CISCO

Official Cert Guide

Learn, prepare, and practice for exam success

CCNA Routing and Switching ICND2 200-105

Academic Edition

ciscopress.com

WENDELL ODOM, CCIE® No. 1624

In addition to the wealth of updated content, this new edition includes a series of free hands-on exercises to help you master several real-world configuration and troubleshooting activities. These exercises can be performed on the CCNA ICND2 200-105 Network Simulator Lite software included for free on the DVD or companion web page that accompanies this book. This software, which simulates the experience of working on actual Cisco routers and switches, contains the following 19 free lab exercises, covering all the topics in Chapters 10 and 11 on EIGRP IPv4 configuration and troubleshooting.

1. EIGRP Serial Configuration I
2. EIGRP Serial Configuration II
3. EIGRP Serial Configuration III
4. EIGRP Serial Configuration IV
5. EIGRP Serial Configuration V
6. EIGRP Serial Configuration VI
7. EIGRP Route Tuning I
8. EIGRP Route Tuning II
9. EIGRP Route Tuning III
10. EIGRP Route Tuning IV
11. EIGRP Neighbors I
12. EIGRP Neighbors II
13. EIGRP Neighbors III
14. EIGRP Auto-Summary Configuration Scenario
15. EIGRP Configuration I Configuration Scenario
16. EIGRP Metric Manipulation Configuration Scenario
17. EIGRP Variance and Maximum Paths Configuration Scenario
18. EIGRP Troubleshooting Scenario
19. Path Troubleshooting Scenario IV

If you are interested in exploring more hands-on labs and practicing configuration and troubleshooting with more router and switch commands, check out our full simulator product offerings at http://www.pearsonitcertification.com/networksimulator.

CCNA ICND2 Network Simulator Lite minimum system requirements:

Windows (Minimum)

- Windows 10 (32/64-bit), Windows 8.1 (32/64-bit), or Windows 7 (32/64-bit)
- 1 gigahertz (GHz) or faster 32-bit (x86) or 64-bit (x64) processor
- 1 gigabyte (GB) RAM (32-bit) or 2 GB RAM (64-bit)
- 16 GB available hard disk space (32-bit) or 20 GB (64-bit)
- DirectX 9 graphics device with WDDM 1.0 or higher driver
- Adobe Acrobat Reader version 8 and above

Mac (minimum)

- OS X 10.11, 10.10, 10.9, or 10.8
- Intel core Duo 1.83 GHz
- 512 MB RAM (1GB recommended)
- 1.5 GB hard disk space
- 32-bit color depth at 1024 x 768 resolution
- Adobe Acrobat Reader version 8 and above

advertise those prefixes to other routers inside the same ASN (Internal BGP), as shown in the center of Figure 12-2.

Basically, eBGP refers to BGP's use in advertising routes between two different ASNs. iBGP refers to using BGP to advertise routes to other routers inside the same ASN. BGP uses slightly different rules and details of what it advertises and how it works based on whether a neighbor is an eBGP or iBGP peer. This chapter deals with eBGP only.

Choosing the Best Routes with BGP

The beginning of this chapter stated that all routing protocols learn prefixes, choose the best route if multiple routes exist, and converge when the network changes. BGP focuses on that first action—advertising reachability for a prefix. Additionally, like the IGPs, BGP has a concept like a metric, so that BGP can choose the best route among competing routes. However, BGP uses a much different approach to defining a metric and a different approach to the logic BGP uses to determine the best route.

First, BGP does not use a single idea of a metric. Instead, it uses *path attributes*. BGP advertises each prefix along with a list of different path attributes. The path attributes are different facts about the route (path) to reach that subnet. By using multiple concepts (multiple path attributes), BGP allows for a much wider variety of decisions about which route is best.

BGP then uses a process called the *best path selection* process to choose the best route between two competing routes. When receiving a BGP update that lists a prefix that already exists, the router has a simple choice to make: Is the old route better, or is this new one better? The best path selection process works through a series of comparisons (about ten) until a comparison shows one of the routes as being better. Understanding the process requires a detailed understanding of the path attributes by the network engineer (and is beyond the scope of this chapter). However, BGP on the routers requires very little work to make the comparisons, much less than, say, OSPF's SPF process, so the BGP best path selection process scales well.

While that general description of best path selection is accurate, a single example can help. BGP uses the AS_Path path attribute (PA) in one step of the best path selection process. AS_Path is a path attribute sent with routes in BGP, and this attribute lists the ASNs in the route. The best path selection process considers a shorter AS_Path as better; you can think of it a little like a hop count, but the hops are entire ASNs rather than single routers.

