

Introduction to Dynamic Routing Protocols

Objectives

Upon completion of this chapter, you should be able to answer the following questions:

- Can you describe the role of dynamic routing protocols and place these protocols in the context of modern network design?
- What are several ways to classify routing protocols?
- How are metrics used by routing protocols, and what are the metric types used by dynamic routing protocols?
- How do you determine the administrative distance of a route, and what is its importance in the routing process?
- What are the different elements in the routing table?
- Given realistic constraints, can you devise and apply subnetting schemes?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary at the end of the book.

scale 157

algorithm 159

autonomous system 163

routing domain 163

Interior gateway protocols 163

Exterior gateway protocols 163

path vector protocol 164

Distance vector 165

vectors 165

link-state 165

link-state router 165

converged 166

Classful routing protocols 166

VLSM 166

discontiguous 166

Classless routing protocols 167

Convergence 168

Administrative distance 173

The data networks that we use in our everyday lives to learn, play, and work range from small, local networks to large, global internetworks. At home, you might have a router and two or more computers. At work, your organization might have multiple routers and switches servicing the data communication needs of hundreds or even thousands of PCs.

In Chapters 1 and 2 you discovered how routers are used in packet forwarding and that routers learn about remote networks using both static routes and dynamic routing protocols. You also know how routes to remote networks can be configured manually using static routes.

This chapter introduces dynamic routing protocols, including how different routing protocols are classified, what metrics they use to determine best path, and the benefits of using a dynamic routing protocol.

Dynamic routing protocols are typically used in larger networks to ease the administrative and operational overhead of using only static routes. Typically, a network uses a combination of both a dynamic routing protocol and static routes. In most networks, a single dynamic routing protocol is used; however, there are cases where different parts of the network can use different routing protocols.

Since the early 1980s, several different dynamic routing protocols have emerged. This chapter begins to discuss some of the characteristics and differences in these routing protocols; however, this will become more evident in later chapters, with a discussion of several of these routing protocols in detail.

Although many networks will use only a single routing protocol or use only static routes, it is important for a network professional to understand the concepts and operations of all the different routing protocols. A network professional must be able to make an informed decision regarding when to use a dynamic routing protocol and which routing protocol is the best choice for a particular environment.

Introduction to Dynamic Routing Protocols

Dynamic routing protocols play an important role in today's networks. The following sections describe several important benefits that dynamic routing protocols provide. In many networks, dynamic routing protocols are typically used with static routes.

Perspective and Background

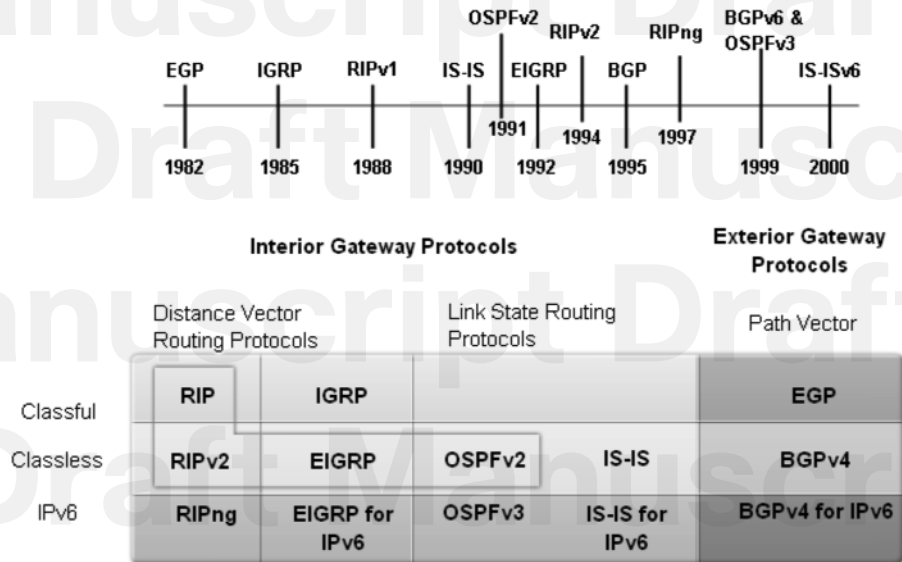
Dynamic routing protocols have evolved over several years to meet the demands of changing network requirements. Although many organizations have migrated to more recent routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP) and Open Shortest Path First (OSPF), many of the earlier routing protocols, such as Routing Information Protocol (RIP), are still in use today.

Evolution of Dynamic Routing Protocols

Dynamic routing protocols have been used in networks since the early 1980s. The first version of RIP was released in 1982, but some of the basic algorithms within the protocol were used on the ARPANET as early as 1969.

As networks have evolved and become more complex, new routing protocols have emerged. Figure 3-1 shows the classification of routing protocols.

Figure 3-1 Routing Protocols' Evolution and Classification



Highlighted routing protocols are the focus of this course.

Figure 3-1 shows a timeline of IP routing protocols, with a chart that helps classify the various protocols. This chart will be referred to several times throughout this book.

One of the earliest routing protocols was RIP. RIP has evolved into a newer version: RIPv2. However, the newer version of RIP still does not *scale* to larger network implementations. To address the needs of larger networks, two advanced routing protocols were developed: OSPF and Intermediate System-to-Intermediate System (IS-IS). Cisco developed Interior Gateway Routing Protocol (IGRP) and Enhanced IGRP (EIGRP). EIGRP also scales well in larger network implementations.

Additionally, there was the need to interconnect different internetworks and provide routing among them. Border Gateway Protocol (BGP) is now used between Internet service providers (ISP) as well as between ISPs and their larger private clients to exchange routing information.

With the advent of numerous consumer devices using IP, the IPv4 addressing space is nearly exhausted. Thus IPv6 has emerged. To support the communication based on IPv6, newer versions of the IP routing protocols have been developed (see the IPv6 row in Figure 3-1).

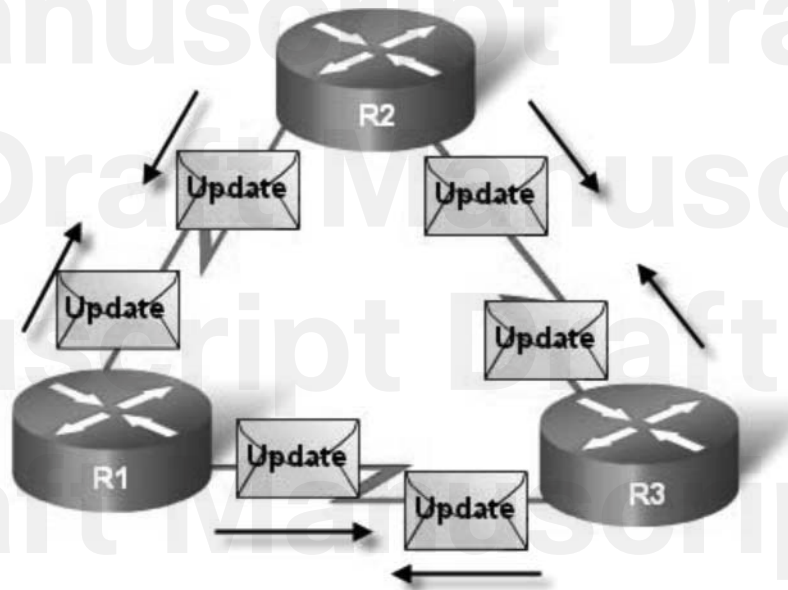
Note

This chapter presents an overview of the different dynamic routing protocols. More details about RIP, EIGRP, and OSPF routing protocols will be discussed in later chapters. The IS-IS and BGP routing protocols are explained in the CCNP curriculum. IGRP is the predecessor to EIGRP and is now considered obsolete.

