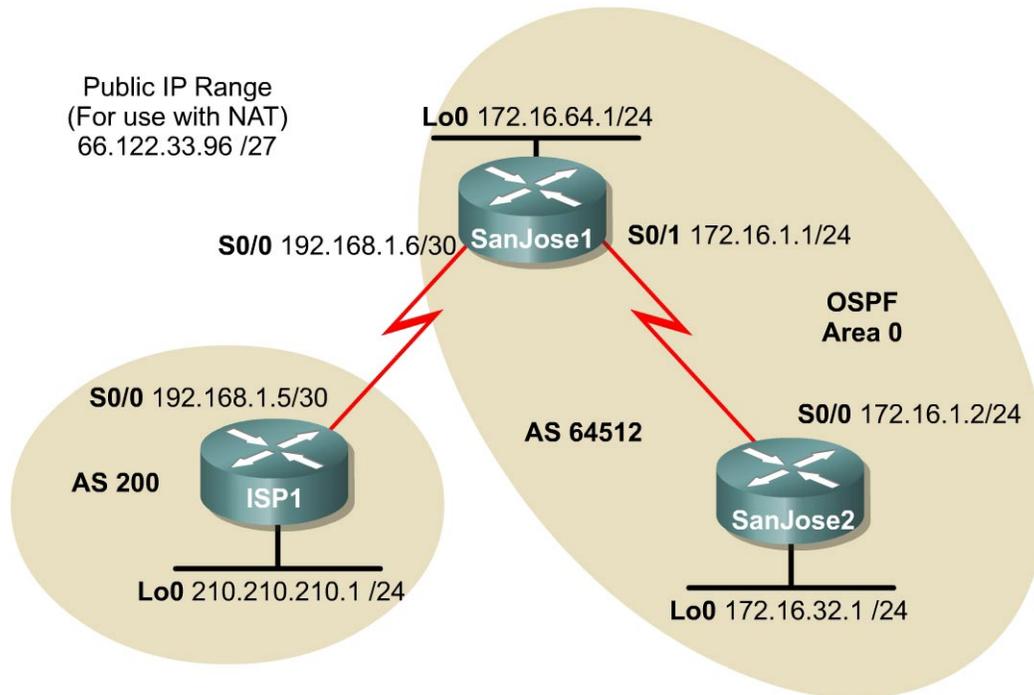


Lab 9.11.2 Configuring BGP with NAT



Objective

In this lab, the student will configure EBGP with ISP1. Configure NAT with the S0/0 link to the ISP using the public IP address range of 66.122.33.96/27. This network range will need to be advertised by way of BGP.

Scenario

The International Travel Agency runs BGP on its SanJose1 router externally with ISP1, AS 200. OSPF will be used as the internal routing protocol between the SanJose1 and SanJose2 routers. EBGP needs to be configured between ISP1 and SanJose1, and finally NAT needs to be configured on the border router in AS 64512.

Step 1

Build and configure the network according to the diagram, but do not configure a routing protocol. Configure a loopback interface on the three routers as shown in the figure. These loopbacks will be used to give a stable router ID for OSPF and BGP as well as simulate an existing network segment.

Use `ping` to test connectivity between the directly connected routers.

Note: The SanJose1 router should be able to ping the neighbor address of 192.168.1.5 and 172.16.1.2. SanJose2 will not be able to ping the ISP1 router and visa versa.

Step 2

Configure OSPF between the SanJose1 and SanJose2 routers with the following commands:

```
SanJose1(config)#router ospf 1
SanJose1(config-router)#network 172.16.64.1 0.0.0.0 area 0
SanJose1(config-router)#network 172.16.1.1 0.0.0.0 area 0

SanJose2(config)#router ospf 1
SanJose2(config-router)#network 172.16.32.1 0.0.0.0 area 0
SanJose2(config-router)#network 172.16.1.2 0.0.0.0 area 0
```

Verify that the SanJose1 and SanJose2 routers have formed an OSPF adjacency. Troubleshoot as necessary, first with the `show ip ospf neighbor` command, secondly with the `show ip ospf interface` command. The following output should be seen:

```
SanJose1#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
172.16.32.1      1    FULL/ -         00:00:37   172.16.1.2    Serial0/1

SanJose1#show ip ospf interface
Loopback0 is up, line protocol is up
  Internet Address 172.16.64.1/24, Area 0
  Process ID 1, Router ID 172.16.64.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
Serial0/1 is up, line protocol is up
  Internet Address 172.16.1.1/24, Area 0
  Process ID 1, Router ID 172.16.64.1, Network Type POINT_TO_POINT, Cost:64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.32.1
  Suppress hello for 0 neighbor(s)
```

The OSPF cost on Serial 0/1 may be different from the 64 listed in the output.

