CHAPTER 5

Routing Information Protocol Version 2 (RIP-2)

In this chapter:

RIP Version 2 is not a new protocol—it is RIP Version 1 with some additional fields in the route update packet, key among them being subnet mask information in each route entry. The underlying DV algorithms in RIP-2 are identical to those in RIP-1, implying that RIP-2 still suffers from convergence problems and the maximum hop-count limit of 16 hops. Hence, RIP-2 may not be your choice as the routing protocol for a large or mid-sized network with multiple paths between segments. However, the new features in RIP-2 may be compelling enough for you to consider migrating an existing RIP-1 network to RIP-2. The new features in RIP-2 are summarized here:

Subnet mask
: RIP-2 updates carry the subnet mask in each route entry, making RIP-2 a classless routing protocol that supports Variable Length Subnet Masks (VLSM), discontiguous address spaces, and CIDR blocks.

Next hop IP address
: RIP-2 updates carry the next hop IP address in each route entry. As we will see later, the next hop IP address is useful when routes are being redistributed between RIP-2 and another routing protocol.

Authentication data
: Every RIP-2 packet can carry authentication data to validate the source of the RIP-2 update. Remember that RIP-1 has no security features—any host transmitting on UDP port 520 will be believed by neighbors running RIP-1.

Route tag
: RIP-2 updates carry a tag in each route entry that is not used by RIP but could be used to represent information such as the source of the route when the route is imported from another AS (for example, BGP).

These additions to the RIP-1 update take the place of the unused or "must be zero" octets in the RIP-1 packet. This strategic placement has a major goal—backward compatibility. Most versions of RIP-1 can process RIP-2 updates by ignoring the new fields.

Configuring and using RIP-2 is similar to RIP-1 and just as easy. A major reason for the long life of RIP may be the simplicity of the protocol and the ease of its use.

The next section gets RIP-2 running on TraderMary's network.

Getting RIP-2 Running

RIP-1—a *classful* routing protocol—does not support VLSM. We'll configure TraderMary's network using RIP-2—a *classless* routing protocol—much like we did using RIP-1, but we will use VLSM. The distinction between classful and classless protocols and the support of VLSM are discussed in detail in the section "Classful Versus Classless Routing Protocols."

TraderMary's network is an ideal candidate for VLSM because of the mix of user segments and serial links in the `172.16.0.0` address space. Using a 24-bit mask (`255.255.255.0`) on Ethernet segments yields 254 addresses per segment for hosts. However, serial links require only 2 IP addresses—using a 24-bit mask on a serial link wastes 252 addresses. A 30-bit mask (`255.255.255.252`) is more appropriate for a serial link, as it yields 2 usable IP addresses. How should `172.16.0.0` be segmented into 24-bit subnets for users on Ethernet segments and 30-bit subnets for serial links?

Using 24-bit masks (`255.255.255.0`) on Ethernet segments will give us 254 host addresses per user segment. Let's first use this mask to subnet `172.16.0.0`. The resulting subnets can be listed as follows:

1. `172.16.1.0/24`
2. `172.16.2.0/24`
3. `172.16.3.0/24`
4. ...

253.`172.16.253.0/24`

254.`172.16.254.0/24`

Let's now take one of these subnets (say, `172.15.250.0`) and segment it further into 30-bit subnets for serial links. The resulting subnets can be listed as follows:

1. `172.16.250.0/30`
2. `172.16.250.4/30`
3. `172.16.250.8/30`
4. `172.16.250.12/30`
5. ...

63. 172.16.250.248/30
64. 172.16.250.252/30

In these two lists we have carved the 172.16.0.0 address space using two subnet masks: 255.255.255.0 for users on Ethernet segments and 255.255.255.252 for serial links. Let's recap the steps we took. First, we used the shorter mask (255.255.255.0) and listed the resulting subnets. Next, we used one subnet from the first step and subnetted it using the longer mask (255.255.255.252). The second step is sometimes referred to as *sub-subnetting*. If we were creating a nightmare of a network and had a third mask to work with as well, we would apply the third mask (the longest mask) on one or more subnets from either of the earlier steps. Following these steps ensures that we do not create overlapping subnets.

If TraderMary's network ran out of all 64 30-bit subnets, another 24-bit subnet (say, 172.16.251.0) could be carved further to yield another 64 subnets.

See Figure 5-1 for the new addresses on TraderMary's network.

