CHAPTER 5

RIP Version 1

Objectives

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

- What are the functions, characteristics, and operation of the RIPv1 protocol?
- Can you configure a device for RIPv1?
- Can you verify proper RIPv1 operations?
- How does RIPv1 perform automatic summarization?
- Can you configure, verify, and troubleshoot default routes propagated in a routed network implanting RIPv1?
- What are the recommended techniques to solve problems related to RIPv1?

Key Terms

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

XNS	*page 233*
automatic summarization	*page 251*
boundary router	*page 254*
discontiguous network	*page 260*

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Over the years, routing protocols have evolved to meet the increasing demands of complex networks. The first protocol used was Routing Information Protocol (RIP). RIP still enjoys popularity because of its simplicity and widespread support.

Figure 5-1 shows a chart of the most common IP routing protocols, with the routing protocols that will be discussed in this book highlighted. Notice in the chart that RIP (RIP version 1) is a classful, distance vector routing protocol.

Figure 5-1 Chart of Routing Protocols

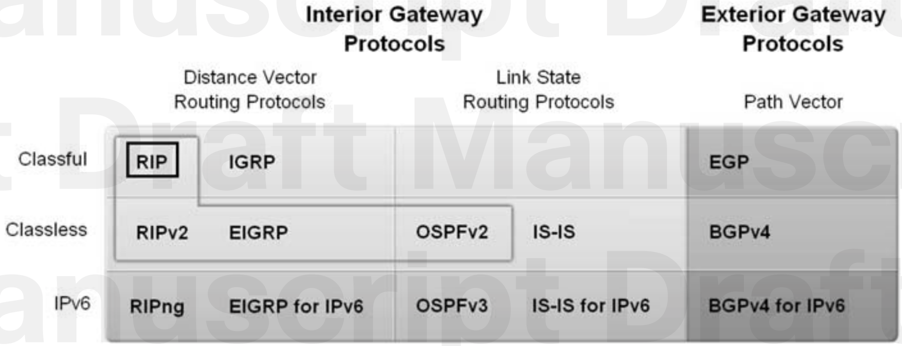

	Interior Gateway Protocols				Exterior Gateway Protocols
	Distance Vector Routing Protocols		Link State Routing Protocols		Path Vector
Classful	RIP	IGRP			EGP
Classless	RIPv2	EIGRP	OSPFv2	IS-IS	BGPv4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGPv4 for IPv6

Understanding RIP is important to your networking studies for two reasons:

- RIP is still in use today. You might encounter a network implementation that is large enough to need a routing protocol, yet simple enough to use RIP effectively.
- Familiarity with many of the fundamental concepts of RIP will help you to compare RIP with other protocols. Understanding how RIP operates and knowing its implementation will make learning other routing protocols easier.

This chapter covers the details of RIP version 1, including a bit of history, RIPv1 characteristics, operation, configuration, verification, and troubleshooting. Throughout the chapter, you can use Packet Tracer Activities to practice what you learn. *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* and the online curriculum provide three hands-on labs and a Packet Tracer Skills Integration Challenge Activity to help you integrate RIPv1 into your growing set of networking knowledge and skills.

RIPv1: Distance Vector, Classful Routing Protocol

RIPv1 is a distance vector routing protocol for IPv4. RIPv1 is also a classful routing protocol. In this chapter, we will begin to examine the limitations of a classful routing protocol. Chapter 6, "VLSM and CIDR," and Chapter 7, "RIPv2," will discuss classless routing protocols and compare them to classful routing protocols.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Draft Manuscript

Background and Perspective

RIP is the oldest of the distance vector routing protocols. Although RIP lacks the sophistication of more advanced routing protocols, its simplicity and continued widespread use is a testament to its longevity. RIP is not a protocol “on the way out.” In fact, an IPv6 form of RIP called RIPng (next generation) is now available.

Figure 5-2 compares RIP and other network protocol developments over time.

Figure 5-2 Overview of RIP Historical Impact

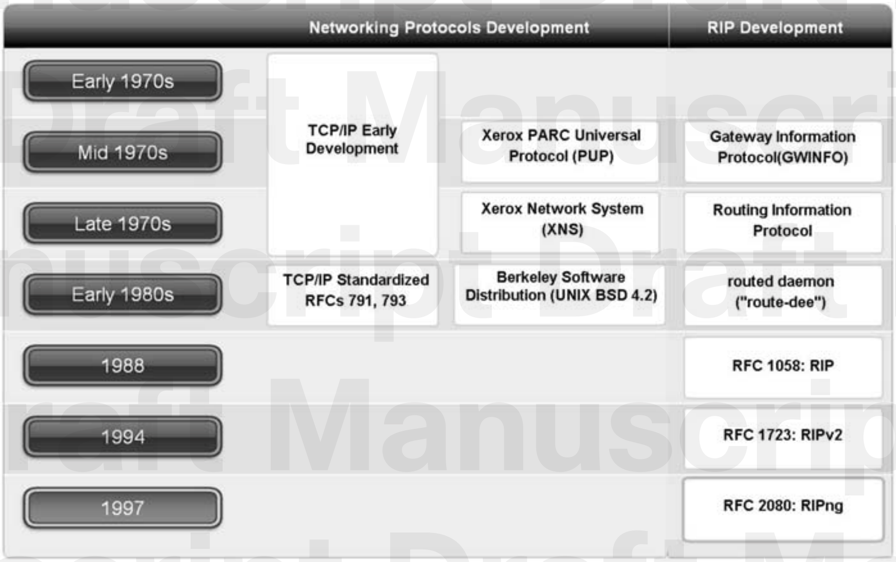

	Networking Protocols Development		RIP Development
Early 1970s	TCP/IP Early Development		
Mid 1970s		Xerox PARC Universal Protocol (PUP)	Gateway Information Protocol(GWINFO)
Late 1970s		Xerox Network System (XNS)	Routing Information Protocol
Early 1980s	TCP/IP Standardized RFCs 791, 793	Berkeley Software Distribution (UNIX BSD 4.2)	routed daemon ("route-dee")
1988			RFC 1058: RIP
1994			RFC 1723: RIPv2
1997			RFC 2080: RIPng

RIP evolved from an earlier protocol developed at Xerox, called Gateway Information Protocol (GWINFO). With the development of Xerox Network System (***XNS***), GWINFO evolved into RIP. It later gained popularity because it was implemented in the Berkeley Software Distribution (BSD) as a daemon named routed (pronounced *route-dee*, not *rout-ed*). Various other vendors made their own, slightly different implementations of RIP. Recognizing the need for standardization of the protocol, Charles Hedrick wrote RFC 1058 in 1988, in which he documented the existing protocol and specified some improvements. RFC 1058 can be found at http://www.ietf.org/rfc/rfc1058.txt. Since then, RIP has been improved with RIPv2 in 1994 and with RIPng in 1997.

Note

The first version of RIP is often called RIPv1 to distinguish it from RIP version 2 (RIPv2). However, both versions share many of the same features. When discussing features common to both versions, we will refer to RIP. When discussing features unique to each version, we will use RIPv1 and RIPv2. RIPv2 is discussed Chapter 7.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

RIPv1 Characteristics and Message Format

RIPv1 is a routing protocol and, like other protocols, has a format with fields containing specific information. For example, the IP protocol has fields containing information such as source IP address and destination IP address. Routing protocols also have fields containing information. One of the fields in the RIPv1 routing protocol is the IP Address field, which contains an IP network address. Using the information in these fields is how routers share routing information. Examining some of these fields can help take some of the mystery out of the protocol and its operations.

RIP Characteristics

As discussed in Chapter 4, "Distance Vector Routing Protocols," RIP has the following key characteristics:

- RIP is a distance vector routing protocol.
- RIP uses hop count as its only metric for path selection.
- Advertised routes with hop counts greater than 15 are considered unreachable.
- Response messages (routing table updates) are broadcast every 30 seconds.

Figure 5-3 shows an encapsulated RIPv1 message.

Figure 5-3 Encapsulated RIPv1 Message

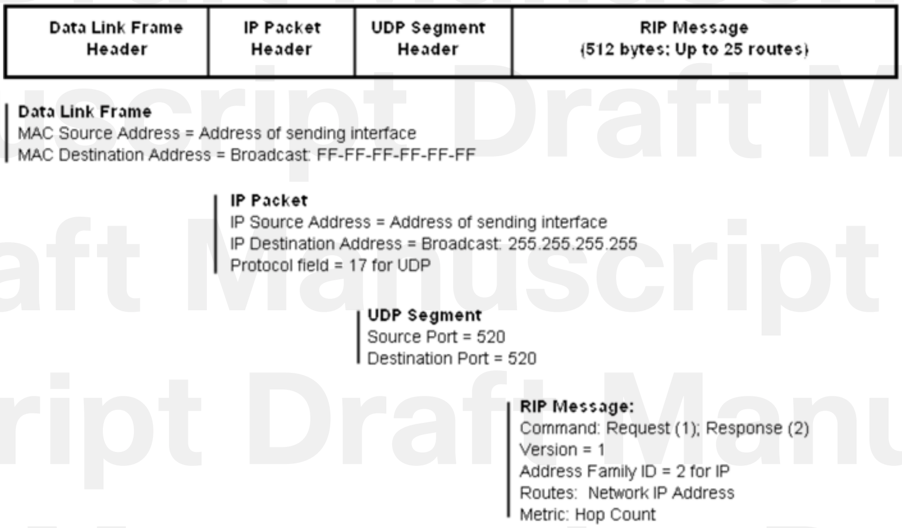

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

The data portion of a RIP message is encapsulated into a User Datagram Protocol (UDP) segment, with both source and destination port numbers set to 520. The IP header and data-link headers add broadcast destination addresses before the message is sent out all RIP-configured interfaces.

RIP Message Format: RIP Header

Figure 5-4 shows the detail of a RIPv1 message. Table 5-1 lists and describes the main fields of the message.

Figure 5-4 RIPv1 Message Format

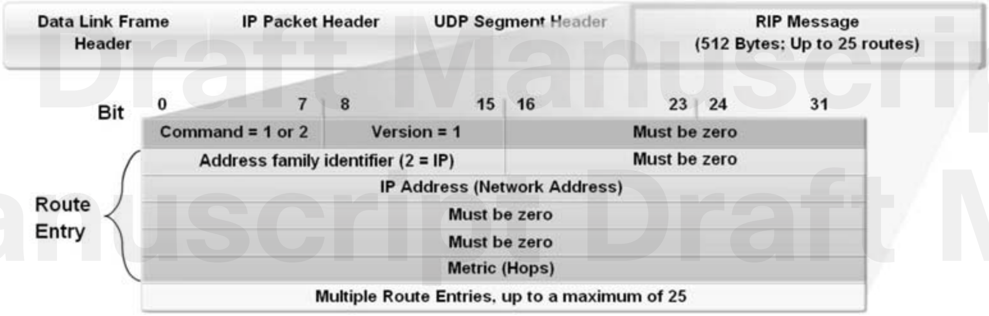

Table 5-1 RIPv1 Message Field Descriptions

Field	Description
Command	1 for a Request or 2 for a Response.
Version	1 for RIPv1 or 2 for RIPv2.
Address Family Identifier	2 for IP unless a Request is for the full routing table, in which case the field is set to 0.
IP Address	The address of the destination route, which can be a network, subnet, or host address.
Metric	Hop count between 1 and 16. The sending router increases the metric before sending out the message.

