

# Troubleshooting IP Routing

---

This troubleshooting chapter has several goals. First, it explains several tools and functions not covered in Chapters 4 through 6—specifically, tools that can be very helpful when you’re analyzing problems. This chapter also reviews concepts from all three of the other chapters in Part II, “IP Routing.” It pulls together the concepts by showing a suggested process for troubleshooting routing problems, as well as examples of how to use the process. The second half of the chapter focuses on a series of troubleshooting tips for many of the specific topics covered in Chapters 4 through 6.

## “Do I Know This Already?” Quiz

The troubleshooting chapters of this book pull in concepts from many other chapters, including some chapters in *CCENT/CCNA ICND1 Official Exam Certification Guide*. They also show you how to approach some of the more challenging questions on the CCNA exams. Therefore, it is useful to read these chapters regardless of your current knowledge level. For these reasons, the troubleshooting chapters do not include a “Do I Know This Already?” quiz. However, if you feel particularly confident about troubleshooting IP routing features covered in this book and *CCENT/CCNA ICND1 Official Exam Certification Guide*, feel free to move to the “Exam Preparation Tasks” section near the end of this chapter to bypass the majority of the chapter.

---

## Foundation Topics

---

This chapter focuses on troubleshooting the IP routing process. To that end, it begins with a section about two important troubleshooting tools: ping and traceroute. Following that, the chapter examines the IP routing process from a troubleshooting perspective, particularly focusing on how to isolate routing problems to identify the root cause of the problem. The final section covers a wide variety of small topics, all of which can be useful when you're troubleshooting IP routing problems.

**NOTE** This chapter, and Chapter 15 in *CCENT/CCNA ICND1 Official Exam Certification Guide*, both explain details of how to troubleshoot the IP routing process. IP routing is vitally important on both the ICND1 and ICND2 exams, as well as on the CCNA exam, so there is overlap between the exams, requiring some overlap in the books. However, this chapter covers many topics that go beyond the details required for the ICND1 exam. To be fully prepared, read this entire chapter, but feel free to skim portions if the chapter seems repetitive with the ICND1 book.

## The ping and traceroute Commands

This section examines a suggested process of troubleshooting IP routing—in other words, the data plane process of how hosts and routers forward IP packets. To that end, this section first examines a set of useful tools and protocols—in particular, ICMP, **ping**, and **traceroute**. Following that, the text suggests a good general troubleshooting process for IP problems, with a few examples to show how to use the processes.

### Internet Control Message Protocol (ICMP)

TCP/IP includes ICMP, a protocol designed to help manage and control the operation of a TCP/IP network. The ICMP protocol provides a wide variety of information about a network's health and operational status. *Control Message* is the most descriptive part of the name. ICMP helps control and manage IP's work by defining a set of messages and procedures about the operation of IP. Therefore, ICMP is considered part of TCP/IP's network layer. Because ICMP helps control IP, it can provide useful troubleshooting information. In fact, the ICMP messages sit inside an IP packet, with no transport layer header, so ICMP is truly an extension of the TCP/IP network layer.

RFC 792 defines ICMP. The following excerpt from RFC 792 describes the protocol well:

Occasionally a gateway (router) or destination host will communicate with a source host, for example, to report an error in datagram processing. For such

purposes, this protocol, the Internet Control Message Protocol (ICMP), is used. ICMP uses the basic support of IP as if it were a higher level protocol; however, ICMP is actually an integral part of IP and must be implemented by every IP module.

ICMP defines several different types of messages to accomplish its varied tasks, as summarized in Table 7-1.

**Table 7-1** *ICMP Message Types*

Message	Description
Destination Unreachable	Tells the source host that there is a problem delivering a packet.
Time Exceeded	The time that it takes a packet to be delivered has expired, so the packet has been discarded.
Redirect	The router sending this message has received a packet for which another router has a better route. The message tells the sender to use the better route.
Echo Request, Echo Reply	Used by the <b>ping</b> command to verify connectivity.



#### **The ping Command and the ICMP Echo Request and Echo Reply**

The **ping** command uses the ICMP Echo Request and Echo Reply messages. In fact, when people say they sent a ping packet, they really mean that they sent an ICMP Echo Request. These two messages are somewhat self-explanatory. The Echo Request simply means that the host to which it is addressed should reply to the packet. The Echo Reply is the ICMP message type that should be used in the reply. The Echo Request includes some data that can be specified by the **ping** command; whatever data is sent in the Echo Request is sent back in the Echo Reply.

The **ping** command itself supplies many creative ways to use Echo Requests and Replies. For instance, the **ping** command lets you specify the length as well as the source and destination addresses, and it also lets you set other fields in the IP header. Chapter 4, “IP Routing: Static and Connected Routes,” shows an example of the extended **ping** command that lists the various options.

#### **The Destination Unreachable ICMP Message**

This book focuses on IP. But if you take a broader view, the role of the entire set of TCP/IP protocols is to deliver data from the sending application to the receiving application. Hosts and routers send ICMP Destination Unreachable messages back to the sending host when that host or router cannot deliver the data completely to the application at the destination host.

To aid in troubleshooting, the ICMP Unreachable message includes five separate unreachable functions (codes) that further identify the reason why the packet cannot be delivered. All five code types pertain directly to an IP, TCP, or UDP feature.

For example, the internetwork shown in Figure 7-1 can be used to better understand some of the Unreachable codes. Assume that Fred is trying to connect to the web server, called Web. (Web uses HTTP, which in turn uses TCP as the transport layer protocol.) Three of the ICMP unreachable codes can possibly be used by Routers A and B. The other two codes are used by the web server. These ICMP codes are sent to Fred as a result of the packet originally sent by Fred.

**Figure 7-1** Sample Network for Discussing ICMP Unreachable Codes

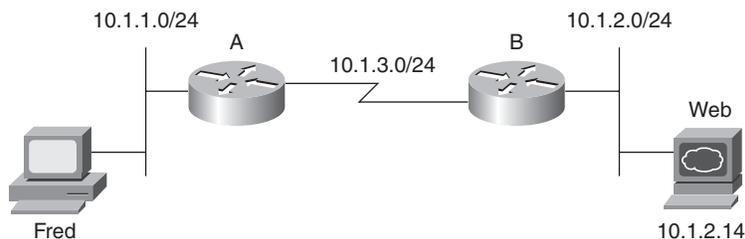


Table 7-2 summarizes the more common ICMP unreachable codes. After the table, the text explains how each ICMP code might be needed for the network shown in Figure 7-1.

**Table 7-2** ICMP Unreachable Codes

Unreachable Code	When It Is Used	What Typically Sends It
Network unreachable	There is no match in a routing table for the packet's destination.	Router
Host unreachable	The packet can be routed to a router connected to the destination subnet, but the host is not responding.	Router
Can't fragment	The packet has the Don't Fragment bit set, and a router must fragment to forward the packet.	Router
Protocol unreachable	The packet is delivered to the destination host, but the transport layer protocol is not available on that host.	Host
Port unreachable	The packet is delivered to the destination host, but the destination port has not been opened by an application.	Host

The following list explains each code in Table 7-2 in greater detail using the network in Figure 7-1 as an example:

- **Network unreachable:** Router A uses this code if it does not have a route telling it where to forward the packet. In this case, Router A needs to route the packet to subnet 10.1.2.0/24. If it cannot, Router A sends Fred the ICMP Destination Unreachable message with the code “network unreachable” in response to Fred’s packet destined for 10.1.2.14.
- **Host unreachable:** This code implies that the single destination host is unavailable. If Router A has a route to 10.1.2.0/24, the packet is delivered to Router B. If Router B’s LAN interface is working, B also has a connected route to 10.1.2.0/24, so B tries to ARP and learn the web server’s MAC address. However, if the web server is down, Router B does not get an ARP reply from the web. Router B sends Fred the ICMP Destination Unreachable message with the code “host unreachable,” meaning that B has a route but cannot forward the packet directly to 10.1.2.14.
- **Can’t fragment:** This code is the last of the three ICMP unreachable codes that a router might send. Fragmentation defines the process in which a router needs to forward a packet, but the outgoing interface allows only packets that are smaller than the packet. The router is allowed to fragment the packet into pieces, but the packet header can be set with the “Do Not Fragment” bit in the IP header. In this case, if Router A or B needs to fragment the packet, but the Do Not Fragment bit is set in the IP header, the router discards the packet and sends Fred an ICMP Destination Unreachable message with the code “can’t fragment.”
- **Protocol unreachable:** If the packet successfully arrives at the web server, two other unreachable codes are possible. One implies that the protocol above IP, typically TCP or UDP, is not running on that host. This is highly unlikely, because most operating systems that use TCP/IP use a single software package that provides IP, TCP, and UDP functions. But if the host receives the IP packet and TCP or UDP is unavailable, the web server host sends Fred the ICMP Destination Unreachable message with the code “protocol unreachable” in response to Fred’s packet destined for 10.1.2.14.
- **Port unreachable:** This final code field value is more likely today. If the server—the computer—is up and running, but the web server software is not running, the packet can get to the server but cannot be delivered to the web server software. In effect, the server is not listening on that application protocol’s well-known port. So, host 10.1.2.14 sends Fred the ICMP Destination Unreachable message with the code “port unreachable” in response to Fred’s packet destined for 10.1.2.14.

