

10

CHAPTER

INTERNET ROUTING

Chapter Outline

- Introduction
- 10-1 Internet Routing—BGP
- 10-2 Configuring BGP
- 10-3 BGP Best Path Selection
- 10-4 IPv6 over the Internet
- 10-5 Configuring BGP on JUNIPER Routers
- Summary
- Questions and Problems

Objectives

- Understand the concept and purpose of the AS number (ASN)
- Understand the concept of the stubby areas
- Examine Internet routing using BGP
- Understand the steps for configuring BGP on a router
- Examine the purpose of the BGP Best Path Selection Algorithm
- Examine the issues of running IPv6 over the Internet using BGP4+
- Understand the issues of configuring BGP on a Juniper router

Key Terms

stubby areas
totally stubby areas
BGP
multi-homed
AS
ASN
EGP
peering
iBGP
eBGP
description *description*
router bgp AS

neighbor *[ip-address]*
remote *[AS-number]*
sh ip bgp sum
BGP Best Path Selection Algorithm
show ip bgp
BGP4+
no bgp default ipv4-unicast
neighbor IPv6_address
remote as AS_Number
neighbor IPv6_address
activate
network IPv6_network

address-family ipv6
show bgp ipv6 unicast summary
autonomous-system *[AS_Number]*
set router-id *[ip_address]*
set group *[BGP_group_name]* **neighbor** *[next_hop_address]*
set group *[BGP_group_name]* **type** *[external/internal]*
family inet6

INTRODUCTION

This chapter introduces the issues of wide-area networks and configuring Border Gateway Protocol (BGP) for Internet routing. The issues of wide-area networks (WAN) are first examined in Section 10-1. The concepts of **stubby areas** and **totally stubby areas** are examined. Connections to WANs and the Internet are similar and examined in this section; however, the routing protocols typically differ. WANs typically use static routes, Open Shortest Path First (OSPF), and Enhanced Interior Gateway Routing Protocol (EIGRP), but Internet routing uses BGP. Section 10-2 introduces the steps for configuring routers to run BGP. This includes the commands to establish and verify that a BGP connection has been established. The section concludes with an opportunity to verify your understanding of configuring BGP using the Net-Challenge Software.

Section 10-3 introduces BGP. It is amazing when you think about all of the possible paths a data packet can take when traveling from a source to a destination on the Internet. BGP addresses this challenge using the Best Path Selection Algorithm to compare the possible routes and select the best route to reach a destination. Section 10-4 looks at running IPv6 over the Internet. BGP4+ is the primary routing protocol for the Internet, and this is examined in this section. This chapter concludes with the steps required for configuring a Juniper router to run BGP.

10-1 INTERNET ROUTING—BGP

This section examines the routing issues for wide-area network (WAN) and Internet routing. WAN connections typically link remote sites and branch offices to the main network. These links usually have slower connection speeds than LAN connections and usually incorporate smaller and less powerful routers. It is critical that the network administrator be aware of these limited resources when choosing a routing protocol for a link to a WAN.

Configuring a WAN Connection

The easiest routing to configure for WAN links is the *static route*. This requires that the main site will need to configure a static route to each subnet at the remote end, as shown in Figure 10-1. At each remote site, you also need to configure a default static route. Each remote site router will have a static route attached to the network's WAN router that goes back to the main network. No routing updates are passed across the link, and the routing table is small and takes up little memory in the router. Static routing works well for single connections to the remote sites, as shown in Figure 10-1. For multiple connections over a WAN connection, you should use a dynamic routing protocol such as OSPF or EIGRP.

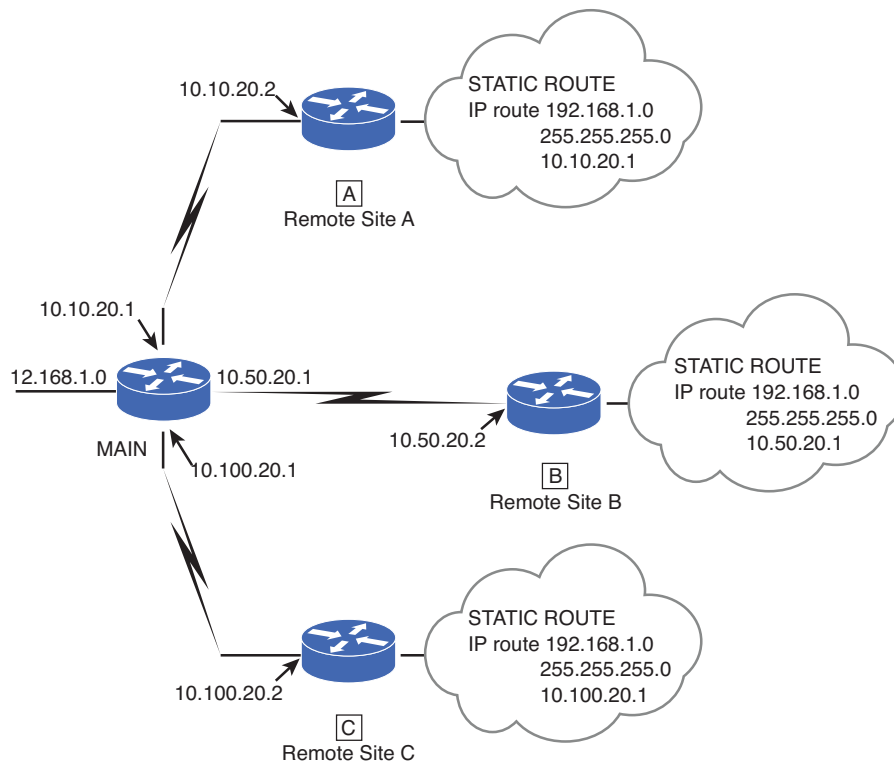


FIGURE 10-1 An example of configuring the static routes for the remote ends on a WAN

When choosing the routing protocol, be cautious about the amount of routing updates traversing the link. Remember that distance vector protocols send the entire routing table at set intervals. Routing Information Protocol (RIP) typically sends routing table updates every 30 seconds. This is an issue when large routing tables are exchanged. In some cases, the exchange of the routing table traffic could consume more than an acceptable amount of data bandwidth. OSPF and EIGRP are more desirable protocols for WAN routing, because they send updates only when routing changes occur.

The size of the router at the remote site will also play a part in the routing protocols you implement. The amount of memory in the remote routers is usually smaller than in the LAN routers. Therefore, the size of the routing table that can be passed to the remote site might need to be smaller than the routing table at the LAN site. Access lists can be used to filter out routes passed to the remote sites. Route filters or access lists are implemented in most modern routers, including Cisco routers. Chapter 7, “Network Security,” examined the procedures for configuring access

Stubby Areas

Stubby areas accept routes from the Internet.

Totally Stubby Areas

Use only a default route to reach destinations external to the autonomous system.

BGP

Border Gateway Protocol.

Multi-Homed

This means the customer has more than one network connection.

lists. Some routing protocols, such as OSPF, have built-in functions to filter out routes. These are called **stubby areas** and **totally stubby areas**. Stubby areas give only inter-area routes and do not accept routes from the external network (that is, routes from the Internet). Totally stubby areas use a default route to reach destinations external to the autonomous system (AS). For a more detailed discussion of route filters and OSPF routing, you should seek out a routing reference book.