Three fields are specified in the 4-byte header portion shown in blue in the figure. The Command field specifies the message type, discussed in more detail in the next section. The Version field is set to 1 for RIP version 1. The third field is labeled “Must be zero.” “Must be zero” fields provide room for future expansion of the protocol.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

RIP Message Format: Route Entry

The route entry portion of the message includes three fields with content:

- Address Family Identifier (set to 2 for IP unless a router is requesting a full routing table, in which case the field is set to 0)
- IP Address
- Metric

This route entry portion represents one destination route with its associated metric. One RIP update can contain up to 25 route entries. The maximum datagram size is 512 bytes, not including the IP or UDP headers.

Why Are So Many Fields Set to Zero?

RIP was developed before IP and was used for other network protocols (like XNS). BSD also had its influence. Initially, the extra space was added with the intention of supporting larger address spaces in the future. As you will see in Chapter 7, RIPv2 has now used most of these empty fields.

RIP Operation

The following sections introduce the basic operations of RIPv1. Later sections will discuss these operations in more detail.

RIP Request/Response Process

RIP uses two message types specified in the Command field: Request message and Response message.

Figure 5-5 shows the RIPv1 request/response process.

Each RIP-configured interface sends out a Request message on startup, requesting that all RIP neighbors send their complete routing tables. A Response message is sent back by RIP-enabled neighbors. When the requesting router receives the responses, it evaluates each route entry. If a route entry is new, the receiving router installs the route in the routing table. If the route is already in the table, the existing entry is replaced if the new entry has a better hop count. The startup router then sends a triggered update out all RIP-enabled interfaces containing its own routing table so that RIP neighbors can be informed of any new routes.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Figure 5-5 RIP Request/Response Example

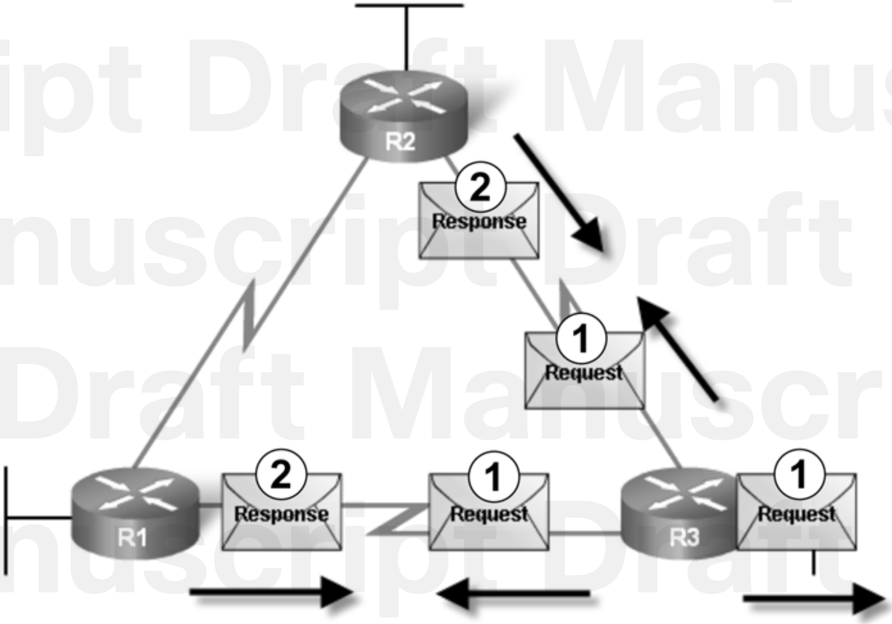

IP Address Classes and Classful Routing

You might recall from previous studies that IP addresses assigned to hosts were initially divided into three classes: Class A, Class B, and Class C. Each class was assigned a default subnet mask, as shown in Figure 5-6. Knowing the default subnet mask for each class is important to understanding how RIP operates.

RIP is a classful routing protocol. As you might have realized from the previous message format discussion, RIPv1 does not send subnet mask information in the update. Therefore, a router either uses the subnet mask configured on a local interface or applies the default subnet mask based on the address class. Because of this limitation, RIPv1 networks cannot be discontiguous, nor can they implement VLSM.

IP addressing is discussed further in Chapter 6.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Figure 5-6 Default Subnet Masks for Address Classes

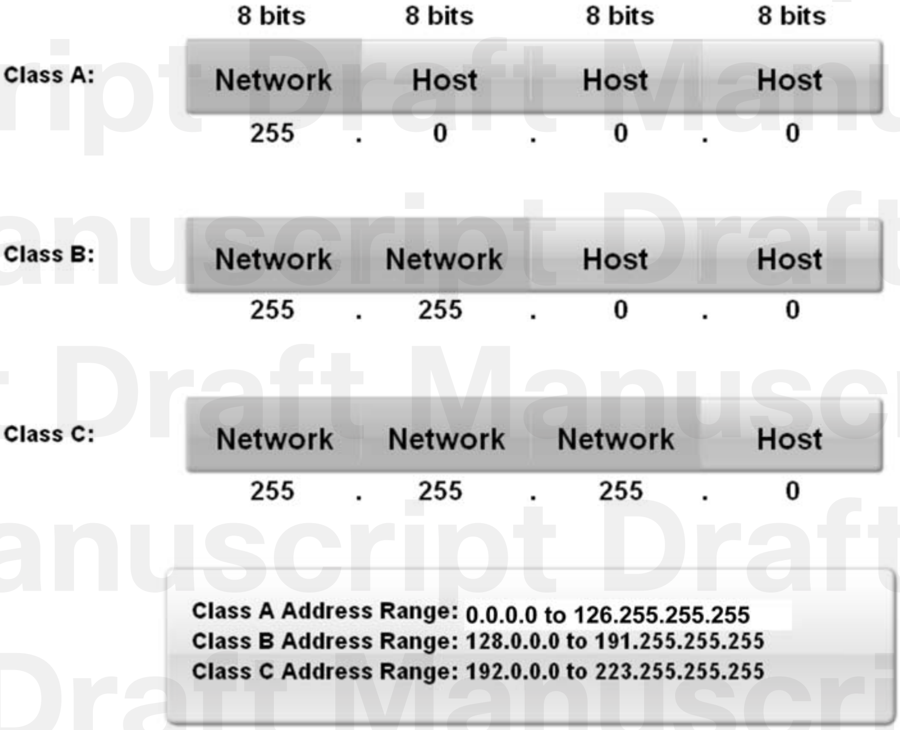

Administrative Distance

As discussed in Chapter 3, "Introduction to Dynamic Routing Protocols," administrative distance (AD) is the trustworthiness (or preference) of the route source. RIP has a default administrative distance of 120. When compared to other interior gateway protocols, RIP is the least-preferred routing protocol. Intermediate System–to–Intermediate System (IS-IS), Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP), and Enhanced IGRP (EIGRP) all have lower default AD values.

Remember, you can check the administrative distance using the **show ip route** command (see Example 5-1) or **show ip protocols** command (see Example 5-2).

Example 5-1 AD Values in the **show ip route** Command

```
R3# 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
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

```
       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

R    192.168.1.0/24 [120/1] via 192.168.6.2, 00:00:05, Serial0/0/0
R    192.168.2.0/24 [120/1] via 192.168.6.2, 00:00:05, Serial0/0/0
                    [120/1] via 192.168.4.2, 00:00:05, Serial0/0/1
R    192.168.3.0/24 [120/1] via 192.168.4.2, 00:00:05, Serial0/0/1
C    192.168.4.0/24 is directly connected, Serial0/0/1
C    192.168.5.0/24 is directly connected, FastEthernet0/0
C    192.168.6.0/24 is directly connected, Serial0/0/0
```

Example 5-2 AD Values in the **show ip protocols** Command

```
R3# show ip protocols

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 22 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip
  Default version control: send version 1, receive any version
    Interface              Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0        1      1 2
    Serial0/0/0              1      1 2
    Serial0/0/1              1      1 2
  Automatic network summarization is in effect
  Routing for Networks:
    192.168.4.0
    192.168.5.0
    192.168.6.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.6.2          120      00:00:10
    192.168.4.2          120      00:00:18
  Distance: (default is 120)
```

Draft Manuscript

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Basic RIPv1 Configuration

The following sections introduce the first of three topologies that will be used in this chapter.

RIPv1 Scenario A

Figure 5-7 shows the three router topologies used in Chapter 2, "Static Routing." Physically, the topology is the same, except that you will not need PCs attached to the LANs. Logically, however, the addressing scheme is different; this topology uses five Class C network addresses.

Figure 5-7 RIP Topology: Scenario A

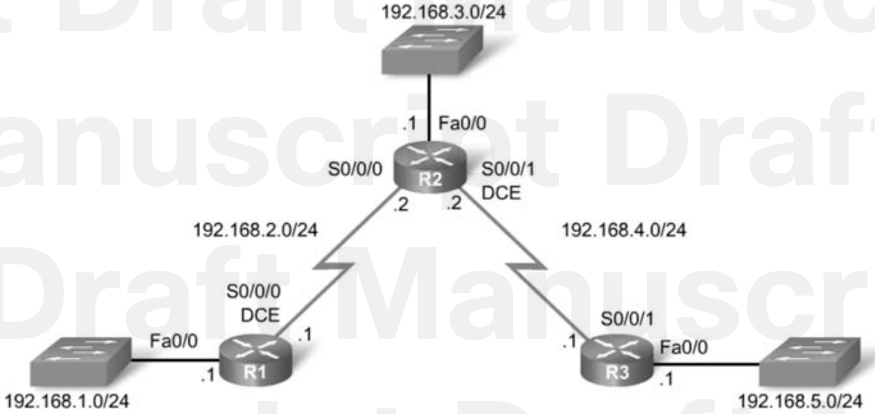

Table 5-2 displays the interface addressing for each router.

Table 5-2 Addressing Table: Scenario A

Device	Interface	IP Address	Subnet Mask
R1	Fa0/0	192.168.1.1	255.255.255.0
	S0/0/0	192.168.2.1	255.255.255.0
R2	Fa0/0	192.168.3.1	255.255.255.0
	S0/0/0	192.168.2.2	255.255.255.0
	S0/0/1	192.168.4.2	255.255.255.0
R3	Fa0/0	192.168.5.1	255.255.255.0
	S0/0/1	192.168.4.1	255.255.255.0

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Packet Tracer ☐ **Activity**

Basic RIPv1 Configuration (5.2.1)

Use the Packet Tracer Activity to configure and activate all the interfaces for the RIP Topology: Scenario A. Detailed instructions are provided within the activity. Use file e2-521.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Enabling RIP: router rip Command

To enable a dynamic routing protocol, enter global configuration mode and use the **router** command. As shown in Example 5-3, if you type a space followed by a question mark, a list of all the available routing protocols supported by IOS displays.

Example 5-3 RIP Router Configuration Mode

```
R1# conf t

Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router ?

  bgp       Border Gateway Protocol (BGP)
  egp       Exterior Gateway Protocol (EGP)
  eigrp     Enhanced Interior Gateway Routing Protocol (EIGRP)
  igrp      Interior Gateway Routing Protocol (IGRP)
  isis      ISO IS-IS
  iso-igrp  IGRP for OSI networks
  mobile    Mobile routes
  odr       On Demand stub Routes
  ospf      Open Shortest Path First (OSPF)
  rip       Routing Information Protocol (RIP)

R1(config)# router rip

R1(config-router)#
```

To enter router configuration mode for RIP, enter **router rip** at the global configuration prompt. Notice that the prompt changes from a global configuration prompt to the following:

```
R1(config-router)#
```

This command does not directly start the RIP process. Instead, it provides access to configure routing protocol settings. No routing updates are sent until additional commands are configured.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

If you need to remove the RIP routing process from a device, negate the command with **no router rip**. This command stops the RIP process and erases all existing RIP configuration commands.

Specifying Networks

By entering RIP router configuration mode, the router is enabled for RIP. But the router still needs to know which local interfaces it should use for communication with other routers, as well as which locally connected networks it should advertise to those routers. To enable RIP routing for a network, use the **network** command in router configuration mode and enter the classful network address for each directly connected network.

