CHAPTER 4

IP Routing Principles

The topics in this chapter are the basis of the BSCN course that feeds the Routing exam. The very name of the exam—the Routing exam—suggests that there are more than likely some questions on this very topic. In this chapter, the concepts of routing with IP and the mechanics of the process are dealt with generically as a foundation for the subsequent chapters, which deal with the individual routing protocols. The topics will directly reflect questions on the Routing exam. If you do not understand the contents of this chapter, it will be impossible for you to pass the exam. Some of this chapter reviews subjects dealt with in the ICND course. The subsequent chapters assume the comprehension of the subjects covered in this chapter.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place, for easy reference.
- When you take a quiz, write down your answers. Studies show that retention significantly increases by writing down facts and concepts, even if you never look at the information again.
- Use the diagram in Figure 4-1 to guide you to the next step.

Figure 4-1 *How to Use This Chapter*

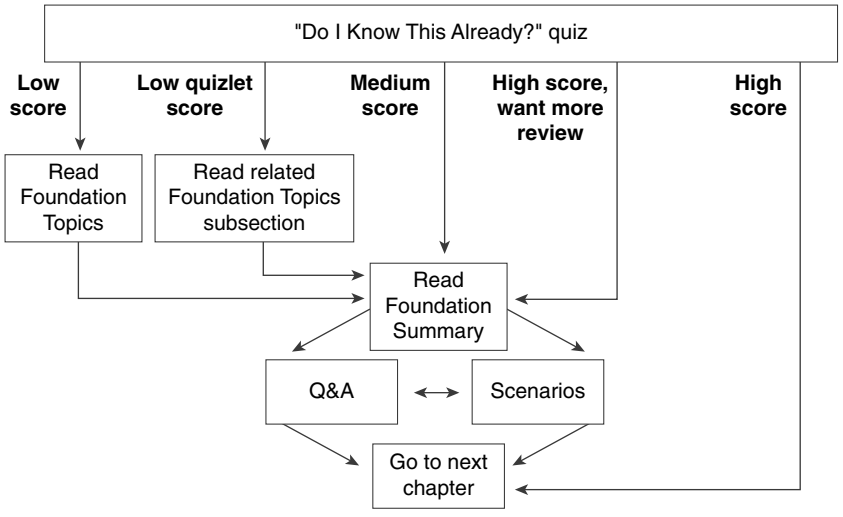

If you skip to the Foundation Summary, Q&A, and scenarios sections and have trouble with the material there, you should go back to the Foundation Topics section.

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

This 16-question quiz helps you determine how to spend your limited study time. The quiz is sectioned into four smaller four-question "quizlets," which correspond to four major topics in the chapter. Figure 4-1 outlines suggestions on how to spend your time in this chapter. Use Table 4-1 to record your scores.

Table 4-1 *Score Sheet for Quiz and Quizlets*

Quizlet Number	Topic	Questions	Score
1	The routing process	1 to 4	
2	Classful and classless routing protocols	5 to 8	
3	Distance vector and link-state routing protocols	9 to 12	
4	The routing table	13 to 16	
All questions	All	1 to 16	

1 Cisco distinguishes between the routing and the switching function—what is the difference?

2 State the two ways that an outgoing interface is selected as the preferred path.

3 What is administrative distance?

4 If IGRP has three paths to a remote network in which each path has an equal metric, what will happen?

5 Name the interior IP routing protocols that send the mask with the routing update.

6 Name the interior routing protocol that sends a routing update on a Cisco router every 30 seconds by default.

7 Does VLSM require a classful or classless routing protocol, and why?

_____

_____

_____

8 State one of the characteristics of a classful routing protocol.

_____

_____

_____

9 A distance vector routing protocol uses the mechanism of poison reverse—what is this?

_____

_____

_____

10 Name two distance vector routing protocols.

_____

_____

_____

11 Name two link-state IP routing protocols.

_____

_____

_____

12 Describe the mechanism of split horizon.

_____

_____

_____

13 What is the command syntax to empty the Cisco routing table of all its routes?

_____

_____

_____

14 What does 0.0.0.0 signify in an IP routing table?

_____

_____

_____

15 What is the command to show whether a specific network, such as 141.131.6.16, is present in the routing table?

_____

_____

_____

16 What is the next logical hop in the routing table?

_____

_____

_____

The answers to this quiz are found in Appendix A, "Answers to Quiz Questions." The suggested choices for your next step are as follows:

- **2 or less on any quizlet**—Review the appropriate sections of the "Foundation Topics" portion of this chapter, based on Table 4-1. Then move on to the "Foundation Summary" section, the "Q&A" section, and the "Scenarios" at the end of the chapter.
- **8 or less overall score**—Read the entire chapter. This includes the "Foundation Topics" and "Foundation Summary" sections, the "Q&A" section, and the "Scenarios" at the end of the chapter.
- **9 to 12 overall score**—Begin with the "Foundation Summary" section, and then go to the "Q&A" section and the "Scenarios" at the end of the chapter. If you have trouble with these exercises, read the appropriate sections in "Foundation Topics."
- **13 or more overall score**—If you want more review on these topics, skip to the "Foundation Summary" section, and then go to the "Q&A" section and the "Scenarios" at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Introduction: What Is a Routing Protocol?

The following case study illustrates some of the common concerns that a small company faces as it introduces routing into a switched/bridged network. This is a major transition and requires careful analysis to make the correct decisions regarding the choice of routing protocol as well as its implementation. This case study illustrates some of these issues with the Mental Merge company.

Case Study

A small startup company, Mental Merge, the company seen in the last chapter, still has a flat network. It has deployed several switches and an addressing scheme that divides the network into logical networks. The company has deployed virtual LANs (VLANs) in some areas, but it is not routing between them. On the end stations that need to communicate with servers on another network, Mental Merge has a problem, and these clients are in the ridiculous situation of having to physically connect from a system connected to the appropriate network. Those systems that are on different logical networks, but in the same VLAN the end systems have, overcome the routing problem by making themselves their own default gateway.

Despite the implementation of VLANs, the traffic is increasing, and congestion is beginning to cause delays.

It is time to purchase routers, to install them, and to turn on routing.

It is important for Mental Merge to understand the routing process to install the routers in the right place and to ensure that the network addressing is correct. The company also needs to choose a routing protocol for IP.

This chapter deals with these subjects and considers the network for Mental Merge in the "Scenarios" section.

What Is a Routing Protocol?

This section covers the definition, purpose, and operation of a routing protocol. It also covers the difference between a routing and routed protocol. It is necessary to understand the meaning of a protocol to understand exactly what a routing protocol is trying to achieve.

The Definition of a Routing Protocol

In simple terms, a protocol is an agreed set of rules that determine how something will operate.

A routing protocol is a set of rules that describes how Layer 3 routing devices will send updates between each other about the available networks. If more than one path to the remote network exists, the protocol also determines how the best path or route is selected.

The Purpose of a Routing Protocol

A routing protocol is the mechanism used to update the Layer 3 routing devices. When they all have the same accurate understanding of the network, they can route the data across the best path.

How the Routing Protocol Works

Participating routers advertise the routes that they know about to their neighbors in routing updates. Routes learned from routing updates are held in the routing table.

Routing and Routed

It is important to distinguish between the datagram and the routing protocol used to determine the path of the datagram.

The distinction is between the *routed* and the *routing* protocol.

The routed protocol is the Layer 3 protocol used to transfer data from one end device to another across the network. The routed protocol is the Layer 3 datagram that carries the application data as well as the upper-layer information.

The routing protocol is the protocol used to send updates between the routers about the networks that exist in the organization, thereby allowing the routing process to determine the path of the datagram across the network.

Table 4-2 provides a list of routed protocols and their corresponding interior routing protocols.

Table 4-2 *Routing and Routed Protocols*

Routed Protocol	Corresponding Interior Routing Protocol[1]
AppleTalk	RTMP, AURP, EIGRP
IPX	RIP, NLSP, EIGRP
Vines	RTP
DECnet IV	DECnet
IP	RIPv1, RIPv2, OSPF, IS-IS, IGRP, EIGRP

[1]IGRP and EIGRP are Cisco Systems proprietary routing protocols.

The router will reference the routing table and make a decision about forwarding data packets to the end destination identified in the destination address of the datagram/packet.

Table 4-3 shows the fields that are present in a typical routing table.