Configuring an Internet Connection

Configuring an Internet connection is similar to configuring a WAN connection. The Internet connection can use the same type of link. For example, the link to the Internet could be made using a T1 connection or a high-speed link, such as a DS-3 (44.7 Mbps), OC-3 (155 Mbps), or even MetroEthernet. WAN connections typically connect sites that belong to the same organization, such as a branch office of a business. Internet connections are usually between an Internet service provider (ISP) and its customers. Typically, the ISP and its customers do not use routing protocols, such as OSPF, for the Internet connection because these protocols do not scale well to this type of implementation. Instead, the two main routing protocol options that are available for making the Internet connection to the ISP are static routes and **BGP**.

Static routes are implemented in the same fashion as in the WAN routing section. The procedure for configuring static routes was presented in Chapter 2, “Advanced Router Configuration I.” Static routes are used only when the customer has a single Internet connection. If the customer is **multi-homed**, meaning the customer has more than one Internet connection, BGP is used. The most current version of BGP is version 4. Figure 10-2 provides an example of a single and multi-homed customer.

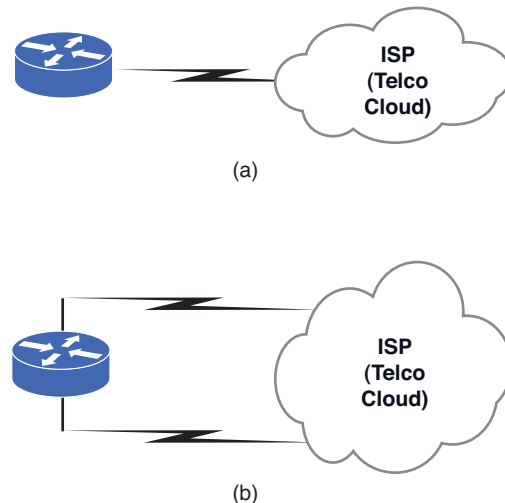


FIGURE 10-2 An example of (a) a single Internet connection and (b) a customer that is multi-homed

BGP is considered to be an external routing protocol. This protocol is designed for routing between separate organizational networks. The BGP term for these networks is *autonomous systems* (**AS**). An AS is assigned an AS number (**ASN**) by the same organization that assigns North American IP addresses, ARIN. The ASN has a different use than the ASN used in such IGP protocols, such as EIGRP. The routing protocols RIP, OSPF, IS-IS, and EIGRP are considered IGPs (Interior Gateway Protocols), while BGP is considered an Exterior Gateway Protocol (**EGP**). The ASN in BGP is used to distinguish separate networks and prevent routing loops. Each router participating in BGP must manually make a **peering** with its BGP neighbor. Peering is an agreement made for the exchange of data traffic between large and small ISPs or, as in this case, between a router and its neighbor router. The agreement on peering is how different networks are joined to form the Internet.

BGP uses TCP as its transport protocol to establish peering and to exchange messages and routes. The network administrator configuring the Internet connection must know the remote IP address and ASN to make this peering. An AS path is created when a network is connected. This is demonstrated in the next subsection.

NOTE

If BGP routers in the same AS peer with each other, that is, have the same ASN, this is called **iBGP**, or internal BGP, whereas the BGP between separate ASs is called **eBGP**, or external BGP. The protocols are collectively referred to as BGP.

ASNs have a set of numbers reserved for private use. These numbers are 64512 through 65535 are not to be propagated to the eBGP world (the Internet). It is the ISP's best practice to remove all private ASs before advertising the BGP routes. The public AS numbers are within the range of 1 to 64511. For this chapter, the configurations presented will use the private ASN numbers.

10-2 CONFIGURING BGP

The objective of this section is to examine the issues associated with configuring an Internet connection. This includes the steps for configuring the router for BGP. The section concludes with the Net-Challenge exercise for configuring a BGP connection to the Internet service provider.

Configuring BGP

This section demonstrates how to configure a router to run BGP for connecting to an ISP. Figure 10-3 illustrates the example network for this section. This is typical of the customer's connection to the ISP. In this example, the steps for configuring both the ISP router (Router-ISP) and the customer's router (Router B) will be shown. In practice, however, the ISP router (Router-ISP) will be configured by the networking personnel with the ISP. This exercise begins with configuring the ISP's router (Router-ISP).

AS

Autonomous systems separate organizational networks.

ASN

Autonomous systems number is used to distinguish separate networks and to prevent routing loops.

EGP

Exterior Gateway Protocol is a routing protocol that's used to exchange routing information between hosts in a network of autonomous systems.

Peering

How an agreement is made for the exchange of data traffic between large and small ISPs or between a router and its neighbor router.

iBGP

Internal Border Gateway Protocol—BGP between the same ASN.

eBGP

External Border Gateway Protocol—BGP between separate ASNs.




FIGURE 10-3 A BGP connection to an ISP

The ISP has assigned an IP address of 192.168.1.1 to the interface that connects to the customer. The customer is assigned an IP address of 192.168.1.2. The first step is to enter the configuration mode on Router-ISP using the **conf t** command. This command is issued from the privilege EXEC prompt, as shown. Next, enter the configuration mode for the router's serial interface that connects to the customer. In this example, serial interface 1/0 (S1/0) on the ISP's router is being used for the connection; therefore, enter the command **interface serial 1/0** [*int serial 1/0*]:

```
Router-ISP# conf t
Router-ISP(config)# int serial 1/0
```

description

The command to enter a description.

Next, enter the IP address for the connection to the customer with a subnet mask of 255.255.255.252. Recall from Chapter 1, "Network Infrastructure Design," that a subnet mask of 255.255.255.252 provides for two usable host IP addresses in the subnet. This is adequate for the connection to the customer. The command **ip address [ip address] [subnet mask]** is used to configure the Serial1/0 interface with the specified IP address and subnet mask. The ISP will typically have multiple customer connections; therefore, it is a good idea to document the IP address entry using the command **description**. This command is used to enter a comment in the router's configuration file about an interface configuration. For example, the serial1/0 interface is being configured for the connection to customer B. The router has been placed in the (*config-if*)# mode. The entry **description Customer A** is entered so that the network administrator of the ISP knows what the serial1/0 interface is being used for. (*Note*: Documentation of the router's configuration file is extremely important.)

```
Router-ISP(config-if)#ip address 192.168.1.1 255.255.255.252
Router-ISP(config-if)#description Customer A
```

router bgp AS

The command entered to run BGP.

The next step is to configure the router to run BGP. The ISP has been assigned the AS number of 65000. (*Note*: This is actually a private AS number. In practice the ISP is assigned a public AS number.) The AS number is used when entering the command for BGP. In this example, the command **router bgp 65000** is entered. The command **network 10.20.20.0 mask 255.255.255.0** follows, which instructs the router to advertise the 10.20.20.0/24 network to its BGP peers. The next command is for specifying the IP address of the BGP neighbor. This is the IP address of the serial0/0 interface on the customer's router (Router B). The format for the command is **neighbor [ip address] remote-as number [neighbor's AS number]**. The last command is for entering a description of the entries. The comment **neighbor 192.168.1.2 description Customer BGP** is used to document the router configura-

tion for the neighbor IP address for the customer and that the routing protocol is BGP:

```
Router-ISP(config)#router bgp 65000
Router-ISP(config-router)#network 10.20.20.0 mask 255.255.255.0
Router-ISP(config-router)#neighbor 192.168.1.2 remote-as 65001
Router-ISP(config-router)#neighbor 192.168.1.2 description Customer A
BGP
```

The command **sh ip int brief** can be used to check the configuration for the interface. This shows that the ISP's router serial1/0 interface has been configured and is connected:

```
Router-ISP#sh ip int brief
Interface          IP-Address  OK?  Method  Status          Protocol
FastEthernet0/0   unassigned YES   unset   administratively
                  down           down
Serial1/0         192.168.1.1 YES   manual  up             up
```

The next example demonstrates the steps for configuring the customer's router (Router B). The customer receives an IP address and subnet mask from the ISP when the connection is requested. This information is used to configure the serial interface that connects to the ISP.

