LEARNING MADE EASY

9th Edition

Networking

ALL-IN-ONE

for dummies®

A Wiley Brand



10 Books in one!

Doug Lowe

Bestselling author of more than 30 *For Dummies* titles



Networking

ALL-IN-ONE

9th Edition

by Doug Lowe

for dummies®
A Wiley Brand

Networking All-in-One For Dummies®, 9th Edition

Published by: **John Wiley & Sons, Inc.,** 111 River Street, Hoboken, NJ 07030-5774, www.wiley.com

Copyright © 2025 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial technologies or similar technologies.

Media and software compilation copyright © 2025 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial technologies or similar technologies.

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Trademarks: Wiley, For Dummies, the Dummies Man logo, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For technical support, please visit https://hub.wiley.com/community/support/dummies.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media that is not included in the version you purchased, you may download this material at http://booksupport.wiley.com. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2024948598

ISBN 978-1-394-27838-1 (pbk); ISBN 978-1-394-27840-4 (ebk); ISBN 978-1-394-27839-8 (ebk)

Contents at a Glance

Introduction 1

Book 1: Getting Started with Networking 5

CHAPTER 1: Welcome to Networking 7
CHAPTER 2: Network Infrastructure 23
CHAPTER 3: Switches, Routers, and VLANs 41
CHAPTER 4: Cybersecurity 57
CHAPTER 5: Servers and Virtualization 69
CHAPTER 6: Cloud Computing 81

Book 2: Understanding Network Protocols 91

CHAPTER 1: Network Protocols and Standards 93
CHAPTER 2: TCP/IP and the Internet 115
CHAPTER 3: IP Addresses 123
CHAPTER 4: Routing 145
CHAPTER 5: DHCP 155
CHAPTER 6: DNS 173
CHAPTER 7: TCP/IP Tools and Commands 207

Book 3: Planning a Network 231

CHAPTER 1: Local Area Networks 233
CHAPTER 2: Wide Area Networks 249
CHAPTER 3: Server Architecture 261
CHAPTER 4: Virtualization Architecture 271
CHAPTER 5: Storage Architecture 283
CHAPTER 6: Backup Architecture 295
CHAPTER 7: Hyperconverged Infrastructure 313

Book 4: Implementing a Network 325

CHAPTER 1: Network Hardware 327
CHAPTER 2: Wireless Networks 339
CHAPTER 3: Windows Clients 357
CHAPTER 4: Mac Networking 367
CHAPTER 5: Network Printers 377
CHAPTER 6: Virtual Private Networks 385

Book 5: Implementing Virtualization 391

CHAPTER 1: Hyper-V 393
CHAPTER 2: VMware 413

CHAPTER 3: Azure 425
CHAPTER 4: Amazon Web Services 441
CHAPTER 5: Desktop Virtualization 459

Book 6: Implementing Windows Server 2025 467
CHAPTER 1: Installing Windows Server 2025 469
CHAPTER 2: Configuring Windows Server 2025 487
CHAPTER 3: Configuring Active Directory 497
CHAPTER 4: Configuring User Accounts 507
CHAPTER 5: Configuring a File Server 529
CHAPTER 6: Using Group Policy 543
CHAPTER 7: Comandeering Windows Commands 555
CHAPTER 8: Using PowerShell 583

Book 7: Administering Microsoft 365 603
CHAPTER 1: Getting Started with Microsoft 365 Administration 605
CHAPTER 2: Configuring Exchange Online 625
CHAPTER 3: Administering Teams 641

Book 8: Implementing Linux 657
CHAPTER 1: Installing a Linux Server 659
CHAPTER 2: Linux Administration 673
CHAPTER 3: Basic Linux Network Configuration 705
CHAPTER 4: Running DHCP and DNS 717
CHAPTER 5: Linux Commands 725

Book 9: Managing a Network 755
CHAPTER 1: Welcome to Network Administration 757
CHAPTER 2: Managing Remotely 771
CHAPTER 3: Managing Network Assets 791
CHAPTER 4: Solving Network Problems 801

Book 10: Dealing with Cybersecurity 815
CHAPTER 1: Securing Your Users 817
CHAPTER 2: Managing Firewalls and Virus Protection 827
CHAPTER 3: Dealing with Spam 847
CHAPTER 4: Managing Disaster Recovery and Business Continuity Planning 861
CHAPTER 5: Planning for Cybersecurity Incident Response 869
CHAPTER 6: Penetration Testing 883

Index 895

Table of Contents

INTRODUCTION 1
About This Book 2
Foolish Assumptions 3
Icons Used in This Book 3
Beyond the Book 4
Where to Go from Here 4

BOOK 1: GETTING STARTED WITH NETWORKING 5

CHAPTER 1: Welcome to Networking 7
Defining a Network 8
Delving into Network Lingo 10
Seeing What You Can Do with a Network 10
Accessing the internet 11
Sharing files 11
Sharing resources 11
Sharing programs 12
Getting Acquainted with Servers and Clients 13
Weighing Your Options: Dedicated Servers versus Peer-to-Peer Networks 14
Understanding What Makes a Network Tick 15
Putting the Pieces Together 17
Considering Network Size 18
Recognizing That Your Personal Computer Isn't Personal When It's on a Network 19
Assigning a Network Administrator 21

CHAPTER 2: Network Infrastructure 23
Introducing Infrastructure 24
Understanding Network Protocols and Standards 25
Recognizing Network Topology 26
Bus topology 26
Star topology 27
Ring topology 29
Mesh topology 30
Considering Cable 30
Twisted-pair cable 30
RJ45 connectors 31
Patch panels and patch cables 31
Repeaters and hubs 32
Switches 32

Perusing Ports, Interfaces, and MAC Addresses 33
Pondering Packets 35
Contemplating Collisions 37
Dealing with Broadcast Packets 38
Examining Wireless Networks 39

CHAPTER 3: **Switches, Routers, and VLANs** 41
Understanding Switches 42
Learning 45
Forwarding 46
Flooding 47
Looking Deeper into Switches 48
Collision domains 48
Bridging 49
SFP ports and uplinks 50
Broadcast domains 51
Managed and unmanaged switches 51
Understanding Routers 52
Network address translation 54
Virtual private networks 54
Understanding VLANs 55

CHAPTER 4: **Cybersecurity** 57
But We're a Small Business — Do We Need Security? 58
The Two Pillars of Cybersecurity 59
Prevention 59
Recovery 61
Cybersecurity Frameworks 63
The NIST Cybersecurity Framework 64

CHAPTER 5: **Servers and Virtualization** 69
Understanding Network Operating Systems 69
Network services 70
File-sharing services 70
Multitasking 71
Directory services 72
Security services 73
Knowing What's Important in a Server 74
Scalability 74
Reliability 74
Availability 74
Service and support 74
Identifying the Components of a Server Computer 75
Motherboard 75
Processor 75

Memory 76
Hard drives 76
Network interfaces 77
Video 77
Power supply 77
Considering Server Form Factors 77
Tower cases 78
Rack-mounted servers 78
Blade servers 78
Tiny servers 79
Understanding Virtualization 79

CHAPTER 6: **Cloud Computing** 81
Introducing Cloud Computing 81
Looking at the Benefits of Cloud Computing 82
Cost 82
Scalability 83
Reliability 83
Accessibility 83
Free of hassles 84
Detailing the Drawbacks of Cloud Computing 84
Entrenched applications 84
Internet connection speed 84
Internet connection reliability 85
Security threats 85
Examining Three Basic Kinds of Cloud Services 85
Applications 86
Platforms 86
Infrastructure 87
Public Clouds versus Private Clouds 87
Introducing Some of the Major Cloud Providers 88
Amazon 88
Google 88
Microsoft 89
Getting into the Cloud 89

BOOK 2: UNDERSTANDING NETWORK PROTOCOLS 91

CHAPTER 1: **Network Protocols and Standards** 93
Understanding Protocols 93
Understanding Standards 94
Seeing the Seven Layers of the OSI Reference Model 95
The physical layer 96
The data link layer 98
The network layer 100

The transport layer 102
The session layer 103
The presentation layer 104
The application layer 105
Following a Packet through the Layers 105
The Ethernet Protocol 106
Standard Ethernet 108
Fast Ethernet 109
Gigabit Ethernet 109
Beyond gigabit 109
The TCP/IP Protocol Suite 110
IP 111
TCP 112
UDP 112
Other Protocols Worth Knowing About 114

CHAPTER 2: **TCP/IP and the Internet** 115
What Is the Internet? 116
A Little Internet History 117
TCP/IP Standards and RFCs 118
The TCP/IP Protocol Framework 120
Network interface layer 121
Network layer 121
Transport layer 122
Application layer 122

CHAPTER 3: **IP Addresses** 123
Understanding Binary 123
Counting by ones 123
Doing the logic thing 126
Working with the binary Windows Calculator 127
Introducing IP Addresses 128
Networks and hosts 129
The dotted-decimal dance 129
Classifying IP Addresses 130
Class A addresses 133
Class B addresses 133
Class C addresses 134
Subnetting 134
Subnets 135
Subnet masks 136
Network prefix notation 137
Default subnets 137
The great subnet roundup 138

IP block parties 139
Private and public addresses 140
Pondering Ports 141
Understanding Network Address Translation 141

CHAPTER 4: **Routing** 145
Considering the Usefulness of Routers 146
Connecting to the internet 146
Connecting remote locations 148
Splitting up large networks 149
Understanding Routing Tables 150

CHAPTER 5: **DHCP** 155
Understanding DHCP 155
Configuration information provided by DHCP 156
DHCP servers 156
How DHCP actually works 157
Understanding Scopes 158
Scopes, subnets, and VLANs 159
Feeling excluded? 160
Reservations suggested 161
How long to lease? 161
Working with a DHCP Server 162
Installing a Windows Server 2025 DHCP server 162
Configuring a new scope 163
How to Configure a Windows DHCP Client 169
Automatic private IP addressing 171
Renewing and releasing leases 171

CHAPTER 6: **DNS** 173
Understanding DNS Names 173
Domains and domain names 174
Fully qualified domain names 176
Top-Level Domains 177
Generic domains 178
Country code domains 179
The Hosts File 180
Understanding DNS Servers and Zones 183
Zones 184
Primary and secondary servers 186
Root servers 187
Caching 190
Understanding DNS Queries 190

Zone Files and Resource Records 192
SOA records 194
NS records 195
A records 195
CNAME records 196
PTR records 197
MX records 197
Reverse Lookup Zones 198
Working with the Windows DNS Server 199
Creating a new zone 200
Creating a new host record 203
How to Configure a Windows DNS Client 204

CHAPTER 7: **TCP/IP Tools and Commands** 207
Using the arp Command 207
Using the hostname Command 208
Using the ipconfig Command 209
Displaying basic IP configuration 209
Displaying detailed configuration information 210
Renewing an IP lease 211
Releasing an IP lease 211
Flushing the local DNS cache 212
Using the nbtstat Command 212
Using the netstat Command 213
Displaying connections 213
Displaying interface statistics 214
Using the nslookup Command 215
Looking up an IP address 215
Using nslookup subcommands 215
Displaying DNS records 217
Locating the mail server for an email address 218
Using the pathping Command 219
Using the ping Command 221
Using the route Command 222
Displaying the routing table 222
Modifying the routing table 225
Using the tracert Command 226

BOOK 3: PLANNING A NETWORK 231

CHAPTER 1: **Local Area Networks** 233
Making a Network Plan 233
Being Purposeful 234
Taking Stock 235
What you need to know 235
Programs that gather information for you 237

Considering Cable 238
Surmising Switches 240
Planning the Network Topology 242
Planning the TCP/IP Implementation 244
Drawing Diagrams 245
More Questions Your Network Plan Should Address 246

CHAPTER 2: **Wide Area Networks** 249
Connecting to the Internet 250
Connecting with cable or DSL 250
Connecting with T1 lines 251
Connecting with fiber 252
Connecting with a cellular network 253
Choosing a Router 253
Choosing a small office router 254
Choosing an enterprise router 255
Choosing a cellular router 255
Securing Your Connection with a Firewall 256
Providing Redundancy for Your Internet Connection 257
Securing Connections to Remote Locations and Remote Users 258
Connecting Remote Offices with an Ethernet Private Line 260

CHAPTER 3: **Server Architecture** 261
Deciding How Many Servers You Need 261
Deciding Which Servers You Need 262
Domain controllers 262
DHCP servers 263
Mail servers 263
File servers 264
Print servers 265
Web servers 266
Database servers 266
Application servers 266
Backup servers 267
Deployment servers 267
Update servers 267
Virtualization management platform 268
Connecting Your Servers 268

CHAPTER 4: **Virtualization Architecture** 271
Understanding Virtualization 272
Understanding Hypervisors 273
Understanding Virtual Disks 275
Understanding Network Virtualization 277

Considering the Benefits of Virtualization 278
Choosing Virtualization Hosts 280
Understanding Windows Server 2025 Licensing 281

CHAPTER 5: **Storage Architecture** 283
Planning Disk Capacity 283
Considering Disk Drive Types 285
Hard disk drives 285
Solid state drives 285
Considering Drive Interfaces 286
SATA 287
SAS 287
Considering RAID 288
RAID 10 288
RAID 5 289
RAID 6 291
Considering Attachment Types 291
Direct attached storage 291
Storage area networks 292
Network-attached storage 293

CHAPTER 6: **Backup Architecture** 295
Backup Basics 296
Considering Three Basic Types of Backup 297
Where to Back Up Your Data 298
Backing Up to Tape 299
Looking closer at LTO 299
Hardware for tape backup 300
A word about tape reliability 301
About cleaning the heads 302
Backing Up to NAS 303
Using a Backup Appliance 303
Understanding File-Based Backup 304
Full backups 306
Copy backups 307
Daily backups 308
Incremental backups 308
Differential backups 309
Understanding Image-Based Backups and Virtualization 310
Backup Security 310

CHAPTER 7: **Hyperconverged Infrastructure** 313
Considering the Headaches of Traditional IT Architecture 314
Defining Hyperconverged Infrastructure 315
Discerning Deduplication 317

Understanding How Deduplication Works 318
Considering Backup 320
Digging into HCI Clusters 322
Incorporating HCI Into Your Plan 323

BOOK 4: IMPLEMENTING A NETWORK 325

CHAPTER 1: Network Hardware 327
Working with Cable 327
Cable categories 327
What's with the pairs? 328
To shield or not to shield 329
When to use plenum cable 329
Sometimes solid, sometimes stranded 330
Installation guidelines 330
Getting the tools that you need 331
Pinouts for twisted-pair cables 332
Attaching RJ-45 connectors 333
Wall jacks and patch panels 335
Server rooms and distribution frames 336
Installing Switches 337

CHAPTER 2: Wireless Networks 339
Installing a Wireless Access Point 339
Configuring a Wireless Access Point 340
Connecting to a Wireless Network 342
Paying Attention to Wireless Network Security 343
Understanding wireless security threats 344
Securing your wireless network 348
Troubleshooting a wireless network 352

CHAPTER 3: Windows Clients 357
Configuring Network Connections 357
Joining a Domain 363

CHAPTER 4: Mac Networking 367
Basic Mac Network Settings 368
Joining a Domain 372
Connecting to a Share 374

CHAPTER 5: Network Printers 377
Configuring Network Printers 377
Adding a network printer 378
Accessing a network printer using a web interface 382

CHAPTER 6: **Virtual Private Networks** 385
Understanding VPN 385
Looking at VPN Security 387
Understanding VPN Servers and Clients 388

BOOK 5: IMPLEMENTING VIRTUALIZATION 391

CHAPTER 1: **Hyper-V** 393
Understanding the Hyper-V Hypervisor 393
Understanding Hyper-V Virtual Disks 394
Enabling Hyper-V 395
Getting Familiar with Hyper-V 396
Creating a Virtual Switch 398
Creating a Virtual Disk 400
Creating a Virtual Machine 404
Installing an Operating System 409

CHAPTER 2: **VMware** 413
Looking at vSphere 414
Getting Started with VMware Workstation Pro 414
Creating a Virtual Machine 416
Installing VMware Tools 423

CHAPTER 3: **Azure** 425
Looking at Azure Services 426
Creating an Azure Account 427
Examining the Azure Portal 428
Creating a Windows Virtual Machine 429
Managing an Azure Virtual Machine 435
Connecting to an Azure Virtual Machine 438

CHAPTER 4: **Amazon Web Services** 441
Looking at What Amazon Web Services Can Do 442
Creating an Amazon Web Services Account 443
Examining the Amazon Web Services Console 444
Creating a Windows Virtual Machine 446
Managing an Amazon Web Services Virtual Machine 454
Connecting to an Amazon Web Services Virtual Machine 455

CHAPTER 5: **Desktop Virtualization** 459
Introducing Desktop Virtualization 459
Considering Two Approaches to Desktop Virtualization 461
Looking at VMware's Horizon View 462
Looking at Citrix XenApp 463

BOOK 6: IMPLEMENTING WINDOWS SERVER 2025 467

CHAPTER 1: Installing Windows Server 2025 469
Planning a Windows Server Installation 469
Checking system requirements 470
Reading the release notes 470
Deciding whether to upgrade or install 470
Considering your licensing options 471
Thinking about multiboot 471
Planning your partitions 472
Deciding your TCP/IP configuration 473
Choosing workgroups or domains 473
Before You Install 474
Backing up 474
Checking the event logs 474
Applying updates 475
Disconnecting UPS devices 475
Running Setup 475
Considering Your Next Steps 481
Adding Server Roles and Features 482

CHAPTER 2: Configuring Windows Server 2025 487
Using the Administrator Account 487
Using Remote Desktop Connection 488
Enabling remote access 488
Connecting remotely 490
Using Microsoft Management Console 491
Working with MMC 492
Taking an overview of the MMC consoles 493
Customizing MMC 495

CHAPTER 3: Configuring Active Directory 497
What Active Directory Does 497
Understanding How Active Directory Is Structured 498
Objects 498
Domains 499
Organizational units 500
Trees 501
Forests 501
Creating a New Domain 502
Creating an Organizational Unit 503

CHAPTER 4: **Configuring User Accounts** 507
Understanding Windows User Accounts 507
Local accounts versus domain accounts 508
User account properties 508
Creating a New User 509
Setting User Properties 512
Changing the user's contact information 513
Setting account options 513
Specifying logon hours 514
Restricting access to certain computers 515
Setting the user's profile information 516
Resetting User Passwords 517
Disabling and Enabling User Accounts 518
Deleting a User 519
Working with Groups 519
Group types 520
Group scope 520
Default groups 521
Creating a group 522
Adding a member to a group 523
Working with User Profiles 525
Types of user profiles 525
Roaming profiles 526
Creating a Logon Script 528

CHAPTER 5: **Configuring a File Server** 529
Understanding Permissions 529
Understanding Shares 531
Considering Best Practices for Setting Up Shares 532
Managing Your File Server 533
Using the New Share Wizard 534
Sharing a folder without the wizard 538
Granting permissions 540

CHAPTER 6: **Using Group Policy** 543
Understanding Group Policy 543
Enabling Group Policy Management on Windows Server 2025 544
Creating Group Policy Objects 545
Filtering Group Policy Objects 552

CHAPTER 7: **Comandeering Windows Commands** 555
Using a Command Window 556
Opening and closing a command window 556
Editing commands 557
Using the Control menu 557

Special Command Tricks 558
Wildcards 558
Chaining commands 559
Redirection and piping 559
Environment variables 560
Batch files 562
The EventCreate Command 563
Net Commands 564
The Net Accounts command 565
The Net Computer command 566
The Net Config command 566
The Net Continue command 567
The Net File command 568
The Net Group command 568
The Net Help command 570
The Net Helpmsg command 570
The Net Localgroup command 571
The Net Pause command 572
The Net Session command 573
The Net Share command 574
The Net Start command 575
The Net Statistics command 575
The Net Stop command 576
The Net Time command 577
The Net Use command 577
The Net User command 579
The Net View command 580
The RunAs Command 581

CHAPTER 8: **Using PowerShell** 583
Using PowerShell 584
Understanding PowerShell Commands 586
Using Cmdlets 587
Using Parameters 587
Getting Help 589
Using Aliases 591
Using the Pipeline 593
Using Providers 597
Using Scripts 598

BOOK 7: ADMINISTERING MICROSOFT 365 603

CHAPTER 1: **Getting Started with Microsoft 365 Administration** 605
Introducing Microsoft 365 606
Considering Microsoft 365 Plans 608

Understanding Tenants 610
Creating an Microsoft 365 Tenant 611
Creating a New User 615
Resetting a User's Password 621
Disabling a User 622

CHAPTER 2: **Configuring Exchange Online** 625
Looking at Exchange Online Recipient Types 625
Examining the Exchange Admin Center 626
Managing Mailboxes 629
Creating an email alias 630
Delegating a mailbox 632
Converting a standard mailbox to a shared mailbox 634
Enabling or disabling mailbox apps 635
Creating a forwarder 636
Creating a Shared Mailbox 637

CHAPTER 3: **Administering Teams** 641
What Is Teams? 642
A Brief Look at How Teams Works 645
Microsoft 365 Group 646
SharePoint 648
OneDrive for Business 650
Using the Teams Admin Center 650
Managing Teams 652

BOOK 8: IMPLEMENTING LINUX 657

CHAPTER 1: **Installing a Linux Server** 659
Planning a Linux Server Installation 659
Checking system requirements 659
Choosing a distribution 660
Going virtual 662
Deciding on your TCP/IP configuration 662
Installing Fedora Server 663

CHAPTER 2: **Linux Administration** 673
On Again, Off Again 673
Logging in 673
Logging out 675
Shutting down 675
Wait, Where's the Desktop? 676
Playing the Shell Game 676
Getting into Virtual Consoles 677
Using a Remote Console 678

Enabling the root User 679
Using the sudo Command 680
Understanding the file system 680
Looking at top-level directories 681
Browsing the file system 682
Using the RPM Package Manager 683
Listing packages 684
Installing packages 686
Removing packages 687
Updating packages 688
Editing Text Files with Vi 689
Starting vi 690
Saving changes and quitting Vi 692
Understanding Vi's operating modes 692
Moving around in a file 693
Inserting text 694
Deleting text 695
Changing text 696
Copying and pasting text 696
Repeating commands 697
Other useful Vi commands 698
Using Cockpit 698
Managing User Accounts 702

CHAPTER 3: **Basic Linux Network Configuration** 705
Using Cockpit to Configure Network Interfaces 705
Working with Network Configuration Files 710
The Network file 710
The interface configuration files 711
The Hosts file 713
The resolv.conf file 714
Displaying Your Network Configuration with the ifconfig Command 714

CHAPTER 4: **Running DHCP and DNS** 717
Running a DHCP Server 717
Installing DHCP 718
Configuring DHCP 718
Starting DHCP 720
Running a DNS Server 720
Installing BIND 720
Editing BIND configuration files 721
named.conf 721
Zone files 723
Restarting BIND 724

CHAPTER 5: **Linux Commands** 725
Command Shell Basics 725
Getting to a shell 726
Editing commands 726
Wildcards 726
Redirection and piping 727
Environment variables 728
Shell scripts 728
Running a command with root-level privileges 729
Directory- and File-Handling Commands 730
The pwd command 730
The cd command 730
The mkdir command 731
The rmdir command 731
The ls command 732
The cp command 733
The rm command 734
The mv command 735
The cat command 735
Commands for Working with Packages and Services 737
The service command 737
The yum and dnf commands 738
Commands for Administering Users 739
The useradd command 739
The usermod command 741
The userdel command 741
The chage command 741
The passwd command 742
The newusers command 742
The groupadd command 743
The groupdel command 743
The gpasswd command 744
Commands for Managing Ownership and Permissions 745
The chown command 745
The chgrp command 746
The chmod command 746
Networking Commands 747
The hostname command 748
The ifconfig command 748
The netstat command 749
The ping command 750
The route command 752
The traceroute command 752

BOOK 9: MANAGING A NETWORK 755

CHAPTER 1: Welcome to Network Administration 757
Knowing What Network Administrators Do 758
Choosing the Part-Time Administrator 759
Establishing Routine Chores 760
Managing Network Users 761
Patching Up Your Operating System and Software 762
Discovering Software Tools for Network Administrators 763
Building a Library 764
Getting Certified 765
CompTIA 766
Microsoft 767
Cisco 767
Gurus Need Gurus, Too 768
Helpful Bluffs and Excuses 769

CHAPTER 2: Managing Remotely 771
Enabling Remote Desktop Connection 772
Connecting Remotely 774
Using Keyboard Shortcuts for Remote Desktop 776
Configuring Remote Desktop Options 777
Setting the Display options 778
Setting the Local Resources options 779
Setting the Experience options 780
Setting the Advanced options 781
Using Remote Assistance 782
Enabling Remote Assistance 783
Inviting Someone to Help You via a Remote Assistance Session 784
Responding to a Remote Assistance Invitation 787

CHAPTER 3: Managing Network Assets 791
Introducing IT Asset Management 792
Why Bother? 793
Getting Organized 793
What to Track 794
Taking Pictures 795
Picking a Number 796
Making Labels 796
Tracking Software 798
Using Asset-Tracking Software 798
Other Sources of Asset-Tracking Information 799

CHAPTER 4: **Solving Network Problems** 801
When Bad Things Happen to Good Computers 802
Fixing Dead Computers 803
Ways to Check a Network Connection 804
A Bunch of Error Messages Just Flew By! 805
Double-Checking Your Network Settings 806
Time to Experiment 806
Who's on First? 807
Restarting a Client Computer 808
Booting in Safe Mode 809
Using System Restore 809
Restarting Network Services 811
Restarting a Network Server 812
Looking at Event Logs 813
Documenting Your Trials and Tribulations 814

BOOK 10: DEALING WITH CYBERSECURITY 815

CHAPTER 1: **Securing Your Users** 817
Knowing the Difference between Authentication and Authorization 818
Following Password Best Practices 818
Securing the Administrator Account 821
Understanding Multifactor Authentication 823
Securing the Human Firewall 824
Establishing cybersecurity policies 824
Training 824
Phish testing 825

CHAPTER 2: **Managing Firewalls and Virus Protection** 827
Firewalls 828
The Many Types of Firewalls 829
Packet filtering 829
Stateful packet inspection (SPI) 831
Circuit-level gateway 832
Application gateway 832
Firewall Best Practices 833
The Built-In Windows Firewall 834
Configuring Windows Defender Firewall with Group Policy 836
Virus Protection 842
What is a virus? 843
Antivirus programs 844
Safe computing 846

CHAPTER 3: **Dealing with Spam** 847
Defining Spam 848
Sampling the Many Flavors of Spam 849
Using Antispam Software 850
Understanding Spam Filters 851
Looking at Three Types of Antispam Software 854
On-premises antispam 854
Antispam appliances 856
Cloud-based antispam services 856
Minimizing Spam 858

CHAPTER 4: **Managing Disaster Recovery and Business Continuity Planning** 861
Assessing Different Types of Disasters 862
Environmental disasters 863
Deliberate disasters 863
Disruption of services 864
Equipment failure 864
Other disasters 865
Analyzing the Impact of a Disaster 865
Developing a Business Continuity Plan 866
Holding a Fire Drill 867

CHAPTER 5: **Planning for Cybersecurity Incident Response** 869
Seeing the Importance of a Cybersecurity Incident Response Plan 870
Preparing Your Cybersecurity Incident Response Plan 872
Assembling Your Response Team 873
Identifying and Reporting a Cybersecurity Incident 874
Triaging Reported Incidents 875
Containing a Cybersecurity Incident 876
Engaging the Eradication Phase 878
Restoring Lost Data and Systems 878
Considering Communication 879
Internal communication 879
Methods of communication 880
Executive leadership communication 880
External communication 881
Closing the Incident 881
Documentation 882
Lessons learned 882

CHAPTER 6: **Penetration Testing** 883
Understanding Ethical Hacking 884
Introducing the Red Team 885
Seeing How Penetration Testing Works 885
Scoping a Penetration Test 888
Establishing Boundaries 889
Examining Tools for Penetration Testing 890
How to set up Kali 890
What you'll find on Kali 891
Looking at one of Kali's tools 892
Knowing What to Expect from a Penetration Test 893

INDEX 895

Introduction

Welcome to the ninth edition of *Networking All-in-One For Dummies*, the one networking book that's designed to replace an entire shelf full of the dull and tedious networking books you'd otherwise have to buy. This book contains all the basic and not-so-basic information you need to know to get a network up and running and to stay on top of the network as it grows, develops problems, and encounters trouble.

If you're just getting started as a network administrator, this book is ideal. As a network administrator, you have to know about a lot of different topics: installing and configuring network hardware and software, planning a network, working with TCP/IP, securing your network, working with mobile devices, virtualizing your servers, backing up your data, managing cloud services, and many others.

You can, and probably eventually will, buy separate books on each of these topics. It won't take long before your bookshelf is bulging with 10,000 or more pages of detailed information about every imaginable nuance of networking. But before you're ready to tackle each of those topics in depth, you need to get a bird's-eye picture. This book is the ideal way to do that.

And if you already own 10,000 pages or more of network information, you may be overwhelmed by the amount of detail and wonder, "Do I really need to read 1,000 pages about BIND to set up a simple DNS server?" or "Do I really need a six-pound book to show me how to install Linux?" Truth is, most 1,000-page networking books have about 100 or so pages of really useful information — the kind you use every day — and about 900 pages of excruciating details that apply mostly to networks at places like NASA and the CIA.

The basic idea of this book is that I've tried to wring out the 100 or so most useful pages of information on nine different networking topics: network basics, building a network, network administration and security, troubleshooting and disaster planning, working with TCP/IP, home networking, wireless networking, Windows server operating systems, and Linux.

So whether you've just been put in charge of your first network or you're a seasoned pro, you've found the right book.

About This Book

Networking All-in-One For Dummies, 9th Edition, is intended to be a reference for all the great things (and maybe a few not-so-great things) that you may need to know when you're setting up and managing a network. You can, of course, buy a huge 1,000-page book on each of the networking topics covered in this book. But then, who would you get to carry them home from the bookstore for you? And where would you find the shelf space to store them? In this book, you get the information you need all conveniently packaged for you in between one set of covers.

This book doesn't pretend to be a comprehensive reference for every detail of these topics. Instead, this book shows you how to get up and running fast so that you have more time to do the things you really want to do. Designed using the easy-to-follow *For Dummies* format, this book helps you get the information you need without laboring to find it.

Networking All-in-One For Dummies, 9th Edition, is a big book made up of ten smaller books — minibooks, if you will. Each of these minibooks covers the basics of one key element of network management, such as setting up network hardware, installing a network operating system, or troubleshooting network problems. Whenever one big thing is made up of several smaller things, confusion is always a possibility. That's why *Networking All-in-One For Dummies,* 9th Edition, is designed to have multiple access points (I hear an acronym coming on — MAP!) to help you find what you want. At the beginning of the book is a detailed table of contents that covers the entire book. Then each minibook begins with a table of contents that shows you at a glance what chapters are included in that minibook. Useful running heads appear at the top of each page to point out the topic discussed on that page. And handy thumb tabs run down the side of the pages to help you find each minibook quickly. Finally, a comprehensive index lets you find information anywhere in the entire book.

This isn't the kind of book you pick up and read from start to finish, as though it were a cheap novel. (If I ever see you reading it at the beach, I'll kick sand in your face.) This book is more like a reference — the kind of book you can pick up, turn to just about any page, and start reading. You don't have to memorize anything in this book. It's a need-to-know book: You pick it up when you need to know something. Need to know how to set up a DHCP server in Windows? Pick up the book. Need to know how to create a user account in Linux? Pick up the book. Otherwise, put it down, and get on with your life.

Within this book, you may note that some web addresses break across two lines of text. If you're reading this book in print and want to visit one of these web pages, simply key in the web address exactly as it's noted in the text, pretending

as though the line break doesn't exist. If you're reading this as an e-book, you've got it easy — just click the web address to be taken directly to the web page.

Foolish Assumptions

As I was writing this book, I made a few assumptions about you, the reader:

- **You are responsible for or would like to be responsible for a computer network.** The network we speak of may be small — just a few computers, or large — consisting of dozens or even hundreds of computers. The network may already exist, or it may be a network you would like to build. But one way or another, I assume that managing the network is, at least in part, your responsibility.
- **You are an experienced computer user.** You don't need to be an expert, but this book assumes a modest level of experience with computers.
- **You are familiar with Windows.** This book touches on Mac and Linux networks, but the primary focus is on creating and managing networks of Windows computers.

Icons Used in This Book

Like any *For Dummies* book, this book is chock-full of helpful icons that draw your attention to items of particular importance. You find the following icons throughout this book:



TECHNICAL STUFF

Hold it — technical stuff is just around the corner. Read on only if you have your pocket protector.



TIP

Pay special attention to this icon; it lets you know that some particularly useful tidbit is at hand.



REMEMBER

Did I tell you about the memory course I took?



WARNING

Danger, Will Robinson! This icon highlights information that may help you avert disaster.

Beyond the Book

In addition to what you're reading right now, this product also comes with a free access-anywhere Cheat Sheet that includes tables where you can record key network and internet connection information, the RJ-45 pin connections, private IP address ranges, and useful websites for networking information. To get this Cheat Sheet, simply go to `www.dummies.com` and type **Networking All-in-One For Dummies Cheat Sheet** in the search box.

To download a directory of useful websites and a glossary of terms used in this book, head to `www.dummies.com/go/networkingaiofd9e`.

Where to Go from Here

Yes, you can get there from here. With this book in hand, you're ready to plow right through the rugged networking terrain. Browse the table of contents, and decide where you want to start. Be bold! Be courageous! Be adventurous! And above all, have fun!

1 Getting Started with Networking

Contents at a Glance

CHAPTER 1: **Welcome to Networking** 7

CHAPTER 2: **Network Infrastructure** 23

CHAPTER 3: **Switches, Routers, and VLANs** 41

CHAPTER 4: **Cybersecurity** 57

CHAPTER 5: **Servers and Virtualization** 69

CHAPTER 6: **Cloud Computing** 81

IN THIS CHAPTER

- Getting a handle on networks
- Considering why networking is useful (and is everywhere)
- Telling the difference between servers and clients
- Seeing how networks change computing life
- Examining network topology
- Identifying (and offering sympathy to) the network administrator

Chapter 1
Welcome to Networking

Computer networks get a bad rap in the movies. In the 1980s, the *Terminator* movies featured Skynet, a computer network that becomes self-aware (a computer network of the future), takes over the planet, builds deadly terminator robots, and sends them back through time to kill everyone unfortunate enough to have the name Sarah Connor. In the *Matrix* movies, a vast and powerful computer network enslaves humans and keeps them trapped in a simulation of the real world. And in the 2015 blockbuster *Spectre,* James Bond goes rogue (again) to prevent the Evil Genius Ernst Blofeld from taking over the world (again) by linking the computer systems of all the world's intelligence agencies together to form a single, all-powerful evil network that spies on everybody.

Fear not. These bad networks exist only in the dreams of science-fiction writers. Real-world networks are much more calm and predictable. Although sophisticated networks do seem to know a lot about you, they don't think for themselves and they don't evolve into self-awareness. And although they can gather a sometimes disturbing amount of information about you, they won't try to kill you, even if your name is Sarah Connor.

Now that you're over your fear of networks, you're ready to breeze through this chapter. It's a gentle, even superficial, introduction to computer networks, with a slant to the concepts that can help you use a computer that's attached to a network. This chapter goes easy on the details; the detailed stuff comes later.

Defining a Network

A *network* is nothing more than two or more computers connected by a cable or by a wireless radio connection so that they can exchange information.

You can create a simple computer network by using a device called a *switch* to connect all the computers in your network to each other. You do that by stringing a *network cable* from the switch to each computer. The cable plugs into a special jack on the back of each computer; this jack is connected to a *network interface*, which is an electronic circuit that resides inside your computer to facilitate networking. *Voilà!* You have a working network.

If you don't want to mess with cables, you can create a wireless network instead. In a wireless network, the computers use wireless network adapters that communicate via radio signals. All modern laptop computers have built-in wireless network adapters, as do most desktop computers. (If yours doesn't, you can purchase a separate wireless network adapter that plugs into one of the computer's USB ports.) You'll need a device called a *wireless access point* (WAP) to enable the computers to properly connect. In small office or home networks, the WAP is bundled with a device called a *router,* which lets you connect your network to the internet. The combination of a WAP and a router is called a *wireless router.*

Figure 1-1 shows a typical network with five computers. This network is a home network used by a family that bears only a totally coincidental similarity to a famous TV family you may or may not have heard of. You can see that each family member has a device that connects to the network — two of them wirelessly, two of them through cables. There's also a printer that connects wirelessly.

In this example, the wireless router also has a built-in switch that provides several jacks for connecting computers via cable. Most wireless routers include this feature, typically with three to five wired network ports.

Although the network is a small one, it has much in common with larger networks that contain dozens, hundreds, or even thousands of connected computers.

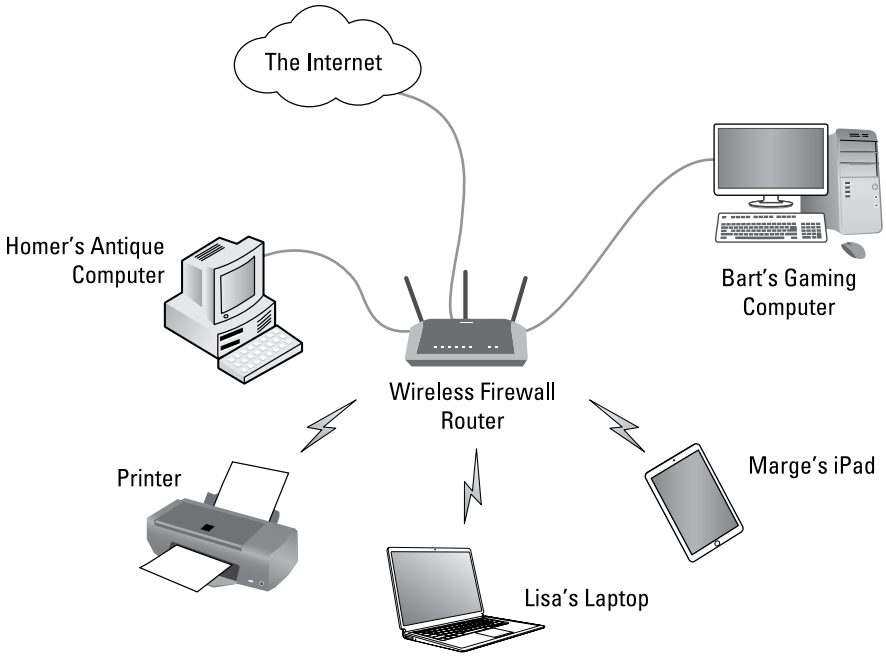


FIGURE 1-1: A typical network.

Here's the rundown for each of the devices connected to this network:

- Lisa has a laptop computer that connects wirelessly. She uses it mostly for school.
- Bart has a fancy gaming computer that's cabled directly to the router.
- Marge doesn't have a full-fledged computer, but she does use an iPad, which is connected wirelessly.
- Homer has an old computer he bought at a garage sale in 1989. He doesn't know how to use it, but he doesn't know that, so no one tells him. Lisa set it up for him and repairs it when Homer breaks it (which happens every few months); she gets repair parts from eBay.
- The printer connects wirelessly to the network and is set up so that any member of the family can print on it.
- The wireless router connects to the internet using the family's cable TV provider. This allows everyone in the family to access the internet.

Delving into Network Lingo

Computer networking has its own strange vocabulary. Although you don't have to know every esoteric networking term, it helps to be acquainted with a few of the basic buzzwords:



- **Local area network (LAN):** Networks are often called LANs, short for *local area network.*

  LAN is the first *three-letter acronym* (TLA) of this book. You don't really need to remember it or any of the many TLAs that follow. You may guess that the acronym for *four-letter acronym* is *FLA.* Wrong! A four-letter acronym is an *ETLA,* which stands for *extended three-letter acronym.* After all, it just wouldn't be right if the acronym for *four-letter acronym* had only three letters.

- **On the network:** Every computer connected to the network is said to be "on the network." The technical term (which you can forget) for a computer that's on the network is a *node.*

- **Online, offline:** When a computer is turned on and can access the network, the computer is *online.* When a computer can't access the network, it's *offline.* A computer can be offline for several reasons. The computer can be turned off, the user may have disabled the network connection, the computer may be broken, the cable that connects it to the network can be unplugged, or a wad of gum can be jammed into the disk drive.

- **Up, down:** When a computer is turned on and working properly, it's *up.* When a computer is turned off, broken, or being serviced, it's *down.* Turning off a computer is sometimes called *taking it down.* Turning it back on is sometimes called *bringing it up.*

- **Local, remote:** A resource such as a disk drive is *local* if it resides in your computer. It's *remote* if it resides in another computer somewhere else on your network.

- **Internet:** The *internet* is a huge amalgamation of computer networks strewn about the entire planet. Networking the computers in your home or office so that they can share information with one another and connecting your computer to the worldwide internet are two separate but related tasks.

Seeing What You Can Do with a Network

Frankly, computer networks are a bit of a pain to set up. So, why bother? Because the benefits of having a network outweigh the difficulties of setting one up.

You don't have to be a PhD to understand the benefits of networking. In fact, you learned everything you need to know in kindergarten: Networks are all about sharing. Specifically, networks are about sharing three things: files, resources, and programs.

Accessing the internet

Probably the main reason most small business and home networks exist is to allow everyone to access the internet through a single shared internet connection. In Figure 1-1, you can see that the wireless router is connected to the internet. By sharing this connection, all the computers on the network, whether connected wirelessly or via cables, can access the internet through the wireless router.

It's important to note that nearly all wireless routers also contain a built-in *firewall.* The firewall helps protect the computers on the network from the imminent dangers of the internet. The moment you connect a home or office network to the internet, cybercriminals will begin trying to break into your network and try to trick you into divulging sensitive information, such as the password to your bank account.



WARNING

Never — and I mean *never* — allow any computer to connect directly to the internet without a firewall in place.

Sharing files

Networks enable you to share information with other computers on the network. Depending on how you set up your network, you can share files with your network friends in several different ways. You can send a file from your computer directly to a friend's computer by attaching the file to an email message and then mailing it. Or you can let your friend access your computer over the network so that your friend can retrieve the file directly from your hard drive. Yet another method is to copy the file to a disk on another computer and then tell your friend where you put the file so that your friend can retrieve it later. One way or the other, the data travels to your friend's computer over the network cable and not on a CD or DVD or flash drive.

Sharing resources

You can set up certain computer resources — such as hard drives or printers — so that all computers on the network can access them. For example, the printer in Figure 1-1 is a *shared resource,* which means that anyone on the network can use it. Without the network, Homer, Marge, Lisa, and Bart would have to buy their own printers.

Hard drives can be shared resources, too. In fact, you must set up a hard drive as a shared resource to share files with other users. Suppose that Lisa wants to share a file with Bart, and a shared folder has been set up on Homer's computer. All Lisa has to do is copy the file to the shared folder in Homer's computer and tell Bart where she put it. Then, when Bart gets around to it, he can copy the file from Homer's computer to his own.

Sharing programs

Instead of keeping separate copies of programs on each person's computer, put programs on a drive that everyone shares. For example, if ten computer users all use a particular program, you can purchase and install ten copies of the program, one for each computer. Or you can purchase a ten-user license for the program and then install just one copy of the program on a shared drive. Each of the ten users can then access the program from the shared hard drive.



WARNING

Purchasing a single-user copy of a program and then putting it on a shared network drive — so that everyone on the network can access it — is illegal. If five people use the program, you need to either purchase five copies of the program or purchase a network license that specifically allows five or more users.



TIP

That being said, many software manufacturers sell their software with a concurrent usage license, which means that you can install the software on as many computers as you want, but only a certain number of people can use the software at any given time. Usually, special licensing software that runs on one of the network's server computers keeps track of how many people are currently using the software. This type of license is frequently used with more specialized (and expensive) software, such as accounting systems or computer drafting systems.

Another common method for software vendors to license their software is through a monthly or yearly subscription. You just give them your credit card number, and they give you the right to use the software. You need a working internet connection so that the software can confirm that you have a valid subscription each time you run the software.

Another benefit of networking is that networks enable computer users to communicate with one another over the network. The most obvious way networks allow computer users to communicate is by passing messages back and forth, using email or instant-messaging programs. Networks also offer other ways to communicate. For example, you can hold online meetings over the network. Network

users who have inexpensive video cameras (webcams) attached to their computers can have videoconferences. You can even play a friendly game of hearts over a network — during your lunch break, of course.

Getting Acquainted with Servers and Clients

The network computer that contains the hard drives, printers, and other resources that are shared with other network computers is a *server.* This term comes up repeatedly, so you have to remember it. Write it on the back of your left hand.

Any computer that's not a server is a *client.* You have to remember this term, too. Write it on the back of your right hand.

Only two kinds of computers are on a network: servers and clients. Look at your left hand and then look at your right hand. Don't wash your hands until you memorize these terms.

The distinction between servers and clients in a network has parallels in sociology — in effect, a sort of class distinction between the "haves" and "have-nots" of computer resources:

- Usually, the most powerful and expensive computers in a network are the servers. There's a good technical reason for this: All users on the network share the server's resources.
- The cheaper and less-powerful computers in a network are the clients. *Clients* are the computers used by individual users for everyday work. Because clients' resources don't have to be shared, they don't have to be as fancy. (The exception to this rule is if the users do work that requires powerful desktop computers — for example, engineering design or video processing.)
- Most networks have more clients than servers. For example, a network with ten clients can probably get by with one server, but larger networks will likely require more servers.
- In most networks, a clean line of demarcation exists between servers and clients. In other words, a computer functions as either a server or a client, not both. For the sake of an efficient network, a server can't become a client, nor can a client become a server.

- Other (usually smaller) networks can be more evenhanded by allowing any computer in the network to be a server and allowing any computer to be both a server and a client at the same time.

Weighing Your Options: Dedicated Servers versus Peer-to-Peer Networks

In most networks, a server computer is a server computer and nothing else. It's dedicated to the sole task of providing shared resources, such as hard drives and printers, to be accessed by the network client computers. This type of server is a *dedicated server* because it can perform no other task than network services.

Some smaller networks take an alternative approach by enabling any computer on the network to function as both a client and a server. Thus, any computer can share its printers and hard drives with other computers on the network. And while a computer is working as a server, you can still use that same computer for other functions, such as word processing. This type of network is a *peer-to-peer network* because all the computers are thought of as *peers*, or equals.

Here are some points to ponder concerning the differences between dedicated-server networks and peer-to-peer networks while you're walking the dog tomorrow morning:

- Peer-to-peer networking features are built into Windows. Thus, if your computer runs Windows, you don't have to buy any additional software to turn your computer into a server. All you have to do is enable the Windows server features.
- The network server features that are built into Windows 11 (the most popular desktop operating system) aren't particularly efficient because this version of Windows wasn't designed primarily to be a network server.

  If you dedicate a computer to the task of being a full-time server, use a special server operating system rather than the standard Windows desktop operating system. A *server operating system* is specially designed to handle networking functions efficiently.

  The most commonly used server operating systems are the server versions of Windows. As of this writing, the current server version of Windows is Windows Server 2025. However, many companies still use the previous version



REMEMBER

(Windows Server 2022), and a few use even earlier versions such as Windows Server 2016 and 2019.

Another popular server operating system is Linux. Linux is popular because it's free. Linux requires more expertise to set up than Windows Server does, but it's just as capable.

- Many networks are both peer-to-peer *and* dedicated-server networks at the same time. These networks have
  - At least one server computer that runs a server operating system such as Windows Server 2025
  - Client computers that use the server features of Windows 11 to share their resources with the network
- Besides being dedicated, your servers should also be sincere.



TIP

Understanding What Makes a Network Tick

To use a network, you don't really have to know much about how it works. Still, you may feel a little bit better about using the network if you realize that it doesn't work by voodoo. A network may seem like magic, but it isn't. The following list describes the inner workings of a typical network:

- **Network interface:** Inside any computer attached to a network is a special electronic circuit called the *network interface.* The network interface has either an external jack into which you can plug a network cable — or, in the case of a wireless network interface, an antenna.
- **Network cable:** The network cable physically connects the computers. It plugs into the network interface card (NIC) on the back of your computer.

  The type of network cable most commonly used is *twisted-pair cable,* so named because it consists of several pairs of wires twisted together in a certain way. Twisted-pair cable superficially resembles telephone cable. However, appearances can be deceiving. Most phone systems are wired using a lower grade of cable that doesn't work for networks.

  For the complete lowdown on networking cables, see Book 1, Chapter 2.



TIP

Network cable isn't necessary when wireless networking is used. For more information about wireless networking, see Book 1, Chapter 2.



TECHNICAL STUFF

- **Network switch:** Networks built with twisted-pair cabling require one or more switches. A *switch* is a box with a bunch of cable connectors. Each computer on the network is connected by cable to the switch. The switch, in turn, connects all the computers to each other.

  Most networks of more than a few dozen computers have more than one switch. In that case, the switches themselves are connected to each other with cable in a manner that allows all the computers to communicate with each other without regard to which switch they're directly connected to.

  In the early days of twisted-pair networking, devices known as *hubs* were used rather than switches. The term *hub* is sometimes used to refer to switches, but true hubs went out of style sometime around the turn of the century.

  I explain much more about switches and hubs in Book 1, Chapters 2 and 3.

- **WAP:** In a wireless network, most cables and switches are moot. Instead, radio takes the place of cables. The device that enables a computer to connect wirelessly to a network is called a *wireless access point.* A WAP is a combination of a radio transmitter and a radio receiver and has an integrated wired network port. The WAP must be connected to the network via a cable, but it allows wireless devices such as laptops, tablets, and phones to connect wirelessly.

- **Router:** A device found in nearly all networks is a *router,* which is used to connect two networks — typically your internal network and the internet. You can find out more about routers in Book 1, Chapters 2 and 3.

- **Firewall:** A *firewall* is an essential component of any network that connects to the internet. The firewall provides security features that help keep cybercriminals out of your network.

  In most cases, the function of a firewall is combined with the function of a router in a single device called a *firewall router.* A firewall is a security wall between two networks (usually the internet and your internal network). So, in a firewall router, the router component links the two networks, while the firewall component provides security.

  In home networks or small office networks, it's also common to combine the functions of firewall, router, WAP, and switch into a single device that's usually called a *wireless router* or a *Wi-Fi router.* When you purchase such a device, check to make sure it has adequate firewall features and the correct number of switch ports for your wired devices.

Putting the Pieces Together

In a small network such as the one that was shown in Figure 1-1, a wireless router combines the function of firewall, router, switch, and WAP. This arrangement is fine for very small networks, but when you exceed the wired switch capacity of the wireless router, you'll need additional components.

Figure 1-2 shows a network with a separate switch to connect multiple computers. Here, you can see that the wireless firewall router connects to both the internet and the switch. Several computers have wired connections to the switch, and wireless devices connect via the WAP that's built in to the Wi-Fi router. The router also provides the firewall function.

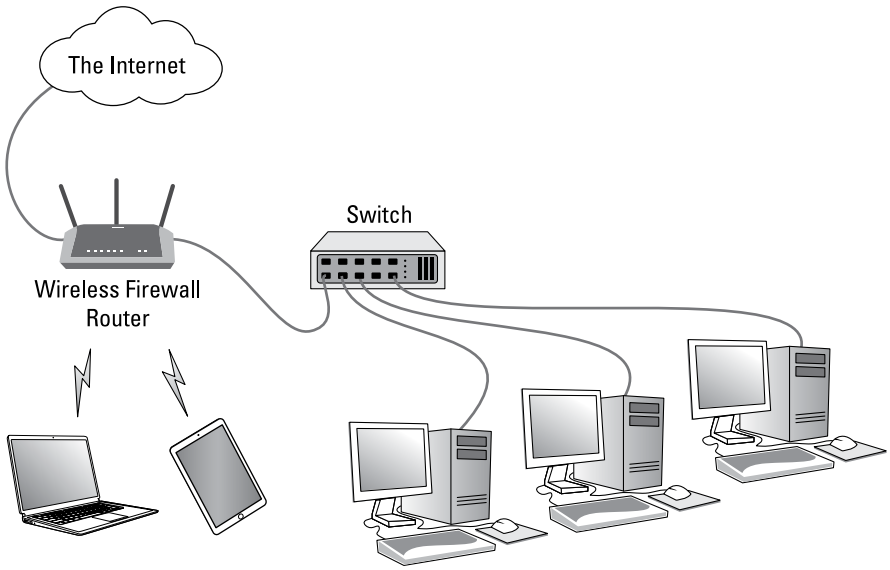


FIGURE 1-2: A network with a wireless firewall router and a switch.

Figure 1-3 shows a more complicated setup, in which the WAP is separated from the router. Here, the router with its built-in firewall connects to the internet and to the switch. As before, several computers have wired connections to the switch. In addition, the WAP has a wired connection to the switch, allowing wireless devices to connect to the network.

In Book 1, Chapter 3, you see examples of more complicated arrangements of these basic network components.

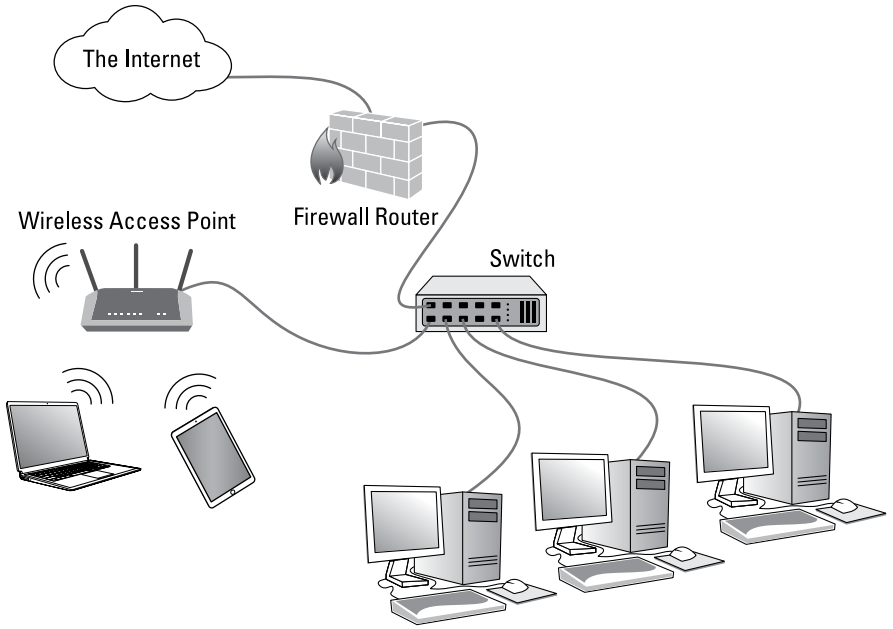


FIGURE 1-3: A network with a separate firewall router, switch, and WAP.

Considering Network Size

Networks come in all sizes and shapes. In fact, networks are commonly based on the geographical size they cover, as described in the following list:

- **LANs:** In this type of network, computers are relatively close together, such as within the same office or building.

  Don't let the descriptor *local* fool you. A LAN doesn't imply that a network is small. A LAN can contain hundreds or even thousands of computers. What makes a network a LAN is that all its connected computers are located within close proximity. Usually a LAN is contained within a single building, but a LAN can extend to several buildings on a campus, provided that the buildings are close to each other (typically within 300 feet, although greater distances are possible with special equipment).

- **Wide area networks (WANs):** These networks span a large geographic territory, such as an entire city, a region, or even a country. WANs are typically used to connect two or more LANs that are relatively far apart. For example, a WAN may connect an office in San Francisco with an office in New York.



REMEMBER

The geographic distance, not the number of computers involved, makes a network a WAN. If an office in San Francisco and an office in New York each has only one computer, the WAN will have a grand sum of two computers — but will span more than 3,000 miles.

- **Metropolitan area networks (MANs):** This kind of network is smaller than a typical WAN but larger than a LAN. Typically, a MAN connects two or more LANs within the same city that are far enough apart that the networks can't be connected via a simple cable or wireless connection.

Recognizing That Your Personal Computer Isn't Personal When It's on a Network

If I had to choose one point that I want you to remember from this chapter more than anything else, it's this: After you hook up your personal computer (PC) to a network, it's not a "personal" computer anymore. You're now part of a network of computers, and in a way, you've given up one of the key concepts that made PCs so successful in the first place: independence.

I got my start in computers back in the days when mainframe computers ruled the roost. *Mainframe computers* are big, complex machines that used to fill entire rooms and had to be cooled with chilled water. My first computer was a water-cooled Acme Hex Core Model 2000. (I'm not making up the part about the water. A plumber was often required to install a mainframe computer. In fact, the really big ones were cooled by liquid nitrogen. I *am* making up the part about the Acme Hex Core 2000.)

Mainframe computers required staffs of programmers and operators in white lab coats just to keep them going. The mainframes had to be carefully managed. A whole bureaucracy grew up around managing them.

Mainframe computers used to be the dominant computers in the workplace. PCs changed all that: They took the computing power out of the big computer room and put it on the user's desktop, where it belongs. PCs severed the tie to the centralized control of the mainframe computer. With a PC, a user could look at the computer and say, "This is mine — all mine!" Mainframes still exist, but they're not nearly as popular as they once were.

But networks have changed everything all over again. In a way, it's a change back to the mainframe-computer way of thinking: central location, distributed resources. True, the network isn't housed in the basement and doesn't have to be installed by a plumber. But you can no longer think of "your" PC as your own. You're part of a network — and like the mainframe, the network has to be carefully managed.

Here are several ways in which a network robs you of your independence:

- **You can't just indiscriminately delete files from the network.** They may not be yours.
- **You're forced to be concerned about network security.** For example, a server computer has to know who you are before it allows you to access its files. So, you have to know your user ID and password to access the network. This precaution prevents some 15-year-old kid from hacking their way into your office network by using its internet connection and stealing all your computer games.
- **You may have to wait for shared resources.** You may need to print a quick page on your way into a meeting that you're already late for, only to discover that someone else sent a 1,000-page document to the printer. You'll have to wait or find a different printer.
- **You may have to wait for access to documents.** You may try to retrieve a Microsoft Excel spreadsheet file from a network drive, only to discover that someone else is using it. You'll just have to wait. (Newer technologies have made it possible for multiple people to edit files at the same time, which is kind of mind-blowing.)
- **You don't have unlimited storage space.** If you copy a 100GB video file to a server's drive, you may get calls later from angry coworkers complaining that no room is left on the server's drive for their important files.
- **Your files can become infected by viruses given to you by someone over the network.** You may then accidentally infect other network users.
- **You have to be careful about saving sensitive files on the server.** If you write an angry note about your boss and save it on the server's hard drive, your boss may find the memo and read it.
- **The server computers may be down for maintenance.** This happens all the time. Servers need to be kept up to date with system updates, or new software may need to be installed. At times, the servers will be taken offline for such purposes. When the servers are offline, you'll have to wait. (Most IT administrators schedule server downtime at weird hours, like midnight, so these outages shouldn't affect you unless you keep odd hours.)

Assigning a Network Administrator

Because so much can go wrong — even with a simple network — designating one person as network administrator is important. This way, someone is responsible for making sure that the network doesn't fall apart or get out of control.

The network administrator doesn't have to be a technical genius. In fact, some of the best network administrators are complete idiots when it comes to technical stuff. What's important is that the administrator is organized. That person's job is to make sure that plenty of space is available on the file server, that the file server is backed up regularly, and that new employees can access the network, among other tasks.

The network administrator's job also includes solving basic problems that the users themselves can't solve — and knowing when to call in an expert (when something really bad happens). It's a tough job, but somebody's got to do it. Here are a few tips that may help:

- In small companies, picking the network administrator by drawing straws is common. The person who draws the shortest straw loses and becomes administrator.
- Of course, the network administrator can't be a *complete* technical idiot. I was lying about that. (For those of you in Congress, the word is *testifying*.) I exaggerated to make the point that organizational skills are more important than technical skills. The network administrator needs to know how to do various maintenance tasks. Although this knowledge requires at least a little technical know-how, the organizational skills are more important.

Because network administration is such an important job, all the chapters in Book 9 are devoted to it.

IN THIS CHAPTER

- Looking at the various elements that make up a typical network infrastructure
- Considering how standards and protocols are used in networking
- Taking a look at network topology
- Examining the elements of a network's cable infrastructure
- Understanding ports, interfaces, and MAC addresses
- Seeing how network data is transmitted via packets
- Understanding collisions in wired and wireless networks
- Introducing broadcast packets
- Perusing wireless networks

Chapter 2
Network Infrastructure

In this chapter, I cover the key concepts of local area networks (LANs) — that is, networks that are contained within a single location. Although this chapter may seem a little abstract, you'll be much better prepared to design and implement a solid LAN if you have a good understanding of these concepts from the very beginning.

I go into more depth on many of the concepts presented in this chapter in Book 2, which dives deeper into the various networking standards and protocols.

Introducing Infrastructure

As I mention in the preceding chapter, a *local area network* is a network that connects computers and other devices that are located in relatively close proximity to one another. Most LANs are contained to a single building, although it's possible to create LANs that span several buildings at a single site, provided the buildings are close to one another. For the purposes of this chapter, I stick to LANs that operate within a single building and support anywhere from a few dozen to a few hundred users.

LANs exist to connect computing devices — such as workstation computers, servers, printers, scanners, cameras, and so on — to one another. The essence of a network is the *physical infrastructure* that enables the connections. The infrastructure is similar to the infrastructure of a city. A city's infrastructure has many physical elements, including roads, stop signs and stoplights, water supply lines, stormwater drains, sewage lines and treatment plants, electrical distribution cables, transformers, and much more.

Similarly, the infrastructure of a network consists of physical elements:

- **Cables:** These run through walls and ceiling spaces, through conduits, between floors, and wherever else they need to go to reach their destinations.
- **Patch panels:** These allow cables to be organized at a central location.
- **Network switches:** A *switch* is an intermediate device that sits between the networked devices that allows those devices to communicate with each other. In a real way, switches are the core of the network; without switches, computers wouldn't be able to talk.
- **Wireless access points (WAPs):** A *wireless access point* lets devices connect wirelessly to the network. Depending on the size of your network and the physical space your users occupy, you may need more than one WAP. Each WAP needs to be connected to the LAN via a cabled switch connection.
- **At least one router:** A *router* connects the network to the outside world. The most common use of a router is to connect the LAN to the internet. However, routers can also be used to connect one LAN to another. (You can find more about routers in Book 1, Chapter 3.)

Understanding Network Protocols and Standards

To operate efficiently, the infrastructure of a network consists of devices that conform to well-known standards and protocols. A *protocol* provides a precise sequence of steps that each element of a network must follow to enable communications. Protocols also define the precise format of all data that is exchanged in a network. For example, the Internet Protocol (IP) defines the format of IP addresses: four 8-bit numbers called *octets*, whose decimal values range from 0 to 255, as in 10.0.101.155 (see "Perusing Ports, Interfaces, and MAC Addresses," later in this chapter for more on octets).

A *standard* is a detailed definition of a protocol that has been established by a standards organization and that vendors follow when they create products. Without standards, it would be impossible for one vendor's products to work with another vendor's products. Because of standards, you can purchase equipment from different vendors with the assurance that they'll work together.

Network standards are organized into a framework called the Open Systems Interconnection (OSI) reference model. The OSI model establishes a hierarchy for protocols so that each protocol can deal with just one part of the overall task of data communications. Table 2-1 shows seven distinct layers at which a protocol may operate as described by the OSI reference model.

TABLE 2-1 The Seven Layers of the OSI Model

Layer	Name	Description
1	Physical	Governs the layout of cables and devices, such as repeaters and hubs.
2	Data link	Provides media access control (MAC) addresses to uniquely identify network nodes and a means for data to be sent over the physical layer in the form of packets. Bridges and switches are layer-2 devices.
3	Network	Handles routing of data across network segments.
4	Transport	Provides for reliable delivery of packets.
5	Session	Establishes sessions between network applications.
6	Presentation	Converts data so that systems that use different data formats can exchange information.
7	Application	Allows applications to request network services.

Although the upper layers of the OSI model (layers 4 through 7) are equally important, in this chapter and the next, I focus on the first three layers of the OSI model — physical, data link, and network. These layers are the ones where the most common types of networking hardware (such as cables, interfaces, switches, and routers) operate.

Although many different network protocols and standards can be used in various layers of the OSI model, the most common standard found at layers 1 and 2 is Ethernet. Similarly, the most common standard at layer 3 is IP. You can find more about Ethernet and IP in Book 2, Chapters 2 and 3.

Recognizing Network Topology

The term *network topology* refers to the shape of how the computers and other network components are connected to each other. Several different types of network topologies exist, each with advantages and disadvantages.



TIP

In the following discussion of network topologies, I use two important terms:

- **Node:** A *node* is a device that's connected to the network. For your purposes here, a node is the same as a computer. Network topology deals with how the nodes of a network are connected to each other.
- **Packet:** A *packet* is a message that's sent over the network from one node to another node. The packet includes the address of the node that sent the packet, the address of the node the packet is being sent to, and data.

Bus topology

In a *bus topology*, nodes are strung together in a line, as shown in Figure 2-1. The key to understanding how a bus topology works is to think of the entire network as a single cable, with each node "tapping" into the cable so it can listen in on the packets being sent over that cable.

In a bus topology, every node on the network can see every packet that's sent on the cable. Each node looks at each packet to determine whether the packet is intended for it. If it is, the node claims the packet. If it isn't, the node ignores the packet. This way, each computer can respond to data sent to it and ignore data sent to other computers on the network.

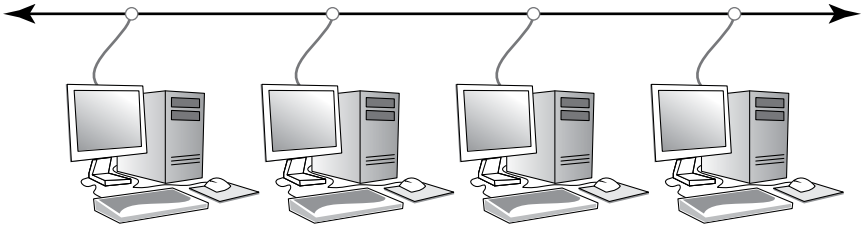


FIGURE 2-1:
Bus topology.

If the cable in a bus network breaks, the entire network is effectively disabled. Obviously, the nodes on opposite sides of the break can't continue to communicate with each other, because data can't span the gap created by the break. But even those nodes that are on the *same* side of the break may not be able to communicate with each other, because the open end of the cable left by the break disrupts the proper transmission of electrical signals.

In the early days of Ethernet networking, bus topology was commonplace. Although, for most networks, bus topology has given way to star topology (see the next section), many networks today still have elements that rely on bus topology.

Star topology

In a *star topology*, each network node is connected to a central device called a *hub* or a *switch*, as shown in Figure 2-2. Star topologies are commonly used with LANs.

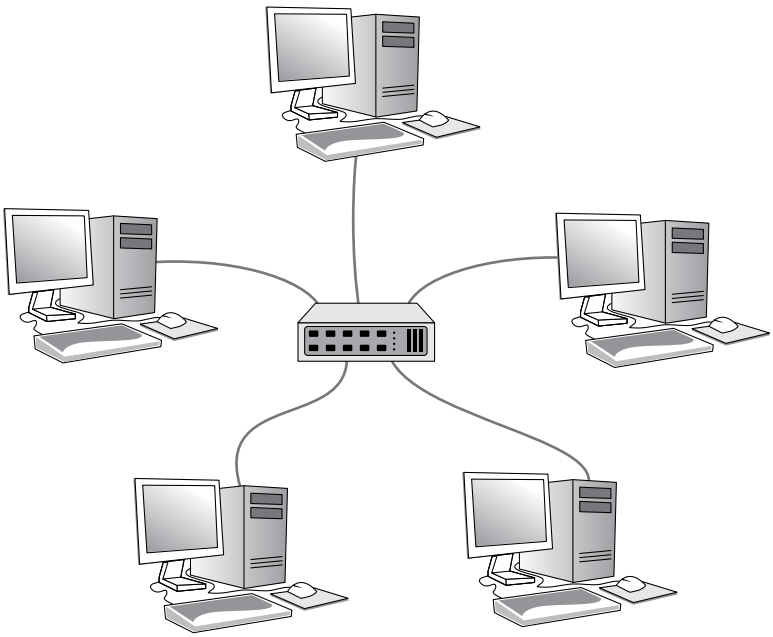


FIGURE 2-2:
Star topology.

If a cable in a star network breaks, only the node connected to that cable is isolated from the network. The other nodes can continue to operate without interruption — unless, of course, the node that's isolated because of the break happens to be the file server.


TECHNICAL STUFF

You should be aware of the somewhat technical distinction between a hub and a switch. Simply put, a *hub* doesn't know anything about the computers that are connected to each of its ports. So, when a computer connected to the hub sends a packet to a computer that's connected to another port, the hub sends a duplicate copy of the packet to all its ports. In contrast, a switch knows which computer is connected to each of its ports. As a result, when a switch receives a packet intended for a particular computer, it sends the packet only to the port that the recipient is connected to.

Strictly speaking, only networks that use switches have a true star topology. If the network uses a hub, the network topology has the physical appearance of a star, but it's actually a bus. That's because when a hub is used, each computer on the network sees all the packets sent over the network, just as in a bus topology. In a true star topology, as when a switch is used, each computer sees only those packets that were sent specifically to it, as well as packets that were specifically sent to all computers on the network (those types of packets are called *broadcast packets;* see "Dealing with Broadcast Packets," later in this chapter).

EXPANDING STARS

Physicists say that the universe is expanding, and network administrators know they're right. A simple bus or star topology is suitable only for small networks, with a dozen or so computers. But small networks inevitably become large networks as more computers are added. For larger networks, it's common to create more complicated topologies that combine stars and buses.

For example, a bus can be used to connect several stars. In this case, two or more hubs or switches are connected to each other using a bus. Each of these hubs or switches is then the center of a star that connects two or more computers to the network. This type of arrangement is commonly used in buildings that have two or more distinct workgroups. The bus that connects the switches is sometimes called a *backbone.*

Another way to expand a star topology is to use a technique called *daisy-chaining.* When you use daisy-chaining, a switch is connected to another switch as if it were one of the nodes on the star. Then this second switch serves as the center of a second star.

Ring topology

In a ring topology, packets are sent around the circle from computer to computer, as shown in Figure 2-3. Each computer looks at each packet to decide whether the packet was intended for it. If it isn't, the packet is passed on to the next computer in the ring.

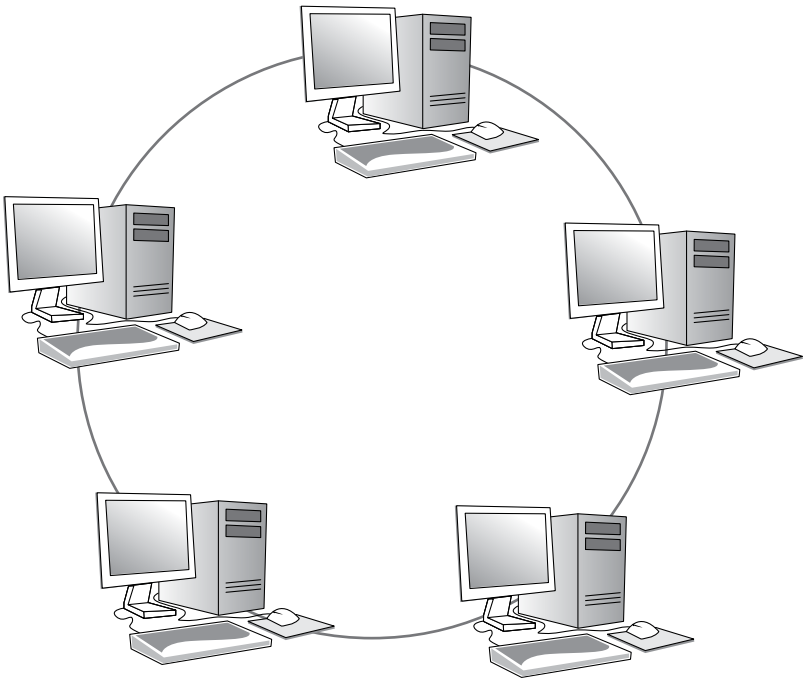


FIGURE 2-3: Ring topology.

Years ago, ring topologies were common in LANs, because two popular networking technologies used rings: Attached Resource Computer Network (ARCNET) and Token Ring. ARCNET is still used for certain applications, such as factory automation, but it's rarely used in business networks. Token Ring is still a popular network technology for IBM midrange computers. Although plenty of Token Ring networks are still in existence, not many new networks use Token Ring any more.

Ring topology was also used by Fiber Distributed Data Interface (FDDI), one of the first types of fiber-optic network connections. FDDI has given way to more efficient fiber-optic techniques, however. So, ring networks have all but vanished from business networks.

Mesh topology

In a *mesh topology,* multiple connections exist between each of the nodes on the network, as shown in Figure 2-4. The advantage of a mesh topology is that if one cable breaks, the network can use an alternative route to deliver its packets.

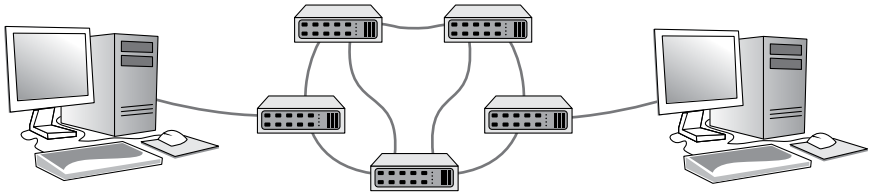


FIGURE 2-4: Mesh topology.

Mesh networks are often used to link switches in a LAN. In Figure 2-4, the mesh has a total of seven connections. If any of these seven connections goes bad, any switch can still reach any other switch while traveling through, at most, one intermediate switch.

Mesh networks are also very common for metropolitan area networks (MANs) or wide area networks (WANs). These networks use routers to route packets from network to network. For reliability and performance reasons, routers are usually arranged in a way that provides multiple paths between any two nodes on the network in a mesh-like arrangement.

Considering Cable

You can find much more about the details of working with network cable in Book 3, Chapter 1, as well as Book 4, Chapter 1. But before we get too far, I want to give you an overview of what's involved with cabling together a network.

For starters, network cable and all the bits and pieces that go along with it are the most important components of layer 1 of the OSI reference model. The following sections describe the most important layer 1 and cabling details you need to know.

Twisted-pair cable

You can choose from several varieties of cable, but the most common is called *twisted-pair.* It's called that because inside the outer sheath of the cable are four pairs of small insulated wire. The wires are 24 gauge, which means they're about

half a millimeter in diameter. These pairs are color coded: blue, green, orange, and brown. For each pair, there is one solid-colored wire and one striped wire — so, the blue pair consists of a solid blue wire and a blue-and-white striped wire.

The two wires that make up each pair are twisted together in a way that prevents the electrical signals within each pair from interfering with the other pairs. To accomplish this, each pair is twisted at a different rate.

The maximum length of a single run of Cat5e cable is 100 meters.

Cat5e cable is able to carry network data at speeds of up to 1 gigabit per second (Gbps). The newer and somewhat more expensive Cat6 cable can carry data at up to 10 Gbps but can sustain that speed for only 55 meters.

RJ45 connectors

Twisted-pair cable is attached to network devices using a special type of connector called an RJ45, which is a small block of plastic with eight metal contacts. RJ45 connectors resemble telephone connectors, but they're larger (telephone connectors have just four electrical contacts). For the cable to meet Cat5e standards, the twists of the individual pairs must be maintained all the way up to the RJ45 connector.

RJ45 connectors come in both male (plug) and female (receptacle) varieties. Typically, the male connector is installed on the cables and the female connectors are installed in equipment. So, in order to connect a cable to a computer, you plug the male RJ45 plug on the cable into the female RJ45 receptacle on the computer.

Patch panels and patch cables

A *patch panel* is a group of RJ45 receptacles on a single metal plate, usually attached to a 19-inch equipment rack. Patch panels are used to bring cables that are run from individual computer locations to a single location where they can then be patched to other equipment using patch cables. A *patch cable* is simply a short length of twisted-pair cable with an RJ45 plug on both ends. Patch cables are usually 3 to 10 feet in length, but longer lengths are occasionally used.

Patch panels typically have either 24 or 48 ports. Depending on the size of your network, you may have more than one patch panel at a single location. For example, a large network may have four 48-port patch panels to support a total of 192 computers.



REMEMBER

A patch panel by itself doesn't actually *do* anything. Its job is simply to provide a central collecting point for all your network cables so that you can easily use patch cables to connect the cables to other devices, such as switches or servers.

Repeaters and hubs

A *repeater* is a layer-1 device that is designed to circumvent the maximum length limitation of twisted-pair network cables. A repeater contains two RJ45 ports, which are connected internally by an amplifier. Electrical signals received on either of the two ports are boosted by the amplifier and sent through the other port. Thus, the cables on both ends of the repeater can be up to 100 meters. The repeater effectively doubles the reach of the cable.

A *hub* is a repeater with more than two ports. For example, a hub may have four or eight ports. Each of these ports can connect to another device on the network, such as a client computer, a server, or a printer. A port on a hub can also connect to another hub, so that (for example) an eight-port hub can connect to seven computers and another eight-port hub, which can connect to seven more computers. In this way, two eight-port hubs can connect 14 computers to each other.

There are two very important things to know about hubs:

- **An electrical signal received on any of the hub's ports is amplified and repeated on all the other ports in the hub.** So, in an eight-port hub, any electrical signals received on port 1 are amplified and then sent out on ports 2 through 8. Any devices that are connected to ports 2 through 8 see the signals that were received on port 1. The same is true for signals received on any of the other ports; for example, any signals received on port 4 will be amplified and repeated on ports 1 through 3, as well as ports 5 through 8.
- **Hubs are almost never used anymore.** That's because simply repeating all incoming signals on all ports is an incredibly bad idea, for reasons that will become apparent later in this chapter and in Book 1, Chapter 3. If your network still has hubs, you should seriously consider replacing them with switches, which I describe in the next section and explain further in the next chapter.

Switches

A *switch* is a layer-2 device that is similar to a hub in that it allows you to connect more than one device, and packets received on one port are relayed to other ports. The difference, however, is that a switch is able to examine the actual contents of the data that it receives. As I explain in the "Pondering Packets" section, later in this chapter, data is sent in units called *packets* that contain a destination address.

A switch looks at this destination address and repeats the incoming packet only on the port that can deliver the packet to the intended destination.

For example, suppose computer A is connected to switch port 1, and computer D is connected to switch port 4. If computer A sends a packet to computer D, that packet is received on switch port 1. The switch knows that computer D is connected to switch port 4, so the switch sends the packet out on switch port 4. In this way, computer D receives the packet. The computers or devices that are connected to the other ports on the switch are not bothered with the packet intended for computer D.

If that doesn't make a lot of sense, don't worry — it will. The next two sections in this chapter explain the concept of MAC addresses, which are how networks identify the intended recipients of data packets, as well as how data packets work. Then, in Book 1, Chapter 3, I dive deeper into how switches do their magic.

Perusing Ports, Interfaces, and MAC Addresses

A *network interface* is the electronic circuitry that allows a device to connect to a network. Each network interface provides a *port,* which is the plug-in point for the interface. Generally speaking, the terms *port* and *interface* are synonymous.

A network interface may be a separate add-on card for a computer, in which case the interface is called a *network interface card* (NIC). On some devices, such as printers, separate NICs are still common. But nearly all desktop and laptop computers have a network interface built into the computer's motherboard, so separate NICs are rarely used on desktop computers or laptops. NICs are still widely used on servers, however, because servers are often configured with two or more interfaces; using a separate card for the interface allows for more flexibility.



TIP

The term *adapter* is often used as a synonym for *interface. Port, interface, adapter* — three words that mean the same thing.

Every network interface must have a unique identifier called a *MAC address.* (*MAC* stands for *media access control,* but that won't be on the test.) Each MAC address is unique throughout the entire world. I have no idea whether MAC addresses are unique throughout the galaxy; it's entirely possible that the computer system on some invading alien spacecraft would have a MAC address that is the same as your laptop, but if that were to happen, I doubt you'd be too concerned about fixing your network.

MAC addresses are important because they provide the means for a network to keep track of the devices that make up the network. Without MAC addresses, it would be impossible to know what devices are on the network. And it would be impossible to send information to a particular device or to know which particular device sent information.



TIP

The term *physical address* is sometimes used as a synonym for *MAC address.* The two terms are interchangeable.



TECHNICAL STUFF

MAC addresses are a part of layer 2 of the OSI reference model, called the link layer. This layer is responsible for the exchange of basic information on a network. The ability to uniquely identify every device on a network is a key component of enabling that to happen.

MAC addresses are 48 binary bits in length, which means that more than 280 trillion devices can be assigned unique MAC addresses before we run out. When written, MAC addresses are written in a peculiar notation that uses six two-digit hexadecimal numbers (each called an *octet*), separated by hyphens. Hexadecimal notation uses a combination of the ten digits 0 through 9, plus the letters A through F, to represent decimal values from 0 through 15 with a single symbol. For example, the decimal value 10 is represented by the letter A, 11 by the letter B, and so on. A typical MAC address looks like this:

```
48-2C-6A-1E-59-3D
```

The details of binary, hexadecimal, and octets are not important for the purposes of this chapter. I take a deep dive into these subjects in Book 2, Chapter 3.

If you want to see the MAC address of your computer's network adapter, open a command prompt and type **ipconfig /all**. Scroll through the output from this command to see the MAC address (ipconfig calls a physical address) for each interface on your computer. For example, here's the ipconfig output for the built-in adapter on my Microsoft Surface Book:

```
Ethernet adapter Ethernet 2:
    Media State ....................   : Media disconnected
    Connection-specific DNS Suffix ...   : lowewriter.com
    Description ....................   : Surface Ethernet Adapter
    Physical Address ...............   : 58-82-A8-9C-A7-28
    DHCP Enabled ...................   : Yes
    Autoconfiguration Enabled ......   : Yes
```

Here, you can see the MAC address is 58-82-A8-9C-A7-28.



TIP

A MAC address is technically associated with a network interface, not with the device that uses that interface. For example, if your computer's motherboard has a network interface built in, the MAC address of the network interface is pretty much married to the motherboard. However, if your computer has a separate NIC, the MAC address is a part of the card, not the computer that the card is plugged into. If you remove the interface card from one computer and install it in another, the MAC address travels with the card.



REMEMBER

The key point to remember here is that in order for a computer, printer, or any other device to connect to a network, that device must contain a network interface. That interface has a unique MAC address, which is the primary way that the network can distinguish one device from another.

Pondering Packets

When two or more devices are connected to a network via cables plugged into their network interfaces, those devices can exchange information with one another. This bit of magic is accomplished through the use of *packets,* which are relatively small units of data that are sent and received through the network interface and cables. A network packet always originates at a single network interface, called the *sender,* and it's usually (but not always) sent to a single network interface, called the *destination.*

A packet is very similar to an envelope that you would send through standard mail delivery. It includes the MAC address of both the sender and the destination, as well as some other interesting header information, along with a *payload* that contains the actual data being sent by the packet. You can think of the payload as what you would put in an envelope you want to send through the mail. You wouldn't dream of dropping an envelope in the mail without writing the recipient's address, as well as your own address, on the envelope. So it is with packets.

The payload of an Ethernet packet may be a packet created by some higher-level protocol, such as IP. This is analogous to putting a letter in an envelope, putting that envelope in a larger envelope, and sending it through the mail. When the recipient receives your mail, they open the envelope only to find another envelope that must be opened. That envelope may itself contain another envelope and so on, like Russian nesting dolls.



The term *frame* is often used instead of *packet*, but technically they're not quite the same. Every packet begins with a *preamble*, which consists of 56 bits of alternating zeros and ones. This preamble is used by the electronic circuitry of the interfaces to get their clocks synchronized properly so they can accurately read the rest of the packet. It's the rest of the packet that is technically called the *frame.* In other words, a *packet* consists of a *preamble* followed by a *frame.* Because the preamble is of concern only to the electronic engineers who design network interfaces, most non-engineers use the terms *packet* and *frame* interchangeably.

Ethernet has a standard packet format that all packets sent on an Ethernet network must follow. An Ethernet packet contains the following information:

- **Preamble (56 bits):** The preamble consists of alternating ones and zeros and is used to synchronize the precise timing required to read packet data.
- **Start-of-frame marker (1 byte):** A start-of-frame marker is a single byte that indicates that the frame is about to begin.
- **Destination MAC address (6 bytes).**
- **Sender MAC address (6 bytes).**
- **Tag (4 bytes):** The tag, which is used to support virtual local area networks (VLANs), is optional. A VLAN lets you divide two or more distinct LANs on a shared physical infrastructure (for example, cables and switches). (For more information about VLANs, see Book 1, Chapter 3, as well as Book 3, Chapter 1.)
- **Ethertype (2 bytes):** This field indicates the specific protocol that is contained in the payload.
- **Payload (46 to 1,500 bytes):** The payload contains the actual data being sent by the packet. If the information that needs to be sent is longer than 1,500 bytes, the information must be broken into two or more packets, sent separately, and then reassembled when the packets reach their destination. (The tasks of breaking up and reassembling the data are handled by protocols at higher layers in the OSI reference model; Ethernet itself has no understanding of what is in the packets it sends.)
- **Frame check sequence (4 bytes):** The frame check sequence (FCS) is used to ensure that the frame data was sent correctly. Basically, the interface that sends the packet uses an algorithm to calculate a 4-byte number based on the contents of the frame and saves this number in the FCS field. When the packet is received, the receiving interface repeats the calculation and then makes sure that the number recorded in the FCS portion of the packet matches the number it calculated. If the numbers disagree, the packet got garbled in transmission and is discarded.

Note that the details of an Ethernet packet are not really of much concern when you design and implement a network. Here are the main points to remember:

- Ethernet packets contain the MAC addresses of the sender and the receiver.
- The payload of an Ethernet packet is almost always a packet created by another higher-level protocol such as IP.
- Ethernet packets can contain a tag field used to implement VLANs, which provide an important means of organizing a large network into smaller parts that can be more easily managed.

Contemplating Collisions

One of the basic principles of Ethernet is that multiple devices can be connected to shared media (that is, cables), and that all devices connected to this media can and should examine every packet that is sent on the media. In other words, Ethernet uses shared media. (You'll find more information about this in Book 1, Chapter 3.)

Every packet contains the MAC address of the intended recipient. So, when an interface detects an incoming packet, it inspects the recipient MAC address and compares it with its own MAC address. If the addresses match, the interface passes the packet up to the next higher protocol on the protocol stack (typically, IP). If the addresses don't match, the interface assumes that the packet doesn't belong to the interface, so the interface simply ignores the packet.

Using hubs on an Ethernet propagates the shared cable through the network. That's because a hub simply amplifies any packet that arrives on any of its ports and then forwards the amplified packet to all the other ports in the hub. So, if you use a 12-port hub to connect 12 computers together, all 12 of the computers will see all the packets generated by any of the other computers. And if two or more of the computers try to transmit a packet at the same time, the packets will collide.

Ethernet has been very successful — in fact, it has become one of the most widely used networking protocols of all time. However, Ethernet's shared media approach has a basic problem: It doesn't scale well. When two or more interfaces are shared on a single cable, there is always the possibility that two or more interfaces will try to send information at the same time. This is called a *collision.* The result of a collision between two packets is that both packets will be destroyed in the process and will need to be sent again.

In a small network with just a few computers, collisions happen now and again but aren't a big deal. However, in a large network with dozens or hundreds of devices, collisions can become a constant annoyance. In fact, collisions can become such a problem that the network slows to a halt and no one is able to get anything done.


TIP

As a result, it's important to design a network in a way that reduces the possibility of collisions becoming a problem. Fortunately, that's easy to do with modern network equipment: All you have to do is use switches instead of hubs. Switches all but eliminate the problem of collisions by forwarding network packets only to the cable segments that the destination devices are connected to rather than forwarding them throughout the entire network.

You can find out more about how switches work and why it's so important in Book 1, Chapter 3.

Dealing with Broadcast Packets

Not all packets on an Ethernet network are intended for a single destination. Instead, some packets, called *broadcast packets*, are intended to be received by every device on the network. To send a broadcast packet, the sending interface sets the destination MAC address to FF-FF-FF-FF-FF-FF, which is the largest possible MAC address. Then all interfaces that receive the packet inspect the destination, see that the packet is a broadcast packet, and pass the packet up to the next higher protocol.

One of the most common users of broadcast packets is Dynamic Host Configuration Protocol (DHCP), which allows computers that join a network to be assigned an IP address. When a network interface is first connected to a network, it sends out a broadcast packet requesting the address of the network's DHCP server. Every device on the network sees this packet. But only the DHCP service will respond.


WARNING

As you see in the next chapter, broadcast packets can sometimes cause serious problems on your network. All networks should be planned in a way that minimizes problems caused by broadcast packets.

Examining Wireless Networks

As I mention in Book 1, Chapter 1, a *wireless network* is a network in which radio signals are used to connect devices to the network instead of using physical cables. You learn much more about wireless networks in Book 2, Chapter 1, as well as Book 4, Chapter 2. But for now, I want you to keep the following points in mind:

- **Just like with a wired network, a device connecting to a wireless network does so via a network interface.** A wireless interface, also known as a wireless adapter, includes a radio transmitter and receiver rather than a physical cable connection.
- **Every wireless network adapter has a MAC address.**
- **Instead of connecting to a switch or a hub, wireless devices connect to a WAP.**
- **Collisions are likely on a WAP, just as they are on a hub.** Unfortunately, there is no equivalent to a wireless switch that reduces the collision problem. WAPs are essentially hubs in that every device that connects to the WAP is competing for the same bandwidth. Whenever the WAP sends a packet, all devices connected to the WAP must inspect the packet to determine the MAC address destination. And if two devices try to send packets at the same time, a collision will occur. This is one of the inherent reasons that wireless networking is slower than wired networking.

IN THIS CHAPTER

- Considering the value of switches
- Understanding how switches do their magic
- Examining the role of routers
- Getting to know VLANs

Chapter 3

Switches, Routers, and VLANs

In this chapter, I dig deeper into two of the most basic and ubiquitous networking devices: switches and routers. Every network has at least one switch and one router, and all but the smallest networks have more than one switch. These components are the basic building blocks of networks, so understanding what they do and how they work is essential to properly designing, implementing, and maintaining a network that functions well.

Besides switches and routers, this chapter also introduces the concept of virtual local area networks (VLANs). A VLAN is a fancy technique that lets you split a single physical network into two or more logical networks. VLANs are one of the key techniques for organizing a network in a way that will allow the network to scale up as your organization grows. Small networks don't need to worry about VLANs, but even in a relatively small network, it pays to know what VLANs are. Introducing VLANs into your network before you actually need them will simplify your life as your network grows.

Understanding Switches

In the preceding chapter, I explain that a *hub* is a layer-1 device that simply repeats all incoming network data to all its output ports. In other words, if a hub has eight ports, any input data that arrives on port 1 will be amplified and repeated on ports 2 through 8. A hub is an unintelligent device — the hub doesn't know or care what the intended destination of the incoming data is. It simply sends the data to all its ports, hoping that the intended recipient is on one of those ports. (Actually, using the term *hoping* here is misleading, because as I said, the hub not only doesn't know who the intended recipient is but also doesn't even care. Hubs have no capacity for hope.)

Figure 3-1 shows a simple network with four computers connected via a hub. In this example, computer 1 is sending data to computer 4. As the figure shows, the hub doesn't know that the intended recipient is computer 3, so it sends the data not just to computer 3, but also to computer 2 and computer 3 as well.

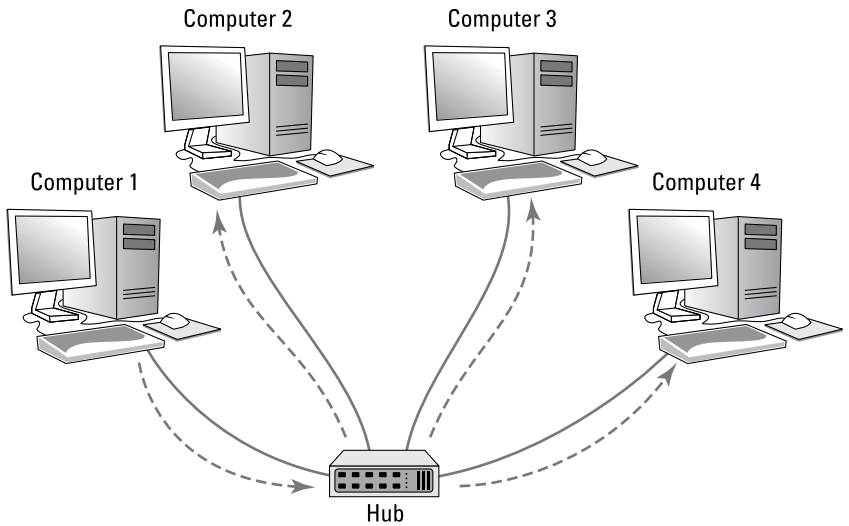


FIGURE 3-1: A hub repeats all incoming data on all its ports.

To understand why hubs even exist, or at least did exist in the distant past, you need a little history lesson. Ethernet was invented in the late 1970s and first became commercially available in 1980. From the very beginning, Ethernet used what is called *shared media* to connect devices in a network. The basic idea of Ethernet is that data is sent over network cables in the form of *packets*, which follow a well-defined structure. The key elements of the original Ethernet were (and still are) as follows:

- **All devices on the network can access all data sent over the network.** That's why the network cable itself is considered to be *shared media.*
- **Every device on the network has a unique identifier called a MAC address.** I cover MAC addresses in the preceding chapter. As a quick reminder, MAC addresses are 48 bits long and are written as six *octets* separated by hyphens. For example, 21-76-3D-7A-F6-1E is a valid MAC address. (See the preceding chapter for more on octets.)
- **A data packet includes the MAC address of the packet's intended recipient, as well as the MAC address of the sender.**
- **Every device on the network receives every packet that is sent on the network and examines the destination MAC address to determine whether the packet is intended for it.** If so, the device says, "Mine!" and stores the packet to be processed by other protocols higher up the food chain (that is, at higher levels in the OSI reference model; see the preceding chapter). If the destination MAC address doesn't match the device's, the device says "Hmph!" and simply ignores the packet.

  All the devices on the network do this examination, keeping only the packets that belong to them and ignoring all the others.
- **If the destination MAC address is all ones (represented as FF-FF-FF-FF-FF-FF), the packet is called a *broadcast packet.*** When a broadcast packet is sent, every device on the network looks at the destination MAC address, sees that the packet is a broadcast packet, and says, "Mine!" Broadcast packets are received by every device on the network.
- **Every once in a while, two devices try to send a packet at the exact same time.** When that happens, both packets are garbled. The result is called a *collision.* When collisions happen, both senders wait for a brief amount of randomly generated time and then try again. The collision probably won't happen again. But if it does, the senders wait and try again later.

So, that's a recap of the basic operation of the Ethernet networking system. Because it was a great system when it was invented, it quickly replaced the two dominant network technologies that were popular at the time, Attached Resource Computer Network (ARCNET) and Token Ring. But unfortunately, Ethernet had a few serious problems lurking under the surface that proved to be a problem for larger networks:

- **The frequency of collisions rises exponentially with the number of devices added to the network.** When you get too many devices, collisions happen all the time, and devices spend way too much time resending packets, sometimes having to resend them over and over again until a collision doesn't happen. This results in the network becoming much slower as it grows larger.

- **The frequency of broadcast packets can quickly increase as more devices are added to the network, further adding to the performance problem and the likelihood of collisions.**
- **Security is difficult to enforce, because every device on the network *must* examine every packet that comes its way.** Even though devices are supposed to ignore packets that aren't meant for them, there is no way to ensure that they do so.

Switches to the rescue!

A switch is essentially an intelligent hub that has the ability to actually look at the contents of the packets it processes and make intelligent decisions about what to do with them. A hub is a layer-1 device, which means that it can do nothing but receive and amplify electrical signals. In contrast, switches are layer-2 devices, which means they can actually inspect the layer-2 packets and act intelligently based on the contents of each packet.

A switch examines the destination MAC address of every packet it receives and forwards the packet only to the port that leads to the packet's intended destination. Thus, packets aren't repeated on ports that don't contain the packets' destination.

Figure 3-2 shows the same simple network that was shown in Figure 3-1, but this time with a switch instead of a hub. As you can see, the switch is smart enough to know that the data sent by computer 1 is intended for computer 3, so it sends the data only to computer 3. The switch leaves computer 2 and computer 3 alone so they can concentrate on other work.

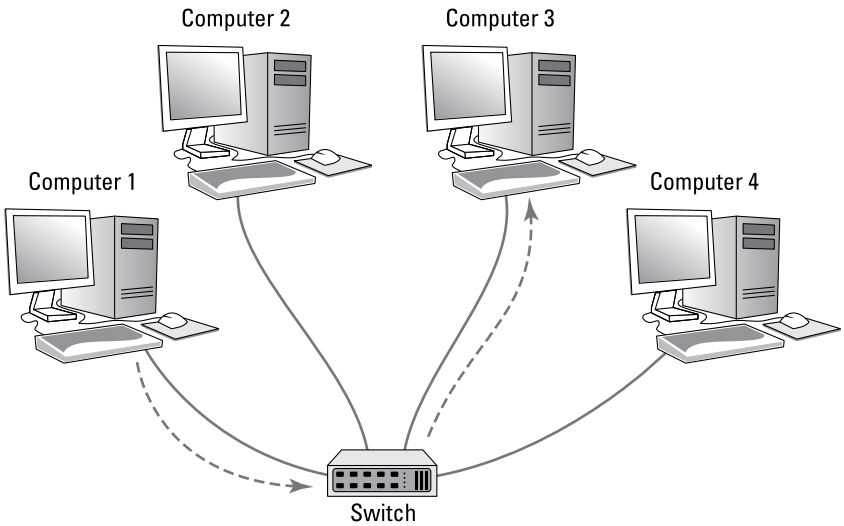


FIGURE 3-2: Unlike a hub, a switch knows where to send its data.

In order to accomplish intelligent forwarding, a switch must know what devices are connected to each of its ports. In the next section, you see how a switch learns what devices are connected to each of its ports.

Learning

For a switch to do its job, it needs to know which devices are connected to each of its ports. More specifically, the switch needs to know what MAC addresses are reachable via each of its ports. It does this in an ingeniously simple way: It simply learns. Whenever a packet is received on any of the switch's ports, the switch examines the sending MAC address in the packet. The switch rightly assumes that if it received a packet from a given MAC address on a given port, the switch can reach that MAC address via that port. For example, if a switch receives a packet from computer C on port 3, the switch has learned that computer C is reachable on port 3. The switch adds this information to the MAC address table. This table is sometimes referred to as a *forwarding database*, because it keeps track of which port packets intended for a given destination should be forwarded to. The MAC address table simply keeps a tally of which MAC addresses are reachable on each port of the switch. Suppose the MAC address for computer C is 21-76-3D-7A-F6-1E. If the switch receives a packet from port 3 with that MAC address, it would add the following entry to the MAC address table:

Port	MAC Address
3	21-76-3D-7A-F6-1E

In this way, the switch has learned that computer C is reachable via port 3.

After a short time, the switch will likely receive packets from all its ports and will associate the sender's MAC address with each port:

Port	MAC Address
1	40-20-08-78-84-52
2	2F-B6-E0-F6-EA-05
3	21-76-3D-7A-F6-1E
4	63-44-E4-A7-4F-E0
5	76-2F-F9-C8-B6-08
6	FC-78-B6-07-52-EA
7	CD-34-E4-B3-2C-76
8	1C-FD-E0-63-21-C0

It's important to keep in mind that a switch port may actually connect to more than one device. For example, suppose port 5 isn't connected to a computer but is connected to another switch, which, in turn, has three other computers connected to it. In that case, the first switch can receive packets from three different computers on port 5. Then the switch records each distinct MAC address in its MAC address table, something like this:

Port	MAC Address
1	40-20-08-78-84-52
2	2F-B6-E0-F6-EA-05
3	21-76-3D-7A-F6-1E
4	63-44-E4-A7-4F-E0
5	76-2F-F9-C8-B6-08
5	D6-4E-69-86-E9-F7
5	06-C1-15-A2-BA-60
6	FC-78-B6-07-52-EA
7	CD-34-E4-B3-2C-76
8	1C-FD-E0-63-21-C0

The process of building the MAC address table is called *learning,* and it's one of the three basic functions of a switch. The other two functions are *forwarding* and *flooding,* and they're described in the next two sections.

Forwarding

Now that you know about the MAC address table, you should have a good idea of how a switch knows which ports to forward incoming packets to: The switch simply looks up the destination MAC address in the table and sends the packet out through the corresponding port.

For example, if the switch receives a packet on port 1 intended for MAC address CD-34-E4-B3-2C-76, the switch looks up that MAC address in the table, finds that the MAC address can be reached on port 7, and forwards the packet out to port 7. This process, called *forwarding,* is the second basic function of a switch.

Switches have memory buffers associated with each port that allow the switch to store a complete packet before forwarding it to the destination port. This allows the switch to hold onto the packet for a bit if necessary before forwarding it.

For example, the destination port may be busy sending out a packet received from a different port, or the destination port may be busy receiving a packet. In either case, when the port becomes free, the switch can transmit the packet to its destination.

It's important to understand that the switch doesn't modify the packet in any way prior to sending it. What gets sent out to the destination port is an exact replica of what was received on the incoming port. When the destination device receives the packet, the device has no idea that the packet passed through the switch. In other words, no tracing information is added to the packet by the switch.

It's also important to know that, at least at this level of operation of the switch, the switch has no idea or concern for the contents of the Ethernet frame's payload. In particular, the switch is not concerned with the possibility that the payload may be an Internet Protocol (IP) packet, which, in turn, contains an IP address. Switching doesn't rely on or even know about IP addresses. Switching is a layer-2 function, and layer 2 is concerned with MAC addresses. IP addresses are a layer-3 concern and, thus, are hidden from switches.



Here's where I have to tell you that I lied. It isn't exactly true that switches don't care about IP addresses. Many advanced switches have layer-3 features that *do* look at the IP address. But when they do, they're acting more like routers than switches. *Routers* work at layer 3 and, therefore, deal with IP addresses. I have more to say about this subject later in this chapter, in the "Understanding Routers" section.

So, to recap, when a switch receives a packet on one of its ports, the switch looks in the Ethernet frame to determine the destination MAC address. The switch then looks up that address in its MAC address table, determines which port is associated with the destination address, and forwards the packet on to that port.

Which begs the question: What happens if the switch doesn't recognize the destination MAC address in the forwarding database? The answer is found in the next section.

Flooding

When a switch receives a packet that is intended for a MAC address that isn't in the switch's internal MAC address table, the switch has no way to know what port to forward the packet to. In that case, the switch has no option but to revert to acting like a hub: The switch simply forwards the packet on all available ports other than the one the packet arrived on, of course. This is called *flooding,* which is the third function of a switch (the first two being *learning* and *forwarding*).

The packet will be forwarded even to ports for which the switch has already learned a MAC address. This is necessary because a single port can be a pathway to more than one MAC address, as is the case when the port is connected to another switch.

Flooding is similar to broadcasting, but it isn't quite the same. A broadcast packet is a packet that is intended for every recipient on the network. Thus, a switch must forward broadcast packets to every port. In contrast, flooding results when the packet has a single destination, but the switch doesn't know how to reach it. Thus, the switch sends the packet to every port in the hopes that one of them will lead to the destination.

Hopefully, flooding doesn't happen too often. There's a very good chance that the destination device will receive the packet and send a reply back to the sender. In that case, the switch will record the MAC address of the recipient in its table. Then, the next time a packet intended for that destination is reached, the switch will be able to forward it to the correct port rather than flood the network again.

Looking Deeper into Switches

In the previous sections, you learned about the three basic functions of a switch:

- **Learning:** The switch learns what devices are reachable on each of its ports.
- **Forwarding:** The switch forwards incoming packets just to the correct port based on the intended destination.
- **Flooding:** The switch forwards incoming packets to all ports when it hasn't yet learned how to reach the intended destination.

In the following sections, I dig deeper into the operation of switches to explain more about how they operate.

Collision domains

One of the main benefits of switches over hubs is that switches minimize the frequency of collisions on the network. Consider a four-port switch in which computers 1, 2, 3, and 4 are connected to ports 1, 2, 3, and 4. If port 1 receives a packet from computer 1 that is intended for computer 2, the switch will forward the packet to port 2. If, at the same time, port 3 receives a packet intended for computer 4, the switch will forward that packet to port 4. Both of these packets can travel on the network at the same time because at no time will they exist on the same set of network interfaces or cables. Thus, the packets will never collide.

In contrast, if these four computers were connected with a hub, the packets would collide because the two packets would be forwarded to all the ports, not just the ports connected to the destination computers.

This reduction of collisions is so fundamental to what a switch does that a common definition of what a switch is reads like this: *A switch is a device that divides collision domains.* A *collision domain* is a segment of a network on which collisions are possible. In an old-style Ethernet network built with hubs, the entire network is a single collision domain because all the network interfaces that connect to the network will see all packets that travel on the network. But when a switch is used, the network is divided into separate collision domains.

In a switched network, each collision domain consists of just two network interfaces: the port on the switch and the port on the destination device (typically, a computer, but possibly another switch). An eight-port switch divides a single collision domain with eight devices into eight separate collision domains, each with only two devices.

Switches don't completely eliminate collisions. For example, suppose a switch has received a packet intended for a computer, and that computer attempts to send a packet at the same moment that the switch attempts to forward the received packet to the computer. In that case, the two packets collide, and both the switch and the computer must wait and try again a bit later.

Bridging

A *bridge* is a device that is very similar to a switch, but it typically has fewer ports — perhaps as few as two. The primary purpose of a bridge is to provide a link between two networks, so some bridges have just two ports. Like a switch, a bridge examines the destination MAC address of every packet it receives and forwards the packet to the other side of the bridge only if the bridge knows that the destination is on the other side.

Technically speaking, a switch is simply a multiport repeaters bridge. The distinction is mostly a historical one, because bridges were invented and widely used before switches. Before switches became inexpensive, large Ethernet networks used multiple hubs to connect computers and other devices, and a few bridges would be introduced into the network to break up large collision domains. Now that switches are common, you don't see separate bridging devices much anymore.

However, one function that a bridge can perform can come in handy: A bridge can be used to connect two different types of networks. For example, suppose your main network uses Cat5e cable, but you also have a smaller network that uses fiber-optic cable. You can use a bridge to link these two types of networks. The

bridge would have two ports: one Cat5e port and one fiber-optic port. When the bridge receives a packet on the Cat5e port, it forwards it to the fiber-optic port, and vice versa.

All switches can perform this type of bridging to connect Cat5e devices that operate at different speeds. For example, most computers have network interfaces that operate at 1 gigabit per second (Gbps). But many printers have slower, 100 megabits per second (Mbps) connections. The ports on a switch can automatically detect the speed of the device on the other end of the cable, so you can plug a 1 Gbps computer or a 100 Mbps printer into a switch port. The switch will automatically take care of buffering and forwarding packets received from the 1 Gbps devices to the slower, 100 Mbps devices.

Some switches also include ports that allow you to connect the switch to even faster networks that use 10 Gbps copper or fiber-optic cable, as described in the next section.

SFP ports and uplinks

Some switches have special ports called *small form-factor pluggable* (SFP) *ports.* You can use an SFP port to connect a variety of different types of high-speed networks, including 10 Gigabit Ethernet (GbE), which uses copper cable, or 8 Gbps Fibre Channel, which uses fiber-optic cables. In this way, the SFP ports allow the switch to bridge 100 Mbps or 1 Gbps Cat5e networks with faster copper or fiber-optic networks.

One of the most common uses of SFP ports is to connect switches to each port at speeds faster than 1 Gbps. The interconnection between two switches is often called an *uplink.* It makes sense to use high-speed uplinks because the uplink ports are likely to be the busiest ports on the switch. For example, suppose you have a network with 80 computers in which 40 of the computers are connected to one switch (switch A) and the other 40 computers are connected to a second switch (switch B). If a computer on switch A sends a packet to a computer on switch B, that packet must travel through the uplink ports to get from switch A to switch B. So, you can expect that the uplink ports will carry as much as 40 times the amount of traffic that the other ports carry.

Another common use of SFP is to connect switches to server computers. This also makes sense, because the ports that connect to your servers will carry much more traffic than the ports that connect to workstations. In order to connect a switch to a server using an SFP port, both the switch and the server must have SFP ports. So, you'll need to make sure both your servers and your switches have SFP ports.

Broadcast domains

Earlier in this chapter (in the "Understanding Switches" section), I mention that packets whose destination MAC addresses are all ones (FF-FF-FF-FF-FF-FF) are intended to be received by all devices that see the packet. Such packets are called *broadcast packets.*

The scope of the devices that broadcast packets are intended for is called the *broadcast domain.* Ordinarily, a switch forwards broadcast packets to all the ports on the switch except the port on which the broadcast packet was received. Thus, the broadcast domain consists of all the devices connected to the switch, either directly or indirectly through another switch.

In many cases, allowing broadcast packets to travel throughout a large network is not a good idea. If the network is large, broadcast packets may consume a significant amount of the total bandwidth available on the network, slowing down other, more important traffic.

You may be surprised to discover just how much broadcast traffic actually happens on a large network. The most common type of broadcast packet is an Address Resolution Protocol (ARP) request. ARP is the protocol used to determine the MAC address of a given IP address. If one IP device wants to send a packet to another IP device, the sender needs to know the MAC address of the recipient. So, the sender broadcasts an ARP request, which is essentially the question "Does anyone know the MAC address of this particular IP address? If so, please let me know."

Reducing the amount of broadcast traffic on a network is a key way to improve the network's overall performance. One of the best ways to do that is to segment the network in a way that splits up the broadcast domains. There are two ways to do this: by using routers, which are described in the next section, or by using VLANs, which are described later in this chapter, in the "Understanding VLANs" section.

Managed and unmanaged switches

Most advanced switches have management features built in to them, which means that you can monitor and configure the switch remotely, usually by logging in to a web console. To accomplish this, the switch has a small web server built into it to provide the management console. In addition, the switch itself must have an IP address.


TIP

In contrast, inexpensive consumer-grade switches that you would purchase at a retail store are usually unmanaged switches. Unmanaged switches are often appropriate for small networks, but if you have more than a few dozen computers on your network, I suggest you invest in managed switches to gain more control over your network.

With a managed switch, you can monitor traffic over the switch, which can be useful when troubleshooting network issues. In addition, you can often configure certain functions for each port of the switch. Among the most important features you can configure are VLANs, which allow you to actually create separate layer-2 networks on a single switch. I cover VLANs in greater detail later in this chapter, in the "Understanding VLANs" section.

Understanding Routers

A *router* is a layer-3 device, which means it works at the network layer of the OSI reference model. In practical terms, that means that routers know about IP addresses. At least one router is a vital component of any modern network.

A router differs from a switch in the following ways:

- **Switches work with MAC addresses and know nothing about IP addresses.** In contrast, routers work with IP addresses.
- **Routers can facilitate communication between IP networks with different subnets.** For example, if your organization has a 10.0.100.x network and a 192.168.0.x network, a router can enable packets to get from the 10.0.100.x network to the 192.168.0.x network, and vice versa. A switch can't do that. (For more about subnets, refer to Book 2, Chapter 3.)
- **Routers also enable a private network to communicate with the internet.** For example, suppose you want to connect your network to the internet via a broadband cable provider such as Comcast. The cable provider will give you a network interface that has a public IP address. You must then use a router to exchange packets from your private network to the internet via the public IP address. A switch can't do that for you.
- **Switches split up collision domains.** The segments created by switches are still part of the same broadcast domain. In contrast, routers split up broadcast domains. So, broadcast packets do not cross the boundaries created by routers. (Actually, as I explain in the "Understanding VLANs" section, later in this chapter, switches can also break up broadcast domains.)
- **Switches typically have a large number of ports — often as many as 48 in a single switch.** Routers usually have fewer ports, typically between two and eight. (However, routers for very large networks may have many more ports. For example, Cisco makes a router that can accommodate as many as 256 ports in a single chassis.)

The basic operation of a router is fairly simple. Consider the simple network depicted in Figure 3-3. Here, an organization has two separate IP networks, one using a 10.0.100.x subnet and the other using 192.168.0.x. (In both cases, the subnet mask is 255.255.255.0. Again, for more information about subnetting, refer to Book 2, Chapter 3.) A router is used to connect these two networks. On either side of the router is a switch, and each switch has just one computer connected. On the 10.0.100.x side, the computer's IP address is 10.0.100.50. On the 192.168.0.x side, the computer's IP address is 192.168.0.50. (For simplicity, I only show one computer connected to the switches on either side of the router, but in the real world there would probably be many more.)

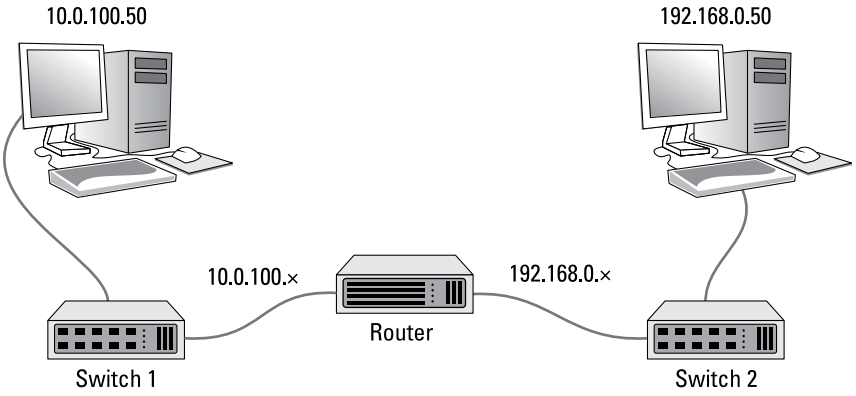


FIGURE 3-3: Two IP networks connected by a router.

Now suppose that the computer on the left side of the figure (10.0.100.50) needs to send a packet over to the computer on the right side of the figure (192.168.0.50). The sending computer forms the packet and sends it to switch 1. Switch 1, in turn, sends the packet to the router. The router examines the destination IP address and determines that the destination computer is on the 192.168.0.50 network, so it forwards the packet over to switch 2. Switch 2, in turn, forwards the packet to the destination computer.

Note that this exchange is actually considerably more complicated than the previous description lets on. For one thing, the switches — which don't know about IP addresses — must determine the MAC addresses not only of the sending and receiving computers, but also of the router. And the router must also know the MAC addresses of the two switches. You'll learn more about how this type of routing actually happens in Book 2, Chapter 4. But for now, I think you get the general idea.

The following sections describe a few of the other features commonly provided by routers.

Network address translation

When a router is used to connect a private network to the internet, one of the router's most important functions is routing traffic from all the computers on the private side of the router to the public side, which usually has just a single public IP address. To accomplish this magic, the router uses network address translation (NAT).

In short, when a computer on the private side of the network sends a packet through the router to the internet, the router substitutes its own public IP address as the sender address and keeps track of the fact that it sent a packet on behalf of a computer on the private side. When the recipient on the internet receives the packet, it sees that the sender was the router. It then sends a response back to the router, which then substitutes the original sender's private IP address for the destination address and forwards the packet to the correct computer on the private network.

For more information about NAT, see Book 2, Chapter 3.

Virtual private networks

A *virtual private network* (VPN) is a secure connection between two private networks over a public network (in other words, over the internet). All the data that flows over the VPN is encrypted, so anyone who steals packets from the VPN will find them unintelligible; only the parties on either end of the VPN are able to decrypt the packets.

VPN connections are often called *tunnels*, because they provide an isolated pathway from one point to another through the internet. The only way to gain meaningful access to a VPN tunnel is at either end.

There are two common uses for VPNs:

- **To provide remote workers with secure access to your company network:** To do that, you set up a VPN on the router, and then provide your remote workers with the credentials necessary to access the VPN. The remote workers can run a software VPN client on their home computers or laptops to connect to your company network.
- **To establish a tunnel directly between routers on two networks that are separated geographically:** For example, suppose you have offices in Los Angeles and Las Vegas. You can use routers on both networks to establish a VPN tunnel between them. This effectively joins the networks together, so that

devices on the Los Angeles network can freely exchange packets with devices on the Las Vegas network, and vice versa.

Figure 3-4 shows this arrangement. As you can see, the routers in both Los Angeles and Las Vegas are connected through the internet via a VPN tunnel. This tunnel enables computers in Los Angeles and Las Vegas to communicate freely and securely with each other.

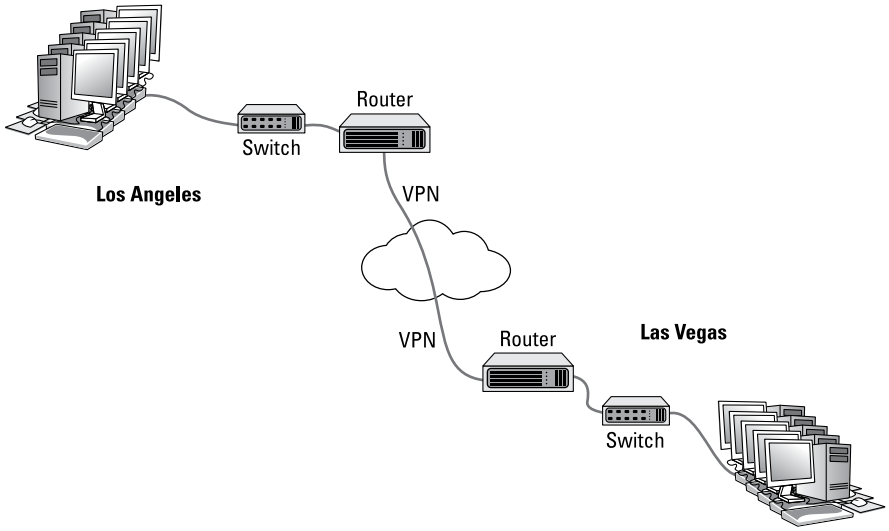


FIGURE 3-4: Connecting offices with a VPN tunnel.

For more information about working with VPN tunnels, refer to Book 4, Chapter 6.

Understanding VLANs

The final topic for this whirlwind introduction to switches and routers is the concept of VLANs. Most advanced switches allow you to create VLANs.

As its name suggests, a VLAN is a virtual network that runs on top of your actual physical network. VLANs work at layer 2 of the OSI model, which means that they're related to MAC addresses, not IP addresses. That said, there is usually a direct correlation between VLANs and IP subnets. If (or when) your network grows large enough that you want to set up two or more subnets to better manage it, you'll probably also want to set up two or more VLANs, one for each of your subnets.

A VLAN can divide a single switch into two virtual switches that behave exactly as if they were separate switches. This means the following:

- If a port on one VLAN receives a packet intended for a destination on the same VLAN, the switch forwards the packet to the destination port, the same as if VLANs were not in use.
- When a port on one VLAN receives a packet intended for a destination on the same VLAN that the switch has not yet learned, the switch will flood only those ports that are on the destination VLAN — not all the ports on the switch. Thus, VLANs can reduce traffic caused by flooding.
- When a broadcast packet is received, the switch will forward the packet only to those ports that are on the same VLAN. In other words, VLANs can break up broadcast domains in the same way that a router can.
- If a port on one VLAN receives a packet intended for a different VLAN, a router is required to link the networks. That's because separate VLANs are, for all intents and purposes, separate networks.

  That being said, most switches that support VLANs also support trunk ports, which can switch traffic between VLANs. A *trunk port* is a port that can handle traffic for two or more VLANs.

To use VLANs, you must manually configure each port of your switches to operate on the appropriate VLAN. By default, all switches regardless of manufacturer are configured out of the box so that all ports operate on a VLAN named VLAN1. To create a new VLAN, you simply create a name for the new VLAN and then configure the ports that will talk on the new VLAN.

In VLAN terminology, a port that is configured to operate on a single VLAN is called an *access port.* Ports that are configured to work on more than one VLAN are called *trunk ports.* By default, all switch ports are configured as access ports on VLAN1.

Note that if you have more than one switch in your network, you can configure VLANs to work across the switches. For example, you can create a VLAN for your company's accounting department — let's call it VLAN-Acct. Then you can configure ports on any of your switches as access ports on VLAN-Acct. In this way, your entire accounting staff can operate on the accounting VLAN.

For more information about working with VLANs, refer to Book 2, Chapter 1 and Book 3, Chapter 1.

IN THIS CHAPTER

- Assessing the risk for security
- Looking at two pillars of cybersecurity
- Identifying the most important protection and recovery measures
- Examining standardized cybersecurity frameworks
- Looking more closely at the NIST Cybersecurity Framework

Chapter 4
Cybersecurity

As an IT professional, cybersecurity is the thing most likely to keep you awake at night. Consider the following scenarios:

- Your phone starts ringing like crazy at three o'clock one afternoon because no one anywhere on the network can access any of their files. You soon discover that your network has been infiltrated by *ransomware,* nefarious software that has encrypted every byte of data on your network, rendering it useless to your users until you pay a ransom to recover the data.
- Your company becomes a headline on CNN because a security breach has resulted in the theft of your customers' credit card information.
- On their last day of work, a disgruntled employee copies your company contact list and other vital intellectual property to a flash drive and walks away with it along with their red Swingline stapler. A few months later, your company loses its biggest contract to the company where this jerk now works.

There is no way you can absolutely prevent such scenarios from ever happening, but with proper security, you can greatly reduce their likelihood. This chapter presents a brief overview of some of the basic principles of securing your network.

Cybersecurity goes hand in hand with networking. In fact, the moment you think of building a network, you should lay the groundwork for how you'll keep it secure. You should consider the security aspects of a network from the very start and throughout the design and implementation of your network. Security will touch every aspect of your network environment — not just network equipment such as firewalls and switches, but also servers, end-user computers, user accounts, data storage, and so on.

But We're a Small Business — Do We Need Security?

It's tempting to think that cybersecurity is important only to large enterprises. In a small business, everyone knows and trusts everyone else. Folks don't lock up their desks when they take a coffee break, and although everyone knows where the petty cashbox is, money never disappears.

Cybersecurity isn't necessary in an idyllic setting like this one — or is it? You bet it is. Here's why any network should be set up with built-in concern for security:

- **Mitts off:** Even in the friendliest office environment, some information is and should be confidential. If this information is stored on the network, you want to store it in a directory that's available only to authorized users.
- **Hmm:** Not all security breaches are malicious. A network user may be routinely scanning files and come across a filename that isn't familiar. The user may then call up the file, only to discover that it contains confidential personnel information, juicy office gossip, or your résumé. Curiosity, rather than malice, is often the source of security breaches.
- **Trust:** Sure, everyone at the office is trustworthy *now.* But what if someone becomes disgruntled, a screw pops loose, and they decide to trash the network files before jumping out the window? Or what if someone decides to print a few $1,000 checks before packing off to Tahiti?
- **Temptation:** Sometimes the mere opportunity for fraud or theft can be too much for some people to resist. Give people free access to the payroll files, and they may decide to vote themselves a raise when no one is looking.

  If you think that your network contains no data worth stealing, think again. For example, your personnel records probably contain more than enough information for an identity thief: names, addresses, phone numbers, Social Security numbers, and so on. Also, your customer files may contain your customers' credit card numbers.

- **Malice:** Hackers who break into your network may not be interested in stealing your data. Instead, they may be looking to plant a *Trojan horse* program on your server, which enables them to use your server for their own purposes. For example, someone may use your server to send thousands of unsolicited spam email messages. The spam won't be traced back to the hackers; it will be traced back to you.
- **Whoops:** Bear in mind that not everyone on the network knows enough about how your operating system and the network work to be trusted with full access to your network's data and systems. One careless mouse click can wipe out an entire directory of network files. One of the best reasons for activating your network's security features is to protect the network from mistakes made by users who don't know what they're doing.

The Two Pillars of Cybersecurity

There are two basic elements that you must consider as part of your cybersecurity plan:

- **Prevention:** The first pillar of cybersecurity is the tools and technology that you can deploy to prevent bad actors from penetrating your network and stealing or damaging your data. This pillar includes firewalls that block unwelcome access, antivirus programs that detect malicious software, patch management tools that keep your software up-to-date, and anti-spam programs that keep suspicious email from reaching your users' inboxes.
- **Recovery:** The second pillar of cybersecurity is necessary because the first pillar isn't always successful. Successful cyberattacks are inevitable, so you need to have technology and plans in place to quickly recover from them when they hit. This pillar includes such things as creating backup copies of all your data and having recovery plans in place to quickly get your organization back up and running.

I cover both of these pillars in greater detail in the following sections.

Prevention

A comprehensive cybersecurity plan is filled with prevention measures.

First and foremost, your prevention measures should start with a complete understanding of your IT environment, the threats it's exposed to, and the vulnerabilities

it presents to would-be attackers. The foundation of this knowledge is an *asset management system* that lets you keep track of absolutely everything that's connected to your network. This inventory includes at least the following:

- **All the hardware connected to your network:** That includes all the desktop computers, mobile devices, servers, switches, Wi-Fi access points, routers, printers, and every other piece of hardware connected to your network.
- **All the software connected to your network:** That includes operating systems, web browsers, Microsoft Office applications, and any other programs your organization uses. It also includes cloud service providers such as Microsoft 365, online meeting platforms, cloud storage providers, and so on. Finally, it includes the software that runs on devices such as routers, switches, printers, and other similar devices.
- **All the people connected to your network, typically represented by Active Directory accounts:** You need to understand who they are, what their jobs are, what permissions they require, and what devices they use.

With the information gleaned from this asset management, you can deploy specific preventive measures to protect each asset. The following list is not complete, but it's a good starting point:

- **Firewalls:** Your internet connection must be protected by a firewall device that's configured to keep dangerous traffic out of your network. (For more information, see Book 10, Chapter 2.)
- **Wi-Fi security:** All wireless access to your network must be encrypted and protected by password access. (For more information, see Book 4, Chapter 2.)
- **Antivirus software:** Every computer on your network must be protected by active antivirus software. That includes *every* computer — workstations, laptops, tablets, and servers. All it takes is one unprotected computer to expose your entire environment to attack. (For more information, see Book 10, Chapter 2.)
- **Anti-spam software:** Most cyberattacks come in through email. Make sure all email is protected by anti-spam software that can block email that contains malicious code or suspicious links. (For more information, see Book 10, Chapter 3.)
- **Strong passwords:** All accounts that have access to your systems should be secured by strong passwords. (For more information, see Book 10, Chapter 1.)
- **Multifactor authentication:** The most critical access, such as for those with administrative control, should be controlled by multifactor authentication, which requires additional verification beyond a username and password. (For more information, see Book 10, Chapter 1.)

- **Data protection:** All shared data on your network should be protected with roll-based security so that only those users who have a demonstrated need for the data are allowed access. This is done by controlling access permissions on files and folders, as well as share permissions. (For more information, see Book 6, Chapter 5.)
- **Encryption:** *Encryption* refers to the process of encoding data so that it can be read only by those who possess the secret encryption key. Encryption is one of the most important aspects of data security and should be employed whenever possible.

  One common way to use encryption is on wireless networks, where all data should be encrypted. This type of encryption is called *data-in-flight encryption* because it encrypts data while it's in transit from one computer or device to another. It's also common to encrypt data that resides on disk drives — this type of encryption is called *data-at-rest encryption* and is especially important if someone physically steals your disk drives (or the computers that contain them).
- **User life-cycle management:** All user accounts should be subject to a documented life-cycle management policy that ensures that when a user leaves the organization, that user's access is terminated.
- **Auditing:** All aspects of your security environment should be regularly audited to ensure everything is operating as expected and is appropriate for the current environment. This includes regularly reviewing your user accounts and file permissions; reviewing firewall, antivirus, and anti-spam software to make sure it's functioning; and reviewing event logs.
- **User training:** The weakest points in any network are its users. Make sure to regularly offer security training for your users. (For more information, see Book 10, Chapter 1.)
- **Physical security:** This aspect of cybersecurity is often overlooked. Any hacker worth their salt can quickly defeat all but the most paranoid security measures if they can gain physical access to a computer on your network. Make sure the server room is locked at all times. Make sure your users lock their computers when they step away from their desks.

Recovery

No matter how good your prevention measures are, cybersecurity events are bound to happen. A user will exercise bad judgement and click a link in a phishing email, an important security patch will be neglected and an intruder will exploit the resulting weakness, or someone's password will be compromised. It's bound to happen, so your cybersecurity plan must include recovery measures in addition to prevention measures.

A recovery plan should also protect you against threats that aren't necessarily malicious. For example, what if a hardware failure takes out a key file server and you lose all its data? Or what if there's a fire in the server room? Disasters like these are unlikely but not impossible. For more information about disaster recovery planning, check out Book 10, Chapter 4.


TIP

The most important aspect of recovery is to plan for it in advance. Don't wait until after a cyberattack has succeeded to start wondering how you can recover. Instead, assume that a cyberattack *will* eventually happen, and plan in advance how you'll recover.

The basis of any recovery plan is a good backup plan. In fact, planning for backup is an integral part of planning any network. I've devoted Book 3, Chapter 6 to this topic, so I won't go into every detail here. But for now, know that backups must be:

- **Comprehensive:** Identify every critical server and data store in your organization, and make sure it's backed up regularly.
- **Up-to-date:** When you're forced to recover from a backup, you'll be rolling your business back to the date the backup was made. If that was three weeks ago, you'll lose three weeks' worth of work.
- **Redundant:** Keep multiple copies of your backups, each representing a different recovery point. At the minimum, keep at least three generations of backups. That way, if the most recent set of backups doesn't work, you can revert to the set before that and, if necessary, the set before that. A key factor to consider is that if your files have been corrupted by a cyberattack and you don't discover the attack right away, your backups may contain copies of the corrupted data. You want to make sure that you have a good backup that was made *before* the attack occurred.
- **Kept off-site:** If a fire burns down your server room and your backups are kept on a shelf next to the servers, you'll lose the backups, too. At that point, you won't be able to restore anything.
- **Offline:** It's not enough to keep backups off-site; they must also be offline. Backing up to the cloud is popular, but keep in mind that a hacker skilled enough to break into your network and delete files on your servers may also be skilled enough to delete your cloud backups as well.
- **Automated:** Don't rely on remembering to run a backup every Friday at the end of the day. You'll forget. Make sure your backup processes are automated.
- **Monitored:** Don't assume backups worked this week just because they worked last week. Monitor your backups regularly to ensure they're working as designed.

- **Tested:** Don't wait until the pressure of a recovery to see if your backups actually work. Regularly test them by restoring individual files and entire servers.

Here are a few other elements your recovery plan should include:

- **Spare computers:** If a cyberattack compromises one of your desktop computers, make sure you have a spare or two that you can quickly configure to quickly get the user back to work.
- **Emergency disk capacity:** Restore operations often require that you have plenty of spare disk capacity available so that you can move data around. Inexpensive network-attached storage (NAS; see Book 3, Chapter 5) may fill the bill, but keep in mind that this type of storage is very slow. If you rely on it, you may find that it takes several days to recover multiple terabytes of data.
- **Communications:** In the midst of a recovery from a cyberattack, communicating with your users is critical. They'll need to know what's going on, how long you expect the recovery to take, and so on. Unfortunately, this communication may be difficult if the normal channels of communication — such as email — have been disrupted by the attack. So, you should plan in advance for alternative methods of communicating with users, such as cloud-based communication platforms like Microsoft Teams or Slack.

Cybersecurity Frameworks

You may think that all you need to do to secure your network is install a firewall, run antivirus software on all your computers, and back up all your data. Those are important first steps, but cybersecurity is much bigger than a checklist of things to do.

In fact, cybersecurity should be baked into your IT systems from the ground up. Every aspect of your system designs should take cybersecurity into account, not as an afterthought but from the very beginning. That includes your servers, storage platforms, desktop computers, network infrastructure (including switches, routers, firewalls, cables, and wireless networks), mobile devices, operating systems, software, and anything else that's part of your IT environment.

It's a daunting task, but fortunately you're not alone in figuring out how to make cybersecurity a top priority in your IT organization. Plenty of resources are available to you — including standardized frameworks that can help you plan and implement your security environment.

You have plenty of cybersecurity frameworks to choose from. Although most of these frameworks are similar, there are subtle differences. Here are five of the most popular cybersecurity frameworks you may want to investigate:

- **NIST:** The NIST Cybersecurity Framework is probably the most commonly used framework in the United States. It's governed by the National Institute of Standards and Technology (NIST). (For more information about this popular framework, refer to the next section, "The NIST Cybersecurity Framework.")
- **ISO/IEC 27001:** This is the most popular international cybersecurity framework. For more information, go to `www.iso.org/standard/27001`.
- **ISA/IEC 62443:** The International Society of Automation (ISA) sponsors a series of standards known as ISA/IEC 62443, which comprise a flexible framework for managing security. For more information, head to `www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards`.
- **CIS Critical Security Controls:** The Center for Internet Security (CIS) provides a list of 18 cybersecurity controls that can be used as a framework for organizing your cybersecurity measures. For more information, go to `www.cisecurity.org/controls/cis-controls-list`.
- **Control Objectives for Information and Related Technologies (COBIT):** Sponsored by the Information Systems Audit and Control Association (ISACA), COBIT is one of the more popular cybersecurity frameworks. For more information, head to `www.isaca.org/resources/cobit`.

The NIST Cybersecurity Framework

In 2014, NIST issued the first version of its cybersecurity framework, officially known as the Framework for Improving Critical Infrastructure Cybersecurity but commonly referred to as the NIST Cybersecurity Framework (and often when speaking in the context of cybersecurity simply NIST). I refer to it simply as "the framework" throughout the rest of this chapter.

The framework was originally intended to apply to critical infrastructure such as the power grid, transportation systems, dams, government agencies, and so on. But it quickly became popular in the private sector as well and is now considered one of the best overall tools for planning cybersecurity for large and small organizations, public and private.

The framework is useful for any organization large enough to have a dedicated IT staff, even if that staff consists of just one person. No organization can or should implement every detail that is spelled out in the framework. Instead, the framework invites you to develop a solid understanding of the cybersecurity risks your organization faces and to implement a risk management strategy based on informed decisions about which security practices make sense for your organization.

In 2018, NIST issued a new version of the framework, known as version 1.1. This version added a section on self-assessment and greatly expanded its coverage of the cybersecurity risk associated with business supply chains.

The current version of the framework, known as version 2.0, was released in February 2024. It added additional information about how an organization can provide oversite for cybersecurity by formally adopting a governance function.

You can find the complete documentation for the NIST Cybersecurity Framework at `www.nist.gov/cyberframework` strongly suggest you download the framework document, print it out, and read it. It's only about 32 pages.

The framework consists of three basic components:

- **Cybersecurity Framework Core:** This section identifies six basic functions of cybersecurity:
  - **Govern:** This function provides a formalized method to ensure that the other five functions are properly planned, implemented, and monitored.
  - **Identify:** This function helps you determine, in detail, exactly what parts of your organization are vulnerable to cyberattack.
  - **Protect:** This function enables you to take specific steps to protect those parts of your organization that you've identified as being vulnerable.
  - **Detect:** This function involves monitoring your systems and environment so that you know as soon as possible when a cyberattack occurs.
  - **Respond:** This function helps you plan in advance how you'll respond when a cybersecurity incident occurs.
  - **Recover:** This function prompts you to develop plans and procedures to restore any parts of your environment that were damaged by a cyberattack. For example, if data was lost, you may need to restore the lost data from backup copies.

Within each of these six basic functions, best practices, guidelines, and standards are presented focusing on specific cybersecurity outcomes, such as "Remote access is managed" or "Removable media is protected and its use restricted according to policy." I offer more detail on the Cybersecurity Framework Core later in this section.

- **Cybersecurity Framework Organizational Profiles:** This section discusses the use of profiles to indicate which specific outcomes in the Cybersecurity Framework Core are implemented. You can create a *current profile,* which documents the current cybersecurity practices at your organization, and then create a *target profile* to represent where you'd like to be. Then you can devise a plan to move from the current profile to the target profile.
- **Cybersecurity Framework Tiers:** This section describes four distinct tiers that represent an increasing level of sophistication in cybersecurity practices. As an organization invests more in cybersecurity, it moves up through the tier levels.

Each of the six functions of the Cybersecurity Framework Core (listed earlier) is divided into several categories, which are in turn divided into subcategories. A simple numbering scheme is used to track the functions, categories, and subcategories. For example, the Identify function is designated by the identifier *ID.* Its first category is Asset Management, which is designated by *ID.AM.* The first subcategory under Asset Management is "Inventories of hardware managed by the organization are maintained," and it's designated *ID.AM-01.* Table 4-1 lists the six functions along with each function's categories and the identifier for each.

TABLE 4-1 The Functions and Categories of the NIST Framework Core

Function	Category	Identifier
Govern (GV)	Organizational Context	GV.OC
	Risk Management Strategy	GV.RM
	Roles, Responsibilities, and Authorities	GV.RR
	Policy	GV.PO
	Oversite	GV.OV
	Cybersecurity Supply Chain Management	GV.SC
Identify (ID)	Asset Management	ID.AM
	Risk Assessment	ID.RA
	Improvement	ID.IM

Function	Category	Identifier
Protect (PR)	Identity Management, Authentication, and Access Control	PR.AA
	Awareness and Training	PR.AT
	Data Security	PR.DS
	Platform Security	PR.PS
	Technology Infrastructure Resilience	PR.IR
Detect (DE)	Continuous Monitoring	DE.CM
	Adverse Event Analysis	DE.AE
Respond (RS)	Incident Management	RS.MA
	Incident Analysis	RS.AN
	Incident Response Reporting and Communication	RS.CO
	Incident Mitigation	RS.MI
Recover (RC)	Incident Recovery Plan Execution	RC.RP
	Incident Recovery Communication	RC.CO

In all, there are 22 categories across the six functions. Each of these categories is broken down into from 2 to 10 subcategories, for a total of 106 subcategories altogether.

The framework doesn't prescribe specific solutions for each of the 106 subcategories; it merely states the outcome to be achieved by each subcategory and invites you to design a solution that produces the desired outcome. For example, the first subcategory of Asset Management (ID.AM-01) is as follows:

> Inventories of hardware managed by the organization are maintained

You can accomplish this goal in many ways. If your organization is small, you may just keep track of all your computer and network devices in a simple Microsoft Excel spreadsheet. If your organization is larger, you may use software that automatically scans your network to create a catalog of all attached devices, and you may want to use inventory tags with barcodes so you can track hardware assets. But one way or another, keeping an inventory of all your physical devices and systems is a vital element of cybersecurity.



REMEMBER

Although the framework doesn't prescribe specific solutions, it does offer a set of links to other cybersecurity frameworks, which it calls *Informative References* (`www.nist.gov/informative-references`). You can cross-reference these Information References to gain additional insight into each of the subcategories.

IN THIS CHAPTER

- Understanding what network operating systems do
- Examining what makes a good server
- Looking at the different packaging options for servers
- Taking a quick look at virtualization

Chapter 5
Servers and Virtualization

Servers are the lifeblood of any network. They provide the shared resources that network users crave, such as file storage, databases, email, web services, and so on. Choosing which servers your network needs and selecting the type of equipment you use to implement your servers are among the key decisions you'll make when you set up a network.

In this chapter, I take a quick look at what's important in a server. First, I cover the basic functions of a server operating system. Then I survey the various types of servers most networks need. Then I turn my attention to important matters to consider when selecting the kind of hardware a server should run on. And finally, I look quickly at the idea of virtualizing your entire server environment.

Understanding Network Operating Systems

The server operating system is what enables your server computers to function as servers rather than as ordinary Windows clients. Server operating systems provide essential functions such as providing basic security services, sharing disk

storage and printers, and so on. The following sections cover some of the core functions of a server operating system.

Network services

Obviously, a server operating system must provide networking capabilities in order for it to function on a network. If your client computers can't connect to your servers, your network will be useless. For this reason, it's a good idea to make sure your server computers are equipped with more than one network interface. That way, if one of the interfaces fails, the other can pick up the slack and keep your server connected to your network.

In addition to basic network connectivity, one of your servers will typically be responsible for providing some essential software services that are required to keep a network operating in an efficient manner. One of these is called Dynamic Host Configuration Protocol (DHCP); it's the service that recognizes computers and other devices that want to join the network, providing each with a unique address so that all the devices on the network can identify one another. For more information about this vital service, refer to Book 2, Chapter 5.

A second basic service that is provided by one of the servers on your network is called Domain Name System (DNS). This service is what enables people to use network names instead of the actual addresses that are handed out by DHCP. It's also the service that enables people to browse the World Wide Web using addresses such as `www.amazon.com` rather than cryptic addresses like 54.239.28.85. For more information about this important service, refer to Book 2, Chapter 6.

File-sharing services

One of the most important functions of a server operating system is to share resources with other network users. The most common resource that's shared is the server's *file system* (organized disk space that a server must be able to share, in whole or in part, with other users). In effect, users can treat the server's disk space as an extension of their own computers' disk space.

The server operating system allows the system administrator to determine which portions of the server's file system to share.


TIP

Although an entire hard drive can be shared, it isn't commonly done. Instead, individual folders are shared. The administrator can control which users are allowed to access each shared folder.

Because file sharing is the reason many network servers exist, server operating systems have more sophisticated disk management features than are found in desktop operating systems. For example, most server operating systems can manage two or more hard drives as though they were a single drive. In addition, most can create a *mirror* (an automatic backup copy of a drive) on a second drive.

Windows server operating systems also provide a feature called *Distributed File System* (DFS). DFS lets you spread your file shares across multiple servers while creating a naming scheme that lets you access the data without referencing a specific server. This makes it much easier to reorganize your file shares to accommodate growing file storage needs. For more about DFS, refer to Book 6, Chapter 5.

Multitasking

Only one user at a time uses a desktop computer; however, multiple users simultaneously use server computers. As a result, a server operating system must provide support for multiple users who access the server remotely via the network.

At the heart of multiuser support is *multitasking,* which is the capability of an operating system to execute more than one program (a task or a process) at a time. Multitasking operating systems are like the guy who used to spin plates balanced on sticks on the old *Ed Sullivan Show* back in the 1950s. He'd run from plate to plate, trying to keep them all spinning so they wouldn't fall off the sticks — and just for grins, he was blindfolded or riding on a unicycle.

Although multitasking creates the appearance that two or more programs are executing on the computer at one time, in reality, a computer with a single processor can execute only one program at a time. The operating system switches the central processing unit (CPU) from one program to another to create the appearance that several programs are executing simultaneously, but at any given moment, only one of the programs is actually executing; the others are patiently waiting for their turns. (However, if the computer has more than one CPU, the CPUs *can* execute programs simultaneously, which is called *multiprocessing.*)

For multitasking to work reliably, the server operating system must completely isolate the executing programs from each other. Otherwise, one program may perform an operation that adversely affects another program. Multitasking operating systems do this by providing each task with its own unique address space that makes it almost impossible for one task to affect memory that belongs to another task.

Directory services

Directories are everywhere — and they were even in the days when they were all hard copies. When you needed to make a phone call, you looked up the number in a phone directory. When you needed to find the address of a client, you looked them up in your Rolodex. And then there were the nonbook versions: When you needed to find the Sam Goody store at a shopping mall, you looked for the mall directory — usually, a lighted sign showing what was where.

Networks have directories, too, providing information about the resources that are available on the network: users, computers, printers, shared folders, and files. Directories are essential parts of any server operating system.

The most popular modern directory service is called *Active Directory.* Active Directory is a standard component of all Windows operating systems, and because it's so popular, most other operating systems support it as well. Active Directory is a database that organizes information about a network and all its computers and users. It's simple enough to use for networks with just a few computers and users, but powerful enough to work with large networks containing tens of thousands of computers and users. Figure 5-1 shows the Active Directory Users and Computers tool, which manages Active Directory user and computer accounts on Windows Server 2025.

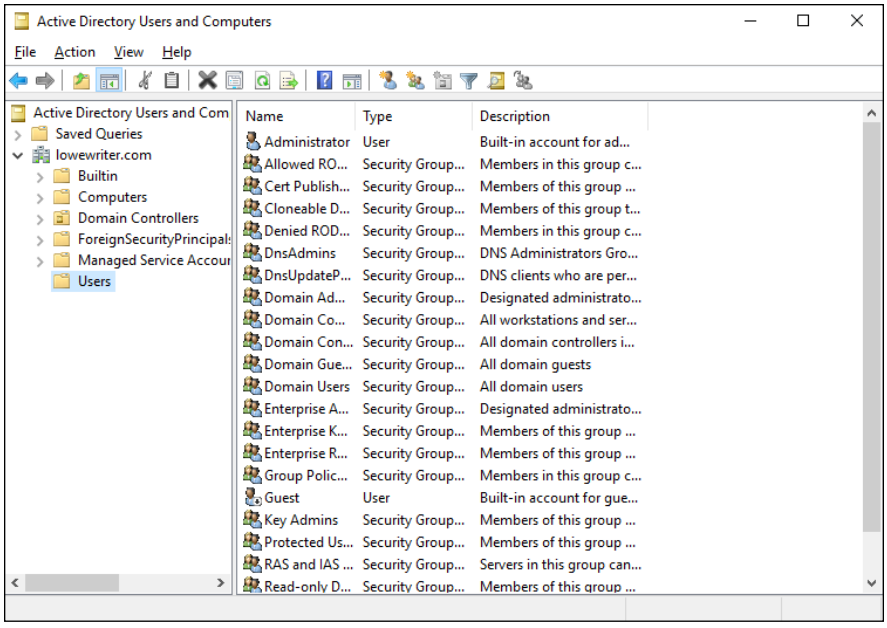


FIGURE 5-1: Managing Active Directory users and computers.

Security services

All server operating systems must provide some measure of security to protect the network from unauthorized access. Hacking seems to be the national pastime these days. With most computer networks connected to the internet, anyone anywhere in the world can — and probably will — try to break into your network.

The most basic type of security is handled through *user accounts*, which grant individual users the right to access the network resources and govern which resources the user can access. User accounts are secured by passwords; therefore, good password policy is a cornerstone of any security system. Most server operating systems give you some standard tools for maintaining network security:

- **Establish password policies.** For example, you can mandate that passwords have a minimum length and include a mix of letters and numerals.
- **Set passwords to expire after a certain number of days.** Network users must change their passwords frequently.
- **Encrypt network data.** A data-encryption capability scrambles data before it's sent over the network or saved on disk, making unauthorized use a lot more difficult.

  Good encryption is the key to setting up a virtual private network (VPN), which enables network users to securely access a network from a remote location by using an internet connection.
- **Manage digital certificates.** Digital certificates are used to ensure that users are who they say they are and files are what they claim to be.


TIP

The overwhelming majority of business networks rely on server versions of Windows, known as Windows Server. Microsoft periodically releases updated versions of Windows Server, so Windows Server is frequently improved, and older versions are occasionally rendered obsolete. Currently, the most commonly used versions are Windows Server 2019 and Windows Server 2012. The newest version is known as Windows Server 2025; you'll learn about Windows Server 2025 throughout this entire book, but Book 6 dives deeply into setting it up and configuring it.

Windows Server is not the only server operating system at your disposal. Many servers — especially those whose primary responsibility is to host websites — use Linux instead of Windows Server. You can find out more about Linux in Book 8.

Knowing What's Important in a Server

The following sections point out some general things to keep in mind when selecting the equipment that a server should run on.

Scalability

Scalability is the ability to increase the size and capacity of the server computer without unreasonable hassle. Purchasing a server computer that just meets your current needs is a major mistake because, rest assured, your needs will double within a year. If at all possible, equip your servers with far more disk space, random access memory (RAM), and processor power than you currently need.

Reliability

The old adage "You get what you pay for" applies especially well to server computers. Why spend $5,000 on a server computer when you can buy one with seemingly similar specifications at a discount electronics store for a mere $1,000? The main reason: reliability. When a client computer fails, only the person who uses that computer is affected. When a server fails, however, everyone on the network is affected. The less-expensive computer is probably made of inferior components that are more likely to fail, and it doesn't have redundant components built in. For example, many server computers have two power supplies, two CPUs, two or more network interfaces, and other redundant components.

Availability

Availability is closely related to reliability. When a server computer fails, how long does it take to correct the problem and get the server up and running again? Server computers are designed so their components can be easily diagnosed and replaced, which minimizes the downtime that results when a component fails. In some servers, components are *hot swappable* (certain components can be replaced without shutting down the server). Some servers are fault tolerant so that they can continue to operate even if a major component fails.

Service and support

Service and support are often overlooked factors when picking computers. If a component in a server computer fails, do you have someone on-site qualified to repair the broken computer? If not, you should get an on-site maintenance contract for the computer.



Don't settle for a maintenance contract that requires you to take the computer in to a repair shop or, worse, mail it to a repair facility. You can't afford to be without your server that long. Get a maintenance contract that provides for on-site service and repair of your server 24 hours a day, seven days a week.

Identifying the Components of a Server Computer

The hardware components that make up a typical server computer are similar to the components used in less-expensive client computers. However, server computers are usually built from higher-grade components than client computers for the reasons given in the "Knowing What's Important in a Server" section, earlier in this chapter. The following sections describe the typical components of a server computer.

Motherboard

A motherboard is the computer's main electronic circuit board to which all the other components of your computer are connected. More than any other component, the motherboard *is* the computer. All other components attach to the motherboard.

The major components on the motherboard include the CPU, supporting circuitry (the chipset), RAM, expansion slots, a hard drive controller, USB ports for devices such as keyboards and mice, a graphics adapter, and one or more network interfaces.

Processor

The CPU is the brain of the computer. Although the processor isn't the only component that affects overall system performance, it's the one that most people think of first when deciding what type of server to purchase. At the time of this writing, most servers used one of several variations of Intel's Xeon processor. These processors are designed specifically for server computers rather than client computers; they offer anywhere from 4 to 22 independent processor cores, depending on the model.

Each motherboard is designed to support a particular type of processor. CPUs come in two basic mounting styles: slot or socket. However, you can choose from several types of slots and sockets, so you have to make sure that the motherboard supports the specific slot or socket style used by the CPU. Some server motherboards have two or more slots or sockets to hold two or more CPUs.



Clock speed refers to how fast the basic clock that drives the processor's operation ticks. In theory, the faster the clock speed, the faster the processor. However, clock speed alone is reliable only for comparing processors within the same family. What matters more in a server is the number of processor cores. The more cores the server has, the more tasks the server can perform simultaneously. Because servers are in the business of supporting many clients, being able to do many tasks simultaneously is a huge benefit for server performance.

What's more, processor cores utilize a technology called *hyperthreading,* which effectively lets each processor core juggle two threads at once. (In general terms, a *thread* is a sequence of instructions that performs a single task.) Because each core can handle two simultaneous threads, a processor with four cores can handle eight concurrent threads.

Many server motherboards can support two separate processors, which doubles the potential workload of the server. For example, if the server has two 14-core processors, the server has a total of 28 cores available for its workload. Because of hyperthreading, each of these 28 cores can handle 2 threads, so the server can handle 56 concurrent threads.

Memory

Don't scrimp on memory. People rarely complain about servers having too much memory. The total memory capacity of the server depends on the motherboard. It isn't unusual to see servers configured with anywhere from 32GB to 512GB of RAM.

Hard drives

Most desktop computers use inexpensive consumer-grade SATA hard drives, which are adequate for individual users. Because of their low cost, SATA drives are sometimes also used in inexpensive servers. But because performance and reliability are important in servers, most servers rely on faster and more reliable SCSI or SAS disk drives instead. For the best performance, solid-state drives (SSDs) can be used; these drives have no mechanical parts, so they're considerably faster than traditional spinning disks.

Network interfaces

The network connection is one of the most important parts of any server. Ideally, your server should have at least two network interfaces. Additional network interfaces not only improve the performance of your server, but also make it more reliable. If one of the network interfaces fails, the other(s) can pick up the ball.

If possible, the server's network interfaces should be 10 gigabit per second (Gbps) interfaces. Then, you can use 10 Gbps switches to connect the servers to each other and to your access switches. With many users contending for access to the servers simultaneously, 1 Gbps interfaces can easily become a performance-limiting bottleneck.

Video

Fancy graphics aren't that important for a server computer. You don't need to equip your server with an expensive video card; the video interface that's built in to the motherboard will suffice. (This is one of the few areas where cutting costs on a server is okay.)

Power supply

Because a server usually has more devices than a typical desktop computer, it requires a larger power supply (typically 600 watts). If the server houses a large number of hard drives, it may require an even larger power supply.

Because the power supply is one of the most likely components to fail, many server computers have two built-in power supplies for redundancy. That way, if one of the power supplies fails, the other can pick up the load and keep the server running.

Considering Server Form Factors

Form factor refers to the size, shape, and packaging of a hardware device. Server computers typically come in one of three form factors: tower cases, rack-mounted servers, or blade servers. You can also get servers in very small packages; these are known as *tiny servers.*

Tower cases

Most servers are housed in a traditional tower case, similar to the tower cases used for desktop computers. A typical server tower case is 18 inches high, 20 inches deep, and 9 inches wide, with room inside for a motherboard, five or more hard drives, and other components.

Some server cases include advanced features specially designed for servers, such as redundant power supplies (so both servers can continue operating if one of the power supplies fails), hot-swappable fans, and hot-swappable disk-drive bays. (*Hot-swappable* components can be replaced without powering down the server.)

Rack-mounted servers

If you need only a few servers, tower cases are fine. You can just place the servers next to each other on a table or in a cabinet that's specially designed to hold servers. If you need more than a few servers, though, space can quickly become an issue. For example, what if your departmental network requires a bank of ten file servers? You'd need a pretty long table!

Rack-mounted servers are designed to save space when you need more than a few servers in a confined area. A rack-mounted server is housed in a small chassis that's designed to fit into a standard 19-inch equipment rack. The rack allows you to vertically stack servers to save space.

Blade servers

Blade servers are designed to save even more space than rack-mount servers. A *blade server* is a server on a single card that can be mounted alongside other blade servers in a blade chassis, which itself fits into a standard 19-inch equipment rack. A typical blade chassis holds six or more servers, depending on the manufacturer.

One of the key benefits of using blade servers is that you don't need a separate power supply for each server. Instead, the blade enclosure provides power for all its blade servers. Some blade server systems provide rack-mounted power supplies that can serve several blade enclosures mounted in a single rack.

In addition, the blade enclosure provides keyboard, video, and mouse (KVM) switching so that you don't have to use a separate KVM switch. You can control any of the servers in a blade server network from a single keyboard, monitor, and mouse. (For more information, see the sidebar, "Saving space with a KVM switch.")

SAVING SPACE WITH A KVM SWITCH

If you have more than two or three servers in one location, consider getting a KVM switch to save space by connecting several server computers to a single keyboard, monitor, and mouse. Then you can control any of the servers from a single keyboard, monitor, and mouse by turning a dial or by pressing a button on the KVM switch.

Simple KVM switches are mechanical and allow you to choose from 2 to 16 (or more) computers. More elaborate KVM switches can control more computers, using a pop-up menu or a special keyboard combination to switch among computers. Some advanced KVMs can even control a mix of PCs and Mac computers from a single keyboard, monitor, and mouse.

Another big benefit of using blade servers is that they drastically cut down the amount of cable clutter. With rack-mounted servers, each server requires its own power, keyboard, video, mouse, and network cables. With blade servers, a single set of cables can service all the servers in a blade enclosure.

Tiny servers

You can also get servers in vey small packages — even as small as a deck of cards. Most of these computers are *single-board computers,* meaning that the entire computer is built on a single small motherboard that houses the CPU, RAM, disks, video, and network components. For Windows servers, NUC computers are an excellent choice; for Linux servers, look into Raspberry Pi.

Tiny servers are typically used for very specific purposes. For example, you might set up a NUC or Raspberry Pi to provide DHCP or routing services for a small network.

Understanding Virtualization

One final subject for this chapter is the concept of virtualization. Throughout this chapter, I use the term *server* to refer both to an operating system that provides services (such as file-sharing or directory services), as well as to the hardware on which that operating system runs. However, in many (if not most) modern network environments, a single physical computer system is used to run more than one *virtual machine* (VM). A VM is a simulation of an actual computer system. This concept is called *virtualization.* When virtualization is used, a single physical server computer actually runs more than one virtual server.

Virtualization is the reason that server computer hardware often has such high performance specifications, such as dual processors with multiple cores each and a large amount of RAM (256GB or more). In most environments, no single server really needs that much capacity. But when a single physical computer is responsible for running multiple virtual servers, the physical server must have sufficient capacity to run all its virtual servers.

Also, note that virtualization isn't just for servers: In many organizations, desktop computers are also virtualized. Virtualizing desktops offers many advantages, especially in organizations where most users need just basic computers to access a few simple business applications.

If this concept seems confusing at first, don't sweat it. You can find out more about virtualization in Book 3, Chapter 4, as well as in Book 5.

IN THIS CHAPTER

- Examining the basics of cloud computing
- Looking at three kinds of cloud computing services
- Understanding the pros and cons of cloud computing
- Perusing a few major cloud computing service providers

Chapter 6
Cloud Computing

The world's two most popular science-fiction franchises — *Star Wars* and *Star Trek* — both feature cities that are suspended in the clouds. In *The Empire Strikes Back*, Han Solo takes the *Millennium Falcon* to Cloud City, hoping that his friend Lando Calrissian can help repair their damaged hyperdrive. And in the original *Star Trek* series episode "The Cloud Minders," the crew of the *Enterprise* visits a city named Stratos, which is suspended in the clouds.

Coincidence? Perhaps. Or maybe Gene Roddenberry and George Lucas both knew that the future would be in the clouds. At any rate, the future of computer networking is rapidly heading for the clouds — cloud computing, to be specific. This chapter is a brief introduction to cloud computing. You discover what it is, the pros and cons of adopting it, and what services are provided by the major cloud-computing providers.

Introducing Cloud Computing

The basic idea behind cloud computing is to outsource one or more of your networked computing resources to the internet. The cloud represents a newish way of handling common computer tasks. Table 6-1 outlines just a few examples of how the cloud way differs from the traditional way.

TABLE 6-1 Traditional versus Cloud Computing

	Traditional	Cloud
Email services	Provide email services by installing Microsoft Exchange on a local server computer.	Contract with an internet-based email provider, such as Gmail (from Google) or Exchange Online (from Microsoft).
Disk storage	Set up a local file server computer with a large amount of shared disk space.	Sign up for Microsoft Office and store your files in Microsoft OneDrive or Microsoft Teams. Or, use Google Drive to store your files in the cloud.
Accounting services	Purchase expensive accounting software and install it on a local server computer.	Sign up for a web-based accounting service.

Looking at the Benefits of Cloud Computing

Cloud computing is a different — and, in many ways, better — approach to networking. The following sections cover a few of the main benefits of moving to cloud-based networking.

Cost

Cloud-based computing typically is less expensive than traditional computing. Consider a typical file server application: To implement a file server, first you have to purchase a file server computer with enough disk space to accommodate your users' needs, perhaps as much as 10TB of disk storage or more. You want the most reliable data storage possible, so you purchase a server-quality computer and fully redundant solid-state disk drives. For the sake of this discussion, figure that the total price of the server — including the disk storage, the operating system license, and the labor cost of setting it up — is about $10,000. Assuming that the server will last for four years, that totals about $2,500 per year.

If you instead acquire your disk storage from a cloud-based file-sharing service, you can expect to pay about one-fourth of that amount for an equivalent amount of storage.

The same economies apply to most other cloud-based solutions. Cloud-based email solutions, for example, typically cost around $5 per month per user — far less than the cost of setting up and maintaining an on-premises Microsoft Exchange Server.

Scalability

So, what happens if you guess wrong about the storage requirements of your file server, and your users end up needing 20TB instead of 10TB? With a traditional file server, you must purchase additional disk drives to accommodate the extra space. Sooner than you want, you'll run out of capacity in the server's cabinet. Then you'll have to purchase an external storage cabinet. Eventually, you'll fill that up, too.

Now suppose that after you expand your server capacity to 20TB, your users' needs contract to just 10TB. Unfortunately, you can't return disk drives for a refund.



REMEMBER

With cloud computing, you pay only for the capacity you're actually using, and you can add capacity whenever you need it. In the file server example, you can write as much data as you need to the cloud storage. Each month, you're billed according to your actual usage. Thus, you don't have to purchase and install additional disk drives to add storage capacity.

Reliability

Especially for smaller businesses, cloud services are much more reliable than in-house services. Just a week before I wrote this chapter, the tape drive that a friend uses to back up his company's data failed. As a result, he was unable to back up data for three days while the tape drive was repaired. Had he been using cloud-based backup, he could've restored his data immediately and wouldn't have been without backups for those four days.

The reason for the increased reliability of cloud services is simply a matter of scale. Most small businesses can't afford the redundancies needed to make their computer operations as reliable as possible. My friend's company can't afford to buy two tape drives so that an extra is available in case the main one fails.

By contrast, cloud services are usually provided by large companies such as Amazon, Google, and Microsoft. These companies have state-of-the-art data centers filled with redundancy. Cloud storage may be kept on multiple servers so that if one server fails, others can take over the load. In some cases, these servers are in different data centers in different parts of the country. Thus, your data will still be available even in the event of a disaster that shuts down an entire data center.

Accessibility

One of the best things about cloud services is that they're available anywhere you have an internet connection. Suppose that you have offices in five cities. Using

traditional computing, each office would require its own servers, and you'd have to carefully design systems that allowed users in each of the offices to access shared data.

With cloud computing, each office simply connects to the internet to access the cloud applications. Cloud-based applications are also great if your users are mobile because they can access the applications anywhere they can find an internet connection.

Free of hassles

IT can be a hassle. With cloud-based services, you can outsource the job of complex system maintenance chores, such as upgrades, patches, hardware maintenance, backups, and so on. You get to consume the services while someone else takes care of making sure that the services run properly.

Detailing the Drawbacks of Cloud Computing

Although cloud computing has many advantages over traditional techniques, it isn't without its drawbacks. The following sections outline some of the most significant roadblocks to adopting cloud computing.

Entrenched applications

Your organization may depend on entrenched applications that don't lend themselves especially well to cloud computing — or that at least require significant conversion efforts to migrate to the cloud. For example, you may use an accounting system that relies on local file storage.

Fortunately, many cloud providers offer assistance with this migration. And in many cases, the same application that you run locally can be run in the cloud, so no conversion is necessary.

Internet connection speed

Cloud computing shifts much of the burden of your network to your internet connection. Your users used to access their data on local file servers over gigabit-speed

connections, and each user had a relatively dedicated network path to those servers. With cloud computing, everyone has to access data over a single internet connection, often slower than the individual pathways within your local network.


REMEMBER

Although you can upgrade your connection to higher speeds, doing so will cost money — money that may well offset the money you would otherwise have saved by migrating to the cloud.

Internet connection reliability

The cloud resources you access may feature all the redundancy in the world, but if your users access the cloud through a single internet connection, that connection becomes a single point of vulnerability. If it fails, everything that depends on it will fail. Your business may come to a halt until the connection is restored.


TIP

Here are two ways to mitigate this risk:

- **Make sure that you have an enterprise-class internet connection.** Enterprise-class connections are more expensive but provide much better fault tolerance and repair service than consumer-class connections do.
- **Provide redundant connections if you can.** That way, if one connection fails, traffic can be rerouted through alternative connections.

Security threats

You can bet your life that hackers around the world are continually probing for ways to break through the security perimeter of all the major cloud providers. When they do, your data may be exposed.


REMEMBER

Always ensure that strong password policies are enforced!

Examining Three Basic Kinds of Cloud Services

Three distinct kinds of services can be provided via the cloud: applications, platforms, and services (infrastructure). The following sections describe these three types of cloud services in greater detail.

Applications

Most often referred to as *Software as a Service* (SaaS), fully functional applications can be delivered via the cloud. One of the best-known examples is Microsoft 365, which includes cloud-based email (Exchange Online), cloud storage (OneDrive), cloud-based collaboration (Teams), and traditional Office applications (Outlook, Word, Excel, PowerPoint, Access, and more).

A popular alternative to Microsoft 365 is Google Workspace (formerly known as G Suite). Google Workspace provides its own office applications that compete with Microsoft's traditional office applications, as well as cloud-based email and storage.

When you use a cloud-based app, you don't have to worry about any of the details that are commonly associated with running an application on your network, such as deploying the app and applying product upgrades and software patches. Cloud-based apps usually charge a small monthly fee based on the number of users running the software, so costs are low.

Also, as a cloud-based app user, you don't have to worry about providing the hardware or operating system platform on which the application will run. The app provider takes care of those details for you, so you can focus simply on deploying the app to best serve your users' needs.

Platforms

Also referred to as *Platform as a Service* (PaaS), this class of service refers to providers that give you access to a remote virtual operating platform on which you can build your own applications.

At the simplest level, a PaaS provider gives you a complete, functional remote virtual machine (VM) that's fully configured and ready for you to deploy your applications to. If you use a web provider to host your company's website, you're already using PaaS: Most web host providers give you a functioning Linux system, fully configured with all the necessary servers, such as Apache or MySQL. All you have to do is build and deploy your web application on the provider's server.

More-complex PaaS solutions include specialized software that your custom applications can tap to provide services such as data storage, online order processing, and credit card payments.

You have many cloud platform providers to choose from. The best known are AWS, Google Cloud, and Microsoft Azure. Azure is described in Book 5, Chapter 3, and AWS is described in Book 5, Chapter 4.



REMEMBER

When you use PaaS, you take on the responsibility of developing your own custom apps to run on the remote platform. The PaaS provider takes care of the details of maintaining the platform itself, including the base operating system and the hardware on which the platform runs.

Infrastructure

Infrastructure as a Service (IaaS) lets you retain control over your virtual servers but delegate control of the hardware platform on which they run. When you use IaaS, you're purchasing raw computing power that's accessible via the cloud. Typically, IaaS provides you access to remotely hosted VMs, which you can configure however you want.

Public Clouds versus Private Clouds

The most common form of cloud computing uses what is known as a *public cloud* (cloud services that are available to anyone in the world via the internet). Google Workspace is an excellent example of a public cloud service. Anyone with access to the internet can access the public cloud services of Google Workspace: Just point your browser to `https://workspace.google.com` to get started.

A public cloud is like a public utility, in that anyone can subscribe to it on a pay-as-you-go basis. One drawback of public cloud services is that when you use a public cloud service, you entrust your valuable data to a third party that you can't control. You can protect access by using strong passwords, but if your usernames and passwords are compromised, public cloud services can be hacked into and your data can be stolen. Every so often, we all hear news stories about how this company's or that company's backdoor security has been compromised.

A *private cloud* mimics many of the features of cloud computing but is implemented on private hardware within a local network, so it isn't accessible to the general public. Private clouds are inherently more secure because the general public can't access them. Also, they're dependent only on private network connections, so they aren't subject to the limits of a public internet connection.



TIP

As a rule, private clouds are implemented by large organizations that have the resources available to create and maintain their own cloud servers.

A compromise between a public and a private cloud is a *hybrid cloud*, which combines the features of both. Typically, a hybrid cloud system uses a small private cloud for local access to some applications and a public cloud for others. For

example, you might maintain your most frequently used data on a private cloud and use the public cloud to store archive data or other less frequently used data.

Introducing Some of the Major Cloud Providers

Although hundreds (even thousands) of companies provide cloud services, most cloud computing is provided by just a few providers.

Amazon

By far the largest provider of cloud services in the world is Amazon. Amazon launched its cloud platform, AWS, in 2006. Since then, hundreds of thousands of customers have signed up. Some of the most notable users of AWS include Airbnb, Netflix, and Pinterest.

AWS includes the following features:

- **Amazon CloudFront:** A PaaS content-delivery system designed to deliver web content to large numbers of users
- **Amazon Elastic Compute Cloud (EC2):** An IaaS system that provides access to raw computing power
- **Amazon Simple Storage Service (S3):** Provides web-based data storage for unlimited amounts of data
- **Amazon Relational Database Service (RDS):** Lets you house your databases in the cloud
- **Amazon Virtual Private Cloud (VPC):** Uses virtual private network (VPN) connections to connect your local network to Amazon's cloud services

For more information about AWS, refer to Book 5, Chapter 4.

Google

Google is also one of the largest providers of cloud services. Its offerings include the following:

- **Google Workspace:** An alternative to Microsoft 365 that provides email, word processing, and spreadsheet functions via the cloud, as well as storage.

Monthly subscription prices range from $7.20 to $21.60 per month, depending on the level of services you need.

- **Google App Engine:** A PaaS interface that lets you develop your own applications that work with Google's cloud services.
- **Google Cloud:** A rich environment in which you can create virtual servers, networks, and storage resources within Google's massive cloud data centers.

Microsoft

Microsoft has its own cloud strategy, designed in part to protect its core business of operating systems and Office applications against competition from other cloud providers, such as Google.

Here are two of Microsoft's cloud offerings:

- **Microsoft 365:** A cloud-based version of Office, which includes Outlook, Teams, SharePoint, Word, Excel, PowerPoint, Access, OneDrive, and many other powerful tools
- **Azure:** A PaaS offering that lets you build websites, deploy VMs that run Windows Server or Linux, or access cloud versions of server applications such as Microsoft SQL Server

For more information about Azure, refer to Book 5, Chapter 3.

Getting into the Cloud


TIP

After you wrap your head around just how cool cloud computing can be, what should you do to move your network toward the cloud? Allow me to make a few recommendations:

- **Don't depend on a poor internet connection.** First and foremost, before you take any of your network operations to the cloud, make sure that you're *not* dependent on a consumer-grade internet connection. Consumer-grade internet connections can be fast, but when an outage occurs, there's no telling how long you'll wait for the connection to be repaired. Invest in a high-speed enterprise-class connection that can scale as your dependence on it increases.

- **Assess which applications you may already have running on the cloud.** If you use Gmail rather than Exchange for your email, congratulations! You've already embraced the cloud. Other examples of cloud services that you may already be using include a remote web host or a file-sharing application such as Dropbox or ShareFile.
- **Don't move to the cloud all at once.** Start by identifying a single application that lends itself to the cloud. For example, if your company archives files when they reach a certain age, look to the cloud for a file storage service.
- **Go with a reputable company.** Amazon, Google, and Microsoft are all huge companies with proven track records in cloud computing. Many other large and established companies also offer cloud services. Don't stake your company's future on a company that didn't exist six months ago.
- **Research, research, research.** Pour yourself into the web, and buy a few books. *Cloud Computing For Dummies,* 2nd Edition, by Daniel Kirsch and Judith Hurwitz (John Wiley & Sons), is a good place to start.

2 Understanding Network Protocols

Contents at a Glance

CHAPTER 1: **Network Protocols and Standards** 93

CHAPTER 2: **TCP/IP and the Internet** 115

CHAPTER 3: **IP Addresses** 123

CHAPTER 4: **Routing** 145

CHAPTER 5: **DHCP** 155

CHAPTER 6: **DNS** 173

CHAPTER 7: **TCP/IP Tools and Commands** 207

IN THIS CHAPTER

- Discovering protocols
- Deciphering the layers of the OSI reference model
- Understanding Ethernet
- Getting the inside scoop on TCP/IP and IPX/SPX
- Finding out about other important protocols

Chapter 1

Network Protocols and Standards

Protocols and standards make networks work together. Protocols make it possible for the various components of a network to communicate with each other, and standards make it possible for different manufacturers' network components to work together. This chapter introduces you to the protocols and standards that you're most likely to encounter when building and maintaining a network.

Understanding Protocols

A *protocol* is simply a set of rules that enable effective communications to occur. You encounter protocols every day and probably don't even realize it. When you pay for groceries with a debit card, the clerk tells you how much the groceries cost, and then you swipe your debit card in the card reader, punch in your security code, indicate whether you want cash back, enter the amount of the cash back if you so indicated, and verify the total amount. You then cross your fingers behind your back and say a quiet prayer while the machine authorizes the purchase. Assuming the amount is authorized, the machine prints out your receipt.



REMEMBER

Computer networks depend upon many different types of protocols. These protocols are very rigidly defined — and for good reason. Network interfaces must know how to talk to other network interfaces to exchange information, operating systems must know how to talk to network interfaces to send and receive data on the network, and application programs must know how to talk to operating systems to know how to retrieve a file from a network server.

Protocols come in many different types. At the lowest level, protocols define exactly what type of electrical signal represents a 1 and what type of signal represents a 0. At the highest level, protocols allow, for example, a computer user in the United States to send an email to another computer user in New Zealand — and in between are many other levels of protocols. You find out more about these levels of protocols (often called *layers*) in the upcoming section, "Seeing the Seven Layers of the OSI Reference Model."



TIP

Protocols tend to be used together in matched sets called *protocol suites.* The two most popular protocol suites for networking are Transmission Control Protocol/Internet Protocol (TCP/IP) and Ethernet. TCP/IP, originally developed for Unix networks, is the protocol of the internet and most local area networks (LANs). Ethernet is a low-level protocol that spells out the electrical characteristics of the network hardware used by most LANs.

Understanding Standards

As I mention earlier, a *standard* is an agreed-upon definition of a protocol. In the early days of computer networking, each computer manufacturer developed its own networking protocols. As a result, you couldn't easily mix equipment from different manufacturers on a single network.

Then along came standards to save the day. Hurrah! Because standards are industry-wide protocol definitions not tied to a particular manufacturer, you can mix and match equipment from different vendors. As long as the equipment implements the standard protocols, it should be able to coexist on the same network.

Many organizations are involved in setting standards for networking. The five most important organizations are

- **American National Standards Institute (ANSI; `www.ansi.org`):** The official standards organization in the United States. *ANSI* is pronounced *an*-see.
- **Institute of Electrical and Electronics Engineers (IEEE; `www.ieee.org`):** An international organization that publishes several key networking

standards — in particular, the official standard for the Ethernet networking system (known officially as IEEE 802.3). IEEE is pronounced eye-triple-E.

- **International Organization for Standardization (ISO; `www.iso.org`):** A federation of more than 100 standards organizations throughout the world.


TECHNICAL STUFF

If you're wondering why the acronym for *International Organization for Standardization* is *ISO* and not *IOS,* the answer is simple: ISO is truly an international organization, and although it is known as the International Organization for Standardization in English-speaking countries, it goes by different names in non-English-speaking countries — for example, in French-speaking countries, it's known as *Organisation internationale de normalisation.* The organization's founders chose *ISO* so that it would have the same short name in all languages.

- **Internet Engineering Task Force (IETF; `www.ietf.org`):** The organization responsible for the protocols that drive the internet.
- **World Wide Web Consortium (W3C; `www.w3.org`):** An international organization that handles the development of standards for the World Wide Web.

Seeing the Seven Layers of the OSI Reference Model

OSI sounds like the name of a top-secret government agency you hear about only in Tom Clancy novels. What it really stands for in the networking world is *Open Systems Interconnection*, as in the Open Systems Interconnection reference model, affectionately known as the OSI model.

The OSI model breaks the various aspects of a computer network into seven distinct layers. These layers are kind of like the layers of an onion: Each successive layer envelops the layer beneath it, hiding its details from the levels above. The OSI model is also like an onion in that if you start to peel it apart to have a look inside, you're bound to shed a few tears.

The OSI model is not a networking standard in the same sense that Ethernet and TCP/IP are networking standards. Instead, the OSI model is a framework into which the various networking standards can fit. The OSI model specifies which aspects of a network's operation can be addressed by various network standards. So, in a sense, the OSI model is sort of a standard of standards.

Table 1-1 summarizes the seven layers of the OSI model.

TABLE 1-1 The Seven Layers of the OSI Model

Layer	Name	Description
1	Physical	Governs the layout of cables and devices, such as repeaters and hubs.
2	Data link	Provides media access control (MAC) addresses to uniquely identify network nodes and a means for data to be sent over the physical layer in the form of packets. Bridges and switches are layer-2 devices.
3	Network	Handles routing of data across network segments.
4	Transport	Provides for reliable delivery of packets.
5	Session	Establishes sessions between network applications.
6	Presentation	Converts data so that systems that use different data formats can exchange information.
7	Application	Allows applications to request network services.

The first three layers are sometimes called the *lower layers.* They deal with the mechanics of how information is sent from one computer to another over a network. Layers 4 through 7 are sometimes called the *upper layers.* They deal with how application software can relate to the network through application programming interfaces (APIs).

The following sections describe each of these layers in greater detail.



REMEMBER

The seven layers of the OSI model are a somewhat idealized view of how networking protocols should work. In the real world, actual networking protocols don't follow the OSI model to the letter. The real world is always messier. Still, the OSI model provides a convenient — if not completely accurate — conceptual picture of how networking works.

The physical layer

The bottom layer of the OSI model is the *physical layer.* It addresses the physical characteristics of the network, such as the types of cables used to connect devices, the types of connectors used, how long the cables can be, and so on. For example, Ethernet spells out the exact layer-1 requirements for twisted-pair cables that can be used at various speeds — 100 Mbps, 1 Gbps, 10 Gbps, and even faster. The star, bus, ring, and mesh network topologies described in Book 1, Chapter 2 apply to the physical layer.

Another aspect of the physical layer is the electrical characteristics of the signals used to transmit data over the cables from one network node to another. The physical layer doesn't assign any meaning to those signals other than the basic binary values of 1 and 0. The higher levels of the OSI model must assign meanings to the bits that are transmitted at the physical layer.

One type of physical layer device commonly used in networks is a *repeater*, which is used to regenerate the signal whenever you need to exceed the cable length allowed by the physical layer standard. In the old days, we used to use physical layer devices called *hubs* to split an Ethernet segment to multiple devices. Technically, hubs are known as *multiport repeaters* because the purpose of a hub is to regenerate every packet received on any port on all the hub's other ports. Repeaters and hubs don't examine the contents of the packets that they regenerate, though. If they did, they would be working at the data link layer, and not at the physical layer.

The *network adapter* (also called a network interface card [NIC]) installed in each computer on the network is a physical layer device. You can display information about the network adapter(s) installed on a Windows computer by displaying the adapter's Ethernet Properties dialog box, as shown in Figure 1-1. To access this dialog box in Windows, open the Control Panel, choose Network and Internet, choose Adapter Settings, right-click the adapter whose settings you want to view, and choose Properties.

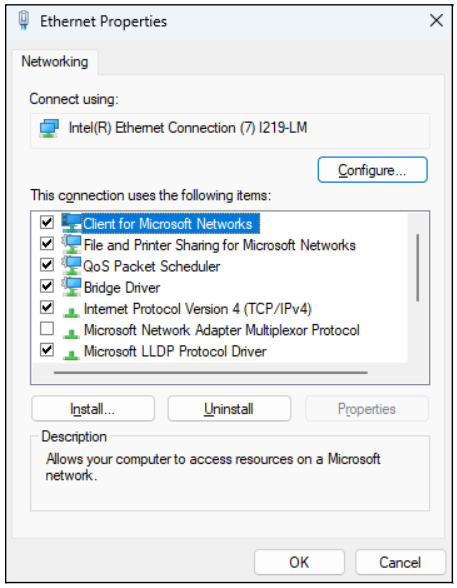


FIGURE 1-1: The Ethernet Properties dialog box for a network adapter.



TIP

Like its predecessor, Windows 10, Windows 11 seems to bury some of the most useful settings pages, making them difficult to find. Even the incredibly useful Control Panel can be a chore to find. I suggest you pin the Control Panel to both the Start menu and the taskbar. You can find the Control Panel by pressing the Windows key, typing **Control Panel**, right-clicking the Control Panel icon, and choosing both Pin to Start and Pin to Taskbar so the Control Panel will always be readily available.

While you're at it, switch Control Panel from Category view to Small Icons view. This step will eliminate a lot of extra navigation trying to get to the settings pages you need. For example, in Small Icons view, you can go directly from Control Panel to Network and Sharing Center without first having to open the Network and Internet link.

The data link layer

The *data link layer* is the lowest layer at which meaning is assigned to the bits that are transmitted over the network. Data link protocols address things such as the size of each packet of data to be sent, a means of addressing each packet so that it's delivered to the intended recipient, and a way to ensure that two or more nodes don't try to transmit data on the network at the same time.

The data link layer also provides basic error detection and correction to ensure that the data sent is the same as the data received. If an uncorrectable error occurs, the data link standard must specify how the node is to be informed of the error so that it can retransmit the data.

At the data link layer, each device on the network has an address: the MAC address. This address is hardwired into every network device by the manufacturer. MAC addresses are unique; no two network devices made by any manufacturer anywhere in the world can have the same MAC address.

You can see the MAC address for a computer's network adapter by opening a command window and running the `ipconfig /all` command, as shown in Figure 1-2. In this example, the MAC address of the network card is 80-E8-2C-CE-4E-0D. (The `ipconfig` command refers to the MAC address as the *physical address.*)



TECHNICAL STUFF

One of the most important functions of the data link layer is to provide a way for packets to be sent safely over the physical media without interference from other nodes trying to send packets at the same time. The most popular ways to do this is Carrier Sense Multiple Access/Collision Detection (CSMA/CD).

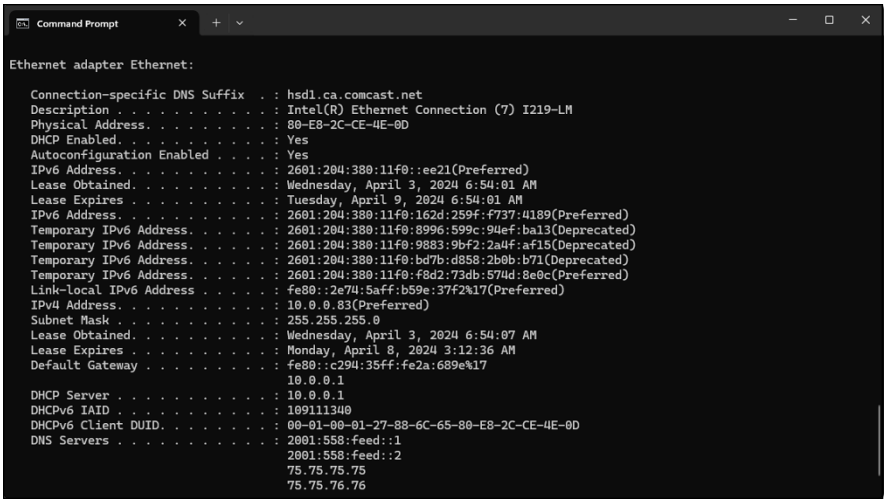


FIGURE 1-2: Display the MAC address of a network adapter.

Two types of data link layer devices are commonly used on networks:

- **Bridge:** An intelligent repeater that's aware of the MAC addresses of the nodes on either side of the bridge and can forward packets accordingly
- **Switch:** An intelligent hub that examines the MAC address of each arriving packet to determine which port to forward the packet to

Another important layer-2 concept is the idea of *virtual local area networks* (VLANs). VLANs allow you to create separate isolated networks that share devices. For example, you can create separate VLANs for your accounting and sales departments but use the same switch to connect computers in each VLAN. Computers in the accounting VLAN won't be able to communicate with computers in the sales VLAN, even though the computers are on the same switch.

CSMA/CD IS A MOUTHFUL!

An important function of the data link layer is to make sure that two computers don't try to send packets over the network at the same time. If they do, the signals will collide with each other, and the transmission will be garbled. Ethernet accomplishes this feat by using CSMA/CD. This phrase is a mouthful, but if you take it apart piece-by-piece, you'll get an idea of how it works.

- *Carrier Sense* means that whenever a device wants to send a packet over the network media, it first listens to the network media to see whether anyone else is

(continued)

(continued)

already sending a packet. If it doesn't hear any other signals on the media, the computer assumes that the network is free, so it sends the packet.

- *Multiple Access* means that nothing prevents two or more devices from trying to send a message at the same time. Sure, each device listens before sending. However, suppose that two devices listen, hear nothing, and then proceed to send their packets at the same time? Picture what happens when you and someone else arrive at a four-way stop at the same time. You wave the other driver on, they wave you on, you wave, they wave, you both wave, and then you both go at the same time.
- *Collision Detection* means that after a device sends a packet, it listens carefully to see whether the packet crashes into another packet. This is kind of like listening for the screeching of brakes at the four-way stop. If the device hears the screeching of brakes, it waits a random period of time and then tries to send the packet again. Because the delay is random, two packets that collide are sent again after different delay periods, so a second collision is unlikely.

CSMA/CD works pretty well for smaller networks. After a network hits about 30 computers, however, packets start to collide like crazy, and the network slows to a crawl. When that happens, the network should be divided into two or more separate sections that are sometimes called *collision domains.*

The network layer

The *network layer* handles the task of routing network messages from one computer to another. The most popular layer-3 protocol is Internet Protocol (IP), which is usually paired with Transmission Control Protocol (TCP).

Network layer protocols provide two important functions: logical addressing and routing. The following sections describe these functions.

Logical addressing

As I mention earlier, every network device has a physical address — a MAC address — assigned to the device at the factory. When you buy a NIC to install into a computer, the MAC address of that NIC is fixed and can't be changed. So, what happens if you want to use some other addressing scheme to refer to the computers and other devices on your network? This is where the concept of logical addressing comes in. With a logical address, you can access a network device by using an address that you assign.

Logical addresses are created and used by network layer protocols such as IP. The network layer protocol translates logical addresses to MAC addresses. For example, if you use IP as the network layer protocol, devices on the network are assigned IP addresses, such as 207.120.67.30. Because the IP protocol must use a data link layer protocol to send packets to devices, IP must know how to translate the IP address of a device to the device's MAC address.



REMEMBER

You can use the `ipconfig` command (refer to Figure 1-2) to see the IP address of your computer. The IP address shown in that figure is 10.0.0.83. Another way to display this information is to use the System Information command, best found by pressing the Windows key and searching for "System Information." The IP address is highlighted in Figure 1-3. Notice that the System Information program displays a lot of other useful information about the network besides the IP address. For example, you can also see the MAC address and what protocols are being used.

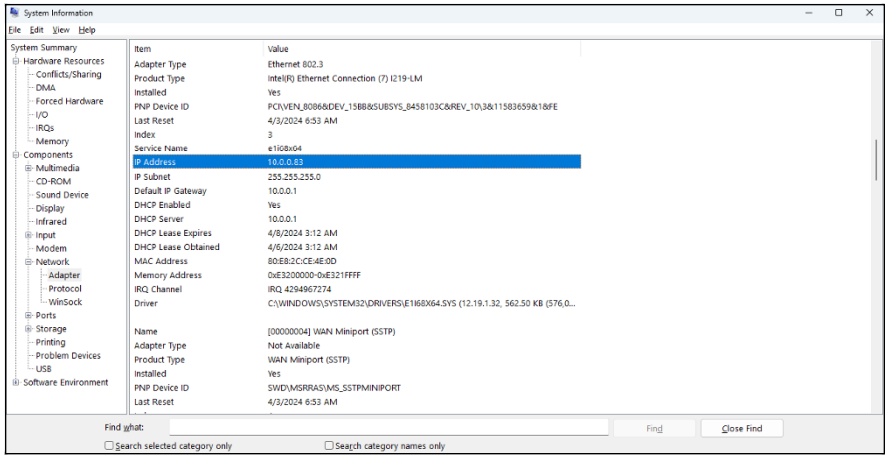


FIGURE 1-3: Find network information from System Information.

Although the exact format of logical addresses varies depending on the protocol being used, most protocols divide the logical address into two parts:

- **Network address:** Identifies which network the device resides on
- **Device address:** Identifies the device on that network

In a typical IP address — say, 192.168.1.102 — the network address is 192.168.1, and the device address (called a *host address* in IP) is 102.

Routing

Routing comes into play when a computer on one network needs to send a packet to a computer on another network. In this case, a router is used to forward the packet to the destination network. In some cases, a packet may have to travel through several intermediate networks in order to reach its final destination network. You can find out more about routers in Book 2, Chapter 4.

An important feature of routers is that you can use them to connect networks that use different layer-2 protocols. For example, a router can be used to send a packet from an Ethernet to a Token Ring network. As long as both networks support the same layer-3 protocol, it doesn't matter whether their layer-1 and layer-2 protocols are different.


TIP

A protocol is considered routable if it uses addresses that include a network part and a host part. Any protocol that uses physical addresses isn't routable because physical addresses don't indicate to which network a device belongs.

The transport layer

The *transport layer* is where you find two of the most well-known networking protocols: TCP (typically paired with IP) and SPX (typically paired with IPX). As its name implies, the transport layer is concerned with the transportation of information from one computer to another.

The main purpose of the transport layer is to ensure that packets are transported reliably and without errors. The transport layer does this task by establishing connections between network devices, acknowledging the receipt of packets, and resending packets that aren't received or are corrupted when they arrive.

In many cases, the transport layer protocol divides large messages into smaller packets that can be sent over the network efficiently. The transport layer protocol reassembles the message on the receiving end, making sure that all the packets that make up a single transmission are received so that no data is lost.

For some applications, speed and efficiency are more important than reliability. In such cases, a connectionless protocol can be used. As you can likely guess, a connectionless protocol doesn't go to the trouble of establishing a connection before sending a packet — it simply sends the packet. TCP is a connection-oriented transport layer protocol. The connectionless protocol that works alongside TCP is User Datagram Protocol (UDP).

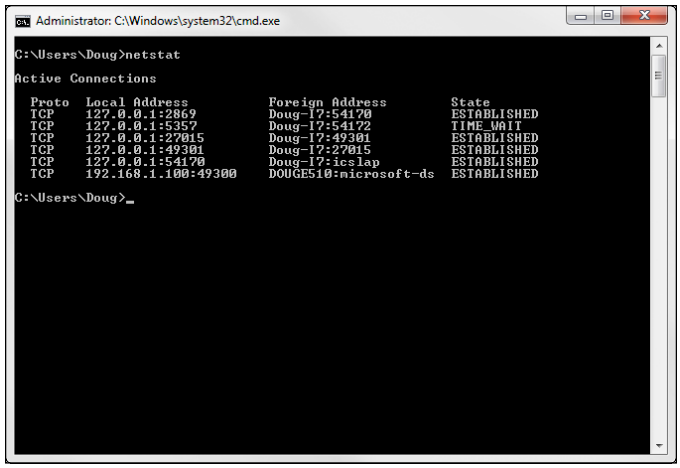


FIGURE 1-4:
TCP connections.

You can view information about the status of TCP and UDP connections by running the `Netstat` command from a command window, as shown in Figure 1-4. In the figure, you can see that several TCP connections are established.

In fact, you can use the command `Netstat /N` to see the numeric network addresses instead of the names. With the `/N` switch, the output in Figure 1-4 would look like this:

```
Active Connections
Proto   Local Address          Foreign Address      State
TCP     127.0.0.1:2869         127.0.0.1:54170      ESTABLISHED
TCP     127.0.0.1:5357         127.0.0.1:54172      TIME_WAIT
TCP     127.0.0.1:27015        127.0.0.1:49301      ESTABLISHED
TCP     127.0.0.1:49301        127.0.0.1:27015      ESTABLISHED
TCP     127.0.0.1:54170        127.0.0.1:2869       ESTABLISHED
TCP     192.168.1.100:49300    192.168.1.101:445    ESTABLISHED
```



REMEMBER

TCP is a connection-oriented transport layer protocol. *UDP* is a connectionless transport layer protocol.

The session layer

The *session layer* establishes *conversations* (sessions) between networked devices. A *session* is an exchange of connection-oriented transmissions between two network devices. Each transmission is handled by the transport layer protocol. The session itself is managed by the session layer protocol.

A single session can include many exchanges of data between the two computers involved in the session. After a session between two computers has been established, it's maintained until the computers agree to terminate the session.

The session layer allows three types of transmission modes:

- **Simplex:** Data flows in only one direction.
- **Half-duplex:** Data flows in both directions, but only in one direction at a time.
- **Full-duplex:** Data flows in both directions at the same time.


TIP

In actual practice, the distinctions in the session, presentation, and application layers are often blurred, and some commonly used protocols actually span all three layers. For example, Server Message Block (SMB) — the protocol that is the basis of file sharing in Windows networks — functions at all three layers.

The presentation layer

The *presentation layer* is responsible for how data is represented to applications. The most common representation for representing character data today is called UTF-8, which uses 8-bit sets to represent most characters found in western alphabets. UTF-8 is compatible with an older standard called ASCII.


TECHNICAL STUFF

UTF-8 is sometimes called *Unicode*, which is a standard for representing the characters found in most of the world's writing systems. Technically, UTF-8 is a particular method of implementing Unicode, so although the two terms are related, they are not identical.


TECHNICAL STUFF

Some computers, in particular IBM mainframe computers, use a different code called EBCDIC. ASCII and EBCDIC aren't compatible. To exchange information between a mainframe computer and a Windows computer, the presentation layer must convert the data from ASCII to EBCDIC, and vice versa.

Besides simply converting data from one code to another, the presentation layer can also apply sophisticated compression techniques so that fewer bytes of data are required to represent the information when it's sent over the network. At the other end of the transmission, the presentation layer then decompresses the data.

The presentation layer can also scramble the data before it's transmitted and then unscramble it at the other end by using a sophisticated encryption technique that even Sherlock Holmes would have trouble breaking.

The application layer

The highest layer of the OSI model, the application layer, deals with the techniques that application programs use to communicate with the network. The name of this layer is a little confusing. Application programs (such as Microsoft Office or QuickBooks) aren't a part of the application layer. Instead, the application layer represents the programming interfaces that application programs use to request network services.

Some of the better-known application layer protocols are

- **Domain Name System (DNS):** For resolving internet domain names
- **File Transfer Protocol (FTP):** For file transfers
- **Simple Mail Transfer Protocol (SMTP):** For email
- **Server Message Block (SMB):** For file sharing in Windows networks
- **Network File System (NFS):** For file sharing in Unix networks
- **Telnet:** For terminal emulation

Following a Packet through the Layers

Figure 1-5 shows how a packet of information flows through the seven layers as it travels from one computer to another on the network. The data begins its journey when an end-user application sends data to another network computer. The data enters the network through an application layer interface, such as SMB. The data then works its way down through the protocol stack. Along the way, the protocol at each layer manipulates the data by adding header information, converting the data into different formats, combining packets to form larger packets, and so on. When the data reaches the physical layer protocol, it's placed on the network media (in other words, the cable) and sent to the receiving computer.

When the receiving computer receives the data, the data works its way up through the protocol stack. Then the protocol at each layer reverses the processing that was done by the corresponding layer on the sending computer. Headers are removed, data is converted back to its original format, packets that were split into smaller packets are recombined into larger packets, and so on. When the packet reaches the application layer protocol, it's delivered to an application that can process the data.

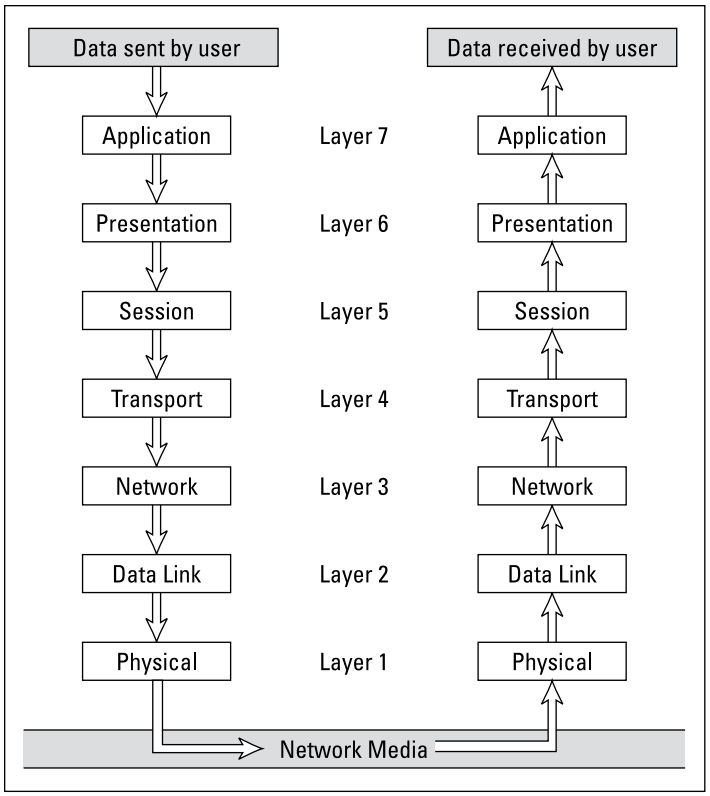


FIGURE 1-5: How data travels through the seven layers.

The Ethernet Protocol

As I mention earlier, the first two layers of the OSI model deal with the physical structure of the network and the means by which network devices can send information from one device on a network to another. By far, Ethernet is the most popular set of protocols for the physical and data link layers.

Ethernet has been around in various forms since the early 1970s. (For a brief history of Ethernet, see the sidebar, "Ethernet folklore and mythology.") The current incarnation of Ethernet is defined by the 802.3 IEEE standard. Various flavors of Ethernet operate at different speeds and use different types of media. However, all the versions of Ethernet are compatible with each other, so you can mix and match them on the same network by using devices such as bridges, hubs, and switches to link network segments that use different types of media.

ETHERNET FOLKLORE AND MYTHOLOGY

The original idea for the Ethernet was hatched in the mind of Robert Metcalfe, a graduate student in computer science at Harvard University. Looking for a thesis idea in 1970, he refined a networking technique used in Hawaii — the AlohaNet (actually a wireless network) — and developed a technique that would enable a network to efficiently use as much as 90 percent of its capacity. By 1973, he had his first Ethernet network up and running at the famous Xerox Palo Alto Research Center (PARC). Bob dubbed his network "Ethernet" in honor of the thick network cable, which he called "the ether." (Xerox PARC was busy in 1973. In addition to Ethernet, PARC developed the first personal computer that used a graphical user interface [GUI], complete with icons, windows, and menus, and the world's first laser printer.)

In 1979, Xerox began working with Intel and DEC (a once-popular computer company) to make Ethernet an industry standard networking product. Along the way, they enlisted the help of the IEEE, which formed committee number 802.3 and began the process of standardizing Ethernet in 1981. The 802.3 committee released the first official Ethernet standard in 1983.

Meanwhile, Bob Metcalfe left Xerox, turned down an offer from Steve Jobs to work at Apple, and started a company called 3Com, which has since become one of the largest manufacturers of Ethernet equipment in the world.



TIP

The actual transmission speed of Ethernet is measured in megabits per second (Mbps) or gigabits per second (Gbps). Ethernet comes in several different speeds:

- **Standard Ethernet:** 10 Mbps; rarely (if ever) used today.
- **Fast Ethernet:** 100 Mbps; still used for devices where speed is not particularly important, such as printers or fax machines.
- **Gigabit Ethernet and beyond:** 1,000 Mbps; the most common speed used to connect user computers to a network. Faster speeds, such as 10 Gbps, 100 Gbps, and even faster, are sometimes used in high-speed networks to connect servers and other critical devices to the network.



REMEMBER

Network transmission speed refers to the maximum speed that can be achieved over the network under ideal conditions. In reality, the actual throughput of an Ethernet network rarely reaches this maximum speed.

Ethernet operates at the first two layers of the OSI model — the physical and data link layers. However, Ethernet divides the data link layer into two separate layers: the logical link control (LLC) layer and the medium access control (MAC) layer. Figure 1-6 shows how the various elements of Ethernet match up to the OSI model.

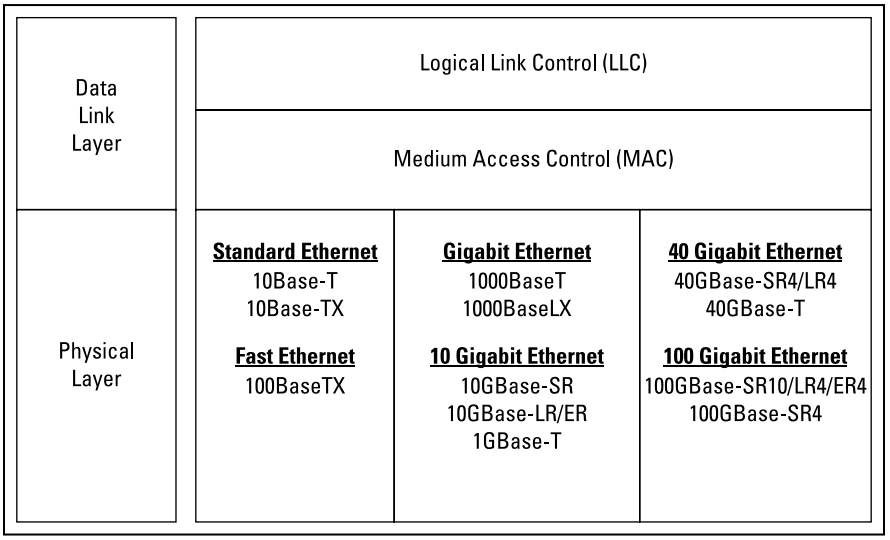


FIGURE 1-6: Ethernet and the OSI model.

The following sections describe the Ethernet standards in more detail.

Standard Ethernet

Standard Ethernet is the original Ethernet. It runs at 10 Mbps, which was considered fast in the 1970s but is excruciatingly slow by today's standards. Although plenty of existing Standard Ethernet is still in use, it's considered obsolete and should be replaced by Gigabit Ethernet as soon as possible.

Standard Ethernet came in three incarnations, depending on the type of cable used to string the network together:

- **10Base5:** This original Ethernet cable was thick (about as thick as your thumb), heavy, and difficult to work with. It's seen today only in museum exhibits.
- **10Base2:** This thinner type of coaxial cable (it resembles TV cable) became popular in the 1980s and lingered into the early 1990s. Plenty of 10Base2 cable is still in use, but it's rarely installed in new networks. 10Base2 (like 10Base5) uses a bus topology, so wiring a 10Base2 network involves running

cable from one computer to the next until all the computers are connected in a segment.

- **10Base-T:** Unshielded twisted-pair (UTP) cable became popular in the 1990s because it's easier to install, lighter, and more reliable, and it also offers more flexibility in how networks are designed. 10Base-T networks use a star topology with a hub at the center of each star. Although the maximum length of 10Base-T cable is only 100 meters, hubs can be chained to extend networks well beyond the 100-meter limit.

  10Base-T cable has four pairs of wires twisted together throughout the entire span of the cable. However, 10Base-T uses only two of these wire pairs, so the unused pairs are spares.



TIP

If you find yourself working with 10 Mbps Ethernet, spend a few moments enjoying your historical find. Then, as quickly as you can, update the entire network to Gigabit Ethernet.

Fast Ethernet

Fast Ethernet refers to Ethernet that runs at 100 Mbps, which is ten times the speed of Standard Ethernet. Although there are several varieties of Fast Ethernet, the most common is 100Base-TX, which transmits at 100 Mbps over just two pairs of UTP cable. 100 Mbps Ethernet requires at least Cat5 cable, but most networks are now wired with Cat5e or Cat6 cable, both of which are capable of gigabit speeds.

Gigabit Ethernet

Gigabit Ethernet is Ethernet running at a 1,000 Mbps, or 1 Gbps. Gigabit Ethernet was once considerably more expensive than Fast Ethernet, so it was used only when the improved performance justified the extra cost. However, today Gigabit Ethernet is the standard for nearly all desktop and laptop PCs. Two grades of cable are commonly used: Cat5e and Cat6. Cat6 is preferred because it can be used for even faster networks.

Beyond gigabit

Several varieties of Ethernet faster than 1 Gbps on copper cable are available:

- **2.5GBase-T:** 2.5 Gbps speed that can operate on Cat5e cable
- **5GBase-T:** 5 Gbps speed that requires Cat6 cable

- **10GBase-T:** 10 Gbps speed that requires Cat6a cable
- **25GBase-T:** 25 Gbps speed that requires Cat8 cable
- **40GBase-T:** 40 Gbps speed that requires Cat8 cable

There are also many varieties of 10 Gbps, 40 Gbps, 100 Gbps, 200 Gbps, 400 Gbps, and even 800 Gbps Ethernet that run on fiber-optic cable. And a standard for 1.6 terabit (1,600 Gbps) Ethernet is currently under development.

The TCP/IP Protocol Suite

TCP/IP, the protocol on which the internet is built, is not a single protocol but rather an entire suite of related protocols. TCP is even older than Ethernet. It was first conceived in 1969 by the U.S. Department of Defense. For more on the history of TCP/IP, see the sidebar, "The fascinating story of TCP/IP," later in this chapter. Currently, the IETF manages the TCP/IP protocol suite.

The TCP/IP suite is based on a four-layer model of networking similar to the seven-layer OSI model. Figure 1-7 shows how the TCP/IP model matches up with the OSI model and where some of the key TCP/IP protocols fit into the model. As you can see, the lowest layer of the model, the network interface layer, corresponds to the OSI model's physical and data link layers. TCP/IP can run over a wide variety of network interface layer protocols, including Ethernet, as well as other protocols, such as Token Ring and Fiber Distributed Data Interface (FDDI), an older standard for fiber-optic networks.

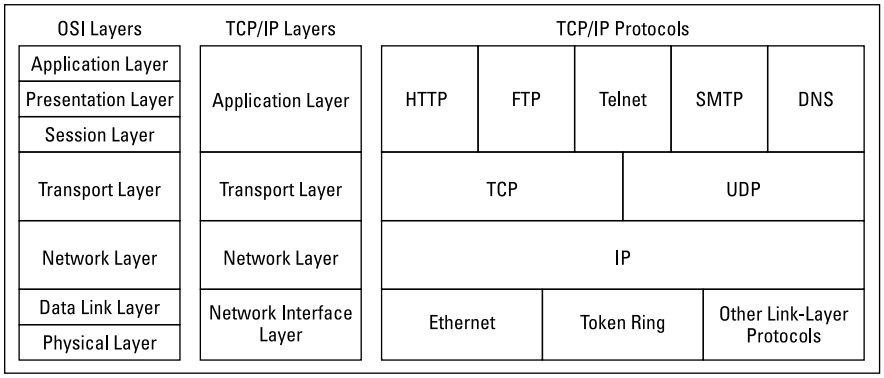


FIGURE 1-7: TCP/IP and the OSI model.

The application layer of the TCP/IP model corresponds to the upper three layers of the OSI model — the session, presentation, and application layers. Many protocols can be used at this level. A few of the most popular are Hypertext Transfer Protocol (HTTP), FTP, Telnet, SMTP, DNS, and Simple Network Management Protocol (SNMP).

In the following sections, I point out a few more details on the three most important protocols in the TCP/IP suite: IP, TCP, and UDP.

IP

IP is a network layer protocol responsible for delivering packets to network devices. The IP protocol uses logical IP addresses to refer to individual devices rather than physical (MAC) addresses. Address Resolution Protocol (ARP) handles the task of converting IP addresses to MAC addresses.

10BASE WHAT?

The names of Ethernet cable standards resemble the audible signals a quarterback might shout at the line of scrimmage. In reality, the cable designations consist of three parts:

- **The first number is the speed of the network in Mbps.** So, 10Base-T is for 10 Mbps networks (Standard Ethernet), 100Base-TX is for 100 Mbps networks (Fast Ethernet), and 1000Base-T is for 1,000 Mbps networks (Gigabit Ethernet). Above 1,000 Mbps, the letter G is added to indicate Gigabit speeds. Thus, 10GBase-T, 40GBase-T, and 100GBase-T are for 10 Gbps, 40 Gbps, and 100 Gbps networks, respectively.
- ***Base* (short for *baseband*) indicates the type of network transmission that the cable uses.** Baseband transmissions carry one signal at a time and are relatively simple to implement. The alternative to baseband is *broadband,* which can carry more than one signal at a time but is more difficult to implement. At one time, broadband incarnations of the 802.*x* networking standards existed, but they have all but fizzled due to lack of use.
- **The tail end of the designation indicates the cable type.** For coaxial cables, a number is used that roughly indicates the maximum length of the cable in hundreds of meters. 10Base5 cables can run up to 500 meters. 10Base2 cables can run up to 185 meters. (The IEEE rounded 185 up to 200 to come up with the name 10Base2.) If the designation ends with a T, twisted-pair cable is used; other letters are used for other types of cables.

Because IP addresses consist of a network part and a host part, IP is a routable protocol. As a result, IP can forward a packet to another network if the host isn't on the current network. After all, the capability to route packets across networks is where IP gets its name. An *internet* is a just a series of two or more connected TCP/IP networks that can be reached by routing.

