

THE PRACTICAL **NETWORK ENGINEER'S** GUIDE

Master routing,
network diagnostics,
and troubleshooting



Robert McMillen

<packt>

The Practical Network Engineer's Guide

Master routing, network diagnostics, and troubleshooting

Robert McMillen

<packt>

The Practical Network Engineer's Guide

Copyright © 2026 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

Portfolio Director: Kartikey Pandey

Relationship Lead: Prachi Rana

Project Manager: Sonam Pandey

Content Engineers: Apramit Bhattacharya and Pragya Mittal

Technical Editor: Aysha Nadeem

Indexer: Rekha Nair

Production Designer: Vijay Kamble

Growth Lead: Vikramaditya Viswanath

First published: September 2026

Production reference: 1290926

Published by Packt Publishing Ltd.

Grosvenor House

11 St Paul's Square

Birmingham

B3 1RB, UK.

ISBN 978-1-80849-527-4

www.packtpub.com

Contributors

About the author

Robert McMillen, MBA, MCT, is an IT consultant, educator, and author with extensive experience across infrastructure, networking, cloud computing, cybersecurity, and AI. Over the course of his career, he has worked on large-scale and high-profile technology initiatives, including restoring email data at Enron, consulting with the White House, and managing complex network mergers and infrastructure projects for organizations valued at more than \$1 billion.

Alongside his consulting work, Robert has built a significant career in technology education. He has authored numerous books and created more than 150 IT courses for leading learning platforms and publishers, including LinkedIn Learning, Pluralsight, Cengage, and O'Reilly. His courses have reached millions of learners and are used by colleges and universities around the world.

Drawing on decades of hands-on industry experience and classroom teaching, Robert focuses on making complex technologies practical, accessible, and relevant to the real-world challenges IT professionals face.

Table of Contents

Preface	xiii
----------------	-------------

Part 1: Understanding How Networks Communicate	1
---	----------

Chapter 1: How Networks Communicate	3
--	----------

Technical requirements.....	4
-----------------------------	---

Why networks exist	4
--------------------------	---

What a network needs	6
----------------------------	---

Endpoints servers and network infrastructure • 6

How infrastructure devices divide the work • 6

Reading the network as a path • 7

Wired and wireless transmission • 8

IP and transport protocols • 9

Transport behavior and security are separate decisions • 11

Operating systems and network stacks • 11

What modern networks must deliver.....	12
---	-----------

Security • 12

Performance reliability and automation • 12

Access without unnecessary complexity • 12

Useful data at a manageable cost • 13

Where the Internet is	13
------------------------------------	-----------

Endpoints services and resilience • 13

How the Internet developed • 14

Who coordinates standards names and addresses • 14

How the responsibilities fit together • 15

How data travels over IP.....	16
--------------------------------------	-----------

How the host prepares data • 16

From the local device to the destination • 17

Using tracer to see the route • 19

What tracer does not prove • 20

Summary	21
 Chapter 2: Seeing the Network One Layer at a Time	 23

Technical requirements.....	24
The OSI model and its seven layers.....	24
Why shared layers and protocols matter • 25	
Physical layer functions and initial host checks • 26	
Data link layer components and operations • 27	
Network layer routing and devices • 27	
Transport layer protocols and ports • 27	
Session presentation and application functions • 28	
Seeing protocol layers with Wireshark • 28	
<i>Reading a captured packet from the outside in • 29</i>	
IPv4 addressing and binary representation	30
Converting an octet between binary and decimal • 30	
Hexadecimal in networking • 32	
Essential network connectivity and routing tools	33
Inspecting configuration with ipconfig • 35	
Testing reachability with ping • 35	
Examining address resolution with arp • 36	
Analyzing a path with pathping and tracert • 37	
Reading and changing the local routing table • 38	
<i>Reading the route print output • 38</i>	
<i>Adding a persistent static route • 39</i>	
Summary	40
 Chapter 3: How IPv4 Addresses Organize Network Traffic	 41

Technical requirements.....	42
Using binary to locate subnet boundaries	42
Reading prefix lengths as subnet masks • 42	
Testing whether a destination is local • 43	
From fixed address classes to flexible subnets	44
Reading the classful address ranges • 44	
Why CIDR replaced fixed classes • 46	
Choosing a prefix for the required host count • 46	

Calculating the network broadcast and host range • 48	
Checking the mask and gateway together • 49	
Crossing network boundaries with NAT and PAT	50
Sharing a public address with Hide NAT • 50	
Publishing a server with static NAT • 52	
Directing services with PAT and port forwarding • 53	
Summary	55
 Chapter 4: From Scarcity to Scale: How IPv6 Changes Addressing	 57
Technical requirements.....	58
Why IPv6 expands the addressing model	58
Reading IPv6 addresses and prefixes	59
From binary bits to hexadecimal hextets • 59	
Shortening an IPv6 address safely • 60	
Using the prefix to make a routing decision • 61	
Recognizing address scope in IPv4 and IPv6.....	62
Private IPv4 ranges and NAT • 62	
Global, unique-local, and link-local IPv6 addresses • 63	
Reading the structure of a unique local address • 64	
Understanding automatic address assignment	66
Reading IPv6 configuration on Windows	68
Summary	70
 Part 2: From Network Infrastructure to a Troubleshooting Method	 71
 Chapter 5: How Network Infrastructure Moves and Protects Traffic	 73
Technical requirements.....	73
Segmenting and routing traffic with VLANs.....	74
Reading a multilayer switch configuration • 75	
Choosing between static and dynamic routes	77
OSPF, EIGRP, and RIP inside an organization • 78	
BGP between autonomous systems • 78	
Managing static routes with PowerShell.....	79
Filtering traffic with firewalls	80
Stateless access control lists • 81	

Application proxies and protocol inspection • 81	
Stateful inspection • 82	
Building virtual network paths with Hyper-V	85
Controlling virtual adapters with PowerShell • 87	
Choosing external, internal, and private switch types • 87	
Managing routes and interfaces on macOS	89
Checking IPv4 settings in the macOS interface • 90	
Reading routes, reachability, paths, and DNS • 91	
Summary	93
 Chapter 6: A Systematic Approach to Network Troubleshooting	 95
Technical requirements.....	96
Network severity triage and stakeholder involvement	96
Criteria for assessing network severity • 96	
Common network issues in organizational environments	97
Network slowness and DNS failures • 98	
Troubleshooting corporate resource access	100
Video and audio conferencing quality issues • 100	
IP structure and addressing in network troubleshooting	102
Subnetting and subnet mask consistency • 104	
DNS and gateway configuration • 104	
Summary	107
 Part 3: Troubleshooting Access at the Network Edge	 109
 Chapter 7: Troubleshooting Remote and Home Networks	 111
Technical requirements.....	112
Common issues in remote and home networks.....	112
Troubleshooting remote desktop and VPN connectivity	112
Troubleshooting printing and local network discovery in remote environments.....	118
Diagnosing slowness, internet access, and dropped connections	120
Addressing dropped connections in wired and wireless networks • 120	
Wireless access point configuration and channel selection • 122	
Guest network configuration and isolation • 124	
Using built-in wireless access point diagnostic tools • 125	

Troubleshooting internet connectivity from the home network	126
Router and ISP troubleshooting steps • 127	
Resetting the ARP cache and router hardware checks • 131	
Summary	131
Chapter 8: Troubleshooting Enterprise Wireless Networks	133
Technical requirements.....	134
Wireless network security concerns and threats	134
Common causes of unstable or slow wireless connections.....	134
Managing user experience and bandwidth utilization	136
Sources of Wi-Fi interference in enterprise environments	137
Interference from appliances and electronic devices • 138	
Challenges of high-density wireless environments	138
Line of sight and device placement • 138	
Summary	139
Chapter 9: Troubleshooting VPNs and Remote Access	141
Technical requirements.....	142
Common VPN protocols and their characteristics	142
VPN protocol suitability in public networks.....	144
VPN connectivity considerations in Windows Server	144
Configuring Windows Server for VPN and remote access	147
Investigating VPN ports and WAN miniport problems	148
Client-side VPN configuration in Windows 11	149
DNS configuration and name resolution for VPN connectivity	150
Troubleshooting common VPN error codes.....	151
Summary	153
Part 4: Diagnosing the Enterprise Network with Evidence	155
Chapter 10: Essential Network Troubleshooting Tools	157
Technical requirements.....	158
Windows command-line network troubleshooting tools	158
IP configuration and management with ipconfig • 158	
Viewing network connections with netstat • 160	

DNS troubleshooting with nslookup • 162	
Address Resolution Protocol for network diagnostics • 163	
PowerShell network troubleshooting tools	165
Testing connectivity with Test-NetConnection • 165	
Tracing network paths with Test-NetConnection • 167	
Resolving names with PowerShell • 168	
Inspecting TCP connections with Get-NetTCPConnection • 169	
Linux network troubleshooting tools.....	170
Viewing network configuration with ifconfig • 170	
Managing routing tables with ip route • 170	
Connectivity testing with ping and traceroute • 171	
DNS diagnostics with dig • 171	
Port scanning and network analysis with nmap and netstat • 172	
Advanced Windows network troubleshooting commands.....	173
Using ping switches deliberately • 173	
Path analysis with pathping and tracert • 174	
Viewing routing tables with route print • 175	
Port scanning and packet analysis tools	175
Capturing relevant traffic with Wireshark • 175	
Comparing a scan with the captured traffic • 176	
Validating a targeted firewall change • 178	
Summary	179
 Chapter 11: Troubleshooting Routing, Switching, and Firewalls	 181
Technical requirements.....	182
Core Layer 3 switches and inter-VLAN routing.....	182
Layer 2 switches and VLAN segmentation	182
Firewall placement and internet connectivity.....	183
Server roles in network infrastructure	184
Email servers and cloud email services • 184	
VPN server placement and port forwarding • 184	
Wireless access points and controller services • 185	
Network printers and office placement.....	185
VoIP phone systems and network integration • 185	
Remote office connectivity: SD-WAN, MPLS, and VPN	185

Cloud server integration • 186	
Internet of Things devices in modern networks	186
Working through an inter-VLAN connection	186
Checking the Layer 3 switch • 188	
Assigning VLANs to switch ports • 188	
Troubleshooting inter-VLAN routing and subnet communication.....	189
Troubleshooting Windows Firewall and port connectivity	190
Testing port status with Test-NetConnection • 190	
Using netstat to identify local listeners • 191	
Summary	194
 Chapter 12: Unlock Your Exclusive Benefits	 195

Unlock this Book's Free Benefits in 3 Easy Steps.....	196
---	-----

Other Books You May Enjoy	200
----------------------------------	------------

Index	203
--------------	------------

Preface

Networks have a habit of disappearing when they work. A user opens an application, joins a meeting, reaches a cloud service, or connects from home without needing to think about the interfaces, addresses, switches, routers, wireless radios, firewalls, DNS servers, and protocols involved in making that interaction possible. It is usually only when something stops working that the network becomes visible.

Unfortunately, failures rarely introduce themselves clearly. A user does not normally report that a subnet mask is incorrect, a return route is missing, a wireless channel is congested, or a service is listening only on the wrong interface. They report that "the network is down," "the VPN won't connect," "the Wi-Fi is slow," or "I can't reach the server." The real skill in networking is therefore not simply knowing what each technology does. It is being able to take an imprecise symptom and turn it into a sequence of questions that can be answered with evidence.

That is the approach of this book.

We begin by building an end-to-end picture of network communication. You will see how endpoints, transmission media, switches, routers, firewalls, addresses, and transport protocols contribute to one exchange, and then use the OSI model to separate that exchange into functions that can be tested independently. IPv4 subnetting and CIDR develop the reasoning needed to decide whether a destination is local or remote, while NAT shows what changes when privately addressed systems communicate across another network. IPv6 extends those same concepts into a larger and increasingly important addressing model.

Once those foundations are established, the focus shifts from individual addresses to the infrastructure that carries them. VLANs, Layer 3 switching, routing protocols, firewalls, operating-system routes, and virtual switches all affect the path traffic is allowed to take. Rather than treating these as isolated configuration topics, the book continually returns to the same question: what path should this traffic follow, and what evidence confirms that it actually follows it?

That question becomes the basis of the troubleshooting method used through the second half of the book. You will learn to establish the impact and scope of an incident before changing anything, compare failing behavior with a known-working path, and use each successful or failed test to shrink the fault domain. Remote users, home networks, enterprise Wi-Fi, VPNs, DNS, routing, and firewalls may create different symptoms, but the underlying investigation remains consistent.

The later chapters make the evidence increasingly visible. Windows utilities, PowerShell, Linux commands, Nmap, and Wireshark provide different views of the same communication path. Their value comes from understanding their limits as well as their output. A successful ping does not prove that an application works. A listening socket does not prove that a remote client can reach it. A resolved DNS name does not prove that the service behind that name is available. A missing traceroute response does not automatically identify a failed router. Effective troubleshooting comes from correlating these observations rather than allowing any single result to become the diagnosis.

The final chapter brings those pieces together across an enterprise network. By that point, the aim is no longer to remember individual commands or definitions. It is to be able to map an expected path, select a test that answers one question, identify where observed behavior stops matching the design, make a controlled change, and verify the user's original operation from end to end.

Networking technologies will continue to change. Interfaces, management platforms, operating systems, wireless standards, and cloud services will evolve. The method in this book is designed to remain useful when they do: understand the path, establish the scope, gather evidence, change only what the evidence supports, and verify the result.

That is the habit this book is intended to build.

Who this book is for

This book is for IT professionals and learners who want to understand networks well enough to troubleshoot them methodically rather than rely on memorized fixes. It is particularly suited to help desk and desktop-support professionals, junior network and system administrators, infrastructure support engineers, and IT professionals whose roles increasingly require them to diagnose connectivity across endpoints, local networks, wireless environments, VPNs, servers, and cloud-connected services.

The book is also suitable for readers moving into networking from broader IT administration. You do not need to arrive with advanced routing, switching, or packet-analysis experience. The early chapters build the communication, addressing, subnetting, and protocol foundations that the later troubleshooting chapters depend on. Familiarity with Windows administration and basic command-line work will make the practical exercises easier to follow.

More experienced administrators can use the book as a structured refresher, particularly for connecting individual technologies into an end-to-end troubleshooting method. Rather than concentrating on one vendor or one network layer, the book repeatedly asks you to follow the user-to-service path, identify what is still working, and use evidence to determine which dependency deserves attention next.

What this book covers

Chapter 1, How Networks Communicate, introduces the end-to-end network path and the devices, media, protocols, and services that make communication possible. You will distinguish the roles of endpoints, switches, routers, firewalls, IP, TCP, and UDP, follow data from an application toward a remote service, and use `tracert` to make part of that otherwise invisible path visible.

Chapter 2, Seeing the Network One Layer at a Time, turns the broad communication path into a structured troubleshooting model using the seven layers of the OSI model. You will connect frames, packets, addresses, protocols, and ports to those layers, develop the binary foundations needed for IP addressing, and begin using Windows utilities to inspect configuration, reachability, address resolution, and routing decisions.

Chapter 3, How IPv4 Addresses Organize Network Traffic, develops IPv4 addressing into a practical routing and troubleshooting skill. You will use subnet masks and CIDR prefixes to identify network boundaries, calculate subnet capacity and usable ranges, determine whether a destination is local or remote, and follow private traffic through NAT, PAT, and port-forwarding boundaries.

Chapter 4, From Scarcity to Scale: How IPv6 Changes Addressing, extends the addressing model into IPv6. You will learn to read and compress hexadecimal addresses, interpret prefix lengths, distinguish global, unique-local, and link-local scope, understand common automatic-assignment mechanisms, and interpret IPv6 information on a Windows endpoint.

Chapter 5, How Network Infrastructure Moves and Protects Traffic, moves from addresses to the infrastructure that forwards and controls them. You will work with VLANs, Layer 3 switching, static and dynamic routing, endpoint routes, security controls, virtual networking, and cross-platform tools while following the same fundamental questions: where should the packet go, which device should forward it, and should the connection be permitted?

Chapter 6, A Systematic Approach to Network Troubleshooting, establishes the troubleshooting method used throughout the rest of the book. You will begin with severity, scope, and stakeholder impact, classify common failure patterns, trace access to corporate resources, and use addressing, DNS, gateways, ipconfig, ping, and tracert to narrow a broad symptom to a specific fault domain.

Chapter 7, Troubleshooting Remote and Home Networks, applies that method where the administrator has less visibility and control. You will trace remote-access, RDP, VPN, printing, discovery, Wi-Fi, router, DNS, and ISP dependencies while learning to use each successful test to decide whether the next investigation belongs on the endpoint, the home network, the corporate path, or farther upstream.

Chapter 8, Troubleshooting Enterprise Wireless Networks, scales wireless troubleshooting from one user to shared enterprise infrastructure. You will investigate security, coverage, interference, access-point capacity, heat maps, RF evidence, QoS, WMM, channel planning, and high-density environments while distinguishing poor coverage from contention and other performance problems.

Chapter 9, Troubleshooting VPNs and Remote Access, treats VPN connectivity as a sequence of dependencies rather than one feature that simply works or fails. You will compare VPN protocols, public-network restrictions, RRAS configuration, certificates, client addressing, WAN miniports, Windows client settings, DNS, routing, and error codes while separating failures that prevent a tunnel from forming from failures that appear after connection.

Chapter 10, Essential Network Troubleshooting Tools, builds a practical diagnostic toolkit across Windows, PowerShell, Linux, Nmap, and Wireshark. The emphasis is on choosing a test that answers the current question and understanding what its output does and does not prove, progressing from endpoint configuration and DNS through routing, socket state, port reachability, scans, and packet captures.

Chapter 11, Troubleshooting Routing, Switching, and Firewalls, brings the entire book together across an enterprise network. You will map Layer 2 and Layer 3 infrastructure, VLANs, firewalls, servers, remote sites, cloud resources, and other endpoints; follow an inter-VLAN exchange in both directions; and correlate remote TCP tests, local listener state, routing, and firewall policy to identify the first boundary where a connection stops working.

To get the most out of this book

You do not need advanced networking experience before starting this book. The early chapters build the foundations of network communication, addressing, subnetting, routing, and protocols before the later chapters apply those concepts to troubleshooting.

Basic familiarity with Windows, command-line tools, and common networking concepts will help, particularly once you reach the practical exercises. Some chapters also use Linux, macOS, Cisco-style networking, virtualization, VPNs, and packet-analysis tools, but you do not need access to every platform to follow the book. The specific software, hardware, permissions, and lab requirements are listed in each chapter.

Where exercises involve changing routes, VLANs, firewall rules, VPN settings, or other connectivity-related configuration, use an authorized test environment and record the existing configuration before making changes. The most useful habit throughout the book is to observe first, make one deliberate change, retest the original problem, and use the result as evidence for what to investigate next.

Download the color images

Your purchase includes a color, DRM-free PDF copy of this book, ideal for viewing color images, screenshots, and diagrams. Refer to *Free benefits with your book section* at the end of the *Preface* to unlock your PDF copy.

Conventions used

There are a number of text conventions used throughout this book.

CodeInText: Indicates code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles. For example: "Finally, we will use the Windows `tracert` command to expose part of that path. "

Any command-line input or output is written as follows:

```
tracert 8.8.8.8
```

Bold: Indicates a new term, an important word, or words that you see on the screen. For instance, words in menus or dialog boxes appear in the text like this. For example: "The **Internet Corporation for Assigned Names and Numbers**, or **ICANN**, coordinates key naming and identifier systems. "

Get in touch

Feedback from our readers is always welcome.

General feedback: If you have questions about any aspect of this book or have any general feedback, please email us at customercare@packt.com and mention the book's title in the subject of your message.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you reported this to us. Please visit <http://www.packt.com/submit-errata>, click **Submit Errata**, and fill in the form.

Piracy: If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at copyright@packt.com with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit <http://authors.packt.com/>.

Free benefits with your book


This book comes with free benefits to support your learning. Activate them now for instant access (see the *"How to Unlock"* section for instructions).

Here's a quick overview of what you can instantly unlock with your purchase:



DRM-Free PDF Version

Download DRM-free PDF and ePub copies of this book.



7-Day Packt Library Access

Get 7-day unlimited access to 8,000+ books and videos. No credit card required.

Available for first-time Packt+ trial users only.



Next-Gen Reader Access

Read this book on Packt Reader with progress sync, dark mode and note-taking.

How to Unlock

Scan the QR code (or go to packtpub.com/unlock). Search for this book by name, confirm the edition, and then follow the steps on the page.



UNLOCK NOW

Note: Keep your invoice handy. Purchases made directly from Packt don't require one

Share your thoughts

Once you've read *The Practical Network Engineer's Guide*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback.



<https://packt.link/r/1808495276>

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Part 1

Understanding How Networks Communicate

Before you can troubleshoot a network, you need a reliable way to picture what is happening when one system communicates with another. This part builds that picture from the ground up. You will begin with the devices, media, protocols, and services that make an end-to-end network path possible, then use the OSI model to separate that path into functions that can be examined individually.

From there, the focus moves to addressing. You will learn how IPv4 addresses and subnet masks define network boundaries, how CIDR and subnetting determine where traffic is local and when a gateway is required, and how NAT changes addresses as traffic crosses a network boundary. The final chapter extends that reasoning into IPv6, where a larger address space, hexadecimal notation, prefix lengths, and different address scopes change the operational details without changing the need to understand where an interface sits and what destinations it can reach.

By the end of this part, you will have the mental model used throughout the rest of the book: communication is an end-to-end path made up of smaller dependencies, and addressing tells each system how the next stage of that path should be reached.

This part of the book includes the following chapters:

- *Chapter 1, How Networks Communicate*
- *Chapter 2, Seeing the Network One Layer at a Time*
- *Chapter 3, How IPv4 Addresses Organize Network Traffic*
- *Chapter 4, From Scarcity to Scale: How IPv6 Changes Addressing*

1

How Networks Communicate

Most people notice a network only when it stops working. A document will not synchronize, a website will not load, or the Wi-Fi icon looks healthy while nothing useful moves. When the network works, it disappears into the background. That apparent simplicity depends on a chain of devices, transmission links, protocols, and services agreeing on how to move data from one place to another.

This chapter builds that chain from the ground up. We will start with the reason networks exist, identify the components required for communication, and examine the qualities a modern network must deliver. We will then look at the Internet as a collection of interconnected networks and follow data from an application, through a local network and an Internet service provider, to a remote service. Finally, we will use the Windows `tracert` command to expose part of that path. The goal is to give you a practical mental model that will support the addressing, routing, and troubleshooting work later in the book.

After going through this chapter, you will understand why individuals and organizations connect devices through networks and be able to identify endpoints, servers, transmission media, and the role of an operating system's network stack. You will also understand how IP, TCP, and UDP support network communication; distinguish the Internet from the World Wide Web; explain how standards, names, addresses, and other shared resources are coordinated; trace the basic path of data from an application to a remote destination; and run `tracert` and interpret the main parts of its output.

We will be covering the following main topics:

- Why networks exist
- What a network needs
- What modern networks must deliver
- Where the Internet is
- How data travels over IP

Note

Download the PDF version of this book

Your purchase includes a DRM-free PDF copy of this book, the code bundle, and a range of exclusive benefits. To unlock everything, follow the *Free benefits with your book* section in the *Preface*.

Technical requirements

This chapter is mainly conceptual, so no software installation is required. A sheet of paper or a diagramming application such as Microsoft Visio is useful if you want to sketch the paths described in the chapter. To follow the final example, you need a Windows computer with access to Command Prompt and an active network connection.

Why networks exist

A network allows one device to reach data, applications, and services held somewhere else. The resource may be on a server in the next room, in another office, or in a cloud data center thousands of miles away. From the user's point of view, the useful outcome is the same: the resource is available without the user having to know its physical location.

Sharing is only part of the job. A network also needs to keep services available when equipment or a connection fails. Important data is often stored or replicated in more than one location so that another copy remains reachable if the first location becomes unavailable. Cloud platforms such as Microsoft Azure and Amazon Web Services use this principle on a large scale. In network design, we describe the broader goal as **high availability**. A highly available system reduces the effect of a single failure by providing another path, another device, or another copy of the service.

Figure 1.1 shows how a load balancer can preserve access to a shared service by directing requests away from an unavailable instance and toward one that remains healthy.

Keeping a shared service available

A load balancer sends new requests only to service instances that pass their health checks.




Figure 1.1: Maintaining access to a shared service after an instance failure

Figure 1.1 shows users reaching shared service data through a network gateway and a firewall with load-balancing capabilities. When service instance A fails its health check, the load balancer stops directing requests to it and sends them to the healthy service instance instead. This illustrates why availability depends on more than making a service reachable under normal conditions: the design must also provide another working instance, path, or copy of the data when a component fails.

This distinction becomes important during troubleshooting. Reaching a service once does not prove that the design is reliable. We also need to ask what happens when a cable is cut, a router fails, or a server goes offline. Networking therefore concerns both communication and continuity.

Meeting both requirements depends on a chain of devices, links, protocols, and software functions working together. Understanding the role of each component turns the network from a black box into a path that can be designed, tested, and troubleshot systematically.

What a network needs

At its simplest, a network needs at least two devices that can exchange data. Those devices do not have to be traditional computers. A phone, tablet, printer, server, sensor, or other IP-capable device can participate. Each device needs a way to connect, a common set of communication rules, and enough addressing information to send data toward the correct destination.

To build a useful mental model, we will separate those requirements into the devices involved, the path between them, the transmission media they use, the protocols that govern the exchange, and the operating-system software that brings those elements together. Understanding these roles will help you locate a failure instead of treating the network as a single black box.

Endpoints servers and network infrastructure

An **endpoint** is a device that consumes or produces data. A laptop opening a website, a phone checking email, and a workstation saving a file are all endpoints in that interaction. A server provides a resource, such as a website, application, file, image, or database record. These roles are not permanent labels. A computer can act as an endpoint in one exchange and provide a service in another.

Communication between the endpoint and server normally depends on other network devices. Switches connect devices within a local network, while routers move traffic between different networks. A firewall controls traffic according to security rules. The exact design changes with the environment, but the practical question remains the same: which devices and links must the data cross to reach its destination?

Distinguishing these functions is the first step toward reading a network path correctly, because the same physical device may perform several jobs while each job can fail for a different reason.

How infrastructure devices divide the work

The devices in a network path do not all perform the same job. A network interface is the endpoint's connection to the network. It works with the operating system and the physical connection, such as an Ethernet cable or Wi-Fi radio, to send and receive data.

A switch moves Ethernet frames within a local network. It learns which devices are reachable through its ports and forwards local traffic toward the appropriate connection. A wireless access point performs a related function for wireless devices by connecting them to the local network over radio. The access point commonly bridges that wireless traffic into the wired network.

A router connects different IP networks. It examines the destination IP address and selects the next network or router that should receive the packet. On a small network, the default gateway is usually the router that carries traffic away from the local network. A firewall performs another job: it evaluates traffic against security rules and decides which communication should be permitted or blocked.

These roles may be combined. A home router often contains a switch, wireless access point, router, and firewall in one device. Enterprise networks are more likely to separate the functions. When troubleshooting, concentrate on the function rather than the number of boxes. A working link to a switch does not prove that the router has a route, the firewall permits the traffic, or the destination service is available.

Once those functions are separated, you can stop viewing the network as a collection of boxes and start following the actual path that traffic must take from its source to its destination.

Reading the network as a path

When someone reports that the network is down, the entire network is rarely one indivisible problem. Communication begins at a source, ends at a destination, and depends on a series of components between them. A laptop may connect to a wireless access point, pass through a switch, reach a router and firewall, cross an Internet provider, and finally arrive at a remote service. A failure at any one of those points can produce the same complaint from the user's perspective: the application does not work.

A useful early habit is to sketch the end-to-end network path you are investigating before changing anything. On that sketch, mark the source device, the destination service, the transmission medium, and every known device or network boundary between them. Then ask where communication is still confirmed and where it stops. You may not yet have the tools to answer every question, but this path diagram keeps the investigation tied to the actual route instead of turning it into a series of guesses. Later chapters will add addresses, routing tables, command-line utilities, packet analysis, and structured troubleshooting to this same model. Because a single path can cross several types of connection, you also need to recognize how wired and wireless media carry different parts of the same exchange.


Wired and wireless transmission

The connection can be wired, wireless, or a combination of both. Ethernet is common on local wired networks, and Wi-Fi connects devices over radio. Cellular services such as 5G and satellite links can also carry IP traffic. A phone may begin on Wi-Fi, pass through a wired router, cross an Internet provider's fiber network, and arrive at a server through another wired connection. The medium can change several times while the communication remains part of one end-to-end exchange.

Figure 1.2 follows two common local connections—a desktop using Ethernet and a laptop using Wi-Fi—as both paths converge on the router and continue toward a remote service.

Network components and transmission media

Traffic crosses devices with different responsibilities before it reaches a remote service.



The role of the device matters more than whether several roles share one physical box.

Figure 1.2: Wired and wireless paths through common network components

The desktop reaches the switch over Ethernet, while the laptop reaches the access point over Wi-Fi. Both local paths then feed into the router and firewall, cross the provider network, and reach the cloud service. The medium and local device may change, but IP communication continues across the complete path.

Physical connectivity can place a device on the path, but it cannot by itself tell the data where to go or how the exchange should behave. Those decisions depend on the protocols carried over the connection.

IP and transport protocols

Internet Protocol, or **IP**, provides addressing and routing information so that data can move across interconnected networks. IP gets the data toward the destination, but applications also need rules for how the exchange should behave. That is where transport protocols such as **Transmission Control Protocol (TCP)** and **User Datagram Protocol (UDP)** enter the picture.

TCP is connection-oriented. It establishes a communication session and provides ordered, reliable delivery. If data is lost during transmission, TCP can retransmit it. This behavior suits many applications that need complete data to arrive in the correct order.

UDP has less protocol overhead and does not provide the same delivery, ordering, or retransmission mechanisms. Applications may choose UDP when low delay is more important or when the application supplies any recovery behavior it needs. Real-time applications and games commonly use it, but UDP is not simply an insecure version of TCP. Security is a separate concern that is provided by the protocols and controls around the transport. When you compare TCP and UDP, focus first on delivery behavior, overhead, and application requirements.

Figure 1.3 illustrates the practical difference by showing how TCP and UDP respond when one unit of data does not reach the receiver.




Figure 1.3: TCP retransmission and UDP delivery after data loss

On the TCP side, the missing data is detected through the acknowledgement process and sent again so that the receiver can obtain the required sequence. On the UDP side, later datagrams can continue without the transport protocol retransmitting the missing one; any recovery must come from the application. The figure compares delivery behavior only and does not imply that either protocol provides encryption by itself.

That distinction matters because reliability and security solve different problems. Choosing TCP for dependable delivery or UDP for lower overhead does not remove the need to protect the application data.

Transport behavior and security are separate decisions

TCP and UDP describe how applications transport data, but the transport protocol alone does not determine whether that data is secure. TCP provides connection management, ordered delivery, and retransmission, but it does not automatically encrypt the application's data. An application must use an additional security protocol, such as **Transport Layer Security**, when confidentiality and authentication are required.

UDP does not provide TCP's built-in delivery and ordering mechanisms, but that does not make UDP inherently insecure. Applications can add security, reliability, and connection behavior above UDP. HTTP/3, for example, uses QUIC over UDP and includes TLS protection as part of the transport design.

The application's needs usually drive the transport choice. A file transfer or database exchange commonly benefits from reliable, ordered delivery. Voice, video, and interactive applications may place greater value on low delay because retransmitting data that is already too late to use may not improve the experience.

During troubleshooting, treat transport behavior and security as separate questions. First determine whether traffic can reach the endpoint and whether the expected transport exchange is occurring. Then determine whether the application is authenticating the other party and protecting the data. A successful TCP connection confirms that a transport path was established, but it does not prove that the application data is encrypted or that the remote service should be trusted.

These protocol choices still need software on each endpoint to implement them, manage the network interface, and pass data between the application and the network. That coordinating role belongs to the operating system's network stack.

Operating systems and network stacks

A device also needs software that knows how to use its network interface and communication protocols. Modern operating systems provide this network stack. Windows, macOS, Linux, Unix, Android, and iOS can all communicate over IP through wired or wireless interfaces. Servers may run specialized editions or distributions, but the essential requirement is not a particular brand of network operating system. It is a functioning network stack, a configured interface, and compatible protocols at both ends.

At this point, you can identify the devices, media, protocols, and software that allow data to move between endpoints. A working collection of components, however, is only the baseline; users and organizations also depend on the network to protect traffic, perform consistently, remain available, and control access without unnecessary friction.

What modern networks must deliver

A network can be connected and still be poor at its job. Modern networks are judged by whether they protect data, perform consistently, remain available, and let users reach the resources they are authorized to use. They also have to achieve those results without unnecessary cost or complexity. This section examines those operational expectations through security, performance, reliability, automation, usability, and cost. Together, they provide a more realistic test of whether a network supports the people and services that depend on it.

Security

Security is an operating requirement, not a feature to add after the network is built. Data moving between devices can be exposed to theft, misuse, malware, and ransomware. The network design therefore needs controls that reduce unauthorized access and limit the damage when a device is compromised. Secure communication also depends on maintaining systems and applying updates. A fast connection is not useful if it gives an attacker the same easy access as the intended user.

Performance reliability and automation

Users expect to connect quickly, remain connected, and receive adequate performance for the application they are running. Reliability means more than speed. It includes stable links, available services, and recovery options when part of the environment fails. Routine maintenance and updates should be automated where appropriate so that essential work does not depend on someone remembering every task at the right moment.

Access without unnecessary complexity

People now expect to reach data from offices, homes, and other locations. That access must be secure, but it must also be usable. If a person needs several confusing applications and repeated manual steps just to reach an approved resource, the design has created another operational problem. Good networking hides much of the underlying complexity while still giving administrators the visibility and control needed to support it.

Useful data at a manageable cost

Networks should save time and make useful information available without requiring unreasonable spending. They do not, however, turn clutter into useful data. They only move the clutter faster. Storage fills up when people keep information they no longer need, and retaining everything increases cost and management effort. The network supports access; sound data management determines whether what users reach is worth keeping.

The same trade-off applies to infrastructure. Running equipment on premises can provide control, but it also introduces maintenance, security, programming, and hardware-lifecycle risks. Using remote or cloud services changes those responsibilities rather than removing them. Network engineers must understand the path to the service, the dependencies along that path, and the effect of a failure at each point.

A modern network therefore has to balance connectivity with security, reliability, usability, and cost. Those same concerns continue beyond the local environment, where traffic becomes part of the much larger system of interconnected networks that we call the Internet.

Where the Internet is

I once asked a group of executives to point to the Internet. No one could. I asked them to hold up their phones and told them, "The Internet is right there in your hands." The phone is not the entire Internet, of course. It is an endpoint connected to a vast collection of networks. The exercise makes the important point that the Internet is not a single machine, building, or cloud. It exists through the connections among many independently operated networks.

The World Wide Web is one service that uses the Internet. Email, file transfer, voice and video communication, cloud storage, and many other services also use the same underlying connectivity. A private IP network can use the same protocols without being directly reachable from the public Internet, so an IP-connected device is not automatically public.

Seen this way, the Internet becomes less abstract: it is a path between endpoints and services, supported by many independently operated networks. Its distributed design creates resilience, but the availability experienced by one user still depends on every critical part of that path.

Endpoints services and resilience

From a practical point of view, Internet communication still involves endpoints requesting resources and systems providing them. The request crosses local equipment, an Internet service provider, other connected networks, and the destination network. The Internet is resilient because it is distributed and can offer alternate routes, but that does not make every connection failure-proof. A local access link, a provider outage, a routing problem, or a failure at the destination can still interrupt a user's service.

That combination of independent networks and shared communication rules did not appear all at once. Understanding how it developed helps explain why the Internet has no single owner and why interoperability remains central to its design.

How the Internet developed

The Internet grew from research into packet-switched communication and the problem of connecting different networks. **ARPANET**, funded by the US Department of Defense's Advanced Research Projects Agency, connected research institutions and demonstrated that computers could exchange data across a packet-switched network. Later work on TCP/IP allowed separate networks to interconnect, and ARPANET moved to TCP/IP in 1983. Academic, government, and commercial networks continued to expand, and public use grew rapidly during the 1990s.

The architecture did not emerge from one company handing a finished system to the world on a single date. It developed through research, shared protocols, interoperable equipment, and cooperation among many organizations. The Web arrived later and made information easier to publish and navigate, which helped bring the Internet into everyday use.

Growth through cooperation created another requirement: someone still had to coordinate standards and globally unique resources so that independently operated networks could communicate without conflicting definitions or identifiers.

Who coordinates standards names and addresses

No single organization controls every network on the Internet. Coordination still matters, because devices must agree on protocols and globally unique identifiers. The **Internet Engineering Task Force**, or **IETF**, develops and maintains many of the technical standards used on the Internet. The **Institute of Electrical and Electronics Engineers**, or **IEEE**, develops standards that include the 802.11 family used for wireless local area networking.

The **Internet Corporation for Assigned Names and Numbers**, or **ICANN**, coordinates key naming and identifier systems. The **Internet Assigned Numbers Authority** functions coordinate global Internet number resources and the Domain Name System root. IP address blocks are distributed through a hierarchy that includes Regional Internet Registries and network operators. These organizations coordinate shared resources; they do not operate every network or decide how every country, provider, or organization permits Internet access. Governments and network owners may still apply legal, policy, and technical restrictions within the environments they control.