Table 4-3 *The Routing Table*

Network	Outgoing Interface	Metric	Next Logical Hop
140.100.100.0 /24	E0	6	131.108.13.15
140.100.110.0 /24	E0	7	131.108.13.15
140.100.120.0 /24	E0	8	131.108.13.15
140.100.130.0 /24	E0	8	131.108.13.15
166.99.0.0 /16	E1	10	131.108.14.11
166.90.0.0 /16	E1	11	131.108.14.11
145.0.88.0 /24	S0	3	131.108.10.9

It is useful to look at each field in the routing table to determine the functionality of the table to the routing process. The next sections cover the following fields of the routing table:

- The Network field
- The Outgoing Interface field
- The Metric field
- The Next Logical Hop field

The Network Field

The Network field contains the networks that the router knows exist in the organization. These entries either were entered manually as *static routes* or *default routes*, or were learned via a routing protocol as *dynamic routes*.

The Purpose of the Network Field

When a datagram comes into the router, the routing process attempts to forward it to the remote network, where it is hoped that it will find the destination host. To achieve this, it must know that the remote network exists. It determines this by looking in the routing table for the remote network.

How the Network Field Is Used

Typically, only the network portion of the address is stored in the table. Using the hierarchical strength of the addressing keeps the routing table small and the lookup short. The routing process makes a decision based on the longest match. This ensures that if VLSM has been deployed, the most specific network is chosen. Cisco IOS code mandates that the longest match can be a /32 or 255.255.255.255 mask. This is a match based on the full host address and is used in specific situations such as an OSPF environment. It is not encouraged as a common configuration because the size of the routing table grows rapidly.

The routes in the table are held in an order that speeds up the lookup process, ensuring that the routing decision is streamlined.

Later in the chapter, in the section "How the Routing Table Is Kept Current and Correct," you will see how the networks are placed in the table and how path selection to a remote network is chosen.

The Outgoing Interface Field

The Outgoing Interface is the interface on the router to which the routing process sends the datagram. This is the first step of its journey, the exit point of the router.

The Purpose of the Outgoing Interface Field

It is necessary for the routing process to know which interface queue to use to send the outbound datagram. It also informs the administrator of the interface through which the network was heard in the routing update—or, more accurately, the interface through which the chosen network was heard.

In summary, the outgoing interface field in the routing table indicates the following:

- Which interface to send the datagram to
- Which interface the routing update came through

The Metric Field

The metric is a value that is assigned to each path based on the criteria specified in the routing protocol. The Metric field is used to determine which path to use if there are multiple paths to the remote network. The metric used depends on the routing protocol.

This value is used to choose between different paths to the same destination network, to select the best path. If the values are the same, either the router selects the path that it heard first, or it uses both paths, sending the datagrams across each route.

It is the responsibility of the end device to reassemble the datagrams before sending them to the application.

Table 4-4 shows the metrics used by the different routing protocols.

Table 4-4 *Routing Protocol Metrics*

Routing Protocol	Metric
RIPv1	Hop count.
IGRP	Bandwidth, delay, load, reliability, MTU.
EIGRP	Bandwidth, delay, load, reliability, MTU.
OSPF	Cost. (The Cisco default states that the cost of an interface is inversely proportional to the bandwidth of that interface. A higher bandwidth indicates a lower cost.)
IS-IS	Cost.

NOTE By default, on a Cisco router, if multiple equal cost paths exist in IP, up to six paths are used in a round-robin manner to load balance the traffic across the network.

The Next Logical Hop Field

The next logical hop is the destination address of the next forwarding router. The address of the next logical hop will be on the same subnet as the outgoing interface.

The Purpose of the Next Logical Hop Field

The purpose of identifying the next logical hop is so that the router can create the Layer 2 frame with the destination address. The reason that the logical address is stored instead of the MAC address of the next hop is to ensure that the information is accurate. The MAC address may change because of changes in the hardware; however, such changes do not affect the logical address. Also, the router is dealing at Layer 3 and just examines the source address of the routing update to determine the next hop. The simplicity of this action reduces the need for extra computation and memory.

TIP It is useful for troubleshooting to remember that the next logical hop address is the address of the router directly connected to the forwarding router. Therefore, the address of the next logical hop shares the same subnet as the determining router.

The following section gives an example of a routing table. In the exam, you may be asked to interpret the output of the **show IP route** command, and it is necessary, therefore, to be able to extrapolate information from this table. The following section does this.

The show ip route Command

```
Router# show ip route
```

This command is used to show the IP routing table on the router. It details the network as known to the router and its sources for the information (such as the routing protocols). This command is excellent for troubleshooting configuration errors and understanding how the network is communicating about its routes.

To see a particular network in the routing table, issue this command:

```
Router# show ip route network number
```

Example 4-1 shows the output of this command. Table 4-5 explains how to read this information.

Example 4-1 **show ip route** *Output*

```
SanJose#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, * - candidate default
       U - per-user static route, o - ODR
       T - traffic engineered route

Gateway of last resort is not set

     140.100.0.0/28 is subnetted, 3 subnets
C       140.100.17.192 is directly connected, FastEthernet3/0
C       140.100.17.128 is directly connected, FastEthernet1/0
C       140.100.32.0 is directly connected, Fddi2/0


Bldg_1#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, * - candidate default
       U - per-user static route, o - ODR

Gateway of last resort is not set

     140.100.0.0/28 is subnetted, 3 subnets
O       140.100.17.192 [110/20] via 140.100.17.129, 00:07:44, Ethernet0
C       140.100.17.128 is directly connected, Ethernet0
O       140.100.32.0 [110/11] via 140.100.17.129, 00:07:44, Ethernet0
```

Table 4-5 explains the meaning of the important fields.

Table 4-5 *Explanation of the* **show IP route** *Command That Was Performed on Router Building 1*

Field	Explanation
O	Indicates the protocol that derived the route. Possible values include the following: **I**—IGRP-derived **R**—RIP-derived **O**—OSPF-derived **C**—Connected **S**—Static **E**—EGP-derived **B**—BGP-derived **i**—IS-IS-derived
140.100.17.192	Indicates the address of the remote network.
[110/20]	The first number in the brackets is the administrative distance of the information source; the second number is the metric for the route.
via 140.100.17.129	Specifies the address of the next router to the remote network.
00:07:44	Specifies the last time that the route was updated in hours:minutes:seconds.
Ethernet0	Specifies the interface through which the specified network can be reached.

These commands are useful to verify that the configuration has worked and that the OSPF network is functioning correctly. In a single-area environment, the full complexity of OSPF is not engaged. The full strength and complexity of OSPF come to the fore in the design and configuration of a multiarea network.

How the Routing Table Is Kept Current and Correct

The capability to send traffic from one end of the network to the other depends on how accurate and current the routing table in every router is within the network. Although all routing protocols have this written into their mission statements, the more recent routing protocols are more efficient, so their networks scale more easily. For example, RIP will send out the entire routing table every 30 seconds, while OSPF updates contain only the change and are sent only when that change occurs. Although OSPF sends the entire table every 30 minutes after the last update, this is far less demanding of network resources than the older protocol, RIP.

The accuracy of the table will be affected by how quickly it responds to changes in the network. These changes include the following:

- Learning new networks
- Learning a better path to an existing network
- Learning that a network is no longer available
- Learning an alternative route to a network

How each of these changes is achieved depends on the routing protocol.

Emptying the contents of the routing table and thus force the router to learn the information about the network is very useful in troubleshooting a network.

This command empties all the routes from the table:

```
Router# clear ip route *
```

This command removes the specific network from the table:

```
Router# clear ip route network
```

Switching Versus Routing

Cisco makes a distinction between the routing function and the switching function of a router. The difference is simple: Two jobs within a router need to be done to move a datagram from an incoming interface to the outgoing interface.

The Routing Function

The *routing function* is responsible for learning the logical topology of the network and then making decisions based on that knowledge. The decisions determine whether the incoming datagram can be routed and, if so, how.

The decisions include these:

- Is the protocol stack configured on the router?
- Is there an entry for the remote network in the routing table?
- Is there a default network configured?
- Is the network reachable?
- Which is the best path to that remote network?
- Are there equal-cost multiple paths?
- To which outgoing interface(s) should the datagram(s) be queued?

The Switching Function

The *switching function* is concerned with moving data across the router. It is responsible for forwarding the datagram. Switching takes over after the routing decisions have been made. Although the router has lookups to make, the few decisions that need to be made are performed in hardware. Therefore, this function is very fast.