TCP

TCP is a connection-oriented transport layer protocol. TCP lets a device reliably send a packet to another device on the same network or on a different network. TCP ensures that each packet is delivered, if at all possible, by establishing a connection with the receiving device and then sending the packets. If a packet doesn't arrive, TCP resends the packet. The connection is closed only after the packet has been successfully delivered or an unrecoverable error condition has occurred.

One key aspect of TCP is that it's always used for one-to-one communications. In other words, TCP allows a single network device to exchange data with another single network device. TCP isn't used to broadcast messages to multiple network recipients. Instead, UDP is used for that purpose.

Many well-known application layer protocols rely on TCP. For example, when a user running a web browser requests a page, the browser uses HTTP to send a request via TCP to a web server. When that web server receives the request, it uses HTTP to send the requested web page back to the browser, again via TCP. Other application layer protocols that use TCP include Telnet (for terminal emulation), FTP (for file exchange), and SMTP (for email).

UDP

UDP is a connectionless transport layer protocol used when the overhead of a connection isn't required. After UDP has placed a packet on the network (via the IP), it forgets about it. UDP doesn't guarantee that the packet arrives at its destination. Most applications that use UDP simply wait for any replies expected as a result of packets sent via UDP. If a reply doesn't arrive within a certain period of time, the application either sends the packet again or gives up.

Probably the best-known application layer protocol that uses UDP is the Domain Name System (DNS). When an application needs to access a domain name (such as `wiley.com`), DNS sends a UDP packet to a DNS server to look up the domain. When the server finds the domain, it returns the domain's IP address in another UDP packet.

THE FASCINATING STORY OF TCP/IP

Some people are fascinated by history. They subscribe to cable TV just to get the History Channel. If you're one of those history buffs, you may be interested in the following chronicle of TCP/IP's humble origins. (For maximum effect, play some melancholy violin music in the background as you read the rest of this sidebar.)

In the summer of 1969, the four mop-topped singers from Liverpool were breaking up. The war in Vietnam was escalating. Astronauts Neil Armstrong and Buzz Aldrin walked on the moon. And the U.S. Department of Defense built a computer network called ARPANET to link its defense installations with several major universities throughout the United States.

By the early 1970s, ARPANET was becoming difficult to manage. So, it was split into two networks: one for military use (called MILNET) and the other for nonmilitary use. The nonmilitary network retained the name ARPANET. To link MILNET with ARPANET, a new method of connecting networks, IP, was invented.

The whole purpose of IP was to enable these two networks to communicate with each other. Fortunately, the designers of IP realized that it wouldn't be too long before other networks wanted to join in the fun, so they designed IP to allow for more than two networks. In fact, their ingenious design allowed for tens of thousands of networks to communicate via IP.

The decision was a fortuitous one, as the internet quickly began to grow. By the mid-1980s, the original ARPANET reached its limits. Just in time, the National Science Foundation (NSF) decided to get into the game. NSF had built a network called NSFNET to link its huge supercomputers. NSFNET replaced ARPANET as the new background for the internet. Around that time, such magazines as *Time* and *Newsweek* began writing articles about this new phenomenon called the internet, and the *Net* (as it became nicknamed) began to grow like wildfire. Soon NSFNET couldn't keep up with the growth, so several private commercial networks took over management of the internet backbone. The internet has grown at a dizzying rate ever since, and nobody knows how long this frenetic growth rate will continue. One thing is sure: TCP/IP is now the most popular networking protocol in the world.

Other Protocols Worth Knowing About

Although the vast majority of networks now use Ethernet and TCP/IP, a few other networking protocols are still in use and are, therefore, worth knowing about:

- **AppleTalk:** An obsolete suite of network protocols introduced by Apple in the 1980s and finally abandoned in 2009. The AppleTalk suite included a physical and data link layer protocol called LocalTalk, but it could also work with standard lower-level protocols, including Ethernet and Token Ring.
- **IPX/SPX:** A protocol suite made popular in the 1980s by Novell for use with its NetWare servers. TCP/IP has become so dominant that IPX/SPX is rarely used now.
- **Network Basic Input/Output System (NetBIOS):** The basic API for network services on Windows computers. It's installed automatically when you install TCP/IP, but it doesn't show up as a separate protocol when you view the network connection properties (refer to Figure 1-1). NetBIOS is a session layer protocol that can work with transport layer protocols, such as TCP, SPX, or NetBEUI (see the following bullet).
- **Network BIOS Extended User Interface (NetBEUI):** A transport layer protocol designed for early IBM and Microsoft networks. NetBEUI is now considered obsolete.
- **Systems Network Architecture (SNA):** An IBM networking architecture dating back to the 1970s, when mainframe computers roamed the earth and PCs had barely emerged from the primordial computer soup. SNA was designed primarily to support huge terminals such as airline reservations and banking systems, with tens of thousands of terminals attached to central host computers. Now that IBM mainframes that support TCP/IP and mainframe terminal systems have all but vanished, SNA is beginning to fade away. Still, many networks that incorporate mainframe computers have to contend with SNA.

IN THIS CHAPTER

- Introducing the internet
- Familiarizing yourself with TCP/IP standards
- Figuring out how TCP/IP lines up with the OSI Reference Model
- Discovering important TCP/IP applications

Chapter 2
TCP/IP and the Internet

Many years ago, Transmission Control Protocol/Internet Protocol (TCP/IP) was known primarily as the protocol of the internet. The biggest challenge of getting a local area network (LAN) connected to the internet was figuring out how to mesh TCP/IP with the proprietary protocols that were the basis of the LANs — most notably Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) used by Novel networks and NetBIOS Extended User Interface (NetBEUI) used by Microsoft networks.

Eventually, both IPX/SPX and NetBIOS gave way to TCP/IP as the basis for local area networking, eliminating the challenge of translating IPX/SPX or NetBEUI to TCP/IP. As a result, TCP/IP is not just the protocol of the internet now, but it's also the protocol on which most LANs are based.

This chapter is a gentle introduction to the internet in general and the TCP/IP suite of protocols in particular. After I get the introductions out of the way, you'll be able to focus more in-depth on the detailed TCP/IP information given in the remaining chapters of Book 2.

What Is the Internet?

The Goliath of all computer networks, the internet links hundreds of millions of computer users throughout the world. Strictly speaking, the internet is a network of networks. It consists of hundreds of thousands of separate computer networks, all interlinked, so that a user on any of those networks can reach out and potentially touch a user on any of the other networks. This network of networks connects more than a billion computers to each other. (That's right, *billion* with a *b*.)

One of the official documents (RFC 2026) of the Internet Engineering Task Force (IETF) defines the internet as "a loosely organized international collaboration of autonomous, interconnected networks." Broken down piece by piece, this definition encompasses several key aspects of what the internet is:

- **Loosely organized:** No single organization has authority over the internet. As a result, the internet is not highly organized. Online services, such as America Online or MSN, are owned and operated by individual companies that control exactly what content appears on the service and what software can be used with the service. No one exercises that kind of control over the internet. As a result, you can find just about any kind of material imaginable on the internet. No one guarantees the accuracy of information that you find on the internet, so you have to be careful as you work your way through the labyrinth.


TECHNICAL STUFF

JUST HOW BIG IS THE INTERNET?

Because the internet is not owned or controlled by any one organization, no one knows how big the internet really is. Several organizations do attempt to periodically determine the size of the internet, including the Internet Systems Consortium (ISC), which completed its last survey in January 2019 and found that well over a billion host computers are connected to the internet. The first year the ISC did the survey (1993), it found only 1.3 million host computers. It passed 10 million hosts in 1996, 100 million hosts in 2000, and edged over 1 billion hosts in 2014.

Unfortunately, no one knows how many actual users are on the internet. Each host can support a single user — or in the case of domains, dozens, hundreds, thousands, or perhaps even millions of users. No one really knows.

In fact, the ISC gave up on trying to count the number of hosts on the internet. The January 2019 survey was its last. If you're interested, you can check its historical survey data at `www.isc.org/network/survey`.

- **International:** More than 300 countries are represented on the internet, from Ascension Island to Zimbabwe.
- **Collaboration:** The internet exists only because many different organizations cooperate to provide the services and support needed to sustain it. For example, much of the software that drives the internet is open source software that's developed collaboratively by programmers throughout the world, who constantly work to improve the code.
- **Autonomous:** The internet community respects that organizations that join the internet are free to make their own decisions about how they configure and operate their networks. Although legal issues sometimes boil up, for the most part, each player on the internet operates independently.
- **Interconnected:** The whole key to the internet is the concept of *interconnection,* which uses standard protocols that enable networks to communicate with each other. Without the interconnection provided by the TCP/IP protocol, the internet would not exist.
- **Networks:** The internet would be completely unmanageable if it consisted of half a billion individual users, all interconnected. That's why the internet is often described as a network of networks. Most individual users on the internet don't access the internet directly. Instead, they access the internet indirectly through another network, which may be a LAN in a business or academic environment, or a dialup or broadband network provided by an internet service provider (ISP). In each case, however, the users of the local network access the internet via a gateway IP router.

  The internet is composed of several distinct types of networks: government agencies, such as the Library of Congress and the White House; military sites (did you ever see *War Games* or any of the *Terminator* movies?); educational institutions, such as universities and colleges (and their libraries); businesses, such as Microsoft and IBM; ISPs, which allow individuals to access the internet; and commercial online services, such as America Online and MSN.

A Little Internet History

The internet has a fascinating history, if such things interest you. There's no particular reason why you should be interested in such things, of course, except that a superficial understanding of how the internet got started may help you to understand and cope with the way this massive computer network exists today. So here goes.

The internet traces its beginnings back to a small network called ARPANET, built by the Department of Defense in 1969 to link defense installations. ARPANET soon expanded to include not only defense installations but universities as well. In the 1970s, ARPANET was split into two networks: one for military use (renamed MILNET) and the original ARPANET (for nonmilitary use). The two networks were connected by a networking link called the *Internet Protocol* (IP), so called because it allowed communication between two networks.

The good folks who designed IP had the foresight to realize that soon, more than two networks would want to be connected. In fact, they left room for tens of thousands of networks to join the game, which is a good thing because it wasn't long before the internet began to take off.

By the mid-1980s, ARPANET was beginning to reach the limits of what it could do. Enter the National Science Foundation (NSF), which set up a nationwide network designed to provide access to huge *supercomputers,* those monolithic computers used to discover new prime numbers and calculate the orbits of distant galaxies. The supercomputers were never put to much use, but the network that was put together to support the supercomputers — NSFNET — was used. In fact, NSFNET replaced ARPANET as the new backbone for the internet.

Then, out of the blue, it seemed as if the whole world became interested in the internet. Stories about it appeared in *Time* and *Newsweek.* Any company that had "dot com" in its name practically doubled in value every month. Al Gore claimed he invented the internet (well, not really — but he was *accused* of taking credit). The Net began to grow so fast that even NSFNET couldn't keep up, so private commercial networks got into the game. The size of the internet nearly doubled every year for most of the 1990s. Then, in the first few years of the millennium, the growth rate slowed a bit. However, the internet still seems to be growing at the phenomenal rate of about 30 to 50 percent per year, and who knows how long this dizzying rate of growth will continue.

TCP/IP Standards and RFCs

The TCP/IP protocol standards that define how the internet works are managed by the IETF. However, the IETF doesn't impose standards. Instead, it simply oversees the process by which ideas are developed into agreed-upon standards.

An internet standard is published in the Request for Comments (RFC) document. When a document is accepted for publication, it is assigned an RFC number by the

IETF. The RFC is then published. After it's published, an RFC is never changed. If a standard is enhanced, the enhancement is covered in a separate RFC.

Thousands of RFCs are available from the IETF's RFC Editor website (`www.rfc-editor.org`). The oldest RFC is RFC 0001, published in April 1969. It describes how the host computers communicated with each other in the original ARPANET. As of this writing, the most recent proposed standard is RFC 9565, entitled "An Update to the tcpControlBits IP Flow Information Export (IPFIX) Information Element."

Not all RFCs represent internet standards. The following paragraphs summarize the various types of RFC documents:

- **Internet Standards Track:** This type of RFC represents an internet standard. Standards Track RFCs have one of three maturity levels, as described in Table 2-1. An RFC enters circulation with Proposed Standard status but may be elevated to Draft Standard status — and, ultimately, to Internet Standard status.
- **Experimental specifications:** These are a result of research or development efforts. They're not intended to be standards, but the information they contain may be of use to the internet community.
- **Informational specifications:** These simply provide general information for the internet community.
- **Historic specifications:** These RFCs have been superseded by a more recent RFC and are thus considered obsolete.
- **Best current practice (BCP):** RFCs are documents that summarize the consensus of the internet community's opinion on the best way to perform an operation or procedure. BCPs are guidelines, not standards.
- **April Fools:** It's a long-standing tradition to post amusing RFCs on April Fools' Day. One of my favorites is RFC 1149, "A Standard for the Transmission of IP Datagrams on Avian Carriers." The specification calls for IP datagrams to be written in hexadecimal on scrolls of paper and secured to "avian carriers" with duct tape.

  The most recent April 1 RFC at the time of this writing is RFC 9564, "Faster Than Light Speed Protocol (FLIP)." It defines a protocol in which the receiver uses AI technology to predict the contents of packets before they have even been sent, thereby achieving faster-than-light internet speeds.

Table 2-2 summarizes the RFCs that apply to the key internet standards described in this book.

TABLE 2-1 Maturity Levels for Internet Standards Track RFCs

Maturity Level	Description
Proposed Standard	Generally stable, have resolved known design choices, are believed to be well understood, have received significant community review, and appear to enjoy enough community interest to be considered valuable.
Draft Standard	Well understood and known to be quite stable. At least two interoperable implementations must exist, developed independently from separate code bases. The specification is believed to be mature and useful.
Internet Standard	Have been fully accepted by the internet community as highly mature and useful standards.

TABLE 2-2 RFCs for Key Internet Standards

RFC	Date	Description
768	August 1980	User Datagram Protocol (UDP)
791	September 1981	Internet Protocol (IP)
792	September 1981	Internet Control Message Protocol (ICMP)
793	September 1981	Transmission Control Protocol (TCP)
826	November 1982	Ethernet Address Resolution Protocol (ARP)
950	August 1985	Internet Standard Subnetting Procedure
959	October 1985	File Transfer Protocol (FTP)
1034	November 1987	Domain Names — Concepts and Facilities (DNS)
1035	November 1987	Domain Names — Implementation and Specification (DNS)
1939	May 1996	Post Office Protocol Version 3 (POP3)
2131	March 1997	Dynamic Host Configuration Protocol (DHCP)
3376	November 1997	Internet Group Management Protocol (IGMP) (Updates RFC 2236 and 1112)
5321	October 2008	Simple Mail Transfer Protocol (SMTP)
7230 through 7235	June 2014	Hypertext Transfer Protocol – HTTP/1.1

The TCP/IP Protocol Framework

Like the seven-layer OSI Reference Model, TCP/IP protocols are based on a layered framework. TCP/IP has four layers, as shown in Figure 2-1. These layers are described in the following sections.

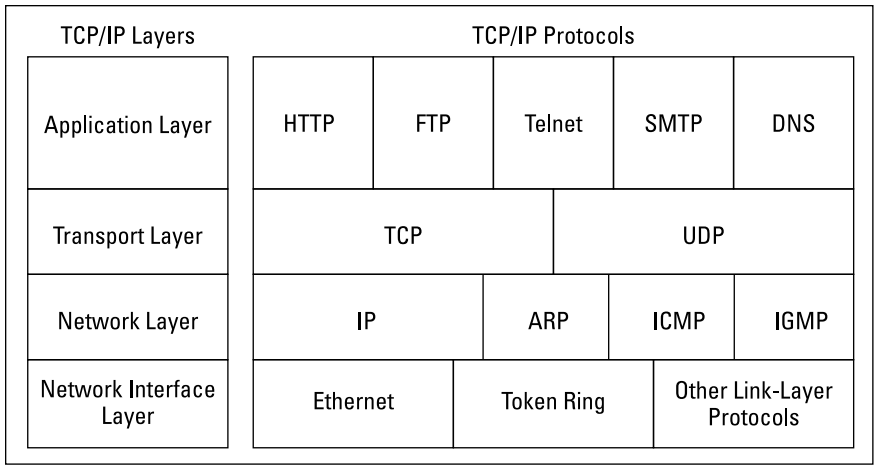


FIGURE 2-1: The four layers of the TCP/IP framework.

Network interface layer

The lowest level of the TCP/IP architecture is the network interface layer. It corresponds to the OSI physical and data link layers. You can use many different TCP/IP protocols at the network interface layer, including Ethernet and Token Ring for LANs and protocols such as X.25, Frame Relay, and ATM for wide area networks (WANs).

The network interface layer is assumed to be unreliable.

Network layer

The network layer is where data is addressed, packaged, and routed among networks. Several important internet protocols operate at the network layer:

- **Internet Protocol (IP):** A routable protocol that uses IP addresses to deliver packets to network devices. IP is an intentionally unreliable protocol, so it doesn't guarantee delivery of information.
- **Address Resolution Protocol (ARP):** Resolves IP addresses to hardware Media Access Control (MAC) addresses, which uniquely identify hardware devices.
- **Internet Control Message Protocol (ICMP):** Sends and receives diagnostic messages. ICMP is the basis of the ubiquitous `ping` command.
- **Internet Group Management Protocol (IGMP):** Used to multicast messages to multiple IP addresses at once.

Transport layer

The transport layer is where sessions are established and data packets are exchanged between hosts. Two core protocols are found at this layer:

- **Transmission Control Protocol (TCP):** Provides reliable connection-oriented transmission between two hosts. TCP establishes a session between hosts, and then ensures delivery of packets between the hosts.
- **User Datagram Protocol (UDP):** Provides connectionless, unreliable, one-to-one or one-to-many delivery.

Application layer

The application layer of the TCP/IP model corresponds to the session, presentation, and application layers of the OSI Reference Model. A few of the most popular application layer protocols are

- **Hypertext Transfer Protocol (HTTP):** The core protocol of the World Wide Web.
- **File Transfer Protocol (FTP):** A protocol that enables a client to send and receive complete files from a server.
- **Telnet:** The protocol that lets you connect to another computer on the internet in a terminal emulation mode.
- **Simple Mail Transfer Protocol (SMTP):** One of several key protocols that are used to provide email services.
- **Domain Name System (DNS):** The protocol that allows you to refer to other host computers by using names rather than numbers.

IN THIS CHAPTER

- Delving into the binary system
- Digging into IP addresses
- Finding out how subnetting works
- Delving into ports
- Looking at network address translation

Chapter 3
IP Addresses

One of the most basic components of TCP/IP is Internet Protocol (IP) addressing. Every device on a TCP/IP network must have a unique IP address. In this chapter, I describe the ins and outs of these IP addresses. Enjoy!

Understanding Binary

Before you can understand the details of how IP addressing works, you need to understand how the binary numbering system works because binary is the basis of IP addressing. If you already understand binary, please skip to the section "Introducing IP Addresses." I don't want to bore you with stuff that's too basic.

Counting by ones

Binary is a counting system that uses only two numerals: 0 and 1. In the decimal system (with which most people are accustomed), you use ten numerals: 0–9. In an ordinary decimal number — such as 3,482 — the rightmost digit represents ones; the next digit to the left, tens; the next, hundreds; the next, thousands; and so on. These digits represent powers of ten: first 10^0 (which is 1); next, 10^1 (10); then 10^2 (100); then 10^3 (1,000); and so on.

In binary, you have only two numerals rather than ten, which is why binary numbers look somewhat monotonous, as in 110011, 101111, and 100001.

The positions in a binary number (called *bits* rather than *digits*) represent powers of two rather than powers of ten: 1, 2, 4, 8, 16, 32, and so on. To figure the decimal value of a binary number, you multiply each bit by its corresponding power of two and then add the results. The decimal value of binary 10111, for example, is calculated as follows:

$$1 \times 2^0 = 1 \times 1 = 1$$

$$1 \times 2^1 = 1 \times 2 = 2$$

$$1 \times 2^2 = 1 \times 4 = 4$$

$$0 \times 2^3 = 0 \times 8 = 0$$

$$1 \times 2^4 = 1 \times 16 = 16$$

$$\text{Total} = 1 + 2 + 4 + 0 + 16 = 23$$

Fortunately, converting a number between binary and decimal is something a computer is good at — so good, in fact, that you're unlikely ever to need to do any conversions yourself. The point of learning binary is not to be able to look at a number such as 1110110110110 and say instantly, "Ah! Decimal 7,606!" (If you could do that, Piers Morgan would probably interview you, and they would even make a movie about you.)

Instead, the point is to have a basic understanding of how computers store information and — most important — to understand how the binary counting system works, which I describe in the following section.

Here are some of the more interesting characteristics of binary and how the system is similar to and differs from the decimal system:

- **In decimal, the number of decimal places allotted for a number determines how large the number can be.** If you allot six digits, for example, the largest number possible is 999,999. Because 0 is itself a number, however, a six-digit number can have any of 1 million different values.

  Similarly, the number of bits allotted for a binary number determines how large that number can be. If you allot eight bits, the largest value that number can store is 11111111, which happens to be 255 in decimal.


TIP

- **To quickly figure how many different values you can store in a binary number of a given length, use the number of bits as an exponent of two.** An eight-bit binary number, for example, can hold 2^8 values. Because 2^8 is 256, an eight-bit number can have any of 256 different values. This is why a *byte* — eight bits — can have 256 different values.
- **This "powers of two" thing is why computers don't use nice, even, round numbers in measuring such values as memory or disk space.** A value of 1K, for example, is not an even 1,000 bytes: It's actually 1,024 bytes because 1,024 is 2^{10}. Similarly, 1MB is not an even 1,000,000 bytes but instead 1,048,576 bytes, which happens to be 2^{20}.


REMEMBER

One basic test of computer nerddom is knowing your powers of two because they play such an important role in binary numbers. Just for the fun of it, but not because you really need to know, Table 3-1 lists the powers of two up to 32.

TABLE 3-1 Powers of Two

Power	Bytes	Kilobytes	Power	Bytes	K, MB, or GB
2^1	2		2^{17}	131,072	128K
2^2	4		2^{18}	262,144	256K
2^3	8		2^{19}	524,288	512K
2^4	16		2^{20}	1,048,576	1MB
2^5	32		2^{21}	2,097,152	2MB
2^6	64		2^{22}	4,194,304	4MB
2^7	128		2^{23}	8,388,608	8MB
2^8	256		2^{24}	16,777,216	16MB
2^9	512		2^{25}	33,554,432	32MB
2^{10}	1,024	1K	2^{26}	67,108,864	64MB
2^{11}	2,048	2K	2^{27}	134,217,728	128MB
2^{12}	4,096	4K	2^{28}	268,435,456	256MB
2^{13}	8,192	8K	2^{29}	536,870,912	512MB
2^{14}	16,384	16K	2^{30}	1,073,741,824	1GB
2^{15}	32,768	32K	2^{31}	2,147,483,648	2GB
2^{16}	65,536	64K	2^{32}	4,294,967,296	4GB

Table 3-1 also shows the common shorthand notation for various powers of two. The abbreviation *K* represents 2^{10} (1,024). The *M* in *MB* stands for 2^{20}, or 1,024K, and the *G* in *GB* represents 2^{30}, which is 1,024MB. These shorthand notations don't have anything to do with TCP/IP, but they're commonly used for measuring computer disk and memory capacities, so I thought I'd throw them in at no charge because the table had extra room.

Doing the logic thing

One of the great things about binary is that it's very efficient at handling special operations: namely, logical operations. Four basic logical operations exist although additional operations are derived from the basic four operations. Three of the operations — `AND`, `OR`, and `XOR` — compare two binary digits (bits). The fourth (`NOT`) works on just a single bit.

The following list summarizes the basic logical operations:

- **AND:** Compares two binary values. If both values are 1, the result of the AND operation is 1. If one or both of the values are 0, the result is 0.
- **OR:** Compares two binary values. If at least one value is 1, the result of the OR operation is 1. If both values are 0, the result is 0.
- **XOR:** Compares two binary values. If one of them is 1, the result is 1. If both values are 0 or if both values are 1, the result is 0.
- **NOT:** Doesn't compare two values but simply changes the value of a single binary value. If the original value is 1, NOT returns 0. If the original value is 0, NOT returns 1.

Table 3-2 summarizes how AND, OR, and XOR work.

TABLE 3-2 Logical Operations for Binary Values

First Value	Second Value	AND	OR	XOR
0	0	0	0	0
0	1	0	1	1
1	0	0	1	1
1	1	1	1	0

Logical operations are applied to binary numbers that have more than one binary digit by applying the operation one bit at a time. The easiest way to do this manually is to line the two binary numbers on top of one another and then write the result of the operation beneath each binary digit. The following example shows how you would calculate 10010100 AND 11011101:

```
     10010100
AND  11011101
     10010100
```

As you can see, the result is 10010100.

Working with the binary Windows Calculator

The Calculator program that comes with all versions of Windows has a special Programmer mode that many users don't know about. When you flip the Calculator into this mode, you can do instant binary and decimal conversions, which can occasionally come in handy when you're working with IP addresses.

To launch the Calculator, press the Windows key, type **Calculator** into the search bar, and then press Enter. When the Calculator is open, you can switch to Programmer mode by clicking the menu icon in the upper-left corner of the Calculator and choosing Programmer. In Programmer mode, you can do calculations in decimal (base 10), hexadecimal (base 16), octal (base 8), or binary (base 2). Figure 3-1 shows the Programmer mode for the latest version of the Calculator (as of this writing).

In the middle left of the Calculator window, you can see the current value displayed by the Calculator in hexadecimal (HEX), decimal (DEC), octal (OCT), and binary (BIN). You can also tell which base the main display shows, because it's highlighted with a bar. In Figure 3-1, the current mode is DEC, so the decimal value 100 is shown in large text in the upper-middle part of the display. You can switch the main display to hexadecimal, octal, or binary by clicking HEX, OCT, or BIN, respectively.

You can also see the current value in all four bases. Thus, in the figure, you can see that decimal 100 is 64 in hexadecimal, 144 in octal, and 01100100 in binary.

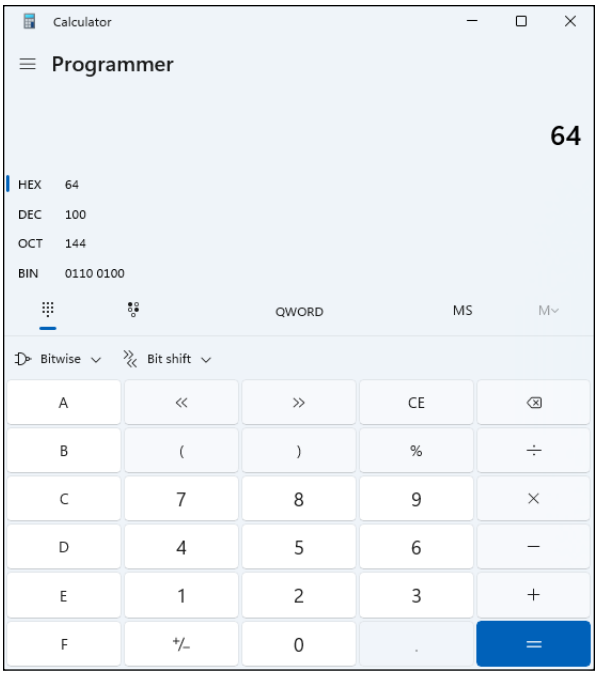


FIGURE 3-1: The Windows Calculator in Programmer mode.

Here are a few other things to note about the Programmer mode of the Calculator:

- Although you can convert decimal values to binary values with the programmer Calculator, the Calculator can't handle the dotted-decimal IP address format that's described later in this chapter. To convert a dotted-decimal address to binary, just convert each octet separately. For example, to convert 172.65.48.120 to binary, first convert 172; then convert 65; then convert 48; and finally, convert 120.
- The Programmer Calculator has several features that are designed specifically for binary operations, such as AND, OR, XOR, and so on.
- The Programmer Calculator also has many other cool features for working with binary values. Spend some time exploring it when you have a few minutes!

Introducing IP Addresses

An *IP address* is a number that uniquely identifies every host on an IP network. IP addresses operate at the network layer of the TCP/IP protocol stack, so they are independent of lower-level data link layer MAC addresses, such as Ethernet MAC addresses.

IP addresses are 32-bit binary numbers, which means that theoretically, a maximum of something in the neighborhood of 4 billion unique host addresses can exist throughout the internet. You'd think that would be enough, but TCP/IP places certain restrictions on how IP addresses are allocated. These restrictions severely limit the total number of usable IP addresses. Many experts predict that we will run out of IP addresses soon. However, new techniques for working with IP addresses have helped to alleviate this problem, and a standard for 128-bit IP addresses has been adopted, though it still is not yet in widespread use.



32-bit IP addresses are technically known as IP4 addresses, because they reflect version 4 of the IP. A newer version of IP, known as IP6, uses 128-bit addresses. For more information about IP6, see the sidebar "What about IP6?"

Networks and hosts

The Internet Protocol's primary purpose is to enable communications between networks. As a result, a 32-bit IP address actually consists of two parts:

- **The network ID (or network address):** Identifies the network on which a host computer can be found
- **The host ID (or host address):** Identifies a specific device on the network indicated by the network ID

Most of the complexity of working with IP addresses has to do with figuring out which part of the complete 32-bit IP address is the network ID and which part is the host ID, as described in the following sections.



As I describe the details of how host IDs are assigned, you may notice that two host addresses seem to be unaccounted for. For example, the Class C addressing scheme, which uses eight bits for the host ID, allows only 254 hosts — not the 256 hosts you'd expect. The host ID can't be 0 (the host ID is all zeros) because that address is always reserved to represent the network itself. And the host ID can't be 255 (the host ID is all ones) because that host ID is reserved for use as a broadcast request that's intended for all hosts on the network.

The dotted-decimal dance

IP addresses are usually represented in a format known as *dotted-decimal notation.* In dotted-decimal notation, each group of eight bits — an *octet* — is represented by its decimal equivalent. For example, consider the following binary IP address:

```
11000000101010001000100000011100
```

To convert this value to dotted-decimal notation, first divide it into four octets, as follows:

```
11000000 10101000 10001000 00011100
```

Then, convert each of the octets to its decimal equivalent:

```
11000000 10101000 10001000 00011100
192      168      136      28
```

Then, use periods to separate the four decimal numbers, like this:

```
192.168.136.28
```

This is the format in which you'll usually see IP addresses represented.

Figure 3-2 shows how the 32 bits of an IP address are broken down into four octets of eight bits each. As you can see, the four octets of an IP address are often referred to as *w*, *x*, *y*, and *z*.

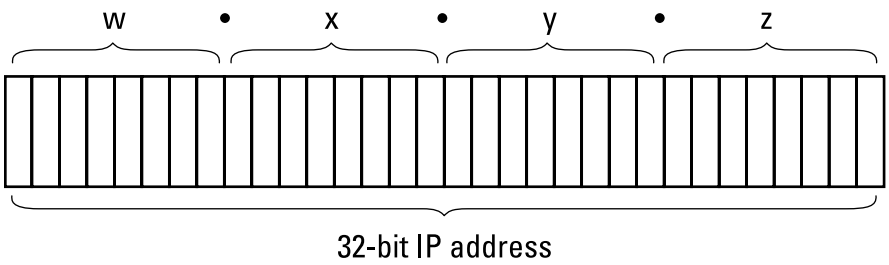


FIGURE 3-2: Octets and dotted-decimal notation.

Classifying IP Addresses

When the original designers of the IP protocol created the IP addressing scheme, they could have assigned an arbitrary number of IP address bits for the network ID. The remaining bits would then be used for the host ID. For example, suppose that the designers decided that half of the address (16 bits) would be used for the network, and the remaining 16 bits would be used for the host ID. The result of that scheme would be that the internet could have a total of 65,536 networks, and each of those networks could have 65,536 hosts.

In the early days of the internet, this scheme probably seemed like several orders of magnitude more than would ever be needed. However, the IP designers realized from the start that few networks would actually have tens of thousands of hosts. Suppose that a network of 1,000 computers joins the internet and is assigned one of these hypothetical network IDs. Because that network will use only 1,000 of its 65,536 host addresses, more than 64,000 IP addresses would be wasted.

WHAT ABOUT IPV6?

Most of the current internet is based on version 4 of the Internet Protocol, also known as IPv4. IPv4 has served the internet well for more than 40 years. However, the growth of the internet has put a lot of pressure on IPv4's limited 32-bit address space. This chapter describes how IPv4 has evolved to make the best possible use of 32-bit addresses. Eventually, though, all the addresses will be assigned, and the IPv4 address space will be filled to capacity. When that happens, the internet will have to migrate to the next version of IP, known as IPv6.

IPv6 is also called *IP next generation,* or *IPng,* in honor of the favorite television show of most internet gurus, *Star Trek: The Next Generation.*

IPv6 offers several advantages over IPv4, but the most important is that it uses 128 bits for internet addresses instead of 32 bits. The number of host addresses possible with 128 bits is a number so large that it would have made Carl Sagan proud. It doesn't just double or triple the number of available addresses, or even a thousand-fold or even a million-fold. Just for the fun of it, here is the number of unique internet addresses provided by IPv6:

340,282,366,920,938,463,463,374,607,431,768,211,456

This number is so large it defies understanding. If the Internet Assigned Numbers Authority (IANA) had been around at the creation of the universe and started handing out IPv6 addresses at a rate of one per millisecond — that is, 1,000 addresses every second — it would now, 15 billion years later, have not yet allocated even 1 percent of the available addresses.

The transition from IPv4 to IPv6 has been slow. IPv6 is available on all new computers and has been supported on Windows since Windows XP Service Pack 1 (released in 2002). However, most ISPs still base their service on IPv4. Thus, the internet will continue to be driven by IPv4 for at least a few more years.

(continued)

(continued)

> ***Note:*** This is now the eighth edition of this book. Every previous edition of this book, all the way back to the very first edition published in 2004, has had this very sidebar. In 2004, I said that the internet would continue to be driven by IPv4 for at least a few more years. "A few more years" has morphed into 14 years, and we're still living in the world of IPv4. Make no mistake: The world will eventually run out of IPv4 addresses, and we'll have to migrate to IPv6 . . . in a few years, whatever that means.

As a solution to this problem, the idea of IP address classes was introduced. The IP protocol defines five different address classes: A, B, C, D, and E. Each of the first three classes, A–C, uses a different size for the network ID and host ID portion of the address. Class D is for a special type of address called a *multicast address.* Class E is an experimental address class that isn't used.

The first four bits of the IP address are used to determine into which class a particular address fits, as follows:

- **Class A:** The first bit is zero.
- **Class B:** The first bit is one, and the second bit is zero.
- **Class C:** The first two bits are both one, and the third bit is zero.
- **Class D:** The first three bits are all one, and the fourth bit is zero.
- **Class E:** The first four bits are all one.

Because Class D and E addresses are reserved for special purposes, I focus the rest of the discussion here on Class A, B, and C addresses. Table 3-3 summarizes the details of each address class.

TABLE 3-3 IP Address Classes

Class	Address Number Range	Starting Bits	Length of Network ID	Number of Networks	Hosts
A	1–126.*x.y.z*	0	8	126	16,777,214
B	128–191.*x.y.z*	10	16	16,384	65,534
C	192–223.*x.y.z*	110	24	2,097,152	254

Class A addresses

Class A addresses are designed for very large networks. In a Class A address, the first octet of the address is the network ID, and the remaining three octets are the host ID. Because only eight bits are allocated to the network ID and the first of these bits is used to indicate that the address is a Class A address, only 126 Class A networks can exist in the entire internet. However, each Class A network can accommodate more than 16 million hosts.

Only about 40 Class A addresses are actually assigned to companies or organizations. The rest are either reserved for use by the Internet Assigned Numbers Authority (IANA) or are assigned to organizations that manage IP assignments for geographic regions such as Europe, Asia, and Latin America.

In case you're interested, you can find a complete list of all the Class A address assignments at `www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml`.

You may have noticed in Table 3-3 that Class A addresses end with 126.*x.y.z*, and Class B addresses begin with 128.*x.y.z*. What happened to 127.*x.y.z*? This special range of addresses is reserved for loop-back testing, so these addresses aren't assigned to public networks.


TIP

The special address 127.0.0.1 is called the *loop-back address.* A device at any IP address that sends a message to 127.0.0.1 is sending a message to itself. This may sound useless, but it actually plays an important role in troubleshooting network problems.

Class B addresses

In a Class B address, the first two octets of the IP address are used as the network ID, and the second two octets are used as the host ID. Thus, a Class B address comes close to my hypothetical scheme of splitting the address down the middle, using half for the network ID and half for the host ID. It isn't identical to this scheme, however, because the first two bits of the first octet are required to be 10, in order to indicate that the address is a Class B address. As a result, a total of 16,384 Class B networks can exist. All Class B addresses fall within the range 128.*x.y.z* to 191.*x.y.z*. Each Class B address can accommodate more than 65,000 hosts.

The problem with Class B networks is that even though they are much smaller than Class A networks, they still allocate far too many host IDs. Very few networks have tens of thousands of hosts. Thus, careless assignment of Class B addresses can lead to a large percentage of the available host addresses being wasted on organizations that don't need them.

Class C addresses

In a Class C address, the first three octets are used for the network ID, and the fourth octet is used for the host ID. With only eight bits for the host ID, each Class C network can accommodate only 254 hosts. However, with 24 network ID bits, Class C addresses allow for more than 2 million networks.

The problem with Class C networks is that they're too small. Although few organizations need the tens of thousands of host addresses provided by a Class B address, many organizations need more than a few hundred. The large discrepancy between Class B networks and Class C networks is what led to the development of *subnetting,* which I describe in the next section.

Subnetting

Subnetting is a technique that lets network administrators use the 32 bits available in an IP address more efficiently by creating networks that aren't limited to the scales provided by Class A, B, and C IP addresses. With subnetting, you can create networks with more realistic host limits.

Subnetting provides a more flexible way to designate which portion of an IP address represents the network ID and which portion represents the host ID. With standard IP address classes, only three possible network ID sizes exist: 8 bits for Class A, 16 bits for Class B, and 24 bits for Class C. Subnetting lets you select an arbitrary number of bits to use for the network ID.

Two reasons compel people to use subnetting. The first is to allocate the limited IP address space more efficiently. If the internet were limited to Class A, B, or C addresses, every network would be allocated 254, 64,000, or 16 million IP addresses for host devices. Although many networks with more than 254 devices exist, few (if any) exist with 64,000, let alone 16 million. Unfortunately, any network with more than 254 devices would need a Class B allocation and probably waste tens of thousands of IP addresses.

The second reason for subnetting is that even if a single organization has thousands of network devices, operating all those devices with the same network ID would slow the network to a crawl. The way TCP/IP works dictates that all the computers with the same network ID must be on the same physical network. The physical network comprises a single *broadcast domain,* which means that a single network medium must carry all the traffic for the network. For performance reasons, networks are usually segmented into broadcast domains that are smaller than even Class C addresses provide.

Subnets

A *subnet* is a network that falls within a Class A, B, or C network. Subnets are created by using one or more of the Class A, B, or C host bits to extend the network ID. Thus, instead of the standard 8-, 16-, or 24-bit network ID, subnets can have network IDs of any length.

Figure 3-3 shows an example of a network before and after subnetting has been applied. In the unsubnetted network, the network has been assigned the Class B address 144.28.0.0. All the devices on this network must share the same broadcast domain.

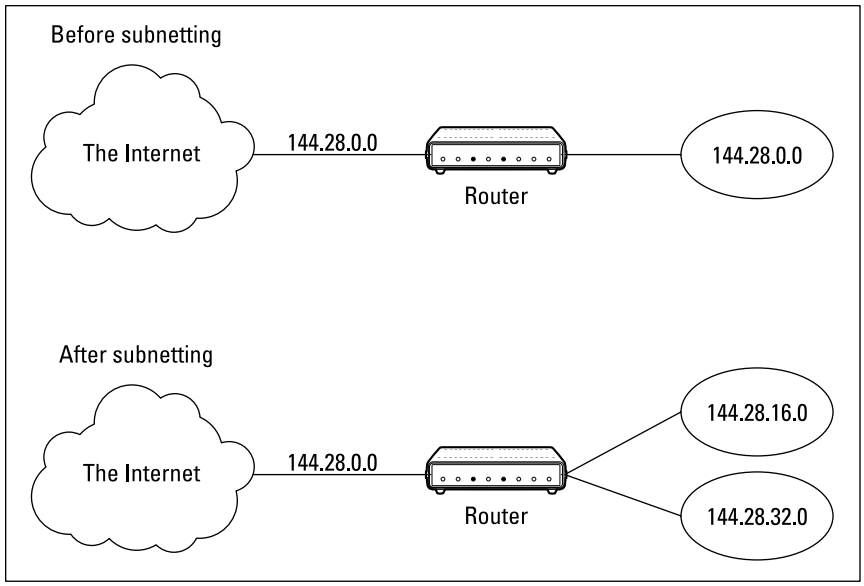


FIGURE 3-3: A network before and after subnetting.

In the second network, the first four bits of the host ID are used to divide the network into two small networks, identified as subnets 16 and 32. To the outside world (that is, on the other side of the router), these two networks still appear to be a single network identified as 144.28.0.0. For example, the outside world considers the device at 144.28.16.22 to belong to the 144.28.0.0 network. As a result, a packet sent to this device will be delivered to the router at 144.28.0.0. The router then considers the subnet portion of the host ID to decide whether to route the packet to subnet 16 or subnet 32.

Subnet masks

For subnetting to work, the router must be told which portion of the host ID should be used for the subnet network ID. This little sleight of hand is accomplished by using another 32-bit number, known as a *subnet mask.* Those IP address bits that represent the network ID are represented by a 1 in the mask, and those bits that represent the host ID appear as a 0 in the mask. As a result, a subnet mask always has a consecutive string of ones on the left, followed by a string of zeros.

For example, the subnet mask for the subnet shown in Figure 3-3, where the network ID consists of the 16-bit network ID plus an additional 4-bit subnet ID, would look like this:

```
11111111 11111111 11110000 00000000
```

In other words, the first 20 bits are ones, and the remaining 12 bits are zeros. Thus, the complete network ID is 20 bits in length, and the actual host ID portion of the subnetted address is 12 bits in length.

To determine the network ID of an IP address, the router must have both the IP address and the subnet mask. The router then performs a bitwise operation called a *logical AND* on the IP address in order to extract the network ID. To perform a logical AND, each bit in the IP address is compared with the corresponding bit in the subnet mask. If both bits are 1, the resulting bit in the network ID is set to 1. If either of the bits are 0, the resulting bit is set to 0.

For example, here's how the network address is extracted from an IP address using the 20-bit subnet mask from the previous example:

```
IP address:  10010000 00011100 00010000 00010001 (144.28.16.17)
Subnet mask: 11111111 11111111 11110000 00000000
Network ID:  10010000 00011100 00010000 00000000 (144.28.16.0)
```

Thus, the network ID for this subnet is 144.28.16.0.

The subnet mask itself is usually represented in dotted-decimal notation. As a result, the 20-bit subnet mask used in the previous example would be represented as 255.255.240.0:

```
Subnet mask: 11111111 11111111 11110000 00000000
 255 . 255 . 240 .  0
```

Don't confuse a subnet mask with an IP address. A subnet mask doesn't represent any device or network on the internet. It's just a way of indicating which portion of an IP address should be used to determine the network ID.

Note that a subnet mask cannot be an arbitrary collection of octets. Instead, a subnet mask always has a certain number of binary 1s on its left side, and the remaining bits of the mask are always 0. This limits the dotted-decimal representation of a subnet mask to certain values.


TIP

You can spot a subnet mask right away because the first octet is always 255, and 255 is not a valid first octet for any class of IP address.

Network prefix notation

Because a subnet mask always begins with a consecutive sequence of ones to indicate which bits to use for the network ID, you can use a shorthand notation — a *network prefix* — to indicate how many bits of an IP address represent the network ID. The network prefix is indicated with a slash immediately after the IP address, followed by the number of network ID bits to use. For example, the IP address 144.28.16.17 with the subnet mask 255.255.240.0 can be represented as 144.28.16.17/20 because the subnet mask 255.255.240.0 has 20 network ID bits.

Network prefix notation is also called *classless interdomain routing* notation (*CIDR*, for short) because it provides a way of indicating which portion of an address is the network ID and which is the host ID without relying on standard address classes.

Default subnets

The *default subnet masks* are three subnet masks that correspond to the standard Class A, B, and C address assignments. These default masks are summarized in Table 3-4.

TABLE 3-4 The Default Subnet Masks

Class	Binary	Dotted-Decimal	Network Prefix
A	11111111 00000000 00000000 00000000	255.0.0.0	/8
B	11111111 11111111 00000000 00000000	255.255.0.0	/16
C	11111111 11111111 11111111 00000000	255.255.255.0	/24


TIP

Keep in mind that a subnet mask is not actually required to use one of these defaults because the IP address class can be determined by examining the first three bits of the IP address. If the first bit is 0, the address is Class A, and the subnet mask 255.0.0 is applied. If the first two bits are 10, the address is Class B,

and 255.255.0.0 is used. If the first three bits are 110, the Class C default mask 255.255.255.0 is used.

The great subnet roundup

You should know about a few additional restrictions that are placed on subnets and subnet masks. In particular

- **The minimum number of network ID bits is eight.** As a result, the first octet of a subnet mask is always 255.
- **The maximum number of network ID bits is 30.** You have to leave at least two bits for the host ID portion of the address to allow for at least two hosts. If you use all 32 bits for the network ID, that leaves no bits for the host ID. Obviously, that won't work. Leaving just one bit for the host ID won't work, either, because a host ID of all ones is reserved for a broadcast address, and all zeros refers to the network itself. Thus, if you use 31 bits for the network ID and leave only 1 for the host ID, host ID 1 would be used for the broadcast address, and host ID 0 would be the network itself, leaving no room for actual hosts. That's why the maximum network ID size is 30 bits.
- **Because the network ID portion of a subnet mask is always composed of consecutive bits set to 1, only eight values are possible for each octet of a subnet mask:** 0, 128, 192, 224, 248, 252, 254, and 255.
- **A subnet address can't be all zeros or all ones.** Thus, the number of unique subnet addresses is two less than two raised to the number of subnet address bits. For example, with three subnet address bits, six unique subnet addresses are possible ($2^3 - 2 = 6$). This implies that you must have at least two subnet bits. (If a single-bit subnet mask were allowed, it would violate the "can't be all zeros or all ones" rule because the only two allowed values would be 0 or 1.)

SUBNETS VERSUS VLANS

All of this talk of subnets might have you wondering: What's the difference between subnets and virtual local area networks (VLANs)? If you've read Book 1, Chapter 2, you know that VLANs are a divide-and-conquer technique for managing large networks. Subnetting is also a divide-and-conquer technique.

So, are they the same thing, and do they serve the same purpose?

The answer is: No, but sort of kind of. But really, no.

Although VLANs and subnets seem similar, VLANs are a layer 2 construct, and subnets are a layer 3 construct.

In other words, VLANs have nothing to do with IP addresses and subnets have nothing to do with MAC addresses.

That being said, it is very common — and usually desirable — to design your network with a one-to-one correspondence between VLANs and IP subnets. This usually simplifies the task of managing both.

As an example, suppose you want to divide a single-office network with just a hundred or fewer users into three groups: End-user devices like computers and printers, servers and network devices, and Voice over Internet Protocol (VoIP) phones. You could use three VLANs to do this — call them VLAN 10, VLAN 20, and VLAN 30.

You could then use three subnets — 192.168.10.*x*, 192.168.20.*x*, and 192.168.30.*x*.

There's a natural correspondence between these three VLANs and the three subnets, and network setup and management will be easier because the VLANs and subnets correspond to one another.

Although you can have a single VLAN that supports multiple subnets, in most networks there is a one-to-one correspondence that allows the benefits of VLANs and subnets to complement one another.

IP block parties

A subnet can be thought of as a range or block of IP addresses that have a common network ID. For example, the CIDR 192.168.1.0/28 represents the following block of 14 IP addresses:

```
192.168.1.1   192.168.1.2   192.168.1.3   192.168.1.4
192.168.1.5   192.168.1.6   192.168.1.7   192.168.1.8
192.168.1.9   192.168.1.10  192.168.1.11  192.168.1.12
192.168.1.13  192.168.1.14
```

Given an IP address in CIDR notation, it's useful to be able to determine the range of actual IP addresses that the CIDR represents. This matter is straightforward when the octet within which the network ID mask ends happens to be `0`, as in the preceding example. You just determine how many host IDs are allowed based on the size of the network ID and count them off.

However, what if the octet where the network ID mask ends is not 0? For example, what are the valid IP addresses for 192.168.1.100 when the subnet mask is 255.255.255.240? In that case, the calculation is a little harder. The first step is to determine the actual network ID. You can do that by converting both the IP address and the subnet mask to binary and then extracting the network ID as in this example:

```
IP address:  11000000 10101000 00000001 01100100 (192.168..100)
Subnet mask: 11111111 11111111 11111111 11110000
Network ID:  11000000 10101000 00000001 01100000 (192.168.1.96)
```

As a result, the network ID is 192.168.1.96.

Next, determine the number of allowable hosts in the subnet based on the network prefix. You can calculate this by subtracting the last octet of the subnet mask from 254. In this case, the number of allowable hosts is 14.

To determine the first IP address in the block, add 1 to the network ID. Thus, the first IP address in my example is 192.168.1.97. To determine the last IP address in the block, add the number of hosts to the network ID. In my example, the last IP address is 192.168.1.110. As a result, the 192.168.1.100 with subnet mask 255.255.255.240 designates the following block of IP addresses:

```
192.168.1.97  192.168.1.98  192.168.1.99  192.168.1.100
192.168.1.101 192.168.1.102 192.168.1.10  192.168.1.104
192.168.1.105 192.168.1.106 192.168.1.107 192.168.1.108
192.168.1.109 192.168.1.110
```

Private and public addresses

Any host with a direct connection to the internet must have a globally unique IP address. However, not all hosts are connected directly to the internet. Some are on networks that aren't connected to the internet. Some hosts are hidden behind firewalls, so their internet connection is indirect.

Several blocks of IP addresses are set aside just for this purpose, for use on private networks that are not connected to the internet or to use on networks that are hidden behind a firewall. Three such ranges of addresses exist, summarized in Table 3-5. Whenever you create a private TCP/IP network, you should use IP addresses from one of these ranges.

TABLE 3-5 Private Address Spaces

CIDR	Subnet Mask	Address Range
10.0.0.0/8	255.0.0.0	10.0.0.1–10.255.255.254
172.16.0.0/12	255.240.0.0	172.16.1.1–172.31.255.254
192.168.0.0/16	255.255.0.0	192.168.0.1–192.168.255.254

Pondering Ports

When you use an IP address, you often associate that IP address with a *port*, which enables a connection to a particular service. The best-known port is port 80, which corresponds to the HTTP of the World Wide Web. The combination of a transport protocol (for example, TCP), an IP address, and a port is called an *internet socket.*

Although IP addresses are defined at layer 3 of the OSI model (the network layer), ports are a layer 4 construct. Layer 4 is the transport layer, so it makes sense that ports would live there.

Ports are commonly combined with IP addresses when used in URLs (also known as web addresses.) I dive deep into URLs in Book 2, Chapter 6, so hold tight.

Ports are represented by 32-bit numbers, so they range from 0 to 65535. There are three ranges of port numbers:

- **0 to 1023:** These are called *well-known ports,* and they're used for the widely used services available on the internet.
- **1024 to 49151:** These are called *registered ports,* and they're assigned by the internet's governing authorities to various service providers. For example, Apple's iTunes uses port 3689 and Adobe's Media Server uses port 8134.
- **49152 to 65535:** These are called *dynamic ports, private ports,* or *ephemeral ports.* These ports cannot be registered and are used only for a specific communication.

Understanding Network Address Translation

Nearly all firewalls use a technique called *network address translation* (NAT) to hide the actual IP address of a computer on the local network from the outside world. When that's the case, the NAT device must use a globally unique IP address to

represent the computer to the internet. Behind the firewall, though, the computer has a private IP address. When packets cross the firewall, the NAT device translates the private IP address to the public IP address and vice versa.

NAT is one of the foundational techniques that enables the internet to work. It's the way an organization can have dozens, hundreds, or thousands of computers on its network without requiring a separate public IP address for each computer. Instead, each organization has a relatively small number of public IP addresses that are assigned to the public-facing interfaces of its firewall(s). NAT enables all the computers behind the firewall to communicate with the internet, piggybacking on the public IP address of the firewall itself.

Consider what typically happens when a user sends a request to a local HTTP server — that is, an HTTP server that is on the same network as the user. Let's assume that the IP address of the local HTTP server is 192.168.0.100, and the IP address of the user's computer is 192.168.0.50. What happens is this:

1. The user's computer sends an HTTP request in the form of an IP packet with the following address information:
   - For the source, the transport protocol is TCP, the IP address is 192.168.0.50. The port number for the source is chosen by client and is typically a high port number. For this example, I'll use port 45444 for the source port.
   - For the destination, the transport protocol is TCP, the IP address is 192.168.0.100, and the port is 80.
2. The HTTP server receives the request, processes it, and sends back an HTTP response in the form of an IP packet with the following address information:
   - For the source, the transport protocol is TCP, the IP address is 192.168.0.100, and the port is 80.
   - For the destination, the transport protocol is TCP, the IP address is 192.168.0.50, and the port is 45444 (the port that was chosen by the client).

This won't work if the user wants to send a request to an HTTP server on the internet, because the IP address of the user's computer is a private address, not a public address. So, the HTTP server won't be able to send a response to 192.168.0.50 because such an address doesn't exist on the public internet.

That's where NAT comes in.

The magic of NAT is handled by the firewall itself. The basic idea of NAT is that the firewall maintains an internal table of outgoing packets so it can remember which computer in the local network has requested information from sites on the

public internet. Because more than one computer may make requests for information from the same internet site, NAT exploits ephemeral ports (see the preceding section) to keep things straight.

Let me walk you through an example. But first, let's assume that the firewall in this example has the following IP addresses:

- **Outside IP address (public):** 75.68.10.201
- **Inside IP address (private):** 192.168.0.1

Let's also assume that the HTTP server is at 99.84.206.125 (which happens to be Wiley's web server), and a user whose private IP address is 192.168.0.50 uses a web browser to request information from the HTTP server. The HTTP request will have the following address information:

- **Source IP:** 192.168.0.50
- **Source port:** 45444
- **Destination IP:** 99.84.206.125
- **Destination port:** 80

Here's how it works:

1. The firewall sees this packet and realizes that it must substitute its own IP address (let's assume 192.168.0.1).
2. The firewall selects a random port number from a pool of ephemeral port numbers, which it will use to keep track of the request.

   For example, let's say it picks port 42003.
3. The firewall records the following information in its NAT table for this request:
   - Source IP: 192.168.0.50
   - Source port: 45444
   - Destination IP: 99.84.206.125
   - Destination port: 80
   - Temporary port: 42003
4. The firewall modifies the packet by substituting its own public IP address for the source IP and the temporary port for the source port.

5. The firewall sends the modified packet to the public internet.

   The modified packet contains the following information:

   - Source IP: 75.68.10.201
   - Source port: 42003
   - Destination IP: 99.84.206.125
   - Destination port: 80

6. A few seconds later, the firewall receives an incoming HTTP response message with the following address information:

   - Source IP: 99.84.206.125
   - Source port: 80
   - Destination IP: 75.68.10.201
   - Destination port: 42003

7. The firewall peruses its NAT table and finds that this response matches the entry it recorded in Step 3.

8. The firewall retrieves the original source IP address and port from the NAT table and substitutes it for the destination IP and port.

   The modified response message now has the following address information:

   - Source IP: 99.84.206.125
   - Source port: 80
   - Destination IP: 192.168.0.50
   - Destination port: 45444

9. The firewall places the modified packet on the inside interface (that is, the local network), where the network can then deliver the packet to the original requestor (the user at 192.168.0.50).

So, what happens if two or more users have requests to the same web server at the same time? NAT is able to figure it out because each of those requests has a different temporary port number. For example, the user in the preceding example got port 42003. Another user sending a request to the same web server might get port 43859. When the reply comes back from the web server, NAT looks at the destination port to determine which local computer should receive the reply.

IN THIS CHAPTER

- Examining the basic operation of a router
- Considering different types of routers
- Looking at the routing table
- Examining how packets flow through a router

Chapter 4
Routing

In the simplest terms, a *router* is a device that works at layer 3 of the OSI Reference Model. Layer 3 is the network layer, which means that layer 3 is responsible for exchanging information between distinct networks. And that's exactly the function of a router: connecting two or more networks so that packets can flow freely between them.

Routing is the general term that describes what routers do. In short, when a router receives a packet from one network that is destined for a device on another network, the router determines the best way to get the packet to its destination.

In a way, routers are the post offices of the internet. When you drop a letter into a public mailbox, a mail carrier collects the mail and delivers it to a nearby post office. There, the mail is sorted and sent off to a regional post office, where the mail is sorted again and maybe sent off to yet another regional post office, and so on until the mail finally arrives at a post office close to the delivery address, where the mail is sorted one last time and given to a mail carrier who delivers the letter to the correct address.

At each step of the way, the best route to move the mail closer to its destination is determined, and the mail is sent along its way. Routers work pretty much just like that.

In this chapter, you learn about three of the most important uses for routers: connecting to the internet, connecting remote offices to each other via the internet, and managing traffic in extremely large networks such as large campuses, huge office buildings, or the internet itself.

Considering the Usefulness of Routers

Routers are an indispensable part of any network. In fact, virtually all networks require at least one router. The following sections describe the most common reasons for introducing a router into your network.

Connecting to the internet

A router is required for any network that needs access to the internet. Such a device is known as an *internet gateway* because it serves as a "gateway" to the internet.

Strictly speaking, the internet isn't a network — it's an enormous collection of millions of networks, all tied together via millions of individual routers. Your internet service provider (ISP) is just one of those millions of networks, and your internet gateway connects your private network to your ISP's network. Your ISP, in turn, provides the routers that connect the ISP's network to other networks, which ultimately connects to the internet backbone and the rest of the world.

For a home network or a very small business network, you can use an inexpensive *residential gateway* device that you can buy at a consumer electronics store such as Best Buy. A residential gateway typically includes five distinct components, all bundled into one neat package:

- A router, used to connect your private network to your ISP's network
- A small switch, typically providing from three to eight ports to connect wired devices such as computers and printers
- A wireless access point (WAP) to connect wireless devices such as laptops or smartphones
- A firewall to provide protection from intruders seeking to compromise your network
- A DHCP server to provide IP addresses for the computers and other devices on your network

Figure 4-1 shows an example of this type of setup. In this example, an ISP provides a cable feed into your house that connects to a *cable modem*, which provides a single Ethernet port to which you can connect a residential gateway. Computers on the home network are connected to the gateway's switch ports, and wireless devices connect via the Wi-Fi network. The gateway's DHCP server hands out IP addresses to any devices connected to the network.

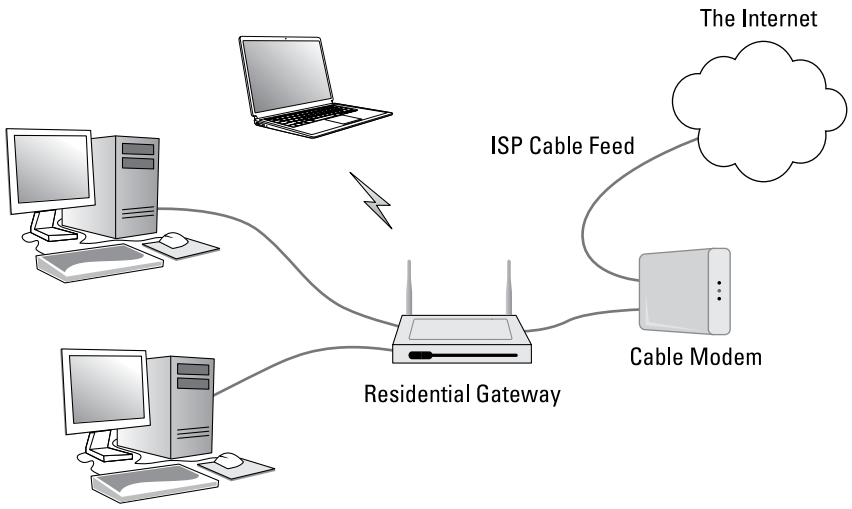


FIGURE 4-1: Connecting to the internet via a residential gateway.

Figure 4-2 shows the type of internet gateway typically used in a larger network. Here, the ISP delivers a high-speed fiber-optic feed to the customer's location and provides an *Ethernet handoff*, which is simply one or more Ethernet ports that the customer can connect to.

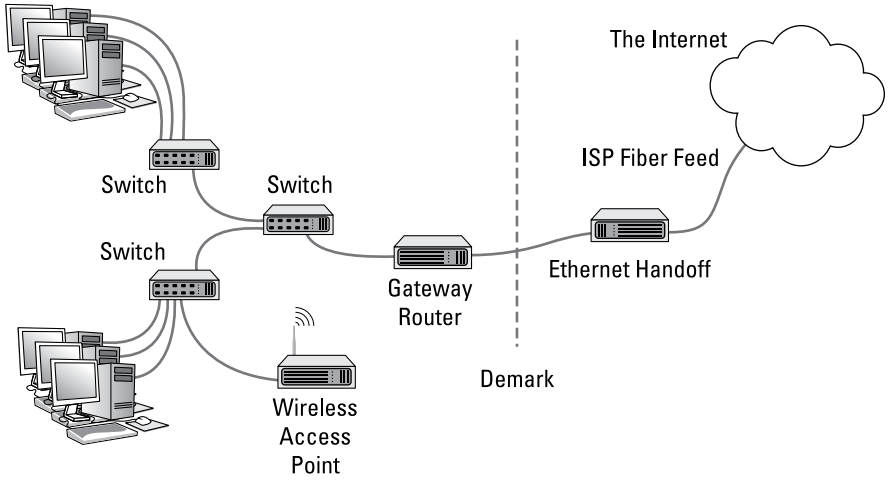


FIGURE 4-2: Connecting a larger business network to the internet.

The Ethernet handoff establishes what is called the *demarcation point*, usually called simply the *demark*. The demark is simply the dividing line that establishes who is responsible for what: The ISP is responsible for everything between the internet and the demark; the customer is responsible for everything on the private network side of the demark.

A *gateway router* is used to connect the private network to the Ethernet handoff. Much like a residential gateway, a gateway router typically provides several features into one combined device, including a router, a small switch, and a firewall. However, most business-class gateway routers don't provide Wi-Fi — the wireless network is provided by dedicated WAPs. And the small switch provided by the gateway doesn't serve the entire network; instead, it is connected to a network of switches that in turn connect the network's computers together.

One of the distinguishing characteristics of a gateway router is that it has a small number of network interfaces. At the minimum, a gateway router needs just two interfaces: an external interface and an internal interface. The *external interface,* often labeled *WAN* on the device, connects to the ISP's feed. The *internal interface* connects to the private network. (If the device includes a switch, it will have more than one internal interface.)

For more information about using a router to connect to the internet, refer to Book 3, Chapter 2.

Connecting remote locations

Another common use for routers is to connect geographically separated offices to form a single network that spans multiple locations. You can do this by using a pair of gateway routers to create a secure virtual private network (VPN) between the two networks. Each network uses its gateway router to connect to the internet, and the routers establish a secure tunnel between themselves to exchange private information.

Figure 4-3 shows how a VPN can be used to establish a site-to-site tunnel between offices in Los Angeles and Las Vegas. As you can see, each site has its own gateway router that connects to the internet. The routers are configured to provide a VPN that securely connects the two networks.

Note that the need for a VPN tunnel between networks is not related to the size of the network, at least not in terms of how many users are on the network. Instead, it's simply a function of geography. A small-town company that has two three-person offices on opposite sides of the same street can benefit from a VPN tunnel just as much as a company that has a 200-person office in Dallas and a second 200-person office in Houston.

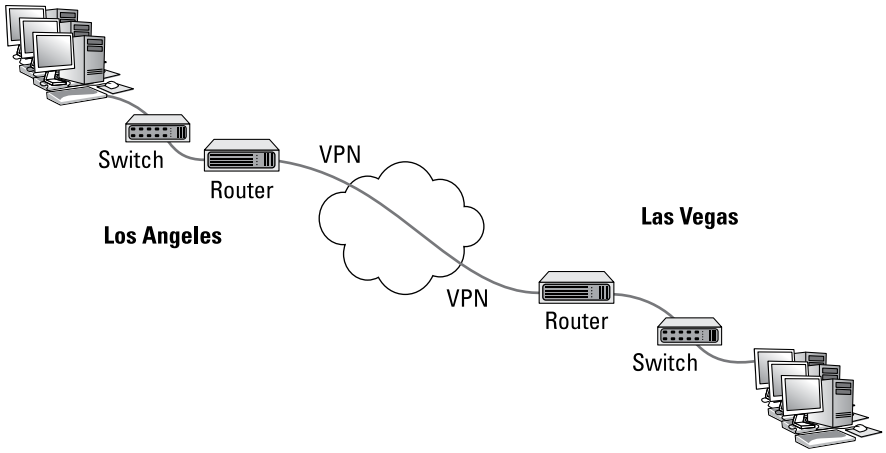


FIGURE 4-3: Connecting two networks via VPN.

Splitting up large networks

Large networks often have need for routers that are internal to the network itself. For example, consider a company that employs several thousand employees on a single campus that consists of several dozen buildings. For a network like this, routers are used to manage the network by dividing it into smaller, more manageable networks all connected with routers.

Figure 4-4 shows a simplified version of how this works. Here, a large network is segmented into two smaller networks, each on a different subnet: one on 10.0.100.*x* (subnet mask 255.255.255.0), the other on 10.0.200.*x* (also 255.255.255.0). A router is used to provide a link between the two subnets, so packets can flow from one subnet to the other.

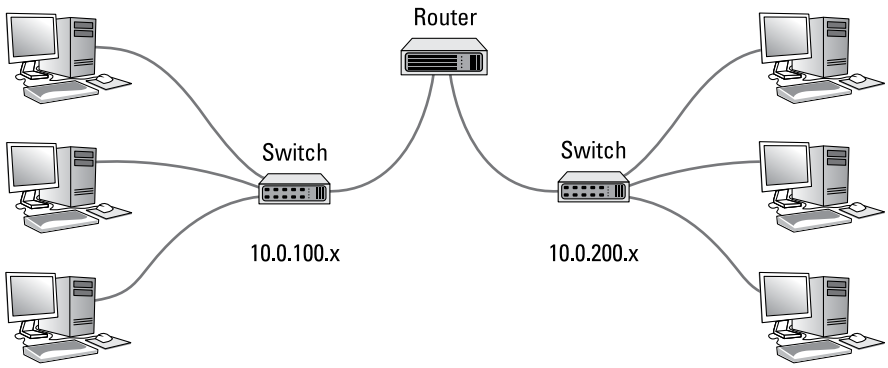


FIGURE 4-4: Using an internal router.

A router used in this way is called an *internal router* because it doesn't connect a private network to a public network. Instead, it connects two portions of a larger network.

Separate internal routers are becoming less commonplace because most switches now have routing functions built in to allow subnets to communicate with one another.

Still, very large networks still require routers to handle the large amount of traffic that must flow between networks. Picture a large college campus, in which each department has its own network each with dozens or possibly hundreds of devices. Each of those networks would have a relatively small departmental router, but a larger internal router would be used in the data center to connect all the departmental routers.

Understanding Routing Tables

Routers work by maintaining an internal list of networks that can be reached via each of the router's interfaces. This list is called a *routing table.* When a packet arrives on one of the router's interfaces, the router examines the destination IP address of the incoming packet, consults the routing table to determine which of its interfaces it should forward the packet to, and then forwards the packet to the correct interface.

Sounds simple enough.

The trick is building the routing table. In simple cases, such as a gateway router that connects a private network to the internet, the routing table is created manually with *static routes.* Configuring a gateway router with static routes isn't much more complicated than configuring a host computer with a static IP address. All you need to know is the IP address, subnet mask, and gateway address provided by your ISP for the external interface and the network address and subnet mask of the private network for the internal interface.

For more complicated environments, where multiple routers are used on the private network, special *routing protocols* are used to build *dynamic routes.* These routing protocols are designed to discover the topology of the network by finding out which routers are present on the network and which networks each router can reach.

Refer back to Figure 4-2, which depicts a small business network that has a fiber-optic connection to the internet provided by an ISP and a gateway router that connects to the Ethernet handoff.

Let's assume that the private network for this business operates on a single subnet, and the IP address for the network is 10.0.1.0 with the subnet mask 255.255.255.0. The six computers in the private network have IP addresses 10.0.101.1 through 10.0.101.6. And the internal interface on the gateway will be configured with the IP address 10.0.1.254.

Now let's assume that the ISP provides you with the following information for your Ethernet handoff:

IP address: 205.186.181.97

Subnet mask: 255.255.255.255

Default gateway: 107.0.65.31

In this example, the gateway router would have the following entries in its routing table:

Entry	Destination Network IP	Subnet Mask	Gateway	Interface
1	10.0.1.0	255.255.255.0	10.0.1.254	Internal
2	205.186.181.97	255.255.255.255	205.186.181.97	External
3	107.0.65.31	255.255.255.255	107.0.65.31	External
4	0.0.0.0	0.0.0.0	107.0.65.31	External

First, let's have a look at each of the columns in the routing table:

- **Entry:** The entry number.
- **Destination network IP:** This is the IP address of the destination network. This column is used in conjunction with the subnet mask column to determine the network to which that packet's destination IP address belongs.
- **Subnet mask:** The subnet mask that is applied to the destination IP address to determine the destination network.
- **Gateway:** The address of the router that the packet should be forwarded to.

- **Interface:** The interface that the packet should be forwarded through. Here, *internal* means the interface to which the internal private network is connected and *external* means the interface on which the ISP's handoff is connected. (In many gateway devices, these interfaces are labeled LAN and WAN, respectively.)

Now let's have a look at the four entries in this routing table:

- **Entry 1:** This entry tells the router what to do with packets whose destination is on the internal network (10.0.1.*x*). The IP address of the internal network is 10.0.1.0 and the subnet mask is 255.255.255.0. These packets will be sent to the internal interface, whose IP address is 10.0.1.254.
- **Entry 2:** This entry handles packets whose destination is the ISP's gateway (107.0.65.31). The 255.255.255.255 subnet mask means that the destination is a specific IP address, not a network. These packets are forwarded to the ISP's gateway on the external interface.
- **Entry 3:** This entry handles packets whose destination is the gateway's external interface, which has been assigned the IP address 205.186.181.97 by the ISP. These packets are forwarded to the gateway address on the external interface.
- **Entry 4:** This entry handles everything else. The network IP address 0.0.0.0 with no subnet mask means that all packets that aren't caught by any of the other rules are forwarded out to the ISP's gateway router (107.0.65.31) on the external interface.

The entries in the routing table are evaluated against each packet's destination IP address to determine where the packet should be sent. The entries are evaluated in order, and the first one that matches is used to send the packet along its way.

For example, suppose a packet is received on the external interface and the destination address is 10.0.1.5. The router will first consider entry 1, applying the subnet mask 255.255.255.0 to consider the 10.0.1.0. Because this matches the network ID in entry 1, the packet will be forwarded over the internal interface, where the switch can hand the packet off to the correct computer.

On the other hand, suppose a packet is received on the internal interface and the destination IP address is 108.211.23.42. When the router tries the first entry, the subnet mask extracts the network address 108.211.23.42. This doesn't match 10.0.1.0, so the router considers the second entry. The subnet mask 255.255.255.255

tells the router to compare the entire destination address with the IP address 107.0.65.31. Because that address doesn't match, the router tries the third entry. Again, the subnet mask 255.255.255.255 tells the router to compare the entire destination address, this time with the IP address 205.186.181.97. Again, the addresses don't match, so the router moves to the fourth and final entry in the router table. The subnet mask 0.0.0.0 reduces the entire destination address to 0.0.0.0, which matches the destination network 0.0.0.0. Therefore, the router forwards the packet on to the ISP's router at 107.0.65.31 via the external interface.

It sounds pretty simple, but in reality there is a lot more going on under the hood. In more complicated networks, there are a lot more than just four entries in the routing table. And in a busy network, a router is likely handling hundreds or even thousands of packets per second. For example, if you have 100 users on your network, all of them browsing the web, accessing email, and using other applications that cross the router, the router has an enormous workload.

IN THIS CHAPTER

- Discovering the basics of DHCP
- Exploring scopes
- Configuring a DHCP server
- Setting up a DHCP client

Chapter 5
DHCP

Every host on a Transmission Control Protocol/Internet Protocol (TCP/IP) network must have a unique Internet Protocol (IP) address. Each host must be properly configured so that it knows its IP address. When a new host comes online, it must be assigned an IP address that's within the correct range of addresses for the subnet but not already in use. Although you can manually assign IP addresses to each computer on your network, that task quickly becomes overwhelming if the network has more than a few computers.

That's where Dynamic Host Configuration Protocol (DHCP) comes into play. DHCP automatically configures the IP address for every host on a network, thus assuring that each host has a valid, unique IP address. DHCP even automatically reconfigures IP addresses as hosts come and go. As you can imagine, DHCP can save a network administrator many hours of tedious configuration work.

In this chapter, you discover the ins and outs of DHCP: what it is, how it works, and how to set it up.

Understanding DHCP

DHCP allows individual computers on a TCP/IP network to obtain their configuration information — in particular, their IP address — from a server. The DHCP server keeps track of which IP addresses are already assigned so that when a

computer requests an IP address, the DHCP server offers it an IP address that's not already in use.

Configuration information provided by DHCP

Although the primary job of DHCP is to dole out IP addresses and subnet masks, DHCP actually provides more configuration information than just the IP address to its clients. The additional configuration information consists of DHCP options. The following is a list of some common DHCP options that can be configured by the server:

- The router address, also known as the Default Gateway address
- The expiration time for the configuration information
- Domain name
- Domain Name Server (DNS) server address
- Windows Internet Name Service (WINS) server address

DHCP servers

A DHCP server can be a server computer located on the TCP/IP network. All modern server operating systems have a built-in DHCP server. To set up DHCP on a network server, all you have to do is enable the server's DHCP function and configure its settings. In the upcoming section, "Working with a DHCP Server," I show you how to configure a DHCP server for Windows Server 2025. (The procedure for previous versions of Windows Server is similar.)

A server computer running DHCP doesn't have to be devoted entirely to DHCP unless the network is very large. For smaller networks, a file server can share duty as a DHCP server. This is especially true if you provide long leases for your IP addresses. (*Lease* is the term used by DHCP to indicate that an IP address has been temporarily given out to a particular computer or other device.)

Many multifunction routers also have built-in DHCP servers. If you don't want to burden one of your network servers with the DHCP function, you can enable the router's built-in DHCP server. An advantage of allowing the router to be your network's DHCP server is that you rarely need to power-down a router. In contrast, you occasionally need to restart or power-down a file server to perform system maintenance, apply upgrades, or perform troubleshooting.



TIP

Most networks require only one DHCP server. Setting up two or more servers on the same network requires that you carefully coordinate the IP address ranges (known as *scopes*) for which each server is responsible. If you accidentally set up two DHCP servers for the same scope, you may end up with duplicate address assignments if the servers attempt to assign the same IP address to two different hosts. To prevent this from happening, just set up one DHCP server unless your network is so large that one server can't handle the workload.

How DHCP actually works

You can configure and use DHCP without knowing the details of how DHCP client configuration actually works. However, a basic understanding of the process can help you to understand what DHCP is actually doing. Not only is this understanding enlightening, but it can also help when you're troubleshooting DHCP problems.

The following paragraphs contain a blow-by-blow account of how DHCP configures TCP/IP hosts. This procedure happens every time you boot up a host computer. It also happens when you release an IP lease and request a fresh lease.

1. **When a host computer starts up, the DHCP client software sends a special broadcast packet, known as a *DHCP Discover message.***

   This message uses the subnet's broadcast address (all host ID bits set to one) as the destination address and 0.0.0.0 as the source address.



TIP

   The client has to specify 0.0.0.0 as the source address because it doesn't yet have an IP address, and it specifies the broadcast address as the destination address because it doesn't know the address of any DHCP servers. In effect, the DHCP Discover message is saying, "Hey! I'm new here. Are there any DHCP servers out there?"

2. **The DHCP server receives the broadcast DHCP Discover message and responds by sending a *DHCP Offer* message.**

   The DHCP Offer message includes an IP address that the client can use.

   Like the DHCP Discover message, the DHCP Offer message is sent to the broadcast address. This makes sense because the client to which the message is being sent doesn't yet have an IP address and won't have one until it accepts the offer. In effect, the DHCP Offer message is saying, "Hello there, whoever you are. Here's an IP address you can use, if you want it. Let me know."

   What if the client never receives a DHCP Offer message from a DHCP server? In that case, the client waits for a few seconds and tries again. The client will try four times — at 2, 4, 8, and 16 seconds. If it still doesn't get an offer, it will try again after five minutes.

3. **The client receives the DHCP Offer message and sends back a message known as a *DHCP Request message.***

   At this point, the client doesn't actually own the IP address: It's simply indicating that it's ready to accept the IP address that was offered by the server. In effect, the DHCP Request message says, "Yes, that IP address would be good for me. Can I have it, please?"

4. **When the server receives the DHCP Request message, it marks the IP address as assigned to the client and broadcasts a *DHCP Ack message.***

   The DHCP Ack message says, in effect, "Okay, it's all yours. Here's the rest of the information you need to use it."

5. **When the client receives the DHCP Ack message, it configures its TCP/IP stack by using the address it accepted from the server.**

Understanding Scopes

A *scope* is simply a range of IP addresses that a DHCP server is configured to distribute. In the simplest case, where a single DHCP server oversees IP configuration for an entire subnet, the scope corresponds to the subnet. However, if you set up two DHCP servers for a subnet, you can configure each with a scope that allocates only one part of the complete subnet range. In addition, a single DHCP server can serve more than one scope, and a DHCP server can (and typically does) serve more than one subnet.

You must create a scope before you can enable a DHCP server. When you create a scope, you can provide it with the following properties:

- A **scope name,** which helps you to identify the scope and its purpose
- A **scope description,** which lets you provide additional details about the scope and its purpose
- A **starting IP address** for the scope
- An **ending IP address** for the scope
- A **subnet mask** for the scope

  You can specify the subnet mask with dotted-decimal notation or with network prefix notation.

- **One or more ranges of excluded addresses**

  These addresses won't be assigned to clients. For more information, see the section "Feeling excluded?" later in this chapter.

- **One or more reserved addresses**

  These are addresses that will always be assigned to particular host devices. For more information, see the section "Reservations suggested" later in this chapter.

- The **lease duration,** which indicates how long the host will be allowed to use the IP address

  The client will attempt to renew the lease when half of the lease duration has elapsed. For example, if you specify a lease duration of eight days, the client will attempt to renew the lease after four days. This allows the host plenty of time to renew the lease before the address is reassigned to some other host.

- The **router address** for the subnet

  This value is also known as the Default Gateway address.

- The **domain name** and the **IP address** of the network's DNS servers and WINS servers


TIP

Scopes, subnets, and VLANs

You might be wondering just how DHCP works in a network with several virtual local area networks (VLANs). Because each VLAN in a network is a separate broadcast domain, a DHCP request from one computer can't cross over to other VLANs on the network.

There are two basic ways to solve this problem. The first is to put a separate DHCP server on each VLAN. On very large networks, that's a sensible solution. But for most networks, there's an easier way, called *DHCP relay*, also known as *IP Helper.*

DHCP relay is a routing function that forwards DHCP traffic across VLANs. Most routers can provide for DHCP routing, and many switches can do it as well. (DHCP relay is a layer 3 function, so switches that provide this feature are considered to be layer 3 switches.)

To configure a router (or switch) for DHCP relay, you simply associate a VLAN with a DHCP router that's in a different VLAN. For example, suppose you have two VLANs — VLAN 20 on subnet 10.0.100.*x* and VLAN 30 on subnet 10.0.200.*x* — and your DHCP server is at 10.0.100.15 on VLAN 20. The router or switch would be configured to forward all DHCP traffic for VLAN 30 to 10.0.100.15. That way, both VLANs get DHCP from the same server.

In this case, you'd also need to ensure that the DHCP server has a scope for the subnets that correspond to the two VLANs. For example, you could set up a scope named VLAN 20 that serves IP addresses in the range 10.0.100.10 to 10.0.

Feeling excluded?

Everyone feels excluded once in awhile. But sometimes being excluded is a good thing. In the case of DHCP scopes, exclusions can help you to prevent IP address conflicts and can enable you to divide the DHCP workload for a single subnet among two or more DHCP servers.

An *exclusion* is a range of addresses that are not included in a scope. The exclusion range falls within the range of the scope's starting and ending addresses. In effect, an exclusion range lets you punch a hole in a scope. The IP addresses that fall within the hole won't be assigned.

Here are a few reasons for excluding IP addresses from a scope:

- **The computer that runs the DHCP service itself must usually have a static IP address assignment.** As a result, the address of the DHCP server should be listed as an exclusion.
- **Some hosts, such as a server or a printer, may need to have a predictable IP address.** In that case, the host will require a static IP address. By excluding its IP address from the scope, you can prevent that address from being assigned to any other host on the network.


TIP

Holding back some IP addresses at the bottom and top of a subnet is always a good idea. After all, the future is hard to predict. Even though you may not need the static IP space now, things change fast in our business. Here's a typical configuration for a subnet that allows for this breathing room:

Start Address	End Address	Description
10.0.100.1	10.0.100.254	Address range for distribution
10.0.100.1	10.0.100.19	Excluded from distribution
10.0.100.220	10.0.100.254	Excluded from distribution

Here, the two exclusion ranges mean that the scope will distribute addresses from 10.0.101.20 to 10.0.101.219.

You could achieve the same thing without the exclusions — just list 10.0.101.20 as the start of the scope and 10.0.101.219 as the end of the scope. But it's a common

practice to start by specifying the entire subnet as the address range for the scope, and then exclude parts of the subnet as needed.

Reservations suggested

In some cases, you may want to assign a particular IP address to a particular host. One way to do this is to configure the host with a static IP address so that the host doesn't use DHCP to obtain its IP configuration. However, here are two major disadvantages to that approach:

- **TCP/IP configuration supplies more than just the IP address.** If you use static configuration, you must manually specify the subnet mask, the Default Gateway address, the DNS server address, and other configuration information required by the host. If this information changes, you have to change it not only at the DHCP server, but also at each host that you configured statically.
- **You must remember to exclude the static IP address from the DHCP server's scope.** Otherwise, the DHCP server won't know about the static address and may assign it to another host. Then, you'll have two hosts with the same address on your network.

A better way to assign a fixed IP address to a particular host is to create a DHCP reservation. A *reservation* simply indicates that whenever a particular host requests an IP address from the DHCP server, the server should provide it the address that you specify in the reservation. The host won't receive the IP address until the host requests it from the DHCP server, but whenever the host does request IP configuration, it will always receive the same address.


TIP

To create a reservation, you associate the IP address that you want assigned to the host with the host's Media Access Control (MAC) address. As a result, you need to get the MAC address from the host before you create the reservation. You can get the MAC address by running the command `ipconfig /all` from a command prompt.


REMEMBER

If you set up more than one DHCP server, each should be configured to serve a different range of IP addresses. Otherwise, the servers might assign the same address to two different hosts.

How long to lease?

One of the most important decisions that you'll make when you configure a DHCP server is the length of time to specify for the lease duration. The default value

is eight days, which is appropriate in many cases. However, you may encounter situations in which a longer or shorter interval may be appropriate:

- **The more stable your network,** the longer the lease duration can safely exist. If you only periodically add new computers to the network or replace existing computers, you can safely increase the lease duration past eight days.
- **The more volatile the network,** the shorter the lease duration should be. For example, a wireless network in a university library is used by students who bring their laptop computers into the library to work for a few hours at a time. For this network, a duration such as one hour may be appropriate.



WARNING

Don't configure your network to allow infinite duration leases. Some administrators feel that this cuts down the workload for the DHCP server on stable networks. However, no network is permanently stable. Whenever you find a DHCP server that's configured with infinite leases, look at the active leases. I guarantee you'll find IP leases assigned to computers that no longer exist.

Working with a DHCP Server

Usually, the best way to understand abstract concepts is to see how they work in the real world. To that end, the next few sections show you a brief overview of how DHCP is managed in a Windows network. First, you see how a DHCP server is installed in Windows Server 2025. Then you see how a DHCP server is configured.

Installing a Windows Server 2025 DHCP server

To install the DHCP server role on Windows Server 2025, follow these steps:

1. **Click Server Manager in the Start menu.**

   The Server Manager application appears.

2. **From the menu near the upper-right, choose Manage ➪ Add Roles & Features.**

   The Before You Begin screen of the Add Roles and Features Wizard appears.

3. **Click Next.**

   The Installation Type screen appears.

4. **Choose Role-Based or Feature-Based Installation and then click Next.**

   The wizard displays a list of available servers.

5. **Select the server on which you want to install the DHCP role on; then click Next.**

   The wizard displays a list of available server roles.

6. **Select DHCP Server from the list of roles and then click Next.**

   The wizard displays a list of required features that must also be installed to support DHCP.

7. **Click Add Features, and then click Next.**

   The wizard displays a list of features you can optionally install.

8. **Click Next.**

   The wizard displays a screen describing what the DHCP role entails.

9. **Click Next.**

   The wizard displays a confirmation page.

10. **Click Install.**

    The wizard installs the DHCP role, which may take a few minutes. When the installation completes, a results screen is displayed to summarize the results of the installation.

11. **Click Close.**

    You're done!

Configuring a new scope

After you install the DHCP role on Windows Server 2025, you'll need to create at least one scope so the server can start handing out IP addresses. Here are the steps:

1. **In Server Manager, choose Tools ➪ DHCP**

   This brings up the DHCP management console, shown in Figure 5-1.

2. **Select the DHCP server you want to define the scope for, click IPv4, and then click the New Scope button on the toolbar.**

   This brings up the New Scope Wizard dialog box, as shown in Figure 5-2.

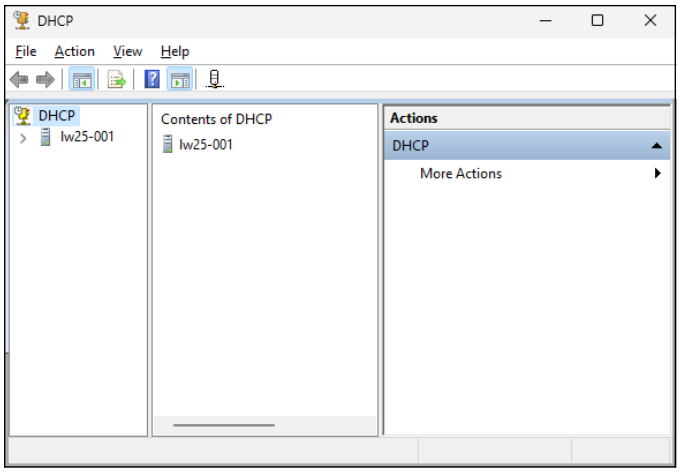


FIGURE 5-1: The DHCP management console.

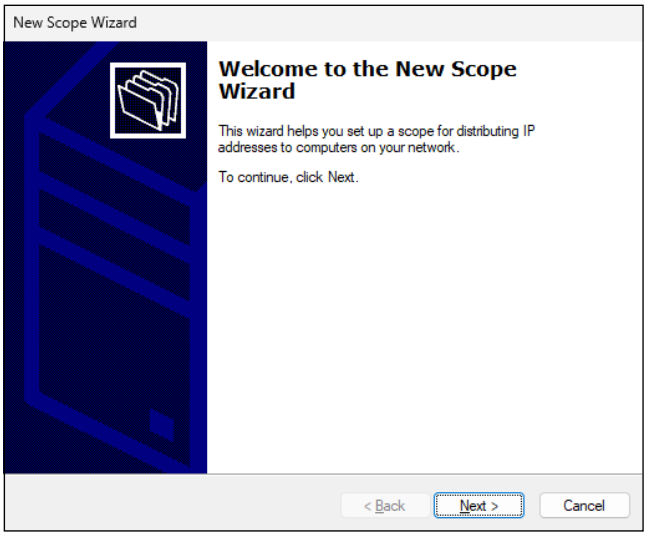


FIGURE 5-2: The New Scope Wizard comes to life.

3. **Click Next.**

   You're prompted for the name of the scope, as shown in Figure 5-3. I suggest you pick a descriptive name, such as "Main Office Computers."

4. **Enter a name and optional description, and then click Next.**

   The wizard asks for information required to create the scope, as shown in Figure 5-4.

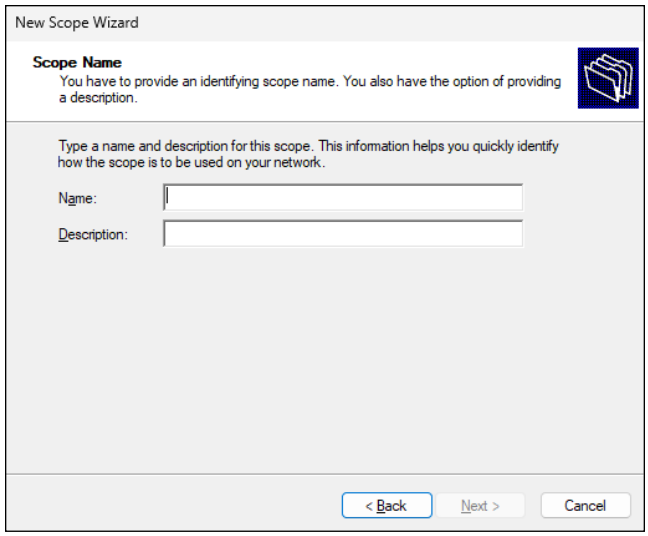


FIGURE 5-3: The wizard asks for a name for the new scope.

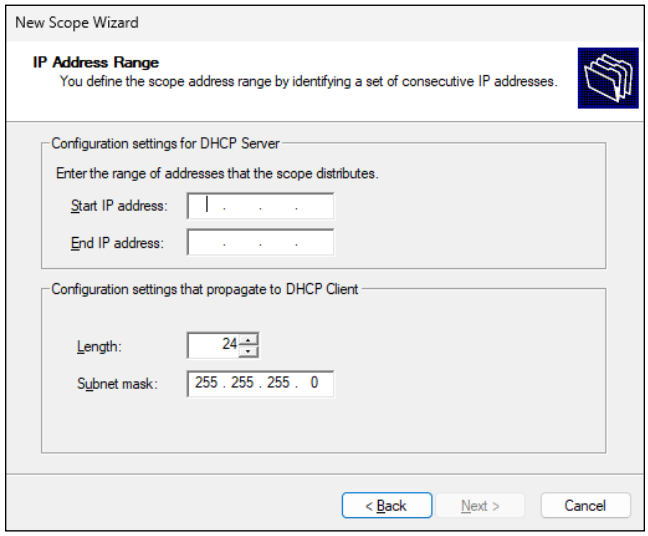


FIGURE 5-4: The wizard asks for scope information.

5. Enter the information for the new scope.

You must enter the following information:

- *Start IP Address:* This is the lowest IP address that will be issued for this scope.
- *End IP Address:* This is the highest IP address that will be issued for this scope.
- *Subnet Mask:* This is the subnet mask issued for IP addresses in this scope.

6. **Click Next.**

   The wizard asks whether you want to exclude any ranges from the scope range, as shown in Figure 5-5.

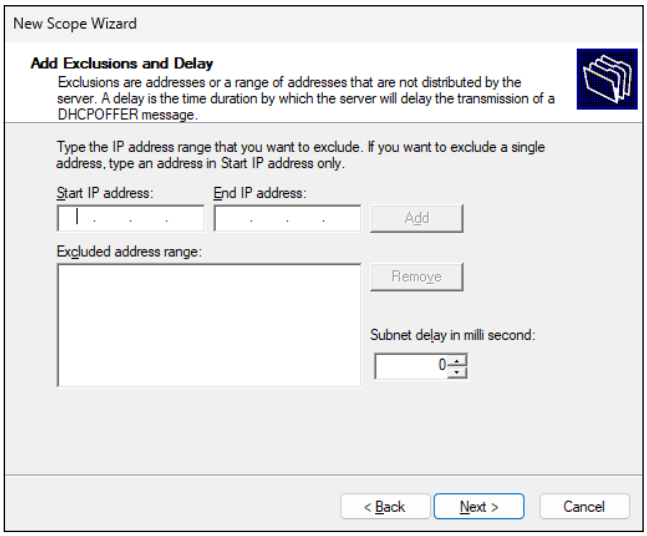


FIGURE 5-5: Do you want to create exclusions?

7. **(Optional) To create an exclusion, enter the IP address range to exclude and then click Add.**

   You can repeat this step as many times as necessary to add any excluded addresses.

8. **Click Next.**

   The wizard asks for the lease duration, as shown in Figure 5-6. The default is set to eight days. You can shorten or lengthen this value if you wish, but most people leave it set to the default.

9. **(Optional) Change the lease duration; then click Next.**

   When the wizard asks whether you want to configure additional DHCP options, leave this option set to Yes to complete your DHCP configuration now.

10. **Click Next.**

    The wizard asks if you'd like to change DHCP options such as the default gateway and DNS servers.

11. **Select Yes; then click Next.**

    The wizard asks for the default gateway information, as shown in Figure 5-7.

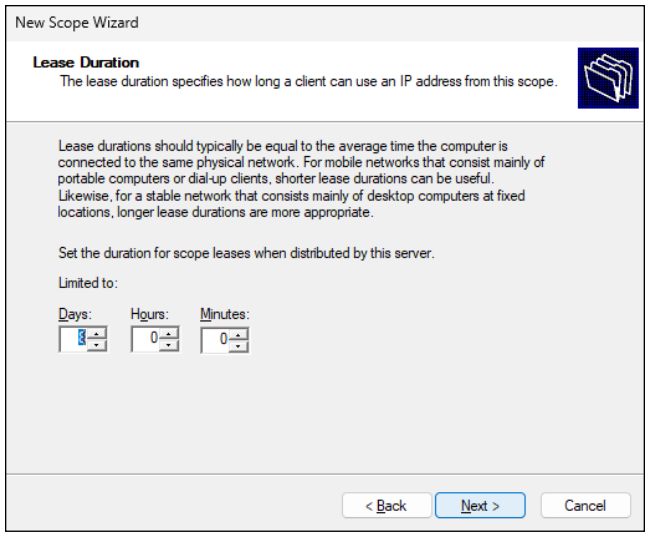


FIGURE 5-6: Set the lease duration.

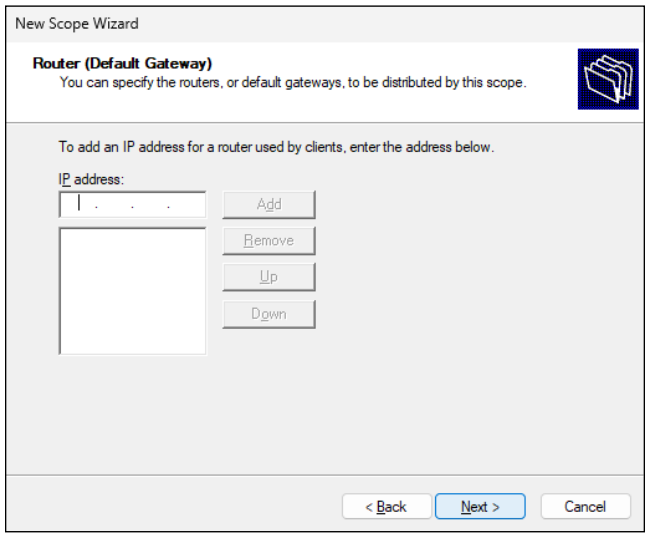


FIGURE 5-7: Provide the Default Gateway address.

12. Enter the address of your network's gateway and click Add; then click Next.

The wizard now asks for additional DNS information, as shown in Figure 5-8.

13. (Optional) If you want to add a DNS server, enter its address and then click Add.

Repeat this step as many times as necessary to add any additional DNS servers.

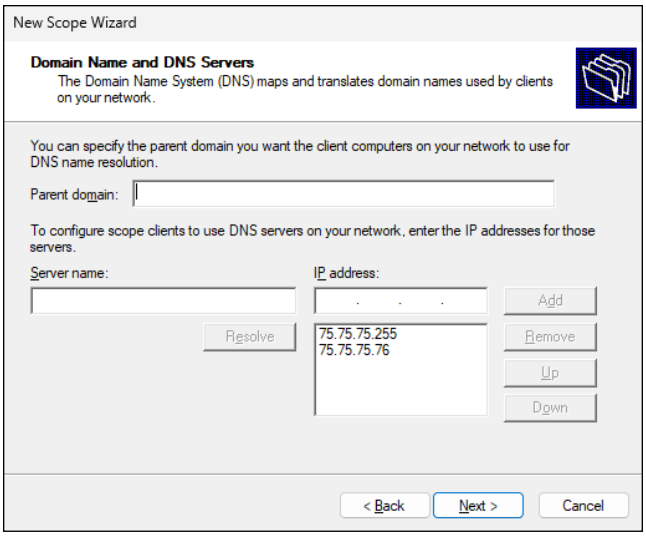


FIGURE 5-8: Provide additional DNS information.

14. **Click Next.**

    The wizard next asks for WINS configuration information.

15. **(Optional) If you want to enable WINS, enter the WINS server configuration.**

    WINS isn't required for most modern networks, so you can usually just leave this screen blank.

16. **Click Next.**

    The wizard now asks whether you want to activate the scope.

17. **Select Yes, I Want to Activate This Scope and then click Next.**

    A final confirmation screen is displayed.

18. **Click Finish.**

    The scope is created and you're returned to the DHCP Management Console.

You can confirm that the scope was set up properly by navigating through the DHCP Management Console to the scope you just created and selecting Address Pool. The IP distribution range and any exclusions will be displayed, as shown in Figure 5-9.

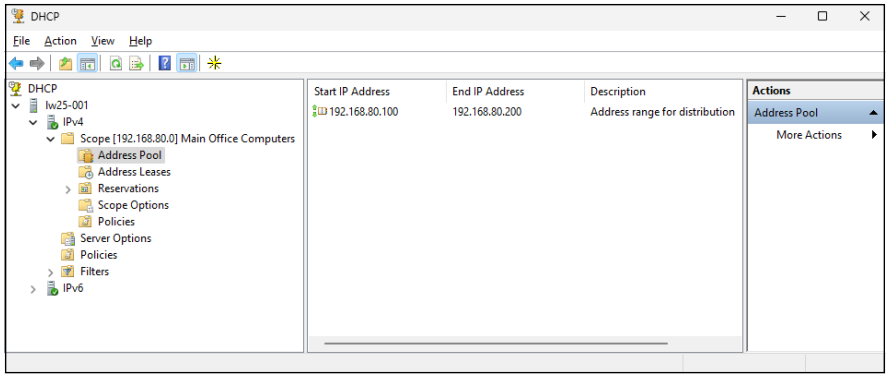


FIGURE 5-9: Viewing the address pool of a DHCP scope.

How to Configure a Windows DHCP Client

Configuring a Windows client for DHCP is easy. The DHCP client is automatically included when you install the TCP/IP protocol, so all you have to do is configure TCP/IP to use DHCP. And in nearly all cases, DHCP is configured automatically when you install Windows.

If you must configure DHCP manually, follow these steps:

1. **Open the Control Panel.**

   If you haven't already, switch to Small Icons view.

2. **Open Network and Sharing Center.**

3. **Click the link for your wired or wireless network adapter.**

   The adapter's Status dialog box shows useful information about the adapter, as shown in Figure 5-10.

   If you just want to find out your IP address, click the Details button.

4. **Click Properties.**

   This brings up the adapter's Properties dialog box, shown in Figure 5-11.

5. **Select Internet Protocol Version 4; then click Properties.**

   The Internet Protocol Version 4 (TCP/IPv4) Properties dialog box, shown in Figure 5-12, appears.

6. **Select Obtain an IP Address Automatically and Obtain DNS Server Address Automatically.**

7. **Click OK to apply the changes and dismiss the Internet Protocol Version 4 (TCP/IPv4) Properties dialog box.**

8. **Keep clicking OK to close all the other dialog boxes you've opened.**

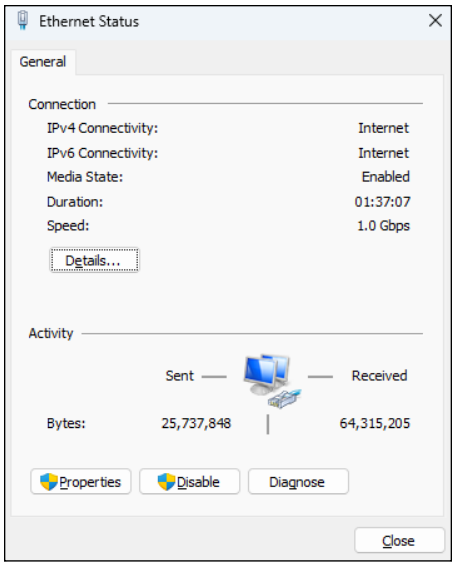


FIGURE 5-10: The adapter's Status dialog box.

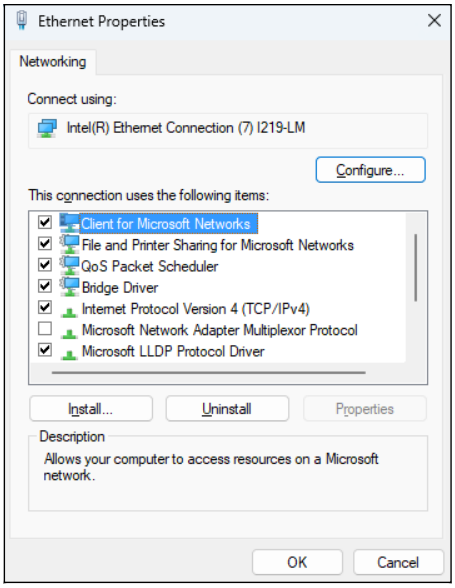


FIGURE 5-11: The adapter's Properties dialog box.

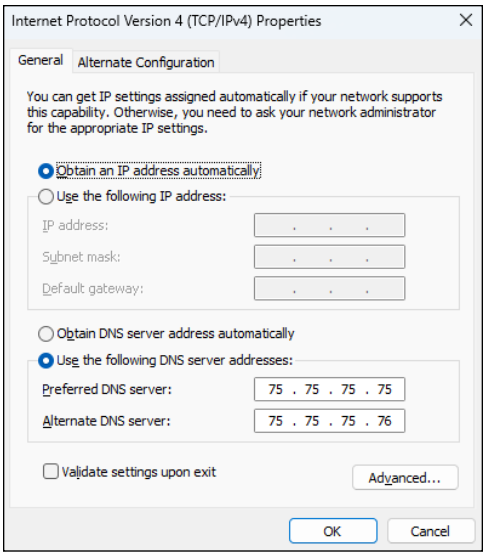


FIGURE 5-12: Enabling DHCP in the Internet Protocol Version 4 (TCP/IPv4) Properties dialog box.

Automatic private IP addressing

If a Windows computer is configured to use DHCP but the computer can't obtain an IP address from a DHCP server, the computer automatically assigns itself a private address by using a feature called Automatic Private IP Addressing (APIPA). APIPA assigns a private address from the 169.254.*x.x* range and uses a special algorithm to ensure that the address is unique on the network. As soon as the DHCP server becomes available, the computer requests a new address, so the APIPA address is used only while the DHCP server is unavailable.

Renewing and releasing leases

Normally, a DHCP client attempts to renew its lease when the lease is halfway to the point of being expired. For example, if a client obtains an eight-day lease, it attempts to renew the lease after four days. However, you can renew a lease sooner by issuing the `ipconfig /renew` command at a command prompt. You may want to do this if you changed the scope's configuration or if the client's IP configuration isn't working correctly.

You can also release a DHCP lease by issuing the `ipconfig /release` command at a command prompt. When you release a lease, the client computer no longer has a valid IP address. When you release an IP lease, you can't communicate with the network by using TCP/IP until you issue an `ipconfig /renew` command to renew the IP configuration or restart the computer.

IN THIS CHAPTER

- Discovering the basics of DNS
- Exploring zones
- Examining resource records
- Configuring a DNS server
- Setting up a DNS client

Chapter 6
DNS

Domain Name System (DNS) is the Transmission Control Protocol/Internet Protocol (TCP/IP) facility that lets you use names rather than numbers to refer to host computers. Without DNS, you'd buy books from 54.239.28.85 instead of from `www.amazon.com` and you'd sell your used furniture at 23.216.149.153 instead of on `www.ebay.com`.

Understanding how DNS works and how to set up a DNS server is crucial to setting up and administering a TCP/IP network. (For more on TCP/IP, see Book 2, Chapter 2.) This chapter introduces you to the basics of DNS, including how the DNS naming system works and how to set up a DNS server.


TECHNICAL STUFF

If you want to review the complete official specifications for DNS, look up RFC 1034 and 1035 at `www.ietf.org/rfc/rfc1034.txt` and `www.ietf.org/rfc/rfc1035.txt`, respectively.

Understanding DNS Names

DNS is a name service that provides a standardized system for providing names to identify TCP/IP hosts as well as a way to look up the IP address of a host, given the host's DNS name. For example, if you use DNS to look up the name `www.ebay.com`, you get the IP address of the eBay web host: 23.216.149.153. Thus, DNS allows you to access the eBay website by using the DNS name `www.ebay.com` instead of the site's IP address.



Note that by the time you get this book, read this chapter, and try to actually look up the IP address for `www.ebay.com`, there's a very good chance you'll discover that eBay has changed its IP address.

The following sections describe the basic concepts of DNS.

Domains and domain names

To provide a unique DNS name for every host computer on the internet, DNS uses a time-tested technique: Divide and conquer. DNS uses a hierarchical naming system that's similar to how folders are organized hierarchically on a Windows computer. Instead of folders, however, DNS organizes its names into domains. Each domain includes all the names that appear directly beneath it in the DNS hierarchy.

For example, Figure 6-1 shows a small portion of the DNS domain tree. At the very top of the tree is the *root domain*, which is the anchor point for all domains. Directly beneath the root domain are four top-level domains, named `edu`, `com`, `org`, and `gov`.

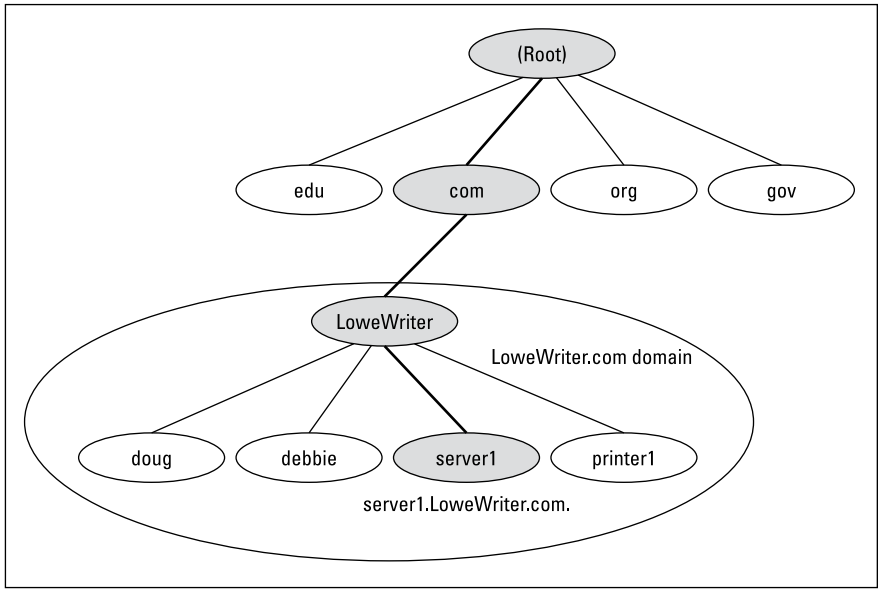


FIGURE 6-1: DNS names.

In reality, many more top-level domains than this exist in the internet's root domain. For more information, see the section "Top-Level Domains," later in this chapter.

Beneath the com domain in Figure 6-1 is another domain called LoweWriter, which happens to be my own personal domain. (Pretty clever, eh?) To completely identify this domain, you have to combine it with the name of its *parent domain* (in this case, com) to create the complete domain name: LoweWriter.com. Notice that the parts of the domain name are separated from each other with periods, which are called dots. As a result, when you read this domain name, you pronounce it *LoweWriter dot com.*

Beneath the LoweWriter node are four host nodes, named doug, debbie, server1, and printer1. Respectively, these correspond to three computers and a printer on my home network. You can combine the host name with the domain name to get the complete DNS name for each of my network's hosts. For example, the complete DNS name for my server is server1.LoweWriter.com. Likewise, my printer is printer1.LoweWriter.com.

Here are a few additional details that you need to remember about DNS names:

- **DNS names are not case sensitive.** As a result, LoweWriter and Lowewriter are treated as the same name, as are LOWEWRITER, LOWEwriter, and LoWeWrItEr. When you use a domain name, you can use capitalization to make the name easier to read, but DNS ignores the difference between capital and lowercase letters.
- **The name of each DNS node can be up to 63 characters long (not including the dot) and can include letters, numbers, and hyphens.**

  No other special characters are allowed.


WARNING

- **A *subdomain* is a domain that's beneath an existing domain.** For example, the com domain is actually a subdomain of the root domain. Likewise, LoweWriter is a subdomain of the com domain.
- **DNS is a hierarchical naming system that's similar to the hierarchical folder system used by Windows.**


TIP

  However, one crucial difference exists between DNS and the Windows naming convention. When you construct a complete DNS name, you start at the bottom of the tree and work your way up to the root. Thus, doug is the lowest node in the name doug.LoweWriter.com. In contrast, Windows paths are the opposite: They start at the root and work their way down. For example, in the path \Windows\System32\dns, dns is the lowest node.
- **The DNS tree can be up to 127 levels deep.** However, in practice, the DNS tree is pretty shallow. Most DNS names have just three levels (not counting the root). And although you'll sometimes see names with four or five levels, you'll rarely see more levels than that.

» **Although the DNS tree is shallow, it's very broad.** In other words, each of the top-level domains has a huge number of second-level domains immediately beneath it. For example, at the time of this writing, the `com` domain had well over 100 million second-level domains beneath it.

Fully qualified domain names

If a domain name ends with a trailing dot, that trailing dot represents the root domain, and the domain name is said to be a *fully qualified domain name* (also known as an FQDN). A fully qualified domain name is also called an *absolute name.* A fully qualified domain name is unambiguous because it identifies itself all the way back to the root domain. In contrast, if a domain name doesn't end with a trailing dot, the name may be interpreted in the context of some other domain. Thus, DNS names that don't end with a trailing dot are called *relative names.*

This is similar to how relative and absolute paths work in Windows. For example, if a path begins with a backslash, such as `\Windows\System32\dns`, the path is absolute. However, a path that doesn't begin with a backslash, such as `System32\dns`, uses the current directory as its starting point. If the current directory happens to be `\Windows`, then `\Windows\System32\dns` and `System32\dns` refer to the same location.

In many cases, relative and fully qualified domain names are interchangeable because the software that interprets them always interprets relative names in the context of the root domain. That's why, for example, you can type `www.wiley.com` (without the trailing dot) — not `www.wiley.com.` — to go to the Wiley home page in a web browser. Some applications, such as DNS servers, may interpret relative names in the context of a domain other than the root.

DNS AND URLS

Whenever you use a web browser to navigate to a resource on the internet, you use a *Uniform Resource Locator* (URL) to find your way. But the format of a URL is much more complex than most people assume.

A URL can consist of five distinct parts, with the following syntax:

```
scheme://authority/path?query#fragment
```

The following paragraphs describe these five parts:

- **Scheme:** The *scheme* identifies the protocol used to access the resource. The scheme is followed by a colon (`:`). When working with HTTP, this should be (obviously) `http:`. Note that the scheme is used to determine the default port for TCP/IP. So, unless you explicitly specify a different port, using `http:` as the scheme will give you port 80.
- **Authority:** The *authority* portion of a URL is strangely named. It generally consists of a host name that can be resolved by DNS. But it can have two additional elements. The host name can be preceded by a username followed by the at symbol (`@`). And it can be followed by a port number, separated from the host by a colon (`:`). So, the authority might be `doug@lowewriter.com`, which looks like an email address but isn't. Or, it might be `www.lowewriter.com:8080`, indicating that instead of the port number suggested by the scheme, I want to use port 8080.
- **Path:** The *path* identifies a specific resource on the host server. It can be the name of a file or, more commonly, a file system path that lists one or more folders separated by slashes (`/`). The path always begins with a slash to separate it from the host name. For example: `files/java/examples.zip`.

  A path may end with a filename, but that's not always necessary. If the filename is omitted, the HTTP server uses a default filename such as `index.html`.
- **Query:** The query part of a URL is optional but very useful. It provides additional information to the server. The query begins with a question mark (`?`) and consists of one or more key-value pairs in the form *`key=value`*. When you need more than one key-value pair in a query, separate the pairs with ampersands (`&`). For example, `?month=10&day=31` could be used to represent a date.
- **Fragment:** The fragment portion of a URL is used less commonly than the other elements but is still useful from time to time. It's usually used to refer to a specific part of the resource. For example, if the resource is an HTML file, the fragment may refer to a section within the file. For example, `#references` could be used to specify a references section within an HTML file.

Top-Level Domains

A *top-level domain* appears immediately beneath the root domain. Top-level domains come in two categories: generic domains and geographic domains. These categories are described in the following sections. (Actually, a third type of top-level domain exists, which is used for reverse lookups. I describe it later in this chapter, in the section "Reverse Lookup Zones.")

Generic domains

Generic domains are the popular top-level domains that you see most often on the internet. Originally, seven top-level organizational domains existed. In 2002, seven more were added to help ease the congestion of the original seven — in particular, the `com` domain.

Table 6-1 summarizes the original seven generic top-level domains.

TABLE 6-1 The Original Seven Top-Level Domains

Domain	Description
`com`	Commercial organizations
`edu`	Educational institutions
`gov`	Government institutions
`int`	International treaty organizations
`mil`	Military institutions
`net`	Network providers
`org`	Noncommercial organizations

Because the `com` domain ballooned to an almost unmanageable size in the late 1990s, the internet authorities approved seven new top-level domains in an effort to take some of the heat off of the `com` domain. Most of these domains, listed in Table 6-2, became available in 2002.

TABLE 6-2 The New Seven Top-Level Domains

Domain	Description
`aero`	Aerospace industry
`biz`	Business
`coop`	Cooperatives
`info`	Informational sites
`museum`	Museums
`name`	Individual users
`pro`	Professional organizations

Country code domains

Although the top-level domains are open to anyone, U.S. companies and organizations dominate them. An additional set of top-level domains corresponds to international country code designations. Organizations outside the United States often use these top-level domains to avoid the congestion of the generic domains.

Table 6-3 lists those country code top-level domains with more than 200 registered subdomains at the time of this writing, in alphabetical order. In all, about 300 geographic top-level domains exist. The exact number varies from time to time as political circumstances change.

TABLE 6-3 Country Code Top-Level Domains with More Than 200 Subdomains

Domain	Description	Domain	Description
ac	Ascension Island	ie	Ireland
ae	United Arab Emirates	in	India
ag	Antigua and Barbuda	is	Iceland
am	Armenia	it	Italy
an	Netherlands Antilles	jp	Japan
as	American Samoa	kz	Kazakhstan
at	Austria	la	Lao People's Democratic Republic
be	Belgium	li	Liechtenstein
bg	Bulgaria	lk	Sri Lanka
br	Brazil	lt	Lithuania
by	Belarus	lu	Luxembourg
bz	Belize	lv	Latvia
ca	Canada	ma	Morocco
cc	Cocos (Keeling) Islands	md	Moldova
ch	Switzerland	nl	Netherlands
cl	Chile	no	Norway
cn	China	nu	Niue

(continued)

TABLE 6-3 *(continued)*

Domain	Description	Domain	Description
cx	Christmas Island	pl	Poland
cz	Czech Republic	pt	Portugal
de	Germany	ro	Romania
dk	Denmark	ru	Russian Federation
ee	Estonia	se	Sweden
es	Spain	si	Slovenia
eu	European Union	sk	Slovakia
fi	Finland	st	São Tomé and Principe
fm	Micronesia	su	Soviet Union
fo	Faroe Islands	to	Tonga
fr	France	tv	Tuvalu
ge	Georgia	tw	Taiwan
gr	Greece	ua	Ukraine
hr	Croatia	us	United States
hu	Hungary	ws	Samoa

The Hosts File

Long ago, in a network far, far away, the entire internet was small enough that network administrators could keep track of it all in a simple text file. This file, called the *Hosts file,* simply listed the name and IP address of every host on the network. Each computer had its own copy of the Hosts file. The trick was keeping all those Hosts files up to date. Whenever a new host was added to the internet, each network administrator would manually update his copy of the Hosts file to add the new host's name and IP address.

As the internet grew, so did the Hosts file. In the mid-1980s, it became obvious that a better solution was needed. Imagine trying to track the entire internet today by using a single text file to record the name and IP address of the millions of hosts on the internet! DNS was invented to solve this problem.

Understanding the Hosts file is important for two reasons:

- **The Hosts file is not dead.** For small networks, a Hosts file may still be the easiest way to provide name resolution for the network's computers. In addition, a Hosts file can coexist with DNS. The Hosts file is always checked before DNS is used, so you can even use a Hosts file to override DNS if you want.
- **The Hosts file is the precursor to DNS.** DNS was devised to circumvent the limitations of the Hosts file. You'll be in a better position to appreciate the benefits of DNS when you understand how the Hosts file works.

The Hosts file is a simple text file that contains lines that match IP addresses with host names. You can edit the Hosts file with any text editor, including Notepad. The exact location of the Hosts file depends on the client operating system, as listed in Table 6-4.

TABLE 6-4 Location of the Hosts File

Operating System	Location of Hosts File
Windows	`c:\windows\system32\drivers\etc`
Unix/Linux	`/etc/hosts`

All TCP/IP implementations are installed with a starter Hosts file. For example, Listing 6-1 shows a typical Windows TCP/IP Hosts file. As you can see, the starter file begins with some comments that explain the purpose of the file.

The Hosts file ends with comments, which show the host mapping commands used to map for the host name `localhost`, mapped to the IP address 127.0.0.1. The IP address 127.0.0.1 is the standard loopback address. As a result, this entry allows a computer to refer to itself by using the name `localhost`.

Note that after the 127.0.0.1 `localhost` entry, another `localhost` entry defines the standard IPv6 loopback address (`::1`). This is required because all versions of Windows since Vista provide built-in support for IPv6.

Prior to Windows 7, these lines were not commented out in the Hosts file. But beginning with Windows 7, the name resolution for `localhost` is handled by DNS itself, so its definition isn't required in the Hosts file.

LISTING 6-1: A Sample Hosts File

```
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each entry
   should be kept
# on an individual line. The IP address should be placed in the first column
   followed by
# the corresponding host name. The IP address and the host name should be
   separated by at
# least one space.
#
# Additionally, comments (such as these) may be inserted on individual lines or
   following
# the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.97     rhino.acme.com          # source server
#       38.25.63.10     x.acme.com              # x client host

# localhost name resolution is handled within DNS itself.
#      127.0.0.1       localhost
#      ::1             localhost
```

To add an entry to the Hosts file, simply edit the file in any text editor. Then, add a line at the bottom of the file, after the `localhost` entry. Each line that you add should list the IP address and the host name that you want to use for the address. For example, to associate the host name `server1.LoweWriter.com` with the IP address 192.168.168.201, you add this line to the Hosts file:

```
192.168.168.201 server1.LoweWriter.com
```

Then, whenever an application requests the IP address of the host name `server1.LoweWriter.com`, the IP address 192.168.168.201 is returned.

You can also add an alias to a host mapping. This enables users to access a host by using the alias as an alternative name. For example, consider the following line:

```
192.168.168.201 server1.LoweWriter.com s1
```

Here, the device at address 192.168.168.201 can be accessed as `server1.LoweWriter.com` or just `s1`.

Listing 6-2 shows a Hosts file with several hosts defined.

LISTING 6-2: A Hosts File with Several Hosts Defined

```
Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each entry
   should be kept
# on an individual line. The IP address should be placed in the first column
   followed by
# the corresponding host name. The IP address and the host name should be
   separated by at
# least one space.
#
# Additionally, comments (such as these) may be inserted on individual lines or
   following
# the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.97     rhino.acme.com          # source server
#       38.25.63.10     x.acme.com              # x client host

# localhost name resolution is handled within DNS itself.
#     127.0.0.1       localhost
#     ::1             localhost
192.168.168.200 doug.LoweWriter.com #Doug's computer
192.168.168.201 server1.LoweWriter.com s1 #Main server
192.168.168.202 kristen.LoweWriter.com #Kristen's computer
192.168.168.203 printer1.LoweWriter.com p1 #HP Laser Printer
```

Note that even if your network uses DNS, every client still has a Hosts file.

Understanding DNS Servers and Zones

A *DNS server* is a computer that runs DNS server software, helps to maintain the DNS database, and responds to DNS name resolution requests from other computers. Although many DNS server implementations are available, the two most popular

are BIND and the Windows DNS service. BIND runs on Unix-based computers (including Linux computers), and Windows DNS (naturally) runs on Windows computers. Both provide essentially the same services and can interoperate.

The key to understanding how DNS servers work is to realize that the DNS database — that is, the list of all the domains, subdomains, and host mappings — is a massively distributed database. No single DNS server contains the entire DNS database. Instead, authority over different parts of the database is delegated to different servers throughout the internet.

For example, suppose that I set up a DNS server to handle name resolutions for my `LoweWriter.com` domain. Then, when someone requests the IP address of `doug.LoweWriter.com`, my DNS server can provide the answer. However, my DNS server wouldn't be responsible for the rest of the internet. Instead, if someone asks my DNS server for the IP address of some other computer, such as `coyote.acme.com`, my DNS server will have to pass the request on to another DNS server that knows the answer.

Zones

To simplify the management of the DNS database, the entire DNS namespace is divided into zones, and the responsibility for each zone is delegated to a particular DNS server. In many cases, zones correspond directly to domains. For example, if I set up a domain named `LoweWriter.com`, I can also set up a DNS zone called `LoweWriter.com` that's responsible for the entire `LoweWriter.com` domain.

However, the subdomains that make up a domain can be parceled out to separate zones, as shown in Figure 6-2. Here, a domain named `LoweWriter.com` has been divided into two zones. One zone, `us.LoweWriter.com`, is responsible for the entire `us.LoweWriter.com` subdomain. The other zone, `LoweWriter.com`, is responsible for the entire `LoweWriter.com` domain except for the `us.LoweWriter.com` subdomain.

Why would you do that? The main reason is to delegate authority for the zone to separate servers. For example, Figure 6-2 suggests that part of the `LoweWriter.com` domain is administered in the United States and that part of it is administered in France. The two zones in the figure allow one server to be completely responsible for the U.S. portion of the domain, and the other server handles the rest of the domain.

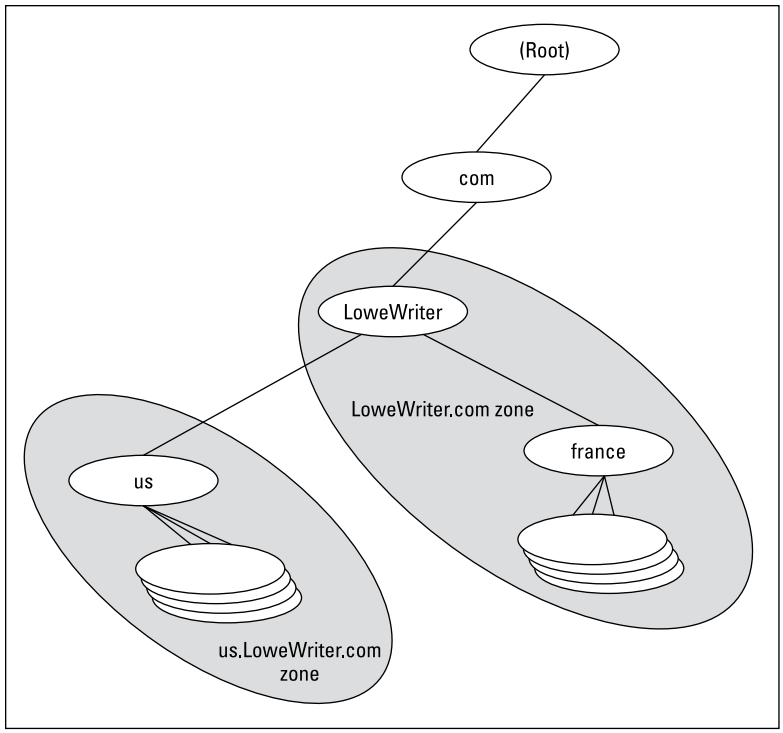


FIGURE 6-2: DNS zones.

THE OLD PHONY HOSTS FILE TRICK

The Hosts file can be the basis of a fun, practical joke. Of course, neither I nor my editors or publishers recommend that you actually do this. If it gets you into trouble, don't send your lawyers to me. This sidebar is here only to let you know what to do if it happens to you.

The idea is to edit your poor victim's Hosts file so that whenever the user tries to access his favorite website, a site of your choosing comes up instead. For example, if you're trying to get your husband to take you on a cruise, add a line to his Hosts file that replaces his favorite website with the website for a cruise line. For example, this line should do the trick:

```
104.92.181.171 www.espn.com
```

Now, whenever your husband tries to call up the ESPN website, he'll get the Carnival Cruise Lines home page instead. (To find out the IP address of the website you want displayed, open a command prompt and type **ping** followed by a space and the URL of the

(continued)

(continued)

desired website. This will usually display the correct IP address. If not, you can use the `nslookup` command, as described in Book 2, Chapter 7.)

Of course, to actually pull a stunt like this would be completely irresponsible, especially if you didn't first make a backup copy of the Hosts file, just in case it somehow gets messed up.

Be warned: If the wrong websites suddenly start coming up, check your Hosts file to see whether it's been tampered with.

The following are the two basic types of zones:

- A **primary zone** is the master copy of a zone. The data for a primary zone is stored in the local database of the DNS server that hosts the primary zone. Only one DNS server can host a particular primary zone. Any updates to the zone must be made to the primary zone.
- A **secondary zone** is a read-only copy of a zone. When a server hosts a secondary zone, the server doesn't store a local copy of the zone data. Instead, it obtains its copy of the zone from the zone's primary server by using a process called *zone transfer.* Secondary servers must periodically check primary servers to see whether their secondary zone data is still current. If not, a zone transfer is initiated to update the secondary zone.

Primary and secondary servers

Each DNS server is responsible for one or more zones. The following are the two different roles that a DNS server can take:

- **Primary server for a zone,** which means that the DNS server hosts a primary zone. The data for the zone is stored in files on the DNS server. Every zone must have one primary server.
- **Secondary server for a zone,** which means that the DNS server obtains the data for a secondary zone from a primary server. Every zone should have at least one secondary server. That way, if the primary server goes down, the domain defined by the zone can be accessed via the secondary server or servers.



WARNING

A secondary server should be on a different subnet than the zone's primary server. If the primary and secondary servers are on the same subnet, both servers will be unavailable if the router that controls the subnet goes down.

Note that a single DNS server can be the primary server for some zones and a secondary server for other zones. A server is said to be *authoritative* for the primary and secondary zones that it hosts because it can provide definitive answers for queries against those zones.

Root servers

The core of DNS comprises the *root servers*, which are authoritative for the entire internet. The main function of the root servers is to provide the address of the DNS servers that are responsible for each of the top-level domains. These servers, in turn, can provide the DNS server address for subdomains beneath the top-level domains.

The root servers are a major part of the glue that holds the internet together. As you can imagine, they're swamped with requests day and night. A total of 13 root servers are located throughout the world. Table 6-5 lists the IP address and the organization that oversees the operation of each of the 13 root servers.

TABLE 6-5 The 13 Root Servers

Server	IP Address	Operator
A	198.41.0.4	VeriSign Global Registry Services
B	170.247.170.2	University of Southern California
C	192.33.4.12	Cogent Communications
D	199.7.91.13	University of Maryland
E	192.203.230.10	NASA Ames Research Center
F	192.5.5.241	Internet Systems Consortium
G	192.112.36.4	U.S. Department of Defense (NIC)
H	198.97.190.53	U.S. Army Research Lab
I	192.36.148.17	Netnod
J	192.58.128.30	Verisign, Inc.
K	193.0.14.129	RIPE NCC
L	199.7.83.42	ICANN
M	202.12.27.33	WIDE Project

DNS servers learn how to reach the root servers by consulting a *root hints* file that's located on the server. In the Unix/Linux world, this file is known as `named.root` and can be found at `/etc/named.root`. For Windows, the root hints are stored within Active Directory rather than as a separate file. Listing 6-3 shows a typical `named.root` file.

LISTING 6-3: **The named.root File**

```
;       This file holds the information on root name servers needed to
   initialize cache of
;       internet domain name servers (for example, reference this file in the
;       "cache  .  <file>" configuration file of BIND domain name servers).
;
;       This file is made available by InterNIC under anonymous FTP as
;           file                /domain/named.cache
;           on server           FTP.INTERNIC.NET
;       -OR-                    RS.INTERNIC.NET
;
;       last update:     August 12, 2020
;       related version of root zone:     2020081201
;
; FORMERLY NS.INTERNIC.NET
;
.                        3600000      NS    A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET.      3600000      A     198.41.0.4
A.ROOT-SERVERS.NET.      3600000      AAAA  2001:503:ba3e::2:30
;
; FORMERLY NS1.ISI.EDU
;
.                        3600000      NS    B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET.      3600000      A     199.9.14.201
B.ROOT-SERVERS.NET.      3600000      AAAA  2001:500:200::b
;
; FORMERLY C.PSI.NET
;
.                        3600000      NS    C.ROOT-SERVERS.NET.
C.ROOT-SERVERS.NET.      3600000      A     192.33.4.12
C.ROOT-SERVERS.NET.      3600000      AAAA  2001:500:2::c
;
; FORMERLY TERP.UMD.EDU
;
.                        3600000      NS    D.ROOT-SERVERS.NET.
D.ROOT-SERVERS.NET.      3600000      A     199.7.91.13
D.ROOT-SERVERS.NET.      3600000      AAAA  2001:500:2d::d
;
```

```
; FORMERLY NS.NASA.GOV
;
.                       3600000      NS    E.ROOT-SERVERS.NET.
E.ROOT-SERVERS.NET.     3600000      A     192.203.230.10
E.ROOT-SERVERS.NET.     3600000      AAAA  2001:500:a8::e
;
; FORMERLY NS.ISC.ORG
;
.                       3600000      NS    F.ROOT-SERVERS.NET.
F.ROOT-SERVERS.NET.     3600000      A     192.5.5.241
F.ROOT-SERVERS.NET.     3600000      AAAA  2001:500:2f::f
;
; FORMERLY NS.NIC.DDN.MIL
;
.                       3600000      NS    G.ROOT-SERVERS.NET.
G.ROOT-SERVERS.NET.     3600000      A     192.112.36.4
G.ROOT-SERVERS.NET.     3600000      AAAA  2001:500:12::d0d
;
; FORMERLY AOS.ARL.ARMY.MIL
;
.                       3600000      NS    H.ROOT-SERVERS.NET.
H.ROOT-SERVERS.NET.     3600000      A     198.97.190.53
H.ROOT-SERVERS.NET.     3600000      AAAA  2001:500:1::53
;
; FORMERLY NIC.NORDU.NET
;
.                       3600000      NS    I.ROOT-SERVERS.NET.
I.ROOT-SERVERS.NET.     3600000      A     192.36.148.17
I.ROOT-SERVERS.NET.     3600000      AAAA  2001:7fe::53
;
; OPERATED BY VERISIGN, INC.
;
.                       3600000      NS    J.ROOT-SERVERS.NET.
J.ROOT-SERVERS.NET.     3600000      A     192.58.128.30
J.ROOT-SERVERS.NET.     3600000      AAAA  2001:503:c27::2:30
;
; OPERATED BY RIPE NCC
;
.                       3600000      NS    K.ROOT-SERVERS.NET.
K.ROOT-SERVERS.NET.     3600000      A     193.0.14.129
K.ROOT-SERVERS.NET.     3600000      AAAA  2001:7fd::1
;
; OPERATED BY ICANN
;
.                       3600000      NS    L.ROOT-SERVERS.NET.
L.ROOT-SERVERS.NET.     3600000      A     199.7.83.42
L.ROOT-SERVERS.NET.     3600000      AAAA  2001:500:9f::42
;
```

(continued)

LISTING 6-3: ***(continued)***

```
; OPERATED BY WIDE
;
.                       3600000      NS    M.ROOT-SERVERS.NET.
M.ROOT-SERVERS.NET.     3600000      A     202.12.27.33
M.ROOT-SERVERS.NET.     3600000      AAAA  2001:dc3::35
; End of file
```

Caching

DNS servers don't really like doing all that work to resolve DNS names, but they're not stupid. They know that if a user visits `www.wiley.com` today, he'll probably do it again tomorrow. As a result, name servers keep a cache of query results. The next time the user visits `www.wiley.com`, the name server is able to resolve this name without having to query all those other name servers.

The internet is constantly changing, however, so cached data can quickly become obsolete. For example, suppose that John Wiley & Sons, Inc., switches its website to a different server. It can update its name servers to reflect the new IP address, but any name servers that have a cached copy of the query will be out of date.

To prevent this from being a major problem, DNS data is given a relatively short expiration time. The expiration value for DNS data is called the *TTL* (Time to Live). TTL is specified in seconds. Thus, a TTL of 60 means the data is kept for one minute.

Understanding DNS Queries

When a DNS client needs to resolve a DNS name to an IP address, it uses a library routine — a *resolver* — to handle the query. The resolver takes care of sending the query message over the network to the DNS server, receiving and interpreting the response, and informing the client of the results of the query.

A DNS client can make two basic types of queries: recursive and iterative. The following list describes the difference between these two query types. (The following discussion assumes that the client is asking the server for the IP address of a host name, which is the most common type of DNS query. You find out about other types of queries later; they, too, can be either recursive or iterative.)

- **Recursive queries:** When a client issues a *recursive DNS query,* the server must reply with either the IP address of the requested host name or an error message indicating that the host name doesn't exist. If the server doesn't have the information, it asks another DNS server for the IP address. When the first server finally gets the IP address, it sends it back to the client. If the server determines that the information doesn't exist, it returns an error message.
- **Iterative queries:** When a server receives an iterative query, it returns the IP address of the requested host name if it knows the address. If the server doesn't know the address, it returns a *referral,* which is simply the address of a DNS server that should know. The client can then issue an iterative query to the server to which it was referred.

Normally, DNS clients issue recursive queries to DNS servers. If the server knows the answer to the query, it replies directly to the client. If not, the server issues an iterative query to a DNS server that it thinks should know the answer. If the original server gets an answer from the second server, it returns the answer to the client. If the original server gets a referral to a third server, the original server issues an iterative query to the third server. The original server keeps issuing iterative queries until it either gets the answer or an error occurs. It then returns the answer or the error to the client.

Confused? I can understand why. An example may help to clear things up. Suppose that a user wants to view the web page `www.wiley.com`. The following sequence of steps occurs to resolve this address:

1. **The browser asks the client computer's resolver to find the IP address of `www.wiley.com`.**
2. **The resolver issues a recursive DNS query to its name server.**

   In this case, I'll call the name server `ns1.LoweWriter.com`.
3. **The name server `ns1LoweWriter.com` checks whether it knows the IP address of `www.wiley.com`.**

   It doesn't, so the name server issues an iterative query to one of the root name servers to see whether it knows the IP address of `www.wiley.com`.
4. **The root name server doesn't know the IP address of `www.wiley.com`, so it returns a list of the name servers that are authoritative for the `com` domain.**

5. **The** ns1.LoweWriter.com **name server picks one of the** com **domain name servers and sends it an iterative query for** www.wiley.com**.**

6. **The** com **name server doesn't know the IP address of** www.wiley.com**, so it returns a list of the name servers that are authoritative for the** wiley.com **domain.**

7. **The** ns1.LoweWriter.com **name server picks one of the name servers for the** wiley.com **domain and sends it an iterative query for** www.wiley.com**.**

8. **The** wiley.com **name server knows the IP address for** www.wiley.com**, so the name server returns it.**

9. **The** ns1.LoweWriter.com **name server shouts with joy for having finally found the IP address for** www.wiley.com**. It gleefully returns this address to the client. It also caches the answer so that the next time the user looks for** www.wiley.com**, the name server won't have to contact other name servers to resolve the name.**

10. **The client also caches the results of the query.**

    The next time the client needs to look for www.wiley.com, the client can resolve the name without troubling the name server.

Zone Files and Resource Records

Each DNS zone is defined by a *zone file* (also known as a *DNS database* or a *master file*). For Windows DNS servers, the name of the zone file is *domain*.zone. For example, the zone file for the LoweWriter.com zone is named LoweWriter.com.zone. For BIND DNS servers, the zone files are named db.*domain*. Thus, the zone file for the LoweWriter.com domain would be db.LoweWriter.com. The format of the zone file contents is the same for both systems, however.

A zone file consists of one or more resource records. Creating and updating the resource records that comprise the zone files is one of the primary tasks of a DNS administrator. The Windows DNS server provides a friendly graphical interface to the resource records. However, you should still be familiar with how to construct resource records.

Resource records are written as simple text lines, with the following fields:

```
Owner TTL Class Type RDATA
```

These fields must be separated from each other by one or more spaces. The following list describes the five resource record fields:

- **Owner:** The name of the DNS domain or the host that the record applies to. This is usually specified as a fully qualified domain name (with a trailing dot) or as a simple host name (without a trailing dot), which is then interpreted in the context of the current domain.

  You can also specify a single at symbol (@) as the owner name. In that case, the current domain is used.



TIP

- **TTL:** Also known as *Time to Live;* the number of seconds that the record should be retained in a server's cache before it's invalidated. If you omit the TTL value for a resource record, a default TTL is obtained from the Start of Authority (SOA) record.
- **Class:** Defines the protocol to which the record applies. You should always specify `IN`, for the Internet Protocol. If you omit the class field, the last class field that you specified explicitly is used. As a result, you'll sometimes see zone files that specify `IN` only on the first resource record (which must be an SOA record) and then allow it to default to `IN` on all subsequent records.
- **Type:** The resource record type. The most commonly used resource types are summarized in Table 6-6 and are described separately later in this section. Like the Class field, you can also omit the Type field and allow it to default to the last specified value.
- **RDATA:** Resource record data that is specific to each record type.

TABLE 6-6 Common Resource Record Types

Type	Name	Description
SOA	Start of Authority	Identifies a zone
NS	Name Server	Identifies a name server that is authoritative for the zone
A	Address	Maps a fully qualified domain name to an IP address
CNAME	Canonical Name	Creates an alias for a fully qualified domain name
MX	Mail Exchange	Identifies the mail server for a domain
PTR	Pointer	Maps an IP address to a fully qualified domain name for reverse lookups



REMEMBER

Most resource records fit on one line. If a record requires more than one line, you must enclose the data that spans multiple lines in parentheses.



TIP

You can include comments to clarify the details of a zone file. A comment begins with a semicolon (;) and continues to the end of the line. If a line begins with a semicolon, the entire line is a comment. You can also add a comment to the end of a resource record. You see examples of both types of comments later in this chapter.

SOA records

Every zone must begin with an SOA record, which names the zone and provides default information for the zone. Table 6-7 lists the fields that appear in the RDATA section of an SOA record. Note that these fields are positional, so you should include a value for all of them and list them in the order specified. Because the SOA record has so many RDATA fields, you'll probably need to use parentheses to continue the SOA record onto multiple lines.

TABLE 6-7 RDATA Fields for an SOA Record

Name	Description
MNAME	The domain name of the name server that is authoritative for the zone.
RNAME	An email address (specified in domain name format; not regular email format) of the person responsible for this zone.
SERIAL	The serial number of the zone. Secondary zones use this value to determine whether they need to initiate a zone transfer to update their copy of the zone.
REFRESH	A time interval that specifies how often a secondary server should check whether the zone needs to be refreshed. A typical value is 3600 (one hour).
RETRY	A time interval that specifies how long a secondary server should wait after requesting a zone transfer before trying again. A typical value is 600 (ten minutes).
EXPIRE	A time interval that specifies how long a secondary server should keep the zone data before discarding it. A typical value is 86400 (one day).
MINIMUM	A time interval that specifies the TTL value to use for zone resource records that omit the TTL field. A typical value is 3600 (one hour).

Note two things about the SOA fields:

- **The email address of the person responsible for the zone is given in DNS format, not in normal email format.** Thus, you separate the user from the

mail domain with a dot rather than an at symbol (@). For example, doug@LoweWriter.com would be listed as doug.lowewriter.com.

- **The serial number should be incremented every time you change the zone file.** If you edit the file via the graphic interface provided by Windows DNS, the serial number is incremented automatically. However, if you edit the zone file via a simple text editor, you have to manually increment the serial number.

Here's a typical example of an SOA record, with judicious comments to identify each field:

```
lowewriter.com. IN SOA (
ns1.lowewriter.com ; authoritative name server
doug.lowewriter.com ; responsible person
148 ; version number
3600 ; refresh (1 hour)
600 ; retry (10 minutes)
86400 ; expire (1 day)
3600 ) ; minimum TTL (1 hour)
```

NS records

Name server (NS) records identify the name servers that are authoritative for the zone. Every zone must have at least one NS record. Using two or more NS records is better so that if the first name server is unavailable, the zone will still be accessible.

The owner field should either be the fully qualified domain name for the zone, with a trailing dot, or an at symbol (@). The RDATA consists of just one field: the fully qualified domain name of the name server.

The following examples show two NS records that serve the lowewriter.com domain:

```
lowewriter.com. IN NS ns1.lowewriter.com.
lowewriter.com. IN NS ns2.lowewriter.com.
```

A records

Address (A) records are the meat of the zone file: They provide the IP addresses for each of the hosts that you want to make accessible via DNS. In an A record, you usually list just the host name in the owner field, thus allowing DNS to add the

domain name to derive the fully qualified domain name for the host. The RDATA field for the A record is the IP address of the host.

The following lines define various hosts for the `LoweWriter.com` domain:

```
doug IN A 192.168.168.200
server1 IN A 192.168.168.201
debbie IN A 192.168.168.202
printer1 IN A 192.168.168.203
router1 IN A 207.126.127.129
www IN A 64.71.129.102
```

Notice that for these lines, I don't specify the fully qualified domain names for each host. Instead, I just provide the host name. DNS will add the name of the zone's domain to these host names in order to create the fully qualified domain names.

If I wanted to be more explicit, I could list these A records like this:

```
doug.lowewriter.com. IN A 192.168.168.200
server1.lowewriter.com. IN A 192.168.168.201
debbie.lowewriter.com. IN A 192.168.168.202
printer1.lowewriter.com. IN A 192.168.168.203
router1.lowewriter.com IN A 207.126.127.129
www.lowewriter.com. IN A 64.71.129.102
```

However, all this does is increase the chance for error. Plus, it creates more work for you later if you decide to change your network's domain.

CNAME records

A *Canonical Name* (CNAME) record creates an alias for a fully qualified domain name. When a user attempts to access a domain name that is actually an alias, the DNS system substitutes the real domain name — known as the *Canonical Name* — for the alias. The owner field in the CNAME record provides the name of the alias that you want to create. Then, the RDATA field provides the Canonical Name — that is, the real name of the host.

For example, consider these resource records:

```
ftp.lowewriter.com. IN A 207.126.127.132
files.lowewriter.com. IN CNAME www1.lowewriter.com.
```

Here, the host name of an FTP server at 207.126.127.132 is `ftp.lowewriter.com`. The CNAME record allows users to access this host as `files.lowewriter.com` if they prefer.

PTR records

A *Pointer* (PTR) record is the opposite of an address record: It provides the fully qualified domain name for a given address. The owner field should specify the reverse lookup domain name, and the RDATA field specifies the fully qualified domain name. For example, the following record maps the address 64.71.129.102 to `www.lowewriter.com`:

```
102.129.71.64.in-addr.arpa. IN PTR www.lowewriter.com.
```

PTR records don't usually appear in normal domain zones. Instead, they appear in special reverse lookup zones. For more information, see the section "Reverse Lookup Zones," later in this chapter.

MX records

Mail Exchange (MX) records identify the mail server for a domain. The owner field provides the domain name that users address mail to. The RDATA section of the record has two fields. The first is a priority number used to determine which mail servers to use when several are available. The second is the fully qualified domain name of the mail server itself.

For example, consider the following MX records:

```
lowewriter.com. IN MX 0 mail1.lowewriter.com.
lowewriter.com. IN MX 10 mail2.lowewriter.com.
```

In this example, the `lowewriter.com` domain has two mail servers, named `mail1.lowewriter.com` and `mail2.lowewriter.com`. The priority numbers for these servers are `0` and `10`. Because it has a lower priority number, mail will be delivered to `mail1.lowewriter.com` first. The `mail2.lowewriter.com` server will be used only if `mail1.lowewriter.com` isn't available.


TIP

The server name specified in the RDATA section should be an actual host name, not an alias created by a CNAME record. Although some mail servers can handle MX records that point to CNAMEs, not all can. As a result, you shouldn't specify an alias in an MX record.



WARNING

Be sure to create a reverse lookup record (PTR, described in the next section) for your mail servers. Some mail servers won't accept mail from a server that doesn't have valid reverse lookup entries.

Reverse Lookup Zones

Normal DNS queries ask a name server to provide the IP address that corresponds to a fully qualified domain name. This kind of query is a *forward lookup.* A *reverse lookup* is the opposite of a forward lookup: It returns the fully qualified domain name of a host based on its IP address.

Reverse lookups are possible because of a special domain called the `in-addr.arpa` domain, which provides a separate fully qualified domain name for every possible IP address on the internet. To enable a reverse lookup for a particular IP address, all you have to do is create a PTR record in a reverse lookup zone (a zone that is authoritative for a portion of the `in-addr.arpa` domain). The PTR record maps the `in-addr.arpa` domain name for the address to the host's actual domain name.

The technique used to create the reverse domain name for a given IP address is pretty clever. It creates subdomains beneath the `in-addr.arpa` domain by using the octets of the IP address, listing them in reverse order. For example, the reverse domain name for the IP address 207.126.67.129 is `129.67.126.207.in-addr.arpa`.

Why list the octets in reverse order? Because that correlates the network portions of the IP address (which work from left to right) with the subdomain structure of DNS names (which works from right to left). The following description should clear this up:

- The 255 possible values for the first octet of an IP address each have a subdomain beneath the `in-addr.arpa` domain. For example, any IP address that begins with 207 can be found in the `207.in-addr.arpa` domain.
- Within this domain, each of the possible values for the second octet can be found as a subdomain of the first octet's domain. Thus, any address that begins with 207.126 can be found in the `126.207.in-addr.arpa` domain.
- The same holds true for the third octet, so any address that begins with 207.126.67 can be found in the `67.126.207.in-addr.arpa` domain.
- By the time you get to the fourth octet, you've pinpointed a specific host. The fourth octet completes the fully qualified reverse domain name. Thus, 207.126.67.129 is mapped to `129.67.126.207.in-addr.arpa`.

As a result, to determine the fully qualified domain name for the computer at 207.126.67.129, the client queries its DNS server for the FQDN that corresponds to `129.67.126.207.in-addr.arpa`.

Working with the Windows DNS Server

Installing and managing a DNS server depends on the network operating system (NOS) that you're using. The following sections are specific to working with a DNS server in Windows Server 2025. Working with BIND in a Unix/Linux environment is similar but without the help of a graphical user interface (GUI), and working with Windows Server 2016 is similar.

You can install the DNS server on a Windows server from the Server Manager application. Open the Server Manager and choose Manage➪Add Roles and Features. Then, follow the wizard's instructions to add the DNS Role.

After you set up a DNS server, you can manage the DNS server from the DNS Manager, as shown in Figure 6-3. From this management console, you can perform common administrative tasks, such as adding additional zones, changing zone settings, adding A or MX records to an existing zone, and so on. The DNS Manager hides the details of the actual resource records from you, thus allowing you to work with a friendly GUI instead.

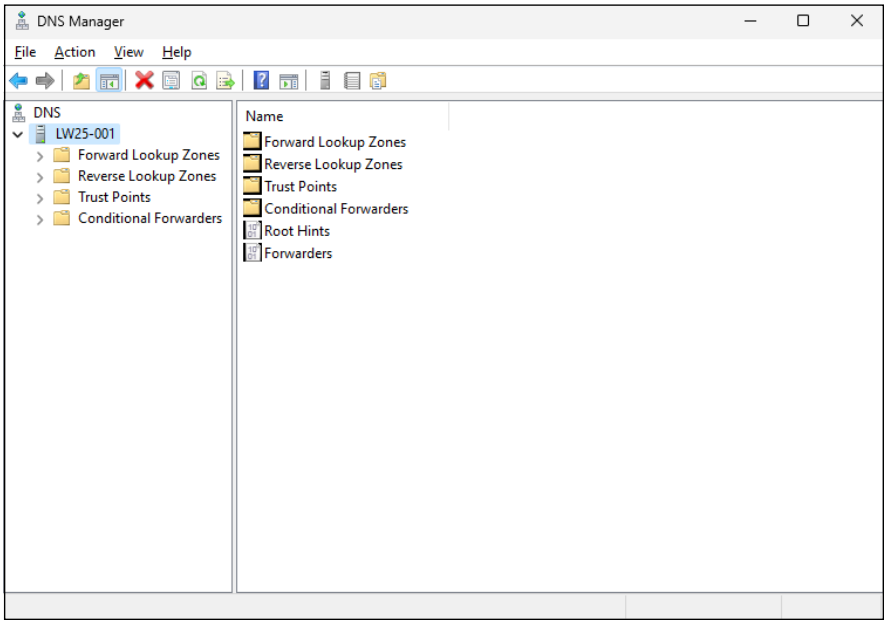


FIGURE 6-3: The DNS Manager.

Creating a new zone

To add a zone to be managed by a DNS server, follow these steps:

1. **Right-click Forward Lookup Zones and choose New Zone.**

   The New Zone Wizard, shown in Figure 6-4, appears.



FIGURE 6-4: The New Zone Wizard greets you.

2. **Click Next.**

   The wizard starts by asking which type of zone you'd like to create (see Figure 6-5). There are three choices:

   - **Primary Zone:** A zone that will be managed directly by this server.
   - **Secondary Zone:** A read-only copy of a zone that is managed by another DNS server.
   - **Stub Zone:** Creates a copy of a zone with just an NS, SOA, and possible H records. A stub zone is not authoritative for the zone.

   For our purposes here, we'll create a primary zone.

3. **Select Primary Zone and click Next.**

   The wizard asks for the name of the zone (see Figure 6-6).

4. **Enter the zone name and click Next.**

   The zone name should be the full DNS name of the zone you want to manage (for example, `lowewriter.com`).

   The wizard offers to create a new zone file, as shown in Figure 6-7.

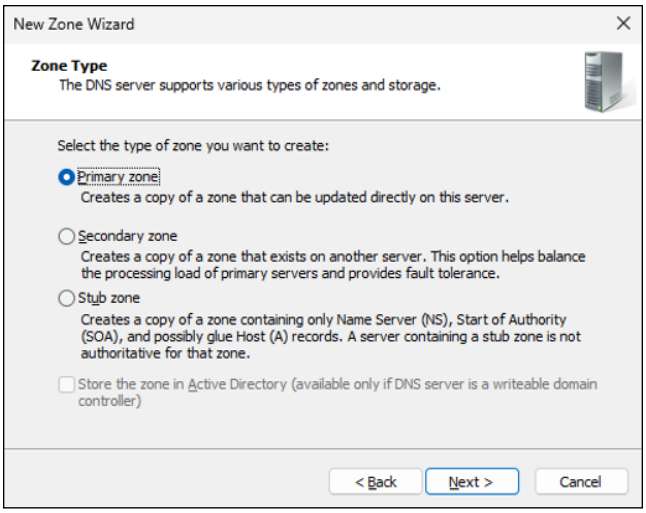


FIGURE 6-5: The New Zone Wizard asks for the zone type.

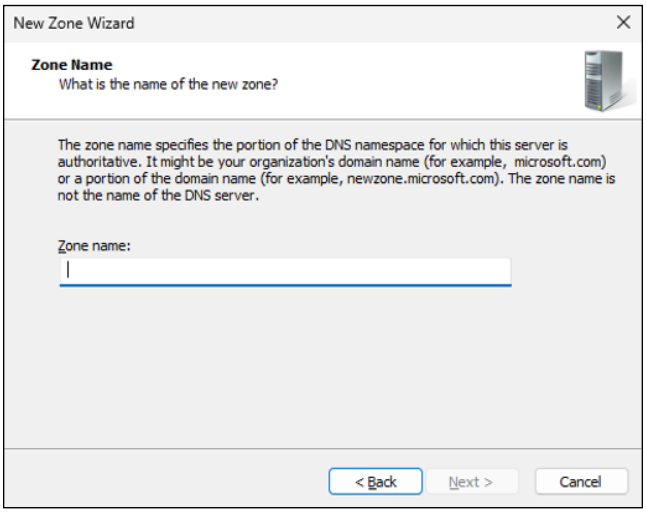


FIGURE 6-6: The New Zone Wizard asks for the zone name.

5. Click Next to let the wizard create a new zone file.

Optionally, you could import an existing zone file. You must first obtain the zone file from an existing DNS server, and then drop it in the `%SystemRoot%\System32\dns` folder.

The wizard offers to let you switch from a secure mode in which DNS must be updated manually to a dynamic update mode, which is terribly insecure because it allows any DNS client to update the zone (see Figure 6-8). Don't enable dynamic updates.

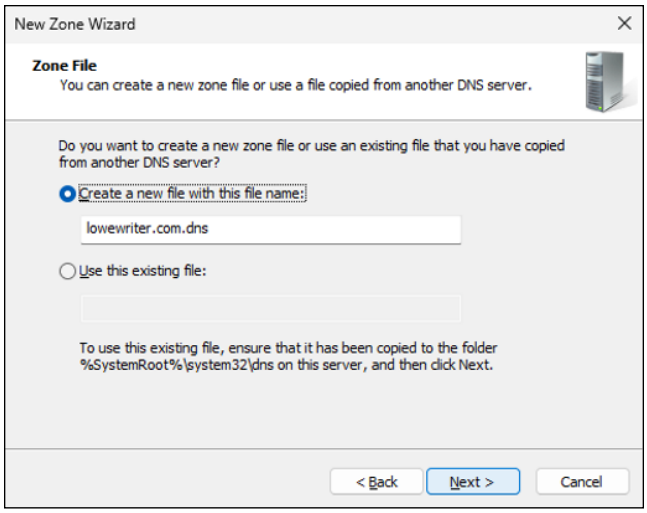


FIGURE 6-7: The New Zone Wizard offers to create a new zone file.

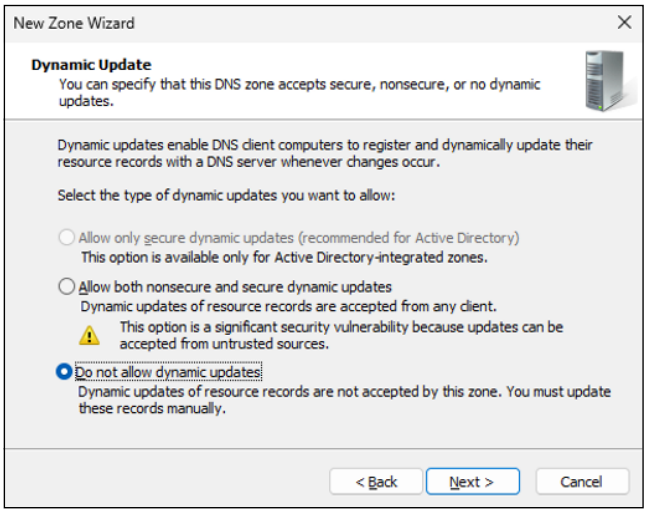


FIGURE 6-8: The New Zone Wizard asks you how the zone should be updated.

6. **Select Do Not Allow Dynamic Updates and click Next.**

   The wizard displays a summary of all your choices.

7. **Click Finish.**

   Presto! Your new zone now appears in the Forward Lookup Zones section of the DNS Manager, as shown in Figure 6-9.

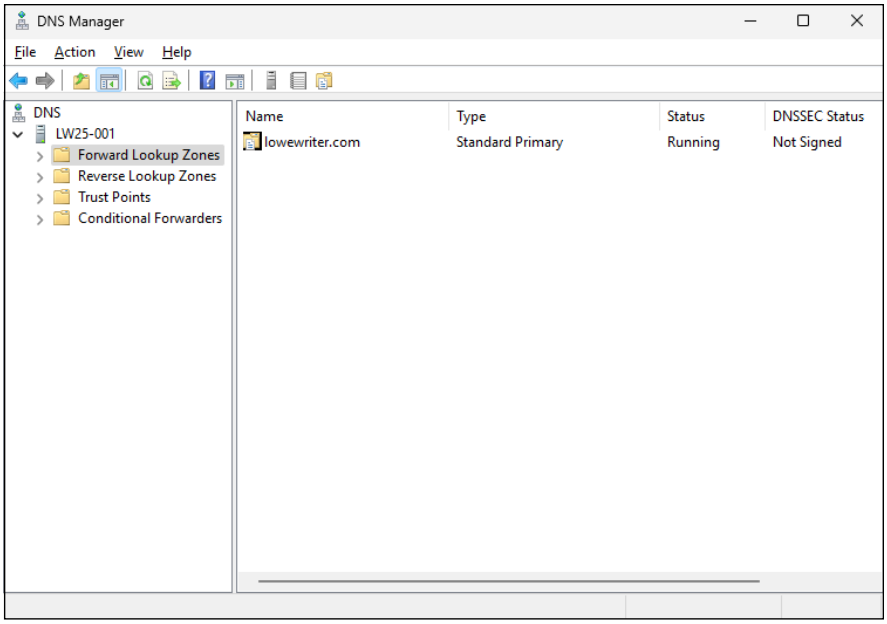


FIGURE 6-9: Congratulations on creating your first DNS zone!

Creating a new host record

To add a new host (that is, an A record) to a zone, right-click the zone in the Forward Management Zones section of the DNS Manager. Then choose the Add New Host command. This brings up the New Host dialog box, shown in Figure 6-10. From this dialog box, specify the following information.

- **Name:** The host name for the new host.
- **IP Address:** The host's IP address.
- **Create Associated Pointer (PTR) Record:** Automatically creates a PTR record in the reverse lookup zone file. Select this option if you want to allow reverse lookups for the host.

You can add other records, such as MX or CNAME records, in the same way.

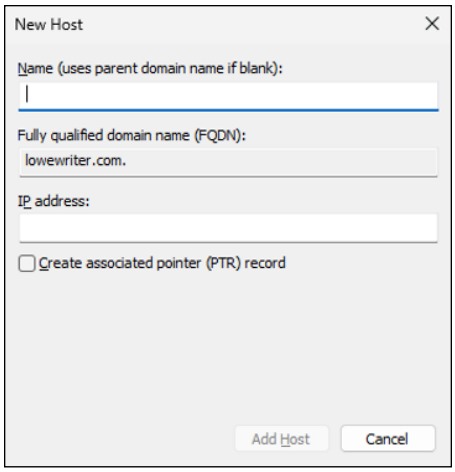


FIGURE 6-10: The New Host dialog box.

How to Configure a Windows DNS Client

Client computers don't need much configuration in order to work properly with DNS. The client must have the address of at least one DNS server. Usually, this address is supplied by DHCP, so if the client is configured to obtain its IP address from a DHCP server, it will also obtain the DNS server address from DHCP.

If you must configure DHCP manually, follow these steps:

1. **Open the Control Panel.**

   If you haven't already, switch to Small Icons view.

2. **Open Network and Sharing Center.**

3. **Click the link for your wired or wireless network adapter.**

   The adapter's Status dialog box appears.

4. **Click Properties.**

   The adapter's Properties dialog box appears.

5. **Select Internet Protocol Version 4 and then click Properties.**

   The Internet Protocol Version 4 (TCP/IPv4) Properties dialog box, shown in Figure 6-11, appears.

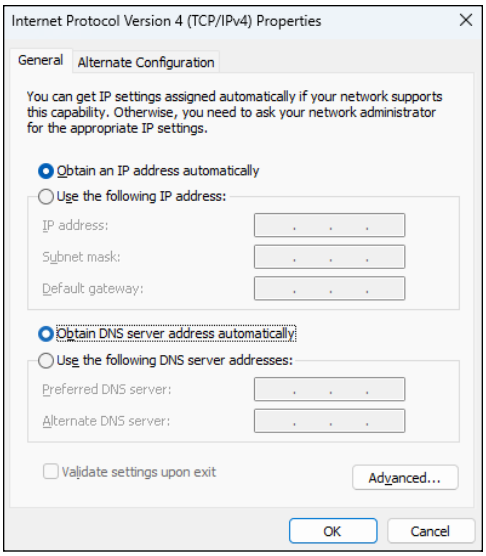


FIGURE 6-11: Configuring a Windows client to obtain its DNS address from DHCP.

6. **Configure the DNS options.**

   To pull DNS from the DHCP server, choose Obtain DNS Server Address Automatically.

   To manually specify the DNS servers, choose Use the Following DNS Server Addresses, and then enter the IP addresses for your preferred and secondary DNS servers.

7. **Click OK several times until all the dialog boxes you've opened are gone.**

IN THIS CHAPTER

- Recognizing tools and commands
- Making all your hosts sing with `IPConfig` and `Ping`

Chapter 7
TCP/IP Tools and Commands

Most client and server operating systems that support Transmission Control Protocol/Internet Protocol (TCP/IP) come with a suite of commands and tools that are designed to let you examine TCP/IP configuration information and diagnose and correct problems. Although the exact form of these commands varies between Windows and Unix/Linux, most are surprisingly similar. This chapter is a reference to the most commonly used TCP/IP commands.

Using the arp Command

Using the `arp` command allows you to display and modify the Address Resolution Protocol (ARP) cache, which is a simple mapping of IP addresses to MAC addresses. Each time a computer's TCP/IP stack uses ARP to determine the Media Access Control (MAC) address for an IP address, it records the mapping in the ARP cache so that future ARP lookups go faster.

If you use the `arp` command without any parameters, you get a list of the command's parameters. To display the ARP cache entry for a specific IP address, use an `-a` switch followed by the IP address. For example:

```
C:\>arp -a 192.168.168.22
Interface: 192.168.168.21 --- 0x10004
Internet Address Physical Address Type
192.168.168.22 00-60-08-39-e5-a1 dynamic
C:\>
```

You can display the complete ARP cache by using `-a` without specifying an IP address, like this:

```
C:\>arp -a
Interface: 192.168.168.21 --- 0x10004
Internet Address Physical Address Type
192.168.168.9 00-02-e3-16-e4-5d dynamic
192.168.168.10 00-50-04-17-66-90 dynamic
192.168.168.22 00-60-08-39-e5-a1 dynamic
192.168.168.254 00-40-10-18-42-49 dynamic
C:\>
```


TIP

ARP is sometimes useful when diagnosing duplicate IP assignment problems. For example, suppose you can't access a computer that has an IP address of 192.168.1.100. You try to ping the computer, expecting the ping to fail, but lo and behold — the ping succeeds. One possible cause for this may be that two computers on the network have been assigned the address 192.168.1.100, and your ARP cache is pointing to the wrong one. The way to find out is to go to the 192.168.1.100 computer that you want to access, run `ipconfig /all`, and make a note of the physical address. Then return to the computer that's having trouble reaching the 192.168.1.100 computer, run `arp -a`, and compare the physical address with the one you noted. If they're different, two computers are assigned the same IP address. You can then check the Dynamic Host Configuration Protocol (DHCP) or static TCP/IP configuration of the computers involved to find out why.

Using the hostname Command

The `hostname` command is the simplest of all the TCP/IP commands presented in this chapter. It simply displays the computer's host name. For example:

```
C:\>hostname
doug
C:\>
```

Here, the host name for the computer is `doug`. The Windows version of the `hostname` command has no parameters. However, the Unix/Linux versions of `hostname` let you set the computer's host name as well as display it. You do that by specifying the new host name as an argument.

Using the ipconfig Command

Using the `ipconfig` command displays information about a computer's TCP/IP configuration. It can also be used to update DHCP and Domain Name System (DNS) settings.

Displaying basic IP configuration

To display the basic IP configuration for a computer, use the `ipconfig` command without any parameters, like this:

```
C:\>ipconfig

Windows IP Configuration

Ethernet adapter Ethernet:

Connection-specific DNS Suffix  .  :
IPv6 Address. . . . . . . . . . . : 2600:6c52:7900:1623:9809:47db:8e66:b0c
Link-local IPv6 Address . . . . . : fe80::9809:47db:8e66:b0c%2
IPv4 Address. . . . . . . . . . . : 192.168.1.15
Subnet Mask . . . . . . . . . . . : 255.255.255.0
Default Gateway . . . . . . . . . : fe80::6238:e0ff:fea2:4f06%2
                                    192.168.1.1
```

When you use `ipconfig` without parameters, the command displays the name of the adapter, the domain name used for the adapter, the IP address, the subnet mask, and the default gateway configuration for the adapter. This is the easiest way to determine a computer's IP address. In this example, the IP address is 192.168.1.15.


TIP

If your computer indicates an IP address in the 169.254.*x.x* block, odds are good that the DHCP server isn't working. 169.254.*x.x* is the Class B address block that Windows uses when it resorts to IP autoconfiguration. This usually happens only when the DHCP server can't be reached or isn't working.

Displaying detailed configuration information

You can display detailed IP configuration information by using an `/all` switch with the `ipconfig` command, like this:

```
C:\>ipconfig /all

Windows IP Configuration

Host Name . . . . . . . . . . . . : WIN1901
Primary Dns Suffix . . . . . . . :
Node Type . . . . . . . . . . . . : Hybrid
IP Routing Enabled. . . . . . . . : No
WINS Proxy Enabled. . . . . . . . : No

Ethernet adapter Ethernet:

Connection-specific DNS Suffix .  :
Description . . . . . . . . . . . : Microsoft Hyper-V Network Adapter
Physical Address. . . . . . . . . : 00-15-5D-00-F5-01
DHCP Enabled. . . . . . . . . . . : No
Autoconfiguration Enabled . . . . : Yes
IPv6 Address. . . . . . . . . . . : 2600:6c52:7900:1623:9809:47db:8e66:b0c(Preferred)
Link-local IPv6 Address . . . . . : fe80::9809:47db:8e66:b0c%2(Preferred)
IPv4 Address. . . . . . . . . . . : 192.168.1.15m  (Preferred)
Subnet Mask . . . . . . . . . . . : 255.255.255.0
Default Gateway . . . . . . . . . : fe80::6238:e0ff:fea2:4f06%2
                                    192.168.1.1
DHCPv6 IAID . . . . . . . . . . . : 33559901
DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-26-DD-9F-1D-00-15-5D-00-F5-01
DNS Servers . . . . . . . . . . . : 192.168.1.10
                                    192.168.1.11
NetBIOS over Tcpip. . . . . . . . : Enabled
C:\>
```

You can determine a lot of information about the computer from the `ipconfig /all` command. For example:

- The computer's host name is WIN1901.
- The computer's IPv4 address is 192.168.1.15, and the subnet mask is 255.255.255.0.
- The default gateway is a router located at 192.168.1.1.
- The DNS servers are at 192.168.1.10 and 192.168.1.11.

Renewing an IP lease

If you're having an IP configuration problem, you can often solve it by renewing the computer's IP lease. To do that, use a `/renew` switch, like this:

```
C:\>ipconfig /renew
Windows IP Configuration
Ethernet adapter Local Area Connection:
Connection-specific DNS Suffix . :
IP Address. . . . . . . . . . : 192.168.1.100
Subnet Mask . . . . . . . . . : 255.255.255.0
Default Gateway . . . . . . . : 192.168.1.1
```

When you renew an IP lease, the `ipconfig` command displays the new lease information.



WARNING

This command won't work if you configured the computer to use a static IP address.

Releasing an IP lease

You can release an IP lease by using an `ipconfig` command with the `/release` parameter, like this:

```
C:\>ipconfig /release
Windows IP Configuration
Ethernet adapter Local Area Connection:
Connection-specific DNS Suffix . :
IP Address . . . . . . . . . . . : 0.0.0.0
Subnet Mask. . . . . . . . . . . : 0.0.0.0
Default Gateway. . . . . . . . . :
```

As you can see, the DNS suffix and default gateway for the computer are blank, and the IP address and subnet mask are set to 0.0.0.0.

After you release the DHCP lease, you can use an `ipconfig /renew` command to obtain a new DHCP lease for the computer.



WARNING

This command won't work if you configured the computer to use a static IP address.

Flushing the local DNS cache

You probably won't need to do this unless you're having DNS troubles. If you've been tinkering with your network's DNS configuration, though, you may need to flush the cache on your DNS clients so that they'll be forced to reacquire information from the DNS server. You can do that by using a `/flushdns` switch:

```
C:\>ipconfig /flushdns
Windows IP Configuration
Successfully flushed the DNS Resolver Cache.
C:\>
```


TIP

Even if you don't need to do this, it's fun just to see the computer read `flushed`. If I worked at Microsoft, you'd be able to revert Windows Vista computers back to XP by using a `/flushVista` switch.

Using the nbtstat Command

`nbtstat` is a Windows-only command that can help solve problems with NetBIOS name resolution. (*nbt* stands for *NetBIOS over TCP/IP.*) You can use any of the switches listed in Table 7-1 to specify what `nbtstat` output you want to display. For example, you can use an `-a` switch to display the cached name table for a specified computer, like this:

```
C:\>nbtstat -a WK07-001
Local Area Connection:
Node IpAddress: [192.168.1.110] Scope Id: []
NetBIOS Remote Machine Name Table
Name Type Status
---------------------------------------------
WK07-001 <00> UNIQUE Registered
WORKGROUP <00> GROUP Registered
WK07-001 <20> UNIQUE Registered
WORKGROUP <1E> GROUP Registered
WORKGROUP <1D> UNIQUE Registered
..__MSBROWSE__.<01> GROUP Registered
MAC Address = 00-12-3F-A7-17-BAC:\>
C:\>
```

Table 7-1 lists the switches that you can use with `nbtstat` and explains the function of each switch.

TABLE 7-1 `nbtstat` Command Switches

Switch	What It Does
`-a name`	Lists the specified computer's name table given the computer's name
`-A IP-address`	Lists the specified computer's name table given the computer's IP address
`-c`	Lists the contents of the NetBIOS cache
`-n`	Lists locally registered NetBIOS names
`-r`	Displays a count of the names resolved by broadcast and via WINS
`-R`	Purges and reloads the cached name table from the LMHOSTS file
`-RR`	Releases and then reregisters all names
`-S`	Displays the sessions table using IP addresses
`-s`	Displays the sessions table and converts destination IP addresses to computer NetBIOS names

Using the netstat Command

Using the `netstat` command displays a variety of statistics about a computer's active TCP/IP connections. It's a useful tool to use when you're having trouble with TCP/IP applications, such as File Transfer Protocol (FTP), HyperText Transport Protocol (HTTP), and so on.

Displaying connections

If you run `netstat` without specifying any parameters, you get a list of active connections, something like this:

```
C:\>netstat
Active Connections
Proto Local Address Foreign Address State
TCP Doug:1463 192.168.168.10:1053 ESTABLISHED
TCP Doug:1582 192.168.168.9:netbios-ssn ESTABLISHED
TCP Doug:3630 192.168.168.30:9100 SYN_SENT
TCP Doug:3716 192.168.168.10:4678 ESTABLISHED
TCP Doug:3940 192.168.168.10:netbios-ssn ESTABLISHED
C:\>
```

This list shows all the active connections on the computer and indicates the local port used by the connection, as well as the IP address and port number for the remote computer.

You can specify the `-n` switch to display both local and foreign addresses in numeric IP form:

```
C:\>netstat -n
Active Connections
Proto Local Address Foreign Address State
TCP 192.168.168.21:1463 192.168.168.10:1053 ESTABLISHED
TCP 192.168.168.21:1582 192.168.168.9:139 ESTABLISHED
TCP 192.168.168.21:3658 192.168.168.30:9100 SYN_SENT
TCP 192.168.168.21:3716 192.168.168.10:4678 ESTABLISHED
TCP 192.168.168.21:3904 207.46.106.78:1863 ESTABLISHED
TCP 192.168.168.21:3940 192.168.168.10:139 ESTABLISHED
C:\>
```

Finally, you can specify the `-a` switch to display all TCP/IP connections and ports that are being listened to. I won't list the output from that command here because it would run several pages, and I want to do my part for the rainforests. Suffice it to say that it looks a lot like the `netstat` output shown previously, but a lot longer.

Displaying interface statistics

If you use an `-e` switch, `netstat` displays various protocol statistics, like this:

```
C:\>netstat -e
Interface Statistics
Received Sent
Bytes 672932849 417963911
Unicast packets 1981755 1972374
Non-unicast packets 251869 34585
Discards 0 0
Errors 0 0
Unknown protocols 1829
C:\>
```


REMEMBER

The items to pay attention to in this output are the Discards and Errors. These numbers should be zero, or at least close to it. If they're not, the network may be carrying too much traffic or the connection may have a physical problem. If no physical problem exists with the connection, try segmenting the network to see whether the error and discard rates drop.

Using the nslookup Command

The nslookup command is a powerful tool for diagnosing DNS problems. You know you're experiencing a DNS problem when you can access a resource by specifying its IP address but not its DNS name. For example, if you can get to www.ebay.com by typing **173.223.234.210** in your browser's address bar but not by typing www.ebay.com, you have a DNS problem.

Looking up an IP address

The simplest use of nslookup is to look up the IP address for a given DNS name. For example, how did I know that 173.223.234.210 was the IP address for www.ebay.com? I used nslookup to find out:

```
C:\>nslookup ebay.com
Server:  cdns01.comcast.net
Address:  2001:558:feed::1

Non-authoritative answer:
Name:    ebay.com
Addresses:  173.223.234.210
          173.223.234.197
C:\>
```

As you can see, just type **nslookup** followed by the DNS name you want to look up, and nslookup issues a DNS query to find out. This DNS query was sent to the server named cdms-1/cp,cast/met. It then displayed two IP addresses associated with ebay.com: namely, 173.222.234.210 and 173.223.234.197.


TIP

In some cases, you may find that using an nslookup command gives you the wrong IP address for a host name. To know that for sure, of course, you have to know with certainty what the host IP address *should* be. For example, if you know that your server is 203.172.182.10 but nslookup returns a completely different IP address for your server when you query the server's host name, something is probably wrong with one of the DNS records.

Using nslookup subcommands

If you use nslookup without any arguments, the nslookup command enters a subcommand mode. It displays a prompt character (>) to let you know that you're in nslookup subcommand mode rather than at a normal Windows command prompt. In subcommand mode, you can enter various subcommands to set options or to perform queries. You can type a question mark (?) to get a list of these commands. Table 7-2 lists the subcommands you'll use most.

GET ME OUT OF HERE!

One of my pet peeves is that it seems as if every program that uses subcommands chooses a different command to quit the application. I can never remember whether the command to get out of `nslookup` is `quit`, `bye`, or `exit`. I usually end up trying them all. And no matter what program I'm using, I always seem to choose the one that works for some other program first. When I'm in `nslookup`, I use `bye` first. When I'm in FTP, I try `exit` first. Arghh! If I were King of the Computer Hill, every program that had subcommands would respond to any of the following commands by exiting the program and returning to a command prompt:

Quit	Sayonara
Exit	Ciao
Bye	Mañana
Leave	Makelikeatree

Of course, the final command to try would be `Andgetouttahere` (in honor of Biff from the *Back to the Future* movies).

TABLE 7-2 The Most Commonly Used `nslookup` Subcommands

Subcommand	What It Does
`name`	Queries the current name server for the specified name.
`server name`	Sets the current name server to the server you specify.
`root`	Sets the root server as the current server.
`set type=x`	Specifies the type of records to be displayed, such as A, CNAME, MX, NS, PTR, or SOA. Specify ANY to display all records.
`set debug`	Turns on Debug mode, which displays detailed information about each query.
`set nodebug`	Turns off Debug mode.
`set recurse`	Enables recursive searches.
`set norecurse`	Disables recursive searches.
`exit`	Exits and returns you to a command prompt.

Displaying DNS records

One of the main uses of `nslookup` is to examine your DNS configuration to make sure that it's set up properly. To do that, follow these steps:

1. **At a command prompt, type** nslookup **without any parameters.**

   `nslookup` displays the name of the default name server and displays the > prompt.

```
C:\>nslookup
Default Server: ns1.orng.twtelecom.net
Address: 168.215.210.50
>
```

2. **Type the subcommand** set type=any**.**

   `nslookup` silently obeys your command and displays another prompt:

```
> set type=any
>
```

3. **Type your domain name.**

   `nslookup` responds by displaying the name servers for your domain:

```
> lowewriter.com
Server: ns1.orng.twtelecom.net
Address: 168.215.210.50
Non-authoritative answer:
lowewriter.com nameserver = NS000.NS0.com
lowewriter.com nameserver = NS207.PAIR.com
lowewriter.com nameserver = NS000.NS0.com
lowewriter.com nameserver = NS207.PAIR.com
>
```

4. **Use a** `server` **command to switch to one of the domain's name servers.**

   For example, to switch to the first name server listed in Step 3, type **server `NS000.NS0.com`**. `nslookup` replies with a message that indicates the new default server:

```
> server ns000.ns0.com
Default Server: ns000.ns0.com
Address: 216.92.61.61
>
```

5. **Type your domain name again.**

   This time, `nslookup` responds by displaying the DNS information for your domain:

```
> lowewriter.com
Server: ns000.ns0.com
Address: 216.92.61.61
lowewriter.com
primary name server = ns207.pair.com
responsible mail addr = root.pair.com
serial = 2001121009
refresh = 3600 (1 hour)
retry = 300 (5 mins)
expire = 604800 (7 days)
default TTL = 3600 (1 hour)
lowewriter.com nameserver = ns000.ns0.com
lowewriter.com nameserver = ns207.pair.com
lowewriter.com MX preference = 50, mail exchanger = sasi.
   pair.com
lowewriter.com internet address = 209.68.34.15
>
```

6. **Type** exit **to leave the** `nslookup` **program.**

   You return to a command prompt.

```
> exit
C:\>
```

Wasn't that fun?

Locating the mail server for an email address

If you're having trouble delivering mail to someone, you can use `nslookup` to determine the IP address of the user's mail server. Then, you can use the `ping` command to see whether you can contact the user's mail server. If not, you can use the `tracert` command to find out where the communication breaks down. (See "Using the `tracert` Command" later in this chapter for more information.)

To find a user's mail server, start `nslookup` and enter the command **set type=MX**. Then, enter the domain portion of the user's email address. For example, if the user's address is `Doug@LoweWriter.com`, enter **`LoweWriter.com`**. `nslookup` will display the MX (Mail Exchange) information for the domain, like this:

```
C:\>nslookup
Default Server: ns7.attbi.com
Address: 204.127.198.19
> set type=mx
> lowewriter.com
Server: ns7.attbi.com
Address: 204.127.198.19
lowewriter.com MX preference = 50, mail exchanger =
   sasi.pair.com
lowewriter.com nameserver = ns000.ns0.com
lowewriter.com nameserver = ns207.pair.com
ns000.ns0.com internet address = 216.92.61.61
ns207.pair.com internet address = 209.68.2.52
>
```

Here, you can see that the name of the mail server for the `LoweWriter.com` domain is `sasi.pair.com`.

Using the pathping Command

`pathping` is an interesting command that's unique to Windows. It's sort of a cross between the `ping` command and the `tracert` command, combining the features of both into one tool. When you run `pathping`, it first traces the route to the destination address much the way `tracert` does. Then, it launches into a 25-second test of each router along the way, gathering statistics on the rate of data loss to each hop. If the route has a lot of hops, this can take a long time. However, it can help you to spot potentially unreliable hops. If you're having intermittent trouble reaching a particular destination, using `pathping` may help you pinpoint the problem.

The following command output is typical of the `pathping` command. Using an `–n` switch causes the display to use numeric IP numbers only, instead of DNS host names. Although fully qualified host names are convenient, they tend to be very long for network routers, which makes the `pathping` output very difficult to decipher.

```
C:\>pathping –n www.lowewriter.com
Tracing route to lowewriter.com [209.68.34.15]
over a maximum of 30 hops:
0 192.168.168.21
1 66.193.195.81
2 66.193.200.5
```

```
3 168.215.55.173
4 168.215.55.101
5 168.215.55.77
6 66.192.250.38
7 66.192.252.22
8 208.51.224.141
9 206.132.111.118
10 206.132.111.162
11 64.214.174.178
12 192.168.1.191
13 209.68.34.15
Computing statistics for 325 seconds...
Source to Here This Node/Link
Hop RTT Lost/Sent = Pct Lost/Sent = Pct Address
0 192.168.168.21
0/ 100 = 0% |
1 1ms 0/ 100 = 0% 0/ 100 = 0% 66.193.195.81]
0/ 100 = 0% |
2 14ms 0/ 100 = 0% 0/ 100 = 0% 66.193.200.5
0/ 100 = 0% |
3 10ms 0/ 100 = 0% 0/ 100 = 0% 168.215.55.173
0/ 100 = 0% |
4 10ms 0/ 100 = 0% 0/ 100 = 0% 168.215.55.101
0/ 100 = 0% |
5 12ms 0/ 100 = 0% 0/ 100 = 0% 168.215.55.77
0/ 100 = 0% |
6 14ms 0/ 100 = 0% 0/ 100 = 0% 66.192.250.38
0/ 100 = 0% |
7 14ms 0/ 100 = 0% 0/ 100 = 0% 66.192.252.22
0/ 100 = 0% |
8 14ms 0/ 100 = 0% 0/ 100 = 0% 208.51.224.141
0/ 100 = 0% |
9 81ms 0/ 100 = 0% 0/ 100 = 0% 206.132.111.118
0/ 100 = 0% |
10 81ms 0/ 100 = 0% 0/ 100 = 0% 206.132.111.162]
0/ 100 = 0% |
11 84ms 0/ 100 = 0% 0/ 100 = 0% 64.214.174.178]
0/ 100 = 0% |
12 --- 100/ 100 =100% 100/ 100 =100% 192.168.1.191
0/ 100 = 0% |
13 85ms 0/ 100 = 0% 0/ 100 = 0% 209.68.34.15
Trace complete.
```

Using the ping Command

`ping` is probably the most basic TCP/IP command line tool. Its main purpose is to determine whether you can reach another computer from your computer. It uses Internet Control Message Protocol (ICMP) to send mandatory `ECHO_REQUEST` datagrams to the specified host computer. When the reply is received back from the host, the `ping` command displays how long it took to receive the response.

You can specify the host to ping by using an IP address, as in this example:

```
C:\>ping 192.168.168.10
Pinging 192.168.168.10 with 32 bytes of data:
Reply from 192.168.168.10: bytes=32 time<1ms TTL=128
Reply from 192.168.168.10: bytes=32 time<1ms TTL=128
Reply from 192.168.168.10: bytes=32 time<1ms TTL=128
Reply from 192.168.168.10: bytes=32 time<1ms TTL=128
Ping statistics for 192.168.168.10:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\>
```

By default, the `ping` command sends four packets to the specified host. It displays the result of each packet sent. Then it displays summary statistics: how many packets were sent, how many replies were received, the error loss rate, and the approximate round-trip time.

You can also ping by using a DNS name, as in this example:

```
C:\>ping www.lowewriter.com
Pinging lowewriter.com [209.68.34.15] with 32 bytes of data:
Reply from 209.68.34.15: bytes=32 time=84ms TTL=53
Reply from 209.68.34.15: bytes=32 time=84ms TTL=53
Reply from 209.68.34.15: bytes=32 time=84ms TTL=53
Reply from 209.68.34.15: bytes=32 time=84ms TTL=53
Ping statistics for 209.68.34.15:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 84ms, Maximum = 84ms, Average = 84ms
C:\>
```

The `ping` command uses a DNS query to determine the IP address for the specified host, and then pings the host based on its IP address.

The ping command has a number of other switches that you'll use rarely, if ever. Some of these switches are available only for some operating systems. To find out which switches are available for your version of Ping, type **ping /?** (Windows) or **man ping** (Unix/Linux).



You can find a very interesting story about the creation of the ping command written by the command's author, Mike Muus, at his website at https://ftp.arl.army.mil/~mike/ping.html. (Sadly, Mr. Muus was killed in an automobile accident in November 2000.)

Using the route Command

Using the route command displays or modifies the computer's routing table. For a typical computer that has a single network interface and is connected to a local area network (LAN) that has a router, the routing table is pretty simple and isn't often the source of network problems. Still, if you're having trouble accessing other computers or other networks, you can use the route command to make sure that a bad entry in the computer's routing table isn't the culprit.

For a computer with more than one interface and that's configured to work as a router, the routing table is often a major source of trouble. Setting up the routing table properly is a key part of configuring a router to work.

Displaying the routing table

To display the routing table (both IPv4 and IPv6) in Windows, use the route print command. In Unix/Linux, you can just use route without any command line switches. The output displayed by the Windows and Unix/Linux commands are similar. Here's an example from a typical Windows client computer:

```
C:\>route print
===========================================================================
Interface List
8....00 12 3f a7 17 ba........ Intel(R) PRO/100 VE Network Connection
1........................... Software Loopback Interface 1
9......02 00 54 55 4e 01...... Teredo Tunneling Pseudo-Interface
10 ....00 00 00 00 00 00 00 e0 isatap.{D0F85930-01E2-402F-B0FC-31DFF887F06F}
===========================================================================
IPv4 Route Table
===========================================================================
```

```
Active Routes:
Network Destination Netmask Gateway Interface Metric
0.0.0.0 0.0.0.0 192.168.1.1 192.168.1.110 276
127.0.0.0 255.0.0.0 On-link 127.0.0.1 306
127.0.0.1 255.255.255.255 On-link 127.0.0.1 306
127.255.255.255 255.255.255.255 On-link 127.0.0.1 306
192.168.1.0 255.255.255.0 On-link 192.168.1.110 276
192.168.1.110 255.255.255.255 On-link 192.168.1.110 276
192.168.1.255 255.255.255.255 On-link 192.168.1.110 276
224.0.0.0 240.0.0.0 On-link 127.0.0.1 306
224.0.0.0 240.0.0.0 On-link 192.168.1.110 276
255.255.255.255 255.255.255.255 On-link 127.0.0.1 306
255.255.255.255 255.255.255.255 On-link 192.168.1.110 276
===========================================================================
Persistent Routes:
Network Address Netmask Gateway Address Metric
0.0.0.0 0.0.0.0 192.168.1.1 Default
===========================================================================
IPv6 Route Table
===========================================================================
Active Routes:
If Metric Network Destination Gateway
9 18 ::/0 On-link
1 306 ::1/128 On-link
9 18 2001::/32 On-link
9 266 2001:0:4136:e38c:2c6c:670:3f57:fe91/128
On-link
8 276 fe80::/64 On-link
9 266 fe80::/64 On-link
10 281 fe80::5efe:192.168.1.110/128
On-link
8 276 fe80::cca:9067:9427:a911/128
On-link
9 266 fe80::2c6c:670:3f57:fe91/128
On-link
1 306 ff00::/8 On-link
9 266 ff00::/8 On-link
8 276 ff00::/8 On-link
===========================================================================
Persistent Routes:
None
C:\>
```

For each entry in the routing table, five items of information are listed:

- **The destination IP address**

  Actually, this is the address of the destination subnet, and must be interpreted in the context of the subnet mask.

- **The subnet mask that must be applied to the destination address to determine the destination subnet**
- **The IP address of the gateway to which traffic intended for the destination subnet will be sent**
- **The IP address of the interface through which the traffic will be sent to the destination subnet**
- **The *metric,* which indicates the number of hops required to reach destinations via the gateway**

Each packet that's processed by the computer is evaluated against the rules in the routing table. If the packet's destination address matches the destination subnet for the rule, the packet is sent to the specified gateway via the specified network interface. If not, the next rule is applied.

The computer on which I ran the `route` command in this example is on a private 192.168.1.0 subnet. The computer's IP address is 192.168.1.100, and the default gateway is a router at 192.168.1.1.

Here's how the rules shown in this example are used. Notice that you have to read the entries from the bottom up:

- **The first rule is for packets sent to 255.255.255.255, with subnet mask 255.255.255.255.** This special IP address is for broadcast packets. The rule specifies that these broadcast packets should be delivered to the local network interface (192.168.1.100).
- **The next rule is for packets sent to 192.168.1.255, again with subnet mask 255.255.255.255.** These are also broadcast packets and are sent to the local network interface.
- **The next rule is for packets sent to 192.168.1.100, again with subnet mask 255.255.255.255.** This is for packets that the computer is sending to itself via its own IP address. This rule specifies that these packets will be sent to the local loopback interface on 127.0.0.1.
- **The next rule is for packets sent to 192.168.1.0, with subnet mask 255.255.255.0.** These are packets intended for the local subnet. They're sent to the subnet via the local interface at 192.169.1.100.

- **The next rule is for packets sent to the loopback address (127.0.0.1, subnet mask 255.0.0.0).** These packets are sent straight through to the loopback interface, 127.0.0.1.
- **The last rule is for everything else.** All IP addresses will match the destination IP address 0.0.0.0 with subnet mask 0.0.0.0 and will be sent to the default gateway router at 192.168.1.1 via the computer's network interface at 192.168.1.100.


TIP

One major difference between the Windows version of `route` and the Unix/Linux version is the order in which they list the routing table. The Windows `route` command lists the table starting with the most general entry and works toward the most specific. The Unix/Linux version is the other way around: It starts with the most specific and works toward the more general. The Unix/Linux order makes more sense — the Windows `route` command displays the routing list upside down.

Modifying the routing table

Besides displaying the routing table, the `route` command also lets you modify it by adding, deleting, or changing entries.


WARNING

Don't try this unless you know what you're doing. If you mess up the routing table, your computer may not be able to communicate with anyone.

The syntax for the `route` command for adding, deleting, or changing a route entry is

```
route [-p] command dest [mask subnet] gateway [-if interface]
```

The following list describes each of the `route` command's parameters:

- `-p`: Makes the entry persistent. If you omit `-p`, the entry will be deleted the next time you reboot. (Use this only with `add` commands.)
- *`command`*: Add, delete, or change.
- *`dest`*: The IP address of the destination subnet.
- `mask` *`subnet`*: The subnet mask. If you omit the subnet mask, the default is 255.255.255.255, meaning that the entry will apply only to a single host rather than a subnet. You usually want to include the mask.
- *`gateway`*: The IP address of the gateway to which packets will be sent.
- `if` *`interface`*: The IP address of the interface through which packets will be sent. If your computer has only one network interface, you can omit this.

Suppose that your network has a second router that serves as a link to another private subnet, 192.168.2.0 (subnet mask 255.255.255.0). The interface on the local side of this router is at 192.168.1.200. To add a static route entry that sends packets intended for the 192.168.2.0 subnet to this router, use a command like this:

```
C:\>route –p add 192.168.2.0 mask 255.255.255.0 192.168.1.200
```

Now, suppose that you later change the IP address of the router to 192.168.1.222. You can update this route with the following command:

```
C:\>route change 192.168.2.0 mask 255.255.255.0 192.168.1.222
```

Notice that I specify the mask again. If you omit the mask from a `route change` command, the command changes the mask to 255.255.255.255!

Finally, suppose that you realize that setting up a second router on this network wasn't such a good idea after all, so you want to just delete the entry. The following command will do the trick:

```
C:\>route delete 192.168.2.0
```

Using the tracert Command

The `tracert` command (`traceroute` in Unix/Linux implementations) is one of the key diagnostic tools for TCP/IP. It displays a list of all the routers that a packet must go through to get from the computer where `tracert` is run to any other computer on the internet. Each one of these routers is called a *hop*, presumably because the original designers of the IP protocol played a lot of hopscotch when they were young. If you can't connect to another computer, you can use `tracert` to find out exactly where the problem is occurring.

`tracert` makes three attempts to contact the router at each hop and displays the response time for each of these attempts. Then, it displays the DNS name of the router (if available) and the router's IP address.

To use `tracert`, type the `tracert` command followed by the host name of the host to which you want to trace the route. For example, here's how you can use `tracert` to find out how many routers are between you and the whitehouse's servers:

```
C:\>tracert whitehouse.gov

Tracing route to whitehouse.gov [2a04:fa87:fffd::c000:42a8]
```

```
over a maximum of 30 hops:

  1     1 ms     1 ms     1 ms  2601:204:380:11f0:c294:35ff:fe2a:689e
  2    12 ms    13 ms    10 ms  2001:558:4010:33::1
  3    11 ms    10 ms    11 ms  po-254-1221-rur01.pinedale.ca.ccal.comcast.net
   [2001:558:212:6801::1]
  4    36 ms    19 ms    16 ms  ae-38-ar01.sacramento.ca.ccal.comcast.net
   [2001:558:210:f4::1]
  5    17 ms    16 ms    37 ms  be-36441-cs04.sunnyvale.ca.ibone.comcast.net
   [2001:558:3:25b::1]
  6    19 ms    16 ms    18 ms  be-1412-cr12.sunnyvale.ca.ibone.comcast.net
   [2001:558:3:38a::2]
  7    19 ms    20 ms    18 ms  be-304-cr12.9greatoaks.ca.ibone.comcast.net
   [2001:558:3:16d::2]
  8    18 ms    27 ms    15 ms  be-1112-cs01.9greatoaks.ca.ibone.comcast.net
   [2001:558:3:421::1]
  9    18 ms    17 ms    17 ms  be-2101-pe01.9greatoaks.ca.ibone.comcast.net
   [2001:558:3:136::2]
 10    15 ms    17 ms    17 ms  2001:559::f16
 11    16 ms    16 ms    18 ms  2a04:fa87:fffd::c000:42a8

Trace complete.
```

As you can see, getting from my computer to `whitehouse.gov` required 11 hops.

Note that some of the hops in the previous output list IP6 addresses. You can force `tracert` to show you IP4 addresses by using the `-4` flag, like this:

```
C:\>tracert -4 whitehouse.gov

Tracing route to whitehouse.gov [192.0.66.168]
over a maximum of 30 hops:

  1     1 ms     1 ms     1 ms  10.0.0.1
  2    13 ms    10 ms    14 ms  96.120.86.213
  3    10 ms     9 ms    10 ms  po-254-1221-rur01.pinedale.ca.ccal.comcast.net
   [96.110.146.25]
  4    15 ms    23 ms    22 ms  ae-38-ar01.sacramento.ca.ccal.comcast.net
   [68.87.221.57]
  5    19 ms    16 ms    17 ms  be-36441-cs04.sunnyvale.ca.ibone.comcast.net
   [96.110.41.109]
  6    16 ms    17 ms    16 ms  be-1412-cr12.sunnyvale.ca.ibone.comcast.net
   [96.110.46.42]
  7    17 ms    17 ms    15 ms  be-303-cr12.9greatoaks.ca.ibone.comcast.net
   [96.110.37.178]
  8    21 ms    17 ms    16 ms  be-1112-cs01.9greatoaks.ca.ibone.comcast.net
   [68.86.166.133]
```

```
  9    18 ms    16 ms    17 ms  be-2101-pe01.9greatoaks.ca.ibone.comcast.net
   [96.110.36.218]
 10    19 ms    42 ms    17 ms  50.248.116.14
 11    15 ms    16 ms    17 ms  192.0.66.168

Trace complete.
```

The most likely problem that you'll encounter when you use `tracert` is a timeout during one of the hops. Timeouts are indicated by asterisks where you'd expect to see a time. For example, the following `tracert` output shows the fourth hop timing out on all three attempts:

```
C:\>tracert -4 whitehouse.gov

Tracing route to whitehouse.gov [192.0.66.168]
over a maximum of 30 hops:

  1     1 ms     1 ms     1 ms  10.0.0.1
  2    13 ms    10 ms    14 ms  96.120.86.213
  3    10 ms     9 ms    10 ms  po-254-1221-rur01.pinedale.ca.ccal.comcast.net
   [96.110.146.25]
  4     *        *        * Request timed out.
```

Sometimes, timeouts are caused by temporary problems, so you should try the `tracert` again to see if the problem persists. If you keep getting timeouts at the same router, the router could be having a genuine problem.

UNDERSTANDING HOW TRACERT WORKS

Understanding how `tracert` works can provide some insight that may help you to interpret the results it provides. Plus, you can use this knowledge to impress your friends, who probably don't know how it works.

The key to `tracert` is a field that's a standard part of all IP packets called TTL, which stands for *Time to Live.* In most other circumstances, a value called TTL would be a time value — not in IP packets, however. In an IP packet, the TTL value indicates how many routers a packet can travel through on its way to its destination. Every time a router forwards an IP packet, it subtracts one from the packet's TTL value. When the TTL value reaches zero, the router refuses to forward the packet.

The `tracert` command sends a series of special messages called ICMP Echo Requests to the destination computer. The first time it sends this message, it sets the TTL value of the packet to 1. When the packet arrives at the first router along the path to the destination, that router subtracts one from the TTL value, sees that the TTL value has become 0, so it sends a Time Exceeded message back to the original host. When the `tracert` command receives this Time Exceeded message, it extracts the IP address of the router from it, calculates the time it took for the message to return, and displays the first hop.

Then the `tracert` command sends another Echo Request message: this time, with the TTL value set to 2. This message goes through the first router to the second router, which sees that the TTL value has been decremented to 0 and then sends back a Time Exceeded message. When `tracert` receives the Time Exceeded message from the second router, it displays the line for the second hop. This process continues, each time with a greater TTL value, until the Echo Request finally reaches the destination.

Pretty clever, eh?

(Note that the Unix/Linux `traceroute` command uses a slightly different set of TCP/IP messages and responses to accomplish the same result.)

3 Planning a Network

Contents at a Glance

CHAPTER 1: **Local Area Networks** 233

CHAPTER 2: **Wide Area Networks** 249

CHAPTER 3: **Server Architecture** 261

CHAPTER 4: **Virtualization Architecture** 271

CHAPTER 5: **Storage Architecture** 283

CHAPTER 6: **Backup Architecture** 295

CHAPTER 7: **Hyperconverged Infrastructure** 313

IN THIS CHAPTER

- Making a network plan
- Taking stock of your computer stock
- Making sure that you know why you need a network
- Making the basic network decisions that you can't avoid
- Looking at additional topics your plan should address

Chapter 1
Local Area Networks

Okay, so you're convinced that you need to network your computers. What now? Do you stop by Computers-R-Us on the way to work, install the network before drinking your morning coffee, and expect the network to be fully operational by noon?

I don't think so.

Networking your computers is just like any other worthwhile endeavor: Doing it right requires a bit of planning. This chapter helps you to think through your network before you start spending money. It shows you how to come up with a networking plan that's every bit as good as the plan that a network consultant would charge thousands of dollars for. See? This book is already saving you money!

Making a Network Plan

Before you begin any networking project, whether a new network installation or an upgrade of an existing network, make a detailed plan *first.* If you make technical decisions too quickly, before studying all the issues that affect the project, you'll regret it. You'll discover too late that a key application won't run over the network, the network has unacceptably slow performance, or key components of the network don't work together.

Here are some general thoughts to keep in mind while you create your network plan:

- **Don't rush the plan.** The most costly networking mistakes are the ones that you make before you install the network. Think things through and consider alternatives.
- **Write down the network plan.** The plan doesn't have to be a fancy, 500-page document. If you want to make it look good, pick up a ½-inch three-ring binder, which is big enough to hold your network plan with plenty of room to spare.
- **Ask someone else to read your network plan before you buy anything.** Preferably, ask someone who knows more about computers than you do.
- **Keep the plan up to date.** If you add to the network, dig up the plan, dust it off, and update it.


TIP

"The best laid schemes of mice and men gang aft agley, and leave us naught but grief and pain for promised joy." Robert Burns lived a few hundred years before computer networks, but his famous words ring true. A network plan isn't chiseled in stone. If you discover that something doesn't work the way you thought it would, that's okay. Just change your plan.

Being Purposeful

One of the first steps in planning your network is making sure that you understand why you want the network in the first place. Here are some of the more common reasons for needing a network, all of them quite valid:

- **My co-worker and I exchange files using flash drives just about every day.** With a network, trading files is easier.
- **I don't want to buy everyone a color laser printer when I know the one we have now just sits there taking up space most of the day.** So wouldn't buying a network be better than buying a color laser printer for every computer?
- **I want everyone to be able to access the internet.** Many networks, especially smaller ones, exist solely for the purpose of sharing an internet connection.
- **Business is so good that one person typing in orders eight hours each day can't keep up.** With a network, more than one person can enter orders, which expedites orders and possibly saves on overtime expenses.

- **My brother-in-law just put in a network at his office.** No one wants to be behind the times.
- **I already have a network, but it's so old that it may as well be made of kite string and tin cans.** An improved network speeds up access to shared files, provides better security, is easier to manage, and is more reliable.


REMEMBER

After you identify all the reasons why you think you need a network, write them down. Don't worry about winning the Pulitzer Prize for your stunning prose. Just make sure that you write down what you expect a network to do for you. If you were making a 500-page networking proposal, you'd place the description of why a network is needed in a tabbed section labeled Justification. In your ½-inch network binder, file the description under Purpose.


TIP

As you consider the reasons why you need a network, you may conclude that you don't need a network after all. That's okay. You can always use the binder for your stamp collection.

Taking Stock

One of the initial challenges of planning a network is figuring out how to work with the computers that you already have. In other words, how do you get from here to there? Before you can plan how to get "there," you have to know where "here" is. In other words, you have to take a thorough inventory of your current computers.

What you need to know

You need to know the following information about each of your computers:

- **The processor type and, if possible, its clock speed:** It would be nice if each of your computers had a shiny new i7 12-Core processor. In most cases, though, you find a mixture of computers: some new, some old, some borrowed, some blue.

  You can't usually tell what kind of processor a computer has just by looking at the computer's case. But you can easily find out by right-clicking Computer on the Start menu and choosing Properties.


TIP

- **The size of the hard drive and the arrangement of its partitions:** To find the size of your computer's hard drive, open the Computer window, right-click the drive icon, and choose the Properties command from the shortcut menu that appears. Figure 1-1 shows the Properties dialog box for a 476GB hard drive that has about 224GB of free space.

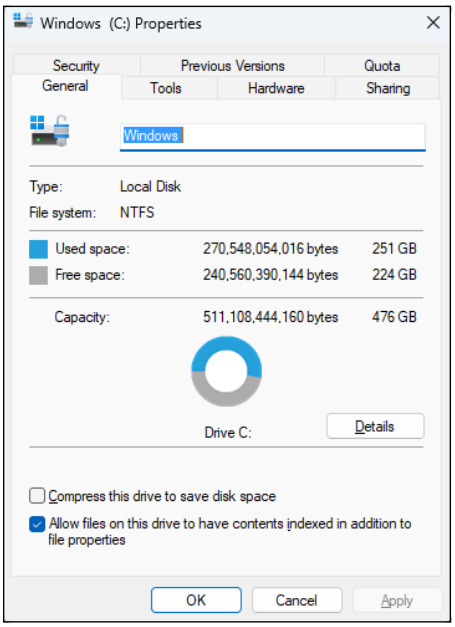


FIGURE 1-1: The Properties dialog box for a disk drive.

If your computer has more than one hard drive, Windows lists an icon for each drive in the Computer window. Jot down the size and amount of free space available on each drive.

- **The amount of memory:** To find this information in Windows, right-click Computer from the Start menu and choose the Properties command. The amount of memory on your computer is shown in the dialog box that appears. For example, Figure 1-2 shows the System Properties page for a computer with 64GB of RAM.
- **The operating system version:** This you can also deduce from the System Properties dialog box. For example, the Properties page shown in Figure 1-2 indicates that the computer is running Windows 11 Professional.

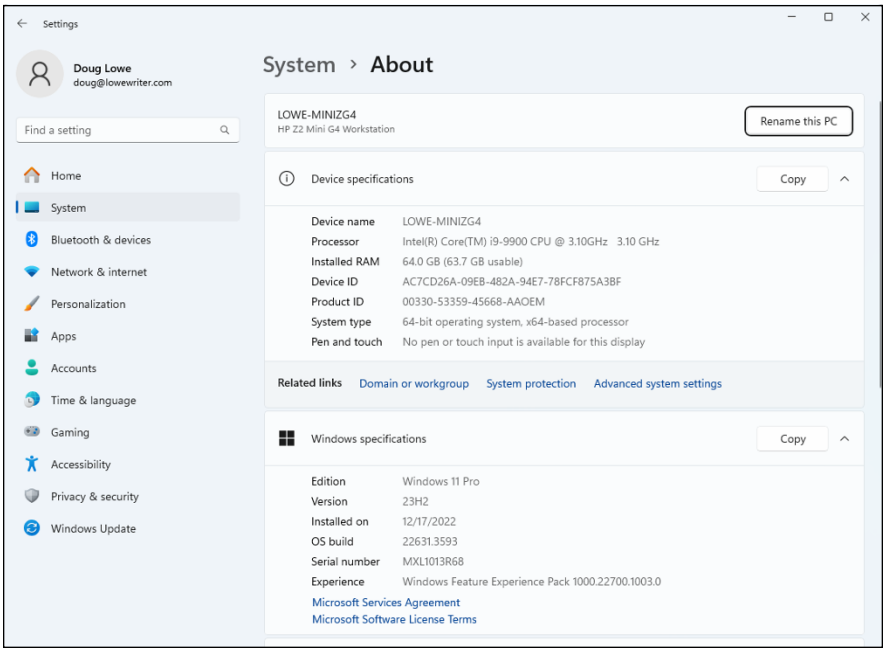


FIGURE 1-2: The Properties page for a computer with 64GB of RAM.

- **What kind of printer, if any, is attached to the computer:** Usually, you can tell just by looking at the printer. You can also tell by double-clicking the Printers icon in Control Panel.
- **Any other devices connected to the computer:** A DVD or Blu-ray drive? Scanner? External disk? Webcam? Battle droid? Hot tub?
- **What software is used on the computer:** Microsoft Office? AutoCAD? QuickBooks? Make a complete list and include version numbers.

Programs that gather information for you

Gathering information about your computers is a lot of work if you have more than a few computers to network. Fortunately, several software programs are available that can automatically gather the information for you. These programs inspect various aspects of a computer, such as the CPU type and speed, amount of RAM, and the size of the computer's hard drives. Then they show the information on the screen and give you the option of saving the information to a hard drive file or printing it.

Windows comes with just such a program: Microsoft System Information. Click Start, type **System Information**, and click the link to the System Information program.

When you fire up Microsoft System Information, you see a window similar to the one shown in Figure 1-3. Initially, Microsoft System Information displays basic information about your computer, such as your version of Microsoft Windows, the processor type, the amount of memory on the computer, and so on. You can obtain more detailed information by clicking Hardware Resources, Components, or other categories in the left side of the window.