**NOTE** Most security policies today filter these various unreachable messages to help bolster the network’s security profile.

The **ping** command lists various responses that in some cases imply that an unreachable message was received. Table 7-3 lists the various unreachable codes that may be displayed by the Cisco IOS Software **ping** command.

**Table 7-3** Codes That the **ping** Command Receives in Response to Its ICMP Echo Request

<b>ping Command Code</b>	<b>Description</b>
!	ICMP Echo Reply received
.	Nothing was received before the <b>ping</b> command timed out
U	ICMP unreachable (destination) received
N	ICMP unreachable (network/subnet) received
M	ICMP Can't Fragment message received
?	Unknown packet received

#### **The Redirect ICMP Message**

The ICMP Redirect message provides a means by which routers can tell hosts to use another router as default gateway for certain destination addresses. Most hosts use the concept of a default router IP address, sending packets destined for subnets to their default router. However, if multiple routers connect to the same subnet, a host's default gateway may not be the best router on that subnet to which to forward packets sent to some destinations. The default gateway can recognize that a different router is a better option. Then it can send ICMP redirect messages to the host to tell it to send the packets for that destination address to this different router.

For example, in Figure 7-2, the PC uses Router B as its default router. However, Router A's route to subnet 10.1.4.0 is a better route. (Assume the use of mask 255.255.255.0 in each subnet in Figure 7-2.) The PC sends a packet to Router B (Step 1 in Figure 7-2). Router B then forwards the packet based on its own routing table (Step 2); that route points through Router A, which has a better route. Finally, Router B sends the ICMP redirect message to the PC (Step 3), telling it to forward future packets destined for 10.1.4.0 to Router A instead. Ironically, the host can ignore the redirect and keep sending the packets to Router B, but in this example, the PC believes the redirect message, sending its next packet (Step 4) directly to Router A.

#### **The ICMP Time Exceeded Message**

The ICMP Time Exceeded message notifies a host when a packet it sent has been discarded because it was "out of time." Packets are not actually timed, but to prevent them from being forwarded forever when there is a routing loop, each IP header uses a Time to Live (TTL) field. Routers decrement the TTL by 1 every time they forward a packet; if a router

decrements the TTL to 0, it throws away the packet. This prevents packets from rotating forever. Figure 7-3 shows the basic process.

Figure 7-2 ICMP Redirect

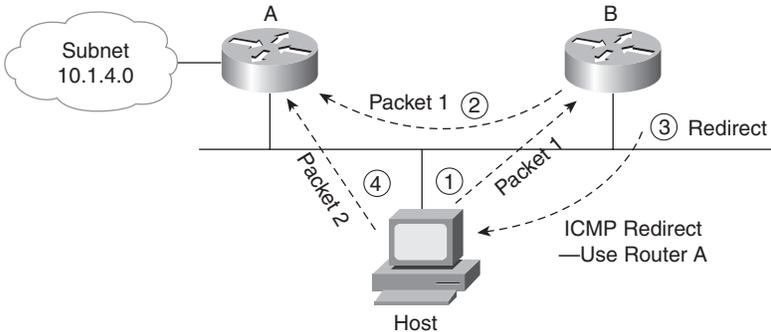
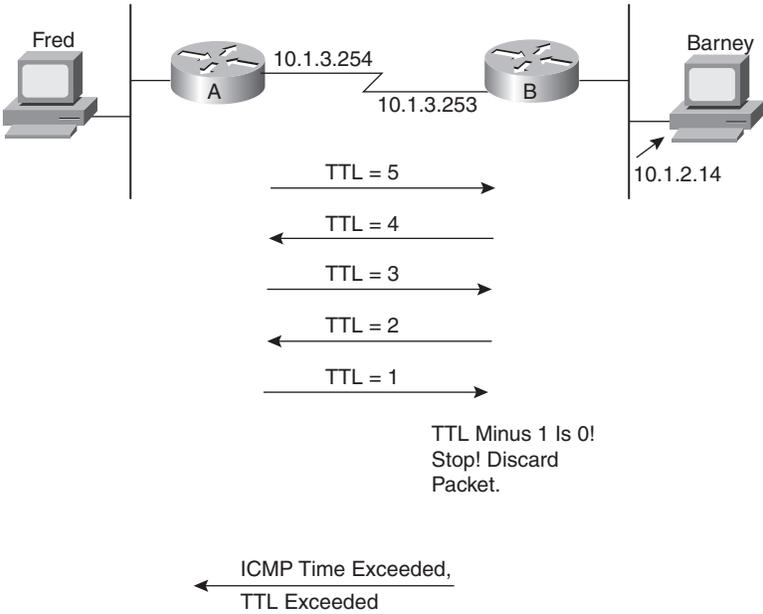


Figure 7-3 TTL Decrement to 0



As you can see in the figure, the router that discards the packet also sends an ICMP Time Exceeded message, with a Code field of “time exceeded” to the host that sent the packet. That way, the sender knows that the packet was not delivered. Getting a Time Exceeded

message can also help you when you troubleshoot a network. Hopefully, you do not get too many of these; otherwise, you have routing problems.

### The traceroute Command

The **ping** command is a powerful troubleshooting tool that can be used to answer the question “Does the route from here to there work?” The **traceroute** command provides an arguably better troubleshooting tool because not only can it determine if the route works, but it can supply the IP address of each router in the route. If the route is not working, **traceroute** can identify the best places to start troubleshooting the problem.

The IOS **traceroute** command uses the Time Exceeded message and the IP TTL field to identify each successive router in a route. The **traceroute** command sends a set of messages with increasing TTL values, starting with 1. The **traceroute** command expects these messages to be discarded when routers decrement the TTL to 0, returning Time Exceeded messages to the **traceroute** command. The source IP addresses of the Time Exceeded messages identify the routers that discarded the messages, which can then be displayed by the **traceroute** command.

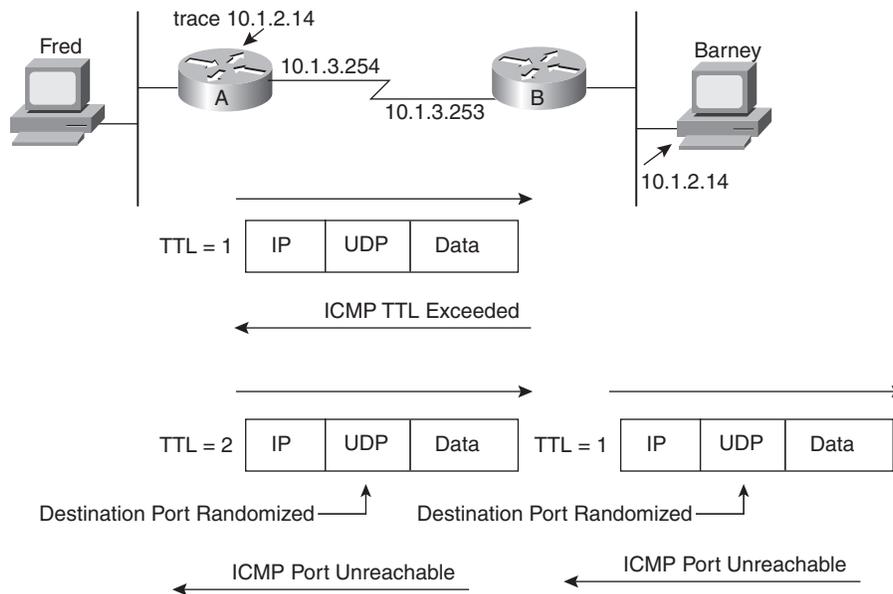
To see how this command works, consider the first set of packets (three packets by default) sent by the **traceroute** command. The packets are IP packets, with a UDP transport layer, and with the TTL set to 1. When the packets arrive at the next router, the router decrements the TTL to 0 in each packet, discards the packet, and sends a Time Exceeded message back to the host that sent the discarded packet. The **traceroute** command looks at the first router’s source IP address in the received Time Exceeded packet.

