Early Release
RAW & UNEDITED

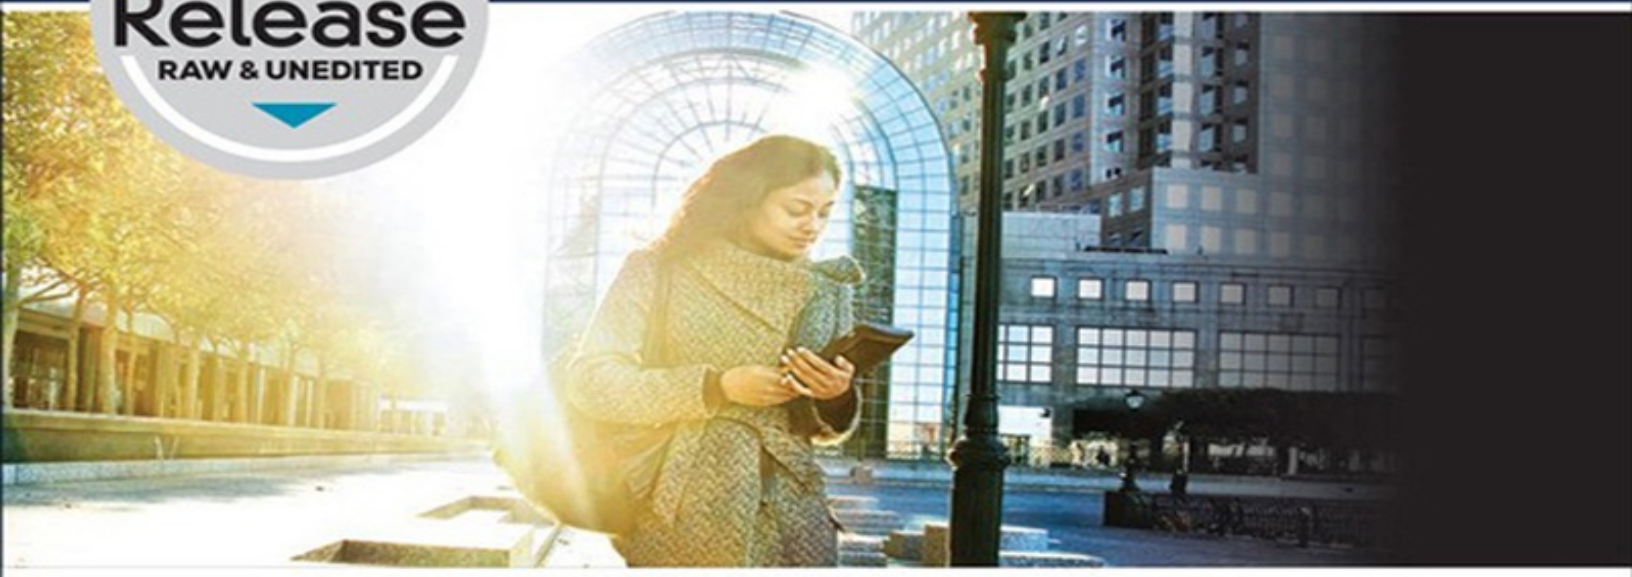

31 Days Before Your CCNA Exam

A Day-By-Day Review Guide for the CCNA 200-301 Certification Exam

ciscopress.com

Allan Johnson

Early Release
RAW & UNEDITED

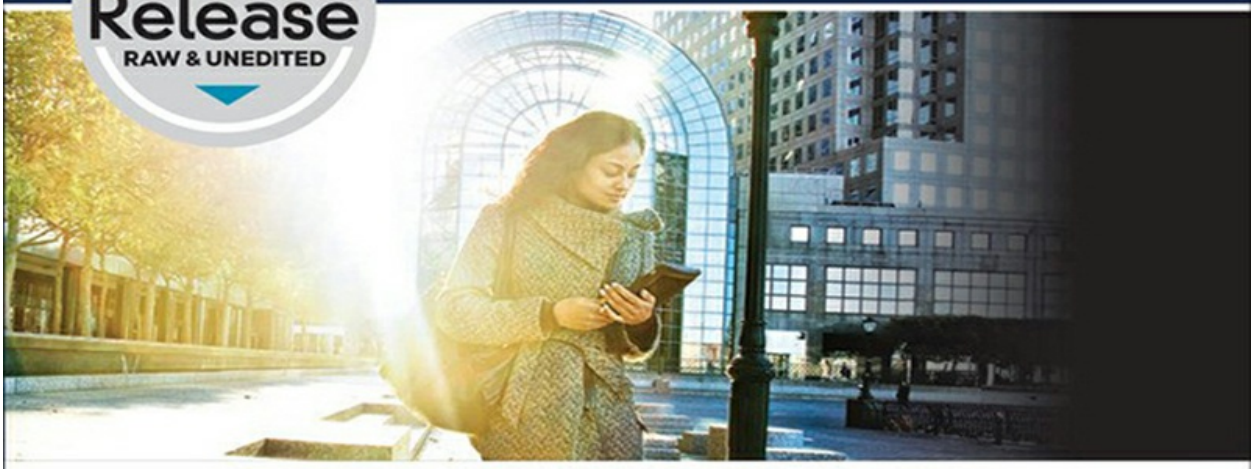

31 Days Before Your CCNA Exam

A Day-By-Day Review Guide for the CCNA 200-301 Certification Exam

ciscopress.com

Allan Johnson

31 Days Before Your CCNA Exam

A Day-By-Day Review Guide for the CCNA 200-301 Certification Exam, Second Edition

Allan Johnson

Cisco Press

Hoboken, New Jersey

A NOTE FOR EARLY RELEASE READERS

With Early Release eBooks, you get books in their earliest form—the author's raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

If you have comments about how we might improve the content and/or examples in this book, or if you notice missing material within this title, please reach out to Pearson at PearsonITAcademics@pearson.com

31 Days Before Your CCNA Exam

Allan Johnson

Copyright© 2024 Cisco Systems, Inc.

Published by:
Cisco Press

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

$PrintCode

Library of Congress Control Number:

ISBN-13: 978-0-13-821425-8

ISBN-10: 0-13-821425-5

Warning and Disclaimer

This book is designed to provide information about exam topics for the Cisco Certified Networking Associate (CCNA) certification. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Please contact us with concerns about any potential bias at https://www.pearson.com/report-bias.html.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems,

Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any

comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Please contact us with concerns about any potential bias at https://www.pearson.com/report-bias.html.

GM K12, Early Career and Professional Learning

Soo Kang

Alliances Manager, Cisco Press

Caroline Antonio

Director, ITP Product Management

Brett Bartow

Executive Editor

James Manly

Managing Editor

Sandra Schroeder

Development Editor

Christopher Cleveland

Senior Project Editor

Mandie Frank

Copy Editor

Kitty Wilson

Editorial Assistant

Cindy Teeters

Designer

Chuti Prasertsith

Composition

codeMantra

Indexer

Proofreader

About the Author

Allan Johnson entered the academic world in 1999 after 10 years as a business owner/operator to dedicate his efforts to his passion for teaching. He holds both an MBA and an MEd in training and development. He taught CCNA courses at the high school level for 7 years and has taught both CCNA and CCNP courses at Del Mar College in Corpus Christi, Texas. In 2003, Allan began to commit much of his time and energy to the CCNA Instructional Support Team, providing services to Networking Academy instructors worldwide and creating training materials. He now splits his time between working as a Curriculum Lead for Cisco Networking Academy and as Account Lead for Unicon (unicon.net), supporting Cisco's educational efforts.

Dedications

For my wife, Becky. Thank you for all your support during this crazy whirlwind of a year. You are the stabilizing force that keeps me grounded.

Acknowledgments

As a technical author, I relied heavily on my technical editor. During the heavy revision of this book back in 2019, Steve Stiles had my back for this work. Steve Stiles did arduous review work necessary to make sure that you get a book that is both technically accurate and unambiguous.

Wendell Odom's *CCNA 200-301 Official Cert Guide, Volume 1* and *Volume 2* were two of my main sources. These two books have the breadth and depth needed to master the CCNA exam topics.

The Cisco Networking Academy authors for the online curriculum and series of Companion Guides take the reader deeper, past the CCNA exam topics, with the ultimate goal of preparing the student not only for CCNA certification but for more advanced college-level technology courses and degrees as well. Thank you especially to Rick Graziani, Bob Vachon, John Pickard, Dave Holzinger, Jane Gibbons, Martin Benson, Suk-Yi Pennock, Allan Reid, Jane Brooke, Anna Bolen, Telethia Willis, and the rest of the ACE team. Their excellent treatment of the material is reflected throughout this book.

James Manly, executive editor, daily juggles multiple projects simultaneously, steering each from beginning to end. James and I have completed many projects together. Thank you again, James, for shepherding this project for me.

Thank you to the professional and thorough review of this work by development editor Christopher Cleveland and project editor Mandie Frank. I've worked with the stellar Chris and Mandie on many previous projects. Their combined efforts ensure that what I authored is ready for publication.

And to the rest of the Pearson family who contributes in countless ways to bring a book to the reader, thank you for all your hard work.

Contents at a Glance

Introduction

Day 31: Networking Models, Devices, and Components

Day 30: Ethernet Switching

Day 29: Switch Configuration Basics

Day 28: IPv4 Addressing

Day 27: IPv6 Addressing

Day 26: VLAN and Trunking Concepts and Configurations

Day 25: STP

Day 24: EtherChannel and HSRP

Day 23: DHCP and DNS

Day 22: Wireless Concepts

Day 21: WLAN Configuration

Day 20: LAN Security and Device Hardening

Day 19: Basic Routing Concepts

Day 18: Basic Router Configuration

Day 17: The Routing Table

Day 16: Inter-VLAN Routing

Day 15: Static and Default Route Configuration

Day 14: OSPF Operation

Day 13: Single-Area OSPF Implementation

Day 12: Fine-Tuning and Troubleshooting OSPF

Day 11: Network Security Concepts

Day 10: ACL Concepts

Day 9: ACL Implementation

Day 8: NAT

Day 7: WAN, VPN, and IPsec

Day 6: QoS

Day 5: CDP and LLDP

Day 4: Device Monitoring, Management, and Maintenance

Day 3: Cloud, Virtualization, and SDN

Day 2: Cisco SD-Access, Ansible, and Terraform

Day 1: Data Formats, REST, and AI

Exam Day

Post-Exam Information

Countdown Calendar

Exam Checklist

Reader Services

Register your copy of this book at www.ciscopress.com/title/9780138214258 for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account. (Be sure to check the box indicating that you would like to hear from us to receive exclusive discounts on future editions of this product.) Enter the product ISBN 9780138214258 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

Contents

- Introduction
  - Study Resources
  - Goals and Methods
  - Who Should Read This Book?
  - Getting to Know the CCNA 200-301 Exam
- Day 31. Networking Models, Devices, and Components
  - CCNA 200-301 Exam Topics
  - Key Points
  - The OSI and TCP/IP Models
  - The TCP/IP Application Layer
  - The TCP/IP Transport Layer
  - The TCP/IP Internet Layer
  - The TCP/IP Network Access Layer
  - Data Encapsulation Summary
  - Networking Icons
  - Devices
  - Physical Layer
  - LAN Device Connection Guidelines
  - LANs and WANs
  - Small Office/Home Office (SOHO)
  - Physical and Logical Topologies
  - Hierarchical Campus Designs

Study Resources
Day 30. Ethernet Switching
CCNA 200-301 Exam Topics
Key Topics
Evolution to Switching
Switching Logic
Collision and Broadcast Domains
Frame Forwarding
Ethernet Overview
Legacy Ethernet Technologies
Current Ethernet Technologies
UTP Cabling
Benefits of Using Switches
Ethernet Addressing
Ethernet Framing
The Role of the Physical Layer
Study Resources
Day 29. Switch Configuration Basics
CCNA 200-301 Exam Topics
Key Topics
Accessing and Navigating the Cisco IOS
Basic Switch Configuration Commands
Half Duplex, Full Duplex, and Port Speed
Verifying Network Connectivity

Troubleshoot Interface and Cable Issues
Study Resources
Day 28. IPv4 Addressing
CCNA 200-301 Exam Topics
Key Topics
IPv4 Addressing
Private and Public IP Addressing
Subnetting in Four Steps
VLSM
Study Resources
Day 27. IPv6 Addressing
CCNA 200-301 Exam Topics
Key Topics
Overview and Benefits of IPv6
The IPv6 Protocol
IPv6 Address Types
Representing the IPv6 Address
IPv6 Subnetting
EUI-64 Concept
Stateless Address Autoconfiguration
Migration to IPv6
Study Resources
Day 26. VLAN and Trunking Concepts and Configurations
CCNA 200-301 Exam Topics

- Key Points
- VLAN Concepts
- Trunking VLANs
- VLAN Configuration and Verification
- Trunking Configuration and Verification
- VLAN Troubleshooting
- Trunking Troubleshooting
- Study Resources

Day 25. STP

- CCNA 200-301 Exam Topic
- Key Topics
- STP Concepts and Operation
- STP Algorithm
- STP Convergence
- STP Varieties
- PVST Operation
- Rapid PVST+ Operation
- Configuring and Verifying Varieties of STP
- Study Resources

Day 24. EtherChannel and HSRP

- CCNA 200-301 Exam Topics
- Key Topics
- EtherChannel Operation
- EtherChannel Protocols

Configuring EtherChannel
Verifying EtherChannel
Troubleshooting EtherChannel
First-Hop Redundancy Concepts
FHRPs
HSRP Operation
HSRP Configuration and Verification
HSRP Load Balancing
Troubleshooting HSRP
VRRP and GLBP
Study Resources
Day 23. DHCP and DNS
CCNA 200-301 Exam Topics
Key Topics
DHCPv4
DHCPv4 Configuration Options
DHCPv6
DHCPv6 Configuration Options
DHCP Troubleshooting
DNS Operation
Troubleshooting DNS
Verifying Host IP Configuration
Study Resources
Day 22. Wireless Concepts

- CCNA 200-301 Exam Topics
- Key Topics
- Wireless Standards
- Wireless Topologies
- AP Architectures
- Wireless Security Protocols
- Study Resources

Day 21. WLAN Configuration

- CCNA 200-301 Exam Topics
- Key Topics
- Logging Into a Cisco WLC
- Configuring a WLC with a WLAN
- Study Resources

Day 20. LAN Security and Device Hardening

- CCNA 200-301 Exam Topics
- Key Topics
- Access Control
- Port Security
- LAN Threat Mitigation
- Study Resources

Day 19. Basic Routing Concepts

- CCNA 200-301 Exam Topics
- Key Topics
- Packet Forwarding

Routing Methods
Classifying Dynamic Routing Protocols
Dynamic Routing Metrics
Administrative Distance
IGP Comparison Summary
Routing Loop Prevention
Link-State Routing Protocol Features
Study Resources
Day 18. Basic Router Configuration
CCNA 200-301 Exam Topics
Key Topics
Basic Router Configuration with IPv4
Basic Router Configuration with IPv6
Verifying IPv4 and IPv6 Network Connectivity
Small Office or Home Office Routers
Basic IP Addressing Troubleshooting
Study Resources
Day 17. The Routing Table
CCNA 200-301 Exam Topics
Key Topics
Two Router Functions
Components of the Routing Table
Study Resources
Day 16. Inter-VLAN Routing

CCNA 200-301 Exam Topics
Key Points
Inter-VLAN Routing Concepts
Router on a Stick Configuration and Verification
Multilayer Switching Inter-VLAN Routing Configuration and Verification
Study Resources
Day 15. Static and Default Route Configuration
CCNA 200-301 Exam Topics
Key Topics
Static and Default Routing Overview
IPv4 Static Route Configuration
Study Resources
Day 14. OSPF Operation
CCNA 200-301 Exam Topics
Key Topics
Single-Area OSPF Operation
OSPFv2 Versus OSPFv3
Multiarea OSPF Operation
Study Resources
Day 13. Single-Area OSPF Implementation
CCNA 200-301 Exam Topics
Key Topics
Single-Area OSPFv2 Configuration

Verifying OSPFv2
Study Resources
Day 12. Fine-Tuning and Troubleshooting OSPF
CCNA 200-301 Exam Topics
Key Topics
OSPFv2 Configuration Example
Modifying OSPFv2
Troubleshooting OSPF
Study Resources
Day 11. Network Security Concepts
CCNA 200-301 Exam Topics
Key Topics
Security Fundamentals
Network Attacks
Security Program
Password Vulnerabilities
Password Alternatives
Study Resources
Day 10. ACL Concepts
CCNA 200-301 Exam Topics
Key Topics
ACL Operation
Planning to Use ACLs
Study Resources

Day 9. ACL Implementation
- CCNA 200-301 Exam Topics
- Key Topics
- Configuring Standard Numbered IPv4 ACLs
- Configuring Extended Numbered IPv4 ACLs
- Configuring Named IPv4 ACLs
- Verifying IPv4 ACLs
- Comparing IPv4 and IPv6 ACLs
- Configuring IPv6 ACLs
- Verifying IPv6 ACLs
- Troubleshooting ACLs
- Study Resources

Day 8. NAT
- CCNA 200-301 Exam Topics
- Key Topics
- NAT Concepts
- Configuring Static NAT
- Verifying NAT
- Troubleshooting NAT
- Study Resources

Day 7. WAN, VPN, and IPsec
- CCNA 200-301 Exam Topics
- Key Topics
- WAN Topologies

WAN Connection Options
VPN Technology
Study Resources
Day 6. QoS
CCNA 200-301 Exam Topic
Key Topics
QoS
Study Resources
Day 5. CDP and LLDP
CCNA 200-301 Exam Topic
Key Topics
CDP Overview
LLDP Overview
Study Resources
Day 4. Device Monitoring, Management, and Maintenance
CCNA 200-301 Exam Topics
Key Topics
SNMP Operation
Configuring SNMP
Verifying SNMP
Syslog
Network Time Protocol
Cisco IOS File System and Devices
Managing Cisco IOS Images

Password Recovery
Study Resources
Day 3. Cloud, Virtualization, and SDN
CCNA 200-301 Exam Topics
Key Topics
Cloud Computing
Software-Defined Networking
Study Resources
Day 2. Cisco SD-Access, Ansible, and Terraform
CCNA 200-301 Exam Topics
Key Topics
Cisco SD-Access
Configuration Management Tools: Ansible and Terraform
Study Resources
Day 1. Data Formats, REST, and AI
CCNA 200-301 Exam Topics
Key Topics
Data Formats
RESTful APIs
AI and ML in IT Operations
Study Resources
Exam Day
What You Need for the Exam
What You Should Receive After Completion

Summary
Post-Exam Information
Receiving Your Certificate
Determining Career Options
Examining Certification Options
If You Did Not Pass the Exam
Summary
CCNA Countdown Calendar
Exam Checklist

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

• **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).

• *Italic* indicates arguments for which you supply actual values.

• Vertical bars (|) separate alternative, mutually exclusive elements.

• Square brackets ([ ]) indicate an optional element.

• Braces ({ }) indicate a required choice.

• Braces within brackets ([{ }]) indicate a required choice within an optional element.

Introduction

If you're reading this introduction, you've probably already spent a considerable amount of time and energy pursuing your CCNA 200-301 certification. Regardless of how you got to this point in your travels through your CCNA studies, *31 Days Before Your CCNA Exam* most likely represents the last leg of your journey on your way to the destination: to become a Cisco Certified Network Associate. However, if you are like me, you might be reading this book at the *beginning* of your studies. If so, this book provides an excellent overview of the material you must now spend a great deal of time studying and practicing. But I must warn you: Unless you are extremely well versed in networking technologies and have considerable experience configuring and troubleshooting Cisco routers and switches, this book will *not* serve you well as the sole resource for your exam preparations. Therefore, let me spend some time discussing my recommendations for study resources.

Study Resources

Cisco Press and Pearson IT Certification offer an abundance of CCNA-related books to serve as your primary source for learning how to install, configure, operate, and troubleshoot small to medium-size routed and switched networks.

Primary Resources

Technet24

First on the list of important resources is Wendell Odom's *CCNA 200-301 Official Cert Guide Library*, *Volume 1* and *Volume 2*, 2nd edition (ISBN: 9780138221393). If you do not buy any other books, buy these. Wendell's method of teaching, combined with his technical expertise and down-to-earth style, is unsurpassed in our industry. As you read through his books, you sense that he is sitting right there next to you, walking you through the material. With your purchase, you get access to practice exams and study materials and other online resources that are worth the price of the book. There is no better resource on the market for a CCNA candidate.

If you are a Cisco Networking Academy student, you are blessed with access to the online version of the CCNA version 7 curriculum and the wildly popular Packet Tracer network simulator. The Cisco Network Academy curriculum has three courses. To learn more about CCNAv7 courses and to find an Academy near you, visit http://www.netacad.com.

However, if you are not an Academy student but want to benefit from the extensive authoring done for these courses, you can buy any or all of CCNAv7 Companion Guides (CGs) and Labs & Study Guides (LSGs) of the Academy's popular online curriculum. Although with this option you will not have access to the Packet Tracer files, you will have access to the tireless work of an outstanding team of Cisco Academy instructors dedicated to providing students with comprehensive and engaging CCNA preparation course material. The titles and ISBNs for the CCNAv7 CGs and LSGs follow:

• *Introduction to Networks v7 Companion Guide* (ISBN: 9780136633662)

• *Introduction to Networks v7 Labs & Study Guide* (ISBN: 9780136634454)

• *Switching, Routing, and Wireless Essentials v7 Companion Guide* (ISBN: 9780136729358)

• *Switching, Routing, and Wireless Essentials v7 Labs & Study Guide* (ISBN: 9780136634386)

• *Enterprise Networking, Security, and Automation v7 Companion Guide* (ISBN: 9780136634324)

• *Enterprise Networking, Security, and Automation v7 Labs & Study Guide* (ISBN: 9780136634690)

You can find these books at http://www.ciscopress.com by clicking the Cisco Networking Academy link.

Supplemental Resources

In addition to the book you hold in your hands, I recommend three supplemental resources to augment your final 31 days of review and preparation.

First is Scott Empson's very popular *CCNA 200-301 Portable Command Guide* (ISBN: 9780135937822). This guide is much more than just a listing

Technet24

of commands and what they do. Yes, it summarizes all the CCNA certification-level IOS commands, keywords, command arguments, and associated prompts. It also provides you with tips and examples of how to apply the commands to real-world scenarios. Configuration examples throughout the book provide you with a better understanding of how these commands are used in simple network designs.

Second, Jason Gooley's *CCNA 200-301 Complete Video Course*, second edition (ISBN: 9780138213596) is a comprehensive training course that brings Cisco CCNA exam topics to life through the use of real-world demonstrations, animations, live instruction, and configurations, making learning these foundational networking topics easy and fun.

Third, Wendell Odom's *IP Subnetting LiveLessons* (ISBN: 9780135497777) and *IP Subnetting Practice Questions Kit* (ISBN: 9780135647288) will help you master this crucial skill. Subnetting is not only an IPv4 address design skill, it is also crucial skill for troubleshooting situations where IPv4 addressing has been misconfigured. You are likely to have both types of questions on the CCNA exam.

The Cisco Learning Network

Finally, if you have not done so already, you should register with The Cisco Learning Network at https://learningnetwork.cisco.com/s/. The Cisco Learning Network, sponsored by Cisco, is a free social learning network where IT professionals can engage in the common pursuit of enhancing and

advancing their IT careers. Here you can find many resources to help you prepare for your CCNA exam, in addition to a community of like-minded people ready to answer your questions, help you with your struggles, and share in your triumphs.

So which resources should you buy? The answer to that question depends largely on how deep your pockets are and how much you like books. If you're like me, you must have it all! I admit it; my bookcase is a testament to my Cisco "geekness." But if you are on a budget, choose one of the primary study resources and one of the supplemental resources (such as Wendell Odom's certification library and Scott Empson's command guide). Whatever you choose, you will be in good hands. Any or all of these authors will serve you well.

Goals and Methods

The main goal of this book is to provide you with a clear and succinct review of the CCNA objectives. Each day's exam topics are grouped into a common conceptual framework and use the following format:

• A title for the day that concisely states the overall topic

• A list of one or more CCNA 200-301 exam topics to be reviewed

• A "Key Topics" section that introduces the review material and quickly orients you to the day's focus

Technet24

• An extensive review section consisting of short paragraphs, lists, tables, examples, and graphics

• A “Study Resources” section that gives you a quick reference for locating more in-depth treatment of the day’s topics

The book counts down starting with Day 31, continues through exam day, and provides post-test information. This book also provides a calendar and a checklist that you can tear out and use during your exam preparation.

Use the calendar to enter each actual date beside the countdown day and the exact day, time, and location of your CCNA exam. The calendar provides a visual for the time you can dedicate to each CCNA exam topic.

The checklist highlights important tasks and deadlines leading up to your exam. Use it to help map out your studies.

Who Should Read This Book?

The audience for this book is anyone finishing preparation for taking the CCNA 200-301 exam. A secondary audience is anyone needing a refresher review of CCNA exam topics—possibly before attempting to recertify or sit for another certification for which the CCNA is a prerequisite.

Getting to Know the CCNA 200-301 Exam

For the current certification announced in August 2024, Cisco made minor

revisions to the CCNA 200-301 exam. This book focuses on the entire list of topics published for the exam. It is a 120-minute exam associated with the CCNA certification. This exam tests a candidate's knowledge and skills related to network fundamentals, network access, IP connectivity, IP services, security fundamentals, and automation and programmability. Use the following steps to access a tutorial at home that demonstrates the exam environment before you go to take the exam:

Visit **https://home.pearsonvue.com/cisco** and go through the signup process or log in.

On the right side, click the **Certification Tutorial** link.

When you get to the testing center and check in, the proctor verifies your identity, gives you some general instructions, and takes you into a quiet room containing a PC. When you're at the PC, you have a few things to do before the timer starts on your exam. For instance, you can take the tutorial to get accustomed to the PC and the testing engine. Every time I sit for an exam, I go through the tutorial even though I know how the test engine works. It helps me settle my nerves and get focused. Anyone who has user-level skills in getting around a PC should have no problem with the testing environment.

Note

You can also take the exam at home. For more information, visit https://home.pearsonvue.com/cisco/online.

Technet24

When you start the exam, you are asked a series of questions. The questions are presented one at a time, and you must answer each one before you can move on to the next question. The exam engine does not let you go back and change any answers. Each exam question is in one of the following formats:

- Multiple choice
- Fill in the blank
- Drag and drop
- Testlet
- Simlet
- Simulation

The multiple-choice format simply requires that you point and click a circle or checkbox next to the correct answer(s). Cisco traditionally tells you how many answers you need to choose, and the testing software prevents you from choosing too many or too few.

Fill-in-the-blank questions usually require you only to type numbers. However, if words are requested, the case does not matter unless the answer is a command that is case sensitive (such as passwords and device names when you’re configuring authentication).

Drag-and-drop questions require you to click and hold, move a button or an icon to another area, and release the mouse button to place the object somewhere else—usually in a list. For some questions, to get the question correct, you might need to put a list of five things in the proper order.

A testlet contains one general scenario and several multiple-choice questions about the scenario. Testlets are ideal if you are confident in your knowledge of the scenario's content because you can leverage your strength over multiple questions.

A simlet is similar to a testlet, in that you are given a scenario with several multiple-choice questions. However, a simlet uses a network simulator to allow you access to a simulation of the Cisco IOS Software command line. You can use **show** commands to examine a network's current behavior and answer the question.

A simulation also involves a network simulator, but you are given a task to accomplish, such as implementing a network solution or troubleshooting an existing network implementation. You do this by configuring one or more routers and switches. The exam grades the question based on the configuration you changed or added. A newer form of the simulation question is the GUI-based simulation, which simulates a graphical interface such as that found on a Linksys router or the Cisco Security Device Manager.

Registering for the CCNA 200-301 Exam

Technet24

If you are starting this book 31 days before you take the CCNA 200-301 exam, register for the exam right now. In my testing experience, there is no better motivator than a scheduled test date staring me in the face. I'm willing to bet the same holds true for you. Don't worry about unforeseen circumstances. You can cancel your exam registration for a full refund up to 24 hours before the scheduled exam time. So if you're ready, gather the following information and register right now!

• Legal name

• Social Security or passport number

• Company name

• Valid email address

• Method of payment

You can schedule your exam at any time by visiting www.pearsonvue.com/cisco/. I recommend that you schedule it for 31 days from now. The process and available test times vary based on the local testing center you choose.

Remember: There is no better motivation for study than an actual test date. *Sign up today*.

Topics Covered on the CCNA Exam

Table I-1 summarizes the seven domains of the CCNA 200-301 exam.

Table I-1 *CCNA 200-301 Exam Domains and Weightings*

Domain	Percentage of Exam
1.0 Network Fundamentals	20%
2.0 Network Access	20%
3.0 IP Connectivity	25%
4.0 IP Services	10%
5.0 Security Fundamentals	15%
6.0 Automation and Programmability	10%

Although Cisco outlines general exam topics, not all topics might appear on the CCNA exam; likewise, topics that are not specifically listed might appear on the exam. The exam topics that Cisco provides and that this book covers provide a general framework for exam preparation. Be sure to check Cisco's website for the latest exam topics.

Exam Topics Cross-reference

Technet24

To give you a quick reference to help find which exam topics are covered on which days, I created two cross-refence tables. You can use Table I-2 to look up an exam topic and find the day or days on which it is covered. Use Table I-3 to find out which exam topics are covered on each day.

Table I-2 *Days on Which Each Exam Topic Is Covered*

Exam Topic	Day(s)
1.1 Explain the role and function of network components	31, 30, 22, 19
1.2 Describe characteristics of network topology architectures	31, 18, 7, 3
1.3 Compare physical interface and cabling types	31
1.4 Identify interface and cable issues (collisions, errors, mismatch duplex, and/or speed)	31, 29, 18
1.5 Compare TCP to UDP	31
1.6 Configure and verify IPv4 addressing and subnetting	29, 28, 18
1.7 Describe private IPv4 addressing	28
1.8 Configure and verify IPv6 addressing and prefix	27, 18
1.9 Describe IPv6 address types	27
1.10 Verify IP parameters for Client OS (Windows, Mac OS, Linux)	23
1.11 Describe wireless principles	22
1.12 Explain virtualization fundamentals (server virtualization, containers, and VRFs)	3
1.13 Describe switching concepts	30
2.1 Configure and verify VLANs (normal range) spanning multiple switches	26, 16
2.2 Configure and verify interswitch connectivity	26, 16
2.3 Configure and verify Layer 2 discovery protocols (Cisco Discovery Protocol and LLDP)	5
2.4 Configure and verify (Layer 2/Layer 3) EtherChannel (LACP)	24
2.5 Interpret basic operations of Rapid PVST+ Spanning Tree Protocol	25
2.6 Describe Cisco Wireless Architectures and AP modes	22
2.7 Describe physical infrastructure connections of WLAN components (AP, WLC, access/trunk ports, and LAG)	22
2.8 Describe network device management access (Telnet, SSH, HTTP, HTTPS, console, TACACS+/RADIUS, and cloud managed)	29
2.9 Interpret the wireless LAN GUI configuration for client connectivity, such as WLAN creation, security settings, QoS profiles, and advanced settings	21
3.1 Interpret the components of routing table	17
3.2 Determine how a router makes a forwarding decision by default	19, 17
3.3 Configure and verify IPv4 and IPv6 static routing	15
3.4 Configure and verify single area OSPFv2	14, 13, 12
3.5 Describe the purpose, functions, and concepts of first hop redundancy protocols	24

Technet24

4.1 Configure and verify inside source NAT using static and pools	8
4.2 Configure and verify NTP operating in a client and server mode	4
4.3 Explain the role of DHCP and DNS within the network	23
4.4 Explain the function of SNMP in network operations	4
4.5 Describe the use of syslog features, including facilities and severity levels	4
4.6 Configure and verify DHCP client and relay	23
4.7 Explain the forwarding per-hop behavior (PHB) for QoS such as classification, marking, queuing, congestion, policing, and shaping	6
4.8 Configure network devices for remote access using SSH	20
4.9 Describe the capabilities and functions of TFTP/FTP in the network	4
5.1 Define key security concepts (threats, vulnerabilities, exploits, and mitigation techniques)	11
5.2 Describe security program elements (user awareness, training, and physical access control)	11
5.3 Configure and verify device access control using local passwords	20
5.4 Describe security password policy elements, such as management, complexity, and password alternatives (multifactor authentication, certificates, and biometrics)	11
5.5 Describe IPsec remote access and site-to-site VPNs	7
5.6 Configure and verify access control lists	10, 9
5.7 Configure and verify Layer 2 security features (DHCP snooping, dynamic ARP inspection, and port security)	20
5.8 Compare authentication, authorization, and accounting concepts	20
5.9 Describe wireless security protocols (WPA, WPA2, and WPA3)	22
5.10 Configure and verify WLAN within the GUI using WPA2 PSK	21
6.1 Explain how automation impacts network management	2
6.2 Compare traditional networks with controller-based networking	3
6.3 Describe controller-based, software defined architecture (overlay, underlay, and fabric)	3, 2
6.4 Explain AI (generative and predictive) and machine learning in network operations	1
6.5 Describe characteristics of REST-based APIs (authentication types, CRUD, HTTP verbs, and data encoding)	1
6.6 Recognize the capabilities of configuration management mechanisms such as Ansible and Terraform	2
6.7 Recognize components of JSON-encoded data	1

Table I-3 *Exam Topics Covered on Each Day*

Day	Exam Topic(s)
31	1.1 Explain the role and function of network components
	1.2 Describe characteristics of network topology architectures
	1.3 Compare physical interface and cabling types
	1.4 Identify interface and cable issues (collisions, errors, mismatch duplex, and/or speed)
	1.5 Compare TCP to UDP
30	1.1 Explain the role and function of network components
	1.13 Describe switching concepts
29	1.4 Identify interface and cable issues (collisions, errors, mismatch duplex, and/or speed)
	1.6 Configure and verify IPv4 addressing and subnetting
	2.8 Describe network device management access (Telnet, SSH, HTTP, HTTPS, console, TACACS+/RADIUS, and cloud managed)
28	1.6 Configure and verify IPv4 addressing and subnetting
	1.7 Describe private IPv4 addressing
27	1.8 Configure and verify IPv6 addressing and prefix
	1.9 Describe IPv6 address types
26	2.1 Configure and verify VLANs (normal range) spanning multiple switches
	2.2 Configure and verify interswitch connectivity
25	2.5 Interpret basic operations of Rapid PVST+ Spanning Tree Protocol
24	2.4 Configure and verify (Layer 2/Layer 3) EtherChannel (LACP)
	3.5 Describe the purpose, functions, and concepts of first hop redundancy protocols

Technet24

23	1.10 Verify IP parameters for Client OS (Windows, Mac OS, Linux)
	4.3 Explain the role of DHCP and DNS within the network
	4.6 Configure and verify DHCP client and relay
22	1.1 Explain the role and function of network components
	1.11 Describe wireless principles
	2.6 Describe Cisco Wireless Architectures and AP modes
	2.7 Describe physical infrastructure connections of WLAN components (AP, WLC, access/trunk ports, and LAG)
	5.9 Describe wireless security protocols (WPA, WPA2, and WPA3)
21	2.9 Interpret the wireless LAN GUI configuration for client connectivity, such as WLAN creation, security settings, QoS profiles, and advanced settings
	5.10 Configure and verify WLAN within the GUI using WPA2 PSK
20	4.8 Configure network devices for remote access using SSH
	5.3 Configure and verify device access control using local passwords
	5.7 Configure and verify Layer 2 security features (DHCP snooping, dynamic ARP inspection, and port security)
	5.8 Compare authentication, authorization, and accounting concepts
19	1.1 Explain the role and function of network components
	3.2 Determine how a router makes a forwarding decision by default
18	1.2 Describe characteristics of network topology architectures
	1.4 Identify interface and cable issues (collisions, errors, mismatch duplex, and/or speed)

	1.6 Configure and verify IPv4 addressing and subnetting
	1.8 Configure and verify IPv6 addressing and prefix
17	3.1 Interpret the components of routing table
	3.2 Determine how a router makes a forwarding decision by default
16	2.1 Configure and verify VLANs (normal range) spanning multiple switches
	2.2 Configure and verify interswitch connectivity
15	3.3 Configure and verify IPv4 and IPv6 static routing
14	3.4 Configure and verify single area OSPFv2
13	3.4 Configure and verify single area OSPFv2
12	3.4 Configure and verify single area OSPFv2
11	5.1 Define key security concepts (threats, vulnerabilities, exploits, and mitigation techniques)
	5.2 Describe security program elements (user awareness, training, and physical access control)
	5.4 Describe security password policy elements, such as management, complexity, and password alternatives (multifactor authentication, certificates, and biometrics)
10	5.6 Configure and verify access control lists
9	5.6 Configure and verify access control lists
8	4.1 Configure and verify inside source NAT using static and pools
7	1.2 Describe characteristics of network topology architectures
	5.5 Describe IPsec remote access and site-to-site VPNs

Technet24

6	4.7 Explain the forwarding per-hop behavior (PHB) for QoS such as classification, marking, queuing, congestion, policing, and shaping
5	2.3 Configure and verify Layer 2 discovery protocols (Cisco Discovery Protocol and LLDP)
4	4.2 Configure and verify NTP operating in a client and server mode
	4.4 Explain the function of SNMP in network operations
	4.5 Describe the use of syslog features, including facilities and severity levels
	4.9 Describe the capabilities and functions of TFTP/FTP in the network
3	1.2 Describe characteristics of network topology architectures
	1.12 Explain virtualization fundamentals (server virtualization, containers, and VRFs)
	6.2 Compare traditional networks with controller-based networking
	6.3 Describe controller-based, software defined architecture (overlay, underlay, and fabric)
2	6.1 Explain how automation impacts network management
	6.3 Describe controller-based, software defined architecture (overlay, underlay, and fabric)
	6.6 Recognize the capabilities of configuration management mechanisms such as Ansible and Terraform
1	6.4 Explain AI (generative and predictive) and machine learning in network operations
	6.5 Describe characteristics of REST-based APIs (authentication types, CRUD, HTTP verbs, and data encoding)
	6.7 Recognize components of JSON-encoded data

Credits

Technet24

Day 31. Networking Models, Devices, and Components

CCNA 200-301 Exam Topics

• 1.1 Explain the role and function of network components

• 1.2 Describe characteristics of network topology architectures

• 1.3 Compare physical interface and cabling types

• 1.4 Identify interface and cable issues (collisions, errors, mismatch duplex, and/or speed

• 1.5 Compare TCP to UDP

Key Points

Both the Open Systems Interconnection (OSI) and Transmission Control Protocol/Internet Protocol (TCP/IP) networking models are important conceptual frameworks for understanding networks. Today we review the layers and functions of each model, along with the process of data flow from source to destination. We also spend some time on Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP). Then we wrap up the day with a look at devices used in today’s networks, the media used to interconnect those devices, and the different types of network topologies.

Note

This day might seem a bit long. However, you need to be very familiar with all of this content. Scan the day, focusing on areas where you feel least confident in your knowledge.

The OSI and TCP/IP Models

To understand how communication occurs across the network, you can use layered models as a framework for representing and explaining networking concepts and technologies. Layered models, such as the TCP/IP and OSI models, support interoperability between competing vendor product lines.

The OSI model principally serves as a tool for explaining networking concepts and troubleshooting. However, the protocols of the TCP/IP suite are the rules by which networks now operate. Because both models are important, you should be well versed in each model's layers and know how the models map to each other. Figure 31-1 summarizes the two models.

Technet24

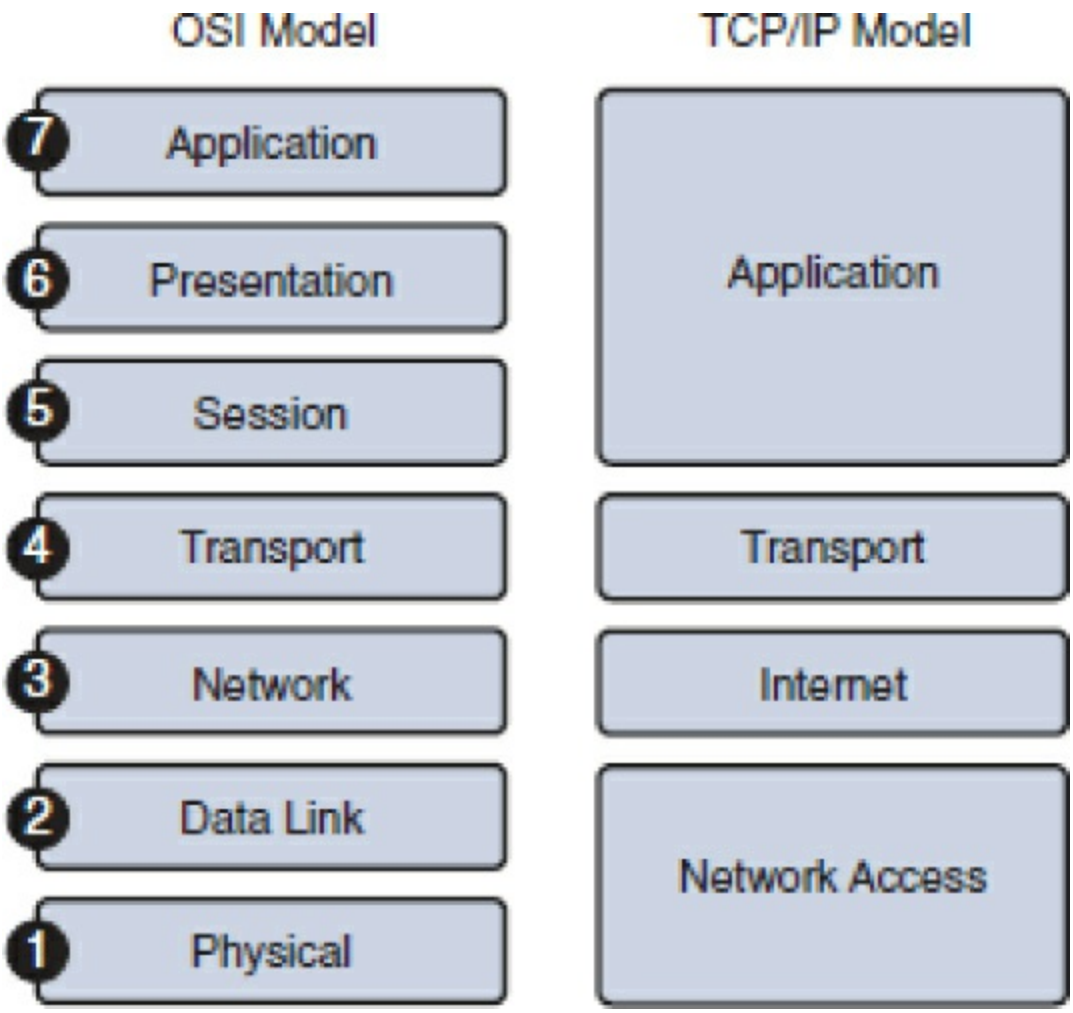

Figure 31-1 *OSI and TCP/IP Models*

Using two models can be confusing; however, these simple guidelines might help:

• When discussing layers of a model, we are usually referring to the OSI model.

• When discussing protocols, we are usually referring to the TCP/IP model.

The next sections quickly review the OSI layers and the TCP/IP protocols.

OSI Layers

Table 31-1 summarizes the layers of the OSI model and provides a brief functional description.

Table 31-1 *OSI Model Layers and Functions*

Technet24

Layer	Functional Description
Application (7)	Refers to interfaces between network and application software. Also includes authentication services.
Presentation (6)	Defines the format and organization of data. Includes encryption.
Session (5)	Establishes and maintains end-to-end bidirectional flows between endpoints. Includes managing transaction flows.
Transport (4)	Provides a variety of services between two host computers, including connection establishment and termination, flow control, error recovery, and segmentation of large data blocks into smaller parts for transmission.
Network (3)	Refers to logical addressing, routing, and path determination.
Data link (2)	Formats data into frames appropriate for transmission onto some physical medium. Defines rules for when the medium can be used. Defines the means by which to recognize transmission errors.
Physical (1)	Defines the electrical, optical, cabling, connectors, and procedural details required for transmitting bits, represented as some form of energy passing over a physical medium.

The following mnemonic phrase, in which the first letter represents the layer (*A* stands for *application*), can help in memorizing the name and order of the layers from top to bottom:

All **P**eople **S**eem **T**o **N**eed **D**ata **P**rocessing

TCP/IP Layers and Protocols

The TCP/IP model defines four categories of functions that must occur for communications to succeed. Most protocol models describe vendor-specific protocol stacks. However, because the TCP/IP model is an open standard, one company does not control the definition of the model.

Table 31-2 summarizes the TCP/IP layers, their functions, and the most common protocols.

Table 31-2 *TCP/IP Layer Functions*

Technet24

TCP/IP Layer	Function	Example Protocols
Application	Represents data to the user and controls dialogue	DNS, Telnet, SMTP, SSH, POP3, IMAP, DHCP, HTTP, FTP, SNMP
Transport	Supports communication between diverse devices across diverse networks	TCP, UDP
Internet	Determines the best path through the network	IP, ARP, ICMP
Network access	Controls the hardware devices and media that make up the network	Ethernet, Wireless

In the coming days, we review these protocols in more detail. For now, a brief description of the main TCP/IP protocols follows:

• **Domain Name System (DNS):** Provides the IP address of a website or domain name so that a host can connect to it

• **Telnet:** Enables administrators to log in to a host from a remote location

• **SSH (Secure Shell):** Enables secure remote access and management of network devices and servers.

• **Simple Mail Transfer Protocol (SMTP), Post Office Protocol (POP3), and Internet Message Access Protocol (IMAP):** Facilitate the sending of email messages between clients and servers

• **Dynamic Host Configuration Protocol (DHCP):** Assigns IP addressing to requesting clients

• **Hypertext Transfer Protocol (HTTP):** Transfers information between web clients and web servers

• **File Transfer Protocol (FTP):** Facilitates the download and upload of files between an FTP client and an FTP server

• **Simple Network Management Protocol (SNMP):** Enables network management systems to monitor devices attached to the network

• **Transmission Control Protocol (TCP):** Supports virtual connections between hosts on the network to provide reliable delivery of data

• **User Datagram Protocol (UDP):** Supports faster, unreliable delivery of lightweight or time-sensitive data

• **Internet Protocol (IP):** Provides a unique global address to computers for communicating over the network

Technet24

• **Address Resolution Protocol (ARP):** Finds a host's hardware address when only the IP address is known

• **Internet Control Message Protocol (ICMP):** Sends error and control messages, including reachability of another host and availability of services

• **Ethernet:** Serves as the most popular LAN standard for framing and preparing data for transmission onto the media

• **Wireless:** Includes both IEEE 802.11 standards for wireless local-area networks (WLANs) and cellular access options.

Protocol Data Units and Encapsulation

As application data is passed down the protocol stack on its way to be transmitted across the network media, various protocols add information to it at each level. This is commonly known as the *encapsulation process*. The data structure at any given layer is called a *protocol data unit* (*PDU*). Table 31-3 lists the PDUs at each layer of the OSI model.

Table 31-3 *PDUs at Each Layer of the OSI Model*

OSI Layer	PDU
Application	Data
Presentation	Data
Session	Data
Transport	Segment
Network	Packet
Data link	Frame
Physical	Bits

The following steps summarize the communication process from any source to any destination:

Data is created at the application layer of the originating source device.

As the data passes down the protocol stack in the source device, it is segmented and encapsulated.

The data is generated onto the media at the network access layer of the stack.

Technet24

The data is transported through the internetwork, which consists of media and any intermediary devices.

The destination device receives the data at the network access layer.

As the data passes up the stack in the destination device, it is decapsulated and reassembled.

The data is passed to the destination application at the application layer of the destination device.

The TCP/IP Application Layer

The application layer of the TCP/IP model provides an interface between software such as a web browser and the network itself. The process of requesting and receiving a web page works like this:

An HTTP request is sent, including an instruction to “get” a file (which is often a website’s home page).

An HTTP response is sent from the web server with a code in the header, usually either 200 (request succeeded, and information is returned in response) or 404 (page not found).

The HTTP request and the HTTP response are encapsulated in headers. The content of the headers allows the application layers on each end device to communicate. Regardless of the application layer protocol (HTTP, FTP,

DNS, and so on), all headers use the same general process for communicating between application layers on the end devices.

The TCP/IP Transport Layer

The transport layer, through TCP, provides a mechanism to guarantee delivery of data across the network. TCP supports error recovery to the application layer through the use of basic acknowledgment logic. Adding to the process for requesting a web page, TCP operation works like this:

The web client sends an HTTP request for a specific web server down to the transport layer.

TCP encapsulates the HTTP request with a TCP header and includes the destination port number for HTTP.

Lower layers process and send the request to the web server.

The web server receives HTTP requests and sends a TCP acknowledgment back to the requesting web client.

The web server sends the HTTP response down to the transport layer.

TCP encapsulates the HTTP data with a TCP header.

Lower layers process and send the response to the requesting web client.

Technet24

The requesting web client sends an acknowledgment back to the web server.

If data is lost at any point during this process, TCP must recover the data. HTTP at the application layer does not get involved in error recovery.

In addition to providing TCP, the transport layer provides UDP, a connectionless, unreliable protocol for sending data that does not require or need error recovery. Table 31-4 lists the main features that the transport protocols support. Both TCP and UDP support the first function; only TCP supports the rest.

Table 31-4 *TCP/IP Transport Layer Features*

Function	Description
Multiplexing using ports	Function that enables receiving hosts to choose the correct application for which the data is destined, based on the destination port number.
Error recovery (reliability)	Process of numbering and acknowledging data with Sequence and Acknowledgment header fields.
Flow control using windowing	Process that involves a sliding window size that the two end devices dynamically agree upon at various points during the virtual connection. The window size, represented in bytes, is the maximum amount of data the source will send before receiving an acknowledgment from the destination.
Connection establishment and termination	Process used to initialize port numbers and Sequence and Acknowledgment fields.
Ordered data transfer and data segmentation	A continuous stream of bytes from an upper-layer process that is "segmented" for transmission and delivered to upper-layer processes at the receiving device, with the bytes in the same order.

TCP Header

Technet24

TCP provides error recovery, but to do so, it consumes more bandwidth and uses more processing cycles than UDP. TCP and UDP rely on IP for end-to-end delivery. TCP is concerned with providing services to the applications of the sending and receiving computers. To provide all these services, TCP uses a variety of fields in its header (see Figure 31-2).

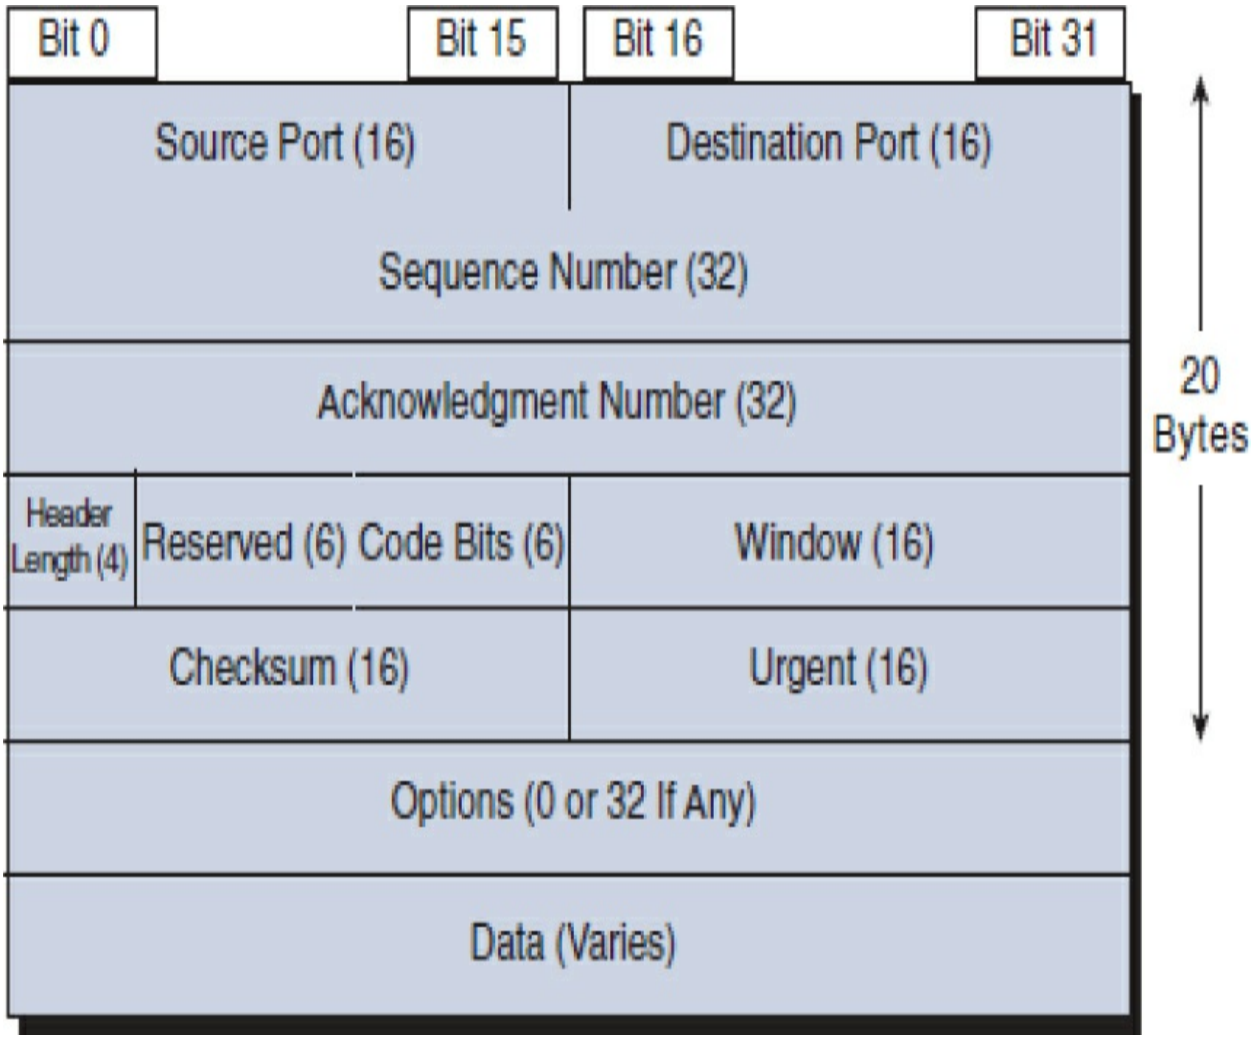

Figure 31-2 *TCP Header*

Port Numbers

The first two fields of the TCP header—the source and destination ports—are

also part of the UDP header (shown later, in Figure 31-7). Port numbers provide TCP (and UDP) with a way to multiplex multiple applications on the same computer. Web browsers now support multiple tabs or pages. Each time you open a new tab and request another web page, TCP assigns a different source port number and sometimes multiple port numbers. For example, you might have five web pages open. TCP almost always assigns destination port 80 for all five sessions. However, the source port for each is different. This is how TCP (and UDP) multiplexes the conversation so that the web browser knows in which tab to display the data.

TCP and UDP usually dynamically assign the source ports, starting at 1024 up to a maximum of 65535. Port numbers below 1024 are reserved for well-known applications. Table 31-5 lists several popular applications and their well-known port numbers.

Table 31-5 *Popular Applications and Their Well-Known Port Numbers*

Technet24

Port Number	Protocol	Application
20	TCP	FTP data
21	TCP	FTP control
22	TCP	SSH
23	TCP	Telnet
25	TCP	SMTP
53	UDP, TCP	DNS
67, 68	UDP	DHCP
69	UDP	TFTP
80	TCP	HTTP (WWW)
110	TCP	POP3
161	UDP	SNMP
443	TCP	HTTPS (SSL)
16384–32767	UDP	RTP-based voice (VoIP) and video

Error Recovery

TCP provides error recovery, also known as *reliability*, during data transfer sessions between two end devices that have established a connection. The Sequence and Acknowledgment fields in the TCP header track every byte of data transfer and ensure that missing bytes are retransmitted.

In Figure 31-3, the Acknowledgment field sent by the web client (4000) implies the next byte to be received; this is called *positive acknowledgment*.

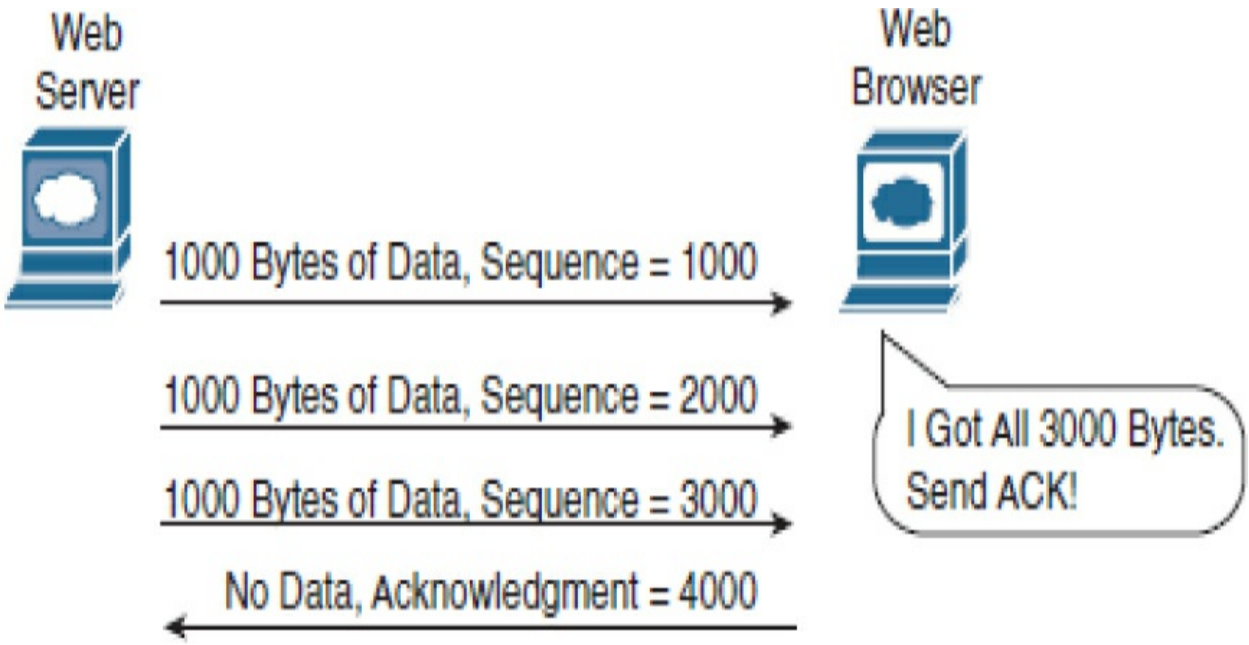

Figure 31-3 *TCP Acknowledgment Without Errors*

Figure 31-4 shows the same scenario, except now with some errors. The second TCP segment was lost in transmission. Therefore, the web client replies with an ACK field set to 2000. This is called a *positive acknowledgment with retransmission* (*PAR*) because the web client is requesting that some of the data be retransmitted. The web server now re-

Technet24

sends data starting at segment 2000. In this way, lost data is recovered.

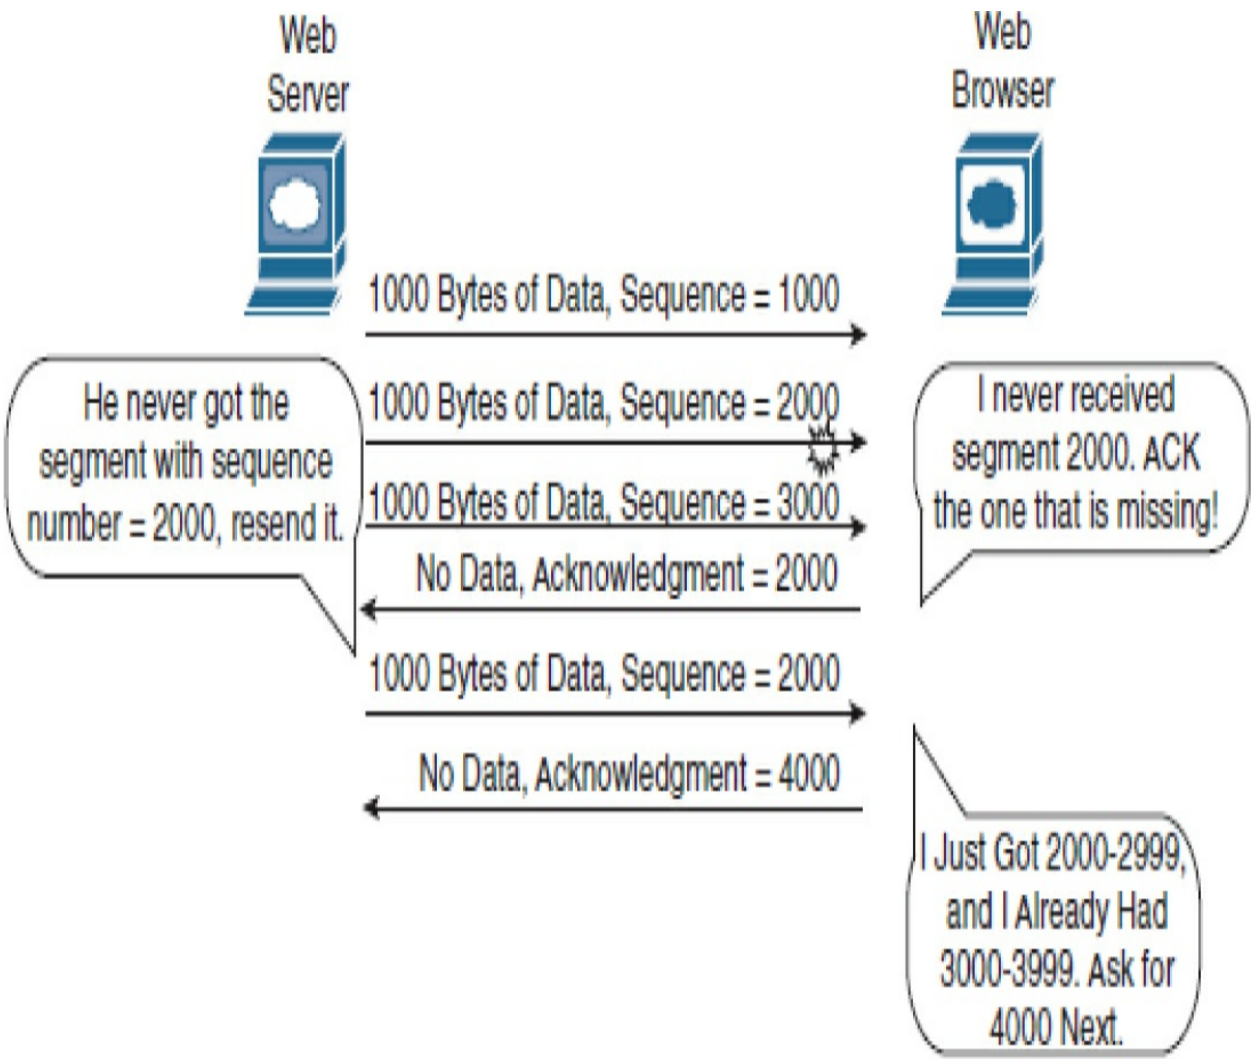

Figure 31-4 *TCP Acknowledgment with Errors*

Although not shown, the web server also sets a retransmission timer and awaits acknowledgment, just in case the acknowledgment is lost or all transmitted segments are lost. If that timer expires, the web server sends all segments again.

Flow Control

TCP handles flow control through a process called *windowing*. The two end devices negotiate the window size when initially establishing the connection; then they dynamically renegotiate window size during the life of the connection, increasing its size until it reaches the maximum window size of 65,535 bytes or until errors occur. Window size is specified in the Window field of the TCP header. After sending the amount of data specified in the window size, the source must receive an acknowledgment before sending the next window size of data.

Connection Establishment and Termination

Connection establishment is the process of initializing Sequence and Acknowledgment fields and agreeing on port numbers and window size. The three-way connection establishment phase shown in Figure 31-5 must occur before data transfer can proceed.

Figure 31-5 *TCP Connection Establishment*

Technet24

In the figure, DPORT and SPORT are the destination and source ports. SEQ is the sequence number. In bold are SYN and ACK, each representing a 1-bit flag in the TCP header used to signal connection establishment. TCP initializes the Sequence Number and Acknowledgment Number fields to any number that fits into the 4-byte fields. The initial Sequence Number is a random 32-bit number generated with each new transmission. The Acknowledgment Number is received back and increments the sender's sequence number by 1.

When data transfer is complete, a four-way termination sequence occurs. This sequence uses an additional flag, called the FIN bit (see Figure 31-6).

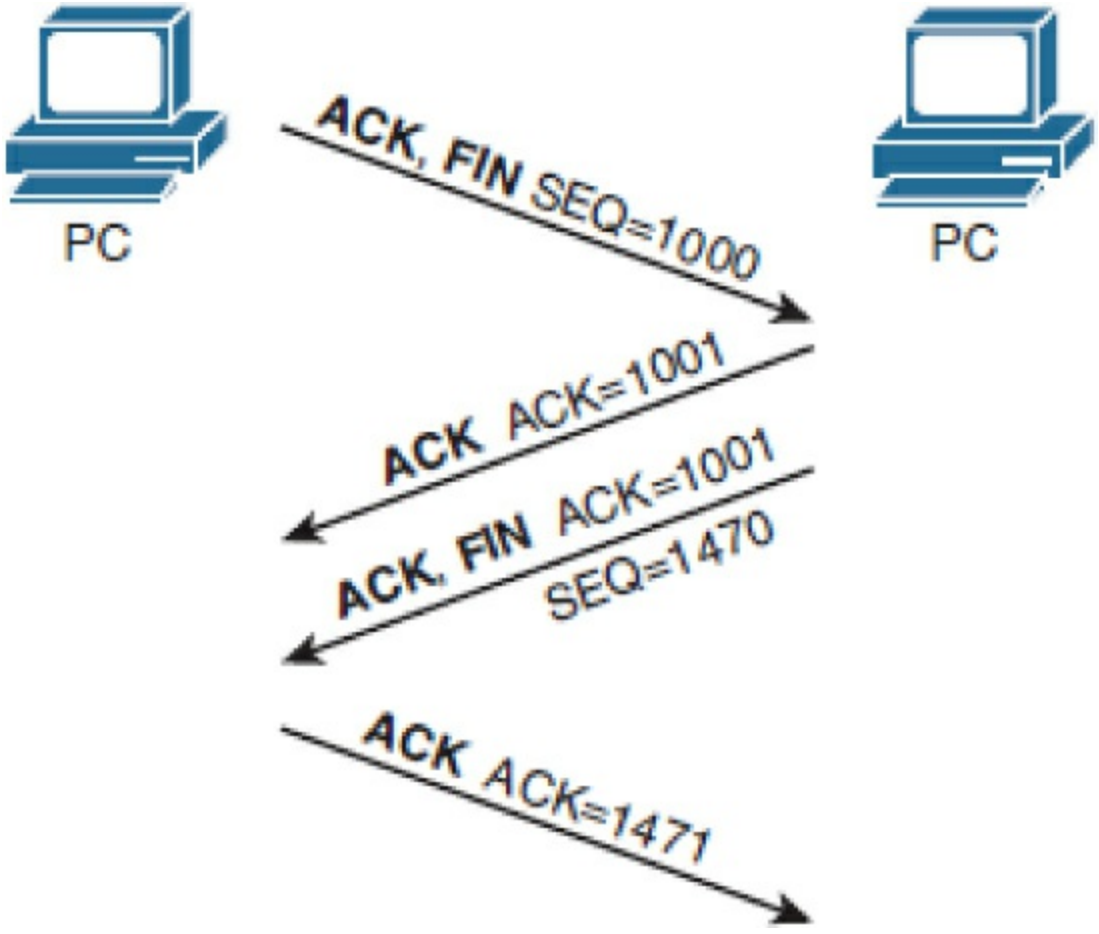

Figure 31-6 *TCP Connection Termination*

UDP

TCP establishes and terminates connections between endpoints, whereas UDP does not. Therefore, UDP is called a *connectionless protocol*. It provides no reliability, no windowing, and no reordering of the data. However, UDP does provide data transfer and multiplexing using port numbers, and it does so with fewer bytes of overhead and less processing than TCP. Applications that use UDP, such as VoIP, trade the possibility of some data loss for less delay. Figure 31-7 compares the two headers.

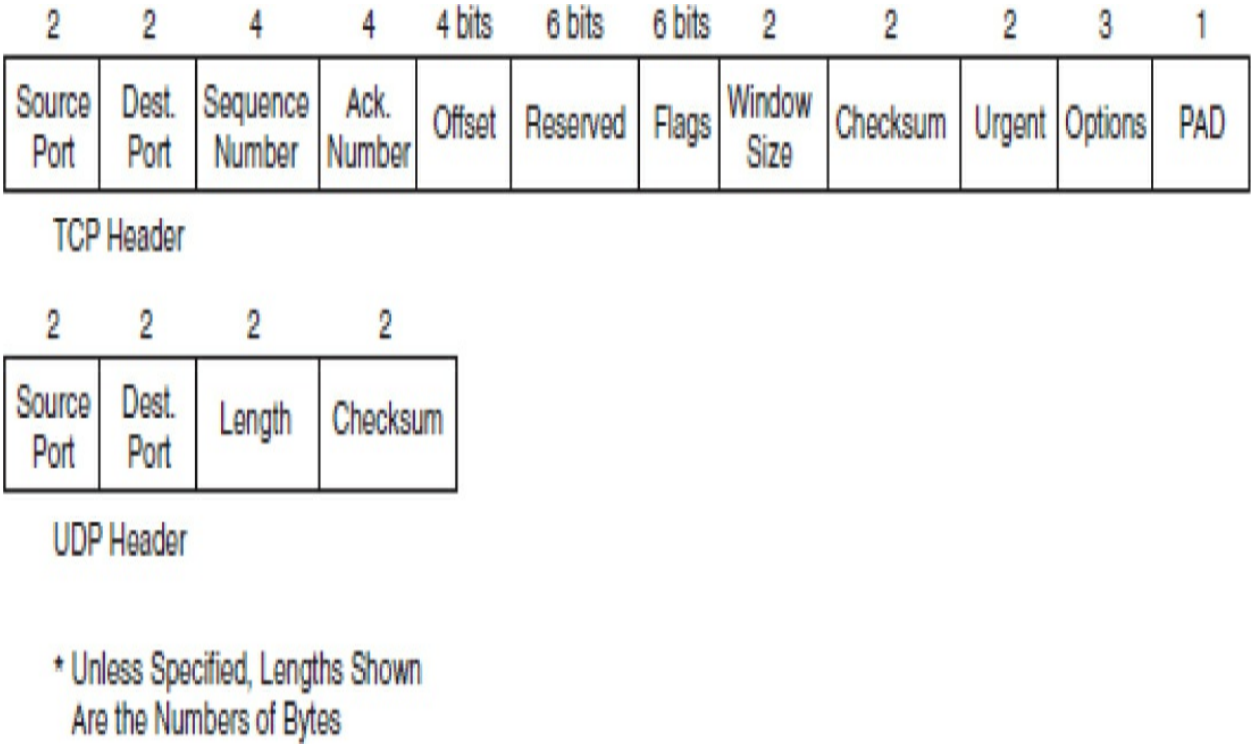

Figure 31-7 *TCP and UDP Headers*

The TCP/IP Internet Layer

Technet24

The Internet layer of the TCP/IP model and its Internet Protocol (IP) define addresses so that each host computer can have a different IP address. In addition, the Internet layer defines the process of routing so that routers can determine the best path for sending packets to the destination. Continuing with the web page example, IP addresses the data as it passes from the transport layer to the Internet layer:

The web client sends an HTTP request.

TCP encapsulates the HTTP request.

IP encapsulates the transport segment into a packet, adding source and destination addresses.

Lower layers process and send the request to the web server.

The web server receives HTTP requests and sends a TCP acknowledgment back to the requesting web client.

The web server sends the HTTP response down to the transport layer.

TCP encapsulates the HTTP data.

IP encapsulates the transport segment into a packet, adding source and destination addresses.

Lower layers process and send the response to the requesting web client.

The requesting web client sends an acknowledgment back to the web server.

The operation of IP includes not only addressing but also the process of routing the data from source to destination. IP is further discussed and reviewed in the upcoming days.

The TCP/IP Network Access Layer

IP depends on the network access layer to deliver IP packets across a physical network. Therefore, the network access layer defines the protocols and hardware required to deliver data across some physical network by specifying exactly how to physically connect a networked device to the physical media over which data can be transmitted.

The network access layer includes many protocols to deal with the different types of media that data can cross on its way from source device to destination device. For example, data might need to travel first on an Ethernet link and then cross a Point-to-Point (PPP) link, then a Frame Relay link, then a Multiprotocol Label Switching (MPLS) link, and then finally an Ethernet link to reach the destination. At each transition from one media type to another, the network access layer provides the protocols, cabling standards, headers, and trailers to send data across the physical network.

Many times, a local link address is needed to transfer data from one hop to the next. For example, in an Ethernet LAN, Media Access Control (MAC) addresses are used between the sending device and its local gateway router.

Technet24

At the gateway router (depending on the needs of the outbound interface), the Ethernet header might be replaced with an MPLS label. The label serves the same purpose as MAC addresses in Ethernet: to get the data across the link from one hop to the next so that the data can continue its journey to the destination. Some protocols, such as PPP, do not need a link address because only one other device on the link can receive the data.

With the network access layer, we can now finalize our web page example. The following greatly simplifies and summarizes the process of requesting and sending a web page:

The web client sends an HTTP request.

TCP encapsulates the HTTP request.

IP encapsulates the transport segment into a packet, adding source and destination addresses.

The network access layer encapsulates the packet in a frame, addressing it for the local link.

The network access layer sends the frame as bits on the media.

Intermediary devices process the bits at the network access and Internet layers and then forward the data toward the destination.

The web server receives the bits on the physical interface and sends them up

through the network access and Internet layers.

The web server sends a TCP acknowledgment back to the requesting web client.

The web server sends the HTTP response down to the transport layer.

TCP encapsulates the HTTP data.

IP encapsulates the transport segment into a packet, adding source and destination addresses.

The network access layer encapsulates the packet in a frame, addressing it for the local link.

The network access layer sends the frame as bits on the media.

Lower layers process and send the response to the requesting web client.

The response travels back to the source over multiple data links.

The requesting web client receives the response on the physical interface and sends the data up through the network access and Internet layers.

The requesting web client sends a TCP acknowledgment back to the web server.

The web page is displayed in the requesting device’s browser.

Technet24

Data Encapsulation Summary

Each layer of the TCP/IP model adds its own header information. As the data travels down through the layers, it is encapsulated with a new header. At the network access layer, a trailer is also added. This encapsulation process is described in five steps:

Create and encapsulate the application data with any required application layer headers. For example, the HTTP OK message can be returned in an HTTP header, followed by part of the contents of a web page.

Encapsulate the data supplied by the application layer inside a transport layer header. For end-user applications, a TCP or UDP header is typically used.

Encapsulate the data supplied by the transport layer inside an Internet layer (IP) header. IP is the only protocol available in the TCP/IP network model at the Internet layer.

Encapsulate the data supplied by the Internet layer inside a network access layer header and trailer. This is the only layer that uses both a header and a trailer.

Transmit the bits. The physical layer encodes a signal onto the medium to transmit the frame.

The numbers in Figure 31-8 correspond to the five steps in the list,

graphically showing the same encapsulation process.

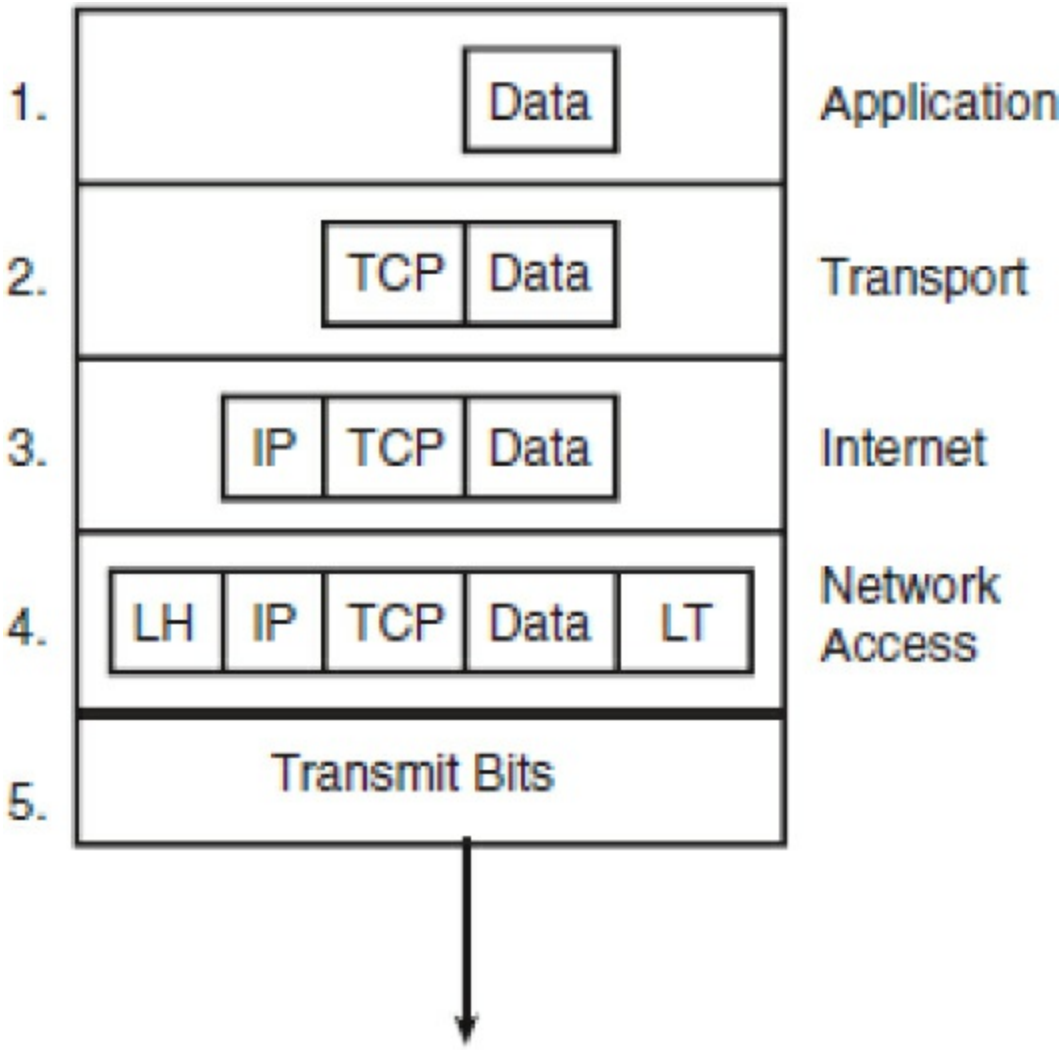

Figure 31-8 *Five Steps of Data Encapsulation*

Note

The letters LH and LT stand for link header and link trailer, respectively, and refer to the data link layer header and trailer.

Networking Icons

Technet24

In order to interpret networking diagrams or topologies, you must understand the symbols or icons used to represent different networking devices and media. The icons in Figure 31-9 are the most common networking symbols for CCNA studies.

Access Point

PC

Laptop

Server

IP Phone

Router

Switch

Layer 3 Switch

Hub

Bridge

Cable (Various)

Serial Line

Virtual Circuit

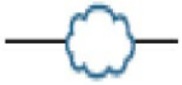

Ethernet WAN

Wireless

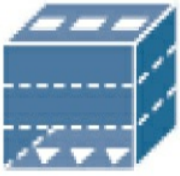

SDN Controller

vSwitch

IPS

ASA

Firewall

Network Cloud

Cable Modem

DSLAM

Technet24

Figure 31-9 *Networking Icons*

Devices

In today's wired networks, switches are almost exclusively used to connect end devices to a single LAN. Occasionally, you might see a hub connecting end devices, but hubs are really legacy devices. The following are the differences between a hub and a switch:

• Hubs were typically chosen as intermediary devices within very small LANs, in which bandwidth usage was not an issue or cost limitations were a factor. In today's networks, switches have replaced hubs.

• Switches replaced hubs as local-area network (LAN) intermediary devices because a switch can segment collision domains and provide enhanced security.

Switches

When choosing a switch, these are the main factors to consider:

• **Cost:** The cost is determined by the number and type of ports, network management capabilities, embedded security technologies, and optional advanced switching technologies.

• **Interface characteristics:** The number of ports must be sufficient both for now and for future expansion. Other characteristics include uplink speeds, a

mixture of UTP and fiber, and modularity.

• **Hierarchical network layer:** Switches at the access layer have different requirements than switches at the distribution or core layers.

Access Layer Switches

Access layer switches facilitate the connection of end devices to the network. Features of access layer switches include the following:

• Port security

• VLANs

• Fast Ethernet/Gigabit Ethernet

• Power over Ethernet (PoE)

• Link aggregation

• Quality of service (QoS)

Distribution Layer Switches

Distribution layer switches receive the data from the access layer switches and forward it to the core layer switches. Features of distribution layer switches include the following:

Technet24

• Layer 3 support

• High forwarding rate

• Gigabit Ethernet/10 Gigabit Ethernet

• Redundant components

• Security policies/access control lists

• Link aggregation

• QoS

Core Layer Switches

Core layer switches make up the backbone and are responsible for handling the majority of data on a switched LAN. Features of core layer switches include the following:

• Layer 3 support

• Very high forwarding rate

• Gigabit Ethernet/10 Gigabit Ethernet

• Redundant components

• Link aggregation

• QoS

Routers

Routers are the primary devices used to interconnect networks—LANs, WANs, and WLANs. When choosing a router, the main factors to consider are as follows:

• **Expandability:** Provides flexibility to add new modules as needs change.

• **Media:** Determines the type of interfaces the router needs to support for the various network connections.

• **Operating system features:** Determines the version of IOS loaded on the router. Different IOS versions support different feature sets. Features to consider include security, QoS, VoIP, and routing complexity, among others.

Figure 31-10 shows a Cisco 4321 router, which provides the following connections:

• **Console ports:** Two console ports for the initial configuration, using a regular RJ-45 port and a USB Type-B (mini-B USB) connector.

• **AUX port:** An RJ-45 port for remote management access.

Technet24

• **LAN interfaces:** Two Gigabit Ethernet interfaces for LAN access (G0/0/0 and G0/0/1). If the RJ-45 G0/0/0 port is used, then the small form-factor pluggable (SFP) port cannot be used. WAN services would then be provided through an expansion card in the network interface module (NIM) slots.

• **Ethernet WAN:** The other G0/0/0 physical port, an SFP port that would support various Ethernet WAN connections, typically fiber. If it is used, the Gi0/0 RJ-45 port is disabled.

• **NIM slots:** Two slots that support different types of interface modules, including serial (shown in Figure 31-10), digital subscriber line (DSL), switch port, and wireless.

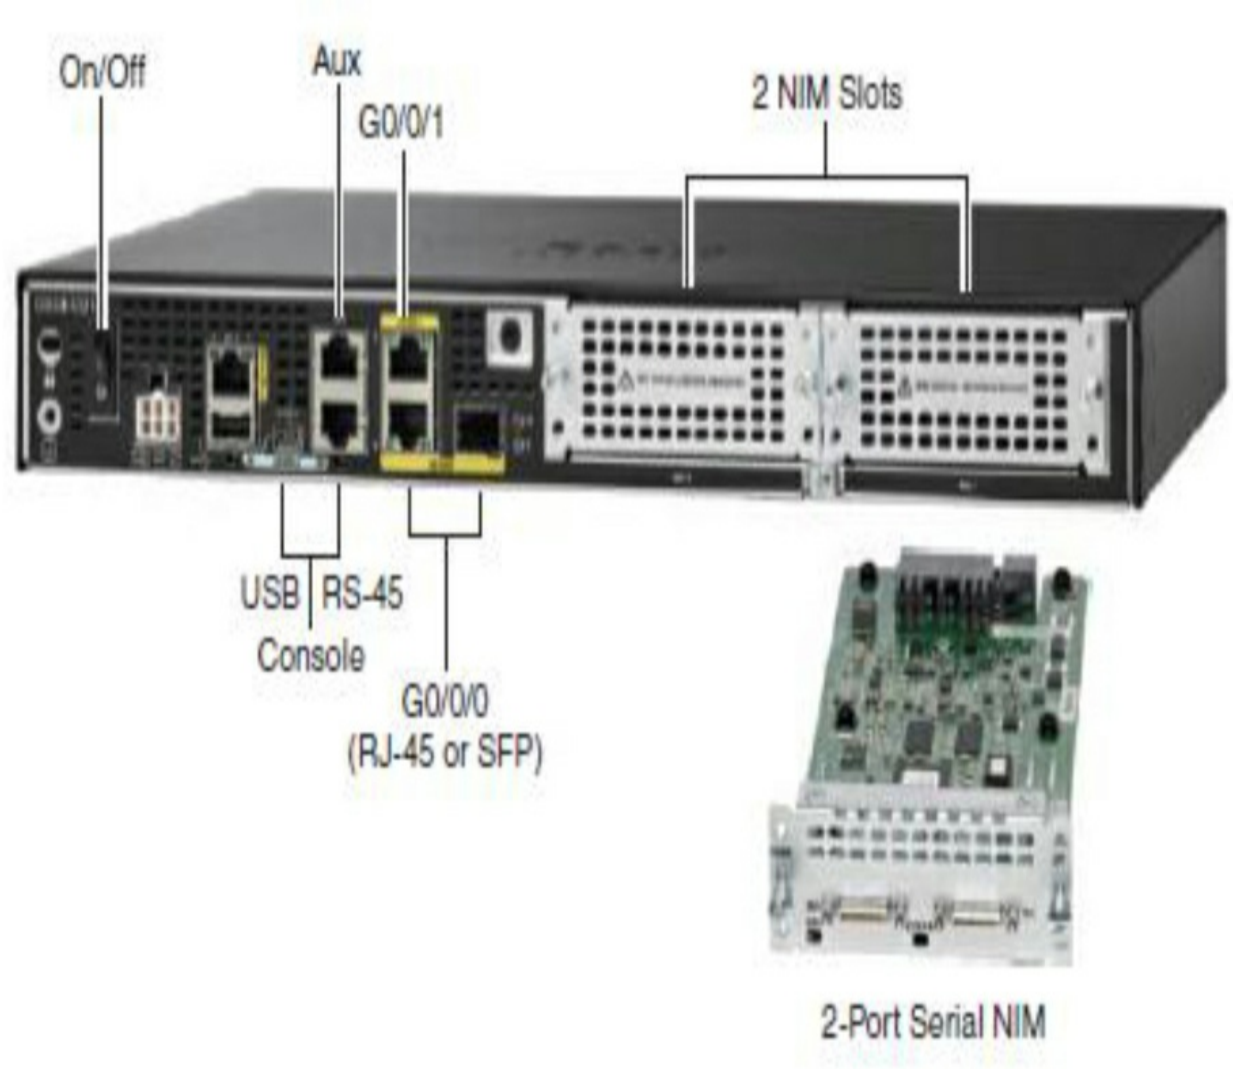

Figure 31-10 *Backplane of the Cisco 4321 Integrated Services Router (ISR)*

Specialty Devices

Switches and routers make up the backbone of a network. In addition, many networks integrate various specialized network devices.

Firewalls

A firewall is a networking device, either hardware or software based, that

Technet24

controls access to the organization’s network. This controlled access is designed to protect data and resources from outside threats.

Organizations implement software firewalls through a network operating system (NOS) such as Linux/UNIX, Windows servers, and macOS servers. The firewall is configured on the server to allow or block certain types of network traffic. Hardware firewalls are often dedicated network devices that can be implemented with little configuration.

Figure 31-11 shows a basic stateful firewall.

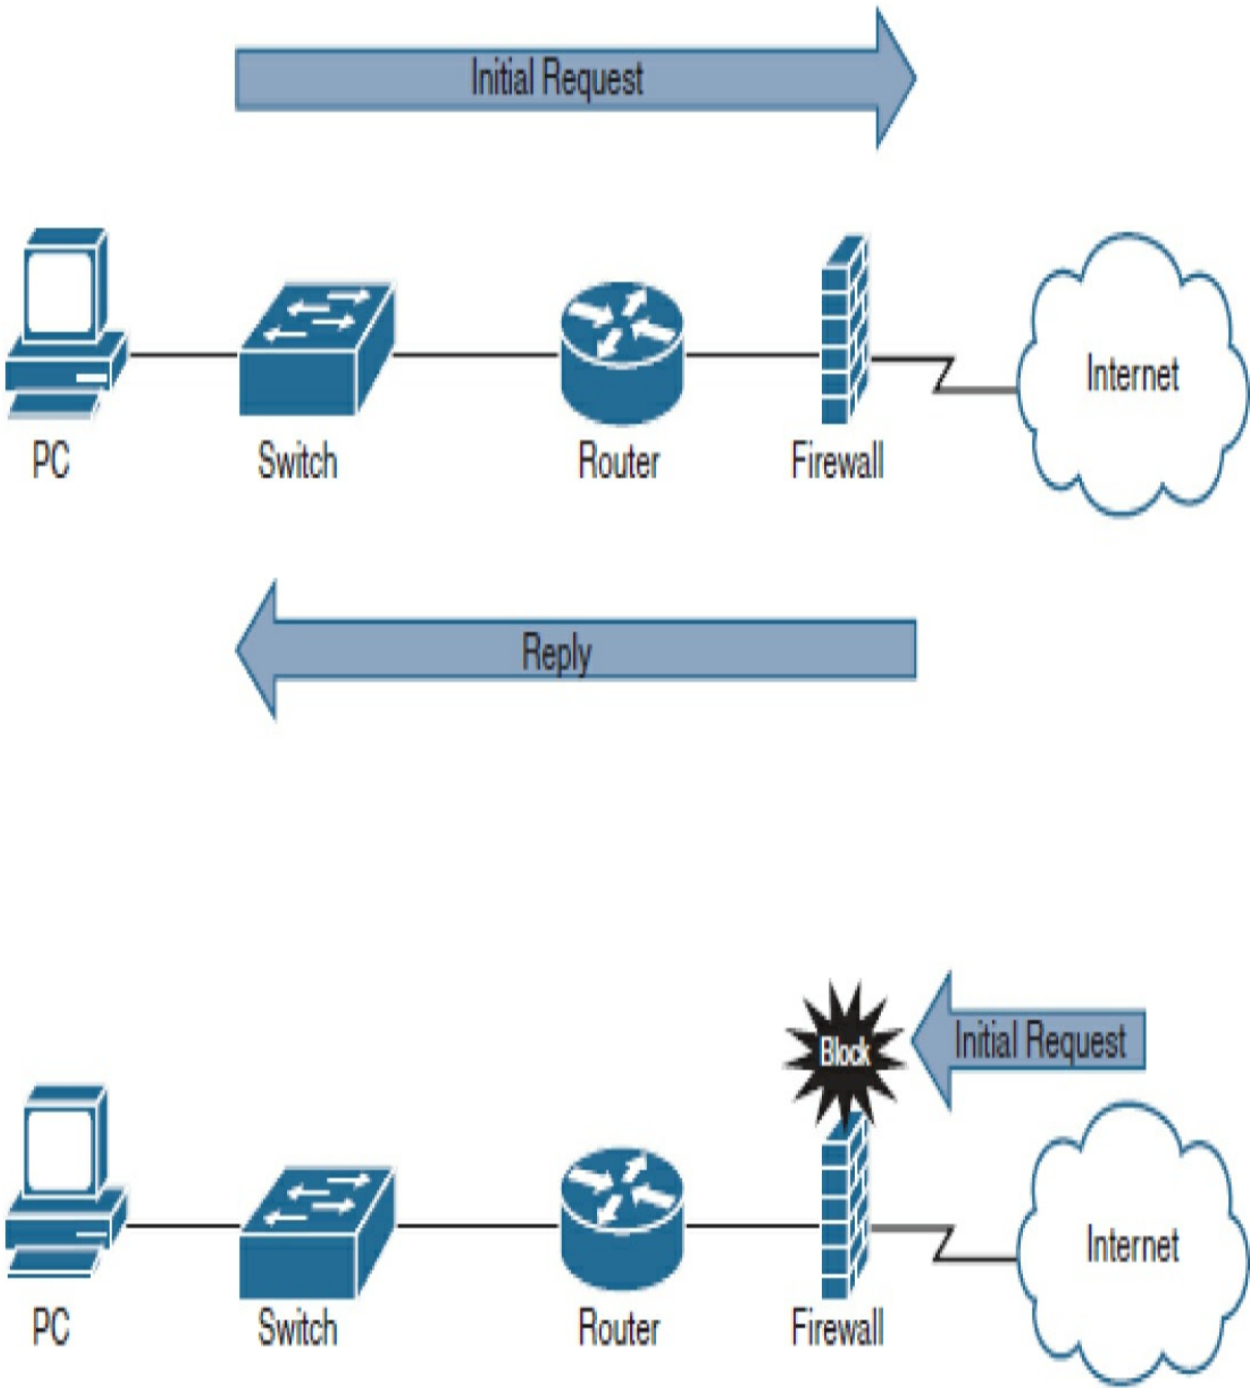

Figure 31-11 *The Function of a Firewall*

A stateful firewall allows traffic to originate from an inside, trusted network and go out to an untrusted network, such as the Internet. The firewall allows return traffic that comes back from the untrusted network to the trusted

Technet24

network. However, the firewall blocks traffic that originates from an untrusted network.

IDS and IPS

Both intrusion detection systems (IDS) and intrusion prevention systems (IPS) can recognize network attacks; they differ primarily in their network placement. An IDS device receives a copy of traffic to be analyzed. An IPS device is placed inline with the traffic, as Figure 31-12 shows.

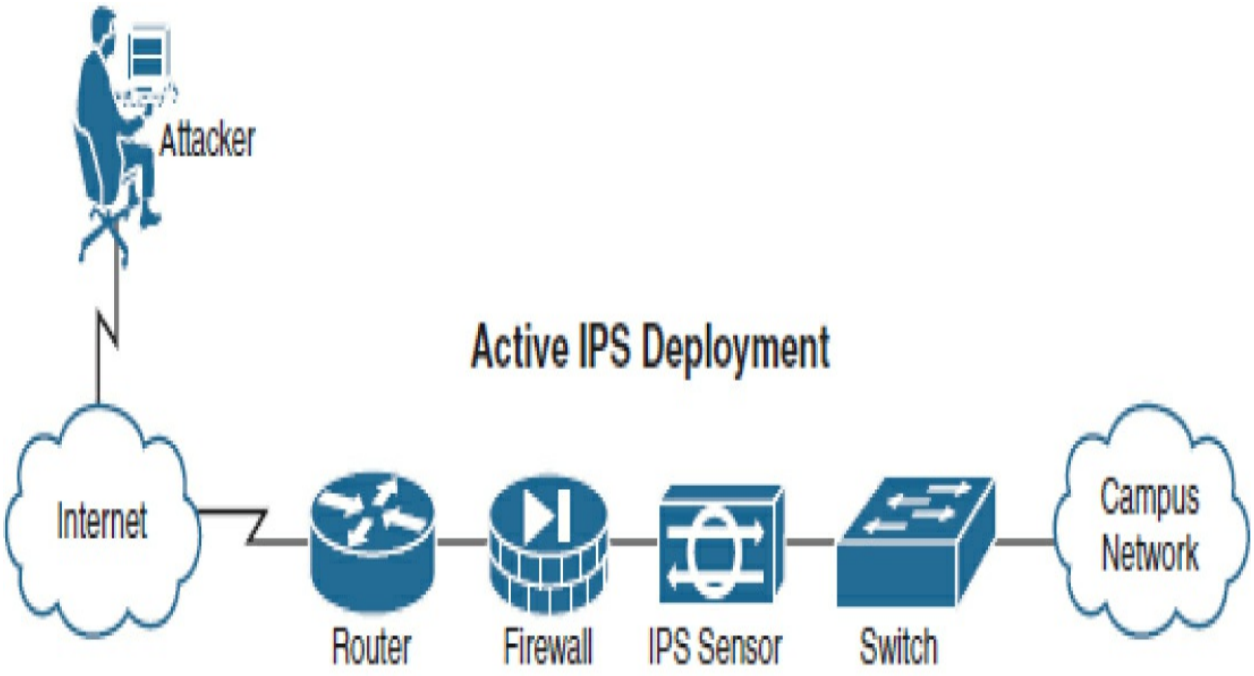
Attacker
Active IPS Deployment
Internet
Router
Firewall
IPS Sensor
Switch
Campus
Network

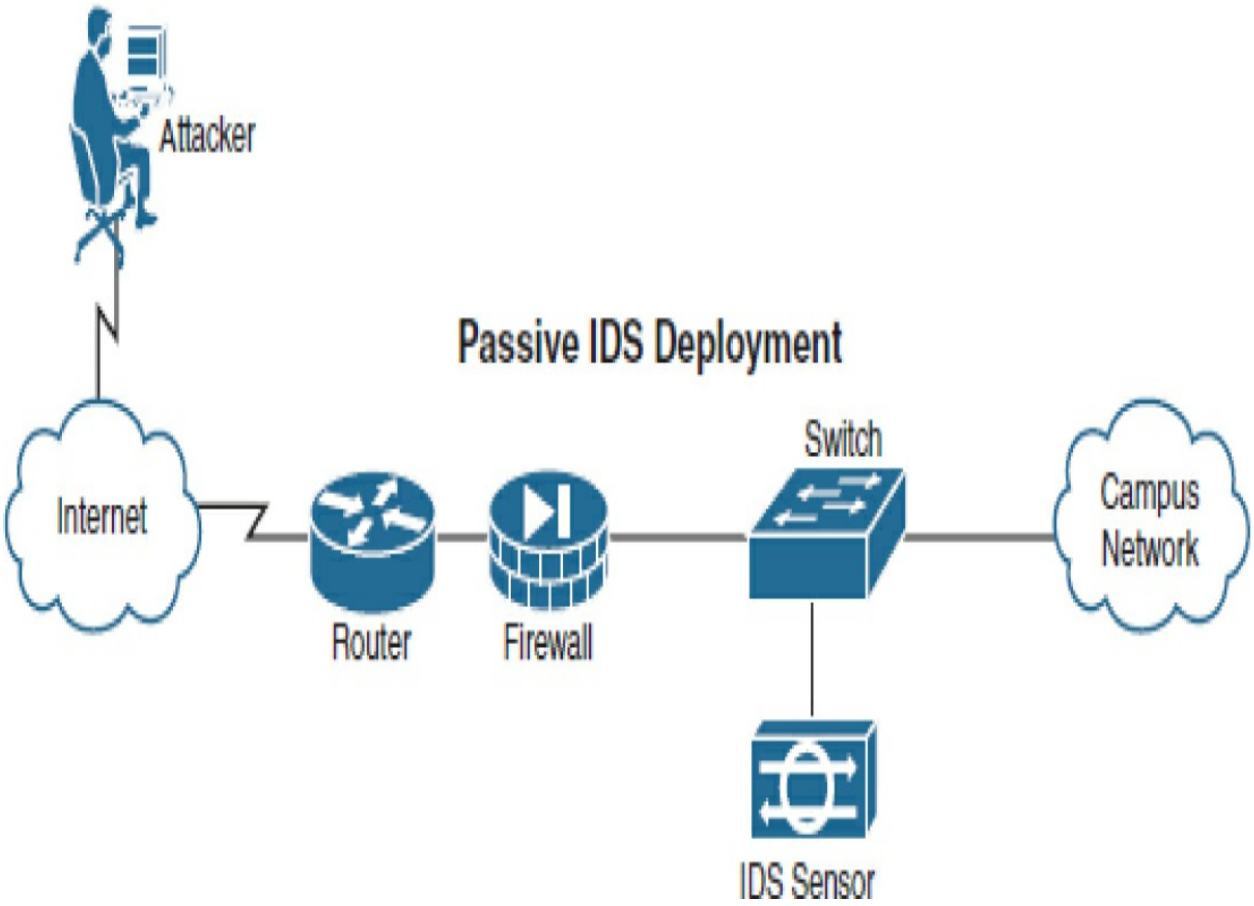
Attacker
Passive IDS Deployment
Internet
Router
Firewall
Switch
Campus
Network
IDS Sensor

Technet24

Figure 31-12 *IPS and IDS Comparison*

An IDS is a passive detection system. It can detect the presence of an attack, log the information, and send an alert.

An IPS has the same functionality as an IDS, but in addition, an IPS is an active device that continually scans the network, looking for inappropriate activity. It can shut down any potential threats. The IPS looks for any known signatures of common attacks and automatically tries to prevent those attacks.

Next-Generation Firewalls

Although the term *next-generation* in relation to firewalls has been around at least since the earlier 2010s, it can be misleading. Next-generation firewalls (NGFWs) or next-generation IPSs (NGIPSs) are actually what Cisco currently sells as its Cisco Adaptative Security Appliance (ASA) and Firepower product lines. Be sure to visit www.cisco.com/go/firewalls for more information on Cisco's current firewall offerings.

An NGFW typically has the following features:

• **Traditional firewall:** An NGFW performs traditional firewall functions, such as stateful firewall filtering, NAT/PAT, and VPN termination.

• **Application Visibility and Control (AVC):** AVC makes it possible to look deeply into the application layer data to identify the application to defend against attacks that use random port numbers.

• **Advanced Malware Protection (AMP):** AMP can block file transfers that would install malware and save copies of files for later analysis.

• **Uniform resource locator (URL) filtering:** URL filtering examines the URLs in each web request, categorizes the URLs, and either filters or rate limits the traffic based on rules. The Cisco Talos security group monitors and creates reputation scores for each domain known in the Internet, and URL filtering can use those scores in its decisions to categorize, filter, or rate limit.

• **NGIPS:** Cisco's NGFW products can also run their NGIPS feature along with the firewall, as shown in Figure 31-13.

Technet24

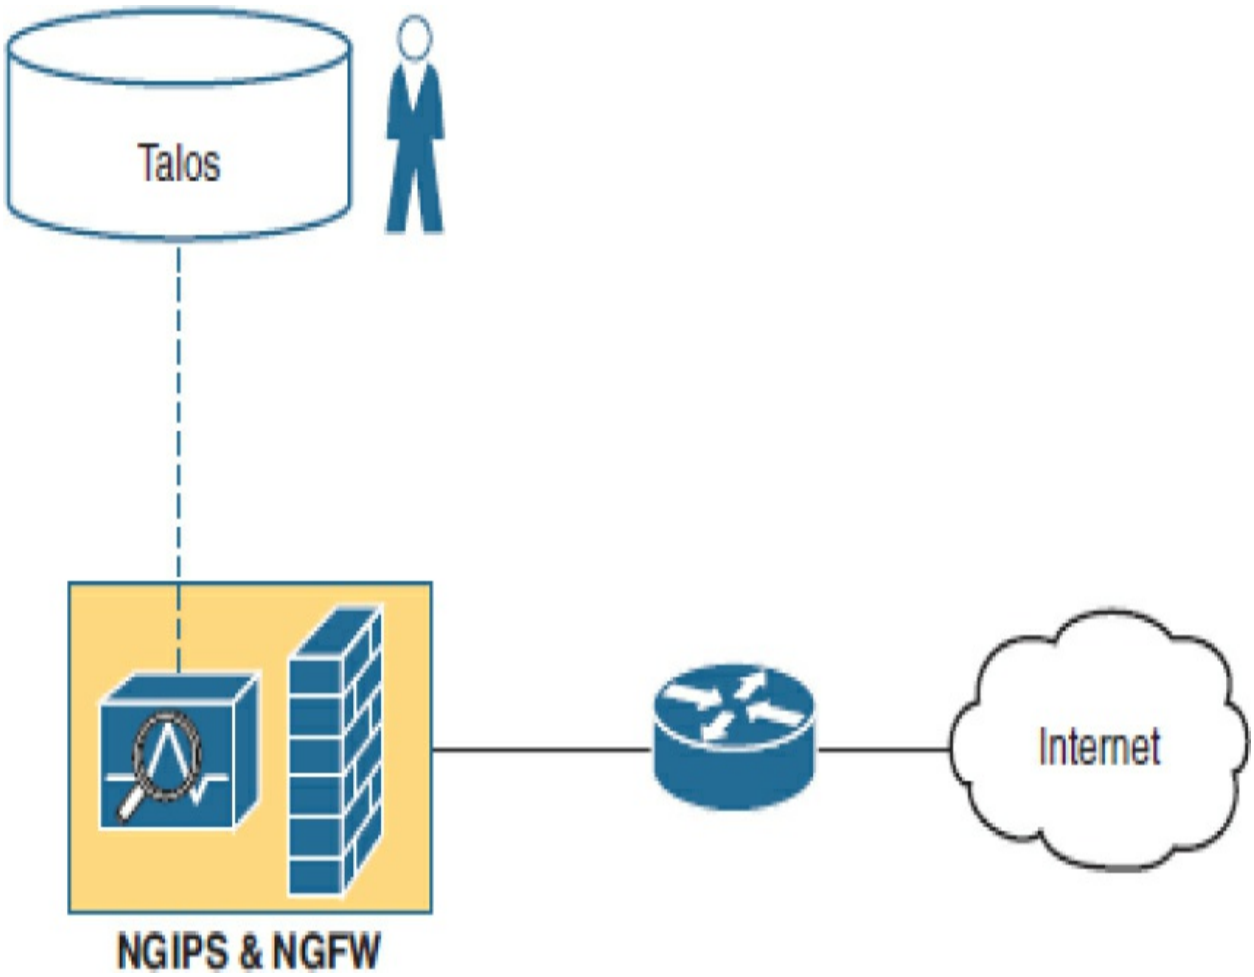

Figure 31-13 *NGFW with NPIPS Module*

Access Points and Wireless LAN Controllers

Wireless LANs (WLANs) are commonly used in networks. Users expect to be able to connect seamlessly as they move from location to location within a home, a small business, or an enterprise campus network. To enable this connectivity, network administrators manage a collection of wireless access points (APs) and wireless LAN controllers (WLCs).

In small networks, APs are typically used when a router is already providing

Layer 3 services, as in Figure 31-14.

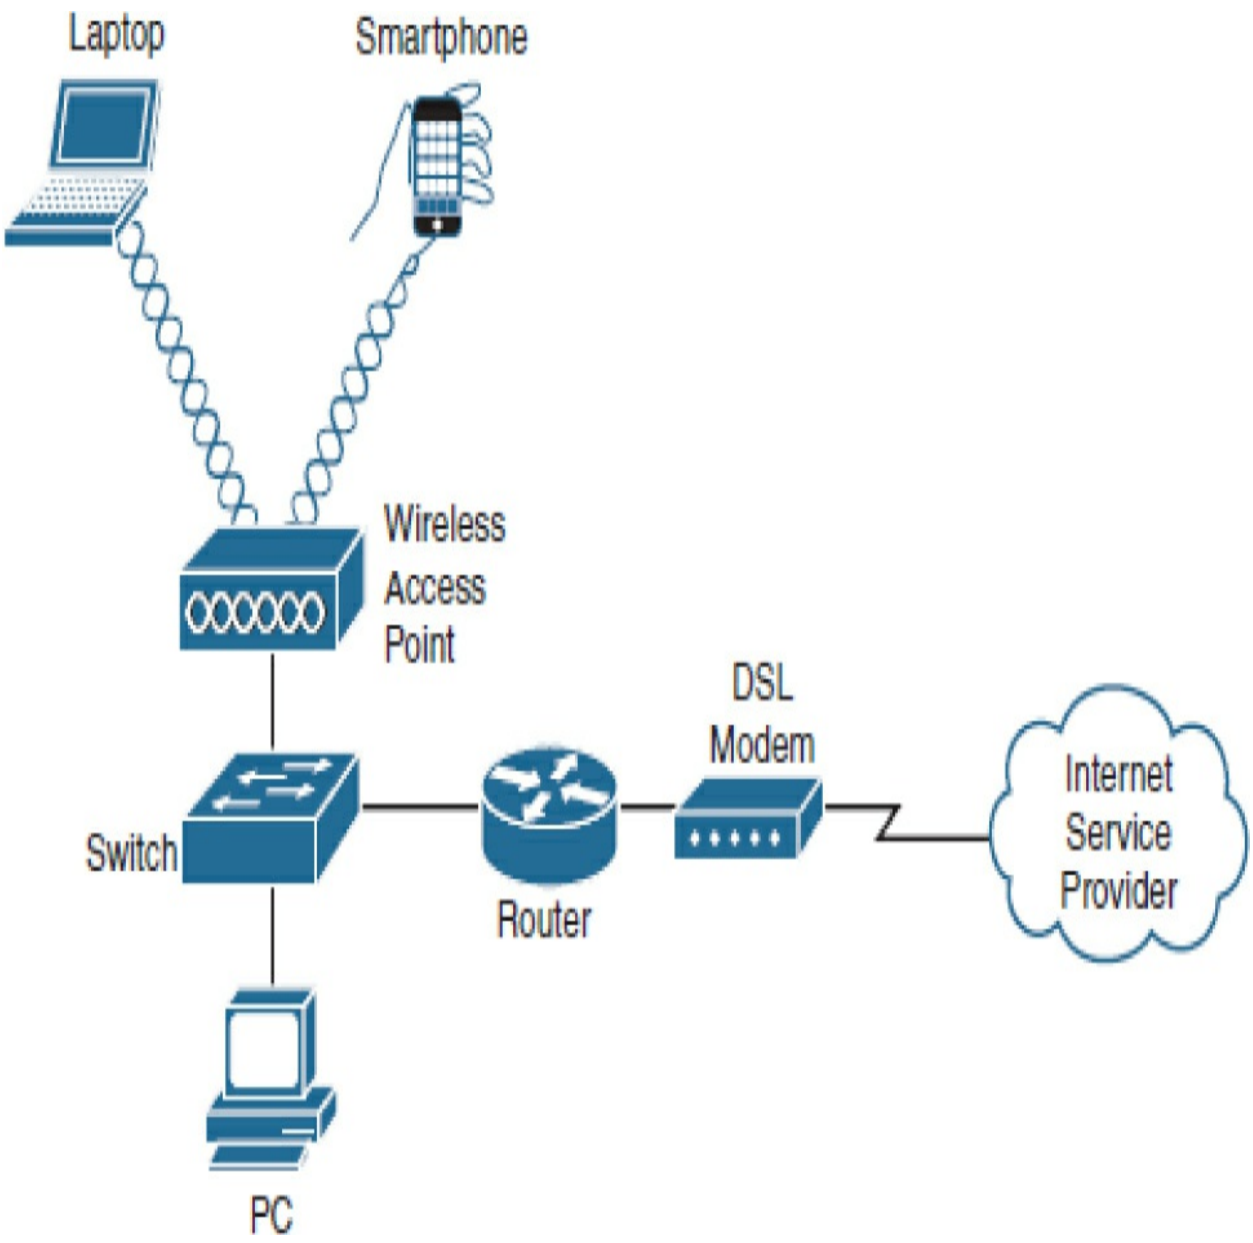

Figure 31-14 *Small Network with an AP*

An AP has an Ethernet port that enables it to be connected to a switch port. In a home or small office network, an AP can be another wireless router with all the Layer 3 services turned off: You simply connect one of the AP's switch

Technet24

ports to one of the switch ports on the wireless router.

APs are also used when the coverage area of an existing WLAN needs to be extended. In larger networks, a wireless LAN controller (WLC) is typically used to manage multiple APs, as in Figure 31-15.

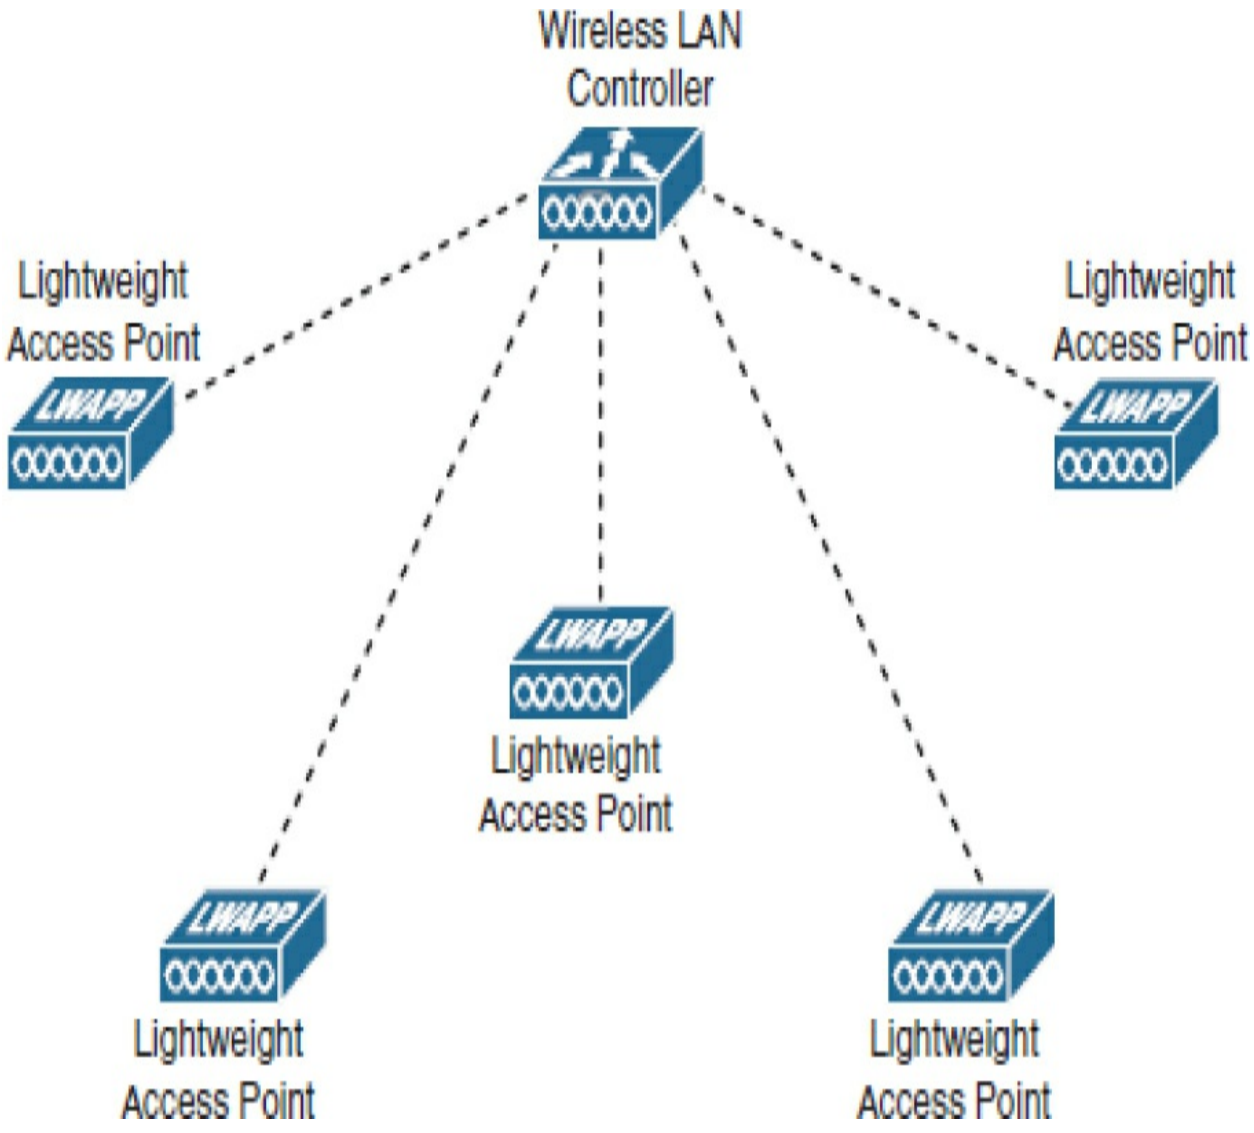

Figure 31-15 *Example of a Wireless LAN Controller Implementation*

WLCs can use the older Lightweight Access Point Protocol (LWAPP) or the

more current Control and Provisioning of Wireless Access Points (CAPWAP). With a WLC, VLAN pooling can be used to assign IP addresses to wireless clients from a pool of IP subnets and their associated VLANs.

Physical Layer

Before any network communications can occur, a wired or wireless physical connection must be established. The type of physical connection depends on the network setup. In larger networks, switches and APs are often two separate dedicated devices. In a very small business (three or four employees) or home network, wireless and wired connections are combined into one device and include a broadband method of connecting to the Internet. These wireless broadband routers offer a switching component with multiple ports and an AP, which allows wireless devices to connect as well. Figure 31-16 shows the backplane of a Cisco WRP500 Wireless Broadband Router.

Technet24

Figure 31-16 *Cisco RV160W Wireless-AC VPN Router*

Network Media Forms and Standards

Three basic forms of network media exist:

• **Copper cable:** The signals are patterns of electrical pulses.

• **Fiber-optic cable:** The signals are patterns of light.

• **Wireless:** The signals are patterns of microwave transmissions.

Messages are encoded and then placed onto the media. Encoding is the process of converting data into patterns of electrical, light, or electromagnetic energy so that it can be carried on the media.

Table 31-6 summarizes the three most common networking media in use today.

Table 31-6 *Networking Media*

Technet24

Media	Physical Components	Frame Encoding Technique	Signaling Methods
Copper cable	UTP Coaxial Connectors NICs Ports Interfaces	Manchester encoding Nonreturn to zero (NRZ) techniques 4B/5B codes used with Multi-Level Transition Level 3 (MLT-3) signaling 8B/10B PAM5	Changes in the electromagnetic field. Intensity of the electromagnetic field. Phase of the electromagnetic wave.
Fiber-optic cable	Single-mode fiber Multimode fiber Connectors NICs Interfaces Lasers and LEDs Photoreceptors	Pulses of light Wavelength multiplexing using different colors	A pulse equals 1. No pulse is 0.
Wireless	Access points NICs Radio Antennas	Direct Sequence Spread Spectrum (DSSS) Orthogonal Frequency Division Multiplexing (OFDM)	Radio waves.

Each media type has advantages and disadvantages. When choosing the media, consider each of the following:

• **Cable length:** Does the cable need to span a room or run from building to building?

• **Cost:** Does the budget allow for using a more expensive media type?

• **Bandwidth:** Does the technology used with the media provide adequate bandwidth?

• **Ease of installation:** Does the implementation team have the capability to install the cable, or is a vendor required?

• **Susceptibility to EMI/RFI:** Will the local environment interfere with the signal?

Table 31-7 summarizes the media standards for LAN cabling.

Table 31-7 *Media Standard, Cable Length, and Bandwidth*

Technet24

Ethernet Type	Bandwidth	Cable Type	Maximum Distance
10BASE-T	10 Mbps	Cat3/Cat5 UTP	100 m
100BASE-TX	100 Mbps	Cat5 UTP	100 m
100BASE-TX	200 Mbps	Cat5 UTP	100 m
100BASE-FX	100 Mbps	Multimode fiber	400 m
100BASE-FX	200 Mbps	Multimode fiber	2 km
1000BASE-T	1 Gbps	Cat5e UTP	100 m
1000BASE-TX	1 Gbps	Cat6 UTP	100 m
1000BASE-SX	1 Gbps	Multimode fiber	550 m
1000BASE-LX	1 Gbps	Single-mode fiber	2 km
10GBASE-T	10 Gbps	Cat6a/Cat7 UTP	100 m
10GBASE-SX4	10 Gbps	Multimode fiber	550 m
10GBASE-LX4	10 Gbps	Single-mode fiber	2 km

LAN Device Connection Guidelines

End devices are pieces of equipment that are either the original source or the final destination of a message. Intermediary devices connect end devices to the network to assist in getting a message from the source end device to the destination end device.

Connecting devices in a LAN is usually done with unshielded twisted pair (UTP) cabling. Although many newer devices have an automatic crossover feature that enables you to connect either a straight-through or a crossover cable, you still need to know the following basic rules:

Use straight-through cables for the following connections:

• Switch to router Ethernet port

• Computer to switch

• Computer to hub

Use crossover cables for the following connections:

• Switch to switch

• Switch to hub

• Hub to hub

Technet24

• Router to router (Ethernet ports)

• Computer to computer

• Computer to router Ethernet port

LANs and WANs

A local-area network (LAN) is a network of computers and other components located relatively close together in a limited area. LANs can vary widely in size, from one computer connected to a router in a home office, to hundreds of computers in a corporate office. However, in general, a LAN spans a limited geographic area. The fundamental components of a LAN include the following:

• Computers

• Interconnections (NICs and the media)

• Networking devices (hubs, switches, and routers)

• Protocols (Ethernet, IP, ARP, DHCP, DNS, and so on)

A wide-area network (WAN) generally connects LANs that are geographically separated. A collection of LANs connected by one or more WANs is called an *internetwork*; thus, we have the Internet. The term *intranet* is often used to refer to a privately owned connection of LANs and

WANs.

Depending on the type of service, connecting to the WAN normally works in one of the following ways:

• 60-pin serial connection to a CSU/DSU (legacy)

• RJ-45 T1 controller connection to a CSU/DSU (legacy)

• RJ-11 connection to a dialup or DSL modem

• Cable coaxial connection to a cable modem

• Fiber Ethernet connection to service provider’s switch

Small Office/Home Office (SOHO)

With the growing number of remote workers, enterprises have an increasing need for secure, reliable, and cost-effective ways to connect people working in small offices or home offices (SOHO) or other remote locations to resources on corporate sites. For SOHO workers, this is typically done through a cable or DSL connection, as shown in Figure 31-17.

Technet24

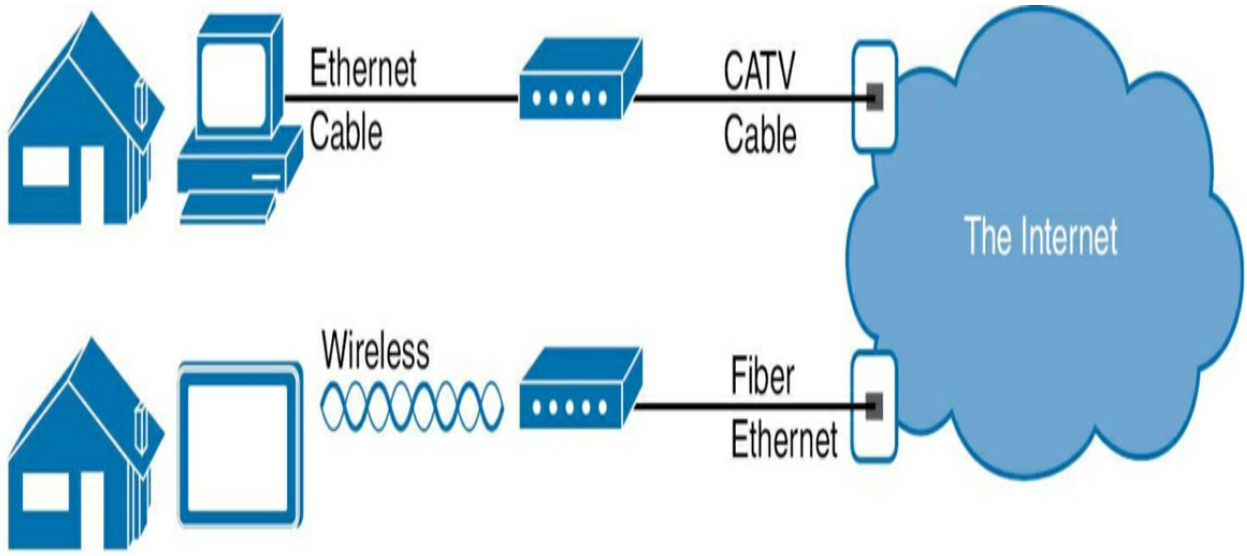

Figure 31-17 *SOHO Connections to the Internet*

Remote connection technologies to support teleworkers include the following:

• Traditional private WAN technologies, including Frame Relay, ATM, and leased lines, although these technologies are now considered legacy

• Remote secure virtual private network (VPN) access through a broadband connection over the public Internet

Components needed for teleworker connectivity include the following:

• **Home office components:** Computer, broadband access (cable or DSL), and a VPN router or VPN client software installed on the computer

• **Corporate components:** VPN-capable routers, VPN concentrators, multifunction security appliances, authentication, and central management

devices for resilient aggregation and termination of the VPN connections

SOHO Routers

The gateway to the Internet for a SOHO is typically an integrated multifunction routers. SOHO routers have the following features:

• Use the Internet and VPN technology for their WAN connections to send data back and forth to the rest of the enterprise.

• Use a multifunction device that handles routing, LAN switching, VPN, wireless, and other features, as shown in Figure 31-18.

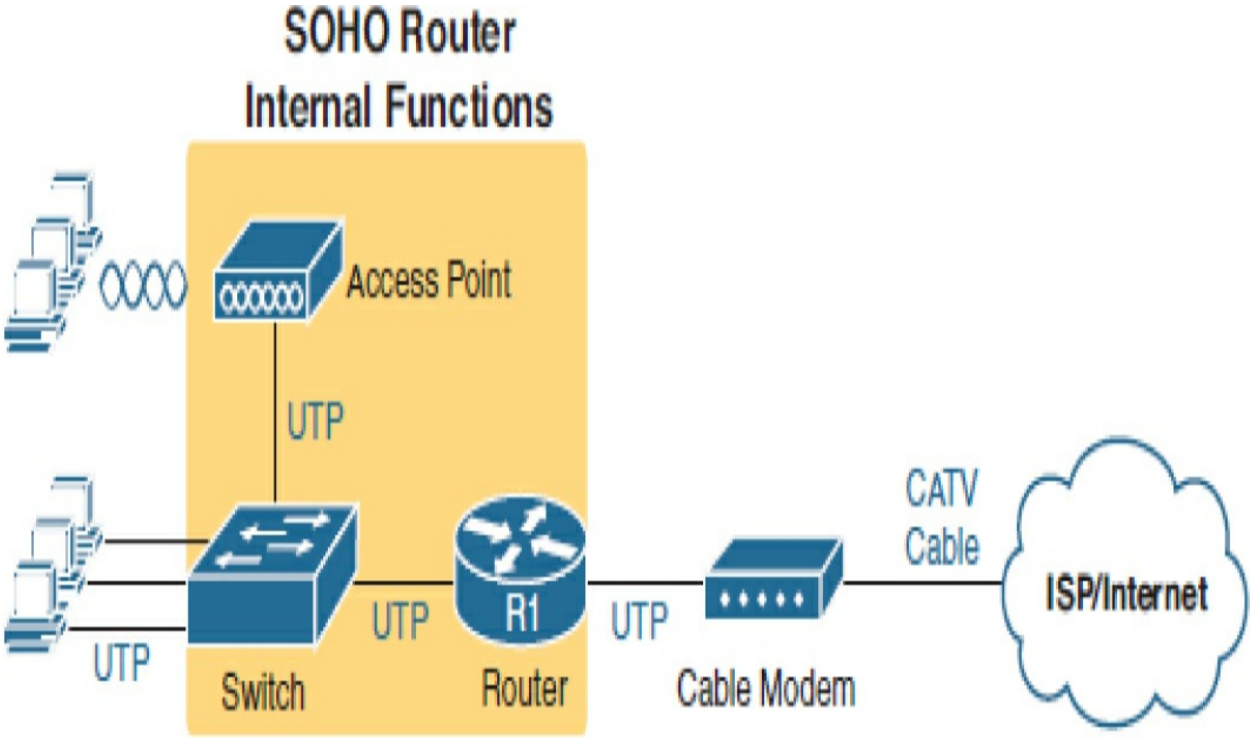

Figure 31-18 *Internal Functions of a SOHO Router*

In reality, the access point and switch are integrated into the router.

Technet24

Physical and Logical Topologies

Network diagrams are usually referred to as *topologies*. A topology graphically displays the interconnection methods used between devices.

Physical topologies refer to the physical layout of devices and how they are cabled. Seven basic physical topologies exist (see Figure 31-19).

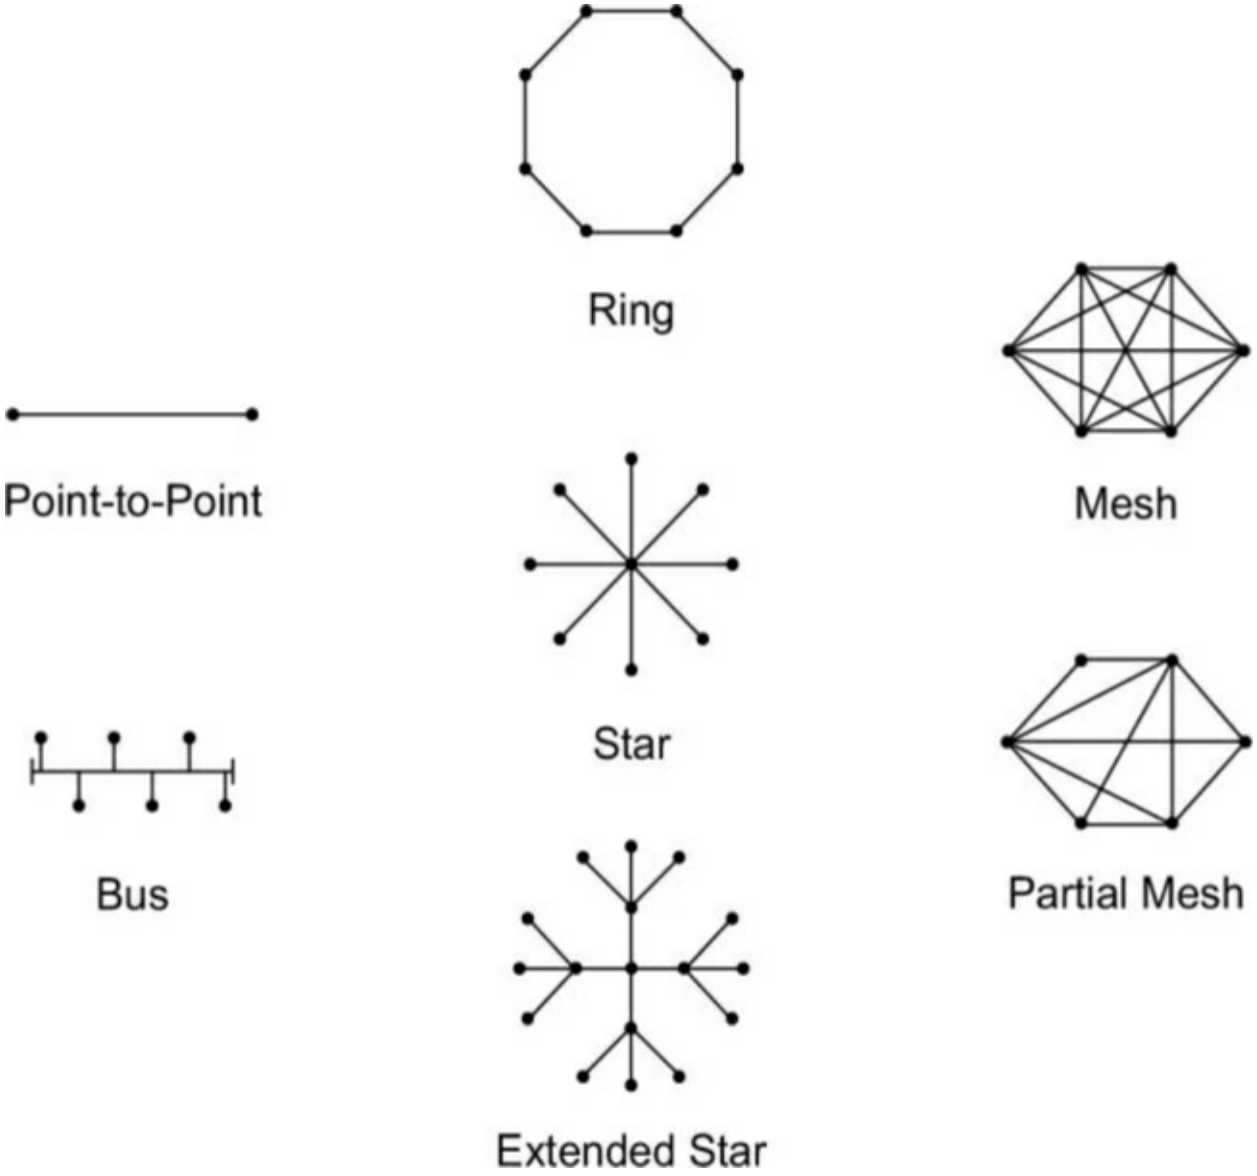

Figure 31-19 *Physical Topologies*

Logical topologies refer to the way that a signal travels from one point on the network to another and are largely determined by the access method—deterministic or nondeterministic. Ethernet is a nondeterministic access method. Logically, Ethernet operates as a bus topology. However, Ethernet networks are almost always physically designed as star or extended star topologies.

Other access methods use a deterministic access method. Token Ring and Fiber Distributed Data Interface (FDDI) both logically operate as ring topologies, passing data from one station to the next. Although these networks can be designed as physical rings, like Ethernet, they are often designed as star or extended star topologies. Logically, however, they operate like ring topologies.

Hierarchical Campus Designs

Hierarchical campus design involves dividing the network into discrete layers. Each layer provides specific functions that define its role within the overall network. By separating the various functions that exist on a network, the network design becomes modular, which facilitates scalability and performance. The hierarchical design model is divided into three layers:

• **Access layer:** Provides local and remote user access

• **Distribution layer:** Controls the flow of data between the access and core layers

Technet24

• **Core layer:** Acts as the high-speed redundant backbone

Figure 31-20 shows an example of the three-tier hierarchical campus network design.

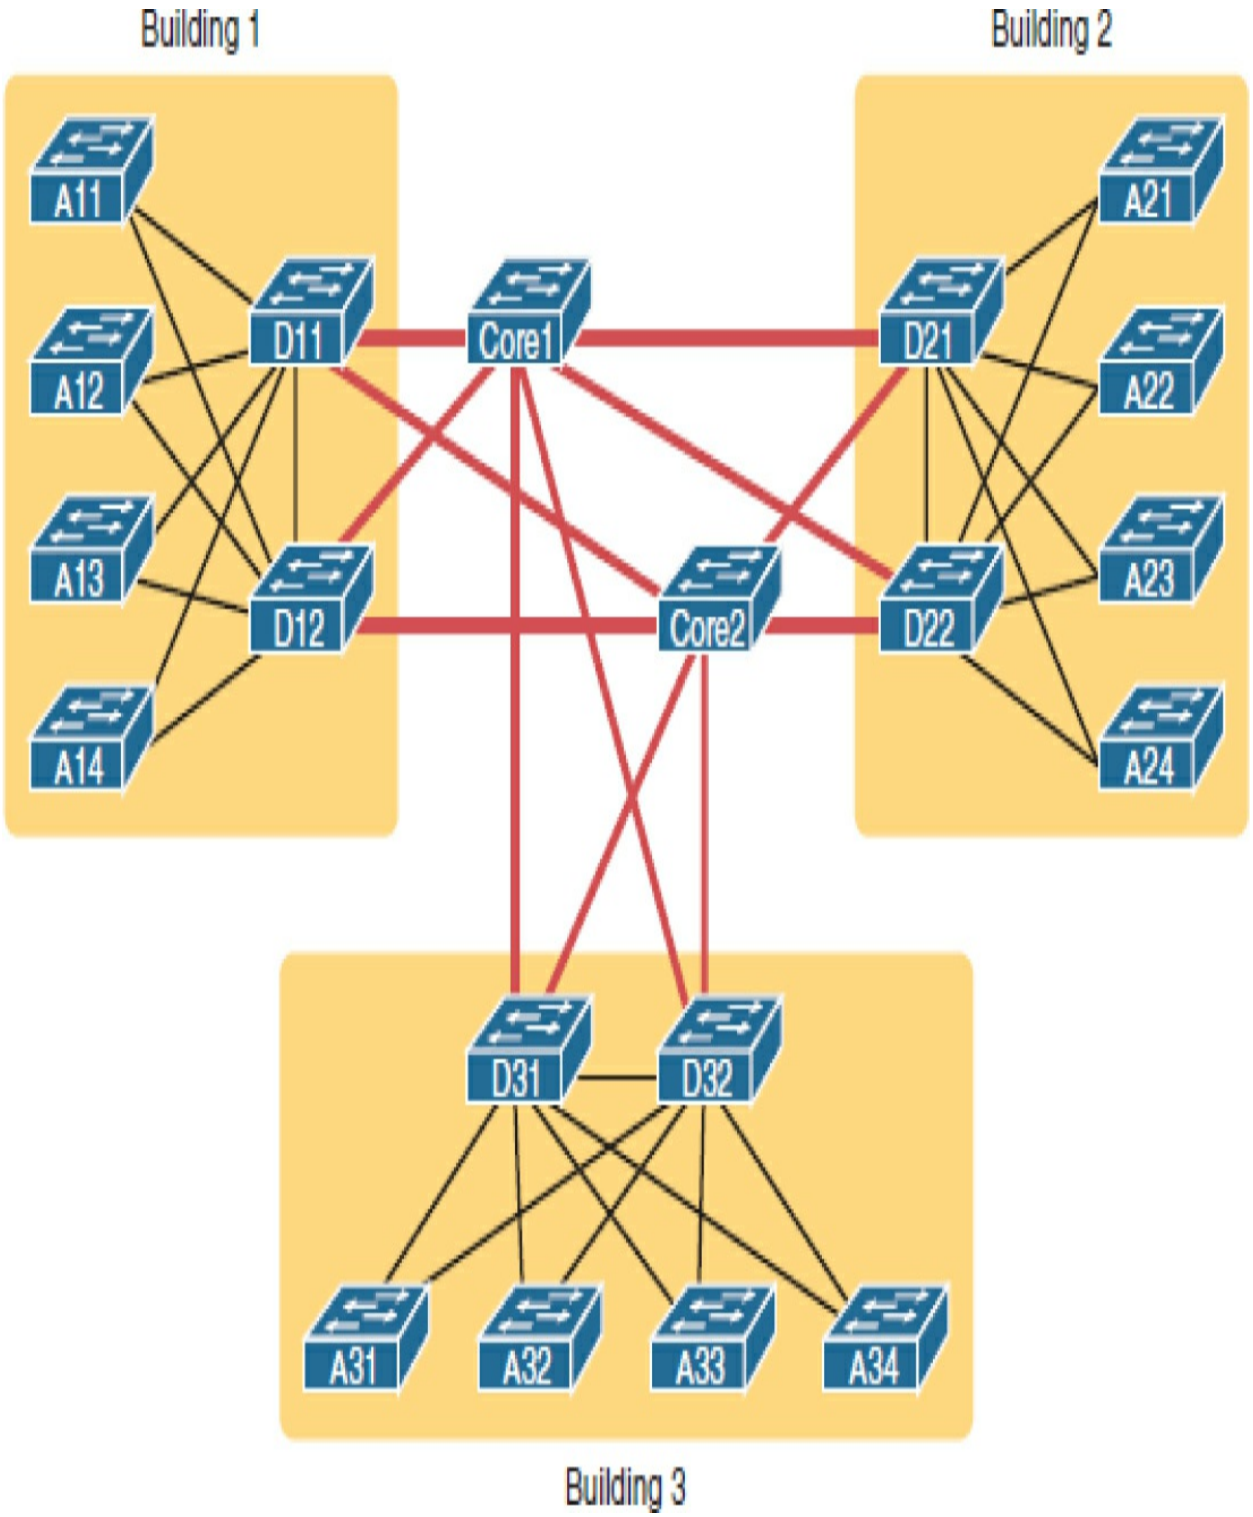

Figure 31-20 *Three-Tier Campus Design*

For smaller networks, the core is often collapsed into the distribution layer

Technet24

for a two-tier design, as in Figure 31-21.

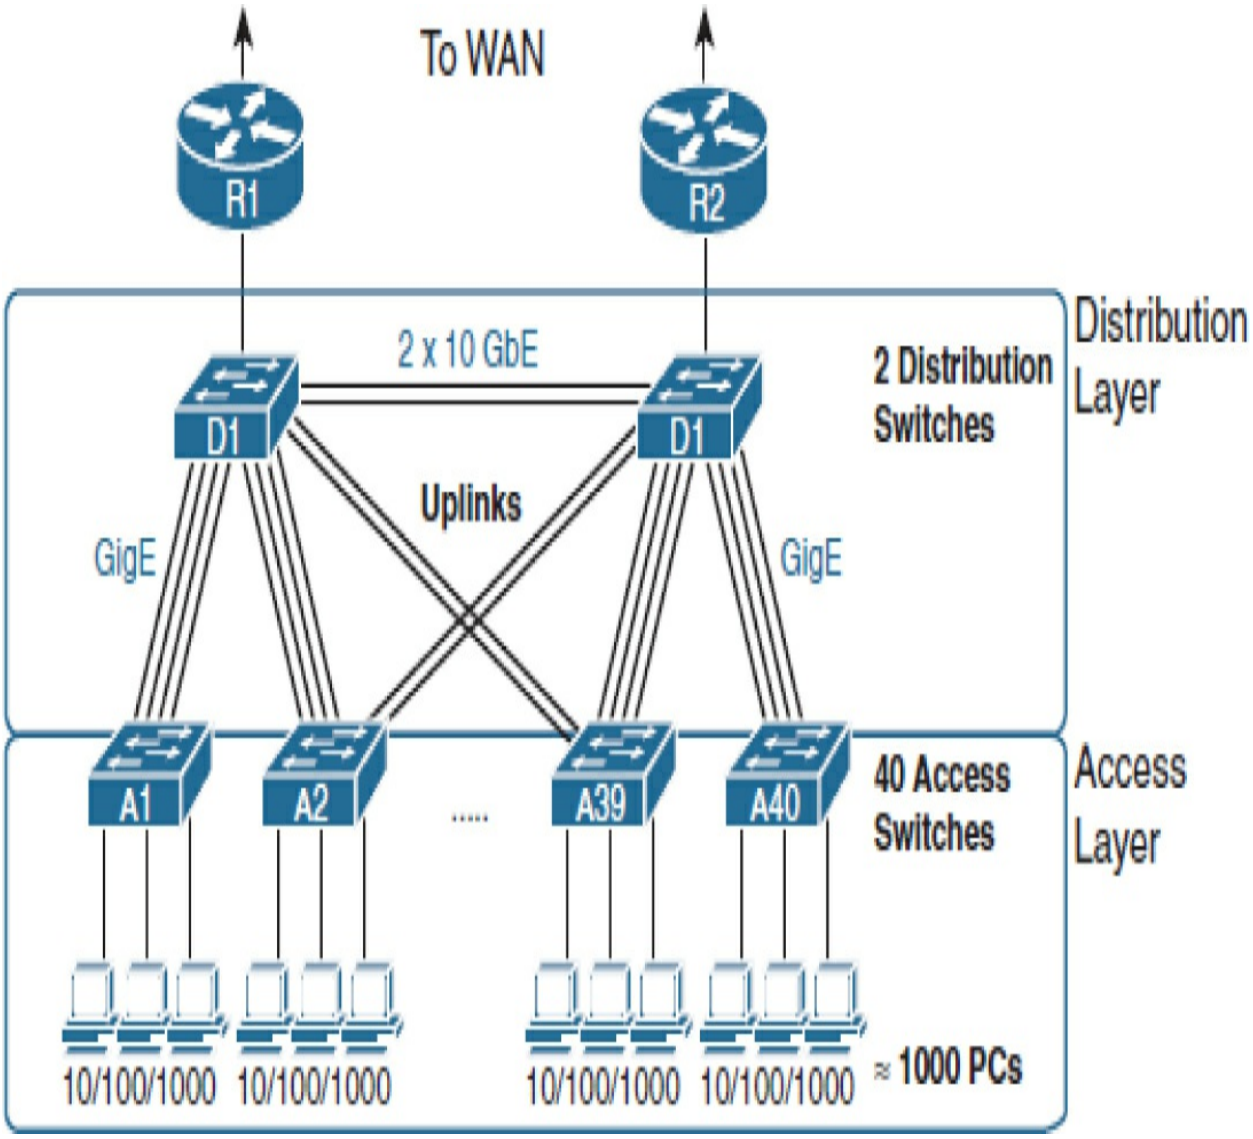

Figure 31-21 *Two-Tier Campus Design*

A two-tier design solves two major design needs:

• Provides a place to connect end-user devices (the access layer, with access switches)

• Connects the switches with a reasonable number of cables and switch ports by connecting all 40 access switches to two distribution switches

For very small networks and home networks, all three tiers can be seen in one device, such as the wireless broadband router shown earlier in Figure 31-16.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Technet24

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	1
	3
	4
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	1–3
	5
	7
	16
	18
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	1
	4
	5
	10
	18–22
Portable Command Guide	6

Day 30. Ethernet Switching

CCNA 200-301 Exam Topics

• 1.1 Explain the role and function of network components

• 1.13 Describe switching concepts

Key Topics

Today we review the concepts behind Ethernet switching, including the history of switching development, how switching actually works, and the variety of switch features. We also review the details of Ethernet operation.

Evolution to Switching

Today’s LANs almost exclusively use switches to interconnect end devices; however, this was not always the case. Initially, devices were connected to a physical bus, a long run of coaxial backbone cabling. With the introduction of 10BASE-T and UTP cabling, the hub gained popularity as a cheaper, easier way to connect devices. But even 10BASE-T with hubs had the following limitations:

• A frame sent from one device can collide with a frame sent by another device attached to that LAN segment. Devices were in the same collision domain sharing the bandwidth.

Technet24

• Broadcasts sent by one device were heard and processed by all other devices on the LAN. Devices were in the same broadcast domain. Much like to hubs, switches forward broadcast frames out all ports except for the incoming port.

Ethernet bridges were soon developed to solve some of the inherent problems in a shared LAN. A bridge basically segmented a LAN into two collision domains, which reduced the number of collisions in a LAN segment. This increased the performance of the network by decreasing unnecessary traffic from another segment.

When switches arrived on the scene, these devices provided the same benefits of bridges, as well as the following:

• A larger number of interfaces to break up the collision domain into more segments

• Hardware-based switching instead of using software to make the decision

In a LAN where all nodes are connected directly to the switch, the throughput of the network increases dramatically. With each computer connected to a separate port on the switch, each is in a separate collision domain and has its own dedicated segment. There are three primary reasons for this increase:

• Dedicated bandwidth to each port

• Collision-free environment

• Full-duplex operation

Switching Logic

Ethernet switches selectively forward individual frames from a receiving port to the port where the destination node is connected. During this instant, the switch creates a full-bandwidth, logical, point-to-point connection between the two nodes.

Switches create this logical connection based on the source and destination Media Access Control (MAC) addresses in the Ethernet header. Specifically, the primary job of a LAN switch is to receive Ethernet frames and then make a decision to either forward the frame or ignore the frame. To accomplish this, the switch performs three actions:

Decides when to forward a frame or when to filter (not forward) a frame, based on the destination MAC address

Learns MAC addresses by examining the source MAC address of each frame the switch receives

Creates a (Layer 2) loop-free environment with other switches by using Spanning Tree Protocol (STP)

To make the decision to forward or filter, the switch uses a dynamically built

Technet24

MAC address table stored in RAM. By comparing the frame's destination MAC address with the fields in the table, the switch decides how to forward and/or filter the frame.

For example, in Figure 30-1, the switch receives a frame from Host A with the destination MAC address OC. The switch looks in its MAC table, finds an entry for the MAC address, and forwards the frame out port 6. The switch also filters the frame by not forwarding it out any other port, including the port on which the frame was received.

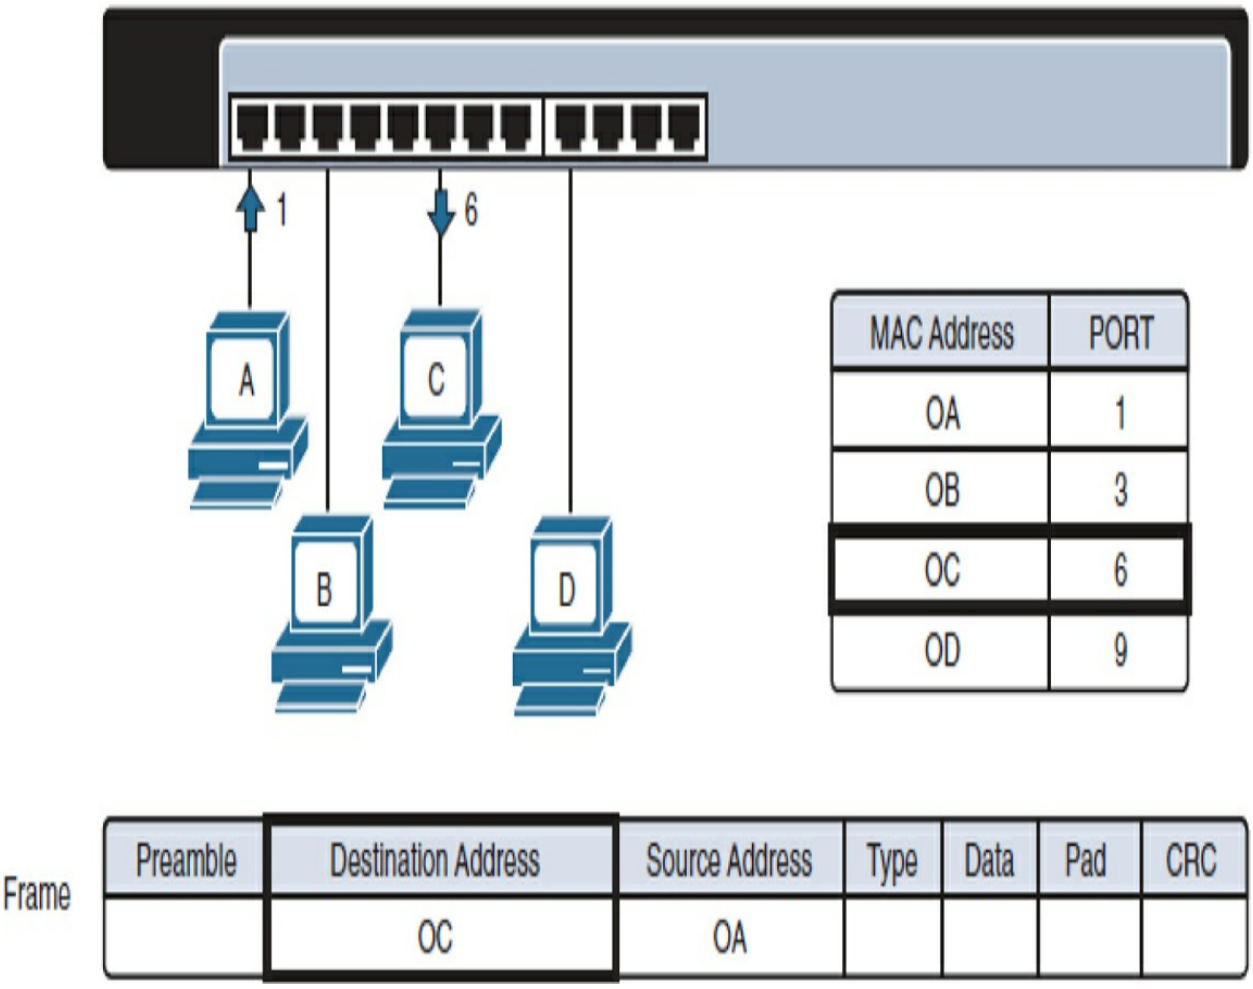

Figure 30-1 *Switch Forwarding Based on MAC Address*

In addition to forwarding and filtering frames, the switch refreshes the timestamp for the source MAC address of the frame. In Figure 30-1, the MAC address for Host A, OA, is already in the MAC table, so the switch refreshes the entry. Entries that are not refreshed eventually are removed (after the default 300 seconds in Cisco IOS).