The first step is to enter the router's configuration mode using the **conf t** command from the privilege EXEC prompt, as shown. Next, enter the configuration mode for the router's serial interface that connects to the ISP. In this example, serial interface 0/0 is being used for the connection; therefore, enter the command **int ser 0/0**:

```
RouterB# conf t
Enter configuration commands, one per line. End with CNTL/Z.
RouterB(config)#int ser 0/0
```

Next, configure the serial interface's IP address and subnet mask. In this case, the ISP assigned an IP address of 192.168.1.2 with a subnet mask of 255.255.255.252. (*Note:* It is a good idea to document your entries into the router's configuration by using the command **description (descr)**.) The entries for the IP address and the description are shown:

```
RouterB(config-if)#ip address 192.168.1.2 255.255.255.252
RouterB(config-if)#descr ISP Connection
```

The customer has been assigned the AS number of 65001. (*Note:* This is actually a private AS number. In practice, the customer is assigned a public AS number.) The AS number is used when entering the command for configuring the BGP routing protocol. In this example, the command **router bgp 65001** is entered. The 10.10.10.0/24 network is a network for Customer A. The command **network 10.10.10.0 mask 255.255.255.0** follows, which instructs the router to advertise the Border Gateway Protocol (BGP) over any interfaces with an IP address in the 10.10.10.0/24 network. The next command is for specifying the IP address of the BGP neighbor. This is the IP address of the Serial1/0 interface on the ISP's router. The format for the command is **neighbor [ip-address] remote-as number [neighbor's AS number]**. The last command is for entering a description of the entries.

**neighbor [ip address]
remote-as number**

The next command is for specifying the IP address of the BGP neighbor.

The comment **neighbor 192.168.1.1 descr ISP BGP** is used to document the router configuration. The description identifies the neighbor IP address (the ISP) and comments that the routing protocol is BGP:

```
RouterB(config)#router bgp 65001
RouterB(config-router)#network 10.10.10.0 mask 255.255.255.0
RouterB(config-router)#neighbor ?
A.B.C.D Neighbor address
WORD Neighbor tag
RouterB(config-router)#neighbor 192.168.1.1 remote-as ?
<l-65535> AS of remote neighbor
RouterB(config-router)#neighbor 192.168.1.1 remote-as 65000
RouterB(config-router)#neighbor 192.168.1.1 descr ISP BGP
```

After completing the router configuration, the next step is for the customer to ping the ISP using the IP address provided by the ISP. In this example, the ISP's IP address is 192.168.1.1:

```
RouterB#ping 192.168.1.1
Type escape sequence to abort.
Sending 5 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max 32/35/36 ms
```

sh ip bgp sum

The command used to see whether the routers are exchanging BGP routes.

This test verifies that the routers are connected. The command **sh ip bgp sum** can be used to see whether the routers are exchanging routes.

```
RouterB#sh ip bgp sum
BGP router identifier 192.168.1.2, local AS number 65001
BGP table version is 1, main routing table version 1
1 network entries using 101 bytes of memory
1 path entries using 48 bytes of memory
1 BGP path attribute entries using 60 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
BGP using 209 total bytes of memory
BGP activity 17/16 prefixes, 17/16 paths, scan interval 60 secs

Neighbor      V  AS    MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/
PfxRcd
192.168.1.1   4  65000  41      34       0     0     0  00:00:14  Idle
```

```
RouterB#sh ip bgp sum
BGP router identifier 192.168.1.2, local AS number 65001
BGP table version is 1, main routing table version 1
1 network entries using 101 bytes of memory
1 path entries using 48 bytes of memory
1 BGP path attribute entries using 60 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
```

```
BGP using 209 total bytes of memory
BGP activity 17/16 prefixes, 17/16 paths, scan interval 60 secs
```

```
Neighbor      V  AS    MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/
PfxRcd
192.168.1.1   4  65000  41      34      0     0     00:00:21
Active
```

The first screen shows a status of **Idle**. The second screen shows a status of **Active**. The screens indicate a local AS number of 65002 and a neighbor IP address of 192.168.1.1. The idle and active indicates the BGP adjacency states of the connection. These states are officially known as a BGP finite state machine (FSM) in RFC 1771, Border Gateway Protocol 4. The BGP FSM consists of six states:

1. **Idle state** is the initial state where the resources are being initialized and a BGP router is initiating a TCP connection with its BGP peer.
2. **Connect state** is the next state after the initialization is done. A BGP router waits for a successful TCP connection with its BGP peer.
3. **Active state** is when the TCP connection between BGP peers cannot be established. The term active is misleading in this case, because it does not imply a positive or successful operation. While in this state, it will continue to try to establish a TCP connection with its peer. If it is still unsuccessful, the state will be reset to idle.
4. **Opensent state** indicates a successful TCP connection between the BGP peers. At this point, the BGP peers exchange the open message to establish the peering and check the validity of the peering setup information.
5. **Openconfirm state** is the state where the BGP router is waiting for a BGP keepalive message from its peer. The keepalive message is the indicator the open message is being accepted by its peer.
6. **Established state** is the final state indicating the BGP peering is established and the routing can be exchanged.

Entering the command **sh ip bgp sum** now shows a *I* in the place where **Idle** and **Active** were present. This indicates the router is exchanging routes. The router is now in the established state and one network route is being received from its BGP peer.

```
RouterB#sh ip bgp sum
BGP router identifier 192.168.1.2, local AS number 65001
BGP table version is 5, main routing table version 5
4 network entries using 404 bytes of memory
4 path entries using 192 bytes of memory
4 BGP path attribute entries using 240 bytes of memory
3 BGP AS-PATH entries using 72 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
BGP using 908 total bytes of memory
BGP activity 16/12 prefixes, 16/12 paths, scan interval 60 secs
```

```
Neighbor      V  AS      MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down
State/PfxRcd
192.168.1.1  4  65000   38        31        5      0    0    00:04:00
1
```

show ip bgp neighbor

This command displays BGP neighbor relationship with its peer.

Another good command used to examine the peering information is **show ip bgp neighbor**. This command displays a BGP neighbor relationship with its peer with more details. This command is generally used to troubleshoot a BGP peering issue. An example of using this command is provided. Notice the amount of detail provided with this command.

```
RouterB# sh ip bgp neighbors
BGP neighbor is 192.168.1.1, remote AS 65000, external link
  BGP version 4, remote router ID 192.168.1.1
  BGP state = Established, up for 00:00:11
  Last read 00:00:11, hold time is 180, keepalive interval is 60
seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                Sent          Rcvd
Opens:                6            6
Notifications:       0            0
Updates:              6           16
Keepalives:          49           49
Route Refresh:        0            0
Total:                61           71

Default minimum time between advertisement runs is 30 seconds
```

For address family: IPv4 Unicast

BGP table version 3, neighbor version 3
Index 2, Offset 0, Mask 0x4

	Sent	Rcvd
Prefix activity:	----	----
Prefixes Current:	1	1 (Consumes 48 bytes)
Prefixes Total:	1	1
Implicit Withdraw:	0	0
Explicit Withdraw:	0	0
Used as bestpath:	n/a	1
Used as multipath:	n/a	0

	Outbound	Inbound
Local Policy Denied Prefixes:	-----	-----
Bestpath from this peer:	1	n/a
Total:	1	0

Number of NLRI in the update sent: max 1, min 0

Connections established 6; dropped 5

```
Last reset 00:00:51, due to User reset
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 192.168.1.2, Local port: 38195
Foreign host: 192.168.1.1, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
```