```
Router(config-router)# network directly-connected-classful-network-address
```

The **network** command performs the following functions:

- Enables RIP on all interfaces that belong to a specific network. Associated interfaces will now both send and receive RIP updates.
- Advertises the specified network in RIP routing updates sent to other routers every 30 seconds.

Note

If you enter a subnet address, IOS automatically converts it to a classful network address. For example, if you enter the command **network 192.168.1.32**, the router will convert it to **network 192.168.1.0**.

Example 5-4 shows the **network** command configured on all three routers for the directly connected networks. Notice that only classful networks were entered.

Example 5-4 Enabling RIP with the **network** Command

```
R1(config)# router rip
R1(config-router)# network 192.168.1.0
R1(config-router)# network 192.168.2.0
R2(config)# router rip
R2(config-router)# network 192.168.2.0
R2(config-router)# network 192.168.3.0
R2(config-router)# network 192.168.4.0
R3(config)# router rip
R3(config-router)# network 192.168.4.0
R3(config-router)# network 192.168.5.0
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

What happens if you enter a subnet address or interface IP address instead of the classful network address when using the **network** command for RIP configurations?

```
R3(config)# router rip
R3(config-router)# network 192.168.4.0
R3(config-router)# network 192.168.5.1
```

This example uses an interface IP address instead of the classful network address. Notice that IOS does not give an error message. Instead, IOS corrects the input and enters the classful network address, as you can see in the following output for verification.

```
R3# show running-config

!
router rip
 network 192.168.4.0
 network 192.168.5.0
!
```

Packet Tracer
☐ **Activity**

Specifying Networks (5.2.3)

Use the Packet Tracer Activity to practice configuring RIP routing on all three routers in the topology. Detailed instructions are provided within the activity. Use file e2-523.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Verification and Troubleshooting

It is important to be able to verify and troubleshoot your routing configuration. Verifying routing operations immediately after configuration will help solve any potential troubleshooting issues that might arise later.

To verify and troubleshoot routing, first use the **show ip route** and **show ip protocols** commands. If you cannot isolate the problem using these two commands, use the **debug ip rip** command to see exactly what is happening. These three commands are discussed in a suggested order that you might use to verify and troubleshoot a routing protocol configuration. Remember, before you configure any routing—whether static or dynamic—make sure that all necessary interfaces are up and up with the **show ip interface brief** command.

Verifying RIP: show ip route Command

Examining the routing table is an easy way to see whether the routing protocol and commands have been properly configured. Be sure to look for any routes that you expect to see in the routing table, along with any routes that should not be there.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Example 5-5 shows the routing tables for R1, R2, and R3 by using the **show ip route** command.

Example 5-5 Verifying RIP Convergence with the **show ip route** Command

```
R1# show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
<output omitted>

Gateway of last resort is not set

R    192.168.4.0/24 [120/1] via 192.168.2.2, 00:00:02, Serial0/0/0
R    192.168.5.0/24 [120/2] via 192.168.2.2, 00:00:02, Serial0/0/0
C    192.168.1.0/24 is directly connected, FastEthernet0/0
C    192.168.2.0/24 is directly connected, Serial0/0/0
R    192.168.3.0/24 [120/1] via 192.168.2.2, 00:00:02, Serial0/0/0
```

```
R2# show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
<output omitted>

Gateway of last resort is not set

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, 00:00:12, Serial0/0/1
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
```

```
R3# show ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
<output omitted>

Gateway of last resort is not set

C    192.168.4.0/24 is directly connected, Serial0/0/1
C    192.168.5.0/24 is directly connected, FastEthernet0/0
R    192.168.1.0/24 [120/2] via 192.168.4.2, 00:00:08, Serial0/0/1
R    192.168.2.0/24 [120/1] via 192.168.4.2, 00:00:08, Serial0/0/1
R    192.168.3.0/24 [120/1] via 192.168.4.2, 00:00:08, Serial0/0/1
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript Draft Manuscript Draft Manuscript

The **show ip route** command verifies that routes received by RIP neighbors are installed in a routing table. An **R** in the output indicates RIP routes. Because this command displays the entire routing table, including directly connected and static routes, it is normally the first command used to check for convergence. Routes might not immediately appear when you execute the command because networks take some time to converge. However, when routing is correctly configured on all routers, the **show ip route** command will reflect that each router has a full routing table, with a route to each network in the topology.

In the topology shown earlier in Figure 5-7, you can see there are five networks. Each router lists five networks in the routing table; therefore, you can say that all three routers are converged because each router has a route to every network shown in the topology.

To better understand the output from the **show ip route** command, focus on one RIP route learned by R1 and interpret the output shown in the routing table:

```
R    192.168.5.0/24 [120/2] via 192.168.2.2, 00:00:23, Serial0/0/0
```

The listing of routes with an R code is a quick way to verify that RIP is running on this router. If RIP is not at least partially configured, you will not see any RIP routes.

Next, the remote network address and subnet mask are listed (**192.168.5.0/24**).

The AD value (**120** for RIP) and the distance to the network (**2** hops) are shown in brackets.

The next-hop IP address of the advertising router is listed (R2 at **192.168.2.2**) as well as how many seconds have passed since the last update (**00:00:23**, in this case).

Finally, the exit interface that this router will use for traffic destined for the remote network is listed (**Serial 0/0/0**).

Table 5-3 lists the output and description of each part.

Table 5-3 Interpreting a Route

Output	Description
R	Identifies the source of the route as RIP.
192.168.5.0	Indicates the address of the remote network.
/24	Indicates the subnet mask used for this network.
[120/2]	Shows the administrative distance (120) and the metric (2 hops).
via 192.168.2.2,	Specifies the address of the next-hop router (R2) to send traffic to for the remote network.
00:00:23,	Specifies the amount of time since the route was updated (here, 23 seconds). Another update is due in 7 seconds.
Serial0/0/0	Specifies the local interface through which the remote network can be reached.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Verifying RIP: show ip protocols Command

Another useful command in verifying RIP or other routing protocols is the **show ip protocols** command. If a network is missing from the routing table, check the routing configuration using **show ip protocols**.

The **show ip protocols** command displays the routing protocol that is currently configured on the router. This output can be used to verify most RIP parameters to confirm that

- RIP routing is configured.
- The correct interfaces send and receive RIP updates.
- The router advertises the correct networks.
- RIP neighbors are sending updates.

This command is also useful when verifying the operations of other routing protocols, as you will see later with EIGRP and OSPF.

Figure 5-8 shows the output from the **show ip protocols** command, with numbers by each portion of the output. The descriptions that follow the figure correspond to the numbers in the figure.

Figure 5-8 Interpreting **show ip protocols** Output

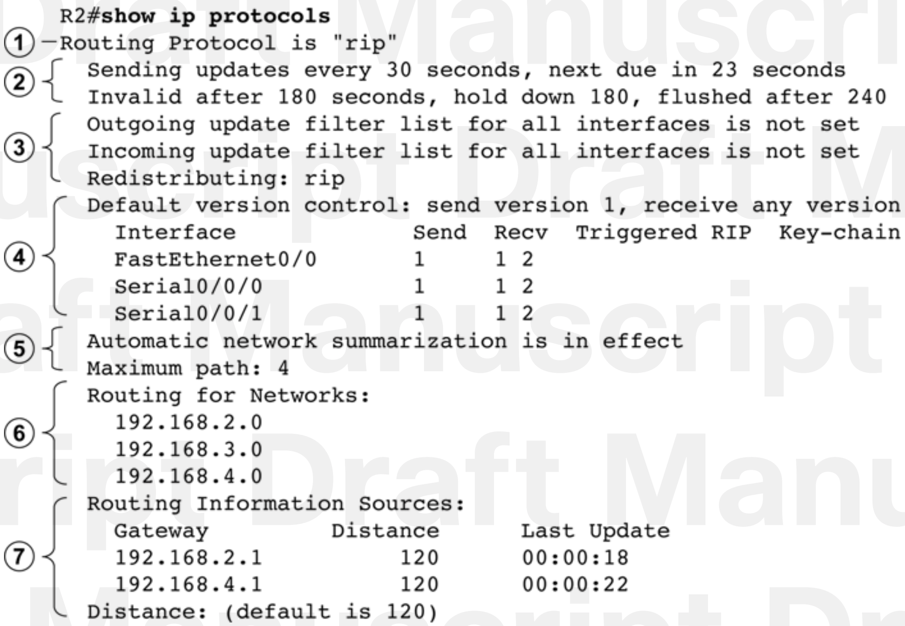

```
R2#show ip protocols
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: rip
  Default version control: send version 1, receive any version
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0       1     1 2
    Serial0/0/0           1     1 2
    Serial0/0/1           1     1 2
  Automatic network summarization is in effect
  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          120      00:00:18
    192.168.4.1          120      00:00:22
  Distance: (default is 120)
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript Draft Manuscript Draft Manuscript

1. The first line of output verifies that RIP routing is configured and running on router R2. As you saw in the section "Basic RIPv1 Configuration," earlier in this chapter, at least one active interface with an associated **network** command is needed before RIP routing will start.
2. These are the timers that show when the next round of updates will be sent out from this router—23 seconds from now, in the example.
3. This information relates to filtering updates and redistributing routes, if configured on this router. Filtering and redistribution are both CCNP-level topics.
4. This block of output contains information about which RIP version is currently configured and which interfaces are participating in RIP updates.
5. This part of the output shows that Router R2 is currently summarizing at the classful network boundary and, by default, will use up to four equal-cost routes to load-balance traffic. Automatic summarization is discussed later in this chapter.
6. The classful networks configured with the **network** command are listed next. These are the networks that R2 will include in its RIP updates.
7. Here, the RIP neighbors are listed as Routing Information Sources. Gateway is the next-hop IP address of the neighbor that is sending R2 updates. Distance is the AD that R2 uses for updates sent by this neighbor. Last Update is the seconds since the last update was received from this neighbor.