Next, the **traceroute** command sends another set of three IP packets, this time with TTL = 2. The first router decrements TTL to 1 and forwards the packets, and the second router decrements the TTL to 0 and discards the packets. This second router sends Time Exceeded messages back to the router where the **traceroute** command was used, and the **traceroute** command now knows the second router in the route.

The **traceroute** command knows when the test packets arrive at the destination host because the host sends back an ICMP Port Unreachable message. The original packets sent by the IOS **traceroute** command use a destination UDP port number that is very unlikely to be used on the destination host, so as soon as the TTL is large enough to allow the packet to arrive at the destination host, the host notices that it does not have an application listening at that particular UDP port. So, the destination host returns a Port Unreachable message, which tells the **traceroute** command that the complete route has been found, and the command can stop.

Figure 7-4 shows an example, but with only one of the three messages at each TTL setting (to reduce clutter). Router A uses the **traceroute** command to try to find the route to Barney. Example 7-1 shows this **traceroute** command on Router A, with debug messages from Router B, showing the three resulting Time Exceeded messages.

**Figure 7-4** Cisco IOS Software **traceroute** Command: Messages Generated



**Example 7-1** ICMP debug on Router B When Running the **traceroute** Command on Router A

```

RouterA#traceroute 10.1.2.14

Type escape sequence to abort.
Tracing the route to 10.1.2.14

 1 10.1.3.253 8 msec 4 msec 4 msec
 2 10.1.2.14 12 msec 8 msec 4 msec
RouterA#
! Moving to Router B now
! The following output occurs in reaction to the traceroute command on A
RouterB#debug ip icmp
RouterB#
ICMP: time exceeded (time to live) sent to 10.1.3.254 (dest was 10.1.2.14)
ICMP: time exceeded (time to live) sent to 10.1.3.254 (dest was 10.1.2.14)
ICMP: time exceeded (time to live) sent to 10.1.3.254 (dest was 10.1.2.14)
    
```

The **tracert** command lists the IP address of Router B in the first line and the IP address of the destination host in the second line. Note that it lists Router B's left-side IP address. B replies with the Time Exceeded message, using B's outgoing interface IP address as the source address in that packet. As a result, the **tracert** command lists that IP address. If the address is known to a DNS server, or if it's in Router A's hostname table, the command can list the hostname instead of the IP address.

Similar to the extended **ping** command as described in the section titled, "The Extended **ping** Command" in Chapter 4, the extended version of the **tracert** command does a much better job of simulating packets sent by end-user hosts, especially for testing reverse routes. For example, in Example 7-1, A's **tracert** command uses A's 10.1.3.254 IP address as the source address of sent packets, because A uses the interface with address 10.1.3.254 to send the packets generated by the **tracert** command. So, the **tracert** command in Example 7-1 tests the forward route toward 10.1.2.14 and the reverse route to 10.1.3.254. By using the extended **tracert** command, the command can be used to test a more appropriate reverse route, such as the route to the LAN subnet on the left side of Router A. Example 7-2, later in this chapter, shows an example of the extended **tracert** command.

**NOTE** The **tracert** command on Microsoft operating systems works much like the IOS **tracert** command. However, it is important to note that the Microsoft **tracert** command sends ICMP Echo Requests and does not use UDP. So, IP ACLs could cause the IOS **tracert** to fail while the Microsoft **tracert** worked, and vice versa.

## Troubleshooting the Packet Forwarding Process

Troubleshooting the IP routing process is one of the more complex tasks faced by network engineers. As usual, using a structured approach can help. Chapter 4 in particular, as well as Chapters 5 and 6, have already explained a lot about the first major part of the troubleshooting process—namely, what should happen in a network. This section focuses on the second major step: problem isolation. (For a more general reference on troubleshooting techniques, refer to Chapter 3, "Troubleshooting LAN Switching.")

**NOTE** This chapter defers any detailed troubleshooting of routing protocols until Chapter 11, "Troubleshooting Routing Protocols."

### Isolating IP Routing Problems Related to Hosts

The troubleshooting process outlined in this chapter separates the troubleshooting steps—one part for the hosts, and one part for the routers. Essentially, for any problem in which two hosts cannot communicate, the first part of this troubleshooting process examines the

issues that might impact each host's ability to send packets to and from its respective default gateway. The second part isolates problems related to how routers forward packets.

The following list outlines the troubleshooting steps focused on testing the host's connectivity to the first router:

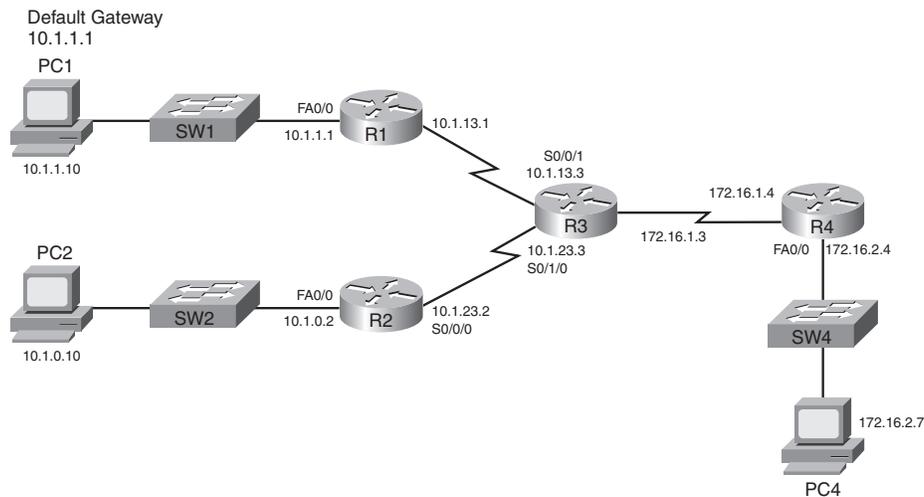
**Step 1** Check the host's ability to send packets inside its own subnet. Either ping the host's default gateway IP address from the host, or ping the host's IP address from the default gateway. If the ping fails, do the following:



- a. Ensure that the router's interface used at the default gateway is in an "up and up" state.
- b. Check the source host's IP address and mask setting as compared to the router's interface used as the default gateway. Ensure that both agree as to the subnet number and mask, and therefore agree to the range of valid addresses in the subnet.
- c. If the router uses VLAN trunking, solve any trunk configuration issues, ensuring that the router is configured to support the same VLAN in which the host resides.
- d. If the other steps do not lead to a solution, investigate Layer 1/2 problems with the LAN, as covered in Chapter 3. For example, look for an undefined VLAN.

**Step 2** Verify the default gateway setting on the host by pinging one of the default router's other interface IP addresses. Or, from the default router, use an extended ping of the host's IP address with a source address from another of the router's interfaces.

For example, in Figure 7-5, the problem symptoms may be that PC1 cannot browse the web server at PC4. To test PC1's ability to send packets over its local subnet, PC1 could use the **ping 10.1.1.1** command to test connectivity to the default router in its same subnet. Or the engineer could simply **ping 10.1.1.10** from R1 (Step 1). Either location for the **ping** works fine, because both ping locations require that a packet be sent in each direction. If the **ping** fails, further problem isolation should uncover the two specific problem areas listed in Steps 1A, 1B, and 1C. If not, the problem is likely to be some Layer 1 or 2 problem, as discussed in Chapter 3.