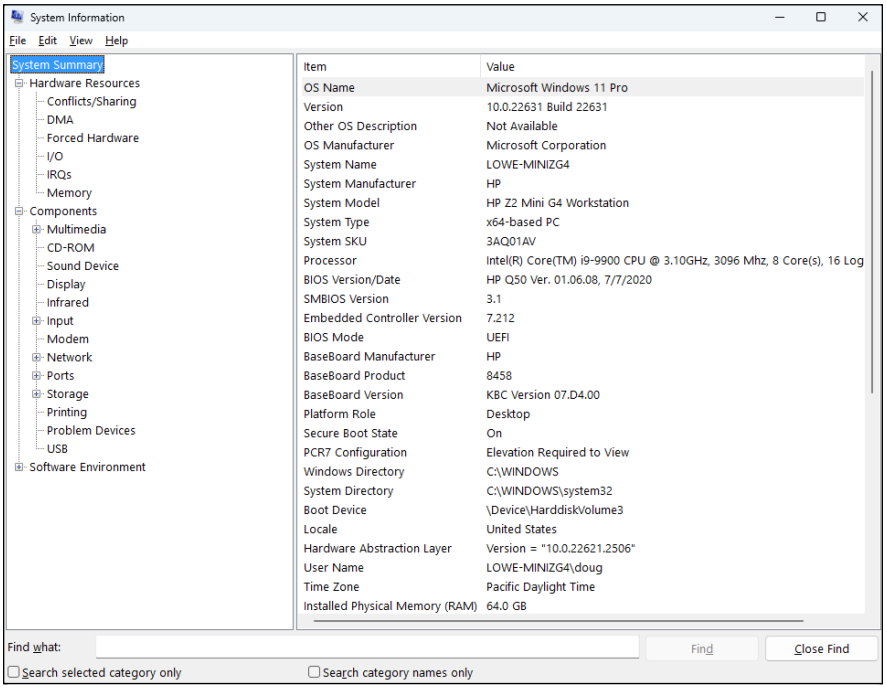


FIGURE 1-3: Let the System Information program gather the data you need.

Considering Cable

Over the years, several different types of cables have been used for networking. But today, almost all cabled networks are built using simple copper-based unshielded twisted pair (UTP) cable. Figure 1-4 shows a twisted pair cable.

When you use UTP cable to construct an Ethernet network, you connect the computers in a starlike arrangement, in which each computer is connected to a central point. At the center of the stars are switches (see Book 1, Chapter 3). Depending on the model, a single switch can connect from 4 to 48 or more devices.

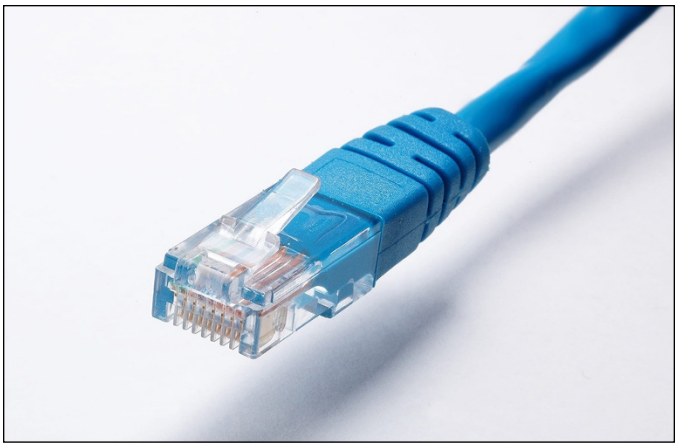


FIGURE 1-4: Twisted-pair cable.

Here are a few additional details that you should know about twisted pair cabling:

- UTP cable consists of four pairs of thin wire twisted around each other; several such pairs are gathered up inside an outer insulating jacket. Ethernet uses two pairs of wires, or four wires altogether.
- UTP cable comes in various grades known as *categories.* Don't use anything less than Category 5e cable for your network; Category 6 is better yet. Although lower-category cables may be less expensive, they won't be able to support faster networks.

  Be prepared for the future. Although higher-category cables are more expensive than lower-category cables, the real cost of installing Ethernet cabling is the labor required to actually pull the cables through the walls. As a result, I recommend that you invest in Category 6.

TIP

- UTP cable connectors look like modular phone connectors but are a bit larger. UTP connectors are officially called RJ-45 connectors.
- UTP cable can be purchased in prefabricated lengths, but for most jobs you'll want to purchase the cable in bulk and have a professional installer attach the connectors. Or, you can attach the connectors yourself using a simple crimping tool you can purchase for about $50.
- The maximum allowable cable length between the switch and the computer is 100 meters (about 328 feet). That should be more than enough for most circumstances, but don't forget that the distance includes the vertical distance required for getting from the floor up to the ceiling, and back down again.
- Always leave at least 5 feet or more of extra cable neatly coiled up in the ceiling space above each location where the cable drops through the wall to floor level. That way, you'll have some flexibility to re-route the cable later on if necessary.

Surmising Switches

As I mention in the previous section, computers and other devices are connected to a network in a starlike configuration, with switches at the center of the star. Figure 1-5 shows a switch with five computers connected to it.

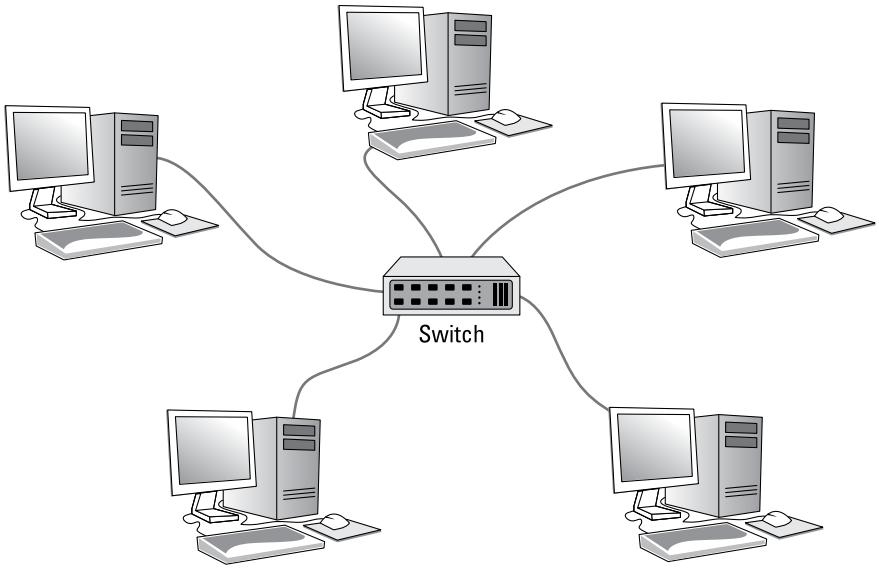


FIGURE 1-5: A switch with five computers connected.

A switch contains a number of *ports*, each of which is a receptacle that can accommodate an RJ-45 jack connected to a UTP cable. In Figure 1-5, there are five UTP cables. One end of each of these cables is plugged into a port on the switch, and the other end is plugged into the computer's network adapter.

Although it may not be obvious from the figure, the switch does not have to be in the same room as the computers. In fact, ideally the switch will be in a separate room from the computers. The cables run through the ceilings and the walls from the location of the switch to the location of the computers, within the 100-meter limit of UTP cable. (The switches are generally located in the same room as the servers.)

Here are some additional ins and outs for working with switches:

- Because you must run a cable from each computer to the switch, find a central location for the switch to which you can easily route the cables.

- The switch requires electrical power, so make sure that an electrical outlet is handy.
- As a general rule, purchase twice as many switch ports as you currently need. Don't buy an eight-port switch if you want to network eight computers because when (not *if*) you add the ninth computer, you'll have to buy another switch.
- You can connect — or *daisy-chain* — switches to one another, as shown in Figure 1-6. You connect one end of a cable to a port on one switch and the other end to a port on the other switch.



WARNING

- Although you can daisy-chain as many switches together as you want, in actual practice you should limit the number of daisy chains in your switch configuration. Daisy-chaining can slow down a network a bit because each switch must fully receive each packet before it begins to forward the packet to the next switch. (However, some switches actually start the packet forwarding before the entire packet is received, which reduces the performance hit a bit.)
- If you need more ports than a single switch can provide, you can use *stackable switches.* Stackable switches have high-speed direct connections that enable two or more switches to be connected in such a way that they behave as if they were a single switch.

  This type of connection is sometimes called a *back-plane connection* because the interconnect may be on the back of the switch, but that's not always the case. If a single switch will suffice for you now, but there is a reasonable chance that you might outgrow it and need a second switch, I suggest you invest in a stackable switch so that you can expand your network later without daisy-chaining.
- Another way to provide a high-speed interconnect between switches is to purchase switches that have a few high-speed SFP ports. You can then equip those ports with 10 Gb connections to route traffic between the switches. (These high-speed connections can also be used to connect switches to servers.)
- Yet another way to create high-speed interconnects between switches is to use a feature called *link aggregation.* If your switches provide this feature, you simply run two or more cables between the switches, using two or more ports on each switch. Then, you use the switch's configuration software to bond the two ports together to create one link with double the port speed.

  Professional-quality network switches have network-management features that allow you to log in to the switch, usually via a web interface, to monitor and configure the switch. Such switches are called *managed switches.* Consumer-grade switches, also called *unmanaged switches,* are less expensive primarily because they do not support this feature. If you have more than a few dozen users, you'll want to invest in managed switches.

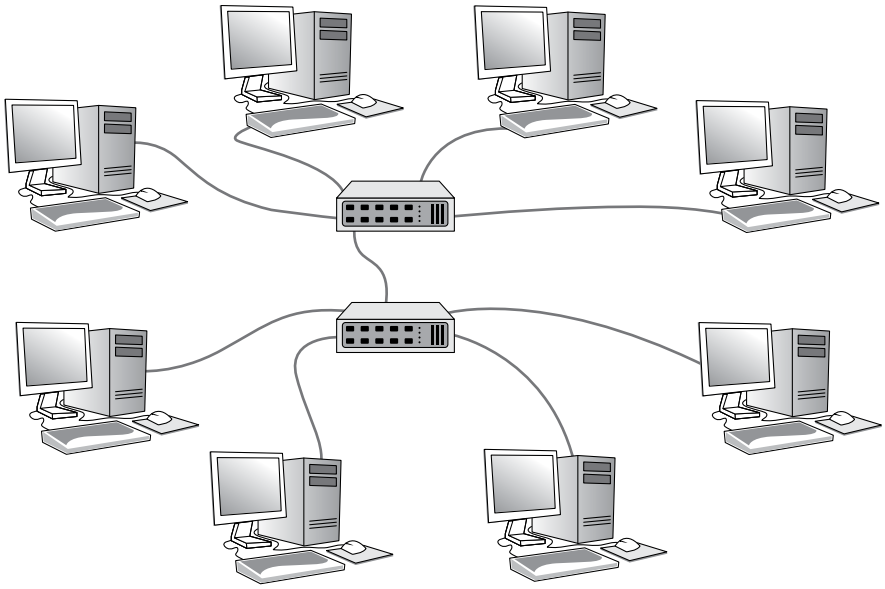


FIGURE 1-6: Daisy-chaining switches.

Planning the Network Topology

Topology refers to the way the devices in your network are connected to each other via network switches. You'll need to determine what kind of switches to use, how many, where to run the cable, where to locate the switches, and so on.

Here are just a few of the questions to consider:

- Where will the servers be located? Will there be just a single switch in the server room, or will there be a central switch with workgroup switches located closer to the computers?
- Where will you place workgroup switches? On a desktop somewhere within the group or in a small wiring closet?
- How many client computers will you place on each workgroup switch, and how many switches will you need?
- If you need more than one switch, how will you connect the switches to each other?
- If you need wireless networking, what kind of wireless networking devices will you need, where will you put them, and how will you connect them to the network?

For midsized networks (say, 50 to 200 users), a common way to design the network topology is to use a two-layer switch architecture as shown in Figure 1-7:

- **Core layer:** The *core layer* contains high-performance switches that connect to the servers, the internet gateway, and to each other. These connections should be as fast as possible — ideally, 10 Gbps fiber or copper connections using SFP ports.
- **Access layer:** The *access layer* consists of switches that are connected to the core layer and to the end-user computers.

In Figure 1-7, there is one switch at the core layer and four switches at the access layer. The two core switches are connected to each other, to the servers, and to the access layer switches using 10 Gbps fiber SFP connections. The access switches connect to the computers using standard 1 Gb Ethernet connections.

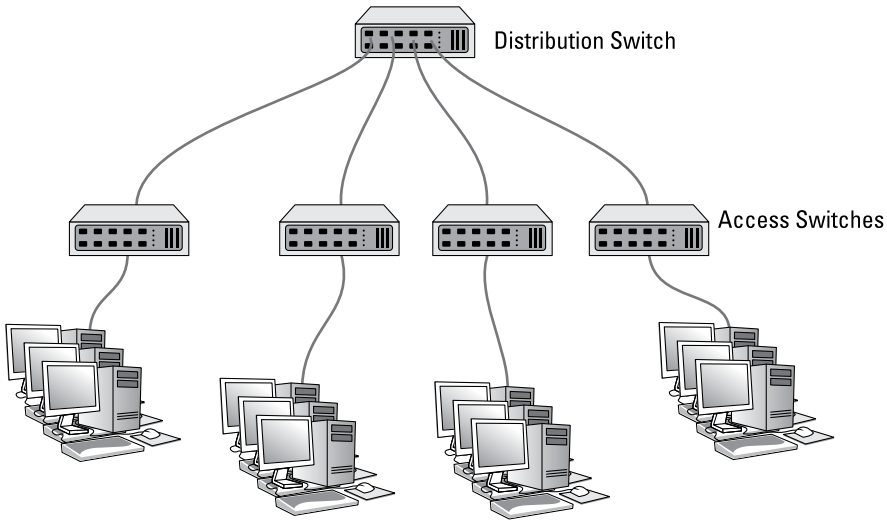


FIGURE 1-7: A two-tiered switch design.

For even larger networks, a three-tier design can be used. In that case, a *distribution layer* is added between the access and core layers. The servers are moved to the distribution layer and the core layer using specialized high-speed switches whose sole purpose is to move large amounts of data between the distribution switches as quickly as possible.

Planning the TCP/IP Implementation

In addition to planning for the physical parts of your network infrastructure (cables, switches, ports, and so on), you'll also need to plan the details of how you'll implement TCP/IP for your network. TCP/IP is the basic networking protocol that your network uses to keep track of the individual computers and other devices on the network. Each computer or device will need an IP address (for example, 10.0.101.65). You'll need to devise a plan for how these addresses will be allocated.

You learn everything you need to know about TCP/IP in Book 2, Chapter 2, so make sure you understand the information in that chapter before you complete this part of your plan. For now, here are some of the main points your plan should address:

- **The subnet and VLAN structure of your network:** Will everything be on a single subnet, or will you use two or more subnets to separate different types of devices?

  Although it isn't impossible, dividing an existing network into separate subnets later on is a bit of a pain. So unless your network is very small, I suggest you plan on using subnets from the very start. In particular, you should consider using separate subnets for the following:

  - **Wireless networks:** In fact, if you create two or more wireless networks (for example, a corporate network for company-owned mobile devices, an employee network for personal smartphones, and a guest network for visitors), I suggest a separate subnet for each of the wireless networks.
  - **Remote locations that will be connected via a VPN tunnel.**

  In addition, if you use IP phones, definitely put the phones on their own subnet. And if your organization has more than a few dozen users, consider dividing them among two or more subnets according to their work groups. For example, you might use one subnet for the sales department and another one for the production department.

  You'll probably need to set up VLANs to manage your subnets. Typically, there's a one-to-one correspondence between subnets and VLANs; in other words, each subnet lives on its own VLAN.

  
  TIP

  Don't go overboard with the subnets, however. Try to find the right balance between running the entire organization on a single subnet versus creating a lot of subnets, each with just a few users.

  
  TECHNICAL STUFF

  Why bother with the subnets? The main reason is to avoid issues that will come up when your organization grows. You may have just 20 employees now, but years from now, when you have 100, and everyone starts bringing their smartphones and tablets and connecting to your Wi-Fi, you'll find that

the limit of 253 devices per subnet on a 255.255.255.0 network is simply not enough. When you run out of DHCP space and your users can't get on the network, you'll wish you had spread things out over a couple of subnets.

- **The DHCP structure:** Speaking of DHCP, what server will be responsible for DHCP? What will be the DHCP scope — that is, the range of addresses that are given out by DHCP? How will the size of your scope accommodate all the devices that will require DHCP on the subnet, with plenty of room for growth? (Again, be especially careful if you connect a wireless access point to the same subnet that your cabled network is on. You'll be surprised how quickly you can run out of IP addresses when every one of your users brings an iPhone and an iPad. Subnets are the best answer to that problem.)
- **The static IP addresses of devices whose IP should never change:** These devices may include servers, printers, firewalls, and other managed devices. You'll be surprised how quickly these can add up, as well. You'll need static IP addresses for each of the network interfaces on your servers, for your switches, printers, copiers, fax machines, firewalls, routers, tape backup devices, and network storage devices. If you use virtualization software, the host processors will also need an IP address for each network interface. Even your UPS battery backups may want an IP address. The list goes on and on.


TIP

It is absolutely imperative that you keep a good record of what static IP addresses you have assigned in your network, and that you configure your DHCP server properly so that it doesn't step on top of static IP addresses. Every time you add a device with a static IP address, be sure to update your list. And, just as important, whenever you retire a device that uses a static IP address, update your master list to remove the IP address.

Drawing Diagrams

One of the most helpful techniques for creating a network plan is to draw a picture of it. The diagram can be a detailed floor plan, showing the actual location of each network component: a *physical map*. If you prefer, the diagram can be a *logical map*, which is more abstract and Picasso-like. Any time you change the network layout, update the diagram. Also include a detailed description of the change, the date that the change was made, and the reason for the change.

You can diagram very small networks on the back of a napkin, but if the network has more than a few computers, you'll want to use a drawing program to help you create the diagram. One of the best programs for this purpose is Microsoft Visio, shown in Figure 1-8.

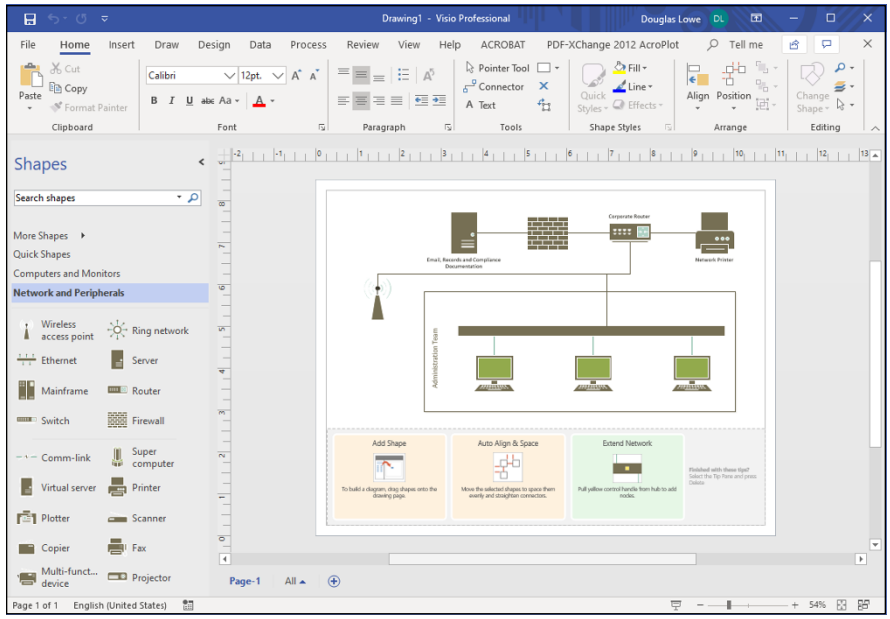


FIGURE 1-8: Using Visio to draw a network diagram.

Here's a rundown of some of the features that make Visio so useful:

- Smart shapes and connectors maintain the connections drawn between network components, even if you rearrange the layout of the components on the page.
- Stencils provide dozens of useful shapes for common network components — not just for client and server computers, but for routers, hubs, switches, and just about anything else you can imagine. If you're really picky about the diagrams, you can even purchase stencil sets that have accurate drawings of specific devices, such as Cisco routers or IBM mainframe computers.
- You can add information to each computer or device in the diagram, such as the serial number or physical location. Then, you can quickly print an inventory that lists this information for each device in the diagram.
- You can easily create large diagrams that span multiple pages.

More Questions Your Network Plan Should Address

In addition to the basic questions of why you need a network, what kind of servers you need to provide, and what kind of infrastructure your network will require, your network plan should address the following questions:

- **Does it allow for growth?** What growth areas do you anticipate over the next few years? Does this network plan provide for such growth? For example, if you currently have 20 devices on the network, a 24-port switch may be adequate for today. But you should consider a 48-port switch instead. It will cost more now, but will simplify your expansion down the road.

  Similarly, if you anticipate that each office will have just one employee, consider what you'll have to do if you run out of offices and end up putting two employees in each office. If you run a single cable to each office now, you'll have to pay to have a second cable run later. Better to spend a little more for extra cable and have the installer pull two cables to each office. (Better yet, have the installer pull three cables to each office: When you move a second employee into the office, you may also put a printer in there.)

- **How will you secure it?** What kind of safety precautions will you take to keep unwanted visitors off your network? You'll need a strong, well-configured firewall to keep intruders from breaking in to your network via your internet connection. If you're installing wireless access points, you'll have to take precautions to secure the wireless networks. And you'll need strong password policies to prevent hackers who do manage to get on to your network from getting at any valuable data.

  For more information about network security, refer to the chapters in Book 10.

- **How will you back it up?** You'll need to include a solid plan to back up your servers and the data that resides on them. That plan will probably require additional hardware, such as separate disk storage to hold the first level of backup data, as well as a means to get the backed-up files off-site so they can survive a true disaster such as a fire or flood.

  You'll also need to be certain that you provide adequate network disk storage so that all users can put all their work on the network, where it can be backed up. In lieu of that, you'll need a plan that backs up not only your servers, but also the client computers.

  For more information about backing up your network, refer to Book 3, Chapter 6.

- **How will you recover from failures?** Make sure you have a plan in place that will protect you from the commonplace maladies of daily life such as occasional power failures, as well as from the unforeseen, such as vandalism, theft, or fire. Every device on your network, no matter how insignificant, should be protected by battery backup. When possible, you should have spares of critical components.

  For more information about disaster recovery, see Book 10, Chapter 4.

IN THIS CHAPTER

- Looking at WAN connection options
- Choosing a router
- Securing your connection with a firewall
- Connecting remote users and branch offices with VPN or EPL

Chapter 2
Wide Area Networks

Obviously, your network needs to be connected to the internet. But that's easy, right? All you have to do is call the cable company and have them send someone out. They'll get you hooked up in a jiffy.

Wrong. Unfortunately, connecting to the internet involves more than just calling the cable company. For starters, you have to make sure that cable is the right way to connect. Then you have to select and configure the right device to connect your network to the internet. And, in all likelihood, you have to figure out how to provide remote access to your network so you can connect from your home office, from a hotel room on a business trip, or from the branch office in Albuquerque. And finally, you have to lie awake at night worrying whether hackers are breaking into your network via its internet connection. Which, of course, they are.

Not to worry. The advice in this chapter helps you decide how to design your wide area network (WAN) architecture. This includes your internet connection, as well as remote access options.

Connecting to the Internet

Connecting to the internet isn't free. For starters, you have to purchase the computer equipment necessary to make the connection. Then you have to obtain a connection from an internet service provider (ISP). The ISP charges you a monthly fee that depends on the speed and capacity of the connection.

Choosing an ISP and negotiating a contract is a basic first step in setting up a WAN connection for your private network. The following sections describe the most commonly used methods of connecting network users to the internet.

Connecting with cable or DSL

For small and home offices, the two most popular methods of connecting to the internet are cable and digital subscriber line (DSL). Cable and DSL connections are often called *broadband connections* for technical reasons you don't really want to know.

Cable internet access works over the same cable that brings 40 billion TV channels into your home, whereas DSL is a digital phone service that works over a standard phone line. Typical cable and DSL speeds range from 50 Mbps to 1 Gbps.

In most areas of the United States, basic cable internet runs about $50 per month for residential users; business users can expect to pay two to three times that for the same speeds, primarily because the providers expect a higher level of usage and offer a slightly better service level for business connections.

The cost for DSL service depends on the access speed you choose. In some areas, residential users can get a relatively slow DSL connection for as little as $30 per month. For higher access speeds or for business users, DSL can cost substantially more.

Besides the cost, there are a few inherent disadvantages with DSL and cable providers:

- **Cable and DSL are *asymmetrical* technologies, which means that their download speeds are much faster than their upload speeds.** For example, a circuit that can download at 100 Mbps is probably limited to about 10 Mbps for upload speeds. For many users, this is acceptable. But if you need to upload data as often as you need to download, the asymmetrical nature of cable and DSL will be a drawback.

- **Business-class cable and DSL provide "best effort" service levels.** The provider will do its best to keep the connection up and respond to issues, but there is no guaranteed service level. When the service goes down, it can be down for a few hours or a few days.

  And it will go down. Most users find that business-class cable and DSL are unreliable. Some users find that short service interruptions are an almost daily experience. The reason is that both cable and DSL service are shared services. The performance you get depends on what else is happening nearby. If all your neighbors suddenly start streaming the latest big thing on Netflix, your performance will suffer. Business-class cable and DSL don't claim to be 100 percent reliable — and they aren't.

- **Cable and DSL access aren't available everywhere.** But if you live in an area where cable or DSL isn't available, you can still get high-speed internet access by using a satellite hookup or a cellular network.

Connecting with T1 lines

Telephone providers such as AT&T, Time Warner, and others offer internet service over dedicated copper phone lines using a time-proven technology called T1. I say "time-proven" because the original T1 service was developed in 1962, the same year *The Beverly Hillbillies* premiered on CBS. This was decades before the internet even existed. T1 is not particularly fast — a single T1 line carries data at a paltry 1.44 Mbps. You can bond multiple T1 lines together to increase the speed, but you'd have to use 35 T1 lines to get 50 Mbps service. Newer versions such as T3 provide faster service (44.184 Mbps) but cost considerably more.

Although T1 is not the best type of service available (see the next section, "Connecting with fiber"), it's an improvement over business-class cable or DSL from a service and reliability perspective. Your carrier will provide a guaranteed service-level agreement (SLA) with a T1 line and will give you priority service if a problem occurs.

In addition, T1 service is symmetrical and predictable. Upload and download speeds are the same, so if you have ten T1 circuits that aggregate to 14.4 Mbps, you'll get that performance level for both uploads and downloads. And because the circuits are dedicated to your network, the performance will be consistent — it won't slow down in the afternoon when school gets out and kids start gaming over the internet with their home cable or DSL connections.

If you don't have enough users to justify the expense of an entire T1 or T3 line, you can lease just a portion of the line. With a fractional T1 line, you can get connections with speeds of 128 Kbps to 768 Kbps; with a fractional T3 line, you can choose speeds ranging from 4.6 Mbps to 32 Mbps.



TIP

You may be wondering whether T1 or T3 lines are really any faster than cable or DSL connections. After all, T1 runs at 1.544 Mbps and T3 runs at 44.184 Mbps, and cable and DSL claim to run at much faster speeds, at least for downloads. But there are many differences that justify the substantial extra cost of a T1 or T3 line. In particular, a T1 or T3 line is a *dedicated* line — not shared by any other users. T1 and T3 are higher-quality connections, so you actually get the 1.544 or 44.184 connection speeds. In contrast, both cable and DSL connections usually run at substantially less than their advertised maximum speeds because of poor-quality connections and because the connections are often shared with other users.

Connecting with fiber

The fastest, most reliable, best, and of course most expensive form of internet connection is fiber-optic. Fiber-optic cable uses strands of glass to transmit data over light signals at very high speeds. Because the light signals traveling within the fiber cables are not subject to electromagnetic interference, fiber connections are extremely reliable; about the only thing that can interrupt a fiber connection is if someone physically cuts the wire.

Fiber connections are typically available starting at 100 Mbps and ranging up to 10 Gbps or more. Obviously, the 10 Gbps service will cost a lot more than the 100 Mbps. But the cost of increased speed is incremental. For example, 100 Mbps might cost $800 per month, but 500 Mbps might be $1,100 per month and 1 Gbps might be $1,400 per month. In other words, the cost per gigabit per second goes down as the speed increases.

Costs vary greatly depending on your location, so the only way to find out for sure is to get quotes from providers in your area.

In most major communities throughout the United States, providers are still building out their fiber-optic networks. The cost to bring fiber to your location may be prohibitive if you're in an area that isn't yet developed. If a provider already has fiber under the street running right past your building, getting fiber to your business will be relatively inexpensive. But if the nearest fiber is five miles away, the cost may be prohibitive.

You may be able to negotiate with the provider if you're willing to commit to a longer term of service, such as three, four, or even five years. That will make their investment more worthwhile. It also helps if you're in a business area where you'll be the first fiber customer but there is a potential customer pool nearby that the provider can tap into. If you're the only business out on the edge of town, you may not be able to convince anyone to bring fiber to you.


TIP

Phone service can be delivered via a fiber connection and bundled for one price. That can work to your advantage, because the provider will be more willing to bargain on the overall deal if the phone service is included.

Connecting with a cellular network

In areas where wired service (such as cable or fiber) is not available, you may be able to find wireless service, which provides internet access using cellular or other wireless technology.

Cellular connections are not particularly fast, but they're getting faster all the time. The most widely used generation of cellular technology (4G) can consistently achieve speeds in the neighborhood of 10 to 12 Mbps for download, with peak speeds approaching 50 Mbps. Upload is a bit slower, usually in the 5 Mbps range.

However, actual performance depends a lot on your location. I've seen 4G service as bad as 0.1 Mbps. You should use a smartphone to test the upload and download speed in your area before committing to a cellular solution.

The next-generation cellular technology (5G) is currently being rolled out throughout the world. It can deliver speeds eight to ten times faster than 4G and is currently available in most major metropolitan areas. Even if 5G is not available in your area yet, I suggest you purchase 5G equipment now on the assumption that 5G is coming to your neighborhood.


TIP

With a cellular connection, the cost isn't so much the speed but the amount of data transferred. Individual cellular contracts run about $50 to $100 per month, but they typically limit the amount of data to about 5GB or 10GB per month. You can expect to pay considerably more than that if you need more data.

Choosing a Router

After you choose a method to connect to the internet, you can turn your attention to setting up the connection so that your private network can access the internet. The provider you select for your internet connection will give you an *Ethernet handoff*, which is simply an Ethernet port that you can use to connect to your private network. You'll need a router to make that connection. The router is the device that provides the link between your private network and the Ethernet handoff that leads to the internet. (For more information about routers, refer to Book 1, Chapters 2 and 3, and Book 2, Chapter 4.)

Because all communications between your network and the internet must go through the router, the router is a natural place to provide the security measures necessary to keep your network safe from the many perils of the internet. As a result, a router used for internet connections often doubles as a firewall, as described in the "Securing Your Connection with a Firewall" section, later in this chapter.

Choosing a small office router

For a small office, you can probably get by with a consumer-grade router that you can purchase at a local electronics retailer such as Best Buy. Figure 2-1 shows one such router, a Linksys Hydra Pro 6E. This router has the following specifications:

- A WAN connection that lets you connect to your ISP's Ethernet handoff.
- A four-port 1 Gbps Ethernet switch. You can use this to connect up to four PCs, or to connect to an external switch for additional computers.
- A Wi-Fi Access Point that works with most 802.11 Wi-Fi standards, including 802.11ac.
- A USB 3.0 port that lets you connect a USB disk drive to provide storage accessible throughout your network.
- Built-in firewall capability.



FIGURE 2-1: A Linksys Hydra Pro 6E router.

Courtesy of Linksys

To learn more about this router and other routers offered by Linksys, visit `www.linksys.com`.

Choosing an enterprise router

For larger networks where greater throughput and more control is needed, you'll want to select an enterprise-grade router with a built-in firewall. There are many brands to choose from, but most professionals use Cisco.

These routers range from small tabletop units to powerful rack-mounted units that are capable of serving networks of all sizes. As the name *Firepower* suggests, these devices aren't just routers but incorporate state-of-the-art firewall capabilities.

Table 2-1 outlines the basic capabilities of six models of the ASA 5500-X that are appropriate for most networks.

TABLE 2-1 Cisco Firepower 1000 Series Models

Model	Throughput	Ethernet Ports	Form Factor
FPR-1010	900 Mbps	8 x 1 Gbps	Desktop
FPR-1120	2.6 Gbps	8 x 1 Gbps 4 x SFP	1U Rackmount
FPR-1140	3.5 Gbps	8 x 1 Gbps 4 x SFP	1U Rackmount
FPR-1150	6.1 Gbps	8 x 1 Gbps 2 x SFP 2 x 10G SFP+	1U Rackmount

As you can see, the main differences between these models is the total throughput that can be supported and the addition of SFP or SFP+ Ethernet ports. To support the higher bandwidth, the higher model numbers have faster CPUs and more RAM than the lower model numbers. Additional models of the ASA series can support substantially more bandwidth, but these models are sufficient for nearly all mid-size networks.

Choosing a cellular router

If you opt to use a cellular connection for internet, either as your office's primary connection or as a fail-over connection in case your primary connection goes down, you'll need a router that can interface with a cellular modem. Cellular modems are usually USB devices, so your router will need to provide a USB external port to connect the cellular modem to.

Securing Your Connection with a Firewall

If your network is connected to the internet, a whole host of security issues bubbles to the surface. You probably connected your network to the internet so that your network's users can get out to the internet. Unfortunately, however, your internet connection is a two-way street. It not only enables your network's users to step outside the bounds of your network to access the internet, but it also enables others to step in and access your network.

And step in they will. The world is filled with hackers who are looking for networks like yours to break into. They may do it just for the fun of it, or they may do it to steal your customers' credit card numbers or to coerce your mail server into sending thousands of spam messages on behalf of the bad guys. Whatever their motive, rest assured that your network will be broken into if you leave it unprotected.

A *firewall* is a security-conscious router that sits between the internet and your network with a single-minded task: preventing *them* from getting to *us*. The firewall acts as a security guard between the internet and your private network. All network traffic into and out of the private network must pass through the firewall, which prevents unauthorized access to the network.



WARNING

Some type of firewall is an absolute must if your network has a connection to the internet, whether that connection is broadband (cable modem or DSL), T1, fiber, cellular modem, smoke signals, carrier pigeon, or anything else. Without it, sooner or later a hacker will discover your unprotected network and tell his friends about it, and within a few hours, your network will be toast.

You can set up a firewall in two basic ways:

- **Firewall appliance:** The easiest way, and usually the best choice. A firewall appliance is basically a self-contained router with built-in firewall features.

  Most firewall appliances include web-based interfaces that enable you to connect to the firewall from any computer on your network by using a browser. You can then customize the firewall settings to suit your needs.

- **Server computer:** Can be set up to function as a firewall computer.

  The server can run just about any network operating system, but most dedicated firewall systems run Linux.

Whether you use a firewall appliance or a firewall computer, the firewall must be located between your network and the internet, as shown in Figure 2-2. Here, one end of the firewall is connected to a network switch, which is, in turn, connected

to the other computers on the network. The other end of the firewall is connected to the internet. As a result, all traffic from the LAN to the internet (and vice versa) must travel through the firewall.

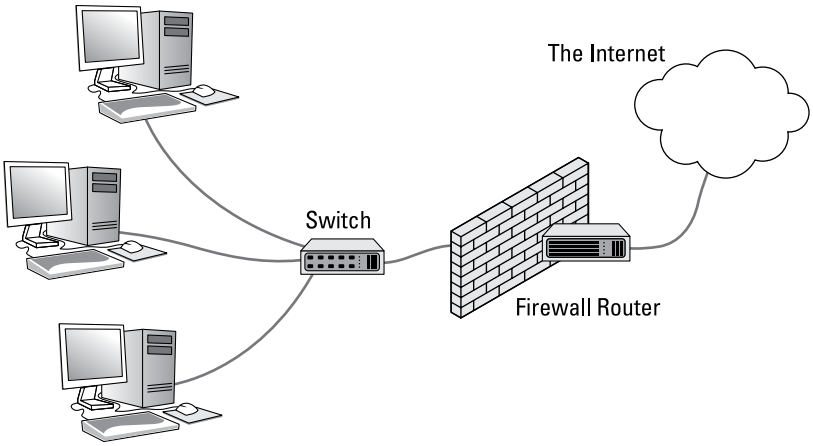


FIGURE 2-2: A firewall router creates a secure link between a network and the internet.

The term *perimeter* or *edge* is sometimes used to describe the location of a firewall on your network. In short, a firewall is like a perimeter fence that completely surrounds and protects the edge of your property and forces all visitors to enter through the front gate.



WARNING

In large networks, figuring out exactly where the perimeter is located can be a little difficult. If your network has two or more internet connections, make sure that every one of those connections connects to a firewall — and not directly to the network. You can do this by providing a separate firewall for each internet connection or by using a firewall with more than one internet port.



TIP

Some firewall routers can also enforce virus protection for your network. For more information about virus protection, see Book 10, Chapter 2.

Providing Redundancy for Your Internet Connection

Important considerations when designing how your private network will connect to the internet are the reliability of your internet connection and the importance to your company for having that connection be reliable. For some companies, an

occasional disruption in internet connectivity is acceptable. For others, it isn't — business grinds to a halt, and money is lost for every minute the internet is down.

If that's the case, you'll want to provide at least two pathways to the internet: a primary internet connection and a backup internet connection. The backup connection is often called a *fail-over connection*, because it comes into play only when the primary connection fails. With the right setup (and proper configuration), fail-over can happen automatically. When the primary internet connection drops, the gateway router can instantly switch over to the backup connection. Then, when the primary connection is re-established, the gateway router can revert to it.

In most cases, you can get away with a slower and less reliable connection for the backup. For example, you might have a fiber-optic connection as your primary connection and use business-class cable as the backup. Fiber-optic connections are very reliable, but they do go down from time to time. Especially when a back-hoe operator doesn't realize that he or she is digging in the middle of a street where your provider's fiber run lies buried.

Business-class cable isn't nearly as reliable as fiber, but what are the odds that both will be down at the same time? Not likely, because most providers use separate routes for their fiber and cable runs. So, a single mishap with a back-hoe is unlikely to take out both.

If you do use a backup internet service, you'll need to ensure that your router can support automatic fail-over. That means you'll need to use an enterprise-grade router.


TIP

If you use a backup internet service with automatic fail-over, be sure to test it periodically. The easiest way to do that is simply to unplug the cable from the primary internet Ethernet handoff to the router, and then see if your router has switched over to the backup connection. If you can still reach the internet, your fail-over is working. (If you want to keep what friends you have at your company, I suggest conducting this test after hours.)

Securing Connections to Remote Locations and Remote Users

One final topic for this chapter is providing secure connections for remote users. These can be individuals who need to occasionally work from home or from the road, telecommuters who have convinced their boss to let them work from a home

office, or branch offices that need a permanent connection to the main office network.

The solution to all these situations is a virtual private network. A VPN works by establishing a secure *tunnel* between two devices that are connected to the internet. For the private network at your main office, the gateway router will provide the VPN capability. Remote users can run VPN software on their computers to connect to the main office VPN; remote sites such as branch offices should use gateway routers that can permanently (and transparently) connect to the VPN.

As part of your WAN network planning, you should identify all the VPN capabilities that your network will require. This will help you choose appropriate routers, because less expensive routers don't usually provide VPN features.

Figure 2-3 shows an example of a network drawing that shows four VPN tunnels — three to remote offices and one for mobile users. To support this network, you'd need a router that can let you create at least four separate VPNs. So, a consumer-grade gateway won't be sufficient for this network. In the figure, I specify various Cisco ASA routers to use for the VPN connections.

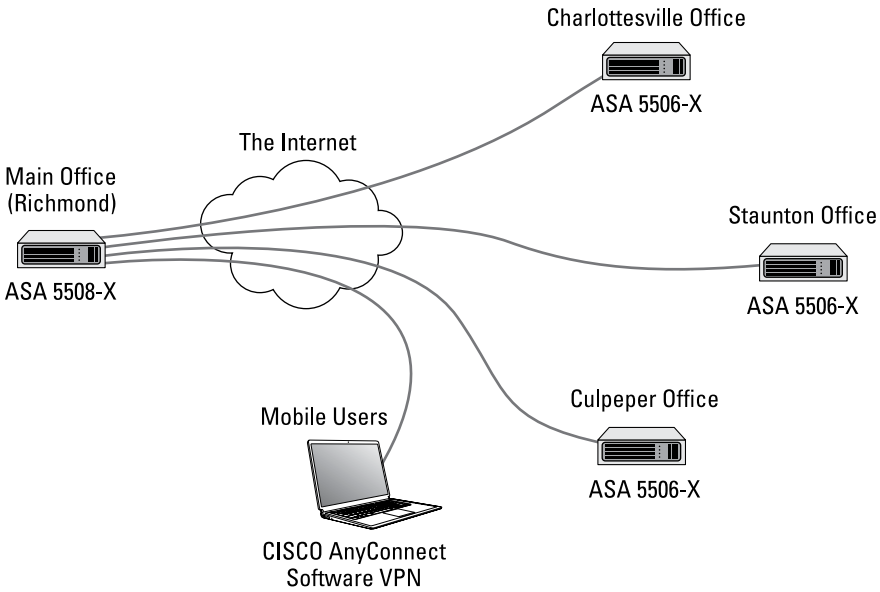


FIGURE 2-3: A network that requires four VPN connections.

For more information about VPN, refer to Book 4, Chapter 6.

Connecting Remote Offices with an Ethernet Private Line

Your ISP may be able to provide you with a point-to-point fiber link between two office locations. This type of connection is called an *Ethernet private line* (EPL). An EPL is effectively the same thing as a VPN, except that the internet provider manages all the details necessary to maintain the privacy and security of the point-to-point link. The circuit presents itself to you as a standard Ethernet connection, which doesn't require a firewall on either end. So, you plug both ends of the EPL circuit into a switch rather than a firewall.

You can set it up so that both offices have their own separate internet connections (which require firewalls), or you can set it up so that only one of the two locations has an internet connection with a firewall; the remote location then accesses the internet via the main office's internet connection.

Figure 2-4 shows a setup in which two remote offices connect to a home office with EPL circuits. The home office in Raleigh connects to two remote offices in Mount Pilot and Mayberry. The Raleigh and Mount Pilot offices have direct internet connections, but users in the Mayberry office must access the internet through the home office's internet connection.

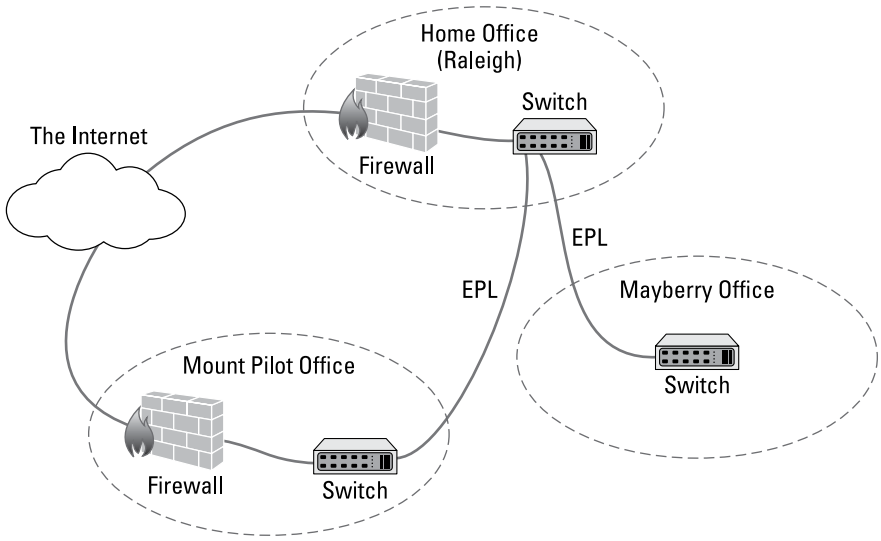


FIGURE 2-4: Two branch offices connected with EPL circuits.

IN THIS CHAPTER

- Knowing how many servers is enough
- Identifying what kinds of servers you need
- Figuring out how to connect all your servers

Chapter 3
Server Architecture

This chapter presents the task of planning the servers that your network will require. Over time, most networks gather servers like squirrels collect nuts. You start with just a few, and within a few years you have dozens of servers. In this chapter, you take stock of what servers you need, consider whether you should combine some of your servers, and look briefly at how to connect all the servers.

In the end, you'll find that the best way to manage your servers is to virtualize them. I look more closely at that topic in the next chapter.

Deciding How Many Servers You Need

A basic decision that you must make when planning a network is determining how many servers your network will require. At a minimum, all but the smallest networks require at least two domain controllers, plus additional servers to satisfy the needs of your users.

When setting up servers, you have the option of creating a bunch of single-purpose servers or a smaller number of multipurpose servers. For example, if you need a printer server and a file server, you can create a separate server for each function,

or you can use a single server to perform both functions. There are advantages and disadvantages to both approaches:

- Consolidating server functions into a smaller number of servers can save in licensing and administration costs. For example, if you can support all your functions on 6 servers rather than 12, you only need to purchase 6 licenses of Windows Server and you only need to keep 6 servers up to date rather than 12.
- On the other hand, overloading several functions on a single server increases risk. For one thing, the more complicated a server is, the more likely it is to malfunction. And if you need to reboot the server or take the server down for repairs, you'll have downtime on multiple functions.

The trick is finding the right balance, and doing so is tricky. As a general rule, mission-critical functions such as Active Directory and email should be isolated on their own servers. Other server functions can be combined, as long as you take into consideration the risks and complexities of doing so.

Note that throughout this chapter, I'm ignoring the implications of virtualization. If you're running a virtual environment, you still need the same servers in place. Yes, virtualization makes it much easier to create more servers with fewer responsibilities on each server. In fact, that's one of the main reasons to virtualize your servers — if you don't, you have to purchase actual server hardware to bring up a new server. Even so, if a server function deserves to be isolated on its own server, it deserves its own server, regardless of whether you virtualize the server or run it on dedicated hardware. I look more closely at the question of virtualization in Book 3, Chapter 4.

Deciding Which Servers You Need

At the outset, you should make a list of the various servers you think your network will need. The following sections describe some of the basic server functions you're likely to require. You may be surprised how many servers you end up with!

Domain controllers

Your Active Directory infrastructure should have at least two domain controllers. You can technically get by with just one, but if it fails, your entire network will be

down until you get it fixed. Running two domain controllers provides an essential safety net.

In addition, you should isolate the domain controller function to single-purpose servers to minimize any required downtime. For example, if a domain controller doubles as a file server, you may have to take the server down to add disk space. That's a routine operation for a file server, but you don't want to impose the need to take down Active Directory just to increase disk space on a file server. Best to keep the domain controllers on separate servers.

Note that DNS and Active Directory are pretty much intertwined and dependent on one another, so it's common to enable DNS on a domain controller. In fact, Microsoft actually recommends that you combine DNS and Active Directory on your domain controllers, because Active Directory depends on DNS for name resolution.

DHCP servers

DHCP is a core service that is required for every network to run smoothly. Without DHCP, your users won't be able to connect to the network unless they all have static IP addresses, which is not a good idea, even on really small networks.

DHCP can run in several different locations, depending on your needs. For small networks, you can configure your router as the network's DHCP server. However, most DHCP implementations on routers aren't suitable for larger networks.

Technically, you can run DHCP on one of your domain controllers, but Microsoft recommends against it for two reasons:

- **Performance:** On a large network, DHCP can suck a lot of performance from a server, which can slow down Active Directory.
- **Security:** Running DHCP on a domain controller can potentially compromise the security of the domain controller.

So, if you can, I recommend you set up a separate server devoted to DHCP. And if you can't, double up DHCP on some other server rather than on a domain controller.

Mail servers

A *mail server* is a server that handles the network's email needs. It's configured with email server software, such as Microsoft Exchange Server. Exchange Server

is designed to work with Microsoft Outlook, the email client software that comes with Microsoft Office.

Most mail servers actually do much more than just send and receive email. For example, here are some of the features that Exchange Server offers beyond simple email:

- Calendaring and scheduling meetings
- Collaboration features that simplify the management of collaborative projects
- Audio and video conferencing
- Chat rooms and instant messaging (IM) services
- Microsoft Exchange Forms Designer, which lets you develop customized forms for applications such as vacation requests or purchase orders

Microsoft Exchange is a major piece of software that requires careful administration and, depending on how many users you have and how long you retain email, tons of disk storage. It is *always* a mistake to combine Exchange with any other server role — email should always be installed on its own dedicated server.

Many organizations are not installing Exchange on their own servers and instead using the Exchange Online feature of Office 365 for their email. In essence, this puts the burden of maintaining Exchange on Microsoft's Azure cloud services. For more information about this, refer to Book 5, Chapter 3 and Book 7, Chapter 2.

File servers

File servers provide centralized disk storage that can be conveniently shared by client computers on the network. The most common task of a file server is to store shared files and programs. For example, members of a small workgroup can use disk space on a file server to store their Microsoft Office documents.

File servers must ensure that two users don't try to update the same file at the same time. The file servers do this by *locking* a file while a user updates the file so that other users can't access the file until the first user finishes. For document files (for example, word processing or spreadsheet files), the whole file is locked. For database files, the lock can be applied just to the portion of the file that contains the record(s) being updated.

Most organizations will have at least one file server, and some may have many file servers to support different applications or departments.

Print servers

Although it isn't necessary, a server computer can be dedicated for use as a *print server,* the sole purpose of which is to collect information being sent to a shared printer by client computers and print it in an orderly fashion.

A single computer may double as both a file server and a print server, but performance is better if you use separate print and file server computers.

With inexpensive inkjet printers running about $100 each, just giving each user his or her own printer is tempting. But you get what you pay for. Instead of buying $100 printers for 15 users, you may be better off buying one high-speed $1,500 color laser printer and sharing it. The $1,500 laser printer will be much faster, will probably produce better-looking output, and will be less expensive to operate.

Better yet, lease a high-speed multifunction copier from a copier vendor. That way, the copier vendor will be responsible for keeping the beast working, and you'll be able to get a high-performance machine in the bargain. The printer function of a multifunction copier can be managed through a print server.

You can get by without setting up a print server. Instead, each user on your network can connect directly to the printer via its IP address. There are several disadvantages to that, however:

- **You have to manage drivers for each computer separately.** If you have 50 users connected to a printer and you need to update the driver, you'll have to update 50 computers.
- **Some users will inevitably mess up their print driver configuration.** You'll get called to fix it.
- **You lose overall control of the printer.** There is no centralized print queue and no ability to manage all your printers from a single point.

For more information about managing network printers, refer to Book 4, Chapter 5.

Web servers

A *web server* is a server computer that runs software that enables the computer to host an internet website. The two most popular web server programs are Microsoft's Internet Information Services (IIS) and Apache, an open-source web server managed by the Apache Software Foundation.

If you're going to use an internal web server to provide external users access to your corporate website, you need to carefully manage the security configuration of both the web server and your firewall to ensure that intruders can't use the website as a way to gain access to your entire network. For that reason, it's a good idea to host your website on a completely separate network if possible. Many companies use web hosting services for that purpose, so the web server for their company's website isn't part of their network at all.

However, it's very common to set up internal web servers for your company's intranet — that is, for web pages that are meant to be used within your company, not by users outside your company. If you intend to support a company intranet, you'll need to set up a separate web server for it.

Database servers

A *database server* is a server computer that runs database software, such as Microsoft's SQL Server 2022. Database servers are usually used along with customized business applications, such as accounting or marketing systems.

Like Exchange, SQL Server is a complicated enough piece of software that you should run it on a dedicated server.

Application servers

An *application server* is a server computer that runs a specific application. For example, you might use an accounting application that requires its own server. In that case, you'll need to dedicate a server to the accounting application.

Again, application servers are typically complicated enough and important enough to merit their own servers. It's not a good idea to bundle your accounting server with your print server; you don't want the entire accounting department calling your desk if you need to reboot the print server.

Backup servers

Depending on the backup software you use, you may need to provide a separate server that is devoted strictly to backing up your other servers. This is especially true if you back up to tape, as most tape devices don't connect via the network, but instead connect directly to a server. Isolating the important backup functions to a separate server is a great idea so backups don't interfere with other server processes, and vice versa.

For more information on different approaches to backing up your network, refer to Book 3, Chapter 6.

Deployment servers

A *deployment server* is a server dedicated to the task of automatically installing Windows images on network computers. You probably don't need this on a small network, but when your network gets to more than 50 or 100 computers, it's nice to have an automated way to deploy images to new computers or to redeploy images to computers that are having difficulty.

Microsoft provides support for this capability built in to Windows Server though a server role called Windows Deployment Services. Other companies that provide similar services that are more comprehensive and easier to use include Symantec Ghost and Acronis Snap Deploy. Search the web for more information about these tools, and consider setting up a deployment server if your network is large enough to merit it.

Update servers

An *update server* is a server devoted to managing updates to Windows computers. If you have just a few computers on your network, you can simply turn on automatic updates for the computers and let them update themselves. However, imagine how much network traffic will be wasted if you have 100 computers, all of them downloading updates directly from Microsoft's servers. With an update server, the updates are downloaded from Microsoft's servers to your update server. Then your computers are directed to your server for their updates, cutting down significantly on the internet traffic required to keep your computers up to date.

The simplest and easiest way to set up an update server is to use Windows Server Update Services (WSUS), a built-in component of all Windows Server operating systems. Simply devote a server computer to the task, install Windows Server, and then activate and configure the WSUS role.

An alternative is to use a third-party patch management tool such as Patch Manager Plus from ManageEngine (`www.manageengine.com`).

Virtualization management platform

If you're using a virtualizing platform (and you should be!), you'll need a server dedicated to managing all your virtual servers, network, and storage. For VMware, this management server is called vCenter. For Hyper-V, it's System Center Virtual Machine Manager. Either way, you'll want to devote one server to the management of your virtualization environment.

You'll find more information about planning for virtualization in the next chapter, and specific information about managing VMware and Hyper-V in Book 5, Chapters 1 and 2, respectively.

Connecting Your Servers

After you've determined just how many servers you need, it's time to figure out how you'll connect all those servers to your network.

If you aren't using virtualization, and instead you're implementing each of your servers as a separate physical server, you're bound to have a mess on your hands. It's generally a good idea to double up the network connection on each server for redundancy's sake (in other words, provide at least two paths to each server in case one path goes down). And it's also best to use the fastest connection speed possible for each server.

Figure 3-1 shows a simplified version of how you might connect seven servers (two domain controllers, two file servers, an Exchange server, a web server, and an update server) to a network. In this example, I'm using two core switches for all the servers and providing a separate connection from each server to each of the core switches. To keep the figure simple, I've omitted any access switches or end-user computers. As you can see, the figure is a rat's nest as it is.

Imagine how much more complex the diagram would be if there were 15 servers! One of the many reasons for using virtualization is to reduce the complexity of the network connections required to integrate all your servers into the network. I take a more detailed look at virtualization and what it has to offer in the next chapter.

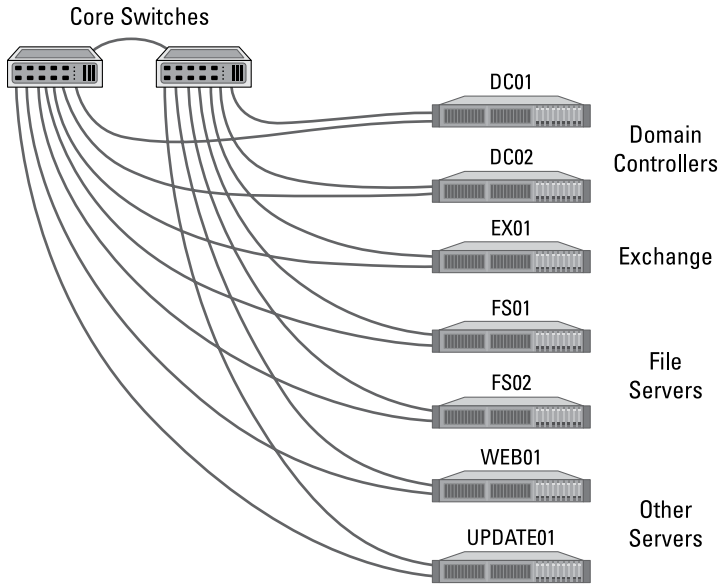


FIGURE 3-1: Connecting separate physical servers to multiple core switches can get complicated.

IN THIS CHAPTER

- Examining the basics of virtualization
- Looking at what a hypervisor does
- Considering how disks and networks are virtualized
- Weighing the benefits of virtualization
- Choosing host servers
- Considering how virtualization affects Microsoft licensing for Windows Server

Chapter 4
Virtualization Architecture

Virtualization is one of the hottest trends in networking today. According to some industry pundits, virtualization is the best thing to happen to computers since the invention of the transistor. If you haven't already begun to virtualize your network, you're standing on the platform watching as the train is pulling out.

This chapter introduces you to the basic concepts of virtualization, with an emphasis on using it to leverage your network server hardware to provide more servers using less hardware. Virtualization can dramatically simplify the design of your network — you can support more servers on less hardware, and with less hardware, your network will have fewer interconnects that link servers to the private network. Win, win!

Understanding Virtualization

The basic idea behind virtualization is to use software to simulate the existence of hardware. This powerful idea enables you to run more than one independent computer system on a single physical computer system. Suppose that your organization requires a total of 12 servers to meet its needs. You could run each of these 12 servers on a separate computer, in which case you would have 12 computers in your server room. Or, you could use virtualization to run these 12 servers on just two computers. In effect, each of those computers would simulate six separate computer systems, each running one of your servers.

Each of the simulated computers is called a *virtual machine* (VM). For all intents and purposes, each VM appears to be a complete, self-contained computer system with its own processor (or, more likely, processors), memory, disk drives, CD-ROM/DVD drives, keyboard, mouse, monitor, network interfaces, USB ports, and so on.

Like a real computer, each virtual machine requires an operating system to do productive work. In a typical network server environment, each virtual machine runs its own copy of Windows Server. The operating system has no idea that it's running on a virtual machine rather than on a real machine.

Here are a few terms you need to be familiar with if you expect to discuss virtualization intelligently:

- **Host:** The actual physical computer on which one or more virtual machines run. Admittedly, this term is kind of confusing, because it's also used to refer to any device that is connected to the network, such as an end-user computer. Context is everything — when discussing servers, *host* usually means the physical computer that virtual servers run on.
- **Bare metal:** Another term for the host computer that runs one or more virtual machines.
- **Guest:** Another term for a virtual machine running on a host.
- **Guest operating system:** An operating system that runs within a virtual machine. By itself, a guest is just a machine; it requires an operating system to run. The guest operating system is what brings the guest to life.



WARNING

As far as licensing is concerned, Microsoft treats each virtual machine as a separate computer. Thus, if you run six guests on a single host, and each guest runs Windows Server, you need licenses to run six servers. Unfortunately, figuring out how to ensure that you have the right number of licenses can be a bit complicated; see the section "Understanding Windows Server 2019 Licensing" later in this chapter for an explanation.

- **Hypervisor:** The virtualization operating system that creates and runs virtual machines. For more information about hypervisors, read the next section, "Understanding Hypervisors."
- **Hardware Abstraction Layer (HAL):** A layer of software that acts as a go-between to separate actual hardware from the software that interacts with it. An operating system provides a hardware abstraction layer, because it uses device drivers to communicate with actual hardware devices so that software running in the operating system doesn't have to know the details of the specific device it's interacting with. A hypervisor also provides a hardware abstraction layer that enables the guest operating systems in virtual machines to interact with virtualized hardware.

Understanding Hypervisors

At the core of virtualization is a *hypervisor*, a layer of software that manages the creation and execution of virtual machines. A hypervisor provides several core functions:

- It provides a HAL, which virtualizes all the hardware resources of the host computer on which it runs. This includes processor cores, RAM, and I/O devices such as disk drives, keyboards, mice, monitors, USB devices, and so on.
- It creates pools of these abstracted hardware resources that can be allocated to virtual machines.
- It creates virtual machines, which are complete implementations of an idealized computer system that has the hardware resources of the host available to it. The hardware for each virtual machine is drawn from the pools of available hardware resources managed by the hypervisor.
- It manages the execution of its virtual machines, allocating host hardware resources as needed to each virtual machine and starting and stopping virtual machines when requested by users.
- It ensures that each virtual machine is completely isolated from all other virtual machines, so that if a problem develops in one virtual machine, none of the other virtual machines is affected.
- It manages communication among the virtual machines over virtual networks, enabling the virtual machines to connect with each other and with a physical network that reaches beyond the host.

There are two basic types of hypervisors you should know about:

- **Type-1:** A type-1 hypervisor runs directly on the host computer, with no intervening operating system. This is the most efficient type of hypervisor because it has direct access to the hardware resources of the host system.

  The two best-known examples of type-1 hypervisors are VMware's ESXi and Microsoft's Hyper-V. ESXi is part of a suite of popular virtualization products from VMware, and Hyper-V is the built-in virtualization platform that is included with recent versions of Windows Server.

- **Type-2:** A type-2 hypervisor runs as an application within an operating system that runs directly on the host computer. Type-2 hypervisors are less efficient than type-1 hypervisors because when you use a type-2 hypervisor, you add an additional layer of hardware abstraction: the first provided by the operating system that runs natively on the host, and the second by the hypervisor that runs as an application on the host operating system.


TIP

For production use, you should always use type-1 hypervisors because they're much more efficient than type-2 hypervisors. Type-1 hypervisors are considerably more expensive than type-2 hypervisors, however. As a result, many people use inexpensive or free type-2 hypervisors to experiment with virtualization before making a commitment to purchase an expensive type-1 hypervisor.

THE LONG TREK OF VIRTUALIZATION

Kids these days think they invented everything, including virtualization.

Little do they know.

Virtualization was developed for PC-based computers in the early 1990s, around the time Captain Picard was flying the Enterprise around in *Star Trek: The Next Generation*.

But the idea is much older than that.

The first virtualized server computers predate Captain Picard by about 20 years. In 1972, IBM released an operating system called simply VM, which had nearly all the basic features found in today's virtualization products.

VM allowed the administrators of IBM's System/370 mainframe computers to create multiple independent virtual machines, each of which was called (you guessed it) a virtual machine, or VM. This terminology is still in use today.

Each VM could run one of the various guest operating systems that were compatible with the System/370 and appeared to this guest operating system to be a complete, independent System/370 computer with its own processor cores, virtual memory, disk partitions, and input/output devices.

The core of the VM system itself was called the *hypervisor* — another term that persists to this day.

The VM product that IBM released in 1972 was actually based on an experimental product that IBM released on a limited basis in 1967.

So whenever someone tells you about this new technology called *virtualization,* you can tell him or her that it was invented when *Star Trek* was on TV. When someone asks, "You mean the one with Picard?" you can say, "No, the one with Kirk."

Understanding Virtual Disks

Computers aren't the only things that are virtualized in a virtual environment. In addition to creating virtual computers, virtualization also creates virtual disk storage. Disk virtualization lets you combine a variety of physical disk storage devices to create pools of disk storage that you can then parcel out to your virtual machines as needed.

Virtualization of disk storage is nothing new. In fact, there are actually several layers of virtualization involved in any disk storage environment. At the lowest level are the actual physical disk drives. Physical disk drives are usually bundled together in arrays of individual drives. This bundling is a type of virtualization in that it creates the image of a single large disk drive that isn't really there. For example, four 2TB disk drives might be combined in an array to create a single 8TB disk drive.

Note that disk arrays are usually used to provide data protection through redundancy. This is commonly called RAID, which stands for *Redundant Array of Inexpensive Disks.*

One common form of RAID, called RAID-10, lets you create mirrored pairs of disk drives so that data is always written to both of the drives in a mirror pair. So, if one of the drives in a mirror pair fails, the other drive can carry the load. With RAID-10, the usable capacity of the complete array is equal to one-half of the total capacity of the drives in the array. For example, a RAID-10 array consisting of four 2TB drives contains two pairs of mirrored 2TB disk drives, for a total usable capacity of 4TB.

Another common form of RAID is RAID-5, in which disk drives are combined and one of the drives in the group is used for redundancy. Then, if any one of the drives in the array fails, the remaining drives can be used to re-create the data that was on the drive that failed. The total capacity of a RAID-5 array is equal to the sum of the capacities of the individual drives, minus one of the drives. For example, an array of four 2TB drives in a RAID-5 configuration has a total usable capacity of 6TB.

In a typical virtual environment, the host computers can be connected to disk storage in several distinct ways:

- **Local disk storage:** In local disk storage, disk drives are mounted directly into the host computer and are connected to the host computer via its internal disk drive controllers. For example, a host computer might include four 4TB disk drives mounted within the same chassis as the computer itself. These four drives might be used to form a RAID-10 array with a usable capacity of 8TB.

  The main drawbacks of local disk storage are that it's limited to the physical capacity of the host computers and is available only to the host computer that it's installed in.

- **Storage Area Network (SAN):** In a SAN, disk drives are contained in a separate device that is connected to the host via a high-speed controller. The difference between a SAN and local storage is that the SAN is a separate device. Its high-speed connection to the host is often just as fast as the internal connection of local disk storage, but the SAN includes a separate storage controller that manages the disk drives.

  A typical SAN can hold a dozen or more disk drives and can allow high-speed connections to more than one host. A SAN can often be expanded by adding one or more expansion chassis, which can contain a dozen or more disk drives each. Thus, a single SAN can manage hundreds of terabytes of disk data.

- **Network Accessible Storage (NAS):** This type of storage skips the high-speed storage controller and instead connects disk drives to host computers via TCP/IP over a standard Ethernet connection. NAS is the least expensive of all forms of disk storage, but it's also the slowest.

Regardless of how storage is attached to the host, the hypervisor consolidates its storage and creates virtual pools of disk storage typically called *data stores.* For example, a hypervisor that has access to three 6TB RAID-5 disk arrays might consolidate them to create a single 18TB data store.

From this data store, you can create *volumes*, which are essentially virtual disk drives that can be allocated to a particular virtual machine. Then, when an operating system is installed in a virtual machine, the operating system can mount the virtual machine's volumes to create drives that the operating system can access.

For example, let's consider a virtual machine that runs Windows Server. If you were to connect to the virtual machine, log in, and use Windows Explorer to look at the disk storage that's available to the machine, you might see a `C:` drive with a capacity of 200GB. That `C:` drive is actually a 200GB volume that is created by the hypervisor and attached to the virtual machine. The 200GB volume, in turn, is allocated from a data store, which might be 8TB in size. The data store is created from disk storage contained in a SAN attached to the host, which might be made up of a RAID-10 array consisting of four 4TB physical disk drives.

So, you can see that there are at least four layers of virtualization required to make the raw storage available on the physical disk drives available to the guest operating system:

- Physical disk drives are aggregated using RAID-10 to create a unified disk image that has built-in redundancy. RAID-10 is, in effect, the first layer of virtualization. This layer is managed entirely by the SAN.
- The storage available on the SAN is abstracted by the hypervisor to create data stores. This is, effectively, a second level of virtualization.
- Portions of a data store are used to create volumes that are then presented to virtual machines. Volumes represent a third layer of virtualization.
- The guest operating system sees the volumes as if they're physical devices, which can be mounted and then formatted to create usable disk storage accessible to the user. This is the fourth layer of virtualization.

Although it may seem overly complicated, these layers of virtualization give you a lot of flexibility when it comes to storage management. New disk arrays can be added to a SAN, or a new NAS can be added to the network, and then new data stores can be created from them without disrupting existing data stores. Volumes can be moved from one data store to another without disrupting their virtual machines. In fact, you can increase the size of a volume on the fly, and the virtual machine will immediately see the increased storage capacity of its disk drives, and without requiring a reboot.

Understanding Network Virtualization

When you create one or more virtual machines on a host system, you need to provide a way for those virtual machines to communicate not only with each other but also with the other physical computers on your network. To enable such connections, you must create a *virtual network* within your virtualization environment. The virtual network connects the virtual machines to each other and to the physical network.

To create a virtual network, you must create a *virtual switch*, which connects the virtual machines to each other and to a physical network via the host computer's network interfaces. Like a physical switch, a virtual switch has ports. When you create a virtual switch, you connect the virtual switch to one or more of the host computer's network interfaces. These interfaces are then connected with network cable to physical switches, which effectively connects the virtual switch to the physical network.

Then, when you create virtual machines, you connect each virtual machine to a port on the virtual switch. When all the virtual machines are connected to the switch, the VMs can communicate with each other via the switch. And they can communicate with devices on the physical network via the connections through the host computer's network interfaces.

Considering the Benefits of Virtualization

You might suspect that virtualization is inefficient because a real computer is inherently faster than a simulated computer. Although it's true that real computers are faster than simulated computers, virtualization technology has become so advanced that the performance penalty for running on a virtualized machine rather than a real machine is only a few percent.

The small amount of overhead imposed by virtualization is usually more than made up for by the simple fact that even the most heavily used servers spend most of their time twiddling their digital thumbs, waiting for something to do. In fact, many servers spend nearly *all* their time doing nothing. As computers get faster and faster, they spend even more of their time with nothing to do.

Virtualization is a great way to recoup all this unused processing power.

Besides efficiency, virtualization has other compelling benefits:

- **Hardware cost:** You typically can save a lot of money by reducing hardware costs when you use virtualization. Suppose that you replace ten servers that cost $4,000 each with one host server. Granted, you'll probably spend more than $4,000 on that server, because it needs to be maxed out with memory, processor cores, network interfaces, and so on. So you'll probably end up spending $15,000 or $20,000 for the host server. Also, you'll end up spending something like $5,000 for the hypervisor software. But that's still a lot less than the $40,000 you would have spent on ten separate computers at $4,000 each.

- **Energy costs:** Many organizations have found that going virtual has reduced their overall electricity consumption for server computers by 80 percent. This savings is a direct result of using less computer hardware to do more work. One host computer running ten virtual servers uses approximately one-tenth the energy that would be used if each of the ten servers ran on separate hardware.
- **Reduced downtime:** Virtual environments typically have less downtime than nonvirtual environments. For example, suppose you need to upgrade the BIOS on one of your server computers. With physical servers, this type of upgrade will ordinarily require that you shut down the operating system that runs on the server, upgrade the BIOS, and then restart the server. During the upgrade, the server will be unavailable.

  In a virtual environment, you don't need to shut down the servers to upgrade the BIOS on the host computer that runs the server. Instead, all you do is move the servers that run on the host that needs the upgrade to another host. When the servers are moved (an operation that can be done without shutting them down), you can shut down the host and upgrade its BIOS. Then, after you restart the host, you can move the servers back to the host — again, without shutting down the servers.
- **Recoverability:** One of the biggest benefits of virtualization isn't the cost savings, but the ability to recover quickly from hardware failures. Suppose that your organization has ten servers, each running on separate hardware. If any one of those servers goes down due to a hardware failure — say, a bad motherboard — that server will remain down until you can fix the computer. On the other hand, if those ten servers are running as virtual machines on two different hosts, and one of the hosts fails, the virtual machines that were running on the failed host can be brought up on the other host in a matter of minutes.

  Granted, the servers will run less efficiently on a single host than they would have on two hosts, but the point is that they'll all be running after only a short downtime.

  In fact, with the most advanced hypervisors available, the transfer from a failing host to another host can be done automatically and instantaneously, so downtime is all but eliminated.
- **Disaster recovery:** Besides the benefit of recoverability when hardware failures occur, an even bigger benefit of virtualization comes into play in a true disaster-recovery situation. Suppose that your organization's server infrastructure consists of 20 separate servers. In the case of a devastating disaster, such as a fire in the server room that destroys all hardware, how long will it take you to get all 20 of those servers back up and running on new hardware? Quite possibly, recovery time will be measured in weeks.

By contrast, virtual machines are actually nothing more than files that can be backed up onto tape. As a result, in a disaster-recovery situation, all you have to do is rebuild a single host computer and reinstall the hypervisor software. Then you can restore the virtual-machine backups from tape, restart the virtual machines, and get back up and running in a matter of days instead of weeks.

Choosing Virtualization Hosts

Having made the decision to virtualize your servers, you're next faced with the task of selecting the host computers on which you'll run your virtual servers. The good news is that you need to purchase fewer servers than if you use physical servers. The not-so-good news is that you need to purchase really good servers to act as hosts, because each host will support multiple virtual servers. Here are some tips to get you started:

- **If possible, purchase at least two hosts, and make sure that each host is independently capable of running all your virtual servers.** That way, if one of the hosts goes down, you can temporarily move all your servers to the good host while the bad one is being repaired. When both hosts are up, you can spread the workload across the two hosts for better performance.
- **Add up the amount of memory you intend to allocate for each server to determine the amount of RAM for each host.** Then give yourself plenty of cushion. If your servers will require a total of 50GB of RAM, get 72GB on each host, for a total of 144GB if you have two hosts. That will give you plenty of room to grow.
- **Do a similar calculation for processor cores.** Like most computers, servers spend an enormous percentage of their time idling. Virtualization makes very efficient use of processor cores for a large number of servers.
- **Get the best network connections you can afford.** Ideally, each host should have a pair of small form-factor pluggable (SFP) ports that you can run 10 Gb fiber over. That way, your hosts can communicate with the core switches at top speed.
- **Provide redundancy in the host's subcomponents.** Most hosts support two processors, two memory banks, two network interfaces, and two power supplies. That provides for a maximum of uptime.

Understanding Windows Server 2025 Licensing

When planning your server architecture, you'll need to account for the fact that you must purchase sufficient licenses of Windows Server to cover all the servers you're running. Before virtualization, this was easy: Each server required its own license. With virtualization, things get tricky — and Microsoft doesn't make it easier by trying to simplify things.

Windows Server 2025 comes in three editions. These editions are as follows:

- **Standard Edition:** Ideal for customers who aren't virtualized or who are virtualized but have a relatively small number of server instances (approximately 12 per host). Each Standard Edition license allows you to run two instances of Windows Server. These licenses are sold based on the number of cores on each physical server, in increments of two or eight. The two-pack costs $243 at the time of this writing. There's no quantity discount for the eight-pack — it costs $972, which is four times the cost of the two-pack. And each physical server must have a minimum of eight core licenses. (Note that these pricings are based on the pricing for Windows Server 2022. Windows Server 2025 was not available at the time I wrote this. Don't be surprised if the price changes.)
- **Datacenter Edition:** This edition is used for larger organizations with more than a few dozen server instances. The Datacenter Edition license costs more than the Standard Edition license ($6,155 for the eight-pack, $1,538.75 for the two-pack) but lets you run an unlimited number of server instances.
- **Essentials Edition:** Designed for small businesses setting up their first server. It's limited to just 25 users. I won't consider this edition further.

Here's where it gets complicated. Each server instance must be licensed for all the cores available on the physical server that will host the server instance. With Datacenter Edition, this isn't a consideration because the license allows for an unlimited number of server instances. But with Standard Edition, you're limited to two server instances per license. So, you need to purchase additional licenses to cover additional server instances beyond two.

For example, suppose you have a host computer that has two 12-core CPUs, for a total of 24 cores on the host and that you want to run a total of ten server instances on this host. Because each license lets you run two server instances, you need five licenses for all 24 cores. So, you need to purchase licenses for 120 cores (5×24), which amounts to 15 eight-packs. At $972 per eight-pack, that works out to a total of $14,580.

At this point, you can see that the Standard Edition is cheaper than the Datacenter Edition because, although you only need one Datacenter license to run ten server instances, you need to purchase three eight-packs of the Datacenter license to cover all 24 cores. At $6,155 each, that works out to $18,465.

It turns out that the break-even point for deciding when to purchase Standard Edition versus Datacenter Edition is at around 12 server instances per physical server. With 12 instances, Standard Edition is just a tad cheaper than Datacenter Edition. At 13 instances, it's the other way around.



TIP

The break-even point doesn't depend on the number of cores per physical server: It's still between 12 and 13 server instances per physical server. The bottom line is that if you anticipate running more than a dozen server instances per physical server, you should consider Datacenter Edition.

Clearly, Microsoft charges more to run Windows Server on more powerful hosts. Soon enough, you'll be hard pressed to purchase hosts that fall below the single-license core limit of eight cores per processor or two processors per host. That's because Intel's dual-socket Xeon processors are getting more and more cores with each successive generation. The current generation of Xeon processors sports up to a whopping 64 cores per processor. While Intel still makes 4-core, 8-core, and 16-core versions of the Xeon processor, who knows what the future will bring.

IN THIS CHAPTER

- Knowing how much is enough
- Figuring out what kind of disks you should use
- Deciding on a drive interface
- Thinking about RAID
- Focusing on attachment types

Chapter 5
Storage Architecture

There's never enough storage, right? What may seem huge today will be laughably small in just a few years. Disk storage is like closet space: No matter how big the space, people find a way to fill it up.

In this chapter, I look at the most important things you need to consider when planning the storage side of your network. I cover various technologies for providing that disk storage, including various types of disk drives, different drive interfaces, and several ways to connect storage to your servers.

Planning Disk Capacity

The first basic decision you need to make when planning your network storage is the most obvious: How much storage do you need?

Unfortunately, that's not an easy decision to make. You can easily add up how much disk storage you currently have, and you can calculate how much is actually in use and how much free capacity you have. But predicting the future is hard.

One thing is for certain: The increase in your disk usage is not a linear function. In other words, you can't calculate some number — let's randomly pick 5TB — and assume that your company will use that much additional storage every year for the next five years. If it were true, your life would be simple. You'd just need to provide 25TB of free space to last for five years.

Unfortunately, experience suggests that disk usage is an exponential function, not a linear function. It's more like, "Every few years we need twice as much as we needed just a few years ago."

Experience bears that out. Most of us old folks remember the first 10MB hard disks in 1981. Soon it was 100MB, then gigabytes. Now terabytes. When will it be petabytes? Over the course of my career, disk capacity on my desktop has increased a million-fold, at a rate of about 1.4 times per year!

The point is, don't underestimate how much your disk needs will grow. Make sure your network plan can accommodate growth.

Here are a few general rules:

- **Plan on about 150GB of disk space for the root drives of all your servers, and allow for twice as many servers as you currently have.** So if you have ten servers now, plan on 3TB of disk space to support server root drives (20 servers times 150GB each).
- **If possible, choose expandable disk subsystems, and don't load them up to capacity.** Make sure the disk subsystem you purchase can be expanded to at least double its current capacity, either by adding additional drives or adding additional enclosures to accommodate more drives.

  Also, make sure you leave some room in your rack (or racks) to accommodate additional storage devices, and plan your rack layout in a way that leaves empty space below your current disk subsystem.
- **Be wary of the desire to increase capacity at the risk of performance or reliability.** For example, always use RAID to provide the protection of redundancy for your server storage. And resist the temptation to use RAID 5 rather than RAID 10 for critical data simply because RAID 5 gives you more space. Opt for more drives instead, and use the best level of protection.
- **Don't be tempted by huge drive capacities.** At the time I wrote this, the largest hard disk commonly available was 26TB. For safety, you could put two of them in a mirrored array for a total of 26TB of redundant storage. But I wouldn't do it. That's just too much data to commit to a single drive. I'd rather build an array from 4TB or 8TB drives to spread the risk and simplify recovery when a failure occur (and it will!).
- **Don't put too much faith in manufacturer claims of the benefits of tricks like compression or deduplication.** These techniques can (and do) work, but not always at the rate that the manufacturers claim.
- **Don't neglect data retention policies and archiving strategies.** They can help keep unneeded files off your server storage.

Considering Disk Drive Types

The next basic decision in planning your disk storage is deciding what type of disk drives to use. There are two basic types of storage to choose from:

Hard disk drives

Hard disk drives (HDDs), also known as *spinning drives*, are traditional magnetic disk drives. Capacities of modern HDDs range anywhere from 500GB to 6TB or more.

HDDs include mechanical components such as the motor that spins the disk platters and the servos that move the read/write heads over the spinning platters to read and write data.

The performance of an HDD depends in large part on how fast the disk platters spin. Disk speed is measured in revolutions per minute (RPM), with three speeds being common: 7.2K, 10K, and 15K. The higher-RPM drives have better performance because the read/write heads must wait less time for data to arrive under the heads. In addition, when the data does arrive at the read/write heads, it can be read or written faster because the magnetic medium of the disk platter travels past the heads at a higher rate of speed.

Higher-RPM drives are also more expensive than slower-spinning drives because greater engineering care is needed to safely spin the platters at higher speeds.

Solid state drives

Solid state drives (SSDs) are all-electronic devices with no moving parts. They're based on memory technology, and they're considerably faster than HDDs. They're also considerably more expensive, and they tend to have smaller capacity — typically 100GB to 4TB, thought some have a capacity of 8TB or more.

SSD storage is significantly faster than HDD storage because no moving parts are involved. SSD technology is on the order of a thousand times faster than HDD technology, though that doesn't necessarily mean a given SSD drive is 1,000 times faster than a given HDD; many other factors combine to determine the overall performance of a storage device. Still, SSD is several orders of magnitude faster than HDD storage. (HDD access speed is measured in milliseconds — thousands of a second — while SSD storage is measured in microseconds — *millionths* of a second.)

FORM FACTORS

Form factor refers to the size of the disk drives you will use. Both HDDs and SSDs come in two basic form factors: 3.5-inch, called *LFF* (for *large form factor*), and 2.5-inch, called *SFF* (for *small form factor*). Because 3.5-inch disk drives are larger, they have potentially higher capacity. At the time I wrote this, the largest 3.5-inch enterprise-class HDD drives held 10TB.

The smaller 2.5-inch drives have smaller capacity (the maximum is currently 2.4TB). However, more 2.5-inch drives can be accommodated in a single enclosure. Typically, a rack enclosure of a given size can hold twice as many 2.5-inch drives as 3.5-inch drives.

SSD storage devices are based on *flash memory*, similar to the memory that is used in USB flash drives, but more reliable and considerably faster.

SSDs are considerably more expensive than HDDs of similar capacity, so for the time being, HDD is more likely to fit within your budget. Most networks include a combination of both SSDs and HDDs, reserving SSDs for data in which the speed benefit of HDD outweighs the price penalty.

The fastest available solid-state storage attaches directly to the motherboard via a special connection called M.2 (pronounced "Em-dot-two"). M.2 storage is popular on laptop computers and desktop computers, but is also becoming available on server computers. Most M.2 drives have a capacity of 1TB or 2TB, but 4TB and even 8TB are available.

Considering Drive Interfaces

Another factor to consider when planning your storage environment is which drive interfaces to use. The drive interface manages the connection between the disk drive itself and the control unit that the drive is attached to. In a desktop or laptop computer, the disk controller is built into the motherboard, and it's almost always the first variety, known as *SATA.* In a network server, the disk controller is often a separate card installed into the server's chassis — or, sometimes, in a separate chassis. In that case, either the SATA interface or the more advanced SAS interface can be used.

The following paragraphs describe the differences between SATA and SAS.

SATA

SATA is the most popular interface for consumer devices. It's an evolution of the original disk interface that was used when hard drives were first introduced on IBM PCs. That interface was originally called *IDE,* which stood for *integrated device electronics.* That was soon replaced by an improved interface called *ATA,* which stood for *AT attachment* because it was designed to work with IBM's PC-AT line of personal computers.

The original IDE and ATA interfaces were *parallel interfaces,* which meant that they transmitted and received 16 bits of data at a time. This arrangement required a total of 40 separate wires on the cables that connected the disk drives to the controllers, and complicated circuitry that kept the data synchronized on all the wires.

Parallel interfaces were increasingly difficult to keep up with increasing disk transfer speeds, so IDE and ATA evolved into a serial interface called *SATA,* which stands for *serial ATA.* In a serial interface, data is transmitted one bit at a time. Intuitively, that sounds less efficient than transmitting data 16 bits at a time, but in reality it's possible to send and receive data much faster using serial transmission than using parallel transmission because of the difficulty of keeping parallel transmission lines in sync.

Today, SATA is used on nearly all desktop and laptop computers and on many low-end server computers. Most SATA disks can transmit data at 6 Gbps (6 billion bits per second).

You also need to be aware that there are actually two classes of SATA disk devices: consumer and enterprise. Consumer-class SATA disks are found in desktop and laptop computers and are the least expensive disk drives available. Enterprise-class SATA drives are preferred for server storage, because they're about ten times as reliable as consumer-class drives. They're a bit more expensive, but the additional cost is well worth it.

SAS

SAS is the preferred drive interface for network storage. It's an evolution of an older drive interface called *SCSI,* which stands for *small computer system interface.* Like IDE and ATA, the original SCSI interface was a parallel interface. *SAS* is the serial version of SCSI; it stands for *Serial Attached SCSI.* (Incidentally, *SCSI* is pronounced "scuzzy.")

The SAS interface is faster than the SATA interface. Most SAS devices transfer data from the disk to the controller at either 6 Gbps or 12 Gbps.

The ability to work at 12 Gbps is one of the main benefits of SAS or SATA, but reliability is another important factor: Enterprise-class SAS drives are about ten times more reliable than enterprise-class SATA drives. (Because enterprise SATA is about ten times more reliable than consumer-class SATA, that makes enterprise-class SAS about 100 times more reliable than consumer-class SATA.)

Other than price, performance, and reliability, there's not much practical difference between SATA and SAS. But because performance and reliability are important considerations for network storage, I recommend you go with 12 Gbps SAS drives whenever your budget will allow.

Considering RAID

Reliability is one of the most important considerations when planning your network storage. All disk devices will eventually fail. This includes SSDs as well as HDDs. In fact, SSDs and HDDs have about the same reliability; both fail at about the same rate.

As a general rule, about 2 percent of your disk drives will fail every year. So if you have 25 disk drives in your server room, you can expect one to fail about every two years. If you have 100 disk drives, expect one to fail every six months. You can do the math: Drive failures are not uncommon.

Fortunately, we have ways to survive disk drive failures. The first line of defense is to use *RAID,* which groups disk drives together into arrays that have built-in redundancy and automatic recovery when one of the drives in an array fails.



RAID stands for either *Redundant Array of Inexpensive Disks* or *Redundant Array of Independent Disks,* depending on who you talk to. Either way, the idea of RAID is to group disks together to provide redundant data storage.

Although there are many different types of RAID configurations, only three are in widespread use: RAID 10, RAID 5, and RAID 6.

RAID 10

In a RAID 10 array, the disks in the array are paired into mirror sets, in which both disks in each set contain the same data. Whenever data is written to one disk in a set, the exact same data is written to the other disk. Thus, if either of the two disks in the set fails, the other disk in the set has a backup copy of the data.

For example, suppose a RAID 10 array has six drives with a capacity of 1TB each. The array has a total of 6TB of disk storage, but because the drives are paired into mirror sets, only 3TB of data can be stored on the array. If any one drive fails, nothing is lost — the data can be retrieved from the surviving drive in the mirror set. When the drive that failed is replaced, the array can heal itself by copying all the data from the surviving disk to the replacement disk.

RAID 10 is generally considered the safest form of RAID, but it's vulnerable to a loss of two disks in the array. If two disks fail at the same time, only luck will determine whether the entire array is lost. If the failing disks are in separate mirror sets, the array will survive. But if both disks in a single mirror set are lost, the entire array will be lost.

RAID 5

In a RAID 5 array, multiple disks are combined into a single array, but the equivalent of one disk's worth of space is set aside for redundancy. (The redundancy data is actually spread across all the disks in the array, but the total amount of disk space needed for the redundancy is equivalent to one full disk in the array.)

If any disk in the array fails, the contents of that disk can be recovered to a new disk by calculating the data that was on the failed disk using the data that is on the surviving disks.

The usable capacity of the array is one drive less than the total number of drives in the array. For example, if you create a RAID 5 array using six 1TB drives, the usable capacity of the array will be 5TB. The sixth terabyte is used for redundancy.

The basic principle of how a RAID 5 array works is actually pretty simple to understand. Suppose I tell you to write down a list of five numbers. For example:

22, 37, 16, 81, 53

If I were to then erase one of the numbers at random, could you reconstruct the list? Not unless you have a really good memory!

But if you know in advance that I might erase one of the numbers, there's an easy trick that will help you recover the erased number: Just add all the numbers up, and write down the sum:

22 + 37 + 16 + 81 + 53 = 209

Now, if I erase any of the original five numbers, you can easily figure out what it was by subtracting the surviving four numbers from the sum.

That's essentially how RAID 5 works. The math is a bit more complicated than that, but the principle is the same. (It's also worth noting that RAID 5 doesn't simply designate one of the drives in the array to hold all the calculated redundancy data; instead, the redundancy data is spread across all the drives in the array.)

RAID 5 is more efficient than RAID 10 in terms of disk capacity. For example, a RAID 10 array of six 1TB drives has a usable capacity of just 3TB, while a RAID 5 array of the same six 1TB drives has a usable capacity of 5TB.

But from a performance perspective, RAID 5 is considerably slower than RAID 10 when writing data to the disk. To write data to a RAID 5 array, first the redundancy data must be calculated. Then both the data initially to be written as well as the redundancy data must be written to the array. The RAID 5 is less efficient because of the calculation and because of the need for multiple writes.

And finally, when one of the drives in a RAID 5 array fails, the array will take much longer to rebuild than when a drive in a RAID 10 array fails. In a RAID 10 rebuild, data from the surviving mirror pair is simply copied to the replacement drive. But in a RAID 5 rebuild, all the data from all the surviving drives must be read. Then the data for the replacement drive must be written. In our six 1TB-drive array examples, recovering a RAID 10 array requires that 1TB of data be read and 1TB of data be written. But for a RAID 5 array, 5TB of data must be read, 1TB of data must be calculated, and 1TB of data must be written.

In short, the rebuild of a failed RAID 5 array often requires several *days*.



WARNING

In fact, many experts and most disk drive manufacturers recommend against RAID 5 altogether because of how long it takes to rebuild a failed drive. The problem with RAID 5 is that disk drive capacity has increased much faster than disk drive speed. We've been stuck at 6 Gbps or 12 Gbps for many years now, but disk capacity has soared. That means that rebuild times for RAID 5 arrays have also soared. Unfortunately, there's a not unreasonable chance that a *second* drive in a RAID 5 will fail during a rebuild operation. If that happens, the entire array will be lost.

You may think it unlikely that a second drive failure will happen during a rebuild, but keep in mind that most RAID arrays are populated with disk drives that were purchased at the same time from a single manufacturer. There's a good chance all the drives came from a single manufacturing lot, will have a similar expected lifetime, and have about the same amount of usage on them. The odds are better than you think. Because of this, most experts now recommend you use RAID 6 instead of RAID 5, as explained in the next section.

RAID 6

RAID 6 is one step more secure than RAID 5. Instead of calculating one set of redundancy data for the entire array, in RAID 6 two sets of redundancy information are calculated. Effectively, two of the disks in the array are set aside for redundancy. This allows the array to survive the loss of any two disks in the array, not just a single disk.

Of course, RAID 6 imposes a greater space penalty than RAID 5. A RAID 6 array of six 1TB drives will have a usable space of 4TB. In addition, RAID 6 is a bit slower than RAID 5 because two sets of redundancy data must be calculated rather than just one.

But RAID 6 is considerably safer than RAID 5.

Considering Attachment Types

To be useful, disk storage must be attached to your servers. It will do you no good to populate your rack with terabytes of disk storage if your users can't access it!

The following sections describe four basic approaches to attaching storage to your servers.

Direct attached storage

Direct attached storage (DAS) is the simplest and most obvious way to attach storage to a server. With DAS, storage is directly connected to a hard disk controller within the server. This provides the fastest possible connection to the computer, but it's also the most limited because the storage can be used only by the computer to which it is directly attached.

In a normal workstation computer, the hard disk controller is on the motherboard, and the disk drive or drives are mounted inside the computer's case in internal drive bays. In a typical rack-mounted server computer, drive bays for DAS are also built into the case, but they're usually accessible from the front of the server and they're usually hot-swappable, which means they can be removed and replaced while the server is powered up.

Most server computers need at least a small amount of DAS installed directly in the server chassis. You can use this storage for the server's operating system — or, if you use virtualization, for the server's hypervisor. Typically, a pair of 72 or 100GB SAS drives in a RAID 10 array are appropriate.

It is generally *not* a good idea to install large amounts of storage directly into a server chassis, because that storage will be accessible only to that server. This doesn't mean that you can't use DAS for large amounts of storage to be shared by several host servers; it just means you shouldn't install that storage in the chassis of one of the servers. Instead, you can use an external storage subsystem that has the ability to directly attach to more than one host. Such systems can typically be attached to anywhere from two to four host servers. The attachments are usually made with external SAS cables. This arrangement requires external SAS adapters on both the host servers and the external storage subsystem.

Figure 5-1 shows how two host servers might be connected to a single storage subsystem. In this case, the storage subsystem contains three enclosures that each hold 12 disks, so a total of 36 disk drives are available via this subsystem. Each of the two host computers has an external SAS adapter that is used to connect to the disk subsystem via external SAS cables.

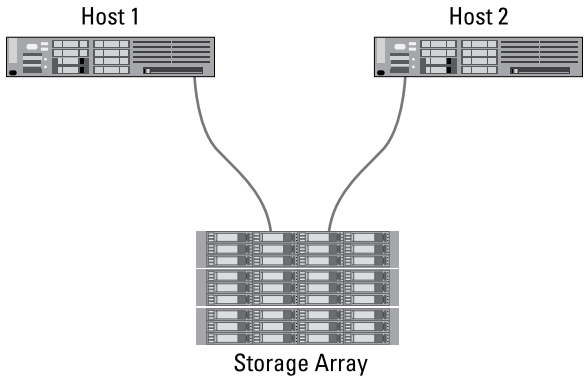


FIGURE 5-1: Two hosts directly attached to a storage subsystem.

Storage area networks

A *storage area network* (SAN) is used when the number of storage devices or host computers makes it impossible to directly connect the storage to the hosts. Instead, a separate network of storage devices is created using a networking technology called *Fibre Channel*. Fibre Channel is similar in many ways to other networking technologies such as Ethernet, but it's designed specifically for connecting huge numbers of storage devices to servers. Fibre Channel networks can support thousands of storage devices.

Fibre Channel is also very fast, with top speeds of up to 128 Gbps. However, most Fibre Channel networks run at a more modest 16 Gbps. Fibre Channel usually operates over fiber-optic cable, but it can also run on copper cable at slower speeds.

Like Ethernet, Fibre Channel relies on switches to interconnect storage devices and hosts. Figure 5-2 shows a small Fibre Channel network in which six hosts are connected to three storage subsystems via a Fibre Channel switch. The cables and connectors for this network are 16 Gbps fiber-optic.

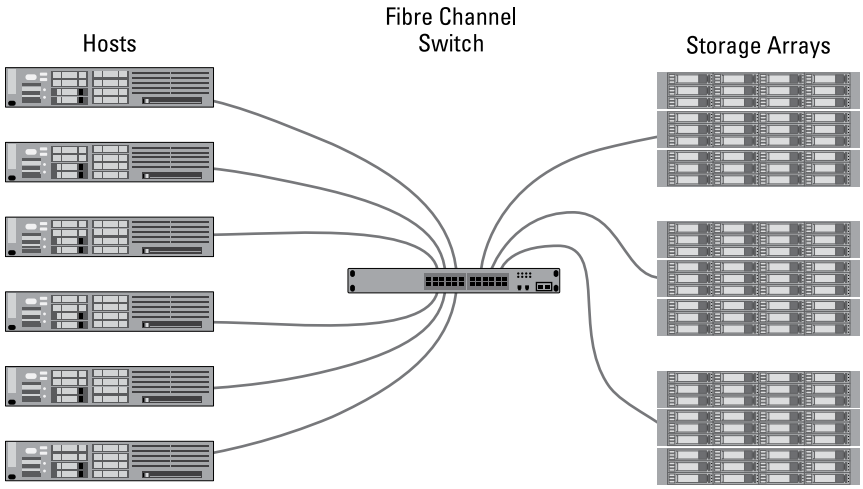


FIGURE 5-2: A Storage Area Network.



Why the British spelling of *Fibre* rather than the American spelling *Fiber*? Originally, the American spelling was used, and Fiber Channel networks could be implemented only using fiber-optic cables. When copper cabling was added to the specification, the spelling was changed to the British variant just for fun.

Network-attached storage

One final form of attaching storage in a network is called *network-attached storage* (NAS). When NAS is used, storage devices are connected directly to the existing Ethernet network and data is accessed over TCP/IP using a variety of protocols that enable normal disk and file handling operations to be encapsulated in IP packets. NAS is one of the easiest ways to add large amounts of storage to a network, but NAS doesn't have nearly the performance that SAN or DAS does. In effect, data accessed via NAS devices is limited to the speed of the underlying network, which is typically 1 Gbps. Figure 5-3 shows how a NAS device can be attached to a network.

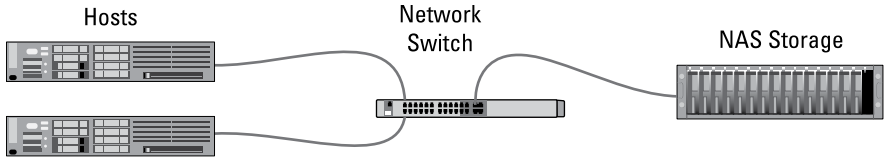


FIGURE 5-3: A NAS system connected to a network.

The most common form of NAS consists of appliance-like devices that are essentially a small computer running as a file server with a large amount of disk storage. Users can access data on a NAS appliance as if it were any other file server on the network. The NAS appliances usually have a web-based administrative console that can be used to set up shares, manage permissions, and so on.



TIP

When you incorporate NAS into your overall storage plan, be sure to account for the backup and recovery requirements of the NAS. It's temptingly easy to add inexpensive terabytes of NAS storage to your network to satisfy your users' increasing appetite for storage. But don't forget that if the data is important enough to save on the network, it's important enough to back up on a regular basis. Your users will be sorely disappointed if they lose data they thought was safely ensconced on NAS if you fail to incorporate it into your backup plans.



WARNING

Another issue to be concerned about with NAS is that it can just randomly appear on your network when a user decides to stop at Best Buy on the way to work one day. Inexpensive, consumer-quality NAS is readily available and can easily be plugged in to any available network port. Keep on the lookout for rogue NAS devices.

IN THIS CHAPTER

- Understanding the need for backups
- Working with tape drives and other backup media
- Understanding the different types of backups
- Finding alternatives for backing up in a virtual environment

Chapter 6
Backup Architecture

One of the most important aspects of designing a network is planning for the backup of data that will be housed on the network's servers. It's best to incorporate a backup architecture into your plans from the very beginning, instead of waiting until after everything else is in place before thinking about how to back it all up.

This chapter introduces you to the most common approaches to backup so that you can develop a backup strategy that will work for your network.

If you're thinking that the topic of backup isn't really a networking topic, think again. Server data capacity is getting bigger and bigger all the time. Not too long ago, data capacity was measured in gigabytes. Now, data is measured in terabytes — thousands of gigabytes. And the term *petabyte* is starting to show up — a petabyte (PB) is a thousand terabytes.

If the servers on your network house 10TB, 20TB, or 100TB of data, you definitely need to take that into consideration when you plan your network. After all, that data has to move over the network to be backed up. The design of your network is influenced by, and can influence, the design of your backup architecture.

Backup Basics

Having data backed up is absolutely essential to any computer system. It is the cornerstone of any disaster recovery plan. Without backups, a simple hard drive failure can set your company back days or even weeks while it tries to reconstruct lost data. In fact, without backups, your company's very existence is in jeopardy.


REMEMBER

The main goal of a backup is existential: Keep a spare copy of your network's critical data so that no matter what happens to your computer systems, your company can survive. Cybercriminals may ransom your data, your server room may burn down, or Disney may decide to bring Jar Jar Binks back. But as long as you have a good backup architecture and stay on top of it, you'll be able to recover from these or even worse disasters.

A secondary goal of backup is to protect your users from themselves. Users will be users — they'll accidentally delete important files, or they'll outright trash a file. You or your Help Desk will get plenty of requests from users to restore a file from a recent version, such as "Can you restore xyz.xls to the way it was last Thursday?" With a good backup architecture, you'll be able to say, "Sure." Or, better yet, "Sure, but it'll cost you a donut."

Conventional wisdom along with common sense says that you should always have at least three backup copies of all essential data:

- **At least one *local* copy, which is a part of your local network and can be accessed quickly:** This backup is the go-to source when an individual file needs to be recovered and is also the backup of choice if an entire server needs to be recovered.
- **At least one *off-site* copy, which is kept at a remote location:** This is essential so that if your servers are damaged because of a physical event such as a flood, a fire, or vandalism, you'll have a copy of everything at another location. The remote location must be secure, and it should be far enough away that it won't be affected by a regional disaster. In other words, across the street isn't far enough.
- **At least one *off line* copy, which is kept completely isolated from the internet and from your network:** This is required because skilled cybercriminals have been known to delete or ransom all of an organization's online data plus all of its backups. The only way to prevent this nightmare scenario is to keep at least one set of backups completely offline, where a hacker can't touch without physical access.

Note that I said you should keep *at least* one copy of each of these three types of backups. Another conventional-wisdom paradigm of backups is to keep three *generations* of backups, typically known as *Grandfather*, *Father*, and *Son*, or *GFS*. The idea is to have three different sets of backups, differentiated by age. For example, one set that was made last night, another that was made two days ago, and a third that was made three days ago. (I'd like to lobby the backup industry to switch this terminology to *GPC*, for *Grandparent, Parent, and Child.*)

Most often, the multigenerational backup scheme maintains a certain number of daily backups (usually five to cover one workweek, or seven if your business is open on the weekends), a certain number of weekly backups (perhaps four sets, to cover one month), and a certain number of monthly backups (perhaps 3, 6, or 12).

So, if you keep 5 daily backups, 4 weekly backups, and 6 monthly backups, you'll have a total of 15 copies of your data backed up.

You may use a different generational scheme for the local, off-site, and offline backups. For example, you might devise a scheme something like this:

- **Local backups:** 10 daily, 4 weekly, and 3 monthly
- **Off-site backups:** 4 weekly
- **Offline backup:** 4 weekly and 3 monthly

There's no right or wrong to this scheme — the point is that you have a strategy that incorporates both the separation of local/off-site/offline and multiple generations in a way that takes into account your organization's work cycles and risk tolerance.

Considering Three Basic Types of Backup

There are two fundamentally different types of backup you can choose for your organization:

- **File-based backup:** The oldest type of backup software backs up the files on your computer one at a time. With this type of software, you can be very selective about exactly which files to back up. For example, you can choose to skip all files with the extension `.bak`, because these are backup files and you don't really need to back up backup files. Or, you can choose to omit specific folders or even entire disk drives from the backup. The drawback is that if a computer contains a lot of files (say, in the millions — not uncommon for a

large file server), the backup program will spend an inordinate amount of time just keeping track of all the files.

- **Image-based backup:** This type of backup software makes copies of entire disk images. The best-known of these is Acronis (`www.acronis.com`). With an image-based backup, it doesn't matter how many files are on a disk; the backup software makes an exact copy of the entire disk, copying data block by block. This is much faster than a file-based backup, not just for backing up but also for restoring data.

An important variation of image-based backups is backup programs that specialize in backing up virtual machines (VMs). In a virtual environment, virtual machines generally store virtual disk drives as a single large file on the actual hardware disk. Thus, virtual-based backup programs back up entire virtual machines by making copies of all the files required to store the VM, which is usually just a few. You can find out more about how these programs work in the section "Understanding Image-Based Backups and Virtualization" later in this chapter.

Where to Back Up Your Data

If you plan on backing up the data on your network server's hard drives, you obviously need some type of media to serve as the destination for your backup data. This is true whether you'll use file-based backups, image-based backups, or VM-aware backups.



TIP

In backup lingo, the destination for your backups is called the *backup target.*

Here are the four most common media options:

- **Tape:** Magnetic tape is the oldest storage medium for backups and is still one of the most widely used. One big advantage of tape backups is that tape cartridges are small and can thus be easily transported to an off-site location. For more information, see the section "Backing Up to Tape" later in this chapter.
- **Network-attached storage (NAS):** A *Network Attached Storage* device connects directly to your network. NAS devices are often used as backup devices because they're relatively inexpensive. Note, however, that the most inexpensive NAS devices have 1GB network connections, and it can take a very long time to back up terabytes of data over a 1GB connection. 10 Gbps NAS units are more expensive but will back up your data much faster.
- **Network-attached backup appliances:** A *network-attached backup appliance* is similar in many ways to a NAS, but it's specifically designed for backup. NAS

devices are designed to work as file servers and can be used for backups, but a purpose-built backup appliance will give better performance and capacity for your backups. Examples of backup appliances include Exagrid (`www.exagrid.com`) and HPE's StoreOnce (`www.hpe.com/us/en/storage/storeonce.html`).

- **Cloud backup:** An increasingly popular option is to use a third-party service to back up data to a remote location via the internet. Cloud backup has the advantage of already being off-site, but keep in mind that it's online, meaning that a very skilled hacker can in theory wipe out your cloud backups.

Backing Up to Tape

Tape backup is one of the oldest and most reliable forms of backup there is. Computers have been using tape for data storage since 1951 (that's over 70 years!), but obviously a lot has changed since the days of reel-to-reel tape drives often seen in old science-fiction movies. The most common format for modern tape storage is the LTO Cartridge.

Looking closer at LTO

LTO has been around since 2000 and has been through several backward-compatible generations. Originally, two physical cartridge formats were planned, but only one — known as *Ultrium* — was ever manufactured. Thus, you'll often see the name *Ultrium* on LTO tape cartridges.

LTO stands for *Linear Tape-Open*, which is important only in that the *Open* means that the format is an open standard, so different vendors can make drives and tape cartridges, which are compatible with one another. Table 6-1 lists all the generations of LTO tape. The current generation is LTO-9; LTO-10 through LTO-12 are still in the planning stages. Note that as a general rule, each generation doubles the capacity of the preceding generation.

Notice that the capacity of LTO tape cartridges has increased at a faster rate than the speed at which data can be written. Between the first generation (LTO-1) and the current generation (LTO-9), capacity has increased by a factor of 184 (that is, the capacity of an LTO-9 tape is 184 times the capacity of an LTO-1 tape), but the speed has increased by a factor of only 20. Capacity has increased almost ten times faster than speed.

TABLE 6-1 LTO Tape Generations

Generation	Year Released	Raw Capacity	Compressed Capacity	Speed (MB/sec)
LTO-1	2000	100GB	200GB	20
LTO-2	2003	200GB	400GB	40
LTO-3	2005	400GB	800GB	80
LTO-4	2007	800GB	1.6TB	120
LTO-5	2010	1.5TB	3.0TB	140
LTO-6	2012	2.5TB	6.25TB	160
LTO-7	2015	6TB	15TB	300
LTO-8	2017	12TB	30TB	360
LTO-9	2020	18TB	45TB	400
LTO-10	2025	36TB	90TB	1,100
LTO-11	?	72TB	180TB	?
LTO-12	?	144TB	360TB	?
LTO-13	?	288B	720B	?
LTO-14	?	578TB	1,440TB	?

What that means is that it takes a lot more time to fill up an entire tape. Although an LTO-1 tape could be filled in less than 90 minutes, it takes more than 12 hours to fill an LTO-9 tape. The amount of time it takes to fill a tape can become a significant factor in how you design your backup architecture.

Hardware for tape backup

At a minimum, you'll need a tape drive to write backups to tape. An LTO tape drive consists of the following internal and external components:

- A carrier that can hold the tape while it's being written to or read from and can also eject the tape when necessary.
- A read/write head that reads data from and writes data to the tape.
- A motor that spins the tape reel inside the cartridge so that the tape moves across the read/write head. The motor is also responsible for rewinding the tape when all its data has been read.

- A tape take-up reel. The tape cartridge itself has just one reel for the tape. Therefore, the tape drive itself includes a take-up reel for the tape that is spun off the cartridge's reel as the tape is read or written.
- Optionally, a bar code reader that can read a bar code printed on a label that's attached to the cartridge. This helps the tape drive verify that the correct tape is inserted.
- A controller interface to connect the drive to a computer. Usually, the interface is Serially Attached SCSI (SAS) or Fibre Channel (FC).

Single-tape LTO drives are fairly inexpensive and are available in desktop models. But if you need to back up more than a single tape of data, you'll need an automated tape changer, usually called a *tape library.* A tape library consists of the following internal components:

- One or more *tape magazines,* each of which can hold several tape cartridges.
- One or more tape drives. When multiple tape drives are available, each drive can read or write tapes simultaneously. Thus, a library with two drives can write tape backups twice as fast as a stand-alone tape drive or a library with just one drive.
- A robotic mechanism to automatically move tapes from the magazines to the drives and vice versa.

The greatest advantage of a tape library is that it can change tapes automatically, so no manual intervention is needed when a tape fills up. But the ability to process multiple tapes simultaneously is also a big advantage, especially when many terabytes of data need to be backed up.

A word about tape reliability

From experience, I've found that although tape drives are very reliable, they do run amok once in a while. The problem is that they don't always tell you when they're not working. A tape drive can spin along for hours, pretending to back up your data — but in reality, your data isn't being written reliably to the tape. In other words, a tape drive can trick you into thinking that your backups are working just fine. Then, when disaster strikes and you need your backup tapes, you may just discover that the tapes are worthless.


TIP

Don't panic! Here's a simple way to assure you that your tape drive is working. Just activate the "compare-after-backup" feature of your backup software. As soon as your backup program finishes backing up your data, it rewinds the tape, reads each backed-up file, and compares it with the original version on the hard drive. If all files compare, you know that your backups are trustworthy.

Here are some additional thoughts about the reliability of tapes:

- The compare-after-backup feature doubles the time required to do a backup, but that doesn't matter if your entire backup fits on one tape. You can just run the backup after hours. Whether the backup and repair operation takes one hour or ten doesn't matter, as long as it's finished by the time the network users arrive at work the next morning.
- If your backups require more than one tape, you may not want to run the compare-after-backup feature every day. Be sure to run it periodically, however, to check that your tape drive is working.
- If your backup program reports errors, throw away the tape, and use a new tape.
- Actually, you should ignore that last comment about waiting for your backup program to report errors. You should discard tapes *before* your backup program reports errors. Most experts recommend that you should use a tape only about 20 times before discarding it. If you use the same tape every day, replace it monthly. If you have tapes for each day of the week, replace them twice yearly. If you have more tapes than that, figure out a cycle that replaces tapes after about 20 uses.

About cleaning the heads

An important aspect of backup reliability is proper maintenance of your tape drives. Every time you back up to tape, little bits and specks of the tape rub off onto the read and write heads inside the tape drive. Eventually, the heads become too dirty to read or write data reliably.

To counteract this problem, clean the tape heads regularly. The easiest way to clean them is to use a cleaning cartridge for the tape drive. The drive automatically recognizes when you insert a cleaning cartridge and then performs a routine that wipes the cleaning tape back and forth over the heads to clean them. When the cleaning routine is done, the tape is ejected. The whole process takes only about 30 seconds.

Because the maintenance requirements of drives differ, check each drive's user's manual to find out how and how often to clean the drive. As a general rule, clean drives once weekly.

The most annoying aspect of tape drive cleaning is that the cleaning cartridges have a limited life span, and unfortunately, if you insert a used-up cleaning cartridge, the drive accepts it and pretends to clean the drive. For this reason, keep

track of how many times you use a cleaning cartridge and replace it as recommended by the manufacturer.

Backing Up to NAS

Book 3, Chapter 5 covers network-attached storage (NAS). In short, NAS is storage that is not attached directly to a server but is instead connected to the network. Any other device on the network can access the storage provided by the NAS, as long as proper credentials are used. Because it's relatively inexpensive, NAS is a popular target for backups.



WARNING

The only caveat for using NAS as a backup target is to consider the speed of the NAS device's network connection. Most NAS devices have a 1GB network interface that can be connected basically anyplace on your network. But data measured in terabytes is being transferred, so a 1GB connect can be painfully slow. You can count on about five hours per terabyte.

You'll be much better off if you can use a 10 Gb connection for your NAS backup target. That will cut the transfer time to about half an hour per terabyte, but of course the 10 Gb interface for the NAS and the 10 Gb cable will increase the cost.

Using a Backup Appliance

Backup appliances are devices that are specifically designed to work as backup targets. They're similar to NAS devices in that they contain storage and are connected to the network, but they have features that make them vastly superior to NAS — and, of course, more expensive. Most backup appliances rely on the technique of *deduplication* to dramatically reduce the storage space required to store your backup data.

Deduplication works by eliminating blocks of data that are identical. In a typical backup environment, deduplication can reduce actual disk space required by surprisingly large factors. That's because most backup architectures save multiple time-based copies of the same data. For example, deduplication can save ten daily backup copies of the same data in just a little more storage than a single daily copy would require. That's because only a small fraction of the data on most file servers actually changes on a daily basis. Deduplication ratios of 20:1 are not uncommon.

Deduplication is a compute-heavy operation, though. That's where a deduplicating backup appliance will beat a NAS every time: The backup appliance has built-in hardware and software that's designed to optimize deduplication.

One of my favorite backup appliances, Exagrid (`www.exagrid.com`), takes a unique approach to optimizing deduplication even more: It doesn't perform deduplication while the backup is running. Instead, it deduplicates the data after the backup has finished.

To accomplish this, Exagrid carves the total storage of the appliance into two regions. The first region is called the *landing zone.* This is where backup software writes its backup data to. The second region is called the *retention zone.* After data has been written to the landing zone, the Exagrid device begins the process of moving that data to the retention zone, applying its deduplication algorithm as it goes.

As a result, deduplication takes place behind the scenes, after the backup is completed. By the time the next backup is scheduled to run, all the data in the landing zone will have been moved to the retention zone, so the backup has the full space of the landing zone to use as its backup target.

Exagrid comes in a variety of sizes to fit your needs, starting at 12TB and ranging up to 168TB. In each model, half of that space is used for the landing zone. Thus, the 12TB model can accommodate backups of up to 6TB. You'll be able to fit dozens of copies of that 6TB backup into the retention zone, where deduplication comes into play.

And Exagrid devices can be expanded in a *cluster,* in which all the storage across all devices in the cluster is aggregated. Thus, two of the 168TB models can be combined to provide 336TB of backup storage.

Understanding File-Based Backup

File-based backup relies on programs that run on individual computers or on backup programs that run on a central server but deploy agents to each computer that needs to be backed up. Either way, the backup program backs up files one at a time. The main advantage of this type of backup program is that you have a lot of flexibility over exactly what files, folders, and volumes should be backed up. The drawback is that processing files individually significantly slows down the backup program.

All versions of Windows come with a built-in file-based backup program. In addition, most tape drives come with backup programs that are often faster or more flexible than the standard Windows backup.

You can also purchase sophisticated backup programs that are specially designed for networks that have multiple servers with data that must be backed up. For a basic Windows file server, you can use the backup program that comes with Windows Server. Server versions of Windows come with a decent backup program that can run scheduled, unattended tape backups.

File-based backup programs do more than just copy data from your hard drive to tape. Backup programs use special compression techniques to squeeze your data so that you can cram more data onto fewer tapes. Compression factors of 2:1 are common, so you can usually squeeze 100GB of data onto a tape that would hold only 50GB of data without compression. (Tape drive manufacturers tend to state the capacity of their drives by using compressed data, assuming a 2:1 compression ratio. Thus, a 200GB tape has an uncompressed capacity of 100GB.)



WARNING

Whether you achieve a compression factor of 2:1 depends on the nature of the data you're backing up:

- **Documents:** If your network is used primarily for Microsoft Office applications and is filled with Word and Excel documents, you'll probably get better than 2:1 compression.
- **Graphics:** If your network data consists primarily of graphic image files, you probably won't get much compression. Most graphic image file formats are already compressed, so they can't be compressed much more by the backup software's compression methods.

Backup programs also help you keep track of which data has been backed up and which hasn't. They also offer options, such as incremental or differential backups, that can streamline the backup process, as I describe in the next section.

You can perform five different types of backups. Many backup schemes rely on full daily backups, but for some networks, using a scheme that relies on two or more of these backup types is more practical.

The differences among the five types of backups involve a little technical detail known as the archive bit. The *archive bit* indicates whether a file has been modified since it was backed up. The archive bit is a little flag that's stored along with the filename, creation date, and other directory information. Any time a program

modifies a file, the archive bit is set to the On position. That way, backup programs know that the file has been modified and needs to be backed up.

The differences among the various types of backups center on whether they use the archive bit to determine which files to back up, as well as whether they flip the archive bit to the Off position after they back up a file. Table 6-2 summarizes these differences, which I explain in the following sections.

TABLE 6-2 How Backup Types Use the Archive Bit

Backup Type	Selects Files Based on Archive Bit?	Resets Archive Bits After Backing Up?
Normal	No	Yes
Copy	No	No
Daily	No*	No
Incremental	Yes	Yes
Differential	Yes	No

** Selects files based on the Last Modified date.*



TIP

Backup programs allow you to select any combination of drives and folders to back up. As a result, you can customize the file selection for a backup operation to suit your needs. For example, you can set up one backup plan that backs up all a server's shared folders and drives, plus its mail server stores, but then leaves out folders that rarely change, such as the operating system folders or installed program folders. You can then back up those folders on a less-regular basis. The drives and folders that you select for a backup operation are collectively called the *backup selection.*

The archive bit would have made a good Abbott and Costello routine. ("All right, I wanna know who modified the archive bit." "What." "Who?" "No, What." "Wait a minute . . . just tell me what's the name of the guy who modified the archive bit!" "Right.")

Full backups

A *full backup* is the basic type of backup. In a full backup, all files in the backup selection are backed up regardless of whether the archive bit has been set. In other words, the files are backed up even if they haven't been modified since the last

time they were backed up. When each file is backed up, its archive bit is reset, so backups that select files based on the archive bit setting won't back up the files.

When a full backup finishes, none of the files in the backup selection has its archive bit set. As a result, if you immediately follow a full backup with an incremental backup or a differential backup, files won't be selected for backup by the incremental or differential backup because no file will have its archive bit set.

The easiest backup scheme is to simply schedule a full backup every night. That way, all your data is backed up on a daily basis. Then, if the need arises, you can restore files from a single tape or set of tapes. Restoring files is more complicated when other types of backups are involved.


REMEMBER

Do full backups nightly if you have the tape capacity to do them unattended — that is, without having to swap tapes. If you can't do an unattended full backup because the amount of data to be backed up is greater than the capacity of your tape drive(s), you have to use other types of backups in combination with full backups.


TIP

If you can't get a full backup on a single tape, and you can't afford a second tape drive or a tape changer, take a hard look at the data that's being included in the backup selection. I recently worked on a network that was difficult to back up onto a single tape. When I examined the data that was being backed up, I discovered a large amount of static data that was essentially an online archive of old projects. This data was necessary because network users needed it for research purposes, but the data was read-only. Even though the data never changed, it was being backed up to tape every night, and the backups required two tapes. After I removed this data from the cycle of nightly backups, the backups were able to squeeze onto a single tape again.

If you remove static data from the nightly backup, make sure that you have a secure backup of the static data on tape, CD-RW, or some other media.

Copy backups

A *copy backup* is similar to a full backup except that the archive bit isn't reset when each file is copied. As a result, copy backups don't disrupt the cycle of normal and incremental or differential backups.

Copy backups usually aren't incorporated into regular, scheduled backups. Instead, you use a copy backup when you want to do an occasional one-shot backup. If you're about to perform an operating system upgrade, for example, you should

back up the server before proceeding. If you do a full backup, the archive bits are reset, and your regular backups are disrupted. If you do a copy backup, however, the archive bits of any modified files remain unchanged. As a result, your regular normal and incremental or differential backups are unaffected.

If you don't incorporate incremental or differential backups into your backup routine, the difference between a copy backup and a full backup is moot.

Daily backups

A *daily backup* backs up just those files that changed the same day that the backup was performed. A daily backup examines the modification date stored with each file's directory entry to determine whether a file should be backed up. Daily backups don't reset the archive bit.


WARNING

I'm not a big fan of this option because of the small possibility that some files may slip through the cracks. Someone may be working late one night and modify a file after the evening's backups have completed — but before midnight — meaning that those files won't be included in the following night's backups. Incremental or differential backups, which rely on the archive bit rather than the modification date, are more reliable.

Incremental backups

An *incremental backup* backs up only those files that were modified since the last time you did a backup. Incremental backups are a lot faster than full backups because your network users probably modify only a small portion of the files on the server on any given day. As a result, if a full backup takes three tapes, you can probably fit an entire week's worth of incremental backups on a single tape.

When an incremental backup copies each file, it resets the file's archive bit. That way, the file will be backed up again before your next full backup only when a user modifies the file again.

Here are some thoughts about using incremental backups:


TIP

» **The easiest way to use incremental backups is the following:**

- A *full backup* every Monday

  If your full backup takes more than 12 hours, you may want to do it on Friday so that it can run over the weekend.

- An incremental backup on each remaining normal business day (for example, Tuesday, Wednesday, Thursday, and Friday)

- **When you use incremental backups, the complete backup consists of the full backup tapes and all the incremental backup tapes that you've made since you did the full backup.**

  If the hard drive crashes, and you have to restore the data onto a new drive, you first restore Monday's full backup and then restore each of the subsequent incremental backups.

- **Incremental backups complicate restoring individual files because the most recent copy of the file may be on the full backup tape or on any of the incremental backups.**

  Backup programs keep track of the location of the most recent version of each file to simplify the process.


TECHNICAL STUFF

- **When you use incremental backups, you can choose whether you want to**
  - Store each incremental backup on its own tape.
  - Append each backup to the end of an existing tape.

  Often, you can use a single tape for a week of incremental backups.


TIP

Differential backups

A *differential backup* is similar to an incremental backup except that it doesn't reset the archive bit when files are backed up. As a result, each differential backup represents the difference between the last full backup and the current state of the disk.

To do a full restore from a differential backup, you first restore the last full backup and then restore the most recent differential backup.

Suppose that you do a full backup on Monday and differential backups on Tuesday, Wednesday, and Thursday, and your hard drive crashes Friday morning. On Friday afternoon, you install a new hard drive. To restore the data, you first restore the full backup from Monday. Then you restore the differential backup from Thursday. The Tuesday and Wednesday differential backups aren't needed.

The main difference between incremental and differential backups is that

- *Incremental* backups result in smaller and faster backups.
- *Differential* backups are easier to restore.

  If your users often ask you to restore individual files, consider using differential backups.


TIP

Understanding Image-Based Backups and Virtualization

If your servers are virtualized using either VMware or Hyper-V, you may want to consider adopting an altogether different approach to backups. Instead of creating complicated schemes of weekly full backups and daily incremental backups that are based on backing up the hundreds of thousands (or even millions) of individual files on all your servers, a virtual backup solution can focus instead on backing up the files that represent entire virtual machines.

Virtualization presents an entirely different set of challenges and opportunities for backup and recovery. In a virtualized environment, each server is represented not by hundreds of thousands (or even millions) of files on a physical disk drive, but by just a few files that represent the contents of the entire server. These files are very large, but software exists that allows you to easily and quickly replicate these files onto other media.

Virtualization platforms such as VMware and Hyper-V have built-in capabilities to manage this replication, but you can also purchase third-party solutions that can turn this replication capability into a full-fledged backup solution. For example, the Swiss-based company Veeam (`www.veeam.com`) has a powerful backup solution that is specifically designed for virtual environments. With Veeam, you can do full and incremental backups of virtual machines in a way that lets you recover either individual files or entire machines. One of the best features of Veeam is that you can run a virtual server directly from a backup image, without the need to first do a time-consuming restore. This can cut your recovery time from hours to minutes. And, while continuing to run the machine from the backup image, you can simultaneously restore the machine to its primary media. After the restore is completed, Veeam will automatically switch over to the restored copy of the machine.

If your environment is virtualized, I definitely recommend you investigate options such as Veeam for your backup solution instead of using traditional backup methods.

Backup Security

Backups create an often-overlooked security exposure for your network: No matter how carefully you set up user accounts and enforce password policies, if any user (including a guest) can perform a backup of the system, that user may make

an unauthorized backup. In addition, your backup tapes themselves are vulnerable to theft. As a result, make sure that your backup policies and procedures are secure by taking the following measures:

- **Set up a user account for the user who does backups.** Because this user account has backup permission for the entire server, guard its password carefully. Anyone who knows the username and password of the backup account can log on and bypass any security restrictions that you place on that user's normal user ID.
- **Counter potential security problems by restricting the backup user ID to a certain client and a certain time of the day.** If you're really clever (and paranoid), you can probably set up the backup user's account so that the only program it can run is the backup program.
- **Use encryption to protect the contents of your backups.** Any backup that includes sensitive data should be encrypted.
- **Secure the backup tapes in a safe location, such as, um, a safe.**

IN THIS CHAPTER

- Understanding hyperconvergence
- Examining how deduplication works
- Integrating backups into your HCI system
- Incorporating HCI into your network plan

Chapter 7

Hyperconverged Infrastructure

One of the newer trends in IT architecture is called *hyperconverged infrastructure* (HCI). HCI is an attempt to solve some of the most vexing problems in IT infrastructure, including the complexity of managing explosive growth in capacity and performance requirements. I don't know about you, but it seems like just a few years ago, 100GB was a lot of data. Now 100GB is just a spoonful of data, and terabytes are the norm. Managing dozens or hundreds of terabytes of data with traditional virtualization based on host servers and storage area network (SAN) storage arrays is getting more and more difficult.

In this chapter, I present just an overview of what HCI is and how it can impact your overall IT architecture.

Considering the Headaches of Traditional IT Architecture

Hyperconvergence isn't just hype: It's designed to solve real-world IT problems. Before I look at the details of what HCI is and how it works, I want to review some of the basic issues that HCI addresses:

- **Storage is difficult to manage.** SAN storage is pretty flexible when it comes to accessing it from multiple servers, but it's still constrained by the limits of physical disks. For example, suppose you start out with a SAN consisting of eight 2TB disk drives in a RAID-5 array, yielding a total usable capacity of 14TB. You soon fill that up, so you add another eight disks, creating a second RAID-5 array with a usable capacity of 14TB. So, you've nearly doubled the capacity of the SAN from 14TB to 24TB. But unfortunately, this doesn't give you a single, uninterrupted span of 24TB of disk to the host servers. Instead, it presents two 12TB arrays. As a result, you'll have to arrange all your virtual disks so they'll fit into these two bins. This task can become like the game of Tetris, trying to fit the various virtual disks into the bins of storage provided by your SAN. It's not uncommon to find yourself in a position where you have plenty of free disk space, but not in the right place.
- **The SAN is itself a network that needs to be managed.** Sometimes the term SAN is used to describe something that technically isn't a SAN — it's actually a SAN-compatible disk controller that's connected to a small number of host servers using SAS connections. In a true SAN, there's a separate network for connecting the disk controllers to each other and to the servers. This network typically uses fiber channel connections and requires fiber channel switches. The storage network is separate from and not accessible to the local area network (LAN). The SAN is just one more level of complexity that must be configured, managed, upgraded, and expanded as needs grow.
- **SAN has limited expansion.** Most SAN products are expandable, meaning you can add additional disks to empty slots and, if you're out of free slots, you can add expansion shelves with 12 or more slots for additional disk. But eventually you'll reach the expansion limit of your SAN and you'll need an additional SAN. That just intensifies the silo problem I mention earlier.
- **Performance is difficult to optimize.** Most SAN arrays consist of a large number of spinning disks or hard disk drives (HDDs). You may have a few solid-state drives (SSDs) for better performance, which forces you to make decisions about which data to store on fast SSDs and which to store on slower HDDs. You can use special caching software that automatically places the most frequently used data on the SSD storage, but that's complicated to set up and administer; plus, it's just one more thing to maintain.

- **Storage outgrows backup capacity.** Disk storage is relatively cheap, but if you keep adding terabytes of storage, you need to figure out how to back it all up. A backup architecture that was designed to back up 10TB of data can be difficult to expand to 20TB, not just because of the capacity of your backup targets but also because of the time it takes to back up so much data.
- **Deduplication is often more trouble than it's worth.** I cover deduplication in more detail later in this chapter, but for now just realize that it works by examining all the data blocks on a disk drive to find blocks that are identical. Then, deduplication eliminates the redundant data blocks, often resulting in space savings of 50 percent or more. The problem with deduplication in a traditional storage environment is that every time you move data from one place to another, you must undo the deduplication (that process is called *hydration*), and then deduplicate the data again when the data arrives at its destination. This process can be very time consuming and can also put you in a capacity bind when you discover that you don't have enough storage to hydrate your deduplicated data.

Defining Hyperconverged Infrastructure

Hyperconvergence to the rescue! HCI is designed to address the problems described in the previous section as well as other problems. It does this by consolidating the three major components of a typical SAN installation — the servers, the storage, and the storage network — into a single component called an *HCI Node.*

Figure 7-1 is a simplified depiction of a small SAN installation. Here, two host processors are connected to two storage devices though a fiber channel switch. Each of the storage devices has a disk controller and two shelves of disks populated with disk drives. The host servers each include compute resources (CPU and RAM) plus network adapters — both Ethernet adaptors to connect to the LAN and fiber channel adapters to connect to the SAN. (Note that the host servers also contain disk storage, but just enough to hold the hypervisor so that the server can run virtual machines.)

Figure 7-2 shows how this equipment can be replaced by HCI. Here are the most obvious differences between the two figures:

- The host servers have been replaced by *HCI nodes,* which are HCI appliances that are typically based on the same hardware platform as a host server but that include additional software and possibly a special hardware card to enable HCI functions.

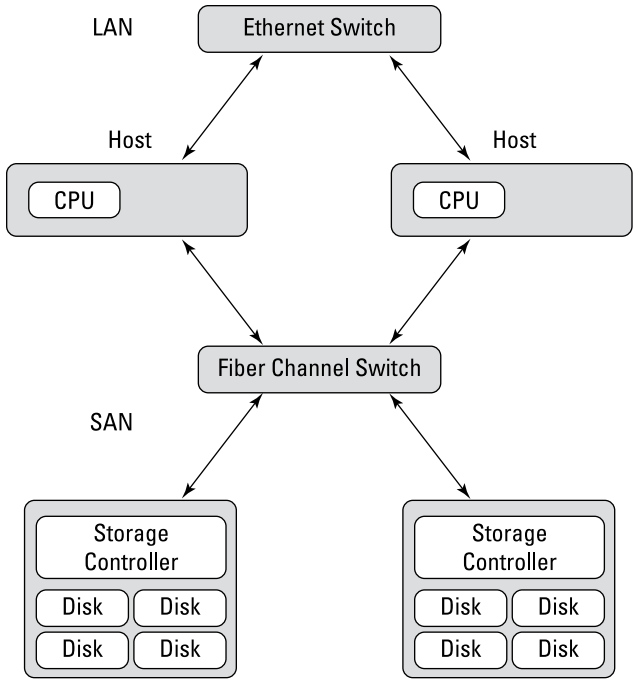


FIGURE 7-1: Traditional storage with separate compute, storage, and a storage network.

- The separate storage devices are gone. Instead, each of the HCI nodes is populated with disk drives that will provide storage for the HCI environment. In other words, the compute resources of the host servers and the storage resources of the SAN have been consolidated into the HCI nodes. In most cases, the storage placed in the HCI nodes is flash storage (SSDs) to provide maximum performance.
- The fiber channel switch is gone. A separate storage area network (SAN) is not needed, so the fiber channel elements can be eliminated. Note that because there is no fiber channel network, disk performance is faster in HCI. Directly attached storage as used by HCI is always faster than even the best tuned SAN.
- The two HCI nodes are combined to create an *HCI cluster,* which can be managed and used as if the two nodes were a single entity.

In short, an HCI node is a server computer that is outfitted with ample disk storage, special software, and possibly special hardware to consolidate storage and compute resources into a single appliance. HCI nodes can be combined into HCI clusters, which operate as a unified system.

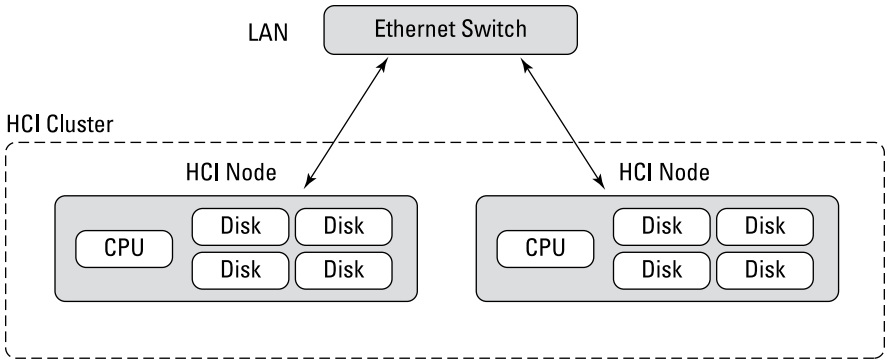


FIGURE 7-2: HCI integrates compute, storage, and the storage network.

Discerning Deduplication

One of the best features about HCI is that it depends heavily on built-in deduplication to massively expand the capacity of the disk storage built into the nodes. For example, a large HCI node might contain 40TB of SSD. That alone may seem like a lot of storage, but when deduplication is employed, often 80TB or more of data can be stored on the node. Here are a few things to note about deduplication in an HCI system:

- **In HCI, deduplication isn't a feature that can be turned on or off; it's an inherent part of the system.** Everything stored on an HCI node is deduplicated.
- **Depending on the make and model of the HCI node, a special deduplication accelerator card may be installed in the node to provide the compute resources needed for deduplication.** That can result in an enormous performance advantage, because deduplication is a compute-intensive task. The accelerator card relieves the main CPU from the burden of deduplication so that it can run the workloads of the virtual machines assigned to the node.
- **In HCI, deduplication is applied to the entire storage resource, not just to the individual virtual hard drives that are allocated on the storage.** For example, suppose a traditional host runs ten servers, each with a separate operating system volume of 100GB. Deduplication can be applied to these operating system volumes, but it won't have much of an effect. However, in HCI, deduplication is applied at the storage level, not the volume level. Because the ten 70GB operating system volumes will likely be nearly identical, HCI can store them in about 10 percent of the space required to store them on a traditional host.

Understanding How Deduplication Works

There are two basic types of deduplication technology:

- **In-line:** In-line deduplication performs its deduplication magic as data is written to the disk. In other words, when in-line deduplication is used, duplicate blocks are never written to the disk. This results in the best savings of disk space, but writing data to disk is slowed down by the deduplication process.
- **Post-process:** In contrast, post-process deduplication performs its magic *after* data is written to the disk. Data is written to the disk without checking for duplication. Then, later, a deduplication process scans blocks on the disk, looking for and removing duplicate blocks. This results in faster disk writes, at the cost of disk capacity.

Windows Server has built-in post-process deduplication that can be used for disk volumes. This feature is often enabled on file servers to increase capacity.

Most HCI solutions use in-line deduplication so that data is immediately deduplicated as it's written to disk. To improve the performance of in-line deduplication, HCI solutions often use SSDs instead of hard disk drives (HDDs) and often use dedicated hardware to assist with the deduplication process.



TECHNICAL STUFF

Deduplication may seem like magic, but it isn't. It's just an algorithm, like everything else in computing. Here's a simplified overview of how it works:

1. When a block of data is written to disk, the deduplication system creates a *hash* of the block by applying a *hashing function* to the block of data. Data blocks are typically 4K, but some HCI systems use longer 8K blocks. The hash values are typically 20 bytes long (160 bits). (If you want to know more about hashing functions, see the sidebar "Hashing for fun and profit.")
2. The hash values for all blocks that are written to the disk are kept in memory in a *hash table,* which records not only the hash value but the disk location of the actual data block that corresponds to the hash.
3. If a hash value calculated for a block of data to be written to the disk already exists in the hash table, the deduplication software does a bit-by-bit comparison of the block to be written with the block already stored for the hash value. This is done to make sure the data blocks are actually identical. Hash functions are not perfect; although it's very unlikely, it is possible that two different blocks of data may produce the same hash value.

4. Assuming the data blocks are, indeed, identical, the new block is not written to the disk. Instead, a link to the corresponding entry in the hash table is written to the disk. Then, when reading data blocks from the disk, if a link is encountered, the hash table is used to find the actual data block to be retrieved.

5. If the comparison in Step 3 indicates that the blocks are different, the new block is written to the disk and the hash table is updated to reflect that the hash value represents more than one actual data block value. (Links subsequently written for the hash indicate which version of the data block should be used.)

Okay, maybe the process isn't all that simple after all. The implementation details are actually more complicated than the preceding description lets on. But the basic idea is that if a data block to be written is identical to a block that's already on the disk, the data block isn't written a second time; instead, a link to the existing block is written. The hash table that helps identify identical blocks and keeps track of where they are is kept in memory so it can be accessed quickly.


REMEMBER

Because the deduplication process requires a lot of computation — applying hash values, hash table lookups, comparing blocks that may be identical to ensure they really are, and so on — most HCI products use advanced hardware to mitigate the performance hit. For starters, SSD storage is usually used instead of slower HDD storage. And in many cases, a separate card in the HCI appliance is used to handle the deduplication process. This deduplication card contains its own CPU and RAM so that the task of deduplication doesn't tax the main CPU and RAM that's used to run the server's virtual machines.

HASHING FOR FUN AND PROFIT

The term *hash* in computing refers to an algorithm that converts data of an arbitrary size to a value of a fixed size. As used in data deduplication, the hash function typically converts 4K data blocks to 160-bit values.

A *hashing algorithm* is a mathematical process that can be used to create a hash value. There are hundreds of different hashing algorithms in use. To be considered effective, a hashing algorithm should be

- **Repeatable:** For a given input value, the hash function must always produce the same output.
- **Uniform:** The hash function should produce hash values that are evenly distributed across the output size.

(continued)

(continued)

- **Efficient:** It shouldn't take an enormous amount of compute time to calculate hash values.

The simplest possible way to generate a 160-bit hash from a 4K data block would be to simply use the first 20 bytes of the data block (160 bits equals 20 bytes). That's a bad way to do it, however, because it's very likely that a large number of distinct 4K data blocks would have the identical values for the first 20 bytes. So, something more sophisticated is required.

You could also just use every 200th bit in the 4K data block. That would give you 160 bits selected from across the entire data block. That may not be too bad, but what about data blocks that have a bunch of zeros at the end of the block?

Clearly something more sophisticated is required — something involving bitwise operations like XOR, or fun math functions like modulo division or Fibonacci numbers. Hashing functions tend to have interesting names like *radix conversion* or *folding hash* or *rolling hash*, which sounds like something you could get at a diner in Texas.

Many HCI solutions use a well-known algorithm called *SHA-1,* which was once used in cryptography but has since fallen from use because it's too easy to crack. However, SHA-1 is still a great hash algorithm when secrecy is not important.

SHA-1 is a great algorithm that works by dividing the input into 512-bit chunks and applying a complex scheme of bitwise operations on those chunks to produce a 160-bit hash that is consistent, uniform, and efficient. If you want to know more about how it works, search for "how SHA-1 works" on the internet.

Considering Backup

HCI is a game-changer for backup. Its built-in deduplication abilities enable an HCI platform to integrate backups into HCI storage. Most HCI platforms have built-in backup capabilities that let you run backups on a regularly scheduled basis, daily or weekly, and even hourly or even more often.

In HCI, there's really no such thing as a full backup: All backups are incremental (or differential) backups based on changed blocks within the HCI storage. When you consider how deduplication works, it's easy to see how deduplication can be leveraged for backups. Backups are really a construct of the hash table that's used to keep track of all the duplicate blocks that are on the storage. When you make a backup, HCI records the backup in the hash table. Then, when a block of data is changed, HCI knows whether to update the actual block of data or whether to write a new block of data, retaining the old block for the backup.

Keeping the backups on the storage itself has major advantages over traditional backups in which data is copied to a separate backup target. First, no data needs to be copied to create an HCI backup. And second, restoring an HCI backup is nearly instantaneous.

Because HCI leverages deduplication to create and maintain backups, you'll see astonishing deduplication ratios. For example, suppose you configure HCI to take a backup every day. After a month, you'll have 30 restore points. In effect, you have 30 copies of your data. Depending on how much of the data has changed over the 30 days, you'll see a deduplication ratio of up to 30:1, because each restore point represents a copy of most of the data on the disk. It's not uncommon to see deduplication ratios of 100:1 or even more.


TECHNICAL STUFF

The enormous deduplication ratios seen in HCI can trick you into thinking that you have more available storage than you really do. For example, suppose you have 10TB of data on a 20TB HCI node, with a deduplication ratio of 10:1. You may think that you're only using 1TB of actual disk and have 19TB available. But if you have ten backup restore points, HCI includes the size of all the restore points, which may indicate that you have as much as 100TB of data on the disk. 100TB of data at a deduplication ratio of 10:1 is actually consuming 10TB of disk, so your storage is actually half-full.

Fortunately, most HCI platforms have excellent management consoles that should give you a true picture of how your storage is being used.


WARNING

The one thing HCI backup does *not* provide is redundancy. Because data and its backup are stored on the same disk, if the disk is damaged, the data *and* its backup will be lost.

That's obviously an unacceptable situation. HCI mitigates that problem by providing redundancy across HCI nodes in a cluster. I explain how this works in the next section, but for now just realize that to provide redundancy, data is always written to at least two nodes. In an HCI cluster of two nodes, those nodes are effectively mirrors of one another. So, if one node is lost, the other node can step in without losing any data. For this reason, you should always plan for at least two HCI nodes. With just one node, you can't provide redundancy.


WARNING

One other point to consider with HCI backups: Data on an HCI system — including your backup data — is always online. You may remember from Book 3, Chapter 6 that a good backup strategy will include three types of backups: local, off-site, and offline. The HCI backup itself meets the local backup requirement. You can place two HCI nodes in different locations to meet the off-site requirement. But HCI on its own cannot meet the offline requirement. So, even with the most advanced HCI environment, I recommend you still incorporate tape backup into your design. Tape is still the only surefire way to provide offline backups.

Digging into HCI Clusters

As I mention earlier, you can combine HCI nodes into clusters to increase capacity. In its simplest incarnation, you can start with one node and add a second node to double capacity. When you do that, the storage on both of the nodes are combined into one unified storage bucket. The HCI platform takes care of deciding exactly where to store data across all the nodes in the cluster. This is a major advantage of HCI: With traditional storage, you often struggle to figure out how to pigeonhole all your data into many discrete buckets of storage, whose boundaries can't be crossed.

In addition to adding nodes to expand capacity, you can add nodes to provide redundancy. In many HCI installations, the second node doesn't add more capacity but replicates the data stored on the first node. That way, if either of the nodes fails, no data is lost.

As you continue to add nodes to a cluster, the HCI platform can spread your data out across all the nodes to provide even more redundancy. And if one or more of the nodes is located at a different site, you can ensure that off-site backups are built in to the HCI system.

For example, consider Figure 7-3. The left side of this figure shows an HCI cluster with two 40TB nodes. Because these nodes are mirrored, the total capacity of the cluster is 20TB. The right side of the figure shows a cluster with three 40TB nodes and a total capacity of 80TB.

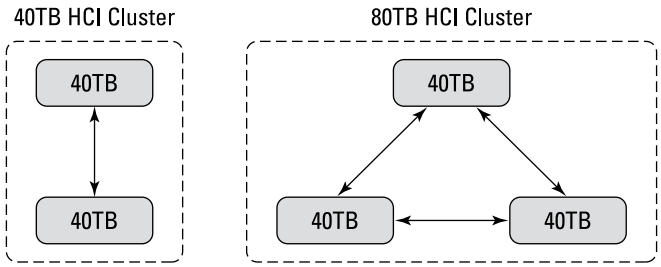


FIGURE 7-3: HCI clusters with two and three nodes.

Some HCI providers make a special type of HCI node that focuses on data only, not on computing power. These nodes, sometimes called disaster recovery (DR) nodes, are not designed to run virtual machine workloads. Instead, they're designed to provide secure off-site replication of your data. Because they don't have large amounts of memory or multiple high-performance CPUs, they can be considerably less expensive than regular HCI nodes.

Incorporating HCI Into Your Plan

One of the challenges of adopting HCI is determining what the actual capacity of your HCI storage will be. With traditional storage, 1TB is 1TB. But because HCI inherently uses deduplication as well as redundancy across nodes in a cluster and stores its backups on its own storage rather than on separate backup target storage, figuring out exactly how much data you'll be able to store can be a challenge.

For example, consider an HCI cluster with three 20TB nodes. The manufacturer of the nodes will probably tell you that you can expect at least 2:1 deduplication, so they'll advertise these nodes as having a capacity of 40TB. With three nodes, you have a theoretical capacity of 120TB. But with replication built in, the actual capacity will be more like 80TB, spread out across the three nodes.

Now if one of those nodes is a DR node and located off-site, that node needs to have enough capacity to hold all the data in case the two nodes at the primary site are lost. And a 40TB node can't hold 80TB of data. But maybe it can, if the deduplication ratio that's actually achieved is more like 4:1.

And if you keep a lot of backup restore points, you'll be dealing with enormous deduplication ratios that can skew your perception of how much actual data can be stored in the HCI cluster.

The bottom line about deduplication is that it's difficult to predict: It depends entirely on the nature of the data you're storing. And because of this, it's important that you engage with your HCI vendor for help in determining how to size your HCI implementation.

4 Implementing a Network

Contents at a Glance

CHAPTER 1: **Network Hardware** 327

CHAPTER 2: **Wireless Networks** 339

CHAPTER 3: **Windows Clients** 357

CHAPTER 4: **Mac Networking** 367

CHAPTER 5: **Network Printers** 377

CHAPTER 6: **Virtual Private Networks** 385

IN THIS CHAPTER

- Installing network cable
- Attaching cable connectors
- Figuring out pinouts for twisted-pair cabling
- Installing switches

Chapter 1
Network Hardware

After you plan your network, then comes the fun of actually putting everything together. In this chapter, I describe some of the important details for installing network hardware, including cables and switches, as well as professional touches like patch panels and cable management.

Working with Cable

Most Ethernet networks are built using twisted-pair cable (also known as UTP cable), which resembles phone cable but isn't the same. For more information about the general characteristics of twisted-pair cable, see of Book 3, Chapter 1.

In the following sections, you find out what you need to know to select and install twisted-pair cable.

Cable categories

Twisted-pair cable comes in various grades, or *categories.*

CAT GOT YOUR TONGUE?

Twisted-pair cable grades are categories specified by the ANSI/TIA standard 568. ANSI stands for American National Standards Institute; TIA stands for Telecommunications Industries Association. The higher the number, the faster the data transfer rate, so Cat-5 is faster than Cat-2. If you want to sound like you know what you're talking about, say "Cat 5e" instead of "Category 5e." Twisted pair is often referred to as UTP (unshielded twisted pair). Now you're hip.

Higher-category cables are more expensive than lower-category cables, but the real cost of installing Ethernet cabling is the labor required to actually pull the cables through the walls. You should never install anything less than Category 5 cable. And if at all possible, invest in Category 5e (the *e* stands for enhanced) or even Category 6 cable to allow for upgrades to your network.

Table 1-1 lists the various categories of twisted-pair cable in common use today.

TABLE 1-1 Twisted-Pair Cable Categories

Category	Maximum Data Rate	Intended Use
5e	1 Gbps	1 Gbps Ethernet, 100 meters
6	10 Gbps	10 Gbps Ethernet, 55 meters
6a	10 Gbps	10 Gbps Ethernet, 100 meters
7	10 Gbps	10 Gbps Ethernet, 100 meters
8	25-40 Gbps	25 or 40 Gbps, 30 meters

What's with the pairs?

Twisted-pair cable used for Ethernet has four pairs of wires, for a total of eight wires. The two wires within each pair are twisted around each other in a precise manner that prevents the electrical signals traveling through them from interfering with signals in the other pair. Without this twisting, the electrical signals would degrade and the connections wouldn't be reliable.

Each wire pair has a distinct color: green, blue, orange, or brown (see "Pinouts for twisted-pair cables," later in this chapter). Within each pair, one of the wires is a solid color and the other has a white stripe.

To shield or not to shield

Unshielded twisted pair cable (UTP) is designed for normal office environments. When you use UTP cable, you must be careful not to route cable close to fluorescent light fixtures, air conditioners, or electric motors (such as automatic door motors or elevator motors). UTP is the least expensive type of cable.

In outdoor environments or other environments with a lot of electrical interference, such as factories, you may want to use shielded twisted pair cable (STP). STP can cost up to three times more than regular UTP, so you won't want to use STP unless you have to. With a little care, UTP can withstand the amount of electrical interference found in a normal office environment.

Most STP cable is shielded by a layer of aluminum foil. For buildings with unusually high amounts of electrical interference, you can use more expensive, braided copper shielding for even more protection.

When to use plenum cable

The outer sheath of both shielded and unshielded twisted pair cable comes in two varieties:

- **PVC:** The most common and least expensive type.
- **Plenum:** A special type of fire-retardant cable designed for use in the plenum space of a building (typically, in the hollows below a floor or above a ceiling).

Plenum cable has a special Teflon coating that not only resists heat, but also gives off fewer toxic fumes if it does burn. Unfortunately, plenum cable costs more than twice as much as ordinary PVC cable.



WARNING

Most local building codes require that you use plenum cable whenever the wiring is installed within the plenum space of the building.

PLENUM SPACE

Plenum space is a compartment in the building's air distribution system — typically, the space above a suspended ceiling or under a raised floor.

The area above a suspended ceiling is *not* a plenum space if both the delivery and return lines of the air conditioning and heating system are ducted. Plenum cable is required only if the air conditioning and heating system are not ducted. When in doubt, have the local inspector look at your facility before you install cable.

Sometimes solid, sometimes stranded

The actual copper wire that composes the cable comes in two varieties: solid and stranded. Your network will have some of each.

- **Stranded cable:** Each conductor is made from a bunch of very small wires twisted together. Stranded cable is more flexible than solid cable, so it doesn't break as easily. However, stranded cable is more expensive than solid cable and isn't very good at transmitting signals over long distances. Stranded cable is best used for patch cables, such as the cable used to connect a computer to a wall jack or the cable used to connect patch panels to hubs and switches.

  Strictly speaking, the cable that connects your computer to the wall jack is a *station cable* — not a patch cable. Patch cables are used in the wiring closet, usually to connect patch panels to switches. However, in common practice, the terms *station cable* and *patch cable* are used interchangeably.

- **Solid cable:** Each conductor is a single solid strand of wire. Solid cable is less expensive than stranded cable and carries signals farther, but it isn't very flexible. If you bend it too many times, it will break. Solid cable is usually used for permanent wiring within the walls and ceilings of a building.

Installation guidelines

The hardest part about installing network cable is the physical task of pulling the cable through ceilings, walls, and floors. This job is just tricky enough that I recommend that you don't attempt it yourself except for small offices. For large jobs, hire a professional cable installer. You may even want to hire a professional for small jobs if the ceiling and wall spaces are difficult to access.

Here are some general pointers to keep in mind if you decide to install cable yourself:

- You can purchase twisted-pair cable in prefabricated lengths, such as 50 feet, 75 feet, or 100 feet. You can also special-order prefabricated cables in any length you need. However, attaching connectors to bulk cable isn't that difficult. I recommend that you use prefabricated cables only for very small networks and only when you don't need to route the cable through walls or ceilings.

- Always use a bit more cable than you need, especially if you're running cable through walls. For example, when you run a cable up a wall, leave a few feet of slack in the ceiling above the wall. That way, you'll have plenty of cable if you need to make a repair later on.

- When running cable, avoid sources of interference, such as fluorescent lights, big motors, x-ray machines, and so on. The most common source of interference for cables that are run behind dropped ceiling panels are fluorescent lights; be sure to give light fixtures a wide berth as you run your cable. Three feet should do it.
- The maximum allowable cable length between a hub and the computer is 100 meters (about 328 feet).
- If you must run cable across the floor where people walk, cover the cable so that no one trips over it. Inexpensive cable protectors are available at most hardware stores.
- When running cables through walls, label each cable at both ends. Most electrical supply stores carry pads of cable labels that are perfect for the job. These pads contain 50 sheets or so of precut labels with letters and numbers. They look much more professional than wrapping a loop of masking tape around the cable and writing on the tape with a marker.
- If nothing else, use a permanent marker to write directly on the cable.
- When several cables come together, tie them with plastic cable ties or — better yet — strips of Velcro. Don't use masking tape or duct tape; the tape doesn't last, but the sticky glue stuff does. It's a mess a year later. Cable ties are available at electrical supply stores. You can purchase rolls of Velcro that you can cut to the desired lengths from online suppliers.


TIP

- Cable ties have all sorts of useful purposes. Once on a backpacking trip, I used a pair of cable ties to attach an unsuspecting buddy's hat to a high tree limb. He wasn't impressed with my innovative use of the cable ties, but my other hiking companions were.


TIP

- When you run cable above suspended ceiling panels, use cable ties, J-hooks, or clamps to secure the cable to the actual ceiling or to the metal frame that supports the ceiling tiles. Don't just lay the cable on top of the tiles.

Getting the tools that you need

Of course, to do a job right, you must have the right tools.

Start with a basic set of computer tools, which you can get for about $15 from any computer store or large office-supply store. These kits include the right screwdrivers and socket wrenches to open up your computers and insert adapter cards. (If you don't have a computer toolkit, make sure that you have several flat-head and Phillips screwdrivers of various sizes.)

If all your computers are in the same room and you're going to run the cables along the floor and you're using prefabricated cables, the computer tool kit should contain everything that you need.

If you're using bulk cable and plan on attaching your own connectors, you need the following tools in addition to the tools that come with the basic computer toolkit:

- **Wire cutters:** You need decent wire cutters to cut UTP cable.
- **Crimp tool:** Use this tool to attach the connectors to the cable. Don't use a cheap $10 crimp tool. A good one will cost $100 but will save you many headaches in the long run. Remember this adage: When you crimp, you mustn't scrimp.
- **Wire stripper:** You need this only if the crimp tool doesn't include a wire stripper.

If you plan on running cables through walls, you need these additional tools:

- **A hammer**
- **A keyhole saw:** This is useful if you plan on cutting holes through walls to route your cable.
- **A flashlight**
- **A ladder**
- **Someone to hold the ladder**
- **Possibly a fish tape:** A *fish tape* is a coiled-up length of stiff metal tape. To use it, you feed the tape into one wall opening and fish it toward the other opening, where a partner is ready to grab it when the tape arrives. Next, your partner attaches the cable to the fish tape and yells something like, "Let 'er rip!" or "Bombs away!" Then you reel in the fish tape and the cable along with it. (You can find fish tape in the electrical section of most well-stocked hardware stores.)

If you plan on routing cable through a concrete subfloor, you need to rent a jackhammer and a back-hoe and hire someone to hold a yellow flag while you work.

Pinouts for twisted-pair cables

Each pair of wires in a twisted-pair cable is one of four colors: green, blue, orange, or brown. The two wires that make up each pair are complementary: One is a solid color, and the other is white with a stripe of the corresponding color. For example, the orange pair has an orange wire and then a white wire with an orange stripe.

Likewise, the blue pair has a blue wire and a white wire with a blue stripe. The wire with the solid color is referred to with just the name of the color (for example, "the blue wire"). The corresponding wire with the same color and the white stripe is referred to using the both colors (for example, "the blue/white" cable).

When you attach a twisted-pair cable to a modular connector or jack, you must match up the right wires to the right pins. You can use either of two different standards to wire the connectors. One is known as ANSI/TIA 568A; the other is ANSI/TIA 568B, also known as AT&T 258A. Table 1-2 shows both wiring schemes.

TABLE 1-2 Pin Connections for Twisted-Pair Cable

Pin Number	Function	EIA/TIA 568A	EIA/TIA 568B AT&T 258A
Pin 1	Transmit +	Green/white	Orange/white
Pin 2	Transmit –	Green	Orange
Pin 3	Receive +	Orange/white	Green/white
Pin 4	Unused	Blue	Blue
Pin 5	Unused	Blue/white	Blue/white
Pin 6	Receive –	Orange	Green
Pin 7	Unused	Brown/white	Brown/white
Pin 8	Unused	Brown	Brown



WARNING

It doesn't matter which of these wiring schemes you use, but pick one and stick with it. If you use one wiring standard on one end of a cable and the other standard on the other end, the cable won't work.

The only difference between the two wiring standards is which pair is used to transmit data and which pair is used to receive data. In the EIA/TIA 568A standard, the green pair is used to transmit and the orange pair is used to receive. In the EIA/TIA 568B and AT&T 258A standards, the orange pair is used to transmit and the green pair to receive.

Attaching RJ-45 connectors

RJ-45 connectors for twisted-pair cables aren't too difficult to attach if you have the right crimping tool. The trick is in both making sure that you attach each wire to the correct pin and pressing the tool hard enough to ensure a good connection.

Here's the procedure for attaching an RJ-45 connector:

1. **Cut the end of the cable to the desired length.**

   Make sure that you make a square cut — not a diagonal cut.

2. **Insert the cable into the stripper portion of the crimp tool so that the end of the cable is against the stop.**

   Squeeze the handles and slowly pull the cable out, keeping it square. This strips off the correct length of outer insulation without puncturing the insulation on the inner wires.

3. **Arrange the wires so that they lay flat and line up according to Table 1-2.**

   You'll have to play with the wires a little bit to get them to lay out in the right sequence.

4. **Slide the wires into the pinholes on the connector.**

   Double-check to make sure that all the wires slip into the correct pinholes.

5. **Insert the plug and wire into the crimping portion of the tool and then squeeze the handles to crimp the plug.**

   Squeeze it tight!

6. **Remove the plug from the tool and double-check the connection.**

   You're done!

Here are a few other points to remember when dealing with RJ-45 connectors and twisted-pair cable:


TIP

- The pins on the RJ-45 connectors aren't numbered, but you can tell which is pin 1 by holding the connector so that the metal conductors are facing up, as shown in Figure 1-1. Pin 1 is on the left.
- Be extra careful to follow the rules of Cat-5 cabling. That means, among other things, making sure that you use Cat-5 components throughout. The cable and all the connectors must be up to Cat-5 specs. When you attach the connectors, don't untwist more than one-half inch of cable. And don't try to stretch the cable runs beyond the 100m maximum. When in doubt, have cable for a 100 Mbps Ethernet system professionally installed.

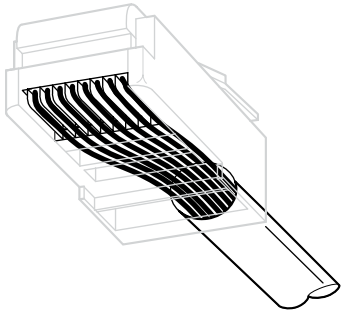


FIGURE 1-1: Attaching an RJ-45 connector to twisted-pair cable.

Wall jacks and patch panels

From the files of "Sure, you could do this, but here's why this is a bad idea," you could run a single length of cable from a network switch in a wiring closet through a hole in the wall, up the wall to the space above the ceiling, through the ceiling space to the wall in an office, down the wall, through a hole, and all the way to a desktop computer. Here's the pitfall, though: Every time someone moves the computer or even cleans behind it, the cable will get moved a little bit. Eventually, the connection will fail, and the RJ-45 plug will have to be replaced. Then the cables in the wiring closet will quickly become a tangled mess.

The smarter path is to put a wall jack at the user's end of the cable and connect the other end of the cable to a patch panel. Then, the cable itself is completely contained within the walls and ceiling spaces. To connect a computer to the network, you plug one end of a patch cable (properly called a *station cable*) into the wall jack and plug the other end into the computer's network interface. In the wiring closet, you use a patch cable to connect the wall jack to the network switch. Figure 1-2 shows how this arrangement works.

Connecting a twisted-pair cable to a wall jack or a patch panel is similar to connecting it to an RJ-45 plug. However, you don't usually need any special tools. Instead, the back of the jack has a set of slots that you lay each wire across. You then snap a removable cap over the top of the slots and press it down. This forces the wires into the slots, where little metal blades pierce the insulation and establish the electrical contact.



TIP

When you connect the wire to a jack or patch panel, be sure to carefully follow the color-coded label on the jack, and untwist as little of the wire as possible. If you untwist too much of the wire, the signals that pass through the wire may become unreliable.

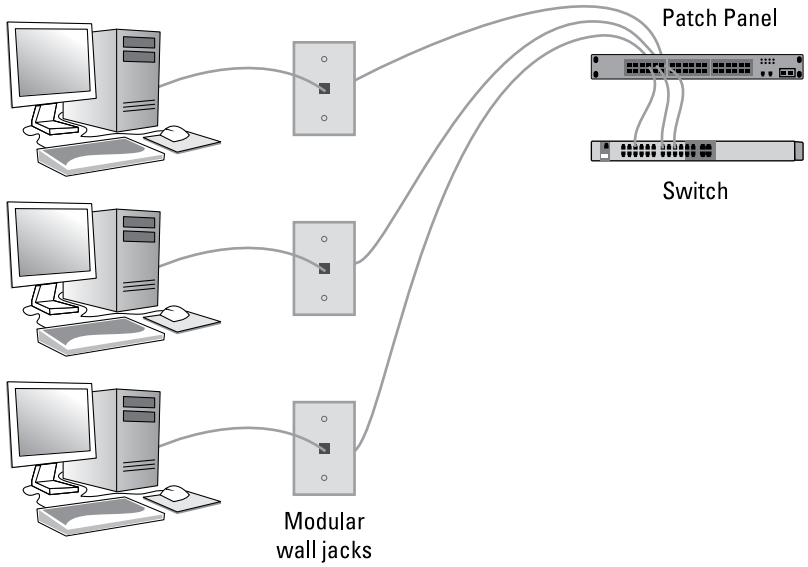


FIGURE 1-2: Using wall jacks and patch panels.

Server rooms and distribution frames

In a very small network, you may have a server computer, a switch, and a firewall router stuck in a corner somewhere. But for larger networks, you'll need a dedicated space for the networking equipment, as well as the servers and other paraphernalia that goes with them, such as uninterruptible power supply (UPS) units and patch panels.

Depending on the size of your network and the layout of your office complex (for example, single- or multi-story, one building or several), you may want to consider various types of locations for your gear:

- **Server room:** The room that contains your servers. Usually, this room will contain one or two equipment racks that hold server computers, disk subsystems, tape backup devices, and UPS units. This room should be equipped with plenty of power outlets on isolated circuits (not shared with wall outlets outside the server room). It will also need plenty of air conditioning, usually a separate A/C unit dedicated to keeping your equipment cool.
- **Main distribution frame (MDF):** The main switching hub of your network. Technically the term *MDF* refers to a rack (or collection of racks) that houses the patch panels and switching devices that manage the core of your network. But the term is sometimes used to refer to the room that contains the MDF. This room often doubles as the server room unless your network is very large.

» **Intermediate distribution frame (IDF):** A remote distribution point for your network. Typically this is for separate floors, suites within a multisuite building, or separate buildings on a campus. Equipment in an IDF is usually connected directly to the main IDF via a single high-speed cable, preferably capable of carrying 10 Gbps traffic even if the equipment in the MDF and IDF can't handle 10 Gbps. If the IDF is more than 100 meters from the MDF, fiber cable should be used.

Figure 1-3 shows an example in which an MDF/server room contains a couple of switches and a couple of servers on a single rack. Then two IDF rooms — one on the first floor, the other on the second floor — contain smaller racks, with just one patch panel and one switch. The IDF switches are connected to the MDF switches using a single cable. Then the IDF switches connect to end-user computers located in offices on each floor.

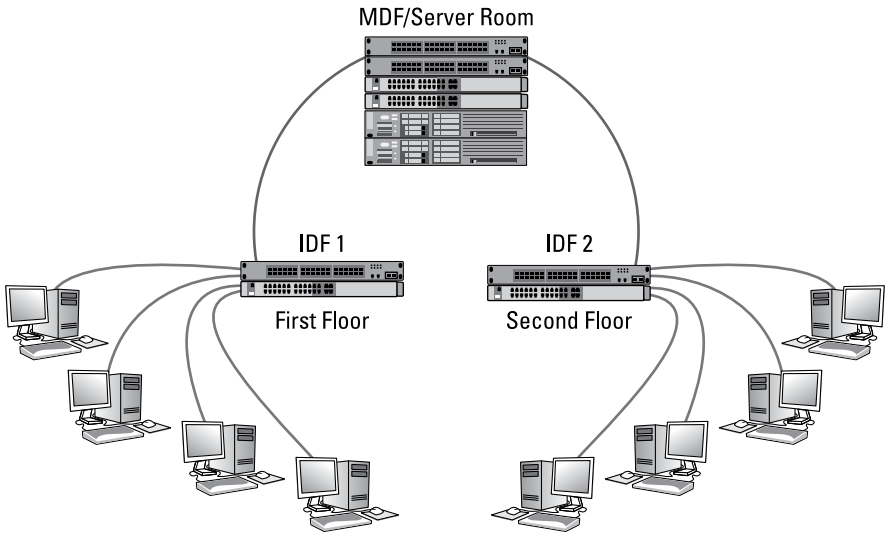


FIGURE 1-3: Switches distributed in IDFs on separate floors in a multistory office.

Installing Switches

Setting up a network switch is remarkably simple. In fact, you need to know only a few details:

» Installing a switch is usually very simple. Just plug in the power cord and then plug in patch cables to connect the network.

- Each port on the switch has an RJ-45 jack and an LED indicator that lights up when a connection has been established on the port. If you plug one end of a cable into the port and the other end into a computer or other network device, the Link light should come on. If it doesn't, something is wrong with the cable, the hub (or switch port), or the device on the other end of the cable.
- Each port may also have an LED indicator that flashes to indicate network activity. If you stare at a switch for a while, you can find out who uses the network most by noting which activity indicators flash the most.
- The ports may also have an LED that indicates the speed of the connection that has been established. For example, the LED might be light green for a 1GB connection but amber for a 100MB connection. You can use this LED to identify ports that have potential cabling problems, as computers that should connect at 1GB will often connect at 100MB instead if the cable is suspect.

IN THIS CHAPTER

- Working with a wireless access point
- Configuring Windows for a wireless network
- Securing a wireless network

Chapter 2
Wireless Networks

Nearly all modern networks include wireless access. In Book 1, Chapter 2, you learn the basics of how a wireless network works and how to incorporate wireless access in your networking plan. In this chapter, you learn the basics of configuring and securing a wireless network.

Installing a Wireless Access Point

When planning your wireless network, you'll find two distinct devices that provide wireless access:

- **Wireless firewall routers:** Wireless firewall routers combine the function of a wireless access point with the function of a router and a firewall. This type of device is also sometimes called a *multifunction wireless router* or a *gateway* device and is often supplied by your internet service provider (ISP).
- **Wireless access point (WAP):** WAPs provide just the wireless access point function. These are sometimes referred to just as *access points* (APs).

For a small office or home network, you'll probably just use a wireless firewall router provided by your ISP. For larger offices, you'll want to use dedicated AP

devices rather than the access point that's built in to the gateway devices provided by your ISP. If your office is large enough, you'll need two or more AP devices, strategically placed throughout the office to provide coverage to every nook and cranny of the space.

The physical setup for a wireless AP is pretty simple: You take it out of the box, mount it in its desired location, and plug it in. You can place the access point on top of a bookcase or shelf, or you can mount it to a wall or the ceiling. Many access points come with mounting hardware designed to hang the access point from a suspended ceiling.

The most important part of installing the access point is choosing the location. Ideally, it should be centrally located within the space it needs to provide access to. And it will need to have access to both power and the wired network. That's why it's important that you plan for wireless installation when you create the network's cabling plan.

If you don't have an electrical outlet at the location where you want to place the access point, you can use a Power over Ethernet (PoE) switch, which injects power on the Ethernet cable. That way, the power and network access are delivered to the access point over a single cable. Of course, you'll need to make sure that your access point is compatible with PoE to use this solution.

An alternative to using a PoE switch is to use a device called a *PoE injector.* A PoE injector has two Ethernet jacks — an input and an output — plus a power cable. The input jack accepts a standard Ethernet cable coming from a switch. You can then connect one end of an Ethernet cable to the injector's output jack and the other end to the access point. The PoE injector adds power to the output jack to provide power for the access point.

Configuring a Wireless Access Point

The software configuration for an access point is a little more involved but still not very complicated. It's usually done via a web interface. To get to the configuration page for the access point, you need to know the access point's Internet Protocol (IP) address. Then you just type that address in the address bar of a browser on any computer on the network. If you aren't certain of the IP address, you can usually find out by checking your Dynamic Host Configuration Protocol (DHCP) server.

Figure 2-1 shows the main Wireless Settings page for a typical WAP. I called up this configuration page by entering the IP address of the access point (10.0.0.145) in the address bar of a web browser and then supplying the login credentials when prompted.

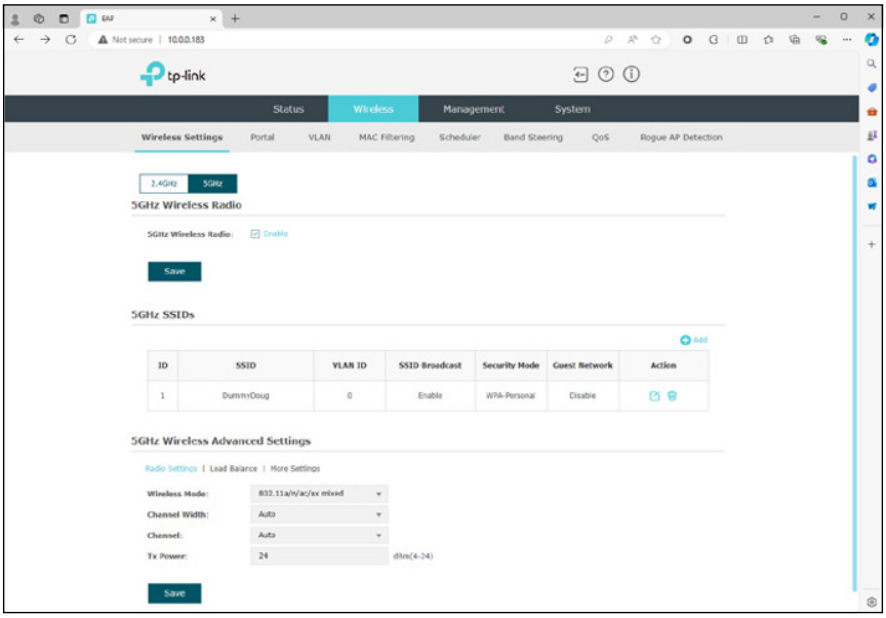


FIGURE 2-1: The Wireless Settings page for a typical wireless router.

When configuring a WAP, you'll want to consider the following options:

- **Broadcast frequency:** Most APs support wireless communication on two frequencies: 2.4 GHz and 5 GHz. 5 GHz wireless networks are faster than 2.4 GHz networks but have a shorter range. In most cases, you'll want to enable both frequencies.
- **SSID:** The SSID is the name of the wireless network that your users will recognize when they connect wirelessly. Choose a name that identifies your network and distinguishes it from wireless networks that originate from nearby businesses or homes. (*SSID* stands for Service Set Identifier, but that won't be on the test.)
- **VLAN:** By default, any new SSID you create will be associated with VLAN 0. If you want, you can associate an SSID with a different VLAN by enabling VLAN for the SSID and specifying a VLANID. For more information about virtual local area networks (VLANs), refer to Book 1, Chapter 3.

» **Security:** Always enable the strongest possible form of security for every one of your SSIDs. The only exception is if you enable a guest wireless network, which will allow visitors to connect to your wireless network for internet access only, with no access to any other resources connected to your local area network (LAN).

Connecting to a Wireless Network

Connecting to a wireless network on a Windows computer is straightforward. Windows automatically detects any wireless networks that are in range and displays them in a list when you tap the Wireless icon at the bottom of the screen, as shown in Figure 2-2.

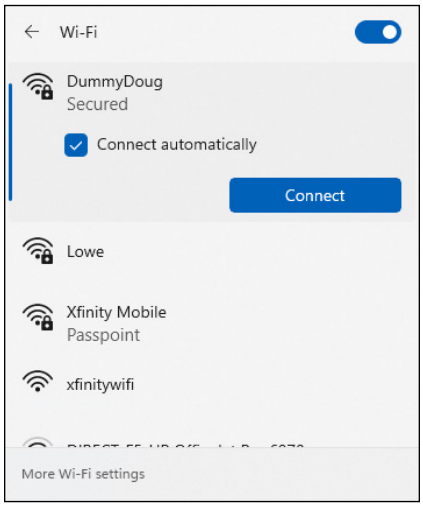


FIGURE 2-2: Choosing a wireless network in Windows 11.

To connect to a network, just tap it, and then enter the security key when prompted. If the key is correct, you'll be connected.

At the time you connect, you can choose to connect to the network automatically whenever it's in range. If you select this option, you won't have to select the network manually or enter the security key; you'll just be connected automatically.

Windows remembers every network you connect to, which is a plus for networks you frequently use but a drawback for networks you'll likely never use again. To tell Windows to forget a network, follow these steps:

1. **Click Start, and then tap Settings.**

   The Settings window appears.

2. **Click Network & Internet.**

   This brings up the Network & Internet page.

3. **Click Wi-Fi.**

   This brings up the Wi-Fi settings page.

4. **Click Manage Known Networks.**

   This brings up the Manage Known Networks page, shown in Figure 2-3.

5. **In the Manage Known Networks section, click the Forget button for the network you want to forget.**

   The network will be forgotten. To log into this network again, you'll have to enter the security key.

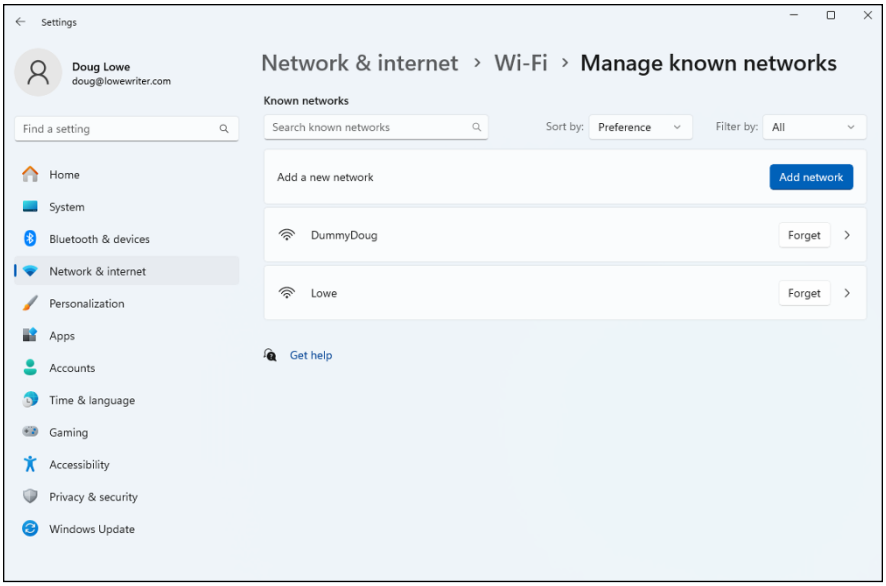


FIGURE 2-3: Forgetting a wireless network in Windows 11.

Paying Attention to Wireless Network Security

Before you dive headfirst into the deep end of the wireless networking pool, you should consider the inherent security risks in setting up a wireless network. With a cabled network, the best security tool that you have is the lock on the front door

of your office. Unless someone can physically get to one of the computers on your network, they can't get into your network.

If you go wireless, an intruder doesn't have to get into your office to hack into your network. They can do it from the office next door. Or the lobby. Or the parking garage below your office. Or the sidewalk outside. In short, when you introduce wireless devices into your network, you usher in a whole new set of security issues to deal with.

Understanding wireless security threats

Wireless networks have the same basic security considerations as wired networks. As a network administrator, you need to balance the need of legitimate users to access network resources against the risk of illegitimate users breaking into your network. That's the basic dilemma of network security. Whether the network uses cables, wireless devices, kite strings and tin cans, or smoke signals, the basic issues are the same.

At one extreme of the wireless network security spectrum is the totally open network, in which anyone within range of your wireless transmissions can log on as an administrator and gain full access to every detail of your network. At the other end is what I call the "cone-of-silence syndrome," in which the network is so secure that no one can gain access to the network — not even legitimate users.

The goal of securing a wireless network is to find the happy medium between these two extremes that meets the access and risk-management needs of your organization.

The following sections describe the types of security threats that wireless networks are most likely to encounter. You should take each of these kinds of threats into consideration when you plan your network's security.

Intruders

With a wired network, an intruder usually must gain access to your facility to physically connect to your network. That's not so with a wireless network. In fact, hackers equipped with notebooks that have wireless network capability can gain access to your network if they can place themselves physically within range of your network's radio signals. Consider these possibilities:

- If you share a building with other tenants, the other tenants' offices may be within range.
- If you're in a multifloor building, the floor immediately above or below you may be in range.

- The lobby outside your office may be within range of your network.
- The parking lot outside or the parking garage in the basement may be in range.

If a would-be intruder can't get within normal broadcast range, they may try one of several tricks to increase the range:

- A would-be intruder can switch to a bigger antenna to extend the range of their wireless computer. Some experiments have shown that big antennas can receive signals from wireless networks miles away. In fact, I once read about someone who listened in on wireless networks based in San Francisco from the Berkeley hills, across San Francisco Bay.
- If a would-be intruder is serious about breaking into your network, they may smuggle a wireless repeater device into your facility — or near it — to extend the range of your wireless network to a location that they *can* get to.


REMEMBER

A *physical* connection to your network isn't the only way an intruder can gain access, of course. You must still take steps to prevent an intruder from sneaking into your network through your internet gateway. In most cases, this means that you need to set up a firewall to block unwanted and unauthorized traffic.

Freeloaders

Freeloaders are intruders who want to piggyback on your wireless network to get free access to the internet. If they manage to gain access to your wireless network, they probably won't do anything malicious: They'll just fire up their web browsers and surf the internet. These are folks who are too cheap to spend $40 per month on their own broadband connection at home, so they'd rather drive into your parking lot and steal yours.

Even though freeloaders may be relatively benign, they can be a potential source of trouble. In particular:

- Freeloaders use bandwidth that you're paying for. As a result, their mere presence can slow down internet access for your legitimate users.
- After freeloaders gain internet access through your network, they can potentially cause trouble for you or your organization. They may use your network to download illegal pornography, or they may try to send spam via your mail server. Most ISPs will cut you off cold if they catch you sending spam, and they won't believe you when you tell them that the spam came from a kid parked in a Pinto out in your parking lot.
- If you're in the business of *selling* access to your wireless network, obviously, freeloaders are a problem.

- Freeloaders may start out innocently looking for free internet access. Once they get in, though, curiosity may get the better of them, leading them to snoop around your network.
- If freeloaders can get in, so can malicious intruders.

Eavesdroppers

Eavesdroppers just like to listen to your network traffic. They don't actually try to gain access via your wireless network — at least, not at first. They just listen.

Unfortunately, wireless networks give them plenty to listen to:

- Most WAPs regularly broadcast their Service Set Identifiers (SSIDs) to anyone who's listening.
- When a legitimate wireless network user joins the network, an exchange of packets occurs as the network authenticates the user. An eavesdropper can capture these packets and, if security isn't set up right, determine the user's logon name and password.
- An eavesdropper can steal files that are opened from a network server. If a wireless user opens a confidential sales report that's saved on the network, the sales-report document is broken into packets that are sent over the wireless network to the user. A skilled eavesdropper can copy those packets and reconstruct the file.
- When a wireless user connects to the internet, an eavesdropper can see any packets that the user sends to or receives from the internet. If the user purchases something online, the transaction may include a credit card number and other personal information. (Ideally, these packets will be encrypted so that the eavesdropper won't be able to decipher the data.)

Spoilers

A *spoiler* is a hacker who gets kicks from jamming networks so that they become unusable. A spoiler usually accomplishes this act by flooding the network with meaningless traffic so that legitimate traffic gets lost in the flow. Spoilers may also try to place viruses or worm programs on your network via an unsecured wireless connection.

Rogue access points

One of the biggest problems that network administrators have to deal with is the problem of rogue access points. A *rogue access point* is an access point that

suddenly appears on your network out of nowhere. What usually happens is that an employee decides to connect a notebook computer to the network via a wireless computer. So this user stops at Computers-R-Us on the way home from work one day, buys a Fisher-Price WAP for $25, and plugs it into the network without asking permission.

Now, in spite of all the elaborate security precautions you've taken to fence in your network, this well-meaning user has opened the barn door. It's *very* unlikely that the user will enable the security features of the WAP; in fact, they probably isn't even aware that wireless access devices *have* security features.

Unless you take some kind of action to find it, a rogue access point can operate undetected on your network for months or even years. You may not discover it until you report to work one day and find that your network has been trashed by an intruder who found their way into your network via an unprotected WAP that you didn't even know existed.



TIP

Here are some steps you can take to reduce the risk of rogue access points appearing on your system:

- Establish a policy prohibiting users from installing WAPs on their own. Then make sure that you inform all network users of the policy, and let them know why installing an access point on their own can be such a major problem.
- If possible, establish a program that quickly and inexpensively grants wireless access to users who want it. Rogue access points show up in the first place for two reasons:
  - Users need the access.
  - The access is hard to get through existing channels.
- If you make it easier for users to get legitimate wireless access, you're less likely to find WAPs hidden behind file cabinets or in flower pots.
- Once in a while, take a walk through the premises, looking for rogue access points. Take a look at every network outlet in the building; see what's connected to it.
- Turn off all your WAPs and then walk around the premises with a wireless-equipped mobile device such as a smartphone and look for wireless networks that pop up. Just because you detect a wireless network, of course, doesn't mean you've found a rogue access point; you may have stumbled onto a wireless network in a nearby office or home. But knowing what wireless networks are available from within your office will help you determine whether or not any rogue access points exist.

Securing your wireless network

I hope you're convinced that wireless networks do, indeed, pose many security risks. In the following sections, I describe some steps that you can take to help secure your wireless network.

Changing the password

Probably the first thing you should do when you install a WAP is change its administrative password. Most access points have a built-in, web-based setup page that you can access from any web browser to configure the access point's settings. The setup page is protected by a username and password, but the username and password are initially set to default values that are easy to guess.

For Linksys access points, for example, the default username is usually either blank or `admin`, and the password is `admin`. If you leave the username and password set to their default values, anyone can access the access point and change its configuration settings, thus bypassing any other security features that you enable for the access point.

So, the first step in securing your WAP is changing the setup password to a value that can't be guessed. I suggest that you use a random combination of numerals and both uppercase and lowercase letters. Be sure to store the password in a secure location. (If you forget the password, you can press the Reset button on the router to restore it to its factory default. Then you can log on by using the default password, which you can find with the documentation that came with the router.)

Securing the SSID

The next step is to secure the SSID that identifies the network. A client must know the access point's SSID to join the wireless network. If you can prevent unauthorized clients from discovering the SSID, you can prevent them from accessing your network.



WARNING

Securing the SSID isn't a complete security solution, so you shouldn't rely on it as your only security mechanism. SSID security can slow down casual intruders who are just looking for easy and free internet access, but it isn't possible to prevent serious hackers from discovering your SSID.



TIP

You can do three things to secure your SSID:

- **Change the SSID from the default.** Most access points come preconfigured with well-known default SSIDs, such as those listed in Table 2-1. By changing your access point's SSID, you can make it more difficult for an intruder to determine your SSID and gain access.

- **Disable SSID broadcast.** Most access points frequently broadcast their SSIDs so that clients can discover the network when they come within range. Clients that receive this SSID broadcast can use the SSID to join the network.

  You can increase network security somewhat by disabling the SSID broadcast feature. That way, clients won't automatically learn the access point's SSID. To join the network, a client computer must figure out the SSID on its own. Then you can tell your wireless network users the SSID to use when they configure their clients.

  
  WARNING

  Unfortunately, when a client computer connects to a wireless network, it sends the SSID to the access point in an unencrypted packet. So a sophisticated intruder who's using a packet sniffer to eavesdrop on your wireless network can determine your SSID as soon as any legitimate computer joins the network.

- **Disable guest mode.** Many access points have a guest-mode feature that enables client computers to specify a blank SSID or to specify "any" as the SSID. If you want to ensure that only clients that know the SSID can join the network, you must disable this feature.

TABLE 2-1 Common Default SSID Values

SSID	Manufacturer
3com	3Com
Compaq	Compaq
Linksys	Linksys
Tsunami	Cisco
Wireless	NetGear
WLAN	D-Link
WLAN	SMC

Using WPA and WPA2

WPA, which stands for *Wi-Fi Protected Access*, is a security protocol for wireless networks. You should always ensure that WPA security is enabled. Figure 2-4 shows a typical configuration screen for a WAP that confirms that WPA has been enabled on the device.

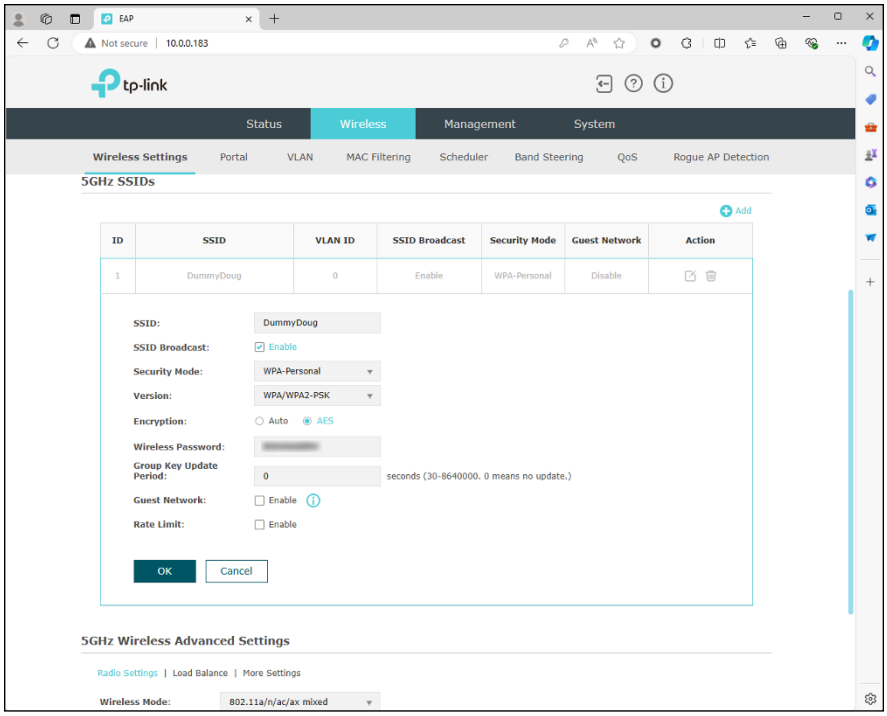


FIGURE 2-4: Changing the WPA settings on a wireless router.

Here are a few additional things to know about WPA:

- A small office and home version of WPA, called WPA-PSK, bases its encryption keys on a passkey value that you supply. True WPA devices, however, rely on a special authentication server to generate the keys.
- All versions of Windows since Windows XP Service Pack 2 have built-in support for WPA.
- The official IEEE standard for WPA is 802.11i. However, WPA devices were widely available before the 802.11i standard was finalized; as a result, not all WPA devices implement every aspect of 802.11i.
- The original version of WPA has been superseded by a newer version, named WPA2.

Using MAC address filtering

MAC address filtering allows you to specify a list of MAC addresses for the devices that are allowed to access the network or are prohibited from accessing the network. If a computer with a different MAC address tries to join the network via the access point, the access point will deny access.

MAC address filtering is a great idea for wireless networks with a fixed number of clients. If you set up a wireless network at your office so that a few workers can connect their notebook computers, you can specify the MAC addresses of those computers in the MAC filtering table. Then other computers won't be able to access the network via the access point.



WARNING

Unfortunately, it isn't difficult to configure a computer to lie about its MAC address. Thus, after a potential intruder determines that MAC filtering is being used, they can just sniff packets to determine an authorized MAC address and then configure their computer to use that address. (This practice is called *MAC spoofing.*) So you shouldn't rely on MAC address filtering as your only means of security.

Figure 2-5 shows the screen used to edit the MAC address table for a typical WAP.

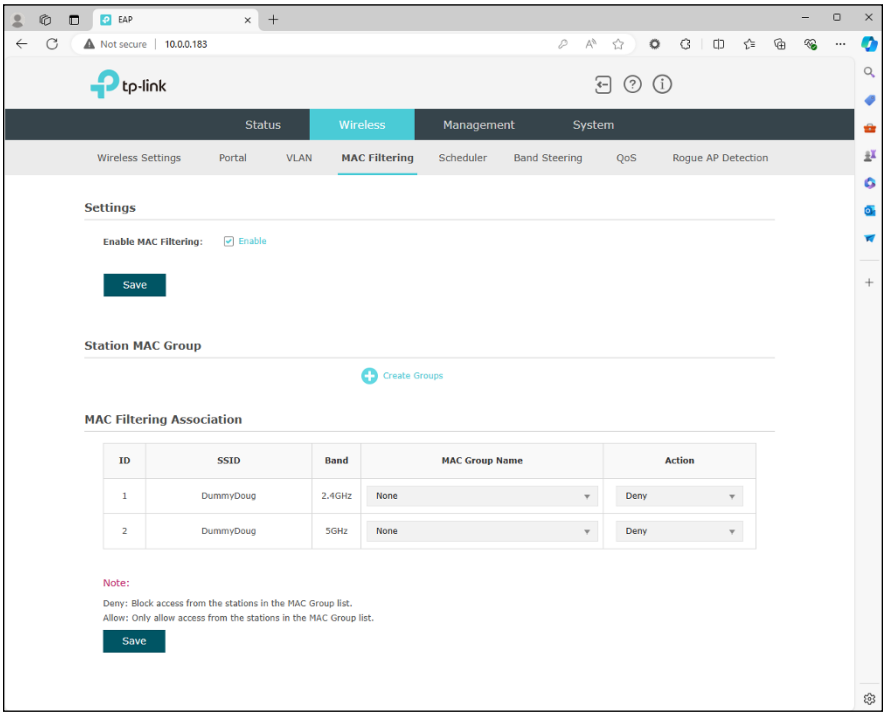


FIGURE 2-5: A MAC address table for a wireless router.

Placing your access points outside the firewall

The most effective security technique for wireless networking is placing all your WAPs *outside* your firewall. That way, all network traffic from wireless users will have to travel through the firewall to access the network.

DON'T NEGLECT THE BASICS

The security techniques described in this chapter are specific to wireless networks. They should be used alongside the basic security techniques that are presented in Book 1, Chapter 4 and in Book 10. In other words, don't forget the basics, such as these:

- Use strong passwords for your user accounts.
- Apply security patches to your servers.
- Change default server account information (especially the administrator password).
- Disable unnecessary services.
- Check your server logs regularly.
- Install virus protection.
- Back up!

As you can imagine, doing this can significantly limit network access for wireless users. To get around those limitations, you can enable a virtual private network (VPN) connection for your wireless users. The VPN will allow full network access to authorized wireless users.

Obviously, this solution requires a bit of work to set up and can be a little inconvenient for your users, but it's an excellent way to fully secure your WAPs.

Troubleshooting a wireless network

Wireless networks are great until something goes haywire. When a regular network doesn't work, you usually know about it right away because the network simply becomes unavailable. You can't display web pages, read email, or access files on shared drives.

The troubleshooting chapters in Book 9, Chapter 4 address the most common problems encountered on cabled networks. But wireless networks can cause problems of their own. And to add to the frustration, wireless networks tend to degrade rather than completely fail. Performance gets slower. Web pages that usually pop up in a second or two take 15 to 20 seconds to appear. Or sometimes they don't appear at all, but if you try again a few minutes later, they download fine.

This following sections offer some troubleshooting tips that can help you restore normalcy to a failing wireless network.

Checking for obvious problems

Before you roll up your sleeves and take drastic corrective action, you should check for a few obvious things if you're having wireless network trouble. The following list highlights some basic things you should check for:

- Is everything turned on? Make sure you have lights on your WAP/router, as well as on your cable or DSL modem.
- Many access point/routers use a power supply transformer that plugs into the wall. Make sure that the transformer is plugged into the wall outlet and that the small cable that comes out of the transformer is plugged into the power connector on the access point/router.
- Are the cables connected? Check the network cable that connects your access point/router to the cable or DSL modem.
- Try restarting everything. Turn off the computer, the access point/router, and your cable or DSL modem. Leave everything off for at least two minutes. Then turn everything back on. Sometimes, simply cycling the power off and back on clears up a connection problem.

Pinpointing the problem

If you can't connect to the internet, one of the first steps (after you've made sure that everything is turned on) is finding out whether the problem is with your access point/router or with your broadband connection. Here is one way you can check to find out whether your wireless connection is working:

1. **Open a command prompt window by choosing Start⇨cmd, and pressing Enter.**
2. **At the command prompt, type ipconfig, and press Enter.**

   You should get a display similar to this:

```
Ethernet adapter Wireless Network Connection:
Connection-specific DNS Suffix . : hsd1.ca.comcast.net)).
IP Address ..................... : 10.0.0.200
Subnet Mask .................... : 255.255.255.0
Default Gateway ................ : 10.0.0.1
```

If the display resembles this but with different numbers, you're connected to the wireless network, and the problem most likely lies with your broadband modem.

But if the IP Address, Subnet Mask, and Default Gateway indicate 0.0.0.0 instead of valid IP addresses, you have a problem with your wireless network.

Changing channels

One of the most common sources of wireless network trouble is interference from other wireless devices. The culprit might be a cordless phone, or it could be a neighbor who also has a wireless network.

The simplest solution to this type of interference is changing channels. 802.11b access points let you select 1 of 11 different channels to broadcast on. If you're having trouble connecting to your access point, try changing the channel. To do that, you must log on to the router with the administrator password. Then hunt around the router's administrator pages until you find the controls that let you change the channel.

You may have to try changing the channel several times before you solve the problem. Unfortunately, 802.11b channels overlap slightly, which means that broadcasts on one channel may interfere with broadcasts on adjacent channels. Thus, if you're having trouble connecting on channel 1, don't bother switching to channel 2. Instead, try switching to channel 5 or 6. If that doesn't work, switch to channel 10 or 11.

Fiddling with the antennas

Sometimes, you can fix intermittent connection problems by fiddling with the antennas on the access point and your computer's wireless adapter. This procedure is similar to playing with old-fashioned rabbit-ear antennas on a TV to get the best reception.

The angles of the antennas sometimes make a difference, so try adjusting the antenna angles. In addition, you usually have better results if you place the access point at a high location, such as on top of a bookshelf.

In some cases, you may actually need to add a high-gain antenna to the access point to increase its range. A high-gain antenna simply snaps or screws onto the access point to provide a bigger antenna. Antennas such as these cost about $70 for a pair.

A more drastic fix is to add a signal booster to your access point. A *signal booster* is a power amplifier that increases the transmission power of most wireless devices by a factor of five. A typical signal booster costs about $100.

Adding another access point

If you have a computer that's out of range of your access point, one solution is to add a second access point closer to the problematic computer. Most likely, the

only difficulty will be getting an Ethernet cable to the location where you want to put your second access point.

If possible, you can simply run a length of cable through your walls or attic to the second access point. If that solution isn't feasible, you can use a HomePlug or HomePNA network connection for the second access point.

An alternative to a second access point is simply adding a range extender. This handy device plugs directly into any electrical outlet and provides a pass-through outlet so you can still use your vacuum cleaner. Just plug this device midway between your access point and the computer that's having trouble connecting. (Note that using a range extender will slow down your Wi-Fi connection.)

Help! I forgot my router's password!

I mention many times throughout this book that you should always change the default passwords that come with computer and operating systems to more secure passwords, usually consisting of a random combination of letters, digits, and special symbols.

Ideally, you've already taken my sage advice and changed the password on your combination WAP/router. Good for you. But what if you forget the password later? Is there any way to get back into your access point/router then?


TIP

Fortunately, there is. Most access point/routers have a Reset button. It's usually located on the back or on the bottom of the router's case. Press this button to restore the access point/router to its factory default settings. That action resets the administrator password to the factory default — and also resets any other custom settings you've applied, so you may have to reconfigure your router to get it working again.

IN THIS CHAPTER

- Configuring network connections for Windows clients
- Changing the computer's name
- Joining a domain

Chapter 3
Windows Clients

Before your network setup is complete, you must configure the network's client computers. In particular, you have to configure each client's network interface card so that it works properly, and you have to install the right protocols so that the clients can communicate with other computers on the network.

Fortunately, the task of configuring client computers for the network is child's play in Windows. For starters, Windows automatically recognizes your network interface card when you start up your computer. All that remains is to make sure that Windows properly installed the network protocols and client software.

With each version of Windows, Microsoft has simplified the process of configuring client network support. In this chapter, I describe the steps for configuring networking for Windows 11. The procedures for previous versions of Windows are similar.

Configuring Network Connections

Windows automatically detects and configures network adapters, so you don't have to manually install device drivers for your network adapters. When Windows detects that a network adapter is present on the system, Windows automatically creates a network connection and configures it to support basic

networking protocols. You may need to change the configuration of a network connection manually, however.

The following steps show you how to configure your network adapter on a Windows 11 system:



1. **Click the Start icon (or press the Start button on the keyboard), and then tap or click the Settings icon (shown in the margin).**

   The Settings page appears, as shown in Figure 3-1.

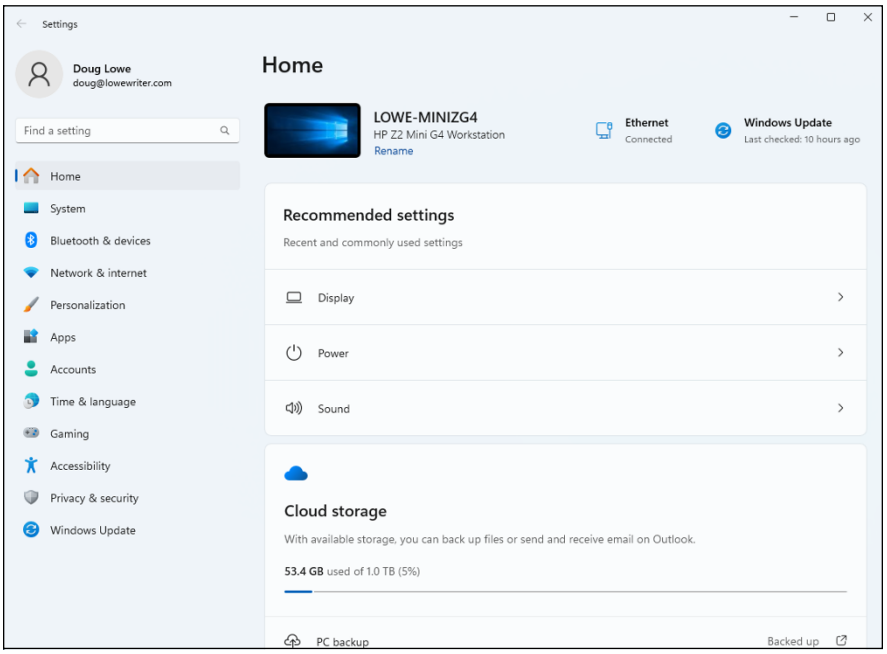


FIGURE 3-1: The Settings page.

2. **Click Network & Internet.**

   The Network & Internet page appears, as shown in Figure 3-2.

3. **Click Advanced Network Settings.**

   The Advanced Network Settings page, shown in Figure 3-3, appears.

4. **Click Ethernet to adjust settings for your Ethernet connection.**

   The Ethernet settings section of the Advanced Network Settings page expands, as shown in Figure 3-4.

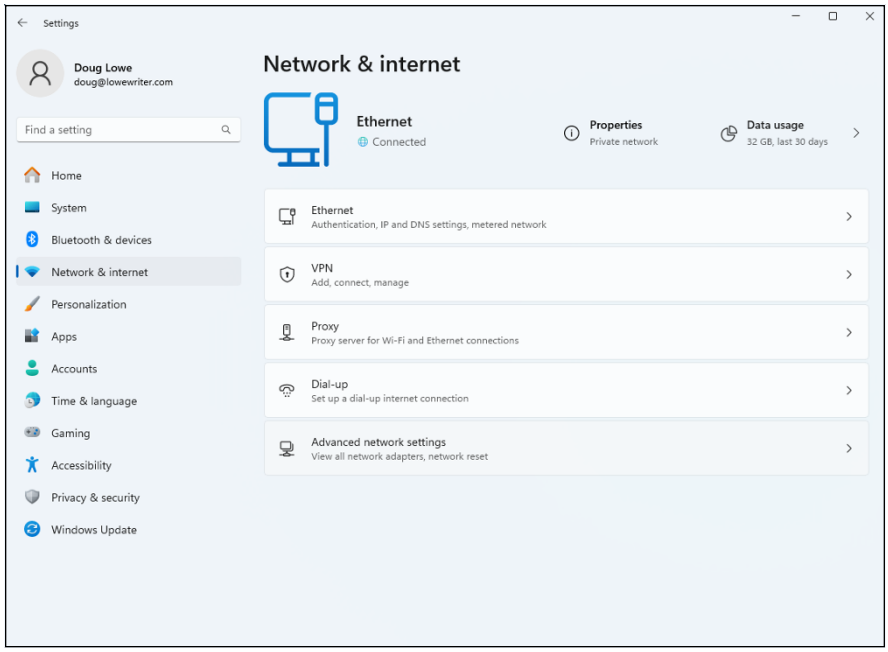


FIGURE 3-2: The Network & Internet page.

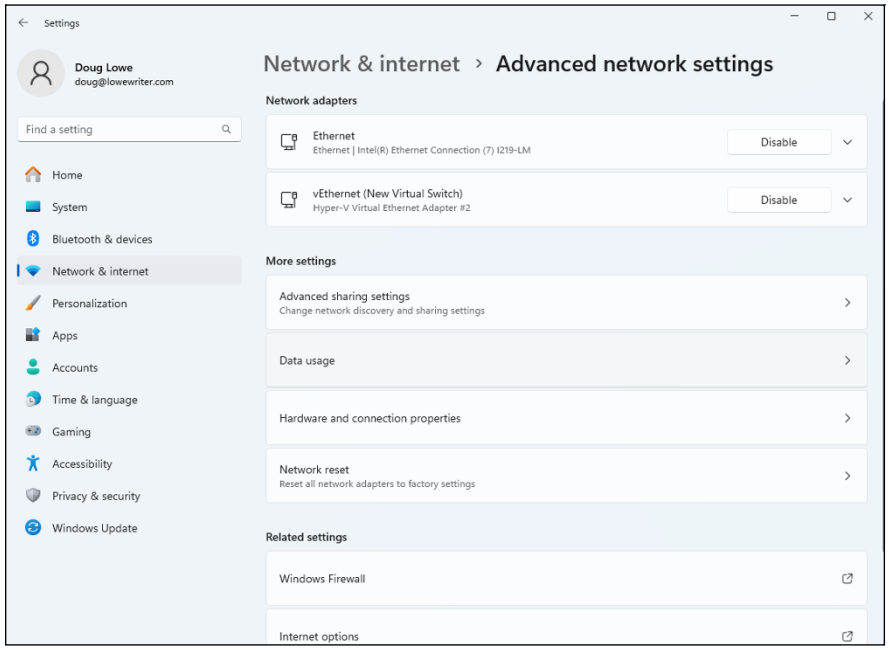


FIGURE 3-3: The Advanced Network Settings page.

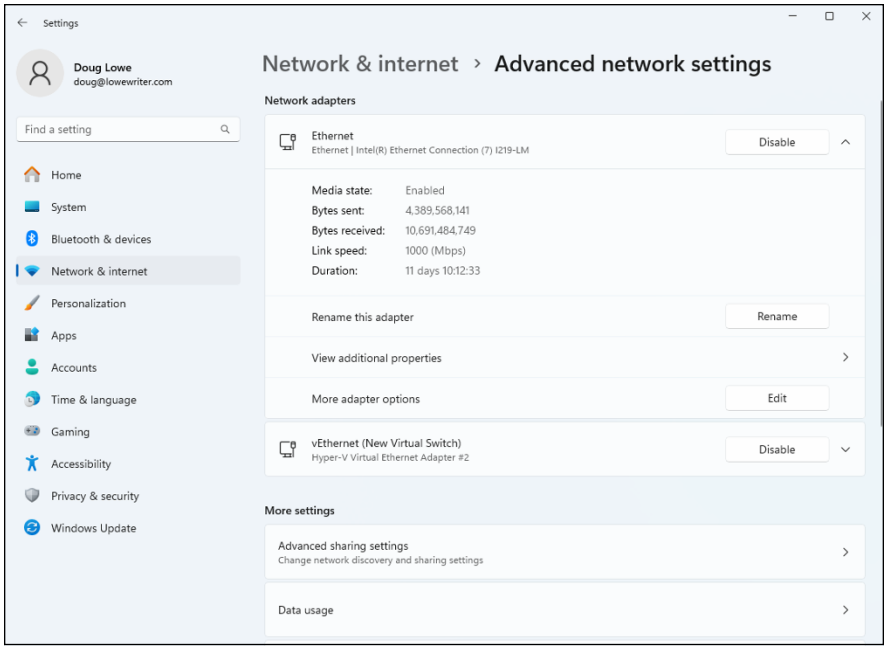


FIGURE 3-4: The Ethernet settings section of the Advanced Network Settings page.

5. **Click the Edit button next to More Adapter Options.**

   This action opens the Properties dialog box for the network adapter, as shown in Figure 3-5.

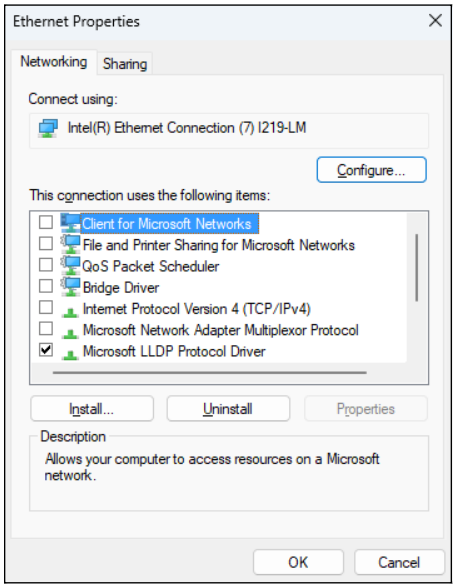


FIGURE 3-5: The network adapter Properties dialog box.

6. **To configure the network adapter card settings, click Configure.**

The Network Connection Properties dialog box appears, as shown in Figure 3-6. This dialog box has seven tabs that let you configure the adapter:

- *General:* This tab shows basic information about the adapter, such as the device type and status.
- *Advanced:* This tab lets you set a variety of device-specific parameters that affect the operation of the adapter.
- *About:* Displays information about the device's patent protection.
- *Driver:* This tab displays information about the device driver that's bound to the NIC and lets you update the driver to a newer version, roll back the driver to a previously working version, or uninstall the driver.
- *Details:* With this tab, you can inspect various properties of the adapter such as the date and version of the device driver. To view the setting of a particular property, select the property name from the drop-down list.
- *Events:* This tab lists recent events that have been logged for the device.
- *Power Management:* This tab lets you configure power management options for the device.

TIP

When you click OK to dismiss the Network Adapter Properties dialog box, the network connection's Properties dialog box closes and you are returned to the Network Connections page (refer to Figure 3-4). Click the Edit button again to continue the procedure.

FIGURE 3-6: The Properties dialog box for a network connection.

7. **Review the list of connection items listed in the Properties dialog box.**

   The most important items you commonly see are:

   - *Client for Microsoft Networks:* This item is required if you want to access a Microsoft Windows network. It should always be present.
   - *File and Printer Sharing for Microsoft Networks:* This item allows your computer to share its files or printers with other computers on the network.

     This option is usually used with peer-to-peer networks, but you can use it even if your network has dedicated servers. If you don't plan to share files or printers on the client computer, however, you should disable this item.

     Internet Protocol Version 4 (TCP/IPv4): This item enables the client computer to communicate by using the version 4 standard TCP/IP protocol.

     Internet Protocol Version 6 (TCP/IPv6): This item enables version 6 of the standard TCP/IP protocol. Typically, both IP4 and IP6 are enabled, even though most networks rely primarily on IP4.


TIP

8. **If a protocol that you need isn't listed, click the Install button to add the needed protocol.**

   A dialog box appears, asking whether you want to add a network client, protocol, or service. Click Protocol and then click Add. A list of available protocols appears. Select the one you want to add; then click OK.

9. **To remove a network item that you don't need (such as File and Printer Sharing for Microsoft Networks), select the item, and click the Uninstall button.**

   For security reasons, you should make it a point to remove any clients, protocols, or services that you don't need.