Role of Dynamic Routing Protocol

What exactly are dynamic routing protocols? Routing protocols are used to facilitate the exchange of routing information between routers. Routing protocols allow routers to dynamically learn information about remote networks and automatically add this information to their own routing tables, as shown in Figure 3-2.

Figure 3-2 Routers Dynamically Pass Updates



Routing protocols determine the best path to each network, which is then added to the routing table. One of the primary benefits of using a dynamic routing protocol is that routers exchange routing information whenever there is a topology change. This exchange allows routers to automatically learn about new networks and also to find alternate paths if there is a link failure to a current network.

Compared to static routing, dynamic routing protocols require less administrative overhead. However, the expense of using dynamic routing protocols is dedicating part of a router's resources for protocol operation, including CPU time and network link bandwidth. Despite the benefits of dynamic routing, static routing still has its place. There are times when static routing is more appropriate and other times when dynamic routing is the better choice. More often than not, you will find a combination of both types of routing in any network that has a moderate level of complexity. You will learn about the advantages and disadvantages of static and dynamic routing later in this chapter.

Network Discovery and Routing Table Maintenance

Two important processes concerning dynamic routing protocols are initially discovering remote networks and maintaining a list of those networks in the routing table.

Purpose of Dynamic Routing Protocols

A routing protocol is a set of processes, algorithms, and messages that are used to exchange routing information and populate the routing table with the routing protocol's choice of best paths. The purpose of a routing protocol includes

- Discovering remote networks
- Maintaining up-to-date routing information
- Choosing the best path to destination networks
- Having the ability to find a new best path if the current path is no longer available

The components of a routing protocol are as follows:

- **Data structures:** Some routing protocols use tables and/or databases for their operations. This information is kept in RAM.
- **Algorithm:** An *algorithm* is a finite list of steps used in accomplishing a task. Routing protocols use algorithms for processing routing information and for best-path determination.
- **Routing protocol messages:** Routing protocols use various types of messages to discover neighboring routers, exchange routing information, and do other tasks to learn and maintain accurate information about the network.

Dynamic Routing Protocol Operation

All routing protocols have the same purpose: to learn about remote networks and to quickly adapt whenever there is a change in the topology. The method that a routing protocol uses to accomplish this depends on the algorithm it uses and the operational characteristics of

that protocol. The operations of a dynamic routing protocol vary depending on the type of routing protocol and the operations of that routing protocol. The specific operations of RIP, EIGRP, and OSPF are examined in later chapters. In general, the operations of a dynamic routing protocol can be described as follows:

1. The router sends and receives routing messages on its interfaces.
2. The router shares routing messages and routing information with other routers that are using the same routing protocol.
3. Routers exchange routing information to learn about remote networks.
4. When a router detects a topology change, the routing protocol can advertise this change to other routers.

Note

Understanding dynamic routing protocol operation and concepts and using these protocols in real networks require a solid knowledge of IP addressing and subnetting. Three subnetting scenarios are available in *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-204-4) for your practice.

Dynamic Routing Protocol Advantages

Dynamic routing protocols provide several advantages, which will be discussed in this section. In many cases, the complexity of the network topology, the number of networks, and the need for the network to automatically adjust to changes require the use of a dynamic routing protocol.

Before examining the benefits of dynamic routing protocols in more detail, you need to consider the reasons why you would use static routing. Dynamic routing certainly has several advantages over static routing; however, static routing is still used in networks today. In fact, networks typically use a combination of both static and dynamic routing.

Table 3-1 compares dynamic and static routing features. From this comparison, you can list the advantages of each routing method. The advantages of one method are the disadvantages of the other.

Table 3-1 Dynamic Versus Static Routing

Feature	Dynamic Routing	Static Routing
Configuration complexity	Generally independent of the network size	Increases with network size
Required administrator knowledge	Advanced knowledge required	No extra knowledge required
Topology changes	Automatically adapts to topology changes	Administrator intervention required

Feature	Dynamic Routing	Static Routing
Scaling	Suitable for simple and complex topologies	Suitable for simple topologies
Security	Less secure	More secure
Resource usage	Uses CPU, memory, and link bandwidth	No extra resources needed
Predictability	Route depends on the current topology	Route to destination is always the same

Static Routing Usage, Advantages, and Disadvantages

Static routing has several primary uses, including the following:

- Providing ease of routing table maintenance in smaller networks that are not expected to grow significantly.
- Routing to and from stub networks (see Chapter 2).
- Using a single default route, used to represent a path to any network that does not have a more specific match with another route in the routing table.

Static routing advantages are as follows:

- Minimal CPU processing
- Easier for administrator to understand
- Easy to configure

Static routing disadvantages are as follows:

- Configuration and maintenance are time-consuming.
- Configuration is error-prone, especially in large networks.
- Administrator intervention is required to maintain changing route information.
- Does not scale well with growing networks; maintenance becomes cumbersome.
- Requires complete knowledge of the entire network for proper implementation.

Dynamic Routing Advantages and Disadvantages

Dynamic routing advantages are as follows:

- Administrator has less work in maintaining the configuration when adding or deleting networks.

- Protocols automatically react to the topology changes.
- Configuration is less error-prone.
- More scalable; growing the network usually does not present a problem.

Dynamic routing disadvantages are as follows:

- Router resources are used (CPU cycles, memory, and link bandwidth).
- More administrator knowledge is required for configuration, verification, and troubleshooting.

Classifying Dynamic Routing Protocols

Figure 3-1 showed how routing protocols can be classified according to various characteristics. This chapter will introduce you to these terms, which will be discussed in more detail in later chapters.

This section gives an overview of the most common IP routing protocols. Most of these routing protocols will be examined in detail later in this book. For now, we will give a very brief overview of each protocol.

Routing protocols can be classified into different groups according to their characteristics:

- IGP or EGP
- Distance vector or link-state
- Classful or classless

The sections that follow discuss these classification schemes in more detail.

The most commonly used routing protocols are as follows:

- **RIP:** A distance vector interior routing protocol
- **IGRP:** The distance vector interior routing protocol developed by Cisco (deprecated from Cisco IOS Release 12.2 and later)
- **OSPF:** A link-state interior routing protocol
- **IS-IS:** A link-state interior routing protocol
- **EIGRP:** The advanced distance vector interior routing protocol developed by Cisco
- **BGP:** A path vector exterior routing protocol

Note

IS-IS and BGP are beyond the scope of this book.

