

CCENT Study Guide

Chapter 9

IP Routing

Chapter 9 Objectives

- The CCENT Topics Covered in this chapter include:
- **3.0 Routing Technologies**
- 3.1 Describe the routing concepts.
- 3.1.a Packet handling along the path through a network
- 3.1.b Forwarding decision based on route lookup
- 3.1.c Frame rewrite
- 3.2 Interpret the components of a routing table.
- 3.2.a Prefix
- 3.2.b Network mask
- 3.2.c Next hop
- 3.2.d Routing protocol code
- 3.2.e Administrative distance
- 3.2.f Metric
- 3.2.g Gateway of last resort
- 3.3 Describe how a routing table is populated by different routing information sources.
- 3.3.a Admin distance
- 3.5 Compare and contrast static routing and dynamic routing.
- 3.6 Configure, verify, and troubleshoot IPv4 and IPv6 static routing.
- 3.6.a Default route
- 3.6.b Network route
- 3.6.c Host route
- 3.6.d Floating static
- 3.7 Configure, verify, and troubleshoot RIPv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution).

Routing Basics

Here's an important list of the minimum factors a router must know to be able to effectively route packets:

- Destination address
- Neighbor routers from which it can learn about remote networks
- Possible routes to all remote networks
- The best route to each remote network
- How to maintain and verify routing information

The router learns about remote networks from neighboring routers or from an administrator.

The router then builds a routing table, which is basically a map of the internetwork, and it describes how to find remote networks.

If a network is directly connected, then the router already knows how to get to it.

Figure 9.1: A simple routing example




Figure 9.1 shows a simple network. Lab A has four interfaces. Can you see which interface will be used to forward an IP datagram to a host with a destination IP address of 10.10.10.30?

Show ip route

By using the command `show ip route` on a router, we can see the routing table (map of the internetwork) that Lab A has used to make its forwarding decisions:

```
Lab_A#sh ip route
```

```
Codes: L - local, C - connected, S - static,
```

```
[output cut]
```

```
    10.0.0.0/8 is variably subnetted, 6 subnets, 4 masks  
C    10.0.0.0/8 is directly connected, FastEthernet0/3  
L    10.0.0.1/32 is directly connected, FastEthernet0/3  
C    10.10.0.0/16 is directly connected, FastEthernet0/2  
L    10.10.0.1/32 is directly connected, FastEthernet0/2  
C    10.10.10.0/24 is directly connected, FastEthernet0/1  
L    10.10.10.1/32 is directly connected, FastEthernet0/1  
S*   0.0.0.0/0 is directly connected, FastEthernet0/0
```

The C in the routing table output means that the networks listed are “directly connected,” and until we add a routing protocol like RIPv2, OSPF, etc., to the routers in our internetwork, or enter static routes, only directly connected networks will show up in our routing table. What about that L in the routing table—that’s new, isn’t it? Yes it is, because in the new Cisco IOS 15 code, Cisco defines a different route, called a local route. Each has a /32 prefix defining a route just for the one address.

The IP Routing Process

The IP routing process is fairly simple and doesn't change, regardless of the size of your network. For a good example of this fact, I'll use Figure 9.2 to describe step-by-step what happens when Host A wants to communicate with Host B on a different network.