Continuing the example in Figure 30-1, assume that another device, Host E, is attached to port 10. Host B then sends a frame to the new Host E. The switch does not yet know where Host E is located, so it forwards the frame out all active ports (in a process known as flooding) except for the port on which the frame was received. The new Host E receives the frame. When it replies to Host B, the switch learns Host E's MAC address and port for the first time and stores it in the MAC address table. Subsequent frames destined for Host E then are sent out only port 10.

Finally, LAN switches must have a method for creating a loop-free path for frames to take within the LAN. STP provides loop prevention in Ethernet networks where redundant physical links exist.

Collision and Broadcast Domains

A collision domain is the set of LAN interfaces whose frames could collide with each other. All shared media environments, such as those created by using hubs, are collision domains. When one host is attached to a switch port, the switch creates a dedicated connection, thereby eliminating the potential

Technet24

for a collision. Switches reduce collisions and improve bandwidth use on network segments because they provide full-duplex, dedicated bandwidth to each network segment.

Out of the box, however, a switch cannot provide relief from broadcast traffic. A collection of connected switches forms one large broadcast domain. If a frame with the destination address FFFF.FFFF.FFFF crosses a switch port, that switch must flood the frame out all other active ports. Each attached device must then process the broadcast frame at least up to the network layer. Routers and VLANs are used to segment broadcast domains. Day 26, "VLAN and Trunking Concepts and Configurations," reviews the use of VLANs to segment broadcast domains.

Frame Forwarding

Switches operate in several ways to forward frames. They can differ in forwarding methods, port speeds, memory buffering, and the OSI layers used to make the forwarding decision. The following sections discuss these concepts in greater detail.

Switch Forwarding Methods

Switches use one of the following forwarding methods to switch data between network ports:

• **Store-and-forward switching:** The switch stores received frames in its

buffers, analyzes each frame for information about the destination, and evaluates the data integrity using the cyclic redundancy check (CRC) in the frame trailer. The entire frame is stored, and the CRC is calculated before any of the frame is forwarded. If the CRC passes, the frame is forwarded to the destination.

• **Cut-through switching:** The switch buffers just enough of the frame to read the destination MAC address so that it can determine which port to forward the data to. When the switch determines a match between the destination MAC address and an entry in the MAC address table, the frame is forwarded out the appropriate port(s). This happens as the rest of the initial frame is still being received. The switch does not perform any error checking on the frame.

• **Fragment-free mode:** The switch waits for the collision window (64 bytes) to pass before forwarding the frame. This means that each frame is checked into the data field to make sure that no fragmentation has occurred. Fragment-free mode provides better error checking than cut-through, with practically no increase in latency.

Symmetric and Asymmetric Switching

Symmetric switching provides switched connections between ports with the same bandwidth, such as all 100-Mbps ports or all 1000-Mbps ports. An asymmetric LAN switch provides switched connections between ports of unlike bandwidth, such as a combination of 10-Mbps, 100-Mbps, and 1000-

Technet24

Mbps ports.

Memory Buffering

Switches store frames for a brief time in a memory buffer. Two methods of memory buffering exist:

• **Port-based memory:** Frames are stored in queues that are linked to specific incoming ports.

• **Shared memory:** Frames are deposited into a common memory buffer that all ports on the switch share.

Layer 2 and Layer 3 Switching

A Layer 2 LAN switch performs switching and filtering based only on MAC addresses. A Layer 2 switch is completely transparent to network protocols and user applications. A Layer 3 switch functions similarly to a Layer 2 switch. But instead of using only the Layer 2 MAC address information for forwarding decisions, a Layer 3 switch can also use IP address information. Layer 3 switches are also capable of performing Layer 3 routing functions, reducing the need for dedicated routers on a LAN. Because Layer 3 switches have specialized switching hardware, they can typically route data as quickly as they can switch data.

Ethernet Overview

802.3 is the IEEE standard for Ethernet, and the two terms are commonly used interchangeably. The terms *Ethernet* and *802.3* both refer to a family of standards that together define the physical and data link layers of the definitive LAN technology. Figure 30-2 shows a comparison of Ethernet standards to the OSI model.

Ethernet separates the functions of the data link layer into two distinct sublayers:

- **Logical Link Control (LLC) sublayer:** Defined in the 802.2 standard
- **Media Access Control (MAC) sublayer:** Defined in the 802.3 standard

Technet24

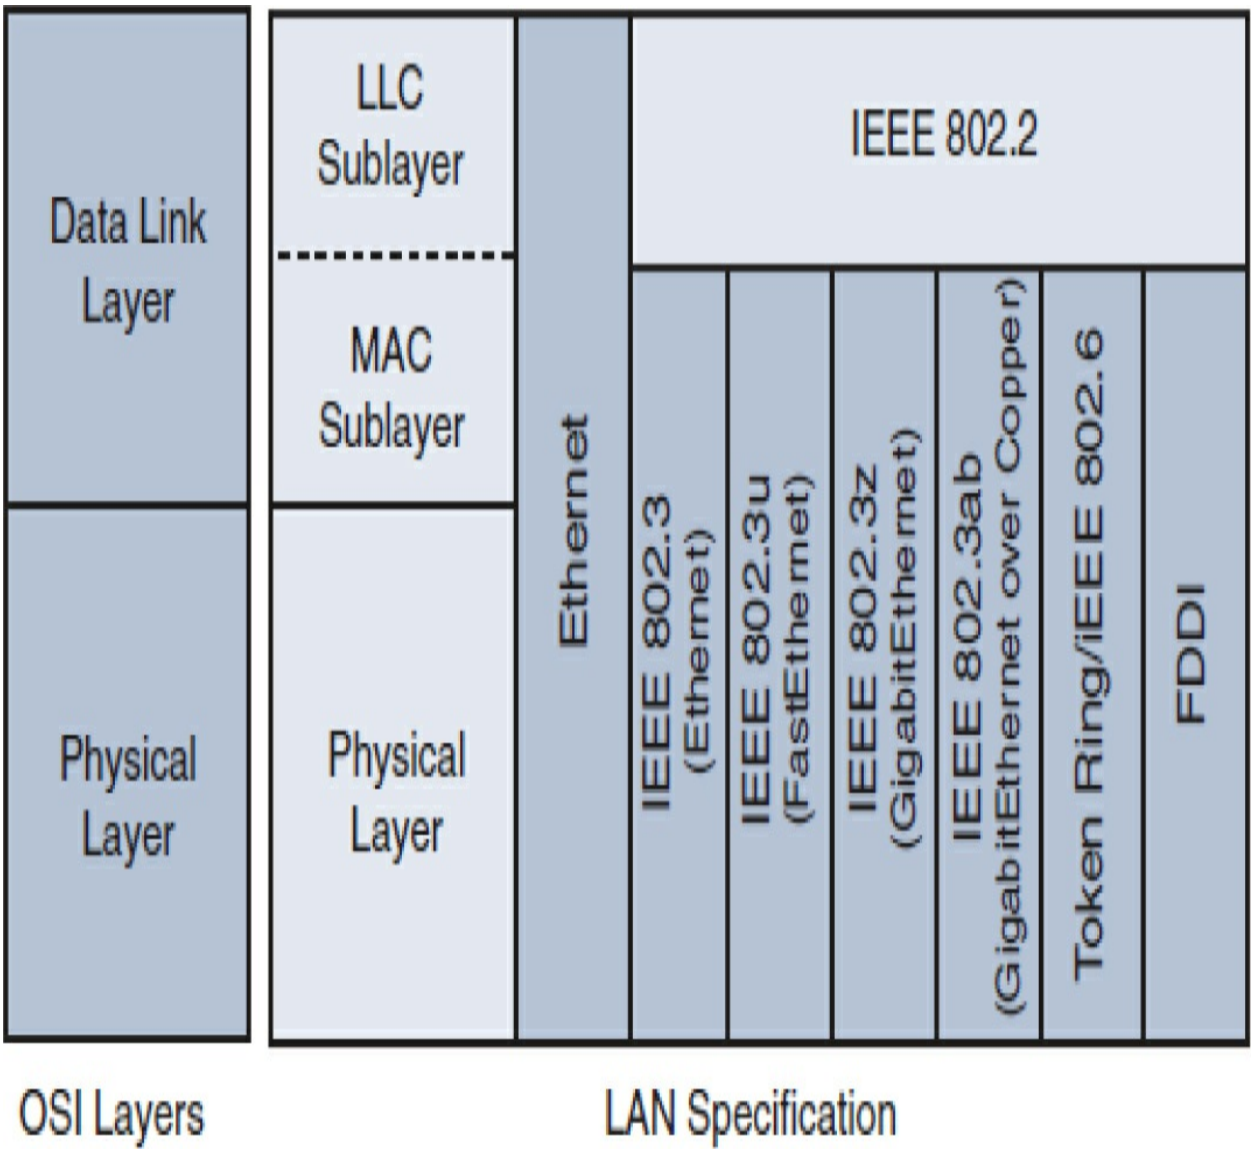

Figure 30-2 *Ethernet Standards and the OSI Model*

The LLC sublayer handles communication between the network layer and the MAC sublayer. In general, LLC provides a way to identify the protocol that is passed from the data link layer to the network layer. In this way, the fields of the MAC sublayer are not populated with protocol type information, as was the case in earlier Ethernet implementations.

The MAC sublayer has two primary responsibilities:

• **Data encapsulation:** Included here is frame assembly before transmission, frame parsing upon reception of a frame, data link layer MAC addressing, and error detection.

• **Media Access Control:** Because Ethernet is a shared medium and all devices can transmit at any time, media access is controlled by a method called Carrier Sense Multiple Access/Collision Detect (CSMA/CD) when operating in half-duplex mode.

At the physical layer, Ethernet specifies and implements encoding and decoding schemes that enable frame bits to be carried as signals across both unshielded twisted pair (UTP) copper cables and optical fiber cables. In early implementations, Ethernet used coaxial cabling.

Legacy Ethernet Technologies

Ethernet is best understood by first considering the two early Ethernet specifications, 10BASE-5 and 10BASE-2. With these two specifications, the network engineer installs a series of coaxial cables connecting each device on the Ethernet network, as in Figure 30-3.

The series of cables creates an electrical circuit, called a *bus*, that is shared among all devices on the Ethernet. When a computer wants to send some bits to another computer on the bus, it sends an electrical signal, and the

Technet24

electricity propagates to all devices on the Ethernet.

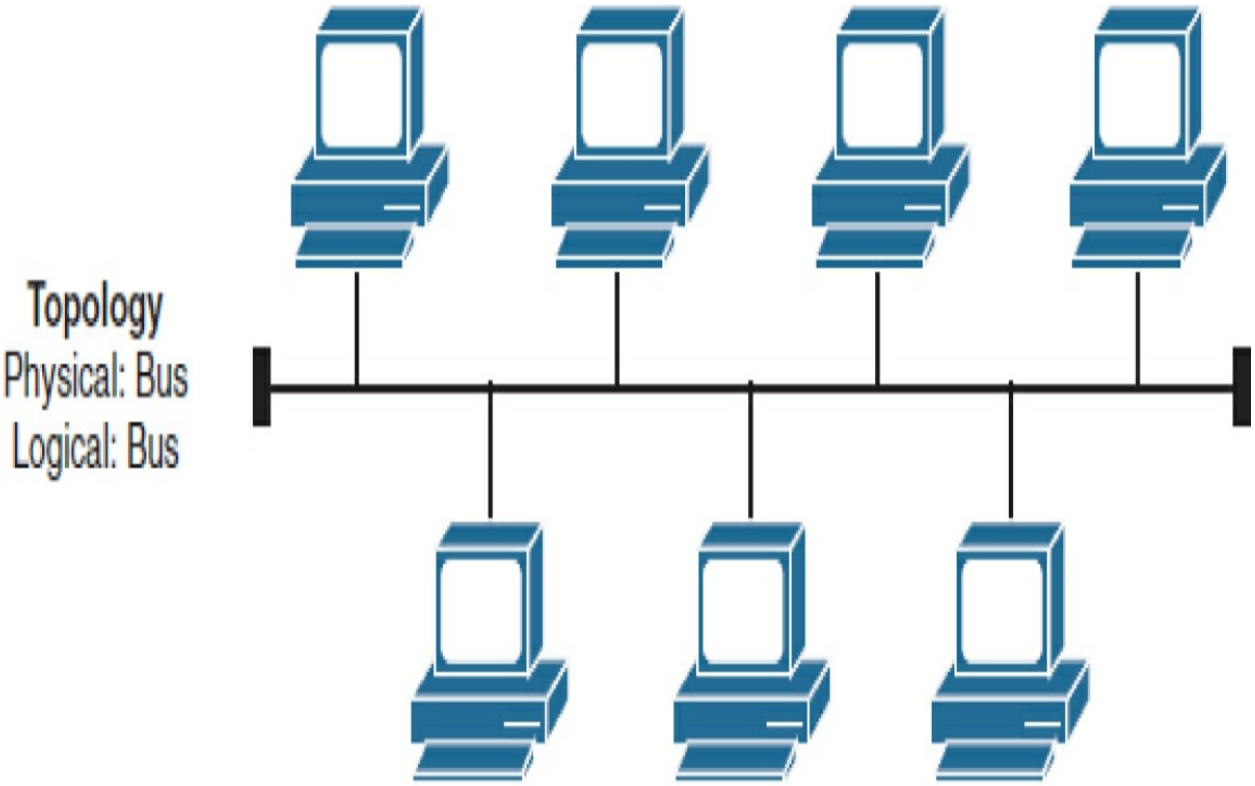

Figure 30-3 *Ethernet Physical and Logical Bus Topology*

With the change of media to UTP and the introduction of the first hubs, Ethernet physical topologies migrated to a star, as shown in Figure 30-4.

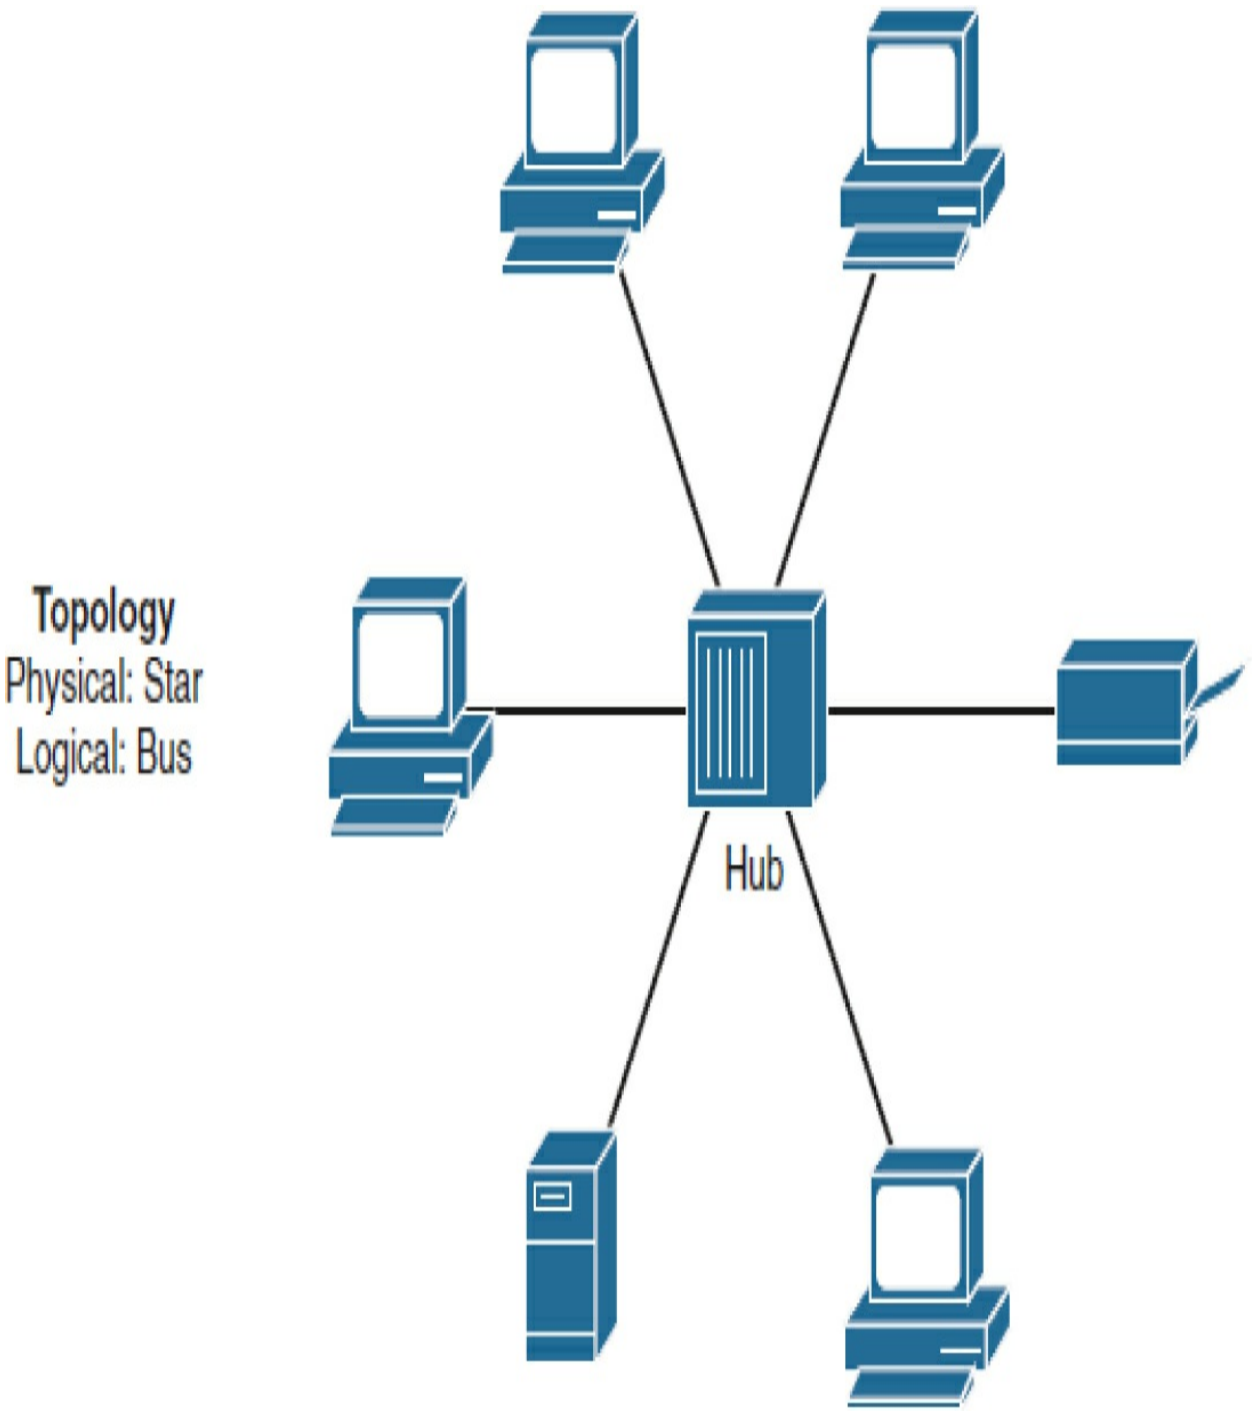

Figure 30-4 *Ethernet Physical Star and Logical Bus Topology*

Regardless of the change in the physical topology from a bus to a star, hubs logically operate similarly to a traditional bus topology and require the use of

Technet24

CSMA/CD.

CSMA/CD

Because Ethernet is a shared medium in which every device has the right to send at any time, it also defines a specification to ensure that only one device sends traffic at a time. The CSMA/CD algorithm defines how the Ethernet logical bus is accessed.

CSMA/CD logic helps prevent collisions and also defines how to act when a collision does occur. The CSMA/CD algorithm works like this:

A device with a frame to send listens until the Ethernet is not busy.

When the Ethernet is not busy, the sender(s) begin(s) sending the frame.

The sender(s) listen(s) to make sure that no collision occurs.

If a collision occurs, the devices that were sending a frame each send a jamming signal to ensure that all stations recognize the collision.

When the jamming is complete, each sender randomizes a timer and waits until the timer expires before trying to resend the collided frame.

When each random timer expires, the process starts again from the beginning.

When CSMA/CD is in effect, a device’s network interface card (NIC)

operates in half-duplex mode, either sending or receiving frames. CSMA/CD is disabled when a NIC autodetects that it can operate in—or is manually configured to operate in—full-duplex mode. In full-duplex mode, a NIC can send and receive simultaneously.

Legacy Ethernet Summary

LAN hubs occasionally appear, but switches generally are used instead of hubs. Keep in mind the following key points about the history of Ethernet:

• The original Ethernet LANs created an electrical bus to which all devices connected.

• 10BASE-2 and 10BASE-5 repeaters extended the length of LANs by cleaning up the electrical signal and repeating it (a Layer 1 function) but without interpreting the meaning of the electrical signal.

• Hubs are repeaters that provide a centralized connection point for UTP cabling—but they still create a single electrical bus that the various devices share, just as with 10BASE-5 and 10BASE-2.

• Because collisions can occur in any of these cases, Ethernet defines the CSMA/CD algorithm, which tells devices how to both avoid collisions and take action when collisions do occur.

Current Ethernet Technologies

Technet24

Refer to Table 30-1 and notice the different 802.3 standards. Each new physical layer standard from the IEEE requires many differences at the physical layer. However, each of these physical layer standards uses the same 802.3 header, and each uses the upper LLC sublayer as well. Table 30-1 lists today's most commonly used IEEE Ethernet physical layer standards.

Table 30-1 *Today's Most Common Types of Ethernet*

Common Name	Speed	Alternative Name	Name of IEEE Standard	Cable Type, Maximum Length
Ethernet	10 Mbps	10BASE-T	802.3	Copper, 100 m
Fast Ethernet	100 Mbps	100BASE-TX	802.3u	Copper, 100 m
Gigabit Ethernet	1000 Mbps	1000BASE-LX	802.3z	Fiber, 550 m
Gigabit Ethernet	1000 Mbps	1000BASE-T	802.3ab	Copper, 100 m
10GigE (Gigabit Ethernet)	10 Gbps	10GBASE-T	802.3an	Copper, 100 m
10GigE (Gigabit Ethernet)	10 Gbps	10GBASE-S	802.3ae	Fiber, 400 m

UTP Cabling

Technet24

The three most common Ethernet standards used today—10BASE-T (Ethernet), 100BASE-TX (Fast Ethernet, or FE), and 1000BASE-T (Gigabit Ethernet, or GE)—use UTP cabling. Some key differences exist, particularly with the number of wire pairs needed in each case and the type (category) of cabling.

The UTP cabling in popular Ethernet standards includes either two or four pairs of wires. The cable ends typically use an RJ-45 connector. The RJ-45 connector has eight specific physical locations into which the eight wires in the cable can be inserted; these are called pin positions or, simply, pins.

The Telecommunications Industry Association (TIA) and the Electronics Industry Alliance (EIA) define standards for UTP cabling, with color coding for wires and standard pinouts on the cables. Figure 30-5 shows two TIA/EIA pinout standards, with the color coding and pair numbers listed.

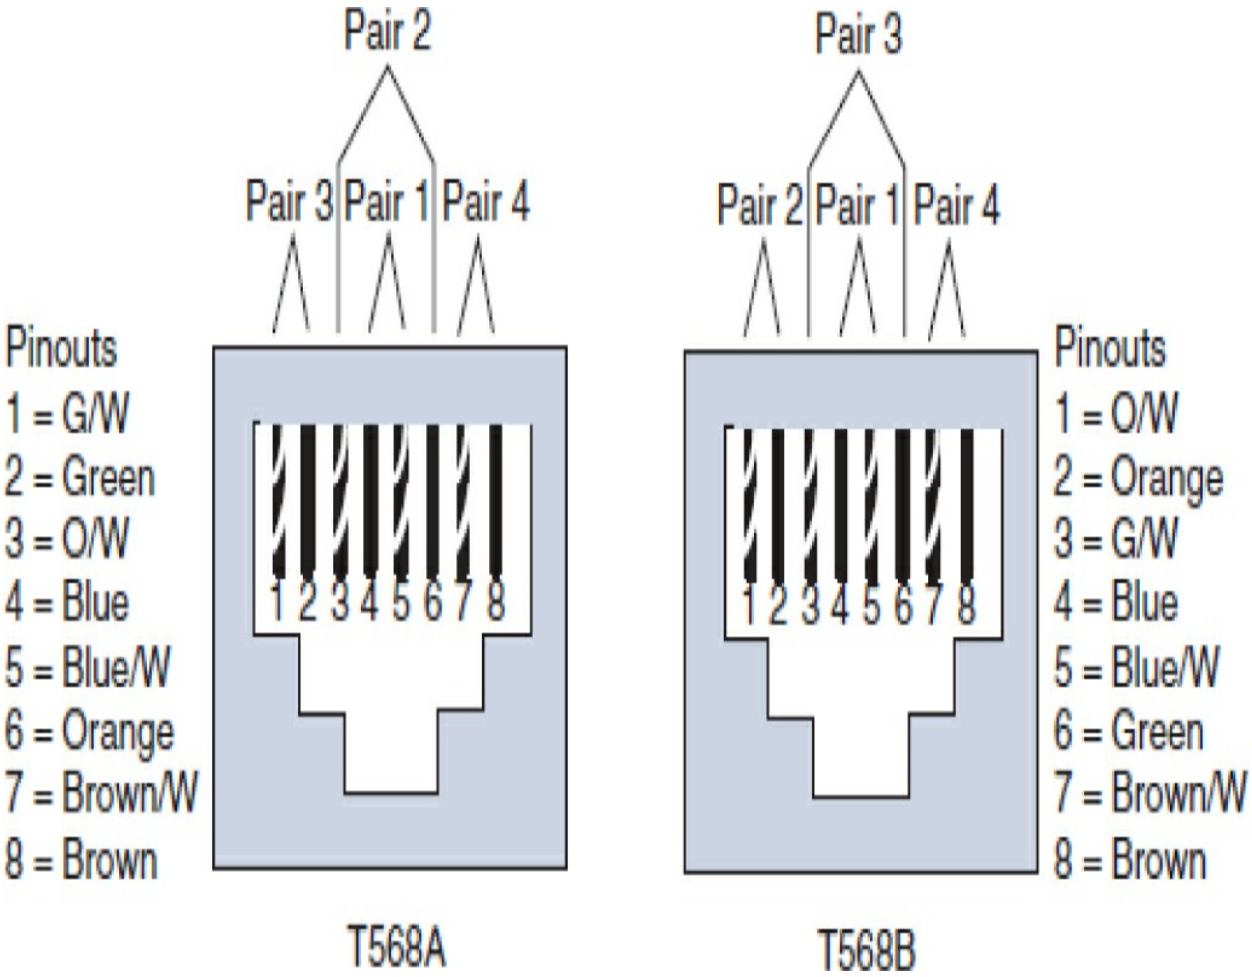

Figure 30-5 *TIA/EIA Standard Ethernet Cabling Pinouts*

For the exam, you should be well prepared to choose which type of cable (straight-through or crossover) is needed in each part of the network. In short, devices on opposite ends of a cable that use the same pair of pins to transmit need a crossover cable. Devices that use an opposite pair of pins to transmit need a straight-through cable. Table 30-2 lists typical devices and the pin pairs they use, assuming that they use 10BASE-T and 100BASE-TX.

Table 30-2 *10BASE-T and 100BASE-TX Pin Pairs Used*

Technet24

Devices That Transmit on 1,2 and Receive on 3,6	Devices That Transmit on 3,6 and Receive on 1,2
PC NICs	Hubs
Routers	Switches
Wireless access points (Ethernet interfaces)	—
Networked printers (printers that connect directly to the LAN)	—

1000BASE-T requires four wire pairs because Gigabit Ethernet transmits and receives on each of the four wire pairs simultaneously.

However, Gigabit Ethernet does have a concept of straight-through and crossover cables, with a minor difference in the crossover cables. The pinouts for a straight-through cable are the same—pin 1 to pin 1, pin 2 to pin 2, and so on.

A crossover cable has the 568A standard on one end and the 568B standard on the other end. This crosses the pairs at pins 1,2 and 3,6.

Benefits of Using Switches

A collision domain is a set of devices whose frames may collide. All devices

on a 10BASE-2, 10BASE-5, or other network using a hub risk collisions between the frames that they send. Thus, devices on one of these types of Ethernet networks are in the same collision domain and use CSMA/CD to detect and resolve collisions.

LAN switches significantly reduce, or even eliminate, the number of collisions on a LAN. Unlike a hub, a switch does not create a single shared bus. Instead, a switch does the following:

• It interprets the bits in the received frame so that it can typically send the frame out the one required port instead of out all other ports.

• If a switch needs to forward multiple frames out the same port, the switch buffers the frames in memory, sending one at a time and thereby avoiding collisions.

In addition, switches with only one device cabled to each port of the switch allow the use of full-duplex operation. Full-duplex operation means that the NIC can send and receive concurrently, effectively doubling the bandwidth of a 100-Mbps link to 200 Mbps—100 Mbps for sending and 100 Mbps for receiving.

These seemingly simple switch features provide significant performance improvements compared with using hubs. In particular, consider these points:

• If only one device is cabled to each port of a switch, no collisions can occur.

Technet24

• Devices connected to one switch port do not share their bandwidth with devices connected to another switch port. Each has its own separate bandwidth, meaning that a switch with 100-Mbps ports has 100 Mbps of bandwidth per port.

Ethernet Addressing

The IEEE defines the format and assignment of LAN addresses. To ensure a unique MAC address, the first half of the address identifies the manufacturer of the card. This code is called the organizationally unique identifier (OUI). Each manufacturer assigns a MAC address with its own OUI as the first half of the address. The second half of the address is assigned by the manufacturer and is never used on another card or network interface with the same OUI. Figure 30-6 shows the structure of a unicast Ethernet address.

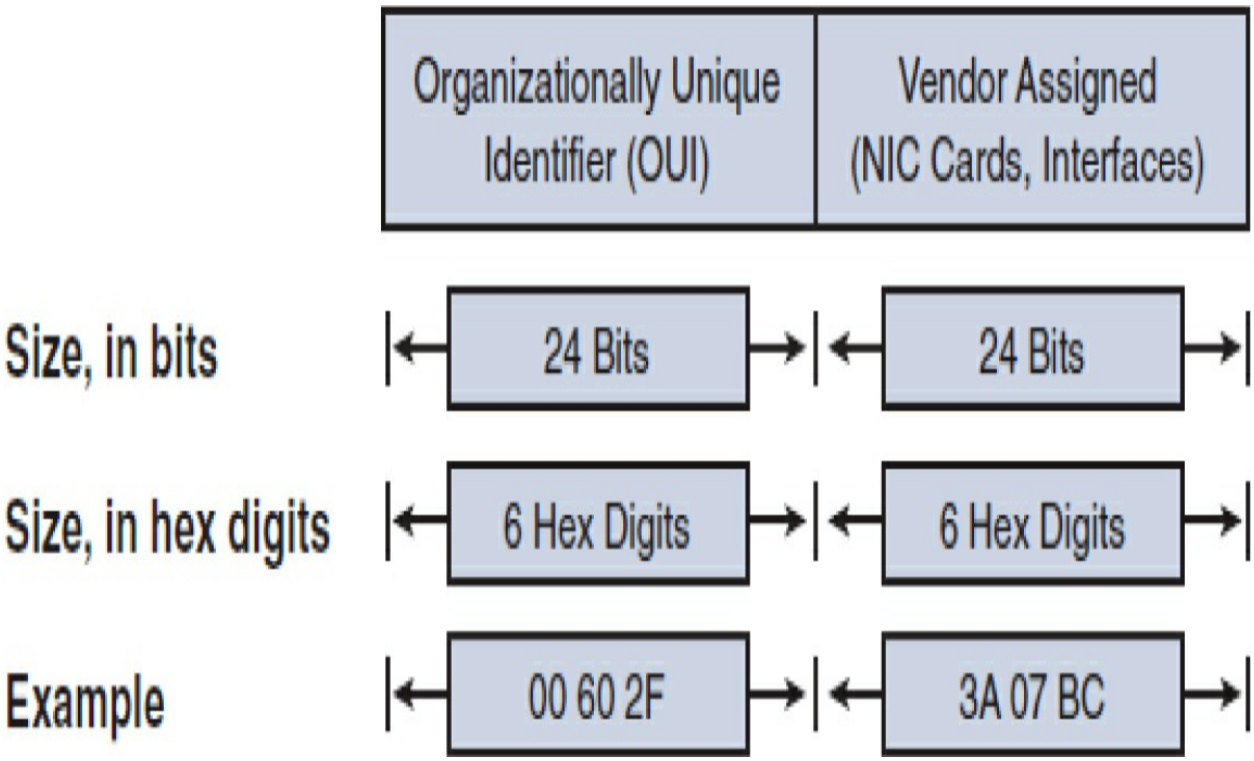

Figure 30-6 *Structure of a Unicast Ethernet Address*

Ethernet also has group addresses, which identify more than one NIC or network interface. The IEEE defines two general categories of group addresses for Ethernet:

• **Broadcast addresses:** A broadcast address implies that all devices on the LAN should process the frame and has the value FFFF.FFFF.FFFF.

• **Multicast addresses:** A multicast address allows a subset of devices on a LAN to communicate. When IP multicasts over an Ethernet network, the multicast MAC addresses that IP uses follow this format: 0100.5exx.xxxx. The xx.xxxx portion is divided between IPv4 multicast (00:0000–7F.FFFF) and MPLS multicast (80:0000–8F:FFFF). Multiprotocol Label Switching

Technet24

(MPLS) is a CCNP topic.

Ethernet Framing

The physical layer helps you get a string of bits from one device to another. The framing of the bits allows the receiving device to interpret the bits. The term *framing* refers to the definition of the fields assumed to be in the data that is received. Framing defines the meaning of the bits transmitted and received over a network.

The framing used for Ethernet has changed a couple times over the years. Figure 30-7 shows each iteration of Ethernet, with the current version shown at the bottom.

DIX

Preamble	Destination	Source	Type	Data and Pad	FCS
8	6	6	2	46 – 1500	4

IEEE 802.3 (Original)

Preamble	SFD	Destination	Source	Length	Data and Pad	FCS
7	1	6	6	2	46 – 1500	4

IEEE 802.3 (Revised 1997)

	Preamble	SFD	Destination	Source	Length/	Data and Pad	FCS
Bytes	7	1	6	6	Type 2	46 – 1500	4

Figure 30-7 *Ethernet Frame Formats*

Table 30-3 further explains the fields in the last version shown in Figure 30-7.

Table 30-3 *IEEE 802.3 Ethernet Field Descriptions*

Technet24

Field	Field Length, in Bytes	Description
Preamble	7	Synchronization
Start Frame Delimiter (SFD)	1	Signifies that the next byte begins the Destination MAC field
Destination MAC Address	6	Identifies the intended recipient of this frame
Source MAC Address	6	Identifies the sender of this frame
Length	2	Defines the length of the data field of the frame (either length or type is present, but not both)
Type	2	Defines the type of protocol listed inside the frame (either length or type is present, but not both)
Data and Pad	46–1500	Holds data from a higher layer, typically a Layer 3 PDU (generic) and often an IP packet
Frame Check Sequence (FCS)	4	Provides a method for the receiving NIC to determine whether the frame experienced transmission errors

The Role of the Physical Layer

We have already discussed the most popular cabling used in LANs: UTP. To fully understand the operation of the network, you should know some additional basic concepts of the physical layer.

The OSI physical layer accepts a complete frame from the data link layer and encodes it as a series of signals that are transmitted onto the local media.

The delivery of frames across the local media requires the following physical layer elements:

• The physical media and associated connectors

• A representation of bits on the media

• Encoding of data and control information

• Transmitter and receiver circuitry on the network devices

Data is represented on three basic forms of network media:

• Copper cable

• Fiber

• Wireless (IEEE 802.11)

Technet24

Bits are represented on the medium by changing one or more of the following characteristics of a signal:

• Amplitude

• Frequency

• Phase

The nature of the actual signals representing the bits on the media depends on the signaling method in use. Some methods use one attribute of a signal to represent a single 0 and use another attribute of a signal to represent a single 1. The actual signaling method and its detailed operation are not important to your CCNA exam preparation.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	4
	6
	7
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	5
	8

Technet24

Day 29. Switch Configuration Basics

CCNA 200-301 Exam Topics

• 1.4 Identify interface and cable issues (collisions, errors, mismatched duplex, and/or speed)

• 1.6 Configure and verify IPv4 addressing and subnetting

• 2.8 Describe network device management access (Telnet, SSH, HTTP, HTTPS, console, TACACS+/RADIUS, and cloud managed)

Key Topics

Today we review Cisco IOS basics and the commands necessary to perform a basic initial configuration of a switch. Although not explicitly called out in the exam topics, you can expect to see questions that assume you have this skill. We review verification techniques such as the **ping**, **traceroute**, and **show** commands. And we review interface and cable issues.

Accessing and Navigating the Cisco IOS

By now, you are very familiar with connecting to Cisco devices and configuring them using the command-line interface (CLI). Here we quickly review methods for accessing and navigating the CLI.

Connecting to Cisco Devices

Cisco IOS Software for Catalyst switches controls the switch's performance and behavior and provides a CLI for user interaction via terminal emulation programs. The switch CLI can be accessed through three popular methods, as shown in Figure 29-1:

• **Console:** A physical port built specifically to allow access to the CLI

• **Telnet:** An unsecured method that involves using the IP network in which the switch resides to reach the switch

• **Secure Shell (SSH):** A secured method that involves using the IP network in which the switch resides to reach the switch

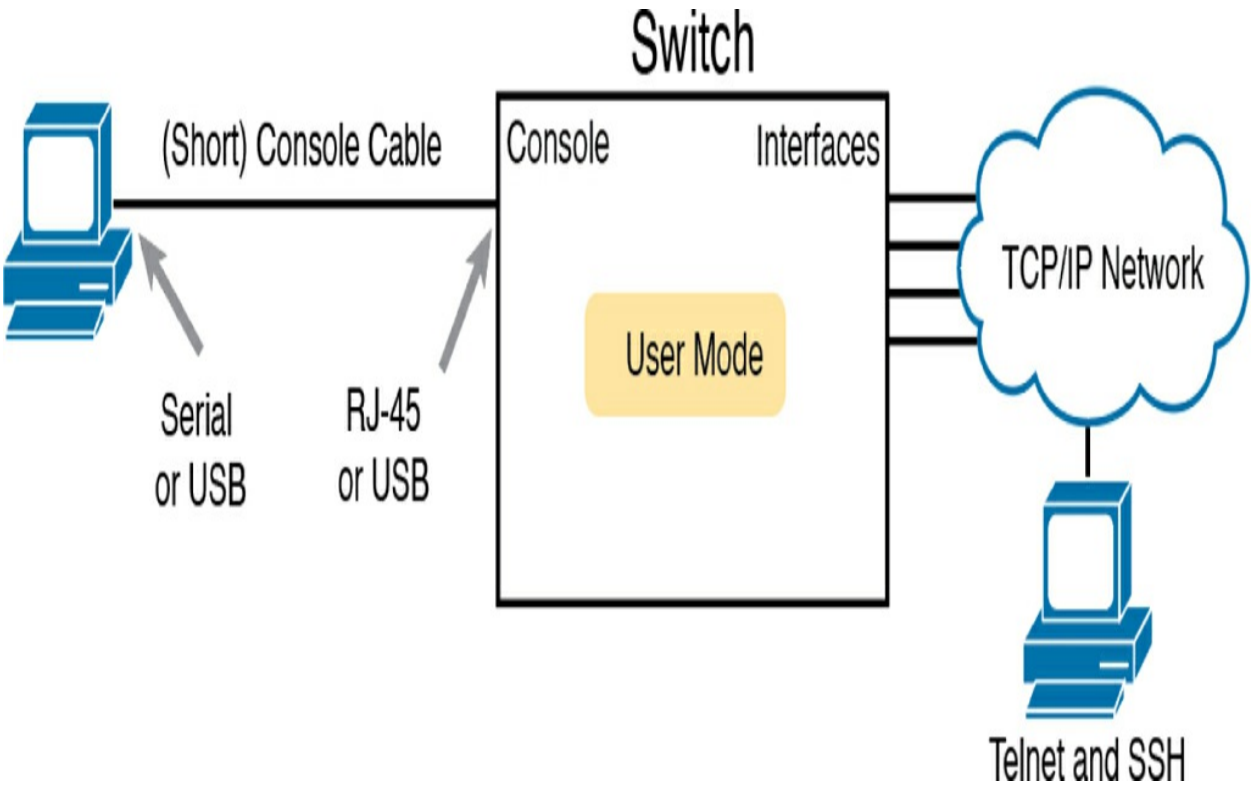

Figure 29-1 *Cisco Device Configuration Options*

Technet24

WebUI

Cisco IOS XE routers and Catalyst switches feature a web-based management interface called WebUI (or web GUI) that enables administrators to configure and monitor devices visually through a web browser. To set up WebUI access, you need to configure a username and use specific IP HTTP commands. It is recommended to secure WebUI with an access control list. For example, the configuration in Example 29-1 allows secure WebUI access to R1 through the 192.168.1.1/24 interface by any user on the 192.168.1.0/24 network.

Example 29-1 *Configuring a Cisco Device for WebUI Access*

```
R1(config)# interface GigabitEthernet0/0/1
R1(config-if)# ip address 192.168.1.1 255.255.255.252
R1(config-if)# no shutdown
R1(config)# exit
R1(config)# username admin privilege 15 secret cisco12
R1(config)# no ip http server
R1(config)# ip http secure-server
R1(config)# ip http authentication local
R1(config)# access-list 10 permit 192.168.1.0 0.0.0.25
R1(config)# ip http access-class ipv4 10
```

By default, IOS enables both HTTP and HTTPS servers, but Cisco recommends disabling the HTTP server for better security. The HTTP server

allows login via enable password, local username/password, or AAA settings, with a preference for local usernames in recent IOS versions. To access all WebUI features, users need a privilege level of 15, which grants full access to CLI configuration, software installation, and other advanced functions.

The **ipv4** argument in the **ip http access-class** command is new in IOS 17. Although it is not required yet, if you enter only **ip http access-class 10**, you will receive the following message:

```
R1(config)# ip http access-class 10
This CLI will be deprecated soon, Please use new CLI '
  ipv4 <1-99> | <access-list-name>'
```

WebUI Connection

After configuring WebUI access, you can enter the IP address of the router from any web browser that is connected to the allowed IP address range. After successfully logging in, you are presented with the Dashboard, which provides a variety of at-a-glance statistics and system information, as shown in Figure 29-2.

Technet24

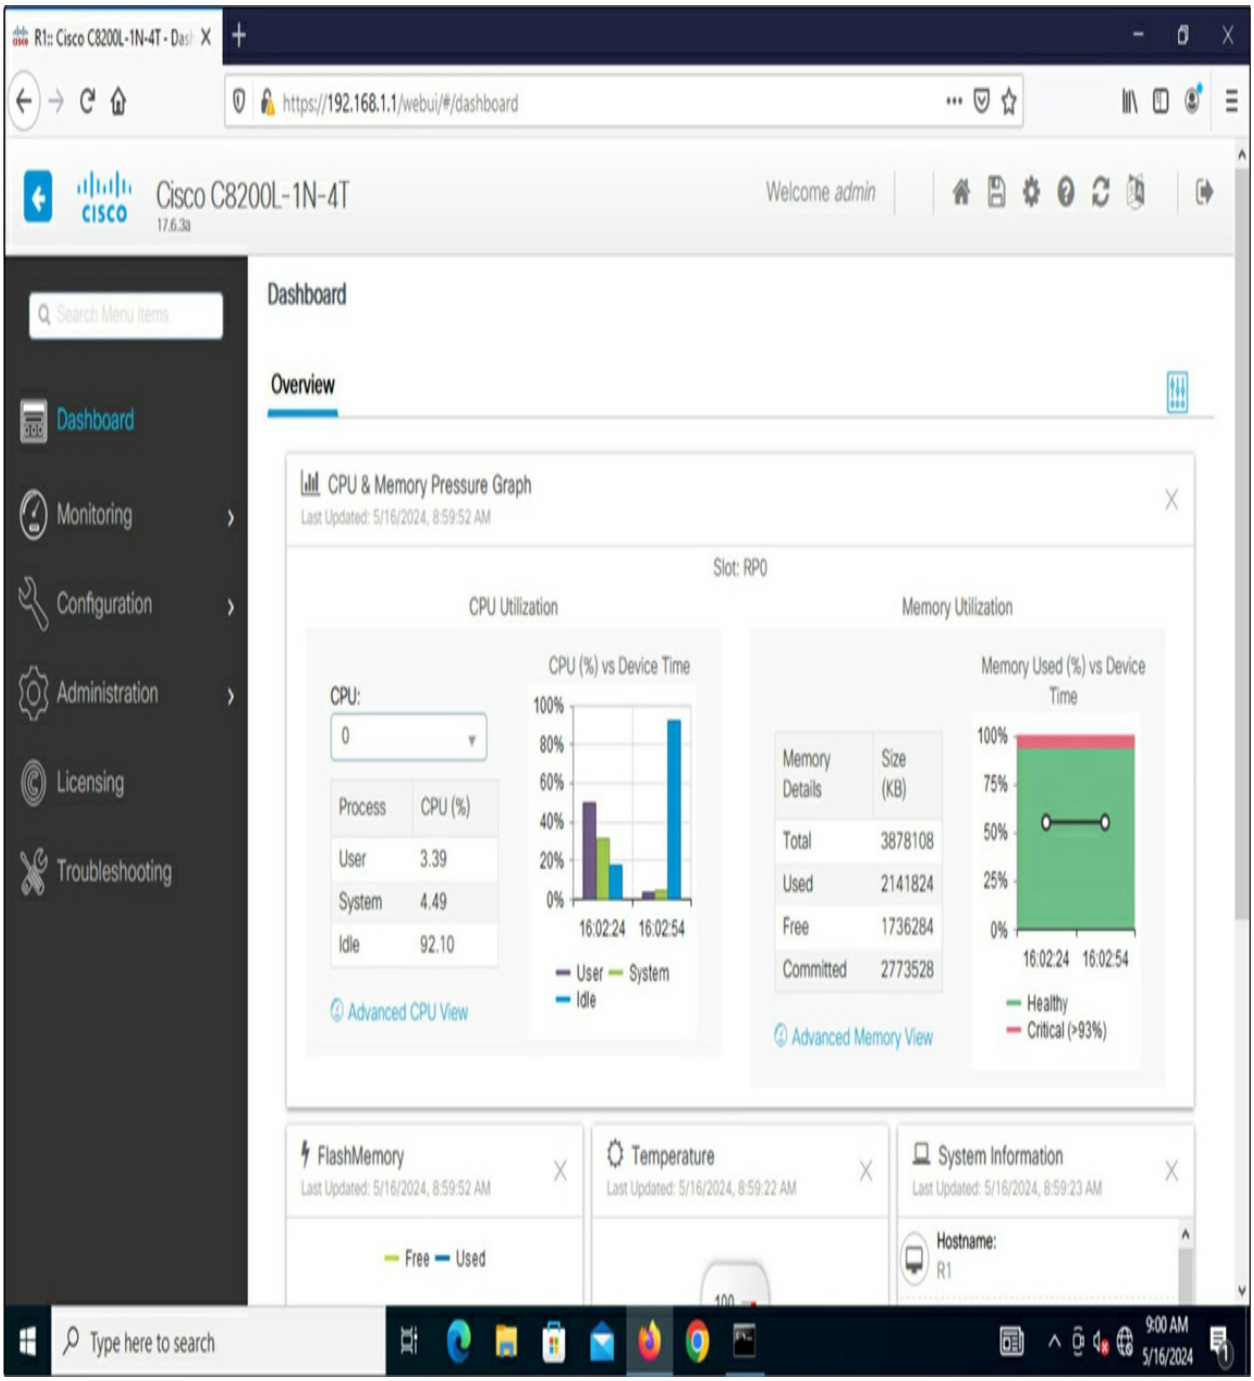

Figure 29-2 *WebUI Dashboard on a Cisco 8200 Router*

CLI EXEC Sessions

Cisco IOS separates the EXEC session into two basic access levels:

• **User EXEC mode:** Access to only a limited number of basic monitoring and troubleshooting commands, such as **show** and **ping**

• **Privileged EXEC mode:** Full access to all device commands, including configuration and management

Using the Help Facility

Cisco IOS has extensive command-line input help facilities, including context-sensitive help. Two types of help are available:

• **Word help:** Enter a character sequence of an incomplete command immediately followed by a question mark (for example, **sh?**) to get a list of available commands that start with the character sequence.

• **Command syntax help:** Enter the **?** command to get command syntax help to see all the available arguments to complete a command (for example, **show ?**). Cisco IOS then displays a list of available arguments.

As part of the help facility, Cisco IOS displays console error messages when incorrect command syntax is entered. Table 29-1 shows sample error messages, what they mean, and how to get help.

Table 29-1 *Console Error Messages*

Technet24

Example Error Message	Meaning	How to Get Help
`switch# cl % Ambiguous command: "cl"`	You did not enter enough characters for your device to recognize the command.	Reenter the command, followed by a question mark (?), without a space between the command and the question mark. The possible keywords that you can enter with the command appear.
`switch# clock % Incomplete command.`	You did not enter all the keywords or values required by this command.	Reenter the command, followed by a question mark (?), with a space between the command and the question mark.
`switch# clock ste ^ % Invalid input detected at '^' marker.`	You entered the command incorrectly. The caret (^) marks the point of the error.	Enter a question mark (?) to display all the available commands or parameters.

CLI Navigation and Editing Shortcuts

Table 29-2 summarizes the shortcuts for navigating and editing commands in

the CLI. Although not specifically tested on the CCNA exam, these shortcuts can save you time when using the simulator during the exam.

Table 29-2 *Hot Keys and Shortcuts*

Technet24

Keyboard Command	What Happens
Navigation Key Sequences	
Up arrow or Ctrl+P	Displays the most recently used command. If you press the sequence again, the next most recent command appears, until the history buffer is exhausted. (The P stands for *previous*.)
Down arrow or Ctrl+N	Moves forward to the more recently entered commands, in case you have gone too far back into the history buffer. (The N stands for *next*.)
Left arrow or Ctrl+B	Moves the cursor backward in the currently displayed command without deleting characters. (The B stands for *back*.)
Right arrow or Ctrl+F	Moves the cursor forward in the currently displayed command without deleting characters. (The F stands for *forward*.)
Tab	Completes a partial command name entry.
Backspace	Moves the cursor backward in the currently displayed command, deleting characters.
Ctrl+A	Moves the cursor directly to the first character of the currently displayed command.
Ctrl+E	Moves the cursor directly to the end of the currently displayed command.

Ctrl+R	Redisplays the command line with all characters. This command is useful when messages clutter the screen.
Ctrl+D	Deletes a single character.
Esc+B	Moves back one word.
Esc+F	Moves forward one word.
At the --More Prompt	
Enter key	Displays the next line.
Spacebar	Displays the next screen.
Any other alphanumeric key	Returns to the EXEC prompt.
Break Keys	
Ctrl+C	When in any configuration mode, ends the configuration mode and returns to privileged EXEC mode. When in setup mode, reverts to the command prompt.
Ctrl+Z	When in any configuration mode, ends the configuration mode and returns to privileged EXEC mode. When in user or privileged EXEC mode, logs you out of the router.
Ctrl+Shift+6	Acts as an all-purpose break sequence. Use to abort DNS lookups, traceroutes, and pings.

Technet24

Command History

By default, the Cisco IOS stores in a history buffer the 10 commands you have most recently entered. This gives you a quick way to move backward and forward in the history of commands, choose one, and then edit it before reissuing the command. To view or configure the command history buffer, use the commands in Table 29-3. Although this table shows the switch prompt, these commands are also appropriate for a router.

Table 29-3 *Command History Buffer Commands*

Command Syntax	Description
`switch# show history`	Displays the commands currently stored in the history buffer.
`switch# terminal history`	Enables terminal history. This command can be run from either user or privileged EXEC mode.
`switch# terminal history size 50`	Configures the terminal history size. The terminal history can maintain 0–256 command lines.
`switch# terminal no history size`	Resets the terminal history size to the default value of 20 command lines in Cisco IOS 15.
`switch# terminal no history`	Disables terminal history.

IOS Examination Commands

To verify and troubleshoot network operation, you use **show** commands. Figure 29-3 delineates the different **show** commands, as follows:

• Commands applicable to Cisco IOS (stored in RAM)

Technet24

• Commands that apply to the backup configuration file stored in NVRAM

• Commands that apply to Flash or specific interfaces

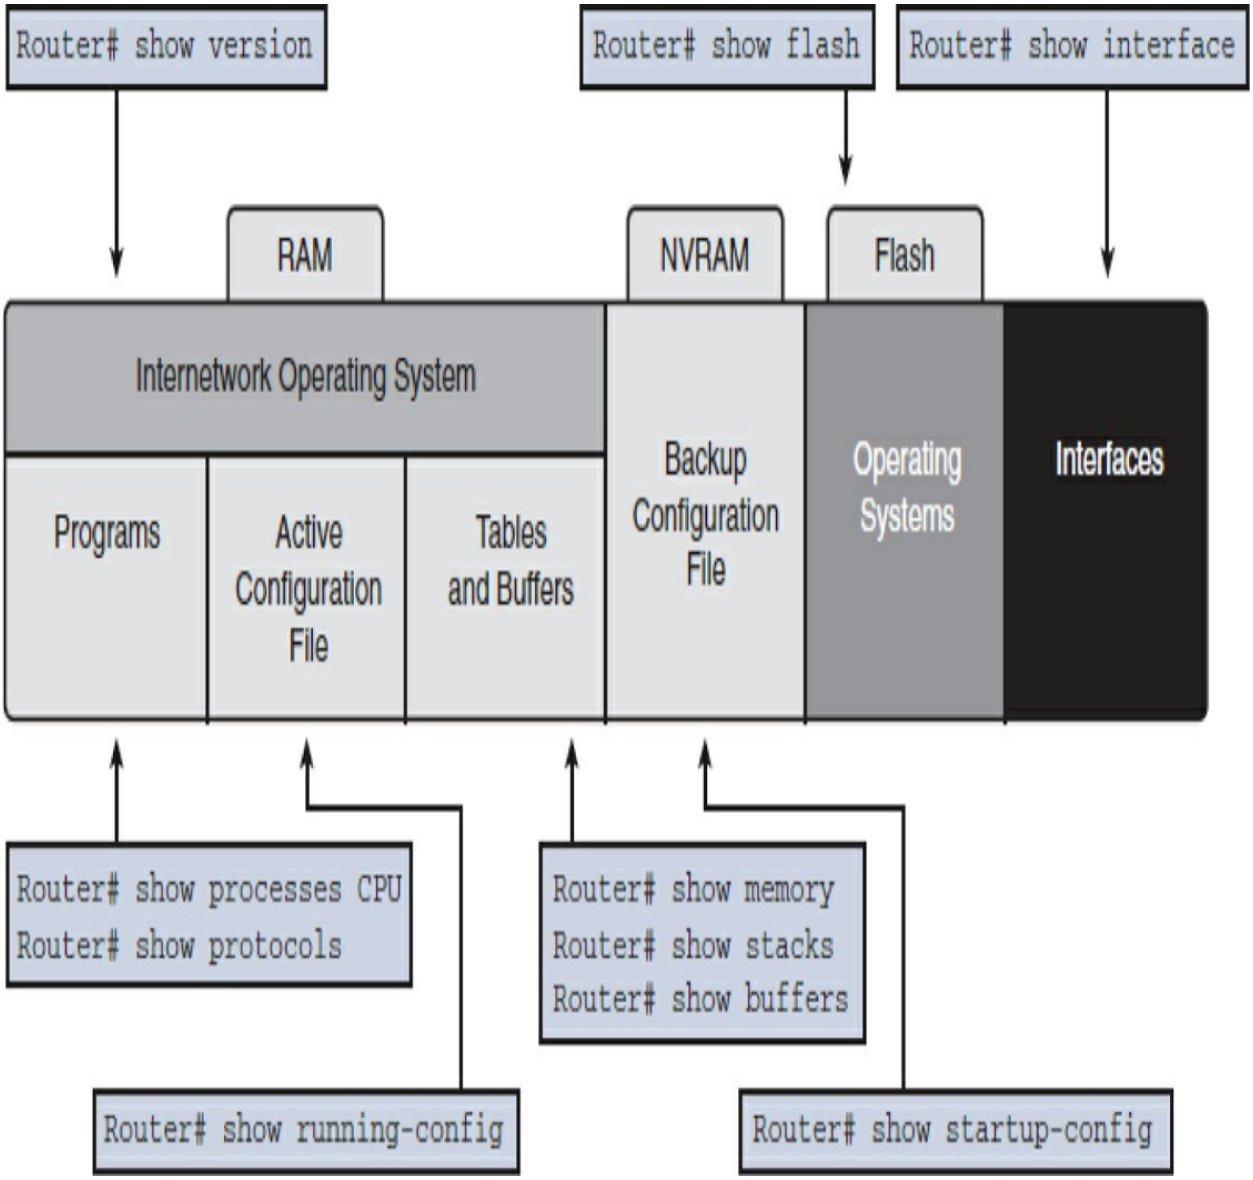

Figure 29-3 *Typical* **show** *Commands and the Information Provided*

Subconfiguration Modes

To enter global configuration mode, enter the **configure terminal** command.

From global configuration mode, Cisco IOS provides a multitude of subconfiguration modes. Table 29-4 summarizes the most common subconfiguration modes pertinent to the CCNA exam.

Table 29-4 *Cisco Device Subconfiguration Modes*

Prompt	Name of Mode	Examples of Commands Used to Reach This Mode
`hostname(config)#`	Global	configure terminal
`hostname(config-line)#`	Line	line console 0 line vty 0 15
`hostname(config-if)#`	Interface	interface fastethernet 0/0
`hostname(config-router)#`	Router	router rip router eigrp 100

Basic Switch Configuration Commands

Table 29-5 reviews basic switch configuration commands.

Table 29-5 *Basic Switch Configuration Commands*

Technet24

Command Description	Command Syntax
Enter global configuration mode.	`Switch# configure terminal`
Configure a name for the device.	`Switch(config)# hostname S1`
Enter the interface configuration mode for the VLAN 1 interface.	`S1(config)# interface vlan 1`
Configure the interface IP address.	`S1(config-if)# ip address 172.17.99.11 255.255.255.0`
Enable the interface.	`S1(config-if)# no shutdown`
Return to global configuration mode.	`S1(config-if)# exit`
Enter the interface to assign the VLAN.	`S1(config)# interface fastethernet 0/6`

Define the VLAN membership mode for the port.	`S1(config-if)# switchport mode access`
Assign the port to a VLAN.	`S1(config-if)# switchport access vlan 123`
Configure the interface duplex mode to enable AUTO duplex configuration.	`S1(config-if)# duplex auto`
Configure the interface speed and enable AUTO speed configuration.	`S1(config-if)# speed auto`
Enable auto-MDIX on the interface.	`S1(config-if)# mdix auto`
Return to global configuration mode.	`S1(config-if)# exit`
Configure the default gateway on the switch.	`S1(config)# ip default-gateway 172.17.50.1`

Technet24

Configure the HTTP server for authentication using the enable password, which is the default method of HTTP server user authentication.	`S1(config)# ip http authentication enable`
Enable the HTTP server.	`S1(config)# ip http server`
Switch from global configuration mode to line configuration mode for console 0.	`S1(config)# line console 0`
Set **cisco** as the password for the console 0 line on the switch.	`S1(config-line)# password cisco`
Set the console line to require the password to be entered before access is granted.	`S1(config-line)# login`
Return to global configuration mode.	`S1(config-if)# exit`
Switch from global configuration mode to line configuration mode for vty terminals 0–15.	`S1(config)# line vty 0 15`

Set **cisco** as the password for the vty lines on the switch.	`S1(config-line)# password cisco`
Set the vty line to require the password to be entered before access is granted.	`S1(config-line)# login`
Return to global configuration mode.	`S1(config-line)# exit`
Configure **cisco** as the enable password to enter privileged EXEC mode.	`S1(config)# enable password cisco`
Configure **class** as the enable secret password to enter privileged EXEC mode. This password overrides **enable password**.	`S1(config)# enable secret class`
Encrypt all the system passwords that are stored in plaintext.	`S1(config)# service password-encryption`
Configure a login banner. The # character delimits the beginning and end of the banner.	`S1 (config)# banner login #Authorized Personnel Only!#`
Configure a message of the day (MOTD) login banner. The # character delimits the beginning and end of the banner.	`S1(config)# banner motd #Device maintenance will be occurring on Friday!#`
Return to privileged EXEC mode.	`S1(config)# end`
Save the running configuration to the switch startup configuration.	`S1# copy running-config startup-config`

Technet24

To configure multiple ports with the same command, use the **interface range** command. For example, to configure ports 6–10 as access ports belonging to VLAN 10, you enter the following:

```
Switch(config)# interface range FastEthernet 0/6 – 10
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
```

Half Duplex, Full Duplex, and Port Speed

Half-duplex communication is unidirectional data flow in which a device can either send or receive on an Ethernet LAN—but not both at the same time. Today's LAN networking devices and end device network interface cards (NICs) operate at full duplex as long as the device is connected to another device capable of full-duplex communication. Full-duplex communication increases the effective bandwidth by allowing both ends of a connection to transmit and receive data simultaneously; this is known as *bidirectional*. This microsegmented LAN is collision free. Gigabit Ethernet and 10-Gbps NICs require full-duplex connections to operate. Port speed is simply the bandwidth rating of the port. The most common speeds today are 100 Mbps, 1 Gbps, and 10 Gbps.

Although the default duplex and speed setting for Cisco Catalyst 2960 and 3560 switches is **auto**, you can manually configure speed with the **speed** and **duplex** commands.

Note

Setting the duplex mode and speed of switch ports can cause issues if one end is mismatched or set to autonegotiation. In addition, all fiber-optic ports, such as 100BASE-FX ports, operate only at one preset speed and are always full duplex.

Automatic Medium-Dependent Interface Crossover (auto-MDIX)

In the past, switch-to-switch or switch-to-router connections required using different Ethernet cables (crossover or straight-through). Using the automatic medium-dependent interface crossover (auto-MDIX) feature on an interface eliminates this problem. When auto-MDIX is enabled, the interface automatically detects the required cable connection type (straight-through or crossover) and configures the connection appropriately.

The auto-MDIX feature is enabled by default starting with IOS 12.2 and the Catalyst 2960 and Catalyst 3560 switches. The Gigabit Ethernet standard requires auto-MDIX, so any 1000-Mbps port has this capability. When using auto-MDIX on an interface, the interface speed and duplex must be set to **auto** so that the feature operates correctly.

Verifying Network Connectivity

Using and interpreting the output of various testing tools is often the first step in isolating the cause of a network connectivity issue. The **ping** command can

Technet24

systematically test connectivity by looking for answers to the following questions, in this order:

Can an end device ping itself?

Can an end device ping its default gateway?

Can an end device ping the destination?

By using the **ping** command in this ordered sequence, you can isolate problems more quickly. If local connectivity is not an issue—in other words, if the end device can successfully ping its default gateway—using the **traceroute** utility can help isolate the point in the path from source to destination where the traffic stops.

As a first step in the testing sequence, verify the operation of the TCP/IP stack on the local host by pinging the loopback address, 127.0.0.1, as Example 29-2 demonstrates.

Example 29-2 *Testing the TCP/IP Stack on a Windows PC*

```
C:\> ping 127.0.0.1

Pinging 127.0.0.1 with 32 bytes of data:

Reply from 127.0.0.1: bytes=32 time<1ms TTL=64
Reply from 127.0.0.1: bytes=32 time<1ms TTL=64
```

```
Reply from 127.0.0.1: bytes=32 time<1ms TTL=64
Reply from 127.0.0.1: bytes=32 time<1ms TTL=64

Ping statistics for 127.0.0.1:
 Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
 Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

This test should succeed regardless of whether the host is connected to the network, so a failure indicates a software or hardware problem on the host itself. Either the network interface is not operating properly or support for the TCP/IP stack has been inadvertently removed from the operating system.

Next, verify connectivity to the default gateway. Determine the default gateway address by using **ipconfig** and then attempt to ping it, as in Example 29-3.

Example 29-3 *Testing Connectivity to the Default Gateway on a Windows PC*

```
C:\> ipconfig

Windows IP Configuration


Ethernet adapter Local Area Connection:
```

Technet24

```
     Connection-specific DNS Suffix . : cisco.com
     IP Address. . . . . . . . . . . .: 192.168.1.25
     Subnet Mask . . . . . . . . . . .: 255.255.255.0
     Default Gateway . . . . . . . . . : 192.168.1.1

C:\> ping 192.168.1.1

Pinging 192.168.1.1 with 32 bytes of data:

Reply from 192.168.1.1: bytes=32 time=162ms TTL=255
Reply from 192.168.1.1: bytes=32 time=69ms TTL=255
Reply from 192.168.1.1: bytes=32 time=82ms TTL=255
Reply from 192.168.1.1: bytes=32 time=72ms TTL=255

Ping statistics for 192.168.1.1:
     Packets: Sent = 4, Received = 4, Lost = 0 (0% los
Approximate round trip times in milli-seconds:
     Minimum = 69ms, Maximum = 162ms, Average = 96ms
```

Failure here can indicate several problems, which must be checked in a systematic sequence. One possible order might be the following:

Is the cabling from the PC to the switch correct? Are link lights lit?

Is the configuration on the PC correct according to the logical map of the network?

Are the affected interfaces on the switch the cause of the problem? Is there a duplex, speed, or auto-MDIX mismatch? Are there VLAN misconfigurations?

Is the cabling from the switch to the router correct? Are link lights lit?

Is the configuration on the router interface correct according to the logical map of the network? Is the interface active?

Finally, verify connectivity to the destination by pinging it. Assume that you are trying to reach a server at 192.168.3.100. Example 29-4 shows a successful **ping** test to the destination.

Example 29-4 *Testing Connectivity to the Destination on a Windows PC*

```
PC> ping 192.168.3.100

Pinging 192.168.3.100 with 32 bytes of data:

Reply from 192.168.3.100: bytes=32 time=200ms TTL=126
Reply from 192.168.3.100: bytes=32 time=185ms TTL=126
Reply from 192.168.3.100: bytes=32 time=186ms TTL=126
Reply from 192.168.3.100: bytes=32 time=200ms TTL=126

Ping statistics for 192.168.3.100:
 Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
```

Technet24

```
    Minimum = 185ms, Maximum = 200ms, Average = 192ms
```

Failure here indicates a failure in the path beyond the default gateway interface because you already successfully tested connectivity to the default gateway. From a Windows PC, the best tool to use to find the break in the path is the **tracert** command (see Example 29-5).

Note

Both macOS and Linux use the **traceroute** command rather than **tracert**.

Example 29-5 *Tracing the Route from a Windows PC*

```
C:\> tracert 192.168.3.100

Tracing route to 192.168.3.100 over a maximum of 30 ho

 1    97 ms    75 ms    72 ms       192.168.1.1
 2    104 ms  119 ms   117  ms      192.168.2.2
 3    *        *        *           Request timed out.
 4    *        *        *           Request timed out.
 5    *        *        *           Request timed out.
 6    ^C
C:\>
```

Note

Failure at hops 3, 4, and 5 in Example 29-5 could indicate that these routers are configured to not send ICMP messages back to the source.

As shown in in Example 29-5, the last successful hop on the way to the destination was 192.168.2.2. If you have administrator rights to 192.168.2.2, you can continue your research by remotely accessing the command line on 192.168.2.2 and investigating why traffic will not go any further. In addition, other devices between 192.168.2.2 and 192.168.3.100 could be the source of the problem. The point is, you want to use your **ping** and **tracert** tests, as well as your network documentation, to proceed in a logical sequence from source to destination.

Regardless of how simple or complex your network is, using **ping** and **tracert** from the source to the destination is a simple yet powerful way to systematically verify end-to-end connectivity, as well as locate breaks in a path from one source to one destination.

Troubleshoot Interface and Cable Issues

The physical layer is often the reason a network issue exists—power outage, disconnected cable, power-cycled devices, hardware failures, and so on. This

Technet24

section looks at some troubleshooting tools, in addition to the approach of actually walking over to the wiring closet or network device and “physically” checking Layer 1.

Media Issues

Besides failing hardware, common physical layer issues occur with media. Consider a few examples:

• New equipment is installed that introduces electromagnetic interference (EMI) sources into the environment.

• Cable runs too close to powerful motors, such as an elevator.

• Poor cable management puts a strain on some RJ-45 connectors, causing one or more wires to break.

• New applications change traffic patterns.

• When new equipment is connected to a switch, the connection operates in half-duplex mode or a duplex mismatch occurs, which can lead to an excessive number of collisions.

Figure 29-4 shows an excellent troubleshooting flowchart that you can use in troubleshooting switch media issues.

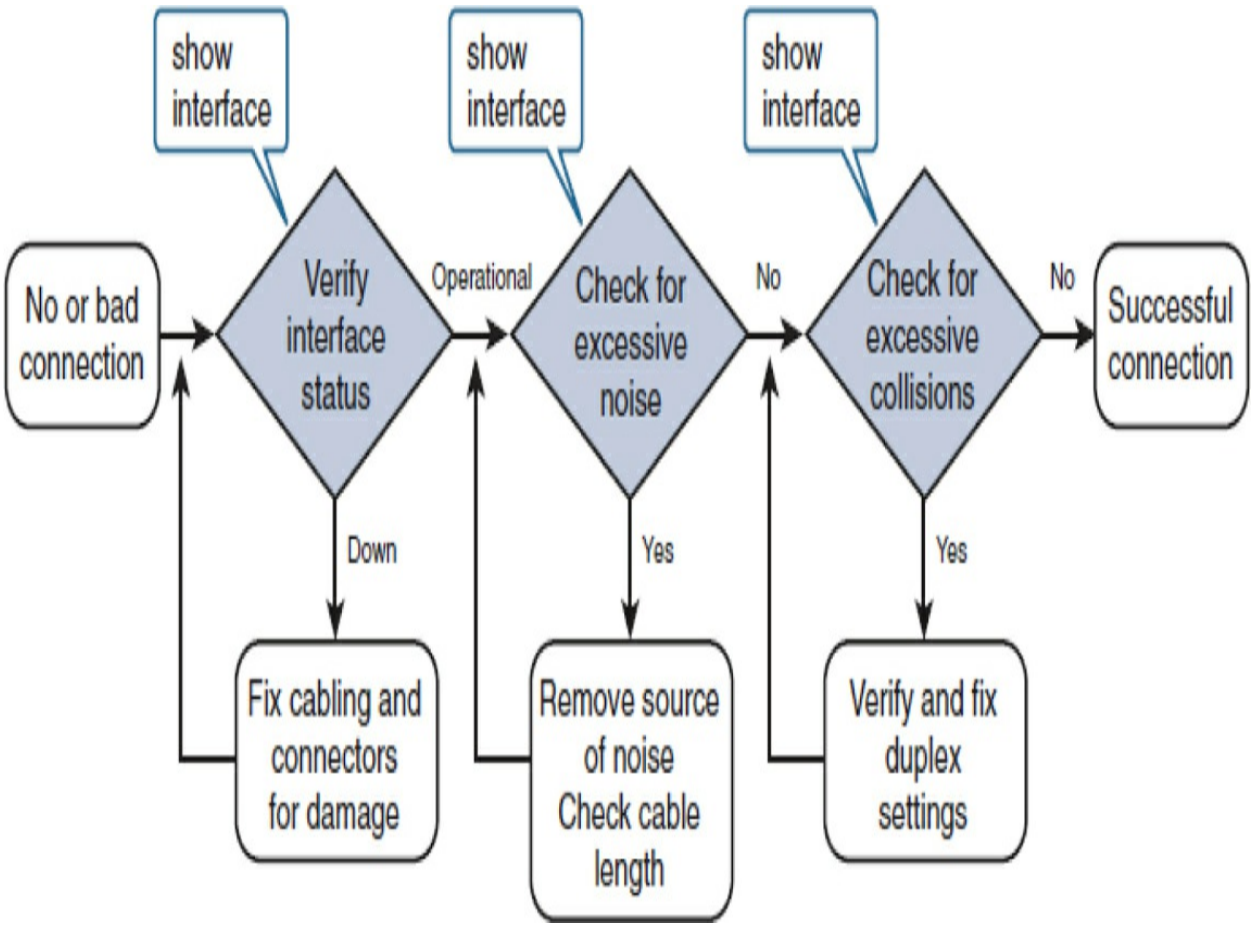

Figure 29-4 *Troubleshooting Switch Media Issues*

Next, examine the output from the **show interface** and **show interface status** commands, as described in the next section.

Interface Status and Switch Configuration

Because today we are focusing on switch troubleshooting, we look at the **show** commands that help in troubleshooting the basic configuration.

Interface Status Codes

In general, interfaces are either “up” or “down.” However, when an interface

Technet24

is “down” and you don’t know why, the code in the **show interfaces** command provides more information to help you determine the reason. Table 29-6 lists the code combinations and some possible causes for each status.

Table 29-6 *LAN Switch Interface Status Codes*

Line Status	Protocol Status	Interface Status	Typical Root Cause
Administratively down	Down	disabled	The interface is configured with the **shutdown** command.
Down	Down	notconnect	No cable exists, the cable is bad, incorrect cable pinouts are used, the two connected devices have mismatched speeds, or the device on the other end of the cable is powered off or the other interface is shut down.
Up	Down	notconnect	An interface up/down state is not expected on LAN switch interfaces. This indicates a Layer 2 problem on Layer 3 devices.
Down	Down (err-disabled)	err-disabled	Port security has disabled the interface. The network administrator must manually reenable the interface.
Up	Up	connect	The interface is working.

Technet24

Duplex and Speed Mismatches

One of the most common problems is issues with speed and/or duplex mismatches. On switches and routers, the **speed** {**10** | **100** | **1000**} interface subcommand and the **duplex** {**half** | **full**} interface subcommand set these values. Note that configuring both speed and duplex on a switch interface disables the IEEE-standard autonegotiation process on that interface.

The **show interfaces status** and **show interfaces** commands list both the speed and duplex settings on an interface, as Example 29-6 shows.

Example 29-6 *Commands to Verify Speed and Duplex Settings*

```
S1# show interface status

Port Name  Status      Vlan   Duplex  Speed   Type
Fa0/1      connected   trunk  full    100     10/100Ba
Fa0/2      connected   1      half    100     10/100Ba
Fa0/3      connected   1      a-full  a-100   10/100Ba
Fa0/4      disabled    1      auto    auto    10/100Ba
Fa0/5      disabled    1      auto    auto    10/100Ba
Fa0/6      notconnect  1      auto    auto    10/100Ba

!Remaining output omitted
S1# show interface fa0/3
FastEthernet0/1 is up, line protocol is up (connected)
   Hardware is Fast Ethernet, address is 001b.5302.4e8
```

```
    MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
 Keepalive set (10 sec)
  Full-duplex, 100Mb/s, media type is 10/100BaseTX
  input flow-control is off, output flow-control is uns
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 1000 bits/sec, 1 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     2745 packets input, 330885 bytes, 0 no buffer
     Received 1386 broadcasts (0 multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignor
     0 watchdog, 425 multicast, 0 pause input
     0 input packets with dribble condition detected
     56989 packets output, 4125809 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 PAUSE output
     0 output buffer failures, 0 output buffers swapped
```

Notice that both commands show the duplex and speed settings of the

Technet24

interface. However, the **show interface status** command is preferred for troubleshooting duplex or speed mismatches because it shows exactly how the switch determined the duplex and speed of the interface. In the duplex column, **a-full** means the switch autonegotiated full duplex. The setting **full** or **half** means that the switch was configured at that duplex setting. Autonegotiation has been disabled. In the speed column, **a-100** means the switch autonegotiated 100 Mbps as the speed. The setting **10** or **100** means that the switch was configured at that speed setting.

Finding a duplex mismatch can be much more difficult than finding a speed mismatch because if the duplex settings do not match on the ends of an Ethernet segment, the switch interface will still be in a connect (up/up) state. In this case, the interface works, but the network might work poorly, with hosts experiencing poor performance and intermittent communication problems. To identify duplex mismatch problems, check the duplex setting on each end of the link and watch for incrementing collision and late collision counters, as highlighted in the output at the end of Example 29-6.

Common Layer 1 Problems On "Up" Interfaces

When a switch interface is "up," it does not necessarily mean that the interface is operating in an optimal state. For this reason, Cisco IOS tracks certain counters to help identify problems that can occur even though the interface is in a connect state. The output in Example 29-6 highlights these counters. Table 29-7 summarizes three general types of Layer 1 interface

problems that can occur while an interface is in the “up,” connected, state.

Table 29-7 *Common LAN Layer 1 Problem Indicators*

Type of Problem	Counter Values Indicating This Problem	Common Root Causes
Excessive noise	Many input errors, few collisions	Wrong cable category (Cat5, Cat5E, Cat6), damaged cables, EMI
Collisions	More than roughly 0.1% of all frames are collisions	Duplex mismatch (seen on the half-duplex side), jabber, DoS attack
Late collisions	Increasing late collisions	Collision domain or single cable too long, duplex mismatch

Study Resources

For today’s exam topics, refer to the following resources for more study.

Technet24

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	2
	17
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	2
	4
	6
	7
Portable Command Guide	8

Day 28. IPv4 Addressing

CCNA 200-301 Exam Topics

• 1.6 Configure and verify IPv4 addressing and subnetting

• 1.7 Describe private IPv4 addressing

Key Topics

Today we focus on reviewing the structure of an IPv4 address, the classes, and private and public IPv4 addresses. Then we turn our focus to IPv4 subnetting.

By now, you should be able to subnet quickly. For example, you should be able to quickly answer a question such as the following: If you are given a /16 network, what subnet mask would you use to maximize the total number of subnets while still providing enough addresses for the largest subnet with 500 hosts? The answer is 255.255.254.0, or /23. This gives you 128 subnets with 510 usable hosts per subnet. You should be able to quickly calculate this information.

The CCNA exam promises to contain many subnetting and subnetting-related questions. Therefore, we devote some time to this necessary skill and also look at designing addressing schemes using variable-length subnet masking (VLSM).

Technet24

IPv4 Addressing

Although IPv6 is rapidly permeating the networks of the world, most networks still have large IPv4 implementations. Especially on private networks, migration away from IPv4 will take years to complete. Clearly, IPv4 and your skill in its use are still in demand.

Header Format

Figure 28-1 shows the layout of the IPv4 header.

Note that each IP packet carries this header, which includes a source IP address and destination IP address.

An IP address consists of two parts:

• **Network ID:** The high-order, or leftmost, bits specify the network address component of the address.

• **Host ID:** The low-order, or rightmost, bits specify the host address component of the address.

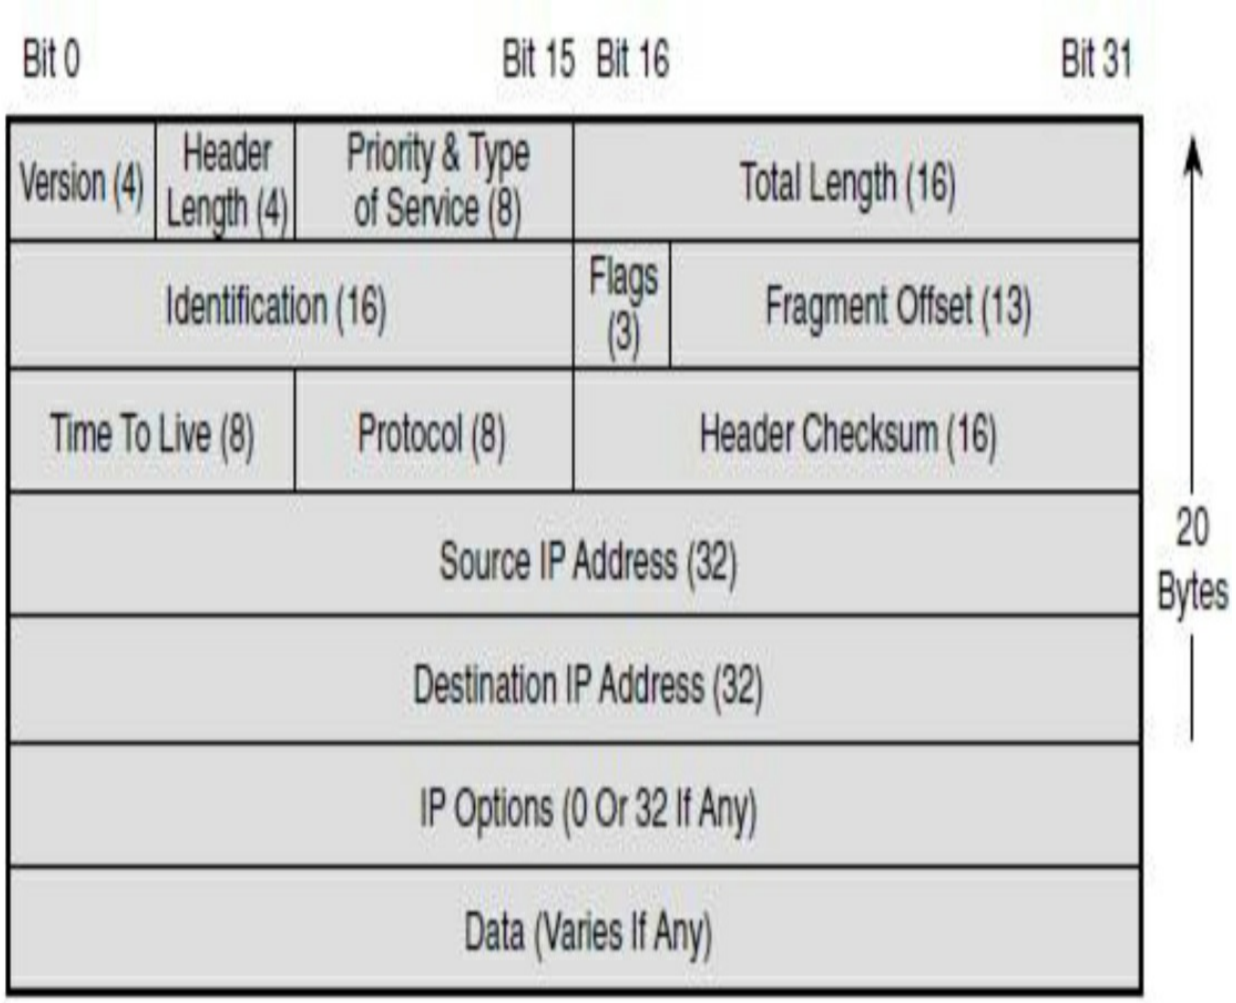

Figure 28-1 *IPv4 Header Format*

Classes of Addresses

From the beginning, IPv4 was designed with class structure: Classes A, B, C, D, and E. Class D is used for multicasting addresses, and Class E is reserved for experimentation. Classes A, B, and C are assigned to network hosts. To provide a hierarchical structure, these classes are divided into network and host portions, as Figure 28-2 shows. The high-order bits specify the network ID, and the low-order bits specify the host ID.

Technet24

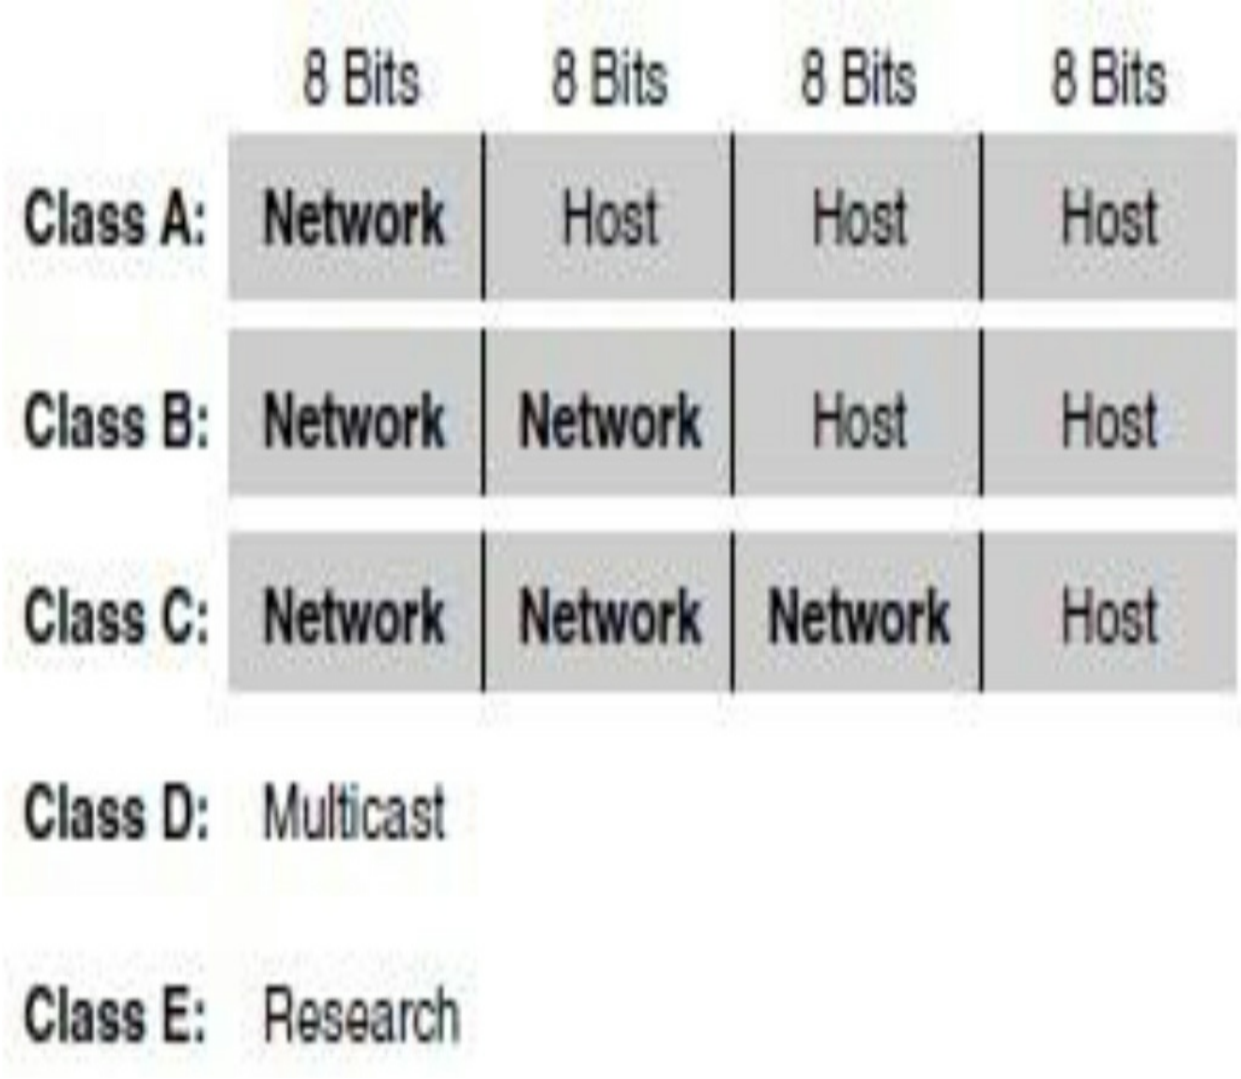

Figure 28-2 *Network/Host Boundary for Each Class of IPv4 Address*

In a classful addressing scheme, devices that operate at Layer 3 can determine the address class of an IP address from the format of the first few bits in the first octet. Initially, this was important so that a networking device could apply the default subnet mask for the address and determine the host address. Table 28-1 summarizes how addresses are divided into classes, the default subnet mask, the number of networks per class, and the number of hosts per classful network address.

Table 28-1 *IPv4 Address Classes*

Address Class	First Octet Range (Decimal)	First Octet Bits (Highlighted Bits Do Not Change)	Network (N) and Host (H) Portions of Addresses	Default Subnet Mask (Decimal and Binary)	Number of Possible Networks and Hosts per Network
A	1–127	00000000–01111111	N.H.H.H	255.0.0.0	2^7, or 128, networks
				11111111.00000000.00000000.00000000	2^{24}–2, or 16,777,214, hosts per network
B	128–191	10000000–10111111	N.N.H.H	255.255.0.0	2^{14}, or 16,384, networks
				11111111.11111111.00000000.00000000	2^{16}–2, or 65,534, hosts per network

Technet24

C	192–223	11000000–11011111	N.N.N.H	255.255.255.0	2^{21}, or 2,097,152, networks
				11111111.11111111.11111111.00000000	2^{8}–2, or 254, hosts per network
D	224–239	11100000–11101111	Not used for host addressing		
E	240–255	11110000–11111111	Not used for host addressing		

In the last column of Table 28-1, the –2 for hosts per network accounts for the reserved network and broadcast addresses for each network. These two addresses cannot be assigned to hosts.

Note

We do not review the process of converting between binary and decimal. At this point in your studies, you should be

comfortable moving between the two numbering systems. If not, take some time to practice this necessary skill. You can search the Internet for binary conversion tricks, tips, and games to practice. The Cisco Learning Network has a fun game you can play, at https://learningnetwork.cisco.com/docs/DOC-1803.

Purpose of the Subnet Mask

Subnet masks are always a series of 1 bits followed by a series of 0 bits. The boundary where the series changes from 1s to 0s is the boundary between the network and the host. This is how a device that operates at Layer 3 determines the network address for a packet: by finding the bit boundary where the series of 1 bits ends and the series of 0 bits begins. The bit boundary for default subnet masks breaks on the octet boundary. Determining the network address for an IP address that uses a default mask is easy.

Say that a router receives a packet destined for 192.168.1.51. By ANDing the IP address and the subnet mask, the router determines the network address for the packet. By the ANDing rules, a 1 AND a 1 equals 1. All other possibilities equal 0. Table 28-2 shows the results of the ANDing operation. Notice that the host bits in the last octet are ignored.

Table 28-2 *ANDing an IP Address and Subnet Mask to Find the Network Address*

Technet24

Destination address	192.168.1.51	11000000.10101000.00000001.00110011
Subnet mask	255.255.255.0	11111111.11111111.11111111.00000000
Network address	192.168.1.0	11000000.10101000.00000001.00000000

The bit boundary can now occur in just about any place in the 32 bits. Table 28-3 summarizes the values for the last nonzero octet in a subnet mask.

Table 28-3 *Subnet Mask Binary Values*

Mask (Decimal)	Mask (Binary)	Network Bits	Host Bits
0	00000000	0	8
128	10000000	1	7
192	11000000	2	6
224	11100000	3	5
240	11110000	4	4
248	11111000	5	3
252	11111100	6	2
254	11111110	7	1
255	11111111	8	0

Private and Public IP Addressing

RFC 1918, “Address Allocation for Private Internets,” eased the demand for

Technet24

IP addresses by reserving the following addresses for use in private internetworks:

• **Class A:** 10.0.0.0/8 (10.0.0.0–10.255.255.255)

• **Class B:** 172.16.0.0/12 (172.16.0.0–172.31.255.255)

• **Class C:** 192.168.0.0/16 (192.168.0.0–192.168.255.255)

If you are addressing a nonpublic intranet, these private addresses are normally used instead of globally unique public addresses. This provides flexibility in your addressing design. Any organization can take full advantage of an entire Class A address (10.0.0.0/8). Forwarding traffic to the public Internet requires translation to a public address using Network Address Translation (NAT). But by overloading an Internet-routable address with many private addresses, a company needs only a handful of public addresses. Day 8, "NAT," reviews NAT operation and configuration in greater detail.

Subnetting in Four Steps

Everyone has a preferred method of subnetting. Each teacher uses a slightly different strategy to help students master this crucial skill, and each of the suggested study resources has a slightly different way of approaching this subject.

The method I prefer consists of four steps:

Determine how many bits to borrow, based on the host requirements.

Determine the new subnet mask.

Determine the subnet multiplier.

List the subnets, including the subnetwork address, host range, and broadcast address.

The best way to demonstrate this method is to use an example. Assume that you are given the network address 192.168.1.0 with the default subnet mask 255.255.255.0. The network address and subnet mask can be written as 192.168.1.0/24. The /24 represents the subnet mask in a shorter notation and means that the first 24 bits are network bits.

Now further assume that you need 30 hosts per network and want to create as many subnets for the given address space as possible. With these network requirements, you can now subnet the address space.

Determine How Many Bits to Borrow

To determine the number of bits you can borrow, you first must know how many host bits you have to start with. Because the first 24 bits are network bits in this example, the remaining 8 bits are host bits.

Technet24

Because our requirement specifies 30 host addresses per subnet, we need to first determine the minimum number of host bits to leave. The remaining bits can be borrowed:

Host Bits = Bits Borrowed + Bits Left

To provide enough address space for 30 hosts, we need to leave 5 bits. Use the following formula:

$2^{BL} – 2$ = number of host addresses

The exponent BL is bits left in the host portion.

Remember, the –2 accounts for the network and broadcast addresses that cannot be assigned to hosts.

In this example, leaving 5 bits in the host portion provides the right number of host addresses:

$2^5 – 2 = 30$

Because we have 3 bits remaining in the original host portion, we borrow all these bits to satisfy the requirement to “create as many subnets as possible.” To determine how many subnets we can create, use the following formula:

2^{BB} = number of subnets

The exponent BB is bits borrowed from the host portion.

In this example, borrowing 3 bits from the host portion creates eight subnets: $2^3 = 8$.

As Table 28-4 shows, the 3 bits are borrowed from the leftmost bits in the host portion. The highlighted bits in the table show all possible combinations of manipulating the 8 bits borrowed to create the subnets.

Table 28-4 *Binary and Decimal Value of the Subnetted Octet*

Technet24

Subnet Number	Last Octet Binary Value	Last Octet Decimal Value
0	00000000	.0
1	00100000	.32
2	01000000	.64
3	01100000	.96
4	10000000	.128
5	10100000	.160
6	11000000	.192
7	11100000	.224

Determine the New Subnet Mask

Notice in Table 28-4 that the network bits now include the 3 borrowed host bits in the last octet. Add these 3 bits to the 24 bits in the original subnet mask, and you have a new subnet mask, /27. In decimal format, you turn on the 128, 64, and 32 bits in the last octet, for a value of 224. The new subnet

mask is thus 255.255.255.224.

Determine the Subnet Multiplier

Notice in Table 28-4 that the last octet decimal value increments by 32 with each subnet number. The number 32 is the subnet multiplier. You can quickly find the subnet multiplier by using one of two methods:

• **Method 1:** Subtract the last nonzero octet of the subnet mask from 256. In this example, the last nonzero octet is 224. The subnet multiplier is therefore 256 – 224 = 32.

• **Method 2:** The decimal value of the last bit borrowed is the subnet multiplier. In this example, we borrowed the 128 bit, the 64 bit, and the 32 bit. The 32 bit is the last bit we borrowed and is, therefore, the subnet multiplier.

By using the subnet multiplier, you no longer have to convert binary subnet bits to decimal.

List the Subnets, Host Ranges, and Broadcast Addresses

Listing the subnets, host ranges, and broadcast addresses helps you see the flow of addresses within one address space. Table 28-5 documents our subnet addressing scheme for the 192.168.1.0/24 address space.

Table 28-5 *Subnet Addressing Scheme for 192.168.1.0/24: 30 Hosts per*

Technet24

Subnet

Subnet Number	Subnet Address	Host Range	Broadcast Address
0	192.168.1.0	192.168.1.1– 192.168.1.30	192.168.1.31
1	192.168.1.32	192.168.1.33– 192.168.1.62	192.168.1.63
2	192.168.1.64	192.168.1.65– 192.168.1.94	192.168.1.95
3	192.168.1.96	192.168.1.97– 192.168.1.126	192.168.1.127
4	192.168.1.128	192.168.1.129– 192.168.1.158	192.168.1.159
5	192.168.1.160	192.168.1.161– 192.168.1.190	192.168.1.191
6	192.168.1.192	192.168.1.193– 192.168.1.222	192.168.1.223
7	192.168.1.224	192.168.1.225– 192.168.1.254	192.168.1.255

Technet24

Following are three examples using the four subnetting steps. For brevity, step 4 lists only the first three subnets.

Subnetting Example 1

Subnet the address space 172.16.0.0/16 to provide at least 80 host addresses per subnet while creating as many subnets as possible.

There are 16 host bits. Leave 7 bits for host addresses ($2^7 – 2 = 126$ host addresses per subnet). Borrow the first 9 host bits to create as many subnets as possible ($2^9 = 512$ subnets).

The original subnet mask is /16, or 255.255.0.0. Turn on the next 9 bits starting in the second octet, for a new subnet mask of /25 or 255.255.255.128.

The subnet multiplier is 128, which can be found as 256 – 128 = 128, or because the 128 bit is the last bit borrowed.

For step 4, Table 28-6 lists the first three subnets, host ranges, and broadcast addresses.

Table 28-6 *Subnet Addressing Scheme for Example 1*

Subnet Number	Subnet Address	Host Range	Broadcast Address
0	172.16.0.0	172.16.0.1– 172.16.0.126	172.16.0.127
1	172.16.0.128	172.16.0.129– 172.16.0.254	172.16.0.255
2	172.16.1.0	172.16.1.1– 172.16.1.126	172.16.1.127

Subnetting Example 2

Subnet the address space 172.16.0.0/16 to provide at least 80 subnet addresses.

There are 16 host bits. Borrow the first 7 host bits to create at least 80 subnets (2^7 = 128 subnets). That leaves 9 bits for host addresses, or $2^9 - 2$ = 510 host addresses per subnet.

The original subnet mask is /16, or 255.255.0.0. Turn on the next 7 bits starting in the second octet, for a new subnet mask of /23, or 255.255.254.0.

The subnet multiplier is 2, which can be found as 256 – 254 = 2, or because the 2 bit is the last bit borrowed.

Technet24

For step 4, Table 28-7 lists the first three subnets, host ranges, and broadcast addresses.

Table 28-7 *Subnet Addressing Scheme for Example 2*

Subnet Number	Subnet Address	Host Range	Broadcast Address
0	172.16.0.0	172.16.0.1– 172.16.1.254	172.16.1.255
1	172.16.2.0	172.16.2.1– 172.16.3.254	172.16.3.255
2	172.16.4.0	172.16.4.1– 172.16.5.254	172.16.5.255

Subnetting Example 3

Subnet the address space 172.16.10.0/23 to provide at least 60 host addresses per subnet while creating as many subnets as possible.

There are 9 host bits. Leave 6 bits for host addresses ($2^6 – 2 = 62$ host addresses per subnet). Borrow the first 3 host bits to create as many subnets as possible ($2^3 = 8$ subnets).

The original subnet mask is /23, or 255.255.254.0. Turn on the next 3 bits starting with the last bit in the second octet, for a new subnet mask of /26, or 255.255.255.192.

The subnet multiplier is 64, which can be found as 256 – 192 = 64, or because the 64 bit is the last bit borrowed.

For step 4, Table 28-8 lists the first three subnets, host ranges, and broadcast addresses.

Table 28-8 *Subnet Addressing Scheme for Example 3*

Subnet Number	Subnet Address	Host Range	Broadcast Address
0	172.16.10.0	172.16.10.1– 172.16.10.62	172.16.10.63
1	172.16.10.64	172.16.10.65– 172.16.10.126	172.16.10.127
2	172.16.10.128	172.16.10.129– 172.16.10.190	172.16.10.191

VLSM

Technet24

You probably noticed that the starting address space in Subnetting Example 3 is not an entire classful address. In fact, it is subnet 5 from Subnetting Example 2. In Subnetting Example 3, therefore, we “subnetted a subnet.” In a nutshell, VLSM is subnetting a subnet.

With VLSM, you can customize subnets to fit your network. Subnetting works the same way. You just have to do it more than once to complete your addressing scheme. To avoid overlapping address spaces, start with your largest host requirement, create a subnet for it, and then continue with the next-largest host requirement.

Consider a small example. Given the address space 172.30.4.0/22 and the network requirements in Figure 28-3, apply an addressing scheme that conserves the most addresses for future growth.

We need five subnets: four LAN subnets and one WAN subnet. Starting with the largest host requirement on LAN 3, begin subnetting the address space.

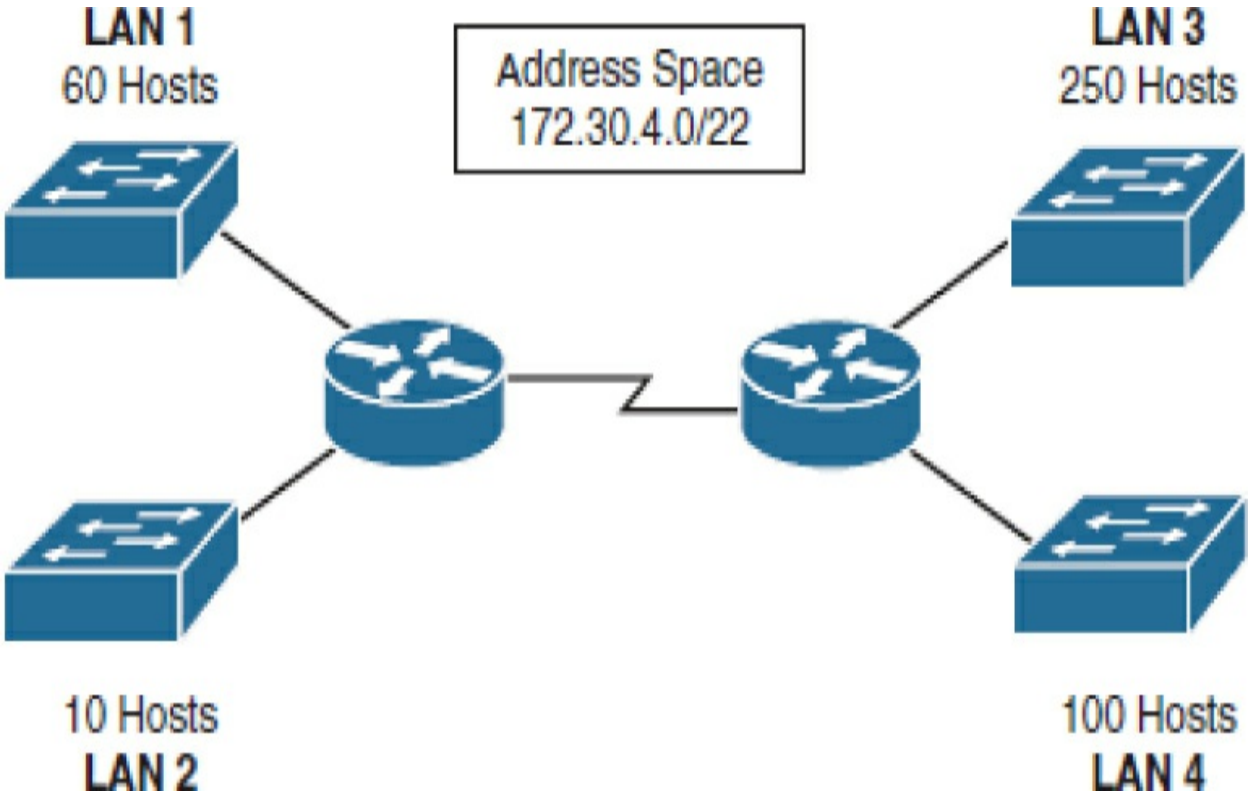

Figure 28-3 *VLSM Example Topology*

To satisfy the 250-host requirement, we leave 8 host bits ($2^8 – 2 = 254$ hosts per subnet). Because we have 10 host bits total, we borrow 2 bits to create the first round of subnets ($2^2 = 4$ subnets). The starting subnet mask is /22, or 255.255.252.0. We turn on the next 2 bits in the subnet mask to get /24, or 255.255.255.0. The multiplier is 1. The four subnets are as follows:

- **Subnet 0:** 172.30.4.0/24
- **Subnet 1:** 172.30.5.0/24
- **Subnet 2:** 172.30.6.0/24
- **Subnet 3:** 172.30.7.0/24

Technet24

Assigning Subnet 0 to LAN 3, we are left with three /24 subnets. Continuing on to the next-largest host requirement on LAN 4, we further subnet Subnet 1, 172.30.5.0/24.

To satisfy the 100-host requirement, we leave 7 bits ($2^7 – 2 = 128$ hosts per subnet). Because we have 8 host bits total, we can borrow only 1 bit to create the subnets ($2^1 = 2$ subnets). The starting subnet mask is /24, or 255.255.255.0. We turn on the next bit in the subnet mask to get /25, or 255.255.255.128. The multiplier is 128. The two subnets are as follows:

• **Subnet 0:** 172.30.5.0/25

• **Subnet 1:** 172.30.5.128/25

Assigning Subnet 0 to LAN 4, we are left with one /25 subnet and two /24 subnets. Continuing on to the next-largest host requirement on LAN 1, we further subnet Subnet 1, 172.30.5.128/25.

To satisfy the 60-host requirement, we leave 6 bits ($2^6 – 2 = 62$ hosts per subnet). Because we have 7 host bits total, we borrow 1 bit to create the subnets ($2^1 = 2$ subnets). The starting subnet mask is /25, or 255.255.255.128. We turn on the next bit in the subnet mask to get /26, or 255.255.255.192. The multiplier is 64. The two subnets are as follows:

• **Subnet 0:** 172.30.5.128/26

• **Subnet 1:** 172.30.5.192/26

Assigning Subnet 0 to LAN 1, we are left with one /26 subnet and two /24 subnets. Finishing our LAN subnetting with LAN 2, we further subnet Subnet 1, 172.30.5.192/26.

To satisfy the 10-host requirement, we leave 4 bits ($2^4 – 2 = 14$ hosts per subnet). Because we have 6 host bits total, we borrow 2 bits to create the subnets ($2^2 = 4$ subnets). The starting subnet mask is /26, or 255.255.255.192. We turn on the next 2 bits in the subnet mask to get /28, or 255.255.255.240. The multiplier is 16. The four subnets are as follows:

• **Subnet 0:** 172.30.5.192/28

• **Subnet 1:** 172.30.5.208/28

• **Subnet 2:** 172.30.5.224/28

• **Subnet 3:** 172.30.5.240/28

Assigning Subnet 0 to LAN 2, we are left with three /28 subnets and two /24 subnets. To finalize our addressing scheme, we need to create a subnet for the WAN link, which needs only two host addresses. We further subnet Subnet 1, 172.30.5.208/28.

To satisfy the two-host requirement, we leave 2 bits ($2^2 – 2 = 2$ hosts per subnet). Because we have 4 host bits total, we borrow 2 bits to create the

Technet24

subnets (2^2 = 4 subnets). The starting subnet mask is /28, or 255.255.255.240. We turn on the next 2 bits in the subnet mask to get /30, or 255.255.255.252. The multiplier is 4. The four subnets are as follows:

- **Subnet 0:** 172.30.5.208/30
- **Subnet 1:** 172.30.5.212/30
- **Subnet 2:** 172.30.5.216/30
- **Subnet 3:** 172.30.5.220/30

We assign Subnet 0 to the WAN link. We are left with three /30 subnets, two /28 subnets, and two /24 subnets.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	11
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	11–15
CCNA 200-301 Official Cert Guide, Volume 2, 2e	14
CCNA Portable Command Guide	1–4

Technet24

Day 27. IPv6 Addressing

CCNA 200-301 Exam Topics

• 1.8 Configure and verify IPv6 addressing and prefix

• 1.9 Describe IPv6 address types

Key Topics

In the early 1990s, the Internet Engineering Task Force (IETF) grew concerned about the exhaustion of IPv4 network addresses and began to look for a replacement for this protocol. This activity led to the development of what is now known as IPv6. Today's review focuses on the IPv6 protocol and IPv6 address types. We also review the various ways to implement IPv6 addressing, including subnetting, autoconfiguring hosts, and running IPv6 and IPv4 in a dual-stack configuration. IPv6 configuration on routers will be reviewed on Day 18, "Basic Router Configuration."

Note

If you have not yet purchased a copy of Rick Graziani's *IPv6 Fundamentals* to add to your library of study tools, now is the time to do so. His book is my definitive source for everything IPv6.

Overview and Benefits of IPv6

Scaling networks today requires a limitless supply of IP addresses and improved mobility that private addressing and NAT alone cannot meet. IPv6 satisfies the increasingly complex requirements of hierarchical addressing that IPv4 does not provide. The main benefits and features of IPv6 include the following:

• **Extended address space:** A 128-bit address space represents about 340 trillion trillion trillion addresses.

• **Stateless address autoconfiguration:** IPv6 provides host devices with a method for generating their own routable IPv6 addresses. IPv6 also supports stateful configuration using DHCPv6.

• **Eliminates the need for NAT/PAT:** NAT/PAT was conceived as part of the solution to IPv4 address depletion. With IPv6, address depletion is no longer an issue. NAT64, however, does play an important role in providing backward compatibility with IPv4.

• **Simpler header:** A simpler header offers several advantages over IPv4:

• Better routing efficiency for performance and forwarding-rate scalability

• No broadcasts and, thus, no potential threat of broadcast storms

• No requirement for processing checksums

Technet24

• Simpler and more efficient extension header mechanisms

• **Mobility and security:** Mobility and security help ensure compliance with mobile IP and IPsec standards:

• IPv4 does not automatically enable mobile devices to move without breaks in established network connections.

• In IPv6, mobility is built in, which means that any IPv6 node can use mobility when necessary.

• IPsec is enabled on every IPv6 node and is available for use, making the IPv6 Internet more secure.

• **Transition strategies:** You can incorporate existing IPv4 capabilities with the added features of IPv6 in several ways:

• You can implement a dual-stack method, with both IPv4 and IPv6 configured on the interface of a network device.

• You can use tunneling, which will become more prominent as the adoption of IPv6 grows.

The IPv6 Protocol

Table 27-1 compares the binary and alphanumeric representations of IPv4 and IPv6 addresses.

Table 27-1 *IPv4 and IPv6 Address Comparison*

	IPv4 (4 Octets)	IPv6 (16 Octets)
Binary representation	11000000.101010 00.00001010. 01100101	10100101.00100100.01110010. 11010011.00101100.10000000. 11011101.00000010.00000000. 00101001.11101100.01111010. 00000000.00101011.11101010. 01110011
Alphanumeric representation	192.168.10.101	2001:0DB8:2C80:DD02:0029:EC7A:002B:EA73
Total IP addresses	4,294,967,296, or 2^{32}	3.4×10^{38}, or 2^{128}

Figure 27-1 compares the IPv4 header with the main IPv6 header. Notice that the IPv6 header is represented in 64-bit words instead of the 32-bit words used by IPv4.

Note

Refer to RFC 2460 and the "Study Resources" section for the full specification of IPv6.

Technet24

IPv4 Header

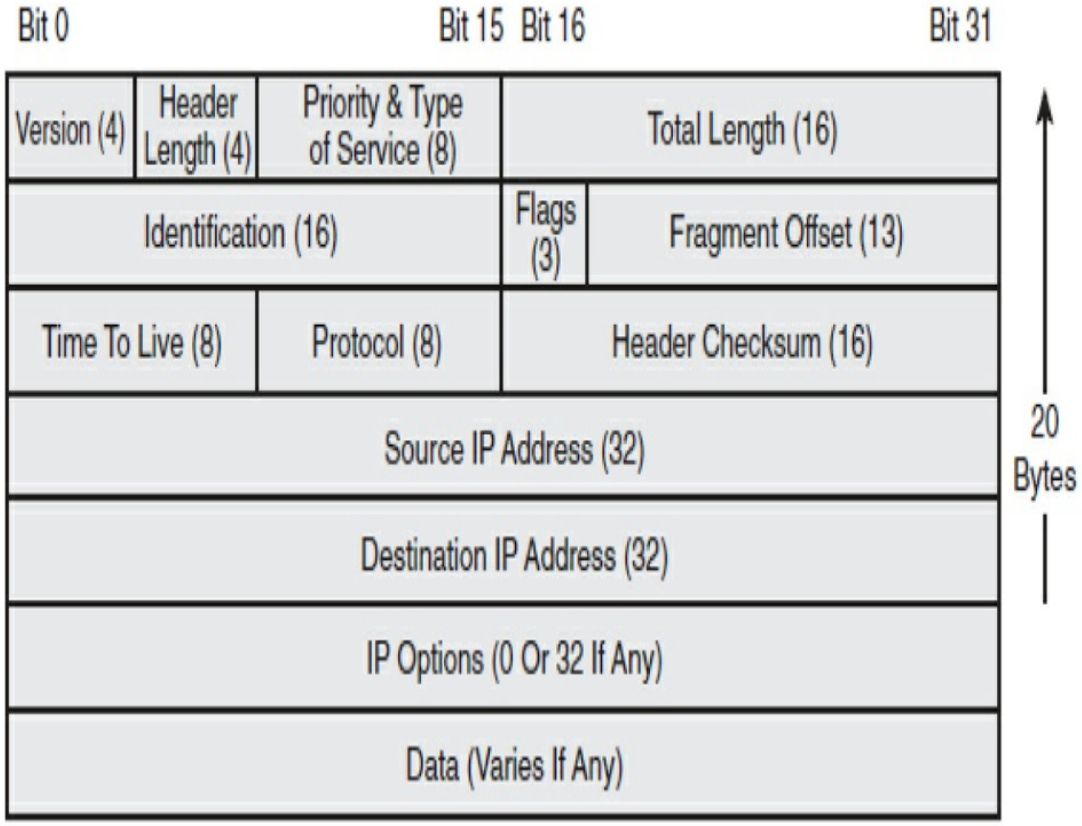

IPv6 Header

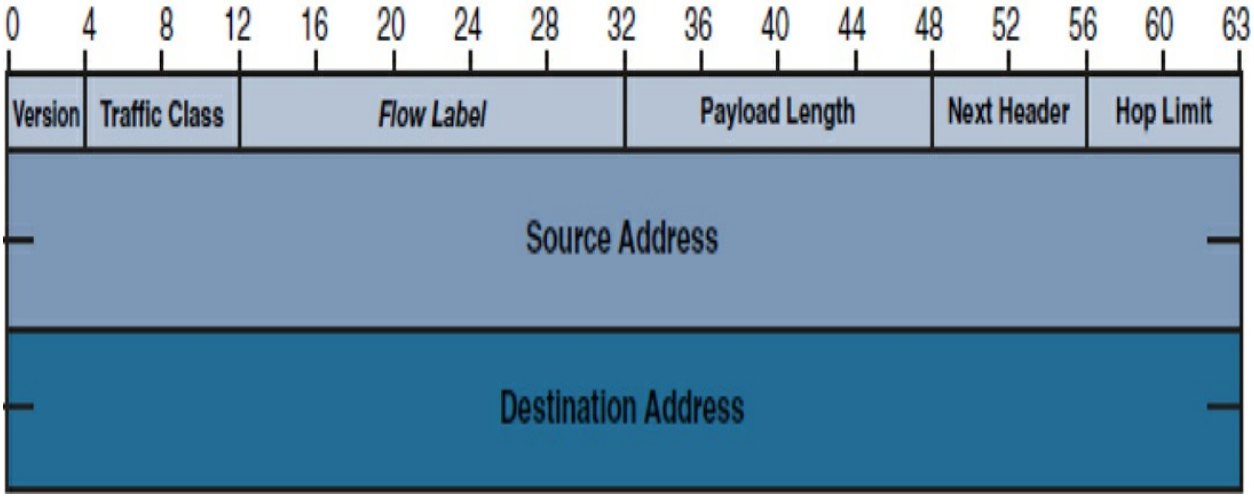

Figure 27-1 *IPv6 Header Format*

IPv6 Address Types

IPv4 has three address types: unicast, multicast, and broadcast. IPv6 does not use broadcast addresses. Instead, IPv6 uses unicast, multicast, and anycast addresses. Figure 27-2 illustrates these three types of IPv6 addresses.

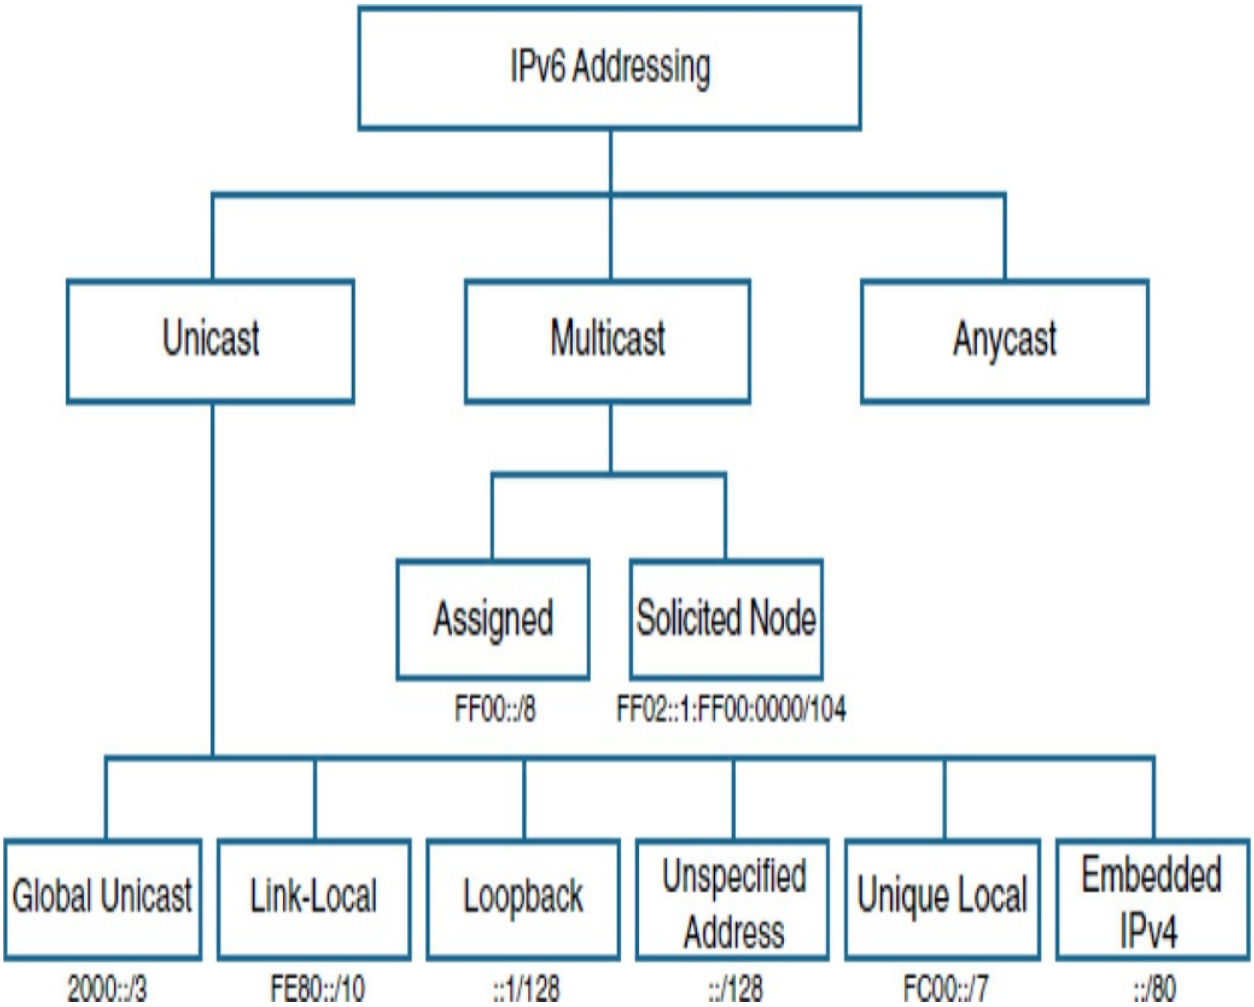

Figure 27-2 *IPv6 Address Types*

Unicast

Technet24

The first classification of IPv6 address types shown in Figure 27-2 is the unicast address. A unicast address uniquely identifies an interface on an IPv6 device. A packet sent to a unicast address is received by the interface that is assigned to that address. Much as with IPv4, source IPv6 addresses must be unicast addresses. Because unicast addressing—as opposed to multicast and anycast addressing—is the major focus for a CCNA candidate, we spend some time reviewing the Unicast branch in Figure 27-2.

Global Unicast Address

IPv6 has an address format that enables aggregation upward, eventually to the ISP. An IPv6 global unicast address is globally unique. Like a public IPv4 address, it can be routed in the Internet without modification. An IPv6 global unicast address consists of a 48-bit global routing prefix, a 16-bit subnet ID, and a 64-bit interface ID. Use Rick Graziani's method of breaking down the IPv6 address with the 3-1-4 rule (also known as the *pi rule*, for 3.14), shown in Figure 27-3.

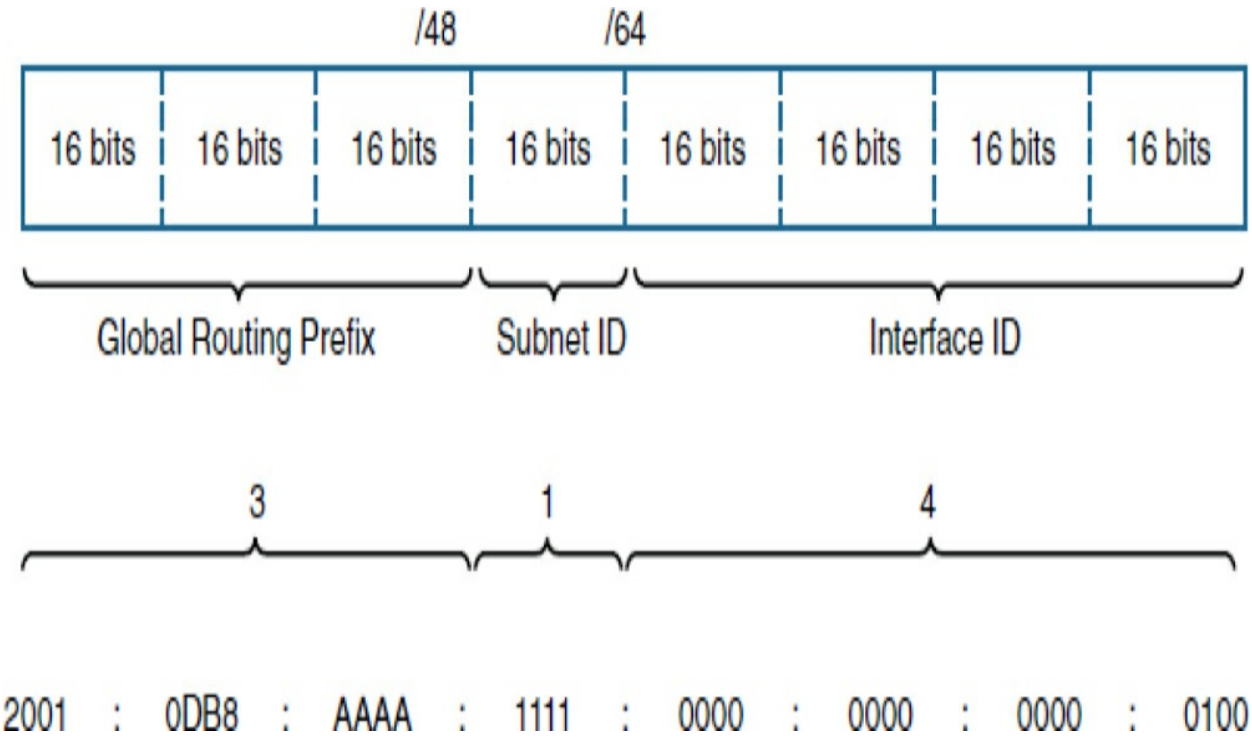

Figure 27-3 *Graziani's 3-1-4 Rule for Remembering the Global Unicast Address Structure*

Each number refers to the number of hextets, or 16-bit segments, of that portion of the address:

• **3:** Three hextets for the global routing prefix

• **1:** One hextet for the subnet ID

• **4:** Four hextets for the interface ID

Global unicast addresses that are currently assigned by the Internet Assigned Numbers Authority (IANA) use the range of addresses that start with binary value 001 (2000::/3). This range represents one-eighth of the total IPv6

Technet24

address space and is the largest block of assigned addresses. Figure 27-4 shows how the IPv6 address space is divided into an eight-piece pie based on the value of the first 3 bits.

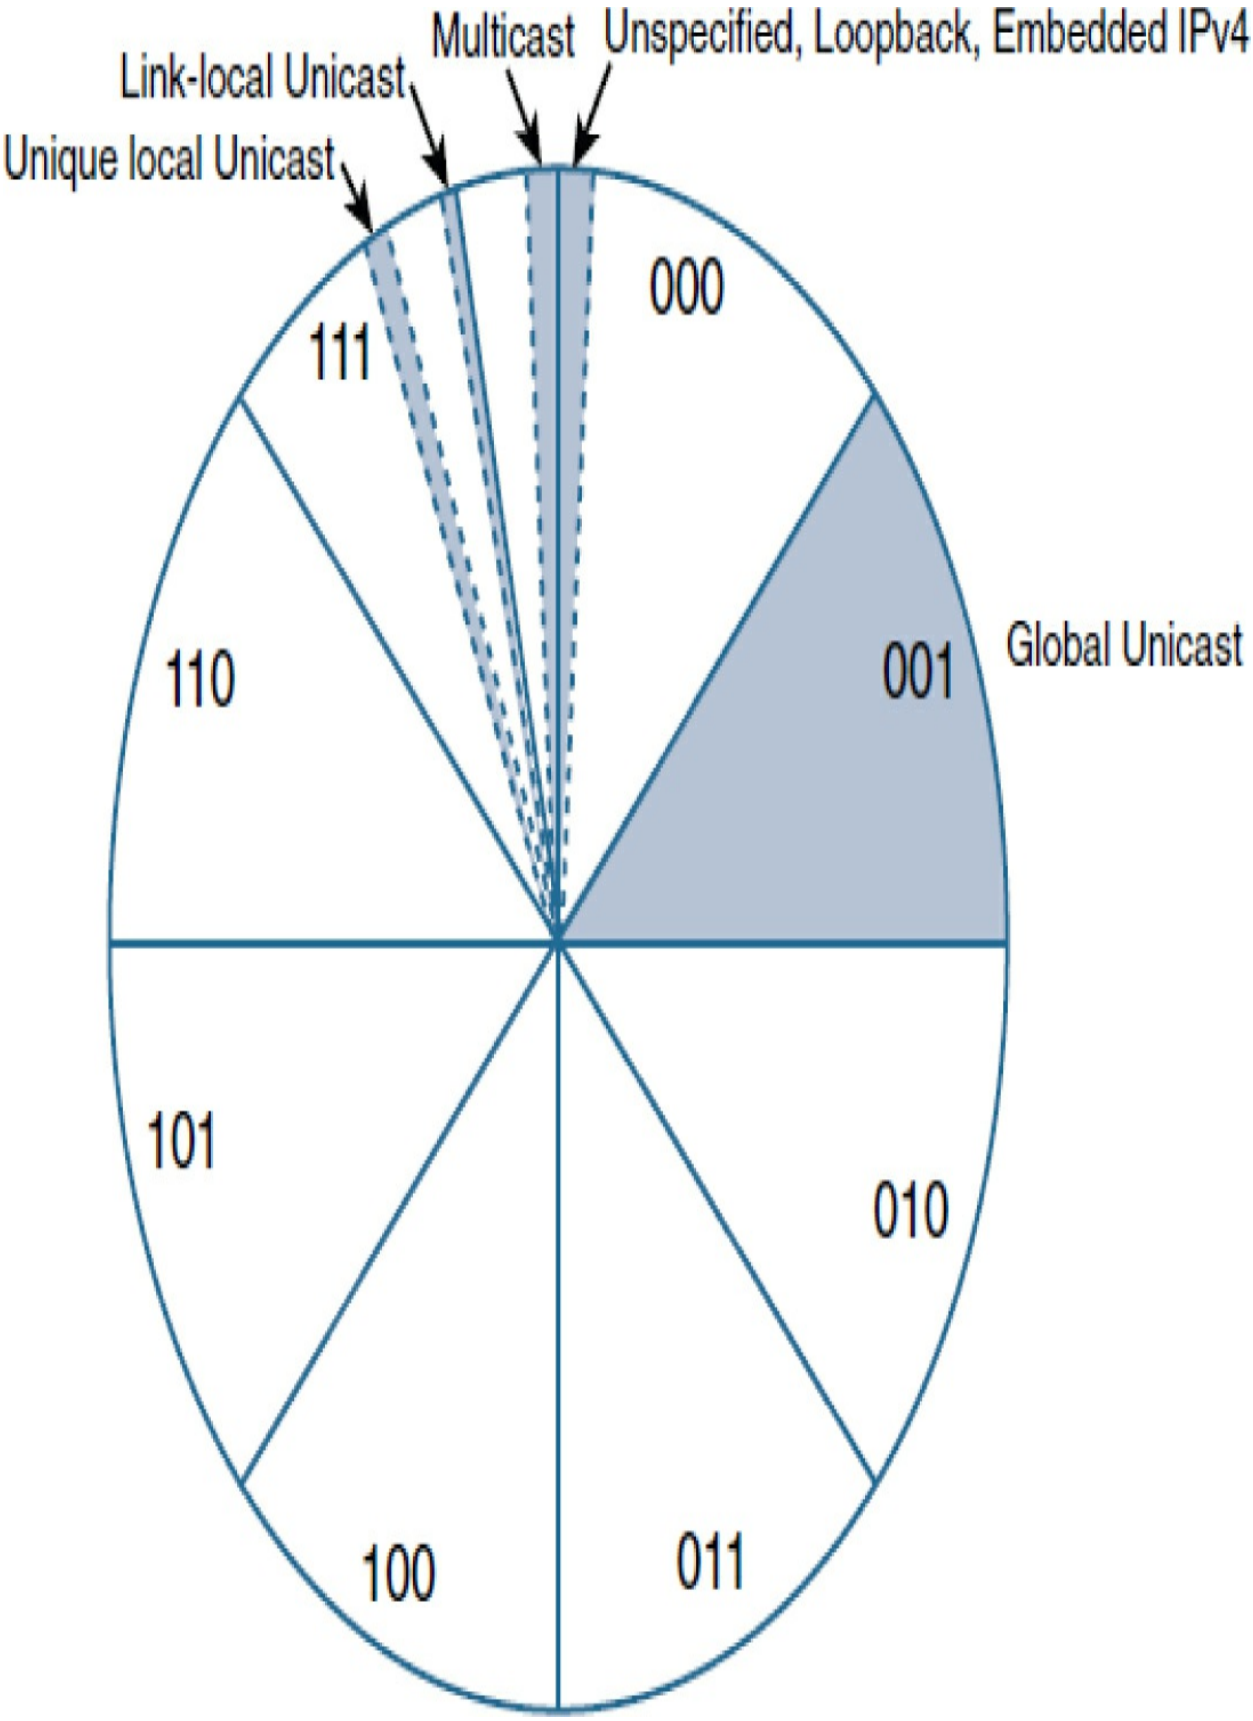
Multicast
Unspecified, Loopback, Embedded IPv4
Link-local Unicast
Unique local Unicast
000
111
001
Global Unicast
110
101
010
100
011

Technet24

Figure 27-4 *Allocation of IPv6 Address Space*

Using the 2000::/3 pie piece, the IANA assigns /23 or shorter address blocks to the five Regional Internet Registries (RIRs). From there, ISPs are assigned /32 or shorter address blocks. ISPs then assign each site—that is, each customer—a /48 or shorter address block. Figure 27-5 shows the breakdown of global routing prefixes.

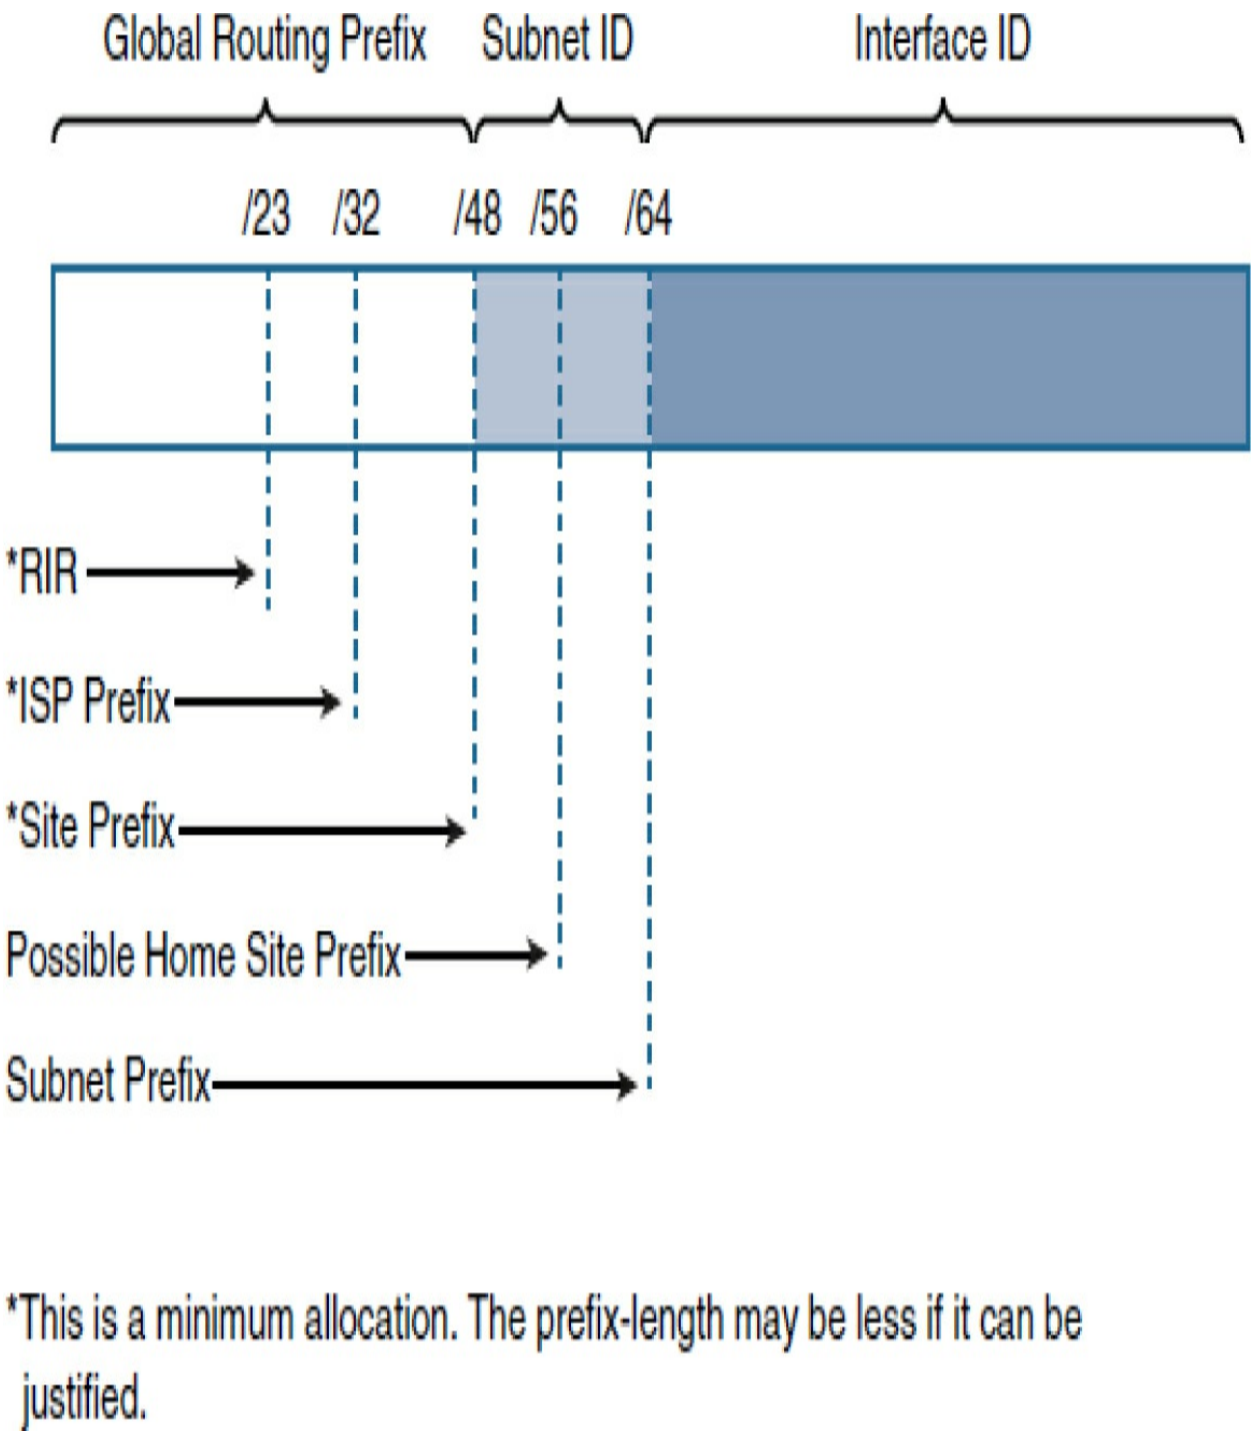

Figure 27-5 *Classification of Global Routing Prefix Sizes*

In IPv6, an interface can be configured with multiple global unicast addresses, which can be on the same or different subnets. In addition, an

Technet24

interface does not have to be configured with a global unicast address, but it must at least have a link-local address.

A global unicast address can be further classified into the various configuration options available, as Figure 27-6 shows.

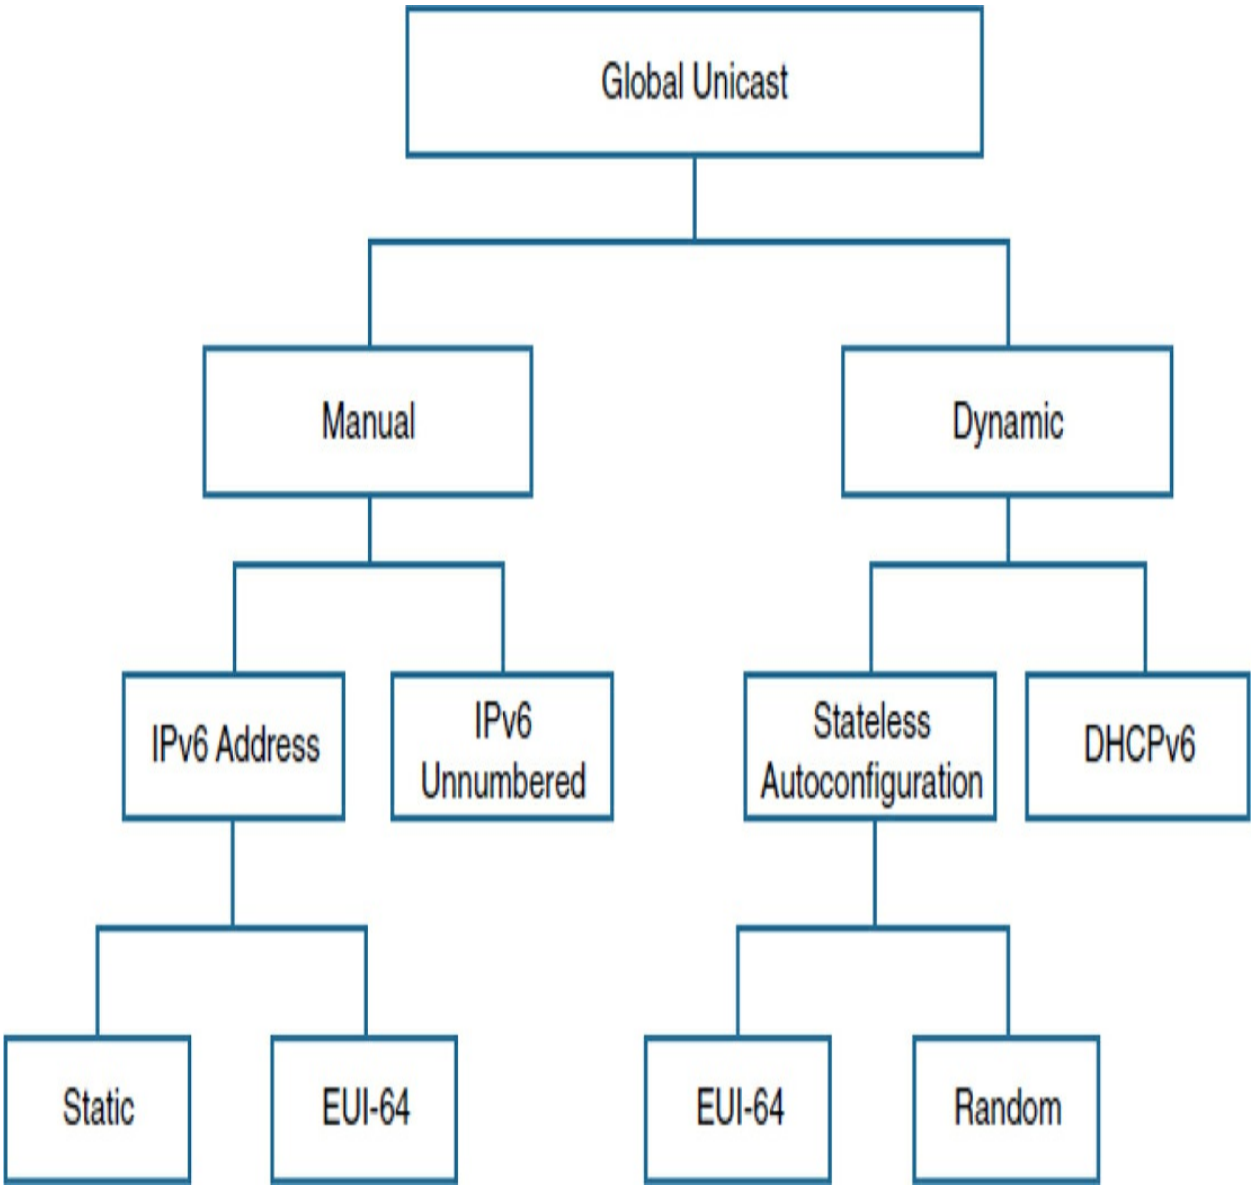

Figure 27-6 *Global Unicast Address Configuration Options*

We review EUI-64 and stateless address autoconfiguration in more detail later in this day. In upcoming days, we review the rest of the configuration options in Figure 27-6 in more detail. For now, Table 27-2 summarizes them.

Table 27-2 *Summary of Global Unicast Configuration Options*

Technet24

Global Unicast	Configuration Option	Description
Manual	Static	Much as with IPv4, the IPv6 address and prefix are statically configured on the interface.
	EUI-64	The prefix is configured manually. The EUI-64 process uses the MAC address to generate the 64-bit interface ID.
	IPv6 unnumbered	Much as with IPv4, an interface can be configured to use the IPv6 address of another interface on the same device.
Dynamic	Stateless address autoconfiguration	SLAAC determines the prefix and prefix length from neighbor discovery router advertisement messages and then creates the interface ID using the EUI-64 method.
	DHCPv6	Much as with IPv4, a device can receive some or all of its addressing from a DHCPv6 server.

Link-Local Address

As Figure 27-2 shows, link-local addresses are a type of unicast address. Link-local addresses are confined to a single link. They need to be unique only to that link because packets with a link-local source or destination address are not routable off the link.

Link-local addresses are configured in one of three ways:

• Dynamically, using EUI-64

• Using a randomly generated interface ID

• Statically, entering the link-local address manually

Link-local addresses provide a unique benefit in IPv6. A device can create its link-local address completely on its own. Link-local unicast addresses are in the range FE80::/10 to FEBF::/10, as Table 27-3 shows.

Table 27-3 *Range of Link-Local Unicast Addresses*

Technet24

Link-Local Unicast Address	Range of First Hextet	Range of First Hextet in Binary
FE80::/10	FE80	1111 1110 10 00 0000
	FEBF	1111 1110 10 11 1111

Figure 27-7 shows the format of a link-local unicast address.

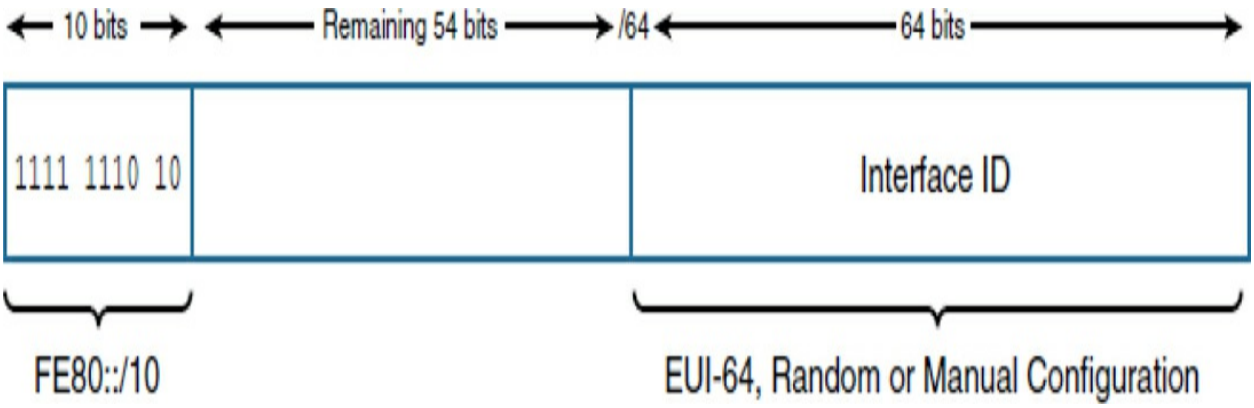

Figure 27-7 *Link-Local Unicast Address*

Loopback Address

The loopback address for IPv6 is an all-0s address except for the last bit, which is set to 1. As in IPv4, an end device uses the IPv6 loopback address to send an IPv6 packet to itself to test the TCP/IP stack. The loopback address cannot be assigned to an interface and is not routable outside the device.

Unspecified Address

The unspecified unicast address is the all-0s address, represented as ::. It cannot be assigned to an interface but is reserved for communications when the sending device does not have a valid IPv6 address yet. For example, a device uses :: as the source address when using the duplicate address detection (DAD) process. The DAD process ensures a unique link-local address. Before a device can begin using its newly created link-local address, it sends out an all-nodes multicast to all devices on the link, with its new address as the destination. If the device receives a response, it knows that link-local address is in use and, therefore, needs to create another link-local address.

Unique Local Address

Unique local addresses (ULA) are defined by RFC 4193, “Unique Local IPv6 Unicast Addresses.” Figure 27-8 shows the format for ULAs.

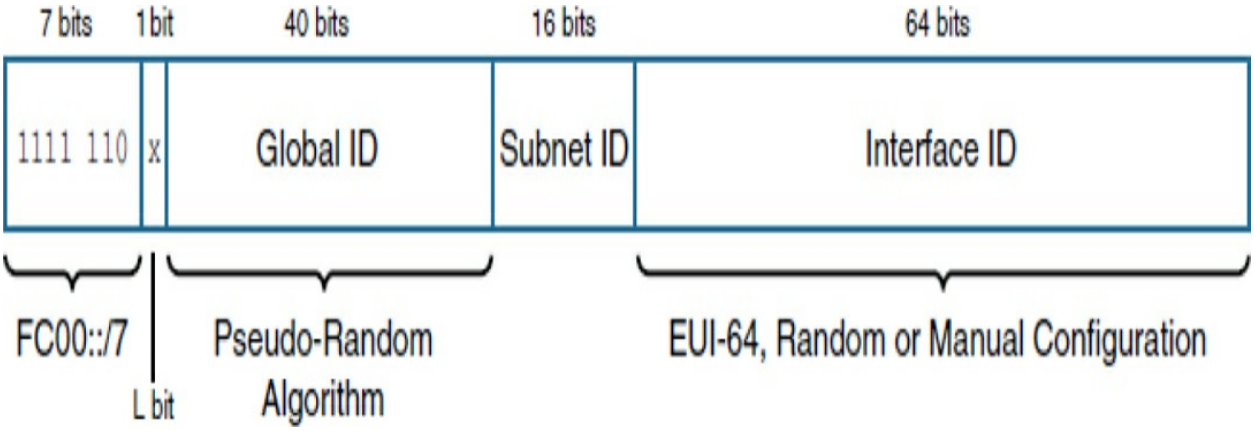

Figure 27-8 *Unique Local Address*

These are private addresses. However, unlike in IPv4, IPv6 ULAs are globally unique. This is possible because of the relatively large amount of

Technet24

address space in the Global ID portion shown in Figure 27-8: 40 bits, or more than 1 trillion unique global IDs. As long as a site uses the pseudo-random global ID algorithm, it will have a very high probability of generating a unique global ID.

Unique local addresses have the following characteristics:

• Possess a globally unique prefix or at least have a very high probability of being unique

• Allow sites to be combined or privately interconnected without address conflicts or addressing renumbering

• Remain independent of any Internet service provider and can be used within a site without having Internet connectivity

• If accidentally leaked outside a site by either routing or the Domain Name System (DNS), don’t cause a conflict with any other addresses

• Can be used just like a global unicast address

IPv4 Embedded Address

IPv4 and IPv6 packets are not compatible. Features such as NAT-PT (now deprecated) and NAT64 are required to translate between the two address families. IPv4-mapped IPv6 addresses are used by transition mechanisms on hosts and routers to create IPv4 tunnels that deliver IPv6 packets over IPv4

networks.

Note

NAT64 is beyond the scope of the CCNA exam topics.

To create an IPv4-mapped IPv6 address, the IPv4 address is embedded within the low-order 32 bits of IPv6. Basically, IPv6 just puts an IPv4 address at the end, adds 16 all-1 bits, and pads the rest of the address. The address does not have to be globally unique. Figure 27-9 illustrates this IPv4-mapped IPv6 address structure.

Technet24

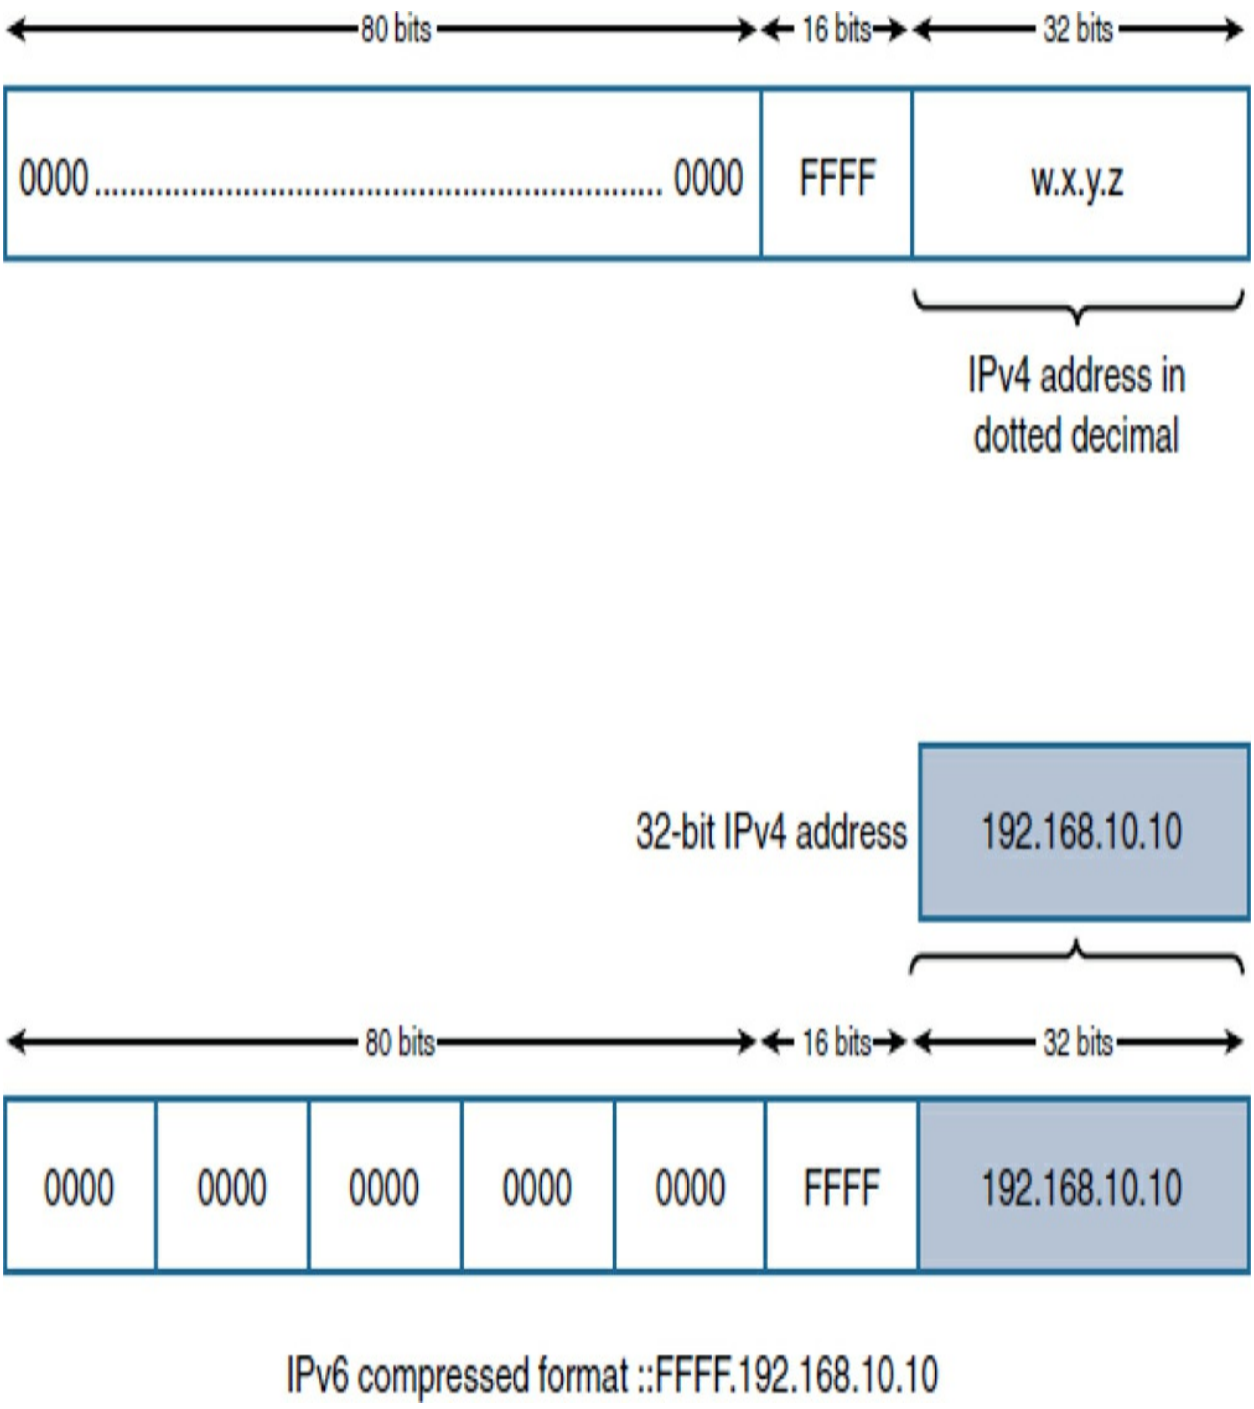

Figure 27-9 *IPv4-Mapped IPv6 Address*

Multicast

The second major classification of IPv6 address types in Figure 27-2 is

multicast. Multicast is a technique by which a device sends a single packet to multiple destinations simultaneously. An IPv6 multicast address defines a group of devices known as a *multicast group* and is equivalent to IPv4 224.0.0.0/4. IPv6 multicast addresses have the prefix FF00::/8.

Two types of IPv6 multicast addresses are used:

• Assigned multicast

• Solicited-node multicast

Assigned Multicast

Assigned multicast addresses are used in context with specific protocols.

Two common IPv6 assigned multicast groups include the following:

• **FF02::1 All-nodes multicast group:** This is a multicast group that all IPv6-enabled devices join. As with a broadcast in IPv4, all IPv6 interfaces on the link process packets sent to this address. For example, a router sending an ICMPv6 Router Advertisement (RA) uses the all-nodes FF02::1 address. IPv6-enabled devices can then use the RA information to learn the link's address information, such as prefix, prefix length, and default gateway.

• **FF02::2 All-routers multicast group:** This is a multicast group that all IPv6 routers join. A router becomes a member of this group when it is enabled as an IPv6 router with the **ipv6 unicast-routing** global configuration

Technet24

command. A packet sent to this group is received and processed by all IPv6 routers on the link or network. For example, IPv6-enabled devices send ICMPv6 Router Solicitation (RS) messages to the all-routers multicast address requesting an RA message.

Solicited-Node Multicast

In addition to every unicast address assigned to an interface, a device has a special multicast address known as a solicited-node multicast address (refer to Figure 27-2). These multicast addresses are automatically created using a special mapping of the device's unicast address with the solicited-node multicast prefix FF02:0:0:0:0:1:FF00::/104.

As Figure 27-10 shows, solicited-node multicast addresses are used for two essential IPv6 mechanisms, both part of Neighbor Discovery Protocol (NDP):

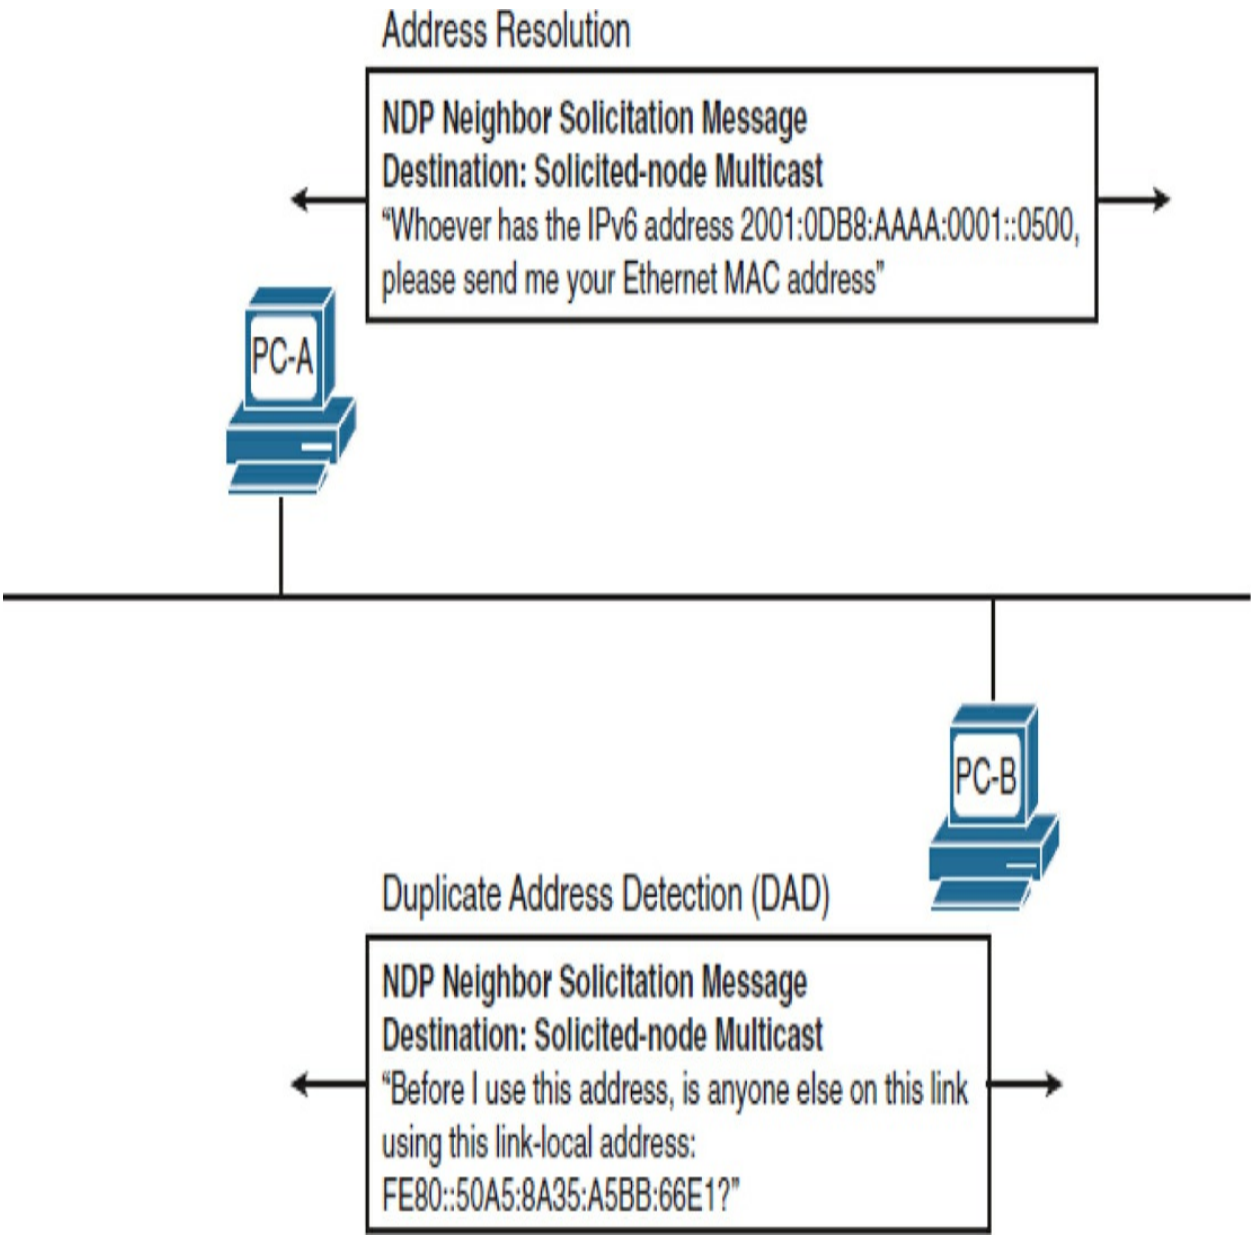

Figure 27-10 *Uses of Solicited-Node Multicasts*

• **Address resolution:** In this mechanism, which is equivalent to ARP in IPv4, an IPv6 device sends an NS message to a solicited-node multicast address to learn the link layer address of a device on the same link. The device recognizes the IPv6 address of the destination on that link but needs to know its data link address.

Technet24

• **Duplicate address detection (DAD):** As mentioned earlier, DAD allows a device to verify that its unicast address is unique on the link. An NS message is sent to the device's own solicited-node multicast address to determine whether anyone else has this same address.

As Figure 27-11 shows, the solicited-node multicast address consists of two parts:

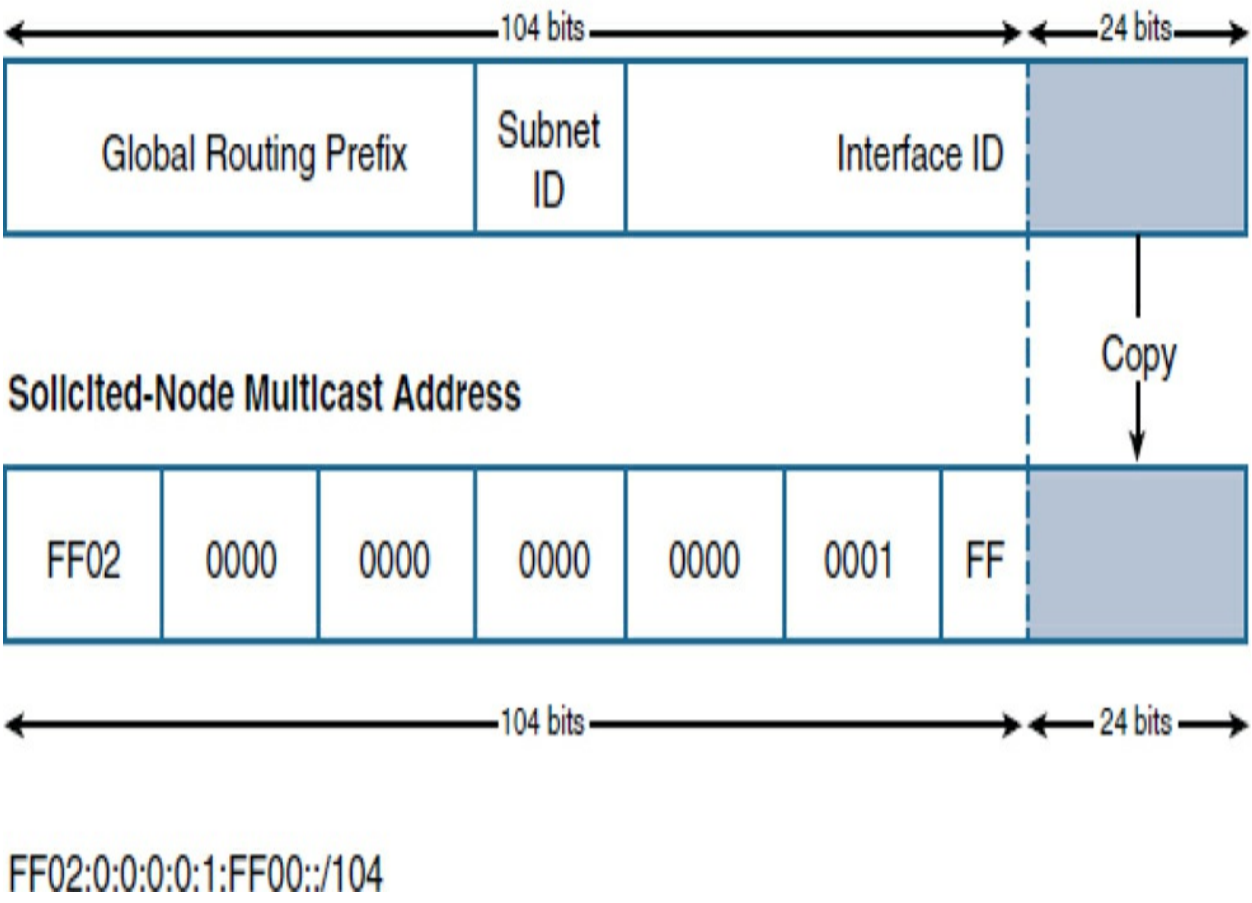

Figure 27-11 *Solicited-Node Multicast Address Structure*

• **FF02:0:0:0:0:FF00::/104 multicast prefix:** This is the first 104 bits of the

all solicited-node multicast address.

• **Least significant 24 bits:** These bits are copied from the far-right 24 bits of the global unicast or link-local unicast address of the device.

Anycast

The last major classification of IPv6 address types in Figure 27-2 is the anycast address. An anycast address can be assigned to more than one device or interface. A packet sent to an anycast address is routed to the “nearest” device that is configured with the anycast address, as Figure 27-12 shows.

Technet24

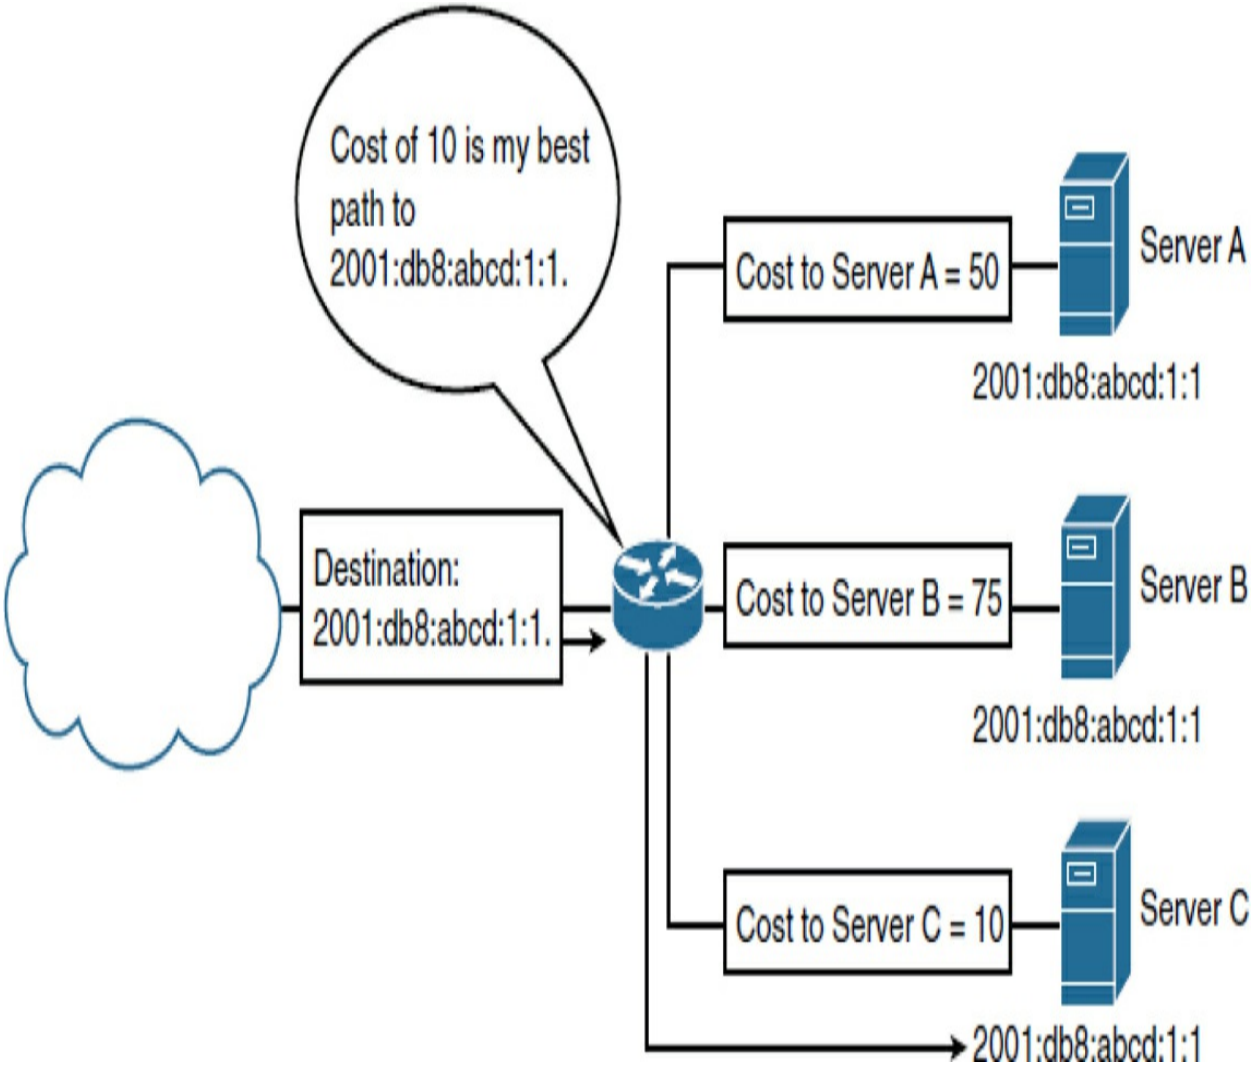

Figure 27-12 *Example of Anycast Addressing*

Representing the IPv6 Address

An IPv6 address can look rather intimidating to someone who is used to IPv4 addressing. However, an IPv6 address can be easier to read and is much simpler to subnet than IPv4.

Conventions for Writing IPv6 Addresses

IPv6 conventions use 32 hexadecimal numbers, organized into eight hextets of four hex digits separated by colons, to represent a 128-bit IPv6 address. For example:

2340:1111:AAAA:0001:1234:5678:9ABC

To make things a little easier, two rules allow you to shorten what must be configured for an IPv6 address:

• **Rule 1:** Omit the leading 0s in any given hextet.

• **Rule 2:** Omit the all-0s hextets. Represent one or more consecutive hextets of all hex 0s with a double colon (::), but only for one such occurrence in a given address.

For example, in the following address, the highlighted hex digits represent the portion of the address that can be abbreviated:

```
FE00:0000:0000:0001:0000:0000:0000:0056
```

This address has two locations in which one or more hextets have four hex 0s, so two main options work for abbreviating this address with the :: abbreviation in one of the locations. The following two options show the two briefest valid abbreviations:

• FE00::1:0:0:0:56

Technet24

• FE00:0:0:1::56

In the first example, the second and third hextets preceding 0001 were replaced with ::. In the second example, the fifth, sixth, and seventh hextets were replaced with ::. In particular, note that the :: abbreviation, meaning "one or more hextets of all 0s," cannot be used twice because that would be ambiguous. Therefore, the abbreviation FE00::1::56 would not be valid.

Conventions for Writing IPv6 Prefixes

An IPv6 prefix represents a range or block of consecutive IPv6 addresses. The number that represents the range of addresses, called a *prefix*, is usually seen in IP routing tables, just as you see IP subnet numbers in IPv4 routing tables.

As with IPv4, when writing or typing a prefix in IPv6, the bits past the end of the prefix length are all binary 0s. The following IPv6 address is an example of an address assigned to a host:

2000:1234:5678:9ABC:1234:5678:9ABC:1111/64

The prefix in which this address resides is as follows:

2000:1234:5678:9ABC:**0000:0000:0000:0000**/64

When abbreviated, this is:

2000:1234:5678:9ABC::/64

If the prefix length does not fall on a hextet boundary (that is, is not a multiple of 16), the prefix value should list all the values in the last hextet. For example, assume that the prefix length in the previous example is /56. By convention, the rest of the fourth hextet is written, after being set to binary 0s, as follows:

2000:1234:5678:9A**00**::/56

The following list summarizes some key points about how to write IPv6 prefixes:

• The prefix has the same value as the IP addresses in the group for the first number of bits, as defined by the prefix length.

• Any bits after the prefix length number of bits are binary 0s.

• The prefix can be abbreviated with the same rules as for IPv6 addresses.

• If the prefix length is not on a hextet boundary, write down the value for the entire hextet.

Table 27-4 shows several sample prefixes, their formats, and a brief explanation.

Table 27-4 *Example IPv6 Prefixes and Their Meanings*

Technet24

Prefix	Explanation	Incorrect Alternative
2000::/3	All addresses whose first 3 bits are equal to the first 3 bits of hex number 2000 (bits are 001)	2000/3 (omits ::) 2::/3 (omits the rest of the first hextet)
2340:1140::/26	All addresses whose first 26 bits match the listed hex number	2340:114::/26 (omits the last digit in the second hextet)
2340:1111::/32	All addresses whose first 32 bits match the listed hex number	2340:1111/32 (omits ::)

IPv6 Subnetting

In many ways, subnetting IPv6 addresses is much simpler than subnetting IPv4 addresses. A typical site is assigned an IPv6 address space with a /48 prefix length. Because the least significant bits are used for the interface ID, that leaves 16 bits for the subnet ID and a /64 subnet prefix length, as Figure 27-13 shows.

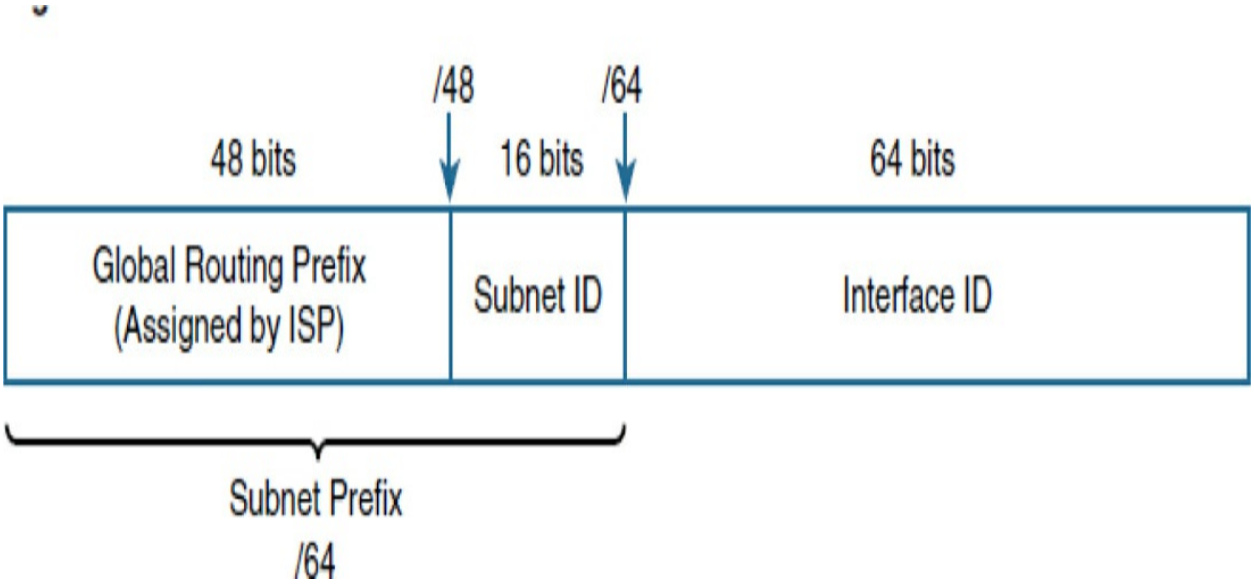

Figure 27-13 */64 Subnet Prefix*

For our subnetting examples, we use 2001:0DB8:000A::/48, or simply 2001:DB8:A::/48, which includes subnets 2001:DB8:A::/64 through 2001:DB8:A:FFFF::/64. That's 2^{16}, or 65,536 subnets, each with 2^{64}, or 18 quintillion, interface addresses.

Subnetting the Subnet ID

To subnet in a small to medium-size business, simply increment the least significant bits of the subnet ID (as in Example 27-1) and assign /64 subnets to your networks.

Example 27-1 *Subnetting the Subnet ID*

```
2001:DB8:A:0001::/64
2001:DB8:A:0002::/64
2001:DB8:A:0003::/64
```

Technet24

```
2001:DB8:A:0004::/64
2001:DB8:A:0005::/64
```

Of course, if you are administering a larger implementation, you can use the four hexadecimal digits of the subnet ID to design a quick and simple four-level hierarchy. Most large enterprise networks have plenty of room to design a logical address scheme that aggregates addresses for an optimal routing configuration. In addition, applying for and receiving another /48 address is not difficult.

Subnetting into the Interface ID

If you extend your subnetting into the interface ID portion of the address, it is a best practice to subnet on the nibble boundary. A nibble is 4 bits, or one hexadecimal digit. For example, let's borrow the first 4 bits from the interface ID portion of the network address 2001:DB8:A:1::/64. That means the network 2001:DB8:A:1::/64 would now have 2^4, or 16, subnets from 2001:DB8:A:1:0000::/68 to 2001:DB8:A:1:F000::/68. Listing the subnets is easy, as Example 27-2 shows.

Example 27-2 *Subnetting into the Interface ID*

```
2001:DB8:A:1:0000::/68
2001:DB8:A:1:1000::/68
2001:DB8:A:1:2000::/68
2001:DB8:A:1:3000::/68
```

```
        thru
2001:DB8:A:1:F000::/68
```

EUI-64 Concept

Day 18 reviews static IPv6 addressing, including how to configure a router to use EUI-64 addressing (EUI stands for Extended Unique Identifier). Today we are reviewing the concept behind the EUI-64 configuration.

Recall from Figure 27-13 that the second half of the IPv6 address is called the *interface ID*. The value of the interface ID portion of a global unicast address can be set to any value, as long as no other host in the same subnet attempts to use the same value. However, the size of the interface ID was chosen to allow easy autoconfiguration of IP addresses by plugging the MAC address of a network card into the interface ID field in an IPv6 address.

MAC addresses are 6 bytes (48 bits) in length. To complete the 64-bit interface ID, IPv6 fills in 2 more bytes by separating the MAC address into two 3-byte halves. It then inserts hex FFFE between the halves and sets the seventh bit in the first byte to binary 1 to form the interface ID field. Figure 27-14 shows this format, called the EUI-64 format.

For example, the following two lines list a host's MAC address and corresponding EUI-64 format interface ID, assuming the use of an address configuration option that uses the EUI-64 format:

Technet24

• **MAC address:** 0034:5678:9ABC

• **EUI-64 interface ID:** 0234:56FF:FE78:9ABC

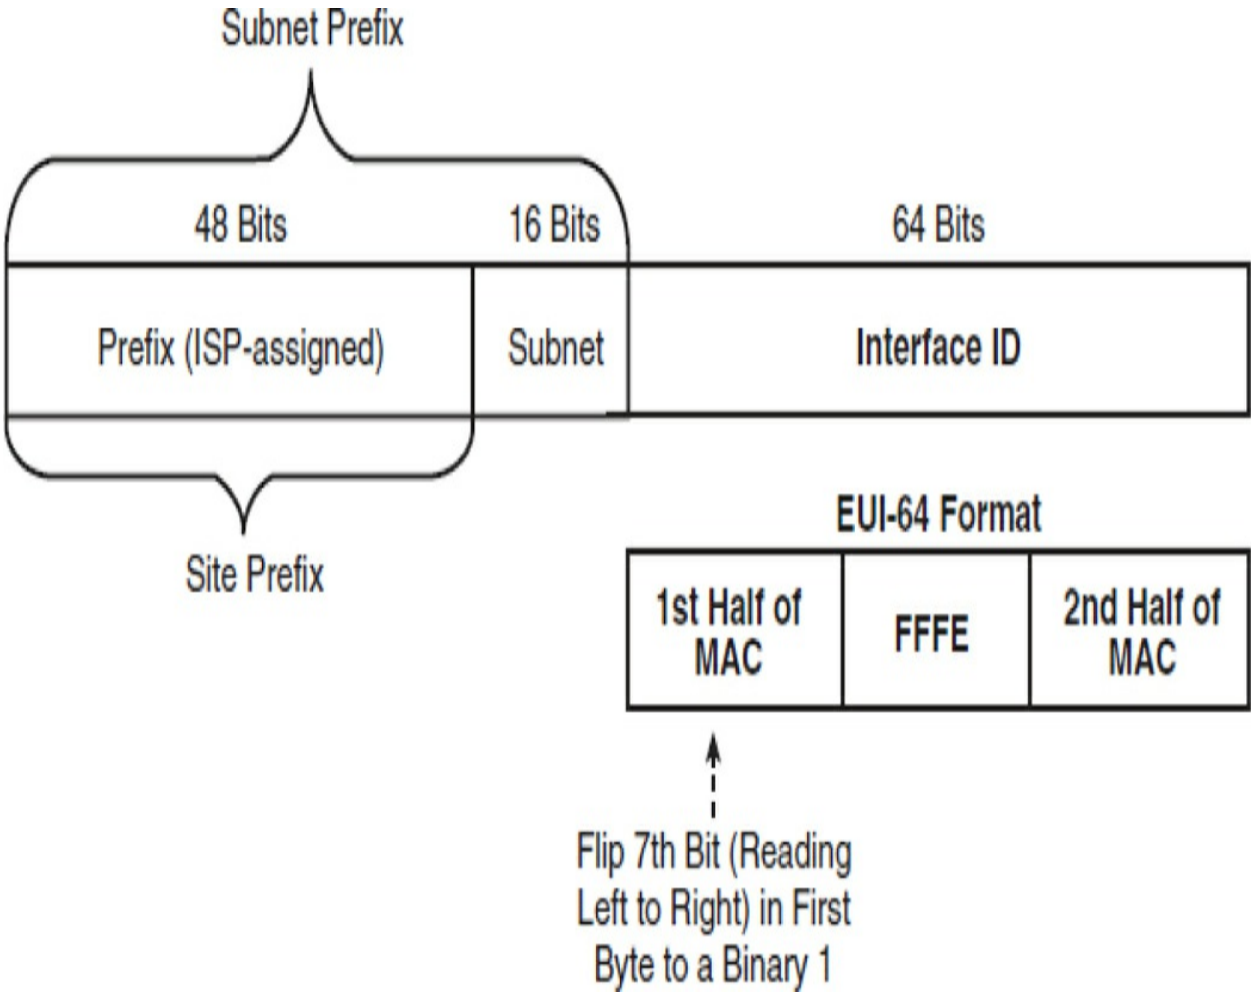

Figure 27-14 *IPv6 Address Format with Interface ID and EUI-64*

Note

To change the seventh bit (reading left to right) in the example, convert hex 00 to binary 00000000, change the seventh bit to 1 (00000010), and then convert back to hex, for

hex 02 as the first two digits.

Stateless Address Autoconfiguration

IPv6 supports two methods of dynamic configuration of IPv6 addresses:

• **Stateless address autoconfiguration (SLAAC):** A host dynamically learns the /64 prefix through the IPv6 Neighbor Discovery Protocol (NDP) and then calculates the rest of its address by using the EUI-64 method.

• **DHCPv6:** This works the same conceptually as DHCP in IPv4. We review DHCPv6 on Day 23, "DHCP and DNS."

By using the EUI-64 process and Neighbor Discovery Protocol (NDP), SLAAC allows a device to determine its entire global unicast address without any manual configuration and without a DHCPv6 server. Figure 27-15 illustrates the SLAAC process between a host and a router configured with the **ipv6 unicast-routing** command, which means it will send and receive NDP messages.

Technet24

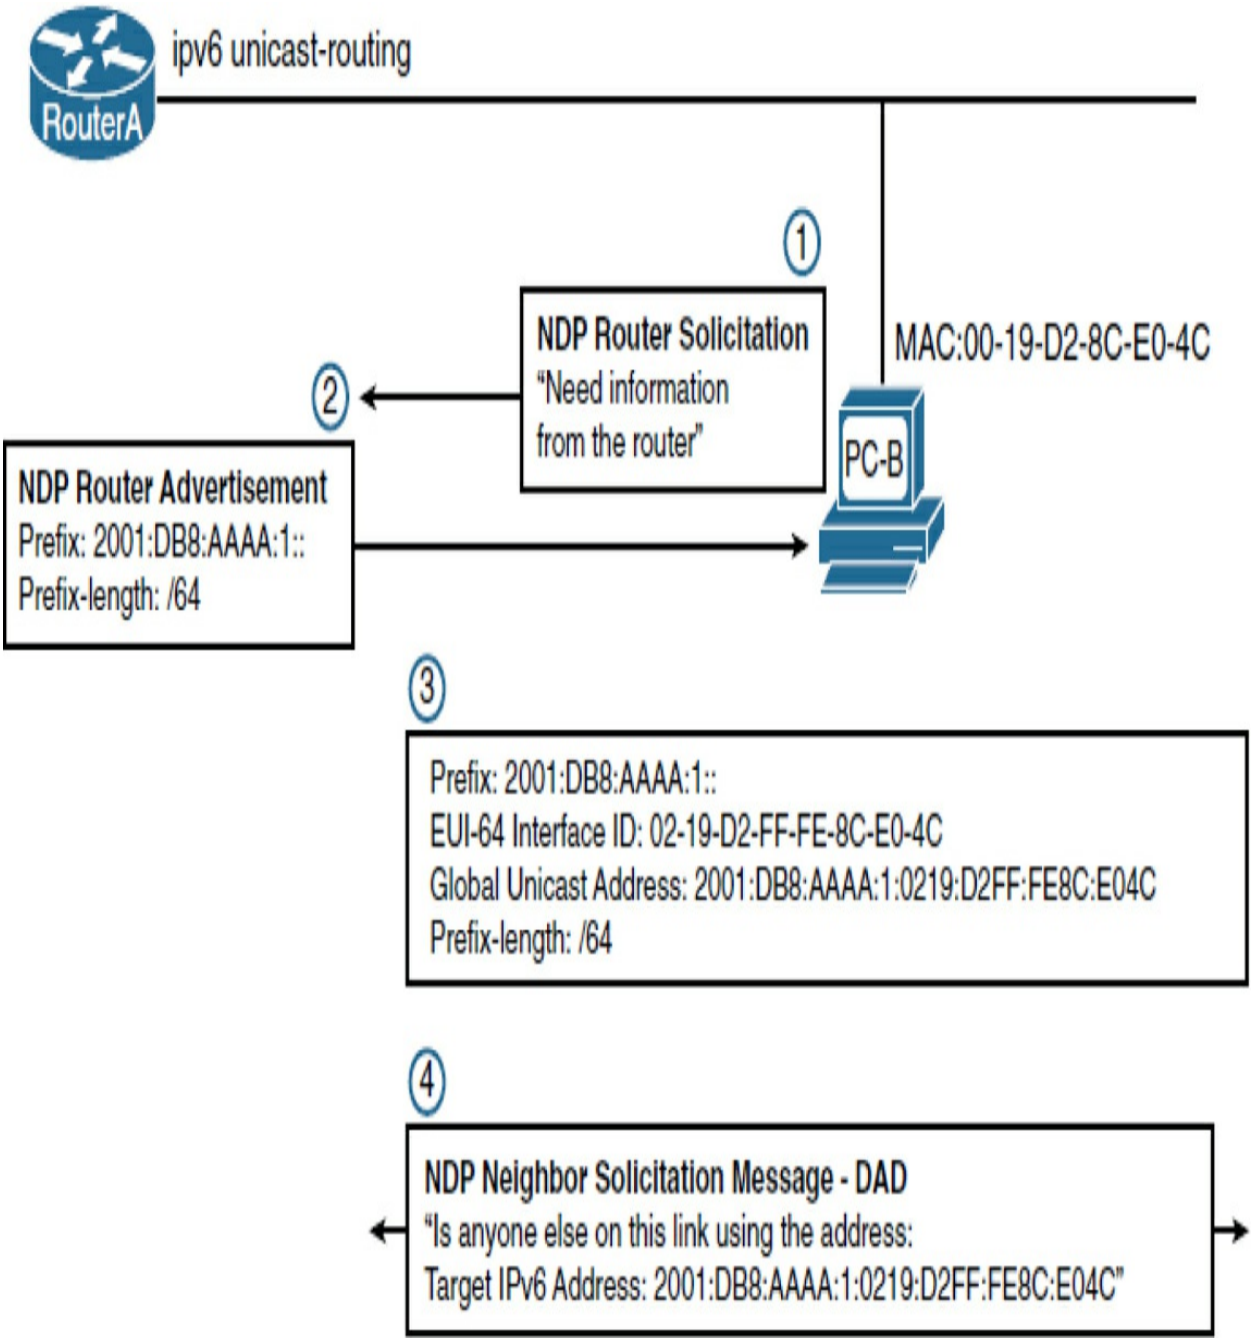

Figure 27-15 *Neighbor Discovery and the SLAAC Process*

Migration to IPv6

Two major transition strategies are currently used to migrate to IPv6:

• **Dual-stacking:** In this integration method, a node has implementation and connectivity to both an IPv4 network and an IPv6 network. This is the recommended option and involves running IPv4 and IPv6 at the same time.

• **Tunneling:** Tunneling is a method for transporting IPv6 packets over IPv4-only networks by encapsulating the IPv6 packet inside IPv4. Several tunneling techniques are available.

Because of the simplicity of running dual-stacking, it will most likely be the preferred strategy as IPv4-only networks begin to disappear. But it will probably still be decades before we see enterprise networks running exclusively IPv6. Figure 27-16 illustrates one way Wendell Odom thinks about the transition to IPv6: “But who knows how long it will take?”

Remember this advice: “Dual-stack where you can; tunnel where you must.” These two methods are the most common techniques to transition from IPv4 to IPv6. Dual-stacking is easy enough: Just configure all your devices to use both IPv4 and IPv6 addressing. Tunneling is more complex and beyond the scope of the CCNA exam topics.

Technet24

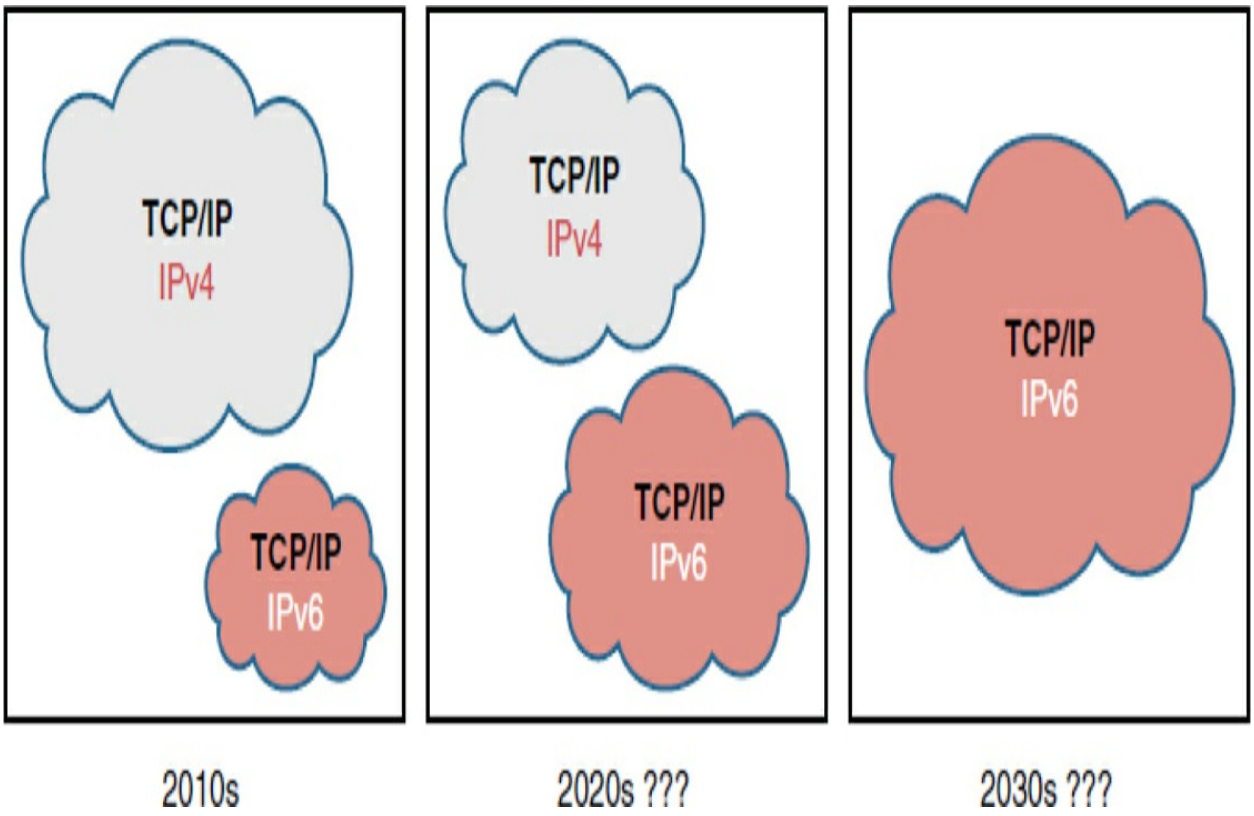

Figure 27-16 *Transition to IPv6 Using Dual-Stacking*

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	4
	6
	7
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	25–28
CCNA Portable Command Guide	5

Technet24

Day 26. VLAN and Trunking Concepts and Configurations

CCNA 200-301 Exam Topics

• 2.1 Configure and verify VLANs (normal range) spanning multiple switches

• 2.2 Configure and verify interswitch connectivity

Key Points

Most large networks today implement virtual local-area networks (VLANs). Without VLANs, a switch considers every port to be in the same broadcast domain. With VLANs, switch ports can be grouped into different VLANs, essentially segmenting the broadcast domain. Today we review VLAN concepts, consider traffic types, discuss VLAN types, and review the concept of trunking, including Dynamic Trunking Protocol (DTP). Then we review the commands to configure and verify VLANs, trunking, and inter-VLAN routing.

VLAN Concepts

Although a switch comes out of the box with only one VLAN, normally a switch is configured to have two or more VLANs. With such a switch, you can create multiple broadcast domains by putting some interfaces into one

VLAN and other interfaces into other VLANs.

Consider these reasons for using VLANs:

• Grouping users by department instead of by physical location

• Segmenting devices into smaller LANs to reduce processing overhead for all devices on the LAN

• Reducing the workload of STP by limiting a VLAN to a single access switch

• Enforcing better security by isolating sensitive data to separate VLANs

• Separating IP voice traffic from data traffic

• Assisting troubleshooting by reducing the size of the failure domain (that is, the number of devices that can cause a failure or that can be affected by one)

Benefits of using VLANs include the following:

• **Security:** Sensitive data can be isolated to one VLAN, separated from the rest of the network.

• **Cost reduction:** Reduced need for expensive network upgrades and more efficient use of existing bandwidth and uplinks lead to cost savings.

• **Higher performance:** Dividing flat Layer 2 networks into multiple logical

Technet24

broadcast domains reduces unnecessary traffic on the network and boosts performance.

• **Problem isolation:** To solve problems more quickly by confining the failure domain to the same set of devices within the same broadcast domain.

STP efficiency: To reduce the workload for Spanning Tree Protocol (STP) by limiting a VLAN to a single access switch.

Traffic Types

The key to successful VLAN deployment is understanding the traffic patterns and the various traffic types in the organization. Table 26-1 lists the common types of network traffic to evaluate before placing devices and configuring VLANs.

Table 26-1 *Traffic Types*

Traffic Type	Description
Network management	Many types of network management traffic can be present on the network. To make network troubleshooting easier, some designers assign a separate VLAN to carry certain types of network management traffic.
IP telephony	Two types of IP telephony traffic exist: signaling information between end devices and the data packets of the voice conversation. Designers often configure the data to and from the IP phones on a separate VLAN designated for voice traffic so that they can apply quality-of-service measures to give high priority to voice traffic.
IP multicast	Multicast traffic can produce a large amount of data streaming across the network. Switches must be configured to keep this traffic from flooding to devices that have not requested it, and routers must be configured to ensure that multicast traffic is forwarded to the network areas where it is requested.
Normal data	Normal data traffic is typical application traffic that is related to file and print services, email, Internet browsing, database access, and other shared network applications.
Scavenger class	Scavenger class includes all traffic with protocols or patterns that exceed their normal data flows. Applications assigned to this class have little or no contribution to the organizational objectives of the enterprise and are typically entertainment oriented.

Technet24

Types of VLANs

Some VLAN types are defined by the type of traffic they support; others are defined by the specific functions they perform. The principal VLAN types and their descriptions follow:

• **Data VLAN:** Configured to carry only user-generated traffic, ensuring that voice and management traffic is separated from data traffic.

• **Default VLAN:** All the ports on a switch are members of the default VLAN when the switch is reset to factory defaults. The default VLAN for Cisco switches is VLAN 1. VLAN 1 has all the features of any VLAN, except that you cannot rename it, and you cannot delete it. It is a security best practice to restrict VLAN 1 to serve as a conduit only for Layer 2 control traffic (for example, CDP) and support no other traffic.

• **Black hole VLAN:** A security best practice is to define a black hole VLAN to be a dummy VLAN distinct from all other VLANs defined in the switched LAN. All unused switch ports are assigned to the black hole VLAN so that any unauthorized device connecting to an unused switch port is prevented from communicating beyond the switch to which it is connected.

• **Native VLAN:** This VLAN type serves as a common identifier on opposing ends of a trunk link. A security best practice is to define a native VLAN to be a dummy VLAN distinct from all other VLANs defined in the switched LAN. The native VLAN is not used for any traffic in the switched network

unless legacy bridging devices happen to be present in the network or a multiaccess interconnection exists between switches joined by a hub.

• **Management VLAN:** The network administrator defines this VLAN as a means to access the management capabilities of a switch. By default, VLAN 1 is the management VLAN. It is a security best practice to define the management VLAN to be a VLAN distinct from all other VLANs defined in the switched LAN. You do this by configuring and activating a new VLAN interface.

• **Voice VLANs:** A voice VLAN enables switch ports to carry IP voice traffic from an IP phone. The network administrator configures a voice VLAN and assigns it to access ports. Then when an IP phone is connected to the switch port, the switch sends CDP messages that instruct the attached IP phone to send voice traffic tagged with the voice VLAN ID.

Voice VLAN Example

Figure 26-1 shows an example of using one port on a switch to connect a user's IP phone and PC. The switch port is configured to carry data traffic on VLAN 20 and voice traffic on VLAN 150. The Cisco IP Phone contains an integrated three-port 10/100 switch to provide the following dedicated connections:

• Port 1 connects to the switch or other VoIP device.

Technet24

• Port 2 is an internal 10/100 interface that carries the IP Phone traffic.

• Port 3 (access port) connects to a PC or other device.

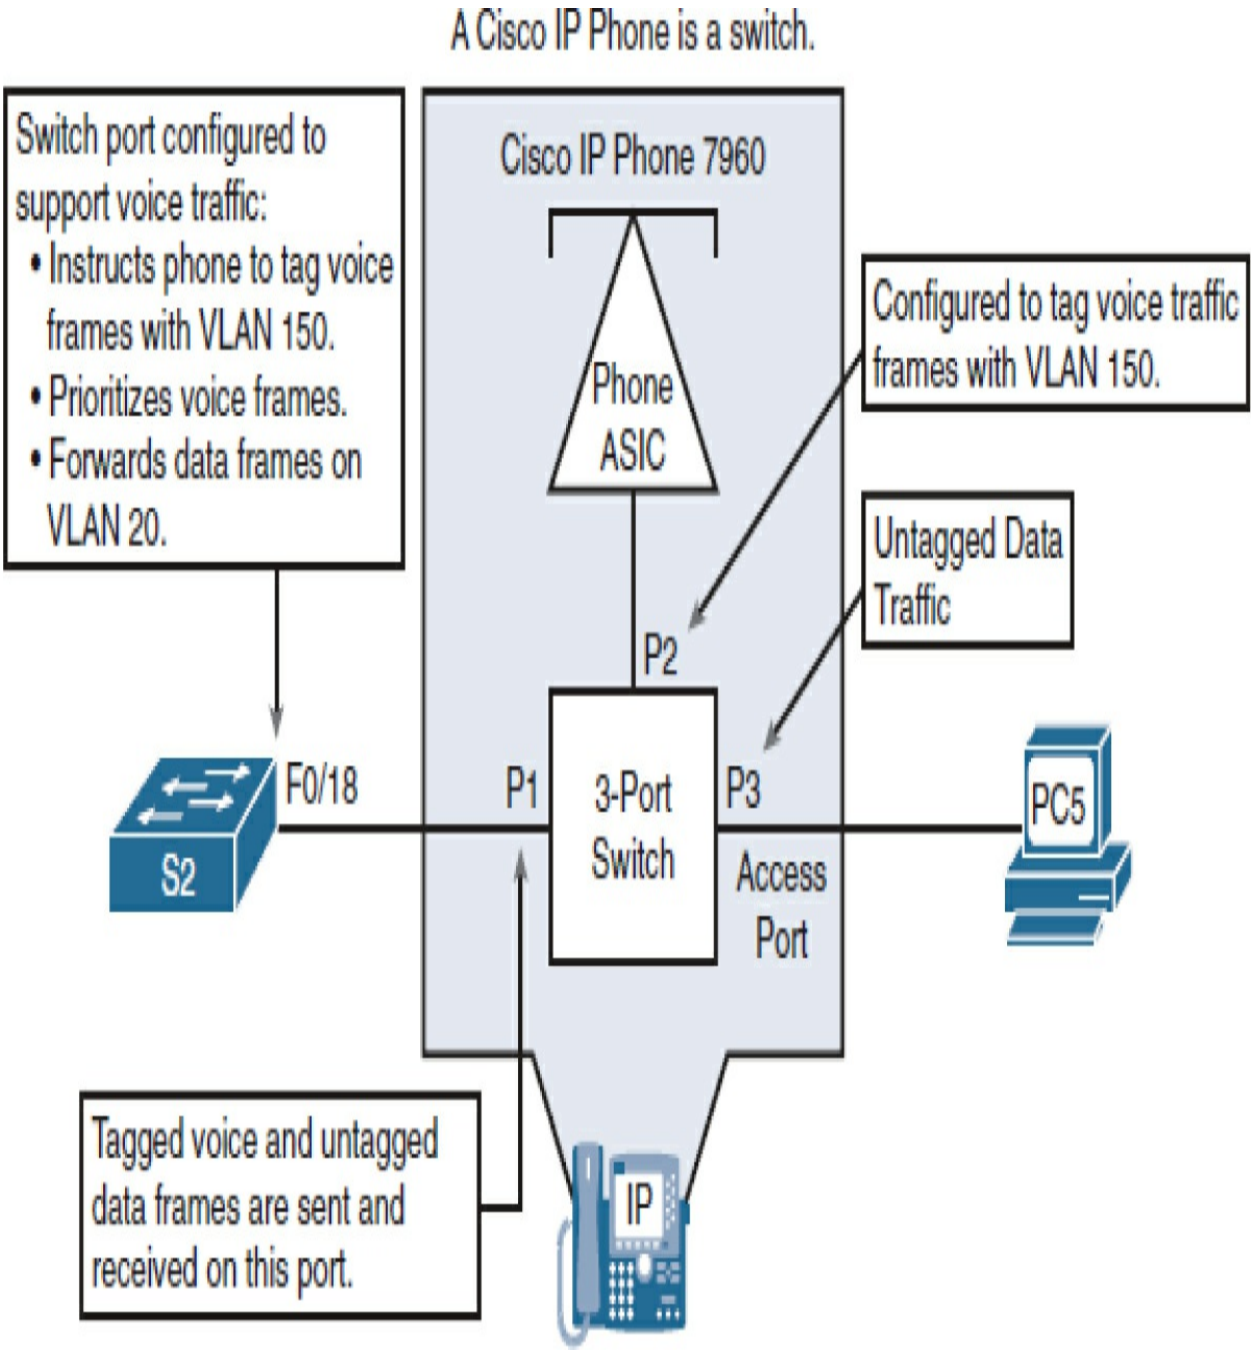

Figure 26-1 *Cisco IP Phone Switching Voice and Data Traffic*

The traffic from PC5 attached to the IP Phone passes through the IP Phone untagged. The link between S2 and the IP Phone acts as a modified trunk to carry both the tagged voice traffic and the untagged data traffic.

Trunking VLANs

A VLAN trunk is an Ethernet point-to-point link between an Ethernet switch interface and an Ethernet interface on another networking device, such as a router or a switch, carrying the traffic of multiple VLANs over the singular link. A VLAN trunk enables you to extend the VLANs across an entire network. A VLAN trunk does not belong to a specific VLAN; instead, it serves as a conduit for VLANs between switches. Figure 26-2 shows a small switched network with a trunk link between S1 and S2 carrying multiple VLAN traffic.

Technet24

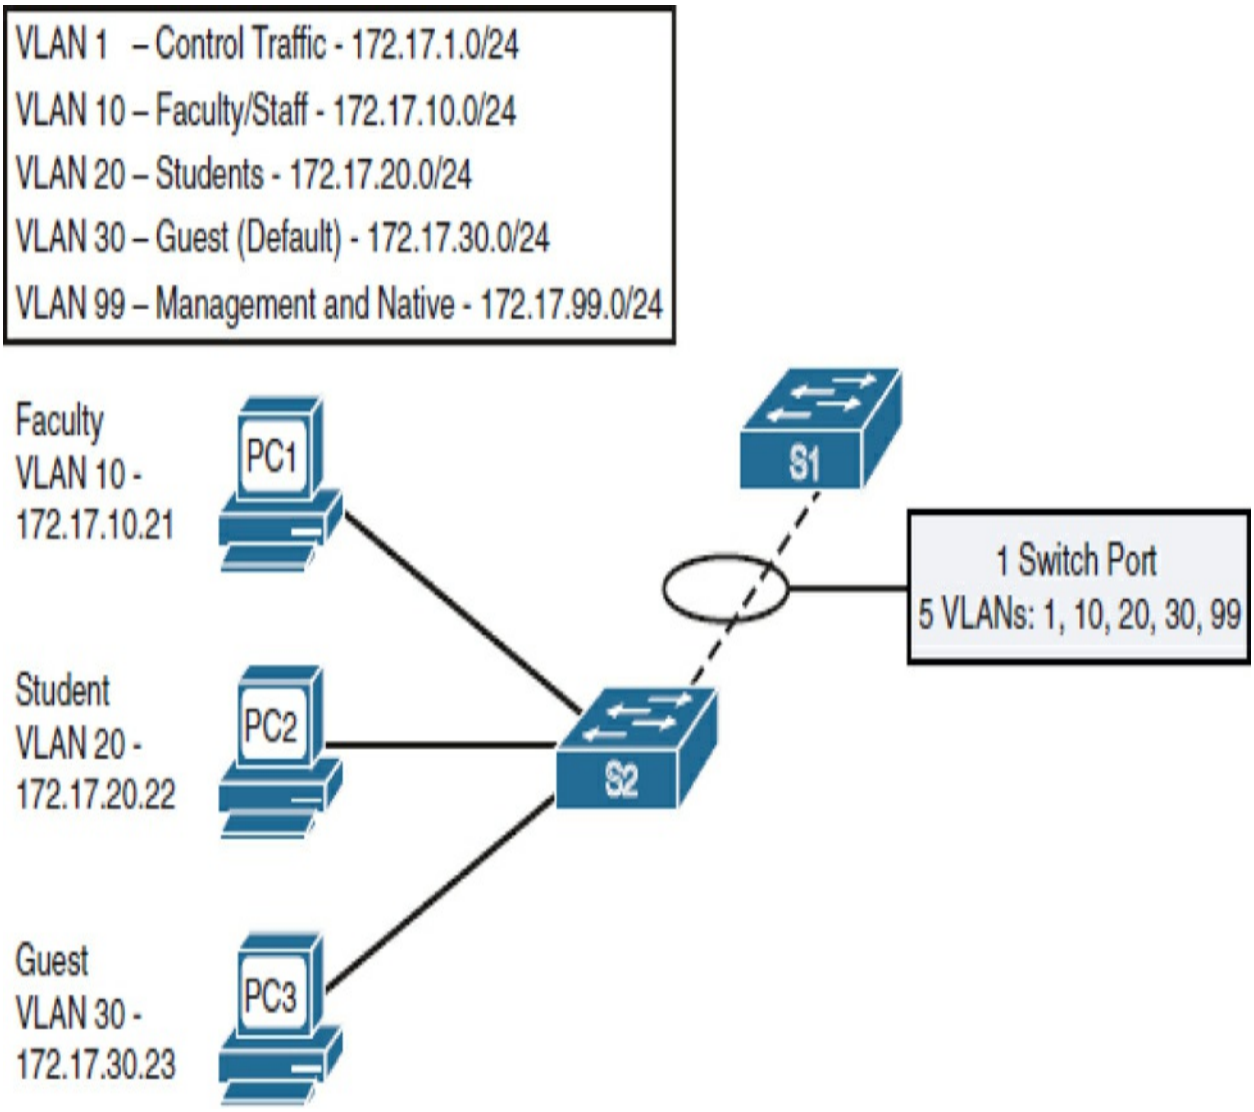

Figure 26-2 *Example of a VLAN Trunk*

When a frame is placed on a trunk link, information about the VLAN it belongs to must be added to the frame. This is accomplished by using IEEE 802.1Q frame tagging. When a switch receives a frame on a port configured in access mode and destined for a remote device through a trunk link, the switch takes apart the frame and inserts a VLAN tag, recalculates the frame check sequence (FCS), and sends the tagged frame out the trunk port. Figure 26-3 shows the 802.1Q tag inserted in an Ethernet frame.

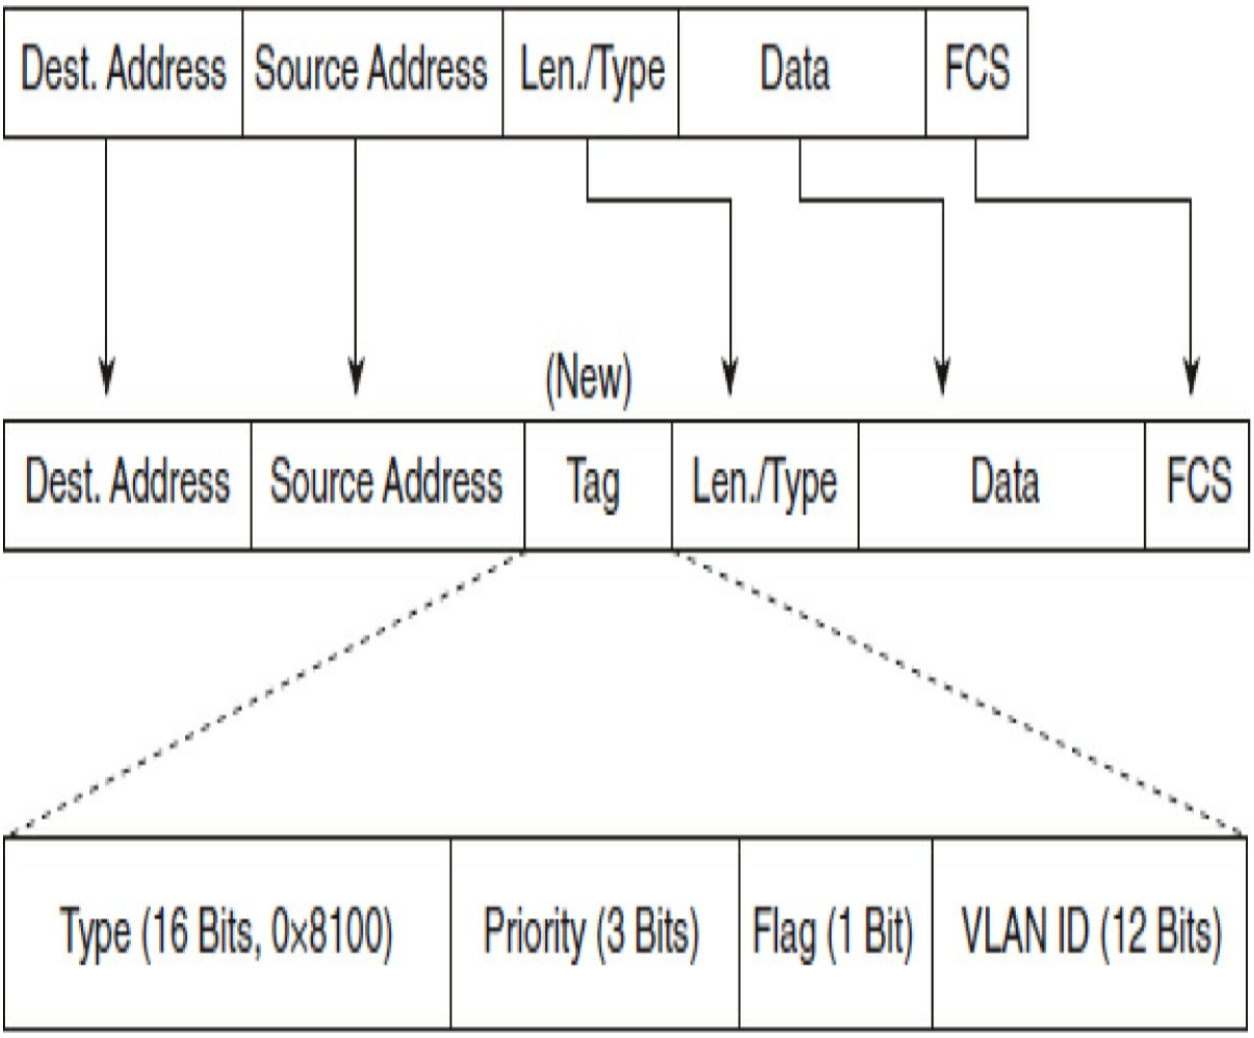

Figure 26-3 *Fields of the 802.1Q Tag Inside an Ethernet Frame*

The VLAN tag field consists of a 16-bit Type field called the EtherType field and a Tag control information field. The EtherType field is set to the hexadecimal value 0x8100. This value is called the tag protocol ID (TPID) value. With the EtherType field set to the TPID value, the switch receiving the frame knows to look for information in the Tag control information field. The Tag control information field contains the following:

• **3 bits of user priority:** Provides expedited transmission of Layer 2 frames,

Technet24

such as voice traffic

• **1 bit of Canonical Format Identifier (CFI):** Enables Token Ring frames to be easily carried across Ethernet links

• **12 bits of VLAN ID (VID):** Provides VLAN identification numbers

---

Note

Although 802.1Q is the recommended method for tagging frames, you should be aware of the Cisco proprietary legacy trunking protocol called Inter-Switch Link (ISL).

Dynamic Trunking Protocol

Dynamic Trunking Protocol (DTP) is a Cisco-proprietary protocol that negotiates both the status of trunk ports and the trunk encapsulation of trunk ports. DTP manages trunk negotiation only if the port on the other switch is configured in a trunk mode that supports DTP. A switch port on a Cisco Catalyst switch supports a number of trunking modes. The trunking mode defines how the port negotiates using DTP to set up a trunk link with its peer port. The following is a brief description of each trunking mode:

• If the switch is configured with the **switchport mode trunk** command, the switch port periodically sends DTP messages to the remote port, advertising that it is in an unconditional trunking state.

• If the switch is configured with the **switchport mode trunk dynamic auto** command, the local switch port advertises to the remote switch port that it is able to trunk but does not request to go to the trunking state. After a DTP negotiation, the local port ends up in the trunking state only if the remote port trunk mode has been configured so that the status is **on** or **desirable**. If both ports on the switches are set to **auto**, they do not negotiate to be in a trunking state. They negotiate to be in the access mode state.

• If the switch is configured with the **switchport mode dynamic desirable** command, the local switch port advertises to the remote switch port that it is able to trunk and asks the remote switch port to go to the trunking state. If the local port detects that the remote port has been configured as **on**, **desirable**, or **auto** mode, the local port ends up in the trunking state. If the remote switch port is in the **nonegotiate** mode, the local switch port remains as a nontrunking port.

• If the switch is configured with the **switchport nonegotiate** command, the local port is considered to be in an unconditional trunking state. Use this feature when you need to configure a trunk with a switch from another switch vendor.

Table 26-2 summarizes the results of DTP negotiations based on the different DTP configuration commands on local and remote ports.

Table 26-2 *Trunk Negotiation Results Between a Local Port and a Remote*

Technet24

Port

	Dynamic Auto	Dynamic Desirable	Trunk	Access
Dynamic Auto	Access	Trunk	Trunk	Access
Dynamic Desirable	Trunk	Trunk	Trunk	Access
Trunk	Trunk	Trunk	Trunk	Not recommended
Access	Access	Access	Not recommended	Access

VLAN Configuration and Verification

Refer to the topology in Figure 26-4 as you review the commands in this section for configuring, verifying, and troubleshooting VLAN and trunking. The packet tracer activity later in the day uses this same topology.

VLAN 1 - Control Traffic - 172.17.1.0/24
VLAN 10 - Faculty/Staff - 172.17.10.0/24
VLAN 20 - Students - 172.17.20.0/24
VLAN 30 - Guest (Default) - 172.17.30.0/24
VLAN 99 - Management and Native - 172.17.99.0/24

Ports
G0/1-2 are 802.1Q trunk interfaces with native VLAN 99

F0/11-17 are in VLAN 15
F0/18-24 are in VLAN 25
F0/6-10 are in VLAN 35

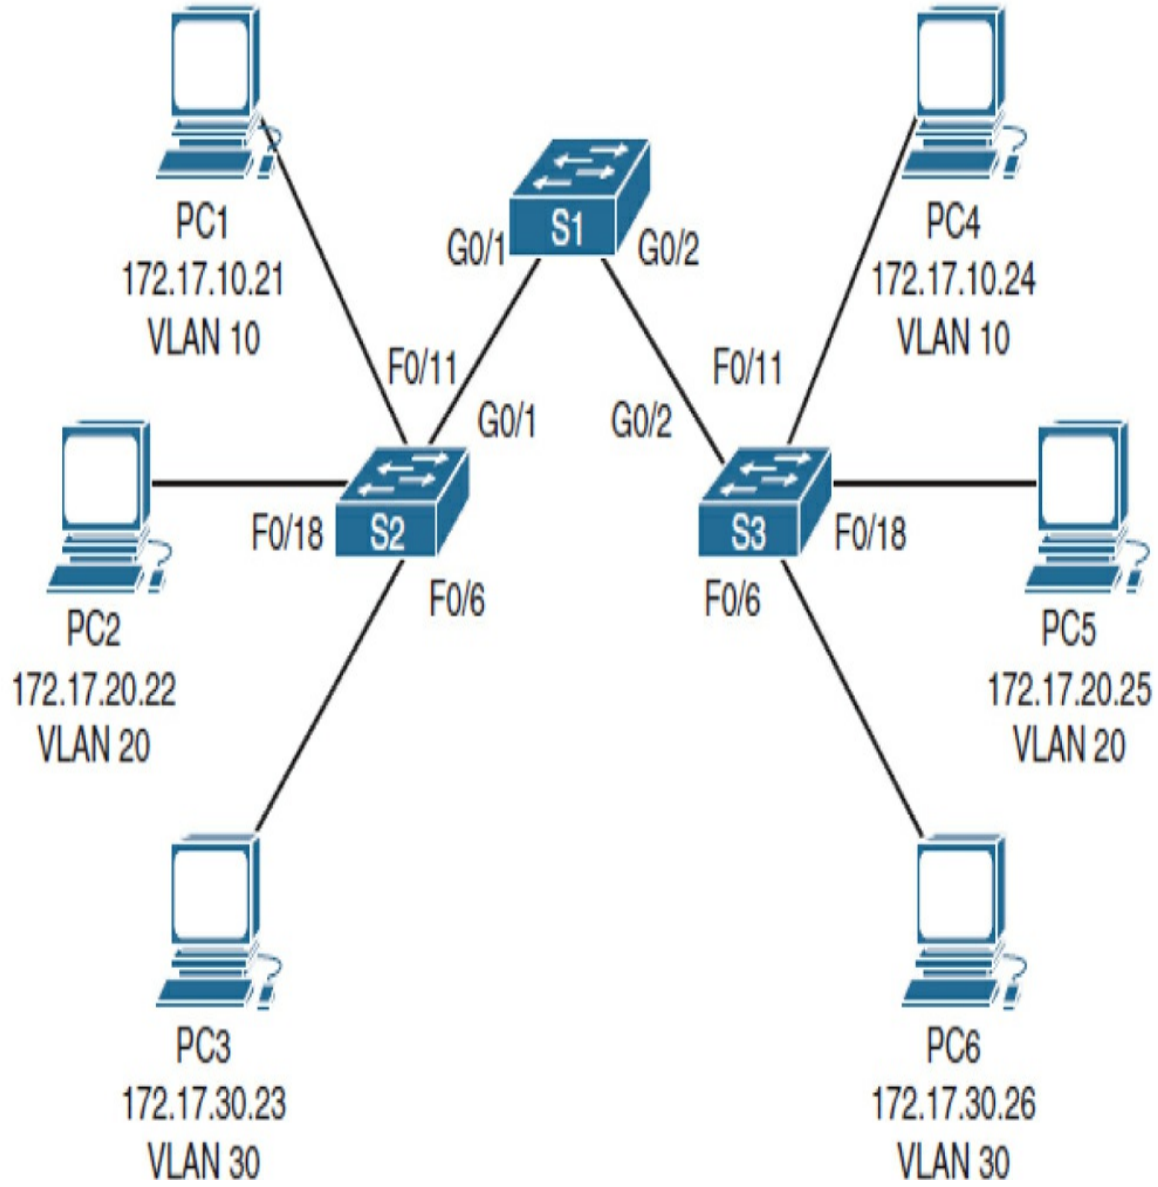

Technet24

Figure 26-4 *Day 26 Sample Topology*

The default configuration of a Cisco switch is to put all interfaces in VLAN 1. You can verify this with the **show vlan brief** command, as demonstrated for S2 in Example 26-1.

Example 26-1 *Default VLAN Configuration*

```
S2# show vlan brief

VLAN   Name                 Status              Ports
----   -------------------------------    ------------
1      default              active              Fa0/1,
                                                    Fa0
                                                 Fa0/9,
                                                 Fa0/13
                                                 Fa0/17
                                                 Fa0/21
                                                 Gig0/1
1002   fddi-default         active
1003   token-ring-default   active
1004   fddinet-default      active
1005   trnet-default        active
S2#
```

A VLAN is created in one of two ways: either in global configuration mode

or directly under the interface. The advantage to configuring in global configuration mode is that you can then assign a name with the **name** *vlan-name* command. The advantage to configuring the VLAN in interface configuration mode is that you assign the VLAN to the interface and create the VLAN with just one command. However, to name the VLAN, you still have to go back to the global configuration method. Example 26-2 shows the creation of VLANs 10 and 20 using these two methods. VLAN 20 is then named, and the remaining VLANs are created in global configuration mode.

Example 26-2 *Creating VLANs*

```
S2# config t
Enter configuration commands, one per line. End with C
S2(config)# vlan 10
S2(config-vlan)# name Faculty/Staff
S2(config-vlan)# interface fa0/18
S2(config-if)# switchport access vlan 20
% Access VLAN does not exist. Creating vlan 20
S2(config-if)# vlan 20
S2(config-vlan)# name Students
S2(config-vlan)# vlan 30
S2(config-vlan)# name Guest(Default)
S2(config-vlan)# vlan 99
S2(config-vlan)# name Management&Native
S2(config-vlan)# end
%SYS-5-CONFIG_I: Configured from console by console
S2#
```

Technet24

Notice in Example 26-3 that all the VLANs are created, but only VLAN 20 is assigned to an interface.

Example 26-3 *Verifying VLAN Creation*

```
S2# show vlan brief

VLAN   Name                 Status        Ports
----   -------------------- ------------- --------------
1      default              active        Fa0/1, Fa0/2,
                                          Fa0/5, Fa0/6,
                                          Fa0/9, Fa0/10,
                                          Fa0/13, Fa0/14
                                          Fa0/17, Fa0/19
                                          Fa0/22, Fa0/23
                                          Gig1/2
10     Faculty/Staff        active
20     Students             active      Fa0/18
30     Guest(Default)       active
99     Management&Native    active
1002  fddi-default          active
1003  token-ring-default    active
1004  fddinet-default       active
1005  trnet-default         active
S2#
```

To assign the remaining interfaces to the VLANs specified in Figure 26-4, either you can configure one interface at a time or you can use the **range** command to configure all the interfaces that belong to a VLAN with one command, as shown in Example 26-4.

Example 26-4 *Assigning VLANs to Interfaces*

```
S2# config t
Enter configuration commands, one per line. End with C
S2(config)# interface range fa 0/11 - 17
S2(config-if-range)# switchport access vlan 10
S2(config-if-range)# interface range fa 0/18 - 24
S2(config-if-range)# switchport access vlan 20
S2(config-if-range)# interface range fa 0/6 - 10
S2(config-if-range)# switchport access vlan 30
S2(config-if-range)# end
%SYS-5-CONFIG_I: Configured from console by console
S2#
```

The **show vlan brief** command in Example 26-5 verifies that all interfaces specified in Figure 26-4 have been assigned to the appropriate VLAN. Notice that unassigned interfaces still belong to the default VLAN 1.

Example 26-5 *Verifying VLAN Assignments to Interfaces*

```
S2# show vlan brief
```

Technet24

```
VLAN  Name                             Status    Ports
----  -------------------------------- --------- ---------
1 default                              active    Fa0/1, Fa0/2, F
                                                 Fa0/5, Gig0/1,
10 Faculty/Staff                       active    Fa0/11, Fa0/12,
                                                 Fa0/15, Fa0/16,
20 Students                            active    Fa0/18, Fa0/19,
                                                 Fa0/22, Fa0/23,
30 Guest(Default)                      active    Fa0/6, Fa0/7, F
                                                 Fa0/10
99 Management&Native                   active
1002 fddi-default                      active
1003 token-ring-default                active
1004 fddinet-default                   active
1005 trnet-default                     active
S2#
```

You can also verify a specific interface's VLAN assignment with the **show interfaces** *type number* **switchport** command, as shown for FastEthernet 0/11 in Example 26-6.

Example 26-6 *Verifying an Interface's VLAN Assignment*

```
S2# show interfaces fastethernet 0/11 switchport
Name: Fa0/11
Switchport: Enabled
```

```
Administrative Mode: dynamic auto
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 10 (Faculty/Staff)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
Protected: false
Appliance trust: none
S2#
```

For the sample topology shown in Figure 26-4, you would configure the VLANs on S1 and S3 as well, but only S3 needs VLANs assigned to interfaces.

Technet24

Trunking Configuration and Verification

Following security best practices, we are configuring a different VLAN for the management and default VLAN. In a production network, you would want to use a different one for each: one for the management VLAN and one for the native VLAN. For expediency here, we are using VLAN 99 for both.

We first define a new management interface for VLAN 99, as in Example 26-7.

Example 26-7 *Defining a New Management Interface*

```
S1# config t
Enter configuration commands, one per line. End with C
S1(config)# interface vlan 99
%LINK-5-CHANGED: Interface Vlan99, changed state to up
S1(config-if)# ip address 172.17.99.31 255.255.255.0
S1(config-if)# end
%SYS-5-CONFIG_I: Configured from console by console
S1#
```

Then we repeat the configuration on S2 and S3. The IP address is used to test connectivity to the switch, as is the IP address the network administrator uses for remote access (Telnet, SSH, SDM, HTTP, and so on).

Depending on the switch model and Cisco IOS version, DTP might have

already established trunking between two switches that are directly connected. For example, the default trunk configuration for 2950 switches is **dynamic desirable**. Therefore, a 2950 initiates trunk negotiations. For our purposes, assume that the switches are all 2960s. The 2960 default trunk configuration is **dynamic auto**, and in this configuration, the interface does not initiate trunk negotiations.

In Example 26-8, the first five interfaces on S1 are configured for trunking. Also notice that the native VLAN is changed to VLAN 99.

Example 26-8 *Trunk Configuration and Native VLAN Assignment*

```
S1# config t
Enter configuration commands, one per line. End with C
S1(config)# interface range g0/1 - 2
S1(config-if-range)# switchport mode trunk
S1(config-if-range)# switchport trunk native vlan 99
S1(config-if-range)# end
%SYS-5-CONFIG_I: Configured from console by console
S1#
%CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch disc
 (99), with S2 FastEthernet0/1 (1).
%CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch disc
 (99), with S3 FastEthernet0/3 (1).
```

If you wait for the next round of CDP messages, you should get the error

Technet24

message shown in Example 26-8. Although the trunk is working between S1 and S2 and between S1 and S3, the switches do not agree on the native VLAN. Repeat the trunking commands on S2 and S3 to correct the native VLAN mismatch.

Note

The encapsulation type—dot1q or isl—might need to be configured, depending on the switch model. The syntax for configuring the encapsulation type is as follows:

```
Switch(config-if)# switchport trunk encapsu
negotiate }
```

The 2960 Series supports only 802.1Q, so this command is not available.

To verify that trunking is operational, use the commands in Example 26-9.

Example 26-9 *Verifying Trunk Configuration*

```
S1# show interfaces trunk
Port        Mode   Encapsulation  Status      Native vlan
Gig0/1      on     802.1q          trunking   99
Gig0/2      on     802.1q          trunking   99
```

```
Port      Vlans allowed on trunk
Gig0/1   1-1005
Gig0/2   1-1005
Port      Vlans allowed and active in management domai
Gig0/1   1,10,20,30,99
Gig0/2   1,10,20,30,99


Port      Vlans in spanning tree forwarding state and
Gig0/1    1,10,20,30,99
Gig0/2    1,10,20,30,99

S1# show interface g0/1 switchport
Name: Gig0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 99 (Management&Native)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
```

Technet24

```
  Administrative private-vlan trunk normal VLANs: none
  Administrative private-vlan trunk private VLANs: none
  Operational private-vlan: none
  Trunking VLANs Enabled: ALL
  Pruning VLANs Enabled: 2-1001
  Capture Mode Disabled
  Capture VLANs Allowed: ALL
  Protected: false
  Appliance trust: none
  S1#
```

Remember that hosts on the same VLAN must be configured with an IP address and subnet mask on the same subnet. The ultimate test of your configuration, then, is to verify that end devices on the same VLAN can now ping each other. If they can't, use the verification commands to systematically track down the problem with your configuration.

VLAN Troubleshooting

If connectivity issues arise between VLANs and you have already resolved potential IP addressing issues, you can use the flowchart in Figure 26-5 to methodically track down any issues related to VLAN configuration errors.

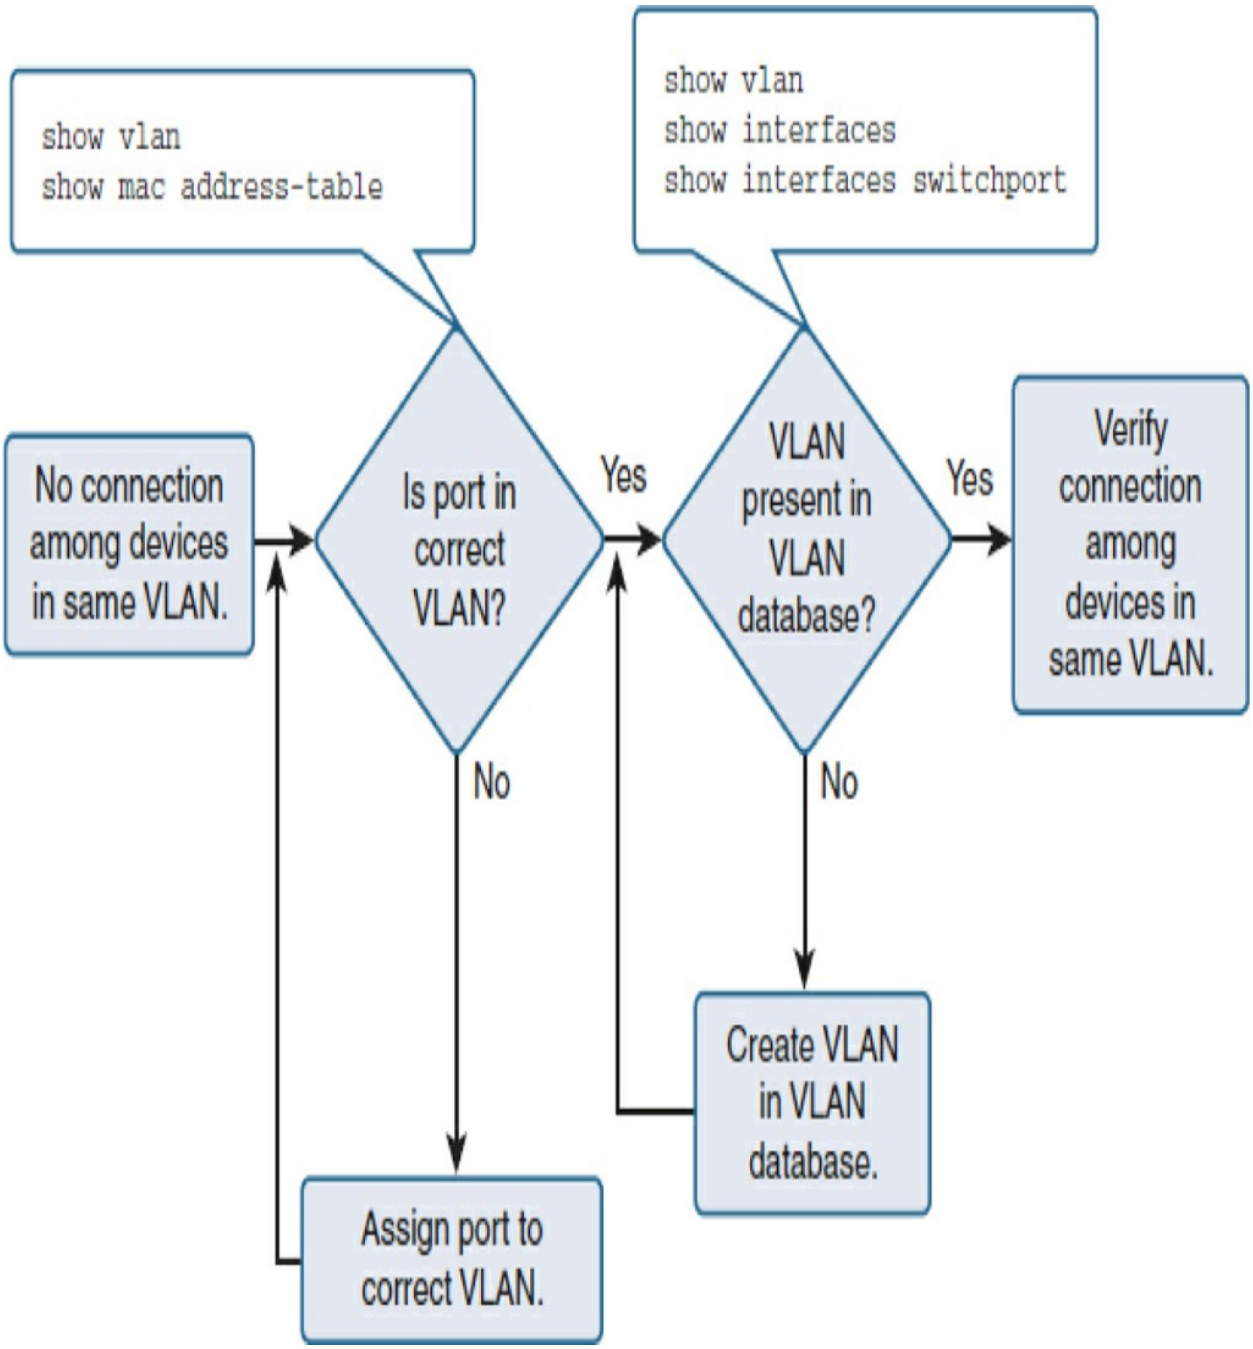

Figure 26-5 *VLAN Troubleshooting Flowchart*

The flowchart in Figure 26-5 works in this way:

Use the **show vlan** command to check whether the port belongs to the expected VLAN. If the port is assigned to the wrong VLAN, use the

Technet24

switchport access vlan command to correct the VLAN membership. Use the **show mac address-table** command to check which addresses were learned on a particular port of the switch and see the VLAN to which that port is assigned.

If the VLAN to which the port is assigned is deleted, the port becomes inactive. Use the **show vlan** or **show interfaces switchport** command to discover issues with deleted VLANs. If the port is inactive, it is not functional until the missing VLAN is created using the **vlan** *vlan_id* command.

Table 26-3 summarizes these commands, which can be particularly helpful in troubleshooting VLAN issues.

Table 26-3 *VLAN Troubleshooting Commands*

EXEC Command	Description
`show vlan` `show vlan brief`	Lists each VLAN and all interfaces assigned to that VLAN (but does not include operational trunks)
`show vlan id` *num*	Lists both access and trunk ports in the VLAN
`show interfaces switchport` `show interfaces` *type number* `switchport`	Identifies the interface's access VLAN and voice VLAN, the configured and operational mode (access or trunk), and the state of the port (up or down)
`show mac address-table`	Lists MAC table entries, including the associated VLAN
`show interface status`	Summarizes the status listing for all interfaces (connected, notconnect, err-disabled), the VLAN, duplex, speed, and type of port

Disabled VLANs

VLANs can be manually disabled. You can verify that VLANs are active by

Technet24

using the **show vlan** command. As Example 26-10 shows, VLANs can be in one of two states: either *active* or *act/lshut*. The second of these states means that the VLAN is shut down.

Example 26-10 *Enabling and Disabling VLANs on a Switch*

```
S1# show vlan brief
VLAN  Name              Status      Ports
----  ----------------- ---------  -------------------
1     default           active      Fa0/1, Fa0/2, Fa0/3
                                     Fa0/5, Fa0/6, Fa0/
                                     Fa0/9, Fa0/10, Fa0
                                     Fa0/14, Fa0/15, Fa
                                     Fa0/18, Fa0/19, Fa
                                     Fa0/22, Fa0/23, Fa
10    VLAN0010          act/lshut   Fa0/13
20    VLAN0020          active
30    VLAN0030          act/lshut
40    VLAN0040          active
S1# configure terminal
Enter configuration commands, one per line. End with C
S1(config)# no shutdown vlan 10
S1(config)# vlan 30
S1(config-vlan)# no shutdown
S1(config-vlan)#
```

The highlighted commands in Example 26-10 show the two configuration methods you can use to enable a VLAN that had been shut down.

Trunking Troubleshooting

To summarize issues with VLANs and trunking, you need to check for four potential issues, in this order:

Identify all access interfaces and their assigned access VLANs and reassign them into the correct VLANs, as needed.

Determine whether the VLANs exist and are active on each switch. If needed, configure and activate the VLANs to resolve problems.

Check the allowed VLAN lists on the switches on both ends of the trunk and ensure that the lists of allowed VLANs are the same.

Ensure that, for any links that should use trunking, one switch does not think it is trunking, while the other switch does not think it is trunking.

The previous section reviewed steps 1 and 2. Next, we review steps 3 and 4.

Check Both Ends of a Trunk

For the CCNA exam, you should be ready to notice a couple oddities that happen with some unfortunate configuration choices on trunks.

It is possible to configure a different allowed VLAN list on the opposite ends

Technet24

of a VLAN trunk. As Figure 26-6 shows, when the VLAN lists do not match, the trunk cannot pass traffic for that VLAN.

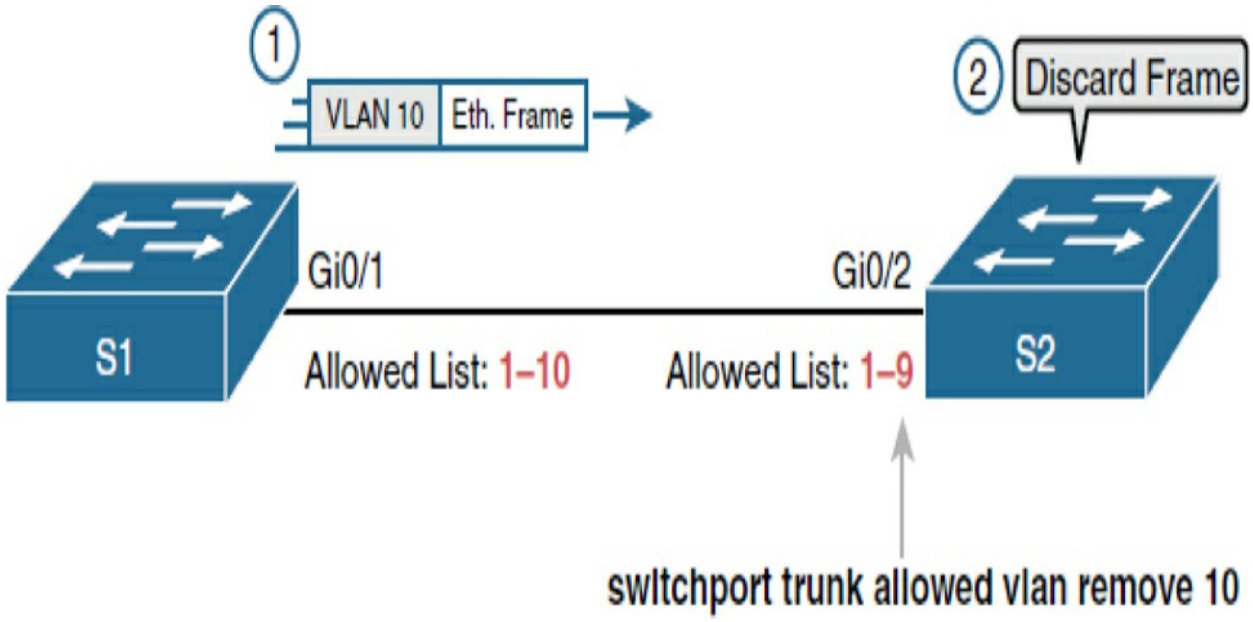

Figure 26-6 *Mismatched VLAN-Allowed Lists on a Trunk*

You can isolate this problem only by comparing the allowed lists on both ends of the trunk. Example 26-9 displays the output of the **show interfaces trunk** command on S2.

To compare the allowed VLANs on each switch, you need to look at the second of three lists of VLANs listed by the **show interfaces trunk** command. See the output in Example 26-11.

Example 26-11 *Verifying the Allowed VLANs on S2*

```
S2# show interfaces trunk
Port    Mode         Encapsulation   Status    Native vla
```

```
Gi0/2  desirable  802.1q          trunking  1

Port Vlans allowed on trunk
Gi0/2 1-4094

Port Vlans allowed and active in management domain
Gi0/2 1-9

Port Vlans in spanning tree forwarding state and not p
Gi0/2 1-9
```

To add VLAN 10 to S2's trunk, enter the following commands:

```
S2(config)# interface g0/2
S2(config-if)# switchport trunk allowed vlan add 10
```

The keyword **add** provides the capability to add one or more VLANs to the trunk without having to specify again all the existing VLANs that are already allowed.

Check Trunking Operational States

Trunks can be misconfigured. In some cases, both switches conclude that their interfaces do not trunk. In other cases, one switch believes that its interface is correctly trunking, while the other switch does not.

Technet24

The most common incorrect configuration—which results in both switches not trunking—is a configuration that uses the **switchport mode dynamic auto** command on both switches on the link. The keyword **auto** does not mean that trunking happens automatically. Instead, both switches passively wait on the other device on the link to begin negotiations.

With this particular incorrect configuration, the **show interfaces switchport** command on both switches confirms both the administrative state (auto) and the fact that both switches operate as static access ports. Example 26-12 highlights those parts of the output for S2.

Example 26-12 *Verifying the Trunking State for a Specific Interface*

```
SW2# show interfaces gigabit0/2 switchport
Name: Gi0/2
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
! lines omitted for brevity
```

Always check the trunk's operational state on both sides of the trunk. The best commands for checking trunking-related facts are **show interfaces trunk** and **show interfaces switchport**.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	3
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	8
Portable Command Guide	9
	10

Technet24

Day 25. STP

CCNA 200-301 Exam Topic

• 2.5 Interpret basic operations of Rapid PVST+ Spanning Tree Protocol

Key Topics

Today's review covers the operation and configuration of Spanning Tree Protocol (STP). The original STP IEEE 802.1D standard allowed for only one instance of STP to run for an entire switched network. Today's network administrators can implement Per-VLAN Spanning Tree (PVST) and Rapid STP (RSTP), both of which improve the original standard.

STP Concepts and Operation

A key characteristic of a well-built communications network is its resiliency. A resilient network is capable of handling a device or link failure through redundancy. A redundant topology can eliminate a single point of failure by using multiple links, multiple devices, or both. STP helps prevent loops in a redundant switched network. Figure 25-1 shows an example of a three-layer topology (core, distribution, access) with redundant links.

Without STP, redundancy in a switched network can introduce the following issues:

• **Broadcast storms:** Each switch floods broadcasts endlessly.

• **Multiple-frame transmission:** Multiple copies of unicast frames are delivered to the destination, causing unrecoverable errors.

• **MAC table instability:** Instability in the content of the MAC address table results from different ports of the switch receiving copies of the same frame.

Technet24

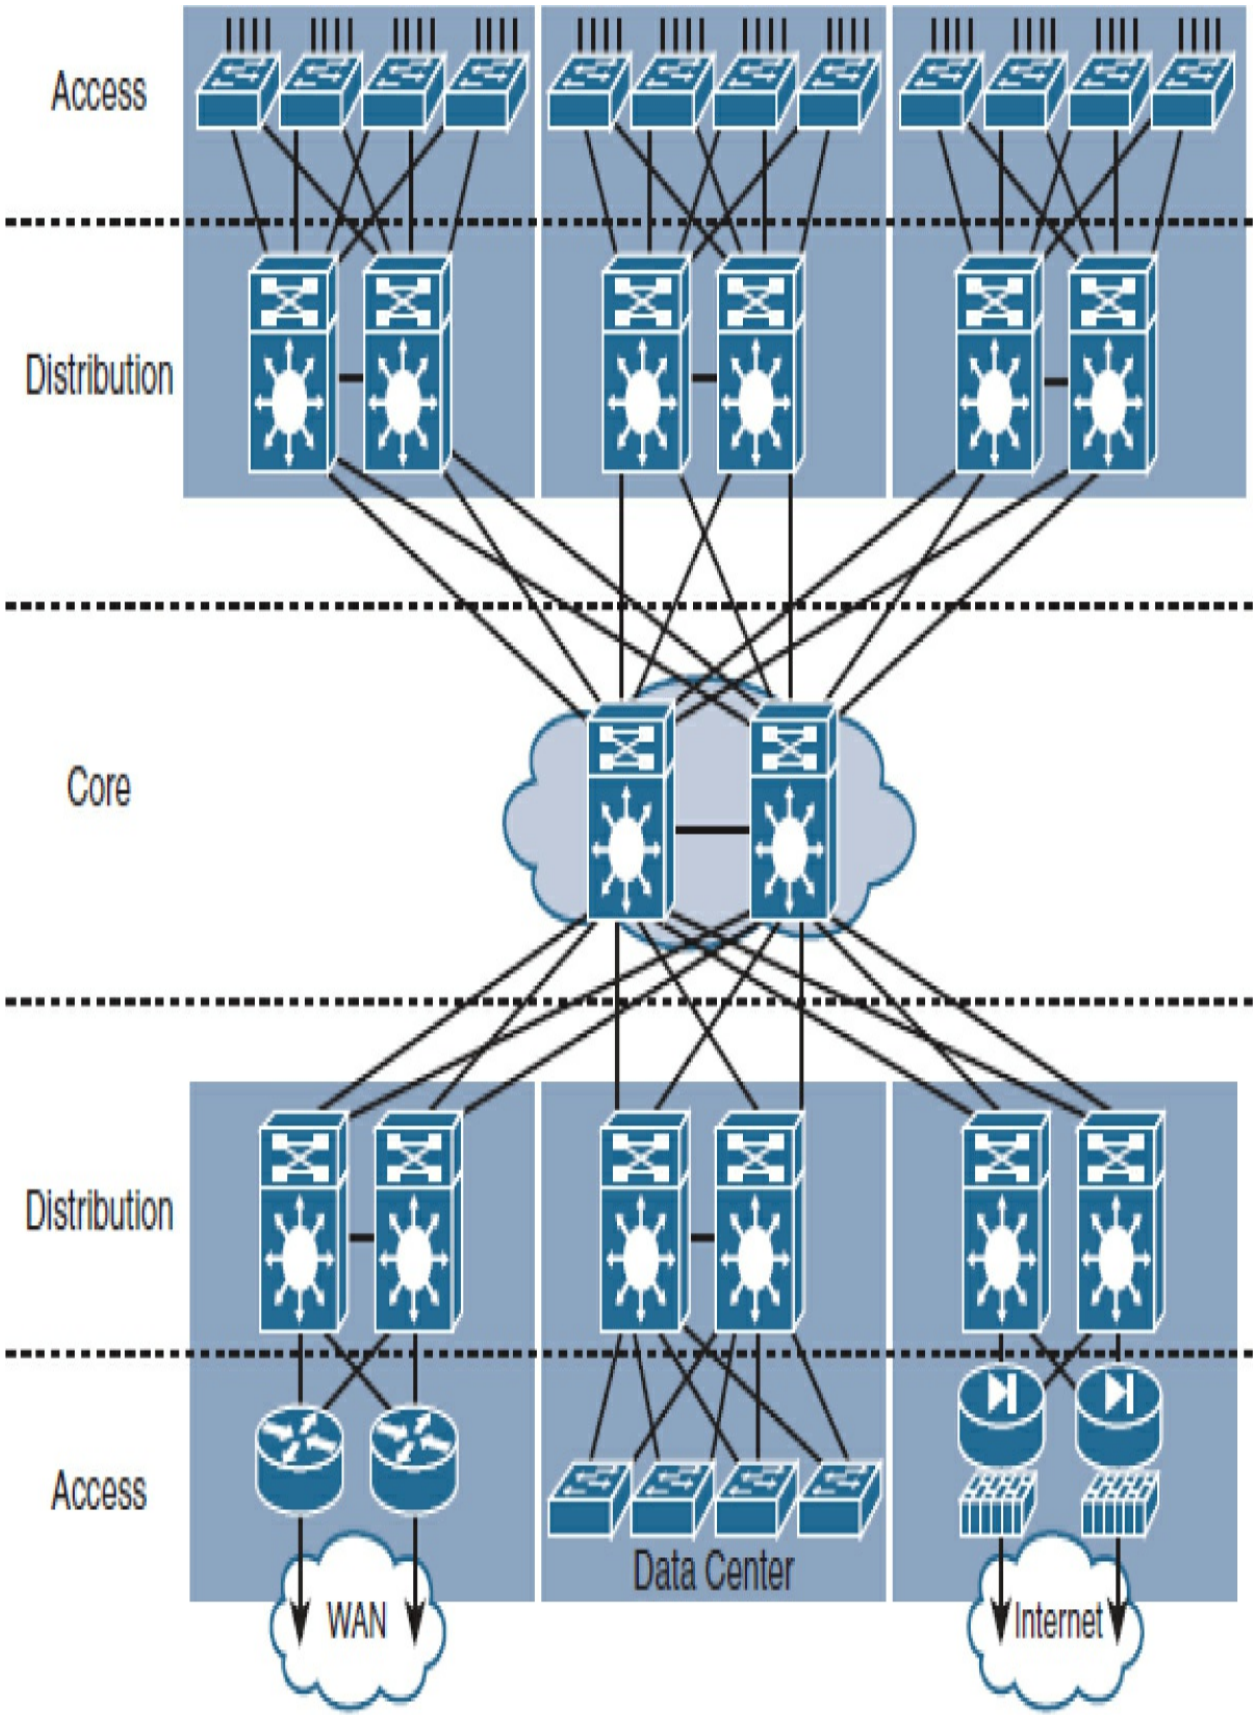
Access
Distribution
Core
Distribution
Access
WAN
Data Center
Internet

Figure 25-1 *Redundant Switched Topology*

STP Algorithm

STP is an IEEE Committee standard defined as 802.1D. STP places certain ports in the blocking state so that they do not listen to, forward, or flood data frames. STP creates a tree that ensures that only one path exists for each network segment at any one time. If any segment experiences a disruption in connectivity, STP rebuilds a new tree by activating the previously inactive but redundant path.

The algorithm STP uses chooses the interfaces that should be placed into a forwarding state. For any interfaces not chosen to be in a forwarding state, STP places the interfaces in blocking state.

Switches exchange STP configuration messages every 2 seconds, by default, using a multicast frame called the bridge protocol data unit (BPDU). Blocked ports listen for these BPDUs to detect whether the other side of the link is down, thus requiring an STP recalculation. One piece of information included in the BPDU is the bridge ID (BID).

As Figure 25-2 shows, the BID is unique to each switch. It consists of a priority value (2 bytes) and the bridge MAC address (6 bytes).

Technet24

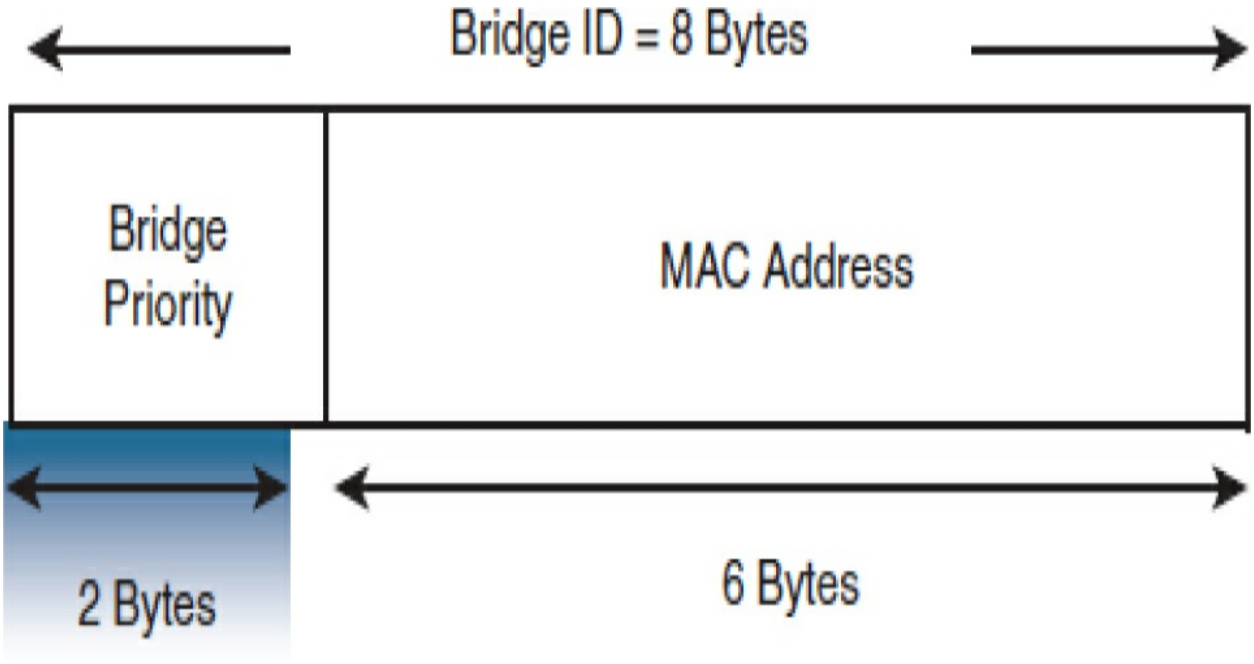

Figure 25-2 *Bridge ID*

The default priority is 32,768. The root bridge is the bridge with the lowest BID. Therefore, if the default priority value is not changed, the switch with the lowest MAC address becomes the root.

STP Convergence

STP convergence is the process by which switches collectively realize that something has changed in the LAN topology. The switches determine whether they need to change which ports block and which ports forward. The following steps summarize the STP algorithm used to achieve convergence:

Elect a root bridge (that is, the switch with the lowest BID). Only one root bridge can exist per network. All ports on the root bridge are forwarding ports.

Elect a root port for each nonroot switch, based on the lowest root path cost. Each nonroot switch has one root port. The root port is the port through which the nonroot bridge has its best path to the root bridge.

Elect a designated port for each segment, based on the lowest root path cost. Each link has one designated port.

The root ports and designated ports transition to the forwarding state, and the other ports stay in the blocking state.

Table 25-1 summarizes the reasons STP places a port in forwarding or blocking state.

Table 25-1 *STP: Reasons for Forwarding or Blocking*

Technet24

Characterization of Port	STP State	Description
All the root switch's ports	Forwarding	The root switch is always the designated switch on all connected segments.
Each nonroot switch's root port	Forwarding	This is the port through which the switch has the least cost to reach the root switch.
Each LAN's designated port	Forwarding	The switch forwarding the lowest-cost BPDU onto the segment is the designated switch for that segment.
All other working ports	Blocking	The port is not used for forwarding frames, nor are any frames received on these interfaces considered for forwarding. BPDUs are still received.

Port bandwidth is used to determine the cost to reach the root bridge. Table 25-2 lists the default port costs defined by the IEEE; these have been revised over time as ports have continued to get faster. You can configure a switch to

use either the old (short) or new (long) defaults with the global command **spanning-tree pathcost method** {**short** | **long**}. However, it is advisable to use the same setting on all switches within the same campus for consistency.

Table 25-2 *Default IEEE Port Costs*

Ethernet Speed	IEEE Cost: 1998 (and Before)	IEEE Cost: 2004 (and After)
10 Mbps	100	2,000,000
100 Mbps	19	200,000
1 Gbps	4	20,000
10 Gbps	2	2000
100 Gbps	N/A	200
1 Tbps	N/A	20
10 Tbps	N/A	2

STP uses the four states in Figure 25-3 as port transitions from blocking to

Technet24

forwarding.

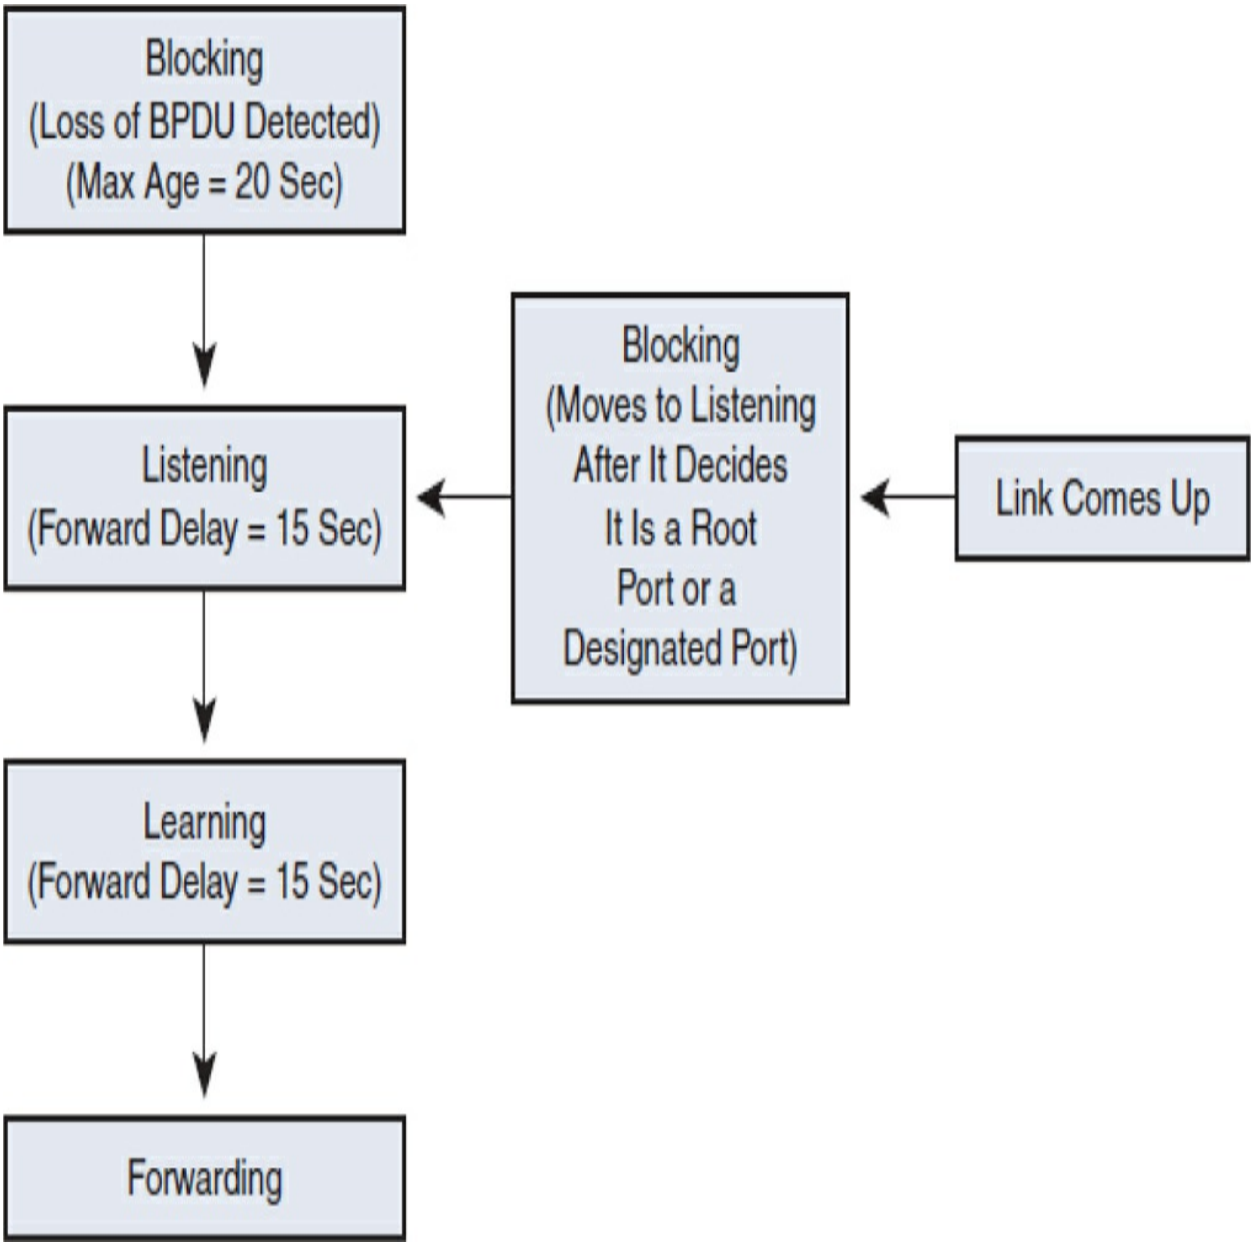

Figure 25-3 *Spanning Tree Port States*

A fifth state, disabled, occurs either when a network administrator manually disables the port or when a security violation disables the port.

STP Varieties

Several varieties of STP emerged after the original IEEE 802.1D:

• **STP:** The original specification of STP, defined in 802.1D, provides a loop-free topology in a network with redundant links. STP is sometimes referred to as Common Spanning Tree (CST) because it assumes one spanning tree instance for the entire bridged network, regardless of the number of VLANs.

• **PVST+:** Per-VLAN Spanning Tree Plus (PVST+) is a Cisco enhancement of STP that provides a separate 802.1D spanning tree instance for each VLAN configured in the network.

• **RSTP:** Rapid STP (RSTP), or IEEE 802.1w, is an evolution of STP that provides faster convergence than STP. However, RSTP still provides for only a single instance of STP.

• **Rapid PVST+:** Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+. Rapid PVST+ provides a separate instance of 802.1w per VLAN.

• **MSTP and MST:** Multiple Spanning Tree Protocol (MSTP) is an IEEE standard inspired by the earlier Cisco-proprietary Multiple Instance STP (MISTP) implementation. MSTP maps multiple VLANs into the same spanning tree instance. The Cisco implementation of MSTP is Multiple Spanning Tree (MST), which provides up to 16 instances of RSTP and combines many VLANs with the same physical and logical topology into a

Technet24

common RSTP instance.

Part of your switch administration skill set is the ability to decide which type of STP to implement. Table 25-3 summarizes the features of the various STP flavors.

Table 25-3 *Features of STP Varieties*

Protocol	Standard	Resources Needed	Convergence	Tree Calculation
STP	802.1D	Low	Slow	All VLANs
PVST+	Cisco	High	Slow	Per VLAN
RSTP	802.1w	Medium	Fast	All VLANs
Rapid PVST+	Cisco	Very high	Fast	Per VLAN
MSTP	802.1s, Cisco	Medium or high	Fast	Per instance

PVST Operation

PVST Plus (PVST+) is the default setting on all Cisco Catalyst switches. In a PVST+ environment, you can tune the spanning-tree parameters so that half

the VLANs forward on each uplink trunk. You do this by configuring one switch to be elected the root bridge for half of the VLANs in the network and a second switch to be elected the root bridge for the other half of the VLANs. In the example in Figure 25-4, S1 is the root bridge for VLAN 10, and S3 is the root bridge for VLAN 20.

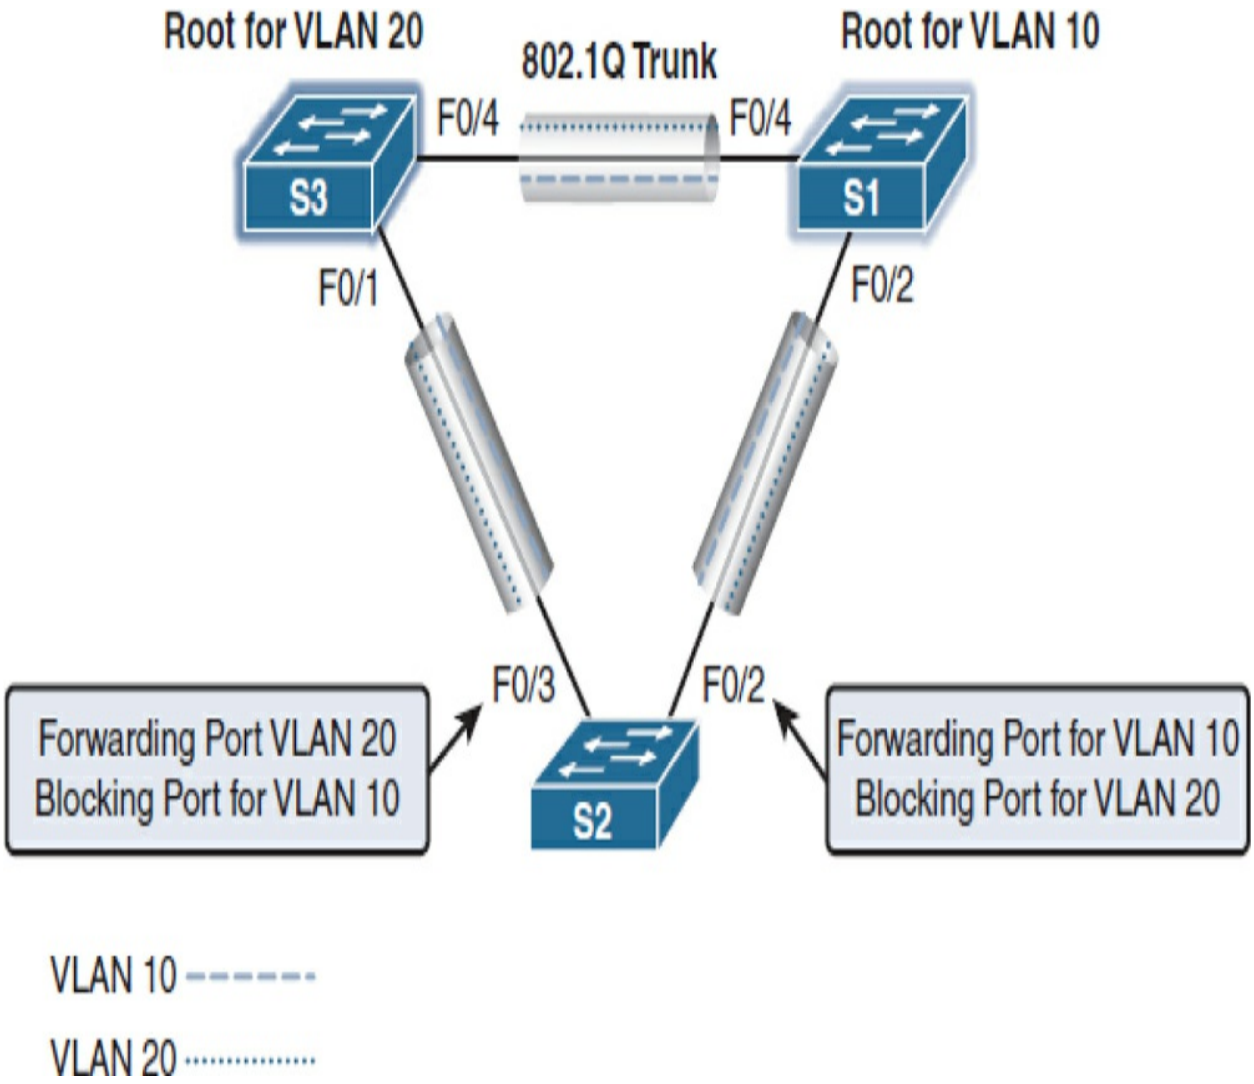

Figure 25-4 *PVST+ Topology Example*

From the perspective of S2, a port is forwarding or blocking depending on the VLAN instance. After convergence, port F0/2 will be forwarding VLAN

Technet24

10 frames and blocking VLAN 20 frames. Port F0/3 will be forwarding VLAN 20 frames and blocking VLAN 10 frames.

Switched networks running PVST+ have the following characteristics:

• Configured PVST per VLAN allows redundant links to be fully utilized.

• Each additional spanning tree instance for a VLAN adds more CPU cycles to all switches in the network.

Port States

The spanning tree is determined immediately after a switch is finished booting. If a switch port transitions directly from the blocking state to the forwarding state without information about the full topology during the transition, the port can temporarily create a data loop. For this reason, STP introduces the five port states. Table 25-4 describes the port states that ensure that no loops are created during the creation of the logical spanning tree.

Table 25-4 *PVST Port States*

Operation Allowed	Blocking	Listening	Learning	Forwarding	Disabled
Can receive and process BPDUs	Yes	Yes	Yes	Yes	No
Can forward data frames received on the interface	No	No	No	Yes	No
Can forward data frames switched from another interface	No	No	No	Yes	No
Can learn MAC addresses	No	No	Yes	Yes	No

Extended System ID

PVST+ requires a separate instance of spanning tree for each VLAN. The BID field in the BPDU must carry VLAN ID (VID) information, as Figure 25-5 shows.

Technet24

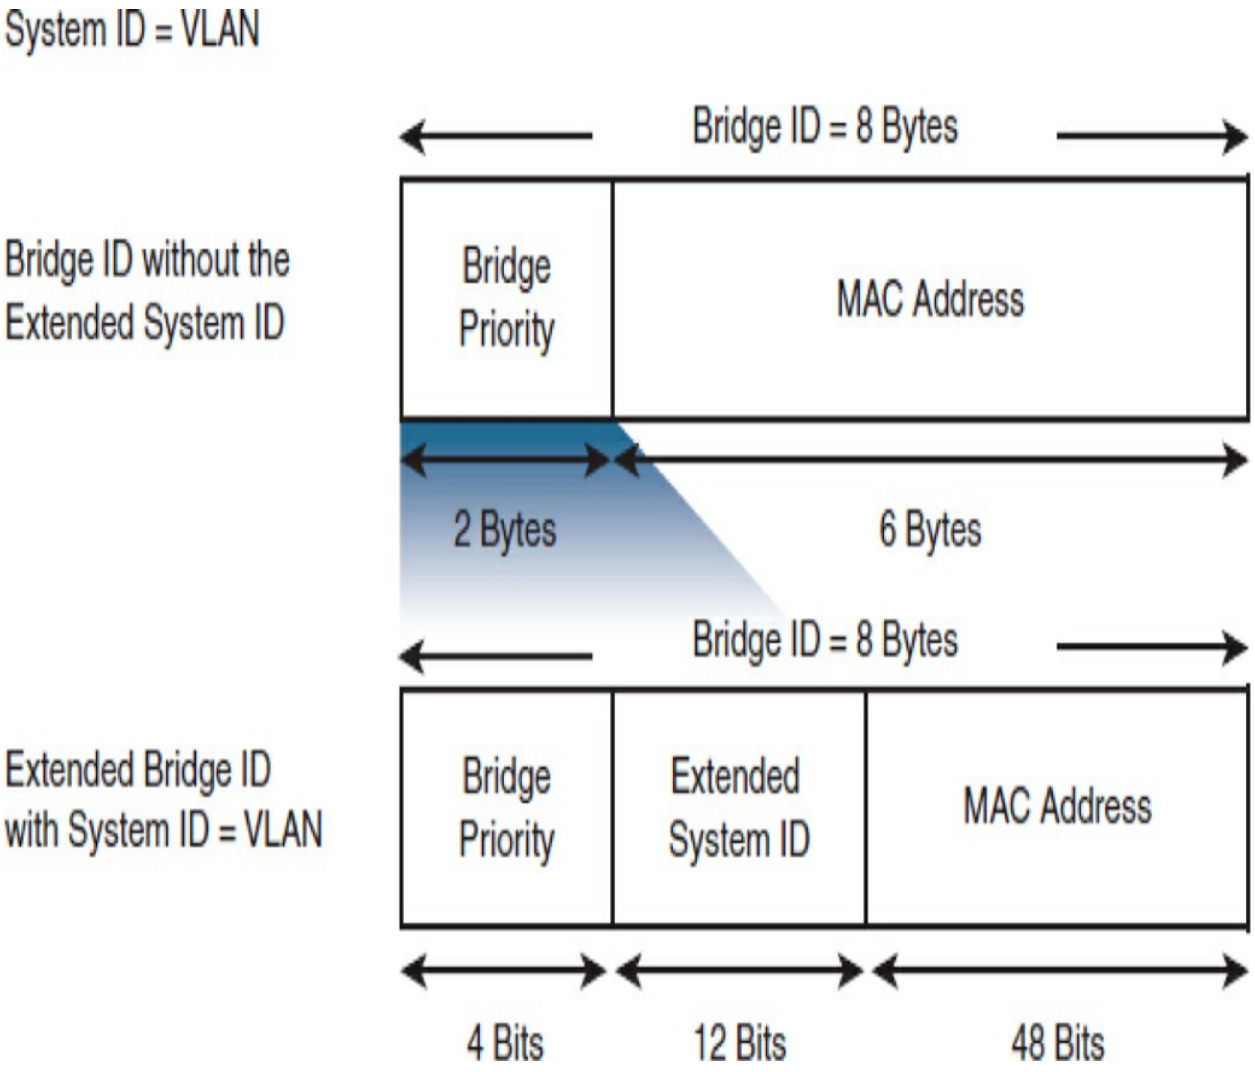

Figure 25-5 *Bridge ID for PVST+ with Extended System ID*

The BID includes the following fields:

• **Bridge Priority:** A 4-bit field is still used to carry bridge priority. However, the priority is conveyed in discrete values in increments of 4096 instead of discrete values in increments of 1 because only the first 4 most-significant bits are available from the 16-bit field.

• **Extended System ID:** A 12-bit field carrying the VID for PVST+.

• **MAC Address:** A 6-byte field with the MAC address of a single switch.

Rapid PVST+ Operation

In Rapid PVST+, a single instance of RSTP runs for each VLAN. This is why Rapid PVST+ has a very high demand for switch resources (CPU cycles and RAM).

Note

Rapid PVST+ is simply the Cisco implementation of RSTP on a per-VLAN basis. The rest of this review uses the terms *RSTP* and *Rapid PVST+* interchangeably.

With RSTP, the IEEE improved the convergence performance of STP from 50 seconds to less than 10 seconds in the RSTP standard 802.1w. RSTP is identical to STP in the following ways:

• It elects the root switch by using the same parameters and tiebreakers.

• It elects the root port on nonroot switches by using the same rules.

• It elects designated ports on each LAN segment by using the same rules.

• It places each port in either forwarding or blocking state, although RSTP calls the blocking state the discarding state.

Technet24

RSTP Interface Behavior

The main changes with RSTP can be seen when changes occur in the network. RSTP acts differently on some interfaces based on what is connected to the interface:

• **Edge-type behavior and PortFast:** RSTP improves convergence for edge-type connections by immediately placing the port in forwarding state when the link is physically active.

• **Link-type shared:** RSTP does not do anything differently from STP on link-type shared links. However, because most links between switches today are full duplex, point-to-point, and not shared, this does not matter.

• **Link-type point-to-point:** RSTP improves convergence over full-duplex links between switches. RSTP recognizes the loss of the path to the root bridge through the root port in 6 seconds (based on three times the hello timer value of 2 seconds). RSTP thus recognizes a lost path to the root much more quickly.

RSTP uses different terminology to describe port states. Table 25-5 lists the port states for RSTP and STP.

Table 25-5 *RSTP and STP Port States*

Operational State	STP State (802.1D)	RSTP State (802.1w)	Forwards Data Frames in This State?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	No
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

RSTP removes the need for listening state and reduces the time required for learning state by actively discovering the network's new state. STP passively waits on new BPDUs and reacts to them during the listening and learning states. With RSTP, the switches negotiate with neighboring switches by sending RSTP messages. The messages enable the switches to quickly determine whether an interface can be immediately transitioned to a forwarding state. In many cases, the process takes only a second or two for the entire RSTP domain.

RSTP Port Roles

Technet24

RSTP adds three more port roles in addition to the root port and designated port roles defined in STP. Table 25-6 lists and defines the port roles.

Table 25-6 *RSTP and STP Port Roles*

RSTP Role	STP Role	Definition
Root port	Root port	A single port on each nonroot switch in which the switch hears the best BPDU out of all the received BPDUs
Designated port	Designated port	Of all switch ports on all switches attached to the same segment/collision domain, the port that advertises the "best" BPDU
Alternate port	—	A port on a switch that receives a suboptimal BPDU
Backup port	—	A nondesignated port on a switch that is attached to the same segment/collision domain as another port on the same switch
Disabled	—	A port that is administratively disabled or that is not capable of working for other reasons

Figure 25-6 shows an example of these RSTP port roles.

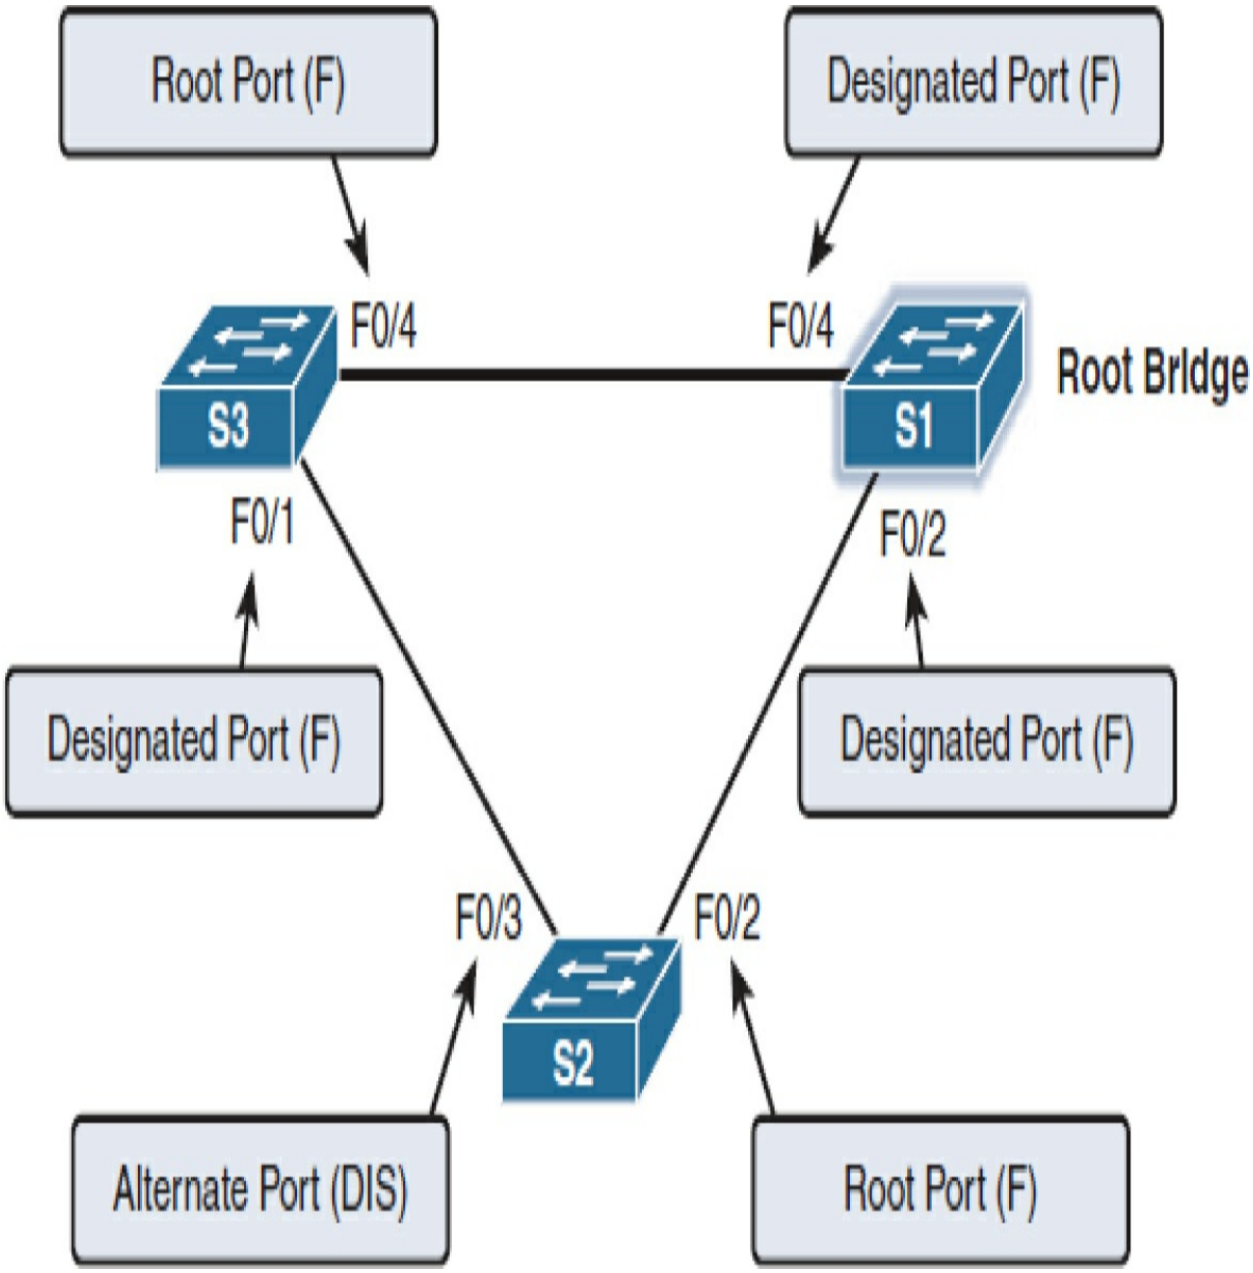

Figure 25-6 *RSTP Port Roles*

Edge Ports

In addition to the port roles just described, RSTP uses an edge port concept that corresponds to the PVST+ PortFast feature. An edge port connects directly to an end device. Therefore, the switch assumes that no other switch

Technet24

is connected to it. RSTP edge ports should immediately transition to the forwarding state, thereby skipping the time-consuming original 802.1D listening and learning port states. The only caveat is that the port must be a point-to-point link. If it is a shared link, the port is not an edge port, and PortFast should not be configured. Why? Another switch could be added to a shared link—on purpose or inadvertently. Figure 25-7 shows examples of edge ports.

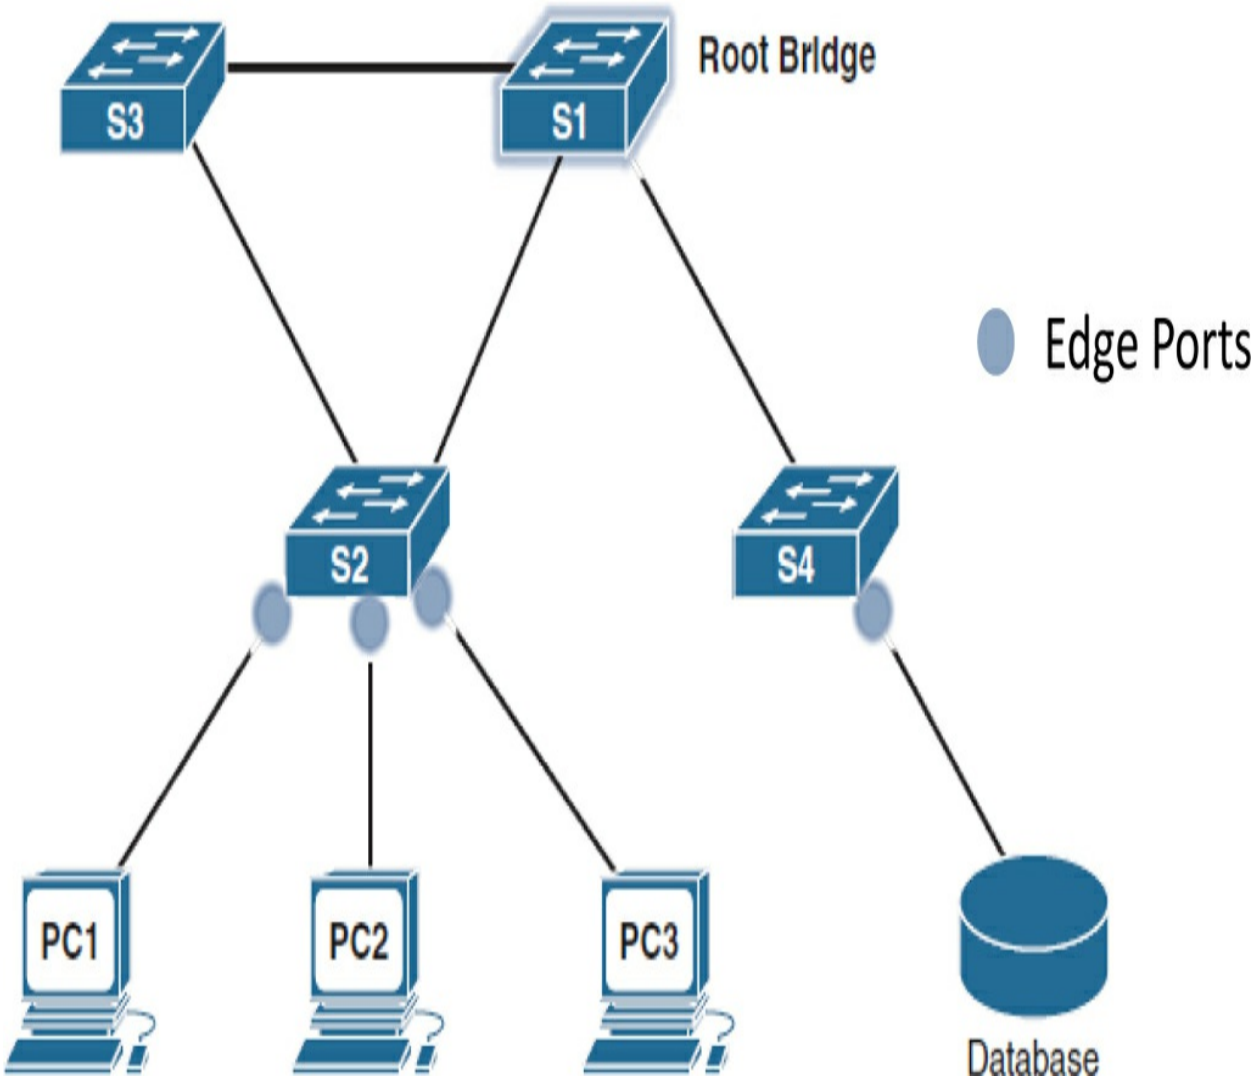

Figure 25-7 *Edge Ports in RSTP*

Configuring and Verifying Varieties of STP

By default, all Cisco switches use STP without any configuration by the network administrator. However, because STP runs on a per-VLAN basis, you can take advantage of several options to load balance traffic across redundant links.

STP Configuration Overview

Before you configure or alter the behavior of STP, it is important to know the current default settings. There are some minor differences, depending on which Cisco Catalyst switch series you are configuring, as highlighted in Table 25-7.

Table 25-7 *Default STP Configuration on Cisco Catalyst Switches*

Technet24

Feature	2900 Series Default Settings	1000 Series Default Settings	9300 Series Default Settings
Enable state	Enables STP on VLAN 1	Enables STP on VLAN 1	Enables STP on VLAN 1
Spanning tree mode	PVST+ (Rapid PVST+ and MSTP disabled)	Rapid PVST+ (PVST+ and MSTP disabled)	Rapid PVST+ (PVST+ and MSTP disabled)
Switch priority	32768	32768	32768
Spanning tree port priority (configurable on a per-interface basis)	128	128	128
Spanning tree port cost (configurable on a per-interface basis)	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100	10 Mbps: 2000000 100 Mbps: 200000 1 Gbps: 20000 10 Gbps: 2000 40 Gbps: 500 100 Gbps: 200 1 Tbps: 20 10 Tbps: 2

Spanning tree VLAN port priority (configurable on a per-VLAN basis)	128	128	128
Spanning tree VLAN port cost (configurable on a per-VLAN basis)	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100	1000 Mbps: 4 100 Mbps: 19 10 Mbps: 100	10 Mbps: 2000000 100 Mbps: 200000 1 Gbps: 20000 10 Gbps: 2000 40 Gbps: 500 100 Gbps: 200 1 Tbps: 20 10 Tbps: 2
Spanning tree timers	Hello time: 2 seconds Forward-delay time: 15 seconds Maximum-aging time: 20 seconds Transmit hold count: 6 BPDUs	Hello time: 2 seconds Forward-delay time: 15 seconds Maximum-aging time: 20 seconds Transmit hold count: 6 BPDUs	Hello time: 2 seconds Forward-delay time: 15 seconds Maximum-aging time: 20 seconds Transmit hold count: 6 BPDUs

Configuring Rapid PVST+

Technet24

Rapid PVST+ is on by default in the later Cisco Catalyst switch models. If you are using a Cisco Catalyst 2900 series switch, you can configuring Rapid PVST+ with the global configuration command **spanning tree mode** { **mst** | **pvst** | **rapid-pvst** }, as shown in Example 25-1.

Example 25-1 *Cisco Switch Spanning Tree Modes*

```
S1(config)# spanning-tree mode ?
  mst            Multiple spanning tree mode
  pvst           Per-Vlan spanning tree mode
  rapid-pvst     Per-Vlan rapid spanning tree mode

S1(config)# spanning-tree mode
```

Configuring and Verifying the BID

Regardless of which PVST you use, two main configuration options can help you achieve load balancing: the bridge ID and the port cost manipulation. The bridge ID influences the choice of root switch and can be configured per VLAN. Each interface's (per-VLAN) STP cost to reach the root influences the choice of designated port on each LAN segment. Because PVST requires that a separate instance of spanning tree run for each VLAN, the BID field is required to carry VLAN ID (VID) information. This is accomplished by reusing a portion of the Priority field as the extended system ID to carry a VID.

To change the bridge ID, use one of the following commands:

```
Switch(config)# spanning-tree vlan vlan-id root {prima
Switch(config)# spanning-tree vlan vlan-id priority pr
```

To change the interface cost, use the following command:

```
Switch(config-if)# spanning-tree vlan vlan-id cost cos
```

Figure 25-8 shows a simple three-switch STP topology without redundant links.

Technet24

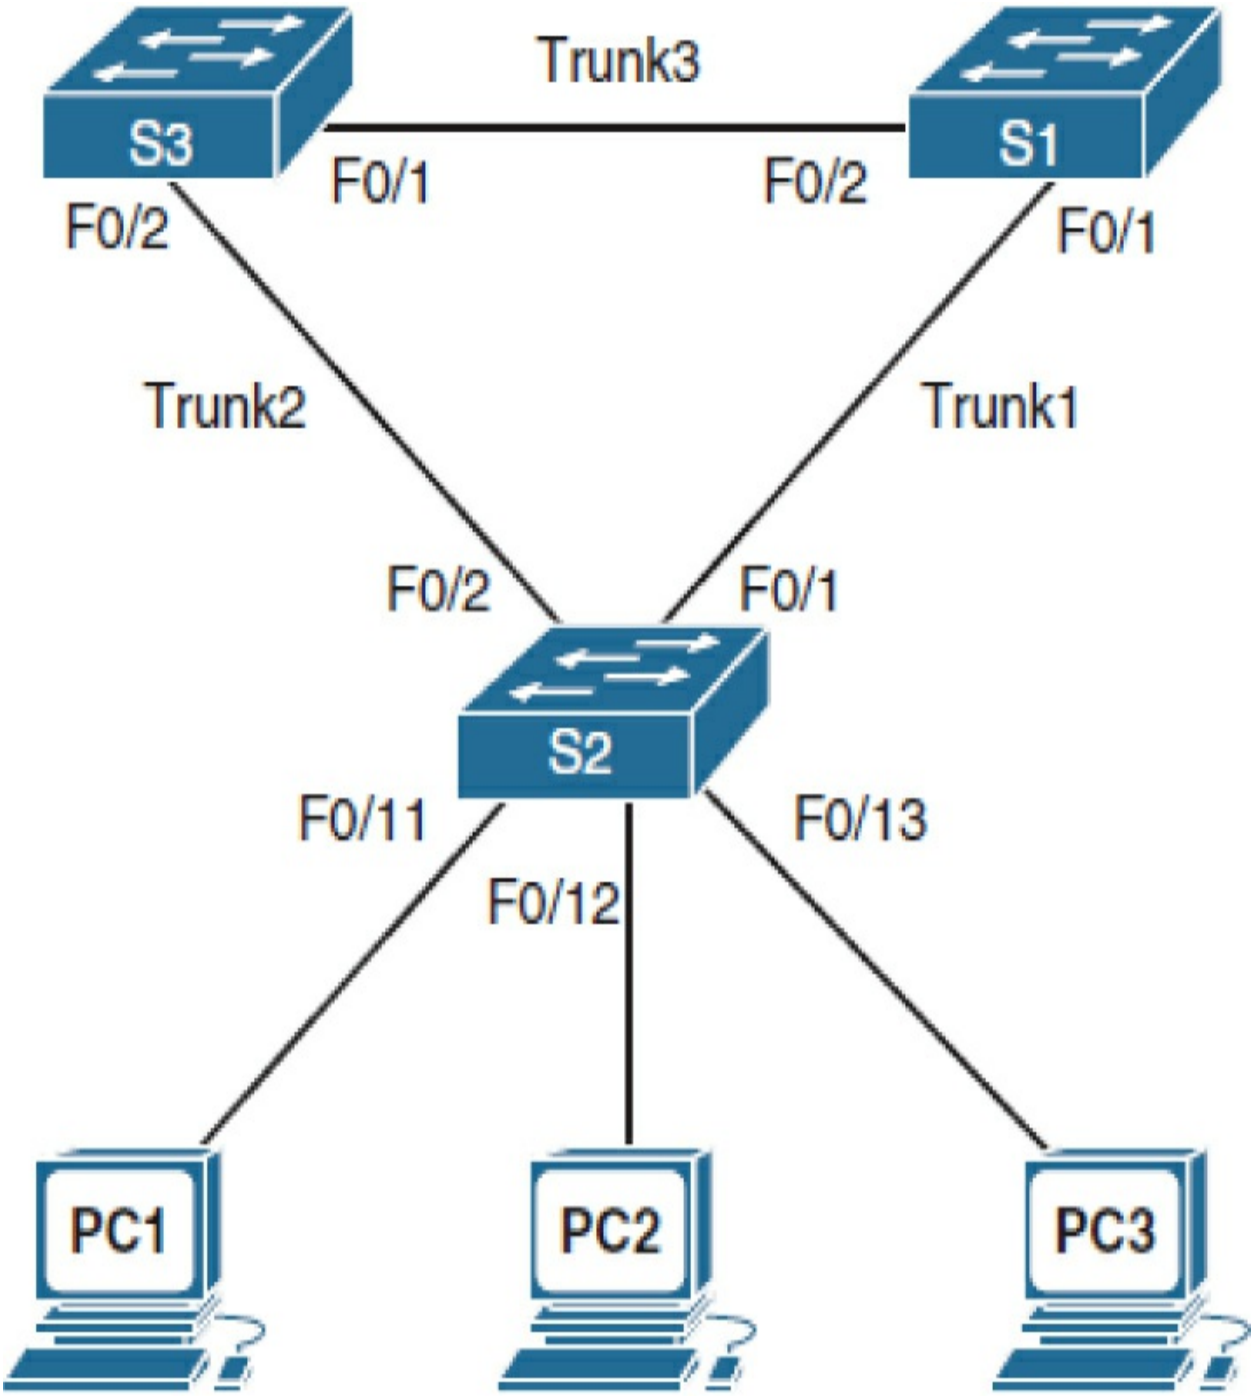

Figure 25-8 *STP Topology*

The network administrator wants to ensure that S1 is always the root bridge and S2 is the backup root bridge. The following commands achieve this objective:

```
S1(config)# spanning-tree vlan 1 root primary
!---------
S2(config)# spanning-tree vlan 1 root secondary
```

The **primary** keyword automatically sets the priority to 24576 or to the next 4096 increment value below the lowest bridge priority detected on the network.

The **secondary** keyword automatically sets the priority to 28672, assuming that the rest of the network is set to the default priority of 32768.

Alternatively, the network administrator can explicitly configure the priority value in increments of 4096 between 0 and 65536 using the following command:

```
S1(config)# spanning-tree vlan 1 priority 24576
!---------
S2(config)# spanning-tree vlan 1 priority 28672
```

Note

In this example, these commands changed the priority values only for VLAN 1. Additional commands must be entered for each VLAN to take advantage of load balancing.

Technet24

To verify the current spanning tree instances and root bridges, use the **show spanning-tree** command (see Example 25-2).

Example 25-2 *Verifying Spanning Tree Configurations*

```
S1# show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority      24577
             Address       001b.5302.4e80
             This bridge  is the root
             Hello Time    2 sec Max Age 20 sec Forwar

  Bridge ID Priority    24577 (priority 24576 sys-id-e
            Address     001b.5302.4e80
            Hello Time  2 sec Max Age 20 sec Forward D
            Aging Time  300

Interface        Role  Sts  Cost      Prio.Nbr   Type
---------------- ----- ---- --------- -------- -------
Fa0/1            Desg  FWD  19        128.1      P2p
Fa0/2            Desg  FWD  19        128.2      P2p
```

Notice in Example 25-2 that the BID is not the default 24576 that was configured on S1 with the **spanning-tree vlan 1 priority 24576** command.

This is because the value of the priority is extended to include the VLAN ID. Therefore, a priority of 24576 plus a VLAN of 1 results in a priority output of 24577.

Configuring PortFast and BPDU Guard

To speed convergence for access ports when they become active, you can use Cisco's proprietary PortFast technology. After PortFast is configured and a port is activated, the port immediately transitions from the blocking state to the forwarding state. This immediate transition is useful for devices that do not participate in STP and helps to minimize network disruption during device startups or reconnections.

In a valid PortFast configuration, BPDUs should never be received because receipt of a BPDU indicates that another bridge or switch is connected to the port, potentially causing a spanning tree loop. When it is enabled, BPDU Guard puts the port in an errdisabled (error-disabled) state upon receipt of a BPDU. This effectively shuts down the port. The BPDU Guard feature provides a secure response to invalid configurations because you must manually put the interface back into service.

Example 25-3 shows the interface commands to configure PortFast and BPDU Guard on S2 in Figure 25-8.

Example 25-3 *Configuring PortFast and BPDU Guard*

Technet24

```
S2# configure terminal
Enter configuration commands, one per line. End with C
S2(config)# interface range f0/11 - f0/13
S2(config-if-range)# switchport mode access
S2(config-if-range)# spanning-tree portfast
S2(config-if-range)# spanning-tree bpduguard enable
```

Alternatively, you can skip the **interface range** command and directly configure the global commands **spanning-tree portfast default** and **spanning-tree bpduguard default**, which enable PortFast and BPDU Guard on all access ports.

Configuring BPDU Filter

Instead of using BPDU Guard, you can use a feature called BPDU Filter to suppress the sending and receiving of BPDUs on a port. When enabled globally, BPDU Filter prevents a PortFast-enabled port from sending BPDUs. If a BPDU is received, PortFast is disabled, and the port starts participating in STP. When configured at the interface level, BPDU Filter prevents the port from both sending and receiving BPDUs, regardless of PortFast.

In general, use BPDU Guard on ports where you want to ensure that no BPDUs are received and where immediate shutdown of the port upon receipt of a BPDU is desired. This is typically used on access ports connected to end devices. Use BPDU Filter when you need to suppress BPDU traffic entirely.

This might be necessary in more controlled environments or specific topologies where you know that BPDU traffic should not exist. For example, in data centers administrators might use BPDU Filter to suppress BPDU traffic. This helps simplify the network topology and prevent unintended participation in STP by devices that do not need to be involved. BPDU Filter ensures that specific ports neither send nor receive BPDUs, maintaining a streamlined and optimized network.

Note

BPDU Guard and BPDU Filter should not be used on the same interface because their functionalities can overlap and potentially cause confusion or unexpected behavior in your network.

To configure BPDU Filter globally for all PortFast interfaces, reverse the default **no spanning-tree portfast bpdufilter default** command by entering **spanning-tree portfast bpdufilter default**. To disable BDPU Filter on an interface, enter the command **spanning-tree bpdufilter disable**. This interface will participate in normal STP operations. To disable STP on an interface, use the command **spanning-tree bpdufilter enable**. Be sure to exercise great care when enabling BPDU Filter on an interface because you are effectively shutting down STP. That one command can create a loop on a LAN that will make it unusable in another switch that is attached to that interface.

Technet24

Configuring Root Guard

To avoid the potential of suboptimal paths, Root Guard is used on ports that are connected to switches that should never become root. It is configured on the interface with the **spanning-tree guard root** command. This feature ensures that a designated switch remains the root bridge, preventing topology changes that could occur if an unauthorized or misconfigured switch tries to become the root bridge.

When Root Guard is enabled on a port, it monitors incoming BPDUs. As shown in Example 25-4, if a superior BPDU (indicating a switch that has a lower bridge ID and is trying to become the root bridge) is received on a Root Guard-enabled port, the port is placed into a "root-inconsistent" state. This state prevents the port from forwarding traffic but allows it to receive BPDUs. Once the superior BPDU ceases, the port automatically recovers and returns to its normal state.

Example 25-4 *Configuring and Verifying Root Guard*

```
S1(config)# interface g1/0/1
S1(config-if)# spanning-tree guard root
S1(config-if)# end

S1# show spanning-tree vlan 1 int g1/0/1
```

```
Vlan            Role Sts Cost      Prio.Nbr  Type
--------------- ---- --- ------- --------  ------------
VLAN0001        Desg FWD 20000     128.1      P2p

S1#! A rogue switch is attached to G1/0/1 and sends a

*May 23 18:39:20.418: %SPANTREE-2-ROOTGUARD_BLOCK: Roo
GigabitEthernet1/0/1 on VLAN0001.