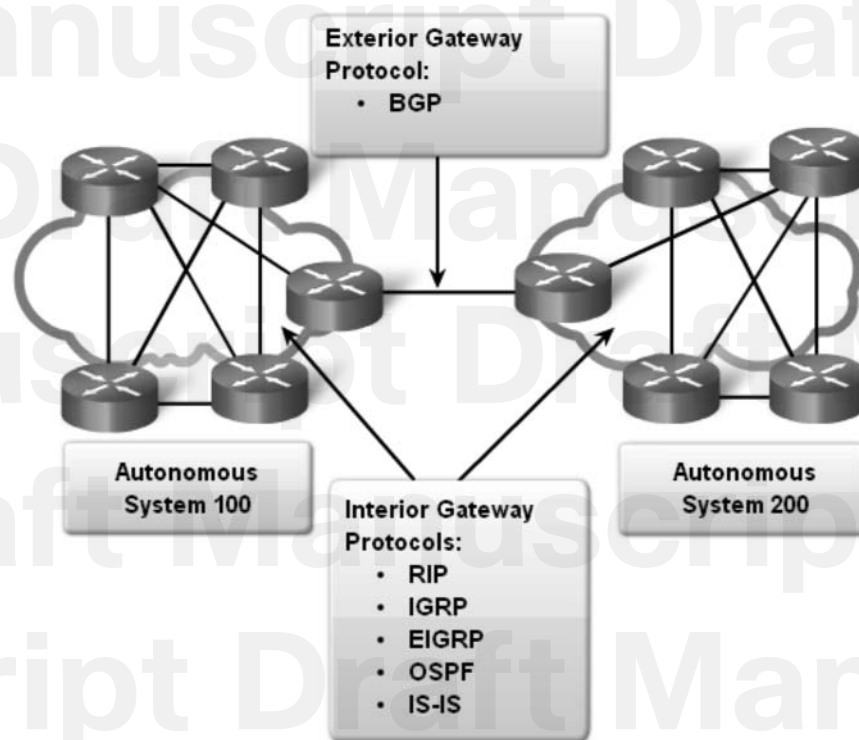
IGP and EGP

An *autonomous system* (AS)—otherwise known as a *routing domain*—is a collection of routers under a common administration. Typical examples are a company's internal network and an ISP's network. Because the Internet is based on the autonomous system concept, two types of routing protocols are required: interior and exterior routing protocols. These protocols are

- **Interior gateway protocols** (IGP): Used for intra-autonomous system routing, that is, routing inside an autonomous system
- **Exterior gateway protocols** (EGP): Used for inter-autonomous system routing, that is, routing between autonomous systems

Figure 3-3 is a simplified view of the difference between IGPs and EGPs. The autonomous system concept will be explained in more detail later in the chapter. Even though this is an oversimplification, for now, think of an autonomous system as an ISP.

Figure 3-3 IGP Versus EGP Routing Protocols



IGPs are used for routing within a routing domain, those networks within the control of a single organization. An autonomous system is commonly comprised of many individual networks belonging to companies, schools, and other institutions. An IGP is used to route within the autonomous system and also used to route within the individual networks themselves. For example, The Corporation for Education Network Initiatives in California (CENIC) operates an autonomous system comprised of California schools, colleges, and universities. CENIC uses an IGP to route within its autonomous system to interconnect all of these institutions. Each of the educational institutions also uses an IGP of its own choosing to route within its own individual network. The IGP used by each entity provides best-path determination within its own routing domains, just as the IGP used by CENIC provides best-path routes within the autonomous system itself. IGP for IP include RIP, IGRP, EIGRP, OSPF, and IS-IS.

Routing protocols (and more specifically, the algorithm used by that routing protocol) use a metric to determine the best path to a network. The metric used by the routing protocol RIP is *hop count*, which is the number of routers that a packet must traverse in reaching another network. OSPF uses *bandwidth* to determine the shortest path.

EGPs, on the other hand, are designed for use between different autonomous systems that are under the control of different administrations. BGP is the only currently viable EGP and is the routing protocol used by the Internet. BGP is a *path vector protocol* that can use many different attributes to measure routes. At the ISP level, there are often more important issues than just choosing the fastest path. BGP is typically used between ISPs and sometimes between a company and an ISP. BGP is not part of this course or CCNA; it is covered in CCNP.

Packet Tracer
Activity

Characteristics of IGP and EGP Routing Protocols (3.2.2)

In this activity, the network has already been configured within the autonomous systems. You will configure a default route from AS2 and AS3 (two different companies) to the ISP (AS1) to simulate the exterior gateway routing that would take place from both companies to their ISP. Then you will configure a static route from the ISP (AS1) to AS2 and AS3 to simulate the exterior gateway routing that would take place from the ISP to its two customers, AS2 and AS3. View the routing table before and after both static routes and default routes are added to observe how the routing table has changed. Use file e2-322.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Distance Vector and Link-State Routing Protocols

Interior gateway protocols (IGP) can be classified as two types:

- Distance vector routing protocols
- Link-state routing protocols

Distance Vector Routing Protocol Operation

Distance vector means that routes are advertised as *vectors* of distance and direction.

Distance is defined in terms of a metric such as hop count, and direction is simply the next-hop router or exit interface. Distance vector protocols typically use the Bellman-Ford algorithm for the best-path route determination.

Some distance vector protocols periodically send complete routing tables to all connected neighbors. In large networks, these routing updates can become enormous, causing significant traffic on the links.

Although the Bellman-Ford algorithm eventually accumulates enough knowledge to maintain a database of reachable networks, the algorithm does not allow a router to know the exact topology of an internetwork. The router only knows the routing information received from its neighbors.

Distance vector protocols use routers as signposts along the path to the final destination. The only information a router knows about a remote network is the distance or metric to reach that network and which path or interface to use to get there. Distance vector routing protocols do not have an actual map of the network topology.

Distance vector protocols work best in situations where

- The network is simple and flat and does not require a hierarchical design.
- The administrators do not have enough knowledge to configure and troubleshoot link-state protocols.
- Specific types of networks, such as hub-and-spoke networks, are being implemented.
- Worst-case convergence times in a network are not a concern.

Chapter 4, “Distance Vector Routing Protocols,” covers distance vector routing protocol functions and operations in greater detail. You will also learn about the operations and configuration of the distance vector routing protocols RIP and EIGRP.

Link-State Protocol Operation

In contrast to distance vector routing protocol operation, a router configured with a *link-state* routing protocol can create a “complete view,” or topology, of the network by gathering information from all the other routers. Think of using a link-state routing protocol as having a complete map of the network topology. The signposts along the way from source to destination are not necessary, because all link-state routers are using an identical “map” of the network. A *link-state router* uses the link-state information to create a topology map and to select the best path to all destination networks in the topology.

With some distance vector routing protocols, routers send periodic updates of their routing information to their neighbors. Link-state routing protocols do not use periodic updates.

After the network has *converged*, a link-state update is only sent when there is a change in the topology.

Link-state protocols work best in situations where

- The network design is hierarchical, usually occurring in large networks.
- The administrators have a good knowledge of the implemented link-state routing protocol.
- Fast convergence of the network is crucial.

Link-state routing protocol functions and operations will be explained in later chapters. You will also learn about the operations and configuration of the link-state routing protocol OSPF in Chapter 11, “OSPF.”

Classful and Classless Routing Protocols

All routing protocols can also be classified as either

- Classful routing protocols
- Classless routing protocols

Classful Routing Protocols

Classful routing protocols do not send subnet mask information in routing updates. The first routing protocols, such as RIP, were classful. This was at a time when network addresses were allocated based on classes: Class A, B, or C. A routing protocol did not need to include the subnet mask in the routing update because the network mask could be determined based on the first octet of the network address.

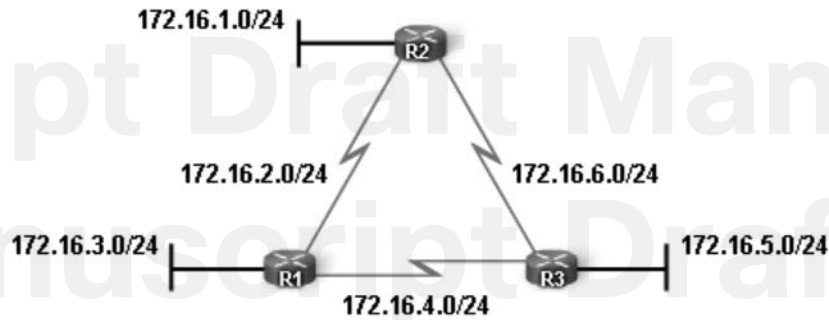
Classful routing protocols can still be used in some of today’s networks, but because they do not include the subnet mask, they cannot be used in all situations. Classful routing protocols cannot be used when a network is subnetted using more than one subnet mask. In other words, classful routing protocols do not support variable-length subnet masks (*VLSM*).