The names of these organizations are easier to remember when you connect each one to the specific problem it solves. Separating protocol standards, local-network standards, domain names, number resources, and network operation prevents their responsibilities from being blurred together.

How the responsibilities fit together

Internet coordination becomes easier to understand when you separate standards, names, addresses, and daily network operation. These responsibilities are related, but they move through different organizations and processes.

Area	Primary responsibility	Practical result
Internet protocols	The IETF develops technical specifications and operational guidance through the RFC process.	Vendors and software developers can implement compatible protocols for addressing, routing, transport, security, and applications.
Local network standards	IEEE develops standards such as 802.3 for Ethernet and 802.11 for wireless networking.	Network interfaces, switches, and access points from different vendors can use compatible link technologies.
Domain-name coordination	ICANN coordinates important parts of the global naming system. Registries maintain information for top-level domains, while registrars process domain registrations for customers.	An organization can register a domain name and publish DNS records that direct users toward its services.
Internet number resources	The IANA functions manage global pools of IP addresses and autonomous system numbers. Regional Internet Registries distribute resources within their regions under applicable policies.	Providers and qualifying network operators can obtain globally unique addressing and routing resources.
Network operation	Service providers, businesses, governments, and other network owners configure and operate their own systems.	Each operator decides how to route, secure, monitor, and provide access through its network.

Table 1.1: How Internet coordination responsibilities are divided

A company registering a domain name does not receive it from the IETF, and a protocol standard does not assign an organization an IP address. Each system handles a different part of the larger process. Together, they provide common technical rules and unique identifiers while allowing individual organizations to operate their own networks.

The Internet depends on shared standards, coordinated identifiers, and independently operated networks. Turning that structure into a practical troubleshooting model requires following what happens to data as it leaves an application, crosses those networks, and reaches its destination.

How data travels over IP

The easiest way to connect the chapter's ideas is to follow one piece of data. Suppose you are working in Microsoft Word and the document is synchronized to cloud storage. Saving a local copy writes data to storage. Synchronizing the document starts a network operation: the application gives data to the operating system's network stack, transport and IP information are added, and the network interface prepares the data for transmission over the available link.

The first part of that journey takes place before any signal reaches the cable or radio, as the host converts application data into forms that each network layer can process.

How the host prepares data

Before the first signal leaves a device, the application's data passes down through the network stack. Each stage adds the control information needed for its part of the delivery process. This is called **encapsulation**.

At the transport layer, TCP or UDP adds information such as source and destination port numbers. These numbers help the receiving operating system deliver the data to the correct application. TCP also adds information used to manage its connection and provide ordered, reliable delivery. TCP carries data in segments, while UDP carries application messages in datagrams.

The network layer places the transport information and its data inside an IP packet. The IP header contains source and destination addresses along with other information needed to move the packet between networks. The data-link layer then places the IP packet inside a frame for the current local link. On an Ethernet network, that frame includes local link-layer addressing used to reach the next device on the path.

The network interface transmits the frame as signals over the selected medium. When a router receives it, the incoming frame does not continue unchanged to the destination. The router removes the frame information, examines the IP packet, selects the next hop, and places the packet into a new frame suitable for the next link.

The destination reverses the process. It receives the frame, extracts the IP packet, processes the transport information, and passes the application data to the correct program. This explains why an application does not communicate directly with a cable or radio. Several layers cooperate to prepare, carry, route, receive, and deliver the data.

With the data now prepared for transmission, the focus can move from the work inside the host to the devices and networks that carry it toward the destination.

From the local device to the destination

On a typical network, the device first sends traffic across Ethernet or Wi-Fi. Local switches or a wireless access point carry it toward the default gateway. A firewall may inspect the traffic, and the router forwards it to the Internet service provider. From there, routers move the data across interconnected networks until it reaches the destination network. The destination may be a cloud service, a server, storage, or another endpoint.

The word cloud can make this path sound mysterious, but the cloud is still made of networks, routers, servers, and storage. Data entering a cloud provider crosses that provider's edge, passes through its network, and reaches the service that will process or store it. At the destination, the receiving network stack passes the data upward to the application. The same protocol layers that prepared the data for transmission help the destination interpret it.

IP operates at the network layer, layer 3 of the **Open Systems Interconnection (OSI)** model. TCP and UDP operate at the transport layer, layer 4. The full OSI model contains seven layers and gives us a structured way to discuss where communication is working and where it is failing. We will examine that model in the next chapter. For now, remember that an application does not place raw data directly onto a cable. Several software and hardware functions prepare, address, transmit, route, receive, and deliver it. *Figure 1.4* brings these operations together by showing encapsulation at the sender, forwarding across the network path, and decapsulation at the destination.

End-to-end packet journey

Encapsulation prepares the data for transmission; decapsulation delivers it to the receiving application.




Figure 1.4: Encapsulation, forwarding, and decapsulation across an end-to-end path

The sender adds transport, IP, and local-link information before placing the frame on the network. Switches and access points carry the frame locally, while routers examine the IP packet and create a new frame for each succeeding link. At the destination, the process is reversed until the original application data reaches the correct program.

The full process is mostly invisible to the user, but network engineers need a way to expose at least part of the routed path. On Windows, `tracert` provides that first layer 3 view and helps establish where traffic is receiving responses.

Using tracert to see the route

Windows includes the `tracert` command, which reveals the sequence of layer 3 hops toward a destination. Open Command Prompt and run the following command to trace a route to Google's public DNS service at 8.8.8.8:

```
tracert 8.8.8.8
```

A sample run produced the following output:

```
Tracing route to dns.google [8.8.8.8]
over a maximum of 30 hops:

  1  <1 ms  <1 ms  <1 ms  router.charter.net [192.168.1.1]
  2   9 ms  10 ms   8 ms  22.54.156.1
  3  14 ms  13 ms  11 ms  acr01csbyor-gbe-1-44.csby.or.charter.com
[96.34.105.202]
  4  15 ms  11 ms  11 ms  dtr01lnbhwa-tge-0-0-0-0.lnbh.wa.charter.com
[96.34.111.208]
  5  17 ms  22 ms  17 ms  096-034-108-175.biz.spectrum.com [96.34.108.175]
  6  21 ms  22 ms  17 ms  096-034-108-160.biz.spectrum.com [96.34.108.160]
  7  31 ms  31 ms  38 ms  bbr02sttlwa-bue-4.sttl.wa.charter.com [96.34.0.56]
  8  25 ms  32 ms  27 ms  prr01sttlwa-bue-2.sttl.wa.charter.com [96.34.3.39]
  9  25 ms  25 ms  27 ms  72.14.194.48
 10  25 ms  25 ms  26 ms  108.170.227.7
 11  27 ms  27 ms  32 ms  142.251.50.177
 12  31 ms  27 ms  27 ms  dns.google [8.8.8.8]

Trace complete.
```

Each numbered line represents a responding router along this route. The three time values are round-trip measurements for the probes sent to that hop, followed by the router's resolved name and IP address when that information is available. The first hop is the local router. Later hops belong to the provider and other networks on the path, and the final line shows the destination.

The phrase maximum of 30 hops describes the command's current limit for this trace, not a rule that all Internet traffic must cross 30 routers. Your output will differ because routes depend on your location, provider, destination, and the state of the network at that moment. Treat a trace as a snapshot of one path. It helps you see where traffic travels and gives you a starting point when you later investigate delay or loss.

That visibility is useful only when the output is interpreted with its limitations in mind. A missing reply or a slow intermediate hop can reflect how a router handles diagnostic traffic rather than a failure in the application path.

What tracert does not prove

A tracert result is useful evidence, but it is not a complete measurement of everything happening between the source and destination. An asterisk (*) means that tracert did not receive the expected response before its timeout expired. It does not automatically mean that the router is down or that ordinary traffic cannot pass through it. Some routers forward traffic normally while filtering or limiting the diagnostic responses used by tracert.

The displayed time is also a round-trip measurement. It includes the path to the responding router and the return path to your computer. Those paths may not be identical. A slow response from one hop therefore does not prove that the router itself is causing an application problem. If later hops respond normally, the earlier router may simply be giving low priority to diagnostic traffic.

Name resolution can add delays while tracert attempts to convert addresses into router names. Load balancing may also cause separate traces, or even separate probes in the same trace, to encounter different routers. Firewalls and destination systems can filter the probes, causing a trace to stop even though the destination remains accessible through the required application protocol.

Use tracert to identify the approximate layer 3 path and to look for patterns that deserve further investigation. Compare its results with the user's symptoms, repeated measurements, and tests of the actual destination service. A single slow hop or timeout should be treated as a clue rather than proof of the fault.

Summary

In this chapter, we built a practical model of network communication. Networks connect endpoints to shared data, applications, and services. They rely on transmission links, network interfaces, operating-system network stacks, IP addressing, and transport protocols. TCP provides connection-oriented, ordered, reliable delivery, while UDP offers a lighter transport with different delivery guarantees. Security must be provided by the wider protocol and system design rather than inferred from the choice between TCP and UDP.

We also separated the Internet from the World Wide Web, examined the roles of endpoints and servers, and saw how technical standards, names, and number resources are coordinated. Finally, we followed data from an application through the local network and across routed networks, then used `tracert` to expose the layer 3 hops toward a destination.

You now have the broad path in view: an application creates a request, the host prepares it, network devices forward it, and a remote system receives it. In the next chapter, we will examine the OSI model, binary and hexadecimal numbering, and essential network utilities. Those tools will give us a more precise vocabulary for describing each stage of the path and diagnosing what happens when communication fails.

In the next chapter, we will turn this broad path into a more precise troubleshooting model by examining the seven layers of the OSI model and the roles they perform. You will also learn how binary and hexadecimal notation represent network values and use essential utilities to test connectivity and identify where communication begins to fail.

2

Seeing the Network One Layer at a Time

A user who cannot reach a service may describe the problem in one sentence: the network is not working. The cause, however, could be a disconnected cable, an incorrect local address, a missing route, a blocked protocol, or an application that is not responding. Effective troubleshooting depends on separating that end-to-end exchange into smaller functions that can be tested one at a time.

This chapter develops that troubleshooting vocabulary. We will use the **Open Systems Interconnection (OSI)** model to organize network communication, connect IPv4 addresses to the binary values computers process, and apply Windows utilities that reveal configuration, reachability, address resolution, network paths, and local routing decisions. The aim is not to memorize seven isolated layer names or a list of commands. It is to understand what each layer and tool can tell you, and just as importantly, what it cannot prove.

In this chapter, we will examine the OSI model and its seven layers, learn how Wireshark can be used to analyze network protocols, and explore IPv4 addressing and binary representation. We will also consider why shared protocols are essential to network communication before working with the connectivity and routing tools used to investigate common network problems.

After going through this chapter, you will be able to explain the purpose of each OSI layer, relate frames, packets, segments, addresses, and ports to the path introduced in *Chapter 1*, and describe how shared protocols allow different systems to communicate. You will also be able to interpret the basic structure of an IPv4 address, convert an octet between binary and decimal, inspect a Windows network configuration, test IP reachability, examine the ARP cache, trace a routed path, and read or add an entry in the local routing table.

This chapter contains the following main headings:

- The OSI model and its seven layers
- IPv4 addressing and binary representation
- Essential network connectivity and routing tools

Technical requirements

The conceptual sections do not require additional software. To follow the command examples, use a Windows computer with Command Prompt and an active network connection. The route command that adds a persistent entry must be run from an elevated Command Prompt, and the gateway value must be replaced with an address that is valid for your environment.

Wireshark is optional; the chapter introduces what a protocol analyzer shows but does not require you to perform a packet capture.

The OSI model and its seven layers

The OSI model is a reference model for discussing how data moves from an application on one device to an application on another. It divides communication into seven layers. Each layer describes a related set of responsibilities, from the signals on a cable or radio link at the bottom to the network services used by applications at the top. Real protocol suites do not always fit those boundaries perfectly, but the model gives administrators a common way to locate and describe a problem.

When a device sends data, the information moves down through these functions. Each stage prepares the data for the next part of the journey by adding the control information it needs. The receiving device processes the information in the opposite direction and ultimately delivers the original data to the correct application. The following figure presents the seven layers in that order.

Open System Interconnection (OSI) Model




Figure 2.1: The OSI model divides network communication into seven related layers

The figure separates the application, presentation, and session layers from the transport, network, data link, and physical layers. For troubleshooting, this separation helps us move from the visible symptom toward the underlying dependency. If an application cannot communicate, we can ask whether the service is functioning, whether a transport exchange is possible, whether a route exists, whether the local frame can reach the next device, and whether the physical link is active.

That layered view helps us locate a problem, but communication succeeds only when both devices interpret the same rules, which makes shared protocols the practical link between the model and a working network.

Why shared layers and protocols matter

A network works only when both ends agree on the rules for communication. Early computer systems and later commercial vendors used different protocols, which made interoperability difficult. The development and adoption of shared protocol suites, particularly TCP/IP, allowed systems from different manufacturers to exchange data across interconnected networks. The

OSI model does not replace those protocols; it provides a structured language for describing the functions they perform.

That distinction matters in practice. A protocol is an implemented set of rules, while a model is a way to organize and discuss those rules. Ethernet, IP, TCP, UDP, and application protocols perform real work. The OSI layers help us identify where that work occurs and choose a test that examines the right part of the path.

With a common set of rules established, troubleshooting can begin at the most fundamental dependency: whether the device can send and receive signals on its local connection.

Physical layer functions and initial host checks

Layer 1 is the physical layer. It covers the signals and media used to carry bits, including copper and fiber cabling, connectors, network interface hardware, and wireless radio transmission. A link light, cable test, interface status, or wireless signal reading can provide evidence about this layer. These checks answer a basic question: can the device transmit and receive signals on the local connection?

Physical-layer checks establish whether signals can travel across the local connection, but they do not confirm that the host's TCP/IP implementation is functioning. A loopback test provides that separate evidence. First, use the following command to view the device's current IP configuration:

```
ipconfig
```

The output identifies the IPv4 and IPv6 addresses, subnet mask, and default gateway assigned to each adapter. An address such as 192.168.1.111 is a logical address. It can be changed by manual configuration or assigned dynamically. It should not be confused with the adapter's **Media Access Control (MAC)** address, which serves a different purpose at layer 2.

Next, use the IPv4 loopback address to test whether the local host can process IP and ICMP traffic through its own network stack:

```
ping 127.0.0.1
```

A successful reply confirms that the local TCP/IP stack is responding. It does not send traffic through the Ethernet cable, Wi-Fi radio, switch, router, or Internet connection, so it cannot prove that those components work. Together, the physical checks and the loopback test separate a link problem from a failure inside the host's own network stack. Once those possibilities have been distinguished, troubleshooting can move to the layer 2 mechanisms that deliver frames across the local link.

Data link layer components and operations

Layer 2 is the data link layer. On an Ethernet network, it organizes data into frames and uses MAC addresses to deliver those frames across the current local link. A network adapter normally has a factory-assigned MAC address, although many operating systems and virtual environments can override it in software. MAC addresses should be unique within the same local network. Duplicate addresses can cause unstable or misdirected communication because the network cannot reliably determine where the device is located.

A switch operates mainly at layer 2. It learns the source MAC addresses seen on its ports and uses that information to forward known unicast frames toward the appropriate port. It still forwards broadcasts, and it may flood traffic when the destination is unknown, but it avoids sending every known unicast frame to every connected device. This is a major improvement over a hub, which repeats incoming signals to all of its ports. Bridges perform a similar layer 2 forwarding role on a smaller scale and now appear more often in certification material than as separate devices in modern networks.

Frames are local to a link. When traffic crosses a router, the layer 2 information is removed and replaced for the next link. The IP packet inside the frame continues toward the destination, which leads directly to the network layer.

Network layer routing and devices

Layer 3 is the network layer. It uses logical addresses, such as IPv4 addresses, and provides the information routers need to move packets between different IP networks. A switch can deliver a frame within the local network, but communication from a 192.168.1.0 network to a 192.168.2.0 network requires a layer 3 device with a route between them.

Routers are the devices most closely associated with layer 3, although many enterprise switches can perform both layer 2 switching and layer 3 routing. When a router receives a frame, it examines the destination IP address in the enclosed packet, consults its routing table, and selects the next hop or outgoing interface. It then places the packet inside a new frame for the next local link. This repeated process is how one IP packet can cross several different networks and media types.

Transport layer protocols and ports

Layer 4 is the transport layer. TCP and UDP carry application data over IP and use port numbers to identify the sending and receiving application services. An IP address identifies a host interface on a network; a port number helps the operating system deliver the arriving data to the correct process on that host.

TCP is connection-oriented and provides mechanisms for ordered delivery, acknowledgments, and retransmission. UDP sends datagrams without TCP's built-in connection and recovery

behavior. Firewalls and access control lists can permit or block traffic according to addresses, protocols, and ports. An open port should therefore be understood as a permitted path to a listening service, not as a physical doorway. Exposing unnecessary services increases the attack surface, but the transport protocol alone does not determine whether application data is secure.

The first four layers account for the media, local delivery, routed delivery, and end-to-end transport of data. The upper three layers describe how applications establish conversations and represent the information they exchange.

Session presentation and application functions

Layer 5 is the session layer. It describes functions that establish, maintain, coordinate, and end conversations between applications. Dialog control and recovery checkpoints are commonly associated with this layer. In modern TCP/IP implementations, these responsibilities may be handled by the application, a library, or another protocol rather than by one clearly separated session-layer protocol. The OSI model still gives us a useful name for the function.

Layer 6 is the presentation layer. It concerns the form in which information is exchanged, including character encoding, data formatting, compression, and encryption. Its purpose is to ensure that the receiving system can interpret the data. Encryption used by an application can be discussed at this layer, but technologies such as virtual private networks may span several functions and should not be forced into one layer without considering the actual protocol.

Layer 7 is the application layer. It provides network services that applications use. Web browsers, office applications, and other programs do not simply become layer 7 in their entirety; their network communication uses application-layer protocols and services. This distinction keeps the model tied to the communication path rather than to the icon a user opens on the desktop.

Because real implementations distribute these functions across protocols and software components, a packet capture provides the clearest way to see how the layers appear in actual communication.

Seeing protocol layers with Wireshark

A **protocol analyzer** such as Wireshark makes the layered structure visible. A captured Ethernet frame can be expanded to reveal its layer 2 addresses, the IP packet carried inside it, the TCP or UDP information inside the packet, and the application data that follows. The analyzer does not merely label a problem as layer 2 or layer 3. It exposes the fields and values that let an administrator determine whether the observed communication matches what the network should be doing.

Reading a captured packet from the outside in

When you select a captured packet in Wireshark, the packet details pane displays a series of expandable sections. The Frame section contains information about the capture itself, while the sections below it represent the protocols that Wireshark identified. A typical packet may contain Ethernet II, Internet Protocol Version 4, Transmission Control Protocol, and an application protocol. Expanding these sections from top to bottom allows you to follow the packet through the layers rather than treating it as a single piece of data.

Begin with the Ethernet II section. Its source and destination fields contain the MAC addresses used to deliver the frame across the local network. If the destination is on the same network, the destination MAC address normally belongs to the receiving device. If the destination is on another network, it normally belongs to the local router or default gateway. This distinction explains why the destination IP address can identify a remote server while the destination MAC address identifies a nearby router: IP handles delivery between networks, whereas Ethernet handles delivery across the current local link.

Next, expand the IPv4 section and compare the source and destination IP addresses with the communication you expected to see. Then examine the TCP or UDP section. The port numbers identify the application endpoints involved in the exchange, while TCP flags can show whether a connection is being established, acknowledged, or closed. For example, if a client repeatedly sends a TCP SYN but no SYN-ACK returns, the capture confirms that the client attempted to begin the connection. It does not, by itself, prove whether the response was blocked by a firewall, lost because of a routing problem, or never generated because the service was unavailable.

This layered reading method prevents a common troubleshooting mistake: focusing on the application before confirming that the lower layers are functioning. Wireshark cannot determine every cause from a single packet, but it can show how far the communication progressed and which part of the exchange did not occur as expected.

Taken together, the OSI layers provide a disciplined way to narrow a fault. We can start with the physical connection, move through local and routed delivery, inspect the transport exchange, and finally evaluate the application service. The next step is to look more closely at the layer 3 address that appears throughout that process.

IPv4 addressing and binary representation

IPv4 uses 32-bit addresses. Those 32 bits provide 2 to the 32nd power possible values, although not every value can be assigned as an ordinary public host address. Some ranges serve private networks, loopback, multicast, and other defined purposes. To make the address easier for people to read, the 32 bits are divided into four 8-bit octets and written as decimal numbers separated by periods. This format is called dotted-decimal notation.

An address may appear as 192.168.0.1. Leading zeros are normally omitted, so an octet written conceptually as 001 is displayed as 1. That shortening applies within each IPv4 octet; it does not remove octets from the address. IPv6 uses a different colon-separated notation and different rules for compressing zeros, which we will examine in *Chapter 4*.

An IPv4 configuration normally presents the address with a subnet mask and default gateway. The address identifies the host interface, the subnet mask helps the host determine which destinations are local, and the default gateway provides a route toward destinations outside that local network. Later chapters will develop these relationships through subnetting and routing. For now, the important point is that every decimal octet represents exactly eight binary positions.

Converting an octet between binary and decimal

Binary uses two digits: 0 and 1. At the physical level, those values can be represented by two distinguishable states, such as on and off. People usually read addresses in decimal, but computers process the underlying bit patterns. Understanding those patterns makes subnet masks, network boundaries, and later routing decisions much easier to interpret.

An IPv4 octet contains eight bits. From left to right, the bit positions have decimal values of 128, 64, 32, 16, 8, 4, 2, and 1. A 1 means that the position's value is included; a 0 means that it is not. To convert a binary octet to decimal, add the values of the positions that contain 1.

For example, the decimal value 1 is 00000001 in binary because only the rightmost position is set. The decimal value 192 is 11000000 because $128 + 64 = 192$. The decimal value 168 is 10101000 because $128 + 32 + 8 = 168$. Using that method, the address 192.168.0.1 can be written as 11000000.10101000.00000000.00000001.

The following figure isolates one octet and shows how its active bit positions produce the decimal value 192.

Reading an IPv4 octet as eight binary positions

A position contributes its place value only when its bit is set to 1.




Figure 2.2: Binary place values showing how an IPv4 octet represents decimal 192

In this example, only the 128 and 64 positions are set to 1, so their values are included in the calculation while the remaining positions contribute zero. The figure makes the conversion method explicit: identify the active positions, add their values, and read the resulting decimal number.

Eight bits provide 256 possible combinations. The numeric range is 0 through 255 because zero is one of those values. This is why no IPv4 octet can be greater than 255. The same place-value method works in reverse: begin with a decimal value, mark the largest binary value that fits, subtract it, and continue until the remainder is zero.

Binary is therefore not a separate topic from addressing. It is the representation underneath every IPv4 address and subnet mask. Once you can see an octet as eight positions, the distinction between the network and host portions of an address becomes easier to understand. Binary also provides a direct bridge to hexadecimal, a more compact numbering system that appears in hardware addresses, packet captures, and IPv6 addresses.

Grouping an eight-bit octet into two four-bit units provides the next useful representation because each group maps directly to a single hexadecimal digit.

Hexadecimal in networking

Hexadecimal, commonly shortened to hex, is a base-16 numbering system. Decimal uses the ten symbols 0 through 9, whereas hexadecimal uses 0 through 9 followed by the letters A through F. In hexadecimal, A represents decimal 10, B represents 11, and the sequence continues through F, which represents decimal 15.

The relationship between binary and hexadecimal is especially useful in networking because one hexadecimal digit represents exactly four binary bits. The following table shows the direct relationship among decimal, four-bit binary, and hexadecimal values.

Decimal	Binary	Hexadecimal
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E

Decimal	Binary	Hexadecimal
15	1111	F

Table 2.1: Decimal, binary, and hexadecimal equivalents

Because an octet contains eight bits, it can be represented by two hexadecimal digits. To perform the conversion, divide the octet into two groups of four bits and convert each group separately. For example, the binary value 11001000 becomes 1100 and 1000. The first group is C in hexadecimal, and the second is 8, producing the hexadecimal value C8. Both values represent decimal 200.

Hexadecimal commonly appears in MAC addresses. An address such as 00-AA-00-4F-2A-9C contains six pairs of hexadecimal digits, with each pair representing one byte. Wireshark and similar tools also use hexadecimal when displaying frame contents and individual field values. IPv6 addresses, covered later in this book, use hexadecimal to represent their much larger address space compactly.

You do not need to convert every hexadecimal value manually during routine troubleshooting. You should, however, be able to recognize the notation and understand its relationship to binary. This makes hardware addresses and packet captures less intimidating and prepares you to interpret IPv6 addresses later. With the major numbering systems established, we can now examine the Windows tools that expose addresses, routes, and connectivity during practical troubleshooting.

Essential network connectivity and routing tools

Windows includes command-line tools that reveal different parts of the communication path. One command shows the host's configuration, another tests whether ICMP replies return, another displays recent IP-to-MAC resolutions, and others expose routed hops or the local routing table. The useful troubleshooting sequence is to choose the command that answers the current question rather than run every tool and treat any response as proof that the entire network works. The following figure organizes the tools into a troubleshooting path that begins with the host and expands outward toward the destination.

A practical path for troubleshooting Windows connectivity

Start close to the host and expand outward only after each dependency has been verified.




Figure 2.3: A practical sequence for narrowing Windows connectivity problems

The sequence starts with the adapter and local TCP/IP stack because remote tests are difficult to interpret when the host's own configuration is uncertain. Once those checks succeed, the investigation can expand to local delivery, remote reachability, and the routed path.

Inspecting configuration with ipconfig

The `ipconfig` command provides the basic IPv4 and IPv6 configuration for each Windows adapter. As shown earlier, the simplest form displays addresses, subnet masks, and default gateways. When you need the fuller configuration, use the following command:

```
ipconfig /all
```

The expanded output includes details such as the physical address, whether DHCP is enabled, the DHCP server, configured DNS servers, and NetBIOS over TCP/IP status. Start by identifying the adapter that is actually in use. A disconnected Ethernet adapter, a virtual adapter, and an active Wi-Fi adapter can all appear in the same output, and reading the wrong section can send the investigation in the wrong direction.

Testing reachability with ping

ping sends ICMP Echo Request messages and reports any Echo Reply messages that return. A reply provides evidence of IP-level reachability for that test and shows the measured round-trip time. A timeout is less conclusive. The destination may be unavailable, a route may be missing, or a firewall may simply be filtering ICMP while allowing the application traffic you actually care about.

The `-a` option attempts to resolve the target IP address to a host name. The following example tests `8.8.8.8` and requests reverse name resolution:

```
ping -a 8.8.8.8
```

If the reverse lookup succeeds, the result includes the corresponding name. That result tests more than basic reachability because name resolution is now part of the operation.

Windows sends four Echo Requests by default. The `-t` option continues until you stop the command, which is useful when watching whether a connection drops during a change or over time:

```
ping -t 8.8.8.8
```

Use `Ctrl + C` to stop the continuous test and display the final statistics. A continuous ping creates ongoing traffic, so use it deliberately rather than leaving it running without a reason.

The `-l` option changes the size of the data field in the Echo Request. Windows uses 32 bytes by default. The following example sends a 500-byte data field:

```
ping 8.8.8.8 -l 500
```

You can also request a 1,500-byte data field:

```
ping 8.8.8.8 -l 1500
```

Packet-size tests can help investigate a suspected **Maximum Transmission Unit (MTU)** problem, but the result needs careful interpretation. The `-l` value controls only the ICMP data field; IP and ICMP headers add bytes to the packet. In addition, an IPv4 packet may be fragmented unless the **Do Not Fragment** option is set. A failed large ping therefore does not by itself identify the MTU or the device causing the problem.

Examining address resolution with arp

The **Address Resolution Protocol (ARP)** maps IPv4 addresses to MAC addresses on the local network. Windows stores recently resolved entries in an ARP cache. To display the current entries for all relevant interfaces, run the following command:

```
arp -a
```

If the output is longer than the Command Prompt window, pipe it through `more` to display one screen at a time:

```
arp -a | more
```

The output associates IPv4 addresses with physical addresses and identifies whether entries are dynamic or static. For a destination on the same subnet, an ARP entry can show that the host resolved that destination's MAC address. For a remote destination, the host normally resolves the MAC address of the next-hop gateway instead of the remote host. An ARP entry is therefore evidence about local link-layer resolution. It does not prove that the destination service is reachable or that every router beyond the gateway has a valid route.

This distinction prevents a common troubleshooting mistake. If ping receives no reply but ARP has resolved the local destination or gateway, the local network path may be functioning while ICMP is blocked or the failure lies farther away. That narrows the investigation, but it is not enough to declare the end-to-end route healthy.

Analyzing a path with pathping and tracert

When a destination is several routed hops away, pathping and tracert expose more of the path than ping. pathping first discovers the intermediate routers and then sends repeated probes over time to calculate latency and packet-loss statistics for the hops and links it observes. Run it with the destination address or name, as shown here:

```
pathping 8.8.8.8
```

Because pathping gathers repeated measurements, it takes longer to complete than tracert. That additional time provides statistical information that a single route trace does not.

The Windows tracert command discovers the approximate layer 3 path by sending probes with increasing Time to Live values. Each router that decrements the value to zero can return an ICMP Time Exceeded response, allowing tracert to list that hop. The following figure shows a trace toward Google's public DNS address at 8.8.8.8.

```

C:\>tracert 8.8.8.8

Tracing route to dns.google [8.8.8.8]
over a maximum of 30 hops:

  1  <1 ms  <1 ms  <1 ms  router.charter.net [192.168.1.1]
  2   9 ms   9 ms   7 ms  22.54.156.1
  3   9 ms   9 ms   9 ms  acr01csbyor-gbe-1-44.csby.or.charter.com [96.34.105.202]
  4  12 ms  11 ms  11 ms  dtr01lnbhwa-tge-0-0-0-0.lnbh.wa.charter.com [96.34.111.208]
  5  18 ms  20 ms  17 ms  096-034-108-175.biz.spectrum.com [96.34.108.175]
  6  25 ms  31 ms  23 ms  096-034-108-160.biz.spectrum.com [96.34.108.160]
  7  47 ms  31 ms  31 ms  bbr02sttlwa-bue-4.sttl.wa.charter.com [96.34.0.56]
  8  25 ms  26 ms  26 ms  prr01sttlwa-bue-2.sttl.wa.charter.com [96.34.3.39]
  9  26 ms  24 ms  25 ms  72.14.194.48
 10  25 ms  29 ms  28 ms  108.170.227.7
 11  26 ms  27 ms  25 ms  142.251.50.177

```

Figure 2.4: tracert lists responding layer 3 hops toward the destination

Each numbered row in the figure represents a hop that returned the expected response. The three time values are round-trip measurements for separate probes. When reverse name resolution succeeds, tracert also displays a host name; otherwise, it shows only the IP address. The first hop is commonly the local default gateway, later hops belong to provider or transit networks, and the final responding row is the destination when the trace completes.

Run the same type of trace with the following command:

```
tracert 8.8.8.8
```

Name lookups can slow the display. If you need the addresses without waiting for `tracert` to resolve names, use the `-d` option:

```
tracert -d 8.8.8.8
```

A missing response appears as an asterisk (*), but that does not automatically identify a failed router. Some devices forward ordinary traffic while filtering or limiting the diagnostic responses used by `tracert`. Routes can also change, and the return path may differ from the forward path. Treat a trace as a snapshot and compare it with repeated tests and the behavior of the actual application.

A trace can reveal responding hops, but the host must select a route before the first probe leaves, so the local routing table shows how every routed path begins.

Reading and changing the local routing table

A Windows host consults its local routing table before sending an IP packet. To display that table, run the following command:

```
route print
```

The command exposes every route available to the host, but effective troubleshooting depends on identifying which of those entries Windows will actually consider for the destination being tested.

Reading the route print output

Windows may display an interface list, an IPv4 routing table, and an IPv6 routing table. When troubleshooting IPv4 communication, begin with the Active Routes section of the IPv4 table. Each row describes a group of destinations and tells Windows where packets matching that destination should be sent. A shortened view of some common entries may resemble the following:

Network	Destination	Netmask	Gateway
	0.0.0.0	0.0.0.0	192.168.1.1
	127.0.0.0	255.0.0.0	On-link
	192.168.1.0	255.255.255.0	On-link
	224.0.0.0	240.0.0.0	On-link

This shortened output illustrates the routes Windows can use, including directly connected networks, the default route, and any static or persistent entries. It does not test whether a

destination is reachable; instead, it shows the forwarding information the host consults when deciding where to send a packet.

The first row represents the default route. The combination of destination `0.0.0.0` and netmask `0.0.0.0` matches any destination for which Windows does not have a more specific route. In this example, those packets are sent to the default gateway at `192.168.1.1`. If a computer can communicate with devices on its local network but cannot reach remote networks, this is one of the first entries to verify.

An entry marked `On-link` identifies a destination that Windows can reach directly through one of its interfaces without sending the packet to another router. The `192.168.1.0` entry represents the locally connected network in this example. Routes beginning with `127` identify loopback communication within the computer, while the `224.0.0.0` entry covers multicast destinations.

The complete output also includes `Interface` and `Metric` columns. The `Interface` column identifies the local address through which the packet will leave. The `Metric` column represents the route's relative cost; when otherwise equivalent routes exist, Windows generally prefers the one with the lower metric. Windows first looks for the most specific matching destination and then uses the metric to choose between routes with the same degree of specificity.

Reading the table in this order—destination, mask, gateway, interface, and metric—turns it from a list of numbers into a forwarding decision. Once the current decision is understood, an administrator can determine whether an additional route is actually required.

Adding a persistent static route

A static route can direct traffic for a destination network through a chosen next-hop router. To add a persistent route for the `192.168.2.0` network, open Command Prompt as an administrator and use the following structure, replacing the bracketed value with a directly reachable gateway in your environment:

```
route /p add 192.168.2.0 mask 255.255.255.0 [gateway IP address]
```

The `/p` option makes the route persistent across TCP/IP restarts; `add` creates the entry; `192.168.2.0` is the destination network, `255.255.255.0` is its subnet mask; and the final address is the next-hop gateway. An incorrect persistent route can direct traffic away from a working path, so verify the destination, mask, and gateway before adding it.

After the command succeeds, run `route print` again and confirm that the new destination appears under the active or persistent routes with the intended gateway. This final check connects the earlier concepts: the destination and mask are binary layer 3 values, the gateway

is the next routed hop, and the local interface must still resolve a layer 2 address and transmit over a working physical link.

These utilities answer different questions, but together they form a practical progression. `ipconfig` establishes what the host believes its configuration to be. `ping` checks whether ICMP replies return. `arp` shows local IPv4-to-MAC resolution. `tracert` and `pathping` reveal information about routed hops, while `route print` explains the host's forwarding choice. Using the tools in that context turns their output into evidence instead of a collection of disconnected command results.

Summary

In this chapter, we used the OSI model as a troubleshooting map for network communication. The physical layer carries signals, the data link layer handles frames and local MAC addressing, the network layer routes IP packets, and the transport layer uses TCP or UDP and port numbers to reach application services. The session, presentation, and application layers describe the conversation, representation, and network services used by applications. A protocol analyzer such as Wireshark can expose these nested structures in captured traffic.

We also connected IPv4 dotted-decimal notation to its 32-bit binary representation. Each octet contains eight bit positions with values from 128 through 1, producing 256 combinations in the range 0 through 255. Finally, we used `ipconfig`, `ping`, `arp`, `pathping`, `tracert`, and `route` to inspect the host configuration, test selected forms of reachability, examine local address resolution, observe routed paths, and understand local route selection.

You can now separate a broad report that the network is not working into more precise questions about the host, the local link, the routed path, the transport exchange, and the application service.

In the next chapter, we will build on that model by examining IPv4 addresses in greater depth, distinguishing network and host portions, and using subnetting and CIDR to design and troubleshoot IPv4 networks.

3

How IPv4 Addresses Organize Network Traffic

An IPv4 address becomes useful only when it is read together with its subnet mask. The address `192.168.1.10` identifies an interface, but it does not reveal which destinations that interface considers local. With a `/24` mask, the address belongs to the `192.168.1.0` network. With a `/28` mask, it belongs to the smaller `192.168.1.0` through `192.168.1.15` block. That boundary determines whether a packet can be delivered on the local network, must be sent to a router, or fails because the host has been configured incorrectly.

Chapter 2 introduced the binary place values beneath dotted-decimal IPv4 addresses. This chapter puts that knowledge to work. We use **subnet masks** to separate network bits from host bits, examine why fixed address classes gave way to **Classless Inter-Domain Routing**, or **CIDR**, and calculate the address range represented by a prefix. We then follow traffic beyond the local subnet and examine how **Network Address Translation** and **Port Address Translation** connect privately addressed systems to public networks. These skills remain useful wherever IPv4 supports enterprise systems, home networks, and growing populations of Internet of Things devices.

After going through this chapter, you will be able to read a prefix length as a subnet mask, determine whether an IPv4 destination is local or remote, distinguish the historical classful model from CIDR, and choose a subnet size for a stated host requirement. You will also be able to calculate network, broadcast, and usable-host ranges and distinguish Hide NAT, static NAT, PAT, and port forwarding by following the address and port changes made at a network boundary.