Figure 12-3 shows an example with two competing routes to reach prefix 192.0.2.0/24, one with a shorter AS_Path length. The enterprise advertises its public prefix 192.0.2.0/24 to ISP1, which then advertises it to other ISPs. ISP3 eventually learns two possible routes to that prefix, one that lists an AS_Path with three ASNs, and the shorter one with two ASNs (the best path, as noted with the >).

Key Topic

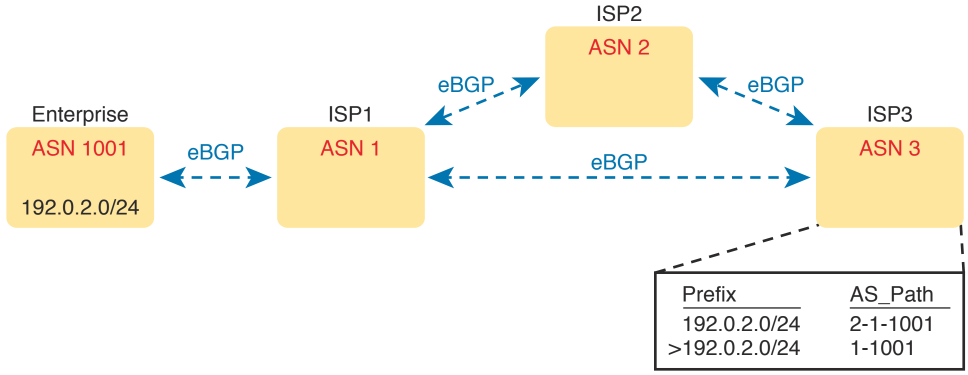

Figure 12-3 *ASNs and Shortest AS_Path as Chosen at ISP3 for Prefix 192.0.2.0/24*

If you continue down the routing and switching track in your Cisco studies, you will learn more and more about BGP and the best path algorithm. Both CCNP ROUTE and CCIE R&S include many details of BGP.

eBGP and the Internet Edge

The term *Internet edge* refers to the connection between an ISP customer and an ISP. The one BGP exam topic focuses the BGP discussion on the Internet edge, specifically the eBGP peering between an enterprise and an ISP, and what BGP can usefully do at the edge.

Internet Edge Designs and Terminology

The term *single homed* (used in the one BGP exam topic) refers to a particular Internet edge design with a single link to one ISP, as shown in Figure 12-4. A single-homed design has a single link between the enterprise and an ISP. You would typically find a single-homed Internet edge design used when connecting an enterprise branch office to the Internet, or for a simple connection from a core site at the enterprise.

Key Topic

Figure 12-4 *Single-Homed Design: Single Link, One Home (ISP)*

A single-homed design allows any kind of WAN link. That is, it could be DSL, cable, fiber Ethernet, or even a wireless LTE connection. The more important points are that the enterprise site connects to a single ISP, and that only one link exists to that ISP. Figure 12-5 expands upon those concepts.

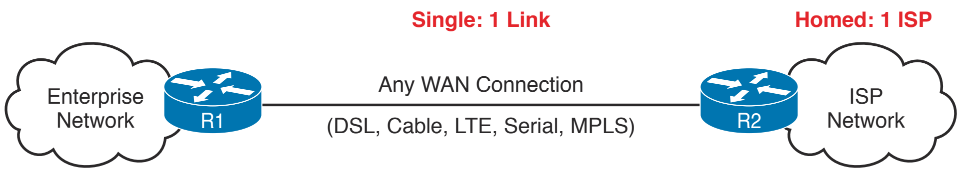

Figure 12-5 *Any Physical Link in Single-Homed Internet Edge*

A single-homed connection has some pros and cons. With only a single link, if the link fails, the ability to reach the Internet fails as well. A failure on either of the routers on the ends of the links also causes the Internet connection to fail. Of course, a single Internet connection, versus multiple, saves money.

Other Internet edge designs do exist, of course, mostly adding redundancy, increasing capacity, and raising the cost and complexity. Figure 12-6 shows three similar designs, in growing levels of complexity. Note that the designs with *dual* in the name refer to designs with two (or more) links to the same router, while *multihomed* refers to having connections to multiple ISPs.

You need basic BGP skills only to understand what happens in single-homed and dual-homed designs—in fact, this chapter focuses mostly on the needs of the single-homed designs, with some extra information useful to dual-homed designs. The added redundancy of the different multi-homed designs requires a much deeper understanding of BGP concepts, path attributes, best path selection, and configuration.