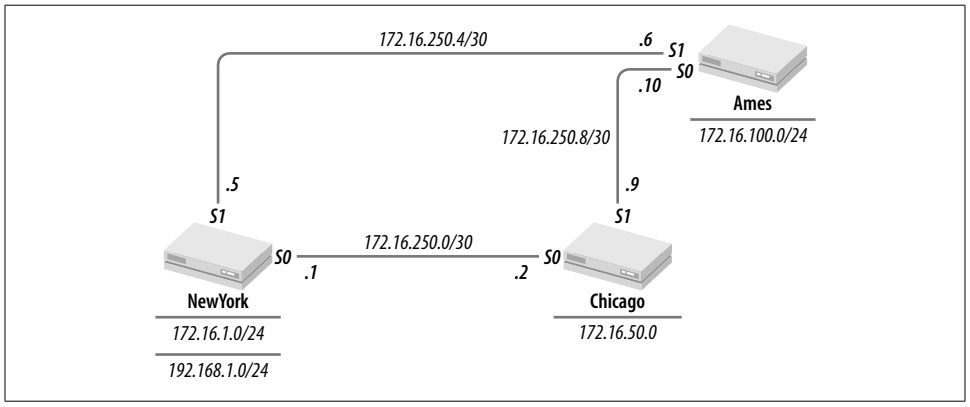

Figure 5-1. TraderMary's network with VLSM

The configuration for this network is as follows:

```
hostname NewYork
...
interface Ethernet0
ip address 172.16.1.1 255.255.255.0
!
interface Ethernet1
ip address 192.168.1.1 255.255.255.0
!
interface Serial0
description New York to Chicago link
ip address 172.16.250.1 255.255.255.252
!
interface Serial1
description New York to Ames link
bandwidth 56
```

```
ip address 172.16.250.5 255.255.255.252
...
router rip
version 2
network 172.16.0.0


hostname Chicago
...
interface Ethernet0
ip address 172.16.50.1 255.255.255.0
!
interface Serial0
description Chicago to New York link
ip address 172.16.250.2 255.255.255.252
!
interface Serial1
description Chicago to Ames link
ip address 172.16.250.9 255.255.255.0
...

router rip
version 2
network 172.16.0.0


hostname Ames
...
interface Ethernet0
ip address 172.16.100.1 255.255.255.0
!
interface Serial0
description Ames to Chicago link
ip address 172.16.250.10 255.255.255.0
!
interface Serial1
description Ames to New York link
bandwidth 56
ip address 172.16.250.6 255.255.255.0
...

router rip
version 2
network 172.16.0.0
```

Next, let's verify that all the routers are seeing all the `172.16.0.0` subnets:

```
NewYork#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default

Gateway of last resort is not set
```

```
C       192.168.1.0 is directly connected, Ethernet1
     172.16.0.0/16 is variably subnetted, 6 subnets, 2 masks
C       172.16.1.0/24 is directly connected, Ethernet0
C       172.16.250.0/30 is directly connected, Serial0
C       172.16.250.4/30 is directly connected, Serial1
R       172.16.50.0/24 [120/1] via 172.16.250.2, 0:00:11, Serial0
R       172.16.100.0/24 [120/1] via 172.16.250.6, 0:00:19, Serial1
R       172.16.250.8 [120/1] via 172.16.250.2, 0:00:11, Serial0
                     [120/1] via 172.16.250.6, 0:00:19, Serial1
```

Note that this routing table shows the mask associated with each subnet: `/24` or `/30`.

RIP-2 is supported in Cisco IOS Versions 11.1 and later.

RIP-2 Packet Format

The additions in the RIP-2 update occupy the "must be zero" or unused fields in the RIP-1 update. This careful selection of fields allows older (pre-RIP-2) implementations of RIP to interpret a RIP-2 update by just ignoring the new fields. Let's look closely at the fields in the RIP-2 update shown in Figure 5-2.