Event Timers (current time is 0x2A88F4):

Timer	Starts	Wakeups	Next
Retrans	4	0	0x0
TimeWait	0	0	0x0
AckHold	2	0	0x0
SendWnd	0	0	0x0
KeepAlive	0	0	0x0
GiveUp	0	0	0x0
PmtuAger	0	0	0x0
DeadWait	0	0	0x0

```
iss: 950556929  snduna: 950557084  sndnxt: 950557084  sndwnd:
16230
irs: 3579589441  rcvnxt: 3579589596  rcvwnd: 16230  delrcvwnd:
154
```

```
SRTT: 124 ms, RTTO: 1405 ms, RTV: 1281 ms, KRTT: 0 ms
minRTT: 16 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle
```

Datagrams (max data segment is 1460 bytes):

```
Rcvd: 4 (out of order: 0), with data: 2, total data bytes: 154
Sent: 5 (retransmit: 0, fastretransmit: 0), with data: 3, total data
bytes: 154
```

There is a lot of information provided by the **show ip bgp neighbor** command. The first section of the output provides the BGP neighbor relation. Specifically, it provides the following:

- The router ID of the neighbor, which can be different than the next hop ip address (router ID of 192.168.1.1 is shown)
- The BGP state (Established state is shown)
- The BGP messages (opens, notifications, updates, keepalives, and route refreshes) being sent and received between the two BGP peers

The second section gives a more in-depth look into the following;

- How many BGP routes or prefixed are being exchanged between the two BGP peers (one prefix was sent and one prefix was received)
- How many BGP routes are considered best paths (one best path was show)

The last section of the output provides TCP connection status and statistics of the two peers. Information provided includes the following:

- How many times the TCP connection is established and dropped (six times since the connection was established and five times since the connection was dropped).
- The reason of the last peering reset (the reason of user reset was shown).
- The information of the TCP port being used for the BGP connection is identified as TCP port 179.
- The rest of this section displays the TCP packet statistics information for the BGP connection.

The command **sh ip route** can be used to examine the customer's routing table, as shown. The table shows that the Border Gateway Protocol (B) is advertising the 10.10.10.0 network via 192.168.1.1. The 192.168.1.0 network is directly attached to the ISP's router. BGP(B) is also running on the customer's 10.20.20.0 network:

```
RouterB#sh ip route
Codes: C connected, S static, I IGRP, R RIP, M mobile, B BGP D
EIGRP, EX EIGRP external, O OSPF, IA OSPF inter area
N1 OSPF NSSA external type 1, N2 OSPF NSSA external type 2
E1 OSPF external type 1, E2 OSPF external type 2, E EGP
i IS-IS, L1 IS-IS level-1, L2 IS-IS level-2, * candidate default
U per-user static route, o ODR T traffic engineered route
Gateway of last resort is not set
10.0.0.0/24 is subnetted, 2 subnets
B      10.20.20.0 [20/0] via 192.168.1.1, 00:03:56
C      10.10.10.0 is directly connected, FastEthernet0/0
      192.168.1.0/30 is subnetted, 1 subnets
C      192.168.1.0 is directly connected, serial0/1
```

The **show running-configuration (sh run)** command can be used to examine the changes made to the router's configuration file. A partial example of the customer's running-configuration file is shown in the following code. Also, don't forget to save the file to NVRAM after the changes have been verified:

```
RouterB# sh run
.
.
interface Serial0/1
  description ISP Connection
  ip address 192.168.1.2 255.255.255.0
  no ip directed-broadcast
!
.
.
router eigrp 100
network 10.10.10.0
!
router bgp 65001
```

```
network 10.10.10.0 mask 255.255.255.0
neighbor 192.168.1.1 remote-as 65000
neighbor 192.168.1.1 description ISP BGP
!
```

When it comes to BGP routing, you must remember an important concept: BGP can only advertise network routes that are installed in the IGP routing table. The network statement seen above under the router BGP section will not mean anything if the network 10.0.0.0 is not in the IGP routing table of the router. So, one must always check the router's IGP table when troubleshooting BGP. As long as IGP routing protocols, such as RIP, OSPF, EIGRP, IS-IS, or even static are used to advertise the network, BGP will synchronize with the IGP routes and advertise them out. In this case, the network 10.10.10.0 is already in the IGP routing table, because it is being advertised via the EIGRP routing protocol.

Networking Challenge: BGP

Use the Net-Challenge Simulator Software included with the text's companion CD-ROM to demonstrate that you can configure BGP for a router connection to an Internet service provider. The network connection is displayed on the screen when the BGP challenge is started. Place the companion CD-ROM in your computer's drive. Open the *Net-Challenge* folder, and click *Net Challenge V3-2.exe*. When the software is running, click the **Select Challenge** button to open a Select Router Challenge drop-down menu. Select **Chapter 10—BGP**. This opens a checkbox that can be used to verify that you have completed all the tasks. This task assumes that a connection to the Internet cloud has already been established:

1. Enter the privileged EXEC mode on the router.
2. Enter the router configuration mode [the **Router(config)#** prompt].
3. Set the hostname to **Border-Router**.
4. Configure the Fast Ethernet 0/0 interface on Border-Router with the following:
IP address: 10.10.1.2
Subnet mask: 255.255.255.0
5. Enable the router's Fast Ethernet 0/0 interface.
6. Configure the Serial0/0 interface with the following:
IP address: 192.168.1.2
Subnet mask: 255.255.255.0.
7. Enable the router's Serial0/0 interface.
8. Use the router's **description** command (**descr**) to indicate that this interface is the ISP connection. (*Note:* The text ISP Connection is case-sensitive.)

9. Enable BGP on the router with an AS number of 65002.
10. Configure Border-Router's BGP neighbor with a remote AS of 65001.
11. Configure a BGP route to the 10.0.0.0 network.
12. Use the **show ip route** command to verify that the route from Border-Router to the ISP is configured.
13. Use the **ping** command to verify that the 192.168.1.1 interface is connected.
14. Use the **sh run** command to view the running-configuration file on Border-Router. Verify that BGP is enabled, the description **ISP Connection** has been entered, and the proper network address is specified for the ISP connection.
15. Use the **sh ip int brief** command to check the interface status.

10-3 BGP BEST PATH SELECTION

It is amazing when you think about all the possible paths a data packet can take when traveling from a source to a destination on the Internet. It is equally amazing that the path to a destination can be quickly determined. Selecting the best path to reach a destination is an incredible challenge. This is the challenge addressed by BGP.

BGP Best Path Selection Algorithm

Used to compare the possible routes and to select the best route.

The BGP routing protocol enables computers to connect to other computers in millisecond time to any location around the world. BGP does this by establishing neighbor connections. Once this is done, each neighbor router shares its routes. It has been estimated that there are upward of 400,000 possible routes on the Internet. BGP works differently than IGP's because it does not make routing decisions based on best path metrics. Instead, BGP is a path vector routing protocol, which makes routing decisions based on path attributes that are added to routing information. BGP uses these attributes to make a routing decision and select the best path. It allows an AS to control traffic flow using multiple BGP attributes. More often than not, BGP routers learn multiple paths or routes to the same destination. The **BGP Best Path Selection Algorithm** is used to compare the possible routes and to select the best route. BGP will assign the first valid path as the best path. It will then compare this path with the next path until it reaches the end of the valid paths.

The following list provides the rules in the order of how the path selection is made. These rules are more Cisco-centric, but they generally apply to most of the BGP routers of other vendors:

1. **Weight:** This is a Cisco's proprietary attribute that is assigned locally to your router. This attribute only applies to Cisco routers. This attribute only affects the local router. In situations where there are multiple routes, BGP will select the route with the highest weight.
2. **Local Preference:** This indicates which path to the AS has local preference. The local preference is shared among the iBGP routers within the same AS. Local preferences are widely used to influence the route or path selection as it

propagates throughout the AS. The highest preference is preferred. The default is 100.

