INTERNET TEACHING LAB: BGP LAB

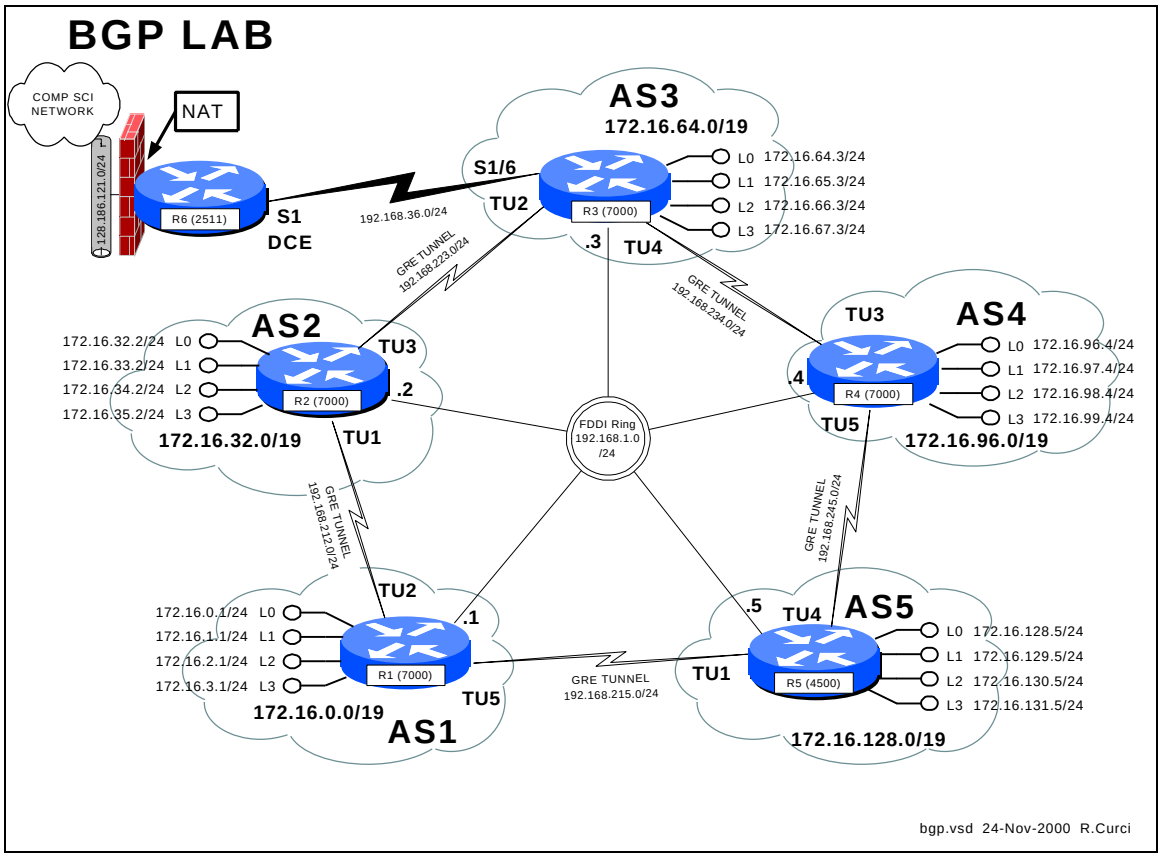

Overview

In this lab, we will explore the Border Gateway Protocol (BGP) and Generic Route Encapsulation (GRE) tunnels. Each router r1 through r5 will physically connect to a common FDDI ring. A set of 5 GRE tunnels will be implemented connecting r1→r2, r2→r3, r3→r4, r4→r5, and r5→r1. These tunnels do not use TCP or UDP, but instead a separate protocol number 47 that operates over IP. Once established, tunnels are treated by the router like any other point-to-point interface. Each router r1 through r5 will be in a separate autonomous system each with its own /19 CIDR block of IP address space. Each router r1 through r5 will be configured to peer using exterior BGP with its two neighbors. BGP version 4 is the exterior routing protocol deployed on the backbone of the Internet. BGP organizes the network into autonomous systems identified by autonomous system numbers (ASNs). ASNs are uniquely assigned by the American Registry for Internet Numbers (ARIN). Only organizations with more than one Internet Service Provider (ISP) who are "multihomed" are eligible to receive a registered ASN. You can find out more about BGP in the Cisco routing protocols configuration guide. As of this writing, the definitive source of information for this protocol is the textbook Internet Routing Architectures by Bassam Halabi published by Cisco Press in 1997.