The switching function does the following:

- Checks the incoming frame for validity
- Checks whether the frame is addressed (at Layer 2) to the router
- Checks whether the frame is within the scope of the framing criteria (too big or too small)
- Checks whether the frame passes CRC
- Strips the Layer 2 header and trailer from the frame, and checks the destination address against the cache entries
- Creates the appropriate frame header and trailer (if there is an entry in cache for the destination address), and forwards the frame to the outbound interface queue

---

TIP Please note that the preceding section refers to the internals of the IOS, which is extremely complex. It has been described at the level required by the Routing exam. If live networks are to be designed and configured, it would be wise to understand in more depth some of the issues concerning caching and the placement of access lists. This information is readily available on the Cisco web page.

Functionality is broken into two components to ensure that the process is as fast as possible. After the routing decisions are made, the Cisco router caches the result, allowing subsequent datagrams to be switched.

The Routing/Switching Relationship in a Cisco Router

A packet transiting the router is accepted into the router if the frame header (of the frame in which the packet resides) contains the Layer 2 address of one of the router's interfaces. If properly addressed, after the framing is checked, the frame and its content (the packet) are buffered, pending further processing. The buffering occurs in main memory or some other specialized memory location.

If the source and destination Layer 3 address of the datagram have not been seen by this router before, the datagram will be process switched (routed). This involves the following actions:

1. When a datagram is to be forwarded, a process initiates a lookup in this routing table and a decision about how the datagram should be forwarded.

2 The packet is then encapsulated.

3 If fast switching is enabled, the packet is then examined again, and an entry is put into a route cache. The entry in this cache consists of the following:

   — An IP prefix

   — The output interface

   — The link-layer header to be used in forwarding the packet

On subsequent packets, if the IP destination matches a prefix found in the route cache, the packet is forwarded using this information. The routing function is not disturbed, nor are the CPU cycles required to feed this monster expended.

The type of route cache used depends on the hardware used. The caches available are called *fast switching*, *autonomous switching*, *silicon switching*, and *Cisco Express Forwarding (CEF)*. If CEF switching is used, then the story changes again. With CEF switching, each card runs its own copy of the express forwarding and has its own copy of a *forwarding information base (FIB)*. In the event of a routing change, this new entry is forwarded by the CPU to each separate line card.

Types of Routing Protocols

Although the switching and routing functions within the router are set, there are many differences to be seen among the different routing protocols.

The routing protocols are essentially applications on the router. Their purpose is to ensure the correct and timely exchange of information about the network between the routers so that the routers can successfully perform the routing and switching functions described previously.

IP routing protocols can be divided into several distinct groups. The first is the difference between protocols that send the mask in the updates and the older protocols that do not. These are labeled classless and classful protocols, respectively.

Classful and Classless Routing Protocols

Classful routing protocols do not carry the subnet or routing mask in the update. The older distance vector protocols tend to be classful. This incapability to carry the subnetting information leads to design constraints in the IP network.

Classful Routing

Classful IP routing protocols include RIPv1 and IGRP. The characteristics of a classful routing protocol are listed here:

- Summarization occurs at the network boundary.

- Routes exchanged between foreign networks are summarized to the NIC number network boundary.
- Within the same network (NIC number), subnet routes are exchanged by routers, without the mask.
- The subnet mask is assumed to be consistent for a NIC number used within a network, so all router interfaces must share the subnet mask for interfaces in the same NIC network.
- The utilization of address space may be inefficient.
- VLSM is not possible within the network.

Classless Routing

Classless routing protocols were designed to overcome the constraints listed previously. The routing protocols that can do this are OSPF, EIGRP, RIPv2, IS-IS, and BGP.

The characteristics of a classless routing protocol are listed here:

- Router interfaces within the same network can have different subnet masks (VLSM).
- Some of the classless routing protocols, including BGP-4 and RIP v2, support the use of classless interdomain routing (CIDR).
- Some routes can be summarized within the major NIC number. This is done manually.

NOTE BGP-4 and EIGRP summarizes at the network boundary automatically. Summarization within the NIC number boundary must be configured manually.

The distinctions of classless and classful are important between routing protocols. Another important distinction is based on the technology that they employ.

IP routing protocols use two main technologies: *link-state* and *distance vector* technologies**.** These are discussed next.

Distance Vector and Link-State Routing Protocols

Distance vector protocols are the earliest protocols, and they include RIPv1 and IGRP. These protocols are classful protocols, but RIPv2 and EIGRP are examples of classless routing protocols.

Distance vector protocols were designed for small networks. As the networks started to expand, enhancements were made to the distance vector protocols (RIPv2 and IGRP). At the same time, link-state protocols such as OSPF were introduced

Distance Vector Routing Protocols

Distance vector protocols send periodic updates. These updates are sent to directly connected neighbors. The update is periodic because it waits for the timer to expire before it sends an update. After receiving a neighbor's routing table, the router updates its table and sends the modified table in subsequent updates. This is the reason that distance vector routing protocols are said to be "routing by rumor."

The purpose of the protocol is to provide accurate, loop-free information to the routers. The update affects the entire routing table, excluding those networks that were learned through the interface through which the update is being sent. This is in accordance to the *split horizon rule*; this reduces network overhead and also prevents information from traveling in circles through the network, which can create *routing loops*.

To prevent routing loops, distance vector routing protocols employ the following techniques, which are described in more depth in the section on RIPv1:

- Split horizon
- Count to infinity
- Poison reverse
- Hold-down
- Triggered updates
- Aging of routes from the routing table

NOTE Although EIGRP is defined by Cisco as an advanced distance vector routing protocol, it has adopted some of the link-state characteristics in favor of the distance vector solution. It does not use either count to infinity or hold-down timers.

The Distance Vector Routing Metrics

The metric used by distance vector protocols is often stated as being distance measured in the number of *hand-off points* or *hops* (routers) encountered on the way to the end device. Cisco defines IGRP and EIGRP as distance vector routing protocols. This muddies the original definition because IGRP and EIGRP use a composite and complex metric.

The path selection is made using the Bellman Ford algorithm based on the metric or value of each available path. RFC 1058 discusses this in depth in reference to RIPv1. EIGRP, however, uses a proprietary algorithm called *Diffusing Update Algorithm (DUAL)*.

TIP If you are asked a question on distance vector metrics, it may be wise to use the original definition of hop count because IGRP and EIGRP are proprietary protocols. Cisco also uses the original definition in its documentation.

Link-State Routing Protocols

A link-state routing protocol is a sophisticated protocol dedicated to maintaining loop-free, accurate tables. It does not send the entire routing table via broadcasts every 30 seconds, as the original distance vector protocols (such as RIPv1) did, but it instead utilizes multicast addressing and incremental updates. Some routing protocols may be sent updates every 30 minutes (not 30 seconds) in addition to the incremental ones. Table 4-6 is a summary of IP routing protocols and the update timers.

Table 4-6 *A Summary of IP Routing Protocols and the Update Timers*

Protocol	Update Timer	Technology
RIPv1	Every 30 seconds, for entire routing table.	Distance vector.
OSPF	Incremental, with only the network change. However, 30 minutes after the last update was received, a compressed version of the table is propagated.	Link state.
EIGRP	Incremental updates, with network change only.	Advanced distance vector.
IGRP	Updates every 90 seconds, with incremental updates as needed.	Distance vector.
BGP-4	Incremental, with only the network change.	Path vector (an exterior routing protocol). The term refers to the list of autonomous system numbers that are carried in the BGP-4 updates, and the vector indicates the direction to send the traffic to find the path to a remote network.
IS-IS	Incremental, with only the network change. However, approximately 15 minutes after the last update was received, a compressed version of the table is propagated.	Link state.

The Meaning of Link State

As with a distance vector router, information is exchanged only with its directly connected neighbor. Unlike distance vector protocols, the information concerns only the local links (not the routes) connected to the router, and these links are propagated, unchanged, to every other router in the network. Therefore, every router has the same image of the network, created from the original updates from every other router in the network.

The purpose of link-state routers is to reduce the network overhead of the routing updates that are both current and thus accurate, allowing it to scale to large networks.

Sending an update about links is more efficient than sending data about routes because one link may effect many routes. Sending the links allows the routers to compute the routes that may be affected. The resources used are router CPU rather than network bandwidth.

Learning About the Network

A link-state routing protocol develops a relationship with an adjacent router, one that is on the same physical network. The two also must have the same subnet mask and have the same hello timers. The routing protocol develops and maintains the relationship by sending a simple message across the medium. When another router replies, it is identified as a *neighbor* for the routing process. This neighbor relationship is maintained as long as the simple message (Hello protocol) is received. Because the neighbor relationship is continuous, information can be exchanged between the routing processes quickly and efficiently. Therefore, changes in the network are realized very quickly.

