

This chapter covers the following exam topics:

Troubleshooting

- Identify and correct common network problems
- Troubleshoot and resolve interVLAN routing problems
  - Connectivity
  - Encapsulation
  - Subnet
  - Native VLAN
  - Port mode trunk status
- Troubleshoot and resolve routing issues
  - Routing is enabled
  - Routing table is correct
  - Correct path selection

CHAPTER 5

Troubleshooting IPv4 Routing Part II

Chapter 4, "Troubleshooting IPv4 Routing Part I," began the discussion of IPv4 troubleshooting, looking at the usual first steps when troubleshooting a problem. This chapter moves on to a later stage, when the problem has been isolated to a smaller part of the network, and to a smaller set of possible causes of the problem. The topics in this chapter get specific and look for those root causes: the causes of network problems that have specific solutions that, once a change is made, will solve the original problem.

This chapter breaks down the discussion based on the two major divisions in how packets are forwarded in an IPv4 internetwork. The first half of the chapter focuses on the root causes of problems between a host and its default router. The second half looks at the routers that forward the packet over the rest of a packet's journey, from the router acting as default router all the way to the destination host.

Note that in addition to Chapters 4 and 5, other chapters in this book discuss troubleshooting topics that help when troubleshooting IPv4 internetworks. In particular, Chapter 11, "Troubleshooting IPv4 Routing Protocols," discusses troubleshooting IPv4 routing protocols, namely Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP). Chapter 3, "Troubleshooting LAN Switching," discussed how to troubleshoot LAN issues. Some topics inside the chapters in Part IV explain how to troubleshoot WAN links. Finally, Chapter 16, "Troubleshooting IPv6 Routing," discusses how to apply these same IPv4 troubleshooting concepts to IPv6.

"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 100-101 Official Cert 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 100-101 Official Cert 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

Problems Between the Host and the Default Router

Imagine that you work as a customer support rep (CSR) fielding calls from users about problems. A user left a message stating that he couldn't connect to a server. You could not reach him when you called back, so you did a series of pings from that host's default router, using some of the problem isolation strategies described in Chapter 4. And at the end of those pings, you think the problem exists somewhere between the user's device and the default router—for instance, between router R1 and host A as shown in Figure 5-1.

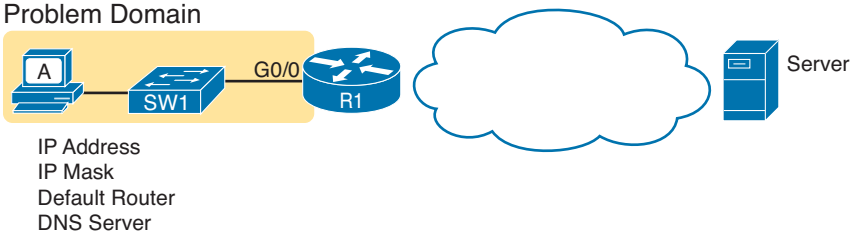


Figure 5-1 *Focus of the Discussions in This Section of the Chapter*

This first major section of the chapter focuses on problems that can occur on hosts, their default routers, and between the two. To begin, this section looks at the host itself, and its four IPv4 settings, as listed in the figure. Following that, the discussion moves to the default router, with focus on the LAN interface, and the settings that must work for the router to serve as a host's default router.

Root Causes Based on a Host's IPv4 Settings

A typical IPv4 host gets its four key IPv4 settings in one of two ways: either through static configuration or by using DHCP. In both cases, the settings can actually be incorrect. Clearly, any static settings can be set to a wrong number just through human error when typing the values. More surprising is the fact that the DHCP can set the wrong values: The DHCP process can work, but with incorrect values configured at the DHCP server, the host can actually learn some incorrect IPv4 settings.

This section first reviews the settings on the host, and what they should match, followed by a discussion of typical issues.

Ensure IPv4 Settings Correctly Match

Once an engineer thinks that a problem exists somewhere between a host and its default router, the engineer should review of the host's IPv4 settings versus the intended settings. That process begins by guiding the user through the GUI of the host operating system or by using command-line commands native to host operating systems, such as **ipconfig** and **ifconfig**. This process should uncover obvious issues, like completely missing parameters, or if using DHCP, the complete failure of DHCP to learn any of the IPv4 settings.

If the host has all its settings, the next step is to check the values to match them with the rest of the internetwork. The Domain Name System (DNS) server IP address—usually a list of at least two addresses— should match the DNS server addresses actually used in the internetwork. The rest of the settings should be compared to the correct LAN interface on the router that is used as this host's default router. Figure 5-2 collects all the pieces that should match, with some explanation to follow.

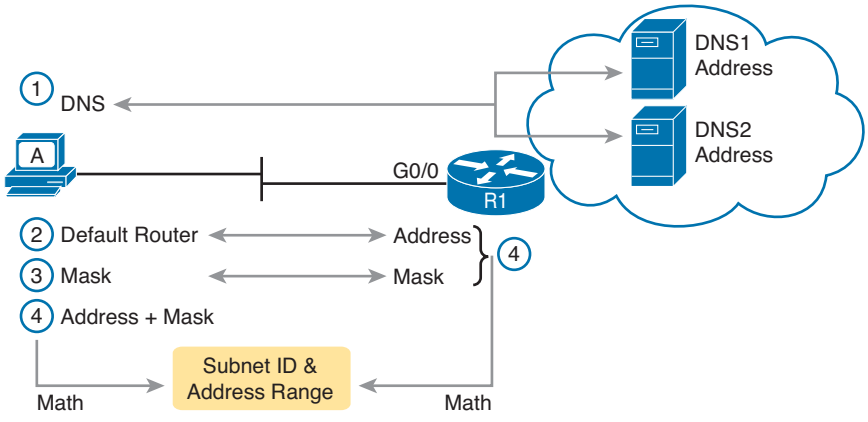


Figure 5-2 *Host IPv4 Settings Compared to What the Settings Should Match*

As numbered in the figure, these steps should be followed to check the host's IPv4 settings:

Step 1. Check the host's list of DNS server addresses against the actual addresses used by those servers.

Step 2. Check the host's default router setting against the router's LAN interface configuration, for the **ip address** command.

Step 3. Check the subnet mask used by the router and the host; if they use a different mask, the subnets will not exactly match, which will cause problems for some host addresses.

Step 4. The host and router should attach to the exact same subnet—same subnet ID and same range of IP addresses. So, use both the router's and host's IP address and mask, calculate the subnet ID and range of addresses, and confirm they are in the same subnet as the subnet implied by the address/mask of the router's **ip address** command.

If an IPv4 host configuration setting is missing, or simply wrong, checking these settings can quickly uncover the root cause. For instance, if you can log in to the router and do a **show interfaces G0/0** command, and then ask the user to issue an **ipconfig /all** (or similar) command and read the output to you, you can compare all the settings in Figure 5-2.

However, although checking the host settings is indeed very useful, some problems related to hosts are not so easy to spot. The next few topics walk through some example problems to show some symptoms that occur when some of these less obvious problems occur.