Figure 3-4 shows an example of a network using the same subnet mask on all its subnets for the same major network address. In this situation, either a classful or classless routing protocol could be used.

There are other limitations to classful routing protocols, including their inability to support *discontiguous* networks. Later chapters discuss classful routing protocols, discontiguous networks, and VLSM in greater detail.

Classful routing protocols include RIPv1 and IGRP.

Figure 3-4 Classful Routing




Classful: Subnet mask is the same throughout the topology

Classless Routing Protocols

Classless routing protocols include the subnet mask with the network address in routing updates. Today's networks are no longer allocated based on classes, and the subnet mask cannot be determined by the value of the first octet. Classless routing protocols are required in most networks today because of their support for VLSM, discontinuous networks, and other features that will be discussed in later chapters.

In Figure 3-5, notice that the classless version of the network is using both /30 and /27 subnet masks in the same topology. Also notice that this topology is using a discontinuous design.

Figure 3-5 Classless Routing



Classless: Subnet mask can vary in the topology

Classless routing protocols are RIPv2, EIGRP, OSPF, IS-IS, and BGP.

Dynamic Routing Protocols and Convergence

An important characteristic of a routing protocol is how quickly it converges when there is a change in the topology.

Convergence is when the routing tables of all routers are at a state of consistency. The network has converged when all routers have complete and accurate information about the network. Convergence time is the time it takes routers to share information, calculate best paths, and update their routing tables. A network is not completely operable until the network has converged; therefore, most networks require short convergence times.

Convergence is both collaborative and independent. The routers share information with each other but must independently calculate the impacts of the topology change on their own routes. Because they develop an agreement with the new topology independently, they are said to *converge* on this consensus.

Convergence properties include the speed of propagation of routing information and the calculation of optimal paths. Routing protocols can be rated based on the speed to converge; the faster the convergence, the better the routing protocol. Generally, RIP and IGRP are slow to converge, whereas EIGRP, OSPF, and IS-IS are faster to converge.

Packet Tracer
Activity

Convergence (3.2.5)

In this activity, the network has already been configured with two routers, two switches, and two hosts. A new LAN will be added, and you will watch the network converge. Use file e2-325.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Metrics

Metrics are a way to measure or compare. Routing protocols use metrics to determine which route is the best path.

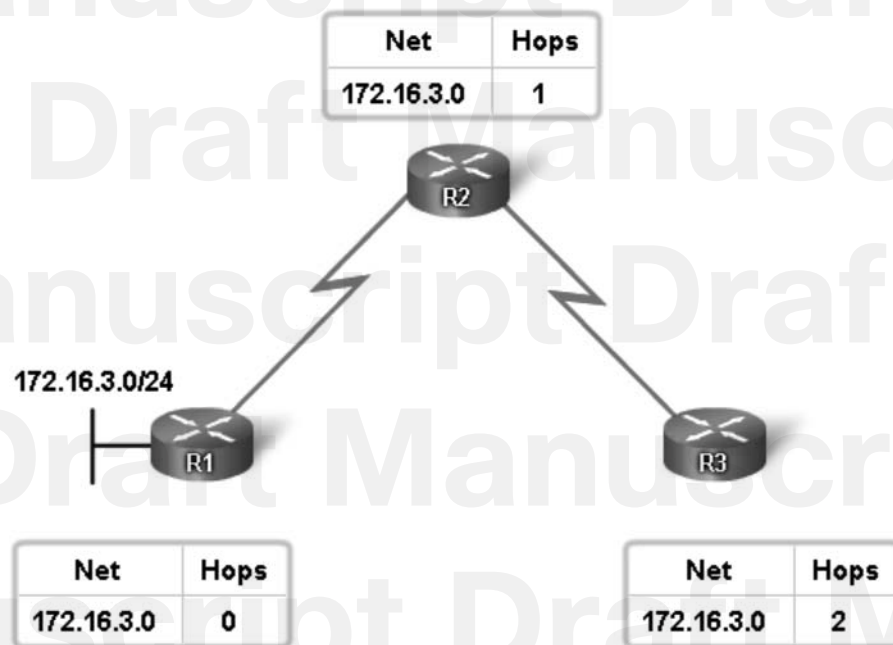
Purpose of a Metric

There are cases when a routing protocol learns of more than one route to the same destination. To select the best path, the routing protocol must be able to evaluate and differentiate among the available paths. For this purpose, a metric is used. A metric is a value used by routing protocols to assign costs to reach remote networks. The metric is used to determine which path is most preferable when there are multiple paths to the same remote network.

Each routing protocol calculates its metric in a different way. For example, RIP uses hop count, EIGRP uses a combination of bandwidth and delay, and the Cisco implementation of OSPF uses bandwidth. Hop count is the easiest metric to envision. The hop count refers to the number of routers a packet must cross to reach the destination network.

For R3 in Figure 3-6, network 172.16.3.0 is two hops, or two routers, away. For R2, network 172.16.3.0 is one hop away, and for R1, it is 0 hops (because the network is directly connected).

Figure 3-6 Metrics



Note

The metrics for a particular routing protocol and a discussion of how they are calculated will be presented in the chapter for that routing protocol.

Metrics and Routing Protocols

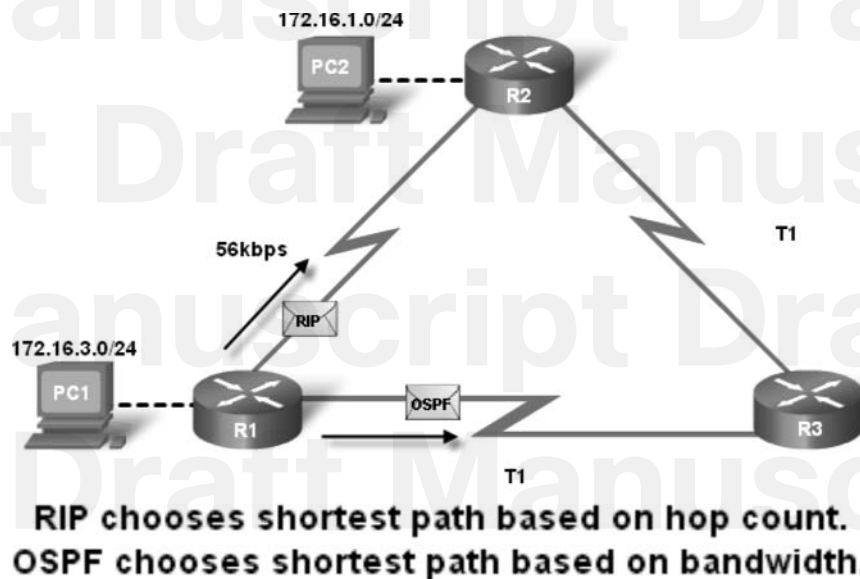
Different routing protocols use different metrics. The metric used by one routing protocol is not comparable to the metric used by another routing protocol.

Metric Parameters

Two different routing protocols might choose different paths to the same destination because of using different metrics.

Figure 3-7 shows how R1 would reach the 172.16.1.0/24 network. RIP would choose the path with the least amount of hops through R2, whereas OSPF would choose the path with the highest bandwidth through R3.