S1# show spanning-tree vlan 1 int g1/0/1

Vlan            Role Sts   Cost       Prio.Nbr  Type
--------------- ---- ---   --------- --------  --------
VLAN0001        Desg  BKN* 20000      128.1      P2p *RO

S1#! The rogue switch in disconnected from G1/0/1

*May 23 19:43:25.452: %SPANTREE-2-ROOTGUARD_UNBLOCK: R
GigabitEthernet1/0/1 on VLAN0001.

S1# show spanning-tree vlan 1 int g1/0/1

Vlan            Role Sts Cost      Prio.Nbr   Type
--------------- ----  --- ------- ---------- ------
VLAN0001        Desg FWD 20000     128.1      P2p
S1#
```

Technet24

Configuring Loop Guard

Loop Guard is used to prevent network loops in situations where a blocking port erroneously transitions to the forwarding state, potentially due to the absence of BPDUs. Loop Guard ensures that if a port should be receiving BPDUs and it suddenly stops receiving them, the port does not transition to the forwarding state. In this way, Loop Guard prevents possible loops.

By ensuring that designated and root ports remain in the blocking state when BPDUs are not received, Loop Guard enhances the redundancy and reliability of the network. It complements other STP features like Root Guard, providing additional protection in complex network topologies with redundant links.

When Loop Guard is enabled on a switch port, it monitors the receipt of BPDUs. If a port configured to receive BPDUs stops receiving them, the port is placed into a “loop-inconsistent” state instead of transitioning to the forwarding state. This state prevents the port from forwarding traffic, thus avoiding a potential loop. The port will remain in this state until BPDUs are detected again, at which point it will automatically recover.

You can enable Loop Guard globally with the **spanning-tree loopguard default** command. You can also configure it on an interface with the **spanning-tree guard loop** command.

Verifying STP

Several commands enable you to verify the state of the current STP implementation. Table 25-9 summarizes commands most likely to appear on the CCNA exam.

Table 25-9 *STP Verification Commands*

Technet24

Description	Command
Displays STP information	`Switch# show spanning-tree`
Displays STP information for active interfaces only	`Switch# show spanning-tree active`
Displays abbreviated information for all STP instances	`Switch# show spanning-tree bridge`
Displays detailed information for all STP instances	`Switch# show spanning-tree detail`
Displays STP information for the specified interface	`Switch# show spanning-tree interface interface-id`
Displays STP information for the specified VLAN	`Switch# show spanning-tree vlan vlan-id`
Displays a summary of STP port states	`Switch# show spanning-tree summary`

Note

Ideally, you should review the output of these commands

today on lab equipment or a simulator. At the very least, refer to the examples in your study resources.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	5
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	5
	9–10
Portable Command Guide	11

Technet24

Day 24. EtherChannel and HSRP

CCNA 200-301 Exam Topics

• 2.4 Configure and verify (Layer 2/Layer 3) EtherChannel (LACP)

• 3.5 Describe the purpose, functions, and concepts of first hop redundancy protocols

Key Topics

EtherChannel technology enables you to bundle multiple physical interfaces into one logical channel to increase the bandwidth on point-to-point links. In addition, EtherChannel provides a way to prevent the need for Spanning Tree Protocol (STP) convergence when only a single port or cable failure occurs.

Most end devices do not store routes to reach remote networks. Instead, an end device is typically configured with a default gateway that handles routing for the device. But what if that default gateway fails? To ensure that a device will still have access to remote networks, you should implement some type of default gateway redundancy in the network. That is the role of first-hop redundancy protocols (FHRPs).

EtherChannel Operation

EtherChannel, a technology that Cisco developed, can bundle up to eight equal-speed links between two switches, as you can see between the two

distribution layer switches in Figure 24-1.

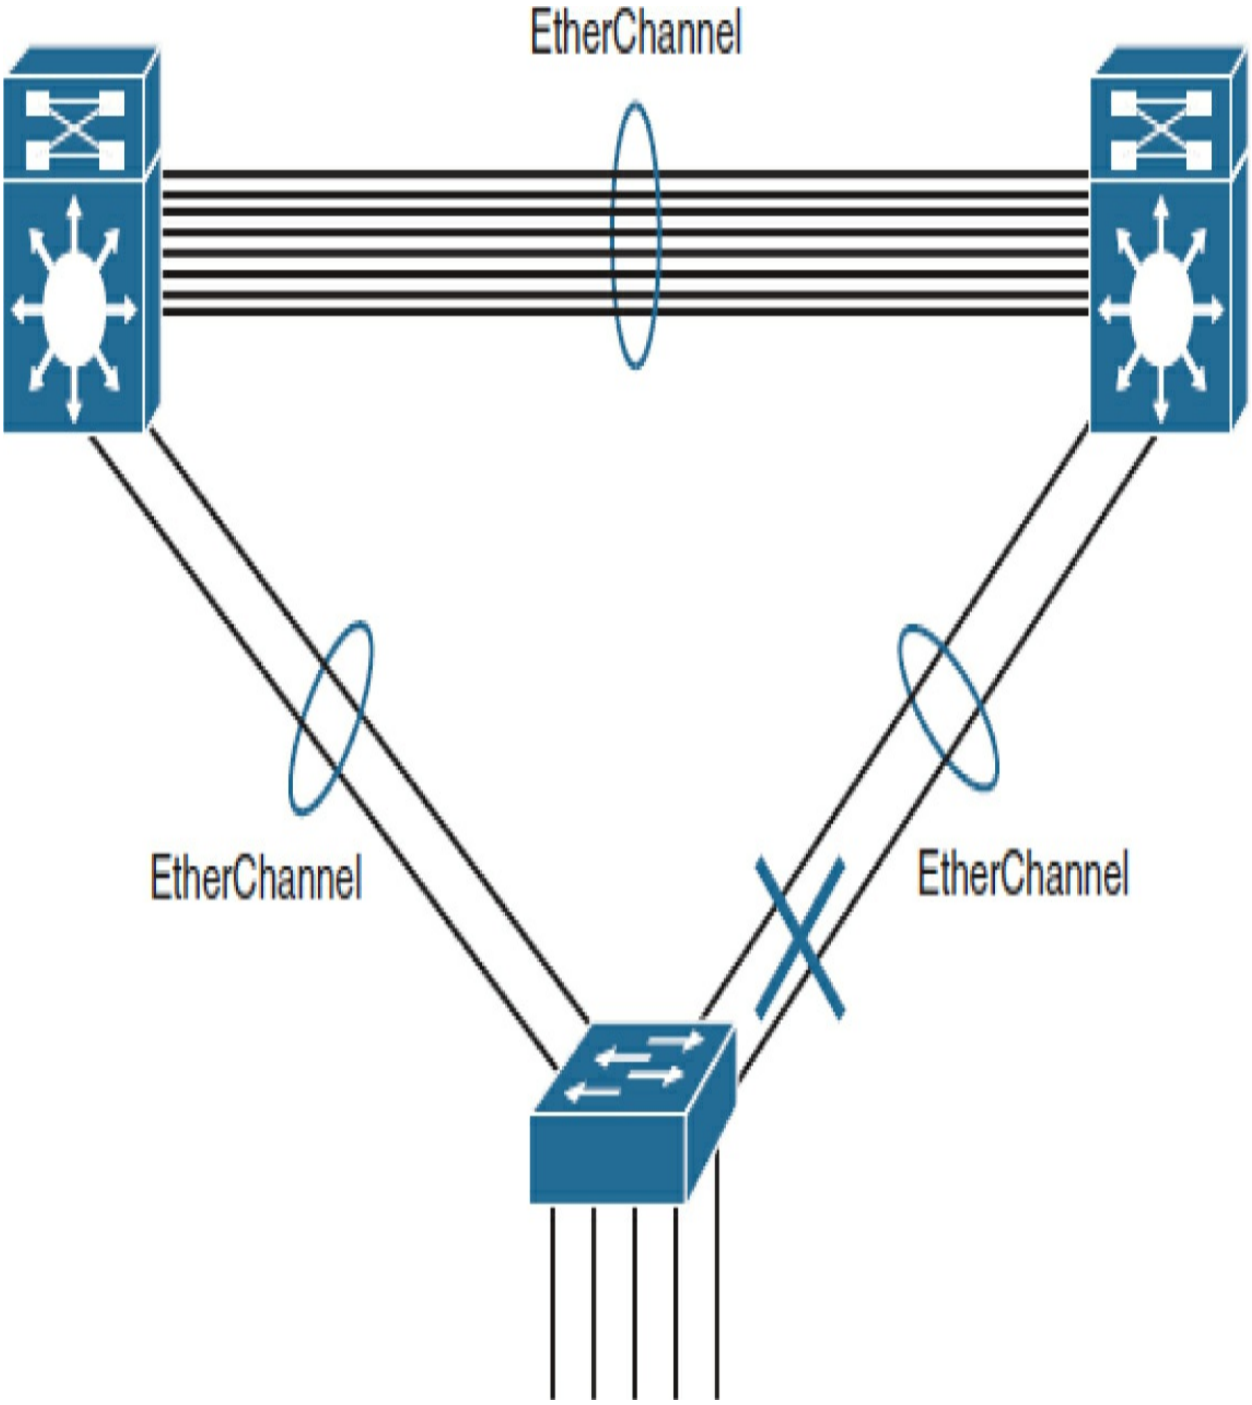

Figure 24-1 *Sample EtherChannel Topology*

STP sees the bundle of links as a single interface. As a result, if at least one

Technet24

of the links is up, STP convergence does not have to occur. This makes much better use of available bandwidth while reducing the number of times STP must converge. Without the use of EtherChannel or modification of the STP configuration, STP would block all the links except one.

Benefits of EtherChannel

When EtherChannel is configured, the resulting virtual interface is called a *port channel*. The physical interfaces are bundled together into a port channel interface. EtherChannel has the following benefits:

• Most configuration tasks can be done on the EtherChannel interface instead of on each individual port, thus ensuring configuration consistency throughout the links.

• EtherChannel relies on the existing switch ports to increase bandwidth. No hardware upgrades are needed.

• Load balancing is possible between links that are part of the same EtherChannel. (Load balancing configuration is beyond the scope of the CCNA exam.)

• EtherChannel creates an aggregation that STP recognizes as one logical link.

• EtherChannel provides redundancy. The loss of one physical link does not create a change in the topology.

Implementation Restrictions

Keep in mind a few limitations when implementing EtherChannel on Cisco 2960 Catalyst switches:

• Interface types, such as Fast Ethernet and Gigabit Ethernet, cannot be mixed within the same EtherChannel.

• Each EtherChannel can consist of up to eight compatibly configured Ethernet ports.

• Cisco IOS Software currently supports up to six EtherChannels.

• Some servers also support EtherChannel to the switch to increase bandwidth; however, the server then needs at least two EtherChannels to provide redundancy because it can send traffic to only one switch through the EtherChannel.

• The EtherChannel configuration must be consistent on the two switches. The trunking configuration (native VLAN, allowed VLANs, and so on) must be the same. All ports also must be Layer 2 ports.

• All ports in the EtherChannel must be Layer 2 ports, or all ports within the EtherChannel must be Layer 3 ports.

Note

Technet24

You can configure Layer 3 EtherChannels on multilayer switches; however, that is beyond the scope of the CCNA exam.

EtherChannel Protocols

You can configure EtherChannel as static or unconditional; however, you also can use two protocols to configure the negotiation process: Port Aggregation Protocol (PAgP, which is Cisco proprietary) and Link Aggregation Control Protocol (LACP, which is IEEE 802.3ad). These two protocols ensure that the two sides of the link have compatible configurations—same speed, duplex setting, and VLAN information. The modes for each differ slightly.

Port Aggregation Protocol

PAgP is a Cisco-proprietary protocol that aids in the automatic creation of EtherChannel links. PAgP checks for configuration consistency and manages link additions and failures between two switches. It ensures that when an EtherChannel is created, all ports have the same type of configuration. PAgP uses the following modes:

• **On:** This mode forces the interface to channel without PAgP.

• **Desirable:** The interface initiates negotiations with other interfaces by sending PAgP packets.

• **Auto:** The interface responds to the PAgP packets that it receives but does not initiate PAgP negotiation.

The modes must be compatible on the two sides of the EtherChannel. For example, Sw1 and Sw2 in Figure 24-2 must be configured with a particular combination of settings, as shown in Table 24-1.

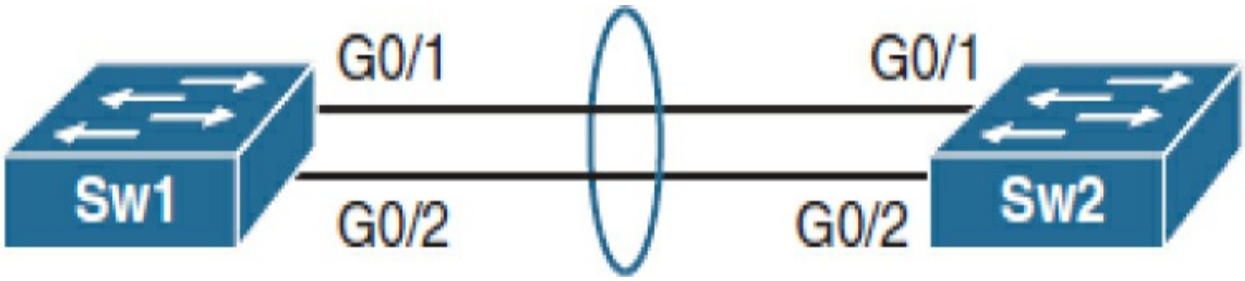

Figure 24-2 *Two-Switch EtherChannel Topology*

Table 24-1 *PAgP Mode Settings*

Technet24

Sw1	Sw2	Channel Established?
On	On	Yes
Auto/Desirable	Desirable	Yes
On/Auto/Desirable	Not configured	No
On	Desirable	No
Auto/On	Auto	No

Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a function similar to PAgP with Cisco EtherChannel. Cisco devices support both PAgP and LACP. LACP uses the following modes:

• **On:** This mode forces the interface to channel without LACP.

• **Active:** The interface initiates negotiations with other interfaces by sending LACP packets.

• **Passive:** The interface responds to the LACP packets that it receives but does not initiate LACP negotiation.

As with PAgP, the LACP modes must be compatible on the two sides of the EtherChannel. For example, Sw1 and Sw2 in Figure 24-2 must be configured with a particular combination of settings, as shown in Table 24-2.

Table 24-2 *LACP Mode Settings*

Sw1	Sw2	Channel Established?
On	On	Yes
Active/Passive	Active	Yes
On/Active/Passive	Not configured	No
On	Active	No
Passive/On	Passive	No

Note

For both the PAgP and LACP protocols, the on mode creates

Technet24

the EtherChannel configuration unconditionally, without PAgP or LACP dynamic negotiation. You should probably memorize the mode settings for both PAgP and LACP in preparation for the CCNA exam.

Configuring EtherChannel

To implement EtherChannel, follow these steps:

Specify the interfaces that you want to bundle together in one link by using the **interface range** *interfaces* command.

Create a port channel by using the **channel-group** *identifier* **mode** *mode* command. *identifier* can be any number between 1 and 6, inclusive, and does not have to match the other switch. The mode is either **on** or one of the PAgP or LACP modes.

Enter interface configuration mode for the new port channel with the **interface port-channel** *identifier* command. *identifier* is the same number used with the **channel-group** command.

Configure the trunking and VLAN settings.

Using the topology in Figure 24-2, assume that Sw1 is already configured for EtherChannel with G0/1 and G0/2 trunking. The native VLAN is 86. The allowed VLANs are 1, 10, 20, and 86. EtherChannel is forced on. No PAgP

or LACP is needed. Example 24-1 shows the configuration for Sw2.

Example 24-1 *EtherChannel Configuration*

```
Sw2(config)# interface range g0/1-2
Sw2(config-if-range)# channel-group 1 mode on
Creating a port-channel interface Port-channel 1
Sw2(config-if-range)# interface port-channel 1
Sw2(config-if)# switchport mode trunk
Sw2(config-if)# switchport trunk native vlan 86
Sw2(config-if)# switchport trunk allowed vlan 1,10,20,
```

In configuring PAgP or LACP, use the appropriate mode keyword for the **channel-group** command. Just ensure that the commands on both sides of the channel are compatible, according to Tables 24-1 and 24-2.

Verifying EtherChannel

If you configured management addressing, you can quickly verify both sides of an EtherChannel bundle by pinging across the trunk. The two switches should be able to ping each other. Devices configured as members of the various VLANs also should be able to ping each other.

To verify the configuration, use the **show run** command (see Example 24-2).

Example 24-2 *Verifying the EtherChannel Configuration*

Technet24

```
Sw2# show run | begin interface Port
interface Port-channel1
  switchport trunk native vlan 86
  switchport trunk allowed vlan 1,10,20,86
  switchport mode trunk
!
<output omitted>
interface GigabitEthernet0/1
  switchport trunk native vlan 86
  switchport trunk allowed vlan 1,10,20,86
  switchport mode trunk
  channel-group 1 mode on
!
interface GigabitEthernet0/2
  switchport trunk native vlan 86
  switchport trunk allowed vlan 1,10,20,86
  switchport mode trunk
  channel-group 1 mode on
```

To get an overall summary of the EtherChannel configuration, use the **show etherchannel summary** command (see Example 24-3).

Example 24-3 *Verifying That EtherChannel Is Operational*

```
Sw2# show etherchannel summary
Flags: D -  down         P - bundled in port-channel
       I -  stand-alone  s - suspended
```

```
        H -  Hot-standby  (LACP only)
        R -  Layer3        S - Layer2
        U -  in use        f - failed to allocate aggre

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port


Number of channel-groups in use:  1
Number of aggregators:            1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
1      Po1(SU)             -     Gig0/1(P) Gig0/2(P)
```

To verify the operational status of a specific interface in the EtherChannel bundle, use the **show interface switchport** command (see Example 24-4).

Example 24-4 *Verifying an Interface's Port Channel Settings*

```
Sw2# show interface fa0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk (member of bundle Po1)
```

Technet24

```
  Administrative Trunking Encapsulation: dot1q
  Operational Trunking Encapsulation: dot1q
  Negotiation of Trunking: On
  Access Mode VLAN: 1 (default)
  Trunking Native Mode VLAN: 86 (VLAN0086)
  Administrative Native VLAN tagging: enabled
  Voice VLAN: none
  Administrative private-vlan host-association: none
  Administrative private-vlan mapping: none
  Administrative private-vlan trunk native VLAN: none
  Administrative private-vlan trunk Native VLAN tagging:
  Administrative private-vlan trunk encapsulation: dot1q
  Administrative private-vlan trunk normal VLANs: none
  Administrative private-vlan trunk associations: none
  Administrative private-vlan trunk mappings: none
  Operational private-vlan: none
  Trunking VLANs Enabled: 1,10,20,86
  Pruning VLANs Enabled: 2-1001
```

Troubleshooting EtherChannel

All interfaces within an EtherChannel must have the same configuration of speed for the duplex mode, native and allowed VLANs on trunks, and access VLAN on access ports:

• Assign all ports in the EtherChannel to the same VLAN or configure them as trunks. Ports with different native VLANs cannot form an EtherChannel.

• When configuring a trunk on an EtherChannel, verify the trunking mode on the EtherChannel. Configuring trunking mode on individual ports that make up the EtherChannel is not recommended. However, if it is done, verify that the trunking configuration is the same on all interfaces.

• An EtherChannel supports the same allowed range of VLANs on all the ports. If the allowed range of VLANs is not the same, the ports do not form an EtherChannel even when PAgP is set to **auto** or **desirable** mode.

• The dynamic negotiation options for PAgP and LACP must be compatibly configured on both ends of the EtherChannel.

Configuration issues with the **channel-group** command include the following:

• Configuring the **on** keyword on one switch and **desirable**, **auto**, **active**, or **passive** on the other switch. The **on** keyword does not enable PAgP or LACP. Both switches should be configured on one of the acceptable PAgP or LACP modes.

• Configuring the **auto** keyword on both switches. This enables PAgP, but each switch waits on the other to begin negotiations.

• Configuring the **passive** keyword on both switches. This enables LACP, but each switch waits on the other to begin negotiations.

Technet24

• Mixing keywords from PAgP and LACP, which are not compatible (for example, configuring **active** (LACP) on one switch and **desirable** or **auto** (PAgP) on the other switch).

First-Hop Redundancy Concepts

Networks require redundant links to enhance availability because failures are inevitable. A router's power supply might fail, a link could break, or a switch might lose power. WAN links are actually complex systems with numerous components that can fail.

A single point of failure in a network can lead to outages affecting users. For example, in Figure 24-3, while LANs show some redundancies among the switches, the WAN does not. If traffic primarily flows between sites, multiple single points of failure exist. Figure 24-3 shows seven single points of failure. Note that the end users have one default gateway, at 10.1.1.9.

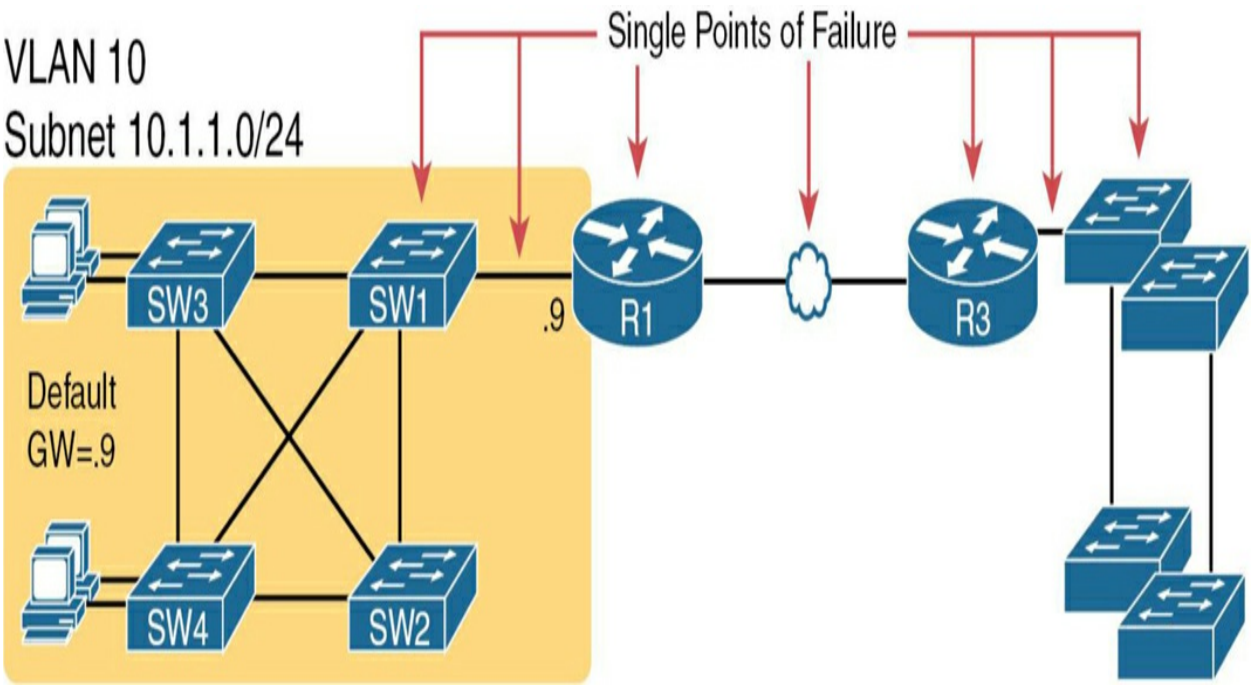

Figure 24-3 *Multiple Single Points of Failure*

To enhance network availability, engineers identify and address these single points of failure. They add redundant devices and links and implement functions that utilize these redundancies to ensure continued network operation. Figure 24-4 shows redundant routers, switches, and WAN links, which would be the ideal setup for an organization that must protect against single points of failure. Note that the end users now have two default gateways. R1 is serving as the primary gateway at 10.1.1.9, and R2 is the backup gateway at 10.1.1.8.

Technet24

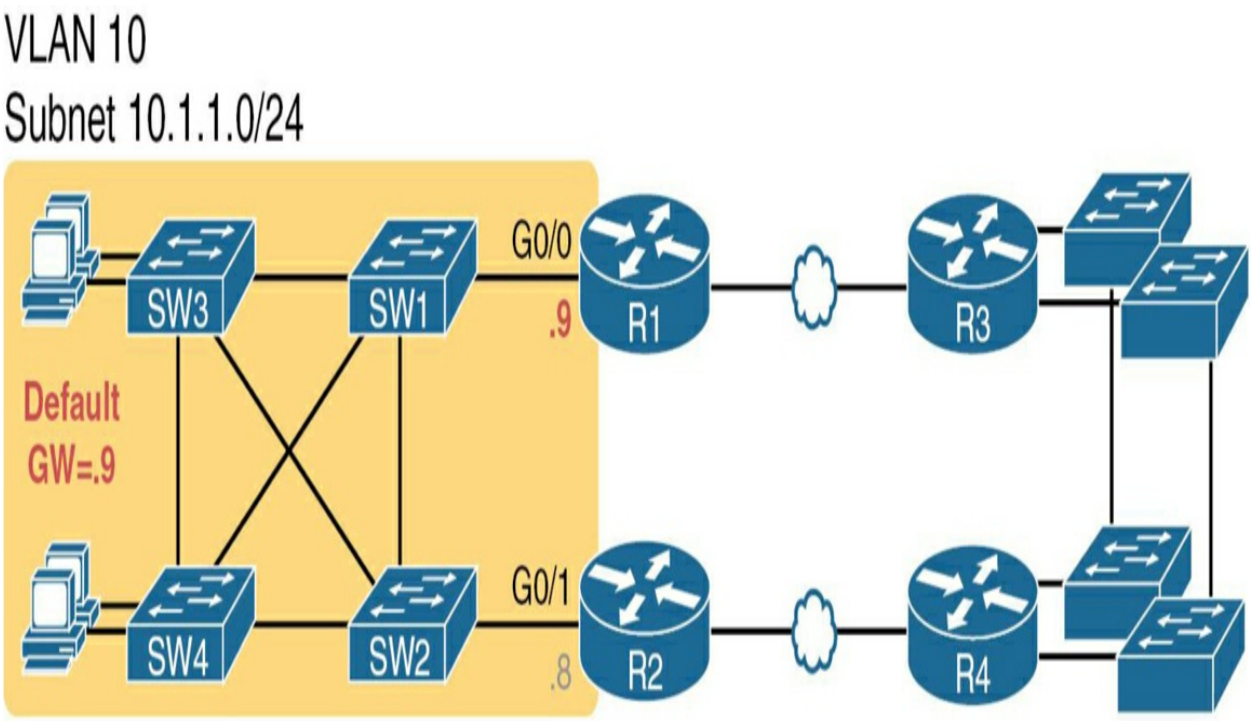

Figure 24-4 *Fully Redundant Network*

FHRPs enable you to install multiple routers in a subnet to collectively act as a single default router. These routers share a virtual IP address, as Figure 24-5 shows.

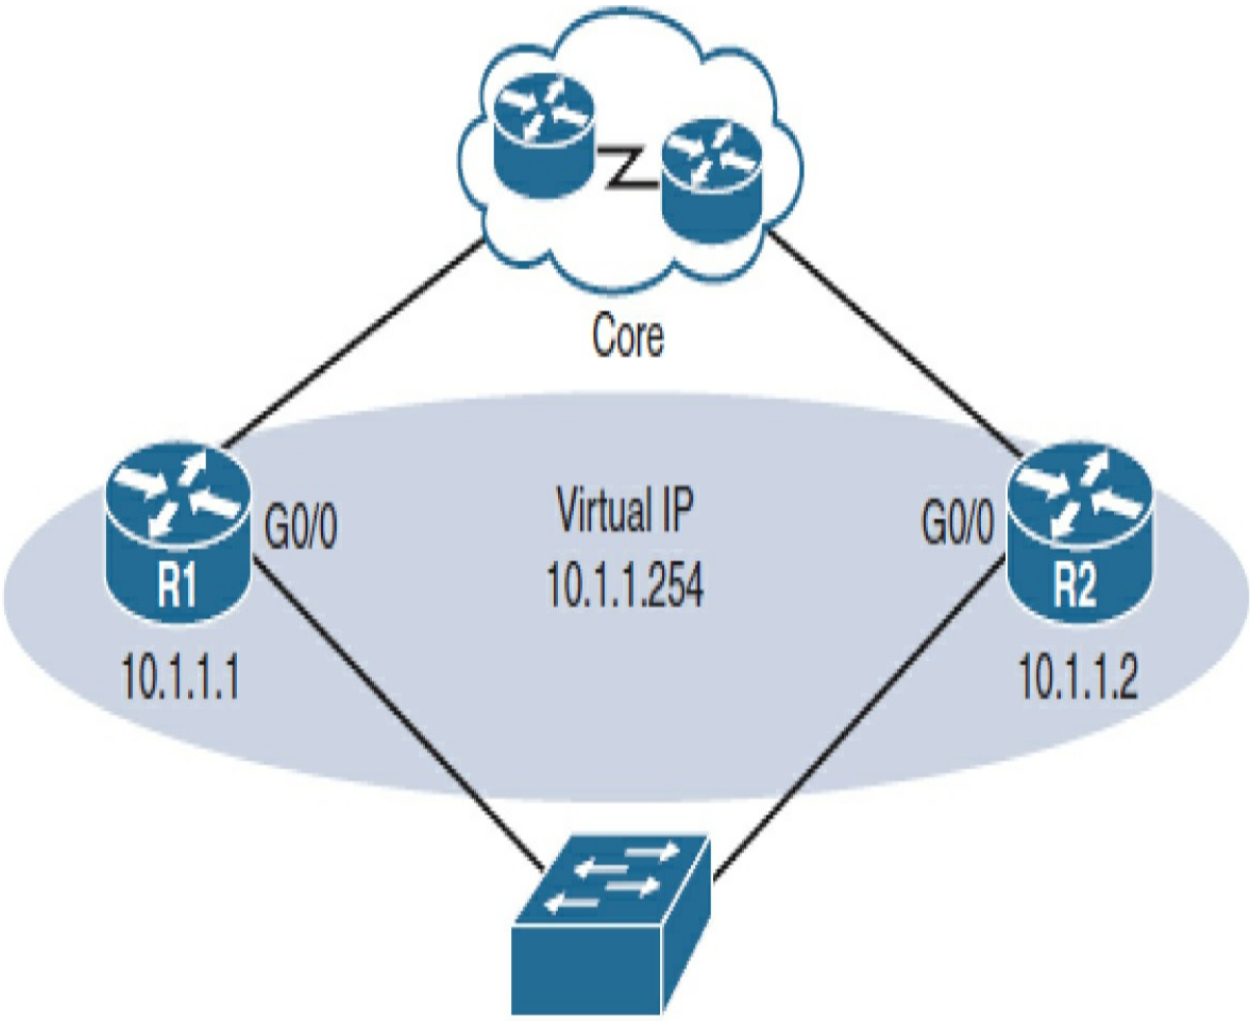

Figure 24-5 *Redundant Default Gateway Example*

In the figure, the G0/0 interfaces on R1 and R2 are configured with the IP addresses shown. However, both routers are also configured with the virtual IP address. This virtual IP address is the default gateway address configured on end devices. A redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic. It also determines when a standby router must take over the forwarding role. The transition from one forwarding router to another is transparent to the end devices. This capability of a network to dynamically recover from the failure

Technet24

of a device acting as a default gateway is known as *first-hop redundancy*.

Regardless of which FHRP is implemented, the following steps take place when the active router fails:

The standby router stops seeing hello messages from the forwarding router.

The standby router assumes the role of the forwarding router.

Because the new forwarding router assumes both the IP and MAC addresses of the virtual router, the end stations do not recognize a disruption in service.

FHRPs

The following list defines the three options available for FHRPs:

• **Hot Standby Router Protocol (HSRP):** A Cisco-proprietary FHRP designed to allow for transparent failover of a first-hop IPv4 device. The function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router fails. HSRP for IPv6 provides support for IPv6 networks.

• **Virtual Router Redundancy Protocol (VRRP):** An IETF standard that dynamically assigns responsibility for one or more virtual routers to the VRRP routers on an IPv4 LAN. Its operation is similar to that of HSRP. VRRPv3 supports IPv4 and IPv6.

• **Gateway Load Balancing Protocol (GLBP):** A Cisco-proprietary FHRP that protects data traffic from a failed router or circuit, as in HSRP and VRRP, while also allowing load balancing (also called load sharing) between a group of redundant routers. GLBP for IPv6 provides support for IPv6 networks.

The CCNA exam covers HSRP.

HSRP Operation

HSRP uses an active/standby model in which one router actively assumes the role of default gateway for devices on the subnet. One or more routers on the same subnet are then in standby mode. The HSRP active router implements a virtual IP address and matching virtual MAC address. This virtual IP address is part of the HSRP configuration and belongs to the same subnet as the physical interface IP address, but it is a different IP address. The router then automatically creates the virtual MAC address. All the cooperating HSRP routers know these virtual addresses, but only the HSRP active router uses these addresses at any one point in time.

Assume that you have two HSRP routers similar to R1 and R2 in Figure 24-5. These HSRP routers send each other messages to negotiate which router should be active. Then they continue to send each other messages so that the standby router can detect when the active router fails. If the active router fails, the standby router automatically assumes the virtual IP and MAC

Technet24

addresses and serves as the default gateway for the LAN. The new active router then sends out a gratuitous ARP so that the switches on the subnet will change their MAC address tables to reflect the correct port to reach the virtual MAC. This failover process is transparent to end devices, which are all configured with the virtual IP address as the default gateway.

So what about load balancing? Aren't we wasting the capacity of the standby router and the links connecting to it? Yes, if the routers are connected to only one subnet. However, if VLANs are configured, the routers can share the load by each serving as the active router for some of the VLANs. For example, in Figure 24-5, R1 is the active router for VLAN 10, and R2 is the active router for VLAN 20. Both routers are configured with subinterfaces for inter-VLAN routing and the two virtual IP addresses so that each can assume the role of active router if the other router fails.

HSRP Versions

Cisco IOS defaults to HSRP version 1. Table 24-3 compares HSRP version 1 and version 2.

Table 24-3 *HSRP Version 1 and Version 2 Features*

HSRP Feature	Version 1	Version 2
Group numbers supported	0–255	0–4095
Authentication	None	MD5
Multicast addresses	IPv4: 224.0.0.2	IPv4: 224.0.0.102 IPv6: FF02::66
Virtual MAC ranges	0000.0C07.AC00 to 0000.0C07.ACFF	IPv4: 0000.0C9F.F000 to 0000.0C9F.FFFF IPv6: 0005.73A0.0000 to 0005.73A0.0FFF

Note

The last three hexadecimal digits of the virtual MAC address indicate the configured group number. Group numbers are important for more advanced HSRP configurations, which are beyond the scope of the CCNA exam.

HSRP Priority and Preemption

Technet24

By default, the router with the numerically highest IPv4 address is elected as the active HSRP router. To configure a router to be the active router, regardless of IPv4 addressing, use the **standby priority** interface configuration command. The default priority is 100. The router with the highest priority will be the active HSRP router, assuming that no election has already occurred.

To force a new HSRP election, preemption must be enabled with the **standby preempt** interface configuration command.

HSRP Configuration and Verification

Let's look at how to configure the topology in Figure 24-5. HSRP requires only one command on both routers:

```
Router(config-if)# standby group ip ip-address
```

The interface must be on the same subnet as the other HSRP router or routers. The *group* number and virtual *ip-address* must be the same on all HSRP routers.

Unless the **priority** command is used, the first router configured becomes the HSRP active router. Therefore, even though in Example 24-5 R1 is configured first, it includes a priority configuration to make sure that R1 is always the active router. Also, to make sure that R1 resumes the active router

role after losing connectivity, the **standby preempt** command is configured.

Example 24-5 *Configuring HSRP*

```
R1(config)# interface g0/0
R1(config-if)# ip address 10.1.1.1 255.255.0.0
R1(config-if)# standby 1 ip 10.1.1.254
R1(config-if)# standby 1 priority 200
R1(config-if)# standby 1 preempt
R2(config)# interface g0/0
R2(config-if)# ip address 10.1.1.2 255.255.0.0
R2(config-if)# standby 1 ip 10.1.1.254
```

To verify that HSRP is up and running, use the **show standby** command or the **brief** version of the command, as in Example 24-6.

Example 24-6 *Verifying HSRP*

```
R1# show standby
GigabitEthernet0/0 - Group 1
 State is Active
    2 state changes, last state change 00:11:51
 Virtual IP address is 10.1.1.254
 Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (v1 defa
 Hello time 3 sec, hold time 10 sec
  Next hello sent in 1.232 secs
```

Technet24

```
 Preemption enabled
 Active router is local
 Standby router is 10.1.1.2, priority 100 (expires in
 Priority 200 (configured 200)
 Group name is "hsrp-Gi0/0-1" (default)
R1# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri  P State    Active  Standby   Virt
Gi0/0         1   200   Active   local   10.1.1.2  10.1
```

```
R2# show standby
GigabitEthernet0/0 - Group 1
 State is Standby
  1 state change, last state change 00:15:23
 Virtual IP address is 10.1.1.254
 Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (v1 defa
 Hello time 3 sec, hold time 10 sec
  Next hello sent in 1.008 secs
 Preemption disabled
 Active router is 10.1.1.1, priority 200 (expires in 8
 Standby router is local
 Priority 100 (default 100)
 Group name is "hsrp-Gi0/0-1" (default)
R2# show standby brief
                   P indicates configured to preempt.
```

```
                     |
  Interface   Grp Pri  P  State     Active    Standby  V
  Gi0/0         1  100     Standby   10.1.1.1  local
```

The **show standby brief** command displays the most pertinent information you might need in a few lines of output. The more verbose **show standby** command provides additional information, such as the number of state changes, the virtual MAC address, hellos, and the group name.

HSRP Load Balancing

As with STP, you might want your HSRP routers to be configured in active/active state, with one router active for one set of VLANs and the other router active for the remaining VLANs. Figure 24-6 shows a topology with multiple VLANs.

Technet24

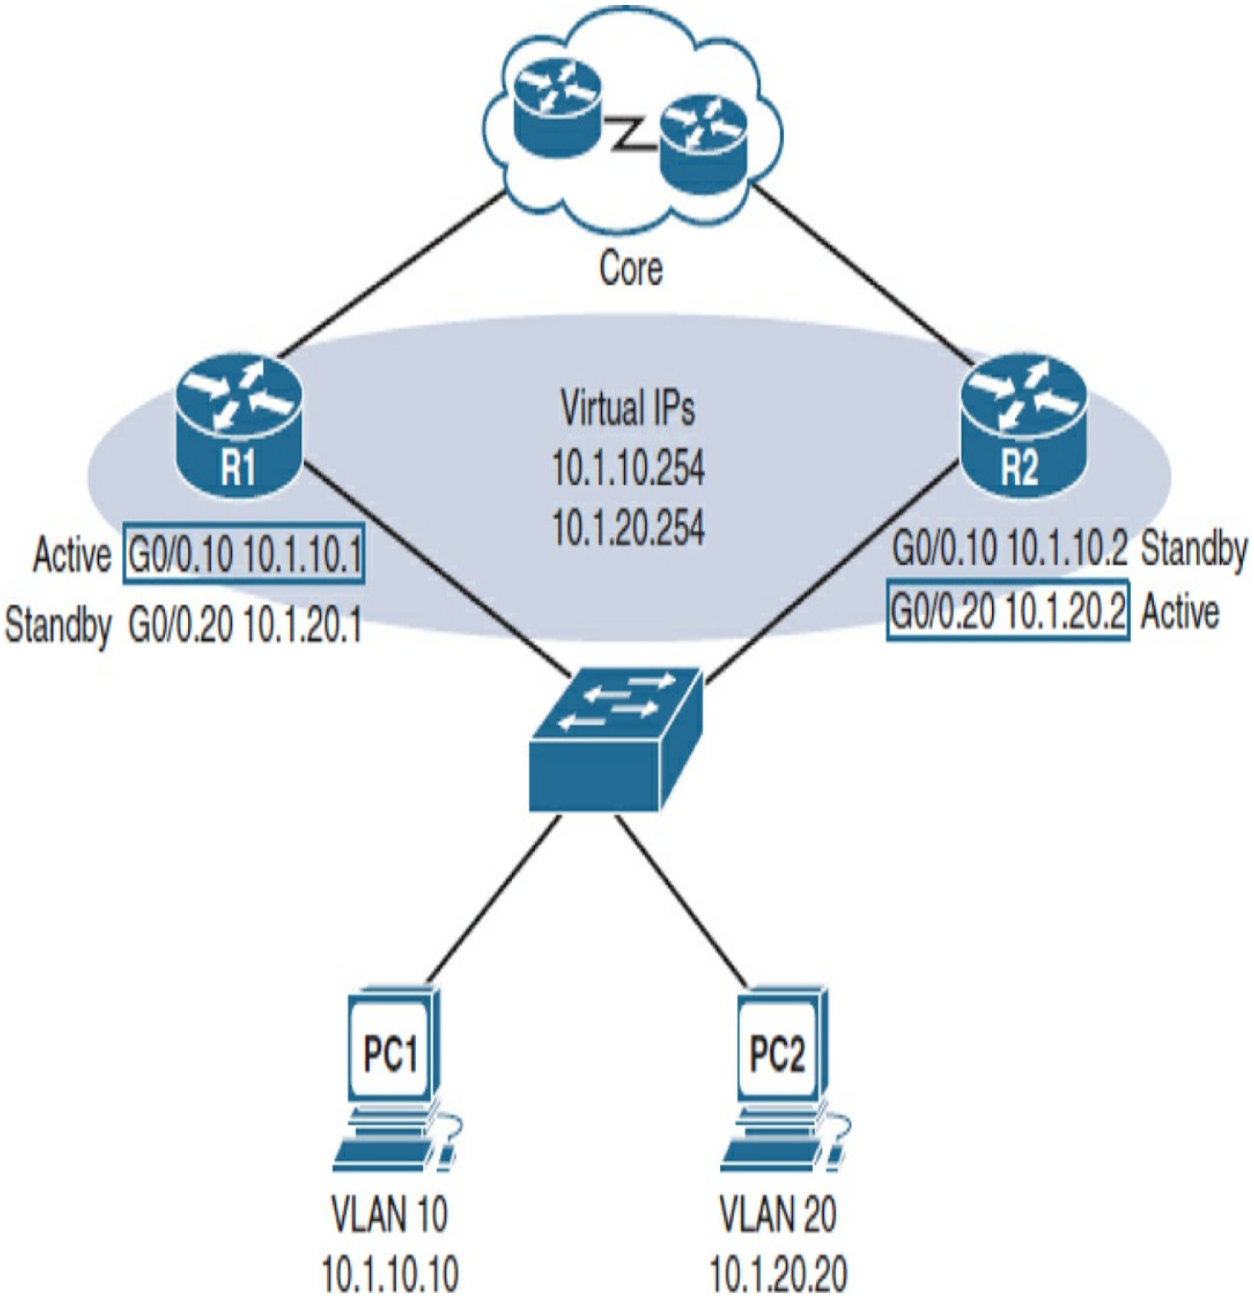

Figure 24-6 *HSRP Load Balancing Example*

To implement HSRP load balancing for different VLANs, configure R1 as the active router for half the VLANs and R2 as the active router for the other half of the VLANs (see Example 24-7).

Example 24-7 *Configuring HSRP Load Balancing*

```
R1# show run | begin interface G

interface GigabitEthernet0/0
  no ip address
  duplex auto
  speed auto
!
interface GigabitEthernet0/0.10
  encapsulation dot1Q 10
  ip address 10.1.10.1 255.255.255.0
  standby version 2
  standby 1 ip 10.1.10.254
  standby 1 priority 150
  standby 1 preempt
!
interface GigabitEthernet0/0.20
  encapsulation dot1Q 20
  ip address 10.1.20.1 255.255.255.0
  standby version 2
  standby 1 ip 10.1.20.254
```

```
R2# show run | begin interface G

interface GigabitEthernet0/0
  no ip address
```

Technet24

```
  duplex auto
  speed auto
!
interface GigabitEthernet0/0.10
  encapsulation dot1Q 10
  ip address 10.1.10.2 255.255.255.0
  standby version 2
  standby 1 ip 10.1.10.254
!
interface GigabitEthernet0/0.20
  encapsulation dot1Q 20
  ip address 10.1.20.2 255.255.255.0
  standby version 2
  standby 1 ip 10.1.20.254
  standby 1 priority 150
  standby 1 preempt
!
```

To verify that HSRP with load balancing is operational, use the **show standby** command or the **brief** version of the command (see Example 24-8).

Example 24-8 *Verifying HSRP Load Balancing*

```
R1# show standby brief
 P indicates configured to preempt.
 |
Interface   Grp  Pri  P  State    Active     Standby
```

```
            1    150     Active   local      10.1. 10.2
            1    100     Standby  10.1.20.2  local
```

```
R2# show standby brief
 P indicates configured to preempt.
 |
Interface   Grp  Pri  P  State    Active       Standby
            1    100     Standby  10.1.10.1    local
            1    150     Active   local        10.1.20
```

Troubleshooting HSRP

Issues with HSRP most likely result from one or more of the following:

• The active router that controls the virtual IP address for the group was not successfully elected.

• The standby router did not successfully keep track of the active router.

• No decision was made regarding when to hand another router control of the virtual IP for the group.

• End devices failed to successfully configure the virtual IP address as the default gateway.

Common HSRP configuration issues include the following:

Technet24

• The HSRP routers are not connected to the same network segment.

• The HSRP routers are not configured with IPv4 addresses from the same subnet.

• The HSRP routers are not configured with the same virtual IPv4 address.

• The HSRP routers are not configured with the same HSRP group number.

• End devices are not configured with the correct default gateway address.

VRRP and GLBP

Virtual Router Redundancy Protocol (VRRP) and Gateway Load Balancing Protocol (GLBP) were both built on the foundational understanding of HSRP. Like HSRP, both VRRP and GLBP aim to provide network redundancy, ensuring continuous availability of the default gateway in a network, but they have distinct characteristics and advantages. In essence, VRRP and GLBP offer robust solutions for network redundancy, with VRRP providing a standardized, HSRP-like protocol and GLBP adding advanced load balancing capabilities to improve overall network performance and reliability. Let’s briefly review each.

VRRP

Developed in response to the lack of a standardized solution for router

redundancy, VRRP was created by the IETF to offer similar functionality to Cisco's proprietary HSRP. VRRP supports the same fundamental functions as HSRP, such as providing a standby router, enabling preemption, and allowing load balancing across multiple VRRP groups. The primary differences lie in protocol specifics, default settings, and addressing methods:

• VRRP can use a virtual IP address (VIP) that matches one of the router's interface addresses, whereas HSRP requires a different IP address within the subnet.

• VRRP uses the multicast address 224.0.0.18 for its communication.

• The virtual MAC address in VRRP follows the pattern 0000.5e00.01*xx*, where *xx* represents the group number in hexadecimal.

• VRRP groups are numbered from 1 to 255, and the protocol supports both IPv4 and IPv6.

GLBP

Cisco's GLBP extends beyond the capabilities of HSRP and VRRP by incorporating advanced load balancing features. GLBP not only provides redundancy but also optimizes network efficiency through load distribution:

• Unlike HSRP and VRRP, GLBP assigns a unique virtual MAC address to each router in a group, facilitating effective load balancing.

Technet24

• GLBP’s load balancing mechanism involves the Active Virtual Gateway (AVG), which manages ARP responses and distributes traffic by alternating the virtual MAC addresses provided in ARP replies. This ensures that different hosts send traffic to different routers, balancing the network load.

• In a GLBP setup, each router can act as an Active Virtual Forwarder (AVF), receiving traffic for its specific virtual MAC address. The AVG also monitors AVF status and reassigns traffic in the event of AVF failure.

• GLBP supports preemption, priority, and tracking much like to HSRP and VRRP, but these features primarily apply to the AVG.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	6
	9
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	9–10
	18
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	16
Portable Command Guide	17

Technet24

Day 23. DHCP and DNS

CCNA 200-301 Exam Topics

• 1.10 Verify IP parameters for Client OS (Windows, Mac OS, Linux)

• 4.3 Explain the role of DHCP and DNS within the network

• 4.6 Configure and verify DHCP client and relay

Key Topics

Imagine that you have to manually configure the IP addressing for every device you want to connect to the network. Furthermore, imagine that you have to type in the IP address for every website you want to visit. Today we review the two protocols that automate this process: Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS). DHCP and DNS make the life of Internet users easier. We also review how to verify IP configuration of end devices for Windows, macOS, and Linux.

DHCPv4

DHCPv4 allows a host to obtain an IP address dynamically when it connects to the network. The DHCPv4 client contacts the DHCPv4 server by sending a request for an IP address. The DHCPv4 server chooses an address from a configured range of addresses called a *pool* and assigns it to the host client for a set period. Figure 23-1 graphically shows the process by which a

DHCPv4 server fulfills a request from a DHCPv4 client.

When a DHCPv4-configured device boots up or connects to the network, the client broadcasts a DHCPDISCOVER packet to identify any available DHCPv4 servers on the network. A DHCPv4 server replies with a DHCPOFFER, which is a lease offer message with an assigned IP address, subnet mask, DNS server, and default gateway information, as well as the duration of the lease.

The client can receive multiple DHCPOFFER packets if the local network has more than one DHCPv4 server. The client chooses the first offer and broadcasts a DHCPREQUEST packet that identifies the explicit server and lease offer that it is accepting.

Assuming that the IP address is still valid, the chosen server returns a DHCPACK (acknowledgment) message, finalizing the lease. If the offer is no longer valid for some reason, the chosen server responds to the client with a DHCPNAK (negative acknowledgment) message. After it is leased, the client renews before the lease expiration through another DHCPREQUEST. If the client is powered down or taken off the network, the address is returned to the pool for reuse.

Technet24

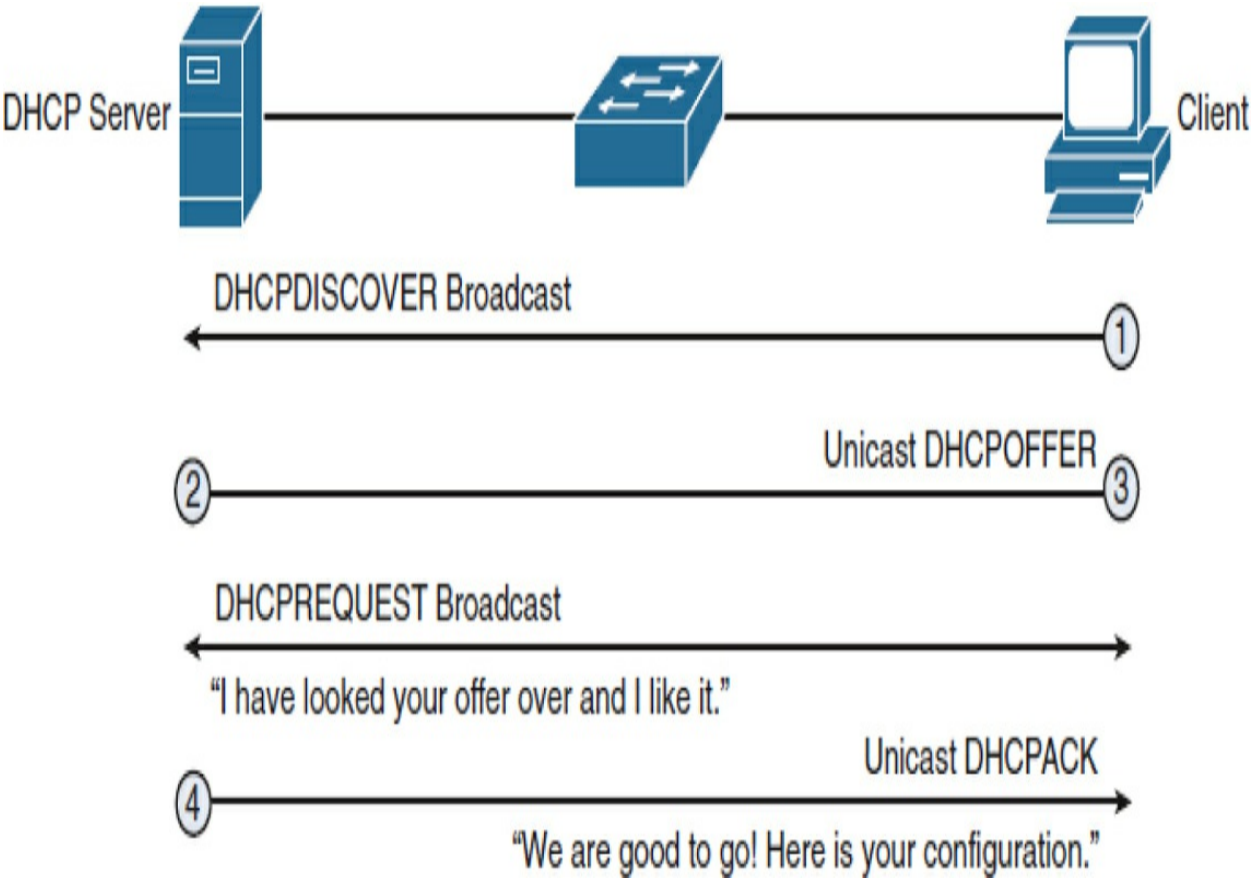

Figure 23-1 *Allocating IP Addressing Information Using DHCPv4*

DHCPv4 Configuration Options

A Cisco router can be configured to handle DHCP requests in two ways: as a DHCP server or as a DHCP relay agent. A Cisco router can also be configured as a DHCP client, requesting an IPv4 address from a DHCP server for one or more of its interfaces. All these options can be configured at the same time on the same device. For example, a router might be the DHCP server for a directly connected LAN while at the same time forwarding DHCP server requests to another DHCP server for other LANs. In addition, the router could have one or more of its interfaces configured to request

DHCP addressing from a remote server.

Configuring a Router as a DHCPv4 Server

A Cisco router running Cisco IOS Software can be configured to act as a DHCPv4 server. The Cisco IOS DHCPv4 server assigns and manages IPv4 addresses from specified address pools within the router to DHCPv4 clients.

The steps to configure a router as a DHCPv4 server follow:

Use the **ip dhcp excluded-address** *low-address* [*high-address*] command to identify an address or range of addresses to exclude from the DHCPv4 pool. For example:

```
R1(config)# ip dhcp excluded-address 192.168.10.1
192.168.10.9
R1(config)# ip dhcp excluded-address 192.168.10.254
```

Create the DHCPv4 pool by using the **ip dhcp pool** *pool-name* command, which places you in DHCP configuration mode:

```
R1(config)# ip dhcp pool LAN-POOL-10
R1(dhcp-config)#
```

Configure the IP addressing parameter you need to automatically assign to requesting clients. Table 23-1 lists the required commands.

Technet24

Table 23-1 *Required DHCPv4 Configuration Commands*

Required Task	Command
Define the address pool	**network** *network-number [mask | /prefix-length]*
Define the default router or gateway	**default-router** *address [address2...address8]*

Table 23-2 lists some of the common optional DHCPv4 tasks.

Table 23-2 *Optional DHCPv4 Configuration Commands*

Optional Task	Command
Define a DNS server	**dns-server** *address* [*address2...address8*]
Define the domain name	**domain-name** *domain*
Define the duration of the DHCPv4 lease	**lease** {*days [hours] [minutes]* | **infinite**}
Define the NetBIOS WINS server	**netbios-name-server** *address [address2...address8]*

Figure 23-2 shows a sample DHCPv4 topology.

Technet24

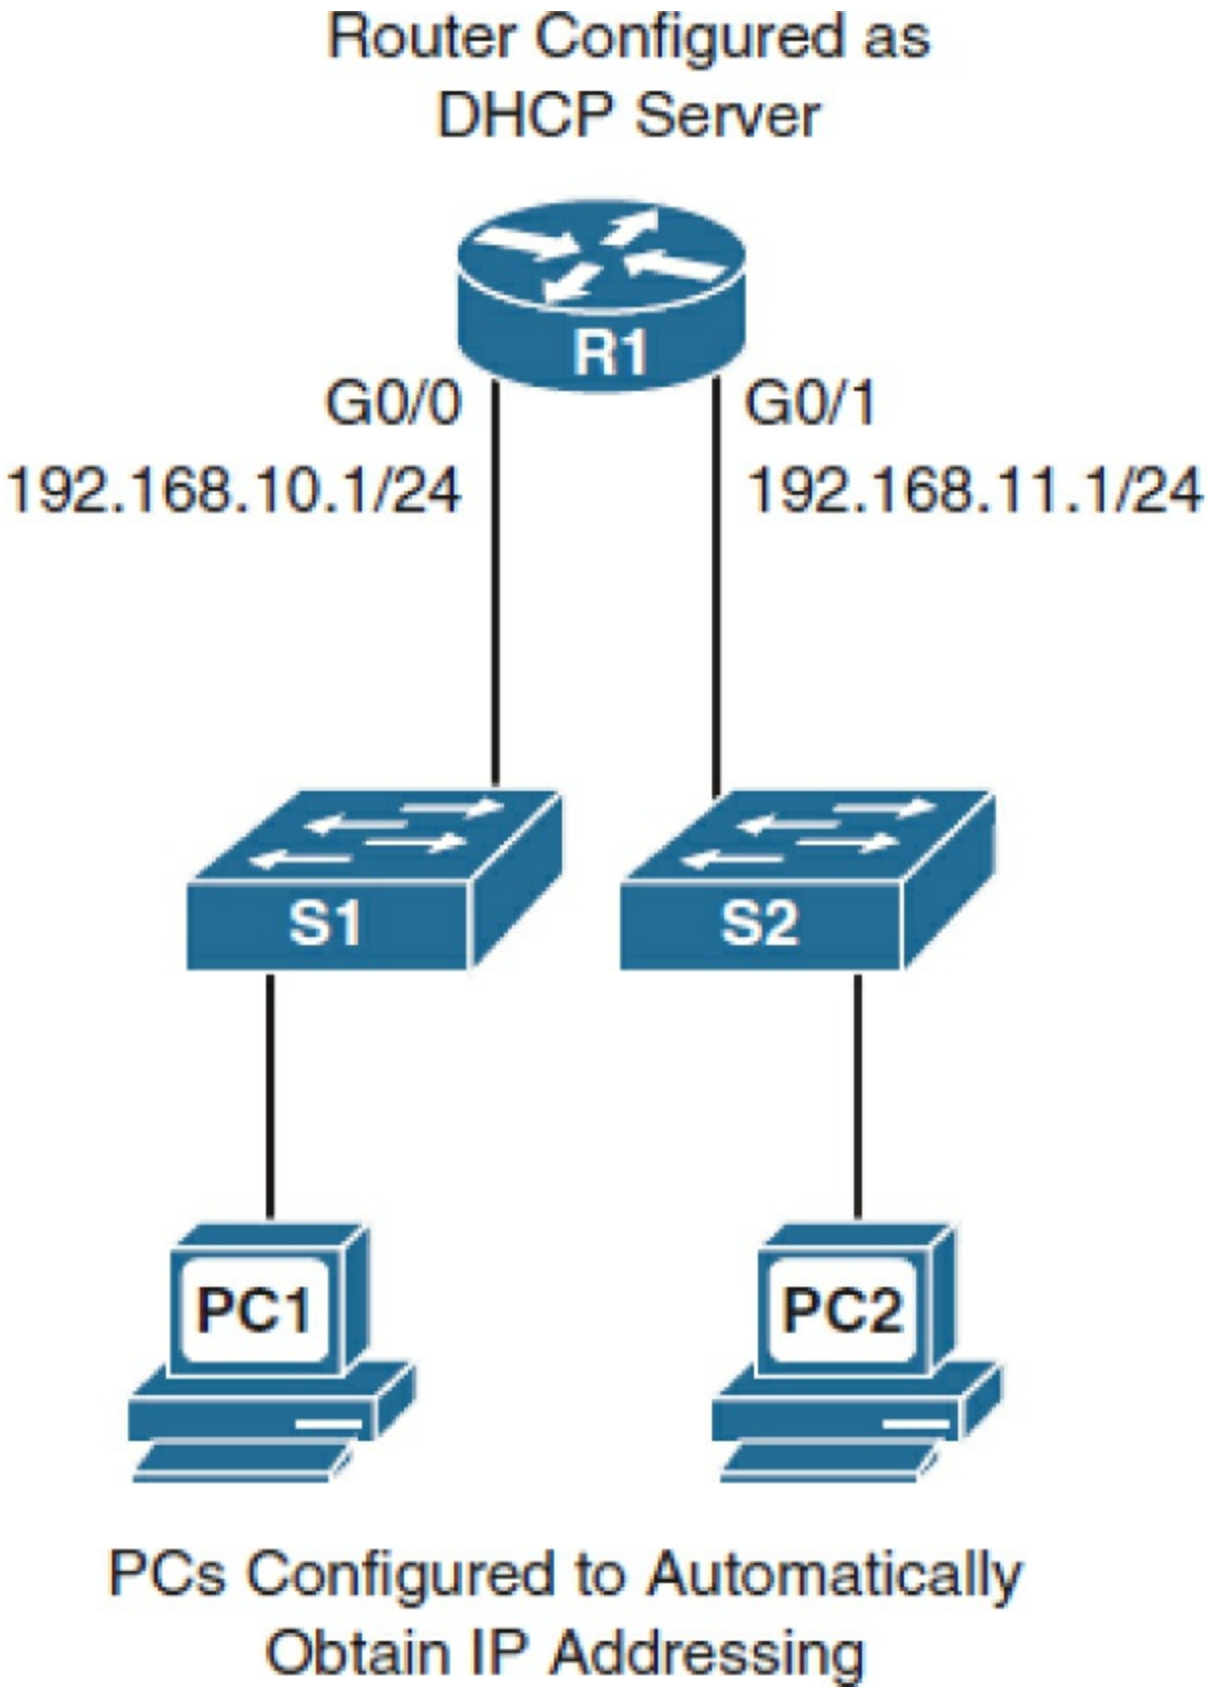
Router Configured as
DHCP Server
R1
G0/0
192.168.10.1/24
G0/1
192.168.11.1/24
S1
S2
PC1
PC2
PCs Configured to Automatically
Obtain IP Addressing

Figure 23-2 *DHCPv4 Sample Topology*

Example 23-1 shows DHCPv4 required and optional commands to configure R1 as the DHCPv4 server for both LANs in Figure 23-2.

Example 23-1 *DHCPv4 Configuration Example*

```
!Configure IP addresses that you want excluded from th
R1(config)# ip dhcp excluded-address 192.168.10.1 192.
R1(config)# ip dhcp excluded-address 192.168.10.254
R1(config)# ip dhcp excluded-address 192.168.11.1 192.
R1(config)# ip dhcp excluded-address 192.168.11.254
!R1 needs two DHCPv4 pools for the two LANs. Each pool
and optional commands.
R1(config)# ip dhcp pool LAN-POOL-10
R1(dhcp-config)# network 192.168.10.0 255.255.255.0
R1(dhcp-config)# default-router 192.168.10.1
R1(dhcp-config)# dns-server 192.168.50.195 209.165.202
R1(dhcp-config)# domain-name cisco.com
R1(dhcp-config)# lease 2
R1(dhcp-config)# netbios-name-server 192.168.10.254
R1(dhcp-config)# ip dhcp pool LAN-POOL-11
R1(dhcp-config)# network 192.168.11.0 255.255.255.0
R1(dhcp-config)# default-router 192.168.11.1
R1(dhcp-config)# dns-server 192.168.50.195 209.165.202
R1(dhcp-config)# domain-name cisco.com
R1(dhcp-config)# lease 2
R1(dhcp-config)# netbios-name-server 192.168.11.254
```

Technet24

```
R1(dhcp-config)# end
```

Cisco IOS Software supports DHCPv4 service by default. To disable it, use the global command **no service dhcp**.

To verify DHCPv4 operations on R1 in Figure 23-2, use the commands in Example 23-2.

Example 23-2 *Verifying DHCPv4 Operation*

```
R1# show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address    Client-ID/              Lease expiration
                Hardware address/
                User name
192.168.10.10  0100.1641.aea5.a7    Jul 18 2008 08:17
192.168.11.10  0100.e018.5bdd.35    Jul 18 2008 08:17

R1# show ip dhcp server statistics
Memory usage         26455
Address pools        2
Database agents      0
Automatic bindings   2
Manual bindings      0
Expired bindings     0

Malformed messages   0
```

```
Secure arp entries   0
Message              Received
BOOTREQUEST          0
DHCPDISCOVER         2
DHCPREQUEST          2
DHCPDECLINE          0
DHCPRELEASE          0
DHCPINFORM           0

Message              Sent
BOOTREPLY            0
DHCPOFFER            2
DHCPACK              2
DHCPNAK              0
R1#
```

Because PC1 and PC2 are connected to the LANs, each automatically receives its IP addressing information from the router's DHCPv4 server. Example 23-3 shows the output from the **ipconfig/all** command on PC1.

Example 23-3 *DHCPv4 Client Configuration*

```
C:\> ipconfig/all

Windows IP Configuration

 Host Name . . . . . . . . . . . .     : ciscolab
```

Technet24

```
   Primary Dns Suffix . . . . . .      :
   Node Type . . . . . . . . . . . .   : Hybrid
   IP Routing Enabled. . . . . . . .   : No
   WINS Proxy Enabled. . . . . . . .   : No

  Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix .    : cisco.com
   Description . . . . . . . . . . .   : Intel(R) PRO/
   Physical Address. . . . . . . . .   : 00-7-41-AE-A5
   Dhcp Enabled. . . . . . . . . . .   : Yes
   Autoconfiguration Enabled . . . .   : Yes
   IP Address. . . . . . . . . . . .   : 192.168.10.11
   Subnet Mask . . . . . . . . . . .   : 255.255.255.0
   Default Gateway . . . . . . . . .   : 192.168.10.1

   DHCP Server . . . . . . . . . . .   : 192.168.10.1
   DNS Servers . . . . . . . . . . .   : 192.168.50.19
   209.165.202.158
  Primary WINS Server . . . . . . .    : 192.168.10.25
   Lease Obtained. . . . . . . . . .     : Wednesday,
   Lease Expires . . . . . . . . . .   : Friday, July

  C:\>
```

To release the DHCPv4 configuration on a Windows-based client, enter the **ipconfig/release** command. To renew the DHCPv4 configuration, enter the

ipconfig/renew command.

Configuring a Router to Relay DHCPv4 Requests

In a complex network, the DHCPv4 servers are usually contained in a server farm. Therefore, clients typically are not on the same subnet as the DHCPv4 server, as in the previous example. To ensure that broadcasted DHCPDISCOVER messages are sent to the remote DHCPv4 server, use the **ip helper-address** command.

For example, in Figure 23-3, the DHCPv4 server is located on the 192.168.11.0/24 LAN and is serving IP addressing information for both LANs.

Technet24

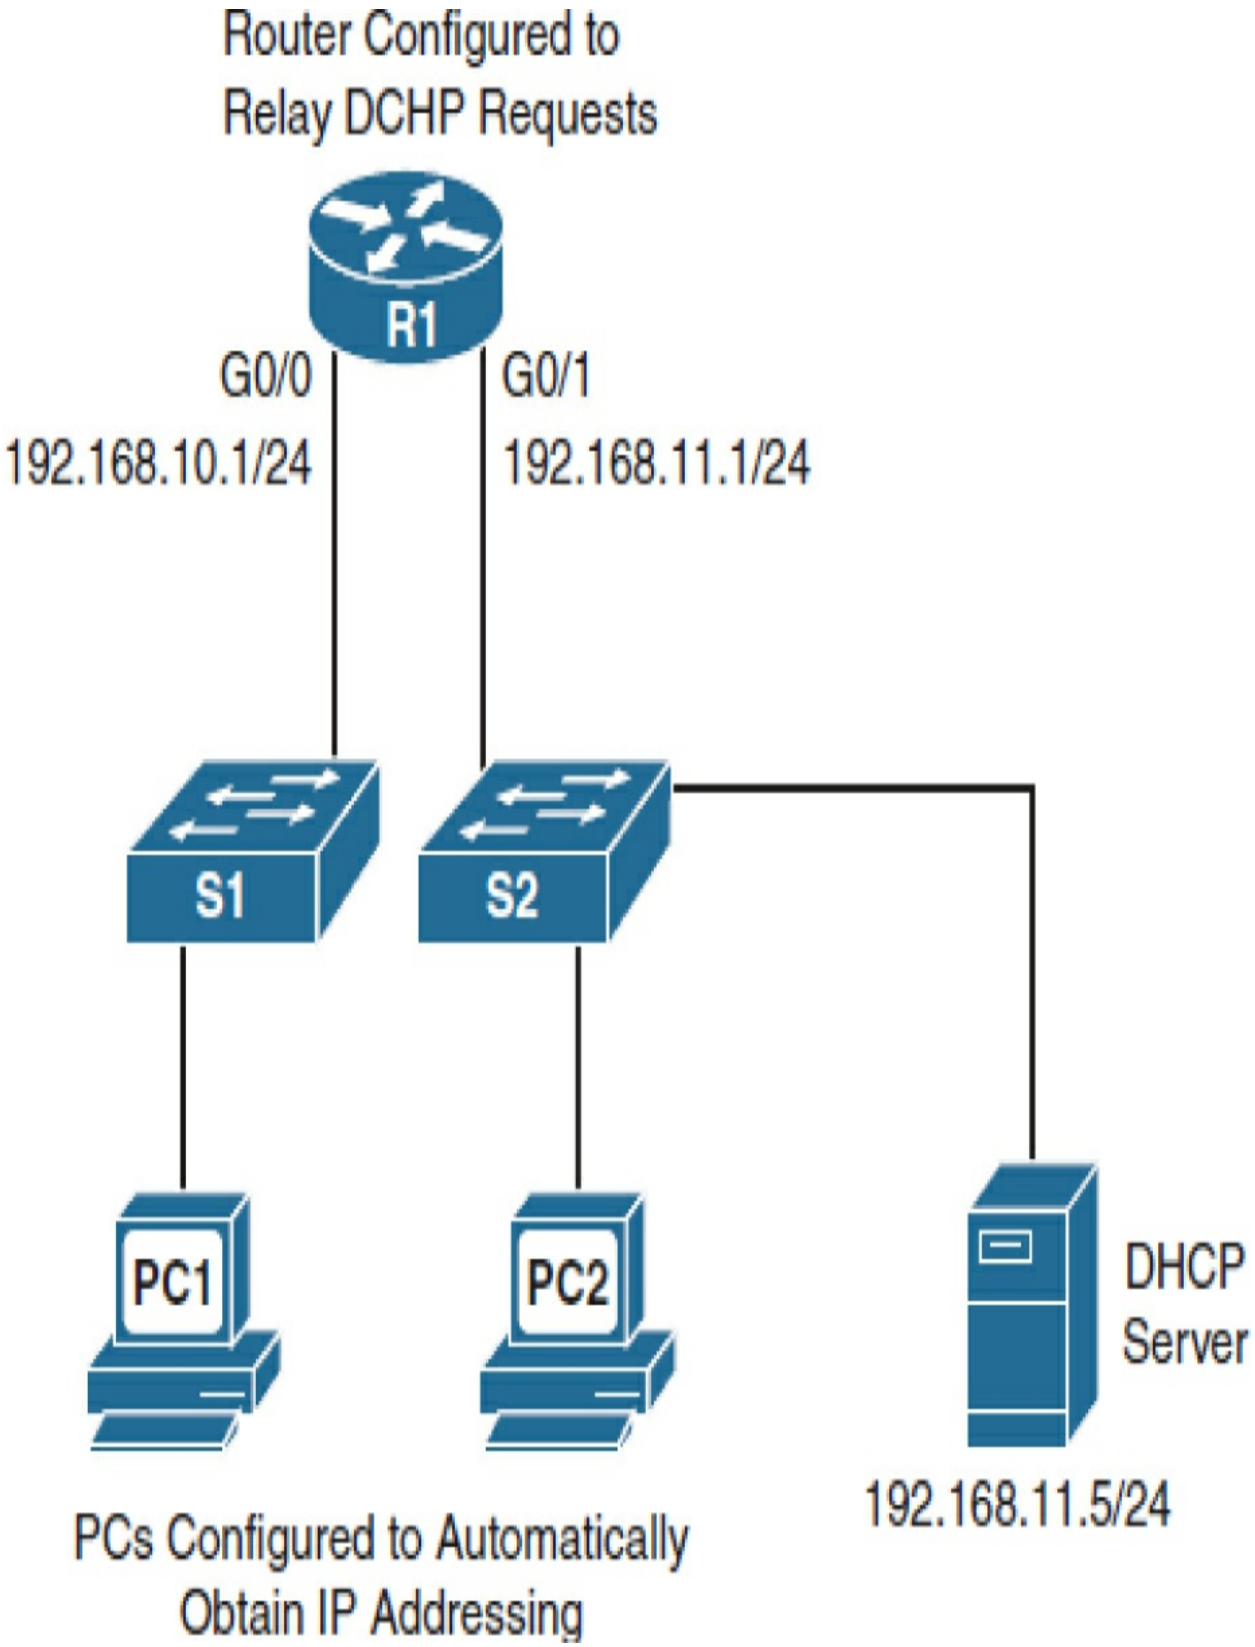

Figure 23-3 *DHCPv4 Relay Topology*

Without the **ip helper-address** command, R1 would discard any broadcasts from PC1 requesting DHCPv4 services. To configure R1 to relay DHCPDISCOVER messages, enter the following commands:

```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ip helper-address 192.168.11.5
```

Notice that the commands are entered on the interface that will receive DHCPv4 broadcasts. R1 then forwards DHCPv4 broadcast messages as a unicast to 192.168.11.5. By default, the **ip helper-address** command forwards the following eight UDP services:

• **Port 37:** Time

• **Port 49:** TACACS

• **Port 53:** DNS

• **Port 67:** DHCP/BOOTP server

• **Port 68:** DHCP/BOOTP client

• **Port 69:** TFTP

• **Port 137:** NetBIOS name service

• **Port 138:** NetBIOS datagram service

Technet24

To specify additional ports, use the global command **ip forward-protocol udp** [*port-number* | *protocol*]. To disable broadcasts of a particular protocol, use the **no** form of the command.

Configuring a Router as a DHCPv4 Client

Cisco routers in small offices or branch sites are often configured as DHCPv4 clients. The method used depends on the ISP. However, in its simplest configuration, the interface used to connect to a cable or DSL modem is configured with the **ip address dhcp** interface configuration command.

For example, in Figure 23-4, the BRANCH router's GigabitEthernet 0/1 interface can be configured to request addressing from the ISP router.

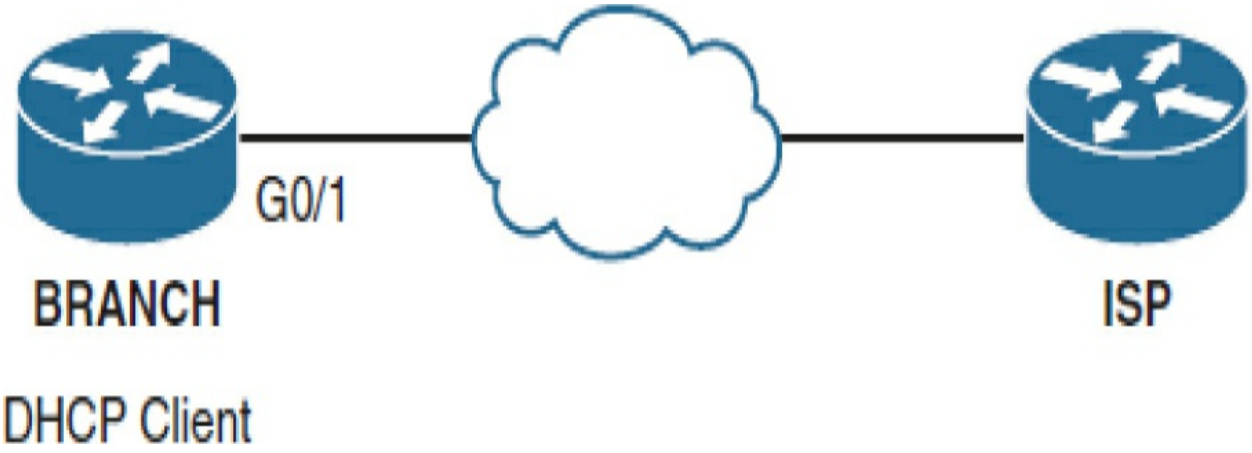

Figure 23-4 *Router as a DHCP Client*

Example 23-4 shows the configuration and verification of DHCP addressing on BRANCH.

Example 23-4 *Configuring a Router as a DHCP Client*

```
BRANCH(config)# interface g0/1
BRANCH(config-if)# ip address dhcp
BRANCH(config-if)# no shutdown
*Mar 15 08:45:34.632: %DHCP-6-ADDRESS_ASSIGN: Interfac
assigned
   DHCP address 209.165.201.12, mask 255.255.255.224,
BRANCH(config-if)# end
BRANCH# show ip interface g0/1
GigabitEthernet0/1 is up, line protocol is up
 Internet address is 209.165.201.12/27
 Broadcast address is 255.255.255.255
 Address determined by DHCP
 <output omitted>
BRANCH#
```

DHCPv6

IPv6 has two methods for automatically obtaining a global unicast address:

• Stateless address autoconfiguration (SLAAC)

• Stateful DHCPv6 (Dynamic Host Configuration Protocol for IPv6)

SLAAC

SLAAC uses ICMPv6 Router Solicitation (RS) and Router Advertisement (RA) messages to provide addressing and other configuration information. A

Technet24

client then uses the RA information to build an IPv6 address and verify it with a special type of Neighbor Solicitation (NS) message through duplicate address detection (DAD). These three message types—RS, RA, and NS—belong to the Neighbor Discovery Protocol:

• **Router Solicitation (RS) message:** When a client is configured to obtain its addressing information automatically using SLAAC, the client sends an RS message to the router. The RS message is sent to the IPv6 all-routers multicast address, FF02::2.

• **Router Advertisement (RA) message:** A client uses this information to create its own IPv6 global unicast address. A router sends RA messages periodically or in response to RS messages. An RA message includes the prefix and prefix length of the local segment. By default, Cisco routers send RA messages every 200 seconds. RA messages are sent to the IPv6 all-nodes multicast address, FF02::1.

• **Neighbor Solicitation (NS) message:** An NS message is normally used to learn the data link layer address of a neighbor on the same network. In the SLAAC process, a host uses DAD by inserting its own IPv6 address as the destination address in an NS message. The NS message is sent out on the network to verify that a newly minted IPv6 address is unique. If a Neighbor Advertisement message is received, the host knows that the IPv6 address is not unique.

Figure 23-5 shows the SLAAC process using three messages of NDP.

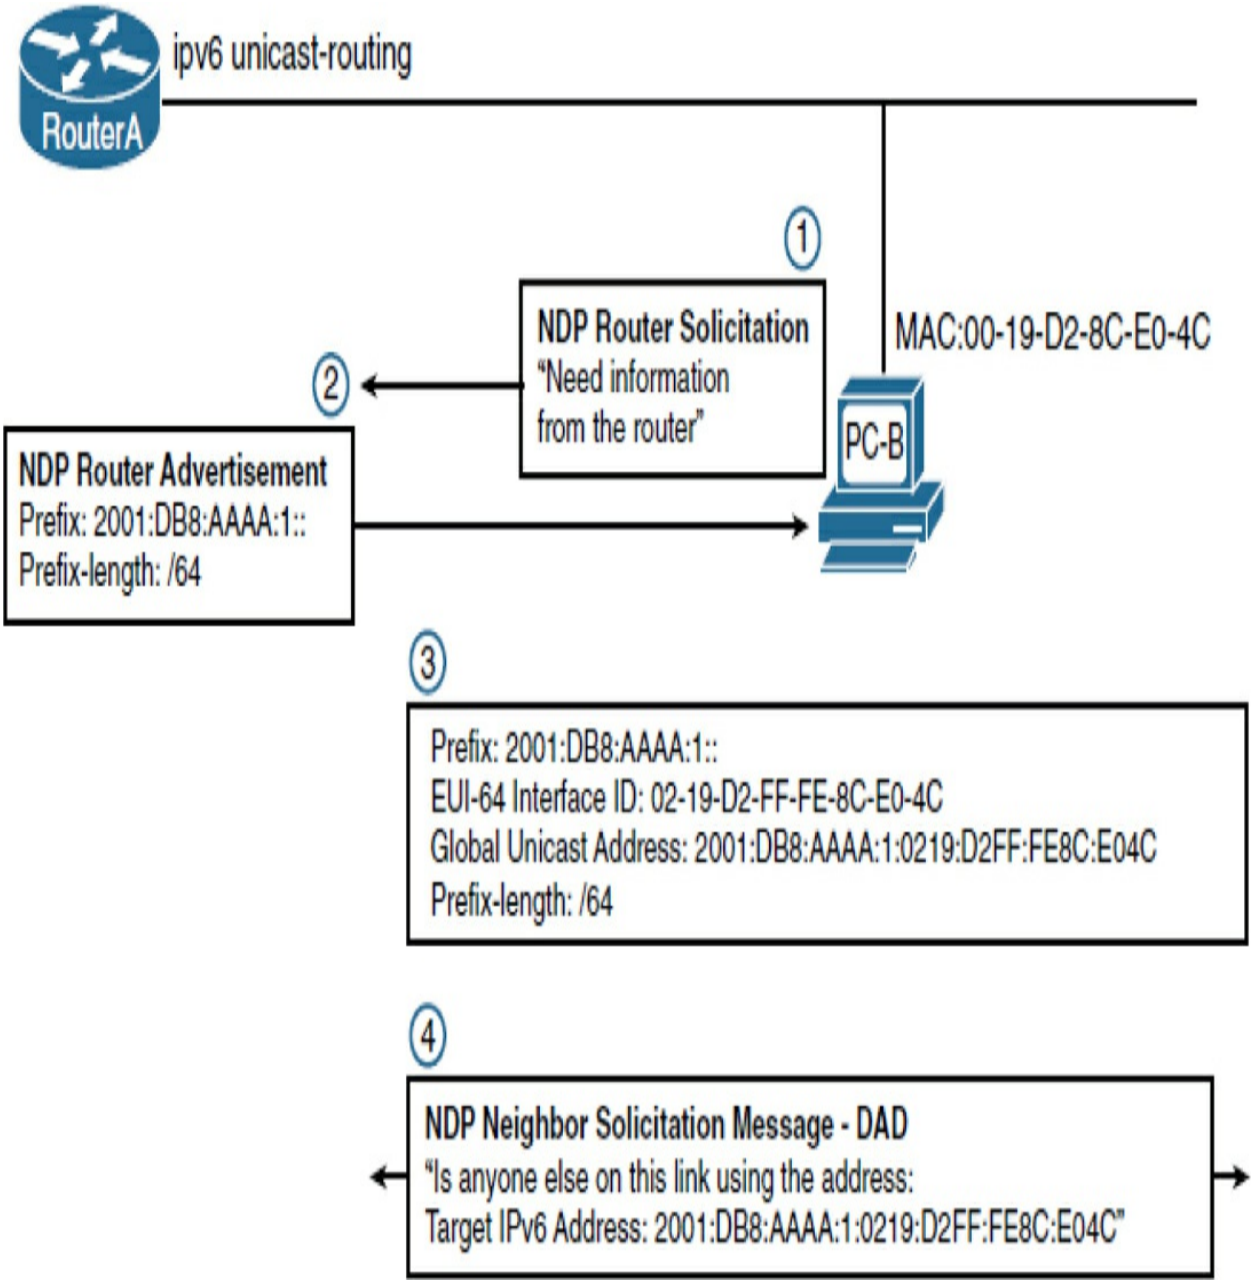

Figure 23-5 *Neighbor Discovery and the SLAAC Process*

Let's briefly review the steps in Figure 23-5.

PC-B sends an RS message to the all-routers multicast address, FF02::2, to

Technet24

inform the local IPv6 router that it needs an RA message.

RouterA receives the RS message and responds with an RA message. Included in the RA message are the prefix and prefix length of the network. The RA message is sent to the IPv6 all-nodes multicast address, FF02::1, with the link-local address of the router as the IPv6 source address.

PC-B uses this information to create its own IPv6 global unicast address. It appends the 64-bit prefix address to its own locally generated 64-bit interface ID, which it creates using either the EUI process (see Figure 23-5) or a random number generator. It uses RouterA's link-local address as the default gateway.

Before PC-B can use this newly created IPv6 address, it uses the DAD process, sending out an NS message to verify that the address is unique.

Note

A client's operating system can be configured to ignore RA messages, in which case the client always opts to use the services of a DHCPv6 server.

An RA message informs a client how to obtain automatic IPv6 addressing: using SLAAC, DHCPv6, or a combination of the two. The RA message contains two flags to indicate the configuration option: the Managed Address

Configuration flag (M flag) and the Other Configuration flag (O flag).

The default setting for these flags is 0, or both bits off. To the client, this means it is to use the SLAAC process exclusively to obtain all of its IPv6 addressing information. If either of these flags is set to 1 for some reason, you can use the **no** form of the following **ipv6 nd** commands in interface configuration mode to reset them to 0:

```
Router(config-if)# no ipv6 nd managed-config-flag
Router(config-if)# no ipv6 nd other-config-flag
```

Stateless DHCPv6

In stateless DHCPv6, the client uses the RA message from the router to generate its global unicast address. However, the client then sends a request to the DHCPv6 server to obtain any additional information that the RA has not already supplied.

For stateless DHCPv6, the O flag is set to 1 so that the client is informed that additional configuration information is available from a stateless DHCPv6 server. Use the following command on the interface to modify the RA message:

```
Router(config-if)# ipv6 nd other-config-flag
```

Technet24

Stateful DHCPv6

For stateful DHCPv6, the RA message tells the client to obtain all its addressing information from a DHCPv6 server. The M flag must be set on the interface with the following command:

```
Router(config-if)# ipv6 nd managed-config-flag
```

Stateless and Stateful DHCPv6 Operation

Figure 23-6 shows the full operation of DHCPv6, regardless of the method used: SLAAC, stateless DHCPv6, or stateful DHCPv6.

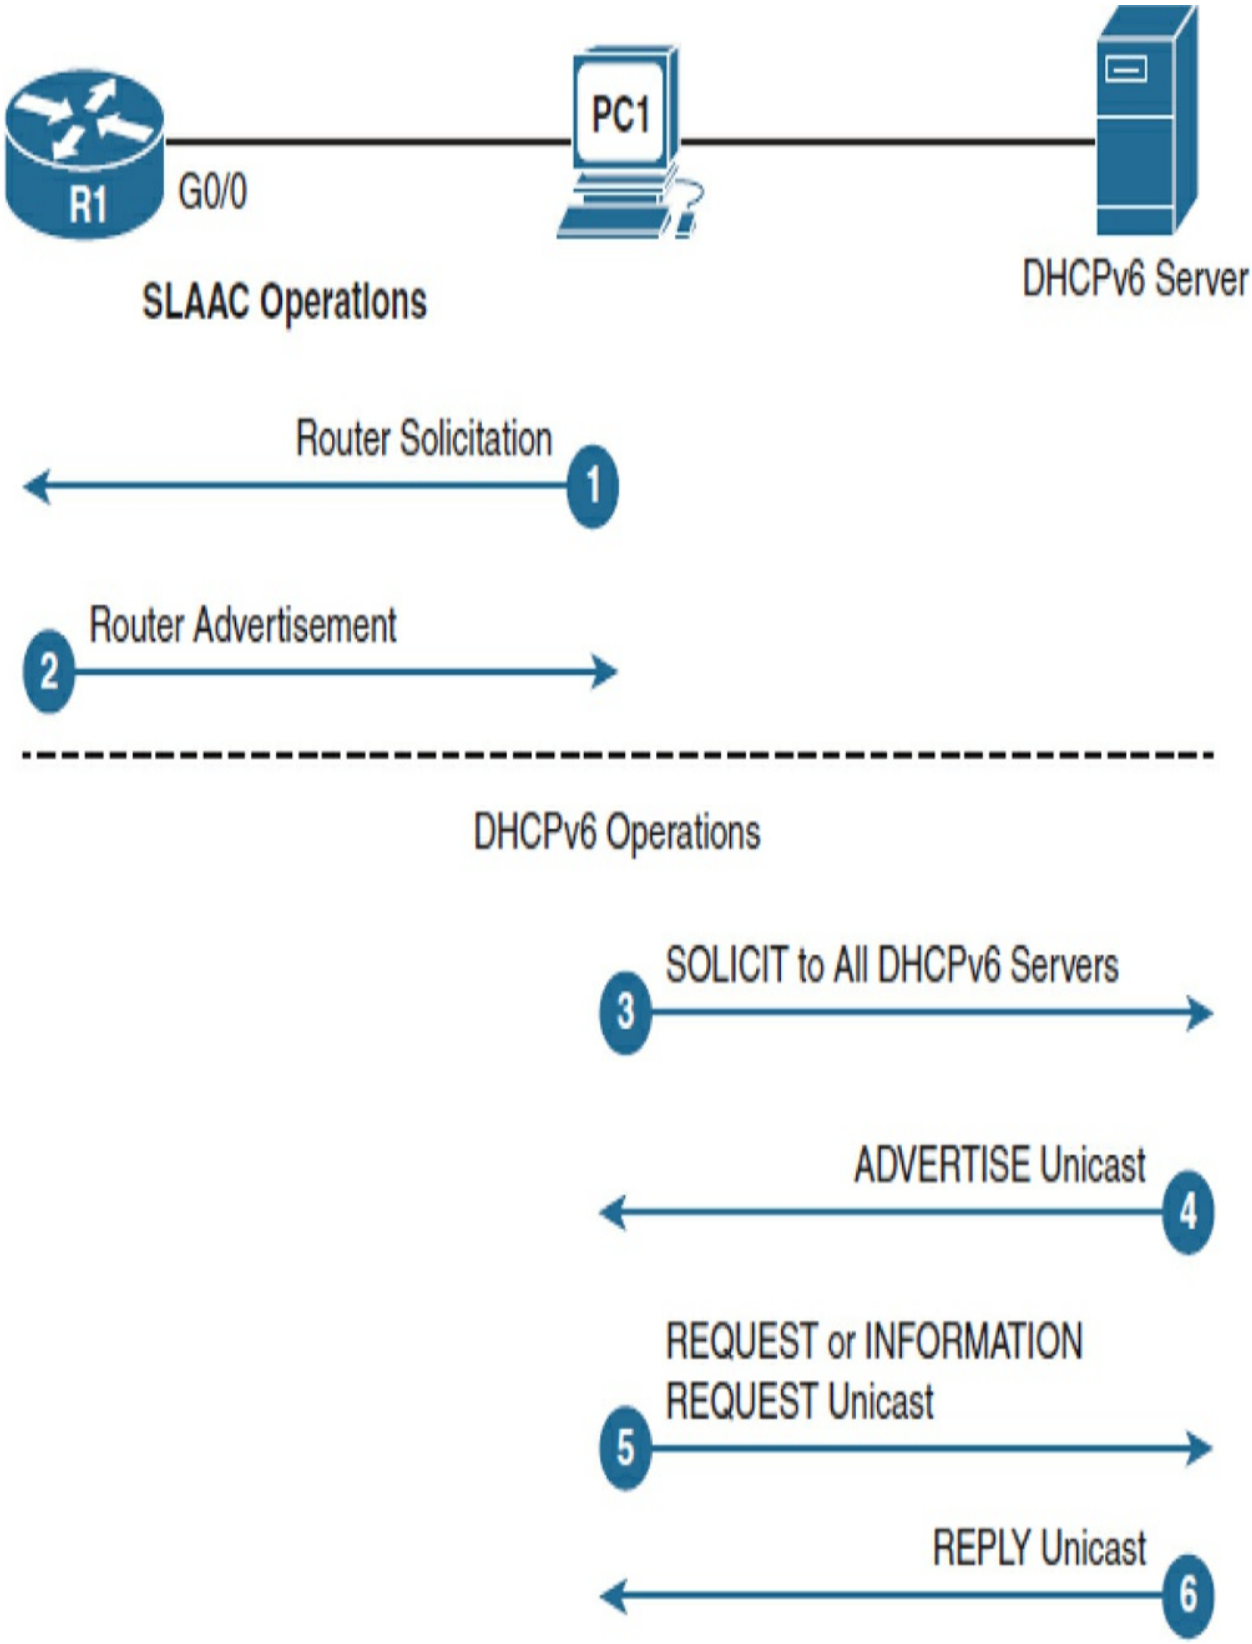

Figure 23-6 *DHCPv6 Operations*

Technet24

The following steps occur in Figure 23-6:

PC1 sends an RS message on bootup to begin the process of obtaining IPv6 addressing.

R1 replies with an RA message. If the M and O flags are not set, PC1 uses SLAAC. If either the M flag or the O flag is set, PC1 begins the DHCPv6 process.

PC1 sends a DHCPv6 SOLICIT message to the all-DHCPv6-servers address, FF02::1:2—a link-local multicast address that will not be forwarded by routers.

A DHCPv6 server responds with a DHCPv6 ADVERTISE unicast message informing the client of its presence.

The client sends either a unicast DHCPv6 REQUEST (the M flag was set, and the client is using stateful DHCPv6) or a unicast DHCPv6 INFORMATION-REQUEST (the O flag was set, and the client is using stateless DHCPv6).

The server replies with the information requested.

DHCPv6 Configuration Options

A router can be configured as a stateless DHCPv6 server, a stateful DHCPv6 server, and a DHCPv6 client. As in DHCPv4, the router can be configured

with all three, depending on what role it plays for its various interfaces.

Configuring a Router as a Stateless DHCPv6 Server

We use Figure 23-7 for all the examples in this section. R1 is the DHCPv6 server, and R3 is the DHCPv6 client.

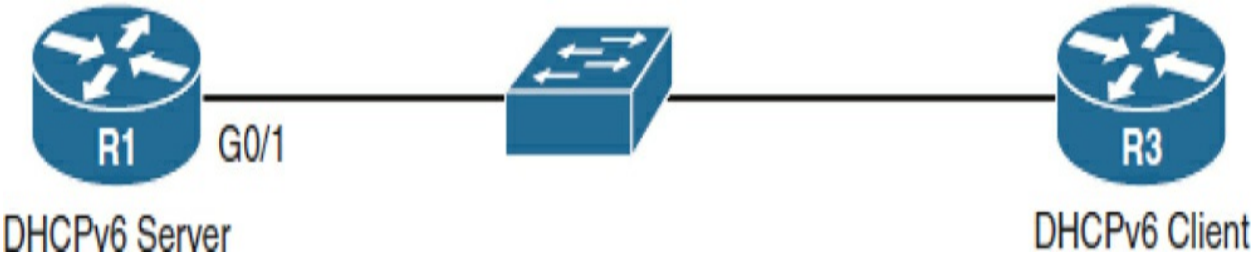

Figure 23-7 *DHCPv6 Server and Client Topology*

To configure R1 as a stateless DHCP server, you need to make sure that **ipv6 unicast-routing** is enabled. Then, in global configuration mode, configure the pool name, DNS server, and domain name. Finally, enable the DHCPv6 pool on the appropriate interface and set the O flag so that clients on that interface know to request DHCPv6 services from the router. Example 23-5 shows the configuration for R1.

Example 23-5 *Configuring a Router as a Stateless DHCPv6 Server*

```
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 dhcp pool O-FLAG-SET
R1(config-dhcpv6)# dns-server 2001:db8:acad:1::5
R1(config-dhcpv6)# domain-name cisco.com
R1(config-dhcpv6)# exit
```

Technet24

```
R1(config)# interface g0/1
R1(config-if)# ipv6 address 2001:db8:1:1::1/64
R1(config-if)# ipv6 dhcp server O-FLAG-SET
R1(config-if)# ipv6 nd other-config-flag
R1(config-if)# end
R1# show ipv6 dhcp pool
DHCPv6 pool: O-FLAG-SET
 DNS server: 2001:DB8:ACAD:1::5
 Domain name: cisco.com
 Active clients: 0
R1#
```

To configure a router interface as a DHCPv6 client, enable IPv6 on the interface and enter the **ipv6 address autoconfig** command, as in Example 23-6. Verify the configuration with the **show ipv6 interface** command.

Example 23-6 *Configuring an Interface as a DHCPv6 Client*

```
R3(config)# interface g0/1
R3(config-if)# ipv6 enable
R3(config-if)# ipv6 address autoconfig
R3(config-if)# end
R3# show ipv6 interface g0/1
GigabitEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::32F7:DF
  No Virtual link-local address(es):
  Stateless address autoconfig enabled
```

```
    Global unicast address(es):
       2001:DB8:1:1:32F7:DFF:FE25:2DE1, subnet is 2001:D
       valid lifetime 2591935 preferred lifetime 604735
    Joined group address(es):
      FF02::1
      FF02::1:FF25:2DE1
    MTU is 1500 bytes
    ICMP error messages limited to one every 100 millise
    ICMP redirects are enabled
    ICMP unreachables are sent
    ND DAD is enabled, number of DAD attempts: 1
    ND reachable time is 30000 milliseconds (using 30000
    ND NS retransmit interval is 1000 milliseconds
    Default router is FE80::D68C:B5FF:FECE:A0C1 on Gigab
  R3#
```

Configuring a Router as a Stateful DHCPv6 Server

The main difference between a stateless configuration and a stateful configuration is that a stateful server includes IPv6 addressing information and keeps a record of the IPv6 addresses that are leased out. Also, for the client side, the **ipv6 address dhcp** command is used instead of the **ipv6 address autoconfig** command. Example 23-7 shows the stateful DHCPv6 server configuration with stateful address information added and the M bit set instead of the O bit.

Technet24

Example 23-7 *Configuring a Router as a Stateful DHCPv6 Server*

```
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 dhcp pool M-FLAG-SET
R1(config-dhcpv6)# address prefix 2001:db8:1:1::/64 li
R1(config-dhcpv6)# dns-server 2001:db8:acad:1::5
R1(config-dhcpv6)# domain-name cisco.com
R1(config-dhcpv6)# exit
R1(config)# interface g0/1
R1(config-if)# ipv6 address 2001:db8:1:1::1/64
R1(config-if)# ipv6 nd managed-config-flag
R1(config-if)# end
!After R3 is configured as a DHCP client, verify DHCP
R1# show ipv6 dhcp pool
DHCPv6 pool: M-FLAG-SET
  Address allocation prefix: 2001:DB8:1:1::/64 valid 4
4294967295 (1 in use, 0 conflicts)
  DNS server: 2001:DB8:ACAD:1::5
  Domain name: cisco.com
  Active clients: 1
R1# show ipv6 dhcp binding
Client: FE80::32F7:DFF:FEA3:1640
 DUID: 0003000130F70DA31640
 Username: unassigned
 IA NA: IA ID 0x00060001, T1 43200, T2 69120
   Address: 2001:DB8:1:1:8902:60D6:E76:6C16
 preferred lifetime INFINITY, , valid lifetime INFINIT
R1#
```

DHCP Troubleshooting

DHCP problems can arise for a multitude of reasons, such as software defects in operating systems, NIC drivers, or DHCP relay agents. However, the most common problems are configuration issues.

Resolving IPv4 Address Conflicts

An IPv4 address lease can expire on a client that is still connected to a network. If the client does not renew the lease, the DHCP server can reassign that IPv4 address to another client. When the client reboots, it requests an IPv4 address. If the DHCP server does not respond quickly, the client uses the last IPv4 address. Then two clients begin using the same IPv4 address, creating a conflict.

The **show ip dhcp conflict** command displays all address conflicts recorded by the DHCP server. The server uses the **ping** command to detect conflicts. The client uses Address Resolution Protocol (ARP) to detect clients. If an address conflict is detected, the address is removed from the pool and is not assigned until an administrator resolves the conflict.

Testing Connectivity Using a Static IP Address

When troubleshooting any DHCP issue, verify network connectivity by configuring static IPv4 address information on a client workstation. If the

Technet24

workstation cannot reach network resources with a statically configured IPv4 address, the root cause of the problem is not the DHCP server. At this point, network connectivity troubleshooting is required.

Verifying Switch Port Configuration

If the DHCP client cannot obtain an IPv4 address from the DHCP server at startup, attempt to obtain an IPv4 address from the DHCP server by manually forcing the client to send a DHCP request. If a switch lies between the client and the DHCP server and the client cannot obtain the DHCP configuration, switch port configuration issues might be the cause. These causes can include issues from trunking and channeling to STP and RSTP. PortFast configuration and edge port configurations resolve the most common DHCPv4 client issues that occur with an initial installation of a Cisco switch.

Testing DHCPv4 Operation on the Same Subnet or VLAN

Distinguishing whether DHCP is functioning correctly is important when the client is on the same subnet or VLAN as the DHCP server. If DHCP is working correctly when the client is on the same subnet or VLAN, the problem might be the DHCP relay agent. If the problem persists even when testing DHCP on the same subnet or VLAN as the DHCP server, the problem might be with the DHCP server.

DNS Operation

DNS is a distributed system of servers that resolve domain names to IP addresses. The domain name is part of the uniform resource identifier (URI), as Figure 23-8 shows.

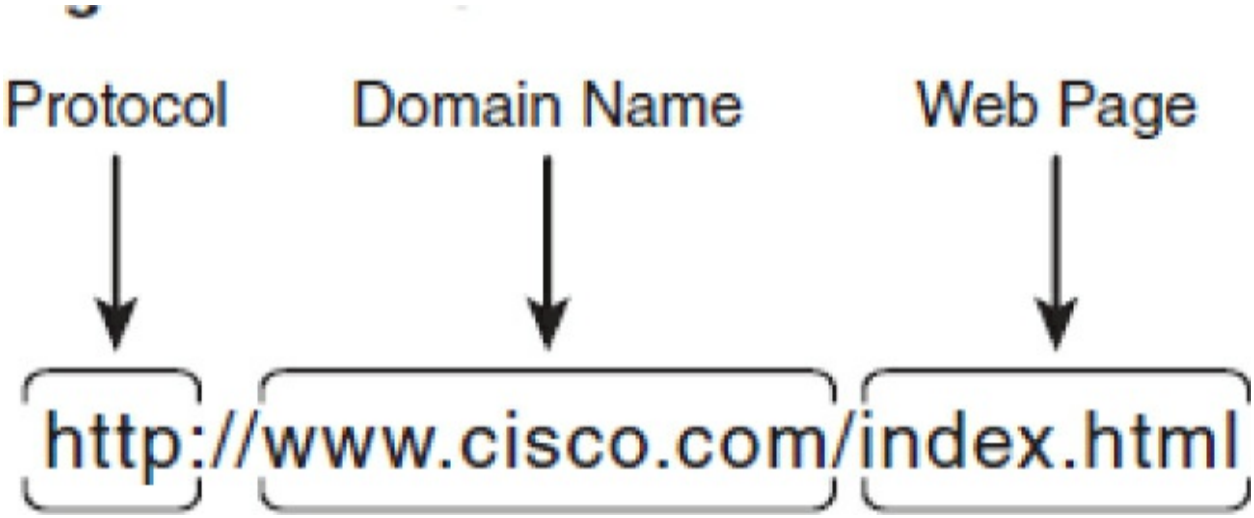

Figure 23-8 *URI Structure*

Note

Many people use the terms *web address* and *universal (or uniform) resource locator (URL)*. However, *uniform resource identifier (URI)* is the correct formal term.

When you type a new URI in your browser, your computer uses DNS to send out a request to resolve the URI into an IP address. Figure 23-9 summarizes the DNS process.

Technet24

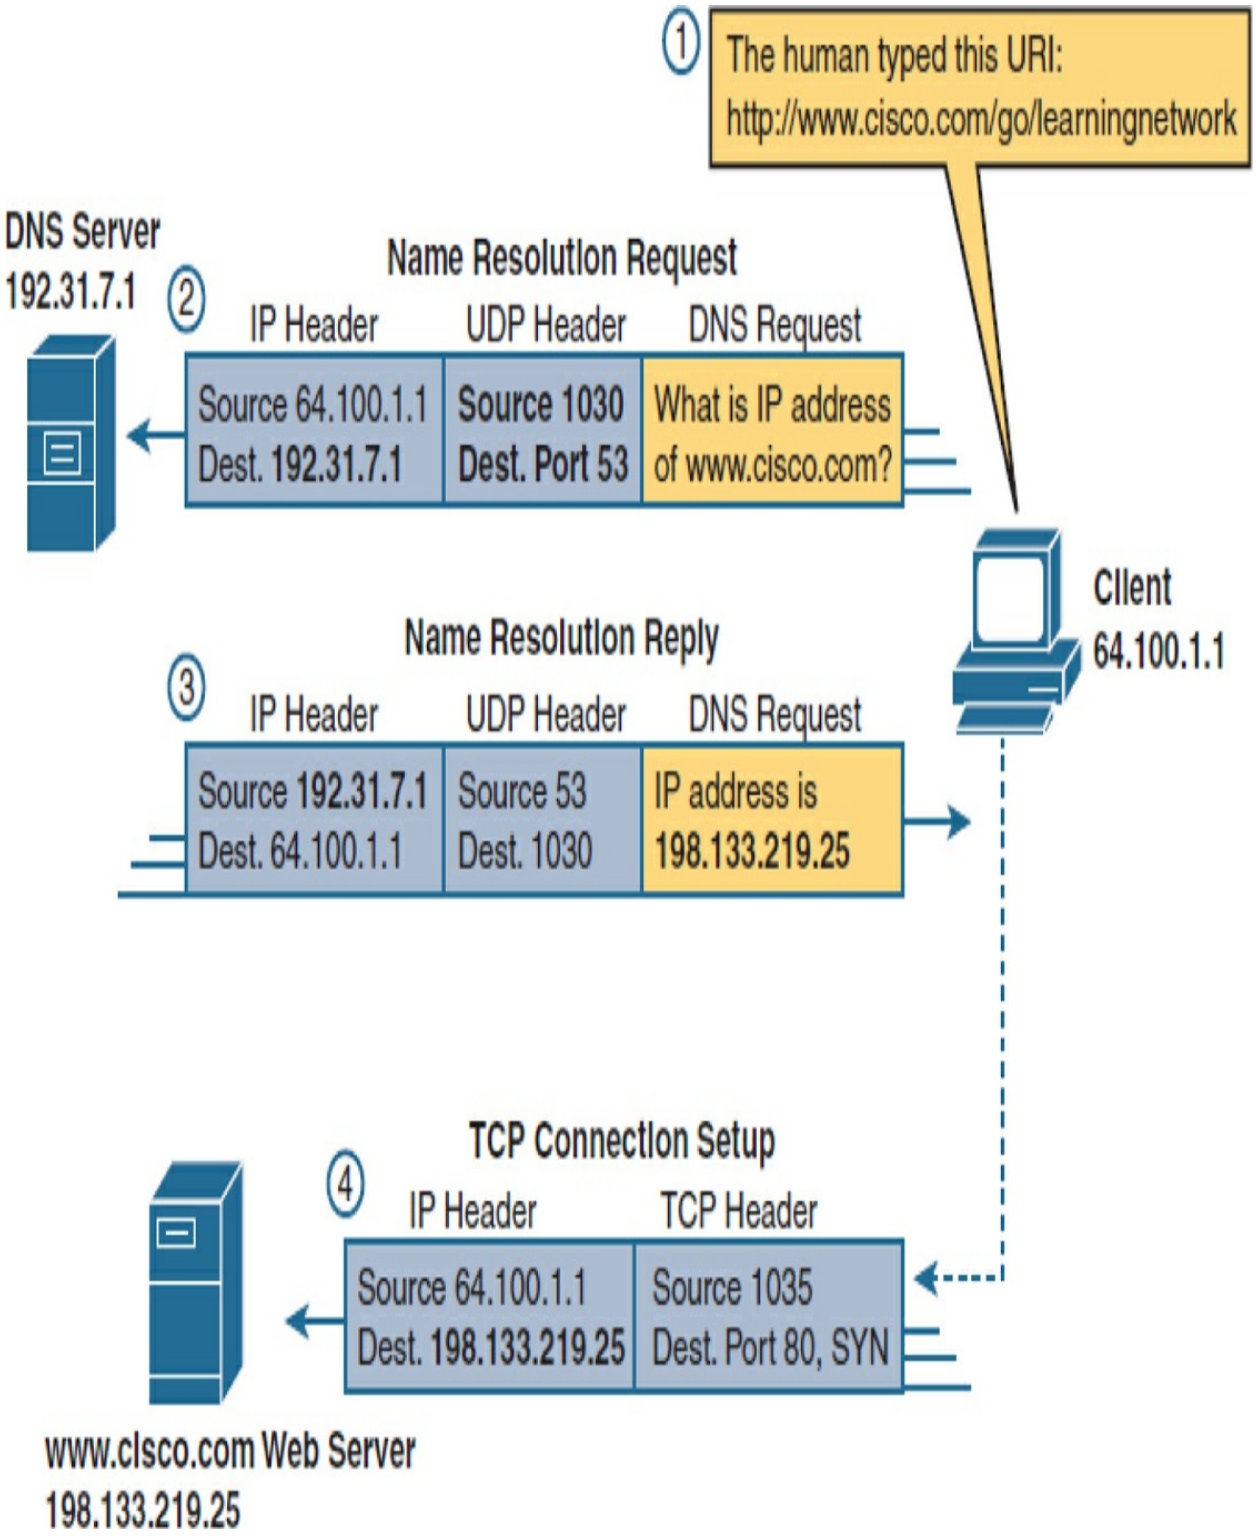

Figure 23-9 *DNS Process*

The DNS server stores different types of resource records used to resolve

names. These records contain the name, address, and type of record. Some of these record types follow:

• **A:** An end device IPv4 address

• **NS:** An authoritative name server

• **AAAA:** An end device IPv6 address (pronounced “quad-A”)

• **MX:** A mail exchange record

When a client makes a query, the server’s DNS process first looks at its own records to resolve the name. If it cannot resolve the name using its stored records, it contacts other servers to resolve the name.

DNS root servers manage the top-domain suffixes, such as these:

• **.com:** Commercial businesses

• **.edu:** Educational organizations

• **.gov:** Government organizations

• **.mil:** Military organizations

• **.net:** Networking organizations, such as ISPs

• **.org:** Noncommercial organizations

Technet24

Top-level DNS servers also exist for each country code, such as .ca (Canada), .de (Germany), .ru (Russia), and .cn (China).

Troubleshooting DNS

As a network administrator, your control over DNS issues is limited to two basic issues: DHCP server configurations and DNS server configurations.

In a small branch office, you are most likely using your ISP for all your DNS resolutions. Therefore, all the clients on your network will most likely have the IP address of the default gateway configured as the DNS server, as shown in the **ipconfig /all** output in Example 23-8.

Example 23-8 *DNS Server As the Default Gateway*

```
C:\> ipconfig /all

Windows IP Configuration
<output omitted>

   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 10.10.10.2(Pref
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Sunday, Novembe
   Lease Expires . . . . . . . . . . : Monday, Novembe
   Default Gateway . . . . . . . . . : 10.10.10.1
```

```
   DHCP Server . . . . . . . . . . . . : 10.10.10.1
   DNS Servers . . . . . . . . . . . . : 10.10.10.1
```

Therefore, issues with DNS are most likely due to issues with the default gateway router or the connection to your ISP. If you know the IP address of a publicly available server, you can verify that DNS is the issue if you can ping the IP address but not the URI.

In larger organizations, the network administrator is responsible for making sure the DHCP server is configured with accurate DNS IP addresses. Those DNS servers are most likely managed in-house to reduce the amount of outbound traffic to the public DNS servers. DNS server misconfiguration could be the cause if end-user devices cannot resolve URIs. Therefore, the hierarchy of DNS servers within the organization should ensure that there are backup DNS servers and that, when a record doesn't exist, the DNS server can accurately forward the request to another DNS server.

Verifying Host IP Configuration

Whether manually configured or dynamically learned, every device on the network must have a valid IP address configuration. The following are some examples of those settings on Windows, Linux, and macOS.

IP Settings

To work correctly, an IP host needs to know these values:

Technet24

• DNS server IP addresses

• Default gateway (router) IP address

• Device’s own IP address

• Device’s own subnet mask

Host IP Settings on Windows

In Windows 10, you can access IP address details from the Network and Sharing Center, as shown in Figure 23-10, to quickly view the four important settings: address, mask, router, and DNS.

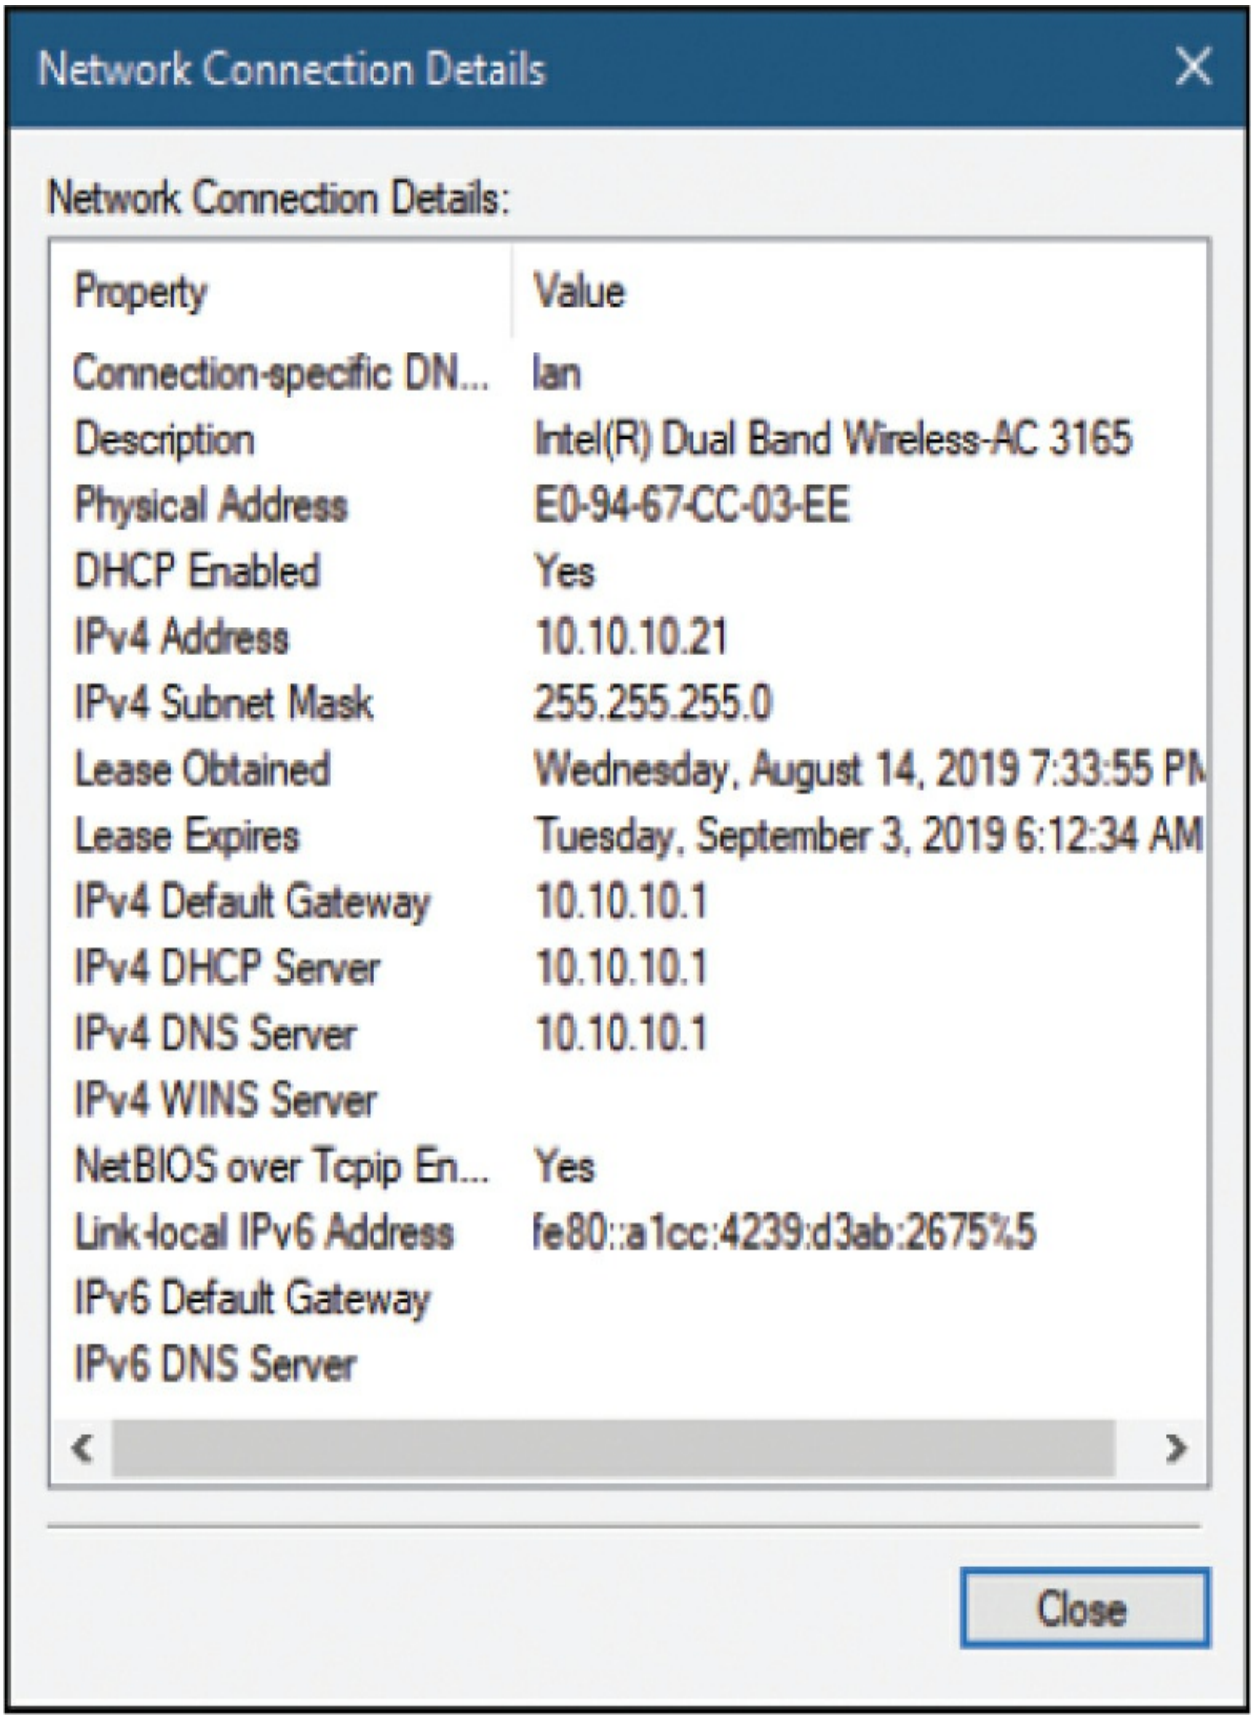
Network Connection Details
Network Connection Details:
Property
Value
Connection-specific DN...
lan
Description
Intel(R) Dual Band Wireless-AC 3165
Physical Address
E0-94-67-CC-03-EE
DHCP Enabled
Yes
IPv4 Address
10.10.10.21
IPv4 Subnet Mask
255.255.255.0
Lease Obtained
Wednesday, August 14, 2019 7:33:55 PM
Lease Expires
Tuesday, September 3, 2019 6:12:34 AM
IPv4 Default Gateway
10.10.10.1
IPv4 DHCP Server
10.10.10.1
IPv4 DNS Server
10.10.10.1
IPv4 WINS Server
NetBIOS over Tcpip En...
Yes
Link-local IPv6 Address
fe80::a1cc:4239:d3ab:2675%5
IPv6 Default Gateway
IPv6 DNS Server
Close

Technet24

Figure 23-10 *Windows 10 Network Connection Details*

However, network administrators typically open a command line window to verify IP settings. All Windows versions support **ipconfig** and **ipconfig /all**, as shown in Example 23-9. Both list the address, mask, and default gateway. But you need **ipconfig /all** to see DNS server settings. Also notice in Example 23-9 that the Ethernet adapter does not have a default gateway. That is because the computer is currently using the wireless adapter for network connectivity.

Example 23-9 *Windows Command Line IP Settings Verification*

```
C:\> ipconfig
<some output omitted>

Windows IP Configuration

Ethernet adapter Ethernet 2:

   Connection-specific DNS Suffix .  : cisco.com
   IPv6 Address. . . . . . . . . . . : 2001:db8:acad:1
   Link-local IPv6 Address . . . . . : fe80::ad66:4abc
   IPv4 Address. . . . . . . . . . . : 10.24.247.53
   Subnet Mask . . . . . . . . . . . : 255.255.255.255
   Default Gateway . . . . . . . . . : ::
                                       0.0.0.0
```

```
Wireless LAN adapter Wi-Fi:

   Connection-specific DNS Suffix .  : lan
   Link-local IPv6 Address . . . . . : fe80::90cb:adf9
   IPv4 Address. . . . . . . . . . . : 10.10.10.73
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 10.10.10.1


C:\> ipconfig /all
<some output omitted>

Windows IP Configuration

 Host Name . . . . . . . . . . . . .   : ALLANJ
 Primary Dns Suffix . . . . . . . .    : cisco.com
 Node Type . . . . . . . . . . . . .   : Hybrid
 IP Routing Enabled. . . . . . . . .   : No
 WINS Proxy Enabled. . . . . . . . .   : No
 DNS Suffix Search List. . . . . .     : cisco.com
 lan
 Wireless LAN adapter Wi-Fi:

 Connection-specific DNS Suffix .     : lan
 Description . . . . . . . . . . . .     : Intel(R) Dual
 Physical Address. . . . . . . . . .    : 88-B1-11-77-4A
 DHCP Enabled. . . . . . . . . . . .    : Yes
 Autoconfiguration Enabled . . . .    : Yes
```

Technet24

```
   Link-local IPv6 Address . . . . .   : fe80::90cb:adf9
   IPv4 Address. . . . . . . . . . .    : 10.10.10.73(Pr
   Subnet Mask . . . . . . . . . . .    : 255.255.255.0
   Lease Obtained. . . . . . . . . .     : Saturday, Aug
   Lease Expires . . . . . . . . . .    : Tuesday, Septe
   Default Gateway . . . . . . . . .   : 10.10.10.1
   DHCP Server . . . . . . . . . . .   : 10.10.10.1
   DHCPv6 IAID . . . . . . . . . . .   : 92844305
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-21-8E
   DNS Servers . . . . . . . . . . . : 10.10.10.1
   NetBIOS over Tcpip. . . . . . . .   : Enabled

C:\Users\allan>
```

Host IP Settings on macOS

On a Mac, open **Network Preferences > Advanced** to get the IP addressing information shown in Figure 23-11. The router IP address is the default gateway and also serves as the DNS server for this device.

Network
Search
Wi-Fi
Wi-Fi
TCP/IP
DNS
WINS
802.1X
Proxies
Hardware
Configure IPv4:
Using DHCP
IPv4 Address:
10.10.10.113
Renew DHCP Lease
Subnet Mask:
255.255.255.0
DHCP Client ID:
(if required)
Router:
10.10.10.1
Configure IPv6:
Automatically
Router:
IPv6 Address:
Prefix Length:
?
Cancel
OK

Technet24

Figure 23-11 *macOS Network Settings*

On the command line for macOS, use the **ifconfig** command to view IP addressing information, as shown in Example 23-10. Other useful commands include **networksetup -listallnetworkservices**, **networksetup -getinfo** <*network service*>, and **networksetup -getdnsservers** <*network service*> (not shown).

Example 23-10 *macOS Command Line IP Settings Verification*

```
MacBook-Air:~ Becky$ ifconfig en0
en0: flags=8863 mtu 1500
        ether c4:b3:01:a0:64:98
        inet6 fe80::c0f:1bf4:60b1:3adb%en0 prefixlen
        inet 10.10.10.113 netmask 0xffffff00 broadcas
        nd6 options=201
        media: autoselect
        status: active
MacBook-Air:~ Becky$ networksetup -listallnetworkservi
An asterisk (*) denotes that a network service is disa
iPhone USB
Wi-Fi
Bluetooth PAN
Thunderbolt Bridge
MacBook-Air:~ Becky$ networksetup -getinfo Wi-Fi
DHCP Configuration
IP address: 10.10.10.113
```

```
Subnet mask: 255.255.255.0
Router: 10.10.10.1
Client ID:
IPv6: Automatic
IPv6 IP address: none
IPv6 Router: none
Wi-Fi ID: c4:b3:01:a0:64:98
MacBook-Air:~ Becky$
```

Host IP Settings on Linux

Verifying IP settings using the GUI on a Linux machine differs depending on the Linux distribution and desktop interface. Figure 23-12 shows the Connection Information dialog box on the Ubuntu distro running the Gnome desktop.

Technet24

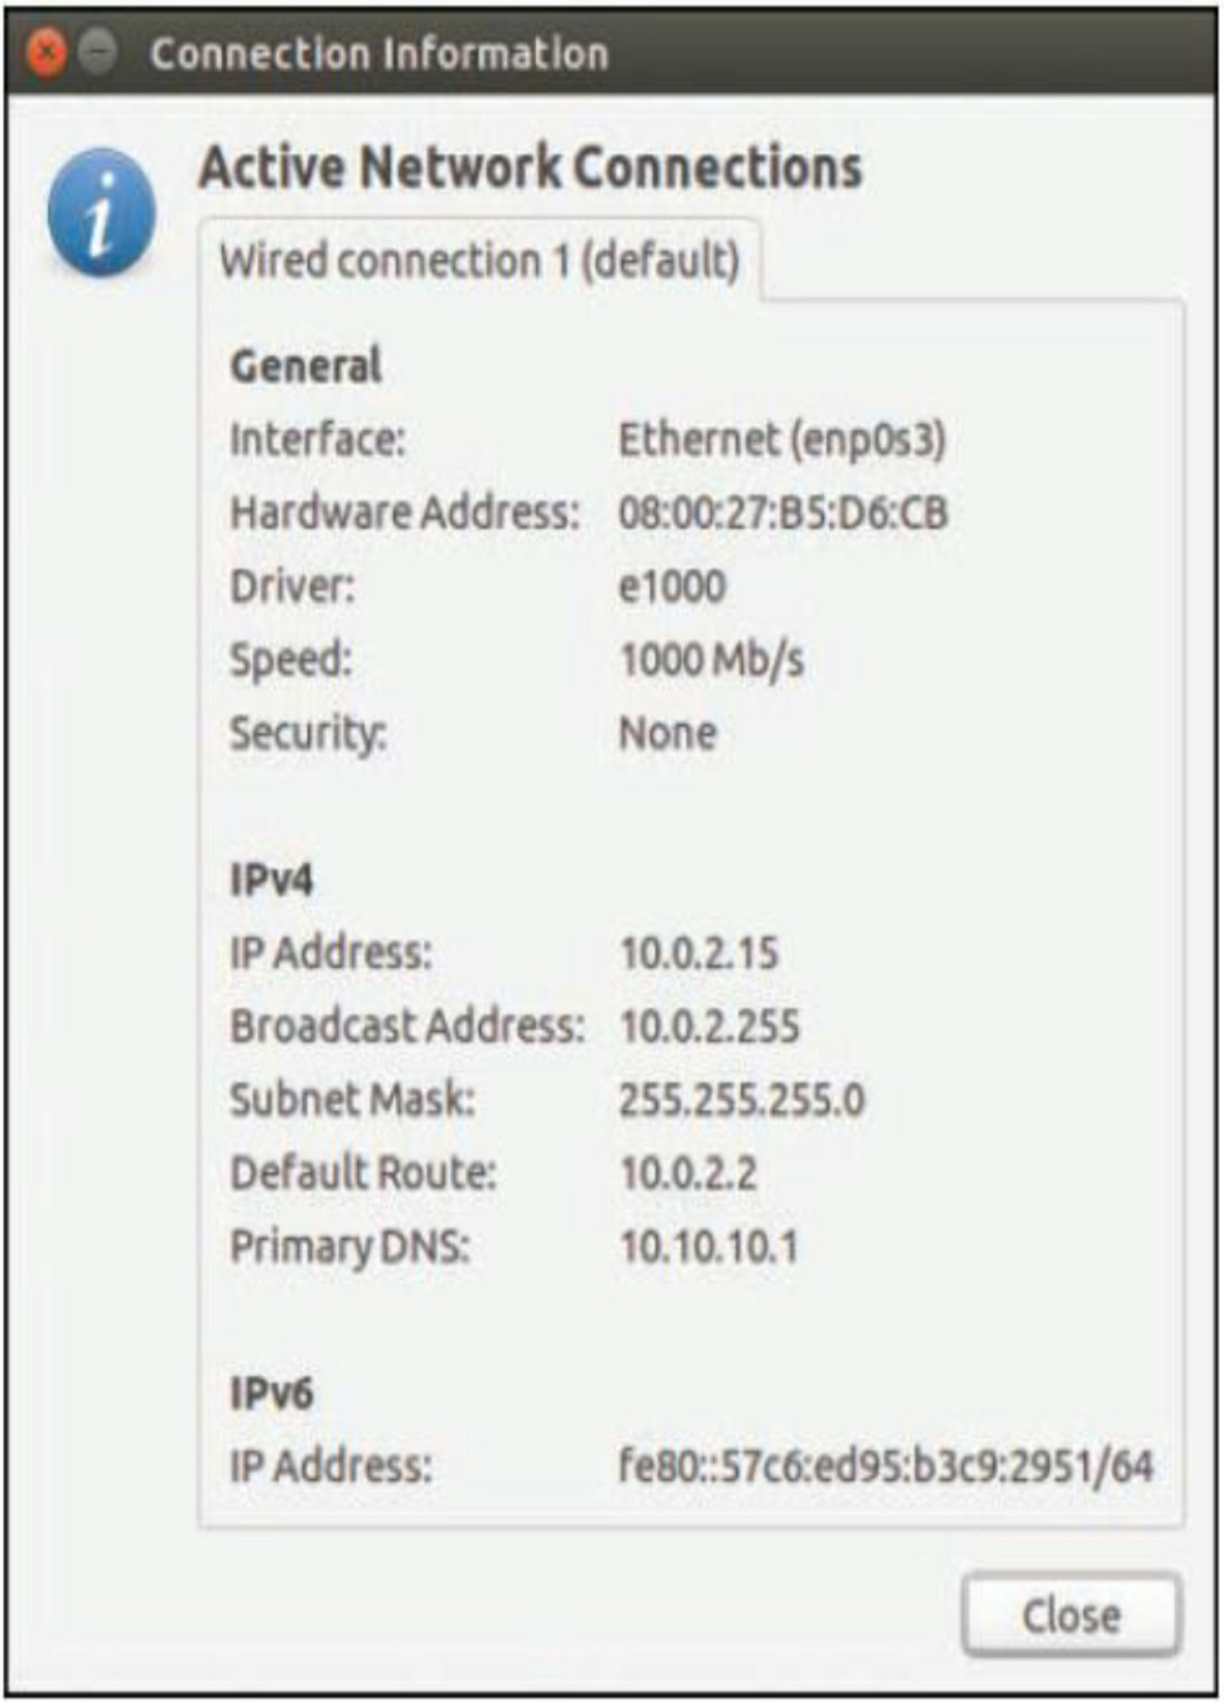
Connection Information
Active Network Connections
Wired connection 1 (default)
General
Interface: Ethernet (enp0s3)
Hardware Address: 08:00:27:B5:D6:CB
Driver: e1000
Speed: 1000 Mb/s
Security: None
IPv4
IP Address: 10.0.2.15
Broadcast Address: 10.0.2.255
Subnet Mask: 255.255.255.0
Default Route: 10.0.2.2
Primary DNS: 10.10.10.1
IPv6
IP Address: fe80::57c6:ed95:b3c9:2951/64
Close

Figure 23-12 *Linux Ubuntu Connection Information*

Example 23-11 shows the commands to verify the IP settings on a Linux machine.

Example 23-11 *Linux OS Command Line IP Settings Verification*

```
allan@allan-VirtualBox:~$ ifconfig enp0s3
enp0s3    Link encap:Ethernet  HWaddr 08:00:27:b5:d6:c
          inet addr:10.0.2.15  Bcast:10.0.2.255  Mask:
          inet6 addr: fe80::57c6:ed95:b3c9:2951/64 Sco
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Me
          RX packets:1332239 errors:0 dropped:0 overru
          TX packets:105910 errors:0 dropped:0 overrun
          collisions:0 txqueuelen:1000
          RX bytes:1855455014 (1.8 GB)  TX bytes:13140

allan@allan-VirtualBox:~$ ip address
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue
qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: enp0s3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500
group default qlen 1000
    link/ether 08:00:27:b5:d6:cb brd ff:ff:ff:ff:ff:ff
```

Technet24

