preferred router as a default gateway.

17.10 IPv6 Routing over IPv4

Go to your “desperation checklist”! No, not PBR!

17.11 QoS

Traffic shaping buffers traffic that is out of profile. Traffic policing drops or marks traffic that is out of profile. You can deploy traffic policing with CAR or the Modular QOS CLI.

17.12 Catalyst Specialties

No SPT, no BPDUs.

17.13 Gateway Redundancy

No hint.

17.14 Multicast

224.0.1.39 and 224.0.1.40 are reserved multicast addresses used by PIM Auto-RP.