Figure 3-7 Hop Count Versus Bandwidth



Metrics used in IP routing protocols include the following:

- **Hop count:** A simple metric that counts the number of routers a packet must traverse.
- **Bandwidth:** Influences path selection by preferring the path with the highest bandwidth.
- **Load:** Considers the traffic utilization of a certain link.
- **Delay:** Considers the time a packet takes to traverse a path.
- **Reliability:** Assesses the probability of a link failure, calculated from the interface error count or previous link failures.
- **Cost:** A value determined either by the IOS or by the network administrator to indicate preference for a route. Cost can represent a metric, a combination of metrics, or a policy.

Note

At this point, it is not important to completely understand these metrics; they will be explained in later chapters.

Metric Field in the Routing Table

The routing table displays the metric for each dynamic and static route. Remember from Chapter 2 that static routes always have a metric of 0.

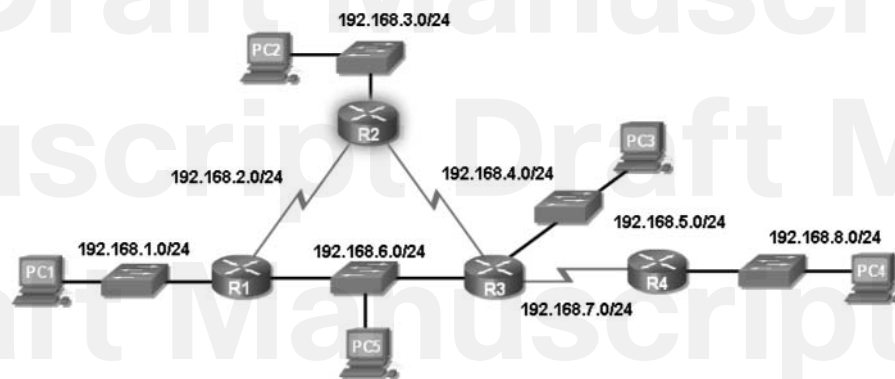
The list that follows defines the metric for each routing protocol:

- **RIP: Hop count:** Best path is chosen by the route with the lowest hop count.
- **IGRP and EIGRP: Bandwidth, delay, reliability, and load:** Best path is chosen by the route with the smallest composite metric value calculated from these multiple parameters. By default, only bandwidth and delay are used.
- **IS-IS and OSPF: Cost:** Best path is chosen by the route with the lowest cost. The Cisco implementation of OSPF uses bandwidth to determine the cost. IS-IS is discussed in CCNP.

Routing protocols determine best path based on the route with the lowest metric.

In Figure 3-8, the routers are all using the RIP routing protocol.

Figure 3-8 Best Path Determined in a Network Using RIP



The metric associated with a certain route can be best viewed using the **show ip route** command. The metric value is the second value in the brackets for a routing table entry. In Example 3-1, R2 has a route to the 192.168.8.0/24 network that is two hops away. The highlighted **2** in the command output is where the routing metric is displayed.

Example 3-1 Routing Table for R2

```
R2# show ip route
```

```
<output omitted>
```

```
Gateway of last resort is not set
```

```
R   192.168.1.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
C   192.168.2.0/24 is directly connected, Serial0/0/0
C   192.168.3.0/24 is directly connected, FastEthernet0/0
C   192.168.4.0/24 is directly connected, Serial0/0/1
R   192.168.5.0/24 [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
R   192.168.6.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
      [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
R   192.168.7.0/24 [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
R   192.168.8.0/24 [120/2] via 192.168.4.1, 00:00:26, Serial0/0/1
```

Load Balancing

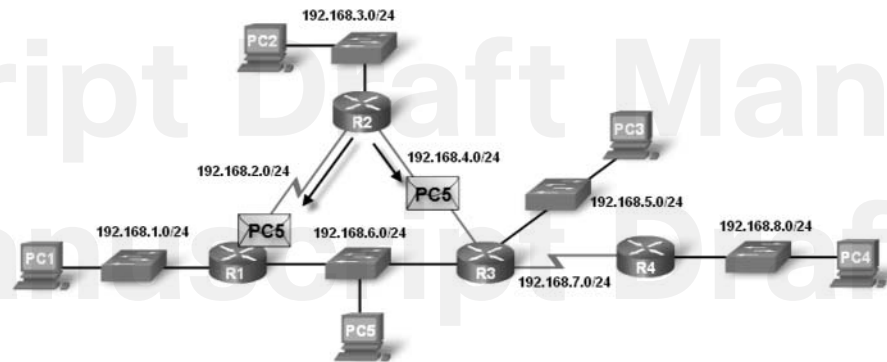
You now know that individual routing protocols use metrics to determine the best route to reach remote networks. But what happens when two or more routes to the same destination have identical metric values? How will the router decide which path to use for packet forwarding? In this case, the router does not choose only one route. Instead, the router *load-balances* between these equal-cost paths. The packets are forwarded using all equal-cost paths.

To see whether load balancing is in effect, check the routing table. Load balancing is in effect if two or more routes are associated with the same destination.

Note

Load balancing can be done either per packet or per destination. How a router actually load-balances packets between the equal-cost paths is governed by the switching process. The switching process will be discussed in greater detail in a later chapter.

Figure 3-9 shows an example of load balancing, assuming that R2 load-balances traffic to PC5 over two equal-cost paths.

Figure 3-9 Load Balancing Across Equal-Cost Paths

R2 load balances traffic destined for the 192.168.6.0/24 network

The **show ip route** command in Example 3-1 reveals that the destination network 192.168.6.0 is available through 192.168.2.1 (Serial 0/0/0) and 192.168.4.1 (Serial 0/0/1). The equal-cost routes are shown again here:

```
R2# show ip route
```

```
<output omitted>
```

```
R    192.168.6.0/24 [120/1] via 192.168.2.1, 00:00:24, Serial0/0/0
      [120/1] via 192.168.4.1, 00:00:26, Serial0/0/1
```

All the routing protocols discussed in this course are capable of automatically load-balancing traffic for up to four equal-cost routes by default. EIGRP is also capable of load-balancing across unequal-cost paths. This feature of EIGRP is discussed in the CCNP courses.

Administrative Distance

The following sections introduce the concept of administrative distance. Administrative distance will also be discussed within each chapter that focuses on a particular routing protocol.

Purpose of Administrative Distance

Before the routing process can determine which route to use when forwarding a packet, it must first determine which routes to include in the routing table. There can be times when a router learns a route to a remote network from more than one routing source. The routing process will need to determine which routing source to use. *Administrative distance* is used for this purpose.

Multiple Routing Sources

You know that routers learn about adjacent networks that are directly connected and about remote networks by using static routes and dynamic routing protocols. In fact, a router might learn of a route to the same network from more than one source. For example, a static route might have been configured for the same network/subnet mask that was learned dynamically by a dynamic routing protocol, such as RIP. The router must choose which route to install.

Note

You might be wondering about equal-cost paths. Multiple routes to the same network can only be installed when they come from the same routing source. For example, for equal-cost routes to be installed, they both must be static routes or they both must be RIP routes.

Although less common, more than one dynamic routing protocol can be deployed in the same network. In some situations, it might be necessary to route the same network address using multiple routing protocols such as RIP and OSPF. Because different routing protocols use different metrics—RIP uses hop count and OSPF uses bandwidth—it is not possible to compare metrics to determine the best path.

So, how does a router determine which route to install in the routing table when it has learned about the same network from more than one routing source? Cisco IOS makes the determination based on the administrative distance of the routing source.