```
      inet 10.0.2.15/24 brd 10.0.2.255 scope global dyna
         valid_lft 86130sec preferred_lft 86130sec
      inet6 fe80::57c6:ed95:b3c9:2951/64 scope link
         valid_lft forever preferred_lft forever
```

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	4
	6
	7
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	6
	19
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	5
Portable Command Guide	17

Technet24

Day 22. Wireless Concepts

CCNA 200-301 Exam Topics

• 1.1 Explain the role and function of network components

• 1.11 Describe wireless principles

• 2.6 Describe Cisco Wireless Architectures and AP modes

• 2.7 Describe physical infrastructure connections of WLAN components (AP, WLC, access/trunk ports, LAG)

• 5.9 Describe wireless security protocols (WPA, WPA2, and WPA3)

Key Topics

Wireless specifications are detailed in the IEEE 802.11 family of standards, including wireless topologies, spectrum allocation, and wireless security. Today we review basic wireless network concepts.

Wireless Standards

The IEEE 802.11 WLAN standards define how radio frequencies (RFs) are used for wireless links. To avoid interference, different channels within an RF can be used.

RF Spectrum

The RF spectrum, shown in Figure 22-1, includes all types of radio communications, including the 2.4-GHz and 5-GHz frequencies used by wireless devices.

Technet24

Frequency (Hz)	Frequency Notation
10^{21}	
10^{20}	
10^{19}	
10^{18}	
10^{17}	
10^{16}	
10^{15}	
10^{14}	
10^{13}	
10^{12}	
10^{11}	100 GHz
10^{10}	10 GHz
10^{9}	1 GHz
10^{8}	100 MHz
10^{7}	10 MHz
10^{6}	1 MHz
10^{5}	100 kHz
10^{4}	10 kHz
10^{3}	1 kHz
10^{2}	100 Hz
10^{1}	10 Hz
0	0 Hz

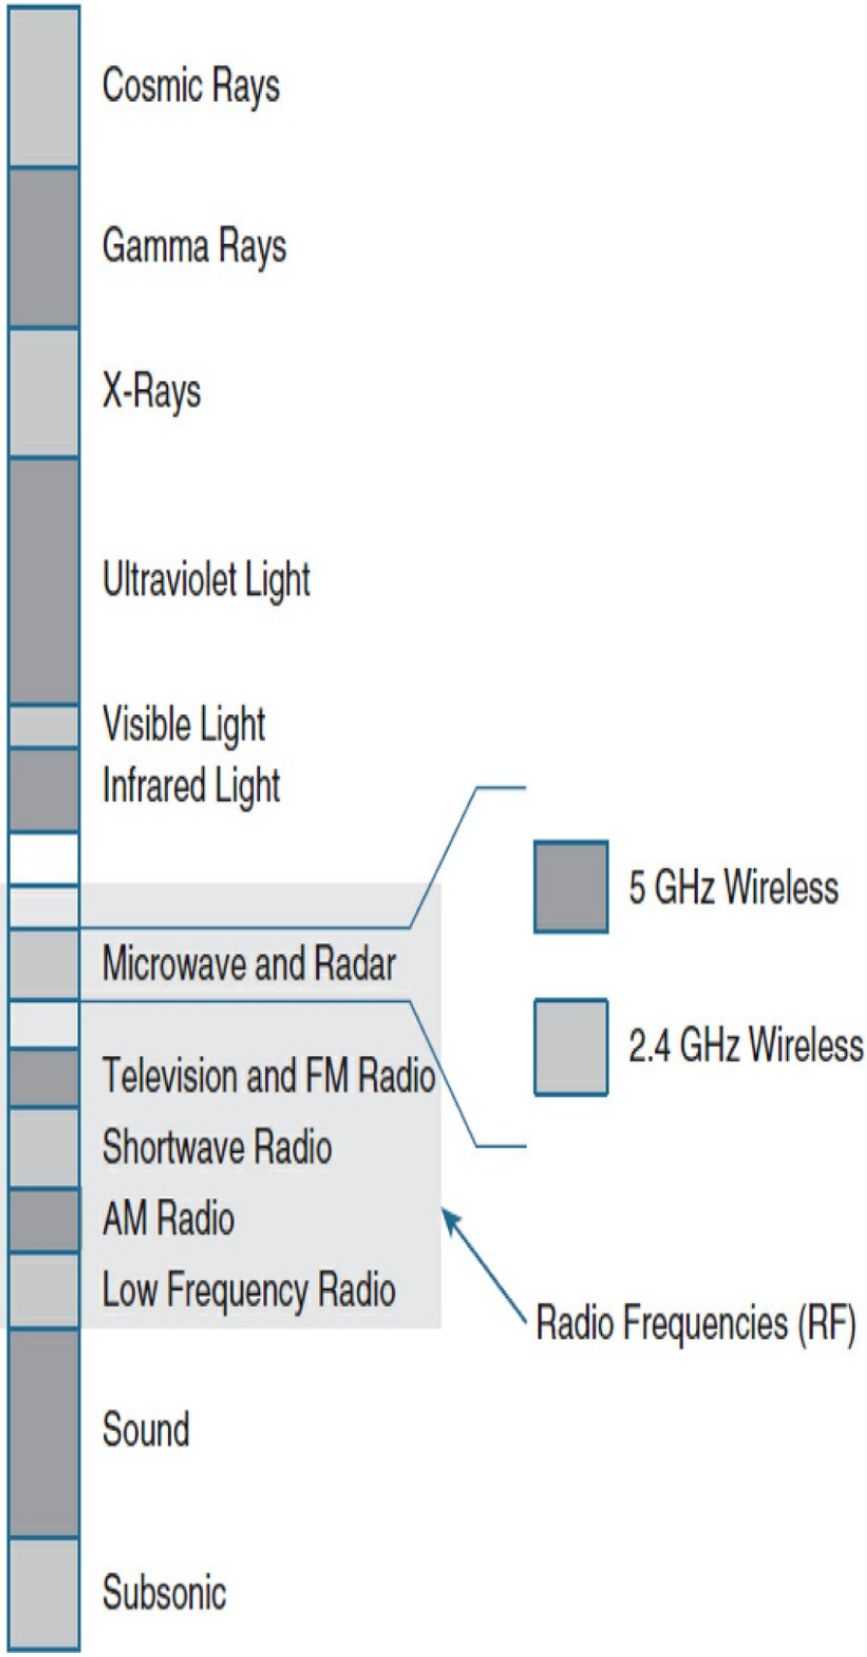

Figure 22-1 *RF Spectrum*

Channels

A frequency range is typically called a *band* of frequencies. For example, a wireless LAN device with a 2.4-GHz antenna can actually use any frequency from 2.4000 to 2.4835 GHz. The 5-GHz band lies between 5.150 and 5.825 GHz.

The bands are further subdivided into frequency channels. Channels become particularly important when the wireless devices in a specific area become saturated. Each channel is known by a channel number and is assigned to a specific frequency. As long as the channels are defined by a national or international standards body, they can be used consistently in all locations. Figure 22-2 and Figure 22-3 show the channel layouts for the 2.4- and 5-GHz bands, respectively.

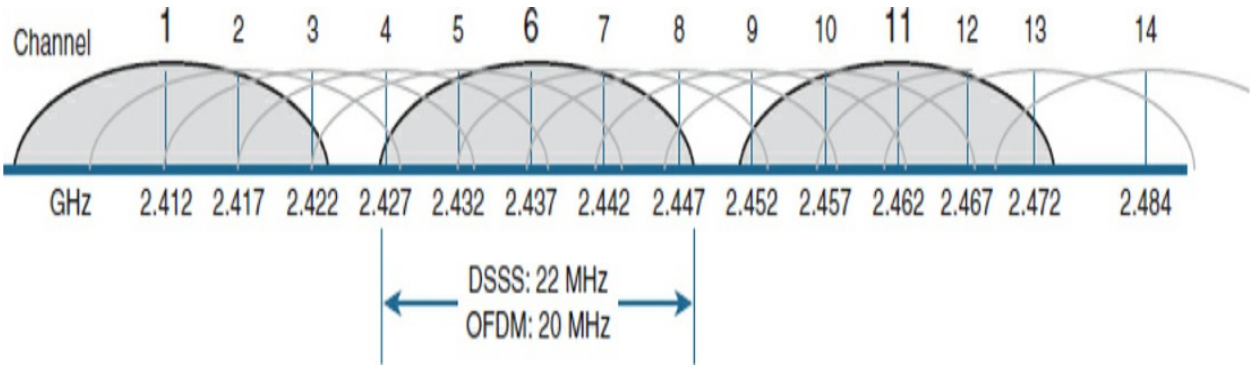

Figure 22-2 *2.4-GHz Channels*

Technet24

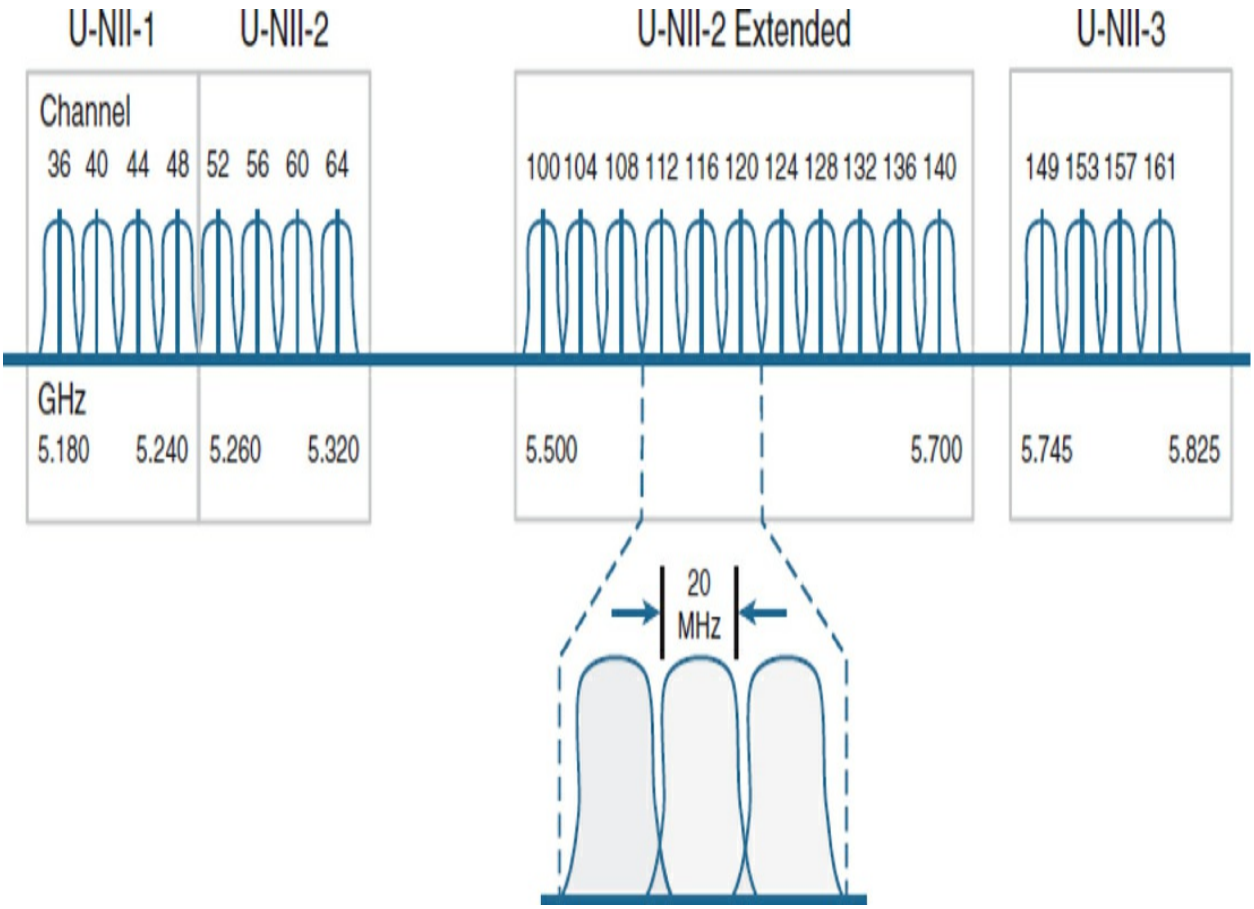

Figure 22-3 *5-GHz Channels*

Notice in Figure 22-3 that the 5-GHz band consists of nonoverlapping channels. Each channel is allocated a frequency range that does not encroach on or overlap the frequencies allocated for any other channel. The same is not true of the 2.4-GHz band in Figure 22-2. The only way to avoid any overlap between adjacent channels is to configure access points (APs) to use only channels 1, 6, and 11.

802.11 Standards

Most of the standards specify that a wireless device must have one antenna to

transmit and receive wireless signals on the specified radio frequency (2.4 GHz or 5 GHz). Some of the newer standards that transmit and receive at higher speeds require APs and wireless clients to have multiple antennas using the multiple input, multiple output (MIMO) technology. MIMO uses multiple antennas as both the transmitter and receiver to improve communication performance. Up to four antennas can be supported.

Various implementations of the IEEE 802.11 standard have been developed over the years. Table 22-1 highlights these standards.

Table 22-1 *Summary of 802.11 Standards*

Technet24

IEEE WLAN Standard	2.4 GHz	5 GHz	6 GHz	Max Data Rate	Description
802.11-1997	Yes	No	No	2 Mbps	The original 802.11 standard, ratified in 1997
802.11a	No	Yes	No	54 Mbps	Introduced in 1999 Small coverage area Less effective at penetrating building structures Not interoperable with 802.11b and 802.11g
802.11b	Yes	No	No	11 Mbps	Introduced in 1999 Longer range than 802.11a Better able to penetrate building structures
802.11g	Yes	No	No	54 Mbps	Introduced in 2003 Backward compatible with 802.11b with reduced bandwidth capacity

802.11n	Yes	Yes	No	600 Mbps	Introduced in 2009 Also known as High Throughput (HT) Distance range of up to 70 m (230 feet) APs and wireless clients require multiple antennas using MIMO technology Backward compatible with 802.11a/b/g devices with limited data rates
802.11ac	No	Yes	No	6.93 Gbps	Introduced in 2013 Also known as Very High Throughput (VHT) Uses MIMO technology

Technet24

					Up to eight antennas can be supported Backward compatible with 802.11a/n devices with limited data rates
802.11ax	Yes	Yes	Yes	4x 802.11ac	Released in 2019 (latest standard) Also known as High-Efficiency Wireless (HEW) Higher data rates and increased capacity Handles many connected devices Improved power efficiency 1 GHz and 7 GHz capable when those frequencies become available

The Wi-Fi Alliance introduced a set of simplified names to identify each Wi-Fi generation, as listed in Table 22-2.

Table 22-2 *Wi-Fi Alliance Generational Names*

Wi-Fi Alliance Designation	Bands Supported	IEEE 802.11 Amendments Supported
Wi-Fi 0	2.4	802.11 (the original)
Wi-Fi 1	2.4	802.11b
Wi-Fi 2	5	802.11a
Wi-Fi 3	2.4	802.11g
Wi-Fi 4	2.4, 5	802.11n
Wi-Fi 5	5	802.11ac
Wi-Fi 6	2.4, 5	802.11ax
Wi-Fi 6E	2.4, 5, 6	802.11ax
Wi-Fi 7	2.4, 5, 6	802.11be

Wireless Topologies

The 802.11 standard identifies two main wireless topology modes: infrastructure mode and Independent Basic Service Set (IBSS). IBSS is also knows as ad hoc mode. With the ubiquity of wireless networks, mesh topologies are now common.

Infrastructure Mode

With infrastructure mode, wireless clients interconnect via an AP. Figure 22-

Technet24

4 illustrates infrastructure mode terminology. Notice that the configuration of the APs to share the same SSID allows wireless clients to roam between BSAs.

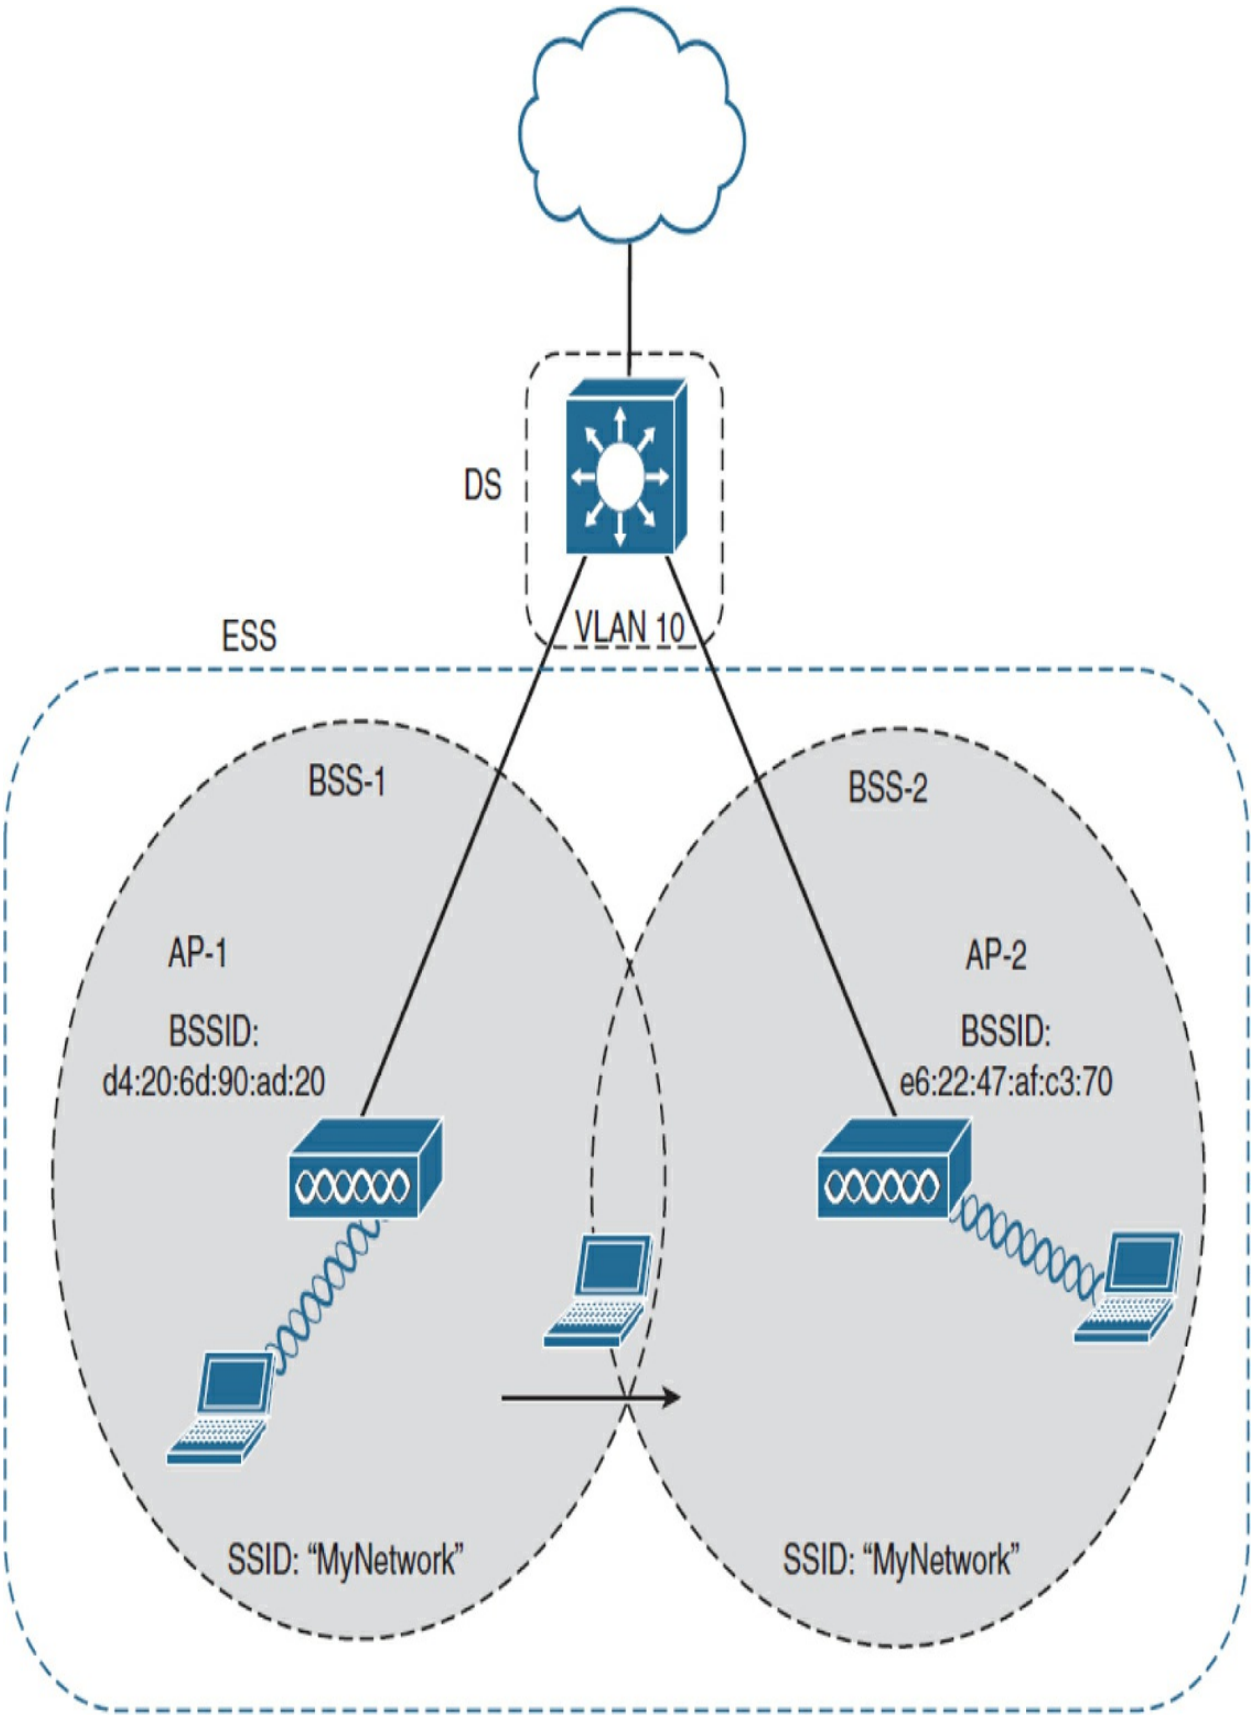
DS
VLAN 10
ESS
BSS-1
BSS-2
AP-1
BSSID:
d4:20:6d:90:ad:20
AP-2
BSSID:
e6:22:47:af:c3:70
SSID: "MyNetwork"
SSID: "MyNetwork"

Technet24

Figure 22-4 *Example of ESS Infrastructure Mode*

Infrastructure mode terminology includes the following:

• **Basic service set (BSS):** This consists of a single AP interconnecting all associated wireless clients.

• **Basic service area (BSA):** This is the area that is bound by the reach of the AP's signal. The BSA is also called a *cell* (the gray area in Figure 22-4).

• **Basic service set identifier (BSSID):** This is the unique, machine-readable identifier for the AP that is in the format of a MAC address and is usually derived from the AP's wireless MAC address.

• **Service set identifier (SSID):** This is a human-readable, non-unique identifier used by the AP to advertise its wireless service.

• **Distribution system (DS):** APs connect to the network infrastructure using the wired DS, such as Ethernet. An AP with a wired connection to the DS is responsible for translating frames between 802.3 Ethernet and 802.11 wireless protocols.

• **Extended service set (ESS):** When a single BSS provides insufficient coverage, two or more BSSs can be joined through a common DS into an ESS. An ESS is the union of two or more BSSs interconnected by a wired DS. Each ESS is identified by its SSID, and each BSS is identified by its

BSSID.

IBSS, or Ad Hoc Mode

In the 802.11 standard, Independent Basic Service Set (IBSS) is defined as two devices connected wirelessly in a peer-to-peer (P2P) manner without the use of an AP. One device takes the role of advertising the wireless network to clients. The IBSS allows two devices to communicate directly without the need for any other wireless devices, as shown in Figure 22-5. IBSSs do not scale well beyond 8 to 10 devices.

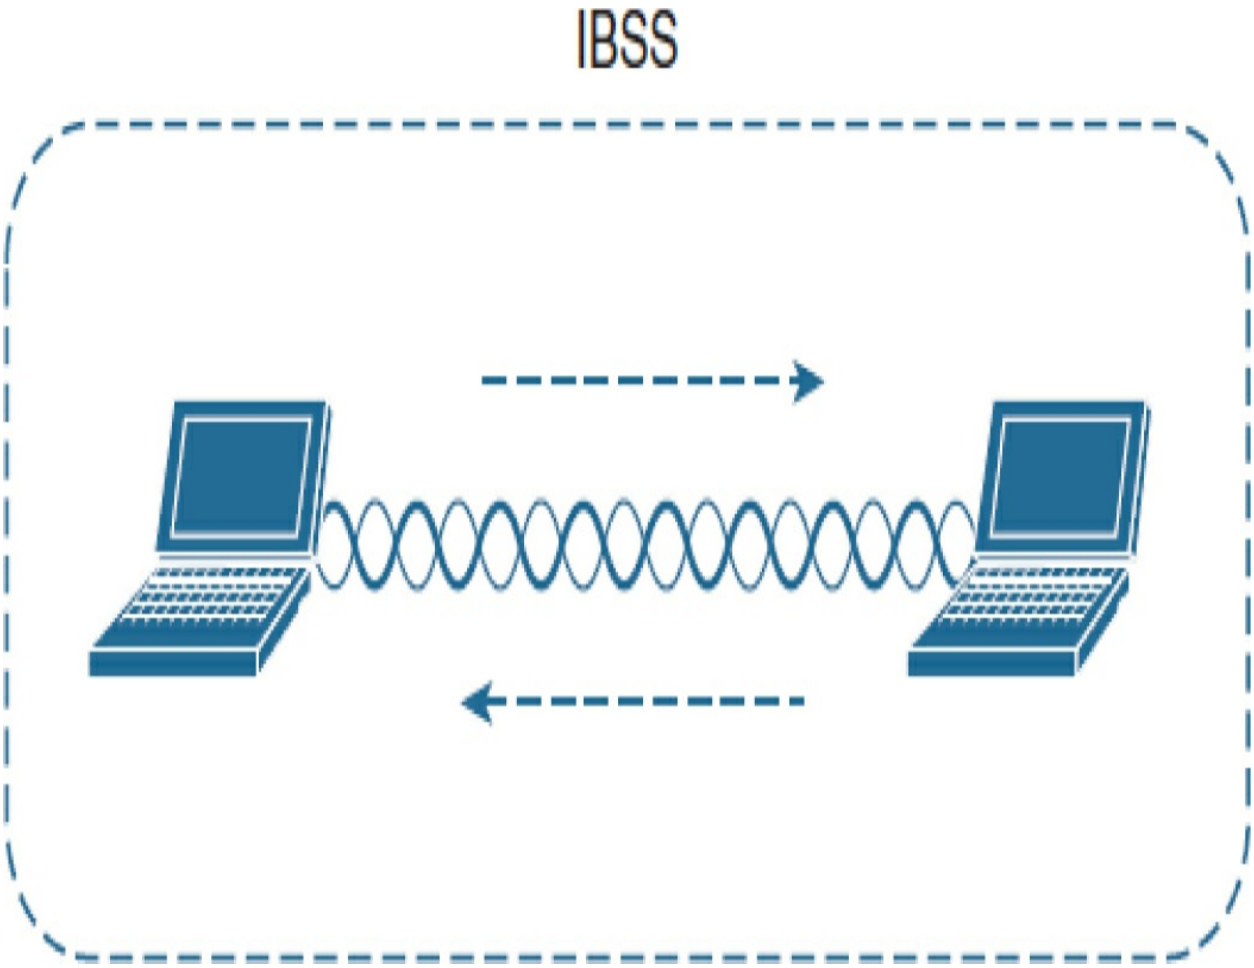

Figure 22-5 *802.11 Independent Basic Service Set*

Technet24

Mesh

Having a wired DS connecting all APs is not always practical or necessary. Instead, APs can be configured to connect in mesh mode. In this mode, APs bridge client traffic between each other, as shown in Figure 22-6.

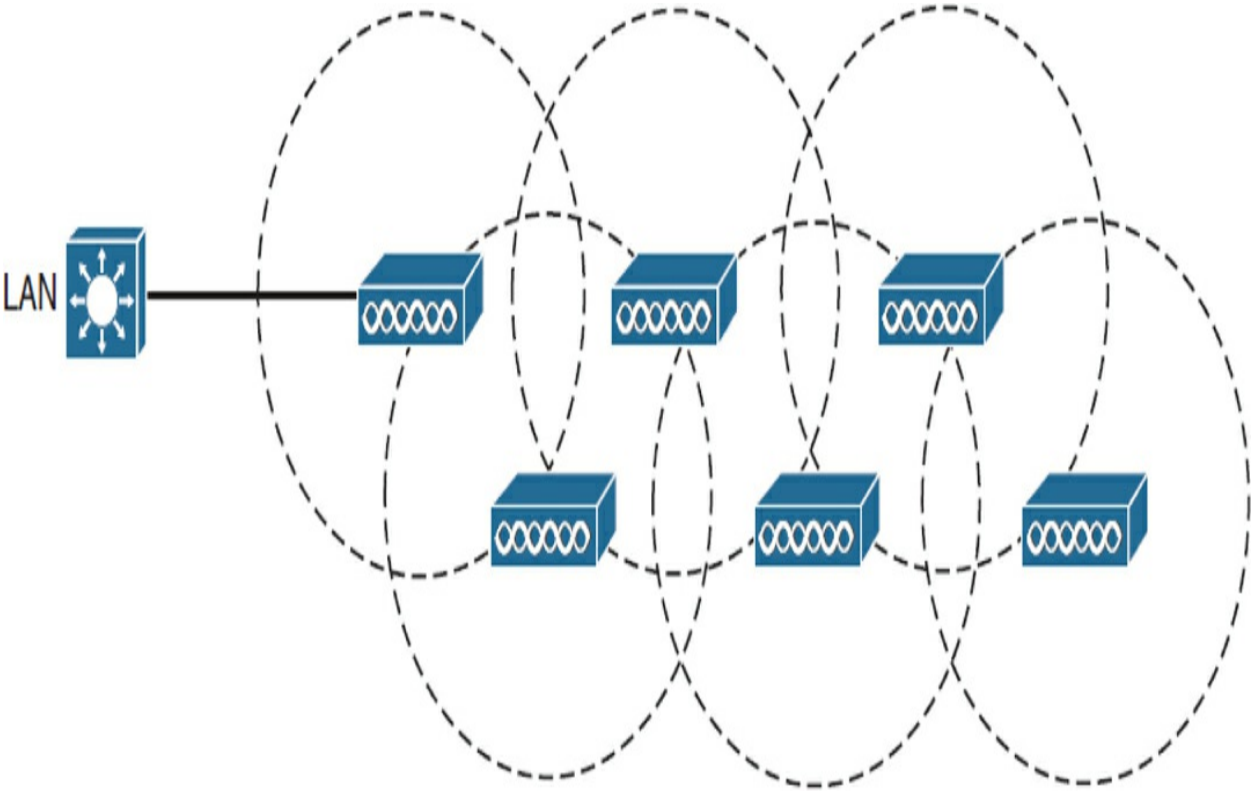

Figure 22-6 *Example of a Wireless Mesh Network*

Each AP in the mesh maintains a BSS on one channel used by wireless clients. Then the APs bridge between each other using other channels. The mesh network runs its own dynamic routing protocol to determine the best path to the wired network.

AP Architectures

APs can be networked together in a variety of architectures. The size and scalability of the network determine which architecture is most suited for a given implementation.

Autonomous AP Architecture

An autonomous AP is a self-contained device with both wired and wireless hardware so that it can bridge to the wired VLAN infrastructure wireless clients that belong to SSIDs, as shown in Figure 22-7. Each autonomous AP must be configured with a management IP address so that it can be remotely accessed using Telnet, SSH, or a web interface. Each AP must be individually managed and maintained unless you use a management platform such as Cisco DNA Center.

Technet24

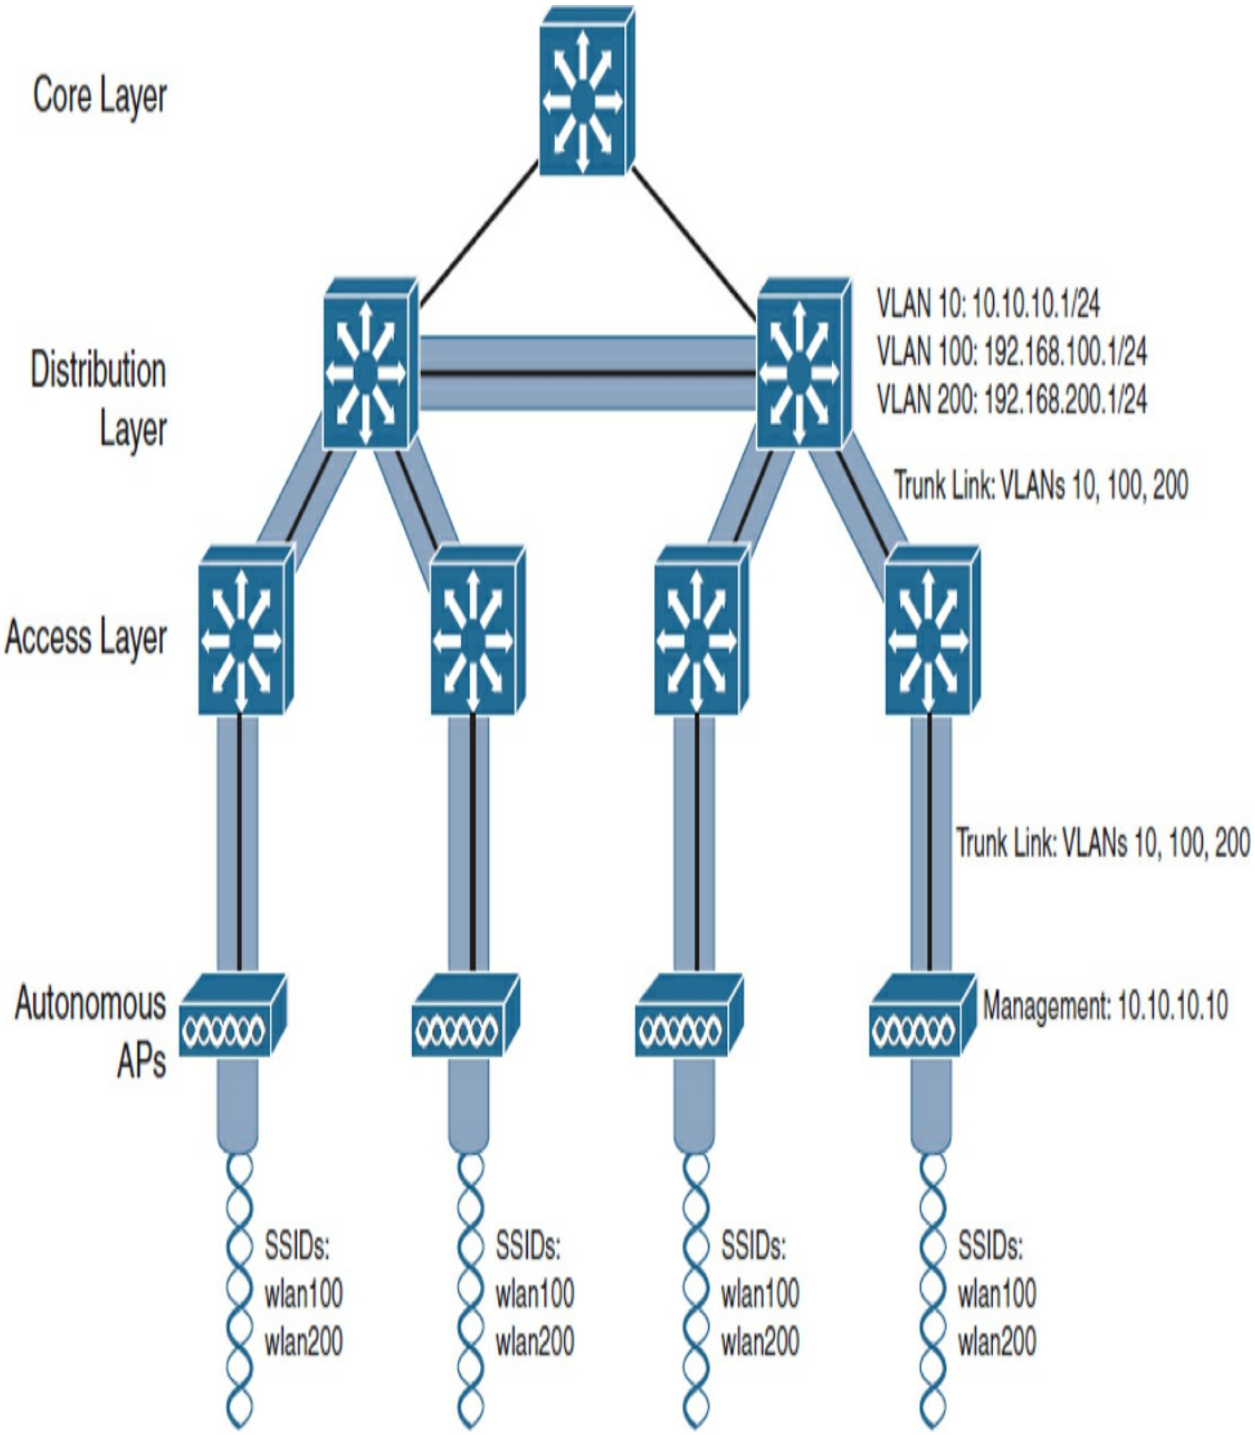

Figure 22-7 *Autonomous APs*

Cloud-Based AP Architecture

Cloud-based AP management is an alternative to purchasing a management

platform. The AP management function is pushed into the Internet cloud. For example, Cisco Meraki is a cloud-based AP management service that allows you to automatically deploy Cisco Meraki APs. These APs can then be managed from the Meraki cloud web interface (dashboard). In Figure 22-8, the same APs shown in Figure 22-7 are now managed in the cloud.

Technet24

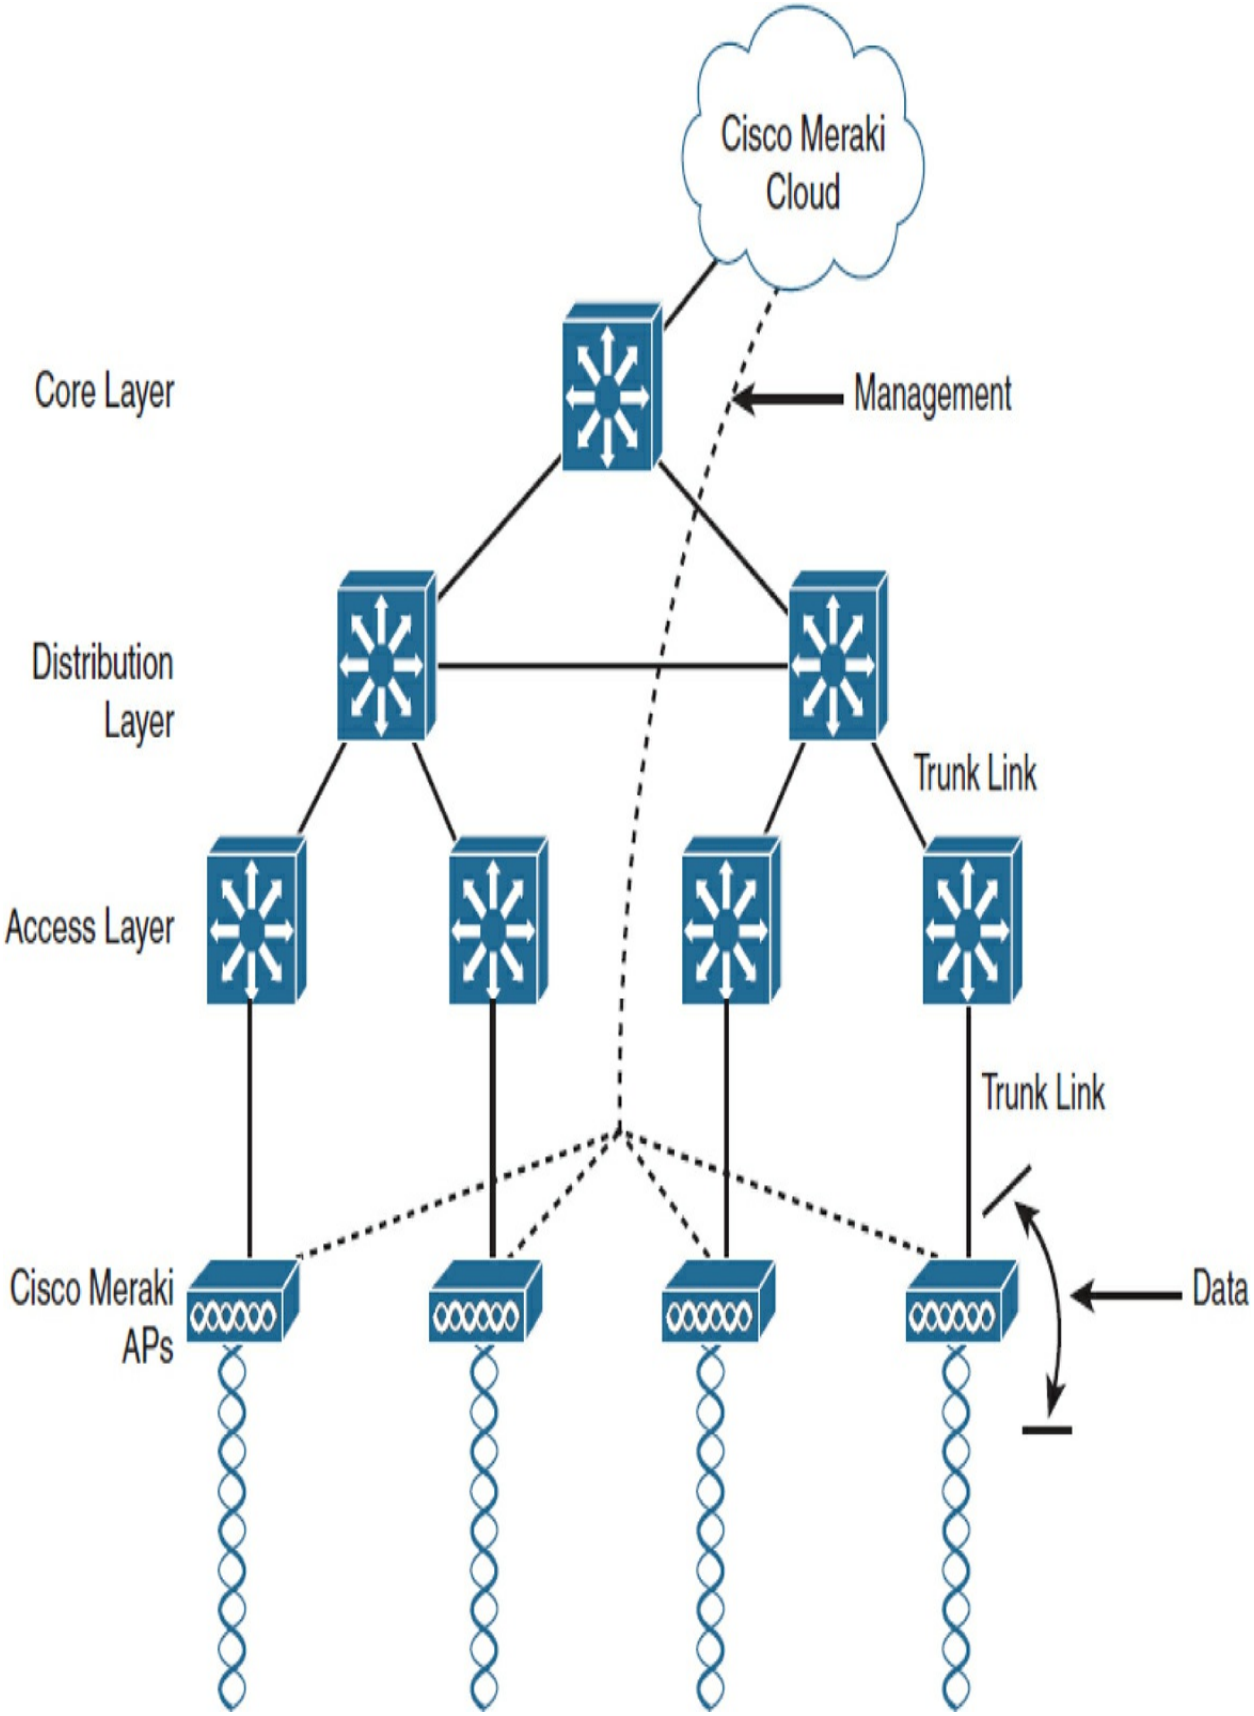
Cisco Meraki
Cloud
Core Layer
Management
Distribution
Layer
Trunk Link
Access Layer
Trunk Link
Cisco Meraki
APs
Data

Figure 22-8 *Cisco Meraki Cloud-Based AP Management*

Notice that there are two distinct paths for data traffic and for management traffic, corresponding to the following two functions:

• **A control plane:** Traffic used to control, configure, manage, and monitor the AP itself

• **A data plane:** End-user traffic passing through the AP

Lightweight AP Architectures

Wireless LAN controllers (WLCs) use Lightweight Access Point Protocol (LWAPP) to communicate with lightweight APs (LAPs), as shown in Figure 22-9. LAPs are useful in situations where many APs are required in the network. They are “lightweight” because they only perform the 802.11 wireless operation for wireless clients. Each LAP is automatically configured and managed by the WLC.

Technet24

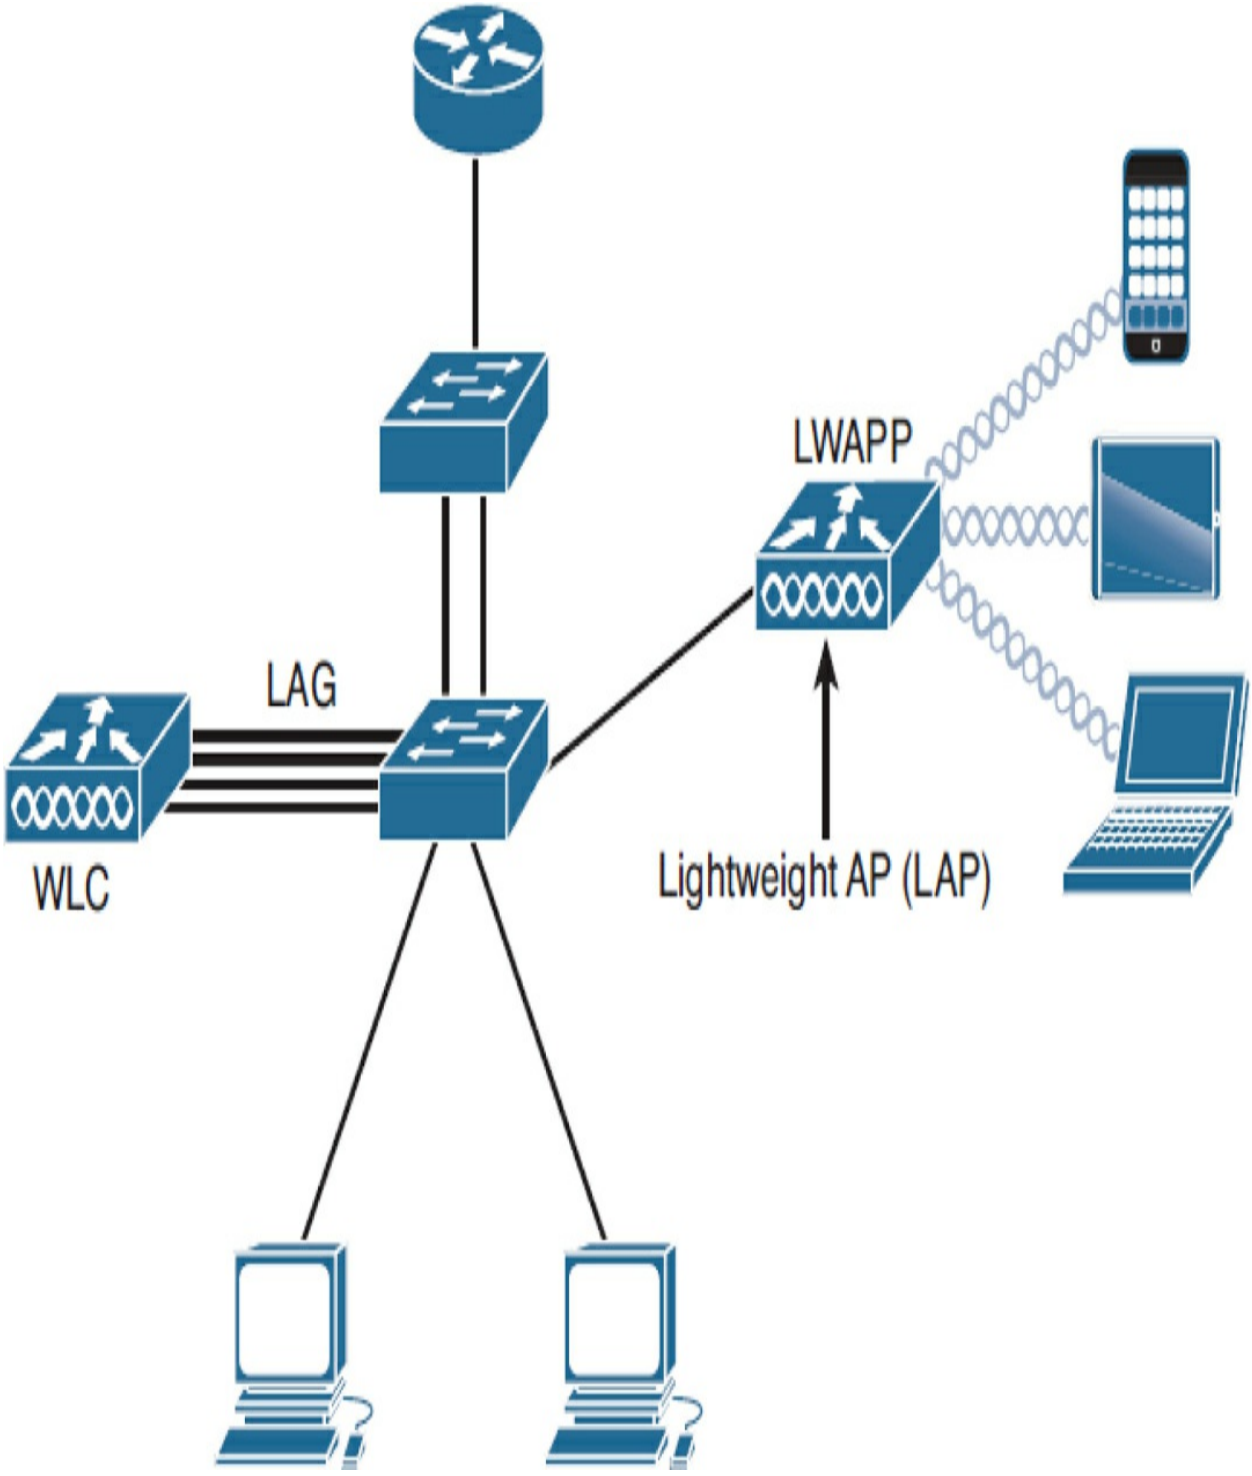

Figure 22-9 *Controller-Based AP Architecture*

Notice in Figure 22-9 that the WLC has four ports connected to the switching infrastructure. These four ports are configured as a link aggregation group

(LAG) so they can be bundled together. Much like EtherChannel, LAG provides redundancy and load balancing.

CAPWAP Operation

The division of labor between the WLC and LAPs is known as *split-MAC architecture*. The LAP must interact with wireless clients on some low level, known as the Media Access Control (MAC) layer. These functions must stay with the LAP hardware, closest to the clients. The management functions are not integral to handling frames but are things that should be centrally administered. Therefore, those functions can be moved to a centrally located platform away from the AP. Table 22-3 summarizes MAC functions of the LAP and WLC.

Table 22-3 *Split-MAC Functions of the AP and WLC*

Technet24

AP MAC Functions	WLC MAC Functions
Beacons and probe responses	Authentication
Packet acknowledgments and retransmissions	Association and re-association of roaming clients
Frame queueing and packet prioritization	Frame translation to other protocols
MAC layer data encryption and decryption	Termination of 802.11 traffic on a wired interface

LWAPP has been replaced with the Control and Provisioning of Wireless Access Points (CAPWAP) tunneling protocol to implement these split-MAC functions. CAPWAP uses two tunnels—one for control and one for data—as shown in Figure 22-10 and described in the list that follows:

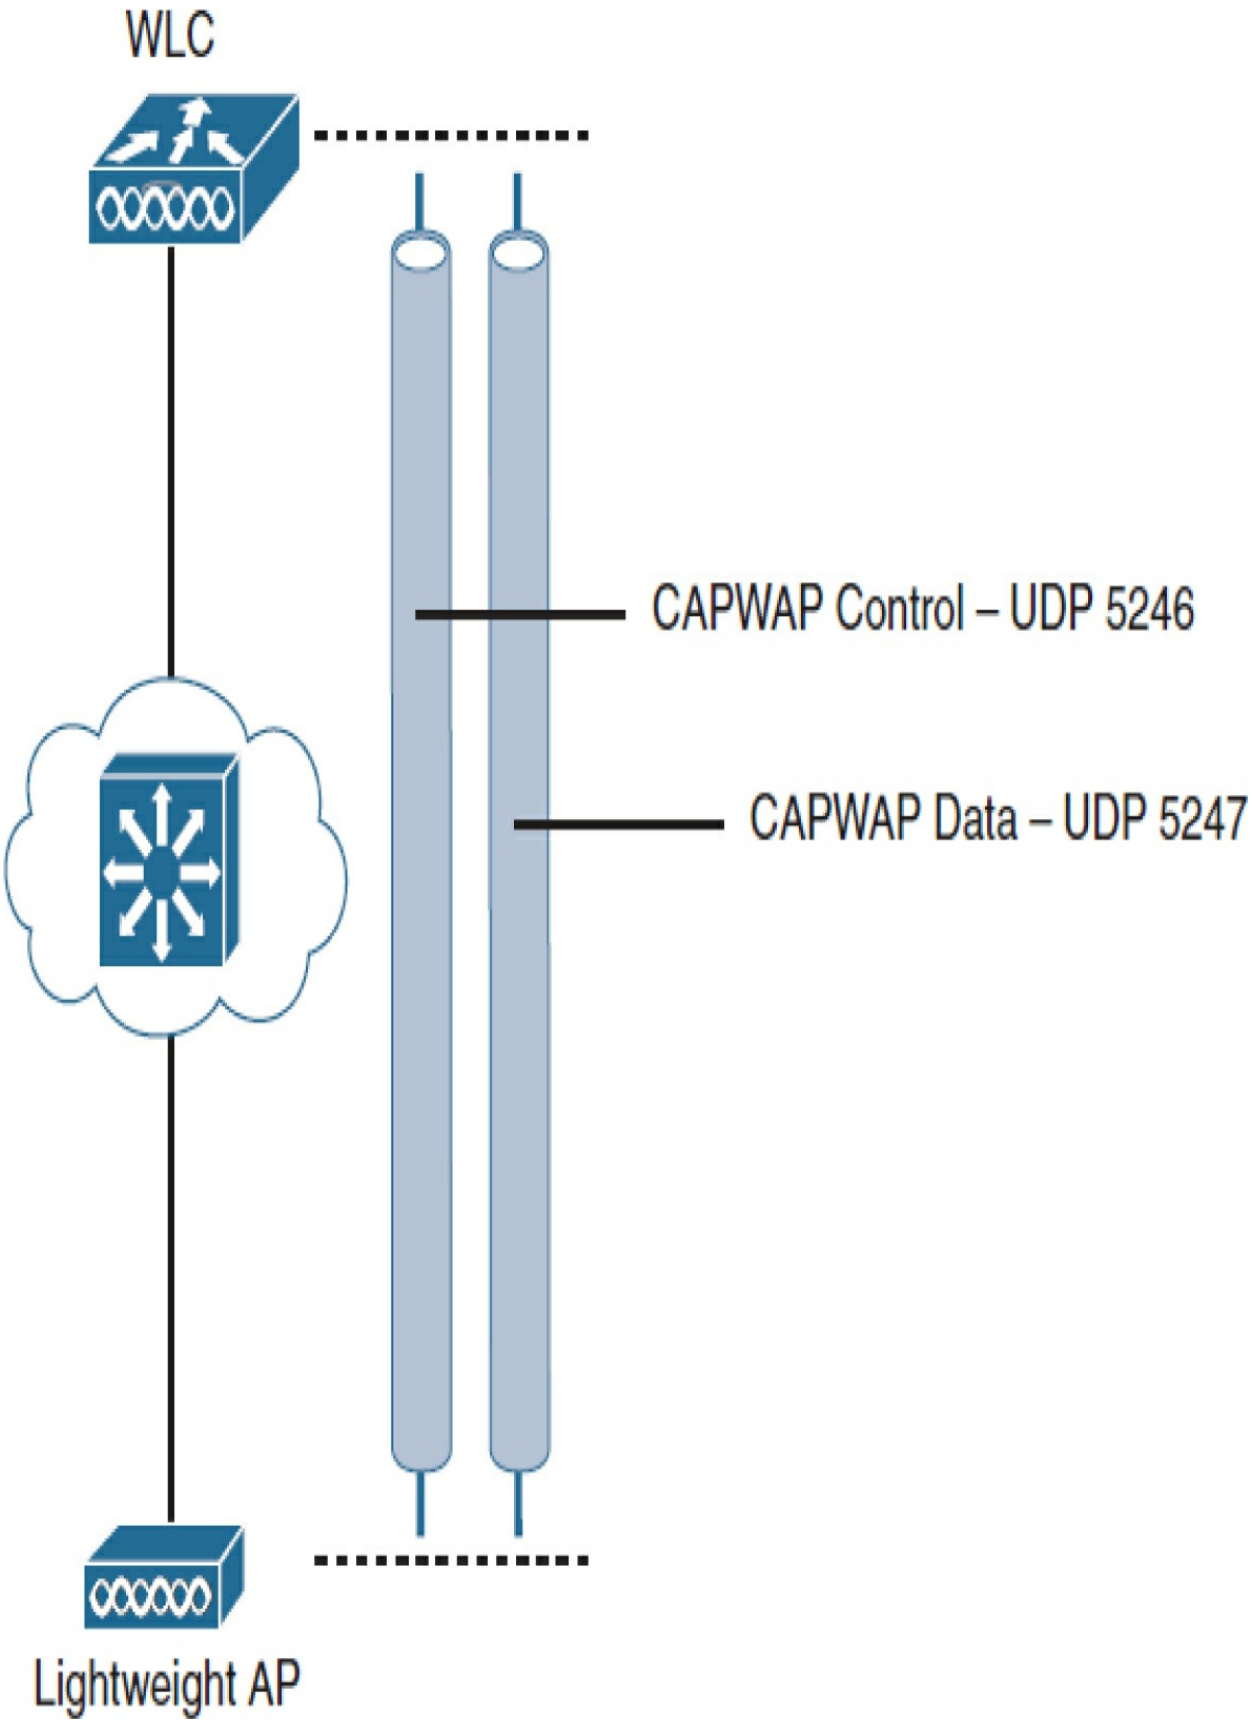
WLC
CAPWAP Control – UDP 5246
CAPWAP Data – UDP 5247
Lightweight AP

Technet24

Figure 22-10 *CAPWAP Control and Data Tunnels*

• **CAPWAP control message tunnel:** Carries exchanges that are used to configure the LAP and manage its operation. The control messages are authenticated and encrypted, so the LAP is securely controlled by only the appropriate WLC and then transported over the control tunnel using UDP port 5246.

• **CAPWAP data tunnel:** Used for packets traveling to and from wireless clients that are associated with the AP. Data packets are transported over the data tunnel using UDP port 5247 but are not encrypted by default. When data encryption is enabled for a LAP, packets are protected with Datagram Transport Layer Security (DTLS).

Wireless Security Protocols

Wireless traffic is inherently different from traffic traveling over a wired infrastructure. Any wireless device operating in the same frequency can hear the frames and potentially read them. Therefore, WLANs need to be secured to allow only authorized users and devices and to prevent eavesdropping and tampering of wireless traffic.

Wireless Authentication Methods

For wireless devices to communicate over a network, they must first associate with the AP. An important part of the 802.11 process is discovering

a WLAN and subsequently connecting to it. During this process, transmitted frames can reach any device within range. If the wireless connection is not secured, then others can read the traffic, as shown in Figure 22-11.

Technet24

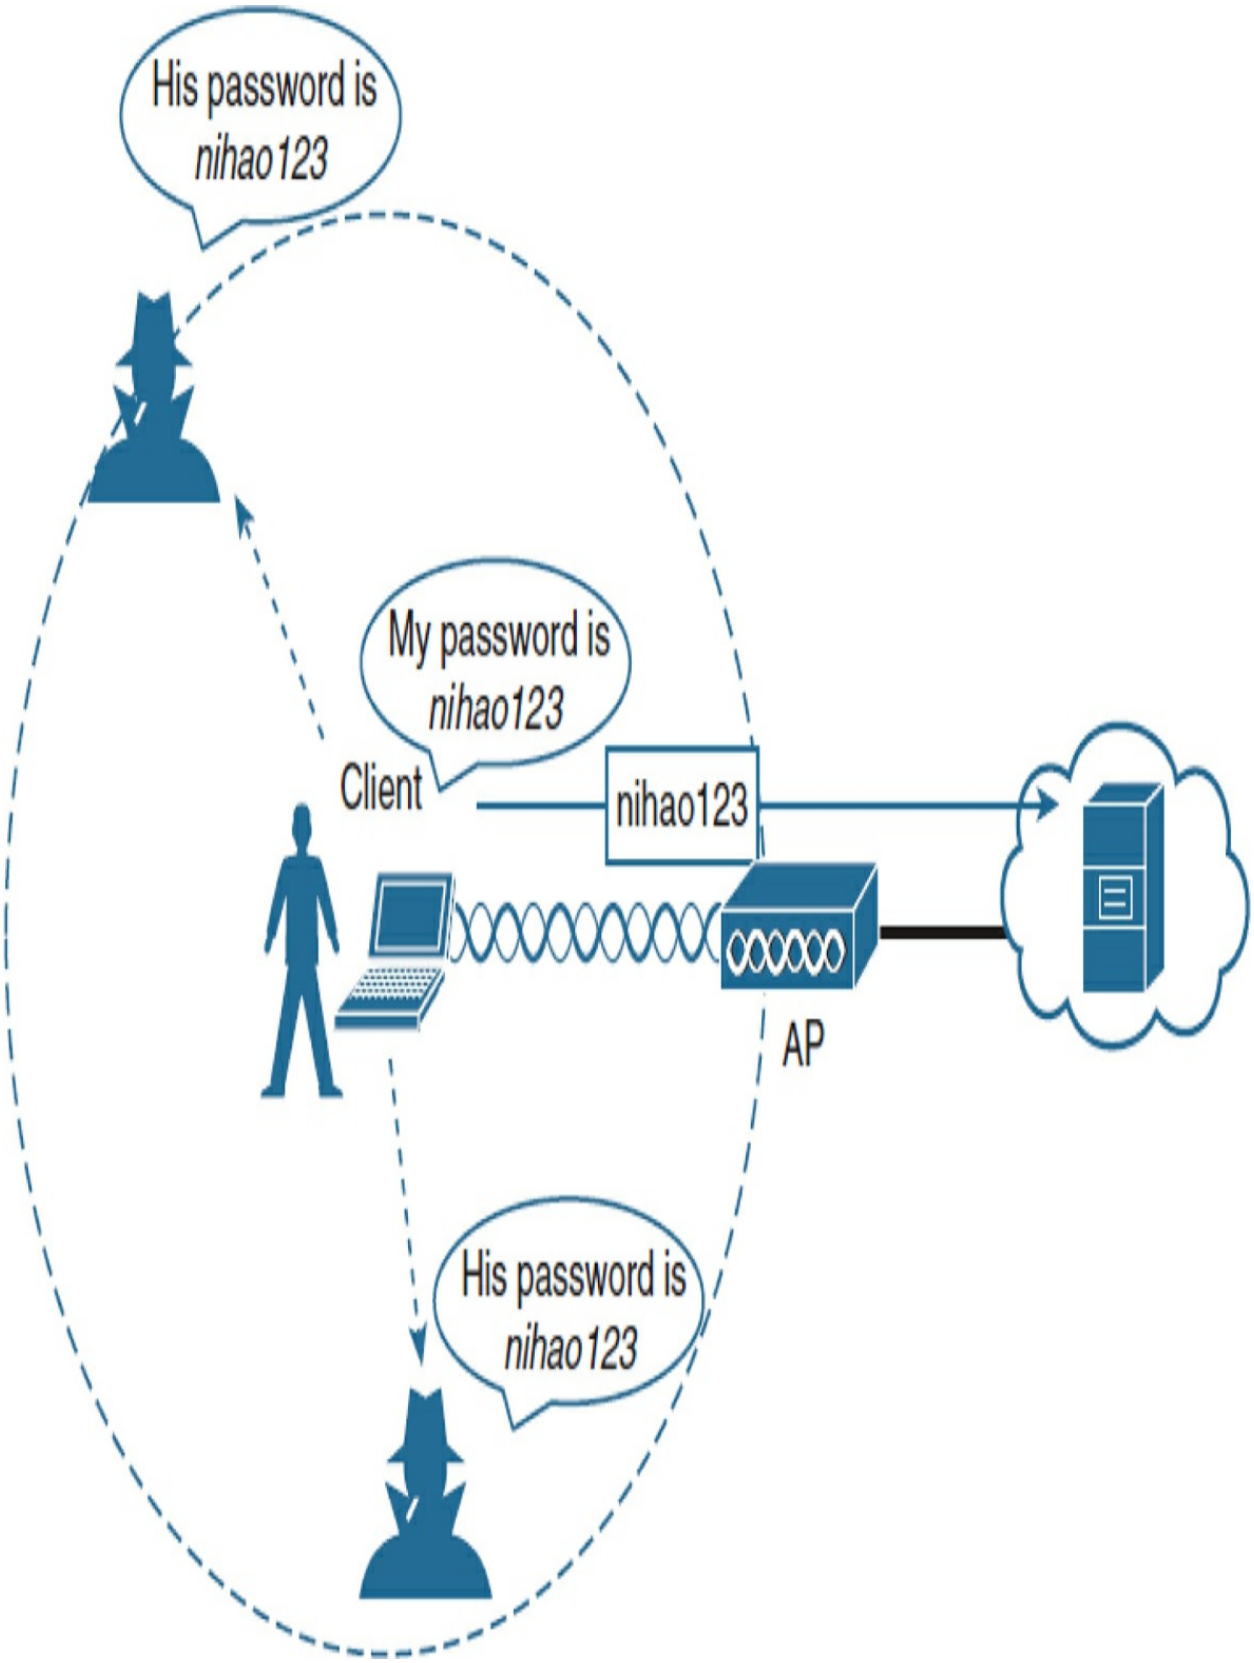

Figure 22-11 *Open Wireless Network*

The best way to secure a wireless network is to use authentication and encryption systems.

Two types of authentication were introduced with the original 802.11 standard:

• **Open system authentication:** Should only be used in situations where security is of no concern. The wireless client is responsible for providing security such as by using a virtual private network (VPN) to connect securely.

• **Shared key authentication:** Provides mechanisms shown in Table 22-4 to authenticate and encrypt data between a wireless client and an AP. However, the password must be pre-shared between the parties to allow connection.

Table 22-4 *Shared Key Authentication Methods*

Technet24

Authentication Method	Description
Wired Equivalent Privacy (WEP)	The original 802.11 specification designed to secure the data using the Rivest Cipher 4 (RC4) encryption method with a static key. However, the key never changes when exchanging packets. This makes WEP easy to hack. WEP is no longer recommended and should never be used.
Wi-Fi Protected Access (WPA)	A Wi-Fi Alliance standard that uses WEP but secures the data with the much stronger Temporal Key Integrity Protocol (TKIP) encryption algorithm. TKIP changes the key for each packet, making it much more difficult to hack.
WPA2	The current industry standard for securing wireless networks. It uses the Advanced Encryption Standard (AES) for encryption. AES is currently considered the strongest encryption protocol.
WPA3	The next generation of Wi-Fi security. All WPA3-enabled devices use the latest security methods, disallow outdated legacy protocols, and require the use of Protected Management Frames (PMF). However, devices with WPA3 are not yet readily available.

WPA and WPA2

Home routers typically have two choices for authentication: WPA and WPA2. WPA2 is the stronger of the two. WPA2 authentication methods included the following:

• **Personal:** Intended for home or small office networks, users authenticate using a pre-shared key (PSK). Wireless clients authenticate with the wireless router using a pre-shared password. No special authentication server is required.

• **Enterprise:** Intended for enterprise networks but requires a Remote Authentication Dial-In User Service (RADIUS) authentication server. Although more complicated to set up, it provides additional security. The device must be authenticated by the RADIUS server, and then users must authenticate using the 802.1X standard, which uses Extensible Authentication Protocol (EAP) for authentication.

802.1X/EAP

With open and WEP authentication, wireless clients are authenticated locally at the AP without further intervention. The scenario changes with 802.1X: The client uses open authentication to associate with the AP, and then the client authentication process occurs at a dedicated authentication server. Figure 22-11 shows the three-party 802.1X arrangement, which consists of the following entities:

Technet24

• **Supplicant:** The client device that is requesting access.

• **Authenticator:** The network device that provides access to the network. In Figure 22-11, the AP forwards the supplicant's message to the WLC.

• **Authentication server (AS):** The device that permits or denies network access based on a user database and policies (usually a RADIUS server).

WPA3

WPA3 includes four features:

• **WPA3-Personal:** In WPA2-Personal, threat actors can listen in on the "handshake" between a wireless client and the AP and use brute-force attacks to try to guess the PSK. WPA3-Personal thwarts such attacks by using Simultaneous Authentication of Equals (SAE), a feature specified in the IEEE 802.11-2016. The PSK is never exposed, making it impossible for the threat actor to guess.

• **WPA3-Enterprise:** WPA3-Enterprise still uses 802.1X/EAP authentication. However, it requires the use of a 192-bit cryptographic suite and eliminates the mixing of security protocols for previous 802.11 standards. WPA3-Enterprise adheres to the Commercial National Security Algorithm (CNSA) suite, which is commonly used in high-security Wi-Fi networks.

• **Open networks:** Open networks in WPA2 send user traffic in

unauthenticated plaintext. In WPA3, open or public Wi-Fi networks still do not use any authentication. However, they do use Opportunistic Wireless Encryption (OWE) to encrypt all wireless traffic.

• **IoT onboarding:** Although WPA2 included Wi-Fi Protected Setup (WPS) to quickly onboard devices that were not previously configured, WPS is vulnerable to a variety of attacks and is not recommended. Furthermore, IoT devices are typically headless, meaning they have no built-in GUI for configuration and need any easy way to get connected to the wireless network. Device Provisioning Protocol (DPP) was designed to address this need. Each headless device has a hard-coded public key. The key is typically stamped on the outside of the device or its packaging as a Quick Response (QR) code. The network administrator can scan the QR code and quickly onboard the device. Although DPP is not strictly part of the WPA3 standard, it will replace WPS over time.

Wireless Encryption Methods

Encryption is used to protect data. An intruder may be able to captured encrypted data, but he or she would not be able to decipher it in any reasonable amount of time. The following encryption protocols are used with wireless authentication:

• **Temporal Key Integrity Protocol (TKIP):** TKIP is the encryption method used by WPA. It provides support for legacy WLAN equipment and addresses the original flaws associated with the 802.11 WEP encryption

Technet24

method. It makes use of WEP but encrypts the Layer 2 payload using TKIP and carries out a message integrity check (MIC) in the encrypted packet to ensure that the message has not been altered.

• **Advanced Encryption Standard (AES):** AES is the encryption method used by WPA2. It is the preferred method because it is a very strong method of encryption. It uses Counter Cipher Mode with Block Chaining Message Authentication Code Protocol (CCMP), which allows destination hosts to recognize if the encrypted and nonencrypted bits have been altered.

• **Counter/CBC-MAC Protocol (CCMP):** CCMP is the encryption protocol used by WPA2. It provides enhanced data confidentiality and integrity by using AES. CCMP operates using Counter Mode for encryption and Cipher Block Chaining Message Authentication Code (CBC-MAC) for message integrity, ensuring that both encrypted and non-encrypted parts of the packet have not been altered. This makes CCMP a highly secure and reliable protocol for protecting wireless communications.

• **The Galois/Counter Mode Protocol (GCMP):** GCMP is the encryption protocol used by WPA3. It enhances security and efficiency compared to CCMP. GCMP combines the Galois Message Authentication Code (GMAC) for data integrity and authenticity with the Counter Mode for encryption. This dual approach ensures that encrypted data remains confidential, and any tampering with the data can be detected. GCMP provides robust protection against attacks and improves performance, making it suitable for modern

wireless networks that require higher security standards.

Table 22-5 summarizes the basic differences between WPA, WPA2, and WPA3. Each successive version is meant to replace prior versions and offer better security features. You should avoid using WPA and use WPA2 instead —at least until WPA3 becomes widely available on wireless client devices, APs, and WLCs.

Table 22-5 *Wireless Authentication and Encryption Comparison*

Feature	WPA	WPA2	WPA3
Authentication with pre-shared keys?	Yes	Yes	Yes
Authentication with 802.1X?	Yes	Yes	Yes
Encryption and MIC with TKIP?	Yes	No	No
Encryption and MIC with AES and CCMP?	Yes	Yes	No
Encryption and MIC with AES and GCMP?	No	No	Yes

Study Resources

Technet24

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	12
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	1–3
Portable Command Guide	23

Day 21. WLAN Configuration

CCNA 200-301 Exam Topics

• 2.9 Interpret the wireless LAN GUI configuration for client connectivity, such as WLAN creation, security settings, QoS profiles, and advanced settings

• 5.10 Configure and verify WLAN within the GUI using WPA2 PSK

Key Topics

Today we review the steps to configure a wireless LAN controller (WLC). The figures show the graphical user interface (GUI) and menus from a Cisco 3504 Wireless Controller (see Figure 21-1). However, other WLC models have similar menus and features.

Figure 21-1 *Cisco 3504 Wireless Controller*

Technet24

Logging Into a Cisco WLC

In order to configure a WLC, you need to be able to access it. The WLC requires an initial configuration and a management IP address before you can access it with a web browser through HTTP or HTTPS. This initial configuration requires a console connection. The WLC can also be further configured from the command-line interface (CLI) using Telnet or SSH. However, the CCNA exam focuses on GUI access to the WLC. Therefore, the rest of this day focuses on logging into and configuring a WLC that already has its basic configuration.

Log in to the WLC web interface, as shown in Figure 21-2.

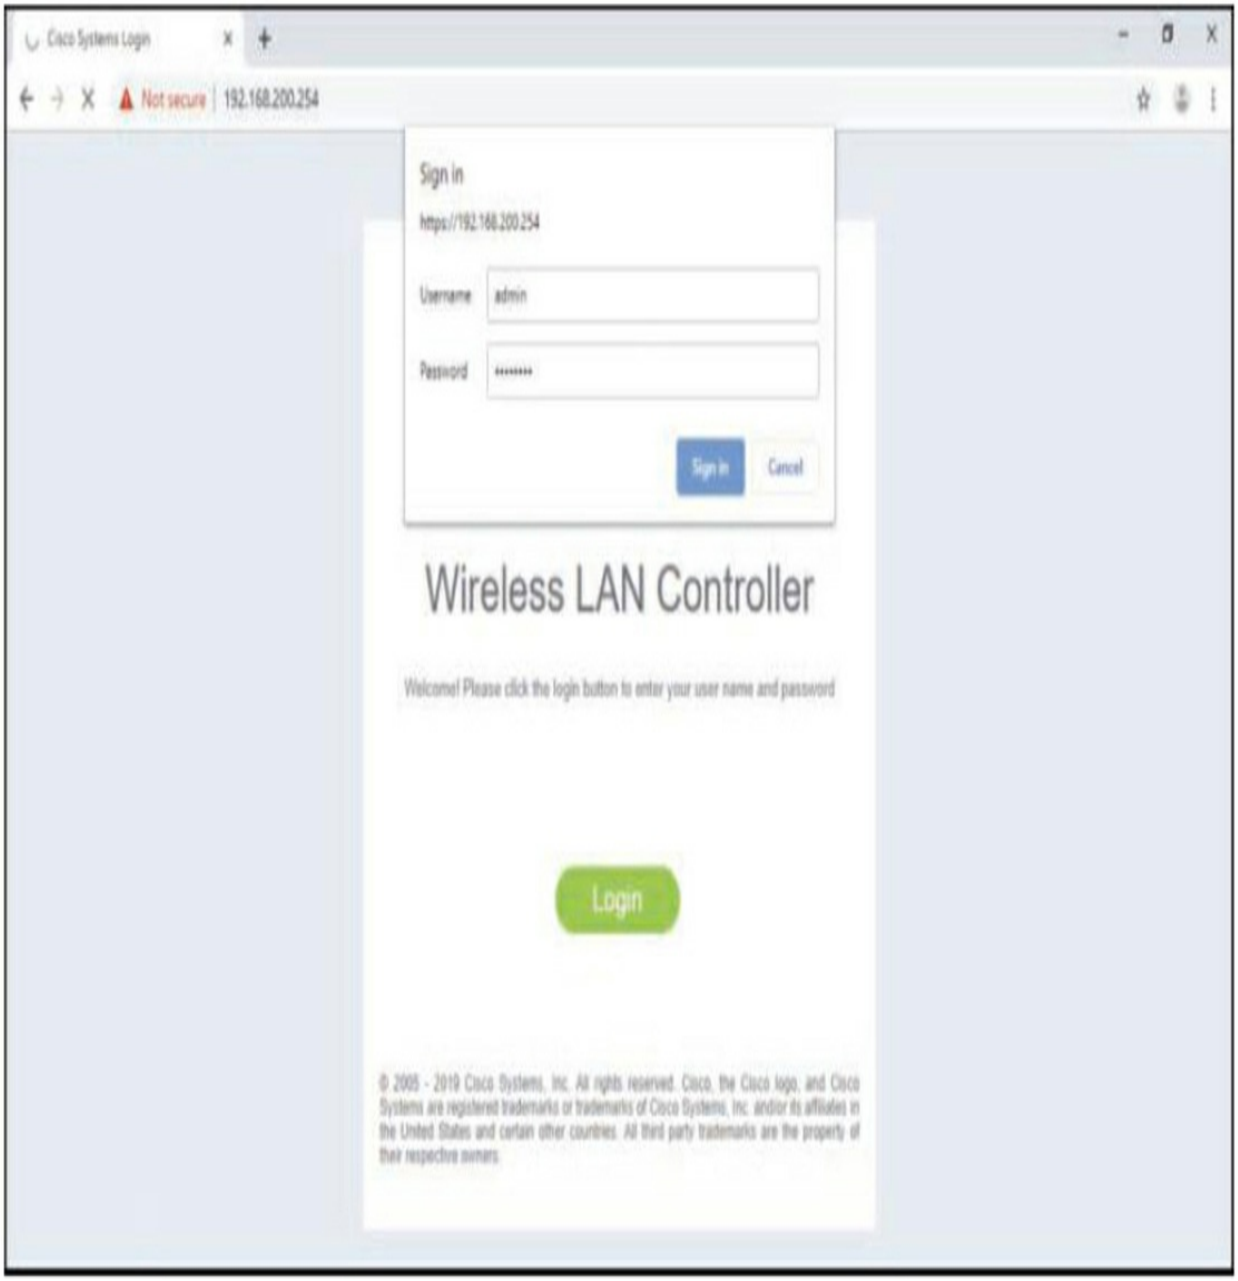

Figure 21-2 *WLC Login Window*

The Network Summary page is a dashboard that provides a quick overview of the number of configured wireless networks, associated access points (APs), and active clients, as shown in Figure 21-3.

Technet24

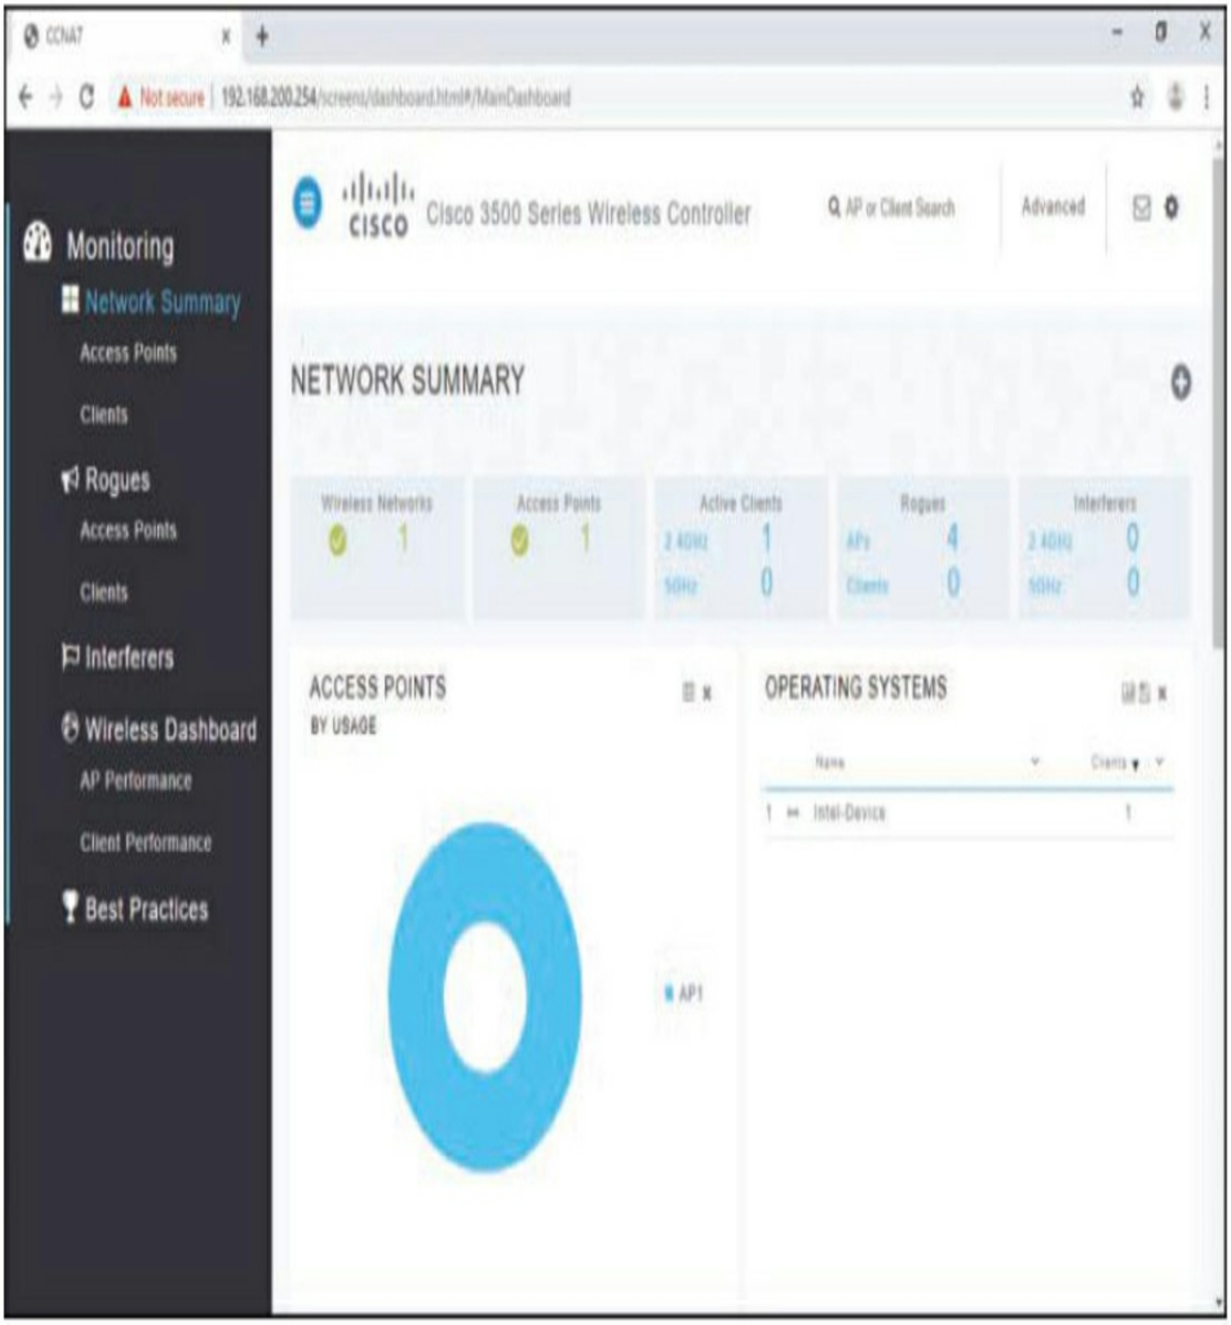

Figure 21-3 *Network Summary Dashboard*

In the menu on the left side of the Network Summary page, click **Access Points** to view an overall picture of AP system information and performance, as shown in Figure 21-4.

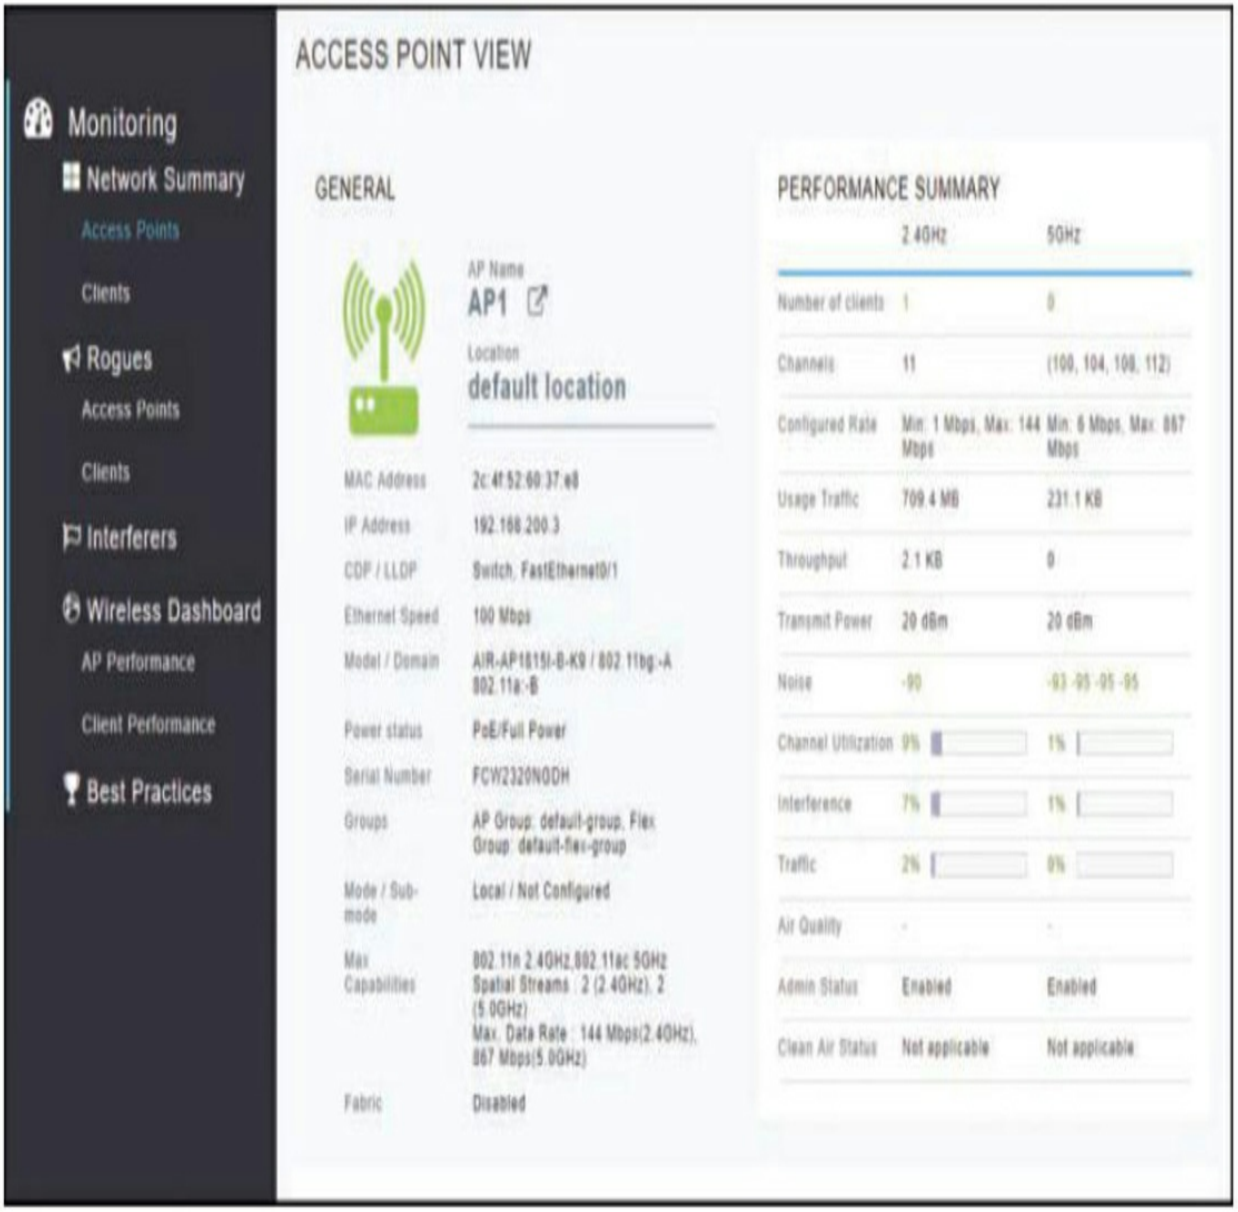

Figure 21-4 *Access Points*

Click **Advanced** to access the advanced Summary page, as shown in Figure 21-5. From here, you can access all the features of the WLC.

Technet24

Figure 21-5 *WLC Advanced Features*

Configuring a WLC with a WLAN

You can configure a WLAN directly on the Cisco 3504 Wireless Controller so that it serves as an AP for wireless clients. However, a WLC is more commonly used in enterprise networks to manage a number of APs.

Configuring a RADIUS Server

An enterprise WLAN typically uses a RADIUS server for user and device authentication before allowing wireless clients to associate with an AP. To

configure the WLC with the RADIUS server information, click the **SECURITY** tab > **RADIUS > Authentication** to navigate to the screen in Figure 21-6. Click **New** to add the RADIUS server.

Figure 21-6 *Accessing a RADIUS Authentication Server's Configuration*

Configuring a New Interface

Each WLAN configured on the WLC needs its own virtual interface. The WLC has five physical ports for data traffic. Each physical port can be configured to support multiple WLANs, each on its own virtual interface. The virtual interface is typically named with a VLAN number and associated to that VLAN. Use the following steps to configure a new interface:

Create a new interface by clicking **CONTROLLER > Interfaces > New**, as

Technet24

shown in Figure 21-7.

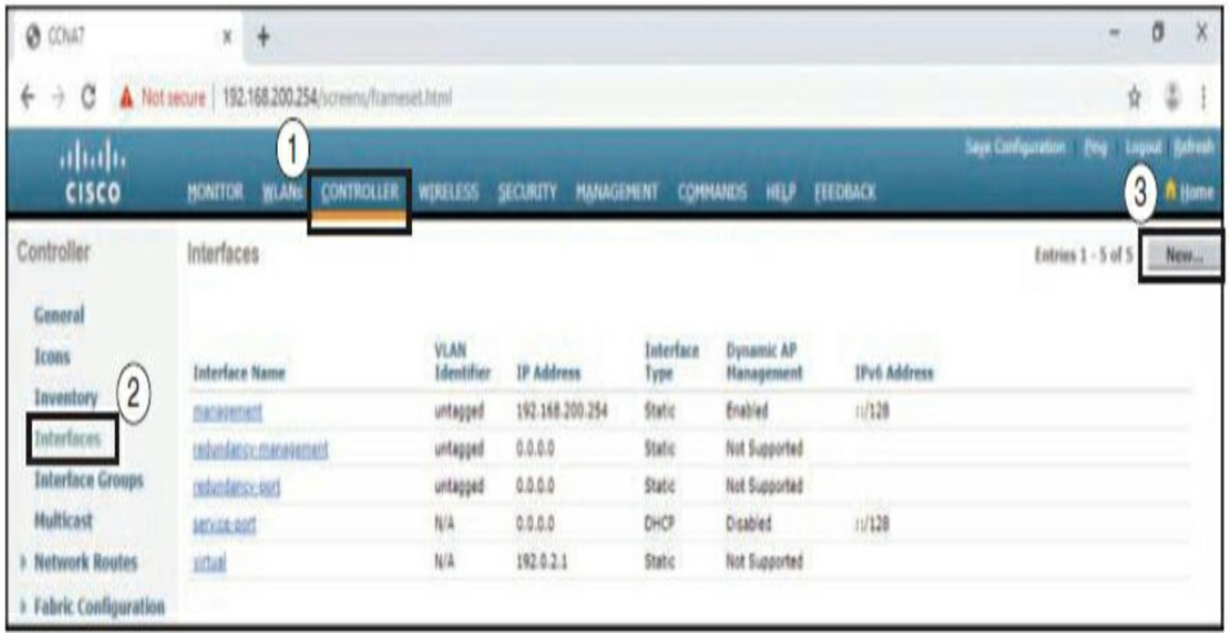

Figure 21-7 *Creating a New Virtual Interface*

Configure an interface name and VLAN ID as shown in Figure 21-8, which shows the interface name being set to **vlan5** and the VLAN ID being set to **5**. Click **Apply** to create the new interface.

Figure 21-8 *Configuring the Interface Name and VLAN ID*

On the Edit page for the interface, configure the physical port number and IP addressing information (see Figure 21-9).

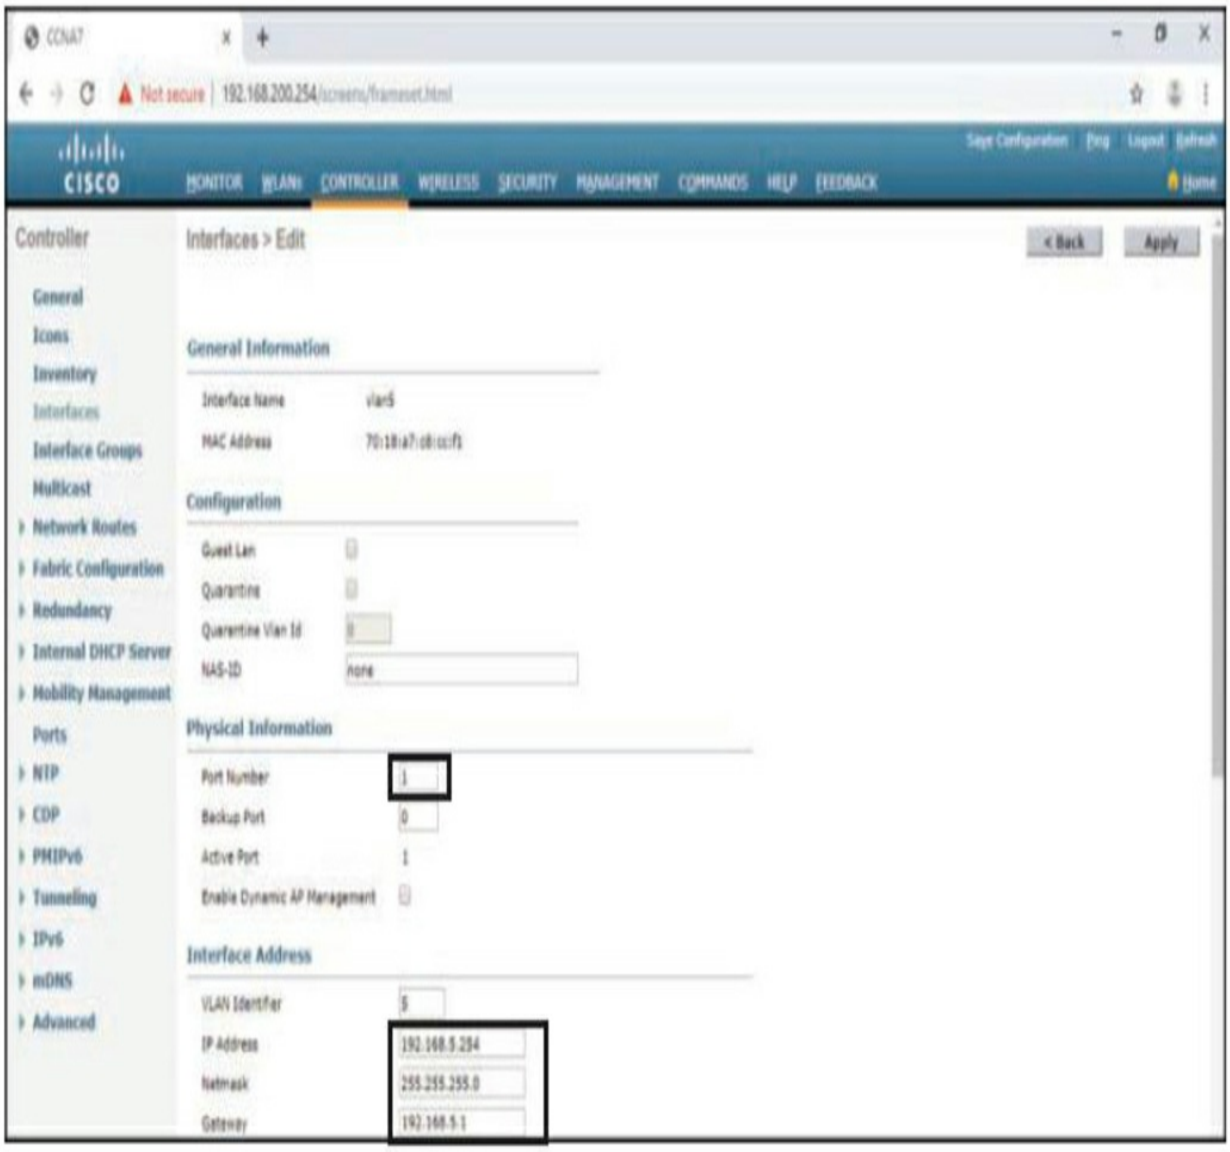

Figure 21-9 *Configuring Port and IP Addressing*

In order to forward DHCP messages to a dedicated DHCP server, configure the DHCP server address as shown in Figure 21-10.

Technet24

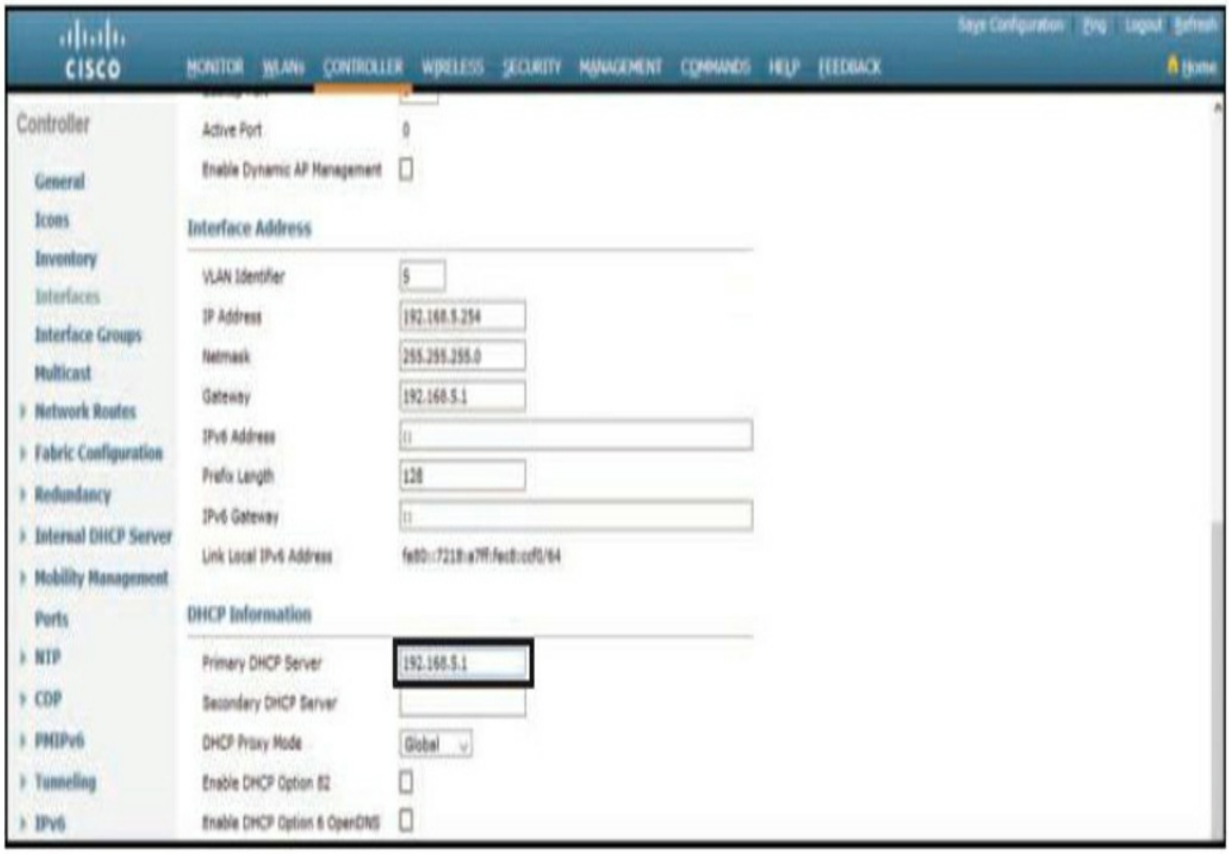

Figure 21-10 *Configuring the DHCP Server Address*

Scroll to the top and click **Apply**, as shown in Figure 21-11. Click **OK** in the warning message.

Figure 21-11 *Applying a New Virtual Interface*

To verify the newly configured virtual interface, click **Interfaces**. The new vlan5 interface is now shown in the list of interfaces with its IPv4 address, as shown in Figure 21-12.

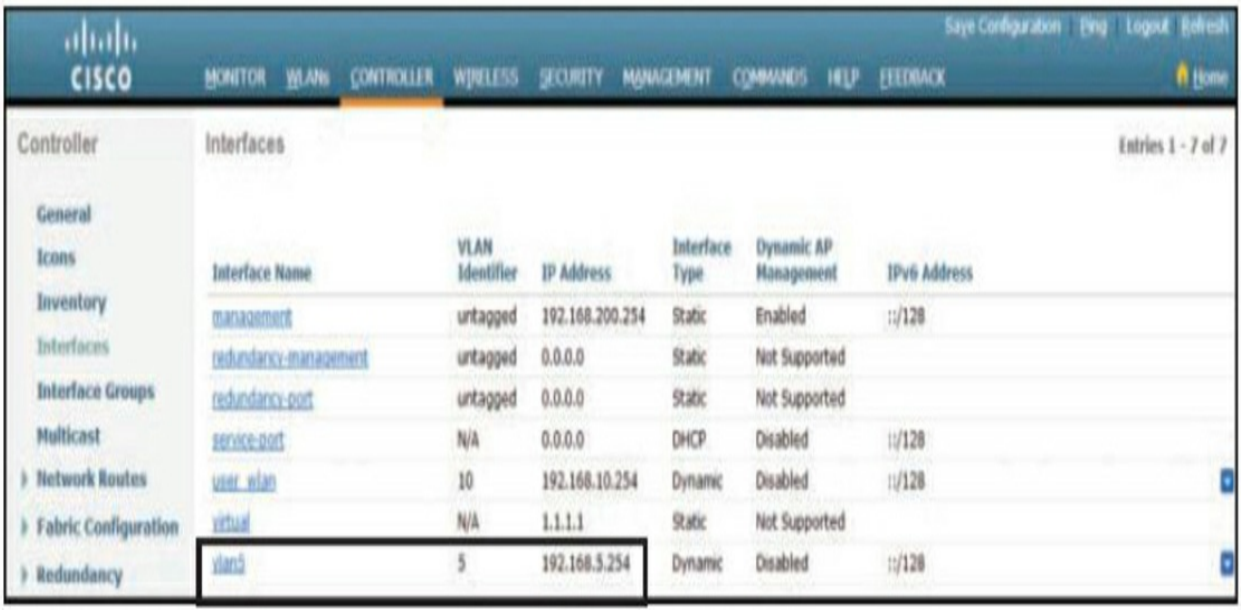

Figure 21-12 *Verifying a New Virtual Interface*

Configuring a WPA2 Enterprise WLAN

By default, all newly created WLANs on the WLC use WPA2 with Advanced Encryption System (AES). 802.1X is the default key management protocol used to communicate with the RADIUS server. The WLC is already configured with the IP address of the RADIUS server.

Configuring a new WLAN for interface **vlan5** on the WLC involves the following steps:

Technet24

To create a new WLAN, click the **WLANs** tab and then **Go**, as shown Figure 21-13.

Figure 21-13 *Creating a New WLAN*

Configure the WLAN name and SSID. In Figure 21-14, the SSID is also used as the profile name and uses the same ID as vlan5, created earlier.

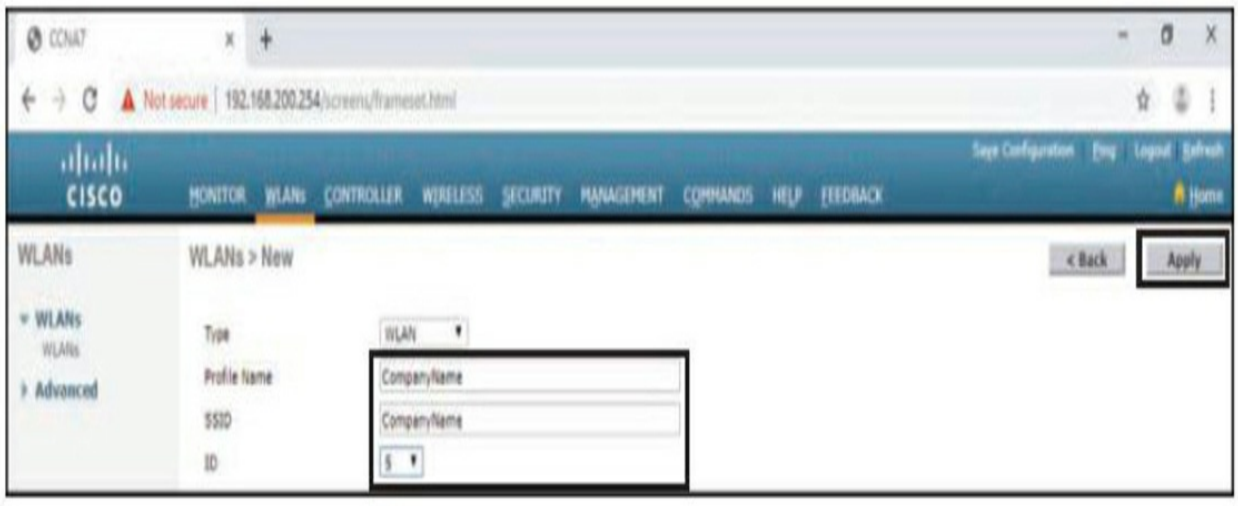

Figure 21-14 *Setting the Profile Name and SSID*

To enable the WLAN for vlan5, change the status to **Enabled** and choose **vlan5** from the Interface/Interface Group(G) dropdown list. Click **Apply** and click **OK** to accept the popup message, as shown Figure 21-15.

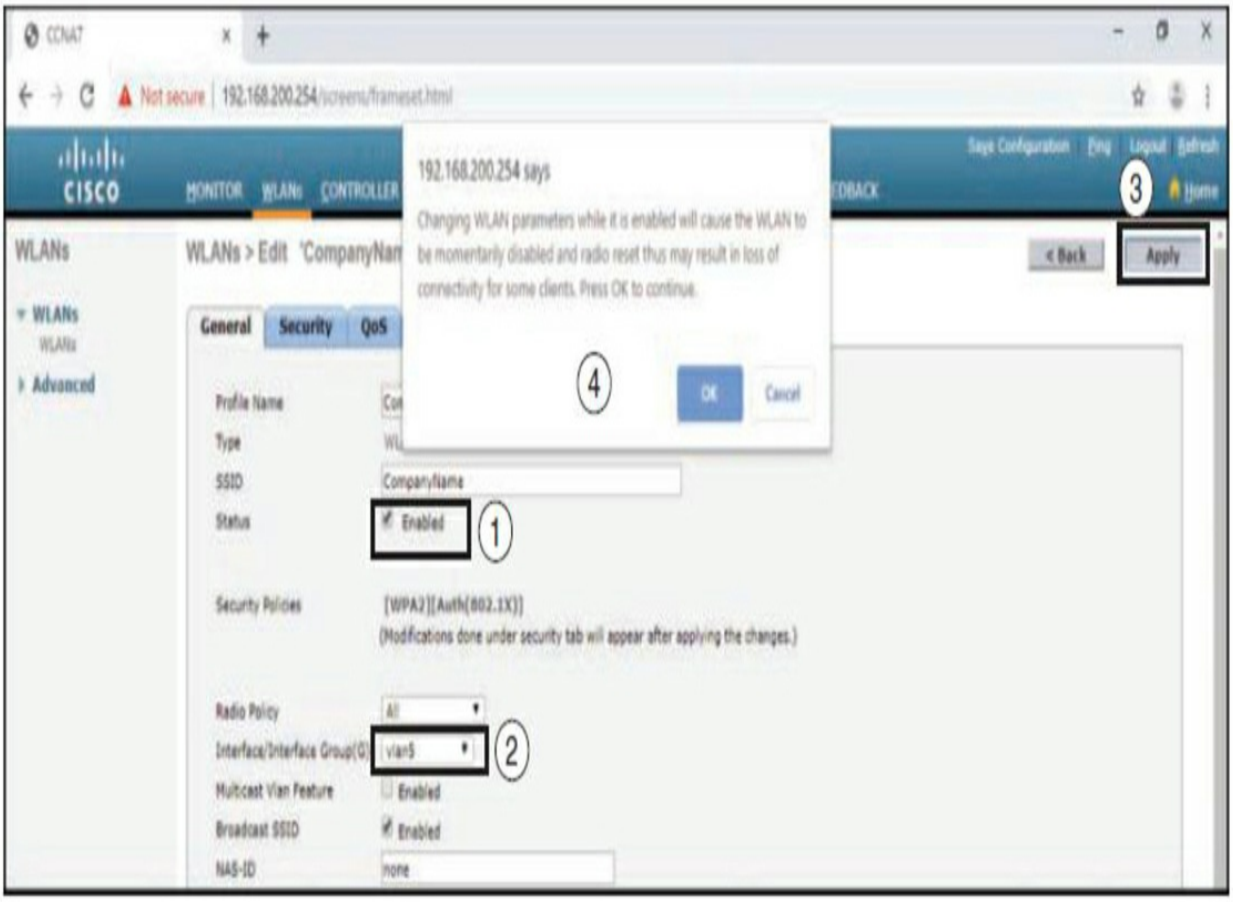

Figure 21-15 *Enabling the WLAN*

To verify AES and the 802.1X defaults, click the **Security** tab to view the default security configuration for the new WLAN, as shown in Figure 21-16. The WLAN should use WPA2 security with AES encryption. Authentication traffic is handled by 802.1X between the WLC and the RADIUS server.

Technet24

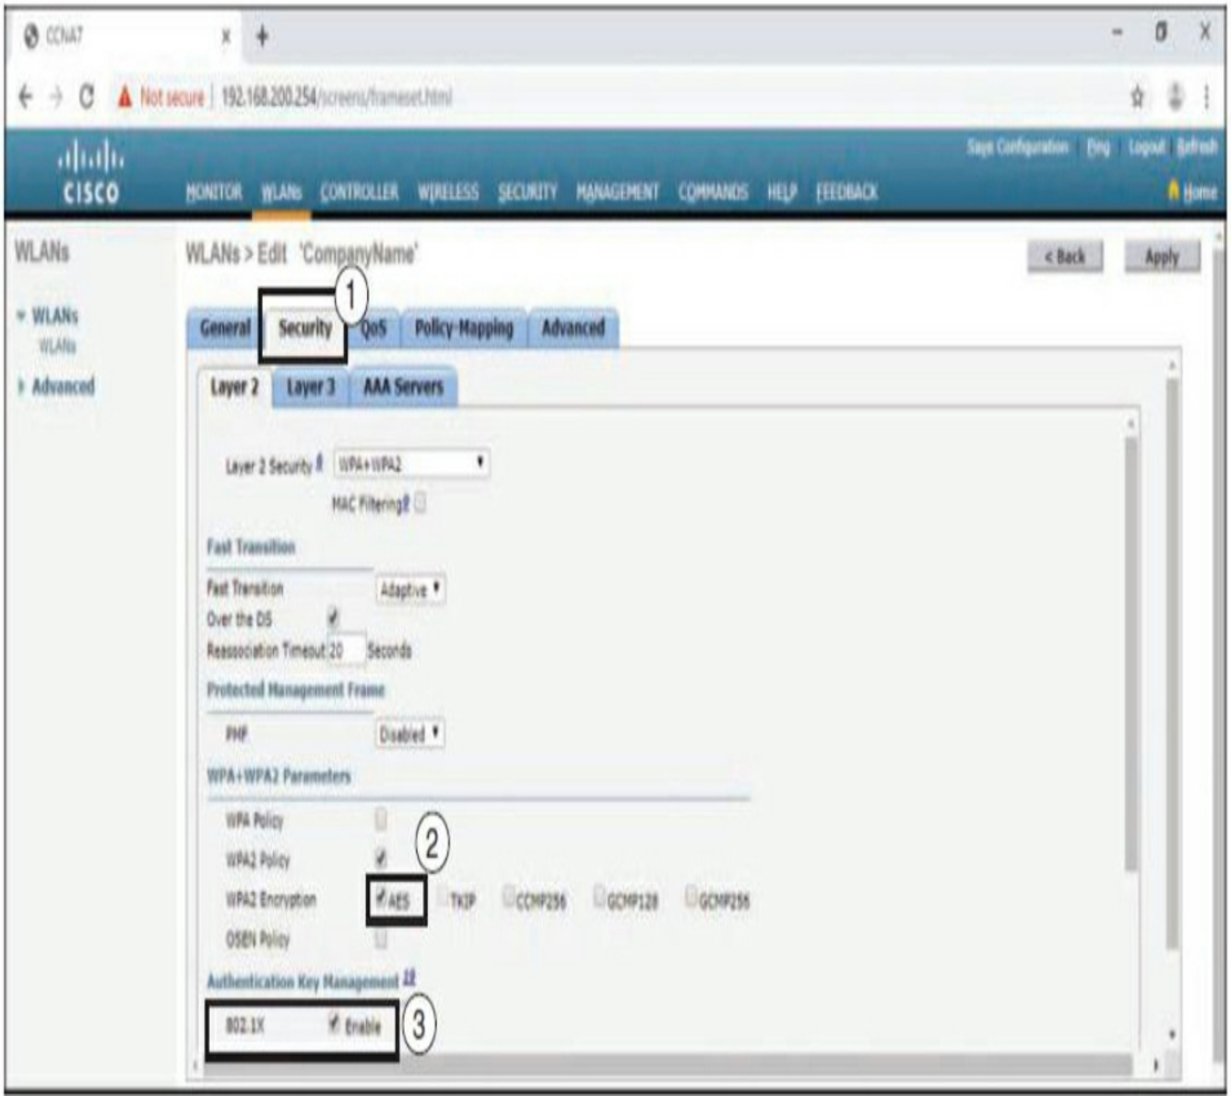

Figure 21-16 *Verifying Security*

To configure WLAN security to use the RADIUS server, click the **AAA Servers** tab, as shown in Figure 21-17. In the dropdown box, select the RADIUS server that was configured on the WLC previously.

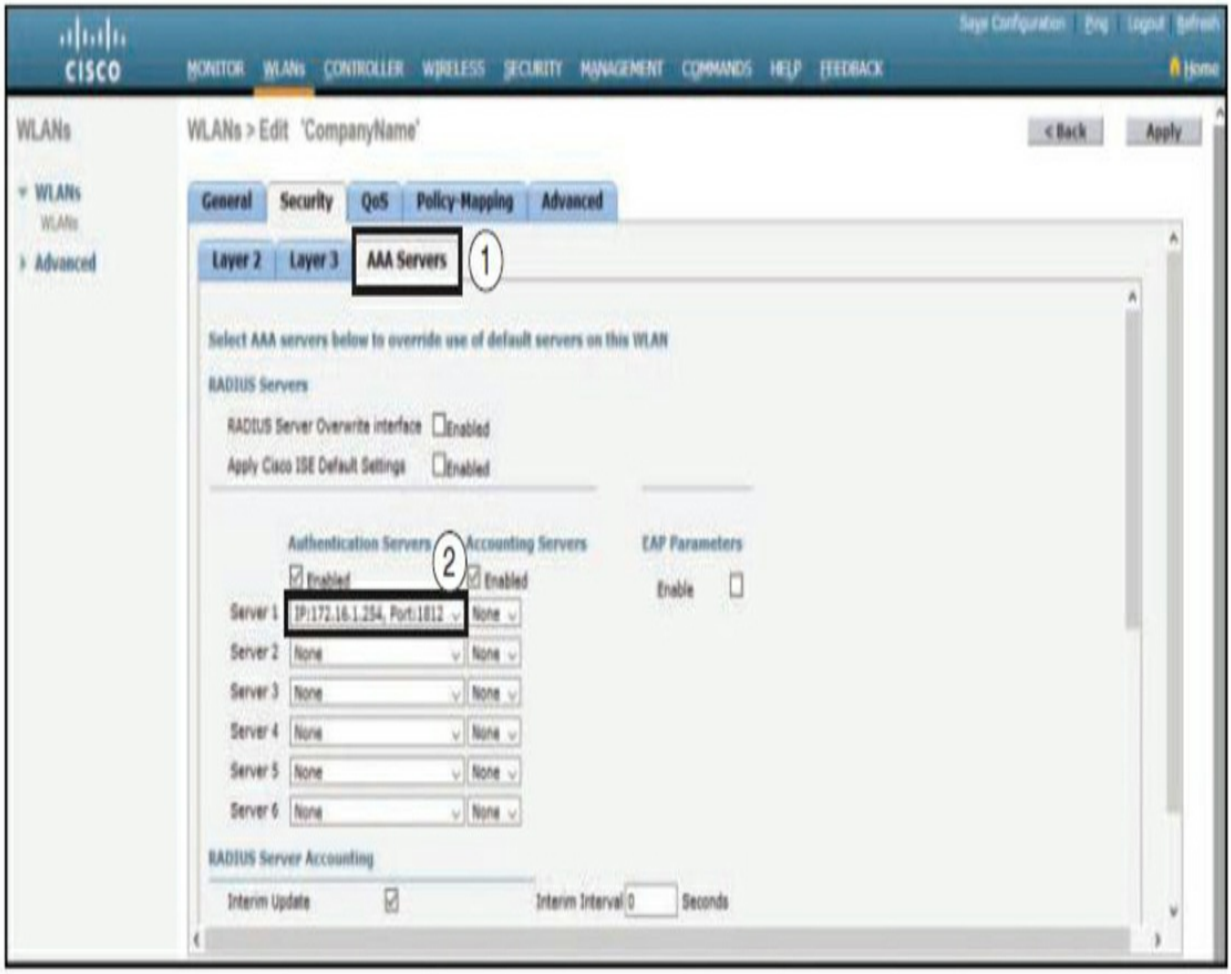

Figure 21-17 *Associating the RADIUS Server to the WLAN*

To configure a QoS profile, click the **QoS** tab, as shown in Figure 21-18. From here, you can configure a QoS profile that adheres to the company policy. Silver (best effort) is currently selected. Click **Apply** to apply the changes.

Technet24

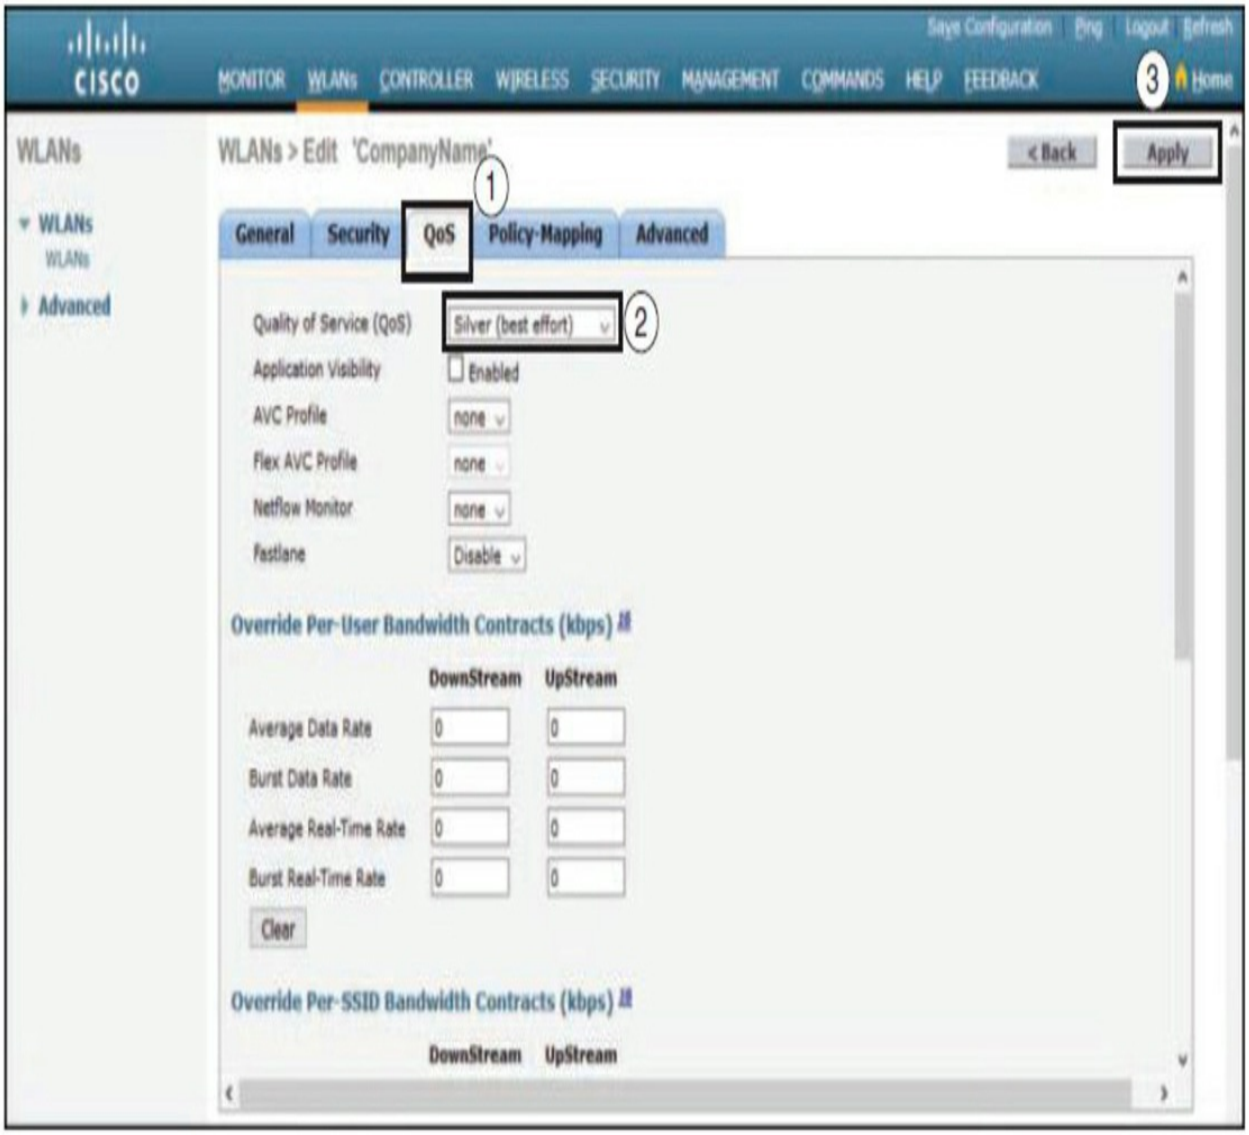

Figure 21-18 *Configuring a QoS Profile*

To verify that the new WLAN is listed and enabled, click the **WLANs** submenu on the left. In Figure 21-19, notice that the WLAN **CompanyName** is enabled and is using WPA2 security with 802.1X authentication.

Figure 21-19 *Verifying the New WLAN*

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	13
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	4
Portable Command Guide	23

Technet24

Day 20. LAN Security and Device Hardening

CCNA 200-301 Exam Topics

• 4.8 Configure network devices for remote access using SSH

• 5.3 Configure and verify device access control using local passwords

• 5.7 Configure and verify Layer 2 security features (DHCP snooping, dynamic ARP inspection, and port security)

• 5.8 Compare authentication, authorization, and accounting concepts

Key Topics

Today’s review is a whirlwind of topics related to LAN security and device hardening. We will review access control, port security, and LAN threat mitigation techniques.

Access Control

Many types of authentication can be performed on networking devices to control access, and each method offers varying levels of security.

Local Authentication

The simplest method of remote access authentication is to configure a login and password combination on console, vty lines, and aux ports, as shown in

Example 20-1. This method, however, provides no accountability, and the password is sent in plaintext. Anyone who has the password can gain entry to the device.

Example 20-1 *Local Password Only Authentication*

```
R1(config)# line vty 0 4
R1(config-line)# password ci5c0
R1(config-line)# login
```

Instead of using a shared password with no usernames, you can use the **username** *username* **secret** *password* command to configure local username/password pairs. Require a username/password pair with the **login local** line configuration command. Use the **no password** line configuration command to remove any configured passwords. In Example 20-2, a username/password pair is configured and applied to the lines, and then Telnet access is tested from a switch. Notice that the password has been hashed using MD5 encryption, indicated by the **5** following **secret** in the output from the **show run** command.

Example 20-2 *Local Username/Password Authentication*

```
R1(config)# username allanj secret 31daysCCNA
R1(config)# line console 0
R1(config-line)# login local
```

Technet24

```
R1(config-line)# no password
R1(config-line)# line vty 0 15
R1(config-line)# login local
R1(config-line)# no password
```

```
S1# telnet 10.10.10.1
Trying 10.10.10.1 ...Open


User Access Verification

Username: allanj
Password:
R1> enable
Password:
R1# show run | include username
username allanj secret 5 $1$mERr$e/edsAr7D0CyM/z3tMvyL
R1#
```

SSH Configuration

Using Secure Shell (SSH) is considered a security best practice because Telnet (port 23) uses insecure plaintext transmission of both the login and the data across the connection. SSH (port 22) is a more secure form of remote access:

• It requires a username and a password, both of which are encrypted during transmissions.

• The username and password can be authenticated using the local database method.

• It provides more accountability because the username is recorded when a user logs in.

Example 20-3 illustrates SSH and local database methods of remote access.

Example 20-3 *Configuring SSH Remote Access on a Switch*

```
S1# show ip ssh
SSH Disabled-version 1.99
%Please create RSA keys to enable SSH (of at least 768
v2.
Authentication timeout: 120 secs; Authentication retri
S1# conf t
S1(config)# ip domain-name cisco.com
S1(config)# crypto key generate rsa
The name for the keys will be: S1.cisco.com
Choose the size of the key modulus in the range of 360
   General Purpose Keys. Choosing a key modulus greate
   a few minutes.

How many bits in the modulus [512]:1024
```

Technet24

```
% Generating 1024 bit RSA keys, keys will be non-expor
[OK] (elapsed time was 4 seconds)
*Mar 1 02:20:18.529: %SSH-5-ENABLED: SSH 1.99 has beer
S1(config)# line vty 0 15
S1(config-line)# login local
S1(config-line)# transport input ssh
S1(config-line)# username allanj secret 31daysCCNA
!The following commands are optional SSH configuration
S1(config)# ip ssh version2
S1(config)# ip ssh authentication-retries 5
S1(config)# ip ssh time-out 60
S1(config)# end
S1# show ip ssh
SSH Enabled - version 2.0
Authentication timeout: 60 secs; Authentication retrie
S1#
```

The following steps occur in Example 20-3:

Verify that the switch supports SSH by using the **show ip ssh** command. If the command is not recognized, you know that SSH is not supported.

Configure a DNS domain name with the **ip domain-name** global configuration command.

Configure the switch by using the **crypto key generate rsa** command to generate an RSA key pair and automatically enable SSH. When generating

RSA keys, you are prompted to enter a modulus length. Cisco recommends a minimum modulus size of 1024 bits, as in Example 20-3.

Note

To remove the RSA key pair, use the **crypto key zeroize rsa** command. This disables the SSH service.

Change the vty lines to use usernames, with either locally configured usernames or an authentication, authorization, and accounting (AAA) server. In Example 20-3, the **login local** vty subcommand defines the use of local usernames, replacing the **login** vty subcommand.

Configure the switch to accept only SSH connections with the **transport input** ssh vty subcommand. (The default is **transport input telnet**.)

Add one or more **username password** global configuration commands to configure username/password pairs.

If desired, modify the default SSH configuration to change the SSH version to 2.0, the number of authentication tries, and the timeout, as in Example 20-3.

Verify your SSH parameters by using the **show ip ssh** command.

Switch Port Hardening

Technet24

Router interfaces must be activated with the **no shutdown** command before they become operational. The opposite is true for Cisco Catalyst switches: An interface is activated when a device is connected to the port. To provide out-of-the-box functionality, Cisco chose a default configuration that includes interfaces that work without any configuration, including automatically negotiating speed and duplex. In addition, all interfaces are assigned to the default VLAN 1.

This default configuration exposes switches to some security threats. The following are security best practices for unused interfaces:

• Administratively disable the interface by using the **shutdown** interface subcommand.

• Prevent VLAN trunking by making the port a nontrunking interface using the **switchport mode** access interface subcommand.

• Assign the port to an unused VLAN by using the **switchport access vlan** *number* interface subcommand.

• Set the native VLAN to not be VLAN 1 but to instead be an unused VLAN, using the **switchport trunk native vlan** *vlan-id* interface subcommand.

Even when you shut down unused ports on the switches, if a device is connected to one of those ports and the interface is enabled, trunking can occur. In addition, all ports are in VLAN 1 by default. A good practice is to

put all unused ports in a black hole VLAN. Example 20-4 demonstrates this best practice, assuming that ports 20–24 are unused.

Example 20-4 *Assigning Unused Ports to a Black Hole VLAN*

```
S1(config)# vlan 999
S1(config-vlan)# name BlackHole
S1(config-vlan)# interface range fa0/20 - 24
S1(config-if-range)# shutdown
S1(config-if-range)# switchport mode access
S1(config-if-range)# switchport access vlan 999
S1(config-if-range)#
```

AAA

Configuring usernames and passwords on all your network devices is not very scalable. A better option is to use an external server to centralize and secure all username/password pairs. To address this issue, Cisco devices support the authentication, authorization, and accounting (AAA) framework to help secure device access.

Cisco devices support two AAA authentication protocols:

• Terminal Access Controller Access Control System Plus (TACACS+, pronounced as “tack-axe plus”)

• Remote Authentication Dial-In User Service (RADIUS)

Technet24

The choice of TACACS+ or RADIUS depends on the needs of the organization. For example, a large ISP might select RADIUS because it supports the detailed accounting required for billing users. An organization with various user groups might select TACACS+ because it requires authorization policies to be applied on a per-user or per-group basis. Table 20-1 compares TACACS+ and RADIUS.

Table 20-1 *Comparison of TACACS+ and RADIUS*

Feature	TACACS+	RADIUS
Most often used for	Network devices	Users
Transport protocol	TCP	UDP
Authentication port number(s)	49	1645, 1812
Protocol encrypts the password	Yes	Yes
Protocol encrypts entire packet	Yes	No
Supports function to authorize each user to a subset of CLI commands	Yes	No
Defined by	Cisco	RFC 2865

Both TACACS+ and RADIUS use a client/server model, where an authenticating device is the client talking to an AAA server. Figure 20-1 shows a simplified view of the process, where a user is attempting to connect to a switch for management purposes.

Technet24

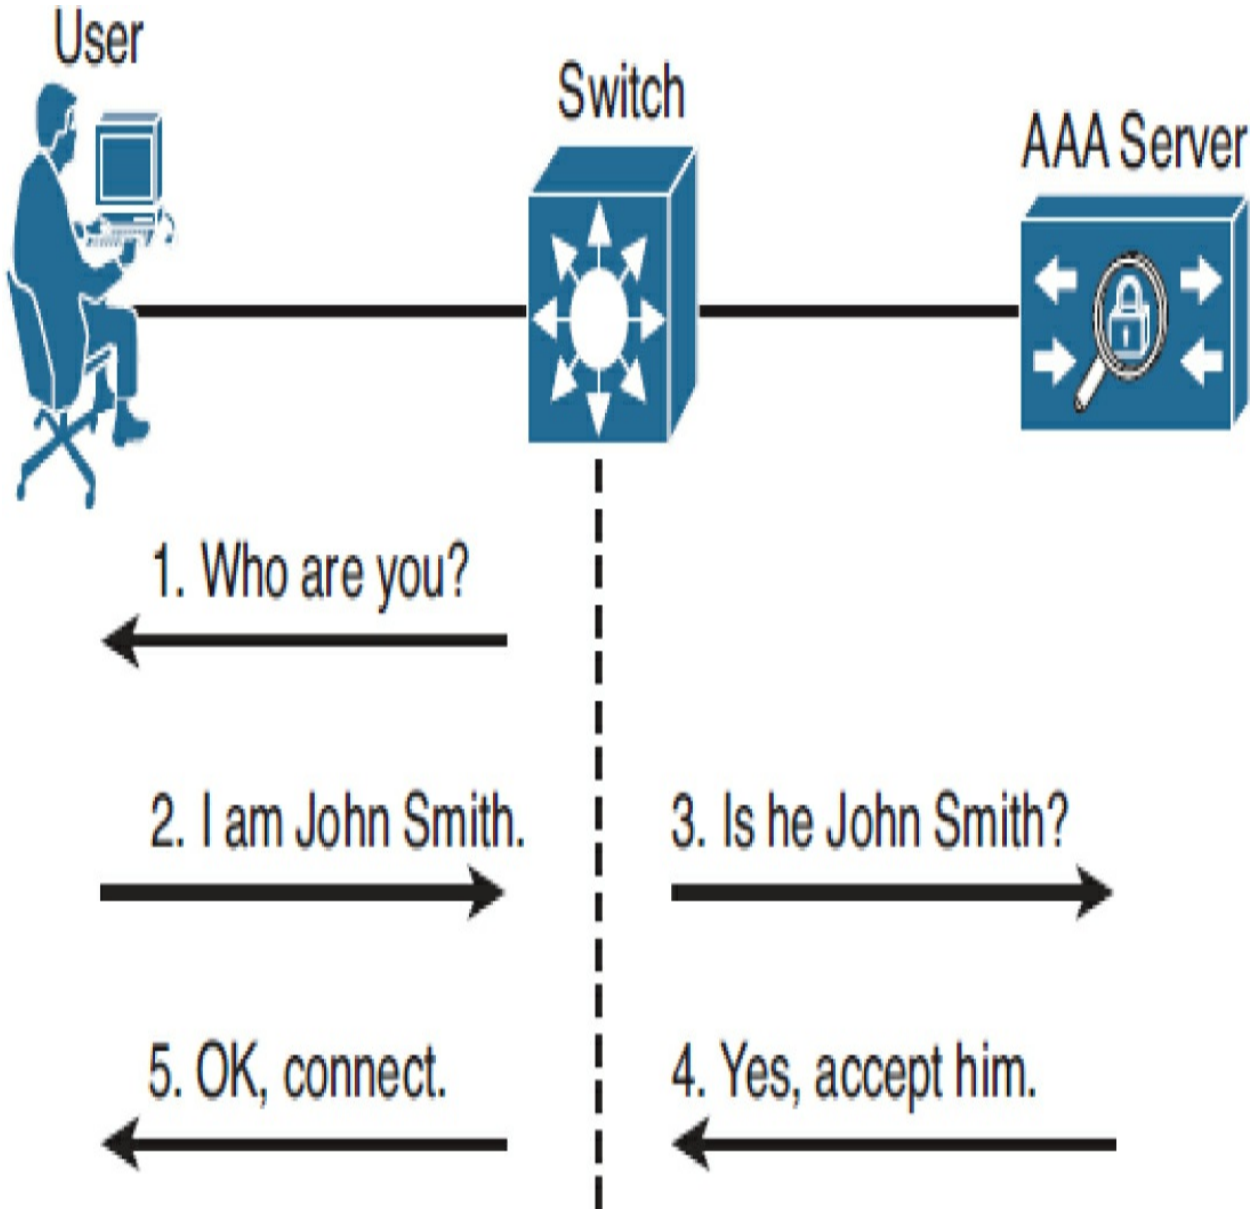

Figure 20-1 *A Simplified View of AAA*

802.1X

IEEE 802.1X is a standard port-based access control and authentication protocol. It is ideal for restricting unauthorized access through publicly available LAN devices, such as switches and wireless access points.

802.1X defines three roles for devices in the network, as Figure 20-2 shows:

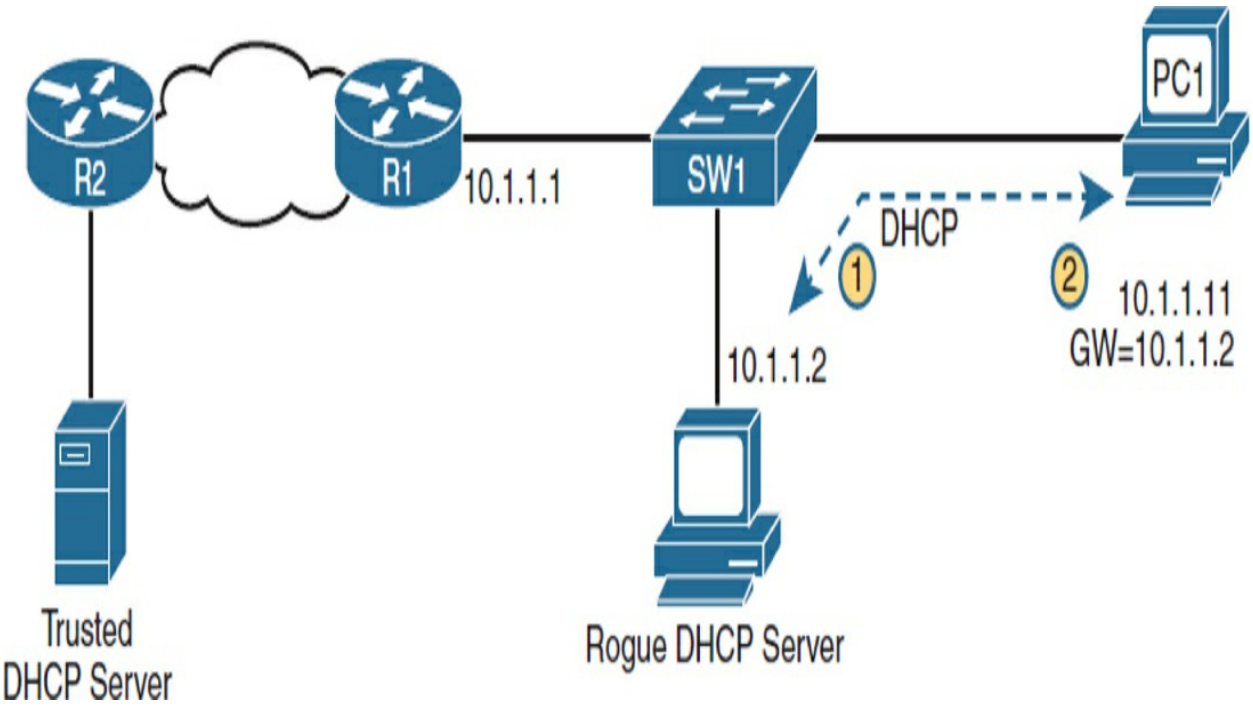

Figure 20-2 *802.1X Roles*

• **Client (supplicant):** This is usually the 802.1X-enabled port on the device that requests access to LAN and switch services and responds to requests from the switch. In Figure 20-2, the device is a PC running 802.1X-compliant client software.

• **Switch (authenticator):** The switch controls physical access to the network, based on the authentication status of the client. The switch acts as a proxy between the client and the authentication server. It requests identifying information from the client, verifies that information with the authentication server, and relays a response to the client.

• **Authentication server:** The authentication server performs the actual

Technet24

authentication of the client. The authentication server validates the identity of the client and notifies the switch about whether the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. RADIUS is the only supported authentication server.

Figure 20-3 shows the authentication flows for a typical 802.1X process.

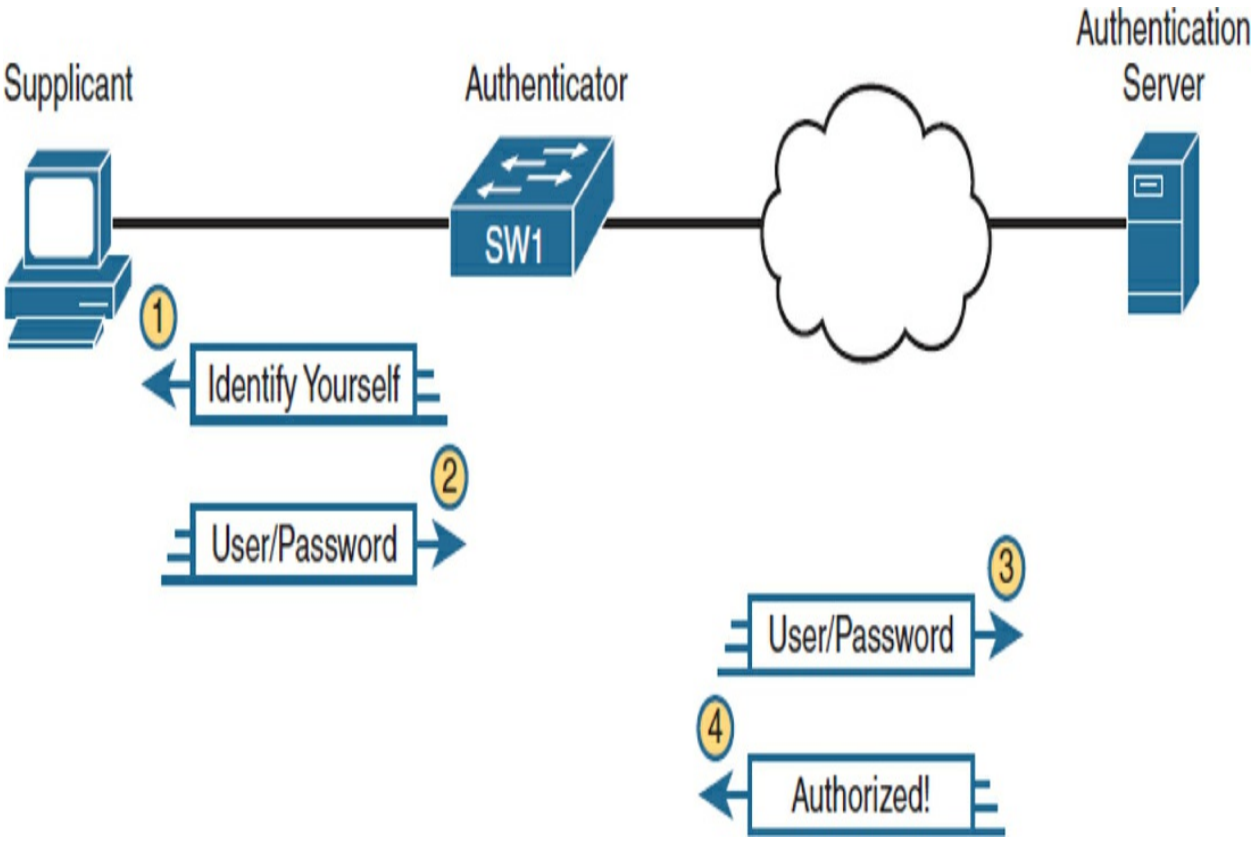

Figure 20-3 *802.1X Authentication Flows*

The 802.1X process is summarized as follows:

• The RADIUS authentication server is configured with usernames and passwords.

• Each LAN switch is enabled as an 802.1X authenticator, is configured with the IP address of the authentication server, and has 802.1X enabled on all required ports.

• A user who connects a device to an 802.1X-enabled port must know the username/password before they can access the network.

Port Security

If you know which devices should be cabled and connected to a particular interface on a switch, you can use port security to restrict that interface so that only the expected devices can use it. This reduces exposure to some types of attacks in which the attacker connects a laptop to the wall socket or uses the cable attached to another end device to gain access to the network.

Port Security Configuration

Port security configuration involves several steps. Basically, you need to make the port an access port, which means the port is not doing any VLAN trunking. You then need to enable port security and configure the Media Access Control (MAC) addresses of the devices that are allowed to use that port. The following list outlines the steps in port security configuration, including the configuration commands used:

Configure the interface for static access mode by using the **switchport mode access** interface subcommand.

Technet24

Enable port security by using the **switchport port-security** interface subcommand.

(Optional) Override the maximum number of allowed MAC addresses associated with the interface (one) by using the **switchport port-security maximum** *number* interface subcommand.

(Optional) Override the default action when there is a security violation (shutdown) by using the **switchport port-security violation {protect | restrict | shutdown}** interface subcommand.

(Optional) Predefine any allowed source MAC address(es) for this interface by using the **switchport port-security mac-address** *mac-address* command. Use the command multiple times to define more than one MAC address.

(Optional) Instead of taking step 5, configure the interface to dynamically learn and configure the MAC addresses of currently connected hosts by configuring the **switchport port-security mac-address sticky** interface subcommand.

When an unauthorized device attempts to send frames to the switch interface, the switch can issue informational messages, discard frames from that device, or even discard frames from all devices by effectively shutting down the interface. Exactly which action the switch port takes depends on the option you configure in the **switchport port-security violation** command. Table 20-2 lists actions that the switch will take based on whether you configure the

option **protect**, **restrict**, or **shutdown** (default).

Table 20-2 *Actions When a Port Security Violation Occurs*

Option on the switchport port-security violation Command	protect	restrict	shutdown
Discards offending traffic	Yes	Yes	Yes
Sends log and SNMP messages	No	Yes	Yes
Disables the interface, discarding all traffic	No	No	Yes

Example 20-5 shows a port security configuration in which each access interface is allowed a maximum of three MAC addresses. If a fourth MAC address is detected, only the offending device's traffic is discarded. If the violation option is not explicitly configured, the traffic for devices that are allowed on the port also is discarded because the port would be shut down by default.

Example 20-5 *Port Security Configuration Example*

```
S1(config)# interface range fa 0/5 - fa 0/24
S1(config-if-range)# switchport mode access
S1(config-if-range)# switchport port-security
S1(config-if-range)# switchport port-security maximum
```

Technet24

```
S1(config-if-range)# switchport port-security violatio
S1(config-if-range)# switchport port-security mac-addr
```

To verify port security configuration, use the more general **show port-security** command or the more specific **show port-security interface** *type number* command. Example 20-6 demonstrates the use of both commands. In the examples, notice that only one device is currently attached to an access port on S1.

Example 20-6 *Port Security Verification Command Output Examples*

```
S1# show port-security
Secure Port MaxSecureAddr CurrentAddr SecurityViolatio
            (Count)         (Count)      (Count)
Fa0/5        3               1             0
Fa0/6        3               0             0
Fa0/7        3               0             0
Fa0/8        3               0             0
Fa0/9        3               0             0
Fa0/10       3               0             0
Fa0/11       3               0             0
Fa0/12       3               0             0
Fa0/13       3               0             0
Fa0/14       3               0             0
Fa0/15       3               0             0
Fa0/16       3               0             0
```

```
Fa0/17       3                0              0
Fa0/18       3                0              0
Fa0/19       3                0              0
Fa0/20       3                0              0
Fa0/21       3                0              0
Fa0/22       3                0              0
Fa0/23       3                0              0
Fa0/24       3                0              0
Total Addresses in System (excluding one mac per port)
Max Addresses limit in System (excluding one mac per p
S1# show port-security interface fastethernet 0/5
Port Security              :Enabled
Port Status                :Secure-down
Violation Mode             :Restrict
Aging Time                 :0 mins
Aging Type                 :Absolute
SecureStatic Address Aging  :Disabled
Maximum MAC Addresses      :3
Total MAC Addresses        :1
Configured MAC Addresses   :0
Sticky MAC Addresses       :1
Last Source Address:Vlan   :0014.22dd.37a3:1
Security Violation Count   :0
```

Port Security Aging

Port security aging can be used to set the aging time for static and dynamic

Technet24

secure addresses on a port. Two types of aging are supported per port:

• **Absolute:** The secure addresses on the port are deleted after the specified aging time.

• **Inactivity:** The secure addresses on the port are deleted only if they are inactive for the specified aging time.

Use the **switchport port-security aging** command to enable or disable static aging for the secure port or to set the aging time or type:

```
Switch(config-if)# switchport port-security aging { st
  {absolute | inactivity}}
```

Table 20-3 describes the parameters for this command.

Table 20-3 *Parameters for the* port-security aging *Command*

Parameter	Description
`static`	Enable aging for statically configured secure addresses on this port.
`time` *time*	Specify the aging time for this port. The range is 0 to 1440 minutes. If the time is 0, aging is disabled for this port.
`type absolute`	Set the absolute aging time. All the secure addresses on this port age out exactly after the time (in minutes) specified and are removed from the secure address list.
`type inactivity`	Set the inactivity aging type. The secure addresses on this port age out only if there is no data traffic from the secure source address for the specified time period.

Example 20-7 shows an administrator configuring the aging type to 10 minutes of inactivity and using the **show port-security interface** command to verify the configuration.

Example 20-7 *Configuring and Verifying Port Security Aging*

```
S1(config)# interface fa0/1
S1(config-if)# switchport port-security aging time 10
```

Technet24

```
S1(config-if)# switchport port-security aging type ina
S1(config-if)# end
S1# show port-security interface fa0/1
Port Security              : Enabled
Port Status                : Secure-shutdown
Violation Mode             : Restrict
Aging Time                 : 10 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 4
Total MAC Addresses        : 1
Configured MAC Addresses   : 1
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 0050.56be.e4dd:1
Security Violation Count   : 1
```

Port Restoration After a Violation

When port security is activated on an interface, the default action when a violation occurs is to shut down the port. A security violation can occur in one of two ways:

• The maximum number of secure MAC addresses has been added to the address table for that interface, and a station whose MAC address is not in the address table attempts to access the interface.

• An address learned or configured on one secure interface is seen on another

secure interface in the same VLAN.

When a violation occurs, a syslog message is sent to the console, stating that the interface is now in the **err-disable** state. The console messages include the port number and the MAC address that caused the violation, as Example 20-8 shows.

Example 20-8 *Port Security Violation Verification and Restoration*

```
S1#
Sep 20 06:44:54.966: %PM-4-ERR_DISABLE: psecure-violat
 Fa0/18,
   putting Fa0/18 in err-disable state
Sep 20 06:44:54.966: %PORT_SECURITY-2-PSECURE_VIOLATIO
   occurred, caused by MAC address 000c.292b.4c75 on p
Sep 20 06:44:55.973: %LINEPROTO-5-PPDOWN: Line protocc
FastEthernet0/18, changed state to down
Sep 20 06:44:56.971: %LINK-3-UPDOWN: Interface FastEth
   to down
!The two following commands can be used to verify the
S1# show interface fa0/18 status
Port     Name   Status         Vlan  Duplex  Speed  Ty
Fa0/18          err-disabled   5     auto    auto   10
S1# show port-security interface fastethernet 0/18
Port Security                       : Enabled
Port Status                         : Secure-shutdown
Violation Mode                      : Shutdown
```

Technet24

```
Aging Time                              : 0 mins
Aging Type                              : Absolute
SecureStatic Address Aging              : Disabled
Maximum MAC Addresses                   : 1
Total MAC Addresses                     : 0
Configured MAC Addresses                : 0
Sticky MAC Addresses                    : 0
Last Source Address:Vlan                : 000c.292b.4c75:1
Security Violation Count                : 1
!To restore a port, manually shut it down and then rea
S1(config)# interface FastEthernet 0/18
S1(config-if)# shutdown
Sep 20 06:57:28.532: %LINK-5-CHANGED: Interface FastEt
to
   administratively down
S1(config-if)# no shutdown
Sep 20 06:57:48.186: %LINK-3-UPDOWN: Interface FastEth
to up
Sep 20 06:57:49.193: %LINEPROTO-5-UPDOWN: Line protoco
FastEthernet0/18, changed state to up
```

You can use the **show interface** *type number* **status** or **show port-security interface** *type number* command to verify the current state of the port. To restore the port, you must first manually shut down the interface and then reactivate it, as in Example 20-8.

LAN Threat Mitigation

This section reviews LAN threats and mitigation techniques for VLAN attacks, DHCP attacks, and ARP attacks.

Native and Management VLAN Modification

The IEEE 802.1Q specification defines a native VLAN to maintain backward compatibility with untagged traffic that is common in legacy LAN scenarios. A native VLAN serves as a common identifier on opposite ends of a trunk link. VLAN 1 is the native VLAN by default.

A management VLAN is any VLAN configured to access the management capabilities of a switch. VLAN 1 is the management VLAN by default. The management VLAN is assigned an IP address and subnet mask, allowing the switch to be managed through HTTP, Telnet, SSH, or SNMP.

It is a best practice to configure the native VLAN as an unused VLAN distinct from VLAN 1 and other VLANs. In fact, it is not unusual to dedicate a fixed VLAN to serve the role of the native VLAN for all trunk ports in the switched domain. Likewise, the management VLAN should be configured as something other than VLAN 1. The management and native VLANs can be configured as the same VLAN, as in Example 20-9.

Example 20-9 *Configuring the Native and Management VLAN*

```
S1(config)# vlan 86
S1(config-vlan)# name Management&Native
```

Technet24

```
S1(config-vlan)# interface vlan 86
*Jul 13 14:14:04.840: %LINEPROTO-5-UPDOWN: Line protoc
S1(config-if)# ip address 10.10.86.10 255.255.255.0
S1(config-if)# no shutdown
S1(config-if)# ip default-gateway 10.10.86.254
S1(config)# interface range fa0/21 - 24
S1(config-if-range)# switchport mode trunk
S1(config-if-range)# switchport trunk native vlan 86
S1(config-if-range)#
*Jul 13 14:15:55.499: %LINEPROTO-5-UPDOWN: Line protoc
 changed state to up
S1(config-if-range)#
```

First, a VLAN is created that will be used for the management and native VLAN. Next, by activating interface VLAN 86, the switch can be remotely managed. Finally, the trunk ports are statically configured, and VLAN 86 is set as the native VLAN for all untagged traffic. After it is configured, the interface VLAN 86 comes up.

VLAN Attacks

VLAN attacks can be launched in one of three ways:

• **Spoofing Dynamic Trunking Protocol (DTP) messages:** Spoofing DTP messages from the attacking host can cause the switch to enter trunking mode. From here, the attacker can send traffic tagged with the target VLAN,

and the switch then delivers the packets to the destination.

• **Introducing a rogue switch and enabling trunking:** After doing this, an attacker can access all the VLANs on the victim switch from the rogue switch.

• **Mounting a double-tagging (or double-encapsulated) attack:** This type of VLAN hopping attack takes advantage of the way hardware on most switches operates. In specific situations, a threat actor could embed a hidden 802.1Q tag inside the frame that already has an 802.1Q tag. This tag allows the frame to go to a VLAN that the original 802.1Q tag did not specify.

VLAN Attack Mitigation

Use the following steps to mitigate VLAN hopping attacks:

Disable DTP (auto trunking) negotiations on non-trunking ports by using the **switchport mode access** interface configuration command.

Disable unused ports and put them in an unused VLAN.

Manually enable the trunk link on a trunking port by using the **switchport mode trunk** command.

Disable DTP (auto trunking) negotiations on trunking ports by using the **switchport nonegotiate** command.

Technet24

Set the native VLAN to a VLAN other than VLAN 1 by using the **switchport trunk native vlan** *vlan_number* command.

For example, assume the following:

• FastEthernet ports 0/1 through fa0/16 are active access ports.

• FastEthernet ports 0/17 through 0/24 are not currently in use.

• FastEthernet ports 0/21 through 0/24 are trunk ports.

VLAN hopping can be mitigated by implementing the following configuration, as shown in Example 20-10:

• Trunking is disabled on FastEthernet ports 0/1 to 0/16.

• FastEthernet ports 0/17 to 0/20 are assigned an unused VLAN.

• FastEthernet ports 0/21 to 0/24 are manually enabled as trunks with DTP disabled. The native VLAN is also changed from the default VLAN 1 to VLAN 86.

Example 20-10 *VLAN Hopping Attack Mitigation*

```
S1(config)# interface range fa0/1 - 16
S1(config-if-range)# switchport mode access
S1(config-if-range)# exit
```