Link-state routing protocols are used in larger networks because the method that they use to update the routing tables requires fewer network resources.

Learning About a Change in the Network

A router knows very quickly whether the neighbor, which may also be the next logical hop, is dead because the router no longer receives Hello protocol messages.

The routing process sends out a message immediately when it identifies a problem, without waiting for the update timer to expire. This is known as an *incremental update*. The update contains only the relevant information. The router also remains silent if there is no change in the network.

The incremental update improves *convergence* time and also reduces the amount of information that needs to be sent across the network. The network overhead on the physical media is eased, and the potential throughput of the network is improved.

Updating Local Network Tables

A link-state protocol holds a topology map of the network and can easily update the map and routing table database, via the incremental updates. In OSPF, these are called link-state advertisements (LSAs). After an update is received and forwarded, the router will compute a new topology map and, from this, a new path. It uses the Dijkstra algorithm to achieve this new understanding of the network.

Path Selection

The metric that is used is stated as cost, although many vendors supply a default that may be overridden manually. This is true of Cisco's implementation of OSPF, which uses bandwidth and delay as its default.

Examples of link-state routing protocols for IP are OSPF and IS-IS.

Another distinction that needs to be made between routing protocols is the difference between interior and exterior protocols. Interior protocols are those that update routers within an organization. Exterior protocols are those that update routers that connect different organizations to each other or to the Internet.

Interior and Exterior Routing Protocols

Routing protocols that operate within an organization are referred to as *interior routing protocols* (for example, RIPv1, IGRP, EIGRP, OSPF, and IS-IS).

Interior Routing Protocols

The boundaries of the organization are defined as the *autonomous system*. The unique number assigned to the autonomous system then identifies the organization. The autonomous system number may be viewed as another layer of hierarchy in the IP addressing scheme because the number can represent a collection of NIC numbers.

Exterior Routing Protocols

Routing protocols that exchange routing information between organizations are known as *exterior routing protocols*. Exterior routing protocols are highly complex. The complexity arises from the need to determine policies between different organizations. Border Gateway Protocol Version 4 (BGP-4) is an example of an exterior gateway protocol.

NOTE This next section deals briefly with an older distance vector routing protocol, RIP; an improved distance vector routing protocol, IGRP; and a link-state routing protocol, OSPF. RIPv1 and IGRP are discussed here because they are not dealt with in the subsequent chapters. Although OSPF is dealt with in greater detail in the following chapters, it lends an interesting contrast to the two distance vector protocols.

RIP Version 1

Routing Information Protocol version 1 (RIPv1) is a simple routing protocol and, as such, works well in a small environment. As a distance vector routing protocol, it sends updates every 30 seconds. These updates comprise the entire routing table.

RIPv1 will support the following:

- **Count to infinity**—A router advertising networks heard from a neighboring router back to the same neighboring router could create a loop. In repeating networks to the router that informed the routing table, when a network goes down, each router may believe that there is an existing path through its neighbor. This problem is limited because each router increments the hop count before it sends out the update. When the hop count reaches 16, the network is rejected as unreachable because the diameter of a RIPv1 network cannot be greater than 15. This is called *counting to infinity*, where infinity equals 16. Although the liability is controlled, it will still slow convergence of the network.
- **Split horizon**—This is a mechanism to prevent loops and, thereby, the necessity of count to infinity. The routing process will not send networks learned through an interface in an update out that interface. It will not repeat information to the router that told of the networks.
- **Split horizon with poison reverse**—Split horizon on its own may not prevent loops. Poison reverse includes all the networks that have been learned from the neighbor, but it sets the metric to infinity (16). By changing the metric value to 16, the networks are reported to be unreachable. It acknowledges the network but denies a valid path. Although this increases network overhead by increasing the update size, it can prevent loops.
- **Holddown**—After deciding that a network in the routing table is no longer valid, the routing process waits for three routing updates (by default) before it believes a routing update with a less-favorable metric. Again, this is to prevent routing loops from generating false information throughout the network.
- **Triggered updates**—As soon as a routing process changes a metric for a network in its routing table, it sends an update. This informs the other routers immediately. If there is a problem in the network, all the affected routers go into holddown immediately instead of waiting for the periodic timer. This increases convergence and helps prevent loops.

- **Load balancing**—If the routing process sees multiple paths of equal cost to a remote network, it distributes the routed (datagram) traffic evenly among the paths. It will allocate datagrams to the different paths on a round-robin basis.

WARNING Because the metric used is hop count, one path may become saturated. A 56-kbps line and a 100-Mbps Fast Ethernet line may both offer paths of equal hop count; dividing the user traffic between them, but, may not optimize the bandwidth of the network.

Cisco has implemented all the preceding options, which are defined in RFC 1058**.**

RIPv1 is useful in small networks and is distributed with *Berkeley Standard Distribution* (BSD), which makes it widely available. It may not be suitable for large environments, however, because the protocol was never designed with the expectation of being used in huge organizations.

As the network grows, problems will be seen with applications timing out and congestion occurring on the network as the routers fail to adapt quickly to changes. When there has been a change in the network, the time that it takes for every router to register that change is known as the convergence time. The longer this timer takes, the greater the likelihood of problems on the network. Therefore, it is necessary either to contain the growth of the network or to use a routing protocol that scales to a larger size. OSPF is designed to scale and has the added advantage of being defined by the Internet Engineering Task Force (IETF), making it an industry standard in the public domain.

IGRP

IGRP is a distance vector routing protocol created by Cisco Systems in the mid-1980s. It is a distance vector routing protocol, but because it is proprietary, it has the advantage of being capable of improving many of the elements seen in RIPv1, including incremental updates, fewer network resources to maintain the routing protocol, a more complex and efficient metric, and no limitation in diameter of the network because of hop count.

Although IGRP can streamline its operation because it does not have to be all things to all people, it can be implemented only on Cisco routers. It is very efficient at sharing its information with other routing protocols, using redistribution.

NOTE It is unlikely that there will be direct questions on either RIPv1 or IGRP. There will be questions on the distance vector and link-state protocols. RIPv1 and IGRP are discussed here only as illustrations of distance vector protocols.

IGRP has the following characteristics of a distance vector routing protocol:

- Periodic updates
- Broadcasting updates
- Full routing table updates
- Count to infinity
- Split horizon
- Triggered updates with route poisoning
- Load balancing on equal paths (up to four, by default)
- Bellman Ford routing algorithm

It differs from RIPv1 in the following ways:

- The metric is a composite calculated from bandwidth, delay, loading, reliability, and MTU. In fact, although MTU was originally designed as part of the metric, it is tracked but not used in the calculation. It is possible to configure the use of all the calculated elements of the metric. If these are not configured, the system will use bandwidth and delay by default.
- The hop count is 100, configurable to 255 (although this is not used as a metric, but to age out datagrams).
- The update timer is set by default to 90 seconds (three times that of RIPv1).
- Unequal load sharing occurs on multiple paths.
- A more efficient packet structure is used.
- Autonomous systems (AS) are used to allow multiple processes within a routing domain, which allows the network to scale.

OSPF

NOTE OSPF is dealt with in great depth in the following chapter. It is considered here so that general comparisons may be made between the array of different IP routing protocols and the technologies underlying them. Unfortunately, this means that there may be some level of repetition between the chapters. Subheadings have been carefully worded so that the reader can read relevant information in context while avoiding repetition.

OSPF is an improvement on RIPv1 for large networks because of the following reasons:

- It utilizes bandwidth more efficiently, sending incremental updates.
- The updates are not broadcast as in RIPv1 but are directed to multicast addresses 224.0.0.5 and 224.0.0.6.
- It propagates changes in the network more quickly, with incremental updates and neighbor relationships.
- It is not limited in size by a maximum hop count of 15.
- It allows for variation in network size throughout the organization, using VLSM.
- It has security options defined in the MD5 specification.
- The metric may be defined manually, allowing for greater sophistication in the path determination.
- It is more responsive to network changes, is flexible in network addressing and design, and scales to a larger size.

The following sections discuss these key points in detail.

Key Attributes of OSPF

OSPF is designed to offer the greatest flexibility for every situation. As an open standard, it is required to offer interoperability in conjunction with this flexibility, while allowing the network to grow. These requirements make OSPF a highly complex routing protocol.