We will be covering the topics through the following main headings:

- Using binary to locate subnet boundaries
- From fixed address classes to flexible subnets
- Crossing network boundaries with NAT and PAT

Note

Download the PDF version of this book

Your purchase includes a DRM-free PDF copy of this book, the code bundle, and a range of exclusive benefits. To unlock everything, follow the *Free benefits with your book* section in the *Preface*.

Technical requirements

This chapter is conceptual and requires no software setup. Keep a sheet of paper or a simple diagramming application nearby so that you can mark network and host bits, write subnet boundaries, and check address ranges. A basic calculator can confirm powers of two, although working through the examples manually will help you recognize the patterns more quickly during configuration and troubleshooting.

Using binary to locate subnet boundaries

Chapter 2 established that each IPv4 octet contains eight binary positions with place values of 128, 64, 32, 16, 8, 4, 2, and 1. Here, those values become a practical tool. A subnet mask places a continuous run of 1s on the network side of an address and 0s on the host side. Reading where that change occurs lets you interpret a CIDR prefix, identify the local network, and decide when a router is required.

Reading prefix lengths as subnet masks

An IPv4 address and its subnet mask are both 32 bits long, but they describe different things. The address identifies an interface within an IPv4 network. The mask identifies which address bits belong to the network prefix and which remain available for host addresses. In a valid conventional mask, the 1s are contiguous from the left; after the mask changes to 0, it does not return to 1. The familiar /24 mask is therefore written as follows:

```
11111111.11111111.11111111.00000000 = 255.255.255.0 = /24
```

The prefix length counts the network-side 1s. A /24 contains 24 of them, leaving eight host bits. Prefixes that end on an octet boundary are easy to recognize:

- /8 is 255.0.0.0
- /16 is 255.255.0.0
- /24 is 255.255.255.0.

The same rule also works when the boundary falls inside an octet, which is where binary place values become useful.

For a /20, the first 16 network bits fill two octets. The next four 1s occupy the 128, 64, 32, and 16 positions of the third octet, which add to 240. The mask is therefore 255.255.240.0. A /26 places two 1s in the final octet, giving $128 + 64 = 192$. A /28 places four 1s there, giving $128 + 64 + 32 + 16 = 240$. The resulting masks can be compared directly:

```
/20: 11111111.11111111.11110000.00000000 = 255.255.240.0
/24: 11111111.11111111.11111111.00000000 = 255.255.255.0
/26: 11111111.11111111.11111111.11000000 = 255.255.255.192
/28: 11111111.11111111.11111111.11110000 = 255.255.255.240
```

The address remains 32 bits long at every prefix. Only the boundary changes. A longer prefix moves that boundary to the right and leaves fewer host bits; a shorter prefix moves it to the left and leaves more. A host uses this boundary whenever it decides whether to communicate directly on the local link or send a packet to its default gateway.

Testing whether a destination is local

A host makes its local-or-remote decision by applying the subnet mask to both its own address and the destination address. At each bit position, a 1 in the mask keeps the corresponding address bit, and a 0 removes it from the network calculation. This bitwise AND operation produces the network address. Consider the host 192.168.1.77/26:

```
Address: 11000000.10101000.00000001.01001101
Mask:    11111111.11111111.11111111.11000000
Network: 11000000.10101000.00000001.01000000 = 192.168.1.64
```

The /26 mask leaves six host bits, so this subnet spans 64 addresses. It begins at 192.168.1.64 and ends at 192.168.1.127. A destination such as 192.168.1.118 produces the same network address, 192.168.1.64, so the host treats it as local and can use the Address Resolution Protocol to discover the destination's MAC address. A destination such as 192.168.1.130

produces the network address 192.168.1.128. Because that result differs, the host sends the packet toward its configured gateway.

```
192.168.1.77/26 -> network 192.168.1.64
192.168.1.118/26 -> network 192.168.1.64 -> local
192.168.1.130/26 -> network 192.168.1.128 -> remote
```

This comparison explains why two devices can have plausible-looking addresses and still fail to communicate. If the first host is mistakenly configured with /24, it will classify 192.168.1.130 as local and attempt to resolve it directly instead of using the gateway. Checking the address and gateway without checking the mask can therefore leave the actual fault hidden.

Modern hosts rely on the configured prefix for this decision. Earlier IPv4 networks used fixed boundaries tied to address classes. Those classes no longer determine a host's subnet, but understanding them explains the notation and assumptions that CIDR was designed to replace.

From fixed address classes to flexible subnets

Early TCP/IP networks divided much of the IPv4 address space into fixed classes. TCP/IP grew from research sponsored by the U.S. Department of Defense's DARPA and became the communication foundation of the modern Internet. Fixed class boundaries initially simplified allocation and routing, but they offered only a few network sizes. As organizations connected more systems, those large jumps between available block sizes wasted addresses and made growth harder to plan.

The move to CIDR replaced those fixed assumptions with an explicit prefix length. The classful ranges still appear in certification material, older documentation, and everyday networking language, so they remain useful historical context. The configured prefix, however, defines the actual network. This section follows that shift and then uses prefix length to choose a capacity and calculate the addresses inside a subnet.

Reading the classful address ranges

Class A addresses historically used a first octet from 1 through 126 and a default /8 mask, written as 255.0.0.0. That left three octets, or 24 bits, for host addressing. Blocks of this size were allocated to governments, universities, and other large scientific or research organizations. The 0/8 block has special meaning, and the entire 127/8 range is reserved for loopback. A **loopback** address tests the local IP stack; it does not prove that a physical adapter, cable, wireless connection, or upstream network is working.

Class B addresses used first-octet values from 128 through 191 and a default /16 mask of 255.255.0.0. Class C addresses used values from 192 through 223 and a default /24 mask of 255.255.255.0. These progressively smaller blocks were intended for organizations with different capacity requirements. *Figure 3.1* places the three default masks together for comparison.




Figure 3.1: Default subnet masks for Class A, Class B, and Class C networks

The masks in the figure differ by one complete octet. Class A fixes the first eight bits as the network portion, Class B fixes the first 16, and Class C fixes the first 24. These are historical defaults. In a modern configuration, an address beginning with 10 does not automatically use /8, and an address beginning with 192 does not automatically use /24. The prefix supplied with the address defines the real boundary.

Two additional classes describe different uses of the address space. Class D covers 224.0.0.0 through 239.255.255.255 and supports multicast rather than ordinary host addressing. **Multicast** lets a sender deliver traffic to a group of interested receivers. For example, OSPF uses 224.0.0.5 to communicate with all OSPF routers and 224.0.0.6 for designated and backup designated routers. Class D values begin with the binary pattern 1110.

Class E covers 240.0.0.0 through 255.255.255.255 and was historically reserved for experimental or future use. These addresses are not treated as ordinary public host space, and Class E values traditionally begin with 1111. The limited broadcast address 255.255.255.255 sits at the top of this range and has a specific purpose of its own.

Knowing the classes helps when you read older material or prepare for certifications such as CompTIA A+, Network+, and some Microsoft networking exams, but a class label does not replace the current mask. It also does not segment a modern LAN by itself. Switches and routers commonly use virtual LANs, or VLANs, to create separate broadcast domains, and each VLAN normally receives its own IP subnet. That separation depends on the configured prefix.

Why CIDR replaced fixed classes

The classful system allocated addresses in coarse blocks. An organization that needed more addresses than a Class C network could supply might receive a Class B network with more than 65,000 addresses, even when it required only a fraction of them. As Internet use expanded, this mismatch wasted address space and filled routing tables with individual classful networks.

Classless Inter-Domain Routing states the prefix length explicitly instead of deriving it from the first octet. A /20 reserves 20 network bits, a /24 reserves 24, and a /28 reserves 28.

Administrators can therefore size a subnet more closely to its host requirement. CIDR also supports route aggregation, in which a router advertises one larger prefix that represents several contiguous networks rather than advertising every smaller route separately.

CIDR turns the boundary into a design choice. Increasing the prefix length allocates more bits to the network and creates smaller subnets. Decreasing it leaves more host bits and creates larger subnets. The next step is to translate that movement into the number of addresses a subnet can support.

Choosing a prefix for the required host count

Every bit transferred from the host portion to the network portion doubles the number of subnets available within the original block and halves the number of addresses in each subnet. Moving the boundary to the right therefore produces more, smaller networks. Moving it to the left produces fewer, larger networks. The correct choice depends on how many addresses a segment must support and how much room it needs for growth.

Figure 3.2 shows how moving the boundary changes the balance between network bits and host bits within 192.168.1.0.

Sliding Scale

Move Mask to Right for More Networks

Move Mask to Left for More Hosts

1	1	1	1	1	1	1	1	Binary Values
128	64	32	16	8	4	2	1	

```
192.168.1.0
11000000.10101000.00000001.00000000
Subnet Mask /24
11111111.11111111.11111111.00000000
```



Figure 3.2: Moving the subnet boundary changes network and host capacity

In the figure, the /24 mask places 24 consecutive 1s on the network side and leaves eight 0s on the host side. Moving the boundary four positions to the right produces /28. The first 28 bits then identify the subnet, while four bits remain available for addresses inside it. The address is still 32 bits long; the mask changes how those bits are interpreted.

If h represents the number of host bits, the subnet contains 2^h total addresses. In an ordinary IPv4 LAN subnet, the all-zero host value identifies the network and the all-one host value identifies the broadcast address. The conventional usable-host calculation is therefore $2^h - 2$. Point-to-point /31 networks and /32 host routes are special cases; the conventional calculation applies to the LAN examples used here.

A /24 leaves eight host bits and contains 256 total addresses, normally providing 254 usable host addresses. A /26 leaves six host bits, giving 64 total addresses and 62 usable hosts. A /28 leaves four host bits, giving 16 total addresses and 14 usable hosts. Comparing these values lets an administrator rule out prefixes that cannot support the required number of devices.

```
/24: 8 host bits -> 256 total addresses -> 254 usable hosts
/26: 6 host bits -> 64 total addresses -> 62 usable hosts
/27: 5 host bits -> 32 total addresses -> 30 usable hosts
/28: 4 host bits -> 16 total addresses -> 14 usable hosts
```

Suppose a department expects 50 addressed devices. A /27 provides only 30 usable addresses, so it is too small; a /26 provides 62 and meets the requirement. A separate sensor network with 12 devices can fit within a /28, which provides 14 usable addresses. Capacity planning should also count the gateway interface, printers, access points, management interfaces, and

reasonable growth. A prefix that barely fits the current device count can force an avoidable renumbering exercise later.

Subnet size also shapes the broadcast domain and the address plan around it. Each VLAN needs a distinct network address, a usable range, a broadcast address, and often a documented gateway address. Recording all four values when a subnet is assigned makes later routing and troubleshooting checks much faster.

Calculating the network broadcast and host range

Knowing a subnet's capacity does not reveal where its address block begins. When the changing portion of the mask is in the final octet, subtract that mask value from 256 to find the block size. For 255.255.255.240, the calculation $256 - 240$ gives 16. The /28 boundaries therefore occur at 0, 16, 32, 48, 64, 80, and continue in increments of 16.

Consider the host address 192.168.1.77/28. The value 77 falls between the boundaries 64 and 80, so the subnet begins at 192.168.1.64. The next subnet begins at 192.168.1.80, making the address immediately before it, 192.168.1.79, the broadcast address. The usable range is 192.168.1.65 through 192.168.1.78, and the original host address falls inside that range.

```
Host address: 192.168.1.77/28
Network:      192.168.1.64/28
First host:   192.168.1.65
Last host:    192.168.1.78
Broadcast:    192.168.1.79
Next subnet:  192.168.1.80/28
```

The changing octet is not always the fourth one. For 10.20.34.117/20, the mask is 255.255.240.0, so the block size in the third octet is $256 - 240 = 16$. The third-octet boundaries are 0, 16, 32, 48, and so on. Because 34 falls in the 32-through-47 block, the network is 10.20.32.0 and the broadcast address is 10.20.47.255.

```
Address:      10.20.34.117/20
Network:      10.20.32.0/20
First host:   10.20.32.1
Last host:    10.20.47.254
Broadcast:    10.20.47.255
```

This example uses the same method as the /28 calculation. The mask identifies the octet in which the boundary changes, the block size identifies the surrounding boundaries, and the addresses immediately inside those boundaries define the usable range. Working from an

assigned host address rather than from an already known network address is especially useful when checking a live configuration.

A broadcast address delivers an IPv4 broadcast to hosts on the local subnet and cannot be assigned to an ordinary host. This function is separate from the way an Ethernet switch learns devices. A switch learns source MAC addresses from received frames, while broadcasts and protocols such as ARP help hosts discover information within a broadcast domain. Keeping these functions separate prevents a correct subnet calculation from leading to an incorrect troubleshooting assumption.

Subnet masks frequently explain failures that initially look like routing problems. A host at 192.168.1.77/28 and another at 192.168.1.82/28 belong to adjacent subnets. Each should use a router to reach the other. If the first host is mistakenly configured with /24, it will treat 192.168.1.82 as local and attempt to resolve it directly. A useful check therefore compares the configured prefix with the intended boundary, then verifies that the host sends remote traffic to the correct gateway.

Checking the mask and gateway together

A subnet calculation becomes a troubleshooting model when the address, mask, and default gateway are checked together. Begin by calculating the network that contains the host. Next, calculate the network that contains the destination. Matching results mean that the host should attempt local delivery; different results mean that it should send the packet to the gateway. This sequence tells you which part of the path to investigate instead of treating every failed connection as a routing problem.

The gateway address must itself be reachable on the host's local subnet. For 192.168.1.77/28, the usable range is 192.168.1.65 through 192.168.1.78. A gateway at 192.168.1.65 is inside that range, whereas 192.168.1.94 belongs to the next /28 subnet. Configuring the latter would leave the host with an address and mask that look valid individually but no usable next hop for remote traffic.

Mismatched masks can also create an asymmetric view of the network. If 192.168.1.77 uses /28 and 192.168.1.82 incorrectly uses /24, the first host sees the second as remote, while the second sees the first as local. One endpoint sends toward a router and the other attempts direct delivery. Checking only one device can miss this disagreement, so both endpoint configurations should be compared whenever a failure affects one direction or only certain address pairs.

The tools introduced in *Chapter 2* supply the evidence for this check. An interface configuration command reveals the address, mask, and gateway; the subnet calculation shows how those values should behave; and ARP, ping, or a route display helps locate the point at which actual

behavior differs. Binary is useful here because it turns the command output into a testable expectation about the packet's next step.

The prefix now explains which destinations remain inside the local network and which must cross a router. At the edge of many IPv4 networks, crossing that boundary creates another problem: privately addressed devices need a controlled way to communicate through the limited public IPv4 address space.

Crossing network boundaries with NAT and PAT

Network Address Translation (NAT) changes address information as traffic crosses a translating router or firewall. It is commonly used between a privately addressed internal network and the public Internet. Internal systems can keep addresses that are not routed globally, while one or more public addresses represent them externally. Translation conserves public IPv4 space and lets an organization manage its internal addressing separately from the addresses assigned by an Internet service provider.

NAT and firewalling often occur on the same device, but they perform different jobs. NAT changes an address or port and records enough state to reverse the change for return traffic. A firewall policy decides whether that traffic is permitted. Hiding an internal address reduces what an external system can see, but translation alone is not the security control. The following examples separate outbound sharing, fixed one-to-one mappings, and port-based forwarding so that each packet can be followed across the correct boundary.

Sharing a public address with Hide NAT

Many firewall products use the term Hide NAT for outbound translation in which multiple private clients share one public IPv4 address. It is commonly described as many-to-one translation from the internal side. In practice, the device normally uses Port Address Translation as part of the process because source ports distinguish simultaneous connections that all use the same public address.

Suppose the client 10.1.1.1 sends a web request to its firewall at 10.1.1.2. Before placing the packet on the Internet, the firewall replaces the private source address with its public address, 199.42.6.23, and records the original address and port. The remote server replies to 199.42.6.23. The firewall consults its translation state, restores 10.1.1.1 as the internal destination, and forwards the response to the client.

Figure 3.3 extends the example across two private networks. It shows source translation at the client edge and destination translation at the server edge, with public addresses carrying the packet across the Internet.




Figure 3.3: Address translation across two private networks

At the source edge, 10.1.1.1 is translated to 199.42.6.23. Across the Internet, the client targets the server network's public address, 8.5.3.2. At the destination edge, the firewall translates 8.5.3.2 to the server's private address, 10.2.2.2. Return traffic reverses both mappings. Following the packet one boundary at a time prevents the two translations from being mistaken for a single operation.

Publishing a server with static NAT

Static NAT creates a fixed one-to-one relationship between a public address and a private address. It is useful when an internal server must remain reachable through a consistent public address. In this example, public address 8.5.3.2 maps to the web and email server at 10.2.2.2. Unlike a temporary outbound translation, this mapping remains defined even when no client session is active.

The mapping tells the translating device which private address corresponds to the public one. It does not require the firewall to permit every protocol and port. The same device can map the address while allowing only the services the server provides, such as TCP 80 and 443 for web traffic or TCP 25 and 587 for email. *Figure 3.4* shows the fixed public-to-private relationship.




Figure 3.4: A static NAT mapping publishes one internal server

An external client sends traffic to 8.5.3.2, and the destination firewall translates that address to 10.2.2.2 before forwarding the packet internally. The server replies through the same firewall, which changes the private source back to 8.5.3.2. The public-private association works in both directions, while the firewall rule still determines which traffic may cross the boundary.

Static NAT is most straightforward when one public address can be dedicated to one internal system. When several internal servers must share a single public address, the address alone can no longer identify the destination. The translating device must also examine the transport-layer port.

Directing services with PAT and port forwarding

Port Address Translation (PAT) uses TCP or UDP port numbers in addition to IP addresses. For outbound Hide NAT, ports allow many clients to share one public address without mixing their sessions. For inbound traffic, a related configuration called port forwarding or destination NAT uses the public destination port to select an internal server. The two uses rely on the same packet fields but solve different traffic direction problems.

Assume that a web server and an email server must share public address 8.5.3.2. The firewall can forward TCP port 443 to the web server at 10.2.2.10 and TCP port 587 to the email server at 10.2.2.20. The public address alone is therefore insufficient to choose the destination; the firewall reads the destination port and applies the corresponding mapping. *Figure 3.5* shows the two paths.




Figure 3.5: Port forwarding directs services to different internal servers

The combination of public address and destination port selects the internal endpoint. One public address can publish several services as long as every public port has an unambiguous mapping. The client in the figure uses 10.1.1.1, which is a private example address. On a real Internet path, that client would use a public source address or be translated at its own network edge; the figure focuses on the destination-side forwarding.

The same public address and public port cannot normally be forwarded to two internal destinations at the same time. If two servers need the same externally visible port, the design

requires another public address, a different public port, or an application-layer intermediary such as a reverse proxy. Whatever mapping is chosen, the firewall should permit only the protocols and ports required by the published service.

Hide NAT, static NAT, and port forwarding solve different addressing problems. Hide NAT lets internal clients share a public address for outbound sessions. Static NAT provides a stable one-to-one address mapping. Port forwarding combines a public address with a port to select an internal service. During troubleshooting, begin with the direction of the first packet, record the address and port before translation, and compare them with the values expected after the boundary.

Summary

In this chapter, you used the binary foundation from *Chapter 2* to interpret IPv4 subnet masks. A CIDR prefix counts the continuous network-side 1s in the mask, and the remaining bits define the host portion. Applying that mask to an address reveals the local network and explains why a host sends some destinations directly while sending others to its default gateway.

You then traced the shift from fixed classful boundaries to CIDR and used host bits to select a subnet capacity. By working through /20, /24, /26, /27, and /28 examples, you calculated block sizes, located network and broadcast addresses, and identified usable-host ranges even when the supplied address was not itself a network address. These calculations also showed how an incorrect mask can make a routing failure look like a local connectivity problem.

Finally, you followed private IPv4 traffic across translated boundaries. Hide NAT supports shared outbound access, static NAT maintains a fixed public-to-private mapping, and port forwarding uses a destination port to publish a particular service. Separating translation from firewall policy provides a clearer model for both design and troubleshooting.

In the next chapter, you will extend this addressing model to IPv6, compare IPv4 subnet masks with IPv6 prefix lengths, and examine public, private, link-local, and automatically assigned addresses. The idea of a network boundary remains useful even though IPv6 uses hexadecimal notation and a much larger address space.

4

From Scarcity to Scale: How IPv6 Changes Addressing

An IP address has two jobs: it identifies an interface and places that interface within a network that routers can reach. IPv4 performs those jobs with 32 bits, but its finite address space was never designed for today's number of computers, phones, cloud workloads, sensors, and virtual systems. IPv6 expands the address to 128 bits and, in doing so, changes how addresses are written, assigned, and interpreted. The familiar reasoning from IPv4 still matters—the prefix identifies the network and the remaining bits identify an interface—but the notation and operating practices are different.

After going through this chapter, you will be able to explain why IPv6 is needed, read and shorten an IPv6 address without changing its value, use a prefix length to distinguish the network portion from the interface portion, and recognize the practical scope of global unicast, unique local, and link-local addresses. You will also compare automatic address assignment in IPv4 and IPv6 and interpret the IPv6 information that Windows displays through `ipconfig`. These skills turn a long hexadecimal string into useful evidence about where a device can communicate and how it obtained its configuration.

We will cover the topics through the following main headings:

- Why IPv6 expands the addressing model
- Reading IPv6 addresses and prefixes
- Recognizing address scope in IPv4 and IPv6
- Understanding automatic address assignment
- Reading IPv6 configuration on Windows

Technical requirements

This chapter assumes that you understand the IPv4 concepts introduced earlier: binary place value, dotted-decimal notation, subnet masks, CIDR prefix lengths, default gateways, DHCP, and the difference between local and remote destinations. You also need access to a Windows computer or virtual machine on which you can open Command Prompt and run `ipconfig /all`. No IPv6 server is required; an IPv6-enabled Windows interface normally creates a link-local address even when no globally reachable IPv6 service is available.

Why IPv6 expands the addressing model

The purpose of this section is to connect the address-space problem to the design choices that appear throughout the rest of the chapter. It begins with the limit of IPv4, then shows why IPv6 is more than a longer version of the same address and why both protocols often appear together on modern networks.

An IPv4 address contains 32 bits, so the complete space has 2^{32} possible values—about 4.29 billion. Not all of those values can be assigned as public host addresses because blocks are reserved for private networks, multicast, loopback, documentation, and other purposes. Address conservation techniques such as private addressing and **Network Address Translation (NAT)** allowed many devices to share a smaller number of public addresses, but they did not increase the underlying IPv4 space. They also introduced translation state that administrators must consider when publishing services or troubleshooting end-to-end traffic.

IPv6 uses 128 bits, producing 2^{128} possible values—approximately 3.4×10^{38} . The practical benefit is not simply an enormous count of individual hosts. Organizations can receive hierarchically structured address space, divide it into many subnets, and preserve end-to-end addressing without depending on address translation as the default design. Firewalls and security policy remain essential: a globally unique IPv6 address does not make a device safe, and IPv6 support for IPsec does not mean that all IPv6 traffic is automatically encrypted or authenticated.

IPv4 and IPv6 commonly operate side by side in a dual-stack network. An application and the operating system select an appropriate destination and source address according to DNS results, route availability, and local policy. As a result, troubleshooting must begin by identifying which protocol and which address were actually used. Before that evidence can be interpreted, however, the hexadecimal notation of IPv6 has to become as readable as dotted decimal.

Reading IPv6 addresses and prefixes

This section develops a repeatable way to read an IPv6 address: separate it into hextets, expand or compress it correctly, and apply the prefix length to find the routing boundary. The goal is not to perform lengthy binary conversions by hand, but to see enough structure in the address to make a sound local-versus-remote decision.

From binary bits to hexadecimal hextets

IPv6 is still processed as binary. Hexadecimal is the compact notation used by people and tools to represent those bits. One hexadecimal digit represents four bits and can have a value from 0 through 9 or A through F. Four hexadecimal digits therefore represent 16 bits. IPv6 writes eight of these 16-bit groups, called hextets, separated by colons: eight hextets $\times$ 16 bits gives the full 128-bit address.

Consider the address 2001:0db8:0000:0042:0000:0000:0000:0015. Its eight hextets are visible, and every hextet contains four hexadecimal digits. The left side usually carries the routing and subnet information; the right side identifies an interface on the subnet. The exact dividing point is not inferred from the colons—it is stated by the prefix length that follows the address.

The following comparison places the 32-bit IPv4 format beside the 128-bit IPv6 format and highlights the operational changes that accompany the larger address space.

IPv4 vs. IPv6 Comparison

Feature	IPv4	IPv6
Address size	32 bits	128 bits
Address format	Dotted decimal (e.g., 192.168.1.1)	Hexadecimal, colon-separated (e.g., 2001:0db8:85a3:0000:0000:8a2e:0370:7334)
Example address	192.168.1.1	2001:0db8:85a3:8a2e:370:7334 (compressed format)
Total number of addresses	Approximately 4.3 billion (2^{32})	Approximately 3.4×10^{38} (2^{128})
Address notation	Decimal (0–255 per octet)	Hexadecimal (0–FFFF per segment)
Address segments	4 octets (e.g., 192 . 168 . 1 . 1)	8 segments (e.g., 2001 : 0db8 : 85a3 : 0000 : 0000 : 8a2e : 0370 : 7334)
Header size	Variable (20–60 bytes)	Fixed at 40 bytes (simpler header)
Checksum	Included in header	Not included (improves performance)
Broadcast	Supported	Not supported (uses multicast and anycast)
Configuration	Manual or DHCP	Manual, DHCPv6, or Stateless Address Autoconfiguration (SLAAC)
NAT (Network Address Translation)	Commonly used (due to address shortage)	Generally not required (large address space)
Security	IPsec optional	IPsec support built-in (optional, but designed for use)
Fragmentation	By routers and hosts	By source host only (uses Path MTU Discovery)
Address resolution	ARP (Address Resolution Protocol)	NDP (Neighbor Discovery Protocol)
DNS record type	A record	AAAA record
Typical use today	Still widely used	Growing adoption (required for future growth)

Figure 4.1: Comparing the address and packet-handling characteristics of IPv4 and IPv6

Figure 4.1 shows that address length is only the most visible difference. IPv6 uses colon-separated hexadecimal notation, removes broadcast in favor of multicast, and relies on Neighbor Discovery rather than ARP. It also uses AAAA records for IPv6 name resolution and changes how packet headers and fragmentation are handled. These differences explain why an IPv6 address cannot be treated merely as an IPv4 address with more digits.

Shortening an IPv6 address safely

A complete IPv6 address can be unwieldy, so two compression rules make it easier to read:

- First, leading zeros may be removed from any hextet: 0db8 becomes db8, and 0042 becomes 42. A hextet that contains only zeros must retain at least one zero unless it is part of the double-colon rule.
- Second, one contiguous sequence of zero-valued hextets may be replaced by ::. The double colon can appear only once because a reader must be able to determine how many omitted hextets are needed to restore a total of eight.

Applying those rules to 2001:0db8:0000:0042:0000:0000:0000:0015 produces 2001:db8:0:42::15. Compression changes the text, not the 128-bit value. When more than

one zero run could be shortened, canonical notation normally compresses the longest run; a tie is resolved by using the first run.

The compression process is illustrated below by separating the removal of leading zeros from the replacement of consecutive zero-valued hexets.

Simplifying an IPv6 Address

You can remove leading zeros in each segment and use :: to replace one consecutive block of all zeros.




Figure 4.2: Compressing an IPv6 address without changing its value

Figure 4.2 reinforces the order of reasoning:

1. Verify that the full address has eight hextets
2. Remove only leading zeros within each hextet
3. Then use :: for a single consecutive run of all-zero hextets

The ability to reverse the process is valuable during troubleshooting because two differently written addresses may represent the same value. Once the notation is under control, the prefix length reveals which part of that value routers use as the network identity.

Using the prefix to make a routing decision

IPv6 writes its network boundary with CIDR notation, just as modern IPv4 does. In an address such as 2001:db8:8635:524e:8b19:ac31:41e0:6512/64, the /64 means that the first 64 bits

form the subnet prefix and the remaining 64 bits form the interface identifier. A /64 therefore contains 2^{64} possible interface-identifier values, not 2^{128} . The entire IPv6 address space is 128 bits; the prefix has already fixed half of those bits in this example.

A /64 is the standard size for most IPv6 LANs and is required for **Stateless Address Autoconfiguration (SLAAC)** on ordinary links, but it is not the only prefix length used anywhere in IPv6. Point-to-point links, delegated prefixes, loopbacks, and routing summaries may use other lengths. The prefix supplied by the network—not an assumption that every address is /64—must guide the calculation.

The local-versus-remote test is familiar from IPv4. A host compares the destination's prefix with the prefix of its own address. If the relevant bits match, the destination is on-link, and the host uses IPv6 Neighbor Discovery to learn how to reach it. If they do not match, the host sends the packet to a router selected from its IPv6 routing information. For example, 2001:db8:10:20::25/64 and 2001:db8:10:20::90/64 share the same first four hexets and are in the same /64 subnet; 2001:db8:10:21::90/64 is in a different subnet. This routing boundary matters more than the visual length of the compressed addresses.

Recognizing address scope in IPv4 and IPv6

An address can be correctly written yet still be unusable for a particular destination. This section connects address patterns to scope: private IPv4 addresses normally depend on NAT for Internet access, while IPv6 uses distinct global, unique-local, and link-local ranges. Recognizing those ranges allows an administrator to set realistic expectations before testing connectivity.

Private IPv4 ranges and NAT

Private IPv4 addresses are intended for internal use and are not routed across the public Internet. The three ranges reserved by RFC 1918 are 10.0.0.0/8 through 10.255.255.255, 172.16.0.0/12 through 172.31.255.255, and 192.168.0.0/16 through 192.168.255.255. A router or firewall commonly translates traffic from these addresses to a public IPv4 address when users access Internet services. NAT conserves public addresses and hides internal addressing details, but it should not be mistaken for a complete security policy.

The following figure brings the three private IPv4 blocks together so that their boundaries can be checked at a glance.

The Three Internal (Private) Subnets in IPv4

These address ranges are defined in RFC 1918 for use within private networks. They are **not routable** on the public Internet.

10.0.0.0/8	172.16.0.0/12	192.168.0.0/16
Address range	Address range	Address range
10.0.0.0 – 10.255.255.255	172.16.0.0 – 172.31.255.255	192.168.0.0 – 192.168.255.255
Total addresses	Total addresses	Total addresses
$2^{24} = 16,777,216$	$2^{20} = 1,048,576$	$2^{16} = 65,536$
Example address	Example address	Example address
10.1.5.25	172.16.10.5	192.168.1.100
A single large private network, commonly used in medium and large organizations.	A medium-sized private network range, often used in enterprise environments.	A smaller private network range, commonly used in home networks and small businesses.



- These ranges are defined in RFC 1918.
- They are intended for internal use only and are **not routable** on the public Internet.
- Many organizations use Network Address Translation (NAT) to allow internal devices to access the Internet using a public IPv4 address.

Figure 4.3: The three private IPv4 address ranges defined by RFC 1918

Figure 4.3 is especially useful for the middle block, which is sometimes misread as all addresses beginning with 172. Only 172.16.0.0 through 172.31.255.255 is private. The precision needed to recognize that /12 boundary carries directly into IPv6, where the first hexadecimal digits indicate whether an address is global, local to an organization, or restricted to one link.

Global, unique-local, and link-local IPv6 addresses

The three unicast scopes most often encountered on an endpoint serve different purposes. A **global unicast address (GUA)** is globally unique and can be routed beyond the local organization when routing and security policy permit; currently allocated global space commonly falls within 2000::/3. A **unique local address (ULA)** comes from fc00::/7 and is intended for communication within an organization or between cooperating sites, not for routing on the public Internet. A link-local address comes from fe80::/10 and works only on the directly attached link; routers do not forward it to another link.

An interface can hold several of these addresses at the same time. It may use a link-local address for Neighbor Discovery and a default-router relationship, a ULA for an internal service, and a GUA for Internet communication. The operating system selects a source address

appropriate for the destination and local policy, so the presence of a GUA alone does not prove that external connectivity works. A route, functioning gateway, DNS information, and permitted traffic path are still required.

The next figure compares the reach and typical use of these three address categories.




Figure 4.4: Comparing global unicast, unique-local, and link-local IPv6 scope

Figure 4.4 ties the visible prefix to the expected communication boundary. GUA, ULA, and link-local addresses are all unicast addresses, but their scopes are not interchangeable. This distinction narrows a fault quickly: a host that has only an `fe80::` address can communicate with neighbors on the link, but that address cannot carry a packet through a router to a remote subnet.

Reading the structure of a unique local address

ULAs are sometimes compared with private IPv4 addresses because both support internal communication, but the comparison has limits. A ULA does not require NAT, and its large randomly generated global ID is intended to reduce collisions if independently addressed networks are later joined. The ULA block is `fc00::/7`. In normal locally assigned use, the L bit is set to 1, so the address begins with `fd`. The next 40 bits form a pseudo-random global ID, the following 16 bits identify a subnet, and the final 64 bits identify an interface on that subnet.

The field layout below shows how the 128 bits of a locally assigned ULA are divided and why two addresses can share an organizational prefix while belonging to different subnets.




Figure 4.5: Field structure of a locally assigned unique local IPv6 address

In Figure 4.5, the fixed fd value and the 40-bit global ID identify the locally created ULA space, while the 16-bit subnet ID lets an organization create internal /64 networks. The 64-bit

interface ID then distinguishes interfaces within a subnet. This hierarchy preserves routing structure without making the address publicly reachable. Address scope explains where an address may be used; address assignment explains how it appears on an interface in the first place.

Understanding automatic address assignment

The aim of this section is to separate three ideas that are often confused: the automatic creation of a link-local address, the use of router advertisements for SLAAC, and the optional use of **DHCPv6**. Understanding which mechanism supplied which information makes the resulting Windows configuration much easier to diagnose.

An IPv6-enabled interface normally creates a link-local address before it receives any global or unique-local prefix. It then performs Duplicate Address Detection to confirm that the tentative address is not already in use on the link. The address is valid only on that interface's local link, which is why Windows appends a zone or scope identifier such as %4 when the same link-local address range could exist on several interfaces.

A router can advertise an IPv6 prefix and other network information in Router Advertisement messages. With SLAAC, the host combines the advertised prefix with an interface identifier and checks the resulting address for duplication. The identifier may be stable, randomized, or temporary; modern systems do not have to expose a hardware MAC address in it. A host may also create a temporary address for outbound connections to reduce long-term tracking. DHCPv6 can provide addresses or additional settings, but it is not required for a host to form an address through SLAAC.

The following sequence illustrates a network in which a host configures IPv6 from router information without an IPv6 DHCP server.

IPv6 Internal Network with No DHCP Server

When there is no DHCP server, IPv6 hosts use Stateless Address Autoconfiguration (SLAAC) to configure their own addresses.




Figure 4.6: Forming an IPv6 address with SLAAC when no DHCPv6 server is present

Figure 4.6 shows the essential SLAAC relationship: the router advertises a prefix, the host supplies an interface identifier, and the two parts form an address for that subnet. The default-router relationship is learned from the Router Advertisement; in an actual Windows display, the default gateway often appears as the router's link-local address rather than as an address from the advertised prefix. The figure therefore represents the configuration process conceptually, while the command output later in the chapter shows how Windows presents the result.

IPv4 has a superficially similar fallback called **Automatic Private IP Addressing (APIPA)**. When a Windows host configured for DHCP cannot obtain an IPv4 lease, it may choose an address from 169.254.0.0/16 and test that the value is not already in use. APIPA supports communication only on the local IPv4 link and does not provide normal routed Internet access. IPv6 link-local addressing also has link scope, but it is not merely a DHCP failure mode: it is a normal part of IPv6 operation and coexists with other IPv6 addresses on the same interface.

Reading IPv6 configuration on Windows

This section turns the address rules into a practical inspection method. Rather than treating `ipconfig` output as a list to memorize, read each line as evidence about scope, assignment, routing, name resolution, and the interface to which the value belongs.

Open Command Prompt and run the following command to display the complete TCP/IP configuration for every adapter:

```
ipconfig /all
```

The excerpt below is the original example. It shows a physical adapter with both IPv4 and IPv6 configuration, followed by a Hyper-V virtual adapter whose addresses and routes serve a different network context.

```
IPv6 Address. . . . . : 2600:6c55:27f:ee35:7562:9a:
4afe:ed7e(Preferred)
Temporary IPv6 Address. . . . . :
2600:6c55:27f:ee35:ed72:1f7d:ee62:c913(Preferred)
Link-local IPv6 Address . . . . . : fe80::9562:9a:4afe:ed7e%4(Preferred)
IPv4 Address. . . . . : 192.168.1.111(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : Monday, May 3, 2021 11:42:07 PM
Lease Expires . . . . . : Tuesday, May 11, 2021 12:00:52 AM
Default Gateway . . . . . : fe80::9691:7fff:fe7a:3120%4
                            192.168.1.1
DHCP Server . . . . . : 192.168.1.1
DHCPv6 IAID . . . . . : 86264038
DHCPv6 Client DUID. . . . . : 00-01-00-01-26-3E-C9-FB-70-54-D2-1A-54-BD
DNS Servers . . . . . : fe80::9691:7fff:fe7a:3120%4
                            8.8.8.8
                            2607:f428:ffff:ffff::1
                            2607:f428:ffff:ffff::2

NetBIOS over Tcpip. . . . . : Enabled

Ethernet adapter vEthernet (Default Switch):

Connection-specific DNS Suffix . :
Description . . . . . : Hyper-V Virtual Ethernet Adapter
Physical Address. . . . . : 96-15-22-90-DD-AE
DHCP Enabled. . . . . : No
```

```

Autoconfiguration Enabled . . . . : Yes
Link-local IPv6 Address . . . . : fe80::7806:deef:2b8:d2ef%18(Preferred)
IPv4 Address. . . . . : 172.17.55.193(Preferred)
Subnet Mask . . . . . : 255.255.255.240
Default Gateway . . . . . :
DHCPv6 IAID . . . . . : 311825698
DHCPv6 Client DUID. . . . . : 00-01-00-01-00-1A-54-BD
DNS Servers . . . . . : fec0:0:0:ffff::1%1
                        fec0:0:0:ffff::2%1
                        fec0:0:0:ffff::3%1
NetBIOS over Tcpip. . . . . : Enabled

```

Start with the physical adapter. The address beginning with 2600 is a GUA in currently allocated global unicast space, and the temporary address shares its /64 prefix while using a different interface identifier for privacy. Both entries are marked Preferred, meaning they are valid for new communication at the time of the capture. Whether either address can reach the Internet still depends on an IPv6 route and the upstream network.