Here is the FSU autonomous system number registration record at ARIN:

```
acns% whois -h whois.arin.net 2553
Florida State University (ASN-FSU)
   Academic Computing & Network Services
   Room 200, Sliger Building
   2035 East Paul Dirac Drive
   Tallahassee, FL 32310

   Autonomous System Name: FSU-AS
   Autonomous System Number: 2553

   Coordinator:
      Garner, Lee [Systems Programmer]  (LG36-ARIN)  garner@ACNS.FSU.EDU
      850-644-2592 (FAX) 850-644-8722

   Record last updated on 25-Jan-1995.
   Database last updated on 24-Nov-2000 18:13:50 EDT.
```

Here is a summary of BGP peering sessions on the FSU BFS-7507 router. Note that our peer at IP address 199.44.5.225 (Sprint) is sending us over 92,000 prefixes.

```
bfs-7507#show ip bgp sum
BGP router identifier 128.186.253.5, local AS number 2553
BGP table version is 10339797, main routing table version 10339797
93124 network entries and 293284 paths using 19684376 bytes of memory
44120 BGP path attribute entries using 2294812 bytes of memory
23517 BGP AS-PATH entries using 634144 bytes of memory
32 BGP community entries using 852 bytes of memory
1772 BGP route-map cache entries using 28352 bytes of memory
34843 BGP filter-list cache entries using 418116 bytes of memory
109503 received paths for inbound soft reconfiguration
BGP activity 657129/958415 prefixes, 6401589/6108305 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
128.186.250.194 4  7202   72889   72879 10339797    0    0 7w1d            1
128.186.250.201 4  3996   73232   72886 10339797    0    0 3w3d           39
128.186.253.7   4  2553 2966677 2230491 10339797    0    0 3w0d        74247
192.80.53.41    4 11537  128228   72861 10339774    0    0 3w2d         4025
192.80.53.62    4  6356   72699   72929 10339792    0    0 5d13h           3
192.80.53.66    4  5661   72870   72878 10339797    0    0 7w1d            1
192.80.53.70    4  7939   72919   72922 10339774    0    0 1w0d            1
192.80.53.106   4  3506  216733 3135856 10339774    0    0 7w1d        12960
199.44.5.225    4  3447 2356372   72883 10339774    0    0 7w1d        92501
```

FSU is only advertising a small number of networks to our ISP (Sprint). This helps prevent us from unintentionally becoming a transit AS:

```
bfs-7507#show ip bgp neighbor 199.44.5.225 advertised-routes
BGP table version is 10339840, local router ID is 128.186.253.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 128.186.0.0      0.0.0.0                  0         32768 i
*> 144.174.0.0      192.80.53.106            0    155      0 3506 i
*> 146.201.0.0      0.0.0.0                 20         32768 i
*> 192.80.53.0      0.0.0.0                  0         32768 i
bfs-7507#
```

PART1 – Basic IGP (RIP) Configuration

Each router r1 through r5 will have only its physical FDDI interface enabled. The only exception is router r3 who will additionally have its serial port enabled to connect with r6 for Internet connectivity. When finished with this part, verify that you can PING the loopback0 IP address on r6, 192.168.66.6. Test by PINGing the FDDI IP broadcast address 192.168.1.255. You should hear responses from the other 4 FDDI-connected routers if all is well.