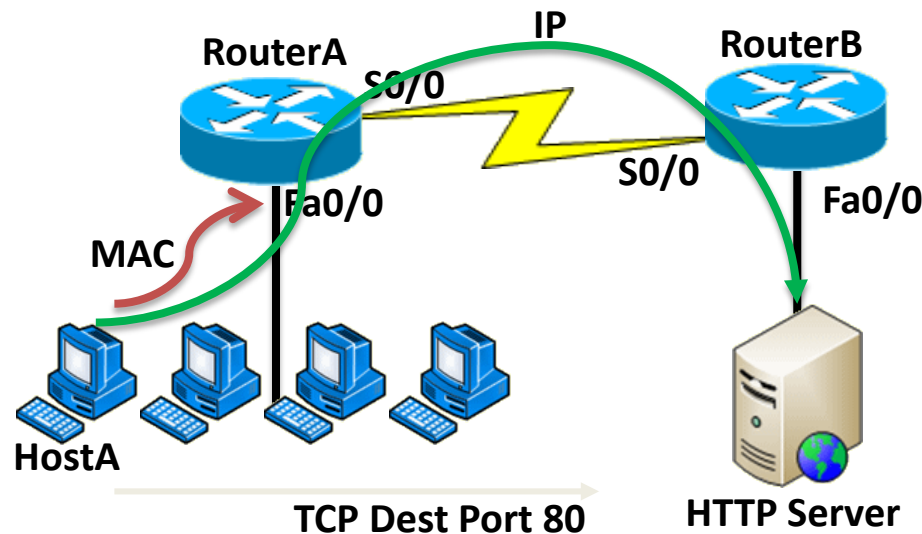


Figure 9.4: IP routing example 1

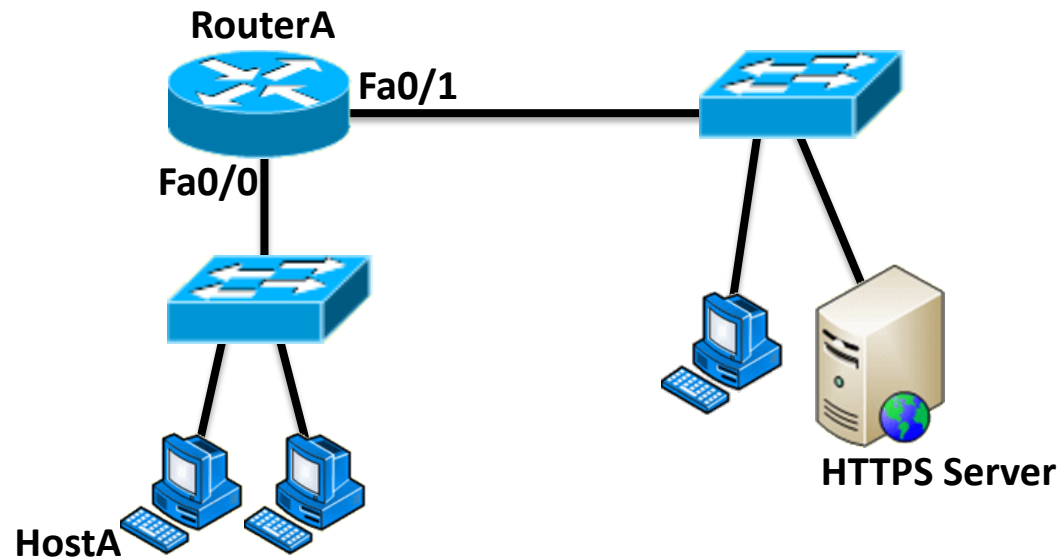
Figure 9.4 shows a LAN connected to RouterA that's connected via a WAN link to RouterB. RouterB has a LAN connected with an HTTP server attached.



1. The destination address of a frame from HostA would be the MAC address of Router A's Fa0/0 interface.
2. The destination address of a packet would be the IP address of the HTTP server's network interface card (NIC).
3. The destination port number in the segment header would be 80.