The `fe80::` address is link-local. The `%4` is the Windows zone identifier that ties the address to interface index 4; it is not part of the 128-bit IPv6 address. The IPv6 default gateway is also link-local and uses the same zone identifier, which is normal because the host reaches its first-hop router directly on the local link. The simultaneous `192.168.1.111` address and `255.255.255.0` mask show that the adapter is dual stack and belongs to the IPv4 subnet `192.168.1.0/24` as well.

The DHCP Server line clearly identifies the IPv4 DHCP server. The DHCPv6 IAID identifies the interface association, while the DUID identifies the DHCPv6 client. Their presence in the output does not by itself prove that the displayed GUA was leased by DHCPv6; Windows may form IPv6 addresses through SLAAC and still maintain DHCPv6 identifiers. Lease information, packet capture, router advertisements, and server records provide stronger evidence about the assignment method.

The DNS Servers list includes both IPv4 and IPv6 addresses. DNS translates hostnames into address records—A records for IPv4 and AAAA records for IPv6—and supports reverse lookup through separate reverse-mapping domains. A client may receive several resolver addresses and choose among them according to operating-system behavior and reachability. If name resolution fails while direct address tests succeed, these lines become more relevant than the host's own IP address.

Finally, the vEthernet adapter represents a separate virtual link created by Hyper-V. Its `fe80::` address has a different zone identifier, `%18`, and its private IPv4 address belongs to

172.17.55.192/28. The blank default-gateway field indicates that this adapter is not presenting a route for ordinary off-link traffic. This contrast is a reminder to identify the correct adapter before drawing conclusions: a valid address on a virtual switch does not describe the path used by the physical network interface.

Reading the output in this order—adapter, address type, prefix, gateway, assignment evidence, and DNS—turns a long configuration listing into a troubleshooting model. It also brings the chapter's main ideas together: notation tells us the value, the prefix defines the subnet, scope limits where the address is useful, and configuration data reveals how the host expects to reach other networks.

Summary

IPv6 responds to the limitations of the 32-bit IPv4 space with 128-bit addressing and a hierarchical design that supports large numbers of subnets and interfaces. Its addresses are written as eight hexadecimal hextets, but leading zeros and one consecutive run of zero hextets can be compressed without changing the underlying value. A prefix length identifies the routing boundary; on a typical /64 LAN, the first 64 bits identify the subnet, and the remaining 64 bits form the interface identifier. Hosts use that boundary to decide whether to reach a neighbor directly or send traffic through a router.

Scope is just as important as syntax. Global unicast addresses may be routed broadly, unique local addresses support internal IPv6 communication, and link-local addresses remain on one link. An interface can use several addresses and obtain them through several mechanisms: link-local creation is normal, SLAAC uses router advertisements, and DHCPv6 may supply addressing or other settings. By reading `ipconfig /all` in terms of scope, prefix, gateway, assignment, and DNS, an administrator can distinguish a correctly formed address from a complete, working path.

The next chapter moves from interpreting network addresses to building the infrastructure that carries them. You will examine how switches, routers, wireless access points, cabling, and management practices work together to turn an addressing plan into a network that can be deployed, monitored, and maintained.

Part 2

From Network Infrastructure to a Troubleshooting Method

Knowing how a host interprets an address is only the beginning. Traffic still has to cross switches, VLANs, routers, virtual networks, and security controls before it reaches the destination. This part turns the addressing foundation from *Part 1* into the infrastructure decisions that actually move and protect packets.

You will first examine how VLANs segment networks, how Layer 3 devices route between them, how static and dynamic routing extend reachability, and how firewalls and virtual switches introduce additional forwarding and security boundaries. You will then make an important shift from configuration to diagnosis. Instead of beginning with a device or command, you will learn to establish the scope and impact of an incident, compare what works with what fails, and use each test to reduce the number of plausible causes.

Together, these chapters provide the bridge between understanding a network and troubleshooting one. The infrastructure tells you where traffic is expected to travel; the troubleshooting method tells you how to determine where that expectation stops matching reality.

This part of the book includes the following chapters:

- *Chapter 5, How Network Infrastructure Moves and Protects Traffic*
- *Chapter 6, A Systematic Approach to Network Troubleshooting*

5

How Network Infrastructure Moves and Protects Traffic

A useful network does more than connect devices. It separates groups that should not share the same broadcast domain, selects a path when a destination lies elsewhere, and applies security policy before traffic crosses a trust boundary. Those tasks may be carried out by a physical router, a multilayer switch, a firewall, an operating-system routing table, or a virtual switch inside a host. Although the interfaces differ, the underlying questions remain consistent: where should this packet go, which device will forward it, and should the connection be allowed?

After going through this chapter, you will be able to explain how VLANs and Layer 3 switching support segmentation and inter-VLAN routing; choose between static routes and dynamic routing protocols; and distinguish the roles of ACLs, application proxies, and stateful inspection. You will also create and inspect routes with Windows PowerShell, configure communication boundaries with Hyper-V virtual switches, and use macOS commands and settings to examine interfaces, routes, paths, and DNS responses. The aim is to connect individual commands to the forwarding and security decisions they produce, so that configuration becomes part of a coherent troubleshooting model rather than a collection of isolated procedures.

Technical requirements

The practical examples span several platforms, so no single workstation must reproduce every exercise. For the Cisco IOS examples, use Cisco Packet Tracer or an authorized Cisco IOS lab containing a multilayer switch and at least two endpoints. Obtaining Packet Tracer through Cisco Skills for All or Networking Academy may require a Cisco sign-in. The ASA inspection

examples require a Cisco Secure Firewall ASA lab or simulator that supports the commands shown; feature coverage in Packet Tracer can vary.

The Windows exercises require an administrator PowerShell session and, for the virtual-switch work, a Windows system with the Hyper-V role, a test virtual machine, and a spare or appropriately configured physical network adapter. The macOS exercises require Terminal access and an administrator account for commands that use `sudo`. Before altering a route, VLAN, or management adapter, record the existing configuration and make sure that you have a recovery path. A mistaken default route or VLAN assignment can disconnect a remote session immediately.

Segmenting and routing traffic with VLANs

This section follows a packet from a host in one logical network to a host in another. It shows why VLANs create separation, why crossing that separation requires Layer 3 forwarding, and how a switched virtual interface turns a multilayer switch into the default gateway for a VLAN. This foundation matters because routing protocols, firewall rules, and virtual switches all depend on the same boundaries being defined correctly.

A VLAN creates a logical Layer 2 broadcast domain on a switch. Two access ports assigned to different VLANs behave as if they were connected to different physical switches: ordinary Ethernet frames do not cross from one VLAN to the other. Each VLAN is therefore normally paired with a different IP subnet. A host in VLAN 10 might use `192.168.10.0/24`, while a server in VLAN 30 uses `192.168.30.0/24`. Because those addresses do not share a subnet, communication between them must pass through a device that can route at Layer 3.

A traditional design connects each subnet to a separate router interface. The design is valid and remains common at WAN or security boundaries, but adding many local VLANs can consume router ports and create an unnecessary bottleneck. A Layer 3 switch combines hardware switching with IP routing. It forwards frames at Layer 2 within a VLAN and packets at Layer 3 between VLANs, often at wire speed. The following figure compares these paths without implying that a multilayer switch replaces every router in a network.




Figure 5.1: Comparing a traditional routed design with inter-VLAN routing on a Layer 3 switch

Figure 5.1 shows that the logical requirement does not change: traffic still needs a gateway in each subnet. In the traditional design, physical router interfaces provide those gateways. In the multilayer design, **switched virtual interfaces (SVIs)** provide them inside the switch. A host in VLAN 10 sends remote-subnet traffic to 192.168.10.1; the switch consults its routing table and forwards the packet through the VLAN 30 SVI toward the destination. This is inter-VLAN routing.

Reading a multilayer switch configuration

The configuration must align three pieces of information: the VLAN database, the access-port assignments, and the SVI addresses. An access port normally carries untagged traffic for one VLAN. A trunk can carry traffic for multiple VLANs between switches, routers, firewalls, or

virtualization hosts, with tags identifying the VLAN. The SVI address becomes useful as a gateway only when the VLAN is active, the SVI is up, and IP routing is enabled on the switch.

On Cisco IOS, `show running-config` displays the saved command structure currently active in memory; it does not display a firewall state table. The illustrative console view below uses two access ports, two VLANs, and two SVIs. It recreates the type of view a video demonstration would show while keeping every label readable.

Reading a multilayer switch configuration

Illustrative Cisco IOS output showing access ports, VLANs, switched virtual interfaces, and IP routing

```
Switch# show running-config
!
ip routing
!
vlan 10
 name USERS
vlan 30
 name SERVERS
!
interface GigabitEthernet1/0/1
 switchport mode access
 switchport access vlan 10
!
interface GigabitEthernet1/0/2
 switchport mode access
 switchport access vlan 30
!
interface Vlan10
 ip address 192.168.10.1 255.255.255.0
 no shutdown
interface Vlan30
 ip address 192.168.30.1 255.255.255.0
 no shutdown
```



Access ports

Gi1/0/1 belongs to VLAN 10 and Gi1/0/2 belongs to VLAN 30.
Untagged outbound traffic enters the VLAN assigned to each port.



VLAN gateways

Vlan10 and Vlan30 are switched virtual interfaces. Their addresses act as the default gateway for hosts in each VLAN.



Routing enabled

The ip routing command enables forwarding between connected subnets. The configuration shown installs directly connected routes and permits inter-VLAN routing.

Figure 5.2: An illustrative Cisco IOS configuration for access ports, VLANs, SVIs, and inter-VLAN routing

In Figure 5.2, the access-port commands determine which broadcast domain receives an endpoint's frames. The interface Vlan10 and Vlan30 sections assign Layer 3 gateway addresses, and ip routing allows the switch to forward between the connected subnets. The addresses are

examples rather than values to copy blindly. They must be unique in their subnets, and endpoints must use the SVI in their own subnet as the default gateway.

The running configuration describes intended settings, but troubleshooting also requires operational evidence. The following commands answer different questions and should be read together:

```
show vlan brief
show interfaces switchport
show ip interface brief
show ip route
```

The first command confirms VLAN existence and access-port membership. The second shows the switching mode and VLAN behavior of a specific port. The third reveals whether the physical interfaces and SVIs are administratively and operationally up. The fourth confirms that the switch has connected routes for the SVI subnets. If VLAN 30 exists but its SVI is down, for example, looking only at the access-port assignment can lead to the wrong conclusion.

VLANs therefore establish the local boundaries, and SVIs make those boundaries routable. Once the network grows beyond directly connected subnets, the forwarding device also needs routes to more distant destinations. That leads to a design choice between routes maintained explicitly by an administrator and routes learned through a protocol.

Choosing between static and dynamic routes

The aim of this section is to connect the size and volatility of a network to the way its routes are maintained. Static routes offer predictability and little protocol overhead; dynamic routing offers automatic discovery and adaptation. Neither is universally better. A useful design assigns each method to the part of the network where its behavior is easiest to understand and operate.

A static route is an administrator-defined instruction for reaching a destination prefix through a particular next hop or interface. It works well for a small, stable branch, a single path to an upstream router, a tightly controlled management network, or a backup route with a less-preferred metric. Static routes do not advertise themselves, and they do not automatically learn a replacement path when the topology changes. That simplicity can be a strength, but only while the number of routes and changes remains manageable.

A dynamic routing protocol allows routers or Layer 3 switches to exchange reachability information. Each device calculates or selects routes according to the protocol's rules and installs preferred routes in its routing table. If a link fails, the protocol can withdraw the lost path and converge on an alternative when one exists. Convergence improves continuity, but it

is not instantaneous and does not guarantee that every application session survives. Redundant paths and sound protocol design are still required.

OSPF, EIGRP, and RIP inside an organization

Open Shortest Path First (OSPF) is a link-state interior gateway protocol. Routers advertise link-state information, build a common view of the topology within an area, and run a shortest-path calculation. OSPF uses cost as its route metric; lower total cost is preferred. Cost is commonly derived from interface bandwidth. Areas reduce the topology information every router must process, making OSPF suitable for a carefully structured enterprise.

Enhanced Interior Gateway Routing Protocol (EIGRP) is a Cisco-developed interior routing protocol that uses the Diffusing Update Algorithm and a composite metric. It can converge quickly and supports unequal-cost load balancing when configured appropriately. EIGRP was historically associated with Cisco-only environments; its core specification was later published as an informational RFC. Platform support and feature parity should still be verified before selecting it for a multivendor design.

Routing Information Protocol (RIP) is a distance-vector protocol whose principal metric is hop count. A route of 15 hops is the largest usable path; a metric of 16 means unreachable. RIP is simple and broadly understood, but its hop limit and convergence behavior make it a poor choice for most growing enterprise networks. RIP version 2 adds classless route information and multicast updates, yet it remains more useful in small labs and legacy environments than in a new campus design.

Comparing these protocols reinforces an important point: "best path" is protocol-specific. OSPF minimizes accumulated cost, EIGRP uses its composite metric, and RIP minimizes hop count. A path with more hops can be better under OSPF if its calculated cost is lower, while a physically shorter path is not automatically preferred by every protocol. The installed route also depends on administrative preference when several routing sources know the same destination.

BGP between autonomous systems

Border Gateway Protocol (BGP) addresses a different scale. It is a path-vector protocol used to exchange reachability between autonomous systems—networks under separate administrative control—and it forms the routing control plane of the public Internet. A BGP advertisement can carry the destination prefix, the sequence of autonomous systems in the path, the next hop, and other attributes. Those attributes support policy, loop prevention, traffic engineering, and route selection.

BGP does not promise the physically shortest or fastest route. An organization may prefer a path because of business relationships, security policy, capacity planning, or traffic-engineering goals. A longer AS path might still be selected if an earlier rule in the local decision

policy gives it preference. This policy-driven behavior is essential on the Internet, where independent networks must cooperate without sharing one internal metric.

Static routing is therefore most attractive when reachability is limited and stable; an interior protocol manages changing paths within an organization; and BGP exchanges policy-controlled reachability between autonomous systems. With those roles separated, the next task is practical: inspecting and changing the static routes held by an endpoint operating system.

Managing static routes with PowerShell

This section uses Windows PowerShell to turn routing concepts into a controlled workflow: inspect the current table, identify the correct interface, add one route, verify it, and remove or adjust it when necessary. The sequence matters. Creating a route before confirming its next hop and interface can produce a convincing configuration that sends traffic nowhere.

Open PowerShell as an administrator for commands that change the routing table. Begin by listing the routes. The interface index (`ifIndex`), destination prefix, next hop, route metric, and policy-store information are especially useful when several adapters or default routes are present. The following commands show a compact view and then every exposed property:

```
Get-NetRoute  
Get-NetRoute | Format-List -Property *
```

A destination of `0.0.0.0/0` is the IPv4 default route. More-specific prefixes take precedence over less-specific ones, so a `192.168.22.0/24` route is selected for that subnet even when a default route also exists. When routes have equally specific prefixes, metrics and policy affect the choice. Confirm the interface index with `Get-NetIPInterface` or by correlating the route output with the adapter's IP configuration.

Suppose interface 8 is attached to `192.168.21.0/24` and the router at `192.168.21.1` knows how to reach `192.168.22.0/24`. The following command creates that route:

```
New-NetRoute -DestinationPrefix "192.168.22.0/24" -InterfaceIndex 8 -NextHop  
"192.168.21.1"
```

The next hop must be reachable through the selected interface; it is not an arbitrary destination address. After creating the route, run `Get-NetRoute` again and test the path with `Test-NetConnection`, ping, or `tracert` as appropriate.

A successful route-table entry proves only that Windows accepted the configuration. It does not prove that the next-hop router has a return route, that a firewall permits the traffic, or that the destination service is listening.

If the route is wrong, use the same identifying values to remove it. PowerShell may request confirmation unless the command includes an appropriate confirmation option:

```
Remove-NetRoute -DestinationPrefix "192.168.22.0/24" -InterfaceIndex 8 -NextHop  
"192.168.21.1"
```

A critical limitation of `Set-NetRoute` is easy to miss: after a route is created, its `DestinationPrefix` and `NextHop` cannot be changed with that cmdlet. To redirect the example route from 192.168.21.1 to 192.168.21.2, remove the old route and create a new one. `Set-NetRoute` remains useful for editable properties such as `RouteMetric`. The following command changes the metric of the route identified by its prefix and next hop:

```
Set-NetRoute -DestinationPrefix "192.168.22.0/24" -NextHop "192.168.21.1" -  
RouteMetric 25
```

PowerShell's discoverability is useful when parameters differ across Windows releases. The next two commands find route-related help and open the detailed help for the creation cmdlet:

```
Get-Help *route*  
Get-Help New-NetRoute -Detailed
```

This workflow is more verbose than a one-line legacy route command, but its named parameters, object output, and pipeline support make it suitable for auditing and automation. It also establishes a habit that carries into firewall and virtual-switch work: first identify the traffic and interface, then change one control, and finally validate the resulting path.

Filtering traffic with firewalls

Routing determines where a packet can go; a firewall determines whether it should be permitted to cross a boundary. This section builds from stateless ACL entries to connection-aware inspection and application proxies. The aim is to show what context each control can evaluate, because security failures often arise when a rule is expected to understand more than the technology actually records.

A firewall can evaluate source and destination addresses, protocols, ports, direction, interface, connection state, application behavior, or user identity, depending on the platform. A rule might permit HTTPS to a public web server while denying unsolicited access to internal

systems. No firewall can compensate for an undefined trust boundary or an exposed management interface.

Stateless access control lists

An **access control list (ACL)** is an ordered set of permit and deny statements. A traditional router ACL evaluates packet fields independently rather than maintaining a complete connection record, which makes it a stateless filter. Rules are processed in order, and the first match determines the result; an implicit deny commonly follows the last explicit entry. Direction and placement therefore matter as much as the rule text.

The following Cisco IOS example permits inbound HTTPS to one internal server, blocks other traffic aimed at the 10.0.0.0/8 space, and then permits traffic to destinations not matched earlier. It is intentionally small so that the ordering is visible:

```
ip access-list extended WEB-IN
  permit tcp any host 10.0.0.10 eq 443
  deny ip any 10.0.0.0 0.255.255.255
  permit ip any any
interface GigabitEthernet1/0/24
  ip access-group WEB-IN in
```

The wildcard mask 0.255.255.255 matches the 10.0.0.0/8 network; it is not a subnet mask. The interface command applies the ACL in the inbound direction, so the rule set evaluates packets as they enter that interface. An ACL can be undermined by a poor rule, an incorrect direction, or missing anti-spoofing controls, but source-address spoofing does not magically bypass every ACL. Filters should reject implausible source addresses at network edges and allow only the minimum required traffic.

Application proxies and protocol inspection

An application proxy terminates the client's connection and opens a separate connection to the server. Because the two sides do not share one end-to-end transport session, the proxy can authenticate clients, normalize requests, enforce application-specific rules, and hide details of the protected service. That additional understanding can improve control, but it also consumes processing resources, introduces latency, and supports only traffic the proxy is designed to understand. A proxy is not automatically "more secure" in every deployment; its value depends on correct policy and protocol support.

Cisco ASA application inspection is related but should not be confused with a generic full proxy. Inspection engines can examine protocol commands, track dynamically negotiated connections, and enforce protocol-specific behavior within the ASA's stateful firewall architecture. The author's following policy-map example enables inspection for DNS, FTP, and ESMTP in the default inspection class:

```
policy-map global_policy
  class inspection_default
    inspect dns preset_dns_map
    inspect ftp
    inspect esmtp
```

The configuration identifies traffic for protocol-aware inspection; it does not mean that every packet is relayed through a standalone proxy process. Before using it, confirm the existing global policy, supported inspection maps, software version, and application requirements. Encrypted traffic also limits what a firewall can examine unless the design intentionally terminates or decrypts the session under an approved security and privacy policy.

Stateful inspection

A stateful firewall records active flows. For a TCP connection, the record normally includes source and destination addresses, source and destination ports, protocol, direction, TCP state, and timers. When an internal client opens an allowed HTTPS connection, the client selects an ephemeral source port and sends to the server's destination port 443. The server continues using port 443; it does not move the established conversation to a random port. The four address-and-port values plus the protocol distinguish this flow from other simultaneous connections.

The next figure shows the two decisions that state makes possible: policy evaluates a new outbound connection, while an existing state entry recognizes the corresponding return traffic. A separate inbound attempt with no state match remains subject to the inbound rule set and can be blocked.




Figure 5.3: How a stateful firewall handles a permitted connection and unsolicited inbound traffic

Figure 5.3 separates authorization from recognition. The outbound policy must first allow the new connection. The state table then allows packets that belong to that connection to return without a separate broad rule permitting every external source. For TCP, the three-way handshake establishes transport state and synchronizes sequence numbers; it does not prove that the human or server has a trusted identity. Authentication and encryption come from other controls such as TLS, IPsec, or application credentials.

A simplified state table makes those fields easier to compare. Real products expose different column names and may store NAT translations, interface identifiers, byte counts, flags, and timeout values in addition to the fields shown here.

Source IP	Source port	Destination IP	Dest. port	Protocol	State	Action
10.1.1.10	49152	172.217.5.110	443	TCP	ESTABLISHED	Allow
10.1.1.25	50120	8.8.8.8	53	UDP	ACTIVE	Allow
10.1.1.30	49321	203.0.113.20	80	TCP	ESTABLISHED	Allow
10.1.1.40	51004	198.51.100.15	443	TCP	SYN SENT	Pending
203.0.113.50	42000	10.1.1.20	3389	TCP	NO MATCH	Block

Table 5.1: A simplified state table used to classify active, pending, and unsolicited traffic

The first and third rows represent established TCP connections from different ephemeral source ports to well-known service ports. The UDP DNS row is marked active even though UDP has no handshake; the firewall creates a temporary pseudo-state after seeing permitted outbound traffic and expires it according to a timer. The SYN SENT row is pending because the TCP handshake has not completed. The final row has no matching state and is blocked. Idle and protocol-specific timeouts remove stale entries so that the table does not grow indefinitely.

Many firewall appliances also route between interfaces and perform Network Address Translation or Port Address Translation. These capabilities often appear in the same interface, but their purposes remain distinct: routing selects a next hop, translation changes address or port information, and policy decides whether traffic is allowed. Keeping those functions separate in the troubleshooting model helps explain why a packet can have a valid route and NAT rule yet still be denied by security policy.

We have now moved from defining network boundaries to controlling traffic across them. Virtualization applies the same ideas inside a single host: virtual switch ports create communication paths, VLAN tags preserve segmentation, and the selected switch type determines whether traffic can reach the host or physical network.

Building virtual network paths with Hyper-V

The purpose of this section is to make the Hyper-V virtual switch part of the same forwarding model as a physical switch. A virtual machine connects through a virtual network adapter to a software switch. Depending on that switch's type and VLAN configuration, the frame may remain between virtual machines, reach the management operating system, or pass through a physical adapter to the external network.

Open an elevated PowerShell session and identify the physical adapter before creating an external switch. The following command creates VMSwitch10 and binds it to the adapter named Ethernet 3:

```
New-VMSwitch -Name "VMSwitch10" -NetAdapterName "Ethernet 3"
```

An external switch changes how the host uses that adapter, so perform this step from a console or through a separate management path whenever possible. The resulting output identifies the switch as External and lists the physical adapter description. If the host operating system must share the adapter, enable that relationship explicitly:

```
Set-VMSwitch -Name "VMSwitch10" -AllowManagementOS $true
```

Hyper-V creates a virtual adapter for the management operating system when sharing is enabled. The following command places that management adapter in access VLAN 100:

```
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName "VMSwitch10" -Access -VlanId 100
```

The `-ManagementOS` parameter is important: this command tags the host's virtual adapter, not every virtual machine connected to the switch. A VM that should use the same access VLAN needs its own assignment, such as `Set-VMNetworkAdapterVlan -VMName "VM1" -Access -VlanId 100`. The physical switch port must also carry VLAN 100 according to the physical design. Otherwise, a syntactically correct Hyper-V setting can isolate the host.

Because the requested video frame is unavailable, the following professional reconstruction shows the settings that the commands produce rather than presenting itself as a literal screenshot of a particular Windows release.

Hyper-V external switch with management VLAN 100

You can configure a management VLAN on the external switch for remote management of the host and switch.

The screenshot shows the 'Virtual Switch Manager' window. On the left, under 'Virtual switches', 'vSwitch0' is selected and labeled 'External'. The main area, 'Virtual switch properties', shows the 'Name' as 'vSwitch0'. Under 'Connection type', 'External network' is selected with a dropdown menu showing 'Ethernet 1'. The checkbox 'Allow management operating system to share this network adapter' is checked. Under 'VLAN ID', the checkbox 'Tag the management operating system with VLAN ID:' is checked, and the 'VLAN ID' is set to '100'. A blue information box at the bottom states: 'This switch provides connectivity to the physical network, and the host operating system is configured with VLAN ID 100 for management access.'

Figure 5.4: Illustrative Hyper-V settings for an external switch and management VLAN 100

Figure 5.4 ties the graphical choices to the PowerShell workflow: switch name and physical adapter define the external path, management sharing gives the host a virtual adapter on that path, and VLAN 100 determines the host adapter's Layer 2 segment. The figure also highlights why this is a coordinated change. The virtual setting, physical switch configuration, IP subnet, and management route must agree before connectivity is tested.

Controlling virtual adapters with PowerShell

PowerShell exposes controls that are easy to overlook in a graphical workflow. Each setting changes a specific risk or performance behavior and should be enabled only after its effect is understood. The following examples use VM1:

- DHCP Guard prevents the VM from sending DHCP server messages:

```
Set-VMNetworkAdapter -VMName "VM1" -DhcpGuard On
```

- Router Guard blocks unauthorized IPv6 Router Advertisement and redirection messages:

```
Set-VMNetworkAdapter -VMName "VM1" -RouterGuard On
```

- Port mirroring makes a virtual adapter a monitoring source or destination:

```
Set-VMNetworkAdapter -VMName "VM1" -PortMirroring Source
```

- MAC address spoofing allows frames whose source MAC differs from the address assigned by Hyper-V:

```
Set-VMNetworkAdapter -VMName "VM1" -MacAddressSpoofing On
```

Hyper-V also supports bandwidth controls, hardware-assisted acceleration, access and trunk VLAN modes, private VLAN modes, and teaming inside supported guests. These are not blanket optimizations: hardware support, security exposure, and the operational reason for each exception must be verified.

Choosing external, internal, and private switch types

The three Hyper-V switch types define different communication boundaries. An external switch connects through a physical network adapter, allowing attached VMs to reach the physical network when addressing and policy permit; the host participates only if management OS sharing is enabled. An internal switch connects attached VMs to one another and to the Hyper-V host but has no direct physical-adapter connection. A private switch connects attached VMs to one another while excluding both the host and the physical network.

The following comparison makes those boundaries explicit. It replaces a missing video screenshot with an instructional view whose purpose is to show reachability rather than reproduce one version of the management console.

Choosing a Hyper-V virtual switch type

Choose the virtual switch type based on whether the switch connects to a physical adapter and the host connectivity requirements.




Figure 5.5: Communication boundaries of external, internal, and private Hyper-V switches

Figure 5.5 corrects an easy source of confusion: internal and private do not mean the same thing. The host can communicate on an internal switch, which makes that type useful for host-based testing, routing, or services. The host cannot communicate directly on a private switch, which creates a narrower VM-only segment. Neither type provides Internet access by itself, although a separately configured router or NAT service could deliberately forward traffic from an internal network.

The following command changes an existing switch to the internal type on supported Hyper-V versions:

```
Set-VMSwitch -Name "VMSwitch10" -SwitchType Internal
```

Changing a switch type can disconnect adapters and alter the host's management path. Returning to an external design also requires a valid physical-adapter binding, so do not assume that replacing Internal with External is sufficient in every version or configuration. Verify with `Get-VMSwitch`, `Get-VMNetworkAdapter`, and `Get-VMNetworkAdapterVlan`, then test from the host and from each affected VM. This validation closes the loop between the intended virtual topology and the traffic that can actually leave it.

Managing routes and interfaces on macOS

The final practical section applies the same evidence-first method to macOS. It begins with the interface and routing table, then moves to a targeted route change, graphical IPv4 configuration, reachability, path tracing, and DNS. The commands differ from Windows, but the questions remain familiar: which address is assigned, which route matches the destination, which gateway is used, and whether name resolution is returning usable data.

Run `ifconfig` to display interface information. Interface names such as `en0` and `en1` do not have a universal mapping to Ethernet or Wi-Fi across every Mac, adapter, and macOS version, so identify the active interface by its status, hardware port, and assigned address rather than memorizing a number. The basic command is:

```
ifconfig
```

Suppose the Mac is attached to `192.168.15.0/24` and a router at `192.168.15.1` provides a path to `192.168.16.0/24`. The following command adds a temporary network route. `sudo` prompts for an administrator password because the routing table affects system-wide traffic:

```
sudo route -n add -net 192.168.16.0/24 192.168.15.1
```

Verify the entry before testing the destination. The traditional `route` command modifies the current kernel table, and the change may not persist across restarts or network-service changes. Use the organization's supported management method for a durable route.

The default route should be changed only in a lab or from a local console with a recovery plan. Removing it can immediately break remote management and Internet access. These commands remove and then restore an IPv4 default route through 192.168.15.1:

```
sudo route delete default
sudo route add default 192.168.15.1
```

A more specific route such as 192.168.16.0/24 normally leaves the existing Internet path intact, which is why a targeted route is safer than replacing the default for a single internal destination. The route must still be paired with a return path and any required firewall permission.

Checking IPv4 settings in the macOS interface

The graphical settings provide a second view of the same configuration. In System Settings, open the active network service, select Details, and then open TCP/IP. The exact labels and layout vary by macOS release, but the pane normally shows how IPv4 is configured and lists the current address, subnet mask, and router. The following source screenshot shows a DHCP configuration.




Figure 5.6: Viewing DHCP-provided IPv4 settings in macOS

Figure 5.6 provides the local inputs used in a routing decision. The subnet mask defines which IPv4 destinations are local, and the router supplies the default next hop for remote destinations. When DHCP is selected, these values normally come from the DHCP server; editing one value manually without understanding the assigned subnet can create a duplicate address or an unreachable gateway.

Selecting **Manually** exposes fields for a fixed address, subnet mask, and router. The next screenshot shows those fields completed before the change is applied.




Figure 5.7: Entering a manual IPv4 address, subnet mask, and router in macOS

Figure 5.7 represents a complete manual IPv4 tuple, not merely an address change. The chosen address must be unique, the mask must match the intended subnet, and the router must be reachable on that subnet. DNS settings may also need to be supplied separately. Record the original DHCP or manual values before clicking **OK** so that the interface can be returned to a known state.

Reading routes, reachability, paths, and DNS

Use the routing table to verify what the system will do with traffic. The numeric option avoids waiting for reverse-DNS lookups and keeps addresses visible:

```
netstat -nr
```

The Internet section describes IPv4 routes and the Internet6 section describes IPv6 routes. An `fe80::/10` entry is link-local, not a private address in the IPv4 sense. It supports essential on-link IPv6 behavior and is not forwarded by routers. Entries beginning with `ff01` or `ff02` represent multicast scope used by the IPv6 stack. They may not be the focus of an IPv4 investigation, but they should not be dismissed as meaningless; their presence reflects normal protocol operation.

The `scutil network-information` view summarizes active reachability and DNS-capable interfaces. Run the following command:

```
scutil --nwi
Network information
IPv4 network interface information
  en0 : flags 0x5 (IPv4,DNS)
        address : 192.168.15.112
        reach   : 0x00000002 (Reachable)
IPv6 network interface information
  No IPv6 states found
Network interfaces: en0 en1
```

This sample reports an active IPv4 state on `en0` and no active IPv6 state at that moment. Reachable describes the system's current network reachability assessment; it is not proof that every remote application is reachable. Also remember that `en0` and `en1` can refer to different hardware on another Mac. Correlate the address and service name before changing an interface.

After the local address and route look correct, trace the path to a known destination. macOS uses the full `traceroute` command name:

```
traceroute 8.8.8.8
```

Each displayed hop is a router that returned a time-exceeded response. Asterisks indicate that a response was not received within the timeout; they do not always mean that forwarding stopped, because routers and firewalls may rate-limit or block diagnostic responses while forwarding later packets. Compare where replies end with the routing and firewall boundaries already identified.

Finally, query DNS separately from the path to distinguish a naming problem from an IP-connectivity problem:

```
nslookup microsoft.com
```

The output identifies the resolver that answered and returns available address records for the name. A missing server name, a non-authoritative response, or the absence of one record type does not by itself prove that the domain is unreachable. Look for timeouts, explicit DNS errors such as NXDOMAIN, returned A or AAAA records, and whether those addresses can be reached. If an IP address works while the name fails, investigate DNS configuration; if neither works, return to the local address, route, gateway, firewall, and upstream path.

Across macOS and Windows, a static route has the same essential meaning: a destination prefix is associated with a next hop or interface. Platform tools merely expose that decision differently. By tracing the decision from interface configuration through the route table and then through security policy, an administrator can locate a fault without treating every failed connection as a generic "network problem."

Summary

Network infrastructure determines where traffic can travel and where it should be stopped. In this chapter, you followed that decision from the local broadcast domain outward: VLANs created Layer 2 boundaries, switched virtual interfaces and routers moved packets between IP networks, and static or dynamically learned routes extended reachability beyond directly connected subnets. OSPF, EIGRP, RIP, and BGP also demonstrated that a route described as the "best" path is always the result of a particular metric, policy, and administrative design rather than a universal property of the network.

You then applied those forwarding decisions across different platforms and security boundaries. PowerShell and macOS tools exposed the routes an endpoint would actually use, while ACLs, application inspection, and stateful filtering showed that knowing where a packet should go is different from deciding whether it should be permitted to continue. Hyper-V virtual switches extended the same reasoning into virtualized networks, where external, internal, and private connectivity create distinct communication boundaries even when the systems share one physical host.

The practical lesson is to follow the traffic rather than the configuration screen. Identify the source interface and subnet, determine which route matches the destination, verify the next hop and return path, and then examine the security control that governs the flow. A configuration that looks correct is only the intended state; operational evidence tells you whether the packet can actually complete the journey. That distinction between intended behavior and observed behavior is what turns network configuration into a troubleshooting model.

The infrastructure is now in place, but real incidents rarely arrive with the failing VLAN, route, or firewall rule already identified. They arrive as reports that email is down, a cloud application is unreachable, or the network is slow. In the next chapter, we will turn the technical model developed so far into a systematic troubleshooting process, beginning with impact and scope before using evidence to narrow a broad symptom to the dependency that is actually failing.

6

A Systematic Approach to Network Troubleshooting

Network troubleshooting is the process of turning an imprecise symptom into a small set of testable causes. A user may report that email is unavailable, a cloud service cannot be reached, or the network is running slowly, but the visible symptom may originate in DNS, IP addressing, routing, wireless connectivity, a firewall, or the service itself. A systematic approach prevents troubleshooting from becoming a series of random configuration changes by establishing the scope of the problem first and then narrowing the fault with evidence.

In this chapter, we begin by looking at network severity triage and the role stakeholders play in establishing the real business impact of an incident. From there, we examine common problems found in organizational networks and learn how their scope and symptoms can point us toward the most useful first checks. We then follow the path users take to corporate resources such as email, cloud platforms, VPN-connected services, and conferencing systems. Finally, we move to the endpoint itself, where IP addressing, subnet masks, DNS, default gateways, and tools such as `ipconfig`, `ping`, and `tracert` allow us to verify whether the problem is local or farther along the network path.

By the end of the chapter, you should be able to take an unstructured report such as "the network is down" and turn it into a deliberate troubleshooting path. The goal is not simply to know more commands, but to know which question to ask and which test to run next, so that you can move from business impact to the failing network dependency without wasting time on components that are already working.

In this chapter, we will cover the following topics:

- Network severity triage and stakeholder involvement
- Common network issues in organizational environments

- Troubleshooting corporate resource access
- IP structure and addressing in network troubleshooting

Technical requirements

To follow the examples in this chapter, you should understand the basic functions of network adapters, switches, routers, wireless access points, and firewalls. Familiarity with the OSI or TCP/IP model is useful because it gives you a structured way to isolate faults at different layers. The command-line examples use Windows Command Prompt, and the gateway demonstration assumes that you can view and modify the IPv4 settings of a test system using a static IP address.