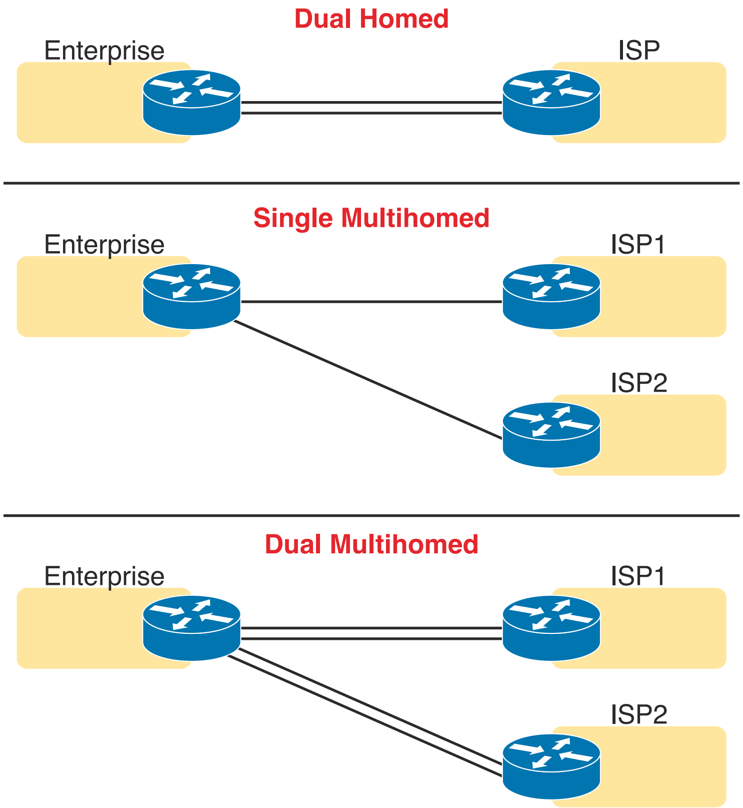

Figure 12-6 *Different Internet Edge Designs*

NOTE The one CCNA R&S exam topic limits BGP specifically to eBGP and to single-homed designs. Frankly, the enterprises that use a single-homed Internet connection often do not even bother to use eBGP, often using static routing instead. However, BGP has many rules and configuration options. Cisco's limits on the BGP exam topic gives CCNA candidates a chance to begin learning a little about BGP, but in a limited way that avoids most of BGP's complexity.

Advertising the Enterprise Public Prefix into the Internet

Next, consider what routes should be advertised between two eBGP peers at the Internet edge. Begin with routes advertised by the enterprise to the ISP, as with a typical enterprise site on the left side of Figure 12-7. In that particular design, the enterprise uses

Private 10.0.0.0/8: Like many companies, this enterprise uses private IP network 10.0.0.0 for most hosts in the enterprise.

Public 192.0.2.0/24: The public IPv4 network assigned to the company. As shown, it has been subnetted, one subnet for use with NAT, the other used for a DMZ with public-facing web servers.

Working through the public addressing in a little more detail, the enterprise uses public IP network 192.0.2.0/24. It uses a NAT configuration on the router, with subnet 192.0.2.0/29 configured on a loopback interface, ready for use by NAT. The design also includes a security demilitarized zone (DMZ), where public-facing servers, such as public web servers, can reside. These hosts also use some of the enterprise's public address range, from subnet 192.0.2.128/26 in this case.

Key Topic

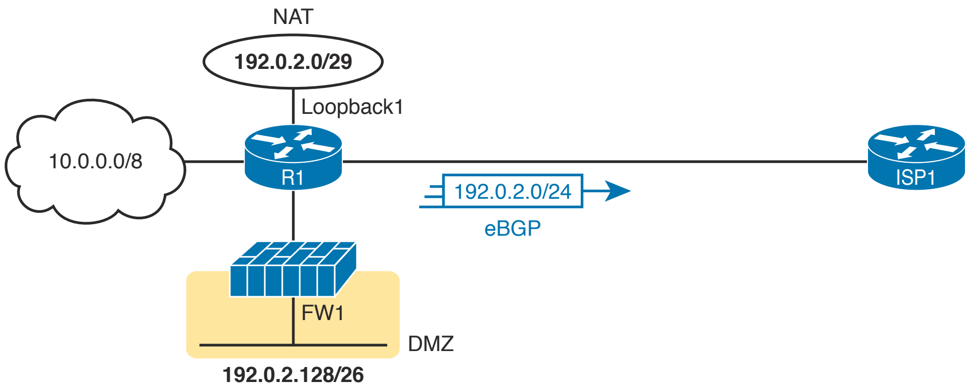

Figure 12-7 *Public and Private IPv4 Address Prefixes Advertised to an ISP*

Note that over the eBGP connection from R1 to ISP1, R1 advertises only one route: a route for public address prefix 192.0.2.0/24. Private IP networks should never be advertised into the Internet; the ISP would filter that advertisement anyway. For the public address range, the ISP wants to receive one route for each public address block, rather than hearing about each subnet, because the ISP does not care how the enterprise subnets the network. Figure 12-7 shows the one prefix the enterprise would normally advertise to the ISP in an eBGP update.