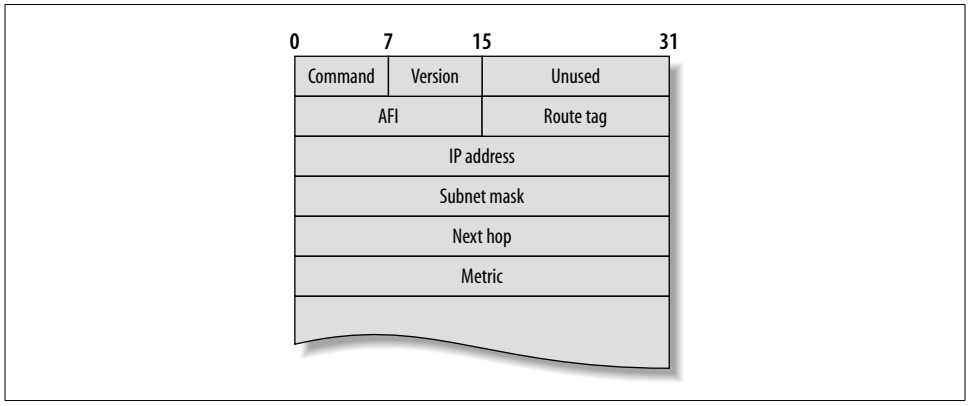

Figure 5-2. Format of RIP-2 update packet

RIP-2 updates are encapsulated in UDP port 520, like RIP-1 updates. However, the destination IP address for a RIP update can be the all-ones broadcast address of `255.255.255.255` or the reserved multicast address of `224.0.0.9`. The use of the reserved multicast address frees devices not listening to RIP-2 from the task of unwrapping RIP-2 updates.

The fields *AFI*, *IP address*, and *metric* have the same semantics as in a RIP-1 update packet. See Chapter 2 for details on these fields. The *version* field in RIP-2 updates is 2.

The *route tag* field is not used by RIP but can be used to carry an attribute assigned to the route, such as the AS number of the EGP (for example, BGP) from which the route was imported. The use of route tags is discussed further in Chapter 8.

The *subnet mask* field in each route entry classifies RIP-2 as a *classless* routing protocol and permits the use of VLSM and the support of discontiguous networks.

The *next hop* IP address is usually identical to the IP address of the source of the RIP update. For example, in TraderMary's network, *NewYork* sends an update to *Ames* for `172.16.1.0`. The source IP address of the RIP update will be `172.16.250.5`, which is identical to the next hop IP address. In such situations, the next hop field will contain no useful information and is set to `0.0.0.0`. However, consider the network shown in Figure 5-3.

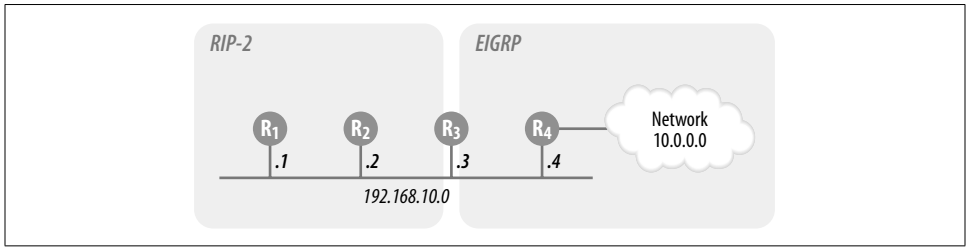

Figure 5-3. Next hop IP address

Routers *R1* and *R2* are running RIP-2. *R4* is running EIGRP, and *R3* is redistributing routes between EIGRP and RIP-2. *R4* learns `10.0.0.0` via EIGRP on interface *Serial0*. *R3* redistributes EIGRP into RIP-2. The next hop field can be used by *R3* to indicate to *R1* and *R2* that the next hop for `10.0.0.0` is `192.168.10.4`. If the next hop field were not available, *R1* and *R2* would have sent traffic for `10.0.0.0` to *R3* (`192.168.10.3`), which would then have to forward the traffic to *R4*.

If authentication is in use, the *authentication fields* take the place of the first route entry in the RIP update packet. An AFI value of 0xFFFF indicates that the route entry contains authentication data (not another route entry). RFC 1723 describes only simple (unencrypted) password authentication. This is indicated by setting the authentication type to 2, which leaves 16 octets for the password. In addition to simple password authentication, Cisco also supports MD5 authentication. When using MD5, Cisco takes the first and last route entries in each update packet to carry cryptographic checksums.

RIP-1/RIP-2 Compatibility

In Chapter 2, we configured RIP as follows on *NewYork* in TraderMary's network:

```
hostname NewYork
...
router rip
network 172.16.0.0
```

This configuration of RIP on a router running IOS 11.1 or later allows the *receipt* of both RIP-1 and RIP-2 updates but the *sending* of only RIP-1 updates.

