CISCO® CCIE Routing and Switching Lab Guide

1ST EDITION

John Kaberna
(CCIE# 7146)

Raymond Fung
(CCIE# 6832)

Table of Contents

SECTION II: LAYER 3 ROUTING PROTOCOLS

C H A P T E R 6: GENERAL ROUTING

C H A P T E R 7: OSPF

CHAPTER 8: BGP

C H A P T E R 9: EIGRP

CHAPTER 1

FRAME-RELAY

Interface Types

Prior to configuring any frame-relay network, you'll need to identify the physical interfaces to be used and their corresponding DLCI's. Figure 1.1 illustrates the topology discussed in the upcoming sections.

Figure 1.1. *Frame Relay full mesh topology*

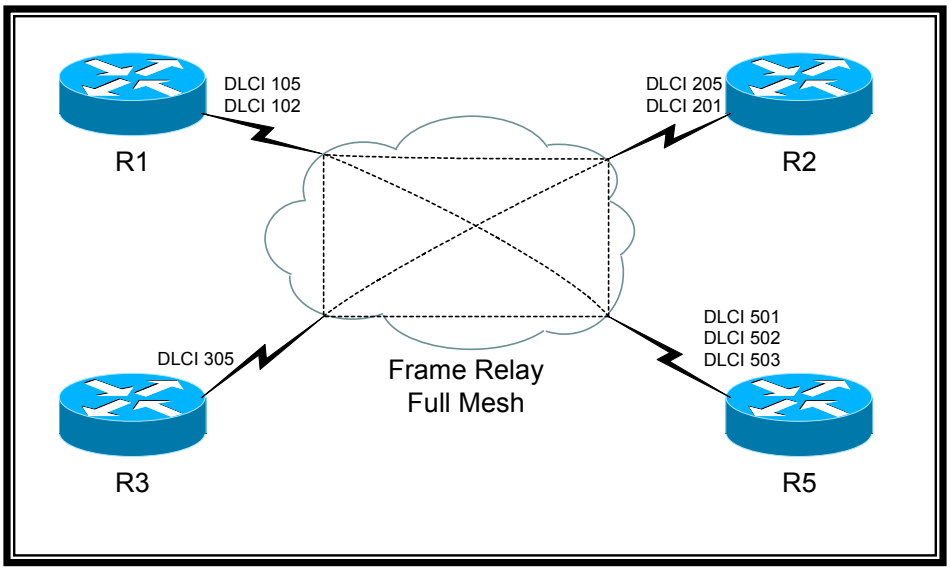

Physical interfaces

Physical interfaces do not use point-to-point or point-to-multipoint subinterfaces. Physical interfaces receive all DLCI's advertised by the switch. If you create a subinterface, you need to tell that subinterface which DLCI it should use. To verify which DLCI's a physical interface receives use the **`show frame pvc`** command.

If your PVC is configured properly, the DLCI USAGE field should be LOCAL and PVC STATUS should be ACTIVE. Remember that both ends of the PVC need to be configured properly in order for the PVC to be active. If you have not yet configured an IP address, your DLCI USAGE should be UNUSED.

```
r5# show frame pvc

PVC Statistics for interface Serial1 (Frame Relay DTE)

DLCI = 501, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE
input pkts 1              output pkts 1          in bytes 30
  out bytes 30            dropped pkts 0         in FECN pkts 0
  in BECN pkts 0   out FECN pkts 0        out BECN pkts 0
  in DE pkts 0            out DE pkts 0
  out bcast pkts 1 out bcast bytes 30
  pvc create time 00:05:51, last time pvc status changed 00:05:52

DLCI = 502, DLCI USAGE = UNUSED, PVC STATUS = ACTIVE, INTERFACE = Serial1

  input pkts 0            output pkts 4          in bytes 0
  out bytes 120    dropped pkts 0         in FECN pkts 0
  in BECN pkts 0   out FECN pkts 0        out BECN pkts 0
  in DE pkts 0            out DE pkts 0
  out bcast pkts 4 out bcast bytes 120    Num Pkts Switched 0
  pvc create time 00:08:53, last time pvc status changed 00:02:43
```

You should also check your PVC status. Active means the PVC is active and information can be exchanged. Inactive means the router's local connection to the switch is working to the frame switch, but there is a problem on the remote end. Both ends of a PVC must be up for it to be active. Deleted means the router is not receiving LMI from the frame switch or there is a layer 1 problem.

CHAPTER 8

BORDER GATEWAY PROTOCOL

BGP is used to route between Autonomous Systems and is the routing protocol for the Internet. Configuration of BGP can be quite complicated and there are many options. We will try to cover most of the BGP topics that may be on the lab exam.

BGP Peers

BGP requires that routers establish a peer relationship. Unlike OSPF, this neighbor (peer) relationship must be manually configured. Routers are considered peers or neighbors whenever they open up a TCP session to exchange routing information. When routers communicate for the first time, they exchange their entire routing table. From then on, they send only incremental updates. BGP uses TCP as its transport protocol, via port 179.