Verifying RIP: debug ip rip Command

The **debug** command is a useful tool to help diagnose and resolve networking problems, providing real-time, continuous information. Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems. Moreover, it is best to use **debug** commands during periods of lower network traffic and fewer users.

Most RIP configuration errors involve an incorrect **network** statement configuration, a missing **network** statement configuration, or the configuration of discontiguous subnets in a classful environment. As shown in Figure 5-9, an effective command used to find issues with RIP updates is **debug ip rip**.

This command displays RIP routing updates as they are sent and received. Because updates are periodic, you need to wait for the next round of updates before seeing any output.

The list that follows corresponds with the numbers in Figure 5-9.

1. First you see an update coming in from R1 on interface Serial 0/0/0. Notice that R1 only sends one route to the 192.168.1.0 network. No other routes are sent because doing so would violate the split horizon rule. R1 is not allowed to advertise networks back to R2 that R2 previously sent to R1.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Figure 5-9 Interpreting **debug ip rip** Output

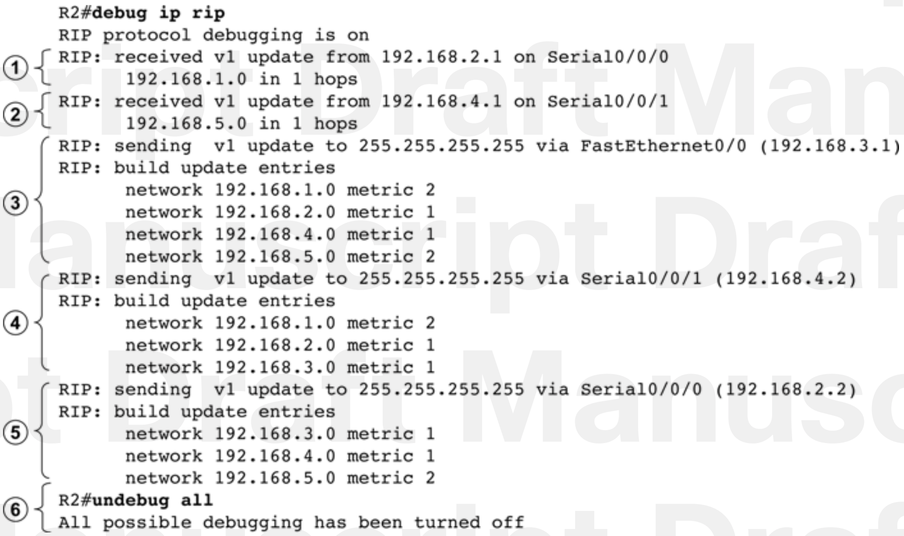

```
R2#debug ip rip
RIP protocol debugging is on
RIP: received v1 update from 192.168.2.1 on Serial0/0/0
      192.168.1.0 in 1 hops
RIP: received v1 update from 192.168.4.1 on Serial0/0/1
      192.168.5.0 in 1 hops
RIP: sending  v1 update to 255.255.255.255 via FastEthernet0/0 (192.168.3.1)
RIP: build update entries
      network 192.168.1.0 metric 2
      network 192.168.2.0 metric 1
      network 192.168.4.0 metric 1
      network 192.168.5.0 metric 2
RIP: sending  v1 update to 255.255.255.255 via Serial0/0/1 (192.168.4.2)
RIP: build update entries
      network 192.168.1.0 metric 2
      network 192.168.2.0 metric 1
      network 192.168.3.0 metric 1
RIP: sending  v1 update to 255.255.255.255 via Serial0/0/0 (192.168.2.2)
RIP: build update entries
      network 192.168.3.0 metric 1
      network 192.168.4.0 metric 1
      network 192.168.5.0 metric 2
R2#undebug all
All possible debugging has been turned off
```

2. The next update that is received is from R3. Again, because of the split horizon rule, R3 only sends one route: the 192.168.5.0 network.
3. R2 sends out its own updates. First, R2 builds an update to send out the FastEthernet 0/0 interface. The update includes the entire routing table except for network 192.168.3.0, which is attached to FastEthernet 0/0.
4. Next, R2 builds an update to send to R3. Three routes are included. R2 does not advertise the network R2 and R3 share, nor does it advertise the 192.168.5.0 network because of split horizon.
5. Finally, R2 builds an update to send to R1. Three routes are included. R2 does not advertise the network that R2 and R1 share, nor does it advertise the 192.168.1.0 network because of split horizon.

Note

If you waited another 30 seconds, you would see all the debug output shown in the figure repeat because RIP sends out periodic updates every 30 seconds.

6. To stop monitoring RIP updates on R2, enter the **no debug ip rip** command or simply **undebug all**, as shown in figure.

Reviewing this debug output, you can verify that RIP routing is fully operational on R2. But do you see a way to optimize RIP routing on R2? Does R2 need to send updates out FastEthernet 0/0? You will see in the next topic how to prevent unnecessary updates.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript Draft Manuscript

Passive Interfaces

Some routers can have interfaces that do not connect to another router; therefore, there is no reason to send routing updates out that interface. You can use the **passive-interface** command with RIP to configure an interface not to send those updates.

Unnecessary RIP Updates Impact Network

As you saw in the previous example, R2 is sending updates out FastEthernet 0/0 even though no RIP router exists on that LAN. R2 has no way of knowing this and, as a result, sends an update every 30 seconds. Sending out unneeded updates on a LAN impacts the network in three ways:

- Bandwidth is wasted transporting unnecessary updates. Because RIPv1 updates are broadcast, switches will forward the updates out all ports.
- All devices on the LAN must process the RIPv1 update up to the transport layers, where the receiving device will discard the update.
- Advertising updates on a broadcast network is a security risk. RIP updates can be intercepted with packet-sniffing software. Routing updates can be modified and sent back to the router, corrupting the routing table with false metrics that misdirect traffic.

Stopping Unnecessary RIP Updates

You might think you could stop the updates by removing the 192.168.3.0 network from the configuration using the **no network 192.168.3.0** command, but then R2 would not advertise this LAN as a route in updates sent to R1 and R3. The correct solution is to use the **passive-interface** command, which prevents the transmission of routing updates through a router interface but still allows that network to be advertised to other routers. Enter the **passive-interface** command in router configuration mode:

```
Router(config-router)# passive-interface interface-type interface-number
```

This command stops routing updates out the specified interface. However, the network that the specified interface belongs to will still be advertised in routing updates that are sent out other interfaces.

In Example 5-6, R2 is first configured with the **passive-interface** command to prevent routing updates on FastEthernet 0/0 because no RIP neighbors exist on the LAN. The **show ip protocols** command is then used to verify the passive interface.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Example 5-6 Disabling Updates with the **passive-interface** Command

```
R2(config)# router rip
R2(config-router)# passive-interface FastEthernet 0/0
R2(config-router)# end
R2# show ip protocols

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 14 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip
  Default version control: send version 1, receive any version
    Interface               Send  Recv  Triggered RIP  Key-chain
    Serial0/0/0              1      1 2
    Serial0/0/1              1      1 2
  Automatic network summarization is in effect
  Routing for Networks:
    192.168.2.0
    192.168.3.0
    192.168.4.0
  Passive Interface(s):
    FastEthernet0/0
  Routing Information Sources:
    Gateway          Distance      Last Update
    192.168.2.1           120      00:00:27
    192.168.4.1           120      00:00:23
Distance: (default is 120)
```

Notice that the interface is no longer listed under **Interface** but under a new section called **Passive Interface(s)**. Also notice that the network 192.168.3.0 is still listed under **Routing for Networks:**, which means that this network is still included as a route entry in RIP updates that are sent to R1 and R3.

All routing protocols support the **passive-interface** command. You will be expected to use the **passive-interface** command when appropriate as part of your normal routing configuration.

Packet Tracer
☐ **Activity**

Passive Interfaces (5.3.4)

Use the Packet Tracer Activity to verify RIP routing and stop RIP updates using the **passive-interface** command. Detailed instructions are provided within the activity. Use file e2-534.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Automatic Summarization

Fewer number of routes in a routing table means that the routing table process can more quickly locate the route needed to forward the packet. Summarizing several routes into a single route is known as *route summarization* or *route aggregation*. Some routing protocols, such as RIP, automatically summarize routes on certain routers. The following sections discuss how RIP performs this ***automatic summarization***.

Modified Topology: Scenario B

To aid the discussion of automatic summarization, refer to the RIP topology shown in Figure 5-10, Scenario B.

Figure 5-10 RIP Topology: Scenario B

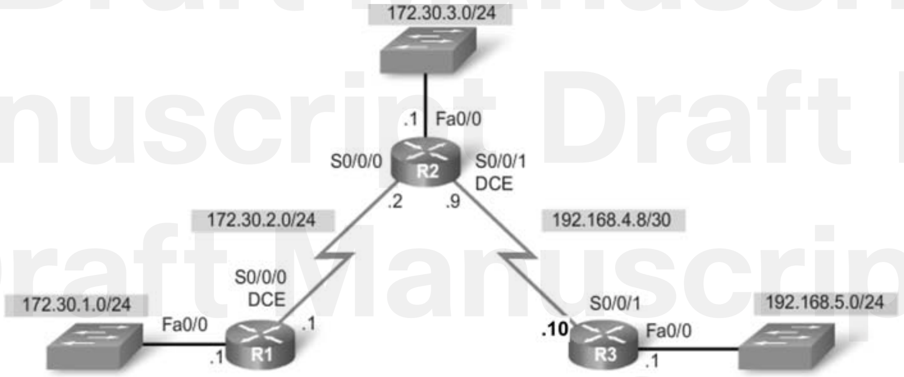

Scenario B is a modification of Scenario A with the following changes:

Three classful networks are used:

- 172.30.0.0/16
- 192.168.4.0/24
- 192.168.5.0/24

The 172.30.0.0/16 network is subnetted into three subnets:

- 172.30.1.0/24
- 172.30.2.0/24
- 172.30.3.0/24

The following devices are part of the 172.30.0.0/16 classful network address:

- All interfaces on R1
- S0/0/0 and Fa0/0 on R2

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

The 192.168.4.0/24 network is subnetted as a single subnet 192.168.4.8/30.

Table 5-4 shows the addressing scheme for this modified topology.

Table 5-4 Addressing Table: Scenario B

Device	Interface	IP Address	Subnet Mask
R1	Fa0/0	172.30.1.1	255.255.255.0
	S0/0/0	172.30.2.1	255.255.255.0
R2	Fa0/0	172.30.3.1	255.255.255.0
	S0/0/0	172.30.2.2	255.255.255.0
	S0/0/1	192.168.4.9	255.255.255.252
R3	Fa0/0	192.168.5.1	255.255.255.0
	S0/0/1	192.168.4.10	255.255.255.252

Examples 5-7, 5-8, and 5-9 show the configuration changes from Scenario A for Routers R1, R2, and R3, respectively.