To modify this configuration to allow the receipt of only RIP-1 updates, specify Version 1 under RIP. In the new configuration, the router will discard any RIP-2 updates it receives and will send only RIP-1 updates:

```
hostname NewYork
...
router rip
version 1
network 172.16.0.0
```

By extension, the following modification allows the receipt of only RIP-2 updates. In this configuration, the router will discard any RIP-1 updates it receives and will send only RIP-2 updates:

```
hostname NewYork
...
router rip
version 2
network 172.16.0.0
```

RIP-1/RIP-2 Interworking

The behavior of RIP can be modified further in interface configuration mode to allow for interworking between RIP-1 and RIP-2 routers.

To send only Version 1 updates out of an interface (for example, when only RIP-1 listeners exist on a network), enter the following command in interface configuration mode:

```
ip rip send version 1
```

To send only Version 2 updates out of an interface (e.g., when only RIP-2 listeners exist on a network), enter the following command in interface configuration mode:

```
ip rip send version 2
```

To send Version 1 and 2 updates out of an interface (e.g., when RIP-1 listeners and RIP-2 listeners coexist on a network), enter the following command in interface configuration mode:

```
ip rip send version 1 2
```

To receive only Version 1 updates on an interface (and to discard any RIP-2 updates), enter the following command in interface configuration mode:

```
ip rip receive version 1
```

To receive only Version 2 updates on an interface (and to discard any RIP-1 updates), enter the following command in interface configuration mode:

```
ip rip receive version 2
```

To receive Version 1 and 2 updates from an interface, enter the following command in interface configuration mode:

```
ip rip receive version 1 2
```

As an example, router *Perth*, configured as follows:

```
hostname Perth
...
router rip
version 2
network 172.22.0.0
```

has RIP-2 routers on all interfaces except *Serial2*, which has a legacy router running RIP-1. To interwork with this RIP-1 router, configure the following on *Serial2*:

```
interface Serial2
ip rip receive version 1
ip rip send version 1
```

When interworking between RIP-1 and RIP-2 and using VLSM, remember that RIP-1 updates do not carry subnet mask information. The RIP-1 portion of your network may end up with improper masks. You may have to resort to static routes or a default route in the event of a discontiguity in the RIP-1 portion of the network.

Classful Versus Classless Routing Protocols

Classful routing protocols do not carry subnet masks; classless routing protocols do. Older routing protocols, including RIP and IGRP, are classful. Newer protocols, including RIP-2, EIGRP, and OSPF, are classless. What are the implications of using classful versus classless routing protocols in your networks?

Let's say that a router *R* received a RIP-1 update with the IP address `172.0.0.0`. *R* would assume that the route being advertised was for the Class B network `172.0.0.0/16`. In other words, since the subnet mask is lacking in the routing update, *R* assumes a natural mask of `/8`, `/16`, and `/24` for Class A, B, and C addresses, respectively. The only time a classful routing protocol can associate a mask other than the natural mask with an update is if *R* has a directly connected network with an IP address belonging to the same class as the IP address received in the update. For example, when *Ames* receives an update of `172.16.1.0` from *NewYork*, *Ames* associates a mask of `/24` with the update because *Ames* is able to deduce the mask from its own interface.

RIP-2 updates carry a subnet mask in each route entry. A routing protocol that carries subnet masks in its updates earns the label "classless routing protocol." The term "classless" implies that routing decisions are not tied to the class of the IP address—A, B, or C—but may be based on any portion of the 32-bit IP address as specified by the mask. Router *R* could receive an update with the address and mask `192.168.0.0` and `255.255.0.0`. This would imply that traffic for all IP addresses with "192.168" in the first two octets should be routed as per the routing advertisement. RIP-2 is thus a classless routing protocol.

Since RIP-2 updates carry subnet masks, it is possible to associate different subnet masks within a single classful network—in other words, RIP-2 supports VLSM. VLSM, a feature of classless routing protocols, is discussed further in the next section.

VLSM

RIP-1 updates do not carry subnet mask information. A router receiving a RIP-1 route deduces the subnet mask from one of its own interfaces, if the router has the same network number. So, for example, when *NewYork* receives the update `172.16.100.0` from *Ames* it assumes that the mask for this network number is `255.255.255.0` because *NewYork* has an interface (*Ethernet0*) with the same mask. When using RIP-1, there is no room for the support of VLSM.