```
S1(config)#
S1(config)# interface range fa0/17 - 20
S1(config-if-range)# switchport mode access
S1(config-if-range)# switchport access vlan 999
S1(config-if-range)# exit
S1(config)#
S1(config)# interface range fa0/21 - 24
S1(config-if-range)# switchport mode trunk
S1(config-if-range)# switchport nonegotiate
S1(config-if-range)# switchport trunk native vlan 86
S1(config-if-range)# end
S1#
```

DHCP Attacks

Two types of DHCP attacks are DHCP starvation and DHCP spoofing. Both attacks are mitigated by implementing DHCP snooping.

DHCP Starvation Attacks

The goal of a DHCP starvation attack is to create a denial of service condition for connecting clients. DHCP starvation attacks require an attack tool such as Gobbler. Gobbler looks at the entire scope of leasable IP addresses and tries to lease them all. Specifically, it creates DHCP discovery messages with bogus MAC addresses.

DHCP Spoofing Attacks

Technet24

A DHCP spoofing attack occurs when a rogue DHCP server is connected to the network and provides false IP configuration parameters to legitimate clients. For example, in Figure 20-4, R1 is configured to relay DHCP requests to the DHCP server attached to R2.

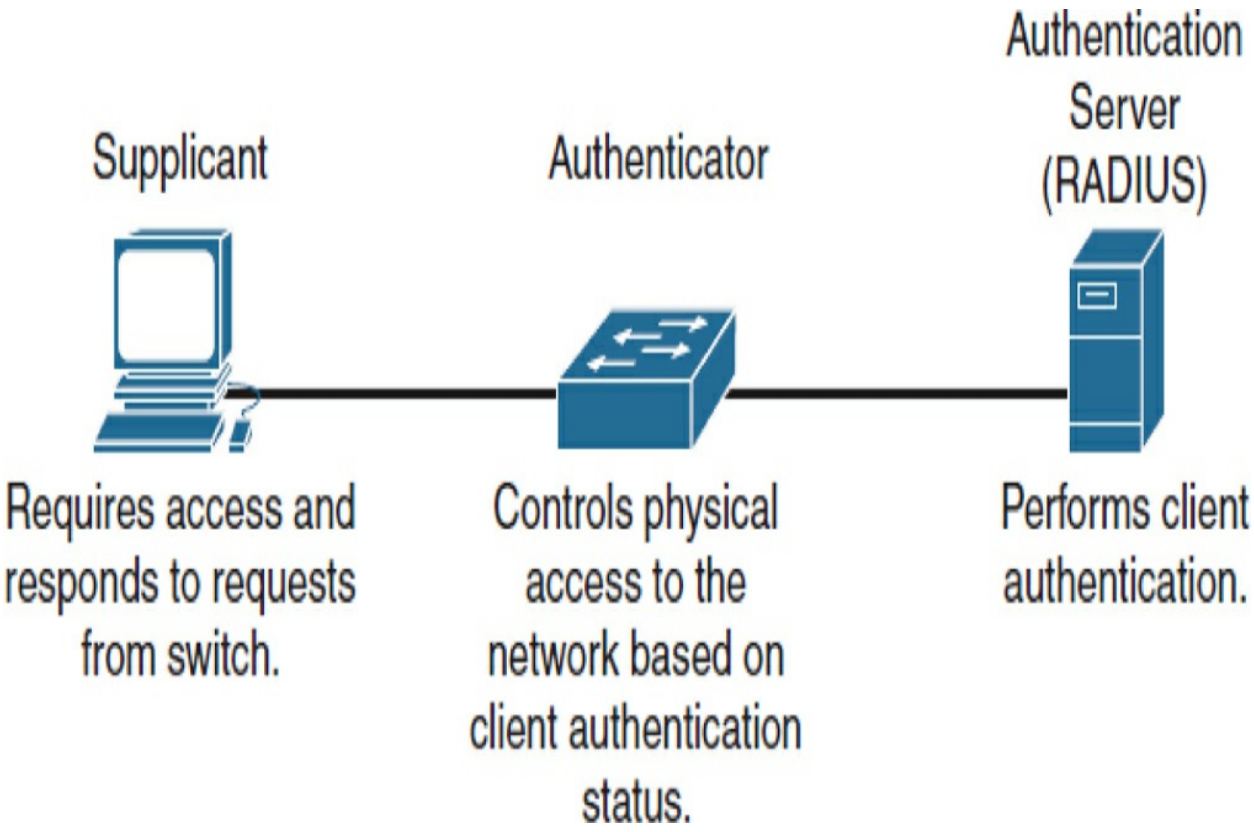

Figure 20-4 *Rogue DHCP Server Intercepting DHCP Requests*

However, the rogue DHCP server attached to SW1 responds to the DHCP request from PC1 first. PC1 accepts the DHCP offer and sets the rogue DHCP server as the default gateway.

DHCP Snooping

To protect against DHCP attacks, DHCP snooping uses the concept of trusted

and untrusted ports. As Figure 20-5 shows, SW2, R1, and the DHCP server are attached to trusted ports on SW1. The other devices, including the wireless access point, are connected to untrusted ports.

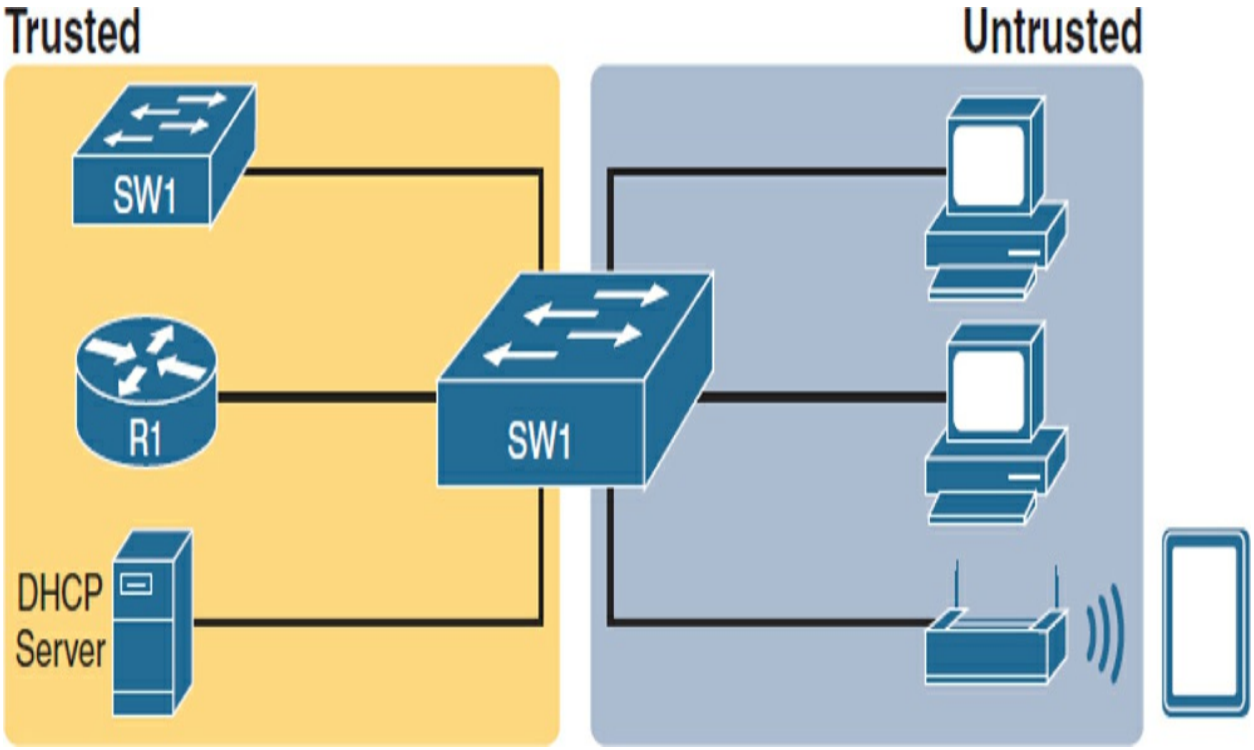

Figure 20-5 *Trusted and Untrusted Ports*

Some critical features of a DHCP snooping configuration include the following:

• **Trusted ports:** Trusted ports allow all incoming DHCP messages.

• **Untrusted ports, server messages:** Untrusted ports discard all incoming messages that are considered server messages.

• **Untrusted ports, client messages:** Untrusted ports apply more complex

Technet24

logic for messages considered client messages. They check whether each incoming DHCP message conflicts with existing DHCP binding table information; if so, they discard the DHCP message. If the message has no conflicts, the switch allows the message through, which typically results in the addition of new DHCP binding table entries.

• **Rate limiting:** This feature optionally limits the number of received DHCP messages per second per port.

Use the following steps to enable DHCP snooping:

Enable DHCP snooping by using the **ip dhcp snooping** global configuration command.

On trusted ports, use the **ip dhcp snooping trust** interface configuration command.

Limit the number of DHCP discovery messages that can be received per second on untrusted ports by using the **ip dhcp snooping limit rate** *number* interface configuration command. This helps mitigate DHCP starvation attacks.

Enable DHCP snooping by VLAN or by a range of VLANs by using the **ip dhcp snooping** *vlan* global configuration command.

For a simple scenario, consider the topology in Figure 20-6.

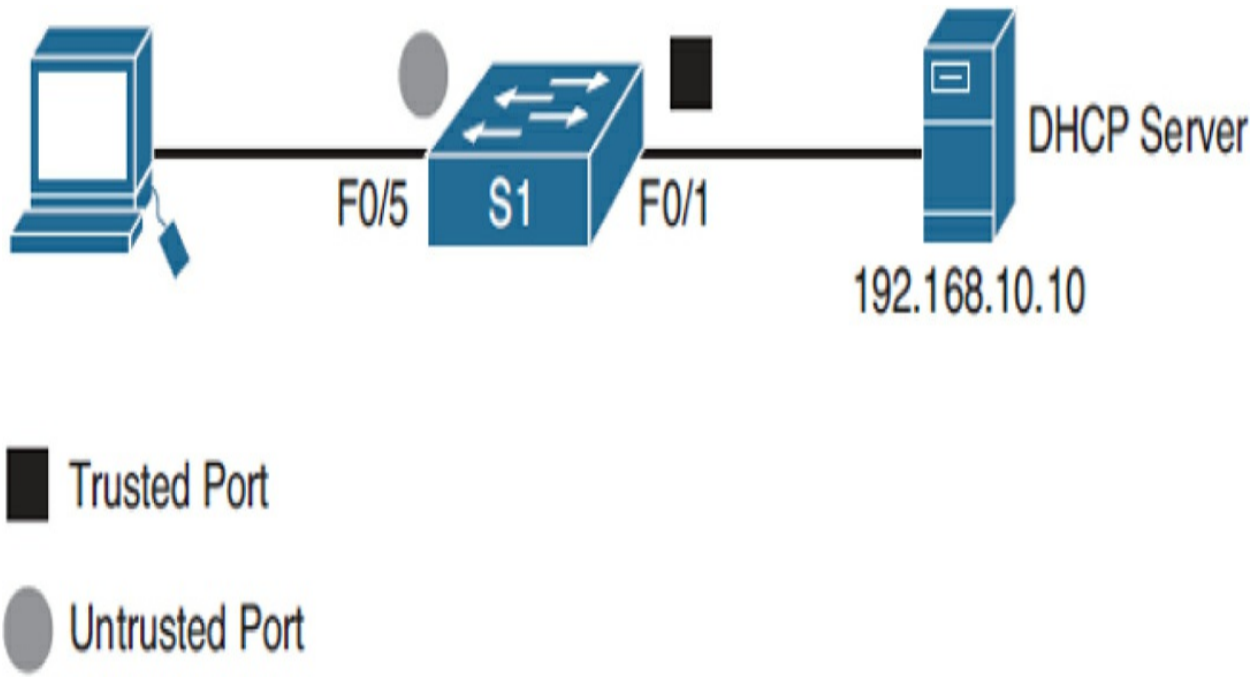

Figure 20-6 *DHCP Snooping Configuration Topology*

Example 20-11 shows how to configure and verify DHCP snooping on S1.

Example 20-11 *Configuring and Verifying DHCP Snooping*

```
S1(config)# ip dhcp snooping
S1(config)# interface f0/1
S1(config-if)# ip dhcp snooping trust
S1(config-if)# exit
S1(config)# interface range f0/5 - 24
S1(config-if-range)# ip dhcp snooping limit rate 6
S1(config-if)# exit
S1(config)# ip dhcp snooping vlan 5,10,50-52
S1(config)# end
S1# show ip dhcp snooping
Switch DHCP snooping is enabled
```

Technet24

```
DHCP snooping is configured on following VLANs:
5,10,50-52
DHCP snooping is operational on following VLANs:
none
DHCP snooping is configured on the following L3 Interf
Insertion of option 82 is enabled
  circuit-id default format: vlan-mod-port
  remote-id: 0cd9.96d2.3f80 (MAC)
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Verification of giaddr field is enabled
DHCP snooping trust/rate is configured on the followir
Interface               Trusted    Allow option   Rate
----------------------- --------   -------------  ----
FastEthernet0/1          yes        yes           unli
  Custom circuit-ids:
FastEthernet0/5          no         no            6
  Custom circuit-ids:
FastEthernet0/6          no         no            6
  Custom circuit-ids:
S1# show ip dhcp snooping binding
MacAddress         IpAddress      Lease(sec)  Type
---------------    -------------  ---------  ---------
00:03:47:B5:9F:AD  192.168.10.10  193185     dhcp-snoo
S1#
```

ARP Attacks

On Ethernet LANs, hosts are allowed to send an unsolicited Address Resolution Protocol (ARP) reply called a *gratuitous ARP* message. These ARP messages cause all other hosts on the LAN to store the MAC address and IP address in their ARP caches. The problem is that an attacker can send a gratuitous ARP message containing a spoofed MAC address to a switch, and the switch would update its MAC table accordingly. Therefore, any host can claim to be the owner of any IP and MAC address combination.

For example, in Figure 20-7, R1 and PC1 have removed the correct entry for each other's MAC address and replaced it with PC2's MAC address. The threat actor has poisoned the ARP caches of all devices on the subnet. ARP poisoning leads to various man-in-the-middle attacks, posing a serious security threat to the network. All traffic between R1 and PC1 will now flow through the threat actor's PC2.

Technet24

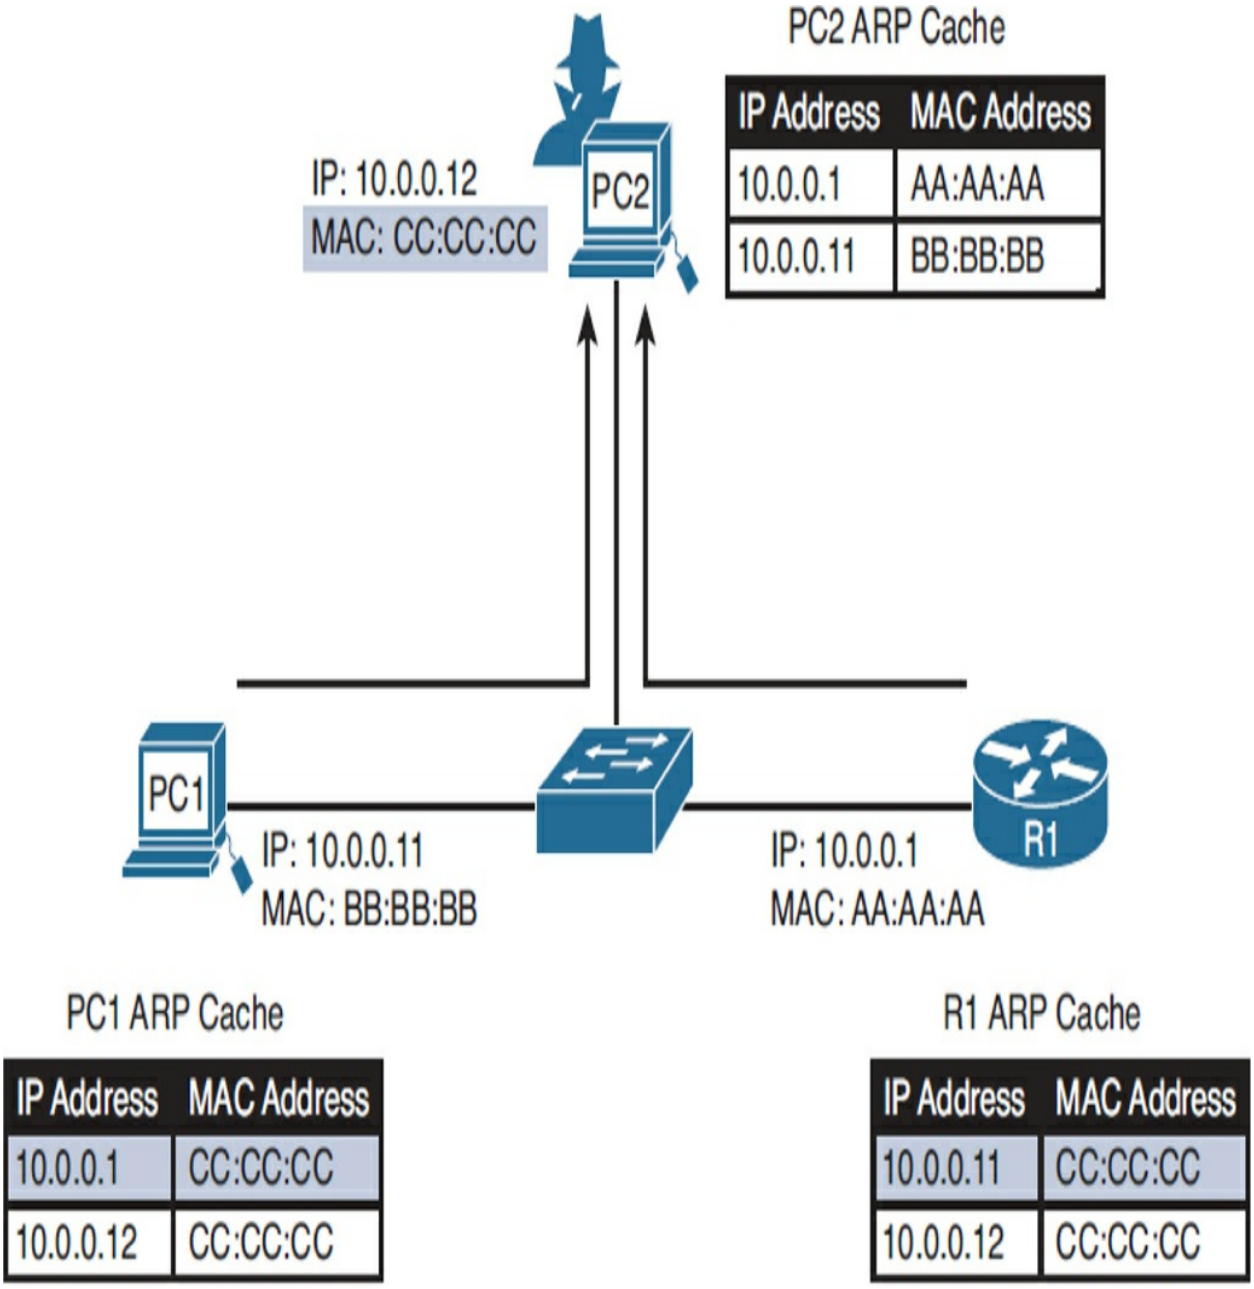

Figure 20-7 *Successful ARP Poisoning Attack*

Dynamic ARP Inspection

To prevent ARP spoofing and then ARP poisoning, a switch must ensure that only valid ARP requests and replies are relayed. Dynamic ARP inspection (DAI) requires DHCP snooping and helps prevent ARP attacks by doing the following:

• Not relaying invalid or gratuitous ARP replies out to other ports in the same VLAN

• Intercepting all ARP requests and replies on untrusted ports

• Verifying each intercepted packet for a valid IP-to-MAC binding

• Dropping and logging ARP replies coming from invalid sources to prevent ARP poisoning

• Error disabling the interface if the configured DAI number of ARP packets is exceeded

To mitigate the chances of ARP spoofing and ARP poisoning, follow these DAI implementation guidelines:

• Enable DHCP snooping globally.

• Enable DHCP snooping on selected VLANs.

• Enable DAI on selected VLANs.

Technet24

• Configure trusted interfaces for DHCP snooping and ARP inspection.

The topology in Figure 20-8 identifies trusted and untrusted ports

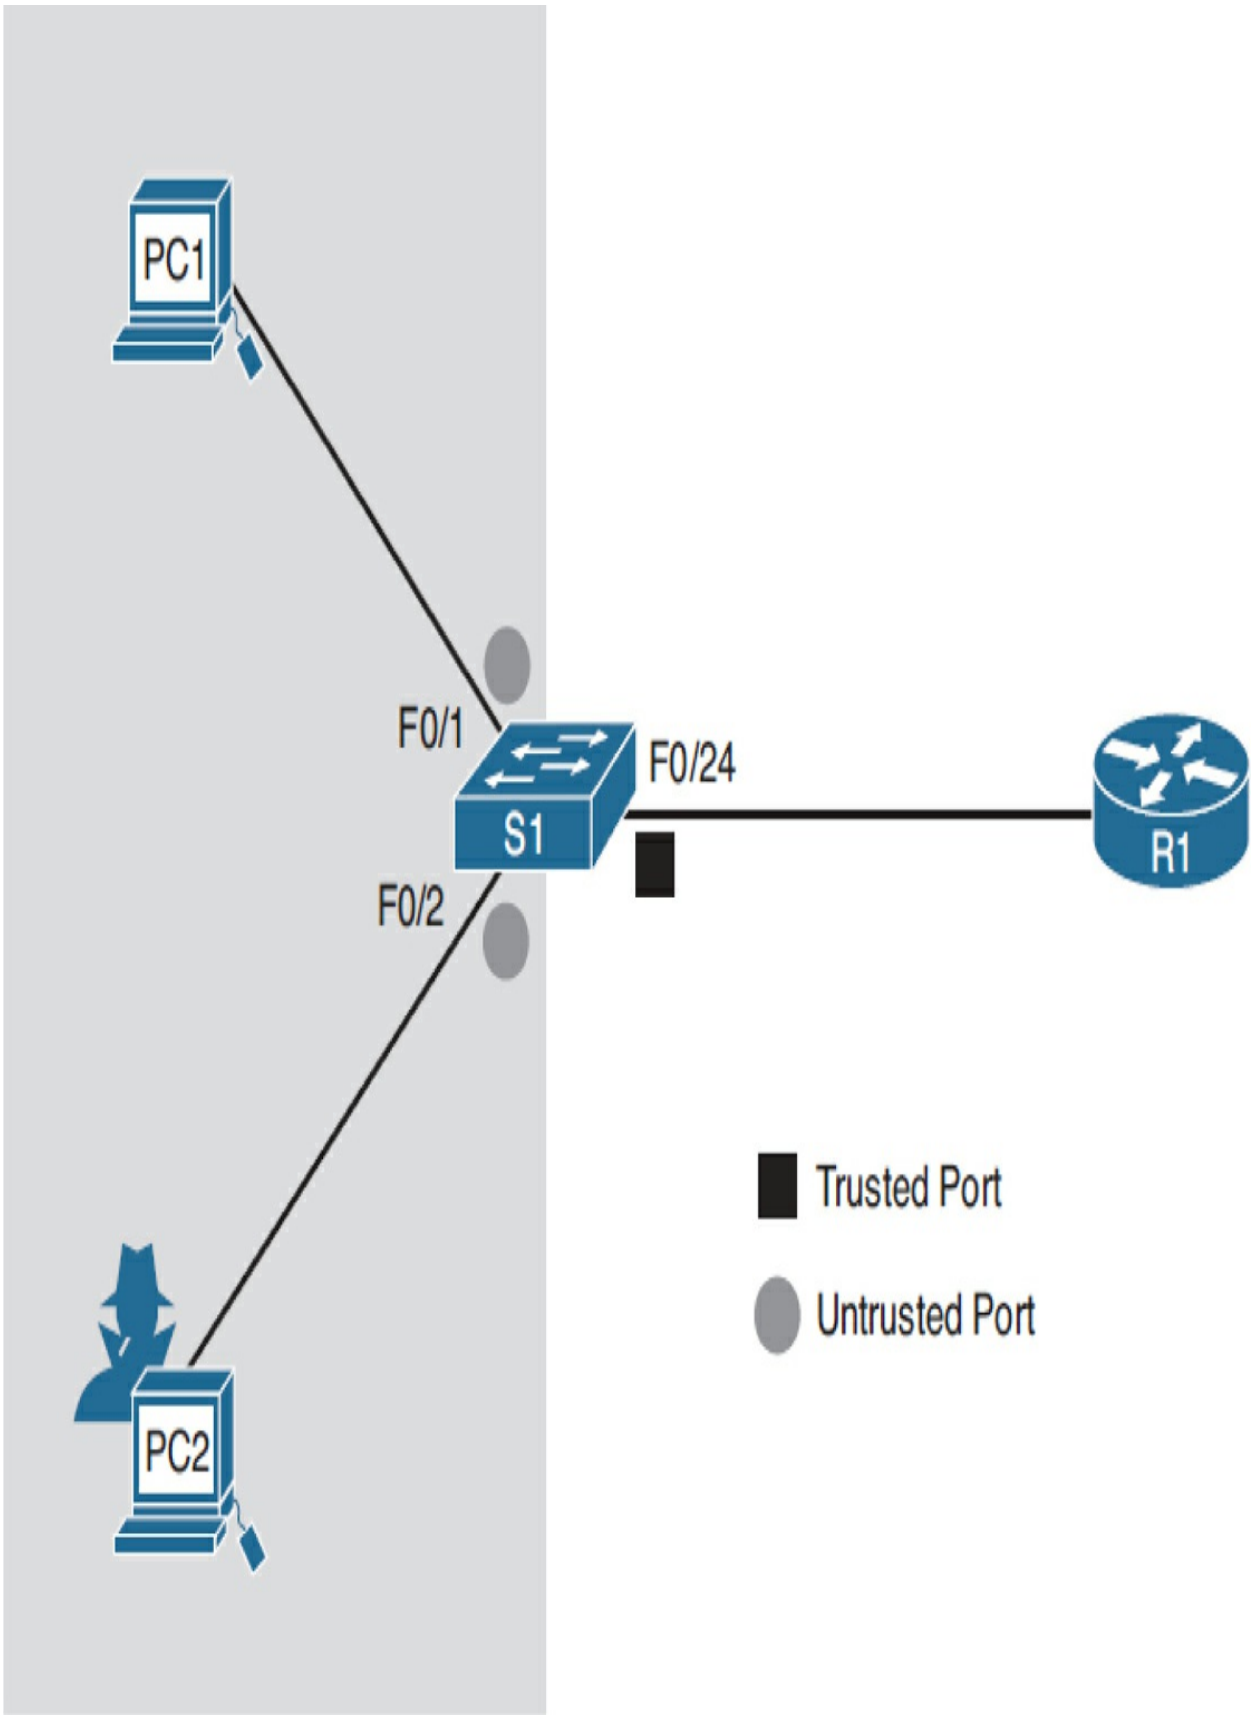

Technet24

Figure 20-8 *Trusted and Untrusted Ports for DAI Configuration*

In Figure 20-8, S1 is connecting two users on VLAN 10. In Example 20-12, DAI is configured to mitigate against ARP spoofing and ARP poisoning attacks. Notice that DHCP snooping is enabled because DAI requires the DHCP snooping binding table to operate.

Example 20-12 *DAI Configuration*

```
S1(config)# ip dhcp snooping
S1(config)# ip dhcp snooping vlan 10
S1(config)# ip arp inspection vlan 10
S1(config)# interface fa0/24
S1(config-if)# ip dhcp snooping trust
S1(config-if)# ip arp inspection trust
```

DAI can also be configured to check for both destination or source MAC and IP addresses with the **ip arp inspection validate** command. Only one command can be configured. Entering multiple **ip arp inspection validate** commands overwrites the previous command. To include more than one validation method, enter them on the same command line, as shown and verified in Example 20-13.

Example 20-13 *Configuring DAI to Validate MAC and IP Addresses*

```
S1(config)# ip arp inspection validate ?
```

```
    dst-mac Validate destination MAC address
    ip    Validate IP addresses
    src-mac Validate source MAC address
S1(config)# ip arp inspection validate src-mac
S1(config)# ip arp inspection validate dst-mac
S1(config)# ip arp inspection validate ip
S1(config)# do show run | include validate
ip arp inspection validate ip
S1(config)# ip arp inspection validate src-mac dst-mac
S1(config)# do show run | include validate
ip arp inspection validate src-mac dst-mac ip
S1(config)#
```

Study Resources

For today’s exam topics, refer to the following resources for more study.

Technet24

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	10–11
CCNA 200-301 Official Cert Guide, Volume 1, 2e	6
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	9–12
Portable Command Guide	20
	22

Day 19. Basic Routing Concepts

CCNA 200-301 Exam Topics

• 1.1 Explain the role and function of network components

• 3.2 Determine how a router makes a forwarding decision by default

Key Topics

Today we review basic routing concepts, including exactly how a packet is processed by intermediary devices (routers) on its way from source to destination. We then review the basic routing methods, including connected, static, and dynamic routes. We conclude the day's review with a deep dive into the operation of dynamic routing protocols.

Packet Forwarding

Packet forwarding by routers is accomplished through path determination and switching functions. The path determination function is the process the router use to determine which path to use when forwarding a packet. To determine the best path, the router searches its routing table for a network address that matches the packet's destination IP address.

This search results in one of three path determinations:

• **Directly connected network:** If the destination IP address of the packet

Technet24

belongs to a device on a network that is directly connected to one of the router's interfaces, that packet is forwarded directly to that device. This means the destination IP address of the packet is a host address on the same network as this router's interface.

• **Remote network:** If the destination IP address of the packet belongs to a remote network, the packet is forwarded to another router. Remote networks can be reached only by forwarding packets to another router.

• **No route determined:** If the destination IP address of the packet does not belong to a connected or remote network and the router does not have a default route, the packet is discarded. The router sends an Internet Control Message Protocol (ICMP) Unreachable message to the source IP address of the packet.

In the first two results, the router completes the process by switching the packet out the correct interface. It does this by reencapsulating the IP packet into the appropriate Layer 2 data-link frame format for the exit interface. The type of interface determines the type of Layer 2 encapsulation. For example, if the exit interface is Fast Ethernet, the packet is encapsulated in an Ethernet frame. If the exit interface is a serial interface configured for PPP, the IP packet is encapsulated in a PPP frame.

Path Determination and Switching Function Example

Let's review the process of path determination and switching functions that

routers perform as a packet travels from source to destination. Consider the topology in Figure 19-1 and the following steps:

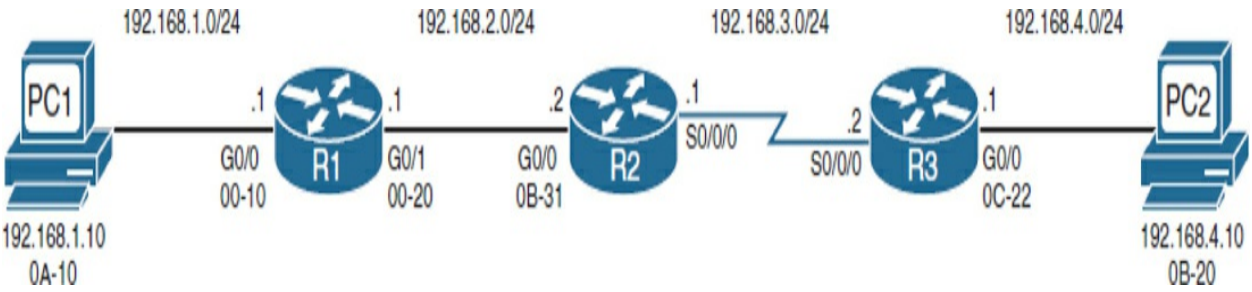

Figure 19-1 *Packet Forwarding Sample Topology*

Note

For brevity, Figure 19-1 shows only the last two octets of the MAC address.

PC1 has a packet to send to PC2. Using the AND operation on the destination's IP address and PC1's subnet mask, PC1 has determined that the IP source and IP destination addresses are on different networks. Therefore, PC1 checks its Address Resolution Protocol (ARP) table for the IP address of the default gateway and its associated MAC address. It then encapsulates the packet in an Ethernet header and forwards it to R1.

Router R1 receives the Ethernet frame. Router R1 examines the destination MAC address, which matches the MAC address of the receiving interface, G0/0. R1 therefore copies the frame into its buffer to be processed.

R1 decapsulates the Ethernet frame and reads the destination IP address.

Technet24

Because it does not match any of R1's directly connected networks, the router consults its routing table to route this packet.

R1 searches the routing table for a network address and subnet mask that include this packet's destination IP address as a host address on that network. It selects the entry with the longest match (longest prefix). R1 encapsulates the packet in the appropriate frame format for the exit interface and switches the frame to the interface (G0/1 in this example). The interface then forwards it to the next hop.

The packet arrives at router R2. R2 performs the same functions as R1, but this time, the exit interface is a serial interface—not Ethernet. Therefore, R2 encapsulates the packet in the appropriate frame format for the serial interface and sends it to R3. For this example, assume that the interface is using High-Level Data Link Control (HDLC), which uses the data-link address 0x8F. Remember that serial interfaces do not use MAC addresses.

The packet arrives at R3. R3 decapsulates the data-link HDLC frame. The search of the routing table results in a network that is one of R3's directly connected networks. Because the exit interface is a directly connected Ethernet network, R3 needs to resolve the destination IP address of the packet with a destination MAC address.

R3 searches for the packet's destination IP address, 192.168.4.10, in its ARP cache. If the entry is not in the ARP cache, R3 sends an ARP request out its

G0/0 interface.

PC2 sends back an ARP reply with its MAC address. R3 updates its ARP cache with an entry for 192.168.4.10 and the MAC address returned in the ARP reply.

The IP packet is encapsulated into a new data-link Ethernet frame and sent out R3's G0/0 interface.

The Ethernet frame with the encapsulated IP packet arrives at PC2. PC2 examines the destination MAC address, which matches the MAC address of the receiving interface—that is, its own Ethernet NIC. PC2 therefore copies the rest of the frame. PC2 sees that the Ethernet Type field is 0x800, which means that the Ethernet frame contains an IP packet in the data portion of the frame. PC2 decapsulates the Ethernet frame and passes the IP packet to its operating system's IP process.

Routing Methods

A router can learn routes from three basic sources:

• **Directly connected routes:** Automatically entered in the routing table when an interface is activated with an IP address

• **Static routes:** Manually configured by the network administrator and entered in the routing table if the exit interface for the static route is active

Technet24

• **Dynamic routes:** Learned by the routers through sharing routes with other routers that use the same routing protocol

In many cases, the complexity of the network topology, the number of networks, and the need for the network to automatically adjust to changes require the use of a dynamic routing protocol. Dynamic routing certainly has several advantages over static routing; however, networks still use static routing. In fact, networks typically use a combination of static and dynamic routing.

Table 19-1 compares dynamic and static routing features. From this comparison, you can list the advantages of each routing method. The advantages of one method are the disadvantages of the other.

Table 19-1 *Dynamic Versus Static Routing*

Feature	Dynamic Routing	Static Routing
Configuration complexity	Generally remains independent of the network size	Increases with network size
Required administrator knowledge	Requires advanced knowledge	Requires no extra knowledge
Topology changes	Automatically adapts to topology changes	Requires administrator intervention
Scaling	Suitable for simple and complex topologies	Suitable for simple topologies
Security	Less secure	More secure
Resource usage	Uses CPU, memory, and link bandwidth	Requires no extra resources
Predictability	Uses a route that depends on the current topology	Always uses the same route to the destination

Classifying Dynamic Routing Protocols

Technet24

Figure 19-2 shows a timeline of IP routing protocols, along with a chart to help you memorize the various ways to classify routing protocols.

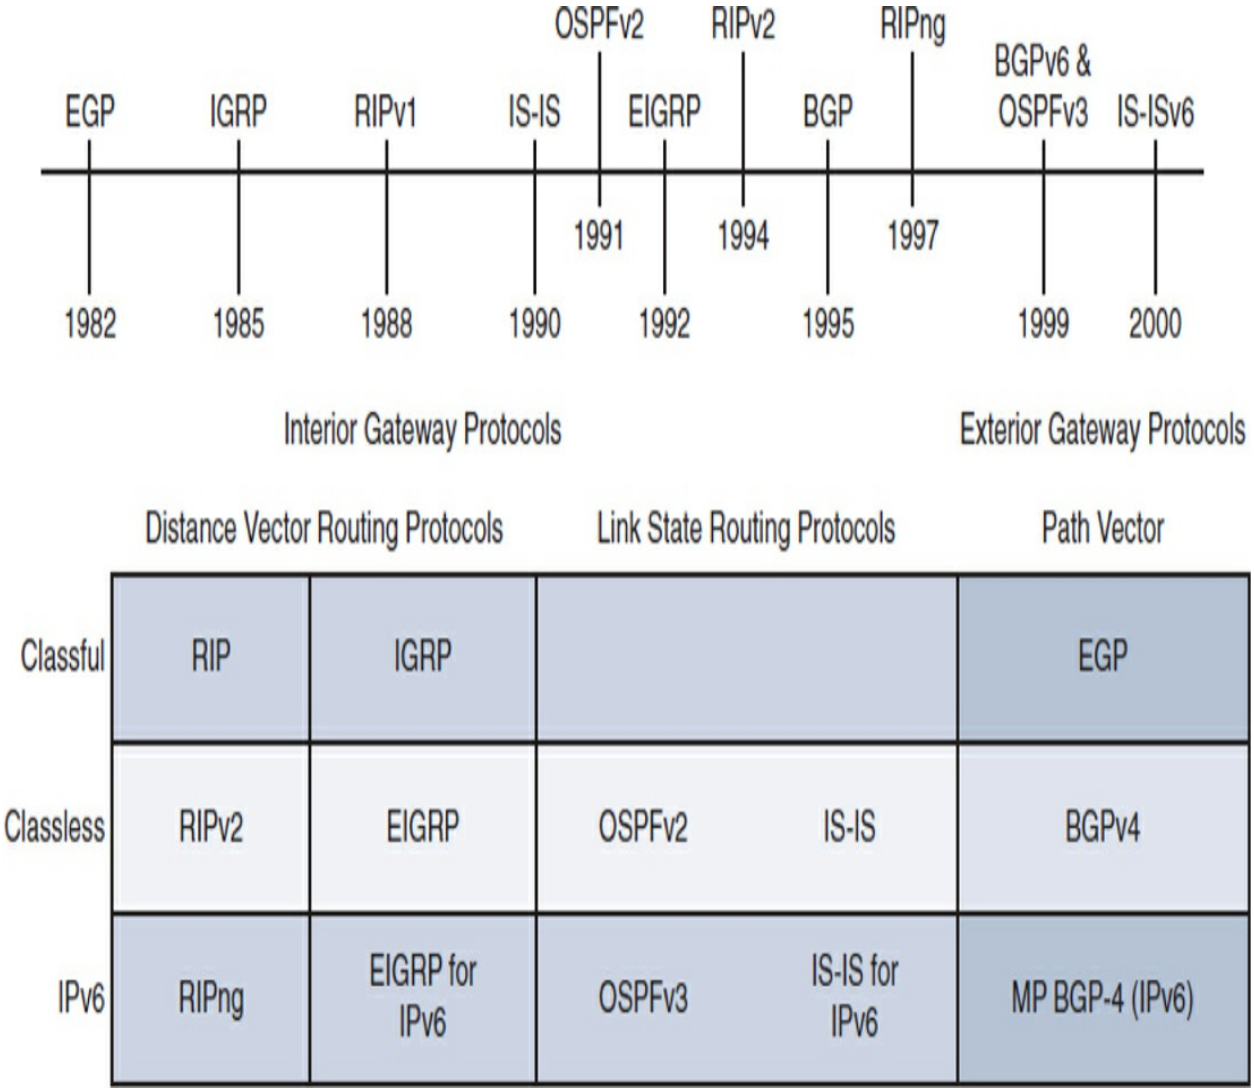

Figure 19-2 *Evolution and Classification of Routing Protocols*

Routing protocols are classified into different groups according to their characteristics:

• IGP or EGP

• Distance vector or link state

• Classful or classless

IGP and EGP

An autonomous system (AS) is a collection of routers under a common administration that presents a common, clearly defined routing policy to the Internet. Typical examples are a large company's internal network and an ISP's network. Most company networks are not autonomous systems; in most cases, a company network is a network within its ISP's autonomous system. Because the Internet is based on the autonomous system concept, two types of routing protocols are required:

• **Interior gateway protocols (IGP):** Used for intra-AS routing—that is, routing inside an AS

• **Exterior gateway protocols (EGP):** Used for inter-AS routing—that is, routing between autonomous systems

Distance Vector Routing Protocols

Distance vector means that routes are advertised as vectors of distance and direction. Distance is defined in terms of a metric such as hop count, and direction is the next-hop router or exit interface. Distance vector protocols typically use the Bellman-Ford algorithm for the best-path route determination.

Technet24

Some distance vector protocols periodically send complete routing tables to all connected neighbors. In large networks, these routing updates can become enormous, causing significant traffic on the links.

Although the Bellman-Ford algorithm eventually accumulates enough knowledge to maintain a database of reachable networks, the algorithm does not allow a router to know the exact topology of an internetwork. The router knows only the routing information received from its neighbors.

Distance vector protocols use routers as signposts along the path to the final destination. The only information a router knows about a remote network is the distance or metric to reach that network and which path or interface to use to get there. A distance vector routing protocol does not have a map of the network topology.

Distance vector protocols work best in these situations:

• When the network is simple and flat and does not require a hierarchical design

• When the administrators do not have enough knowledge to configure and troubleshoot link-state protocols

• When specific types of networks, such as hub-and-spoke networks, are being implemented

• When worst-case convergence times in a network are not a concern

Link-State Routing Protocols

In contrast to distance vector routing protocol operation, a router configured with a link-state routing protocol can create a complete view, or topology, of the network by gathering information from all the other routers. Think of a link-state routing protocol as having a complete map of the network topology. The signposts along the way from source to destination are not necessary because all link-state routers are using an identical map of the network. A link-state router uses the link-state information to create a topology map and to select the best path to each destination network in the topology.

With some distance vector routing protocols, routers periodically send updates of their routing information to their neighbors. Link-state routing protocols do not use periodic updates. After the network has converged, a link-state update is sent only when the topology changes.

Link-state protocols work best in these situations:

• When the network design is hierarchical, which is typically the case in large networks

• When the administrators have good knowledge of the implemented link-state routing protocol

Technet24

• When fast convergence of the network is crucial

Classful Routing Protocols

Classful routing protocols do not send subnet mask information in routing updates. The first routing protocols, such as Routing Information Protocol (RIP), were classful. When those protocols were created, network addresses were allocated based on class: Class A, B, or C. A routing protocol did not need to include the subnet mask in the routing update because the network mask could be determined based on the first octet of the network address.

Classful routing protocols can still be used in some of today's networks, but because they do not include the subnet mask, they cannot be used in all situations. Classful routing protocols cannot be used when a network is subnetted using more than one subnet mask. In other words, classful routing protocols do not support variable-length subnet masking (VLSM).

Other limitations come into play with classful routing protocols, including their inability to support discontiguous networks and supernets. Classful routing protocols include Routing Information Protocol version 1 (RIPv1) and Interior Gateway Routing Protocol (IGRP). CCNA exam topics do not include either RIPv1 or IGRP.

Classless Routing Protocols

Classless routing protocols include the subnet mask with the network address

in routing updates. Today's networks are no longer allocated based on class, and the subnet mask cannot be determined by the value of the first octet. Classless routing protocols are required in most networks today because of their support for VLSM and discontiguous networks and supernets. Classless routing protocols include Routing Information Protocol version 2 (RIPv2), Enhanced IGRP (EIGRP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), and Border Gateway Protocol (BGP).

Dynamic Routing Metrics

In some cases, a routing protocol learns of more than one route to the same destination from the same routing source. To select the best path, the routing protocol must be capable of evaluating and differentiating among the available paths. A *metric* is used for this purpose. Two different routing protocols might choose different paths to the same destination because they use different metrics. Metrics used in IP routing protocols include the following:

• **RIP—Hop count:** The best path is chosen by the route with the lowest hop count.

• **IGRP and EIGRP—Bandwidth, delay, reliability, and load:** The best path is chosen by the route with the smallest composite metric value calculated from these multiple parameters. By default, only bandwidth and delay are used.

Technet24

• **IS-IS and OSPF—Cost:** The best path is chosen by the route with the lowest cost. The Cisco implementation of OSPF uses bandwidth to determine the cost.

The metric associated with a certain route can best be viewed using the **show ip route** command. The metric value is the second value in the brackets for a routing table entry. In Example 19-1, R2 has a route to the 192.168.8.0/24 network that is two hops away.

Example 19-1 *Routing Table for R2*

```
R2# show ip route

<output omitted>

Gateway of last resort is not set

R   192.168.1.0/24 [120/1] via 192.168.2.1, 00:00:20,
    192.168.2.0/24 is variably subnetted, 2 subnets, 2
C     192.168.2.0/24 is directly connected, Serial0/0/
L     192.168.2.2/32 is directly connected, Serial0/0/
    192.168.3.0/24 is variably subnetted, 2 subnets, 2
C     192.168.3.0/24 is directly connected, GigabitEth
L     192.168.3.1/32 is directly connected, GigabitEth
   192.168.4.0/24 is variably subnetted, 2 subnets, 2
C     192.168.4.0/24 is directly connected, Serial0/0/
L     192.168.4.2/32 is directly connected, Serial0/0/
```

```
R  192.168.5.0/24 [120/1] via 192.168.4.1, 00:00:25, S
R 192.168.6.0/24  [120/1] via 192.168.2.1, 00:00:20, S
                  [120/1] via 192.168.4.1, 00:00:25, S
R 192.168.7.0/24 [120/1] via 192.168.4.1, 00:00:25, Se
R 192.168.8.0/24 [120/2] via 192.168.4.1, 00:00:25, Se
```

Notice in the output that one network, 192.168.6.0/24, has two routes. RIP will load balance between these equal-cost routes. All the other routing protocols are capable of automatically load balancing traffic for up to four equal-cost routes, by default. EIGRP is also capable of load balancing across unequal-cost paths.

Administrative Distance

Sometimes a router learns a route to a remote network from more than one routing source. For example, a static route might have been configured for the same network/subnet mask that was learned dynamically by a dynamic routing protocol, such as RIP. The router must choose which route to install.

Although it is less common, more than one dynamic routing protocol can be deployed in the same network. In some situations, it might be necessary to route the same network address using multiple routing protocols, such as RIP and OSPF. Because different routing protocols use different metrics—for example, RIP uses hop count and OSPF uses bandwidth—it is not possible to compare metrics to determine the best path.

Technet24

Administrative distance (AD) defines the preference of a routing source. Each routing source—including specific routing protocols, static routes, and even directly connected networks—is prioritized in order of most preferable to least preferable, using an AD value. Cisco routers use the AD feature to select the best path when they learn about the same destination network from two or more different routing sources.

The AD value is an integer value from 0 to 255. The lower the value, the more preferred the route source. An administrative distance of 0 is the most preferred. Only a directly connected network has an AD of 0, which cannot be changed. An AD of 255 means the router will not believe the source of that route, and it will not be installed in the routing table.

In the routing table in Example 19-1, the AD value is the first value listed in the brackets. You can see that the AD value for RIP routes is 120. You can also verify the AD value with the **show ip protocols** command, as Example 19-2 demonstrates.

Example 19-2 *Verifying the AD Value with the* show ip protocols *Command*

```
R2# show ip protocols

Routing Protocol is "rip"
   Outgoing update filter list for all interfaces is n
   Incoming update filter list for all interfaces is n
   Sending updates every 30 seconds, next due in 21 se
```

```
    Invalid after 180 seconds, hold down 180, flushed a
    Redistributing: rip
    Default version control: send version 1, receive an
       Interface          Send Recv Triggered RIP Key-c
       GigabitEthernet0/0  1    1 2
       Serial0/0/0         1    1 2
       Serial0/0/1         1    1 2
    Automatic network summarization is in effect
    Maximum path: 4
    Routing for Networks:
      192.168.2.0
      192.168.3.0
      192.168.4.0
    Routing Information Sources:
      Gateway      Distance  Last Update
      192.168.2.1       120   00:00:01
      192.168.4.1       120   00:00:01
    Distance: (default is 120)

  R2#
```

Table 19-2 shows the different administrative distance values for various routing protocols.

Table 19-2 *Default Administrative Distances*

Technet24

Route Source	AD
Connected	0
Static	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200

IGP Comparison Summary

Table 19-3 compares several features of the IGPs that are currently most popular: RIPv2, OSPF, and EIGRP.

Table 19-3 *Comparing Features of IGPs: RIPv2, OSPF, and EIGRP*

Technet24

Features	RIPv2	OSPF	EIGRP
Metric	Hop count	Bandwidth	Function of bandwidth, delay
Sends periodic updates	Yes (30 seconds)	No	No
Full or partial routing updates	Full	Partial	Partial
Where updates are sent	(224.0.0.9)	(224.0.0.5, 224.0.0.6)	(224.0.0.10)
Route considered unreachable	16 hops	Depends on MaxAge of LSA, which is never incremented past 3600 seconds	A delay of all 1s
Supports unequal-cost load balancing	No	No	Yes

Routing Loop Prevention

Without preventive measures, distance vector routing protocols can cause severe routing loops in a network. A routing loop is a condition in which a packet is continuously transmitted within a series of routers without ever reaching its intended destination network. A routing loop can occur when two or more routers have inaccurate routing information to a destination network.

Several mechanisms are available to eliminate routing loops, primarily with distance vector routing protocols. These mechanisms include the following:

• **A maximum metric to prevent count to infinity:** To eventually stop the incrementing of a metric during a routing loop, infinity is defined by setting a maximum metric value. For example, RIP defines infinity as 16 hops, an unreachable metric. When the routers “count to infinity,” they mark the route as unreachable.

• **Hold-down timers:** Routers are instructed to hold any changes that might affect routes for a specified period of time. If a route is identified as down or possibly down, any other information for that route containing the same status, or worse, is ignored for a predetermined amount of time (the hold-down period) so that the network has time to converge.

• **Split horizon:** A routing loop is prevented by not allowing advertisements to be sent back through the interface where they originated. The split horizon rule stops a router from incrementing a metric and then sending the route

Technet24

back to its source.

• **Route poisoning or poison reverse:** The route is marked as unreachable in a routing update that is sent to other routers. *Unreachable* is interpreted as a metric that is set to the maximum.

• **Triggered updates:** A routing table update is sent immediately in response to a routing change. Triggered updates do not wait for update timers to expire. The detecting router immediately sends an update message to adjacent routers.

• **TTL field in the IP header:** The Time To Live (TTL) field avoids a situation in which an undeliverable packet circulates endlessly on the network. With TTL, the source device of the packet sets the 8-bit field with a value. This TTL value is decreased by 1 by every router in the path until the packet reaches its destination. If the TTL value reaches 0 before the packet arrives at its destination, the packet is discarded, and the router sends an ICMP error message back to the source of the IP packet.

Link-State Routing Protocol Features

Just as distance vector protocols send routing updates to their neighbors, link-state protocols send link-state updates to neighboring routers, which then forward that information to their neighbors, and so on. Also as with distance vector protocols, at the end of the process, routers that use link-state protocols add the best routes to their routing tables, based on metrics.

However, beyond this level of explanation, these two types of routing protocol algorithms have little in common.

Building the LSDB

Link-state routers flood detailed information about the internetwork to all the other routers so that every router has the same information about the internetwork. Routers use this link-state database (LSDB) to calculate the current best routes to each subnet.

OSPF, the most popular link-state IP routing protocol, advertises information in routing update messages of various types. The updates contain information called link-state advertisements (LSA).

Figure 19-3 shows the general idea of the flooding process. R8 is creating and flooding its router LSA. Note that Figure 19-3 shows only a subset of the information in R8's router LSA.

Figure 19-3 shows the basic flooding process. R8 is sending the original LSA for itself, and the other routers are flooding the LSA by forwarding it until every router has a copy.

Technet24

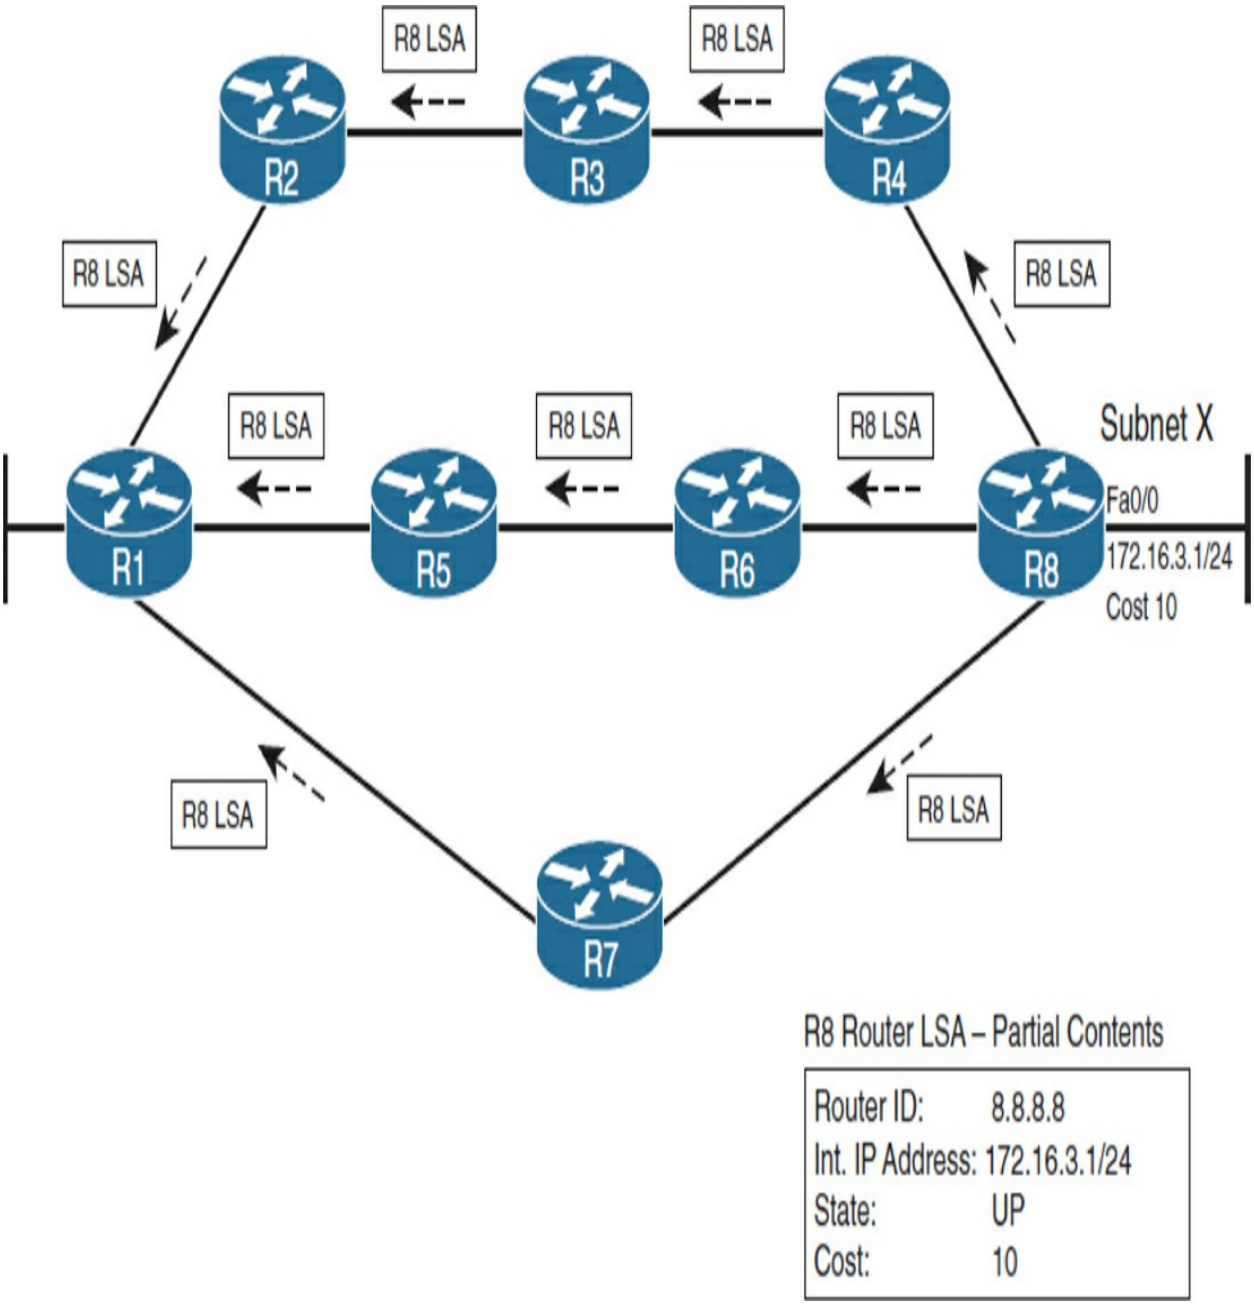

Figure 19-3 *Flooding LSAs Using a Link-State Routing Protocol*

After the LSA has been flooded, even if the LSAs do not change, link-state protocols require periodic reflooding of the LSAs by default every 30 minutes. However, if an LSA changes, the router immediately floods the changed LSA. For example, if Router R8's LAN interface failed, R8 would

need to reflood the R8 LSA, stating that the interface is now down.

Calculating the Dijkstra Algorithm

The flooding process alone does not cause a router to learn what routes to add to the IP routing table. Link-state protocols must then find and add routes to the IP routing table by using the Dijkstra shortest path first (SPF) algorithm.

The SPF algorithm is run on the LSDB to create the SPF tree. The LSDB holds all the information about all the possible routers and links. Each router must view itself as the starting point and each subnet as the destination, and it must use the SPF algorithm to build its own SPF tree to pick the best route to each subnet.

Figure 19-4 shows a graphical view of route possibilities from the results of the SPF algorithm run by router R1 when trying to find the best route to reach subnet 172.16.3.0/24 (based on Figure 19-3).

To pick the best route, a router's SPF algorithm adds the cost associated with each link between itself and the destination subnet over each possible route. Figure 19-4 shows the costs associated with each route beside the links. The dashed lines show the three routes R1 finds between itself and subnet X (172.16.3.0/24).

Technet24

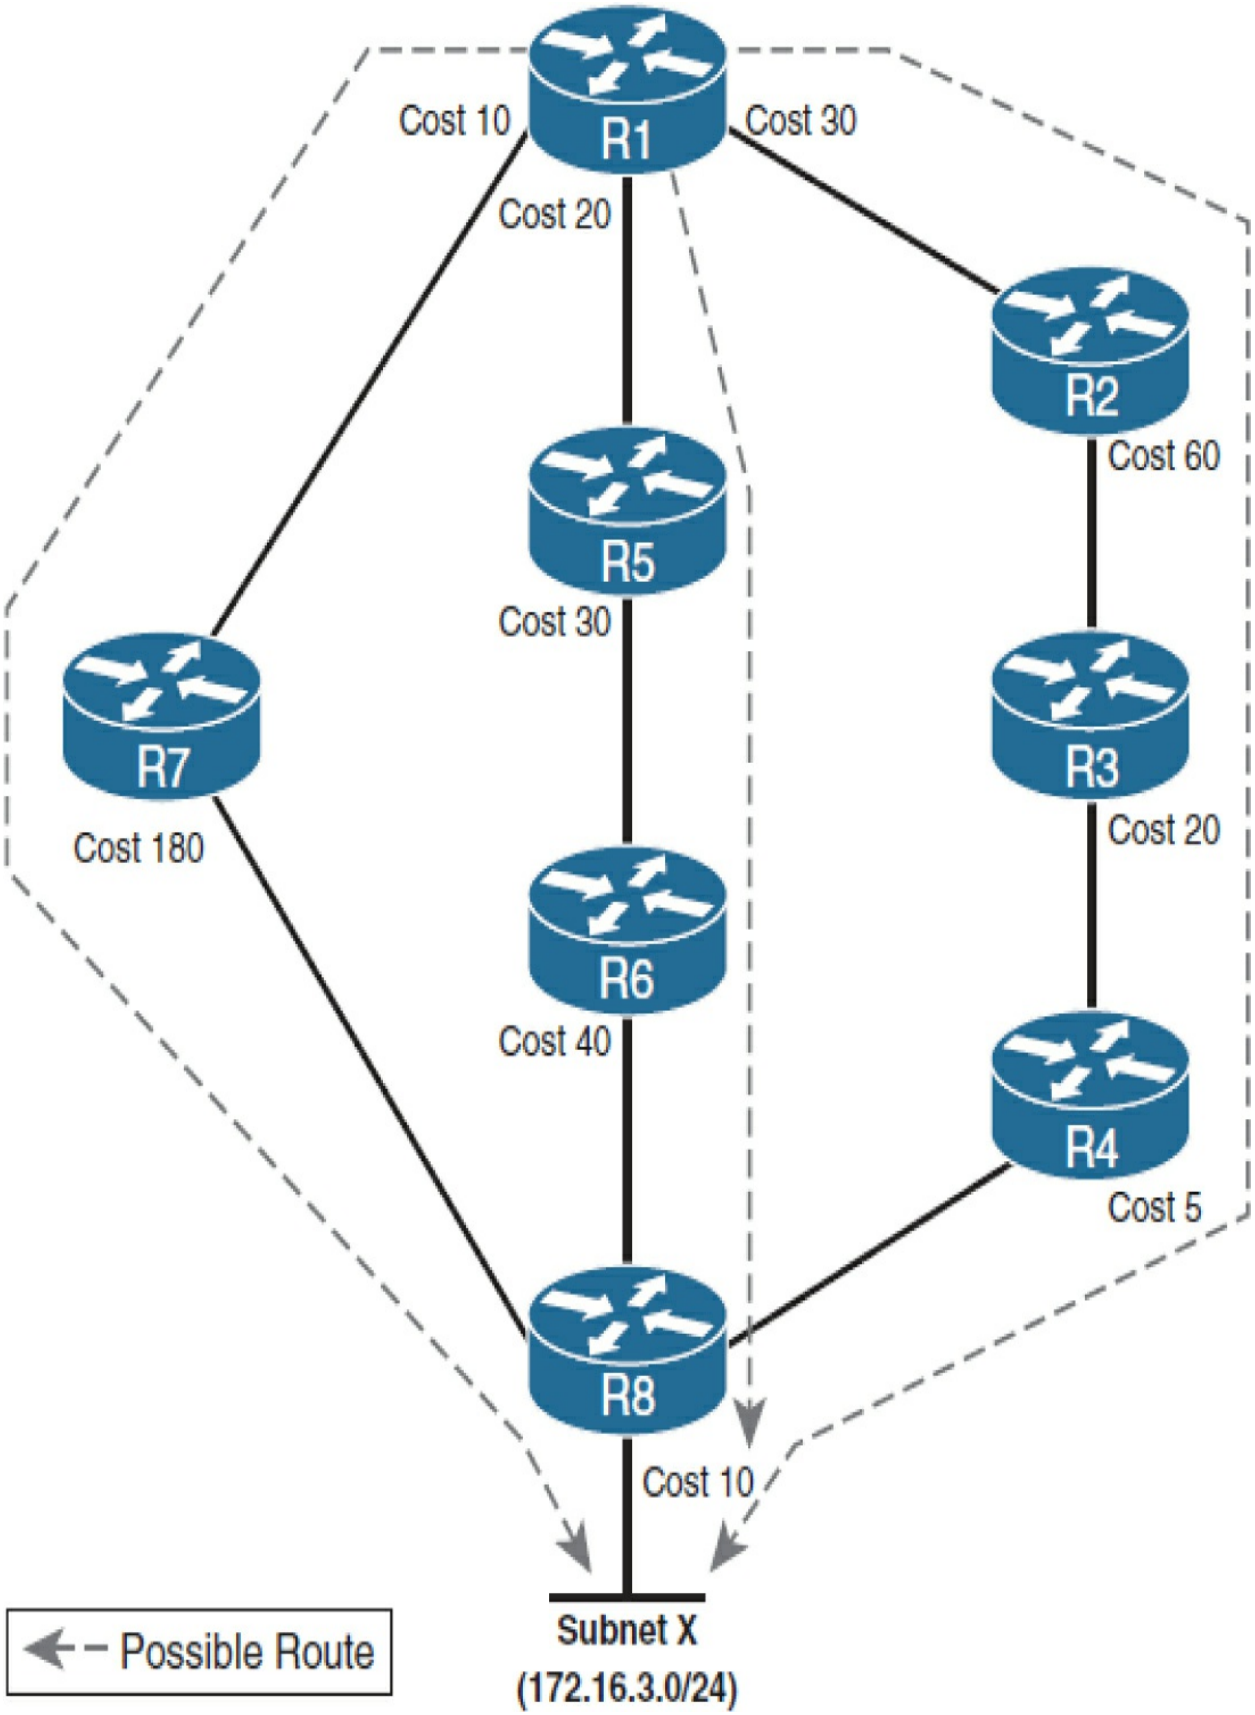
R1
Cost 10
Cost 30
Cost 20
R2
Cost 60
R5
Cost 30
R7
Cost 180
R3
Cost 20
R6
Cost 40
R4
Cost 5
R8
Cost 10
Subnet X
(172.16.3.0/24)
Possible Route

Figure 19-4 *SPF Tree to Find R1's Route to 172.16.3.0/24*

Table 19-4 lists the three routes shown in Figure 19-2, with their cumulative costs. You can see that R1's best route to 172.16.3.0/24 starts by going through R5.

Table 19-4 *Comparing R1's Three Alternatives for the Route to 172.16.3.0/24*

Route	Location in Figure 19-2	Cumulative Cost
R1–R7–R8	Left	10 + 180 + 10 = 200
R1–R5–R6–R8	Middle	20 + 30 + 40 + 10 = 100
R1-R2–R3–R4–R8	Right	30 + 60 + 20 + 5 + 10 = 125

As a result of the SPF algorithm's analysis of the LSDB, R1 adds to its routing table a route to subnet 172.16.3.0/24, with R5 as the next-hop router.

Convergence with Link-State Protocols

Remember that when an LSA changes, link-state protocols react swiftly, converging the network and using the current best routes as quickly as possible. For example, imagine that the link between R5 and R6 fails in the

Technet24

internetwork in Figures 25-3 and 25-4. R1 then uses the following process to switch to a different route:

R5 and R6 flood LSAs, stating that their interfaces are now in a down state.

All routers run the SPF algorithm again to see if any routes have changed.

All routers replace routes, as needed, based on the results of SPF. For example, R1 changes its route for subnet X (172.16.3.0/24) to use R2 as the next-hop router.

These steps allow the link-state routing protocol to converge quickly—much more quickly than distance vector routing protocols.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	8
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	14
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	1
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	3
	17
	21–24

Technet24

Day 18. Basic Router Configuration

CCNA 200-301 Exam Topics

• 1.2 Describe characteristics of network topology architectures

• 1.4 Identify interface and cable issues (collisions, errors, mismatch duplex, and/or speed)

• 1.6 Configure and verify IPv4 addressing and subnetting

• 1.8 Configure and verify IPv6 addressing and prefix

Key Topics

Today we review basic router configuration. First, we focus on configuring and verifying initial settings, including IPv4 addressing. We look at the details of the **show interface** command to understand how it can help in identifying interface and cable issues. Then we review IPv6 addressing and network connectivity verification. Most of this should be very familiar at this point in your studies because these skills are fundamental to all other router configuration tasks. We also review small office or home office (SOHO) setup and configuration.

Basic Router Configuration with IPv4

Figure 18-1 shows the topology and IPv4 addressing scheme that we use

today to review basic router configuration and verification tasks.

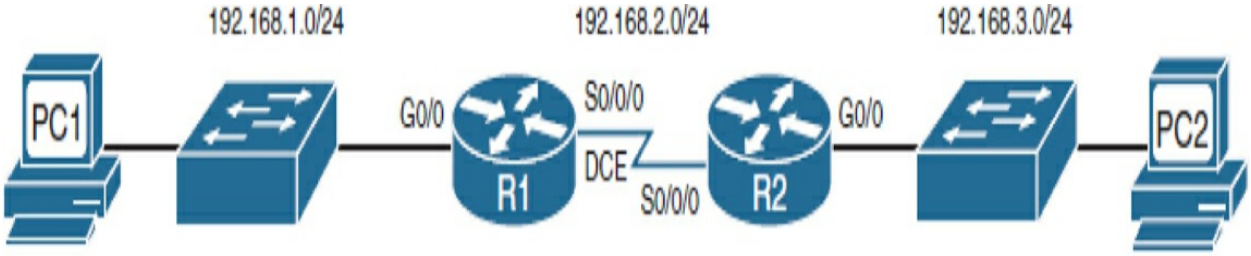

Device	Interface	IP Address	Subnet Mask	Default Gateway
R1	G0/0	192.168.1.1	255.255.255.0	N/A
	S0/0/0	192.168.2.1	255.255.255.0	N/A
R2	G0/0	192.168.3.1	255.255.255.0	N/A
	S0/0/0	192.168.2.2	255.255.255.0	N/A
PC1	N/A	192.168.1.10	255.255.255.0	192.168.1.1
PC2	N/A	192.168.3.10	255.255.255.0	192.168.3.1

Figure 18-1 *IPv4 Example Topology*

When configuring a router, certain basic tasks are performed:

• Naming the router

• Setting passwords

• Configuring interfaces

• Configuring a banner

Technet24

• Saving changes on a router

• Verifying basic configuration and router operations

Command Syntax

Table 18-1 shows the basic router configuration command syntax used to configure R1 in the following example.

Table 18-1 *Basic Router Configuration Command Syntax*

Configuration Task	Commands
Naming the router	`Router(config)# hostname name`
Setting passwords	`Router(config)# enable secret password`
	`Router(config)# line console 0`
	`Router(config-line)# password password`
	`Router(config-line)# login`
	`Router(config)# line vty 0 15`
	`Router(config-line)# transport input ssh`
	`Router(config-line)# login local`
	`Router(config)# username name password password`

Technet24

Configuring a message-of-the-day banner	`Router(config)# banner motd # message #`
Configuring an interface	`Router(config)# interface type number`
	`Router(config-if)# ip address address mask`
	`Router(config-if)# description description`
	`Router(config-if)# no shutdown`
Saving changes on a router	`Router# copy running-config startup-config`
Examining the output of **show** commands	`Router# show running-config`
	`Router# show ip route`
	`Router# show ip interface brief`
	`Router# show interfaces`

Configuration Example

Let’s walk through a basic configuration for R1. First, enter privileged EXEC mode and then global configuration mode:

```
Router> enable
Router# config t
```

Next, name the router and enter the encrypted password for entering privileged EXEC mode. This command overrides the older **enable password** *password* command, so you are not entering that one:

```
Router(config)# hostname R1
R1(config)# enable secret class
```

Next, configure the console password and require that it be entered with the login password:

```
R1(config)# line console 0
R1(config-line)# password cisco
R1(config-line)# login
```

Configuring SSH and disabling Telnet are security best practices, so configure the vty lines to use only SSH:

Note

SSH configuration is not shown here; assume that it is already configured.

Technet24

```
R1(config)# line vty 0 15
R1(config-line)# transport input ssh
R1(config-line)# login local
R1(config-line)# exit
R1(config)# username admin password cisco
```

Encrypt all the plaintext passwords in the running configuration by using the **service-password encryption** command:

```
R1(config)# service-password encryption
```

Configure the message-of-the-day (MOTD) banner. A delimiting character such as a # is used at both the beginning and the end of the message. At a minimum, a banner should warn against unauthorized access. A good security policy prohibits configuring a banner that welcomes an unauthorized user:

```
R1(config)# banner motd #
Enter TEXT message.  End with the character '#'.
**************************************
WARNING!! Unauthorized Access Prohibited!!
**************************************
#
```

Now configure the individual router interfaces with IP addresses and other

information. First, enter interface configuration mode by specifying the interface type and number and then configure the IP address and subnet mask:

```
R1(config)# interface Serial0/0/0
R1(config-if)# ip address 192.168.2.1 255.255.255.0
```

It is good practice to configure a description on each interface to help document the network information:

```
R1(config-if)# description Ciruit#VBN32696-123 (help d
```

Activate the interface:

```
R1(config-if)# no shutdown
```

Assuming that the other side of the link is activated on R2, the serial interface is now up. Finish R1 by configuring the GigabitEthernet 0/0 interface:

```
R1(config-if)# interface GigabitEthernet0/0
R1(config-if)# ip address 192.168.1.1 255.255.255.0
R1(config-if)# description R1 LAN
R1(config-if)# no shutdown
```

Technet24

Assume that R2 is fully configured and can route back to the 192.168.1.0/24 LAN attached to R1. You need to add a static route to R1 to ensure connectivity to R2's LAN. Static routing is reviewed in more detail on Day 19, "Basic Routing Concepts." For now, enter the following command to configure a directly attached static route to R2's LAN:

```
R1(config)# ip route 192.168.3.0 255.255.255.0 Serial
```

Note

Using a next-hop address is generally recommended when configuring static routes. Directly connected static routes should be used only with point-to-point serial interfaces, as in this example.

To save the configuration, enter the **copy running-config startup-config** command or the **copy run start** command.

Verification Example

You can use the **show running-config** command to verify the full current configuration on the router. However, a few other basic commands can help you verify your configuration and also begin troubleshooting any potential problems.

Make sure that the networks for your interfaces are now in the routing table by using the **show ip route** command, as shown in Example 18-1.

Example 18-1 *The* show ip route *Command*

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP,
       D - EIGRP, EX - EIGRP external, O - OSPF, IA -
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA
       E1 - OSPF external type 1, E2 - OSPF external t
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level
      ia - IS-IS inter area, * - candidate default, U
      o - ODR, P - periodic downloaded static route, H
      + - replicated route, % - next hop override
Gateway of last resort is not set

     192.168.1.0/24 is variably subnetted, 2 subnets,
C      192.168.1.0/24 is directly connected, GigabitEt
L      192.168.1.1/32 is directly connected, GigabitEt
    192.168.2.0/24 is variably subnetted, 2 subnets, 2
C      192.168.2.0/24 is directly connected, Serial0/0
L      192.168.2.1/32 is directly connected, Serial0/0
S   192.168.3.0/24 is directly connected, Serial0/0/0
R1#
```

If a network is missing, check your interface status with the **show ip**

Technet24

interface brief command, as shown in Example 18-2.

Example 18-2 *The* show ip interface brief *Command*

```
R1# show ip interface brief
Interface                  IP-Address   OK?   Method
Embedded-Service-Engine0/0   unassigned YES unset admi
GigabitEthernet0/0         192.168.1.1 YES manual up u
GigabitEthernet0/1         unassigned YES unset admini
Serial0/0/0                192.168.2.1 YES manual up up
Serial0/0/1                unassigned YES unset adminis
R1#
```

The output from the **show ip interface brief** command provides three important pieces of information:

• IP address

• Line status (column 5)

• Protocol status (column 6)

The IP address should be correct, and the status codes should be up and up. Table 18-2 summarizes the two status codes and their meanings.

Table 18-2 *Interface Status Codes*

Code	Location	General Meaning
Line status	First status code	Refers to the Layer 1 status—for example, is the cable installed, is it the right/wrong cable, is the device on the other end powered on?
Protocol status	Second status code	Refers generally to the Layer 2 status. It is always down if the line status is down. If the line status is up, a protocol status of down is usually caused by mismatched data link layer configuration.

Four combinations of settings are possible for the status codes when troubleshooting a network. Table 18-3 lists the four combinations and explains the typical reasons an interface might be in that state.

Table 18-3 *Combinations of Interface Status Codes*

Technet24

Line and Protocol Status	Typical Reason(s)
Administratively down, down	The interface has a **shutdown** command configured on it.
down, down	The interface has a **no shutdown** command configured, but the physical layer has a problem. For example, no cable has been attached to the interface (or with Ethernet), the switch interface on the other end of the cable is shut down, or the switch is powered off.
up, down	This almost always refers to data link layer problems, most often configuration problems. For example, serial links have this combination when one router was configured to use PPP and the other defaults to use HDLC. However, a clocking or hardware issue can also be to blame.
up, up	All is well and the interface is functioning.

If necessary, use the more verbose **show interface** command if you need to track down a problem with an interface and get the output for every physical

and virtual interface. You can also specify one interface. Example 18-3 shows the output for GigabitEthernet 0/0.

Example 18-3 *The show interface gigabitethernet 0/0 Command*

```
R1# show interface gigabitethernet 0/0
GigabitEthernet0/0 is up, line protocol is up
  Hardware is CN Gigabit Ethernet, address is 30f7.0da
  Description: R1 LAN
  Internet address is 192.168.1.1/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
      reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full Duplex, 100Mbps, media type is RJ45
  output flow-control is unsupported, input flow-contr
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:01, output hang ne
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Tota
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    387 packets input, 59897 bytes, 0 no buffer
     Received 252 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignc
```

Technet24

```
       0 watchdog, 86 multicast, 0 pause input
      281 packets output, 35537 bytes, 0 underruns
       0 output errors, 0 collisions, 1 interface resets
       56 unknown protocol drops
       0 babbles, 0 late collision, 0 deferred
       0 lost carrier, 0 no carrier, 0 pause output
       0 output buffer failures, 0 output buffers swappe
R1#
```

This command has a lot of output. However, wading through all this information is sometimes the only way to find a problem. Table 18-4 parses and explains each important part of the **show interface** output.

Table 18-4 show *interface Output Explanation*

Output	Description
GigabitEthernet…is {up | down | administratively down}	Whether the interface hardware is currently active or down or whether an administrator has taken it down.
line protocol is {up | down}	Whether the software processes that handle the line protocol consider the interface usable (that is, whether keepalives are successful). If the interface misses three consecutive keepalives, the line protocol is marked as down.
Hardware	Hardware type (for example, MCI Ethernet, serial communications interface [SCI], cBus Ethernet) and address.
Description	Text string description configured for the interface (with a maximum of 240 characters).
Internet address	IP address followed by the prefix length (subnet mask).

Technet24

MTU	Maximum transmission unit (MTU) of the interface.
BW	Bandwidth of the interface, in kilobits per second. The BW parameter is used to compute routing protocol metrics and other calculations.
DLY	Delay of the interface, in microseconds.
rely	Reliability of the interface as a fraction of 255 (where 255/255 is 100% reliability), calculated as an exponential average over 5 minutes.
load	Load on the interface as a fraction of 255 (where 255/255 is completely saturated), calculated as an exponential average over 5 minutes.
Encapsulation	Encapsulation method assigned to an interface.

Loopback	Whether the loopback is set. Can indicate a problem with the carrier.
Keepalive	Whether keepalives are set.
ARP type	Type of Address Resolution Protocol (ARP) assigned.
Last input	Number of hours, minutes, and seconds since the last packet was successfully received by an interface. Useful for knowing when a dead interface failed.
output	Number of hours, minutes, and seconds since the last packet was successfully transmitted by an interface. Useful for knowing when a dead interface failed.
output hang	Number of hours, minutes, and seconds (or never) since the interface was last reset because of a transmission that took too long. When the number of hours in any of the previous fields exceeds 24, the number of days and hours is printed. If that field overflows, asterisks are printed.

Technet24

Last clearing	Time at which the counters that measure cumulative statistics shown in this report (such as number of bytes transmitted and received) were last reset to 0. Note that variables that might affect routing (for example, load and reliability) are not cleared when the counters are cleared. Asterisks indicate elapsed time too large to be displayed. Reset the counters with the **clear interface** command.
Output queue, input queue, drops queue	Number of packets in output and input queues. Each number is followed by a slash (/), the maximum size of the queue, and the number of packets dropped because of a full queue.
Five minute input rate, Five minute output rate	Average number of bits and packets transmitted per second in the past 5 minutes. If the interface is not in promiscuous mode, it senses network traffic that it sends and receives (instead of all network traffic). The 5-minute input and output rates should be used only as an approximation of traffic per second during a given 5-minute period. These rates are exponentially weighted averages with a time constant of 5 minutes. A period of four time constants must pass before the average will be within 2% of the instantaneous rate of a uniform stream of traffic over that period.

packets input	Total number of error-free packets the system received.
bytes input	Total number of bytes, including data and MAC encapsulation, in the error-free packets received by the system.
no buffers	Number of received packets discarded because the main system had no buffer space. Compare with ignored count. Broadcast storms on Ethernet are often responsible for no input buffer events.
Received…broadcasts	Total number of broadcast or multicast packets received by the interface. The number of broadcasts should be kept as low as practicable. An approximate threshold is less than 20% of the total number of input packets.
runts	Number of Ethernet frames that are discarded because they are smaller than the minimum Ethernet frame size. Any Ethernet frame that is less than 64 bytes is considered a runt. Runts are usually caused by collisions. If more than one runt per million bytes is received, it should be investigated.
giants	Number of Ethernet frames discarded because they exceed the maximum Ethernet frame size. Any Ethernet frame that is larger than 1518 bytes is considered a giant.
input error	Runts, giants, no buffer, cyclic redundancy check (CRC), frame, overrun, and ignored counts. Other input-related errors can also increase the input error count, and some datagrams can have more than one error. Therefore, this sum might not balance with the sum of enumerated input error counts.

Technet24

CRC	CRC generated by the originating LAN station or far-end device not matching the checksum calculated from the data received. On a LAN, this usually indicates noise or transmission problems on the LAN interface or the LAN bus itself. A high number of CRCs is usually the result of collisions or a station transmitting bad data.
frame	Number of packets received as incorrectly having a CRC error and a noninteger number of octets. On a LAN, this is usually the result of collisions or a malfunctioning Ethernet device.
overrun	Number of times the receiver hardware could not hand-receive data to a hardware buffer because the input rate exceeded the capability of the receiver to handle the data.
ignored	Number of received packets ignored by the interface because the interface hardware ran low on internal buffers. These buffers are different from the system buffers mentioned in the buffer description. Broadcast storms and bursts of noise can cause the ignored count to increase.

input packets with dribble condition detected	Dribble bit error indicates that a frame is slightly too long. This frame error counter is incremented just for informational purposes; the router accepts the frame.
packets output	Total number of messages transmitted by the system.
bytes	Total number of bytes, including data and MAC encapsulation, transmitted by the system.
underruns	Number of times the transmitter has been running faster than the router can handle. This might never be reported on some interfaces.
output errors	Sum of all errors that prevented the final transmission of datagrams out the interface being examined. Note that this might not balance with the sum of the enumerated output errors because some datagrams might have more than one error and others might have errors that do not fall into any of the specifically tabulated categories.

Technet24

collisions	Number of messages retransmitted because of an Ethernet collision. This is usually the result of an overextended LAN (too-long Ethernet or transceiver cable, more than two repeaters between stations, or too many cascaded multiport transceivers). A packet that collides is counted only once in output packets.
interface resets	Number of times an interface has been completely reset. This could happen if packets queued for transmission were not sent within several seconds. On a serial line, this can be caused by a malfunctioning modem that is not supplying the transmit clock signal, or it can be caused by a cable problem. If the system notices that the carrier detect line of a serial interface is up but the line protocol is down, it periodically resets the interface in an effort to restart it. Interface resets can also occur when an interface is looped back or shut down.

Basic Router Configuration with IPv6

In this section, we use the topology shown in Figure 18-2 to review the basic commands for enabling IPv6 on a router.

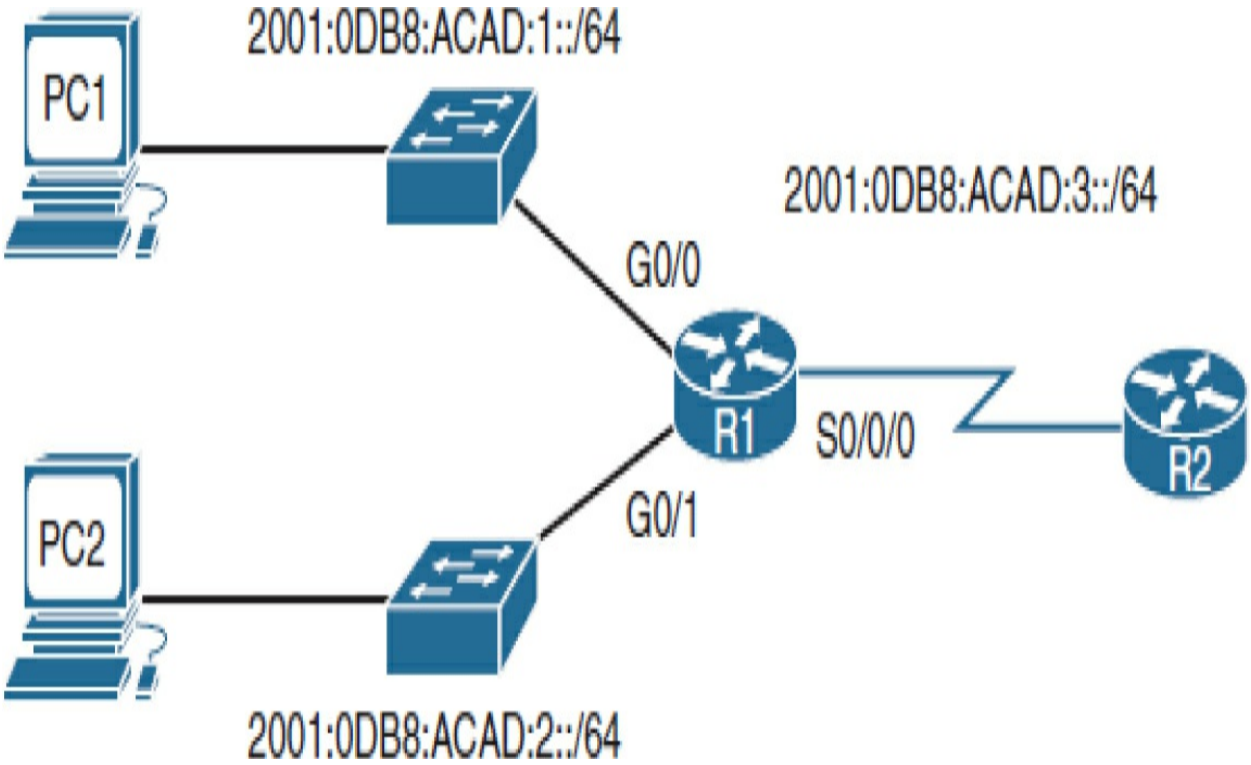

Figure 18-2 *IPv6 Sample Topology*

Command Syntax

You enable IPv6 routing by using the following command in global configuration mode:

```
R1(config)# ipv6 unicast-routing
```

Among other actions, this command configures the router to begin listening for and responding to Neighbor Discovery (ND) messages on all active IPv6 interfaces.

To configure an IPv6 address on a router's interface, you have one of several

Technet24

options:

• Configure the interface to use the EUI-64 method of addressing:

```
Router(config)# ipv6 address ipv6-prefix/prefix-length
```

• Configure the full global unicast address. To manually configure a full IPv6 address, use the following command syntax:

```
Router(config)# ipv6 address ipv6-address/prefix-lengt
```

• Configure the interface as unnumbered (see Day 27, “IPv6 Addressing”).

• Configure the interface as a DHCPv6 client (see Day 23, “DHCP and DNS”).

Note

To manually configure an interface’s link-local address, use the following command syntax:

```
Router(config)# ipv6 address ipv6-address/prefix-lengt
```

Configuration Example

The preferred IPv6 configuration method often is to manually configure the full IPv6 address because you can control the number of hexadecimal digits you must type when testing connectivity or troubleshooting a problem. You can see this by comparing the EUI-64 method to a full configuration. In Example 18-4, the interfaces on R1 are all configured using the EUI-64 method.

Example 18-4 *Configuring Interfaces Using the EUI-64 Method*

```
R1(config)# interface g0/0
R1(config-if)# ipv6 address 2001:db8:acad:1::/64 eui-6
R1(config-if)# interface g0/1
R1(config-if)# ipv6 address 2001:db8:acad:2::/64 eui-6
R1(config-if)# interface s0/0/0
R1(config-if)# ipv6 address 2001:db8:acad:3::/64 eui-6
R1(config-if)# do show ipv6 interface brief
GigabitEthernet0/0 [up/up]
    FE80::2D0:97FF:FE20:A101
    2001:DB8:ACAD:1:2D0:97FF:FE20:A101
GigabitEthernet0/1 [up/up]
    FE80::2D0:97FF:FE20:A102
    2001:DB8:ACAD:2:2D0:97FF:FE20:A102
Serial0/0/0 [down/down]
    FE80::20C:CFFF:FE77:A401
 2001:DB8:ACAD:3:20C:CFFF:FE77:A401
```

Technet24

```
<output omitted>
```

Notice the number of hexadecimal digits in the IPv6 addresses highlighted in the output from the **show ipv6 interface brief** command. Imagine having to ping the GigabitEthernet 0/0 address 2001:DB8:ACAD:1:2D0:97FF:FE20:A101.

Furthermore, notice that the link-local addresses are also rather complex. To reduce the complexity of the router's configuration, verification, and troubleshooting, it is a good practice to manually configure the link-local address as well as the IPv6 global unicast address. In Example 18-5, R1 is reconfigured with simpler IPv6 addresses and with FE80::1 as the link-local address on all interfaces. Remember that the link-local address needs to be unique only on that interface's link.

Example 18-5 *Full IPv6 Address and Link-Local Address Configuration*

```
R1(config-if)# interface g0/0
R1(config-if)# no ipv6 address 2001:db8:acad:1::/64 eu
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# interface g0/1
R1(config-if)# no ipv6 address 2001:db8:acad:2::/64 eu
R1(config-if)# ipv6 address 2001:db8:acad:2::1/64
R1(config-if)# ipv6 address fe80::1 link-local
```

```
R1(config-if)# interface s0/0/0
R1(config-if)# no ipv6 address 2001:db8:acad:3::/64 eu
R1(config-if)# ipv6 address 2001:db8:acad:3::1/64
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# do show ipv6 interface brief
GigabitEthernet0/0 [up/up]
    FE80::1
    2001:DB8:ACAD:1::1
GigabitEthernet0/1 [up/up]
    FE80::1
    2001:DB8:ACAD:2::1
Serial0/0/0 [down/down]
    FE80::1
    2001:DB8:ACAD:3::1
<output omitted>
```

Note

If you do not remove the previous IPv6 address configuration, each interface will have two IPv6 global unicast addresses. This is different than in IPv4, where simply configuring another IPv4 address with the **ip address** command overwrites any previous configuration. However, only one link-local address can exist per interface.

Compare the highlighted output from the **show ipv6 interface brief**

Technet24

command in Example 18-5 with the output in Example 18-4. You can see that simplifying the IPv6 addressing implementation can make your verification and troubleshooting job much easier.

To verify the full configuration of an interface, use the **show ipv6 interface** command. Example 18-6 shows the output for R1's GigabitEthernet 0/0 interface.

Example 18-6 *The show ipv6 interface gigabitethernet 0/0 Command*

```
R1# show ipv6 interface gigabitethernet 0/0
GigabitEthernet0/0 is up, line protocol is up
   IPv6 is enabled, link-local address is FE80::1
   No Virtual link-local address(es):
   Global unicast address(es):
   2001:DB8:ACAD:1::1, subnet is 2001:DB8:ACAD:1::/64
   Joined group address(es):
    FF02::1
    FF02::1:FF00:1
   MTU is 1500 bytes
   ICMP error messages limited to one every 100 millis
   ICMP redirects are enabled
   ICMP unreachables are sent
   ND DAD is enabled, number of DAD attempts: 1
   ND reachable time is 30000 milliseconds
   ND advertised reachable time is 0 milliseconds
   ND advertised retransmit interval is 0 milliseconds
```

```
    ND router advertisements are sent every 200 seconds
    ND router advertisements live for 1800 seconds
    ND advertised default router preference is Medium
    Hosts use stateless autoconfig for addresses.
```

Focus on the highlighted output in Example 18-6. IPv6 is enabled on this interface with a nice, short link-local address. The global unicast address and its subnet are listed, as is the address of multicast groups that this interface automatically joined. Do you remember what the FF02::1 and FF02::1:FF00:1 addresses are used for? If not, revisit Day 27.

That's all the IPv6 configurations for today. As we continue to review the exam topics in the upcoming days, we will incorporate IPv6 topics.

Verifying IPv4 and IPv6 Network Connectivity

As reviewed on Day 29, "Switch Configuration Basics," **ping** and **traceroute** are helpful tools for verifying network connectivity. Example 18-7 demonstrates successful **ping** output on the router.

Example 18-7 *Successful* ping *Output on a Router*

```
R1# ping 192.168.3.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.3.10, timeou
```

Technet24

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/
!Pinging an IPv6 destination
R1# ping 2001:db8:acad:1:290:dff:fee5:8095

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:1:290:
2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/

R1#
```

Unsuccessful **ping** output shows periods (**.**) instead of exclamation points (**!**), as Example 18-8 demonstrates. The output would be the same in IPv6.

Example 18-8 *Unsuccessful* ping *Output on a Router*

```
R1# ping 192.168.3.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.3.2, timeout
.....
Success rate is 0 percent (0/5)
R1#
```

Example 18-9 shows output from a successful **traceroute** command.

Example 18-9 *Successful* traceroute *Output on a Router*

```
R1# traceroute 192.168.3.10
Type escape sequence to abort.
Tracing the route to 192.168.3.10

  1  192.168.2.2    71 msec   70 msec  72 msec
  2  192.168.3.10  111 msec  133 msec  115 msec
R1#
!Tracing to an IPv6 destination.
R2# traceroute 2001:db8:acad:1:290:cff:fee5:8095
Type escape sequence to abort.
Tracing the route to 2001:DB8:ACAD:1:290:CFF:FEE5:8095

  1 2001:DB8:ACAD:3::1 1 msec 1 msec 1 msec
  2 2001:DB8:ACAD:1:290:CFF:FEE5:8095 1 msec 1 msec 0
R2#
```

Unsuccessful traces show the last successful hop and the asterisks for each attempt until the user cancels. To cancel the **traceroute** command on a router, use the key combination **Ctrl+Shift+6** and then press the **x** key. Example 18-10 shows unsuccessful **traceroute** output. The output would be the same with IPv6.

Technet24

Example 18-10 *Unsuccessful* traceroute *Output on a Router*

```
R1# traceroute 192.168.3.2
Type escape sequence to abort.
Tracing the route to 192.168.3.2

 1 192.168.2.2 71 msec 70 msec 72 msec
 2 *     *      *
 3 *     *      *
 4 *     *      *
 5 *
R1#
```

Using Telnet or SSH to remotely access another device also tests connectivity. More important, these remote access methods test whether a device has been correctly configured so that you can access it for management purposes. This can be important when a device is truly remote (for example, across town or in another city). Day 20, "LAN Security and Device Hardening" reviews SSH configuration and verification in greater detail.

During the basic configuration tasks earlier, you entered the commands to properly configure the vty lines for SSH remote access. If you are accessing a device configured with SSH from a PC, you use the SSH setting in your terminal client. However, you can use the **ssh** command on a router or switch to access another device configured with SSH. Example 18-11 shows how to

use SSH to remotely access R2 from R1.

Example 18-11 *Remote Access Using SSH*

```
R1# ssh ?
  -c Select encryption algorithm
  -l Log in using this user name
  -m Select HMAC algorithm
  -o Specify options
  -p Connect to this port
  -v Specify SSH Protocol Version
  -vrf Specify vrf name
   WORD IP address or hostname of a remote system

R1# ssh –l ?
  WORD Login name
R1# ssh -l admin ?
  -c Select encryption algorithm
  -m Select HMAC algorithm
  -o Specify options
  -p Connect to this port
  -v Specify SSH Protocol Version
  -vrf Specify vrf name
  WORD IP address or hostname of a remote system
R1# ssh -l admin 192.168.2.2
Password:

**************************************
```

Technet24

```
WARNING!! Unauthorized Access Prohibited!!
**************************************

R2>
```

Note

During your CCNA studies and lab practice, you most likely used a Telnet configuration to remotely access your lab equipment. Although Telnet is easier to use than SSH, remember that using SSH is considered best practice. Therefore, during the CCNA exam, be ready to use SSH to remotely access devices on simulation questions because Telnet might not be configured or allowed.

Small Office or Home Office Routers

Figure 18-3 shows the common options for small office or home office (SOHO) Internet connections.

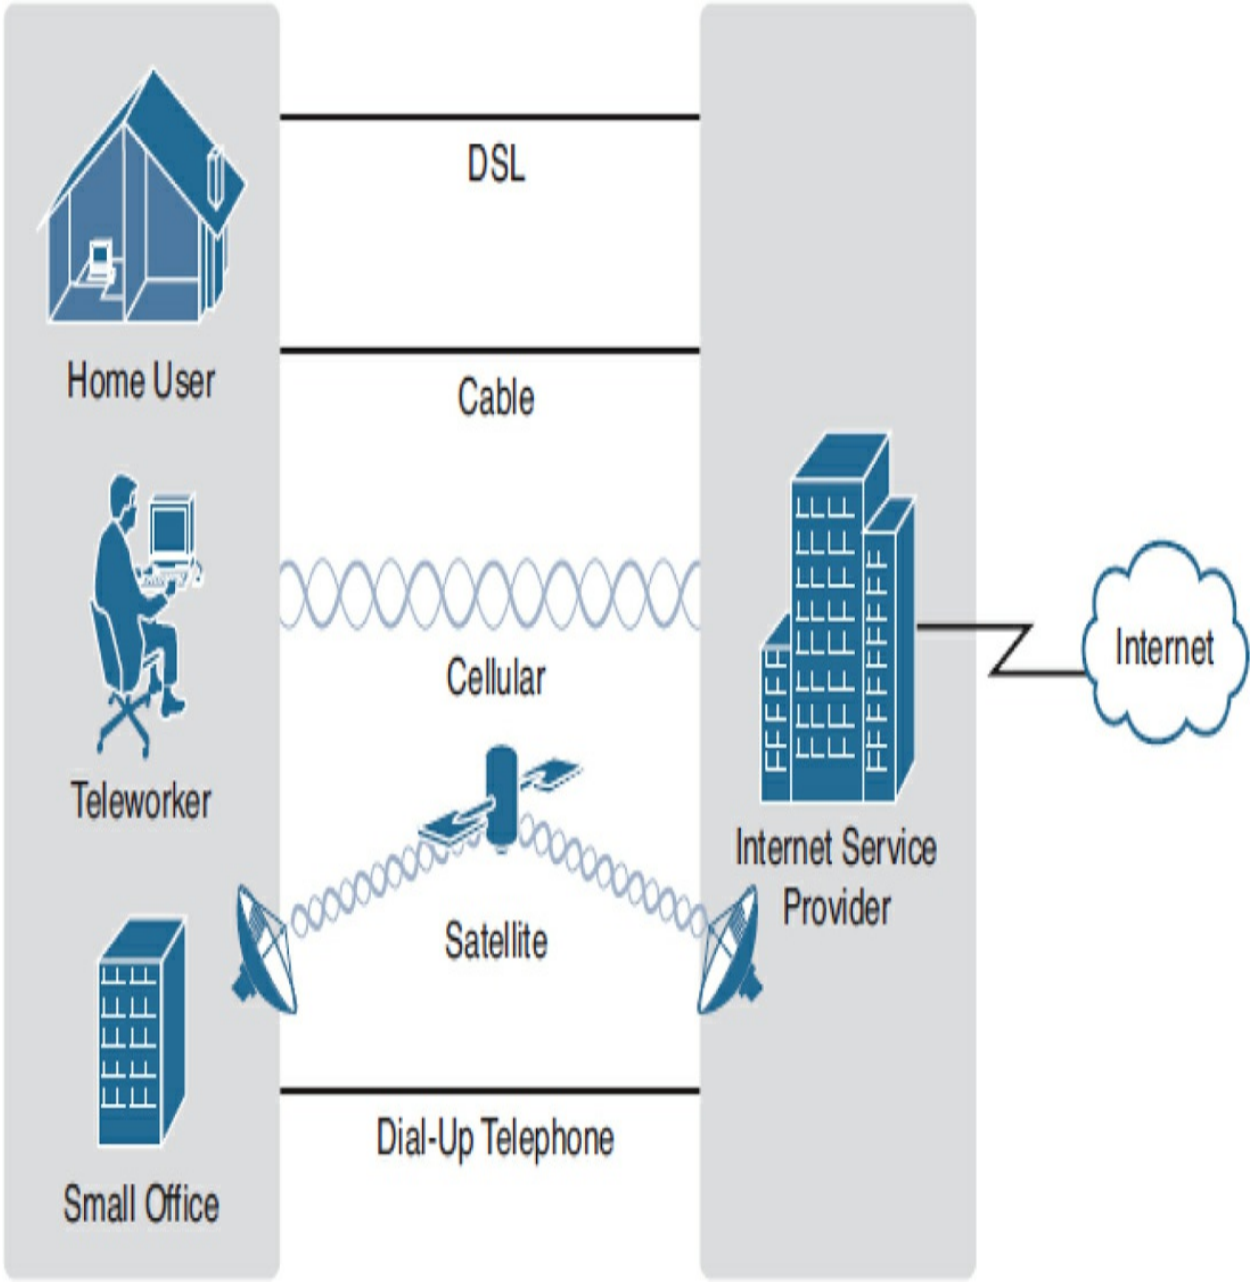

Figure 18-3 *Common SOHO Internet Connections*

The connection options shown in Figure 18-3 are as follows:

• **Cable:** Typically offered by cable television (CATV) service providers, cable transmits the Internet data signal on the same cable that delivers cable television. It provides high bandwidth, high availability, and an always-on

Technet24

connection to the Internet.

• **DSL:** Digital Subscriber Line, which runs over telephone lines, provides high bandwidth, high availability, and an always-on connection to the Internet.

• **Cellular:** Cellular Internet access uses a cell phone network to connect. Wherever you can get a cellular signal, you can get cellular Internet access. Performance is limited by the capabilities of the phone and the cell tower to which it is connected.

• **Satellite:** Satellite Internet access is used in areas that would otherwise have no Internet connectivity at all. Satellite dishes require a clear line of sight to the satellite.

• **Dial-up telephone:** Dial-up is a low-bandwidth option that uses any phone line and a modem. Dial-up is considered a legacy technology, but you might see it on the exam.

A SOHO router is typically used to create the connection to the home user and small office connections in Figure 18-4. SOHO routers typically have two features that an enterprise router would be less likely to have:

• SOHO routers almost always use the Internet and virtual private network (VPN) technology for their WAN connections to send data back and forth to the rest of the enterprise.

• A SOHO router is almost always a multifunction device that does routing, LAN switching, VPN, wireless, and maybe other features.

Figure 18-4 shows a typical SOHO site. The three icons that represent a router, a switch, and an access point actually all exist inside one box. The UTP cables are shown only to indicate that these devices are connected. The actual connection is in the hardware of the SOHO router. On the left, the SOHO router provides wired and wireless LAN servers, and on the right, it provides WAN access through a cable Internet connection.

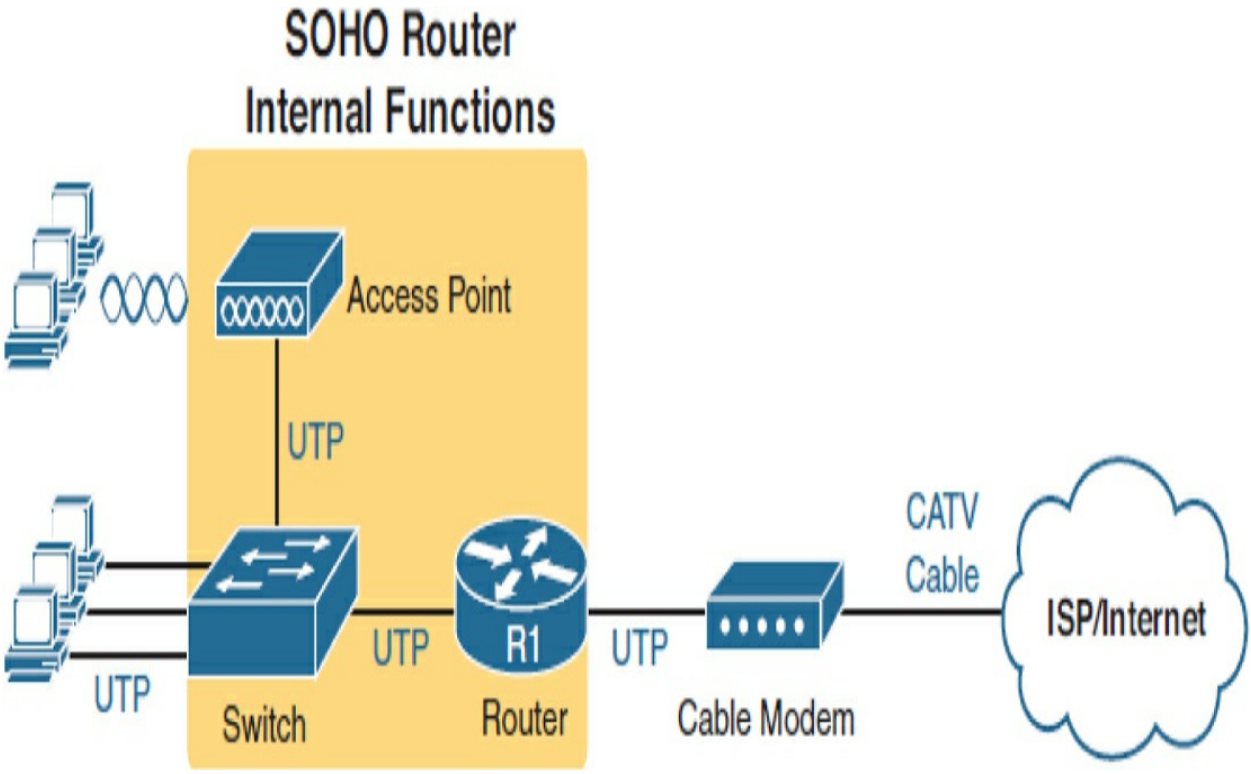

Figure 18-4 *Internal Functions SOHO Router*

Basic IP Addressing Troubleshooting

If you are sure you manually configured the correct IP address and subnet

Technet24

mask (IPv4) or network prefix (IPv6), then any basic IP addressing issue is likely to be the result of a misconfigured default gateway or duplicate addresses.

Default Gateway

A misconfigured default gateway is one of the most common problems in either a static or dynamically assigned IP addressing scheme. For a device to communicate across multiple networks, it must be configured with an IP address, a subnet mask or network prefix, and a default gateway.

The default gateway is used when the host wants to send a packet to a device on another network. The default gateway address is generally the router interface address attached to the local network to which the host is connected.

To resolve a default gateway that was manually configured incorrectly, consult the topology and addressing documentation to verify what the device's default gateway should be; it is normally a router attached to the same LAN.

Note

A misconfigured DHCP server can also cause a default gateway issue. Some DHCP server configurations, such as the Easy IP IOS feature, might require the administrator to manually configure the default gateway address. If this is

configured incorrectly, no devices will have access beyond the LAN. DHCP is reviewed on Day 23.

Duplicate IP Addresses

Under some circumstances, duplicate IP address conflicts can occur between a statically configured network device and a PC obtaining automatic IP addressing information from the DHCP server. To resolve such an IP addressing conflict, you can do one of the following:

• Convert the network device with the static IP address to a DHCP client

• On the DHCP server, exclude the static IP address of the end device from the DHCP pool of addresses

The first solution is a quick fix that you can do in the field. However, the device more than likely needs a static configuration. The second solution might be the better long-term choice. However, it requires that you have administrative privileges to configure the DHCP server.

You might also encounter IP addressing conflicts when manually configuring IP on an end device in a network that uses only static IP addresses. In this case, you must determine which IP addresses are available on the particular IP subnet and configure accordingly. This case illustrates why it is so important for a network administrator to maintain detailed documentation, including IP address assignments and topologies, for end devices.

Technet24

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	10
	17
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	1
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	16
Portable Command Guide	14

Day 17. The Routing Table

CCNA 200-301 Exam Topics

• 3.1 Interpret the components of routing table

• 3.2 Determine how a router makes a forwarding decision by default

Key Topics

Today we review the two router functions: path determination and packet forwarding. Routers use routing tables to determine the best path. A router either uses a directly connected route, a route to a remote network, or a default route. Today, we review the structure of the routing table and its entries.

Two Router Functions

When a router receives an IP packet on one interface, it determines which interface to use to forward the packet to the destination. The primary functions of a router are to

• Determine the best path for forwarding packets, based on the information in its routing table

• Forward packets toward their destinations

Technet24

Longest Match Determines Best Path

The best path in the routing table is also known as the *longest match*. The router uses the longest match process to find a match between the destination IP address of the packet and a routing entry in the routing table. The prefix length of the route in the routing table is used to determine the minimum number of far-left bits that must match. The longest match is the route in the routing table that has the greatest number of far-left bits matching the destination IP address of the packet. The route with the greatest number of equivalent far-left bits, or the longest match, is always the preferred route.

In Table 17-1, an IPv4 packet has the destination IPv4 address 172.16.0.10. The router has three route entries in its IPv4 routing table that match this packet: 172.16.0.0/12, 172.16.0.0/18, and 172.16.0.0/26. Of the three routes, 172.16.0.0/26 has the longest match and would be chosen to forward the packet.

Table 17-1 *IPv4 Address Longest Match Example*

Destination IPv4 Address		Address in Binary
172.16.0.10		10101100.00010000.00000000.00001010
Route Entry	**Prefix/Prefix Length**	**Address in Binary**
1	172.16.0.0/**12**	10101100.0001**0000.00000000.00001010**
2	172.16.0.0/**18**	10101100.00010000.00**000000.00001010**
3	172.16.0.0/**26**	10101100.00010000.00000000.00**001010**

In Table 17-2, an IPv6 packet as the destination IPv6 address 2001:db8:c000::99. This example shows three route entries, but only two of them are valid matches; one of those is the longest match. The first two route entries have prefix lengths that have the required number of matching bits, as indicated by the prefix length. The third route entry is not a match because its /64 prefix requires 64 matching bits. For the prefix 2001:db8:c000:5555::/64 to be a match, the first 64 bits must match the destination IPv6 address of the packet. Only the first 48 bits match, so this route entry is not considered a match.

Table 17-2 *IPv6 Address Longest Match Example*

Technet24

Route Entry	Prefix/Prefix Length	Does it match?
1	**2001:db8:c000::/40**	Match of 40 bits
2	2001:db8:c000::/48	Match of 48 bits (longest match)
3	**2001:db8:c000**:5555::/**64**	Does not match 64 bits

Three Packet Forwarding Decisions

After a router has determined the best path based on the longest match in the routing table, it can do one of three things:

• Forward the packet to a device on a directly connected network

• Forward the packet to a next-hop router

• Drop the packet because there is no match in the routing table

The primary responsibility of the packet forwarding function is to encapsulate packets in the appropriate data link frame type for the outgoing interface. For example, the data link frame format for a serial link could be Point-to-Point Protocol (PPP), High-Level Data Link Control (HDLC) protocol, or some other Layer 2 protocol.

Components of the Routing Table

A router examines the destination IP address of a packet and searches its routing table to determine where to forward the packet. The routing table contains a list of all known network addresses (prefixes) and where to forward the packet. These entries are known as *route entries*, or *routes*. The router forwards a packet using the best (longest) matching route entry.

Recall that a routing table stores three types of routing entries:

• **Directly connected networks:** These network route entries are active router interfaces. In Figure 17-1, the directly connected networks in the R1 IPv4 routing table are 10.0.1.0/24, 10.0.2.0/24, and 10.0.3.0/24.

• **Remote networks:** These network route entries are connected to other routers. Routers learn about remote networks either by being explicitly configured by an administrator or by exchanging route information using a dynamic routing protocol. In Figure 17-1, the remote networks in the R1 IPv4 routing table are 10.0.4.0/24 and 10.0.5.0/24.

• **Default route:** The default route is used when there is no better (longer) match in the IP routing table. In Figure 17-1, the R1 IPv4 routing table has a default route to forward all packets to R2 for any remote network for which it does not have a more explicit route.

Technet24

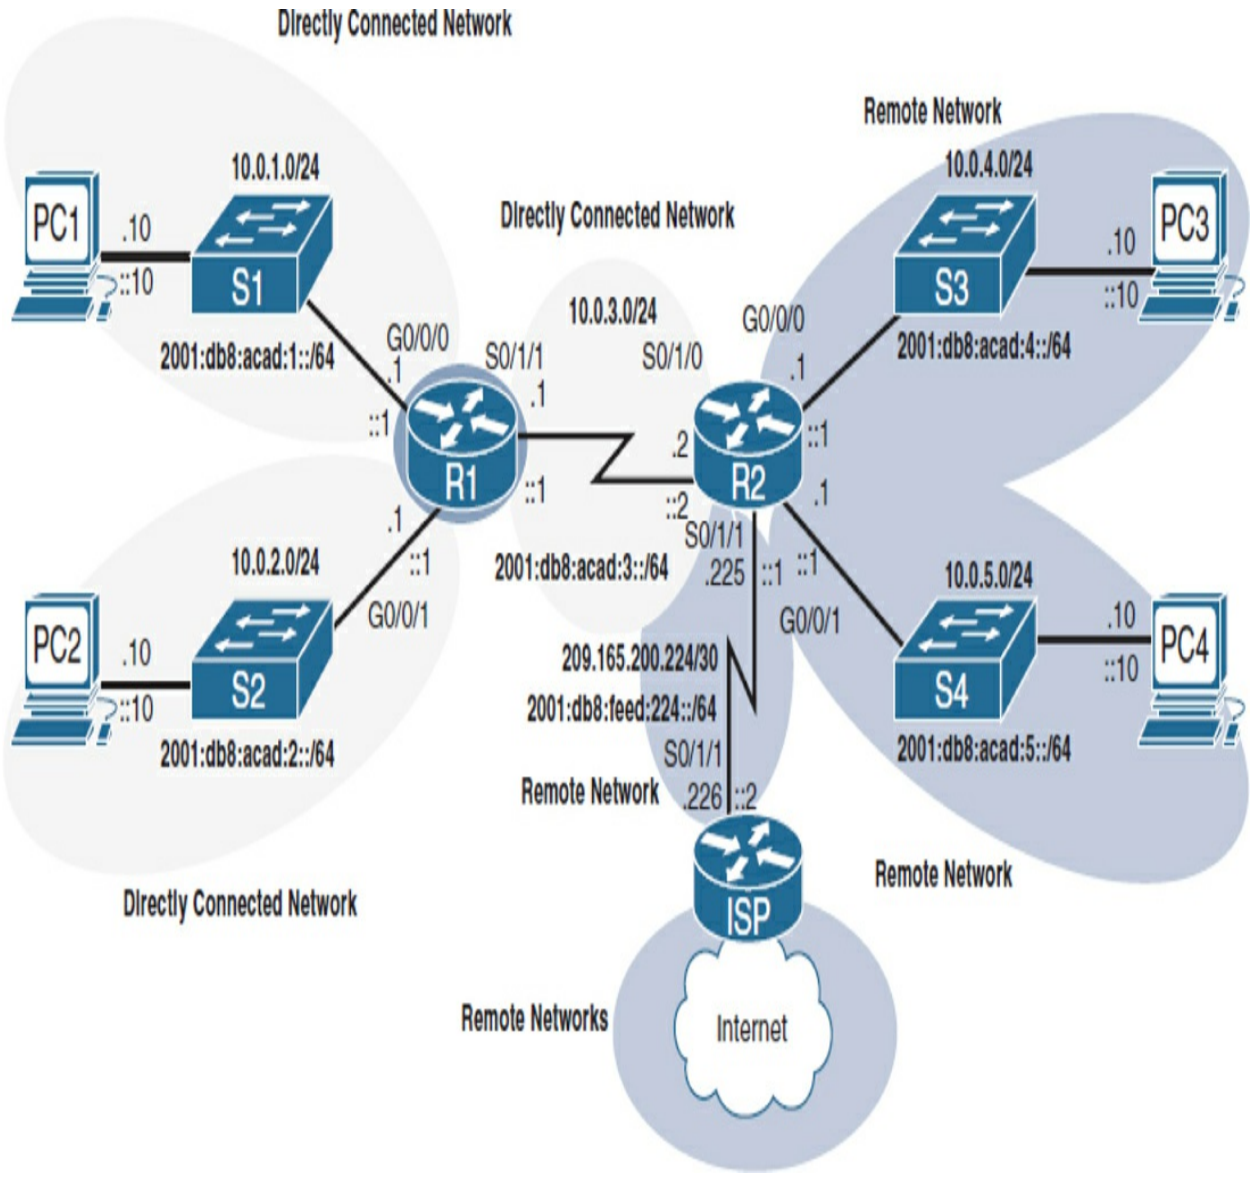

Figure 17-1 *Topology for Route Types*

In Figure 17-1, R1 and R2 are using OSPF routing to advertise directly connected networks. R2 is connected to the Internet. The administrator configured a default route on R2 and propagated it to R1 in the OSPF routing process. R1 uses this propagated default route (O*E2) to send packets to R2 when there is not a more specific entry in the routing table that matches the destination IP address. The routing table in Example 17-1 displays all the

known IPv4 destination routes for R1.

Example 17-1 *IPv4 Routing Table for R1*

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP,
       D - EIGRP, EX - EIGRP external, O - OSPF, IA -
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA
       E1 - OSPF external type 1, E2 - OSPF external t
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level
       ia - IS-IS inter area, * - candidate default, U
       o - ODR, P - periodic downloaded static route,
       a - application route
       + - replicated route, % - next hop override, p
Gateway of last resort is 209.165.200.226 to network 0
O*E2  0.0.0.0/0 [110/1] via 10.0.3.2, 00:51:34, Serial
      10.0.0.0/8 is variably subnetted, 8 subnets, 2 m
C        10.0.1.0/24 is directly connected, GigabitEth
L        10.0.1.1/32 is directly connected, GigabitEth
C        10.0.2.0/24 is directly connected, GigabitEth
L        10.0.2.1/32 is directly connected, GigabitEth
C        10.0.3.0/24 is directly connected, Serial0/1/
L        10.0.3.1/32 is directly connected, Serial0/1/
O        10.0.4.0/24 [110/50] via 10.0.3.2, 00:24:22,
O        10.0.5.0/24 [110/50] via 10.0.3.2, 00:24:15,
R1#
```

Technet24

The IPv6 routing table for R1 is shown in Example 17-2.

Example 17-2 *IPv6 Routing Table for R1*

```
R1# show ipv6 route
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP,
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OE2 ::/0 [110/1], tag 2
     via FE80::2:C, Serial0/0/1
C   2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
C   2001:DB8:ACAD:2::/64 [0/0]
     via GigabitEthernet0/0/1, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
     via GigabitEthernet0/0/1, receive
C   2001:DB8:ACAD:3::/64 [0/0]
     via Serial0/1/1, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
     via Serial0/1/1, receive
O   2001:DB8:ACAD:4::/64 [110/50]
     via FE80::2:C, Serial0/1/1
```

```
O   2001:DB8:ACAD:5::/64 [110/50]
     via FE80::2:C, Serial0/1/1
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

At the beginning of each routing table entry is a code that is used to identify the type of route or how the route was learned. Common route sources (codes) include these:

• **L:** Directly connected local interface IP address

• **C:** Directly connected network

• **S:** Static route manually configured by an administrator

• **O:** OSPF

• **D:** EIGRP

For directly connected routes, R1 adds three route entries with the codes **C** (for the connected network) and **L** (for the local interface IP address of the connected network). The route entries also identify the exit interface to use to reach the network.

R1 and R2 are also using the OSPF dynamic routing protocol to exchange

Technet24

router information. Therefore, R1 has a route entry, designated with the code **O**, for the 10.0.4.0/24 and 10.0.5.0/24 networks.

A default route has a network address of all zeros. For example, the IPv4 network address is 0.0.0.0. Instead of being statically configured, the default route was learned through OSPF and coded as O*E2 for IPv4 and OE2 for IPv6. The asterisk (*) in IPv4 means that this is a candidate for a default route. The E2 in IPv6 designates this route as an external type 2 route. In OSPF, this means the route is to another routing domain outside OSPF. In this case, the route is to the ISP router connected to the network, as shown in the topology in Figure 17-1. R2 is configured with a static default route and is propagating that route in OSPF with the **default-information originate** command configured in routing configuration mode.

Routing Table Principles

Table 17-3 describes three routing table principles. These issues are addressed by the proper configuration of dynamic routing protocols or static routes on all the routers between the source and destination devices. The examples in the table refer to the R1 and R2 in Figure 17-1.

Table 17-3 *Routing Principles and Examples*

Routing Table Principle	Example
Every router makes its decision alone, based on the information it has in its own routing table.	R1 can only forward packets using its own routing table. R1 does not know what routes are in the routing tables of other routers.
The information in the routing table of one router does not necessarily match the information in the routing table of another router.	Just because R1 has routed in its routing table to a network on the Internet through R2 does not mean that R2 knows about that same network.
Routing information about a path does not provide return routing information.	R1 receives a packet with the destination IP address of PC1 and the source IP address of PC3. Just because R1 knows to forward the packet out its G0/0/0 interface doesn't necessarily mean that R1 knows how to forward packets originating from PC1 back to the remote network of PC3.

Route Entry Structure

Figure 17-2 shows IPv4 and IPv6 routing table entries on R1 for the route to

Technet24

remote network 10.0.4.0/24 and 2001:db8:acad:4::/64. Both of these routes were learned dynamically from the OSPF routing protocol.

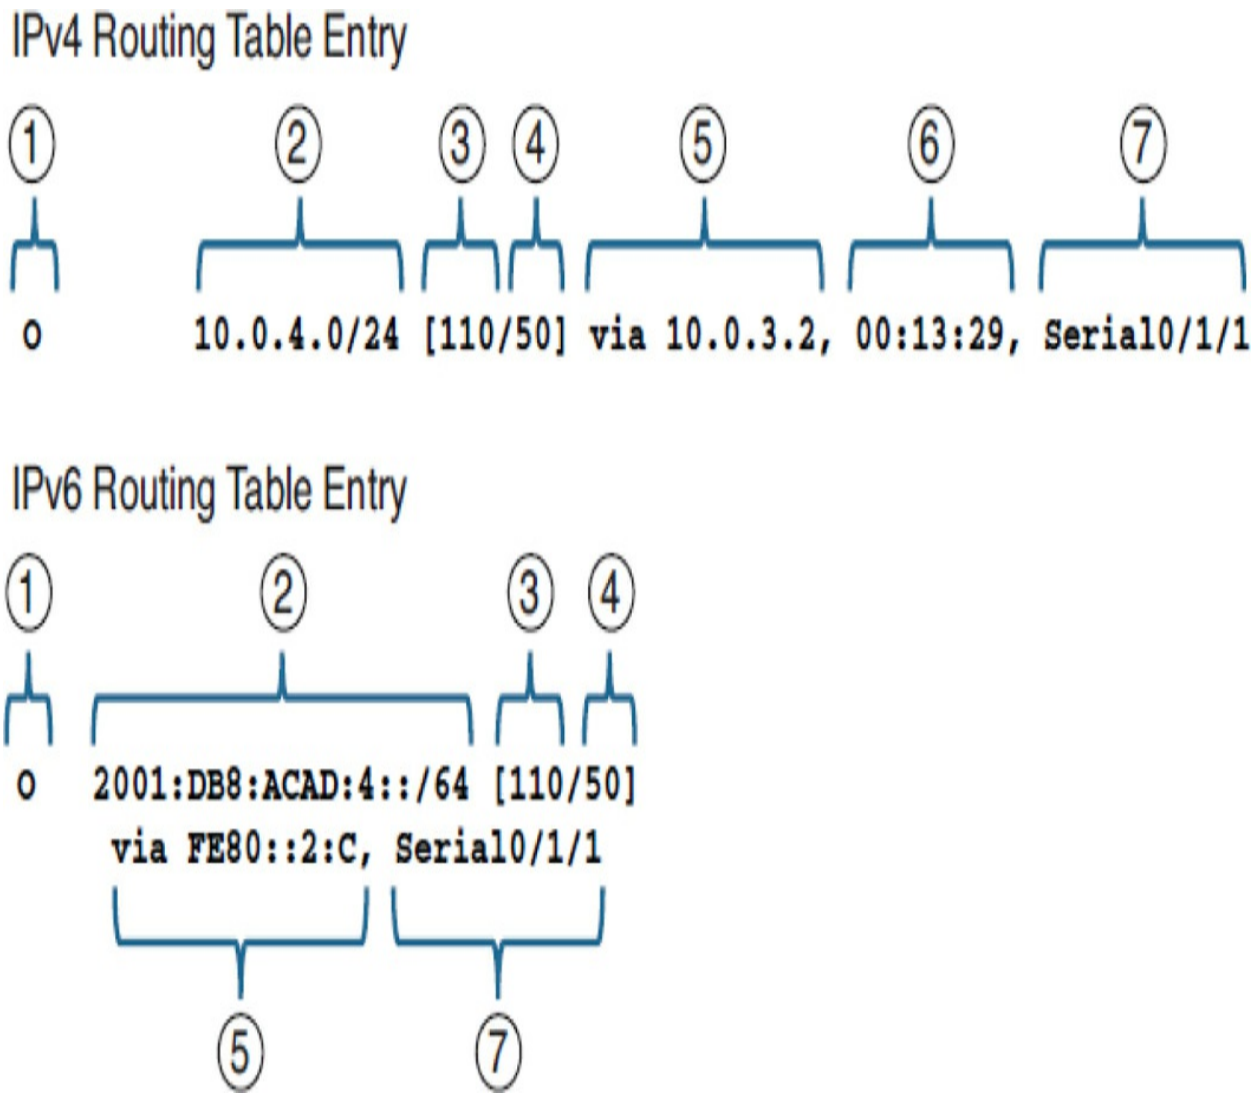

Figure 17-2 *IPv4 and IPv6 Route Entry Examples*

In the Figure 17-2, the numbers identify the following information:

1. Route source: This indicates how the route was learned.

2. Destination network (prefix and prefix length): This identifies the

address of the remote network.

3. Administrative distance: This identifies the trustworthiness of the route source. Lower values indicate preferred route sources.

4. Metric: This identifies the value assigned to reach the remote network. Lower values indicate preferred routes.

5. Next hop: This identifies the IP address of the next router to forward the packet to.

6. Route timestamp: This identifies how much time has passed since the route was learned.

7. Exit interface: This identifies the egress interface to use for outgoing packets to reach their final destination.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Technet24

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	6
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	14
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	17
	24
	29

Day 16. Inter-VLAN Routing

CCNA 200-301 Exam Topics

• 2.1 Configure and verify VLANs (normal range) spanning multiple switches

• 2.2 Configure and verify interswitch connectivity

Key Points

Today we review inter-VLAN routing. Because Layer 2 switches cannot perform the routing function, it is necessary to implement a Layer 3 device to route between VLANs.

Inter-VLAN Routing Concepts

Inter-VLAN communications cannot occur without a Layer 3 device. Three options are available when implementing inter-VLAN routing:

• Traditional or legacy inter-VLAN routing

• Router on a stick

• Multilayer switching

Let’s briefly review the concept of each method.

Technet24

Legacy Inter-VLAN Routing

Legacy inter-VLAN routing requires multiple physical interfaces on both the router and the switch. When using a router to facilitate inter-VLAN routing, the router interfaces can be connected to separate VLANs. Devices on those VLANs send traffic through the router to reach other VLANs. For example, in Figure 16-1, each S2 interface connected to R1 is assigned to a VLAN. The router is already configured with the appropriate IP addressing on each of its interfaces, so no additional configuration is required. However, you can see that if you used a separate interface for each VLAN on a router, you would quickly run out of interfaces.

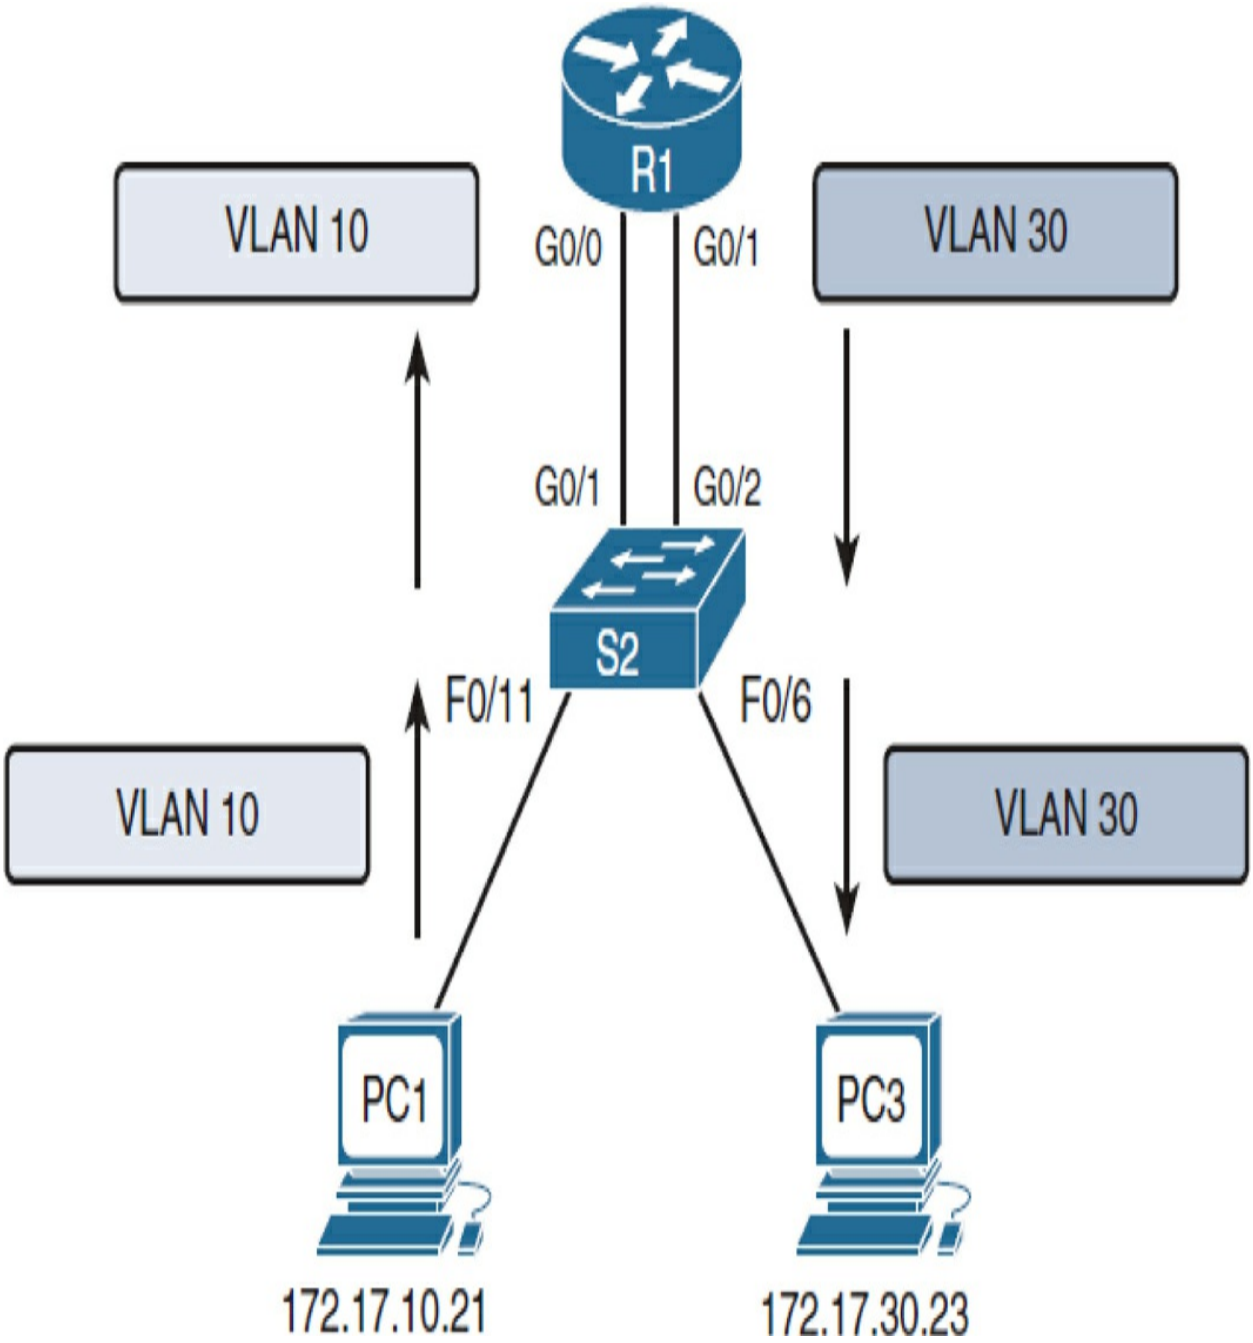

Figure 16-1 *Legacy Inter-VLAN Routing*

Router on a Stick

Today router software makes it possible to configure one router interface as multiple trunks by using subinterfaces. In Figure 16-2, the physical

Technet24

GigabitEthernet 0/0 interface is logically subdivided into two logical interfaces. The one switch trunk is configured to trunk both VLAN 10 and VLAN 30, and each subinterface on the router is assigned a separate VLAN. The router performs inter-VLAN routing by accepting VLAN-tagged traffic on the trunk interface coming from the adjacent switch. The router then forwards the routed traffic, VLAN-tagged for the destination VLAN, out the same physical interface that it used to receive the traffic.

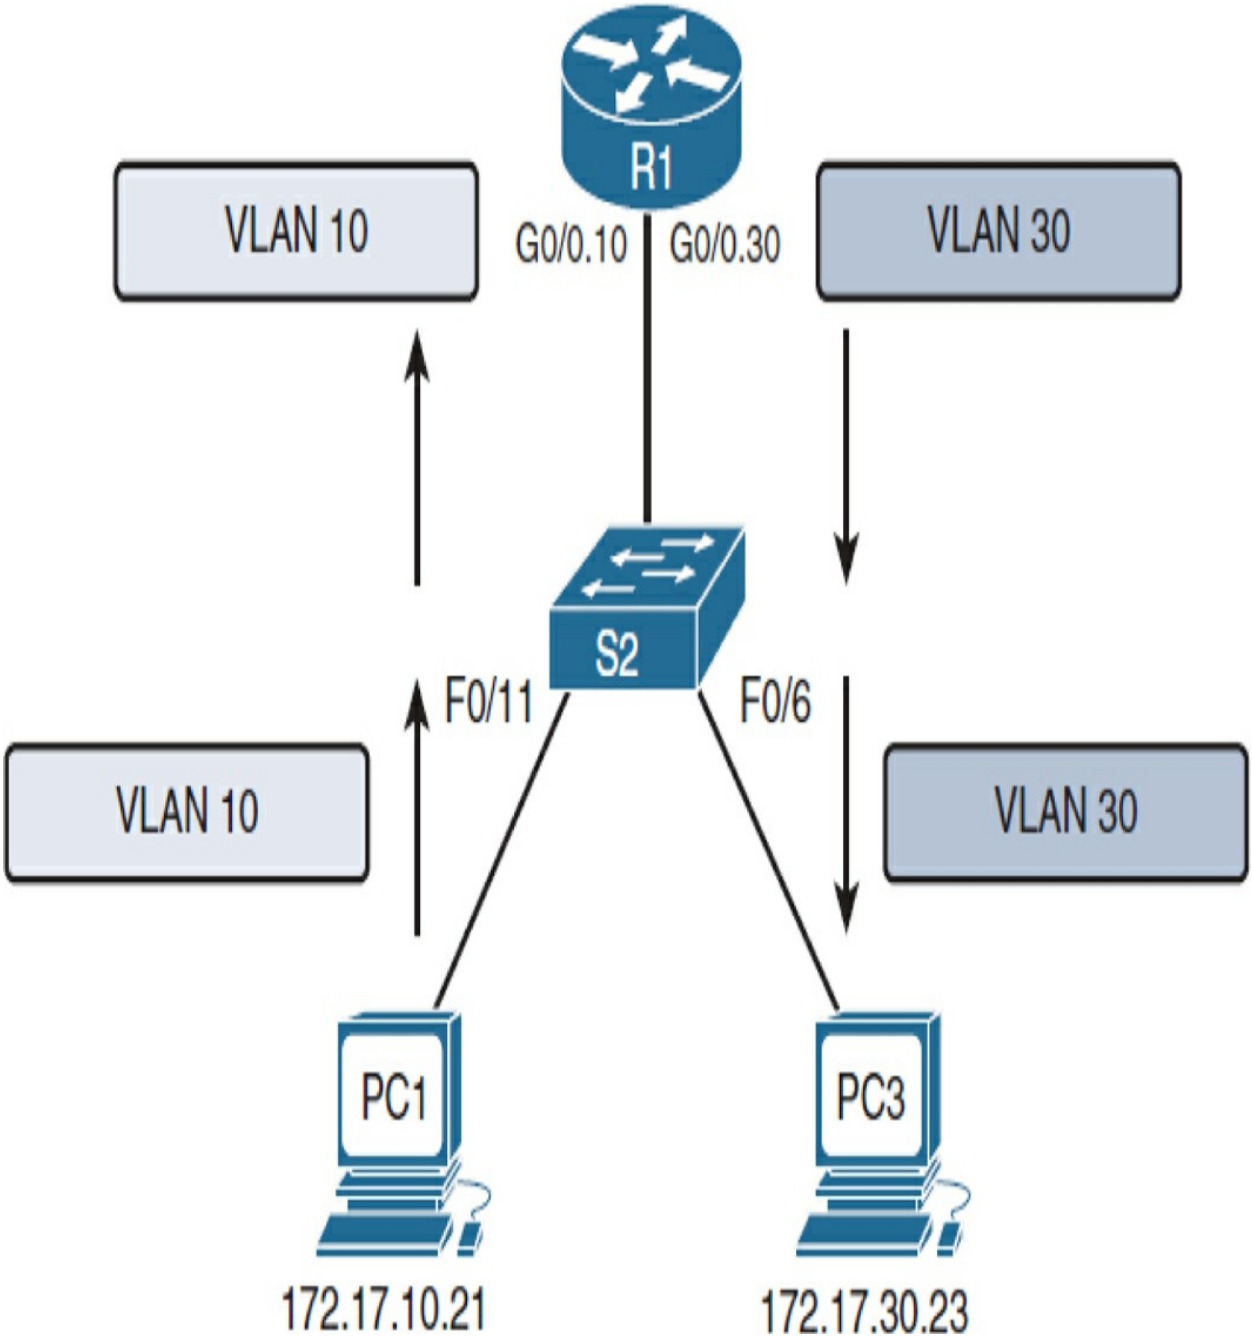

Figure 16-2 *Router on a Stick*

Multilayer Switching

Router on a stick works fine in a small business with one or two routers. But the most scalable solution in enterprise networks today is to use a multilayer

Technet24

switch to replace both the router and the switch, as in Figure 16-3. A multilayer switch performs both functions: switching traffic within the same VLAN and routing traffic between VLANs.

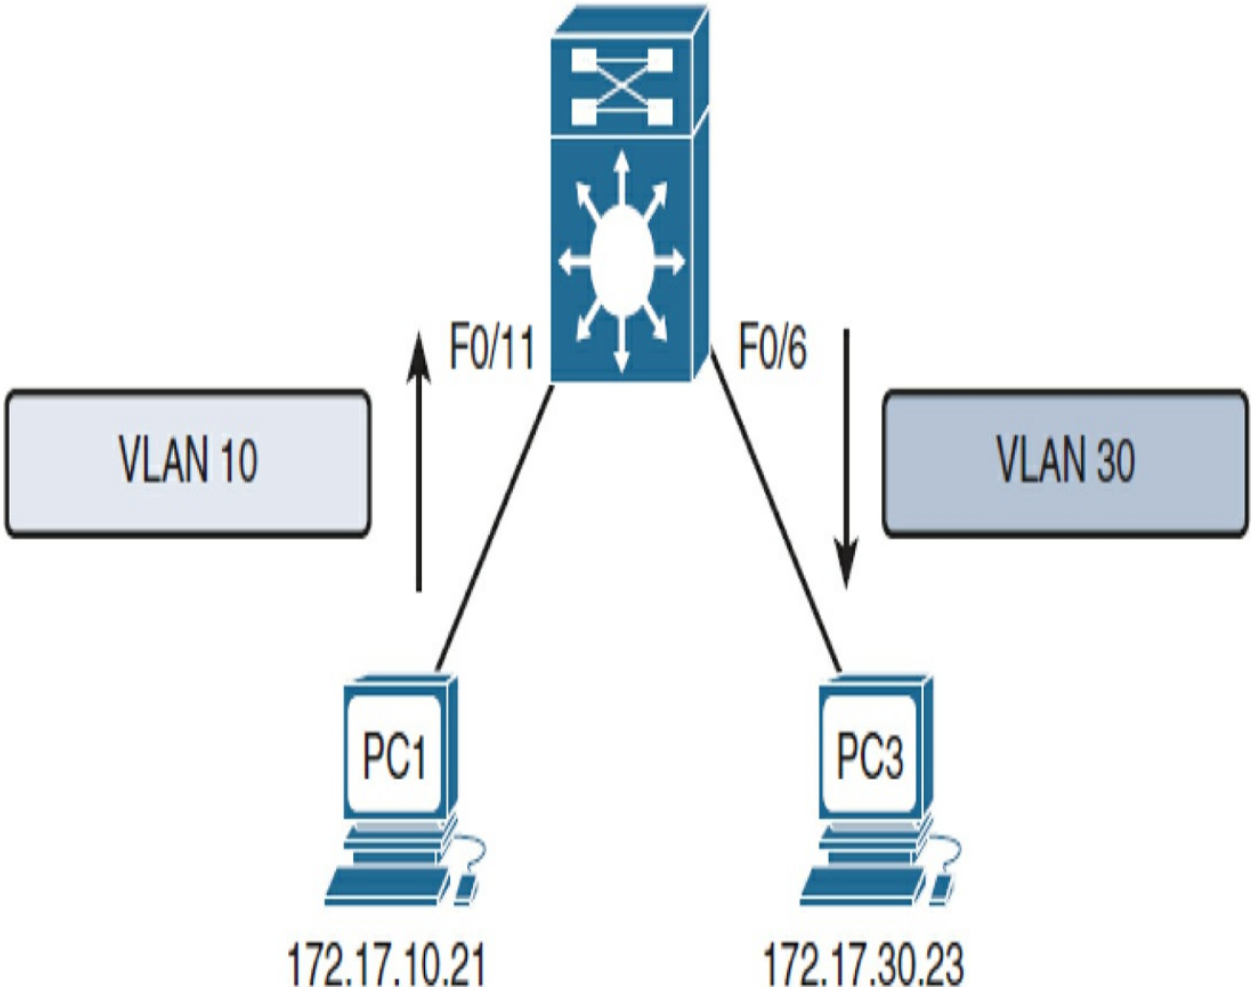

Figure 16-3 *Multilayer Switching*

Multilayer switching is more scalable than any other inter-VLAN routing implementation for two main reasons:

• Routers have a limited number of available interfaces to connect to networks.

• Limited amounts of traffic can be accommodated on the physical link at one time.

With a multilayer switch, packets are forwarded down a single trunk line to obtain new VLAN tagging information. A multilayer switch does not completely replace the functionality of a router but can be thought of as a Layer 2 device that is upgraded to have some routing capabilities.

Router on a Stick Configuration and Verification

When configuring inter-VLAN routing using the router on a stick model, the physical interface of the router must be connected to a trunk link on the adjacent switch. On the router, subinterfaces are created for each unique VLAN on the network. Each subinterface is assigned an IP address specific to its subnet/VLAN and is also configured to tag frames for that VLAN. This way, the router can keep the traffic from the different subinterfaces separated as it traverses the trunk link back to the switch.

Configuring inter-VLAN routing is pretty straightforward. Refer to the sample topology in Figure 16-4 to review the commands.

Technet24

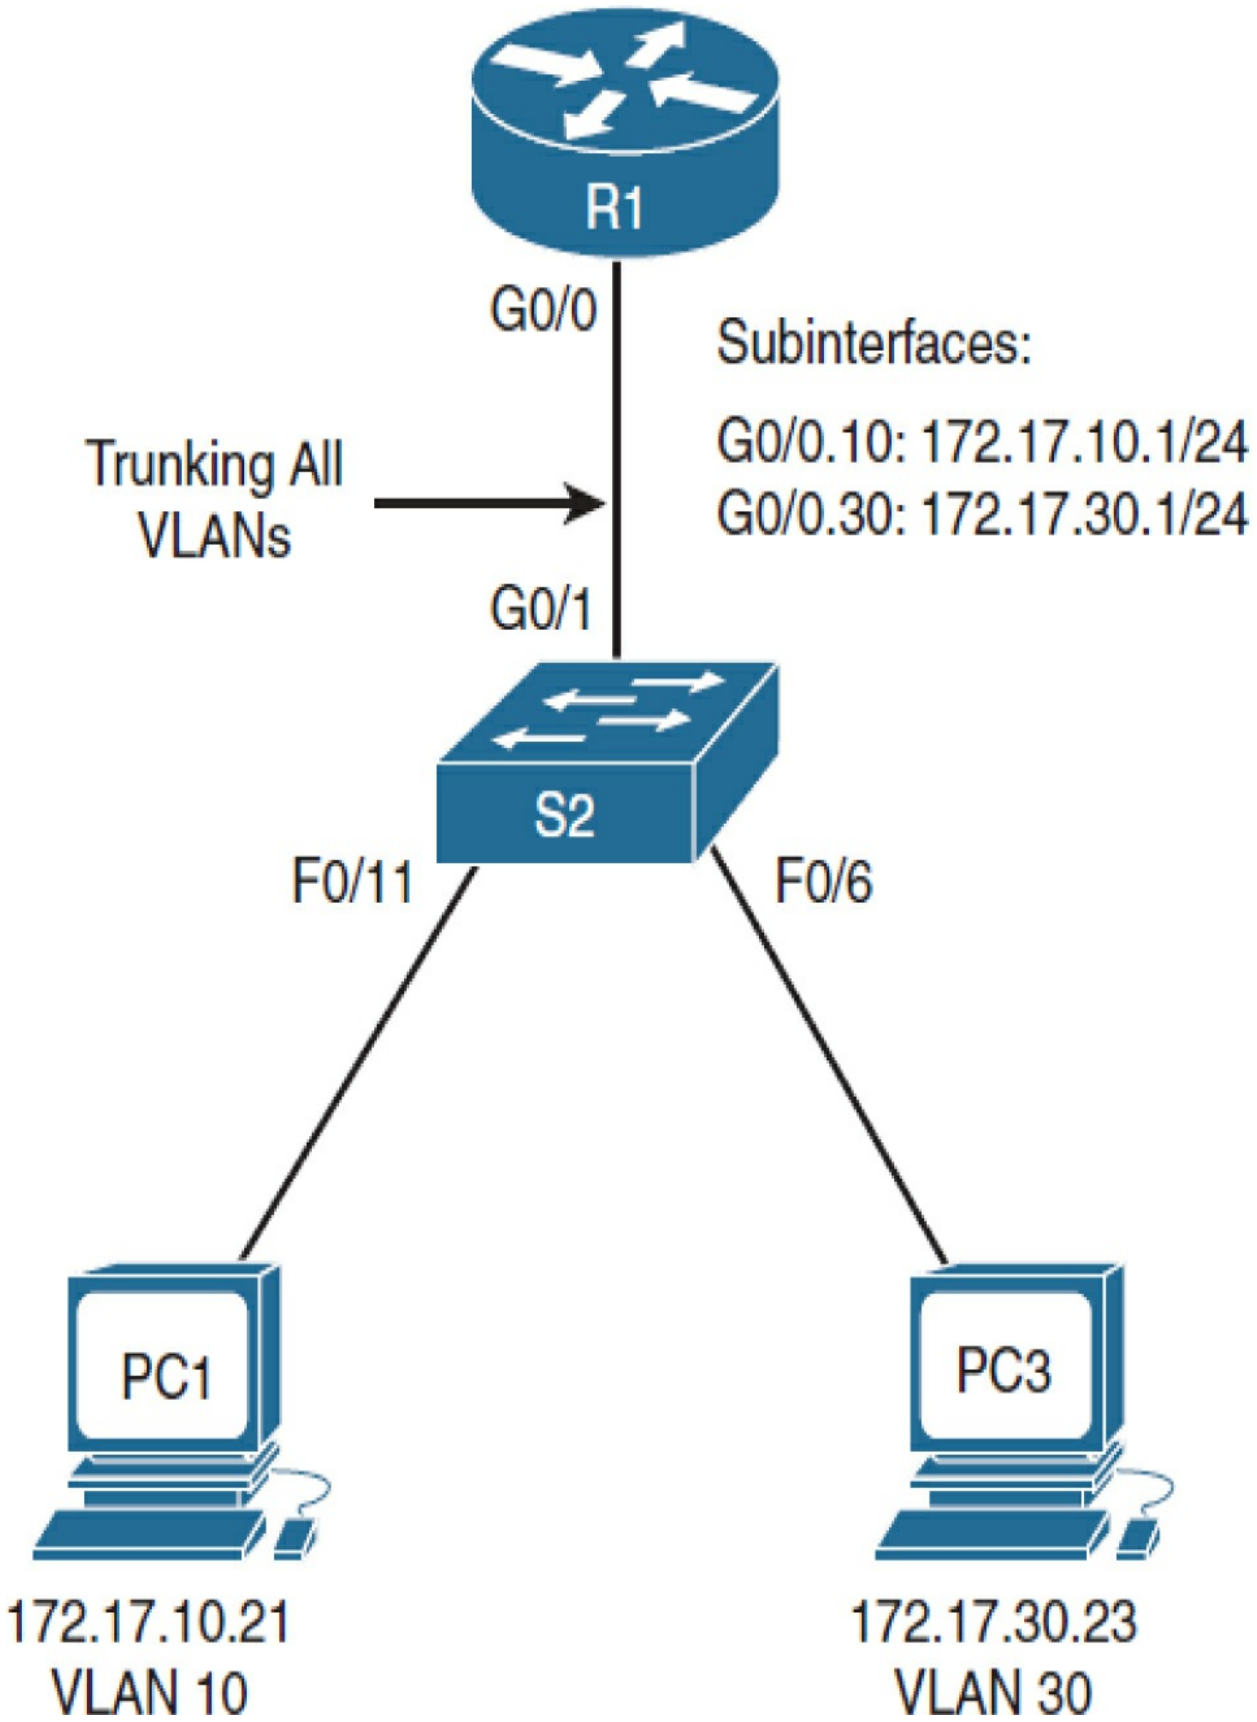
R1
G0/0
Subinterfaces:
G0/0.10: 172.17.10.1/24
G0/0.30: 172.17.30.1/24
Trunking All
VLANs
G0/1
S2
F0/11
F0/6
PC1
PC3
172.17.10.21
VLAN 10
172.17.30.23
VLAN 30

Figure 16-4 *Topology for Inter-VLAN Routing*

This router on a stick topology is configured using the following steps on the router:

Activate the physical interface that is trunking with the switch by using the **no shutdown** command.

Enter subinterface configuration mode for the first VLAN that needs routing. One convention is to use the VLAN number as the subinterface number. For example, the **interface g0/1.10** command enters subinterface configuration mode for VLAN 10.

Configure the trunking encapsulation type by using the subinterface configuration command **encapsulation** {**dot1q** | **isl**} *vlan-number* [**native**]. Set the encapsulation to **dot1q**.

▪ Inter-switch link (ISL) encapsulation, a Cisco proprietary trunking method, existed before the IEEE 802.1Q standard, which is now the recommended best practice. However, older switches that are still in use might support only ISL. In those cases, you substitute the **dot1q** keyword for **isl**.

▪ On some routers, the optional keyword **native** must be configured for the native VLAN before the router will route native VLAN traffic. The following examples do not use native VLAN routing; refer to your study resources for more on this topic.

Technet24

Configure the IP address and subnet mask.

Repeat steps 2–4 for each additional VLAN that needs routing.

Assuming that the switch is already configured with VLANs and trunking, Example 16-1 shows the commands to configure R1 to provide routing between VLAN 10 and VLAN 30.

Example 16-1 *Configuring R1 to Route Between VLANs*

```
R1(config)# interface g0/0
R1(config-if)# no shutdown
R1(config-if)# interface g0/0.10
R1(config-subif)# encapsulation dot1q 10
R1(config-subif)# ip address 172.17.10.1 255.255.255.0
R1(config-subif)# interface g0/0.30
R1(config-subif)# encapsulation dot1q 30
R1(config-subif)# ip address 172.17.30.1 255.255.255.0
```

To verify the configuration, use the **show vlans**, **show ip route**, and **show ip interface brief** commands to make sure the new networks are in the routing table and the subinterfaces are up and up, as in Example 16-2.

Example 16-2 *Verifying the Inter-VLAN Routing Configuration*

```
R1# show vlans
```

```
<output omitted>
Virtual LAN ID: 10 (IEEE 802.1Q Encapsulation)
   vLAN Trunk Interface: GigabitEthernet0/0.10
   Protocols Configured: Address:      Received:  Tran
          IP             172.17.10.1      0
<output omitted>
Virtual LAN ID: 30 (IEEE 802.1Q Encapsulation)
   vLAN Trunk Interface: GigabitEthernet0/0.30
   Protocols Configured: Address:      Received:  Tran
            IP           172.17.30.1        0
<output omitted>
R1# show ip route
<output omitted>

Gateway of last resort is not set

   172.17.0.0/16 is variably subnetted, 4 subnets, 2 m
C    172.17.10.0/24 is directly connected, GigabitEthe
L    172.17.10.1/32 is directly connected, GigabitEthe
C    172.17.30.0/24 is directly connected, GigabitEthe
L    172.17.30.1/32 is directly connected, GigabitEthe
R1# show ip interface brief
Interface              IP-Address   OK?  Method    Stat
GigabitEthernet0/0     unassigned  YES   unset     up
GigabitEthernet0/0.10  172.17.10.1  YES  manual    up
GigabitEthernet0/0.30  172.17.30.1  YES  manual    up
GigabitEthernet0/1      unassigned  YES  unset     admi
Serial0/0/0            unassigned  YES   manual    admi
```

Technet24