**Figure 7-5** *Sample Network for Troubleshooting Scenarios*

Step 2 stresses an often-overlooked troubleshooting concept to verify that the default gateway setting is working. Neither **ping** option listed in Step 1 requires the host to use its default gateway setting, because the source and destination address in each packet are in the same subnet. Step 2 forces the host to send a packet to an IP address in another subnet, thereby testing the host's default gateway setting. Also, by pinging an IP address on the default gateway (router), instead of some faraway host IP address, this step removes much of the IP routing complexity from the test. Instead, the focus is on whether the host's default gateway setting works. For example, in Figure 7-5, a **ping 10.1.13.1** command on PC1 forces PC1 to use its default gateway setting because 10.1.13.1 is not in PC1's subnet (10.1.1.0/24). But the IP address is on router R1, which removes most of the rest of the network as being a possible cause if the ping fails.

### Isolating IP Routing Problems Related to Routers

When the host problem isolation process is complete, and the pings all work, on both the sending and receiving hosts, any remaining IP routing issues should be between the first and last router in both the forward and reverse route between the two hosts. The following list picks up the troubleshooting process with the source host's default gateway/router, relying on the **tracert** command on the router. (Note that the host's equivalent command, such as **tracert** on Microsoft operating systems, can also be used.)

**NOTE** Although the following list may be useful for reference, it is rather long. Do not get bogged down in the details, but do read the examples of its use that follow this list; that should clarify many of the steps. As usual, you do not need to memorize any troubleshooting processes listed here. They are meant as learning tools to help you build your skills.

- Step 3** Test connectivity to the destination host by using the extended **tracert** command on the host's default gateway, using the router's interface attached to the source host for the source IP address of the packets. If the command successfully completes:
- a. No routing problems exist in the forward route or reverse route directions.
  - b. If the end-user traffic still does not work (even though the **tracert** worked), troubleshoot any ACLs on each interface on each router in the route, in both directions.



- Step 4** If the **tracert** command in Step 3 does not complete, test the *forward route* as follows:
- a. **telnet** to the last traced router (the last router listed in the **tracert** command).
  - b. Find that router's route that matches the destination IP address that was used in the original **tracert** command (**show ip route, show ip route ip-address**).
  - c. If no matching route is found, investigate why the expected route is missing. Typically it's either a routing protocol issue or a static route misconfiguration. It could also be related to a missing connected route.
  - d. If a matching route is found, and the route is a default route, confirm that it will be used based on the setting for the **ip classless/no ip classless** commands.
  - e. If a matching route is found, **ping** the next-hop IP address listed in the route. Or, if the route is a connected route, **ping** the true destination IP address.
    - If the **ping** fails, investigate Layer 2 problems between this router and the IP address that was pinged, and investigate possible ACL problems.
    - If the **ping** works, investigate ACL issues.

- f. If a matching route is found, and no other problems are found, confirm that the route is not errantly pointing in the wrong direction.

**Step 5** If Step 4 does not identify a problem in the forward route, test the *reverse route*:

- a. If the forward route on the last traced router refers to another router as the next-hop router, repeat the substeps of Step 3 from that next-hop router. Analyze the reverse route—the route to reach the source IP address used by the failed **tracert** command.
- b. If the forward route on the last traced router refers to a connected subnet, check the destination host's IP settings. In particular, confirm the settings for the IP address, mask, and default gateway.

For example, if PC1 cannot communicate with PC4 in Figure 7-5, and the hosts can both communicate through their respective default gateways, Step 3 of the router-oriented problem isolation process could start with a **tracert 172.16.2.7**, using R1's Fa0/0 IP address (10.1.1.1) as the source IP address. If that **tracert** command lists 10.1.13.3 as the last IP address in the command output, rather than completing, you would then start Step 4, which examines R3's forward route toward 172.16.2.7. If the analysis at Step 4 does not uncover the problem, Step 5 would then move on to the next-hop router, R4 in this case, and examine R4's reverse route—its route back to the original source address of 10.1.1.1.

Next, two separate scenarios show how to use these troubleshooting steps to isolate some sample problems.

### Troubleshooting Scenario 1: Forward Route Problem

This first example of the router troubleshooting process uses the same internetwork shown in Figure 7-5. In this case, PC1 cannot use a web browser to connect to the web service running on PC4. After further investigation, PC1 cannot ping 172.16.2.7 (PC4). Example 7-2 shows the commands used on R1 and R4 for the host-oriented Steps 1 and 2, as well as a beginning of the router-oriented Step 3.

#### Example 7-2 Troubleshooting Scenario 1: Steps 1 and 2 and Part of Step 3

```
R1#ping 10.1.1.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.10, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R1#ping
Protocol [ip]:
Target IP address: 10.1.1.10
Repeat count [5]:
```

**Example 7-2** *Troubleshooting Scenario 1: Steps 1 and 2 and Part of Step 3 (Continued)*

```
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.1.13.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.10, timeout is 2 seconds:
Packet sent with a source address of 10.1.13.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R1#
```

---

```
! Now moving to R4 to repeat the test
R4#ping 172.16.2.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R4#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	172.16.2.4	YES	manual	administratively down	down
FastEthernet0/1	172.16.1.4	YES	manual	up	up
Serial0/0/0	unassigned	YES	unset	administratively down	down
Serial0/0/1	unassigned	YES	unset	administratively down	down
Serial0/1/0	unassigned	YES	unset	administratively down	down

The standard and extended pings on R1 at the beginning of the example essentially perform Steps 1 and 2, the host-oriented steps, to confirm that PC1 seems to be working well. However, the example next shows that R4 cannot reach PC4 because R4’s LAN interface has been shut down, as shown at the end of the example. Although this scenario may seem a bit simple, it provides a good starting point for troubleshooting a problem.

To get a fuller view of the troubleshooting process, next consider this same scenario, with the same root problem, but now you do not have access to router R4. So, you can only perform Steps 1 and 2 for PC1, which work, but you cannot do those same steps for PC4 from R4. So, Example 7-3 moves on to Steps 3 and 4. The beginning of the example shows Step 3, where R1 uses **traceroute 172.16.2.7**, with a source IP address of 10.1.1.1. This

command does not complete, referencing 10.1.13.3 (R3) as the last router. Step 4 proceeds by looking at how R3 then routes packets destined for 172.16.2.7.

**Example 7-3** *Troubleshooting Scenario 1: Step 4*

```

R1#tracert
Protocol [ip]:
Target IP address: 172.16.2.7
Source address: 10.1.1.1
Numeric display [n]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 172.16.2.7

 1 10.1.13.3 0 msec 4 msec 0 msec
 2 10.1.13.3 !H * !H
! Note above that the command did stop by itself, but it does not list the
! destination host 172.16.2.7
R3#show ip route 172.16.2.7
% Subnet not in table
R3#show ip route
Codes: C - connected, S - static, 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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