10. **To configure TCP/IP settings, click Internet Protocol Version 4 (TCP/IPv4); click Properties to display the TCP/IP Properties dialog box (shown in Figure 3-7); adjust the settings; and then click OK.**

    The TCP/IP Properties dialog box lets you choose among these options:

    - *Obtain an IP Address Automatically:* Choose this option if your network has a DHCP server that assigns IP addresses automatically. Choosing this option dramatically simplifies administering TCP/IP on your network. (See Book 2, Chapter 5 for more information about DHCP.)
    - *Use the Following IP Address:* If your computer must have a specific IP address, choose this option and then type the computer's IP address, subnet mask, and default gateway address. (For more information about these settings, see Book 2, Chapter 3.)

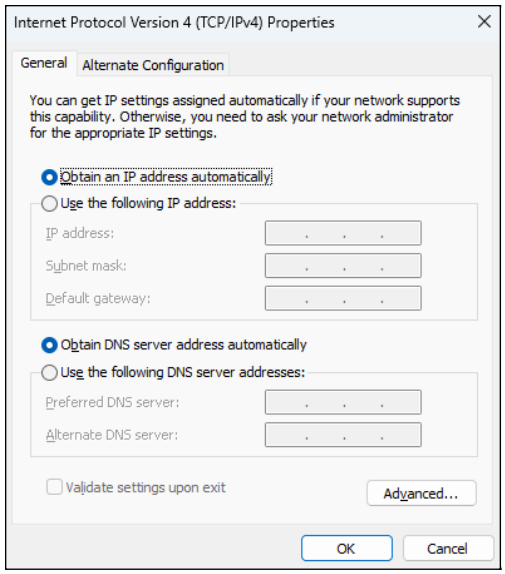


FIGURE 3-7: Configuring TCP/IP.

- *Obtain DNS Server Address Automatically:* The DHCP server can also provide the address of the Domain Name System (DNS) server that the computer should use. Choose this option if your network has a DHCP server. (See Book 2, Chapter 6, for more information about DNS.)
- *Use the Following DNS Server Addresses:* Choose this option if a DNS server isn't available. Then type the IP addresses of the primary and secondary DNS servers.

Joining a Domain

When Windows first installs, it isn't joined to a domain network. Instead, it's available as part of a workgroup, which is an unmanaged network suitable only for the smallest of networks with just a few computers and without dedicated servers. To use a computer in a domain network, you must join the computer to the domain. Here are the steps for Windows 11:



1. **Click the Start icon (or press the Start button on the keyboard), and then tap or click the Settings icon (shown in the margin).**

   The Settings page appears (refer to Figure 3-1).

2. **Click System.**

   The System settings page appears, as shown in Figure 3-8.

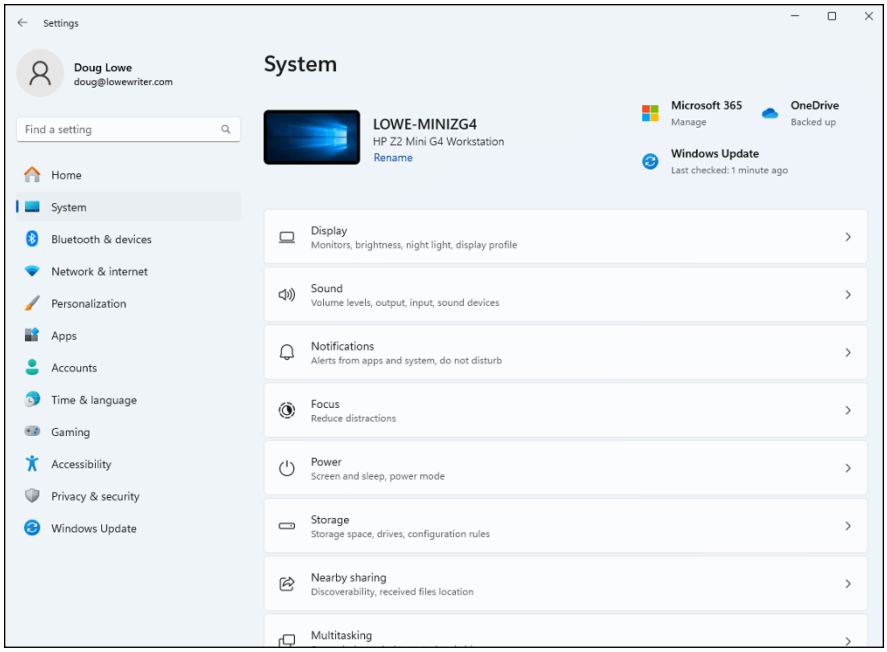


FIGURE 3-8: The System settings page.

3. **Click About.**

   You may have to scroll down a bit to find About — it's near the bottom of the page. Clicking About summons the page shown in Figure 3-9.

4. **Click Domain or Workgroup.**

   This option is listed among the Related Links beneath the Device Specifications such as the Device Name, Processor, Installed RAM, and other general information about your computer.

   When you click Domain or Workgroup, the System Properties dialog box, shown in Figure 3-10, appears.

   Before you join a domain, you should ensure that the computer's name won't be the same as the name of a computer that's already a member of the domain. If it is, you should select a different name.

5. **Click the Change button.**

   The Change Name/Domain Changes dialog box, shown in Figure 3-11, appears.

6. **If you need to change the computer name, enter it in the Computer Name field.**

   Before you join a domain, make sure that the computer's name does not conflict with an existing name that's already in the domain.

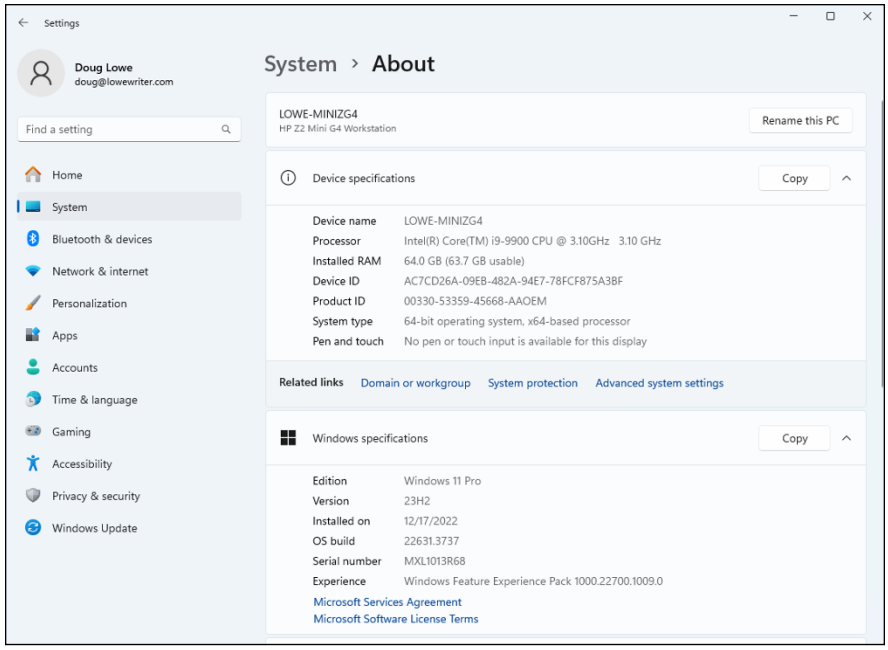


FIGURE 3-9: The About page.

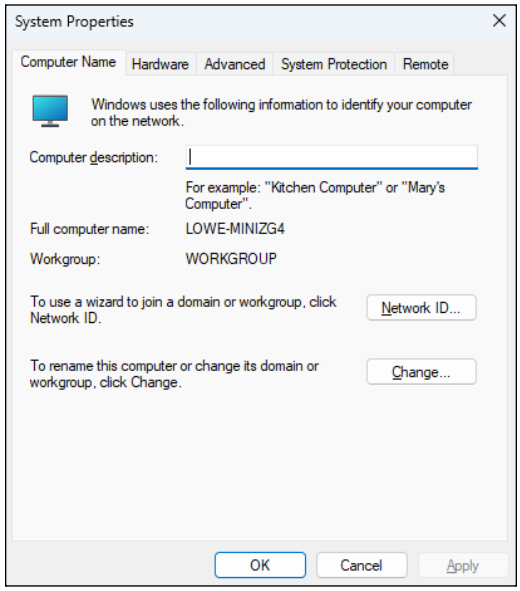


FIGURE 3-10: The System Properties dialog box.

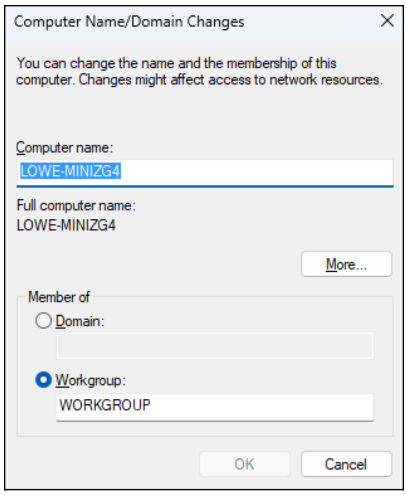


FIGURE 3-11: Joining a domain.

7. **Enter the domain name and click OK.**

   You're prompted for the username and password of a user who has administration privileges on the domain.

8. **Enter the domain administrator credentials, then click OK.**

9. **When informed that you need to restart the computer, click Restart Now.**

   The computer is restarted and added to the domain.

IN THIS CHAPTER

- Hooking up a Mac network
- Using a Mac network
- Mixing Macs and PCs

Chapter 4
Mac Networking

This book dwells on networking Windows-based computers, as though Microsoft were the only game in town. I'm sure plenty of people in Redmond, Washington (where Microsoft is headquartered) wished that it were so. But alas, there is an entirely different breed of computer: the Apple Macintosh, more commonly referred to simply as *Mac.*

Every Mac ever built, even an original 1984 model, includes networking support. Newer Mac computers have better built-in networking features than older Mac computers, of course. The newest Macs include either built-in Gigabit Ethernet connections or 802.11ax wireless connections, or both. Support for these network connections is pretty much automatic, so all you have to do is plug your Mac into a network or connect to a wireless network, and you're ready to go.

This chapter presents what you need to know to network Mac computers. You learn how to control basic Mac network options such as TCP/IP and file sharing. And you learn how to join a Mac to a Windows domain network.

Basic Mac Network Settings

Most network settings on a Mac are automatic. If you prefer, you can look at and change the default network settings by following these steps:

1. **Choose System Settings, and then choose Network.**

   The Network page appears, as shown in Figure 4-1. This page lists your available network connections.

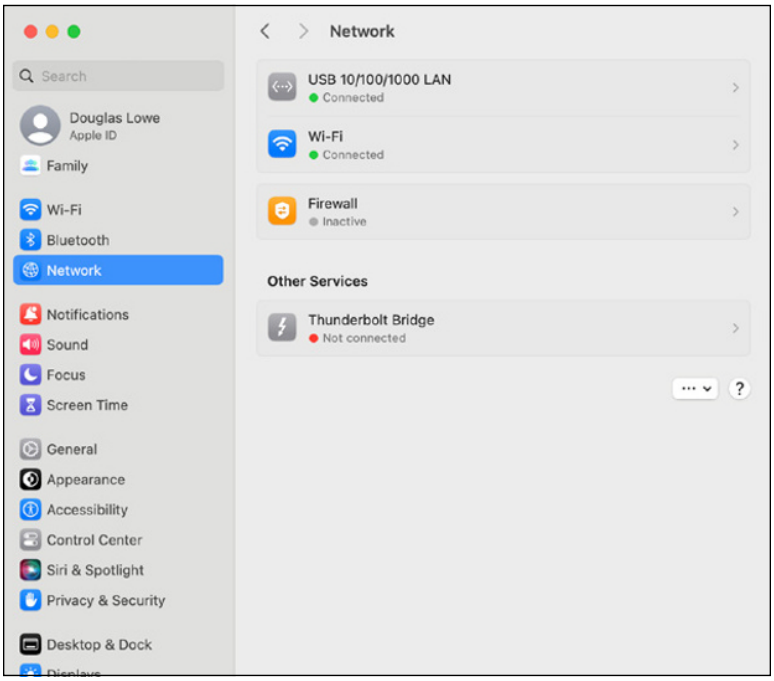


FIGURE 4-1: Network settings.



TIP

On previous versions of the macOS, *System Settings* was called *System Preferences*.

2. **Click the network you want to configure.**

   The settings for the selected network appear, as shown in Figure 4-2.

3. **Click the Details button.**

   A configuration page, shown in Figure 4-3, appears. Here, you can adjust the various settings for the connection.

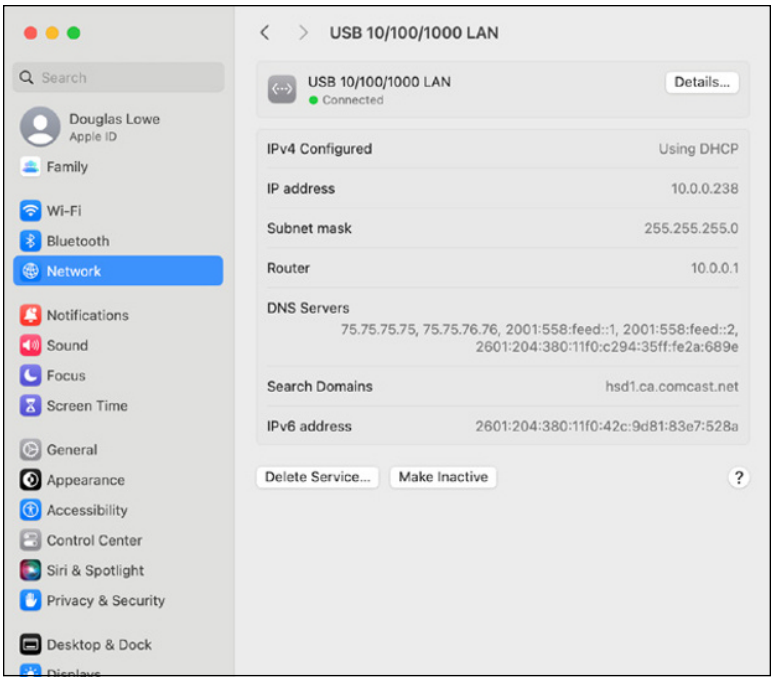


FIGURE 4-2: Connection information.

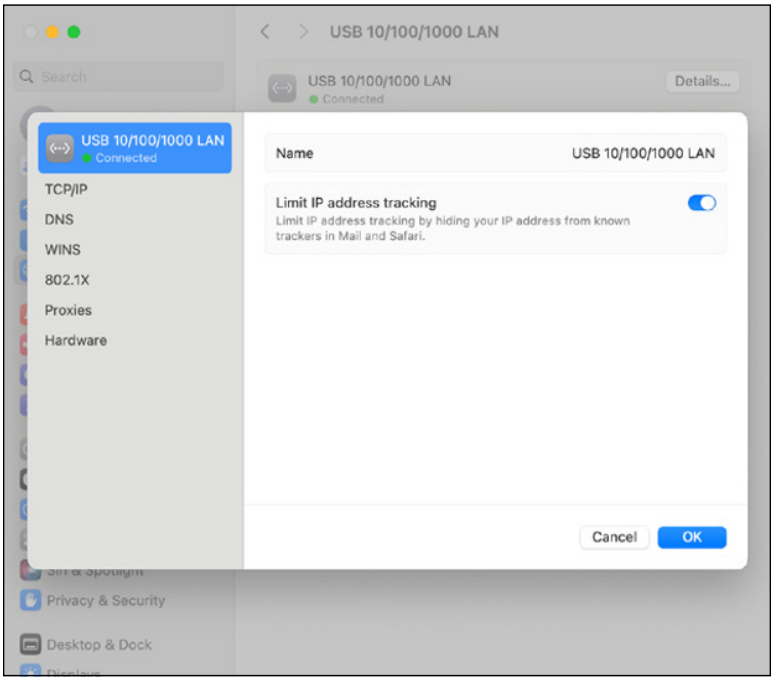


FIGURE 4-3: Adjusting the configuration for a network connection.

4. **Click the TCP/IP tab to view or change the TCP/IP settings.**

   This brings up the TCP/IP settings, as shown in Figure 4-4. From this page, you can view the currently assigned IP address for the computer. And, if you want, you can assign a static IP address by changing the Configure IPv4 drop-down setting from Using DHCP to Manually. Then, you can enter your own IP address, subnet mask, and router address.

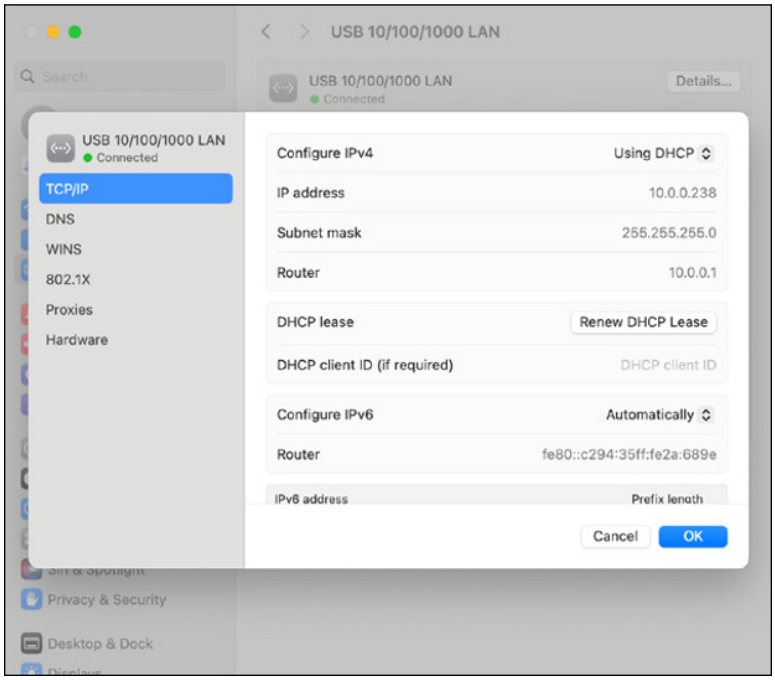


FIGURE 4-4: TCP/IP settings.

5. **Click the DNS tab to view or change the DNS settings.**

   This brings up the DNS settings shown in Figure 4-5. Here, you can see the DNS servers currently being used, and you can add additional DNS servers if you want.

6. **Click the Hardware tab to view hardware information.**

   This brings up the settings shown in Figure 4-6. The most useful bit of information on this tab is the MAC address, which is sometimes needed to set up wireless network security. (For more information, refer to Book 2, Chapter 1.)

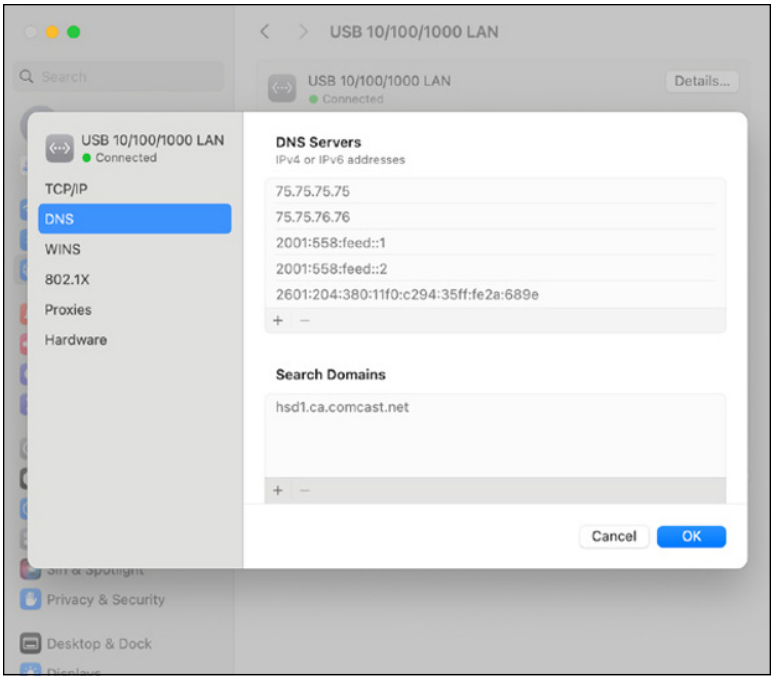


FIGURE 4-5: DNS settings.

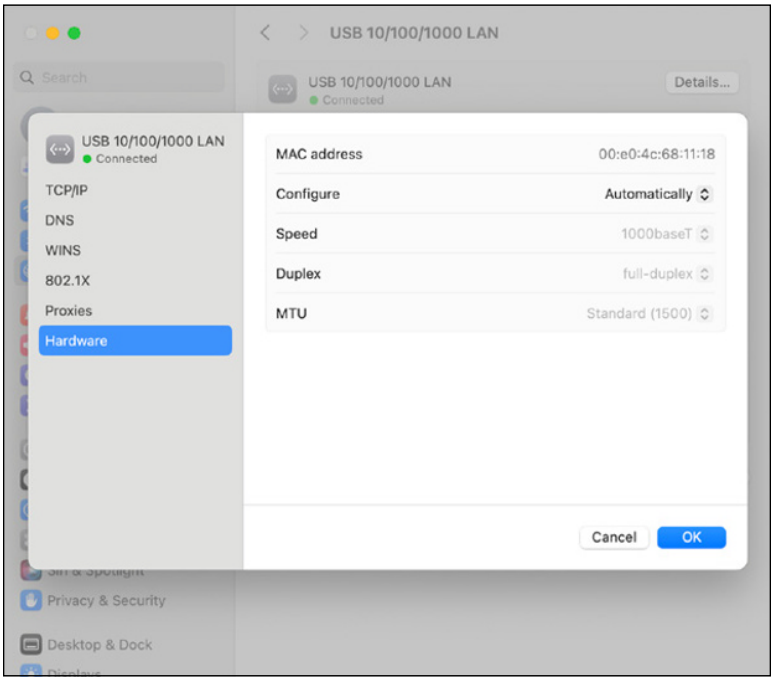


FIGURE 4-6: Hardware settings.

Joining a Domain

If you are using a Mac in a Windows domain environment, you can join the Mac to the domain by following these steps:

1. **Choose Settings, then choose Users & Groups.**

   This brings up the Users & Groups page, as shown in Figure 4-7.

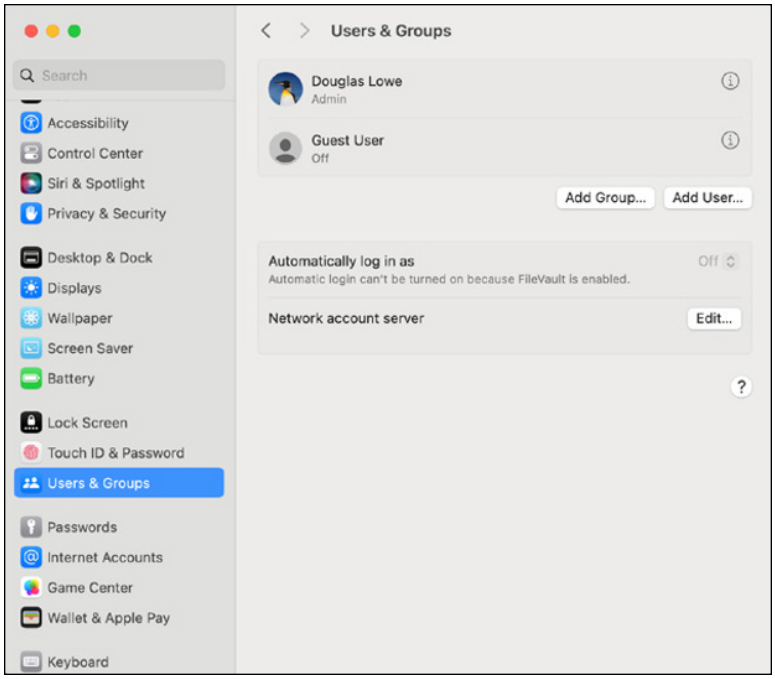


FIGURE 4-7: Users & Groups.

2. **Click the Edit button (next to Network Account Server), and then click the Open Directory Utility button.**

   When you click the Edit button, a small dialog box prompts you to add a network server. Instead of adding a server here, just click the Open Directory Utility button. The Directory Utility page, shown in Figure 4-8, appears.

3. **Click the lock icon at the lower-left corner of the page and enter your password when prompted.**

   By default, the user login options are locked to prevent unauthorized changes. This step unlocks the settings so that you can join the domain.

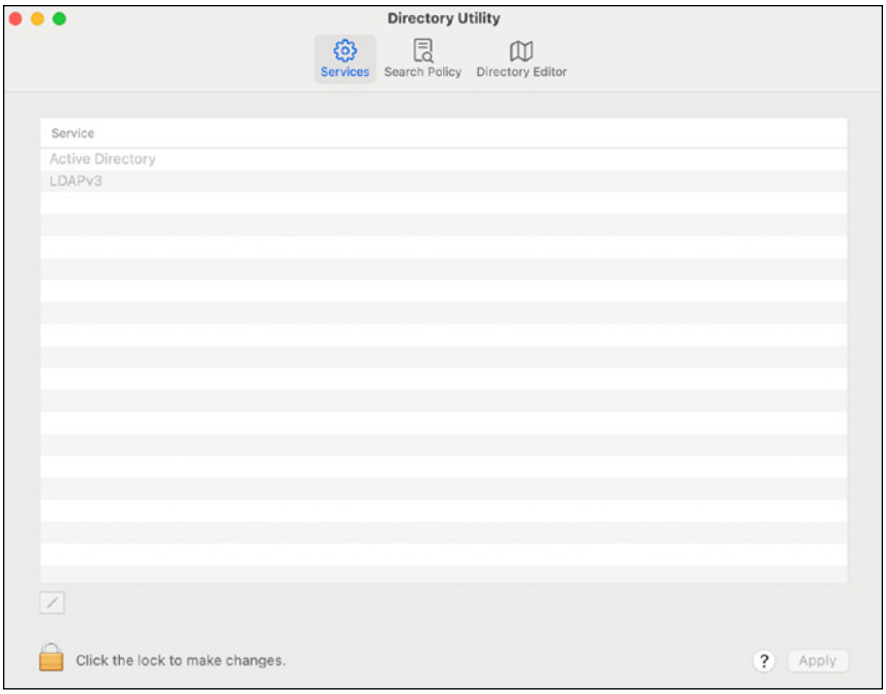


FIGURE 4-8: The Directory Utility.

4. **Double-click Active Directory.**

   You are prompted to enter the name of the domain you want to join, as shown in Figure 4-9.

5. **Enter the name of the domain you want to join and click the Bind button.**

   You're prompted to enter domain administrator credentials to allow you to join the domain, as shown in Figure 4-10.

6. **Enter the name and password of a domain administrator account, then click OK.**

   You're returned to the Login Options page, which displays a final congratulatory screen to let you know that you've successfully joined the domain.

7. **Close the Users & Groups window.**

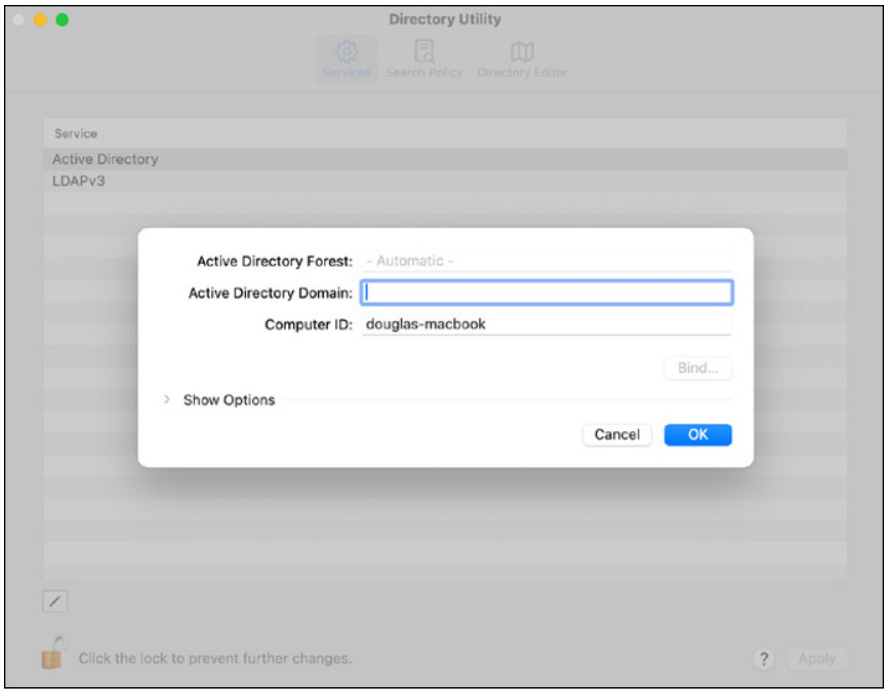


FIGURE 4-9: Joining a domain.

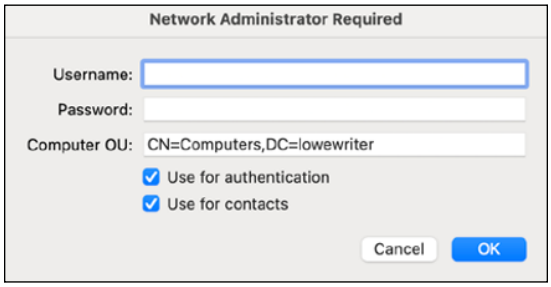


FIGURE 4-10: Authenticating with the domain.

Connecting to a Share

Once you have joined a domain, you can access its network shares via the Finder. Just follow these steps:

1. **Click Finder.**

   This opens the Finder, as shown in Figure 4-11.

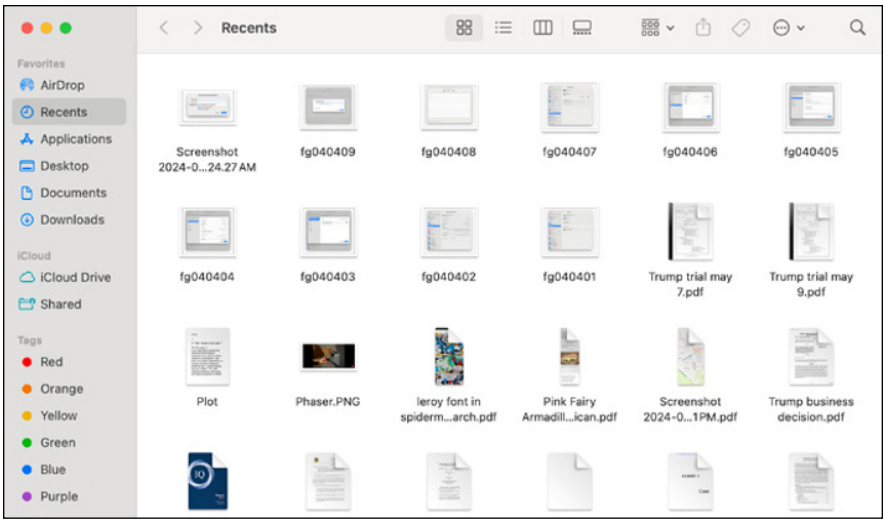


FIGURE 4-11: Welcome to the Finder.

2. Choose Go➪Finder.

The Connect to Server dialog box appears, as shown in Figure 4-12.

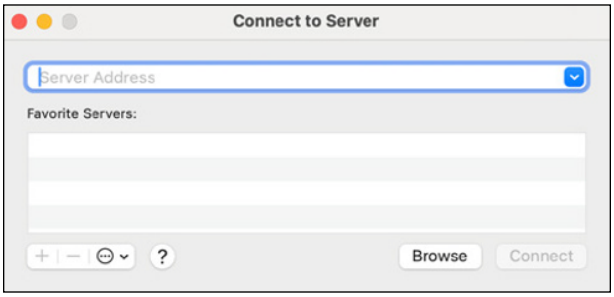


FIGURE 4-12: The Connect to Server dialog box.

3. Type the smb path that leads to the server share you want to connect to.

To type a smb path, follow this syntax:

```
smb://server-name/share-name
```

Replace the *server-name* with the name of the server that contains the share and *share-name* with the name of the share. For example, to connect to a share named `files` on a server named `lowe01`, type **smb://lowe01/files**.

4. Click Connect.

You'll be prompted for login credentials.

5. **Enter a domain username and password, then click OK.**

   Precede the username with the domain name, separated by a backslash. For example, if the domain name is `lowewriter.pri` and the username is Doug, enter **lowewriter.pri\Doug** as the username.

   Once connected, the files in the share will be displayed in the Finder window. You can then open files directly from the share (provided you have the right software, such as Microsoft Office, to read the files). You can also drag and drop files between the Mac and the file shares.

IN THIS CHAPTER

- Setting up network printers
- Using a network printer's web interface

Chapter 5
Network Printers

After you have your network servers and clients up and running, you still have many details to attend to before you can pronounce your network "finished." In this chapter, you discover a few more configuration chores that have to be done: configuring internet access, setting up network printers, configuring email, and configuring mapped network drives.

Configuring Network Printers

Before network users can print on the network, the network's printers must be properly configured. For the most part, this task is a simple one. All you have to do is configure each client that needs access to the printer.



TIP

Before you configure a network printer to work with network clients, read the client configuration section of the manual that came with the printer. Many printers come with special software that provides more advanced printing and networking features than the standard features provided by Windows. If so, you may want to install the printer manufacturer's software on your client computers rather than use the standard Windows network printer support.

Adding a network printer

The exact procedure for adding a network printer varies a bit, depending on the Windows version that the client runs. The following steps describe the procedure for Windows 11; the procedure for previous versions of Windows is similar:



1. **Click the Start icon (or press the Start button on the keyboard), and then tap or click Settings.**

   The Settings page appears, as shown in Figure 5-1.

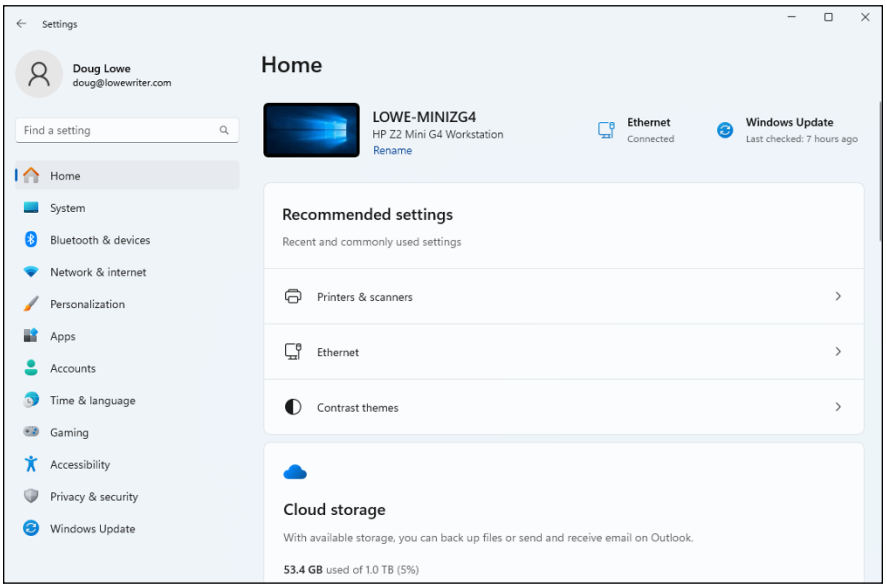


FIGURE 5-1: The Settings page.

2. **Click Printers and Scanners.**

   If you don't see the Printers and Scanners option, type **Printer** in the Find a Setting text box; Printers and Scanners will appear as a selection.

   Either way, the Printers and Scanners page, shown in Figure 5-2, appears.

3. **If the printer you want to add appears in the list, click its Add Device button, wait a moment for the printer to be added, and then skip the rest of this procedure.**

   You're done!

4. **Otherwise, click Add Manually (next to The Printer That I Want Isn't Listed).**

   The first page of the Add Printer Wizard appears, as shown in Figure 5-3.

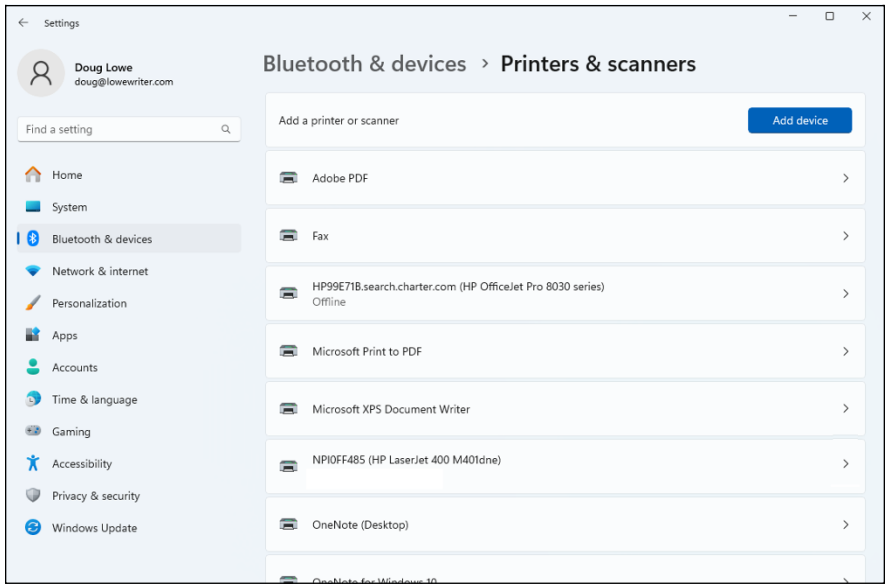


FIGURE 5-2: The Printers and Scanners page.

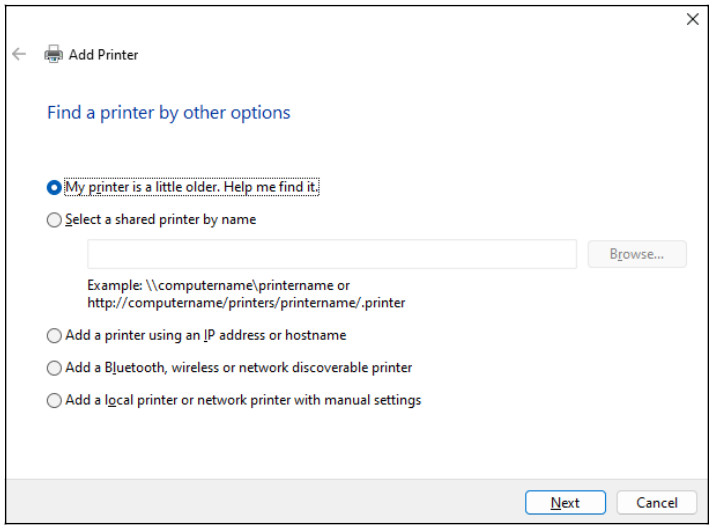


FIGURE 5-3: The Add Printer wizard comes to life.

5. Determine the IP address of the printer you want to add.

If you're uncertain, use the display panel that's physically on the printer to find the IP address. For this example, the IP address of the printer I want to add is 10.0.0.145.

6. **Select Add a Printer Using an IP Address or Host Name, and then click Next.**

   The Add Printer wizard prompts you for the printer's host name or IP address, as shown in Figure 5-4. (Note that Figure 5-4 shows this page after I've already selected TCP/IP Device from the Device Type drop-down, as directed in the next step.)

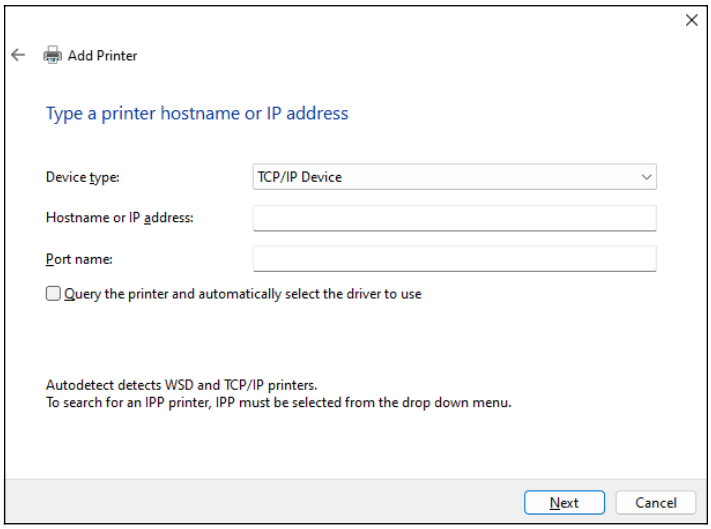


FIGURE 5-4: The Add Printer wizard asks for a hostname or IP address.

7. **Select TCP/IP Device from the Device Type drop-down list, enter the IP address in the Hostname or IP Address text box, and click Next.**

   The Add Printer wizard installs the printer.

8. **If the Install the Printer Driver page appears, select the correct printer driver and then click Next.**

   This page of the wizard, shown in Figure 5-5, appears only if the driver isn't already installed on your computer. Select the printer driver, and then click Next. The Type a Printer Name page appears, as shown in Figure 5-6.

9. **Enter a name for the printer (or accept the default name provided) and click Next.**

   The next page of the wizard invites you to share the printer with other users, as shown in Figure 5-7.

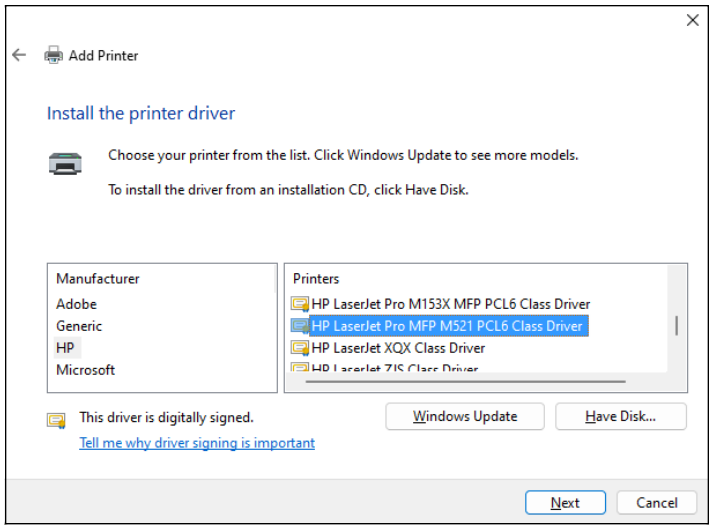


FIGURE 5-5: The Add Printer wizard asks you to select a printer driver.

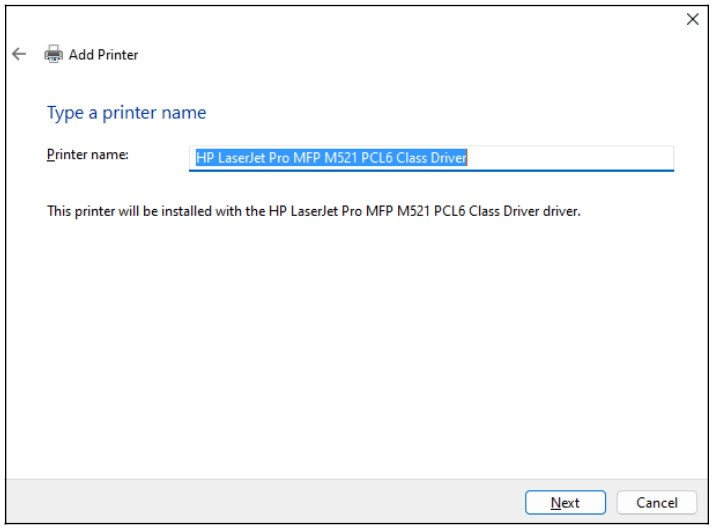


FIGURE 5-6: The Add Printer wizard asks you to provide a name for the printer.

10. Click Next to ignore printer sharing.

If other users need access to the printer, they should install it themselves. There's no reason to share a printer you've connected to via TCP/IP.

When you click Next, a confirmation page appears, congratulating you for installing the printer.

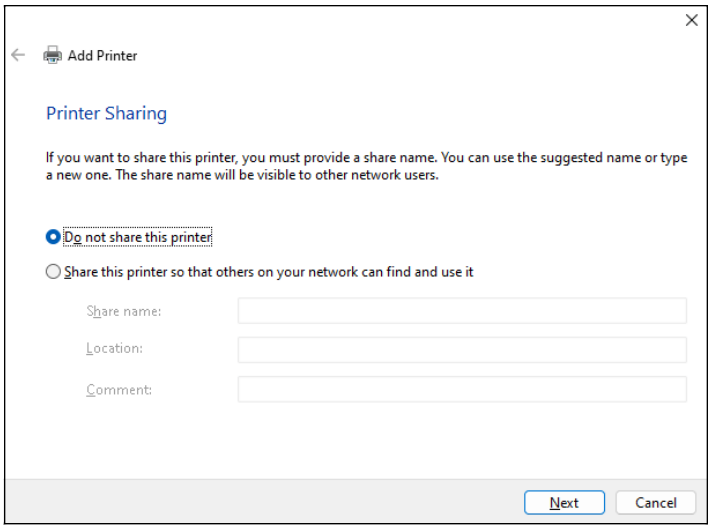


FIGURE 5-7: Do you want to share the printer?



TIP

If you want to, you can click the Print a Test Page button to confirm that the printer is working. The test page will include helpful information such as the name and IP address of the printer, the date and time the printer was installed, and details about the printer driver.

You're done! The network printer has been added to the client computer.

Accessing a network printer using a web interface



TIP

Printers that have a direct network connection often include a built-in web server that lets you manage the printer from any browser on the network. Figure 5-8 shows the home page for an HP LaserJet 400 M401dne printer. This web interface lets you view status information about the printer and check the printer's configuration. You can even view error logs to find out how often the printer jams.

To call up a printer's web interface, enter its IP address or host name in the address bar of any web browser.

In addition to simply displaying information about the printer, you can adjust the printer's configuration from a web browser. Figure 5-9 shows the networking settings page for the HP printer. Here, you can view and change the network configuration details, such as the TCP/IP host name, IP address, subnet mask, domain name, and so on.

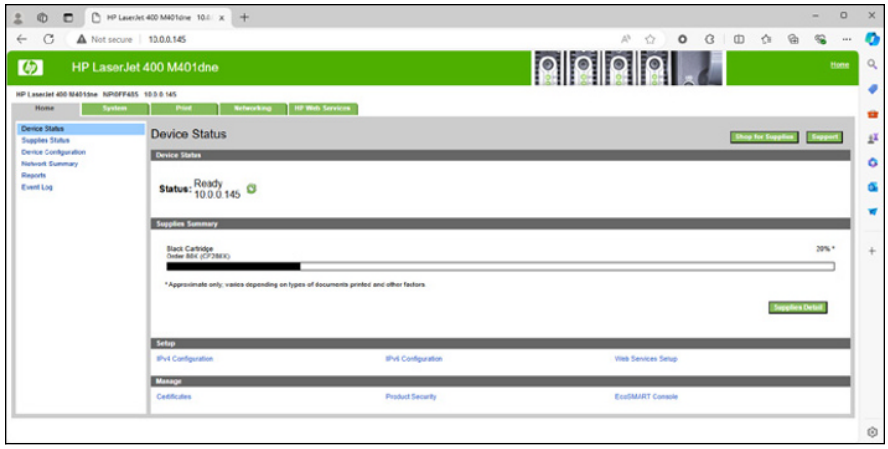


FIGURE 5-8: Using a printer's web interface.

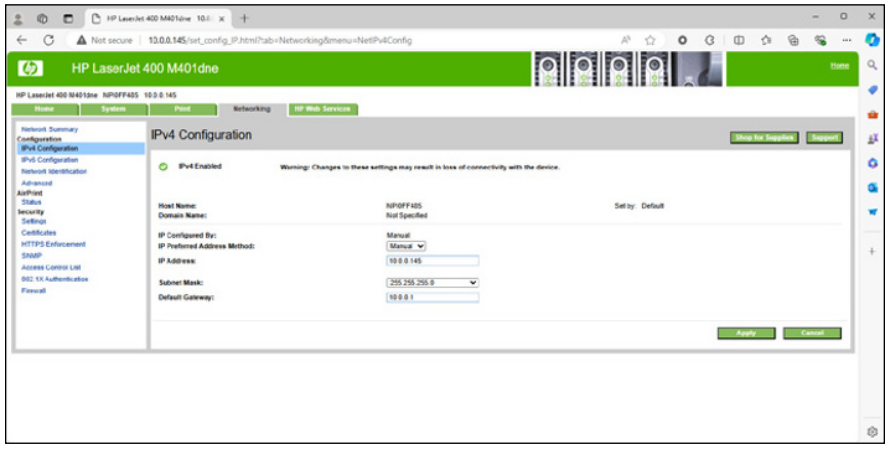


FIGURE 5-9: Changing network settings via a printer's web interface.



TIP

As the network administrator, you may need to visit the printer's web page frequently. I suggest that you add it to your browser's Favorites menu so that you can get to it easily. If you have several printers, add them to a folder named Network Printers.

IN THIS CHAPTER

- Examining VPN uses
- Looking at how VPN works
- Considering VPN clients and servers
- Pondering VPN hardware and software

Chapter 6
Virtual Private Networks

Today's network users frequently need to access their networks from remote locations: home offices, hotel rooms, beach villas, and their kid's soccer fields. In the early days of computer networking, the only real option for remotely accessing a network was to set up dialup access with telephone lines and modems, which was slow and unreliable. Today, enabling remote access to a local area network (LAN) is easily done with a virtual private network. Simply put, a virtual private network (VPN) enables remote users to access a LAN via any internet connection.

This chapter is a short introduction to VPNs. You find out the basics of what a VPN is, how to set one up, and how to access one remotely. Enjoy!

Understanding VPN

A *virtual private network* (VPN) is a type of network connection that creates the illusion that you're directly connected to a network when in fact, you're not. For example, suppose you set up a LAN at your office, but you also occasionally work from home. But how will you access the files on your work computer from home?

» You could simply copy whatever files you need from your work computer onto a flash drive and take them home with you, work on the files, copy the updated files back to the flash drive, and take them back to work with you the next day.

- You could email the files to your personal email account, work on them at home, and then email the changed files back to your work email account.
- You could get a laptop and use the Windows Offline Files feature to automatically synchronize files from your work network with files on the laptop.

Or, you could set up a VPN that allows you to log on to your work network from home. The VPN uses a secured internet connection to connect you directly to your work network, so you can access your network files as if you had a really long Ethernet cable that ran from your home computer all the way to the office and plugged directly into the work network.

There are at least three situations in which a VPN is the ideal solution:

- One or more workers need to occasionally work from home (as in the scenario just described). In this situation, a VPN connection establishes a connection between the home computer and the office network.
- Mobile users — who may not ever actually show up at the office — need to connect to the work network from mobile computers, often from locations like hotel rooms, clients' offices, airports, or coffee shops. This type of VPN configuration is similar to the home user's configuration, except that the exact location of the remote user's computer is not fixed.
- Your company has offices in two or more locations, each with its own LAN, and you want to connect the locations so that users on either network can access each other's network resources. In this situation, the VPN doesn't connect a single user with a remote network; instead, it connects two remote networks to each other.

ACCESSING YOUR COMPUTER REMOTELY

One of the most common reasons for setting up a VPN connection is to allow you to access a computer that is on your work network from a computer that is outside of your work network. For example, as a network administrator, you can use a VPN to connect to your work network from home. You can then use *Remote Desktop Connection* (RDC) to access your servers.

Before you can use RDC to connect remotely, you must enable remote access on the server computer. Right-click Computer on the Start menu and choose Properties, and then click Advanced System Settings. Click the Remote tab, and then select Allow Connections Only from Computers Running Remote Desktop with Network Level

Authentication. This option lets you grant remote access only to specific users, whom you can designate by clicking the Select Users button.

After remote access has been granted, you can access the computer remotely by connecting to the network with a VPN. Then choose Start➪ All Programs➪ Remote Desktop Connection. Enter the name of the computer you want to connect to, and then click Connect. You're prompted for your Windows username and password. After it's all connected, you can access the remote computer's desktop in a window.

Looking at VPN Security

The *V* in VPN stands for *virtual*, which means that a VPN creates the appearance of a local network connection when in fact the connection is made over a public network — the internet. The term "tunnel" is sometimes used to describe a VPN because the VPN creates a tunnel between two locations, which can only be entered from either end. The data that travels through the tunnel from one end to the other is secure as long as it is within the tunnel — that is, within the protection provided by the VPN.

The *P* in VPN stands for *private*, which is the purpose of creating the tunnel. If the VPN didn't create effective security so that data can enter the tunnel only at one of the two ends, the VPN would be worthless; you may as well just open your network and your remote computer up to the internet and let the hackers have their way.

Prior to VPN technology, the only way to provide private remote network connections was through actual private lines, which were (and still are) very expensive. For example, to set up a remote office you could lease a private T1 line from the phone company to connect the two offices. This private T1 line provided excellent security because it physically connected the two offices and could be accessed only from the two endpoints.

VPN provides the same point-to-point connection as a private leased line, but does it over the internet instead of through expensive dedicated lines. To create the tunnel that guarantees privacy of the data as it travels from one end of the VPN to the other, the data is encrypted using special security protocols.

The most important of the VPN security protocols is *Internet Protocol Security* (IPSec), which is a collection of standards for encrypting and authenticating packets that travel on the internet. In other words, it provides a way to encrypt the contents of a data packet so that only a person who knows the secret encryption keys can decode the data. And it provides a way to reliably identify the source

of a packet so that the parties at either end of the VPN tunnel can trust that the packets are authentic.



Referring to the OSI Reference Model (see Book 2, Chapter 1), the IPSec protocol operates at layer 3 of the OSI model (the network layer). What that means is that the IPSec protocol has no idea what kind of data is being carried by the packets it encrypts and authenticates. The IPSec protocol concerns itself only with the details of encrypting the contents of the packets (sometimes called the *payload*) and ensuring the identity of the sender.

Another commonly used VPN protocol is Layer 2 Tunneling Protocol (L2TP). This protocol doesn't provide data encryption. Instead, it's designed to create end-to-end connections — *tunnels* — through which data can travel. L2TP is actually a combination of two older protocols: Layer 2 Forwarding Protocol (L2FP, from Cisco) and Point-to-Point Tunneling Protocol (PPTP, from Microsoft).

Many VPNs today use a combination of L2TP and IPSec: L2TP Over IPSec. This type of VPN combines the best features of L2TP and IPSec to provide a high degree of security and reliability.

Understanding VPN Servers and Clients

A VPN connection requires a VPN server — the gatekeeper at one end of the tunnel — and a VPN client at the other end. The main difference between the server and the client is that the client initiates the connection with the server, and a VPN client can establish a connection with just one server at a time. However, a server can accept connections from many clients.

Typically, the VPN server is a separate hardware device, most often a security appliance such as a Cisco ASA security appliance. VPN servers can also be implemented in software. For example, Windows Server 2008 includes built-in VPN capabilities even though they're not easy to configure. And a VPN server can be implemented in Linux as well.

Figure 6-1 shows one of the many VPN configuration screens for a typical Sophos firewall. This screen provides the configuration details for an IPSec VPN connection. The most important item of information on this screen is the Pre-Shared Key, which is used to encrypt the data sent over the VPN. The client will need to provide the identical key in order to participate in the VPN.

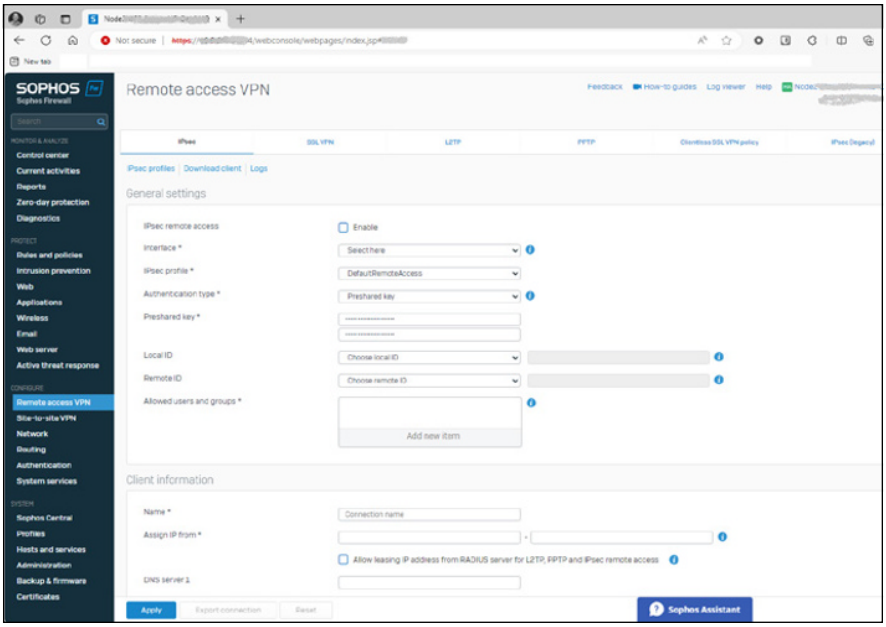


FIGURE 6-1: An IPSec configuration page on a Sophos firewall.



REMEMBER

A VPN client is usually software that runs on a client computer or other device such as a smartphone or tablet that wants to connect to the remote network. The VPN client software must be configured with the IP address of the VPN server as well as authentication information such as a username and the Pre-Shared Key that will be used to encrypt the data. If the key used by the client doesn't match the key used by the server, the VPN server will reject the connection request from the client.

Figure 6-2 shows a typical VPN software client. When the client is configured with the correct connection information (which you can do by importing a configuration file you download from your organization's web-based VPN portal), you just click Connect. After a few moments, the VPN client will announce that the connection has been established and the VPN is connected.

A VPN client can also be a hardware device, like another security appliance. This is most common when the VPN is used to connect two networks at separate locations. For example, suppose your company has an office in Pixley and a second office in Hooterville. Each office has its own network with servers and client computers. The easiest way to connect these offices with a VPN would be to put an identical security appliance at each location. Then, you could configure the security appliances to communicate with each other over a VPN.

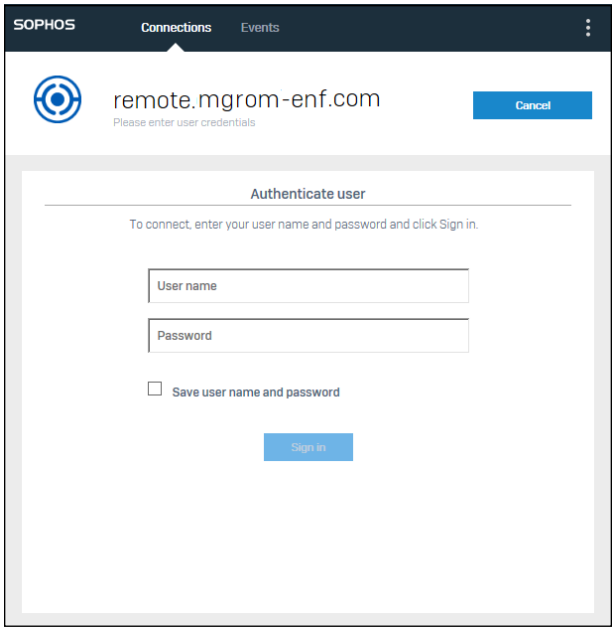


FIGURE 6-2:
A VPN client.

5 Implementing Virtualization

Contents at a Glance

CHAPTER 1: **Hyper-V** 393

CHAPTER 2: **VMware** 413

CHAPTER 3: **Azure** 425

CHAPTER 4: **Amazon Web Services** 441

CHAPTER 5: **Desktop Virtualization** 459

IN THIS CHAPTER

- Introducing Hyper-V
- Working with Hyper-V Manager
- Creating a virtual network to support your virtual machines
- Creating virtual machines in Hyper-V

Chapter 1
Hyper-V

Hyper-V is a virtualization platform that comes as a standard part of all versions of Windows Server since version 2008 and all versions of desktop Windows since Windows 8.

On server versions of Windows, Hyper-V provides an enterprise-grade true Type-1 hypervisor that can manage huge virtualization farms with thousands of virtual machines. The version of Hyper-V that comes with desktop Windows is called Client Hyper-V. Client Hyper-V uses the same Type-1 hypervisor as the server-grade Hyper-V. However, it doesn't have the same enterprise-level management capabilities, because it's intended for use on client computers, not for production servers.

In this chapter, I first explain some of the details of how Hyper-V works. Then I show you how to set up virtual machines using the client version of Hyper-V. That way, you can build your own virtual machines to experiment with Hyper-V without the need for expensive hardware or server software.

Understanding the Hyper-V Hypervisor

Hyper-V is a built-in component of all modern versions of Windows. So, to use Hyper-V, you don't need to purchase any additional software from Microsoft. If you own a modern Microsoft operating system, you already own Hyper-V.



REMEMBER

Don't be confused by the fact that Hyper-V is an integral part of Windows: Although Hyper-V is built into Windows, Hyper-V is *not* a Type-2 hypervisor that runs as an application within Windows. Instead, Hyper-V is a true Type-1 hypervisor that runs directly on the host computer hardware. This is true even for the Client Hyper-V versions that are included with desktop versions of Windows.

In Hyper-V, each virtual machine runs within an isolated space called a *partition.* Each partition has access to its own processor, RAM, disk, network, and other virtual resources.

There are two types of partitions in Hyper-V: a *parent partition* and one or more *child partitions.* The parent partition is a special partition that hosts the Windows operating system that Hyper-V is associated with. Child partitions host additional virtual machines that you create as needed.

When you activate the Hyper-V feature, the hypervisor is installed and the existing Windows operating system is moved into a virtual machine that runs in the parent partition. Then, whenever you start the host computer, the hypervisor is loaded, the parent partition is created, and Windows is started in a virtual machine within the parent partition.

Although it may appear that the hypervisor is running within Windows, actually the reverse is true: Windows is running within the hypervisor.

In addition to the Windows operating system, the parent partition runs software that enables the management of virtual machines on the hypervisor. This includes creating new virtual machines, starting and stopping virtual machines, changing the resources allocated to existing virtual machines (for example, adding more processors, RAM, or disk storage), and moving virtual machines from one host to another.

Understanding Hyper-V Virtual Disks

Every Hyper-V virtual machine must have at least one virtual disk associated with it. A *virtual disk* is nothing more than a disk file that resides in the file system of the host operating system. The file has one of two file extensions, depending on which of two data formats you choose for the virtual disk:

- `.vhd`: An older format that has a maximum virtual disk size of 2TB
- `.vhdx`: A newer format that can support virtual disks up to 64TB

For either of these virtual disk formats, Hyper-V lets you create two different types of virtual disks:

- **Fixed-size disk:** A virtual disk whose disk space is pre-allocated to the full size of the drive when you create the disk. For example, if you create a 100GB fixed-size disk using the .vhdx format, a .vhdx file of 100GB will be allocated to the drive. Even if the drive contains only 10GB of data, it will still consume 100GB of space on the host system's disk drive.
- **Dynamically expanding disk:** A virtual disk that has a maximum disk space, but that actually consumes only the amount of disk space that is required to hold the data on the disk. For example, if you create a dynamically expanding disk with a maximum of 100GB but only put 10GB of data on it, the .vhdx file for the disk will occupy just 10GB of the host system's disk drive.


TECHNICAL STUFF

Actually, there's a third type of disk called a *differencing disk*, which can be used to track changes made to another virtual disk. But this is an advanced topic that I don't cover in this chapter.


TIP

Don't be confused by the names *fixed-size* and *dynamically expanding*. Both types of disks can be expanded later if you run out of space. The main difference is whether the maximum amount of disk space allowed for the drive is allocated when the drive is first created or as needed when data is added to the drive. Allocating the space when the drive is created results in better performance for the drive, because Hyper-V doesn't have to grab more disk space every time data is added to the drive. Both types of drives can be expanded later if necessary.

Enabling Hyper-V

Hyper-V is not automatically enabled when you install Windows; you must first enable this feature before you can use Hyper-V.

To enable Hyper-V on a server version of Windows, call up the Server Manager and open the Add Roles and Features Wizard. Then enable the Hyper-V role. When you complete the wizard, Hyper-V will install the Type-1 hypervisor and move the existing Windows Server operating system into the parent partition. You can then start building virtual machines.

To enable Hyper-V on a desktop version of Windows, follow these steps:

1. **Open the Control Panel.**
2. **Choose Programs and Features.**

   The Programs and Features window comes to life.
3. **Click Turn Windows Features On or Off.**

   The Windows Features dialog box appears, as shown in Figure 1-1.

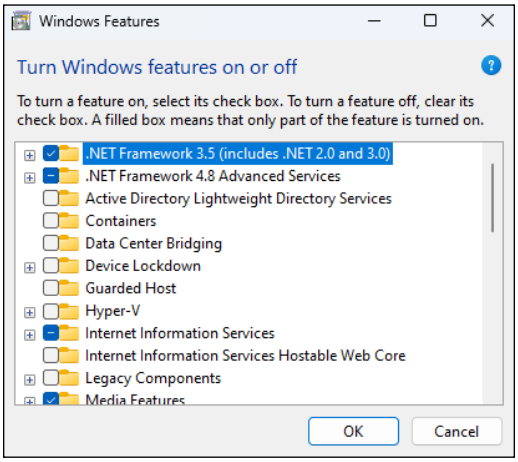


FIGURE 1-1: Enabling Hyper-V on a desktop version of Windows.

4. **Select the Hyper-V feature and click OK.**

   The Client Hyper-V hypervisor is installed as an application on the existing desktop Windows operating system, and you can begin using Hyper-V.
5. **When prompted, restart the computer.**

   The reboot is required to start the Hyper-V hypervisor. When your computer restarts, it's actually the Hyper-V hypervisor that starts, not Windows. The hypervisor then loads your desktop Windows into the parent partition.

Getting Familiar with Hyper-V

To manage Hyper-V, you use the Hyper-V Manager, shown in Figure 1-2. To start this program, click the Start button, type **Hyper-V**, and then choose Hyper-V Manager.

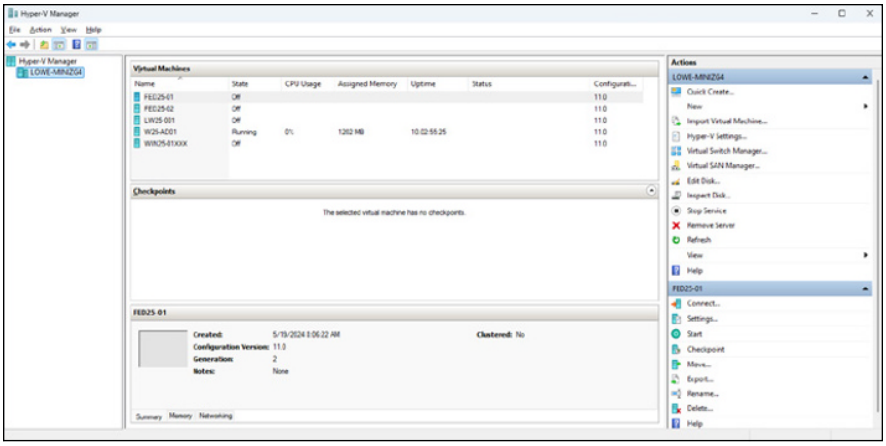


FIGURE 1-2: Hyper-V Manager.

The Hyper-V Manager window is divided into five panes:

- **Navigation:** On the left side of the window is a navigation pane that lists the Hyper-V hosts, which Hyper-V calls *virtualization servers.* In Figure 1-2, just one host is listed: my Windows computer. In an enterprise environment where you have more than one host, each of the hosts will be listed in this pane.
- **Virtual Machines:** This pane lists the virtual machines that are defined for the selected host. In Figure 1-2, you can see several of the Hyper-V virtual machines that I created while I wrote this book, including several Linux machines and several Windows Server 2019 machines. All of these machines are currently turned off.
- **Checkpoints:** In Hyper-V, a *checkpoint* is a recovery point for a virtual machine. You can create a checkpoint when you're going to make a modification to a virtual machine. Then, if something goes wrong, you can revert to the checkpoint. There are no checkpoints shown in Figure 1-2.
- **Virtual machine summary pane:** Below the Checkpoints pane is a pane that provides summary information for the virtual machine selected in the Virtual Machines pane. In Figure 1-2, you can see the summary information for a Fedora virtual machine. This pane has three tabs: Summary, Memory, and Networking. In the figure, the Memory tab is selected so you can see the memory that has been allocated to the machine.
- **Actions:** The Actions tab contains buttons you can click to initiate actions for the selected host (DOUG-WIN10) and the selected machine (FED32-01).

Creating a Virtual Switch

Before you start creating virtual machines in Hyper-V, you should create a virtual switch so that your virtual machines can communicate with each other and with the outside world. To do that, you use the Virtual Switch Manager. Here are the steps:

1. **In Hyper-V Manager, click Virtual Switch Manager.**

   This brings up the Virtual Switch Manager window, as shown in Figure 1-3.

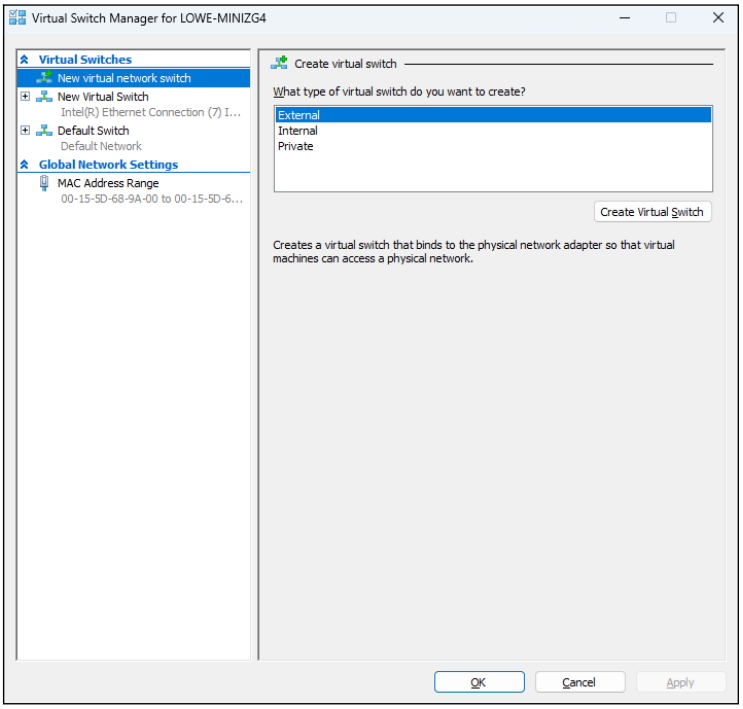


FIGURE 1-3: The Virtual Switch Manager window.

2. **Select the type of virtual switch you want to create.**

   Hyper-V lets you create three types of switches:

   - **External:** A virtual switch that binds to a physical network adapter, which allows virtual machines to communicate with each other, as well as with other computers on your physical network. This is usually the type of switch you should create.

- **Internal:** A virtual switch that does not bind with a physical network adapter. This type of switch lets the virtual machines on this computer communicate with each other and with the host computer, but not with other computers on your physical network.
- **Private:** A virtual switch that lets virtual machines communicate with each other but not with the host computer or with any computers on your physical network.

3. **Click Create Virtual Switch.**

   The settings for the new virtual switch appear, as shown in Figure 1-4.

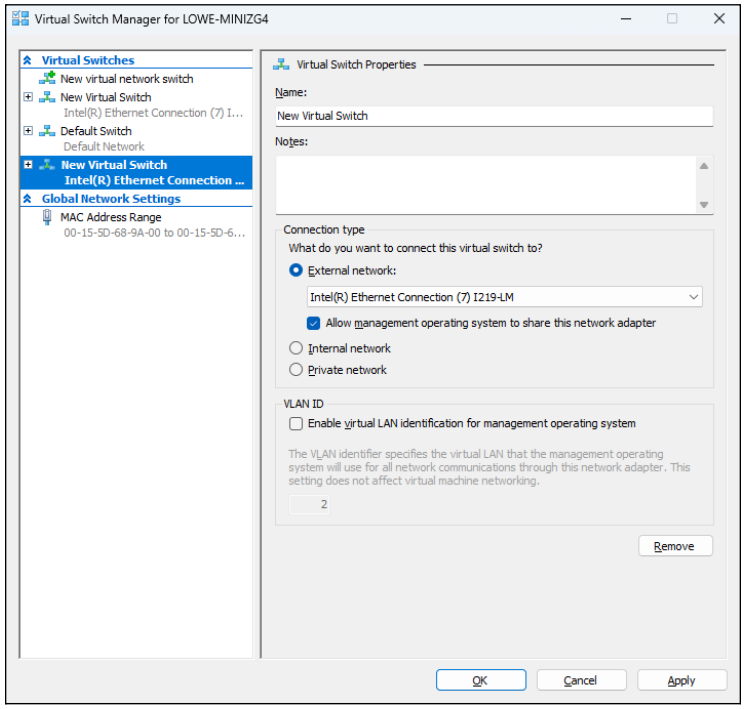


FIGURE 1-4: Creating a new virtual switch.

4. **Type a name for the new virtual switch in the Name field.**

   Use any name you want.

5. **Select the physical network adapter you want to bind the virtual switch to.**

   If your computer has more than one network adapter, select the one you want to use. Binding the virtual switch to a physical network adapter allows the virtual machines to communicate not only with each other but also with other computers connected via the adapter you select.

6. **If your network has multiple VLANs, click the Enable Virtual LAN Identification check box and enter the VLAN ID for the VLAN you want this switch to connect to.**

   If your network does not have multiple VLANs, you can skip this step.

7. **Click OK.**

   The virtual switch is created. Your Hyper-V environment now has a virtual network in place, so you can start creating virtual machines.

Creating a Virtual Disk

Before you create a virtual machine, it's best to first create a virtual disk for the machine to use. Note that you can create a virtual disk at the same time that you create a virtual machine. However, creating the virtual disk first gives you more flexibility. So, I recommend you create virtual disks and virtual machines separately. Here are the steps to create a virtual disk:

1. **In Hyper-V Manager, click New and then choose Hard Disk.**

   This brings up the New Virtual Hard Disk Wizard, as shown in Figure 1-5.

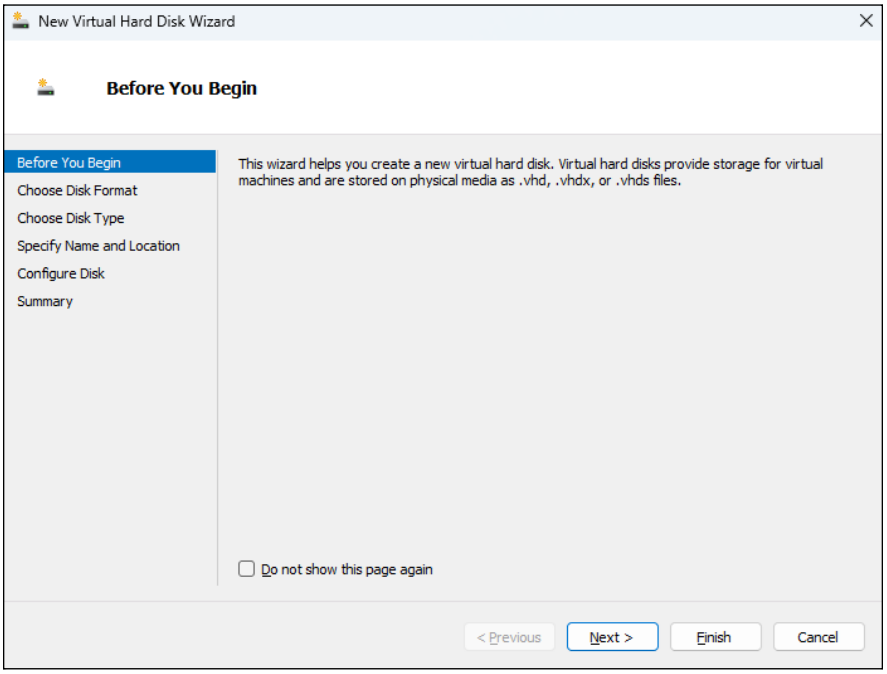


FIGURE 1-5: The New Virtual Hard Disk Wizard.

2. **Click Next.**

   You're asked which disk format to use, as shown in Figure 1-6.

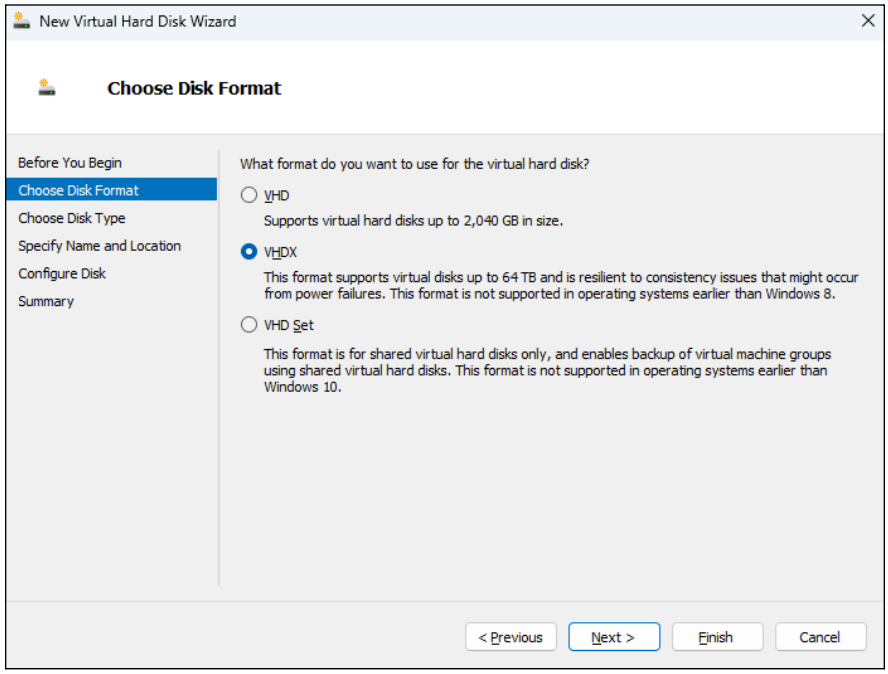


FIGURE 1-6: Choose your disk format.

3. **Select VHDX, and then click Next.**

   I recommend you always use the VHDX format, which can support drives as large as 64TB.

   When you click Next, the Choose Disk Type option page is displayed, as shown in Figure 1-7.

4. **Select the disk type you want to use.**

   The options are Fixed Size, Dynamically Expanding, or Differencing. Choose Fixed Size if you're concerned about the performance of the disk; otherwise, choose Dynamically Expanding.

5. **Click Next.**

   The page shown in Figure 1-8 appears.

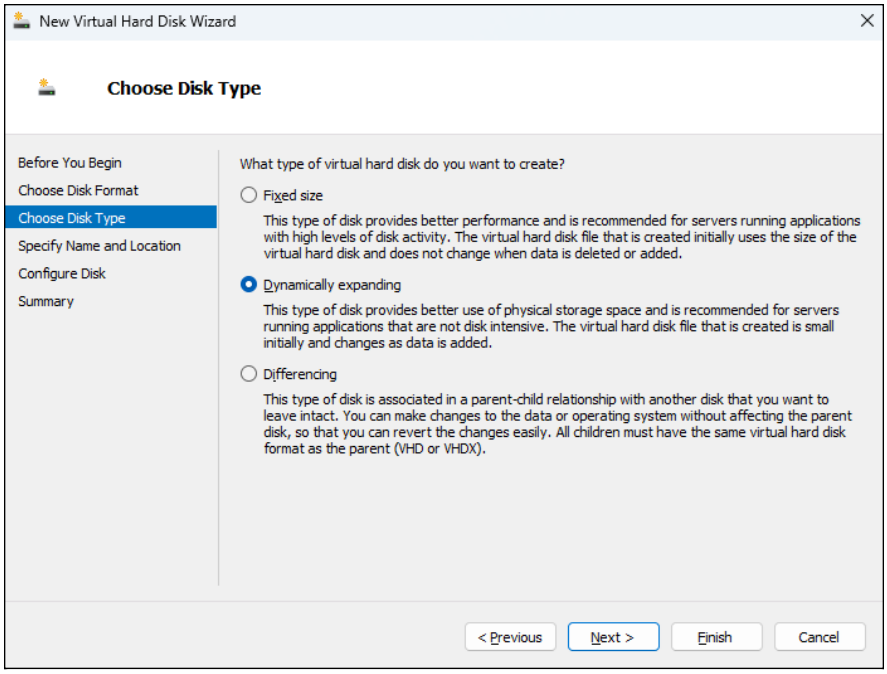


FIGURE 1-7: Choose your disk type.

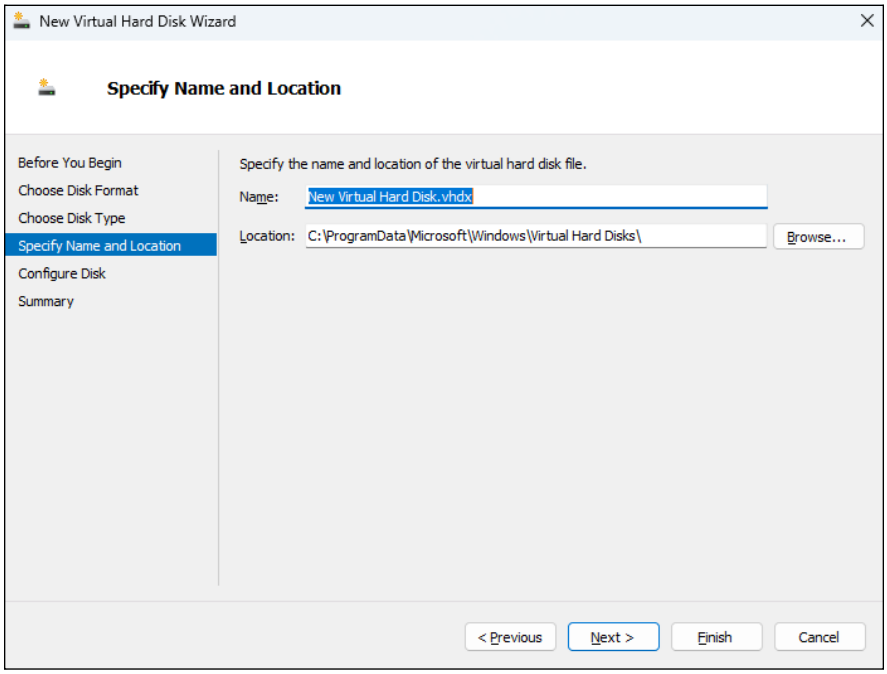


FIGURE 1-8: Specify the name and location of the disk.

6. **Specify the name and location of the new disk.**

Type any name you want for the virtual disk drive. Then, click the Browse button to browse to the disk location where you want Hyper-V to create the `.vhdx` file.



TIP

Make sure you choose a location that has enough disk space to create the `.vhdx` file. If you're creating a dynamically expanding disk, you should ensure that the location has enough space to accommodate the drive as it grows.

7. **Click Next.**

The Configure Disk page appears, as shown in Figure 1-9.

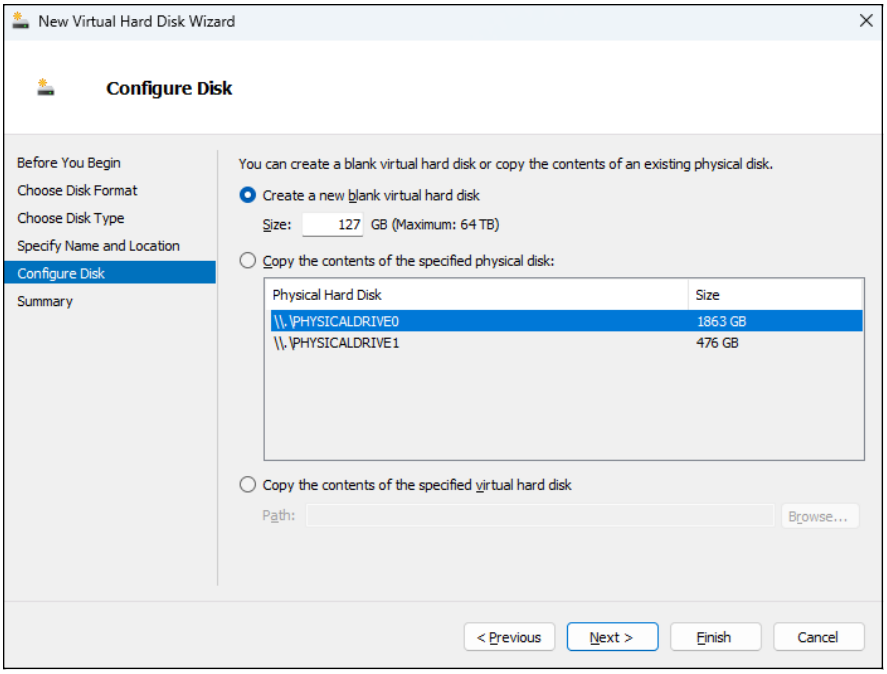


FIGURE 1-9: Specify the size of the disk.

8. **Specify the maximum size for the disk drive.**



TIP

This page also allows you to copy data either from an existing physical disk drive or from an existing virtual disk drive. Copying data from an existing physical drive is a quick way to convert a physical computer to a virtual computer; just copy the physical disk to a virtual disk, and then use the new virtual disk as the basis for a new virtual machine.

9. **Click Next.**

   A confirmation screen appears, summarizing the options you've selected for your new disk.

10. **Click Finish.**

    The new disk is created. Note that if you selected Fixed Disk as the disk type, creating the disk can take a while because the entire amount of disk storage you specified is allocated to the disk. Be patient.

You're done! You've now created a virtual disk that can be used as the basis for a new virtual machine.

Creating a Virtual Machine

After you've created a virtual disk, creating a virtual machine to use it is a straightforward affair. Follow these steps:

1. **From Hyper-V Manager, choose New and then choose Virtual Machine.**

   This brings up the New Virtual Machine Wizard, as shown in Figure 1-10.

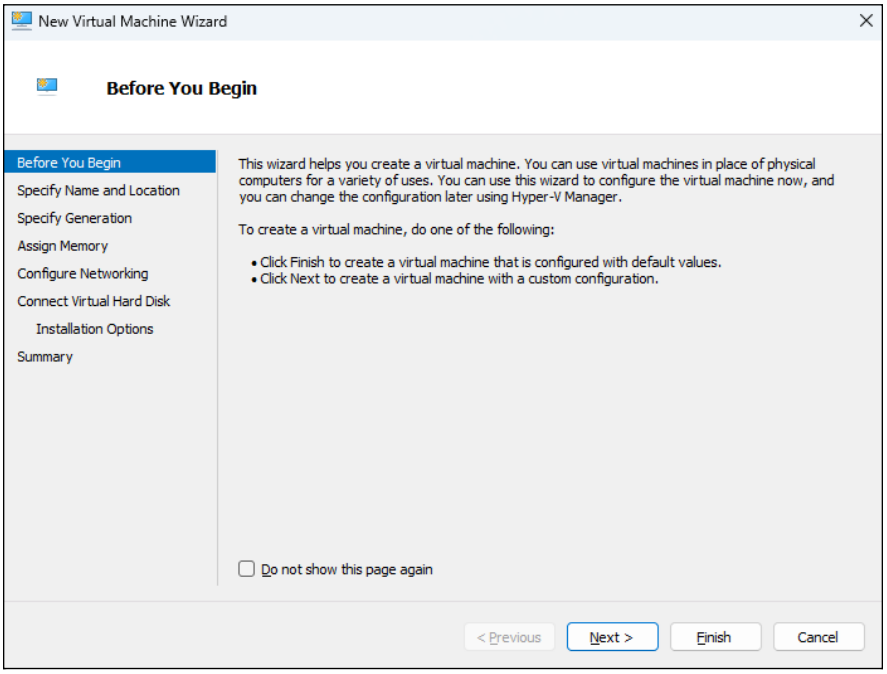


FIGURE 1-10: Say hello to the New Virtual Machine Wizard.

2. **Click Next.**

   The Specify Name and Location page appears, as shown in Figure 1-11.

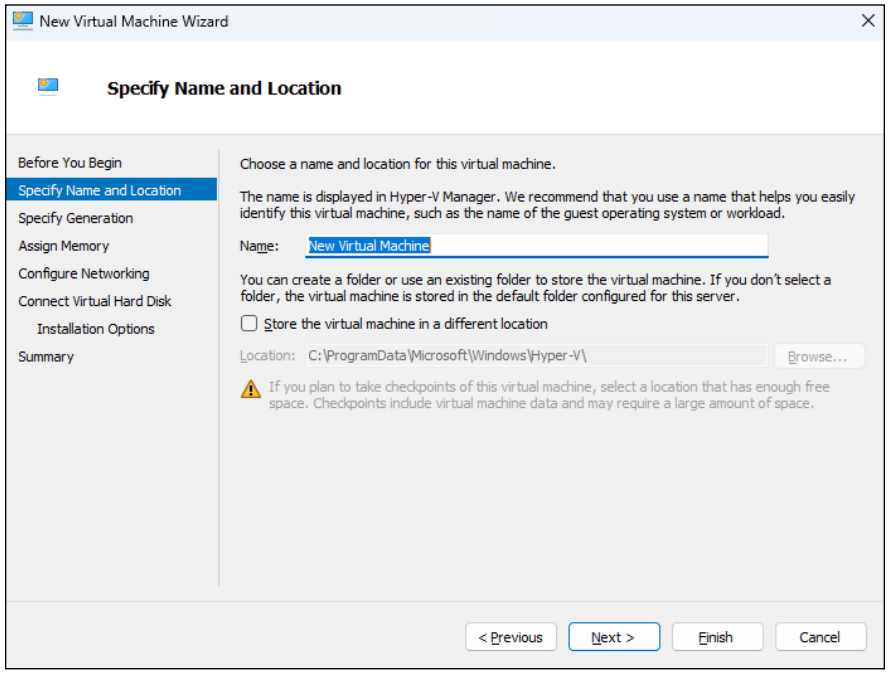


FIGURE 1-11: Specify the name and location of the virtual machine.

3. **Enter the name you want to use for your virtual machine.**

   You can choose any name you want here.

4. **Specify the location of the virtual machine's configuration file.**

   Every virtual machine has an XML file associated with it that defines the configuration of the virtual machine. You can allow this file to be stored in the default location, or you can override the default and specify a custom location.

5. **Click Next.**

   The Specify Generation page appears, as shown in Figure 1-12.

6. **Specify the generation you want to use for the new virtual machine.**

   In most cases, you should opt for Generation 2, which uses newer technology than Generation 1 machines. Use Generation 1 only if the guest operating system will be earlier than Windows Server 2012 or Windows 8.

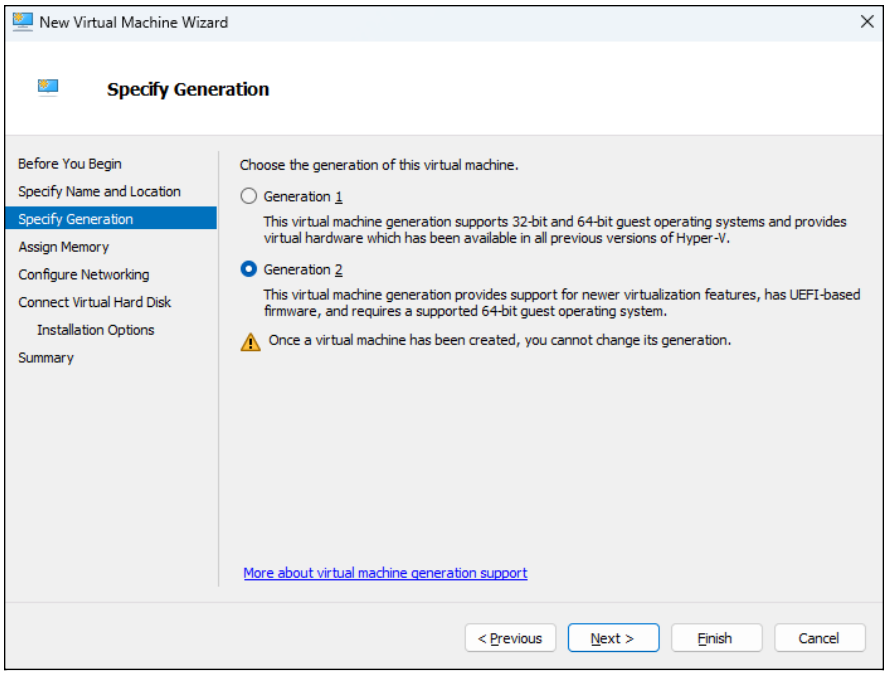


FIGURE 1-12: Specify the generation of the new virtual machine.

7. **Click Next.**

   The Assign Memory page appears, as shown in Figure 1-13.

8. **Indicate the amount of RAM you want to allocate for the new machine.**

   The default is 4096MB (which is equivalent to 4GB), but you may want to increase that.

   I also recommend that you click the Use Dynamic Memory for This Virtual Machine option, which improves memory performance.

9. **Click Next.**

   The Configure Networking page appears, as shown in Figure 1-14.

10. **Select the virtual switch you want to use for the virtual machine.**

    This is the point where you realize why you needed to create a virtual switch before you start creating virtual machines. Use the Connection drop-down list to select the virtual switch you want to connect to this VM.

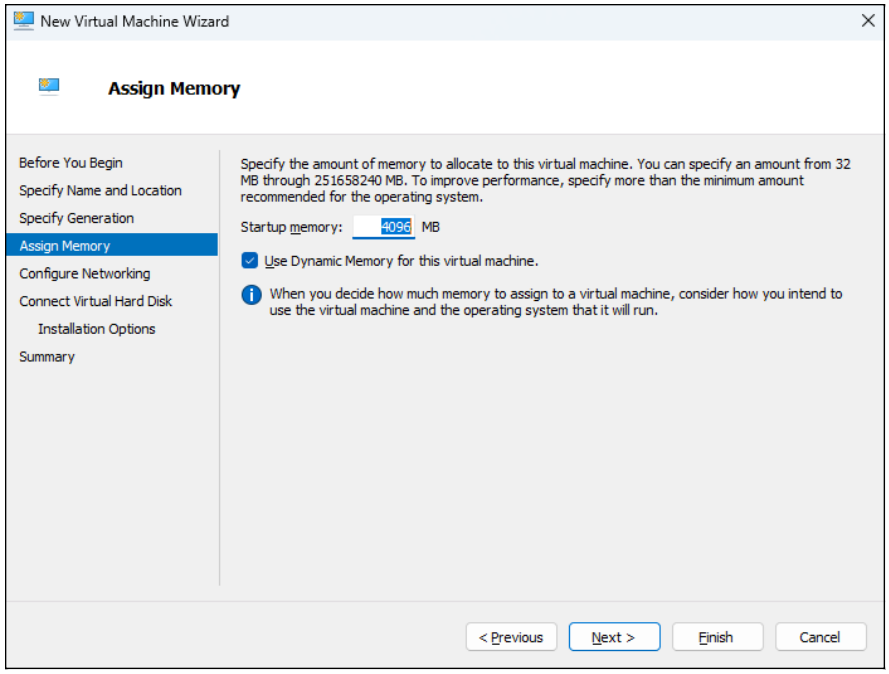


FIGURE 1-13: Specify the memory for the new virtual machine.

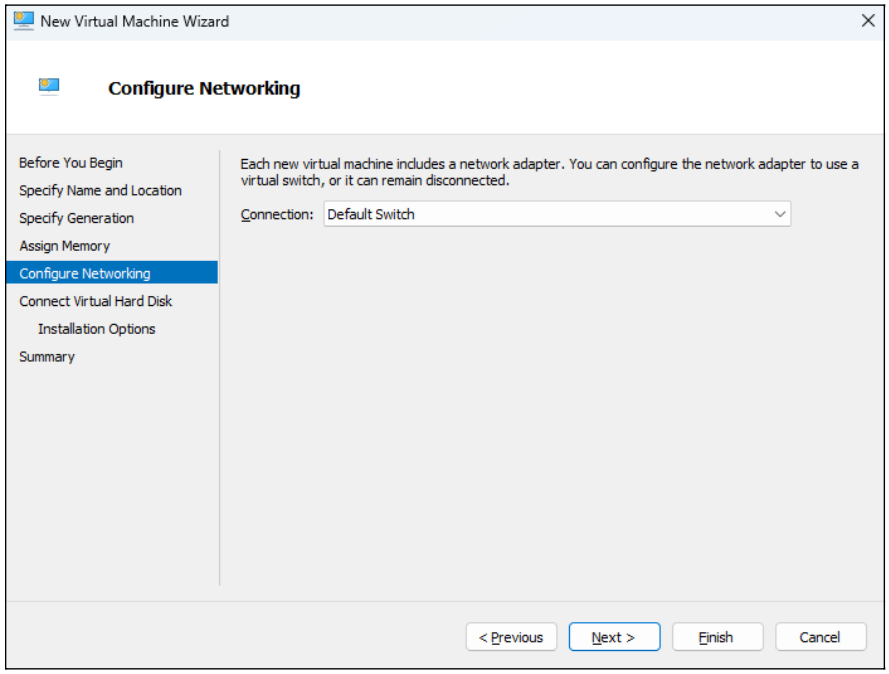


FIGURE 1-14: Configure the networking for the new virtual machine.

11. **Click Next.**

The Connect Virtual Hard Disk page appears, as shown in Figure 1-15.

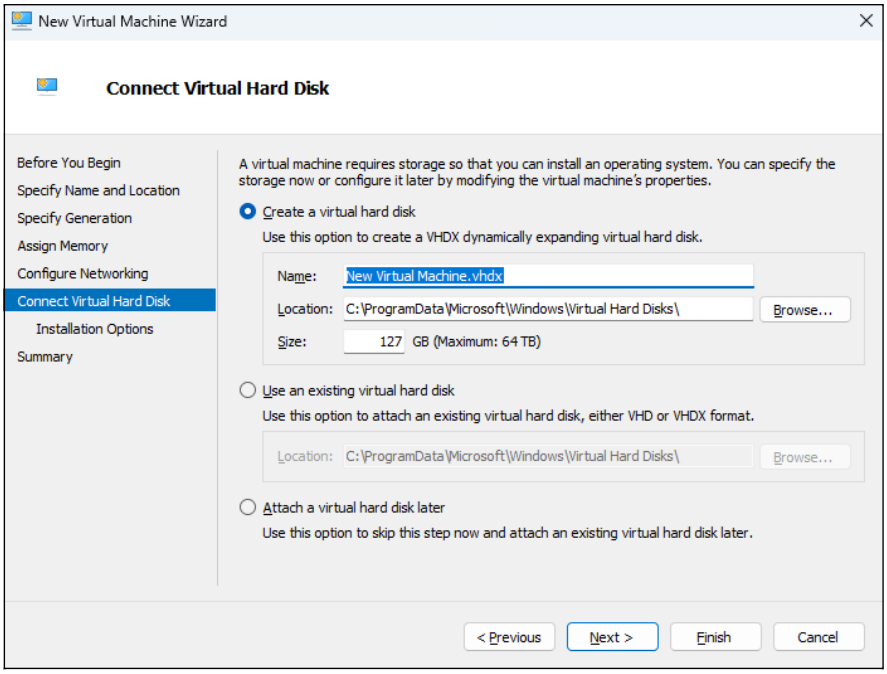


FIGURE 1-15: Connecting a virtual disk.

Assuming you've already created a virtual disk for the virtual machine, choose the Use an Existing Virtual Hard Disk option, click Browse, and locate and select the virtual disk.

If you haven't already created a virtual disk, you can use the Create a Virtual Hard Disk option and create one now.

12. **Click Next.**

The Installation Options page appears, as shown in Figure 1-16. You can specify an operating system installation image here, but I'll skip that for now and install the operating system later (see the next section, "Installing an Operating System").

13. **Click Next.**

A summary page is displayed indicating the selections you've made.

14. **Click Finish.**

The virtual machine is created.

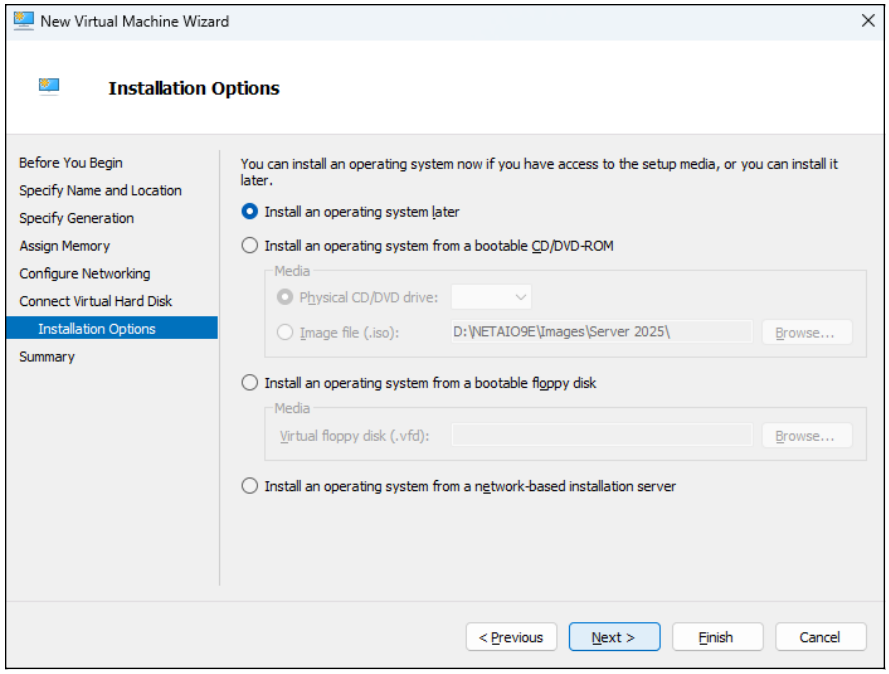


FIGURE 1-16: The Installation Options page.

Installing an Operating System

After you've created a virtual machine, the next step is to configure it to install an operating system. First, you'll need to get the installation media in the form of an `.iso` file (an `.iso` file is a disk image of a CD or DVD drive). After you have the `.iso` file in place, follow these steps:

1. **From the Hyper-V Manager, choose the new virtual machine and click Settings.**

   The Settings dialog box appears, as shown in Figure 1-17.

2. **Click SCSI Controller in the Hardware list. Then select DVD Drive, and click Add.**

   The configuration page shown in Figure 1-18 appears.

3. **Click the Image File option, click Browse, and select the `.iso` file that contains the operating system's installation program.**

4. **Click OK.**

   You're returned to the Hyper-V Manager screen.

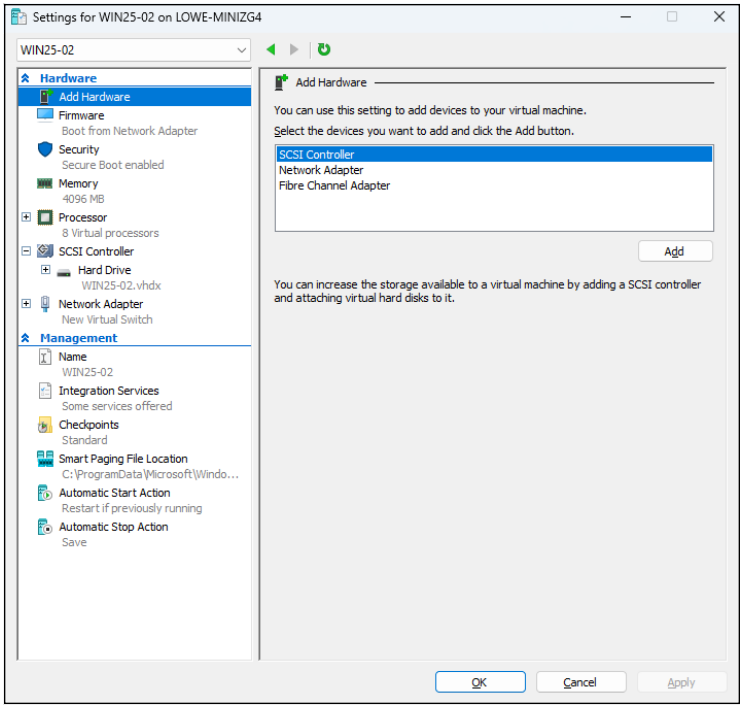


FIGURE 1-17: Editing the settings for a virtual machine.

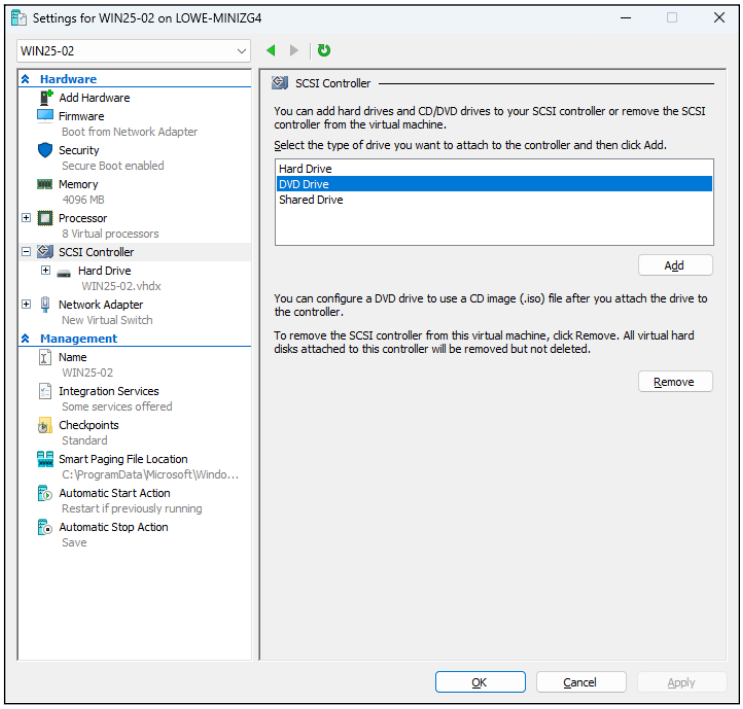


FIGURE 1-18: Configuring a DVD drive.

5. **With the new virtual machine still selected, click Connect.**

   A console window opens, showing that the virtual machine is currently turned off (see Figure 1-19).

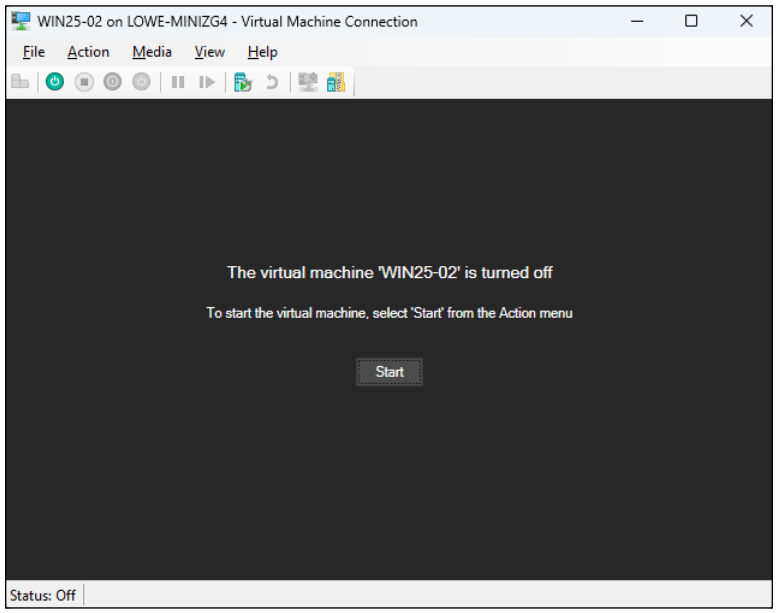


FIGURE 1-19: Viewing a virtual machine through a console.

6. **Click Connect.**
7. **Click Start.**

   The virtual machine powers up.

8. **When prompted to press a key to boot from the CD or DVD, press any key.**

   The operating system's installation program starts.

9. **Follow the instructions of the installation program to install the operating system.**

   That's all there is to it. You've now successfully created a Hyper-V virtual machine!

IN THIS CHAPTER

- Installing VMware Workstation Pro
- Creating and using virtual machines
- Installing VMware Tools

Chapter 2
VMware

Virtualization is a complex subject, and mastering the ins and outs of working with a full-fledged virtualization system like VMware Infrastructure is a topic that's beyond the scope of this book. You can dip your toes into the shallow end of the virtualization pond, however, by downloading and experimenting with VMware's virtualization product, called VMware Workstation Pro, which is free for personal use. You can download this software from `www.vmware.com`.


TECHNICAL STUFF

VMware was recently acquired by Broadcom, Inc., so VMWare is now officially known as VMware by Broadcom. I just refer to it as VMware throughout this chapter, though, because that's what everybody except for VMware by Broadcom's legal team calls it.

This chapter is a brief introduction to VMware's virtualization platform and creating and using virtual machines (VMs) with VMware Workstation Pro.

Looking at vSphere

vSphere is an umbrella term for VMware's virtualization platform. The term *vSphere* encompasses several distinct products and technologies that work together to provide a complete infrastructure for virtualization. These products and technologies include the following:

- **ESXi:** ESXi is the core of vSphere; it is a Type-1 hypervisor that runs on host computers to manage the execution of VMs, allocating resources to the VMs as needed. ESXi comes in two basic flavors:
  - *Installable:* The Installable version of software can be installed onto the hard drive on a host computer, much as any other operating system can be installed.
  - *Embedded:* The Embedded version runs as firmware that is actually built into the host computer. It's preinstalled into read-only memory by the manufacturer of the host computer.
- **vCenter Server:** vCenter Server is a server application that runs on Windows Server installed in a VM. vCenter is the central point for creating new VMs, starting and stopping VMs, and performing other management tasks in a vSphere environment.
- **vCenter Client:** vCenter Client is a Windows application that you use to access the features of a vCenter Server remotely. vCenter Client is the tool you'll work with most when you manage a vSphere environment.
- **VMFS:** VMFS, which stands for *Virtual Machine File System,* is the file system used by vSphere to manage disk resources that are made available to VMs. With VMFS, you can create *data stores* to access physical disk devices, and you can then create *volumes* on these data stores to make disk storage available to VMs.

Getting Started with VMware Workstation Pro

VMware Workstation Pro is a simplified version of the vSphere environment that provides many of vSphere's features but does not utilize the Type-1 ESXi hypervisor. Instead, VMware Workstation Pro is a Type-2 hypervisor that runs within a Windows environment. VMware Workstation Pro is useful for learning about virtualization and for simple desktop virtualization and is handy for creating VMs that run on a workstation machine. But it's not designed to run servers in a production

environment. I'm covering it here because experimenting with VMware Workstation Pro is a great way to learn some of the basic concepts of virtualization.

Figure 2-1 shows VMware Workstation Pro's main screen. From this screen, you can create a new VM or run one of the VMs you've already created. As you can see in the figure, I've created several VMs: two that run Windows Server 2019 and a Linux machine.

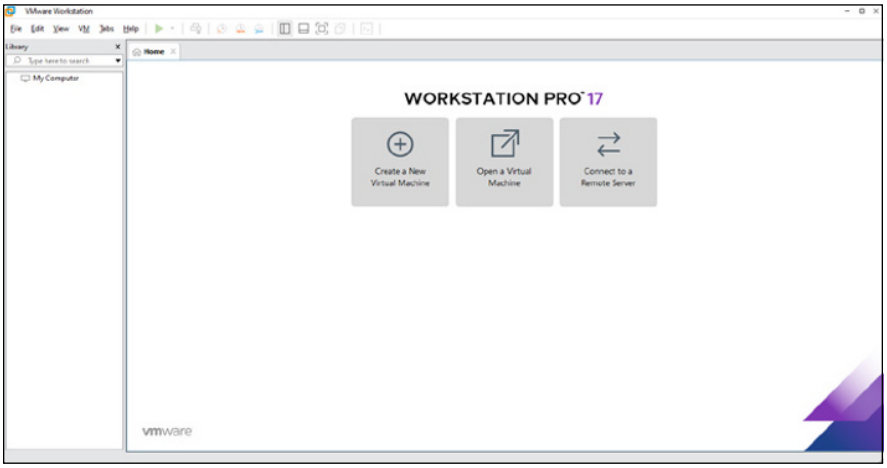


FIGURE 2-1: VMware Workstation Pro lets you experiment with virtualization.

You can run an existing VM by selecting the VM and clicking Play Virtual Machine. This launches the VM, which opens in a new window, as shown in Figure 2-2. When you launch a VM, the VM behaves exactly as a real computer would when you power it up: First, it initializes its virtual hardware devices; then it loads the guest operating system that has been installed in the VM.

In Figure 2-2, Windows Server 2019 has booted up and is waiting for you to press Ctrl+Alt+Del to log on.

The prompt to press Ctrl+Alt+Del shown in Figure 2-2 illustrates one of the peculiar details of running a VM within a host operating system. When you press Ctrl+Alt+Del, which operating system — the host or the guest — responds? The answer is that the host operating system responds to the Ctrl+Alt+Del before the guest operating system ever sees it.

To get around this limitation, VMware uses the special keyboard shortcut Ctrl+Alt+End to send a Ctrl+Alt+Del to the guest operating system. Alternatively, you can use the VM drop-down menu that appears in the menu bar above the VM menu. This menu lists several actions that can be applied to the VM, including Send Ctrl+Alt+Del.

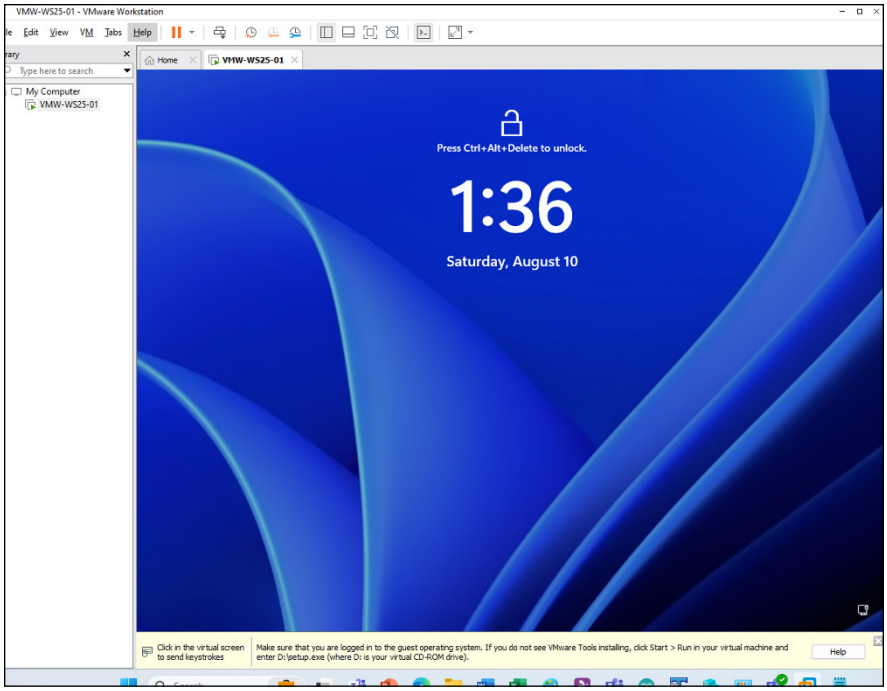


FIGURE 2-2: A VM running Windows Server 2019.

Another detail you should know about when working with a VM is that when you click in the VM's window, the VM captures your mouse and keyboard, so your input will be directed to the VM rather than the host computer. If you want to break the bonds of the VM and return to the host computer, press Ctrl+Alt.

Creating a Virtual Machine

Creating a new VM in VMware Workstation Pro is relatively easy. In fact, the most challenging part is getting hold of the installation disc for the operating system you want to install on the VM. Remember that a VM is useless without a guest operating system, so you need to have the installation disc available before you create the VM.

If you just want to experiment with virtualization and don't have extra licenses of a Windows Server operating system, you can always download an evaluation copy of Windows Server 2025 from `www.microsoft.com`. The evaluation period is six months, so you'll have plenty of time to experiment.

The downloadable trial version of Windows Server 2025 comes in the form of an `.iso` file, which is an image of a DVD file that you can mount within your VM as though it were a real disc.

When you have your `.iso` file or installation disc ready to go, you can create a new VM by following these steps:

1. **Click Create a New Virtual Machine on the VMware Workstation Pro home screen (refer to Figure 2-1).**

   This brings up the New Virtual Machine Wizard, shown in Figure 2-3.



FIGURE 2-3: The first page of the New Virtual Machine Wizard.

2. **Choose the installation option you want to use.**

   You have two choices:

   - *Typical:* Select this option to create a VM using the most common options. This is the default choice and the one I demonstrate here.
   - *Custom:* Select this option to specify advanced options for your VM.

3. **Click Next.**

   The screen in Figure 2-4 appears, which lets you select the operating system installation image.

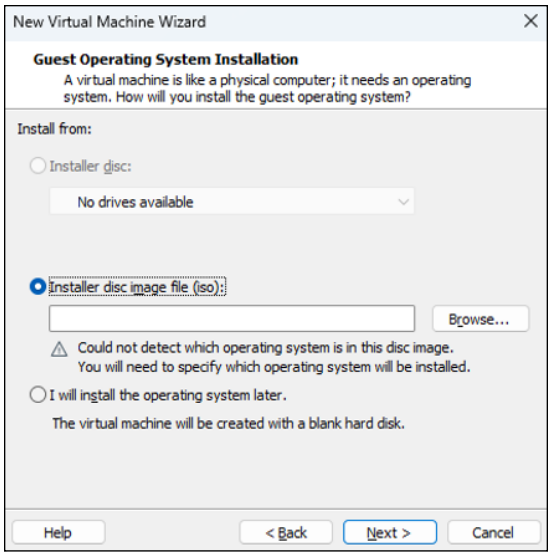


FIGURE 2-4: Selecting the operating system image file to install for the new VM.

4. **Choose your installation mode and select the installation image.**

   You have three choices:

   - *Installer Disc:* Select this option and then choose from the drop-down list the drive you'll install from if you want to install from an actual CD or DVD.
   - *Install Disc Image File (iso):* Select this option, click the Browse button, and browse to the `.iso` file that contains the installation image.
   - *I Will Install the Operating System Later:* Select this option if you want to create the VM now but install the operating system later.

   For this example, select Installer Disk Image File, and then click the Browse button and select the `.iso` file for Windows Server 2025.

5. **Click Next.**

   The Select Guest Operating System screen, shown in Figure 2-5, appears.

6. **Select the version of Windows you want to install, and then click Next.**

   The Name the Virtual Machine screen, shown in Figure 2-6, appears. Although this screen leads you to believe you'll have to install Windows Server Standard Core — the version without the Windows desktop graphical user interface (GUI) — you can select to install the GUI later.

7. **Enter a name for the VM and the location of the VM's disk, and then click Next.**

   The name you enter here is the name by which the VM will be known within VMware Workstation Pro. When you click Next, the Specify Disk Capacity screen, shown in Figure 2-7, appears.

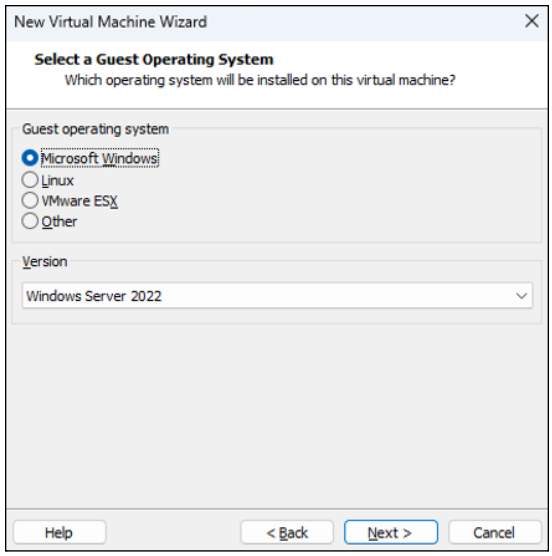


FIGURE 2-5: Selecting the Guest Operating System version.

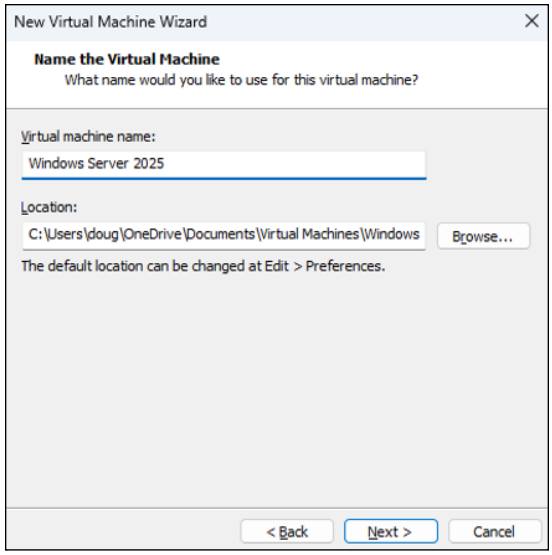


FIGURE 2-6: Naming the VM.

8. Set the size of the disk you want to create, and then click Next.

The default size is 60GB. I usually increase that to 120GB. When you click next, the configuration options you've chosen will be displayed, as shown in Figure 2-8.

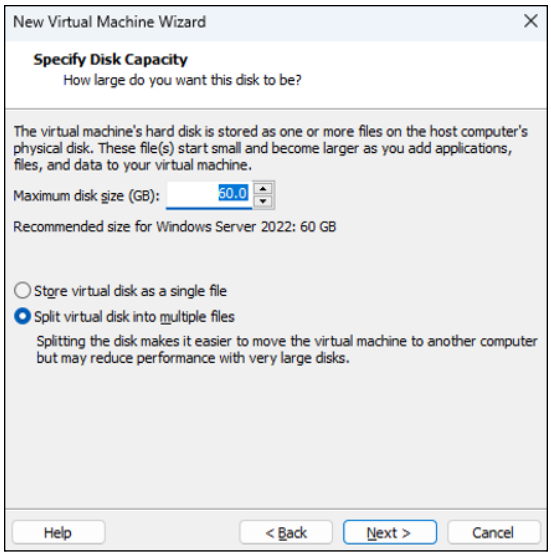


FIGURE 2-7: Specifying the VM disk size.

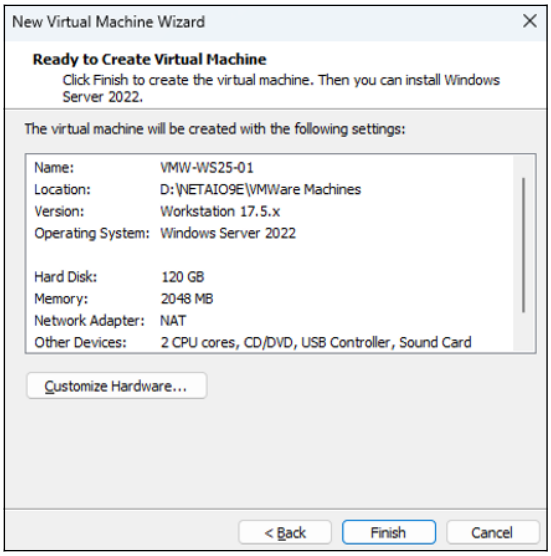


FIGURE 2-8: VMware is ready to create the VM.

9. **Click Finish.**

   The wizard creates the VM and then displays it, as shown in Figure 2-9.

10. **Click Power On This Virtual Machine to start the machine.**

    The VM starts and loads the Windows installer from the DVD.

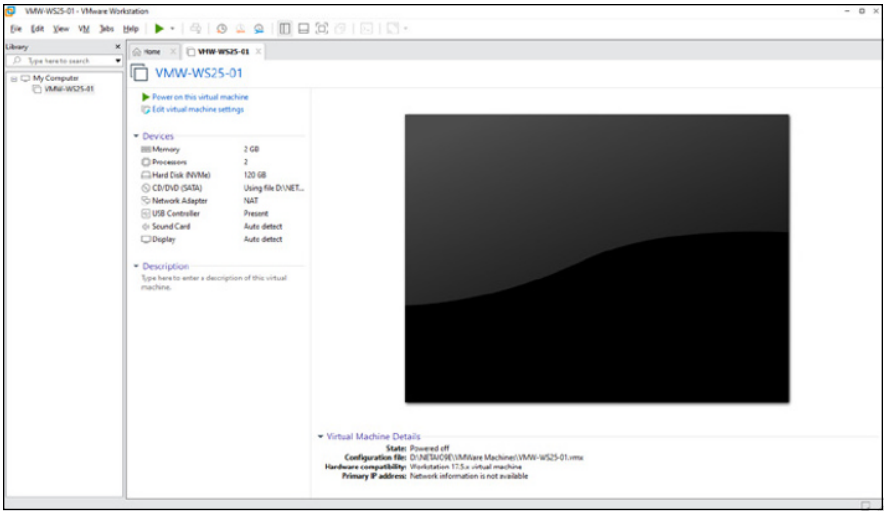


FIGURE 2-9: The new VM is ready to start.

11. **Press any key within five seconds to start the Windows installer.**

    You may need to click within the VM display console to direct focus to the VM, and then press a key.

    If you don't successfully press a key within the allotted time, just power down the VM, and then power it back on and try again.



TIP

12. **Follow the steps to install the operating system.**

    Installing an operating system on a VM is exactly the same as installing it on a physical computer, except that the installation screens appear within a VM window, as shown in Figure 2-10.

    When the operating system is installed, you're done! You can start using the VM.

You can adjust the hardware configuration of a VM by clicking Edit Virtual Machine Settings while the computer is powered off. This action brings up the Virtual Machine Settings dialog box, shown in Figure 2-11. From here, you can adjust the VM's hardware configuration, including the amount of RAM available to the VM and the number of processor cores. You can also adjust the disk drive size; add CD, DVD, or floppy drives; and configure network adapters, USB connections, and sound and display settings.

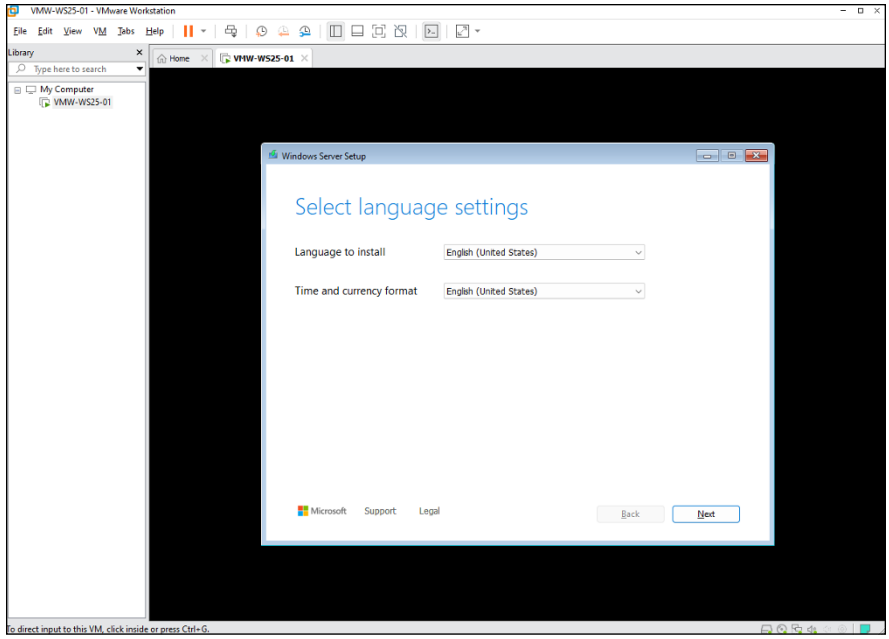


FIGURE 2-10: Installing Windows Server 2025 on a VM.

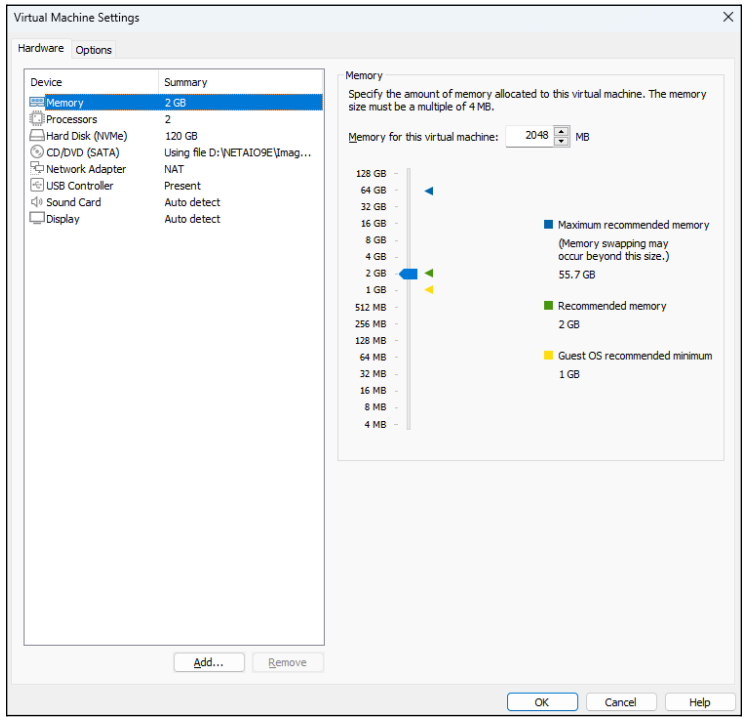


FIGURE 2-11: Configuring VM settings.

Installing VMware Tools

When you've installed an operating system into a VMware VM, you should install an important application called *VMware Tools* before you do anything else. VMware Tools provides a variety of important functions for a VMware VM:

- Significantly improved graphics performance
- Shared folders, which lets you share folders between the VM and the host, making it easy to exchange files between the two
- A shared clipboard, which lets you copy and paste between the VM and the host
- Synchronized time between the guest and host
- Better control of the mouse between guest and host

To install VMware Tools, follow this procedure:

1. **Start the VM.**
2. **On the menu of the VMware console window, choose VM➪Install VMware Tools.**

   The installation image for VMware tools is mounted on the VM's DVD drive. Figure 2-12 shows the server's This PC window with the VMware Tools mounted and ready to install.
3. **Double-click the DVD drive with the VMware Tools installer mounted.**

   This starts the installation of the VMware Tools, as shown in Figure 2-13.

   
   TIP

   If the Setup program doesn't automatically start, press Windows+R, type **d:\setup.exe**, and press Enter. This will manually start the Setup program.
4. **Follow the instructions in the Setup program to install the VMware Tools.**

   When the tools are installed, you'll be prompted to restart the VM.

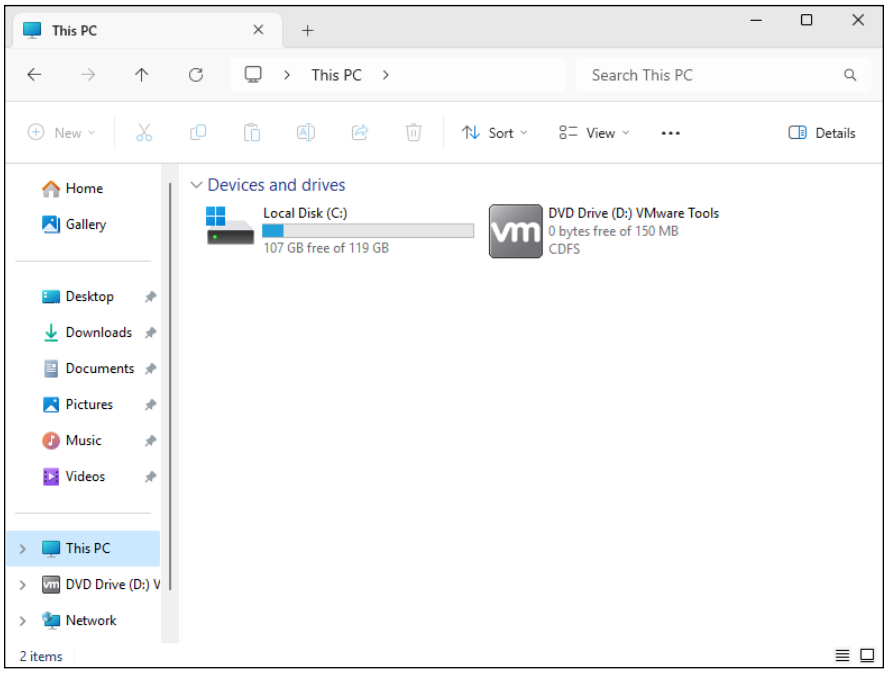


FIGURE 2-12: Installing VMware Tools.

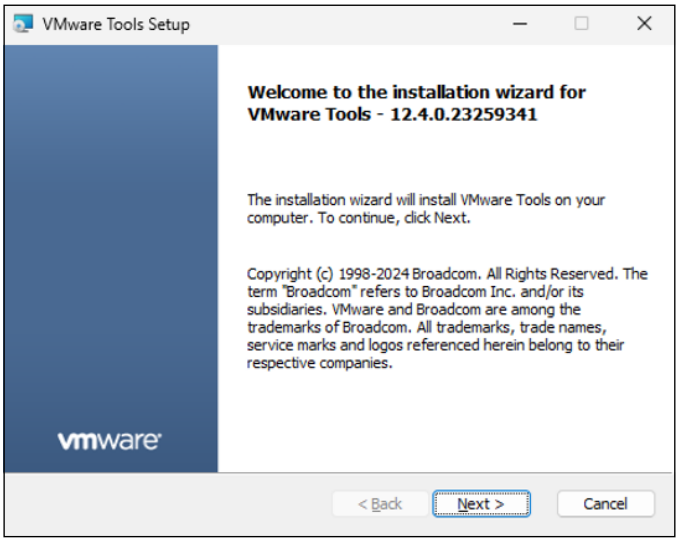


FIGURE 2-13: The VMware Tools Setup program.

IN THIS CHAPTER

- Getting acquainted with Azure
- Signing up for an Azure account
- Creating a virtual machine in Azure
- Managing an Azure virtual machine
- Connecting to an Azure virtual machine

Chapter 3
Azure

Microsoft Azure is a cloud computing service that provides a ton of alternatives for virtualizing your IT infrastructure in the cloud. Azure has been around since 2010 and is now one of the most robust platforms for building IT in the cloud.

Azure provides three basic types of cloud services:

- **Infrastructure as a Service (IaaS):** Cloud-hosted versions of the most basic types of virtual infrastructure components, including virtual machines (VMs), virtual disk storage, and virtual networks. With IaaS, Microsoft provides the physical hardware needed to run basic Hyper-V components.
- **Platform as a Service (PaaS):** A step up from IaaS, PaaS provides cloud-based platforms on which you can run applications. For example, you can deploy web-based ASP.NET applications to Azure and let Microsoft worry about managing the infrastructure needed to run the application.
- **Software as a Service (SaaS):** Fully functional software applications that you simply subscribe to.

In this chapter, I provide a brief overview of the most important services provided by Azure. Then I cover the nuts and bolts of setting up an Azure account and creating a VM in Azure's cloud.



In case you're interested, the color *azure* is technically midway between blue and cyan. It's generally thought of as the color of the sky on a beautiful clear day. It gets its name from the mineral lapis lazuli, which has a striking blue color. The first use of the word *azure* in English came from none other than Geoffrey Chaucer in his 1374 poem *Troilus and Criseyde*.

Looking at Azure Services

At the time of this writing, Microsoft lists hundreds of distinct services provided by Azure. Here are the most popular and useful of these services:

- **VMs:** Allow you to create and run fully functional VMs on Azure's cloud platform. The VMs can run Microsoft Windows Server operating systems or Linux. For Windows Server machines, the operating system license cost is included in the price of the VM.

  When you provision an Azure VM, you select the size of the machine in terms of the number of processors, RAM, and disk space. You can easily adjust these settings later if your needs change.

  You can also select a preconfigured disk image to use for the machine. This means that you don't have to go through the steps of installing an operating system; instead, you just choose an image that has the correct operating system already installed. You can also download server images from the Azure marketplace that are preconfigured by third-party vendors.

  Pricing is based on actual usage at rates that vary depending on the class of service provided. The pricing structure is a bit complicated, so you'll want to explore it in depth to get an accurate understanding of what your VMs will actually cost. Fortunately, Microsoft has a nice estimating tool that can help you eliminate surprises.
- **Storage:** Allows you to create various types of storage resources, including simple disks, file systems, network-attached storage (NAS), and backup storage.
- **Networking:** Lets you create virtual networks that enable your Azure cloud components to communicate with one another and also with your physical network.
- **Web applications:** Azure can host complete web apps that can be developed using many popular programming platforms, including ASP.NET, Java, PHP, Node.js, Python, and HTML5. When you deploy an Azure Web App, Azure takes care of the details of managing the web server VM, so you can focus on the application itself.

- **Mobile applications:** Enables you to create and deploy native or cross-platform mobile apps for iOS and Android. If you're looking for a way to create custom mobile applications for your company, the Azure Mobile App service is a good place to start.
- **SQL database:** Allows you to create a SQL Server database without setting up a dedicated VM to run SQL Server.

Creating an Azure Account

Before you can begin using Azure, you must set up an Azure account. Microsoft offers a free account that's designed to let you familiarize yourself with Azure's capabilities and learn how it works. Setting up the free account is easy, but you have to provide your credit card information. Don't worry — Microsoft won't actually charge your card until you upgrade to paid features.

To set up your account, go to `http://azure.microsoft.com`, click the Get Started with Azure link, click Try Azure for Free, and follow the on-screen instructions.

Here's a general overview of what's available with a free account:

- 750 hours of Windows Server General Purpose machines (known as Azure B1S), both Windows and Linux
- 128GB of managed disk space
- 250GB of SQL database
- 15GB of outbound data bandwidth per month (unlimited inbound)
- 10 web or mobile apps
- An assortment of other free services, too detailed to list here



REMEMBER

Note that all these services are free for a period of one year. After the first year, normal charges are incurred.

When you set up your free account, Microsoft also gives you a $200 credit against any services that aren't included with the free services; this credit must be used within the first 30 days.



TIP

The nuances of what services are free under the free account are detailed and, of course, subject to change. So check out the FAQ available during the signup process to be sure you understand what's free and what isn't.

Examining the Azure Portal

After you've created your free Azure account, you can access the Azure portal by browsing to `http://azure.microsoft.com` and clicking the Portal link in the upper right. After you enter your username and password, the portal page will appear, as shown in Figure 3-1.

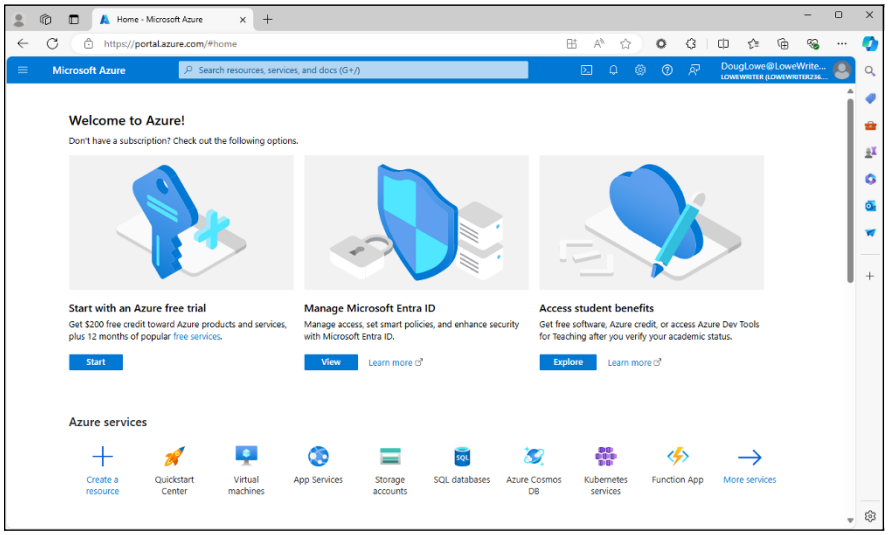


FIGURE 3-1: The Azure Portal.

Across the top of the Azure portal is a toolbar that contains a search box, as well as several useful icons:



- **Shell:** Displays the Azure command shell, which lets you manage Azure using PowerShell or Linux Bash commands.
- **Notifications:** Displays any recent notifications generated by Azure.
- **Settings:** Lets you configure portal settings.
- **Help:** Provides access to useful Azure Help pages.
- **Feedback:** Allows you to send feedback to the Azure team.

Your username appears to the right of these icons. You can hover the mouse over your username to see more specific details about your user account (such as your email address and domain name), or you can click your username to display a drop-down menu that, among other things, lets you sign out or change your password.



At the far left of the toolbar is a button you can click to reveal a menu that leads to additional pages you can display in the portal. For example, + Create a Resource lets you create new resources in Azure, and All Resources displays a list of all the resources you've created in Azure. I recommend you spend some time exploring this menu to familiarize yourself with the range of features that are available in Azure.

Creating a Windows Virtual Machine

After you've had some fun exploring the various pages available within the Azure portal, you're ready to get down to business by creating your first VM. Follow these steps:

1. **Click + Create a Resource.**

   The Azure Marketplace page is displayed, as shown in Figure 3-2. This page lists the range of Azure resources you can create. The menu on the left lists marketplace categories; specific resource types within the selected category are listed on the right.

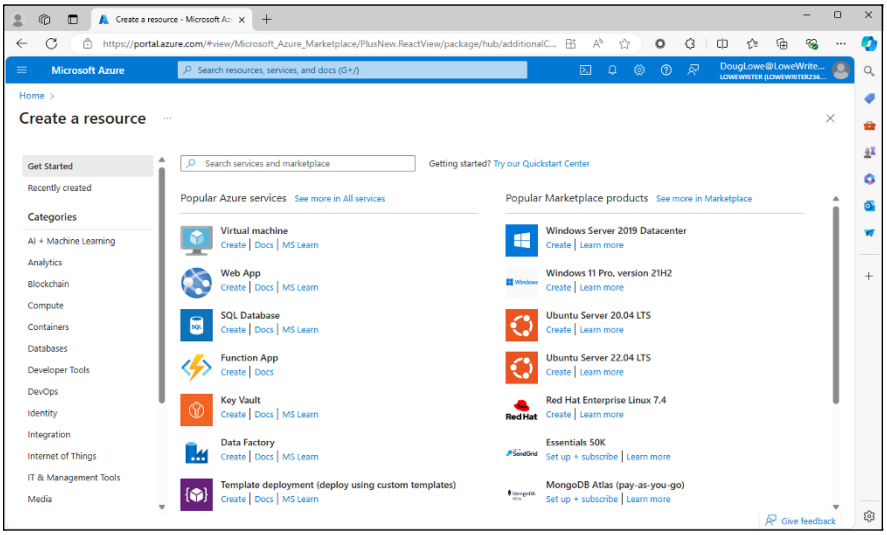


FIGURE 3-2: The Azure Marketplace.

2. **Click Windows Server 2019 Datacenter.**

   You're taken to the Create a Virtual Machine page, shown in Figure 3-3.



TIP

At the time of this writing, only Windows Server 2019 was available from the Create a Resource page. However, as you see in Step 5, you can change the installation image if you want.

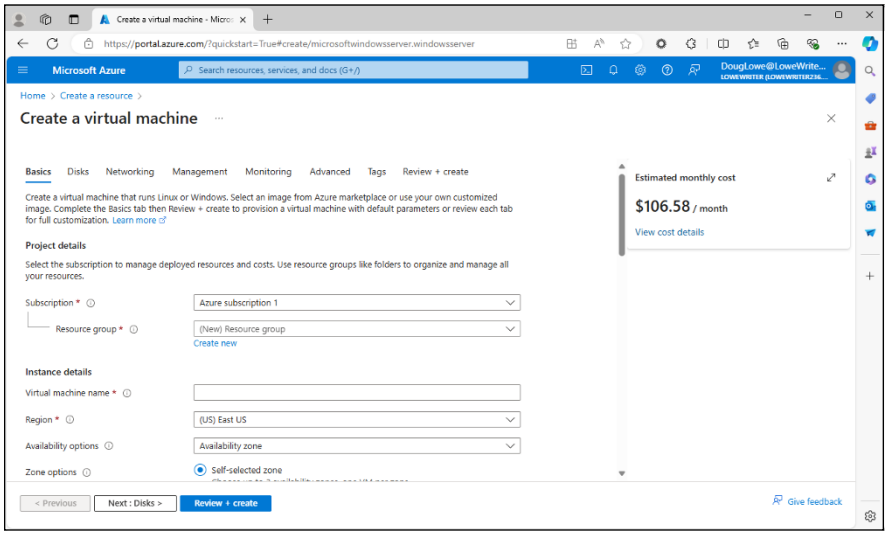


FIGURE 3-3: Creating a virtual machine.

3. Enter the project details.

The project details include the following information:

- *Subscription:* If you have more than one subscription, select it from the drop-down list.
- *Resource Group:* Each subscription can have multiple resource groups, which can simplify management if you have numerous VMs. You need at least one resource group, so click the Create New link and enter a name to use for the resource group; then click OK in the pop-up that appears.

4. Enter the instance details.

The instance details include the following information:

- *Name:* Enter the computer name you want to use for this VM.
- *Region:* The region where you would like your VM to be located. Scroll through the list to see the available options. In most cases, you'll want to choose the location that best describes your own location.
- *Availability Options:* This setting lets you control how much redundancy should be provided for your VM. For our purposes, change this setting to No Infrastructure Redundancy Required.

- *Image:* This setting lets you choose the Windows Server installation image to use. For our purposes, change this setting to Windows Server 2022 DataCenter Azure Edition – X64 Gen2. (Windows Server 2025 wasn't available in Azure at the time I wrote this, but by the time you read this book it should be.)
- *VM Architecture:* For some images, such as various Linux versions, you can choose Arm64 instead of x64. But x64 is the only available option for Windows.
- *Run with Azure Spot discount:* This check box lets you select a cost-saving option that runs your VM only when capacity is available. Your VM may occasionally be interrupted by other workloads. Because we're running under a free trial period, we'll skip this option.
- *Size:* This option lets you select the VM size. For this example, choose Standard.

5. **Enter the administrator account details.**

   These details include the following:

   - *Username:* The username for the administrator account.
   - *Password:* The password for the administrator account. (This password must meet standard Windows complexity requirements.)

6. **Enter the inbound port rules.**

   By default, only Remote Desktop Protocol (RDP) is allowed. You'll need this port open to connect to your VM via Remote Desktop Connection. If you need additional ports, select them from the drop-down.

7. **Select Licensing details.**

   If you already own a Windows Server license you want to apply to this VM, you can select the check box.

8. **Click Next: Disks >.**

   This takes you to the Disks tab, shown in Figure 3-4, where you configure the disks used by the VM. At a minimum, you'll need an operating system disk. You can also create additional disks here if your VM requires them.


TIP

You can select from one of three disk types: Premium SSD, Standard SSD, or Standard HDD. For this example, Premium SSD is the choice. After all, because we're running on a free subscription, why not choose the best?

The Encryption Type drop-down list lets you choose the type of encryption to use for the disk. For simplicity, choose Encryption at Rest with Platform-Managed Key. That way, you won't have to fuss with managing the encryption key.

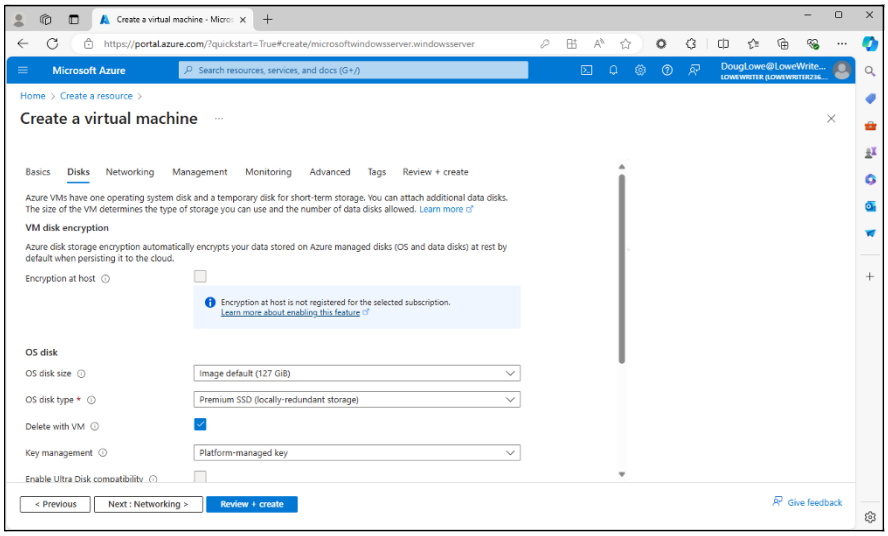


FIGURE 3-4: The Disks tab.

If you want to create additional data disks for your VM, click Create and Attach a New Disk. For this VM, we'll add a 1TB disk for data.

9. **Click Create and Attach a New Disk.**

The Create a New Disk tab, shown in Figure 3-5, appears.

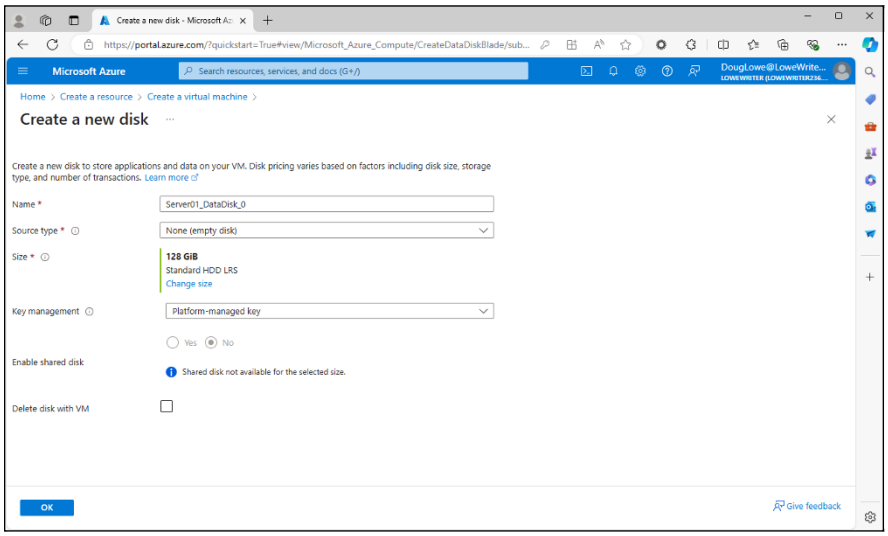


FIGURE 3-5: Creating a new disk.

The following information is required to create a data disk:

- *Name:* The default value supplied is usually appropriate.
- *Source type:* Choose None (Empty Disk).
- *Size:* The default is 1TB of Premium SSD. If you want to change the size or use a less expensive disk type, click Change Size.
- *Encryption type:* The default is Encryption at Rest with Platform-Managed Key, which is the easiest to manage.
- *Enable shared disk*: This advanced feature lets the disk be attached to more than one VM. Unless you have experience with this feature, you should leave it set to No.

Adjust the disk settings as you see fit; then click OK to return to the Disks tab. You can repeat this step if you need additional data disks.

10. **Click Next: Networking>.**

The Networking tab, shown in Figure 3-6, appears.

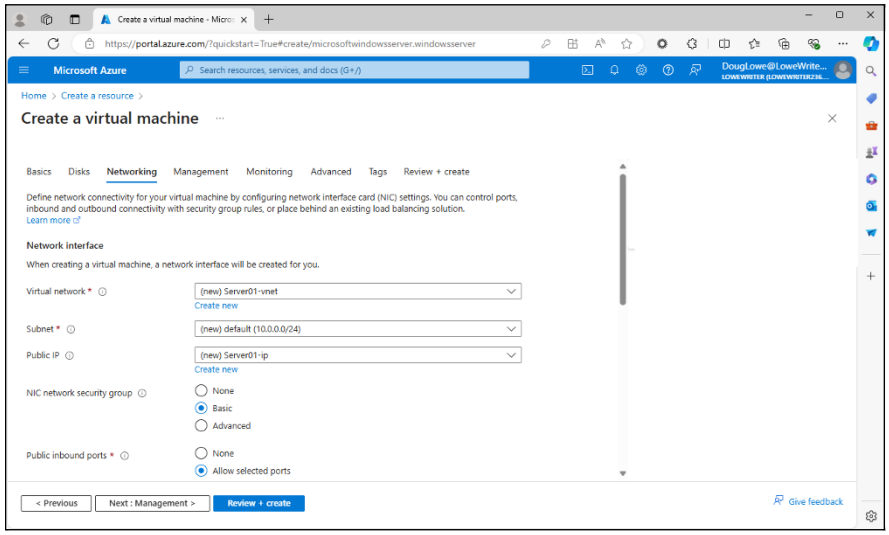


FIGURE 3-6: The Networking tab.

This tab lets you specify various network settings for the VM's network interface. In most cases, you can leave the defaults as they are.

Although the Create a Virtual Machine wizard has several more tabs, their defaults are adequate for demonstration purposes. You can peruse those tabs if you want, but for this example, I'll skip them.

11. Click Review + Create.

The Review + Create tab, shown in Figure 3-7, appears.

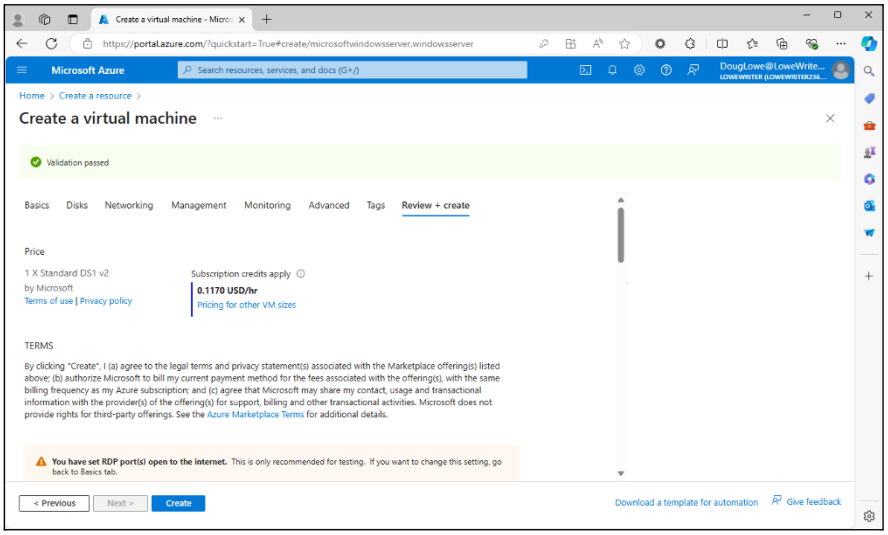


FIGURE 3-7: Ready to create!

Read over the details on this tab to ensure that you've configured the VM correctly and that you understand the subscription terms and pricing. Be sure to verify that "Free Trial" is listed for the Subscription option, unless you're ready to pay for the VM.

12. Click Create.

Azure grinds and whirs for a moment while it creates your VM. You'll see a "Deployment is in progress" notice, shown in Figure 3-8, while Azure is busy deploying your VM.

When the deployment is finished, you'll see the page shown in Figure 3-9. Spend a few minutes reviewing the details on this page, including the options listed under Next Steps. If you want, you can click Go to Resource now and see the details of the VM you just created. But I suggest instead that you proceed to the next section, "Managing an Azure Virtual Machine."

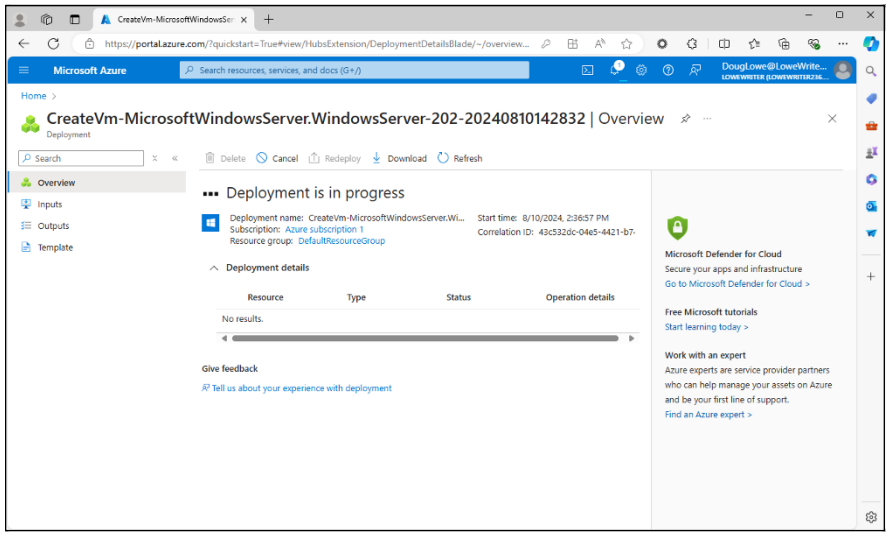


FIGURE 3-8: Azure is deploying the VM.

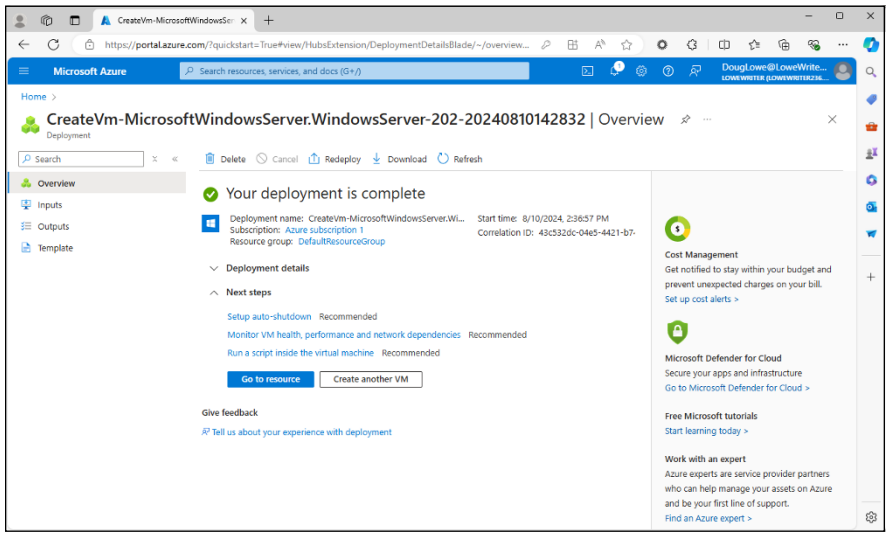


FIGURE 3-9: The deployment is complete!

Managing an Azure Virtual Machine



The Azure Dashboard is your hub for managing your Azure environment, including VMs. To get to the Dashboard, just choose Dashboard from the menu that appears when you click the Show Portal Menu icon at the upper left of any Azure page (shown in the margin). Figure 3-10 shows the Dashboard with the VM created in the previous section.

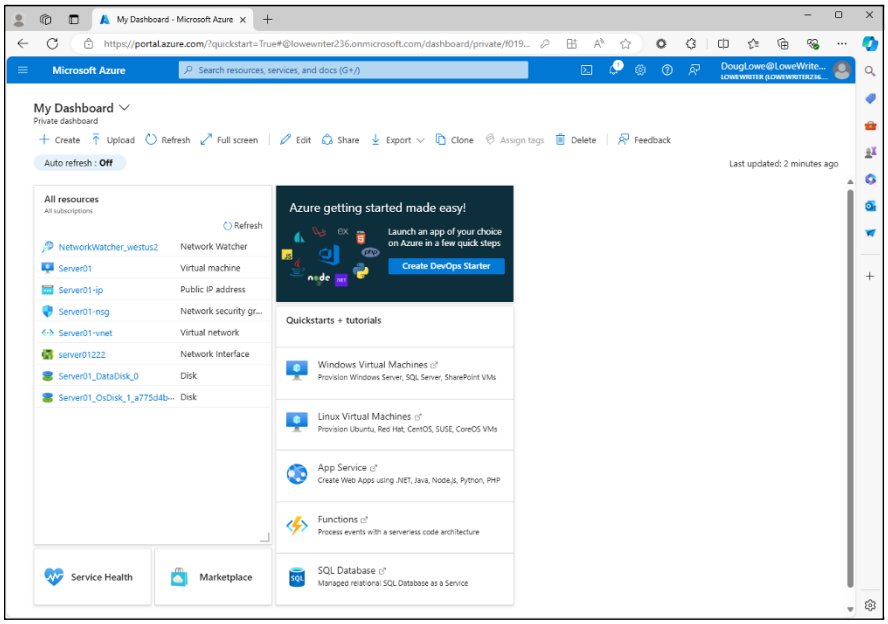


FIGURE 3-10: The Azure Dashboard.

When you create an Azure VM, the machine is automatically pinned to the Dashboard, which means that it appears in a box at the right side of the page (refer to Figure 3-7); the box indicates that the server is running.

In the All resources section, you can see the resources that were created for the VM created in the previous section:

- **NETAIO-01:** The VM that was just created
- **NETAIO-01-ip:** A public IP address that has been allocated to the VM
- **NETAIO-01-nsg:** A network security group that was automatically created to manage the server's security
- **netaio-01362:** The server's network interface
- **NETAIO-01_DataDisk_0:** The data disk that was created for the VM
- **NETAIO-01_OsDisk_1_5d28b. . .:** The OS disk that was created for the VM
- **NETAIO-vnet:** A virtual network that was created and to which the VM was added

You can manage any of these resources from the portal. For example, Figure 3-11 shows the resource page for the NETAIO-01server.

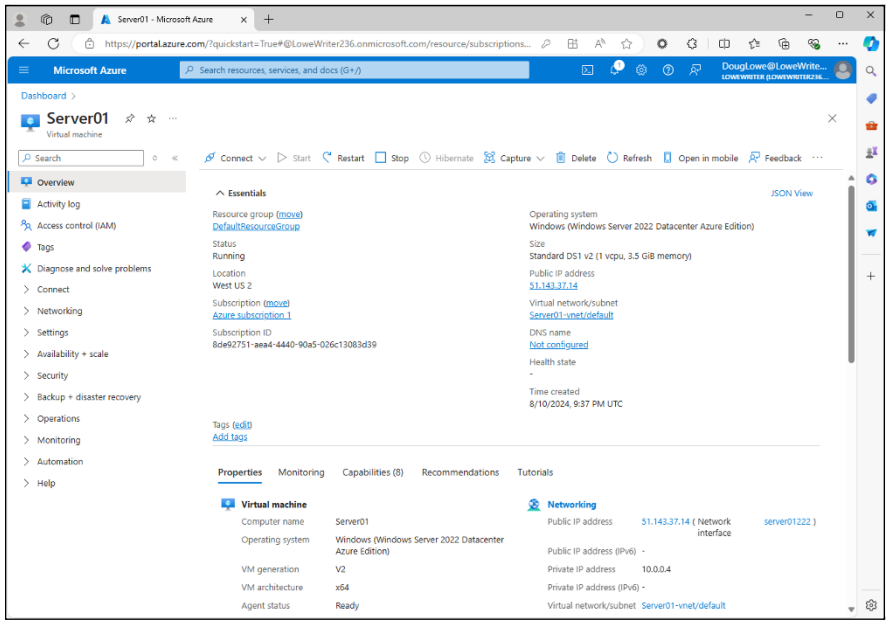


FIGURE 3-11: Managing a server from the Azure portal.

Across the top of the VM's resource page, you'll find a strip of icons useful for managing the machine:



- **Connect:** Lets you connect to the VM (see "Connecting to an Azure Virtual Machine," later in this chapter).
- **Start:** Starts the VM, if it isn't already running.
- **Restart:** Restarts the VM.
- **Stop:** Stops the VM.
- **Capture:** Creates an image of the VM. You can later use this image to create a new VM.
- **Delete:** Deletes the VM.
- **Refresh:** Updates the display with current information.
- **Open in Mobile:** Displays a QR code which you can scan to open the VM in the Azure mobile app, which you should first install on your mobile phone.

Notice also that along the left side of the resource page for a VM is a menu of other activities you can use to manage the machine. For example, you can adjust the machine's size configuration by clicking the Size option (found under Availability + Scale). This displays a list of the VM size options, as shown in Figure 3-12.

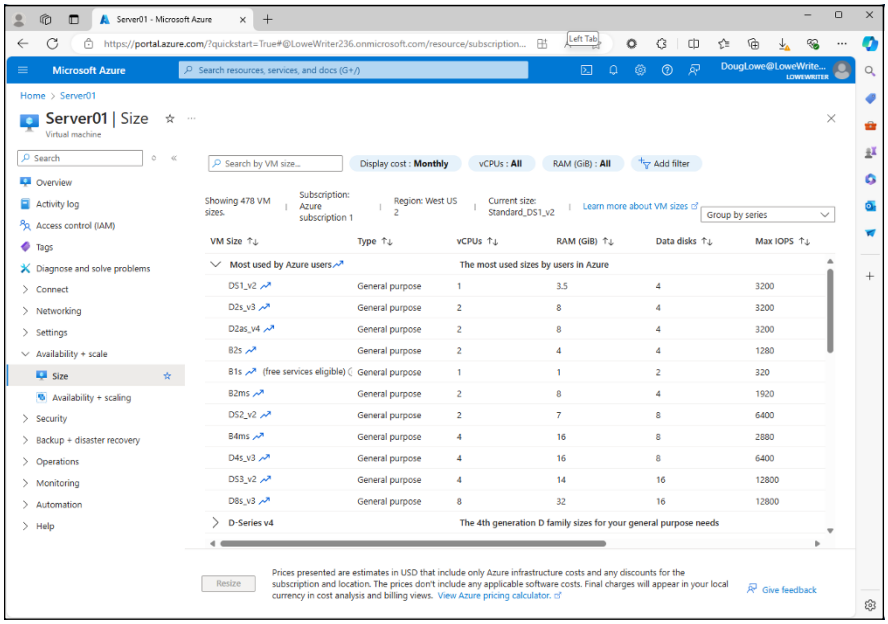


FIGURE 3-12: Resizing an Azure virtual machine.

Connecting to an Azure Virtual Machine

When an Azure VM is up and running, you can connect to it remotely using Remote Desktop Connection, just as you can connect to any other VM. The easiest way to do so is to follow these steps:

1. **Open the VM in the Dashboard (refer to Figure 3-11).**
2. **Click Connect in the toolbar and choose RDP.**

   The Connect screen, shown in Figure 3-13, appears.
3. **Click Download RDP file.**

   This downloads a remote desktop connection file (`.rdp`).
4. **Save the downloaded RDP file to a convenient location.**

   The procedure for doing this depends on the browser you're using.
5. **Click the downloaded RDP file to open the connection.**

   Remote Desktop Connection fires up, connects to the Azure VM, and then asks for credentials.
6. **Enter the username and password for the administrator account you specified when you created the machine, and then click OK.**

   You are connected to the VM, as shown in Figure 3-14.

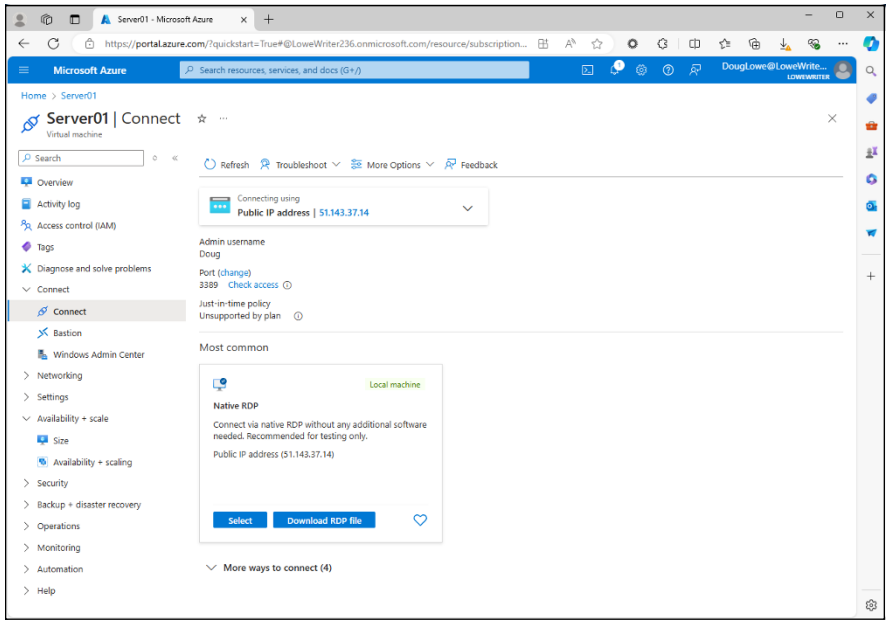


FIGURE 3-13: The Connect page.

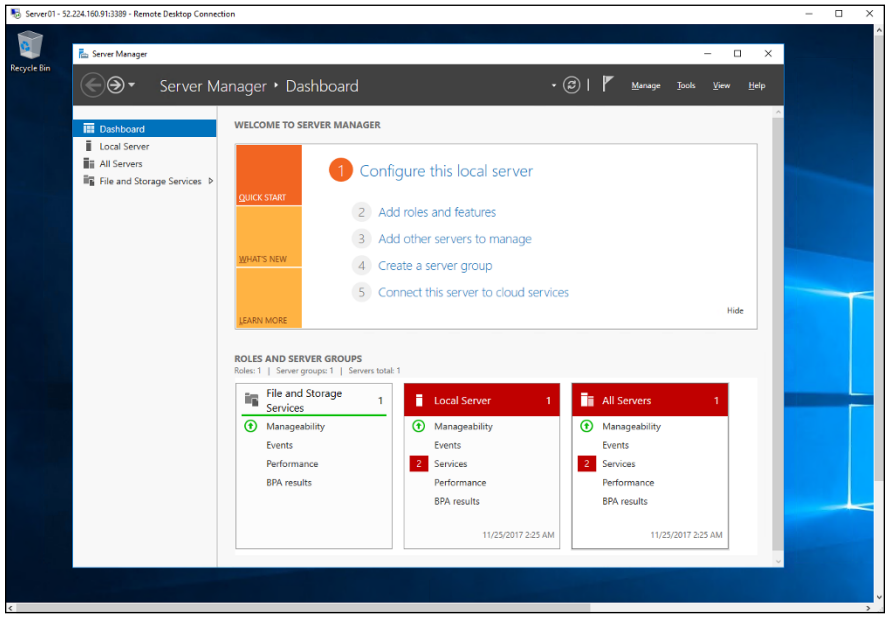


FIGURE 3-14: Connecting to an Azure virtual machine.



WARNING

When you're finished exploring the Azure VM and you've learned what you need to know, don't forget to delete the VM from the Azure environment before the free subscription period expires. If you forget, you may find unexpected charges on your credit card!

IN THIS CHAPTER

- Looking at AWS cloud service offerings
- Signing up for an AWS account
- Creating a virtual machine in AWS
- Managing an AWS virtual machine instance
- Connecting to an AWS virtual machine instance

Chapter 4
Amazon Web Services

Like Microsoft Azure, Amazon Web Services (AWS) is a cloud computing service that includes numerous ways to virtualize your IT infrastructure in the cloud. AWS is the grandfather of cloud-based infrastructure providers — it got its start way back in 2002. Since then, AWS has developed into the largest cloud provider in the world. Amazon's online retail space itself is hosted on AWS, as are many other familiar services, including Airbnb, Netflix, and Pinterest.

AWS provides services that span the full range of cloud-based services, including Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). In this chapter, I give you a brief look at the range of cloud services that AWS provides. Then you'll learn how to create and manage a virtual machine (VM) on AWS.

Looking at What Amazon Web Services Can Do

AWS has well over 2,000 distinct services available on its platform. They're organized into several categories, including (but not limited to), the following:

- **Compute:** Provides cloud-base virtual computing resources. The main service in this category is Amazon's cloud-based virtualization platform, known as *Amazon Elastic Compute Cloud* (EC2). With EC2, you can create and manage VMs that run at Amazon's data centers. You can select from several different pricing models, depending on your needs. Multiple operating system (OS) choices are available, including Windows Server 2022, several variations of Linux, and even macOS. And you can configure VMs with a single processor and as little as 1GB of RAM to as many as 192 processor cores and a whopping 24TB of RAM.

  Naturally, the larger the machine configuration and the more it is used, the more it will cost. The smallest server can be run for just a few dollars per month (in fact, free for 12 months) — perfect for experimenting with AWS to familiarize yourself with its features.
- **Networking & Content Delivery:** Lets you set up virtual networks that enable your AWS cloud components to communicate with one another and also with your physical network. Amazon Virtual Private Cloud (VPC) lets you set up a private network at Amazon's data centers so that you can extend your own private network into Amazon's cloud, allowing EC2 machines to seamlessly integrate into your own network. In addition, this category includes several other specialized features for delivering various types of content.
- **Storage:** Amazon Elastic File System (EFS) is a cloud-based storage system designed to work with EC2 VMs to provide cloud-based data storage.
- **Database:** Relational Database Service (RDS) provides basic relational database capabilities similar to Microsoft SQL Server, and several other database offerings provide non-relational database services.
- **Business Applications:** Features such as Alexa for Business, email, and Chime (a cloud-based phone system).
- **Security, Identity & Compliance:** Provides basic directory and security services for AWS.
- **Developer Tools:** Provides features for developers creating and managing custom applications on AWS.
- **Management & Governance:** Features for managing the AWS environment.

- **Machine Learning:** A variety of tools for artificial intelligence (AI).
- **Analytics:** Features for managing and analyzing large data sets.
- **End-User Computing:** Provides several features aimed at end users, including a desktop virtualization solution called Workspaces and a document-management solution called WorkDocs.

But wait, there's more! AWS also includes Internet of Things (IoT) solutions for managing your coffee pots and toasters, features for game development features, and who knows what else!

Creating an Amazon Web Services Account

Before you can use AWS, you must first set up an AWS account. The good news is that Amazon offers a free account you can use to experiment with AWS. Basic AWS services are free for 12 months, which gives you plenty of time to familiarize yourself with the many capabilities and features of AWS.

Setting up the free account is easy, but you'll have to fork over a credit card number. So you'll want to keep a good eye on your account, just in case you step over the line of what's free and begin incurring monthly charges.

To set up your account, just browse to `http://aws.amazon.com` and follow the links to set up a free account.

Here's what you get your first year with the free account:

- 750 hours per month of compute usage on a small VM called a *micro instance*, which has just one processor core and 1GB of RAM)
- 5GB of EFS storage
- 750 hours per month of Amazon RDS relational database
- An assortment of other free services, too detailed to list here

Note that all these services are free for a period of one year. After the first year, normal charges are incurred.


TIP

The full list of what is free for 12 months is detailed and definitely subject to change. I suggest you examine the details to be sure you understand exactly what's free and what will be charged. And check your billing summary frequently to avoid surprises.

Examining the Amazon Web Services Console

When you've created your free AWS account, you can access the AWS Console by following these steps:

1. **Go to** `https://aws.amazon.com/console`.

   The AWS home page appears, as shown in Figure 4-1.

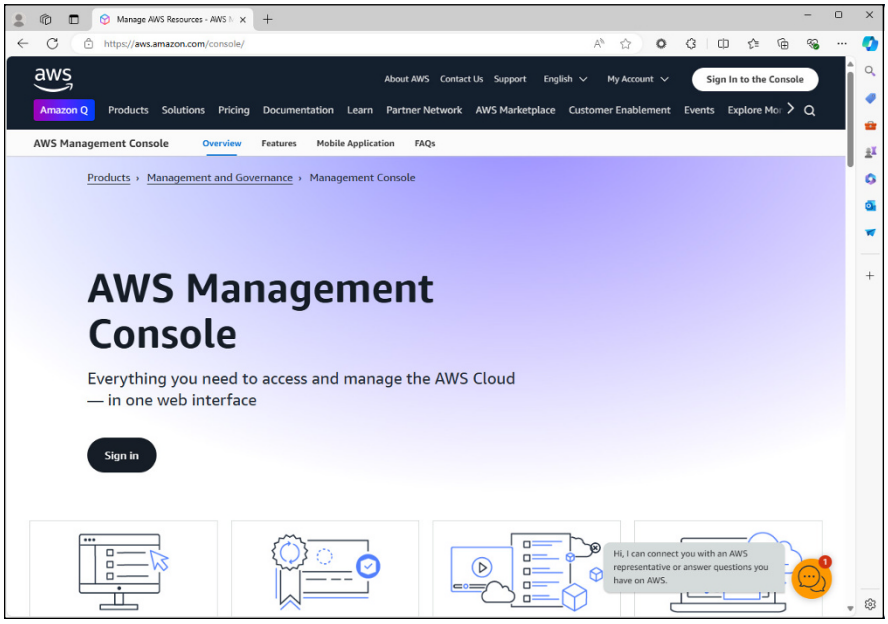


FIGURE 4-1: The AWS home page.

2. **Click the Sign In button.**

   You're taken to a sign-in screen.

3. **Enter your username and click Next.**

   You're prompted for your password.

4. **Enter your password and click Sign In.**

   The AWS Console appears, as shown in Figure 4-2.

Explore the console! Start by clicking Services in the menu bar at the top of the page. This reveals a menu of services you can access via the console, as shown in Figure 4-3.

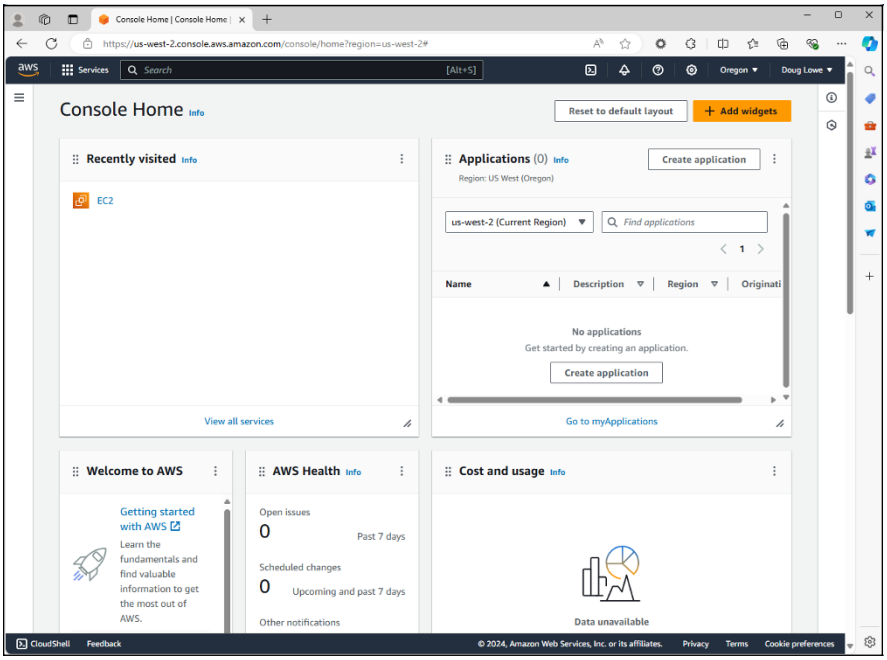


FIGURE 4-2: The AWS Console.

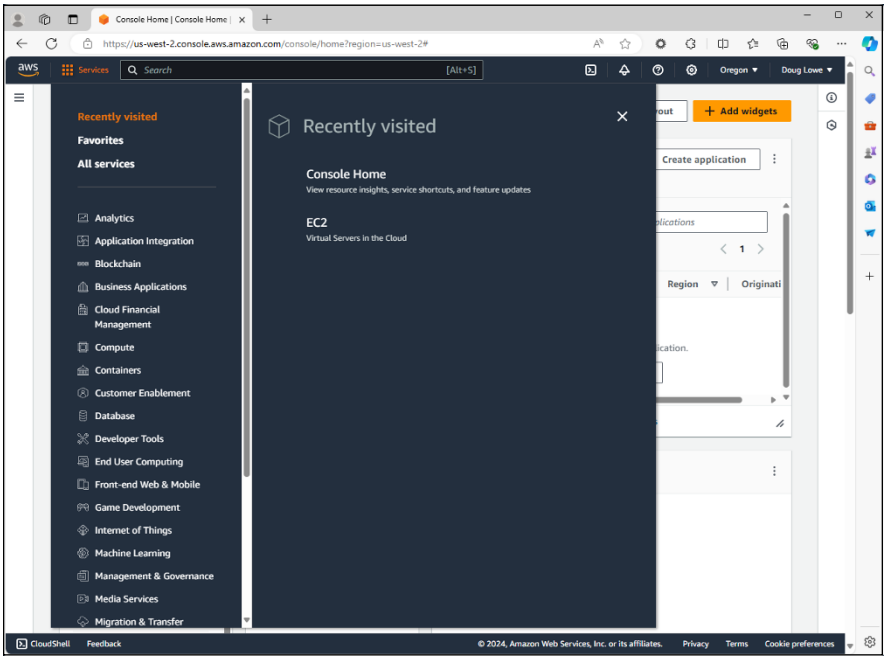


FIGURE 4-3: The Services menu.

You can click any of the services listed in this menu to view the dashboards for the various services. For example, Figure 4-4 shows the EC2 Dashboard, which shows information about EC2 VMs.

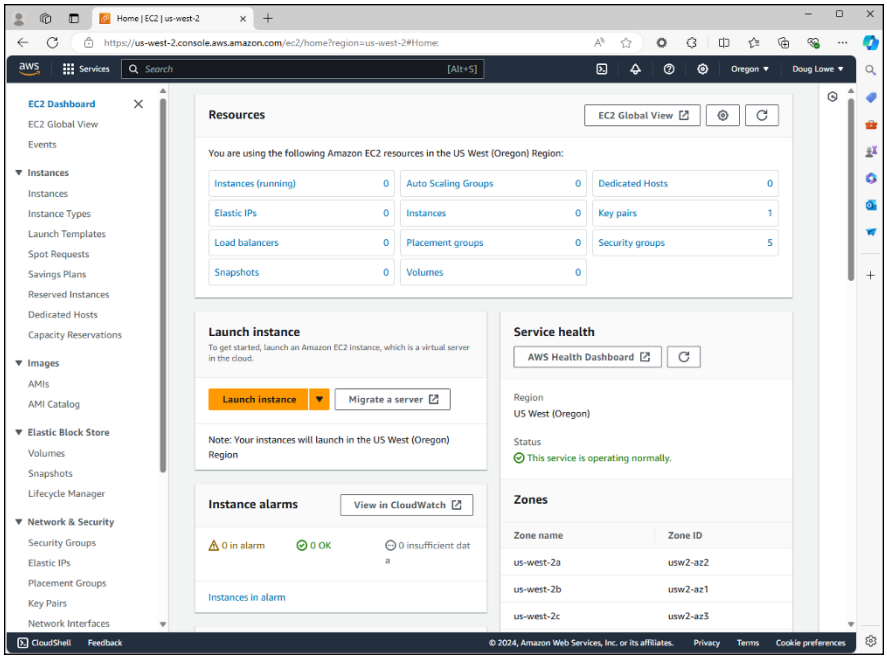


FIGURE 4-4: The EC2 Dashboard.

Creating a Windows Virtual Machine

To create a VM, navigate to the EC2 Dashboard as described in the previous section (refer to Figure 4-4). Then follow these steps:

1. **Click the Launch Instance button.**

   The Launch an Instance page, shown in Figure 4-5, appears. Note that this is a really long page, so you'll need to scroll to view and configure all the settings. Each setting section can be collapsed if you want to limit the amount of scrolling needed.

2. **Type a name for your instance.**

   For this example, I use the name Server01.

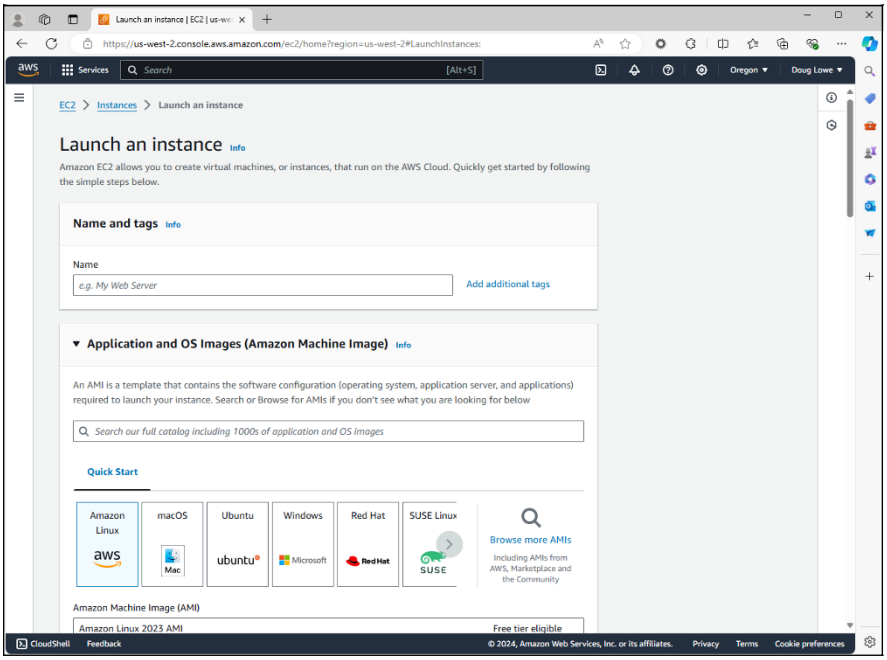


FIGURE 4-5: The Launch an Instance page.

3. **Type** Windows Server **in the Search box and press Enter.**

   You're taken to the Choose an Amazon Machine Image (AMI) page, shown in Figure 4-6. Here, you can see the predefined server images provided by Amazon for various versions of Windows Server. At the time I wrote this, the most current version was Windows Server 2022, but by the time you read this, Amazon may have made Server 2025 images available.

4. **Locate the Microsoft Windows Server 2022 Base image and click it's Select button.**

   You're taken back to the Launch an Instance page with the AMI image set to the Windows Server 2022 Base image you selected.

5. **Collapse the Application and OS Images section so you can access the Instance Type setting.**

   Refer to Figure 4-7 to see the Instance Type setting. As you can see, the instance named t2.micro instance is selected by default. This instance type provides a single processor core and 1GB of memory, which is the largest configuration that's eligible for free usage. If you want, you can select a more powerful image, but you'll have to pay a bit for the beefier configuration.

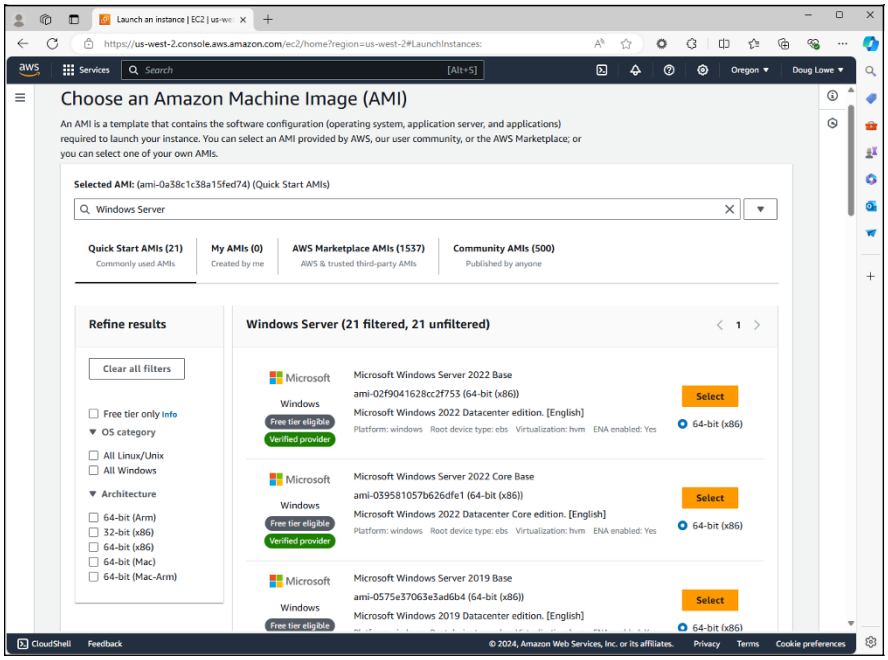


FIGURE 4-6: Choosing an Amazon Machine Image.

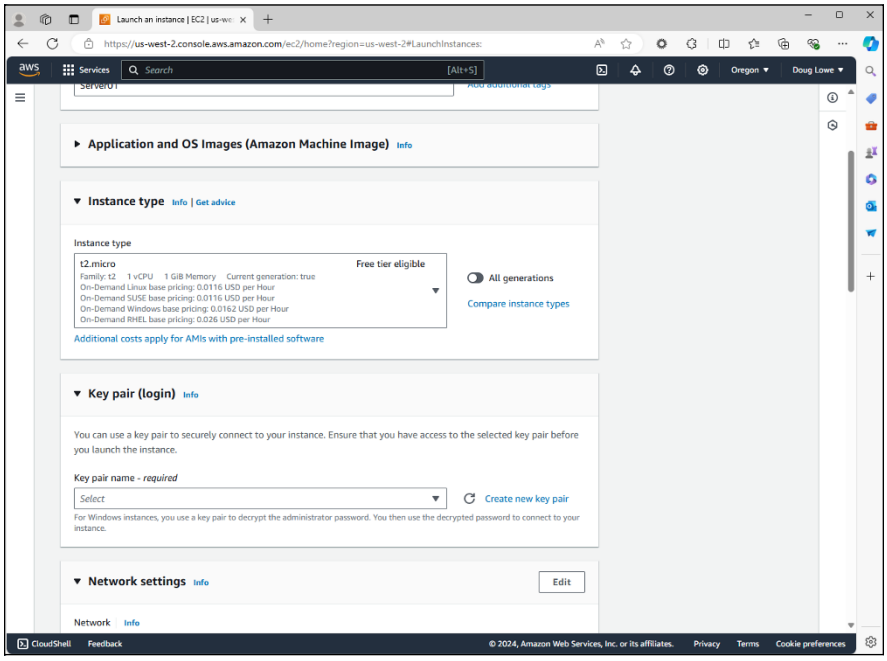


FIGURE 4-7: Choosing the instance type.

6. **Collapse the Instance Type section so you can access the Key Pair (Login) section.**

   You can see this section in Figure 4-8.

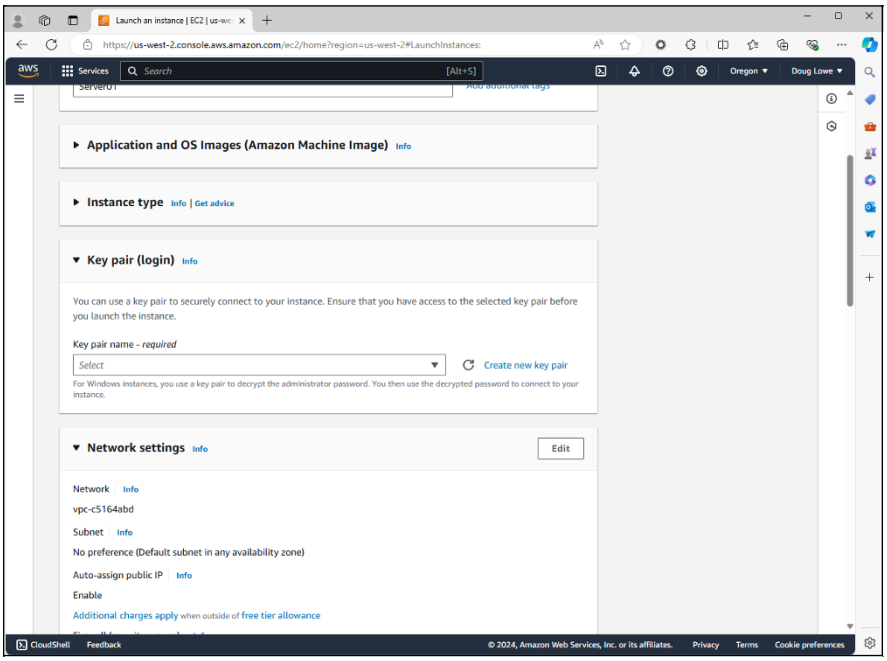


FIGURE 4-8: Setting the key pair.

A *key pair* is a combination of a public key that AWS keeps and a private key stored in a file that you're responsible for storing. You *must* keep the private key file in a safe place; without it, you won't be able to access your instance!

The Key Pair section provides a drop-down list from which you can select existing key pairs you've previously created. But when you're first starting out, you won't have any available key pairs. You'll remedy that in the next step.

7. **Click Create New Key Pair.**

   The Create Key Pair page, shown in Figure 4-9, appears.

8. **Enter a name for the key pair, and then click Create Key Pair.**

   Use whatever name you want, but make sure it's memorable. When you click Create Key Pair, AWS creates the key pair and automatically downloads it as a `.pem` file. You're then returned to the Launch an Instance page with the name of your newly created key pair selected in the drop-down list.

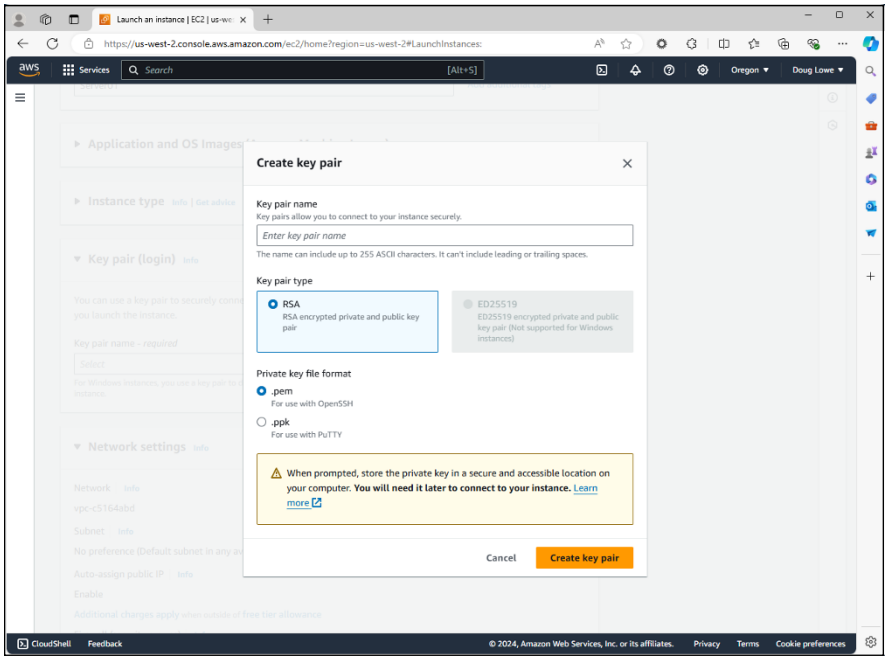


FIGURE 4-9: Creating a new key pair.



WARNING

Save the key pair file in a safe location! If you lose it, you won't be able to access the new virtual machine.

9. **Collapse the Key Pair section and scroll down so you can access the Network Settings section.**

   Figure 4-10 shows the network settings options.

10. **Review the network settings and apply any appropriate changes.**

    Most of the time, the default network settings will be appropriate. For example, you may want to apply a static IP address rather than use DHCP. Click the Edit button if you want to make changes.

11. **Collapse the Network Settings section to reveal the Configure Storage section (see Figure 4-11).**

12. **Configure the storage for the VM.**

    By default, a single disk volume called the *root volume* is created for the instance. You can change the amount of space allocated for the root volume (the default is 30GB, which I usually increase to 60GB), and you can change the disk type. The default is General Purpose SSD, but you can change it to less expensive Magnetic Disk if you want.

    You can add additional disk volumes by clicking the Add New Volume button. When you click this button, an additional disk volume is added; you can then specify the size and volume type for the new volume.

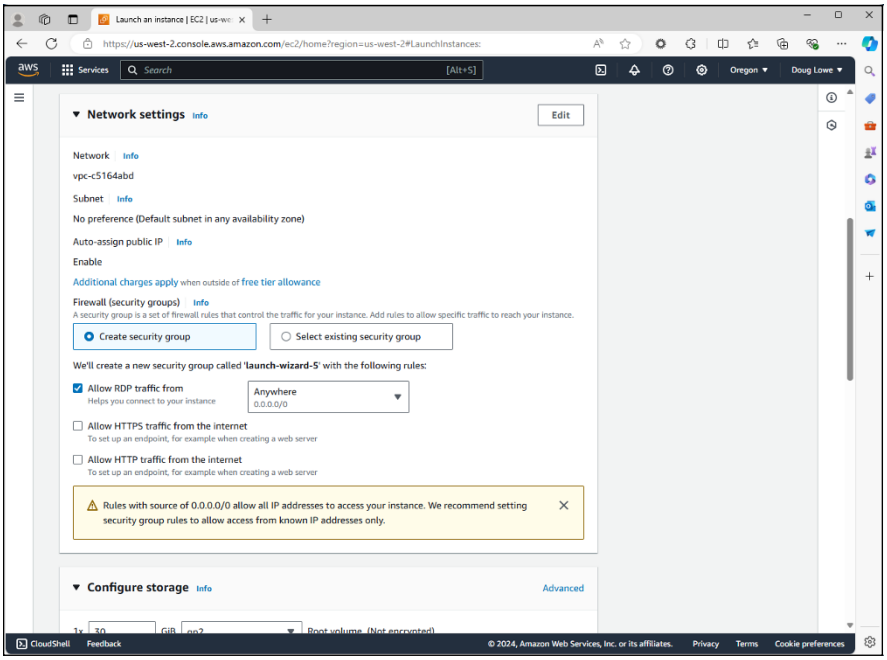


FIGURE 4-10: Configuring the network settings.

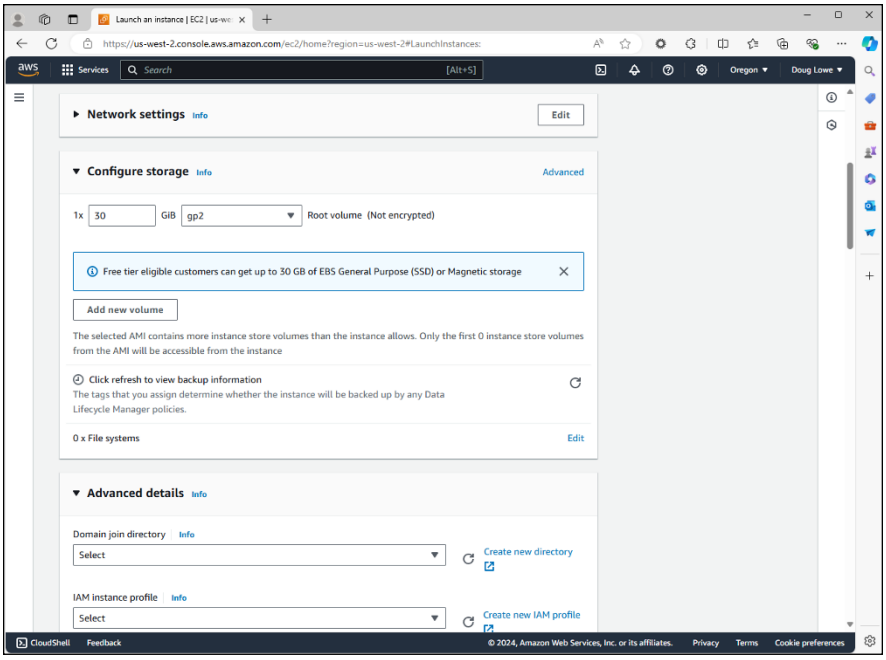


FIGURE 4-11: Configuring storage.

13. **Collapse the Configure Storage section to view the Advanced Details section (see Figure 4-12).**

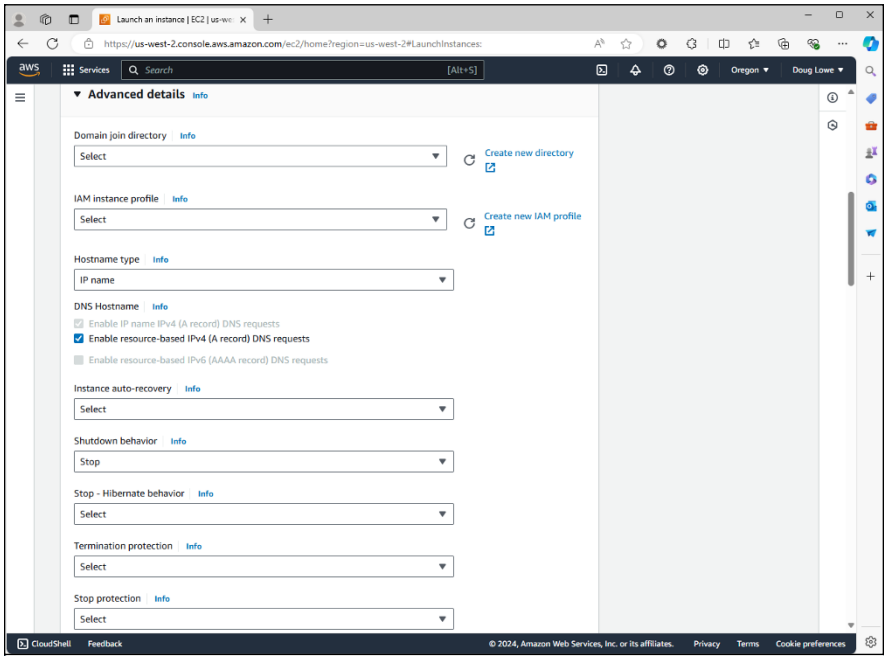


FIGURE 4-12: The Advanced Details section.

14. **Review the Advanced Details section and change any settings you want.**

    Most of these options can be left at their default settings.

15. **Collapse the Advanced Details section to view the Summary Section.**

    This section, shown in Figure 4-13, provides a summary of the instance you're about to create. It also allows you to create more than one instance, if you want.

16. **Click Launch Instance.**

    AWS creates your new instance and displays the Success page shown in Figure 4-14 when the instance is ready for you to use.

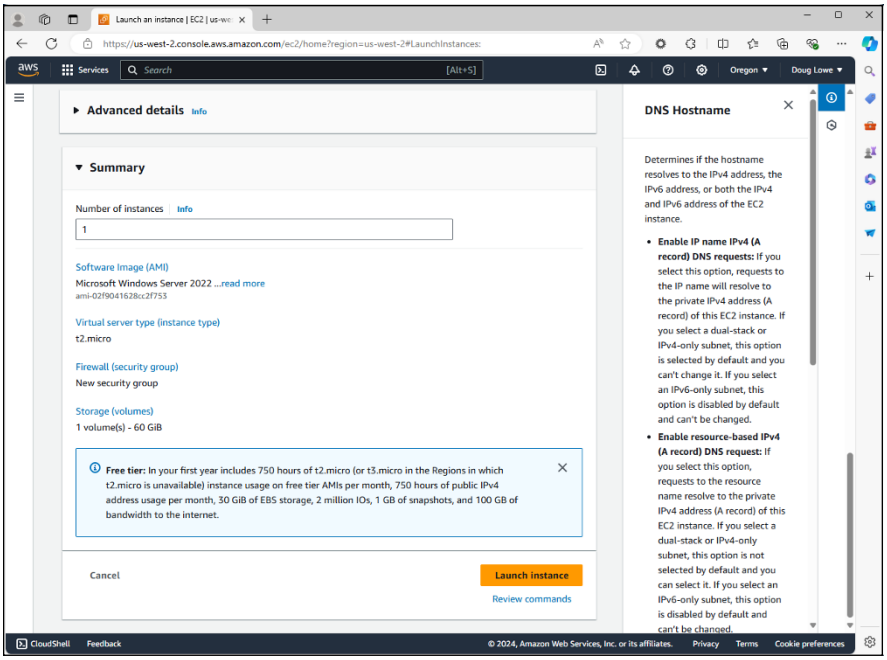


FIGURE 4-13: The Summary section.

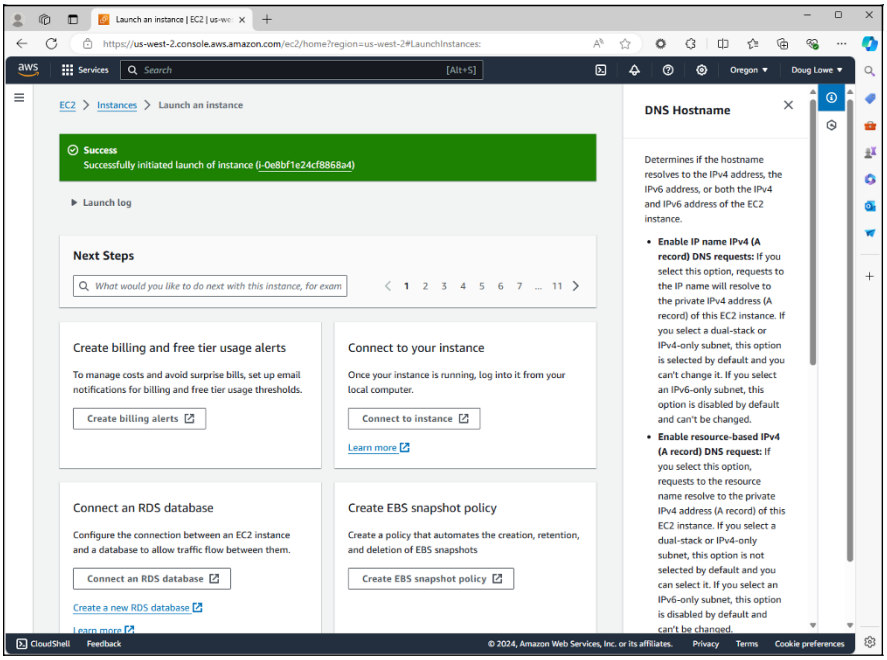


FIGURE 4-14: Success!

Managing an Amazon Web Services Virtual Machine

You can manage your VM instances by opening the EC2 Dashboard and then clicking Instances in the menu that appears at the left side of the Dashboard page. This brings up a list of all EC2 VM instances, as shown in Figure 4-15.

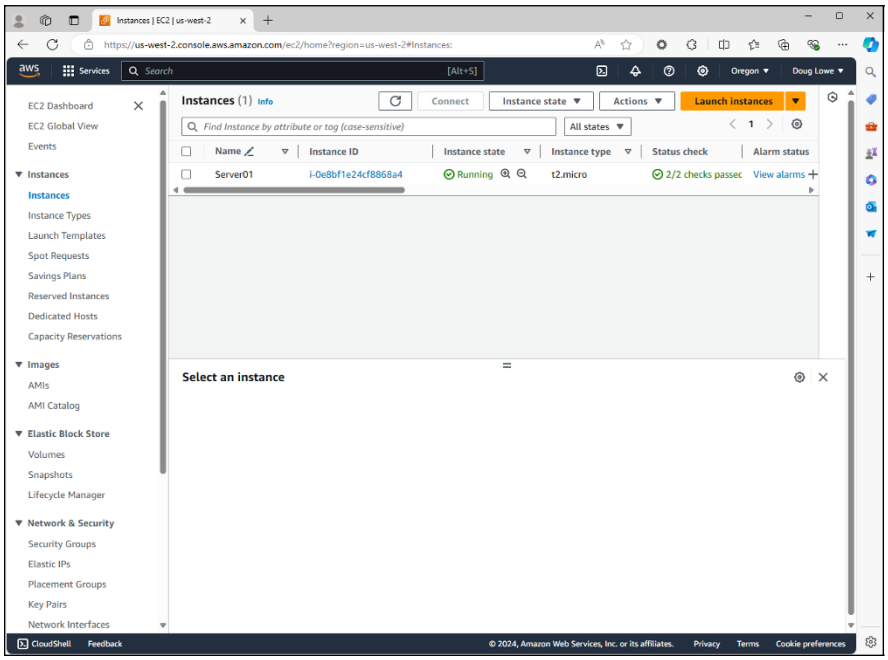


FIGURE 4-15: Managing EC2 instances.



TIP

To add or change the name of an instance, hover the mouse over the Name column for the instance you want to rename and click the pencil icon that appears. You can then type a name for the instance. (In Figure 4-14, I've already changed the name of the instance created in the previous section to WIN19-01.)

To manage an instance, right-click anywhere in the row for the instance you want to manage. This brings up a context menu with the following commands:

- **Launch Instances:** Lets you create a new AWS instance.
- **Launch Instances from Template:** Lets you create a new AWS instance from a template.