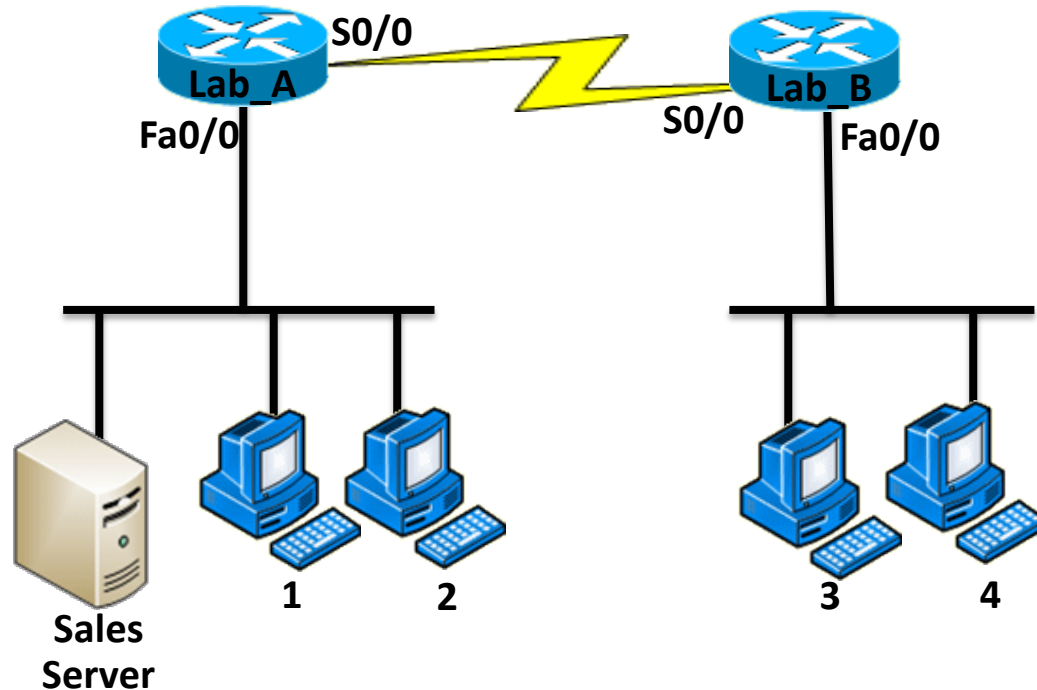
Figure 9.5: IP routing example 2

Figure 9.5 shows a network with only one router but two switches.



1. The destination address of a frame from Host A would be the MAC address of Router A's Fa0/0 interface.
2. The destination address of a packet is the IP address of the HTTPS server's network interface card (NIC).
3. The destination port number in the segment header will have a value of 443.

Figure 9.6: Basic IP routing using MAC and IP addresses



1. In order to begin communicating with the Sales server, Host 4 sends out an ARP request. How will the devices exhibited in the topology respond to this request?
2. Host 4 has received an ARP reply. Host 4 will now build a packet, then place this packet in the frame. What information will be placed in the header of the packet that leaves Host 4 if Host 4 is going to communicate to the Sales server?
3. The Lab A router has received the packet and will send it out Fa0/0 onto the LAN toward the server. What will the frame have in the header as the source and destination addresses?
4. Host 4 is displaying two web documents from the Sales server in two browser windows at the same time. How did the data find its way to the correct browser windows?

IP Routing Configuration

These are the three routing methods I'm going to cover with you:

- Static routing
- Default routing
- Dynamic routing

We're going to start with the first way and implement static routing on our network, because if you can implement static routing *and* make it work, you've demonstrated that you definitely have a solid understanding of the internetwork.

Static Routing

Starting at the beginning, here's the command syntax you use to add a static route to a routing table from global config:

```
ip route [destination_network] [mask] [next-hop_address or  
exitinterface] [administrative_distance] [permanent]
```

ip route

The command used to create the static route

destination_network

The network you're placing in the routing table

mask

The subnet mask being used on the network

next-hop_address

This is the IP address of the next-hop router that will receive packets and forward them to the remote network, which must signify a router interface that's on a directly connected network. You must be able to successfully ping the router interface before you can add the route. Important note to self is that if you type in the wrong next-hop address or the interface to the correct router is down, the static route will show up in the router's configuration but not in the routing table

exitinterface

Used in place of the next-hop address if you want, and shows up as a directly connected route

administrative_distance

By default, static routes have an administrative distance of 1 or 0 if you use an exit interface instead of a next-hop address. You can change the default value by adding an administrative weight at the end of the command. I'll talk a lot more about this later in the chapter when we get to the section on dynamic routing.

permanent

If the interface is shut down or the router can't communicate to the next-hop router, the route will automatically be discarded from the routing table by default. Choosing the permanent option keeps the entry in the routing table no matter what happens.