172.16.0.0/24 is subnetted, 1 subnets
C    172.16.1.0 is directly connected, FastEthernet0/0
10.0.0.0/24 is subnetted, 4 subnets
C    10.1.13.0 is directly connected, Serial0/0/1
R    10.1.1.0 [120/1] via 10.1.13.1, 00:00:04, Serial0/0/1
R    10.1.0.0 [120/1] via 10.1.23.2, 00:00:01, Serial0/1/0
C    10.1.23.0 is directly connected, Serial0/1/0

```

The extended **tracert** command at the beginning of the example shows output identifying R3 (10.1.13.3) as the last listed device in the command output (Step 3). Step 4

then proceeds with an examination of the forward route on R3 toward IP address 172.16.2.7. The **show ip route 172.16.2.7** command gets right to the point. The message "subnet not in table" means that R3 does not have a route matching destination address 172.16.2.7. If the question does not supply access to a simulator, only the output of the **show ip route** command, you would need to examine the routes to determine that none of them refer to a range of addresses that includes 172.16.2.7.

Any time the problem isolation process points to a missing route, the next step is to determine how the router should have learned about the route. In this case, R3 should have used RIP-2 to learn the route. So, the next steps would be to troubleshoot any problems with the dynamic routing protocol.

The root cause of this problem has not changed—R4 has shut down its Fa0/0 interface—but the symptoms are somewhat interesting. Because the interface is shut down, R4 does not advertise a route for subnet 172.16.2.0/24 to R3. However, R3 advertises an autosummarized route to network 172.16.0.0/16 to both R1 and R2, so both R1 and R2, because of RIP-2's default autosummary setting, can forward packets destined for 172.16.2.7 to R3. As a result, the **traceroute** command on R1 can forward packets to R3.

**Troubleshooting Scenario 2: Reverse Route Problem**

This next example uses the same network diagram as shown in Figure 7-5, with all the information shown in the figure still being true. However, the details mentioned in the previous section may have changed—particularly the problem that exists to make the example more interesting. So, approach this second problem only relying on the figure as being true.

In this scenario, PC1 again cannot ping 172.16.2.7 (PC4). The host default gateway checks suggested in Steps 1 and 2 again work for PC1, but the tests cannot be performed for the reverse direction, because the engineer cannot access PC4 or router R4. So, Example 7-4 picks up the suggested troubleshooting process at Step 3, showing the result of the extended **traceroute** command on R1. Note that the command does not even list R3's 10.1.13.3 IP address in this case. So, the rest of Example 7-4 shows the investigations into the specific substeps of Step 4.

**Example 7-4** *Troubleshooting Scenario 2: Steps 3 and 4*

```
R1#traceroute ip 172.16.2.7 source fa0/0
Type escape sequence to abort.
Tracing the route to 172.16.2.7
 1  *  *  *
 2  *  *  *
 3  *
```

*continues*

**Example 7-4** *Troubleshooting Scenario 2: Steps 3 and 4 (Continued)*

```

R1#show ip route 172.16.2.7
Routing entry for 172.16.0.0/16
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 10.1.13.3 on Serial0/1/0, 00:00:05 ago
  Routing Descriptor Blocks:
  * 10.1.13.3, from 10.1.13.3, 00:00:05 ago, via Serial0/1/0
    Route metric is 1, traffic share count is 1

R1#ping 10.1.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R1#show ip access-lists

! Switching to router R3 next
R3#show ip access-lists

R3#

```

The example starts by showing the Step 3 part of the process, with the **tracert** command only listing lines of asterisks. This means that the command did not successfully identify even the very next router in the route.

Next, moving on to Step 4, the following list outlines the substeps of Step 4 as applied to this example:

- Step 4a** The example had already begun with a Telnet into R1, so no extra work is required.
- Step 4b** The next command, **show ip route 172.16.2.7**, shows that R1 has a nondefault route for network 172.16.0.0, pointing to R3 (10.1.13.3) as the next hop.
- Step 4c** This step does not apply in this case, because a matching route was found in Step 4B.
- Step 4d** This step does not apply in this case, because the matching route is not a route to 0.0.0.0/0 (the default route).
- Step 4e** The next listed command, **ping 10.1.13.3**, tests R1's ability to send packets over the link to the next-hop router identified in Step 4B. The ping works.

**Step 4f** On both R1 and the next-hop router (R3), the **show ip access-lists** command confirms that neither router has any IP ACLs configured.

Because all the steps to examine the forward route passed, the process then moves on to Step 5. The original **traceroute** command in Example 7-4 used R1's Fa0/0 interface IP address, 10.1.1.1, as the source IP address. For Step 5, the process begins at R3 with an analysis of R3's reverse route to reach 10.1.1.1. Examine the output in Example 7-5, and look for any problems before reading the explanations following the example.

**Example 7-5** *Troubleshooting Scenario 2: Step 5*

```
! The next command shows the matched route, for subnet 10.1.1.0/26,
! with next-hop 10.1.23.2.
R3#show ip route 10.1.1.1
Routing entry for 10.1.1.0/26
  Known via "static", distance 1, metric 0
  Routing Descriptor Blocks:
    * 10.1.23.2
      Route metric is 0, traffic share count is 1

! The next command shows the overlapping subnets - 10.1.1.0/26 and 10.1.1.0/24.
R3#show ip route
Codes: C - connected, S - static, 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
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 2 subnets
C      172.16.1.0 is directly connected, FastEthernet0/0
R      172.16.2.0 [120/1] via 172.16.1.4, 00:00:18, FastEthernet0/0
    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C      10.1.13.0/24 is directly connected, Serial0/0/1
S      10.1.1.0/26 [1/0] via 10.1.23.2
R      10.1.1.0/24 [120/1] via 10.1.13.1, 00:00:10, Serial0/0/1
R      10.1.0.0/24 [120/1] via 10.1.23.2, 00:00:11, Serial0/1/0
C      10.1.23.0/24 is directly connected, Serial0/1/0
```

R3 has an incorrectly configured static route for subnet 10.1.1.0/26. This subnet includes the address range 10.1.1.0–10.1.1.63, which includes IP address 10.1.1.1. When R3 attempts to send a packet back to 10.1.1.1, R3 has two routes that match the destination address. But R3 picks the more specific (longer prefix) route for subnet 10.1.1.0/26. This route causes R3 to forward packets intended for 10.1.1.1 out R3's link to R2, instead of to R1.

Although you cannot necessarily determine the true intent of this static route, this process has identified the root cause—the static route to 10.1.1.0/26 on R3. If the LAN off R1 should include all addresses between 10.1.1.0 and 10.1.1.255, the static route should just be deleted.

### **An Alternative Problem Isolation Process for Steps 3, 4, and 5**

The router-oriented steps of the IP routing problem isolation process depend on the **tracert** command, relying on this command's ability to identify on which router the router-oriented troubleshooting should begin. As an alternative, the **ping** and **telnet** commands can be used. However, because these commands cannot quickly identify the most likely routers on which the problem exists, using **ping** and **telnet** requires that you perform a set of tasks on the first router (the host's default gateway/router) in a route, and then the next router, and the next, and so on, until the problem is identified.

So, just to be complete, note that you can do the same specific subtasks as already explained in Steps 4 and 5, but when using **ping**, just repeat the steps at each successive router. For example, to apply this revised process to the first of the two just-completed scenarios, the process would begin with router R1, PC1's default router. In the first scenario, R1 did not have any forward route issues for forwarding packets to 172.16.2.7 (PC4), and R1 had no reverse route issues and no ACLs. This new alternative process would then suggest moving on to the next router (R3). In this example, R3's forward route problem—not having a route that matches destination address 172.16.2.7—would be found.

## **Troubleshooting Tools and Tips**

The second half of this chapter covers a wide variety of troubleshooting tools and tips that can be helpful when you're troubleshooting real networks. Some of the information in this section may apply directly to the CCNA exams. Other parts of this section will be indirectly useful for the exams. The information may help you learn as you work with networks in your job, making you better prepared for the unique scenarios on the exams.

### **Host Routing Tools and Perspectives**

This section covers two short topics related to how hosts process IP packets. The first topic lists several tips for troubleshooting hosts. The second topic reviews information covered in *CCENT/CCNA ICND1 Official Exam Certification Guide* on how a LAN switch's IP configuration works like a host.

#### **Host Troubleshooting Tips**

When you're trying to isolate the cause of networking problems, the tips in Table 7-4 may help you more quickly find problems related to hosts. The tips are organized by typical symptoms, along with common root causes. Note that the table does not list all possible causes, just the more common ones.

**Table 7-4** *Common Host Problem Symptoms and Typical Reasons*

Symptom	Common Root Cause
The host can send packets to hosts in the same subnet, but not to other subnets.	The host does not have a default gateway configured, or the default gateway IP address is incorrect.
The host can send packets to hosts in the same subnet, but not to other subnets.	The host's default gateway is in a different subnet than the host's IP address (according to the host's perception of the subnet).
Some hosts in a subnet can communicate with hosts in other subnets, but others cannot.	This may be caused by the default gateway (router) using a different mask than the hosts. This may result in the router's connected route not including some of the hosts on the LAN.
Some hosts on the same VLAN can send packets to each other, but others cannot.	The hosts may not be using the same mask.

When troubleshooting networking problems in real life, it's helpful to get used to thinking about the symptoms, because that's where the problem isolation process typically begins. However, for the exams, most host communication problems are caused by just a handful of issues:

- Step 1** Check all hosts and routers that should be in the same subnet to ensure that they all use the same mask and that their addresses are indeed all in the same subnet.
- Step 2** Compare each host's default gateway setting with the router's configuration to ensure that it is the right IP address.
- Step 3** If the first two items are correct, next look at Layer 1/2 issues, as covered in Chapters 1 through 3.



**LAN Switch IP Support**

Ethernet switches do not need to know anything about Layer 3 to perform their basic Layer 2 function of forwarding Ethernet frames. However, to support several important features, such as the ability to telnet and SSH to the switch to troubleshoot problems, LAN switches need an IP address.

Switches act like hosts when it comes to IP configuration. As compared to a PC, a Cisco switch does not use a NIC. Instead, it uses an internal virtual interface associated with VLAN 1 that essentially gives the switch itself an interface in VLAN 1. Then, the same kinds of items that can be configured on a host for IP can be configured on this VLAN interface: IP address, mask, and default gateway. DNS server IP addresses can also be configured.

The following list repeats the LAN switch IP configuration checklist from *CCENT/CCNA ICND1 Official Exam Certification Guide*. Following the list, Example 7-6 shows the IP address configuration for switch SW1 in Figure 7-5 from earlier in the chapter.



- Step 1** Enter VLAN 1 configuration mode using the **interface vlan 1** global configuration command (from any config mode).
- Step 2** Assign an IP address and mask using the **ip address ip-address mask** interface subcommand.
- Step 3** Enable the VLAN 1 interface using the **no shutdown** interface subcommand.
- Step 4** Add the **ip default-gateway ip-address** global command to configure the default gateway.

#### Example 7-6 Switch Static IP Address Configuration