Purpose of Administrative Distance

Administrative distance (AD) defines the preference of a routing source. Each routing source—including specific routing protocols, static routes, and even directly connected networks—is prioritized in order of most to least preferable using an administrative distance value. Cisco routers use the AD feature to select the best path when they learn about the same destination network from two or more different routing sources.

Administrative distance is an integer value from 0 to 255. The lower the value, the more preferred the route source. An administrative distance of 0 is the most preferred. Only a directly connected network has an administrative distance of 0, which cannot be changed.

Note

It is possible to modify the administrative distance for static routes and dynamic routing protocols. This is discussed in CCNP courses.

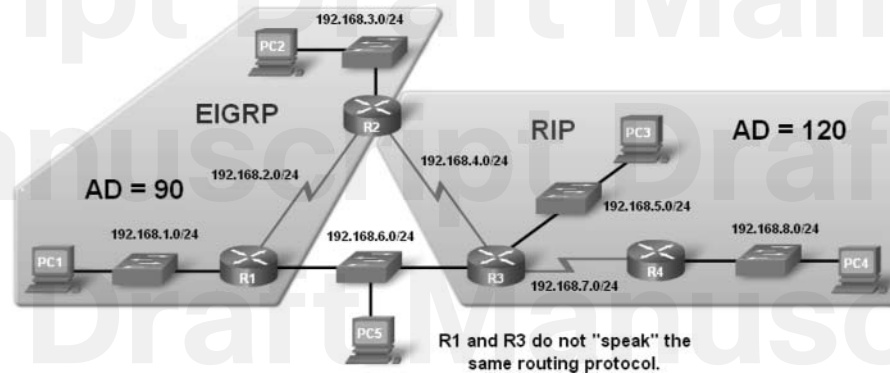
An administrative distance of 255 means the router will not believe the source of that route, and it will not be installed in the routing table.

Note

The term *trustworthiness* is commonly used when defining administrative distance. The lower the administrative distance value, the more trustworthy the route.

Figure 3-10 shows a topology with R2 running both EIGRP and RIP. R2 is running EIGRP with R1 and RIP with R3.

Figure 3-10 Comparing Administrative Distances



Example 3-2 displays the `show ip route` command output for R2.

Example 3-2 Routing Table for R2

R2# `show ip route`

<output omitted>

Gateway of last resort is not set

```
D 192.168.1.0/24 [90/2172416] via 192.168.2.1, 00:00:24, Serial0/0
C 192.168.2.0/24 is directly connected, Serial0/0/0
C 192.168.3.0/24 is directly connected, FastEthernet0/0
C 192.168.4.0/24 is directly connected, Serial0/0/1
R 192.168.5.0/24 [120/1] via 192.168.4.1, 00:00:08, Serial0/0/1
D 192.168.6.0/24 [90/2172416] via 192.168.2.1, 00:00:24, Serial0/0/0
R 192.168.7.0/24 [120/1] via 192.168.4.1, 00:00:08, Serial0/0/1
R 192.168.8.0/24 [120/2] via 192.168.4.1, 00:00:08, Serial0/0/1
```

The AD value is the first value in the brackets for a routing table entry. Notice that R2 has a route to the 192.168.6.0/24 network with an AD value of 90.

```
D 192.168.6.0/24 [90/2172416] via 192.168.2.1, 00:00:24, Serial0/0/0
```

R2 is running both RIP and EIGRP routing protocols. Remember, it is not common for routers to run multiple dynamic routing protocols, but is used here to demonstrate how administrative distance works. R2 has learned of the 192.168.6.0/24 route from R1 through EIGRP updates and from R3 through RIP updates. RIP has an administrative distance of 120, but EIGRP has a lower administrative distance of 90. So, R2 adds the route learned using EIGRP to the routing table and forwards all packets for the 192.168.6.0/24 network to Router R1.

What happens if the link to R1 becomes unavailable? Would R2 not have a route to 192.168.6.0? Actually, R2 still has RIP route information for 192.168.6.0 stored in the RIP database. This can be verified with the **show ip rip database** command, as shown in Example 3-3.

Example 3-3 Verifying RIP Route Availability

```
R2# show ip rip database
```

```
192.168.3.0/24    directly connected, FastEthernet0/0
192.168.4.0/24    directly connected, Serial0/0/1
192.168.5.0/24
    [1] via 192.168.4.1, Serial0/0/1
192.168.6.0/24
    [1] via 192.168.4.1, Serial0/0/1
192.168.7.0/24
    [1] via 192.168.4.1, Serial0/0/1
192.168.8.0/24
    [2] via 192.168.4.1, Serial0/0/1
```

The **show ip rip database** command shows all RIP routes learned by R2, whether or not the RIP route is installed in the routing table. Now you can answer the question as to what would happen if the EIGRP route to 192.168.6.0 became unavailable. RIP has a route, and it would be installed in the routing table. If the EIGRP route is later restored, the RIP route would be removed and the EIGRP route would be reinstalled because it has a better AD value.

Dynamic Routing Protocols and Administrative Distance

You already know that you can verify AD values with the **show ip route** command, as shown previously in Example 3-2.

Example 3-4 shows that the AD value can also be verified with the **show ip protocols** command. This command displays all pertinent information about routing protocols operating on the router.

Example 3-4 Verify Administrative Distance with the **show ip protocols** CommandR2# **show ip protocols**

Routing Protocol is "eigrp 100 "

Outgoing update filter list for all interfaces is not set

Incoming update filter list for all interfaces is not set

Default networks flagged in outgoing updates

Default networks accepted from incoming updates

EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0

EIGRP maximum hopcount 100

EIGRP maximum metric variance 1

Redistributing: eigrp 100

Automatic network summarization is in effect

Automatic address summarization:

Maximum path: 4

Routing for Networks:

192.168.2.0

192.168.3.0

192.168.4.0

Routing Information Sources:

Gateway	Distance	Last Update
192.168.2.1	90	2366569

Distance: internal 90 external 170

Routing Protocol is "rip"

Sending updates every 30 seconds, next due in 12 seconds

Invalid after 180 seconds, hold down 180, flushed after 240

Outgoing update filter list for all interfaces is not set

Incoming update filter list for all interfaces is not set

Redistributing: rip

Default version control: send version 1, receive any version

Interface	Send	Recv	Triggered RIP	Key-chain
Serial0/0/1	1	2	1	
FastEthernet0/0	1	2	1	

Automatic network summarization is in effect

Maximum path: 4

Routing for Networks:

192.168.3.0

192.168.4.0

Passive Interface(s):

Routing Information Sources:

Gateway	Distance	Last Update
192.168.4.1	120	

Distance: (default is 120)

You will see additional coverage of the **show ip protocols** command many times during the rest of the course. However, for now, notice the highlighted output: R2 has two routing protocols listed, and the AD value is called Distance.

Table 3-2 shows the different administrative distance values for various routing protocols.

Table 3-2 Default Administrative Distances

Route Source	AD
Connected	0
Static	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200

Static Routes and Administrative Distance

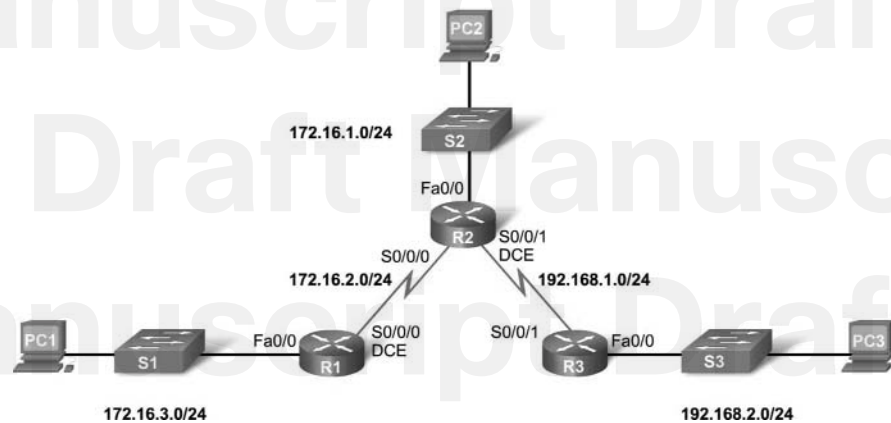
As you know from Chapter 2, static routes are entered by an administrator who wants to manually configure the best path to the destination. For that reason, static routes have a default AD value of 1. This means that after directly connected networks, which have a default AD value of 0, static routes are the most preferred route source.