Note: Remember that OSPF calculates its costs by dividing 100,000,000 by the bandwidth.

1. What must be the bandwidth configured on the Serial 0/1 interface?
-

Router SanJose1 should have a route to the 172.16.32.0 /24 network that is learned through OSPF. Also, SanJose2 should have a route to the 172.16.64.0 /24 network also learned through OSPF.

Step 3

Configure EBGP between the ISP1 and SanJose1 routers. On the ISP1 router, enter the following configuration:

```
ISP1(config)#router bgp 200
ISP1(config-router)#network 210.210.210.0
ISP1(config-router)#neighbor 192.168.1.6 remote-as 64512

SanJose1(config)#router bgp 64512
SanJose1(config-router)#neighbor 192.168.1.5 remote-as 200
```

```
SanJose1#clear ip bgp *
```

The 172.16.0.0 /16 network has purposely not been advertised through BGP. This is because it is a private address that will not be able to route through the network. After a few minutes pass, issue the `show ip bgp summary` command on the SanJose1 router.

2. Has a BGP conversation with ISP1 router been formed?
-

Because BGP will eventually advertise outside networks that are not part of the OSPF area, the following command on the SanJose1 router must be entered:

```
SanJose1(config)#router bgp 64512
SanJose1(config-router)#no synchronization
```

The `no synchronization` command permits BGP to advertise networks without caring whether the IGP, in this case, OSPF, has the route. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP.

The routing table of SanJose1 should look similar to the output that follows:

```
Gateway of last resort is not set

B    210.210.210.0/24 [20/0] via 192.168.1.5, 00:03:24
     172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
O    172.16.32.1/32 [110/65] via 172.16.1.2, 00:41:55, Serial0/1
C    172.16.1.0/24 is directly connected, Serial1
C    172.16.64.0/24 is directly connected, Loopback0
     192.168.1.0/30 is subnetted, 1 subnets
C    192.168.1.4 is directly connected, Serial0/0
```

The 150.150.150.1 host should be able to ping the 210.210.210.1 host from SanJose1 with a standard ping.

An extended ping to the 210.210.210.1 host sourced from the Loopback 0, 172.16.64.1, should not be successful.

3. Why not?
-

Step 4

It is now time to configure Network Address Translation, NAT, to change the source IP addresses on packets from our internal private numbered network, 172.16.0.0, into public addresses that are routable on the Internet.

```
SanJose1(config)#int s0/0
SanJose1(config-if)#ip nat outside
SanJose1(config-if)#int lo 0
SanJose1(config-if)#ip nat inside
SanJose1(config-if)#int s0/1
SanJose1(config-if)#ip nat inside
```

Next, configure the NAT pool name and NAT translation policy as follows:

```
SanJose1(config)#ip nat pool NAT_PUBLIC_SUBNET 66.122.33.98 66.122.33.126
netmask 255.255.255.224
SanJose1(config)#ip nat inside source list 10 pool NAT_PUBLIC_SUBNET
overload
```

```
SanJose1(config)#access-list 10 permit 172.16.0.0 0.0.255.255
```

Access-list 10 is used to declare eligibility for NAT. Only those addresses that are permitted in this access-list will be eligible for NAT translation. Any IP address that is not included in the permit range will not be translated but will still be routed. This ACL is not intended as the only means of securing the network. Imagine SanJose2 added a 172.20.1.0 /24 network and access-list 10 had not been updated to reflect this additional network, perhaps for security reasons. Each packet that originates from the new 172.20.1.0 /24 network, that is destined for either network 210.210.210.0 or the Internet, will not be translated. Therefore, they are dropped by the ISP as unroutable. While this will keep these nodes from being able to connect to the Internet, it does so at a considerable cost. Each of the packets that are dropped will cross the WAN link to the ISP, but they will never return.