```
Serial0/0/1            unassigned  YES   manual   admi
Vlan1                  unassigned  YES   manual   admi
R1#
```

Assuming that the switch and PCs are configured correctly, the two PCs should now be able to ping each other. R1 should route the traffic between VLAN 10 and VLAN 30.

Multilayer Switching Inter-VLAN Routing Configuration and Verification

Most enterprise networks use multilayer switches to achieve high-packet processing rates using hardware-based switching. All Catalyst multilayer switches support the following types of Layer 3 interfaces:

• **Switch virtual interface (SVI):** Virtual VLAN interface used for inter-VLAN routing

• **Routed port:** Similar to a physical interface on a Cisco IOS router

All Layer 3 Cisco Catalyst switches (3500, 4500, and 6500 Series) support routing protocols. Catalyst 2960 Series switches running Cisco IOS Release 12.2(55) or later support static routing.

Creating Additional SVIs

The SVI for the default VLAN (VLAN 1) already exists to permit remote switch administration. For a topology such as the one in Figure 16-5, additional SVIs must be explicitly created.

Technet24

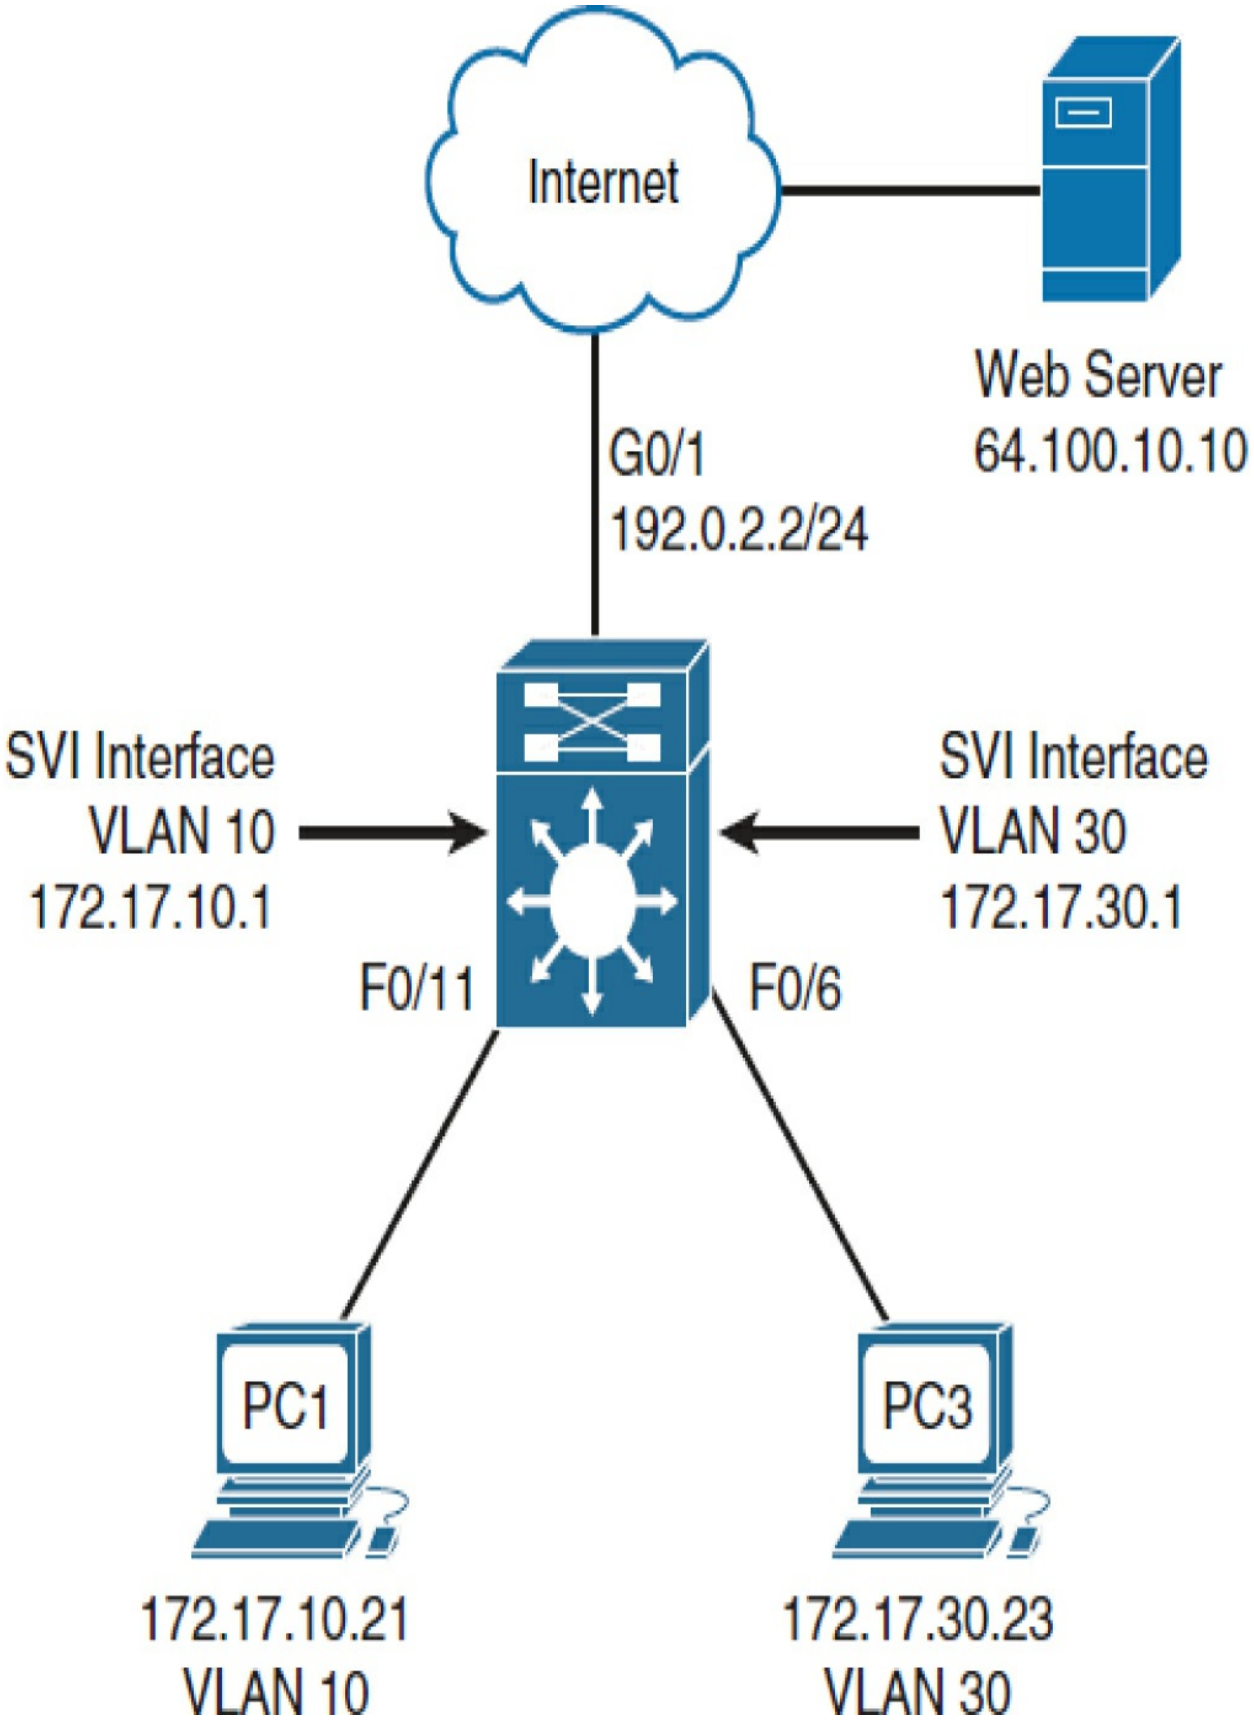
Internet
Web Server
64.100.10.10
G0/1
192.0.2.2/24
SVI Interface
VLAN 10
172.17.10.1
SVI Interface
VLAN 30
172.17.30.1
F0/11
F0/6
PC1
PC3
172.17.10.21
VLAN 10
172.17.30.23
VLAN 30

Figure 16-5 *Switched Virtual Interfaces*

Create an SVI by using the **interface vlan** *vlan-id* command. The *vlan-id* used corresponds to the VLAN tag associated with data frames coming from that VLAN. For example, when creating an SVI as a gateway for VLAN 10, use the **interface VLAN 10** command. Assign an IP address and enable the new SVI with the **no shutdown** command.

In addition, the switch must be configured to do Layer 3 routing with the **ip routing** global configuration command.

The following are some advantages of SVIs (and the only disadvantage is that multilayer switches are more expensive):

• They are much faster than routers on a stick because everything is hardware switched and routed.

• No external links are needed from the switch to the router for routing.

• They are not limited to one link. Layer 2 EtherChannels can be used between the switches to get more bandwidth.

• Latency is much lower because it does not need to leave the switch.

Example 16-3 shows the configuration for the Layer 3 switch in Figure 16-5.

Example 16-3 *Configuring a Switch to Use SVIs for Routing*

Technet24

```
MLS(config)# ip routing
MLS(config)# vlan 10
MLS(config)# vlan 30
MLS(config)# interface vlan 10
MLS(config-if)# ip address 172.17.10.1 255.255.255.0
MLS(config-if)# interface vlan 30
MLS(config-if)# ip address 172.30.10.1 255.255.255.0
MLS(config-if)# interface f0/11
MLS(config-if)# switchport mode access
MLS(config-if)# switchport access vlan 10
MLS(config-if)# interface f0/6
MLS(config-if)# switchport mode access
MLS(config-if)# switchport access vlan 30
MLS(config-if)# end
MLS# show ip route
<Code output omitted>

Gateway of last resort is not set

    172.17.0.0/24 is subnetted, 2 subnets
C   172.17.10.0 is directly connected, Vlan10
C   172.17.30.0 is directly connected, Vlan30
MLS#
```

Because the command **ip routing** is configured, MLS has a routing table.

PC1 and PC3 can now ping each other.

Configuring a Layer 3 Routed Port

In Figure 16-5, notice that the GigabitEthernet 0/1 interface has an IP address assigned to it. To configure this interface as a routed port, turn off switching with the **no switchport** interface configuration command. Then configure the IP address as normal. The **ip routing** command was enabled in the previous step. However, the Layer 3 switch still needs a default route to send traffic to the Internet. Example 16-4 shows the commands to configure the routed port and default route.

Example 16-4 *Configuring a Switch with a Routed Port*

```
MLS(config)# interface g0/1
MLS(config-if)# no switchport
MLS(config-if)# ip address 192.0.2.2 255.255.255.0
!The no shutdown command is not required because switc
MLS(config-if)# exit
MLS(config)# ip route 0.0.0.0 0.0.0.0 g0/1
MLS(config)# exit
MLS# show ip route
<Code output omitted>

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

      172.17.0.0/24 is subnetted, 2 subnets
C     172.17.10.0 is directly connected, Vlan10
C     172.17.30.0 is directly connected, Vlan30
```

Technet24

```
C     192.0.2.0/24 is directly connected, GigabitEther
S*    0.0.0.0/0 is directly connected, GigabitEthernet
MLS#
```

The **show ip route** command verifies that the Layer 3 switch has a route to the Internet. PC1 and PC3 can now access the web server.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	4
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	18
Portable Command Guide	10

Day 15. Static and Default Route Configuration

CCNA 200-301 Exam Topics

• 3.3 Configure and verify IPv4 and IPv6 static routing

Key Topics

Today we focus on static and default routing for IPv4 and IPv6. Static routes are a common part of an enterprise's routing policy. Static routes can be used to force traffic to use a specific path or to establish a default route out of the enterprise. The network administrator hard-codes static routes into the routing table. Thus, a network administrator must monitor and maintain static routes to ensure connectivity.

Static and Default Routing Overview

When a router configured with a dynamic routing protocol can learn routes from other routers without additional input from the network administrator, why would you use static routing? Situations vary, and other reasons might be unique to a particular implementation, but, in general, you use static routing in these cases:

• In a small network that requires only simple routing

• In a hub-and-spoke network topology

Technet24

• When you want to create a quick ad hoc route

• As a backup when the primary route fails

In general, you do not use static routes in these cases:

• In a large network

• When the network is expected to scale

Static routes are commonly used when you are routing from a larger network to a stub network (a network that is accessed by a single link). Static routes can also be useful for specifying a default route or gateway of last resort. For example, in Figure 15-1, R2 is attached to a stub network.

In Figure 15-1, no other route out of the stub network exists except to send packets to HQ. Therefore, it makes sense to configure R2 with a default route pointing out the interface attached to HQ. Similarly, HQ has only one way to route packets destined for the stub network attached to R2. Therefore, it makes sense to configure HQ with a static route pointing out the interface attached to R2. Yes, you could configure both routers with a dynamic routing protocol, but that could introduce a level of complexity that might not be necessary in a stub network situation.

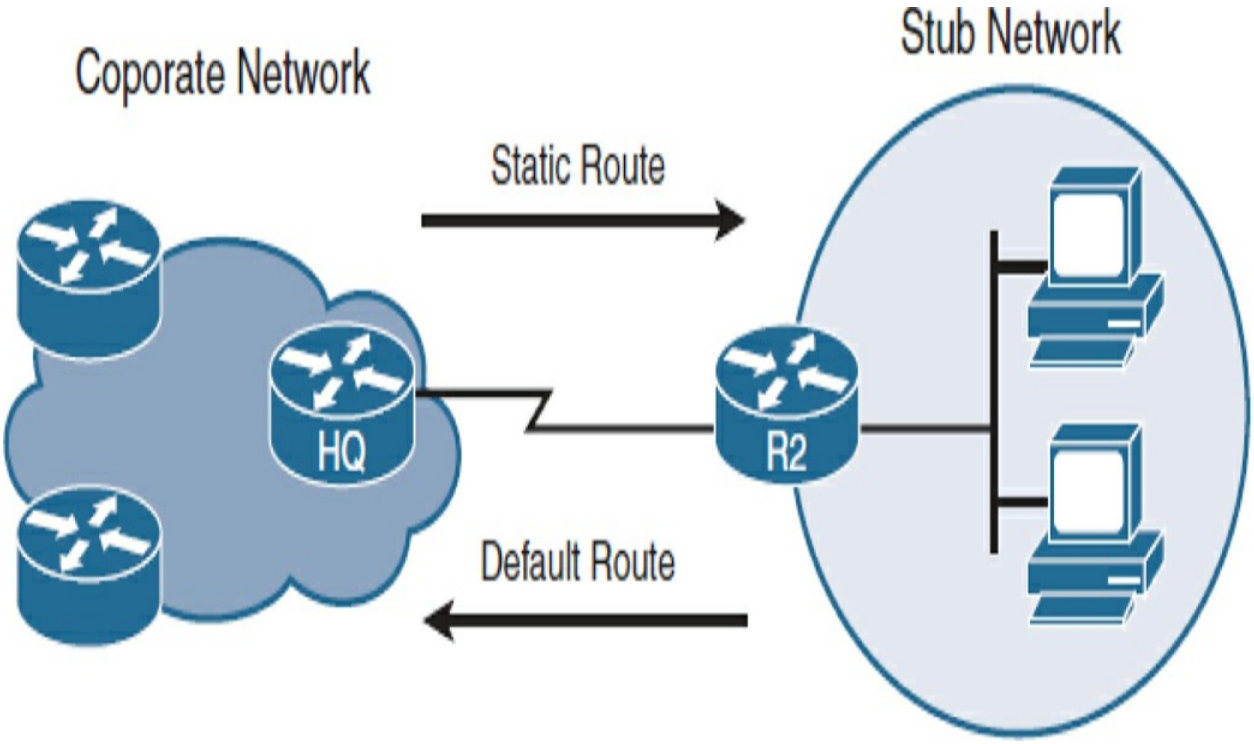

Figure 15-1 *Example of a Stub Network*

IPv4 Static Route Configuration

To configure a static route, use the **ip route** command with the following relevant syntax:

```
Router(config)# ip route network-address subnet-mask {
exit-interface} [administrative-distance]
```

The explanation for each parameter follows:

• **network-address:** The destination network address of the remote network to be added to the routing table.

Technet24

• **subnet-mask:** The subnet mask of the remote network to be added to the routing table. The subnet mask can be modified to summarize a group of networks.

One or both of the following parameters are used:

• **ip-address:** Commonly referred to as the next-hop router's IP address

• **exit-interface:** The outgoing interface used in forwarding packets to the destination network

In addition, the optional *administrative-distance* parameter is used when configuring a floating static route, as you see later in today's review.

Figure 15-2 shows the topology we use today in reviewing IPv4 static and default routing.

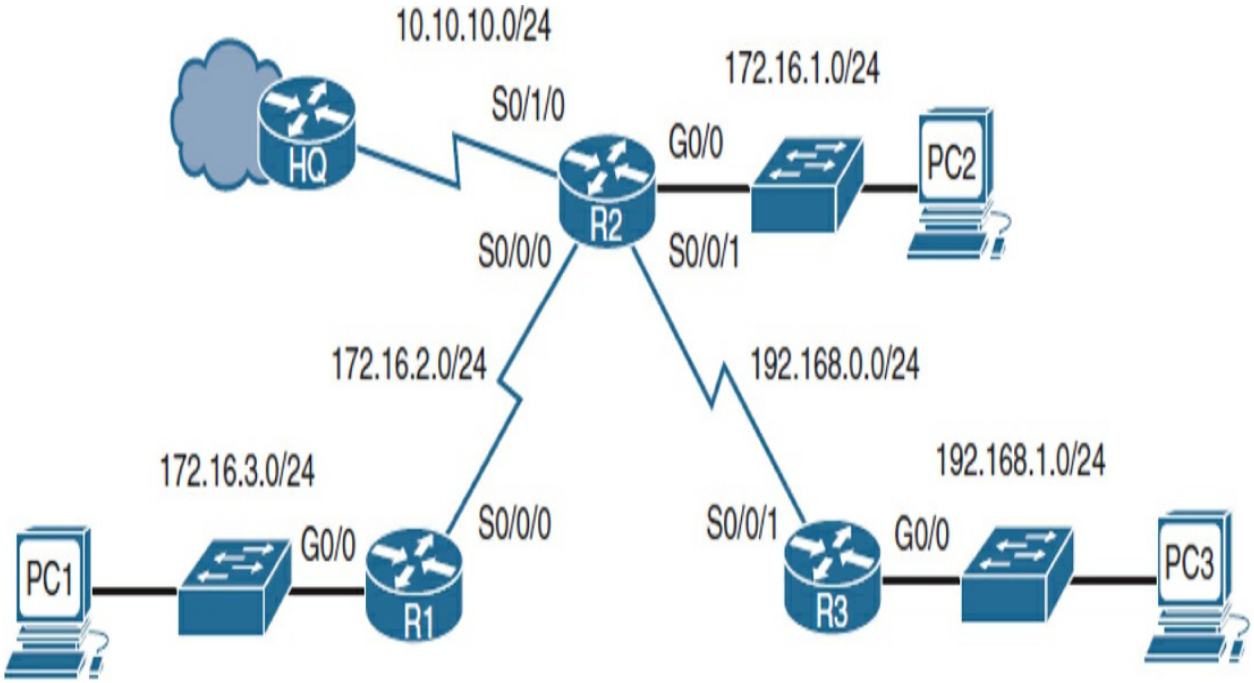

Figure 15-2 *IPv4 Static and Default Routing Topology*

Table 15-1 shows the IPv4 addressing scheme used with the topology in Figure 15-2.

Table 15-1 *IPv4 Addressing Scheme*

Technet24

Device	Interface	IP Address	Subnet Mask	Default Gateway
HQ	S0/0/0	10.10.10.1	255.255.255.0	—
R1	G0/0 S0/0/0	172.16.3.1 172.16.2.2	255.255.255.0 255.255.255.0	— —
R2	G0/0 S0/0/0 S0/0/1 S0/1/0	172.16.1.1 172.16.2.1 192.168.0.1 10.10.10.2	255.255.255.0 255.255.255.0 255.255.255.0 255.255.255.0	— — — —
R3	G0/0 S0/0/1	192.168.1.1 192.168.0.2	255.255.255.0 255.255.255.0	— —
PC1	NIC	172.16.3.10	255.255.255.0	172.16.3.1
PC2	NIC	172.16.1.10	255.255.255.0	172.16.1.1
PC3	NIC	192.168.2.10	255.255.255.0	192.168.2.1

Assume that R1 is configured and knows about its own directly connected

networks. Example 15-1 shows the routing table for R1 before any static routing is configured.

Example 15-1 *R1 Routing Table Before Static Routes Are Configured*

```
R1# show ip route
<output omitted>

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 4 subnets, 2
C       172.16.2.0/24 is directly connected, Serial0/0/0
L       172.16.2.2/32 is directly connected, Serial0/0/0
C       172.16.3.0/24 is directly connected, GigabitEthe
L       172.16.3.1/32 is directly connected, GigabitEthe
R1#
```

R1 does not know about these remote networks:

- **172.16.1.0/24:** The LAN on R2
- **192.168.0.0/24:** The serial network between R2 and R3
- **192.168.1.0/24:** The LAN on R3
- **10.10.10.0/24:** The serial network between R2 and HQ

Technet24

• **0.0.0.0/0:** All other networks accessible through HQ

IPv4 Static Routes Using the Next-Hop Parameter

Using the next-hop parameter, R1 can be configured with three static routes —one for each network R1 does not yet know about. Example 15-2 shows the command syntax.

Example 15-2 *Static Route Configuration with the Next-Hop Parameter*

```
R1(config)# ip route 172.16.1.0 255.255.255.0 172.16.2
R1(config)# ip route 192.168.0.0 255.255.255.0 172.16.
R1(config)# ip route 192.168.1.0 255.255.255.0 172.16.
R1(config)# ip route 10.10.10.0 255.255.255.0 172.16.2
```

The interface that routes to the next hop must be up and up before the static routes can be entered in the routing table. Example 15-3 verifies that the static routes are now in the routing table.

Example 15-3 *R1 Routing Table After Static Routes Are Configured*

```
R1# show ip route
<output omitted>

Gateway of last resort is not set
```

```
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0/24 [1/0] via 172.16.2.1
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 m
S        172.16.1.0/24 [1/0] via 172.16.2.1
C        172.16.2.0/24 is directly connected, Serial0/0/0
L        172.16.2.2/32 is directly connected, Serial0/0/0
C        172.16.3.0/24 is directly connected, GigabitEthe
L        172.16.3.1/32 is directly connected, GigabitEthe
S     192.168.0.0/24 [1/0] via 172.16.2.1
S     192.168.1.0/24 [1/0] via 172.16.2.1
R1#
```

When using the next-hop parameter, the router must have a route in the table to the network that the next-hop address belongs to. In the highlighted line in Example 15-3, we see that R1 does indeed have a route to the 172.16.2.0/24 network, which includes the next-hop address 172.16.2.1. However, configuring a next-hop address requires the router to perform a recursive lookup to find the exit interface before it can send the packet out the Serial 0/0/0 interface.

IPv4 Static Routes Using the Exit Interface Parameter

To avoid a recursive lookup and have a router immediately send packets to the exit interface, configure the static route using the *exit-interface* parameter instead of the *ip-address* next-hop parameter.

Technet24

For example, on R2, we can configure static routes to the R1 and R3 LANs by specifying the exit interface:

```
R2(config)# ip route 172.16.3.0 255.255.255.0 serial 0
R2(config)# ip route 192.168.1.0 255.255.255.0 serial
```

Any previous static routes to this network using a next-hop IP address should be removed. R2 now has two static routes in its routing table (see Example 15-4) that it can use immediately to route to the 172.16.3.0/24 and 192.168.1.0/24 networks without having to do a recursive route lookup.

Example 15-4 *R2 Routing Table After the Static Route Is Configured*

```
R2# show ip route
<output omitted>

Gateway of last resort is not set

   10.0.0.0/8 is variably subnetted, 2 subnets, 2 mask
C     10.10.10.0/24 is directly connected, Serial0/1/0
L     10.10.10.2/32 is directly connected, Serial0/1/0
  172.16.0.0/16 is variably subnetted, 5 subnets, 2 ma
C     172.16.1.0/24 is directly connected, GigabitEther
L     172.16.1.1/32 is directly connected, GigabitEther
C     172.16.2.0/24 is directly connected, Serial0/0/0
L     172.16.2.1/32 is directly connected, Serial0/0/0
```

```
S    172.16.3.0/24 is directly connected, Serial0/0/0
  192.168.0.0/24 is variably subnetted, 2 subnets, 2 m
C    192.168.0.0/24 is directly connected, Serial0/0/1
L    192.168.0.1/32 is directly connected, Serial0/0/1
S  192.168.1.0/24 is directly connected, Serial0/0/1
R2#
```

Note

Although the highlighted output in Example 15-4 shows that the routes are directly connected, technically, that is not true. However, as far as R2 is concerned, the exit interface is the way to get to the destination, much as with truly directly connected routes. Another benefit to using the exit interface configuration instead of the next-hop address configuration is that the static route does not depend on the IP address stability of the next hop. Most of the time, using the exit interface configuration is the best practice, so we use the exit interface configuration for all static and default routes as we continue with the reviews.

IPv4 Default Route Configuration

A default route is a special kind of static route used to represent all routes with zero or no bits matching. In other words, when no routes have a more

Technet24

specific match in the routing table, the default route is a match.

The destination IP address of a packet can match multiple routes in the routing table. For example, consider having the following two static routes in the routing table:

```
     172.16.0.0/24 is subnetted, 3 subnets
S  172.16.1.0 is directly connected, Serial0/0/0
S  172.16.0.0/16 is directly connected, Serial0/0/1
```

A packet destined for 172.16.1.10, the packet's destination IP address, matches both routes. However, the 172.16.1.0 route is the more specific route because the destination matches the first 24 bits, whereas the destination matches only the first 16 bits of the 172.16.0.0 route. Therefore, the router uses the route with the most specific match.

A default route is a route that matches all packets. Commonly called a quad-zero route, a default route uses 0.0.0.0 (thus the term *quad-zero*) for both the *network-address* and the *subnet-mask* parameters, as in this syntax:

```
Router(config)# ip route 0.0.0.0 0.0.0.0 {ip-address |
```

Referring to the topology in Figure 15-2, assume that HQ has a connection to the Internet. From the perspective of R2, all default traffic can be sent to HQ

for routing outside the domain known to R2.

The following command configures R2 with a default route pointing to HQ:

```
R2(config)# ip route 0.0.0.0 0.0.0.0 serial 0/1/0
```

R2 now has a gateway of last resort listed in the routing table—a candidate default route indicated by the asterisk (`*`) next to the S code (see Example 15-5).

Example 15-5 *R2 Routing Table After the Default Route Is Configured*

```
R2# show ip route
<some codes omitted>
       * - candidate default, U - per-user static rout
       P - periodic downloaded static route

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

    10.0.0.0/8 is variably subnetted, 2 subnets, 2 mas
C      10.10.10.0/24 is directly connected, Serial0/1/0
L      10.10.10.2/32 is directly connected, Serial0/1/0
   172.16.0.0/16 is variably subnetted, 5 subnets, 2 m
C      172.16.1.0/24 is directly connected, GigabitEthe
L      172.16.1.1/32 is directly connected, GigabitEthe
C      172.16.2.0/24 is directly connected, Serial0/0/0
L      172.16.2.1/32 is directly connected, Serial0/0/0
```

Technet24

```
S     172.16.3.0/24 is directly connected, Serial0/0/0
   192.168.0.0/24 is variably subnetted, 2 subnets, 2
C     192.168.0.0/24 is directly connected, Serial0/0/
L     192.168.0.1/32 is directly connected, Serial0/0/
S     192.168.1.0/24 is directly connected, Serial0/0/
S*  0.0.0.0/0 is directly connected, Serial0/1/0
R2#
```

From R1's and R3's perspective, R2 is the default route. The following commands configure R1 and R3 with a default route pointing to R2:

```
R1(config)# ip route 0.0.0.0 0.0.0.0 serial 0/0/0
!
R3(config)# ip route 0.0.0.0 0.0.0.0 serial 0/0/1
```

Again, we can verify that the default route is now in the routing table for R1 (see Example 15-6).

Example 15-6 *R1 and R3 Routing Tables After the Default Route Is Configured*

```
R1# show ip route
<some codes omitted>
       * - candidate default, U - per-user static rou
       P - periodic downloaded static route
```

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

   172.16.0.0/16 is variably subnetted, 4 subnets, 2 m
C    172.16.2.0/24 is directly connected, Serial0/0/0
L    172.16.2.2/32 is directly connected, Serial0/0/0
C    172.16.3.0/24 is directly connected, GigabitEther
L    172.16.3.1/32 is directly connected, GigabitEther
S* 0.0.0.0/0 is directly connected, Serial0/0/0
R1#
```

```
R3# show ip route
<some codes omitted>
    * - candidate default, U - per-user static route,
    P - periodic downloaded static route

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

   192.168.0.0/24 is variably subnetted, 2 subnets, 2
C     192.168.0.0/24 is directly connected, Serial0/0/
L     192.168.0.2/32 is directly connected, Serial0/0/
  192.168.1.0/24 is variably subnetted, 2 subnets, 2 m
C     192.168.1.0/24 is directly connected, GigabitEth
L     192.168.1.1/32 is directly connected, GigabitEth
S* 0.0.0.0/0 is directly connected, Serial0/0/1
R3#
```

Technet24

After evaluating the complete routing tables for R1, R2, and R3 shown in Examples 15-5 and 15-6, you can see that R1 and R3 need only one route out—a default route. R2 acts as a hub router to the R1 and R3 spokes. Therefore, it needs two static routes pointing to the R1 and R3 LANs. R2 also has a route out to HQ for any destinations it does not know about. But what about HQ? Currently, HQ does not have routes back to any of the networks accessible through R2. Any traffic from PC1, PC2, and PC3 is thus currently confined to the R1, R2, and R3 networks. None of these PCs can ping the HQ interface address 10.10.10.1. In the traceroute output in Example 15-7, failure occurs after R2 responds. This is because HQ receives the ICMP requests from PC1 but does not have a route back to the 172.16.3.0/24 network. Therefore, HQ drops the packets.

Example 15-7 *A Failed traceroute from PC1 to HQ*

```
C:\> tracert 10.10.10.1

Tracing route to 10.10.10.1 over a maximum of 30 hops:

  1  0 ms   0 ms   1 ms   172.16.3.1
  2  0 ms   0 ms   1 ms   172.16.2.1
  3  *      *     *      Request timed out.
  4  ^C
C:\>
```

In the next section, we configure HQ with static routes to complete the static route configuration for the topology in Figure 15-2.

IPv4 Summary Static Route Configuration

Before configuring five separate static routes for each of the networks in Figure 15-2, notice that the 172.16 networks can be summarized into one route and that the 192.168 networks can be summarized into one route. Example 15-8 shows the five routes in binary, with the bits in common highlighted.

Example 15-8 *Summary Route Calculation for HQ Static Routes*

```
Summary calculation for the 172.16 networks:
10101100.00010000.00000001.00000000
10101100.00010000.00000010.00000000
10101100.00010000.00000011.00000000
Summary calculation for the 192.168 networks:
11000000.10101000.00000000.00000000
11000000.10101000.00000001.00000000
```

The summary route for the 172.16 networks is 172.16.0.0/22 because the three network addresses have 22 bits in common. Although this summary static route is not part of the current addressing scheme, it also includes the route 172.16.0.0/24. The summary route for the 192.168 networks is 192.168.0.0/23 because the two network addresses have 23 bits in common.

Technet24

We can now configure HQ with two summary static routes instead of five individual static routes:

```
HQ(config)# ip route 172.16.0.0 255.255.252.0 serial 0
HQ(config)# ip route 192.168.0.0 255.255.254.0 Serial0
```

Now PC1 can successfully trace a route to the HQ interface, as Example 15-9 shows.

Example 15-9 *A Successful traceroute from PC1 to HQ*

```
C:\> tracert 10.10.10.1

Tracing route to 10.10.10.1 over a maximum of 30 hops:

 1      1 ms     0 ms       0 ms    172.16.3.1
 2      0 ms     1 ms       2 ms    172.16.2.1
 3      1 ms     2 ms       1 ms    10.10.10.1

Trace complete.

C:\>
```

The trace is successful because HQ now has a route back to PC1's network, as shown in Example 15-10.

Example 15-10 *HQ Routing Table with IPv4 Summary Static Routes*

```
HQ# show ip route
<output omitted>

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 2 subnets, 2 mas
C       10.10.10.0/24 is directly connected, Serial0/0/0
L       10.10.10.1/32 is directly connected, Serial0/0/0
    172.16.0.0/22 is subnetted, 1 subnets
S       172.16.0.0/22 is directly connected, Serial0/0/0
S       192.168.0.0/23 is directly connected, Serial0/0/
HQ#
```

IPv6 Static Routing

Static routing with IPv6 is similar to static routing with IPv4. We can use the same topology but change the addressing to IPv6, as shown in Figure 15-3.

Technet24

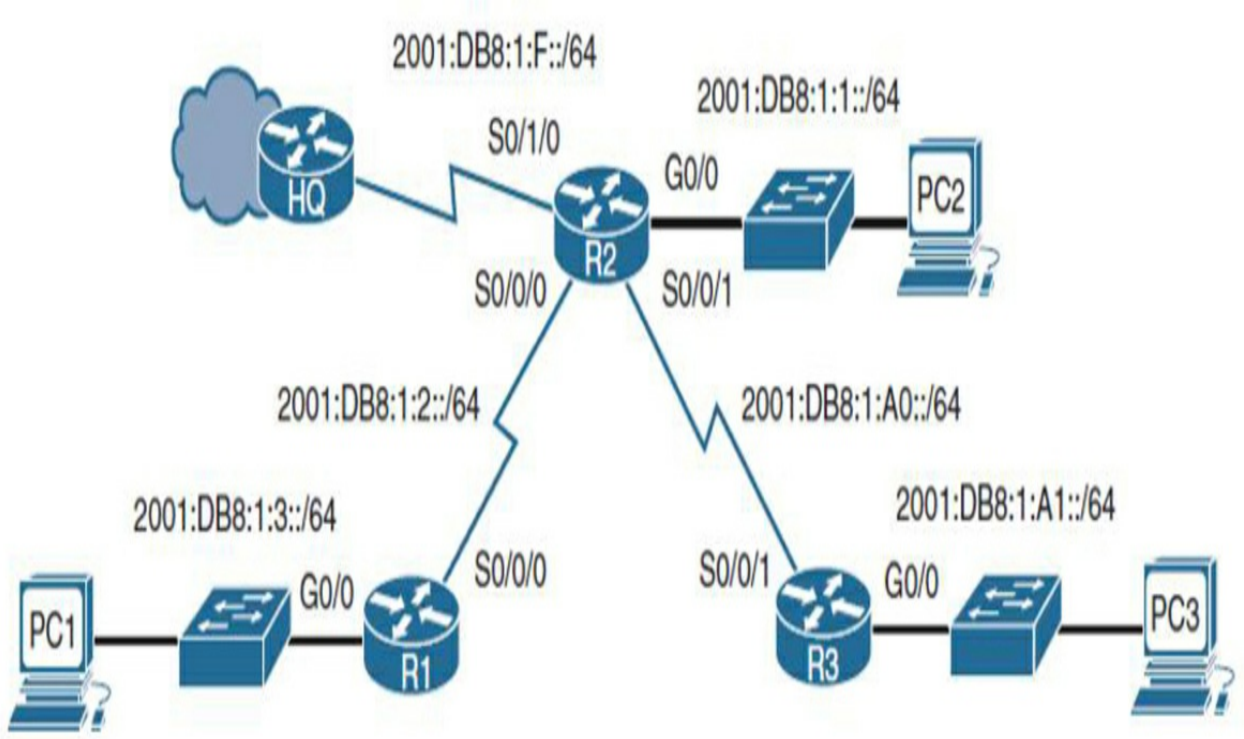

Figure 15-3 *Static and Default Routing IPv6 Topology*

Table 15-2 shows the IPv6 addressing scheme used with the topology in Figure 15-3.

Table 15-2 *IPv6 Addressing Scheme*

Device	Interface	IPv6 Address/Prefix	Default Gateway
HQ	S0/0/0 Link-local	2001:DB8:1:F::1/64 FE80::F	—
R1	G0/0 S0/0/0	2001:DB8:1:3::1/64 2001:DB8:1:2::2/64 FE80::1	— —
R2	G0/0 S0/0/0 S0/0/1 S0/1/0 Link-local	2001:DB8:1:1::1/64 2001:DB8:1:2::1/64 2001:DB8:1:A0::1/64 2001:DB8:1:F::2/64 FE80::2	— — — —
R3	G0/0 S0/0/1 Link-local	2001:DB8:1:A1::1/64 2001:DB8:1:A0::2/64 FE80::3	— — —
PC1	NIC	2001:DB8:1:3:209:7CFF:FE9 A:1A87/64	FE80::1
PC2	NIC	2001:DB8:1:1:204:9AFF:FEE 3: C943/64	FE80::2
PC3	NIC	2001:DB8:1:A1:201:C9FF:FE E5: D3A/64	FE80::3

Technet24

Note

The IPv6 addressing for the PCs is set to autoconfiguration. Pinging from PC to PC would not really be much fun. However, the IPv6 addresses are not manually set so that you can practice your knowledge of how EUI-64 works. Can you figure out the MAC address for each PC? If not, review Day 27, "IPv6 Addressing." (Hint: FFFE and flip the bit.) If you are following along using a simulator, you might want to consider manually configuring the PCs with easier IPv6 addresses—such as 2001:DB8:1:3::A/64 on PC1. Doing so will greatly improve your pinging experience.

IPv6 Static Route Configuration

The command syntax for IPv6 static routing is similar to the syntax for IPv4:

```
Router(config)# ipv6 route ipv6-prefix/prefix-length {
exit-interface} [administrative-distance]
```

Therefore, the following commands configure R2 with static routes to the R1 and R3 LANs:

```
R2(config)# ipv6 route 2001:DB8:1:3::/64 Serial0/0/0
```

```
R2(config)# ipv6 route 2001:DB8:1:A1::/64 Serial0/0/1
```

As highlighted in the output from the **show ipv6 route** command in Example 15-11, R2 now has routes in the routing table to the R1 and R3 LANs.

Example 15-11 *R2 IPv6 Routing Table After Static Routes Are Configured*

```
R2# show ipv6 route
IPv6 Routing Table - 11 entries
<code output omitted>
C   2001:DB8:1:1::/64 [0/0]
      via ::, GigabitEthernet0/0
L   2001:DB8:1:1::1/128 [0/0]
      via ::, GigabitEthernet0/0
C   2001:DB8:1:2::/64 [0/0]
      via ::, Serial0/0/0
L   2001:DB8:1:2::1/128 [0/0]
      via ::, Serial0/0/0
S   2001:DB8:1:3::/64 [1/0]
      via ::, Serial0/0/0
C   2001:DB8:1:F::/64 [0/0]
      via ::, Serial0/1/0
L   2001:DB8:1:F::2/128 [0/0]
      via ::, Serial0/1/0
C   2001:DB8:1:A0::/64 [0/0]
      via ::, Serial0/0/1
L   2001:DB8:1:A0::1/128 [0/0]
```

Technet24

```
      via ::, Serial0/0/1
S   2001:DB8:1:A1::/64 [1/0]
      via ::, Serial0/0/1
L  FF00::/8 [0/0]
     via ::, Null0
```

IPv6 Default Route Configuration

The following is the command syntax for an IPv6 default route:

```
Router(config)# ipv6 route ::/0 {ipv6-address | exit-i
```

Just as with the quad-zero in IPv4, the double colon (::) means all 0s or any address, and the /0 means any prefix length.

Continuing with the example in Figure 15-3, we can configure R1, R2, and R3 with the following default routes:

```
R1(config)# ipv6 route ::/0 serial 0/0/0
R2(config)# ipv6 route ::/0 serial 0/1/0
R3(config)# ipv6 route ::/0 serial 0/0/1
```

The highlights in Example 15-12 show the default routes for R1, R2, and R3.

Example 15-12 *Default Routes in the Routing Tables for R1, R2, and R3*

```
R1# show ipv6 route
IPv6 Routing Table - 6 entries
<code output omitted>
S   ::/0 [1/0]
     via ::, Serial0/0/0
<output for connected and local routes omitted>
```

```
R2# show ipv6 route
IPv6 Routing Table - 12 entries
<code output omitted>
S   ::/0 [1/0]
     via ::, Serial0/1/0
S   2001:DB8:1:3::/64 [1/0]
      via ::, Serial0/0/0
S   2001:DB8:1:A1::/64 [1/0]
      via ::, Serial0/0/1
<output for connected and local routes omitted>
```

```
R3# show ipv6 route
IPv6 Routing Table - 6 entries
<code output omitted>
S   ::/0 [1/0]
     via ::, Serial0/0/1
<output for connected and local routes omitted>
```

Technet24

IPv6 Summary Static Route Configuration

Much as in the IPv4 static routing scenario, HQ can be configured with two summary static routes to the R1, R2, and R3 LANs. Example 15-13 shows the first four hextets (64 bits) of the five routes in binary, with the bits in common highlighted.

Example 15-13 *Summary Route Calculation for HQ Static Routes*

```
Summary calculation for the first four hextets of 2001
2001:DB8:1:2::/64, and 2001:DB8:1:3::/64 networks:
0010000000000001:0000110110111000:0000000000000001:000
0010000000000001:0000110110111000:0000000000000001:000
0010000000000001:0000110110111000:0000000000000001:000
Summary calculation for the first four hextets of 2001
 2001:DB8:1:A1::/64 networks:
0010000000000001:0000110110111000:0000000000000001:000
0010000000000001:0000110110111000:0000000000000001:000
```

The first summary route is 2001:DB8:1::/62 because the three network addresses have 62 bits in common. Although this summary static route is not part of the current addressing scheme, it also includes the network 2001:DB8:1::/64. The second summary route is 2001:DB8:1:A0::/63 because the two network addresses have 63 bits in common.

You can now configure HQ with the following two summary static routes:

```
HQ(config)# ipv6 route 2001:DB8:1::/62 Serial0/0/0
HQ(config)# ipv6 route 2001:DB8:1:A0::/63 Serial0/0/0
```

Now HQ has two summary routes, as you can see in the highlighted entries in Example 15-14.

Example 15-14 *HQ Routing Table with IPv6 Summary Static Routes*

```
HQ# show ipv6 route
IPv6 Routing Table - 5 entries
<output omitted>
S 2001:DB8:1::/62 [1/0]
       via ::, Serial0/0/0
C 2001:DB8:1:F::/64 [0/0]
      via ::, Serial0/0/0
L 2001:DB8:1:F::1/128 [0/0]
      via ::, Serial0/0/0
S 2001:DB8:1:A0::/63 [1/0]
      via ::, Serial0/0/0
L FF00::/8 [0/0]
 via ::, Null0
HQ#
```

Study Resources

Technet24

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Switching, Routing, and Wireless Essentials v7	15
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	17
	29
Portable Command Guide	15

Day 14. OSPF Operation

CCNA 200-301 Exam Topics

• 3.4 Configure and verify single-area OSPFv2

Key Topics

Today we review the basic operation of OSPF. OSPFv2 is used for IPv4 routing, and OSPFv3 is used for IPv6 routing. Although the two versions share the same basic operation principles, we also review how they differ. Tomorrow, we will review single-area OSPFv2 configuration.

Single-Area OSPF Operation

The Internet Engineering Task Force (IETF) chose OSPF over Intermediate System-to-Intermediate System (IS-IS) as its recommended interior gateway protocol (IGP). In 1998, the OSPFv2 specification was updated in RFC 2328, "OSPF Version 2" (see http://www.ietf.org/rfc/rfc2328). Cisco IOS Software chooses OSPF routes over RIP routes because OSPF has an administrative distance (AD) of 110 versus RIP's AD of 120.

OSPF Message Format

The data portion of an OSPF message is encapsulated in a packet. This data field can include one of five OSPF packet types. Figure 14-1 shows an encapsulated OSPF message in an Ethernet frame.

Technet24

The OSPF packet header is included with every OSPF packet, regardless of its type. The OSPF packet header and packet type–specific data are then encapsulated in an IP packet. In the IP packet header, the protocol field is set to 89 to indicate OSPF, and the destination address is typically set to one of two multicast addresses: 224.0.0.5 or 224.0.0.6. If the OSPF packet is encapsulated in an Ethernet frame, the destination MAC address is also a multicast address: 01-00-5E-00-00-05 or 01-00-5E-00-00-06.

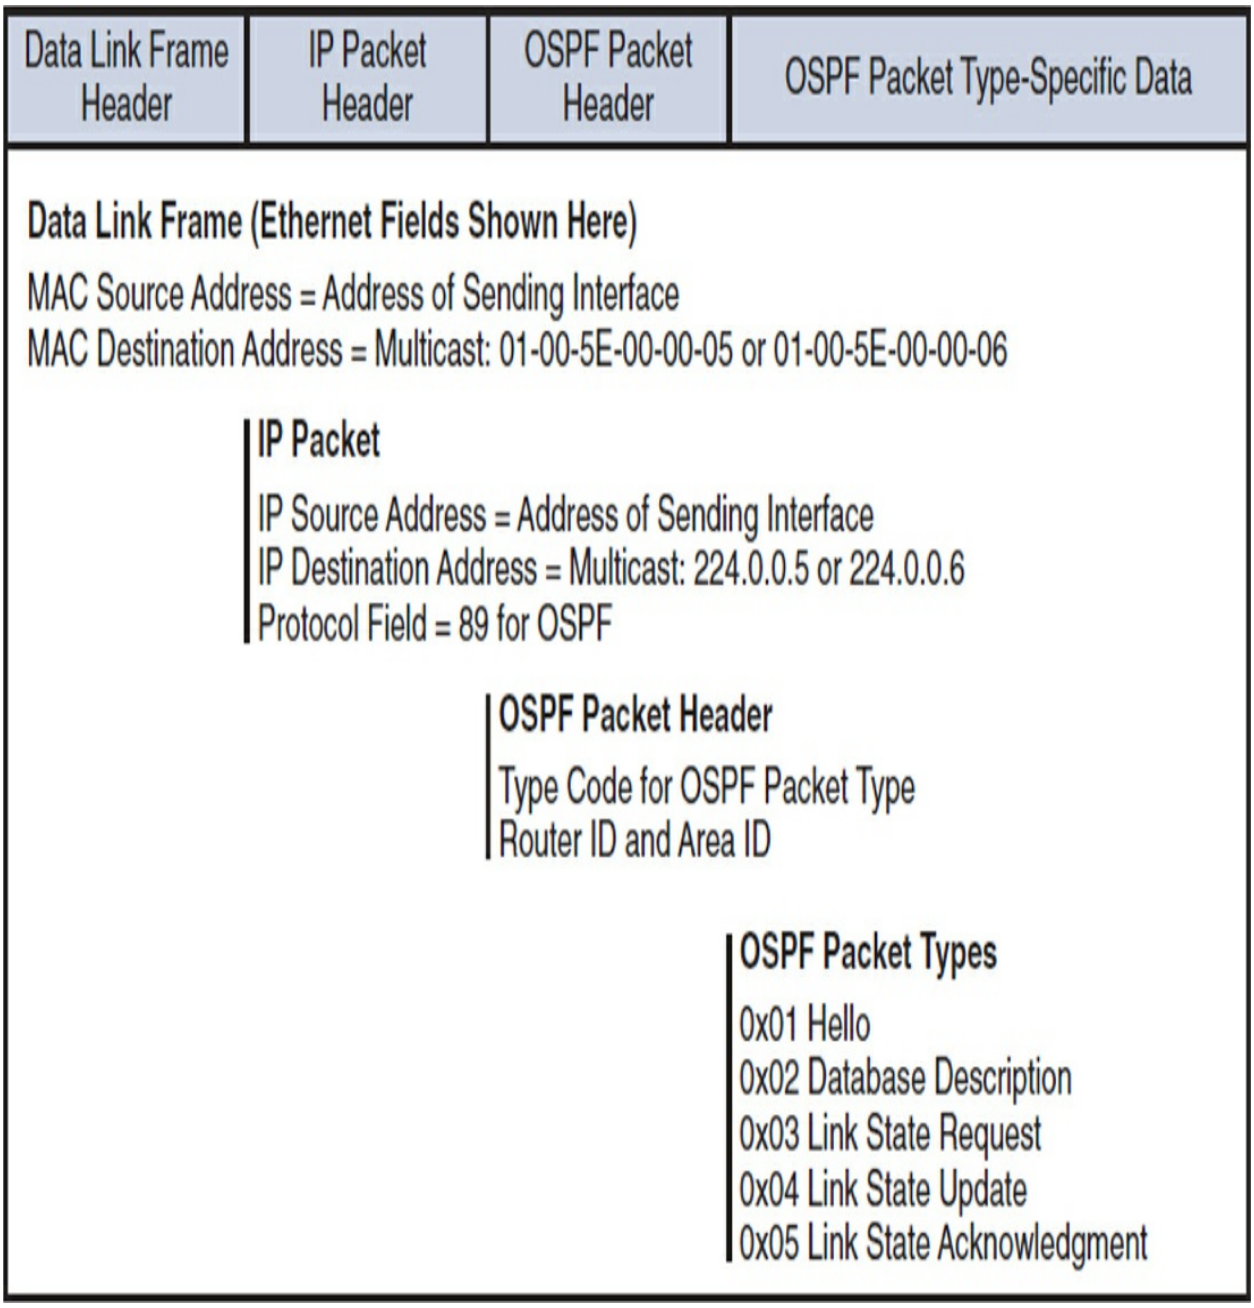

Figure 14-1 *Encapsulated OSPF Message*

OSPF Packet Types

Each of the five OSPF packet types serves a specific purpose in the routing process:

Technet24

• **Hello:** Hello packets establish and maintain adjacency with other OSPF routers.

• **DBD:** The database description (DBD) packet contains an abbreviated list of the sending router's link-state database. Receiving routers use it to check against the local link-state database.

• **LSR:** Receiving routers can request more information about any entry in the DBD by sending a link-state request (LSR).

• **LSU:** Link-state update (LSU) packets reply to LSRs and announce new information. LSUs contain 11 types of link-state advertisements (LSAs).

• **LSAck:** When an LSU is received, the router sends a link-state acknowledgment (LSAck) to confirm receipt of the LSU.

Neighbor Establishment

OSPF neighbors exchange hello packets to establish adjacency. Figure 14-2 shows the OSPF header and hello packet.

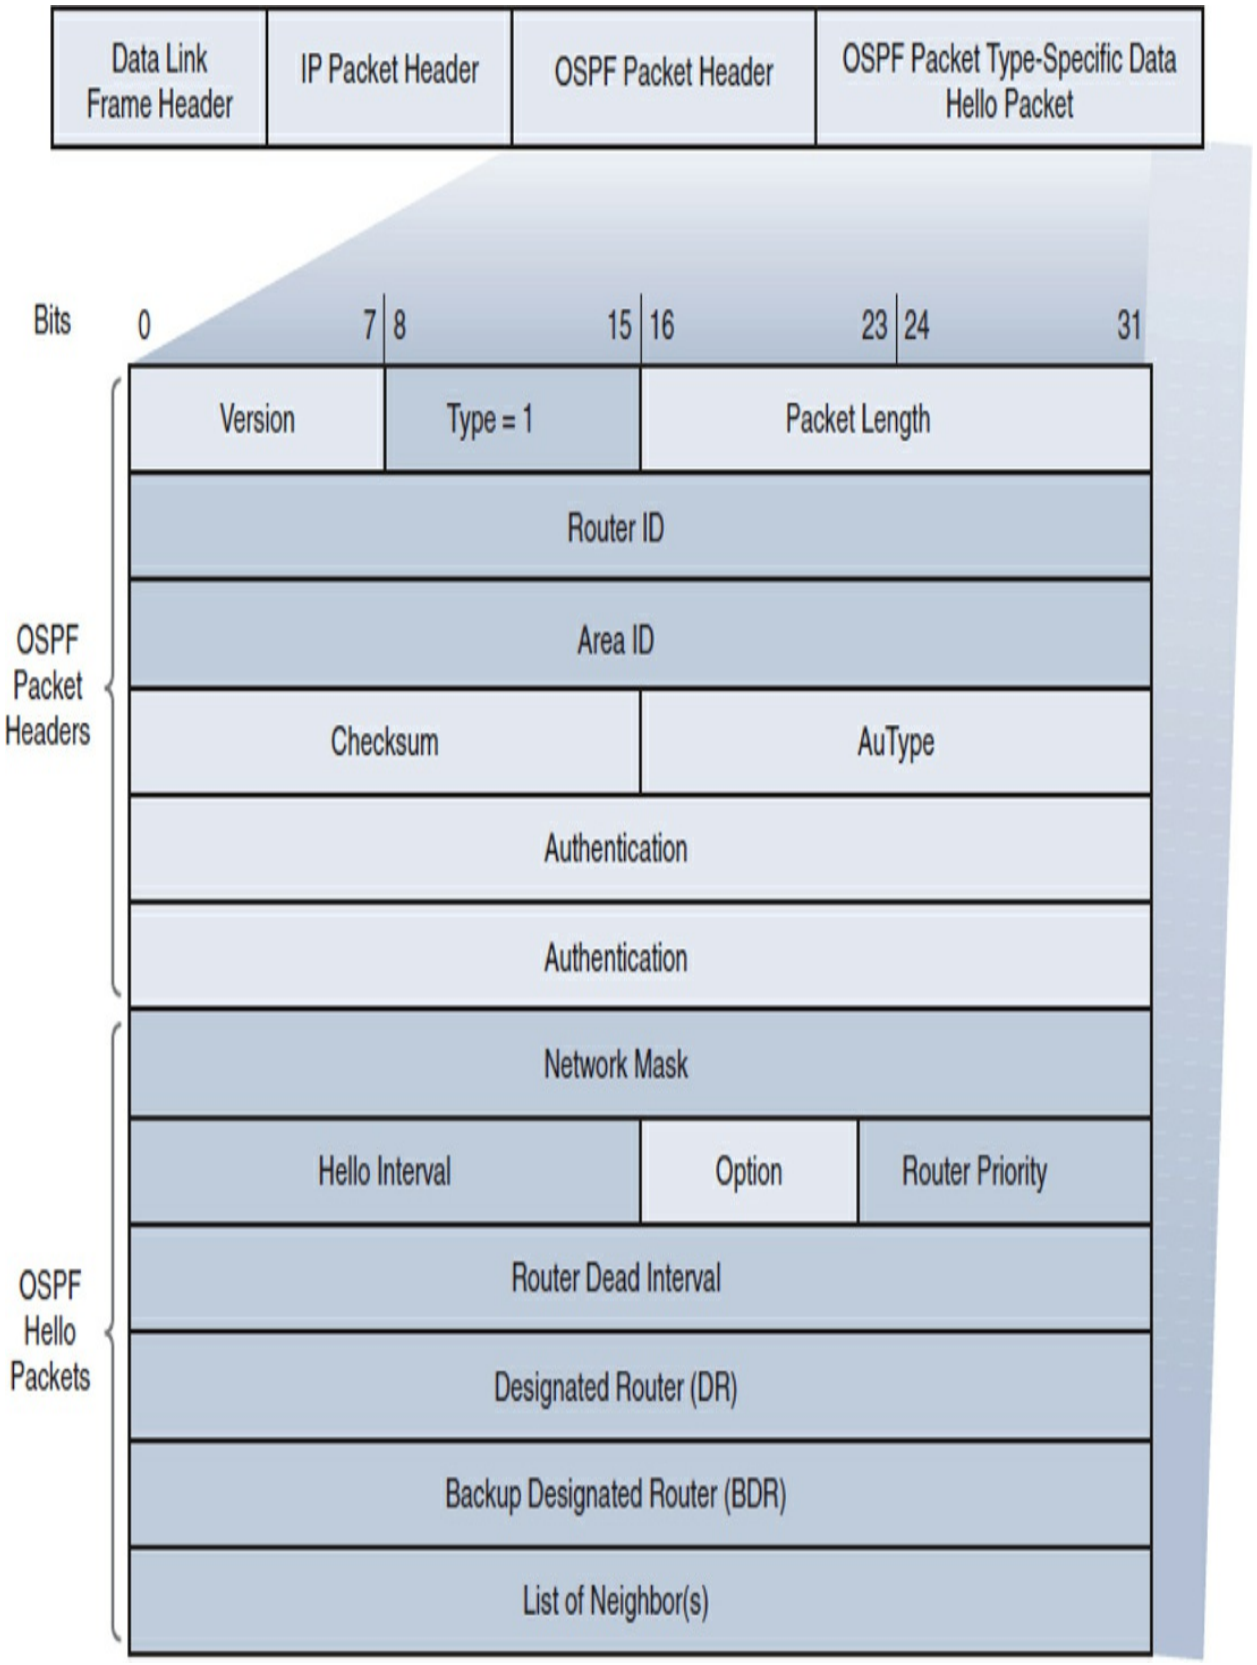

Figure 14-2 *OSPF Packet Header and Hello Packet*

Technet24

Important fields shown in the figure include the following:

• **Type:** OSPF packet type: Hello (Type 1), DBD (Type 2), LS Request (Type 3), LS Update (Type 4), LS ACK (Type 5)

• **Router ID:** ID of the originating router

• **Area ID:** Area from which the packet originated

• **Network Mask:** Subnet mask associated with the sending interface

• **Hello Interval:** Number of seconds between the sending router's hellos

• **Router Priority:** Used in DR/BDR election

• **Designated Router (DR):** Router ID of the DR, if any

• **Backup Designated Router (BDR):** Router ID of the BDR, if any

• **List of Neighbors:** The OSPF router ID of the neighboring router(s)

Hello packets are used to do the following:

• Discover OSPF neighbors and establish neighbor adjacencies

• Advertise parameters on which two routers must agree to become neighbors

• Elect the DR and BDR on multiaccess networks such as Ethernet and Frame

Relay

Receiving an OSPF hello packet on an interface confirms for a router that another OSPF router exists on this link. OSPF then establishes adjacency with the neighbor. To establish adjacency, two OSPF routers must have the following matching interface values:

• Hello Interval

• Dead Interval

• Network Type

• Area ID

Before the two routers can establish adjacency, both interfaces must be part of the same network, including the same subnet mask. Full adjacency happens after the two routers have exchanged any necessary LSUs and have identical link-state databases. By default, OSPF hello packets are sent to the multicast address 224.0.0.5 (*ALLSPFRouters*) every 10 seconds on multiaccess and point-to-point segments and every 30 seconds on nonbroadcast multiaccess (NBMA) segments (Frame Relay, X.25, ATM). The default dead interval is four times the hello interval.

Link-State Advertisements

LSUs are the packets used for OSPF routing updates. An LSU packet can

Technet24

contain 11 types of LSAs, as Figure 14-3 shows.

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them.
2	DBD	Checks for database synchronization between routers.
3	LSR	Requests specific link-state records from router to router.
4	LSU	Sends specifically requested link-state records.
5	LSAck	Acknowledges the other packet types.

The acronyms LSA and LSU are often used interchangeably.

An LSU contains one or more LSAs.

LSAs contain route information for destination networks.

LSA specifics are discussed in CCNP.

LSA Type	Description
1	Router LSAs
2	Network LSAs
3 or 4	Summary LSAs
5	Autonomous System External LSAs
6	Multicast OSPF LSAs
7	Defined for Not-So-Stubby Areas
8	External Attributes LSA for Border Gateway Protocol (BGP)
9, 10, 11	Opaque LSAs

Figure 14-3 *LSUs Contain LSAs*

OSPF DR and BDR

Multiaccess networks create two challenges for OSPF regarding the flooding of LSAs:

• Creation of multiple adjacencies, with one adjacency for every pair of routers

• Extensive flooding of LSAs

The solution to managing the number of adjacencies and the flooding of LSAs on a multiaccess network is the designated router (DR). To reduce the amount of OSPF traffic on multiaccess networks, OSPF elects a DR and a backup DR (BDR). The DR is responsible for updating all other OSPF routers when a change occurs in the multiaccess network. The BDR monitors the DR and takes over as DR if the current DR fails. All other routers become DROTHERs. A DROTHER is a router that is neither the DR nor the BDR.

OSPF Algorithm

Each OSPF router maintains a link-state database containing the LSAs received from all other routers. When a router has received all the LSAs and built its local link-state database, OSPF uses Dijkstra’s shortest path first (SPF) algorithm to create an SPF tree. This algorithm accumulates costs along each path, from source to destination. The SPF tree is then used to

Technet24

populate the IP routing table with the best paths to each network.

For example, in Figure 14-4, each path is labeled with an arbitrary value for cost. The cost of the shortest path for R2 to send packets to the LAN attached to R3 is 27 (20 + 5 + 2 = 27). Notice that this cost is not 27 for all routers to reach the LAN attached to R3. Each router determines its own cost to each destination in the topology. In other words, each router uses the SPF algorithm to calculate the cost of each path to a network and determines the best path to that network from its own perspective.

Table 14-1 lists the shortest path to each LAN for R1, along with the cost.

Table 14-1 *SPF Tree for R1*

Destination	Shortest Path	Cost
R2 LAN	R1 to R2	22
R3 LAN	R1 to R3	7
R4 LAN	R1 to R3 to R4	17
R5 LAN	R1 to R3 to R4 to R5	27

You should be able to create a similar table for each of the other routers in

Figure 14-4.

Technet24

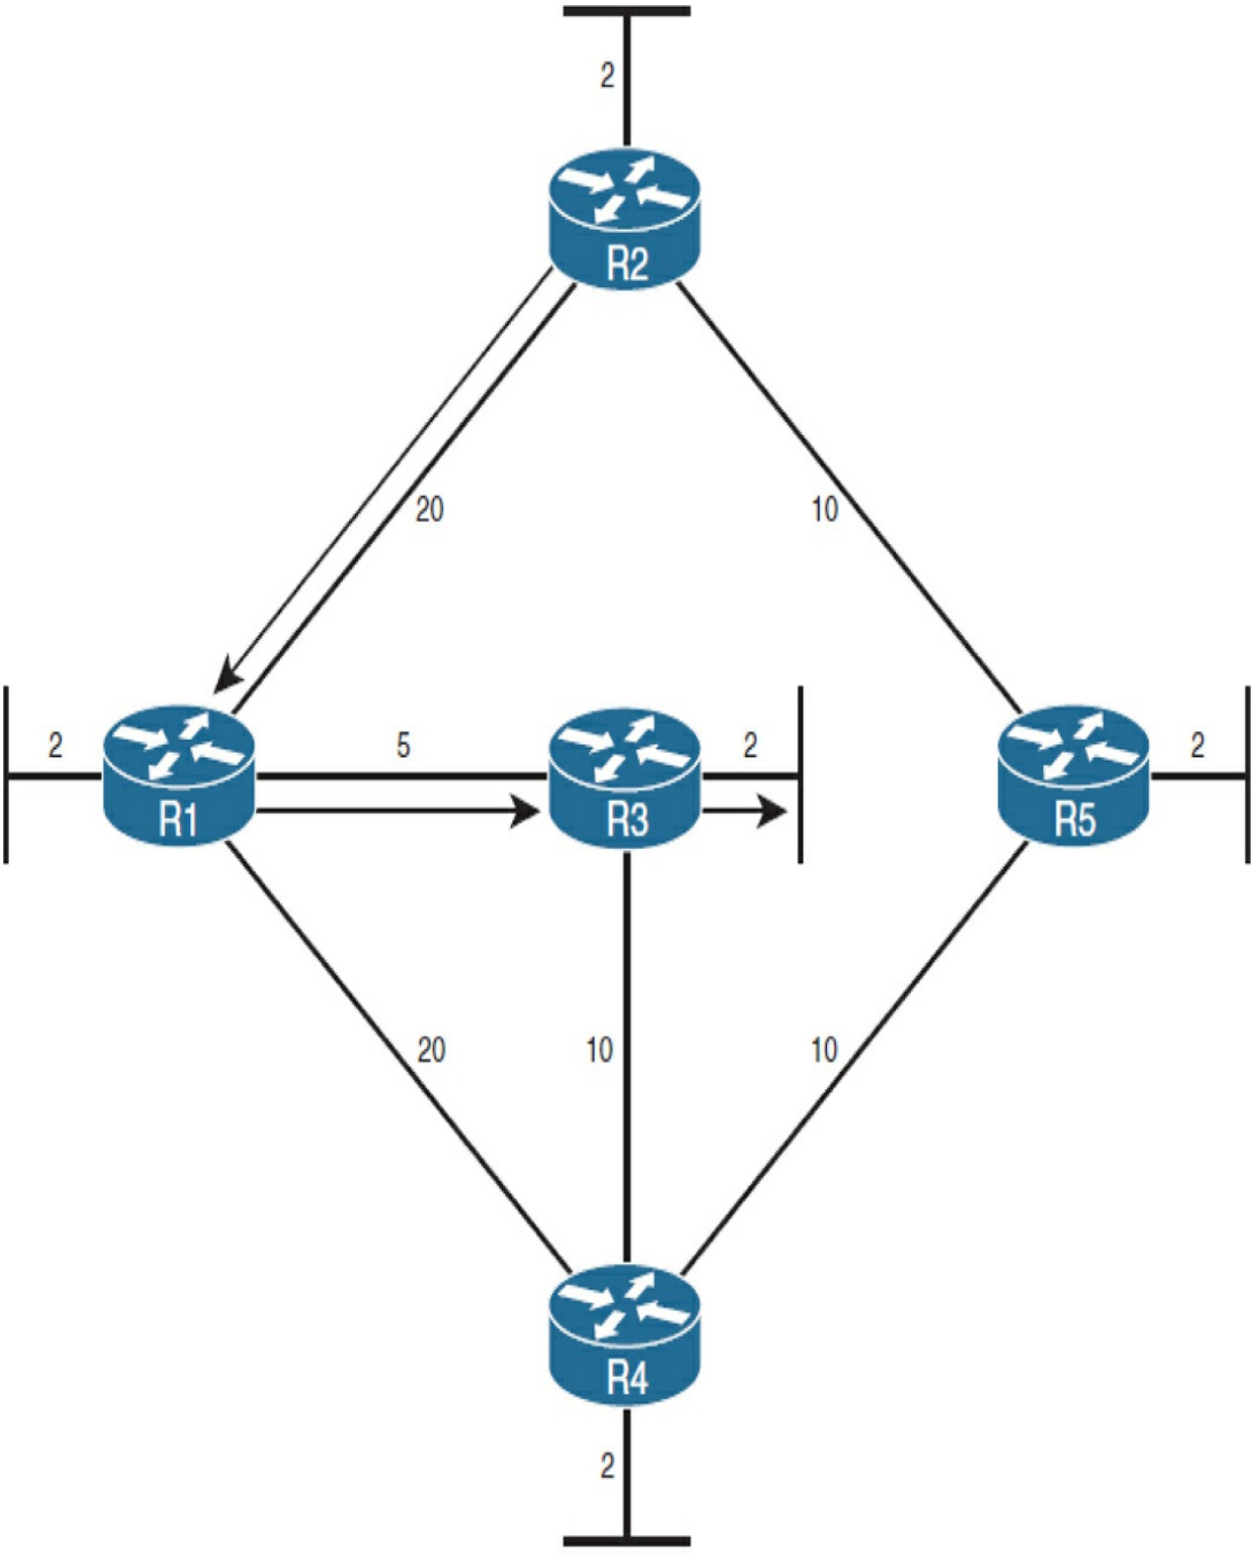

Shortest Path for Host on R2 LAN to Reach Host on R3 LAN:
R2 to R1 (20) + R1 to R3 (5) + R3 to LAN (2) = 27

Figure 14-4 *Dijkstra's Shortest Path First Algorithm*

Link-State Routing Process

The following list summarizes the link-state routing process OSPF uses. All OSPF routers complete the following generic link-state routing process to reach a state of convergence:

Each router learns about its own links and its own directly connected networks. This is done by detecting that an interface has a Layer 3 address configured and is in the up state.

Each router is responsible for establishing adjacency with its neighbors on directly connected networks by exchanging hello packets.

Each router builds a link-state packet (LSP) containing the state of each directly connected link. This is done by recording all the pertinent information about each neighbor, including neighbor ID, link type, and bandwidth.

Each router floods the LSP to all neighbors, which then store all LSPs received in a database. Neighbors then flood the LSPs to their neighbors until all routers in the area have received the LSPs. Each router stores a copy of each LSP received from its neighbors in a local database.

Each router uses the database to construct a complete map of the topology and computes the best path to each destination network. The SPF algorithm is

Technet24

used to construct the map of the topology and determine the best path to each network. All routers have a common map or tree of the topology, but each router independently determines the best path to each network within that topology.

OSPFv2 Versus OSPFv3

In 1999, OSPFv3 for IPv6 was published in RFC 2740. In 2008, OSPFv3 was updated in RFC 5340 as OSPF for IPv6. However, it is still referred to as OSPFv3.

OSPFv3 has the same functionality as OSPFv2 but uses IPv6 as the network layer transport, communicating with OSPFv3 peers and advertising IPv6 routes. OSPFv3 also uses the SPF algorithm as the computation engine to determine the best paths throughout the routing domain.

As with all other IPv6 routing protocols, OSPFv3 has separate processes from its IPv4 counterpart. OSPFv2 and OSPFv3 each have separate adjacency tables, OSPF topology tables, and IP routing tables.

Similarities Between OSPFv2 and OSPFv3

OSPFv3 operates much like OSPFv2. Table 14-2 summarizes the operational features that OSPFv2 and OSPFv3 share.

Table 14-2 *OSPFv2 and OSPFv3 Similarities*

Feature	OSPFv2 and OSPFv3
Link state	Yes
Routing algorithm	SPF
Metric	Cost
Areas	Support the same two-level hierarchy
Packet types	Use the same hello, DBD, LSR, LSU, and LSAck packets
Neighbor discovery	Transition through the same states using hello packets
LSDB synchronization	Exchange contents of their LSDB between two neighbors
DR and BDR	Use the same function and election process
Router ID	Use a 32-bit router ID and the same process in determining the 32-bit router ID

Differences Between OSPFv2 and OSPFv3

Technet24

Table 14-3 lists the major differences between OSPFv2 and OSPFv3.

Table 14-3 *OSPFv2 and OSPFv3 Differences*

Feature	OSPFv2	OSPFv3
Advertising	IPv4 networks	IPv6 prefixes
Source address	IPv4 source address	IPv6 link-local address
Destination address	Choice of Neighbor IPv4 unicast address 224.0.0.5, all-OSPF-routers multicast address 224.0.0.6, DR/BDR multicast address	Choice of Neighbor IPv6 link-local address FF02::5, all-OSPFv3-routers multicast address FF02::6, DR/BDR multicast address
Advertising networks	Configured using the **network** router configuration command	Configured using the **ipv6 ospf area** interface configuration command
IP unicast routing	IPv4 unicast routing enabled by default	Requires configuration of the **ipv6 unicast-routing** global configuration command
Authentication	Plain text and MD5	IPsec

Technet24

Multiarea OSPF Operation

Single-area OSPF works fine in smaller networks in which the number of links is manageable. However, consider an OSPF single-area network with 900 routers and several thousand subnets. In this situation, the single-area design causes the following problems:

• **Large routing tables:** By default, OSPF does not summarize routing updates.

• **Large link-state database (LSDB):** In a single area, each router must maintain a database of all active links in the routing domain, regardless of whether that router is currently using a particular link.

• **Frequent SPF calculations:** In a large network, changes to the LSDB can cause routers to spend many CPU cycles recalculating the SPF algorithm and updating the routing table.

To address these issues, OSPF supports hierarchical design through the uses of multiple OSPF areas. Multiarea OSPF is useful in larger network deployments to reduce processing and memory overhead. This involves breaking the one large LSDB into several smaller LSDBs by using multiple OSPF areas.

Multiarea OSPF Design

Multiarea OSPF design follows a couple basic rules:

• Put all interfaces connected to the same subnet inside the same area.

• An area should be contiguous.

• Some routers might be internal to an area, with all interfaces assigned to that single area.

• Some routers might be area border routers (ABRs) because some interfaces connect to the backbone area and some connect to nonbackbone areas.

• All nonbackbone areas must connect to the backbone area (Area 0) by having at least one ABR connected to both the backbone area and the nonbackbone area.

Figure 14-5 shows a simple multiarea OSPF design with two areas (Area 1 and Area 2) connected to a backbone, Area 0.

Technet24

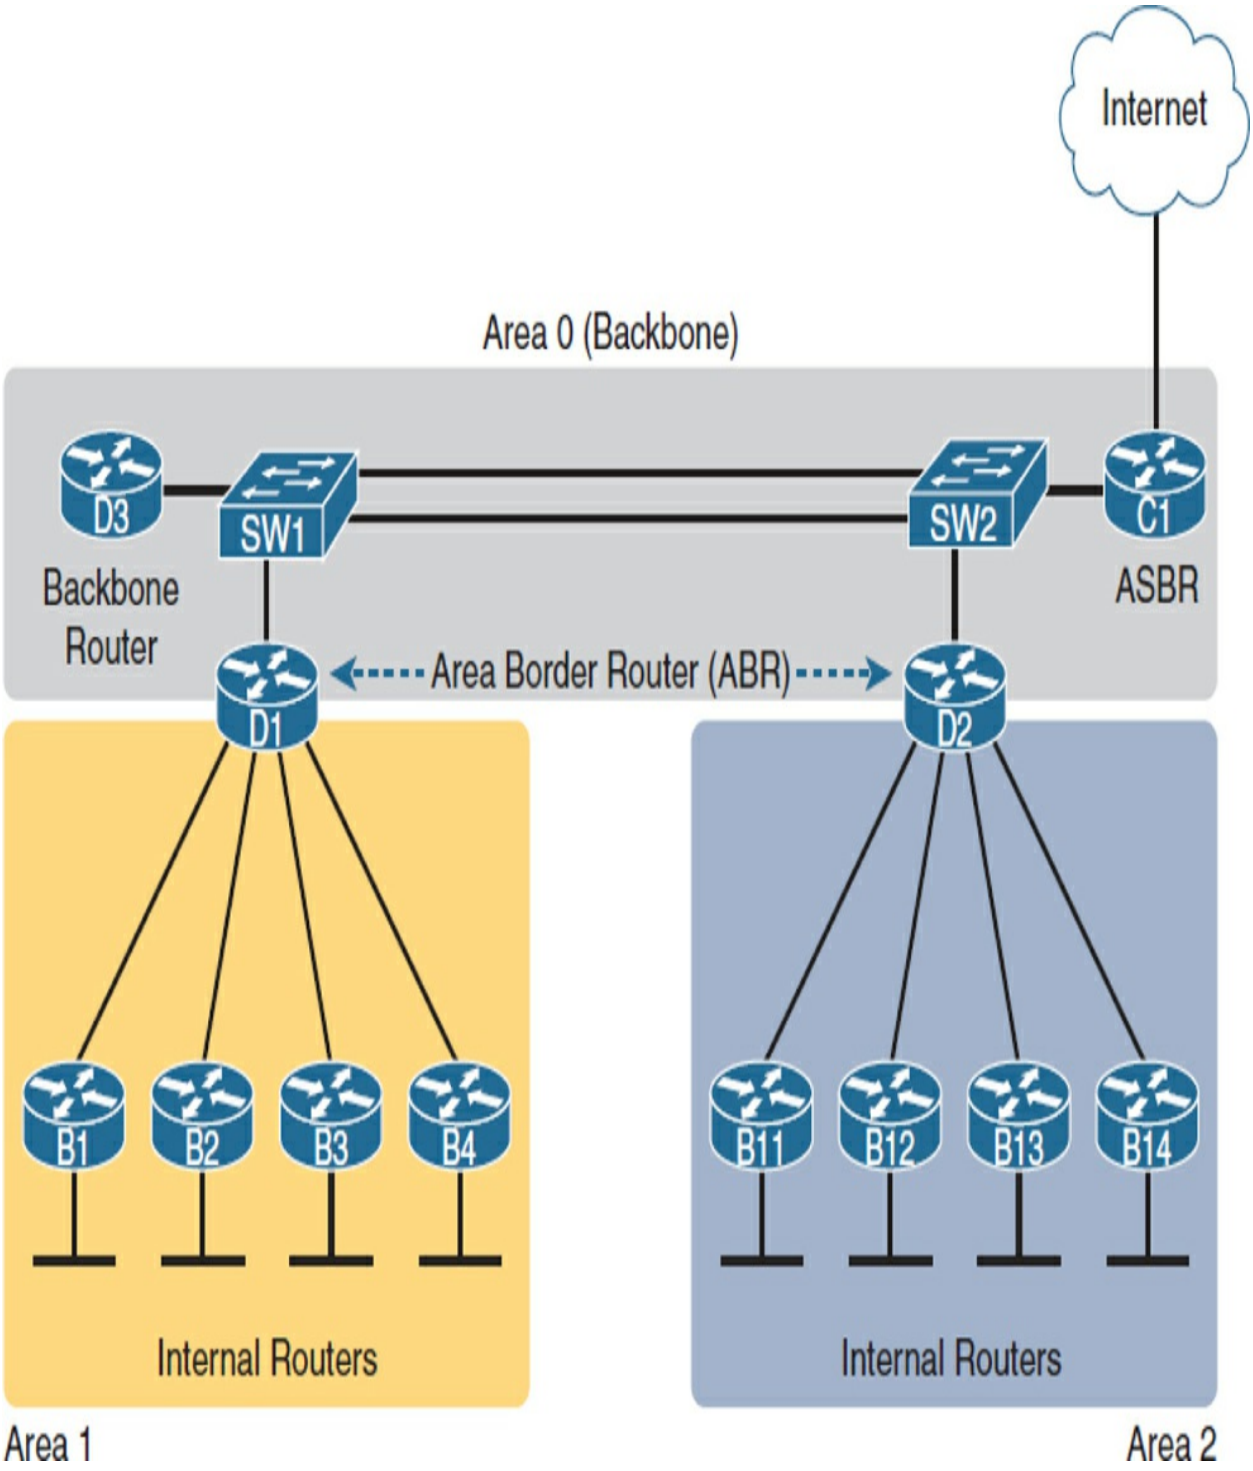

Figure 14-5 *Sample Multiarea OSPF Design*

The figure also shows a few important multiarea OSPF design terms. Table 14-4 describes these terms.

Table 14-4 *Multiarea OSPF Design Terminology*

Technet24

Term	Description
Area border router (ABR)	An OSPF router with interfaces connected to the backbone area and to at least one other area.
Backbone router	A router connected to the backbone area (includes ABRs).
Internal router	A router in one area (not the backbone area).
Autonomous system boundary router (ASBR)	A router that has at least one interface connected to an external network. An external network is a network that is not part of the routing domain, such as EIGRP, BGP, or one with static routing to the Internet, as Figure 14-5 shows.
Area	A set of routers and links that shares the same detailed LSDB information—but not with routers in other areas—for better efficiency
Backbone area	A special OSPF area to which all other areas must connect, such as Area 0
Intra-area route	A route to a subnet inside the same area as the router
Interarea route	A route to a subnet in an area the router is not a part of

Multiarea OSPF Improves Performance

In multiarea OSPF, all areas must connect to the backbone area. Routing still occurs between the areas. ABRs send interarea routes between areas. However, the CPU intensive routing operation of recalculating the SPF algorithm is done only for routes within an area. A change in one area does not cause an SPF algorithm recalculation in other areas.

In Figure 14-5, assume that a link fails in Area 1. Only the routers in Area 1 exchange LSAs. D1, the ABR for Area 1, will send one update to Area 0 after Area 1 has converged on the new information.

The following list summarizes how multiarea OSPF improves OSPF performance:

• The smaller per-area LSDB requires less memory.

• Routers require fewer CPU cycles to process the smaller per-area LSDB with the SPF algorithm, reducing CPU overhead and improving convergence time.

• Changes in the network (for example, links failing and recovering) require SPF calculations only on routers connected to the area where the link changed state, reducing the number of routers that must rerun SPF.

• Less information must be advertised between areas, reducing the bandwidth required to send LSAs.

Technet24

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	1
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	21
Portable Command Guide	16

Day 13. Single-Area OSPF Implementation

CCNA 200-301 Exam Topics

• 3.4 Configure and verify single area OSPFv2

Key Topics

In Day 14, "OSPF Operation," we reviewed the basic implementation of multiarea OSPF. Today we review the basic implementation of OSPFv2 for single-area network designs. We review the commands to configure and verify the implementation.

Single-Area OSPFv2 Configuration

To review the single-area OSPFv2 configuration commands, we use the topology in Figure 13-1 and the addressing scheme in Table 13-1.

Technet24

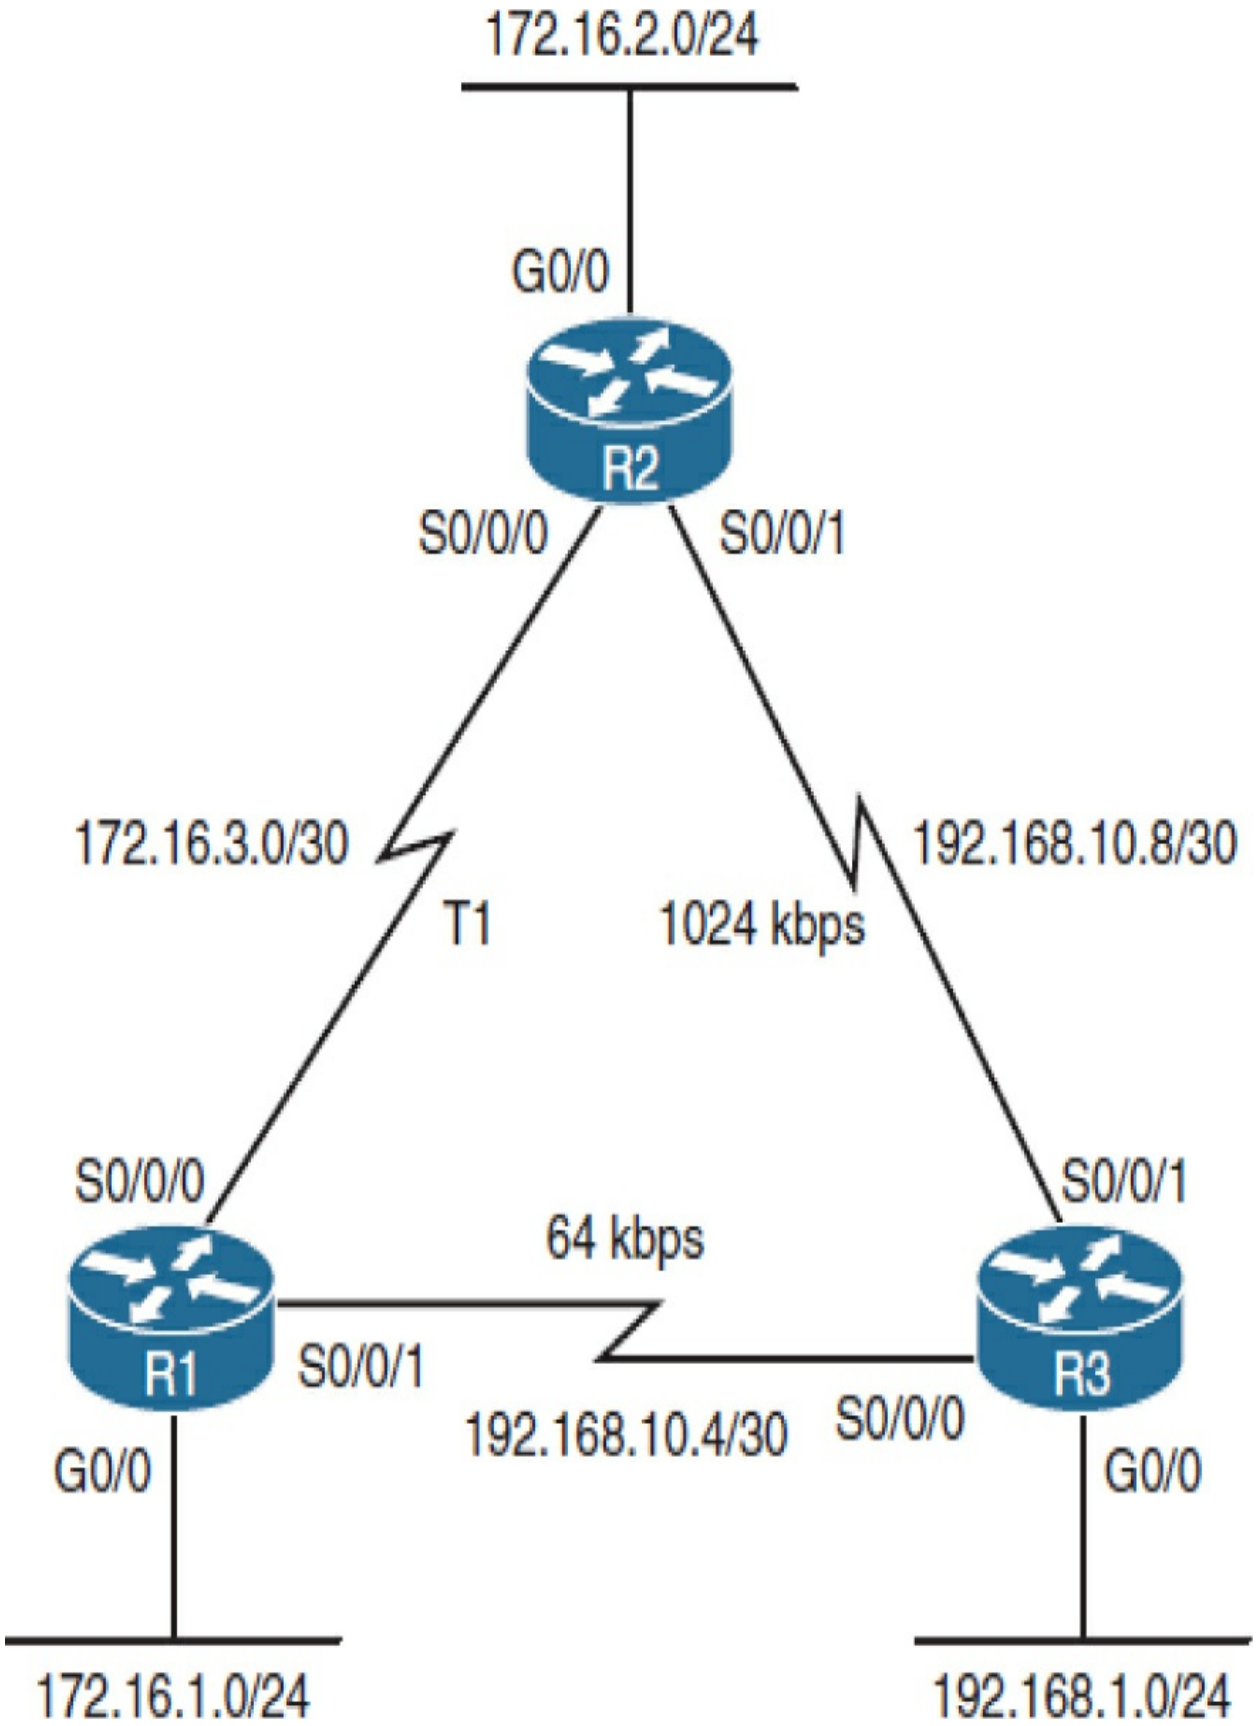
172.16.2.0/24
G0/0
R2
S0/0/0
S0/0/1
172.16.3.0/30
T1
1024 kbps
192.168.10.8/30
S0/0/0
S0/0/1
64 kbps
R1
S0/0/1
R3
192.168.10.4/30
S0/0/0
G0/0
G0/0
172.16.1.0/24
192.168.1.0/24

Figure 13-1 *OSPFv2 Configuration Topology*

Table 13-1 *Addressing Scheme for OSPFv2*

Device	Interface	IP Address	Subnet Mask
R1	G0/0	172.16.1.1	255.255.255.0
	S0/0/0	172.16.3.1	255.255.255.252
	S0/0/1	192.168.10.5	255.255.255.252
R2	G0/0	172.16.2.1	255.255.255.0
	S0/0/0	172.16.3.2	255.255.255.252
	S0/0/1	192.168.10.9	255.255.255.252
R3	G0/0	192.168.1.1	255.255.255.0
	S0/0/0	192.168.10.6	255.255.255.252
	S0/0/1	192.168.10.10	255.255.255.252

The router ospf Command

Technet24

OSPF is enabled with the **router ospf** *process-id* global configuration command:

```
R1(config)# router ospf process-id
```

process-id is a number between 1 and 65535 and is chosen by the network administrator. The process ID is locally significant. It does not have to match other OSPF routers to establish adjacencies with those neighbors. This differs from Enhanced Interior Gateway Routing Protocol (EIGRP). The EIGRP process ID and autonomous system number must match before two EIGRP neighbors can become adjacent.

For our review, we enable OSPF on all three routers using the same process ID, 10.

Router ID

The router ID plays an important role in OSPF, uniquely identifying each router in the OSPF routing domain. Cisco routers derive the router ID as follows:

The router uses the IP address configured with the OSPF **router-id** command.

If the router ID is not configured, the router chooses the highest IP address of any of its loopback interfaces.

If no loopback interfaces are configured, the router chooses the highest active IP address of any of its physical interfaces.

The router ID can be viewed with several commands, including **show ip ospf interfaces**, **show ip protocols**, and **show ip ospf**.

Because the network administrator can control the OSPF **router-id** command and because loopback interfaces clutter the routing table, it is a best practice to configure the **router-id** command. The **router-id** command accepts an IPv4 address as its only argument. Example 13-1 shows the router ID configurations for the routers in our topology.

Example 13-1 *Router ID Configurations*

```
R1(config-router)# router-id 1.1.1.1
```

```
R2(config-router)# router-id 2.2.2.2
```

```
R3(config-router)# router-id 3.3.3.3
```

The router ID is selected when OSPF is configured with its first OSPF **network** command, so the **router-id** command should already be configured. However, you can force OSPF to release its current ID and use the configured router ID by clearing the OSPF routing process:

Technet24

```
Router# clear ip ospf process
```

The network Command

The **network** command is used in router configuration mode:

```
Router(config-router)# network network-address wildcar
```

The OSPF **network** command uses a combination of *network-address* and *wildcard-mask*. The network address, along with the wildcard mask, specifies the interface or range of interfaces that will be enabled for OSPF using this **network** command.

The wildcard mask is customarily configured as the inverse of a subnet mask. For example, R1's Serial 0/0/0 interface is on the 172.16.3.0/30 network. The subnet mask for this interface is /30, or 255.255.255.252. The inverse of the subnet mask results in the wildcard mask 0.0.0.3.

area *area-id* refers to the OSPF area. An OSPF area is a group of routers that share link-state information. All OSPF routers in the same area must have the same link-state information in their link-state databases. Therefore, all the routers within the same OSPF area must be configured with the same area ID on all routers. By convention, the area ID is 0.

As an alternative, OSPFv2 can be enabled using the **network** *intf-ip-address*

0.0.0.0 area *area-id* router configuration mode command.

Example 13-2 shows the **network** commands for all three routers, enabling OSPF on all interfaces. For R1, the alternative interface IP address and quad zero wildcard mask is configured.

Example 13-2 *Configuring OSPF Networks*

```
R1(config)# router ospf 10
R1(config-router)# network 172.16.1.1 0.0.0.0 area 0
R1(config-router)# network 172.16.3.1 0.0.0.0 area 0
R1(config-router)# network 192.168.10.5 0.0.0.0 area 0
```

```
R2(config)# router ospf 10
R2(config-router)# network 172.16.2.0 0.0.0.255 area 0
R2(config-router)# network 172.16.3.0 0.0.0.3 area 0
R2(config-router)# network 192.168.10.8 0.0.0.3 area 0
```

```
R3(config)# router ospf 10
R3(config-router)# network 192.168.1.0 0.0.0.255 area
R3(config-router)# network 192.168.10.4 0.0.0.3 area 0
R3(config-router)# network 192.168.10.8 0.0.0.3 area 0
```

Passive Interfaces

Technet24

By default, OSPF messages are forwarded out all OSPF-enabled interfaces. However, these messages really need to be sent out only interfaces that connect to other OSPF-enabled routers. Sending out unneeded messages on a LAN affects the network in three ways:

• **Inefficient use of bandwidth:** Available bandwidth is consumed by transporting unnecessary messages.

• **Inefficient use of resources:** All devices on the LAN must process the message.

• **Increased security risk:** OSPF messages can be intercepted, and routing updates can be modified, corrupting the routing table.

Use the **passive-interface** command to prevent OSPF updates from being sent out unnecessary interfaces. For our topology in Figure 13-1, each router's GigabitEthernet 0/0 interface should be set to passive with the following command:

```
Router(config)# passive-interface gigabitethernet 0/0
```

As an alternative, you can make all interfaces passive by using the **passive-interface default** command. Then you can reenable interfaces that should not be passive by using the **no passive-interface** *interface* command.

Modifying the OSPF Metric

Cisco IOS Software uses the cumulative bandwidths of the outgoing interfaces from the router to the destination network as the cost value. At each router, the cost for an interface is calculated using the following formula:

```
Cisco IOS Cost for OSPF = 10^8/bandwidth in bps
```

In this calculation, the value 10^8 is known as the *reference bandwidth*. Table 13-2 shows the default OSPF costs using the default reference bandwidth for several types of interfaces.

Table 13-2 *Cisco Default OSPF Cost Values*

Technet24

Interface Type	10^8/bps = Cost	Cost
10 Gigabit Ethernet (10 Gbps)	10^8/10,000,000,000 bps = 1	1
Gigabit Ethernet (1 Gbps)	10^8/1,000,000,000 bps = 1	1
Fast Ethernet (100 Mbps)	10^8/100,000,000 bps = 1	1
Ethernet (10 Mbps)	10^8/10,000,000 bps = 10	10
T1 (1.544 Mbps)	10^8/1,544,000 bps = 64	64
128 kbps	10^8/128,000 bps = 781	781
64 kbps	10^8/64,000 bps = 1562	1562

In Table 13-2, 10GigE, Gigabit Ethernet, and Fast Ethernet all have the same cost. That is because the OSPF cost value must be an integer. This was not an issue before the introduction of gigabit and higher data rates.

However, today's networks are certainly running at gigabit speeds. Therefore, as a matter of policy, you should change the reference bandwidth to accommodate networks with links faster than 100,000,000 bps (100 Mbps). Use the following command to change the reference bandwidth:

```
Router(config-router)# auto-cost reference-bandwidth M
```

Because the value entered is in megabits per second, changing the reference bandwidth to 10000 ensures that all OSPF routers are ready to accurately calculate the cost for 10GigE networks. When used, this command should be entered on all routers so that the OSPF routing metric remains consistent. In fact, the Cisco IOS replies with the follow syslog message when you configure the **auto-cost reference-bandwidth** command:

```
% OSPF: Reference bandwidth is changed.
        Please ensure reference bandwidth is consiste
```

For our topology in Figure 13-1, we enter the commands shown in Example 13-3.

Example 13-3 *Changing the OSPF Reference Bandwidth*

```
R1(config-router)# auto-cost reference-bandwidth 10000
```

```
R2(config-router)# auto-cost reference-bandwidth 10000
```

Technet24

```
R3(config-router)# auto-cost reference-bandwidth 10000
```

Table 13-3 shows the modified cost values with the new reference bandwidth of 10,000,000,000 bps, or 10^{10}.

Table 13-3 *OSPF Cost Values with Modified Reference Bandwidth = 10000*

Interface Type	10^{10}/bps = Cost	Cost
10 Gigabit Ethernet (10 Gbps)	10^{10}/10,000,000,000 bps = 1	1
Gigabit Ethernet (1 Gbps)	10^{10}/1,000,000,000 bps = 1	10
Fast Ethernet (100 Mbps)	10^{10}/100,000,000 bps = 1	100
Ethernet (10 Mbps)	10^{10}/10,000,000 bps = 10	1000
T1 (1.544 Mbps)	10^{10}/1,544,000 bps = 64	6477
128 kbps	10^{10}/128,000 bps = 781	78125
64 kbps	10^{10}/64,000 bps = 1562	156250 (see the Note)

Note

Although the cost for a 64-kbps speed calculates to 156250, the maximum OSPF cost for a Cisco router interface is 65535.

But we are not done: We still have one more adjustment to make to ensure

Technet24

that OSPF is using accurate costs. On Cisco routers, the default bandwidth on most serial interfaces is set to T1 speed, or 1.544 Mbps. But in our topology in Figure 13-1, we have the following actual speeds:

• The link between R1 and R2 is running at 1544 kbps (default value).

• The link between R2 and R3 is running at 1024 kbps.

• The link between R1 and R3 is running at 64 kbps.

You can modify the OSPF metric in two ways:

• Use the **bandwidth** command to modify the bandwidth value the Cisco IOS Software uses in calculating the OSPF cost metric.

• Use the **ip ospf cost** command, which enables you to directly specify the cost of an interface.

An advantage of configuring a cost over setting the interface bandwidth is that the router does not have to calculate the metric when the cost is manually configured. Also, the **ip ospf cost** command is useful in multivendor environments, where non-Cisco routers can use a metric other than bandwidth to calculate the OSPF costs.

Table 13-4 shows the two alternatives that can be used in modifying the costs of the serial links in the topology in Figure 13-1. The right side of the figure shows the **ip ospf cost** command equivalents of the **bandwidth** commands

on the left.

Table 13-4 *Comparing the* bandwidth *and* ip ospf cost *Commands*

Adjusting the Interface Bandwidth	=	Manually Setting the OSPF Cost
R1(config)# **interface S0/0/1** R1(config-if)# **bandwidth 64**	=	R1(config)# **interface S0/0/1** R1(config-if)# **ip ospf cost 65535**
R2(config)# **interface S0/0/1** R2(config-if)# **bandwidth 1024**	=	R2(config)# **interface S0/0/1** R2(config-if)# **ip ospf cost 9765**
R3(config)# **interface S0/0/0** R3(config-if)# **bandwidth 64**	=	R3(config)# **interface S0/0/0** R3(config-if)# **ip ospf cost 65535**
R3(config)# **interface S0/0/1** R3(config-if)# **bandwidth 1024**	=	R3(config)# **interface S0/0/1** R3(config-if)# **ip ospf cost 9765**

Note

The 64-kbps interface is set to the maximum cost of 65535.

Verifying OSPFv2

Technet24

To verify any routing configuration, you will most likely depend on the **show ip interface brief**, **show ip route**, and **show ip protocols** commands. All the expected interfaces should be up and up, and they should be configured with the correct IP addresses. The routing table should have all the expected routes. The protocol status should show routing for all expected networks, and it should also show all expected routing sources. Example 13-4 shows R1's output from these three basic commands.

Example 13-4 *R1 Basic Routing Verification Commands*

```
R1# show ip route
<output omitted>

Gateway of last resort is not set

   172.16.0.0/16 is variably subnetted, 5 subnets, 3 m
C     172.16.1.0/24 is directly connected, GigabitEther
L       172.16.1.1/32 is directly connected, GigabitEth
O       172.16.2.0/24 [110/6576] via 172.16.3.2, 00:04:
C     172.16.3.0/30 is directly connected, Serial0/0/0
L     172.16.3.1/32 is directly connected, Serial0/0/0
O  192.168.1.0/24 [110/16341] via 172.16.3.2, 00:00:41
   192.168.10.0/24 is variably subnetted, 3 subnets, 2
C     192.168.10.4/30 is directly connected, Serial0/0
L     192.168.10.5/32 is directly connected, Serial0/0
O     192.168.10.8/30 [110/16241] via 172.16.3.2, 00:00
```

```
R1# show ip interface brief
Interface                  IP-Address OK? Method Status P
Embedded-Service-Engine0/0 unassigned YES unset  admin
GigabitEthernet0/0         172.16.1.1  YES  manual   up
GigabitEthernet0/1         unassigned   YES unset  admin
Serial0/0/0                172.16.3.1   YES manual   up
Serial0/0/1               192.168.10.5  YES manual   up
R1# show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is “ospf 10”
  Outgoing update filter list for all interfaces is no
  Incoming update filter list for all interfaces is no
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub
  Maximum path: 4
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    172.16.3.0 0.0.0.3 area 0
    192.168.10.4 0.0.0.3 area 0
  Passive Interface(s):
    GigabitEthernet0/0
  Routing Information Sources:
    Gateway Distance Last Update
    3.3.3.3 110 00:09:00
    2.2.2.2 110 00:09:00
  Distance: (default is 110)
```

Technet24

You can verify that expected neighbors have established adjacency with the **show ip ospf neighbor** command. Example 13-5 shows the neighbor tables for all three routers.

Example 13-5 *Verifying Neighbor Adjacency*

```
R1# show ip ospf neighbor

Neighbor ID Pri State      Dead Time   Address
3.3.3.3     0   FULL/ -     00:00:37    192.168.10.6
2.2.2.2     0   FULL/ -     00:00:37     172.16.3.2
```

```
R2# show ip ospf neighbor

Neighbor ID Pri State      Dead Time   Address
3.3.3.3      0  FULL/ -     00:00:38    192.168.10.10
1.1.1.1      0  FULL/ -     00:00:37    172.16.3.1
```

```
R3# show ip ospf neighbor

Neighbor ID Pri State     Dead Time   Address
2.2.2.2      0  FULL/ -   00:00:37    192.168.10.9
1.1.1.1      0  FULL/ -   00:00:30    192.168.10.5
```

For each neighbor, the **show ip ospf neighbor** command displays the following output:

• **Neighbor ID:** The router ID of the neighboring router.

• **Pri:** The OSPF priority of the interface. Example 13-5 shows 0 for this column because point-to-point links do not elect a DR or BDR.

• **State:** The OSPF state of the interface. FULL state means that the router's interface is fully adjacent with its neighbor, and they have identical OSPF link-state databases.

• **Dead Time:** The amount of time remaining that the router will wait to receive an OSPF hello packet from the neighbor before declaring the neighbor down. This value is reset when the interface receives a hello packet.

• **Address:** The IP address of the neighbor's interface to which this router is directly connected.

• **Interface:** The interface on which this router has formed adjacency with the neighbor.

The **show ip ospf** command in Example 13-6 for R1 can also be used to examine the OSPF process ID and router ID. In addition, this command displays the OSPF area information and the last time that the SPF algorithm was calculated.

Technet24

Example 13-6 *The* show ip ospf *Command*

```
R1# show ip ospf
Routing Process “ospf 10” with ID 1.1.1.1
  Start time: 00:29:52.316, Time elapsed: 00:45:15.760
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  Supports Link-local Signaling (LLS)
  Supports area transit capability
  Supports NSSA (compatible with RFC 3101)
  Event-log enabled, Maximum number of events: 1000, M
  Router is not originating router-LSAs with maximum m
  Initial SPF schedule delay 5000 msecs
  Minimum hold time between two consecutive SPFs 10000
  Maximum wait time between two consecutive SPFs 10000
  Incremental-SPF disabled
  Minimum LSA interval 5 secs
  Minimum LSA arrival 1000 msecs
  LSA group pacing timer 240 secs
  Interface flood pacing timer 33 msecs
  Retransmission pacing timer 66 msecs
  Number of external LSA 0. Checksum Sum 0x000000
  Number of opaque AS LSA 0. Checksum Sum 0x000000
  Number of DCbitless external and opaque AS LSA 0
  Number of DoNotAge external and opaque AS LSA 0
  Number of areas in this router is 1. 1 normal 0 stub
  Number of areas transit capable is 0
  External flood list length 0
```

```
  IETF NSF helper support enabled
  Cisco NSF helper support enabled
  Reference bandwidth unit is 10000 mbps
    Area BACKBONE(0)
        Number of interfaces in this area is 3
        Area has no authentication
        SPF algorithm last executed 00:18:32.788 ago
        SPF algorithm executed 7 times
        Area ranges are
        Number of LSA 3. Checksum Sum 0x01BB59
        Number of opaque link LSA 0. Checksum Sum 0x000
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

The quickest way to verify OSPF interface settings is to use the **show ip ospf interface brief** command. As shown in the output for R1 in Example 13-7, this command provides a detailed list for every OSPF-enabled interface. The command is also useful for quickly viewing the cost of each interface and determining whether the network statements were correctly configured.

Example 13-7 *The* **show ip ospf interface brief** *Command*

```
R1# show ip ospf interface brief
Interface  PID  Area   IP Address/Mask   Cost    State
```

Technet24

```
Se0/0/1    10    0     192.168.10.5/30   65535   P2P
Se0/0/0    10    0     172.16.3.1/30      6476   P2P
Gi0/0      10    0     172.16.1.1/24       100   DR
```

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	2
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	22
Portable Command Guide	16

Day 12. Fine-Tuning and Troubleshooting OSPF

CCNA 200-301 Exam Topics

• 3.4 Configure and verify single area OSPFv2

Key Topics

Today’s review focuses on fine-tuning and troubleshooting OSPFv2. Fine-tuning OSPF involves modifying timers, conducting DR/BDR elections, and propagating a default route. We also turn our focus to troubleshooting the OSPF process.

OSPFv2 Configuration Example

To fine-tune OSPFv2, we use the topology in Figure 12-1 and the addressing scheme in Table 12-1.

Technet24

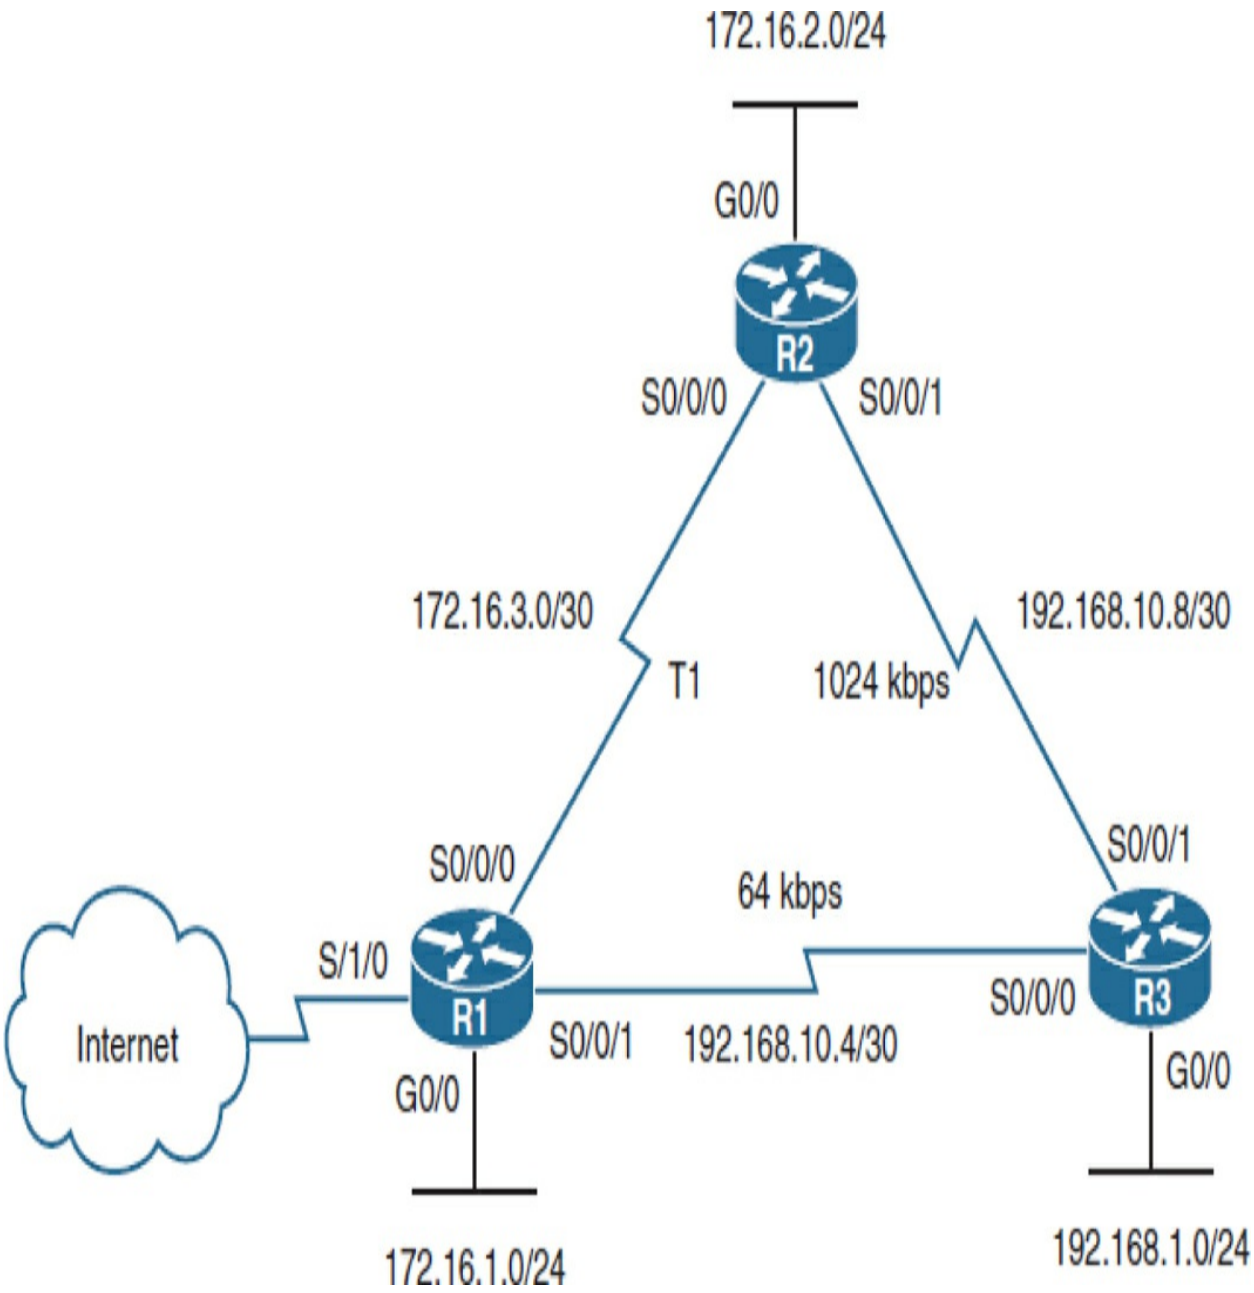

Figure 12-1 *OSPFv2 Configuration Topology*

Table 12-1 *Addressing Scheme for OSPFv2*

Device	Interface	IP Address	Subnet Mask
R1	G0/0	172.16.1.1	255.255.255.0
	S0/0/0	172.16.3.1	255.255.255.252
	S0/0/1	192.168.10.5	255.255.255.252
	S0/1/0	209.165.200.226	255.255.255.224
R2	G0/0	172.16.2.1	255.255.255.0
	S0/0/0	172.16.3.2	255.255.255.252
	S0/0/1	192.168.10.9	255.255.255.252
R3	G0/0	192.168.1.1	255.255.255.0
	S0/0/0	192.168.10.6	255.255.255.252
	S0/0/1	192.168.10.10	255.255.255.252

Example 12-1 shows the network commands for all three routers to enable OSPFv2 on all interfaces.

Technet24

Example 12-1 *Configuring OSPF Networks*

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# network 172.16.1.0 0.0.0.255 area 0
R1(config-router)# network 172.16.3.0 0.0.0.3 area 0
R1(config-router)# network 192.168.10.4 0.0.0.3 area 0
R1(config-router)# passive-interface g0/0
R1(config-router)# auto-cost reference-bandwidth 10000
R1(config-router)# interface S0/0/1
R1(config-if)# bandwidth 64
```

```
R2(config)# router ospf 10
R2(config-router)# router-id 2.2.2.2
R2(config-router)# network 172.16.2.0 0.0.0.255 area 0
R2(config-router)# network 172.16.3.0 0.0.0.3 area 0
R2(config-router)# network 192.168.10.8 0.0.0.3 area 0
R2(config-router)# passive-interface g0/0
R2(config-router)# auto-cost reference-bandwidth 10000
R2(config-router)# interface S0/0/1
R2(config-if)# bandwidth 1024
```

```
R3(config)# router ospf 10
R3(config-router)# router-id 3.3.3.3
R3(config-router)# network 192.168.1.0 0.0.0.255 area
```

```
R3(config-router)# network 192.168.10.4 0.0.0.3 area 0
R3(config-router)# network 192.168.10.8 0.0.0.3 area 0
R3(config-router)# passive-interface g0/0
R3(config-router)# auto-cost reference-bandwidth 10000
R3(config-router)# interface S0/0/0
R3(config-if)# bandwidth 64
R3(config-if)# interface S0/0/1
R3(config-if)# bandwidth 1024
```

Modifying OSPFv2

This section reviews concepts related to and commands for redistributing a default route, tuning OSPF interfaces, and manipulating the designated router/backup designated router (DR/BDR) election process.

Redistributing a Default Route

In Figure 12-1, R1 has a link to the Internet that makes R1 an autonomous system boundary router (ASBR). We therefore configure a default route to the Internet and redistribute the default static route to R2 and R3 with the **default-information originate** command, as in Example 12-2.

Example 12-2 *ASBR Static Default Route Configuration*

```
R1(config)# ip route 0.0.0.0 0.0.0.0 Serial 0/1/0
R1(config)# router ospf 10
```

Technet24

```
R1(config-router)# default-information originate
```

Both R2 and R3 should now have default routes identified with the O*E2 code, as in Example 12-3.

Example 12-3 *R2 and R3 OSPF Routes with Default Route*

```
R2# show ip route ospf
     172.16.0.0/16 is variably subnetted, 5 subnets, 3
O       172.16.1.0 [110/6477] via 172.16.3.1, 00:02:45
O    192.168.1.0 [110/6486] via 192.168.10.10, 00:00:5
     192.168.10.0/24 is variably subnetted, 3 subnets,
O       192.168.10.4 [110/12952] via 192.168.10.10, 00
O*E2 0.0.0.0/0 [110/1] via 172.16.3.1, 00:00:09, Seria
```

```
R3# show ip route ospf
 172.16.0.0/16 is variably subnetted, 3 subnets, 2 mas
    172.16.0.0/16 is variably subnetted, 3 subnets, 2
O     172.16.1.0 [110/6477] via 192.168.10.5, 00:26:01
O     172.16.2.0 [110/6486] via 192.168.10.9, 00:26:01
O     172.16.3.0 [110/6540] via 192.168.10.5, 00:26:01
O*E2 0.0.0.0/0 [110/1] via 192.168.10.9, 00:01:19, Ser
```

Modifying Hello and Dead Intervals

The default hello interval on multiaccess and point-to-point networks is 10 seconds. Nonbroadcast multiaccess (NBMA) networks default to a 30-second hello interval. The default dead interval is four times the hello interval.

It might be desirable to change the OSPF timers so that routers detect network failures in less time. Doing this increases traffic, but sometimes a need for quick convergence outweighs the extra traffic.

You can modify OSPF hello and dead intervals manually by using the following interface commands:

```
Router(config-if)# ip ospf hello-interval seconds
Router(config-if)# ip ospf dead-interval seconds
```

Although the dead interval defaults to four times the hello interval and does not have to be explicitly configured, it is a good practice to document the new dead interval in the configuration. Example 12-4 shows the hello interval and dead interval modified to 5 seconds and 20 seconds, respectively, on the Serial 0/0/0 interface for R1.

Example 12-4 *Modifying Hello and Dead Intervals on R1*

```
R1(config)# interface serial 0/0/0
R1(config-if)# ip ospf hello-interval 5
R1(config-if)# ip ospf dead-interval 20
R1(config-if)# end
```

Technet24

Remember that, unlike with Enhanced Interior Gateway Routing Protocol (EIGRP), with OSPF hello and dead intervals must be equivalent between neighbors. Therefore, R2 should be configured with the same intervals as R1.

OSPF Network Types

OSPF defines five network types:

• **Point-to-point:** Two routers interconnected over a common link. No other routers are on the link. This is often the configuration in WAN links.

• **Broadcast multiaccess:** Multiple routers interconnected over an Ethernet network.

• **NBMA:** Multiple routers interconnected in a network that does not allow broadcasts, such as Frame Relay.

• **Point-to-multipoint:** Multiple routers interconnected in a hub-and-spoke topology over an NBMA network. Often used to connect branch sites (spokes) to a central site (hub).

• **Virtual links:** Special OSPF network used to interconnect distant OSPF areas to the backbone area.

Multiaccess networks create two challenges for OSPF regarding the flooding of LSAs:

• **Creation of multiple adjacencies:** Ethernet networks can potentially interconnect many OSPF routers over a common link. Using the formula $n(n – 1) / 2$, where n equals the number of routers, 5 routers would require 10 separate neighbor adjacencies; 10 routers would require 45.

• **Extensive flooding of LSAs:** Link-state routers flood their link-state packets when OSPF is initialized or when the topology changes. This flooding can become excessive without a mechanism to reduce the number of adjacencies.

DR/BDR Election

The solution to managing the number of adjacencies and the flooding of LSAs on a multiaccess network is the designated router (DR). To reduce the amount of OSPF traffic on multiaccess networks, OSPF elects a DR and backup DR (BDR). The DR is responsible for updating all other OSPF routers when a change occurs in the multiaccess network. The BDR monitors the DR and takes over as DR if the current DR fails.

The following criteria are used to elect the DR and BDR:

• The DR is the router with the highest OSPF interface priority.

• The BDR is the router with the second-highest OSPF interface priority.

• If OSPF interface priorities are equal, the highest router ID breaks the tie.

Technet24

When the DR is elected, it remains the DR until one of the following conditions occurs:

• The DR fails.

• The OSPF process on the DR fails.

• The multiaccess interface on the DR fails.

If the DR fails, the BDR assumes the role of DR, and an election is held to choose a new BDR. If a new router enters the network after the DR and BDR have been elected, it will not become the DR or the BDR even if it has a higher OSPF interface priority or router ID than the current DR or BDR. The new router can be elected the BDR if the current DR or BDR fails. If the current DR fails, the BDR becomes the DR, and the new router can be elected the new BDR.

Without additional configuration, you can control the routers that win the DR and BDR elections by doing either of the following:

• Boot the DR first, followed by the BDR, and then boot all other routers.

• Shut down the interface on all routers and then issue **no shutdown** on the DR, then the BDR, and then all other routers.

The recommended way to control DR/BDR elections, however, is to change the interface priority.

Controlling the DR/BDR Election

Because the DR becomes the focal point for the collection and distribution of LSAs in a multiaccess network, this router must have sufficient CPU and memory capacity to handle the responsibility. Instead of relying on the router ID to decide which routers are elected the DR and BDR, it is better to control the election of these routers with the **ip ospf priority** interface command:

```
Router(config-if)# ip ospf priority {0 - 255}
```

The priority value defaults to 1 for all router interfaces, which means the router ID determines the DR and BDR. If you change the default value from 1 to a higher value, however, the router with the highest priority becomes the DR, and the router with the next highest priority becomes the BDR. A value of 0 makes the router ineligible to become a DR or BDR.

All the routers in Figure 12-2 booted at the same time with a complete OSPF configuration. In such a situation, R3 is elected the DR, and R2 is elected the BDR, based on the highest router IDs, as you can see in the output for the neighbor table on R1 in Example 12-5.

Technet24

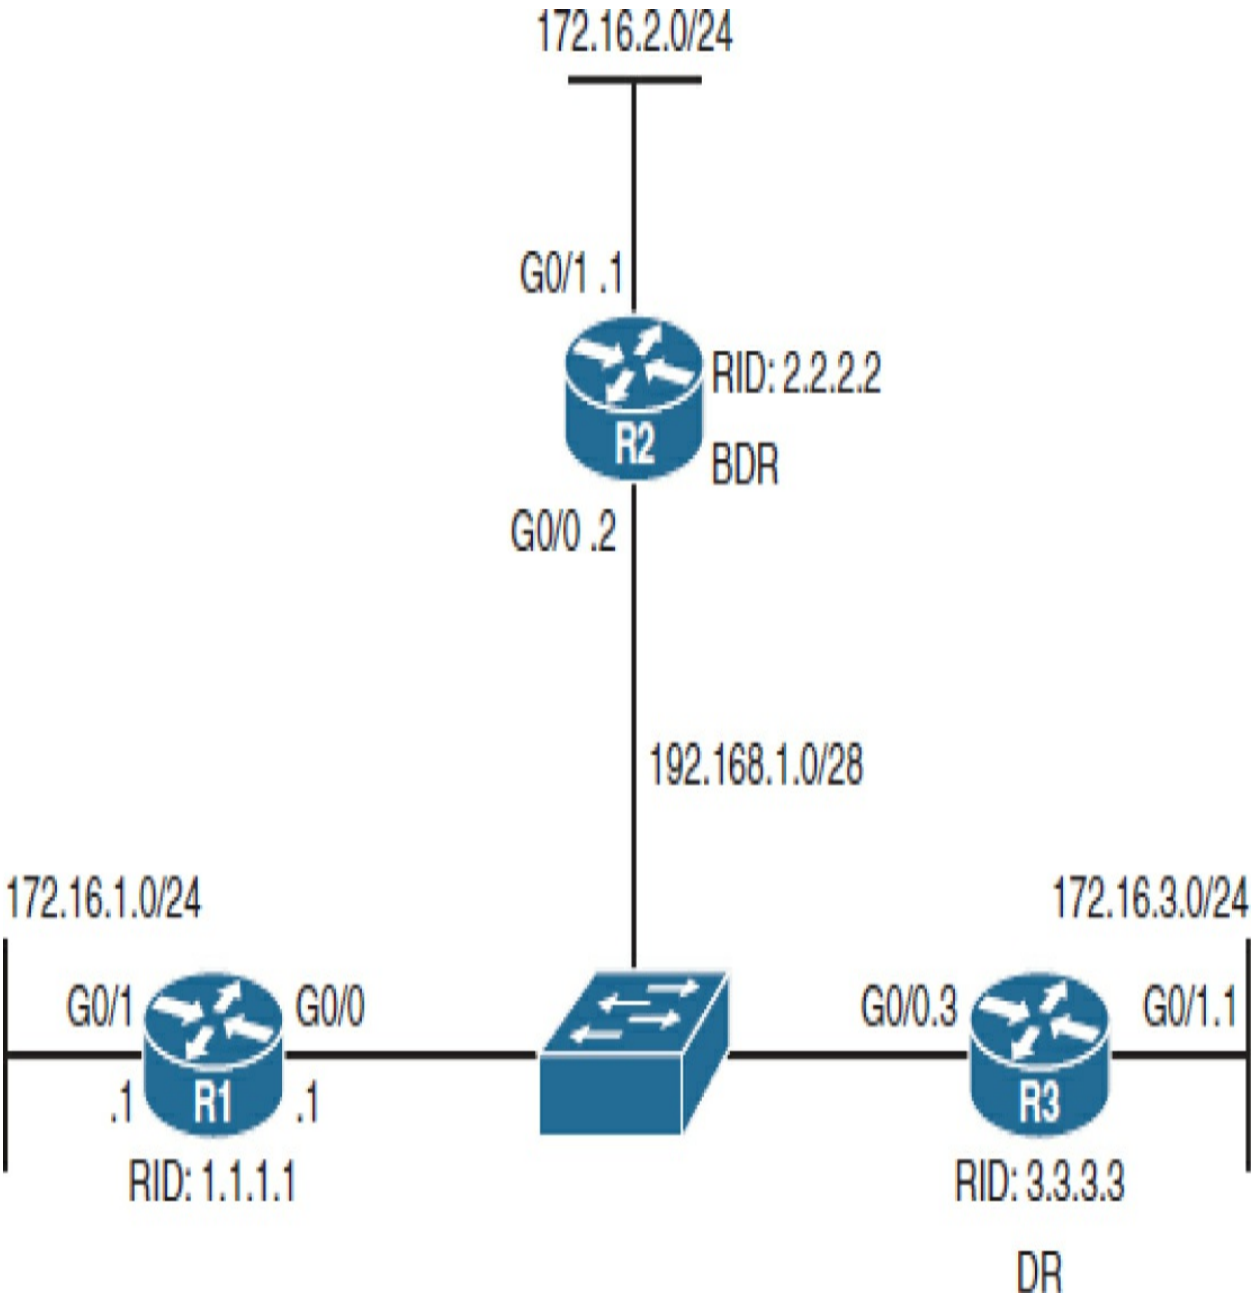

Figure 12-2 *Multiaccess Topology*

Example 12-5 *Verifying the DR and BDR*

```
R1# show ip ospf neighbor
Neighbor ID   Pri State      Dead Time Address       I
```

```
2.2.2.2         1   FULL/BDR    00:00:32 192.168.1.2    G
3.3.3.3         1   FULL/DR     00:00:38 192.168.1.3    G
R1#
```

Assume that R1 is the better candidate to be DR and that R2 should be BDR. Example 12-6 shows a way to control the DR/BDR election in the topology in Figure 12-2.

Example 12-6 *Modifying the OSPF Interface Priority*

```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ip ospf priority 200
```

```
R2(config)# interface gigabitethernet 0/0
R2(config-if)# ip ospf priority 100
```

Notice that we changed both routers. Although R2 was the BDR without doing anything, it would lose this role to R3 if we had not configured R2's priority to be higher than the default.

Before R1 can become DR, the OSPF process needs to restart. We can make this happen by shutting down the interfaces or simply by entering the **clear ip ospf process** command in privileged EXEC mode, as in Example 12-7. The neighbor table on R3 shows that R1 is now the DR, and R2 is the BDR.

Technet24

Example 12-7 *Restarting the OSPF Process and Verifying New DR and BDR*

```
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R1#
```

```
R2# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R2#
```

```
R3# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R3#
```

```
R3# show ip ospf neighbor

Neighbor ID Pri  State      Dead Time  Address      Ir
2.2.2.2     100  FULL/BDR   00:00:38   192.168.1.2  Gi
1.1.1.1     200  FULL/DR    00:00:30   192.168.1.1  Gi
R3#
```

Troubleshooting OSPF

Understanding how OSPF operates is fundamental to troubleshooting any

OSPF issues. Key to this understanding is the concept of the states OSPF transitions through on its way to adjacency with a neighbor.

OSPF States

Figure 12-3 lists the OSPF states. When troubleshooting OSPF neighbors, be aware that the FULL and TWO-WAY states are normal. All other states are transitory.

Technet24

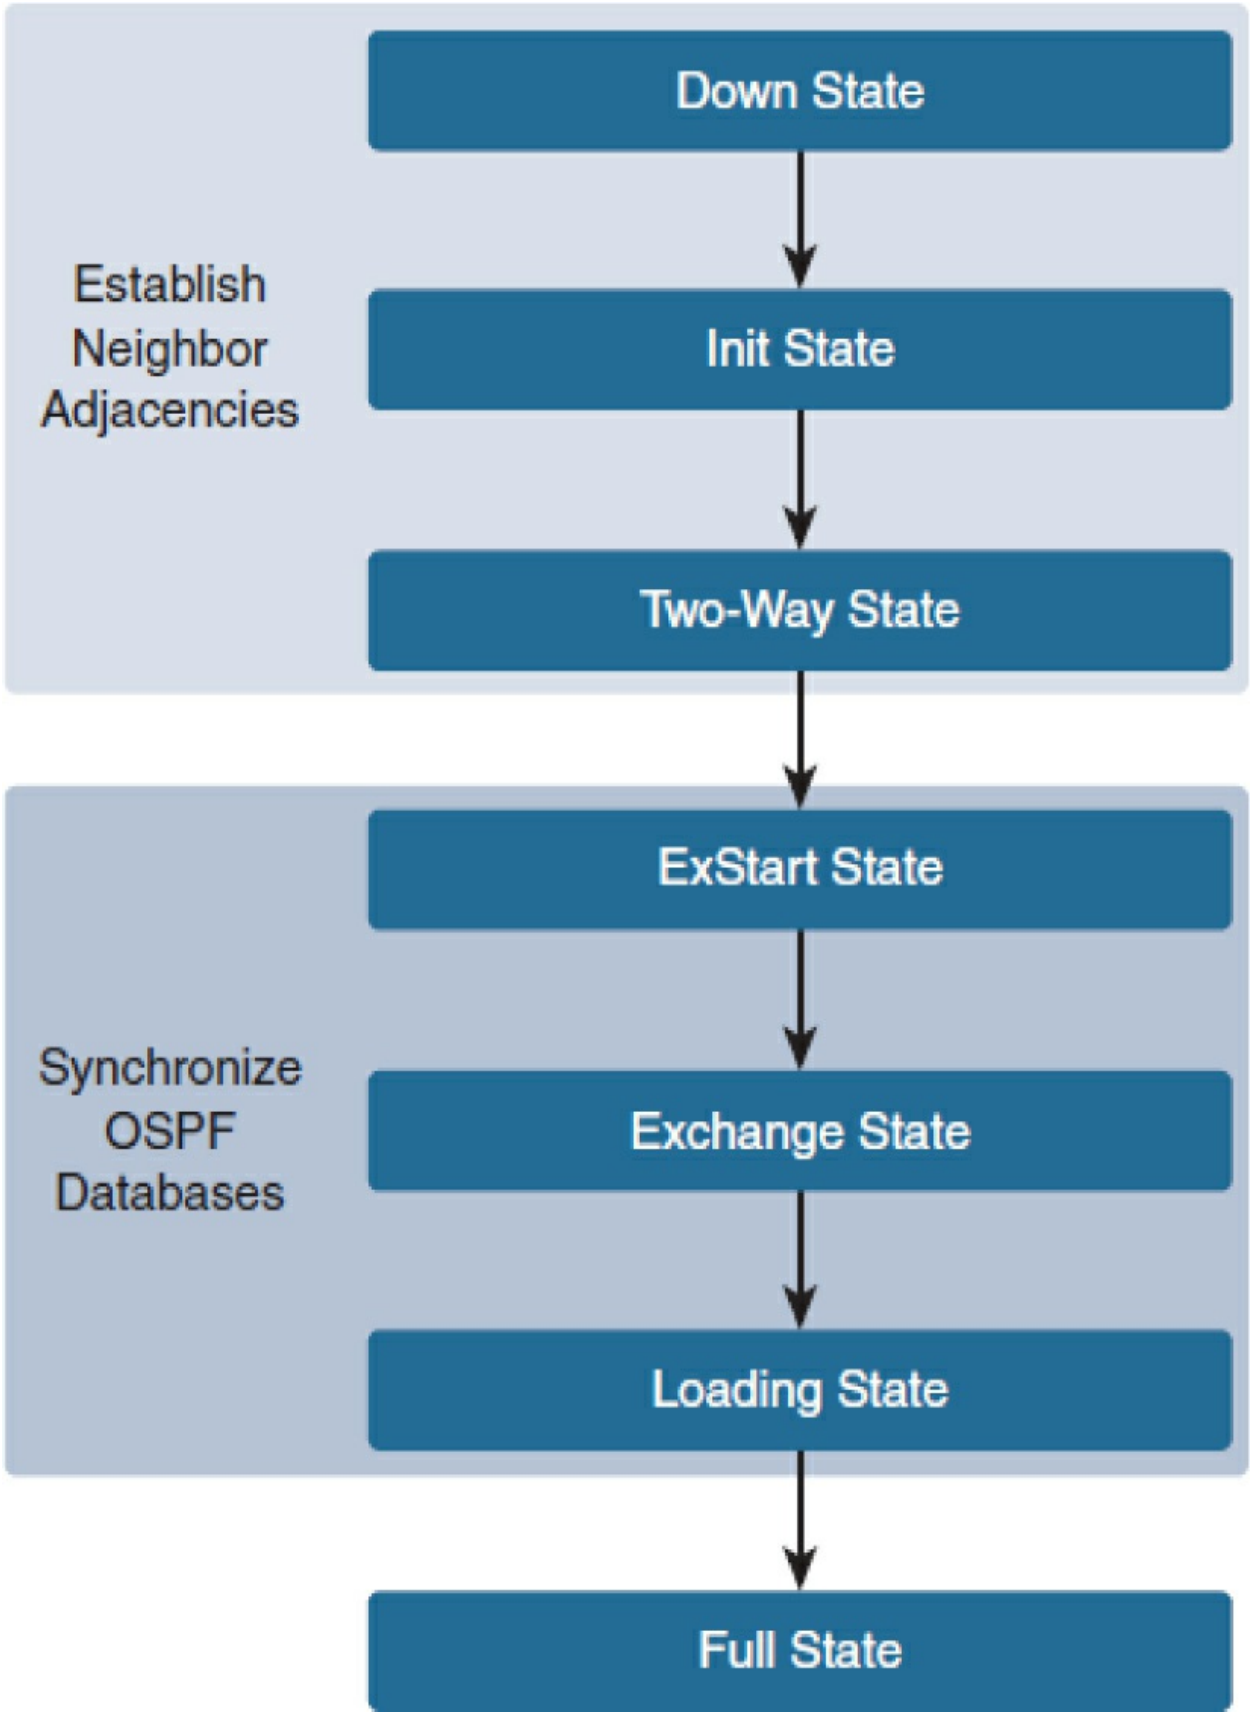
Down State
Establish Neighbor Adjacencies
Init State
Two-Way State
ExStart State
Synchronize OSPF Databases
Exchange State
Loading State
Full State

Figure 12-3 *Transitioning Through the OSPF States*

OSPF Adjacency

Lack of adjacency is a common issue in OSPF troubleshooting because the two OSPF neighbors must agree on several settings. OSPF adjacencies do not form for several reasons:

• The interfaces are not on the same network.

• OSPF network types do not match.

• OSPF hello or dead timers do not match.

• The interface to the neighbor is incorrectly configured as passive.

• An OSPF **network** command is missing or incorrect.

• Authentication is misconfigured.

OSPF Troubleshooting Commands

When trying to isolate an OSPFv2 routing issue, the following commands are useful:

• **show ip protocols:** Verifies vital OSPF configuration information, including the OSPF process ID, the router ID, networks the router is advertising, neighbors the router is receiving updates from, and the default

Technet24

administrative distance, which is 110 for OSPF.

• **show ip ospf neighbor:** Verifies that the router has formed an adjacency with its neighboring routers.

• **show ip ospf interface:** Displays the OSPF parameters configured on an interface, such as the OSPF process ID, area, cost, and timer intervals.

• **show ip ospf:** Examines the OSPF process ID and router ID. This command also displays the OSPF area information and the last time the SPF algorithm was calculated.

• **show ip route ospf:** Displays only the OSPF learned routes in the routing table.

• **clear ip ospf process:** Resets the OSPFv2 neighbor adjacencies.

Figure 12-4 illustrates a method for using these commands in a systematic way.

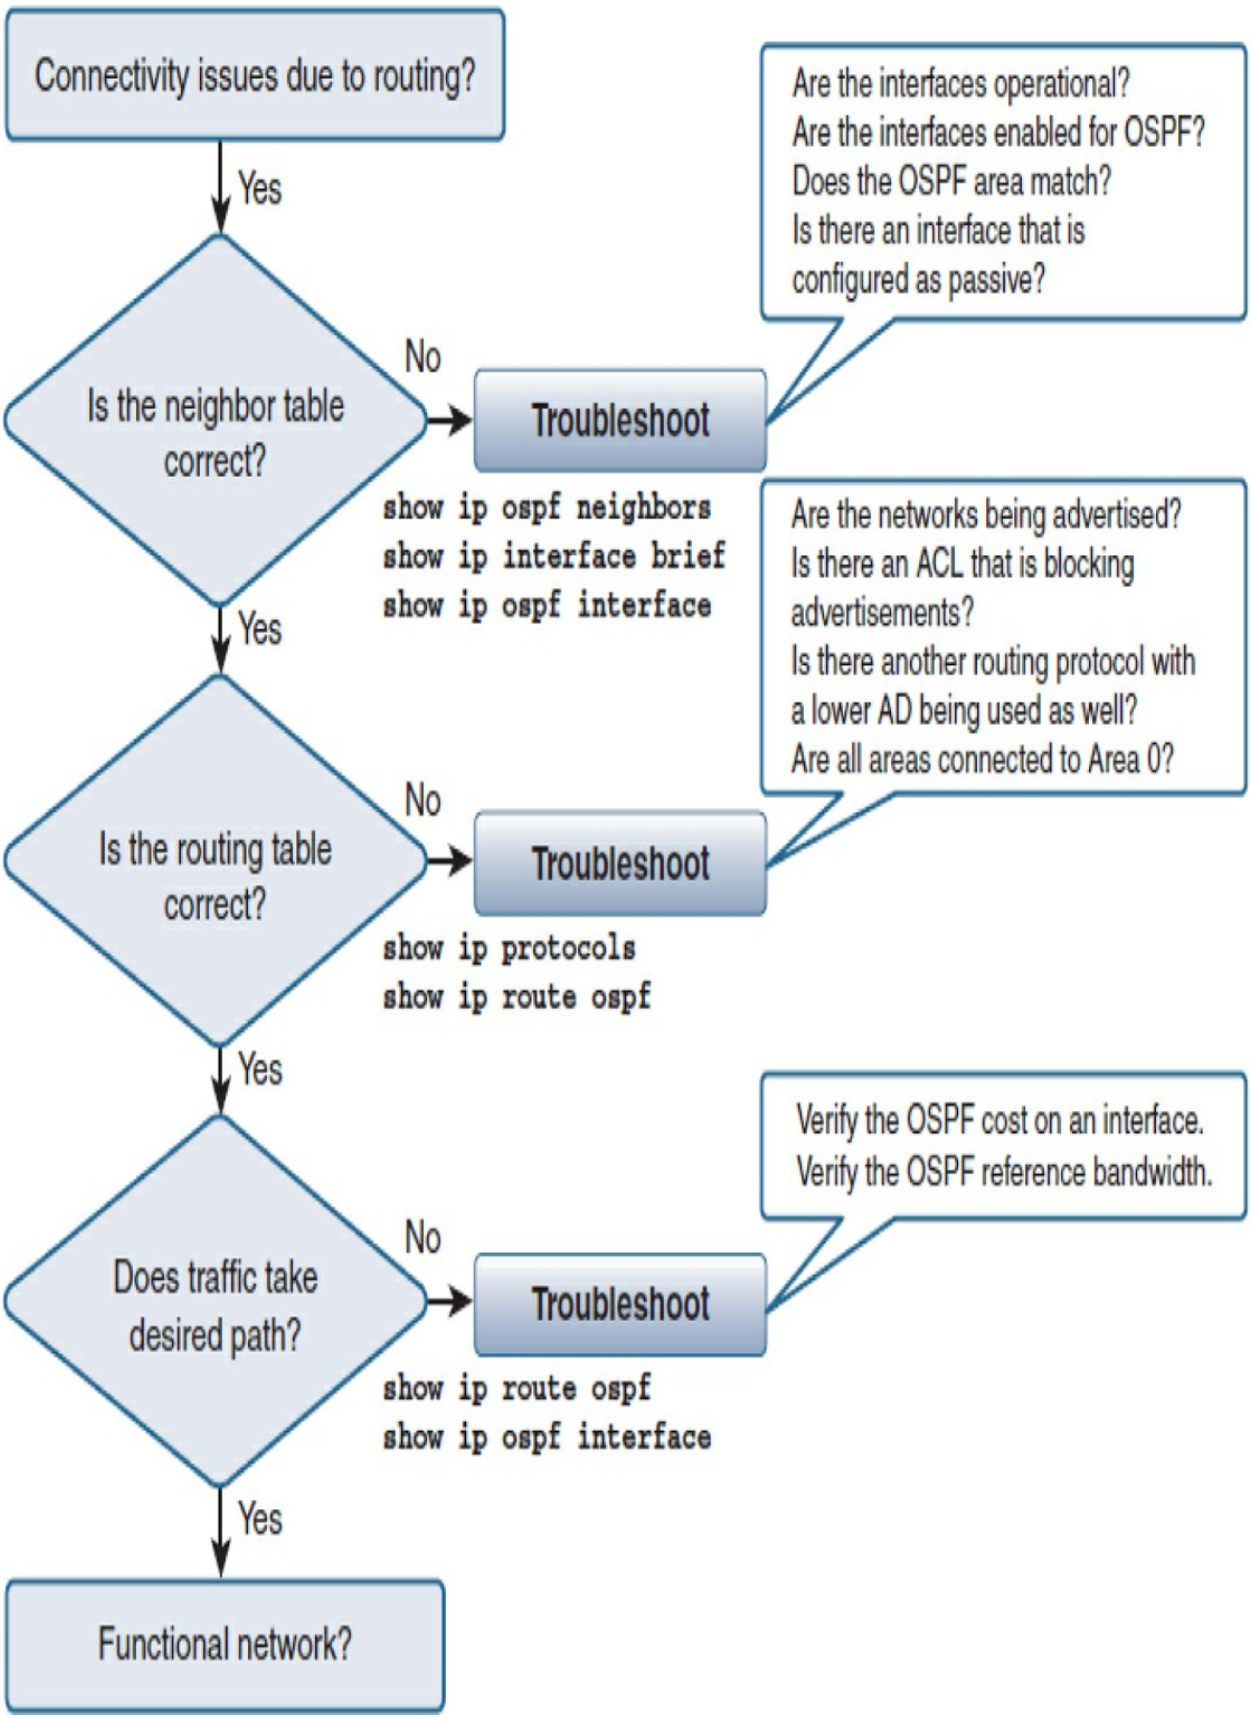
Connectivity issues due to routing?
Yes
Is the neighbor table correct?
No
Troubleshoot
Are the interfaces operational?
Are the interfaces enabled for OSPF?
Does the OSPF area match?
Is there an interface that is configured as passive?
show ip ospf neighbors
show ip interface brief
show ip ospf interface
Yes
Is the routing table correct?
No
Troubleshoot
Are the networks being advertised?
Is there an ACL that is blocking advertisements?
Is there another routing protocol with a lower AD being used as well?
Are all areas connected to Area 0?
show ip protocols
show ip route ospf
Yes
Does traffic take desired path?
No
Troubleshoot
Verify the OSPF cost on an interface.
Verify the OSPF reference bandwidth.
show ip route ospf
show ip ospf interface
Yes
Functional network?

Technet24

Figure 12-4 *Systematic Method for Troubleshooting OSPFv2*

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	2
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	22–24
Portable Command Guide	16

Day 11. Network Security Concepts

CCNA 200-301 Exam Topics

• 5.1 Define key security concepts (threats, vulnerabilities, exploits, and mitigation techniques)

• 5.2 Describe security program elements (user awareness, training, and physical access control)

• 5.4 Describe security password policy elements such as management, complexity and password alternatives (multifactor authentication, certificates, and biometrics)

Key Topics

Wired and wireless computer networks are essential to everyday activities. Individuals and organizations depend on their computers and networks. Intrusion by an unauthorized person can result in costly network outages and loss of work. Attacks on a network can be devastating and can result in lost time and money due to damage or theft of important information or assets. Access management is crucial for defending against some of the most common attacks that target weak authentication methods and passwords. Today we review security fundamentals including threats, vulnerabilities, and attacks. Then we spend some time reviewing best practices for passwords and their alternatives.

Technet24

Security Fundamentals

Cyber criminals now have the expertise and tools needed to take down critical infrastructure and systems. Specific terminology is used to describe their tools and attacks.

Security Terms

Assets must be identified and protected. Vulnerabilities must be addressed before they become threats and are exploited. Mitigation techniques are required before, during, and after an attack. Review the security terms in Table 11-1.

Table 11-1 *Security Terms*

Term	Description
Assets	Anything of value to an organization, including people, equipment, resources, and data.
Vulnerability	A weakness in a system or its design that could be exploited by a threat.
Threat	A potential danger to a company's assets, data, or network functionality.
Exploit	A mechanism that takes advantage of a vulnerability.
Mitigation	The process of taking countermeasures to reduce the likelihood or severity of a potential threat or risk.
Risk	The likelihood of a threat exploiting the vulnerability of an asset, with the aim of negatively affecting an organization.

Attack Vectors and Data Exfiltration

An *attack vector* is a path by which a threat actor can gain access to a server, host, or network. Attack vectors originate outside or inside a network. For example, threat actors may target a network through the Internet to disrupt network operations and create a denial of service (DoS) attack. An internal

Technet24

user, such as an employee, might accidentally or intentionally disrupt the network or steal confidential data. Internal threats have the potential to cause greater damage than external threats because internal users have direct access to the building and its infrastructure devices.

Employees may also have knowledge of the corporate network, its resources, and its confidential data. Data loss or data exfiltration occurs when data is intentionally or unintentionally lost, stolen, or leaked to the outside world. Network security professionals must protect the organization's data. Various data loss prevention (DLP) controls must be implemented, combining strategic, operational, and tactical measures. Common data loss vectors are shown in Table 11-2.

Table 11-2 *Data Loss Vectors*

Vector	Description
Email/social networking	Intercepted email or IM messages could be captured and reveal confidential information.
Unencrypted devices	If data is not stored using an encryption algorithm, the thief may be able to retrieve valuable confidential data.
Cloud storage devices	Sensitive data can be lost if access to the cloud is compromised due to weak security settings.
Removable media	An employee could perform an unauthorized transfer of data to a USB drive, or a USB drive containing valuable corporate data could be lost.
Hard copy	Confidential data should be shredded when no longer required.
Improper access control	Passwords or weak passwords that have been compromised can provide a threat actor with easy access to corporate data.

Penetration Testing Tools

Technet24

To validate the security of a network and its systems, many network penetration testing tools have been developed (see Table 11-3). Unfortunately, threat actors can also use many of these tools for exploitation.

Table 11-3 *Types of Penetration Testing Tools*

Tool	Description
Password crackers	Password cracking tools, often referred to as password recovery tools, can be used to crack, or recover, a password. Password crackers repeatedly make guesses in order to crack the password.
Wireless hacking tools	Wireless hacking tools are used to intentionally hack into a wireless network to detect security vulnerabilities.
Network scanning and hacking tools	Network scanning tools are used to probe network devices, servers, and hosts for open TCP or UDP ports.
Packet crafting tools	These tools are used to probe and test a firewall's robustness, using specially crafted forged packets.
Packet sniffers	These tools are used to capture and analyze packets in traditional Ethernet LANs or WLANs.
Rootkit detectors	White hat hackers use these directory and file integrity checkers to detect installed rootkits.

Technet24

Forensic tools	White hat hackers use these tools to sniff out any trace of evidence existing in a computer.
Debuggers	Black hat hackers use these tools to reverse engineer binary files when writing exploits. White hat hackers also use them for analyzing malware.
Hacking operating systems	These are specially designed operating systems preloaded with tools optimized for hacking.
Encryption tools	Encryption tools use algorithm schemes to encode data to prevent unauthorized access to the encrypted data.
Vulnerability exploitation tools	These tools identify whether a remote host is vulnerable to security attack.
Vulnerability scanners	These tools scan a network or system to identify open ports. They can also be used to scan for known vulnerabilities and scan virtual machines (VMs), devices brought to work by individuals in a bring-your-own-device (BYOD) situation, and client databases.

Attack Types

Threat actors can use tools to create a variety of attacks. Table 11-4 describes common types of attacks.

Table 11-4 *Common Types of Attacks*

Attack Type	Description
Eavesdropping attack	A threat actor captures and "listens" to network traffic. This attack is also referred to as *sniffing* or *snooping*.
Data modification attack	If threat actors have captured enterprise traffic, they can alter the data in a packet without the knowledge of the sender or receiver.
IP address spoofing attack	A threat actor constructs an IP packet that appears to originate from a valid address inside the corporate intranet.
Password-based attack	A threat actor who discovers a valid user account has the same rights as the real user. A threat actor can use a valid account to obtain lists of other users or network information, change server and network configurations, and modify, reroute, or delete data.

Technet24

Denial of service (DoS) attack	A DoS attack prevents normal use of a computer or network by valid users. A DoS attack can flood a computer or an entire network with traffic until a shutdown occurs because of the overload. A DoS attack can also block traffic, which results in a loss of access to network resources by authorized users.
Man-in-the-middle attack	This type of attack occurs when threat actors have positioned themselves between a source and destination. The threat actors can actively monitor, capture, and control the communication transparently.
Compromised-key attack	If a threat actor obtains a secret key, that key is referred to as a *compromised key*. A compromised key can be used to gain access to secured communication without the sender or receiver being aware of the attack.
Sniffer attack	A sniffer is an application or a device that can read, monitor, and capture network data exchanges and read network packets. If packets are not encrypted, a sniffer provides a full view of the data inside the packet.

Types of Malware

Malware, which is short for *malicious software*, is code or software specifically designed to damage, disrupt, steal, or inflict “bad” or illegitimate actions on data, hosts, or networks. Viruses, worms, and Trojan horses are types of malware:

• A *worm* executes arbitrary code and installs copies of itself in the memory of the infected computer. The main purpose of a worm is to automatically replicate itself and spread across the network from system to system.

• A *virus* is malicious software that executes a specific, unwanted, often harmful function on a computer.

• A *Trojan horse* is a non-self-replicating type of malware. It often contains malicious code that is designed to look like something else, such as a legitimate application or file. When an infected application or file is downloaded and opened, the Trojan horse can attack the end device from within.

Table 11-5 describes some other types of malware.

Table 11-5 *Other Types of Malware*

Technet24

Malware	Description
Adware	Adware is usually distributed through downloads of software. Adware can display unsolicited advertising using popup web browser windows or new toolbars, or it can unexpectedly redirect a user from a web page to a different website. Popup windows may be difficult to control as new windows can pop up faster than the user can close them.
Ransomware	Ransomware typically denies a user access to their files by encrypting the files and then displaying a message demanding a ransom for the decryption key. Users without up-to-date backups must pay the ransom to decrypt their files. Payment is usually made using wire transfer or cryptocurrency such as bitcoin.
Rootkit	Threat actors use rootkits to gain administrator account–level access to a computer. They are very difficult to detect because they can alter firewall, antivirus protection, system files, and even OS commands to conceal their presence. A rootkit can provide a backdoor to threat actors, giving them access to the PC and allowing them to upload files and install new software to be used in a distributed DoS (DDoS) attack. Special rootkit removal tools must be used to remove them, or a complete OS reinstallation may be required.
Spyware	Spyware is similar to adware but is used to gather information about the user and send it to threat actors without the user's consent. Spyware can be a low threat, gathering browsing data, or it can be a high threat, capturing personal and financial information.

Network Attacks

Network attacks include reconnaissance attacks, access attacks, DoS attacks, social engineering attacks, and attacks to exploit the vulnerabilities of the TCP/IP protocol suite.

Reconnaissance Attacks

Reconnaissance is information gathering. Threat actors use reconnaissance (or recon) attacks to do unauthorized discovery and mapping of systems, services, or vulnerabilities. Recon attacks precede access attacks or DoS attacks. Table 11-6 describes some common reconnaissance attack techniques.

Table 11-6 *Reconnaissance Attack Techniques*

Technet24

Technique	Description
Perform an information query of a target	The threat actor looks for initial information about a target. Various tools can be used, including a Google search, the organization's website, and whois.
Initiate a ping sweep of the target network	The information query usually reveals the target's network address. The threat actor can then initiate a ping sweep to determine which IP addresses are active.
Initiate a port scan of active IP addresses	A port scan can be used to determine which ports or services are available. Examples of port scanners include Nmap, SuperScan, Angry IP Scanner, and NetScanTools.
Run vulnerability scanners	A vulnerability scanner can query the identified ports to determine the type and version of the application and operating system running on the host. Examples of such tools include Nipper, Secunia PSI, Core Impact, Nessus v6, SAINT, and Open VAS.
Run exploitation tools	The threat actor attempts to discover vulnerable services that can be exploited. A variety of vulnerability exploitation tools exist, including Metasploit, Core Impact, sqlmap, Social-Engineer Toolkit, and Netsparker.

Access Attacks

The purpose of access attacks is to gain entry to web accounts, confidential databases, and other sensitive information. Threat actors use access attacks on network devices and computers to retrieve data, gain access, or escalate access privileges to administrator status. Table 11-7 describes access attacks.