There are situations when an administrator will configure a static route to the same destination that is learned using a dynamic routing protocol, but using a different path. The static route will be configured with an AD greater than that of the routing protocol. If there is a link failure in the path used by the dynamic routing protocol, the route entered by the routing protocol is removed from the routing table. The static route will then become the only source and will automatically be added to the routing table. This is known as a *floating static route* and is discussed in CCNP courses.

A static route using either a next-hop IP address or an exit interface has a default AD value of 1. However, the AD value is not listed in the **show ip route** output when you configure a static route with the exit interface specified. When a static route is configured with an exit interface, the output shows the network as directly connected through that interface.

Using the topology shown in Figure 3-11 and the **show ip route** command for R2 shown in Example 3-5, you can examine the two types of static routes.

Figure 3-11 Administrative Distances and Static Routes



Example 3-5 Routing Table for R2

```
R2# show ip route
```

```
<output omitted>
```

```
Gateway of last resort is not set
```

```
172.16.0.0/24 is subnetted, 3 subnets
```

```
C 172.16.1.0 is directly connected, FastEthernet0/0
```

```
C 172.16.2.0 is directly connected, Serial0/0/0
```

```
S 172.16.3.0 is directly connected, Serial0/0/0
```

```
C 192.168.1.0/24 is directly connected, Serial0/0/1
```

```
S 192.168.2.0/24 [1/0] via 192.168.1.1
```

The static route to 172.16.3.0 is listed as directly connected. However, there is no information on what the AD value is. It is a common misconception to assume that the AD value of this route must be 0 because it states “directly connected.” However, that is a false assumption. The default AD of any static route, including those configured with an exit interface, is 1. Remember, only a directly connected network can have an AD of 0. This can be verified by extending the **show ip route** command with the *[route]* option. Specifying the *[route]* reveals detailed information about the route, including its distance, or AD value.

The **show ip route 172.16.3.0** command in Example 3-6 reveals that, in fact, the administrative distance for static routes—even with the exit interface specified—is 1.

Example 3-6 show ip route Command with the [route] Option

```
R2# show ip route 172.16.3.0
Routing entry for 172.16.3.0/24
Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
    * directly connected, via Serial0/0/0
      Route metric is 0, traffic share count is 1
```

Directly Connected Networks and Administrative Distance

Directly connected networks appear in the routing table as soon as the IP address on the interface is configured and the interface is enabled and operational. The AD value of directly connected networks is 0, meaning that this is the most preferred routing source. There is no better route for a router than having one of its interfaces directly connected to that network. For that reason, the administrative distance of a directly connected network cannot be changed, and no other route source can have an administrative distance of 0.

The output of the **show ip route** command in Example 3-7 highlights the directly connected networks with no information about the AD value.

Example 3-7 Directly Connected Networks in Routing Table Do Not Show AD Value

```
R2# show 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, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set

172.16.0.0/24 is subnetted, 3 subnets
C    172.16.1.0 is directly connected, FastEthernet0/0
C    172.16.2.0 is directly connected, Serial0/0/0
S    172.16.3.0 is directly connected, Serial0/0/0
C    192.168.1.0/24 is directly connected, Serial0/0/1
S    192.168.2.0/24 [1/0] via 192.168.1.1

```

The output is similar to the output for static routes that point to an exit interface. The only difference is the letter C at the beginning of the entry, which indicates that this is a directly connected network.

To see the AD value of a directly connected network, use the `[route]` option, as shown in Example 3-8.

Example 3-8 Directly Connected Route with AD Value Shown

```

R2# show ip route 172.16.3.0
Routing entry for 172.16.1.0/24
Known via "connected", distance 0, metric 0 (connected, via interface)
Routing Descriptor Blocks:
* directly connected, via FastEthernet0/0
Route metric is 0, traffic share count is 1

```

The `show ip route 172.16.1.0` command reveals that the distance is 0 for that directly connected route.

Packet Tracer Activity

Viewing Routing Table Information—show ip route (3.4.4)

In this activity, you will use a version of the `show ip route` command to see details of routing table entries. Use file e2-344.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Summary

Dynamic routing protocols are used by routers to automatically learn about remote networks from other routers. In this chapter, you were introduced to several different dynamic routing protocols.

You learned the following about routing protocols:

- They can be classified as classful or classless.
- They can be a distance vector, link-state, or path vector type.
- They can be an interior gateway protocol or an exterior gateway protocol.

The differences in these classifications will become better understood as you learn more about these routing concepts and protocols in later chapters.

Routing protocols not only discover remote networks but also have a procedure for maintaining accurate network information. When there is a change in the topology, it is the function of the routing protocol to inform other routers about this change. When there is a change in the network topology, some routing protocols can propagate that information throughout the routing domain faster than other routing protocols.

The process of bringing all routing tables to a state of consistency is called convergence. Convergence is when all the routers in the same routing domain or area have complete and accurate information about the network.

Metrics are used by routing protocols to determine the best path or shortest path to reach a destination network. Different routing protocols can use different metrics. Typically, a lower metric means a better path. Five hops to reach a network is better than ten hops.

Routers sometimes learn about multiple routes to the same network from both static routes and dynamic routing protocols. When a Cisco router learns about a destination network from more than one routing source, it uses the administrative distance value to determine which source to use. Each dynamic routing protocol has a unique administrative value, along with static routes and directly connected networks. The lower the administrative value, the more preferred the route source. A directly connected network is always the preferred source, followed by static routes and then various dynamic routing protocols.

All the classifications and concepts in this chapter will be discussed more thoroughly in the rest of the chapters of this course. At the end of this course, you might want to review this chapter to get a review and overview of this information.

Activities and Labs

The activities and labs available in the companion *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-204-4) provide hands-on practice with the following topics introduced in this chapter:



Activity 3-1: Subnetting Scenario 1 (3.5.2)

In this activity, you have been given the network address 192.168.9.0/24 to subnet and provide the IP addressing for the network shown in the topology diagram.



Activity 3-2: Subnetting Scenario 2 (3.5.3)

In this activity, you have been given the network address 172.16.0.0/16 to subnet and provide the IP addressing for the network shown in the topology diagram.



Activity 3-3: Subnetting Scenario 3 (3.5.4)

In this activity, you have been given the network address 192.168.1.0/24 to subnet and provide the IP addressing for the network shown in the topology diagram.



Many of the hands-on labs include Packet Tracer Companion Activities, where you can use Packet Tracer to complete a simulation of the lab. Look for this icon in *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-204-4) for hands-on labs that have a Packet Tracer Companion.

Check Your Understanding

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The section “Check Your Understanding and Challenge Questions Answer Key” at the end of this chapter lists the answers.

1. What are two advantages of static routing over dynamic routing?
 - A. Configuration is less error prone.
 - B. More secure because routers do not advertise routes.
 - C. Growing the network usually does not present a problem.
 - D. No computing overhead.
 - E. Administrator has less work maintaining the configuration.