```
SW1#configure terminal
SW1(config)#interface vlan 1
SW1(config-if)#ip address 10.1.1.200 255.255.255.0
SW1(config-if)#no shutdown
00:25:07: %LINK-3-UPDOWN: Interface Vlan1, changed state to up
00:25:08: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed
state to up
SW1(config-if)#exit
SW1(config)#ip default-gateway 10.1.1.1
```

**NOTE** The VLAN interface on a switch stays in an administratively down state until the user issues the **no shutdown** command; the switch cannot send IP packets until the VLAN 1 interface is up.

A common oversight when configuring or troubleshooting IP connectivity problems to LAN switches relates to VLAN trunking. Cisco generally suggests that you avoid putting end-user devices into VLAN 1, but the switch IP address may well be configured in VLAN 1. To support the ability for the switch to send and receive packets to hosts in different subnets, thereby supporting Telnet into the switch from those end-user subnets, the router's trunking configuration must include configuration for VLAN 1 as well as the end-user VLANs.

#### show ip route Reference

The **show ip route** command plays a huge role in troubleshooting IP routing and IP routing protocol problems. Many chapters in this book and in the ICND1 book mention various facts about this command. This section pulls the concepts together in one place for easier reference and study.

Figure 7-6 shows the output of the **show ip route** command from back in Example 7-3. The figure numbers various parts of the command output for easier reference, with Table 7-5 describing the output noted by each number.

**Figure 7-6** **show ip route** Command Output Reference

```

    ① ② ③
    ④ ⑤ ⑥ ⑦ ⑧ ⑨ ⑩
    C 10.0.0.0/24 is subnetted, 4 subnets
    R 10.1.13.0 is directly connected, Serial0/0/1
    C 10.1.1.0 [120/1] via 10.1.13.1, 00:00:04, Serial0/0/1
    R 10.1.23.0 is directly connected, Serial0/1/0
    R 10.1.0.0 [120/1] via 10.1.23.2, 00:00:01, Serial0/1/0
  
```

**Table 7-5** Descriptions of the **show ip route** Command Output

Item Number	Item	Value in the Figure	Description
1	Classful network	10.0.0.0	The routing table is organized by classful network. This line is the heading line for classful network 10.0.0.0.
2	Prefix length	/24	When this router knows only one subnet mask for all subnets of the network, this location lists that one mask, by default in prefix notation.
3	Number of subnets	4 subnets	Lists the number of routes for subnets of the classful network known to this router.
4	Legend code	R, C	A short code that identifies the source of the routing information. R is for RIP, and C is for connected. The figure omits the legend text at the top of the <b>show ip route</b> command output, but it can be seen in Example 7-3.
5	Subnet number	10.1.0.0	The subnet number of this particular route.
6	Administrative distance	120	If a router learns routes for the listed subnet from more than one source of routing information, the router uses the source with the lowest AD.
7	Metric	1	The metric for this route.
8	Next-hop router	10.1.23.2	For packets matching this route, the IP address of the next router to which the packet should be forwarded.
9	Timer	00:00:01	Time since this route was learned in a routing update.
10	Outgoing interface	Serial0/1/0	For packets matching this route, the interface out which the packet should be forwarded.

The output of the command differs slightly when VLSM is used. The figure shows an example in which VLSM is not used in network 10.0.0.0, with mask /24 used for all subnets of that network. So, IOS lists the mask once, in the heading line (/24 in this case). If VLSM were in use, the heading line would simply note that the network is variably subnetted, and each route would list the mask. For an example, see Example 5-1 in Chapter 5, “VLSM and Route Summarization.”

### Interface Status

One of the steps in the IP routing troubleshooting process described earlier, in the “Troubleshooting the Packet Forwarding Process” section, says to check the interface status, ensuring that the required interface is working. For a router interface to be working, the two interface status codes must both be listed as “up,” with engineers usually saying the interface is “up and up.”

This chapter does not explain the troubleshooting steps for router interfaces, simply assuming that each interface is indeed in an up/up state. Chapter 12’s section titled “Troubleshooting Serial Links” covers many of the details for troubleshooting router interfaces. For router LAN interfaces connected to a LAN switch, the main items to check on routers are that the router and switch match each other’s duplex and speed settings, and that if trunking is configured, both the router and switch have been manually configured for trunking, because routers do not dynamically negotiate LAN trunking.

### VLSM Issues

This section examines several issues when using VLSM:

- Recognizing whether VLSM is used and, if so, which routing protocols can be used
- Understanding the conditions in which routers can allow the misconfiguration of overlapping VLSM subnets
- Understanding the outward symptoms that can occur when overlapping VLSM subnets exist

### Recognizing When VLSM Is Used

One common oversight when troubleshooting a problem in an unfamiliar internetwork is failing to recognize whether VLSM is used. As defined in Chapter 5, an internetwork uses VLSM when multiple subnet masks are used for different subnets of a *single classful network*. For example, if in one internetwork all subnets of network 10.0.0.0 use a 255.255.240.0 mask, and all subnets of network 172.16.0.0 use a 255.255.255.0 mask, the design does not use VLSM. If multiple masks were used for subnets of network 10.0.0.0, VLSM would be in use.

The follow-on concept is that only classless routing protocols (RIP-2, EIGRP, OSPF) can support VLSM; classful routing protocols (RIP-1, IGRP) cannot. So, a quick determination of whether VLSM is actually used can then tell you whether a classless routing protocol is required. Note that the routing protocol does not require any special configuration to support VLSM. It is just a feature of the routing protocol.

### Configuring Overlapping VLSM Subnets

IP subnetting rules require that the address ranges in the subnets used in an internetwork should not overlap. IOS can recognize when a new **ip address** command creates an overlapping subnet, but only in some cases. This section examines the conditions under which overlapping subnets can be configured, beginning with the following general statements about when the overlaps cannot and can be configured:

- **Preventing the overlap:** IOS detects the overlap when the **ip address** command implies an overlap with another **ip address** command *on the same router*. If the interface being configured is up/up, IOS rejects the **ip address** command. If not, IOS accepts the **ip address** command, but IOS will never bring up the interface.
- **Allowing the overlap:** IOS cannot detect an overlap when an **ip address** command overlaps with an **ip address** command on another router.



The router shown in Example 7-7 prevents the configuration of an overlapping VLSM subnet. The example shows router R3 configuring Fa0/0 with IP address 172.16.5.1/24, and Fa0/1 with 172.16.5.193/26. The ranges of addresses in each subnet are:

Subnet 172.16.5.0/24: 172.16.5.1– 172.16.5.254  
Subnet 172.16.5.192/26: 172.16.5.193–172.16.5.254

**Example 7-7** *Single Router Rejects Overlapped Subnets*

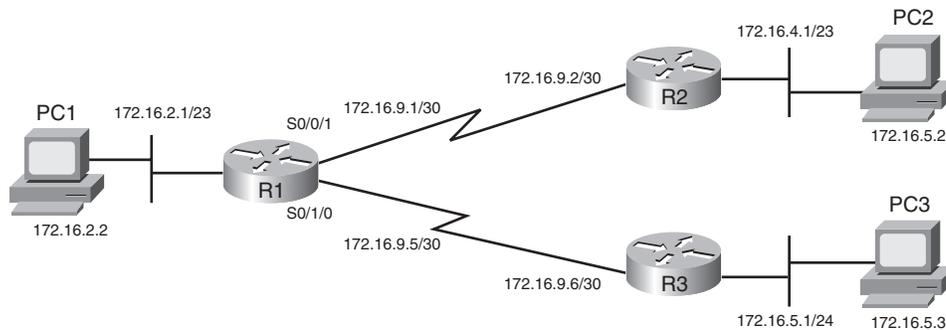
```