Mismatched Masks Impact Route to Reach Subnet

A host and its default router should agree about the range of addresses in the subnet. Sometimes, people are tempted to skip over this check, ignoring the mask either on the host or the router and assuming that the mask used on one device must be the same mask as on the other device. However, if the host and router have different subnet mask values, and therefore each calculates a different range of addresses in the subnet, problems happen.

To see one such example, consider the network in Figure 5-3. Host A has IP address/mask 10.1.1.9/24, with default router 10.1.1.150. Some quick math puts 10.1.1.150—the default router address—inside host A's subnet, right? Indeed it does, and it should. Host A's math for this subnet reveals subnet ID 10.1.1.0, with a range of addresses from 10.1.1.1 through 10.1.1.254, and subnet broadcast address 10.1.1.255.

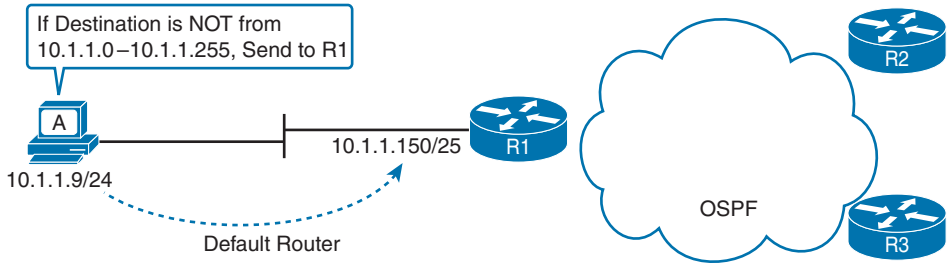


Figure 5-3 *Mismatched Subnet Calculations Appear Workable from Host Toward Network*

In this case, the host routing of packets, to destinations outside the subnet, works well. However, the reverse direction, from the rest of the network back toward the host, does not. A quick check of router R1's configuration reveals the IP address/mask as shown in Figure 5-3, which results in the connected route for subnet 10.1.1.128/25, as shown in Example 5-1.

Example 5-1 *R1's IP Address, Mask, Plus the Connected Subnet That Omits Host A's Address*

```
R1# show running-config interface g0/0
Building configuration...

Current configuration : 185 bytes
!
interface GigabitEthernet0/0
 description LAN at Site 1
 mac-address 0200.0101.0101
 ip address 10.1.1.150 255.255.255.128
 ip helper-address 10.1.2.130
 duplex auto
 speed auto
end
```

```
R1# show ip route connected
! Legend omitted for brevity

      10.0.0.0/8 is variably subnetted, 9 subnets, 4 masks
C        10.1.1.128/25 is directly connected, GigabitEthernet0/0
L        10.1.1.150/32 is directly connected, GigabitEthernet0/0
! Other routes omitted for brevity
```

Because of this particular mismatch, R1's view of the subnet puts host A (10.1.1.9) outside R1's view of the subnet (10.1.1.128/25, range 10.1.1.129 to 10.1.1.254). R1 adds a connected route for subnet 10.1.1.128/25 into R1's routing table, and even advertises this route (with OSPF in this case) to the other routers in the network, as seen in Figure 5-4. All the routers know how to route packets to subnet 10.1.1.128/25, but unfortunately, that route does not include host A's 10.1.1.9 IP address.

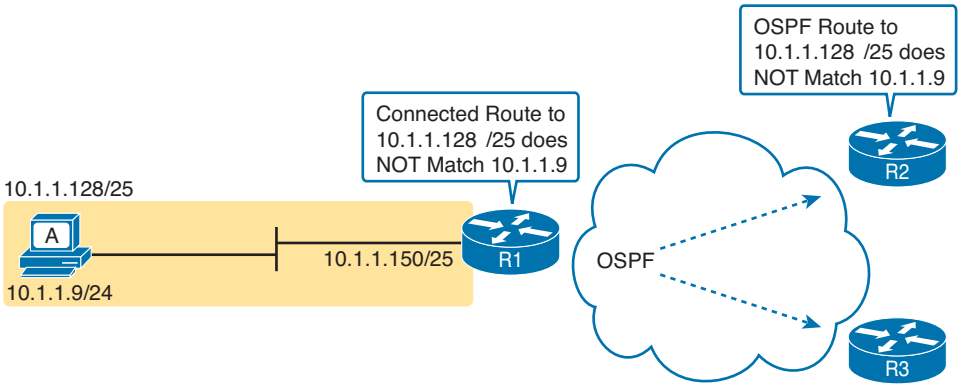


Figure 5-4 *Routers have No Route that Matches Host A's 10.1.1.9 Address*

Hosts should use the same subnet mask as the default router, and the two devices should agree as to what subnet exists on their common LAN. Otherwise, problems may exist immediately, as in this example, or they might not exist until other hosts are added later.

Typical Root Causes of DNS Problems

When a host lists the wrong IP addresses for the DNS servers, the symptoms are somewhat obvious: Any user actions that require name resolution fail. Assuming that the only problem is the incorrect DNS setting, any network testing with commands like **ping** and **traceroute** fails when using names, but it works when using IP addresses instead of names.

When a ping of another host's hostname fails, but a ping of that same host's IP address works, some problem exists with DNS. For example, imagine a user calls the help desk complaining that he cannot connect to Server1. The CSR issues a **ping server1** command from the CSR's own PC, which both works and identifies the IP address of Server1 as 1.1.1.1. Then the CSR asks the user to try two commands from the user's PC: both a **ping Server1** command (which fails), and a **ping 1.1.1.1** command (which works). Clearly, the DNS name resolution process on the user's PC is having some sort of problem.

This book does not go into much detail about how DNS truly works behind the scenes, but the following two root causes of DNS problems do fit within the scope of the CCENT and CCNA:

Key Topic

- An incorrect DNS server setting
- An IP connectivity problem between the user's host and the DNS server

Although the first problem may be more obvious, note that it can happen both with static settings on the host and with DHCP. If a host lists the wrong DNS server IP address, and the setting is static, just change the setting. If the wrong DNS server address is learned with DHCP, you need to examine the DHCP server configuration. (If using the IOS DHCP server feature, you make this setting with the **dns-server** *server-address* command in DHCP pool mode.)

The second bullet point brings up an important issue for troubleshooting any real-world networking problem. Most every real user application uses names, not addresses, and most hosts use DNS to resolve names. So, every connection to a new application involves two sets of packets: packets that flow between the host and the DNS server, and packets that flow between the host and the real server, as shown in Figure 5-5.

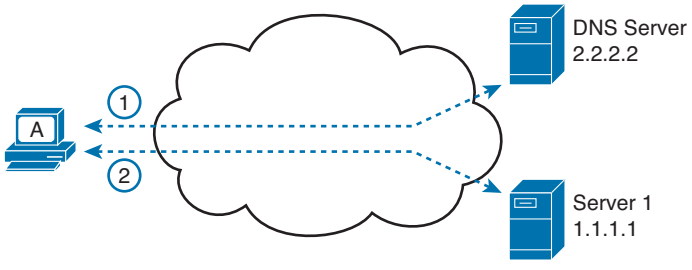


Figure 5-5 *DNS Name Resolution Packets Flow First; Then Packets to the Real Server*

Finally, before leaving the topic of name resolution, note that the router can be configured with the IP addresses of the DNS servers, so that router commands will attempt to resolve names. For instance, a user of the router command-line interface (CLI) could issue a command **ping server1** and rely on a DNS request to resolve server1 into its matching IP address. To configure a router to use a DNS for name resolution, the router needs the **ip name-server** *dns1-address dns2-address…* global command. It also needs the **ip domain-lookup** global command, which is enabled by default.