- **Migrate a Server:** Lets you migrate existing servers such as on-premises or Azure servers to AWS.
- **Connect:** Connects to the instance using Remote Desktop Connection.
- **Stop Instance:** Shuts down the OS.
- **Start Instance:** Starts the VM.
- **Reboot Instance:** Reboots the VM.
- **Hibernate Instance:** Hibernates the instance, if allowed.
- **Terminate Instance:** Terminates the instance. This permanently deletes the instance, so use this option only when you no longer need the instance.
- **Instance Settings:** Lets you change settings for the instance. Note that if the instance is stopped, you can change the instance type to increase the amount of RAM or the processor resources available to the instance.
- **Networking:** Changes network settings for the instance.
- **Image and Templates:** Create an image of the instance that you can later use to create new instances. Or, create a template from this image so you can easily create other similar instances.
- **Monitor and Troubleshooting:** Enables monitoring services for the instance.

Connecting to an Amazon Web Services Virtual Machine

When an AWS VM is up and running, you can connect to it remotely using Remote Desktop Connection, just as you can connect to any other VM. The easiest way to do so is to follow these steps:

1. **In the EC2 Instance Dashboard, right-click the instance you want to connect to and choose Connect, and then click the RDP Client tab.**

   The Connect to Instance page appears (see Figure 4-16).

2. **Click Download Remote Desktop File.**

   This downloads a remote desktop connection file (`.rdp`).

3. **Save the RDP file to your computer.**

   The procedure to do this varies depending on the browser you're using.

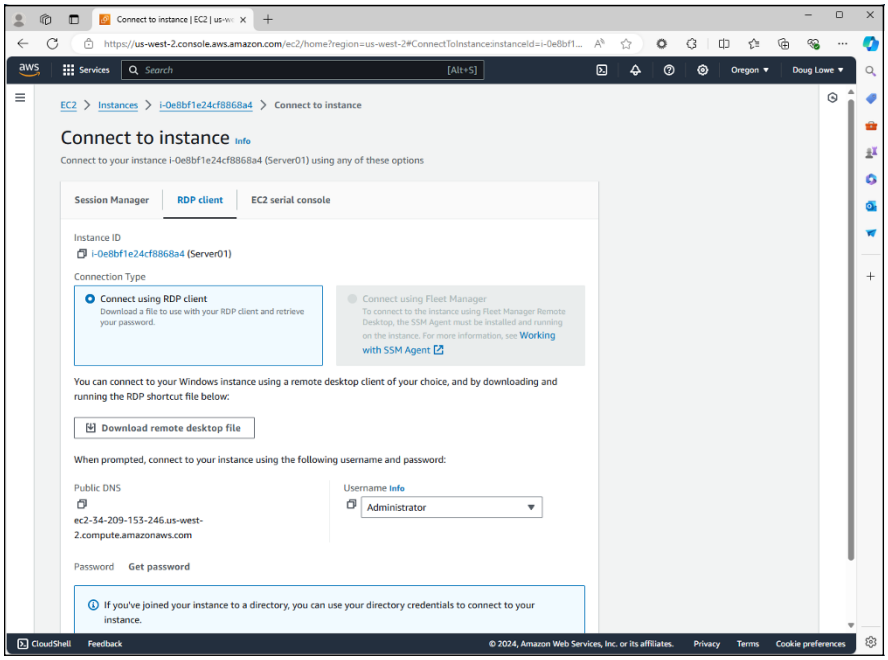


FIGURE 4-16: Connecting to an instance.

4. **Back in the Connect to Your Instance dialog box, click Get Password.**

   The Get Windows Password page appears. Here, you're asked to upload the Private Key File you saved when you created the instance.

5. **Click the Upload Private Key file button, navigate to your private key file, select it, and click Open.**

   AWS shows the contents of the key path file in the text box.

6. **Click Decrypt Password.**

   AWS decrypts the password and displays it, as shown in Figure 4-17. (Well, sort of — I airbrushed out the actual password. Don't get all excited, though. I've already terminated this machine, so don't waste your time trying to hack into it!)



TIP

Notice that the password generated by AWS consists of 32 random characters. You'll never in a lifetime commit that to memory, and you should under no circumstances copy and paste this password into a document on your computer. My recommendation is that when you log in to the server, you change its Administrator password to something you can remember without writing down.

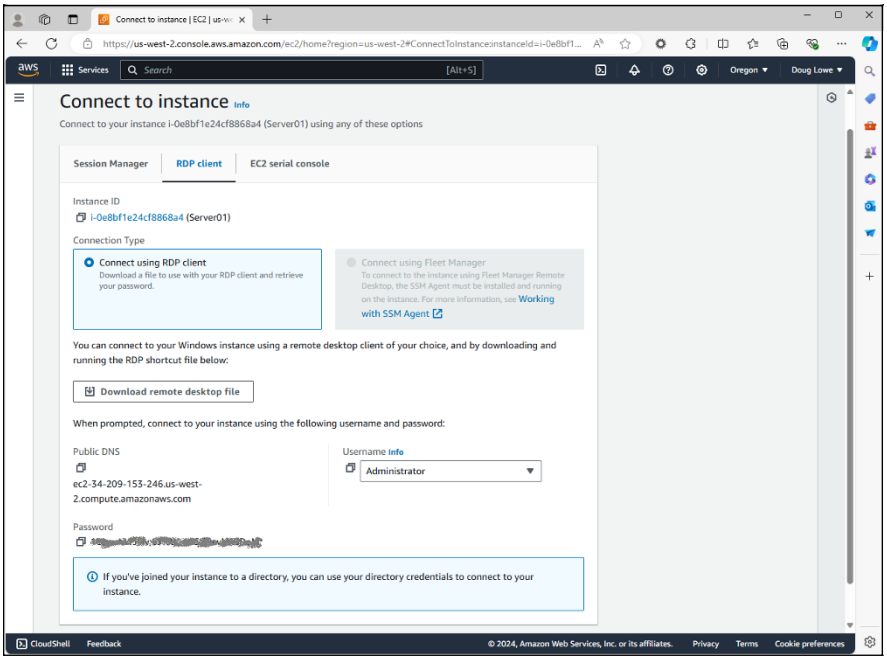


FIGURE 4-17: AWS shows you the Administrator password.



Notice also the Copy to Clipboard icon next to the decrypted password. Click this button to copy the password to the clipboard. Then, in Step 9 (coming up!), you can just paste the random password to log in to the server.

7. **Navigate to the `.rdp` file you saved in Step 4 and double-click to open it.**

   Remote Desktop Connection fires up, connects to the instance, and prompts you for credentials to log in.

8. **Enter the username (Administrator) and password (see Step 7) and click OK.**

   Congratulations! You've successfully logged in to your first AWS EC2 instance, as shown in Figure 4-18!



FIGURE 4-18: The console of an EC2 instance.

IN THIS CHAPTER

- Considering the possibility of virtualizing desktops instead of servers
- Looking at two approaches to virtualizing desktops
- Working with VMware's Horizon View
- Using Citrix XenApp

Chapter 5
Desktop Virtualization

Virtualization is most often applied to server computers. However, more and more organizations are also applying virtualization to desktops, replacing traditional Windows desktop computers with virtual desktops that are accessed from a nontraditional type of device.

In this chapter, I present a basic overview of desktop virtualization and its benefits and introduce you to some of the technologies that can be used to implement virtual desktop infrastructure (VDI).

Introducing Desktop Virtualization

The term *desktop virtualization* refers to any software that separates an end-user's Windows desktop environment from the hardware that the environment runs on. Desktop virtualization is meant to address some of the fundamental weaknesses of the traditional practice of giving each user their own Windows workstation.

Here are just a few of the problems that desktop virtualization addresses:

- Windows workstations must be configured individually for each user. If your organization has 100 workstations and you decide to update your accounting software, you have to figure out how to deploy the update to 100 computers.

- Windows software frequently needs to be updated. Updates are normally delivered via Windows Update. However, deploying Windows updates separately to all your desktop computers is fraught with peril. A particular Windows update might work well on 99 percent of all computers, which means that if your organization has 100 computers, that update is likely to not work on at least one of them. That means a trip to that computer to diagnose the problem caused by the update and get the user back up and running.
- If a user's computer fails, that computer must be repaired or replaced. To replace the computer, you'll need to rebuild the user's profile, reinstall the user's applications, and perform other configuration work to restore the user's desktop environment.
- Windows computers have a dreaded thing called the `C:` drive. Any data stored on the `C:` drive belongs to that computer alone and is not generally backed up to the network. Thus, if the user's `C:` drive dies, its data is likely to die with it.
- If a user moves to another desk or office, the user must take their computer with them.
- If a user wants to work from home, the user can't easily access their desktop environment from their home computer. There are solutions for this problem, such as virtual private network (VPN) or remote access software like GoToMyPC (`www.gotomypc.com`), but those solutions introduce problems of their own.
- If a user has a laptop computer in addition to a desktop computer, the user must make a special effort to ensure that the data on the desktop computer is synchronized with the data on the laptop. (Microsoft's OneDrive is a good way to do that.)
- The user may have devices with different platforms than their desktop computer. For example, a user might have a Windows computer at work, a MacBook Pro at home, and an Apple iPad for the road. These platforms aren't compatible with one another, so not all software can run on all three.

Desktop virtualization addresses all these problems (and more) by moving the user's desktop environment from a desktop computer to a central host computer. Then the user can access the desktop environment from any device that is compatible with the VDI technology chosen to virtualize the desktop. The advantages of this arrangement are many:

- If the user's computer dies, the user's desktop does not die with it. You can replace the failed computer with any other computer and simply reconnect to the virtual desktop.
- Operating systems and application software can be centrally managed. There is no need to visit a user's desk to install or update software.

- The user's desktop can be accessed from different types of devices. So, a user can access their desktop from a Windows computer, a MacBook, an iPad, an Android table, or even from an iPhone or Android phone.
- You can use thin clients at users' desks rather than full-blown Windows computers. A *thin client* is a small computer that has just enough processing power (CPU, RAM, and disk) to run the client piece of the desktop virtualization platform. Typically, the thin client runs an embedded version of Linux that is specially configured to run the client software that accesses the virtual desktop. In most cases, the end user is not even aware that this is happening — to the user, the experience is identical to having a standard Windows computer at their desk.
- In some desktop virtualization environments, multiple users share a common Windows environment, which means that an application needs to be installed only once for it to be available for multiple users, and operating system patches need to be applied just once rather than to multiple computers.
- All data is kept on the host computers, which means the data can be centrally managed and backed up.

Considering Two Approaches to Desktop Virtualization

There are at least two distinct approaches to implementing desktop virtualization. The first approach is to simply create a separate virtual machine for each user and provide a way for the users to efficiently connect to their virtual machines. This approach is usually referred to as *virtual desktop infrastructure* (VDI). VDI solutions are usually built using traditional virtualization products such as Microsoft's Hyper-V or VMware's ESXi hypervisor.

The second approach is to use a single server that is designed to support multiple users and provide a way for each user to connect to their session on the server. This approach is often called *terminal services*, because it's based on the terminal services role that is a standard part of all versions of Windows Server.



Technically, with Windows Server 2008, Microsoft changed the name of Terminal Services to *Remote Desktop Services* to emphasize the role of Terminal Services for virtualizing desktops. The IT industry is pretty reluctant to change its phraseology, however, so most IT professionals still call it *terminal services* even though that term has been obsolete for almost a decade.

The remaining sections of this chapter describe two popular desktop virtualization products that use these two approaches. The first is VMware's Horizon View, which builds on VMware's virtualization platform. The second is Citrix XenApp, which builds on Windows Terminal Services.

Looking at VMware's Horizon View

With VMware's virtualization infrastructure, you could easily implement desktop virtualization by simply creating virtual machines for each of your users' desktops and having your users connect to their virtual machines using Remote Desktop Connection (RDC). However, you'll quickly start to realize some of the limitations of this approach.

First, you'll probably discover that the RDC client is not very efficient when it comes to intensive graphics applications. Watching YouTube videos over RDC can be frustrating, as can working with graphically oriented programs such as Adobe Photoshop.

You'll also discover that managing users' access to virtual desktop machines is difficult with vSphere. vSphere is designed to create and manage virtual servers that are typically accessed only by IT personnel. Access to those servers is controlled through Active Directory credentials; in other words, if you don't know the password, you can't log in. But vSphere isn't really designed to create hundreds of desktop VMs and make them available to hundreds of users.

To address these and other issues, VMware offers a product called VMware Horizon View that builds on the core functions of vSphere and adds features specifically designed for desktop virtualization. Here's a short list of some of the more important features of Horizon:

- **vSphere Desktop:** A version of vSphere specifically designed for running up to tens of thousands of desktop virtual machines
- **vCenter Desktop:** A version of vCenter designed specifically for managing virtual resources such as hosts, RAM, processors, and disk storage in a virtual desktop environment
- **Horizon View:** A management tool designed to provision and deploy virtual desktops
- **Horizon View Client:** Client software for accessing virtual desktops on a variety of platforms, including Windows, Mac, iOS, and Amazon devices
- **Horizon View Composer:** A tool for cloning desktop VMs and for managing software and operating system updates on pools of similar desktop VMs

Looking at Citrix XenApp

Citrix XenApp is a desktop virtualization environment that uses Windows Terminal Services to enable multiple users to access remote desktops from a variety of client devices, including Windows, Mac, iOS, and Amazon devices. Unlike VMware's Horizon View, XenApp does not create a separate virtual machine for each user. Instead, when users connect to XenApp, they log in to separate terminal services sessions on a common Windows Server. The users then have access to all the resources and applications that are available to the Windows Server.

Users connect with XenApp by using a client application called the Citrix Receiver, which can be run on Windows, Mac, iOS, or Android devices. Figure 5-1 shows Citrix Receiver running on a Windows 10 system.

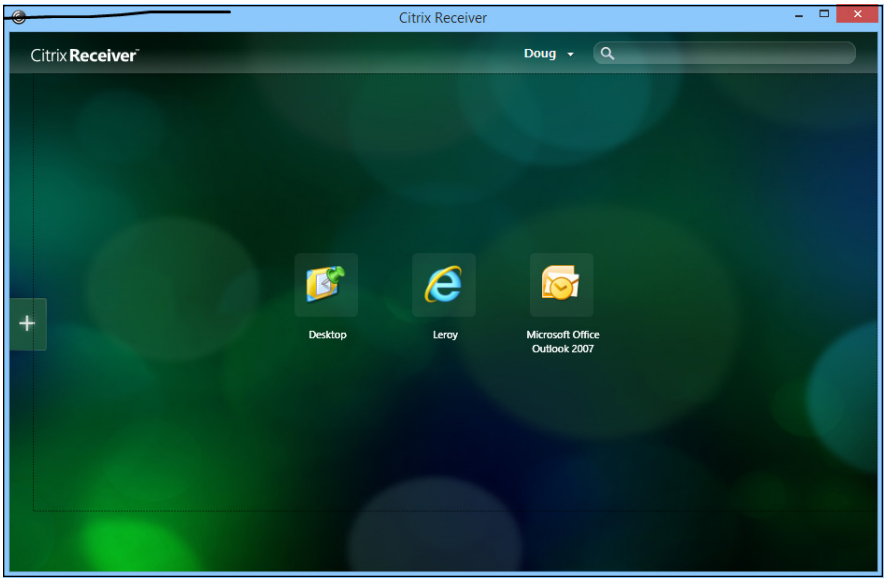


FIGURE 5-1: Citrix Receiver.

When you configure a XenApp server, you create and publish *desktops* and *applications* that users can remotely connect to. Then, when the user connects to the XenApp server, the user is shown the applications that they're authorized to use. For example, Figure 5-1 shows a typical Citrix Receiver screen connected to a XenApp server. In this case, the user is authorized to open a desktop and two applications.

If the user connects to a desktop, the user sees an entire Windows desktop environment, complete with a Start menu that grants access to applications, as well as Explorer to browse disk resources. Figure 5-2 shows the Citrix Receiver connected to a desktop.

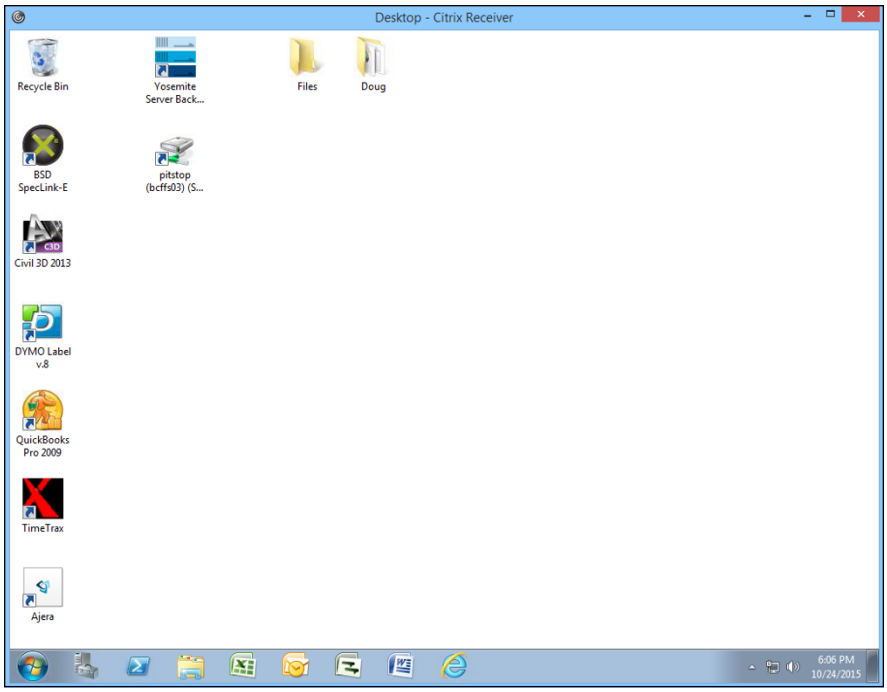


FIGURE 5-2: Connecting to a desktop.

On the other hand, a user can connect to an individual application rather than to a desktop. This feature is called *application virtualization.* When you connect to an application, that application runs on the server but only that one application's window is shown on the user's device. In other words, the user sees the individual application rather than a complete desktop. The user can then use the application as if the application were natively running on their device, even if the user's device is a non-Windows device.

Figure 5-3 shows an example of Microsoft Excel running on an iPhone 6 Plus via Citrix Receiver. As you can see, Citrix Receiver makes the desktop Microsoft Excel application available on my iPhone exactly as it's available on my Windows desktop. In fact, I could choose File➪Open to summon an Open dialog box which would allow me to browse the network to open any file that would be available to me from my Windows desktop. This feature effectively extends my desktop applications to my iPhone.

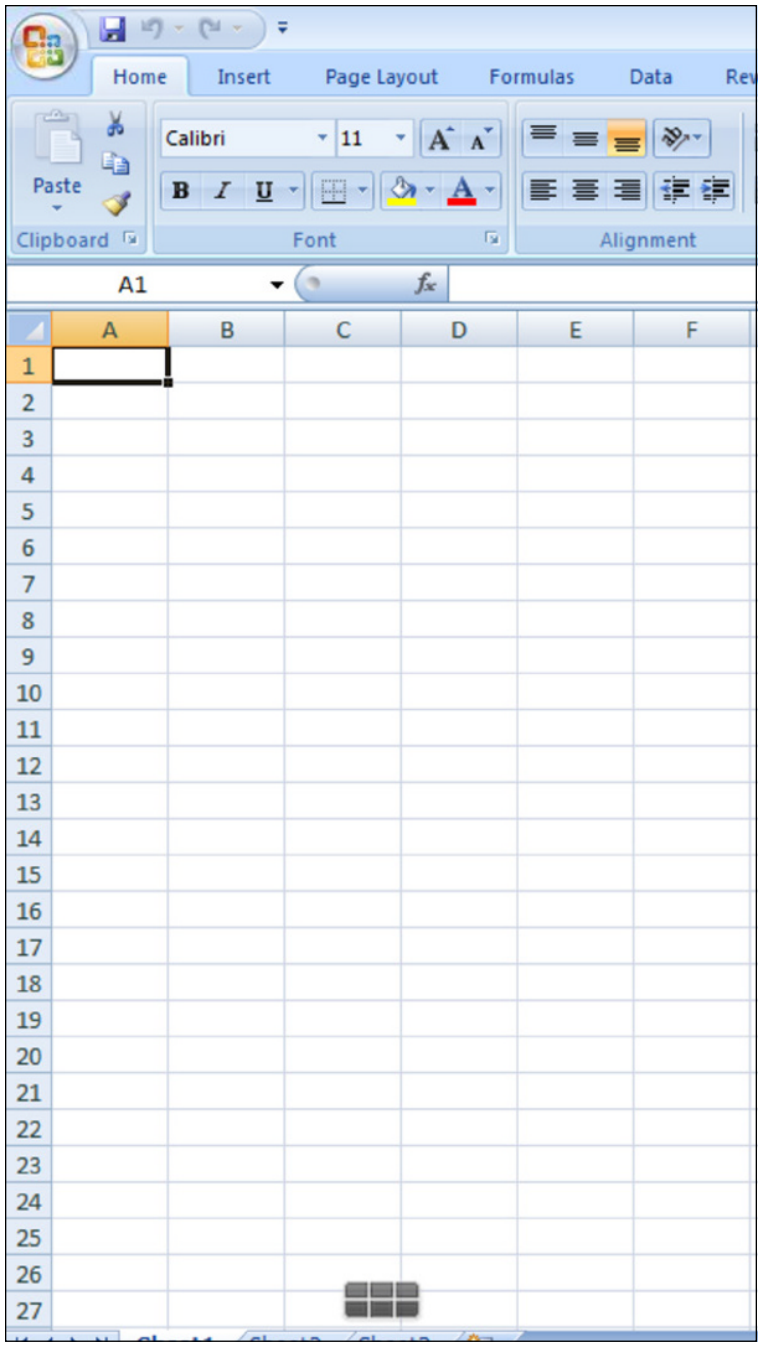


FIGURE 5-3: Viewing Excel on an iPhone.

6

Implementing Windows Server 2025

Contents at a Glance

CHAPTER 1: **Installing Windows Server 2025** 469

CHAPTER 2: **Configuring Windows Server 2025** 487

CHAPTER 3: **Configuring Active Directory** 497

CHAPTER 4: **Configuring User Accounts** 507

CHAPTER 5: **Configuring a File Server** 529

CHAPTER 6: **Using Group Policy** 543

CHAPTER 7: **Comandeering Windows Commands** 555

CHAPTER 8: **Using PowerShell** 583

IN THIS CHAPTER

- Making sure you have everything you need
- Planning how to install Windows Server 2025
- Installing Windows Server 2025
- Knowing what to do immediately after you install Windows Server 2025
- Helping your server perform different roles

Chapter 1

Installing Windows Server 2025

This chapter presents the procedures that you need to follow to install Windows Server — specifically, Windows Server 2025. Note that although the specific details provided are for Windows Server 2025, installing a previous version is very similar. So you won't have any trouble adapting these procedures if you're installing an older version.

Planning a Windows Server Installation

Before you begin the Setup program to actually install a Windows Server operating system, you need to make several preliminary decisions, as the following sections describe.

Checking system requirements

Before you install a Windows Server operating system, you should make sure that the computer meets the minimum requirements. Table 1-1 lists the official minimum requirements for Windows Server 2025. (The minimums for Windows Server 2016 are the same.) Table 1-1 also lists what I consider to be more realistic minimums if you expect satisfactory performance from the server as a moderately used file server.

TABLE 1-1 **Minimum Hardware Requirements for Windows Server 2025 (Standard Edition)**

Item	Official Minimum	A More Realistic Minimum
CPU	1.4 GHz	3 GHz
RAM	512MB	4GB
Free disk space	32GB	100GB

Note that there is no 32-bit version of Windows Server 2025. A 64-bit processor is required, but that shouldn't be a problem, as nearly all computers manufactured since around 2007 have 64-bit processors.

Note also that if you're installing Windows Server 2025 on a virtual machine, you'll need at least 800MB of RAM to complete the installation. After Windows Server 2025 is installed, you can back the RAM down to 512MB if you wish.

Reading the release notes

Like all versions of Windows Server, Windows Server 2025 provides a set of release notes that you should read before you start Setup, just to check whether any of the specific procedures or warnings it contains applies to your situation.

The release notes weren't yet available as I was writing this, but you can use your favorite search engine to search the web for "Windows Server 2025 release notes."

Deciding whether to upgrade or install

Windows offers two installation modes: full installation or upgrade installation.

A *full installation* deletes any existing operating system(s) it finds on the computer and configures the new operating system from scratch. If you do a full installation

onto a disk that already has an operating system installed, the full installation offers to keep any existing data files that it finds on the disk.

An *upgrade installation* assumes that you already have a previous Windows Server 2016 installation in place. The operating system is upgraded to Windows Server 2025, preserving as many settings from the previous installation as possible. You cannot upgrade Windows Server 2008 or earlier to Windows Server 2025.

Here are some points to ponder before you perform an upgrade installation:

- You can't upgrade a client version of Windows to a server version.
- With an upgrade installation, you don't have to reinstall any applications that were previously installed on the disk.
- Always perform a full backup before doing an upgrade installation! Or, if you're using a virtualization platform such as VMware or Hyper-V, make a snapshot copy of the VM.

Considering your licensing options

Two types of licenses are required to run a Windows Server operating system: a *server license*, which grants you permission to run a single instance of the server, and *Client Access Licenses* (CALs), which grant users or devices permission to connect to the server. When you purchase Windows Server, you ordinarily purchase a server license plus some number of CALs.

To complicate matters, there are two distinct types of CALs: per-user and per-device. *Per-user* CALs limit the number of users who can access a server simultaneously, regardless of the number of devices (such as client computers) in your organization. By contrast, *per-device* CALs limit the number of unique devices that can access the server, regardless of the number of users in your organization.

Thinking about multiboot

Windows includes a *multiboot* feature that lets you set up the computer so that it has more than one operating system. When you boot up the computer, you can select the operating system you want to boot up from a menu.


TIP

Although you may be tempted to use the multiboot features to maintain previous operating system installations, I recommend against it. A much better alternative is to install Windows Server into a virtual machine using virtualization technology such as Microsoft's Hyper-V or VMware. Virtualization allows you to install

a complete operating system such as Windows Server 2025 within an already-installed operating system. (Note that the Windows Server 2025 installation illustrated in this chapter and throughout this book is installed on a Hyper-V virtual machine running within Windows 10.)

Planning your partitions

Partitioning enables you to divide a physical disk into one or more separate units called *partitions.* Each disk can have up to four partitions. All four of the partitions can be primary partitions. A *primary partition* contains one — and only one — file system. Alternatively, you can create up to three primary partitions and one extended partition, which can be subdivided into one or more logical drives. Then each logical drive can be formatted with a file system.

Windows Server 2025 offers you two file systems: NTFS and ReFS. NTFS is the tried-and-true file system that has been the standard file system for going on 20 years. The partition that Windows Server 2025 boots from must be NTFS. However, other partitions can be either NTFS or ReFS.

Although you can set up partitions for a Windows Server in many ways, the following two approaches are the most common:

- **Allocate the entire disk as a single partition that will be formatted with NTFS.** The operating system is installed into this partition, and disk space that isn't needed by the operating system or other network applications can be shared.
- **Divide the disk into two partitions.** Install the operating system and any other related software (such as Exchange Server or a backup utility) on the first partition. If the first partition will contain just the operating system, 100GB is a reasonable size, although you can get by with as little as 32GB if space is at a premium. Then use the second partition for application data or network file shares.


TIP

Note that the disk partitioning scheme is independent of any hardware-based RAID configuration your server may employ. Your server may actually include five physical hard drives that are combined by the hardware disk controller to form a single logical drive, for example. Within this logical drive, you can create one or more operating-system partitions.

Deciding your TCP/IP configuration

Before you install the operating system, you should have a plan for implementing TCP/IP on the network. Here are some of the things you need to decide or find out:

- What are the IP subnet address and mask for your network?
- What is the domain name for the network?
- What is the host name for the server?
- Will the server obtain its address from DHCP?
- Will the server have a static IP address? If so, what?
- Will the server be a DHCP server?
- What is the Default Gateway for the server (that is, what is the IP address of the network's internet router)?
- Will the server be a DNS server?



TIP

In almost all cases, you should assign the server a static IP address.

For more information about planning your TCP/IP configuration, see Book 3, Chapter 1.

Choosing workgroups or domains

A *domain* is a method of placing user accounts and various network resources under the control of a single directory database. Domains ensure that security policies are applied consistently throughout a network and greatly simplify the task of managing user accounts on large networks.

A *workgroup* is a simple association of computers on a network that makes it easy to locate shared files and printers. Workgroups don't have sophisticated directory databases, so they can't enforce strict security.

Workgroups should be used only for very small networks with just a few users. Truthfully, any network that is large enough to have a server running Windows Server 2025 is too large to use workgroups. So, if you're installing a Windows Server, you should always opt for domains.

After you decide to use domains, you have to make two basic decisions:

- **What will the domain name be?** If you have a registered internet domain name, such as `mydomain.com`, you may want to use it for your network's domain name. Otherwise, you can make up any name you want.
- **What computer or computers will be the domain controllers for the domain?** If this server is the first server in a domain, you must designate it as a domain controller. If you already have a server acting as a domain controller, you can either add this computer as an additional domain controller or designate it a member server.


TIP

You can always change the role of a server from a domain controller to a member server, and vice versa, if the needs of your network change. If your network has more than one server, it's always a good idea to create at least two domain controllers. That way, if one fails, the other one can take over.

Before You Install

After you've made the key planning decisions for your Windows Server installation, but before you actually start the Setup program, you should take a few precautionary steps. The following sections describe what you should do before you perform an upgrade installation.

Note: The first three steps apply only to upgrades. If you're installing a Windows Server on a new system, you can skip those steps.

Backing up

Do a complete backup of the server before you begin. Although Windows Setup is reliable, sometimes, something serious goes wrong, and data is lost.


TIP

You don't have to back up the drive to external media, such as tape. If you can find a network disk share with enough free space, back up to it.

Checking the event logs

Look at the event logs of the existing server computer to check for recurring errors. You may discover that you have a problem with a SCSI device or your current TCP/IP configuration. It's better to find out now rather than in the middle of setup.

Applying updates

Make sure that the server you're upgrading to Windows Server 2025 is current with all the most recent Windows updates. Click Start, search for "Updates," and select Check for Updates. Then apply any updates that are pending.

Disconnecting UPS devices

If you've installed an uninterruptible power supply (UPS) device on the server and connected it to your computer via a USB cable, you should temporarily disconnect the serial cable before you run Setup. When Setup is complete, you can reconnect the UPS device.

Running Setup

Now that you've planned your installation and prepared the computer, you're ready to run the Setup program. The following procedure describes the steps that you must follow to install Windows Server 2025 on a virtual machine using an ISO image of the installation DVD. Before you begin this procedure, you'll need to download the ISO file from Microsoft's website. Search for "Windows Server 2025 Download" to find the download site.

1. **Configure the new virtual machine with the specifications you want to use, mount the installation ISO file on the virtual DVD drive, and start the VM.**

   The Setup program spends a few moments loading files. Then it asks you to select your language, as shown in Figure 1-1.

2. **Select the correct language, and then click Next.**

   You're asked to choose a keyboard format, as shown in Figure 1-2.

3. **Select your keyboard format, and then click Next.**

   You're presented with two setup options, as shown in Figure 1-3. Here, you can install a new instance of Windows Server, or you can repair an existing installation of Windows Server.

4. **Select Install Windows Server, and then click Next.**

   Now you must enter a product key, as shown in Figure 1-4.

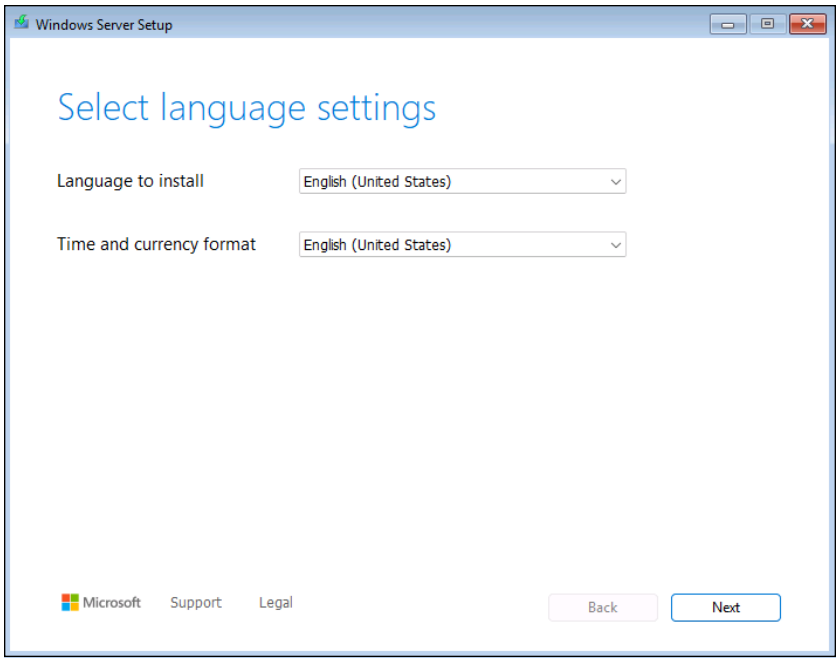


FIGURE 1-1: Select your language.

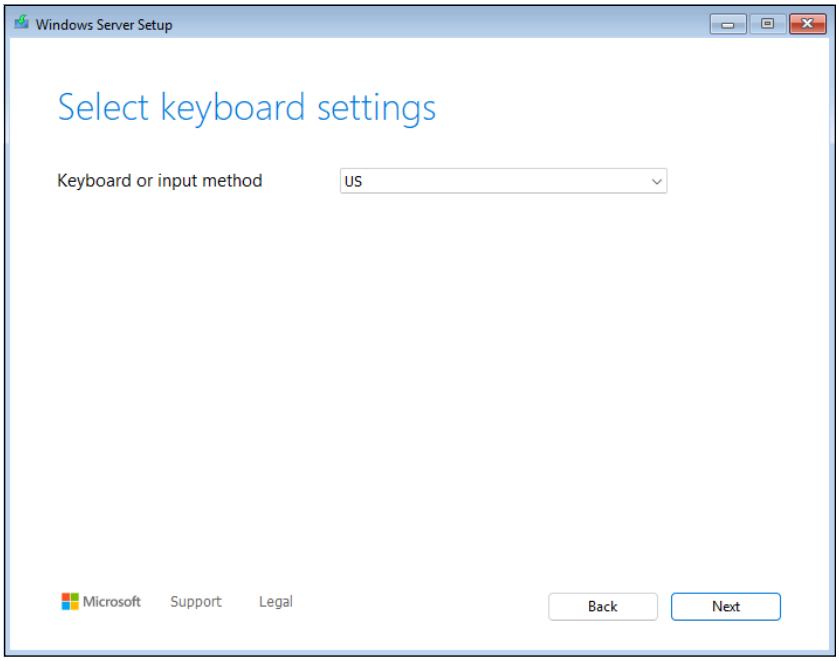


FIGURE 1-2: Select your desired keyboard format.

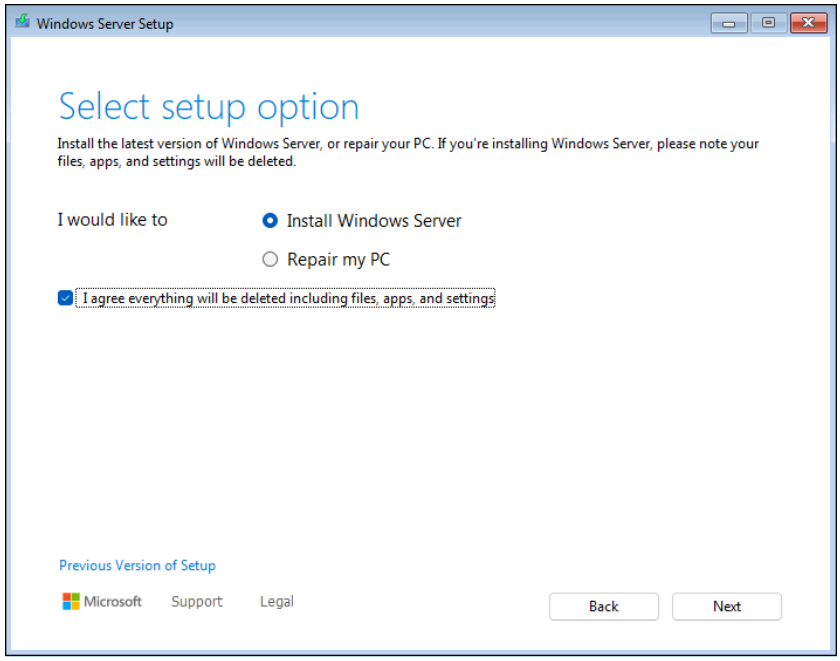


FIGURE 1-3: Do you want to install a new instance of Windows Server or repair an existing one?

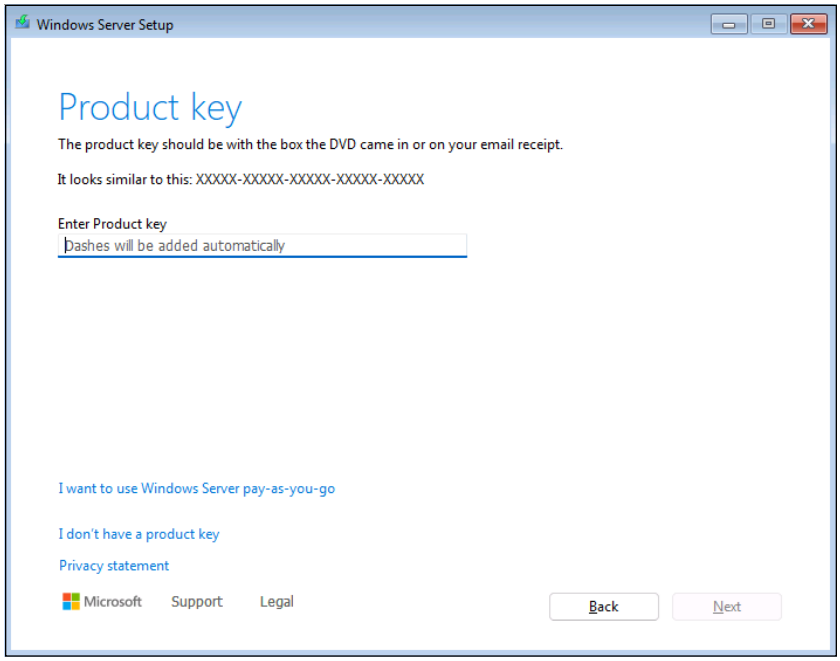


FIGURE 1-4: Enter a product key.

5. **Enter your product key, and then click Next.**

   As shown in Figure 1-5, you can choose whether you want to install Windows Server 2025 Standard, which does not have a graphical user interface (GUI), or Windows Server 2025 Standard (Desktop Experience), which has a GUI similar to desktop Windows. I recommend you choose the Desktop Experience.

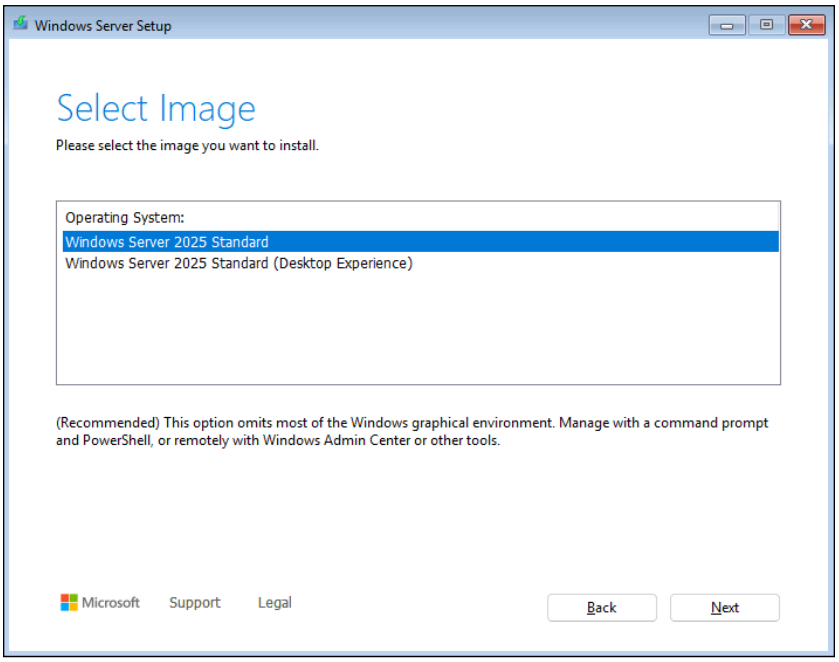


FIGURE 1-5: Choosing the Desktop Experience or the command-line-only version.

6. **Select the Desktop Experience option, and then click Next.**

   Now it's time to accept the license. Try to stay awake.

7. **Read the license agreement, and then click Next.**

   The Setup program now asks where you would like to install the operating system, as shown in Figure 1-6. You can accept the default, or you can use the various options on this page to create a new partition to hold Windows Server.

8. **Select the disk location for the operating system, and then click Next.**

   At last, the Setup Program is ready to install Windows. A final confirmation screen appears, awaiting your consent.

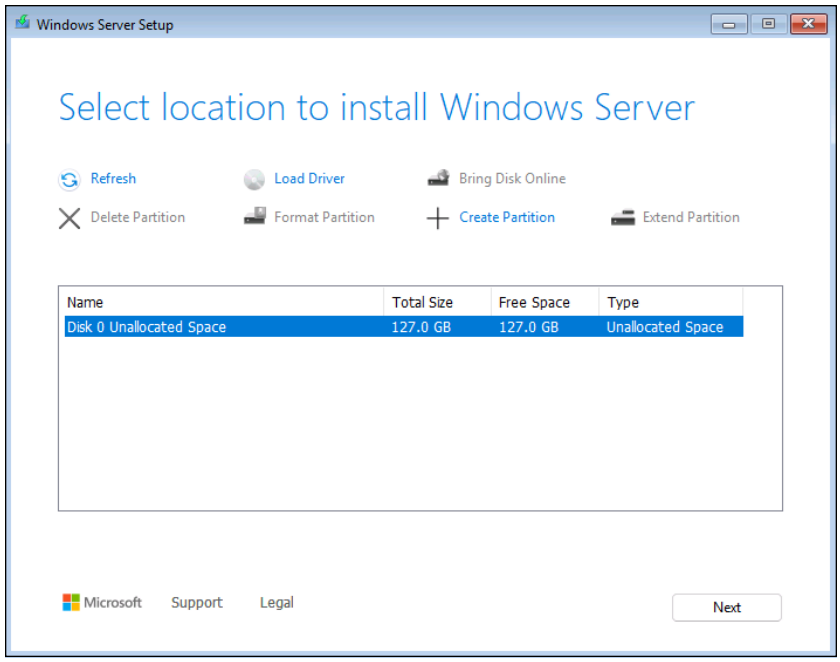


FIGURE 1-6: Choosing the operating system disk location.

9. **Take a deep breath, and then click Install.**

   Windows begins the installation. While it installs, you'll see a progress screen that gives you a hint of how much time remains. Note that the installation may take more than a few minutes and your computer may (or may not) reboot a few times during the installation.

10. **Take a walk, get lunch, or read a book. Then come back.**

    When the installation finally completes, Windows will prompt you for the password to the server's local Administrator account, as shown in Figure 1-7. Pick a good one!

11. **Enter the Administrator account password, and then click Finish.**

    Windows wraps things up. Then it displays its familiar login prompt, shown in Figure 1-8.

12. **Use the Administrator's password to log in to your new server.**

    Congratulations! You've successfully installed Windows Server 2025.

    After you successfully log in, Windows displays the helpful Server Manager Dashboard, shown in Figure 1-9.

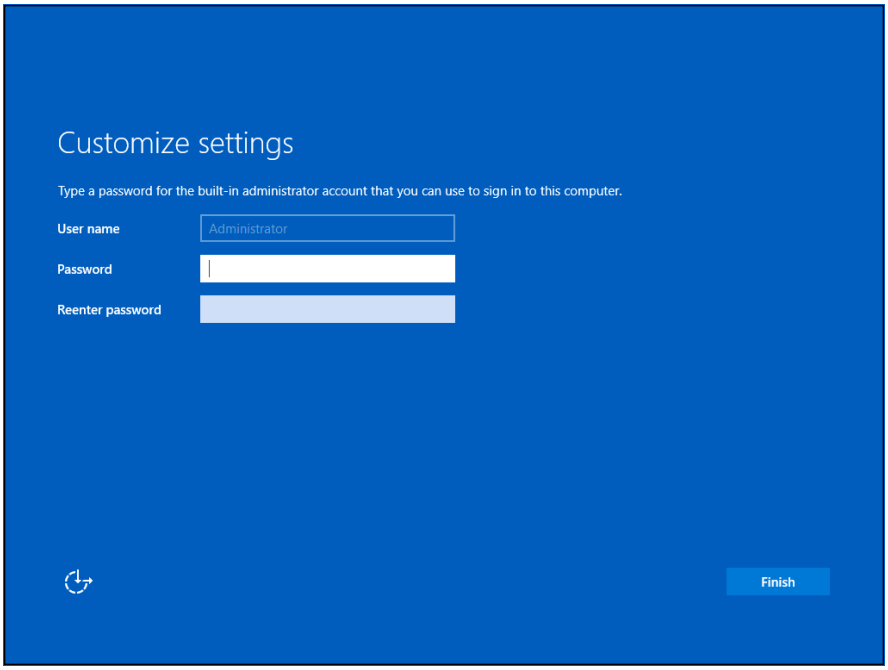


FIGURE 1-7: Choose a good password for the server's Administrator account.

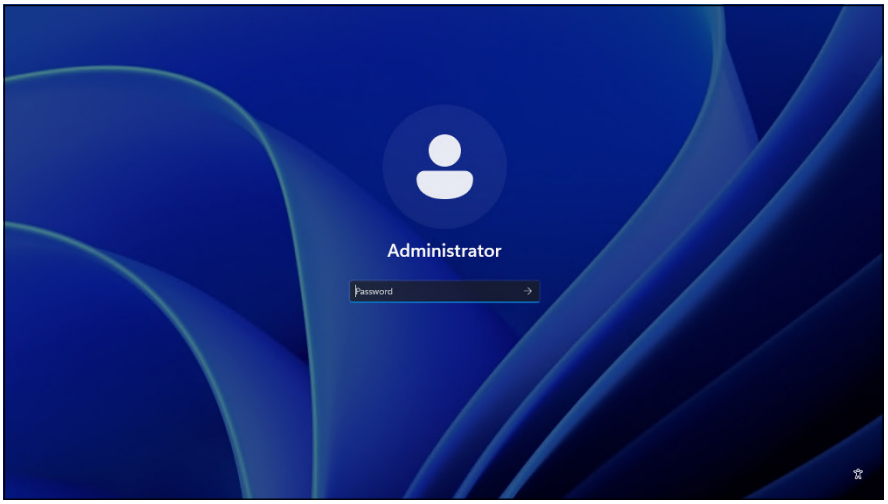


FIGURE 1-8: It's time to log in!

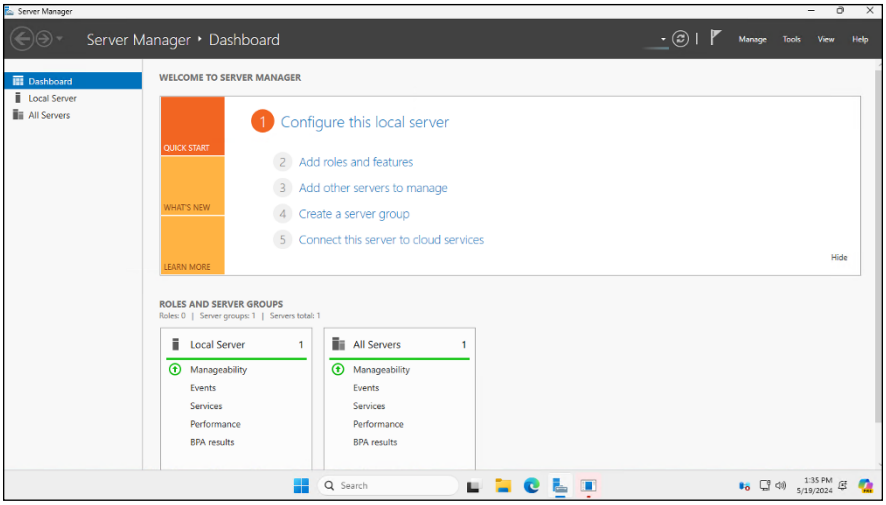


FIGURE 1-9: Behold! The Server Manager Dashboard appears!

The Server Manager Dashboard provides links that let you complete the configuration of your server. Specifically, you can

- Click Configure This Local Server to configure server settings such as the computer's name and the domain it belongs to, network settings such as the static Internet Protocol (IP) address, and so on.
- Click Add Roles and Features to add server roles and features. (For more information, see "Adding Server Roles and Features," later in this chapter.)
- Click Add Other Servers to Manage to manage other servers in your network.
- Click Create a Server Group to create a customized group of servers.
- Click Connect This Server to Cloud Services if you integrate cloud services with your server.

Considering Your Next Steps

After you've finished installing Windows Server, it's time to consider the next steps as you begin to configure the new server for its ultimate role in your organization. I suggest you take the following steps before doing anything else:

1. **Rename the server.**

   By default, the server will have a random name, such as WIN-I42BRD8HT8F. That's probably not going to be the name you want the server known by on

your network. Open File Explorer, right-click This PC, and choose Properties; then change the name to something more meaningful. (This requires a reboot.)

2. **Join it to your domain.**

   While you're at it, you'll want to join the new server to your domain. You can do this at the same time you rename the server. (This, too, requires a reboot.)

3. **Apply updates.**

   The setup image you installed the operating system from is probably not up to date with all the current updates, so be sure to run Windows Update to apply all critical updates. (This may require one or more reboots, depending on the updates.)

4. **Deploy your antivirus solution.**

   Don't let the server (or any other computer) run on your network for more than a few minutes without installing antivirus software. (This may require a reboot, depending on your antivirus software.)

5. **Assign a static IP.**

   By default, new computers use DHCP. For a server, you'll almost always want to assign a static IP (see Book 2, Chapter 5).

Adding Server Roles and Features

Server roles are the roles that your server can play on your network — roles such as file server, web server, or DHCP or DNS server. *Features* are additional capabilities of the Windows operating system itself, such as the .NET Framework or Windows Backup. Truthfully, the distinctions between roles and features are a bit arbitrary. The web server is considered to be a role, for example, but the Telnet server is a feature. Go figure.

The following procedure describes how to install server roles. The procedure for installing server features is similar.

1. **Click Add Roles and Features on the Server Manager Dashboard.**

   The Add Roles and Features Wizard, shown in Figure 1-10, appears.

2. **Click Next.**

   The wizard asks which of two installation types you want to perform. In most cases, you want to leave the default choice (Role-Based or Feature-Based Installation) selected. Select the alternative (Remote Desktop Services Installation) only if you're configuring a remote virtual server.

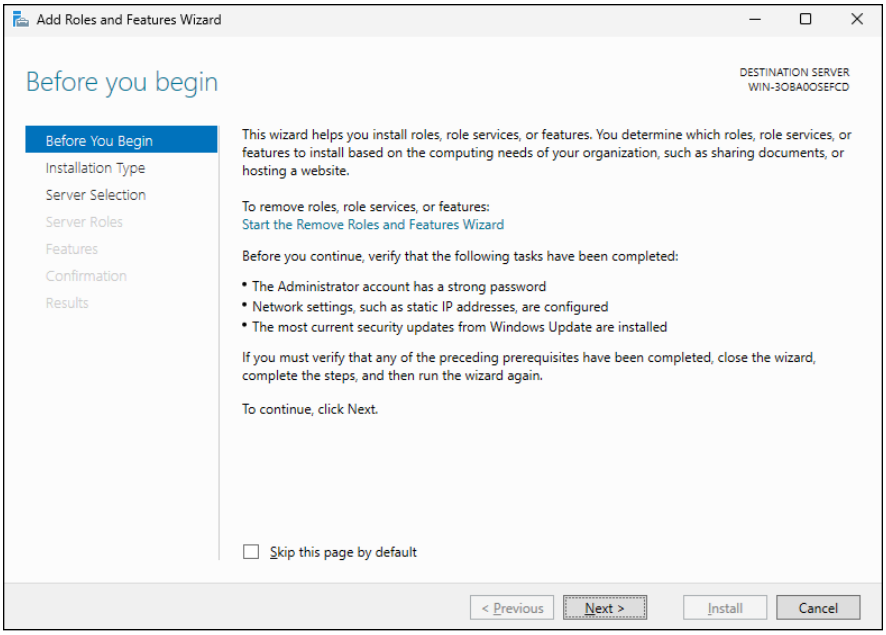


FIGURE 1-10: The Add Roles and Features Wizard.

3. **Click Next.**

   The wizard lets you select the server you want to install roles or features for, as shown in Figure 1-11. In this example, only one server is listed. If you'd chosen Add Other Servers to Manage in the Server Manager Dashboard to add other servers, those servers would appear on this screen as well.

4. **Select the server you want to manage and then click Next.**

   The Select Server Roles screen, shown in Figure 1-12, appears. This screen lets you select one or more roles to add to your server.

5. **Select one or more roles to install.**

   You can click each role to display a brief description of it. If you click DHCP Server, for example, the following text is displayed:

   - `Dynamic Host Configuration Protocol (DHCP) Server enables you to centrally configure, manage, and provide temporary IP addresses and related information for client computers.`

6. **Click Next.**

   The Select Features screen appears, as shown in Figure 1-13. This screen lists additional server features that you can install.

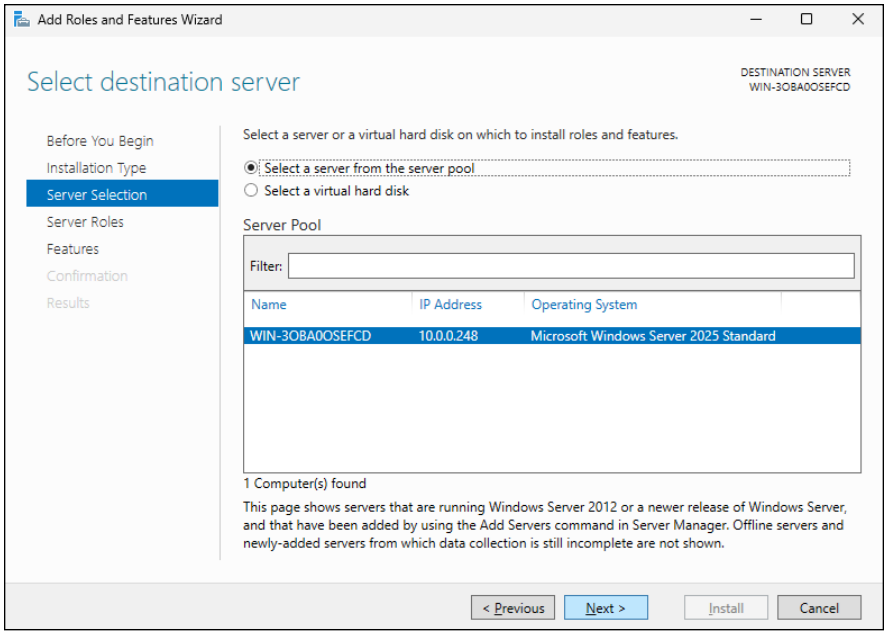


FIGURE 1-11: Selecting the server to manage.

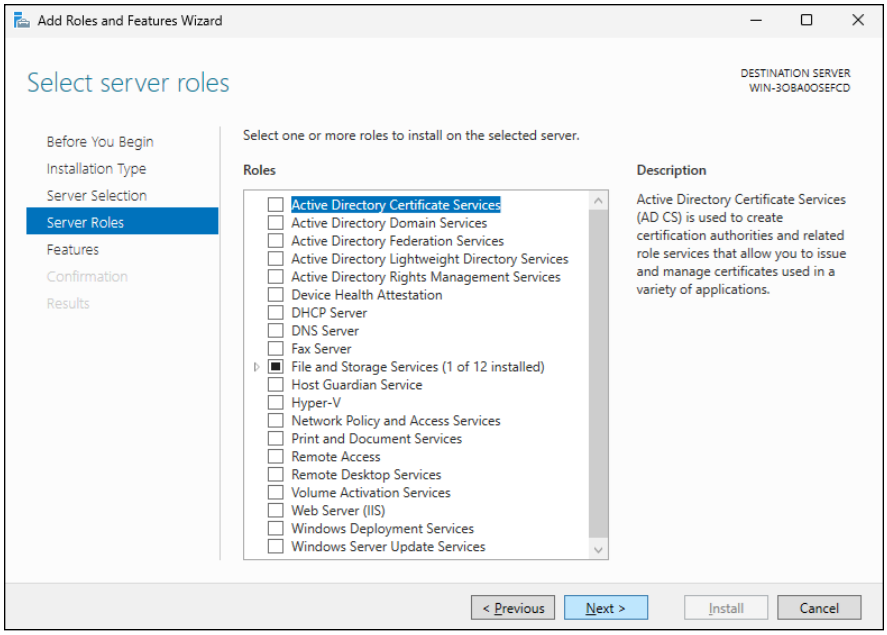


FIGURE 1-12: The Select Server Roles screen.

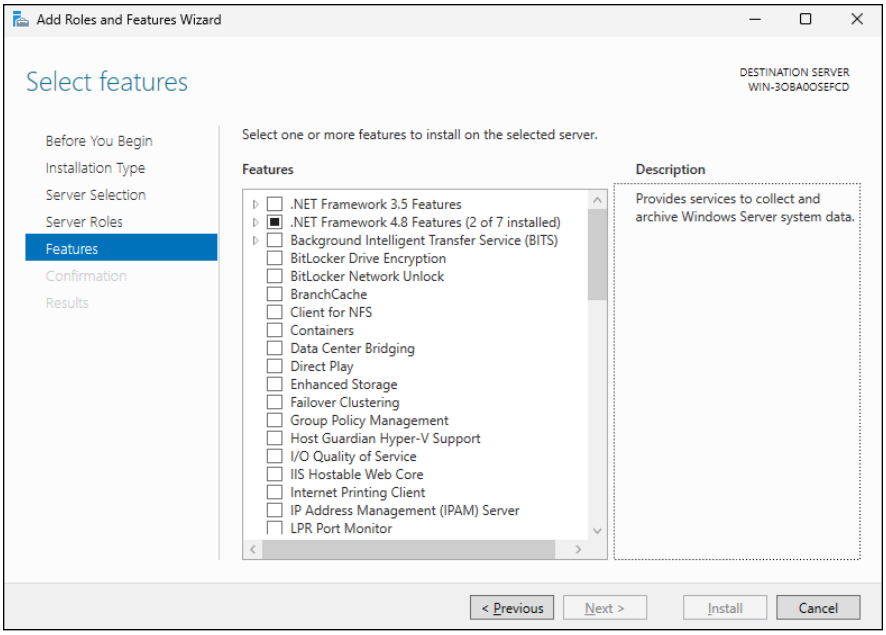


FIGURE 1-13: The Select Features screen.

7. **Select the features you want to install.**

   Again, you can select each feature to see a brief text description of the service.

8. **Click Next.**

   A confirmation screen appears, listing the roles and features you've selected.

9. **Click Install.**

   Windows installs the server role and its features. A progress screen is displayed during the installation so that you can gauge the installation's progress. When the installation finishes, a final results screen is displayed.

10. **Click OK.**

    You're done!

IN THIS CHAPTER

- Working with the trusty Administrator account
- Using Remote Desktop Connection to administer a server from the comfort of your desktop
- Perusing the list of Microsoft Management Consoles
- Customizing your own management console

Chapter 2
Configuring Windows Server 2025

This chapter provides an introduction to the most important tools that you'll use to administer Windows Server 2025.

Using the Administrator Account

Windows comes with a built-in account named *Administrator* that has complete access to all the features of the server. As a network administrator, you frequently log on using the Administrator account to perform maintenance chores.

Because the Administrator account is so powerful, you should always enforce good password practices for it. In other words, don't use your dog's name as the Administrator account password. Instead, pick a random combination of letters and numbers. Then change the password periodically.



REMEMBER

Do not write down the Administrator account password. Instead, keep it in a secure location such as a password manager. If you must write it down on paper, keep the *only* copy in a bank safe-deposit box, with access limited to an absolute minimum number of trusted souls.

Note that you cannot delete or disable the Administrator account. If Windows allowed you to do that, you could potentially find yourself locked out of your own system.



TIP

As much as possible, you should avoid using the Administrator account. Instead, you should create accounts for each of your system administrators and grant them administrator privileges by assigning their accounts to the Administrators group.

Although you can't delete or disable the Administrator account, you can rename it. Some network managers use this ability to hide the true Administrator account. To do this, just follow these steps:

1. **Rename the Administrator account.**

   Write down the new name you use for the Administrator account, along with the password, and store it in a top-secret secure location.

2. **Create a new account named Administrator, and assign it a strong password, but don't give this account any significant privileges.**

   This new account will become a "decoy" Administrator account. The idea is to get hackers to waste time trying to crack this account's password. Even if a hacker does manage to compromise this account, he won't be able to do anything when he gets in.

Using Remote Desktop Connection

One of the most useful tools available to system administrators is a program called *Remote Desktop Connection*, or *RDC* for short. RDC lets you connect to a server computer from your own computer and use it as though you were actually sitting at the server. In short, RDC lets you administer your server computers from your own office.

Enabling remote access

Before you can use Remote Desktop Connection to access a server, you must enable remote access on the server. To do that, follow these steps (on the server computer, not your desktop computer):

1. **Open the Control Panel and then click System.**

   The System settings page appears.

2. **Click the Advanced System Settings link.**

   The Systems Properties dialog box appears.

3. **Click the Remote tab.**

   The remote access options appear, as shown in Figure 2-1.

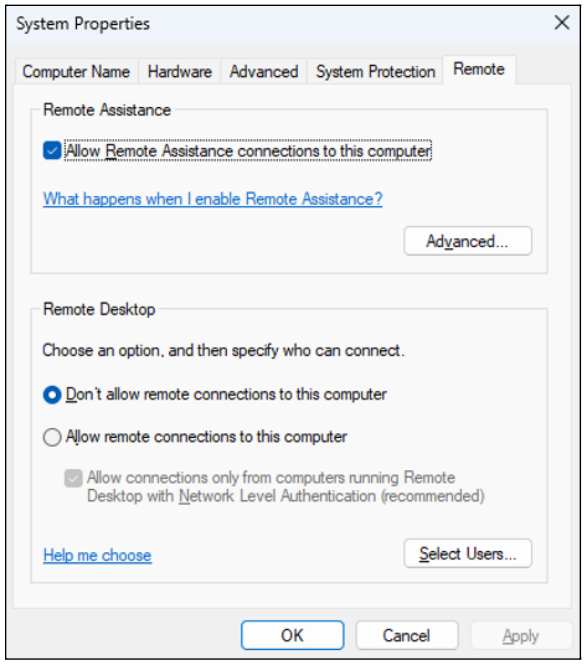


FIGURE 2-1: Configuring remote access.

4. **Select the Allow Remote Connections to This Computer radio button.**
5. **Click OK.**

You're done! Repeat this procedure for each server computer you want to allow access to.

Here are a few other points to ponder concerning remote access:

- You can click the Select Users button to create a list of users who are authorized to access the computer remotely. Note that all members of the Administrators group are automatically granted access, so you don't have to add administrators to this list.



WARNING

» There's no question that RDC is convenient and useful. And it operates over secure, encrypted channels. However, after you enable RDC, it's even more important that you take adequate precautions to secure your Administrator accounts by using strong passwords. Also, you should already have a firewall installed to keep unwanted visitors out of your network. For more information on account security, see Book 9, Chapter 1.

Connecting remotely

After you've enabled remote access on a server, you can connect to the server by using the Remote Desktop Client that's automatically installed with Windows. Here's the procedure:

1. **On a desktop computer, click Start, type** Remote**, and then choose Remote Desktop Connection.**

   The Remote Desktop Connection client comes to life, as shown in Figure 2-2.

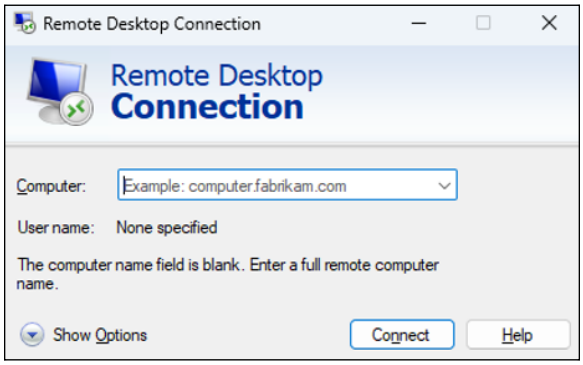


FIGURE 2-2: Connecting with Remote Desktop Connection.

2. **Enter the name of the computer you want to connect to.**

   Alternatively, you can use the drop-down list to select the computer from the list of available computers.

3. **Click the Connect button.**

   You're connected to the computer you selected, and the computer's logon screen is displayed.

4. **Log on, and use the computer.**

   After you log on, you can use the computer as though you were sitting right in front of it.

Here are a few other tips for working with the Remote Desktop Connection client:


TIP

- When you're using the Remote Desktop Connection client, you can't just Alt+Tab to another program running on the client computer. Instead, you must first minimize the RDC client's window by clicking its minimize button. Then you can access other programs running on your computer.
- If you minimize the RDC client window, you have to provide your logon credentials again when you return. This security feature is there in case you forget that you have an RDC session open.
- If you use RDC a lot on a particular computer (such as your own desktop computer), I suggest that you create a shortcut to RDC and place it on the desktop or pin the Remote Desktop Connection program to your Start menu or to your taskbar (or both).
- RDC has several useful configuration options that you can access by clicking the Options button.

Using Microsoft Management Console

Microsoft Management Console, also known as *MMC*, is a general-purpose management tool that's used to administer many different types of objects on a Windows system. Throughout this minibook, you see many examples of MMC for working with objects such as user accounts, disk drives, event logs, and so on. This section provides a general overview of how to use MMC.

By itself, MMC doesn't actually manage anything. Instead, it's a framework that accepts management snap-ins, which do the actual managing. The main point of MMC is that it provides a consistent framework for building management snap-ins so that all the snap-ins behave in similar ways. As a result, you don't have to struggle to learn completely different tools to manage various aspects of Windows Server 2025.

Another advantage of MMC is that you can create your own custom management consoles with just the right combination of snap-ins. Suppose that you spend most of your time managing user accounts, disk devices, and Internet Information Services (IIS, the web server that comes with Windows Server 2025), and studying event logs. You can easily craft a management console with just these four snap-ins. For more information, see the section "Customizing MMC," later in this chapter.

Working with MMC

There are several ways to open a Microsoft Management Console window. The easiest is to open one of the predefined consoles that come with Windows Server 2025. To access these consoles, press the Windows key and then select Administrative Tools.

You can also start MMC from a command prompt. To start MMC without opening a snap-in, just type **mmc** at a command prompt. To open a specific console, type the path to the console file after `mmc`. The following command, for example, opens the Computer Management console:

```
mmc \Windows\System32\compmgmt.msc
```

Figure 2-3 shows a typical Microsoft Management Console window, displaying the Active Directory Users and Computers snap-in. As you can see, the MMC window consists of two panes. The pane on the left is a tree pane that displays a hierarchical tree of the objects that you can manage. The pane on the right is a Details pane that shows detailed information about the object that's selected in the tree pane.

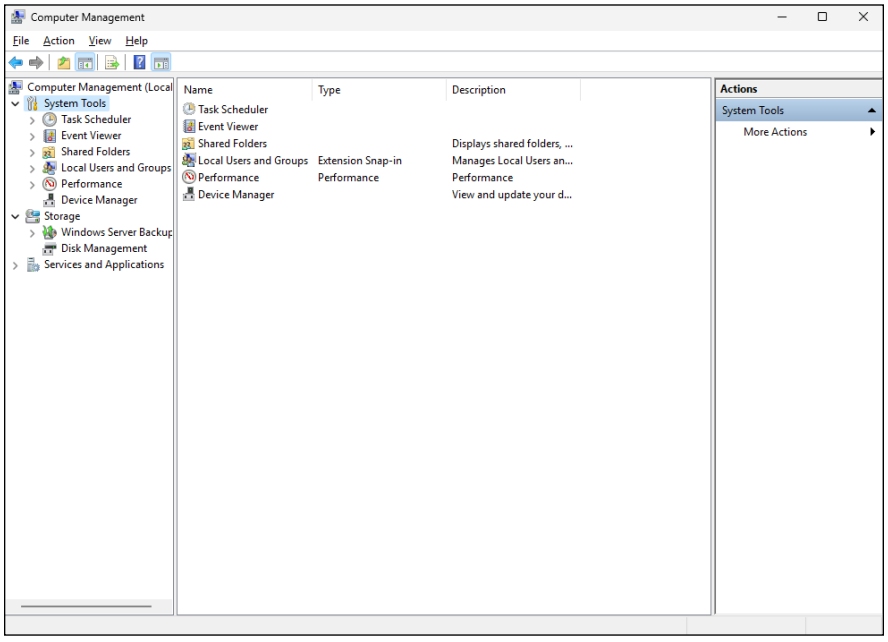


FIGURE 2-3: A typical MMC window.

The procedures for working with the information in the Details pane vary depending on the console you're viewing. Most of the consoles, however, display lists of some kind, such as settings or user accounts. Double-clicking an item usually brings up a Properties dialog box that lets you view or set properties for the object. In most cases, you can click the column headings at the top of the list to change the order in which the list items are displayed.

MMC also includes a menu and toolbar with commands and buttons that vary depending on the item selected in the tree. In particular, the Action menu contains commands that apply to the current item. The Action menu includes a New User command when you're working with the Active Directory Users and Computers console, for example, and a Defragment command when you view the Disk Defragmenter item in the Computer Management Console. As you work with different items within the different consoles, be sure to check the Action menu frequently to find out what commands are available.

Taking an overview of the MMC consoles

The Tools menu in Server Manager Dashboard gives you direct access to many useful management consoles. You find detailed descriptions of several of these tools later in this minibook. The following paragraphs give you a brief overview of the most important of these consoles:

- **Active Directory Domains and Trusts:** Manages the domain and trust relationships for the server.
- **Active Directory Sites and Services:** Manages Active Directory services.
- **Active Directory Users and Computers:** Lets you create and modify user accounts.
- **Component Services:** Lets you manage how COM+ (Component Object Model) services work on the server. You mess with this console only if you're involved in developing applications that use COM+ services.
- **Computer Management:** Provides access to several useful tools for managing a server. In particular, the Computer Management console provides the following management tools:
  - *Event Viewer:* Lets you view event logs.
  - *Shared Folders*: Lets you manage shared folders for a file server. In addition to finding out what shares are available, you can use this tool to find out which users are connected to the server and which files are open.

- *Local Users and Groups* (available only on servers that aren't domain controllers): Lets you manage local user and group accounts. For a domain controller, you use the Active Directory Users and Computers console to manage user accounts.
- *Performance*: Lets you monitor system performance counters.
- *Device Manager*: Lets you manage the hardware devices connected to a server. You'll probably use it only if you're having a problem with the server that you suspect may be hardware-related.
- *Disk Management*: Lets you view the physical disks and volumes that are available to the system. You can also use this tool to create and delete partitions, set up RAID volumes, format disks, and so on.
- *Services:* Lets you manage system services. You can use this tool to start or stop services such as Exchange email services, TCP/IP services such as DNS and DHCP, and so on.
- *WMI Control:* Lets you configure *Windows Management Instrumentation services,* which track management data about computers, users, applications, and other objects in large Enterprise networks.

- **DHCP:** Manages the DHCP server.
- **DNS:** Manages the DNS server.
- **Domain Controller Security Policy:** Lets you set security policy for a domain controller.
- **Event Viewer:** Lets you view event logs.
- **Group Policy Management:** Lets you set system policies that can be applied to objects such as users and groups.
- **IIS Manager:** Lets you manage the services provided by IIS (Microsoft's web server) if IIS is installed on the server.
- **ODBC Data Sources:** Manages database connections that use ODBC. You'll probably use this console only if you're a developer or database administrator.
- **Performance Monitor:** Lets you monitor a server's performance and twiddle with various settings that can have positive or negative effects on performance.
- **Services:** Lets you start and stop Windows services. (It's also available via the Computer Management console.)

Customizing MMC

One of the best things about Microsoft Management Console is that you can customize it so that the tools you use most often are grouped together in whatever combination you choose. To create a custom console, first start Microsoft Management Console without loading a console by pressing the Windows key, typing **cmd** and pressing Enter to open a command prompt, and then entering the command **mmc**. This action creates an empty console, as shown in Figure 2-4.

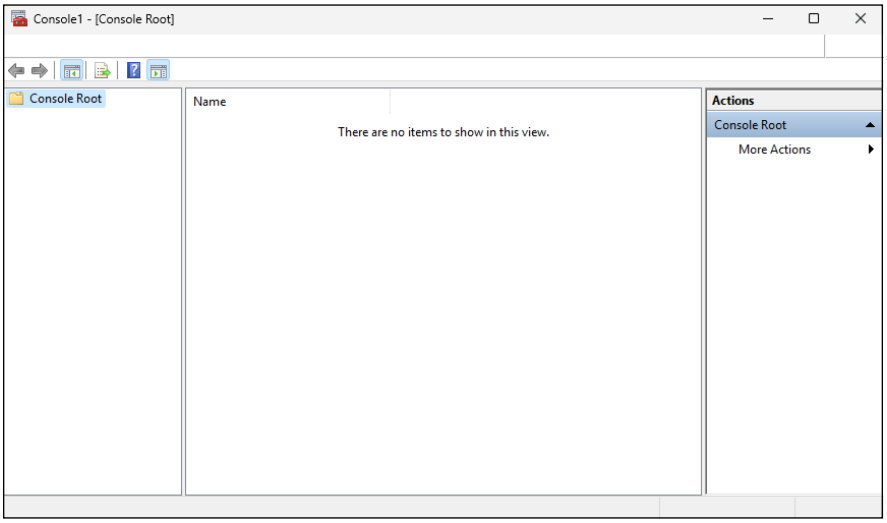


FIGURE 2-4: An empty MMC console.

After you've created an empty console, you can customize it by adding whatever snap-ins you want to make use of in the console. To add a snap-in, follow these steps:

1. **Choose File➪Add/Remove Snap-in.**

   This command brings up the Add or Remove Snap-ins dialog box, shown in Figure 2-5.

2. **Select the snap-in you want to add and then click the Add button.**

   Depending on which snap-in you select, a dialog box appears, asking whether you want to use the add-in to manage settings on your own computer or on a local computer.

3. **Repeat Step 2 if you want to add other snap-ins to the console.**

4. **Click OK.**

   The console is equipped with the snap-ins you've selected.

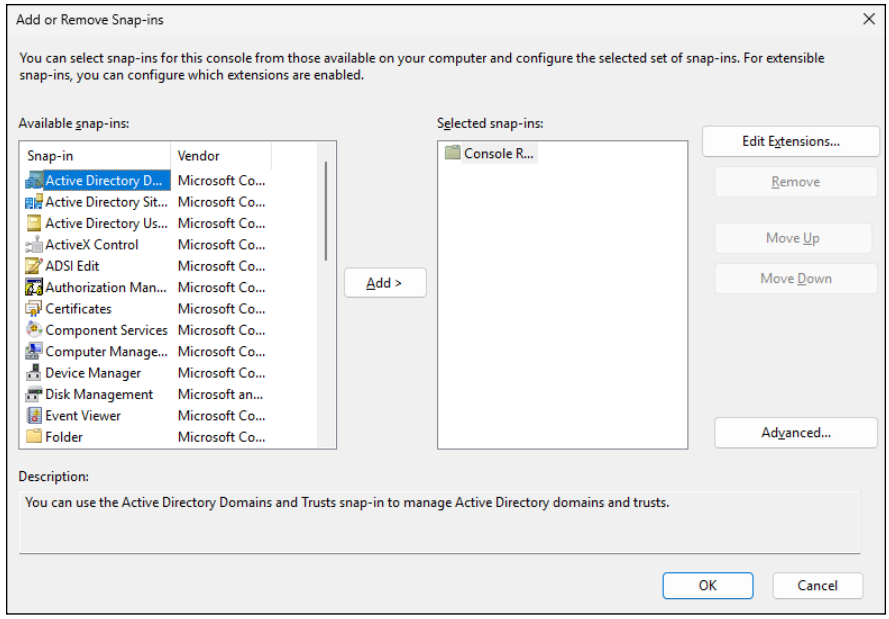


FIGURE 2-5: The Add or Remove Snap-ins dialog box.

IN THIS CHAPTER

- Discovering directories
- Examining how Active Directory is structured
- Setting up a domain controller
- Creating organizational units

Chapter 3
Configuring Active Directory

Active Directory is among the most important features of Windows Server, and much of your time as a network administrator will be spent keeping Active Directory neat and tidy. This chapter lays some foundation by explaining what Active Directory is and how it works.

What Active Directory Does

Everyone uses directory services of one type or another every day. When you look up someone's name in a phone book, you're using a directory service. But you're also using a directory service when you make a call: When you enter someone's phone number into your touch-tone phone, the phone system looks up that number in its directory to locate that person's phone.

Almost from the very beginning, computers have had directory services. When I got started in the computer business back in the 1970s, I used IBM mainframe computers and a transaction-processing system called CICS that's still in widespread use today. CICS relied on many different directories to track such things as files available to the system, users that were authorized to access the system, and application programs that could be run.

But the problem with this directory system, and with most other directory systems that were popular in those days, is that it was made up of many small directory systems that didn't know how to talk to one another. I have the very same problem at home. I have my own little personal address book that has phone numbers and addresses for my friends and family members. I have a Day-Timer book with a bunch of other phone numbers and addresses. Then I have a church directory that lists everyone who goes to my church. Oh, and there's the list of players on the softball team I coach, and of course, my cellphone has a directory.

All counted, I probably have a dozen sources for phone numbers that I routinely call. So when I need to look up someone's phone number, I first have to decide which directory to look in. Some of my friends are listed in two or three of these sources, which raises the possibility that their listings are out of sync.

That's exactly the type of problem that Active Directory is designed to address. *Active Directory* is a comprehensive directory management system that tracks just about everything worth tracking in a Windows network, including users, computers, files, folders, applications, and much more. Much of your job as a network administrator involves working with Active Directory, so it's vital that you have a basic understanding of how it works.

Note that Active Directory uses the same naming scheme that's used on the internet: Domain Name System (DNS). Thus, an Active Directory domain might have a name like `sales.mycompany.com`.

Understanding How Active Directory Is Structured

Like all directories, Active Directory is essentially a database management system. The Active Directory database is where the individual objects tracked by the directory are stored. Active Directory uses a *hierarchical* database model, which groups items in a treelike structure.

The terms *object, organizational unit, domain, tree,* and *forest* are used to describe the way Active Directory organizes its data. The following sections explain the meaning of these important Active Directory terms.

Objects

The basic unit of data in Active Directory is called an *object.* Active Directory can store information about many kinds of objects. The objects you work with most are users, groups, computers, and printers.

Figure 3-1 shows the Active Directory Manager displaying a list of built-in objects that come preconfigured with Windows Server 2025. To get to this management tool, choose Start➪Administrative Tools➪Active Directory Users and Computers. Then click the Builtin node to show the built-in objects.

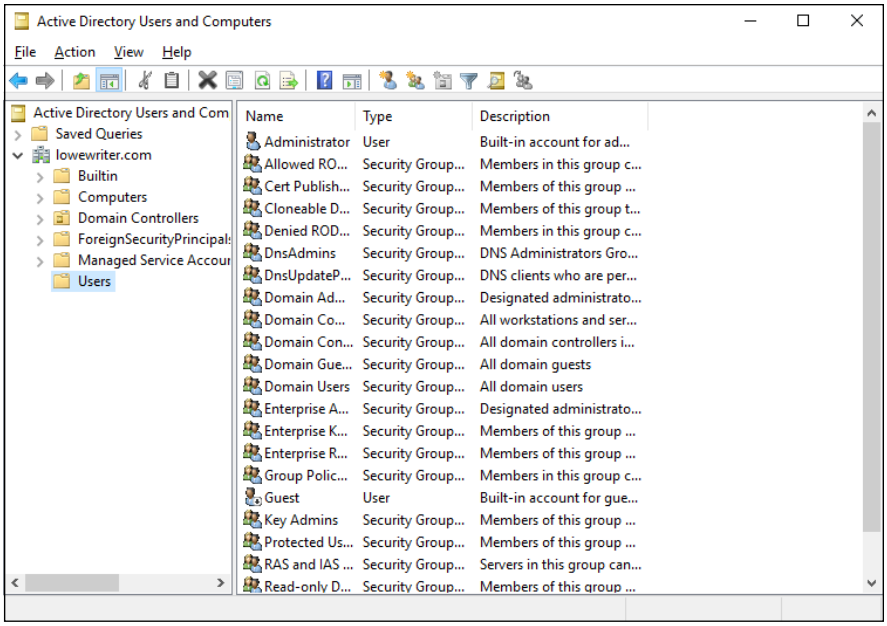


FIGURE 3-1: Objects displayed by the Active Directory Manager console.

Objects have descriptive characteristics called *properties* or *attributes.* You can call up the properties of an object by double-clicking the object in the management console.

Domains

A *domain* is the basic unit for grouping related objects in Active Directory. Typically, domains correspond to departments in a company. A company with separate Accounting, Manufacturing, and Sales departments might have domains named (you guessed it) `Accounting`, `Manufacturing`, and `Sales`. Or the domains may correspond to geographical locations. A company with offices in Detroit, Dallas, and Denver might have domains named `det`, `dal`, and `den`.

Note that because Active Directory domains use DNS naming conventions, you can create subdomains that are considered to be child domains. You should always create the top-level domain for your entire network before you create any other domain. If your company is named Nimbus Brooms, and you've registered `nimbusbroom.com` as your domain name, you should create a top-level

domain named `nimbusbroom.com` before you create any other domains. Then you can create subdomains such as `accounting.nimbusbroom.com`, `manufacturing.nimbusbroom.com`, and `sales.nimbusbroom.com`.


TIP

If you have Microsoft Visio, you can use it to draw diagrams for your Active Directory domain structure. Visio includes several templates that provide cool icons for various types of Active Directory objects. Figure 3-2 shows a diagram that shows an Active Directory with four domains created with Visio.

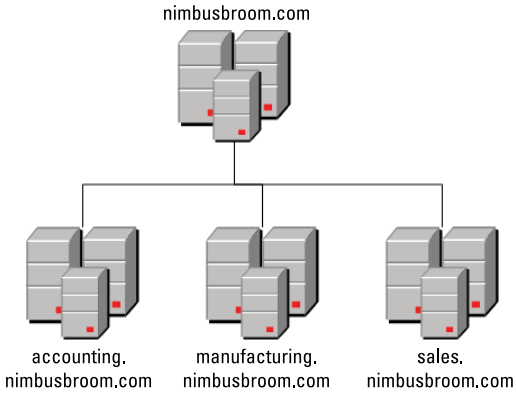


FIGURE 3-2: Domains for a company with three departments.

Note that these domains have little to do with the physical structure of your network. In Windows Server, domains usually are related to the network's physical structure.

Every domain must have at least one *domain controller*, which is a server that's responsible for the domain. Unlike a Windows Server PDC, however, an Active Directory domain controller doesn't have unique authority over its domain. In fact, a domain can have two or more domain controllers that share administrative duties. A feature called *replication* works hard at keeping all the domain controllers in sync.

Organizational units

Many domains have too many objects to manage together in a single group. Fortunately, Active Directory lets you create one or more *organizational units* (OUs). OUs let you organize objects within a domain, without the extra work and inefficiency of creating additional domains.

One reason to create OUs within a domain is to assign administrative rights to each OU of different users. Then these users can perform routine administrative tasks such as creating new user accounts or resetting passwords.

Suppose that the domain for the Denver office, named `den`, houses the Accounting and Legal departments. Rather than create separate domains for these departments, you could create organizational units for the departments.

Trees

A *tree* is a set of Active Directory names that share a namespace. The domains `nimbusbroom.com`, `accounting.nimbusbroom.com`, `manufacturing.nimbusbroom.com`, and `sales.nimbusbroom.com` make up a tree that's derived from a common root domain, `nimbusbroom.com`.

The domains that make up a tree are related to one another through *transitive trusts.* In a transitive trust, if DomainA trusts DomainB and DomainB trusts DomainC, DomainA automatically trusts DomainC.



TIP

Note that a single domain all by itself is still considered to be a tree.

Forests

As its name suggests, a *forest* is a collection of trees. In other words, a forest is a collection of one or more domain trees that do *not* share a common parent domain.

Suppose that Nimbus Brooms acquires Tracorum Technical Enterprises, which already has its own root domain named `tracorumtech.com`, with several subdomains of its own. You can create a forest from these two domain trees so that the domains can trust each other. Figure 3-3 shows this forest.

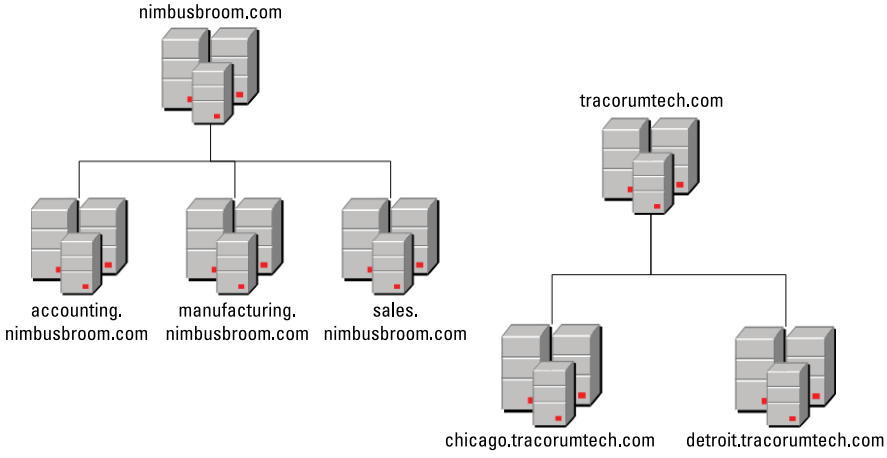


FIGURE 3-3: A forest with two trees.

The key to Active Directory forests is a database called the global catalog. The *global catalog* is sort of a superdirectory that contains information about all the objects in a forest, regardless of the domain. If a user account can't be found in the current domain, the global catalog is searched for the account. The global catalog provides a reference to the domain in which the account is defined.

Creating a New Domain

To create a domain, you start by designating a Windows Server 2025 system to be the new domain's controller. You can do that by using the Server Manager to install the Active Directory Domain Services role. (Refer to Book 6, Chapter 1 for instructions on installing server roles.) After you've installed Active Directory services, click the Notifications icon near the top-right corner of the Server Manager, and choose Promote This Server to a Domain Controller. This command launches the wizard shown in Figure 3-4.

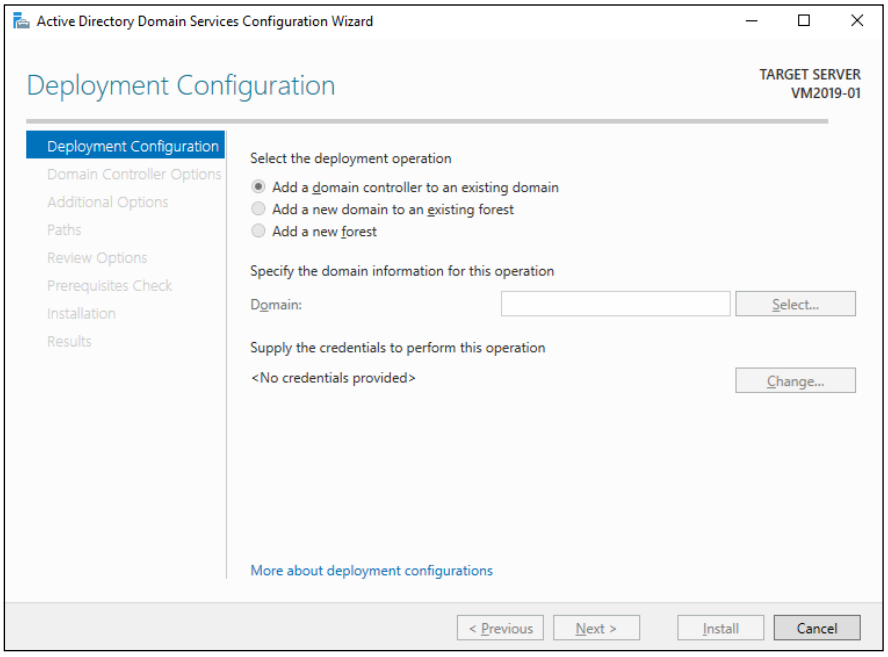


FIGURE 3-4: Creating a domain controller.

This wizard lets you designate the server as a domain controller. As you can see, the wizard gives you three options:

- **Add a Domain Controller to an Existing Domain:** Choose this option if you've already created the domain and want to add this server as a domain controller.
- **Add a New Domain to an Existing Forest:** If you've already created a forest but want to create a new domain within the existing forest, choose this option.
- **Add a New Forest:** This option is the one to choose if you're setting up a new domain in a brand-new forest.

When you create a new domain, the configuration wizard asks you for a name for the new domain. If you're creating the first domain for your network, use your company's domain name, such as `nimbusbroom.com`. If you're creating a subdomain, use a name such as `sales.nimbusbroom.com`.

Creating an Organizational Unit

Organizational units can simplify the task of managing large domains by dividing users, groups, and other objects into manageable collections. By default, Active Directory domains include several useful OUs. The Domain Controllers OU, for example, contains all the domain controllers for the domain.

If you want to create additional organizational units to help manage a domain, follow these steps:

1. **In Server Manager, choose Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers console appears, as shown in Figure 3-5.

2. **Right-click the domain you want to add the OU to, and choose New➪Organizational Unit.**

   The New Object – Organizational Unit dialog box appears, as shown in Figure 3-6.

3. **Type a name for the new organization unit.**

4. **Click OK.**

   You're done!

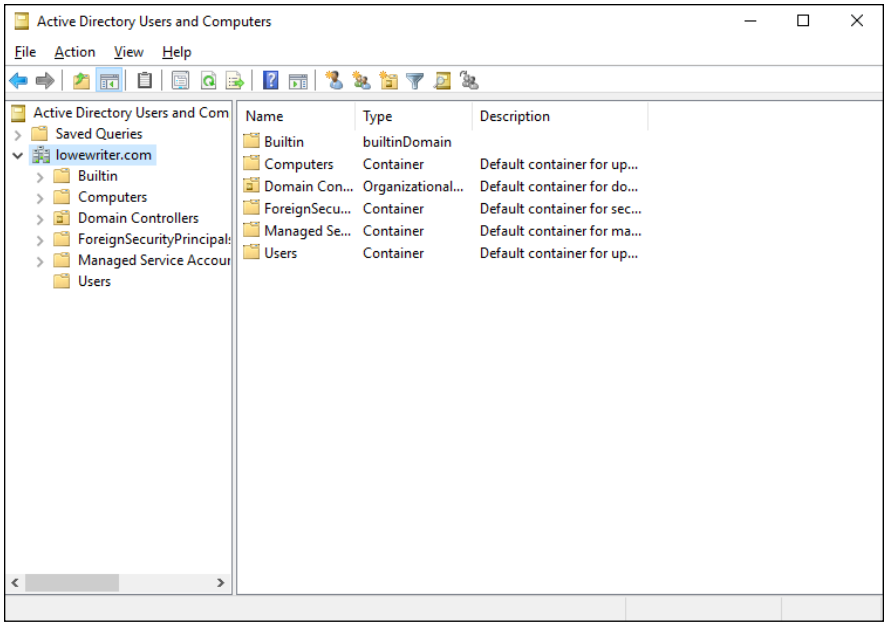


FIGURE 3-5: The Active Directory Users and Computers console.

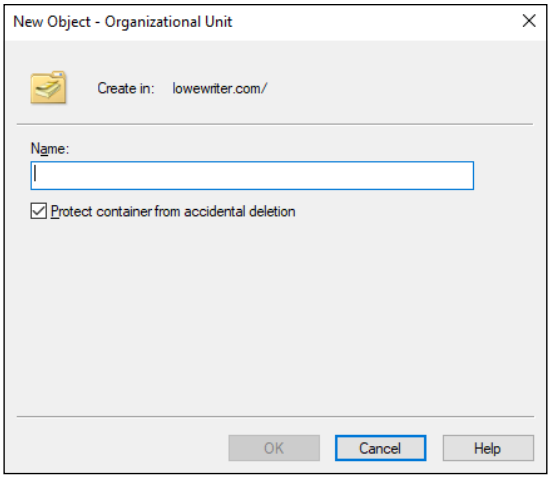


FIGURE 3-6: Creating a new organizational unit.



TIP

Each domain you create in Active Directory has a handful of default OUs that are set up automatically by Active Directory:

- Builtin
- Computers
- Domain Controllers

- ForeignSecurityPrincipals
- Managed Service Accounts
- Users

Two of these OUs — Computers and Users — seem like the natural place for you to store Active Directory records for the computers and users on your domain. I recommend against doing that, though. Instead, you should create your own arrangement of OUs to manage the Active Directory objects you'll create for your network.

I suggest you start by creating a single OU immediately beneath your domain. You can call it whatever you want — use your domain name (without the `.com` or other top-level domain), or use a short nickname for your organization. For example, for the `lowewriter.com` domain, I'll call the OU simply `Lowe`.

Then, within that, create OUs to contain the three most common types of Active Directory objects you'll create: Computers, Groups, and Users. Thus, for my Active Directory, I use the following structure:

```
lowewriter.com
    Lowe
        Computers
        Groups
        Users
```

Within those OUs, you can create additional OUs as needed.

Here are just a few more thoughts about OUs to ponder as you drift off to sleep:

- You can delegate administrative authority for an OU to another user by right-clicking the OU and choosing Select Delegate Control from the contextual menu. Then you can select the user or group that will have administrative authority over the OU. You can also choose which administrative tasks will be assigned to the selected user or group.
- Remember that OUs aren't the same as groups. *Groups* are security principals, which means that you can assign them rights. Thereafter, when you assign a user to a group, the user is given the rights of the group. By contrast, an OU is merely an administrative tool that lets you control how user and group accounts are managed.
- For more information about how to create user and group accounts as well as other Active Directory objects, turn to Book 6, Chapter 4.

IN THIS CHAPTER

- Understanding user accounts
- Creating user accounts
- Setting account options
- Working with groups
- Creating a roaming profile

Chapter 4

Configuring User Accounts

Every user who accesses a network must have a *user account.* User accounts let you control who can access the network and who can't. In addition, user accounts let you specify what network resources each user can use. Without user accounts, all your resources would be open to anyone who casually dropped by your network.

Understanding Windows User Accounts

User accounts are among the basic tools for managing a Windows server. As a network administrator, you'll spend a large percentage of your time dealing with user accounts — creating new ones, deleting expired ones, resetting passwords for forgetful users, granting new access rights, and so on. Before I get into the specific procedures of creating and managing user accounts, this section presents an overview of user accounts and how they work.

Local accounts versus domain accounts

A *local account* is a user account that's stored on a particular computer and applies only to that computer. Typically, each computer on your network will have a local account for each person who uses that computer.

By contrast, a *domain account* is a user account that's stored by Active Directory and can be accessed from any computer that's a part of the domain. Domain accounts are centrally managed. This chapter deals primarily with setting up and maintaining domain accounts.

User account properties

Every user account has several important account properties that specify the characteristics of the account. The three most important account properties are

- **Username:** A unique name that identifies the account. The user must enter the username when logging on to the network. The username is public information. In other words, other network users can (and often should) find out your username.
- **Password:** A secret word that must be entered to gain access to the account. You can set up Windows so that it enforces password policies, such as the minimum length of the password, whether the password must contain a mixture of letters and numerals, and how long the password remains current before the user must change it.
- **Group membership:** The group or groups to which the user account belongs. Group memberships are the key to granting access rights to users so that they can access various network resources, such as file shares or printers, or perform certain network tasks, such as creating new user accounts or backing up the server.

Many other account properties record information about the user, such as the user's contact information, whether the user is allowed to access the system only at certain times or from certain computers, and so on. I describe these features in later sections of this chapter.

Creating a New User

To create a new domain user account in Windows Server 2025, follow these steps:

1. **Choose Start➪Windows Administrative Tools➪Active Directory Users and Computers.**

   This command fires up the Active Directory Users and Computers management console, as shown in Figure 4-1.

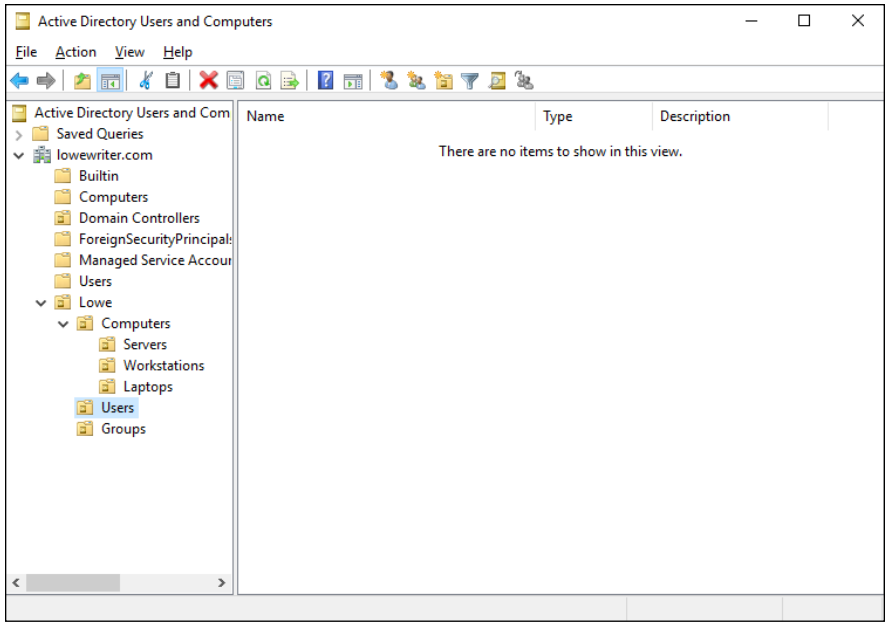


FIGURE 4-1: The Active Directory Users and Computers management console.

2. **Right-click the organizational unit that you want to add the user to and then choose New➪User.**

   This command summons the New Object – User Wizard, as shown in Figure 4-2.

3. **Type the user's first name, middle initial, and last name.**

   As you type the name, the New Object Wizard automatically fills in the Full Name field.

4. **Change the Full Name field if you want it to appear different from what the wizard proposes.**

   You may want to reverse the first and last names so the last name appears first, for example.

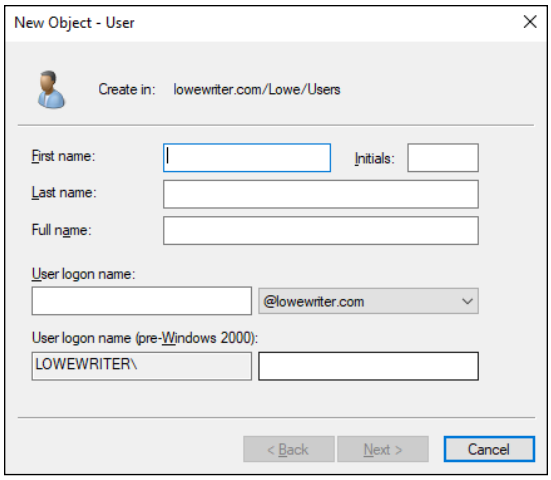


FIGURE 4-2: Creating a new user.

5. **Type the user logon name.**

   This name must be unique within the domain.



TIP

Pick a naming scheme to follow when creating user logon names. You can use the first letter of the first name followed by the complete last name, the complete first name followed by the first letter of the last name, or any other scheme that suits your fancy.

6. **Click Next.**

   The second page of the New Object – User Wizard appears, as shown in Figure 4-3.

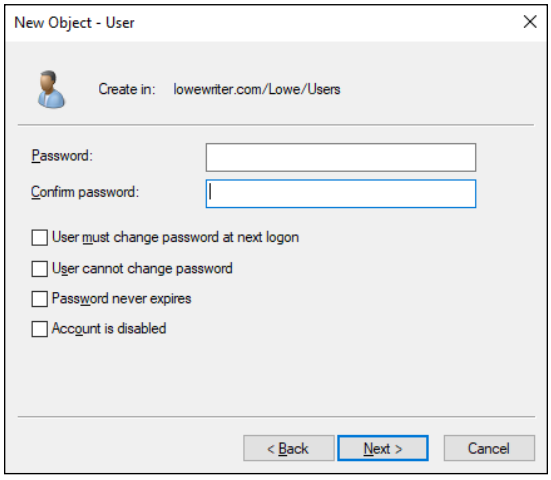


FIGURE 4-3: Setting the user's password.

7. **Type the password twice.**

   You're asked to type the password twice, so type it correctly. If you don't type it identically in both boxes, you're asked to correct your mistake.

8. **Specify the password options that you want to apply.**

   The following password options are available:

   - User Must Change Password at Next Logon.
   - User Cannot Change Password.
   - Password Never Expires.
   - Account Is Disabled.

   For more information about these options, see the section "Setting account options," later in this chapter.

9. **Click Next.**

   You're taken to the final page of the New Object – User Wizard, as shown in Figure 4-4.

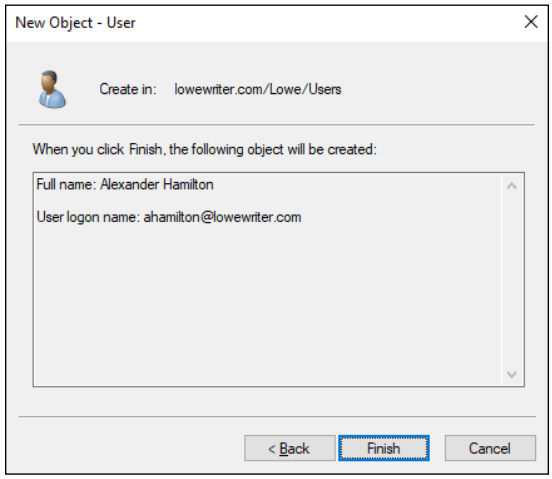


FIGURE 4-4: Verifying the user account information.

10. **Verify that the information is correct and then click Finish to create the account.**

    If the account information isn't correct, click the Back button, and correct the error.

You're done! Now you can customize the user's account settings. At minimum, you'll probably want to add the user to one or more groups. You may also want to add contact information for the user or set up other account options.



TIP

An alternative way to create a new user is to simply copy an existing user. When you copy an existing user, you provide a new username and password and Windows copies all the other property settings from the existing user to the new user.

Setting User Properties

After you've created a user account, you can set additional properties for the user by right-clicking the new user and choosing Properties from the contextual menu. This command brings up the User Properties dialog box, which has about a million tabs that you can use to set various properties for the user. Figure 4-5 shows the General tab, which lists basic information about the user, such as the user's name, office location, and phone number.

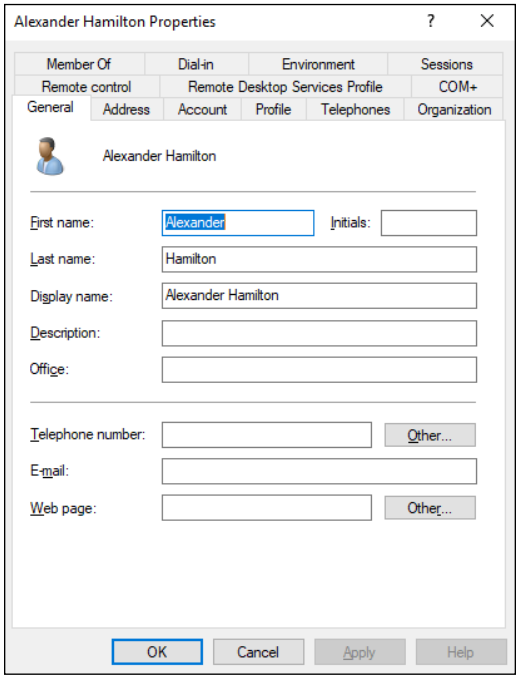


FIGURE 4-5: The General tab.

The following sections describe some of the administrative tasks that you can perform via the various tabs of the User Properties dialog box.

Changing the user's contact information

Several tabs of the User Properties dialog box contain contact information for the user:

- **Address:** Lets you change the user's street address, post office box, city, state, zip code, and so on
- **Telephones:** Lets you specify the user's phone numbers
- **Organization:** Lets you record the user's job title and the name of his or her boss

Setting account options

The Account tab of the User Properties dialog box, shown in Figure 4-6, features a variety of interesting options that you can set for the user. From this dialog box, you can change the user's logon name. In addition, you can change the password options that you set when you created the account, and you can set an expiration date for the account.

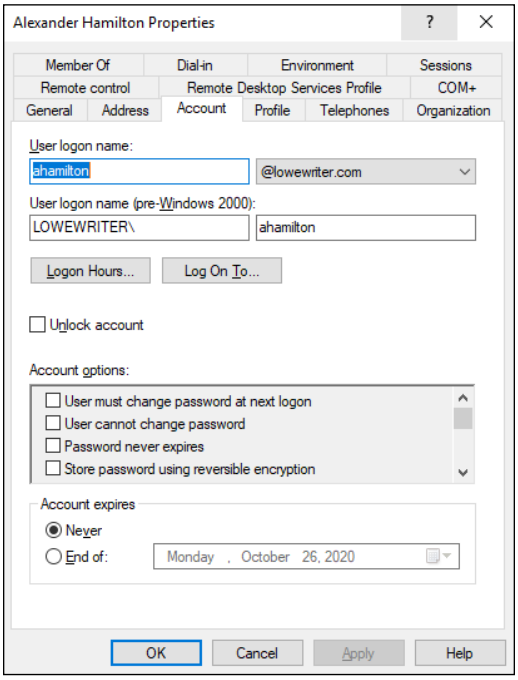


FIGURE 4-6: The Account tab.

The following account options are available in the Account Options list box:

- **User Must Change Password at Next Logon:** This option, which is selected by default, allows you to create a one-time-only password that can get the user started with the network. The first time the user logs on to the network, he or she is asked to change the password.
- **User Cannot Change Password:** Use this option if you don't want to allow users to change their passwords. (Obviously, you can't use this option and the preceding one at the same time.)
- **Password Never Expires:** Use this option if you want to bypass the password-expiration policy for this user so that the user will never have to change his or her password.
- **Store Password Using Reversible Encryption:** This option stores passwords by using an encryption scheme that hackers can easily break, so you should avoid it like the plague.
- **Account Is Disabled:** This option allows you to create an account that you don't yet need. As long as the account remains disabled, the user won't be able to log on. See the section "Disabling and Enabling User Accounts," later in this chapter, to find out how to enable a disabled account.
- **Smart Card Is Required for Interactive Logon:** If the user's computer has a smart card reader to read security cards automatically, check this option to require the user to use it.
- **Account Is Trusted for Delegation:** This option indicates that the account is trustworthy and can set up delegations. This advanced feature usually is reserved for Administrator accounts.
- **Account Is Sensitive and Cannot Be Delegated:** This option prevents other users from impersonating this account.
- **Use DES Encryption Types for This Account:** This option beefs up the encryption for applications that require extra security.
- **Do Not Require Kerberos Preauthentication:** Select this option if you use a different implementation of the Kerberos protocol.

Specifying logon hours

You can restrict the hours during which the user is allowed to log on to the system by clicking the Logon Hours button on the Account tab of the User Properties dialog box. This button brings up the Logon Hours for [User] dialog box, shown in Figure 4-7.

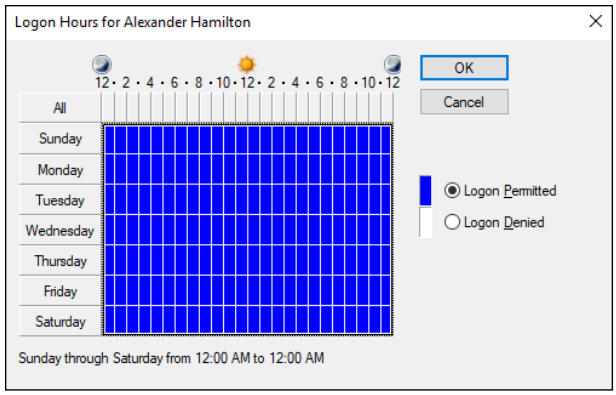


FIGURE 4-7: Restricting the user's logon hours.

Initially, the Logon Hours dialog box is set to allow the user to log on at any time of day or night. To change the hours that you want the user to have access, click a day and time or a range of days and times; choose either Logon Permitted or Logon Denied; and click OK.

Restricting access to certain computers

Normally, a user can use his or her user account to log on to any computer that's part of the user's domain. You can restrict a user to certain computers, however, by clicking the Log On To button on the Account tab of the User Properties dialog box. This button brings up the Logon Workstations dialog box, as shown in Figure 4-8.

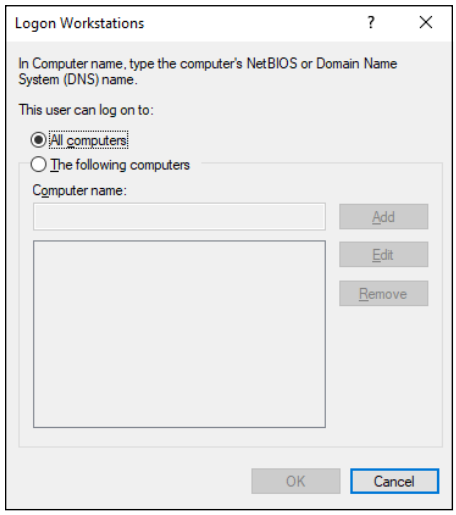


FIGURE 4-8: Restricting the user to certain computers.

To restrict the user to certain computers, select the radio button labeled The Following Computers. Then, for each computer you want to allow the user to log on from, type the computer's name in the text box, and click Add.


TIP

If you make a mistake, you can select the incorrect computer name and then click Edit to change the name or Remove to delete the name.

Setting the user's profile information

The Profile tab, shown in Figure 4-9, lets you configure the user's profile information. This dialog box lets you configure three bits of information related to the user's profile:

- **Profile Path:** This field specifies the location of the user's roaming profile. For more information, see the section "Working with User Profiles," later in this chapter.
- **Logon Script:** This field is the name of the user's logon script. A *logon script* is a batch file that's run whenever the user logs on. The main purpose of the logon script is to map the network shares that the user requires access to. Logon scripts are carryovers from early versions of Windows NT Server. In Windows Server 2025, profiles are the preferred way to configure the user's computer when the user logs on, including setting up network shares. Many administrators still like the simplicity of logon scripts, however. For more information, see the section "Creating a Logon Script," later in this chapter.
- **Home Folder:** This section is where you specify the default storage location for the user.


TIP

The Profile tab lets you specify the location of an existing profile for the user, but it doesn't actually let you set up the profile. For more information about setting up a profile, see the section "Working with User Profiles," later in this chapter.

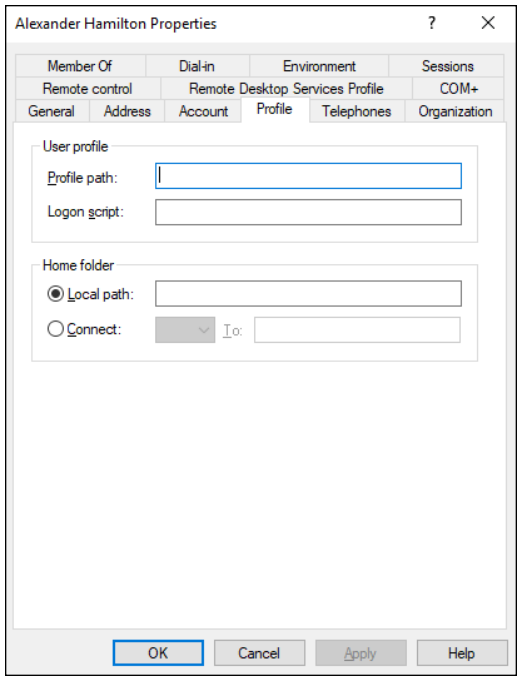


FIGURE 4-9: The Profile tab.

Resetting User Passwords

By some estimates, the single most time-consuming task of most network administrators is resetting user passwords. It's easy to think that users are forgetful idiots, but put yourself in their shoes. Administrators insist that they set their passwords to something incomprehensible, such as 94kD82leL384K; that they change it a week later to something more unmemorable, such as dJUQ63DWd8331; and that they don't write it down. Then administrators get mad when they forget their passwords.

So when a user calls and says that he or she forgot his or her password, the least the administrator can do is be cheerful when resetting it. After all, the user probably spent 15 minutes trying to remember it before finally giving up and admitting failure.

Here's the procedure to reset the password for a user domain account:

1. **Log on as an administrator.**

   You must have administrator privileges to perform this procedure.

2. **In Server Manager, choose Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers management console appears.

3. **Drill down to the organizational unit that contains the user's Active Directory object.**

4. **In the Details pane, right-click the user who forgot his or her password, and choose Reset Password from the contextual menu.**

5. **Type the new password in both password boxes.**

   You have to type the password twice to ensure that you type it correctly.

6. **If desired, select the User Must Change Password at Next Logon option.**

   If you select this option, the password that you assign will work for only one logon. As soon as the user logs on, he or she will be required to change the password.

7. **Click OK.**

   That's all there is to it! The user's password is reset.

Disabling and Enabling User Accounts

If you want to temporarily prevent a user from accessing the network, you can disable his or her account. Then you can enable the account later, when you're ready to restore the user to full access. Here's the procedure:

1. **Log on as an administrator.**

   You must have administrator privileges to perform this procedure.

2. **From Server Manager, choose Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers management console appears.

3. **Drill down to the organizational unit that contains the user's Active Directory object.**

4. **In the Details pane, right-click the user that you want to enable or disable; then choose either Enable Account or Disable Account from the contextual menu to enable or disable the user.**

Deleting a User

Deleting a user account is surprisingly easy. Just follow these steps:

1. **Log on as an administrator.**

   You must have administrator privileges to perform this procedure.

2. **Choose Start➪Administrative Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers management console appears.

3. **Drill down to the organizational unit that contains the user's Active Directory object.**

4. **In the Details pane, right-click the user that you want to delete and then choose Delete from the contextual menu.**

   Windows asks whether you really want to delete the user, just in case you're kidding.

5. **Click Yes.**

   Poof! The user account is deleted.



WARNING

Deleting a user account is a permanent, nonreversible action. Do it only if you're absolutely sure that you'll never want to restore the user's account. If there's any possibility of restoring the account later, you should disable the account rather than delete it.

Working with Groups

A *group* is a special type of account that represents a set of users who have common network access needs. Groups can dramatically simplify the task of assigning network access rights to users. Rather than assign access rights to each user individually, you can assign rights to the group itself. Then those rights automatically extend to any user you add to the group.

The following sections describe some of the key concepts that you need to understand to use groups, along with some of the most common procedures you'll employ when setting up groups for your server.

Group types

Two distinct types of groups exist:

- **Security groups:** Most groups are security groups, which extend access rights to members of the group. If you want to allow a group of users to access your high-speed color laser printer, for example, you can create a group called ColorPrintUsers. Then you can grant permission to use the printer to the ColorPrintUsers group. Finally, you can add individual users to the ColorPrintUsers group.
- **Distribution groups:** Distribution groups aren't used as much as security groups are. They're designed as a way to send email to a group of users by specifying the group as the recipient.

Group scope

A group can have any of three distinct *scopes*, which determine what domains the group's members can belong to:

- **Domain local:** A group with *domain local scope* can have members from any domain. The group can be granted permissions only from the domain in which the group is defined, however.
- **Global:** A group with *global scope* can have members only from the domain in which the group is defined. The group can be granted permissions in any domain in the forest, however. (For more information about forests, refer to Book 6, Chapter 3.)
- **Universal scope:** Groups with *universal scope* are available in all domains that belong to the same forest.

As you can probably guess, universal scope groups are usually used only on very large networks.

One common way you can use domain local and global groups is as follows:

1. **Use domain local groups to assign access rights for network resources.**

   To control access to a high-speed color printer, for example, create a domain local group for the printer. Grant the group access to the printer, but don't add any users to the group.

2. **Use global groups to associate users with common network access needs.**

   Create a global group for users who need to access color printers, for example. Then add each user who needs access to a color printer membership to the group.

3. **Finally, add the global group to the domain local group.**

   That way, access to the printer is extended to all members of the global group.

This technique gives you the most flexibility when your network grows.

Default groups

Windows Server 2025 comes with several predefined groups that you can use. Although you shouldn't be afraid to create your own groups when you need them, there's no reason to create your own group if you find a default group that meets your needs.

Some of these groups are listed in the Builtin container in the Active Directory Users and Computers management console. Others are listed in the Users container. Table 4-1 lists the most useful default groups in Builtin, and Table 4-2 lists the default groups in the Users container.

TABLE 4-1 Default Groups Located in the Builtin Container

Group	Description
Account Operators	This group is for users who should be allowed to create, edit, or delete user accounts but shouldn't be granted full administrator status.
Administrators	This group is for the system administrators who have full control of the domain. The Administrator account is a default member of this group. You should create only a limited number of accounts that belong to this group.
Backup Operators	This group is for users who need to perform backup operations. Because this group must have access to the files that are backed up, it presents a security risk, so you should limit the number of users that you add to this group.
Guests	This group allows members to log on but little else. The default Guest account is a member of this group.
Network Configuration	This group is allowed to twiddle with network configuration settings, including releasing and renewing DHCP leases.
Print Operators	This group grants users access to printers, including the ability to create and share new printers and to manage print queues.
Remote Desktop Users	This group can remotely log on to domain controllers in the domain.
Replicator	This group is required to support directory replication. Don't add users to this group.
Server Operators	These users can log on locally to a domain controller.
Users	These users can perform common tasks, such as running applications and using local and network printers.

TABLE 4-2 Default Groups Located in the Users Container

Group	Description
Cert Publishers	These users can publish security certificates for users and computers.
DnsAdmins	This group is installed if you install DNS. It grants administrative access to the DNS Server service.
DnsUpdateProxy	This group is installed if you install DNS. It allows DNS clients to perform dynamic updates on behalf of other clients, such as DHCP servers.
Domain Admins	These users have complete control of the domain. By default, this group is a member of the Administrators group on all domain controllers, and the Administrator account is a member of this group.
Domain Computers	This group contains all computers that belong to the domain. Any computer account created becomes a member of this group automatically.
Domain Controllers	This group contains all domain controllers in the domain.
Domain Guests	This group contains all domain guests.
Domain Users	This group contains all domain users. Any user account created in the domain is added to this group automatically.
Group Policy	These users can modify group policy for the domain.
IIS_WPG	This group is created if you install IIS. It's required for IIS to operate properly.
RAS and IAS Servers	This group is required for RAS and IAS servers to work properly.

Creating a group

If none of the built-in groups meets your needs, you can create your own group by following these steps:

1. **Log on as an administrator.**

   You must have administrator privileges to perform this procedure.

2. **From Server Manager, choose Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers management console appears.

3. **Right-click the domain to which you want to add the group and then choose New➪Group from the contextual menu.**

   The New Object – Group dialog box appears, as shown in Figure 4-10.

4. **Type the name for the new group.**

   Enter the name in both text boxes.

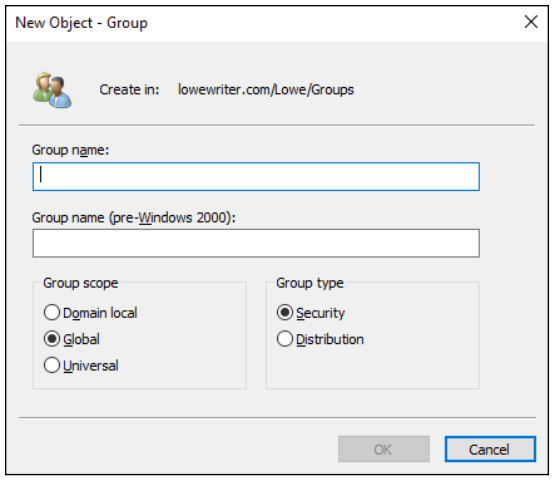


FIGURE 4-10: Creating a new group.

5. **Choose the group scope.**

   The choices are Domain Local, Global, and Universal. For groups that will be granted access rights to network resources, choose Domain Local. Use Global for groups to which you'll add users and Domain Local groups. Use Universal groups only if you have a large network with multiple domains.

6. **Choose the group type.**

   The choices are Security and Distribution. In most cases, choose Security.

7. **Click OK.**

   The group is created.

Adding a member to a group

Groups are collections of objects, called *members.* The members of a group can be user accounts or other groups. When you create a group, it has no members. As a result, the group isn't useful until you add at least one member.

Follow these steps to add a member to a group:

1. **Log on as an administrator.**

   You must have administrator privileges to perform this procedure.

2. **Choose Start➪Administrative Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers management console appears.

3. **Open the folder that contains the group to which you want to add members and then double-click the group.**

   The Group Properties dialog box appears.

4. **Click the Members tab.**

   The members of the group are displayed, as shown in Figure 4-11.

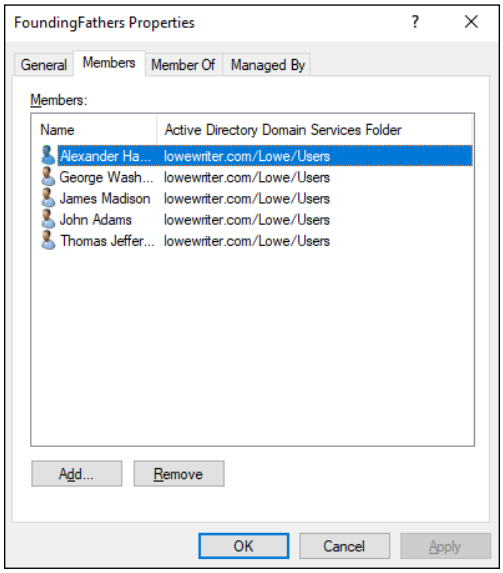


FIGURE 4-11: Adding members to a group.

5. **Click Add, type the name of a user or another group that you want to add to this group, and click OK.**

   The member is added to the list.

6. **Repeat Step 5 for each user or group that you want to add.**

   Keep going until you've added everyone!

7. **Click OK.**

That's all there is to it.



TIP

The Group Properties dialog box also has a Member Of tab that lists each group that the current group is a member of.

Adding members to a group is only half the process of making a group useful. The other half is adding access rights to the group so that the members of the group can actually *do* something. The procedures for doing that are covered in Book 6, Chapter 5.

Working with User Profiles

User profiles automatically maintain desktop settings for Windows users. By default, a user profile is stored on the user's local computer. The following items are just some of the settings that are stored as part of the user profile:

- **Desktop settings** in the Display Properties dialog box, including wallpaper, screen savers, and color schemes
- **Start-menu programs** and Windows toolbar options
- **Favorites,** which provide easy access to the files and folders that the user accesses frequently
- **Application Data,** such as option settings, custom dictionaries, and so on
- **Cookies,** used for web browsing
- **Recent Documents,** which keeps shortcuts to the documents most recently accessed by the user
- **Templates,** which stores user templates
- **Network,** which keeps shortcuts to the user's network locations
- **Send To,** which keeps shortcuts to document-handling utilities
- **Local Settings,** such as history and temporary files
- **Printers,** which keeps shortcuts to the user's printers
- **Documents,** which stores the user's local documents

Types of user profiles

Four types of user profiles exist:

- **Local user profile:** A local user profile is stored on the user's local computer and is applied only when the user logs on to that computer. A local user profile is created automatically when a new user logs on.

- **Roaming user profile:** A roaming user profile is created on a network share. That way, the user can access the roaming profile when he or she logs on to any computer on the network.
- **Mandatory user profile:** A mandatory user profile is a roaming user profile that the user is not allowed to change. One benefit of mandatory user profiles is that users can't mess up their desktop settings. Another benefit is that you can create a single mandatory profile that can be used by multiple users.
- **Temporary user profile:** If a roaming or mandatory profile isn't available for some reason, a temporary user profile is automatically created for the user. The temporary profile is deleted when the user logs off, so any changes that the user makes while using a temporary profile are lost at the end of the session.

Roaming profiles

A *roaming user profile* is simply a user profile that has been copied to a network share so that it can be accessed from any computer on the network.

Before you can create roaming user profiles, you should create a shared folder on the server to hold the profiles. You can name the shared folder anything you like, but most administrators call it Users. For information on the procedure to create a shared folder, see Book 6, Chapter 5.

After you've created the shared Users folder, you can copy the profile to the server by following these steps at the user's local computer:

1. **Log on to the computer by using an account other than the one you want to make a user account.**

   Windows won't let you copy the profile that you're logged on with.

2. **Open File Explorer, right-click This PC, choose Properties, click Advanced System Settings, and then click Settings in the User Profile section.**

   This step brings up the User Profiles dialog box, shown in Figure 4-12.

3. **Select the profile that you want to copy and then click Copy To.**

   A Copy To dialog box appears.

4. **Type the path and name for the roaming profile in the Copy Profile To box.**

   To copy a profile named Doug to the Users share on a server named Server01, for example, type **\\Server01\Users\Doug**.

5. **Click OK.**

   The profile is copied.

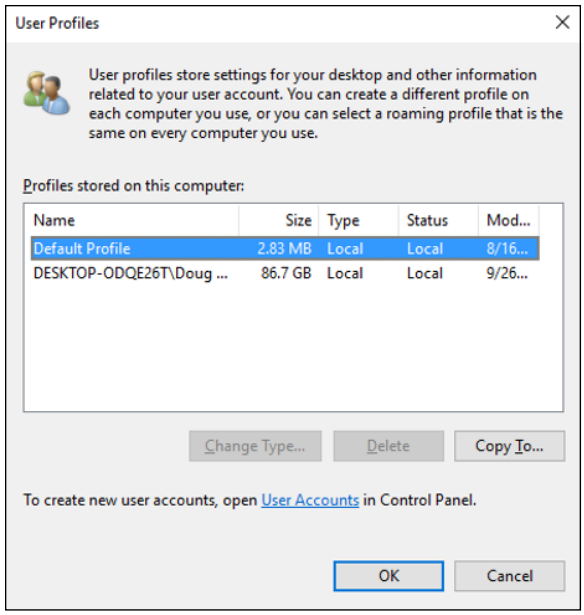


FIGURE 4-12: The User Profiles dialog box.

Now you can go back to the server, log on as an administrator, and follow these steps to designate a roaming profile for the user's domain account:

1. **From the Server Manager, choose Tools➪Active Directory Users and Computers.**

   The Active Directory Users and Computers management console appears.

2. **Right-click the user account, and choose Properties from the contextual menu.**

   The User Properties dialog box appears.

3. **Click the Profile tab.**

   The Profile tab appears. (This tab is shown in Figure 4-9, earlier in this chapter, so I won't repeat it here.)

4. **Type the path and name of the profile in the Profile Path text box.**

   The path and name that you type here should be the same path and name that you used to copy the profile to the server.

5. **Click OK.**

Creating a Logon Script

A *logon script* is a batch file that's run automatically whenever a user logs on. The most common reason for using a logon script is to map the network shares that the user needs access to. Here's a simple logon script that maps three network shares:

```
echo off
net use m: \\server1\shares\admin
net use n: \\server1\shares\mktg
net use o: \\server2\archives
```

Here, two shares on server1 are mapped to drives `M:` and `N:`, and a share on server2 is mapped as drive `O:`.

If you want, you can use the special variable `%username%` to get the user's username. This variable is useful if you've created a folder for each user, and you want to map a drive to each user's folder, as follows:

```
net use u: \\server1\users\%username%
```

If a user logs on with the username `dlowe`, for example, drive `U:` is mapped to `\\server1\users\dlowe`.


TIP

Scripts should be saved in the Scripts folder, which is buried deep in the bowels of the SYSVOL folder — typically, `c:\Windows\SYSVOL\Sysvol\`*`domainname`*`\Scripts`, where *`domainname`* is your domain name. Because you need to access this folder frequently, I suggest creating a shortcut to it on your desktop.

After you've created a logon script, you can assign it to a user by using the Profile tab of the User Properties dialog box. For more information, see the section "Setting the user's profile information," earlier in this chapter.

IN THIS CHAPTER

- Looking at file server settings
- Sharing folders
- Setting permissions

Chapter 5
Configuring a File Server

In this chapter, you discover how to set up and manage file and print servers for Windows Server 2025. Because the features for file and print servers are essentially the same for previous versions of Windows Server, the techniques presented in this chapter should work for older versions as well.

Understanding Permissions

Before I get into the details of setting up a file server, you need to have a solid understanding of the concept of permissions. *Permissions* allow users to access shared resources on a network. Simply sharing a resource such as a disk folder or a printer doesn't guarantee that a given user is able to access that resource. Windows makes this decision based on the permissions that have been assigned to various groups for the resource and group memberships of the user. If the user belongs to a group that has been granted permission to access the resource, the access is allowed. If not, access is denied.

In theory, permissions sound pretty simple. In practice, however, they can get pretty complicated. The following paragraphs explain some of the nuances of how access control and permissions work:

- Every object — that is, every file and folder — on an NTFS volume has a set of permissions called the *access control list* (ACL) associated with it.

- The ACL identifies the users and groups who can access the object and specifies what level of access each user or group has. A folder's ACL may specify that one group of users can read files in the folder, whereas another group can read and write files in the folder, and a third group is denied access to the folder.
- Container objects — files and volumes — allow their ACLs to be inherited by the objects that they contain. As a result, if you specify permissions for a folder, those permissions extend to the files and child folders that appear within it.

  Table 5-1 describes the six permissions that can be applied to files and folders on an NTFS volume.


TIP

- Actually, the six file and folder permissions comprise various combinations of *special permissions* that grant more detailed access to files or folders. Table 5-2 lists the special permissions that apply to each of the six file and folder permissions.
- It's best to assign permissions to groups rather than to individual users. Then if a particular user needs access to a particular resource, add that user to a group that has permission to use the resource.

TABLE 5-1 File and Folder Permissions

Permission	Description
Full Control	The user has unrestricted access to the file or folder.
Modify	The user can change the file or folder's contents, delete the file or folder, read the file or folder, or change the attributes of the file or folder. For a folder, this permission allows you to create new files or subfolders within the folder.
Read & Execute	For a file, this permission grants the right to read or execute the file. For a folder, this permission grants the right to list the contents of the folder or to read or execute any of the files in the folder.
List Folder Contents	This permission applies only to folders; it grants the right to list the contents of the folder.
Read	This permission grants the right to read the contents of a file or folder.
Write	This permission grants the right to change the contents of a file or its attributes. For a folder, this permission grants the right to create new files and subfolders within the folder.

TABLE 5-2 Special Permissions

Special Permission	Full Control	Modify	Read & Execute	List Folder Contents	Read	Write
Traverse Folder/Execute File	X	X	X	X		
List Folder/Read Data	X	X	X	X	X	
Read Extended Attributes	X	X	X	X	X	
Create Files/Write Data	X	X				X
Create Folders/Append Data	X	X				X
Write Attributes	X	X				X
Write Extended Attributes	X	X				X
Delete Subfolders and Files	X					
Delete	X	X				
Read Permissions	X	X	X	X	X	X
Change Permissions	X					
Take Ownership	X					
Synchronize	X	X	X	X	X	X

Understanding Shares

A *share* is simply a folder that is made available to other users via the network. Each share has the following elements:

- **Share name:** The name by which the share is known over the network. To make the share name compatible with older computers, you should stick to eight-character share names whenever possible.
- **Path:** The path to the folder on the local computer that's being shared, such as `D:\Accounting`.
- **Description:** An optional one-line description of the share.
- **Permissions:** A list of users or groups who have been granted access to the share.

When you install Windows and configure various server roles, special shared resources are created to support those roles. You shouldn't disturb these special shares unless you know what you're doing. Table 5-3 describes some of the most common special shares.

TABLE 5-3 Special Shares

Share Name	Description
drive$	The root directory of a drive; for example, C$ is the root share for the C: drive.
ADMIN$	Used for remote administration of a computer. This share points to the operating system folder (usually, C:\Windows).
IPC$	Used by named pipes, a programming feature that lets processes communicate with one another.
NETLOGON	Required for domain controllers to function.
SYSVOL	Another required domain controller share.
PRINT$	Used for remote administration of printers.
FAX$	Used by fax clients.

Notice that some of the special shares end with a dollar sign ($). These shares are *hidden shares* that aren't visible to users. You can still access them, however, by typing the complete share name (including the dollar sign) when the share is needed. The special share C$, for example, is created to allow you to connect to the root directory of the C: drive from a network client. You wouldn't want your users to see this share, would you? (Shares such as C$ are also protected by privileges, of course, so if an ordinary user finds out that C$ is the root directory of the server's C: drive, he or she still can't access it.)

Considering Best Practices for Setting Up Shares

Before you actually share any data on a file server, you should consider the following common practices:

- Never share data from a file server's system volume (that is, the C: drive). Doing so is a bad idea because if the operating system becomes corrupted, you may need to restore the system volume from a backup. If the system volume contains data that's shared on the network, you'll have to revert that data to the version stored on the backup being restored. By placing shared data on a separate volume, you can restore the system volume without affecting data.
- It's a common practice to use the same name for both the folder being shared and the name of the share itself. So, a shared folder named Accounting is usually shared using the share name Accounting.

» It's also a common practice to create a folder named `Shares` at the root level of the volume that contains the shared data. Then create individual shared folders within the `Shares` folder.

For example, if the `D:` drive will contain shared folders named `Accounting`, `Administration`, and `Marketing`, you would create a folder structure like this:

```
D:\Shares
      Accounting
      Administration
      Marketing
```

Managing Your File Server

To manage shares on a Windows Server 2025 system, open the Server Manager, and select File and Storage Services in the task pane on the left side of the window. Then click Shares to reveal the management console shown in Figure 5-1.

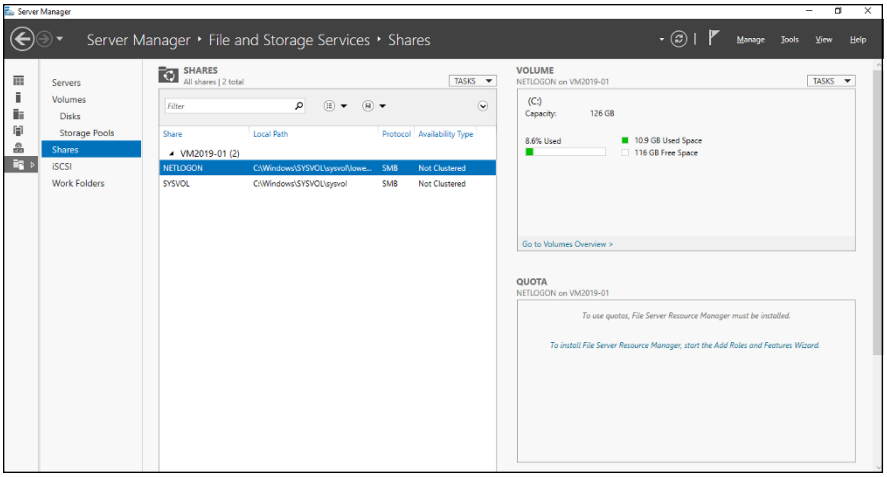


FIGURE 5-1: Managing shares in Windows Server 2025.

The following sections describe some of the most common procedures that you'll use when managing your file server.

Using the New Share Wizard

To be useful, a file server should offer one or more *shares* — folders that have been designated as publicly accessible via the network. To create a new share, use the New Share Wizard. The following procedure shows how to share a folder named `Accounting` on a file server's `D:` drive; the name of the share will be `Accounting`:

1. **In Server Manager, click File and Storage Services, click Shares, and then choose New Share from the Tasks drop-down menu.**

   The opening screen of the New Share Wizard appears, as shown in Figure 5-2. Here, the wizard asks you what folder you want to share.

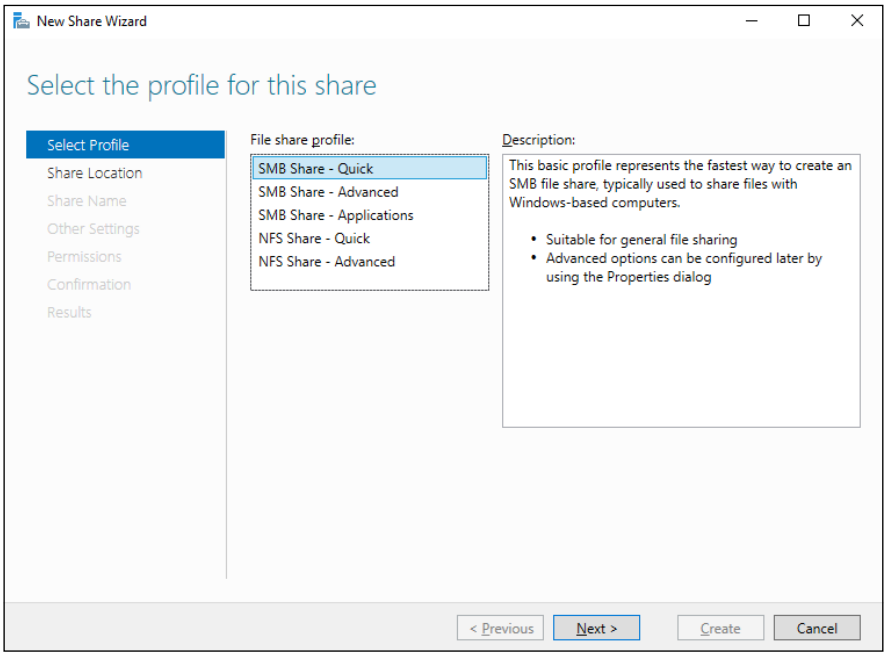


FIGURE 5-2: The New Share Wizard comes to life.

2. **Select SMB Share – Quick in the list of profiles and then click Next.**

   Next, the New Share Wizard asks for the location of the share, as shown in Figure 5-3.

3. **Select the server you want the share to reside on.**

   For this example, I chose the server named `VM2025-01`.

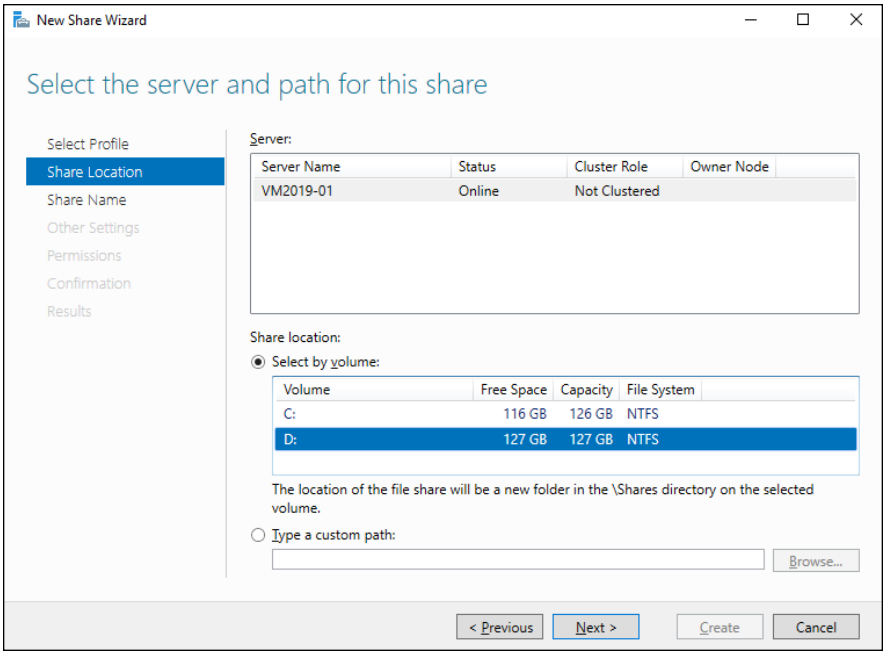


FIGURE 5-3: The wizard asks where you'd like to locate the share.

4. **Select the location of the share by choosing one of these two options:**

   - *Select by Volume:* This option selects the volume on which the shared folder will reside while letting the New Share Wizard create a folder for you. If you select this option, the wizard will create the shared folder on the designated volume. Use this option if the folder doesn't yet exist and you don't mind Windows placing it in the default location, which is inside a folder called Shares on the volume you specify.
   - *Type a Custom Path:* Use this option if the folder exists or if you want to create one in a location other than the Shares folder.

   For this example, I chose the Select by Volume example to allow the wizard to create the share from a folder on the `D:` drive.

5. **Click Next.**

   The screen shown in Figure 5-4 appears.

6. **Type the name that you want to use for the share in the Share Name box.**

   For this example, I entered the share name `Accounting`. Note that as you type the name of the share, the local path to the share is updated to reflect name of the share. By default, the shared folder is in a root-level folder named `Shares`, and the share name and the folder name are the same.

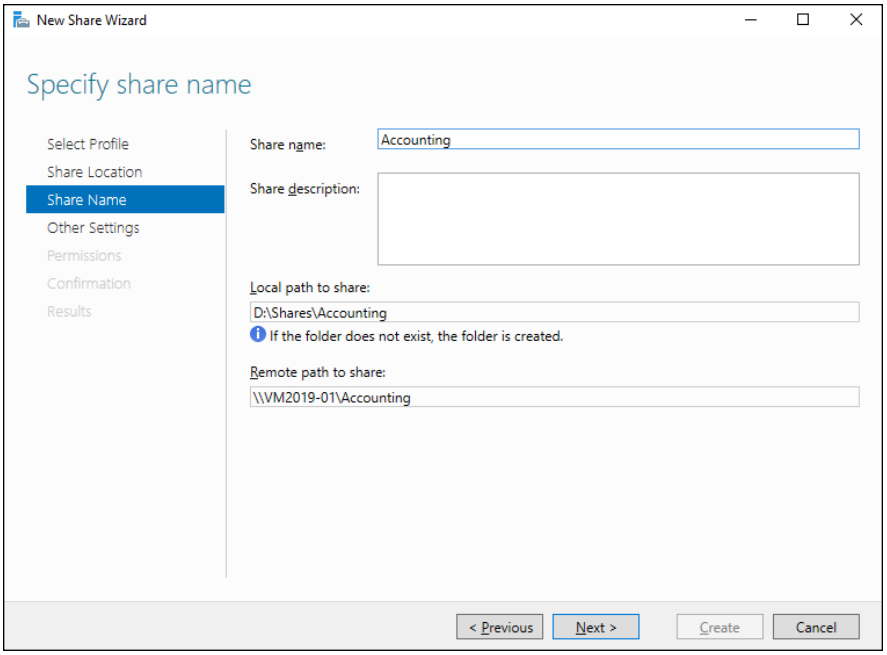


FIGURE 5-4: The wizard asks for the share name and description.

7. **Enter a description for the share.**

   For this example, I left the description blank.

8. **Click Next.**

   The screen shown in Figure 5-5 appears.

9. **Select the share settings you'd like to use:**

   - *Enable Access-Based Enumeration:* Hides files that the user does not have permission to access
   - *Allow Caching of Share:* Makes the files available to offline users
   - *Encrypt Data Access:* Encrypts files accessed via the share

10. **Click Next.**

    The wizard displays the default permissions that will be used for the new share, as shown in Figure 5-6.

11. **If you want to customize the permissions, click Customize Permissions.**

    This button summons the Advanced Security Settings for Data dialog box, which lets you customize both the NTFS and the share permissions.

    You can always customize the NTFS and share permissions later. For more information, refer to the section "Granting permissions," later in this chapter.

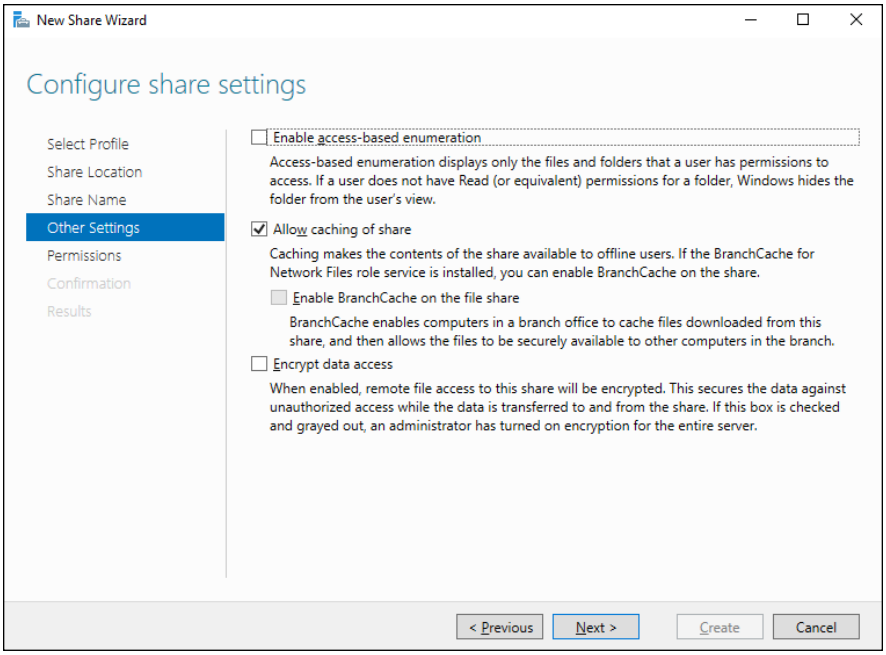


FIGURE 5-5: Specifying the share settings.

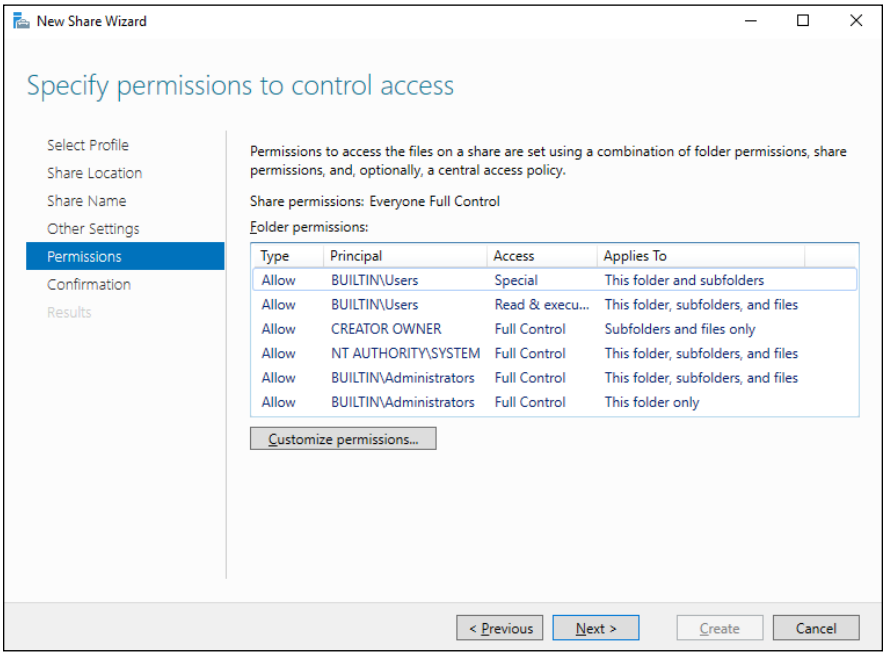


FIGURE 5-6: Setting the share permissions.

12. **Click Next.**

The confirmation screen appears, as shown in Figure 5-7.

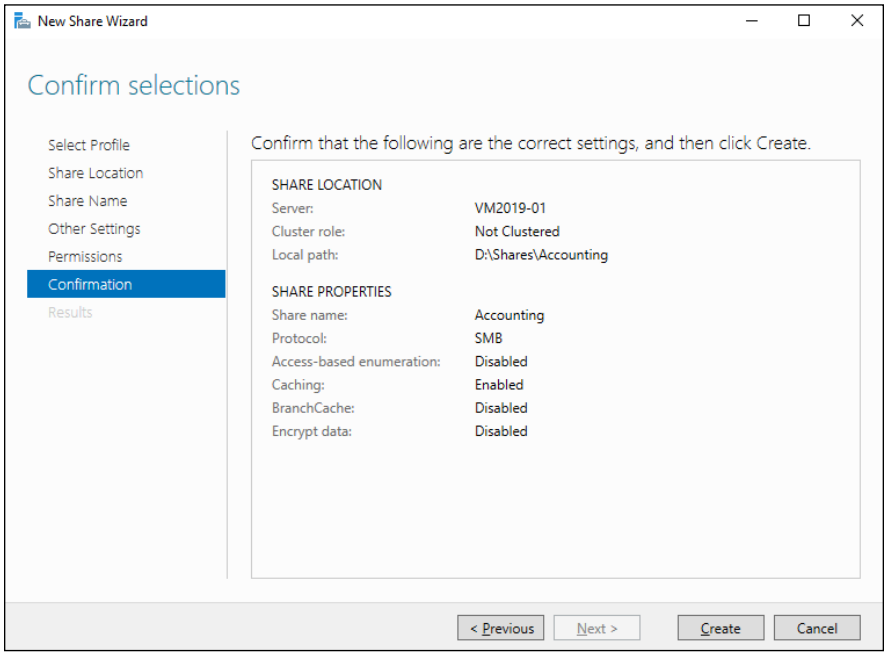


FIGURE 5-7: Confirming your share settings.

13. **Verify that all the settings are correct and then click Create.**

The share is created! A results page is displayed.

Sharing a folder without the wizard

If you think wizards should be confined to *Harry Potter* movies, you can set up a share without bothering with the wizard. Just follow these steps:

1. **Open File Explorer and navigate to the folder that you want to share.**
2. **Right-click the folder, and choose Properties from the contextual menu.**

   This action brings up the Properties dialog box for the folder.
3. **Select the Sharing tab (shown in Figure 5-8).**
4. **Click the Advanced Sharing button.**

   The Advanced Sharing dialog box appears.

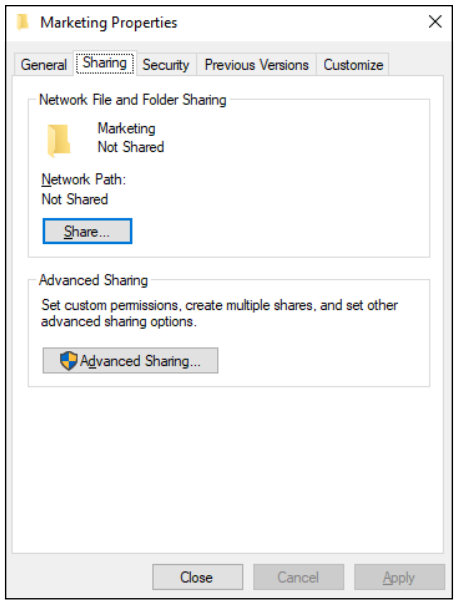


FIGURE 5-8: Manually sharing a folder.

5. **Select the Share This Folder check box to designate the folder as shared.**

   The rest of the controls in this dialog box will be unavailable until you select this check box (see Figure 5-9).

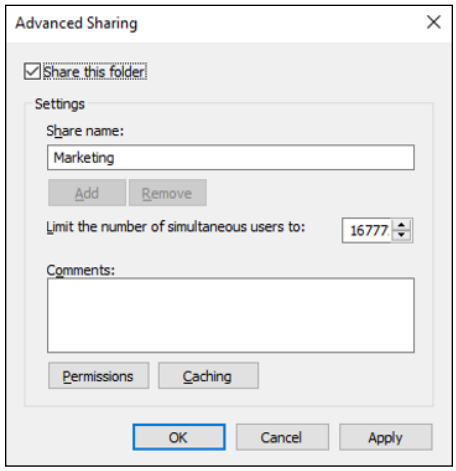


FIGURE 5-9: Setting the share name.

6. **Type the name that you want to use for the share in the Share Name box, and type a description of the share in the Comments box.**

   The default name is the name of the folder being shared.

The description is strictly optional but sometimes helps users determine the intended contents of the folder.

7. **If you want to specify permissions now, click Permissions.**

   This button brings up a dialog box that lets you create permissions for the share. For more information, see the next section, "Granting permissions."

8. **Click OK.**

   The folder is now shared and you're returned to the Properties dialog box for the folder.

9. **Click Close.**

   The folder Properties dialog box vanishes.

Granting permissions

When you first create a file share, all users are granted read-only access to the share. If you want to allow users to modify files in the share or allow them to create new files, you need to add permissions. Here's how to do this via File Explorer:

1. **Open File Explorer and browse to the folder whose permissions you want to manage.**

2. **Right-click the folder you want to manage, and choose Properties from the contextual menu.**

   The Properties dialog box for the folder appears.

3. **Select the Sharing tab; then click Advanced Sharing.**

   The Advanced Sharing dialog box appears.

4. **Click Permissions.**

   The dialog box shown in Figure 5-10 appears. This dialog box lists all the users and groups to whom you've granted permission for the folder. When you select a user or group from the list, the check boxes at the bottom of the list change to indicate which specific permissions you've assigned to each user or group.

5. **Click Add.**

   The dialog box shown in Figure 5-11 appears.

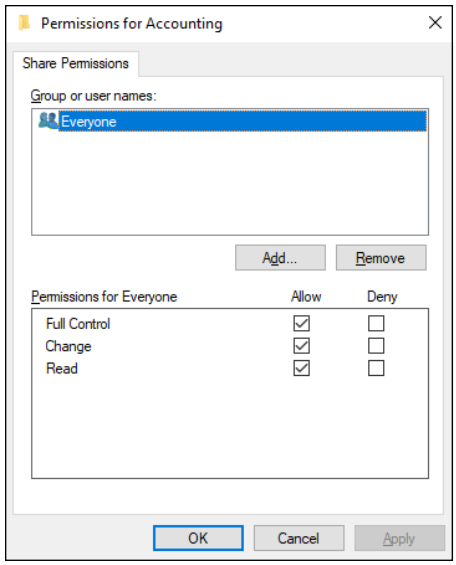


FIGURE 5-10: Setting the share permissions.

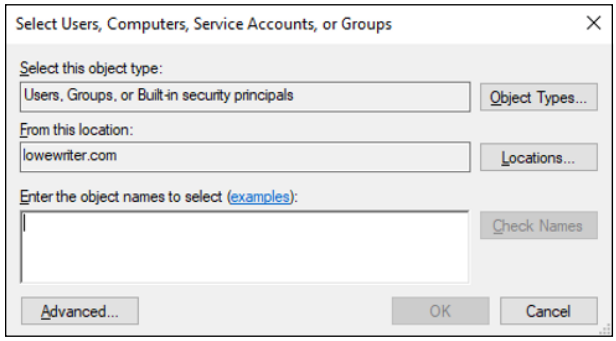


FIGURE 5-11: The Select Users, Computers, Service Accounts, or Groups dialog box.

6. **Type the name of the user or group to whom you want to grant permission and then click OK.**



TIP

If you're not sure of the name, click Advanced. This action brings up a dialog box that lets you search for existing users. You can click the Find Now button to display a list of all users and groups in the domain. Alternatively, you can enter the first part of the name that you're looking for before you click Find Now to search more specifically.

When you click OK, you return to the Share Permissions tab, with the new user or group added.

7. **Select the appropriate Allow and Deny check boxes to specify which permissions to allow for the user or group.**
8. **Repeat Steps 5 through 7 for any other permissions that you want to add.**
9. **When you're done, click OK.**

Here are a few other thoughts to ponder concerning adding permissions:



TIP

- If you want to grant full access to everyone for this folder, don't bother adding another permission. Instead, select the Everyone group and then select the Allow check box for each permission type.
- You can remove a permission by selecting the permission and then clicking Remove.



REMEMBER

- The permissions assigned in this procedure apply only to the share itself. The underlying folder also has a set of NTFS permissions. The effective permissions granted to a particular user are a combination of the share permissions and the NTFS permissions, with the more restrictive permissions always overriding less restrictive permissions. For example, if the share permissions grant a user Full Control permission but the folder's NTFS permissions grant the user only Read permission, the user has only Read permission for the folder.

  To set the NTFS permissions for a folder, open the folder's Properties dialog box and select the Security tab.

IN THIS CHAPTER

- Looking at group policy concepts
- Enabling group policy on a Windows Server
- Editing group policy objects

Chapter 6
Using Group Policy

Group policy refers to a feature of Windows operating systems that lets you control how certain aspects of Windows and other Microsoft software work throughout your network. Many features that you might expect to find in a management console, such as Active Directory Users and Computers, are — or at least can be — controlled by group policy instead. You must use group policy to control how often users must change their passwords, for example, and how complicated their passwords must be. And you can use group policy to set the default home page users will encounter when they launch a web browser. In short, group policy is an important tool for any Windows network administrator.

Unfortunately, group policy can be a confusing beast. In fact, it's one of the most confusing aspects of Windows network administration. So don't be put off if you find this chapter more confusing than other chapters in this minibook. Group policy becomes clear after you spend some time actually working with it.

Understanding Group Policy

Here it is in a nutshell: Group policy consists of a collection of *group policy objects* (GPOs) that define individual policies. These policy objects are selectively applied to both users and computers. Each policy object specifies how some aspect of Windows or some other Microsoft software should be configured. A group policy object might specify the home page that's initially displayed when any user

launches Internet Explorer, for example. When a user logs on to the domain, that policy object is retrieved and applied to the user's Internet Explorer configuration.

Group policy objects can apply to either computers or users. A policy that applies to a computer will be enforced for any user of the computer, and a policy that applies to a user will be enforced for that user no matter what computer he or she logs on to. As a network administrator, you'll be concerned mostly with policies that apply to users. But computer policies are useful from time to time as well.

To use group policy, you have to know how to do two things: (1) create individual group policy objects, and (2) apply — or *link* — those objects to user and computer objects. Both tasks can be a little tricky.

The trick to creating group policy objects is finding the particular setting you want to employ. Trying to find a specific group policy among the thousands of available policies can be frustrating. Suppose that you want to force all network users to change their passwords every 30 days. You know that a group policy controls the password-expiration date. But where is it? You'll find help with this aspect of working with group policy in the section titled "Creating Group Policy Objects," later in this chapter.

After you've created a group policy object, you then are faced with the task of linking it to the users or computers you want it to apply to. Creating a policy that applies to all users or computers is simple enough. But things get more complicated if you want to be more selective — for example, if you want the policy to apply only to users in a particular organizational unit (OU) or to users that belong to a particular group. You'll find help for this aspect of working with group policy in the section "Filtering Group Policy Objects," later in this chapter.

Enabling Group Policy Management on Windows Server 2025

Before you can work with group policy on a Windows Server 2025, you must enable group policy on the server. The procedure is simple enough and needs to be done only once for each server. Here are the steps:

1. **In the Server Manager, click Add Roles and Features.**
2. **Follow the wizard until you get to the Select Features screen, which is shown in Figure 6-1.**
3. **If the Group Policy Management check box is not already checked, select it.**
4. **Click Next.**

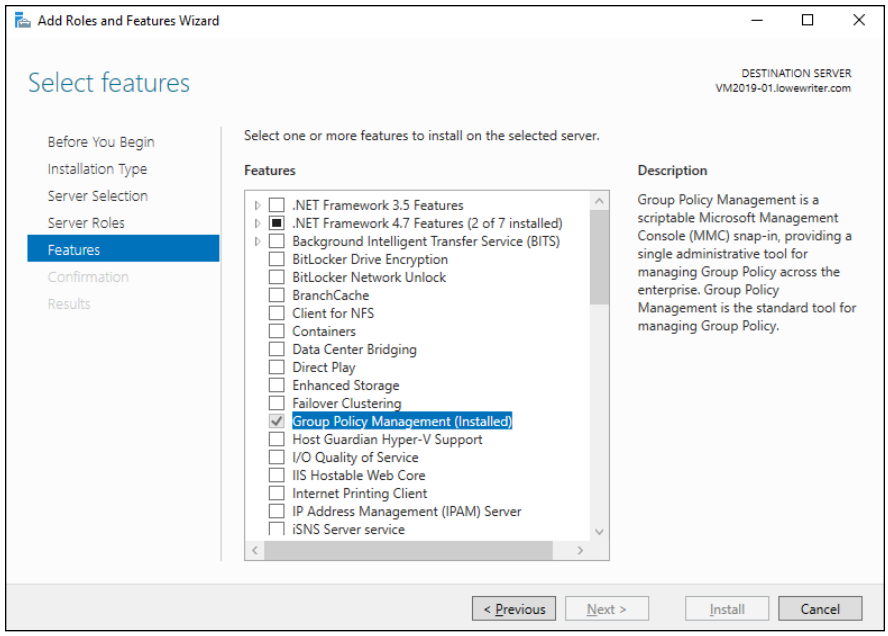


FIGURE 6-1: Enabling group policy management on Windows Server 2025.

5. **When the confirmation page appears, click Install.**

   Be patient; installation may take a few minutes.

6. **Click Close.**

   You're done!

After you've completed this procedure, a new command titled Group Policy Management appears on the Tools menu in the Server Manager.

Creating Group Policy Objects

The easiest way to create group policy objects is to use the Group Policy Management console, which you can run from the Server Manager by choosing Tools➪Group Policy Management.

A single group policy object can consist of one setting or many individual group policy settings. The Group Policy Management console presents the thousands of group policy settings that are available for your use in several categories. The more you work with group policy, the more these categories will begin to make sense. When you get started, you can expect to spend a lot of time hunting through the lists of policies to find the specific one you're looking for.

The easiest way to learn how to use the Group Policy Management console is to use it to create a simple group policy object. In the following procedure, I show you how to create a GPO that defines a group policy enabling Windows Update for all computers in a domain so that users can't disable Windows Update.

1. **In the Server Manager, choose Tools➪Group Policy Management.**

   The Group Policy Management console appears, as shown in Figure 6-2.

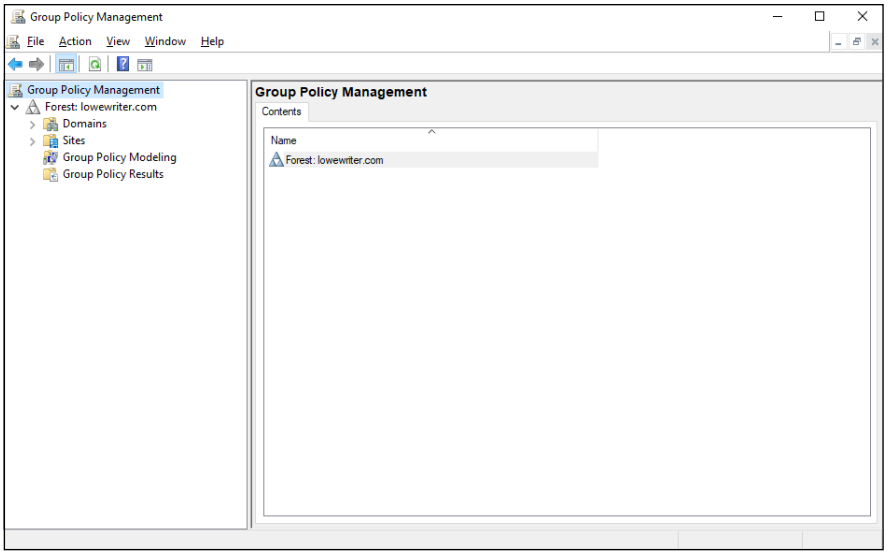


FIGURE 6-2: The Group Policy Management console.

2. **In the Navigation pane, drill down through the Domains node to the node for your domain, then select the Group Policy Objects node for your domain.**

3. **Right-click the Group Policy Objects node and then choose New from the contextual menu that appears.**

   This command brings up the dialog box shown in Figure 6-3.

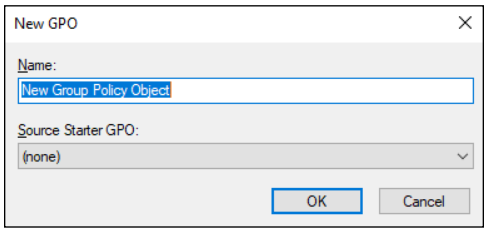


FIGURE 6-3: Creating a new group policy object.

4. **Type a name for the group policy object and then click OK.**

   For this example, type something like Windows Update for a policy that will manage the Windows Update feature.

   When you click OK, the group policy object is created and appears in the Group Policy Objects section of the Group Policy Management window.

5. **Double-click the new group policy.**

   The group policy opens, as shown in Figure 6-4. Note that at this stage, the Location section of the group policy doesn't list any objects. As a result, this policy is not yet linked to any Active Directory domains or groups. I get to that topic in a moment. First, I create the policy settings.

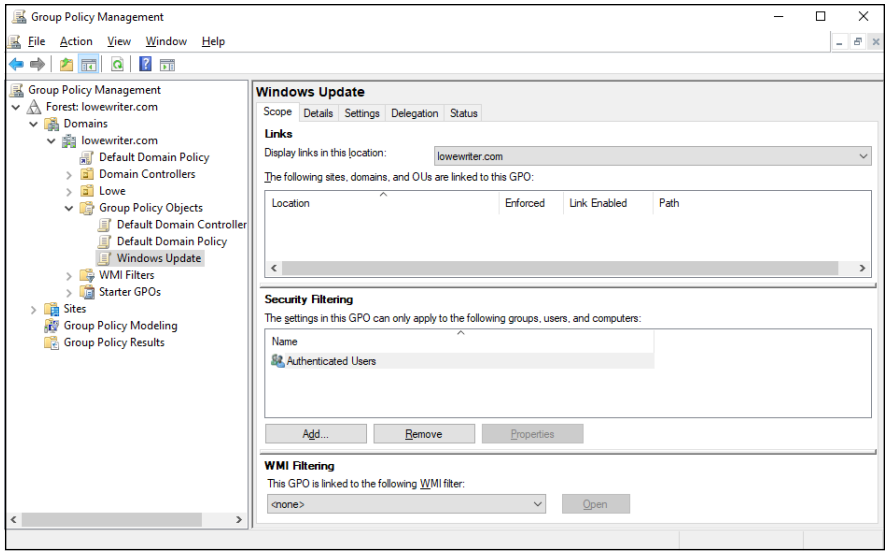


FIGURE 6-4: A new group policy object.

6. **Click the Settings tab.**

   The message "Generating Report" appears for a moment, and then the group policy settings are displayed, as shown in Figure 6-5.



TIP

   If you get a message indicating that the content has been blocked by security settings, click Add to enable access. You'll see the policy settings.

7. **Right-click Computer Configuration and then choose Edit from the contextual menu.**

   This command opens the Group Policy Management Editor, as shown in Figure 6-6, where you can edit the Computer Configuration policies.

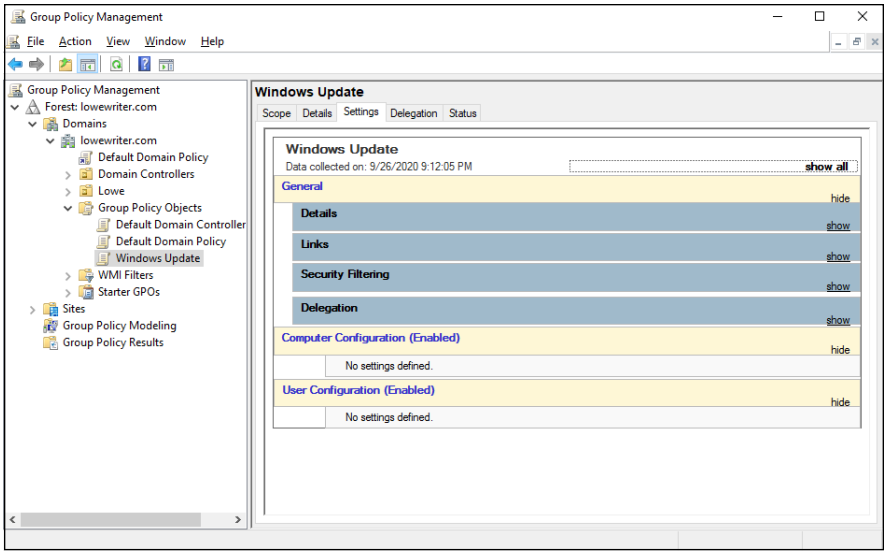


FIGURE 6-5: Group policy settings.

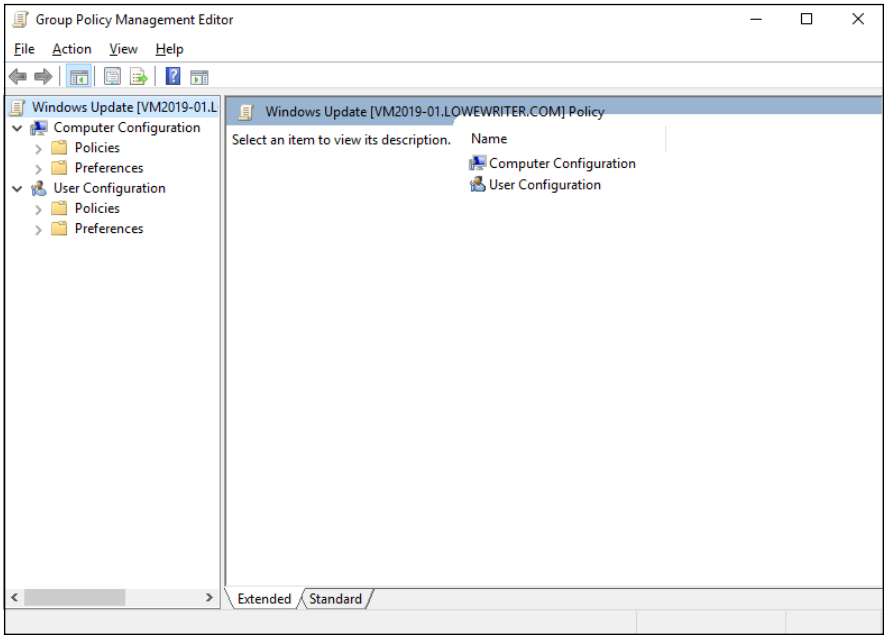


FIGURE 6-6: Editing group policy.

8. In the Navigation pane, navigate to Computer Configuration➪Policies➪Administrative Templates➪Windows Components➪Windows Update.

This step brings up the Windows Update policy settings, as shown in Figure 6-7.

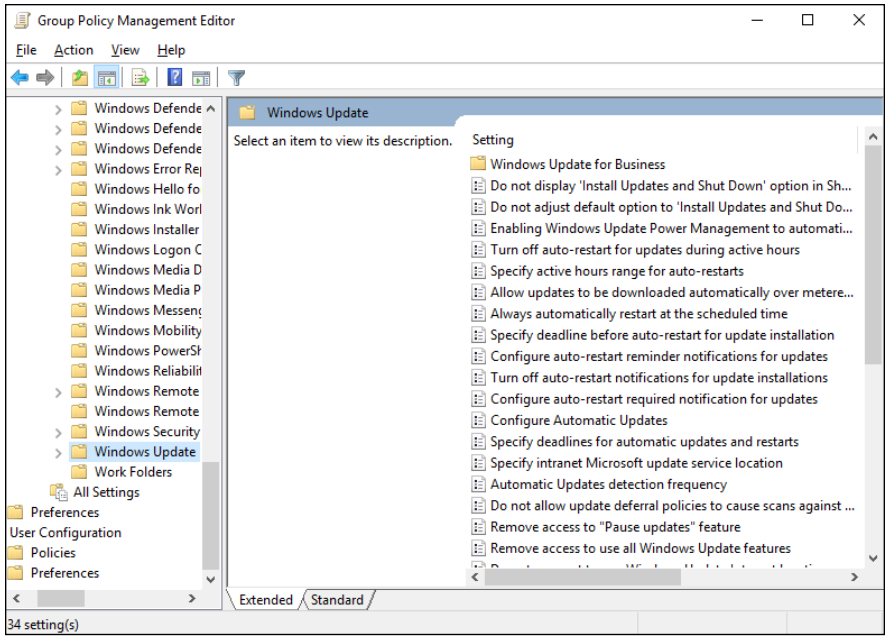


FIGURE 6-7: The Windows Update policy settings.

9. **Double-click Configure Automatic Updates.**

   This step brings up the Configure Automatic Updates dialog box, as shown in Figure 6-8.

10. **Select Enabled to enable the policy.**

11. **Configure the Windows Update settings however you want.**

    For this example, I configure Windows Update so that updates are automatically downloaded every day at 3 a.m. (Figure 6-8 reflects those settings.)

12. **Click OK.**

    You return to the Group Policy Management Editor.

13. **Close the Group Policy Management Editor window.**

    This step returns you to the Group Policy Management settings window.

14. **Right-click Computer Configuration, and choose Refresh from the contextual menu.**

    The Windows Update policy is visible, as shown in Figure 6-9. (To show the full details of the policy, I expanded the Administrative Templates and Windows Components/Windows Update sections of the policy report.)

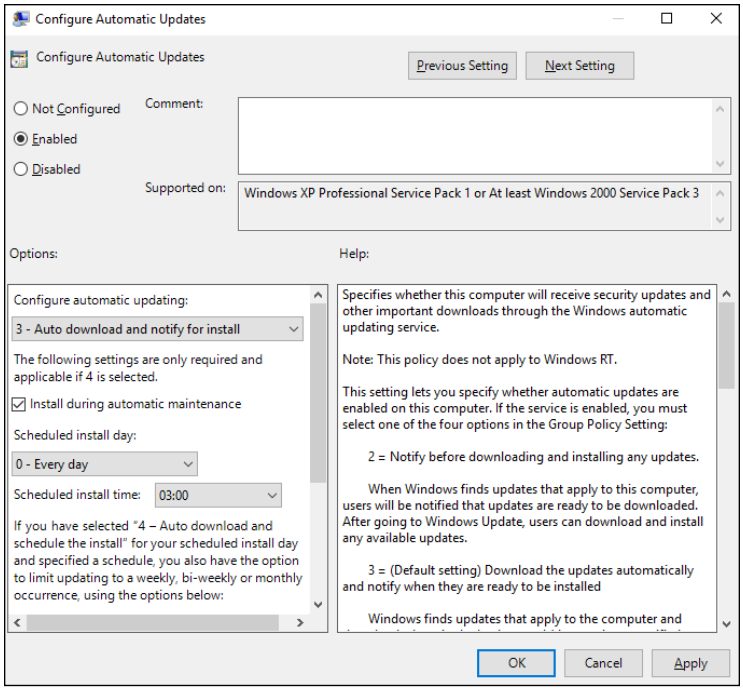


FIGURE 6-8: The Configure Automatic Updates dialog box.

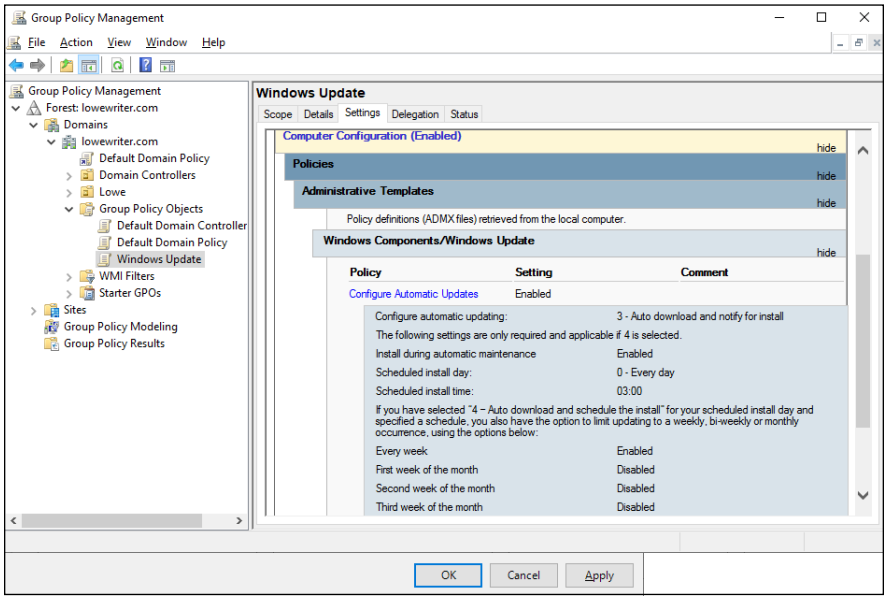


FIGURE 6-9: The Windows Update policy.

15. In the Navigation pane of the Group Policy Management window, drag the new Windows Update policy object to the top-level domain (in this case, `lowewriter.com`).

When you release the mouse button, the dialog box shown in Figure 6-10 appears.

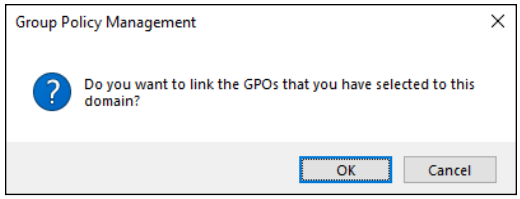


FIGURE 6-10: Confirming the scope.

16. Click OK.

The domain is added to the scope. You can verify this by selecting the Scope tab, as shown in Figure 6-11.

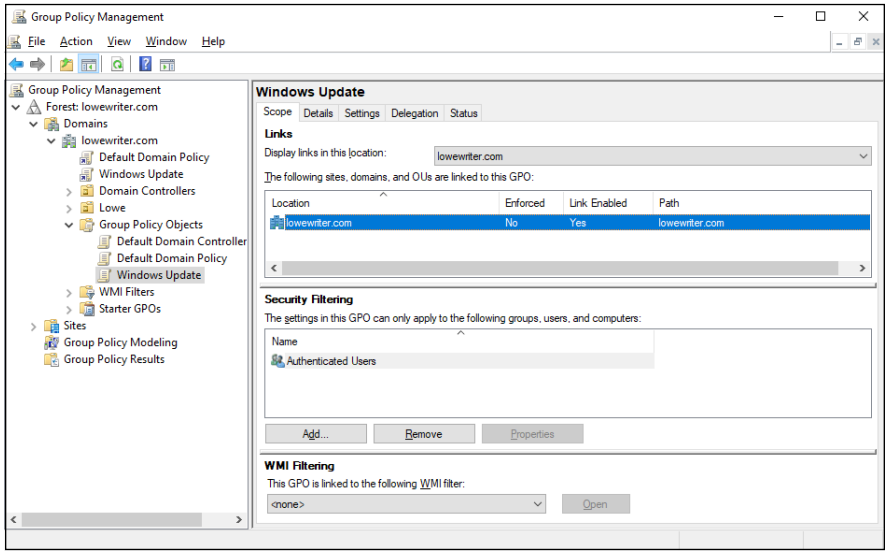


FIGURE 6-11: The policy is finished.

17. Close the Group Policy Management window.

The new group policy is now active.



TIP

You may notice in Figure 6-11 that there is both an Enforced column and a Link Enabled column on the Scope tab. In this example, the Windows Update GPO is linked to the `lowewriter.com` domain, but not enforced. Don't let this confuse you: The GPO is indeed applied to all Active Directory objects in the domain. Here's the distinction between Link Enabled and Enforced:

- **Link Enabled:** When a GPO is linked to an Active Directory container, the settings of the GPO are applied to all objects in the container.
- **Enforced:** This setting provides an extra layer of assurance that the GPO settings will be applied by preventing other GPOs that may also be applied to the object from overriding the settings of this GPO. Enforced is rarely used.

Filtering Group Policy Objects

One of the most confusing aspects of group policy is that even though it applies to users and computers, you don't associate group policy objects with users or computers. Instead, you link GPOs to sites, domains, or organizational units (OUs). At first glance, this aspect may seem to limit the usefulness of group policy. For most simple networks, you'll work with group policy mostly at the domain level and occasionally at the OU level. Site-level group policy objects are used only for very large or complex networks.

Group policy wouldn't be very useful if you had to assign exactly the same policy to every user or computer within a domain. And although OUs can help break down group policy assignments, even that capability is limiting, because a particular user or computer can be a member of only one OU. Fortunately, group policy objects can have *filters* that further refine which users or computers the policy applies to. Although you can filter policy objects so that they apply only to individual users or computers, you're more likely to use groups to apply your group policy objects.

Suppose that you want to use group policy to assign two different default home pages for Internet Explorer. For the Marketing department, you want the default home page to be `www.dummies.com`, but for the Accounting department, you'd like the default home page to be `www.beancounters.com`. You can easily accomplish this task by creating two groups named Marketing and Accounting in Active Directory Users and Computers, and assigning the marketing and accounting users to the appropriate groups. Next, you can create two group policy objects: one for the Marketing department's home page and the other to assign the Accounting department's home page. Then you can link both of these policy objects to the domain and use filters to specify which group each policy applies to.

For the following procedure, I've created two group policies, named Home Page Dummies and Home Page BeanCounters, as well as two Active Directory groups, named Marketing and Accounting. Here are the steps for filtering these policies to link correctly to the groups:

1. **Choose Start➪Administrative Tools➪Group Policy Management.**

   The Group Policy Management console appears. (Refer to Figure 6-2 for a refresher on what it looks like.)

2. **In the Navigation pane, navigate to the group policy object you want to apply the filter to.**

   For this example, I navigated to the IE Home Page BeanCounters policy, as shown in Figure 6-12.

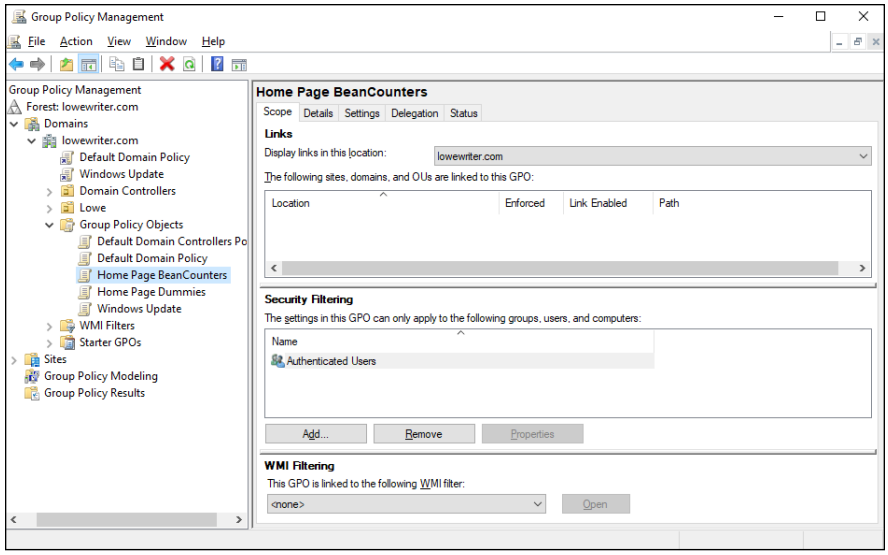


FIGURE 6-12: The IE Home Page Dummies policy.

3. **In the Security Filtering section, click Authenticated Users and then click Remove.**

   This step removes Authenticated Users so that the policy won't be applied to all users.

4. **Click Add.**

   This step brings up the Select User, Computer, or Group dialog box, as shown in Figure 6-13.

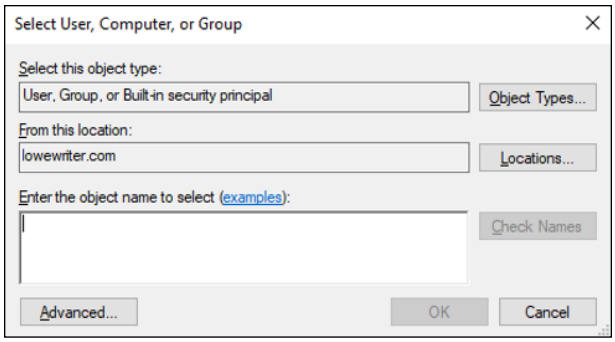


FIGURE 6-13: The Select User, Computer, or Group dialog box.

5. **Type Accounting in the text box and then click OK.**

   The policy is updated to indicate that it applies to members of the Accounting group, as shown in Figure 6-14.

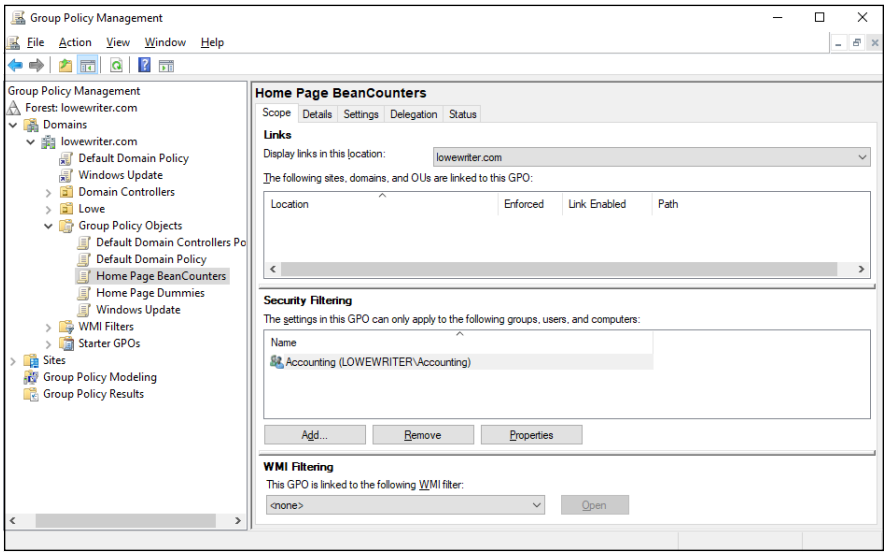


FIGURE 6-14: A policy that uses a filter.

6. **Repeat Steps 2 through 5 for the Home Page Dummies policy, applying it to the Marketing group.**

   You're done!

IN THIS CHAPTER

- Getting started with a command window
- Taking advantage of command tricks and techniques
- Looking at batch files
- Using the amazing Net commands

Chapter 7

Comandeering Windows Commands

Although Windows sports a fancy graphical interface that makes it possible to perform most network management tasks by pointing and clicking, you can also do almost any network management task from a command prompt. Whether you choose to do so is largely a matter of personal style. Some network administrators pride themselves on being able to type Windows commands blindfolded and with two fingers on each hand tied behind their backs. Others have fully embraced the graphical user interface and think the command line is for administrators with Unix envy.

So the choice is yours. Skip this chapter if the thought of typing commands causes you to lose sleep. If you're willing to venture forth, this chapter begins with an overview of working from the command prompt. Then it describes some of the more useful Windows commands. Finally, this chapter introduces the fine (and almost lost) art of writing batch files.


TIP

Windows Server 2025 also includes an alternative command environment known as PowerShell. *PowerShell* is an advanced command processor that has many sophisticated features that are designed especially for creating powerful scripts. For more information, see Book 6, Chapter 8.

Using a Command Window

Command prompts are even older than video monitors. The first computer I worked on used a teletype machine as its terminal, so the command prompt was printed on paper rather than displayed onscreen. Surprisingly, though, the concept of the command prompt hasn't changed much since those days. The system displays a prompt to let you know it's waiting for a command. When you type the command and press the Enter key, the system reads your command, interprets it, executes it, displays the results, and then displays the prompt again so that you can enter another command.

Opening and closing a command window

To get to a command prompt on a Windows server, follow these steps:

1. **Press the Windows key on your keyboard, and then type** cmd**.**
2. **Press the Enter key.**

   The command prompt window appears, as shown in Figure 7-1.

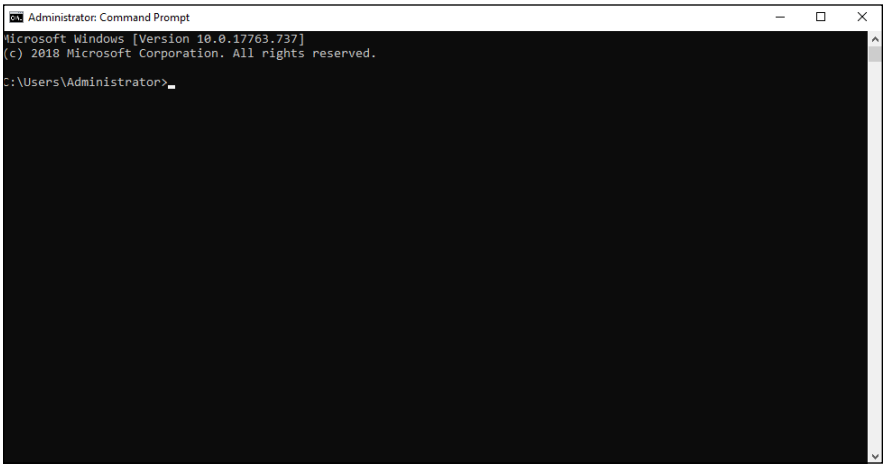


FIGURE 7-1: The command prompt.

You can type any commands you want in the window.



TIP

To exit the command prompt, type **Exit**, and press Enter. This action properly terminates `cmd.exe` and closes the command prompt window. If you try to close the command prompt window by clicking its Close button, Windows is forced to shut down `cmd.exe`. The process works, but you have to click your way

through an intervening dialog box and wait a few seconds while Windows terminates `cmd.exe`. Entering the `Exit` command is a much faster method.


TIP

Sometimes it's helpful to open a command prompt in *elevated mode*, which means that you have full administrative privileges while in the prompt window. To do that, click Start and type **cmd**; then right-click the Command Prompt icon and choose Run As Administrator.

Editing commands

Most of the time, you just type commands by using the keyboard. If you make a mistake, you just retype the command, being careful not to repeat the mistake. `cmd.exe`, however, has several built-in editing features that can simplify the task of correcting a mistaken command or entering a sequence of similar commands:

- Press the right-arrow key to recall the text of the last command that you entered, one letter at a time. When you get to the spot where the new command should differ from the previous command, start typing.
- Press F3 to recall all the previous commands from the current cursor position to the end of the line.
- If you want to repeat a command that you've used recently, press the up-arrow key. This action recalls up to 50 of the most recently executed commands. You can press Enter to execute a command as is, or you can edit the command before you execute it.

Using the Control menu

Although the command window has no menu bar, it does have a menu that you can access via the control box in the top-left corner of the window. Besides the commands found on this menu for all windows (such as Move, Size, and Minimize), this menu includes three additional commands:

- **Edit:** The Edit command leads to a submenu with several choices. Several of these commands work together so that you can copy information from the command window to the clipboard, and vice versa. If you choose Edit ➪ Mark, you're placed in a special editing mode that lets you highlight text in the command window with the mouse. (Normally, the mouse doesn't do anything in the command window.) Then you can choose Edit ➪ Copy or just press Enter to copy the text that you selected to the clipboard.

You can also use the Edit menu to paste text from the clipboard, to scroll the window, and to search the window for text.

- **Default:** This command lets you set default properties for the command window.
- **Properties:** This command displays a Properties dialog box that you can use to change the appearance of the window. You can change the font size, choose background colors, and make other adjustments to make the command window look good on your computer.

Special Command Tricks

Before I get into the details of using specific commands, I want to describe some techniques you should familiarize yourself with. In many cases, these techniques can let you accomplish in a single command what would otherwise take dozens of separate commands.

Wildcards

Wildcards are among the most compelling reasons to use the command prompt. With wildcards, you can process all the files that match a particular naming pattern with a single command. Suppose that you have a folder containing 500 files, and you want to delete all the files that contain the letters *Y2K* and end with the extension `.doc`, which happens to be 50 files. If you open a Documents window, you'll spend ten minutes picking these files out from the list. From a command window, you can delete them all with the single command `del *Y2K*.doc`.

You can use two wildcard characters. An asterisk stands for any number of characters, including zero, and an exclamation point stands for just one character. Thus, `!Text.doc` would match files with names like `aText.doc`, `xText.doc`, and `4Text.doc`, but not `abcText.doc` or just `Text.doc`. `*Text.doc`, however, would match any of the names mentioned in the previous sentence.

Wildcards work differently in Windows than they did in MS-DOS. In MS-DOS, anything you typed after an asterisk was ignored. Thus, `ab*cd.doc` was the same as `ab*.doc`. In Windows, the asterisk wildcard can come before static text, so `ab*cd.doc` and `ab*.doc` are *not* the same.

Chaining commands

You can enter two or more commands on the same line by separating the commands with an ampersand (&), like this:

```
C:\>copy *.doc a: & del *.doc
```

Here, the `copy` command copies all the `.doc` files to the `A:` drive. Then, the `del` command deletes the `.doc` files.

Although that technique may be convenient, it's also dangerous. What if the `A:` drive fills up so that all the files can't be copied? In that case, the `del` command executes anyway, deleting the files that didn't get copied.

A safer alternative is to use two ampersands, telling Windows to execute the second command only if the first command finishes successfully:

```
C:\>copy *.doc a: && del *.doc
```

Now the `del` command will be executed only if the `copy` command succeeds.

You can also use two pipe characters (the *pipe* is the vertical-bar character that's above the backslash on the keyboard) to execute the second command only if the first command fails. Thus,

```
C:\>copy *.doc a: || echo Oops!
```

displays the message `Oops!` if the `copy` command fails.

Finally, you can use parentheses to group commands. Then you can use the other symbols in combination:

```
C:\>(copy *.doc a: && del *.doc) || echo Oops!
```

Here, the files are copied and then deleted if the copy was successful. If either command fails, the message is displayed.

Redirection and piping

Redirection and piping are related techniques. *Redirection* lets you specify an alternative destination for output that will be displayed by a command or an

alternative source for input that should be fed into a command. You can save the results of an `ipconfig /all` command to a file named `myconfig.txt` like this:

```
C:\>ipconfig /all > myconfig.txt
```

Here, the greater-than sign (`>`) is used to redirect the command's console output.

If a command accepts input from the keyboard, you can use input redirection to specify a file that contains the input you want to feed to the command. You can create a text file named `lookup.txt` with subcommands for a command such as `nslookup`. Then you can feed those scripted subcommands to the `nslookup` command, like this:

```
C:\>nslookup < lookup.txt
```

Piping is a similar technique. It takes the console output from one command and feeds it into the next command as input. Piping is often used with special commands called *filters*, which are designed to read input from the console, modify the data in some way, and then write it to the console.

Suppose that you want to display the contents of a file named `users.txt` sorted into alphabetical order. You can use the `Type` command, which displays a file on the console, and then pipe the output into the `Sort` command, a filter that sorts its input and displays the sorted output on the console. The resulting command looks like this:

```
C:\>type users.txt | sort
```

The vertical bar is often called the *pipe character* because it's the symbol used to indicate piping.

Environment variables

The command shell makes several *environment variables* available to commands. Environment variables all begin and end with percent signs. You can use an environment variable anywhere in a command. The command

```
C:\>echo %OS% running on a %PROCESSOR_IDENTIFIER%
```

displays a line such as this:

```
Windows_NT running on an x86 Family 15 Model 2 Stepping 8,
     GenuineIntel
```

Interestingly, later versions of Windows Server all display `Windows_NT` for the operating-system name.

If the environment variable represents a path, you may need to enclose it in quotation marks, like this:

```
C:\>dir "%HOMEPATH%"
```

This command displays the contents of the user's home directory. The quotation marks are required here because the environment variable expands to a pathname that may include spaces, and the command shell requires that long filenames that include spaces be enclosed in quotation marks.

Table 7-1 lists the environment variables that are available to you and your commands.

TABLE 7-1 Environment Variables

Variable	Description
%ALLUSERSPROFILE%	The location of the All Users profile
%APPDATA%	The path where applications store data by default
%CD%	The path to the current directory
%CMDCMDLINE%	The command line that was used to start the command shell
%CMDEXTVERSION%	The version number of the command shell
%COMPUTERNAME%	The computer's name
%COMSPEC%	The path to the command shell executable (cmd.exe)
%DATE%	The current date in the format generated by the date /t command
%ERRORLEVEL%	The error returned by the most recent command
%HOMEDRIVE%	The drive letter of the user's home directory
%HOMEPATH%	The path to the user's home directory
%HOMESHARE%	The network path to the user's shared home directory
%LOGONSERVER%	The name of the domain controller the user logged on to
%NUMBER_OF_PROCESSORS%	The number of processors on the computer
%OS%	The name of the operating system

(continued)

TABLE 7-1 *(continued)*

Variable	Description
%PATH%	The current search path
%PATHEXT%	A list of the extensions the operating system treats as executable files
%PROCESSOR_ARCHITECTURE%	The chip architecture of the processor
%PROCESSOR_IDENTIFIER%	A description of the processor
%PROCESSOR_REVISION%	The revision level of the processor
%PROMPT%	The current prompt string
%RANDOM%	A random number between 1 and 32,767
%SYSTEMDRIVE%	The drive containing the operating system
%SYSTEMROOT%	The path to the operating system
%TEMP%	The path to a temporary folder for temporary files
%TMP%	Same as %TEMP%
%TIME%	The time in the format produced by the time /t command
%USERDOMAIN%	The name of the user's domain
%USERNAME%	The user's account name
%USERPROFILE%	The path to the user's profile
%WINDIR%	The path to the operating-system directory

Batch files

A *batch file* is simply a text file that contains one or more commands. Batch files are given the extension .bat and can be run from a command prompt as though they were commands or programs. You can also run a batch file from the Start menu by choosing Start ➪ Run, typing the name of the batch file, and clicking OK.

As a network administrator, you'll find plenty of uses for batch files. Most of them won't be very complicated. Here are some examples of very simple batch files I've used:

» I once used a one-line file to copy the entire contents of an important shared network drive to a user's computer every night at 10 p.m. The user wanted a quick-and-dirty backup solution that would complement the regular tape backups that ran every night.

- I've also used a pair of short batch files to stop and then restart an Exchange server before and after nightly backups.
- If I frequently need to work with several related folders at the same time, I create a short batch file that opens Explorer windows for each of the folders. (You can open an Explorer window from a batch file simply by typing the path to the folder that you want to open as a command.) Then I place the batch file on my desktop so that I can get to it quickly.

You can also use batch files to create logon scripts that are executed whenever a user logs on. Microsoft keeps trying to get users to use profiles instead of logon scripts, but many networks still use logon scripts.

The EventCreate Command

The `EventCreate` command lets you create an event that's added to one of the Windows event logs. This command can be useful if you want to make a note of something unusual that's happened. It's often used in batch files to mark the start or completion of a task such as a nightly backup.

Here's the basic syntax:

```
eventcreate [options]
eventcreate /T type /D "description" /ID eventide
[/L logname] [/SO sourcename]
[/S system [/U username [/P password]]]
```

Here's a description of the options:

- `/T`: Specifies the type. The options are `Information`, `Warning`, and `Error`.
- `/D`: Provides a descriptive message that's saved in the log. Use quotes if the message contains more than one word.
- `/ID`: A number from 1 to 1,000.
- `/L`: The name of the log to write the event to. The default is `Application`.
- `/SO`: A string that represents the source of the event. The default is `EventCreate`. If you specify this option, you must also specify the `/L` option.
- `/S`: The name of the system on which the event should be recorded.

- /U: The user account to use when logging the event. You can specify this option only if you also specify /S.
- /P: The password. You can specify this option only if you also specify /U.

Here's an example that writes an informational message to the Application log:

```
eventcreate /t information /id 100 /d "Nightly processing
     completed" /L Application /SO Nightly
```

Figure 7-2 shows an event created by the preceding command.

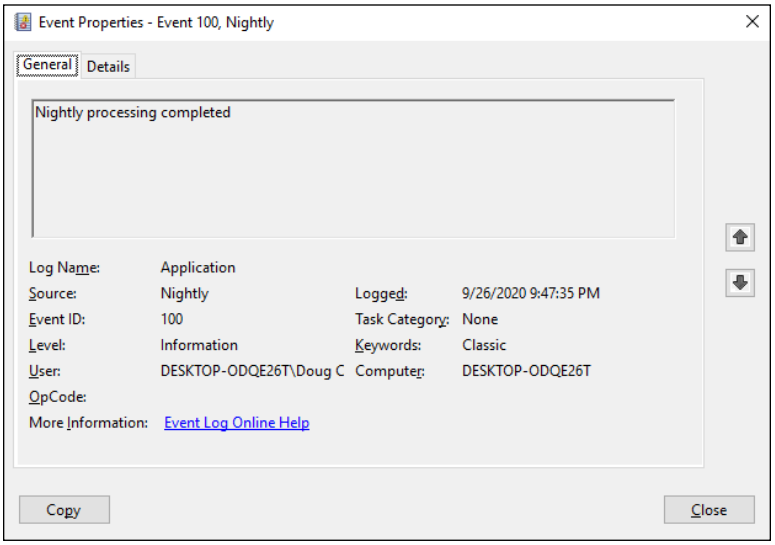


FIGURE 7-2: An event generated by the EventCreate command.

Net Commands

Among the most useful commands for network administrators are the Net Services commands. All these commands are two-word commands beginning with Net — such as Net Use and Net Start. In the following sections, I present the Net commands in alphabetical order for handy reference. First, though, I want to point out a few details about the Net commands:

- You can get a quick list of the available Net commands by typing **net /?** at a command prompt.

- You can get brief help for any `Net` command by typing **net help *command***. To display help for the `Net Use` command, for example, type **net help use**. (Yes, we all could use some help.)
- Many of the `Net` commands prompt you for confirmation before completing an operation. For these commands, you can specify `/Y` or `/N` to bypass the confirmation prompt. You'll want to do that if you include these commands in a batch file that runs unattended. Note that you can use `/Y` or `/N` on any `Net` command, even if it doesn't prompt you for confirmation. So I suggest that you place `/Y` on every `Net` command in a batch file that you intend to run unattended.

The Net Accounts command

This command updates user account policies for password requirements. Here's the command syntax:

```
net accounts [/forcelogoff:{minutes | no}]
    [/minpwlen:length] [/maxpwage:{days | unlimited}]
    [/minpwage:days] [/uniquepw:number]
    [/domain]
```

The following paragraphs describe the parameters for the `Net Accounts` command:

- `forcelogoff`: Specifies how long to wait before forcing a user off the system when the user's logon time expires. The default value, `no`, prevents users from being forced to log off. If you specify a number, the user will be warned a few minutes before being forcibly logged off.
- `minpwlen`: Specifies the minimum length for the user's password. *`Length`* can be 0 through 127. The default is 6.
- `maxpwage`: Specifies the number of days a user's password is considered to be valid. `Unlimited` means that the password will never expire. *`Days`* can be from 1 through 49,710, which is about 135 years. The default is 90.
- `minpwage`: Specifies the minimum number of days after a user changes a password before the user can change it again. The default value is 0. You usually should set this value to 1 day to prevent users from bypassing the `Uniquepw` policy.
- `uniquepw`: Indicates how many different passwords the user must use before he or she is allowed to reuse the same password. The default setting is 5. The range is 0 through 24.
- `domain`: Specifies that the operation should be performed on the primary domain controller rather than on the local computer.

If you enter `Net Accounts` without any parameters, the command simply displays the current policy settings.

Here's an example that sets the minimum and maximum password ages:

```
C:\>net accounts /minpwage:7 /maxpwage:30
```

The Net Computer command

This command creates or deletes a computer account. Here's the syntax:

```
net computer \\computername {/add | /del}
```

The following paragraphs describe the parameters for the `Net Computer` command:

- *Computername*: Specifies the computer to add or delete
- `add`: Creates a computer account for the specified computer
- `del`: Deletes the specified computer account

Here's an example that adds a computer named Theodore:

```
C:\>net computer \\theodore /add
```

The Net Config command

This command lets you view or configure various network services. Here's the syntax:

```
net config [{server|workstation}] [options]
```

To configure server settings, use this syntax:

```
net config server [/autodisconnect:time] [/srvcomment:"text"]
     [/hidden:{yes | no}]
```

The following paragraphs describe the parameters for the `Net Config` command:

- `server`: Lets you display and configure the Server service while it's running.
- `workstation`: Lets you display and configure the Workstation service while it's running.

- `autodisconnect`: Specifies how long a user's session can be inactive before it's disconnected. Specify –1 to never disconnect. The range is –1 to 65,535 minutes, which is about 45 days. The default is 15 minutes.
- `srvcomment`: Specifies a description of the server. The comment can be up to 48 characters long and should be enclosed in quotation marks.
- `hidden`: Specifies whether the server appears in screens that display available servers. Hiding a server doesn't make the server unavailable; it just means that the user will have to know the name of the server to access it. The default is `No`.

Here's an example that sets a server's descriptive comment:

```
C:\>net config server /srvcomment:"DHCP Server"
```

The Net Continue command

This command continues a service you've suspended with the `net pause` command. Here's the syntax:

```
net continue service
```

Here are some typical services that you can pause and continue:

- `netlogon`: The Net Logon service.
- `schedule`: The Task Scheduler service.
- `server`: The Server service.
- `workstation`: The Workstation service.

Here's an example that continues the Workstation service:

```
C:>net continue workstation
```

If the service name has embedded spaces, enclose the service name in quotation marks. This command continues the NT LM Security Support Provider service:

```
C:\>net continue "nt lm security support provider"
```

The Net File command

This command lists all open shared files and the number of file locks placed on each file. You can also use this command to close files and remove locks, which is a useful procedure when a user manages to accidentally leave a file open or locked. Here's the syntax:

```
C:\>net file [id [/close]]
```

The following paragraphs describe the `Net File` command's parameters:

- *id*: The file's identification number.
- `close`: Closes an open file and releases any locks that were placed on the file.

To close a file, you must issue the command from the server where the file is shared.



TIP

`net files` works, too.

To close an open file, first run `net file` without any parameters to list the open files. Here's a sample of the output that you can expect from `net file`:

```
File Path Username #locks
--------------------------------------------
0 C:\BUDGET.DOC WARD 0
1 C:\RECIPE.MDF JUNE 4
```

Next, run `net file` again, specifying the file number displayed for the file that you want to close. To close the `RECIPE.MDF` file, for example, use this command:

```
C:\>net file 1 /close
```

The Net Group command

This command lets you add, display, or change global groups. This command has several different syntaxes, depending on how you intend to use it.

To display information about a group or to change a group's comment, use this syntax:

```
net group groupname [/comment:"text"] [/domain]
```

To create a new group, use this syntax:

```
net group groupname /add [/comment:"text"] [/domain]
```

To delete a group, use this syntax:

```
net group groupname /delete [/domain]
```

Finally, to add or remove users from a group, use this syntax:

```
net group groupname username[ ...] {/add | /delete} [/domain]
```

The following paragraphs describe the parameters that you can use with the `net group` command:

- *groupname*: Specifies the name of the group to add, change, or delete. If you specify this parameter and no others, a list of users in the group appears.
- `comment`: Specifies a comment for the group. The comment can be up to 48 characters in length and should be enclosed in quotation marks.
- `domain`: Specifies that the operation should be performed on the primary domain controller rather than on the local computer.
- `add`: Creates a new group or adds users to an existing group. Before you add a user to a group, you must create a user account for the user.
- `delete`: Removes a group or removes users from the group.
- *username*: Specifies one or more usernames to be added to or removed from the group. If you list more than one name, separate the names with spaces.



TIP

Windows isn't picky: You can specify `net groups` rather than `net group` if you want.

This example lists all the groups on a server:

```
C:\>net group
```

This example adds a group named Admin:

```
C:\>net group Admin /add
```

This example adds three users to the Admin group:

```
C:\>net group Admin Ward Wally June /add
```

This example lists the users in the Admin group:

```
C:\>net group Admin
```

The Net Help command

This command displays help for the `net` command or for a specific `net` subcommand. Here's the basic syntax:

```
net help [command]
```

The *command* parameter can be any of the following commands:

```
accounts
computer
config
continue
file
group
help
helpmsg
localgroup
pause
session
share
start
statistics
stop
time
use
user
view
```



TIP

You can type **net help services** to display a list of services that you can start via the `Net Start` command.