To understand this complexity, it is useful to identify the main characteristics of OSPF. These key attributes of OSPF include the following:

- Maintaining a connection-oriented relationship with other routers on the same physical segment. These are known as *adjacent neighbors*. This is a TCP connection maintained by keepalives.
- Sending the minimum amount of information in an incremental update when there has been a change in the network. This allows for fast network convergence.
- Adding another level of hierarchy to the IP address, by designing networks into *areas*.
- Using VLSM and summarization.
- Assigning specific functionality to different routers to streamline the process of communication change in the network.
- Operating within an organization as an interior routing protocol.

Path Selection Between Routing Protocols

Clearly there are many IP routing protocols from which to choose. It is better if a single routing protocol can be chosen because this gives a consistency that relates directly to the strength of the network. It complicates the network to have more than a single routing protocol attempting to perform the same job. The routing table, in particular, sometimes deals with confusion in how to select one path to place into the routing table.

When more than one routing protocol is running on the router, the routing process must make a decision to have one entry per network in the routing table. The choice cannot be based on the metric because metrics differ between protocols. Instead, another method, called administrative distance, was devised to solve the problem.

NOTE When it is necessary to have more than one routing protocol within an organization, *redistribution* is configured. However, the router that is responsible for redistribution will have more than one routing protocol informing the routing table.

Administrative Distance

The administrative distance will select one path to enter the routing table from several paths offered by multiple routing protocols.

In Figure 4-2, for example, both RIP and EIGRP have paths to the network 140.100.6.0. RIP is offering a metric of 2 hops, and EIGRP is tendering a metric of 768. Without redistribution, no conversion or choice is possible because there is no similar criteria for distinguishing the two paths. Therefore, the metric is ignored, and the administrative distance is used to make the selection.

In Figure 4-2, despite the speed of Frame Relay being set at 56 kbps as opposed to the 100 Mbps of FDDI, Router D would select the Frame Relay path based on administrative distance. In this case, manually configuring the administrative distance on Router D would be advisable.

Administrative distance is a rather arbitrary set of values placed on the different sources of routing information. The defaults can be changed, but care should be taken when subverting the natural path selection, and any manual configuration must be done with careful reference to the network design of the organization and its traffic flow.

Administrative distance reflects the preferred choice. The defaults are listed in Table 4-7.

Figure 4-2 *Path Selection Using Administrative Distance*

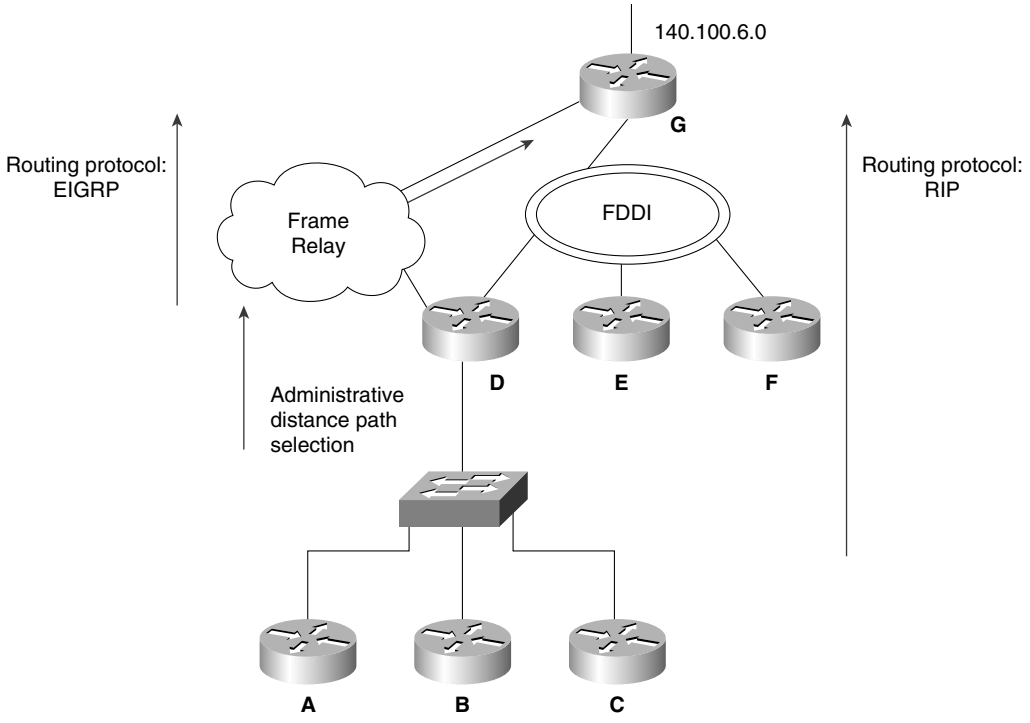

Table 4-7 *Default Administrative Distance*

Routing Source	Administrative Distance
Connected interface or static route that identifies the outgoing interface rather than the next logical hop	0
Static route	1
EIGRP summary route	5
External BGP	20
EIGRP	90
IGRP	100
OSPF	110
RIP	120
External EIGRP	170
Internal BGP	200
An unknown network	255 or infinity

The administrative distance can be manually configured. The reason for manually configuring the administrative distance for a protocol such as EIGRP is that it may have a less desirable path compared to one offered by another protocol such as RIP, which has a higher default AD.

Because the administrative distance is looked at with total disregard of the metrics, however, the EIGRP path will be selected. The other reason is that a directly connected network, which has precedence, is being used as a backup link for redundancy because the directly connected network is not used on a daily basis. Backup links for redundancy are often implemented on serial connections where the network charges are based on usage. This design is called a *floating static route.*

Convergence

Convergence occurs when all the routers in the routing domain agree on the routes that are available. Convergence time is the time that it takes for every router's routing table to synchronize after there has been a change in the network topology.

It is important to ensure that the time taken is as short as possible because while the routers disagree on the available networks, they cannot route data correctly or efficiently.

Each routing protocol has chosen a different method of updating the routing table. This affects convergence time. Some new concepts are introduced in the following comparison in protocol convergence methods. This is simply to show the relative merits of each approach. The concepts are explained in depth in the chapters that concentrate on the specific protocols. The terms shown in italics are defined in the glossary at the end of this chapter, as well as in the final glossary at the end of the book.

RIPv1 Convergence

The steps for RIPv1 convergence are as follows:

Step 1 When the local router sees a connected route disappear, it sends a *flash update* and removes the route entry from its table. This is called a *triggered update* with *poison reverse*.

Step 2 The receiving routers send flash updates and put the affected route in *holddown*.

Step 3 The originating router queries its neighbor for alternative routes. If the neighbor has an alternative route, it is sent; otherwise, the *poisoned route* is sent.

Step 4 The originating router installs the best alternative route that it hears because it has purged the original routes.

Step 5 Routers that are in holddown ignore the alternative route.

Step 6 When the other routers emerge from holddown, they will accept the alternative route.

Step 7 Convergence takes the time for detection, plus holddown, plus the number of routing updates (equal to the hop-count diameter of the network). This could take a long time.

IGRP Convergence

The steps for IGRP convergence are as follows:

Step 1 When the local router sees a connected route disappear, it sends a flash update and removes the route entry from its table. This is called a triggered update with poison reverse.

Step 2 The receiving routers send flash updates and put the affected route in holddown.

Step 3 The originating router queries its neighbor for alternative routes. If the neighbor has an alternative route, it is sent; otherwise, the poisoned route is sent.

Step 4 The originating router installs the best alternative route that it hears because it has purged the original routes. It sends a new flash update. This is either the routing table, with or without the network available, stating the higher metric.

Step 5 Routers that are in holddown ignore the alternative route.

Step 6 When the routers come out of holddown, they accept the alternative route.

When the other routers emerge from holddown, they will accept the alternative route.

Step 7 Convergence takes the time for detection, plus holddown, plus the number of routing updates (equal to the hop-count diameter of the network). Because the update timer is 90 seconds, this could take a very long time.

EIGRP Convergence

The steps for EIGRP convergence are as follows:

Step 1 When the local router sees a connected route disappear, it checks the *topology table* for a feasible successor.

Step 2 If no *feasible successor* exists, it moves into active state.

Step 3 The originating router queries its neighbor for alternative routes, the receiving router acknowledges.

Step 4 If an alternative exists, it is sent.

Step 5 If the router receives an acceptable successor, it adds the route to the table.

Step 6 A flash update of the path with the higher metric is sent out.