Verify that NAT is working properly. An attempt to ping from the inside interface on SanJose1 to the ISP will result in dropped packets since the ISP still does not know about the 66.122.33.96 /27 network. Issue the `debug ip nat` command and then issue an extended `ping` from an inside address to an outside address as follows:

```
SanJose1#debug ip nat
IP NAT debugging is on

SanJose1#ping
Protocol [ip]:
Target IP address: 192.168.1.6
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 172.16.64.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 44/48/52 ms

00:22:13: NAT: s=172.16.64.1->66.122.33.98, d=192.168.1.6 [0]
00:22:13: NAT: s=192.168.1.6, d=66.122.33.98->172.16.64.1 [0]
00:22:13: NAT: s=172.16.64.1->66.122.33.98, d=192.168.1.6 [1]
00:22:13: NAT: s=192.168.1.6, d=66.122.33.98->172.16.64.1 [1]
00:22:13: NAT: s=172.16.64.1->66.122.33.98, d=192.168.1.6 [2]
00:22:13: NAT: s=192.168.1.6, d=66.122.33.98->172.16.64.1 [2]
00:22:13: NAT: s=172.16.64.1->66.122.33.98, d=192.168.1.6 [3]
00:22:13: NAT: s=192.168.1.6, d=66.122.33.98->172.16.64.1 [3]
00:22:13: NAT: s=172.16.64.1->66.122.33.98, d=192.168.1.6 [4]
00:22:13: NAT: s=192.168.1.6, d=66.122.33.98->172.16.64.1 [4]
```

This debug output shows that source 172.16.64.1 is re-encapsulated as public (external) host 66.122.33.98 destined for host 192.168.1.6. The next line shows that the return packet sources from 192.168.1.6 and is sent to destination, external, host 66.122.33.98. This is actually re-encapsulated to internal host IP 172.16.64.1.

Connectivity to the ISP is sometimes lost not because of a routing issue, but rather because NAT is not translating properly. When in doubt, issue `debug ip nat` and look for output that is similar to the previous output. Notice that this never left the SanJose1 router and therefore was not influenced by routing or BGP issues that may exist. Attempt to localize your troubleshooting.

Step 5

Configure SanJose1 router to advertise the NAT pool via BGP. Reset the BGP conversation and test connectivity with an extended ping from SanJose1 to the ISP1 router as follows:

```
SanJose1(config)#router bgp 64512
SanJose1(config-router)#network 66.122.33.96 mask 255.255.255.224
```

```
SanJose1#clear ip bgp *
```

4. Was the extended ping successful from either Lo0 or S0/1 on SanJose1 to either interface IP address of the ISP1 router?

-
5. Issue the `show ip bgp` command to investigate. What was found? Is there anything missing from this output that was expected to be seen here?

```
BGP table version is 4, local router ID is 172.16.64.1
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	0.0.0.0	0		32768	i
*> 210.210.210.0	192.168.1.5	0		0	200i

If the local routes listed in the `show ip bgp` output cannot be seen by the user, then the ISP1 router, BGP Peer, will not see these routes either.

Note: Remember that BGP will advertise routes that are directly connected, identified by a static route. Routes that exist in the routing table by way of a dynamic routing protocol that has been redistributed into BGP may also be advertised. There are two options here. Either a loopback interface can be created using one of the IP addresses from the NAT pool or a static route referencing the NAT pool can be created. Review the following solutions from these options:

- Using a loopback interface results in the following:

```
SanJose1(config)#interface loopback 100
SanJose1(config-if)#ip addr 66.122.33.97 255.255.255.224
SanJose1(config-if)#ip nat outside
```

```
SanJose1#show ip bgp
BGP table version is 4, local router ID is 172.16.64.1
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
*> 66.122.33.96/27	0.0.0.0	0		32768	i
*> 172.16.0.0	0.0.0.0	0		32768	i
*> 210.210.210.0	192.168.1.5	0		0	200 i

- Using a static route results in the following:

```
SanJose1(config)#ip route 66.122.33.96 255.255.255.224 null0 230
SanJose1#show ip bgp
```

```
BGP table version is 6, local router ID is 172.16.64.1
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
---------	----------	--------	--------	--------	------

```
*> 66.122.33.96/27 0.0.0.0 0 32768 i
*> 172.16.0.0 0.0.0.0 0 32768 i
*> 210.210.210.0 192.168.1.5 0 0 200 i
```

Note: Using the loopback address assigned to IP NAT outside can be very useful for testing. Imagine that the link to the ISP had not been provisioned. Using an outside loopback address allows for configuration and testing prior to this link being put into production. The negative is that it uses one of the IP addresses from the public range, in this case it was a subnet of 32 hosts. The NAT pool started at X.X.X.98 instead of X.X.X.97 in anticipation of this loopback address.

Testing connectivity to the 210.210.210.1 host off ISP1 shows that any interface, or subnet, on SanJose1 can ping successfully.

Step 6

SanJose2 cannot reach the 210.210.210.1 when attempting communication from any one of its three interfaces.

6. Why is this the case?

Issue the `debug ip packet` and reattempt the `ping` from the SanJose2 router to 210.210.210.1. Notice that each attempted packet was dropped because of unroutable error messages?