Network severity triage and stakeholder involvement

Network severity triage determines what receives attention first. Before opening a command prompt or changing a configuration, establish who is affected, which services are unavailable, and which business processes are blocked. A practical triage process combines technical impact with communication and an agreed priority for restoring service.

Before technical diagnosis begins, the team needs a shared view of what deserves attention first. Frequency, impact, persistence, and stakeholder priorities provide that view, helping you decide what must be restored first and how urgently the response should proceed.

Making that priority defensible requires more than deciding that an outage feels serious. A small set of practical criteria gives the team a consistent way to judge whether an incident is isolated and manageable or recurring, disruptive, and worthy of immediate escalation.

Criteria for assessing network severity

Several factors help determine the severity of a network incident.

- **Frequency** tells you whether the problem is rare or recurring.
- **Impact** describes how many users and services are affected and how difficult the problem is to work around.
- **Persistence** distinguishes a short-lived event from an issue that will continue to interrupt users until the underlying cause is corrected.

These factors should be considered together: a recurring fault that affects a revenue-critical service, for example, deserves a different response from an isolated issue that has an immediate workaround.

Severity is more reliable when it is established collectively rather than by the IT team alone. The technical symptoms describe what has failed, but the people who depend on the affected service can explain the operational consequence.

Different groups can assign very different priorities to the same outage. The IT team may consider email the most visible service, while company leadership may be more concerned about an accounting or line-of-business system. During a significant outage, bring together the stakeholders who understand the business impact, such as department heads, owners, and critical contractors. The organization should also maintain a notification method that does not depend entirely on the affected network, because email or messaging may be unavailable at the point when coordination matters most. Combining technical evidence with stakeholder input produces a more accurate restoration order.

At this point, you know how to establish the scope and business priority of an incident before touching the network. That matters because the same technical symptom can demand a very different response depending on who and what it affects. With the priority established, the next useful question is what pattern the failure follows, because recognizing that pattern will determine where the technical investigation should begin.

Common network issues in organizational environments

Once severity has been established, the next step is to classify the type and scope of the failure. Some incidents affect one endpoint; others affect a department, a site, or the entire organization. Recognizing common patterns helps you select the first useful test instead of treating every report as a unique problem.

This section therefore focuses on recognizing the shape of a problem before choosing a troubleshooting tool. By comparing isolated and widespread failures across malware, wireless access, addressing, remote connectivity, and shared services, you can reduce the number of plausible causes before making any change to the environment.

A malware outbreak that affects core business systems should be treated as a high-severity incident because it can spread and disrupt multiple services at once. Infection can begin through email attachments, malicious software, compromised advertisements, or links delivered through social platforms. In this situation, containment and communication may be more urgent than troubleshooting an individual user's connectivity.

Loss of access is another common symptom, but its severity depends on scope. Determine quickly whether the problem affects one user, a small group, a particular network segment, or everyone. A single endpoint may point to local addressing, adapter, or firewall settings, whereas a wider outage suggests that you should look at shared infrastructure or services first.

Wireless problems require the same scope check. Wi-Fi uses shared radio bandwidth and can be affected by interference, congestion, coverage, or access-point failures. In environments with centrally managed wireless infrastructure, the controller or management platform can help you compare affected users, access points, and locations before you begin changing settings on individual devices.

Not every failure results in a clean loss of connectivity. When the network remains reachable but performs poorly, or when resources work by IP address but fail by name, the investigation shifts from basic availability to two less obvious causes: performance degradation and name resolution.

Network slowness and DNS failures

Slowness can be harder to diagnose than a complete outage because traffic is still flowing, just not at the expected rate. DNS failures are also a common cause of apparent outages: a user may have IP connectivity but still be unable to reach a resource by name. When symptoms suggest name-resolution problems, verify DNS early so that you do not spend time troubleshooting a working physical or routing path.

IP address exhaustion can create another broad class of access failures. A user may now connect several devices, such as a laptop, phone, tablet, or VoIP handset, and IoT devices add further demand. If a DHCP scope has no leases available, new clients cannot obtain normal addressing until addresses are released or the available pool is expanded.

Remote access issues often add another layer because the user must first reach the internet and then establish the VPN path. A VPN can fail when a required protocol or port is blocked by a local, upstream, or corporate firewall. SSL-based VPNs may be easier to use across restrictive networks because they commonly operate over HTTPS. During triage, separate general internet connectivity from the VPN-specific path so that you know which part of the connection is actually failing.

These symptoms can usually be grouped by whether they cause a partial loss of resources or a total loss of connectivity. That distinction, combined with the number of affected users, gives you a useful starting point for deciding how broadly to investigate.

Do not begin by guessing at the technical root cause. First confirm the impact and severity with the appropriate stakeholders, then move into diagnosis. This prevents a team from spending valuable time on a visible but low-impact symptom while a more important service remains unavailable.

Once the scope and priority are clear, use the affected resource to choose the first tests. Email, web applications, internal servers, cloud services, and conferencing tools depend on overlapping network components, but each provides different clues about where the path may be breaking down.

Most organizations rely on a familiar set of shared services even when the underlying applications differ. As a result, troubleshooting often starts by checking how the user reaches the resource, how the resource is named, and whether any firewall, routing, or remote-access component sits between the two endpoints.

For email problems, first determine where the mail system is hosted. With an on-premises service, verify that users can reach the local mail server and resolve its name correctly. With Microsoft 365 or another cloud-hosted platform, confirm internet connectivity and make sure the organization's DNS records and local name resolution direct clients to the correct service. This distinction immediately tells you whether to focus on the internal network, the external path, or the service configuration.

Also identify how the user is connecting. A client may reach email from the LAN, over Wi-Fi, through a remote VPN, or through Microsoft Outlook using MAPI. If one path works and another does not, that comparison helps narrow the fault. For example, successful access from the LAN but not through the VPN points you toward the remote-access path rather than the mailbox itself.

Email can also be accessed through Outlook on the web over HTTPS or, in some environments, through protocols such as POP3 and IMAP. Verify that the required traffic is allowed by both the endpoint firewall and the network firewall, especially for users connecting from outside the corporate network. The goal is to confirm each dependency in the path rather than assuming that an email symptom must originate on the mail server.

DNS deserves particular attention when users can reach some resources but not others. If the public DNS name of the organization is also used internally, the internal DNS service must return the correct records for local resources while public DNS must return the records intended for external users. A wrong answer on either side can direct a client to the wrong destination even when the network itself is functioning.

When external sites are unavailable, rule out the broader internet path as well. Test more than one known external destination and check whether the ISP is experiencing an outage. This simple comparison helps distinguish an organization-wide connectivity problem from a failure that belongs to a single external service.

You now have two pieces of information that materially narrow an incident: its business impact and the pattern of the failure. The next step is to follow the user's actual path to the affected resource and identify the first dependency that stops behaving as expected. Doing that converts a broad symptom such as "I cannot access the service" into a testable question about DNS, addressing, routing, firewalling, remote access, or the service itself.

Troubleshooting corporate resource access

When a user cannot reach a corporate resource, identify the starting point, the destination, and the first point at which the expected path stops working. Routing and DNS problems can look similar from the user's perspective, so check both reachability and name resolution. Internal DNS servers resolve private resources and may forward other queries to an upstream or public resolver. If internal names work but an external corporate service does not, verify the relevant public DNS record and the external path rather than continuing to change internal records.

Once the symptom has been classified, troubleshooting becomes an exercise in tracing the path to the affected resource. Looking at the user's starting point, the destination, and the dependencies between them helps you decide whether the failure belongs to the endpoint, DNS, DHCP, routing, a firewall, a VPN or cloud connection, or the destination service.

Address allocation and endpoint controls are also common failure points. If a DHCP scope has exhausted its available addresses, expand the pool or add another appropriately configured scope. On the endpoint, check whether the built-in operating system firewall or a security product is blocking the required traffic. Remote devices introduce an additional challenge: if the system is offline and cannot be reached through remote-support tools, hands-on access may be required before troubleshooting can continue.

Cloud resources should be approached as another end-to-end path. Confirm that the required port is open and, where the environment uses a VPN tunnel or other private connection to Azure, AWS, or another provider, verify that the tunnel is operational. Compare the failing cloud resource with other external resources. If general internet access and other vendors' services work normally, concentrate on the route, tunnel, DNS, or service configuration associated with the affected cloud platform.

Most resource-access problems can initially be judged by a simple question: can the user reach the service or not? Real-time communication adds another dimension, because a connection can remain technically available while latency, jitter, or insufficient bandwidth makes it unusable. That makes conferencing a useful example of why successful connectivity does not always mean a healthy network path.

Video and audio conferencing quality issues

Video and audio conferencing can fail completely or degrade into choppy sound, frozen video, or high delay. Check whether **quality of service (QoS)** is configured to prioritize real-time traffic and whether large downloads, media streaming, or other bandwidth-heavy activity is competing for the same link. Pay attention to upload capacity as well as download capacity; many access links are asymmetric, so the upstream path can become the bottleneck even when download speed appears adequate.

The following figure condenses these resource-level checks into a troubleshooting map, helping you choose the next test according to what is failing, what still works, and how the affected resource is reached.




Figure 6.1: Troubleshooting paths for common corporate resource access issues

The diagram moves from incident scope into service-specific troubleshooting paths, showing that email, cloud resources, VPN-connected services, and conferencing problems do not all point to the same dependency. Use it as a branching checklist: compare the failing path with a known-working path and let the first meaningful difference determine which layer or service to test next.

So far, you have worked from the outside in: first establishing the impact of the incident, then classifying the symptom, and finally tracing the resource path. Those steps tell you where to investigate. The endpoint's IP configuration now gives you the first concrete set of values you can verify to determine whether the device is correctly participating in that path.

IP structure and addressing in network troubleshooting

After you have established the scope of the outage and the resource path involved, inspect the endpoint's IP configuration. Addressing problems can prevent a device from reaching only remote networks, only named resources, or everything beyond the local segment, so the fields on the network adapter provide a useful checkpoint before you investigate deeper infrastructure.

At this point, the investigation reaches the endpoint configuration that underpins nearly every network connection. Verifying the protocol and address, subnet information, DNS, gateway, and command-line results helps distinguish a local problem from a failure farther along the network path.

Before examining those fields individually, use the following figure as a checklist of the IP configuration elements that must work together for a host to communicate locally, resolve names, and reach networks beyond its own subnet.




Figure 6.2: Key IP configuration elements to verify during network troubleshooting

The figure brings together the address, subnet or prefix information, default gateway, and DNS configuration that you will keep returning to throughout this section. A problem with any one of these elements produces a different pattern of failure, which is why effective troubleshooting checks them as related dependencies rather than treating "IP configuration" as a single setting.

First determine whether the connection you are troubleshooting uses IPv4, IPv6, or both. Modern operating systems can enable both protocol stacks on the same network adapter, and

each stack maintains its own addressing information. The examples in this chapter focus on IPv4 because the private addressing scheme is easy to illustrate, but the same troubleshooting principle applies to either protocol: confirm that the address, subnet or prefix, gateway, and name-resolution settings match the network the device is expected to use.

The first relationship to validate is between the IP address and the subnet mask. Together, they determine which destinations the host believes are local, so an incorrect subnet boundary can cause traffic to be sent in the wrong direction even when the IP address itself appears valid.

Subnetting and subnet mask consistency

Subnetting errors are a frequent source of confusing connectivity problems. Consider the address 192.168.1.5 with the subnet mask 255.255.255.0, or /24. In that example, 192.168.1.0 is the network, and .5 identifies the host. Devices that are intended to communicate directly on that subnet need compatible addressing and unique host addresses. If one device uses an incorrect mask, the hosts may disagree about which destinations are local and which must be sent to a router. Check the address and subnet mask together before assuming that the problem belongs to DNS or an application.

A correct address and subnet mask establish where the host sits on the network, but they do not yet tell it how to reach another network or translate a resource name into an address. Those two responsibilities belong to the default gateway and DNS, making them the next logical checks once local addressing has been validated.

DNS and gateway configuration

A usable IP configuration normally includes more than an address and subnet mask. The client also needs a valid DNS server for name resolution and, when it must reach another subnet or the internet, a default gateway. If local communication works but named resources do not, investigate DNS. If local communication works but remote networks do not, the gateway and routing path become the next places to look.

The default gateway is the router a host uses to reach destinations outside its local subnet. On Windows, the `ipconfig` command provides a quick view of the adapter's addressing information and is a useful first check when you suspect a configuration problem. Start by establishing a baseline for the endpoint's current addressing configuration with the following command:

```
ipconfig
```

The command displays the current IP configuration, including the IPv4 address, subnet mask, and default gateway. In a typical result, you might see an IPv4 address such as 192.168.21.110 with a subnet mask of 255.255.255.0. With a /24 mask, the network is 192.168.21.0 and the

final octet identifies the host on that subnet. Some adapters can have more than one IP address, so verify which address belongs to the network path you are troubleshooting. Pay particular attention to the default gateway because it provides the route beyond the local subnet. Once the local configuration looks reasonable, use `tracert` to inspect the path to a remote destination. The `-d` switch skips reverse DNS lookups and keeps the output focused on IP addresses:

```
tracert -d 8.8.8.8
```

Each line in the trace represents a hop along the route to the destination. A sudden increase in latency or a timeout can be useful evidence, but an individual router may also ignore or deprioritize traceroute probes, so a missed reply does not automatically mean that forwarding has failed. Focus on whether the trace ultimately reaches the destination and where a consistent failure begins. If the destination is not reached within the maximum hop count, investigate routing problems such as an unreachable path or a loop.

You can demonstrate the role of the gateway on a test system that uses a static IPv4 address. Open the adapter's IPv4 settings from Control Panel and Network and Sharing Center, temporarily clear the default gateway, save the change, and then rerun the same `tracert` command. Perform this only on a lab or test system where intentionally removing external connectivity will not disrupt users.

Without a default gateway, the adapter has no route to destinations outside its local subnet. A trace to 8.8.8.8 therefore fails immediately. The following output shows the behavior in this example:

```
C:\>tracert -d 8.8.8.8

Tracing route to 8.8.8.8 over a maximum of 30 hops

  1  Transmit error: code 1231.

Trace complete.

C:\>
```

The absence of a gateway does not prevent the host from communicating with devices that it considers local. Because another address on the 192.168.21.0/24 network can be reached directly, the system can still exchange traffic with that device even though remote networks are unavailable.

The following ping test confirms that local connectivity still works:

```
C:\>ping 192.168.21.8

Pinging 192.168.21.8 with 32 bytes of data:
Reply from 192.168.21.8: bytes=32 time=1ms TTL=128
```

This comparison is useful because it separates local addressing from off-subnet routing. If the local ping succeeds but the remote `tracert` fails after the gateway is removed, the test confirms that the adapter and local subnet are still functioning while the route beyond that subnet is missing. Restore the correct gateway after the test, then verify the complete IP configuration before moving on to more complex routing, firewall, or application checks. The following figure shows the addressing, gateway, DNS, and IPv6 information that should be reviewed together before the investigation moves farther into the network.

```
Ethernet adapter vEthernet (External):

Connection-specific DNS Suffix . . : 
Description . . . . . : Hyper-V Virtual Ethernet Adapter #2
Physical Address. . . . . : 04-92-26-D6-B2-FA
DHCP Enabled. . . . . : No
Autoconfiguration Enabled . . . . : Yes
Link-local IPv6 Address . . . . . : fe80::ec7f:e14d:1aa9:a2f1%16(Preferred)
IPv4 Address. . . . . : 192.168.15.2(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 192.168.15.1
DHCPv6 IAID . . . . . : 402952742
DHCPv6 Client DUID. . . . . : 00-01-00-01-29-D0-65-D0-04-92-26-D6-B2-FA
DNS Servers . . . . . : 1.1.1.1
DoH: https://cloudflare-dns.com/dns-query
```

Figure 6.3: Example adapter configuration showing IPv4 address, subnet mask, default gateway, DNS, and IPv6 information

The adapter output provides the addressing checklist in one place. In this example, the IPv4 address, subnet mask, default gateway, and DNS server are all present, while IPv6 also has a link-local address. A link-local IPv6 address can exist even when no IPv6 DHCP service is providing a routable address. Reading these fields together helps you decide whether the next troubleshooting step belongs on the endpoint or farther along the network path.

Summary

Effective network troubleshooting begins before the first command is entered. The most important habit developed in this chapter is to resist treating the user's symptom as the cause. Instead, establish how many users and services are affected, understand the business consequence with the relevant stakeholders, and then use the scope of the failure to decide where the technical investigation should begin.

From there, the problem becomes a process of narrowing. A failure affecting one endpoint suggests a different starting point from a failure affecting an entire department. A resource that works by IP address but not by name points toward DNS, while successful communication within the local subnet combined with failed remote access directs attention toward the default gateway or routing path. Email, VPNs, cloud resources, conferencing systems, DHCP, firewalls, and wireless services may appear to be separate problems, but the same method applies to each: identify what still works, compare it with what fails, and find the first dependency at which their paths diverge.

The `ipconfig`, `ping`, and `tracert` examples turned that method into concrete tests. By reading the IP address, subnet mask, gateway, and DNS information together, and then comparing local and off-subnet reachability, you can determine whether the investigation should remain on the endpoint or move deeper into the network. That is the central troubleshooting skill to carry forward: **reduce the problem deliberately until the evidence tells you where the failure actually belongs.**

The troubleshooting method developed here assumes that you can inspect at least some of the network path and progressively test its dependencies. Remote and home environments make that harder because the user, endpoint, Wi-Fi connection, router, ISP path, and corporate access method may all sit partly outside the administrator's direct control. In the next chapter, we will apply the same discipline of establishing scope, comparing working and failing paths, and testing from the endpoint outward to troubleshoot remote and home networks where visibility and control are much more limited.

Part 3

Troubleshooting Access at the Network Edge

The troubleshooting method becomes more demanding when parts of the network are shared, remote, or outside the administrator's direct control. This part applies the method from *Chapter 6* to three environments where the visible symptom can sit several dependencies away from the actual fault.

You will begin with remote and home networks, tracing problems from the user's endpoint through Wi-Fi, the home router, the ISP, remote-access services, and the corporate environment. You will then scale wireless troubleshooting into the enterprise, where multiple access points, interference, client density, channel use, and shared airtime can create performance problems even when clients remain connected. Finally, you will isolate VPN failures by separating tunnel establishment from the addressing, DNS, routing, and resource-access dependencies that determine whether the tunnel is actually useful.

The common thread is comparison. In each environment, you will learn to identify what still works, locate the first boundary where behavior changes, and investigate the dependencies at that boundary instead of resetting unrelated components.

This part of the book includes the following chapters:

- *Chapter 7, Troubleshooting Remote and Home Networks*
- *Chapter 8, Troubleshooting Enterprise Wireless Networks*
- *Chapter 9, Troubleshooting VPNs and Remote Access*

7

Troubleshooting Remote and Home Networks

Remote and home network troubleshooting extends the same diagnostic principles used inside an organization into an environment where much of the path may be outside your direct control. The endpoint may be in a user's home, the connection may cross Wi-Fi and a consumer router before reaching an ISP, and access to corporate resources may depend on VPNs, firewalls, authentication, and services that the user cannot see. The challenge is therefore not simply restoring connectivity, but determining which part of this longer path is actually failing.

In this chapter, we begin with the common symptoms encountered when supporting users remotely and then follow the remote-access path through Remote Desktop permissions, VPN configuration, authentication, certificates, and firewall rules. We then examine printing and local network discovery before moving into slow or intermittent connections, where Wi-Fi association, addressing, DNS, access-point configuration, and channel selection become important. Finally, we work outward from the home computer through its adapter, router, and ISP connection using tools such as `ipconfig`, `pathping`, and `ARP`.

By the end of the chapter, you should be able to take a remote user's report and identify the first boundary at which a known-working connection becomes a failing one. The goal is to determine whether the problem belongs to the endpoint, the home network, the remote-access path, the corporate environment, or the upstream provider before changing configuration unnecessarily.

In this chapter, we will cover the following topics:

- Common issues in remote and home networks
- Troubleshooting remote desktop and VPN connectivity

- Troubleshooting printing and local network discovery in remote environments
- Diagnosing slowness, internet access, and dropped connections
- Troubleshooting internet connectivity from the home network

Technical requirements

To follow the examples in this chapter, you should be familiar with **Remote Desktop Protocol (RDP)**, VPN concepts, Windows printing, and basic home networking technologies such as Ethernet, Wi-Fi, routers, DHCP, and NAT. You should also be comfortable using common troubleshooting utilities such as `ipconfig`, `ping`, `tracert`, and `nslookup`. The chapter assumes that you can move between the Windows graphical network settings and Command Prompt while tracing a problem from the endpoint toward the destination.

Common issues in remote and home networks

Remote and home networks produce a familiar set of symptoms: a user cannot reach the corporate environment, a printer does not appear in a remote session, nearby devices are missing from the network view, Wi-Fi is slow or unstable, or internet access disappears completely. The symptom is only the starting point. The underlying cause may be local to the computer, related to addressing or DNS, caused by the home router or wireless link, or located farther upstream with the ISP or corporate network.

Supporting a home user also changes the troubleshooting process because you may not be able to inspect cables, indicators, logs, or error messages directly. Begin by establishing what still works: can the user reach the home router, browse the internet, resolve names, connect to the VPN, or reach the corporate service? Each successful test narrows the path.

Troubleshooting remote desktop and VPN connectivity

Remote access depends on a chain of components rather than a single setting. In this section, you will follow that chain from Remote Desktop availability and user permissions through VPN authentication and certificates to the firewall and NAT rules that ultimately determine whether traffic can reach the intended system.

Start Remote Desktop troubleshooting with access on the destination system. Confirm that Remote Desktop is enabled and that the user is permitted to sign in. On systems that expose the classic System Properties controls, open the remote settings and verify that **Allow remote connections to this computer** is enabled. If the service is disabled at the destination, no client-side change will restore the session.

Where the environment supports it, use **Network Level Authentication (NLA)** so that authentication occurs before a full Remote Desktop session is created. Then review the permitted users or groups. Select **Select Users** and grant access through an appropriate group, such as **Remote Desktop Users**, rather than adding accounts individually. In a domain environment, also confirm that the required users are members of that group in Active Directory.

A dedicated Active Directory group can make this easier to administer. For example, create a group named RDP, add the users who require remote access, and then grant that group permission on the Remote Desktop host. Group-based access gives you one place to add or remove users as responsibilities change.

After the group is populated, add it to the Remote Desktop permissions on the destination system. The user must belong to the approved group, and the destination must permit that group to use Remote Desktop.

Next, verify the host firewall. RDP normally uses TCP port 3389, so the appropriate Remote Desktop inbound rule must be enabled. Also check the rule scope; opening the port does not mean it should be reachable from every source.

Open **Windows Defender Firewall with Advanced Security**, select **Inbound Rules**, and inspect the Remote Desktop rules. Confirm that the protocol and port match the service configuration.

The following figure shows the protocol and local-port settings for a Remote Desktop user-mode inbound rule:




Figure 7.1: Remote Desktop inbound rule configured for TCP port 3389

In this example, the rule is configured for TCP and local port 3389, which is the standard RDP listening port.

In **Inbound Rules**, locate the Remote Desktop entries and verify that the required rule is enabled. An enabled rule normally shows the corresponding status indicator in the firewall console. If the organization uses a nonstandard listening port, the firewall rule and the service configuration must agree; testing one value while the host listens on another will look like a network failure.

The same basic checks apply to Windows Server and supported Windows desktop editions that can accept Remote Desktop connections. The session model differs between a server and a desktop operating system, but the troubleshooting sequence is the same: confirm the service is enabled, verify the user's authorization, and then validate the firewall path.

If Windows Server also provides VPN access, open **Routing and Remote Access** and confirm that the service is running. A stopped VPN service can block users before Remote Desktop is ever reached.

If the VPN uses multi-factor authentication, verify that the additional authentication method is available and correctly integrated. Valid primary credentials alone will not complete the connection if the second factor fails.

For L2TP or IKEv2 deployments, verify the authentication material required by the configuration, such as the pre-shared key or certificate. A client/server mismatch can stop the VPN negotiation.

For SSTP, verify the server certificate and confirm that it is valid and has not expired. The network path can be healthy while certificate validation still prevents the VPN from connecting.

Third-party products can add another authentication dependency. Cisco Duo, for example, can provide multi-factor authentication for Windows remote access. Include that component in the troubleshooting path rather than checking only Windows settings.

Windows Server supports several VPN technologies, including L2TP, SSTP, and IKEv2. Older PPTP deployments may still exist, but PPTP should not be used for new secure deployments. Whatever VPN type is selected, the client settings, server configuration, certificate or shared-secret requirements, and firewall rules must match. Use the vendor or Microsoft documentation for the detailed protocol-specific setup rather than changing several settings at once during an outage.

The Routing and Remote Access properties shown in the following figure provide a checkpoint for two of those dependencies: authentication configuration and the certificate selected for SSTP.




Figure 7.2: Configuring VPN authentication and SSTP certificate binding in Routing and Remote Access

Once the host configuration is correct, follow the connection through the firewalls. A perimeter device such as a Cisco ASA must allow the required remote-access traffic and direct it toward the correct internal system or VPN endpoint.

The following Cisco ASA excerpt illustrates NAT and routing configuration you may encounter while tracing that path. Use it to identify the translated addresses, interfaces, and routes involved in the connection:

```
mtu gig 1500
icmp unreachable rate-limit 1 burst-size 1
no asdm history enable
arp timeout 14400
no arp permit-nonconnected
!
object network obj_any
  nat (inside,outside) dynamic interface
object network obj-192.168.21.23
  nat (inside,outside) static 199.107.166.116
```

```

object network hyperv1
  nat (inside,outside) static 199.107.166.115
object network WillAlli
  nat (inside,outside) static 199.107.166.118
object network obj-192.168.21.112
  nat (inside,outside) static 199.107.166.117
access-group nat_in in interface outside
route outside 0.0.0.0 0.0.0.0 199.107.166.113 1
route outside 0.0.0.0 0.0.0.0 2.2.2.2 2
route outside 0.0.0.0 0.0.0.0 3.3.3.3 3
timeout xlate 3:00:00
timeout pat-xlate 0:00:30
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02

```

The corresponding access-list entries show which inbound TCP and UDP services are permitted to particular internal hosts. For RDP, several rules restrict TCP 3389 to named source addresses, while the VPN-related entries allow the ports required by the configured VPN service. Read these rules as a path: source address, destination address, protocol, and destination port must all match the connection you expect. A VPN does not automatically make every internal service reachable; the firewall and internal access policy still determine what the connected user can access. The following access-list entries show how those decisions are expressed in the firewall configuration, with individual rules identifying the permitted source, destination, protocol, and service:

```

access-list nat_in extended permit tcp host 73.190.100.226 host 192.168.21.8 eq 3389
access-list nat_in extended permit tcp host 73.190.100.226 host 192.168.21.23 eq 3389
access-list nat_in extended permit tcp host 47.7.3.119 host 192.168.21.23 eq 3389
access-list nat_in extended permit tcp host 47.7.3.119 host 192.168.21.8 eq 3389
access-list nat_in extended permit tcp any host 192.168.21.19 eq 3389
access-list nat_in extended permit tcp any host 192.168.21.19 eq https
access-list nat_in extended permit tcp any host 192.168.21.19 eq www
access-list nat_in extended permit tcp any host 192.168.21.19 eq ssh
access-list nat_in extended permit tcp host 24.20.141.11 host 192.168.21.23 eq 3389
access-list nat_in extended permit tcp host 24.20.141.11 host 192.168.21.23 eq ftp
access-list nat_in extended permit tcp host 24.20.141.11 host 192.168.21.23 eq ftp-data
access-list nat_in extended permit tcp any host 192.168.21.114 eq https

```

```
access-list nat_in extended permit tcp host 47.7.26.78 host 192.168.21.8 eq 3389
access-list nat_in extended permit tcp host 47.7.26.78 host 192.168.21.23 eq 3389

access-list nat_in extended permit udp any host 192.168.21.112 eq 1701
access-list nat_in extended permit udp any host 192.168.21.112 eq isakmp
access-list nat_in extended permit udp any host 192.168.21.112 eq 4500
access-list nat_in extended permit tcp any host 192.168.21.112 eq https
access-list nat_in extended permit tcp any host 192.168.21.112 eq www
access-list nat_in extended permit tcp any host 192.168.21.112 eq smtp
pager lines 24
mtu outside 1500
mtu inside 1500
mtu gig 1500
icmp unreachable rate-limit 1 burst-size 1
```

Taken together, the configuration shows several separate dependencies in the remote-access path. The NAT statements map internal systems to translated addresses, the routes determine where traffic leaves the firewall, and the access-list entries restrict particular protocols and destination ports to the intended hosts. The key troubleshooting lesson is that a successful VPN or RDP connection depends on all of these decisions agreeing on where the traffic should go; a correct endpoint configuration cannot compensate for an incorrect translation, route, or firewall rule.

For an on-premises network, traffic from a home user first crosses the user's local router or firewall, then the internet, and then the organization's perimeter firewall. The corporate firewall must forward or terminate the traffic at the correct destination. Direct RDP uses its configured listening port, while VPN protocols use their own required ports and negotiation traffic. If the perimeter rule, NAT translation, or VPN listener points to the wrong host, the connection fails before Remote Desktop permissions become relevant.

After a VPN or Remote Desktop session is established, continue to apply least-privilege rules to internal systems. Permit only the servers, ports, and source networks required for the user's task. This limits the effect of a compromised remote account and also makes troubleshooting clearer because the intended path is explicitly defined.

Troubleshooting printing and local network discovery in remote environments

Remote printing introduces several dependencies: the local printer must work, the remote session must redirect it as expected, and the remote system must be able to use it. First,

determine whether the problem exists locally or only inside the remote session. If a local test page fails, troubleshoot the printer or driver before RDP redirection.

On the client computer, confirm that the intended printer is installed and selected as the default when that is appropriate for the user's workflow. If printing began failing after an operating-system update, check the printer vendor for a current driver. The same principle applies to macOS and Linux clients: verify the platform-specific driver or supported print path before blaming the remote host.

If local printing works but the printer is missing inside a Remote Desktop session, inspect the session's printer-redirection settings and the remote host. Depending on the deployment, the session may rely on redirected printing or a matching vendor driver. Compare a working user's session with the failing one and confirm that the printer is actually being presented to the remote system before replacing drivers unnecessarily.

Device visibility on the home LAN is a separate but related support issue. Open **File Explorer** and select **Network**. If other computers and devices are missing, Windows may report that network discovery is turned off.




Figure 7.3: Windows warning indicating that network discovery is turned off

Enable network discovery and file sharing only on networks where that visibility is intended, such as a trusted private network.




Figure 7.4: Enabling network discovery in Windows sharing settings

If network discovery is enabled but devices still do not appear, open `services.msc` and inspect the discovery-related services. Verify that **Function Discovery Provider Host** and **Function Discovery Resource Publication** are running as required before moving on to name-resolution or firewall checks.




Figure 7.5: Function Discovery services running on a Windows system

Up to this point, the failures have been tied to identifiable services or Windows features: remote access, printing, and network discovery. Confirming those dependencies is important because it prevents a service-specific problem from being mistaken for a wider network outage. When the connection itself is slow, intermittent, or unable to reach the internet, the investigation must now move beneath those services and examine the underlying wired or wireless path.

Diagnosing slowness, internet access, and dropped connections

When a user reports slowness or no internet access, start with the endpoint's IP configuration. Confirm that the interface has a usable IP address, subnet information, a default gateway for off-subnet traffic, and a DNS server for name resolution. Test each dependency separately: reachability to the gateway tells you something different from successful DNS resolution, and both are different from application performance.

Addressing dropped connections in wired and wireless networks

Dropped connections require you to separate the physical or radio link from higher-layer services. On Wi-Fi, compare signal and coverage, move the client, or test another access point or channel. On Ethernet, inspect the cable and link state; an intermittent cable can look like software instability. A cable tester or known-good cable is often the fastest check.

Because laptops and other endpoints rely heavily on Wi-Fi both at home and in offices, wireless troubleshooting deserves its own sequence. First determine whether the client can associate with the access point, then whether it receives an address, then whether it can reach the local gateway, resolve names, and reach external destinations. This keeps radio, DHCP, DNS, and internet problems from being mixed together. The following troubleshooting path

keeps those dependencies in sequence, allowing you to identify the first stage at which a wireless client's connection stops working.




Figure 7.6: Troubleshooting sequence for isolating wireless connectivity failures

A wireless client that connects to the SSID but does not receive a usable IP address may be failing to reach DHCP, may have an incorrect static configuration, or may not be communicating reliably with the access point. Signal quality also matters. Walls, floors, distance, and other obstructions can weaken or distort the radio path, so compare the client's location and signal quality with a known-working device before changing addressing settings.

If the client receives a valid address and can reach a resource by IP address but not by hostname, shift your attention to DNS. Test whether the configured resolver responds and whether it returns the expected address for the resource. Internal and external DNS services may answer different names, so identify which resolver the client is actually using instead of assuming that all name lookups follow the same path.

Wireless access point configuration and channel selection

At the access point, review both the radio band and the channel in use. Nearby access points can compete for airtime, especially when channels overlap, and the client must support the band and wireless mode that the access point is advertising. Interfaces may label legacy and modern modes with standards such as 802.11a, b, g, n, ac, or ax. Check the wireless adapter's capabilities before forcing the access point to a mode that some clients cannot use.

Hardware age can also contribute to compatibility problems. An older wireless adapter may not support the same bands or modes as a newer access point. Check supported combinations, firmware, and drivers before replacing otherwise functional equipment.

Use the access point's management interface to compare the configured band, channel, security mode, and client capabilities with the environment you are troubleshooting.

The following screenshot shows the security-mode choices exposed by one access point. Use it as a configuration checkpoint rather than assuming that every client failure is caused by signal strength.




Figure 7.7: Wireless security-mode options on an access point

Many access points expose separate settings for the 2.4 GHz and 5 GHz radios, and newer equipment may add 6 GHz. The available wireless modes differ by hardware and band. During troubleshooting, verify that the band is enabled, that the client supports it, and that the selected mode does not exclude the device you are testing. If you temporarily disable a band to isolate a problem, restore the intended production configuration after the test rather than treating one band as universally preferable.

Security settings matter as much as radio settings. A disabled security mode may simplify association, but it leaves the network exposed. Confirm that the client and access point use compatible authentication and encryption rather than leaving security disabled as a workaround.

In environments using WPA2, use an AES-based configuration rather than legacy TKIP compatibility modes. A pre-shared-key deployment authenticates clients with a shared passphrase, while an enterprise deployment can authenticate individual users or devices. The important troubleshooting check is consistency: the client must be configured for the same security mode and credentials that the access point expects.

With RADIUS-backed enterprise Wi-Fi, verify the authentication service as well as the access point. Even with a healthy radio link, a user may fail to join because the RADIUS service, account, or access policy is failing.

Successful authentication answers only whether a device is allowed to join the wireless network. The next question is what that device should be able to reach once connected, which makes network separation and access policy especially important for guest users.

Guest network configuration and isolation

A guest network should keep untrusted or visitor devices away from production resources while still providing the intended internet access. The separation can be implemented with a dedicated subnet, firewall policy, or VLAN carried through the access point and switching infrastructure. Whichever method is used, test from a guest device and confirm both sides of the policy: internet access should work, while protected internal resources should remain unreachable.

Channel selection is another useful variable when the wireless link is unstable. Many access points default to automatic selection, but the management interface may also allow you to choose a channel manually. If several nearby networks occupy the same part of the spectrum, testing a different permitted channel can show whether interference is contributing to the problem.

The number of usable channels depends on the frequency band and channel width. *Table 7.1* summarizes the channel counts used in this example across the 2.4 GHz, 5 GHz, and 6 GHz bands.

Channel width	2.4 GHz	5 GHz	6 GHz
20 MHz	3 non-overlapping	~25	59
40 MHz	Generally 1 practical	~12	29
80 MHz	Not supported	~6	14
160 MHz	Not supported	~2	7
320 MHz	Not supported	Not supported	3
Common Wi-Fi generations	Wi-Fi 4–7	Wi-Fi 4–7	Wi-Fi 6E & 7

Table 7.1: Approximate Wi-Fi channel availability by band and channel width

Wider channels consume more spectrum, leaving fewer independent channel choices. Use the table as a planning aid alongside the options actually offered by the access point in your region.

For 2.4 GHz networks using 20 MHz channels, 1, 6, and 11 are commonly used as non-overlapping choices where available. A significant change after switching channels points toward interference rather than addressing.

Using built-in wireless access point diagnostic tools

Many access points also include diagnostic tools that run from the access point itself. Typical options include Ping, traceroute, NSLookup, Speedtest, and Device Discovery. These tools are valuable because they test connectivity from the infrastructure device rather than from the user's computer, giving you a comparison point when the client and access point behave differently.

Enter a local or external IP address in the access point's **Ping** tool and send a few probes. If the access point reaches the target but the client cannot, focus on the client-to-access-point path or client-specific policy.

A typical ping interface asks for the target, packet size, and number of probes, as in the following example:

```
Target IP / Domain Name: 192.168.1.1
Ping Packet Size: 64 Bytes
Number of Pings: 4
[Start]
```

A successful result confirms that the access point can reach the selected destination. You can then move farther along the path by running traceroute to an external address such as Google's public DNS service at 8.8.8.8.