Step 7 Updates are acknowledged.

Convergence is very quick because it is the detection time, plus query, reply, and update time. If there is a feasible successor, then convergence is almost instantaneous.

OSPF Convergence

The steps for OSPF convergence are as follows:

Step 1 When a router detects a link failure, an LSA is sent to its neighbors. If the router is on a multiaccess link, then the update is sent to the DR and BDR, not to all neighbors.

Step 2 The path is removed from the originating router's tables.

Step 3 On receipt of the LSA, all routers update the topology table and flood the LSA out its interfaces.

Step 4 The *Dijkstra algorithm* is run to rebuild the routing table.

Convergence is detection time, plus LSA flooding, plus 5 seconds before computing the topology table. This comes to a few seconds. If convergence is deemed to be the topology table being updated, this could take longer.

Conclusion

This chapter reviewed IP routing protocols. It has shown how to untangle the various categories of routing protocols, how to analyze the IP routing table, and how to understand how routing decisions are made.

This information is crucial to the Routing exam because it is essentially an exam on IP routing. This chapter is a foundation to the other chapters, which deal in depth with particular technologies.

Foundation Summary

The "Foundation Summary" is a collection of quick reference information that provides a convenient review of many key concepts in this chapter. For those of you who already feel comfortable with the topics in this chapter, this summary will help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparations before the exam, these tables and figures will be a convenient way to review the day before the exam.

Table 4-8 shows the metrics used by the IP routing protocols.

Table 4-8 *Routing Protocol Metrics*

Routing Protocol	Metric
RIPv1	Hop count.
IGRP	Bandwidth, delay, load, reliability, MTU.
EIGRP	Bandwidth, delay, load, reliability, MTU.
OSPF	Cost. (The Cisco default states that the cost of an interface is inversely proportional to the bandwidth of that interface. A higher bandwidth indicates a lower cost.)
IS-IS	Cost.

Table 4-9 explains how to read the information in the routing table, as explained in the **show ip route** command.

Table 4-9 *Explanation of the* **show ip route** *Command*

Code	Protocol That Derived the Route
I	IGRP.
D	EIGRP.
EX	External EIGRP.
R	RIP.
C	Connected.
S	Static.
E	EGP.
B	BGP.
I	IS-IS.
L1	IS-IS level 1.
L2	IS-IS level 2.

Table 4-9 *Explanation of the* **show ip route** *Command (Continued)*

Code	Protocol That Derived the Route
M	Mobile.
U	Per-user static route.
O	ODR.
T	Traffic-engineered route.
O	OSPF networks from within the same area as the router. These are networks learned from router and network LSAs.
IA	OSPF interarea. This is sent out by the ABRs and is created from the summary link LSA (type 3 and type 4). These routes will not be seen on a router within a totally stubby area because it will not receive LSAs external to the area.
N1	OSPF NSSA external type 1.
N2	OSPF NSSA external type 2.
E1	OSPF external type 1. These routes are generated by the ASBR and show routes that are external to the autonomous system. The cost of this external route is the summarization of the external cost, plus the cost of the path to the ASBR. These routes will not be seen in a stub or totally stubby area.
E2	OSPF external type 2. These routes do not take into account the cost of the path to the ASBR. They consider only the external cost.

Table 4-10 summarizes the major differences between distance vector routing protocols and link-state routing protocols.

Table 4-10 *Distance Vector Routing Protocols Versus Link-State Routing Protocols*

Distance Vector	Link-State
Sends its entire routing table at periodic intervals out of all interfaces (typically, this is based in seconds). It will also send triggered updates to reflect changes in the network.	Sends incremental updates when a change is detected. OSPF will send summary information every 30 minutes, regardless of whether incremental updates have been sent in that time.
Typically involves updates sent using a broadcast address to everyone on the link.	Typically involves updates sent to those routers participating in the routing protocol domain, via a multicast address.
Uses a metric based on how distant the remote network is to the router. (IGRP does not conform to this as a proprietary solution.)	Is capable of using a complex metric, referred to by OSPF as cost.
Has knowledge of the network based on information learned from its neighbors.	Has knowledge of the network based on information learned from every router in the area.

continues

Table 4-10 *Distance Vector Routing Protocols Versus Link-State Routing Protocols (Continued)*

Distance Vector	Link-State
Includes a routing table that is a database viewed from the perspective of each router.	Has a topological database that is the same for every router in the area. The routing table that is built from this database is unique to each router.
Uses Bellman Ford algorithm for calculating the best path.	Uses the Dijkstra algorithm.
Does not consume many router resources, but is heavy in the use of network resources.	Uses many router resources, but is relatively low in its demand for network resources.
Maintains one domain in which all the routes are known.	Has a hierarchical design of areas that allow for summarization and growth.
Is not restricted by addressing scheme.	For effective use, the addressing scheme should reflect the hierarchical design of the network.
Involves slower convergence because information of changes must come from the entire network (but indirectly). Each routing table on every intervening router must be updated before the changes reach the remote end of the network.	Involves quicker convergence because the update is flooded immediately throughout the network.

Table 4-11 summarizes the differences between RIPv1 and OSPF. Because RIPv1 is a distance vector routing protocol and OSPF is a link-state routing protocol, you will find it helpful to keep in mind the information in Table 4-10.

Table 4-11 *RIPv1 Versus OSPF*

RIPv1	OSPF
Is a simple protocol to design, configure, and maintain.	Is a complex protocol to design and, in some instances, to configure and maintain.
Does not require a hierarchical addressing scheme.	If full benefits of the protocol are to be harnessed, should use a hierarchical IP addressing scheme.
Does not pass the subnet mask in the routing update, and therefore is not capable of classless routing or VLSM.	Carries the mask in the update, and therefore can implement VLSM, summarization, and classless routing.
Is limited to a 15-hop diameter network.	Is unlimited in the diameter of the network, although it is suggested that an area not exceed more that 50 networks.
Does not acknowledge routing updates; just repeats them periodically (every 30 seconds).	Acknowledges updates.
Has a routing table that is sent out of every interface every 30 seconds (by default).	Involves updates sent as required (when changes are seen) and every 30 minutes after no change has been seen.

Table 4-11 *RIPv1 Versus OSPF (Continued)*

RIPv1	OSPF
Can transmit information about the network in two messages: the routing update and the triggered update.	Has protocols for discovering neighbors and forming adjacencies, as well as protocols for sending updates through the network. These protocols alone add up to nine message types.
Uses hop count as a metric, the number of routers to process the data.	Uses cost as a metric. Cost is not stated in the RFCs, but it has the capacity to be a complex calculation, as seen in Cisco's implementation.

Table 4-12 summarizes default administrative distances.

Table 4-12 *Default Administrative Distance*

Routing Source	Administrative Distance
Connected interface or static route that identifies the outgoing interface rather than the next logical hop	0
Static route	1
EIGRP summary route	5
External BGP	20
EIGRP	90
IGRP	100
OSPF	110
RIP	120
External EIGRP	170
Internal BGP	200
An unknown network	255 or infinity

Chapter Glossary

This glossary provides an official Cisco definition for key words and terms introduced in this chapter. I have supplied my own definition for terms that the Cisco glossary does not contain. The words listed here are identified in the text by italics. A complete glossary, including all the chapter terms and additional terms, can be found in Appendix C, "Glossary."

adjacent neighbors—A neighbor is a router that is directly connected to another router. They must also have same mask and hello parameters. An adjacent router is a router that has exchanged routing information with its neighbor.

area—A logical set of network segments and their attached devices. Areas are usually connected to other areas via routers, making up a single autonomous system. See also *AS*. Used in DECnet, IS-IS, and OSPF.

autonomous switching—Feature on Cisco routers that provides faster packet processing by allowing the ciscoBus to switch packets independently without interrupting the system processor.

autonomous system—A collection of networks under a common administration sharing a common routing strategy. Autonomous systems may be subdivided into areas.

Berkeley Standard Distribution (BSD)—Term used to describe any of a variety of UNIX-type operating systems based on the UC Berkeley BSD operating system.

Cisco Express Forwarding (CEF)—Advanced, Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive web-based applications or interactive sessions.

classful routing protocol—A protocol that does not carry the subnet mask. A distance vector routing protocol that will not allow VLSM or route summarization.

classless routing protocol—A routing protocol that carries the subnet mask in the routing update. This allows the implementation of VLSM and summarization.

convergence—Speed and capability of a group of internetworking devices running a specific routing protocol to agree on the topology of an internetwork after a change in that topology.

count to infinity—Problem that can occur in routing algorithms that are slow to converge, in which routers continuously increment the hop count to particular networks. Typically, some arbitrary hop count limit is imposed to prevent this problem.

default routes—A route that should be used if the destination network is not present in the routing table.