```
SanJose2#ping 210.210.210.1

Sending 5, 100-byte ICMP Echos to 210.210.210.1, timeout is 2 seconds:

00:35:29: IP: s=172.16.32.1 (local), d=210.210.210.1, len 100, unroutable.
00:35:31: IP: s=172.16.32.1 (local), d=210.210.210.1, len 100, unroutable.
00:35:33: IP: s=172.16.32.1 (local), d=210.210.210.1, len 100, unroutable.
00:35:35: IP: s=172.16.32.1 (local), d=210.210.210.1, len 100, unroutable.
00:35:37: IP: s=172.16.32.1 (local), d=210.210.210.1, len 100, unroutable
Success rate is 0 percent (0/5)
```

Look at the routing table on SanJose2 and check to see whether a Gateway of Last Resort has been defined. Now check the SanJose1 router for the same information.

7. Does either router have a default route or Gateway of Last Resort?
-

Step 7

Establish a default route on SanJose1 and allow it to be advertised within OSPF to SanJose2.

```
SanJose1(config)#ip default-network 210.210.210.0
SanJose1(config)#router ospf 1
SanJose1(config-router)#default-information originate always metric 2500
metric-type 1
```

Use the `show ip route` command on SanJose2 to verify that SanJose1 is advertising the candidate default route by way of OSPF.

```
Gateway of last resort is 172.16.1.1 to network 0.0.0.0

    172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       172.16.32.0/24 is directly connected, Loopback0
```

```
C      172.16.1.0/24 is directly connected, Serial0/1
O      172.16.64.1/32 [110/65] via 172.16.1.1, 00:03:15, Serial0/1
O*E1 0.0.0.0/0 [110/2564] via 172.16.1.1, 00:02:24, Serial0/1
```

From any of the interface IP addresses on SanJose2, ping the 210.210.210.1, ISP1, address. This should now be successful.

```
SanJose2#ping 210.210.210.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 210.210.210.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 80/82/84 ms
```

Also note, that each packet being sent from SanJose2 to ISP1 is also being translated once on the way out and again on the way back as shown below. This is an appropriate area to investigate had we seen unsuccessful ping attempts after having had established the Gateway of Last Resort on SanJose2.

```
SanJose1# (Output from Debug IP NAT while pinging from SanJose2)

01:03:57: NAT: s=172.16.1.2->66.122.33.98, d=210.210.210.1 [20]
01:03:57: NAT: s=210.210.210.1, d=66.122.33.98->172.16.1.2 [20]
01:03:58: NAT: s=172.16.1.2->66.122.33.98, d=210.210.210.1 [21]
01:03:58: NAT: s=210.210.210.1, d=66.122.33.98->172.16.1.2 [21]
01:03:58: NAT: s=172.16.1.2->66.122.33.98, d=210.210.210.1 [22]
01:03:58: NAT: s=210.210.210.1, d=66.122.33.98->172.16.1.2 [22]
01:03:58: NAT: s=172.16.1.2->66.122.33.98, d=210.210.210.1 [23]
01:03:58: NAT: s=210.210.210.1, d=66.122.33.98->172.16.1.2 [23]
01:03:58: NAT: s=172.16.1.2->66.122.33.98, d=210.210.210.1 [24]
01:03:58: NAT: s=210.210.210.1, d=66.122.33.98->172.16.1.2 [24]
01:04:02: NAT: expiring 66.122.33.98 (172.16.1.2) icmp 4548 (4548)
01:04:03: NAT: expiring 66.122.33.98 (172.16.1.2) icmp 4549 (4549)
01:04:03: NAT: expiring 66.122.33.98 (172.16.1.2) icmp 4550 (4550)
01:04:03: NAT: expiring 66.122.33.98 (172.16.1.2) icmp 4551 (4551)
01:04:03: NAT: expiring 66.122.33.98 (172.16.1.2) icmp 4552 (4552)
```