R3#configure terminal
R3(config)#interface Fa0/0
R3(config-if)#ip address 172.16.5.1 255.255.255.0
R3(config-if)#interface Fa0/1
R3(config-if)#ip address 172.16.5.193 255.255.255.192
% 172.16.5.192 overlaps with FastEthernet0/0
R3(config-if)#

```

IOS knows that it is illegal to overlap the ranges of addresses implied by a subnet. In this case, because both subnets would be connected subnets, this single router knows that these two subnets should not coexist, because that would break subnetting rules, so IOS rejects the second command.

However, it is possible to configure overlapping subnets if they are connected to different routers. Figure 7-7 shows a figure very similar to Figure 5-2 in Chapter 5—used in that chapter to explain the problem of overlapping subnets. Example 7-8 shows the configuration of the two overlapping subnets on R2 and R3, with the resulting routing table on R2.

**Figure 7-7** *Internetwork That Allows the Configuration of Overlapped Subnets***Example 7-8** *Two Routers Accept Overlapped Subnets*

```

R2#configure terminal
R2(config)#interface Fa0/0
R2(config-if)#ip address 172.16.4.1 255.255.254.0
R3#configure terminal
R3(config)#interface Fa0/0
R3(config-if)# ip address 172.16.5.1 255.255.255.0

```

For the exams, keep in mind that overlapped subnets can be configured if the subnets do not connect to the same router. So, if a question asks you to pick a new subnet number and configure an interface to be in that subnet, the router's acceptance of your **ip address** command does not necessarily tell you that you did the math correctly.

The next topic explains some of the problem symptoms you might see if such an overlap exists.

### Symptoms with Overlapping Subnets

**NOTE** Although this section is included for the sake of completeness, the types of problems described here may well be beyond the scope of the CCNA exams.

The outward problem symptoms differ depending on whether the address in question is in the overlapped portion of the subnets and if multiple hosts are attempting to use the exact same IP address. The addresses in the nonoverlapped parts of the subnet typically work fine, whereas those in the overlapped area may or may not work at all. For example, continuing with the overlapped subnets shown in Figure 7-6, subnets 172.16.4.0/23 and 172.16.5.0/24 overlap—specifically, addresses 172.16.5.0–172.16.5.255. Hosts in the nonoverlapped range of 172.16.4.0–172.16.4.255 probably work fine.

For the addresses in the overlapped address range, in many cases, hosts in the smaller of the two overlapped subnets work fine, but hosts in the larger of the two subnets do not. To see why, consider the case in which PC1 in Figure 7-7 tries to ping both 172.16.5.2 (PC2, off R2) and 172.16.5.3 (PC3, off R3). (For the sake of this example, assume that PC2's and PC3's IP addresses are not duplicated in the opposite overlapped subnet.) As you can see from the routing tables on R1 and R3 and the **traceroute 172.16.5.2** command in Example 7-9, the packet sent by PC1 to PC2 would actually be delivered from R1 to R3, and then onto R3's LAN.

#### Example 7-9 Two Routers Accept Overlapped Subnets

```
! R1's route to reach 172.16.5.2, off R2, points to R3
R1#show ip route 172.16.5.2
Routing entry for 172.16.5.0/24
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 172.16.9.6 on Serial0/1/0, 00:00:25 ago
  Routing Descriptor Blocks:
    * 172.16.9.6, from 172.16.9.6, 00:00:25 ago, via Serial0/1/0
      Route metric is 1, traffic share count is 1
! R1's route to reach 172.16.5.3, off R3, points to R3
R1#show ip route 172.16.5.3
Routing entry for 172.16.5.0/24
  Known via "rip", distance 120, metric 1
```

*continues*

**Example 7-9** *Two Routers Accept Overlapped Subnets (Continued)*

```

Redistributing via rip
Last update from 172.16.9.6 on Serial0/1/0, 00:00:01 ago
Routing Descriptor Blocks:
* 172.16.9.6, from 172.16.9.6, 00:00:01 ago, via Serial0/1/0
  Route metric is 1, traffic share count is 1

! The traceroute to PC2 shows R3, not R2, as the first router, so the packet never
! reaches PC2, and the command never completes until stopped by the user.
R1#traceroute 172.16.5.2