2. Match the description to the proper routing protocol.

Routing protocols:

RIP

IGRP

OSPF

EIGRP

BGP

Description:

- A. Path vector exterior routing protocol:
 - B. Cisco advanced interior routing protocol:
 - C. Link-state interior routing protocol:
 - D. Distance vector interior routing protocol:
 - E. Cisco distance vector interior routing protocol:
3. Which statement best describes convergence on a network?
- A. The amount of time required for routers to share administrative configuration changes, such as a password change, from one end of a network to the other end
 - B. The time required for the routers in the network to update their routing tables after a topology change has occurred
 - C. The time required for the routers in one autonomous system to learn routes to destinations in another autonomous system
 - D. The time required for routers running disparate routing protocols to update their routing tables
4. Which of the following parameters are used to calculate metrics? (Choose two.)
- A. Hop count
 - B. Uptime
 - C. Bandwidth
 - D. Convergence time
 - E. Administrative distance

5. Which routing protocol has the most trustworthy administrative distance by default?
 - A. EIGRP internal routes
 - B. IS-IS
 - C. OSPF
 - D. RIPv1
 - E. RIPv2
6. How many equal-cost paths can a dynamic routing protocol use for load balancing by default?
 - A. 2
 - B. 3
 - C. 4
 - D. 6
7. Which command will show the administrative distance of routes?
 - A. R1# **show interfaces**
 - B. R1# **show ip route**
 - C. R1# **show ip interfaces**
 - D. R1# **debug ip routing**
8. When do directly connected networks appear in the routing table?
 - A. When they are included in a static route
 - B. When they are used as an exit interface
 - C. As soon as they are addressed and operational at layer 2
 - D. As soon as they are addressed and operational at layer 3
 - E. Always when a **no shutdown** command is issued
9. Router1 is using the RIPv2 routing protocol and has discovered multiple unequal paths to reach a destination network. How will Router1 determine which path is the best path to the destination network?
 - A. Lowest metric.
 - B. Highest metric.
 - C. Lowest administrative distance.
 - D. Highest administrative distance.
 - E. It will load-balance between up to four paths.

10. Enter the proper administrative distance for each routing protocol.
 - A. eBGP:
 - B. EIGRP (Internal):
 - C. EIGRP (External):
 - D. IS-IS:
 - E. OSPF:
 - F. RIP:
11. Designate the following characteristics as belonging to either a classful routing protocol or a classless routing protocol.
 - A. Does not support discontinuous networks:
 - B. EIGRP, OSPF, and BGP:
 - C. Sends subnet mask in its routing updates:
 - D. Supports discontinuous networks:
 - E. RIP version 1 and IGRP:
 - F. Does not send subnet mask in its routing updates:
12. Explain why static routing might be preferred over dynamic routing.
13. What are four ways of classifying dynamic routing protocols?
14. What are the most common metrics used in IP dynamic routing protocols?
15. What is administrative distance, and why is it important?

Challenge Questions and Activities

These questions and activities require a deeper application of the concepts covered in this chapter. You can find the answers at the end of this chapter.

1. It can be said that every router must have at least one static route. Explain why this statement might be true.
2. Students new to routing sometimes assume that bandwidth is a better metric than hop count. Why might this be a false assumption?

To Learn More

Border Gateway Protocol (BGP) is an inter-autonomous routing protocol—the routing protocol of the Internet. Although BGP is only briefly discussed in this course (it is discussed more fully in CCNP), you might find it interesting to view routing tables of some of the Internet core routers.

Route servers are used to view BGP routes on the Internet. Various websites provide access to these route servers, for example, <http://www.traceroute.org>. When choosing a route server in a specific autonomous system, you will start a Telnet session on that route server. This server is mirroring an Internet core router, which is most often a Cisco router.

You can then use the **show ip route** command to view the actual routing table of an Internet router. Use the **show ip route** command followed by the public or global network address of your school, for example, **show ip route 207.62.187.0**.

You will not be able to understand much of the information in this output, but these commands should give you a sense of the size of a routing table on a core Internet router.

Check Your Understanding and Challenge Questions Answer Key

Check Your Understanding

1. B, D. Static routes are considered more secure because they are not propagated between routers and therefore are not susceptible to snooping or malicious attacks. Dynamic routes can be secured using authentication. Static routes require no computing overhead because they are not propagated between routers. Note: There is some computing overhead with static routes, but it is minimal.
2. Answers:
 - A. Path vector exterior routing protocol: BGP
 - B. Cisco advanced interior routing protocol: EIGRP
 - C. Link-state interior routing protocol: OSPF
 - D. Distance vector interior routing protocol: RIP
 - E. Cisco distance vector interior routing protocol: IGRP
3. B. Convergence is the time required by routers to have complete and accurate information about the network.
4. A, C. Hop count is used by RIP. Bandwidth is used by IGRP, EIGRP, and OSPF. The other choices are not valid routing protocol metrics.
5. A. Given these choices, EIGRP internal routes are the most trustworthy with the lowest administrative distance of 90. IS-IS has an administrative distance of 115, OSPF has an administrative distance of 110, and RIP, which includes both version 1 and 2, has an administrative distance of 120.
6. C. By default, Cisco routers can load-balance up to four equal-cost paths. The maximum number of equal-cost paths depends on the routing protocol and IOS version.
7. B. The **show ip route** command displays route entry information, including the administrative distance. The administrative distance is the first number in brackets, followed by the metric. For example, [120/2] shows an administrative distance of 120 (RIP) and a metric of 2 (hop count).
8. D. A directly connected network will appear in the routing table when it is addressed and operational at Layer 3, in other words, when it has been configured with an IP address and subnet mask, and the interface and line protocol are both in the up state.
9. A. Whenever any routing protocol has multiple paths to the same network, it will choose the path with the lowest metric. This is the route that is added to the routing table.

10. Answers:

- A. eBGP: 20
- B. EIGRP (Internal): 90
- C. EIGRP (External): 170
- D. IS-IS: 115
- E. OSPF: 110
- F. RIP: 120

11. Answers:

- A. Does not support discontinuous networks: classful routing protocol
- B. EIGRP, OSPF, and BGP: classless routing protocol
- C. Sends subnet mask in its routing updates: classless routing protocol
- D. Supports discontinuous networks: classless routing protocol
- E. RIP version 1 and IGRP: classful routing protocol
- F. Does not send subnet mask in its routing updates: classful routing protocol

12. Static routing is more secure, uses less router computational power, and is easier to understand. It is more secure because routers do not advertise routing information to other routers. It uses less router resources than dynamic routing, which requires the implementation of algorithms and the processing of update packets. It is often easier to understand than some of the more complex routing protocols.

13. Dynamic routing protocols can be classified as either interior or exterior, distance vector or link-state, classful or classless, and by speed of convergence.

14. Hop count, bandwidth, delay, and cost.

15. Administrative distance is a measure of the trustworthiness of a route source. It is used when a router has learned routes to the same destination from two different route sources. It is important because not all route sources are equal. For example, you certainly would not want a router sending traffic to another router if the destination is a directly connected network! Administrative distance ensures that this does not happen because directly connected routes are trusted over all other route sources.

Challenge Questions and Activities

- 1.** Every router that forwards user traffic to the Internet will have at least one static route. That static route would be a default route. Every household that has a router to connect to the Internet uses a static default route to send all traffic to the ISP.
- 2.** Hop count can use the better path if the path chosen by a bandwidth metric is saturated with traffic.