RIP-2 updates carry subnet masks, so a router receiving the update does not have to guess the mask. RIP-2 updates can carry masks of any length. This permits the network engineer to assign subnet masks that match the true size of the host population. The RIP-2 configuration of TraderMary's network used 24-bit masks for user segments and 30-bit masks for serial links.

When carving a network number into subnets of varying length, it is key that the two subnet populations not overlap. One way to tackle this is to first carve the address space using the shorter mask and then use one or more of the resulting subnets and carve it further using the longer mask, as we did for TraderMary's network.

Use of Subnet Zero

A zero subnet has all zeros in the subnet portion of the IP address. For example, `172.16.0.0/24` (with host addresses in the range `172.16.0.1` through `172.16.0.254`) is a zero subnet. `192.168.100.0/26` is also a zero subnet: the subnet bits are bits 25 and 26 in the IP address, and both are zero.

Zero subnets cannot be used with classful routing protocols. This is because an update for the subnet (without the mask) is indistinguishable from an update for the entire network number. If router *R* received an update for `172.16.0.0`, it could not tell if the update was for the entire Class B or just a zero subnet, such as `172.16.0.0/24`. Similarly, an update for `192.168.100.0` could mean a path to the entire Class C or just to a zero subnet, such as `192.168.100.0/28`. Because of this ambiguity, zero subnets are not permitted to be configured by Cisco IOS. However, a classless routing protocol clearly distinguishes between a zero subnet and the entire network. So, `172.16.0.0 255.255.255.0` would represent a zero subnet, whereas `172.16.0.0 255.255.0.0` would represent the entire network. To configure subnet zero on a router interface, a special command has to be turned on in global configuration mode:

```
ip subnet zero
```

This command relaxes the IOS restriction on configuring zero subnets.

Classless Inter-Domain Routing (CIDR)

Another feature of classless routing protocols is the support of CIDR. The primary use of CIDR is to reduce the size of routing tables by aggregating several classful

addresses in a single route entry. All Class C addresses in the range 192.168.0.0 through 192.168.255.0 can be represented by the single route 192.168.0.0/16.

The use of CIDR is most relevant in the Internet, where Class C addresses have been allocated to various service providers in blocks. We will thus reserve further discussion of CIDR to Chapter 7, where we discuss BGP and Internet routing.

Classful Versus Classless Route Lookup

To route a packet, all routers must extract the destination IP address in the packet header. Older (or "classful") routers take this address and compute its major Class A, B, or C network number (for example, the address 172.16.1.1 belongs to the major network 172.16.0.0). This major network number is matched in the routing table. If there is no matching major network number (and there is no default route in the routing table), the packet is dropped. If there is a match against the major network number, the router proceeds to match the subnet field. If there is no matching subnet field in the routing table, the packet is dropped. If there is a matching subnet field, the packet is routed as specified in the route entry. This "classful" routing behavior is described in more detail in Chapter 3.

Classless route lookups also refer to the destination IP address in the packet header. However, classless route lookups do not compute the major Class A, B, or C network number for the destination IP address. Instead, classless routing protocols use a rule called *longest prefix match*. By this rule, the destination IP address from the packet header is matched bit-by-bit against every destination IP address in the routing table. The route entry that has the longest bitwise match with the destination IP address is chosen for routing the packet.

To turn on classless route lookups, enter the following command in global configuration mode:

```
ip classless
```

To turn on classful route lookups, enter the following command in global configuration mode:

```
no ip classless
```

Authentication

There are two reasons to authenticate a routing update. First, for security. After all, if an intruder gains access to a network and begins announcing RIP routes, she will at least disrupt traffic and, in a worse scenario, may maliciously reroute traffic to steal critical data. The second reason for authenticating routing updates is to guard against misconfiguration. For example, using a password on a network backbone will ensure that if a router is attached to the backbone by mistake, it won't begin participating in the backbone routing protocol.

Cisco's implementation of RIP-2 supports two authentication modes: plain-text and MD5. Plain-text authentication works well to guard against misconfigurations but is not a great security solution, since plain-text passwords can be gleaned with a network sniffer.