Type escape sequence to abort.
Tracing the route to 172.16.5.2

  1 172.16.9.6 4 msec 0 msec 4 msec
  2 * * *
  3 * * *
  4

R1#traceroute 172.16.5.3

Type escape sequence to abort.
Tracing the route to 172.16.5.3

  1 172.16.9.6 0 msec 4 msec 0 msec
  2 172.16.5.3 4 msec * 0 msec

```

The example shows that R1 forwards packets to hosts 172.16.5.2 (PC2) and 172.16.5.3 (PC3) by sending them to R3 next. R3 then tries to send them onto R3's LAN subnet, which works well for PC3 but not so well for PC2. So, PC3, in the smaller of the two overlapped subnets, works fine, whereas PC2, in the larger of the two overlapped subnets, does not.

The symptoms can get even worse when addresses are duplicated. For example, imagine that PC22 has been added to R2's LAN subnet, with IP address 172.16.5.3 duplicating PC3's IP address. Now when the PC22 user calls to say that his PC cannot communicate with other devices, the network support person uses a **ping 172.16.5.3** to test the problem—and the ping works! The ping works to the wrong instance of 172.16.5.3, but it works. So, the symptoms may be particularly difficult to track down.

Another difficulty with overlapped VLSM subnets is that the problem may not show up for a while. In this same example, imagine that all addresses in both subnets were to be assigned by a DHCP server, beginning with the smallest IP addresses. For the first six months, the server assigned only IP addresses that began with 172.16.4.x on the R2 LAN subnet. Finally, enough hosts were installed on the R2 LAN to require the use of addresses that begin with 172.16.5, like PC2's address of 172.16.5.2 used in the preceding example.

Unfortunately, no one can send packets to those hosts. At first glance, the fact that the problem showed up long after the installation and configuration were complete may actually cloud the issue.

### VLSM Troubleshooting Summary

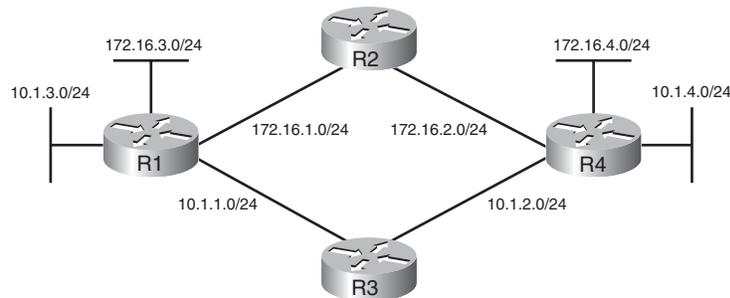
The following list summarizes the key troubleshooting points to consider when you're troubleshooting potential VLSM problems on the exams:

- Pay close attention to whether the design really uses VLSM. If it does, note whether a classless routing protocol is used.
- Be aware that overlapping subnets can indeed be configured.
- The outward problem symptoms may be that some hosts in a subnet work well, but others cannot send packets outside the local subnet.
- Use the **tracert** command to look for routes that direct packets to the wrong part of the network. This could be a result of the overlapped subnets.
- On the exams, you might see a question you think is related to VLSM and IP addresses. In that case, the best plan of attack may well be to analyze the math for each subnet and ensure that no overlaps exist, rather than troubleshooting using **ping** and **tracert**.



### Discontiguous Networks and Autosummary

Chapter 5 explained the concept of discontiguous networks, along with the solution: using a classless routing protocol with autosummarization disabled. This section examines one particular case in which a discontiguous network exists only part of the time. Figure 7-8 shows an internetwork with two classful networks: 10.0.0.0 and 172.16.0.0. The design shows two contiguous networks because a route consisting of only subnets of each network exists between all subnets of that network.

**Figure 7-8** Internetwork with (Currently) Contiguous Networks

In this figure, with all links up and working, using a routing protocol with autosummary enabled by default, all hosts can ping all other hosts. In this design, packets for network 172.16.0.0 flow over the high route, and packets for network 10.0.0.0 flow over the low route.

Unfortunately, a problem can occur later when one of the four links between routers fails. If any link between the routers fails, one of the two classful networks becomes discontinuous. For example, if the link between R3 and R4 fails, the route from R1 to R4 passes through subnets of network 172.16.0.0, so network 10.0.0.0 is discontinuous. Even with a classless routing protocol, but with autosummarization enabled, both R1 and R4 advertise a route for 10.0.0.0/8 to R2, and R2 sees two routes to all of network 10.0.0.0—one through R1, and another through R4. The solution, as always, is to use a classless routing protocol with autosummary disabled.

Although the design in Figure 7-8 may seem a bit contrived, it happens more often than you might think—particularly as companies are bought and sold. Both for real life and the exams, keep the concept of discontinuous networks in mind for normal working cases and for cases in which redundant links fail.

### Access List Troubleshooting Tips

Troubleshooting problems that are impacted by ACLs may well be one of the most difficult tasks for real networking jobs. One of the major difficulties is that the traditional troubleshooting tools such as **ping** and **traceroute** do not send packets that look like the packets matched by the variety of fields in extended ACLs. So, although a **ping** may work, the end-user host may not be able to get to the right application, or vice versa.

This section summarizes some tips for attacking ACL-related problems in real life and on the exams:



**Step 1** Determine on which interfaces ACLs are enabled, and in which direction (**show running-config**, **show ip interfaces**).

- Step 2** Determine which ACL statements are matched by test packets (**show access-lists**, **show ip access-lists**).
- Step 3** Analyze the ACLs to predict which packets should match the ACL, focusing on the following points:
- a. Remember that the ACL uses first-match logic.
  - b. Consider using the (possibly) faster math described in Chapter 6, “IP Access Control Lists,” which converts ACL address/wildcard mask pairs into address/subnet mask pairs that allow the use of the same math as subnetting.
  - c. Note the direction of the packet in relation to the server (going to the server, coming from the server). Make sure that the packets have particular values as either the source IP address and port, or as the destination IP address and port, when processed by the ACL enabled for a particular direction (in or out).
  - d. Remember that the **tcp** and **udp** keywords must be used if the command needs to check the port numbers. (See Table 6-5 in Chapter 6 for a list of popular TCP and UDP port numbers.)
  - e. Note that ICMP packets do not use UDP or TCP. ICMP is considered to be another protocol matchable with the **icmp** keyword (instead of **ip**, **tcp**, and **udp**).
  - f. Instead of using the implicit deny any at the end of each ACL, use an explicit configuration command to deny all traffic at the end of the ACL so that the **show** command counters increment when that action is taken.

Chapter 6 covered the background information behind the tips listed in Step 3. The remainder of this section focuses on commands available for you to investigate problems in the first two steps.

If a problem in forwarding IP packets is occurring, and existing ACLs may be impacting the problem, the first problem isolation step is to find the location and direction of the ACLs. The fastest way to do this is to look at the output of the **show running-config** command and to look for **ip access-group** commands under each interface. However, in some cases, enable mode access may not be allowed, and **show** commands are required. The only way to find the interfaces and direction for any IP ACLs is the **show ip interfaces** command, as shown in Example 7-10.

**Example 7-10** *Sample show ip interface Command*

```

R1>show ip interface s0/0/1
Serial0/0/1 is up, line protocol is up
  Internet address is 10.1.2.1/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.9
  Outgoing access list is not set
  Inbound access list is 102
! roughly 26 more lines omitted for brevity

```

Note that the command output lists whether an ACL is enabled, in both directions, and which ACL it is. The example shows an abbreviated version of the **show ip interface S0/0/1** command, which lists messages for just this one interface. The **show ip interface** command would list the same messages for every interface in the router.

Step 2 then says that the contents of the ACL must be found. Again, the most expedient way to look at the ACL is to use the **show running-config** command. If enable mode is not allowed, the **show access-lists** and **show ip access-lists** commands give the same output. The only difference is that if other non-IP ACLs have been configured, the **show access-lists** command lists the non-IP ACLs as well. The output provides the same details shown in the configuration commands, as well as a counter for the number of packets matching each line in the ACL. Example 7-11 shows an example.

**Example 7-11** *Sample show ip access-lists Command*

```

R1#show ip access-lists
Extended IP access list 102
  10 permit ip 10.1.2.0 0.0.0.255 10.1.1.0 0.0.0.255 (15 matches)

```

After the locations, directions, and configuration details of the various ACLs have been discovered in Steps 1 and 2, the hard part begins—interpreting what the ACL really does. Of particular interest is the last item in the troubleshooting tips list, item 3E. In the ACL shown in Example 7-11, some packets (15 so far) have matched the single configured **access-list** statement in ACL 102. However, some packets have probably been denied because of the implied deny all packets logic at the end of an ACL. By configuring the **access-list 102 deny ip any any** command at the end of the ACL, which explicitly matches all packets and discards them, the **show ip access-lists** command would then show the number of packets being denied at the end of the ACL. Cisco sometimes recommends adding the explicit deny all statement at the end of the ACL for easier troubleshooting.

## Exam Preparation Tasks

### Review All the Key Topics

Review the most important topics from this chapter, noted with the key topics icon. Table 7-6 lists these key topics and where each is discussed.



**Table 7-6** *Key Topics for Chapter 7*

Key Topic Element	Description	Page Number
Table 7-1	Popular ICMP messages and their purpose	271
Figure 7-3	Diagram of how the TTL IP header field and the ICMP Time Exceeded message work	275
Figure 7-4	Demonstration of how the <b>tracert</b> command uses the TTL field and Time Exceeded message	277
List	Two major steps and several substeps in a suggested host routing problem isolation process	279
List	Three major steps for problem isolation with IP routing in routers, with the list numbered as a continuation of the host routing problem isolation list	281
List	Three tips for general items to check when troubleshooting host connectivity problems	289
List	Configuration step list for LAN switch IP details	290
List	Conditions under which overlapping subnets can be configured, and when IOS can prevent this error	293
List	Summary of troubleshooting tips for questions in which VLSM may be causing a problem	297
List	Three steps for troubleshooting ACL problems, particularly when the configuration cannot be displayed	298-299

### Complete the Tables and Lists from Memory

Print a copy of Appendix J, “Memory Tables” (found on the CD), or at least the section for this chapter, and complete the tables and lists from memory. Appendix K, “Memory Tables Answer Key,” also on the CD, includes completed tables and lists for you to check your work.

### Definitions of Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

Forward route, reverse route