Diffusing Update Algorithm (DUAL)—A convergence algorithm used in Enhanced IGRP that provides loop-free operation at every instant throughout a route computation. This allows routers involved in a topology change to synchronize at the same time, while not involving routers that are unaffected by the change.

Dijkstra algorithm—Routing algorithm that iterates on length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms. Sometimes called *shortest path first algorithm*.

distance vector routing protocol—Class of routing algorithms that iterate on the number of hops in a route to find a shortest-path spanning tree. Distance vector routing algorithms call for each router to send its entire routing table in each update, but only to its neighbors. Distance vector routing algorithms can be prone to routing loops but are computationally simpler than link-state routing algorithms. These routing protocols also use the Bellman-Ford routing algorithm.

dynamic routes—Automatic rerouting of traffic based on sensing and analyzing current actual network conditions, not including cases of routing decisions taken on predefined information.

exterior routing—A routing protocol used to exchange information between autonomous systems or organizations, used to connect organizations into the Internet. BGP and EGP are examples of exterior routing protocols.

fast switching—A cache in the Cisco router that contains routing decisions. After the routing decision for a packet has been made, it can be cached in any one of a variety of caches. This means that the forwarding of traffic through the router is greatly enhanced.

feasible successor—A term used by EIGRP to describe a next-hop router that has a path to the remote network that EIGRP considers a viable route.

flash update—A routing update sent asynchronously in response to a change in the network topology. If there is a change in the metric, the update is sent immediately without waiting for the update timer to expire. Sometimes known as triggered updates.

floating static route—A route that has been manually configured. Manually configured routes will be chosen as the routing path first. A floating static route is a route that, although manually configured, has been identified as a route to choose only if the dynamically learned routes fail. These routes need to have a higher administrative distance than the routing protocol that you are using.

flooding—A traffic-passing technique used by switches and bridges in which traffic received on an interface is sent out to all the interfaces of that device, except the interface on which the information was originally received.

incremental update—A routing update that is sent only when there is a change in the topology, not periodically when a timer expires.

interior routing protocol—A routing protocol used to route information between routers within an autonomous system or organization.

link-state advertisement (LSA)—Broadcast packet used by link-state protocols that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables.

link-state routing algorithm—A routing algorithm in which each router broadcasts or multicasts information regarding the cost of reaching each of its neighbors to all nodes in the internetwork. Compare with *distance vector routing protocol*.

neighbor—In OSPF or EIGRP, two routers that have interfaces to a common network.

poison reverse—Routing updates that specifically indicate that a network or subnet is unreachable, rather than implying that a network is unreachable by not including it in updates.

redistribution—Allowing routing information discovered through one routing protocol to be distributed in the update messages of another routing protocol. This is sometimes called route redistribution.

routed protocol—Protocol that can be routed by a router. A router must be capable of interpreting the logical internetwork as specified by that routed protocol. Examples of routed protocols include AppleTalk, DECnet, and IP.

routing function—Process of finding a path to a destination host. Routing is very complex in large networks because of the many potential intermediate destinations that a packet might traverse before reaching its destination host.

routing loop—A loop in which the routing information is fed back to the originating router as if from another router. This often happens when redistribution is configured. It can lead to confusion in the network because when the originating router loses the route, it may well believe that there is an alternative path.

routing protocol—Protocol that accomplishes routing through the implementation of a specific routing algorithm. Examples of routing protocols include IGRP, OSPF, and RIP.

routing table—Table stored in a router or some other internetworking device that keeps track of routes to particular network destinations and, in some cases, metrics associated with those routes.

silicon switching—Switching based on the Silicon Switch Engine (SSE), which allows the processing of packets independent of the Silicon Switch Processor (SSP) system processor. Silicon switching provides high-speed, dedicated packet switching.

split horizon rules—Routing technique in which information about routes is prevented from exiting the router interface through which that information was received. Split-horizon updates are useful in preventing routing loops.

static route—A route that is explicitly configured and entered into the routing table.

switching function—Forwarding packets from an inbound interface to an outbound interface.

topology table—Used by EIGRP and OSPF, the table that records all the routes in the network before determining which will be entered into the routing table.

triggered update—See *flash update*.

Q&A

The following questions test your understanding of the topics covered in this chapter. The final questions in this section repeat of the opening "Do I Know This Already?" questions. These are repeated to enable you to test your progress. After you have answered the questions, find the answers in Appendix A. If you get an answer wrong, review the answer and ensure that you understand the reason for your mistake. If you are confused by the answer, refer to the appropriate text in the chapter to review the concepts.

1 Name one routing protocol that sends periodic updates.

_____

_____

_____

2 What is an incremental update, and how often is it sent out?

_____

_____

_____

3 What is the routing algorithm used in OSPF?

_____

_____

_____

4 State one method by which a link-state routing protocol attempts to reduce the network overhead.

_____

_____

_____

5 Distance vector routing protocols naturally summarize at which boundary?

_____

_____

_____

6 Which routing protocol technology uses the Bellman Ford algorithm?

7 Give three reasons why RIPv1 has problems with working in a large network.

8 What is the Dijkstra algorithm used for?

9 What is the destination address of the distance vector periodic update?

10 State one major difference between a classful and classless routing protocol.

11 In the routing table, a field indicates the source of the routing information. If the field showed the letter C, what would this mean?

12 In the routing table, how is the next logical hop indicated?

13 Cisco distinguishes between the routing and the switching function—what is the difference?

14 State the two ways that an outgoing interface is selected as the preferred path.

15 What is administrative distance?

16 If IGRP has three paths to a remote network in which each path has an equal metric, what will happen?

17 Name the interior IP routing protocols that send the mask with the routing update.

18 Name the interior routing protocol that sends a routing update on a Cisco router every 30 seconds by default.

19 Does VLSM require a classful or classless routing protocol, and why?

20 State one of the characteristics of a classful routing protocol.

21 A distance vector routing protocol uses the mechanism of poison reverse—what is this?

22 Name two distance vector routing protocols.

23 Name two link-state IP routing protocols.

24 Describe the mechanism of split horizon.

25 What is the command syntax to empty the Cisco routing table of all its routes?

26 What does 0.0.0.0 signify in an IP routing table?

27 What is the command to show whether a specific network, such as 141.131.6.16, is present in the routing table?

28 What is the next logical hop in the routing table?

Scenarios

The following scenarios and questions are designed to draw together the content of the chapter and to exercise your understanding of the concepts. There is not necessarily a right answer. The thought process and practice in manipulating the concepts is the goal of this section. The answers to the scenario questions are found at the end of this chapter.

Scenario 4-1

Mental Merge, our poor congested company, with many ideas but no bandwidth with which to develop or communicate them, is badly in need of a routed solution. Taking the scenario solution provided for the addressing scheme, outlined in the previous chapter, it is now necessary to implement routing.

Using the network in Figure 4-3, answer the following questions.

1. Using the addressing skills developed in Chapter 3, "IP Addressing," and in reference to Figure 4-3, state where the routers should be placed.
2. The administrator has decided that a link-state routing protocol is the best solution for this network design. Justify his choice, explaining which characteristics of the link-state routing protocol would benefit this network.
3. The administrator must create an implementation plan for his team. List the IP routing protocol requirements for every router that may be used as a checklist for the installation staff.
4. The links between the various sites are leased lines with a backup link using a dialup line. Should the administrator be aware of any considerations?