The following commands may be helpful in debugging this part:

- show cdp neighbor
- ping w.x.y.z
- show ip protocol
- show ip route
- show ip route RIP

For each router, you will need both the common part of the configuration and router specific portion as appropriate that follows:

COMMON:

```
service udp-small-servers
service tcp-small-servers
enable password cisco
no ip domain-lookup
ip classless
ip subnet-zero
logging buffered
clock timezone EST -5
clock summer-time EDT recurring
ntp server 192.168.66.6
snmp-server community public RO
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 password cisco
 login
```

R1:

```
hostname r1
interface Fddi0/0
 ip address 192.168.1.1 255.255.255.0
 no shutdown
interface E2/0
 description Vlan 10 to cat1 FA0/1
 ip address 192.168.10.1 255.255.255.0
 no shutdown
interface E2/1
 description Vlan 20 to cat1 FA0/2
 ip address 192.168.20.1 255.255.255.0
 no shutdown
interface E2/2
 description Vlan 30 to cat1 FA0/3
 ip address 192.168.30.1 255.255.255.0
 no shutdown
interface E2/3
 description Vlan 40 to cat1 FA0/4
 ip address 192.168.40.1 255.255.255.0
 no shutdown
interface E2/4
 description Vlan 50 to cat1 FA0/5
 ip address 192.168.50.1 255.255.255.0
 no shutdown
interface E2/5
 description Vlan 60 to cat1 FA0/6
 ip address 192.168.60.1 255.255.255.0
 no shutdown
router rip
 network 192.168.1.0
 network 192.168.10.0
 network 192.168.20.0
 network 192.168.30.0
 network 192.168.40.0
 network 192.168.50.0
 network 192.168.60.0
```

R2:

```
hostname r2
interface Fddi0/0
 ip address 192.168.1.2 255.255.255.0
 no shutdown
router rip
 network 192.168.1.0
```

R3:

```
hostname r3
interface Fddi0/0
 ip address 192.168.1.3 255.255.255.0
 no shutdown
interface Serial1/6
 description Link to R6 S1
 ip address 192.168.36.3 255.255.255.0
 bandwidth 2000
 no shutdown
router rip
 network 192.168.36.0
 network 192.168.1.0
```

R4:

```
hostname r4
interface Fddi0/0
 description Link to R5 FDDI0
 ip address 192.168.1.4 255.255.255.0
 no shutdown
router bgp 4
 network 172.16.96.0 mask 255.255.224.0
 neighbor 192.168.234.3 remote-as 3
 neighbor 192.168.234.3 version 4
 neighbor 192.168.245.5 remote-as 5
 neighbor 192.168.245.5 version 4
ip route 172.16.96.0 255.255.224.0 null0
router rip
 network 192.168.1.0
```

R5:

```
hostname r5
interface FastEthernet0
 description Vlan70 to cat1 FA0/7
 ip address 192.168.70.1 255.255.255.0
 media-type 100BaseX
 no shutdown
interface Ethernet0
 description Vlan80 to cat1 FA0/8
 ip address 192.168.80.1 255.255.255.0
 media-type 10BaseT
 no shutdown
interface Ethernet1
 description Vlan90 to cat1 FA0/9
 ip address 192.168.90.1 255.255.255.0
 media-type 10BaseT
 no shutdown
interface Fddi0
 description Link to R4 FDDI0/0
 ip address 192.168.1.5 255.255.255.0
 no keepalive
 no shutdown
router rip
 network 192.168.70.0
 network 192.168.80.0
 network 192.168.90.0
 network 192.168.1.0
```

PART2 – GRE Tunnel and Loopback Interfaces

GRE tunnel and loopback interfaces are virtual interfaces created in the Cisco IOS software. On each router, establish two GRE tunnel interfaces and four loopback interfaces as shown on your network diagram and table below. GRE Tunnel interfaces are normally used to encapsulate non-IP traffic through an IP-only core network or to encapsulate private IP addresses through the public Internet. Recent versions of the Linux operating system also support GRE tunnels. The tunnel interfaces in this lab will encapsulate IP traffic in frames that will physically traverse the FDDI ring but will appear to the routers as point-to-point interfaces. You will assign an IP address to each tunnel interface just like a serial point-to-point interface. Anchor the tunnels using the FDDI IP addresses as specified in the following table. Be sure you can PING both your tunnel endpoints and the IP address assigned to the tunnel interfaces on the other side. Do **NOT** enable RIP on any tunnel or loopback interfaces (**NOT** on any 172.16.x.y interfaces). We will use BGP for routing across the tunnels in the next part. Note that CDP does not work across tunnel interfaces. The following commands may be helpful in debugging this section:

- ping
- show ip interface
- show ip interface brief
- clear counters
- show interface