Example 5-7 Configuration Changes for R1

```
R1(config)# interface fa0/0
R1(config-if)# ip address 172.30.1.1 255.255.255.0
R1(config-if)# interface S0/0/0
R1(config-if)# ip address 172.30.2.1 255.255.255.0
R1(config-if)# no router rip
R1(config)# router rip
R1(config-router)# network 172.30.1.0
R1(config-router)# network 172.30.2.0
R1(config-router)# passive-interface FastEthernet 0/0
R1(config-router)# end
R1# show run

<output omitted>
!
router rip
 passive-interface FastEthernet0/0
 network 172.30.0.0
!
<output omitted>
```

Draft Manuscript

Draft Manuscript Draft Manuscript

Example 5-8 Configuration Changes for R2

```
R2(config)# interface S0/0/0
R2(config-if)# ip address 172.30.2.2 255.255.255.0
R2(config-if)# interface fa0/0
R2(config-if)# ip address 172.30.3.1 255.255.255.0
R2(config-if)# interface S0/0/1
R2(config-if)# ip address 192.168.4.9 255.255.255.252
R2(config-if)# no router rip
R2(config)# router rip
R2(config-router)# network 172.30.0.0
R2(config-router)# network 192.168.4.8
R2(config-router)# passive-interface FastEthernet 0/0
R2(config-router)# end
R2# show run

<output omitted>
!
router rip
 passive-interface FastEthernet0/0
 network 172.30.0.0
 network 192.168.4.0
!
<output omitted>
```

Example 5-9 Configuration Changes for R3

```
R3(config)# interface fa0/0
R3(config-if)# ip address 192.168.5.1 255.255.255.0
R3(config-if)# interface S0/0/1
R3(config-if)# ip address 192.168.4.10 255.255.255.252
R3(config-if)# no router rip
R3(config)# router rip
R3(config-router)# network 192.168.4.0
R3(config-router)# network 192.168.5.0
R3(config-router)# passive-interface FastEthernet 0/0
R3(config-router)# end
R3# show run

<output omitted>
!
router rip
```

Draft Manuscript Draft Manuscript

```
 passive-interface FastEthernet0/0
 network 192.168.4.0
 network 192.168.5.0
!
<output omitted>
```

Notice that the **no shutdown** and **clock rate** commands are not needed because these commands are still configured from Scenario A. However, because new networks were added, the RIP routing process was removed with the **no router rip** command before enabling it again.

In the configuration for R1 (Example 5-7), notice that both subnets were configured with the **network** command. This configuration is technically incorrect because RIPv1 sends the classful network address in its updates and not the subnet. Therefore, IOS changed the configuration to reflect the correct, classful configuration, as you can see from the **show run** output.

In the configuration for R2 (Example 5-8), notice that the subnet 192.168.4.8 was configured with the **network** command. Again, this configuration is technically incorrect and IOS changed it to 192.168.4.0 in the running configuration.

The routing configuration for R3 is correct (Example 5-9). The running configuration matches what was entered in router configuration mode.

Note

On assessment and certification exams, entering a subnet address instead of the classful network address in a **network** command is considered an incorrect answer, even though Cisco IOS will make the correction.

Boundary Routers and Automatic Summarization

As you know, RIP is a classful routing protocol that automatically summarizes classful networks across major network boundaries. In Figure 5-11, you can see that R2 has interfaces in more than one major classful network. This makes R2 a ***boundary router*** in RIP. Serial 0/0/0 and FastEthernet 0/0 interfaces on R2 are both inside the 172.30.0.0 boundary. The Serial 0/0/1 interface is inside the 192.168.4.0 boundary.

Because boundary routers summarize RIP subnets from one major network to the other, updates for the 172.30.1.0, 172.30.2.0, and 172.30.3.0 networks will automatically be summarized into 172.30.0.0 when sent out R2's Serial 0/0/1 interface.

The next two sections examine how boundary routers perform this summarization.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Figure 5-11 RIP Boundary Router

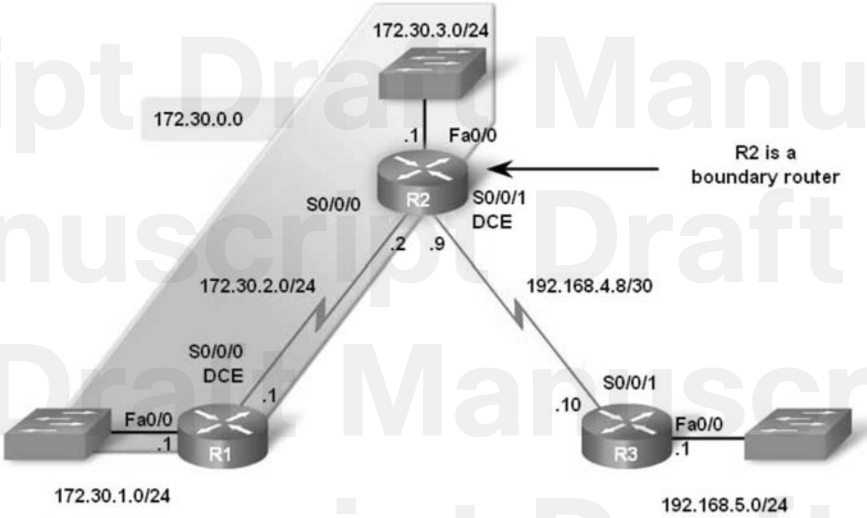

Processing RIP Updates

Classful routing protocols such as RIPv1 do not include the subnet mask in the routing update. However, the routing table includes RIPv1 routes with both the network address and the subnet mask. So how does a router running RIPv1 determine what subnet mask it should apply to a route when adding it to the routing table? The following section explains this process.

Rules for Processing RIPv1 Updates

The following two rules govern RIPv1 updates:

- If a routing update and the interface on which it is received belong to the same major network, the subnet mask of the interface is applied to the network in the routing update.
- If a routing update and the interface on which it is received belong to different major networks, the classful subnet mask of the network is applied to the network in the routing update.

Example of RIPv1 Processing Updates

In Example 5-10, R2 receives an update from R1 and enters the network in the routing table.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Example 5-10 Boundary Router Receiving RIP Updates

```
R2# debug ip rip

RIP protocol debugging is on
RIP: received v1 update from 172.30.2.1 on Serial0/0/0
      172.30.1.0 in 1 hops
<output omitted>
R2# undebug all

All possible debugging has been turned off
R2# show ip route

<output omitted>

Gateway of last resort is not set

     172.30.0.0/24 is subnetted, 3 subnets
R        172.30.1.0 [120/1] via 172.30.2.1, 00:00:18, Serial0/0/0
C        172.30.2.0 is directly connected, Serial0/0/0
C        172.30.3.0 is directly connected, FastEthernet0/0
     192.168.4.0/30 is subnetted, 1 subnets
C        192.168.4.8 is directly connected, Serial0/0/1
R    192.168.5.0/24 [120/1] via 192.168.4.10, 00:00:16, Serial0/0/1
```

How does R2 know that this subnet has a /24 (255.255.255.0) subnet mask? It knows because

- R2 received this information on an interface that belongs to the same classful network (172.30.0.0) as that of the incoming 172.30.1.0 update.
- The IP address for which R2 received the **172.30.1.0 in 1 hops** message was on Serial 0/0/0 with an IP address of 172.30.2.2 and a subnet mask of 255.255.255.0 (/24).
- R2 uses its own subnet mask on this interface and applies it to this and all other 172.30.0.0 subnets that it receives on this interface—in this case, 172.30.1.0.
- The 172.30.1.0 /24 subnet was added to the routing table.

Routers running RIPv1 are limited to using the same subnet mask for all subnets with the same classful network.

As you will learn in later chapters, classless routing protocols such as RIPv2 allow the same major (classful) network to use different subnet masks on different subnets, better known as variable-length subnet masking (VLSM).

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Sending RIP Updates: Using debug to View Automatic Summarization

To verify the network addresses sent and received by a router, you can use the **debug ip rip** command. Then by examining the routing tables, you can see the subnet mask that the receiving router applied to the RIPv1 routes.

Example 5-11 again shows the **debug ip rip** output for R2.

Example 5-11 Additional R2 Debug Output

```
R2# debug ip rip

RIP protocol debugging is on
RIP: sending  v1 update to 255.255.255.255 via Serial0/0/0 (172.30.2.2)
RIP: build update entries
       network 172.30.3.0 metric 1
       network 192.168.4.0 metric 1
       network 192.168.5.0 metric 2
RIP: sending  v1 update to 255.255.255.255 via Serial0/0/1 (192.168.4.9)
RIP: build update entries
       network 172.30.0.0 metric 1
R2# undebug all

All possible debugging has been turned off
```

When sending an update, boundary router R2 will include the network address and associated metric. If the route entry is for an update sent out a different major network, the network address in the route entry is summarized to the classful or major network address. This is exactly what R2 does for 192.168.4.0 and 192.168.5.0. It sends these classful networks to R1.

R2 also has routes for the 172.30.1.0/24, 172.30.2.0/24, and 172.30.3.0/24 subnets. In R2's routing update to R3 on Serial 0/0/1, R2 sends only a summary of the classful network address of 172.30.0.0.

If the route entry is for an update sent within a major network, the subnet mask of the outbound interface is used to determine the network address to advertise. R2 sends the 172.30.3.0 subnet to R1 using the subnet mask on Serial 0/0/0 to determine the subnet address to advertise.

R1 receives the 172.30.3.0 update on the Serial 0/0/0 interface, which has an interface address of 172.30.2.1/24. Because the routing update and interface both belong to the same major network, R1 applies its /24 mask to the 172.30.3.0 route.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Compare the routing tables for R1 and R3 in Examples 5-12 and 5-13.

Example 5-12 Routing Table for R1

```
R1# show ip route

<output omitted>

Gateway of last resort is not set

     172.30.0.0/24 is subnetted, 3 subnets
C       172.30.1.0 is directly connected, FastEthernet0/0
C       172.30.2.0 is directly connected, Serial0/0/0
R       172.30.3.0 [120/1] via 172.30.2.2, 00:00:17, Serial0/0/0
R    192.168.4.0/24 [120/1] via 172.30.2.2, 00:00:17, Serial0/0/0
R    192.168.5.0/24 [120/2] via 172.30.2.2, 00:00:17, Serial0/0/0
```

Example 5-13 Routing Table for R3

```
R3# show ip route

<output omitted>

Gateway of last resort is not set

R    172.30.0.0/16 [120/1] via 192.168.4.9, 00:00:15, Serial0/0/1
     192.168.4.0/30 is subnetted, 1 subnets
C       192.168.4.8 is directly connected, Serial0/0/1
C    192.168.5.0/24 is directly connected, FastEthernet0/0
```

Notice that R1 has three routes for the 172.30.0.0 major network, which has been subnetted to /24 or 255.255.255.0. R3 has only one route to the 172.30.0.0 network, and the network has not been subnetted. R3 has the major network in its routing table. However, it would be a mistake to assume that R3 does not have full connectivity. R3 will send any packets destined for the 172.30.1.0/24, 172.30.2.0/24, and 172.30.3.0/24 networks to R2 because all three of those networks belong to 172.30.0.0/16 and are reachable through R2.