The Net Helpmsg command

This command displays an explanation of network error codes. Here's the syntax:

```
net helpmsg message#
```

The *message#* parameter should be the four-digit number displayed when the error occurred. If you get an error with message 2180, for example, use this command to see an explanation of the error:

```
C:\>net helpmsg 2180
The service database is locked.
EXPLANATION
Another program is holding the service database lock.
ACTION
Wait for the lock to be released and try again later. If it is
     possible to determine which program is holding the lock,
     then end that program.
```


The Net Localgroup command

This command lets you add, display, or change local groups. This command has several different syntaxes, depending on how you intend to use it.

To display information about a local group or to change a local group's comment, use this syntax:

```
net localgroup groupname [/comment:"text"] [/domain]
```

To create a new group, use this syntax:

```
net localgroup groupname /add [/comment:"text"] [/domain]
```

To delete a group, use this syntax:

```
net localgroup groupname /delete [/domain]
```

Finally, to add users to or remove users from a group, use this syntax:

```
net localgroup groupname username[ ...] {/add | /delete}
     [/domain]
```

The following paragraphs describe the parameters that you can use with the `net localgroup` command:

- *groupname*: Specifies the name of the group to add, change, or delete. If you specify this parameter and no others, a list of users in the group appears.
- `comment`: Specifies a comment for the group. The comment can be up to 48 characters in length and should be enclosed in quotation marks.

- `domain`: Specifies that the operation should be performed on the primary domain controller rather than on the local computer.
- `add`: Creates a new group or adds users to an existing group. Before you add a user to a group, you must create a user account for the user.
- `delete`: Removes a group or removes users from the group.
- *`username`*: Specifies one or more usernames to be added to or removed from the group. If you list more than one name, separate the names with spaces.

This example lists all the local groups:

```
C:\>net localgroup
```

This example adds a local group named Admin:

```
C:\>net localgroup Admin /add
```

This example adds three users to the Admin local group:

```
C:\>net localgroup Admin Ward Wally June /add
```

This example lists the users in the Admin group:

```
C:\>net localgroup Admin
```

The Net Pause command

This command temporarily pauses a service. It's a good idea to pause a service for a while before you stop the service altogether. That gives users who are currently using the service a chance to finish any pending tasks, while at the same time preventing other users from beginning new sessions with the service. To reactivate the service later, use the `net continue` command.

The syntax to pause a service is

```
net pause service
```

Here are some typical services that you can pause:

- `netlogon`: The Net Logon service
- `schedule`: The Task Scheduler service

- `server`: The Server service
- `workstation`: The Workstation service

Here's an example that pauses the Workstation service:

```
CL>net pause workstation
```

If the service name has embedded spaces, enclose the service name in quotation marks. This command pauses the NT LM Security Support Provider service, for example:

```
C:\>net pause "nt lm security support provider"
```


The Net Session command

This command lets you view current server connections and kick users off, if you feel inclined. Here's the syntax:

```
net session [\\ComputerName] [/delete]
```

Here's what the parameters do:

- *computerName*: Indicates which computer's session you want to view or disconnect. If you omit this parameter, all sessions are listed.
- `delete`: Disconnects the computer's session. Any open files are immediately closed. If you use this parameter without specifying a computer name, all computers currently connected to the server are disconnected.

  This command is an obviously dangerous one. If you disconnect users while they're updating files or before they have a chance to save their work, they'll be hopping mad.


WARNING

To find out who is connected to a computer, use this command:

```
C:\>net session
Computer User name Client type Opens Idle time
-------------------------------------------------------------
\\DEN Ward Windows XP 1 00:00:4
\\BEDROOM Administrator Windows 2008 0 02:15:17
```

The Net Share command

This command lets you manage shared resources. To display information about all shares or a specific share, use this syntax:

```
net share [ShareName]
```

To create a new share, use this syntax:

```
net share ShareName=path [{/users:number|/unlimited}]
     [/remark:"text"] [/cache: {manual|automatic|no}]
```

To change the properties of an existing share, use this syntax:

```
net share ShareName [{/users:number|unlimited}] [/remark:"text"]
     [/cache: {manual|automatic|no}]
```

To delete an existing share, use this syntax:

```
net share {ShareName|drive:path} /delete
```

Here's what the parameters do:

- *ShareName*: Specifies the share name. Use this parameter by itself to display information about the share.
- *path*: Specifies the path to the folder to be shared. The path should include a drive letter. If the path includes spaces, enclose it in quotation marks.
- `users`: Specifies how many users can access the share concurrently.
- `unlimited`: Specifies that an unlimited number of users can access the share concurrently.
- `remark`: Creates a descriptive comment for the share. The comment should be enclosed in quotation marks.
- `cache`: Specifies the caching option for the share.
- `delete`: Stops sharing the folder.

If you use `net share` without any parameters, all the current shares are listed, as shown in this example:

```
Share name    Resource                        Remark
-------------------------------------------------------------
C$            C:\                             Default share
IPC$                                          Remote IPC
ADMIN$        C:\WINDOWS                      Remote Admin
Users         C:\Users
The command completed successfully.
```

The following example creates a share named Docs:

```
C:\>net share Docs=C:\SharedDocs /remark:"Shared documents"
```


The Net Start command

This command lets you start a networking service or display a list of all the services that are currently running. The syntax is

```
net start [service]
```

In most cases, you'll use this command to start a service that you've previously stopped with the `net stop` command. In that case, you should first run the `net start` command without any parameters to find the name of the service that you want to stop. Make a note of the exact spelling of the service that you want to stop. Then use the `net stop` command to stop the service. When you want to restart the service, use the `net start` command again — this time specifying the service to start.

Suppose that you need to stop your DNS server. Using `net start`, you discover that the name of the service is DNS Server, so you use the following command to stop it:

```
C:\>net stop "DNS Server"
```

Later, you can use this command to restart the service:

```
C:\>net start "DNS Server"
```

The Net Statistics command

This command lists the statistics log for the local Workstation or Server service. The syntax is

```
net statistics [{workstation | server}]
```

You can specify `workstation` or `server` to indicate the service for which you'd like to view statistics.

If you use `net statistics workstation`, the following information appears:

- The computer name
- The date and time when the statistics were last updated
- The number of bytes and server message blocks (SMB) received and transmitted
- The number of read and write operations that succeeded or failed
- The number of network errors
- The number of sessions that failed, disconnected, or were reconnected
- The number of connections to shared resources that succeeded or failed

If you use `Net Statistics Server`, the following information is listed:

- The computer name
- The date and time when the statistics were last updated
- The number of sessions that have been started, disconnected automatically, and disconnected because of errors
- The number of kilobytes sent and received, and the average response time
- The number of password and permission errors and violations
- The number of times the shared files, printers, and communication devices were used
- The number of times the size of the memory buffer was exceeded

The Net Stop command

This command lets you stop a networking service. The syntax is

```
net stop service
```

To use this command, first run the `net start` command to determine the exact spelling of the service that you want to stop. If the service name includes spaces, enclose it in quotation marks.

You can restart the service later by using the `net start` command.

The following example stops the DNS service:

```
C:\>net stop "DNS Server"
```

The Net Time command

This command synchronizes the computer's clock with the clock on another computer. To access a clock on another computer in the same domain or workgroup, use this form:

```
net time \\ComputerName [/set]
```

To synchronize time with a domain, use this form:

```
net time /domain[:DomainName] [/set]
```

To use an RTS time server, use this syntax:

```
net time /rtsdomain[:DomainName] [/set]
```

To specify the computer to use for Network Time Protocol, use this syntax:

```
net time [\\ComputerName] [/querysntp] [/
   setsntp[:NTPServerList]]
```

To set the computer's clock to match the Server01 clock, use this command:

```
C:\>net time \\Server01 /set
```

The Net Use command

This command connects to or disconnects from a shared resource on another computer and maps the resource to a drive letter. Here's the complete syntax:

```
net use [{drive | *}]
[{\\computername\sharename]
[{password | *}]]
[/user:[domainname\]username]
[/savecred]
[/smartcard]
[{/delete | /persistent:{yes | no}}]
```

To set up a home directory, use this syntax:

```
net use [drive [/home[{password | *}]
[/delete:{yes | no}]]
```

And to control whether connections should be persistent, use this:

```
net use [/persistent:{yes | no}]
```

Here's what the parameters do:

- *drive*: Specifies the drive letter. (Note that for a printer, you should specify a printer device such as LPT1: here instead of a drive letter.) If you specify an asterisk, Windows will determine what drive letter to use.
- *computername**sharename*: Specifies the server and share name to connect to.
- *password*: Provides the password needed to access the shared resource. If you use an asterisk, you're prompted for the password.
- user: Specifies the username to use for the connection.
- savecred: Saves the credentials for reuse later if the user is prompted for a password.
- smartcard: Specifies that the connection should use a smart card for authorization.
- delete: Deletes the specified connection. If you specify an asterisk (*), all network connections are canceled.
- persistent: Specifies whether connections should be persistent.
- home: Connects to the home directory.

To display all current connections, type **net use** with no parameters.

The following example shows how to create a persistent connection to a drive named Acct on a server named Server01, using drive K::

```
C:\>net use k: \\Server01\Acct /persistent: yes
```

The following example drops the connection:

```
C:\>net use k: /delete
```

The Net User command

This command creates or changes user accounts. To display a user's information, use this form:

```
net user username
```

To update user information, use this form:

```
net user [username [password | *] [options]] [/domain]
```

To add a new user, use this form:

```
net user username [password | *] /add [options] [/domain]
```

To delete a user, use this form:

```
net user username /delete [/domain]
```

Most of the parameters for this command are straightforward. The `options` parameters, however, can have a variety of settings. Table 7-2 lists the descriptions of these options as presented by the `Net Help Users` command.

TABLE 7-2 The Options Parameters

Options	Description
/ACTIVE:{YES | NO}	Activates or deactivates the account. If the account isn't active, the user can't access the server. The default is YES.
/COMMENT:"*text*"	Provides a descriptive comment about the user's account (maximum of 48 characters). Enclose the text in quotation marks.
/COUNTRYCODE:*nnn*	Uses the operating-system country code to implement the specified language files for a user's help and error messages. A value of 0 signifies the default country code.
/EXPIRES:{*date* | NEVER}	Causes the account to expire if date is set. NEVER sets no time limit on the account. An expiration date is in the form *mm/dd/yy* or *dd/mm/yy*, depending on the country code. The month can be a number, spelled out, or abbreviated with three letters. The year can be two or four numbers. Use slashes (/), not spaces, to separate parts of the date.
/FULLNAME:" *name*"	Is a user's full name (rather than a username). Enclose the name in quotation marks.
/HOMEDIR:*pathname*	Sets the path for the user's home directory. The path must exist.

(continued)

TABLE 7-2 *(continued)*

Options	Description
/PASSWORDCHG:{YES | NO}	Specifies whether users can change their own passwords. The default is YES.
/PASSWORDREQ:{YES | NO}	Specifies whether a user account must have a password. The default is YES.
/PROFILEPATH[:*path*]	Sets a path for the user's logon profile.
/SCRIPTPATH:*pathname*	Is the location of the user's logon script.
/TIMES:{times | ALL}	Is the logon hours. TIMES is expressed as *day*[-*day*][,*day*[-*day*]],*time*[-*time*] [,*time*[-*time*]], limited to one-hour increments. Days can be spelled out or abbreviated. Hours can be 12- or 24-hour notation. For 12-hour notation, use am or pm (without periods) or a.m. or p.m. ALL means that a user can always log on, and a blank value means that a user can never log on. Separate day and time entries with a comma, and separate multiple day and time entries with a semicolon.
/USERCOMMENT:"*text*"	Lets an administrator add or change the User Comment for the account.
/WORKSTATIONS:	Lists as many as eight computers from which a user {*ComputerName*[,...] | *} can log on to the network. If /WORKSTATIONS has no list or if the list is *, the user can log on from any computer.

To display information for a particular user, use the command like this:

```
C:\>net user Doug
```

To add a user account for Theodore Cleaver with the username Beaver, use this command:

```
C:\>net user Beaver /add /fullname:"Theodore Cleaver"
```

The Net View command

This command displays information about your network. If you use it without parameters, it displays a list of the computers in your domain. You can use parameters to display resources that are being shared by a particular computer. Here's the syntax:

```
net view [\\computername] [/domain[:domainname]]
net view /network:nw [\\computername]
```

Here's what the parameters do:

- *computername*: Specifies the computer whose shared resources you want to view.
- *domainname*: Specifies the domain you want to view, if it's other than the current domain.

Here's typical output from a `net view` command:

```
C:\>net view
Server Name Remark
------------------------------------------------
\\Server01 Main file server
\\Print01 Main print server
```


The RunAs Command

The `runas` command lets you run a program from a command prompt by using the credentials of another user account. Here's the basic syntax:

```
runas /user:username [other parameters] program
```

To run the Microsoft Management Console with the dom1 domain's administrator account, for example, you can use this command:

```
runas /user:dom1\administrator mmc
```

Assuming that the username is valid, you'll be prompted for the user's password. Then the program will be run using the specified user's account.

Here are some of the parameters you can use with the `RunAs` command:

- `/user`: Specifies the domain and username. You can use either of two forms to specify the domain and username: `domain\username` or `user-name@domain`.
- `/profile`: Specifies that the user's profile should be loaded. (This option is on by default, so you don't have to specify it explicitly.)
- `/noprofile`: Doesn't load the user's profile. Although this parameter can cause the application to load faster, it can also prevent some applications from functioning properly.

- `/env`: Uses the current environment instead of the user's.
- `/netonly`: Indicates that the user account isn't valid in the current domain. (If you use `/netonly`, the username must be specified in the form `domain\username`; the `username@domain` form won't work.)
- `/savecred`: Saves the password so that it has to be entered only the first time the RunAs command is used.

  
  WARNING

  Using the `/savecred` parameter is an extremely bad idea, as it creates a gaping security hole. In short, after you've used `/savecred`, any user at the computer can use the `RunAs` command to run any program with administrator privileges.
- `/smartcard`: Specifies that the user's credentials will be supplied by a smart card device.

IN THIS CHAPTER

- Getting started with PowerShell
- Using PowerShell cmdlets
- Working with parameters
- Understanding the pipeline
- Using scripts

Chapter 8
Using PowerShell

In the preceding chapter, you learn how to use a variety of Windows commands from a standard command prompt to perform various Windows administrative chores. In this chapter, you learn how to use a significantly more advanced command-line interface called *PowerShell.* PowerShell is to the Windows command prompt what a Tesla is to a Model A. Both cars look good, but the Model A was popular 90 years ago and took forever to get to its top speed of 65 miles per hour. The Tesla will get you to 155 miles per hour in less than 30 seconds.

Truth be told, I'd rather drive the Model A. But that's because I like things that are old like me. For real command-line performance, however, I suggest you spend some time learning PowerShell. You won't regret it.


TIP

This short chapter can't possibly cover everything there is to know about PowerShell. For more information, see Microsoft's PowerShell site at `www.microsoft.com/powershell`.

Using PowerShell

PowerShell runs in a command window that's very similar to the standard Windows command prompt. However, the procedure to open it is a bit different:

1. **Press the Windows key on your keyboard, and then type** PowerShell**.**
2. **Press the Enter key.**

   The PowerShell window appears, as shown in Figure 8-1.

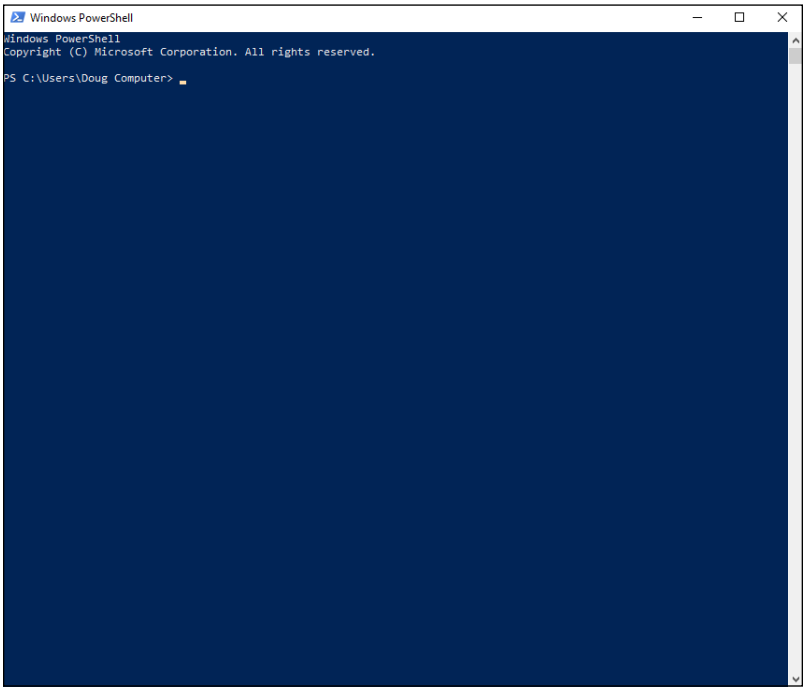


FIGURE 8-1: The PowerShell window.



TIP

If you want to perform administrative functions while in PowerShell, you'll need to open PowerShell as an administrator. To do that, right-click the Windows PowerShell icon and choose Run As Administrator.



WARNING

PowerShell is an incredibly powerful administrative tool, and opening a PowerShell window as an administrator can be risky if you aren't sure what you're doing. I suggest that while you're learning PowerShell, you do so on your desktop computer rather than on one of your servers. At least if you make serious mistakes

on your own desktop computer, you won't bring down everyone else! (I'm half joking here. As a network administrator, you know how to be careful. But it's a good idea to experiment on your own computer or one you've designated as a sandbox rather than on a critical server!)

As with a standard command shell, you can type any commands you want in the window. However, you'll notice a few differences right off the bat:

- Though it's not apparent in the figure, the background for a PowerShell window is blue rather than black.
- The title bar indicates that you're in Windows PowerShell.
- The welcome banner at the top of the window announces that you're using PowerShell.
- Although the command prompt is similar to the prompt for a standard command shell, the current directory is prefixed with *PS* to remind you that you're in PowerShell.


TIP

You can formally exit PowerShell by typing **exit** and pressing Enter. Or, you can just close the window.

You can enter and edit commands within a PowerShell window pretty much the same as you do in a standard command shell. However, one difference you'll find useful is that the Tab key serves as an auto-complete feature: If you type a partial command and then press the Tab key, Windows will try to finish the command for you. Because most PowerShell commands are a bit long compared to their standard Windows command counterparts, this can be a real timesaver.

Give it a try: Open a PowerShell window and then type **get-r** and press the Tab key. PowerShell automatically expands this to the first standard PowerShell command that starts with `get-r`:

```
PS C:\Users\Doug> Get-Random
596196043
PS C:\Users\Doug>
```

As you can see, the shell expands your text to `Get-Random`, a PowerShell command that returns a random number. (Run this command several times; you'll see a different number each time.)

Understanding PowerShell Commands

PowerShell supports four distinctly different kinds of commands you can run directly from the PowerShell prompt:

- **Native commands:** A *native command* is a standard Windows command that you would ordinarily run at a normal command prompt. Traditional commands such as `xcopy`, `ipconfig`, and `ping` can all be run from a PowerShell prompt.
- **Cmdlets:** A *cmdlet* (which is short for *commandlet*) is the basic built-in PowerShell command. All cmdlet names follow a consistent *verb-noun* format. In the example in the preceding section, the verb is `Get` and the noun is `Random`. (For more information about cmdlets, see the section "Using Cmdlets" later in this chapter.)
- **Scripts**: A *script* is a collection of PowerShell commands saved to a text file with the extension `.ps1`. You can run a script at a PowerShell prompt simply by typing the name of the script, without the extension. In short, scripts are PowerShell's equivalent for batch files. (For more information about scripts, see the section "Using Scripts" later in this chapter.)
- **Functions:** A *function* is a set of PowerShell commands that you give a name. Then you can run the named set of commands using the function's name as if it were a command. Here's a simple example that creates and then calls a function:

```
PS C:\Users\Doug> function rnd {Get-Random}
PS C:\Users\Doug> rnd
948203949
PS C:\Users\Doug>
```

  The first command entered above creates a function named `rnd`. The list of commands to be executed for the function is enclosed within curly braces; in this example, just a single command is used (`Get-Random`). This line creates a function named `rnd` that runs the `Get-Random` cmdlet.

  The second command calls the `rnd` function. As you can see, PowerShell responds to the `rnd` function by displaying another random number.

  Note that functions are usually used within scripts. And because they're a somewhat advanced topic, I won't be covering them further in this chapter. (For more information, you can refer to Microsoft's PowerShell website at `www.microsoft.com/powershell`.)

Using Cmdlets

Cmdlets are the bread and butter of PowerShell. At first glance, cmdlets seem similar to native commands, but actually they're quite different. The most obvious difference is how they're named. All cmdlets are named using a simple *verb-name* convention, where the first word is one of several standardized verbs (such as *get*, *create*, or *show*) and the noun is a somewhat less standardized name of the object that the verb acts on. For example, in the `Get-Random` cmdlet, the verb is `Get` and the noun is `Random`.

Nouns in PowerShell are always singular. For example, the command that retrieves a list of all system services is called `get-service`, not `get-services`.

All cmdlets follow this naming convention, which makes it easy to remember cmdlet names, at least once you work with PowerShell long enough to learn the most common verb and noun names.

You can see a listing of all the verbs by running the `get-verb` cmdlet (again, singular: `get-verb`, not `get-verbs`). The `get-verb` command displays a list of 98 different verbs that can be used in cmdlets — too many to show here. Run the command at a PowerShell prompt to get a feel for the types of verbs that are used in cmdlets.

Incidentally, PowerShell names are not case-sensitive. Thus, you don't have to capitalize the verbs and nouns when you type PowerShell commands. `Get-Random` and `get-random` have exactly the same effect.

Using Parameters

Most cmdlets let you use *parameters* that allow you to customize the behavior of the cmdlet. Parameter names are preceded by a hyphen and followed by the parameter value. For example:

```
PS C:\Users\Doug> Get-Random -Minimum 1 -Maximum 10
2
PS C:\Users\Doug>
```

This cmdlets returns a random number between 1 and 10; in the above example, the number returned happens to be 2, but each time you run the cmdlet you'll get a different random value.

Many parameters have default values, so if you omit a parameter, the default value is used. For the `Get-Random` cmdlet, the default value for `-Minimum` is 1, so you can omit it, as in this example:

```
Get-Random -Maximum 10
```

Some parameters can accept two or more values. In that case, you simply separate the values by commas.

You don't always have to type the full name of a parameter; PowerShell will do its best to figure out which parameter you intend. For example, the following command works:

```
Get-Random 1 -Max 10
```

In fact, this command works, too, because `Ma` is enough to distinguish `-Maximum` from `-Minimum`:

```
Get-Random -Ma 10
```

However, the following command doesn't work:

```
Get-Random -M 10
```

Here, PowerShell can't tell whether you mean to use the `-Minimum` parameter or the `-Maximum` parameter.

Some parameters don't have values; in that case, you just list the parameter name without a subsequent value, as in this example:

```
PS C:\Users\Doug> Get-ChildItem -recurse
```

This cmdlet, `Get-ChildItem`, is PowerShell's equivalent to the `dir` command: It lists the contents of the current directory. The `-recurse` parameter tells `Get-ChildItem` to list not just the current directory, but all subdirectories as well.

PowerShell defines a set of *common parameters* that work in a consistent way across many different cmdlets. These common parameters are listed in Table 8-1. Note that not all cmdlets implement all the common parameters. But the point is that when a cmdlet provides the feature indicated by one of these common parameters, the name of the common parameter will be used. (Of special interest is the `-whatif` parameter, which lets you check out what a complicated cmdlet will do before you actually run it.)

TABLE 8-1 PowerShell Common Parameters

Parameter	What It Does
`-WhatIf`	Displays a message that indicates what the cmdlet will do without actually doing anything.
`-Confirm`	Prompts the user before proceeding.
`-Verbose`	Displays additional explanatory information.
`-ErrorAction`	Indicates what to do if an error occurs. Possible actions are `Continue`, `Ignore`, `Inquire`, `SilentlyContinue`, `Stop`, and `Suspend`.
`-ErrorVariable`	Provides the name of a variable used to hold error information.
`-WarningAction`	Indicates what to do if a warning message is generated. Actions are the same as for `-ErrorAction`.
`-WarningVariable`	Provides the name of a variable used to hold warning information.
`-OutVariable`	Provides the name of a variable used to hold the cmdlet's output.
`-Debug`	Displays messages that are sometimes helpful when debugging a cmdlet.

Some cmdlet parameters are positional, which means that you can omit the parameter name and just list the parameter values. For the `Get_Random` cmdlet, `-Maximum` is the first positional parameter. So, if you simply specify a value without a parameter name, `Get-Random` uses the value as the `-Maximum` parameter:

```
Get-Random 10
```

One final bit about parameters: If you omit a required parameter, PowerShell will prompt you to enter its value. You'll see an example of this in the section, "Using Aliases" later in this chapter.

Getting Help

PowerShell includes an extensive collection of help information that you can access via the `Get-Help` cmdlet. Simply provide the name of the cmdlet you need help with as a positional parameter. For example, here is the `Get-Help` output for the `Get-Random` cmdlet:

```
PS C:\Users\Doug> Get-Help Get-Random

NAME
    Get-Random
```

```
SYNOPSIS
    Gets a random number, or selects objects randomly from a collection.


SYNTAX
    Get-Random [-InputObject] <Object[]> [-Count Int32
    [-SetSeed Int32] [<CommonParameters>]


    Get-Random [[-Maximum] <Object>] [-Minimum <Object>]
    [-SetSeed Int32] [<CommonParameters>]

DESCRIPTION
    The Get-Random cmdlet gets a randomly selected number. If you submit a
    collection of objects to Get-Random , it gets one or more randomly selected
    objects from the collection.

    Without parameters or input, a Get-Random command returns a randomly
    selected 32-bit unsigned integer between 0 (zero) and Int32.MaxValue
    (0x7FFFFFFF, 2,147,483,647).

    You can use the parameters of Get-Random to specify a seed number, minimum
    and maximum values, and the number of objects returned from a submitted
    collection.


RELATED LINKS
    Online Version:
    http://go.microsoft.com/fwlink/?LinkId=821799

REMARKS
    To see the examples, type: "get-help Get-Random -examples".
    For more information, type: "get-help Get-Random -detailed".
    For technical information, type: "get-help Get-Random -full".
    For online help, type: "get-help Get-Random -online"

PS C:\Users\Doug>
```

You can use several additional parameters to get even more help information:

- `-Examples`: Displays examples of the cmdlet, along with a detailed explanation of what each example does.
- `-Detailed`: Provides more detailed help.
- `-Full`: Displays all available help information.

» –Online: Opens a web browser homed on Microsoft's help page for the cmdlet. (This is actually the most useful form of help; see Figure 8-2 for an example of the help page for the Get_Random cmdlet.)

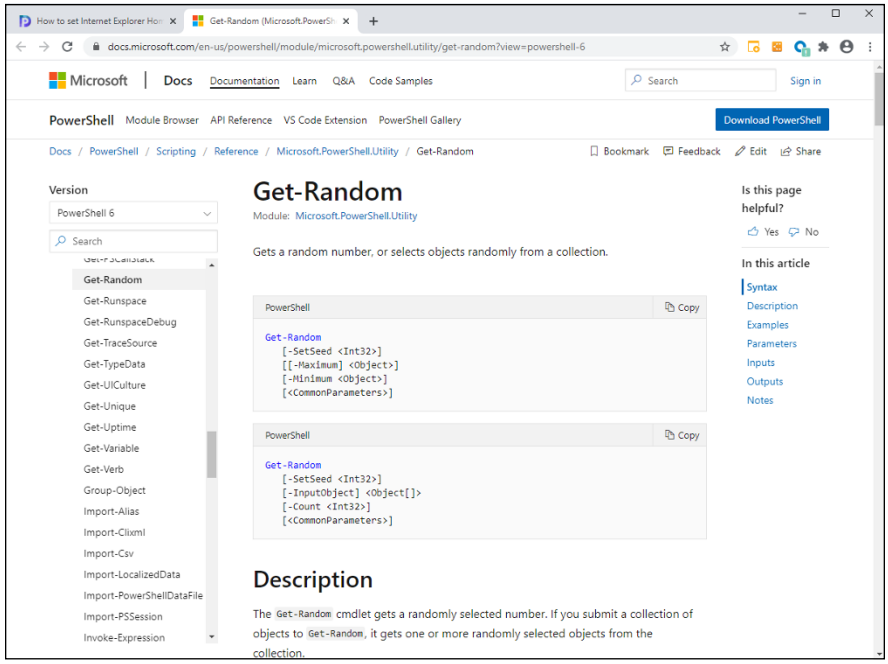


FIGURE 8-2: Online help for a PowerShell cmdlet.

Using Aliases

By now, you may be grumbling that PowerShell is a bit verbose. Indeed, most cmdlet names are pretty long. And although PowerShell will attempt to figure out what parameter names you mean if you don't spell them out completely, PowerShell doesn't give you the same grace with cmdlet names: If you don't spell out a cmdlet name in its entirety, PowerShell displays a rude error message.

Fortunately, PowerShell does provide relief in the form of aliases. An *alias* is an alternative name for a cmdlet. You can make up your own aliases, or you can use a somewhat large number of built-in aliases.

For example, earlier I show an example of the `Get-ChildItem` cmdlet used to list the contents of a folder. `Get-ChildItem` is the PowerShell equivalent of the `dir` command, and PowerShell provides `dir` as an alias for the `Get-ChildItem`

command. So although `Get-ChildItem` is the cmdlet to list the contents of a folder, you can call the `Get-ChildItem` cmdlet by entering `dir` at a PowerShell prompt.

Thus, you can display the contents of a folder like this:

```
PS C:\Users\Doug> dir

                    Directory: C:\Users\Doug

Mode            LastWriteTime                 Length Name
----            -------------                 ------ ----
d-r---      1/4/2018 10:18 PM                        Contacts
d-r---      1/4/2018 10:18 PM                        Desktop
d-r---      1/4/2018 10:18 PM                        Documents
d-r---      1/4/2018 10:18 PM                        Downloads
d-r---      1/4/2018 10:18 PM                        Favorites
d-r---      1/4/2018 10:18 PM                        Links
d-r---      1/4/2018 10:18 PM                        Music
d-r---      1/4/2018 10:19 PM                        OneDrive
d-r---      1/4/2018 10:18 PM                        Pictures
d-r---      1/4/2018 10:18 PM                        Saved Games
d-r---      1/4/2018 10:18 PM                        Searches
d-r---      1/4/2018 10:18 PM                        Videos

PS C:\Users\Doug>
```

If you want to see a list of all the aliases that are available, use the `Get-Alias` cmdlet. To narrow the list down to show just the aliases for a particular cmdlet, use the `-Definition` parameter, as in this example:

```
PS C:\Users\Doug> Get-Alias -Definition Get-ChildItem

CommandType Name                                Version Source

----------- ----                                ------- ------
Alias       dir -> Get-ChildItem
Alias       gci -> Get-ChildItem
Alias       ls -> Get-ChildItem

PS C:\Users\Doug>
```

Here, you can see that three aliases are defined for the `Get-ChildItem` cmdlet: `dir`, `gci`, and `ls`. `Dir` is the Windows equivalent to `Get-ChildItem`, `gci` is simply an abbreviation for `Get-ChildItem`, and `ls` is the Linux equivalent.

If you want to create your own aliases, you can use the `Set-Alias` command. This cmdlet requires two parameters: `-name`, which provides the name of the alias, and `-value`, which indicates the cmdlet that will be aliases. For example:

```
Set-Alias -Name ListFiles -Value Get-ChildItem
```

This creates a new alias for the `Get-ChildItem` cmdlet named `ListFiles`.

To remove an alias, you have to use the `Remove-Item` cmdlet, as in this example:

```
Remove-Item Alias:ListFiles
```

In this case, you indicate that you want to remove an Alias item, followed by a colon and the name of the alias you want to remove.

Using the Pipeline

The verb-noun naming convention isn't the most important difference between PowerShell and other command shells. The real difference is how PowerShell cmdlets deal with piped input and output. PowerShell takes the idea of piping to a new level.

In Book 6, Chapter 7, I explain how to use piping to chain two standard Windows commands together so that the output from the first command is piped into the second command. For example:

```
C:\>type users.txt | sort
```

Here, the `type` command displays the contents of a text file named `users.txt`. But instead of being displayed on the screen, the output from the `type` command is fed into the `sort` command, which sorts the text and then displays it on the screen. The result is that the contents of the `users.txt` file are displayed on the screen in sorted order.

With a standard Windows command, the input and output for commands that can use piping is always simple text. Thus, the `type` command creates text output, and the `sort` command reads text input and creates more text output. When the shell gets to the end of a sequence of piped commands, the output from the last command is displayed on the screen.

With cmdlets, the information that is piped is not simple text but complete objects. An *object* is an amalgamation of data, as well as executable code. Objects have

properties, which are named characteristics of the object, and *methods,* which are named functions that the object can perform. Methods are important in PowerShell, but using them is an advanced topic that's beyond the scope of this chapter. So I'm focusing here on properties.

Consider the `Get-ChildItem` cmdlet, which lists the contents of a folder:

```
PS C:\Users\Doug> Get-ChildItem

    Directory: C:\Users\Doug

Mode         LastWriteTime           Length Name
----          -------------          ------ ----
d-r---    1/4/2018 10:18 PM                 Contacts
d-r---    1/4/2018 10:18 PM                 Desktop
d-r---    1/4/2018 10:18 PM                 Documents
d-r---    1/4/2018 10:18 PM                 Downloads
d-r---    1/4/2018 10:18 PM                 Favorites
d-r---    1/4/2018 10:18 PM                 Links
d-r---    1/4/2018 10:18 PM                 Music
d-r---    1/4/2018 10:19 PM                 OneDrive
d-r---    1/4/2018 10:18 PM                 Pictures
d-r---    1/4/2018 10:18 PM                 Saved Games
d-r---    1/4/2018 10:18 PM                 Searches
d-r---    1/4/2018 10:18 PM                 Videos

PS C:\Users\Doug>
```

This cmdlet doesn't actually produce the text that is displayed in the PowerShell window. Instead, it returns a collection of file system objects. These file system objects have a number of important properties, among them `Name`, `Length`, `LastWriteTime`, and `Mode`.

The `Get-ChildItem` cmdlet puts this collection in the *pipeline,* which is a repository for objects that passed from one cmdlet to another. In most cases, you only invoke one cmdlet at a time in PowerShell. In that case, the output from the cmdlet you invoke is passed to the end of the pipeline, which automatically renders the contents of the pipeline as text. Hence, the list of file system objects is converted to text form and displayed in the PowerShell window.

You can easily manipulate the output displayed for a cmdlet by piping the output to one of several commonly used cmdlets that sort, filter, or otherwise format the objects in the pipeline. For example, if you want to show the contents of a folder in reverse alphabetical order, you can pipe the `Get-ChildItem` cmdlet's output into the `Sort-Object` cmdlet and use the `-Descending` parameter to reverse the order:

```
PS C:\Users\Doug> Get-ChildItem | Sort-Object -Descending

    Directory: C:\Users\Doug

Mode          LastWriteTime              Length Name
----          -------------              ------ ----
d-r---     1/4/2018 10:18 PM                    Videos
d-r---     1/4/2018 10:18 PM                    Searches
d-r---     1/4/2018 10:18 PM                    Saved Games
d-r---     1/4/2018 10:18 PM                    Pictures
d-r---     1/4/2018 10:19 PM                    OneDrive
d-r---     1/4/2018 10:18 PM                    Music
d-r---     1/4/2018 10:18 PM                    Links
d-r---     1/4/2018 10:18 PM                    Favorites
d-r---     1/4/2018 10:18 PM                    Downloads
d-r---     1/4/2018 10:18 PM                    Documents
d-r---     1/4/2018 10:18 PM                    Desktop
d-r---     1/4/2018 10:18 PM                    Contacts

PS C:\Users\Doug>
```

As you can see, PowerShell uses the vertical-bar character (also known as the pipe character) to indicate piping.

If you want to pick and choose which properties to display when you use `Get-ChildItem`, you can use the `Select-Object` cmdlet. For example:

```
PS C:\Users\Doug> Get-ChildItem | Select-Object -Property Name

Name
----
Contacts
Desktop
Documents
Downloads
Favorites
Links
Music
OneDrive
Pictures
Saved Games
Searches
Videos

PS C:\Users\Doug>
```

In this example, the `Select-Object` cmdlet's `Property` method indicates that you want to include only the `Name` property. The result is a list of filenames.

You can select more than one property by separating the property names with commas, as in this example:

```
PS C:\Users\Doug> Get-ChildItem | Select-Object -Property Name, LastWriteTime

Name           LastWriteTime
----           -------------
Contacts       1/4/2018 10:18:49 PM
Desktop        1/4/2018 10:18:49 PM
Documents      1/4/2018 10:18:49 PM
Downloads      1/4/2018 10:18:49 PM
Favorites      1/4/2018 10:18:49 PM
Links          1/4/2018 10:18:49 PM
Music          1/4/2018 10:18:49 PM
OneDrive       1/4/2018 10:19:05 PM
Pictures       1/4/2018 10:18:49 PM
Saved Games    1/4/2018 10:18:49 PM
Searches       1/4/2018 10:18:49 PM
Videos         1/4/2018 10:18:49 PM

PS C:\Users\Doug>
```

Here's an example that invokes three cmdlets: The first gets the contents of the current folder, the second selects just the `Name` property, and the third sorts the list in descending order:

```
PS C:\Users\Doug> Get-ChildItem | Select-Object -Property Name | Sort-Object
   -Descending
Name
----
Pictures
OneDrive
Music
Videos
Searches
Saved Games
Documents
Desktop
Contacts
Links
Favorites
Downloads
PS C:\Users\Doug>
```

The more you learn about PowerShell, the more you'll come to rely on the pipeline to tailor PowerShell to meet your precise needs.

Using Providers

One of the most interesting things about PowerShell is the concept of providers. A *provider* is a source of data that is consumed by many of PowerShell's commands. For example, the `Get-ChildItem` command consumes information from a provider called `FileSystem`, which represents the host computer's file system.

PowerShell provides several providers besides `FileSystem`. To see them all, you can use the `Get-PSProvider` command:

```
PS C:\Users\Doug> Get-PSProvider

Name          Capabilities                          Drives
----          ------------                          ------
Registry      ShouldProcess, Transactions           {HKLM, HKCU}
Alias         ShouldProcess                         {Alias}
Environment   ShouldProcess                         {Env}
FileSystem    Filter, ShouldProcess, Credentials    {C}
Function      ShouldProcess                         {Function}
Variable      ShouldProcess                         {Variable}

PS C:\Users\Doug>
```

Depending on the environment in which you run the `Get-PSProvider` cmdlet, you may see additional providers as well.

All providers are modeled on the concept of a file system, meaning that providers present their data to PowerShell cmdlets through one or more drives which contain items organized into folders. This might seem confusing at first, but you'll get used to it once you start to work with it.

Looking at the output from the `Get-PSProvider` cmdlet, you can see that the `FileSystem` provider lists just one drive, identified as `C`. If more disk drives were available on the computer, additional drive letters would appear.

Other providers list their drives using short words or abbreviations rather than single letters. For example, the `Alias` provider has a single drive, named `Alias`.

Similarly, the `Registry` provider has two drives, named `HKLM` and `HKCU`. (If you're familiar with the Windows Registry, you'll recognize these as the common abbreviations for `HKEY_Local_Machine` and `HKEY_Current_User`, respectively.)

By default, the `Get-ChildItem` cmdlet uses the `FileSystem` provider, starting at the current folder location. However, you can easily switch the provider for `Get-Children` by specifying an alternative path. For example, to see a list of all available aliases, you can use this command:

```
Get-Children Alias:\
```

Notice that the drive name `Alias` is followed by a colon and a single backslash in much the same way that the root folder of the `FileSystem C` drive would be written as `C:\`.

You can change the default location for cmdlets that work with providers by using the `Set-Location` cmdlet. For example:

```
Set-Location Alias:\
```

Having set the default location to the root of the `Alias` drive, subsequent cmdlets such as `Get-ChildItem` will automatically pull data from the `Alias` provider rather than from the `FileSystem` provider.

Using Scripts

A *script* is a collection of PowerShell commands saved to a text file with the extension `.ps1`. You can run a script at a PowerShell prompt simply by typing the name of the script, without the extension. Thus, scripts are PowerShell's equivalent for batch files.

Scripts are a great way to simplify routine Windows administration tasks. Any time you find yourself entering the same cmdlets over and over again, consider placing the cmdlets in a script. Then you can simply run the script, and let the script take care of the details of each command.

For example, suppose you routinely want to know what processes are consuming the most memory resources. You can do that using a combination of several cmdlets:

- Use `Get-Process` to get a list of active processes.
- Use `Select-Object` to select just the `ProcessName` and `WS` properties. (`WS` stands for *working set,* which is one of the key memory indicators in a Windows system.)
- Use `Sort-Object` to sort the result in descending order on the `WS` property.
- Use `Select-Object` again to select just the top ten results.

The resulting command would look like this:

```
PS C:\Users\Doug> Get-Process | Select-Object -Property ProcessName, WS | Sort-
   Object -descending WS |
     Select-Object -first 10

ProcessName            WS
-----------            --
MicrosoftEdgeCP 269713408
MicrosoftEdgeCP 243097600
WINWORD         169254912
SearchUI        158670848
MsMpEng         120512512
SelfService      98873344
EXCEL            89866240
explorer         87056384
MicrosoftEdge    78503936
powershell       70873088

PS C:\Users\Doug>
```

But that's a lot to type. To save all the wear and tear on your fingers and your keyboard, you can create a `.ps1` file with the command, as shown in Figure 8-3.

```
Get-Process | Select-Object -Property ProcessName, WS | Sort-Object
   -descending WS |
     Select-Object -first 10
```

Then you can invoke the whole thing just by running the `.ps1` file.

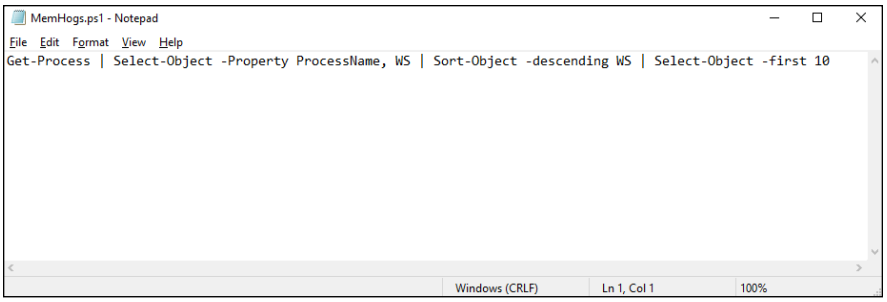


FIGURE 8-3: A PowerShell script.

However, before you can run scripts in PowerShell, you have to make a few preparations:

- **Run PowerShell with administrator permissions.** You can do that by right-clicking the PowerShell icon and choosing Run As Administrator.
- **Enable script execution by using the** `Set-ExecutionPolicy` **cmdlet.** For example:

  ```
  Set-ExecutionPolicy unrestricted
  ```

  When you run this command, the cmdlet will ask for your permission to enable unrestricted script execution.
- **Save your scripts to a location you can easily access.** For example, use `C:\Scripts`.

After you've enabled scripting, you can run a script by entering the script filename (including path) at the prompt:

```
PS C:\Windows\system32> c:\scripts\memhogs
ProcessName            WS
-----------            --
MicrosoftEdgeCP 269381632
MicrosoftEdgeCP 237064192
WINWORD         179699712
SearchUI        164970496
MsMpEng         117723136
SelfService      98873344
explorer         89452544
MicrosoftEdge    77590528
EXCEL            75345920
powershell       74674176

PS C:\Windows\system32>
```

There is much more to scripting than the limits of this short chapter allows me to go into. Here are some additional features you can explore on Microsoft's PowerShell website:

- Variables, which let you store and later retrieve values and objects.
- Functions, which let you create a set of PowerShell commands that you give a name to. Then you can run the named set of commands using the function's name as if it were a command.
- Advanced Functions, which let you create functions with parameters.
- Logic statements, including `While`, `Do...While`, `Do...Until`, `For`, `Foreach`, `If`, and `Switch`.

7 Administering Microsoft 365

Contents at a Glance

CHAPTER 1: **Getting Started with Microsoft 365 Administration** 605

CHAPTER 2: **Configuring Exchange Online** 625

CHAPTER 3: **Administering Teams** 641

IN THIS CHAPTER

- Finding out about Microsoft 365 and its various incarnations
- Understanding the options for Microsoft 365 subscriptions
- Creating an Microsoft 365 tenant
- Avoiding the mistake of hastily choosing a tenant name
- Adding users
- Resetting passwords
- Blocking users

Chapter 1

Getting Started with Microsoft 365 Administration

This chapter is a gentle introduction to administering Microsoft 365 for your organization. I start by presenting some basic Microsoft 365 concepts — like the difference between Microsoft 365 and Office 365, what a *tenant* is, how Microsoft 365 is licensed, and so on. Then I show you how to set up a new Microsoft 365 tenant. And finally, I show you how to do basic user management in Microsoft 365, such as creating new user accounts, assigning licenses, resetting passwords and deleting users.

Introducing Microsoft 365

Microsoft 365 — also known as M365 — is a suite of cloud applications designed to aid business productivity. M365 has its origins in the original Microsoft Office, which has been around for more than 30 years — the first version, Microsoft Office for Windows, was released in 1990. That version contained just three applications: Word, Excel, and PowerPoint. (Fun fact: Office was released for Mac computers a year before it was released for Windows.) Microsoft Access was added in 1995, and Outlook was added in 2000. Office continues to grow and flourish, with new versions coming every few years.

But in 2010, Office 365 arrived on the scene. Office 365 — also known as O365 — is an alternative to regular Office, which is distributed on CDs and intended for stand-alone use on individual computers. Office 365, on the other hand, is distributed via the internet and includes features integrated with Microsoft's cloud platform, Azure. (For more information about Azure, turn to Book 5, Chapter 3.)

In 2020, Microsoft rebranded Office 365 as Microsoft 365. Check out the sidebar, "Wait, is it Microsoft 365 or Office 365?" for a rundown on the confusion surrounding this name change.

As I write this in 2024, you can still purchase stand-alone versions of Office (the current version is 2021), but most businesses use Office 365 because of the additional features it provides.

Speaking of features, here are the most popular features available in Microsoft 365 (the specific features available depend on your subscription plan):

- Azure-based Active Directory synchronized with your on-premises Active Directory
- All the standard applications that traditionally came with Office — Access, Excel, PowerPoint, Publisher, OneNote, Outlook, and Word
- Exchange Online, a cloud-hosted version of Exchange Server that provides email services (see Book 7, Chapter 2)
- Microsoft Teams for online meetings, collaboration, chat, and calling (see Book 7, Chapter 3)
- SharePoint for building a company Intranet
- Yammer to provide a company social network

- OneDrive for cloud-based file storage (limited to 1TB with some plans, unlimited with others)
- Various work-management apps, including Forms, Planner, Power Apps, Power Automate, Power Virtual Agents, and To Do

WAIT, IS IT MICROSOFT 365 OR OFFICE 365?

In 2010, Microsoft introduced a subscription-based version of Office called Office 365, also known as O365. Office 365 allowed users to pay for Office on a month-to-month basis rather than all at once, and it entitled users to update to the latest version of Office as long as they maintained their subscription. The only real drawback was that if they let their subscription lapse, they could no long use Office applications.

The idea caught on, and in 2017 Microsoft introduced a new subscription product called Microsoft 365, also known as M365. M365 cost a bit more per month than O365 but included a subscription to Windows as well. So, when new versions of Windows were released, M365 subscribers did not have to pay separately for the Windows upgrade.

Microsoft liked M365 so much that, in an effort to get everyone to everyone to upgrade, they pulled a switcheroo in the spring of 2020 and rebranded some of its Office 365 products as Microsoft 365, raising the price in the process. No one really noticed, probably because it happened during the COVID-19 pandemic, and what Microsoft did with its product branding didn't seem to matter much in the grand scheme of things.

Nevertheless, the name change has proven somewhat confusing, because Microsoft is not entirely consistent about how it uses the terms *Microsoft 365* and *Office 365* when it comes to home usage versus business usage.

For home usage, Microsoft offers two versions of M365: Personal and Family. Neither version includes upgrades to Windows — they're exactly the same as the older O365 home version, just with the new name.

But at the Enterprise subscription level, Microsoft kept the name Office 365 for all the traditional Office 365 Enterprise subscriptions and added a slate of new subscription levels called Microsoft 365, which includes Office applications plus Windows upgrades, as well as a host of other features depending on your subscription level.

This is all very confusing, and I'm not sure why they didn't stick with the name Office 365 for the home and small business versions. But I'm not a marketing genius, just a lowly book author trying to figure out how to explain this mess of a naming fiasco.

BUT WAIT AGAIN . . . WHAT ABOUT TEAMS?

In the spring of 2024, Microsoft decided to confuse everyone even more by removing Microsoft Teams from all versions of Microsoft 365 and Office 365. Teams is now a separate subscription that must be added on top of your basic O365 or M365 subscription.

That is, except for the Business versions of M365, which can be purchased with or without Teams bundled in. If you don't need Teams in your small business, you can save $2.25 per month by subscribing to the no-Teams version.

For Enterprise versions of M365 and O365, a Teams for Enterprise subscription costs $5.25 pre month on top of the M365 or O365 subscription. There is also a premium level for an additional $7 per month that provides advanced features.

Microsoft made this change due to its fears about antitrust concerns. When Teams was included with Office 365 and Microsoft 365 subscriptions, competitors such as Slack, Google, and others had an unfair disadvantage. Microsoft had already removed Teams from subscriptions in Europe due to similar concerns.

Considering Microsoft 365 Plans

Before you set up Microsoft 365, you need to consider which of several subscription plans you want to use. The first decision is whether you want a consumer plan or an Enterprise plan. Consumer plans are designed for home use or very small businesses (those with, say, no more than ten employees). Enterprise plans provide more-advanced features and are designed for larger businesses.

Table 1-1 lists the features and pricing for the various small business plans as of January 2021, and Table 1-2 lists the Enterprise plans. Note that, per Microsoft's new branding, the small business plans are all called Microsoft 365 and the Enterprise plans are called Microsoft 365. For the Enterprise plans, be aware that each plan also has an equivalent Microsoft 365 plan that includes Windows 10 at a cost of an additional $12 per user per month.

For the purposes of this chapter, I focus on the four Microsoft plans that include both desktop versions of the applications as well as online email:

- Microsoft 365 Business Standard
- Microsoft 365 Business Premium
- Microsoft 365 E3
- Microsoft 365 E5

TABLE 1-1 Small Business Plans for Microsoft 365

	Microsoft 365 Business Basic	Microsoft 365 Business Standard	Microsoft 365 Business Premium	Microsoft 365 Apps
Price per user per month with Teams (as of August 2024)	$6.00	$12.50	$22.00	$8.25
Price per user per month without Teams (as of August 2024)	$4.75	$10.25	$19.75	n/a
Applications	Web and mobile versions only (no desktop versions): Excel, OneNote, Outlook, PowerPoint, Word	Access, Excel, OneNote, Outlook, PowerPoint, Publisher, Word	Access, Excel, OneNote, Outlook, PowerPoint, Publisher, Word	Access, Excel, OneNote, Outlook, PowerPoint, Publisher, Word
Cloud Services	Exchange, OneDrive, SharePoint, Teams	Exchange, OneDrive, SharePoint, Teams	Azure Information Protection, Device Management, Exchange, Intune, OneDrive, SharePoint, Teams	OneDrive
Maximum number of users	300	300	300	300

TABLE 1-2 Enterprise Plans for Microsoft 365 (No Teams)

	Microsoft 365 E3	Microsoft 365 E5
Price per user per month (as of August 2024)	$33.75	$54.75
Applications	Access, Excel, Outlook, PowerPoint, Publisher, Word, Visio	Access, Excel, Outlook, PowerPoint. Publisher, Word
Cloud Services	Azure Information Protection, Device Management, Exchange, Online, OneDrive, SharePoint, Teams, Stream, Viva Engage, Viva Insights	Azure Information Protection, Device Management, Exchange Online, Intune, Online, OneDrive, SharePoint, Teams, Stream, Viva Engage, Viva Insights
Work Management	Forms, Planner, Power Apps, Power Automate, To Do, Loop, Clipchamp	Forms, Planner, Power Apps, Power Automate, To Do, Loop, Clipchamp, Power BI Pro
Other Features	Basic Security and Identity	Advanced Security and Identity

The other plans don't include the desktop applications (which are a must for most organizations), Exchange Online, and other cloud features.


REMEMBER

The good news is that you don't have to fret much over the decision of which plan to select when you're just getting started. It's easy to upgrade the plan later. Microsoft is always willing to take more of your money! You can easily start with Microsoft 365 Business Basic, and then upgrade later to Business Premium or Microsoft 365 E3 or E5.

Understanding Tenants

In the world of Microsoft 365, the environment that represents a single organization is called a *tenant.* Technically, a *tenant* is a dedicated instance of Active Directory running in Azure. The tenant manages all the users, apps, and data associated with the organization.

When you first set up a tenant, you choose a *tenant name* that is a subdomain under `onmicrosoft.com`. For example, you may choose `whatchamacallit.onmicrosoft.com` for your tenant name. The tenant name must be unique, so if someone has already used the name you choose, you'll have to pick another.

After you've created a tenant, you can add a custom domain to the tenant. For example, if you already own `whatchamacallit.com`, you could add that domain to the tenant and set it as the default domain for the tenant.

You can add additional domains to a tenant, and you can easily change the default domain. However, you *cannot* change the tenant name after you've selected it. When a tenant is created, the name is permanent.


WARNING

Be very deliberate about the tenant name you choose when you set up a tenant. One of the absolute worst things you can do is set up an experimental or test tenant using your company's primary domain name as the `onmicrosoft.com` name. If you do, you'll never be able to use your company's domain name for the tenant later on.

So, if you're just fooling around, add a number or some other variation to your domain name when you set up a test account. For example, use `whatchamacallit1.onmicrosoft.com` or `testwhatchamacallit.onmicrosoft.com`. That way, when it comes time to set up your actual Microsoft 365 tenant, your real name will still be available (unless someone grabs it before you can get around to it).



REMEMBER

After you set up an Microsoft 365 tenant, you cannot change the tenant name, so choose wisely!

Creating an Microsoft 365 Tenant

Now that you've had a whirlwind tour of Microsoft 365 plans and a stern warning about choosing your tenant name wisely, let's get on with the procedure for creating a new tenant for Microsoft 365. In this example, I'll select the Enterprise-level Microsoft 365 E3 plan and use `lowewriter.onmicrosoft.com` as the tenant name.



TIP

Before you begin, have a credit card handy. Microsoft 365 E3 costs $33.75 per month with a 12-month commitment, so to follow this procedure you need to be prepared to spend $240 for the experience. If you want to see the full range of features that are available, E3 is the plan to get. But if you want to save some money, you can go with the Microsoft 365 Business Standard plan instead. That will cost $12.50 per month for a one-year commitment of $150, and if you want, you can try it free for one month.

Follow these steps:

1. **Go to `www.microsoft.com/en-us/microsoft-365/enterprise` and click See Microsoft 365 Plans and Pricing.**

   The Microsoft 365 Enterprise Plans page is displayed, as shown in Figure 1-1.

2. **Under Microsoft 365 E3, click Try Free for One Month.**

   The page shown in Figure 1-2 appears. This page shows the first of five steps you'll need to complete to create a new tenant.

3. **Enter an email address and click Next.**

   The email address can be an address that already exists, such as an existing Gmail account. Or, if can be an email address that doesn't yet exist, such as an `onmicrosoft.com` address for the tenant you want to create.



TIP

Whichever email you use, it shouldn't be an email that's already associated with Microsoft 365.

For this example, I will use an email address that doesn't exist yet for the tenant I want to create: `doug@lowewriter.onmicrosoft.com`.

When you enter an address that doesn't exist, you'll be informed that you need to set up a new account, as shown in Figure 1-3.

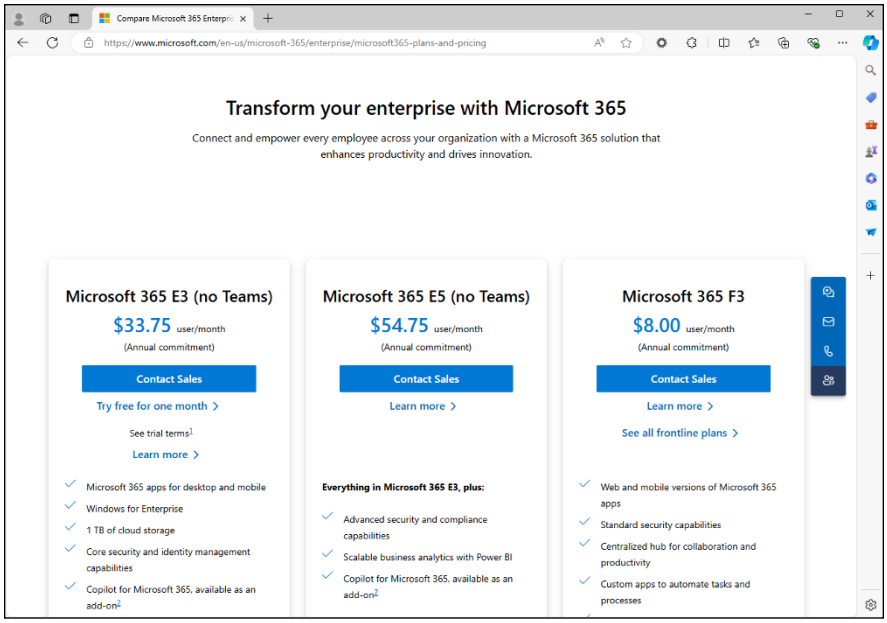


FIGURE 1-1: The available Microsoft 365 Enterprise plans.

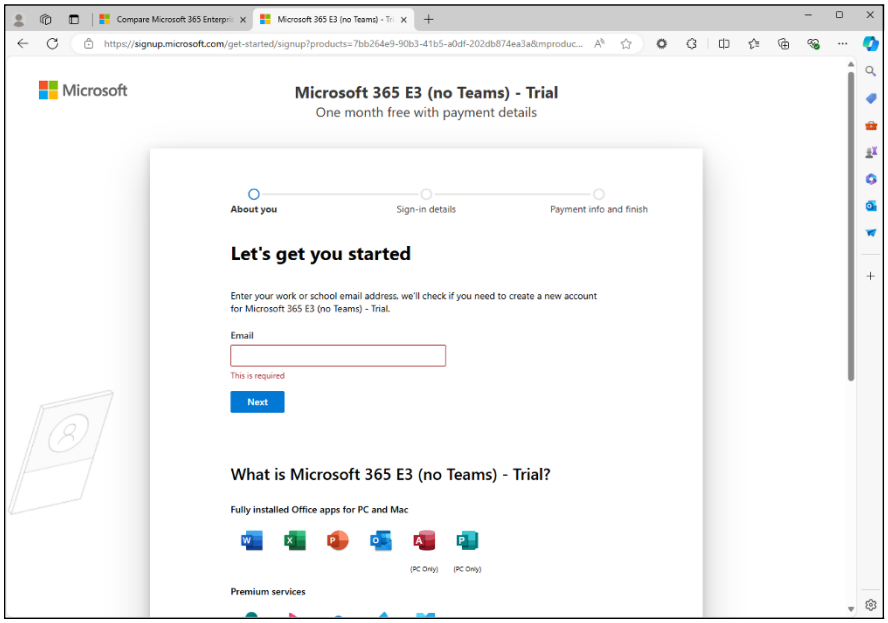


FIGURE 1-2: Entering an email address to create a new tenant.

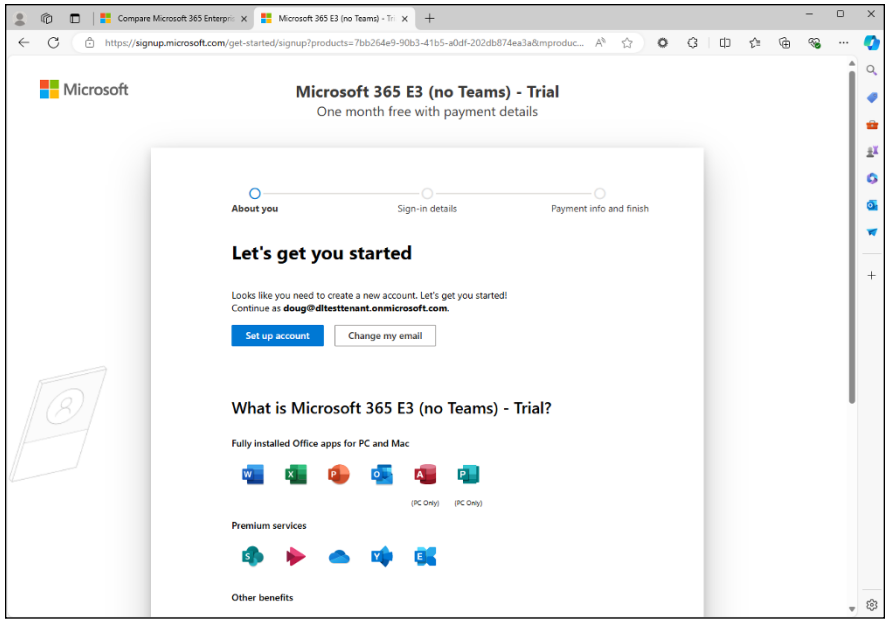


FIGURE 1-3: Getting started with a new account.

4. Click Set Up Account.

This takes you to the page where you enter your personal information (see Figure 1-4).

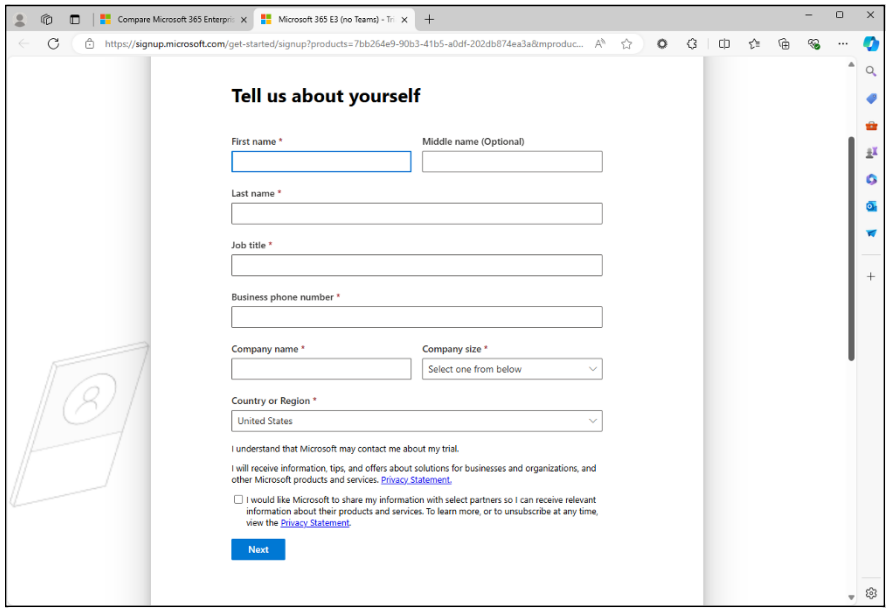


FIGURE 1-4: Entering your personal information.

5. **Enter your personal information.**

   You must enter a first and last name, your phone number, and company name. You must also select the correct country from the drop-down list.


WARNING

   The country is another one of those settings that you can't change later. Choose correctly now or forever hold your peace.

6. **Click Next.**

   A Security check screen prompts you to send a test code or to call you. Text is usually the easiest option.

7. **Click Send Verification Code.**

   Microsoft sends you a text message that includes a verification code, and a text box appears on the page.

8. **When you receive the verification code, enter it and click Verify.**

   You're taken to the next step, shown in Figure 1-5.

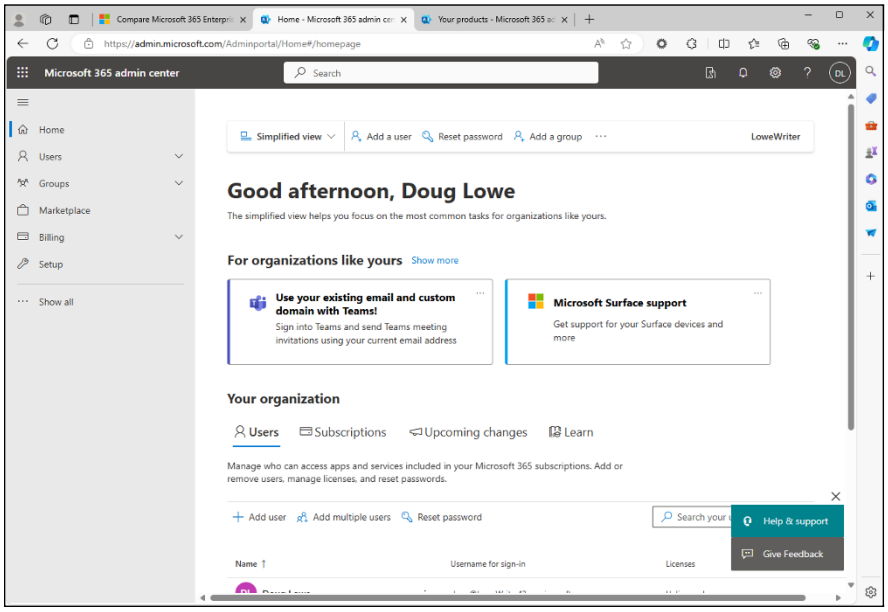


FIGURE 1-5: The Microsoft 365 Admin Center.

9. **Follow the steps to complete the remainder of the signup process.**

   You'll be asked to confirm your tenant name, provide a username and password for the tenant owner account, and complete a few other steps.



WARNING

After you create the tenant name, you can't change it, so choose wisely! (For more information, refer to the section "Understanding Tenants" earlier in this chapter.)



REMEMBER

You can always add a custom domain name later and choose it to be your primary domain.

After you've completed all the steps and placed your order, you'll see a final confirmation page congratulating you for successfully creating a new Microsoft 365 tenant.

10. **Click Go to Setup to get started.**

    You're taken to the Microsoft 365 Admin Center (refer to Figure 1-5). If you aren't taken to this page, close your browser, reopen it, and go to `https://admin.microsoft.com`.

Congratulations! You are now the proud owner of an Microsoft 365 tenant with one E3 license.

Creating a New User

To create a new user account in Microsoft 365, follow these steps:



1. **Go to the Admin Center by browsing to `https://admin.microsoft.com` and logging in.**

   The Admin Center appears (refer to Figure 1-5).

2. **In the menu pane at the left, select Users and then select Active Users.**

   The Active Users page appears, as shown in Figure 1-6.



3. **Click Add User (shown in the margin), and then choose Single User.**

   Other choices here are Guest User or Multiple Users. When you choose Single User, the page shown in Figure 1-7 appears.

4. **Enter the user's first and last name, display name, and username.**

5. **Either let Microsoft generate a random password or click Let Me Create the Password and enter a password yourself.**

   Additional options are Require This User to Change Their Password When They First Sign In and Send Password in Email Upon Completion.

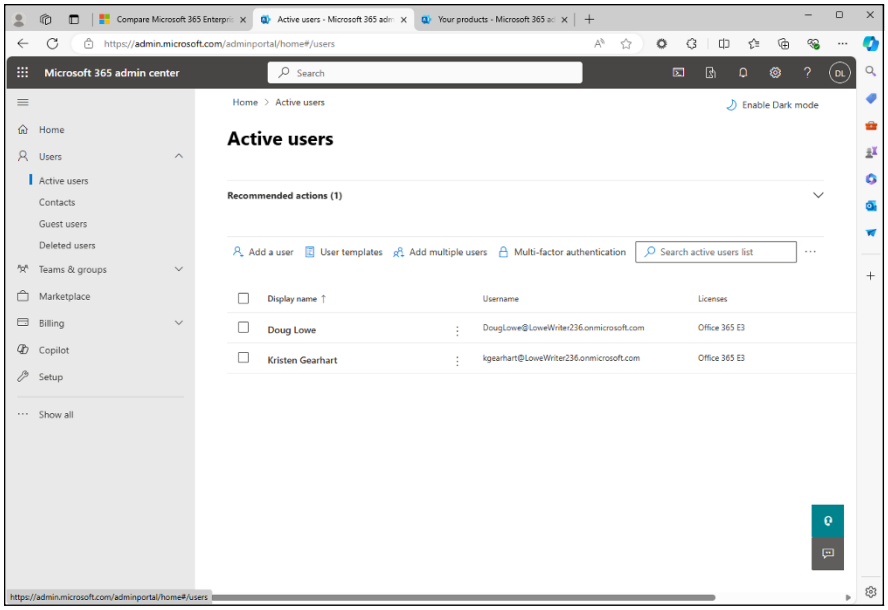


FIGURE 1-6: The Active Users page.

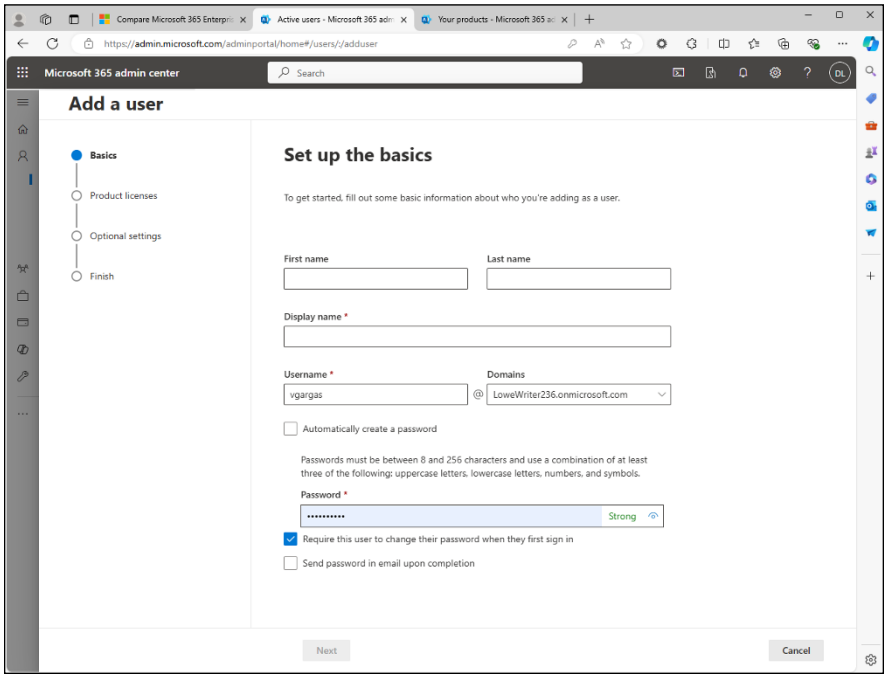


FIGURE 1-7: Adding a user.

6. **Click Next.**

   The next page, shown in Figure 1-8, allows you to select a license for the new user.

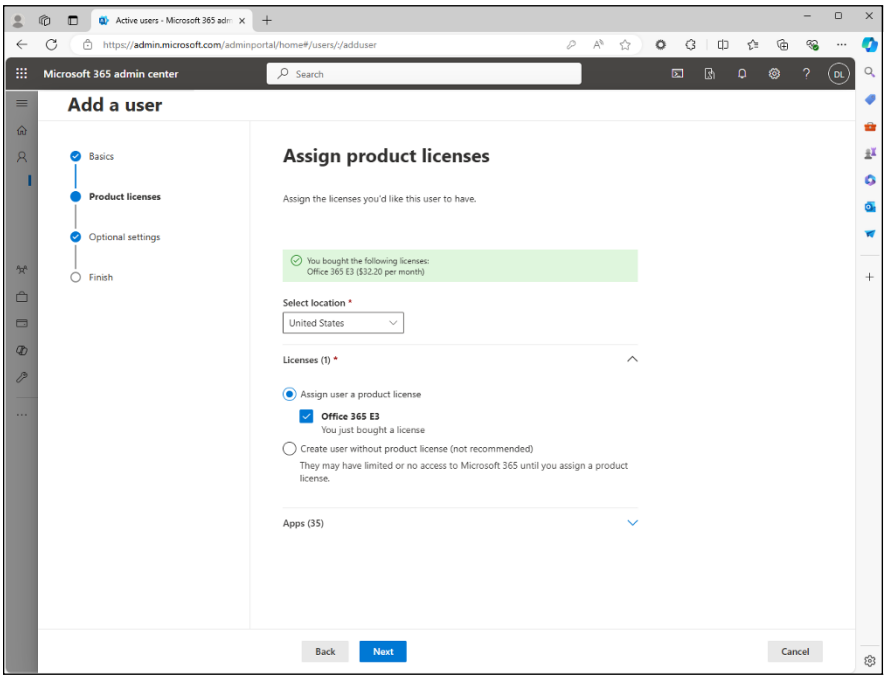


FIGURE 1-8: Assigning a new license.

7. **Select the product license to assign to the new user.**

   In this case, I don't have any E3 licenses available, but I'll let the wizard buy one for me.

8. **Click Next.**

9. **If you're prompted to buy another license, click Yes.**

10. **Click Apps near the bottom of the wizard page to reveal all the applications that are available in the E3 license, as shown in Figure 1-9.**

    If you want, uncheck applications you don't want the user to have access to.

11. **Click Next.**

    The Optional Settings page, as shown in Figure 1-10, appears.

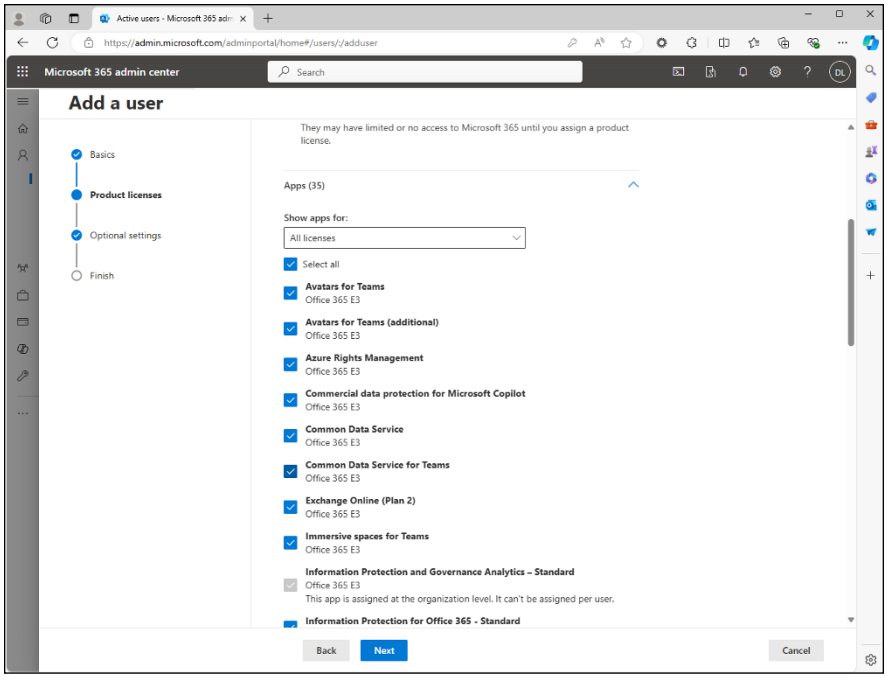


FIGURE 1-9: Assigning apps.

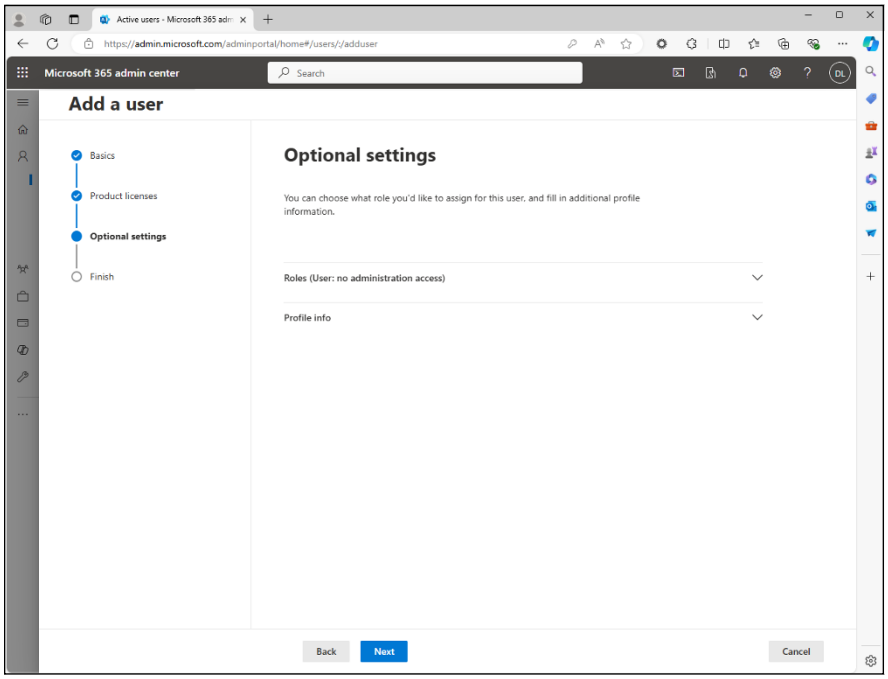


FIGURE 1-10: Configuring optional settings.

12. Click Roles and then click Admin Center Access to assign admin roles to the user, if desired.

The Roles page, along with the various administrator roles, is shown in Figure 1-11. If the user is simply a normal user, leave User (No Admin Center Access) checked. Otherwise, check the specific administrative features the user should have access to.

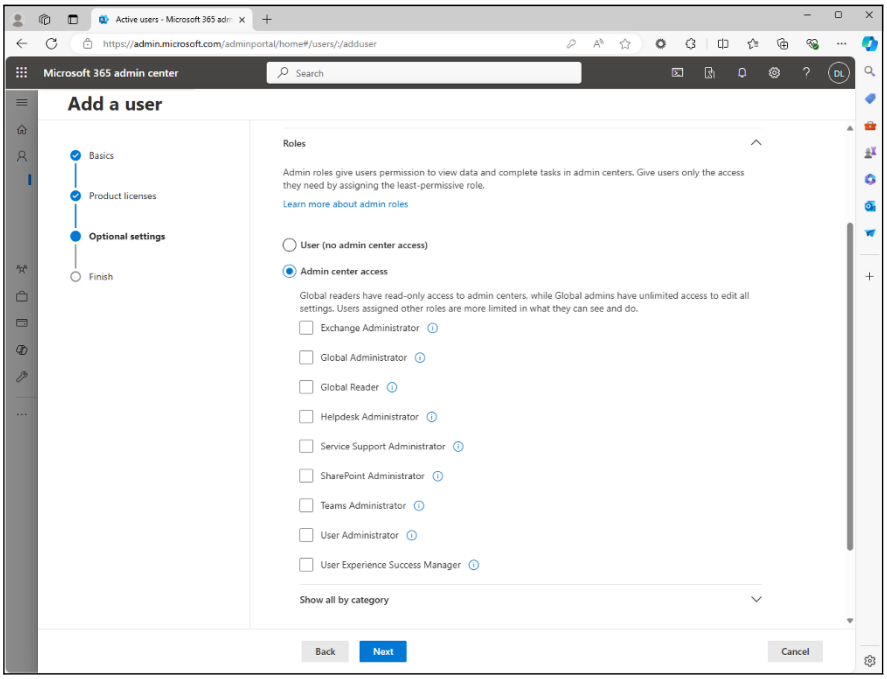


FIGURE 1-11: Assigning administrative roles.

13. Click Profile Info to fill out the user's profile information.

The Profile Info page is shown in Figure 1-12. Here, you can set the following values:

- Job title
- Department
- Office
- Office phone and fax number
- Mobile phone
- Street address, city, and state

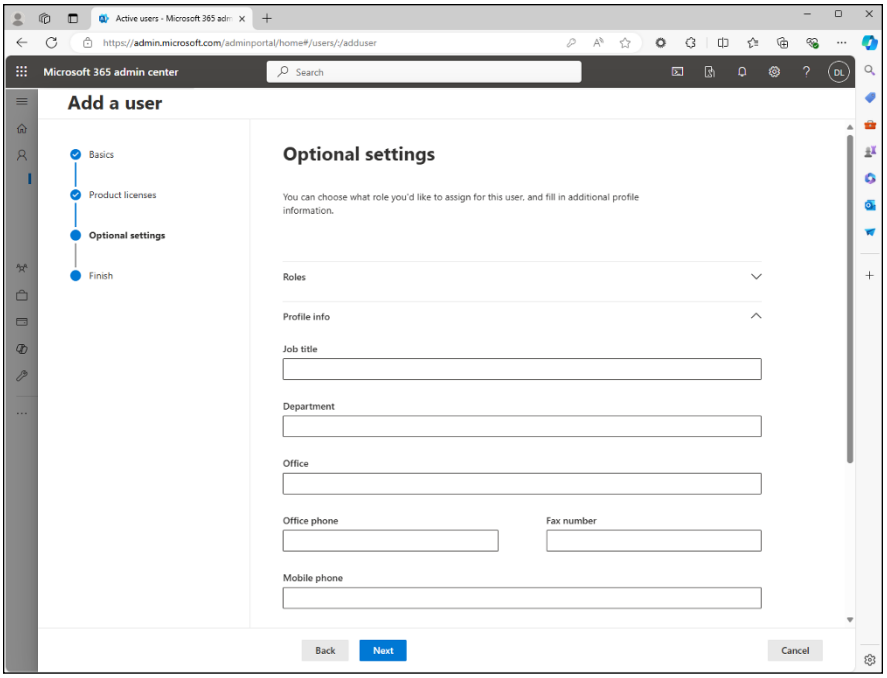


FIGURE 1-12: Entering profile details.

14. Click Next

The Review and Finish page appears. Review this page to make sure everything looks right.

15. Click Finish Adding.

A confirmation page appears to show that the user has been created.

16. Click Close.

A congratulatory page appears, as shown in Figure 1-13, confirming that the new user has been added. From this page, you can also send an email to the new user to let them know they've been added.

That's all there is to it! The new user account is now ready for use.

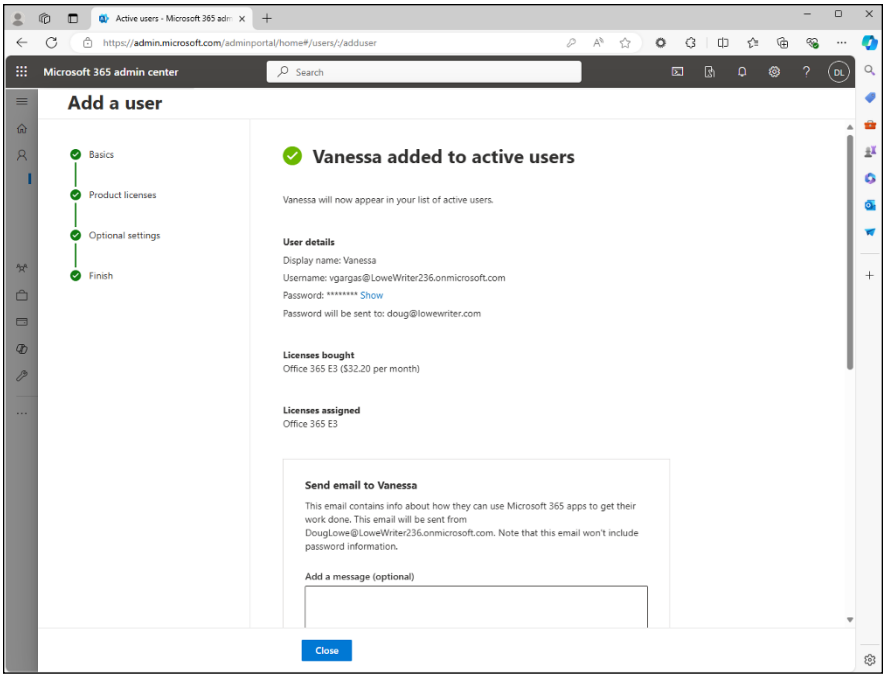


FIGURE 1-13: You've created a new user!

Resetting a User's Password

We all know that the most common support request is, "I forgot my password!" You can easily remedy that situation in Microsoft 365 by following these simple steps:

1. **At the Active Users page, select the user account.**
2. **Click the Reset Password button, shown in the margin.**

   You're prompted to change the password. You can either allow Microsoft to set the password or enter your own password. You can also require that the user change the password.
3. **Click Reset Password.**

   The password is reset and a confirmation message is shown.

TIP

You'll also find a check box that allows you to send the new password via email. You can send the email to yourself or to the user, but sending it to the user won't be much help because the user presumably can't sign in, having forgotten their password. So, send it to yourself or the user's manager, and then let the user know what the new password is.

4. **You're done!**

   You've successfully set the user's password, and you can now consider yourself an accomplished Microsoft 365 administrator.

Disabling a User



Microsoft 365 provides two ways to disable a user. The most drastic is to simply delete the user account, which you can do at the Active Users page by clicking the Delete a User button, shown in the margin.

The more temperate option is to block the user from accessing Microsoft 365. This allows you to restore access later on. When you block access, the user will be barred from all Microsoft 365 applications and be kicked out of any active Microsoft 365 applications within an hour.

Follow these steps to disable a user:

1. **In Active Users, click the user name.**

   This brings up the user properties, shown in Figure 1-14.



2. **Click the Block Sign-in button (shown in the margin).**

   A confirmation appears, as shown in Figure 1-15.

3. **Click Block This User from Signing In, and then click Save Changes.**

   The user is locked out.

To restore access for the user, follow the procedure again but click Unblock Sign-in instead of Block Sign-in in Step 2.

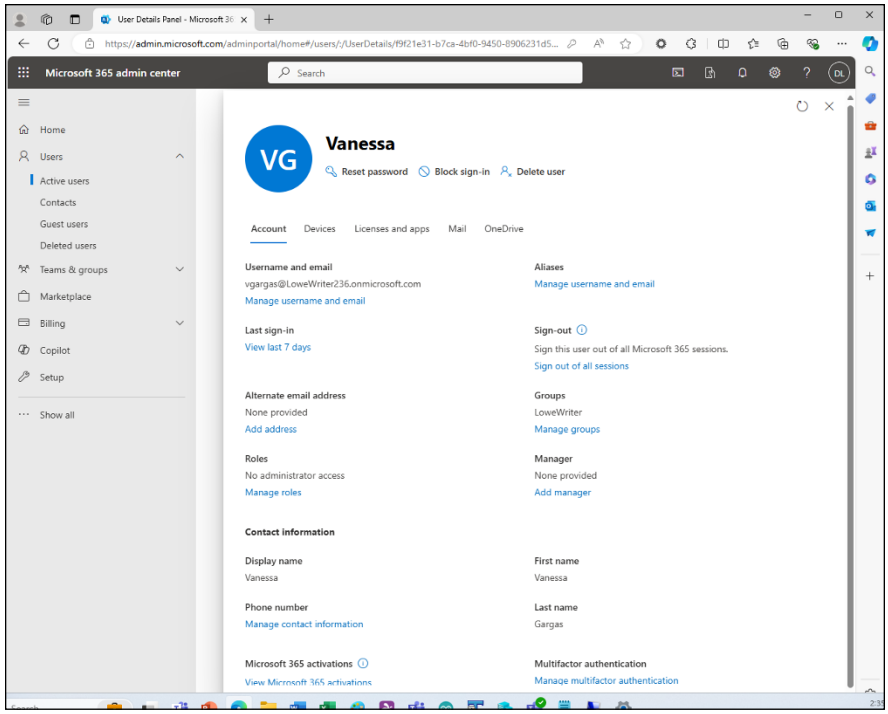


FIGURE 1-14: Displaying user information.

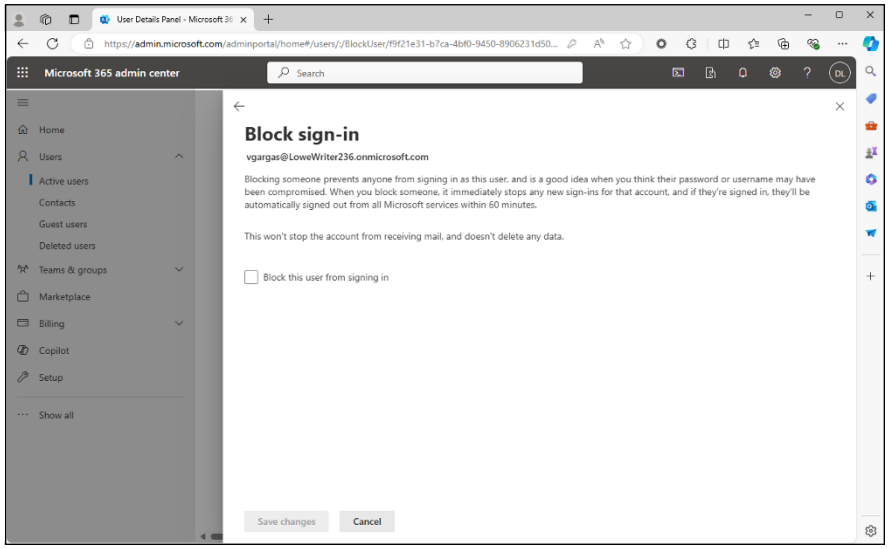


FIGURE 1-15: Blocking a user's access.

IN THIS CHAPTER

- Exploring the various types of mailboxes you can create in Exchange
- Working with the Exchange Server admin consoles
- Managing mailboxes
- Granting mailbox permissions
- Creating a shared mailbox

Chapter 2
Configuring Exchange Online

Exchange Online is the email server for Microsoft 365. It's included in all Microsoft 365 and Office 365 Enterprise plans, as well as in Microsoft 365 Business Basic, Business Standard, and Business Premium.

In this chapter, I show you how to perform the most commonly requested maintenance chores for Exchange Online, such as how to create a new mailbox, grant a user access to an additional mailbox, and deal with mailbox size limits. But first, I explain a few essential basics about how Exchange Online works.

Looking at Exchange Online Recipient Types

Exchange Online supports a variety of different types of recipients for different purposes:

- **Mailbox:** A standard *mailbox* provides email and other features for an Office 365 user. This mailbox contains all the features you normally associate

with an Outlook account: email, calendar, contacts, notes, and to-do list. Unlike on-premises Exchange, in Office 365 a standard mailbox is automatically created for every new user. So, you don't ever have to manually create standard mailboxes in Exchange Online.

- **Shared mailbox:** A *shared mailbox* is a mailbox that doesn't have a specific user associated with it. Instead, existing users can be granted access to the shared mailbox. A shared mailbox is a great way to provide a common email address that's monitored and used by several people. And it's also a great way to preserve the email for someone who no longer works for your company: Exchange Online makes it easy to convert a standard mailbox to a shared mailbox.

  The best thing about a shared mailbox is that it's free, provided it contains less than 50GB of data. If it's larger than 50GB, you'll need to purchase a license.

- **Microsoft 365 Group:** A *Microsoft 365 Group* represents a team of users who can collaborate using Microsoft Teams, which provides an online space for meetings, chat, file storage, and calendars. This type of group also includes a group email address so anyone can send an email to all members of the group.

- **Distribution group:** A *distribution group* is a simpler type of group that provides a group email address but no other features.

- **Mail-enabled security group:** A *mail-enabled security group* (often just called a *security group*) is like a distribution group that lets you assign access rights to members of the group.

- **Dynamic group:** A *dynamic group* is like a distribution group, but the membership isn't defined by a static list of users but by a set of rules. The rules are evaluated each time email is sent to the group.

- **Resource:** A *resource* is a mailbox that isn't associated with a user but instead is associated with a resource within your organization, such as a conference room or a vehicle. The most important feature of a resource is its calendar; users can reserve the resource by including it in a meeting request or other calendar item. The resource can be configured to automatically accept all invitations (provided the resource is available) to accept reservations or you can designate one or more people who must approve booking requests.

Examining the Exchange Admin Center

To get to the Exchange Admin Center, follow these steps:

1. **Log in to the Microsoft Admin page at `https://admin.microsoft.com`.**

   This takes you to the Microsoft Admin home page, shown in Figure 2-1.

2. **Click the Show All icon at the bottom of the list of icons on the left.**

   All the available Office administration centers are shown.



3. **Click the Exchange icon (shown in the margin).**

   The Exchange Admin Center, shown in Figure 2-2, appears.

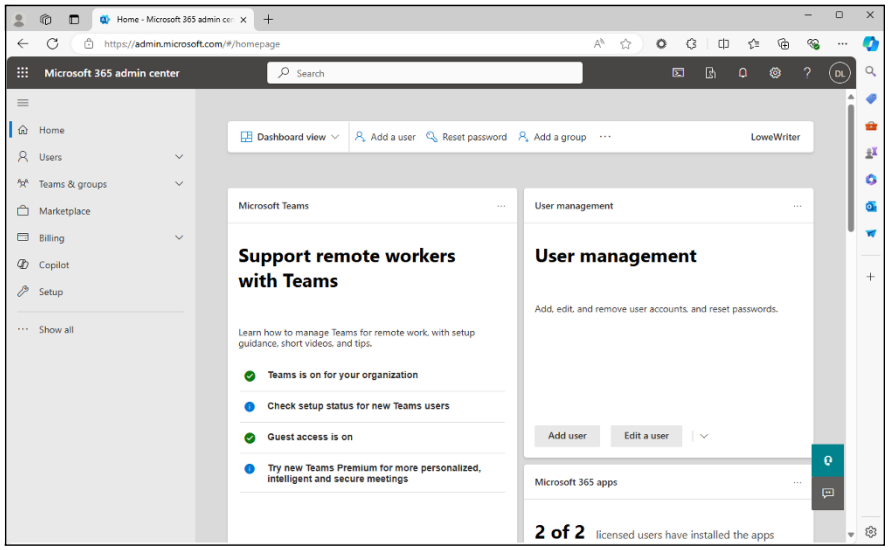


FIGURE 2-1: The Microsoft 365 Admin home page.

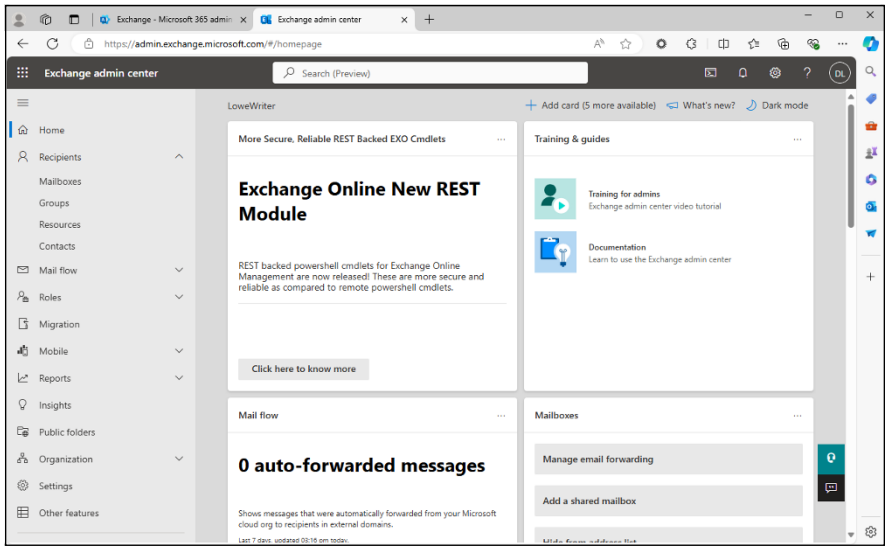
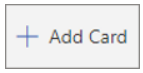


FIGURE 2-2: The Exchange Admin Center.



The information displayed in the home page dashboard is organized into groups called *cards*. You can customize your dashboard with a range of available cards by clicking the Add Card button (shown in the margin). This reveals a list of cards you can add, as shown in Figure 2-3.

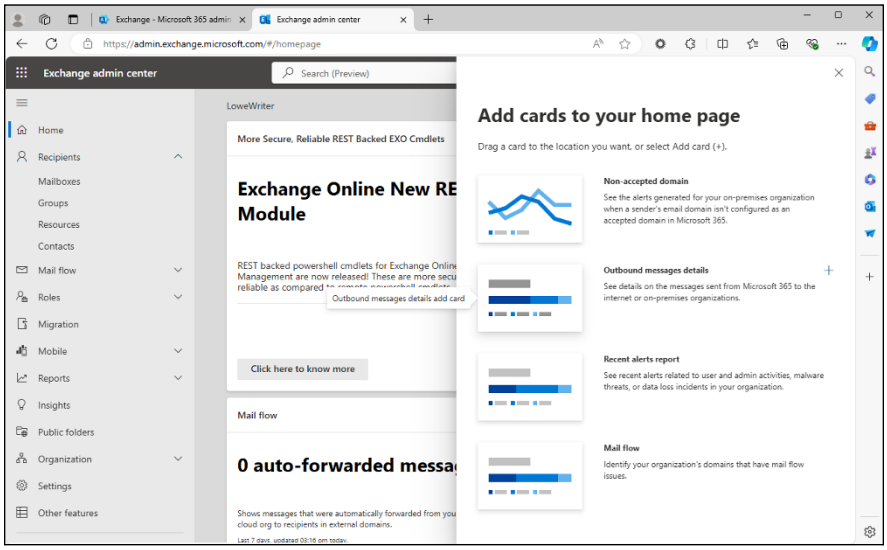


FIGURE 2-3: Adding a card to the new Exchange Admin Center.

The following cards are available:

- **Training & Guide:** Displays links to Microsoft documentation and tutorial information to help you learn more about administering Exchange. I suggest you add this card to your dashboard when you're first starting with Exchange. You can remove it later when you no longer need it.
- **Auto-Forwarded Messages:** Displays information about messages that have been forwarded outside of your organization.
- **Mailboxes:** Provides quick access to common tasks for user mailboxes.
- **Inbound Message Details:** Tracks incoming email.
- **Migration Batch Report:** If you're undergoing a migration from on-premises Exchange to Exchange Online, this card can help you track your progress.
- **Non-Accepted Domain:** This card can help you track email from your own domain where the sender's domain isn't listed as an acceptable domain for your server.
- **Non-Delivery Report:** Tracks error codes for the latest five days of undeliverable email.

- **Outbound Message Details:** Tracks outgoing email.
- **Recent Alerts Report:** Tracks recent Exchange alerts.
- **Mail Flow:** Alerts you to mail flow issues.

To add one of these cards to your dashboard, hover the mouse over the card until a plus sign is displayed (shown in the margin), then click the plus sign. When the card is added to the dashboard, you can drag it to change its location. To remove a card, click the ellipses (. . .) at the upper right of the card and choose Remove.

Managing Mailboxes

Your main task in the Exchange Admin Center will be managing various types of recipients. You do that by using one of the options under Mailboxes in the Exchange Admin Center menu. Figure 2-4 shows the Mailboxes page, used to manage user mailboxes. As you can see, this page lists all the mailboxes for the organization. Both User and Shared mailboxes are listed here.

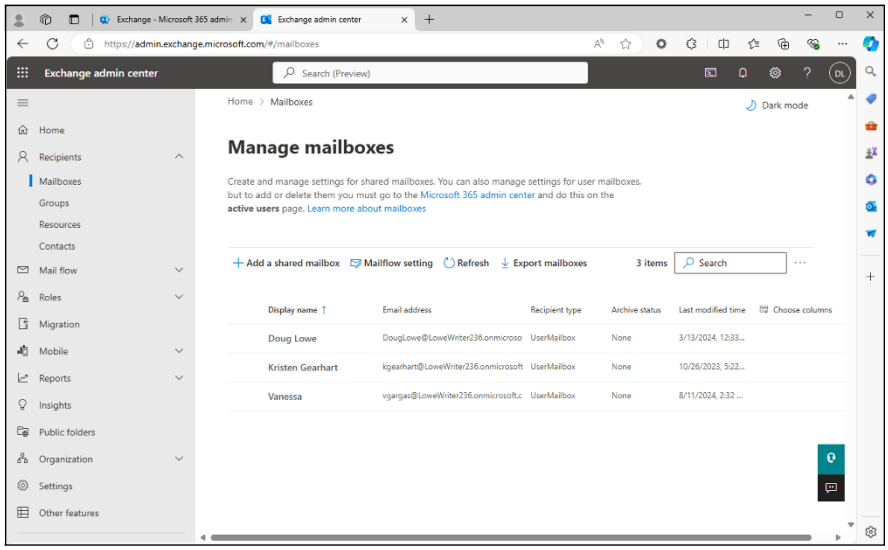


FIGURE 2-4: Managing mailboxes.

If you click any of the mailboxes in the list, you see the details pane for that mailbox, as shown in Figure 2-5.

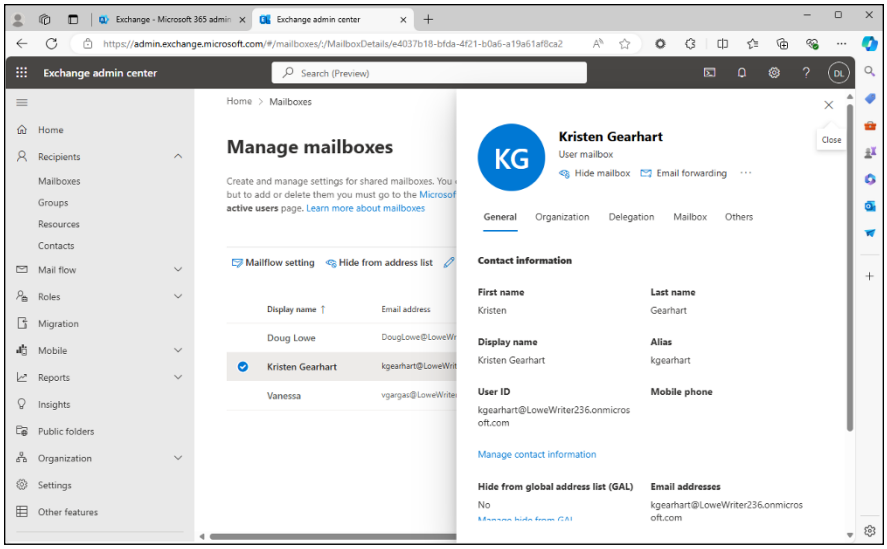


FIGURE 2-5: The details pane for a user mailbox.

The following sections describe several common features you can set from the mailbox details pane.

Creating an email alias

An *email alias* is an alternative email address for a mailbox. You can set up an alias by follow these steps:

1. **Click Manage Email Address Types under Email Addresses.**

   The Manage Email Address Types page, shown in Figure 2-6, appears. (You may need to scroll down a bit to find the Manage Email Address Types link.)

2. **Click + Add Email Address Type.**

   The New Email Address page, shown in Figure 2-7, appears.

3. **Enter the new email address and click OK.**

   The new address is added to the list of valid email addresses for the user.

4. **Click Save to save your changes.**

   The user's mailbox is updated with the new address.

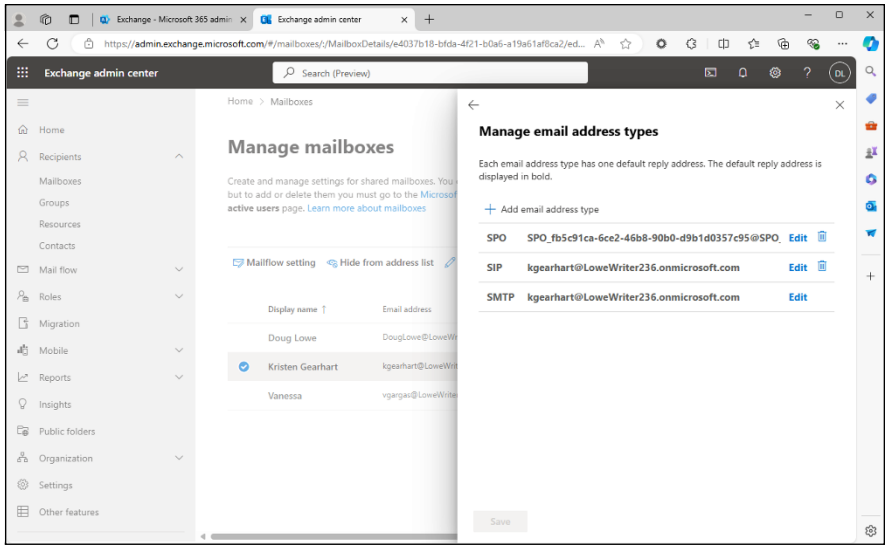


FIGURE 2-6: Managing email addresses.

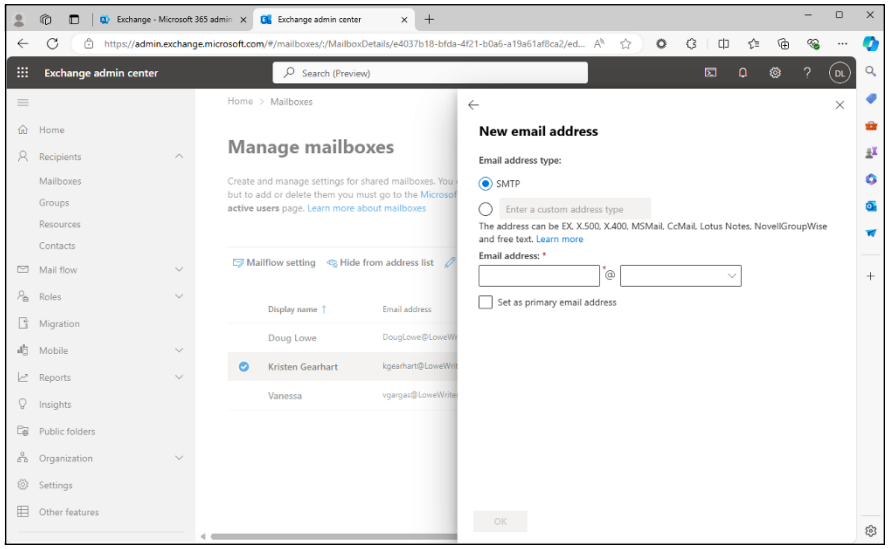


FIGURE 2-7: Creating a new email address.



TIP

Here are a few additional tidbits about the Manage Email Address Types page:

- When you create a new email alias, you can make the alias act as the default reply address for outgoing email sent from the mailbox.
- The email address must use a domain that is authorized for the Exchange server.

- Although SMTP is the most common type of email address, you can also add other address types such as X.500, X.400, Lotus Notes, and Novel GroupWise.
- You can also use the Manage Email Address Types page to remove an email alias or to change the primary email address used for the mailbox.

Delegating a mailbox

Sometimes you want to grant other users the right to access a mailbox. That's commonly done when an administrative assistant requires access to an executive's mailbox. It's also common when an employee takes an extended vacation and someone else needs to monitor the mailbox until the employee returns.

To delegate access to a mailbox, open the mailbox in the Mailboxes page of the Exchange Admin Center, and then follow these steps:

1. **Select Manage Mailbox Delegation.**

   The Manage Mailbox Delegation pane is displayed, as shown in Figure 2-8.

   There are three types of delegation you can set:

   - **Read and Manage:** Also referred to as *Full Access* permissions. A user who has Full Access has complete control over the mailbox.
   - **Send As:** This option provides more limited access; the delegated user cannot fully control the mailbox but can send messages as if they were sent by the mailbox owner.
   - **Send on Behalf:** This option allows the delegated user to send email from the mailbox, but the From address will clearly indicate that the message was sent on behalf of the mailbox owner by the delegated user.

2. **Click Edit next to the type of permission you want to grant.**

   The pane shown in Figure 2-9 appears.

3. **Click Add Permissions.**

   The pane shown in Figure 2-10 appears.

4. **When asked to confirm the delegation, click Confirm.**

   You will be notified that the mailbox permissions were updated.

5. **Click the X at the upper right of the Manage Mailbox Delegation page to close the page.**

Congratulations! You're now an expert at delegation.

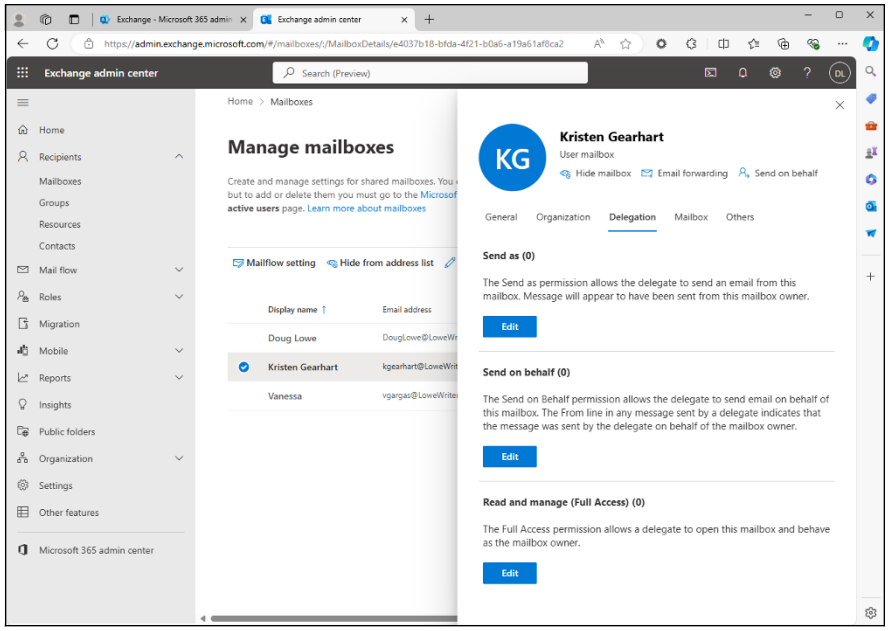


FIGURE 2-8: Managing mailbox delegation.

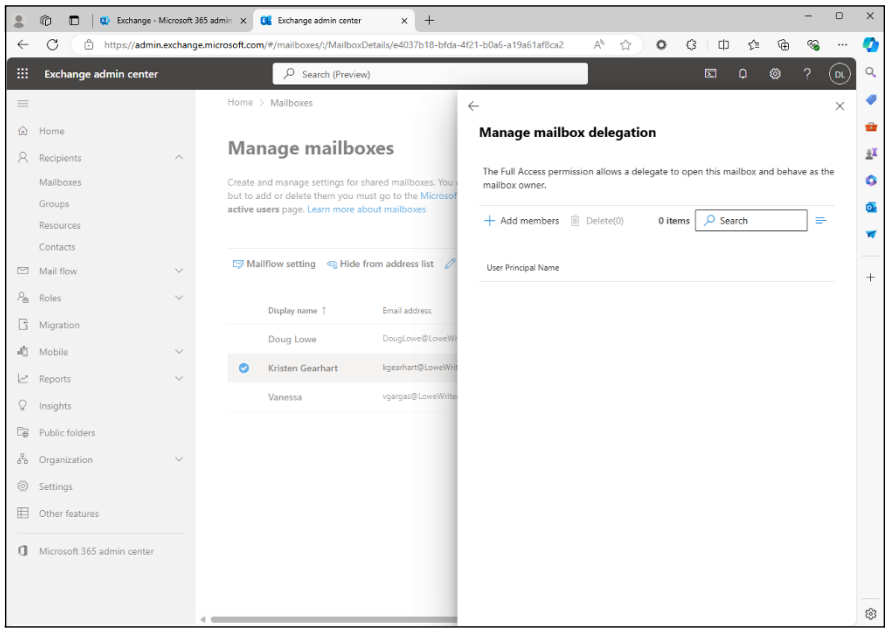


FIGURE 2-9: Granting permission.

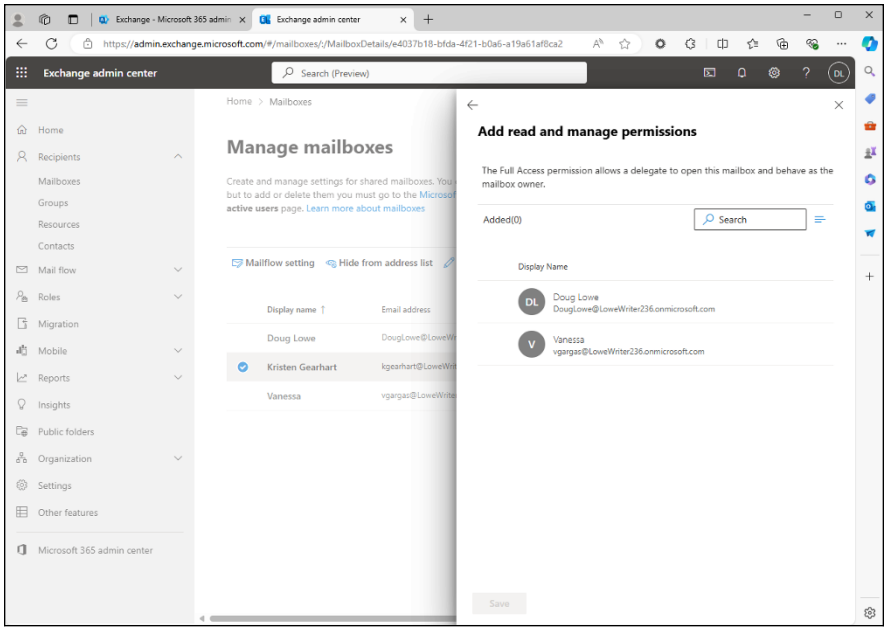


FIGURE 2-10: Selecting the user to which permission will be applied.

Converting a standard mailbox to a shared mailbox

When an employee leaves your organization, it's common to convert that employee's mailbox to a shared mailbox so that email can continue to flow in and other users can access it. This saves you the expense of paying for an Office 365 subscription for a user who no longer works for your organization.

The procedure to convert a Mailbox to a shared mailbox is simple:

1. **Select the user in the Mailboxes administration page.**

   When you select a mailbox, you'll notice the options in the menu bar above the mailbox lists. One of the new options that appears is Convert to Shared Mailbox. (If this option doesn't appear, click the three dots to the right of the Search box.)

2. **Click Convert to Shared Mailbox.**

   Because this is a rather drastic action, Exchange displays a confirmation page asking you to confirm the change.

3. **Click Confirm.**

   The mailbox is converted and a message is displayed to indicate as such.

4. **Click Close.**

   The user's mailbox will now appear as SharedMailbox rather than UserMailbox on the Mailboxes admin page. (It may take a few moments for the mailbox to be converted. You may need to refresh the page to see this change.)



TIP

You can convert the shared mailbox back to a regular mailbox by following the same procedure, but choose Convert to Regular Mailbox rather than Convert to Shared Mailbox in Step 2.

Enabling or disabling mailbox apps

Exchange Mailbox Apps are different ways a user can connect to an Exchange mailbox. To manage this option, click Manage Email Apps Settings in the mailbox details pane. This brings up the Manage Settings for Email Apps pane, shown in Figure 2-11.

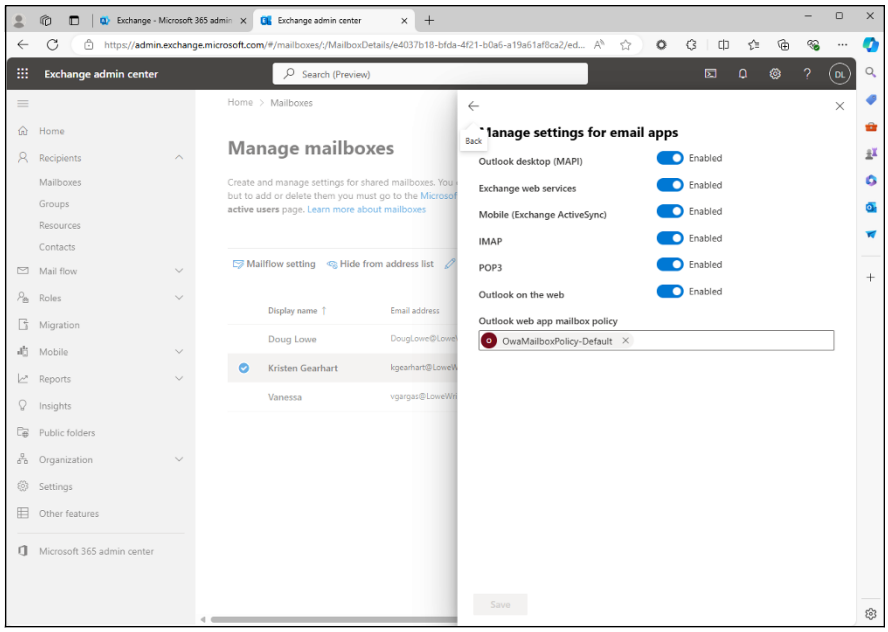


FIGURE 2-11: Managing email apps.

Here are the apps that are controlled from this pane:

- **Outlook Desktop (MAPI):** Enables email using the MAPI protocol, which is used by the Outlook desktop application.

- **Exchange Web Services:** Exchange Web Services (EWS) is a protocol that allows developers to connect to an Exchange mailbox.
- **Mobile (Exchange ActiveSync):** Activates the ActiveSync feature, which allows Exchange data to synchronize with mobile devices such as iPhones or Android devices.
- **IMAP:** Enables email using the IMAP protocol.
- **POP3:** Enables email using the POP protocol.
- **Outlook on the Web:** Lets the user access their Exchange mailbox from a web browser rather than from an Outlook client. With this feature enabled, the user can read email from any computer that has an internet connection.

Creating a forwarder

A *forwarder* is a feature that automatically forwards any incoming email to another email address. This feature is most often used when an employee is on vacation or leave and the employee's manager has requested that someone else temporarily handle the absent employee's email.

To configure a forwarder, follow these steps:

1. **In Exchange Administrative Center, open the User Mailbox details pane for the user (refer to Figure 2-5).**
2. **Click Manage Email Forwarding.**

   The Mail Flow Settings are displayed, as shown in Figure 2-12.
3. **Flip the switch to turn on email forwarding.**

   Additional options appear. (Note that in Figure 2-12, I have already flipped the switch, so the options are already visible.)
4. **Set the desired forwarding options.**

   You can forward to an internal email address or forward to an external address. You can also retain the forwarded email in the user's mailbox or have it deleted after it has been forwarded.
5. **Click Save.**

   The message `Your Preferences Were Saved Successfully` is displayed.

   You're done!

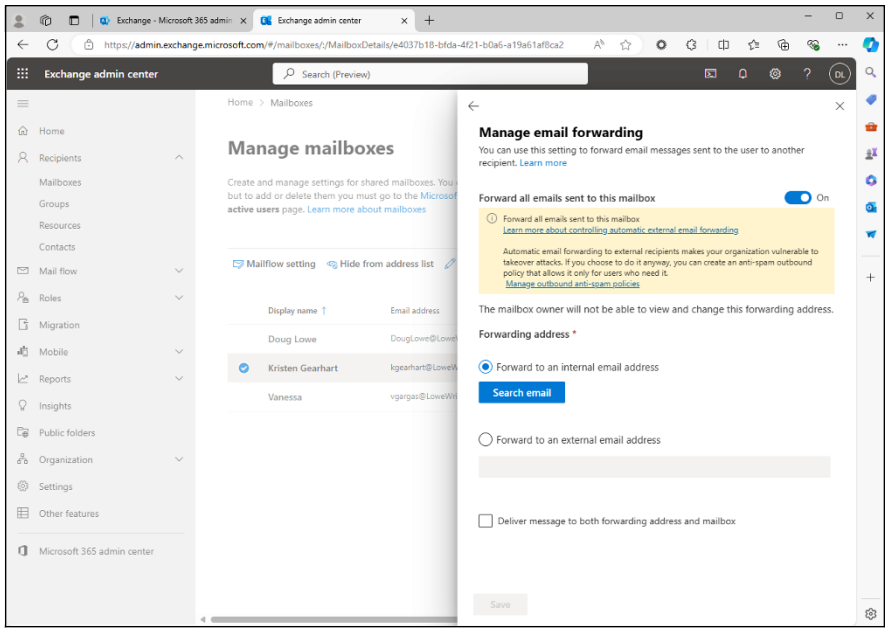


FIGURE 2-12: Managing email forwarding.

Creating a Shared Mailbox

As I mentioned earlier, a *shared mailbox* is a mailbox that can be accessed by multiple people. You can create a shared mailbox from the Recipients page by following these steps:

1. **Open the Mailbox administration page in Exchange Admin Center.**
2. **Click Add a Shared Mailbox.**

   The Create New Shared Mailbox pane appears, as shown in Figure 2-13.
3. **Enter the Display Name, Email Address, and Domain for the new shared mailbox.**

   You can also enter an alias if you want.

   For this example, I used Support Mailbox as the Display Name and `support@lowewriter.onmicrosoft.com` as the Email Name and Domain.
4. **Click Create.**

   The Shared Mailbox Created Successfully pane is displayed, as shown in Figure 2-14.

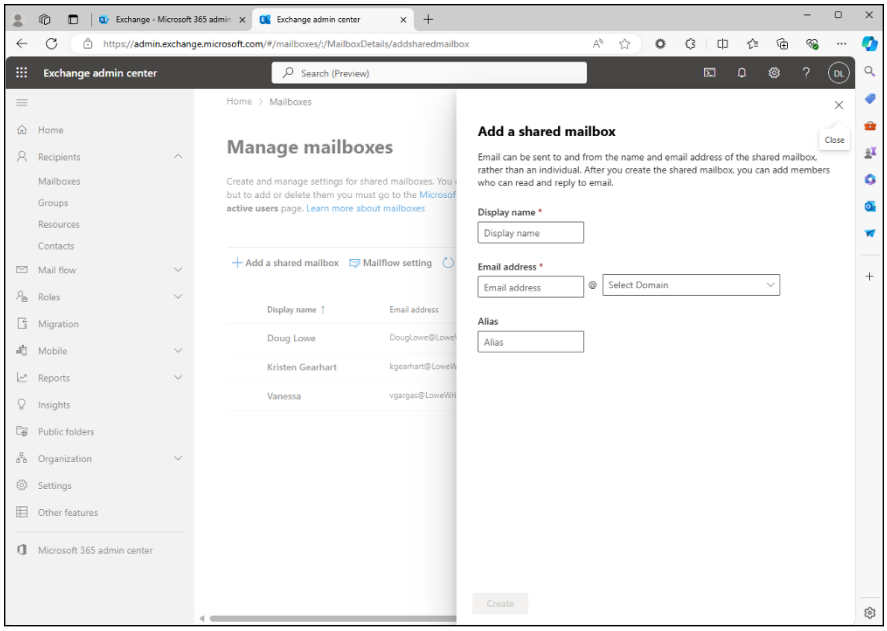


FIGURE 2-13: Creating a shared mailbox.

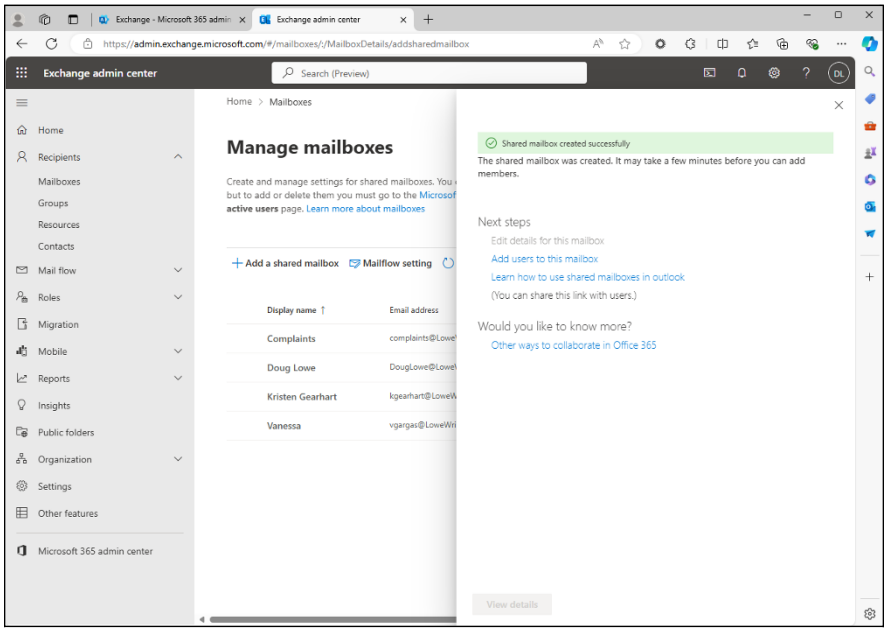


FIGURE 2-14: The shared mailbox has been created!

5. **Choose Add Users to This Mailbox.**

   The Manage Shared Mailbox Members pane appears, as shown in Figure 2-15.

6. **Click Add Members.**

   This displays a list of available users you can add, as shown in Figure 2-16.

7. **Select the users you want to add to the mailbox and click Save.**

   If necessary, use the search box to narrow the list of names displayed.

   When you click Save, a confirmation message will be displayed to ensure that you really want to add the selected users to the shared mailbox.

8. **Click Confirm.**

   After the users have been added, a confirmation message is displayed.

9. **Click the X at the upper-right corner to close the pane.**

   The shared mailbox now appears in the mailbox list.

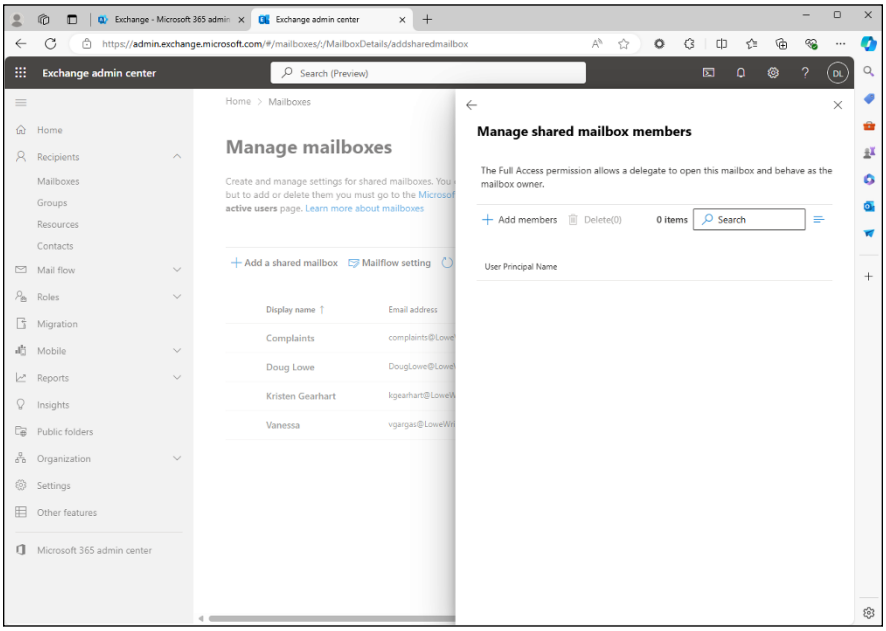


FIGURE 2-15: Managing shared mailbox members.

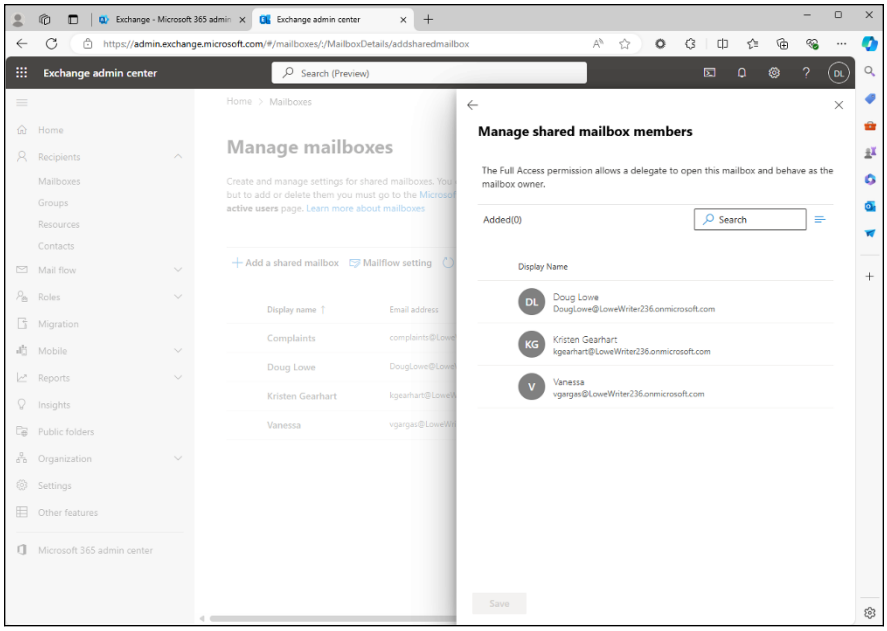


FIGURE 2-16: Adding users to a shared mailbox.

IN THIS CHAPTER

- Looking at Teams
- Understanding how Teams works
- Exploring the Teams admin center
- Setting restrictions on Teams features

Chapter 3
Administering Teams

This chapter covers the basics of administering Microsoft Teams. If you've worked with Teams as an end user, you may not be aware that there's an administrative side to Teams. After all, most Teams environments are set up so that users can create and fully customize new teams, adding public and private channels, building out a channel with new tabs and apps, and admitting members.

All of that's true when Teams is left in its default configuration, wide open so that anyone who uses Teams is sort of their own Teams administrator. But there's a lot going on behind the scenes with Teams. And many organizations have discovered that leaving it wide open creates an environment like the Wild West. It won't take long before you have hundreds or even thousands of teams, with no consistency in naming standards and little concern for security.

That's where the Teams Admin Center comes in. This chapter is a brief introduction to this helpful administration tool. I start with a bird's-eye view of how Teams works behind the curtain so you'll understand what really happens when you (or a user) creates a team. Then I show you how to tame your Teams environment so that it doesn't become the Wild West.

What Is Teams?

In short, Teams is a chat-based online collaboration platform that is designed for people who work together on a common project.

Teams has two distinct modes of operation. The first is ad-hoc chat, in which two or more people connect for a brief conversation. In an ad-hoc chat, users can exchange files, start an audio call or an online meeting, and invite other people into the chat. Ad-hoc chats are by nature ephemeral; when the topic of the chat has been resolved, the chat comes to an end. Figure 3-1 shows an ad-hoc chat in progress.

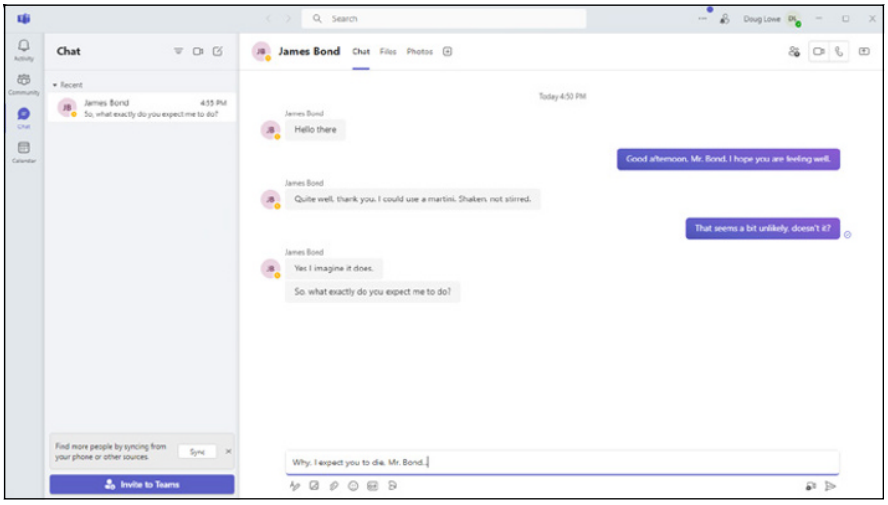


FIGURE 3-1: An ad-hoc chat in Teams.

The second mode for Teams is working in organized teams. In Microsoft Teams, a *team* is a designated space where a group of users come together for ongoing work. Figure 3-2 shows a team named Super Important Project.

Teams are organized around three basic concepts:

- **Team:** A *team* is a virtual workspace where a group of people gather to collaborate and communicate. Some teams are organization-wide, so the entire company participates in them. Other teams are group-specific, so membership is limited to specific people.
- **Channel:** A channel is an organized topic within a team. All activities within a team take place in a channel. The most noticeable activity in a channel is chatting, which provides a way to have online conversations about the

channel's topic. But that's not all that can happen in a channel — you can use channel tabs to access other channel features as well (see the next bullet point).

All teams have a top-level channel called General, which is where conversations about the team itself occur. For example, messages are automatically generated in the General channel when new people or features are added to a team.

Many teams also have a channel devoted to off-topic water-cooler-type conversations. This channel will often have a whimsical name like *Fun* or *Water Cooler* or *Shoot the Breeze.* This channel is for discussions that aren't work-related, such as what happened at the big game the night before, or whether anyone wants to grab lunch today.

Beyond that, individual teams can have additional channels pertaining to specific topics. The team shown in Figure 3-2 has a an additional channel called *Announcements,* which will be used for announcements pertaining to the team.

- **Tabs:** Each channel has a row of tabs that provide access to different features of the channel. All channels have a Posts tab, which is where chatting occurs, and a Files tab, which is a place to save files of interest to the channel. Other tabs may appear as well. For example, you may find a tab for OneNote, Wiki, Lists, and other useful ways to organize information and collaborate within a channel.



REMEMBER

Tabs are specific to channels. Each channel within a team has its own set of tabs.

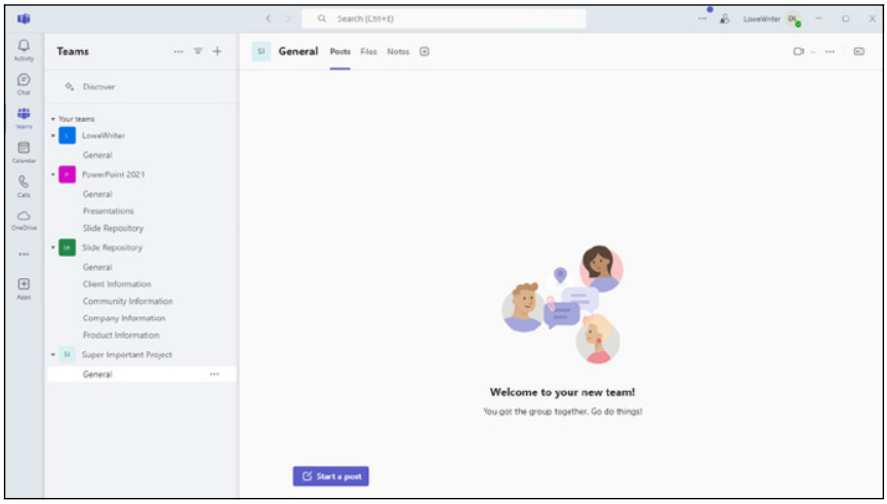


FIGURE 3-2: A typical team.

Out of the gate, any user on Teams can create a new team and add channels to it. In Figure 3-2, you may have noticed the icon in the center of the page labeled Create More Channels. If you click that icon, you see the page shown in Figure 3-3.

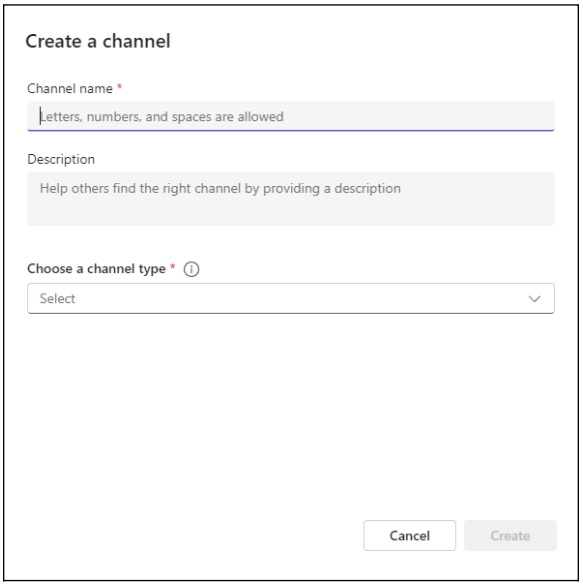


FIGURE 3-3: Anyone can create a new channel.

Anyone can create a new tab in a channel as well, simply by clicking the plus sign (+) to the right of the tabs. Clicking the plus sign brings up the page shown in Figure 3-4, which offers a plethora of choices for creating new tabs.

And finally, anyone can create an entirely new team by clicking the Join or Create a Team link at the bottom of the menu on the left side of the Teams window. Doing so leads you to a list of templates you can choose to act as a starting point for your new team.

This has been an incredibly simplified introduction to Teams, but it's enough to understand the main problem with letting Teams go unmanaged: It won't take long before your teams are overwhelmed with clutter. Too many teams, each having too many channels, each having too many tabs.

Solid organization is essential to the success of Teams in your company, and to achieve that you'll need to lock teams down a bit so that only designated people can manage the structure of your teams.

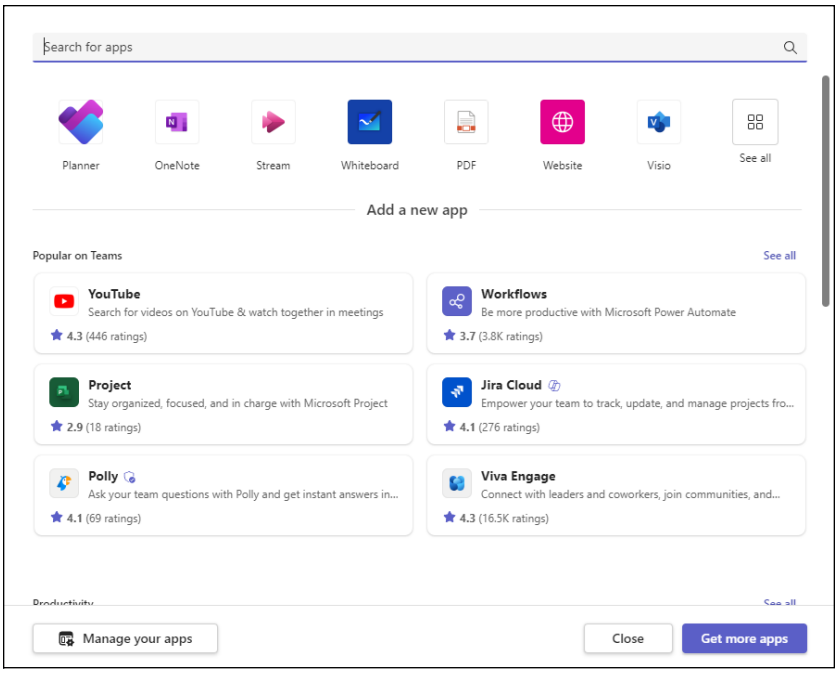


FIGURE 3-4: Anyone can add a tab to a channel.

Before I delve into using the Teams Admin Center to add such controls, however, I want to make sure you have a solid understanding of what's going on behind the scenes when someone creates a team, a channel, or a tab. The next section dives into the shallow end of the Teams architecture pool.

A Brief Look at How Teams Works

Teams is an exquisitely complicated product that is built on a solid layer of various features provided by Office 365 and Azure. It would take many pages — more than we have here — to give a complete explanation of how Teams works. But I want to give you a bird's-eye view of the basic architecture of Teams. This information will help you better understand the administration tasks required to make Teams a success in your organization.

Figure 3-5 shows the components of Office 365 that are involved when you create a team.

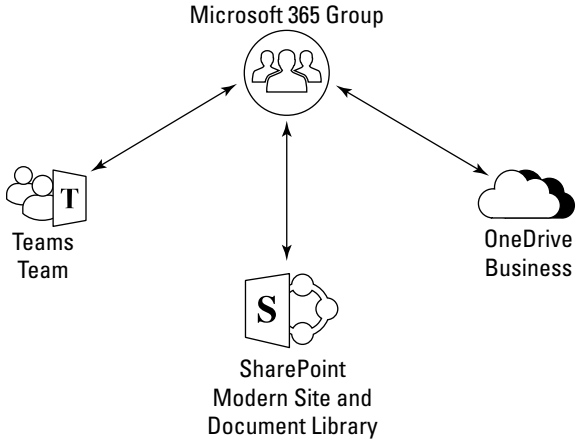


FIGURE 3-5: An overview of Microsoft Teams architecture.

The following sections describe the role each of these Office 365 components plays in Microsoft Teams.

Microsoft 365 Group

Contrary to what you may expect, Teams itself is not at the center of this hub: Microsoft 365 Groups are. A Microsoft 365 Group is the organizing principle around which teams are created.

So, to understand Teams, you need to first understand how Microsoft 365 Groups work. You're probably familiar with standard Active Directory security and distribution groups. A Microsoft 365 Group is similar, but it provides additional capabilities.

With a Microsoft 365 Group, you can provide access to a collection of shared resources, including the following:

- A team in Microsoft Teams
- A modern SharePoint site in SharePoint
- The SharePoint Document Library
- A shared mailbox in Outlook
- A shared calendar
- A OneNote notebook

Other shared resources are available to a Microsoft 365 Group beyond these, but the ones in the preceding list are most relevant to Teams.

When you create a team in Microsoft Teams, a Microsoft 365 Group is automatically created, and the resources detailed in the preceding list are automatically provisioned for the group. As you can see, a lot of work goes on behind the scenes when you create a team!

Figure 3-6 shows a Microsoft 365 Group that was created for a team named Super Important Project. The team is selected and the Membership details are shown so you can see who the owner and members of the team are.

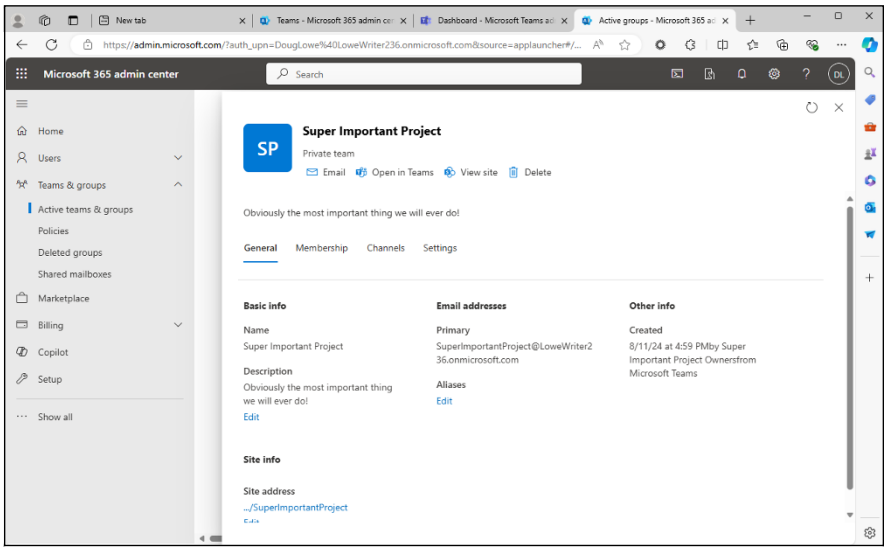


FIGURE 3-6: A Microsoft 365 Group created for a team.

Note that if you create a Microsoft 365 Group directly from the M365 Admin Center, the corresponding resources are not automatically created. However, you can easily create a team from an existing Microsoft 365 Group, either in Teams or in the Teams Admin Center. I'll show you how to do it from the Admin Center later in this chapter, in the "Managing Teams" section.

In addition to automatically provisioning all the services required to support Microsoft Teams, the Microsoft 365 Group automatically manages permissions for those services.

The membership in Teams and the membership of its corresponding Microsoft 365 Group are automatically kept in sync with one another. So, if you add or remove members for the team, they're added or removed for the Group, and vice versa.

SharePoint

When you create a team in Microsoft Teams, two critical components of SharePoint are also created:

- **A modern SharePoint site,** which provides a basic SharePoint page for the team
- **The SharePoint Document Library,** which provides a library for storing files related to the team

Users can access the SharePoint site for a team by clicking the ellipses (. . .) at the upper right of the team and choosing Open in SharePoint. For example, Figure 3-7 shows the modern SharePoint site that was created for the Super Important Project team. (When you click Open in SharePoint, you're actually taken to the site's Documents page. For this figure, I clicked Home to display the site's home page, so you can see a few of the SharePoint features included in the site, such as News and Activity.)

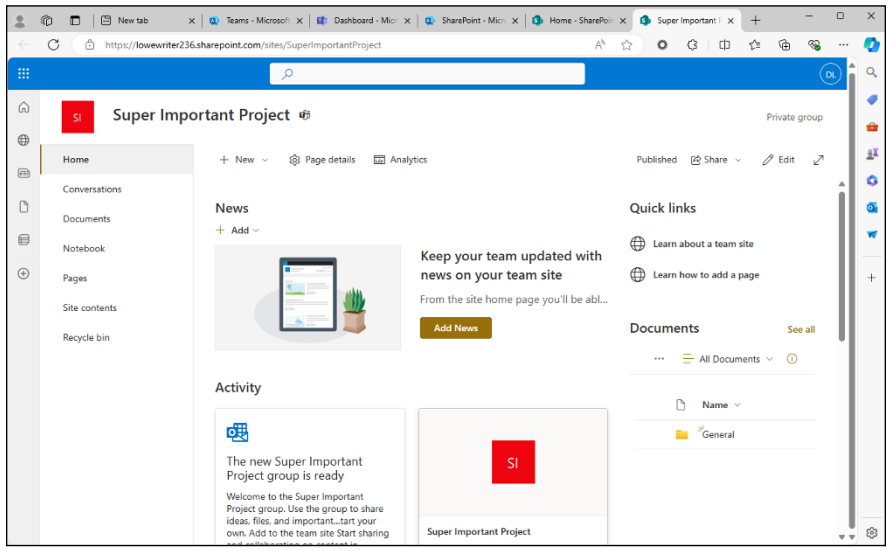


FIGURE 3-7: A modern SharePoint site is created for every team.

The SharePoint Document Library plays a key role in the operation of Microsoft Teams, because it's the repository used for all files stored for the team. Figure 3-8 shows the Document Library for the Super Important Project Team. As you can see, it includes a folder for each of the team's two channels: General and Something or Other.

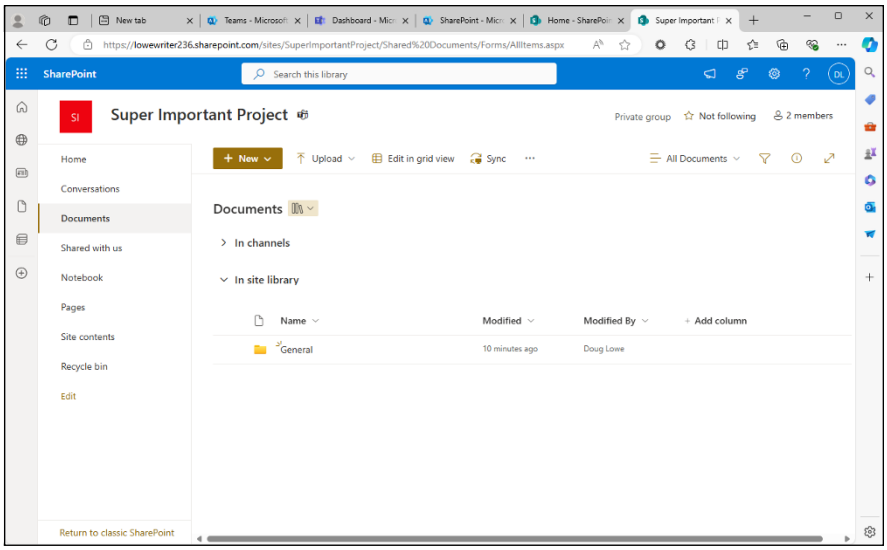


FIGURE 3-8: A SharePoint Document Library is created for every team.

When you create a new channel in a team, a new folder is automatically created in the Document Library. The Files tab in each channel points directly to the corresponding folders in the Document Library.

Note that you can also add the home page from the SharePoint site created for a team as a tab to any channel in the team. Figure 3-9 shows an example of this. In addition, you can add any arbitrary SharePoint page as a tab.

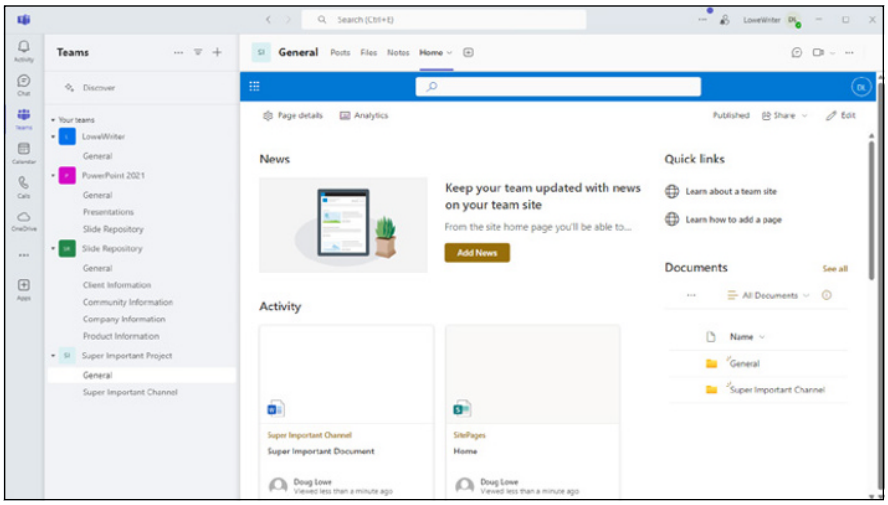


FIGURE 3-9: The SharePoint home page can be viewed in a channel tab.

OneDrive for Business

The SharePoint Document Library provides file storage for teams in Microsoft Teams. But what about storage for files that are exchanged in an-hoc chats? That's where OneDrive for Business comes in.

When a user posts a file to the Files tab in a chat, the file is uploaded to that user's OneDrive for Business. Then, the file is shared with everyone else on the chat.

Figure 3-10 shows how this works. This figure shows the Files tab of a chat between two Teams users. As you can see, the file `menu.png` has been posted to the chat. The figure also shows the user's OneDrive for Business account in `www.office.com`. As you can see, Teams created a folder in OneDrive for Business named Microsoft Teams Chat Files and uploaded the `menu.png` file to that folder. The figure also shows that the file has been shared with the participants in the chat.

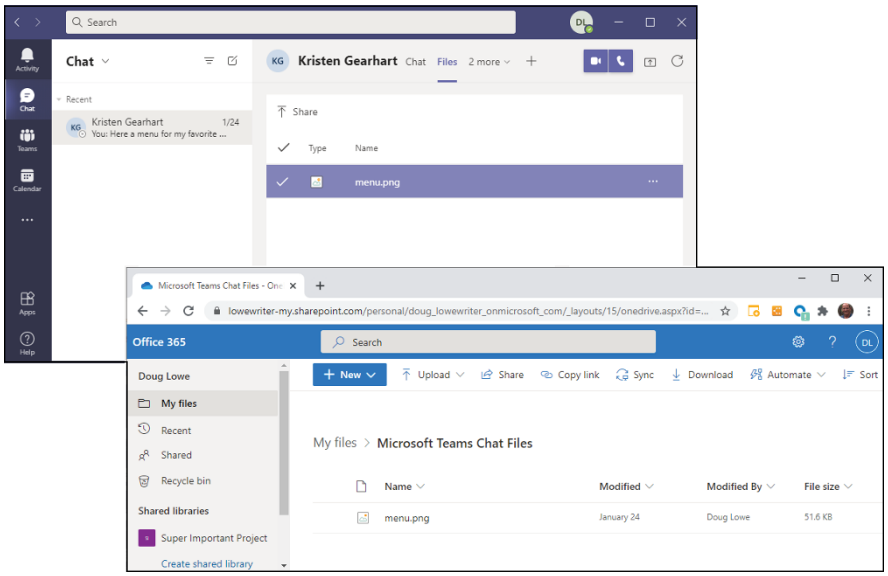


FIGURE 3-10: A file shared in a chat.

Using the Teams Admin Center

Although you can do a lot of Teams administration directly from Teams itself, many tasks can only be accomplished from the Teams Admin Center. To access it, sign in to `https://admin.microsoft.com`, click Show All, and then click Teams. Figure 3-11 shows the Teams Admin Center Dashboard.

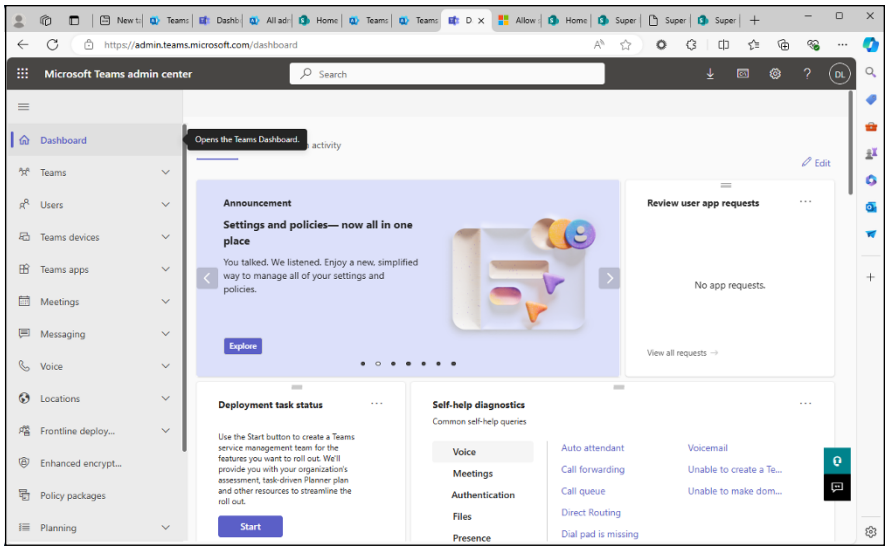


FIGURE 3-11: The Teams Admin Center Dashboard.

As you can see, there are a slew of choices in the menu on the left side of the Admin Center:

- **Dashboard:** Displays a summary of Teams tasks to be done, as well as useful information about Teams activity.
- **Teams:** Manage individual teams, create and update policies, and create templates for new teams.
- **Users:** Manage Teams users.
- **Teams Devices:** Manage devices including phones, useful if you use Teams as your phone solution.
- **Teams Apps:** Manage apps that can be used in Teams.
- **Meetings:** Manage meeting policies and settings.
- **Messaging:** Manage policies for chat and channel messaging.
- **Voice:** Manage voice settings. This option is useful if you use Teams as your phone solution.
- **Locations:** Manage locations. This option is useful if you use Teams as your phone solution, allowing you to set Emergency 911 information.
- **Frontline Deployment:** Manage Teams features that are designed to support frontline workers who generally don't have computers. Instead, they interact with Teams primarily via a smartphone.
- **Enhanced Encryption:** Manage advanced encryption features.

- **Policy Packages:** Manage packages of policies for various user types.
- **Planning:** Plan your Teams deployment.
- **Analytics & Reports:** Access useful information about your organization's utilization of Teams.
- **Notifications & Alerts:** Manage settings that control various notifications and alerts.

Managing Teams

Figure 3-12 shows the Manage Teams page, which you can reach by selecting Teams and then Manage Teams in the left menu pane. This page lists all the teams in your organization, along with useful information such as the number of standard and private channels, the number of members and owners, and so on.

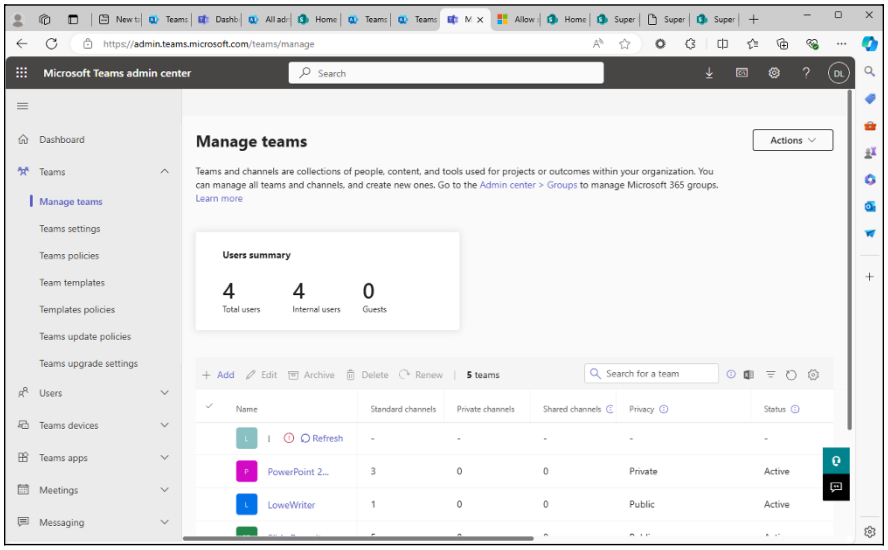


FIGURE 3-12: Managing teams.

Note that you can change the columns that are displayed on this page by clicking the gear icon to the right of the list (shown in the margin). You can select from among the following columns:

- Team name
- Standard channels

- Private channels
- Team members
- Owners
- Guests
- Privacy
- Status
- Description
- Sensitivity
- Classification
- Group ID
- Expiration date



To create a new team, click the + Add button at the left side of the toolbar above the list of teams (shown in the margin). This brings up the panel shown in Figure 3-13, which allows you to enter the name and description for the new team. You can also specify team owners (you'll be added as a team owner by default), and you can use the Privacy drop-down to specify whether the group is private or public. After you've entered this information, click Apply to create the team.

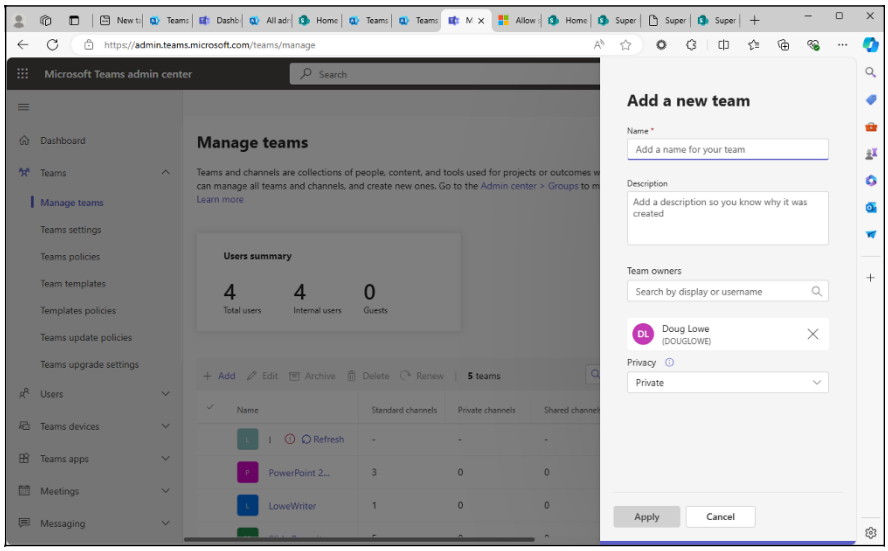


FIGURE 3-13: Creating a new team.

To see the details for a particular team, just click the team in the list. Figure 3-14 shows the detailed information for the Super Important Project team.

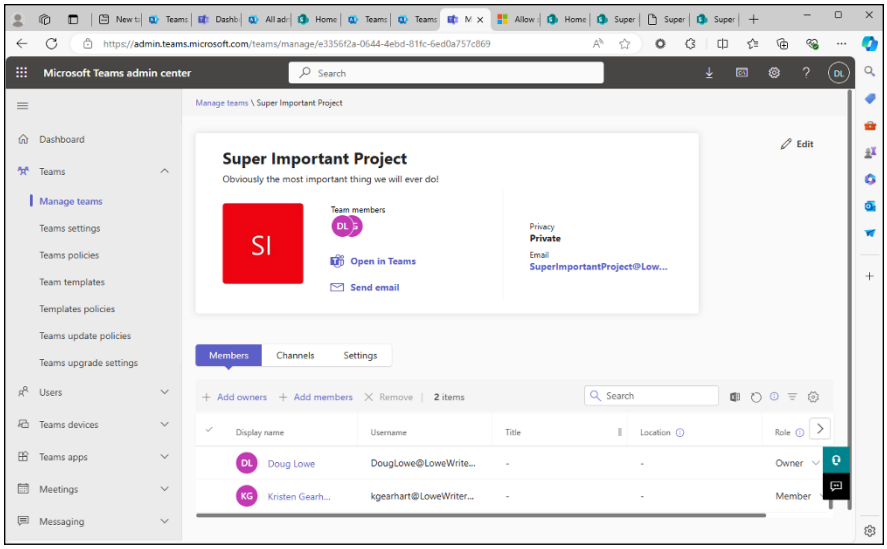


FIGURE 3-14: The Super Important Project team details.

To change information for the team, click Edit on the right side of the screen. The Edit Team pane, shown in Figure 3-15, appears. This pane lets you change the name, description, and private settings for the team.

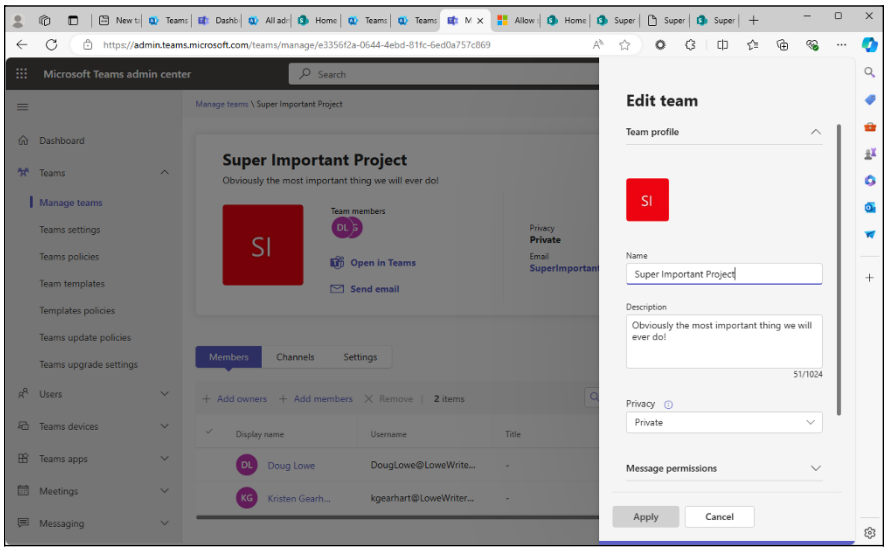


FIGURE 3-15: Editing the Super Important Project team.

To change settings for team channels, click the Channels tab on the team details page. This displays a list of the channels for the team, as shown in Figure 3-16.

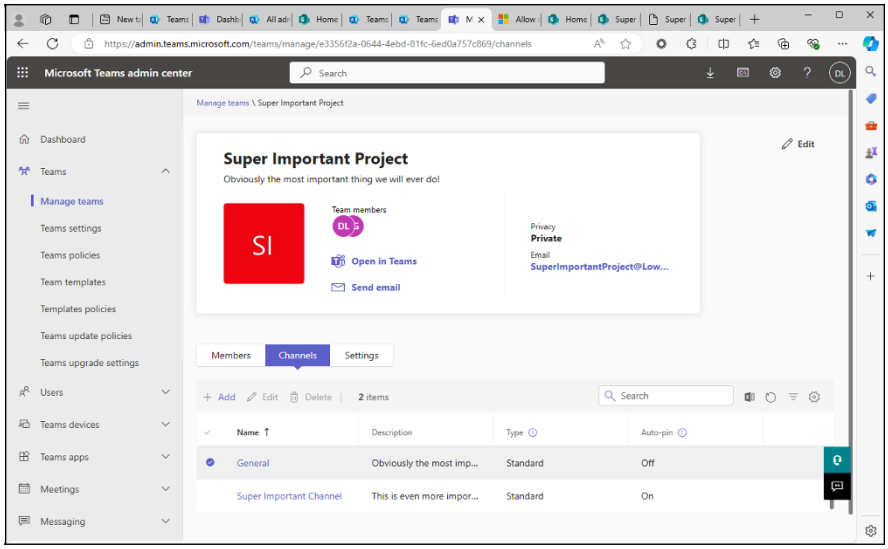


FIGURE 3-16: Team channels.

The Edit Team pane also includes Conversations and Channels groups you can expand to restrict certain features for the teams. Figure 3-16 shows these two groups expanded.

To see details about a particular channel, click the channel to open the channel page. Here, you can see the channel name, description, and members.

8 Implementing Linux

Contents at a Glance

CHAPTER 1: **Installing a Linux Server** 659

CHAPTER 2: **Linux Administration** 673

CHAPTER 3: **Basic Linux Network Configuration** 705

CHAPTER 4: **Running DHCP and DNS** 717

CHAPTER 5: **Linux Commands** 725

IN THIS CHAPTER

- Getting ready to install Linux
- Installing Linux
- Completing the setup

Chapter 1
Installing a Linux Server

This chapter presents the procedures that you need to follow to install Linux on a server computer. The details provided are specifically for Fedora 40, a free Linux distribution sponsored by Red Hat. However, the procedures for installing other distributions of Linux are similar, so you won't have any trouble adapting these procedures if you're using a different distribution.

Planning a Linux Server Installation

Before you begin the installation program, you need to make a number of preliminary decisions. The following sections describe the decisions that you need to make before you install Linux.

Checking system requirements

Before you install Linux, make sure that the computer meets the minimum requirements. Although the minimum requirements for Linux are considerably less than those for the latest version of Windows Server, you can't run Linux on an abacus. The following paragraphs summarize the minimum capabilities you need:

- **A 1GHz processor or faster:** You probably won't find a computer that doesn't meet this requirement.

- **1GB of RAM or more:** Today most computers have at least 4GB of RAM, so the 1GB minimum isn't a problem.
- **A hard drive with enough free space to hold the packages that you need to install:** A suitable minimum is 10GB.
- **A CD or DVD-ROM drive from which to install the operating system**
- **Just about any video card and monitor combination:** You don't need anything fancy for a server. In fact, fancy video cards often lead to hardware compatibility issues. Stick to a basic video card.
- **An Ethernet network interface**

For the purposes of this chapter, I'll be installing Fedora into a Hyper-V virtual machine configured with 4GB of RAM, two processor cores, and 100GB of disk space. (For information about setting up a Hyper-V virtual machine, refer to Book 5, Chapter 1.)

Choosing a distribution

Because the *kernel* (that is, the core operating functions) of the Linux operating system is free, several companies have created their own *distributions* of Linux, which include the Linux OS along with a bundle of packages, such as administration tools, web servers, and other useful utilities, as well as printed documentation.

The following are some of the more popular Linux distributions:

- **Fedora:** One of the popular Linux distributions. You can download Fedora free from `http://getfedora.org`. You can also obtain it by purchasing any of several books on Fedora that include the Fedora distribution on DVD or CD-ROM.

  Fedora comes in two editions: a Workstation edition and a Server edition. The main difference between the two is that the Workstation edition includes a graphical user interface (GUI), whereas the Server edition relies on the command-line interface. I recommend you use the Server edition unless you plan on using Fedora as a desktop workstation rather than as a server.

  All the examples in this book are based on Fedora Server 40.
- **Ubuntu:** A Linux distribution that has gained popularity in recent years. It focuses on ease of use. For more information, go to `www.ubuntu.com`.
- **SUSE:** Pronounced *SOO-zuh,* like the name of the famous composer of marches; a popular Linux distribution sponsored by Novell. You can find more information at `www.suse.com`.

» **Slackware:** One of the oldest Linux distributions and still popular, especially among Linux old-timers. A full installation of Slackware gives you all the tools that you need to set up a network or internet server. See `www.slackware.com` for more information.

All distributions of Linux include the same core components: the Linux kernel, an X Server, popular windows managers such as GNOME and KDE, internet programs such as Apache, Sendmail, and so on. However, not all Linux distributions are created equal. In particular, the manufacturer of each distribution creates its own installation and configuration programs to install and configure Linux.

The installation program is what makes or breaks a Linux distribution. All the distributions I list in this section have easy-to-use installation programs that automatically detect the hardware that's present on your computer and configure Linux to work with that hardware, thus eliminating most — if not all — manual configuration chores. The installation programs also let you select the Linux packages that you want to install and let you set up one or more user accounts besides the root account.

I CAN'T SEE MY C: DRIVE!

Linux and Windows have a completely different method of referring to your computer's hard drives and partitions. The differences can take some getting used to for experienced Windows users.

Windows uses a separate letter for each drive and partition on your system. For example, if you have a single drive formatted into three partitions, Windows identifies the partitions as drives `C:`, `D:`, and `E:`. Each of these drives has its own root directory, which can, in turn, contain additional directories used to organize your files. As far as Windows is concerned, drives `C:`, `D:`, and `E:` are completely separate drives even though the drives are actually just partitions on a single drive.

Linux doesn't use drive letters. Instead, Linux combines all the drives and partitions into a single directory hierarchy. In Linux, one of the partitions is designated as the *root* partition. The root partition is roughly analogous to the root directory of the `C:` drive on a Windows system. Then, the other partitions can be *mounted* on the root partition and treated as if they were directories on the root partition. For example, you may designate the first partition as the root partition and then mount the second partition as `/user` and the third partition as `/var`. Then, any files stored in the `/user` directory would actually be stored in the second partition, and files stored in the `/var` directory would be stored in the third partition.

(continued)

(continued)

The directory to which a drive mounts is called the drive's *mount point.*

Notice that Linux uses regular forward slash characters (/) to separate directory names rather than the backward slash characters (\) used by Windows. Typing backslashes instead of regular slashes is one of the most common mistakes made by new Linux users.

While I'm on the subject, Linux uses a different convention for naming files, too. In Windows, filenames end in a three- or four-letter extension that's separated from the rest of the filename by a period. The extension is used to indicate the file type. For example, files that end in `.exe` are program files, but files that end in `.doc` are word-processing documents.

Linux doesn't use file extensions, but periods are often used in Linux filenames to separate different parts of the name — and the last part often indicates the file type. For example, `ldap.conf` and `pine.conf` are both configuration files.

Going virtual

Another common way to install Linux is in a virtual machine running within the Windows operating system. In fact, all the examples in this minibook were tested using Hyper-V, the built-in virtualization platform that comes with Windows. You can also use VMware Workstation Player, a free virtualization platform you can download from `www.vmware.com/products/workstation-player/workstation-player-evaluation.html`. For more information, see Book 5.

Deciding on your TCP/IP configuration

Before you install the OS, you should have a plan for how you will implement TCP/IP on the network. Here are some of the things you need to decide or find out:

- The public IP subnet address and mask for your network
- The domain name for the network
- The host name for the server
- Whether the server obtains its address from DHCP
- Whether the server has a static IP address — and if so, what
- Whether the server is a DHCP server

- The default gateway for the server — that is, the IP address of the network's internet router
- Whether the server is a DNS server



TIP

If the server will host TCP/IP servers (such as DHCP or DNS), you'll probably want to assign the server a static IP address.

For more information about planning your TCP/IP configuration, see Book 2. And for more information about reconfiguring network settings after you install Linux, refer to Book 8, Chapter 3.

Installing Fedora Server

After you plan your installation and prepare the computer, you're ready to actually install Linux. The following procedure describes the steps you must follow to install Fedora Server (version 40) on a virtual machine using a downloaded `.iso` file containing the Fedora installation media.

Note that, for this example, I chose to install the server version of Fedora. This version does not include a GUI, so you won't have the convenience of a Windows-like GUI. However, you will have the benefit of a web-based configuration program that lets you configure many server features remotely using a web browser. And you'll have the rich features of Linux command-line configuration at your disposal.

1. **Download the Fedora Server 40 `.iso` file from the Fedora project's download page, connect it to the virtual machine's optical drive, and start the virtual machine.**

   The download is located at `https://getfedora.org`.

   When you start the VM, you'll see the prompt shown in Figure 1-1. Select Install Fedora 40, and then press the Enter key.

2. **Wait a moment while Fedora begins its installation process.**

   A bunch of text messages fly across the screen. Eventually, the first page of the setup program is displayed, as shown in Figure 1-2.

3. **Choose your language, and then click Continue.**

   The Installation Summary page is displayed, as shown in Figure 1-3.



FIGURE 1-1: Are you ready to install Fedora Server?

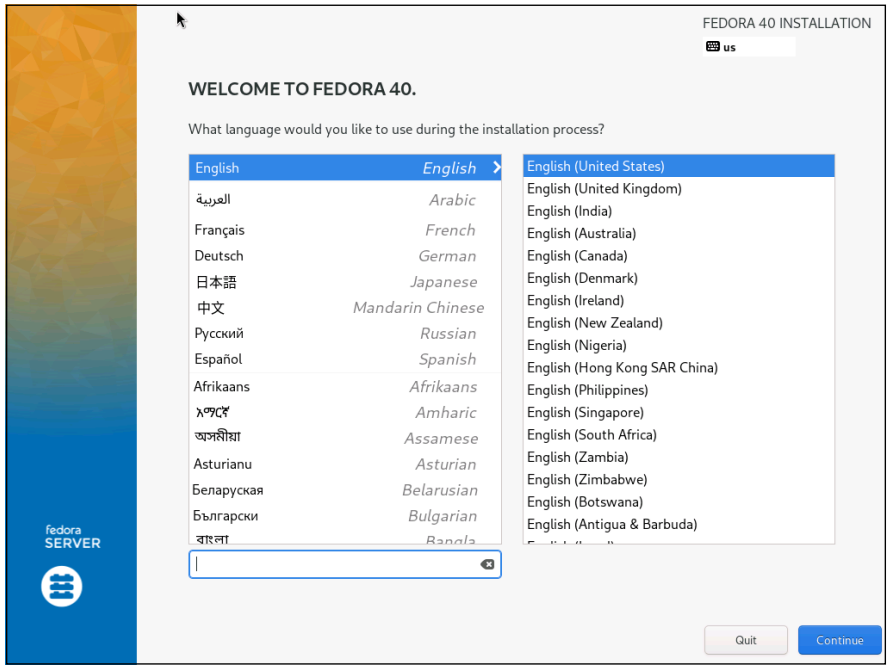


FIGURE 1-2: Choose a language to begin the Fedora installation.

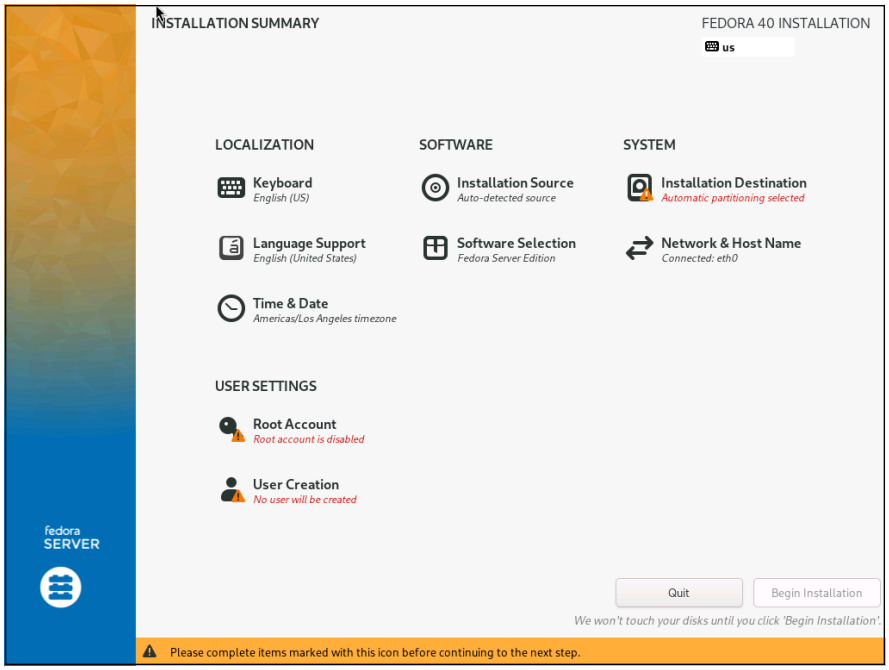


FIGURE 1-3: The Installation Summary page.

This page lets you access a variety of configuration options for your Linux server installation. Here are the ones I cover in this example:

- **Installation Destination:** Lets you configure your disk partitions.
- **Network and Host Name:** Lets you set networking information.
- **Time & Date:** Lets you change the time zone.
- **Software Selection:** Lets you choose optional features for the server.
- **Root Password:** Lets you enable the root account and assign it a password. (In Linux, the root account is equivalent to the Administrator account in Windows.)
- **User Creation:** Lets you create a user account for yourself.

4. **Click Installation Destination.**

   The installation program displays the Installation Destination screen, shown in Figure 1-4.

5. **If necessary, adjust the installation destination settings.**

   The Installation Destination screen indicates the disk volume on which Fedora will be installed. The default setting is to automatically create a partition on the empty primary drive and then install Fedora into the new partition. You can use this page to add additional disks to the server, to change the way storage is managed, and to encrypt disk data.

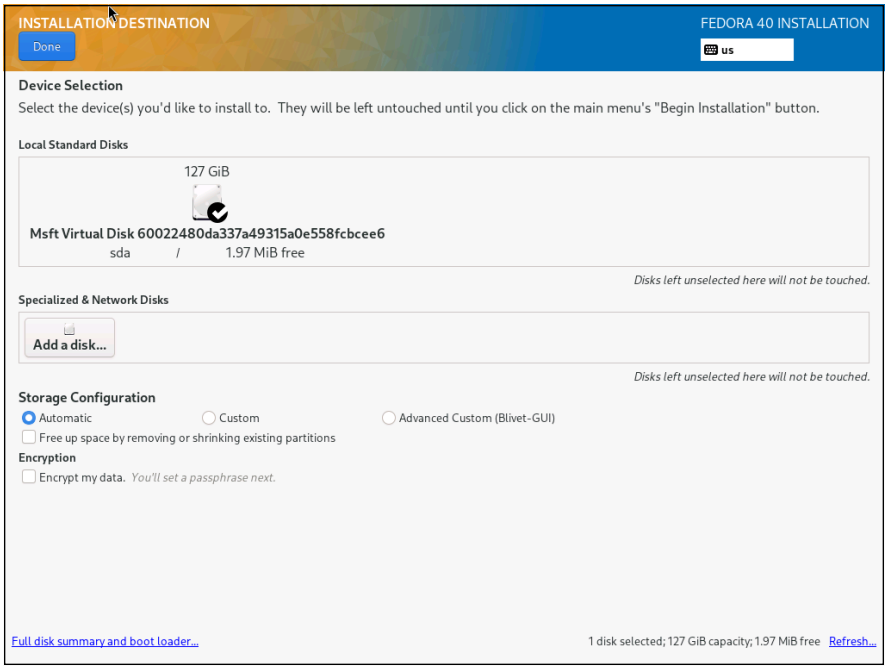


FIGURE 1-4: The Installation Destination screen.

6. **Click Done.**

   You're returned to the Installation Summary screen.

7. **Click Time & Date.**

   The Time & Date page appears, as shown in Figure 1-5.

8. **Select the correct time zone and location, and then click Done.**

   You're returned to the Installation Summary page.

9. **Click Software Selection.**

   The Software Selection page appears, as shown in Figure 1-6. This page lets you select optional software components for the server. The one you're most likely to use is Domain Membership, which allows you to join an Active Directory domain.

10. **Select any optional add-ons you want to include, and then click Done.**

    You're returned again to the Installation Summary page.

11. **Click Network & Host Name.**

    You're taken to the page shown in Figure 1-7, which shows the network interfaces available to the server. In our example, there is just one interface available, named `eth0`.

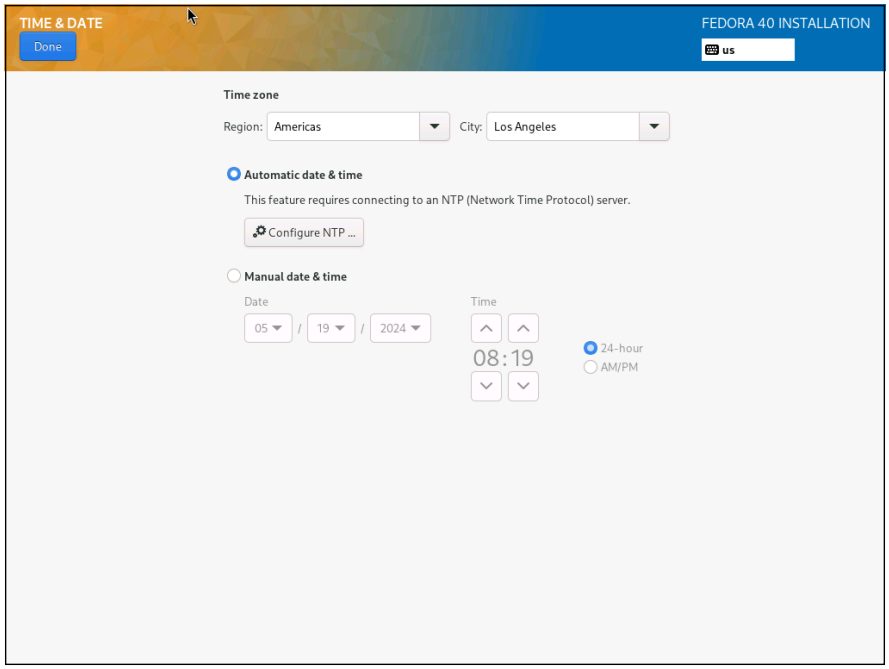


FIGURE 1-5: Choose a time zone.

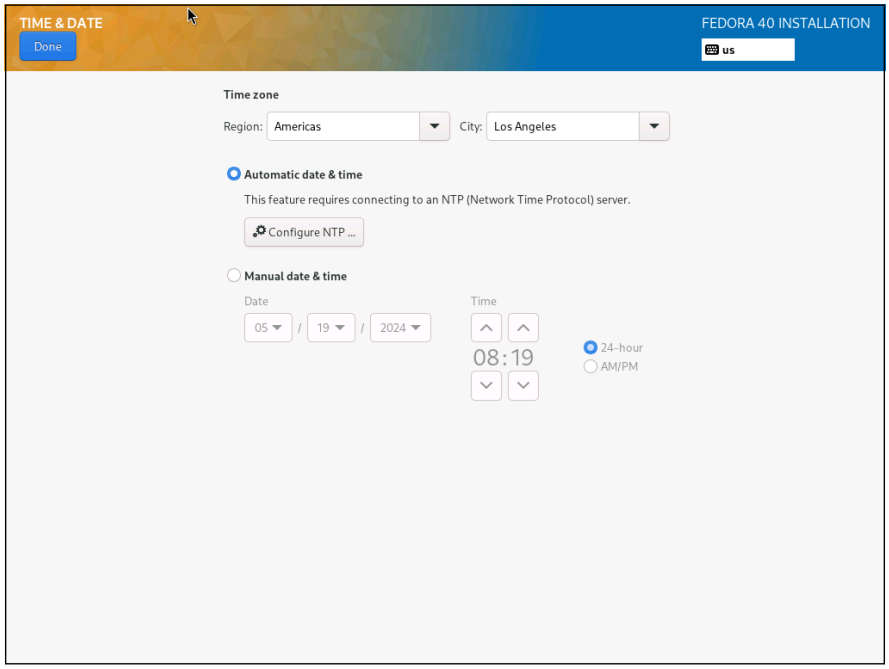


FIGURE 1-6: Choose any optional add-ons you want to install.

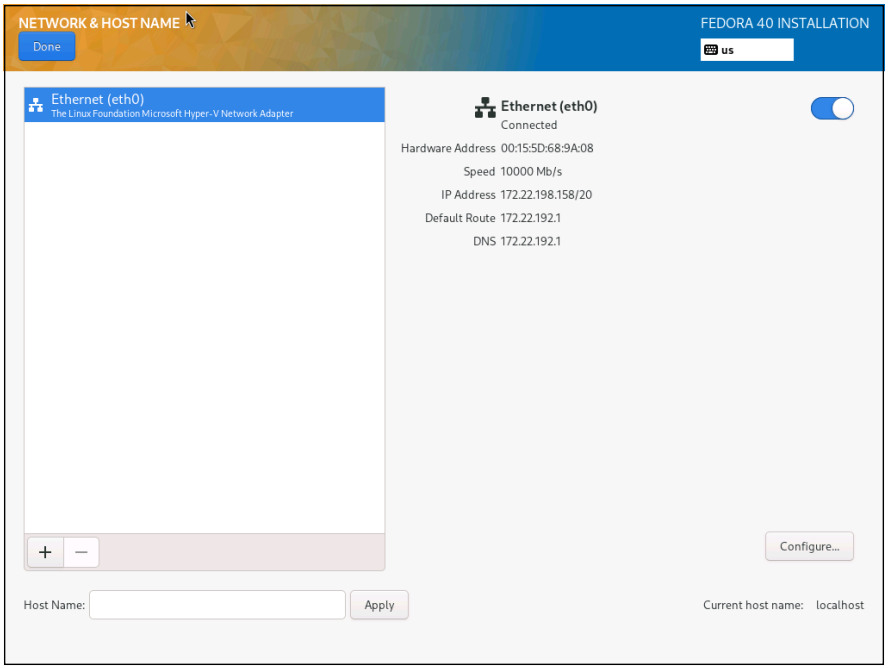


FIGURE 1-7: The Network & Host Name page.

12. **Click Configure.**

    The network configuration editor appears, as shown in Figure 1-8. This editor allows you to change settings for the network interface. In this editor, you can open the IPv4 Settings tab to switch the interface from DHCP to a static IP address. It's a good idea to do this during installation, but I'll skip that step for this example. I'll circle back to applying a static IP address in Book 8, Chapter 3.

13. **Close the editor window, and then click Done.**

    The Installation Summary page comes back into focus.

14. **Click Root account.**

    This takes you to the page that lets you enable the root account and set its password, as shown in Figure 1-9.

15. **To enable the root user, uncheck the Disable Root Account check box, and then enter a strong password for the root user.**

    You'll have to enter the password twice.

16. **Click Done.**

    You're swept back to the Installation Summary page once more.

17. **Click User Creation.**

    The Create User page, shown in Figure 1-10, appears. This page lets you create a new user account.

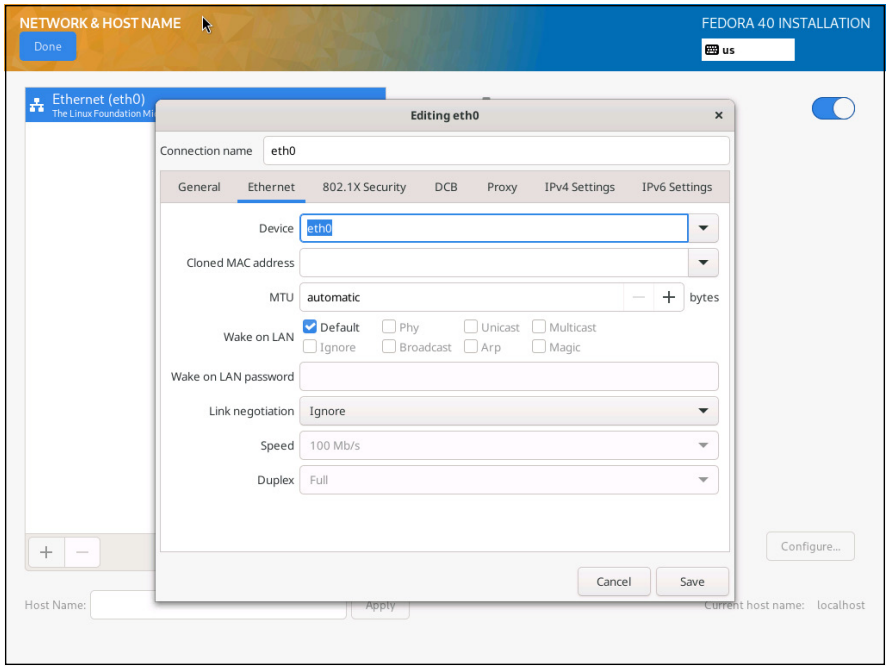


FIGURE 1-8: Editing a network interface.

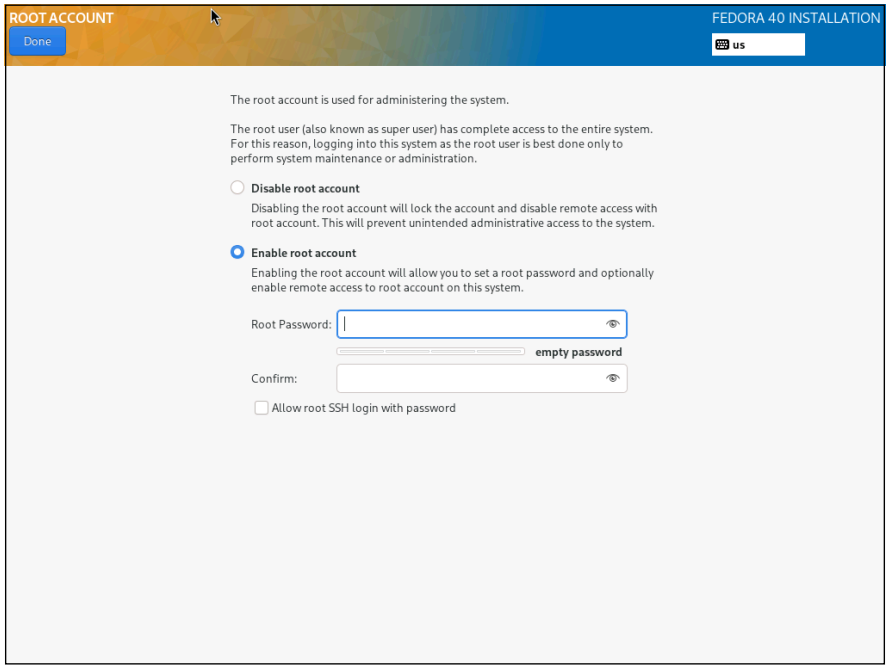


FIGURE 1-9: Enabling the root account.

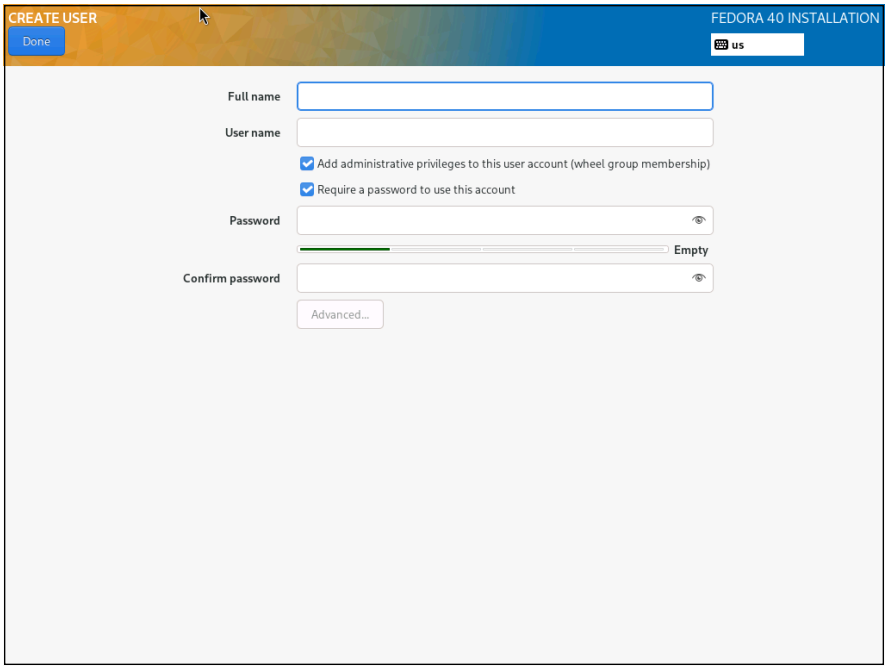


FIGURE 1-10: Creating a new user.

18. **Enter the information for the new user.**

    As you enter the user's full name, Fedora will fill in the username field automatically. If you don't like the username Fedora suggests, you can change it.

    If you want to make the user an administrator, check the Make This User Administrator check box. It's a good idea to designate this user as an administrator so you can avoid using the root account for routine administrative tasks.

    If you want, click the Advanced button to configure additional settings for the user.

19. **Click Done.**

    The Installation Summary page comes back up.

20. **Click Begin Installation.**

    The installation process begins. When the installation is complete, you'll see the completion message shown in Figure 1-11.

21. **Click Reboot System.**

    You're done!

When the server finishes its reboot, you're greeted with a stark Linux login prompt in a command window, as shown in Figure 1-12. From here, you can log in to the server and begin configuring it, as described in the next chapter.

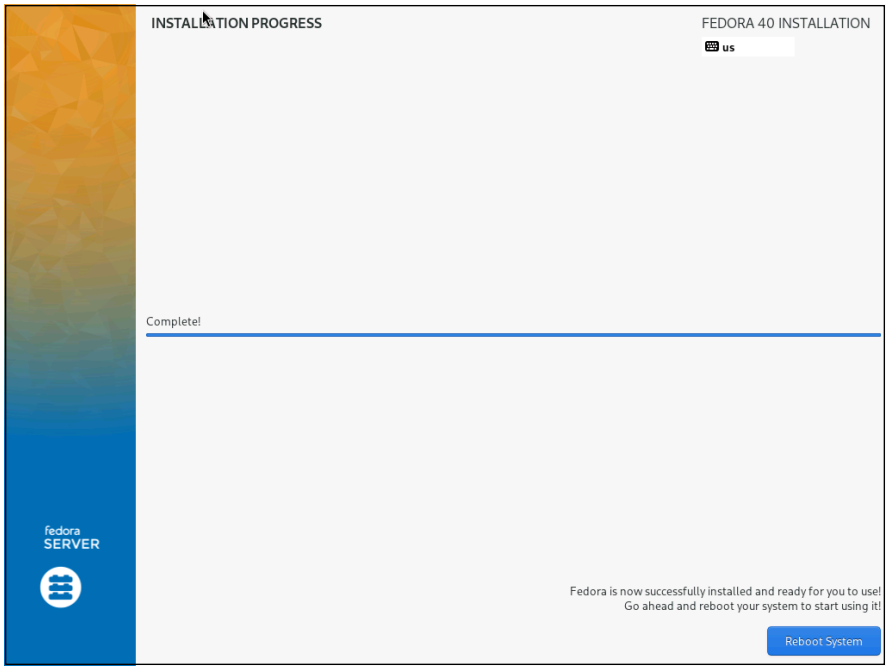


FIGURE 1-11: Fedora has been successfully installed.

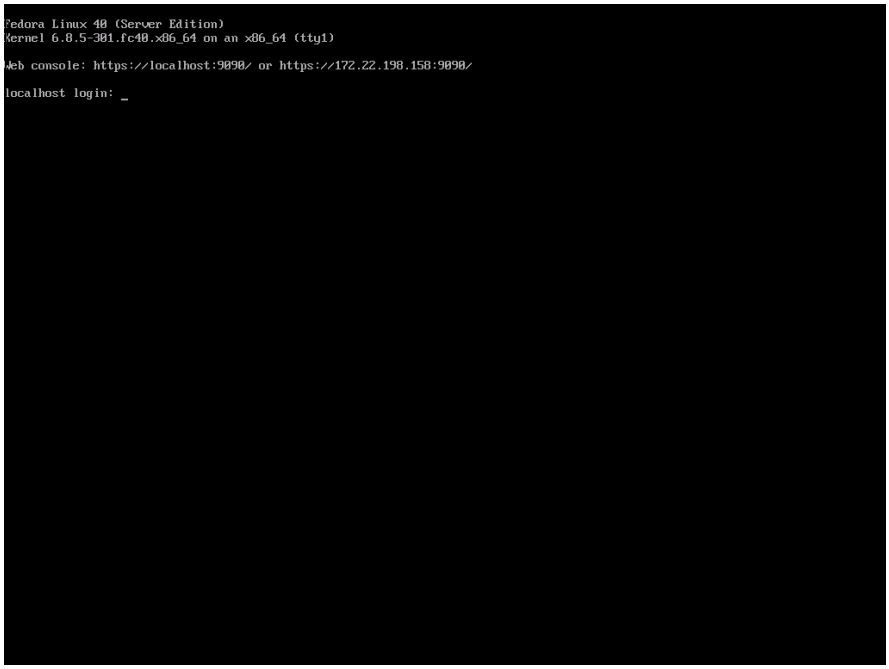


FIGURE 1-12: The stark Linux login prompt.

VIRTUAL CONSOLES AND THE INSTALLATION PROGRAM

Linux is inherently a command-line–oriented OS. Graphical user interfaces — including the installation program's GUI — are provided by an optional component called *X Window System.* However, while you're working with the GUI of the installation program, Linux keeps several additional command-line consoles open. Normally, you don't need to use every one of these consoles during installation. However, if something goes wrong during installation, these consoles may be useful:

- **Console 1: The Installation dialog box.** This is the main installation console. You see it when Setup first starts. After the GUI takes over, it's hidden in the background. You can call it up by pressing Ctrl+Alt+F1.
- **Console 2: Shell prompt.** This console provides you with a shell prompt, from which you can enter Linux commands. If you need to do something manually during installation, you can do it from this console. The keyboard shortcut is Ctrl+Alt+F2.
- **Console 3: Install log.** This console lists messages generated by the installation program. You can get to it by pressing Ctrl+Alt+F3.
- **Console 4: System log.** This console displays system-related messages. You can get to it by pressing Ctrl+Alt+F4.
- **Console 5: Other messages.** Still more messages may appear in this console, which you can open by pressing Ctrl+Alt+F5.
- **Console 6: X graphical display.** This is the console where the GUI of the installation program is displayed. If you use a Ctrl+Alt keyboard combination to view any of the other logs, press Ctrl+Alt+F7 to return to the installation GUI.

IN THIS CHAPTER

- Getting used to the Linux way of thinking
- Logging on and off Linux
- Using Linux commands in a console
- Configuring a Linux server by editing configuration files
- Administering a Linux server remotely with Cockpit

Chapter 2
Linux Administration

Before you can set up Linux to do serious networking, you need to discover the basics of getting around a Linux server. In this chapter, you learn those basics. You see how to log on and off Linux, how the Linux file system works, and how to use commands. I also introduce you to Vi, a text editor commonly used to edit Linux commands. And I introduce you to Cockpit, a tool for administering a Linux server remotely using a web interface. Finally, I show you the basics of setting up a Linux user account.

On Again, Off Again

Any user who accesses a Linux system, whether locally or over a network, must be authenticated by a valid user account on the system. In the following sections, you find out how to log on and off of a Linux system and how to shut down the system.

Logging in

When Linux boots up, it displays a series of startup messages as it starts the various services that compose a working Linux system. Then it displays a login prompt, as shown in Figure 2-1.

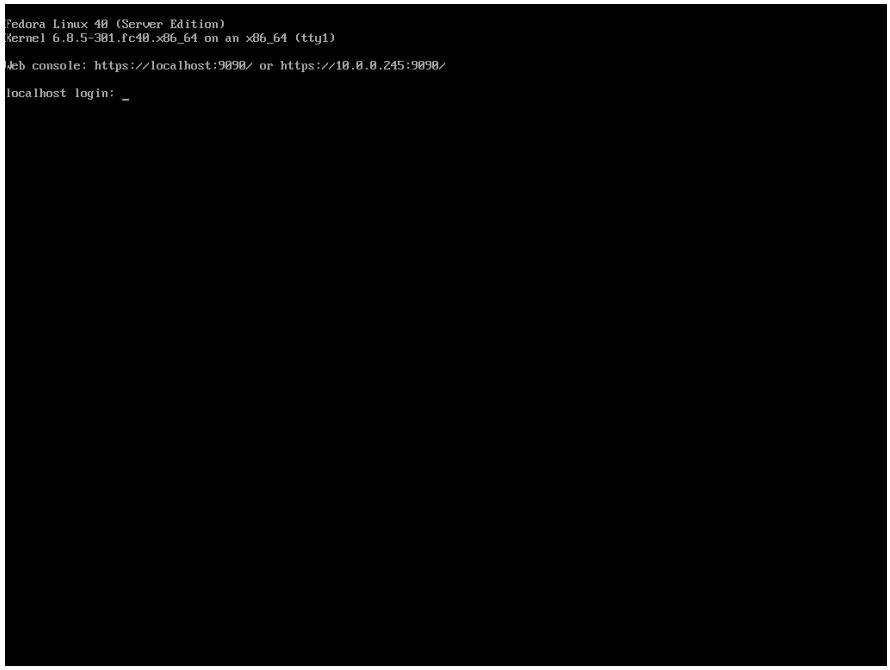


FIGURE 2-1: Begin by logging on.

The console window in Figure 2-1 displays the following text:

```
Fedora 40 (Server Edition)
Kernel 6.8.5-301.fc40.x86_64 on an x86_64 (tty1)

Web console: https://localhost:9090 or https://10.0.0.245:9090

Localhost login:
```

Before you actually log in, take a moment to review the information that's displayed in the console:

- The first line indicates you're running the server edition of Fedora version 40.
- The second line indicates the version of the Linux Kernel (6.8.5-301). The *kernel* is the core part of Linux that distributions such as Fedora version 40 are built upon.
- The third line gives you the URL of a web console you can use for many administration tasks. The web console is provided by a package called Cockpit, which you'll learn about that later in this chapter, in the section "Using Cockpit."
- The final line is the actual login prompt.

To log in to the Linux server, type your user name and press Enter. You'll be prompted for your password; type it and press Enter. You'll be greeted by a command prompt similar to this:

```
[dlowe@localhost ~]$
```


TIP

As a part of the installation process (described in Book 8, Chapter 1), you created a user account other than the `root` account. You should use this user account rather than the `root` user account whenever possible. Use the `root` account only when making major changes to the system's configuration. When you're doing routine work, log on as an ordinary user in order to avoid accidentally corrupting your system.

Logging out

After you've logged on, you'll probably want to know how to log out. You can do that by typing **logout** and pressing Enter. If you prefer, you can press the key combination Ctrl+D. Either way, you're returned to the login prompt shown in Figure 2-1.

Shutting down

As with any OS, you should never turn off the power to a Linux server without first properly shutting down the system. To shut down a Linux system, just type **shutdown** and press Enter.

By default, entering shutdown by itself will schedule a shutdown in one minute. You'll be informed of the exact time of the shutdown, like this:

```
[dlowe@localhost ~]$ shutdown
Shutdown scheduled for Sun 2024-05-19 10:51:16 PDT, use 'shutdown -c' to cancel
[dlowe@localhost ~]$
```

You'll have to wait a full minute before the system shuts down. (If you change your mind before the minute elapses, type **shutdown -c** and press Enter.)

To shut down immediately (instead of having to wait a minute), type **shutdown now** (instead of **shutdown**) and press Enter.

You can also schedule a shutdown for any time by entering the time you want the system to shut down, using the 24-hour clock format. For example, to shut down at 11:30 p.m., type **shutdown 23:30** and press Enter. Again, you can cancel the shutdown any time prior to the time you entered by typing **shutdown -c** and pressing Enter.

Wait, Where's the Desktop?

The first thing you'll notice upon logging in to a Linux server is that there's no desktop. In fact, Linux servers have no graphical user interface (GUI) at all. There are no windows, scroll bars, buttons, check boxes, or frankly anything that resembles Microsoft Windows. Just a bleak command prompt.

Why no desktop? Because a GUI consumes a lot of system resources (CPU and RAM) that could be better spent on true server functions such as sharing files, running web services or database services, or handling other mundane server tasks like Domain Host Configuration Protocol (DHCP) and Domain Name System (DNS). Without the overhead of a GUI, Linux can devote all its resources to server processes.

It's true, the lack of a GUI makes administering a Linux server more difficult. That's part of its charm. Sure, there's a steep learning curve you must climb to learn the nuances of the commands that you'll use to configure Linux services. But after you master the commands, you'll join the ranks of Linux administrators who giggle at Windows administrators because they don't know how to do anything but click buttons and check boxes.

Playing the Shell Game

Most Linux server administration is done from a *console,* also called a *terminal* or a *command line interface* (CLI), which is the Linux equivalent to a Windows command window. A console is a text-based screen that lets you interact with a program called a *shell,* which reads and executes commands.

In the early days of Unix (the 1970s), computer operators communicated with Unix via consoles. At first, these consoles were teletype machines that combined a keyboard with a printer. As the operator typed a command, the command was printed. Then the command was executed and the computer printed the results of the command on the console's printer.

Teletype consoles gave way to CRT consoles, which worked pretty much like teletype machines but didn't waste paper. The operator typed commands, which appeared on the CRT screen. The computer then executed the commands and displayed the results on the screen. Unix didn't really care whether the operator was using a teletype or a CRT console.

Today, a console is the combination of a keyboard and a monitor connected to a computer running Linux. If you're running Linux in a virtual machine, the console

is in a window, but Linux doesn't know that — it just thinks of the console as a physical monitor and a physical keyboard.

By itself, a console doesn't do anything except provide the hardware (physical or virtual) that connects a user to a Linux computer. Within the console, you interact with the shell, which interprets commands you enter, executes them, and manages the display of results. The prompt you see in a command window is generated by the shell.



A variety of shells are available in Linux, but the most common is called *Bash.* The name *Bash* has an interesting history. In the 1980s, the most popular shell on Unix computers was called the Bourne Shell, named for the person who developed it (Stephen Bourne, not Jason). The executable file for the Bourne shell was named `sh`. By the end of the '80s, Brian Fox developed a replacement for the Bourne shell. He named it Bash, which stands for *Bourne Again SHell*. Very clever indeed!

Bash is the default shell used on most Linux distributions, including Fedora Server. Bash is the shell I use throughout this minibook.

Getting into Virtual Consoles

Because Linux is a multiuser system, it lets you work with more than one console. In fact, you actually have six virtual consoles at your disposal. By default, you start in console 1. You can switch to a particular virtual console by pressing Ctrl+Alt+F1 through Ctrl+Alt+F6. For example, to switch to virtual console 3, press Ctrl+Alt+F3.

The console you're in is indicated in the Kernel line that's displayed as part of the login prompt in a Linux console. For example:

```
Kernel 6.8.5-301.fc40.x86_64 on an x86_64 (tty1)
```

Here, `tty5` means you're in console 5. (*Tty* is a common abbreviation for *teletype*.)

You can also find out which console you're in by using the `tty` command. For example:

```
[dlowe@localhost ~}$ tty
/dev/tty1
```

Here, the `tty` command indicates that you are in tty1, which is console 1.

Using a Remote Console

You can easily access the console of a Linux server remotely (that is, from another computer on the same network) by using a handy terminal emulator program called PuTTY. You can find PuTTY at many sources on the internet. If you're working on a Windows computer, the easiest (and safest) way to get PuTTY is to open the Windows Store, search for PuTTY, and then install it.

When you run PuTTY, the configuration screen shown in Figure 2-2 appears. Enter the IP address of your Linux server, and then click Open. PuTTY opens a console window that prompts you to log in, as shown in Figure 2-3.

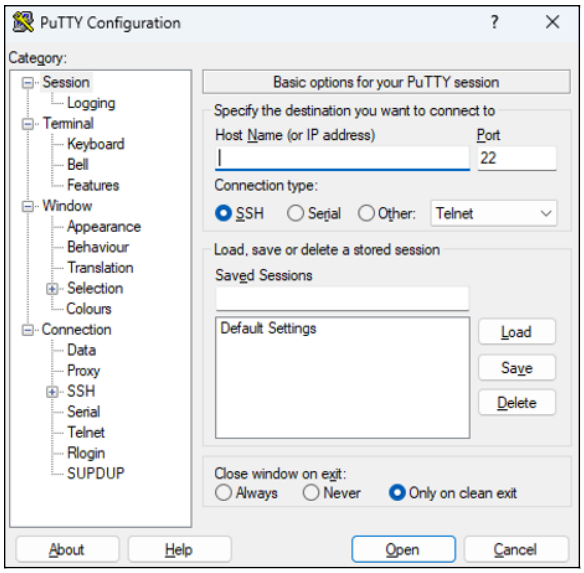


FIGURE 2-2: Set the IP address of your Linux server in PuTTY.

When you're logged in, you can run the `tty` command to see what console you're connected to, as follows:

```
[dlowe@localhost ~]$ tty
/dev/pts/0
```

Here, you can see that you aren't connected to one of the tty consoles, but rather to `/dev/pts/0`. The `pts` folder lists the serial ports that are available; your terminal session lives on port 0.