For troubleshooting, it can be helpful to set a router or switch DNS settings to match that of the local hosts. However, note that these settings have no impact on the user DNS requests.

NOTE On a practical note, IOS defaults with the **ip domain-lookup** command, but with no DNS IP address known. Most network engineers either add the configuration to point to the DNS servers or disable DNS using the **no ip domain-lookup** command.

Wrong Default Router IP Address Setting

Clearly, having a host that lists the wrong IP address as its default router causes problems. Hosts rely on the default router when sending packets to other subnets, and if a host lists the wrong default router setting, the host may not be able to send packets to a different subnet.

Figure 5-6 shows just such an example. In this case, hosts A and B both misconfigure 10.1.3.4 as the default router due to the same piece of bad documentation. Router R3 uses IP address 10.1.3.3. (For the sake of discussion, assume that no other host or router in this subnet currently uses address 10.1.3.4.)

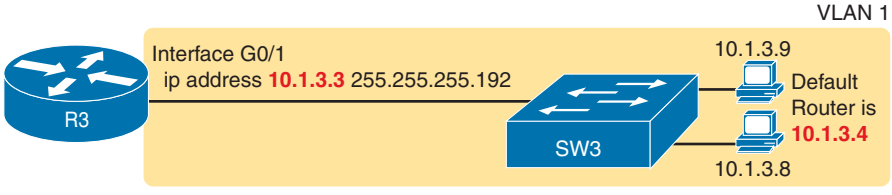


Figure 5-6 *Incorrect Default Router Setting on Hosts A and B*

In this case, several functions do work. For instance, hosts A and B can send packets to other hosts on the same LAN. The CSR at the router CLI can issue a **ping 10.1.3.9** and **ping 10.1.3.8** command, and both work. As a result of those two working pings, R3 would list the MAC address of the two PCs in the output of the **show arp** command. Similarly, the hosts would list R3's 10.1.3.3 IP address (and matching MAC address) in their ARP caches (usually displayed with the **arp –a** commThe one big problem in this case happens when the hosts try to send packets off-subnet. In that case, try to send the packets to IP address 10.1.3.4 next, which fails.

Root Causes Based on the Default Router's Configuration

While hosts must have correct IPv4 settings to work properly, having correct settings does not guarantee that a LAN-based host can successfully send a packet to the default router. The LAN between the host and the router must work. In addition, the router itself must be working correctly, based on the design of the internetwork.

This next topic looks at problems between hosts and their default router in which the root cause exists on the router. In particular, this topic looks at three main topics. The first topic looks at the trunking configuration required on a router to support multiple VLANs (known as router on a stick, or ROAS). Following that, the text examines typical DHCP issues. The final root cause discussed here is the status of the router interface and what causes that interface to fail.

Mismatched VLAN Trunking Configuration with Router on a Stick

Examples that teach configuration details often focus on one topic at a time. For instance, IPv4 configuration examples may show a host and its default router setting with the IP address configured on the router's LAN interface, as shown earlier in Example 5-1. However, the details of the LAN to which the host and router attach may be completely omitted, to focus on the IPv4 details.

Troubleshooting, both in real life and on the exams requires that you put all the pieces together. This next example shows a great case of how the troubleshooting process suffers if you forget to think about both the router and switch part of the problem. This example shows a valid router configuration that, unfortunately, does not match the configuration on the neighboring LAN switch like it should.

The next example focuses on how to connect routers to the subnets on multiple VLANs in the same campus LAN. Today, most sites in an enterprise LAN use at least two VLANs. To make routing work today, one of two options is typically used:

- **Router on a Stick (ROAS):** A router connects to the LAN, with one physical interface configured for VLAN trunking. The router has an IP address in each subnet, with one subnet per VLAN. The router configuration adds each matched subnet and associated VLAN to a subinterface.
- **Layer 3 switch:** Also called a multilayer switch, a Layer 3 switch performs the same job as a router using ROAS, but the switch has routing functions built in. The switch configuration adds each matched subnet and associated VLAN to a VLAN interface.

This example happens to use ROAS, but many of the same kinds of mistakes shown here can be made with Layer 3 switch configurations, as well.

First, the following list outlines the rules for configuring ROAS, using 802.1Q, on both the router and the neighboring switch.



Step 1. On the router, for each VLAN that is not the native VLAN, do the following:

A. Create a unique subinterface for each VLAN that needs to be routed (**interface** *type number.subint*)

B. Enable 802.1Q, and associate one specific VLAN with the subinterface in subinterface config mode (**encapsulation dot1q** *vlan_id*)

C. Configure IP settings (address and mask) in subinterface config mode (**ip address** *address mask*)

Step 2. On the router, for the native VLAN, if using it, use one of the two following options:

A. Configure just like for other VLANs, except add the native keyword to the encapsulation command (**encapsulation dot1q** *vlan_id* **native**).

Or

B. Configure the IP address on the physical LAN interface, without a subinterface and without the **encapsulation dot1q** command.

Step 3. On the switch, enable trunking (because the router will not negotiate to enable 802.1Q trunking):

A. Enable trunking with the **switchport mode trunk** interface subcommand.

B. Set the native VLAN to the same VLAN expected on the router, using the **switchport trunk native vlan** *vlan-id* interface subcommand.

Keeping that long list handy for reference, let's next walk through a brief example of the router configuration. First, imagine that previously a site used a single VLAN; so, the router configuration ignored VLAN trunking, with the IP address configured on the physical LAN interface on the router. All hosts sat in default VLAN 1. The router could ignore the VLAN details, not use trunking, and act as default router for all hosts in VLAN 1, as shown in Figure 5-7.

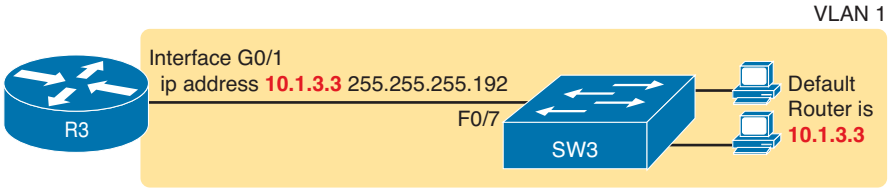


Figure 5-7 *Router IP Address Configuration, Without Trunking*

Then, management planned an expansion in which a second VLAN will be used. This particular company has one network engineer in charge of routers and the other in charge of switches. When planning the changes with the switch engineer, the two engineers did not listen to each other very well, and then the router engineer went off to plan the changes to the router. The router engineer planned to make the following changes to use ROAS:

- Use ROAS on interface G0/1 to support both users in old subnet 10.1.3.0/26, in VLAN 1, and users in new subnet 10.1.3.64/26, in VLAN 2.
- To support VLAN 1 users, leave 10.1.3.3/26 configured as is on the physical interface. This takes advantage of the option to configure the native VLAN IP address on the physical interface because VLAN 1 is the default native VLAN.
- Add a ROAS subinterface to the router configuration to support VLAN 2, using address 10.1.3.65/26 as the router IP address/mask in that subnet.