Static Route Examples

Let's take a look at a sample static route to see what we can find out about it:

```
Router(config)#ip route 172.16.3.0 255.255.255.0 192.168.2.4
```

The ip route command tells us simply that it's a static route.
172.16.3.0 is the remote network we want to send packets to.
255.255.255.0 is the mask of the remote network.
192.168.2.4 is the next hop, or router, that packets will be sent to.

But what if the static route looked like this instead?

```
Router(config)#ip route 172.16.3.0 255.255.255.0 192.168.2.4 150
```

That 150 at the end changes the default administrative distance (AD) of 1 to 150. As said, I'll talk much more about AD when we get into dynamic routing, but for now, just remember that the AD is the trustworthiness of a route, where 0 is best and 255 is worst.

One more example:

```
Router(config)#ip route 172.16.3.0 255.255.255.0 s0/0/0
```

Instead of using a next-hop address, we can use an exit interface that will make the route show up as a directly connected network. Functionally, the next hop and exit interface work exactly the same.

Default Routing

A *stub* indicates that the networks in this design have only one way out to reach all other networks, which means that instead of creating multiple static routes, we can just use a single default route. This default route is used by IP to forward any packet with a destination not found in the routing table, which is why it is also called a gateway of last resort.

```
LA#config t
```

```
LA(config)#ip route 0.0.0.0 0.0.0.0 172.16.10.5
```

```
LA(config)#do sho ip route
```

```
[output cut]
```

```
Gateway of last resort is 172.16.10.5 to network 0.0.0.0
```

```
172.16.0.0/30 is subnetted, 1 subnets
```

```
C    172.16.10.4 is directly connected, Serial0/0/1
```

```
L    172.16.10.6/32 is directly connected, Serial0/0/1
```

```
C   192.168.20.0/24 is directly connected, FastEthernet0/0
```

```
L   192.168.20.0/32 is directly connected, FastEthernet0/0
```

```
S*  0.0.0.0/0 [1/0] via 172.16.10.5
```

Can you see the default route listed last in the routing table? The S* shows that as a candidate for the default route. And I really want you to notice that the gateway of last resort is now set too.

Dynamic Routing

Administrative Distances

The *administrative distance (AD)* is used to rate the trustworthiness of routing information received on a router from a neighbor router. An administrative distance is an integer from 0 to 255, where 0 is the most trusted, and 255 means no traffic will be passed via this route.

If a router receives two updates listing the same remote network, the first thing the router checks is the AD. If one of the advertised routes has a lower AD than the other, then the route with the lowest AD will be chosen and placed in the routing table.

Table 8.1: Default administrative distances

Route Source	Default AD
Connected interface	0
Static route	1
EIGRP	90
OSPF	110
RIP	120
External EIGRP	170
Unknown	255 (This route will never be used.)

Routing Protocols

There are three classes of routing protocols:

Distance vector

The distance-vector protocols in use today find the best path to a remote network by judging distance. In RIP routing, each instance where a packet goes through a router is called a hop, and the route with the least number of hops to the network will be chosen as the best one. The vector indicates the direction to the remote network. RIP is a distance-vector routing protocol and periodically sends out the entire routing table to directly connected neighbors.

Link state

In link-state protocols, also called shortest-path-first protocols, the routers each create three separate tables. One of these tables keeps track of directly attached neighbors, one determines the topology of the entire internetwork, and one is used as the routing table. OSPF is an IP routing protocol that's completely link-state. Link-state protocols send updates containing the state of their own links to all other directly connected routers on the network. This is then propagated to their neighbors.

Hybrid

Hybrid protocols use aspects of both distance-vector and link-state protocols, and EIGRP is a great example—even though Cisco typically just calls EIGRP an advanced distance-vector routing protocol!

Written Labs and Review Questions

- Read through the Exam Essentials section together in class.
- Open your books and go through all the written labs and the review questions.
- Review the answers in class.