You can perform any Linux administrative task from a remote PuTTY session that you could perform on a local console.

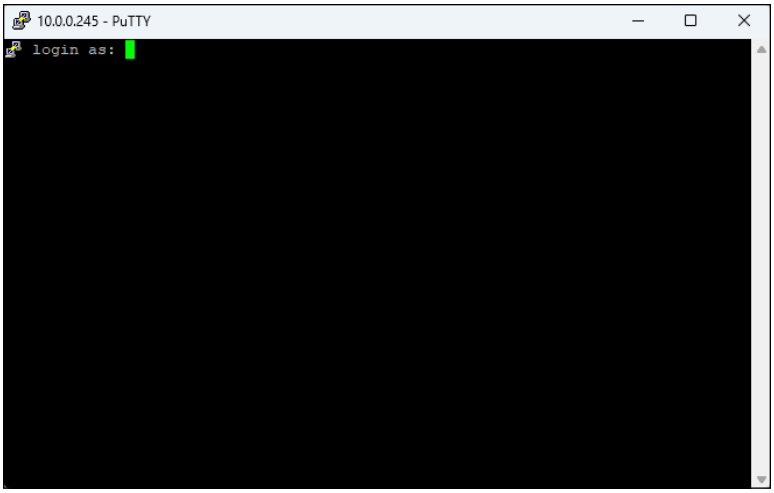


FIGURE 2-3: PuTTY prompts you to log in to your Linux server.

Enabling the root User

The `root` user is the ultimate administrator account in a Linux system. By default, the `root` account is disabled when you install Fedora Server, but you can enable it during the setup wizard. If you didn't enable `root` when you installed Linux, you can enable it at any time by resetting the `root` account's password.

At the console prompt, enter the following command:

```
sudo passwd root
```

Linux responds by prompting you for your own password. When you enter that, you're prompted to enter the new password for the `root` account twice. If the passwords match, you'll see this message:

```
passwd: all authentication tokens updated successfully.
```

This message is Linux's way of telling you the password has been reset. The root account is now enabled and can be used when necessary.



WARNING

Avoid logging in to Linux as the `root` user. Doing so grants total control of the Linux system to the session. And that can be dangerous. Instead, log in using a standard user account, and use the `sudo` command (described in the next section) to do things that require `root` privileges.

Using the sudo Command

The sudo command is required because many Linux administrative commands can only be run by the root user. You can simply log in as the root user to run such commands, but that's considered a risky practice because the root user can do virtually anything in a Linux environment. It's safer to log in with an ordinary user account and use sudo to enable access to administrative functions. As a result, sudo is an essential tool for Linux administration.

For example, you'll often use the dnf command to install new software on a Linux system. The dnf command is one of those commands that can only be run by the root user. So you'll need to use sudo to run the dnf command. To use sudo, you simply prefix the command you want to run with the word sudo, as in the following example:

```
sudo dnf install dhcp
```

Here, the command dnf install dhcp will be run as the root user. Note that for security purposes, the sudo command prompts you for your own password before it runs the dnf command.

Note that your account must be configured with permission to run the sudo command. In other words, not all Linux accounts have access to sudo. The account that you create when you run the setup wizard is automatically granted sudo rights.

Understanding the file system

The Linux file system is a bit different from the Windows file system. Two of the most obvious differences are actually superficial:

- Linux uses forward slashes rather than backward slashes to separate directories. Thus, /home/doug is a valid path in Linux; \Windows\System32 is a valid path in Windows.
- Linux filenames don't use extensions. You can use periods within a filename, but unlike Windows, the final period doesn't identify a file extension.

The fundamental difference between the Linux and Windows file system is that Linux treats everything in the entire system as a file, and it organizes everything into one gigantic tree that begins at a single root. When I say, "Everything is treated as a file," I mean that hardware devices such as floppy drives, serial ports, Ethernet adapters, even consoles are treated as files.

The root of the Linux file system is the root partition from which the OS boots. Additional partitions, including other devices that support file systems such as CD-ROM drives, floppy drives, or drives accessed over the network, can be grafted into the tree as directories called *mount points.* Thus, a directory in the Linux file system may actually be a separate hard drive.

Looking at top-level directories

Another important aspect of the Linux file system is that the directories that compose a Linux system are governed by a standard called the Filesystem Hierarchy Standard (FHS). This standard spells out which directories a Linux file system should have. Because most Linux systems conform to this standard, you can trust that key files will always be found in the same place.

Table 2-1 lists the top-level directories that are described in the FHS.

TABLE 2-1 **Top-Level Directories in a Linux File System**

Directory	Description
`/bin`	Essential command binaries
`/boot`	Static files of the boot loader
`/dev`	Devices
`/etc`	Configuration files for the local computer
`/home`	Home directories for users
`/lib`	Shared libraries and kernel modules
`/lib64`	Shared libraries with 64-bit code
`/media`	Media files
`/mnt`	Mount point for file systems mounted temporarily
`/opt`	Add-on applications and packages
`/proc`	Processes that are currently running on the server
`/root`	Home directory for the root user
`/sbin`	Essential system binaries
`/tmp`	Temporary files
`/usr`	Read-only, shared files such as binaries for commands and libraries that can be used by any user in the system
`/var`	Variable data files

There are two other directories you should know about:

- **The home directory:** Every user in Linux has a *home directory,* which is located at `/home/`. Thus, `/home/dlowe` is the home directory for user `dlowe`.

  Whenever a user logs in to Linux, the user's home folder is set as the current directory. The home directory is represented by the tilde character (`~`).

  
  TECHNICAL STUFF

  The one exception to this rule is the root user. This special user's home directory is a top-level directory named `root` rather than a directory beneath the home directory.

- **The root directory:** The *root directory* is the directory that's at the very top of the file system.

  
  TECHNICAL STUFF

  In Windows, each mounted drive has its own root directory, indicated by a backslash following the drive letter. For example, `C:\` is the root directory of the `C:` drive. In short, each drive in Windows has its own file system. In contrast, a Linux computer has a single file system.

  Try not to be confused by the many uses of the word *root* in Linux. In Linux, the ultimate administrator account is named `root`. The `root` user has a special home directory named `root`. The word *root* is also used to refer to the top directory in the Linux file system. This directory technically doesn't have a name — it's just represented by a single forward slash (`/`).

Browsing the file system

As you work with Linux, you'll often find yourself needing to browse the file system. To do that, you can use any of the following commands:

- `pwd`: Displays the name of the current directory.
- `ls`: Displays the contents of the current directory.
- `cd`: Navigates to a different directory.


TIP

You can always tell what directory you're in by looking at the console prompt.

Let's take these commands for a ride:

1. **Log in to Linux.**

   You'll see the console prompt:

   ```
   [dlowe@localhost ~]$
   ```

You can see in the command prompt (in the first line above) that the tilde character (~) is used to indicate that the user is currently in the home directory.

2. **Type** pwd **and press Enter.**

   Linux displays the path for the current directory:

```
[dlowe@localhost ~]$ pwd
/home/dlowe
[dlowe@localhost ~]$
```

   The output from the `pwd` command shows that the actual directory location is `/home/dlowe`.

3. **Type** cd / **and press Enter.**

   This command switches the current directory to the root directory:

```
[dlowe@localhost ~]$ cd /
[dlowe@localhost /]$
```

   Note that the current directory as indicated in the prompt has changed from ~ to /.

4. **Type** ls **and press Enter.**

   This command lists the contents of the current directory (in this case, the root directory):

```
[dlowe@localhost /]$ ls
afs  boot  etc   lib    media  opt   root  sbin  sys  usr
bin  dev   home  lib64  mnt    proc  run   srv   tmp  var
[dlowe@localhost /]$
```

   Notice that the directories listed by this command correspond to the top-level directories in Table 2-1 earlier in this chapter.

Using the RPM Package Manager

One of the basic tasks of configuring and administrating a Linux system is installing software. Workstation versions of Linux include a GUI tool to perform this task, but in server versions, you'll need to do this from the command line.

The tool you use to deploy software on a Linux system is called a *package manager.* A package manager is a collection of tools that are designed to install *packages*, which are files that contain all the elements needed to install a specific application

or feature. This includes (but isn't limited to) executable program files and documentation files. In addition, packages contain metadata that provides information about the files in the package, including the publisher of the package, version information, dependencies, and so on.

Over the years, two distinct package managers have become popular on Linux systems:

- **DEB:** The standard software deployment tool used on Debian variations of Linux.
- **RPM:** The standard used on Red Hat variations of Linux. Fedora is in the Red Hat camp, so I use RPM deployment tools in this book.

RPM has gone through three major evolutions as it continues to improve. In the original version, you used the `rpm` command to manage packages. This was eventually replaced by a more powerful version called `yum`. And recently, `yum` has been supplanted by `dnf`, which is the command I use in this book.

Note that all three commands are installed by default on Fedora Server, so you can use any of the three commands you want. But `dnf` is preferred because it's faster and it has additional features.

The `dnf` package manager deals with two kinds of packages: those that are installed on the server, and those that are not installed but are available via the *repository*, which is a vast library of packages you can download and install on your server.

Listing packages

Let's start by using `dnf` to see a list of all packages that have been installed on your server. Type **dnf list installed** and press Enter. You get something like this:

```
[dlowe@localhost ~]$ dnf list installed
Installed Packages
ModemManager.x86_64                1.22.0-3.fc40       @anaconda
ModemManager-glib.x86_64           1.22.0-3.fc40       @anaconda
NetworkManager.x86_64              1:1.46.0-2.fc40     @anaconda
NetworkManager-bluetooth.x86_64    1:1.46.0-2.fc40     @anaconda
NetworkManager-libnm.x86_64        1:1.46.0-2.fc40     @anaconda
.
.
.
zchunk-libs.x86_64                 1.4.0-2.fc40        @anaconda
zip.x86_64                         3.0-40.fc40         @anaconda
zlib-ng-compat.x86_64              2.1.6-2.fc40        @anaconda
```

```
zram-generator.x86_64              1.1.2-9.fc40        @anaconda
zram-generator-defaults.noarch     1.1.2-9.fc40        @anaconda
```

The three lines with dots represent a few hundred additional packages that I removed — otherwise, the output from this command would've filled a couple of pages here. For each package installed on the server, you can see the complete package name, the package version, and the source of the package.

If you want to see a list of all packages that are available via the repository, type **dnf list available** and press Enter. You get a much larger list that resembles this:

```
[dlowe@localhost ~]$ dnf list available
Last metadata expiration check: 0:05:29 ago on Sat 15 Jun 2024 10:40:10 AM PDT.
Available Packages
0ad.x86_64                         0.0.26-21.fc40      fedora
0ad-data.noarch                    0.0.26-7.fc40       fedora
0xFFFF.x86_64                      0.10-7.fc40         fedora
2048-cli.x86_64                    0.9.1-20.fc4        fedora
2048-cli-nocurses.x86_64           0.9.1-20.fc40       fedora
.
.
.
zziplib-devel.i686                 0.13.69-7.fc32      fedora
zziplib-devel.x86_64               0.13.69-7.fc32      fedora
zziplib-utils.x86_64               0.13.69-7.fc32      fedora
zzuf.x86_64                        0.15-12.fc32        fedora
```

This command lists literally thousands of packages you can install.

Fortunately, `dnf` lets you narrow down your package searches by using pattern searches. For example, suppose you want to install an FTP package, but you're not sure which FTP packages are available. To do that, you can search for packages using the wildcard pattern `*ftp*`.

First, let's make sure we don't already have an FTP package installed:

```
[dlowe@localhost ~]$ dnf list installed *ftp*
Error: No matching Packages to list
```

Seeing that no installed package matches the wildcard pattern, we can search for available FTP packages:

```
[dlowe@localhost ~]$ dnf list available *ftp*
Last metadata expiration check: 0:09:05 ago on Sat 15 Jun 2024 10:40:10 AM PDT.
Available Packages
compat-golang-goftp-server-2-devel.noarch   2.0.1-3.fc40      fedora
```

```
curlftpfs.x86_64                          0.9.2-37.fc40     fedora
edg-gridftp-client.x86_64                 1.2.9.2-28.fc40   fedora
erlang-ftp.x86_64                         26.2.5-1.fc40     updates
erlang-tftp.x86_64                        26.2.5-1.fc40     updates.
.
.
.
tnftp.x86_64                              20230507-3.fc40   fedora
uberftp.x86_64                            2.9.1-4.fc40      fedora
vsftpd.x86_64                             3.0.5-6.fc40      fedora
```

This command lists more than 100 packages to choose from. But by searching the web for "Fedora FTP," you'll find that the most commonly used packages are `proftpd` and `vsftpd`. And according to Fedora, the preferred FTP server is `vsftpd`.

Installing packages

To install a package, you use the `dnf install` command. For example, to install the `vsftpd` package, type **sudo dnf install vsftpd.x86_64** and press Enter.

This command uses the complete package name I discovered when I listed the available FTP packages. Note that you'll need to run `dnf` under `sudo` because installing packages requires `root` privileges.

As always when you use `sudo`, you'll be prompted for your password. And during the installation, `dnf` will ask you to confirm that you want to proceed; type **Y** and press Enter to continue.

Here's the complete output from the above `dnf` command:

```
[dlowe@localhost /]$ sudo dnf install vsftpd.x86_64
[sudo] password for dlowe:
Last metadata expiration check: 0:10:56 ago on Sat 15 Jun 2024 10:42:12 AM PDT.
Dependencies resolved.
================================================================
 Package       Architecture      Version        Repository      Size
================================================================
Installing:
 vsftpd        x86_64            3.0.5-6.fc40   fedora         168 k

Transaction Summary
================================================================
Install  1 Package

Total download size: 168 k
```

```
Installed size: 344 k
Is this ok [y/N]: y
Downloading Packages:
vsftpd-3.0.5-6.fc40.x86_64.rpm          206 kB/s | 168 kB     00:00
--------------------------------------------------------------------------
Total                                   124 kB/s | 168 kB     00:01
Running transaction check
Transaction check succeeded.
Running transaction test
Transaction test succeeded.
Running transaction
  Preparing        :                                                1/1
  Installing       : vsftpd-3.0.5-6.fc40.x86_64                     1/1
  Running scriptlet: vsftpd-3.0.5-6.fc40.x86_64                     1/1

Installed:
  vsftpd-3.0.5-6.fc40.x86_64

Complete!
```

The `Complete!` message indicates that the package was successfully installed.

Removing packages

If you want to remove a package from your server, use the `dnf remove` command. For example:

```
sudo dnf remove vsftpd
```

You'll be prompted by `sudo` for your password, and you'll be asked to confirm that you want to delete the package. Here's the complete output from this command:

```
[dlowe@localhost /]$ sudo dnf remove vsftpd
[sudo] password for dlowe:
Dependencies resolved.
================================================================
Package       Architecture      Version        Repository      Size
================================================================
Removing:
 vsftpd       x86_64            3.0.5-6.fc40   @fedora        344 k

Transaction Summary
================================================================
Remove  1 Package

Freed space: 344 k
```

```
Is this ok [y/N]: y
Running transaction check
Transaction check succeeded.
Running transaction test
Transaction test succeeded.
Running transaction
  Preparing        :                                                1/1
  Running scriptlet: vsftpd-3.0.5-6.fc40.x86_64                     1/1
  Erasing          : vsftpd-3.0.5-6.fc40.x86_64                     1/1
  Running scriptlet: vsftpd-3.0.5-6.fc40.x86_64                     1/1

Removed:
  vsftpd-3.0.5-6.fc40.x86_64

Complete!
```

Again, the `Complete!` message indicates that the package was successfully removed.

Updating packages

An important task for any administrator is to keep your software up to date. The `dnf` command lets you do that via the `dnf check-update` and `dnf update` commands.

The `dnf check-update` command lets you know what updates need to be applied:

```
  [dlowe@localhost ~]$ sudo dnf check-update
Last metadata expiration check: 0:20:41 ago on Sat 15 Jun 2024
    10:42:12 AM PDT.
PackageKit.x86_64                    1.2.8-5.fc40         updates
PackageKit-glib.x86_64               1.2.8-5.fc40         updates
alternatives.x86_64                  1.27-1.fc40          updates
.
.
.
```

When you first install a Linux server, you should run `dnf update` to ensure that all the server software is up to date. You'll see a list of packages that need updating, followed by a summary of what needs to be done and a prompt to confirm that you want to proceed:

```
[dlowe@localhost /]$ sudo dnf update
Last metadata expiration check: 0:22:45 ago on Sat 15 Jun 2024 10:42:12 AM PDT.
Dependencies resolved.
```

```
================================================================================
 Package          Architecture   Version            Repository     Size
================================================================================
Installing:
 kernel           x86_64         6.8.11-300.fc40    updates       160 k
Upgrading:
 PackageKit       x86_64         1.2.8-5.fc40       updates       653 k
 PackageKit-glib  x86_64         1.2.8-5.fc40       updates       154 k
 alternatives     x86_64         1.27-1.fc40        updates        41 k
.
.
.
 zlib-ng-compat   x86_64         2.1.6-5.fc40       updates        77 k
 zram-generator   x86_64         1.1.2-11.fc40      updates       442 k
Installing dependencies:
 kernel-core      x86_64         6.8.11-300.fc40    updates        17 M
 kernel-modules   x86_64         6.8.11-300.fc40    updates        63 M
Installing weak dependencies:
 reportd          x86_64         0.7.4-13.fc40      fedora         47 k
 sqlite           x86_64         3.45.1-2.fc40      fedora        856 k
 tpm2-tools       x86_64         5.7-1.fc40         updates       810 k

Transaction Summary
================================================================================
Install    9 Packages
Upgrade  213 Packages

Total download size: 465 M
Is this ok [y/N]:
```

Type **y** and press Enter to update everything. Then grab a cup of coffee and a good book to read, because the updates will take a while.

Editing Text Files with Vi

Most Linux configuration settings are stored in configuration files located in the `/etc` folder. To properly configure a Linux system, you need to know how to edit those files. And to do that, you need to know how to use text-editing software.

If you were running a desktop version of Linux, you could open a configuration file in a GUI-based text editor, make your configuration changes, save the file, and reboot the system so the changes will take effect. But editing a file in a text-only Linux console is more difficult: You'll need to use a text-based editor rather

than a GUI-based editor. And text-based editors are notoriously difficult to learn because they rely heavily on arcane commands.

The most popular text on Linux, called Vi, is no different. You'll need to invest some time to learn how to use this powerful editor, and it will be frustrating at first because it isn't very intuitive. But stick with it. Learning Vi is kind of an initiation rite for Linux administration, so after you've mastered it, you'll have bragging rights.



TIP

As tempting as it is to pronounce `vi` with a single syllable ("vee"), the proper way to pronounce it is to actually say the letters separately ("vee eye").

The following sections review the basics of using Vi to edit configuration files.

Starting vi

To create a new text file with Vi, enter a command similar to the following:

```
[dlowe@localhost ~]$ vi newfile
```

This opens `vi`, which fills all available lines of the console, as shown in Figure 2-4.



FIGURE 2-4: Vi creating a new file.

On the bottom line, you can see a status message that indicates you're editing a file called `newfile` and that this is a new file.

The rest of the screen is filled with tilde characters (~). The tildes are simply placeholders that indicate that these lines don't actually exist in the file. As you add lines of text to the file, the tildes will be replaced by the text you create.

If the filename you provide on the `vi` command represents a file that already exists, Vi will open the file and display its content. For example, suppose you enter this command to open a configuration file named `group` that resides in the `/etc` directory:

```
[dlowe@localhost ~]$ vi /etc/group
```

Then, Vi opens the group file and displays its contents, as shown in Figure 2-5.

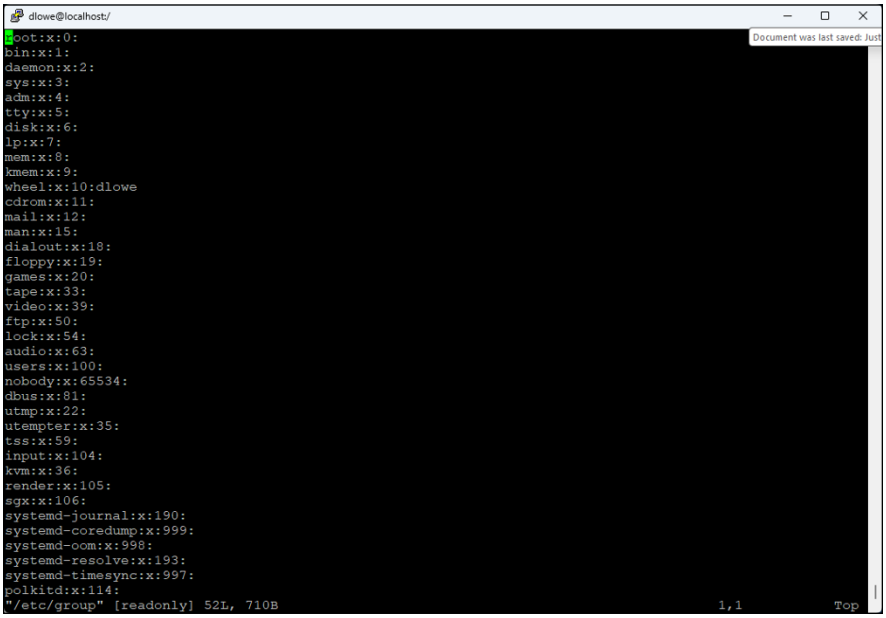


FIGURE 2-5: Vi editing an existing file.

In the status message on the bottom line, you can see that the file is named `group`, but you can also see that the file is opened in read-only mode. That's because all the files in the `/etc` directory are restricted; you need root privileges to change them.

sudo to the rescue! When you edit a configuration file, you need to prefix the vi command with sudo, like this:

```
[dlowe@localhost ~]$ sudo vi /etc/group
```

This removes the read-only limitation so you can make changes to the file.



The files in the /etc directory require root privileges for a reason! These files are vital to the proper operation of Linux. Tread lightly when you edit them.



You can open a file in read-only mode by using the command view instead of the vi command. For example:

```
[dlowe@localhost ~]$ view testfile
```

Here, a file named testfile is opened in read-only mode.

Saving changes and quitting Vi

Before we get into the details of editing files with Vi, I want to be sure you understand how to save changes and exit out of Vi. Changes you make to a file are not automatically saved back to the file, so you need to explicitly tell Vi when you want it to save your changes. You can wait until the end of your editing session to do so, but that's a bad idea. It's better to save your changes periodically as you edit a file so that you don't lose your work if something goes haywire.

On the other hand, it's important to know how to get out of vi without saving changes, just in case you've messed up the file. This is especially important when you're just getting started with vi and the editing commands seem confusing.

Table 2-2 lists the four commands you can use to save changes and exit out of Vi. Note that all four of these commands start with a colon (:), unlike the rest of the vi commands presented in this chapter.

Understanding Vi's operating modes

There are two basic operating modes in Vi: command mode and input mode.

The Vi editor doesn't have menus from which you can select commands. Nor does it recognize the mouse. Instead, all of Vi's editing and other functions are invoked via commands you enter while the editor is in command mode. The Vi editor starts up in *command mode,* which means that if you just start typing, you won't be entering text into the file. Instead, you'll be entering commands.

TABLE 2-2 Commands to Save and Exit

Command	Description
`:w`	Saves changes
`:w filename`	Saves changes to a new file
`:wq`	Saves changes and quit
`:q`	Quits, when no changes have been made since the file was last saved
`:q!`	Quits without saving changes

Input mode lets you type text that becomes a part of the file. To switch to input mode, you must use one of several commands that switch the editor to input mode. For example, the `a` command switches `vi` to input mode so you can type text at the current cursor location.

To switch from input mode back to command mode, just press Escape.

Moving around in a file

The easiest way to move the cursor around in Vi is to use the arrow keys. You can also use the End key to move to the end of the current line or the Home key to move to the start of the current line.

The Vi editor has a handful of other commands you can use to move around (see Table 2-3). Note that the first few commands in this table are single-letter equivalents to the arrow keys, the Home key and the End key.

TABLE 2-3 Commands for Moving the Cursor

Command	Description
`h` or Left Arrow	Moves the cursor one character to the left
`l` or Right Arrow	Moves the cursor one character to the right
`j` or Down Arrow	Moves the cursor one line down
`k` or Up Arrow	Moves the cursor one line up
`0` or Home	Moves the cursor to the start of the current line
`$`	Moves the cursor to the end of the current line

(continued)

TABLE 4-3 *(continued)*

Command	Description
w	Moves the cursor forward one word
W	Moves the cursor forward one word, ignoring punctuation within words
b	Moves the cursor backward one word
B	Moves the cursor backward one word, ignoring punctuation within words
e	Moves the cursor to the end of the current word
E	Moves the cursor to the end of the current word, ignoring punctuation within the word
+ or Enter key	Moves the cursor to the first non-blank character in the next line
–	Moves the cursor to the first non-blank character in the previous line
H	Moves the cursor to the first character in the first line on the screen
M	Moves the cursor to the first character in the middle line on the screen
L	Moves the cursor to the first character in the last line on the screen
1G	Moves the cursor to the first character in the first line of the file
G	Moves the cursor to the first character in the last line of the file
*n*G	Moves the cursor to the first character of line *n*

Inserting text

Table 2-4 lists the Vi commands for inserting text into a file. When you use these commands, Vi is switched to input mode, where you can type the text to be inserted. When you're finished, press the Escape key to return to command mode.

Note that you can insert text at the cursor location by simply pressing the Insert key on your keyboard. This is equivalent to using the `i` command.

TABLE 2-4

Commands for Inserting Text

Command	Description
`i` or Insert key	Inserts text to the left of the cursor
`I`	Inserts text before the first non-blank character on the current line
`a`	Inserts text to the right of the cursor
`A`	Inserts text at the end of the current line
`o`	Inserts a new line following the current line
`O`	Inserts a new line before the following line

Deleting text

Deleting text one character at a time is easy: Just move the cursor to the character you want to delete and press the Delete key.

You can also use any of the commands listed in Table 2-5 to delete text.

TABLE 2-5

Commands for Deleting Text

Command	Description
`x` or Delete key	Deletes the character at the cursor location
`dw`	Deletes from the cursor position to the end of the word
`db`	Deletes from the cursor position to the start of the word
`d0`	Deletes from the cursor position to the start of the line
`D`	Deletes from the cursor position to the end of the line
`dd`	Deletes all of the current line
`dG`	Deletes the current line and all subsequent lines to the end of the file
`dH`	Deletes the current line and all previous lines to the start of the file

Changing text

Changing existing text is one of the trickier parts of using Vi. One way to do it is to use any of the commands in Table 2-5 to delete the text you want to change, and then press the Insert key to flip Vi to insert mode and start typing the new text.

The commands in Table 2-6 are shortcuts that combine both operations into a single command. For example, if you use the command `C`, all characters from the cursor position to the end of the line are deleted; then, Vi switches to input mode so you can enter the text you want to replace the text you deleted.

TABLE 2-6 Commands for Changing Text

Command	Description
`c`	Deletes the character at the cursor location and enters insert mode
`cw`	Deletes from the cursor position to the end of the word and enters insert mode
`cb`	Deletes from the cursor position to the start of the word and enters insert mode
`c0`	Deletes from the cursor position to the start of the line and enters insert mode
`C`	Deletes from the cursor position to the end of the line and enters insert mode
`cc`	Deletes all of the current line and enters insert mode
`cG`	Deletes the current line and all subsequent lines to the end of the file and enters insert mode
`cH`	Deletes the current line and all previous lines to the start of the file and enters insert mode

Copying and pasting text

In Vi, the common operations of copying and pasting text are called *yank* and *put*, respectively. And what we would normally call the clipboard is called the *buffer*. (If those terms seem quaint, remember that they were invented in 1976, long before modern Windows PCs.)

Whenever you delete text using any of the commands in Table 2-5 or Table 2-6, the text that is deleted is placed in the buffer. This is the equivalent of what we call *cutting* today.

You can also yank text — that is, copy it to the buffer — by using any of the commands in Table 2-7.

TABLE 2-7 Commands for Yanking (Copying) and Putting (Pasting)

Command	Description
y1	Yanks the character at the cursor location
yw	Yanks from the cursor position to the end of the word
yb	Yanks from the cursor position to the start of the word
y0	Yanks from the cursor position to the start of the line
y$	Yanks from the cursor position to the end of the line
Y or yy	Yanks all of the current line
yG	Yanks the current line and all subsequent lines to the end of the file
yH	Yanks the current line and all previous lines to the start of the file
P	Pastes the copied text below the cursor
p	Pastes the copied text above the cursor

Repeating commands

There are two ways you can repeat a command in Vi. The first is to press the period (.) key. This repeats the command most recently executed. You can press the period key as many times as you want to keep repeating the previous command.

Another way to repeat a command is to prefix any of the commands presented in Table 2-4 or Table 2-5 with a number that indicates how many times the command should be repeated. Here are some examples and their effects:

Command	Effect
8 Down Arrow (or 8j)	Moves down five lines
3dd	Deletes three complete lines
12cc	Deletes 12 complete lines and enters insert mode
5yy	Yanks (copies) five lines
2p	Puts (pastes) the contents of the buffer twice

Other useful Vi commands

So far, I've introduced you to most of the basic editing commands available in Vi, but there are many more commands at your disposal. Table 2-8 presents some of the more useful `vi` commands that aren't covered in detail here, but can be useful.

TABLE 2-8 Other Commands

Directory	Description
`u`	Undoes the most recent change.
`/text`	Searches for the specified text in a forward direction.
`?text`	Searches for the specified text in a reverse direction.
`n` or `N`	Finds the next occurrence of the search text.
`:g/search/s//replace/g`	Finds every occurrence of the *search* text and replaces it with the *replace* text.
`:g/search/s//replace/gc`	Finds every occurrence of the *search* text and asks whether to replace it with the *replace* text. Reply `y` to replace and `n` to skip.
`:set ic`	Ignores case when searching.
`:set noic`	Does not ignore case when searching.
`:set nu`	Shows line numbers.
`:set nonu`	Hides line numbers.

Using Cockpit

Cockpit is a web-based administration tool that dramatically simplifies the job of managing a Linux server. Because a Fedora Server does not have a GUI such as Gnome, it does not have a web browser that can access Cockpit. But, fortunately, you can using Cockpit to manage the server remotely (that is, from another computer — even a Windows computer) via that computer's browser.

Cockpit is installed and enabled by default in Fedora Server 32. You can confirm that by the message that's displayed when you first log in to Fedora:

```
login as: dlowe
dlowe@10.0.0.244's password:
Web console: https://localhost:9090/ or https://10.0.0.244:9090/

Last login: Sun Oct 4 12:30:21 2020 from 10.0.0.84
[dlowe@localhost ~]$
```

Here, you can see that the web console (that's Cockpit) is available at `https://10.0.0.244:9090`.

If you don't see this message when you log in, you'll need to install and enable Cockpit manually. Fortunately, the procedure is simple:

1. **Install Cockpit with the following command:**

   ```
   sudo dnf install cockpit
   ```

2. **Enable Cockpit with the following command:**

   ```
   sudo systemctl enable --now cockpit.socket
   ```

   Cockpit should now be ready to use.

To access Cockpit, open a web browser on another computer in the same network and browse to the IP URL listed in the login message. You may get the standard security warning indicating that the connection is not private. If so, click Advanced, and then click Proceed. Cockpit will fire up and display a login page, as shown in Figure 2-6.

You can now log in using your Linux username and password. Cockpit then presents its overview page, as shown in Figure 2-7. This page provides an overview of the status of your computer.

The navigation pane on the left side of Cockpit lets you access several administrative centers:

- **Overview:** The overview page.
- **Logs:** Review events recorded in the Linux logs.

- **Storage:** Review and manage storage devices attached to your server.
- **Networking:** Review and manage network devices on your server. (For more information, refer to Book 8, Chapter 3.)
- **Accounts:** Review and manage user accounts on the server. (For more information, see the next section in this chapter, "Managing User Accounts.")
- **Services:** Review and manage Linux services.
- **Applications:** Review and manage third-party additions to Cockpit.
- **Software Updates:** Update packages installed on your server. (In Figure 2-8, you can see that updates are available. You can open the Software Updates page to download and install updates.)
- **Terminal:** Opens a terminal window through which you can enter Linux commands, as shown in Figure 2-8.

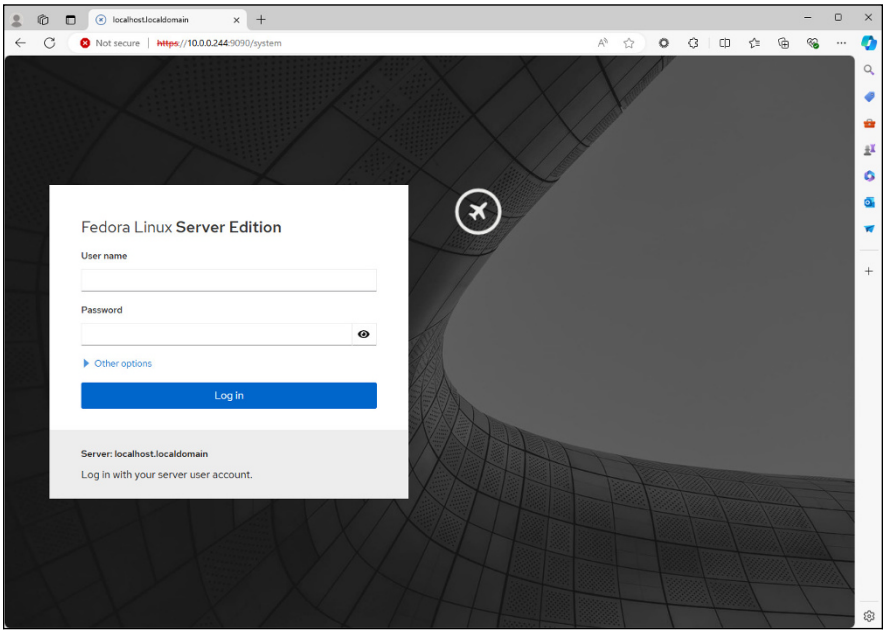


FIGURE 2-6: Logging in to Cockpit.

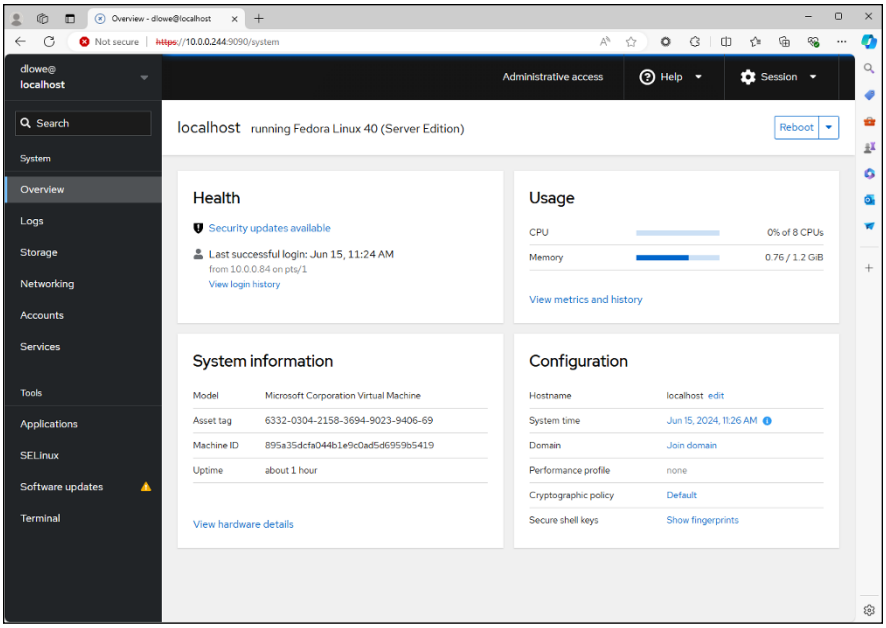


FIGURE 2-7: Cockpit's home page.

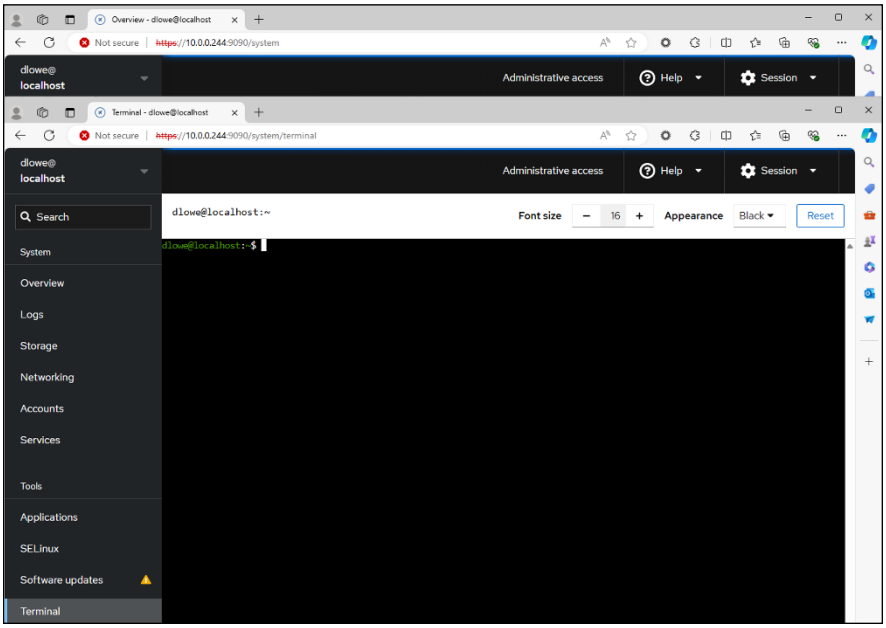


FIGURE 2-8: Using the Cockpit terminal.

Managing User Accounts

One of the most common network administration tasks is adding a user account. The Setup Agent prompts you to create a user account the first time you start Linux after installing it. However, you'll probably need to create additional accounts.

Each Linux user account has the following information associated with it:

- **Username:** The name the user types to log on to the Linux system.
- **Full name:** The user's full name.
- **Home directory:** The directory that the user is placed in when they log on. In Fedora, the default home directory is `/home/username`. For example, if the username is `dlowe`, the home directory is `/home/dlowe`.
- **Shell:** The program used to process Linux commands. Several shell programs are available. In most distributions, the default shell is `/bin/bash`.
- **Group:** You can create group accounts, which make it easy to apply identical access rights to groups of users.
- **User ID:** The internal identifier for the user.

You can add a new user by using the `useradd` command. For example, to create a user account named `kgearhart`, using default values for the other account information, open a terminal window or switch to a virtual console and type this command:

```
sudo useradd kgearhart
```

The `useradd` command has many optional parameters that you can use to set account information, such as the user's home directory and shell. In most cases, however, the defaults are adequate.

Note that before the account can be used, you must set the account password. To do that for the new `kgearhart` user, enter the following command:

```
sudo passwd kgearhart
```

You'll be prompted to enter your password, and then you'll be prompted twice to enter the user's new password.

If you prefer, you can use Cockpit to manage user accounts. Figure 2-9 shows the Accounts page in Cockpit. Here, you can see that the server has two accounts already: `dlowe` and `root`.

You can click an account to see the account details, as shown in Figure 2-10. Here, you can change the user's full name, lock the account, reset the password, or force the user to reset the password on their next login.

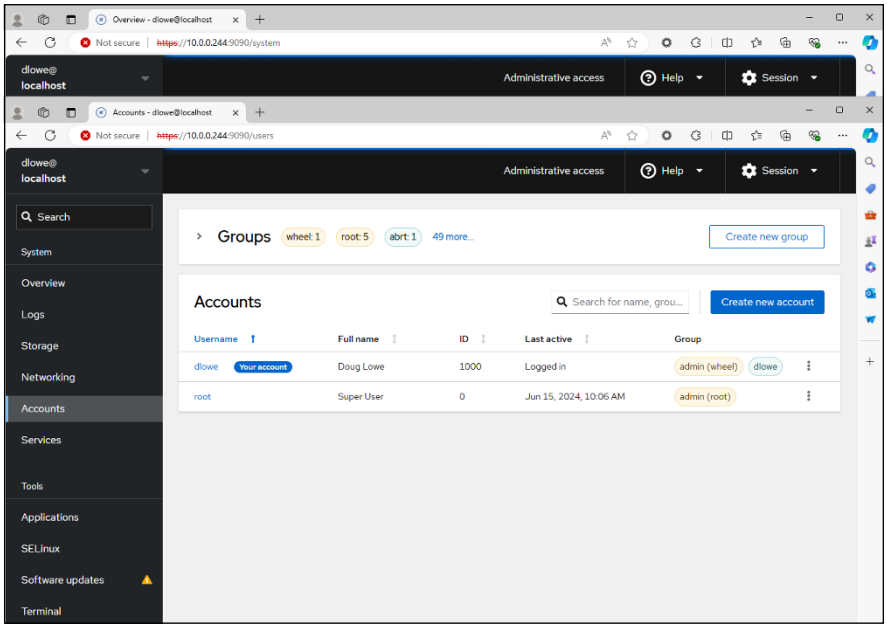


FIGURE 2-9: The Accounts page in Cockpit.

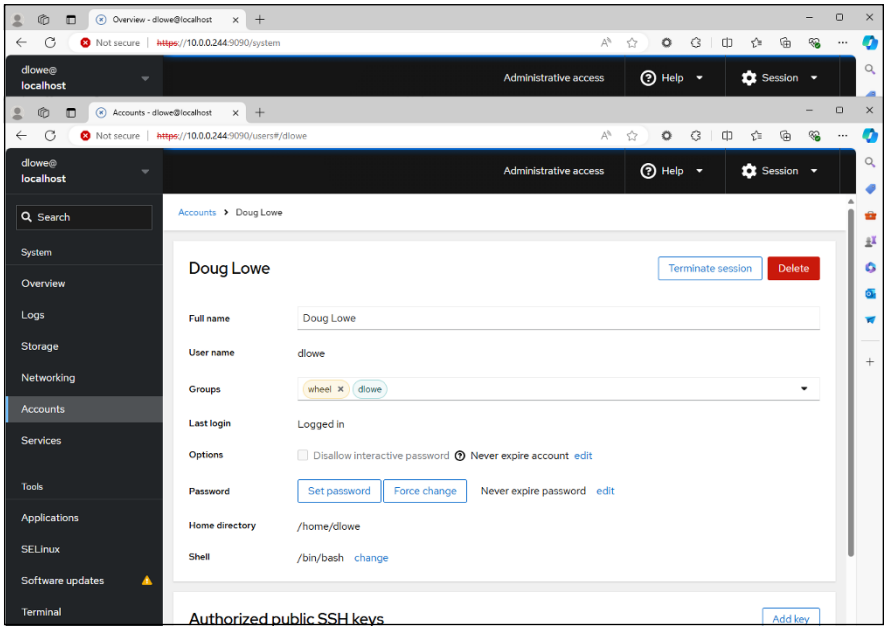


FIGURE 2-10: Displaying account details in Cockpit.

You can also create a new user account by clicking Create New Account on the Accounts page. This brings up the dialog box shown in Figure 2-11, where you can enter the information for the new user. Click Create to create the new user.

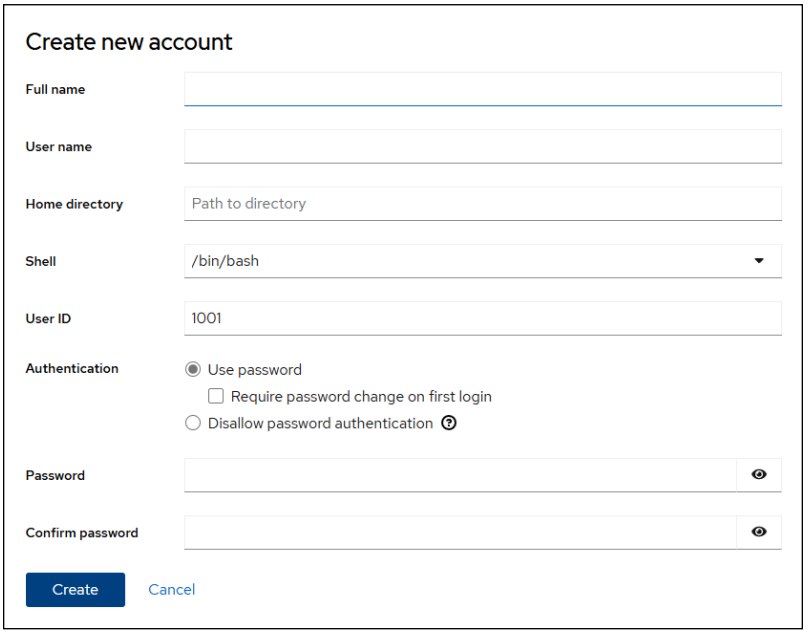


FIGURE 2-11: Displaying account details in Cockpit.

IN THIS CHAPTER

- Configuring network interfaces with Cockpit
- Looking directly at network configuration files
- Using the `ifconfig` command to display network status

Chapter 3

Basic Linux Network Configuration

In many cases, configuring a Linux server for networking is a snap. When you install Linux, the installation program automatically detects your network adapters and installs the appropriate drivers. Then, you're prompted for basic network configuration information, such as the computer's IP address, host name, and so on.

However, you may need to manually change your network settings after installation. You may also need to configure advanced networking features that aren't configured during installation. In this chapter, you discover the basic procedures for configuring Linux networking services.

Using Cockpit to Configure Network Interfaces

Before you can use a network interface to access a network, you have to configure the interface's basic TCP/IP options, such as its IP address, host name, Domain Name System (DNS) servers, and so on. By default, Linux server is configured to

use Dynamic Host Configuration Protocol (DHCP), which means that your server has a dynamic and unpredictable IP address. Most servers should have a static IP address, so unless you explicitly specified a static IP when you installed Linux, you'll need to do it now.

In this section, I show you how to configure your server for a static IP address remotely using Cockpit. (If you're not familiar with Cockpit, refer to Book 8, Chapter 2.)

Before you begin, you'll need to decide what the static IP should be for your server. For this example, I'll set the server's IP address to 10.0.0.30. Then follow these steps:

1. **Open Cockpit by browsing to the address shown in the web console line after you log in to the server.**

   The login page appears.

2. **Log in to Cockpit.**

   Use the username and password for your Linux account. When the login completes, the Cockpit Overview page appears.

3. **Click Networking in the navigation bar on the left.**

   This brings up the Network page, as shown in Figure 3-1.

   Near the center of the page, the Interfaces section lists all the network adapters on the server. In this example, there is just one, named `eth0`.

4. **Click the `eth0` interface.**

   This brings up the details page for the `eth0` interface, as shown in Figure 3-2. Here, you can see that the interface's IPv4 configuration is set to Automatic.

5. **Click the Edit link for the IPv$ setting.**

   This brings up the IPv4 Settings dialog box, shown in Figure 3-3.

6. **Use the drop-down list to switch from Automatic to Manual.**

   The IPv4 Settings dialog box assumes a different appearance, as shown in Figure 3-4.

7. **Enter the static IP, Netmask (the subnet mask), and Gateway.**

   For our example, we'll enter the following:

   - Address: 10.0.0.30
   - Netmask: 255.255.255.0
   - Gateway: 10.0.0.1

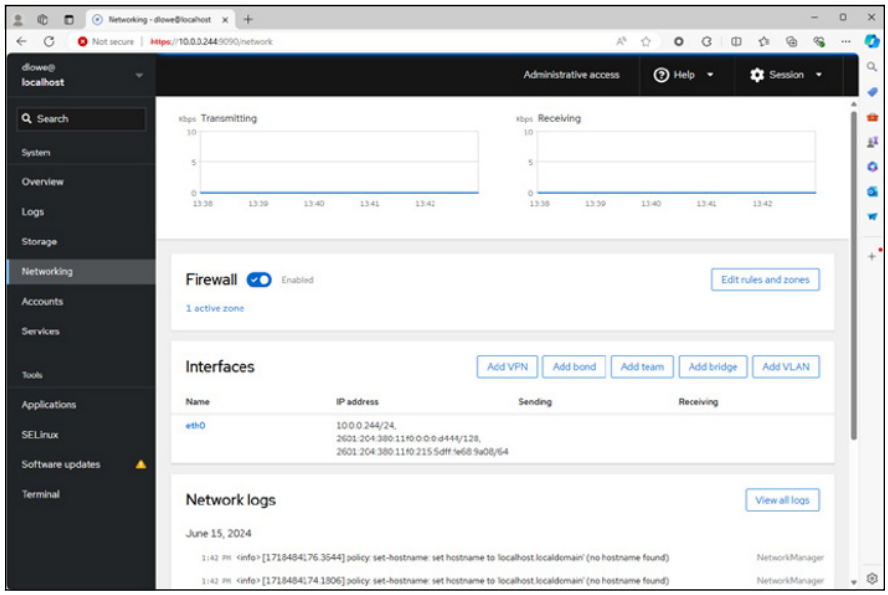


FIGURE 3-1: The Cockpit Network management page.

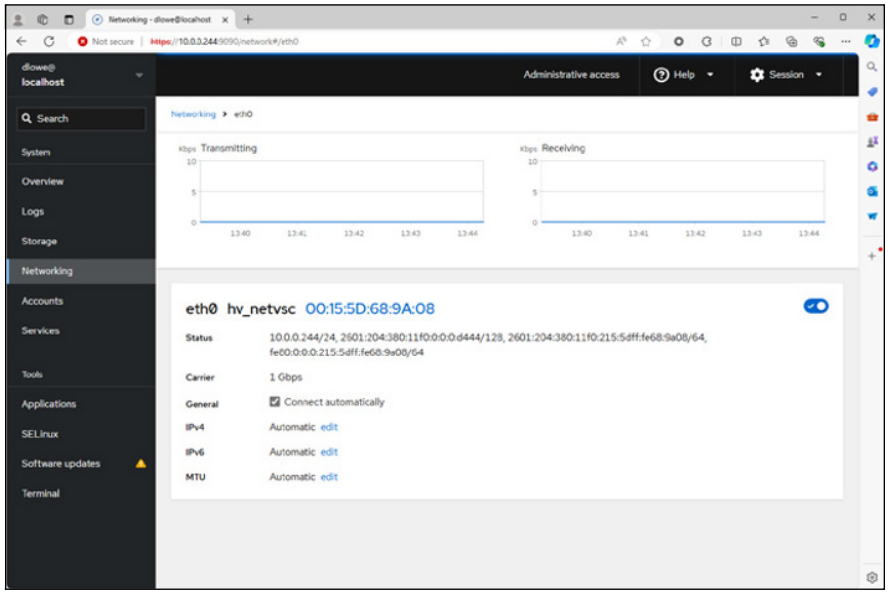


FIGURE 3-2: Viewing the details for a network interface.

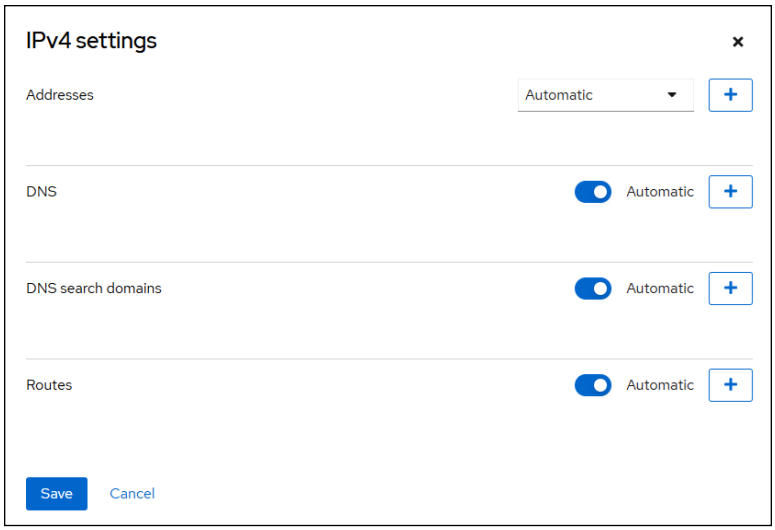


FIGURE 3-3: Configuring IPv4 settings.

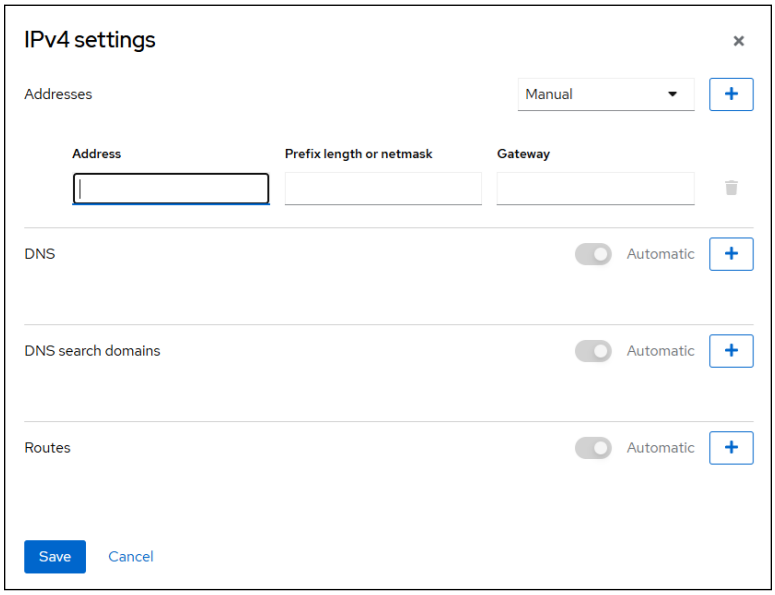


FIGURE 3-4: Setting a manual IP address.

If you're setting up this computer to be the gateway router that will manage traffic between your local network and the internet, use a static address that's easy to remember.

The subnet mask should be the mask that's appropriate for the IP address you choose. For a 10.0.0.*x* address, use 255.255.255.0.

The default gateway address should be the address of the gateway router that links your network to the internet. If this computer *is* the gateway router, specify the gateway address provided to you by your internet service provider (ISP).

8. **Add DNS servers.**

   Click the plus sign (+) in the DNS section once for each DNS server you want to add. Then fill in the name or IP address of the DNS servers.

   For our example, we'll use two DNS servers with the following addresses:

   - 75.75.75.75
   - 75.75.75.76


TIP

   If your network runs its own DNS server, you can specify its address here. Otherwise, you have to get the DNS server addresses from your ISP. If you're not running your own DNS server, you can use Google's public DNS server at 8.8.8.8.

   To enter more than one DNS server address, enter the first DNS address; then click the plus sign (+) button located beneath and to the right of the DNS server address to display a text box for the second DNS server address.

9. **Enter the name of your local domain (if you have one) in the DNS Search Domains text box.**

   For this example, we'll use `lowewriter.com`.

10. **Click Apply.**

    Cockpit tests the connection and then applies the configuration changes. You'll receive the dire warning shown in Figure 3-5, stating that if you proceed you won't be able to access the administration UI. Truthfully, Cockpit is overreacting just a bit: All you have to do is reconnect to Cockpit using the new IP address.

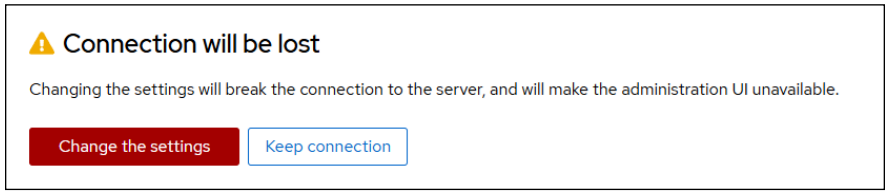


FIGURE 3-5: A dire warning from Cockpit!

Working with Network Configuration Files

Like other OS services, the Linux network is configured by settings that are specified in configuration files that you can find in the `/etc` directory or in one of its subdirectories. Graphical configuration programs, such as Fedora Network Configuration, are actually little more than glorified text editors that enable you to select network configuration options from user-friendly screens and then save your configuration changes to the standard configuration files. If you prefer to do the grunt work yourself, you can open the configuration files in a text editor and make changes to them directly.



WARNING

Any time you open a configuration file in `vi`, you run the risk of messing up your system's configuration. So be careful!

Table 3-1 lists the main Linux network configuration files and describes what each file does. The details of these files are described in the sections that follow.

TABLE 3-1 Linux Network Configuration Files

File	Location	Description
network	`/etc/sysconfig`	Global network settings
hostname	`/etc`	Specifies the host name (obsolete, but should still be present)
ifcfg-*xxxx*	`–/etc/sysconfig/network–scripts`	IP settings for the network adapter named *xxxx*
hosts	`/etc`	Lists host address mappings
resolv.conf	`/etc`	Lists DNS nameservers
nsswitch.conf	`/etc`	Specifies the name search order
xinetd.conf	`/etc`	Specifies which network services are started automatically

The Network file

The Network file, which lives in `/etc/sysconfig`, specifies network-wide configuration settings for your network. Here's a typical Network file:

```
NETWORKING=yes
HOSTNAME=LSERVER
GATEWAY=10.0.0.1
```

This file specifies that networking is enabled, the computer's host name is LSERVER, and the default gateway address is 10.0.0.1.

The following paragraphs describe all the settings that are valid for this file:

- NETWORKING: Specifies YES or NO to enable or disable networking for the computer.
- HOSTNAME: Specifies the host name for this computer. You should also specify this name in /etc/hostname, although that file is considered obsolete and is used only by some old programs. Note that this can be a simple host name (like LSERVER) or a fully qualified domain name (like Lserver.LoweWriter.com).
- FORWARD_IPV4: Specifies YES or NO to enable or disable IP forwarding. Specify FORWARD_IPV4=YES to set up a router.
- GATEWAY: Specifies the IP address of the computer's Default Gateway. If the network has a gateway router, specify its address here. If this computer is the network's gateway router, specify the gateway IP address provided by your ISP.
- GATEWAYDEV: Specifies the interface (such as eth0) that should be used to reach the gateway.

The interface configuration files

Each network interface has a configuration file located in /etc/NetworkManager/system-connections. The file name is the device name (for example, eth0) followed by .nmconnection. So, for example, the configuration file for the etho interface is named eth0.nmconnection.



TIP

This file is created and updated by the network configuration feature in program, so you don't have to edit it directly (if you don't want to).

Here's a typical interface configuration file for an interface that has a static IP address:

```
[connection]
id=eth0
uuid=7453df59-e470-3185-a917-bdb7d80168fa
type=ethernet
autoconnect-priority=-999
interface-name=eth0
timestamp=1718484881
```

```
[ethernet]

[ipv4]
address1=10.0.0.30/8,10.0.0.1
dns=75.75.75.75;75.75.75.76;
dns-search=loweriter.com;
method=manual

[ipv6]
addr-gen-mode=eui64
method=auto

[proxy]
```

Here's an example for an interface that uses DHCP:

```
[connection]
id=eth0
uuid=7453df59-e470-3185-a917-bdb7d80168fa
type=ethernet
autoconnect-priority=-999
interface-name=eth0
timestamp=1718486266

[ethernet]

[ipv4]
method=auto

[ipv6]
addr-gen-mode=eui64
method=auto

[proxy]
```

Here, the configuration file doesn't have to specify the IP address information because the interface gets that information from a DHCP server.

The following paragraphs describe the settings that you're most likely to modify in this file:

- `id`: The device to be configured, such as `eth0` or `eth1`.
- `interface-name`: The displayed name of the interface, usually set the same as the `id`.

- `method`: How the interface gets an IP address. To use DHCP, specify `method=auto`. To manually configure the IP address, specify `method=manual`.
- `address1`: The IP address for the adapter in slash notation, followed by the gateway address.
- `dns`: One or two DNS server addresses.
- `dns-search`: Your domain name.

The Hosts file

The `Hosts` file is a simple list of IP addresses and the host names associated with each address. You can find it in the `/etc` directory. Think of the Hosts file as a local DNS database of sorts. Whenever Linux needs to resolve a DNS name, it first looks for the name in the Hosts file. If Linux finds the name there, it doesn't have to do a DNS lookup; it simply uses the IP address found in the Hosts file.

For small networks, common practice is to list the host name for each computer on the network in the Hosts file on each computer. Then, whenever you add a new computer to the network, you just update each computer's Hosts file to include the new computer. That's not so bad if the network has just a few computers, but you wouldn't want to do it that way for a network with 1,000 hosts. That's why other name resolution systems are more popular for larger networks.

The default Linux Hosts file looks something like this:

```
# Loopback entries; do not change.
# For historical reasons, localhost precedes localhost.localdomain:
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
# See hosts(5) for proper format and other examples:
# 192.168.1.10 foo.example.org foo
# 192.168.1.13 bar.example.org bar
```

Here, the names `localhost.localdomain` and `localhost` both resolve to 127.0.0.1, which is the standard local loopback address.

Here's an example of a Hosts file that has some additional entries:

```
# Do not remove the following line, or various programs that
# require network functionality will fail.
127.0.0.1   LServer localhost.localdomain localhost
10.0.0.1    linksys
10.0.0.100  ward.cleaver.com ward
```

```
10.0.0.101  june.cleaver.com june
10.0.0.102  wally.cleaver.com wally
10.0.0.103  theodore.cleaver.com theodore beaver
```

Here, I defined host names for each of the Cleaver family's four computers and their Linksys router. Each computer can be accessed by using one of two names (for example, `ward.cleaver.com` or just `ward`), except the last one, which has three names.

The resolv.conf file

The `resolv.conf` file lists the DNS nameservers that can be consulted to perform DNS lookups. A typical `resolv.conf` file looks like this:

```
# Generated by NetworkManager
nameserver 75.75.75.75
nameserver 75.75.75.76
nameserver 2001:558:feed::1
# NOTE: the libc resolver may not support more than 3
# nameservers. The nameservers listed below may not be
# recognized.
nameserver 2001:558:feed::2
```

If you have set up a nameserver on your own network, its IP address should be the first one listed.

Displaying Your Network Configuration with the ifconfig Command

Linux doesn't have an `ipconfig` command like Windows. Instead, the command that you use to display information about your network configuration is `ifconfig`. You can also use this command to set network configuration options, but in most cases, using the Network Configuration program or directly editing the network configuration files is easier.

If you enter `ifconfig` without any parameters, you get output similar to the following:

```
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.0.0.30  netmask 255.255.255.0  broadcast 10.0.0.255
        inet6 fe80::215:5dff:fe68:9a08  prefixlen 64  scopeid 0x20<link>
```

```
        inet6 2601:204:380:11f0::d444  prefixlen 128  scopeid 0x0<global>
        inet6 2601:204:380:11f0:215:5dff:fe68:9a08  prefixlen 64
            scopeid 0x0<global>
        ether 00:15:5d:68:9a:08  txqueuelen 1000  (Ethernet)
        RX packets 253824  bytes 42873812 (40.8 MiB)
        RX errors 0  dropped 33  overruns 0  frame 0
        TX packets 245179  bytes 267178117 (254.8 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1000  (Local Loopback)
        RX packets 4  bytes 240 (240.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 4  bytes 240 (240.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

From this output, you can tell that the IP address of the Ethernet adapter (`eth0`) is 10.0.0.30, the netmask is 255.255.255.0, and the broadcast address is 10.0.0.255. You can also see transmit and receive statistics as well as information about the hardware configuration, such as the MAC address and the adapter's interrupt and memory base address assignments.

There are two important bits of information that unfortunately aren't reported by `ifconfig`: the gateway address and the DNS servers.

You can find the gateway by running the route command, which displays output similar to this:

```
[dlowe@localhost etc]$ route -n
Kernel IP routing table
Destination  Gateway   Genmask        Flags Metric Ref  Use Iface
0.0.0.0      10.0.0.1  0.0.0.0        UG    100    0      0 eth0
10.0.0.0     0.0.0.0   255.255.255.0  U     100    0      0 eth0
```

The gateway will be on the line where the destination is 0.0.0.0 — in this case, 10.0.0.1.

To find the gateway address, enter the following command:

```
cat /etc/resolv.conf
```

This lists the contents of the `resolv.conf` file, which identifies the DNS servers.

Linux offers many other commands that can help you configure and troubleshoot a network. Many of these commands are described in detail in Book 8, Chapter 5.

IN THIS CHAPTER

- Dealing with DHCP
- Running a DNS server
- Understanding BIND configuration files

Chapter 4
Running DHCP and DNS

One of the main reasons why many network administrators add Linux servers to their networks is to run internet services, such as DHCP and DNS. These services were originally developed for the Unix environment, so they tend to run better under Linux than they do under Windows.

Well, that's the theory, at least. The most recent versions of Windows are probably just as good at running these services as Linux. Still, if you prefer to set up these services on a Linux server, this chapter is for you.

Running a DHCP Server

DHCP is the TCP/IP protocol that automatically assigns IP addresses to hosts as they come on the network. (DHCP stands for Dynamic Host Configuration Protocol, but that won't be on the test.) For a very small network (say, fewer than ten hosts), you don't really need DHCP: You can just configure each computer to have a static IP address. For larger networks, however, DHCP is almost a must. Without DHCP, you have to manually plan your entire IP address scheme and manually configure every computer with its IP information. Then, if a critical address — such as your internet gateway router or your DNS server address — changes, you have to manually update each computer on the network. As you can imagine, DHCP can save you a lot of time.



TIP

For the complete lowdown on DHCP, please read Book 2, Chapter 5. In the following sections, I show you how to install and configure a DHCP server on the Fedora 12 Linux distribution.

Installing DHCP

You can quickly find out whether DHCP is installed on your system by entering the following command from a shell prompt:

```
sudo dnf install dhcp-server
```

If DHCP has already been installed, the `dnf` command will let you know that the package is already installed and that it has nothing to do. Otherwise, the `dnf` command will ask your permission to install the package:

```
Total download size: 1.3 M
Installed size: 4.1 M
Is this ok [y/N]:
```

Enter **y** to proceed with the installation. After a few moments, `dnf` will announce that the installation is complete.

Configuring DHCP

You configure DHCP settings through a file called `dhcpd.conf` that lives in the `/etc/dhcp` directory. Fedora provides you with a sample configuration file located at the following path:

```
/usr/share/doc/dhcp-server/dhcpd.conf.example
```

Open this file in `vi` to review it. Listing 4-1 shows a portion of the sample configuration file. Note that the exact contents of this file vary from release to release. For brevity, I've omitted portions of the configuration file that are for less common DHCP situations.

LISTING 4-1: **The Example dhcpd.conf File**

```
# dhcpd.conf
#
# Sample configuration file for ISC dhcpd
#

# option definitions common to all supported networks...
```

```
option domain-name "example.org";
option domain-name-servers ns1.example.org, ns2.example.org;

default-lease-time 600;
max-lease-time 7200;

# If this DHCP server is the official DHCP server for the local
# network, the authoritative directive should be uncommented.
#authoritative;

# This is a very basic subnet declaration.

subnet 10.254.239.0 netmask 255.255.255.224 {
  range 10.254.239.10 10.254.239.20;
  option routers rtr-239-0-1.example.org, rtr-239-0-2.example.org;
}
```

The following paragraphs describe some of the key points of this file:

- `option domain-name`: This line provides the domain name for the network.
- `option domain-name-servers`: This line provides the name or IP addresses of your DNS servers.
- `subnet`: This line specifies a subnet that's managed by this DHCP server. Following the subnet ID and netmask is an opening bracket; all the options that appear between this bracket and the closing bracket in the last line of the file belong to this subnet. In some cases, your DHCP server may dole out IP configuration information for two or more subnet groups. In that case, you need additional subnet groups in the configuration file.
- `range`: This line specifies the range of addresses that the DHCP server will assign for this subnet.
- `option routers`: This line provides the IP address of the Default Gateway.
- `default-lease-time`: This line determines the default lease time in seconds. The default in the example file (600 seconds, or 10 minutes) is far too short for anything other than a testing situation. More common settings are:
  - 86400 (1 day)
  - 604800 (7 days)
  - 2592000 (30 days)
- `max-lease-time`: This line determines the maximum life of a lease.
- `host`: This line specifies a reservation. The host group specifies the MAC address for the host and the fixed IP address to be assigned.

Starting DHCP

After you set up the configuration file, you can start DHCP by opening a terminal window or virtual console and entering the following command:

```
systemctl start dhcpd
```

If an error exists in the configuration file, a message to that effect is displayed. You have to edit the file in order to correct the error and then start the DHCP service again.

You should also restart the service whenever you make a change to the configuration file. To restart DHCP, enter this command:

```
systemctl restart dhcpd
```

To automatically start DHCP whenever you start the computer, run this command:

```
chkconfig –level 35 dhcpd on
```

Running a DNS Server

Linux comes with BIND, the best DNS server that money can buy. BIND is an extremely powerful program. Some people make entire careers of setting up and configuring BIND. In these few short pages, I just touch on the very basics of setting up a DNS server on your network.


TIP

You can find plenty of details about DNS in Book 2, Chapter 6. Please review that chapter before playing with BIND on your Linux system.

Installing BIND

You can quickly find out whether BIND is installed on your system by entering the following command from a shell prompt:

```
sudo dnf install bind
```

If BIND has already been installed, the `dnf` command will let you know that the package is already installed and that it has nothing to do. Otherwise, the `dnf` command will ask your permission to install the package. Enter **y** to install the package.

After BIND has been installed, you can start its service (which is called `named`) by entering this command:

```
systemctl start named
```

You can also go to Cockpit to start the service.

Editing BIND configuration files

Like most things Linux, BIND is configured by editing configuration files. These files live in one of two places: `/etc` or `/var/named`. The following sections describe the most important configuration files.

named.conf

This file, found in the `/etc` directory, is the basic BIND configuration file. This file contains global properties and links to the other configuration files.

Listing 4-2 shows a typical `named.conf` file.

LISTING 4-2: The named.conf File

```
Cd //
// named.conf
//
// Provided by Red Hat bind package to configure the ISC BIND named(8) DNS
// server as a caching only nameserver (as a localhost DNS resolver only).
//
// See /usr/share/doc/bind*/sample/ for example named configuration files.
//

options {
        listen-on port 53 { 127.0.0.1; };
        listen-on-v6 port 53 { ::1; };
        directory       "/var/named";
        dump-file       "/var/named/data/cache_dump.db";
        statistics-file "/var/named/data/named_stats.txt";
        memstatistics-file "/var/named/data/named_mem_stats.txt";
        secroots-file   "/var/named/data/named.secroots";
        recursing-file  "/var/named/data/named.recursing";
        allow-query     { localhost; };

        /*
         - If you are building an AUTHORITATIVE DNS server, do NOT
```

```
           enable recursion.
         - If you are building a RECURSIVE (caching) DNS server,
           you need to enable
           recursion.
         - If your recursive DNS server has a public IP address,
           you MUST enable access control to limit queries to your
           legitimate users. Failing to do so will cause your server
           to become part of large scale DNS amplification attacks.
           Implementing BCP38 within your network would greatlyreduce
           such attack surface
        */
        recursion yes;

        dnssec-validation yes;

        managed-keys-directory "/var/named/dynamic";
        geoip-directory "/usr/share/GeoIP";

        pid-file "/run/named/named.pid";
        session-keyfile "/run/named/session.key";

        /* https://fedoraproject.org/wiki/Changes/CryptoPolicy */
        include "/etc/crypto-policies/back-ends/bind.config";
};

logging {
        channel default_debug {
                file "data/named.run";
                severity dynamic;
        };
};

zone "." IN {
        type hint;
        file "named.ca";
};

include "/etc/named.rfc1912.zones";
include "/etc/named.root.key";
```

The `zone` lines name the zone files for each domain for which the server is responsible. Initially, the file contains just one zone, which is used for all lookups that aren't covered by local zones you add to the file. This zone is defined by the following lines:

```
zone "." IN {
        type hint;
        file "named.ca";
};
```

In the first line (`zone`), a single period is used as the domain name. The second line indicates that the zone type is `hint`, which means that the zone will refer to the DNS root servers. Finally, the third line provides the name of the file that identifies the root servers. We'll look at this file in a moment. But first, I want to show you how to edit this file to include your own zones.

To do that, you can add additional `zones` statements to the `named.conf` file. But I prefer to define your own zones in a separate configuration file and simply add an `include` statement to the `named.conf` file to copy in your zone configurations. You can name the file anything you want, but `named.custom` is a common choice. Simply add the following line to the end of the `named.conf` file:

```
Include "/etc/named.custom";
```

Then, you can create this file in the `/etc` directory and add a `zone` statement for your domain's zone.

Here's an example of a `zone` statement to create a new zone named `lowewriter.com`:

```
zone "lowewriter.com" IN {
    type master;
    file "lowewriter.com.zone";
};
```

Note that this zone specifies the type as `master`, which means that this server is authoritative for the zone. Then, the `file` option provides the name of the zone file that actually defines the zone. This file (`lowewriter.com.zone`) resides in `/var/named`. You learn how to create it in the next section.

Zone files

Each zone for which your DNS server is authoritative should have a zone file, named *`domain`*`.zone` and located in the `/var/named` directory. The name of this file must correspond to the filename you specified in either the `named.conf` file or in a file such as `named.custom`, which you included in the `named.conf` file.

Here's a typical zone file, named `lowewriter.com.zone`:

```
$TTL 86400
@ IN SOA dns01.lowewriter.local. root.lowewriter.local. (
     20201004 ;Serial
     3600     ;Refresh
     1800     ;Retry
     604800   ;Expire
     86400    ;Minimum TTL
)
@    IN NS      dns01.lowewriter.local.
@    IN A       10.0.0.30
dns01     IN A  10.0.0.30
fs01      IN A  10.0.0.40
mail      IN A  10.0.0.20
```

Table 4-1 lists the most common types of records that appear in zone files. For a complete description of each of these record types, see Book 2, Chapter 6.

TABLE 4-1 Common Resource Record Types

Type	Name	Description
SOA	Start Of Authority	Identifies a zone and provides settings for the zone such as the serial number (often derived from the date), refresh, retry, and expiration intervals, and a minimum
NS	Name Server	Identifies a name server that is authoritative for the zone
A	Address	Maps a fully qualified domain name to an IP address
CNAME	Canonical Name	Creates an alias for a fully qualified domain name
MX	Mail Exchange	Identifies the mail server for a domain
PTR	Pointer	Maps an IP address to a fully qualified domain name for reverse lookups

Restarting BIND

Whenever you make changes to BIND configuration files, you should restart the `named` service to apply the changes. To do that, use this command:

```
systemctl restart named
```

You can also restart the `named` service in Cockpit; just summon the Services page, locate the `named` service, and restart it.

IN THIS CHAPTER

- Discovering the basics of command shells
- Identifying file and directory commands
- Discovering commands that help with packages and services
- Figuring out commands for managing users and groups
- Becoming familiar with networking commands

Chapter 5
Linux Commands

Linux is, at its core, a command-line-driven operating system. You can add a nice graphical user interface (GUI) such as GNOME, and you can use a web-based administration tool such as Cockpit, but most server administration is done through the command line. In this chapter, I explain the ins and outs of using the command line and explore the details of using some of the most common Linux commands.

Command Shell Basics

A *shell* is a program that accepts commands from a command prompt and executes them. The shell displays a prompt to let you know it's waiting for a command. When you type the command and press Enter, the system reads your command, interprets it, executes it, displays the results, and then displays the prompt again so that you can enter another command.



TIP

Linux commands are case sensitive, so be careful about capitalization when you type Linux commands.

Getting to a shell

You can work with Linux commands directly from one of the six virtual consoles. If you like the responsiveness of text mode, virtual consoles are for you. To switch to a virtual console, press Ctrl+Alt+Fx. For example, press Ctrl+Alt+F1 to switch to virtual console 1. When you're in a virtual console, you have to answer the logon prompt with a valid username and password.

You can also remotely access a shell using a TTY program such as PuTTY. For details on setting up this useful program, refer to Book 8, Chapter 2.



WARNING

For normal Linux users, the command shell prompt character is a dollar sign ($). If you see a hash mark (#) as the prompt character, it means you're logged on as `root`. Whenever you see a hash prompt, you should be extra careful about what you do because you can easily get yourself into trouble by deleting important files or otherwise corrupting the system.

Editing commands

Most of the time, you just type commands using the keyboard. If you make a mistake, you just type the command again, being careful not to repeat the mistake. However, Linux shells have several built-in editing features that can simplify the task of correcting a mistaken command or entering a sequence of similar commands:

- **Repeat:** If you want to repeat a command that you've used recently, press the up-arrow key. This action recalls your most recently executed commands. You can press Enter to execute a command as is, or you can edit the command before you execute it. You can also press the up-arrow key multiple times to scroll back through your recent commands, and if you overshoot the command you want to repeat, you can use the down-arrow key to scroll forward.
- **Cursor movement:** You can use the Home and End keys to move to the beginning and ending of a line. Alternatively, you can press Ctrl+A to move to the start of a line and Ctrl+E to move to the end of a line. You can also press Alt+F to move forward one word at a time and Alt+B to move backward one word at a time.
- **Clear the screen:** Press Ctrl+L to clear the screen.

Wildcards

Wildcards are one of the most powerful features of command shells. With wildcards, you can process all the files that match a particular naming pattern with

a single command. For example, suppose that you have a folder with 500 files in it, and you want to delete all the files that contain the letters Y2K and end with `.doc`, which happens to be 50 files. If you try to do this in GNOME, you'll spend ten minutes picking these files out from the list. From a shell, you can delete them all with the single command `rm *Y2K*.doc`.

You can use two basic wildcard characters. An asterisk (`*`) stands for any number of characters, including zero, and an exclamation mark (`!`) stands for just one character. Thus, `!Text.doc` matches files with names like `aText.doc`, `xText.doc`, and `4Text.doc`, but not `abcText.doc` or just `Text.doc`. However, `*Text.doc` matches any of those filenames.

You can also use brackets to indicate a range of characters to choose from. For example, `report[123]` matches the files `report1`, `report2`, or `report3`. You can also specify `report[1-5]` to match `report1`, `report2`, `report3`, `report4`, or `report5`. The wildcard `r[aeiou]port` matches files named `raport`, `report`, `riport`, `roport`, or `ruport`. As you can see, the possibilities are almost endless.

Redirection and piping

Redirection and piping are related techniques. *Redirection* lets you specify an alternative destination for output that will be displayed by a command or specify an alternative source for input that should be fed into a command. For example, you can save the results of an `ifconfig` command to `/home/doug/myconfig` like this:

```
$ ifconfig > /home/doug/myconfig
```

Here, the greater-than sign (`>`) is used to redirect the command's console output.

You can use two greater-than signs (`>>`) to redirect output to an existing file, writing the redirected output to the end of the file. For example:

```
$ ifconfig >> /home/doug/myconfigs
```

If a command accepts input from the keyboard, you can use input redirection to specify a file that contains the input that you want to feed to the command. For example, you can create a text file named `lookup.commands` with subcommands for a command such as `nslookup`. Then, you can feed those scripted subcommands to the `nslookup` command, like this:

```
$ nslookup < /home/doug/lookup.commands
```

Piping is a similar technique. It takes the console output from one command and feeds it into the next command as input. One of the most common uses of piping is to send the output of a command that displays a lot of information to the `more` program, which displays the output one page at a time. For example:

```
$ ifconfig | more
```

The vertical bar (`|`) is often called the *pipe character* because it's the symbol used to indicate piping.

Environment variables

The shell makes several environment variables available to commands. An *environment variable* is a predefined value you can use in your commands to provide commonly used information, such as the name of the current user or the operating system version. You can use an environment variable anywhere in a command by typing **$** (dollar sign) followed by the environment variable name. For example, this command

```
$ echo This is $HOSTNAME running on an $HOSTTYPE
```

displays a line such as

```
This is LSERVER running on an i386
```

Table 8-1 lists some of the more useful environment variables that are available to you and your commands.

Shell scripts

A *shell script* is simply a text file that contains one or more commands which you can execute in sequence by running the script. The simplest shell scripts are just lists of commands, but advanced shell scripts can include complicated scripting statements that border on a full-featured programming language.


TIP

You can create shell scripts by using any text editor. One of the easiest editors is `nano`, which you can invoke from a shell prompt. But if you want to be a real Linux guru, take a few moments to learn how to use Vi, a powerful text-mode editor. To create or edit a file in Vi, type the command **vi** followed by a filename. Then, type away. For more information about Vi, refer to Book 8, Chapter 2.

TABLE 8-1 Environment Variables

Variable	Description
HOME	The current user's home directory
HOSTNAME	The computer's host name
HOSTTYPE	The host computer type
OSTYPE	The operating system
PATH	The search order for executable programs
PROMPT_COMMAND	The command used to generate the prompt
PWD	The present working directory
SHELL	The shell being used
USERNAME	The current username

After you create a shell script, you have to grant yourself execute permission to run the script. For example, to grant yourself permission to run a script named `myscript`, use this command:

```
$ chmod 755 myscript
```

To run a shell script, you use the `sh` command and provide the name of the script file. For example:

```
$ sh myscript
```

Running a command with root-level privileges

Many Linux commands perform actions that can only be done by the root user. To avoid the need to switch back and forth between your normal user account and the root user account, you can use the `sudo` command to perform any command using root-level permissions. To do so, just prefix the command you want to execute with the word `sudo`. For example:

```
$ sudo dnf install httpd
```

Here, the command `dnf install httpd` is executed using root-level permissions.

To use `sudo`, your Linux account must be configured by the root user to allow root-level access. For more information about how to do that, refer to Book 8, Chapter 2.

Directory- and File-Handling Commands

Because much of Linux administration involves working with configuration files, you frequently need to use the basic directory- and file-handling commands presented in this section.

The pwd command

This command displays the current directory, which is called the *present working directory* — hence the command name `pwd`. Here's the syntax:

```
pwd
```

Enter this command, and you get output similar to the following:

```
$ pwd
/home/doug
```

The cd command

The `cd` command changes the current working directory. The syntax is as follows:

```
cd directory
```

You may want to follow the `cd` command with a `pwd` command to make sure that you changed to the right directory. For example:

```
$ cd /etc/mail
$ pwd
/etc/mail
```

To change to a subdirectory of the current directory, omit the leading slash from the directory name. For example:

```
$ pwd
/home
```

```
$ cd doug
$ pwd
/home/doug
```

You can also use the double-dot (`..`) to represent the parent of the current directory. Thus, to move up one level, use the command `cd ..`, as follows:

```
$ pwd
/home/doug
$ cd ..
$ pwd
/home
```

The mkdir command

To create a new directory, use the `mkdir` command. It has the following syntax:

```
mkdir directory
```

Here's an example that creates a subdirectory named `images` in the current directory:

```
$ mkdir images
```

This example creates a directory named `/home/doug/images`:

```
$ mkdir /home/doug/images
```

The rmdir command

The `rmdir` command removes a directory. It has the following syntax:

```
rmdir directory
```

Here's an example:

```
$ rmdir /home/doug/images
```

Here, the `/home/doug/images` directory is deleted. Note that the directory must be empty to be removed, so you have to first delete any files in the directory.

The ls command

The `ls` command lists the contents of the current directory. Here's the syntax:

```
ls [options] directory
```

The following paragraphs describe the more important options for the `ls` command:

- `-a`: Lists all the files in the directory, including files that start with a period
- `-c`: Sorts entries by the time the files were last modified
- `-d`: Lists only directory names
- `-l`: Displays in long format
- `-r`: Displays files in reverse order
- `-R`: Lists the contents of all subdirectories, and subdirectories of subdirectories, and subdirectories of subdirectories of subdirectories; in other words, lists subdirectories recursively
- `-s`: Displays file sizes
- `-S`: Sorts files by size
- `-t`: Sorts files by timestamp.
- `-u`: Sorts files by the time the files were last accessed.
- `-X`: Sorts files by their extensions.

Without arguments, the `ls` command lists all the files in the current directory, like this:

```
$ pwd
/etc/mail
$ ls
access helpfile Makefile submit.cf virtusertable
access.db local-host-names sendmail.cf submit.cf.bak virtusertable.db
domaintable mailertable sendmail.mc submit.mc
domaintable.db mailertable.db statistics trusted-users
```

You can limit the display to certain files by typing a filename, which can include wildcards. For example:

```
$ ls a*
access access.db
```

You can also specify the directory that you want to display, like this:

```
$ ls /etc/httpd
conf conf.d logs modules run
```

To display detailed information about the files in the directory, use the `-l` switch, as in this example:

```
$ ls /etc/mail/s* -l
-rw-r--r-- 1 root root 57427 Jul 19 16:35 sendmail.cf
-rw-r--r-- 1 root root 5798 Feb 24 16:15 sendmail.mc
-rw------- 1 root root 628 Jul 24 17:21 statistics
-rw-r--r-- 1 root root 39028 Jul 19 17:28 submit.cf
-r--r--r-- 1 root root 39077 Feb 24 16:15 submit.cf.bak
-rw-r--r-- 1 root root 953 Feb 24 16:15 submit.mc
```

The cp command

The `cp` command copies files. Here's the basic syntax:

```
cp [options] source-file destination-file
```

The following list describes the more important options for the `cp` command:

- `-a`: The same as `-dpR`.
- `-b`: Makes backup copies of existing files before they're overwritten. Sounds like a good plan to me.
- `-d`: Copies links rather than the files the links point to.
- `-f`: Removes files that will be overwritten.
- `-i`: Interactively prompts for each file to be overwritten.
- `-l`: Creates links to files rather than actually copying file contents.
- `-p`: Preserves ownership and permissions.
- `-R`: Copies the contents of subdirectories recursively.
- `-s`: Creates symbolic links to files rather than actually copying file contents.
- `-u`: Replaces destination files only if the source file is newer.

To make a copy of a file within the same directory, use `cp` like this:

```
$ cp sendmail.cf sendmail.cf.backup
```

If you want to copy a file to another directory without changing the filename, use `cp` like this:

```
$ cp sendmail.cf /home/doug
```

You can use wildcards to copy multiple files:

```
$ cp send* /home/doug
```

To include files in subdirectories of the source file, use the `–R` (for *recursive*) switch, like this:

```
$ cp –R /etc/*.cf /home/doug
```

In this example, all files in the `/etc` directory or any of its subdirectories that end with `.cf` are copied to `/home/doug`.

The rm command

The `rm` command deletes files. The syntax is as follows:

```
rm [options] file
```

The options are described in the following paragraphs:

- `–f`: Removes files that will be overwritten
- `–i`: Interactively prompts for each file to be overwritten
- `–R`: Copies the contents of subdirectories recursively

To delete a single file, use it like this:

```
$ rm any.old.file
```

To delete multiple files, use a wildcard:

```
$ rm any.*
```

To delete an entire directory, use the `-r` switch:

```
$ rm -r /doug/old.files
```

The mv command

The `mv` command moves files or renames them. In Linux, moving and renaming a file is essentially the same thing. Moving a file changes the file's directory location but leaves its name the same. Renaming a file leaves the file in the same directory but changes the file's name.

The syntax of the `mv` command is

```
mv [options] source-file destination
```

The following paragraphs describe the options:

- `-b`: Makes backup copies of existing files before they're overwritten. Still sounds like a good plan to me.
- `-f`: Removes files that will be overwritten.
- `-i`: Interactively prompts for each file to be overwritten.
- `-u`: Replaces destination files only if the source file is newer.

To move a file to another directory, provide a filename for the first argument and a directory for the second, like this:

```
$ mv monthly.report /home/Kristen/
```

To rename a file, provide filenames for both arguments:

```
$ mv monthly.report august.monthly.report
```

The cat command

The `cat` command displays the contents of a file. It has the following syntax:

```
cat [options] [filename...]
```

The filename is optional. If you omit the filename, the `cat` command obtains its input from the console, which you can redirect if you want.

And, you can specify more than one filename. If you do, the files are combined to create a single output stream.

Here are some of the options you can use:

- –A: Displays new line characters as $, tab characters as ^I, and control characters with a caret (^)
- –b: Numbers all nonblank lines as they're displayed
- –e: Displays new line characters as $ and control characters with a caret (^)
- –E: Displays new line characters as $
- –n: Numbers lines as they are displayed
- –s: Squeezes multiple spaces down to a single space

Here's a basic example:

```
$ cat /etc/hosts
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1 LSERVER localhost.localdomain localhost
$
```

If you don't provide any filename arguments, the `cat` command copies text from the keyboard and displays it on the console. You can use the `cat` command along with output redirection to quickly create a short text file, like this:

```
$ cat >mytext
This is line one.
This is line two.
This is line three.
‹ctrl+D›
$
```

For the last line, press Ctrl+D. This signals the end of the input to the `cat` command.

Commands for Working with Packages and Services

As a Linux administrator, you frequently need to start and stop services and check the status of installed packages or install new packages. The following sections describe the Linux commands that help you to perform these tasks.

The service command

You use the `service` command to check the status of services and to start, stop, or restart services. You need to restart a service whenever you make a configuration change in order for your changes to take effect. Here's the basic syntax:

```
service [service] [ start | stop | restart ]
```

The following paragraphs describe some typical uses of the `service` command:

- To check the status of the `httpd` service (Apache), use this command:

```
$ service httpd status
```

- To stop the `httpd` service:

```
$ service httpd stop
```

- To start the `httpd` service:

```
$ service httpd start
```

- To restart the `httpd` service:

```
$ service httpd restart
```

The only trick to using the `service` command is that you have to know the name of the service. If you're not sure of the name, you can run the `service` command to display the status of all services, like this:

```
$ service --status-all
```

It will take a few moments to list all the services, but after the command is done, you can scroll through the list to find the service that you're looking for.

Table 8-2 lists some of the more common services.

TABLE 8-2 Common Linux Services

Service	Description
atd	Runs commands scheduled by the at command
autof	Automatically mounts file systems
crond	Runs programs at specified times
dhcpd	The DHCP server
finger	The internet finger service
httpd	The Apache web server
imap	The IMAP mail protocol
imaps	Secure IMAP service (SSL)
ipop3	The POP3 mail protocol
iptables	Automatic packet filtering for firewalls
isdn	ISDN services
named	The BIND DNS server
netf	The network file system
network	Activates and deactivates all network interfaces
nfs	Native Unix/Linux network file sharing
pop3s	Secure POP3 service (SSL)
sendmail	The Sendmail service
smb	The Samba file- and printer-sharing service
snmpd	SNMP
telnet	The Telnet server

The yum and dnf commands

yum, which stands for *Yellowdog Updater Modified*, is a tool for installing and updating packages on a Linux system. Until recently, yum was the preferred way

to install packages on a Linux system. One of the chief advantages of `yum` over earlier package managers is that `yum` not only installs packages you tell it to install, but also automatically installs any packages that are prerequisites to the packages you ask `yum` to install.

`yum` has recently been superseded by a faster package manager called `dnf`. `dnf` and `yum` share pretty much the same command line options, so `yum` and `dnf` commands are mostly interchangeable.

Although `dnf` has many command line options, the most common way to use it is as follows:

```
dnf install package-name
```

For example, to install a package named `postfix`, you would use this command:

```
$ dnf install postfix
```

Note that installing a package is an action that requires root-level permissions to perform, so you'll usually use `dnf` or `yum` along with the `sudo` command, like this:

```
$ sudo dnf install postfix
```


Commands for Administering Users

The following sections describe the Linux commands that you can use to create and manage user accounts from a command shell.


TIP

Prefix these commands with `sudo` so they run with root privileges.

The useradd command

The useradd command creates a user account. Here's the basic syntax for adding a new user:

```
useradd [options] username
```

You can also use this command to change the default options for new users. In that case, the syntax is more like this:

```
useradd -D [options]
```

If you use `-D` with no options, a list of the current default settings will be shown.

Here are some of the more commonly used options:

- `-c` *`comment`*: Typically the user's full name
- `-d` *`home-dir`*: The home directory of the new user
- `-g` *`group`*: The initial logon group for the user
- `-G` *`groups`*: Additional groups the user should belong to
- `-m`: Creates the new user's home directory if it doesn't exist already
- `-s` *`shell-path`*: Specifies the user's logon shell

The following option is valid only with `-D` (which sets a default option):

- `-b` *`base-dir`*: Provides the default base directory if a home directory is not specified

In its most basic form, the `useradd` command creates a user with default option settings:

```
$ useradd theodore
```

This command creates a user named `theodore`.

Here's a command that specifies the user's full name in the comment option:

```
$ useradd -c 'Theodore Cleaver' theodore
```

The following command creates a temporary account named `ghost` that expires on Halloween 2026:

```
$ useradd -e 2026-10-31 ghost
```

If you want to see what the default values are for account options, use the `-D` option without any other parameters:

```
$ useradd -D
GROUP=100
HOME=/home
INACTIVE=-1
EXPIRE=
SHELL=/bin/bash
SKEL=/etc/skel
```

The usermod command

The `usermod` command modifies an existing user. It has the following syntax:

```
usermod [options] username
```

Here are some of the more commonly used options:

- `-c` *`comment`*: Typically the user's full name
- `-d` *`home-dir`*: The home directory of the new user
- `-g` *`group`*: The initial logon group for the user
- `-G` *`groups`*: Additional groups the user should belong to
- `-m`: Creates the new user's home directory if it doesn't exist already
- `-s` *`shell-path`*: Specifies the user's logon shell
- `-l`: Locks an account
- `-u`: Unlocks an account

Here's an example that changes a user's full name:

```
$ usermod -c 'The Beave' theodore
```

The userdel command

The `userdel` command deletes a user. It has a simple syntax:

```
userdel [-r] username
```

If you specify `-r`, the user's home directory is deleted along with the account.

The chage command

The `chage` command modifies date policies for a user's passwords. It has the following syntax:

```
chage [options] username
```

The following paragraphs describe the options you can use:

- `-m` *days*: Specifies the minimum number of days allowed between password changes.
- `-M` *days*: Specifies the maximum number of days allowed between password changes.
- `-d` *date*: The date of the last password change.
- `-E` *date*: The date on which the account will expire.
- `-W` *days*: The number of days prior to the password expiring that the user will be warned the password is about to expire.
- `-I` *days*: The number of days of inactivity after the password has expired that the account is locked out. Specify 0 to disable this feature.

Here's an example that sets an account to expire on Halloween 2026:

```
$ chage -E 2026-10-31 ghost
```



TIP

If you specify a username but no other options, you're prompted to enter each option. This is a lot easier than trying to remember all the switches!

The passwd command

The `passwd` command changes the password for a user account. Its syntax is

```
passwd [user]
```

If you don't supply a user, the password for the current user is changed.

The `passwd` command prompts you to enter the new password twice to prevent the possibility of mistyping the password.

The newusers command

The `newusers` command provides an easy way to create a group of new user accounts. It reads a text file that contains one line for each new user, listing the user's name and password.

Here's the syntax of the `newusers` command:

```
newusers [filename]
```

If you omit the filename, `newusers` accepts input from the console.

Suppose that you have a file named `/root/island.users` that contains these lines:

```
gilligan ml9jiedr
skipper 1hiecr8u
professor dr0uxiaf
maryann choe7rlu
ginger jiuqled5
mrhowell j1emoaf1
lovey zo2priak
```

You can then create these seven stranded user accounts by issuing this command:

```
$ newusers /root/island.users
```



WARNING

Because the `newusers` file contains unencrypted passwords, you shouldn't leave it lying around. Require these new users to change their passwords immediately and delete the file you used to create the users.

The groupadd command

The `groupadd` command creates a new group. It has the following syntax:

```
groupadd [options] group
```

Although you have several possible options to use, the only one you're likely to need is `-r`, which creates a system group that has special privileges.

Here's an example that creates a group named `castaways`:

```
$ groupadd castaways
```

That's all you have to do to create a new group. To administer the group, you use the `gpasswd` command.

The groupdel command

The `groupdel` command deletes a group. It has the following syntax:

```
groupdel group
```

Here's an example that deletes a group named `castaways`:

```
$ groupdel castaways
```

Poof! The group is gone.

The gpasswd command

You use the `gpasswd` command to administer a group. This command has several different syntax options.

To change the group password:

```
gpasswd [ –r | –R ] group
```

To add a user:

```
gpasswd –a user group
```

To remove a user:

```
gpasswd –d user group
```

To create group administrators and/or members:

```
gpasswd [–A administrators ...] [–M members... ] group
```

The options are as follows:

- `–r`: Removes the password from the group.
- `–R`: Disables access to the group via the `newgrp` command.
- `–a`: Adds the specified user to the group.
- `–d`: Deletes the specified user from the group.
- `–A`: Specifies one or more group administrators. Use commas with no intervening spaces to separate the administrators from each other. Each administrator must be an existing user.
- `–M`: Specifies one or more group members. Use commas with no intervening spaces to separate the members from each other. Each member must be an existing user.

The following example adds seven group members and one administrator to a group called `castaways`:

```
$ gpasswd -A skipper -M skipper,gilligan,professor,maryann,ginger,mrhowell,lovey
   castaways
```

If the rest of the group finally decides to throw Gilligan off the island, they can remove him from the group with this command:

```
$ gpasswd -d gilligan castaways
```

Commands for Managing Ownership and Permissions

This section presents the details of the `chown` and `chmod` commands, which are the essential tools for assigning file system rights in the Linux environment.

The chown command

The `chown` command changes the owner of a file. Typically, the user who creates a file is the owner of the file. However, the owner can transfer the file to someone else via this command. The basic syntax of this command is

```
chown user file
```

For example, to change the owner of a file named `rescue.plans` to user `professor`, use this command:

```
$ chown professor rescue.plans
```

To change ownership of all the files in the directory named `/home/island` to `professor`, use this command:

```
$ chown professor /home/island
```



TIP

Issuing the following command would be a really bad idea:

```
$ chown gilligan rescue.plans
```

The chgrp command

Every file has not only an individual owner, but also a group owner. You can change the group ownership by using the `chgrp` command, which has the following basic syntax:

```
chgrp group file
```

For example, to grant the `castaways` group ownership of the file `rescue.plans`, use this command:

```
$ chgrp castaways rescue.plans
```

To change group ownership of all the files in the directory named `/home/island` to `castaways`, use this command:

```
$ chgrp castaways /home/island
```

The chmod command

The `chmod` command lets you change the permissions for a Linux file. Before explaining the syntax of the `chmod` command, you need to look at the cryptic way that Linux reports file permissions. Linux grants three different types of permissions — read, write, and execute — for three different scopes: owner, group, and everyone. That's a total of nine permissions.

When you use the `ls` command with the `-l` option, the permissions are shown as a ten-character string that begins with a hyphen if the entry is for a file or a `d` if the entry is for a directory. Then, the next nine letters are the nine permissions, in this order:

- Read, write, execute for the owner
- Read, write, execute for the group
- Read, write, execute for everyone

The letters `r`, `w`, or `x` appear if the permission has been granted. If the permission is denied, a hyphen appears.

For example, suppose the `ls -l` command lists these permissions:

```
-rw-r--r--
```

You interpret this permission string like this:

- The first hyphen indicates that this is a file, not a directory.
- The next three positions are `rw-`. Therefore, the owner has read and write permission on this file, but not execute permission.
- The next three positions are `r--`. That means the group owner has read permissions but not write or execute permission.
- The last three positions are also `r--`. That means that everyone else has read permission but not write or execute permission.

The full syntax of the `chmod` command is pretty complex. However, you can do most of what you need to do with this form:

```
chmod specification file
```

Here, `specification` is in the form `u=rwx`, `g=rwx`, or `o=rwx` to set the permissions for the user (owner), group, and others (everyone). You don't have to specify `r`, `w`, and `x`; you just list the permissions that you want to grant. For example, to grant read and write permission for the user to a file named `rescue.plans`, use this command:

```
$ chmod u=rw rescue.plans
```

You can also combine specifications, like this:

```
$ chmod u=rw,g=rw,o=r rescue.plans
```

To revoke all rights for the user, group, or others, don't type anything after the equal sign. For example, this command revokes all rights for others:

```
$ chmod o= rescue.plans
```

Networking Commands

The following sections present Linux commands that are used to display information about the network or configure its settings.

The hostname command

The `hostname` command simply displays the computer's host name. It has the following syntax:

```
hostname [name]
```

If you use this command without any parameters, the computer's host name is displayed. If you specify a name, the computer's host name is changed to the name you specify.

The ifconfig command

`ifconfig` displays and sets configuration options for network interfaces. Although you can configure an Ethernet adapter using this command, you'll rarely have to. Linux does a pretty good job of automatically configuring network adapters, and the GNOME-based Network Configuration tool supplied with the Red Hat distribution should be able to handle most network configuration chores. So you'll use `ifconfig` mostly to display network configuration settings.

The basic syntax for `ifconfig` is

```
ifconfig interface [address] [netmask mask] [broadcast
  broadcast]
```

Here are the options that you can use on the `ifconfig` command:

- *interface*: The symbolic name for your network adapter. It's typically `eth0` for the first Ethernet adapter or `eth1` for the second adapter.
- *address*: The IP address you want to assign to the interface, such as 192.168.1.100.
- *netmask*: The subnet mask to use, such as 255.255.255.0.
- *broadcast*: The broadcast, which should be the highest address on the subnet. For example: 192.168.1.255.

If you enter `ifconfig` without any parameters, the `ifconfig` command displays the current status of your network adapters, like this:

```
$ ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.0.0.30  netmask 255.255.255.0  broadcast 10.0.0.255
        inet6 fe80::215:5dff:fe68:9a08  prefixlen 64  scopeid 0x20<link>
```

```
        inet6 2601:204:380:11f0::d444  prefixlen 128  scopeid 0x0<global>
        inet6 2601:204:380:11f0:215:5dff:fe68:9a08  prefixlen 64  scopeid
   0x0<gl  obal>
        ether 00:15:5d:68:9a:08  txqueuelen 1000  (Ethernet)
        RX packets 677481  bytes 236476898 (225.5 MiB)
        RX errors 0  dropped 88  overruns 0  frame 0
        TX packets 533569  bytes 515063384 (491.2 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1000  (Local Loopback)
        RX packets 4  bytes 240 (240.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 4  bytes 240 (240.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

To change the IP address of an adapter, use `ifconfig` like this:

```
$ ifconfig eth0 192.168.1.200
```

The netstat command

The `netstat` command lets you monitor just about every aspect of a Linux server's network functions. This command can generate page after page of interesting information — if you know what it all means.

The two most common reasons to use `netstat` are to display the routing table and to display open TCP/IP connections. The syntax for displaying the routing table is

```
$ netstat -r
```

This results in a display similar to this:

```
Kernel IP routing table
Destination Gateway Genmask Flags MSS Window irtt Iface
192.168.1.0 * 255.255.255.0 U 0 0 0 eth1
192.168.1.0 * 255.255.255.0 U 0 0 0 eth0
127.0.0.0 * 255.0.0.0 U 0 0 0 lo
```

To display TCP/IP connections, use this syntax:

```
$ netstat -l
```

This results in a display similar to the following:

```
Active internet connections (only servers)
Proto Recv-Q Send-Q Local Address Foreign Address State
tcp 0 0 *:1024 *:* LISTEN
tcp 0 0 LSERVER:1025 *:* LISTEN
tcp 0 0 *:netbios-ssn *:* LISTEN
tcp 0 0 *:sunrpc *:* LISTEN
tcp 0 0 *:http *:* LISTEN
tcp 0 0 *:x11 *:* LISTEN
tcp 0 0 *:ssh *:* LISTEN
tcp 0 0 LSERVER:ipp *:* LISTEN
tcp 0 0 LSERVER:smtp *:* LISTEN
tcp 0 0 *:https *:* LISTEN
udp 0 0 *:1024 *:*
udp 0 0 LSERVER:1026 *:*
udp 0 0 192.168.1.20:netbios-ns *:*
udp 0 0 192.168.1.20:netbios-ns *:*
udp 0 0 *:netbios-ns *:*
udp 0 0 192.168.1.2:netbios-dgm *:*
udp 0 0 192.168.1.2:netbios-dgm *:*
udp 0 0 *:netbios-dgm *:*
udp 0 0 *:940 *:*
udp 0 0 *:sunrpc *:*
udp 0 0 *:631 *:*
.
.
.
```

From this display, you can tell which Linux services are actively listening on TCP/IP ports.

The ping command

The `ping` command is the basic troubleshooting tool for TCP/IP. You use it to determine whether basic TCP/IP connectivity has been established between two computers. If you're having any kind of network trouble between two computers, the first troubleshooting step is almost always to see whether the computers can ping each other.

The basic syntax of `ping` is straightforward:

```
ping [options] address
```

The options can be

- –c: The number of packets to send. If you omit this, `ping` continues to send packets until you interrupt it.
- –d: Floods the network with packets, as many as 100 per second. Use with care!
- `–i`: Specifies how many seconds to wait between sending packets. The default is one second. If you're having intermittent connection problems, you may try letting `ping` run for a while with this option set to a higher value, such as `60`, to send a packet every minute.
- –R: Displays the route the packets take to get to the destination computer.


WARNING


TIP

`ping` will continue to ping the destination computer until you interrupt it by pressing Ctrl+Z.

You can specify the host to ping by using an IP address, as in this example:

```
$ ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1) 56(84) bytes of data.
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=2.72 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.99 ms
64 bytes from 10.0.0.1: icmp_seq=3 ttl=64 time=1.38 ms
64 bytes from 10.0.0.1: icmp_seq=4 ttl=64 time=1.02 ms
64 bytes from 10.0.0.1: icmp_seq=5 ttl=64 time=2.59 ms
.
.
.
```

You can also ping by using a DNS name, as in this example:

```
$ ping www.dummies.com
PING www.dummies.com (2606:4700:4400::ac40:975c) 56 data bytes
64 bytes from 2606:4700:4400::ac40:975c: icmp_seq=1 ttl=59 time=13.9 ms
64 bytes from 2606:4700:4400::ac40:975c: icmp_seq=2 ttl=59 time=15.4 ms
64 bytes from 2606:4700:4400::ac40:975c: icmp_seq=3 ttl=59 time=15.5 ms
64 bytes from 2606:4700:4400::ac40:975c: icmp_seq=4 ttl=59 time=13.4 ms
64 bytes from 2606:4700:4400::ac40:975c: icmp_seq=5 ttl=59 time=16.2 ms
.
.
.
```

The route command

The `route` command displays or modifies the computer's routing table. To display the routing table, use `route` without any parameters. To add an entry to the routing table, use this syntax:

```
route add [ –net | –host ] address [options]
```

To delete an entry, use this syntax:

```
route del [ –net | –host ] address [options]
```

The available options are as follows:

- `netmask` *`mask`*: Specifies the subnet mask for this entry
- `gw` *`address`*: Specifies the gateway address for this entry
- `dev` *`if`*: Specifies an interface (such as `eth0` or `eth1`) for this entry

If you enter `route` by itself, with no parameters, you'll see the routing table, as in this example:

```
$ route
Kernel IP routing table
Destination     Gateway   Genmask       Flags Metric Ref    Use  Iface
192.168.1.0     *         255.255.255.0 U     0      0      0    eth1
192.168.1.0     *         255.255.255.0 U     0      0      0    eth1
169.254.0.0     *         255.255.0.0   U     0      0      0    eth1
127.0.0.0       *         255.0.0.0     U     0      0      0    lo
default         _gateway  0.0.0.0       UG    0      0      0    eth1
```

Suppose that your network has a second router that serves as a link to another private subnet, 192.168.2.0 (subnet mask 255.255.255.0). The interface on the local side of this router is at 192.168.1.200. To add a static route entry that sends packets intended for the 192.168.2.0 subnet to this router, use a command like this:

```
$ route add 192.168.2.0 netmask 255.255.255.0 gw 192.168.1.200
```

The traceroute command

The `traceroute` command displays a list of all the routers that a packet must go through to get from the local computer to a destination on the internet. Each one of these routers is a *hop.* If you're unable to connect to another computer, you can use `traceroute` to find out exactly where the problem is occurring.

Here's the syntax:

```
traceroute [-i interface] host
```

Although several options are available for the `traceroute` command, the one you're most likely to use is `-i`, which lets you specify an interface. This is useful if your computer has more than one network adapter.

9 Managing a Network

Contents at a Glance

CHAPTER 1: **Welcome to Network Administration** 757

CHAPTER 2: **Managing Remotely** 771

CHAPTER 3: **Managing Network Assets** 791

CHAPTER 4: **Solving Network Problems** 801

IN THIS CHAPTER

- Deciphering the many jobs of the network administrator
- Dusting, vacuuming, and mopping
- Managing the network users
- Choosing the right tools
- Getting certified

Chapter 1

Welcome to Network Administration

Help wanted! Network administrator to help small business get control of a network run amok. Must have sound organizational and management skills. Only moderate computer experience required. Part-time only.

Does this ad sound like one that your company should run? Every network needs a network administrator, whether the network has 2 computers or 200. Of course, managing a 200-computer network is a full-time job for probably more than one person, whereas managing a 2-computer network isn't. At least, it shouldn't be.

This chapter introduces you to the boring job of network administration. Oops . . . you're probably reading this chapter because you've been elected to be the network manager, so I'd better rephrase that: This chapter introduces you to the wonderful, exciting world of network management! Oh, boy! This is going to be fun!

Knowing What Network Administrators Do

Simply put, network administrators administer networks, which means that they take care of the tasks of installing, configuring, expanding, protecting, upgrading, tuning, and repairing the network. Network administrators take care of the network hardware, such as cables, hubs, switches, routers, servers, and clients, as well as network software, such as network operating systems, email servers, backup software, database servers, and application software. Most importantly, network administrators take care of network users by answering their questions, listening to their troubles, and solving their problems.

On a big network, these responsibilities constitute a full-time job. Large networks tend to be volatile: Users come and go, equipment fails, cables break, and life in general seems to be one crisis after another.

Smaller networks are much more stable. After you get your network up and running, you probably won't have to spend much time managing its hardware and software. An occasional problem may pop up, but with only a few computers on the network, problems should be few and far between.

Regardless of the network's size, all network administrators must attend to several common chores:

- **Equipment upgrades:** The network administrator should be involved in every decision to purchase new computers, printers, or other equipment. In particular, the network administrator should be prepared to lobby for the most network-friendly equipment possible, including professional-grade network switchers and firewalls, ample storage capacity, reliable servers, and an adequate backup system.
- **Configuration:** The network administrator must put on the pocket protector whenever a new computer is added to the network. The network administrator's job includes considering what changes to make to the cabling configuration, what computer name to assign to the new computer, how to integrate the new user into the security system, what rights to grant the user, and so on.
- **Software upgrades:** Every once in a while, your trusty network vendors release new versions of their software. The network administrator must study up on the new version and decide whether its new features are beneficial enough to warrant an upgrade. In most cases, the hardest part of upgrading to a new version of network software is determining the *migration path* — that is, how to upgrade your entire network to the new version while disrupting the network or its users as little as possible. Plan carefully!

- **Patches:** Between upgrades, Microsoft releases patches and service packs that fix minor problems with its server operating systems. For more information, see the section "Patching Up Your Operating System and Software" later in this chapter. (Other software vendors also regularly release patches and service packs, so it isn't only Microsoft software that must be kept up to date.)
- **Performance maintenance:** One of the easiest traps that you can get sucked into is the quest for network speed. The network is never fast enough, and users always blame the hapless network manager. So the administrator spends hours and hours tuning and tweaking the network to squeeze out that last 2 percent of performance.
- **Ho-hum chores:** Network administrators perform routine chores, such as backing up the servers, archiving old data, freeing up server hard drive space, and so on. Much of network administration is making sure that things keep working and finding and correcting problems before any users notice that something is wrong. In this sense, network administration can be a thankless job.
- **Software inventory:** Network administrators are also responsible for gathering, organizing, and tracking the entire network's software inventory. You never know when something is going to go haywire on Joe in Marketing's ancient Windows XP computer and you're going to have to reinstall that old copy of WordPerfect. Do you have any idea where the installation discs are?

Choosing the Part-Time Administrator

The larger the network, the more technical support it needs. Most small networks — with just a dozen or so computers — can get by with a part-time network administrator. Ideally, this person should be a closet computer geek: someone who has a secret interest in computers but doesn't like to admit it. Someone who will take books home and read them over the weekend. Someone who enjoys solving computer problems just for the sake of solving them.

The job of managing a network requires some computer skills, but it isn't entirely a technical job. Much of the work that the network administrator does is routine housekeeping. Basically, the network administrator dusts, vacuums, and mops the network periodically to keep it from becoming a mess.

Here are some additional ideas on picking a part-time network administrator:

- The network administrator needs to be an organized person. Conduct a surprise office inspection and place the person with the neatest desk in charge of the network. (Don't warn anyone in advance, or everyone may mess up his or her desk intentionally the night before the inspection.)
- Allow enough time for network administration. For a small network (say, no more than 20 or so computers), an hour or two each week is enough. More time is needed upfront as the network administrator settles into the job and discovers the ins and outs of the network. After an initial settling-in period, though, network administration for a small office network doesn't take more than an hour or two per week. (Of course, larger networks take more time to manage.)
- Make sure that everyone knows who the network administrator is and that the network administrator has the authority to make decisions about the network, such as what access rights each user has, what files can and can't be stored on the server, how often backups are done, and so on.
- Pick someone who is assertive and willing to irritate people. A good network administrator should make sure that backups are working *before* a hard drive fails and make sure that antivirus protection is in place *before* a virus wipes out the entire network. This policing will irritate people, but it's for their own good.
- In most cases, the person who installs the network is also the network administrator. This is appropriate because no one understands the network better than the person who designs and installs it.
- The network administrator needs an understudy — someone who knows almost as much about the network, is eager to make a mark, and smiles when the worst network jobs are delegated.
- The network administrator has some sort of official title, such as Network Boss, Network Czar, Vice President in Charge of Network Operations, or Dr. Network. A badge, a personalized pocket protector, or a set of Spock ears helps, too.

Establishing Routine Chores

Much of the network administrator's job is routine stuff — the equivalent of vacuuming, dusting, and mopping. Or if you prefer, changing the oil and rotating the tires every 3,000 miles. Yes, it's boring, but it has to be done.

- **Backup:** The network administrator needs to make sure that the network is properly backed up. If something goes wrong and the network isn't backed up, guess who gets the blame? On the other hand, if disaster strikes, yet you're able to recover everything from yesterday's backup with only a small amount of work lost, guess who gets the pat on the back, the fat bonus, and the vacation in the Bahamas?
- **Protection:** Another major task for network administrators is sheltering your network from the evils of the outside world. These evils come in many forms, including hackers trying to break into your network and virus programs arriving through email.
- **Clean-up:** Users think that the network server is like the attic: They want to throw files up there and leave them forever. No matter how much storage your network has, your users will fill it up sooner than you think. So the network manager gets the fun job of cleaning up the attic once in a while. Oh, joy. The best advice I can offer is to constantly complain about how messy it is up there and warn your users that spring cleaning is coming up.

Managing Network Users

Managing network technology is the easiest part of network management. Computer technology can be confusing at first, but computers aren't nearly as confusing as people. The real challenge of managing a network is managing the network's users.

The difference between managing technology and managing users is obvious: You can figure out computers, but you can never really figure out people. The people who use the network are much less predictable than the network itself. Here are some tips for dealing with users:

- **Training:** Training is a key part of the network manager's job. Make sure that everyone who uses the network understands it and knows how to use it. If the network users don't understand the network, they may unintentionally do all kinds of weird things to it.
- **Respect:** Never treat your network users like they're idiots. If they don't understand the network, it isn't their fault. Explain it to them. Offer a class. Buy them each a copy of this book and tell them to read it during their lunch hour. Hold their hands. But don't treat them like idiots.

- **Aids:** Make up a network cheat sheet that contains everything that the users need to know about using the network on one page. Make sure that everyone gets a copy.
- **Responsive:** Be as responsive as possible when a network user complains of a network problem. If you don't fix the problem soon, the user may try to fix it. You probably don't want that.


TIP

The better you understand the psychology of network users, the more prepared you'll be for the strangeness they often serve up. Toward that end, I recommend that you read the *Diagnostic and Statistical Manual of Mental Disorders* (also known as *DSM-V*) cover to cover.

Patching Up Your Operating System and Software

One of the annoyances that every network manager faces is applying software patches to keep your OS and other software up to date. A software *patch* is a minor update that fixes small glitches that crop up from time to time, such as minor security or performance issues. These glitches aren't significant enough to merit a new version of the software, but they're important enough to require fixing. Most patches correct security flaws that computer hackers have uncovered in their relentless attempts to prove that they're smarter than security programmers.

Periodically, all the recently released patches are combined into a *service pack.* Although the most diligent network administrators apply all patches as they're released, many administrators just wait for the service packs.

For all versions of Windows, you can use Windows Update to apply patches to keep your operating system and other Microsoft software up to date. You can find Windows Update in the Start menu. Windows Update automatically scans your computer's software and creates a list of software patches and other components that you can download and install. You can also configure Windows Update to automatically notify you of updates so that you don't have to remember to check for new patches.

Discovering Software Tools for Network Administrators

Network administrators need certain tools to get their jobs done. Administrators of big, complicated, and expensive networks need big, complicated, and expensive tools. Administrators of small networks need small tools.

Some of the tools that the administrator needs are hardware tools, such as screwdrivers, cable crimpers, and hammers. The tools that I'm talking about here, however, are software tools. Here's a sampling of the tools you'll need:

- **A diagramming tool:** A diagramming tool lets you draw pictures of your network. Visio (from Microsoft) is great for drawing the types of diagrams you'll want to make as a network administrator.
- **A network discovery program:** For larger networks, you may want to invest in a network discovery program such as Spiceworks (`www.spiceworks.com`) that can automatically document your network's structure for you. These programs scan the network carefully, looking for computers, printers, routers, and other devices. They then create a database of the network components, draw diagrams for you, and chug out helpful reports.
- **The network's built-in tools:** Many software tools that you need to manage a network come with the network itself. As the network administrator, read through the manuals that come with your network software to see what management tools are available. For example, Windows includes a `net diag` command that you can use to make sure that all the computers on a network can communicate with each other. (You can run `net diag` from an MS-DOS prompt.) For TCP/IP networks, you can use the TCP/IP diagnostic commands summarized in Table 1-1.
- **System Information:** This program that comes with Windows is a useful utility for network managers.
- **A network scanning tool:** These tools scan your network to find devices. The best known is *nmap* (also known as *zenmap*, which is a graphical user interface for nmap). For more information, see `https://nmap.org`.
- **A protocol analyzer:** A *protocol analyzer* monitors and logs the individual packets that travel along your network. (Protocol analyzers are also called *packet sniffers*.) You can configure the protocol analyzer to filter specific types of packets, watch for specific types of problems, and provide statistical analysis of the captured packets. Most network administrators agree that Wireshark (`www.wireshark.org`) is the best protocol analyzer available. And it's free!

- **Network Monitor:** All current versions of Windows include Network Monitor, which provides basic protocol analysis and can often help solve pesky network problems.

TABLE 1-1 TCP/IP Diagnostic Commands

Command	What It Does
`arp`	Displays address resolution information used by the Address Resolution Protocol (ARP)
`hostname`	Displays your computer's host name
`ipconfig`	Displays current TCP/IP settings
`nbtstat`	Displays the status of NetBIOS over TCP/IP connections
`netstat`	Displays statistics for TCP/IP
`nslookup`	Displays Domain Name System (DNS) information
`ping`	Verifies that a specified computer can be reached
`route`	Displays the PC's routing tables
`tracert`	Displays the route from your computer to a specified host

Building a Library

One of Scotty's best lines in the original *Star Trek* series was when he refused to take shore leave so he could get caught up on his technical journals. "Don't you ever relax?" asked Kirk. "I am relaxing!" Scotty replied.

To be a good network administrator, you need to read computer books. Lots of them. And you need to enjoy doing it. If you're the type who takes computer books with you to the beach, you'll make a great network administrator.

You need books on a variety of topics. I'm not going to recommend specific titles, but I do recommend that you get a good, comprehensive book on each of the following topics:

- Network security and hacking
- Wireless networking
- Network cabling and hardware

- Ethernet
- Windows Server, including books on every version that is running in your environment
- Desktop Windows, again including books on every version in your environment
- Linux
- TCP/IP
- DNS
- Microsoft 365/Office 365 and Microsoft Azure administration

In addition to books, you may also want to subscribe to some magazines to keep up with what's happening in the networking industry. Here are a few you should probably consider, along with their web addresses:

- ***InformationWeek:*** `www.informationweek.com`
- ***InfoWorld:*** `www.infoworld.com`
- ***Network Computing:*** `www.networkcomputing.com`
- ***Network World:*** `www.networkworld.com`
- ***2600 The Hacker Quarterly*** **(a great magazine on computer hacking and security):** `www.2600.com`



TIP

The internet is one of the best sources of technical information for network administrators. You'll want to stock your browser's Favorites menu with plenty of websites that contain useful networking information. In addition, you may want to subscribe to one of the many online newsletters that deliver fresh information on a regular basis via email.

Getting Certified

Remember the scene near the end of *The Wizard of Oz* when the Wizard grants the Scarecrow a diploma, the Cowardly Lion a medal, and the Tin Man a testimonial?

Network certifications are kind of like that. I can picture the scene now:

> *The Wizard:* "And as for you, my network-burdened friend, any geek with thick glasses can administer a network. Back where I come from, there are people

who do nothing but configure Cisco routers all day long. And they don't have any more brains than you do. But they do have one thing you don't have: certification. And so, by the authority vested in me by the Universita Committeeatum E Pluribus Unum, I hereby confer upon you the coveted certification of CND."

You: "CND?"

The Wizard: "Yes, that's, uh, *Certified Network Dummy*."

You: "The Seven Layers of the OSI Reference Model are equal to the Sum of the Layers on the Opposite Side. Oh, joy, rapture! I feel like a network administrator already!"

My point is that certification in and of itself doesn't guarantee that you really know how to administer a network. That ability comes from real-world experience — not exam crams.

Nevertheless, certification is becoming increasingly important in today's competitive job market. So you may want to pursue certification, not just to improve your skills, but also to improve your résumé. Certification is an expensive proposition. Each test can cost several hundred dollars, and depending on your technical skills, you may need to buy books to study or enroll in training courses before you take the tests.

You can pursue two basic types of certification: vendor-specific certification and vendor-neutral certification. The major software vendors such as Microsoft and Cisco provide certification programs for their own equipment and software. CompTIA, a nonprofit industry trade association, provides the best-known vendor-neutral certification.

The following sections describe some of the certifications offered by CompTIA, Microsoft, and Cisco.

CompTIA

CompTIA (`http://certifications.comptia.org`) is a source for a variety of well-respected network certifications, including the following:

- **A+** is a basic certification for an entry-level computer technician. To attain A+ certification, you have to pass two exams: one on computer hardware, the other on operating systems.
- **Linux+** covers basic Linux skills such as installation, operations, and troubleshooting. This certification is vendor neutral, so it doesn't depend on any particular version of Linux.

- **Network+** is a popular vendor-neutral networking certification. It covers four major topic areas: Media and Topologies, Protocols and Standards, Network Implementation, and Network Support.
- **Server+** covers network server hardware. It includes details such as installing and upgrading server hardware, installing and configuring an NOS, and so on.
- **Cloud+** covers building cloud infrastructure.
- **Security+** is for security specialists. The exam topics include general security concepts, communication security, infrastructure security, basics of cryptography, and operational/organizational security.

Microsoft

Microsoft (`www.microsoft.com/certifications`) offers dozens of certifications that cover all aspects of Microsoft software offerings. The solutions are *role based*, targeting nine types of IT roles:

- Administrator
- AI engineer
- Data engineer
- Data scientist
- Developer
- DevOps engineer
- Functional consultant
- Security engineer
- Solutions architect

Within each of these roles, a variety of certifications are available. Most fall into one of three categories: Fundamentals, Associate, and Expert. For example, a coveted certification within the Administrator role is called Microsoft 365 Certified: Enterprise Administrator Expert.

Cisco

Cisco offers some of the most respected certifications available for those who need to work extensively with Cisco equipment. You can find out more at

www.cisco.com/web/learning/certifications. Here are some of the more popular Cisco certifications:

- **CCNA** (Cisco Certified Network Associate) is an entry-level apprentice certification. A CCNA should be able to install, configure, and operate Cisco equipment for small networks (under 100 nodes).
- **CCNP** (Cisco Certified Network Professional) is a professional-level certification for Cisco equipment. A CCNP should be able to install, configure, and troubleshoot Cisco networks of virtually any size. Several variants are available for cloud, data center, routing, security, service provider, and routing.
- **CCDE** (Cisco Certified Design Expert) is an expert-level design certification.
- **CCIE** (Cisco Certified Internetwork Expert) is an expert-level certification, which can be had in several varieties, including routing, security, wireless, and service provider.
- **CCAr:** (Cisco Certified Architect) is the top dog of Cisco certifications.
- **CCT** (Cisco Certified Technician) is a certification for those who can diagnose and repair Cisco equipment.
- **And much more!** There are many more Cisco certifications to choose from, including certification for security, voice technology, wireless networking, and more.

Gurus Need Gurus, Too

No matter how much you know about computers, plenty of people know more than you do. This rule seems to apply at every rung of the ladder of computer experience. I'm sure that a top rung exists somewhere, occupied by the world's best computer guru. However, I'm not sitting on that rung, and neither are you. (Not even Bill Gates is sitting on that rung. In fact, Bill Gates got to where he is today by hiring people on higher rungs.)

As the local computer guru, one of your most valuable assets can be a knowledgeable friend who's a notch or two above you on the geek scale. That way, when you run into a real stumper, you have a friend to call for advice. Here are some tips for handling your own guru:

- In dealing with your own guru, don't forget the Computer Geek's Golden Rule: "Do unto your guru as you would have your own users do unto you." Don't pester your guru with simple stuff that you just haven't spent the time to think

through. If you have thought it through and can't come up with a solution, however, give your guru a call. Most computer experts welcome the opportunity to tackle an unusual computer problem. It's a genetic defect.

- If you don't already know someone who knows more about computers than you do, consider joining your local PC users' group. The group may even have a subgroup that specializes in your networking software or may be devoted entirely to local folks who use the same networking software that you use. Odds are good that you're sure to make a friend or two at a users' group meeting. Also, you can probably convince your boss to pay any fees required to join the group.
- If you can't find a real-life guru, try to find an online guru. Check out the various computing newsgroups on the internet. Subscribe to online newsletters that are automatically delivered to you via email.

Helpful Bluffs and Excuses

As network administrator, you just won't be able to solve a problem sometimes, at least not immediately. You can do two things in this situation. The first is to explain that the problem is particularly difficult and that you'll have a solution as soon as possible. The second solution is to look the user in the eyes and, with a straight face, try one of these phony explanations:

- Blame it on the version of whatever software you're using. "Oh, they fixed that with version 8.67.3.509."
- Blame it on cheap, imported memory chips.
- Blame it on Democrats. Or Republicans. Doesn't matter.
- Blame it on oil company executives.
- Blame it on global warming.
- Hope that the problem wasn't caused by stray static electricity. Those types of problems are very difficult to track down. Tell your users that not properly discharging themselves before using their computers can cause all kinds of problems.
- You need more memory.
- You need a bigger hard drive.
- You need a faster processor.

- Blame it on Jar Jar Binks.
- You can't do that in Windows 11.
- You can only do that in Windows 11.
- Could be a virus.
- Or sunspots.
- No beer and no TV make Homer something something something. . . .

IN THIS CHAPTER

- Enabling the Remote Desktop feature
- Connecting to a computer remotely
- Using keyboard shortcuts in a Remote Desktop session
- Setting options for Remote Desktop
- Enabling the Remote Assistance feature
- Helping someone remotely
- Working in the Remote Assistance window

Chapter 2
Managing Remotely

One of the first things you'll realize when you become a network administrator is that you quickly tire of walking all over the building to get to your servers or to users' computers. As a result, you'll soon discover that one of your best secret weapons is the ability to manage computers remotely.

Windows provides two distinct built-in ways to do this:

- **Remote Desktop Connection (RDC)** is designed to let you log into a Windows computer from a remote location. This is useful for managing a server without having to actually go to the computer room and accessing the server's console, or for accessing a virtual server that has no physical console.
- **Remote Assistance** is used to connect to an existing session of another Windows computer so that you can provide technical support for the other user.

This chapter shows you how to use both of these features.

Enabling Remote Desktop Connection

Before you can use RDC to access a server, you must enable remote access on the server. To do that, follow these steps (on the server computer, not your desktop computer):

1. **Open Server Manager.**

   The Server Manager, shown in Figure 2-1, appears. (Note that the Server Manager usually opens automatically when you log in to a Windows server, so you can skip this step if Server Manager is already running.)

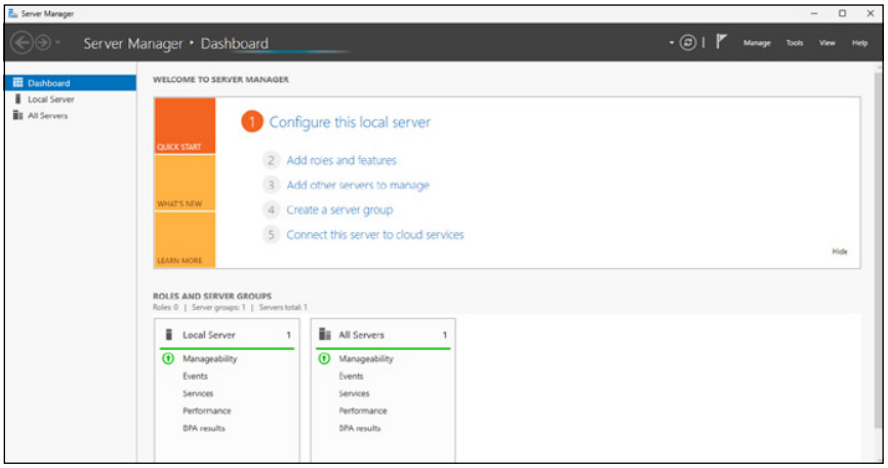


FIGURE 2-1: Using Server Manager to enable Remote Desktop Connection.

2. **Click Local Server in the pane on the left side of the Server Manager window.**

   The Local Server properties page, shown in Figure 2-2, appears.

3. **Locate the Remote Desktop setting on the Local Server Properties page and click Unknown.**

   Note that instead of Unknown, the Remote Desktop setting may indicate Disabled. Either way, clicking this setting will summon the System Properties page with the Remote tab opened, as shown in Figure 2-3.

4. **Select the Allow Remote Connections to This Computer radio button and click OK.**

You're done! Repeat this procedure for each server computer you want to allow access to.

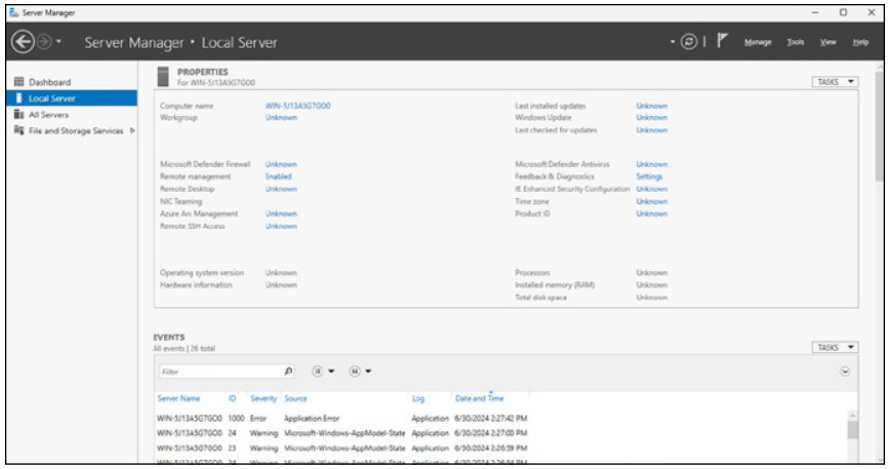


FIGURE 2-2: The Local Server properties page.

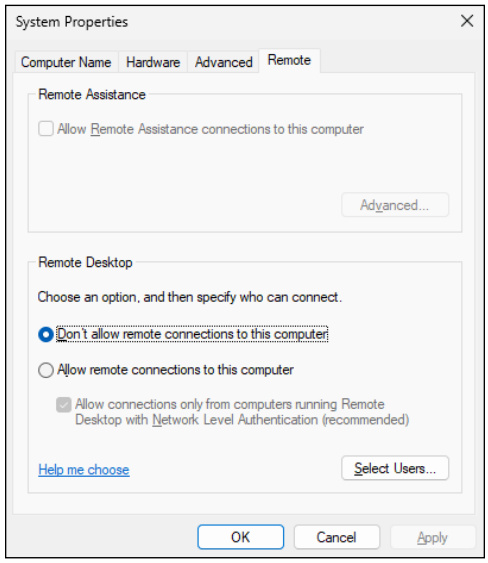


FIGURE 2-3: Enabling Remote Desktop Connection.

You can click the Select Users button to create a list of users who are authorized to access the computer remotely. Note that all members of the Administrators group are automatically granted access, so you don't have to add administrators to this list.



WARNING

There's no question that RDC is convenient and useful. It's also inherently dangerous, however. Don't enable it unless you've taken precautions to secure your Administrator accounts by using strong passwords. Also, you should already have a firewall installed to keep unwanted visitors out of your network.

Connecting Remotely

After you've enabled remote access on a server, you can connect to the server by using the Remote Desktop Client that's automatically installed with Windows. Here's the procedure:

1. **Click the Start button and type the word Remote. Then click the Remote Desktop Connection icon.**

   The Remote Desktop Connection client comes to life, as shown in Figure 2-4.

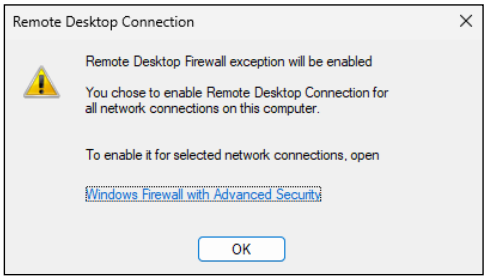


FIGURE 2-4: Connecting with Remote Desktop Connection.

2. **Enter the name of the computer you want to connect to.**

   Alternatively, you can use the drop-down list to select the computer from the list of available computers.

   You can also enter the IP address of the computer you want to connect to.

3. **Click the Connect button.**

   You're connected to the computer you selected, and then prompted for login credentials, as shown in Figure 2-5.

4. **Enter your username and password, and then click OK.**

   Assuming you enter valid credentials, the desktop of the remote computer is displayed, as shown in Figure 2-6.

5. **Use the remote computer!**

You may notice in Figure 2-6 that the Remote Desktop window is not large enough to display the entire desktop of the remote computer. As a result, scroll bars appear to allow you to scroll the desktop horizontally or vertically. You can always maximize the Remote Desktop window to see the entire desktop of the remote computer.

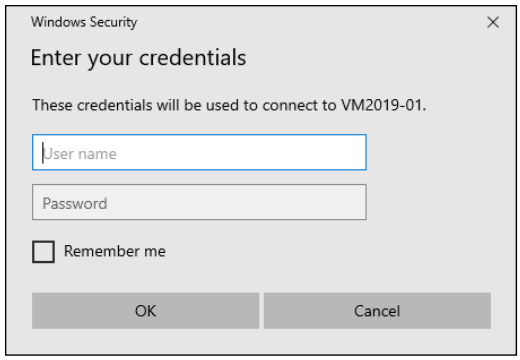


FIGURE 2-5: Logging in to a remote computer.

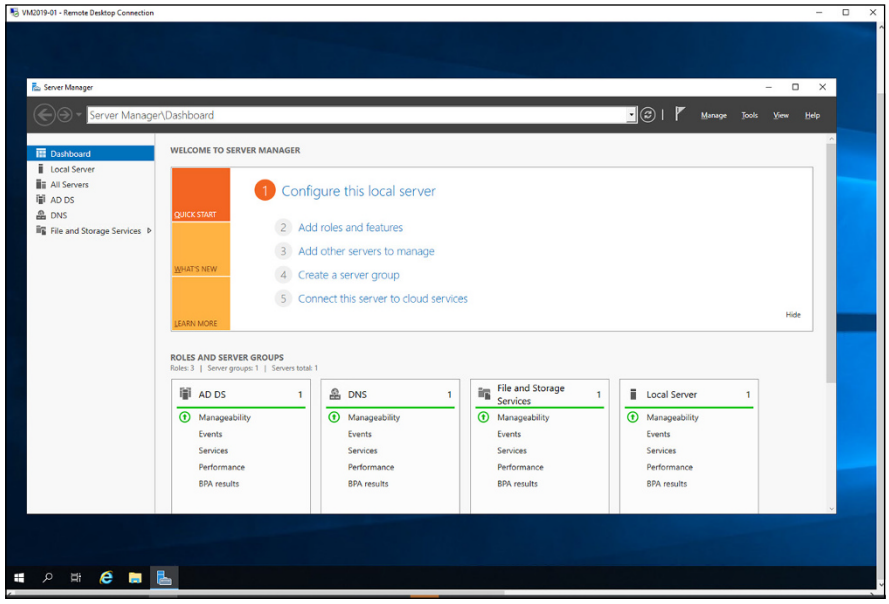


FIGURE 2-6: The desktop of the remote computer.



TIP

Here are a few other tips for working with Remote Desktop Connection:

- **Remote Desktop allows only one user at a time to log in to the remote computer.** If another user is remotely logged in when you try to connect, you'll get a notice informing you that another user is already logged on. You can either cancel or attempt to barge in on the other user's remote session. If you choose the latter option, the other user will see a message stating that you want to connect. If the other user accepts your request, the other user is logged off and you're logged in. If the other user denies your request, your attempt to log in is canceled. If the user does not respond, eventually Windows will kick the other user off and let you in.

- **When you're using the Remote Desktop Connection client, you can't just press Alt+Tab to get another program running on the client computer.** Instead, you must first minimize the RDC client's window by clicking its minimize button. Then you can access other programs running on your computer.
- **If you minimize the RDC client window, you have to provide your logon credentials again when you return.** This security feature is there in case you forget that you have an RDC session open.


TIP

If you use RDC a lot on a particular computer (such as your own desktop computer), I suggest that you create a shortcut to RDC and place it on the desktop, at the top of the Start menu, or in the Quick Launch portion of the taskbar.

- **RDC has several useful configuration options that you can access by clicking the Options button.**

  For more information about these options, refer to the section "Configuring Remote Desktop Options" later in this chapter.

Using Keyboard Shortcuts for Remote Desktop

When you're working in a Remote Desktop session, some of the standard Windows keyboard shortcuts don't work exactly as you expect them to. Table 2-1 lists the special keyboard shortcuts you can use in a Remote Desktop session.

TABLE 2-1 Keyboard Shortcuts for Remote Desktop

Shortcut	What It Does
Ctrl+Alt+Break	Toggles between full-screen and windowed views.
Ctrl+Alt+Pause	Similar to Ctrl+Alt+Break, but instead of maximizing the remote window to full screen, it displays the remote window against a black background.
Alt+Insert	Cycles between applications running on the remote desktop, the same as Alt+Tab on your local machine.
Alt+PageUp	Same as Alt+Insert.
Alt+PageDown	Similar to Alt+Insert, but reverses the order in which applications are cycled. This is the same as Alt+Shift+Tab on your local machine.

Shortcut	What It Does
Ctrl+Alt+End	Sends a Ctrl+Alt+Del to the remote desktop.
Alt+Home	Brings up the Start menu on the remote system.
Alt+Delete	Opens the Windows menu on a window in the remote desktop. (The Windows menu is the one at the top left of every window, with options to move, resize, minimize, maximize, and close the window.)
Ctrl+Alt+Plus Sign (+)	Captures a screen image of the entire remote desktop and saves it to the clipboard. This is the same as pressing Print Screen on your local machine.
Ctrl+Alt+Minus Sign (–)	Captures an image of the current window and saves it to the clipboard. This is the same as pressing Alt+Tab on your local machine.

Configuring Remote Desktop Options

Before you connect to a remote desktop session, you can set a variety of options that affect how the remote desktop session will behave. To summon these options, click the Start button, type the word **Remote**, and then click the Remote Desktop Connection icon. When the Remote Desktop Connection window appears (refer to Figure 2-4), click the Show Options button at the bottom left of the window. This brings up the Remote Desktop options, as shown in Figure 2-7.

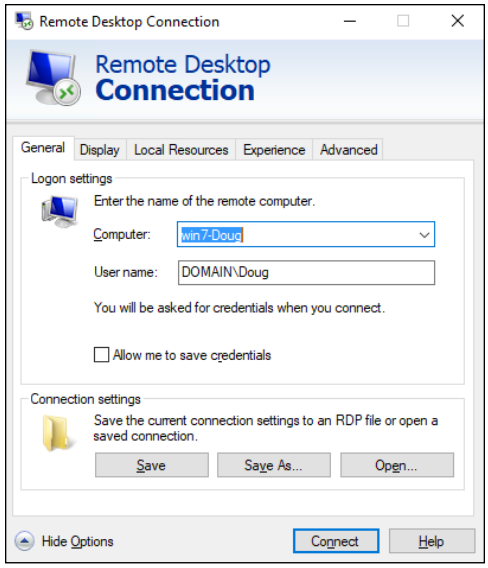


FIGURE 2-7: Setting the options for Remote Desktop Connection.

Figure 2-7 shows the General tab of the Remote Desktop options. On this tab, you can enter the name or IP address of the computer you want to access remotely and the username you want to connect with. You can also save the settings you've created via the other tabs or open a settings file you've previously saved.


TIP

When you click the Save or Save As buttons, your settings are saved to an `.rdp` file. `.rdp` files are associated with the Remote Desktop Connection program, so you can double-click an `.rdp` file to open a saved connection. In other words, an `.rdp` file is a handy shortcut to a remote connection.

The following sections describe the additional options that are available on the other Remote Desktop Connection tabs.

Setting the Display options

Figure 2-8 shows the Display tab of the Remote Desktop Connection dialog box. The following paragraphs describe each of the options available from this tab:

- **Display Configuration size slider:** Use this control to set the size of your remote desktop display. If you drag the slider all the way to the right, the remote desktop will be displayed in full-screen mode. As you move the slider to the left, various display resolutions will appear. Choose the display size you want the remote computer to open up in.
- **Use All My Monitors for the Remote Session:** If your computer has more than one monitor, select this check box to use all your monitors for the remote session. If you want to open the remote desktop in full-screen mode on just one of your monitors, leave this box deselected.
- **Choose the Color Depth of the Remote Session drop-down list:** Use this drop-down list to choose the color quality of the remote desktop. Over slower connections, reducing the color depth will help performance.
- **Display the Connection Bar When I Use the Full Screen:** The connection bar is a thin bar at the top of the screen that enables you to switch back to normal windowed mode. If you deselected this box, you'll have to remember that you can press Ctrl+Alt+Break to switch between full-screen and windowed mode.

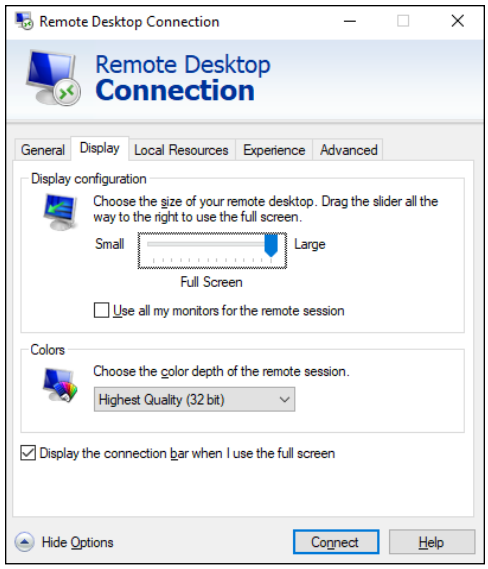


FIGURE 2-8: Setting the Display options for Remote Desktop Connection.

Setting the Local Resources options

The Local Resources tab, shown in Figure 2-9, lets you set the following options:

- **Remote Audio:** Click the Settings button to bring up a dialog box that lets you choose whether to play audio on the remote computer using the remote computer's sound card, the local computer's sound card, or not at all. This display box also lets you choose whether to allow audio recording during the remote session.
- **Apply Windows Key Combinations drop-down list:** This drop-down list lets you specify how Windows keyboard shortcuts are to be interpreted — on the remote computer, on the local computer, or on the remote computer only when it's running in full-screen mode.
- **Local Devices and Resources:** The Printers check box lets you access local printers from the remote session. The Clipboard check box synchronizes the local and remote clipboards, so that when you copy something to the clipboard in the remote session, you can return to the local session and paste the clipboard contents (or vice versa).

  You can also share local drives with the remote session. To do so, click More, and then select the drives that you want to make available to the remote session, as shown in Figure 2-10.

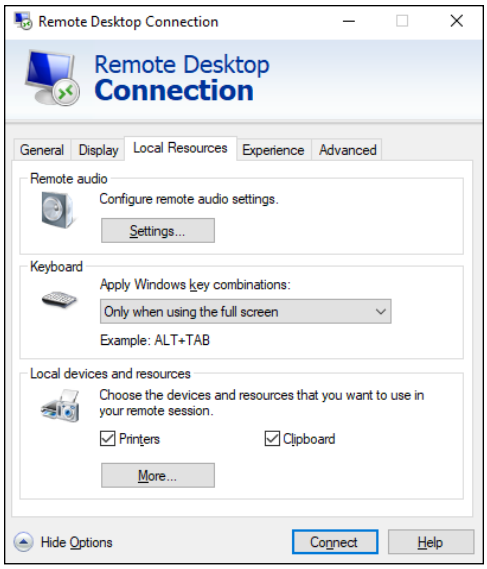


FIGURE 2-9: Setting the Local Resources options for Remote Desktop Connection.

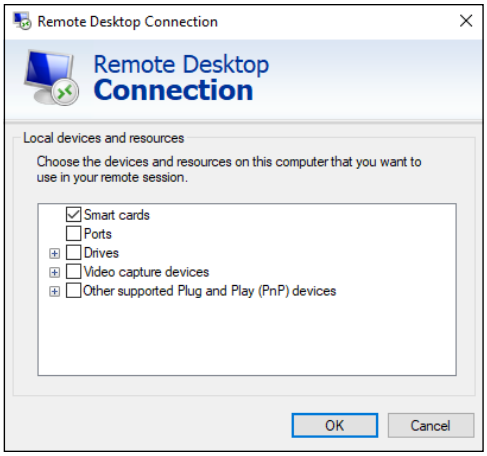


FIGURE 2-10: Sharing drives with the remote computer.

Setting the Experience options

The options on the Experience tab, shown in Figure 2-11, control various settings that affect the responsiveness of your remote connection. The options are as follows:

- **Choose Your Connection Speed to Optimize Performance drop-down list:** This allows you to optimize the amount of information sent back and forth over the network based on your expected connection speed. At slower speeds,

features such as the desktop background, font smoothing, window animations, and so on, will be suppressed. The default setting is to let Windows choose which features to use based on the actual speed of the connection.

- **Persistent Bitmap Caching:** If you select this box, copies of bitmap images are stored on the local computer so they don't have to be transferred across the network every time they're needed.
- **Reconnect If the Connection Is Dropped:** If you select this box, the connection will be automatically reestablished if the connection is broken.

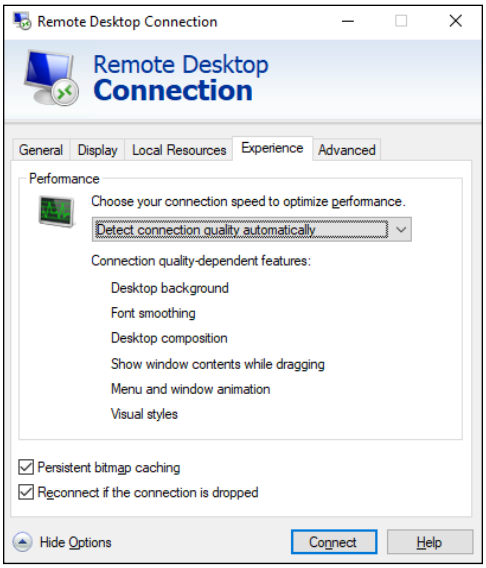


FIGURE 2-11: Setting the Experience options for a Remote Desktop session.

Setting the Advanced options

The Advanced tab of the Remote Desktop Connections window, shown in Figure 2-12, lets you control two features:

- **Server Authentication:** Determines what to do if an authentication problem such as an unknown security certificate is encountered when connecting to the server. The default action is to warn the user, but allow the user to continue if desired. You can change this setting to always connect in spite of the authentication problem, or to never connect when a problem is encountered.

» **Connect from Anywhere:** These settings are used only when you use an advanced server role called Remote Desktop Gateway to manage remote access to computers on your network. For more information about this server role, search the web for *Remote Desktop Gateway*. (This subject is beyond the scope of this book.)

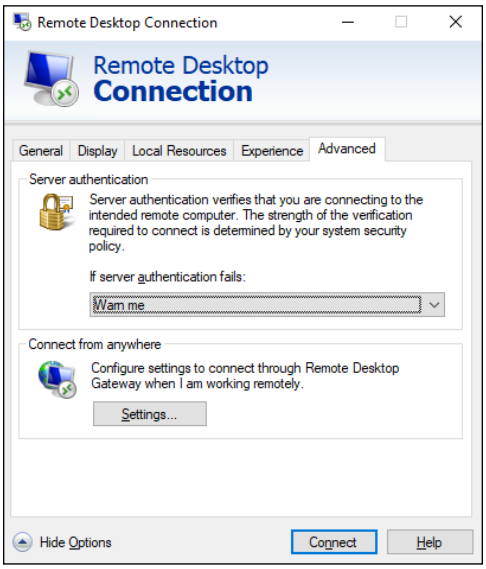


FIGURE 2-12: Setting the Advanced options for a Remote Desktop session.

Using Remote Assistance

One of the most annoying aspects of providing technical support for network users is that you often have to go to the user's desk to see what's going on with his or her computer. That's annoying enough if the other user's desk is across the room or down the hall, but it's almost unworkable if the user you need to support is across town or in a different city or state altogether.

Fortunately, Windows includes a handy feature called Remote Assistance, which is designed to let you provide technical support to an end user without going to the user's location. With Remote Assistance, you can see the user's screen in a window on your own screen, so you can watch what the user is doing. You can even take control when necessary to perform troubleshooting or corrective actions to help solve the user's problems.

Note that there are commercial alternatives to Windows Remote Assistance that do a much better job at this task. This chapter shows you how to use Remote Assistance because it's free and all Windows computers since Windows XP have it.

Enabling Remote Assistance

Before you can lend assistance to a remote computer, Remote Assistance must be enabled on that computer. It is best to enable Remote Assistance before you need it, so that when the time comes, you can easily access your users' computers. But the procedure is simple enough that you can probably walk a user through the steps over the phone so that you can then gain access.

Here are the steps:

1. **Click the Start button, type Remote Assistance, and click Allow Remote Assistance Invitations to Be Sent from This Computer.**

   This brings up the System Properties dialog box, shown in Figure 2-13.

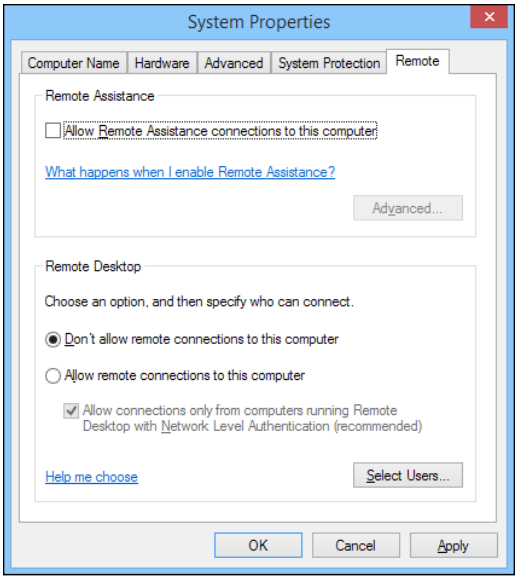


FIGURE 2-13: Enabling Remote Assistance.

2. **Select the Allow Remote Connections to This Computer option.**
3. **Click the Advanced button.**

   The Remote Assistance Settings dialog box appears, as shown in Figure 2-14.

4. **Select the Allow This Computer to Be Controlled Remotely check box.**

   This option will allow you to later take control of this computer remotely.

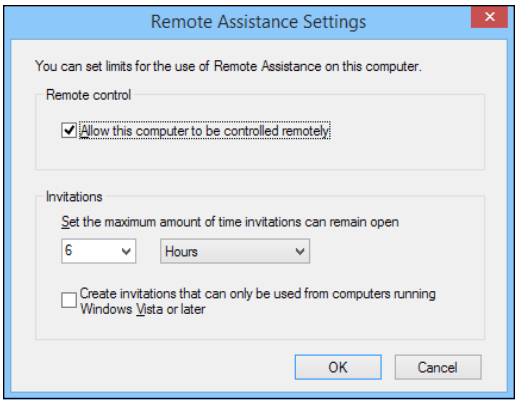


FIGURE 2-14: Setting the advanced Remote Assistance options.

5. **Click OK.**

   You're returned to the System Properties dialog box.

6. **Click OK.**

You're done. You can now initiate Remote Assistance sessions from this computer.

Inviting Someone to Help You via a Remote Assistance Session

The user requesting remote assistance must initiate the request before you can connect to the user's computer to provide help. You may need to guide your user through this procedure over the phone. Here are the steps:

1. **Click the Start button, type Invite, and then click** Invite **Someone to Connect to Your PC.**

   The Windows Remote Assistance window appears, as shown in Figure 2-15.

2. **Click Invite Someone You Trust to Help You.**

   The screen shown in Figure 2-16 appears.

3. **Click Save This Invitation as a File, and save the invitation file to a convenient disk location.**

   This option creates a special file that you can save to disk. A Save As dialog box appears, allowing you to save the invitation file in any disk location you want. You can then email the invitation to your helper, who can use the invitation to connect to your PC.

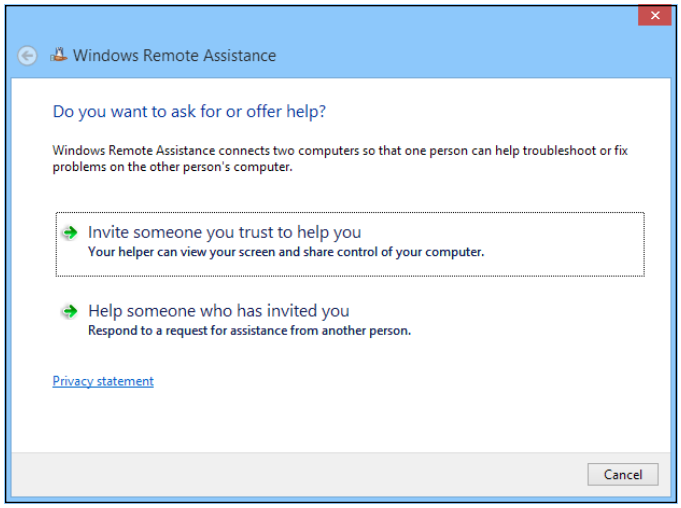


FIGURE 2-15: The Windows Remote Assistance window.

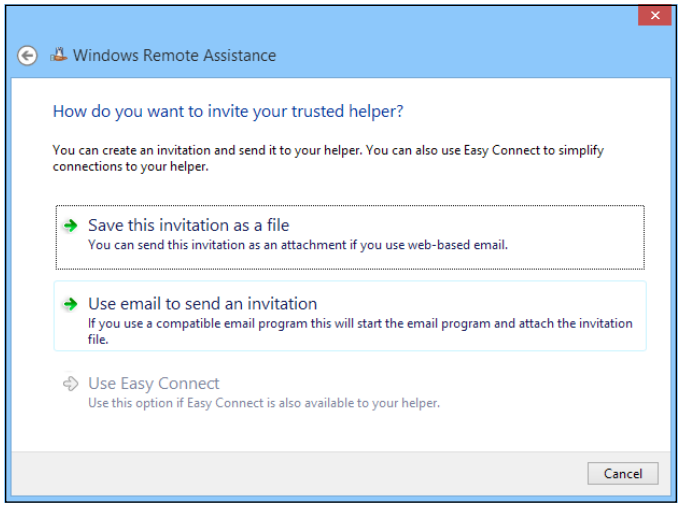


FIGURE 2-16: Inviting someone to help you.

4. **If you use Microsoft Outlook, you can alternatively choose Use Email to Send an Invitation. This option fires up Outlook and creates an email with the invitation file already attached.**

   Either way, a password will be displayed, as shown in Figure 2-17. You'll need to provide this password to your helper when requested. (Typically, you'll do that over the phone.)

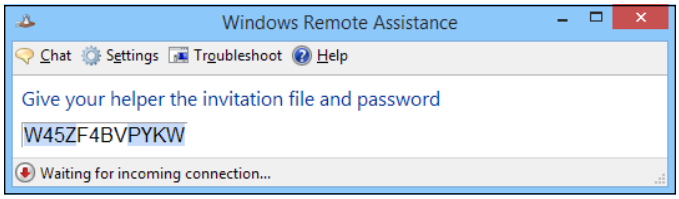


FIGURE 2-17: You'll need to tell your helper the password.

5. **Email the invitation file to the person you want to help you.**

   Use your preferred email program to send the invitation file as an attachment.

6. **Wait for your helper to request the password.**

   When your helper enters the password into his or her Remote Assistance screen, the Remote Assistance session is established. You'll be prompted for permission to allow your helper to take control of your computer, as shown in Figure 2-18.

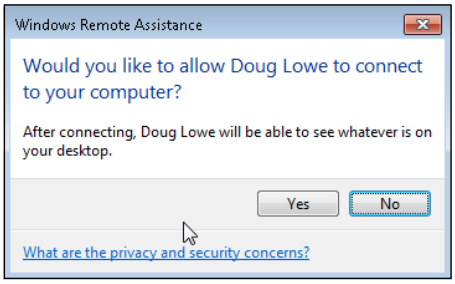


FIGURE 2-18: Granting your helper permission to take control.

7. **Click Yes.**

   Your helper now has access to your computer. To facilitate the Remote Assistance session, the toolbar shown in Figure 2-19 appears. You can use this toolbar to chat with your helper or to temporarily pause the Remote Assistance session.

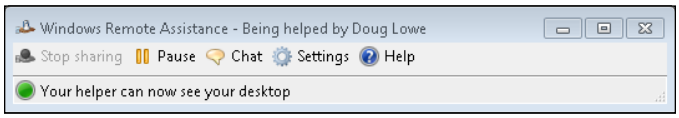


FIGURE 2-19: The Windows Remote Assistance toolbar.

Responding to a Remote Assistance Invitation

If you've received an invitation to a Remote Assistance session, you can establish the session by following these steps:

1. **Click the Start button, type Invite, and then click Invite Someone to Connect to Your PC.**

   This brings up the Windows Remote Assistance window (refer to Figure 2-16).

2. **Click Help Someone Who Has Invited You.**

3. **Click Use an Invitation File.**

   An Open dialog box appears.

4. **Locate the invitation file you were sent, select it, and click Open.**

   You're prompted to enter the Remote Assistance password, as shown in Figure 2-20.

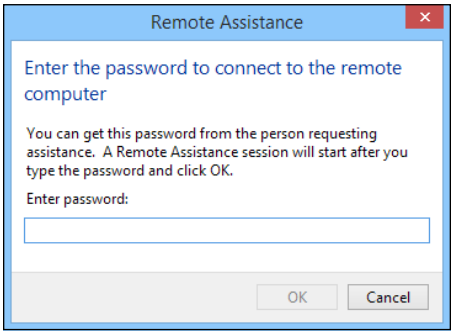


FIGURE 2-20: Enter the Remote Assistance password.



TIP

As an alternative to Steps 1 through 3, you can simply double-click the invitation file you received. This will launch Windows Remote Assistance and prompt you for the password.

5. **Enter the password given to you by the user requesting help, and then click OK.**

   The remote user is prompted to grant you permission to start the Remote Assistance session (this is where the remote user sees the screen that was shown in Figure 2-19). When the user grants permission, the Remote Assistance session is established. You can now see the user's screen in the Remote Assistance window, as shown in Figure 2-21.

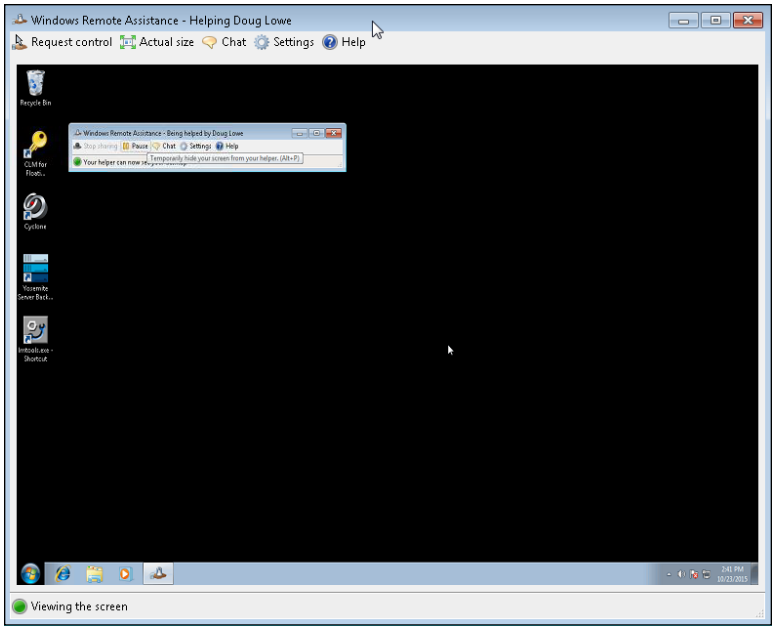


FIGURE 2-21: A Remote Assistance session in progress.

6. **To take control of the remote user's computer, click Request Control.**

   The remote user will be prompted to allow control. Assuming that permission is granted, you can now control the other computer.

7. **Do your thing.**

   Now that you're connected to the remote computer, you can perform whatever troubleshooting or corrective actions are necessary to solve the user's problems.

8. **If necessary, use the Chat window to communicate with the user.**

   You can summon the Chat window by clicking the Chat button in the toolbar. Figure 2-22 shows a chat in progress.

9. **To conclude the Remote Assistance session, simply close the Remote Assistance window.**

   The remote user is notified that the Remote Assistance session has ended.

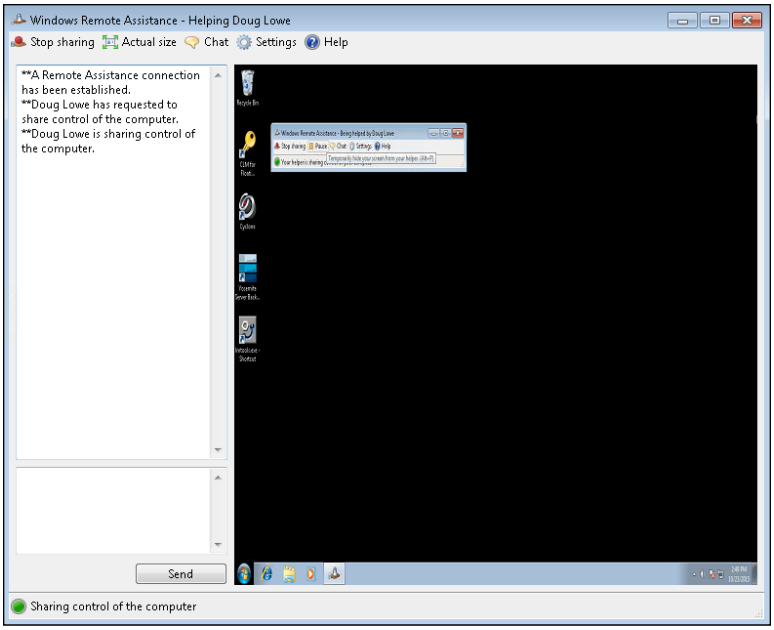


FIGURE 2-22: Using the chat window.

IN THIS CHAPTER

- Understanding why you should keep track of your network assets
- Identifying what kind of information you should track
- Using a spreadsheet to track assets
- Using specialized asset-tracking software

Chapter 3
Managing Network Assets

In a small network, keeping track of the computers that live on the network is easy. For example, in a small office with four or five employees, you can just walk around the room and make sure that everyone's computer is still there. If one of the computers is missing, you know you've lost one!

But in a larger network, keeping track of the assets on the network can become a major burden. With dozens, hundreds, or even thousands of devices to keep track of, you need some kind of organized system to keep track of everything. Otherwise, you'll soon find yourself asking questions like these, and struggling to come up with answers:

- When Bob in Accounting left the company two years ago, did we get his laptop back? If so, where is it?
- When we hired Paige last summer, did we set her up with a MacBook or a Dell laptop? Or did we give her Bob's laptop?
- What kind of printer did we set up in the Redmond office?
- When Richard became a remote worker, did we give him a mobile hotspot?

- Analisa's car was broken into last night and her Surface Pro was stolen. Was it the 8GB i5 model or the 32GB i7 model?
- Juan Carlos is going to be out of town for a conference next week. Do we have a laptop he can borrow?

This chapter gives you some suggestions on what to do when it becomes impossible to keep all this stuff in your head.

Introducing IT Asset Management

IT asset management (ITAM) is a program that assumes the responsibility for managing the life cycle of all IT assets within an organization. IT assets include *hardware assets* (such as computers, tablets, and mobile phones) and *software assets* (such as licenses of Office 365, Adobe Acrobat Standard, and Windows Server 2019).

Ideally, an ITAM system should track all of an organization's IT assets through each asset's entire life cycle, from acquisition to retirement. For a hardware asset such as a computer, the life cycle includes the following:

- **Acquisition:** Consideration of alternative proposals for the asset, planned usage of the asset, development of new or application of existing policies related to the asset, and the actual purchase of the asset
- **Deployment:** Installation and configuration of the asset for use by a specific user
- **Support:** Any necessary reconfiguration, repair, upgrades, or installation of new software that may be done during the asset's useful life
- **Redeployment:** For example, when a desktop computer is redeployed to a different user or application
- **Retirement:** When an asset is taken out of service due to obsolescence or changing business needs
- **Disposition:** When the asset is physically disposed of (e-wasted or sold) and is no longer owned by the organization

Software assets have a similar but distinct characteristic life cycle. The main distinction is that software typically falls under one of several licensing models that permit deployment to a varying number of users and have a predictable maintenance cost that is set by the vendor. For example, a software asset such

as Office 365 is a subscription that is charged per user, per month. Some software, such as AutoDesk's AutoCAD, can be obtained with multiuser licenses that are administered by AutoDesk's cloud-based licensing services and paid annually. Each software vendor provides its own licensing options. The ITAM program should be able to accommodate a variety of licensing models so that you can track your software assets.

Why Bother?

Whew, this sounds like a lot of work. It's true: Maintaining an accurate and detailed record of all your IT assets is a bit of an undertaking. Before you commit to it, you should understand some of the benefits of doing so. Here are some of the top reasons I suggest you take the time:

- To prevent loss due to theft or neglect of equipment
- To reduce cost by not purchasing unnecessary equipment and by repurposing existing equipment for new applications
- To improve performance by retiring and replacing obsolete equipment
- To comply with auditing or regulatory requirements
- To control the cost of maintaining equipment
- To ensure proper recovery of lost or damaged equipment, especially if an insurance claim is necessary
- To prevent the serious security threat that occurs when assets are lost

One of the most important, and often overlooked, benefits is that your users will take much better care of their equipment if they know that *you* care about the equipment enough to keep a record of it.

Getting Organized

The first step for creating a system to track the assets on your network is to get organized. Start by making a list of the various types of IT assets that exist within your organization. This list might include

- Desktop computers
- Portable computers, tablets, smartphones, and Mi-Fi devices

- Printers and scanners
- IP phones
- Servers
- Switches, routers, and wireless access points

After you've compiled this list, you can start to focus on which types of assets will most benefit from tracking in an asset management system. For example, you may decide that parts of the infrastructure (such as servers, switches, and routers) don't need to be tracked. After all, there are probably a limited number of them, and they don't move around much. The other types of devices, on the other hand, move around as often as employees are hired or leave, change jobs, or have new requirements.

After you've identified what types of assets you need to track, you can home in on the specifics. For example, you can make a list of each of your desktop computers, including the make and model of the computer, basic specifications (such as CPU type, RAM, and disk capacity), when it was acquired, and who is using it.

What to Track

When putting together a database for tracking IT assets, you should carefully consider the type of information you want to track. Strive to find the right balance between tracking so much information that the recordkeeping becomes a burden versus tracking so little information that the database is useless.

At the minimum, I recommend you track the following:

- **Asset identifier:** See "Picking a Number" later in this chapter for tips.
- **Type of asset:** Develop a list of asset types and use a standard name for each type of asset (for example, "WK" for "workstation").
- **Manufacturer, product name, model number, and serial number.**
- **Asset's status:** For example, "Deployed," "In Inventory," or "Retired."
- **Date the asset was acquired.**
- **The user currently assigned to the asset.**
- **The date the asset was assigned to the current user.**

You can easily track this information in a simple spreadsheet; just create a column for each data point and start entering your inventory.

But ideally, you may want to track more than just this basic information for each asset. In particular, you may want to track the entire history of the asset. To do that, you'll need detailed records to record each event during the asset's life cycle. These records should include the following:

- Asset identifier
- Event type (for example, "Deployed," "Serviced," "Returned to Inventory," "Retired," and so on)
- Event date
- Name of the user
- Name of the technician
- Description

A spreadsheet won't do for tracking this type of life-cycle information. Instead, you'll need a bona-fide database. If you're smart with Microsoft Access or SQL databases, you can easily design your own database. But you may find it easier in the long run (and ultimately less expensive) to use professionally designed IT asset management software to track this level of detail. For more information, refer to the section "Using Asset-Tracking Software" later in this chapter.

Taking Pictures

Besides all the useful data spelled out in the previous section, consider including a photograph of each asset in your asset database. As they say, a picture is worth a thousand words!

Including a picture can be especially useful if the asset is stolen or damaged and an insurance claim is made. A good photograph coupled with an accurate description can go a long ways toward establishing the legitimacy of a loss claim.

Picking a Number

To keep track of your computer inventory, you'll need to assign each device in your inventory a unique identifier. You can call this identifier an *asset identifier*, *asset number*, *tracking number*, or anything else that makes sense to you. The most important aspect of this identifier is that it is unique. To avoid confusion, every IT asset tracked by your asset management system must have a unique identifier.

If your initial goal is to keep track of Windows computers, you may be tempted to use each computer's Active Directory computer name to uniquely identify your assets. However, that's a shortsighted choice, for two reasons:

- Even if you initially start with Windows computers, you'll eventually realize that you also need to track other types of devices, such as smartphones, tablets, and cellular hotspots. Not all these devices lend themselves to Active Directory naming conventions.
- More important, it's easy to change the name of a Windows computer. If you rename a device, you'll have to remember to rename the device in your asset database.

Instead of using the Active Directory computer name, I suggest you create your own asset ID for every device in your inventory. In keeping with good database design practices, this ID should *not* indicate the type of device. In other words, I recommend against using something like "NB-002" for "Notebook #2" or "18-001 for "the first device purchased in 2018." Instead, use simple sequential identification numbers that have no meaning or purpose other than to uniquely identify each device.

Making Labels

One of the basic steps in keeping track of your computer equipment is to physically apply a label on every piece of equipment maintained in your asset inventory. That way, you'll always be able to correlate a specific piece of equipment with its record in the database.

At the minimum, the label should include the unique asset identifier as well as identifying information for your organization such as the company name and perhaps a phone number, street address, and URL. Any other information is superfluous and will probably make your task more difficult.

Here are a few tips for making labels:

- Don't use a handheld or desktop label maker unless you have only a few assets to track. You'll quickly grow tired of manually keying in the information for each label. Instead, use a computer-attached label printer. Several companies make excellent label printers; check out DYMO (`www.dymo.com`), Brother (`www.brother-usa.com`), and Zebra (`www.zebra.com`).
- Consider printing the asset ID number in bar code format. Then you can use a bar code reader whenever you need to read the label. This can save considerable time when servicing the asset.
- Use tamper-proof label tape when printing labels. It costs a bit more but is worth it.
- Affix the labels in a consistent location for each device type.
- Always affix the label to the device itself, not to a protective case.
- If possible, incorporate your company's branding into the design of the label. Get help from the marketing or publicity department to make sure your labels are consistent with company branding.

DEVELOPING AN EQUIPMENT LOSS POLICY

At the same time you're developing an asset management system for your organization, you may also want to consider developing a policy on how lost or damaged equipment is handled. This is especially true if you intend to hold employees responsible for loss or damage of equipment issued to them.

The policy should spell out exactly what the expectations of the company are regarding the proper care of the equipment, as well as the consequences of losing or damaging the equipment. The policy should specifically address what happens in the following circumstances:

- Theft from a locked vehicle versus an unlocked vehicle
- Misplaced equipment
- Damage as a result of abuse versus accidental damage or ordinary wear and tear

If the policy is to require reimbursement by the employee or a payroll deduction to cover the loss, make sure that the employee is aware of the policy and signs a written agreement at the time the equipment is issued. You should also check with your attorney to make sure that payroll deductions are allowed in your state.

Tracking Software

So far, we've only talked about keeping track of hardware assets. Besides hardware, you'll also need to keep track of the software assets within your organization. Make a thorough list of all the software used by your organization, along with information about the software version, how the software is licensed, how many and what types of seats you own, and who in your organization uses the software. If you don't do this, you'll find yourself unnecessarily purchasing software you already own or installing more copies of a program than your license permits.

Many software products offer their own license management portals you can use to manage their licenses. Examples include the following:

- If you use Microsoft Volume Licensing for your Microsoft software, you can use the Volume Licensing Center (available via `www.microsoft.com/licensing`).
- You can manage Office 365 from the Office 365 admin center. Just log in at `www.office.com`, and then choose Admin from the app launcher. (You must have admin permissions, of course.)
- For Adobe Creative Suite products (such as Acrobat, Photoshop, InDesign, and so on), log in to the team management portal at `http://accounts.adobe.com`. This portal allows you to view all your Adobe subscription licenses and purchase additional seats if needed.

Most other software companies that offer subscription-based software licensing offer their own portals.

Using Asset-Tracking Software

A simple spreadsheet or Sharepoint list may suffice for a very basic asset-tracking system, but eventually you'll outgrow the spreadsheet or list. When that happens, you'll need to move to a bona-fide database system. If you're a skilled database developer, you can create a system on your own. But why spend so much time inventing something that many software vendors have already developed? A simple web search will reveal that there are plenty of IT asset-tracking systems available. Some are low-cost or even free and provide just the basics. Others are capable of tracking inventory for the largest IT systems. The trick, as with any software, is to find the product that meets your current needs, can grow with you as your organization grows, and fits your budget.

When searching for asset management software, here are a few features I suggest you look for:

- **Simple installation and management:** Asset management is an ideal application for a cloud-based service rather than an on-premises install.
- **Secure access:** If you do go with a cloud-based option, it's imperative that the data be secure.
- **Web-based interface:** That way, a client install is not required. It will be very helpful to access your asset management software from any computer on your network so that you don't have to return to your office to look up or update an asset's record.
- **Mobile app support:** This allows you to easily scan a bar code label with your smartphone or tablet to call up an asset's record.
- **Customizable fields:** This allows you to set up the tracking records to meet your own unique needs.
- **Customizable reports:** These let you meet your organization's reporting standards.
- **Import capabilities:** This allows you to convert your current asset-tracking spreadsheets to the new system. Spreadsheet imports can also sometimes help with bulk entry of data; it's often easier to enter data in spreadsheet format than it is to work through a multitude of data-entry screens.
- **Pricing that scales with volume:** That way, you don't pay for capacity you don't need.
- **Responsive support and good training resources.**

Other Sources of Asset-Tracking Information

In addition to dedicated asset-tracking service, there are several other places you can go for information that can be helpful for tracking IT assets. Here are a few:

- **Cellphone vendor portals:** All major mobile phone providers have online portals for managing your accounts. These portals typically let you view and edit all the phone lines on your account, with details including who the phone is assigned to, how much usage the phone has incurred, the exact make and

model of the phone, and when the phone is eligible for upgrade. If a phone is lost or stolen, you can disable or deactivate the device.

These vendor portals also track MiFi hotspot devices as well.

- **Copier vendor portals:** If your organization leases its copiers and has a support plan, your provider probably has a service and support portal you can use to manage your devices. This portal can give you up-to-date information about the make, model, and location of each of your copiers. And you can keep track of maintenance and supplies needed for the copiers.
- **Switch or router management pages:** Modern switches and routers include management pages that can be useful for managing network assets. For example, the management page for a switch can reveal information about the devices that are attached to each of the switch's ports. Some switches provide basic information such as IP address, but others provide detailed information including the computer name for Windows computers connected to the switch.

IN THIS CHAPTER

- Checking the obvious things
- Fixing computers that have expired
- Pinpointing the cause of trouble
- Restarting client and server computers
- Reviewing network event logs
- Keeping a record of network woes

Chapter 4

Solving Network Problems

Face it: Networks are prone to breaking.

They have too many parts. Cables. Connectors. Cards. Switches. Routers. All these parts must be held together in a delicate balance, and the network equilibrium is all too easy to disturb. Even the best-designed computer networks sometimes act as if they're held together with baling wire, chewing gum, and duct tape.

To make matters worse, networks breed suspicion. After your computer is attached to a network, users begin to blame the network every time something goes wrong, regardless of whether the problem has anything to do with the network. You can't get columns to line up in a Word document? Must be the network. Your spreadsheet doesn't add up? The @#$% network's acting up again. The stock market's down? Arghhh!!!!!!

The worst thing about network failures is that sometimes they can shut down an entire company. It's not so bad if just one user can't access a particular shared folder on a file server. If a critical server goes down, however, your network users

may be locked out of their files, applications, email, and everything else they need to conduct business as usual. When that happens, they'll be beating down your doors and won't stop until you get the network back up and running.

In this chapter, I review some of the most likely causes of network trouble and suggest some basic troubleshooting techniques that you can employ when your network goes on the fritz.

When Bad Things Happen to Good Computers

Here are some basic troubleshooting steps explaining what you should examine at the first sign of network trouble. In many (if not most) of the cases, one of the following steps can get your network back up and running:

1. **Make sure that your computer and everything attached to it is plugged in.**


TECHNICAL STUFF

Computer geeks love it when a user calls for help, and they get to tell the user that the computer isn't plugged in or that its power strip is turned off. They write it down in their geek logs so that they can tell their geek friends about it later. They may even want to take your picture so that they can show it to their geek friends. (Most "accidents" involving computer geeks are a direct result of this kind of behavior. So try to be tactful when you ask a user whether he or she is sure the computer is actually turned on.)

2. **Make sure that your computer is properly connected to the network.**
3. **Note any error messages that appear on the screen.**
4. **Try restarting the computer.**


TIP

An amazing number of computer problems are cleared up by a simple restart of the computer, at least on Windows computers. Of course, in many cases, the problem recurs, so you'll have to eventually isolate the cause and fix the problem. But many problems are only intermittent, and a simple reboot is all that's needed.

5. **Try the built-in Windows network troubleshooter.**
6. **Check the free disk space on your computer and on the server.**

When a computer runs out of disk space or comes close to it, strange things can happen. Sometimes you get a clear error message indicating such a

situation, but not always. Sometimes the computer just grinds to a halt; operations that used to take a few seconds now take a few minutes.

7. **Do a little experimenting to find out whether the problem is indeed a network problem or just a problem with the computer itself.**

   See the section "Time to Experiment," later in this chapter, for some simple things that you can do to isolate a network problem.

8. **Try restarting the network server.**

   See the section "Restarting a Network Server," later in this chapter.

Fixing Dead Computers

If a computer seems totally dead, here are some things to check:

- **Make sure that the computer is plugged in.**
- **If the computer is plugged into a surge protector or a power strip, make sure that the surge protector or power strip is plugged in and turned on.** If the surge protector or power strip has a light, it should be glowing. Also, the surge protector may have a reset button that needs to be pressed.
- **Make sure that the computer's On/Off switch is turned on.** This advice sounds too basic to even include here, but many computers have two power switches: an on/off switch on the back of the computer, and a push-button on the front that actually starts the computer. If you push the front button and nothing happens, check the switch on the back to make sure it's in the ON position.

  
  REMEMBER

  To complicate matters, newer computers have a Sleep feature, in which they appear to be turned off but really they're just sleeping. All you have to do to wake such a computer is jiggle the mouse a little. (I used to have an uncle like that.) It's easy to assume that the computer is turned off, press the power button, wonder why nothing happened, and then press the power button and hold it down, hoping it will take. If you hold down the power button long enough, the computer will actually turn itself off. Then, when you turn the computer back on, you get a message saying the computer wasn't shut down properly. Arghhh! The moral of the story is to jiggle the mouse if the computer seems to have nodded off.
- **If you think the computer isn't plugged in but it looks like it is, listen for the fan.** If the fan is running, the computer is getting power, and the problem is more serious than an unplugged power cord. (If the fan isn't running but the computer is plugged in and the power is on, the fan may be out to lunch.)

- **If the computer is plugged in and turned on but still not running, plug a lamp into the outlet to make sure that power is getting to the outlet.** You may need to reset a tripped circuit breaker or replace a bad surge protector. Or you may need to call the power company. (If you live in California, don't bother calling the power company. It probably won't do any good.)
- **Check the surge protector.** Surge protectors have a limited life span. After a few years of use, many surge protectors continue to provide electrical power for your computer, but the components that protect your computer from power surges no longer work. If you're using a surge protector that is more than two or three years old, replace the old surge protector with a new one.
- **Make sure that the monitor is plugged in and turned on.** The monitor has a separate power cord and switch. (The monitor actually has two cables that must be plugged in. One runs from the back of the monitor to the back of the computer; the other is a power cord that comes from the back of the monitor and must be plugged into an electrical outlet.)
- **Make sure that all cables are plugged in securely.** Your keyboard, monitor, mouse, and printer are all connected to the back of your computer by cables.

  Make sure that the other ends of the monitor and printer cables are plugged in properly, too.


REMEMBER

- **If the computer is running but the display is dark, try adjusting the monitor's contrast and brightness.** Some monitors have knobs that you can use to adjust the contrast and brightness of the monitor's display. They may have been turned down all the way.

Ways to Check a Network Connection

The cables that connect client computers to the rest of the network are finicky beasts. They can break at a moment's notice, and by "break," I don't necessarily mean "to physically break in two." Although some broken cables look like someone got to the cable with pruning shears, most cable problems aren't visible to the naked eye.

- **Twisted-pair cable:** If your network uses twisted-pair cable, you can quickly tell whether the cable connection to the network is good by looking at the back of your computer. Look for a small light located near where the cable plugs in; if this light is glowing steadily, the cable is good. If the light is dark or it's flashing intermittently, you have a cable problem (or a problem with the network card or the hub or switch that the other end of the cable is plugged in to).



TIP

If the light isn't glowing steadily, try removing the cable from your computer and reinserting it. This action may cure the weak connection.

- **Patch cable:** Hopefully, your network is wired so that each computer is connected to the network with a short (six feet or so) patch cable. One end of the patch cable plugs into the computer, and the other end plugs into a cable connector mounted on the wall. Try quickly disconnecting and reconnecting the patch cable. If that doesn't do the trick, try to find a spare patch cable that you can use.
- **Switches:** Switches are prone to having cable problems, too — especially switches that are wired in a less-than-professional manner, featuring a rat's nest of patch cables. Be careful whenever you enter the lair of the rat's nest. If you need to replace a patch cable, be very careful when you disconnect the suspected bad cable and reconnect the good cable in its place.

A Bunch of Error Messages Just Flew By!

Error messages that display when your computer boots can provide invaluable clues to determine the source of the problem.

If you see error messages when you start up the computer, keep the following points in mind:

- **Don't panic if you see a lot of error messages.** Sometimes, a simple problem that's easy to correct can cause a plethora of error messages when you start your computer. The messages may look as if your computer is falling to pieces, but the fix may be very simple.
- **If the messages fly by so fast that you can't see them, press your computer's Pause key.** Your computer comes to a screeching halt, giving you a chance to catch up on your error-message reading. After you've read enough, press the Pause key again to get things moving. (On keyboards that don't have a Pause key, pressing Ctrl+Num Lock or Ctrl+S does the same thing.)



TIP

- **If you miss the error messages the first time, restart the computer and watch them again.**
- **Better yet, press F8 when you see the Starting Windows message.** This displays a menu that allows you to select from several startup options. (Note that this won't work on Windows 8, 8.1, or 10.)

Double-Checking Your Network Settings

I swear that there are little green men who sneak into offices at night, turn on computers, and mess up TCP/IP configuration settings just for kicks. These little green men are affectionately known as *networchons.*

Remarkably, network configuration settings sometimes get inadvertently changed so that a computer, which enjoyed the network for months or even years, one day finds itself unable to access the network. So one of the first things you do, after making sure that the computers are actually on and that the cables aren't broken, is a basic review of the computer's network settings. Check the following:

- At a command prompt, run `ipconfig` to make sure that TCP/IP is up and running on the computer and that the IP addresses, subnet masks, and default gateway settings look right.
- Call up the network connection's Properties dialog box and make sure that the necessary protocols are installed correctly.
- Open the System Properties dialog box (double-click System in Control Panel) and check the Computer Name tab.
- Make sure that the computer name is unique and also that the domain or workgroup name is spelled properly.
- Double-check the user account to make sure that the user really has permission to access the resources he needs.



TIP

For more information about network configuration settings, see Book 4, Chapter 3.

Time to Experiment

If you can't find some obvious explanation for your troubles — like the computer is unplugged — you need to do some experimenting to narrow down the possibilities. Design your experiments to answer one basic question: Is it a network problem or a local computer problem?

Here are some ways you can narrow down the cause of the problem:

- **Try performing the same operation on someone else's computer.** If no one on the network can access a network drive or printer, something is probably wrong with the network. On the other hand, if the error occurs on only one computer, the problem is likely with that computer. The wayward

computer may not be reliably communicating with the network or configured properly for the network, or the problem may have nothing to do with the network at all.

- **If you're able to perform the operation on another computer without problems, try logging on to the network with another computer using your own username.** Then see whether you can perform the operation without error. If you can, the problem is probably on your computer. If you can't, the problem may be with the way your user account is configured.
- **If you can't log on at another computer, try waiting for a bit.** Your account may be temporarily locked out. This can happen for a variety of reasons — the most common of which is trying to log on with the wrong password several times in a row. If you're still locked out an hour later, call the network administrator and offer a doughnut.

Who's on First?

When troubleshooting a networking problem, it's often useful to find out who is actually logged on to a network server. For example, if a user can't access a file on the server, you can check whether the user is logged on. If so, you know that the user's account is valid, but the user may not have permission to access the particular file or folder that he's attempting to access. On the other hand, if the user isn't logged on, the problem may lie with the account itself or how the user is attempting to connect to the server.

It's also useful to find out who's logged on in the event that you need to restart the server. For more information about restarting a server, see the section, "Restarting a Network Server," later in this chapter.

To find out who is currently logged on to a Windows server, right-click the Computer icon on the desktop and choose Manage from the menu that appears. This brings up the Computer Management window. Open System Tools in the tree list and then open Shared Folders and select Sessions. A list of users who are logged on appears.


TIP

You can immediately disconnect all users by right-clicking Sessions in the Computer Management window and choosing All Tasks➪Disconnect All. Be warned, however, that this can cause users to lose data.

Restarting a Client Computer

Sometimes, trouble gets a computer so tied up in knots that the only thing you can do is reboot. In some cases, the computer just starts acting weird. Strange characters appear on the screen, or Windows goes haywire and doesn't let you exit a program. Sometimes, the computer gets so confused that it can't even move. It just sits there, like a deer staring at oncoming headlights. It won't move, no matter how hard you press Esc or Enter. You can move the mouse all over your desktop, or you can even throw it across the room, but the mouse pointer on the screen stays perfectly still.

When a computer starts acting strange, you need to reboot. If you must reboot, you should do so as cleanly as possible. I know this procedure may seem elementary, but the technique for safely restarting a client computer is worth repeating, even if it is basic:

1. **Save your work if you can.**

   Choose File➪Save to save any documents or files that you were editing when things started to go haywire. If you can't use the menus, try clicking the Save button on the toolbar. If that doesn't work, try pressing Ctrl+S (the standard keyboard shortcut for the Save command).

2. **Close any running programs if you can.**

   Choose File➪Exit or click the Close button in the upper-right corner of the program window. Or press Alt+F4.

3. **Restart the computer.**

   For Windows 10 or 11, click the Start button, click the Power Options button or the power icon, and then choose Restart.

If restarting your computer doesn't seem to fix the problem, you may need to turn your computer off and then turn it on again. To do so, follow the previous procedure but choose Shut Down instead of Restart.

Here are a few things to try if you have trouble restarting your computer:

- **If your computer refuses to respond to the Start➪Shut Down command, try pressing Ctrl+Alt+Delete.**

  This is called the "three-finger salute." It's appropriate to say, "Queueue" while you do it.

  When you press Ctrl+Alt+Delete, Windows displays a dialog box that enables you to close any running programs or shut down your computer entirely.



WARNING

» **If pressing Ctrl+Alt+Delete doesn't do anything, you've reached the last resort. The only thing left to do is turn off the computer by pressing the power On/Off button and holding it down for a few seconds.**

Turning off your computer by pressing the power button is a drastic action that you should take only after your computer becomes completely unresponsive. Any work you haven't yet saved to disk is lost. (Sniff.) (If your computer doesn't have a Reset button, turn off the computer, wait a few moments, and then turn the computer back on again.)



REMEMBER

If at all possible, save your work before restarting your computer. Any work you haven't saved is lost. Unfortunately, if your computer is frozen or locked up, you probably can't save your work. In that case, you have no choice but to push your computer off the digital cliff.

Booting in Safe Mode

Windows provides a special startup mode called *Safe Mode* that's designed to help fix misbehaving computers. When you start your computer in Safe Mode, Windows loads only the most essential parts of itself into memory — the bare minimum required for Windows to work. Safe Mode is especially useful when your computer has developed a problem that prevents you from using the computer at all.

To boot in Safe Mode on a Windows 7 or earlier computer, first restart the computer. Then, as soon as the computer begins to restart, start pressing the F8 key — just tap away at it until a menu titled Advanced Boot Options appears. One of the options on this menu is Safe Mode; use the up- or down-arrow keys to select that option and then press Enter to boot in Safe Mode.

On a Windows 8, 8.1, or 10 computer, you can reboot into Safe Mode by holding down the Shift key when you choose the Restart command.

Using System Restore

System Restore is a Windows feature that periodically saves important Windows configuration information and allows you to later return your system to a previously saved configuration. This can often fix problems by reverting your computer to a time when it was working.

By default, Windows saves restore points whenever you install new software on your computer or apply a system update. Restore points are also saved automatically every seven days.

Although System Restore is turned on by default, you should verify that System Restore is active and running to make sure that System Restore points are being created. To do that, right-click Computer in the Start menu, choose Properties, and then click the System Protection tab. The dialog box shown in Figure 4-1 is displayed. Verify that the Protection status for your computer's `C:` drive is On. If it isn't, select the `C:` drive and click the Configure button to configure System Restore for the drive.

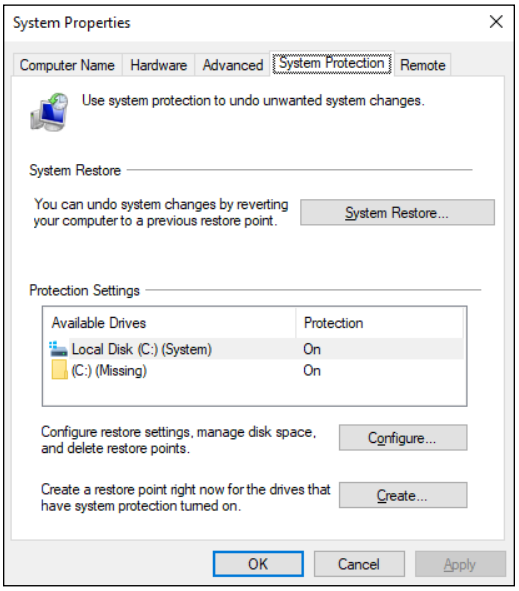


FIGURE 4-1: The System Protection tab of the System Properties dialog box.

If your computer develops a problem, you can restore it to a previously saved restore point by clicking System Restore on the System Protection tab. This brings up the System Restore Wizard, which lets you select the restore point you want to use, as shown in Figure 4-2.

Here are a few additional thoughts to remember about System Restore:

- System Restore *does not* delete data files from your system. Thus, files in your Documents folder won't be lost.

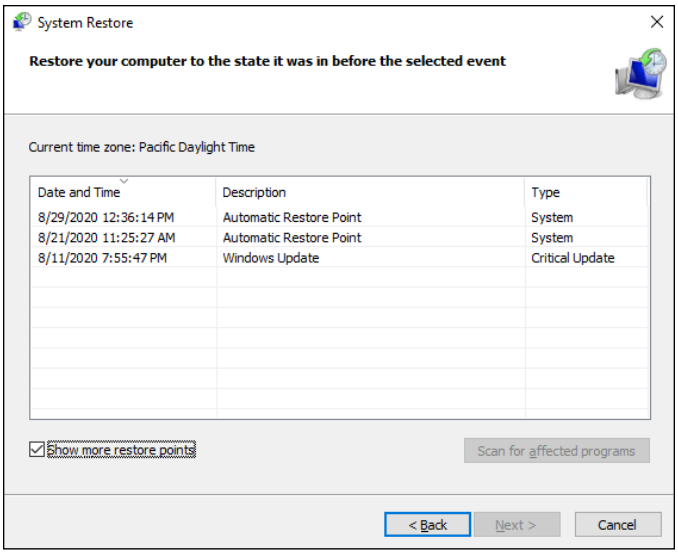


FIGURE 4-2: Use System Restore to restore your system to an earlier configuration.

- System Restore *does* remove any applications or system updates you've installed since the time the restore point was made. Thus, you need to reinstall those applications or system updates — unless, of course, you determine that an application or system update was the cause of your problem in the first place.
- System Restore automatically restarts your computer. The restart may be slow because some of the changes made by System Restore happen after the restart.
- Do *not* turn off or cut power to your computer during System Restore. Doing so may leave your computer in an unrecoverable state.
- After completing a System Restore, you may discover that the user can't log on to the computer because the computer complains about the "domain trust relationship" being lost. This happens because the internal password that Active Directory uses to verify the identity of the computer has been reset by the System Restore to a previous version. The only solution is to log in to a local account on the computer, leave the domain, reboot the computer, log in again using a local account, rejoin the domain, and reboot the computer again. Sigh.

Restarting Network Services

Once in awhile, the OS service that supports the task that's causing you trouble inexplicably stops or gets stuck. If users can't access a server, it may be because one of the key network services has stopped or is stuck.

You can review the status of services by using the Services tool, which you can access by clicking the Start button, typing **services**, and pressing Enter. Review the list of services to make sure that all key services are running. If an important service is paused or stopped, restart it.

Which services qualify as "important" depends on what roles you define for the server. Table 4-1 lists a few important services that are common to most versions of Windows. However, many servers require additional services besides these. In fact, a typical server will have many dozens of services running simultaneously.

TABLE 4-1 Key Windows Services

Service	Description
Computer Browser	Maintains a list of computers on the network that can be accessed. If this service is disabled, the computer won't be able to use browsing services, such as My Network Places.
DHCP Client	Enables the computer to obtain its IP address from a Dynamic Host Configuration Protocol (DHCP) server. If this service is disabled, the computer's Internet Protocol (IP) address won't be configured properly.
DNS Client	Enables the computer to access a Domain Name Server (DNS) server to resolve DNS names. If this service is disabled, the computer won't be able to handle DNS names, including internet addresses and Active Directory names.
Server	Provides basic file- and printer-sharing services for the server. If this service is stopped, clients won't be able to connect to the server to access files or printers.
Workstation	Enables the computer to establish client connections with other servers. If this service is disabled, the computer won't be able to connect to other servers.



WARNING

Key services usually stop for a reason, so simply restarting a stopped service probably won't solve your network's problem — at least, not for long. You should review the System log to look for any error messages that may explain why the service stopped in the first place.

Restarting a Network Server

Sometimes, the only way to flush out a network problem is to restart the network server that's experiencing trouble.



WARNING

Restarting a network server is something you should do only as a last resort. Windows Server is designed to run for months or even years at a time without rebooting. Restarting a server invariably results in a temporary shutdown of the network. If you must restart a server, try to do it during off hours if possible.


TIP

Before you restart a server, check whether a specific service that's required has been paused or stopped. You may be able to just restart the individual service rather than the entire server. For more information, see the section "Restarting Network Services," earlier in this chapter.

Here's the basic procedure for restarting a network server:

1. **Make sure that everyone is logged off the server.**

   The easiest way to do that is to restart the server after normal business hours, when everyone has gone home for the day. Then, you can just shut down the server and let the shutdown process forcibly log off any remaining users.

   To find out who's logged on, refer to the earlier section, "Who's on First?"

2. **After you're sure the users have logged off, shut down the network server.**

   You want to do this step behaving like a good citizen if possible — decently, and in order. Choose Start➪Shut Down to shut down the server. This brings up a dialog box that requires you to indicate the reason for the shutdown. The information you supply here is entered into the server's System log, which you can review by using the Event Viewer.

3. **Reboot the server computer or turn it off and then on again.**

   Watch the server start up to make sure that no error messages appear.

4. **Tell everyone to log back on and make sure that everyone can now access the network.**

   Remember the following when you consider restarting the network server:


WARNING

- Restarting the network server is more drastic than restarting a client computer. Make sure that everyone saves his or her work and logs off the network before you do it! You can cause major problems if you blindly turn off the server computer while users are logged on.
- Obviously, restarting a network server is a major inconvenience to every network user. Better offer treats.

Looking at Event Logs

One of the most useful troubleshooting techniques for diagnosing network problems is to review the network operating system's built-in event logs. These logs contain information about interesting and potentially troublesome events that occur during the daily operation of your network. Ordinarily, these logs run in the

background, quietly gathering information about network events. When something goes wrong, you can check the logs to see whether the problem generated a noteworthy event. In many cases, the event logs contain an entry that pinpoints the exact cause of the problem and suggests a solution.

To display the event logs in a Windows server, use Event Viewer, which is available from the Administrative Tools menu. When you open the Event Viewer, you'll see that a tree listing on the left side of Event Viewer lists five categories of events that are tracked: Application, Security, System, Directory Service, and File Replication Service. Select one of these options to see the log that you want to view. For details about a particular event, double-click the event to display a dialog box with detailed information about the event.

Documenting Your Trials and Tribulations

For a large network, you probably want to invest in problem-management software that tracks each problem through the entire process of troubleshooting, from initial report to final resolution. For small- and medium-sized networks, it's probably sufficient to put together a three-ring binder with pre-printed forms. Or record your log in a Word document or an Excel spreadsheet.

Regardless of how you track your network problems, the tracking log should include the following information:

- The real name and the network username of the person reporting the problem
- The date the problem was first reported
- An indication of the severity of the problem
- Is it merely an inconvenience, or is a user unable to complete his or her work because of the problem? Does a work-around exist?
- The name of the person assigned to resolve the problem
- A description of the problem
- A list of the software involved, including versions
- A description of the steps taken to solve the problem
- A description of any intermediate steps that were taken to try to solve the problem, along with an indication of whether those steps were "undone" when they didn't help solve the problem
- The date the problem was finally resolved

10 Dealing with Cybersecurity

Contents at a Glance

CHAPTER 1: **Securing Your Users** 817

CHAPTER 2: **Managing Firewalls and Virus Protection** 827

CHAPTER 3: **Dealing with Spam** 847

CHAPTER 4: **Managing Disaster Recovery and Business Continuity Planning** 861

CHAPTER 5: **Planning for Cybersecurity Incident Response** 869

CHAPTER 6: **Penetration Testing** 883

IN THIS CHAPTER

- Securing user accounts
- Keeping passwords safe
- Improving the human firewall

Chapter 1
Securing Your Users

While all the technology you can throw at security is vitally important, the weakest link in any security structure are the humans who live behind it. Consider the following actual cases:

- An attacker gains complete access to an accounting company's servers by guessing the name and password of a user with elevated permissions. The user's password was obvious and easy to guess.
- The data systems of a city government falls victim to a ransomware attack when an employee opens an attachment in a malicious email.
- An employee at a securities firm for a hospital clicks a link in a malicious email, leading to the theft of personal information from the firm's databases.
- An employee at a community college accidentally posts detailed student records on the college's website, exposing the information to the entire student body.

In every one of these cases, and in hundreds of thousands of cases like these, IT administrators had technology safeguards — firewalls, antispam and anti-virus software, permissions-based access, and so on — in place. Yet careless users caused serious damage by not paying attention to what they were doing.

In this chapter, I explore some of the things you can do to make your users more secure.

Knowing the Difference between Authentication and Authorization

Before I get far into securing user accounts, I want to make sure you understand the distinction between two important concepts of user account security:

- **Authentication:** *Authentication* refers to the process of determining that the person using an account is the person who *should be* using the account. At the most basic level, this involves entering the correct username and password. But because passwords can be compromised, sometimes additional measures are taken to authenticate a user to ensure that the user isn't actually someone else. These measures are usually called *two-factor authentication* or *multifactor authentication*. (More on this subject later, in the "Understanding Multifactor Authentication" section.)
- **Authorization:** In contrast to authentication, *authorization* refers to the process of establishing whether a given user is allowed to use a given resource on the network. Most often, the resource in question is a shared network folder, or perhaps a specific file or folder within a shared network folder. But authorization can also extend to other resources such as printers, web files, mailboxes (including calendars), and even specific computers.

In this chapter, I focus mostly on authentication. For more about authorization, refer to Book 6, Chapters 4 and 5.

Following Password Best Practices

The most basic way to ensure proper authentication of user accounts is through passwords. Usernames aren't usually considered secret. Passwords, on the other hand, are. Your network password is the one thing that keeps an impostor from logging on to the network by using your username and receiving the same access rights that you ordinarily have. *Guard your password with your life.*



TIP

Here are some tips for creating good passwords:

- **Don't use obvious passwords.** Your last name, your kid's name, or your dog's name are all easy for other people to track down.
- **Don't pick passwords based on your hobbies.** A friend of mine is into boating, and his password is the name of his boat. Anyone who knows him

can guess his password after a few tries. Five lashes for naming your password after your boat.

- **Store your password in your head, not on paper.** Especially bad: writing down your password on a sticky note and sticking it on your computer's monitor. Ten lashes for that. (If you must write down your password, write it on digestible paper that you can swallow after you memorize the password.)
- **Set expiration times for passwords.** For example, you can specify that passwords expire after 30 days. When a user's password expires, the user must change it. Your users may consider this process a hassle, but it helps to limit the risk of someone swiping a password and then trying to break into your computer system later.
- **Configure user accounts so that when they change passwords, they can't specify a password that they've used recently.** For example, you can specify that the new password can't be identical to any of the user's past three passwords.
- **Configure security policies so that passwords must include a mixture of uppercase letters, lowercase letters, numerals, and special symbols.** So, passwords like *DIMWIT* or *DUFUS* are out. Passwords like *87dIM@wit* or *duF39&US* are in.
- **Use a biometric ID device, like a fingerprint reader, as a way to keep passwords.** These devices store your passwords in a secret encoded file, and then supply them automatically to whatever programs or websites require them — but only after the device has read your fingerprint. Fingerprint readers, which used to be exotic and expensive, are available for as little as $50.

Recent research is suggesting that much of what we've believed about password security for the last 30 or so years may actually be counterproductive. Why? Two reasons:

- The requirement to change passwords frequently and making them too complicated to memorize simply encourages users to write their passwords down, which makes them easy to steal.
- A common way that passwords are compromised is by theft of the encrypted form of the password database, which can then be attacked using simple dictionary methods. Even the most complex passwords can be cracked using a dictionary attack if the password is relatively short; the most important factor in making passwords difficult to crack is not complexity but length.

As a result, the National Institute for Standards and Technology (NIST) recommends new guidelines for creating secure passwords:

- Encourage longer passwords.
- Drop the complexity requirement. Instead, encourage users to create passwords that they can easily remember. A simple sentence or phrase consisting of ordinary words will suffice, as long as the sentence or phrase is long. For example, "My password is a simple sentence" would make a good password.
- Drop the requirement to change passwords periodically; it only encourages users to write down their passwords.

Old ways are difficult to change, and it will take a while for these new guidelines to catch on. Personally, I wouldn't drop the requirement to change passwords periodically without also increasing the minimum length to at least 15 characters.

COMING UP WITH A GREAT PASSWORD

How do you come up with passwords that no one can guess but that you can remember? Most security experts say that the best passwords don't correspond to any words in the English language, but they consist of a random sequence of letters, numbers, and special characters. But how in the heck are you supposed to memorize a password like Dks4%DJ2, especially when you have to change it three weeks later to something like 3pQ&X(d8?

Here's a compromise solution that enables you to create passwords that consist of two four-letter words back to back: Take your favorite book (if it's this one, you need to get a life) and turn to any page at random. Find the first four- or five-letter word on the page. Suppose that word is *When.* Then repeat the process to find another four- or five-letter word; say you pick the word *Most* the second time. Now combine the words to make your password: *WhenMost.* I think you agree that *WhenMost* is easier to remember than *3PQ&X(D8* and is probably just about as hard to guess. I probably wouldn't want the folks at the Los Alamos National Laboratory using this scheme, but it's good enough for most of us.

Here are some additional thoughts on concocting passwords from your favorite book:

- If the words end up being the same, pick another word. And pick different words if the combination seems too commonplace, such as *WestWind* or *FootBall.*
- For an interesting variation, insert the page numbers on which you found both words either before or after the words (for example, *135Into376Cat* or *87Tree288Wing*). The resulting password will be a little harder to remember, but you'll have a password worthy of a Dan Brown novel.
- To further confuse your friends and enemies, use archaic language (for example, medieval words from Chaucer's *Canterbury Tales*). Chaucer is a great source for passwords because he lived before the days of word processors with spellcheckers. He wrote *seyd* instead of *said, gret* instead of *great,* and *litel* instead of *little.* And he used lots of seven-letter and eight-letter words suitable for passwords, such as *glotenye* (gluttony), *benygne* (benign), and *opynyoun* (opinion). And he got an A in English.
- If you do decide to go with passwords such as *KdI22UR3xdkL,* you can find random password generators on the internet. Just go to a search engine, such as Google, and search for "password generator." You can find web pages that generate random passwords based on criteria that you specify, such as how long the password should be, as well as whether it should include letters, numbers, punctuation, uppercase and lowercase letters, and so on.

If you use any of these password schemes and someone breaks into your network, don't blame me. You're the one who's too lazy to memorize *D#Sc$h4@bb3xaz5.*

Securing the Administrator Account

At least one network user must have the authority to use the network without any of the restrictions imposed on other users. This user — the *administrator* — is responsible for setting up the network's security system. To do that, the administrator must be exempt from all security restrictions.


WARNING

Windows Active Directory domains always have a master administrator account named, appropriately, Administrator. When you create a domain, the Administrator account is automatically created and you're required to enter a password for this account.


WARNING

It's absolutely vital that you not lose the Administrator password, because it's your final fail-safe method of regaining access to your domain should other accounts become compromised. And it's essential that this password can survive the person who created it. After all, that person may leave the company or — heaven forbid — get hit by a proverbial bus.

Here are some best practices for ensuring the safety of this password:

- It should be very long — perhaps 20 characters or more — and preferably randomly generated.
- Write it down, place it in a sealed envelope, and put the envelope in a safe-deposit vault at your bank. Only a few select individuals should have access to the password. It should *not* be stored in the same safe-deposit box as other corporate papers that may need to be accessed occasionally.
- No one should know the password. It should be used when the domain is created or in a disaster-recovery situation, and no one in the organization should memorize it, write it down, or store it in a file no matter how carefully encrypted.

Do not casually give out Administrative privileges to your IT staff. Windows domains include a group called Domain Admins, who have the keys to the kingdom just like the actual Administrator account does. It would be tempting to add your IT staff accounts to this group. Don't.

Instead, create *two* Active Directory accounts for your IT staff. The first is for daily work, checking email, troubleshooting, completing support tickets, and searching online for solutions to problems. The second account is an account with elevated privileges that give the user access to specific administrative tasks.