Figure 4-3 *The Mental Merge Company Network for Scenario 4-1*

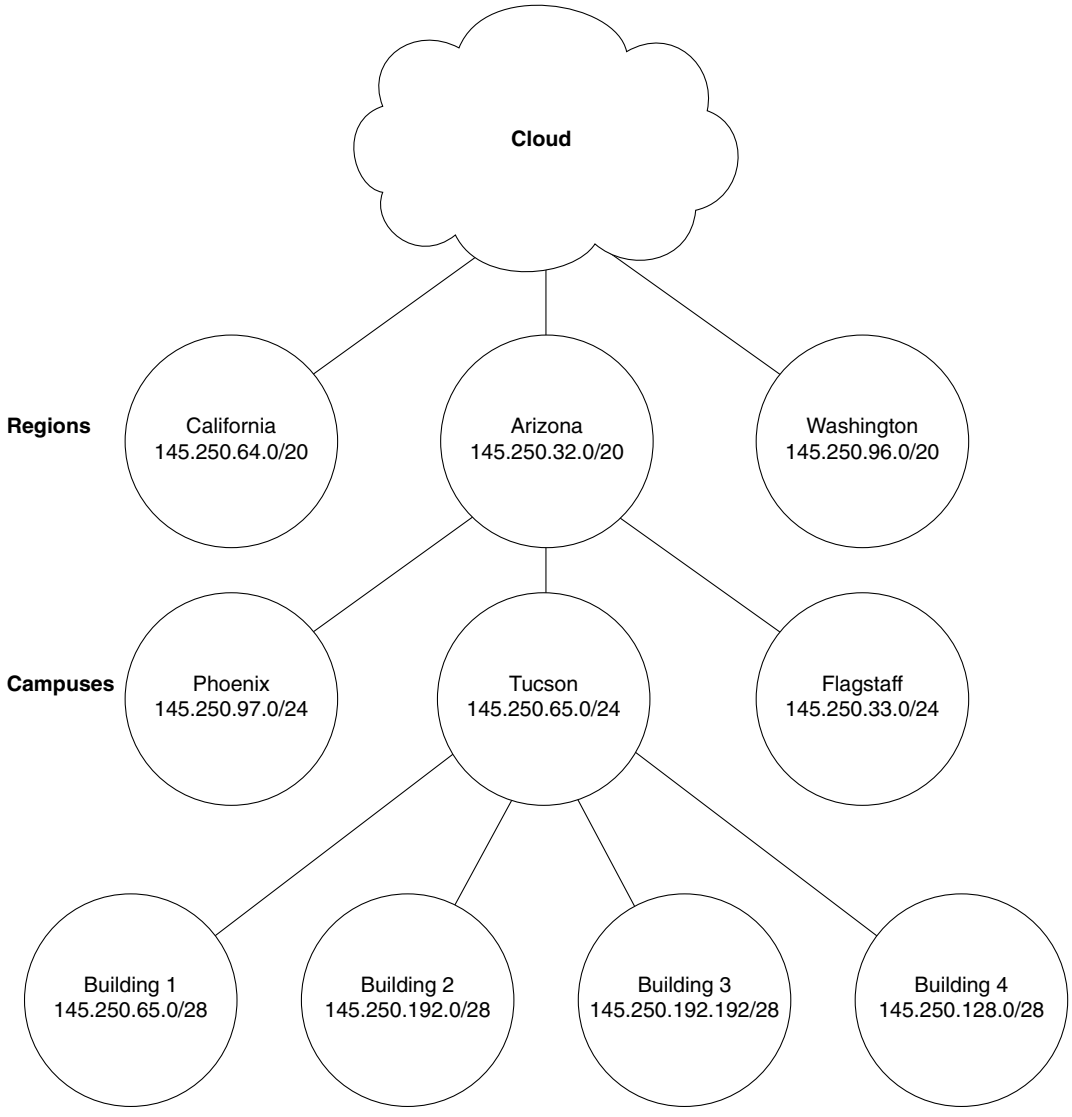

Scenario 4-2

Scenario 4-2 provides you with the data produced with the Cisco router IOS command **show ip route**. A legend defining the fields of the sample output is also provided to assist in answering the questions for Scenario 4-2.

Example 4-2 contains sample output from the **show ip route** command.

Example 4-2 **show ip route** *Command Output*

```
Codes: I - IGRP derived, R - RIP derived, H - Hello derived, O - OSPF derived
       C - connected, S - static, E - EGP derived, B - BGP derived
       * - candidate default route, IA - OSPF inter area route
       E1 - OSPF external type 1 route, E2 - OSPF external type 2 route

Gateway of last resort is 131.119.254.240 to network 129.140.0.0

O E2 150.150.0.0   [160/5]     via 131.119.254.6,     0:01:00, Ethernet2
E       192.67.131.0 [200/128] via 131.119.254.244, 0:02:22, Ethernet2
O E2 192.68.132.0 [160/5]     via 131.119.254.6,     0:00:59, Ethernet2
O E2 130.130.0.0   [160/5]     via 131.119.254.6,     0:00:59, Ethernet2
E       128.128.0.0   [200/128] via 131.119.254.244, 0:02:22, Ethernet2
E       129.129.0.0   [200/129] via 131.119.254.240, 0:02:22, Ethernet2
E       192.65.129.0 [200/128] via 131.119.254.244, 0:02:22, Ethernet2
E       131.131.0.0   [200/128] via 131.119.254.244, 0:02:22, Ethernet2
E       192.75.139.0 [200/129] via 131.119.254.240, 0:02:23, Ethernet2
```

The following information defines the fields reported in the **show ip route** command:

- The first column lists the protocol that derived the route.
- The second column may list certain protocol-specific information as defined in the display header.
- The third column lists the address of the remote network. The first number in the brackets is the administrative distance of the information source; the second number is the metric for the route.
- The fourth column specifies the address of the router that can build a route to the specified remote network.
- The fifth column specifies the last time that the route was updated, in hours:minutes:seconds.
- The final column specifies the interface through which the specified network can be reached.

Answer the following questions by using the output from the preceding **show ip route** command.

1. What routing protocol derived the route 130.130.0.0?
2. What router interface IP address is used to reach IP network 192.67.131.0?
3. When was the last time that the route 192.65.129.0 was updated?
4. Through which router interface can the IP network 128.128.0.0 be reached?

Scenario Answers

The answers are in **bold**. The answers provided in this section are not necessarily the only possible answers to the questions. The questions are designed to test your knowledge and to give practical exercise in certain key areas. This section is intended to test and exercise skills and concepts detailed in the body of this chapter.

If your answer is different, ask yourself whether it follows the tenets explained in the answers provided. Your answer is correct not if it matches the solution provided in the book, but rather if it has included the principles of design laid out in the chapter.

In this way, the testing provided in these scenarios is deeper: It examines not only your knowledge, but also your understanding and ability to apply that knowledge to problems.

If you do not get the correct answer, refer back to the text and review the subject tested. Be certain to also review your notes on the question to ensure that you understand the principles of the subject.

Scenario 4-1 Answers

1 Using the addressing skills developed in Chapter 3, and in reference to Figure 4-3, state where the routers should be placed.

The routers should be placed in each location, with the option of adding routers within each building if the network grows considerably.

2 The administrator has decided that a link-state routing protocol is the best solution for this network design. Justify his choice, explaining which characteristics of the link-state routing protocol would benefit this network.

A link-state routing protocol would be a good choice because of the large number of WAN interfaces. A distance vector routing protocol would increase congestion across these low-bandwidth links. The capability to use VLSM and to summarize these points would be an added advantage.

3 The administrator must create an implementation plan for his team. List the IP routing protocol requirements for every router that may be used as a checklist for the installation staff.

Each person implementing the routing protocol on the router would have to ensure the following:

- **That the appropriate interfaces had IP addresses that were on the same subnet as the other devices on the segment.**

- **That the routing protocol was configured correctly with the correct network addresses.**
- **That the routing table reflected the logical topology map of the network and that all the remote networks were present.**
- **If there were multiple paths available of equal cost, that the routing protocol was load sharing between the paths. This would mean all the paths were present in the routing table.**

4 The links between the various sites are leased lines with a backup link using a dialup line. Should the administrator be aware of any considerations?

The leased lines to the remote sites could be configured to be the primary link; as such, no traffic would traverse the dialup links. However, routing updates would be propagated out the dialup links so that the routing table would be aware of the potential path. To prevent this (and, thus, the dialup line being raised), the path could be manually entered into the routing table. However, this would render it the preferred path. Configuring the dialup paths as floating static routes would ensure that they were used only if the primary line failed, without having to generate network traffic across the link to maintain the routing table.

Scenario 4-2 Answers

1 What routing protocol derived the route 130.130.0.0?

OSPF.

2 What router interface IP address is used to reach IP network 192.67.131.0?

131.119.254.244. The fourth column of the sample output specifies the address of the router that can build a route to the specified remote network.

3 When was the last time that the route 192.65.129.0 was updated?

0:02:22. The fifth column of the sample output specifies the last time the route was updated, in hours:minutes:seconds.

4 Through which router interface can the IP network 128.128.0.0 be reached?

Ethernet2. The last column in the sample output specifies the interface through which the specified network can be reached.