Internal BGP (IBGP)

Exchanges routing information within the same AS between routers.
IBGP routers must be fully meshed
All IBGP routers must have the same BGP routing table (only EBGP links can adjust or filter BGP routes)

External BGP (EBGP)

Used when routers belong to different AS's and exchange BGP updates.
BGP must be synchronized with the IGP (IGP's include such routing protocols as OSPF, RIP, EIGRP, etc.) if the AS provides transit service for other AS's. Synchronization helps prevent BGP from advertising an internal route that is no longer available via the IGP.

When to disable synchronization:
Your AS does not transfer traffic from one AS to another (transit AS)
All the transit routers on your AS are running BGP

Basic BGP Configuration

Enable BGP using a local BGP AS number assigned by InterNIC (for a production environment). During the lab exam, you will use AS numbers assigned by the exam instructions.

There are a few rules when configuring BGP. Neighbors must be configured on both sides. Also, neighbors must be directly connected or have a specific route (a default route will not work) to the neighbor. Multihop must be configured if the neighbors are not directly connected. Networks configured must have a match in the routing table in order for BGP to advertise the route

To configure BGP, first start the BGP routing process. Then advertise networks in to BGP (if applicable). Finally, configure your BGP peers.

Figure 8.1 *Basic BGP topology*

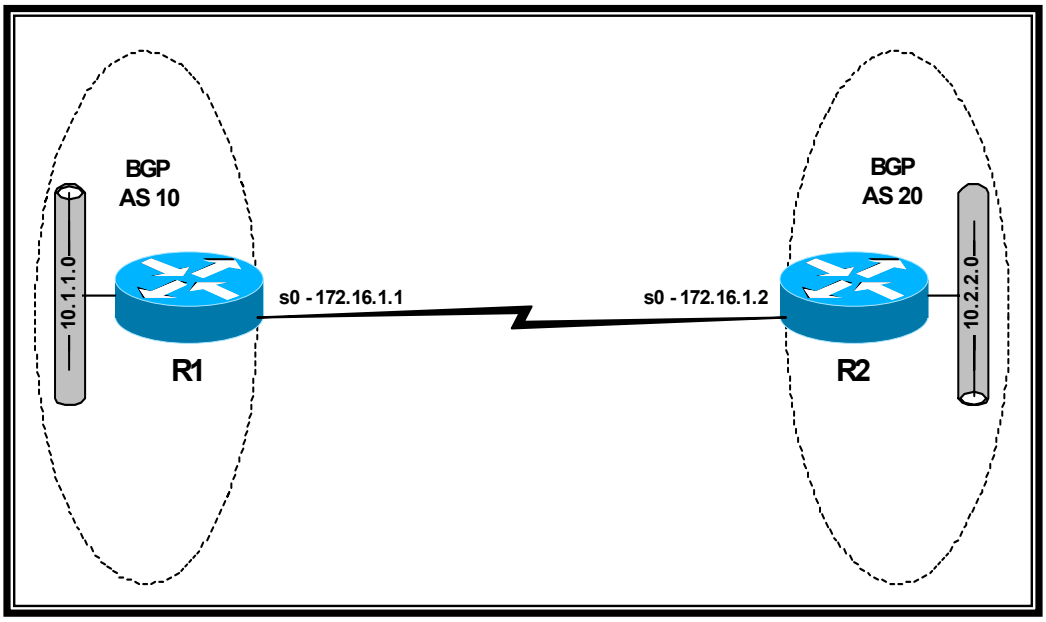


```
r1(config)# router bgp 10
r2(config)# router bgp 20
```

Configure the networks you want to advertise.

```
r1(config-router)# network 10.1.1.0 mask 255.255.255.0
r2(config-router)# network 10.2.2.0 mask 255.255.255.0
```

Specify BGP neighbors and IP address.

```
r1(config-router)# neighbor 172.16.1.2 remote-as 20
r2(config-router)# neighbor 172.16.1.1 remote-as 10
```

Note Once you have configured basic BGP, you'll typically need to clear the BGP session for any new changes to take effect by entering the **`clear ip bgp*`** command.

****WARNING**** While we recommend using the command **clear ip bgp *** in a lab environment, we highly suggest avoiding this command in a production environment!

BGP Neighbor Verification

Once neighbors are configured, verify that you have a valid TCP and BGP connection.

```
r1# show ip bgp neighbors
BGP neighbor is 172.16.1.2,  remote AS 20, external link
  BGP version 4, remote router ID 10.2.2.22
  BGP state = Established, up for 00:01:20
  Last read 00:00:19, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 8 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  1 accepted prefixes consume 36 bytes
  Prefix advertised 1, suppressed 0, withdrawn 0
  Number of NLRIs in the update sent: max 1, min 0
  Connections established 2; dropped 1
  Last reset 00:01:59, due to User reset
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 172.16.1.1, Local port: 11000
Foreign host: 172.16.1.2, Foreign port: 179
```