Notice that the loopback and tunnel interfaces have status=up and protocol=up:

```
r1#show ip int brief
Interface              IP-Address      OK? Method Status   Protocol
Fddi0/0                192.168.1.1     YES manual up       up
Loopback0              172.16.0.1      YES manual up       up
Loopback1              172.16.1.1      YES manual up       up
Loopback2              172.16.2.1      YES manual up       up
Loopback3              172.16.3.1      YES manual up       up
Tunnel2                192.168.212.1   YES manual up       up
Tunnel5                192.168.215.1   YES manual up       up
r1#
```

Here is an example “show interface” command on a GRE tunnel:

```
r1#sh int tunnel2
Tunnel2 is up, line protocol is up
  Hardware is Tunnel
  Description: Tunnel to R2
  Internet address is 192.168.212.1/24
  MTU 1500 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
  Encapsulation TUNNEL, loopback not set, keepalive set (10 sec)
  Tunnel source 192.168.1.1, destination 192.168.1.2
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  ...
```

Rtr	Interface	IP Address	Tunnel Src	Tunnel Dest
r1	fddi0/0	192.168.1.1/24		
	loopback0	172.16.0.1/24		
	loopback1	172.16.1.1/24		
	loopback2	172.16.2.1/24		
	loopback3	172.16.3.1/24		
	tunnel2	192.168.212.1/24	192.168.1.1	192.168.1.2
	tunnel5	192.168.215.1/24	192.168.1.1	192.168.1.5
	ethernet2/0	192.168.10.1/24		
	ethernet2/1	192.168.20.1/24		
	ethernet2/2	192.168.30.1/24		
	ethernet2/3	192.168.40.1/24		
	ethernet2/4	192.168.50.1/24		
	ethernet2/5	192.168.60.1/24		
r2	fddi0/0	192.168.1.2/24		
	loopback0	172.16.32.2/24		
	loopback1	172.16.33.2/24		
	loopback2	172.16.34.2/24		
	loopback3	172.16.35.2/24		
	tunnel1	192.168.212.2/24	192.168.1.2	192.168.1.1
	tunnel3	192.168.223.2/24	192.168.1.2	192.168.1.3
r3	fddi0/0	192.168.1.3/24		
	loopback0	172.16.64.3/24		
	loopback1	172.16.65.3/24		
	loopback2	172.16.66.3/24		
	loopback3	172.16.67.3/24		
	tunnel2	192.168.223.3/24	192.168.1.3	192.168.1.2
	tunnel4	192.168.234.3/24	192.168.1.3	192.168.1.4
	serial1/6	192.168.36.3/24		
r4	fddi0/0	192.168.1.4/24		
	loopback0	172.16.96.4/24		
	loopback1	172.16.97.4/24		
	loopback2	172.16.98.4/24		
	loopback3	172.16.99.4/24		
	tunnel3	192.168.234.4/24	192.168.1.4	192.168.1.3
	tunnel5	192.168.245.4/24	192.168.1.4	192.168.1.5
r5	fddi0	192.168.1.5/24		
	loopback0	172.16.128.5/24		
	loopback1	172.16.129.5/24		
	loopback2	172.16.130.5/24		
	loopback3	172.16.131.5/24		
	tunnel1	192.168.215.5/24	192.168.1.5	192.168.1.1
	tunnel4	192.168.245.5/24	192.168.1.5	192.168.1.4
	fastethernet0	192.168.70.1/24		
	ethernet0	192.168.80.1/24		
	ethernet1	192.168.90.1/24		

PART3 – BGP Peering

On each router r1 through r5, establish a BGP peering session through each tunnel interface to your neighbor. You will be using exterior BGP or EBGP since each router is in a different ASN. On each router, you will need to advertise the networks on your loopback addresses. Instead of advertising these /24 blocks individually, you should advertise only a single prefix with a /19 network mask as documented in the diagram. When everything is working, each router r1 through r5 should have two BGP peering sessions. You should be receiving 3 network advertisements from each of your peers. We will be using the AS path length to determine the best BGP route. For example, on router r1, the BGP route to network 172.16.0.0/19 and 172.16.64.0/19 should be via Tunnel2, while the best route to networks 172.16.96.0/19 and 172.16.128.0/19 should be via Tunnel5.