Advantages and Disadvantages of Automatic Summarization

Automatic summarization has both advantages and disadvantages. Classful routing protocols such as RIPv1 do not allow you to modify this behavior. However, classless routing

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

protocols such as RIPv2 do permit automatic summarization to be disabled. For now, we examine the advantages and disadvantages of automatic summarization.

Advantages of Automatic Summarization

As you saw with R2 in Example 5-11, RIP automatically summarizes updates between classful networks. Because the 172.30.0.0 update is sent out an interface (Serial 0/0/1) on a different classful network (192.168.4.0), RIP sends out only a single update for the entire classful network instead of one for each of the different subnets. This process is similar to what you did when you summarized several static routes into a single static route. Why is automatic summarization an advantage? It is for the following reasons:

- Smaller routing updates are sent and received, which uses less bandwidth for routing updates between R2 and R3.
- R3 has a single route for the 172.30.0.0/16 network, regardless of how many subnets there are or how it is subnetted. Using a single route results in a faster lookup process in the routing table for R3.

Is there a disadvantage to automatic summarization? Yes, when discontiguous networks are configured in the topology.

Disadvantage of Automatic Summarization

Figure 5-12 shows a topology with discontiguous networks. Discontiguous networks will be explained shortly.

Figure 5-12 Discontiguous Topology

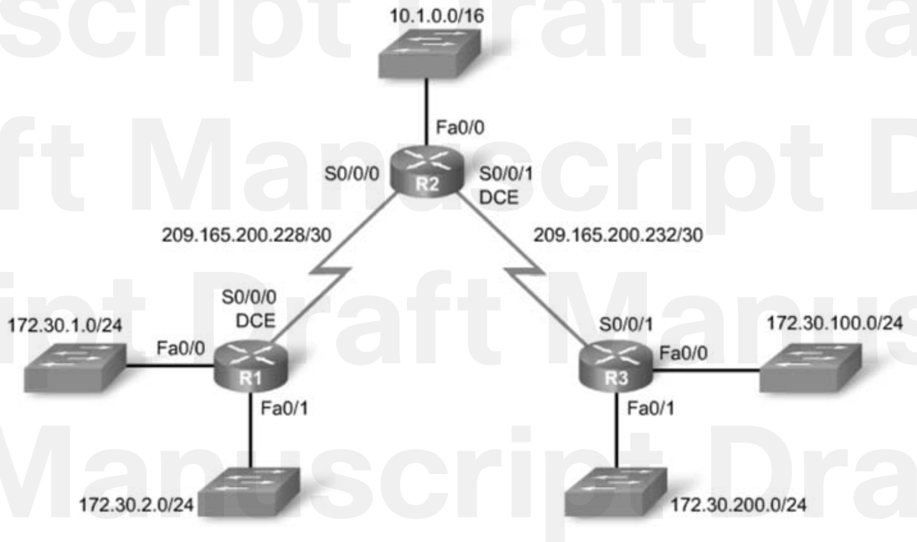

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

As you can see in Table 5-5, the addressing scheme has been changed.

Table 5-5 Addressing Scheme for Discontiguous Topology

Device	Interface	IP Address	Subnet Mask
R1	Fa0/0	172.30.1.1	255.255.255.0
	Fa0/1	172.30.2.1	255.255.255.0
	S0/0/0	209.165.200.229	255.255.255.252
R2	Fa0/0	10.1.0.1	255.255.0.0
	S0/0/0	209.165.200.230	255.255.255.252
	S0/0/1	209.165.200.233	255.255.255.252
R3	Fa0/0	172.30.100.1	255.255.255.0
	Fa0/0	172.30.200.1	255.255.255.0
	S0/0/1	209.165.200.234	255.255.255.252

This topology will be used to show a main disadvantage with classful routing protocols such as RIPv1: their lack of support for discontiguous networks.

Classful routing protocols do not include the subnet mask in routing updates. Networks are automatically summarized across major network boundaries because the receiving router in unable to determine the mask of the route. This is because the receiving interface might have a different mask than the subnetted routes.

Notice in Figure 5-12 that R1 and R3 both have subnets from the 172.30.0.0/16 major network, whereas R2 does not. Essentially, R1 and R3 are boundary routers for 172.30.0.0/16 because they are separated by another major network, 209.165.200.0/24. This separation creates a ***discontiguous network***, because two groups of 172.30.0.0/24 subnets are separated by at least one other major network. 172.30.0.0/16 is a discontiguous network.

Discontiguous Topologies Do Not Converge with RIPv1

Example 5-14 shows the RIP configuration for each router based on the topology shown in Figure 5-12.

Example 5-14 RIP Configuration for the Discontiguous Topology

```
R1(config)# router rip
R1(config-router)# network 172.30.0.0
R1(config-router)# network 209.165.200.0
R2(config)# router rip
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

```
R2(config-router)# network 10.0.0.0
R2(config-router)# network 209.165.200.0
```

```
R3(config)# router rip
R3(config-router)# network 172.30.0.0
R3(config-router)# network 209.165.200.0
```

The RIPv1 configuration is correct, but it is unable to determine all the networks in this discontiguous topology. To understand why, remember that a router will only advertise major network addresses out interfaces that do not belong to the advertised route. As a result, R1 will not advertise 172.30.1.0 or 172.30.2.0 to R2 across the 209.165.200.0 network. R3 will not advertise 172.30.100.0 or 172.30.200.0 to R2 across the 209.165.200.0 network. Both routers, however, will advertise the 172.30.0.0 major network address, a summary route to R3.

What is the result? Without the inclusion of the subnet mask in the routing update, RIPv1 cannot advertise specific routing information that will allow routers to correctly route for the 172.30.0.0/24 subnets.

Examine the routing tables for R1, R2, and R3 in Examples 5-15, 5-16, and 5-17, respectively.

Example 5-15 Routing Table for R1

```
R1# show ip route

<output omitted>

Gateway of last resort is not set

R    10.0.0.0/8 [120/1] via 209.165.200.230, 00:00:26, Serial0/0/0
     172.30.0.0/24 is subnetted, 3 subnets
R       172.30.0.0 [120/2] via 209.165.200.230, 00:00:26, Serial0/0/0
C       172.30.1.0 is directly connected, FastEthernet0/0
C       172.30.2.0 is directly connected, FastEthernet0/1
     209.165.200.0/30 is subnetted, 2 subnets
C       209.165.200.228 is directly connected, Serial0/0/0
R       209.165.200.232 [120/1] via 209.165.200.230, 00:00:26, Serial0/0/0
```

Draft Manuscript

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Example 5-16 Routing Table for R2

```
R2# show ip route

<output omitted>

Gateway of last resort is not set

     10.0.0.0/16 is subnetted, 1 subnets
C       10.1.0.0 is directly connected, FastEthernet0/0
R    172.30.0.0/16 [120/1] via 209.165.200.234, 00:00:14, Serial0/0/1
                   [120/1] via 209.165.200.229, 00:00:19, Serial0/0/0
     209.165.200.0/30 is subnetted, 2 subnets
C       209.165.200.228 is directly connected, Serial0/0/0
C       209.165.200.232 is directly connected, Serial0/0/1
```

Example 5-17 Routing Table for R3

```
R3# show ip route

<output omitted>

Gateway of last resort is not set

R    10.0.0.0/8 [120/1] via 209.165.200.233, 00:00:24, Serial0/0/1
     172.30.0.0/24 is subnetted, 3 subnets
R       172.30.0.0 [120/2] via 209.165.200.233, 00:00:22, Serial0/0/1
C       172.30.100.0 is directly connected, FastEthernet0/0
C       172.30.200.0 is directly connected, FastEthernet0/1
     209.165.200.0/30 is subnetted, 2 subnets
R       209.165.200.228 [120/1] via 209.165.200.233, 00:00:24, Serial0/0/1
C       209.165.200.232 is directly connected, Serial0/0/1
```

- R1 does not have any routes to the LANs attached to R3.
- R3 does not have any routes to the LANs attached to R1.
- R2 has two equal-cost paths to the 172.30.0.0 network.
- R2 will load-balance traffic destined for any subnet of 172.30.0.0. This means that R1 will get half of the traffic and R3 will get the other half of the traffic, whether or not the destination of the traffic is for one of their LANs.

In Chapter 7, you will see a version of this topology. It will be used to show the difference between classful and classless routing.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Packet Tracer ☐ Activity

Advantages and Disadvantages of Automatic Summarization (5.4.5)

Use the Packet Tracer Activity to implement the Scenario B addressing scheme and explore the advantages and disadvantages of automatic summarization. Detailed instructions are provided within the activity. Use file e2-545.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Default Route and RIPv1

Default routes are used by routers to represent all routes that are not specifically in the routing table. A default route is commonly used to represent routes that are not in the locally administered network, such as the Internet.

Modified Topology: Scenario C

Figure 5-13 shows a modified topology, Scenario C, to demonstrate the use of a default route and show how it is propagated by RIPv1 to other routers.

Figure 5-13 RIP Topology: Scenario C

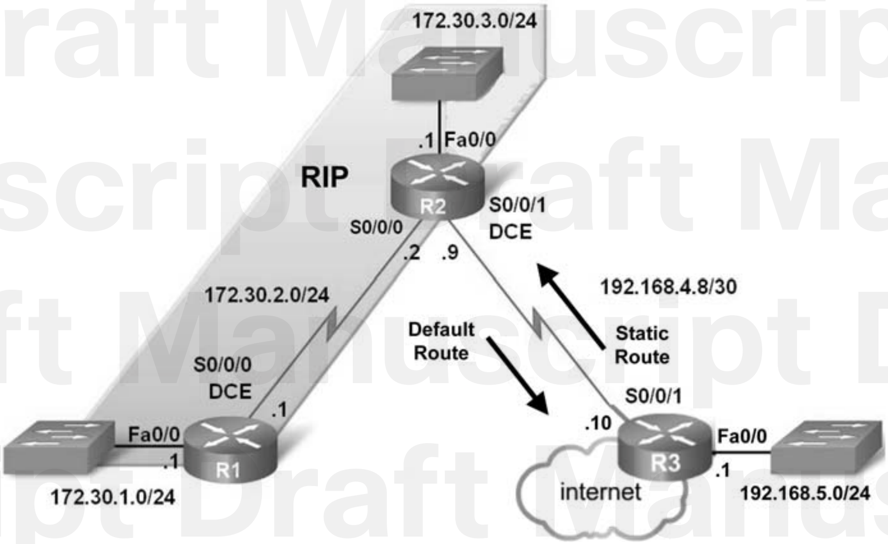

RIP was the first dynamic routing protocol and was used extensively in early implementations between customers and Internet service providers (ISP), as well as between different ISPs. But in today's networks, customers do not necessarily have to exchange routing updates with their ISP. Customer routers that connect to an ISP do not need a listing for

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

every route on the Internet. Instead, these routers have a default route that sends all traffic to the ISP router when the customer router does not have a route to a destination. The ISP configures a static route pointing to the customer router for addresses inside the customer's network.

In Scenario C, R3 is the service provider with access to the Internet, as signified by the cloud. R3 and R2 do not exchange RIP updates. Instead, R2 uses a default route to reach the R3 LAN and all other destinations that are not listed in its routing table. R3 uses a summary static route to reach the subnets 172.30.1.0, 172.30.2.0, and 172.30.3.0.

To prepare the topology, you can leave the addressing in place; it is the same as was used in Scenario B. However, you also need to complete the following steps.

Example 5-18 shows the configuration commands used in the following steps:

Step 1. Disable RIP routing for network 192.168.4.0 on R2.

Step 2. Configure R2 with a static default route to send default traffic to R3.

Step 3. Disable RIP routing on R3.

Step 4. Configure R3 with a static route to the 172.30.0.0 subnets.

Example 5-18 Configuration Changes for R2 and R3