Figure 5-8 shows the concepts and configuration.

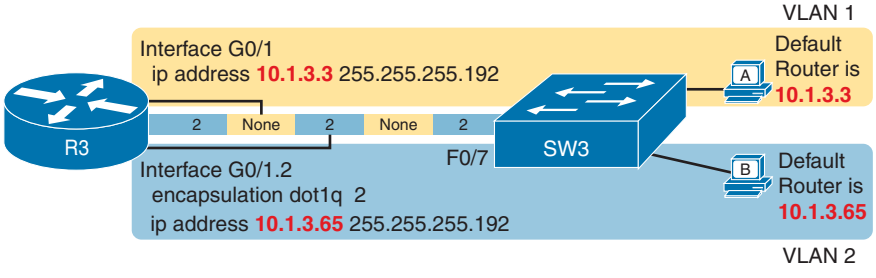


Figure 5-8 *Router IP Address Configuration, with ROAS, and Native VLAN 1*

This configuration could work perfectly well—as long as the switch has a matching correct VLAN trunking configuration. The router configuration implies a couple of things about VLAN trunking, as follows:

- With the IP address listed on physical interface G0/1, the configuration implies that the router intends to use the native VLAN, sending and receiving untagged frames.

- The router intends to use VLAN 2 as a normal VLAN, sending and receiving frames tagged as VLAN 2.

The switch (SW3) needs to configure VLAN trunking to match that logic. In this case, that means to enable trunking on that link, support VLANs 1 and 2, and make sure VLAN 1 is the native VLAN. Instead, in this case, the switch engineer actually added the trunk configuration to the wrong port, with the F0/7 port, connected to router R3, having these settings:

switchport mode access—The port does not trunk.

switchport access v7lan 7—The port is assigned to VLAN 7.

The first command confirms, without a doubt, that the link from R3 to SW3 does not trunk. SW1 will not pass any VLAN 2 traffic over that link at all. A standard ping of host B's IP address from R3 fails; likewise, a **ping 10.1.3.65** command from host B fails.

The second command states that the access VLAN on F0/7 is VLAN 7, which means that SW1 will not forward VLAN 1's traffic over the link to R3, either. Again, pings between R3 and hosts in VLAN 1 will fail as well.

In summary, for ROAS configurations, take the time to verify the matching configuration on the neighboring switch. In particular



- Make sure the switch enables trunking (**switchport mode trunk**).
- Make sure the switch sets the correct VLAN as that trunk's native VLAN (**switchport trunk native vlan** *vlan-id*).
- Make sure the switch knows about all the VLANs the router has configured (**vlan** *vlan-id*).

DHCP Relay Issues

Hosts that use DHCP to lease an IP address (and learn other settings) rely on the network to pass the DHCP messages. In particular, if the internetwork uses a centralized DHCP server, with many remote LAN subnets using the DHCP server, the routers have to enable a feature called *DHCP Relay* to make DHCP work. Without DHCP Relay, DHCP requests from hosts never leave the local LAN subnet.

Figure 5-9 shows the big ideas behind how DHCP Relay works. In this example, a DHCP client (Host A) sits on the left, with the DHCP server (172.16.2.11) on the right. The client begins the DHCP lease process by sending a DHCP Discover message, one that would flow only across the local LAN without DHCP Relay configured on router R1. To be ready to forward the Discover message, R1 enables DHCP Relay with the **ip helper-address 172.16.2.11** command configured under its G0/0 interface.

The steps in the figure point out the need for DHCP Relay. At Step 1, host A sends a message, with destination IP and L2 broadcast address of 255.255.255.255 and ff:ff:ff:ff:ff:ff, respectively. Packets sent to this IP address, the "local subnet broadcast address," should never be forwarded past the router. All devices on the subnet receive and process the frame, and because the **ip helper-address** command configured on R1, router R1 will continue to de-encapsulate the frame and packet to identify that it is a DHCP request and takes action. Step 2 shows the results of DHCP Relay, where R1 changes both the source and destination IP address, with R1 routing the packet to the address listed in the command: 172.16.2.11.

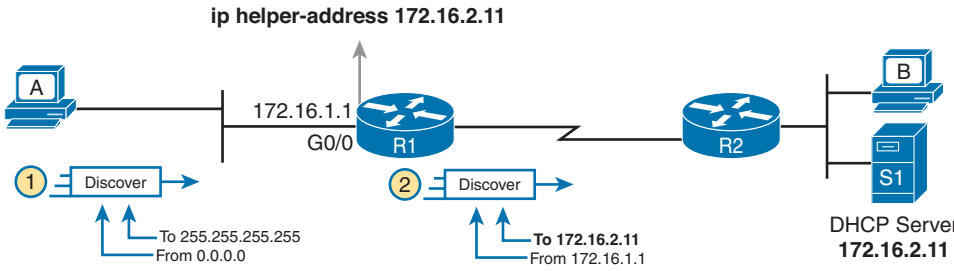


Figure 5-9 *IP Helper Address Effect*

Now, back to troubleshooting. Messages sent by a DHCP client can reach the DHCP server if the following are true:

Key Topic

- The server is in the same subnet as the client, with connectivity working between the two.
- The server is on another subnet, with the router on the same subnet as the client correctly implementing DHCP Relay and with IP connectivity from that router to the DHCP server.

5

Two common mistakes can be made with DHCP Relay, both of which are fairly obvious. If the router omits the **ip helper-address** command on a LAN interface (or subinterface when using ROAS, or VLAN interface with a multilayer switching [MLS] configuration), DHCP fails for those clients. If the configuration includes the **ip helper-address** command but lists the wrong DHCP server IP address, again DHCP fails completely.

The symptom in both cases is that the client learns nothing with DHCP.

For instance, Example 5-2 shows an updated configuration for ROAS on router R3, based on the same scenario as in Figure 5-8. The router configuration works fine for supporting IPv4 and making the router reachable. However, only one subinterface happens to list an **ip helper-address** command.

Example 5-2 *Forgetting to Support DHCP Relay on a ROAS Subinterface*

```
interface GigabitEthernet0/1
 ip address 10.1.3.3 255.255.255.192
 ip helper-address 10.1.2.130
!
interface GigabitEthernet0/1.2
 encapsulation dot1q 2
 ip address 10.1.3.65 255.255.255.192
```

In this case, hosts in VLAN 1 that want to use DHCP can, assuming the host at address 10.1.2.130 is indeed the DHCP server. However, hosts in VLAN 2 will fail to learn settings with DHCP because of the lack of an **ip helper-address** command.

Router LAN Interface and LAN Issues

At some point, the problem isolation process may show that a host cannot ping its default router and vice versa. That is, neither device can send an IP packet to the other device on the same subnet. This basic test tells the engineer that the router, host, and LAN between them,

for whatever reasons, cannot pass the packet encapsulated in an Ethernet frame between the two devices.