The following commands may be helpful in debugging this section:

- show ip route
- show ip bgp sum
- show ip bgp neighbor w.x.y.z
- show ip bgp neighbor w.x.y.z routes
- show ip bgp neighbor w.x.y.z advertised-routes
- show ip bgp regexp .*

The following are some sample SHOW command executed from router r1 to give you an idea of what you can expect when everything is working. Note that there are two active BGP peering sessions:

```
r1#sh ip bgp sum
BGP table version is 26, main routing table version 26
5 network entries (7/15 paths) using 1092 bytes of memory
7 BGP path attribute entries using 800 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State
192.168.212.2   4     2     100     102       26    0    0 01:10:40
192.168.215.5   4     5     111     120       26    0    0 00:01:55
```

These are the BGP routes we are advertising to our BGP neighbors. The only internal route we are advertising is 172.16.0.0/19. Note that the other advertised routes are learned from our BGP peers and have ASPATH “2 3 4”, “5 4”, and “5” which all begin with one of our peer’s ASNs:

```
r1#sh ip bgp nei 192.168.212.2 advertised-routes
BGP table version is 26, local router ID is 172.16.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 172.16.0.0/19    0.0.0.0                0         32768 i
*  172.16.96.0/19   192.168.212.2                        0 2 3 4 i
*>                  192.168.215.5                        0 5 4 i
*> 172.16.128.0/19  192.168.215.5          0             0 5 i
```

Here are the routes we are receiving from our neighbor 192.168.212.2:

```
r1> sh ip bgp nei 192.168.212.2 routes
BGP table version is 26, local router ID is 172.16.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 172.16.32.0/19   192.168.212.2          0             0 2 i
*> 172.16.64.0/19   192.168.212.2                        0 2 3 i
*  172.16.96.0/19   192.168.212.2                        0 2 3 4 i
r1>
```

Here is our routing table. The first character indicates which routing protocol inserted each route where B=BGP, C=connected, and S=static. You can see the /19 CIDR block advertisements learned from BGP only for the other routers.

```
r1#show ip route 172.16.0.0
Routing entry for 172.16.0.0/16, 8 known subnets
  Attached (4 connections)
  Variably subnetted with 2 masks

B       172.16.128.0/19 [20/0] via 192.168.215.5, 00:03:19
B       172.16.32.0/19 [20/0] via 192.168.212.2, 01:12:04
S       172.16.0.0/19 is directly connected, Null0
C       172.16.1.0/24 is directly connected, Loopback1
C       172.16.2.0/24 is directly connected, Loopback2
C       172.16.3.0/24 is directly connected, Loopback3
B       172.16.96.0/19 [20/0] via 192.168.215.5, 00:03:19
B       172.16.64.0/19 [20/0] via 192.168.212.2, 01:12:04
```

Here are our BGP routes to network 172.16.64.0/19. We have two routes, each with a different ASPATH, “2 3” and “5 4 3”. The former is selected as “best” because the ASPATH is shorter.

```
r1#sh ip bgp 172.16.64.0
BGP routing table entry for 172.16.64.0/19, version 4
Paths: (2 available, best #1, advertised over EBGP)
  2 3
    192.168.212.2 from 192.168.212.2 (172.16.35.2)
      Origin IGP, valid, external, best
  5 4 3
    192.168.215.5 from 192.168.215.5 (172.16.131.5)
      Origin IGP, valid, external
```

Here are all our known BGP routes including the ASPATH for each. The argument “.*” is a regular expression matching all ASPATHs.

```
r1#sh ip bgp regexp .*
BGP table version is 26, local router ID is 172.16.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 172.16.0.0/19    0.0.0.0                0         32768 i
*> 172.16.32.0/19   192.168.212.2          0             0 2 i
*> 172.16.64.0/19   192.168.212.2                        0 2 3 i
*                   192.168.215.5                        0 5 4 3 i
*  172.16.96.0/19   192.168.212.2                       0 2 3 4 i
*>                  192.168.215.5                        0 5 4 i
*> 172.16.128.0/19  192.168.215.5          0             0 5 i
r1#
```