```
R2(config)# router rip
R2(config-router)# no network 192.168.4.0
R2(config-router)# exit
R2(config)# ip route 0.0.0.0 0.0.0.0 serial 0/0/1
R3(config)# no router rip
R3(config)# ip route 172.30.0.0 255.255.252.0 serial 0/0/1
```

Examples 5-19, 5-20, and 5-21 show the new routing tables for R1, R2, and R3, respectively.

Example 5-19 Routing Table for R1

```
R1# show ip route

<output omitted>

Gateway of last resort is not set

     172.30.0.0/24 is subnetted, 3 subnets
C       172.30.1.0 is directly connected, FastEthernet0/0
C       172.30.2.0 is directly connected, Serial0/0/0
R       172.30.3.0 [120/1] via 172.30.2.2, 00:00:05, Serial0/0/0
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Draft Manuscript

Example 5-20 Routing Table for R2

```
R2# show ip route

<output omitted>

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

     172.30.0.0/24 is subnetted, 3 subnets
R       172.30.1.0 [120/1] via 172.30.2.1, 00:00:03, Serial0/0/0
C       172.30.2.0 is directly connected, Serial0/0/0
C       172.30.3.0 is directly connected, FastEthernet0/0
     192.168.4.0/30 is subnetted, 1 subnets
C       192.168.4.8 is directly connected, Serial0/0/1
S*   0.0.0.0/0 is directly connected, Serial0/0/1
```

Example 5-21 Routing Table for R3

```
R3# show ip route

<output omitted>

Gateway of last resort is not set

     172.30.0.0/22 is subnetted, 1 subnets
S       172.30.0.0 is directly connected, Serial0/0/1
     192.168.4.0/30 is subnetted, 1 subnets
C       192.168.4.8 is directly connected, Serial0/0/1
C    192.168.5.0/24 is directly connected, FastEthernet0/0
```

Propagating the Default Route in RIPv1

To provide Internet connectivity to all other networks in the RIP routing domain, the default static route needs to be advertised to all other routers that use the dynamic routing protocol. You could configure a static default route on R1 pointing to R2, but this technique is not scalable. With every router added to the RIP routing domain, you would have to configure another static default route. Why not let the routing protocol do the work for you?

In many routing protocols, including RIP, you can use the **default-information originate** command in router configuration mode to specify that this router is to originate default information, by propagating the static default route in RIP updates. In Example 5-22, R2 has been configured with the **default-information originate** command. Notice from the **debug ip rip** output that it is now sending a “quad-zero” static default route to R1.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Example 5-22 Configuring Default Route Propagation

```
R2(config)# router rip
R2(config-router)# default-information originate
R2(config-router)# end
R2# debug ip rip

RIP protocol debugging is on
RIP: sending v1 update to 255.255.255.255 via Serial0/0/0 (172.30.2.2)
RIP: build update entries
        subnet 0.0.0.0 metric 1
        subnet 172.30.3.0 metric 1
R2# undebug all

All possible debugging has been turned off
```

In the routing table for R1 (Example 5-23), you can see that there is a candidate default route, as denoted by the R* code. Cisco IOS uses the concept of *candidate default routes*, which are one or more routes marked manually or automatically as a candidate to be the default route. The actual default route installed in the routing table depends on factors such as administrative distance of the candidate. For example, a static default route will have precedence over a default route learned through a dynamic routing protocol.

Example 5-23 Verifying Default Route Propagation

```
R1# show ip route

<output omitted>
        * - candidate default, U - per-user static route, o - ODR

Gateway of last resort is 172.30.2.2 to network 0.0.0.0

     172.30.0.0/24 is subnetted, 3 subnets
C        172.30.2.0 is directly connected, Serial0/0/0
R        172.30.3.0 [120/1] via 172.30.2.2, 00:00:16, Serial0/0/0
C        172.30.1.0 is directly connected, FastEthernet0/0
R*    0.0.0.0/0 [120/1] via 172.30.2.2, 00:00:16, Serial0/0/0
```

The static default route on R2 has been propagated to R1 in a RIP update. R1 has connectivity to the LAN on R3 and any destination on the Internet.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Packet Tracer
☐ **Activity**

Propagating the Default Route in RIP (5.5.2)

Use the Packet Tracer Activity to implement Scenario C with static and default routing, and configure R2 to propagate a default route. Detailed instructions are provided within the activity. Use file e2-552.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Draft Manuscript

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Summary

RIP (version 1) is a classful distance vector routing protocol. RIPv1 was one of the first routing protocols developed for routing IP packets. RIP uses hop count for its metric, with a metric of 16 hops meaning that a route is unreachable. As a result, RIP can only be used in networks where there are no more than 15 routers between any two networks.

RIP messages are encapsulated in a UDP segment, with source and destination ports of 520. RIP routers send their complete routing tables to their neighbors every 30 seconds except for those routes that are covered by the split horizon rule.

RIP is enabled by using the **router rip** command at the global configuration prompt. The **network** command is used to specify which interfaces on the router will be enabled for RIP along with the classful network address for each directly connected network. The **network** command enables the interface to send and receive RIP updates and also advertises that network in RIP updates to other routers.

The **debug ip rip** command can be used to view the RIP updates that are sent and received by the router. To prevent RIP updates from being sent out an interface, such as on a LAN where there are no other routers, the **passive-interface** command is used.

RIP entries are displayed in the routing table with the source code of R and have an administrative distance of 120. Default routes are propagated in RIP by configuring a static default route and using the **default-information originate command** in RIP.

RIPv1 automatically summarizes subnets to their classful address when sending an update out an interface that is on a different major network than the subnetted address of the route. Because RIPv1 is a classful routing protocol, the subnet mask is not included in the routing update. When a router receives a RIPv1 routing update, RIP must determine the subnet mask of that route. If the route belongs to the same major classful network as the update, RIPv1 applies the subnet mask of the receiving interface. If the route belongs to a different major classful network than the receiving interface, RIPv1 applies the default classful mask.

The **show ip protocols** command can be used to display information for any routing protocol enabled on the router. Regarding RIP, this command displays timer information, status of automatic summarization, the networks that are enabled on this router for RIP, and other information.

Because RIPv1 is a classful routing protocol, it does not support discontiguous networks or VLSM. Both of these topics are discussed in Chapter 7.

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.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

Lab 5-1: Basic RIP Configuration (5.6.1)

In this lab, you will work through the configuration and verification commands discussed in this chapter using the same three scenarios. You will configure RIP routing, verify your configurations, investigate the problem with discontiguous networks, observe automatic summarization, and configure and propagate a default route.

Lab 5-2: Challenge RIP Configuration (5.6.2)

In this lab activity, you will be given a network address that must be subnetted to complete the addressing of the network shown in the topology diagram. A combination of RIPv1 and static routing will be required so that hosts on networks that are not directly connected will be able to communicate with each other.

Lab 5-3: RIP Troubleshooting (5.6.3)

In this lab, you will begin by loading configuration scripts on each of the routers. These scripts contain errors that will prevent end-to-end communication across the network. You will need to troubleshoot each router to determine the configuration errors, and then use the appropriate commands to correct the configurations. When you have corrected all the configuration errors, all the hosts on the network should be able to communicate with each other.

Packet Tracer
☐ **Companion**

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. Which statement is true about the **debug ip rip** command?

   A. It searches through the running configuration and shows possible errors in the RIP configuration.

   B. It displays RIP routing updates as they are sent and received.

   C. It automatically identifies routing loops.

   D. It shows the history of RIP updates over the previous 90 seconds.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

2. What problem does the **passive-interface** command help resolve?
   A. Prevents confusion if both RIPv1 and RIPv2 are being advertised on a network
   B. Prevents wasted bandwidth and processing from unnecessary updates
   C. Prevents routing loops
   D. Prevents updates from being sent out without a password

3. What makes a router a boundary router in RIP?
   A. If it is on the edge of an autonomous system
   B. If a router has multiple interfaces in more than one major classful network
   C. If it runs both RIP and EIGRP at the same time
   D. If it is configured to be a boundary router by an administrator

4. What command is used with RIP to propagate default routes to neighbors?
   A. **network 0.0.0.0**
   B. **ip summary-address rip address mask**
   C. **ip default-network address**
   D. **default-information originate**

5. What command will create a candidate default route on a RIP router?
   A. **default-information originate**
   B. **ip default-network 0.0.0.0**
   C. **ip default-gateway 192.168.0.1**
   D. **ip route 0.0.0.0 0.0.0.0 serial0/0/0**

Draft Manuscript

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

6. Refer to Figure 5-14. All routers are running RIPv1. The interfaces on all routers are up and stable. Users on the 10.16.1.0 network cannot access services on the 10.16.1.64 network. What is the cause of this problem?

Figure 5-14 Check Your Understanding: Question #6

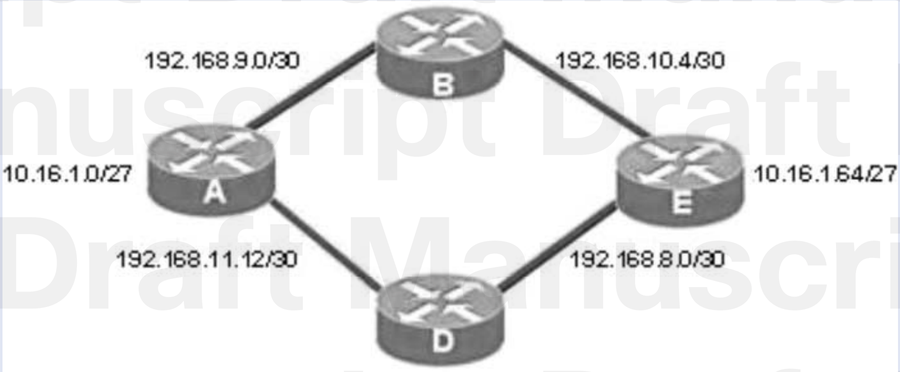

A. The RIP hold-down timer in Router A is not allowing the 10.16.1.64 network into routing updates.

B. The network uses variable-length subnet masking, and RIPv1 does not allow this.

C. The 10.16.1.x subnets are discontiguous.

D. Routers A and B need to have their interfaces configured as passive interfaces.

7. How does a router running RIPv1 determine the subnet mask of the routes that are received in routing updates?

A. The subnet mask is included in the routing update.

B. The router sends a request for the subnet mask to the sending router.

C. The router uses the subnet mask of the local interface or the default subnet mask for the address class in the routing update.