The root causes for this basic LAN connectivity issue fall into two categories:

- Problems that cause the router LAN interface to fail
- Problems with the LAN itself

A router's LAN interface must be in a working state before the router will attempt to send packets out that interface (or receive packets in that interface). Specifically, the router LAN interface must be in an up/up state; if in any other state, the router will not use the interface for packet forwarding. So, if a ping from the router to a LAN host fails (or vice versa), check the interface status, and if not up, find the root cause for the router interface to not be up.

Alternatively, the router interface can be in an up/up state, but problems can exist in the LAN itself. In this case, every topic related to Ethernet LANs may be a root cause. In particular, all the topics reviewed in Chapter 3, such as Ethernet cable pinouts, port security, and even Spanning Tree Protocol, may be root causes of LAN issues.

For instance, in Figure 5-10, router R3 connects to a LAN with four switches. R3's LAN interface (G0/1) can reach an up/up state if the link from R3 to SW1 works. However, many other problems could prevent R3 from successfully sending an IP packet, encapsulated in an Ethernet frame, to the hosts attached to switches SW3 and SW4.

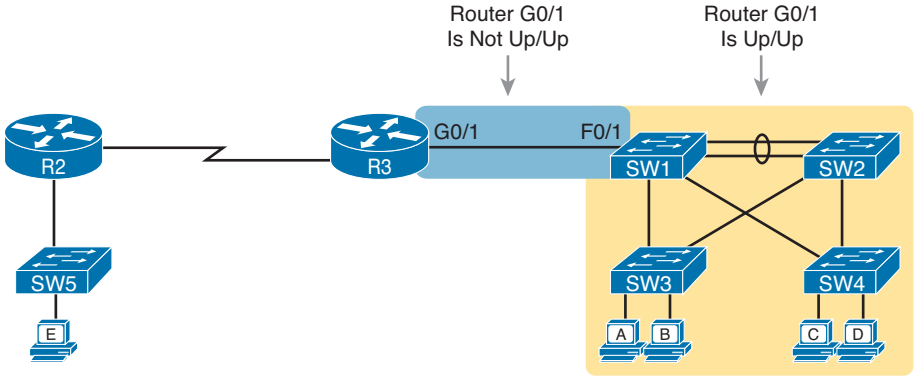


Figure 5-10 *Where to Look for Problems Based on Router LAN Interface Status*

NOTE This book leaves the discussion of LAN issues, as shown on the right side of Figure 5-10, to Part I of this book.

Router LAN interfaces can fail to reach a working up/up state for several reasons. Table 5-1 lists the common reasons discussed within the scope of the CCNA exam.

Key Topic

Table 5-1 Common Reasons Why Router LAN Interfaces Are Not Up/Up

Reason	Description	Router Interface State
Speed mismatch	The router and switch can both use the **speed** interface subcommand to set the speed, but to different speeds.	down/down
Shutdown	The router interface has been configured with the **shutdown** interface subcommand.	Admin down/down
Err-disabled switch	The neighboring switch port uses port security, which has put the port in an err-disabled state.	down/down
No cable/bad cable	The router has no cable installed, or the cable pinouts are incorrect.*	down/down

* Cisco switches use a feature called auto-mdix, which automatically detects some incorrect cabling pinouts and internally changes the pin logic to allow the cable to be used. As a result, not all incorrect cable pinouts result in an interface failing.

Using the speed mismatch root cause as an example, you could configure Figure 5-10's R3's G0/1 with the **speed 1000** command and SW1's F0/1 interface with the **speed 100** command. The link simply cannot work at these different speeds, so the router and switch interfaces both fall to a down/down state. Example 5-3 shows the resulting state, this time with the **show interfaces description** command, which lists one line of output per interface.

Example 5-3 *show interfaces description Command with Speed Mismatch*

```
R3# show interfaces description
Interface                      Status         Protocol Description
Gi0/0                          up             up
Gi0/1                          down           down       link to campus LAN
Se0/0/0                        admin down     down
Se0/0/1                        up             up
Se0/1/0                        up             up
Se0/1/1                        admin down     down
```

Problems with Routing Packets Between Routers

The first half of this chapter focused on the first hop that an IPv4 packet takes when passing over a network. This second major section now looks at issues related to how routers forward the packet from the default router to the final host.

In particular, this section begins by looking at the IP routing logic inside a single router. These topics review how to understand what a router currently does. Following that, the discussion expands to look at some common root causes of routing problems, causes that come from incorrect IP addressing, particularly when the addressing design uses variable-length subnet masks (VLSM).

The end of this section turns away from the core IP forwarding logic, looking at other issues that impact packet forwarding, including issues related to router interface status (which needs to be up/up) and how IPv4 access control lists (ACL) can filter IPv4 traffic.

IP Forwarding by Matching the Most Specific Route

Any router's IP routing process requires that the router compare the destination IP address of each packet with the existing contents of that router's IP routing table. Often, only one route matches a particular destination address. However, in some cases, a particular destination address matches more than one of the router's routes.

The following CCENT and CCNA features can create overlapping subnets:

- Autosummary (as discussed in Chapter 10, "Implementing EIGRP for IPv4")
- Manual route summarization
- Static routes
- Incorrectly designed subnetting plans that cause subnets overlap their address ranges

In some cases, overlapping routes cause a problem; in other cases, the overlapping routes are just a normal result of using some feature. This section focuses on how a router chooses which of the overlapping routes to use, for now ignoring whether the overlapping routes are a problem. The section "Routing Problems Caused by Incorrect Addressing Plans," later in this chapter, discusses some of the problem cases.

Now on to how a router matches the routing table, even with overlapping routes in its routing table. If only one route matches a given packet, the router uses that one route. However, when more than one route matches a packet's destination address, the router uses the "best" route, defined as follows:



When a particular destination IP address matches more than one route in a router's IPv4 routing table, the router uses the most specific route—in other words, the route with the longest prefix length mask.

Using **show ip route** and Subnet Math to Find the Best Route

We humans have a couple of ways to figure out what choice a router makes for choosing the best route. One way uses the **show ip route** command, plus some subnetting math, to decide the route the router will choose. To let you see how to use this option, Example 5-4 shows a series of overlapping routes.

Example 5-4 *show ip route Command with Overlapping Routes*

```
R1# show ip route ospf
Codes: L - local, 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, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 172.16.25.129 to network 0.0.0.0

     172.16.0.0/16 is variably subnetted, 9 subnets, 5 masks
O       172.16.1.1/32 [110/50] via 172.16.25.2, 00:00:04, Serial0/1/1
O       172.16.1.0/24 [110/100] via 172.16.25.129, 00:00:09, Serial0/1/0
O       172.16.0.0/22 [110/65] via 172.16.25.2, 00:00:04, Serial0/1/1
O       172.16.0.0/16 [110/65] via 172.16.25.129, 00:00:09, Serial0/1/0
O       0.0.0.0/0 [110/129] via 172.16.25.129, 00:00:09, Serial0/1/0
!
```

NOTE As an aside, the **show ip route ospf** command lists only OSPF-learned routes, but the statistics for numbers of subnets and masks (9 and 5 in the example, respectively) are for all routes, not just OSPF-learned routes.