Why advertise the enterprise's public prefix to the ISP at all? So that the ISP can then advertise the prefix to other ISPs, and eventually all ISPs have a route to deliver packets to the enterprise's public IP address. Figure 12-8 shows the idea. At Step 1, R1, in the enterprise, uses eBGP to advertise its public IPv4 prefix to ISP1. At Step 2, ISP1 advertises that same prefix to two different ISPs.

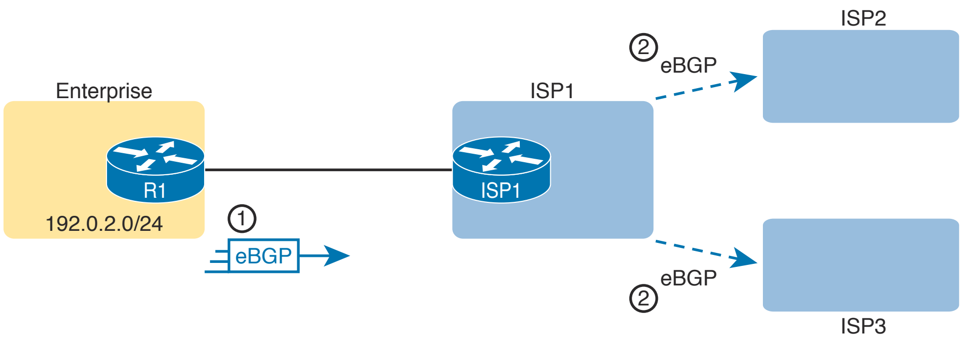

Figure 12-8 *Enterprise Advertises the Public Prefix; ISP Propagates*

Learning Default Routes from the ISP

Now turn the logic around, and think about what the ISP can and should advertise with eBGP to the enterprise. The Internet core routers have almost 600,000 IPv4 routes. An enterprise router at the Internet edge could learn all those routes from the ISP, and put those in its IPv4 routing table. Then, through a process called *redistribution*, that router could take all those BGP-learned routes and advertise them into the rest of the enterprise routers using the IGP.

However, what would all those new IP routes list as their next-hop addresses? All those new routes would send packets toward R1, and R1 would send them toward ISP1. Typically, the routers normally used by enterprises would not be designed to perform well with hundreds of thousands of routes in their routing tables.

Using a default route makes much more sense for a single-homed Internet edge design as compared with learning all those extra IP routes. The enterprise router that connects to the ISP needs one default route, so that for any unknown destinations, the packets are sent to the ISP. The idea is exactly what is discussed in regard to how OSPF advertises default routes back in Chapter 8.

When using a default routing strategy at the Internet edge, you can statically configure the default route, or learn it with eBGP. The OSPF chapter showed the static default route configuration, so of course, this chapter shows the effects of using eBGP.

Figure 12-9 shows the basic concept. At Step 1, the ISP advertises a default route with eBGP. At Step 2, router R1 takes the steps necessary with its IGP so that the IGP then reacts to the eBGP-learned default route, advertising a default route into the enterprise with the IGP.

Key Topic

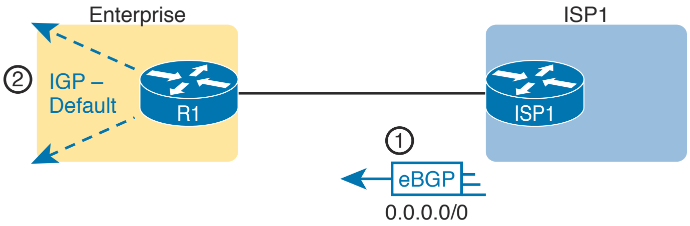

Figure 12-9 *ISP Advertises a Default Route; IGP Propagates*

That concludes the overview of BGP and eBGP in preparation for the configuration and verification topics. The rest of the chapter focuses on the eBGP peering between the enterprise and the ISP, particularly on the enterprise side of the configuration. That includes the details of configuring the peering relationship with the ISP router, advertising the public prefix (with the **network** BGP subcommand), and verifying the eBGP advertised and received routes.