The traceroute view records the sequence of routers that respond on the way to the destination, allowing you to compare the access point's upstream path with the path observed from a client:

```
Target IP / Domain Name: 8.8.8.8
[Start]
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 38 byte packets
 1  192.168.1.1  0.421 ms
 2  22.54.156.1  16.196 ms
 3  96.34.105.202  10.657 ms
```

```
...  
13 8.8.8.8 27.107 ms
```

The completed trace shows that the access point itself can reach the external destination across the upstream network. That does not prove that every wireless client has the same connectivity, but it gives you an important comparison point: if the access point reaches the destination while the client does not, the failure is more likely to sit between the client and the access point, in client-specific addressing or policy, or in the wireless association itself.

These infrastructure-side results still do not prove that a particular client can associate successfully. If the computer still cannot connect, return to the SSID, wireless security settings, signal quality, DHCP process, and client-adaptor compatibility. If the SSID is intentionally hidden, make sure the user has the exact network name and security settings.

A wired test is useful when Wi-Fi is suspect because it removes the radio link from the path. Use the comparison to decide whether the failure belongs to wireless connectivity or to the network beyond the access point.

The wireless checks have now separated client association from the network beyond the access point, and the wired comparison gives you a way to remove the radio link from the equation altogether. If connectivity still fails once the local wireless path has been ruled out, the troubleshooting boundary moves outward to the host's IP configuration, the home router, and ultimately the ISP.

Troubleshooting internet connectivity from the home network

When internet access fails at home, locate the boundary of the outage before changing configuration. The problem may be limited to one computer, affect every device behind the home router, or originate beyond the router with the ISP. Windows network settings and command-line tools let you work outward from the local adapter to the gateway and then to remote destinations.

In Windows 11, review the Ethernet connection in **Settings** before changing values. If you need to verify or modify the IP address, subnet mask, gateway, or DNS server, use the following sequence:

1. Open **Start** and select **Settings**.
2. Select **Network & internet**.
3. Select **Ethernet** for the wired connection you are troubleshooting.
4. Under **IP assignment**, select **Edit**.

5. Choose **Automatic (DHCP)** to obtain settings from the network or **Manual** when the environment requires a static configuration.
6. If you select **Manual**, enable the IP version you need to configure, such as IPv4.
7. Enter the assigned IP address, subnet mask or prefix, default gateway, and DNS information exactly as required by the network.
8. Select **Save**.
9. Open Command Prompt and run `ipconfig /all` to verify the resulting configuration.
10. For adapter-specific options such as **Speed & Duplex**, VLAN settings, power management, or driver properties, open **Device Manager**, expand **Network adapters**, open the Ethernet adapter's **Properties**, and review the **Advanced** tab.

A Windows computer may have Wi-Fi and Ethernet enabled at the same time. If you are specifically testing the wired path, temporarily disable Wi-Fi so that you know which adapter is carrying the traffic. This is a troubleshooting control rather than a permanent performance fix: once the test is complete, restore the normal adapter configuration.

For a wired adapter, open its properties and review **Internet Protocol Version 4 (TCP/IPv4)**. If **Use the following IP address** is selected, the interface is statically configured; verify the address, subnet mask, gateway, and DNS values together. With a DHCP configuration, the same fields should be supplied automatically by the network.

The adapter's **Status** and **Details** views provide another way to inspect the current configuration. Use them to confirm the assigned IPv4 address, subnet mask, default gateway, and DNS servers. Compare these values with a working device on the same network rather than judging each field in isolation.

If the interface has no usable address, gateway, or expected DHCP information, investigate the path to the DHCP server. In many home networks, the router provides DHCP; in a corporate environment, the service may run on a Windows server, router, switch, or another network appliance. The location of the DHCP service determines where you test next.

Router and ISP troubleshooting steps

In the adapter status or settings, verify that the link is connected and check the negotiated link speed. Activity counters should change as the system sends and receives traffic, confirming that the interface is passing frames even if internet access is still unavailable. A healthy local link narrows the problem, but it does not by itself prove that the ISP path is working.

The following figure shows an Ethernet connection with manual IPv4 and DNS assignments. These values provide the first checkpoint before you move from the local adapter to command-line tests.

IP assignment:	Manual	
IPv4 address:	192.168.21.50	Edit
IPv4 mask:	255.255.255.0	
IPv4 gateway:	192.168.21.1	
DNS server assignment:	Manual	Edit
IPv4 DNS servers:	8.8.8.8 (Unencrypted)	
Link speed (Receive/Transmit):	1000/1000 (Mbps)	Copy
Link-local IPv6 address:	fe80::8c7:10fb:5ece:414f%13	
IPv4 address:	192.168.21.50	
IPv4 DNS servers:	8.8.8.8 (Unencrypted)	
Manufacturer:	Microsoft	
Description:	Microsoft Hyper-V Network Adapter	
Driver version:	10.0.22000.3079	
Physical address (MAC):	00-15-5D-0F-70-31	

Figure 7.8: Ethernet properties showing manual IPv4, gateway, DNS, and link-speed information

In this example, the adapter has an IPv4 address, subnet mask, gateway, DNS server, and link speed, so the next test is reachability beyond the host.

Open Command Prompt and collect the complete adapter configuration with the following command:

```
ipconfig /all
```

`ipconfig /all` lists details for each adapter, including MAC address, IPv4 and IPv6 settings, DNS servers, gateways, and DHCP state. Make sure you are reading the adapter that is actually carrying the connection you are testing.

The graphical and command-line views should tell the same story. If an unexpected VPN, wireless, virtual, or disconnected adapter appears in the output, do not assume it is the active path. Identify the adapter with the relevant address and gateway before using its values in later tests.

After validating the local configuration, use `pathping` to examine the route toward a remote destination and gather loss information along that path. The command first discovers the sequence of hops and then sends additional probes, which helps distinguish an isolated nonresponsive router from a pattern of loss that continues toward the destination.


For example, run `pathping` against a known destination:

```
pathping google.com
```

`pathping` combines route discovery with per-hop packet-loss statistics, so it takes longer to complete than a simple `tracert`. Read the output in context: a router may decline to answer diagnostic probes even while it continues forwarding traffic, so one missing response is not automatically evidence that the router is down.

If the trace consistently stops before the destination and later hops never respond, compare the result with another destination before escalating. A repeatable failure at the same part of the upstream path is stronger evidence than a single timeout, especially when local gateway and DNS tests are already successful.

The following `pathping` example includes a hop that does not answer the diagnostic probe. The important question is whether subsequent hops and the destination continue to respond.



```
C:\>pathping google.com

Tracing route to google.com [172.217.3.206]
over a maximum of 30 hops:
  0  DC1.TechPub.us [192.168.21.110]
  1  199.107.166.113
  2  206.72.111.74
  3  xerxes.sterling.net [206.72.100.2]
  4  206.72.110.1
  5  206.72.100.18
  6  206.72.100.8
  7  100.xe-4-0-1.mpr2.pdx1.us.zip.zayo.com [208.185.182.57]
  8  ae9.cs1.sea1.us.zip.zayo.com [64.125.25.180]
  9  ae27.mpr1.sea1.us.zip.zayo.com [64.125.29.1]
 10  72.14.208.172
 11  108.170.245.113
 12  108.170.233.155
 13  sea15s12-in-f206.1e100.net [172.217.3.206]
```

Figure 7.9: Tracing the route to google.com with pathping


Because the route continues after the nonresponsive hop, the missing reply at that line does not by itself show that traffic was rerouted or that the router failed. The device may simply be filtering or deprioritizing ICMP diagnostic traffic while still forwarding packets. If later hops also disappear and the destination cannot be reached, the evidence becomes more consistent with a path failure that the ISP or upstream provider may need to investigate.

Run the same test to more than one destination before drawing a conclusion. If pathping google.com fails at one point but another unrelated destination succeeds through a different route, the issue may be path-specific. If several destinations fail after the same upstream hop, save the results and provide them to the ISP; the comparison gives the provider more useful evidence than a single screenshot.

Return to the default gateway as a local checkpoint. In the earlier example, it is 192.168.21.1; pinging it tests whether the host can reach its first router. Also inspect the router or modem indicators. An abnormal link, internet, or error light may point to a physical or upstream fault.

Resetting the ARP cache and router hardware checks

ARP provides one more local check when the host can reach the link, but address-to-hardware mappings may be stale. The command `arp -a` displays the IPv4 addresses currently known to the host together with their corresponding physical (MAC) addresses. The following output shows several dynamic entries learned on the local network:



```
Interface: 192.168.1.119 --- 0xd
Internet Address      Physical Address      Type
192.168.1.1          70-db-98-3e-a7-e9    dynamic
192.168.1.11         20-0b-74-a7-a7-25    dynamic
192.168.1.103        f8-54-b8-af-e9-2c    dynamic
192.168.1.106        7c-d5-66-61-7f-0e    dynamic
192.168.1.109        3c-8d-20-03-35-5b    dynamic
192.168.1.111        20-fe-00-a0-9e-eb    dynamic
192.168.1.115        b8-5f-98-36-ff-3b    dynamic
```

Internet Address	Physical Address	Type
192.168.1.1	70-db-98-3e-a7-e9	dynamic
192.168.1.11	20-0b-74-a7-a7-25	dynamic
192.168.1.103	f8-54-b8-af-e9-2c	dynamic
192.168.1.106	7c-d5-66-61-7f-0e	dynamic
192.168.1.109	3c-8d-20-03-35-5b	dynamic
192.168.1.111	20-fe-00-a0-9e-eb	dynamic
192.168.1.115	b8-5f-98-36-ff-3b	dynamic

Figure 7.10: ARP cache entries showing IPv4-to-MAC address mappings

After a router, firewall, or endpoint is replaced, a stale ARP entry can temporarily associate an IP address with an old MAC address. If suspected, refresh the relevant ARP state or restart the affected device, then test again. Use this only after IP configuration and gateway tests suggest a local Layer 2 mapping problem.

Some outages ultimately belong to the ISP rather than the home or corporate endpoint. Provider equipment may be restarted, replaced, or reconfigured during an incident, and an upstream failure can affect every local device even when their settings are correct. By tracing the path from the computer to the gateway and then outward with `ping`, `ipconfig`, and `pathping`, you can show where your own network stops working and provide the ISP with evidence that is specific enough to act on.

Summary

Troubleshooting a remote or home user requires you to work with less visibility and less control than you usually have inside the corporate network. The most useful response is therefore not to begin changing settings, but to identify what still works and find the first boundary where the connection stops behaving as expected. A working internet connection but failed VPN, a successful VPN but failed Remote Desktop session, or a working wired connection but unstable Wi-Fi each places the fault in a different part of the path.

We applied that method first to remote access, where Remote Desktop permissions, VPN authentication, certificates, firewall rules, NAT, and internal access policies must all agree before the user can reach the intended system. The same principle carried into printing and network discovery: determine whether the underlying feature works locally before troubleshooting how it behaves across a remote session or trusted network. For wireless connections, the diagnostic sequence becomes especially useful: verify association with the SSID, then addressing, gateway reachability, DNS, and finally access to the required external or corporate resource.

When the problem extends beyond the local device or access point, tools such as `ipconfig /all`, `pathping`, gateway tests, and the ARP cache let you work progressively outward toward the router and ISP. The important lesson is the same one established throughout this part of the book: every successful test removes part of the network from suspicion. By comparing a failing path with a known-working one, you can determine whether the next action belongs on the endpoint, the home network, the remote-access infrastructure, or with the upstream provider.

A home wireless problem often involves one router or access point and a relatively small number of clients, making comparison tests an effective way to isolate the failure. Enterprise wireless networks apply the same principles at a much larger scale, where multiple access points, overlapping coverage, shared radio spectrum, security policies, and many simultaneous clients can produce symptoms that are harder to localize. In the next chapter, we will take the same evidence-driven troubleshooting approach into enterprise Wi-Fi and examine how interference, performance, capacity, and wireless infrastructure change the investigation.

8

Troubleshooting Enterprise Wireless Networks

Enterprise wireless networks extend network access across shared radio spectrum rather than dedicated cabling. That changes the nature of troubleshooting: a client can remain connected while interference, poor coverage, channel contention, security configuration, or excessive demand steadily degrades the experience. At enterprise scale, the first challenge is therefore to determine whether a complaint belongs to one client or reflects a wider problem with the wireless environment.

In this chapter, we begin with the security concerns that shape enterprise Wi-Fi before examining the common causes of slow or unstable wireless connections. We then look at how coverage, heat maps, centralized monitoring, RF-analysis tools, QoS, WMM, channel planning, and client capabilities help explain what users are experiencing. From there, we examine environmental sources of interference and finish with high-density environments, where a strong signal does not necessarily translate into good performance because many clients and access points may be competing for the same airtime.

By the end of the chapter, you should be able to take a broad complaint such as "the Wi-Fi is slow" and narrow it into a specific wireless fault domain. The goal is to determine whether the evidence points to the client, coverage, interference, an access point, shared capacity, security policy, or the wider RF environment before changing configuration.

In this chapter, we will cover the following topics:

- Wireless network security concerns and threats
- Common causes of unstable or slow wireless connections
- Managing user experience and bandwidth utilization

- Sources of Wi-Fi interference in enterprise environments
- Challenges of high-density wireless environments

Technical requirements

You should already understand basic networking concepts and the role of wireless access points, SSIDs, and the 2.4 GHz, 5 GHz, and 6 GHz Wi-Fi bands. You should also understand the difference between authentication, which establishes identity, and encryption, which protects data in transit. Familiarity with Windows network settings and the process of connecting a client to a wireless network will also help when separating endpoint problems from wider wireless infrastructure issues.

Wireless network security concerns and threats

Wireless networks extend connectivity through radio signals, so they should not be treated as risk-free. Administrators need strong authentication and encryption, controlled access, and awareness of threats such as spoofed gateways or access points, man-in-the-middle attacks, and wireless sniffing. Modern Wi-Fi security relies on strong encryption such as AES, but encryption is only one part of the design. In locations where signal leakage outside the building is a concern, organizations may also use RF attenuation or shielding measures to reduce how far usable signals travel.

Common causes of unstable or slow wireless connections

Slow or unstable wireless connectivity can be difficult to diagnose because a client may remain connected while performance degrades. Begin by asking whether the symptom affects one device, one area, one access point, or many users. That scope tells you whether to investigate the client first or move outward toward coverage, interference, and access-point capacity. The following figure maps different patterns of impact to the fault domains that should receive attention first.




Figure 8.1: Scoping an enterprise wireless performance problem by affected clients and areas

The figure converts a broad wireless complaint into an initial fault domain. A problem isolated to one device points toward the client, whereas failures clustered around a location or access point shift attention toward coverage, interference, AP health, or capacity. Similar symptoms across several access points justify moving farther outward to shared wireless infrastructure and the wider RF or upstream network.

Interference is a frequent cause of unstable wireless performance. Nearby radio-frequency activity can compete with Wi-Fi, while obstacles such as thick walls, concrete floors, and multiple stories can weaken or distort the signal. The result may appear as dead spots, retransmissions, higher latency, or dropped connections rather than a complete outage.

Heat maps and centralized wireless management can turn those symptoms into evidence. A heat map can reveal weak coverage, excessive overlap, and areas where access points are competing for the same radio space. Co-channel interference occurs when nearby Wi-Fi networks compete on the same channel, while adjacent-channel interference occurs when transmissions overlap in frequency; both can increase contention and retransmissions.

For example, two nearby 2.4 GHz access points using overlapping channels can interfere with each other's transmissions. Reassigning channels, adjusting access-point placement, and moving capable clients to 5 GHz or 6 GHz where appropriate can reduce this contention. Verify the RF environment before assuming that poor performance is caused by the Internet connection or the user's device.

Enterprise wireless platforms from vendors such as Cisco and Aruba commonly provide centralized monitoring and RF-management capabilities across many access points.

Centralized dashboards are useful, but sometimes you need to observe the radio environment directly. Handheld or portable tools can help identify interference and suspicious wireless activity. Common examples include:

- **RF spectrum analyzers:** Detect radio-frequency activity and interference, including transmissions that are not necessarily Wi-Fi.
- **Wireless security analyzers or WIDS tools:** Identify rogue access points, unexpected clients, unusual associations, and certain wireless attacks.
- **Laptop or handheld packet analyzers:** Use a Wi-Fi adapter capable of monitor mode with tools such as Wireshark to examine 802.11 traffic.

At this point, you can use scope, centralized monitoring, and RF evidence to distinguish a client-specific complaint from a broader wireless problem. That distinction matters because good coverage alone does not guarantee good performance: once clients are connected, they still have to share finite airtime and bandwidth. The investigation therefore moves from whether the wireless path exists to whether that shared capacity can deliver the experience users and applications require.

Managing user experience and bandwidth utilization

Enterprise wireless performance is measured by whether users can work reliably, not simply by whether the Wi-Fi icon shows a connection. Applications place different demands on the

network: voice and video are sensitive to latency and jitter, while large file transfers can consume substantial bandwidth. Troubleshooting therefore requires you to examine both user experience and shared wireless capacity.

Usage can vary sharply between clients. One user may be transferring large files or streaming media while another generates very little traffic. When many devices share an access point, heavy background traffic can affect everyone associated with it.

Quality of Service (QoS) and bandwidth policies can reduce this impact by preventing less time-sensitive traffic from overwhelming delay-sensitive applications. Additional access points can also increase capacity in high-demand areas when they are planned with suitable channels and placement. Simply adding access points without considering RF overlap can create more contention.

QoS classifies and prioritizes traffic so important or delay-sensitive applications receive preferential treatment when the wireless network is congested. Voice calls and video conferences generally need lower latency and jitter than software downloads or bulk file transfers.

Wi-Fi commonly uses **Wi-Fi Multimedia (WMM)** to place traffic into categories such as voice, video, best effort, and background. WMM does not create additional bandwidth, but it can improve how limited airtime is used during periods of contention.

Channel planning is part of this user-experience problem. Each access point should be checked to avoid unnecessary overlap with nearby access points. The 2.4 GHz band is particularly constrained, while 5 GHz and 6 GHz provide more room for channel reuse. Vendor planning and management tools can help identify where access points are competing for the same spectrum.

Client capability also matters. Enterprise environments often contain a mixture of Wi-Fi generations, including 802.11a/b/g/n/ac/ax/be devices. Access points must support required clients, but retaining older compatibility modes can reduce efficiency in some deployments. Before disabling a legacy mode, confirm that no required client depends on it.

Sources of Wi-Fi interference in enterprise environments

After considering client load and channel use, ask what else is competing with or obstructing the signal. Enterprise buildings contain many devices, materials, and neighboring networks that can affect radio performance. Identify the source, determine which band is affected, and then change placement, channel use, or client distribution based on the evidence.

Interference from appliances and electronic devices

Appliances and electronic equipment can contribute to interference, particularly in the 2.4 GHz band. Microwaves, some lighting equipment, power systems, and other electronics may generate radio-frequency noise that affects nearby clients. Where coverage and client support allow it, moving capable devices to 5 GHz can reduce exposure to some 2.4 GHz interference.

Interference also comes from other wireless devices using the same spectrum. Building materials can absorb or reflect radio signals, and reflected signals can make the connection less predictable. These effects may appear as poor throughput or roaming problems even when the client reports that it is still connected.

Changing bands should therefore be treated as one troubleshooting option rather than a universal fix. Moving suitable clients from 2.4 GHz to 5 GHz may reduce contention, but the result still depends on coverage, device capability, and the physical environment.

Wireless headsets, keyboards, mice, cordless phones, baby monitors, security cameras, game controllers, and other radio devices can also contribute to a crowded RF environment. If intermittent problems occur in a specific place or at a predictable time, look for nearby equipment that becomes active at the same time.

Challenges of high-density wireless environments

High-density environments combine many clients competing for airtime with many nearby access points or networks sharing limited spectrum. This can occur in offices, campuses, conference spaces, apartment buildings, and other locations where networks overlap. A client may show strong signal strength but still have poor throughput because the problem is contention rather than coverage.

Some organizations use RF attenuation or shielding techniques to reduce unwanted signal leakage into or out of a controlled space. These measures are site-specific and should be considered alongside building design, access-point placement, and the coverage legitimate users still require.

Concrete, metal, and stone can weaken wireless signals, while reflective surfaces such as mirrors and foil-backed insulation can create unpredictable reflections. Adding access points may improve coverage, but placement and channel planning still need to be checked so the new access point does not add more interference.

Line of sight and device placement

Wi-Fi performs best when the radio path between the client and access point is relatively clear. In real buildings, signals travel through or around walls, floors, furniture, people, and other

obstacles. When troubleshooting a weak or unstable client, compare its location with the access point, nearby obstructions, and the results of a heat map or signal measurement.

Interference and density are additional fault domains within the same systematic troubleshooting process used earlier in the book. Start with the scope of the complaint, verify the client and access point, examine signal quality and channel use, and then look for environmental or capacity-related causes before changing the configuration.

Summary

Enterprise wireless troubleshooting starts by establishing scale. A single client with poor performance suggests a different fault domain from an entire floor of affected users, and a strong Wi-Fi signal does not by itself prove that the network has enough usable airtime or capacity. By identifying who and where the problem affects first, you can decide whether to investigate the client, an access point, coverage, interference, shared infrastructure, or the wider RF environment.

We used that principle to examine both access and performance. Authentication, encryption, and wireless security controls determine whether a client should be able to join safely, while heat maps, centralized management, spectrum analysis, channel planning, QoS, and WMM provide evidence about what happens once clients are connected. Environmental interference and high-density deployments add another layer: walls, electronic devices, overlapping networks, client load, and poorly planned access points can all degrade service even when the WLAN remains technically available.

The practical lesson is to avoid treating every complaint of "slow Wi-Fi" as the same problem. Establish the scope, compare affected and known-working clients or areas, examine the RF and capacity evidence, and make a change only when that evidence identifies the fault domain. That keeps wireless troubleshooting consistent with the systematic process used throughout the book.

Enterprise wireless troubleshooting establishes whether users can reliably reach the network and the services beyond it. A healthy Wi-Fi or internet connection, however, does not guarantee that a remote user can reach the corporate environment when a VPN introduces its own protocol, authentication, certificate, addressing, and routing dependencies. In the next chapter, we will apply the same dependency-by-dependency method to VPNs and remote access, distinguishing failures that prevent a tunnel from forming from problems that appear after the tunnel is connected.

9

Troubleshooting VPNs and Remote Access

A **virtual private network (VPN)** extends access to a private network across another network, usually the internet, by creating a protected tunnel between the client and a VPN endpoint. For remote users, however, a working Wi-Fi or internet connection proves only that the first part of the path is available. The VPN adds its own dependencies, including the tunnel protocol, server reachability, authentication, certificates, client addressing, DNS, routing, and access to the intended corporate resource.

In this chapter, we begin by comparing the VPN protocols you may encounter and the network requirements that distinguish them. We then examine how public networks can affect VPN connectivity before moving to the Windows Server side, where RRAS, authentication, certificates, address allocation, and VPN ports determine whether a client can establish a usable connection. From there, we investigate WAN miniport problems and the Windows 11 client profile, then follow a connected VPN into DNS and internal resource access. Finally, we use common VPN error codes as evidence to identify which stage of the connection has failed.

By the end of the chapter, you should be able to distinguish between a VPN that cannot establish a tunnel and a tunnel that connects but cannot carry useful traffic. The goal is to identify whether the failure belongs to the access network, protocol negotiation, authentication, certificates, client configuration, addressing, DNS, routing, or resource access before changing unrelated settings.

In this chapter, we will cover the following topics:

- Common VPN protocols and their characteristics
- VPN protocol suitability in public networks
- VPN connectivity considerations in Windows Server

- Configuring Windows Server for VPN and remote access
- Investigating VPN ports and WAN miniport problems
- Client-side VPN configuration in Windows 11
- DNS configuration and name resolution for VPN connectivity
- Troubleshooting common VPN error codes

Technical requirements

You should understand TCP/IP addressing, DNS, default gateways, and the troubleshooting method established in *Chapter 6*. To follow the Windows examples, use a Windows 11 client and a Windows Server lab, such as Windows Server 2025, with the Remote Access role and **Routing and Remote Access Service (RRAS)** configured. Administrative access is required for server, certificate, and adapter changes. Use authorized test accounts and a certificate appropriate to your VPN endpoint. Record existing settings before making changes, and arrange local or console access before restarting services that support your remote session.

Common VPN protocols and their characteristics

Protocol identification is the first technical check because it determines which ports, authentication settings, and certificates to inspect. IPsec is a suite that protects IP traffic; **Internet Key Exchange version 2 (IKEv2)** negotiates the security associations used by an IKEv2/IPsec VPN. **Layer 2 Tunneling Protocol (L2TP)** provides tunneling and is commonly combined with IPsec for encryption. These are related components, not interchangeable names.

Secure Socket Tunneling Protocol (SSTP) carries PPP traffic through an HTTPS connection protected by TLS. Although vendors often use the broader term *SSL VPN*, not every TLS-based VPN implements SSTP. OpenVPN, for example, uses its own TLS-based VPN implementation rather than IPsec. Confirm the actual protocol supported by the client and gateway instead of assuming compatibility from a product category.

Table 9.1 summarizes the four Windows VPN types discussed here. Performance depends on configuration, hardware, traffic, and the underlying connection; a protocol name alone does not establish a speed ranking.

VPN type	Transport and firewall considerations	Security and troubleshooting considerations
PPTP	TCP 1723 and GRE, IP protocol 47	Legacy tunneling with known security weaknesses; do not select it for a new secure deployment.
L2TP/ IPsec	UDP 500 and 4500 for IPsec negotiation/NAT traversal; L2TP uses UDP 1701 within the protected exchange	Check IPsec authentication, the pre-shared key or certificate, and user authentication separately.
SSTP	TCP 443	Check endpoint reachability, TLS trust, certificate identity, and RRAS certificate binding.
IKEv2/ IPsec	UDP 500 and 4500; native ESP may also be relevant without UDP encapsulation	Check IKE credentials and IPsec settings. Mobility support can help when the client's network changes.

Table 9.1: VPN protocols and their principal troubleshooting dependencies

IKEv2 mobility and reconnection behavior can help clients move between networks, but it does not guarantee that every interruption preserves a session. L2TP/IPsec likewise has separate machine/tunnel authentication and user authentication stages: a pre-shared key is not the user's account password. Diagnose the stage that failed rather than resetting credentials indiscriminately.

You may also encounter DirectAccess in existing Windows deployments. Its IPv6-based architecture can involve transition mechanisms such as Teredo and IP-HTTPS, adding dependencies on certificates and network traversal. Microsoft has deprecated DirectAccess and recommends migration planning. Always On VPN provides automatically triggered connectivity through configured client profiles and supporting VPN infrastructure; it is not simply a server feature renamed or moved to the client.

Knowing which protocol the VPN uses tells you what the tunnel requires, but those requirements still have to survive the network between the client and the corporate endpoint. This becomes especially important on hotel, airport, and guest networks, where ordinary web access may work even though the ports, protocols, or negotiation traffic required by the VPN are filtered.

VPN protocol suitability in public networks

Hotels, airports, and guest networks can allow ordinary browsing while restricting VPN traffic. Before changing a corporate VPN configuration, verify that the client has completed any captive-portal sign-in and can reach the internet. Compare the same client and profile on another approved connection. A failure confined to one public network points toward that access path, although it does not by itself identify the blocking device.

SSTP often traverses restrictive networks because it uses TCP 443, a port commonly permitted for HTTPS. This improves compatibility but is not a guarantee: proxies, TLS inspection, filtering policies, and captive portals can still interrupt the exchange. Other vendors' SSL VPN products may provide a browser portal, a full client tunnel, or both. A successful portal login does not necessarily establish the tunnel needed for internal applications.

For IKEv2/IPsec or L2TP/IPsec, check UDP 500 and UDP 4500 along the path, including the corporate firewall. Native ESP uses IP protocol 50, which is different from a TCP or UDP port. L2TP/IPsec is not inherently exempt from public-network restrictions. Select an approved protocol supported by both endpoints, and use the relevant port requirements to investigate failures rather than switching to an insecure protocol just to obtain a connection.

Once the access network has been checked, inspect the VPN server. A reachable endpoint still needs an enabled listener, a valid authentication path, and usable client addressing before it can provide remote access.

VPN connectivity considerations in Windows Server

After installing and configuring the Remote Access role, open **Server Manager > Tools > Routing and Remote Access**. Confirm that RRAS is running and that the intended VPN protocol is enabled. *Figure 9.1* shows the Server Manager menu used to open the management console.




Figure 9.1: Opening Routing and Remote Access from Server Manager

In the server's **Properties**, separate the authentication provider from the authentication method and accounting configuration. Windows Authentication and RADIUS Authentication determine where authentication requests are handled. Methods such as EAP and MS-CHAP v2 determine how credentials are exchanged. RADIUS accounting records session activity; selecting it does not choose L2TP, IKEv2, or another tunnel protocol.

When using **Network Policy Server (NPS)** or another RADIUS service, check the applicable access policy and permitted authentication methods. The client, RRAS configuration, and authentication service must agree. Do not assume every option shown in an existing installation is enabled in a new one. In particular, new Windows Server 2025 RRAS configurations do not accept PPTP or L2TP by default, although upgraded installations may retain earlier settings. Keep only the protocols required by your approved design enabled.

For SSTP, inspect the server certificate selected in RRAS. The illustrated lab uses **admin.techpublishing.net**; substitute your organization's actual VPN hostname. Check the certificate's validity dates, private key, Server Authentication usage, trusted chain, and matching DNS identity. Check the certificate currently installed on the server rather than relying on a previous validity check. *Figure 9.2* identifies the certificate-binding area of the server properties.




Figure 9.2: Selecting the SSTP server certificate in RRAS properties

A private certification authority can issue the server certificate when managed clients trust that authority. A publicly trusted certificate may simplify trust for clients outside the domain, but it still needs the correct hostname and usage. Compare the client profile with the selected certificate before changing the binding. Similar-looking hostnames can lead you to inspect the wrong certificate, especially after a renewal or migration.

Apply certificate or protocol changes during a maintenance window and follow any prompt to restart RRAS. A restart disconnects active sessions. Reconnect a test client afterward and verify both tunnel establishment and access to an authorized internal resource before returning the service to normal use.

At this stage, the server can have the correct protocol, authentication path, and certificate and still produce a VPN that is unusable after connection. Successful authentication only establishes who the client is; the client must also receive addressing, DNS information, and a viable return path to the resources it needs. Those post-authentication dependencies are the next part of the server configuration to verify.

Configuring Windows Server for VPN and remote access

A successfully authenticated client also needs an address. In the RRAS server properties, examine the **IPv4** tab and confirm whether addresses are obtained through DHCP or allocated from a static pool. With DHCP, verify server reachability and available addresses. With a static pool, reserve a range that does not conflict with DHCP leases, static hosts, or another VPN allocation.

To configure a pool, select the static address pool option, choose **Add**, and supply the planned start and end addresses. Check the resulting range before applying it. RRAS assigns addresses from this pool to connecting clients; a pool is not a general-purpose DHCP service. Whichever assignment method you choose, verify the DNS servers and routes actually received by a connected client. Address allocation alone does not prove that internal names or remote subnets are reachable.

Use a connected test client to verify the result of address assignment. Compare its allocated address with the configured pool or DHCP range, and check whether another active device already uses it. Record the assigned address with the connection time so you can find the corresponding session in the server logs. If clients authenticate but cannot obtain an address, concentrate on allocation and pool availability before changing the authentication method.

Keep the client network and corporate ranges in mind. Overlapping home and corporate subnets can send traffic toward the wrong interface. Also check that internal routers can return traffic to the VPN client range. A tunnel can connect successfully while resource access fails because only the outbound path is correct. Retest with one known internal destination before investigating several applications at once.

Investigating VPN ports and WAN miniport problems

VPN ports and WAN miniport adapters are related, but they are not the physical network adapter carrying the server's management connection. Start with the RRAS **Ports** configuration: verify the required protocol, inbound access setting, and available ports. Then inspect the failure's event log, Windows build, and adapter status. A port-related symptom is not proof of a universal Windows Server 2025 or Windows 11 defect.

Microsoft documents client-side WAN miniport repair for specific failures, including error 720. First, check the server's address-pool availability and the affected client's adapter bindings. When the documented repair applies, uninstall only the WAN Miniport devices covered by that procedure on the affected client, then use **Action > Scan for hardware changes** to reinstall them. Do not delete unrelated physical, virtual-switch, or management adapters. *Figure 9.3* shows the hardware-rescan action, not a blanket instruction to remove network devices.




Figure 9.3: Scanning for hardware changes in Device Manager


Perform disruptive repairs with console access and a recovery plan. Afterward, verify that the miniports have returned, reconnect the VPN, and repeat the resource-access test. Preserve the original error and log entries so that a temporary improvement does not conceal a recurring problem.

Server-side checks establish that the VPN service can listen, authenticate users, and provide usable client addressing, while the miniport checks confirm that the Windows networking components needed by the tunnel are present. The other endpoint of that negotiation is the

client profile itself, where an incorrect server name, VPN type, or sign-in method can prevent an otherwise healthy server from ever receiving the expected connection.

Client-side VPN configuration in Windows 11

For an SSTP profile, distinguish the **Connection name** from the **Server name or address**. The connection name is a descriptive label. The server hostname must match an identity covered by the server certificate. *Figure 9.4* shows these as separate fields.



The screenshot shows the 'Add a VPN connection' window. It has a title bar and a light gray background. The fields are as follows:

- VPN provider:** A dropdown menu with 'Windows (built-in)' selected.
- Connection name:** A text box containing 'admin.techpublishing.net'.
- Server name or address:** A text box containing 'admin.techpublishing.net'.
- VPN type:** A dropdown menu with 'Secure Socket Tunneling Protocol (SSTP)' selected.
- Type of sign-in info:** A dropdown menu with 'User name and password' selected.
- User name (optional):** An empty text box.
- Buttons:** 'Save' and 'Cancel' buttons at the bottom.

Figure 9.4: VPN profile fields for the connection label and server address

Create the profile using the values supplied by your VPN administrator:

1. Open **Settings > Network & internet > VPN > Add VPN** and select **Windows (built-in)** as the provider.
2. Enter a recognizable connection name, such as Corporate VPN.
3. Enter the VPN server's DNS hostname in Server name or address.

4. Select the configured VPN type, such as Secure Socket Tunneling Protocol (SSTP). An explicit type makes troubleshooting clearer than automatic selection.
5. Choose the required sign-in method and enter credentials when the profile calls for them.
6. Select **Save**, then **Connect**, and complete any additional authentication prompts.

The profile should display **Connected** once the tunnel is established. *Table 9.2* distinguishes that status from the further checks needed to prove useful access.

Check	Evidence to verify
Tunnel status	The intended profile displays Connected.
Client configuration	The VPN interface has the expected address and DNS settings.
Internal access	A permitted resource is reachable by its intended name and service.

Table 9.2: Verifying a connected VPN beyond its status indicator

The table separates tunnel establishment from usable VPN access. A Connected status confirms only that the profile has established the tunnel; the VPN interface must still receive the expected network configuration, and an authorized internal resource must be reachable through that connection. Reading these checks in sequence prevents a successfully negotiated tunnel from being mistaken for a fully working remote-access path.

Once the tunnel is connected and the VPN interface has usable addressing, resource access depends on how the client finds the destination as well as how packets reach it. That makes DNS the next useful boundary to test: name resolution must work once for the public VPN endpoint before the tunnel forms and again for private corporate resources after the client is connected.

DNS configuration and name resolution for VPN connectivity

There are two DNS questions: can the client resolve the public VPN endpoint before connecting, and can it resolve internal resources afterward? For the first, check the public record for the certificate-matching server hostname. Use a DNS query from outside the corporate network. A failed ping alone does not prove that DNS failed, because ICMP replies may be blocked even when the name resolves correctly.

For internal resolution, connect the VPN and open Command Prompt. Inspect the adapter configuration with:

```
ipconfig /all
```

Find the VPN interface and compare its DNS information with the organization's intended resolver settings. Do not replace internal DNS with a public resolver as a generic fix: public DNS may not contain the organization's private records. Confirm whether the problem affects all internal names, only short names, or one resource. That distinction helps separate resolver reachability from a missing record or DNS suffix.

The lab server `li-dc01-2025` provides a useful named destination. A path such as `\li-dc01-2025` in File Explorer requests the server's shared resources; a specific share adds its share name. If an address-based test succeeds but the equivalent name-based test fails, investigate DNS. If neither succeeds, check routing and service access as well. Successful name resolution does not grant permission to open a share.

Troubleshooting common VPN error codes

Capture the complete error, protocol, timestamp, and scope of impact before applying a fix. Compare client RasClient events with RRAS and, where relevant, NPS records. An error affecting one client suggests a different starting point from simultaneous failures across the organization. *Figure 9.5* collects the codes discussed below; treat them as clues rather than one-to-one diagnoses.

Error Codes	
800	609
633	0x80072746
809	13801
691	789

Figure 9.5: VPN error codes to correlate with client and server logs

Errors 800 and 809 point first toward tunnel establishment. Error 800 reports that the attempted tunnels failed; verify the destination, selected protocol, and negotiation settings. Error 809 reports a failure to establish the connection because the remote server did not

respond. Inspect the server and intervening firewall/NAT path, including the appropriate IPsec ports. Neither code, on its own, proves that a certificate has expired.