To predict which of its routes a router will match, two pieces of information are required: the destination IP address of the packet and the contents of the router's routing table. The subnet ID and mask listed for a route defines the range of addresses matched by that route. With a little subnetting math, a network engineer can find the range of addresses matched by each route. For instance, Table 5-2 lists the five subnets listed in Example 5-4 and the address ranges implied by each.

Table 5-2 Analysis of Address Ranges for the Subnets in Example 5-4

Subnet / Prefix	Address Range
172.16.1.1/32	172.16.1.1 (just this one address)
172.16.1.0/24	172.16.1.0–172.16.1.255
172.16.0.0/22	172.16.0.0–172.16.3.255
172.16.0.0/16	172.16.0.0–172.16.255.255
0.0.0.0/0	0.0.0.0–255.255.255.255 (all addresses)

NOTE The route listed as 0.0.0.0/0 is the default route.

As you can see from these ranges, several of the routes' address ranges overlap. When matching more than one route, the route with the longer prefix length is used. That is, a route with /16 is better than a route with /10; a route with a /25 prefix is better than a route with a /20 prefix; and so on.

For example, a packet sent to 172.16.1.1 actually matches all five routes listed in the routing table in Example 5-4. The various prefix lengths range from /0 to /32. The longest prefix (largest /P value, meaning the best and most specific route) is /32. So, a packet sent to 172.16.1.1 uses the route to 172.16.1.1/32, and not the other routes.

The following list gives some examples of destination IP addresses. For each address, the list describes the routes from Table 5-2 that the router would match, and which specific route the router would use.

- **172.16.1.1:** Matches all five routes; the longest prefix is /32, the route to 172.16.1.1/32.
- **172.16.1.2:** Matches last four routes; the longest prefix is /24, the route to 172.16.1.0/24.
- **172.16.2.3:** Matches last three routes; the longest prefix is /22, the route to 172.16.0.0/22.
- **172.16.4.3:** Matches the last two routes; the longest prefix is /16, the route to 172.16.0.0/16.

Using **show ip route** *address* to Find the Best Route

A second way to identify the route a router will use, one that does not require any subnetting math, is the **show ip route** *address* command. The last parameter on this command is the IP address of an assumed IP packet. The router replies by listing the route it would use to route a packet sent to that address.

For example, Example 5-5 lists the output of the **show ip route 172.16.4.3** command on the same router used in Example 5-4. The first line of (highlighted) output lists the matched route: the route to 172.16.0.0/16. The rest of the output lists the details of that particular route, like the outgoing interface of S0/1/0 and the next-hop router of 172.16.25.129.

Example 5-5 *show ip route Command with Overlapping Routes*

```
R1# show ip route 172.16.4.3
Routing entry for 172.16.0.0/16
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 10.2.2.5 on Serial0/1/0, 14:22:06 ago
  Routing Descriptor Blocks:
  * 172.16.25.129, from 172.16.25.129, 14:22:05 ago, via Serial0/1/0
      Route metric is 65, traffic share count is 1
```

Certainly, if you have an option, just using a command to check what the router actually chooses is a much quicker option than doing the subnetting math.

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 5-11 shows the output of a sample **show ip route** command. The figure numbers various parts of the command output for easier reference, with Table 5-3 describing the output noted by each number.

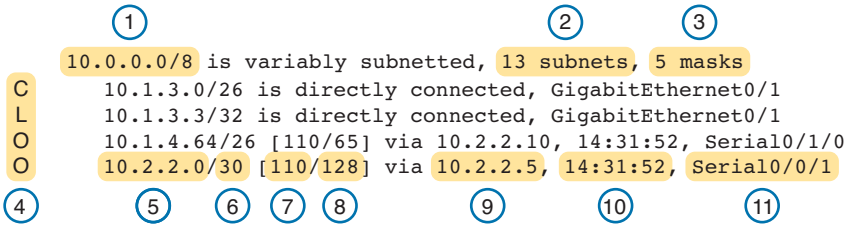


Figure 5-11 *show ip route Command Output Reference*

Table 5-3 Descriptions of the show ip route Command Output

Item	Idea	Value in the Figure	Description
1	Classful network	10.0.0.0/8	The routing table is organized by classful network. This line is the heading line for classful network 10.0.0.0; it lists the default mask for class A networks (/8).
2	Number of subnets	13 subnets	Lists the number of routes for subnets of the classful network known to this router, from all sources, including local routes - the /32 routes that match each router interface IP address.
3	Number of masks	5 masks	The number of different masks used in all routes known to this router inside this classful network.
4	Legend code	C, L, O	A short code that identifies the source of the routing information. *O* is for OSPF, *D* for EIGRP, *C* for Connected, *S* for Static, and *L* for Local. (See Example 5-4 for a sample of the legend.)
5	Subnet ID	10.2.2.0	The subnet number of this particular route.
6	Prefix length	/30	The prefix mask used with this subnet.
7	Administrative distance	110	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.
8	Metric	128	The metric for this route.
9	Next-hop router	10.2.2.5	For packets matching this route, the IP address of the next router to which the packet should be forwarded.
10	Timer	14:31:52	For OSPF and EIGRP routes, this is the time since the route was first learned.
11	Outgoing interface	Serial0/0/1	For packets matching this route, the interface out which the packet should be forwarded.

Routing Problems Caused by Incorrect Addressing Plans

The existence of overlapping routes in a router's routing table does not necessarily mean a problem exists. Both automatic and manual route summarization result in overlapping routes on some routers, with those overlaps not causing problems. However, some overlaps, particularly those related to addressing mistakes, can cause problems for user traffic. So, when troubleshooting, if overlapping routes exist, the engineer should also look for the specific reasons for overlaps that actually cause a problem.

Simple mistakes in either the IP addressing plan or the implementation of that plan can cause overlaps that also cause problems. In these cases, one router claims to be connected to a subnet with one address range, while another router claims to be connected to another subnet with an overlapping range, breaking IP addressing rules. The symptoms are that the routers sometimes forward the packets to the right host, but sometimes not.

This problem can occur whether or not VLSM is used. However, the problem is much harder to find when VLSM is used. This section reviews VLSM, shows examples of the problem both with and without VLSM, and discusses the configuration and verification commands related to these problems.

Recognizing When VLSM Is Used or Not

An internetwork is considered to be using VLSM when multiple subnet masks are used for different subnets of *a single classful network*. For example, if in one internetwork all subnets come from network 10.0.0.0, and masks /24, /26, and /30 are used, the internetwork uses VLSM.

Sometimes people fall into the trap of thinking that any internetwork that uses more than one mask must be using VLSM, but that is not always the case. For instance, if an internetwork uses subnets of network 10.0.0.0, all of which use mask 255.255.240.0, and subnets of network 172.16.0.0, all of which use a 255.255.255.0 mask, the design does not use VLSM. Two different masks are used, but only one mask is used in any single classful network. The design must use more than one mask for subnets of a single classful network to be using VLSM.

Only classless routing protocols can support VLSM. The current CCENT and CCNA certifications cover only classless routing protocols (OSPF and EIGRP), so in all routing protocol discussions for this book, VLSM should be supported. However, for real life, note that RIPv2 (as a classless routing protocol) also supports VLSM, whereas classful routing protocols RIPv1 and Interior Gateway Routing Protocol (IGRP) cannot.