eBGP Configuration and Verification

IGP configuration focuses on enabling the routing protocol on interfaces. The configuration may use the **network** router subcommand, which the router then compares to interface IPv4 addresses. Alternately, the IGP configuration may use an interface subcommand like **ip ospf** *process-id* **area** *number*, which enables the IGP directly on the interface. Once enabled on an interface, the IGP

- Dynamically discovers neighbors that share the same data link off that interface
- Advertises known routes to those dynamically discovered neighbors
- Advertises about the subnet connected to the interface

BGP configuration, in contrast, does none of the actions in the preceding list. (BGP does use a **network** command, but for a different purpose.) BGP does form neighbor relationships, but BGP has no concept of being enabled on an interface, or of dynamically discovering neighbors. Instead, with BGP:

12

Key Topic

- Predefine neighbors with the **neighbor** *ip-address* **remote-as** *asn* BGP subcommand
- Advertise about prefixes that have been added to the BGP table using
  - The BGP **network** command
  - Route redistribution
  - By learning prefixes from a neighbor

In some ways, you have to unlearn some of what you know about IGP configuration to learn BGP configuration. This section walks you through the basics.

BGP Configuration Concepts

BGP uses TCP to transport its messages between two BGP peers, using well-known port 179. When you configure BGP, it opens port 179, waiting for incoming connection requests from other routers. Once a peer connects, the TCP connection is formed.

Once two routers form a TCP connection for use by BGP, the BGP process on each router must decide whether the two routers should become neighbors or not. The idea is much like the overall process with OSPF and EIGRP neighbors exchanging messages to decide if they should become neighbors. The two BGP routers send BGP messages that do some basic checks of parameters to make sure the two routers should become peers. If all checks are passed, the two routers become BGP peers (neighbors). At that point, the two routers can exchange routing information.

BGP uses the update message to exchange information. Once the BGP peer has been established, BGP peers send update messages which hold prefix/length (NLRI) information and the associated path attributes (PA). Those PAs include the AS_Path introduced earlier in the chapter. Figure 12-10 shows an example, the four steps of which are described as follows:

Step 1. Because of proper neighbor configuration in the two routers, the two routers create a TCP connection with each other and become BGP peers.

Step 2. Because of additional configuration, possibly the BGP **network** subcommand, Router R1 adds one or more NLRI and associated PAs to its local BGP table.

Step 3. eBGP on R1 advertises all the best routes in its BGP table—that is, the routes it considers as best for each NLRI—in a BGP update message sent to router ISP1. In this case, R1 advertises about NLRI 192.0.2.0/24.

Step 4. BGP on ISP1 processes that received update, adding a BGP table entry for 192.0.2.0/24.

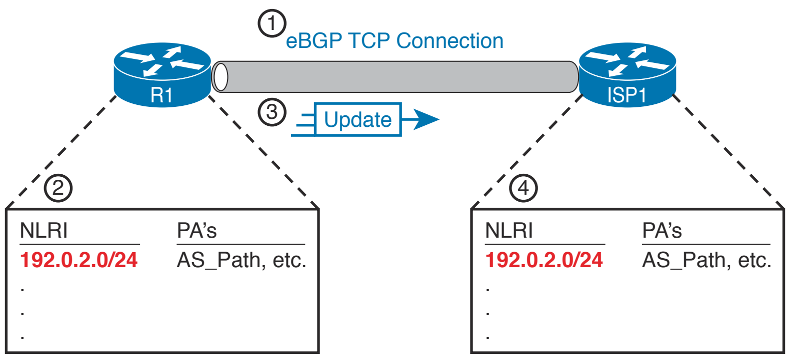

Figure 12-10 *Steps for Advertising a Public Prefix from the Enterprise to an ISP*

This process requires two key configuration steps: configuring neighbors, and configuration that adds entries to a router's BGP table. The next few pages explain both.

Configuring eBGP Neighbors Using Link Addresses

BGP configuration begins with a familiar type of command: the **router bgp** *asn* command, where *asn* is the AS number used by that enterprise or ISP. That command moves the CLI user into BGP configuration mode, much like the **router ospf**, **router eigrp**, and other similar commands.

In a single-homed eBGP design, one link exists between the two routers that need to be eBGP peers. As a result, the two routers can use their interface IP addresses. The **neighbor** *peer-ip-address* BGP subcommand defines the IP address of the neighbor, as shown in Figure 12-11.