D. The router calculates the subnet mask based on the variable-length subnetting in its own configuration.

E. The router defaults to 255.255.255.0 for all updates.

Draft Manuscript

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

8. Refer to the following output. What is the administrative distance of the route to the 172.30.3.0 network?

```
<output omitted>
C        172.30.1.0 is directly connected, FastEthernet0/0
C        172.30.2.0 is directly connected, Serial0/0/0
R        172.30.3.0 [120/1] via 172.30.2.2, 00:00:05, Serial0/0/0
<output omitted>
```

A. 0
B. 1
C. 12
D. 24
E. 120

9. What is the purpose of the **network** command when RIP is being configured as the routing protocol?

A. It identifies the networks connected to the neighboring router.
B. It restricts networks from being used for static routes.
C. It identifies all the destination networks that the router is allowed to install in its routing table.
D. It identifies the directly connected networks that will be included in the RIP routing updates.

10. To ensure proper routing in a network, the network administrator should always check the router configuration to verify that appropriate routes are available. Match the commands with their appropriate function.

Command:

debug ip rip

show ip protocols

show running-config

show ip route

show interfaces

Function:

A. Displays current configuration information for configured routing protocols and interfaces
B. Checks to see that the interfaces are up and operational
C. Displays the networks advertised in the updates as the updates are sent and received
D. Verifies that the routing protocol is running and advertising the correct networks
E. Verifies that the routes received are installed in the routing table

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

11. What are the key characteristics of RIPv1?

12. Refer to Figure 5-15. HQ has connections to three branch routers (BR1, BR2, and BR3) and to the Internet through ISP. RIPv1 is configured between HQ and the branch routers. List the commands used to configure RIPv1 routing on the BR1 router.

Figure 5-15 Summary Topology

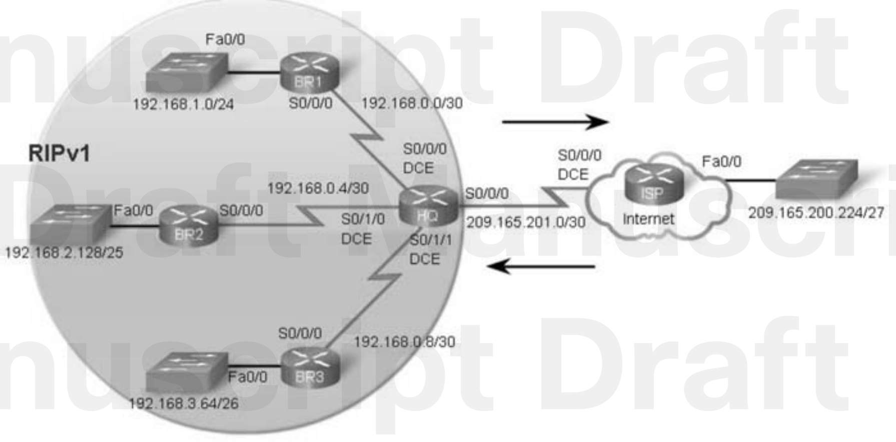

13. List the three commands used to verify and troubleshoot a RIP configuration.

14. What is the purpose of the **passive-interface** command? What is the configuration for BR1 shown in Figure 5-15, including router prompt, for this command?

15. Why would you not want to configure a dynamic routing protocol to exchange updates with your ISP?

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. What is the full routing configuration for HQ in Figure 5-15, including RIPv1, default routing, and propagating the default route to the branch routers?

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

2. The network in Figure 5-15 is not fully converged. Using only the following output from **show ip route**, determine the problem and either suggest a solution or suggest the next step in determining the source of the problem.

```
HQ# show ip route

<output omitted>

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

     192.168.0.0/30 is subnetted, 3 subnets
C       192.168.0.0 is directly connected, Serial0/0/1
C       192.168.0.4 is directly connected, Serial0/1/0
C       192.168.0.8 is directly connected, Serial0/1/1
R    192.168.1.0/24 [120/1] via 192.168.0.2, 00:00:04, Serial0/0/1
R    192.168.2.0/24 [120/1] via 192.168.0.6, 00:00:22, Serial0/1/0
     209.165.201.0/30 is subnetted, 1 subnets
C       209.165.201.0 is directly connected, Serial0/0/0
S*   0.0.0.0/0 is directly connected, Serial0/0/0
```

```
BR1# show ip route

<output omitted>

Gateway of last resort is 192.168.0.1 to network 0.0.0.0

     192.168.0.0/30 is subnetted, 3 subnets
C       192.168.0.0 is directly connected, Serial0/0/0
R       192.168.0.4 [120/1] via 192.168.0.1, 00:00:05, Serial0/0/0
R       192.168.0.8 [120/1] via 192.168.0.1, 00:00:05, Serial0/0/0
C    192.168.1.0/24 is directly connected, FastEthernet0/0
R    192.168.2.0/24 [120/2] via 192.168.0.1, 00:00:05, Serial0/0/0
R*   0.0.0.0/0 [120/1] via 192.168.0.1, 00:00:05, Serial0/0/0
```

```
BR2# show ip route

<output omitted>

Gateway of last resort is 192.168.0.5 to network 0.0.0.0

     192.168.0.0/30 is subnetted, 3 subnets
R       192.168.0.0 [120/1] via 192.168.0.5, 00:00:06, Serial0/0/0
C       192.168.0.4 is directly connected, Serial0/0/0
R       192.168.0.8 [120/1] via 192.168.0.5, 00:00:06, Serial0/0/0
R    192.168.1.0/24 [120/2] via 192.168.0.5, 00:00:01, Serial0/0/0
     192.168.2.0/25 is subnetted, 1 subnets
C       192.168.2.128 is directly connected, FastEthernet0/0
R*   0.0.0.0/0 [120/1] via 192.168.0.5, 00:00:06, Serial0/0/0
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

```
BR3# show ip route

<output omitted>

Gateway of last resort is 192.168.0.9 to network 0.0.0.0

     192.168.0.0/30 is subnetted, 3 subnets
R       192.168.0.0 [120/1] via 192.168.0.9, 00:00:08, Serial0/0/0
R       192.168.0.4 [120/1] via 192.168.0.9, 00:00:08, Serial0/0/0
C       192.168.0.8 is directly connected, Serial0/0/0
R    192.168.1.0/24 [120/2] via 192.168.0.9, 00:00:02, Serial0/0/0
R    192.168.2.0/24 [120/2] via 192.168.0.9, 00:00:08, Serial0/0/0
R*   0.0.0.0/0 [120/1] via 192.168.0.9, 00:00:08, Serial0/0/0
```

3. What static route command on ISP will summarize all the networks (and only those networks) accessible through HQ?
4. Using Packet Tracer, build and configure the topology shown in Figure 5-15.

To Learn More

Requests For Comments (RFC) are a series of documents submitted to the IETF (Internet Engineering Task Force) to propose an Internet standard or convey new concepts, information, or occasionally even humor. RFC 1058 is the original RFC for RIP written by Charles Hedrick.

RFCs can be accessed from several websites, including http://www.ietf.org/rfc/rfc1058.txt. Read all or parts of RFC 1058. Much of this information will now be familiar to you, along with some additional information.

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Check Your Understanding and Challenge Questions Answer Key

Check Your Understanding

1. B. The **debug ip rip** command displays RIP updates sent and received by the router's interfaces. This can be a useful command to verify whether the router is sending or receiving RIP updates.
2. B. The **passive-interface** command prevents RIP updates from being sent out an interface. However, this command does not prevent RIP updates from being received on that interface.
3. B. A boundary router has interfaces on more than one major classful network. For example, a router with the Serial 0/0/0 interface on the 172.16.1.0/24 network and the Serial 0/0/1 interface on the 172.30.1.0/24 network would be considered a boundary router. The 172.16.1.0/24 network is part of the 172.16.0.0/16 major classful network, whereas 172.30.1.0/24 is part of the 172.30.0.0/16 major network. Both are Class B networks but are different major networks.
4. D. The **default-information originate** command will propagate a candidate default route, a default static route, to other routers through RIP.
5. D. The candidate default route is configured using a default static route.
6. C. The 10.16.1.0/27 network and the 10.16.1.64/27 network are both 10.0.0.0 Class A networks separated by other major classful networks, such as 192.16.9.0/30 and 192.168.10.4/30.
7. C. If the route belongs to the same major classful network as the interface, RIPv1 will use the subnet mask of the interface. If the route belongs to a different major classful network than the interface address, RIPv1 will use the default classful mask of the route.
8. E. The route source of this route entry is RIP, and the administrative distance of RIP is 120. The second value in the bracket, 1, is the metric.
9. D. The **network** command is used to enable RIP on any interfaces that belong to that network address and to include those interface addresses in the RIP updates.
10. **debug ip rip**: C

    show ip protocols: D

    show running-config: A

    show ip route: E

    show interfaces: B

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

11. Explanation:

    - RIP is a distance vector routing protocol.
    - RIP uses hop count as its only metric for path selection.
    - Advertised routes with hop counts greater than 15 are unreachable.
    - Update messages are broadcast every 30 seconds.

12. Explanation:

```
BR1(config)# router rip
BR1(config-router)# network 192.168.0.0
BR1(config-router)# network 192.168.1.0
```

13. Explanation:

    - **show ip route**
    - **show ip protocols**
    - **debug ip rip**

14. The **passive-interface** command is used to stop RIP updates from being sent out an interface where RIP updates are not needed. For example, a LAN interface would only need to send out RIP updates if there is another RIP-enabled device on the LAN.

```
BR1(config-router)# passive-interface fa 0/0
```

15. ISP routers have routes to all other destinations on the Internet. Because the ISP is the default router for all traffic for which you do not have routes, the best solution is to configure a default route pointing to the ISP. The alternative is to exchange routing updates, have your routers build huge routing tables, and end up sending externally bound traffic to the ISP anyway.

Challenge Questions and Activities

1. Explanation:

```
HQ(config)# router rip
HQ(config-router)# network 192.168.0.0
HQ(config-router)# default-information originate
HQ(config-router)# exit
HQ(config)# ip route 0.0.0.0 0.0.0.0 s0/0/0
```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.

Draft Manuscript

2. The LAN for R3 is not being advertised in RIP updates. Because this LAN is also missing from the routing table for BR3, the interface is not active. Activating the interface would be the first step to solving the convergence problem. As long as BR3 is configured to advertise 192.168.3.0 in RIP updates, convergence should be achieved.

   ```
   R3(config)# interface fa0/0
   R3(config-if)# ip address 192.168.3.65 255.255.255.192
   R3(config-if)# no shutdown
   ```

   If necessary, add 192.168.3.0 to the RIP process.

   ```
   R3(config)# router rip
   R3(config-router)# network 192.168.3.0
   ```

3. Explanation:

   ```
   ISP(config)# ip route 192.168.0.0 255.255.252.0 s0/0/0
   ```

This is a prepublication draft of the manuscript. The final book will publish in December and will be available for purchase at http://www.ciscopress.com/title/9781587132063.