Overlaps When Not Using VLSM

Even when you are not using VLSM, addressing mistakes that create overlapping subnets can occur. For instance, Figure 5-12 shows a sample network with router LAN IP address/mask information. An overlap exists, but it might not be obvious at first glance.

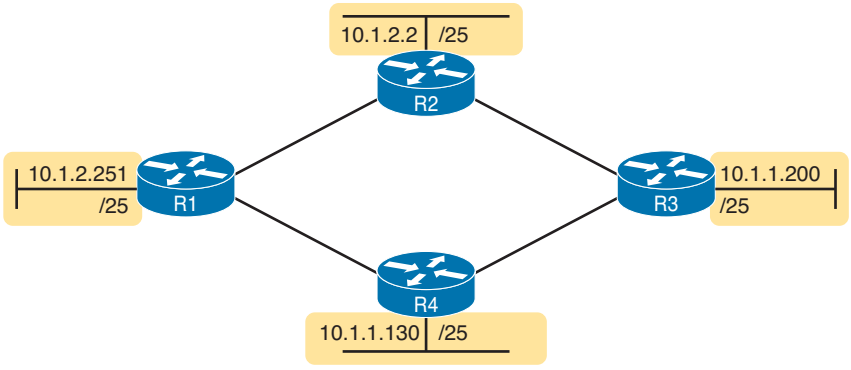


Figure 5-12 *IP Addresses on LAN Interfaces, with One Mask (/25) in Network 10.0.0.0*

If an overlap exists when all subnets use the same mask, the overlapping subnets have the exact same subnet ID, and the exact same range of IP addresses in the subnet. To find the overlap, all you have to do is calculate the subnet ID of each subnet and compare the numbers. For instance, Figure 5-13 shows an updated version of Figure 5-12, with subnet IDs shown and with identical subnet IDs for the LANs off R3 and R4.

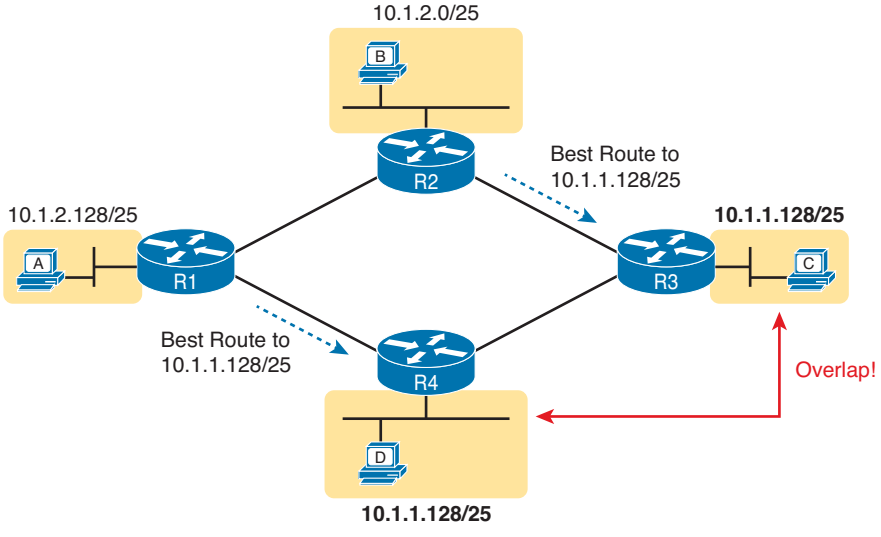


Figure 5-13 *Subnet IDs Calculated from Figure 5-12*

Using the same subnet in two different places (as is done in Figure 5-13) breaks the rules of IPv4 addressing because the routers get confused about where to send packets. In this case, for packets sent to subnet 10.1.1.128/25, some routers send packets so they arrive at R3, whereas others think the best route points toward R4. Assuming all routers use a routing protocol, such as OSPF, both R3 and R4 advertise a route for 10.1.1.128/25.

In this case, R1 and R2 will likely send packets to two different instances of subnet 10.1.1.128/25. With these routes, hosts near R1 will be able to communicate with 10.1.1.128/25 hosts off R4's LAN, but not those off R3's LAN, and vice versa.

Finally, although the symptoms point to some kind of routing issues, the root cause is an invalid IP addressing plan. No IP addressing plan should use the same subnet on two different LANs, as was done in this case. The solution: Change R3 or R4 to use a different, non-overlapping subnet on its LAN interface.

Overlaps When Using VLSM

When using VLSM, the same kinds of addressing mistakes can lead to overlapping subnets; they just may be more difficult to notice.

First, overlaps between subnets that have different masks will cause only a partial overlap. That is, two overlapping subnets will have different sizes and possibly different subnet IDs. The overlap occurs between all the addresses of the smaller subnet, but with only part of the larger subnet. Second, the problems between hosts only occur for some destinations (specifically the subset of addresses in the overlapped ranges), making it even tougher to characterize the problem.

For instance, Figure 5-14 shows an example with a VLSM overlap. The figure shows only the IP address/mask pairs of router and host interfaces. First, look at the example and try to find the overlap by looking at the IP addresses.

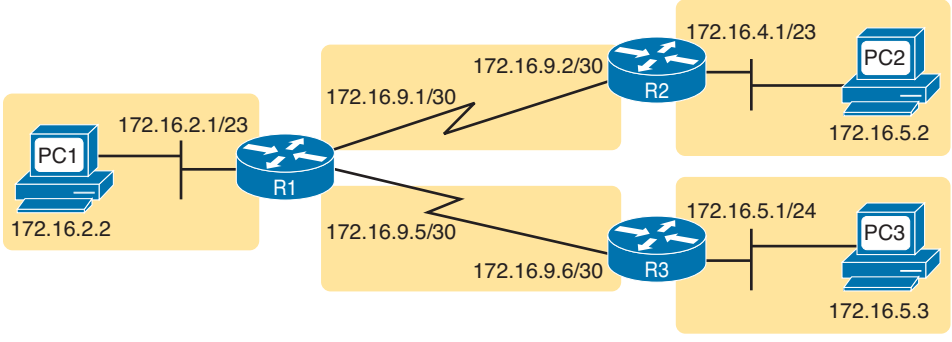


Figure 5-14 *VLSM IP Addressing Plan in Network 172.16.0.0*

To find the overlap, the person troubleshooting the problem needs to analyze each subnet, finding not only the subnet ID but also the subnet broadcast address and the range of addresses in the subnet. If the analysis stops with just looking at the subnet ID, the overlap may not be noticed (as is the case in this example).

Figure 5-15 shows the beginning analysis of each subnet, with only the subnet ID listed. Note that the two overlapping subnets have different subnet IDs, but the lower-right subnet (172.16.5.0/24) completely overlaps with part of the upper-right subnet (172.16.4.0/23). (Subnet 172.16.4.0/23 has a subnet broadcast address of 172.16.5.255, and subnet 172.16.5.0/24 has a subnet broadcast address of 172.16.5.255.)