3. **Self-Originated:** The routes that are locally originated by BGP running on this router would always be preferred over routes learned from other routers or other BGP networks.
4. **Shortest AS_PATH:** This selection is made based on which route has the shortest AS path. Every route records every AS it has to go through. This makes up an AS path. Every BGP router examines the path of autonomous systems the route has to take before reaching the BGP router. The shortest AS path is preferred. This is used by BGP only when there is a “tie” when comparing weight, local preference, and locally originated vs. aggregate addresses.
5. **Lowest origin type:** The origin attribute that defines the origin of the path. If the path is originated from protocols such as Interior Gateway Protocol (IGP), then it has lower preference to the Exterior Gateway Protocol (EGP).
6. **Lowest multi-exit discriminator (MED):** This is an external metric of a route. The metric with a lower MED value is preferred over a higher value. The MED value is another widely used attribute to manipulate route selection. Unlike local preference that influences the routing decision within the local AS, the MED value is used to influence the routing decision of an adjacent AS.
7. **eBGP over iBGP:** This simply states that the paths learned via eBGP is preferred over the paths learned via iBGP, which is a way to prevent the loop.
8. **Lowest IGP metric:** The path with the lowest IGP metric is preferred for the BGP next hop.
9. **Multiple paths:** This is used to determine if multiple paths require installation in the routing table.
10. **External paths:** If two paths are external, then BGP prefers the path that was received first.
11. **Lowest router ID:** The route that comes from the BGP router with the lowest router ID (highest IP address on the router) is preferred.
12. **Minimum cluster list:** If the router ID or originator is the same for multiple paths, the path with the minimum cluster list length is preferred.
13. **Lowest neighbor address:** The path that comes from the lowest neighbor address is preferred.

A command that is very useful in examining BGP routes and their attributes is **show ip bgp**. This command by itself displays the summary contents of the BGP routing table. Some of the BGP attributes are also displayed. The following example shows the output of the **show ip bgp** command on Customer’s Router B:

```
RouterB#sh ip bgp
BGP table version is 5, local router ID is 192.168.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
                r RIB-failure, S Stale
```

show ip bgp

Displays the summary contents of the BGP routing table.

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.20.20.0/24	192.168.1.1	0			
0 65000 i					
*> 192.168.30.0	192.168.1.1				
0 65000 65002 i					
*> 192.168.200.0	192.168.1.1				
0 65000 65002 65111 i					

This time, Router B is receiving more BGP routes from its peer, 192.168.1.1. There are three BGP routes. The * and > indicate that these BGP routes are valid and their paths are the best path to reach these networks. The Next Hop column shows the BGP next hop or peer. The Metric column shows MED value, if it exists. The LocPrf column displays the BGP local preference value, if it is manually configured. The default value is 100. Weight is the Cisco's BGP weight attribute. The result shows that the network 10.20.20.0/24 is from AS 65000. The network 192.168.30.0 was originated from an IGP within the AS 65002 and then it went through AS 65000 before it reaches Router A. The network 192.168.200.0 has the longest AS path of them all. It has to go through three ASs before reaching Router A. The command **sh ip bgp** can be used to display more specific BGP information, such as a specific network prefix or route, as demonstrated in the following example:

```
RouterB#sh ip bgp 192.168.200.0
BGP routing table entry for 192.168.200.0/24, version 5
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
    65000 65002 65111
      192.168.1.1 from 192.168.1.1 (192.168.1.1)
        Origin IGP, localpref 100, valid, external, best
```

The **sh ip bgp 192.168.200.0** command examines a specific prefix. Instead of having to interpret this route from the output of **sh ip bgp**, this gives a more explanatory version of the route.

10-4 IPV6 OVER THE INTERNET

This section addresses the issue of running IPv6 over the Internet. This section also covers BGP4+, which allows BGP to run IPv6 over the Internet. An example on how to configure a Cisco router for running BGP over the Internet is also presented. Troubleshooting IPv6 and BGP are also presented.

BGP4+

The version of BGP for running IPv6.

The Internet routing is dominated by the BGP routing protocol. This is true for both IPv4 and IPv6 routing over the Internet. The current version of BGP that is used by IPv4 is BGP4. The multiprotocol BGP extensions or **BGP4+** allows BGP4 to be used for IPv6. BGP4+ for IPv6 supports the same features and functionality as IPv4

BGP as well as the additional support for the IPv6 address family and the IPv6 address for the BGP next hop.

The steps to configure IPv6 BGP on Cisco routers are similar to the steps used in IPv4. The first step is to configure the interface on the router that will run IPv6. This requires that the next hop IPv6 address or its IPv6 BGP peer must be reachable by the router. The router BGP process must be configured by issuing the command **router bgp AS_Number**. This is not required if a BGP process already exists for IPv4, since the same BGP process will be used.

The **no bgp default ipv4-unicast** command is used to allow protocols other than IPv4 to be activated within the multiprotocol BGP (BGP4+). By default, only the IPv4 unicast is enabled. Next, an IPv6 peer is specified using the command **neighbor IPv6_address remote as AS_Number**. Within the multiprotocol BGP, a protocol must be specified using the command **address-family**. This command is used to specify that IPv6 is selected. Specifically, for IPv6, the command **address-family ipv6** is used.

Within the address family group, the BGP peering with the neighbor can be established using the command **neighbor IPv6_address activate**. Also, inside the same group, one must specify the IPv6 networks that will be advertised to its peer with the command **network IPv6_network**. The following example demonstrates a sample configuration of an IPv6 section of a Cisco router.

```
!
router bgp 65203
  no bgp default ipv4-unicast
  neighbor 2001:DB8:1:128::2 remote-as 65200
  neighbor 2001:DB8:1:128::2 description ISP
!
  address-family ipv6
    neighbor 2001:DB8:1:128::2 activate
    neighbor 2001:DB8:1:128::2 soft-reconfiguration inbound
    network 2001:D00::/32
  exit-address-family
!
```

The concepts for troubleshooting IPv6 BGP are still the same as in IPv4 BGP. The peering relationship with the BGP neighbor has to be established before any IPv6 routes can be exchanged. In IPv4 BGP, commands like **sh ip bgp sum** and **sh ip bgp neighbor** were used. However, the IPv6 BGP commands are different in Cisco routers. The IPv6 BGP commands start with the syntax **show bgp ipv6 [unicast/multicast]**. The type **unicast** or **multicast** must be specified when using the command. The following is the command to view the summary status of the IPv6 BGP neighbor using the command **show bgp ipv6 unicast summary**. This is analogous to the command **show ip bgp summary** in IPv4 that was used to show the summary status of the BGP peering. The output displays almost the same information as what would be in the IPv4 command with the exception of the neighbor IP address, which is shown in IPv6 format:

```
RouterA#sh bgp ipv6 unicast summary
BGP router identifier 10.100.100.100, local AS number 65203
```

no bgp default ipv4-unicast

Used to allow protocols other than IPv4 to be activated within the multiprotocol BGP (BGP4+).

neighbor IPv6_address remote as AS_Number

Used to specify a BGP peer.

address-family ipv6

The command used to specify that IPv6 is specified.

neighbor IPv6_address activate

Used to establish BGP peering with the neighbor.

network IPv6_network

Used to specify the IPv6 networks that will be advertised to its peer.