Passwords must first be defined on each router in global configuration mode. Cisco uses the construct of a "key chain" to define passwords. Let's define a key chain with the name *EmpireStateBldg* on router *NewYork*. The passwords on this key chain are *2000feet* and *1782 feet*.

```
key chain EmpireStateBldg
  key 1
  key-string 2000feet
  key 2
  key-string 1782 feet
```

Routers *Chicago* and *Ames* in TraderMary's network must also be configured with the passwords *2000feet* and *1782 feet*. *Chicago* may be configured as follows:

```
key chain SearsTower
  key 1
  key-string 2000feet
  key 2
  key-string 1782 feet
```

Note that the names of the key chains are not significant: the names of the key chains can be different on each router. The passwords—*2000feet* and *1782 feet*—are significant and must match.

To configure these passwords on an interface, apply the key chain to the interface:

```
hostname NewYork
...
interface Ethernet0
ip address 172.16.1.1 255.255.255.0
!
interface Ethernet1
ip address 192.168.1.1 255.255.255.0
!
interface Serial0
description Link to Chicago
ip address 172.16.250.1 255.255.255.0
ip rip authentication key-chain EmpireStateBldg
!
interface Serial1
description Link to Ames
ip address 172.16.251.1 255.255.255.0
ip rip authentication key-chain EmpireStateBldg
ip rip authentication mode md5
...
router rip
version 2
network 172.16.0.0
```

In this configuration, *Serial1* (to *Ames*) is configured for encryption using MD5, whereas *Serial0* (to *Chicago*) is configured for plain-text authentication, which is the default. *Ames* and *Chicago* would have to be configured for MD5 and plain-text authentication, respectively.

A password encrypted using MD5 cannot be read in plain text, but someone could still copy the encrypted string and play it back. Hence, Cisco introduced the concept of key management, which allows you to define several passwords. The password used at any given time can be defined as follows:

```
key chain EmpireStateBldg
  key 1
  key-string 2000feet
  accept-lifetime 13:00:00 Dec 19 1999 13:00:00 Jan 14 2000
  send-lifetime 13:00:00 Dec 19 1999 13:00:00 Jan 14 2000
  key 2
  key-string 1782 feet
  accept-lifetime 12:00:00 Jan 14 2000 infinite
  send-lifetime 12:00:00 Dec 19 2000 infinite
```

In this example, *2000feet* is a valid password from 1:00 P.M., December 19, 1999 until 1:00 P.M., January 14, 2000. Note that there is an overlap of 1 hour on January 14 (12:00:00 to 13:00:00) during which both *2000feet* and *1782 feet* are valid passwords. This overlap is important to allow for differences in the clocks on the routers, although a time-synchronization protocol such as the Network Time Protocol can also be used to address this issue.

If the lifetime of a key is not specified, the password is always valid.

To check which passwords are active on a router at any given time, use the following command:

```
Chicago#sh key chain
Key-chain SearsTower:
    key 1 -- text "1782feet"
        accept lifetime (13:00:00 Dec 19 1999) - (13:00:00 Jan 14 2000) [valid now]
        send lifetime (13:00:00 Dec 19 1999) - (13:00:00 Jan 14 2000) [valid now]
```

Remember that authentication is available only in RIP Version 2; authentication is not an option when interworking between RIP-1 and RIP-2 routers.

Route Summarization

RIP-2 summarizes on route boundaries just like RIP-1. However, given that RIP-2 is a classless protocol and carries subnet mask information in its updates, it makes sense to allow the network engineer to turn off route summarization to support discontiguous networks. The following command in global configuration mode turns off route summarization:

```
router rip
no auto-summary
```

Summing Up

Why bother with RIP-2? RIP-2, after all, is still RIP. There are still the issues of convergence times and a maximum diameter of 15 hops. Routing updates are sent every 30 seconds and consume network resources. The metric does not account for link bandwidth or delay. These issues with RIP may loom large in your mind if you are building a network from scratch. You have the choice of other, newer routing protocols that do not present these headaches (although they do present other headaches). However, if you are building a small, homogenous network and are not too concerned about occasional convergence problems, RIP-2 may be ideal for you.

RIP-2 may also be a good choice if you are currently running RIP-1 and are happy with it. Maybe your network is small and likely to remain that way. Maybe the link types and speeds in your network are homogenous, so the issue of RIP metrics hasn't bothered you. And maybe there aren't so many paths between any pair of nodes that RIP gets lost during convergence. If you are happy with RIP-1, migrating to RIP-2 may be an excellent solution if you need VLSM, discontiguous address spaces, or authentication. You would still be dealing with RIP—familiar, easy to configure, and reliable—but would have the added benefits of Version 2.