Errors 609 and 633 concern the connection device or port. Error 609 indicates a missing device type; error 633 indicates that the port is already in use or is not configured for remote-access dialing. Review the profile, miniport devices, and service state. These messages do not specifically direct you to change an IIS certificate binding.

Error 13801 means IKE authentication credentials were rejected. For certificate-based IKEv2, inspect the relevant certificates, trust, identity, validity, and intended usage. **Error 691** concerns authentication or an authentication protocol that is not permitted; check account credentials and access policy. **Error 789** indicates a failure during the initial L2TP security negotiation, so investigate IPsec settings, including the pre-shared key or certificate, as applicable. Keep tunnel authentication separate from user authentication.

Error 0x80072746 represents a connection reset by the remote side. Correlate it with server and TLS-related events where appropriate. A shared hexadecimal prefix such as 0x800 does not identify a firewall fault. Read the full code and message before choosing a diagnostic path.

Use the scope of failure to prioritize these checks. If every SSTP client fails after a certificate renewal, compare the replacement certificate and binding with the last working configuration. If one user's connection fails while an equivalent test account works, inspect that user's policy and profile alongside the error. If the same laptop works on one network but not another, preserve the profile and investigate the access path before rebuilding it.

Avoid changing the protocol, credentials, certificate, and address pool in one attempt. That removes the comparison needed to explain an improvement or regression. Record one hypothesis, make the related change, and repeat the same test. When a change has no effect, restore it unless it is independently required by the approved configuration.

After correcting the identified cause, reconnect and retest the original operation. Record what changed and whether other clients or networks show the same result. Automated diagnostics can provide additional evidence, but they do not replace checking the configuration and confirming that the user's actual task now succeeds.

Summary

VPN troubleshooting becomes much easier when the connection is treated as a sequence of dependencies rather than as a single feature that either works or fails. The protocol determines which transport, authentication, and certificate requirements apply; the access network must allow that traffic to reach the VPN endpoint; and the server must then be listening with the expected protocol, authentication path, and certificate configuration.

Establishing the tunnel is only the first half of the investigation. A successfully authenticated client still needs usable addressing, DNS information, routing, and access to the intended corporate resource. That is why a Connected status cannot be treated as proof that the VPN is fully working. Comparing tunnel status with the client's assigned configuration and an actual internal-resource test tells you whether the problem belongs to negotiation or to the network path created after the connection is established.

Error codes help narrow those stages, but they are evidence rather than diagnoses. By correlating the complete error with the protocol, client and server logs, scope of impact, and a known-working comparison, you can decide whether to investigate the access path, authentication, certificates, miniports, addressing, DNS, or routing. The key principle is to preserve those comparisons: change one relevant dependency at a time and retest the operation that originally failed.

The VPN examples in this chapter relied on several kinds of evidence: interface configuration, DNS results, service state, ports, logs, and tests from both ends of the connection. Those same tools are useful far beyond VPN troubleshooting. In the next chapter, we will bring them together into a broader diagnostic toolkit, using Windows, PowerShell, Linux utilities, port scans, and packet captures to choose the right test for the network behavior you are trying to explain.

Part 4

Diagnosing the Enterprise Network with Evidence

The final part brings the book's troubleshooting method and technical foundations together. Once you know which fault domain needs investigation, you need tools that can expose the endpoint configuration, routing decision, resolver response, socket state, remote service, and packets involved in the failing exchange.

Chapter 10 develops that toolkit across Windows, PowerShell, Linux, Nmap, and Wireshark. The emphasis is not on memorizing commands, but on understanding what each result proves and what it leaves unanswered. *Chapter 11* then places those individual observations back into an enterprise network containing Layer 2 and Layer 3 switches, VLANs, firewalls, servers, remote offices, cloud resources, and other endpoints.

The book therefore finishes where real troubleshooting happens: correlating evidence from several points along one path. You will move from the source endpoint through routing and segmentation to the destination service, validate both the forward and return paths, make the smallest justified change, and retest the operation that originally failed.

This part of the book includes the following chapters:

- *Chapter 10, Essential Network Troubleshooting Tools*
- *Chapter 11, Troubleshooting Routing, Switching, and Firewalls*

10

Essential Network Troubleshooting Tools

Network troubleshooting tools give you different views of the same connection. One command may show the endpoint's address and gateway, another may reveal which DNS server answered, while others expose socket state, routing decisions, remote port reachability, or the packets actually exchanged on the wire. No single tool proves that the entire network path is healthy, so effective troubleshooting depends on choosing the test that answers the question you currently have.

In this chapter, we begin with Windows command-line utilities for inspecting addressing, DNS, sockets, and local neighbor information. We then move to PowerShell, where the same checks can return structured information that is easier to compare and automate. From there, we apply equivalent diagnostic questions on Linux before returning to more advanced Windows tests for latency, route tracing, and routing tables. Finally, we combine remote port scanning with Wireshark packet captures to examine what a connection is actually doing and validate a controlled firewall change.

By the end of the chapter, you should be able to choose a troubleshooting tool based on the fault domain you are investigating and interpret what its result does and does not prove. The goal is to move from observation to increasingly specific tests while preserving a baseline, changing as little as possible, and correlating evidence from the endpoint, network path, destination, and application.

In this chapter, we will cover the following topics:

- Windows command-line network troubleshooting tools
- PowerShell network troubleshooting tools
- Linux network troubleshooting tools

- Advanced Windows network troubleshooting commands
- Port scanning and packet analysis tools

Technical requirements

Use a Windows client or server with Command Prompt and Windows PowerShell, plus a Linux machine or virtual machine for the cross-platform examples. The Windows examples are suitable for a lab using Windows 11 and Windows Server; command availability and output can vary by version. Administrative privileges are required for some changes, but not for every read-only check.

For the final exercise, install Wireshark with the appropriate packet-capture support, including Npcap on Windows, and Nmap/Zenmap on an authorized scanning machine. Capture traffic and scan ports only on systems you own or have permission to test. The illustrated addresses belong to the chapter's lab; substitute your own endpoints. Output examples are snapshots, not results that every machine should reproduce. Keep console access available before changing connectivity on a remotely managed host.

Windows command-line network troubleshooting tools

Begin with the endpoint's own view of the network. An incorrect address or resolver can make a functioning server appear unavailable, while a local listener may be running even when a firewall blocks remote access. Windows utilities let you check these layers independently. Preserve the relevant output with the time, interface, and destination so later tests can be compared meaningfully.

IP configuration and management with `ipconfig`

Open **Start**, search for **Command Prompt**, and launch it. Use **Run as administrator** when a command requires elevation, supplying an authorized administrator account when prompted. Changing to the root directory is not necessary for the network commands in this section. You can use `cls` to clear the display without altering the network configuration.

Start with the basic address information:

```
ipconfig
```

Identify the active adapter and read its IPv4 address, subnet mask, and default gateway. A machine may have Ethernet, Wi-Fi, VPN, and virtual adapters, so a plausible address on one interface does not establish that it is the interface being used for the failing connection. Compare the values with the intended subnet before changing anything.

For the complete adapter configuration, run:

```
ipconfig /all
```

This adds details such as the physical address, DHCP status, lease information, and DNS servers. Check **DHCP Enabled** to distinguish a dynamically configured interface from a static configuration. For a VPN issue, compare the physical and VPN interfaces; for a DNS issue, note which resolvers are configured rather than assuming the default gateway also provides DNS.

On a DHCP-configured interface, the following commands release and then request an address lease:

```
ipconfig /release  
ipconfig /renew
```

Releasing an address interrupts connectivity. These unqualified commands can affect multiple DHCP-enabled adapters, so use them carefully during remote support. Renewal may return the same address; it does not guarantee a new one or permanently resolve an address conflict. If two devices are using the same address, investigate the DHCP scope, reservations, and any conflicting static assignment. Statically configured addresses are not repaired by renewing a DHCP lease.

DNS cache operations address a different problem. To remove locally cached resolver entries, run:

```
ipconfig /flushdns
```

Windows reports the local operation as follows:

```
Windows IP Configuration  
  
Successfully flushed the DNS Resolver Cache.
```

This removes stale local entries, but does not correct an inaccurate authoritative record or flush an upstream resolver's cache. Retest the same name afterward and compare the answer. Repeated cache clearing without checking the returned address does not identify where incorrect information originated.

To initiate dynamic registration of the computer's DNS records, use:

```
ipconfig /registerdns
```

The message confirms that registration was initiated, not that every DNS server has already been updated:

Windows IP Configuration

```
Registration of the DNS resource records for all adapters of this computer has
been
initiated. Any errors will be reported in the Event Viewer in 15 minutes.
```

Registration depends on DNS configuration and update permissions. Check the relevant event log and query the record afterward. Keep the two operations distinct:

- `/flushdns` affects this client's cached answers
- `/registerdns` requests registration of this computer's records.

Neither is a substitute for correcting its address, DNS suffix, or configured DNS server.

Viewing network connections with netstat

Once addressing is understood, inspect the local transport endpoints. `netstat` shows socket information, including active TCP connections and, with the appropriate options, listening TCP ports and UDP endpoints. It does not inspect the effective firewall policy or prove that a remote client can reach a displayed port.


Use numeric output to avoid name-resolution delays:

```
netstat -an
```

The `-a` option includes listeners, and `-n` displays numerical addresses and ports. To review a long result one screen at a time, pipe it to `more`:

```
netstat -an | more
```

Press the space bar to advance a page. *Figure 10.1* shows listening sockets in the lab; compare the protocol, local address, foreign address, and state rather than treating the output as a list of firewall permissions.



The screenshot shows a Windows Command Prompt window titled "Administrator: Command Prompt". The output of the netstat command is displayed, showing active connections. The title "Active Connections" is at the top. Below it is a table with four columns: Proto, Local Address, Foreign Address, and State. The table lists several TCP connections in a LISTENING state, all with a local address of 0.0.0.0 and a foreign address of 0.0.0.0. The ports listed are 80, 81, 88, 135, 389, 443, 445, 464, 593, 636, 1723, 3268, and 3269. At the bottom, there is a "-- More --" prompt with a cursor pointing to it.

Proto	Local Address	Foreign Address	State
TCP	0.0.0.0:80	0.0.0.0:0	LISTENING
TCP	0.0.0.0:81	0.0.0.0:0	LISTENING
TCP	0.0.0.0:88	0.0.0.0:0	LISTENING
TCP	0.0.0.0:135	0.0.0.0:0	LISTENING
TCP	0.0.0.0:389	0.0.0.0:0	LISTENING
TCP	0.0.0.0:443	0.0.0.0:0	LISTENING
TCP	0.0.0.0:445	0.0.0.0:0	LISTENING
TCP	0.0.0.0:464	0.0.0.0:0	LISTENING
TCP	0.0.0.0:593	0.0.0.0:0	LISTENING
TCP	0.0.0.0:636	0.0.0.0:0	LISTENING
TCP	0.0.0.0:1723	0.0.0.0:0	LISTENING
TCP	0.0.0.0:3268	0.0.0.0:0	LISTENING
TCP	0.0.0.0:3269	0.0.0.0:0	LISTENING

-- More --

Figure 10.1: Inspecting local TCP listeners with paged netstat output

A TCP LISTENING row represents a local socket waiting for connection requests. The foreign address is normally a wildcard because the listener is not tied to one remote peer. This does not mean the service has no clients: accepted connections appear as separate rows. An ESTABLISHED row indicates that the TCP connection has been established, although it may be idle at the instant you inspect it.

For example, a directory service can listen on TCP 389, the usual LDAP port, while maintaining many established connections to that same local port. **LDAP** means **Lightweight Directory Access Protocol**. Clients generally use ephemeral source ports, and TCP distinguishes connections using both endpoint addresses and ports. The server does not move an accepted LDAP connection to a high-numbered port merely to free port 389 for the next client.

Read wildcard and loopback bindings carefully. A listener on 127.0.0.1 is not the same as a listener on all IPv4 interfaces, shown as 0.0.0.0. IPv6 has corresponding bindings such as ::1 and ::. Preserve the address family in your investigation. A service listening only on an unintended interface can fail remote tests even when its port number is correct.

DNS troubleshooting with nslookup

Use nslookup to ask a DNS server for records and identify which server answered. Running it without a query enters interactive mode and displays the default resolver:

```
nslookup
```

The lab session begins with:

```
Default Server:  DC1.TechPub.us
Address:  192.168.21.110
>
```

The > prompt belongs to interactive nslookup, not the normal command shell. Enter exit to return to Command Prompt. Opening this session does not refresh DNS records. The displayed server identifies the resolver selected for the query; its presence alone does not prove that all subsequent queries will succeed.

To request address records directly from the command line, supply a name:

```
nslookup google.com
```

The following captured result contains both IPv6 and IPv4 answers:

```
Server:  DC1.TechPub.us
Address:  192.168.21.110

Non-authoritative answer:
Name:     google.com
Addresses: 2607:f8b0:400a:809::200e
          172.217.3.206
```

Here, 192.168.21.110 is the responding resolver. The addresses below Name are records for the requested name, not the addresses of its DNS servers. Non-authoritative answer means that the responding server is not authoritative for that zone. It can return information obtained through recursion or caching; the client is not necessarily redirected to contact another DNS server itself.

Authoritative servers publish a zone's records. Recursive resolvers obtain and cache answers according to DNS behavior; those cached answers are not a hierarchy of writable copies replicated down from the root. This distinction helps when one client receives an old answer while another receives the intended one.

For an internal failure, compare the returned address with the expected host and repeat the query against the intended internal resolver when necessary. A timeout, a negative answer, and an unexpected but valid address are different results. Also distinguish name resolution from service availability: receiving an address does not prove that the server's web, file-sharing, or directory service can be reached.


Address Resolution Protocol for network diagnostics

For IPv4 communication on a local link, Address Resolution Protocol (ARP) maps a next-hop IP address to a link-layer address. Display the local cache with:

```
arp -a
```

Read the interface heading first on a multihomed computer. Each list relates to an interface, and the entries include IP addresses, physical addresses, and types. For an off-subnet destination, the relevant local neighbor is usually the gateway, not the remote server. IPv6 uses Neighbor Discovery rather than ARP, so this command is not a complete view of both address families.

Figure 10.2 shows dynamic entries alongside broadcast and multicast mappings. Those rows have different purposes and should not all be described as recently contacted individual computers.



```

Administrator: C:\Windows\system32\cmd.exe

Interface: 192.168.21.29 --- 0x3
  Internet Address      Physical Address      Type
  192.168.21.1          70-df-2f-32-4e-03     dynamic
  192.168.21.10         68-05-ca-24-71-34     dynamic
  192.168.21.45         00-15-5d-15-0f-4d     dynamic
  192.168.21.110        00-15-5d-15-0f-47     dynamic
  192.168.21.113        00-15-5d-15-0f-43     dynamic
  192.168.21.114        00-15-5d-15-0f-44     dynamic
  192.168.21.126        00-15-5d-15-0f-47     dynamic
  192.168.21.199        00-15-5d-15-0f-47     dynamic
  192.168.21.255        ff-ff-ff-ff-ff-ff     static
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
  224.0.0.252           01-00-5e-00-00-fc     static
  224.85.0.0            01-00-5e-55-00-00     static
  224.167.0.0           01-00-5e-27-00-00     static
  224.250.0.0           01-00-5e-7a-00-00     static
-- More --

```

Figure 10.2: ARP cache entries grouped under the local IPv4 interface

A dynamic entry shows that a mapping was learned and retained; it is not proof that the device remains reachable or that an application works. A MAC address may be associated with a physical or virtual interface and can be locally administered. Avoid assuming every displayed address is an immutable identifier burned into a separate physical card.

Consider this unsuccessful probe from the lab:

```
ping 192.168.21.126
```

The reply is generated by the source computer, not the intended destination:

```

Pinging 192.168.21.126 with 32 bytes of data:
Reply from 192.168.21.110: Destination host unreachable.

```

This message does not prove that the target firewall is blocking ICMP. It can reflect a local delivery problem, including failure to resolve a next-hop link-layer address. Check the relevant interface, subnet, route, and ARP information, then compare the result with another known local device.

An ARP entry for 192.168.21.126 observed at another time is additional evidence, but it may be stale or learned under different conditions. Conversely, a device can answer ARP while declining ICMP echo requests. Use the observations together without turning either one into

proof of application health. This same distinction between a local observation and an end-to-end result will recur with port tests and packet captures.

PowerShell network troubleshooting tools

Windows PowerShell provides structured objects that can be filtered, selected, and reused in scripts. Open a Windows PowerShell session, including through the appropriate Windows Terminal profile. A PS prompt identifies the shell; it does not indicate that the session is elevated. Use an administrator session only when required by the operation.

Traditional executables such as `ping`, `nslookup`, and `netstat` can also run from PowerShell. They do not become PowerShell cmdlets simply because they are launched there. Begin with a familiar IP reachability test:

```
ping 8.8.8.8
```

The following lab result shows four successful replies:

```
Pinging 8.8.8.8 with 32 bytes of data:
Reply from 8.8.8.8: bytes=32 time=4ms TTL=114
Reply from 8.8.8.8: bytes=32 time=4ms TTL=114
Reply from 8.8.8.8: bytes=32 time=4ms TTL=114
Reply from 8.8.8.8: bytes=32 time=4ms TTL=114

Ping statistics for 8.8.8.8:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 4ms, Maximum = 4ms, Average = 4ms
```

The statistics describe those four probes, not a guarantee of ongoing availability. ICMP response policy varies by host and firewall configuration; domain-controller status alone does not determine whether a machine answers. Compare an ICMP test with the actual service being investigated when the results appear inconsistent.

Testing connectivity with Test-NetConnection

Test-NetConnection combines several useful checks, including ICMP reachability, TCP connection testing, and route diagnostics. Start with a named internal destination:

```
Test-NetConnection -ComputerName li-dc01-2025
```

The sample result identifies the resolved destination and the selected local interface:

```
ComputerName      : li-dc01-2025
RemoteAddress     : 192.168.21.92
InterfaceAlias    : Ethernet
SourceAddress     : 192.168.21.110
PingSucceeded     : True
PingReplyDetails (RTT) : 0 ms
```

Read `RemoteAddress` and `SourceAddress` before `PingSucceeded`. A successful test through the wrong VPN or physical interface may not reproduce the user's problem. `PingReplyDetails (RTT)` reports the round-trip time for the echo in milliseconds; a displayed zero reflects the reporting resolution, not an instantaneous exchange.

Request the detailed view with the full parameter name:

```
Test-NetConnection -ComputerName li-dc01-2025 -InformationLevel Detailed
```

The additional route information in the lab is:

```
ComputerName      : li-dc01-2025
RemoteAddress     : 192.168.21.92
NameResolutionResults : 192.168.21.92
InterfaceAlias    : Ethernet
SourceAddress     : 192.168.21.110
NetRoute (NextHop) : 0.0.0.0
PingSucceeded     : True
PingReplyDetails (RTT) : 0 ms
```

A next hop of `0.0.0.0` in this on-link example does not mean that the default gateway is missing. It indicates that this route does not require an intermediate gateway. Compare that result with an external destination:

```
Test-NetConnection -ComputerName 8.8.8.8 -InformationLevel Detailed
```

The external test selects the lab gateway:

```
ComputerName      : 8.8.8.8
RemoteAddress     : 8.8.8.8
NameResolutionResults : 8.8.8.8
                  : dns.google
```

```
InterfaceAlias      : Ethernet
SourceAddress       : 192.168.21.110
NetRoute (NextHop)  : 192.168.21.1
PingSucceeded       : True
PingReplyDetails (RTT) : 5 ms
```

The reported next hop is now 192.168.21.1. These two tests demonstrate why a working local connection does not prove that the gateway or wider path works. Compare repeated measurements with a known baseline; one high reply time can reflect congestion, scheduling, or ICMP handling and does not by itself establish a routing loop.

Tracing network paths with Test-NetConnection

To examine the path toward the external destination, enable route tracing:

```
Test-NetConnection -ComputerName 8.8.8.8 -TraceRoute
```

The recorded output contains responding hops and unresolved positions:

```
ComputerName      : 8.8.8.8
RemoteAddress     : 8.8.8.8
InterfaceAlias    : Ethernet
SourceAddress     : 192.168.21.110
PingSucceeded     : True
PingReplyDetails (RTT) : 5 ms
TraceRoute        : 199.107.166.113
                   206.72.111.74
                   206.72.100.18
                   0.0.0.0
                   208.185.182.57
                   64.125.20.212
                   0.0.0.0
                   0.0.0.0
                   0.0.0.0
                   142.250.169.166
                   108.170.255.179
                   209.85.254.171
                   8.8.8.8
```

Entries shown as 0.0.0.0 do not identify real routers at that address. They indicate positions where a usable hop address was not obtained. A router can forward traffic while declining or

limiting diagnostic replies, so a missing response is not automatically a broken path. The successful destination response is important evidence, even when intermediate hops are incomplete.

Compare traces from the affected and an unaffected machine, preserving the source interface and time. Do not assume every probe follows a permanently fixed path. Routing changes, multiple paths, and different ICMP policies can alter the output without reproducing an application failure.

Resolving names with PowerShell

For a reverse lookup, a traditional DNS query can still be useful:

```
nslookup 8.8.8.8
```

A returned name such as `dns.google` is a reverse DNS result, not a test of the resolver service on that address. Similarly, `nslookup microsoft.com` normally asks for address records for that name. It does not enumerate all DNS servers responsible for the domain.

The PowerShell alternative returns structured DNS information:

```
Resolve-DnsName -Name microsoft.com
```

Inspect the record type, name, and returned data. A and AAAA records represent IPv4 and IPv6 addresses; NS records serve a different purpose. Compare the answer with the application's intended destination. When a query fails, distinguish an unreachable DNS server from a negative response or an inappropriate resolver for an internal name.

For a reproducible comparison, keep the requested name and record type unchanged while selecting the resolver explicitly. The following example asks the lab's internal DNS server for an IPv4 address record:

```
Resolve-DnsName -Name li-dc01-2025 -Type A -Server 192.168.21.110
```

Substitute the resolver appropriate to your network. Comparing this result with the default lookup can reveal a difference in resolver selection, but it does not automatically explain why that difference exists. Check the adapter configuration and intended DNS policy before changing them.

PowerShell output can be incorporated into repeatable diagnostics, but retain the question being tested. A name resolving successfully is only one step toward a functioning connection. Use the resulting address in a reachability or service test rather than declaring the issue resolved after a DNS answer appears.

Inspecting TCP connections with Get-NetTCPConnection

To inspect local connections, run the familiar utility or its TCP-focused PowerShell counterpart:

```
netstat  
Get-NetTCPConnection
```

Get-NetTCPConnection exposes fields such as local and remote addresses, ports, state, and owning process. It covers TCP, not UDP endpoints. IPv4, IPv6, wildcard, and loopback entries can all be relevant; filtering out everything except conventional IPv4 addresses can hide the service binding that explains a failure.

Read CLOSE_WAIT as a connection whose peer has closed its sending side while the local application still needs to close its socket. It is not simply a promise that the connection will disappear shortly. Many persistent entries can justify investigating the owning application. TIME_WAIT, by contrast, is part of TCP's normal connection teardown. Interpret states in context instead of treating every non-established row as a fault.

A local ephemeral port such as 51020 connecting to remote TCP 443 is ordinary client behavior. The service remains on its destination port while endpoint pairs distinguish concurrent sessions. A port number convention suggests a service but does not prove the traffic is encrypted or the application is healthy. Use process information and, where authorized, packet analysis to resolve that question.

For a focused local inspection, request only listening TCP sockets and retain their owning process identifiers:

```
Get-NetTCPConnection -State Listen |  
Select-Object LocalAddress, LocalPort, OwningProcess
```

Run this on the machine whose service you are investigating. The output provides a manageable set of bindings to compare with a remote connection test. It still does not establish firewall reachability, and a process identifier should be checked on that same machine rather than compared across unrelated hosts.

The value of PowerShell is consistency: the same fields can be selected for repeated comparisons. Record the source, destination, port, and state together so that the output can be related to a specific user operation.

Linux network troubleshooting tools

The diagnostic questions remain the same on Linux even when command names and output differ. Inspect the interface, route, name resolution, and service in that order. Some utilities in this section are optional packages rather than preinstalled components. Use the distribution's approved package manager and avoid assuming that every Linux system has the same interface names or administrative configuration.

Viewing network configuration with `ifconfig`

Where the `net-tools` package is installed, display active interfaces with:

```
ifconfig
```

Read the `inet` address, netmask, broadcast address, and interface flags together. In the lab, the Ethernet interface uses `192.168.21.158` with a `/24` network. The IPv4 address is not the Ethernet/MAC address; the latter is a separate link-layer value. IPv6 entries and the loopback interface are separate parts of the output, not alternative names for the same connection.

Transmit and receive counters help establish whether an interface is carrying traffic. Increasing error or drop counters merit investigation, while a zero-error snapshot does not prove that the entire path is healthy. Compare counters over time and relate them to the actual interface used by the failing application.

On systems without `ifconfig`, the corresponding modern address-inspection command is:

```
ip address show
```

Keep `ifconfig` recognizable for existing environments, but use the available supported tool rather than installing an older utility solely to reproduce a screenshot. Both approaches should lead to the same essential checks: address, prefix, operational state, and the correct interface.

Managing routing tables with `ip route`

Inspect the IPv4 routing table with:

```
ip route
```

Identify the default route, any via next hop, and the `dev` interface associated with the relevant destination. Connected routes and more-specific remote routes can explain why a host reaches one network but not another. Compare this with `route print` on Windows: the display differs, but both show the local forwarding choices.

The `ip route` command can also modify routes, but inspection should come first. A host route is not a substitute for a missing return route on the remote side. Before adding one, establish the intended next hop, ensure that it is reachable, and decide how the change will be persisted by the distribution's network configuration. Do not treat an ad hoc command as a complete network design.

Connectivity testing with ping and traceroute

Linux ping normally continues until interrupted:

```
ping 8.8.8.8
```

Press `Ctrl+C` to stop and review the statistics, or limit the sample explicitly:

```
ping -c 4 8.8.8.8
```

Record loss and latency alongside the source and destination. A successful echo reply establishes ICMP reachability at that moment, not access to a particular TCP or UDP service. If a public destination does not answer, compare with an approved known-responsive target rather than immediately declaring the internet connection down.

For a hop-by-hop view, use the full Linux command name:

```
traceroute 8.8.8.8
```

Linux implementations and their probe methods can differ from Windows `tracert`. An asterisk marks an unanswered probe within its timeout; filtering, rate limiting, or return-path behavior may explain it. Compare later hops and the destination before assigning blame to the first unanswered position. Repeated traces and application-specific tests provide stronger evidence than a single incomplete path.

DNS diagnostics with dig

`dig` can query a chosen DNS server and report the response in detail. The `@` prefix identifies the server to query; the following name and type identify the requested record. For example:

```
dig @8.8.8.8 microsoft.com A
```

This requests an IPv4 address record from the specified resolver. It does not configure forwarding or demonstrate that the computer's normal resolver forwards requests there. To examine the system's configured resolver path, query without specifying `@server` and compare the reported server with your intended configuration.

Review the response status, question, answer records, query time, and server. A successful DNS transaction may return no record of the requested type, while a negative answer differs from a timeout. An SOA record in a response relates to the authoritative zone information; it is not evidence that every unresolved query is sent to a particular public forwarder. Use the output to establish what was asked and answered before drawing conclusions about DNS configuration.

Port scanning and network analysis with nmap and netstat

nmap tests a remote host from the scanning machine's position. After installing it through the appropriate distribution package manager, perform a basic scan only against an authorized lab target:

```
nmap 192.168.21.110
```

A default scan examines a commonly used set of TCP ports, not every possible port or vulnerability. An open result indicates that an application is accepting connections on that port from this vantage point. Service names based on port conventions are clues rather than proof of the exact application. Ports associated with DNS, LDAP, HTTP, or HTTPS can help identify which service-specific checks to make next.

On the target itself, inspect local connections separately:

```
netstat
```

Linux netstat is also supplied by net-tools on many systems. Where it is unavailable, ss provides modern socket inspection. Do not confuse a scan of a remote endpoint with a list of sockets on the current machine. Comparing the two can separate a missing listener from a filtering or path issue.

To view the interface statistics provided by netstat, use:

```
netstat -i
```

Examine transmit, receive, and error counters in the context of the interface's workload. Unexpected connections merit investigation, but neither an unfamiliar address nor an open port proves compromise. Validate the service owner and intended communication before changing policy.

Advanced Windows network troubleshooting commands

Return to Windows to extend the earlier reachability tests. Longer samples, address-family selection, packet-size changes, and route inspection answer different questions. Keep a record of each test's parameters; otherwise, two results that appear contradictory may simply have measured different traffic.

Using ping switches deliberately

A continuous probe can reveal intermittent loss during an authorized maintenance test:

```
ping -t 192.168.21.1
```

Stop with *Ctrl+C*. The following recorded summary shows a short loss-free sample:

```
Ping statistics for 192.168.21.1:
    Packets: Sent = 44, Received = 44, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
Control-C
^C
```

Use timestamps and an agreed maintenance window when correlating interruptions with a cable or configuration change. Do not unplug unidentified production links to see which ping stops. A loss-free sample only establishes what happened during the sample; it does not exclude earlier or application-specific failures.

Specify the destination as well as the address-family switch:

```
ping -4 8.8.8.8
```

The *-l* switch changes the ICMP data payload size, not the size of the entire IP packet. For example:

```
ping -l 1500 192.168.21.1
```

Headers add to that payload. Depending on the path MTU and fragmentation behavior, this test may produce multiple IP fragments. A large successful echo therefore does not prove that

the path carries an equally large unfragmented packet. The lab also demonstrates a 5,000-byte payload:

```
ping -l 5000 192.168.21.1
```

The recorded replies are:

```
Pinging 192.168.21.1 with 5000 bytes of data:
Reply from 192.168.21.1: bytes=5000 time=1ms TTL=255
Reply from 192.168.21.1: bytes=5000 time=1ms TTL=255
Reply from 192.168.21.1: bytes=5000 time<1ms TTL=255
Reply from 192.168.21.1: bytes=5000 time=1ms TTL=255

Ping statistics for 192.168.21.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

Compare repeated small- and large-payload tests without assuming that increased latency identifies one particular device. Fragmentation, congestion, processing, and ICMP handling can affect the results. Change one test parameter at a time so that differences remain interpretable.

Path analysis with pathping and tracert

Use `tracert` when you need an initial view of responding hops, and `pathping` when you also want a longer sample of response and loss statistics:

```
tracert -d 8.8.8.8
pathping -n 8.8.8.8
```

Disabling intermediate name lookups helps separate tracing from reverse DNS delays. `pathping` takes time to collect its statistics; allow the sampling stage to finish before comparing results. Loss shown only for an intermediate router can reflect how that router handles probes rather than loss of forwarded traffic. Check whether the destination and later hops show a corresponding problem.

For escalation, record the affected operation, destination, time, repeated measurements, and a comparison with a working path. These observations help an ISP or network team investigate; an incomplete trace does not identify an exact faulty router with certainty. Apply the same caution when tracing between internal VLANs.

Viewing routing tables with route print

Display the host's routing table with:

```
route print
```

For IPv4, examine **Network Destination**, **Netmask**, **Gateway**, **Interface**, and **Metric**. The all-zero destination and mask identify a default route. A more-specific matching prefix takes precedence over a less-specific one; metrics help choose among otherwise comparable routes. Check persistent entries as well as active routes when a problem returns after a restart.

Compare the route with `ipconfig` and an actual connection test. A valid-looking route through the wrong interface may explain a VPN failure, while an on-link route calls for local-link investigation. When the route appears correct but the service still fails, the next useful question is whether traffic reaches the destination and how that destination responds. Port scans and packet captures provide that evidence from different perspectives.

Port scanning and packet analysis tools

A connection test tells you whether an exchange succeeded. A packet capture can show the requests, replies, retries, and resets visible at the capture point. Wireshark and Nmap complement rather than replace one another: Nmap generates probes and summarizes responses, while Wireshark records traffic for closer inspection.

Capturing relevant traffic with Wireshark

Choose the capture point before starting. Install Wireshark and its capture support on the system where traffic must be observed; installation on both client and server is useful for a two-sided investigation, but is not mandatory for every capture. On a switched network, capturing on one endpoint does not automatically reveal unrelated traffic between other machines.

Select the interface that carries the failing connection, not merely the first adapter named **Ethernet**. A host may also use Wi-Fi, VPN, or virtual interfaces. Start a capture, reproduce one known operation, and stop promptly. *Figure 10.3* shows the packet list, decoded details, and raw-byte panes for a lab capture.




Figure 10.3: Packet list, protocol details, and packet bytes in Wireshark

Expect background traffic such as ARP, DNS, and ordinary TCP exchanges. Its presence is not an error. Identify the endpoints and time of the operation you reproduced, then inspect the corresponding packets. A capture on the outside of a VPN may show encrypted tunnel traffic rather than the original application exchange; encrypted payloads are not automatically readable because Wireshark can display the packet headers.

Store capture files securely because names, addresses, and unencrypted content may disclose sensitive information. Collect the traffic needed for the investigation rather than an unrestricted record of unrelated user activity.

Comparing a scan with the captured traffic

In Zenmap, choose an authorized lab target and review the command associated with the selected scan profile. Figure 10.4 shows results from a Windows server. Port labels such as DNS, LDAP, SMB, and RDP help direct investigation, but the services enabled on your own server may differ.




Figure 10.4: Reviewing discovered ports in the Zenmap interface

The **Intense scan plus UDP** profile adds UDP probing and other discovery activity. It can take longer and generate substantially more traffic than a quick scan, so use it only within the approved scope. UDP is connectionless, but that does not mean UDP applications cannot provide encryption or authentication. A lack of a reply can also make UDP port states less conclusive than a completed TCP connection.

While the lab scan runs, capture traffic at the target and correlate probes with the known scanning address. Use a display filter for that complete address rather than searching visually for an abbreviated final octet. Stop the capture and select individual exchanges to see which ports were probed and how the target responded. An open port is an exposed service to review, not proof that a vulnerability has been exploited.

You may encounter SMB alongside the scan traffic. **Server Message Block** supports file and printer sharing, administrative shares, and named-pipe communication. Windows clients also use SMB to access domain-controller shares such as SYSVOL and NETLOGON. Modern server workloads can use SMB features such as Multichannel and SMB Direct. Investigate the endpoints and purpose before treating all SMB traffic as suspicious or blocking it wholesale.

Other decoded protocols can include DNS, **Session Initiation Protocol (SIP)** for signaling, and **Datagram Transport Layer Security (DTLS)** for protected datagram exchanges. Expand a relevant packet to inspect its link-layer, IP, and transport fields. Ethernet II is a frame format,

not an indicator that the network is small. For DNS, compare the question and response data; for TCP, examine the handshake or reset that relates to the failed test.

Validating a targeted firewall change

Suppose the lab review confirms that inbound TCP 80 should not be available from the test client. First, confirm that no required application depends on it. Then create a narrowly scoped rule instead of disabling the firewall or blocking unrelated services:

1. Open **Windows Defender Firewall with Advanced Security** and select **Inbound Rules > New Rule**.
2. Select **Port**, choose **TCP**, and enter **80** as the specific local port.
3. Select **Block the connection** and choose the profiles required by the lab policy.
4. Give the rule a meaningful name, complete the wizard, and verify its enabled state and scope.

Windows Firewall does not use the visual list position as a simple first-match rule order. An applicable explicit block generally overrides a conflicting allow rule. Verify the actual profile and scope instead of moving a rule to the top and assuming that this controls precedence.

Repeat a comparable scan from the same authorized client. The quick-scan syntax supplies the target directly:

```
nmap -T4 -F li-dc01-2025
```

Here, **-F** selects a reduced set of commonly scanned ports and **-T4** selects a timing template. It is not a passive scan. If scanning triggers a security control, coordinate with the responsible team and adjust the approved test plan; do not attempt to bypass detection.

The lab output lists these open ports after the change:

PORT	STATE	SERVICE
53/tcp	open	domain
135/tcp	open	msrpc
389/tcp	open	ldap
445/tcp	open	microsoft-ds
3389/tcp	open	ms-wbt-server

Port 80 is absent from this excerpt, but an abbreviated list alone should not replace a targeted retest. A blocked probe may be reported as **filtered** rather than **closed**. Compare the full scan

result with the host's listener and firewall configuration, and restore the previous configuration when the lab exercise is complete. This closes the loop between observation, a controlled change, and validation from the original source.

Summary

Network troubleshooting tools are most useful when each one is treated as a way to answer a specific question rather than as another command to run. `ipconfig` establishes how the endpoint is configured, DNS tools show what names resolve to, ARP and routing information describe local forwarding decisions, and socket tools reveal what the operating system is listening for or communicating with. None of those observations alone proves that an application can complete an end-to-end exchange.

PowerShell and Linux utilities gave us alternative views of the same fault domains, reinforcing that the troubleshooting method should survive a change of operating system. Reachability tests, route traces, resolver queries, and socket inspection become more useful when their source interface, destination, timing, and parameters are preserved so that affected and known-working paths can be compared. The same caution applies to diagnostic gaps: an unanswered ICMP probe does not automatically identify a failed router, just as an open local listener does not prove that a remote firewall permits access to it.

Finally, Nmap and Wireshark moved the investigation from host configuration into observed network behavior. A scan shows how a remote service appears from the scanning host, while a packet capture records the requests and responses visible at a particular point in the path. Used together, they allow you to test a hypothesis, make a narrowly scoped change, and repeat the same operation to verify whether the evidence changed. That progression from baseline to controlled test to validation is the troubleshooting discipline to carry into the enterprise infrastructure scenarios that follow.

The tools in this chapter have mostly examined a connection from the perspective of individual hosts and specific points along the path. In an enterprise network, however, packets may cross access switches, VLAN boundaries, routed interfaces, and firewalls before reaching the destination. In the next chapter, we will apply the same evidence-driven method across that infrastructure, combining local configuration, routing information, remote connection tests, and service state to determine exactly where an enterprise path stops behaving as intended.

11

Troubleshooting Routing, Switching, and Firewalls

Routing, switching, and firewalling determine how traffic moves once it leaves an endpoint. A client can have a valid IP configuration and a server can have the correct service listening, yet communication can still fail because the source is attached to the wrong VLAN, a required route is missing, the return path differs from the forward path, or a firewall policy prevents the intended connection. Troubleshooting at this level therefore means following the expected traffic path and finding the first point where the network stops behaving as designed.

In this chapter, we begin by mapping the major components of an enterprise network, including Layer 2 and Layer 3 switches, VLANs, firewalls, servers, wireless infrastructure, printers, VoIP systems, remote offices, cloud resources, and IoT devices. We then follow traffic between two VLANs to see how host routes, switch virtual interfaces, access ports, trunks, and return paths work together. From there, we turn that known path into a troubleshooting sequence before finishing at the application boundary, where remote TCP tests, local listeners, and firewall policy help distinguish a routing failure from a service that is unavailable or blocked.

By the end of the chapter, you should be able to take the diagnostic evidence developed throughout this book and apply it across an enterprise network rather than at a single endpoint. The goal is to map the expected path, test each dependency in sequence, identify the first failing boundary, make the smallest justified change, and verify the original user operation from end to end.

In this chapter, we will cover the following topics:

- Core Layer 3 switches and inter-VLAN routing
- Layer 2 switches and VLAN segmentation

- Firewall placement and internet connectivity
- Server roles in network infrastructure
- Network printers and office placement
- Remote office connectivity: SD-WAN, MPLS, and VPN
- Internet of Things devices in modern networks
- Working through an inter-VLAN connection
- Troubleshooting inter-VLAN routing and subnet communication
- Troubleshooting Windows Firewall and port connectivity

Technical requirements

You should understand IP addressing, subnet masks, gateways, VLANs, and the diagnostic commands from *Chapter 10*. For the practical examples, use two authorized Windows lab hosts and a Layer 3 switch capable of routing between VLANs. Cisco-style commands are shown; verify interface names and syntax against your switch model. Administrative access is required for persistent host routes, switch configuration, and firewall changes. Keep a configuration backup and console access available, and test on an isolated lab rather than altering a production uplink without a change plan.

Core Layer 3 switches and inter-VLAN routing

Start by identifying where routing occurs. In the example network, the core Layer 3 switch forwards traffic within VLANs and routes between IP subnets through configured Layer 3 interfaces. This is the point where an otherwise local-looking connection may cross both a subnet boundary and an access-control policy.

A redundant core can reduce the impact of a device failure, but high availability depends on the actual design, gateway redundancy, and failover configuration. Two installed switches do not guarantee interruption-free service. Record which device or virtual gateway currently forwards the affected traffic and whether a recent failover changed the path.

During troubleshooting, distinguish the logical topology from the rack layout. A nearby server can still be in another VLAN, while devices on different access switches can share a VLAN. The destination's subnet and the forwarding configuration determine the route, not physical proximity.

Layer 2 switches and VLAN segmentation

A Layer 2 switch forwards Ethernet frames within VLANs; it does not perform IP routing between them. An access port typically carries an endpoint's assigned VLAN, while an

appropriately configured trunk carries multiple VLANs between switches. For an off-subnet destination, the host sends the frame toward its next-hop router through this Layer 2 path.

Check the endpoint's port membership and the VLANs allowed across each uplink. A correct IP address cannot compensate for a port placed in the wrong VLAN. Similarly, a VLAN configured at the core is not sufficient when the access-switch uplink does not carry it. The core and access-switch grouping in Figure 11.1 provides the context for these checks.

Firewall placement and internet connectivity

In this design, the firewall connects the internal network to the upstream internet service. Internal traffic may pass from an access switch through the core and then the firewall, whereas traffic between two internal VLANs may remain at the core unless the design deliberately routes it through a firewall. Establish the real path before troubleshooting a device that the packets never traverse.

The internet cloud in a diagram represents upstream connectivity, not a claim that every external service is a cloud-hosted application. Check the firewall's internal and external interfaces, relevant route, translation where required, and access policy. Return traffic also needs a valid path. *Figure 11.1* places the firewall, remote-office connectivity, and cloud resources alongside the internal components.




Figure 11.1: Enterprise components to include when mapping a traffic path

Use the overview as an inventory, then annotate the specific source, destination, VLANs, gateways, and services for the incident. It is not a complete cable map or a substitute for the

actual running configuration. The same infrastructure can carry several paths with different security requirements.

Before testing, write down the source address, destination address, and required service port. Add the source gateway, the routing device between the subnets, and any firewall on that path. Then make the same check in the reverse direction. This small record prevents a serious mistake: testing one destination or protocol successfully and assuming that the result applies to a different application. It also gives another administrator enough context to reproduce the failure without repeating the entire investigation.

Server roles in network infrastructure

Identify service dependencies before blaming forwarding. Active Directory domain controllers provide directory services, DHCP supplies client configuration, and DNS resolves names. These roles may share a server or run separately. A client that cannot resolve an application name needs a different investigation from one that resolves the name correctly but cannot establish a connection.

File and print servers provide shared resources. Some server environments use iSCSI to present remote block storage to a host, where it can appear as a locally attached disk. A **storage area network (SAN)** is the storage-networking infrastructure, not simply another name for a single storage box. When storage access fails, a file service can become unavailable even though the client can still reach the server's network interface.

Database and application servers introduce further dependencies. Verify the front-end service and its required back-end connections separately. The network test should follow the application's actual communication path rather than stop at the first server that answers ping.

Email servers and cloud email services

Email may be hosted internally or provided through a service such as Microsoft 365. Establish the deployment and access path rather than assuming that organization size determines where email runs. For local services, investigate the internal listener and policies; for hosted services, include external DNS and the permitted internet path. A general internet test does not prove that a particular mail service is reachable.

VPN server placement and port forwarding

A VPN can terminate on a firewall or on a separate Windows, Linux, or appliance-based gateway. Where the VPN server sits behind a translating firewall, the relevant traffic must reach the correct internal endpoint through the intended forwarding and security configuration. Reuse the protocol-specific checks from *Chapter 9* instead of opening unrelated ports.

Successful tunnel establishment is only the first part of the path. Internal routes, return routes, DNS, and access policies must also support communication from the VPN client range to the requested resource. Keep the tunnel endpoint distinct from the application server when comparing logs.

Wireless access points and controller services

Wireless clients reach the wired infrastructure through access points. Centralized management or a controller can provide another place to inspect AP and client state, depending on the deployment. Use the wireless checks from *Chapter 8* to distinguish a radio or association problem from a VLAN or upstream forwarding failure. Then follow the client's assigned network into the same switching and routing infrastructure used by wired endpoints.

Network printers and office placement

Network printing can involve a client, a print server, and the printer itself, potentially on different VLANs or in different offices. Establish which leg fails. A client reaching the print server does not prove that the server can contact the printer, and a reachable printer does not prove that the queue or driver is functioning.

Printers need predictable addressing and an allowed path from the relevant server or clients. Physical placement serves users, but logical placement determines which routes and policies apply. Reuse *Chapter 7*'s endpoint and printing checks after confirming the network path rather than repeating them for an unrelated subnet fault.

VoIP phone systems and network integration

VoIP phones also need addressing and access to their required services. Some deployments use separate cabling; others carry voice and data through shared switching infrastructure with distinct VLANs. Verify the phone's intended VLAN, addressing, gateway, and uplink before investigating call quality. When calls connect but sound poor, use the performance and prioritization checks introduced earlier rather than assuming that a completed connection proves sufficient capacity.

Remote office connectivity: SD-WAN, MPLS, and VPN

A remote office can connect through an MPLS service, an internet-based VPN, or a design managed through SD-WAN. These describe different aspects of connectivity: the provider service, protected tunnel, and path-selection or management layer should not be treated as interchangeable physical media. Document the actual underlay and the path selected for the failing destination.

For a site-to-site incident, compare local connectivity at each office before investigating the inter-site connection. Check the routes to the remote subnet and the return routes, then any tunnel or policy that carries the traffic. The transport may use fiber, cable, or another provider circuit, but a working circuit does not prove that the remote subnet is reachable.

A Layer 2 service can extend a broadcast domain between locations. That may be intentional, but it should not happen accidentally because the site boundaries were left undefined. A routed design separates the sites' subnets and makes forwarding decisions explicit. Establish which design is in use before changing a gateway or assuming that both offices share the same LAN.

Cloud server integration

Cloud resources introduce additional routing and policy boundaries. Azure and AWS connectivity is not inherently a proprietary VPN: supported site-to-site VPN designs can use standards-based IPsec. Distinguish a private routed connection from access to a deliberately published public endpoint. Do not use broad port forwarding as a substitute for the intended security design. Verify the cloud-side return route and access rules as well as the local firewall.

Internet of Things devices in modern networks

Time-recording clocks, music players, and other IoT devices add endpoints with their own service dependencies. Do not assume that a newer device is automatically more secure or that it needs unrestricted access to the rest of the network. Identify its intended subnet and the specific destinations it must reach.

Use segmentation and access policy to separate these devices where the design requires it. When an IoT application fails, check the device's local connectivity and gateway, then the inter-VLAN or external path. The same troubleshooting method applies, even when the endpoint offers fewer diagnostic tools than a Windows or Linux computer.

You now have a map of the main places where enterprise traffic can originate, terminate, or cross another network boundary. That map becomes useful only when you can follow one connection through it. The two-VLAN lab that follows makes the forwarding path explicit, giving you a known route against which each later troubleshooting test can be compared.

Working through an inter-VLAN connection

The following lab uses two subnets: Laptop 1 is on VLAN 1 and Laptop 2 is on VLAN 2. *Table 11.1* sets out the addressing relationship needed to follow the procedure. VLAN 1 is used here to match the lab, not as a general recommendation for production user-network design.

Component	VLAN/ network	Address or next hop	Role in this example
Laptop 1	VLAN 1	192.168.1.4/24	Source; existing default gateway is 192.168.1.1.
Core VLAN 1 interface	192.168.1.0/24	192.168.1.254	Next hop for Laptop 1's route to VLAN 2.
Core VLAN 2 interface	192.168.2.0/24	192.168.2.254	Router interface and Laptop 2's gateway.
Laptop 2	VLAN 2	192.168.2.4/24	Destination; returns traffic through 192.168.2.254.

Table 11.1: Addressing and forwarding roles in the inter-VLAN lab

The intended exchange is Laptop 1 to the core's VLAN 1 interface, through the core's routing function, and out VLAN 2 to Laptop 2. Replies traverse the corresponding return path. Layer 2 switches may carry either side of this exchange without routing it themselves.

On Laptop 1, verify the configuration:

```
ipconfig
```

Confirm 192.168.1.4, mask 255.255.255.0, and the existing gateway 192.168.1.1. This example uses a separate next hop, 192.168.1.254, for VLAN 2. If the configured default gateway already routes correctly to that subnet, an additional static route is unnecessary. Otherwise, the lab's persistent route can be added from an elevated Command Prompt:

```
route -p add 192.168.2.0 mask 255.255.255.0 192.168.1.254
```

This more-specific route directs traffic for 192.168.2.0/24 to the core without changing the host's default path. Verify it instead of adding duplicate entries:

```
route print
```

Check both the active route and its persistence. The next hop must be reachable on the source network. A route on Laptop 1 does not create a VLAN, enable routing on the core, or supply a return route.

Checking the Layer 3 switch

Connect through an authorized management session, using a console connection when necessary. Inspect the current configuration before changing it:

```
show running-config
show ip interface brief
show ip route
```

Confirm that IP routing is enabled and that the VLAN interfaces have the intended addresses. A minimal Cisco-style lab configuration illustrates the required relationship:

```
configure terminal
ip routing
vlan 2
exit
interface Vlan1
 ip address 192.168.1.254 255.255.255.0
 no shutdown
exit
interface Vlan2
 ip address 192.168.2.254 255.255.255.0
 no shutdown
end
```

Do not paste this over an existing production configuration. The switch virtual interfaces must also be operational, which depends on the VLANs and associated Layer 2 connectivity. A configured IP address or successful ping to one interface does not prove that every access port and trunk is correct.

Assigning VLANs to switch ports

On the switch carrying Laptop 2, verify the actual endpoint port. For the lab interface FastEthernet1/0/2, the intended access-port configuration is:

```
configure terminal
interface FastEthernet1/0/2
 switchport mode access
 switchport access vlan 2
 spanning-tree portfast
end
```

PortFast is appropriate here only because the port connects to an endpoint, not another switch. It allows the port to enter forwarding promptly; it does not create the VLAN or enable routing. Do not apply it indiscriminately to inter-switch links. Where an access switch connects to the core through a trunk, verify that the required VLAN is permitted and active on that trunk.

Inspect the resulting configuration with `show running-config`. A newly configured access VLAN affects this port, not every port on the switch. Do not assume that other ports retain a particular VLAN based solely on a manufacturer's defaults; inspect the running configuration.

On Laptop 2, verify `192.168.2.4/24` and gateway `192.168.2.254`. From Laptop 1, test the destination:

```
ping 192.168.2.4
tracert 192.168.2.4
```

When diagnostic replies are permitted, the simple path should pass through the core and reach Laptop 2. Missing ICMP replies require further investigation rather than automatic changes to routing. Test an intended application service as well when validating the actual user requirement.

For VLAN 2 to reach the internet, inspect the core's onward default route and the upstream gateway's return route to `192.168.2.0/24`. In a design where the firewall is on VLAN 1, the return next hop can be `192.168.1.254`. Translation and access policy must also permit the traffic. Working inter-VLAN routing alone does not establish internet access.

Troubleshooting inter-VLAN routing and subnet communication

The earlier walkthrough established the intended configuration. Now revisit the two-subnet arrangement as a diagnostic sequence rather than assuming that every configuration step succeeded. Work from the source toward the destination, and include the reply path:

1. **Verify the source configuration.** Run `ipconfig` on Laptop 1 and compare its address, mask, and gateway with *Table 11.1*. An incorrect mask can change whether the host considers the destination local.
2. **Inspect the selected route.** Use `route print` to check the route to `192.168.2.0/24`. In this lab, it uses `192.168.1.254`; a network whose default gateway already provides the path may not need the additional route.

3. **Check the source VLAN path.** Inspect the endpoint's access VLAN and intervening trunks. Confirm that the source can reach the intended local gateway using the permitted diagnostics.
4. **Check Layer 3 forwarding.** Inspect the core's running configuration, operational VLAN interfaces, and connected routes. Include any access controls between the subnets.
5. **Check the destination attachment and return path.** Confirm Laptop 2's port membership, 192.168.2.4/24 address, and gateway 192.168.2.254. Successful forwarding in one direction is insufficient.
6. **Retest the required service.** Repeat the original application operation after the network checks. If addressing and routing are correct but the service still fails, examine the destination listener and applicable firewalls.

This sequence separates a configuration example from fault isolation. Record the failed test and the evidence behind each change. Repeatedly adding routes without checking the existing path can hide the original problem and introduce a different one.

Troubleshooting Windows Firewall and port connectivity

The final boundary is the service itself. Successful IP or ICMP communication does not prove that a particular TCP port is reachable, and a firewall allowance does not start the application that should answer. Use a remote connection test and a local listener inspection together.

Testing port status with Test-NetConnection

From an authorized Windows client, test SMB's usual TCP port on the lab server:

```
Test-NetConnection -ComputerName li-dc01-2025 -Port 445
```

The relevant success indicator is:

```
TcpTestSucceeded : True
```

A successful TCP test means the connection was established to that port. It does not prove that the user has permission to read a share. To illustrate the distinction between host reachability and a specific service, the following separate test uses port 4450, not 445:

```
Test-NetConnection -ComputerName li-dc01-2025 -Port 4450
```

The failure example reports:

```
WARNING: TCP connect to (192.168.21.92 : 4450) failed

ComputerName      : li-dc01-2025
RemoteAddress     : 192.168.21.92
RemotePort        : 4450
InterfaceAlias    : Ethernet
SourceAddress     : 192.168.21.110
PingSucceeded     : True
PingReplyDetails (RTT) : 0 ms
TcpTestSucceeded  : False
```

Here, PingSucceeded is true while TcpTestSucceeded is false. The host answered the echo test, but the requested TCP connection failed. Possible causes include no listener on that port, a listener bound elsewhere, a host firewall, or filtering along the path. The result does not identify which cause applies; inspect the server next.

Using netstat to identify local listeners

On the destination server, review local socket state:

```
netstat -an | more
```

Figure 11.2 shows the lab's TCP states and UDP endpoints. The utility displays existing endpoints; it does not test every port from 1 through 65,535 or report the effective firewall permissions.

```

Administrator: Window x
TCP [::]:49678 [::]:0 LISTENING
TCP [::]:49687 [::]:0 LISTENING
TCP [::]:49688 [::]:0 LISTENING
TCP [::]:49715 [::]:0 LISTENING
TCP [::]:49739 [::]:0 LISTENING
TCP [::]:51905 [::]:0 LISTENING
TCP [::1]:53 [::1]:0 LISTENING
TCP [::1]:389 [::1]:49679 ESTABLISHED
TCP [::1]:389 [::1]:49680 ESTABLISHED
TCP [::1]:389 [::1]:50850 ESTABLISHED
TCP [::1]:49679 [::1]:389 ESTABLISHED
TCP [::1]:49680 [::1]:389 ESTABLISHED
TCP [::1]:50850 [::1]:389 ESTABLISHED
TCP [::1]:52503 [::1]:47001 TIME_WAIT
TCP [fe80::d742:c8cd:8572:1b07%8]:389 [fe80::d742:c8cd:8572:1b07%8]:50846 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:389 [fe80::d742:c8cd:8572:1b07%8]:50852 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:389 [fe80::d742:c8cd:8572:1b07%8]:50854 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:49190 [fe80::d742:c8cd:8572:1b07%8]:49677 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:49677 [fe80::d742:c8cd:8572:1b07%8]:49190 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:49677 [fe80::d742:c8cd:8572:1b07%8]:52636 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:49677 [fe80::d742:c8cd:8572:1b07%8]:52656 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:50846 [fe80::d742:c8cd:8572:1b07%8]:389 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:50852 [fe80::d742:c8cd:8572:1b07%8]:389 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:50854 [fe80::d742:c8cd:8572:1b07%8]:389 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:52497 [fe80::d742:c8cd:8572:1b07%8]:135 TIME_WAIT
TCP [fe80::d742:c8cd:8572:1b07%8]:52636 [fe80::d742:c8cd:8572:1b07%8]:49677 ESTABLISHED
TCP [fe80::d742:c8cd:8572:1b07%8]:52656 [fe80::d742:c8cd:8572:1b07%8]:49677 ESTABLISHED
UDP 0.0.0.0:123 *: *
UDP 0.0.0.0:389 *: *
UDP 0.0.0.0:3389 *: *
UDP 0.0.0.0:5353 *: *
UDP 0.0.0.0:5355 *: *
UDP 0.0.0.0:49152 *: *
UDP 0.0.0.0:49153 *: *
UDP 0.0.0.0:49154 *: *
-- More --

```

Figure 11.2: Comparing local TCP listeners, connections, and UDP endpoints

A LISTENING row identifies a local TCP socket waiting for requests. It does not prove that Windows Firewall or an upstream firewall allows the remote client to reach it. Conversely, allowing TCP 3389 does not make Remote Desktop available if the service is not enabled or listening. UDP rows do not use TCP's connection states.

The following lab output includes established connections, teardown states, wildcard IPv6 listeners, and loopback traffic:

```

TCP 192.168.21.92:51905 192.168.21.93:52485 ESTABLISHED
TCP 192.168.21.92:52494 192.168.21.93:49678 TIME_WAIT
TCP 192.168.21.92:52500 192.168.21.93:49678 TIME_WAIT
TCP 192.168.21.92:52501 192.168.21.93:49678 ESTABLISHED
TCP 192.168.21.92:59698 192.168.21.93:58329 ESTABLISHED
TCP [::]:80 [::]:0 LISTENING
TCP [::]:88 [::]:0 LISTENING

```

TCP	[::]:135	:::0	LISTENING
TCP	[::]:389	:::0	LISTENING
TCP	[::]:445	:::0	LISTENING
TCP	[::]:464	:::0	LISTENING
TCP	[::]:593	:::0	LISTENING
TCP	[::]:636	:::0	LISTENING
TCP	[::]:3268	:::0	LISTENING
TCP	[::]:3269	:::0	LISTENING
TCP	[::]:3389	:::0	LISTENING
TCP	[::]:5985	:::0	LISTENING
TCP	[::]:9389	:::0	LISTENING
TCP	[::]:47001	:::0	LISTENING
TCP	[::]:49664	:::0	LISTENING
TCP	[::]:49665	:::0	LISTENING
TCP	[::]:49666	:::0	LISTENING
TCP	[::]:49667	:::0	LISTENING
TCP	[::]:49670	:::0	LISTENING
TCP	[::]:49671	:::0	LISTENING
TCP	[::]:49672	:::0	LISTENING
TCP	[::]:49674	:::0	LISTENING
TCP	[::]:49677	:::0	LISTENING
TCP	[::]:49678	:::0	LISTENING
TCP	[::]:49679	:::0	LISTENING
TCP	[::]:49680	:::0	LISTENING
TCP	[::]:49688	:::0	LISTENING
TCP	[::]:49715	:::0	LISTENING
TCP	[::]:49739	:::0	LISTENING
TCP	[::]:51905	:::0	LISTENING
TCP	:::1:53	:::0	LISTENING
TCP	:::1:389	:::1:49679	ESTABLISHED
TCP	:::1:389	:::1:49680	ESTABLISHED

Read an endpoint such as 192.168.21.92:51905 as an address and port together. ESTABLISHED means a TCP session exists, not that data is continuously flowing. TIME_WAIT can be normal after a connection closes. A [::] listener and an [::1] loopback connection have different reachability implications, so do not discard the IPv6 rows.

If the intended listener is absent, investigate the service and its bindings. If it is present, compare the client's destination address and port with that binding, then inspect the applicable firewall profile, rule scope, and upstream policy. For example, a successful local connection alongside a failed remote TCP test narrows the next investigation to the remote-

facing binding or path; it does not automatically prove that the firewall is at fault. Compare another authorized client on the same subnet and, where relevant, a client across the routed boundary. Keep the destination address and port constant. A difference between those tests can help distinguish source-specific policy from a service problem affecting everyone.

Use the targeted change-and-retest approach from *Chapter 10*. Retest from the original client, because a successful local connection bypasses parts of the path that remote users depend on.

Before closing the incident, confirm that the original user operation succeeds and that the change has not opened access beyond the intended scope. Record the verified path, the failed condition, the corrective change, and the successful retest. That record is more useful than simply noting that connectivity was restored.

Summary

This chapter combined topology, addressing, VLAN membership, routing, service dependencies, and firewall policy into one investigation. The inter-VLAN lab demonstrated that both the forward and return paths must work. The port tests demonstrated that a reachable host, a listening service, and an allowed connection are separate conditions. Correlating evidence from the client, network, and destination is what distinguishes those failures.

You began this book by examining how networks communicate and how addresses, protocols, and routes determine where traffic goes. You then applied that foundation to corporate resources, remote users, wireless networks, VPNs, and the tools that make their behavior visible. The final step is to use those pieces together without jumping from a symptom to an unsupported conclusion.

For the next incident, establish the scope, map the expected path, select a test that answers one question, and follow the evidence. Make a controlled change, verify the original user operation, and document the result so the next investigation starts with better information. The network will continue to change; a disciplined method lets you troubleshoot it without starting again from guesswork.

12


Unlock Your Exclusive Benefits

Your copy of this book includes the following exclusive benefits:



DRM-Free PDF Version

Download DRM-free PDF and ePub copies of this book.



7-Day Packt Library Access

Get 7-day unlimited access to 8,000+ books and videos. No credit card required.

Available for first-time Packt+ trial users only.



Next-Gen Reader Access

Read this book on Packt Reader with progress sync, dark mode and note-taking.

Follow the guide below to unlock them. The process takes only a few minutes and needs to be completed once.

Unlock this Book's Free Benefits in 3 Easy Steps

Step 1

Keep your purchase invoice ready for *Step 3*. If you have a physical copy, scan it using your phone and save it as a PDF, JPG, or PNG.

For more help on finding your invoice, visit <https://www.packtpub.com/en-us/unlock?step=1>.

Note

Note: If you bought this book directly from Packt, no invoice is required. After *Step 2*, you can access your exclusive content right away.

Step 2

Scan the QR code or go to [packtpub.com/unlock](https://www.packtpub.com/unlock).



On the page that opens (similar to *Figure 12.1* on desktop), search for this book by name and select the correct edition.

Unlock Your Book's Free Benefits

Bought a Packt book from Amazon or one of our channel partners? Unlock your free benefits in 3 easy steps.

The screenshot shows the 'Unlock Your Book's Free Benefits' landing page. At the top, a progress bar indicates three steps: 'Find Your Book' (active, highlighted with an orange circle), 'Sign Up or Sign In', and 'Upload Purchase Proof'. Below the progress bar, the '1. Find Your Book' section is expanded, showing a search bar with the placeholder text 'Search by title or ISBN'. A purple arrow points to the search bar with the text 'Search for your book here'. Below the search bar, there are two sections: '2. Sign up (Free) or Sign In' and '3. Upload Purchase Proof', each with a dropdown arrow. A 'Need Help?' link is located in the top right corner of the page.

Figure 12.1: Packt unlock landing page on desktop

Step 3

After selecting your book, sign in to your Packt account or create one for free. Then upload your invoice (PDF, PNG, or JPG, up to 10 MB). Follow the on-screen instructions to finish the process.

Need Help

If you get stuck and need help, visit <https://www.packtpub.com/unlock-benefits/help> for a detailed FAQ on how to find your invoices and more. This QR code will take you to the help page.



Note

Note: If you are still facing issues, reach out to customer@packt.com.



packtpub.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.


Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

At www.packtpub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

Other Books You May Enjoy

If you enjoyed this book, you may be interested in these other books by Packt:



AI Networking Cookbook

Eric Chou

ISBN: 978-1-80580-798-8

- Understand the AI LLM landscape and key parameters for networking tasks
- Create OpenAI-enabled scripts for daily network engineering workflows
- Master prompt engineering techniques for improved AI outputs
- Build local LLMs using Ollama for network applications
- Chain language models with LangChain for complex network solutions
- Develop AI application frontends using the Streamlit framework
- Design robust backends for network AI applications
- Build an end-to-end network copilot by integrating all the techniques you've learned



Building AI Agents for Network Operations

Sif Baksh

ISBN: 978-1-80834-683-5

- Run local LLM workflows with Ollama and Python
- Shape reliable NetOps prompts using RACE
- Parse CLI and BGP output into structured JSON
- Build chatbots that remember troubleshooting context
- Connect AI agents to approved network tools
- Create evidence-based troubleshooting loops
- Expose reusable network tools with MCP
- Plan read-only pilots with safety controls

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit authors.packt.com and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Share your thoughts

Now you've finished *The Practical Network Engineer's Guide*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback or leave a review on the site that you purchased it from.



<https://packt.link/r/1808495276>

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Index

A

ARP cache	
resetting	131
ARPANET	14
Address Resolution Protocol (ARP)	36
for network diagnostics	163, 164
Automatic Private IP Addressing (APIPA)	67
access control list (ACL)	81
address scope, in IPv4 and IPv6	
link-local IPv6 addresses	63, 64
private IPv4 ranges and NAT	62, 63
unique local address structure, reading	64 – 66
automatic address assignment	66, 67

B

Border Gateway Protocol (BGP)	
between autonomous systems	78
bandwidth utilization	
managing	136
built-in wireless access point diagnostic tools	
using	125, 126

C

Classless Inter-Domain Routing (CIDR)	46
captured traffic	
scan, comparing with	176 – 178
client-side VPN configuration	
in Windows 11	149, 150
cloud server integration	186
core Layer 3 switch	182

corporate resource access	
troubleshooting	100
video and audio conferencing quality issues	100 – 102

D

DHCPv6	66
DNS	93
Datagram Transport Layer Security (DTLS)	177
Device Discovery	125
data, travels over IP	16
data preparation	16, 17
local device, to destination	17, 18
tracert result	20
tracert, using to route	19
dig	
DNS diagnostics	171, 172
dropped connections	
addressing, in wired and wireless networks	120 – 122
dynamic routes	
versus static routes	77

E

Enhanced Interior Gateway Routing Protocol (EIGRP)	78
encapsulation	16
endpoint	6
enterprise wireless networks	133
error 0x80072746	152
error 13801	152
error 789	152
errors 609 and 633	152
errors 800 and 809	151
external switch	87, 88

F**Function Discovery
Provider Host** 120**Function Discovery
Resource Publication** 120**firewall placement and
internet connectivity** 183, 184**firewalls**

- application proxy 81
- protocol inspection 82
- stateful inspection 82 – 84
- stateless access control list
(ACL) 81
- traffic, filtering with 80

flexible subnets 44

- CIDR 46
- classful address ranges,
reading 44, 45
- mask and gateway,
checking 49
- network broadcast and
host range, calculating 48, 49
- prefix 46, 47

G**Get-NetTCPConnection**
TCP connections, 169
inspecting**global unicast address
(GUA)** 63**guest network**
configuration 124, 125
isolation 124, 125**H****Hide NAT** 50**Hyper-V**
virtual network paths, 85, 86
building with**Hyper-V switch types**
external switch 87
internal switch 87
private switch 87**hextets** 59**high availability** 4**high-density wireless environments**

- challenges 138
- line of sight and device 138
- placement

home network

- internet connectivity, 126, 127
- troubleshooting from
issues 112

I**IPv4 addressing and
binary representation** 30

- hexadecimal, in networking 32, 33
- octet between binary and
decimal, converting 30, 31

IPv4 settings

- checking, in macOS 90, 91
- interface

IPv6

- address scope, recognizing 62
- addresses and prefixes,
reading 59
- used, for expanding 58
- addressing model

IPv6 addresses and prefixes

- binary bits, to hexadecimal 59, 60
- hextets
- prefix, used to make 62
- routing decision
- shortening 60, 61

IPv6 configuration

- reading, on Windows 68 – 70

ISP

- troubleshooting steps 127 – 130

**Institute of Electrical and
Electronics Engineers
(IEEE)** 14**Intense scan plus UDP
profile** 177**Internet** 13

- data, travels over IP 16
- development 14
- endpoints services and
resilience 13
- Institute of Electrical and
Electronics Engineers 14

Internet Corporation for Assigned Names and Numbers	14
Internet Engineering Task Force	14
responsibilities	15, 16
Internet Assigned Numbers Authority	14
Internet Corporation for Assigned Names and Numbers (ICANN)	14
Internet Engineering Task Force (IETF)	14
Internet Key Exchange version 2 (IKEv2)	142
Internet Protocol	9
Internet of Things devices	186
ifconfig	
network configuration,	170
viewing	
inter-VLAN connection	186, 187
Layer 3 switch, checking	188
troubleshooting	189, 190
VLANs, assigning to switch ports	188, 189
inter-VLAN routing	182
internal switch	87, 88
internet connectivity	
troubleshooting, from home network	126, 127
ip route	
routing tables, managing	170
ipconfig command	35
IP configuration and management	158 – 160

L

Layer 2 Tunneling Protocol (L2TP)	142
Layer 2 switch	182
Lightweight Directory Access Protocol (LDAP)	161
local listeners	
identifying, with netstat	191 – 194
loopback	44

M

Maximum Transmission Unit (MTU) problem	36
Media Access Control (MAC) address	26
macOS	
routes and interfaces, managing	89
macOS interface	
IPv4 settings, checking	90, 91
modern networks	12
access, without	12
unnecessary complexity	
data, at manageable cost	13
Internet of Things devices	186
performance reliability and automation	12
security	12
multicast	45
multilayer switch configuration	
reading, with VLAN	76, 77

N

NSLookup	125
Network Address Translation (NAT)	50, 58
Network Level Authentication (NLA)	113
Network Policy Server (NPS)	145
netstat	
network analysis	172
network connections, viewing	160, 161
using, to identify local listeners	191 – 194
network	4, 5
endpoints servers	6
infrastructure	6, 7
Internet Protocol	9
network stacks	11
operating systems	11
reading, as path	7
requirements	6
transport behavior and security	11

transport protocols	9, 10	captured packet, reading	29
wired transmission	8, 9	from outside	
wireless transmission	8, 9	data link layer	27
network boundaries		network layer	27
crossing, with NAT and PAT	50	physical layer	26
network connectivity and routing tools	33, 34	presentation layer	28
address resolution,	36	protocol layers, with	28
examining with arp		Wireshark	
configuration, inspecting	35	session layer	28
with ipconfig		shared layers	25
local routing table, reading	38	transport layer	27, 28
path, analyzing with	37	organizational environments	
pathping and tracert		network issues	97
persistent static route,	39		
adding		P	
reachability, testing with	35	Port Address Translation (PAT)	53
ping tool		PortFast	189
route print output, reading	38, 39	PowerShell	
network issues		names, resolving	168
in organizational	97	static routes, managing	79
environments		virtual adapters,	87
slowness and DNS failures	98, 99	controlling with	
network printers and office placement	185	pathping	37, 129
network severity		path analysis	174
assessing, criteria	96, 97	ping switches	173, 174
triage and stakeholder	96	ping tool	125
involvement		connectivity, testing	171
network troubleshooting		reachability, testing with	35
DNS and gateway	104 – 106	port status	
configuration		testing, with Test-NetConnection	190, 191
IP structure and addressing	102, 103	private IP network	13
subnet mask consistency	104	private switch	87, 88
subnetting	104	public address	
nmap		sharing, with Hide NAT	50, 51
port scanning	172	Q	
nslookup		Quality of Service (QoS)	137
DNS troubleshooting	162, 163	R	
O		Routing Information Protocol (RIP)	78
Open Shortest Path First (OSPF)	78	remote desktop	
Open Systems Interconnection (OSI) model	17, 23 – 25	troubleshooting	112 – 118
application layer	28	remote local network discovery	

troubleshooting	119, 120	versus dynamic routes	77
remote network		storage area network (SAN)	184
issues	112	subnet boundaries	
remote office connectivity	185	destination, testing	43, 44
remote printing		locating, with binary	42
troubleshooting	118 – 120	prefix lengths, reading as subnet masks	42, 43
route print		subnet communication	
routing tables, viewing	175	troubleshooting	189, 190
router		switched virtual interfaces (SVIs)	75
troubleshooting steps	127 – 130		
router hardware checks		T	
resetting	131	Test-NetConnection	
routes		connectivity, testing	165 – 167
reading	92	network paths, tracing	167, 168
		port status, testing	190, 191
S		Transmission Control Protocol (TCP)	9
Secure Socket Tunneling Protocol (SSTP)	142	Transport Layer Security	11
Server Message Block	177	targeted firewall change	
Session Initiation Protocol (SIP)	177	validating	178
Speedtest	125	tracert	37, 129
Stateless Address Autoconfiguration (SLAAC)	62	path analysis	174
scan		U	
comparing, with captured traffic	176, 177	User Datagram Protocol (UDP)	9
server		unique local address (ULA)	63
publishing, with static NAT	52, 53	user experience	
server roles, in network infrastructure	184	managing	136
email servers and cloud	184		
email services		V	
port forwarding	184	VLAN segmentation	183
VPN server placement	184	VLANS	
wireless access points and controller services	185	multilayer switch	75 – 77
services		configuration, reading	
directing, with PAT and	53, 54	used, for routing traffic	74, 75
port forwarding		used, for segmenting	74, 75
slow or unstable wireless connectivity		VPN connectivity	
common causes	134 – 136	DNS configuration	150, 151
static NAT	52	name resolution	150, 151
static routes			
managing, with PowerShell	79		

troubleshooting	112 – 118	Wi-Fi Multimedia (WMM)	137
VPN error codes		Wi-Fi interference	
troubleshooting	151, 152	from appliances and	138
VPN ports		electronic devices	
issue, investigating	148	in enterprise environments	137
VPN protocols	142, 143	Windows	
suitability, in public	144	IPv6 configuration, reading	68, 69
networks		Windows 11	
VoIP phone systems and network integration	185	client-side VPN	149, 150
virtual adapters		configuration	
controlling, with	87	Windows Firewall and port connectivity	
PowerShell		troubleshooting	190
virtual network paths		Windows Server	
building, with Hyper-V	85, 86	configuring, for VPN and	147
virtual private network (VPN)		remote access	
connectivity	144 – 147	VPN connectivity	144 – 147
considerations, in		considerations	
Windows Server		Wireshark	
virtual private network (VPN)	141	relevant traffic, capturing	175, 176
W		World Wide Web	13
WAN miniport		wireless access point	
issue, investigating	148	channel selection	122 – 124
		configuration	122 – 124
		wireless network security	
		concerns	134
		threats	134