To be clear, the design with actual subnets whose address ranges overlap is incorrect and should be changed. However, once implemented, the symptoms show up as routing problems, like the similar case without VLSM. **ping** commands fail, and **traceroute** commands do complete for only certain hosts (but not all).

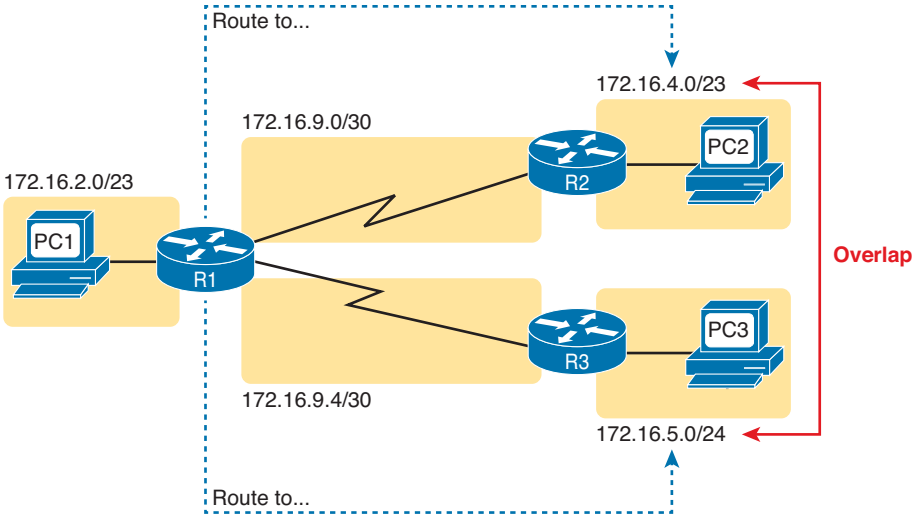


Figure 5-15 *A VLSM Overlap Example, but with Different Subnet IDs*

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 sometimes not, as follows:

Key Topic

- **Preventing the overlap on a single router:** IOS detects the overlap when the **ip address** command implies an overlap with another **ip address** command *on the same router*.
- **Allowing the overlap on different routers:** IOS cannot detect an overlap when an **ip address** command overlaps with an **ip address** command on another router.

The router shown in Example 5-6 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 attempting to configure Fa0/1 with 172.16.5.193/26. The ranges of addresses in each subnet are as follows:

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 5-6 *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.

As an aside of how IOS handles these errors, IOS only performs the subnet overlap check for interfaces that are not in a shutdown state. When configuring an interface in shutdown state, IOS actually accepts the **ip address** command that would cause the overlap. Later, when the **no shutdown** command is issued, IOS checks for the subnet overlap and issues the same error message shown in Example 5-6. IOS leaves the interface in the shutdown state until the overlap condition has been resolved.

IOS cannot detect the configuration of overlapping subnets on different routers, as shown in Example 5-7. The example shows the configuration of the two overlapping subnets on R2 and R3 from Figure 5-15.

Example 5-7 *Two Routers Accept Overlapped Subnets*

```
! First, on router R2
R2# configure terminal
R2(config)# interface G0/0
R2(config-if)# ip address 172.16.4.1 255.255.254.0
! Next, on router R3
R3# configure terminal
R3(config)# interface G0/0
R3(config-if)# ip address 172.16.5.1 255.255.255.0
```

Router WAN Interface Status

One of the steps in the IP routing troubleshooting process described earlier, in the "Router LAN Interface and LAN Issues" 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."

So far, the ICND1 and ICND2 books have explored only basic information about how serial links work. For now, know that both routers must have working serial interfaces in an up/up state before they can send IPv4 packets to each other. The two routers should also have serial IP addresses in the same subnet.

Later, the chapters in Part IV further develop the details of WAN links, including what is required for routers to use these links to forward IP packets.

Filtering Packets with Access Lists

Access control lists (ACL) cause some of the biggest challenges when troubleshooting problems in real networking jobs. End-user packets sent by user applications do not look exactly like packets sent by testing tools such as ping and traceroute. The ACLs sometimes filter the ping and traceroute traffic, making the network engineer think some other kind of problems exists when no problems exist at all. Or, the problem with the end-user traffic really is

caused by the ACL, but the ping and traceroute traffic works fine, because the ACL filters the user traffic but not the ping and traceroute traffic.

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 the ICND1 book, Chapter 22, "Basic IP Access Control Lists," to find the range of addresses matched by an ACL command: dd the address and wildcard mask to find the end of the numeric ranAge.

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.

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 **tcp** or **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.

If you suspect ACLs are causing a 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. In that case, another way to find the interfaces and direction for any IP ACLs is the **show ip interfaces** command, as shown in Example 5-8.

Example 5-8 *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 quickest way to look at the ACL is to use the **show running-config** command. If not available, the **show access-lists** and **show ip access-lists** commands list the same details shown in the configuration commands and a counter for the number of packets matching each line in the ACL. Example 5-9 shows an example.

Example 5-9 *show ip access-lists Command Example*

```
R1# show ip access-lists
Extended IP access list 102
    10 permit ip 10.1.2.0 0.0.0.255 10.1.4.0 0.0.1.255 (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 5-9, 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. If you configure 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.

NOTE Cisco sometimes recommends adding the explicit **deny any** statement at the end of the ACL for easier troubleshooting.

Finally, as a reminder about interpreting ACL commands, when you know the command comes from a router, it is easy to decide the range of addresses matched by an address and wildcard mask. The low end of the range is the address (the first number), and the high end of the range is the sum of the address and wildcard mask. For instance, with ACL 102 in Example 5-9, which is obviously configured in some router, the ranges are as follows:

Source 10.1.2.0, wildcard 0.0.0.255: Matches from 10.1.2.0 through 10.1.2.255

Destination 10.1.4.0, wildcard 0.0.1.255: Matches from 10.1.4.0 through 10.1.5.255

Exam Preparation Tasks

Review All the Key Topics

Review the most important topics from this chapter, noted with the Key Topic icon. Table 5-4 lists these key topics and where each is discussed.

Key Topic

Table 5-4 Key Topics for Chapter 5

Key Topic Element	Description	Page Number
List	Two root causes of DNS problems.	
List	The rules for configuring ROAS.	
List	Items to verify when troubleshooting ROAS configurations.	
List	Items to verify for switch trunking configuration to match a router's ROAS configuration.	
List	Conditions that must be true for DHCP messages to be able to flow from a client to a DHCP server.	
Table 5-1	Common reasons why router LAN interfaces are not up/up.	
Definition	When more than one route matches a packet's destination address, the router uses the "best" (most specific) route.	
List	Types of overlapping IP address configuration issues that IOS can and cannot recognize.	

Complete the Tables and Lists from Memory

Print a copy of DVD Appendix D, "Memory Tables," or at least the section for this chapter, and complete the tables and lists from memory. DVD Appendix E, "Memory Tables Answer Key," includes completed tables and lists to check your work.

Definitions of Key Terms

After your first reading of the chapter, try to define these key terms, but do not be concerned about getting them all correct at that time. Chapter 22 directs you in how to use these terms for late-stage preparation for the exam.