show bgp ipv6 unicast summary

Command to view the summary status of the IPv6 BGP neighbor.

```

BGP table version is 1843311, main routing table version 1843311
8356 network entries using 1178196 bytes of memory
8356 path entries using 635056 bytes of memory
5835/5831 BGP path/bestpath attribute entries using 933600 bytes of
memory
5649 BGP AS-PATH entries using 232282 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
BGP using 2979134 total bytes of memory
BGP activity 202608/194251 prefixes, 242511/234155 paths, scan
interval 60 secs

```

```

Neighbor          V          AS MsgRcvd MsgSent  TblVer  InQ OutQ
Up/Down  State/PfxRcd
2001:DB8:1:128::2 4          65200 1773365 219989   1843311  0    0
3d13h      8355

```

The IPv6 routes can be examined after the IPv6 BGP peering is successfully established. The same rules for the BGP path algorithm still apply for IPv6. The following is a sample of the IPv6 BGP routes received by Router B from its BGP neighbor: 2001:DB8:1:128::2. The command used to accomplish this task is **sh bgp ipv6 unicast**:

```

RouterB#sh bgp ipv6 unicast
BGP table version is 1843542, local router ID is 10.100.100.100
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
                r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

```

      Network          Next Hop          Metric LocPrf Weight
Path
*> ::/0                2001:DB8:1:128::2          0 65200 i
*> 2001::/32           2001:DB8:1:128::2          0 65200 65152
65164 65939 i
*> 2001:200::/32       2001:DB8:1:128::2          0 65200 65152
65153 65388 65660 65500 i
*> 2001:200:900::/40   2001:DB8:1:128::2          0 65200 65152
65153 65537 65660 65660 65660 i
*> 2001:200:C00::/40   2001:DB8:1:128::2          0 65200 65152
651164 65725 65607 65530 i
*> 2001:200:C000::/35  2001:DB8:1:128::2          0 65200 65152
65153 65388 65660 65500 65634 i
*> 2001:200:E000::/35  2001:DB8:1:128::2          0 65200 65152
65153 65388 65660 i
*> 2001:208::/32       2001:DB8:1:128::2          0 65200 65401
65610 65472 i
*> 2001:208:1:1::/64   2001:DB8:1:128::2          0 65200 65401
65610 65472 i

```

The command **sh bgp ipv6 unicast ?** can be used to list more options for looking at more BGP-specific information, such as prefixes or neighbors. The following shows the specific options that can be specified:

```
ict-gate#sh bgp ipv6 unicast ?
X:X:X:X::X/<0-128> IPv6 prefix <network>/<length>
community          Display routes matching the communities
community-list     Display routes matching the community-list
dampening          Display detailed information about dampening
extcommunity-list  Display routes matching the extcommunity-list
filter-list        Display routes conforming to the filter-list
inconsistent-as   Display only routes with inconsistent origin
ASs
injected-paths     Display all injected paths
labels             Display BGP labels for prefixes
neighbors          Detailed information on TCP and BGP neighbor
connections
paths              Path information
peer-group         Display information on peer-groups
pending-prefixes   Display prefixes pending deletion
prefix-list        Display routes matching the prefix-list
quote-regexp      Display routes matching the AS path "regular
expression"
regexp             Display routes matching the AS path regular
expression
replication        Display replication status of update-group(s)
rib-failure        Display bgp routes that failed to install in
the routing table (RIB)
route-map          Display routes matching the route-map
summary           Summary of BGP neighbor status
update-group       Display information on update-groups
|                 Output modifiers
<cr>
```

10-5 CONFIGURE BGP ON JUNIPER ROUTERS

So far, the examples shown in this chapter are for Cisco routers. However, the BGP concepts discussed so far can also be applied to any type of router. BGP is a universal Internet Protocol, and most of the BGP Internet routers comply with the same standard for optimum interoperability. The steps to configure BGP on Juniper routers are similar to Cisco's implementation. First and foremost, the network interface between the router and its BGP neighbor must be configured. This requires that the interface connection has to be up and operational. After the interface connection is established, the BGP configuration can begin.

The first step of the BGP configuration is to define the BGP AS of the router. On Cisco routers, this is accomplished with the command **router bgp AS_Number**. On

Juniper routers, the BGP AS is defined within the **routing-options** section. The following section demonstrates the configuration steps:

```
net-admin@j-router> configure
[edit]
admin@j-router# edit routing-options
[edit routing-options]
admin@j-router# set autonomous-system 65555
[edit routing-options]
admin@j-router# set router-id 172.20.1.1
```

autonomous-system **[AS_Number]**

This command is used in JUNOS to define the BGP AS for the router.

set router-id **[ip_address]**

This command is used in JUNOS to define the BGP router BGP identifier.

The command set **autonomous-system [AS_Number]** is used to define the BGP AS for the router. The command **set router-id [ip_address]** is used to define the BGP router BGP identifier. This command is optional, just like with Cisco routers. If this command is not specified, the router will use the highest IP address of the loopback addresses or the highest IP address of the interface addresses as the router ID.

The next step of BGP configuration is to define the BGP neighbor or peer. This is configured under the protocol section rather than the routing-options section. The following section demonstrates the configuration steps:

```
[edit]
admin@j-router# edit protocols bgp
[edit protocols bgp]
admin@j-router# set group myISP neighbor 10.200.200.2
[edit protocols bgp]
admin@j-router# set group myISP neighbor 10.200.200.2 description
"Connection to my ISP"
[edit protocols bgp]
admin@j-router# set group myISP neighbor 10.200.200.2 peer-as 65000
[edit protocols bgp]
admin@j-router# set group myISP type external
[edit protocols bgp]
admin@j-router# show
group myISP {
    type external;
    neighbor 10.200.200.10 {
        description "Connection to my ISP";
        peer-as 65000;
    }
}
```

set group [BGP_ **group_name]** **neighbor [next_hop_** **address]**

Used to define the BGP neighbor on a Juniper router.

In JUNOS, the BGP neighbors are placed within a group, and a group can contain one or multiple BGP neighbors. The routing policy or BGP configuration can be applied to a group of BGP neighbors. The command **set group [BGP_group_name] neighbor [next_hop_address]** is used to define the BGP neighbor. The command can be extended with keywords such as **peer-as** to specify the AS number of the neighbor and **description** to provide short descriptive information. In JUNOS, one

must specify the type of BGP (external or internal); therefore, the command **set group [BGP_group_name] type [external/internal]** is used. To verify the BGP configuration, the **show** command displays the configuration under the protocol bgp section.

At this point, the Juniper router is ready to establish external BGP peering with its ISP neighbor router. If the BGP peering is established successfully, it will receive BGP routes from its neighbor. However, no network prefixes will be advertised from this router yet. On Cisco routers, one must issue a network statement command to specify a network prefix that is to be advertised over BGP. This is also required on Juniper routers; however, this will not be a one-line command as in Cisco IOS. In JUNOS, this is done via a routing policy. The routing policy must be defined to include the network prefixes to be advertised over BGP. Then, it must be applied within the BGP configuration with an export statement. The following is a sample configuration of a routing policy to define network prefixes:

```
policy-options {
  prefix-list network-for-BGP {
    192.168.100.0/24;
    192.168.200.0/24;
    172.16.100.0/24;
  }
  policy-statement Export-BGP {
    term local-prefixes {
      from {
        prefix-list network-for-BGP;
      }
      then {
        accept;
      }
    }
    then reject;
  }
}
```

Then, the routing policy must be applied within the BGP configuration with an export statement. The following example shows the routing policy **Export-BGP** being applied to the BGP group **myISP**. This enables the network prefixes of 192.168.100.0/24, 192.168.200.0/24, and 172.16.100.0/24 to be advertised over BGP:

```
[edit]
admin@j-router# edit protocols bgp
[edit protocols bgp]
admin@j-router# set group myISP export Export-BGP
```

This completes the basic BGP configuration for Juniper routers. On Cisco routers, one would use a command like **show ip bgp sum** to verify the summary status of its BGP neighbors after the BGP configuration is complete. A similar command

set group [BGP_group_name] type [external/internal]

Used in JUNOS to specify the type of BGP (external or internal).

exists in JUNOS as well. The command is **show bgp sum** and its output is shown here:

```
admin@j-router> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0      10      10      0      0      0      0
Peer          AS  InPkt  OutPkt  OutQ  Flaps Last  Up/Dwn
State|#Active/Received/Damped...
10.200.200.10 65000 53138  53122   0     3     2w4d10h
4/4/0          10/10/0
```

The output shows that the connection with the BGP neighbor of 10.200.200.10 of the AS number 65000 is established. There are ten prefixes received and active from its peer. The command **show bgp neighbor** can be used for an in-depth view of the neighbor status and statistics similar to the command **show ip bgp neighbor** in Cisco IOS. Another useful command for BGP troubleshooting is **show route protocol bgp**. This command shows the BGP routing table. These are the routes received via BGP and their BGP attributes. This is equivalent to the command **show ip bgp** in Cisco IOS. The following is sample output from the **show route protocol bgp** command:

```
admin@j-router> show route protocol bgp

inet.0: 378081 destinations, 378141 routes (378071 active, 8 holddown,
2 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[BGP/10] 1d 11:29:00, MED 0, localpref 100
                   AS path: 65000 I
                   > to 172.20.1.1 via ge-0/0/0.0
10.10.4.0/22       *[BGP/170] 2d 06:12:12, MED 0, localpref 90, from
10.200.200.10
                   AS path: 65000 65215 65164 65323 65545 65620 I
                   > to 172.20.1.1 via ge-0/0/0.0
10.10.16.0/23      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10
                   AS path: 65000 65215 65516 65519 I
                   > to 172.20.1.1 via ge-0/0/0.0
10.10.18.0/23      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10
                   AS path: 65000 65215 65516 65519 I
                   > to 172.20.1.1 via ge-0/0/0.0
10.10.20.0/23      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10
                   AS path: 65000 65215 65516 65519 I
                   > to 172.20.1.1 via ge-0/0/0.0
10.10.22.0/23      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10
                   AS path: 65000 65215 65516 65519 I
                   > to 172.20.1.1 via ge-0/0/0.0
```

```

10.10.25.0/24      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10

    AS path: 65000 65215 65516 65519 I
> to 172.20.1.1 via ge-0/0/0.0
10.10.26.0/23      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10

    AS path: 65000 65215 65516 65519 I
> to 172.20.1.1 via ge-0/0/0.0
10.10.28.0/22      *[BGP/170] 1w4d 07:58:48, MED 0, localpref 90,
from 10.200.200.10

    AS path: 65000 65215 65516 65519 I
> to 172.20.1.1 via ge-0/0/0.0
10.10.64.0/18      *[BGP/170] 1w4d 07:59:01, MED 0, localpref 90,
from 10.200.200.10

    AS path: 65000 65215 65164 65725 65670 65144 I
> to 172.20.1.1 via ge-0/0/0.0

```

The * before the BGP information indicates that the path is the best BGP path. This command displays the BGP attributes such as AS_path, BGP local preference value, and the Multi-Exit Discriminator value.

Configuring IPv6 BGP on Juniper routers is not complicated. The following is an example IPv6 BGP configuration. It is similar to how IPv4 BGP is configured. The configuration is placed inside a BGP group. The type of BGP is defined as external. The big differences are the use of IPv6 as the BGP neighbor next-hop address and the addition of **family inet6**. The **family inet6** enables the IPv6 family address. Within the **family inet6**, there is an option for **unicast**, **multicast**, or any (both). In this example, the option specifies support for both **unicast IPv6** and **multicast IPv6**.

family inet6

JUNOS command to enable the IPv6 family address.

```

protocols {
  bgp {
    group myISP6 {
      type external;
      export Export-BGP6;
      family inet6 {
        any;
      }
      neighbor 2001:db8:feed:beef::1 {
        description "Connection to my ISP";
        peer-as 65000;
      }
    }
  }
}

```

Also, you must keep in mind that the routing policy applied to this BGP group, Export-BGP6, must be written to include IPv6 prefixes instead of IPv4 prefixes. The example policy statement, Export-BGP6, is slightly different than the one shown in the IPv4 example because it is using a route-filter statement not prefix-list. Both will accomplish the same goal, which is to advertise the prefixes into BGP. This

example shows the route-filter is used to filter prefixes of 2001:db8:11::/48 and 2001:db8:22::/48. Only these prefixes will be exported into BGP routing:

```
policy-options {
  policy-statement Export-BGP6 {
    term local-prefixes {
      from {
        route-filter 2001:db8:11::0/48 exact;
        route-filter 2001:db8:22::0/48 exact;
      }
      then {
        accept;
      }
    }
    then reject;
  }
}
```

On Juniper routers, the same BGP troubleshooting commands, such as **show bgp summary**, **show bgp neighbor**, or **show route protocol bgp**, can be used for both IPv4 and IPv6.

SUMMARY

This chapter presented the fundamentals of wide-area networking. The student should understand and appreciate the role that the PSTN (public switched telephone network—telco) plays in wide-area networking. This chapter only introduced a fraction of the technologies and issues needed to be understood by the network administrator. However, this chapter addressed the fundamental or base knowledge needed for a networking administrator to start working in this field.

The student should understand the following:

1. The concept and purpose of the AS number (ASN)
2. The concept of the stubby areas
3. Internet routing using BGP
4. The steps for configuring BGP on a router
5. The purpose of the BGP Best Path Selection Algorithm
6. The issues of running IPv6 over the Internet using BGP4+
7. The issues of configuring BGP on a Juniper router

QUESTIONS AND PROBLEMS

Section 10-1

1. What is the purpose of a wide-area network connection?
2. What is the easiest routing protocol to use for WAN links? What if there are multiple connections to the remote sites?
3. Define the following:
 - a. Stubby areas
 - b. Totally stubby areas
4. A multi-homed customer has
 - a. A single Internet connection
 - b. More than one Internet connection
 - c. Static routes
 - d. None of these answers are correct
5. BGP is considered to be
 - a. An external routing protocol
 - b. Used for routing between the same networks
 - c. Outdated

6. Each router participating in BGP manually does what?
 - a. Makes an AS path
 - b. Makes a peering with its BGP neighbor
 - c. Sets the remote router's IP address
 - d. Sets the remote router's ASN
7. The network administrator must know this when configuring the Internet connection.
 - a. The local address and the ASN
 - b. The loopback address and the ASN
 - c. The remote-as and the ASN
 - d. The remote IP address and ASN

Section 10-2

8. The router command used for entering a description is
 - a. Comment
 - b. !
 - c. *
 - d. Description
9. What does the router command **router bgp 65003** mean?
10. Write the router command for specifying the IP address of the BGP neighbor. Also, show the router prompt.
11. How many IP addresses are provided on each subnet with a 255.255.255.252 subnet mask?
12. List the command that can be used to enter a description that identifies the 10.10.200.2 interface as the connection to the ISP. List the command and the router prompt.
13. The following command is entered on the router:

```
Router-ISP#sh ip int brief
Interface      IP-Address  OK?  Method  Status  Protocol
FastEthernet0/0 unassigned YES  unset  administratively down
down
```

What does it mean for the protocol to be down?

- a. The routing protocol has not been configured.
- b. BGP routing is not enabled.
- c. The interface is not communicating with another networking device.
- d. All of these answers are correct.

14. What is the purpose of the **sh ip bgp sum** command?
15. Which of the following is the code for BGP routes if the command **show ip route** is entered?
- a. O
 - b. B
 - c. C
 - d. D

16. The command **show ip route** is entered on a router. Describe what the following information provides:

```
Gateway of last resort is not set
B 192.10.20.0/24 [20/0] via 192.168.10.1, 00:03:25
172.16.0.0/16 is subnetted, 1 subnets
D 172.16.0.0 is directly connected, Serial1
C 192.168.10.0/24 is directly connected, Serial0
```

17. The following information is displayed after entering the **show running-config** command:

```
router bgp 65020
network 192.168.20.0
neighbor 10.10.100.1 remote-as 65010
neighbor 10.10.100.1 description ISP BGP
```

What is the AS for the neighbor?

What network(s) is(are) running BGP?

18. BGP can only advertise network routes that are installed where on a router?
19. List the six states of the BGP finite state machine in the order of their operation?
20. Match the term with the correct BGP state.

While in this state, the router will continue trying to establish a TCP connection.

- a. Idle State
 - b. Active State
 - c. Established State
 - d. Connect State
21. Match the term with the correct BGP state.
- This state indicates a successful TCP connection between the BGP peers.
- a. Idle State
 - b. Active State
 - c. Opensent State
 - d. Openconfirm State

22. Match the term with the correct BGP state.

This state is where a BGP router is initiating a TCP connection with its BGP peer.

- a. Idle State
- b. Active State
- c. Openconfirm State
- d. Connect State

23. Match the term with the correct BGP state.

This state indicates the BGP peering is established.

- a. Idle State
- b. Active State
- c. Established State
- d. Connect State

24. Match the term with the correct BGP state.

In this state the router is waiting for the keepalive message.

- a. Idle State
- b. Active State
- c. Opensent State
- d. Openconfirm State

25. Match the term with the correct BGP state.

This is the next state after the initialization is done:

- a. Idle State
- b. Active State
- c. Openconfirm State
- d. Connect State

26. The **sh ip bgp sum** command is issued and a 1 shows in the place where Idle and Active were displayed. What does this indicate?

27. What command is generally used to troubleshoot a BGP peering issue?

Section 10-3

28. What is the purpose of the best path selection algorithm?

29. What does BGP use to make routing decisions?

30. What path attributes does the BGP Best Path Algorithm use to make its route selection?

31. In regards to the Cisco rules that are used to select the best path, match the term to the definition.

1. Shortest AS_PATH	a. This indicates which path to the AS has local preference.
2. Lowest Origin Type	b. This is used by BGP only when there is a "tie."
3. Local Preference	c. This is the first rule for selecting the best path.
4. Weight	d. That the BGP AS Path prefers eBGP over iBGP.
5. eBGP over iBGP	e. Protocols, such as Interior Gateway Protocol (IGP), have lower preference to the Exterior Gateway Protocol (EGP).

32. BGP is a path vector routing protocol. What does this mean?

33. What command is used to examine BGP routes and their attributes?

34. How do you know if BGP routes are valid?

35. How do you know if BGP routes are the best path to reach a network?

36. The **sh ip bgp** command shows that following routes to multiple networks:

```
*> 192.20.20.0/24    192.168.10.1          0
0 65000 i
*> 10.20.30.0       192.168.10.1
0 65000 65002 i
*> 10.10.20.0       192.168.10.1
0 65000 65002 65111 i
```

How can you modify the **sh ip bgp** command so that only the information for the 10.20.30.0 network is displayed?

37. The **sh ip bgp** command is issued. Where is the BGP peer information displayed?

Section 10-4

38. What version of BGP is used to run IPv6 over the Internet?

39. What is the command for configuring the BGP process for IPv6 on a router?

40. What command is issued to allow protocols other than IPv4 to be activated within the multiprotocol BGP.

41. What is the purpose of the neighbor IPv6_address remote as AS_Number command?

42. How is a protocol specified in the multiprotocol BGP?

43. How is IPv6 specified in BGP?

44. What information does the following command display?

```
RouterA#sh bgp ipv6 unicast summary
```

45. List the command that is used to establish BGP peering.

46. What is the purpose of the command **network IPv6_network**?

Section 10-5

47. How is the BGP AS configured on a Juniper router? List the required configuration steps. Assume an AS of 65000 and a router ID of 192.168.20.10.
48. What happens if the **set router-id** command is not issued on a Juniper router?
49. What is this information showing?

```
admin@j-router# set group myISP neighbor 10.200.200.2 peer-as 65000
```

50. How is a network prefix advertised on a Juniper router?
51. What is this information showing on a Juniper router?

```
policy-options {  
  prefix-list network-for-BGP {  
    172.16.100.0/24;  
    192.168.100.0/24;  
  }  
}
```

52. What does the following command do on a Juniper router?

```
admin@j-router# set group myISP export Export-BGP
```

53. The command **is show bgp sum** is issued on a Juniper router and its output is shown here. Is a connection to AS 65300 established? How do you know?

```
admin@j-router> show bgp summary  
Groups: 1 Peers: 1 Down peers: 0  
Table Tot Paths Act Paths Suppressed History Damp State Pending  
inet.0 10 10 0 0 0 0 0  
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn  
State| #Active/Received/Damped...  
192.100.10.10 65300 45138 45122 0 3 2w4d10h  
4/4/0 8/8/0
```

54. How is external or internal BGP specified in JUNOS?
55. How is a BGP neighbor defined in JUNOS?
56. What command is used in Juniper to display the BGP routing table?
57. What does the * mean for the following information?

```
10.10.4.0/22 * [BGP/170] 2d 06:12:12, MED 0, localpref 90,  
from 10.200.200.10
```

58. What command is used to specify support for both unicast IPv6 and multicast IPv6?

Critical Thinking

59. Answer the following questions for the following information:

```
RouterA#sh ip bgp  
BGP table version is 5, local router ID is 192.168.1.2  
Status codes: s suppressed, d damped, h history, * valid, > best,  
i - internal,  
r RIB-failure, S Stale  
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network Weight	Path	Next Hop	Metric	LocPrf
*> 192.20.20.0/24		192.168.10.1	0	0
65100	i			
*> 10.20.30.0		192.168.10.1	0	65100
65102	i			
*> 10.10.20.0		192.168.10.1	0	65100
65102	65111 i			

- Is the route to the 10.10.20.0 network valid? How do you know?
 - What AS is network 192.20.20.0 from?
 - What network did 10.20.30.0 originate from?
 - Which network has the longest path? How many ASs does it go through? Identify the ASs.
60. What is the following information showing?

```
router bgp 65203
no bgp default ipv4-unicast
neighbor 2001:DB8:1:128::2 remote-as 65200
neighbor 2001:DB8:1:128::2 description ISP
```

61. What command is used to verify that BGP is exchanging routes? How does this change for IPv6?
62. What is a *cloud*?
63. The running configuration for a router shows the following information. What does this indicate? What is soft-reconfiguration inbound?

```
neighbor 2001:DB8:1:128::2 soft-reconfiguration inbound
```

64. What type of IPv6 address is this and how do you know? What is the purpose of the double colon (::)?
- 2001:DB8:1:128::2
65. The following information is displayed on Cisco router. What information is being displayed for the IPv6 settings?

```
address-family ipv6
neighbor 2001:DB8:1:128::2 activate
neighbor 2001:DB8:1:128::2 soft-reconfiguration inbound
network 2001:D00::/32
exit-address-family
```