The elevated account can be named the same as the user's regular account, with something like *-admin* added to the end. For example, if my account name was dlowe, I could have an elevated account called dlowe-admin. I would only use the dlowe-admin account when performing tasks that require administrative privileges.

If your staff is small, you could just place these elevated accounts in the Domain Admins group. For a larger staff, you should hand out permissions to each staff member only as needed. For example, one person may need just the ability to make Active Directory changes, while another person focuses on maintaining share permissions.

There are huge advantages to the two-account idea. Among the most important is that when someone on your IT staff leaves, you can simply disable that person's account and be done with it. If someone who knows the Administrator account password leaves, you have to change that password immediately, and that could cause a major disruption.

Understanding Multifactor Authentication

Multifactor authentication is an authentication process in which a username and password are not enough to gain access to a system. In addition to the username and password, you must provide at least one other type of evidence that you are who you claim to be.

In general, there are three types of evidence you can present to verify your identity:

- **Something you *know:*** For example, your username and password
- **Something you *have:*** For example, a phone that can be sent a verification number or a security card you can swipe
- **Something you *are:*** For example, your fingerprint or retina pattern

Multifactor authentication requires you to submit at least two of these three categories of evidence to be authenticated. For example, when you log in using your username and password, the system may send a text message to your phone with an identification number, which you must then enter. If you don't have the phone, you can't be authenticated. Or, you may be required to swipe your finger on a fingerprint reader to gain access.

Typically, only two factors are required. When that's the case, the term *two-factor authentication* is sometimes used. If all three factors are used, the term *three-factor authentication* may be used. But that's uncommon.



WARNING

There are several common types of authentication checks that may call themselves multifactor but aren't really. For example:

- **Security questions:** These questions may ask for the name of the school you attended for the sixth grade or the city in which your parents met. Security questions aren't actually all that secure, and they certainly aren't a form of multifactor authentication. After all, the username, the password, and the name of your favorite teacher are all things you know.
- **Email verification codes:** Some authentication systems email you a confirmation number that you must enter to gain access. In my view, this is next to worthless. If a hacker has learned your username and password, there's no reason to think the hacker can't access your email and steal the verification code as well.

Securing the Human Firewall

Security techniques and technology — physical security, user account security, server security, and locking down your servers — are child's play compared with the most difficult job of network security: securing your network's users. All the best-laid security plans are for naught if your users write down their passwords on sticky notes and post them on their computers and click every link that shows up in their email inbox.

The key to securing your network users is to empower them to realize that they're an important part of your company's cybersecurity plan, and then show them what they can do to become an effective human firewall.

The following sections provide a few tips for strengthening the most important part of your security posture.

Establishing cybersecurity policies

Users will be on their own to decide what is and what isn't acceptable behavior if you don't have any policies in place. So, it's essential that you set at least basic security policies and make sure your entire staff knows them.

At the minimum, security policies should address the following:

- Password management
- Acceptable use
- Remote access
- Confidentiality
- Privacy
- Email and document retention
- Asset management and tracking

Training

Improving cybersecurity awareness involves training, and IT training is usually the most dreaded type of training there is. So, do your best to make the training fun and engaging rather than dull and boring.



TIP

If training isn't your thing, search the web. You'll find plenty of inexpensive options for online cybersecurity training, ranging from simple and short videos to full-length online courses.

You also need to establish a written cybersecurity policy and stick to it. Have a meeting with everyone to go over the security policy to make sure that everyone understands the rules. Also, make sure to have consequences when violations occur.

Phish testing

A final idea you should consider is to conduct *phish testing,* which means that you regularly send bogus emails to your staff to see how they respond. The easiest way to do that is to use one of the many cloud-based services that let you regularly send phish test emails to your staff using predefined templates. You can start with easy templates that are simple to spot because they contain obvious spelling errors and outrageous requests such as "Hi, this is your bank and we forgot your Social Security number." You can then progress to more difficult tests that send emails that look like they came from actual organizations such as banks, shipping companies, Microsoft or other software vendors, or even from your IT department.

The phishing emails contain attachments or links that lead the user to a landing page, which is basically the guy from Jurassic Park shaking his finger and saying, "Ah, ah, ahhh!" The landing pages can be customized with your logo and with whatever information you feel is relevant to help your users understand why they fell for the bait.

Phish-testing services aggregate the results of each phish test so you can measure how well your team does over time and isolate repeat offenders for additional training.



REMEMBER

The purpose of phish testing is not to humiliate or embarrass your users, but to help them become better at protecting your organization from a devastating cyberattack. Always follow up a phish test with an email that lets your staff know about the phish test, how you did as an organization, and what clues were in the email they could have noticed.

I always tell my users that the best thing to do when they get caught by a phish test is to call me before I call them. That way, we can have a conversation about how to do better the next time.

IN THIS CHAPTER

- Understanding what firewalls do
- Examining the different types of firewalls
- Considering best practices for using firewalls
- Using Group Policy to enforce firewall protection on all your computers
- Looking at virus protection
- Selecting the right antivirus software

Chapter 2

Managing Firewalls and Virus Protection

If your network is connected to the Internet, a whole host of security issues bubble to the surface. You probably connected your network to the Internet so that your network's users could access the Internet. Unfortunately, however, your Internet connection is a two-way street. Not only does it enable your network's users to step outside the bounds of your network to access the Internet, but it also enables others to step in and access your network.

And step in they will. The world is filled with hackers looking for networks like yours to break into. They may do it just for fun, or they may do it to steal your customer's credit card numbers or to coerce your mail server into sending thousands of spam messages on their behalf. Whatever their motive, rest assured that your network will be broken into if you leave it unprotected.

This chapter presents an overview of two basic techniques for securing your network's Internet connection: firewalls and virus protection.

Firewalls

A *firewall* is a security-conscious router that sits between the Internet and your network with a single-minded task: preventing *them* from getting to *us.* The firewall acts as a security guard between the Internet and your local area network (LAN). All network traffic into and out of the LAN must pass through the firewall, which prevents unauthorized access to the network.



WARNING

Some type of firewall is a must-have if your network has a connection to the Internet, whether that connection is broadband (cable modem or digital subscriber line; DSL), T1, or a high-speed fiber-optic connection. Without it, a hacker will quickly discover your unprotected network and tell his friends about it. Within a few hours, your network will be toast.

You can set up a firewall two basic ways. The easiest way is to purchase a *firewall appliance,* which is basically a self-contained router with built-in firewall features. Most firewall appliances include a web-based interface that enables you to connect to the firewall from any computer on your network using a browser. You can then customize the firewall settings to suit your needs.

Alternatively, you can set up a server computer to function as a firewall computer. The server can run just about any network operating system (NOS), but most dedicated firewall systems run Linux.

Whether you use a firewall appliance or a firewall computer, the firewall must be located between your network and the Internet, as shown in Figure 2-1. Here, one end of the firewall is connected to a network hub, which is in turn connected to the other computers on the network. The other end of the firewall is connected to the Internet. As a result, all traffic from the LAN to the Internet and vice versa must travel through the firewall.



TECHNICAL STUFF

The term *perimeter* is sometimes used to describe the location of a firewall on your network. In short, a firewall is like a perimeter fence that completely surrounds your property and forces all visitors to enter through the front gate.

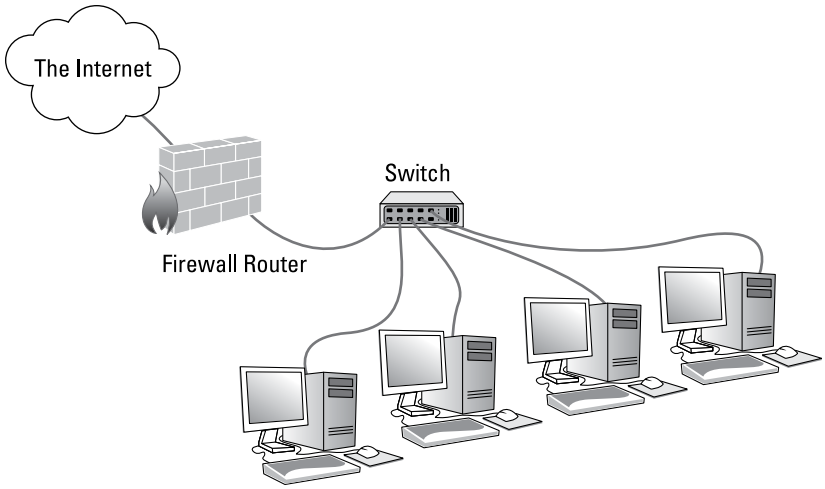


FIGURE 2-1: Using a firewall appliance.

The Many Types of Firewalls

Firewalls employ four basic techniques to keep unwelcome visitors out of your network. The following sections describe these basic firewall techniques.

Packet filtering

A *packet-filtering* firewall examines each packet that crosses the firewall and tests the packet according to a set of rules that you set up. If the packet passes the test, it's allowed to pass. If the packet doesn't pass, it's rejected.

Packet filters are the least expensive type of firewall. As a result, packet-filtering firewalls are very common. However, packet filtering has a number of flaws that knowledgeable hackers can exploit. As a result, packet filtering by itself doesn't make for a fully effective firewall.

Packet filters work by inspecting the source and destination IP and port addresses contained in each Transmission Control Protocol/Internet Protocol (TCP/IP) packet. TCP/IP *ports* are numbers assigned to specific services that help to identify for which service each packet is intended. For example, the port number for the HTTP protocol is 80. As a result, any incoming packets headed for an HTTP server will specify port 80 as the destination port.

Port numbers are often specified with a colon following an IP address. For example, the HTTP service on a server whose IP address is 192.168.10.133 would be 192.168.10.133:80.

Literally thousands of established ports are in use. Table 2-1 lists a few of the most popular ports.

TABLE 2-1 Some Well-Known TCP/IP Ports

Port	Description
20	File Transfer Protocol (FTP)
21	File Transfer Protocol (FTP)
22	Secure Shell Protocol (SSH)
23	Telnet
25	Simple Mail Transfer Protocol (SMTP)
53	Domain Name Server (DNS)
80	World Wide Web (Hypertext Transfer Protocol; HTTP)
110	Post Office Protocol (POP3)
119	Network News Transfer Protocol (NNTP)
137	NetBIOS Name Service
138	NetBIOS Datagram Service
139	NetBIOS Session Service
143	Internet Message Access Protocol (IMAP)
161	Simple Network Management Protocol (SNMP)
194	Internet Relay Chat (IRC)
389	Lightweight Directory Access Protocol (LDAP)
396	NetWare over IP
443	HTTP over TLS/SSL (HTTPS)

The rules that you set up for the packet filter either permit or deny packets that specify certain IP addresses or ports. For example, you may permit packets that are intended for your mail server or your web server and deny all other packets. Or, you may set up a rule that specifically denies packets that are heading for the ports used by NetBIOS. This rule keeps Internet hackers from trying to access NetBIOS server resources, such as files or printers.

One of the biggest weaknesses of packet filtering is that it pretty much trusts that the packets themselves are telling the truth when they say who they're from and who they're going to. Hackers exploit this weakness by using a hacking technique called *IP spoofing*, in which they insert fake IP addresses in packets that they send to your network.

Another weakness of packet filtering is that it examines each packet in isolation without considering what packets have gone through the firewall before and what packets may follow. In other words, packet filtering is *stateless.* Rest assured that hackers have figured out how to exploit the stateless nature of packet filtering to get through firewalls.

In spite of these weaknesses, packet-filter firewalls have several advantages that explain why they are commonly used:

- **Efficient:** They hold up each inbound and outbound packet for only a few milliseconds while they look inside the packet to determine the destination and source ports and addresses. After these addresses and ports are determined, the packet filter quickly applies its rules and either sends the packet along or rejects it. In contrast, other firewall techniques have a more noticeable performance overhead.
- **Almost completely transparent to users:** The only time a user will be aware that a packet-filter firewall is being used is when the firewall rejects packets. Other firewall techniques require that clients and/or servers be specially configured to work with the firewall.
- **Inexpensive:** Even consumer-grade routers include built-in packet filtering.

Stateful packet inspection (SPI)

Stateful packet inspection (SPI) is a step up in intelligence from simple packet filtering. A firewall with stateful packet inspection looks at packets in groups rather than individually. It keeps track of which packets have passed through the firewall and can detect patterns that indicate unauthorized access. In some cases, the firewall may hold on to packets as they arrive until the firewall gathers enough information to make a decision about whether the packets should be authorized or rejected.

Stateful packet inspection was once found only on expensive, enterprise-level routers. Now, however, SPI firewalls are affordable enough for small- or medium-sized networks to use.

Circuit-level gateway

A *circuit-level gateway* manages connections between clients and servers based on TCP/IP addresses and port numbers. After the connection is established, the gateway doesn't interfere with packets flowing between the systems.

For example, you can use a Telnet circuit-level gateway to allow Telnet connections (port 23) to a particular server and prohibit other types of connections to that server. After the connection is established, the circuit-level gateway allows packets to flow freely over the connection. As a result, the circuit-level gateway can't prevent a Telnet user from running specific programs or using specific commands.

Application gateway

An *application gateway* is a firewall system that is more intelligent than a packet-filtering firewall, stateful packet inspection, or circuit-level gateway firewall. Packet filters treat all TCP/IP packets the same. In contrast, application gateways know the details about the applications that generate the packets that pass through the firewall. For example, a web application gateway is aware of the details of HTTP packets. As a result, it can examine more than just the source and destination addresses and ports to determine whether the packets should be allowed to pass through the firewall.

In addition, application gateways work as proxy servers. Simply put, a *proxy server* is a server that sits between a client computer and a real server. The proxy server intercepts packets that are intended for the real server and processes them. The proxy server can examine the packet and decide to pass it on to the real server, or it can reject the packet. Or, the proxy server may be able to respond to the packet itself without involving the real server at all.

For example, web proxies often store copies of commonly used web pages in a local cache. When a user requests a web page from a remote web server, the proxy server intercepts the request and checks whether it already has a copy of the page in its cache. If so, the web proxy returns the page directly to the user. If not, the proxy passes the request on to the real server.

Application gateways are aware of the details of how various types of TCP/IP servers handle sequences of TCP/IP packets to make more intelligent decisions about whether an incoming packet is legitimate or is part of an attack. As a result, application gateways are more secure than simple packet-filtering firewalls, which can deal with only one packet at a time.

The improved security of application gateways, however, comes at a price. Application gateways are more expensive than packet filters, both in terms of their purchase price and in the cost of configuring and maintaining them. In addition, application gateways slow network performance because they do more detailed checking of packets before allowing them to pass.

Firewall Best Practices

Here's what I consider to be best practices for deploying firewalls in your organization:

- **Always protect external connections with a firewall appliance.** This is rule number one. Never allow any type of connection to the outside world that isn't protected by a firewall.
- **Don't skimp when it comes to firewalls.** There are plenty of areas in your budget where you can make compromises to cut operating costs, but firewalls are not one of them. Firewalls are expensive, but they're far less expensive than the cost of a successful cyberattack. In addition to the security features provided by the firewall, also consider the throughput capabilities of the firewall. Usually, more expensive models within a particular vendor's firewall offerings have the same features but at higher performance. If your Internet connection can support 10 Gbps, don't hamper it with a firewall that can only support 1 Gbps of net throughput — you won't be getting the benefit of that 10 Gbps Internet pipe.
- **Use firewall appliances in pairs for redundancy.** If your firewall appliance dies, your entire organization will be without Internet access until the firewall is repaired. To reduce or eliminate this downtime, use firewalls in pairs, with one designated as the primary firewall and the other as a standby that can step in if the primary firewall fails. If possible, configure these firewalls with automatic fail-over. If that's not possible, at least make sure that the procedure for manually flipping the firewall is readily available (post it on the wall near the firewall) so that you can get back online quickly. (Usually, this procedure is simply a matter of switching the cable that carries the external Internet feed from the bad router over to the standby router.)
- **Keep your firewalls up to date.** Firewalls are computers, and like all computers they periodically need software updates. Having redundant firewalls (as I suggest earlier) helps you minimize downtime during an update, because you can switch to the standby firewall while updating the primary firewall, and then you can switch back to the primary firewall while you update the standby firewall.

- **Block everything by default.** Block everything, then explicitly allow only those services that are used by your organization. Newer firewall appliances have web-based interfaces that make this process easy.
- **Document your firewall rules.** Whenever you create a firewall rule to allow a specific type of traffic, document the reason for the rule. Rules to allow traffic are created for a specific purpose — for example, your accounting department may use an application that requires you to open a specific port. Years later, when the accounting department switches to a different application, the rule that opened that port will still exist. And if you don't document the reason that the rule was created, you won't know whether you can remove the rule.
- **Periodically review your firewall logs and configuration.** Firewalls keep logs that can help you understand your network traffic. Review them regularly to ensure your firewall is performing as designed. You may discover rules that aren't being used, and you may discover gaps in your configuration that create risky exposure.
- **Enable the built-in Windows Defender Firewall on your endpoint computers**. This practice may seem redundant, because all your computers are behind advanced firewall appliances. But when it comes to cybersecurity, redundancy is a good thing: The Windows Defender Firewall may block something that slipped through your firewall router. (You'll find more information about Windows Defender Firewall in the next section.)

The Built-In Windows Firewall

Windows comes with a built-in packet-filtering firewall called *Windows Defender Firewall.* Windows has included a built-in firewall for decades, but most IT experts have for years considered it to be ineffective. However, with Windows 10, the built-in firewall has become much more capable, and is now on par with third-party firewalls that can be installed on Windows.

It's generally considered a best practice to enable Windows Defender Firewalls on all your computers, even though they're deployed behind a network firewall. The simple reason for this is that cyberattacks may originate from inside your perimeter, and such attacks won't be caught by firewalls deployed at the perimeter.

Here are the steps to activate the firewall in Windows:

1. **Choose Start ➪ Settings ➪ Privacy & Security ➪ Windows Security ➪ Firewall & Network Protection.**

   The Firewall & Network Protection page appears, as shown in Figure 2-2.

   Note that this page shows the firewall status for three types of connections: Domain, Private, and Public. By default, Windows Defender Firewall is on for all three types of networks.

2. **To change the firewall status for Domain, Private, or Public networks, click Domain, Private, or Public.**

   Figure 2-3 shows the page that appears for Domain networks. The pages for Private and Public networks are essentially the same.

3. **Use the slider button to turn Windows Defender Firewall On or Off.**

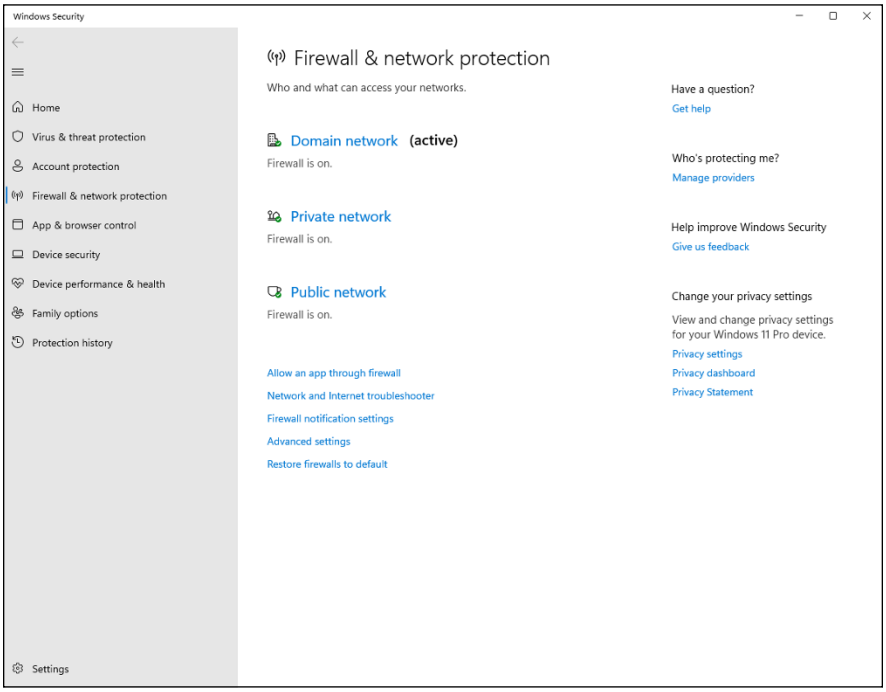


FIGURE 2-2: Windows Defender Firewall settings.

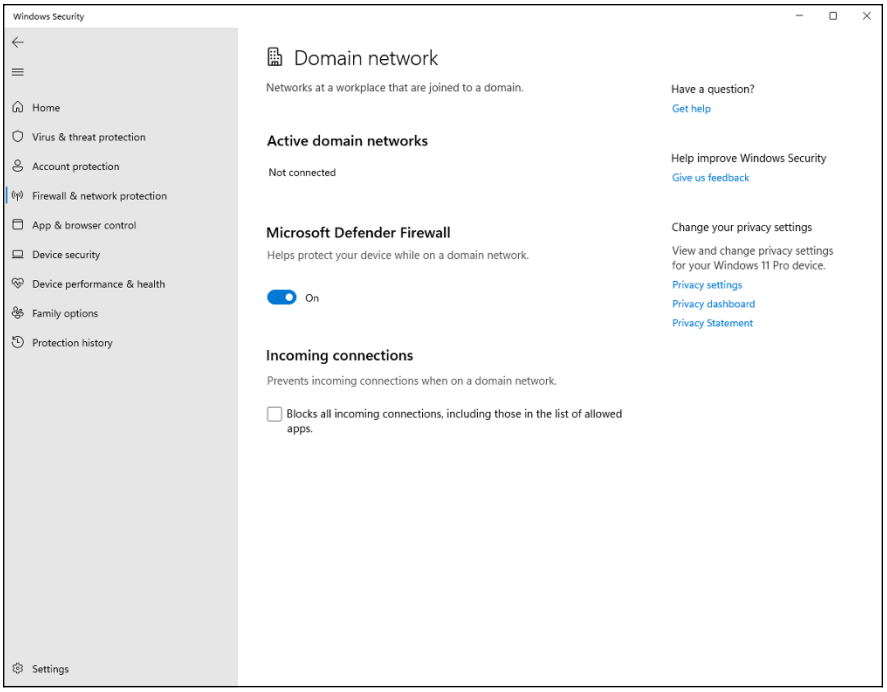


FIGURE 2-3: Activating or deactivating the firewall.

Configuring Windows Defender Firewall with Group Policy

Although you can manually configure Windows Defender Firewall as described in the preceding section, if your network has more than a few computers, you'll want to do it with Group Policy. Fortunately, Windows provides Group Policy templates that allow you to configure every aspect of Windows Defender Firewall centrally. So, you can easily fashion one or more Group Policy Objects (GPOs) to apply firewall configurations to your computers.

The Group Policy settings for Windows Defender Firewall can be found at the following location in the Group Policy Management Editor:

Computer Configuration

Policies

Windows Settings

Security Settings

Windows Defender Firewall with Advanced Security

You'll find four collections of policy settings:

Policy Collection	What It Does
Windows Defender Firewall with Advanced Security	Lets you turn on or off the firewall for Domain, Private, and Public networks.
Inbound rules	Lets you create rules that are applied to incoming traffic.
Outbound rules	Lets you create rules that are applied to outgoing traffic.
Connection Security Rules	Lets you set options that govern how other devices may connect to this computer.

The following procedure shows how to create a GPO to activate the firewall on all networks:

1. **Use Remote Desktop Connect to connect to an Active Directory server and open the Group Policy Management console.**
2. **In the Navigation pane, right-click the Group Policy Objects for your domain, and choose New.**
3. **Enter a name for the new GPO and click OK.**

   I suggest using a name like "Windows Firewall Settings."
4. **Scroll down to the GPO you just created, right-click it, and choose Edit.**

   This brings up the Group Policy Management Editor, shown in Figure 2-4.
5. **Under Computer Configuration, open Policies, Windows Settings, Security Settings, and Windows Defender Firewall with Advanced Security.**

   In Figure 2-4, I've already navigated down to the firewall policies.
6. **Right-click Windows Defender Firewall with Advanced Security and choose Properties.**

   This brings up the Properties dialog box shown in Figure 2-5.
7. **Enable the firewall on the Domain tab.**

   Change the Firewall State to Enabled.
8. **Repeat Step 7 for the Private and Public tabs.**

   The firewall is enabled for all networks.
9. **Click OK, and then close the Group Policy Management Editor.**

   You're returned to the Group Policy Management Console.

10. **Link the new GPO to enforce the policy on your end-user computers.**

    Usually the easiest way to do this is to link it to an Organizational Unit that encompasses the computers.

11. **You're done!**

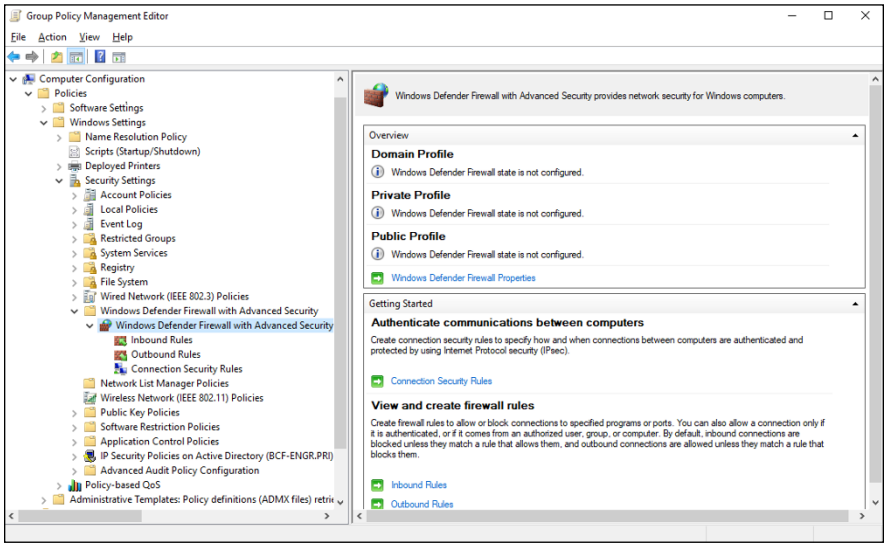


FIGURE 2-4: The Group Policy Management Editor.

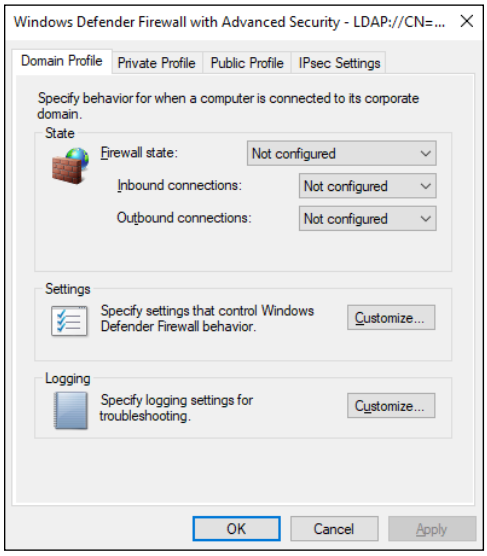


FIGURE 2-5: Enabling the Windows Defender Firewall in Group Policy.

After you've enabled the firewall on end-user computers, you'll need to add rules to open ports that are required by the applications your users need. Here are the steps to open a specific port or group of ports in the firewall:

1. **Open the Group Policy Management Editor.**
2. **Under Computer Configuration, open Policies, Windows Settings, Security Settings, and Windows Defender Firewall with Advanced Security (refer to Figure 2-4).**
3. **Select Inbound Rules, and then right-click and choose New Policy.**

   This brings up the New Inbound Rule Wizard, shown in Figure 2-6.
4. **Select the type of rule you want to create.**

   The choices are:

   - **Program:** A rule that allows a specific program to connect.
   - **Port:** A rule that opens a specific TCP or UDP port.
   - **Predefined:** Predefined rules for well-known services such as DHCP, DNS, and Remote Desktop.
   - **Custom:** A completely customizable rule that lets you choose not only protocols and ports but also ranges of IP address, specific programs that the rule applies to, and more.

   For our purposes, select Port and move on.

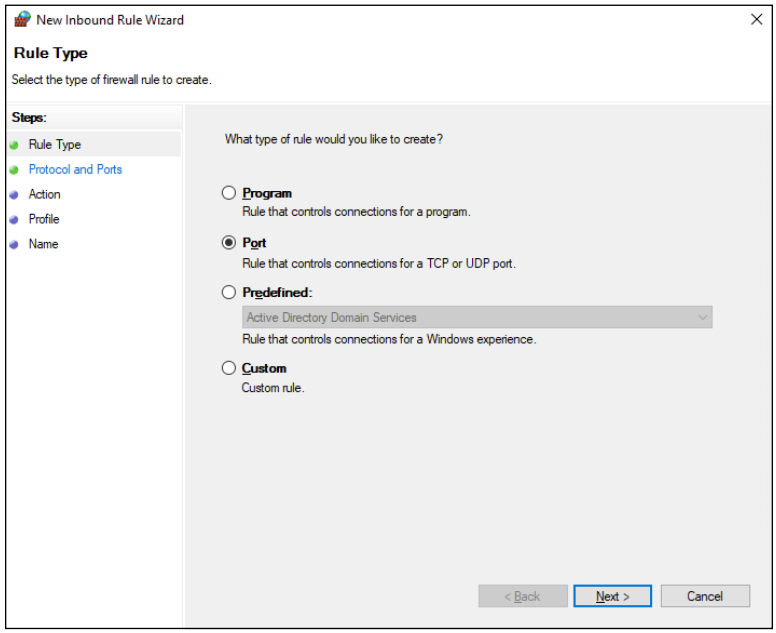


FIGURE 2-6: The New Inbound Rule Wizard.

5. **Click Next.**

   The Protocol and Ports page, shown in Figure 2-7, appears.

6. **Select whether the rule applies to TCP or UDP, and then enter the port or ports that should be opened.**

   You can enter a single port, multiple ports separated by commas (for example, **60063,60072**), or a range of ports such as **60000-60999**.

7. **Click Next.**

   The Action page, shown in Figure 2-8, appears.

8. **Choose how you want the rule to treat traffic on the ports you selected in Step 6.**

   The options are to allow the connection, allow it only if it's secure, or block the connection.

   For this example, I'll allow the connection.

9. **Click Next.**

   The Profile page, shown in Figure 2-9, appears.

10. **Choose the profile that this rule should apply to.**

    The default is to apply the rule to Domain, Private, and Public networks.

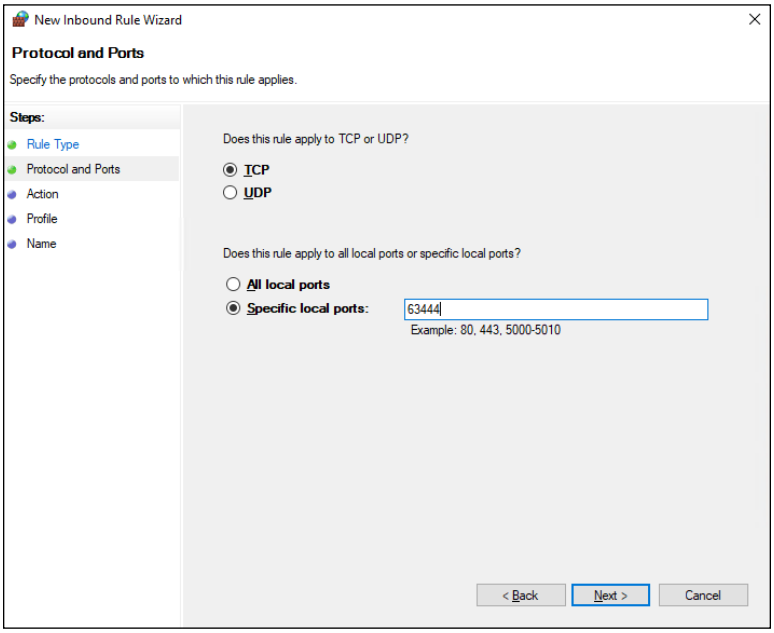


FIGURE 2-7: The Protocol and Ports page.

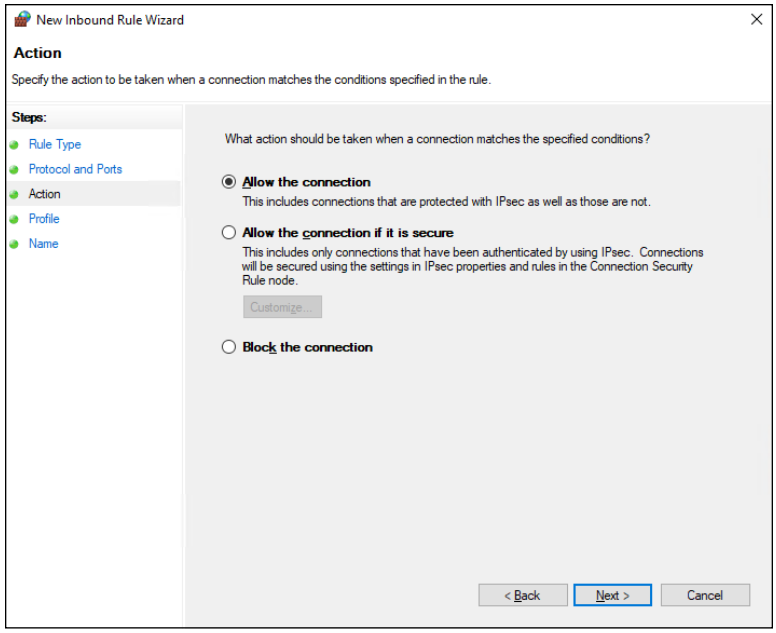


FIGURE 2-8: The Action page.

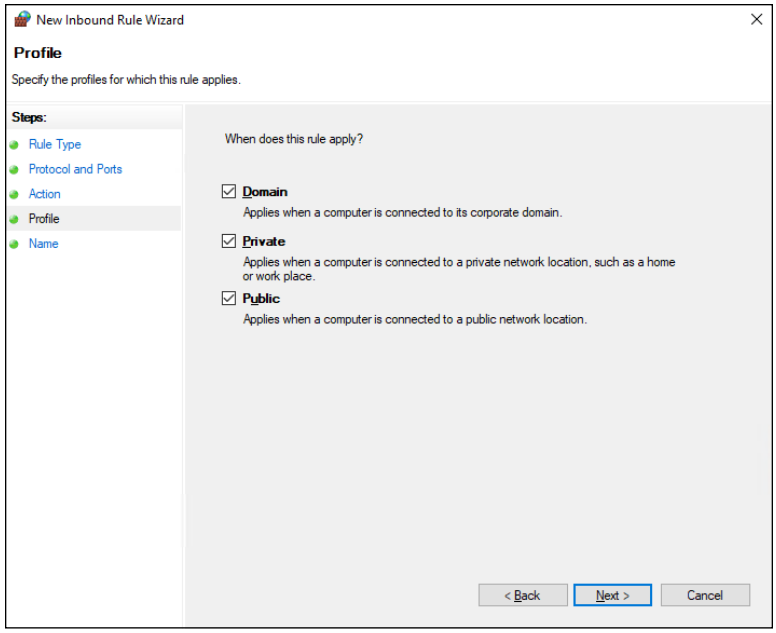


FIGURE 2-9: The Profile page.

11. **Click Next.**

    The Name page, shown in Figure 2-10, appears.

12. **Enter a Name and a Description.**

    Although the description is optional, I highly recommend you enter it. That way, you'll be able to identify why and when this firewall rule was created.

13. **Click Finish.**

    The rule is added to the policy.

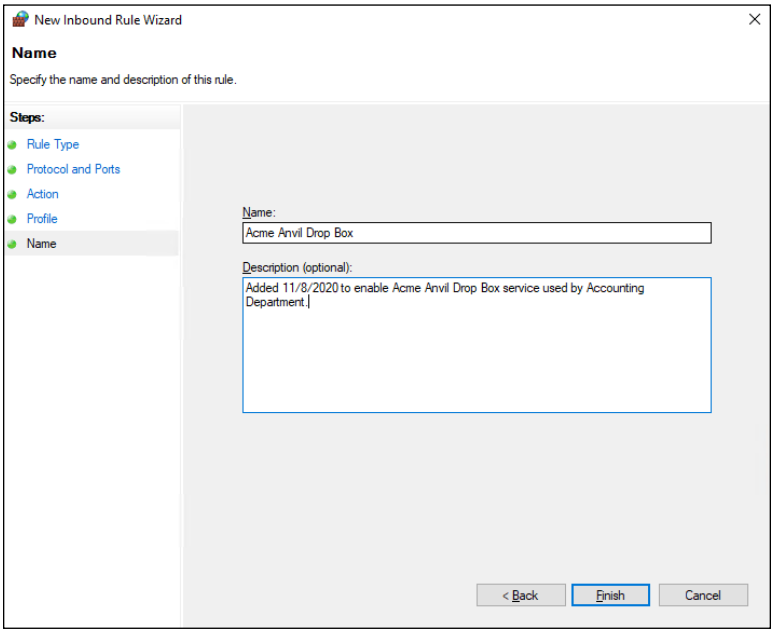


FIGURE 2-10: The Name page.

For more information about working with Group Policy, refer to Book 6, Chapter 6.

Virus Protection

Viruses are one of the most misunderstood computer phenomena around these days. What is a virus? How does it work? How does it spread from computer to computer? I'm glad you asked.

What is a virus?

Make no mistake — viruses are a serious and constant threat to your network. Every computer user is susceptible to attacks by computer viruses, and once a virus works its way onto one computer, it will try to spread itself to every computer on your network.

Viruses don't just spontaneously appear out of nowhere. Viruses are computer programs that are created by malicious programmers who've lost a few screws and should be locked up.

What makes a virus a virus is its capability to make copies of itself that can be spread to other computers. These copies, in turn, make still more copies that spread to still more computers, and so on, ad nauseam.

Then, the virus patiently waits until something triggers it — perhaps when you type a particular command or press a certain key, when a certain date arrives, or when the virus creator sends the virus a message. What the virus does when it strikes also depends on what the virus creator wants the virus to do. Some viruses harmlessly display a "gotcha" message. Some send an email to everyone it finds in your address book. Some wipe out all the data on your hard drive. Ouch.



Many years ago, in the prehistoric days of computers, viruses were passed from one computer to another by latching themselves onto floppy disks. Whenever you borrowed a floppy disk from a buddy, you ran the risk of infecting your own computer with a virus that may have stowed away on the disk.

Virus programmers have discovered that email is a very efficient method to spread their viruses. Typically, a virus masquerades as a useful or interesting email attachment, such as instructions on how to make $1,000,000 in your spare time, pictures of naked celebrities, or a Valentine's Day greeting from your long-lost sweetheart. When a curious but unsuspecting user opens the attachment, the virus springs to life, copying itself onto the user's computer — sometimes sending copies of itself to all the names in the user's address book.

After the virus works its way onto a networked computer, the virus can then figure out how to spread itself to other computers on the network. It can also spread itself by burrowing into a flash drive so that when the flash drive is inserted into another computer, that computer may become infected as well.

Here are some more tidbits about protecting your network from virus attacks:

- The term *virus* is often used to refer not only to true virus programs (which are able to replicate themselves) but also to any other type of program that's

designed to harm your computer. These programs include so-called *Trojan horse* programs that usually look like games but are, in reality, ransomware.

- A *worm* is similar to a virus, but it doesn't actually infect other files. Instead, it just copies itself onto other computers on a network. After a worm has copied itself onto your computer, there's no telling what it may do there. For example, a worm may scan your hard drive for interesting information, such as passwords or credit card numbers, and then email them to the worm's author.
- Computer virus experts have identified several thousand "strains" of viruses. Many of them have colorful names, such as the I Love You virus, the Stoned virus, and the Michelangelo virus.
- Antivirus programs can recognize known viruses and remove them from your system, and they can spot the telltale signs of unknown viruses. Unfortunately, the idiots who write viruses aren't idiots (in the intellectual sense), so they're constantly developing new techniques to evade detection by antivirus programs. New viruses are frequently discovered, and antivirus programs are periodically updated to detect and remove them.

Antivirus programs

The best way to protect your network from virus infection is to use an antivirus program. These programs have a catalog of several thousand known viruses that they can detect and remove. In addition, they can spot the types of changes that viruses typically make to your computer's files, thus decreasing the likelihood that some previously unknown virus will go undetected.

Windows comes with a built-in antivirus program called Windows Defender. Historically, Windows Defender has not had the best of reputations when compared to third-party antivirus programs. However, Defender has come a long way and is now considered as good as the competition.

If you prefer a third-party solution, here are a few popular vendors you should consider for a comprehensive antivirus platform:

- Avast (`www.avast.com`)
- McAfee (`www.mcafee.com`)
- Sophos (`www.sophos.com`)
- Symantec Endpoint Security by Broadcom (`www.broadcom.com/products/cyber-security/endpoint/end-user/enterprise`)



REMEMBER

The people who make antivirus programs have their fingers on the pulse of the virus world and frequently release updates to their software to combat the latest viruses. Because virus writers are constantly developing new viruses, your antivirus software is next to worthless unless you keep it up to date by downloading the latest updates.

Here are several approaches to deploying antivirus protection on your network:

- **Install antivirus software on each network user's computer.** This technique would be the most effective if you could count on all your users to keep their antivirus software up to date. Because that's an unlikely proposition, you may want to adopt a more reliable approach to virus protection.
- **Managed antivirus services place antivirus client software on each client computer in your network.** Then, an antivirus server automatically updates the clients on a regular basis to make sure that they're kept up to date.
- **Server-based antivirus software protects your network servers from viruses.** For example, you can install antivirus software on your mail server to scan all incoming mail for viruses and remove them before your network users ever see them.
- **Some firewall appliances include antivirus enforcement checks that don't allow your users to access the Internet unless their antivirus software is up to date.** This type of firewall provides the best antivirus protection available.



TIP

Here are a few tips for choosing an antivirus program for your environment:

- **Cost should be the last thing you consider when choosing an antivirus platform.** Antivirus software ranges from free to several hundred dollars per year per endpoint. But the cost of a successful cyberattack can be tens or hundreds of thousands of dollars, or more. In fact, a successful cyberattack could be a terminal event for your company if you aren't able to recover from backup.
- **Chose antivirus software that performs well on benchmark tests that indicate how thorough the software is at catching malware.** You can find ratings of antivirus products on the interwebs by searching for *antivirus rating*.
- **Choose antivirus software that has centralized management.** This allows you to monitor and configure your antivirus software from a central management console, which also lets you verify that all your computers are protected.
- **Choose antivirus software that can stop an active ransomware attack.** Ransomware is among the most dangerous types of cyberattacks, and some

antivirus products are actually able to detect a ransomware attack in process and stop it before it encrypts more than a few of your files. I've seen such software work well during actual ransomware attacks, stopping ransomware dead in its tracks. This is an essential feature.

- **Consider antivirus software that integrates well with your firewall.** For example, Sophos (www.sophos.com) has an excellent antivirus product that coordinates with its firewall products; the synergy between the firewall and the antivirus results in protection that is better than firewalls and antivirus software that don't coordinate their efforts.

Safe computing

Besides using an antivirus program, you can take a few additional precautions to ensure virus-free computing. If you haven't talked to your kids about these safe-computing practices, you had better do so soon.

- **Regularly back up your data.** If a virus hits you, and your antivirus software can't repair the damage, you may need the backup to recover your data. Make sure that you restore from a backup that was created before you were infected by the virus!
- **If you buy software from a store and discover that the seal has been broken on the disc package, take the software back.** Don't try to install it on your computer. You don't hear about tainted software as often as you hear about tainted beef, but if you buy software that's been opened, it may well be laced with a virus infection.
- **Use your antivirus software to scan your disk for virus infection after your computer has been to a repair shop or worked on by a consultant.** These guys don't intend harm, but they occasionally spread viruses accidentally, simply because they work on so many strange computers.
- **Don't open email attachments from people you don't know or attachments you weren't expecting.**
- **Use your antivirus software to scan any floppy disk or CD that doesn't belong to you before you access any of its files.**

IN THIS CHAPTER

- Understanding what spam is
- Seeing how antispam filters can block spam
- Options for setting up an antispam solution
- Dealing with the spam that still gets through

Chapter 3
Dealing with Spam

Spam, spam, spam, spam, spam, spam, and spam.

So goes the famous Monty Python sketch, in which a woman at a restaurant just wants to order something that doesn't have spam in it.

That pretty much sums up the situations with most people's inboxes these days. The legitimate email gets lost among the spam emails. Wouldn't you like to look at an inbox that wasn't filled with spam?

Nobody likes spam. You don't like it, and your users don't like it either. And believe me, they'll let you know if they're getting too much spam in their inboxes. They'll hold you personally responsible for every email with an offensive subject line, every email that tries to sell them stuff they aren't interested in, and every email that attempts to get them to provide their bank account password or credit card number.

As a network administrator, part of your job is protecting your users from spam. The holy grail of antispam is a solution that never allows a single piece of spam into anyone's inbox, but at the same time never mistakenly identifies a single legitimate piece of email as spam.

Good luck. This level of perfection doesn't exist. The best thing you can hope for is to find the right balance: a happy medium that lets only a small amount of actual spam through to users' inboxes and only occasionally misidentifies legitimate email as spam.

In this chapter, I explain what you need to know to find and deploy such a solution. I fill you in on the various kinds of spam, where spam comes from, how spammers get people's email addresses, and — most important — the many effective techniques you can employ to keep spam out of your users' inboxes.

Defining Spam

The most basic definition of *spam* is any email that arrives in your inbox that you didn't ask for. Spam is unsolicited email. It's email that isn't welcome, email that you aren't expecting. It's email from people you don't know or haven't heard of, usually trying to sell you something you aren't interested in or can't possibly need, and often trying to trick you into parting with either your money or your valuable personal information, or both.

One of the defining characteristics of spam is that it's sent out in bulk, often to thousands or even millions of recipients all at once. Most spam is not particularly well targeted. Instead of taking the time to figure out who might be interested in a particular product, spammers find it easier and cheaper to pitch their products to every email address they can get their hands on.

Spam is often compared to junk mail of the physical kind — the brochures, catalogs, and other solicitations that show up in your mailbox every day. In fact, spam is often called "junk email."

However, there is a crucial difference between physical junk mail and junk email. With physical junk mail, the sender must pay the cost of postage. As a result, even though junk mail can be annoying, most junk mail is carefully targeted. Junk mailers don't want to waste their money on postage to send mail to people who aren't interested in what they have to sell. They carefully measure response rates to ensure that their mailings are profitable.

In contrast, it costs very little money to send huge numbers of emails. To be sure, spam is expensive. But the bulk of the cost of spam is borne by the recipients, who must spend time and money to receive, store, and manage the unwelcome email, and by the network providers, who must build out their networks with ever greater capacity and speed to accommodate the huge volumes of spam emails that their networks must carry.

Estimates vary, but most studies indicate that as much as three-quarters of all the email sent via the internet is spam. At the time that I wrote this, there were indications that spam was actually becoming less common, accounting for closer to half of all the emails sent. But some organizations report that 80 percent or 90 percent of the email that they receive is actually spam.



WARNING

One thing is sure: Spam is not just annoying; it's dangerous. Besides filling up your users' inboxes with unwanted email, spam emails often carry attachments that harbor viruses or other malware, or entice your users into clicking links that take them to websites that can infect your network. If your network is ever taken down by a virus, there's a very good chance that the virus entered your network by way of spam.

So, understanding spam and taking precautions to block it are an important part of any network administrator's job.

Sampling the Many Flavors of Spam

Spam is unsolicited and/or unwanted email. That's a pretty broad definition, but there are several distinct categories of spam:

- **Advertisements:** Most spam is advertising from companies you've never heard of, trying to sell you products you aren't interested in. The most common type of product pitched by spam emails are pharmaceuticals, but spam also commonly promotes food supplements, knock-offs of expensive products such as watches or purses, weight-loss products, and so on.
- **Phishing emails:** Among the most annoying and dangerous types of spam are phishing emails, which try to get you to divulge private information such as credit card account numbers or passwords. Phishing email masquerades as legitimate email from a bank or other well-known institution and often includes a link to a phony website that resembles the institution's actual website. For example, you might get an email informing you that there was a suspicious charge on your credit card, with a link you can click to log in to verify that the charge is legitimate. When you click the link, you're taken to a page that looks exactly like your credit card company's actual page. However, the phony page exists solely to harvest your username and password.

  Another type of phishing email includes an attachment that claims to be an unpaid invoice or a failed parcel delivery notice. The attachment contains a Trojan that attempts to infect your computer with malware.

- **Scams:** The most common type of email scam is called an *advance-fee scam,* in which you're promised a large reward or prize in the future for advancing a relatively small amount of money now in the form of a wire transfer or money order. You may have heard of or actually received the classic scam known as the Nigerian prince scam, in which a person claiming to be a Nigerian prince needs your help to transfer a huge amount of money (for example, $40 million) but can't use an African bank account. The prince needs to use your personal bank account, and will pay you a percentage — perhaps $1 million — for your help. But you must first open a Nigerian account with a minimum balance — of perhaps $1,000 or $10,000 — to facilitate the transfer. All you have to do is wire the money, and they'll take care of the rest.

  There are many variations of this story, but they all have one thing in common: They're too good to be true. They offer you a huge amount of money later, in exchange for a relatively small amount of money now.

- **Ads for pornographic websites:** Such websites are notorious for being top sources of viruses and other malware.

- **Get-rich-quick schemes:** Pyramid schemes, multilevel marketing schemes, phony real-estate schemes, you name it — they're all in a category of spam that promises to make you rich.

- **Backscatter:** Backscatter is a particularly annoying phenomenon in which your inbox becomes flooded with dozens or perhaps hundreds of nondelivery reports (NDRs), indicating that an email that you allegedly sent didn't arrive. When you examine the NDRs, you can easily determine that you never sent an email to the intended recipient. What's actually going on here is that your email address has been used as the From address in a spam campaign, and you're receiving the NDRs from the mail servers of those spam emails that were not deliverable.



TIP

Though technically not spam, many users consider advertisements and newsletters from companies they *have* dealt with in the past to be a form of spam. An important element of the definition of spam is the word *unsolicited.* When you register at a company's website, you're effectively inviting that company to send you email.

Using Antispam Software

The most effective way to eliminate spam from your users' inboxes is to use antispam software. *Antispam software* examines all incoming email with the intent of distinguishing between spam and legitimate email. Depending on how the software is configured, email identified as spam is deleted, moved to a separate location, or simply marked as possible spam by adding a tag to the email's subject line.

Antispam software works by analyzing every piece of incoming email using sophisticated techniques that determine the likelihood that the email is, indeed, spam. When a certain threshold of probability is reached, the email is deemed to be spam and deleted, moved, or tagged. If the threshold is not reached, the email is passed on to the user as usual.


TIP

Microsoft Exchange mailboxes include a Junk folder that is often the ultimate destination of email identified as spam. You should always check your Junk folder whenever you can't find an email you're expecting.

Not all antispam programs use the Junk folder. Some programs store spam email outside of the user's mailbox, in a separate location on the network or perhaps on the cloud. These programs usually deliver a daily email (often called a *digest*) that lists the emails that were identified as spam. You should review this email whenever you can't find an email you're expecting.

Determining whether an email is spam is not an exact science. As a result, *false positives* (in which a legitimate piece of email is mistakenly identified as spam) and *false negatives* (in which a spam email is not detected as spam and makes it into the user's inbox) are not uncommon. False positives can result in your users not receiving emails they're expecting. False negatives can leave users scratching their heads wondering how in the world the spam filter didn't catch the spam. Sometimes email that to a human is obviously spam slips right by the antispam software.

The challenge of any antispam tool is finding the right balance of not too many false positives and not too many false negatives. Most antispam tools let you tune the filters to some degree, setting them to be more or less permissive — that is, erring on the side of more false negatives or more false positives. The stricter the filters are set, the more false positives you'll have. Loosening the filters will result in more false negatives.


TIP

The possibility of false negatives is one of the main reasons that it's rarely a good idea to configure an antispam program to simply delete spam. Most programs can be configured to delete only the most obvious spam emails — the ones that can be identified as spam with 100 percent certainty. Email that is probably spam but with less than 100 percent certainty should be marked as spam but not deleted.

Understanding Spam Filters

Antispam programs use a variety of different techniques to determine the probability of a given piece of email being spam. These techniques are employed by *filters*, which examine each piece of email; each filter uses a specific technique.

Here are some of the most commonly used filter types:

- **Keyword checking:** The most obvious way to identify spam is to look for certain words that appear either in the email's subject line or in the email body. For example, a keyword checking filter might look for profanity, sexual terms, and other words or phrases such as "Get rich quick!"

  Although this is the most obvious way to identify spam, it's also the least reliable. Spammers learned long ago to leave common words out of their spams to avoid these types of filters. Often they intentionally misspell words or substitute numbers or symbols for letters, such as the numeral 0 for the letter *o,* or the symbol ! for the letter *l.*

  The biggest problem with keyword checking is that it often leads to false positives. Friends and relatives might intentionally or inadvertently use any of the banned words in their emails. Sometimes, the banned words appear in the middle of otherwise completely innocent words. For example, if you list *Cialis* as a keyword that you want blocked, you'll also block the words *spe*cialis*t* or *so*cialis*t.*

  For these reasons, keyword filters are typically used only for the most obvious and offensive words and phrases, if they're used at all.

- **Bayesian analysis:** One of the most trusted forms of spam filtering is *Bayesian analysis,* which works by assuming that certain words occur more often in spam email than in other email. This sounds a lot like keyword checking, but Bayesian analysis is much more sophisticated than simple keyword checking. The Bayesian filter maintains an index of words that are likely to be encountered in spam emails. Each word in this index has a probability associated with it, and each word in the email being analyzed is looked up in this index to determine the overall probability of the email being spam. If the probability calculated from this index exceeds a certain threshold, the email is marked as spam.

  Here's where the magic of Bayesian analysis comes in: The index is self-learning, based on the user's actual email. Whenever the filter misidentifies an email, the user trains the filter by telling the filter that it was incorrect. The user typically does this by clicking a button labeled "This is spam" or "This is not spam." When the user clicks either of these buttons, the filter adjusts the probability associated with the words that led it to make the wrong conclusion. So, when the filter encounters a similar email in the future, it's more likely to make the correct determination.

- **Sender Policy Framework (SPF):** Surprisingly, SMTP (the internet email protocol) has very poor built-in security. In particular, any email server can easily send email that claims to be from any domain. This makes it easy to forge the From address in an email. SPF lets you designate via DNS which

specific email servers are allowed to send email from your domain. An antispam SPF filter works by looking up the sending email server against the SPF records in the DNS of the domain specified by the email's From address.

- **Block list:** Another trusted form of spam filtering is a *block list*, which uses a list of known spammers to block email from sources that aren't trustworthy. There are two types of block lists: private and public. A private block list is a list that you set up yourself to designate sources you don't want to accept email from. A public block list is a list that is maintained by a company or organization and is available for others to use.

  Note that simply block listing a sender email address isn't much help. That's because the sender email address is easy to forge. Instead, block lists track individual email servers that are known to be sources of spam.


WARNING

  Unfortunately, spammers don't usually set up their own servers to send out their spam. Instead, they hijack other servers to do their dirty work. Legitimate email servers can be hijacked by spammers and, thus, become spam sources, often without the knowledge of their owners. This raises the unfortunate possibility that your own email server might be taken over by a spammer, and you might find your email server listed on a public block list. If that happens, you won't be able to send email to anyone who uses that block list until you have corrected the problem that allowed your server to be hijacked and petitioned the block list owners to have your server removed.


TECHNICAL STUFF

  Historically, the block list has been called the *blacklist,* and the friends list (described in the next paragraph) has been called the *whitelist.* Words matter, and even subtle slights like these that associate negative connotations with blackness and positive connotations with whiteness are ultimately harmful to all of us. The change in terminology was long overdue.

- **Friends list:** One of the most important elements of any antispam solution is a *friends list,* which ensures that email from known senders will never be blocked. Typically, the friends list consists of a list of email addresses that you trust. When the antispam tool has confirmed that the From address in the email has not been forged (perhaps by use of an SPF filter), the friends list filters looks up the address in the friends list database. If the address is found, the email is immediately marked as legitimate email, and no other filters are applied. So, if the email is marked as legitimate by the friends list filter, the other filters are not used.


TIP

  Most friends list filters will let you friends list entire domains, as well as individual email addresses. You most certainly do *not* want to friends list domains of large email providers such as `gmail.com` or `comcast.net`. But you should friends list the domains of all your business partners and clients to ensure that emails from new employees at these key companies are never marked as spam.



TIP

Some antispam programs automatically add the recipient addresses of all outgoing emails to the friends list. In other words, anyone that you send an email to is automatically added to the friends list. Over time, this feature can drastically reduce the occurrence of false positives.

Use the friends list to preemptively allow important email that you're expecting from new customers, vendors, or service providers. For example, if you switch payroll providers, find out in advance what email addresses the new provider will be using so that your payroll staff doesn't miss important emails.

- **Graylisting:** Graylisting is an effective antispam technique that exploits the fact that if a legitimate email server can't successfully deliver an email on its first attempt, the server will try again later, typically in 30 minutes. A graylist filter automatically rejects the first attempt to deliver a message but keeps track of the details of the message it rejected. Then, when the same message is received a second time, the graylist filter accepts the message and makes note of the sender so that future messages from the sender are accepted on the first attempt.

  Graylisting works because spammers usually configure their servers to not bother with the second attempt. Thus, the graylist filter knows that if a second copy of the email arrives after the initial rejection, the mail is probably legitimate.

  The drawback of graylisting is that the first time you receive an email from a new sender, the email will be delayed. Many users find that the benefit of graylisting is not worth the cost of the delayed emails, so they simply disable the graylist filter.

Looking at Three Types of Antispam Software

The many different antispam programs that are available fall into three broad categories: on-premises, appliance, and cloud-based (hosted). The following sections describe the relative merits of each of these approaches to providing antispam for your organization.

On-premises antispam

An on-premises antispam program runs on a server on your network and interacts directly with your email server. Email that arrives at your server is passed

over to the antispam program, which evaluates the email to determine whether it's spam or legitimate mail. The antispam software uses a variety of techniques to identify spam and can usually be configured for optimal performance. Email that is identified as legitimate is handed back to the email server for normal processing. Depending on how you configure the software, email that is identified as spam may be sent to your users' Junk folders or stored in some other location.

In smaller organizations, the antispam software can run on the same server as the email server (for example, Microsoft Exchange). In larger organizations, the antispam software can be configured to run on its own dedicated server, separate from the mail server(s).

Here are some of the advantages of using an on-premises antispam product:

- **You have complete control over the configuration and operation of the software.** Most on-premises antispam software is highly configurable, often providing a dozen or more distinct filtering methods, which you can customize in many different ways. (For more information, see the section "Understanding Spam Filters," earlier in this chapter.)
- **On-premises antispam software is usually tightly integrated not only with Microsoft Exchange but also with Microsoft Outlook.** Spam email typically appears in the users' Junk folders, and the software often provides an Outlook add-in that makes it easy for users to mark incorrectly identified email.
- **On-premises software is relatively inexpensive.** Typically, you pay an upfront fee to purchase the license, as well as an annual maintenance fee to receive regular updates not only to the software but also to the spam filters.

Here are the main disadvantages of on-premises antispam software:

- **You're responsible for installing, patching, configuring, updating, and otherwise maintaining the software.**
- **Because the relationship between the email server and the antispam software is complicated, on-premises antispam software periodically malfunctions.** Such a malfunction usually halts mail flow throughout your organization. It then becomes your responsibility to correct the problem so that mail begins flowing again. (This usually happens just at the moment when your boss is expecting an important email, and you find yourself diagnosing and fixing the problem while your boss watches over your shoulder.)
- **On-premises antispam software increases the workload on your servers, requiring additional resources in the form of processor time, RAM, disk storage, and network bandwidth.**

Antispam appliances

An *antispam appliance* is essentially an on-premises server in a dedicated box that you install at your location. The appliance is usually a self-contained Linux-based computer running antispam software that is pre-installed on the appliance. This makes the appliance essentially plug-and-play; you just set it up, connect it to your network, turn it on, and configure it using a simple web-based interface. When the appliance is up and running, it can provide many, if not all, of the features of on-premises antispam software.

Here are some of the main advantages of using an antispam appliance:

- **Because the appliance includes its own hardware and pre-installed operating system, you don't have to worry about purchasing hardware separately, installing an operating system, installing software, or any of the other tasks associated with setting up a server.**
- **After it's set up, an appliance will pretty much take care of itself.** You'll need to check on it once in a while, but appliances are designed to be self-sufficient.
- **The appliance may provide other security features, such as antivirus and firewall protection.** Thus, a single appliance can handle many of your network's security and protection needs.

Using an antispam appliance is not without its disadvantages:

- **Eventually, you'll outgrow the appliance.** For example, if the number of users on your network doubles, you may run out of disk space.
- **If the appliance fails, you may have trouble getting it back up and running.** When a normal Windows server fails, you can usually troubleshoot the problem and get the server back up and running. Because of the self-contained nature of an appliance, troubleshooting it can be difficult when it's nonresponsive.

Cloud-based antispam services

A cloud-based antispam service (also called *hosted antispam*) is an internet-based service that filters your email before it ever arrives at your mail server. When you use hosted antispam, you reconfigure your public DNS so that your mail server (the MX record) points to the cloud-based antispam server rather than to your mail server. That way, all email sent to your organization is first processed by

the servers at the antispam service before it ever arrives at your mail server. Only those emails that are deemed to be legitimate are forwarded to your mail server; spam emails are stored in the cloud, where they can be reviewed and retrieved by your users if necessary.

Typically, you pay for hosted antispam based on how many users you have. For example, you might pay a monthly fee of $2 per user. As your organization grows, you simply purchase additional subscriptions.



TIP

If you use Microsoft Exchange Online as your email provider, you already have built-in antispam protection available.

Here are some of the main advantages of using cloud-based antispam:

- **You get to skip the hassle of installing and configuring software, integrating the software with Exchange, maintaining and patching the software, and all the other chores associated with hosting your own server on your own premises.** Your monthly subscription charges cover the cost of someone else doing all that work.
- **Because you don't have to buy software or hardware, there is no initial investment. You simply subscribe to the service and pay the monthly service charges.** (As an added bonus, if you're dissatisfied with the service, you can easily move to a different one. Switching to a different antispam appliance or on-premises solution is a much more complicated and expensive affair.)
- **A cloud-based antispam solution scales easily with your organization. If you double the number of users, you simply pay twice as much per month.** You don't have to worry about running out of disk space, RAM, clock cycles, or network bandwidth.
- **Cloud-based antispam takes a huge load off your network and your mail server.** Because someone else filters your spam for you, spam never enters your network. In most organizations, email is one of the most taxing applications running on the network. Using cloud-based antispam can easily cut incoming network traffic in half; in some cases, it might cut traffic by as much as 90 percent.

As you would expect, there are drawbacks to using cloud-based antispam:

- **You give up some control.** Cloud-based services usually have fewer configuration options than on-premises software. For example, you'll probably have fewer options for customizing the spam filters.

- **If the service goes down, so does your incoming email.** You won't be able to do anything about it except call technical support. And you can count on getting a busy signal, because when the service goes down, it isn't just you that's affected; it's all its customers. (Of course, this gives such services plenty of motivation to ensure that they fix the problem right away.)

Minimizing Spam


TIP

No antispam program is perfect, so you need to understand and expect that a certain amount of spam will get through to your inbox. Here are some tips that you (and your users) should keep in mind to minimize the amount of spam that gets through undetected:

- **Never trust email that requests your password or credit card.** A bank will *never* send you an email notifying you of a potential problem and containing a link to its online portal's login page. Nor will a credit card company ever send you an email alerting you to potential fraud and containing a link to a page that requests your credit card number to verify the transaction. Such emails may look very convincing, but you can rest assured they're fraudulent.

  If you're in doubt, do *not* click the link. Instead, open a browser window and navigate to the address you know for a fact to be the legitimate login page for your bank or credit card company's web portal.
- **Never open attachments in spam.** Attachments in a spam email almost certainly contain malware. Often, the malware in a spam email harvests all the contacts from your computer and sends them to the spammer, or hijacks your computer so the spammer can use it to send spam email.
- **Do not reply to spam.** If you reply to spam email, you merely confirm to the spammers that they've found a legitimate email address. You'll get even more spam.
- **Use your antispam program's "This is spam" feature.** If your antispam program has a "This is spam" or similar button, be sure to use it. Doing so alerts the antispam program that it has missed a spam message, which helps improve the filters the antispam program uses to detect spam.
- **Unsubscribe from legitimate emails.** Much of what many users consider to be spam is actually mail from legitimate organizations. If the spam is from a reputable organization, it probably isn't really spam; you probably at one time signed up to receive emails from the organization. Click the unsubscribe link on these types of emails to remove yourself from the mailing list.



WARNING

Spammers often include an unsubscribe link on their spam emails. If the email is actually spam, clicking the unsubscribe link is akin to replying to the spam — it simply confirms to the spammers that they've found a legitimate email address, and you'll just get more spam. Worse yet, the link may take you to a malicious website that will attempt to install malware on your computer. So, before you click the unsubscribe link, make sure that the email is indeed from a legitimate sender.

- **Protect your email address.** Be careful who you give your email address to, especially when you fill out forms online. Make sure you give your email address only to trusted websites. And read the fine print when you sign up for an account — you'll often find check boxes that allow you to opt out of mailings such as newsletters or announcements about product updates and so on.
- **Use an alternative email address.** One useful technique to manage the amount of spam you get is to set up a free email account with a provider such as Gmail. Then use this email account for websites that require an email address for registration when you don't want to use your real email address. You can delete or change the alternative email address if it becomes the target of spam.
- **Don't publish your email address.** If you have a personal website or are on social media, don't publish your email address there. Spammers use scanning software that trolls the internet looking for email addresses.

IN THIS CHAPTER

- Realizing the need for backups
- Making a plan
- Practicing disaster recovery
- Remembering tape rotation and other details

Chapter 4

Managing Disaster Recovery and Business Continuity Planning

On April Fools' Day about 30 years ago, my colleagues and I discovered that some loser had broken into the office the night before and pounded our computer equipment to death with a crowbar. (I'm not making this up.)

Sitting on a shelf right next to the mangled piles of what used to be a Wang minicomputer system was an undisturbed disk pack that contained the only complete backup of all the information that was on the destroyed computer. The vandal didn't realize that one more swing of the crowbar would have escalated this major inconvenience into a complete catastrophe. Sure, we were up a creek until we could get the computer replaced. And in those days, you couldn't just walk into your local Computers-R-Us and buy a new computer off the shelf — this was a Wang minicomputer system that had to be specially ordered and took weeks to deliver and install. After we had the new computer, though, a simple restore from the backup disk brought us right back to where we were on March 31. Without that backup, getting back on track would have taken months.

I've been paranoid about disaster planning ever since. Before then, I thought that disaster planning meant doing good backups. That's a part of it, but I can never forget the day we came within one swing of the crowbar of losing everything. Vandals are probably much smarter now: They know to smash the backup tapes as well as the computers themselves. Being prepared for disasters entails much more than just doing regular backups.

Nowadays, the trendy term for disaster planning is *business continuity planning* (BCP). I suppose the term *disaster planning* sounded too negative, like we were planning for disasters to happen. The new term refocuses attention on the more positive aspect of preparing a plan that will enable a business to carry on with as little interruption as possible in the event of a disaster.

Assessing Different Types of Disasters

Disasters come in many shapes and sizes. Some types of disasters are more likely than others. For example, your building is more likely to be struck by lightning than to be hit by a comet. In some cases, the likelihood of a particular type of disaster depends on where you're located. For example, crippling snowstorms are more likely in New York than in Florida.

In addition, the impact of each type of disaster varies from company to company. What may be a disaster for one company may only be a mere inconvenience for another. For example, a law firm may tolerate a disruption in telephone service for a day or two. Loss of communication via phone would be a major inconvenience but not a disaster. To a telemarketing firm, however, a day or two with the phones down is a more severe problem because the company's revenue depends on the phones.

One of the first steps in developing a business continuity plan is to assess the risk of the various types of disasters that may affect your organization. Weigh the likelihood of a disaster happening with the severity of the impact that the disaster would have. For example, a meteor crashing into your building would probably be pretty severe, but the odds of that happening are miniscule. On the other hand, the odds of your building being destroyed by fire are much higher, and the consequences of a devastating fire would be about the same as those from a meteor impact.

The following sections describe the most common types of risks that most companies face. Notice throughout this discussion that although many of these risks are related to computers and network technology, some are not. The scope of business continuity planning is much larger than just computer technology.

Environmental disasters

Environmental disasters are what most people think of first when they think of disaster recovery. Some types of environmental disasters are regional. Others can happen pretty much anywhere.

- **Fire:** Fire is probably the first disaster that most people think of when they consider disaster planning. Fires can be caused by unsafe conditions; carelessness, such as electrical wiring that isn't up to code; natural causes, such as lightning strikes; or arson.
- **Earthquakes:** Not only can earthquakes cause structural damage to your building, but they can also disrupt the delivery of key services and utilities, such as water and power. Serious earthquakes are rare and unpredictable, but some areas experience them with more regularity than others. If your business is located in an area known for earthquakes, your BCP should consider how your company would deal with a devastating earthquake.
- **Weather:** Weather disasters can cause major disruption to your business. Moderate weather may close transportation systems so that your employees can't get to work. Severe weather may damage your building or interrupt delivery of services, such as electricity and water.
- **Water:** Flooding can wreak havoc with electrical equipment, such as computers. If floodwaters get into your computer room, chances are good that the computer equipment will be totally destroyed. Flooding can be caused not only by bad weather but also by burst pipes or malfunctioning sprinklers.
- **Lightning:** Lightning storms can cause electrical damage to your computers and other electronic equipment from lightning strikes as well as surges in the local power supply.

Deliberate disasters

Some disasters are the result of deliberate actions by others. For example:

- **Intentional damage:** Vandalism or arson may damage or destroy your facilities or your computer systems. The vandalism or arson may be targeted at you specifically, by a disgruntled employee or customer, or it may be random. Either way, the effect is the same.

  Don't neglect the possibility of sabotage. A disgruntled employee who gets hold of an administrator's account and password can do all sorts of nasty things to your network.



REMEMBER

- **Theft:** Theft is always a possibility. You may come to work someday to find that your servers or other computer equipment have been stolen.
- **Terrorism:** Terrorism used to be something that most Americans weren't concerned about, but September 11, 2001, changed all that. No matter where you live in the world, the possibility of a terrorist attack is real.

Disruption of services

You may not realize just how much your business depends on the delivery of services and utilities. A BCP should take into consideration how you will deal with the loss of certain services:

- **No juice:** Electrical power is crucial for computers and other types of equipment. During a power failure once (I live in California, so I'm used to it), I discovered that I can't even work with pencil and paper because all my pencil sharpeners are electric. Electrical outages are not uncommon, but the technology to deal with them is readily available. Uninterruptible power supply (UPS) equipment is reliable and inexpensive.
- **No communications:** Communication connections can be disrupted by many causes. A few years ago, a railroad overpass was constructed across the street from my office. One day, a back-hoe cut through the phone lines, completely cutting off our phone service — including our internet connection — for a day and a half.
- **No water:** An interruption in the water supply may not shut down your computers, but it can disrupt your business by forcing you to close your facility until the water supply is reestablished.

Equipment failure

Modern companies depend on many different types of equipment for their daily operations. The failure of any of these key systems can disrupt business until the systems are repaired:

- **Computer equipment failure can obviously affect business operations.**
- **Air-conditioning systems are crucial to regulate temperatures, especially in computer rooms.** Computer equipment can be damaged if the temperature climbs too high.
- **Elevators, automatic doors, and other equipment may also be necessary for your business.**

Other disasters

You should assess many other potential disasters. Here are just a few:

- Labor disputes
- Loss of key staff because of resignation, injury, sickness, or death
- Workplace violence
- Public health issues, such as epidemics, mold infestations, and so on
- Loss of a key supplier
- Nearby disaster, such as a fire or police action across the street that results in your business being temporarily blocked off

Analyzing the Impact of a Disaster

With a good understanding of the types of disasters that can affect your business, you can turn your attention to the impact that these disasters can have on your business. The first step is to identify the key business processes that can be impacted by different types of disasters. These business processes are different for each company. For example, here are a few of the key business processes for a publishing company:

- **Editorial,** such as managing projects through the process of technical editing, copyediting, and production
- **Acquisition,** such as determining product development strategies, recruiting authors, and signing projects
- **Human resource,** such as payroll, hiring, employee review, and recruiting
- **Marketing,** including sales tracking, developing marketing materials, sponsoring sales conferences, and exhibiting at trade events
- **Sales and billing,** such as filling customer orders, maintaining the company website, managing inventory, and handling payments
- **Executive and financial,** such as managing cash flow, securing credit, raising capital, deciding when to go public, and deciding when to buy a smaller publisher or sell out to a bigger publisher

The impact of a disruption to each of these processes will vary. One common way to assess the impact of business process loss is to rate the impact of various

degrees of loss for each process. For example, you may rate the loss of each process for the following time frames:

- 0 to 2 hours
- 2 to 24 hours
- 1 to 2 days
- 2 days to 1 week
- More than 1 week

For some business processes, an interruption of two hours or even one day may be minor. For other processes, even the loss of a few hours may be very costly.

Developing a Business Continuity Plan

A BCP is simply a plan for how you will continue operation of your key business processes should the normal operation of the process fail. For example, if your primary office location is shut down for a week because of a major fire across the street, you won't have to suspend operations if you have a business continuity plan in place.

The key to a BCP is redundancy of each component that is essential to your business processes. These components include:

- **Facilities:** If your company already has multiple office locations, you may be able to temporarily squeeze into one of the other locations for the duration of the disaster. If not, you should secure arrangements in advance with a real-estate broker so that you can quickly arrange an alternative location. By having an arrangement made in advance, you can move into an emergency location on a moment's notice.
- **Computer equipment:** It doesn't hurt to have a set of spare computers in storage somewhere so that you can dig them out to use in an emergency. Preferably, these computers would already have your critical software installed. The next best thing would be to have detailed plans available so that your IT staff can quickly install key software on new equipment to get your business up and running.

  
  WARNING

  Always keep a current set of backup media at an alternative location.
- **Phones:** Discuss emergency phone services in advance with your phone company. If you're forced to move to another location on 24-hour notice, how

quickly can you get your phones up and running? And can you arrange to have your incoming toll-free calls forwarded to the new location?

Disaster planning is a very good reason to use a cloud phone provider rather than maintain an on-premises phone system (often called a *private branch exchange,* or PBX, system). With a cloud provider, your phone service won't be interrupted even if your building burns to the ground.

- **Staff:** Unless you work for a government agency, you probably don't have redundant employees. However, you can make arrangements in advance with a temp agency to provide clerical and administrative help on short notice.
- **Stationery:** This sounds like a small detail, but you should store a supply of all your key stationery products (letterhead, envelopes, invoices, statements, and so on) in a safe location. That way, if your main location is suddenly unavailable, you don't have to wait a week to get new letterhead or invoices printed.
- **Hard copy files:** Keep a backup copy of important printed material (incorporation documents, lease documents, and other legal documents) at an alternative location.

Holding a Fire Drill

Remember in grade school when the fire alarm would go off and your teacher would tell you and the other kids to calmly put down your work and walk out to the designated safe zone in an orderly fashion? Drills are important so that if a real fire occurs, you don't run and scream and climb all over each other in order to be the first one to get out.

Any disaster recovery plan is incomplete unless you test it to see whether it works. Testing doesn't mean that you should burn your building down one day to see how long it takes you to get back up and running. You should, though, periodically simulate a disaster in order to prove to yourself and your staff that you can recover.

The most basic type of disaster recovery drill is a simple test of your network backup procedures. You should periodically attempt to restore key files from your backup tapes just to make sure that you can. You achieve several benefits by restoring files on a regular basis:

- **Tapes are unreliable.** The only way to be sure that your tapes are working is to periodically restore files from them.

- **Backup programs are confusing to configure.** I've seen people run backup jobs for years that don't include all the data they think they're backing up. Only when disaster strikes and they need to recover a key file do they discover that the file isn't included in the backup.
- **Restoring files can be a little confusing, especially when you use a combination of normal and incremental or differential backups.** Add to that the pressure of having the head of the company watching over your shoulder while you try to recover a lost file. If you regularly conduct file restore drills, you'll familiarize yourself with the restore features of your backup software in a low-pressure situation. Then, you can easily restore files for real when the pressure's on.

You can also conduct walk-throughs of more serious disaster scenarios. For example, you can set aside a day to walk through moving your entire staff to an alternate location. You can double-check that all the backup equipment, documents, and data are available as planned. If something is missing, it's better to find out now rather than while the fire department is still putting water on the last remaining hot spots in what used to be your office.

IN THIS CHAPTER

- Creating an official cybersecurity incident response plan
- Knowing when you've been attacked
- Recovering from a cybersecurity attack
- Communicating in the midst of a crisis

Chapter 5
Planning for Cybersecurity Incident Response

In July 2017, credit reporting giant Equifax discovered that hackers had penetrated its database systems and stolen the personal information, including Social Security numbers, of more than 150 million of its customers. The hackers had unfettered access to Equifax's systems for 76 days before Equifax discovered the intrusion. Equifax waited until September of that year before notifying its customers that their personal data had been exposed. The incident significantly damaged Equifax's reputation and resulted in a settlement of more than half a billion dollars.

Equifax did not handle the incident well.

In contrast, in that same year, shipping giant Maersk fell victim to a ransomware attack that crippled its entire shipping operation. Maersk's IT experts soon discovered that the ransomware was spreading throughout its entire massive infrastructure, so they made the decision to shut all of it down. In total, nearly 50,000 computers and thousands of servers were infected at hundreds of sites. Shipping operations came to a halt. But because of the team's quick action, their

entire IT infrastructure was rebuilt in just ten days. The total cost was more than $300 million, but it could have been a lot worse.

Maersk *did* handle the incident well.

In both cases, the attack itself was devastating. The difference was in how the companies were able to respond. One company worked slowly, as if they had no plan. The other company worked quickly and followed a well-devised plan.

Be like Maersk, not Equifax.

This chapter guides you through one of the most important things you can do to protect your network from a cyberattack: preparing a detailed plan for how you'll deal with a cybersecurity attack when it comes. And rest assured, an attack *will* come. No matter how good of a job you do at securing your users, patching your computers, filtering out phishing emails, and ensuring your firewalls and anti-virus software are working, the cybercriminals will eventually figure out how to weasel their way into your systems. So, make a plan now on how you'll deal with that when it happens.

Seeing the Importance of a Cybersecurity Incident Response Plan

A *cybersecurity incident response plan* (CSIRP) is a detailed, step-by-step plan outlining what you'll do in a security emergency. The plan shouldn't be in your head. Instead, it should be written down and made readily accessible so you can find it when you need it. Every member of your IT team should be familiar with your plan, and your boss should be familiar with it as well.



TIP

If you're an IT department of one, enlist some computer-savvy colleagues to lend a hand — it takes more than one person to create an effective plan, and you'll definitely need some help to actually respond to an incident.

Or, better yet, engage the services of a consulting company that specializes in cybersecurity. Regardless of the size of your IT department, it's a good idea to have a relationship with a cybersecurity consultant *before* you have a bona fide incident. You don't want to be scouring the Yellow Pages in the middle of a crisis!

When you sit down to craft your plan, start by focusing on an overall framework for the plan. Within the cybersecurity industry, a well-established framework for a response plan includes the following steps:

1. **Preparation.**

   You want to be ready when a cyber emergency occurs so you don't waste time trying to figure out how to get started with your response. The most important thing you can do to prepare your organization to respond to a cybersecurity incident is to create a CSIRP.

2. **Identification.**

   You should know how to assess the severity of a cybersecurity incident. Many incidents are minor and can be easily dealt with. But some are bona fide emergencies that require a full-throated response.

3. **Containment.**

   Develop strategies to contain an attack after you've detected it so you can limit the damage. It's bad enough if one of your servers is successfully attacked by ransomware. It's much worse if *all* of them have been ransomed. Ideally, you want to stop an attack in its tracks before the attack spreads throughout your network.

4. **Eradication.**

   You want to get the attackers out of your systems and remove all trace of what they've left behind. That means removing all malware, wherever it may be lurking.

5. **Restoration.**

   When you've eradicated all traces of the attack, you can then get started on restoring your systems to working order. That may mean recovering backup data or rebuilding key pieces of your infrastructure that have been damaged beyond repair.

6. **Communication.**

   All along the way, you'll want to communicate clearly and transparently with your employees, customers, vendors, and other stakeholders. Often, managing communication in the midst of a crisis is difficult, because everyone who knows what's actually going on is up to their shoulders trying to resolve the crisis. But ongoing communication is a must.

7. **Closeout.**

   An essential part of any cybersecurity response is ensuring that the response has a proper ending. Make sure that all relevant information about the response is written down somewhere and saved for future reference. And be sure to evaluate your response to discover if there are ways you can handle the incident better in the future.

A ROSE BY ANY OTHER NAME

The term *cybersecurity incident response plan* is a bit unwieldy, and *CSIRP* sounds officious. If you prefer, feel free to come up with a simpler name for your plan. Here are a few suggestions:

- **Security incident playbook:** Short and to the point. Plus, it's a sports metaphor.
- **Security outbreak script:** It lays out the script you should follow in an outbreak of bad cyber juju. And SOS is hard to beat as an acronym.
- **Binary battle blueprint:** I'm not a real fan of this one. If it sounds like something ChatGPT or Microsoft Copilot would come up with, that's because it is.

Preparing Your Cybersecurity Incident Response Plan

Your CSIRP needs to be written down, not just saved in your head. And it shouldn't be written on the back of a napkin. A good CSIRP may run 20 or 30 pages or even more for a larger organization. It should be detailed, well organized, reviewed by several individuals, and approved by management.

If your organization has a chief information officer (CIO) or equivalent, that person should be thoroughly familiar with the plan. In the absence of a CIO, a chief operating officer (COO), chief financial officer (CFO), or even the chief executive officer (CEO) should be aware of the plan.

You should keep multiple printed copies of the CSIRP in various locations throughout your organization. Keep one copy readily available in your office. Keep another copy pinned to the wall in the main computer room. Keep copies in the glove box of your car, in your sock drawer, and under your pillow. You get the idea.

Conduct a thorough review of the plan at least once a year — more often if your environment changes frequently. Make sure to print a revision date on the front page and dispose of outdated copies every time you revise the plan.

If you use a managed service provider (MSP) or an outsourced help desk, make sure those partners have a copy of your plan and receive copies of any updates. You should review the plan regularly with these partners, because they may have suggestions for improvements or offer ideas you haven't thought of.

Do it monthly, quarterly, annually, or at whatever time interval makes sense for your organization.

Assembling Your Response Team

As part of creating your response plan, you should spend some time putting together a team of people who will work together to manage a cybersecurity crisis. Identify these people in advance and include them in your written plan along with their team roles and their contact information. Update your plan whenever this team changes.

Here are the most important people to designate for your response team:

- **Incident response manager:** This person is responsible for overseeing your response to a cybersecurity incident. They should be an experienced IT professional, preferably with a background in cybersecurity, but more important, someone who has leadership skills.

  Considering that you're the one reading this book, it's likely that this person will be you.

  The responsibilities of this role include the following:

  - Preforming a preliminary assessment of any reported security events and determining whether the event is serious enough to warrant a formal response
  - Convening the incident response team when necessary to manage an incident response
  - Coordinating the actions of the incident response team, with ultimate responsibility and authority for all decisions
  - Communicating with upper management about the nature of the incident and the progress of the response

- **Technical specialists:** Your response team will need people with technical knowledge of all systems that may be affected by a security incident. These folks should be highly skilled systems engineers who understand your systems inside and out.

  They'll research and assess the cause and impacts of the incident, recommend solutions to remove malware, mitigate the damage, restore your systems to normal operations, and implement the solutions to remove malware.

- **Cybersecurity experts:** If you don't have experts in cybersecurity within your organization, find a good cybersecurity consultant and engage them for your response. (They can also review your response plan to ensure that you haven't missed any important details!)
- **Communications czar:** If at all possible, an incident response team should have a single person acting as the mouthpiece for the team. In most cases, the incident response manager will be busy assessing and mitigating damage.

  People will line up outside the IT office in the midst of a cyber crisis with a never-ending stream of questions, advice, and comments. Having someone handle communications so the experts can focus on their work can be very useful.


TIP

Meet with your response team on a regular basis — say, quarterly — to remind them of their potential roles in a cybersecurity event, to review the response plan, and to discuss any changes that may pose new security risks.

Identifying and Reporting a Cybersecurity Incident

Let's start with a definition. The U.S. federal government officially defines *cybersecurity incident* as:

> An occurrence that actually or potentially jeopardizes the confidentiality, integrity, or availability of an information system or the information the system processes, stores, or transmits or that constitutes a violation or imminent threat of violation of security policies, security procedures, or acceptable use policies.

This definition is very broad, and it includes many occurrences that probably happen on a daily basis in even the smallest organizations. We all get spam and phishing emails. Even with the best antispam filters, some nasty emails manage to get through. Each of those emails poses a potential threat to the confidentiality, integrity, and availability of your systems. However, the vast majority of them don't rise to the level of threat that requires a formalized response. In most cases, users simply delete the offending emails and move on.

The first step in responding to an actual cybersecurity incident is determining whether an incident has risen to a level that merits a response. And the prerequisite to that first step is educating your users on when they should notify the IT department that a bona fide incident may have occurred. Here are some basic guidelines on what should be reported to the IT department:

- Any unusual or suspicious behavior that results from clicking a link in an email or on a web page or from opening an email attachment
- Any notification from your antivirus or other security management software that indicates malicious software, unauthorized access, intrusion, or any other compromise has been detected
- Any notification by email or any other means that one of your user accounts or systems may have been compromised
- The presence of files anywhere on any of your systems that are unexpected, unrecognized, or appear to have been tampered with

All users within your organization should know to report any of these occurrences to the IT department.

Keep in mind that after they've penetrated your network, the top priority of a cybercriminal will be to prevent you from finding out. Attackers often spend *months* lurking about within your systems before being discovered. That's why it is especially important to respond to any indication that something may be compromised.

Although spam email is likely to get through your antispam filters and can usually be deleted and ignored, your users should also be advised to report any suspicious email that convincingly appears to be from a legitimate source, including the following:

- Suspicious email that appears to be from another user at your company or from the IT department, the HR department, the accounting department, or any executive officer of the company
- Suspicious email that appears to be from a trusted source, such as a known customer, client, vendor, or other business partner
- Suspicious email that convincingly appears to be from Microsoft or other software vendors your company works with

Triaging Reported Incidents

After an incident has been reported, it's time to determine whether the incident is serious enough to merit a formal response. Here are some factors to consider:

- Who reported the event
- Who the event was reported to

- The essential characteristics of the event
- Whether the event actually constitutes a cybersecurity incident and merits a response
- The impact, scope, and potential severity of the incident

As part of the triage process, rate the apparent or potential severity of the incident. Here's a suggested rating scale:

- **Low:** Little if any interruption to operations is likely to occur. The incident can be mitigated with minimal disruption.
- **Medium:** The impact is significant. Some IT resources are already inaccessible and will probably remain so until the incident has been resolved. Resolution is expected to take no more than a few hours.
- **High:** Significant interruption to normal operations has already occurred. One or more major IT resources can't be accessed. Resolution is expected to take a full day or more.
- **Critical:** Many if not all mission-critical IT resources are unavailable. It's impossible to estimate the time required to resolve the issue.

Depending on the severity of the incident, you may need to take immediate action. For anything above a low severity rating, you'll need to immediately convene the incident response team to begin a formal incident response according to your plan.

For high or critical severity incidents, you may also need to perform an immediate network lockdown. That may involve disconnecting your network from the outside world, isolating your users from the network servers, disabling your wireless network, and so on.

Containing a Cybersecurity Incident

After you've initiated a cybersecurity incident response and convened the team, the first order of business is to limit the extent of the incident's impact. Containment includes isolating your IT assets from adversaries, as well as isolating impacted systems from one another to prevent further damage.

Containment requires critical decision-making related to the nature of the incident. The incident response manager will need to review appropriate containment strategies with the team and may need to consult with executive leadership to select and prioritize the containment strategies that best fit the situation.

Common containment strategies include the following:

- Restricting individual account access by disabling accounts, resetting passwords, revoking licenses, disconnecting active sessions, and so on
- Isolating individual computer access by powering down the computers, physically removing network cables, disabling Wi-Fi access, and so on
- Isolating individual virtual servers from the network by disconnecting virtual network connections
- Severing the internet by logically or physically disconnecting internet connections at the firewall
- Disabling Wi-Fi access by logically or physically disconnecting Wi-Fi access points

When you evaluate which containment strategies to put into play, weigh considerations such as the following:

- The impact on business operations.
- The impact on critical network infrastructure such as Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS)
- The likelihood that a particular containment strategy will succeed
- The ability of your incident response team to access critical systems such as network control, your virtualization platform, Active Directory, Microsoft 365, and so on
- Potential damage to systems, databases, and so on
- The possibility that the containment strategy will alert the attacker to your response and whether that is desirable

Engaging the Eradication Phase

In the eradication phase of a cybersecurity incident response, all impacts resulting from the incident are fully mitigated and removed from your system. Everything is cleaned up and no trace of the attack is left behind. For example:

- All breached user accounts have been secured. Affected users have changed their passwords and been forced to re-enroll in multifactor authentication.
- If appropriate, passwords for all accounts with administrator-level access have been reset.
- All impacted end-user systems have been fully cleaned of malware, and a full malware scan has been performed on all systems.
- Operating systems have been rebuilt if necessary.
- Firewall rules have been confirmed, and no unnecessary ports have been left open.
- All systems have been checked for patch and update status and are verifiably up-to-date.
- The antivirus status on all systems has been checked to verify that every system on your network is protected.

Restoring Lost Data and Systems

After you've completed the eradication phase, you can begin the process of restoring any data or systems that have been permanently damaged or lost. Now is the time to make your backup systems earn their keep.



WARNING

Before you begin restoring systems to normal operation, be absolutely certain that you've successfully eradicated all traces of the attack and that your network is secure. If traces of malware are left on your systems, the attack may quickly recur.



TIP

Here are some tips for successfully recovering lost systems and data:

- If affected, key IT support components such as Active Directory, DNS, DHCP, and so on should be recovered first. You'll need these services up and running before you can restore other systems.

- Prioritize data volume restoration based on need balanced against the time expected for the restoration to complete. Don't be surprised by how long it can take to restore data volumes that contain many terabytes of data.
- Some backup systems (for example, Veeam) can directly mount and run virtual machines from their backups while the underlying systems are restored. This option should be considered carefully. The advantage is that it can get your systems back up and running very quickly. The drawback is that if it turns out that eradication has not been fully successful, an ongoing attack can extend damage to the backup copies.

Considering Communication

Good communication is a key factor to the success of your response to a cyber crisis. Whether internal to your company or with the outside world, communication must take place in a timely and accurate manner. Don't allow communication to be an afterthought, dealt with only after everything has been restored. Instead, incorporate communication throughout the incident response, from the very beginning of the incident.

Internal communication

Clear and timely communications within your company are vital to the success of your incident response. The initial staff communication is important because it sets the initial expectations for the incident response.

A single person should be designated to handle the initial staff-wide notification that a cybersecurity incident has occurred. This person should meet briefly with the response team to determine the key points that must be communicated. Then they should manage the initial communication efforts independently of the team, so the rest of the team can focus on containment.

The initial communication should include the following:

- A very brief and general description of what is known about the nature and scope of the incident. For example:

  We are currently investigating a possible security breach that we believe has affected some of our file servers.

- Instructions regarding what you need staff to do immediately. For example:

  > We need everyone to immediately log off of their computers and then turn the computers off.

- Instructions on what you would like staff to do until more information is available. For example:

  > Please remain here. The IT department is actively working on the problem — please don't interrupt them.
  >
  > If you have a smartphone with a cellular connection, you can access email and Microsoft Teams.
  >
  > Other than that, please don't attempt to access any other of our systems.

- An admonition to not publicize the incident outside the company. For example:

  > Please refrain from sharing any information about this with anyone outside of our company, directly or on social media. We need more time to determine exactly what's going on, and we don't want to start rumors.

- An expectation for when more information will be available. For example:

  > We'll try to get more information to you within 30 minutes.

Methods of communication

If your internet connection is still working, the easiest way to deliver your communications is via email, Microsoft Teams, or any other chatting platform you use at your company.

Direct verbal communication may be required with staff if you have no internet. Even if you still have internet access, direct verbal communication may be preferrable to other methods.

Executive leadership communication

Don't neglect your company's executive leadership during an incident response. They need to know what's going on, and you may need their approval to take drastic actions, such as restoring systems to a previous backup, which will result in lost data.

Here are some strategies for communicating with executive leadership:

- Report all information about impacts of the incident with executive leadership as soon as possible.
- Regularly update executive leadership about the status of the incident response.
- If possible, get executive leadership approval for any containment actions that will have significant operational impact prior to implementation. If that's not possible, make sure in advance that you have the authority to implement any containment measures you deem necessary to preserve your company's critical systems or data.

External communication

As a general rule, it's best to avoid sending any information about a cybersecurity incident to anyone outside of your company without first seeking approval from executive leadership. Executives will be appropriately concerned about protecting the company's reputation, which will likely take a hit when word of a cybersecurity incident gets out. This can damage trust in the company and may ultimate affect the bottom line. As a result, executives will want to ensure that all external communications are carefully managed.

Factual accuracy is very important in any public communication, but it's especially critical in the early stages of an incident response when the full scope of the incident may not be fully understood. Stick with what you know so far, what steps you've taken so far, and what you plan to do next. Don't speculate.


TIP

Your best role may be to advocate for transparency and accuracy. In general, those companies who try to hide security incidents end up worse off than those who are transparent from the start. But ultimately, the executive team controls the messaging.

Closing the Incident

The close-out phase includes reporting and post-incident analysis on the systems that were impacted by the incident and other potentially vulnerable systems. You want to have a formal close-out to every incident response in order to improve your response the next time an incident occurs.

Documentation

You should document and file all details of an incident response process so it can be easily accessed later. Here are some suggested items to document:

- How the incident was discovered and reported
- How you determined that the incident justified a formal incident response
- What immediate actions were taken when the incident was reported, and the outcomes of those actions
- A description of the exact sequence of events that led to the incident, if known
- What steps were taken to ensure that all malware, configuration changes, and other damage caused by the incident were eradicated
- What restoration steps were taken
- What preventive measures, if any, were put in place to prevent a similar attack in the future
- Any relevant event logs and audit reports
- All internal and external communications
- Any notes taken by team members during the response

Lessons learned

There's an old saying that the only thing we can truly learn from history is that we never truly learn anything from history. When it comes to the human condition, that may be true. But it doesn't have to be true in the world of cybersecurity. Every time you respond to a cyberattack, you can review your response and learn how to do it better next time.

So, make sure your team meets one final time to talk about what, if anything, they learned during the response. Did anything surprise them? Were obvious preventive measures overlooked? Did they do anything along the way that made the situation worse? Did they act fast enough? Or did they act *too* fast? And, most important, would they do anything differently next time?

Then revise your response plan accordingly.

IN THIS CHAPTER

- Testing your network to see how secure it is
- Understanding what a penetration test looks like
- Looking at some of the tools used by penetration testers
- Preparing yourself for what the testers will find

Chapter 6
Penetration Testing

A *penetration test*, also known as a *pentest*, is an attempt to gain access to a computer system in an effort to discover vulnerabilities in the system. It's a form of hacking called *ethical hacking* (hacking that is done for good purposes by the good guys to prevent the bad guys from doing it for nefarious purposes).

A pentest is not the same thing as conducting a thorough review of your cybersecurity practices. A security review involves confirming that your organization is using the best security practices it can. For example, you should

- Review your Microsoft Windows update policies to ensure that they're working and that all your computers are up-to-date.
- Confirm that multifactor authentication (MFA) is in place for all your users and that everyone has enrolled at least one MFA factor.
- Confirm that your antivirus software is working properly so that all your computers are protected.

All these checks are done passively, simply by reviewing your security configuration to confirm that it follows best practices. And it's done by IT staff who know their way around your systems and already have access.

In contrast, a pentest is done actively, by actually testing whether your security configuration works as it's designed. The person conducting a pentest pretends to be a cybercriminal attacking your systems to see if they can get in. In addition, a pentest can extend inside the system to see if, after they get in, they can exploit vulnerabilities to gain control of your systems.

Proper penetration testing requires substantial technical skill. The tester must be at least as knowledgeable as the cybercriminal; otherwise, they'll miss things that the criminal won't. As a case in point, Equifax, the credit rating company that was breached back in 2017 (see the previous chapter), actually had conducted a penetration test *prior* to the breach. Unfortunately, the tester missed an obvious flaw in one of the company's web servers, which was easily exploited by the attackers. Had Equifax hired a better penetration tester, the incident (which cost them well over half a billion dollars) could have been averted.

This chapter provides a very brief introduction to pentesting. You won't learn how to conduct your own pentest from this chapter — not even close. But you will learn what penetration testing is, why you need it, and what to expect when you hire an experienced cybersecurity consultant to perform a pentest for you.

Understanding Ethical Hacking

Penetration testers use hacking techniques to penetrate systems, but they do so in an ethical manner. Thus, pentesters are sometimes called *ethical hackers,* and what they do is called *ethical hacking.*

Hacking culture is as old as computers. Older, in fact: The first hackers were people who exploited the fact that a tone frequency of 2600 Hertz (cycles per second) played an important role in AT&T's switching networks dating back to the 1950s, and could be hacked to enable free long-distance phone calling. To this day, *2600* is code for hacking, and one of the most popular hacking publications is called *2600: The Hacker Quarterly.*

When computer networks were developed in the late 1960s and through the 1970s, hackers developed techniques to break into those networks so they could access free computer time, long before the days of inexpensive phone computers. In the 1980s, hacking culture was popularized by movies such as *Tron, War Games,* and *Ferris Bueller's Day Off.*

Interesting, isn't it, that Mathew Broderick is in two of those movies? I guess he was the quintessential hacker of the 1980s: an awkward kid who just wanted to have fun.

Fast-forward to today, hackers aren't just kids looking for a good time anymore. Instead, they're criminals looking to steal your data or extort you out of thousands or even millions of dollars.

In the movies, the typical "ethical" hacker is the tech-savvy cybercriminal who gets caught, is convicted, and then agrees to go work for the government to prevent World War III from breaking out. But in the real world, ethical hackers are skilled, well trained, and highly respected professionals who prefer to work *for* you rather than *against* you.

Introducing the Red Team

The term *red team* is often used to describe the group of penetration testers who perform a penetration test. Referring to the pentesters as the red team makes the whole endeavor feel a bit like a spy thriller, and with good reason: The term actually back to the 1960s, at the height of the cold-war era, when war games and simulated attacks were perceived as a matter of national survival. The color red, associated with the Soviet Union, was naturally used to refer to the adversaries in these games and simulations.

In cybersecurity testing, there isn't much danger of triggering global thermonuclear war. But the stakes are still pretty high. The red team plays the role of the adversaries in a simulated cyberattack and tries every technique at their disposal to break into your systems. When the test is complete, they report back to you what they found so you can plug the holes before an actual adversary finds them.

In some cases, the red team engages in a contest against a *blue team,* which is aware of the red team's game and actively tries to thwart their penetration efforts. But more often than not, a blue team is not used. Instead, the red team quietly attempts to penetrate the system and, when the test is completed, submits a report on their efforts and the effectiveness of the system's defenses.

Seeing How Penetration Testing Works

Penetration testing works by simulating the same activities that an actual adversary would use to attack a system. A penetration test usually follows a well-defined sequence of steps similar to the following:

1. **Engagement.**

   In this step, you, as a potential customer, contact a cybersecurity company with experience conducting penetration tests to discuss a possible engagement.

This conversation will mostly focus on setting the goals for the test, the scope of what will be tested, the methodology the testers will use, and any limits you want to place on what they may or may not do. All these details are vital to a successful testing engagement.

2. **Open-source research.**

   In this step, the pentesters learn as much as they possibly can about your organization using information widely available on the internet. The most common sources for this information include your company's own website, social media platforms (especially LinkedIn), news articles about your organization, and so on. In some cases, they may actually visit your physical work sites, sit in the parking lot for a while to see who comes and goes, find out where your staff gets their coffee or eats lunch so they can hang out within earshot, and so on. If your company publishes a blog or has a newsletter, the pentesters may subscribe to it. They may pretend to be a customer and call your sales or support lines to see what they can learn. They may even apply for a job to gain insights into your HR department.

3. **Scanning.**

   Having gathered readily available information about your organization, the pentesters will use specialized software to find out whatever they can about your computer networks. That starts with readily available information such as domain registration tools such as Whois (`www.whois.com`) and detailed Domain Name System (DNS) information such as `https://dnschecker.org/mx-lookup.php`. They also use scanning tools, which can poke your public IP addresses for open ports and then look deeper for vulnerabilities on those ports. They can use advanced tools such as Nmap (`https://nmap.org`) to map your network and Wireshark (`www.wireshark.org`) to sniff network packets traveling into and out of your network.

4. **Penetration.**

   After the scanning phase is complete, the pentesters will begin attacks designed to exploit weaknesses they've found. Here are just a few of the techniques they'll try:

   - **Known vulnerabilities:** Having found open ports on your firewall or public-facing web servers, they'll use specialized tools to exploit known vulnerabilities on those ports, hoping that your firewalls or servers are not fully patched.
   - **Password cracking:** Knowing the email addresses of some of your employees, they'll attempt to extrapolate usernames and then use sophisticated tools that attempt to break passwords, trying lists of commonly used passwords a few at a time to avoid having the account blocked for too many password failures.

- **Targeted phishing:** Having collected the email addresses of some of your employees, they may send carefully crafted phishing emails. If they manage to get a response from your HR director to a job application, those phishing emails will use an email address similar to the HR director's email address and will have your HR director's actual email signature. If they find out what college an employee attended, the phishing emails may be advertising a reunion event for that college. And so on. The exploitations they attempt will be much more carefully crafted than the typical random phishing attempt.

5. **Escalation.**

   If they succeed in getting in, they'll briefly celebrate. Very briefly. Then they'll begin the process of escalation. After they get in, they'll conduct internal scans of your network to see what other resources they can attack. They'll start mapping out your computers, subnets, switches, routers, Wi-Fi access points, servers, printers, phones, and anything else that has an IP address. They'll find your domain controllers, your Exchange servers, your SQL servers, your web servers, and the coveted file servers, which contain your priceless data.

   Most important, they'll find your Active Directory controllers, and if they do, they'll grab one of the best prizes of all: the `NTDS.dit` file, which contains (among other things) the password hash database. If they get that, it's game over: They'll be able to use password hash cracking tools to get the passwords for any users who use any of many thousands of well-known passwords. And they can even use pass-the-hash tools to break into user accounts using the password hash without even knowing the passwords.

   Along the way, they'll capture screenshots of just about everything they see.

6. **Reporting.**

   When they've gotten as far into your systems as they can go, or — more likely — when they run out of time, they'll prepare a detailed report. The report will outline all the techniques they used during the test, what they learned, and what worked and what didn't. They'll include many of those screenshots they captured. And they'll list the specific vulnerabilities they exploited, such as "This server was not patched" or "This port was left wide open" or "This user's password was P@55Word."

   Most important, they'll condense all these findings into a set of recommendations, categorized by the degree of importance.

   Some of the recommendations will be urgent: "Fix this problem before you go home from work tonight." Some will be low-hanging fruit, like "Make this user change their password right now."

Some will be important but not urgent — things you should fix as soon as you can, given your resources and balanced against your other priorities.

And some will be minor — things you can fix when you have some free time and have already solved the day's Wordle.

After the report has been delivered, the test will officially be over.

Scoping a Penetration Test

An important first step in planning a penetration test is determining the scope of the test: Where do you want the test to begin, where do you want it to end, and exactly what limits do you want to impose on the testing team?

One key factor in determining the nature of the test is how much information about your organization and the systems being tested you want to provide to the testing team. Maybe you want the team to know next to nothing about your organization, other than the business's name and perhaps its web address. If that's the case, you may want to conduct your initial meetings with the testing team at their offices so they learn absolutely nothing by entering your facilities. This type of test is often called a *black-box test* because the testing team knows nothing about the system they're attempting to penetrate.

On the other hand, you may want to limit the scope of a test to a specific system. For example, if you want to know whether your firewall is impenetrable, you may tell the pentesters the make and model of the firewall you're using, where it's located on your network, and even the full configuration details. Then, armed with this knowledge, they can attack your firewall repeatedly, using every attack vector they can think of, to see if the firewall is as secure as you think it is. This type of test is called a *white-box test.*

Or, you may prefer something in between, called a *gray-box test.* For this type of test, you give the testers a limited amount of information about the system to be tested. In particular, you may set up a user account and give them the credentials. Then they can proceed with an *internal pentest,* as if an adversary has already obtained a user's credentials. The purpose of this type of test is not to determine whether it's possible to break into your systems, but instead to determine how much damage to expect when an adversary succeeds in breaking in.

Establishing Boundaries

Setting boundaries for the pentest — boundaries that the pentesters may not cross — is crucial. Here are some examples to consider:

- **Executives' accounts:** You may want to insist that your executives' accounts are off-limits. To do that, of course, you'll need to let them know who the executives are. Just provide the usernames for the accounts you want left alone so they'll know not to exploit those accounts, even if they're able to discover login credentials.
- **Sensitive systems or data:** If you have mission-critical systems (such as an accounting or payroll system) that you want left alone, be sure to designate them as off-limits. If there are databases or file servers that contain personal information like credit-card numbers or Social Security numbers, let them know where that data is so they can leave it alone.
- **Physical penetration:** Be sure to discuss whether it's okay for the pentesters to try to get into your building. Unless you have a very good reason (like you're the CIA), avoid making this restriction. Pentesters love to play dress-up, and most of them know that if you wear a shirt with a fake company logo and carry a box of tools — or better yet, a ladder — you can bluff your way into almost any company. After they get in, they can wander the halls looking for empty offices with computers carelessly left unlocked. All they need is a few minutes to put a sniffing device that looks just like a USB flash drive into the computer. Then that device can catch every keystroke, including the user's password the next time they log in.

  Fun story: I once helped with a Halloween event at my local zoo (the Chaffee Zoo in Fresno, California, which is an amazing zoo for a town the size of Fresno). To plan for the event, I took an afternoon off and wandered around the zoo with a clipboard to take notes. I quickly learned that if you walk around a zoo wearing khaki pants and a hat and carrying a clipboard, you'd better know what time the bird show starts because you *will* be asked.
- **Confidentiality:** It goes without saying, but anything the pentesters learn during their test must be held in absolute confidence and shared only with you and others you designate, such as the company's executive leadership or other IT security staff. If the pentesters are somehow discovered by any other employee during the test, they must immediately notify you and back off. They shouldn't disclose that they're conducting a cybersecurity test.

Examining Tools for Penetration Testing

You may be surprised to learn that there is an entire cottage industry that builds tools to conduct penetration testing, and that most of it is open source and, therefore, completely free to anyone who wants to learn how to use it. Even though the overwhelming majority of penetration tests are targeted at Windows environments, the tools themselves are almost exclusively Linux-based. So, pentesters are usually Linux gurus.

For information about working with Linux, refer to Book 8.

How to set up Kali

Several complete Linux distributions come preconfigured with a large battery of pentesting tools, so you can get all the most valuable tools in a single download. The best known of these distributions is called Kali, and it's available for download at `www.kali.org`.

Kali is remarkably simple to set up. You can download it in the following formats:

- **Installer images (ISO):** Installer images are available for 64-bit Intel, 32-bit Intel, or Apple (ARM64) architecture. Just download the ISO image, write it to an optical disk, and install it on the target computer. Or, better yet, mount the ISO image in a virtual machine (VM) and install it that way.
- **VM image:** Even easier, you can download a pre-installed VM image for any of four different virtualization platforms: Hyper-V, QEMU, VirtualBox, or VMware. Just download the VM image and follow the provided instructions to import it into your virtualization platform. This is the fastest way to get going with Kali.
- **ARM:** The ARM version is designed to be run on single-board computers such as Raspberry Pi.
- **Mobile:** Run Kali on an Android device.
- **Cloud:** Run Kali on several cloud platforms, including Amazon Web Services (AWS) and Microsoft Azure.
- **USB:** If you prefer, you can download a bootable image that you can copy to a USB flash drive. Throw the flash drive into a USB port, configure the device to boot from USB, and you're up and running.



TIP

The default administrative credentials for Kali are as follows:

Username: kali

Password: kali

What you'll find on Kali

Kali comes preconfigured with about 100 tool packages that are specially designed for various aspects of penetration testing. For complete documentation on all the tools, go to `www.kali.org/tools`.

The tools are available from the Kali desktop by clicking the Kali menu icon shown in the margin. You'll find this icon at the upper-left corner of the display. Figure 6-1 shows this menu in action. Kali's tools are organized into 14 sections, as listed in Table 6-1.

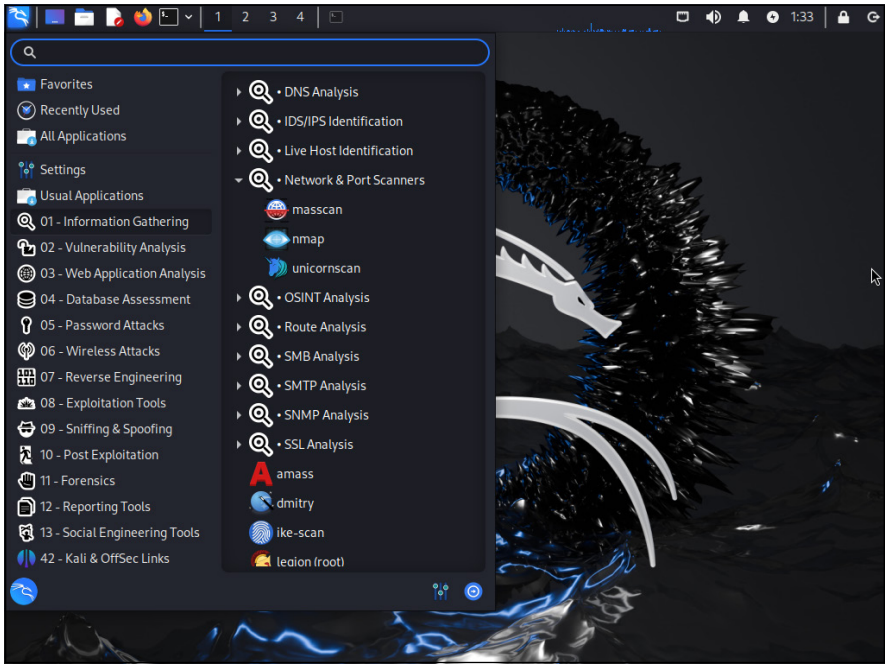


FIGURE 6-1: The Kali desktop provides quick access to Kali's pentesting tools.

TABLE 6-1 Kali Pentesting Tools

Category	Description
01 - Information Gathering	Tools to gather general information about target systems and networks
02 - Vulnerability Analysis	Tools for identifying vulnerabilities
03 - Web Application Analysis	Tools for analyzing web servers and web applications
04 - Database Assessment	Tools for finding vulnerabilities in database servers

(continued)

TABLE 6-1 *(continued)*

Category	Description
05 - Password Attacks	Tools for cracking passwords
06 - Wireless Attacks	Tools for finding vulnerabilities in Wi-Fi networks
07 - Reverse Engineering	Tools for reverse-engineering software
08 - Exploitation Tools	Tools for exploiting vulnerabilities
09 - Sniffing & Spoofing	Tools for packet sniffing and spoofing, including Wireshark
10 - Post Exploitation	Tools to use after a vulnerability has been exploited
11 - Forensics	Tools for doing forensic analysis on computer systems
12 - Reporting Tools	Tools to assist the creation of pentest reports
13 - Social Engineering Tools	Tools to assist with social engineering
42 - Kali & OffSec Links	Links to helpful websites

Looking at one of Kali's tools

To give you a feel for what the tools in Kali are capable of, let's look at one of the most useful: the `netdiscover` tool, which discovers devices that are available on an unknown network. This tool is especially useful after a pentester has gained access to a network, because it can search for other devices on the network and provides their IP addresses, as well as their MAC addresses and the device manufacturer.

To use this tool, open the menu, choose 01 – Information Gathering, and then choose netdiscover. A command window opens up, displaying the help text for the `netdiscover` command followed by a command prompt. Enter the following command:

```
$ sudo netdiscover –i eth0
```

Note that `sudo` is required because you must be an administrator to run this command. You'll be prompted to enter the password for your account.

The `netdiscover` command will immediately start scanning your network and will probably find your gateway router and a few other devices fairly soon. Let it run indefinitely, and it will eventually find all the devices available on your network.

Here's an example of the output from `netdiscover`:

```
Currently scanning: Finished !  |   Screen View: Unique Hosts

567 Captured ARP Req/Rep packets, from 12 hosts.  Total size: 33876
_____________________________________________________________________________
  IP            At MAC Address     Count    Len  MAC Vendor / Hostname
_____________________________________________________________________________
192.168.100.1   c0:94:33:2a:68:9a      1      60 ARRIS Group, Inc.
10.0.0.1        c0:94:33:2a:68:9e    185   11100 ARRIS Group, Inc.
10.0.0.140      10:62:3e:83:a2:f5     33    1980 Hewlett Packard
10.0.0.145      29:80:23:0f:f4:85    113    6780 Hewlett Packard
10.0.0.112      90:dd:5d:3f:10:9b    194   11640 Apple, Inc.
10.0.0.189      78:8c:c5:11:5e:4d      7     420 TP-Link Corporation Limit
10.0.0.177      72:ad:8e:a5:cc:e1      4     240 Unknown vendor
10.0.0.84       80:e8:2c:cc:4e:0d      8     336 Hewlett Packard
10.0.0.89       7e:45:4a:00:47:6b      5     300 Unknown vendor
10.0.0.198      98:e8:fb:08:fa:d8      8     480 Nintendo Co.,Ltd
10.0.0.187      7c:bb:8a:d5:99:db      7     420 Nintendo Co.,Ltd
10.0.0.219      4e:54:88:5d:92:ad      2     120 Unknown vendor
```

As you can see, netdiscover has found a total of 12 devices on this network. For each device, you can see the IP address, the MAC address, and the manufacturer. There are a total of two Arris devices, three Hewlett-Packard devices, an Apple device, two Nintendo devices, a TP-Link device, and a few devices whose manufacturers couldn't be identified.

With this information in hand, an attacker could do a port scan of each of these IP addresses to determine which ports are open. If the attacker gets a response on port 80, the attacker could try connecting with a web browser. For example, 10.0.0.1 opens a login page. It turns out to be an Xfinity residential gateway router. I wonder if the owner bothered to change the default password?

Knowing What to Expect from a Penetration Test

As you can imagine, there is a lot more to penetration testing than this short chapter has described. But I want to leave you with a parting bit of advice about what to expect if you engage a professional to perform a pentest on your network: Prepare to be embarrassed.

No one, no matter how good they are, survives a pentest unscathed. Even organizations that have hundreds of dedicated IT staff with an entire department devoted to cybersecurity use outside professionals to do penetration tests. Why? Because the testers *always* find something. That's the purpose of the testing, after all.

So, don't be afraid to pay for a penetration test because you're afraid of what they may find. You *want* them to find everything they can. Testing the security of your network doesn't make the network *less* secure; it makes the network *more* secure.

Go into the pentest with that attitude, and you'll get the result you're looking for: a more secure network.

Index

Symbols and Numerics

| (pipe character), 560, 595, 728
! (exclamation mark), 727
(hash mark), 726
$ (dollar sign), 726, 728
* (asterisk), 727
> (greater-than sign), 560, 727
2.5GBase-T Ethernet cable, 109
5G cellular technology, 253
5GBase-T Ethernet, 109
10Base2 Ethernet cable, 108–109
10Base5 Ethernet cable, 108
10Base-T Ethernet cable, 109
10GBase-T Ethernet cable, 110
25GBase-T Ethernet cable, 110
40GBase-T Ethernet cable, 110
100-Base-TX Fast Ethernet cable, 109

A

A (address) records, 195–196
A+ certification, 766
access control lists (ACLs), 529
access layer, 243
access points (APs), 339
accessibility, cloud computing, 83
Account Operators group, Windows Server 2025, 525
accounting services, 82
ACLs (access control lists), 529
Acronis, 298
Active Directory
  defined, 72, 498
  domains, 499–500, 502–503
  forests, 501–502
  objects, 498–499
  organizational units, 500–501, 503–505
  overview, 497–498
  trees, 501
Active Directory Domains and Trusts console, Windows Server 2025, 493
Active Directory Sites and Services console, Windows Server 2025, 493
Active Directory Users and Computers console, Windows Server 2025, 493
adapters. *See* ports
`add` parameter
  `Net Computer` command, 566
  `Net Group` command, 569
  `Net Localgroup` command, 572
address (A) records, 195–196
Address Resolution Protocol (ARP)
  `arp` command, 207–208
  defined, 51
  IP address conversion, 111
  network layer, 121
Admin Center
  Exchange Online, 626–629
  Microsoft Teams, 650–652
administration, network
  ITAM (IT asset management)
    asset identifier, 796
    asset-tracking software, 798–799
    cellphone vendor portals, 799–800
    copier vendor portals, 800
    equipment loss policy, 797
    identifying and listing assets, 793
    importance of, 793
    labeling hardware, 796–797
    overview, 792–793
    photographing assets, 795
    router management page, 800
    software, 798
    switch management page, 800
    what to track, 794–795
  network administrators
    applying software patches, 762
    assigning, 21
    bluffs and excuses, 769–770
    certifications, 765–768
    choosing, 759–760
    clean-up, 761

administration, network *(continued)*
defined, 21
gurus and, 768–769
managing users, 761–762
overview, 757–758
routine chores, 761
software tools, 763–764
staying up-to-date, 764–765
overview, 757
RDC (Remote Desktop Connection)
configuring options for, 777–782
connecting to server, 774–776
defined, 771
enabling, 772–773
keyboard shortcuts, 776–777
Remote Assistance
accepting invitation, 787–789
defined, 771
enabling, 783–784
overview, 782
requesting, 784–786
troubleshooting
basic steps, 802–803
booting in Safe Mode, 809
checking network settings, 806
dead computers, 803–804
error messages, 805
event logs, 813–814
narrowing down possibilities, 806–807
network connections, 804–805
overview, 801–802
restarting client computer, 808–809
restarting network server, 812–813
System Restore, 809–811
tracking log, 814
viewing current network users, 807
Administrator account, 487–488, 821–822
Administrators group, Windows Server 2025, 525
Adobe Creative Suite, 798
Advanced options, RDC, 781–782
advertisements, spam, 849
aliases, PowerShell, 591–593
Amazon Web Services (AWS)
Amazon CloudFront, 88
Amazon Elastic Compute Cloud (EC2), 88, 442
Amazon Elastic File System (EFS), 442
Amazon Relational Database Service (RDS), 88, 442
Amazon Simple Storage Service (S3), 88
Amazon Virtual Private Cloud (VPC), 88, 442
Analytics, 443
Business Applications, 442
Console, 444–446
creating account, 443
Developer Tools, 442
End-User Computing, 443
Machine Learning, 443
Management & Governance, 442
overview, 441
Security, Identity & Compliance, 442
VMs
connecting to, 455–457
managing, 454–455
Windows Virtual Machine, 446–453
American National Standards Institute (ANSI), 94
American Standard Code for Information Interchange (ASCII), 104
Analytics & Reports option, Microsoft Teams Admin Center, 652
`AND` logical operation, 126
ANSI (American National Standards Institute), 94
antispam appliances, 856
antispam software
antispam appliances, 856
cloud-based antispam services, 856–858
in cybersecurity, 60
overview, 850–851
on-premises antispam, 854–855
antivirus programs, 844–846
AppleTalk, 114
application layer
OSI model, 25, 105
TCP/IP, 122
application servers, 266
application virtualization, 464
APs (access points), 339
archive bit, 304–305
ARCNET (Attached Resource Computer Network), 29, 43
ARP (Address Resolution Protocol)
`arp` command, 207–208
broadcast packets, 51
defined, 51
IP address conversion, 111
TCP/IP network layer, 121

arp command, 207–208
ARPANET, 113, 117
arson, 863
ASCII (American Standard Code for Information Interchange), 104
asset identifier, 796
asset management system, 60
asset number, 796
asset-tracking software, 798–799
assumptions, in this book, 3
asterisk (*), 727
asymmetrical technologies, 250
AT attachment (ATA) interface, 287
ATA (AT attachment) interface, 287
Attached Resource Computer Network (ARCNET), 29, 43
attributes, Active Directory objects, 499
auditing, 61
authority, URL, 177
autodisconnect parameter, Net Config command, 567
Auto-Forwarded Messages card, Exchange Admin Center, 628
automatic private IP addressing, 171
AWS (Amazon Web Services)
  Amazon CloudFront, 88
  Amazon Elastic Compute Cloud (EC2), 88, 442
  Amazon Elastic File System (EFS), 442
  Amazon Relational Database Service (RDS), 88, 442
  Amazon Simple Storage Service (S3), 88
  Amazon Virtual Private Cloud (VPC), 88, 442
  Analytics, 443
  Business Applications, 442
  Console, 444–446
  creating account, 443
  Developer Tools, 442
  End-User Computing, 443
  Machine Learning, 443
  Management & Governance, 442
  overview, 441
  Security, Identity & Compliance, 442
  VMs
    connecting to, 455–457
    managing, 454–455
    Windows Virtual Machine, 446–453
Azure
  Azure Portal, 428–429
  creating account, 427
  IaaS, 425
  mobile applications, 427
  networking, 426
  overview, 89, 425–426
  PaaS, 425
  SaaS, 425
  services, 426–427
  SQL database, 427
  storage, 426
  VMs
    connecting to, 438–439
    managing, 435–438
    overview, 426
    Windows Virtual Machine, 429–435
  Web applications, 426

B

backbone, 27
backscatter, 850
backup appliance, 303–304
Backup Operators group, Windows Server 2025, 525
backup servers, 267
backups
  backup selection, 306
  backup servers, 267
  cloud backups, 299
  cybersecurity, 62–63
  file-based
    archive bit, 304–305
    copy backups, 307–308
    daily backups, 308
    differential backups, 309
    full backups, 306–307
    incremental backups, 308–309
    overview, 297–298, 304–305
  generations of, 297
  HCI, 320–321
  image-based, 298, 310
  local copy, 296–297
  NAS device, 298, 303
  network administrators and, 761
  network-attached backup appliances, 298–299, 303–304
  offline copy, 296–297
  off-site copy, 296–297
  overview, 295–297
  security, 310–311

backups *(continued)*
tape
cleaning magazines, 302–303
hardware, 300–301
LTO, 299–300
overview, 298
reliability, 301–302
virtual-based, 298, 310
Windows Server 2025, 474
bare metal, 272
Bash shell, 676
batch files, 562–563
Bayesian analysis, 852
BCP (business continuity planning), 862, 866–867. *See also* disasters
binary battle blueprint, 872
binary system
counting by ones, 123–126
logical operations, 126–127
Windows Calculator, 127–128
BIND DNS server
editing configuration files, 721
installing, 720–721
`named.conf` file, 721–723
restarting, 724
zone files, 723–724
bits, 123
blacklist, 853
blade servers, 78–79
block list, 853
booting
host computer, 157
in Safe Mode, 809
Windows Server 2025 multiboot feature, 471–472
Bourne, Stephen, 676
Bourne Shell, 676
bridges
data link layer, OSI model, 99
switches, 49–50
broadband connections, 250–251
broadcast domains, 51, 134
broadcast frequency, WAP, 341
broadcast packets, 38, 43, 51
Broadcom, Inc., 413
bus topology, 26–27
business continuity planning (BCP), 862, 866–867

C

C: drive, 661–662
cable internet, 250–251
cable modem, 146
cables
10Base2 Ethernet cable, 108–109
10Base5 Ethernet cable, 108
10Base-T Ethernet cable, 109
Cat5e cable, 109
Cat6 cable, 31
defined, 8, 15, 24
hubs, 32
patch cables, 31–32, 330, 805
overview, 31–32
patch panels and, 31–32
repeaters, 32
RJ45 connectors, 31
switches, 32–33
twisted-pair cable. *See* twisted-pair cable
`cache` parameter, `Net Share` command, 574
caching, 190
CALs (Client Access Licenses), 471
Canonical Name (CNAME) record, 196–197
Carrier Sense Multiple Access/Collision Detection (CSMA/CD), 98, 99–100
`cat` command, 735–736
Cat5e cable, 109
Cat6 cable, 31
CCAr (Cisco Certified Architect), 768
CCIE (Cisco Certified Internetwork Expert), 768
CCNA (Cisco Certified Network Associate), 768
CCNP (Cisco Certified Network Professional), 768
CCT (Cisco Certified Technician), 768
`cd` command, 730–731
cellphone vendor portals, 799–800
cellular network, 253
cellular router, 255
Center for Internet Security (CIS), 64
Cert Publishers group, Windows Server 2025, 526
`chage` command, 741–742
channels, Microsoft Teams, 642–643
checkpoints, 397
`chgrp` command, 746
child partition, 394
`chmod` command, 746–747

chown command, 745
CIDR (classless interdomain routing notation), 137
circuit-level gateway, 832
CIRP (cybersecurity incident response plan)
  close-out phase, 881–882
  communication, 879–881
  containing cybersecurity incident, 876–877
  cybersecurity incident, 874–875
  defined, 870
  documentation, 882
  eradication phase, 878
  framework, 870–871
  hackers, 869–870
  lessons learned, 882
  names, 872
  preparing, 872–873
  rating scale, 876
  response team, 873–874
  restoring lost data and systems, 878–879
  triaging reported incidents, 875–876
CIS Critical Security Controls, 64
Cisco certifications
  CCAr (Cisco Certified Architect), 768
  CCDE (Cisco Certified Design Expert), 768
  CCIE (Cisco Certified Internetwork Expert), 768
  CCNA (Cisco Certified Network Associate), 768
  CCNP (Cisco Certified Network Professional), 768
  CCT (Cisco Certified Technician), 768
Cisco Certified Architect (CCAr), 768
Cisco Certified Design Expert (CCDE), 768
Cisco Certified Internetwork Expert (CCIE), 768
Cisco Certified Network Associate (CCNA), 768
Cisco Certified Network Professional (CCNP), 768
Cisco Certified Technician (CCT), 768
Cisco Firepower 1000 Series routers, 255
Citrix XenApp, 463–465
Class A addresses, 133
Class B addresses, 133
Class C addresses, 134
classless interdomain routing notation (CIDR), 137
Client Access Licenses (CALs), 471
clients
  defined, 13
  VPNs, 388–390
  Windows clients
    configuring for DHCP, 169–171
    configuring network connections, 357–363
    joining domain, 363–366
clock speed, processor, 76
close parameter, Net File command, 568
cloud backups, 299
cloud computing
  accessibility, 83
  cost effectiveness, 82
  drawbacks of
    entrenched applications, 84
    internet connection reliability, 85
    internet connection speed, 84–85
    security threats, 85
  IaaS, 87
  overview, 81–82
  PaaS, 86–87
  private clouds, 87–88
  providers
    AWS, 88
    Google, 88–89
    Microsoft, 89
  public clouds, 87–88
  recommendations, 89–90
  reliability, 83
  SaaS, 86
  scalability, 83
  traditional computing vs., 82
Cloud Computing For Dummies (Kirsch/Hurwitz), 90
Cloud+ certification, 767
cloud-based antispam services, 856–858
clusters, HCI, 316, 322
cmdlets, PowerShell
  naming conventions, 587
  overview, 586
  parameters, 587–589
CNAME (Canonical Name) record, 196–197
COBIT (Control Objectives for Information and Related Technologies), 64
Cockpit
  administrative centers, 688–701
  configuring Linux for networking, 705–709
  enabling, 699
  installing, 699
  overview, 698

collision domains, 48–49, 52
collisions, Ethernet packets, 37–38, 43
command line interface (CLI), 676
command mode, Vi text editor, 692
command parameter Net Help command, 570
commands
  Linux
    administering users, 739–745
    cat command, 735–736
    cd command, 730–731
    chage command, 741–742
    chgrp command, 746
    chmod command, 746–747
    chown command, 745
    cp command, 733–734
    directory-handling commands, 730–733
    dnf command, 739
    editing commands, 726
    environmental variables, 728
    file-handling commands, 733–736
    gpasswd command, 744
    groupadd command, 743
    groupdel command, 743–744
    hostname command, 748
    ifconfig command, 715, 748–749
    ls command, 732–733
    mkdr command, 731
    mv command, 735
    netstat command, 749–750
    networking, 747–753
    newusers command, 742–743
    ownership and permissions, 745–747
    packages, 738–739
    passwd command, 742
    ping command, 750–751
    piping, 727–728
    pwd command, 730
    redirection, 727–728
    rm command, 734–735
    rmdir command, 731
    with root-level privileges, 729–730
    route command, 752
    service command, 737
    services, 737–738
    shell scripts, 728–729
    shells, 723–724
    sudo command, 680
    traceroute command, 752–753
    useradd command, 739–740
    userdel command, 741
    usermod command, 741
    wildcards, 726–727
    yum command, 738–739
  PowerShell
    aliases, 591–593
    cmdlets, 585, 587
    common parameters, 588–589
    FileSystem provider, 597–598
    functions, 585
    Get-ChildItem cmdlet, 591–595, 598
    Get-Process cmdlet, 599
    Get-PSProvider cmdlet, 597
    native commands, 585
    parameters, 587–589
    piping technique, 593–597
    scripts, 585
    Select-Object cmdlet, 595–596, 599
    Set-ExecutionPolicy cmdlet, 598–601
    Sort-Object cmdlet, 599
  Vi text editor
    changing text, 695
    deleting text, 694
    exiting, 691
    inserting text, 693–695
    moving cursor, 692–694
    operating modes, 691–693
    overview, 688–690
    repeating, 696
    saving changes, 691
    yanking and putting text, 695–697
  Windows Server 2025
    batch files, 562–563
    chaining commands, 559
    command prompt window, 556–558
    Control menu, 557–558
    editing commands, 557
    environmental variables, 560–562
    EventCreate command, 563–564
    Net Accounts command, 565
    Net Computer command, 566

Net Config command, 566–567
Net Continue command, 567
Net File command, 568
Net Group command, 568–570
Net Help command, 570
Net Helpmsg command, 570–571
Net Localgroup command, 571–572
Net Pause command, 572–573
Net Session command, 573
Net Share command, 574–575
Net Start command, 575
Net Statistics command, 575–576
Net Stop command, 576–577
Net Time command, 577
Net Use command, 577–578
Net User command, 579–580
Net View command, 580–581
overview, 555
piping, 559–560
redirection, 559–560
RunAs command, 581–582
wildcards, 558
comment parameter
Net Group command, 569
Net Localgroup command, 571
common parameters, PowerShell, 588–589
communication
executive leadership, 880–881
internal, 879–880
methods, 880
communications, cybersecurity, 63
Component Services console, Windows Server 2025, 493
CompTIA certifications
A+, 766
Cloud+, 767
Linux+, 766
Network+, 767
Securiiy+, 767
Server+, 767
Computer Management console, Windows Server 2025, 493–494
ComputerName parameter
Net Computer command, 566
Net Session command, 573
Net View command, 581
*computername**sharename*, Net Use command, 578
configuring
DHCP for Windows clients
automatic private IP addressing, 171
manually, 169–171
releasing leases, 171
renewing leases, 171
scopes, 163–169
WAPs, 340–342
Windows clients
for DHCP, 169–171
network connections, 357–363
Windows Server 2025
Administrator account, 487–488
Microsoft Management Console, 491–496
Remote Desktop Connection, 488–491
consoles, Linux
defined, 676–677
remote consoles, 678–679
virtual consoles, 672, 677
consoles, Windows Server 2025, 493–494
Control Objectives for Information and Related Technologies (COBIT), 64
conversations, 103
copier vendor portals, 800
copy backups, 307–308
core layer, 243
country code domains, 179–180
cp command, 733–734
crimp tool, 332
CSMA/CD (Carrier Sense Multiple Access/Collision Detection), 98, 99–100
current profile, cybersecurity, 66
cybersecurity
asset management system, 862
backups, 310–311
business continuity planning, 862, 866–867
cloud computing, 85
current profile, 66
cybersecurity incident response plan
close-out phase, 881–882
communication, 879–881
containing cybersecurity incident, 876–877
cybersecurity incident, 874–875
defined, 870
documentation, 882
eradication phase, 878

cybersecurity *(continued)*
framework, 870–871
hackers, 869–870
lessons learned, 882
names, 872
preparing, 872–873
rating scale, 876
response team, 873–874
restoring lost data and systems, 878–879
triaging reported incidents, 875–876
disasters
analyzing impact of, 865–866
deliberate disasters, 863–864
disruption of services, 864
environmental disasters, 863
equipment failure, 864
fire drill, 867–868
overview, 861–862
firewalls
application gateway, 832–833
best practices, 833–834
circuit-level gateway, 832
defined, 827
packet filtering, 829–831
stateful packet inspection, 831
Windows Defender Firewall, 834–842
frameworks
CIS Critical Security Controls, 64
COBIT, 64
ISA/IEC 62443, 64
ISO/IEC 27001, 64
NIST Framework, 64–67
overview, 63–64
network administrators and, 761
NIST Cybersecurity Framework
Framework Core section, 65–67
Framework Organizational Profiles section, 66
Framework Tiers section, 66
overview, 64–65
overview, 57–58
penetration testing
confidentiality, 889
establishing boundaries, 889
ethical hacking, 884–885
executives' accounts, 889
Kali tool, 891–893
overview, 883–884
red team, 885
scoping, 888
sensitive sytems or data, 889
steps in, 885–888
what to expect from, 893
prevention measures
anti-spam software, 60
auditing, 61
data protection, 61
encryption, 61
firewalls, 60
multifactor authentication, 60
overview, 59–60
passwords, 60
physical security, 61
user life-cycle management, 61
user training, 61
Wi-Fi security, 60
recovery measures
backups, 62–63
communications, 63
emergency disk capacity, 63
overview, 59, 61–62
spare computers, 63
for small businesses, 58–59
spam
ads for pornographic websites, 850
advertisements, 849
antispam software, 850–851, 854–858
backscatter, 850
defined, 848–849
filters, 851–854
get-rich-quick schemes, 850
minimizing, 858–859
overview, 847–848
phishing emails, 849
scams, 850
users
Administrator account, 821–822
authentication, 818
authorization, 818
cybersecurity policies, 824
multifactor authentication, 823

overview, 817
passwords, 818–820
phish testing, 825
training, 824–825
virus protection
antivirus programs, 844–846
overview, 842–844
safe-computing practices, 846
VPN, 387–388
WAP, 342
cybersecurity incident, 874–875
cybersecurity incident response plan (CIRP)
close-out phase, 881–882
communication, 879–881
containing cybersecurity incident, 876–877
cybersecurity incident, 874–875
defined, 870
documentation, 882
eradication phase, 878
framework, 870–871
hackers, 869–870
lessons learned, 882
names, 872
preparing, 872–873
rating scale, 876
response team, 873–874
restoring lost data and systems, 878–879
triaging reported incidents, 875–876

D

daily backups, 308
daisy-chaining, 27, 241
DAS (direct attached storage), 291–292
Dashboard option, Microsoft Teams Admin Center, 651
data link layer, OSI model, 25, 98–99
data protection, 61
data stores, 276, 414
data-at-rest encryption, 61
database servers, 266
data-in-flight encryption, 61
DCHP relay, 159
DEB package manager, Linux, 684
dedicated lines, 252
dedicated servers, 14
deduplication
backup appliances, 303–304
in HCI system, 317–320
in-line, 317–318
post-process, 318
default groups, Windows Server 2025
Builtin container, 521
overview, 521
Users container, 522
`del` parameter, `Net Computer` command, 566
`delete` parameter
`Net Group` command, 569
`Net Localgroup` command, 572
`Net Session` command, 573
`Net Share` command, 574
`Net Use` command, 578
deliberate disasters, 863–864
demarcation point, 147
deployment servers, 267
desktop virtualization
Citrix XenApp, 463–465
defined, 459
overview, 459–461
terminal services, 461
virtual desktop infrastructure, 461
VMware's Horizon View, 462
destination, 35
device address, 101
DHCP (Dynamic Host Configuration Protocol)
automatic private IP addressing, 171
basic configuration process, 157–158
configuration information provided by, 156
configuring Windows client for, 169–171
DHCP servers, 156–157, 162–169
lease duration, 159, 161–162
overview, 70, 155–156
releasing leases, 171
renewing leases, 171
scopes
exclusions, 160–161
overview, 158–159
reservations, 161
subnets and, 159–160
VLANs and, 159–160
DHCP Ack message, 158
DHCP console, Windows Server 2025, 494

DHCP Discover message, 157
DHCP Offer message, 157
DHCP Request message, 158
DHCP servers
  configuring, 718–719
  configuring scope, 163–169
  installing
    on Linux, 718
    on Windows Server 2025, 162–163
  overview, 156–157, 263, 717
  performance, 263
  security, 263
  starting, 720
dhcpd.conf file, 718
diagramming tools, 763
differential backups, 309
digest, email, 851
digital certificates, 73
digital subscriber line (DSL), 250–251
direct attached storage (DAS), 291–292
directories, Linux
  directory-handling commands, 730
  home directory, 682
  root directory, 682
  top-level, 681–682
directory services, 72
disaster planning, 862
disaster recovery, virtualization, 279–280
disasters
  analyzing impact of, 865–866
  business continuity plan, 866–867
  deliberate disasters, 863–864
  disruption of services, 864
  environmental disasters, 863
  equipment failure, 864
  fire drill, 867–868
  overview, 861–862
disk drives
  form factors, 286
  hard disk drives, 285
  solid state drives, 285–286
disk storage, 82
Display options, RDC, 778–779
Distributed File System (DFS), 71
distribution frames, 336–337
distribution groups
  Exchange Online, 626
  Windows Server 2025, 520
distribution layer, 243
distributions, Linux, Linux, 659–661
dnf command, 718, 739
DNS (Domain Name System)
  A records, 195–196
  CNAME records, 196–197
  configuring Windows DNS Client, 204–205
  defined, 70, 105, 112, 122
  DNS tree, 175–176
  domain names, 174–176
  fully qualified domain names, 176
  Hosts file, 180–183
  MX records, 197
  NS records, 195
  overview, 173–174
  parent domain, 175
  PTR records, 197
  queries, 190–191
  relative names, 177
  resource records, 192–194
  reverse lookup zones, 198–199
  root domain, 174
  SOA records, 194–195
  subdomains, 175
  top-level domains
    country code domains, 179–180
    defined, 177
    generic domains, 178
  URLs and, 176–177
  Windows DNS Server, 199–204
  zone files, 192–194
DNS console, Windows Server 2025, 494
DNS databases, 192–194
DNS servers
  BIND
    editing configuration files, 721
    installing, 720–721
    named.conf file, 721–723
    restarting, 724
    zone files, 723–724
  caching, 190
  defined, 183
  overview, 183–184
  primary servers, 186–187

root servers, 187–190
secondary servers, 186–187
Windows DNS Server
host records, creating, 203–204
overview, 199
zone, creating, 200–203
zones, 184–185
DnsAdmins group, Windows Server 2025, 526
DnsUpdateProxy group, Windows Server 2025, 526
document access, 20
dollar sign ($), 726, 728
domain accounts, 508
Domain Admins group, Windows Server 2025, 526
Domain Computers group, Windows Server 2025, 526
Domain Controller Security Policy console, Windows Server 2025, 494
domain controllers, 500
Domain Controllers group, Windows Server 2025, 526
Domain Guests group, Windows Server 2025, 526
domain local scope, 520
Domain Name System (DNS)
A records, 195–196
CNAME records, 196–197
configuring Windows DNS Client, 204–205
defined, 70, 105, 112, 122
DNS tree, 175–176
domain names, 174–176
fully qualified domain names, 176
Hosts file, 180–183
MX records, 197
NS records, 195
overview, 173–174
parent domain, 175
PTR records, 197
queries, 190–191
relative names, 177
resource records, 192–194
reverse lookup zones, 198–199
root domain, 174
SOA records, 194–195
subdomains, 175
top-level domains
country code domains, 179–180
defined, 177
generic domains, 178
URLs and, 176–177
Windows DNS Server, 199–204
zone files, 192–194
`domain` parameter
`Net Accounts` command, 565–566
`Net Group` command, 569
`Net Localgroup` command, 572
Domain Users group, Windows Server 2025, 526
domainname parameter, `Net View` command, 581
domains
Active Directory, 499–500, 502–503
joining Mac computer to, 372–374
joining Windows computer to, 363–366
Windows Server 2025, 473–474
dotted-decimal notation, 129–130
down, 10
drive interfaces
ATA, 287
IDE, 287
SCSI, 287
drive parameter, `Net Use` command, 578
DSL (digital subscriber line), 250–251
Dynamic Host Configuration Protocol (DHCP)
automatic private IP addressing, 171
basic configuration process, 157–158
configuration information provided by, 156
configuring Windows client for, 169–171
DHCP servers, 156–157, 162–169
lease duration, 159, 161–162
overview, 70, 155–156
releasing leases, 171
renewing leases, 171
scopes
exclusions, 160–161
overview, 158–159
reservations, 161
subnets and, 159–160
VLANs and, 159–160
dynamic ports, 141
dynamic routes, 150
dynamically expanding disk, 395

E

eavesdroppers, wireless networking, 346
EBCDIC (Extended Binary Coded Decimal Interchange Code), 104

edge, firewall, 257
email
  phishing, 849
  traditional vs. cloud computing, 82
email alias, 630–631
emergency disk capacity, 63
encryption, 61, 73
Enhanced Encryption option, Microsoft Teams Admin Center, 651
enterprise router, 255
/env parameter, runas command, 582
environmental disasters, 863
environmental variables, 560–562, 728
ephemeral ports, 141
EPL (Ethernet private line), 260
Equifax, 869
equipment loss policy, 797
eradication phase, CIRP, 878
error messages, 805
ESXi hypervisor, 414
Ethernet
  2.5GBase-T, 109
  5GBase-T, 109
  10Base2 Ethernet cable, 108–109
  10Base5 Ethernet cable, 108
  10Base-T Ethernet cable, 109
  10GBase-T, 110
  25GBase-T, 110
  40GBase-T, 110
  beyond gigabit, 109–110
  Fast Ethernet, 109
  Gigabit Ethernet, 109
  overview, 106–108
  Standard Ethernet, 108–109
Ethernet handoff, 147, 253
Ethernet packets
  collisions, 37–38, 43
  frame check sequence, 36
  payload, 36
  preamble, 36
  start-of-frame marker, 36
  tags, 36
Ethernet private line (EPL), 260
ethical hacking, 884–885
event logs, 474, 813–814
Event Viewer console, Windows Server 2025, 494
EventCreate command, 563–564
Exagrid, 299, 304
Exchange Online
  Admin Center, 626–629
  mailboxes
    apps, 635–636
    converting to shared mailbox, 634–635
    defined, 625–626
    delegating, 631–634
    email alias, 630–631
    forwarder, 636–637
    overview, 629–630
    shared, 626
  recipients, 625–626
  shared mailbox, 626, 637–640
exclamation mark (!), 727
exclusions, 160–161
executive leadership communication, 880–881
Experience options, RDC, 780–781
Extended Binary Coded Decimal Interchange Code (EBCDIC), 104
external interface, 148

F

fail-over connection, 258
false negatives, 851
false positives, 851
Fast Ethernet, 109
FDDI (Fiber Distributed Data Interface), 29
Fedora distribution, Linux, 659
Fedora Server, 663–671
Fiber Distributed Data Interface (FDDI), 29
fiber-optic internet, 252–253
file and folder permissions, 530
file servers, 264–265
file servers, Windows Server 2025
  granting permissions, 540–542
  New Share Wizard, 533–538
  overview, 533
  sharing folders without wizards, 538–540
file sharing, 11, 70–71
file system
  defined, 70–71
  Linux

browsing, 682–683
overview, 680–681
NFS, 105
VMFS, 414
File Transfer Protocol (FTP), 105, 122
file-based backups
copy backups, 307–308
daily backups, 308
differential backups, 309
full backups, 306–307
incremental backups, 308–309
overview, 297–298, 304–305
file-handling commands, Linux
`cat` command, 735–736
`cp` command, 733–734
`mv` command, 735
`rm` command, 734–735
`FileSystem` provider, PowerShell, 597–598
filters
group policy objects, 552–554
piping and, 560
spam
Bayesian analysis, 852
block list, 853
defined, 851
friends list, 853–854
graylisting, 854
Sender Policy Framework, 852–853
fire, 863
fire drill, 867–868
firewall appliance, 827–828, 833
firewall routers, 16
firewalls
application gateway, 832–833
best practices, 833–834
circuit-level gateway, 832
cybersecurity, 60
defined, 11, 16, 256, 827
firewall appliance, 256
packet filtering, 829–831
server computer, 256
stateful packet inspection, 831
WANs, 256–257
Windows Defender Firewall, 834–842
fish tape, 332
5G, 253
fixed-size disk, 395
flooding, 47–48, 863
`forcelogoff` parameter, `Net Accounts` command, 565
forests, Active Directory, 501–502
form factors, 286
blade servers, 78–79
defined, 77
rack-mounted servers, 78
tiny servers, 79
tower cases, 78
forward lookup, 198
forwarder, 636–637
forwarding, 46–47, 48
forwarding database, 45
fragment, URL, 177
frame check sequence, 36
frames, 35. *See also* packets
freeloaders,wireless networking, 345–346
friends list, 853–854
Frontline Deployment option, Microsoft Teams Admin Center, 651
FTP (File Transfer Protocol), 105, 122
full backups, 306–307
Full Control permission, Windows Server 2025, 530
full installation, 470–471
full-duplex transmission mode, OSI model session layer, 104
fully qualified domain names, 176
functions, PowerShell, 586

G

G Suite, 86
gateway routers
creating VPNs with, 148–149
defined, 148, 339
external interface, 148
internal interface, 148
gateways
application, 832–833
circuit-level, 832
generic domains, 178
`Get-ChildItem` cmdlet, PowerShell, 591–595, 598
`Get-Help` cmdlet, PowerShell, 589–591
`Get-Process` cmdlet, PowerShell, 599

Get-PSProvider cmdlet, PowerShell, 597
Get-Random command, PowerShell, 585, 588–589
get-rich-quick schemes, 850
Gigabit Ethernet, 109
global scope, 520
GNOME, 725
Google
  Google App Engine, 89
  Google Cloud, 89
  Google Workspace, 86, 88–89
GoToMyPC, 460
gpasswd command, 744
GPOs (group policy objects)
  creating, 545–552
  defined, 543
  filtering, 552–554
graphical user interface (GUI), 676
graylisting, 854
greater-than sign (>), 560, 727
group membership, 508
group policy
  configuring Windows Defender Firewall with, 836–842
  defined, 543
  enabling, 544–545
  group policy objects, 545–554
  overview, 543–544
Group Policy group, Windows Server 2025, 526
Group Policy Management console, Windows Server 2025, 494
group policy objects (GPOs)
  creating, 545–552
  defined, 543
  filtering, 552–554
groupadd command, Linux, 743
groupdel command, Linux, 743–744
groupname parameter
  Net Group command, 569
  Net Localgroup command, 571
groups
  Linux
    passwords, 744
    user accounts, 702
  Window Server 2025
    adding members to, 523–525
    creating, 522–523
    default groups, 521–522
    distribution groups, 520
    overview, 519
    scope, 520–521
    security groups, 520
guest operating system, 272
guests, virtualization, 272
Guests group, Windows Server 2025, 525
GUI (graphical user interface), 676

H

hackers, 869–870
HAL (Hardware Abstraction Layer), 273
half-duplex transmission mode, OSI model session layer, 104
hard disk drives (HDDs), 285
hard drives, 76
  partitions, 236
  size, 236
hardware. *See also* network hardware
  2.5GBase-T Ethernet cable, 109
  5GBase-T Ethernet, 109
  10Base2 Ethernet cable, 108–109
  10Base5 Ethernet cable, 108
  10Base-T Ethernet cable, 109
  10GBase-T Ethernet cable, 110
  25GBase-T Ethernet cable, 110
  40GBase-T Ethernet cable, 110
  100-Base-TX Fast Ethernet cable, 109
  Cat5e cable, 109
  Cat6 cable, 31
  distribution frames, 336–337
  hubs, 16, 32, 42, 97
  labeling, 796–797
  patch cables, 330, 805
  pinouts, 332–333
  RJ-45 connectors, 31, 333–335
  switches, 337–338
  for tape backups, 300–301
  twisted-pair cable
    categories, 327–328
    distribution frames, 336–337
    installing, 330–331
    pairs, 328
    patch panels, 335–336
    pinouts, 332–333

RJ-45 connectors, 333–335
server rooms, 336–337
tools, 331–332
unshielded, 329
wall jacks, 335–336
Hardware Abstraction Layer (HAL), 273
hardware assets, 792
hash mark (#), 726
hash table, 318
hashing algorithm, 319–320
HCI (hyperconverged infrastructure)
backups, 320–321
clusters, 316, 322
deduplication, 317–320
incorporating into plan, 323
issues addressed by, 314–315
nodes, 315–317
overview, 313–315
HDDs (hard disk drives), 285
`hidden` parameter, `Net Config` command, 567
hidden shares, 532
home directory, Linux, 682, 702
`home` parameter, `Net Use` command, 578
Horizon View, VMware
Horizon View Client, 462
Horizon View Composer, 462
Horizon View tool, 462
vCenter Desktop, 462
vSphere Desktop, 462
host ID, IP addresses, 129
`hostname` command, 208–209, 748
hosts, virtualization, defined, 272
Hosts file
DNS, 180–183
Linux, 713–714
hot-swappable components, 74, 78
HTTP (Hypertext Transfer Protocol), 122
hubs, defined, 16, 32, 42, 97
Hurwitz, Judith, 90
hyperconverged infrastructure (HCI)
backups, 320–321
clusters, 316, 322
deduplication, 317–320
incorporating into plan, 323
issues addressed by, 314–315
nodes, 315–317
overview, 313–315
Hypertext Transfer Protocol (HTTP), 122
hyperthreading, 76
Hyper-V
enabling, 395–396
installing operating system, 409–411
Manager Window
Actions pane, 397
Checkpoints pane, 397
Navigation Pane, 397
Virtual machine summary pane, 397
Virtual Machines pane, 397
New Virtual Hard Disk Wizard, 400–404
New Virtual Machine Wizard, 404–409
overview, 393–394
System Center Virtual Machine Manager, 268
virtual disks, 394–395
Virtual Switch Manager window
External switch, 398
Internal switch, 399
Private switch, 300
hypervisors
defined, 273
ESXi, 414
functions, 273
type-1, 274
type-2, 274

I

IaaS (Infrastructure as a Service), 425
`id` parameter, `Net File` command, 568
IDE (integrated device electronics) interface, 287
IDF (intermediate distribution frame):, 337
IEEE (Institute of Electrical and Electronics Engineers), 94–95
IETF (Internet Engineering Task Force), 95
`ifconfig` command, 715, 748–749
IIS_WPG group, Windows Server 2025, 526
image-based backups, 298, 310
Inbound Message Details card, Exchange Admin Center, 628
incremental backups, 308–309
Information Systems Audit and Control Association (ISACA), 64

infrastructure
defined, 24
elements of, 24
HCI
backups, 320–321
clusters, 316, 322
deduplication, 317–320
incorporating into plan, 323
issues addressed by, 314–315
nodes, 315–317
overview, 313–315
LANs
broadcast packets, 38
collisions, 37–38
MAC addresses, 34–35
network cables, 30–33
network interfaces, 33
network topology, 26–30
packets, 35–37
protocols, 25–26
standards, 25–26
wireless networking, 39
WANs
defined, 18–19
Ethernet private line, 260
internet connection, 250–253, 257–258
routers, 253–255
securing connections for remote users, 258–259
VPN, 259
wireless networking
connecting to, 342–343
overview, 39
security, 343–352
troubleshooting, 352–355
WAP, 339–342
Infrastructure as a Service (IaaS), 425
in-line deduplication, 317–318
input mode, Vi text editor, 693
installing
BIND DNS server, 720–721
Cockpit, 699
DHCP servers
on Linux, 718
on Windows Server 2025, 162–163
Linux
distributions, 659–661
Fedora Server, 663–671
installation program, 672
system requirements, 659–660
TCP/IP configuration, 662–663
on VM, 662
switches, 337–338
twisted-pair cable, 330–331
Windows Server 2025
backing up, 474
checking system requirements, 470
disconnecting UPS devices, 475
domains, 473–474
event logs, 474
full installation, 470–471
licensing options, 471
multiboot features, 471–472
overview, 469
partitions, 472
release notes, 470
server roles and features, 482–485
Setup program, 475–481
TCP/IP configuration, 473
updates, 475
upgrade installation, 471
workgroups, 473–474
wireless access points (WAPs), 339–340
Institute of Electrical and Electronics Engineers (IEEE), 94–95
integrated device electronics (IDE) interface, 287
intentional damage, 863
interface, 34
interface configuration files, 711–713
intermediate distribution frame (IDF), 337
internal communication, 879–880
internal interface, 148
internal routers, 150
International Organization for Standardization (ISO), 95
International Society of Automation (ISA), 64
internet. *See also* TCP/IP
accessing, 10
connecting to, 146–148
defined, 112, 116–117

history of, 117–118
overview, 10
internet connection
cloud computing and, 84–85
WANs
broadband connections, 250–251
cellular network, 253
fiber-optic cable, 252–253
providing redundancy for, 257–258
T1 lines, 251–252
Internet Engineering Task Force (IETF), 95
internet gateway, 146
Internet Information Services (IIS) Manager console, Windows Server 2025, 494
Internet Protocol (IP), 25, 110–111
internet service providers (ISPs), 146
internet socket, 141
intruders, wireless networking, 344–345
IP addresses
binary
counting by ones, 123–126
logical operations, 126–127
Windows Calculator, 127–128
classifying
Class A addresses, 133
Class B addresses, 133
Class C addresses, 134
loop-back address, 134
multicast address, 132
overview, 130–133
dotted-decimal notation, 129–130
host ID, 129
IPv4, 129, 131–132
IPv6, 129, 131–132
looking up, 215–216
network address translation, 141–144
network ID, 129
overview, 128–129
ports, defined, 141
subnetting
default subnet masks, 137–138
IP block parties, 139–140
network prefix notation, 137
overview, 134
private addresses, 140–141
public addresses, 140–141
restrictions, 138
subnet masks, 136–137
subnets, 135, 138–139
IP block parties, 139–140
IP Helper, 159
IP lease
releasing, 211
renewing, 211
IP spoofing, 831
`ipgconfig` command
displaying basic IP configuration, 209
displaying detailed configuration information, 210
flushing local DNS cache, 212
releasing IP lease, 211
renewing IP lease, 211
IPv6 (IP next generation [IPing]), 131–132
IPX/SPX, 114
ISA/IEC 62443 standard, 64
ISO (International Organization for Standardization), 95
ISO/IEC 27001, 64
ISPs (internet service providers), 146
ITAM (IT asset management)
asset identifier, 796
asset-tracking software, 798–799
cellphone vendor portals, 799–800
copier vendor portals, 800
equipment loss policy, 797
identifying and listing assets, 793
importance of, 793
labeling hardware, 796–797
overview, 792–793
photographing assets, 795
router management page, 800
software, 798
switch management page, 800
what to track, 794–795

K

Kali tool, 891–893
key pair, 449
keyboard, video, and mouse (KVM) switch, 78–79
keyboard shortcuts, for Remote Desktop, 776–777
keyhole saw, 332
Kirsch, Daniel, 90
known vulnerabilities, 886

L

landing zone, 304
LANs (local area networks)
  anticipating growth areas, 247
  backups, 247
  broadcast packets, 38
  collisions, 37–38
  computer requirements, 235–237
  defined, 10, 18
  diagrams, 245–246
  disaster recovery, 247
  MAC addresses, 34–35
  network cables
    Cat5e cable, 31
    hubs, 32
    patch cables, 31–32
    patch panels, 31–32
    planning, 238–239
    repeaters, 32
    RJ45 connectors, 31
    switches, 32–33
    twisted-pair cable, 30–31
  network interfaces, 33
  network plan, 233–234
  network topology
    bus topology, 26–27
    defined, 26
    mesh topology, 30
    nodes, 26
    packets, 26
    planning, 242–243
    ring topology, 29
    star topology, 27–28
  overview, 24
  packets, 35–37
  protocols, 25–26
  purpose, 234–235
  security, 247
  standards, 25–26
  star topology, 27–28
  switches, 240–242
  system information program, 237–238
  TCP/IP implementation, 244–245
  wireless networking, 39
learning, 45–46, 48
leases, DHCP
  duration, 159, 161–162
  releasing, 171
  renewing, 171
lightning, 863
Linear Tape-Open (LTO) backup, 299–300
link aggregation, 241
LinkedIn, 886
Linksys Hydra Pro 6E router, 254
Linux
  BIND DNS server
    editing configuration files, 721
    installing, 720–721
    named.conf file, 721–723
    restarting, 724
    zone files, 723–724
  Cockpit, 698–701
  commands
    administering users, 739–745
    cat command, 735–736
    cd command, 730–731
    change command, 741–742
    chgrp command, 746
    chmod command, 746–747
    chown command, 745
    cp command, 733–734
    directory-handling commands, 730–733
    dnf command, 739
    editing commands, 726
    environmental variables, 728
    file-handling commands, 733–736
    gpasswd command, 744
    groupadd command, 743
    groupdel command, 743–744
    hostname command, 748
    ifconfig command, 715, 748–749
    ls command, 732–733
    mkdr command, 731
    mv command, 735
    netstat command, 749–750
    networking, 747–753
    newusers command, 742–743
    ownership and permissions, 745–747
    packages, 738–739
    passwd command, 742

ping command, 750–751
piping, 727–728
pwd command, 730
redirection, 727–728
rm command, 734–735
rmdir command, 731
with root-level privileges, 729–730
route command, 752
service command, 737
services, 737–738
shell scripts, 728–729
shells, 723–724
sudo command, 680
traceroute command, 752–753
useradd command, 739–740
userdel command, 741
usermod command, 741
wildcards, 726–727
yum command, 738–739
configuring for networking
Cockpit, 705–709
ifconfig command, 714–715
network configuration files, 710–713
consoles
defined, 676–677
remote consoles, 678–679
virtual consoles, 672, 677
DHCP server
configuring, 718–719
installing, 718
overview, 717
starting, 72
directories
home directory, 682, 702
root directory, 682
top-level, 681–682
file system
browsing, 682–683
overview, 680–681
installing
distributions, 659–661
Fedora Server, 663–671
installation program, 672
system requirements, 659–660
TCP/IP configuration, 662–663
on VM, 662
logging in, 673–675
logging out, 675
mount points, 662, 681
package managers
DEB, 684
defined, 683–684
installing packages, 686–687
listing packages, 684–686
removing packages, 687–688
RPM, 684
updating packages, 688–689
root user, 679
shells, 676–677, 702, 725–726
shutting down, 675
sudo command, 680
user accounts, 702–704
Vi text editor
changing text, 696
deleting text, 695
exiting, 692
inserting text, 694–695
moving cursor, 693–694
operating modes, 692–693
overview, 689–690
repeating commands, 697
saving changes, 692
starting, 690–692
yanking and putting, 696–697
X Window System, 672
List Folder Contents permission, Windows Server 2025, 530
local, 10
local accounts, 508
local area networks (LANs)
anticipating growth areas, 247
backups, 247
broadcast packets, 38
collisions, 37–38
computer requirements, 235–237
defined, 10, 18
diagrams, 245–246
disaster recovery, 247
MAC addresses, 34–35
network cables
Cat5e cable, 31
hubs, 32

local area networks (LANs) *(continued)*
patch cables, 31–32
patch panels, 31–32
planning, 238–239
repeaters, 32
RJ45 connectors, 31
switches, 32–33
twisted-pair cable, 30–31
network interfaces, 33
network plan, 233–234
network topology
bus topology, 26–27
defined, 26
mesh topology, 30
nodes, 26
packets, 26
planning, 242–243
ring topology, 29
star topology, 27–28
overview, 24
packets, 35–37
protocols, 25–26
purpose, 234–235
security, 247
standards, 25–26
star topology, 27–28
switches, 240–242
system information program, 237–238
TCP/IP implementation, 244–245
wireless networking, 39
local copy, backups, 296–297
local disk storage, 276
Local Resources options, RDC, 779–780
local user profiles, 525
LocalTalk, 114
Locations option, Microsoft Teams Admin Center, 651
logical addressing, 100–101
logical link control (LLC) layer, 108
logical operations, 126–127
logon hours, 514–515
logon script, 516, 527
loop-back address, 134
lower layers, OSI model, 96
`ls` command, 732–733
LTO (Linear Tape-Open) backup, 299–300

M

MAC (media access control), 33
MAC address, 33–35, 43
MAC address filtering, 350–351
Mac computers
configuring network connections, 368–371
joining domain, 372–374
sharing network, 374–376
MAC spoofing, 351
Maersk, 869–870
Mail Exchange (MX) records, 197
Mail Flow card, Exchange Admin Center, 628
mail servers, 218–219, 263–264
mailboxes, Exchange Online
apps, 635–636
converting to shared mailbox, 634–635
defined, 625–626
delegating, 631–634
email alias, 630–631
forwarder, 636–637
overview, 629–630
shared, 626
Mailboxes card, Exchange Admin Center, 628
mail-enabled security group, 626
main distribution frame (MDF), 336
mainframe computers, 19
managed switches, 51–52, 241
ManageEngine, 268
mandatory user profiles, 526
MANs (metropolitan area networks), 19
master files, 192–194
Matrix movies, 7
`maxpwage` parameter, `Net Accounts` command, 565
MDF (main distribution frame), 336
media access control (MAC) addresses, 33–35, 43
medium access control (MAC) layer, 108
Meetings option, Microsoft Teams Admin Center, 651
members, Windows Server 2025, 523–525
memory
amount of, 236
servers, 76
mesh topology, 30
`message` parameter, `Net Helpmsg` command, 570–571
Messaging option, Microsoft Teams Admin Center, 651
metropolitan area networks (MANs), 19

Microsoft
certifications, 767
cloud offerings, 89
software licenses, 798
Microsoft 365
creating new users, 615–621
disabling users, 622–623
overview, 606–607
plans, 608–610
resetting user passwords, 621–622
tenants, 605, 611–615
Microsoft 365 Group, 626, 646–647
Microsoft Management Console (MMC)
Active Directory Domains and Trusts console, 493
Active Directory Sites and Services console, 493
Active Directory Users and Computers console, 493
Component Services console, 493
Computer Management console, 493–494
consoles, 493–494
customizing, 495–496
DHCP console, 494
DNS console, 494
Domain Controller Security Policy console, 494
Event Viewer console, 494
Group Policy Management console, 494
Internet Information Services (IIS) Manager console, 494
ODBC Data Sources console, 494
overview, 491–494
Performance Monitor console, 494
Services console, 494
Microsoft Office 365, overview, 86, 89
Microsoft Teams
architecture of
Admin Center, 650–652
Microsoft 365 Group, 646–647
OneDrive for Business, 650
overview, 645–646
SharePoint, 648–649
channels, 642–643
managing, 652–655
overview, 608, 641
tabs, 643–645
teams, 642
Microsoft Volume Licensing, 798
Migration Batch Report card, Exchange Admin Center, 628
MILNET, 113, 118
`minpwage` parameter, `Net Accounts` command, 565
`minpwlen` parameter, `Net Accounts` command, 565
mirror, 71
`mkdr` command, 731
MMC (Microsoft Management Console)
Active Directory Domains and Trusts console, 493
Active Directory Sites and Services console, 493
Active Directory Users and Computers console, 493
Component Services console, 493
Computer Management console, 493–494
consoles, 493–494
customizing, 495–496
DHCP console, 494
DNS console, 494
Domain Controller Security Policy console, 494
Event Viewer console, 494
Group Policy Management console, 494
Internet Information Services (IIS) Manager console, 494
ODBC Data Sources console, 494
overview, 491–494
Performance Monitor console, 494
Services console, 494
Modify permission, Windows Server 2025, 530
motherboard, servers, 75–76
mount points, 662, 681
multiboot, 471–472
multicast address, 132
multifactor authentication, 60, 823
multifunction wireless routers, 339
multiport repeaters, 97
multiprocessing, 71
multitasking, 71
`mv` command, 735
MX (Mail Exchange) records, 197

N

name server (NS) records, 195
`named.conf` file, 721–723
NAS (network-attached storage), 276, 293–294
NAS device, 298
NAT (network address translation), 54, 141–144
National Institute for Standards and Technology (NIST), 820

National Institute of Standards and Technology (NIST), 64
National Science Foundation (NSF), 113
native commands, PowerShell, 586
nbtstat command, 212–213
NDRs (nondelivery reports), 848
Net Accounts command, 565
Net Computer command, 566
Net Config command, 566–567
Net Continue command, 567
Net File command, 568
Net Group command, 568–570
Net Help command, 570
Net Helpmsg command, 570–571
Net Localgroup command, 571–572
Net Pause command, 572–573
Net Session command, 573
Net Share command, 574–575
Net Start command, 575
Net Statistics command, 575–576
Net Stop command, 576–577
Net Time command, 577
Net Use command, 577–578
Net User command, 579–580
Net View command, 580–581
NetBEUI (Network BIOS Extended User Interface), 114
NetBIOS (Network Basic Input/Output System), 114
netlogon service
  Net Continue command, 567
  Net Pause command, 567
/netonly parameter, runas command, 582
netstat command
  Linux, 749–750
  TCP/IP
    displaying connections, 213–214
    displaying interface statistics, 214
network
  defined, 8
  illustration of typical, 9
  lingo, 10
  uses of, 10–13
network adapters, 97
network address, 101
network address translation (NAT), 54, 141–144
network administration
  ITAM (IT asset management)
    asset identifier, 796
    asset-tracking software, 798–799
    cellphone vendor portals, 799–800
    copier vendor portals, 800
    equipment loss policy, 797
    identifying and listing assets, 793
    importance of, 793
    labeling hardware, 796–797
    overview, 792–793
    photographing assets, 795
    router management page, 800
    software, 798
    switch management page, 800
    what to track, 794–795
network administrators
  applying software patches, 762
  assigning, 21
  bluffs and excuses, 769–770
  certifications, 765–768
  choosing, 759–760
  clean-up, 761
  defined, 21
  gurus and, 768–769
  managing users, 761–762
  overview, 757–758
  routine chores, 761
  software tools, 763–764
  staying up-to-date, 764–765
overview, 757
RDC (Remote Desktop Connection)
  configuring options for, 777–782
  connecting to server, 774–776
  defined, 771
  enabling, 772–773
  keyboard shortcuts, 776–777
Remote Assistance
  accepting invitation, 787–789
  defined, 771
  enabling, 783–784
  overview, 782
  requesting, 784–786
troubleshooting
  basic steps, 802–803
  booting in Safe Mode, 809
  checking network settings, 806

dead computers, 803–804
error messages, 805
event logs, 813–814
narrowing down possibilities, 806–807
network connections, 804–805
overview, 801–802
restarting client computer, 808–809
restarting network server, 812–813
System Restore, 809–811
tracking log, 814
viewing current network users, 807
network administrators
applying software patches, 762
assigning, 21
bluffs and excuses, 769–770
certifications
Cisco, 768
CompTIA, 766–767
Microsoft, 767
overview, 765–766
choosing, 759–760
clean-up, 761
defined, 21
gurus and, 768–769
managing users, 761–762
overview, 757–758
routine chores
backups, 761
security, 761
software tools
diagramming tool, 763
network built-in tools, 763
network discovery program, 763
Network Monitor, 764
network scanning tool, 763
protocol analyzer, 763
System Information program, 763
TCP/IP diagnostic commands, 764
staying up-to-date, 764–765
Network Basic Input/Output System (NetBIOS), 114
Network BIOS Extended User Interface (NetBEUI), 114
network built-in tools, 763
network cables, 8, 15. *See also* cables
network configuration files
Hosts file, 713–714
interface configuration files, 711–713
Network file, 710–711
overview, 710
resolv.conf file, 714
Network Configuration group, Windows Server 2025, 525
network connections
configuring for Mac, 368–371
configuring for Windows, 357–363
troubleshooting, 804–805
network discovery programs, 763
Network file, Linux, 710–711
Network File System (NFS), 105
network hardware
distribution frames, 336–337
server rooms, 336–337
switches, 337–338
twisted-pair cable
categories, 327–328
installing, 330–331
pairs, 328
patch panels, 335–336
pinouts, 332–333
plenum cable, 329
RJ-45 connectors, 333–335
shielded, 329
solid cable, 330
stranded cable, 330
switches, 337–338
tools, 331–332
unshielded, 329
wall jacks, 335–336
network ID, IP addresses, 129
network interface card (NIC), 33, 97
network interface layer, TCP/IP, 121
network interfaces
defined, 8, 15, 33
IP addresses
dynamic ports, 141
ephemeral ports, 141
internet socket, 141
private ports, 141
registered ports, 141
well-known ports, 141
servers, 77

network layer
OSI model
logical addressing, 100–101
overview, 25
routing, 102
TCP/IP, 121
Network Monitor, 764
network plan, 233–234
network prefix notation, 137
network printers
accessing with web interface, 382–383
adding, 378–382
configuring, 377–383
network protocols. *See* protocols
network scanning tools, 763
network servers, restarting, 812–813
network services
key Windows services, 812
restarting, 811–812
server operating system, 70
network topology
access layer, 243
bus topology, 26–27
core layer, 243
defined, 26
distribution layer, 243
mesh topology, 30
nodes, 26
packets, 26
planning, 242–243
ring topology, 29
star topology, 27–28
network transmission speed, 107
network virtualization, 277–278
Network+ certification, 767
network-attached backup appliances, 298–299, 303–304
network-attached storage (NAS), 276, 293–294
networking
clients, 13
defining a network, 8–9
end of personal computer, 19–20
file sharing, 11
firewalls, 11
hubs, 16
LANs, 18
Linux commands for
`hostname` command, 748
`ifconfig` command, 748–749
`netstat` command, 749–750
`ping` command, 750–751
`route` command, 752
`traceroute` command, 752–753
mainframe computers, 18
MANs, 19
network administrator, 21
network cable, 15
network interfaces, 15
router, 16
servers
dedicated, 14
defined, 13
setups, 17–18
shared resources, 11–12
switch, 16
terminology, 10
WANs, 18–19
WAPs, 16
wireless networks, 39
New Share Wizard, Windows Server 2025, 533–538
`newusers` command, 742–743
NFS (Network File System), 105
NIST (National Institute for Standards and Technology), 820
NIST Cybersecurity Framework
Framework Core section, 65–67
Framework Organizational Profiles section, 66
Framework Tiers section, 66
overview, 64–65
nmap, 763, 886
nodes
defined, 10, 26
HCI, 315–317
Non-Accepted Domain card, Exchange Admin Center, 628
Non-Delivery Report card, Exchange Admin Center, 628
nondelivery reports (NDRs), 850
`/noprofile` parameter, `runas` command, 581
`NOT` logical operation, 126
Notifications & Alerts option, Microsoft Teams Admin Center, 652
NS (name server) records, 195

NSF (National Science Foundation), 113
NSFNET, 113, 118
nslookup command
  displaying DNS records, 217–218
  locating mail server for email address, 218–219
  looking up IP address, 215
  subcommands, 215–216
NTFS file system, 472

O

objects
  Active Directory, 498–499
  PowerShell, 593
ODBC Data Sources console, Windows Server 2025, 494
offline, 10
offline copy, backups, 296–297
off-site copy, backups, 296–297
on the network, 10
OneDrive for Business, 650
online, 10
on-premises antispam, 854–855
Open Systems Interconnection (OSI) reference model
  application layer, 105
  data link layer, 98–99
  network layer, 100–102
  overview, 25–26, 95–96
  physical layer, 96–98
  presentation layer, 104
  session layer, 103–104
  transport layer, 102–103
operating system, 236
OR logical operation, 126
organizational units, 500–501, 503–505
OSI model
  application layer, 105
  data link layer, 98–99
  network layer, 100–102
  overview, 25–26, 95–96
  physical layer, 96–98
  presentation layer, 104
  session layer, 103–104
  transport layer, 102–103
Outbound Message Details card, Exchange Admin Center, 628
ownership, Linux commands for, 745–747

P

PaaS (Platform as a Service), 86–87, 425
package managers, Linux
  DEB, 684
  defined, 683–684
  installing packages, 686–687
  listing packages, 684–686
  removing packages, 687–688
  RPM, 684
  updating packages, 688–689
packet filtering, 829–831
packet sniffers, 763
packets
  broadcast packets, 38, 43, 51
  collisions, 37–38, 43
  defined, 26
  destination, 35
  Ethernet, 36–37
  OSI model layers, 105–106
  senders, 35
parameters, PowerShell, 587–589
parent domain, 175
parent partition, 394
partitions, 394, 472
passwd command, 742
password parameter, Net Use command, 578
passwords
  best practices, 818–821
  cracking, 886
  groups, Linux, 744–745
  Linux user accounts, 742
  Microsoft 365, 621–622
  resetting, 517–518
  server operating system, 73
  strong, 60
  Windows Server 2025 user accounts, 508
  wireless network, 348
patch cables, 31–32, 330, 805
Patch Manager Plus, 268
patch panels, 24, 31–32, 335–336
path, URL, 177
path parameter, Net Share command, 574
pathping command, 219–220
payload, 36
peer-to-peer networks, 14–15

penetration testing
confidentiality, 889
establishing boundaries, 889
ethical hacking, 884–885
executives' accounts, 889
Kali tool, 891–893
overview, 883–884
red team, 885
scoping, 888
sensitive systems or data, 889
steps in, 885–888
what to expect from, 893
per-device CALs, 471
Performance Monitor console, Windows Server 2025, 494
perimeter, firewall, 257
permissions
Linux commands for, 745–747
Windows Server 2025
access control list, 529–530
file and folder permissions, 530
granting, 540–542
special permissions, 531
`persistent` parameter, `Net Use` command, 578
personal computers, 19–20
per-use CALs, 471
petabte, 295
phish testing, 825
phishing emails, 849, 887
physical address, 34, 98
physical infrastructure, 24
physical layer, OSI model, 25, 96–98
physical security, 61
`ping` command, 220–221, 750–751
pinouts, 332–333
pipe character (|), 560, 595, 728
pipeline, 594
piping, 559–560, 593–597, 727–728
Planning option, Microsoft Teams Admin Center, 652
Platform as a Service (PaaS), 86–87, 425
plenum cable, 329
plenum space, 329
Pointer (PTR) records, 197
Policy Packages option, Microsoft Teams Admin Center, 652
ports
defined, 33
dynamic, 141
ephemeral, 141
IP addresses, 141
private, 141
registered, 141
small form-factor pluggable, 50
TCP/IP, 830
well-known, 141
post-process deduplication, 318
power supply, servers, 77
PowerShell
aliases, 591–593
cmdlets, 586, 587
common parameters, 588–589
`FileSystem` provider, 597–598
functions, 586
`Get-ChildItem` cmdlet, 591–595, 598
`Get-Process` cmdlet, 599
`Get-PSProvider` cmdlet, 597
help information, 589–591
native commands, 586
overview, 583
parameters, 587–589
piping technique, 593–597
providers, 597–598
scripts, 586, 598–601
`Select-Object` cmdlet, 595–596, 599
`Set-ExecutionPolicy` cmdlet, 598–601
`Sort-Object` cmdlet, 599
using, 584–585
preamble, 36
present working directory, 730
presentation layer, OSI model, 25, 104
prevention measures, cybersecurity
anti-spam software, 60
auditing, 61
data protection, 61
encryption, 61
firewalls, 60
multifactor authentication, 60
overview, 59–60
passwords, 60
physical security, 61

user life-cycle management, 61
user training, 61
Wi-Fi security, 60
primary partitions, 472
primary servers, DNS, 186–187
primary zones, 186, 200
Print Operators group, Windows Server 2025, 525
print servers, 265
printers, 236. *See also* network printers
private addresses, 140–141
private clouds, 87–88
private ports, 141
processors (CPUs)
clock speed, 76, 235
hyperthreading, 76
server computers, 75–76
`/profile` parameter, `runas` command, 591
program sharing, 12–13
properties, PowerShell, 594
properties (attributes), Active Directory objects, 499
protocol analyzers, 763
protocol suites, 94
protocols
AppleTalk, 114
defined, 24, 93
Ethernet
2.5GBase-T, 109
5GBase-T, 109
10Base2 Ethernet cable, 108–109
10Base5 Ethernet cable, 108
10Base-T Ethernet cable, 109
10GBase-T, 110
25GBase-T, 110
40GBase-T, 110
beyond gigabit, 109–110
Fast Ethernet, 109
Gigabit Ethernet, 109
overview, 106–108
Standard Ethernet, 108–109
IPX/SPX, 114
NetBEUI, 114
NetBIOS, 114
overview, 93–94
SNA, 114
providers, PowerShell, 597–598
PTR (Pointer) records, 197
public addresses, 140–141
public clouds, 87–88
`pwd` command, 730

Q

queries, DNS, 190–191
query, URL, 177

R

rack-mounted servers, 78
RAID (Redundant Array of Inexpensive Disks)
overview, 275–277
RAID 5 array, 276, 289–290
RAID 6 array, 291
RAID 10 array, 275, 288–289
ransomware, 57, 869–870
RAS and IAS Servers group, Windows Server 2025, 526
RDC (Remote Desktop Connection)
configuring options for
Advanced options, 781–782
Display options, 778–779
Experience options, 780–781
Local Resources options, 779–780
overview, 777–778
connecting to server, 774–776
defined, 771
enabling, 772–773
keyboard shortcuts, 776–777
overview, 386
virtual machines, 462
VMs and, 462
VPNs and, 386
Windows Server 2025, 488–491
Read & Execute permission, Windows Server 2025, 530
Read permission, Windows Server 2025, 530
rebooting
client computer, 808–809
Linux server, 670
VMs, 455
Windows Server 2025, 482
Recent Alerts Report card, Exchange Admin Center, 628
recipients, Exchange Online, 625–626
recovery measures, cybersecurity, overview, 59

red team, 885
redirection, 559–560, 727–728
Redundant Array of Inexpensive Disks (RAID)
  overview, 275–277
  RAID 5 array, 276, 289–290
  RAID 6 array, 291
  RAID 10 array, 275, 288–289
ReFS file system, 472
registered ports, 141
relative names, 176
release notes, 470
reliability
  cloud computing, 83
  servers, 74
`remark` parameter, `Net Share` command, 574
remote, 10
Remote Assistance
  accepting invitation, 787–789
  defined, 771
  enabling, 783–784
  overview, 782
  requesting, 784–786
remote consoles, Linux, 678–679
Remote Desktop Connection (RDC)
  configuring options for
    Advanced options, 781–782
    Display options, 778–779
    Experience options, 780–781
    Local Resources options, 779–780
    overview, 777–778
  connecting to server, 774–776
  defined, 771
  enabling, 772–773
  keyboard shortcuts, 776–777
  overview, 386
  virtual machines, 462
  VMs and, 462
  VPNs and, 386
  Windows Server 2025, 488–491
Remote Desktop Users group, Windows Server 2025, 525
repeaters
  defined, 32
  multiport, 97
replication, 500
Replicator group, Windows Server 2025, 525
Request for Comments (RFCs)
  April Fools, 119
  best current practices, 119
  experimental specifications, 119
  historic specifications, 119
  informational specifications, 119
  Internet Standards Track, 119
  for key Internet standards, 120
  maturity levels for, 120
  overview, 118–120
reservations, 161
residential gateway, 146
resolv.conf file, 714
resource records, 192–194
resources, Exchange Online, 626
retention zone, 304
reverse lookup zones, 198–199
RFCs (Request for Comments)
  April Fools, 119
  best current practices, 119
  experimental specifications, 119
  historic specifications, 119
  informational specifications, 119
  Internet Standards Track, 119
  for key Internet standards, 120
  maturity levels for, 120
  overview, 118–120
ring topology, 29
RJ-45 connectors, 31, 333–335
`rm` command, 734–735
`rmdir` command, 731
roaming user profile, 526–527
rogue access points, wireless networking, 346–347
root directory, Linux, 682
root domain, 174, 174–176
root servers, 187–190
root user, 679
root volume, 450
root-level privileges, 729–730
`route` command, 752
  displaying routing table, 222–225
  modifying routing table, 225–226
routers
  connecting remote locations, 148–149
  connecting to internet, 146–148
  defined, 8, 16, 24, 52, 145

firewall, 16
gateway, 148
internal, 150
linking subnets, 149–150
management page, 800
overview, 52–53, 145
routing tables, 150–152
switches vs., 52–53
WANs
cellular router, 255
enterprise router, 255
overview, 253–255
small office router, 254
wireless, 8, 16
routing, 102, 145
routing protocols, 150
routing tables
defined, 150
displaying, 222–225
dynamic routes, 150
modifying, 225–226
overview, 150–152
static routes, 150
RPM package manager, Linux, 684
`RunAs` command, 581–582

S

SaaS (Software as a Service), 86, 425
Safe Mode, 809
safe-computing practices, 846
SAN (storage area network), 276, 292–293
SAS (serial attached SCSI) interface, 287–288
SAS disk drives, 76
SATA (serial ATA) interface, 287
SATA hard drives, 76
`/savecred` parameter, `runas` command, 582
`savecred` parameter, `Net Use` command, 578
scalability
cloud computing, 83
servers, 74
scams, 850
scanning, 886
`schedule` service
`Net Continue` command, 567
`Net Pause` command, 567
scheme, URL, 177
scopes, DHCP
configuring, 163–169
exclusions, 160–161
overview, 158–159
reservations, 161
subnets and, 159–160
VLANs and, 159–160
scripts, PowerShell, 586, 598–601
SCSI (small computer system interface) interface, 287
secondary servers, DNS, 186–187
secondary zones, 186, 200
security
asset management system, 60, 862
backups, 310–311
business continuity planning, 862, 866–867
cloud computing, 85
cybersecurity incident response plan
close-out phase, 881–882
communication, 879–881
containing cybersecurity incident, 876–877
cybersecurity incident, 874–875
defined, 870
documentation. *See* twisted-pair cable
eradication phase, 878
framework, 870–871
hackers, 869–870
lessons learned, 882
names, 872
preparing, 872–873
rating scale, 876
response team, 873–874
restoring lost data and systems, 878–879
triaging reported incidents, 875–876
disasters
analyzing impact of, 865–866
deliberate disasters, 863–864
disruption of services, 864
environmental disasters, 863
equipment failure, 864
fire drill, 867–868
overview, 861–862
firewalls
application gateway, 832–833
best practices, 833–834

security *(continued)*
circuit-level gateway, 832
defined, 827
packet filtering, 829–831
stateful packet inspection, 831
Windows Defender Firewall, 834–842
frameworks
CIS Critical Security Controls, 64
COBIT, 64
ISA/IEC 62443, 64
ISO/IEC 27001, 64
NIST Framework, 64–67
overview, 63–64
network administrators and, 761
NIST Cybersecurity Framework
Framework Core section, 65–67
Framework Organizational Profiles section, 66
Framework Tiers section, 66
overview, 64–65
overview, 57–58
penetration testing
confidentiality, 889
establishing boundaries, 889
ethical hacking, 884–885
executives' accounts, 889
Kali tool, 891–893
overview, 883–884
red team, 885
scoping, 888
sensitive sytems or data, 889
steps in, 885–888
what to expect from, 893
prevention measures
anti-spam software, 60
auditing, 61
data protection, 61
encryption, 61
firewalls, 60
multifactor authentication, 60
overview, 59–60
passwords, 60
physical security, 61
user life-cycle management, 61
user training, 61
Wi-Fi security, 60
recovery measures
backups, 62–63
communications, 63
emergency disk capacity, 63
overview, 59, 61–62
spare computers, 63
for small businesses, 58–59
spam
ads for pornographic websites, 850
advertisements, 849
antispam software, 850–851, 854–858
backscatter, 850
defined, 848–849
filters, 851–854
get-rich-quick schemes, 850
minimizing, 858–859
overview, 847–848
phishing emails, 849
scams, 850
users
Administrator account, 821–822
authentication, 818
authorization, 818
cybersecurity policies, 824
multifactor authentication, 823
overview, 817
passwords, 818–820
phish testing, 825
training, 824–825
virus protection
antivirus programs, 844–846
overview, 842–844
safe-computing practices, 846
VPN, 387–388
WAP, 342
security groups, 520, 626
security incident response plan, 872
security outbreak script, 872
security services, 73
Security+ certification, 767
`Select-Object` cmdlet, PowerShell, 595–596, 599
Sender Policy Framework, 852–853
senders, 35
sensitive files, 20
serial ATA (SATA) interface, 287

serial attached SCSI (SAS) interface, 287–288
server architecture
  application servers, 266
  backup servers, 267
  connecting servers, 268–269
  database servers, 266
  deciding on number of servers needed, 261–262
  deployment servers, 267
  DHCP servers, 263
  domain controllers, 262–263
  file servers, 264–265
  mail servers, 263–264
  print servers, 265
  update servers, 267–268
  virtualization management platform, 268
  web servers, 266
server features, 482–485
server license, 471
Server Message Block (SMB), 104, 105
server operating system
  defined, 14
  digital certificates, 73
  directory services, 72
  encryption, 73
  file-sharing services, 70–71
  multitasking, 71
  network services, 70
  overview, 69–70
  passwords, 73
  security services, 73
Server Operators group, Windows Server 2025, 525
`server` parameter, `Net Config` command, 566
server roles, 482–485
server rooms, 336–337
`server` service
  `Net Continue` command, 567
  `Net Pause` command, 567
Server+ certification, 767
servers
  availability, 74
  connecting, 268–269
  dedicated, 14
  defined, 13
  DHCP, 156–157
  firewalls, 256
  form factors
    blade servers, 78–79
    overview, 77
    rack-mounted servers, 78
    tiny servers, 79
    tower cases, 78
  hard drives, 76
  maintenance, 20
  memory, 76
  motherboard, 75
  network interfaces, 77
  power supply, 77
  processor, 75–76
  reliability, 74
  restarting, 812–813
  scalability, 74
  service and support, 74
  video, 77
`service` command, 737
Service Set Identifiers (SSIDs), 341, 348–349
Services console, Windows Server 2025, 494
services disruption, 864
session layer, OSI model, 25, 103–104
sessions, 103
`Set-ExecutionPolicy` cmdlet, PowerShell, 598–601
SFP (small form-factor pluggable) ports, 50
SHA-1 algorithm, 320
shared mailbox
  converting to, 634–635
  creating, 637–640
  defined, 626
shared media, 42
shared resources, 11–12, 20
ShareName parameter, `Net Share` command, 574
SharePoint, 648–649
shares, Windows Server 2025
  defined, 531
  hidden shares, 532
  setting up, 532–533
  special shares, 532
shell scripts, 728–729
shells, Linux, 676–677, 702, 725–726
shielded twisted-pair cable (STP), 329
Simple Mail Transfer Protocol (SMTP), 105, 122
simplex transmission mode, OSI model session layer, 104

single-board computers, 79
Slackware distribution, Linux, 660
small computer system interface (SCSI) interface, 287
small form-factor pluggable (SFP) ports, 50
small office router, 254
`smartcard` parameter, `Net Use` command, 578
`/smartcard` parameter, `runas` command, 582
SMB (Server Message Block), 104, 105
SMTP (Simple Mail Transfer Protocol), 105, 122
SNA (Systems Network Architecture):, 114
SOA records, DNS, 194–195
software
  licenses, 12, 798
  subscription, 12
  tracking, 798
Software as a Service (SaaS), 86, 425
software assets, 792
software patches, 762
software tools
  diagramming tool, 763
  network built-in tools, 763
  network discovery program, 763
  Network Monitor, 764
  network scanning tool, 763
  protocol analyzer, 76
  System Information program, 763
  TCP/IP diagnostic commands, 764
solid cable, 330
solid-state drives (SSDs), 76, 285–286
`Sort-Object` cmdlet, PowerShell, 599
spam
  filters
    Bayesian analysis, 852
    block list, 853
    defined, 851
    friends list, 853–854
    graylisting, 854
    Sender Policy Framework, 852–853
  minimizing, 858–859
spare computers, 63
special permissions, 529, 531
special shares, 532
Spectre movie, 7
SPI (stateful packet inspection), 831
Spiceworks, 763
spinning drives, 285
spoilers, wireless networking, 346
`srvcomment` parameter, `Net Config` command, 567
SSDs (solid state drives), 285–286
SSIDs (Service Set Identifiers), 341, 348–349
stackable switches, 241
Standard Ethernet, 108–109
standards
  ANSI, 94
  defined, 24, 94
  IEEE, 94–95
  IETF, 95
  ISO, 95
  W3C, 95
star topology, 27–28
Star Wars movie, 79
Start Trek movie, 79
start-of-frame marker, 36
stateful packet inspection (SPI), 831
static IP addresses, 245
static routes, 150
station cable, 330, 335
storage
  attachment types
    DAS, 291–292
    NAS, 293–294
    SAN, 292–293
  disk drives
    form factors, 286
    hard disk drives, 285
    solid state drives, 285–286
  drive interfaces
    overview, 286
    SAS, 287–288
    SATA, 287
  planning disk capacity, 283–284
  RAID
    RAID 5 array, 289–290
    RAID 6 array, 291
    RAID 10 array, 288–289
storage area network (SAN), 276, 292–293
storage space, 20
StoreOne, 299
STP (shielded twisted-pair cable), 329
stub zones, 200
subdomains, 175
subnet masks

default, 137–138
defined, 136–137
subnets
defined, 135
linking with routers, 149–150
scopes and, 159–160
VLANs vs., 138–139
subnetting
IP addresses, 134–141
IP block parties, 139–140
network prefix notation, 137
overview, 134
private addresses, 140–141
public addresses, 140–141
restrictions, 138
subnet masks
default, 137–138
defined, 136–137
subnets
defined, 135
VLANs vs., 138–139
sudo command, Linux, 680, 729–730
supercomputers, 118
surge protector, 804
SUSE distribution, Linux, 659
switches
bridging, 49–50
broadcast domains, 51
collision domains, 48–49, 52
daisy-chaining, 241, 242
data link layer, OSI model, 99
defined, 8, 16, 24
flooding function, 47–48
forwarding function, 46–47, 48
installing, 337–338
learning function, 45–46, 48
link aggregation, 241
managed switches, 51–52, 241
management page, 800
overview, 32–33, 42–45, 240–242
routers vs., 52–53
small form-factor pluggable (SFP) ports, 50
stackable switches, 241
troubleshooting, 805
unmanaged switches, 241
uplinks, 50
System Center Virtual Machine Manager, 268
System Information program, 763
System Restore, 809–811
Systems Network Architecture (SNA):, 114

T

T1 lines, 251–252
T3 lines, 252
tabs, Microsoft Teams, 643–645
tags, 36
tape backups
cleaning magazines, 302–303
hardware, 300–301
LTO, 299–300
overview, 298
reliability, 301–302
tape library, 301
tape magazines, 301
target profile, cybersecurity, 66
targeted phishing, 886
TCP (Transmission Control Protocol), 102–103, 112, 122
TCP/IP
application layer, 122
arp command, 207–208
configuring for Linux, 662–663
diagnostic commands, 764
historical background, 113
hostname command, 208–209
IP layer, 110–111
ipgconfig command
displaying basic IP configuration, 209
displaying detailed configuration information, 210
flushing local DNS cache, 212
releasing IP lease, 211
renewing IP lease, 211
nbtstat command, 212–213
netstat command
displaying connections, 213–214
displaying interface statistics, 214
network interface layer, 121
network layer, 121
nslookup command
displaying DNS records, 217–218
locating mail server for email address, 218–219

TCP/IP *(continued)*
looking up IP address, 215
subcommands, 215–216
overview, 110–111
`pathping` command, 219–220
`ping` command, 220–221
planning implementation, 244–245
ports, 830
RFCs and, 118–120
`route` command
displaying routing table, 222–225
modifying routing table, 225–226
TCP layer, 112
`tracert` command, 226–229
transport layer, 122
UDP layer, 112
Windows Server 2025, 473
Teams, Microsoft
architecture of
Admin Center, 650–652
Microsoft 365 Group, 646–647
OneDrive for Business, 650
overview, 645–646
SharePoint, 648–649
channels, 642–643
managing, 652–655
overview, 608, 641
tabs, 643–645
teams, 642
Teams Apps option, Microsoft Teams Admin Center, 651
Teams Devices option, Microsoft Teams Admin Center, 651
Teams option, Microsoft Teams Admin Center, 651
Telnet, 105, 122
temporary user profiles, 526
tenants, Microsoft 365, 605, 611–615
terminal services, 461
terminals. *See* consoles, Linux
Terminator movies, 7
terrorism, 864
text editing, Linux
changing text, 695
deleting text, 694
inserting text, 693–695
yanking and putting, 695–697
theft, 864
thin client, 461
threads, 76
Time to Live (TTL), 190, 193
tiny servers, 77, 79
top-level, Linux, 681–682
top-level domains
country code domains, 179–180
defined, 177
generic domains, 178
topology, network
bus topology, 26–27
defined, 26
mesh topology, 30
nodes, 26
packets, 26
ring topology, 29
tower cases, 78
`traceroute` command, 752–753
`tracert` command, 226–229
tracking number, 796
Training & Guide card, Exchange Admin Center, option, Microsoft Teams Admin Center, 628
transitive trust, 501
Transmission Control Protocol (TCP), 102–103, 112, 122
transport layer
OSI model, 25, 102–103
TCP/IP, 122
trees, Active Directory, 501
Trojan horses, 59, 844
troubleshooting
basic steps, 802–803
booting in Safe Mode, 809
checking network settings, 806
dead computers, 803–804
error messages, 805
event logs, 813–814
narrowing down possibilities, 806–807
network connections, 804–805
overview, 801–802
restarting client computer, 808–809
restarting network server, 812–813
restarting network services, 811–812
System Restore, 809–811
tracking log, 814
viewing current network users, 807
wireless network

adding access points, 354–355
antennas, 354
changing channels, 354
checking for obvious problems, 353
overview, 352
pinpointing problems, 353
router's passwords, 355
TTL (Time to Live), 190, 193
tunnels, 54. *See also* virtual private networks (VPNs)
twisted-pair cable
categories, 327–328
installing, 330–331
overview, 15, 30–31
pairs, 328
patch panels, 335–336
pinouts, 332–333
plenum cable, 329
RJ-45 connectors, 333–335
shielded, 329
solid cable, 330
stranded cable, 330
tools, 331–332
troubleshooting, 804–805
unshielded, 329
UTP, 238–239
wall jacks, 335–336

U

Ubuntu distribution, Linux, 659
UDP (User Datagram Protocol), 102–103, 112, 122
Unicode (UTF-8), 104
Uniform Resource Locator (URL), 176–177
Uninterruptible power supply (UPS) device, 475
`uniquepw` parameter, `Net Accounts` command, 565
universal scope, 520
Unix, 676
`unlimited` parameter, `Net Share` command, 574
unmanaged switches, 241
unshielded twisted-pair cable (UTP), 238–239, 329
up, 10
updates, 475
upgrade installation, 471
uplinks, 50
upper layers, OSI model, 96
UPS (Uninterruptible power supply) device, 475
URLs (Uniform Resource Locators), 176–177
user accounts
defined, 73
Linux commands for
`chage` command, 741–742
`gpasswd` command, 744–745
`groupadd` command, 742–743
`groupdel` command, 743–744
`newusers` command, 742–743
`passwd` command, 742
`useradd` command, 739–740
`userdel` command, 741
`usermod` command, 741
Windows Server 2025
account properties, 508
contact information, 513
creating new users, 509–512
deleting, 518
disabling and enabling, 518
domain accounts, 508
groups, 519–525
home folder, 516
local accounts, 508
logon script, 516, 527
overview, 507
profile information, 515–516
profile path, 516
resetting passwords, 517–518
restricting access, 515–516
setting account options, 513–514
setting user properties, 512
specifying logon hours, 514–515
user profiles, 525–527
User Datagram Protocol (UDP), 102–103, 112, 122
user ID, :Linux, 702
user life-cycle management, 71
`user` parameter, `Net Use` command, 578
`/user` parameter, `runas` command, 591
user profiles
local, 525
mandatory, 526
overview, 525
roaming, 526–527
temporary, 526
user training, 61

useradd command, 739–740
userdel command, 741
usermod command, 741
username, 508
username parameter
  Net Group command, 569
  Net Localgroup command, 572
users
  cybersecurity
    Administrator account, 821–822
    authentication, 818
    authorization, 818
    cybersecurity policies, 824
    multifactor authentication, 823
    overview, 817
    passwords, 818–820
    phish testing, 825
    training, 824–825
  Linux, 679
  Microsoft 365
    creating, 615–621
    disabling, 622–623
Users group, Windows Server 2025, 525
Users option, Microsoft Teams Admin Center, 651
users parameter, Net Share command, 574
UTF-8 (Unicode), 104
UTP (unshielded twisted-pair cable), 238–239, 329

V

vandalism, 863
vCenter, 268
VDI (virtual desktop infrastructure), 461
Veeam, 310
Vi text editor
  changing text, 695
  deleting text, 694
  exiting, 691
  inserting text, 693–695
  moving cursor, 692–694
  operating modes, 691–693
  overview, 688–690
  repeating commands, 696
  saving changes, 691
  starting, 689–692
  yanking and putting, 695–697
video, servers, 77
virtual consoles, Linux, 672, 677
virtual desktop infrastructure (VDI), 461
virtual disks
  Hyper-V
    creating, 400–404
    dynamically expanding disk, 395
    fixed-size disk, 395
    formats, 394
  local disk storage, 276
  network accessible storage, 276
  storage area network, 276
virtual local area networks (VLANs)
  data link layer, OSI model, 99
  defined, 41
  overview, 55–56
  scopes and, 159–160
  SSID, 341
  subnets vs., 138–139
Virtual Machine File System (VMFS), 414
virtual machines (VMs)
  AWS
    connecting to, 455–457
    managing, 454–455
    Windows Virtual Machine, 446–453
  Azure
    connecting to, 438–439
    managing, 435–438
    Windows Virtual Machine, 429–435
  backups, 298
  defined, 79, 272
  Hyper-V, 404–409
  installing Linux on, 662
  VMware
    creating virtual machines, 416–422
    overview, 413
    VMware Tools, 423–424
    vSphere, 414
    Workstation Pro, 413, 414–416
virtual network, 277–278
virtual private networks (VPNs)
  accessing computer remotely, 386–387

clients, 388–390
creating with gateway routers, 148–149
defined, 54
overview, 54–55, 385–386
security, 387–388
servers, 388–389
WANs, 259
WAPs, 351–352
virtual switches, 398–400
virtual-based backups, 298, 310
virtualization
Azure
Azure Portal, 428–429
creating account, 427
IaaS, 425
mobile applications, 427
networking, 426
PaaS, 425
SaaS, 425
services, 426–427
SQL database, 427
storage, 426
VMs, 426, 429–439
Web applications, 426
Windows Virtual Machine, 429–435
backups, 310
bare metal, 272
benefits of
disaster recovery, 279–280
energy costs, 279
hardware costs, 278
recoverability, 279
reduced downtime, 279
defined, 79–80
guest, 272
guest operating system, 272
historical background, 274–275
hosts
choosing, 280
defined, 272
Hyper-V
enabling, 395–396
installing operating system, 409–411
Manager Window, 397
overview, 393–394
virtual disks, 394–395, 400–404
virtual machine, 404–409
virtual switches, 398–400
hypervisors
defined, 273
functions, 273
type-1, 274
type-2, 274
network virtualization, 277–278
overview, 271–273
virtual disks
local disk storage, 276
network accessible storage, 276
RAID, 275–277
storage area network, 276
Windows Server 2025 licensing, 281–282
virtualization management platform, 268
virtualization servers, 397
virus protection
antivirus programs, 844–846
overview, 842–844
safe-computing practices, 846
viruses, 20
VMFS (Virtual Machine File System), 414
VMs (virtual machines)
AWS
connecting to, 455–457
managing, 454–455
Windows Virtual Machine, 446–453
Azure
connecting to, 438–439
managing, 435–438
Windows Virtual Machine, 429–435
backups, 298
defined, 79, 272
Hyper-V, 404–409
installing Linux on, 662
VMware, 413–424
VMware
creating virtual machines, 416–422
Horizon View, 462
overview, 413
VMware Tools, 423–424
vSphere, 414
Workstation Pro, 413, 414–416

Voice option, Microsoft Teams Admin Center, 651
volumes, virtual disk, 276
vSphere
  ESXi hypervisor, 414
  vCenter Client, 414
  vCenter Server, 414
  VMFS, 414

W

W3C (World Wide Web Consortium), 95
wall jacks, 335–336
WANs (wide area networks)
  defined, 18–19
  Ethernet private line, 260
  internet connection
    broadband connections, 250–251
    cellular network, 253
    fiber-optic cable, 252–253
    providing redundancy for, 257–258
    T1 lines, 251–252
  routers
    cellular router, 255
    enterprise router, 255
    overview, 253–255
    small office router, 254
  securing connections for remote users, 258–259
  VPN, 259
WAPs (wireless access points)
  collisions, 39
  configuring, 340–342
  defined, 8, 16, 24, 339
  installing, 339–340
water disasters, 863
weather disasters, 863
web servers, 266
well-known ports, 141
whitelist, 853
Whois, 886
wide area networks (WANs)
  defined, 18–19
  Ethernet private line, 260
  firewalls, 256–257
  internet connection
    broadband connections, 250–251
    cellular network, 253
    fiber-optic cable, 252–253
    providing redundancy for, 257–258
    T1 lines, 251–252
  routers
    cellular router, 255
    enterprise router, 255
    overview, 253–255
    small office router, 254
  securing connections for remote users, 258–259
  VPN, 259
Wi-Fi Protected Access (WPA), 349–350
Wi-Fi security, 60
wildcards, 558, 726–727
Windows Calculator, 127–128
Windows clients
  configuring for DHCP
    automatic private IP addressing, 171
    manually, 169–171
    releasing leases, 171
    renewing leases, 171
  configuring network connections, 357–363
  joining domain, 363–366
Windows Defender Firewall
  activating, 834–836
  configuring with Group Policy, 836–842
  overview, 834
Windows Deployment Services server role, 267
Windows DNS Server
  host records, creating, 203–204
  overview, 199
  zone, creating, 200–203
Windows Server 2012, 73
Windows Server 2016, 15
Windows Server 2019, 15, 73
Windows Server 2022, 15
Windows Server 2025, 14, 73
  Active Directory
    defined, 72, 498
    domains, 499–500, 502–503
    forests, 501–502
    objects, 498–499
    organizational units, 500–501, 503–505
    overview, 497–498
    trees, 501
  commands
    batch files, 562–563
    chaining commands, 559

command prompt window, 556–558
Control menu, 557–558
editing commands, 557
environmental variables, 560–562
EventCreate command, 563–564
Net Accounts command, 565
Net Computer command, 566
Net Config command, 566–567
Net Continue command, 567
Net File command, 568
Net Group command, 568–570
Net Help command, 570
Net Helpmsg command, 570–571
Net Localgroup command, 571–572
Net Pause command, 572–573
Net Session command, 573
Net Share command, 574–575
Net Start command, 575
Net Statistics command, 575–576
Net Stop command, 576–577
Net Time command, 577
Net Use command, 577–578
Net User command, 579–580
Net View command, 580–581
overview, 555
piping, 559–560
redirection, 559–560
RunAs command, 581–582
wildcards, 558
configuring
Administrator account, 487–488
Microsoft Management Console, 491–496
Remote Desktop Connection, 488–491
Datacenter Edition, 281
Essentials Edition, 281
group policy
defined, 543
enabling, 544–545
group policy objects, 545–554
overview, 543–544
installing
backing up, 474
checking system requirements, 470
disconnecting UPS devices, 475
domains, 473–474
event logs, 474
full installation, 470–471
licensing options, 471
multiboot features, 471–472
overview, 469
partitions, 472
release notes, 470
server roles and features, 482–485
Setup program, 475–481
TCP/IP configuration, 473
updates, 475
upgrade installation, 471
workgroups, 473–474
installing DHCP role on, 162–163
licensing, 281–282
permissions
access control list, 529–530
file and folder permissions, 530
granting, 540–542
special permissions, 531
shares
defined, 531
hidden shares, 532
setting up, 532–533
special shares, 532
Standard Edition, 281
user accounts
account properties, 508
contact information, 513
creating new users, 509–512
deleting, 518
disabling and enabling, 518
domain accounts, 508
groups, 519–525
home folder, 516
local accounts, 508
logon script, 516, 527
overview, 507
profile information, 515–516
profile path, 516
resetting passwords, 517–518
restricting access, 515–516
setting account options, 513–514
setting user properties, 512
specifying logon hours, 514–515
user profiles, 525–527
Windows Server Update Services (WSUS), 267

wire cutters, 332
wire stripper, 332
wireless access points (WAPs)
  collisions, 39
  configuring, 340–342
  defined, 8, 16, 24, 339
  installing, 339–340
wireless firewall routers, 339
wireless network
  connecting to, 342–343
  overview, 39
  security
    eavesdroppers, 346
    freeloaders, 345–346
    intruders, 344–345
    MAC address filtering, 350–351
    overview, 343–344
    passwords, 348
    rogue access points, 346–347
    spoilers, 346
    SSIDs, 348–349
    VPN, 351–352
    WPA, 349–350
  troubleshooting
    adding access points, 354–355
    antennas, 354
    changing channels, 354
    checking for obvious problems, 353
    overview, 352
    pinpointing problems, 353
    router's passwords, 355
  WAP
    configuring, 340–342
    installing, 339–340
wireless routers, 8, 16
Wireshark, 763
workgroups, 473–474
working set, 599
`workstation` parameter, `Net Config` command, 566
workstation service
  `Net Continue` command, 567
  `Net Pause` command, 567
World Wide Web Consortium (W3C), 95
worm, 844
WPA (Wi-Fi Protected Access), 349–350
WPA2, 350
WPA-PSK, 350
Write permission, Windows Server 2025, 530
WSUS (Windows Server Update Services), 267

X

X Window System, Linux, 672
XenApp, 463–465
`XOR` logical operation, 126

Y

`yum` command, 738–739

Z

zenmap, 763
zone files (DNS databases; master files), 192–194, 723–724
zone transfer, 186
zones
  DNS servers and, 184–185
  primary, 185
  secondary, 186

About the Author

Doug Lowe has written a whole bunch of computer books, including more than 50 *For Dummies* books, among them *Networking For Dummies,* 8th Edition; *Electronics All-in-One For Dummies,* 3rd Edition; *Java All-in-One For Dummies,* 7th Edition; and *PowerPoint For Dummies,* 2nd Edition. He lives in sunny Fresno, California, where the only thing worse than the temperature in the summer is the smoke from the fires. He is the Information Technology director at Blair, Church & Flynn Consulting Engineers in nearby Clovis, California.

Dedication

For Kristen. And for my dad, Kenneth Duane Lowe, Jr. I really miss you.

Author's Acknowledgments

I'd like to thank everyone who was involved with the seventh edition of this book, especially the most excellent editor Elizabeth Kuball, who championed this book through all the editorial details needed to put a book of this scope together on time. Thanks also to Ken Hess, who gave the manuscript a thorough review to ensure the technical accuracy of every sentence and in the process offered many excellent suggestions for improvements. And as always, thanks to all the behind-the-scenes people who chipped in with help I'm not even aware of.

Publisher's Acknowledgments

Executive Editor: Steven Hayes
Managing Editor: Ajith Kumar
Editor: Elizabeth Kuball
Technical Editor: Ken Hess
Editorial Assistant: Hanna Sytsma
Production Editor: Saikarthick Kumarasamy
Cover Image: © Jason marz/Getty Images

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley’s ebook EULA.