Table 11-7 *Types of Access Attacks*

Technet24

Access Attack	Description
Password attack	The threat actor attempts to discover critical system passwords using various methods. Password attacks are very common and can be launched using a variety of password cracking tools.
Spoofing attack	The threat actor has a device pose as another device by falsifying data. Common spoofing attacks include IP spoofing, MAC spoofing, and DHCP spoofing.
Trust exploitation	The threat actor uses unauthorized privileges to gain access to a system, possibly compromising the target.
Port redirection	The threat actor uses a compromised system as a base for attacks against other targets.
Man-in-the-middle attack	The threat actor is positioned between two legitimate entities in order to read or modify the data that passes between the two parties.
Buffer overflow attack	The threat actor exploits the buffer memory and overwhelms it with unexpected values. This usually renders the system inoperable, creating a DoS attack.

Social Engineering Attacks

In social engineering attacks, threat actors attempt to manipulate individuals into performing actions or divulging confidential information. Table 11-8 describes social engineering techniques.

Table 11-8 *Types of Social Engineering Attacks*

Technet24

Social Engineering Attack	Description
Pretexting	An attack in which a threat actor pretends to need personal or financial data to confirm the identity of the target.
Phishing	An attack in which a threat actor sends fraudulent email that is disguised as being from a legitimate, trusted source to trick the recipient into installing malware on their device or into sharing personal or financial information.
Spear phishing	An attack in which a threat actor creates a targeted phishing attack tailored for a specific individual or organization.
Spam	Unsolicited email, also known as junk mail, that often contains harmful links, malware, or deceptive content.

Something for something	Sometimes called *quid pro quo*, an attack in which a threat actor requests personal information from a party in exchange for something such as a gift.
Baiting	An attack in which a threat actor leaves a malware-infected flash drive in a public location. A victim finds the drive and inserts it into a laptop, unintentionally installing malware.
Impersonation	An attack in which a threat actor pretends to be someone they are not to gain the trust of a victim.
Tailgating	An attack in which a threat actor quickly follows an authorized person into a secure location to gain access to a secure area.
Shoulder surfing	An attack in which a threat actor inconspicuously looks over someone's shoulder to steal passwords or other information.
Dumpster diving	An attack in which a threat actor rummages through trash bins to discover confidential documents.

DoS and DDoS Attacks

Technet24

A DoS attack creates some sort of interruption of network services to users, devices, or applications. DoS attacks are created in two ways:

• **Overwhelming quantity of traffic:** The threat actor sends an enormous quantity of data at a rate that the network, host, or application cannot handle. This causes transmission and response times to slow down. It can also crash a device or service.

• **Maliciously formatted packets:** The threat actor sends a maliciously formatted packet to a host or an application, and the receiver is unable to handle it. This causes the receiving device to run very slowly or crash.

DoS attacks are relatively simple to conduct, even by an unskilled threat actor. A DDoS attack is similar to a DoS attack, but it originates from multiple, coordinated sources. For example, a threat actor may build a network of infected hosts, known as *zombies*. A network of zombies is called a *botnet*. The threat actor can then use a command-and-control (CnC) program to instruct the botnet of zombies to carry out a DDoS attack.

IP Attacks

IP does not validate whether the source IP address contained in a packet actually came from that source. For this reason, threat actors can send packets using a spoofed source IP address. Threat actors can also tamper with the other fields in the IP header to carry out their attacks. Security analysts must understand the different fields in both the IPv4 and IPv6 headers. Table 11-9

describes some of the most common IP-related attacks.

Table 11-9 *Types of IP Attacks*

IP Attack Technique	Description
ICMP attacks	Threat actors use Internet Control Message Protocol (ICMP) echo packets (pings) to discover subnets and hosts on a protected network, to generate DoS flood attacks, and to alter host routing tables.
Amplification and reflection attack	Threat actors use DoS and DDoS attacks in an attempt to prevent legitimate users from accessing information or services. In one type of amplification and reflection attack, the threat actor forwards ICMP echo request messages to many hosts. These messages contain the source IP address of the victim. The hosts all reply to the spoofed IP address of the victim and overwhelm it.

Technet24

Address spoofing attack	Threat actors spoof the source IP address in an IP packet to perform blind spoofing or non-blind spoofing. In non-blind spoofing, the threat actor can see the traffic that is being sent between the host and the target. The threat actor uses non-blind spoofing to inspect the reply packet from the target victim. Non-blind spoofing determines the state of a firewall and sequence-number prediction. It can also be done to hijack an authorized session. In blind spoofing, the threat actor cannot see the traffic that is being sent between the host and the target. Blind spoofing is used in DoS attacks.
Man-in-the-middle (MITM) attack	Threat actors position themselves between a source and destination to transparently monitor, capture, and control the communication. They can eavesdrop by inspecting captured packets or alter packets and forward them to their original destination.
Session hijacking	Threat actors gain access to the physical network and then use an MITM attack to hijack a session.

Transport Layer Attacks

Threat actors conduct port scans of target devices to discover which services are available. A threat actor can exploit TCP and UDP in the following ways:

• **TCP SYN flood attack:** This type of attack exploits the TCP three-way handshake. The threat actor continually sends TCP SYN session request packets with a randomly spoofed source IP address to a target. The target device replies with a TCP SYN-ACK packet to the spoofed IP address and waits for a TCP ACK packet. The responses never arrive. Eventually the target host is overwhelmed with half-open TCP connections, and TCP services are denied to legitimate users.

• **TCP reset attack:** A threat actor could use a TCP reset attack to send a spoofed packet containing a TCP RST to one or both endpoints. This creates a DoS condition for the connection.

• **TCP session hijacking:** A threat actor takes over an already-authenticated host as it communicates with the target. The threat actor must spoof the IP address of one host, predict the next sequence number, and send an ACK to the other host. If successful, the threat actor could send, but not receive, data from the target device.

• **UDP flood attack:** The threat actor uses a tool to send a flood of UDP packets, often from a spoofed host, to a server on the subnet. The program sweeps through all the known ports, trying to find closed ports. This causes the server to reply with an ICMP port unreachable message. Because there are many closed ports on the server, there is a lot of traffic on the segment, which uses up most of the bandwidth. The result is very similar to the result of a DoS attack.

Technet24

Security Program

An organization should educate its user community through a security program. An effective security program includes the following basic elements:

• **User awareness:** All users should be made aware of the need for data confidentiality to protect corporate information, as well as their own credentials and personal information. They should also be made aware of potential threats, schemes to mislead, and proper procedures to report security incidents. Users should also be instructed to follow strict guidelines regarding data loss.

• **User training:** All users should be required to participate in periodic formal training so that they become familiar with all corporate security policies.

• **Physical access control:** Infrastructure locations, such as network closets and data centers, should remain securely locked. Administrators should control physical access and quickly remove access when an employee is dismissed.

Password Vulnerabilities

Most organizations use some type of authentication to control access, typically requiring a username and password. Usernames are often predictable, and default or simple passwords can be easily guessed. Consider

these common password guesses: password, password123, and 123456. Attackers might also try using admin for both the username and password. Further complicating the matter, network equipment often comes with preconfigured default passwords. These passwords are easily found using a simple Internet search.

Types of Attacks

Attacks can be classified as online or offline.

• **Online attacks:** The attacker manually enters password guesses directly into the system.

• **Offline attacks:** The attacker obtains encrypted or hashed passwords and uses software to crack them externally.

Attackers might use dictionary attacks (guessing passwords using a word list) or brute-force attacks (trying every possible character combination). Brute-force attacks require significant computing power and time.

Mitigation Strategies

Implement password policies that require:

• Long passwords with a mix of upper- and lowercase letters, numbers, and special characters

Technet24

• Regularly changing passwords to limit the effectiveness of brute-force attacks

Password Guidelines

To secure network devices, follow these password guidelines:

• Use at least 8 characters and preferably 10 or more.

• Use a mix of uppercase and lowercase letters, numbers, symbols, and spaces.

• Avoid repetition, common words, sequences, and personal information.

• Misspell words deliberately (for example, 5mYth instead of Smith).

• Change passwords regularly.

• Don’t write passwords down in obvious places.

Table 11-10 lists some examples of weak and strong passwords.

Table 11-10 *Examples of Weak and Strong Passwords*

Weak Password	Why It Is Weak
secret	Simple dictionary password
smith	Maiden name of mother
toyota	Make of a car
bob1967	Name and birthday of the user
Blueleaf23	Simple words and numbers
Strong Password	**Why It Is Strong**
b67n42d39c	Combines alphanumeric characters
12^h u4@1p7	Combines alphanumeric characters, symbols, and spaces

Another strong type of password is a passphrase—a combination of multiple words (such as correct horse battery staple).

Password Alternatives

A password is a single factor for authentication, representing “something you know.” However, more secure methods include the following:

• **Multifactor authentication (MFA):** MFA requires users to provide multiple factors, such as:

• **Something you know:** For example, a password

• **Something you have:** For example, a dynamic cryptographic key or a time-

Technet24

limited code sent via text message

• **Digital certificates:** A digital certificate is a trusted form of identification that contains encrypted information. It must be issued and signed by a trusted certificate authority (CA). Certificates have specific time ranges and can be revoked if necessary.

• **Biometric credentials:** Biometric methods use physical attributes ("something you are") to identify users, such as fingerprints, face recognition, palm prints, voice recognition, iris recognition, and retinal scans. These methods are usually unique to individuals and difficult to replicate, though some, like facial recognition, can be more easily fooled.

Table 11-11 summarizes the alternatives to password authentication.

Table 11-11 *Summary of Password Authentication and Alternatives*

Characteristic	Password Only	Two-Factor	Digital Certificates	Biometric
Something you know	Yes	Yes		
Something you have		Yes	Yes	
Something you are				Yes

Additional Password Security

To keep passwords secret, follow these best practices on Cisco routers and switches:

• **Encrypt plaintext passwords:**

• Use the **service password-encryption** command to prevent unauthorized viewing of plaintext passwords.

• Example: R1(config)# **service password-encryption**

• **Set minimum password length:**

• Use the **security passwords min-length** length command to enforce a minimum password length.

• Example: R1(config)# **security passwords min-length 8**

• **Deter brute-force attacks:**

• Use the **login block-for** # **attempts** # **within** # command to block login attempts after a certain number of failed attempts.

• Example: R1(config)# **login block-for 120 attempts 3 within 60**

• **Disable inactive sessions:**

Technet24

• Use the **exec-timeout** *minutes seconds* command to disconnect inactive sessions.

• Example: R1(config)# **line vty 0 4**

• R1(config-line)# **exec-timeout 5 30**

Example 11-1 shows how these best practices are implemented on a Cisco router.

Example 11-1 *R1 Password Security Configuration*

```
R1(config)# service password-encryption
R1(config)# security passwords min-length 8
R1(config)# login block-for 120 attempts 3 within 60
R1(config)# line vty 0 4
R1(config-line)# password cisco123
R1(config-line)# exec-timeout 5 30
R1(config-line)# transport input ssh
R1(config-line)# end
R1#
R1# show running-config | section line vty
line vty 0 4
  password 7 094F471A1A0A
  exec-timeout 5 30
  login
  transport input ssh
R1#
```

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Introduction to Networks v7	16
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	3
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	9

Technet24

Day 10. ACL Concepts

CCNA 200-301 Exam Topics

• 5.6 Configure and verify access control lists

Key Topics

One of the most important skills a network administrator needs is mastery of access control lists (ACLs). Administrators use ACLs to stop traffic or permit only specified traffic on their networks. Standard and extended ACLs can apply a number of security features, including policy-based routing, quality of service (QoS), Network Address Translation (NAT), and Port Address Translation (PAT).

You can also configure standard and extended ACLs on router interfaces to control the type of traffic that is permitted through a given router. Today we review ACL concepts, including what they are, how a router uses them to filter traffic, and what types of ACLs are available.

ACL Operation

A router’s default operation is to forward all packets, as long as a route exists for the packet and the link is up. ACLs can help implement a basic level of security. However, they are not the only security solution a large organization should implement. In fact, ACLs increase the latency of routers. If an

organization is very large, with routers managing the traffic of hundreds or thousands of users, the administrator more than likely will use a combination of other security implementations that are beyond the scope of the CCNA.

Defining an ACL

An ACL is a router configuration script (that is, a list of statements) that controls whether a router permits or denies packets to pass, based on criteria in the packet header. To determine whether a packet is permitted or denied, it is tested against the ACL statements in sequential order. When a statement matches, no more statements are evaluated; the packet is either permitted or denied. There is an implicit deny any statement at the end of an ACL. If a packet does not match any of the statements in the ACL, it is dropped.

Processing Interface ACLs

ACLs can be applied to an interface for inbound and outbound traffic. However, you need a separate ACL for each direction. The flowchart in Figure 10-1 details the steps a router takes when evaluating an ACL on inbound and outbound interfaces.

Technet24

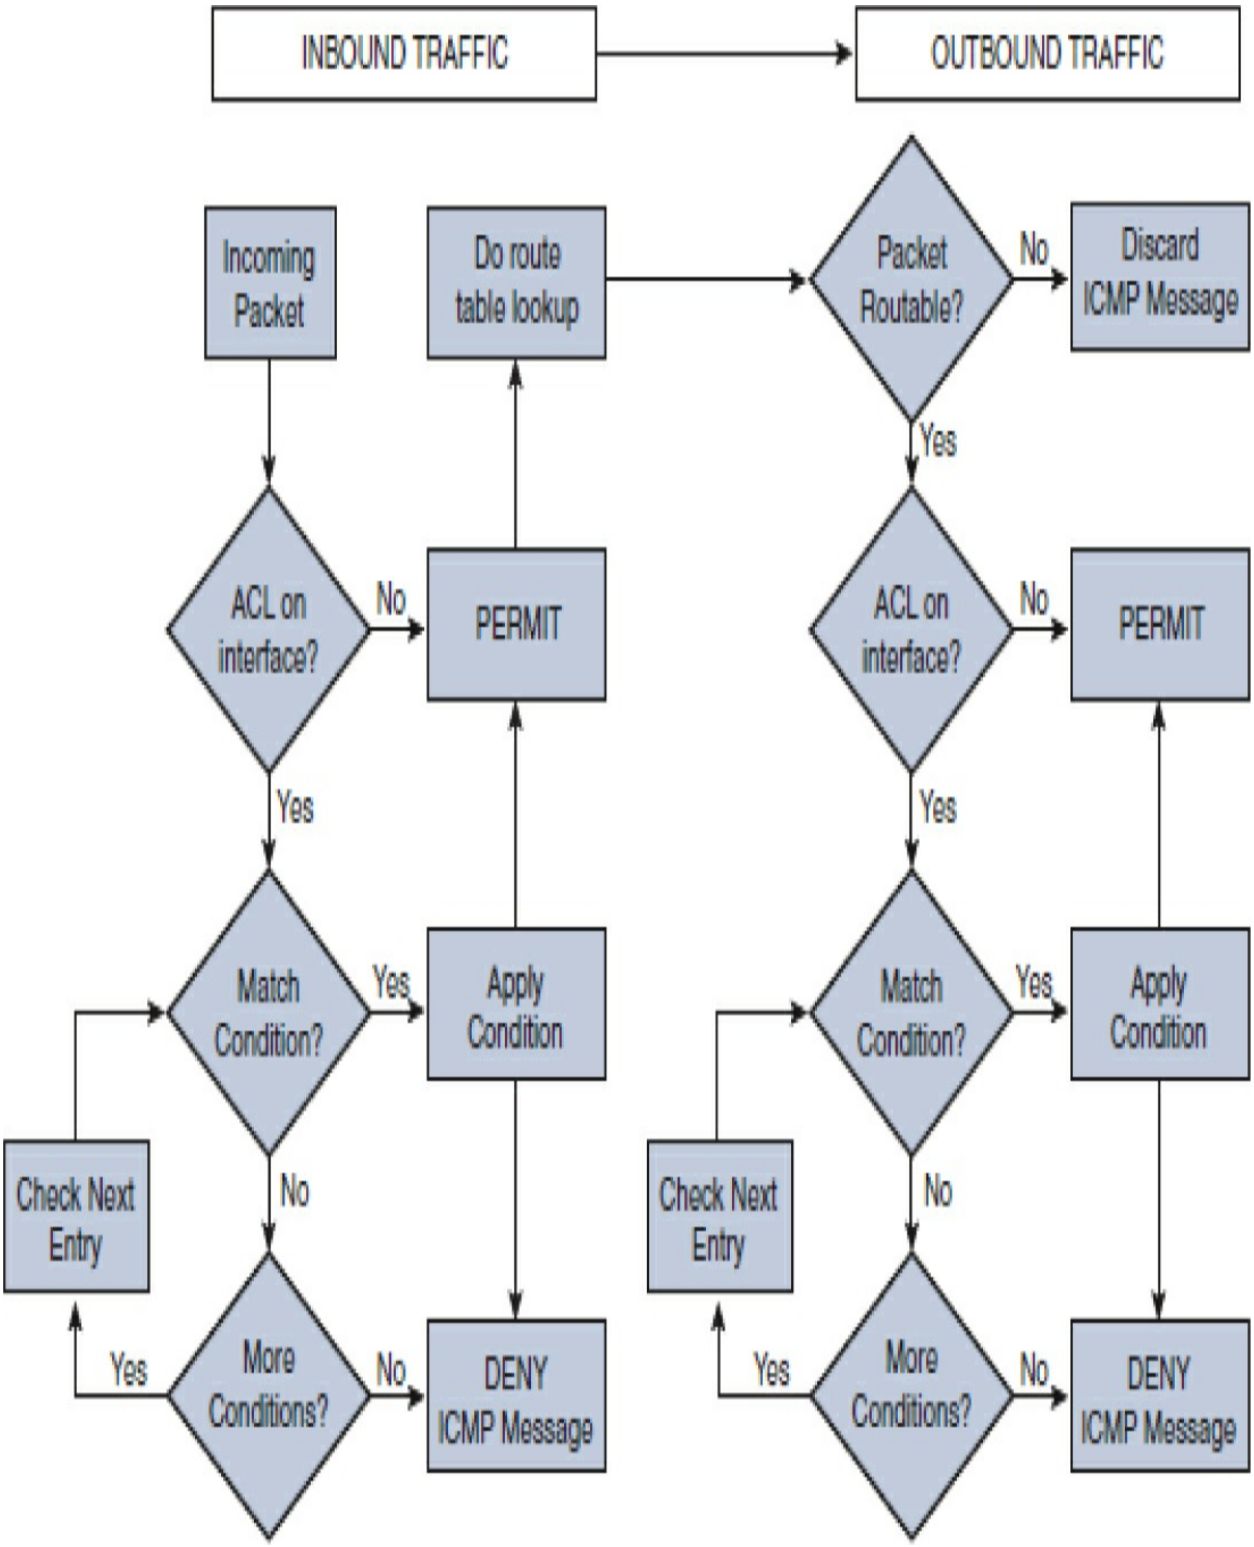

Figure 10-1 *ACL Interface Processing for Inbound and Outbound Traffic*

For inbound traffic, the router checks for an inbound ACL applied to the

interface before doing a route table lookup. Then, for outbound traffic, the router makes sure that a route to the destination exists before checking for ACLs. Finally, if an ACL statement results in a dropped packet, the router sends an ICMP destination unreachable message.

The choice of using an inbound or outbound ACL is easy to make if, first, you place yourself inside the router—that is, be the router. From such a stance, you can visualize processing a packet coming into a router interface (inbound), deciding what to do with the packet (Is there an inbound ACL? Is there a route to the destination?), and forwarding the packet (What is the outbound interface? Is there an ACL on the interface?).

List Logic with IP ACLs

An ACL is a list of commands that are processed in order, from the first statement in the list to the last statement. Each command has different matching logic that the router must apply to each packet when filtering is enabled. ACLs use first-match logic. If a packet matches one line in the ACL, the router takes the action listed in that line of the ACL and ignores the rest of the ACL statements.

For example, Figure 10-2 shows ACL 1 with three lines of pseudocode. The ACL is applied to R2's S0/0/1 interface, as the arrow indicates. Inbound traffic from R1 will be filtered using ACL 1.

Technet24

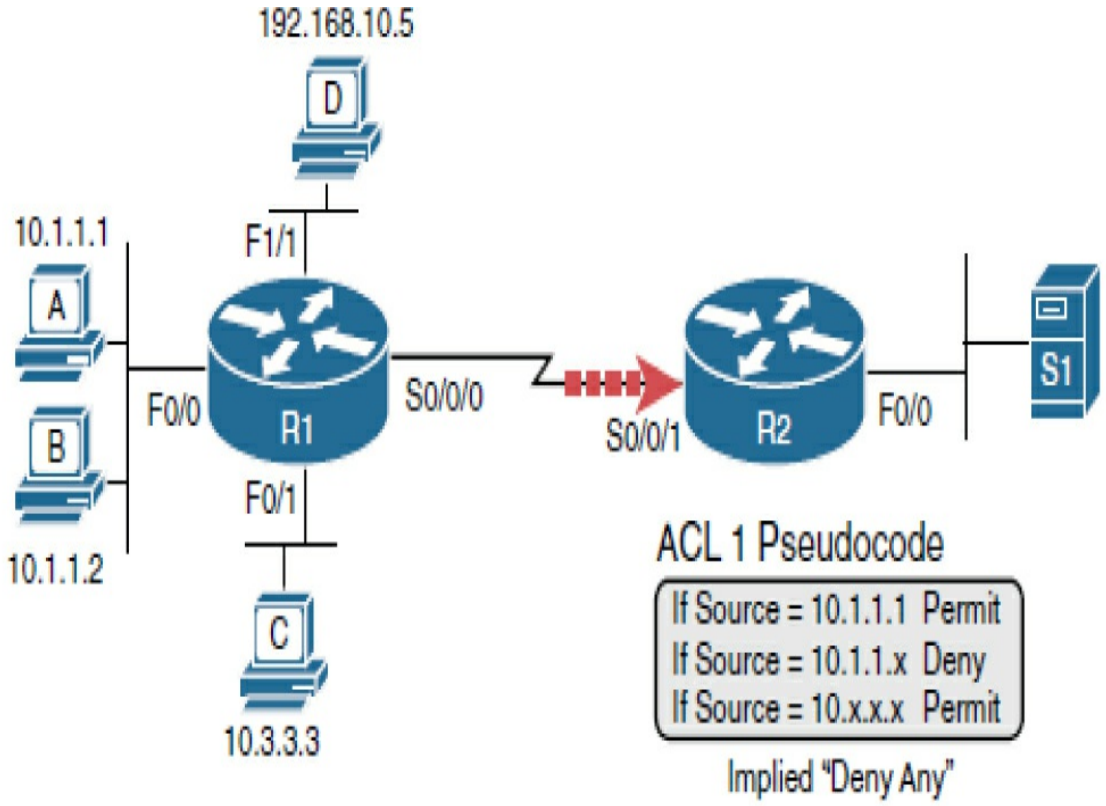
192.168.10.5
D
F1/1
10.1.1.1
A
B
10.1.1.2
F0/0
R1
S0/0/0
S0/0/1
R2
F0/0
S1
F0/1
C
10.3.3.3
ACL 1 Pseudocode
If Source = 10.1.1.1 Permit
If Source = 10.1.1.x Deny
If Source = 10.x.x.x Permit
Implied "Deny Any"

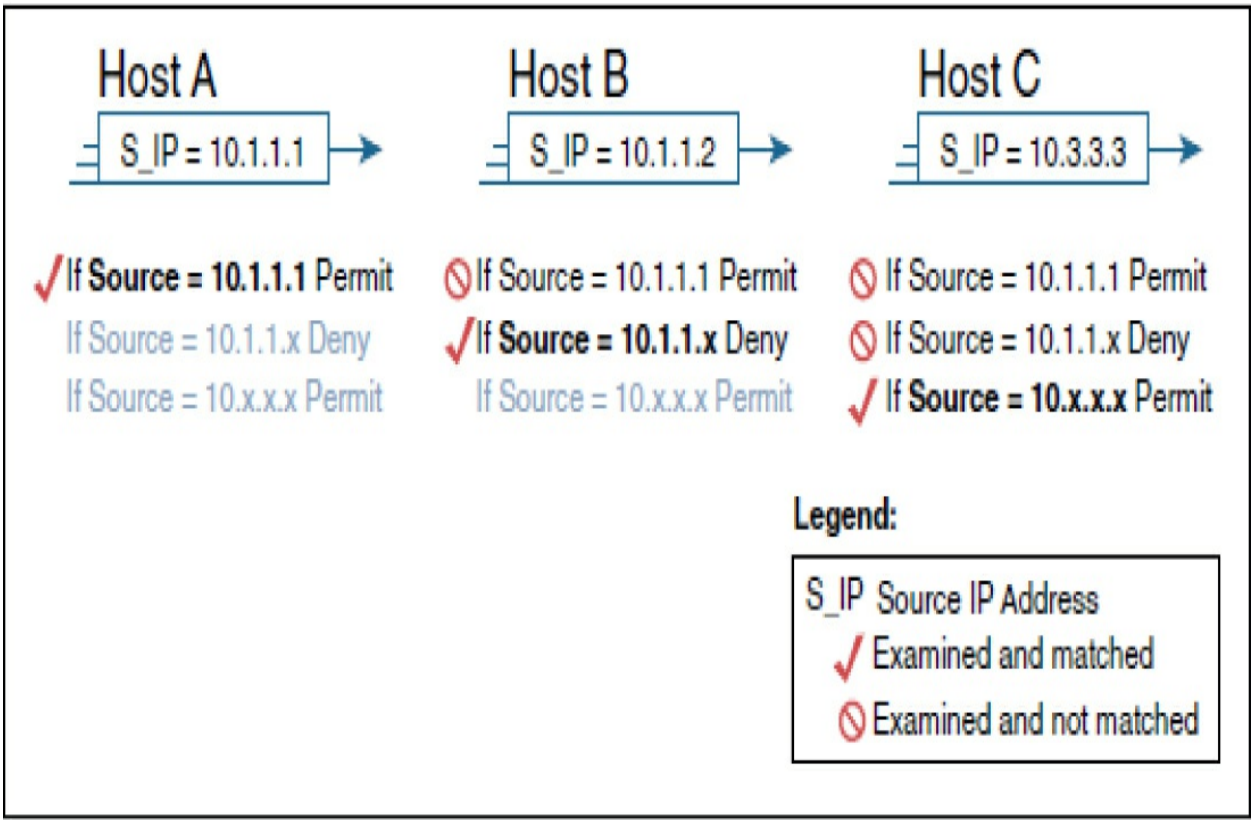
Host A
S_IP = 10.1.1.1
✓ If Source = 10.1.1.1 Permit
If Source = 10.1.1.x Deny
If Source = 10.x.x.x Permit
Host B
S_IP = 10.1.1.2
⊘ If Source = 10.1.1.1 Permit
✓ If Source = 10.1.1.x Deny
If Source = 10.x.x.x Permit
Host C
S_IP = 10.3.3.3
⊘ If Source = 10.1.1.1 Permit
⊘ If Source = 10.1.1.x Deny
✓ If Source = 10.x.x.x Permit
Legend:
S_IP Source IP Address
✓ Examined and matched
⊘ Examined and not matched

Figure 10-2 *Example of ACL Matching Logic*

The box below the topology shows the logic for how packets from each host source address (labeled S_IP in the figure) are processed. Notice that when a match is made for Host A and Host B, the condition is applied (Host A is permitted and Host B is denied), and no further statements are evaluated. Host C matches the last statement in the list and is permitted.

Host D does not match any of the items in the ACL, so the packet is discarded. The reason is that every IP ACL has an implied deny any at the end of the ACL.

Planning to Use ACLs

Because an ACL can be used to filter traffic, it is important that you thoroughly plan the implementation of an ACL before actually configuring it.

Types of ACLs

ACLs can be configured to filter any type of protocol traffic, including other network layer protocols such as AppleTalk and IPX. For the CCNA exam, we focus on IPv4 and IPv6 ACLs, which come in the following types:

• **Standard IPv4 ACLs:** Filter traffic based on source address only

• **Extended IPv4 and IPv6 ACLs:** Can filter traffic based on source and destination address, specific protocols, and source and destination TCP and

Technet24

UDP ports

You can use two methods to identify both standard and extended ACLs:

• **Numbered IPv4 ACLs:** Use a number for identification

• **Named IPv4 and IPv6 ACLs:** Use a descriptive name or number for identification

Named ACLs must be used with some types of Cisco IOS configurations, including IPv6 ACLs. However, they provide two basic benefits for standard and extended IPv4 ACLs:

• By using a descriptive name (such as BLOCK-HTTP), a network administrator can more quickly determine the purpose of an ACL. This is particularly helpful in larger networks, where a router can have many ACLs with hundreds of statements.

• They reduce the amount of typing you must do to configure each statement in a named ACL, as you see in Day 9, "ACL Implementation."

Both numbered and named ACLs can be configured for standard as well as extended ACL implementations. Figure 10-3 summarizes the categories of IPv4 ACLs.

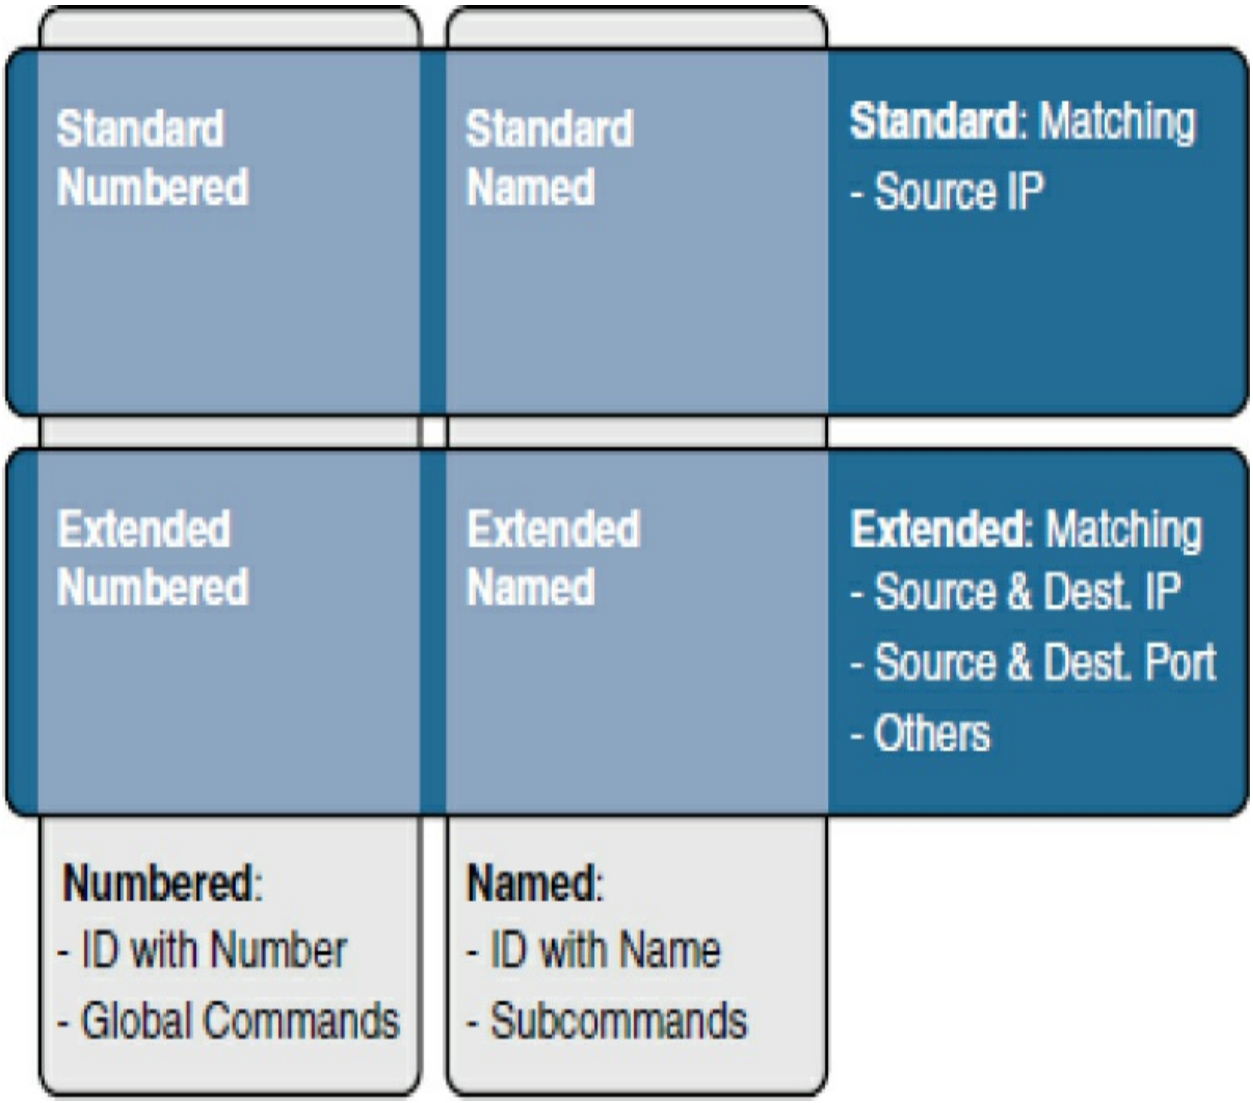

Figure 10-3 *Comparisons of IPv4 ACL Types*

ACL Identification

Table 10-1 lists the different ACL number ranges for the IPv4 protocol. The table is not exhaustive. Other ACL numbers are available for other types of protocols that are either rarely used or beyond the scope of the CCNA. IPv6 uses only named ACLs.

Table 10-1 *IPv4 ACL Numbers*

Technet24

Protocol	Range
IP	1–99
Extended IP	100–199
Standard IP (expanded)	1300–1999
Extended IP (expanded)	2000–2699

Named IP ACLs give you more flexibility in working with the ACL entries. In addition to using more memorable names, using named ACLs instead of numbered ACLs enables you to delete individual statements in a named IP access list.

Cisco IOS Software Release 12.3 introduced IP access list entry sequence numbering for both numbered and named ACLs. IP access list entry sequence numbering provides the following benefits:

• You can edit the order of ACL statements.

• You can remove individual statements from an ACL.

• You can use the sequence number to insert new statements into the middle of the ACL.

Sequence numbers are automatically added to the ACL if they are not entered explicitly at the time the ACL is created.

ACL Design Guidelines

Well-designed and -implemented ACLs add an important security component to your network. Follow these general principles to ensure that the ACLs you create have the intended results:

• Based on the test conditions, choose a standard or extended, numbered, or named ACL.

• Only one ACL is allowed per protocol, per direction, and per interface.

• Organize the ACL to enable processing from the top down. Organize your ACL so that more specific references to a network, subnet, or host appear before more general ones. Place conditions that occur more frequently before conditions that occur less frequently.

• All ACLs contain an implicit deny any statement at the end.

• Create the ACL before applying it to an interface.

• Depending on how you apply the ACL, the ACL filters traffic either going through the router or going to and from the router, such as traffic to or from the vty lines.

Technet24

• You typically should place extended ACLs as close as possible to the source of the traffic that you want to deny. Because standard ACLs do not specify destination addresses, you must put the standard ACL as close as possible to the destination of the traffic you want to deny so that the source can reach intermediary networks.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	4
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	6
Portable Command Guide	21

Day 9. ACL Implementation

CCNA 200-301 Exam Topics

• 5.6 Configure and verify access control lists

Key Topics

On Day 10, "ACL Concepts," we reviewed ACL concepts. Today we focus on the configuration, verification, and troubleshooting of IPv4 and IPv6 ACLs.

Configuring Standard Numbered IPv4 ACLs

Standard IPv4 ACLs, which are numbered ACLs in the ranges 1 to 99 and 1300 to 1999 or are named ACLs, filter packets based on a source address and mask. They permit or deny the entire TCP/IP protocol suite. Configuring an ACL requires two steps:

Create the ACL.

Apply the ACL.

Let's use the simple topology in Figure 9-1 to demonstrate how to configure both standard and extended IPv4 ACLs.

Technet24

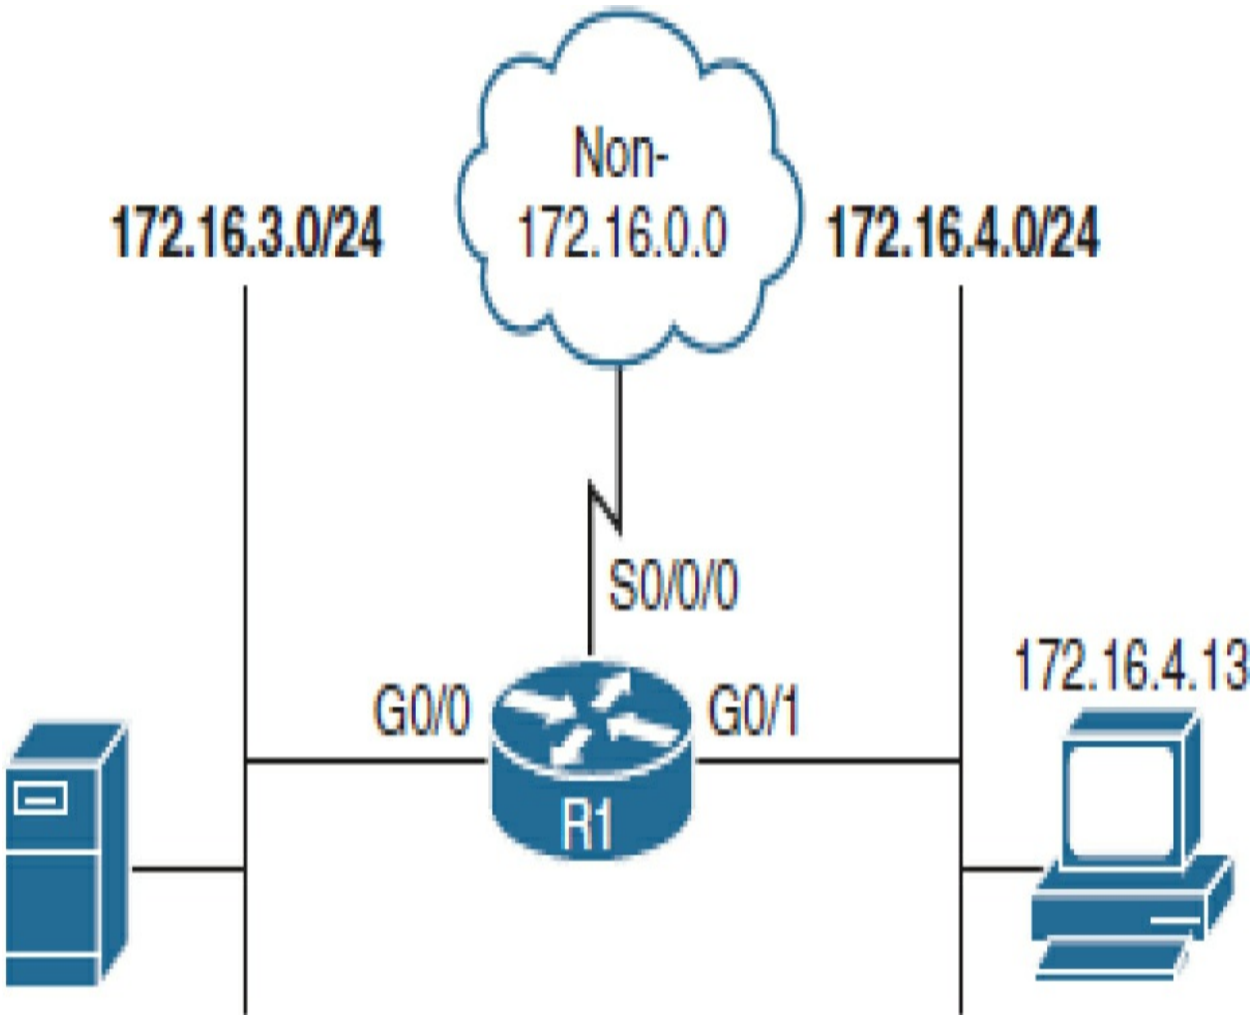

Figure 9-1 *IPv4 ACL Configuration Topology*

Standard Numbered IPv4 ACL: Permit Specific Network

Create an ACL to prevent traffic that is not part of the internal networks (172.16.0.0/16) from traveling out either of the Gigabit Ethernet interfaces:

Create the ACL. Use the **access-list** global configuration command to create an entry in a standard IPv4 ACL:

```
R1(config)# access-list 1 permit 172.16.0.0
```

```
  0.0.255.255
```

The sample statement matches any address that starts with 172.16.x.x. You can use the **remark** option to add a description to the ACL.

Apply the ACL. Use the interface configuration command to select an interface to which to apply the ACL. Then use the **ip access-group** interface configuration command to activate the existing ACL on an interface for a specific direction (in or out):

```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ip access-group 1 out
R1(config-if)# interface gigabitethernet 0/1
R1(config-if)# ip access-group 1 out
```

This step activates the standard IPv4 ACL 1 on both the interfaces as an outbound filter.

This ACL allows only traffic from source network 172.16.0.0 to be forwarded out on G0/0 and G0/1. Traffic from networks other than 172.16.0.0 is blocked with the implied deny any.

Standard Numbered IPv4 ACL: Deny a Specific Host

Create an ACL to prevent traffic that originates from host 172.16.4.13 from traveling out G0/0. Create and apply the ACL with the commands in

Technet24

Example 9-1.

Example 9-1 *ACL Preventing Traffic Originating from a Specific Host*

```
R1(config)# access-list 1 deny 172.16.4.13 0.0.0.0
R1(config)# access-list 1 permit 0.0.0.0 255.255.255.2
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ip access-group 1 out
```

This ACL is designed to block traffic from a specific address, 172.16.4.13, and to allow all other traffic to be forwarded on interface G0/0. The first statement can also be written with the keyword **host** replacing the 0.0.0.0 wildcard mask, as follows:

```
R1(config)# access-list 1 deny host 172.16.4.13
```

In fact, starting with Cisco IOS Software Release 12.3, you can enter the following:

```
R1(config)# access-list 1 deny 172.16.4.13
```

The second statement can be written with the keyword **any** replacing the source address 0.0.0.0 and wildcard mask 255.255.255.255, as follows:

```
R1(config)# access-list 1 permit any
```

Standard Numbered IPv4 ACL: Deny a Specific Subnet

Create an ACL to prevent traffic that originates from the subnet 172.16.4.0/24 from traveling out the G0/0 interface. Create and apply the ACL with the commands in Example 9-2.

Example 9-2 *ACL Preventing Traffic Originating from a Specific Subnet*

```
R1(config)# access-list 1 deny 172.16.4.0 0.0.0.255
R1(config)# access-list 1 permit any
R1(config)# interface g0/0
R1(config-if)# ip access-group 1 out
```

This ACL is designed to block traffic from a specific subnet, 172.16.4.0, and to allow all other traffic to be forwarded out G0/0.

Standard Numbered IPv4 ACL: Deny Telnet or SSH Access to the Router

For traffic into and out of the router (not through the router), filter Telnet or SSH access to the router by applying an ACL to the vty ports. Restricting vty access is primarily a technique for increasing network security and defining which addresses are allowed Telnet access to the router EXEC process.

Technet24

Create and apply the ACL with the commands in Example 9-3.

Example 9-3 *Access List Allowing One Host Only Remote Access to R1*

```
R1(config)# access-list 12 permit host 172.16.4.13
R1(config)# line vty 0 15
R1(config-line)# access-class 12 in
```

In this example, only host 172.16.4.13 is allowed to Telnet into R1. All other IP addresses are denied implicitly.

Configuring Extended Numbered IPv4 ACLs

For more precise traffic filtering control, use extended IPv4 ACLs. Extended IPv4 ACLs can be named or numbered in the ranges 100 to 199 and 2000 to 2699. Extended ACLs check for source and destination IP addresses. In addition, at the end of the extended ACL statement, you can specify the protocol and optional TCP or UDP application to filter more precisely. To configure numbered extended IPv4 ACLs on a Cisco router, create an extended IP ACL and activate that ACL on an interface. For CCNA exam purposes, the extended IPv4 ACL command syntax is as follows:

```
Router(config)# access-list access-list-number {permit
source source-wildcard [operator port] destination des
[operator port] [established] [log]
```

Table 9-1 explains the syntax of the command.

Table 9-1 *Command Parameters for a Numbered Extended IPv4 ACL*

Technet24

Command Parameter	Description
access-list-number	Identifies the list using a number in the range 100–199 or 2000–2699.
permit | deny	Indicates whether this entry allows or blocks the specified address.
protocol	If **ip** is specified, the entire TCP/IP protocol suite is filtered. Other protocols you can filter include TCP, UDP, ICMP, EIGRP, and OSPF. Use **?** after the **permit** | **deny** argument to see all the available protocols.
source and *destination*	Identifies source and destination IP addresses.
source-wildcard and *destination-wildcard*	Wildcard mask. 0s indicate positions that must match, and 1s indicate "don't care" positions.
operator [port | app_name]	The operator can be **lt** (less than), **gt** (greater than), **eq** (equal to), or **neq** (not equal to). The port number referenced can be either the source port or the destination port, depending on where in the ACL the port number is configured. As an alternative to the port number, well-known application names can be used, such as Telnet, FTP, and SMTP.
established	For inbound TCP only. Allows TCP traffic to pass if the packet is a response to an outbound-initiated session. This type of traffic has the acknowledgment (ACK) bits set.
log	Sends a logging message to the console.

Extended Numbered IPv4 ACL: Deny FTP from Subnets

For the network in Figure 9-1, we now create an ACL to prevent FTP traffic originating from the subnet 172.16.4.0/24 and going to the 172.16.3.0/24 subnet from traveling out G0/0. Create and apply the ACL with the commands in Example 9-4.

Example 9-4 *Access List Preventing FTP Traffic from Specific Subnets*

```
R1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.
eq 21
R1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.
eq 20
R1(config)# access-list 101 permit ip any any
R1(config)# interface g0/0
R1(config-if)# ip access-group 101 out
```

The **deny** statements block FTP traffic originating from subnet 172.16.4.0 to subnet 172.16.3.0. The **permit** statement allows all other IP traffic out interface G0/0. Two statements must be entered for the FTP application because port 21 is used to establish, maintain, and terminate an FTP session, while port 20 is used for the actual file transfer task.

Extended Numbered IPv4 ACL: Deny Only Telnet from Subnet

Create an ACL to prevent Telnet traffic that originates from the subnet

Technet24

172.16.4.0/24 from traveling out interface G0/0. Create and apply the ACL with the commands in Example 9-5.

Example 9-5 *Access List Preventing Telnet Traffic from a Specific Subnet*

```
R1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.
R1(config)# access-list 101 permit ip any any
R1(config)# interface g0/0
R1(config-if)# ip access-group 101 out
```

This example denies Telnet traffic from 172.16.4.0 that is being sent out interface G0/0. All other IP traffic from any other source to any destination is permitted out G0/0.

Configuring Named IPv4 ACLs

With a named ACL, you can identify standard and extended ACLs with an alphanumeric string (name) instead of the current numeric representations.

Because you can delete individual entries with named ACLs, you can modify an ACL without having to delete and then reconfigure the entire ACL. With Cisco IOS Software Release 12.3 and later, you can insert individual entries using an appropriate sequence number.

Standard Named IPv4 ACL Steps and Syntax

The following steps and syntax are used to create a standard named ACL:

Name the ACL. Starting from global configuration mode, use the **ip access-list standard** command to name the standard ACL. ACL names are alphanumeric and must be unique:

```
Router(config)# ip access-list standard name
```

Create the ACL. From standard named ACL configuration mode, use **permit** or **deny** statements to specify one or more conditions for determining whether a packet is forwarded or dropped:

```
Router(config-std-nacl)# [sequence-number]
{permit | deny} source source-wildcard
```

If you do not specify a sequence number, Cisco IOS increments the sequence number by 10 for every statement you enter.

Apply the ACL. Activate the named ACL on an interface with the **ip access-group** *name* command:

```
Router(config-if)# ip access-group name [in | out]
```

Standard Named IPv4 ACL: Deny a Single Host from a Given Subnet

Technet24

For the network shown previously in Figure 9-1, create a standard ACL named TROUBLEMAKER to prevent traffic that originates from the host 172.16.4.13 from traveling out interface G0/0. Create and apply the ACL with the commands in Example 9-6.

Example 9-6 *Named ACL Preventing Traffic from a Specific Host*

```
R1(config)# ip access-list standard TROUBLEMAKER
R1(config-std-nacl)# deny host 172.16.4.13
R1(config-std-nacl)# permit 172.16.4.0 0.0.0.255
R1(config-std-nacl)# interface g0/0
R1(config-if)# ip access-group TROUBLEMAKER out
```

Extended Named IPv4 ACL Steps and Syntax

The following steps and syntax are used to create an extended named ACL:

Name the ACL. Starting from global configuration mode, use the **ip access-list extended** command to name the extended ACL:

```
Router(config)# ip access-list extended name
```

Create the ACL. From extended named ACL configuration mode, use **permit** or **deny** statements to specify one or more conditions for determining whether a packet is forwarded or dropped:

```
Router(config-ext-nacl)# [sequence-number] {deny | permit} protocol source source-wildcard [operator port] destination destination-wildcard [operator port] [established] [log]
```

Apply the ACL. Activate the named ACL on an interface with the **ip access-group** *name* command:

```
Router(config-if)# ip access-group name [in | out]
```

Adding Comments to Named or Numbered IPv4 ACLs

You can add comments to ACLs by using the **remark** argument in place of **permit** or **deny**. Remarks are descriptive statements that you can use to better understand and troubleshoot either named or numbered ACLs.

Example 9-7 shows how to add a comment to a numbered ACL.

Example 9-7 *Adding Comments to a Numbered ACL*

```
R1(config)# access-list 101 remark Permitting John to
R1(config)# access-list 101 permit tcp host 172.16.4.1
```

Example 9-8 shows how to add a comment to a named ACL.

Technet24

Example 9-8 *Adding Comments to a Named ACL*

```
R1(config)# ip access-list standard PREVENTION
R1(config-std-nacl)# remark Do not allow Jones subnet
R1(config-std-nacl)# deny 172.16.4.0 0.0.0.255
```

Verifying IPv4 ACLs

When you finish configuring an ACL, use **show** commands to verify the configuration. Use the **show access-lists** command to display the contents of all ACLs, as in Example 9-9. By entering the ACL name or number as an option for this command, you can display a specific ACL.

Example 9-9 *Verifying Access List Configuration*

```
R1# show access-lists
Standard IP access list SALES
    10 permit 10.3.3.1
    20 permit 10.4.4.1
    30 permit 10.5.5.1
    40 deny 10.1.1.0, wildcard bits 0.0.0.255
    50 permit any
Extended IP access list ENG
    10 permit tcp host 10.22.22.1 any eq telnet (25 mat
    20 permit tcp host 10.33.33.1 any eq ftp
    30 permit tcp host 10.33.33.1 any eq ftp-data
```

Notice in the output from the **show access-lists** command in Example 9-9 that sequence numbers are incremented by 10—most likely because the administrator did not enter a sequence number. Also notice that this command tells you how many times Cisco IOS has matched a packet to a statement—25 times, in the case of the first statement in the named ACL ENG.

The **show ip interface** command displays IP interface information and indicates whether any IP ACLs are set on the interface. In the **show ip interface g0/0** command output in Example 9-10, IP ACL 1 has been configured on the G0/0 interface as an inbound ACL. No outbound IP ACL has been configured on the G0/0 interface.

Example 9-10 *Verifying Access List Configuration on a Specific Interface*

```
R1# show ip interface g0/0
GigabitEthernet0/0 is up, line protocol is up
    Internet address is 10.1.1.11/24
    Broadcast address is 255.255.255.255
    Address determined by setup command
    MTU is 1500 bytes
    Helper address is not set
    Directed broadcast forwarding is disabled
    Outgoing access list is not set
    Inbound access list is 1
```

Technet24

```
     Proxy ARP is enabled
      <output omitted>
```

Finally, you can also verify your ACL creation and application with the **show running-config** command (see Example 9-11).

Example 9-11 *Verifying ACL Creation and Application in the Running Configuration*

```
R1# show running-config
Building configuration...
!
<output omitted>
!
interface GigabitEthernet0/0
   ip address 10.44.44.1 255.255.255.0
   ip access-group ENG out
!
<output omitted>
!
interface Serial0/0/0
ip address 172.16.2.1 255.255.255.252
   ip access-group SALES in
!
<output omitted>
ip access-list standard SALES
   permit 10.3.3.1
```

```
    permit 10.4.4.1
    permit 10.5.5.1
    deny 10.1.1.0 0.0.0.255
    permit any
!
ip access-list extended ENG
    permit tcp host 10.22.22.1 any eq telnet
    permit tcp host 10.33.33.1 any eq ftp
    permit tcp host 10.33.33.1 any eq ftp-data
!
<output omitted>
```

Comparing IPv4 and IPv6 ACLs

IPv4 and IPv6 ACLs have some subtle differences (see Table 9-2).

Table 9-2 *IPv4 and IPv6 ACLs*

Technet24

Feature	IPv4 Only	IPv6 Only	Both
Match source and/or destination address			X
Match host addresses or subnets/prefixes			X
Applied directionally on an interface			X
Match TCP or UDP source and/or destination ports			X
Match ICMP codes			X
Include implicit deny at end of ACL			X
Match IPv4 packets only	X		
Match IPv6 packets only		X	
Use numbers to identify the ACL	X		
Use names to identify the ACL			X
Include some implicit **permit** statements at end of ACL		X	

Configuring IPv6 ACLs

The basic steps to configure IPv6 ACLs are the same as for named IPv4 ACLs:

Name the ACL.

Create the ACL.

Apply the ACL.

Step 1: Name the IPv6 ACL

To name an IPv6 ACL, enter the **ipv6 access-list** command in global configuration mode:

```
Router(config)# ipv6 access-list name
```

Notice that the command syntax to name an IPv6 ACL is the same whether you are configuring standard or extended IPv6 ACLs. However, standard and extended IPv6 ACLs are different from standard and extended IPv4 ACLs.

Step 2: Create the IPv6 ACL

A standard IPv6 ACL includes both source and destination address information, but it does not include TCP, UDP, or ICMPv6 information. The syntax for a standard IPv6 ACL follows:

Technet24

```
Router(config-ipv6-acl)# [permit | deny] ipv6 {source-
  | any | host source-ipv6-address} {destination-ipv6-
  | host destination-ipv6-address} [log]
```

Extended IPv6 ACLs match on many more IPv6 packet header fields, as well as TCP, UDP, and ICMPv6 messages and IPv6 extension headers. The syntax for extended IPv6 ACLs follows:

```
Router(config-ipv6-acl)# [permit | deny] protocol {sou
  prefix-length | any | host source-ipv6-address} [ope
  {destination-ipv6- prefix/prefix-length | any | host
  destination-ipv6-address} [operator [port- number]]
  [doh-number |doh-type]] [dscp value] [flow-label val
  [log-input] [mobility][mobility-type [mh-number | mh
  [timeout value]] [routing][routing-type routing-numb
  [time-range name]
```

If you choose **icmp**, **tcp**, or **udp** as the *protocol*, additional filtering options are available to match those specific headers. For example, configuring **icmp** as the *protocol* enables you to filter the *icmp-type*. Configuring **tcp** as the *protocol* enables you to filter the six TCP flags, ACK, FIN, PSH, RST, SYN, and URG. Configuring **udp** as the *protocol* enables you to filter IPv6 extension headers, such as IPsec.

Step 3: Apply the IPv6 ACL

The syntax to apply an IPv6 ACL to an interface follows:

```
Router(config-if)# ipv6 traffic-filter access-list-nam
```

The syntax to apply an IPv6 ACL to vty lines is similar to that of IPv4. Just replace **ip** with **ipv6**, as follows:

```
Router(config-line)# ipv6 access-class access-list-nam
```

Standard IPv6 ACL: Allow SSH Remote Access

The topology in Figure 9-2 is used here for IPv6 ACL configuration scenarios.

Technet24

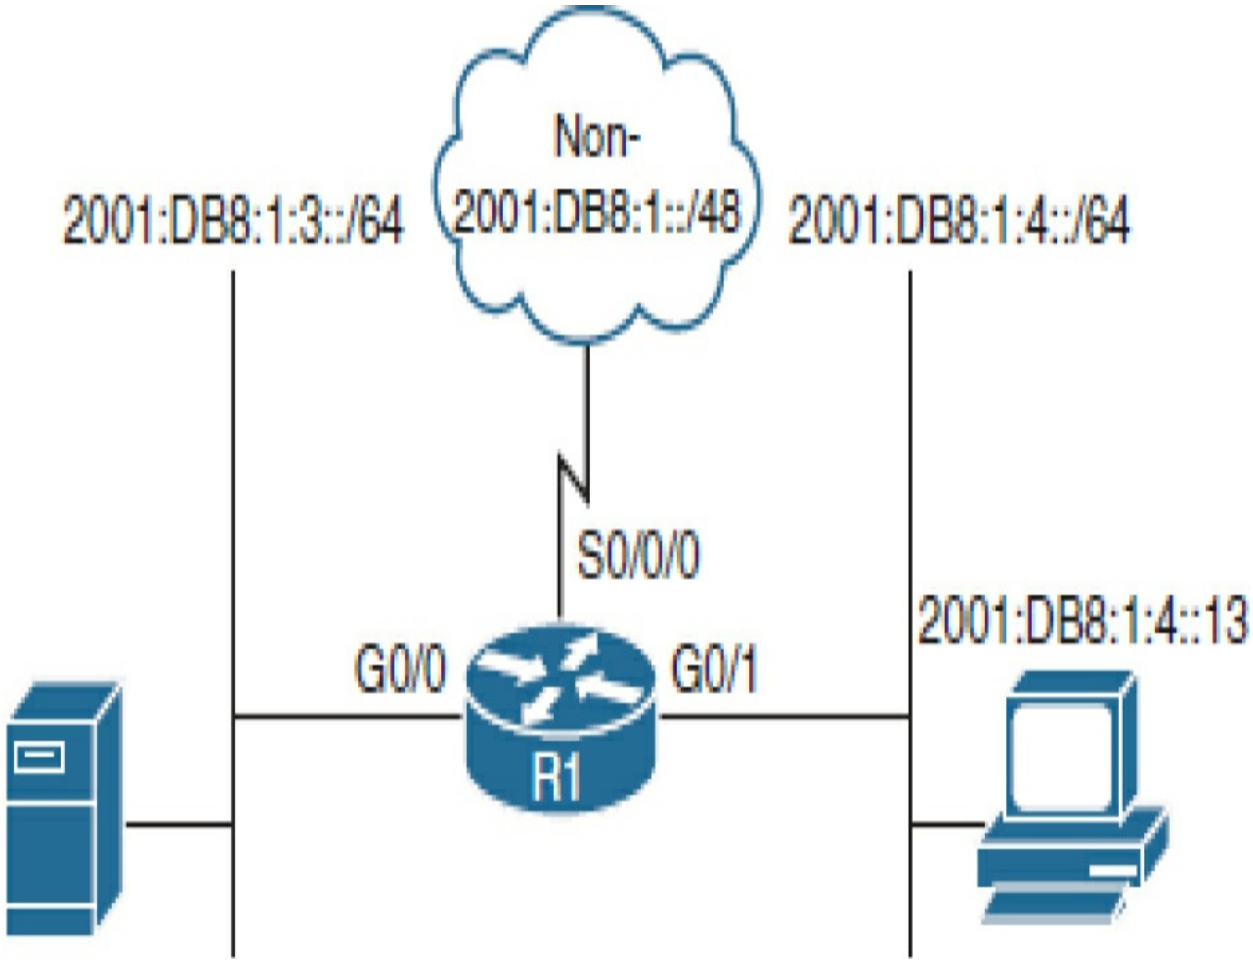

Figure 9-2 *IPv6 ACL Configuration Topology*

Example 9-12 demonstrates how to create and apply an IPv6 ACL to allow only 2001:DB8:1:4::13 to remotely connect to R1 vty lines.

Example 9-12 *Access List Allowing Only One Host to Access R1*

```
R1(config)# ipv6 access-list SSH-HOST
R1(config-ipv6-acl)# permit ipv6 host 2001:db8:1:4::13
R1(config-ipv6-acl)# deny ipv6 any any
R1(config-ipv6-acl)# exit
R1(config)# line vty 0 4
```

```
R1(config-line)# ipv6 access-class SSH-HOST in
```

The **permit** statement allows only one host, 2001:DB8:1:4::13. All other IPv6 traffic is denied. The IPv6 ACL is then applied to the first five vty lines with the **ipv6 access-class** command.

Extended IPv6 ACL: Allow Only Web Traffic

Example 9-13 demonstrates how to create and apply an IPv6 ACL to allow only web traffic originating from the subnet 2001:DB8:1:4::/64 and going to the 2001:DB8:1:3::/64 subnet.

Example 9-13 *Access List Preventing Web Traffic from Specific Subnets*

```
R1(config)# ipv6 access-list WEB-ONLY
R1(config-ipv6-acl)# permit tcp 2001:DB8:1:4::/64 2001
R1(config-ipv6-acl)# deny ipv6 2001:DB8:1:4::/64 2001:
R1(config-ipv6-acl)# permit ipv6 2001:DB8:1:4::/64 any
R1(config-ipv6-acl)# exit
R1(config)# interface g0/1
R1(config-if)# ipv6 traffic-filter WEB-ONLY in
```

The first **permit** statement allows traffic from prefix 2001:DB8:1:4::/64 to access web services on any device in prefix 2001:DB8:1:3::/64. The **deny** statement ensures that all other traffic from 2001:DB8:1:4::/64 to

Technet24

2001:DB8:1:3::/64 is blocked. The last **permit** statement allows all other inbound traffic from 2001:DB8:1:4::/64 to any destination.

Verifying IPv6 ACLs

As with IPv4 ACLs, you can view the configuration and application of IPv6 ACLs with the **show run** command (see Example 9-14).

Example 9-14 *Examining ACLs in the Configuration*

```
R1# show run
Building configuration...
<some output omitted>
!
interface GigabitEthernet0/1
  ipv6 traffic-filter WEB-ONLY in
  ipv6 address FE80::1 link-local
  ipv6 address 2001:DB8:1:4::1/64
!
ipv6 access-list WEB-ONLY
  permit tcp 2001:DB8:1:4::/64 2001:DB8:1:3::/64 eq ww
  deny ipv6 2001:DB8:1:4::/64 2001:DB8:1:3::/64
  permit ipv6 2001:DB8:1:4::/64 any
ipv6 access-list SSH-HOST
  permit ipv6 host 2001:DB8:1:4::13 any
  deny ipv6 any any
!
line vty 0 4
```

```
   ipv6 access-class SSH-HOST in
   login local
   transport input ssh
!
R1#
```

However, the configuration of production routers is usually long and complex. For the simulation questions on the CCNA exam, you might not even have access to the **show run** command. Therefore, you should use verification commands that more precisely and efficiently provide the information you need.

For example, **show access-lists** quickly reveals all IPv4 and IPv6 ACLs configured on the device, as Example 9-15 shows.

Example 9-15 *Verifying All ACLs Configured*

```
R1# show access-lists
Standard IP access list SALES
   10 permit 10.3.3.1
   20 permit 10.4.4.1
   30 permit 10.5.5.1
   40 deny 10.1.1.0, wildcard bits 0.0.0.255
   50 permit any
Extended IP access list ENG
   10 permit tcp host 10.22.22.1 any eq telnet
```

Technet24

```
    20 permit tcp host 10.33.33.1 any eq ftp
    30 permit tcp host 10.33.33.1 any eq ftp-data
IPv6 access list SSH-HOST
    permit ipv6 host 2001:DB8:1:4::13 any sequence 10
    deny ipv6 any any sequence 20
IPv6 access list WEB-ONLY
    permit tcp 2001:DB8:1:4::/64 2001:DB8:1:3::/64 eq w
    deny ipv6 2001:DB8:1:4::/64 2001:DB8:1:3::/64 seque
    permit ipv6 2001:DB8:1:4::/64 any sequence 30
R1#
```

Notice that Cisco IOS added sequence numbers to the end of the IPv6 ACLs instead of at the beginning, as it does for IPv4 ACLs.

In Example 9-16, only the IPv6 ACLs are shown. This output was generated after multiple packets matched each of the statements in the ACLs.

Example 9-16 *Verifying IPv6 ACL Match Statistics*

```
R1# show ipv6 access-list
IPv6 access list SSH-HOST
     permit ipv6 host 2001:DB8:1:4::13 any (1 match(es
      deny ipv6 any any (5 match(es)) sequence 20
IPv6 access list WEB-ONLY
      permit tcp 2001:DB8:1:4::/64 2001:DB8:1:3::/64 e
sequence 10
      deny ipv6 2001:DB8:1:4::/64 2001:DB8:1:3::/64 (4
```

```
        permit ipv6 2001:DB8:1:4::/64 any (75 match(es))
R1#
```

To verify the placement of an IPv6 ACL on an interface, you can use the **show ipv6 interface** command. If an ACL is applied, the output will have a line entry like the one highlighted in Example 9-17.

Example 9-17 *Verifying an IPv6 ACL Applied to an Interface*

```
R1# show ipv6 interface g0/1
GigabitEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::1
  No Virtual link-local address(es):
  Global unicast address(es):
    2001:DB8:1:4::1, subnet is 2001:DB8:1:4::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 millise
  ICMP redirects are enabled
  ICMP unreachables are sent
  Input features: Access List
  Inbound access list WEB-ONLY
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds (using 30000
```

Technet24

```
    ND advertised reachable time is 0 (unspecified)
    ND advertised retransmit interval is 0 (unspecified)
    ND router advertisements are sent every 200 seconds
    ND router advertisements live for 1800 seconds
    ND advertised default router preference is Medium
    Hosts use stateless autoconfig for addresses.
  R1#
```

Troubleshooting ACLs

Your network can be configured correctly with all hosts receiving DHCP addressing, fully populated routing tables, and a fully operating physical layer, but an ACL somewhere in the data path can still be causing a problem. Troubleshooting a problem caused by an ACL can make your job more difficult.

ACLs can block normal troubleshooting tools such as **ping** and **traceroute** while still allowing normal traffic. Therefore, a network administrator might need to rely on other tools to find a problem.

Say that you have determined that a problem is with the ACL configuration. The following three steps summarize a structured troubleshooting process you can use to track down the issue:

Because ACL configurations cannot cause a problem until they are applied, determine what interfaces are impacted by ACLs by using the **show run** or

show ip interfaces command.

Verify the ACL configuration by using the **show access-lists**, **show ip access-lists**, or **show run** command.

Analyze the ACLs to determine which packets will match. The **show access-lists** and **show ip access-lists** commands help by identifying the number of times packets have matched a statement.

Some common ACL configuration errors include the following:

• ACL statements are out of order.

• The source and destination addresses and/or ports are reversed.

• The ACL is applied in the wrong direction.

• Syntax or spelling errors cause the ACL to have the wrong intended effect or no effect.

• Standard ACLs are close to the source instead of to the destination.

Refer to your study resources for several excellent troubleshooting examples for both IPv4 and IPv6 ACLs.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Technet24

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	5
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	7–8
Portable Command Guide	21

Note

In the final section of Chapter 8, “Applied IP ACLs,” in the *CCNA 200-301 Official Cert Guide, Volume 2*, 2nd edition, Wendell Odom explains:

This final section of the chapter summarizes some of the current differences between the recent versions of IOS and IOS XE in early 2023 when I most recently revised this chapter. I did not choose to add this section because of any exam topic. More generally, the exam topics do not mention any need to compare ACLs on IOS versus IOS XE. Cisco historically avoids asking anything to do with IOS software versions. However, you may practice ACLs with IOS, but then

an exam question may use output from IOS XE, or vice versa.

Of course by now you know that I strongly recommend that you have Odom’s two-volume certification library. So go now and read through his content on the difference between IOS and IOS XE ACL implementations.

Technet24

Day 8. NAT

CCNA 200-301 Exam Topics

• 4.1 Configure and verify inside source NAT using static and pools

Key Topics

To cope with the depletion of IPv4 addresses, several short-term solutions were developed. One short-term solution is to use private addresses and Network Address Translation (NAT). NAT enables inside network hosts to borrow a legitimate Internet IPv4 address while accessing Internet resources. When the requested traffic returns, the legitimate IPv4 address is repurposed and available for the next Internet request by an inside host. Using NAT, network administrators need only one or a few IPv4 addresses for the router to provide to the hosts instead of one unique IPv4 address for every client joining the network. Although IPv6 ultimately solves the problem of IPv4 address space depletion that NAT was created to address, it is still in wide use in current network implementation strategies. Today we review the concepts, configuration, and troubleshooting of NAT.

NAT Concepts

NAT, defined in RFC 3022, has many uses. Its key use is to conserve IPv4 addresses by allowing networks to use private IPv4 addresses. NAT translates nonroutable, private, internal addresses into routable public addresses. NAT

also has the benefit of hiding internal IPv4 addresses from outside networks.

A NAT-enabled device typically operates at the border of a stub network. Figure 8-1 shows the master topology used during today's review. R2 is the border router and is the device used for today's example configurations.

Technet24

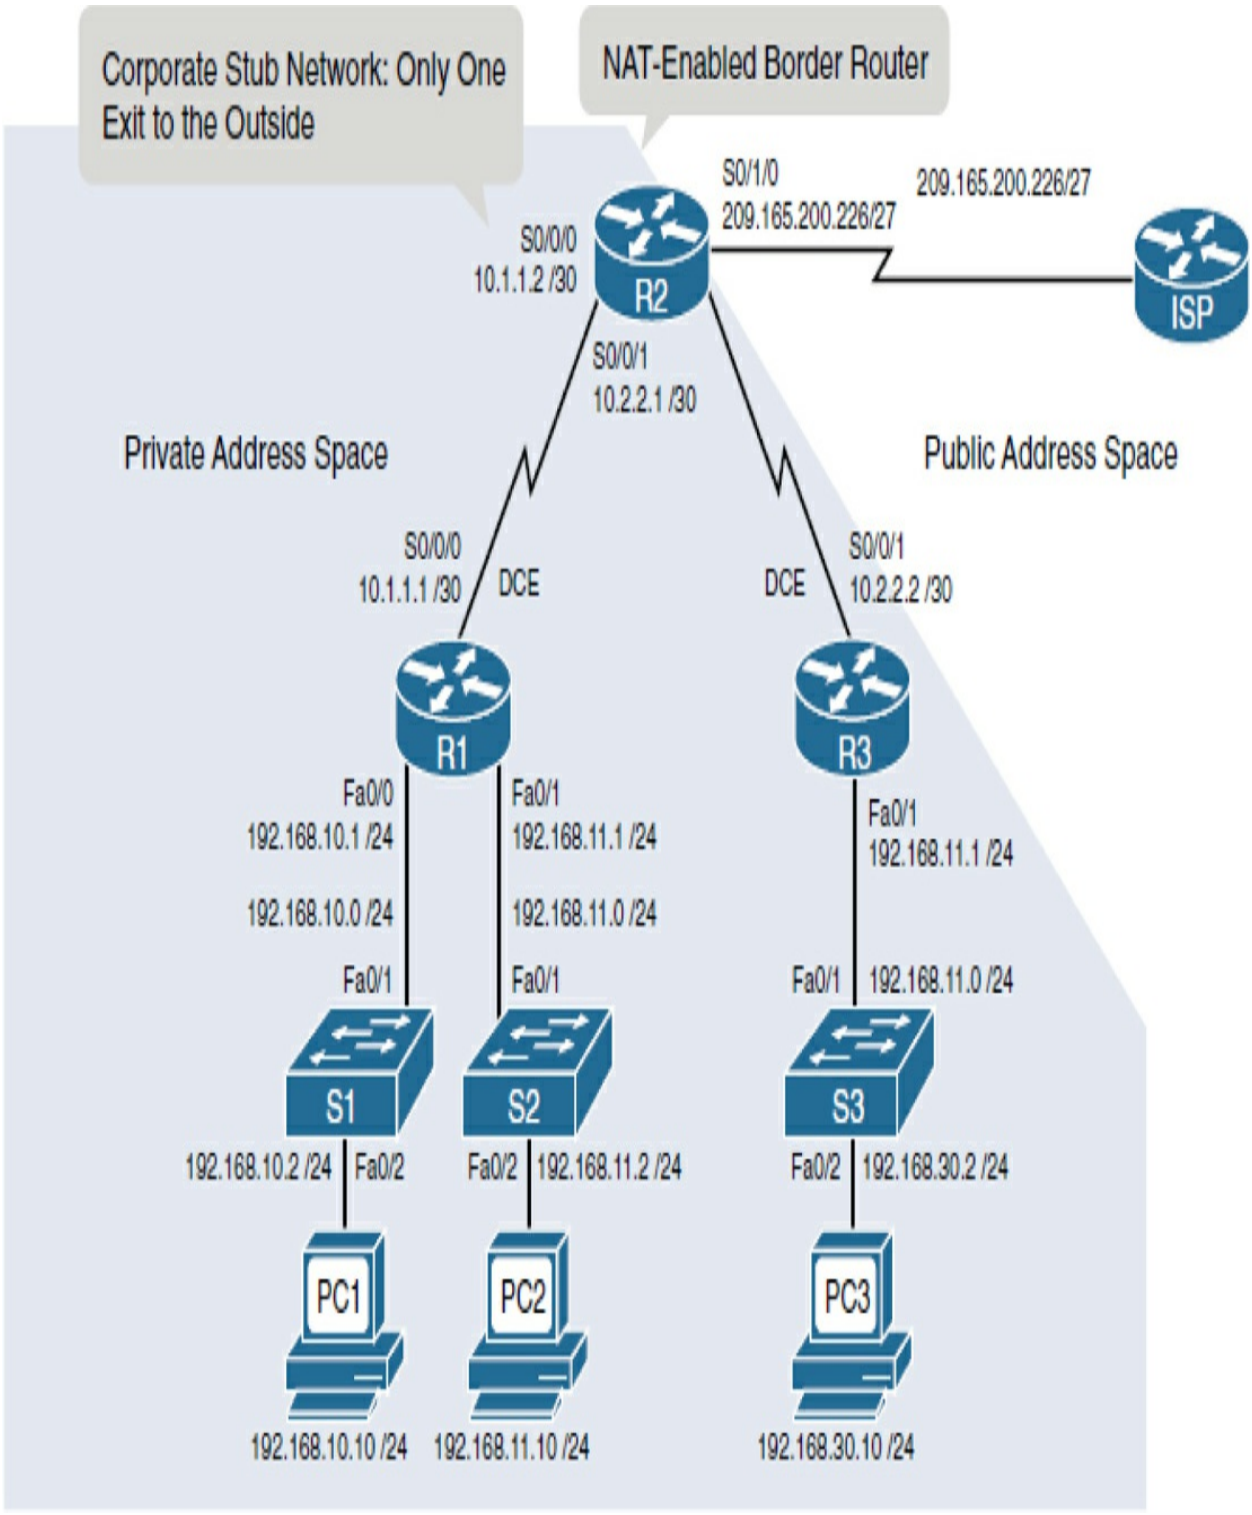

Figure 8-1 *NAT Topology*

In NAT terminology, the *inside network* is the set of networks that are subject

to translation (every network in the shaded region in Figure 8-1). The *outside network* is all other addresses. Figure 8-2 shows how to refer to the addresses when configuring NAT:

• **Inside local address:** Most likely a private address. In the figure, the IPv4 address 192.168.10.10 assigned to PC1 is an inside local address.

• **Inside global address:** A valid public address that the inside host is given when it exits the NAT router. When traffic from PC1 is destined for the web server at 209.165.201.1, R2 must translate the inside local address to an inside global address, which is 209.165.200.226 in this case.

• **Outside global address:** A reachable IPv4 address assigned to a host on the Internet. For example, the web server can be reached at IPv4 address 209.165.201.1.

• **Outside local address:** The local IPv4 address assigned to a host on the outside network. In most situations, this address is identical to the outside global address of that outside device. (Outside local addresses are beyond the scope of the CCNA.)

Technet24

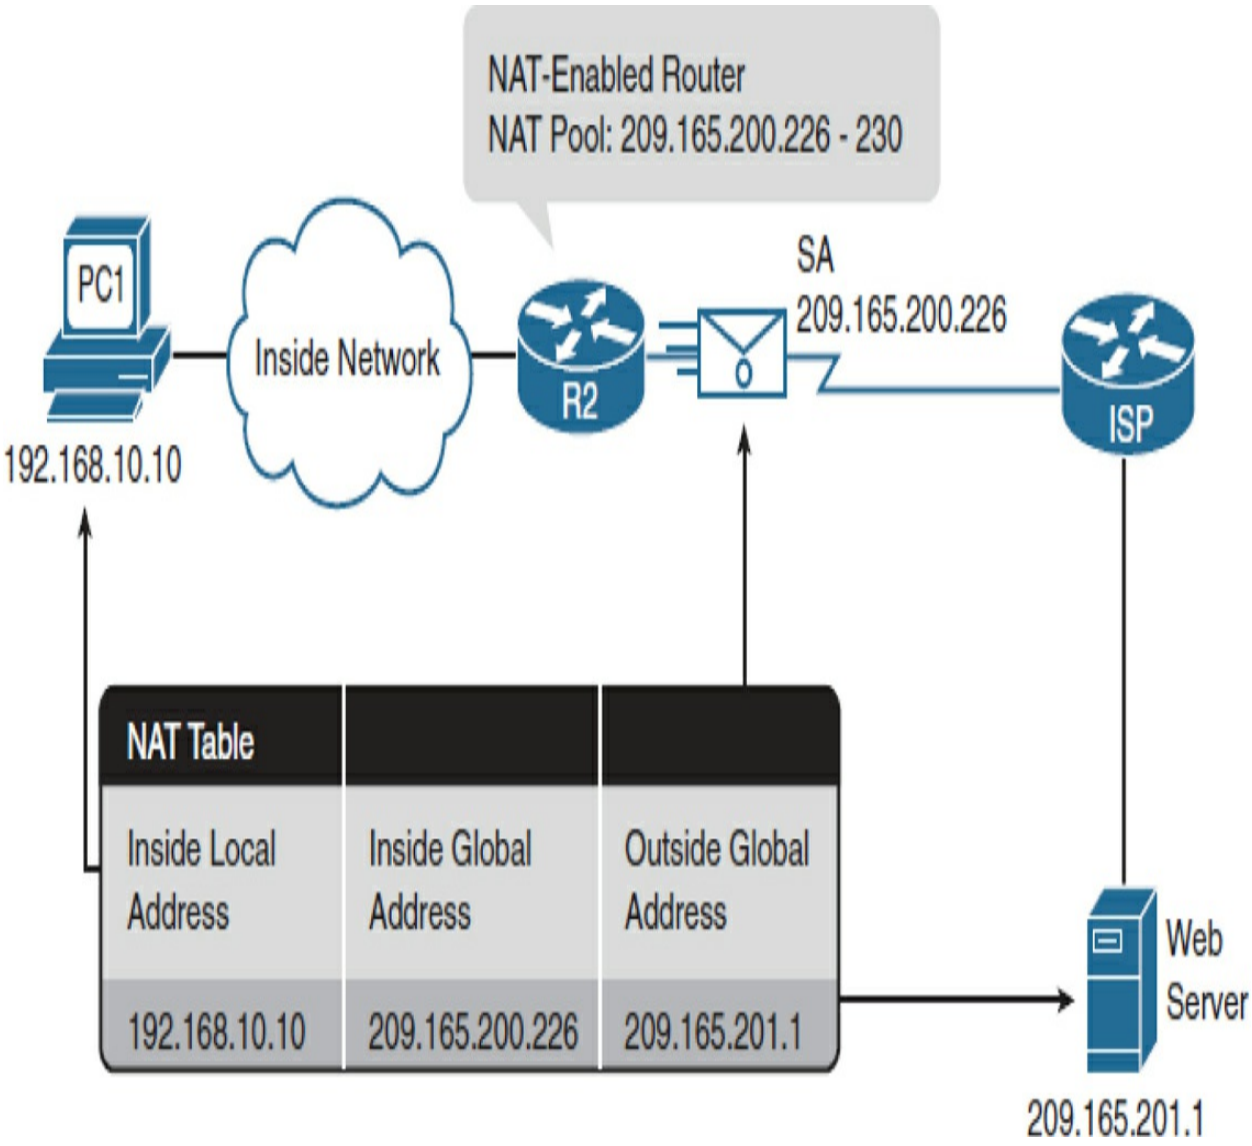

Figure 8-2 *NAT Terminology*

A NAT Example

Referring to Figure 8-1, the following steps illustrate the NAT process when PC1 sends traffic to the Internet:

PC1 sends a packet destined for the Internet to R1, the default gateway.

R1 forwards the packet to R2, as directed by its routing table.

R2 refers to its routing table and identifies the next hop as the ISP router. It then checks to see whether the packet matches the criteria specified for translation. R2 has an ACL that identifies the inside network as a valid host for translation. Therefore, it translates an inside local IPv4 address to an inside global IPv4 address, which, in this case, is 209.165.200.226. It stores this mapping of the local address to global address in the NAT table.

R2 modifies the packet with the new source IPv4 address (the inside global address) and sends it to the ISP router.

The packet eventually reaches its destination, which then sends its reply to the inside global address 209.165.200.226.

When R2 receives replies from the destination, it consults the NAT table to match the inside global address to the correct inside local address. R2 then modifies the packet, inserting the inside local address (192.168.10.10) as the destination address and sending it to R1.

R1 receives the packet and forwards it to PC1.

Dynamic and Static NAT

The two types of NAT translation are as follows:

• **Dynamic NAT:** Uses a pool of public addresses and assigns them on a first-

Technet24

come, first-served basis or reuses an existing public address configured on an interface. When a host with a private IPv4 address requests access to the Internet, dynamic NAT chooses an IPv4 address from the pool that another host is not already using. Instead of using a pool, dynamic NAT can be configured to overload an existing public address configured on an interface.

• **Static NAT:** Uses a one-to-one mapping of local and global addresses. These mappings remain constant. Static NAT is particularly useful for web servers or hosts that must have a consistent address that is accessible from the Internet.

NAT Overload

NAT overloading (also called Port Address Translation [PAT]) maps multiple private IPv4 addresses to a single public IPv4 address or a few addresses. To do this, a port number also tracks each private address. When a response comes back from the outside, source port numbers determine the correct client for the NAT router to translate the packets.

Figure 8-3 and the following steps illustrate the NAT overload process:

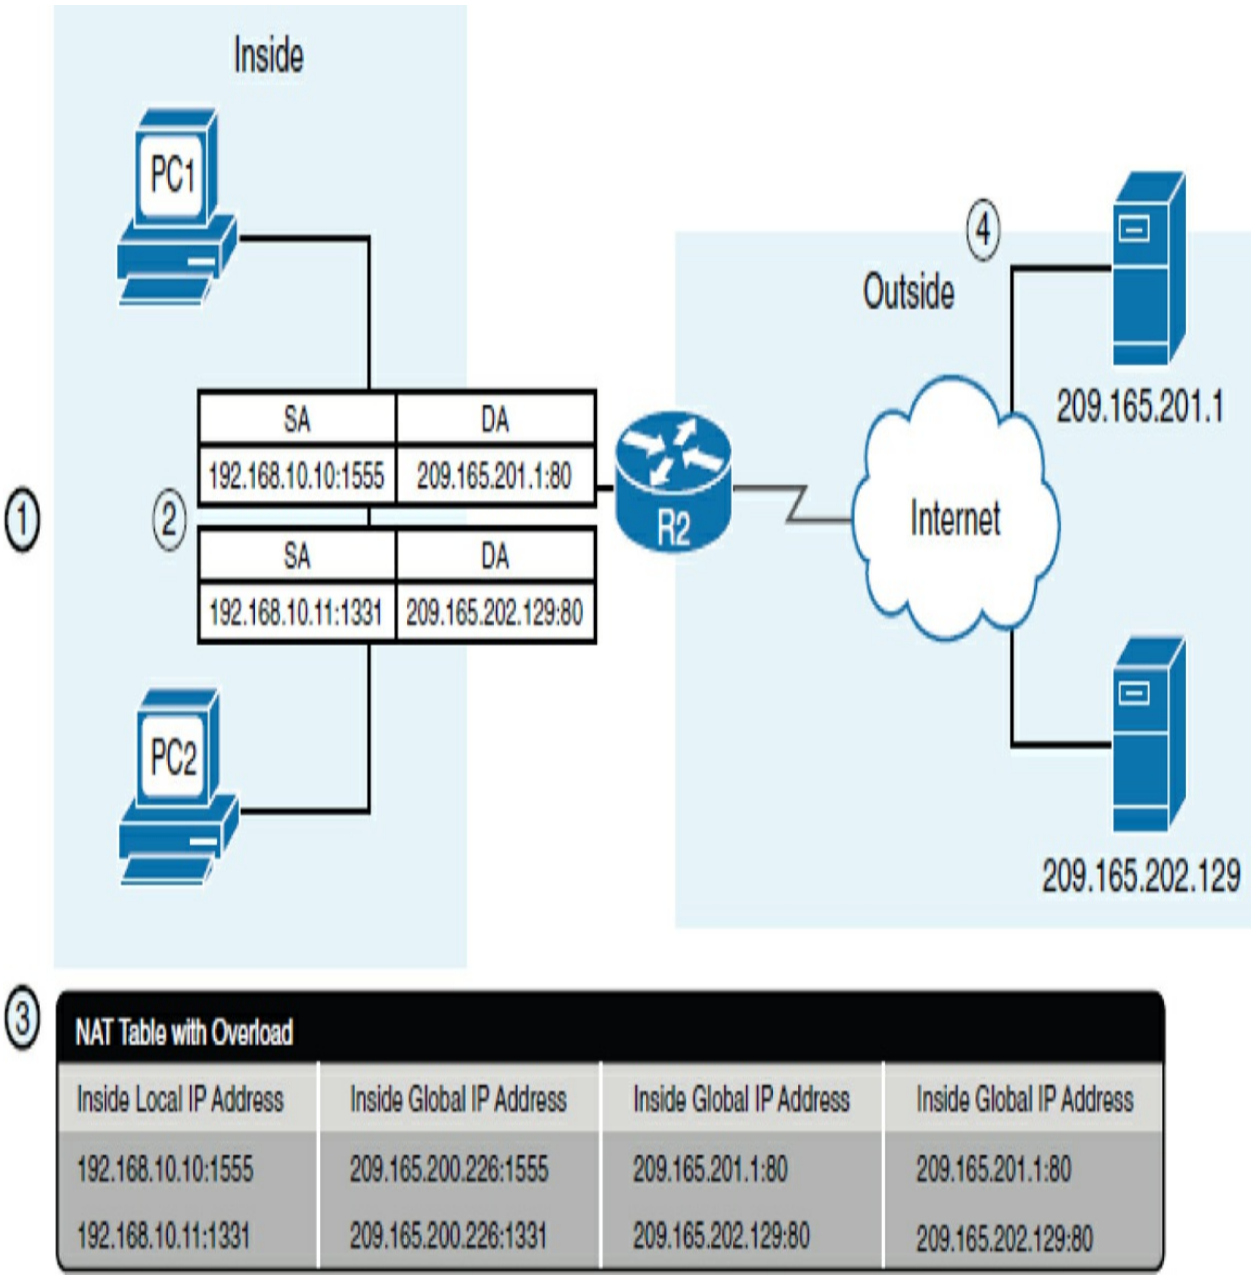

Figure 8-3 *NAT Overload Example*

PC1 and PC2 send packets destined for the Internet.

When the packets arrive at R2, NAT overload changes the source address to the inside global IPv4 address and keeps a record of the assigned source port numbers (1555 and 1331, in this example) to identify the client from which

Technet24

the packets originated.

R2 updates its NAT table. Notice the assigned ports. R2 then routes the packets to the Internet.

When the web server replies, R2 uses the destination source port to translate the packet to the correct client.

NAT overload attempts to preserve the original source port. However, if this source port is already used, NAT overload assigns the first available port number, starting from the beginning of the appropriate port group 0–511, 512–1023, or 1024–65535.

NAT Benefits

Using NAT offers the following benefits:

• NAT conserves registered IPv4 address space because, with NAT overload, internal hosts can share a single public IPv4 address for all external communications.

• NAT increases the flexibility of connections to the public network. Multiple pools, backup pools, and load-balancing pools can be implemented to ensure reliable public network connections.

• NAT allows the existing scheme to remain while supporting a new public addressing scheme. This means that an organization can change ISPs without

needing to change any of its inside clients.

• NAT provides a layer of network security because private networks do not advertise their inside local addresses outside the organization. However, the phrase *NAT firewall* is misleading; NAT does not replace firewalls.

NAT Limitations

The limitations of using NAT include the following:

• **Performance is degraded:** NAT increases switching delays because translating each IPv4 address within the packet headers takes time.

• **End-to-end functionality is degraded:** Many Internet protocols and applications depend on end-to-end functionality, with unmodified packets forwarded from the source to the destination.

• **End-to-end IP traceability is lost:** Tracing packets that undergo numerous packet address changes over multiple NAT hops becomes much more difficult, making troubleshooting challenging.

• **Tunneling is more complicated:** Using NAT also complicates tunneling protocols, such as IPsec, because NAT modifies values in the headers that interfere with the integrity checks that IPsec and other tunneling protocols do.

• **Services can be disrupted:** Services that require the initiation of TCP connections from the outside network, or stateless protocols such as those

Technet24

using UDP, can be disrupted.

Configuring Static NAT

Static NAT is a one-to-one mapping between an inside address and an outside address. Static NAT allows connections initiated by external devices to access inside devices. For example, you might want to map an inside global address to a specific inside local address that is assigned to your inside web server. The steps and syntax to configure static NAT follow:

Configure the static translation of an inside local address to an inside global address:

```
Router(config)# ip nat inside source static
local-ip global-ip
```

Specify the inside interface:

```
Router(config)# interface type number
Router(config-if)# ip nat inside
```

Specify the outside interface:

```
Router(config)# interface type number
Router(config-if)# ip nat outside
```

Figure 8-4 shows a sample static NAT topology.

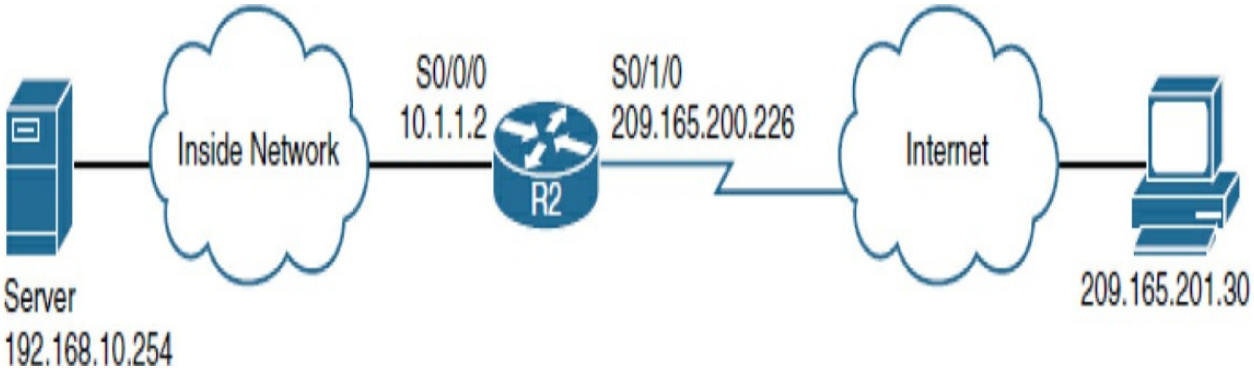

Figure 8-4 *Static NAT Topology*

Example 8-1 shows the static NAT configuration.

Example 8-1 *Static NAT Configuration*

```
R2(config)# ip nat inside source static 192.168.10.254
R2(config)# interface serial0/0/0
R2(config-if)# ip nat inside
R2(config-if)# interface serial 0/1/0
R2(config-if)# ip nat outside
```

This configuration statically maps the inside private IPv4 address 192.168.10.254 to the outside public IPv4 address 209.165.200.254. This allows outside hosts to access the internal web server by using the public IPv4 address 209.165.200.254.

Configuring Dynamic NAT

Technet24

Dynamic NAT maps private IPv4 addresses to public addresses drawn from a NAT pool. The steps and syntax to configure dynamic NAT are as follows:

Define a pool of global addresses to be allocated:

```
Router(config)# ip nat pool name start-ip end-ip {
netmask | prefix-length prefix-length}
```

Define a standard access list permitting addresses that are to be translated:

```
Router(config)# access-list
access-list-number source source-wildcard
```

Bind the pool of addresses to the access list:

```
Router(config)# ip nat inside source list
access-list-number pool name
```

Specify the inside interface:

```
Router(config)# interface type number
Router(config-if)# ip nat inside
```

Specify the outside interface:

```
Router(config)# interface type number
Router(config-if)# ip nat outside
```

Figure 8-5 shows a sample dynamic NAT topology.

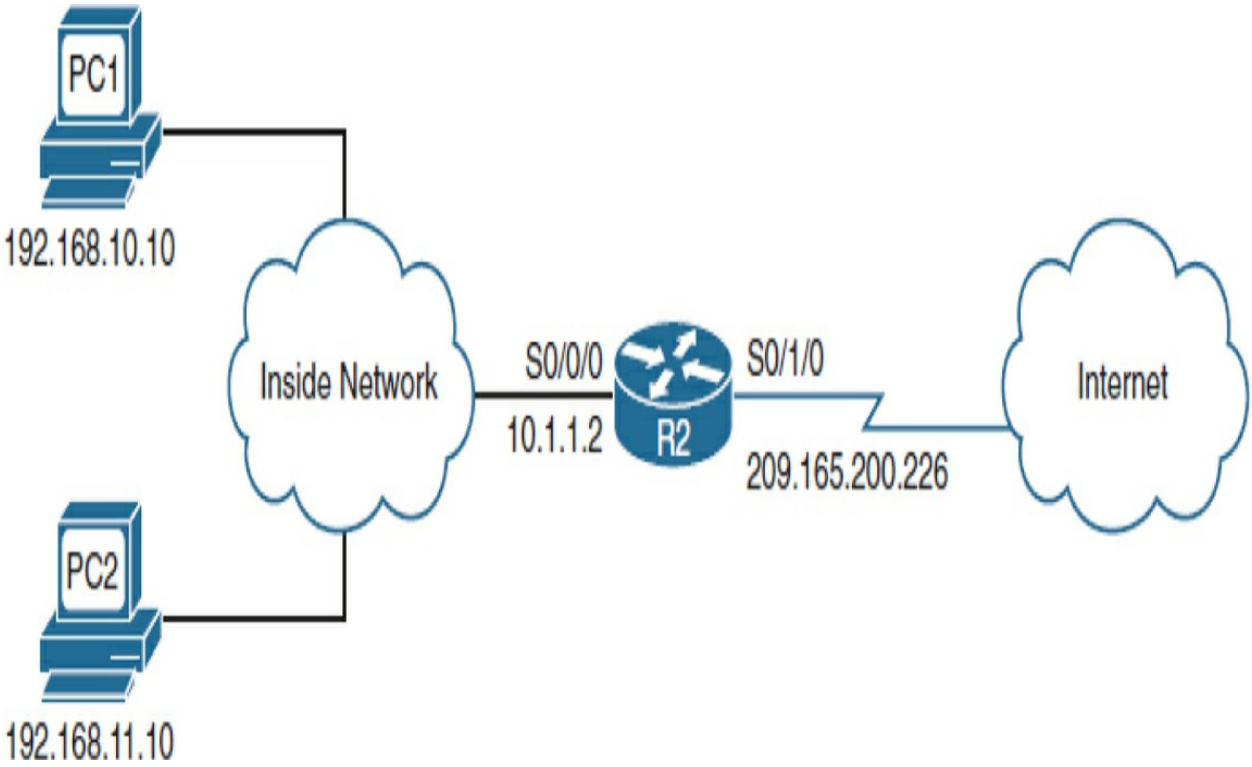

Figure 8-5 *Dynamic NAT Topology*

Example 8-2 shows the dynamic NAT configuration.

Example 8-2 *Dynamic NAT Configuration*

```
R2(config)# ip nat pool NAT-POOL1 209.165.200.226 209.
255.255.255.224
R2(config)# access-list 1 permit 192.168.0.0 0.0.255.2
```

Technet24

```
R2(config)# ip nat inside source list 1 pool NAT-POOL1
R2(config)# interface serial 0/0/0
R2(config-if)# ip nat inside
R2(config-if)# interface serial s0/1/0
R2(config-if)# ip nat outside
```

Configuring NAT Overload

Commonly with home networks and small to medium-size businesses, the ISP assigns only one registered IPv4 address to the router. Therefore, it is necessary to overload that one IPv4 address so that multiple inside clients can use it simultaneously.

The configuration is similar to the use of dynamic NAT, except that instead of using a pool of addresses, the interface keyword is used to identify the outside IPv4 address. The **overload** keyword enables PAT so that source port numbers are tracked during translation.

Example 8-3 shows how R2 in Figure 8-5 is configured to overload its registered IPv4 address on the serial interface.

Example 8-3 *Configuring NAT to Overload an Interface Address*

```
R2(config)# access-list 1 permit 192.168.0.0 0.0.255.2
R2(config)# ip nat inside source list 1 interface seri
R2(config)# interface serial 0/0/0
```

```
R2(config-if)# ip nat inside
R2(config-if)# interface serial s0/1/0
R2(config-if)# ip nat outside
```

You can also overload a NAT pool of addresses, which might be necessary in organizations that potentially have many clients that simultaneously need translations. In Example 8-2, NAT is configured with a pool of 15 addresses (209.165.200.226–209.165.200.240). If R2 is translating all 15 addresses at any given moment, packets for the 16th client will be queued for processing and will possibly time out. To avoid this problem, add the keyword **overload** to the command that binds the access list to the NAT pool, as follows:

```
R2(config)# ip nat inside source list 1 pool NAT-POOL1
```

Interestingly, Cisco IOS uses the first IPv4 address in the pool until it runs out of available port numbers. Then it moves to the next IPv4 address in the pool.

Verifying NAT

Assume that both the static and dynamic NAT topologies in Figures 6-4 and 6-5 are configured on R2, with the inside server statically translated to 209.165.200.254 and **NAT-POOL1** configured with the **overload** keyword. Furthermore, assume that two inside clients have connected to an outside

Technet24

host. You can use the **show ip nat translations** command to verify the current translations in the R2 NAT table, as Example 8-4 shows.

Example 8-4 *Verifying NAT Operations with* **show ip nat translations**

```
R2# show ip nat translations
Pro Inside global          Inside local         Outsid
--- 209.165.200.254        192.168.10.254       ---
tcp 209.165.200.226:47392  192.168.10.10:47392  209.16
tcp 209.165.200.226:50243  192.168.11.10:50243  209.16
```

The static entry is always in the table. Currently, the table has two dynamic entries. Notice that both inside clients received the same inside global address, but the source port numbers are different (47392 for PC1 and 50243 for PC2).

The **show ip nat statistics** command, shown in Example 8-5, displays information about the total number of active translations, NAT configuration parameters, the number of addresses in the pool, and how many have been allocated.

Example 8-5 *Verifying NAT Operations with* **show ip nat statistics**

```
R2# show ip nat statistics
Total translations: 3 (1 static, 2 dynamic, 2 extended
```

```
  Outside Interfaces: Serial0/1/0
  Inside Interfaces: FastEthernet0/0 , Serial0/0/0 , Ser
  Hits: 29 Misses: 7
  Expired translations: 5
  Dynamic mappings:
  -- Inside Source
  access-list 1 pool NAT-POOL1 refCount 2
    pool NAT-POOL1: netmask 255.255.255.224
          start 209.165.200.226 end 209.165.200.240
             type generic, total addresses 3 , allocated
```

Alternatively, use the **show run** command and look for NAT, access command list, interface, or pool-related commands with the required values. Carefully examine the output from these commands to discover any errors.

It is sometimes useful to clear the dynamic entries sooner than the default. This is especially true when testing the NAT configuration. To clear dynamic entries before the timeout has expired, use the **clear ip nat translation *** privileged EXEC command.

Troubleshooting NAT

When you have IP connectivity problems in a NAT environment, determining the cause of a problem is often difficult. The first step in solving your problem is to rule out NAT as the cause. Follow these steps to verify that NAT is operating as expected:

Technet24

Based on the configuration, clearly define what NAT is supposed to achieve. This might reveal a problem with the configuration.

Verify that correct translations exist in the translation table by using the **show ip nat translations** command.

Use the **clear** and **debug** commands to verify that NAT is operating as expected. Check to see whether dynamic entries are re-created after they are cleared.

Review in detail what is happening to the packet and verify that routers have the correct routing information to forward the packet.

Use the **debug ip nat** command to verify the operation of the NAT feature by displaying information about every packet that the router translates, as in Example 8-6.

Example 8-6 *Troubleshooting NAT with* **debug ip nat**

```
R2# debug ip nat
IP NAT debugging is on
R2#
NAT: s=192.168.10.10->209.165.200.226, d=209.165.201.3
NAT*: s=209.165.201.30, d=209.165.200.226->192.168.10.
NAT: s=192.168.10.10->209.165.200.226, d=209.165.201.3
NAT: s=192.168.10.10->209.165.200.226, d=209.165.201.3
NAT*: s=209.165.201.30, d=209.165.200.226->192.168.10.
```

```
NAT*: s=209.165.201.30, d=209.165.200.226->192.168.10.
NAT: s=192.168.10.10->209.165.200.226, d=209.165.201.3
NAT: s=192.168.10.10->209.165.200.226, d=209.165.201.3
NAT*: s=209.165.201.30, d=209.165.200.226->192.168.10.
NAT*: s=209.165.201.30, d=209.165.200.226->192.168.10.
NAT: s=192.168.10.10->209.165.200.226, d=209.165.201.3
R2#
```

You can see that inside host 192.168.10.10 initiated traffic to outside host 209.165.201.30 and has been translated into address 209.165.200.226.

When decoding the debug output, note what the following symbols and values indicate:

• ***:** The asterisk next to NAT indicates that the translation is occurring in the fast-switched path. The first packet in a conversation is always process switched, which is slower. The remaining packets go through the fast-switched path if a cache entry exists.

• **s=:** This refers to the source IPv4 address.

• **a.b.c.d->w.x.y.z:** This indicates that source address **a.b.c.d** is translated into **w.x.y.z**.

• **d=:** This refers to the destination IPv4 address.

Technet24

• **[xxxx]:** The value in brackets is the IP identification number. This information can be useful for debugging because it enables correlation with other packet traces from protocol analyzers.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	6
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	14
Portable Command Guide	18

Day 7. WAN, VPN, and IPsec

CCNA 200-301 Exam Topics

• 1.2 Describe characteristics of network topology architectures

• 5.5 Describe IPsec remote access and site-to-site VPNs

Key Topics

Today is a whirlwind review of WAN topologies, WAN connection options, virtual private networks (VPNs), and Internet Protocol Security (IPsec). Most of these exam topics are conceptual in nature and require no configuration skills, so read through this review several times and refer to your study resources for more in-depth review.

WAN Topologies

Figure 7-1 shows the four basic WAN topology options that a business can select for its WAN infrastructure:

• **Point-to-point:** Typically uses a dedicated leased-line connection, such as T1/E1.

• **Hub-and-spoke:** Offers a single-homed, point-to-multipoint topology in which a single interface on the hub router can be shared with multiple spoke routers through the use of virtual interfaces.

Technet24

• **Full mesh:** Gives each router a connection to every other router. Requires a large number of virtual interfaces.

• **Dual-homed:** Provides redundancy for a single-homed hub-and-spoke topology by providing a second hub to connect to spoke routers.

A business can choose to implement a variety of these topologies. For example, an enterprise might choose to implement a full mesh topology between its regional headquarters. It might use a hub-and-spoke topology between regional headquarters and branch offices. If two of the branch offices communicate frequently, the network administrators might contract for a point-to-point link to reduce the traffic load on the hub routers. Using dual-homed connections to the Internet ensures that customers, partners, and teleworkers can always access the enterprise’s resources.

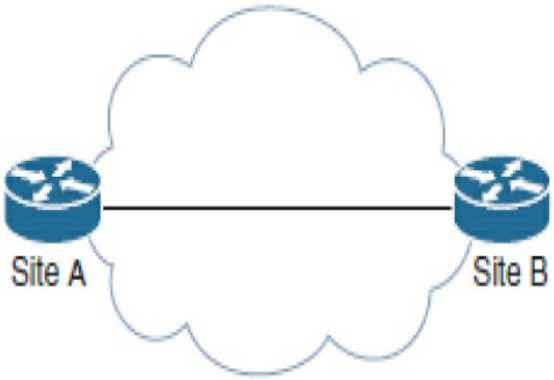

Point-to-Point

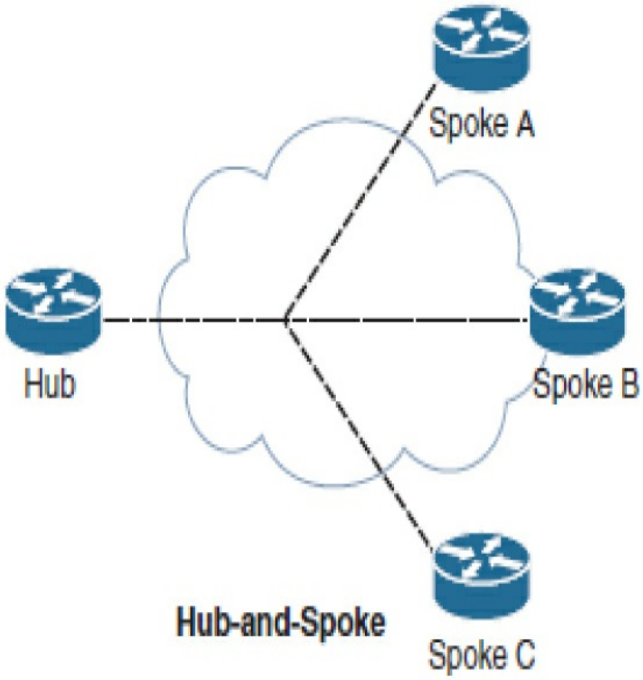

Hub-and-Spoke

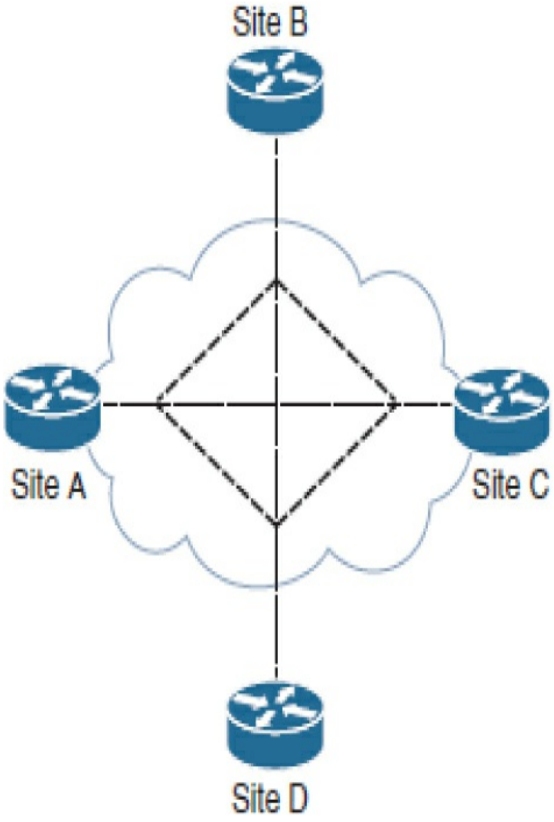

Full Mesh

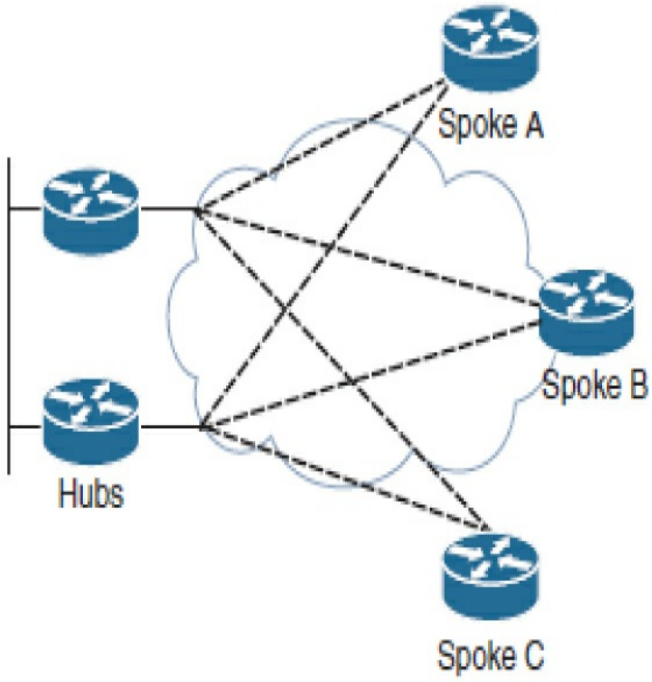

Dual-Homed

Technet24

Figure 7-1 *WAN Topology Options*

WAN Connection Options

Many options for implementing WAN solutions are currently available. They differ in technology, speed, and cost. Figure 7-2 provides a high-level view of the various WAN link connection options. The following subsections describe these options in more detail.

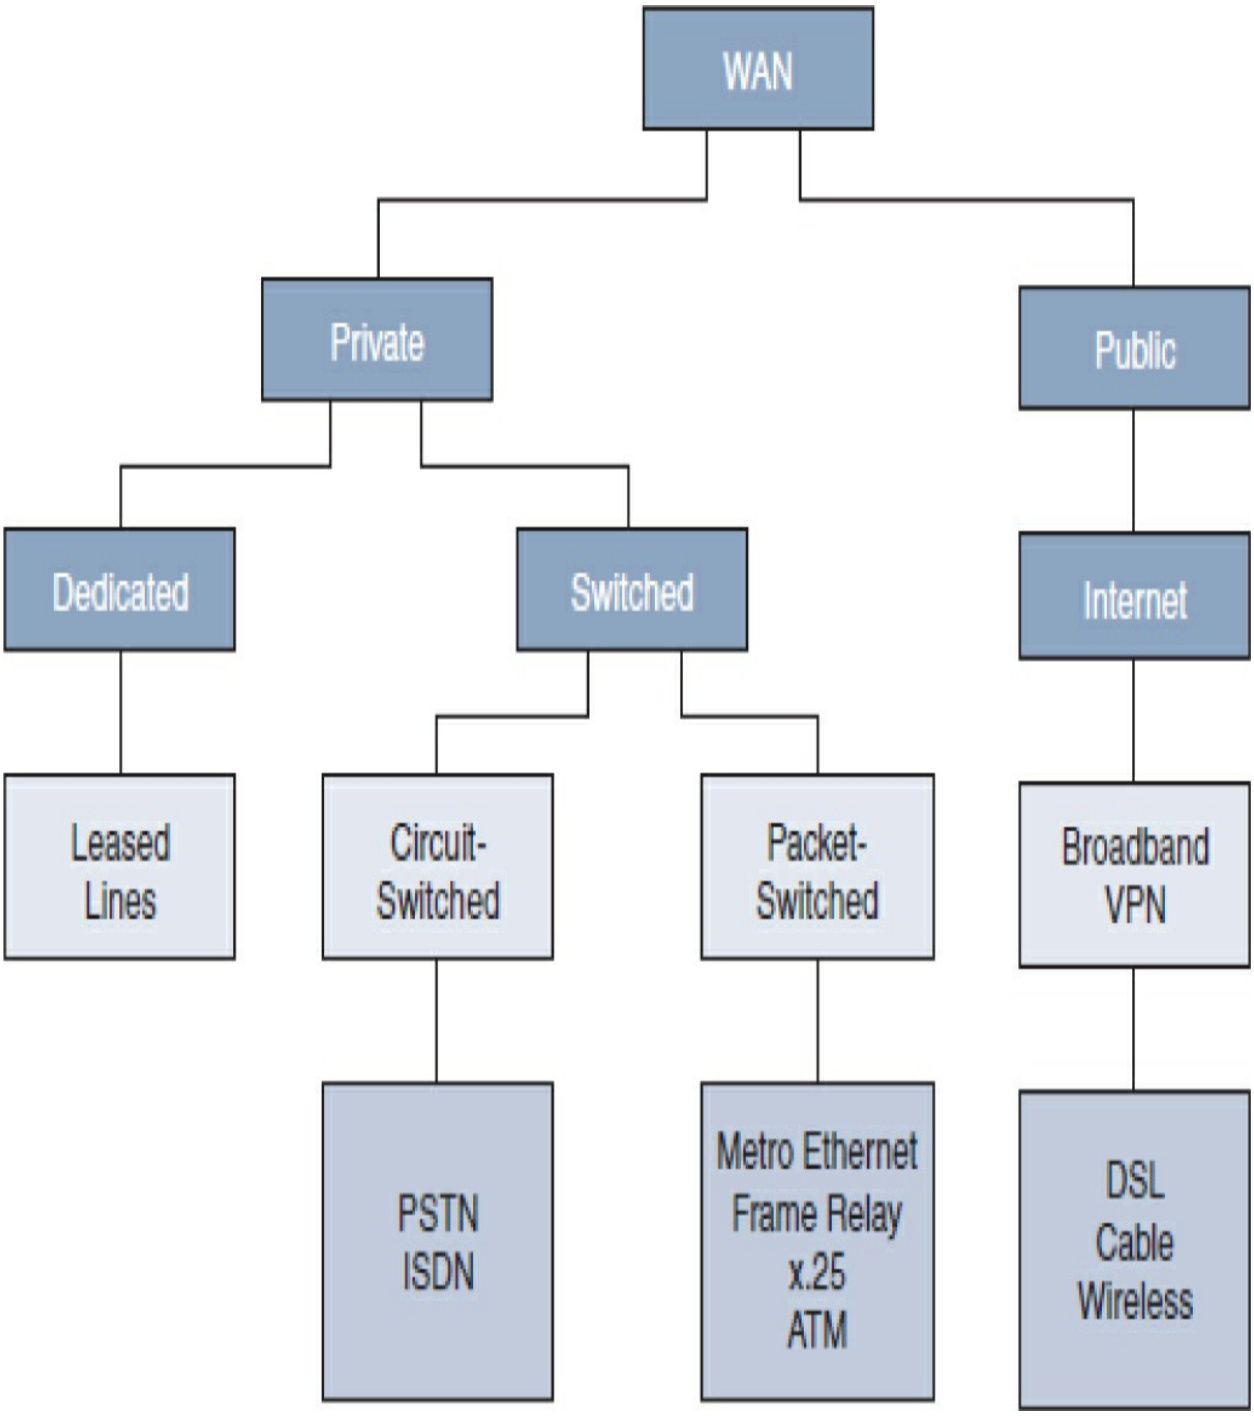

Figure 7-2 *WAN Link Connection Options*

Dedicated Connection Options

Also called leased lines, dedicated connections are pre-established point-to-

Technet24

point WAN connections from the customer premises through the provider network to a remote destination (see Figure 7-3).

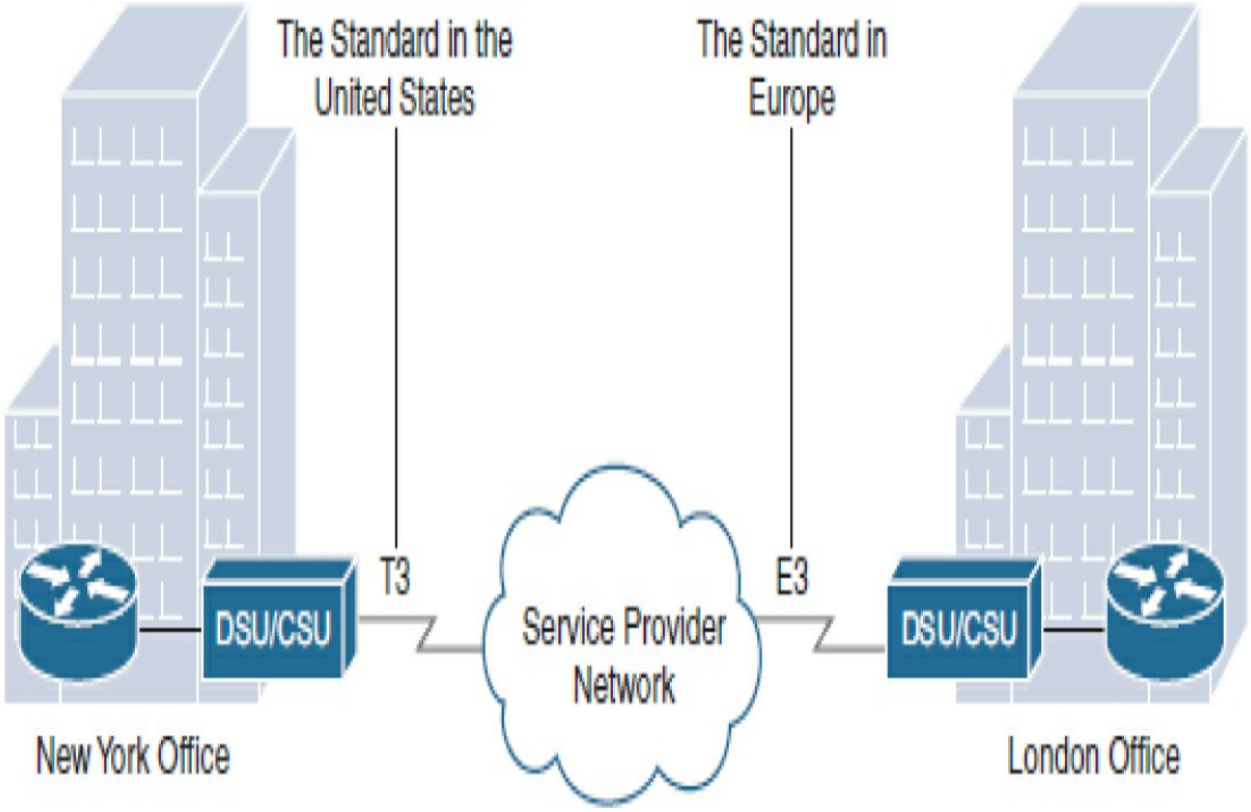

Figure 7-3 *Leased Lines*

Leased lines are usually more expensive than switched services because of the dedicated, always-on cost of providing WAN service to the customer. The dedicated capacity removes latency and jitter and provides a layer of security because only the customer's traffic is allowed on the link. Table 7-1 lists the available leased line types and their bit-rate capacities.

Table 7-1 *Leased Line Types and Capacities*

Line Type	Bit-Rate Capacity	Line Type	Bit-Rate Capacity
56k	56 kbps	OC-9	466.56 Mbps
64k	64 kbps	OC-12	622.08 Mbps
T1	1.544 Mbps	OC-18	933.12 Mbps
E1	2.048 Mbps	OC-24	1244.16 Mbps
J1	2.048 Mbps	OC-36	1866.24 Mbps
E3	34.064 Mbps	OC-48	2488.32 Mbps
T3	44.736 Mbps	OC-96	4976.64 Mbps
OC-1	51.84 Mbps	OC-192	9953.28 Mbps
OC-3	155.54 Mbps	OC-768	39,813.12 Mbps

Circuit-Switched Connection Options

The two main types of circuit-switched connections are analog dialup and ISDN. Both technologies have limited implementation bases in today's

Technet24

networks. However, they are both still used in remote rural areas and other areas of the globe where more recent technologies are not yet available.

Analog dialup uses modems at very low-speed connections that might be adequate for the exchange of sales figures, prices, routine reports, and email, or as an emergency backup link.

ISDN turns the local loop into a TDM digital connection, which enables it to carry digital signals that result in higher-capacity switched connections than are available with analog modems. Two types of ISDN interfaces exist:

• **Basic Rate Interface (BRI):** Provides two 64-kbps B-channels for voice or data transfer and a 16-kbps D-channel for control signaling.

• **Primary Rate Interface (PRI):** Provides 23 B-channels with 64 kbps and 1 D-channel with 64 kbps in North America, for a total bit rate of up to 1.544 Mbps. Europe uses 30 B-channels and 1 D-channel, for a total bit rate of up to 2.048 Mbps.

Figure 7-4 illustrates the various differences between ISDN BRI and PRI lines.

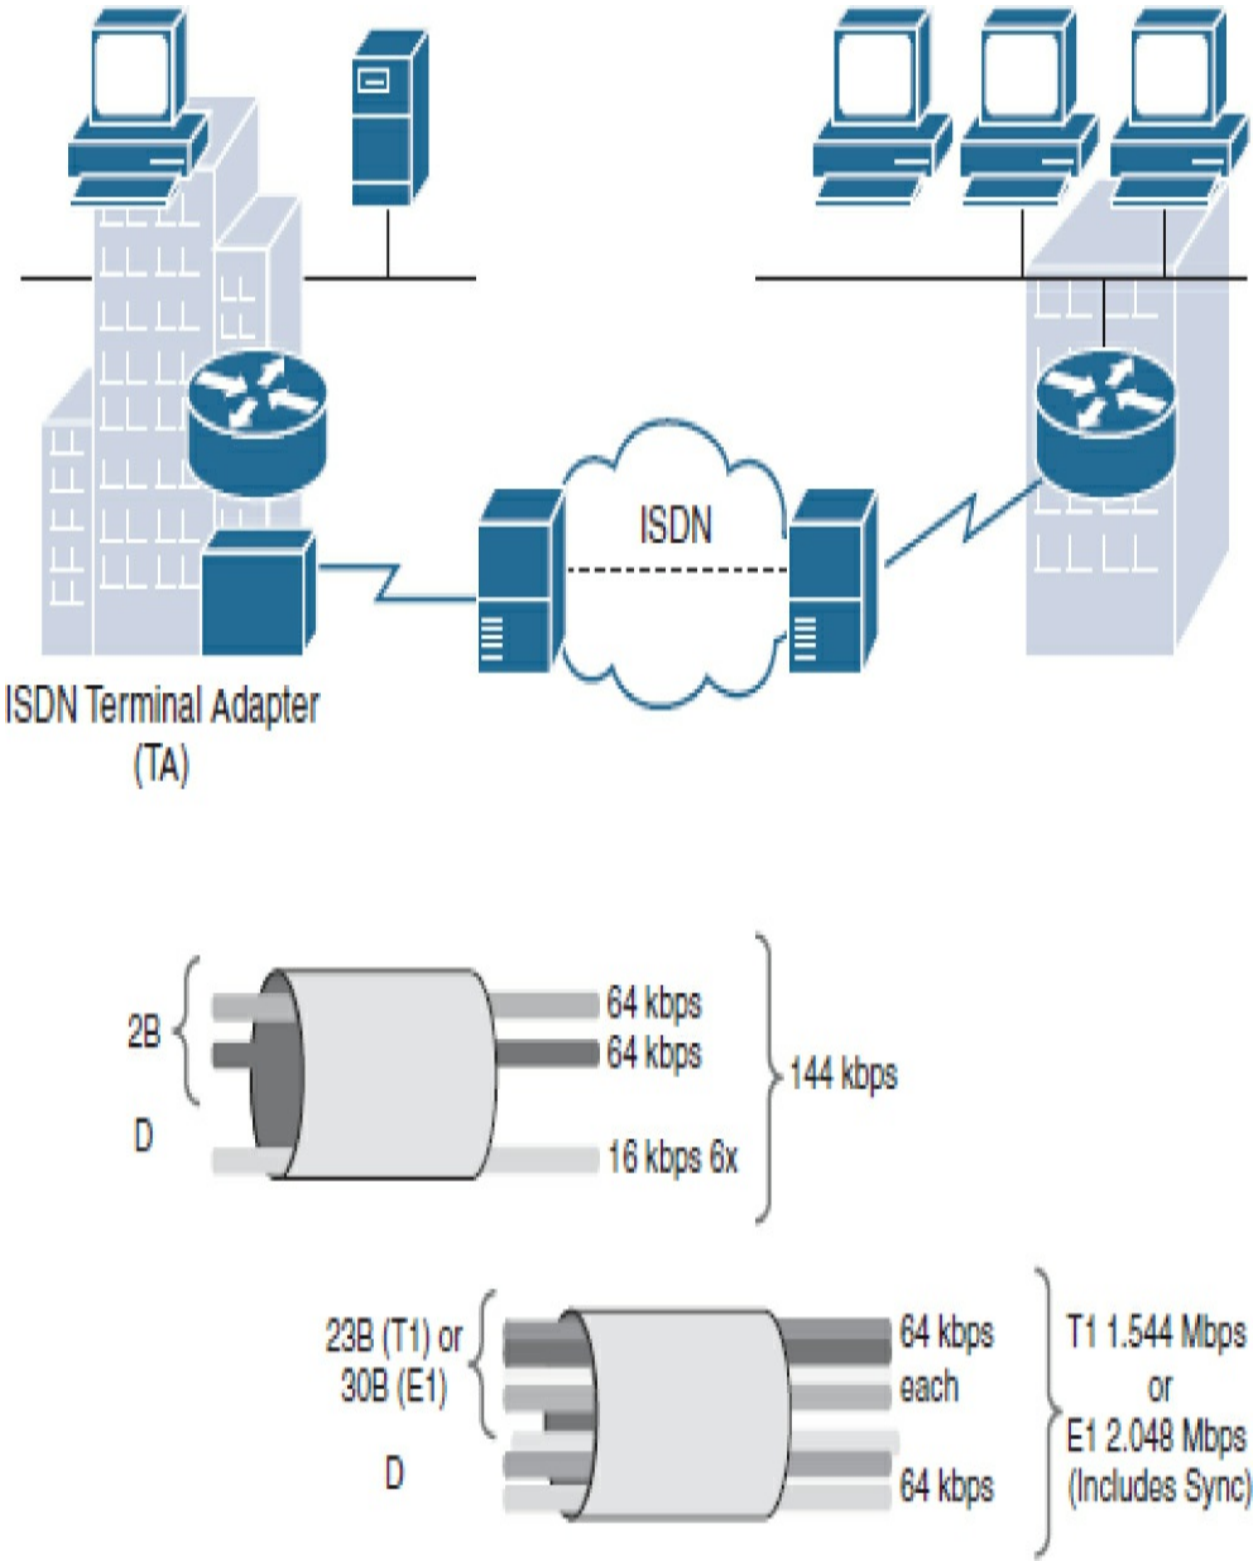

Figure 7-4 *ISDN Network Infrastructure and PRI/BRI Line Capacity*

Technet24

Packet-Switched Connection Options

The most common packet-switching technologies used in today's enterprise WANs include Metro Ethernet and MPLS. Legacy technologies include X.25 and ATM.

Note

Frame Relay is also a packet-switched technology that still has some market share. However, the CCNA exam no longer covers it.

Metro Ethernet

Metro Ethernet (MetroE) uses IP-aware Ethernet switches in the service provider's network cloud to offer enterprises converged voice, data, and video services at Ethernet speeds. Consider some benefits of Metro Ethernet:

• **Reduced expenses and administration:** Enables businesses to inexpensively connect numerous sites in a metropolitan area to each other and to the Internet without the need for expensive conversions to ATM or Frame Relay

• **Easy integration with existing networks:** Connects easily to existing Ethernet LANs

• **Enhanced business productivity:** Enables businesses to take advantage of productivity-enhancing IP applications that are difficult to implement on TDM or Frame Relay networks, such as hosted IP communications, VoIP, and streaming and broadcast video

MPLS

Multiprotocol Label Switching (MPLS) has the following characteristics:

• **Multiprotocol:** MPLS can carry any payload, including IPv4, IPv6, Ethernet, ATM, DSL, and Frame Relay traffic.

• **Labels:** MPLS uses labels inside the service provider's network to identify paths between distant routers instead of between endpoints.

• **Switching:** MPLS actually routes IPv4 and IPv6 packets, but everything else is switched.

As Figure 7-5 shows, MPLS supports a wide range of WAN technologies, including serial leased lines, Metro Ethernet, ATM, Frame Relay, and DSL (not shown).

Technet24

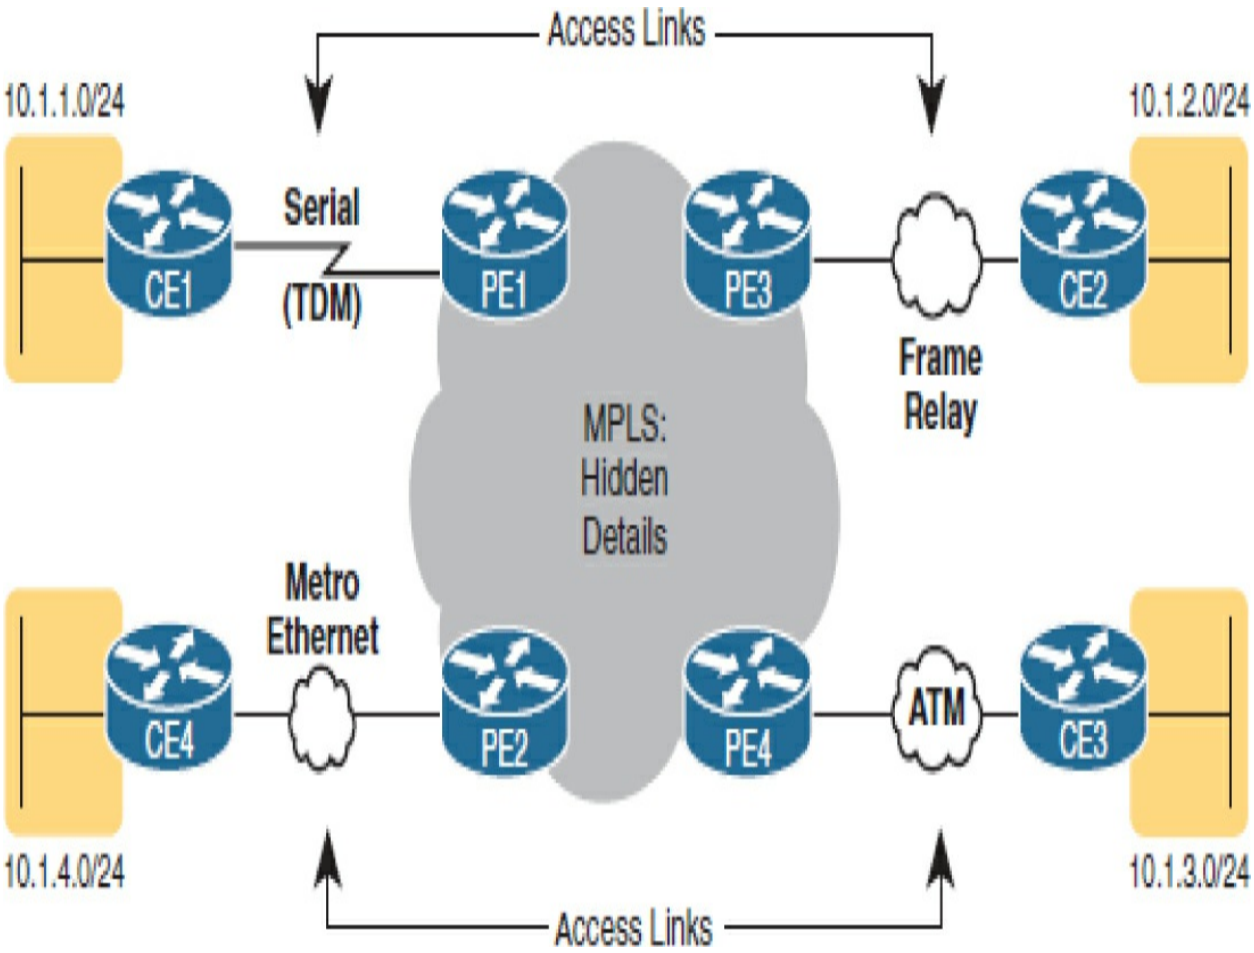

Figure 7-5 *Popular MPLS Connection Options*

In Figure 7-5, CE refers to the customer edge routers. PE is the provider edge routers that add and remove labels.

Note

MPLS is primarily a service provider WAN technology.

Internet Connection Options

Broadband connection options typically are used to connect telecommuting

employees to a corporate site over the Internet. These options include Digital Subscriber Line (DSL), cable, and wireless.

DSL

DSL technology, shown in Figure 7-6, is an always-on connection technology that uses existing twisted-pair telephone lines to transport high-bandwidth data and provides IP services to subscribers.

Technet24

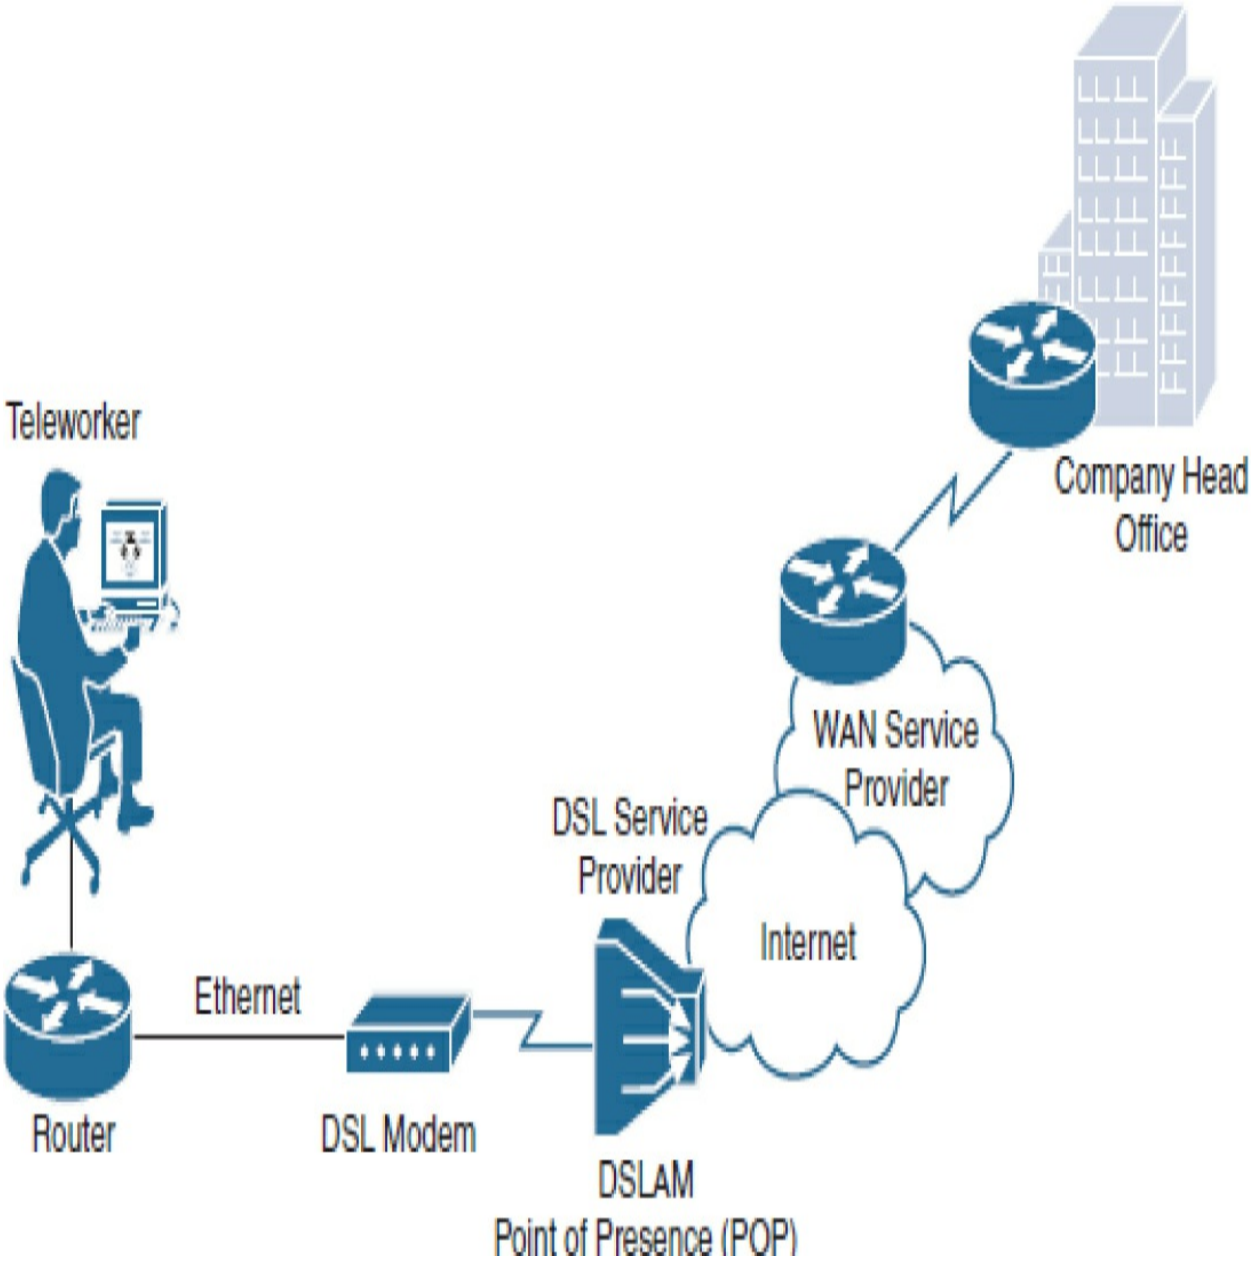

Figure 7-6 *Teleworker DSL Connection*

Current DSL technologies use sophisticated coding and modulation techniques to achieve data rates of up to 8.192 Mbps. A variety of DSL types, standards, and emerging technologies exist. DSL is a popular choice for enterprise IT departments to support home workers.

Cable Modem

A cable modem provides an always-on connection and simple installation. Figure 7-7 shows how a subscriber connects a computer or LAN router to the cable modem, which translates the digital signals into the broadband frequencies used for transmitting on a cable television network.

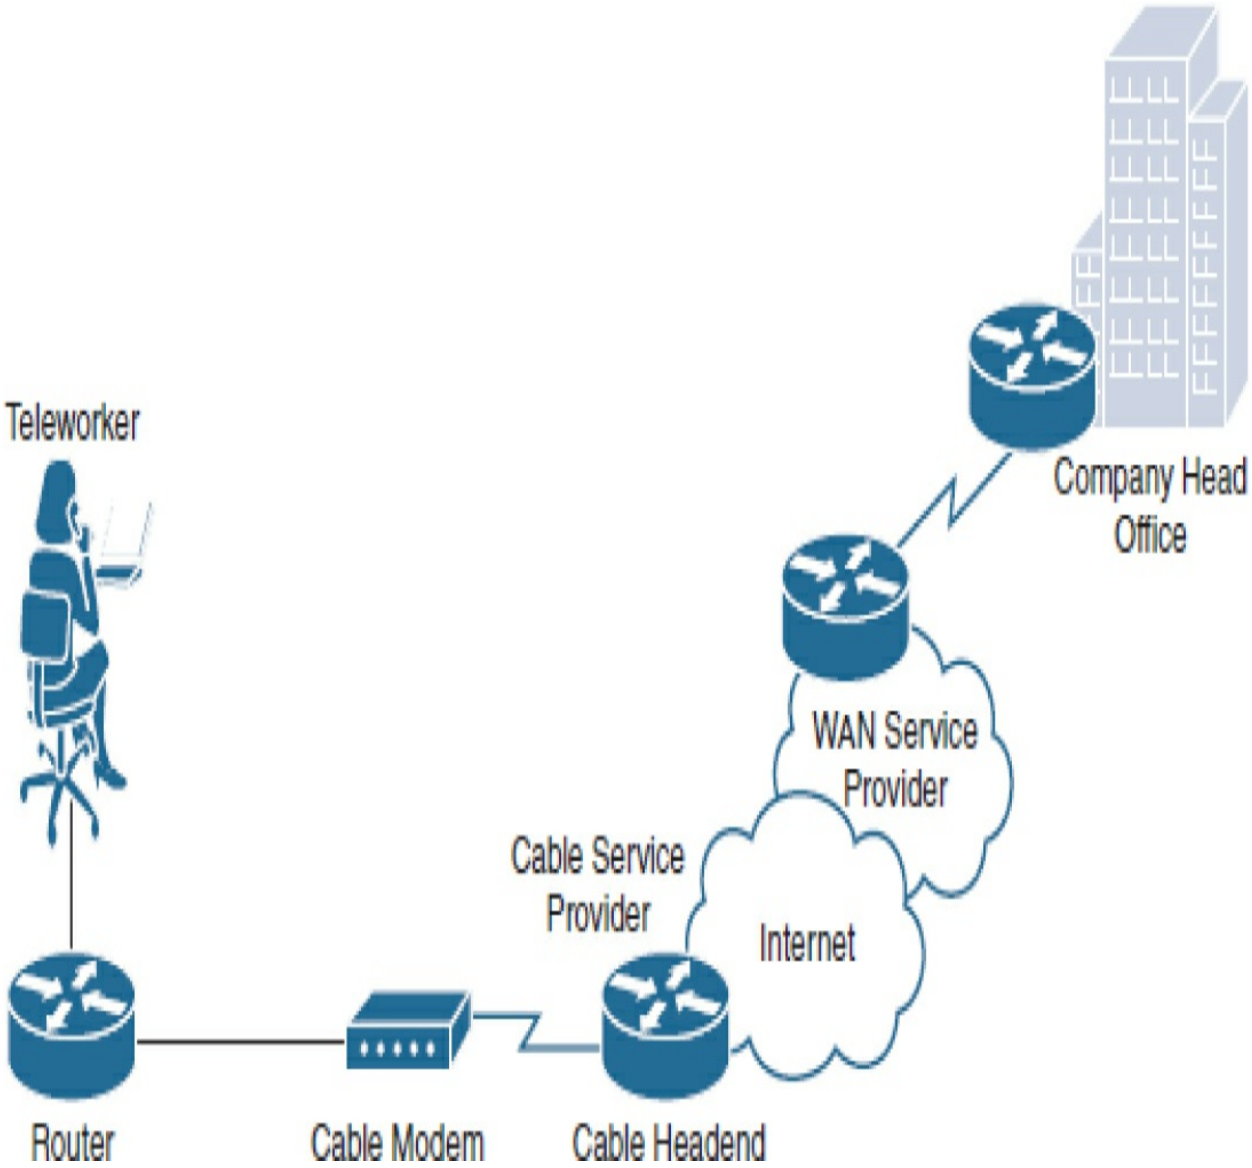

Figure 7-7 *Teleworker Cable Modem Connection*

Technet24

Wireless

In the past, the main limitation of wireless access was the need to be within range of a wireless router or a wireless modem with a wired connection to the Internet; however, the following wireless technologies enable users to connect to the Internet from almost any location:

• **Municipal Wi-Fi:** Many cities have begun setting up municipal wireless networks. Some of these networks provide high-speed Internet access for free or for substantially less than the price of other broadband services.

• **WiMAX:** Worldwide Interoperability for Microwave Access (WiMAX) is an IEEE 802.16 technology that is just beginning to come into use. It provides high-speed broadband service with wireless access and provides broad coverage similar to a cell phone network instead of through small Wi-Fi hotspots.

• **Satellite Internet:** This technology is typically used in rural areas where cable and DSL are unavailable.

• **Cellular service:** Cellular service is an option for connecting users and remote locations where no other WAN access technology is available. Common cellular access methods include 3G/4G (third generation and fourth generation) and Long-Term Evolution (LTE) cellular access.

Choosing a WAN Link Option

Table 7-2 compares the advantages and disadvantages of the various WAN connection options reviewed.

Table 7-2 *Choosing a WAN Link Connection*

Technet24

Option	Description	Advantages	Disadvantages	Sample Protocols
Leased line	Point-to-point connection between two LANs.	Most secure	Expensive	PPP, HDLC, SDLC
Circuit switching	Dedicated circuit path created between endpoints. The best example is dialup connections.	Inexpensive	Call setup	PPP, ISDN
Packet switching	Devices transporting packets via a shared single point-to-point or point-to-multipoint link across a carrier internetwork. Variable-length packets are transmitted over PVCs or SVCs.	Highly efficient use of bandwidth	Shared media across link	Frame Relay, MetroE
Internet	Connectionless packet switching using the Internet as the WAN infrastructure. Uses network addressing to deliver packets. Because of security issues, VPN technology must be used.	Least expensive, globally available	Least secure	DSL, cable modem, wireless

VPN Technology

A virtual private network (VPN) is an encrypted connection between private networks over a public network such as the Internet. Instead of using a dedicated Layer 2 connection such as a leased line, a VPN uses virtual connections called VPN tunnels, which are routed through the Internet from the company's private network to the remote site or employee host.

VPN Benefits

Benefits of VPN include the following:

• **Cost savings:** Eliminates the need for expensive dedicated WAN links and modem banks

• **Security:** Uses advanced encryption and authentication protocols that protect data from unauthorized access

• **Scalability:** Can add large amounts of capacity without adding significant infrastructure

• **Compatibility with broadband technology:** Supported by broadband service providers, so mobile workers and telecommuters can take advantage of their home high-speed Internet service to access their corporate networks

Types of VPN Access

Technet24

The following describes the types of VPN access methods:

• **Site-to-site VPNs:** Site-to-site VPNs connect entire networks to each other. For example, a site-to-site VPN can connect a branch office network to a company headquarters network, as in Figure 7-8. Each site is equipped with a VPN gateway, such as a router, firewall, VPN concentrator, or security appliance. In the figure, a remote branch office uses a site-to-site VPN to connect with the corporate head office.

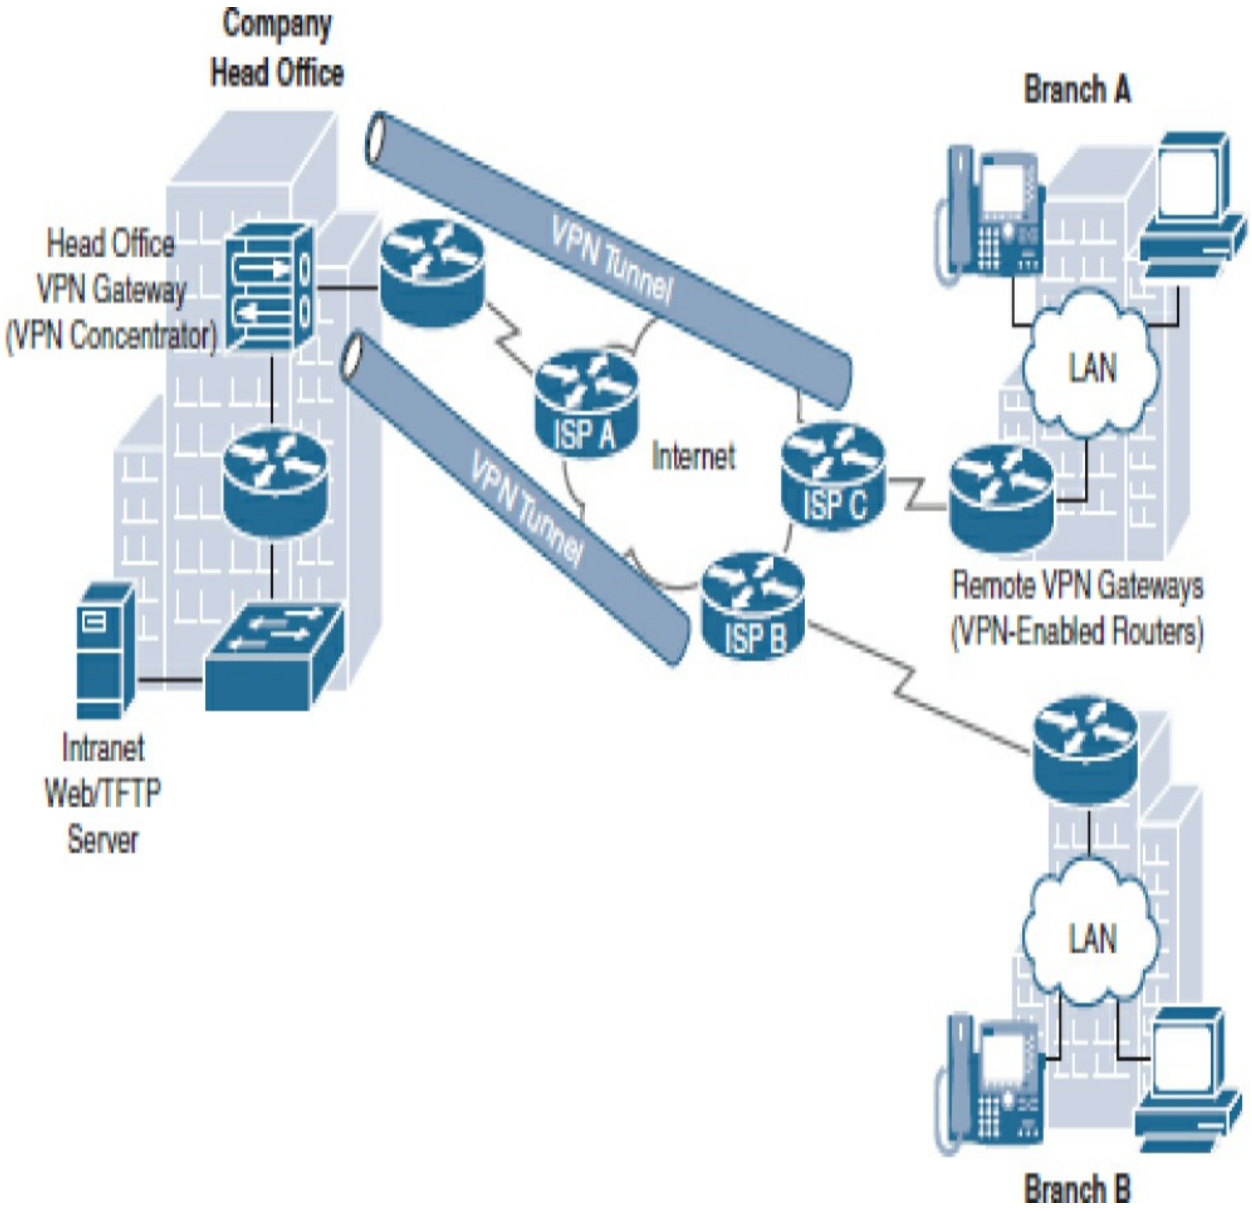

Figure 7-8 *Site-to-Site VPNs*

• **Remote-access VPNs:** Remote-access VPNs enable individual hosts, such as telecommuters, mobile users, and extranet consumers, to access a company network securely over the Internet, as in Figure 7-9. Each host typically has client software for a client-based VPN connection or uses a web browser for clientless VPN connection. Web-based clientless VPNs are also typically

Technet24

called *clientless Secure Sockets Layer (SSL) connections*. However, the VPN is actually established using Transport Layer Security (TLS). TLS is the newer version of SSL and is sometimes expressed as SSL/TLS.

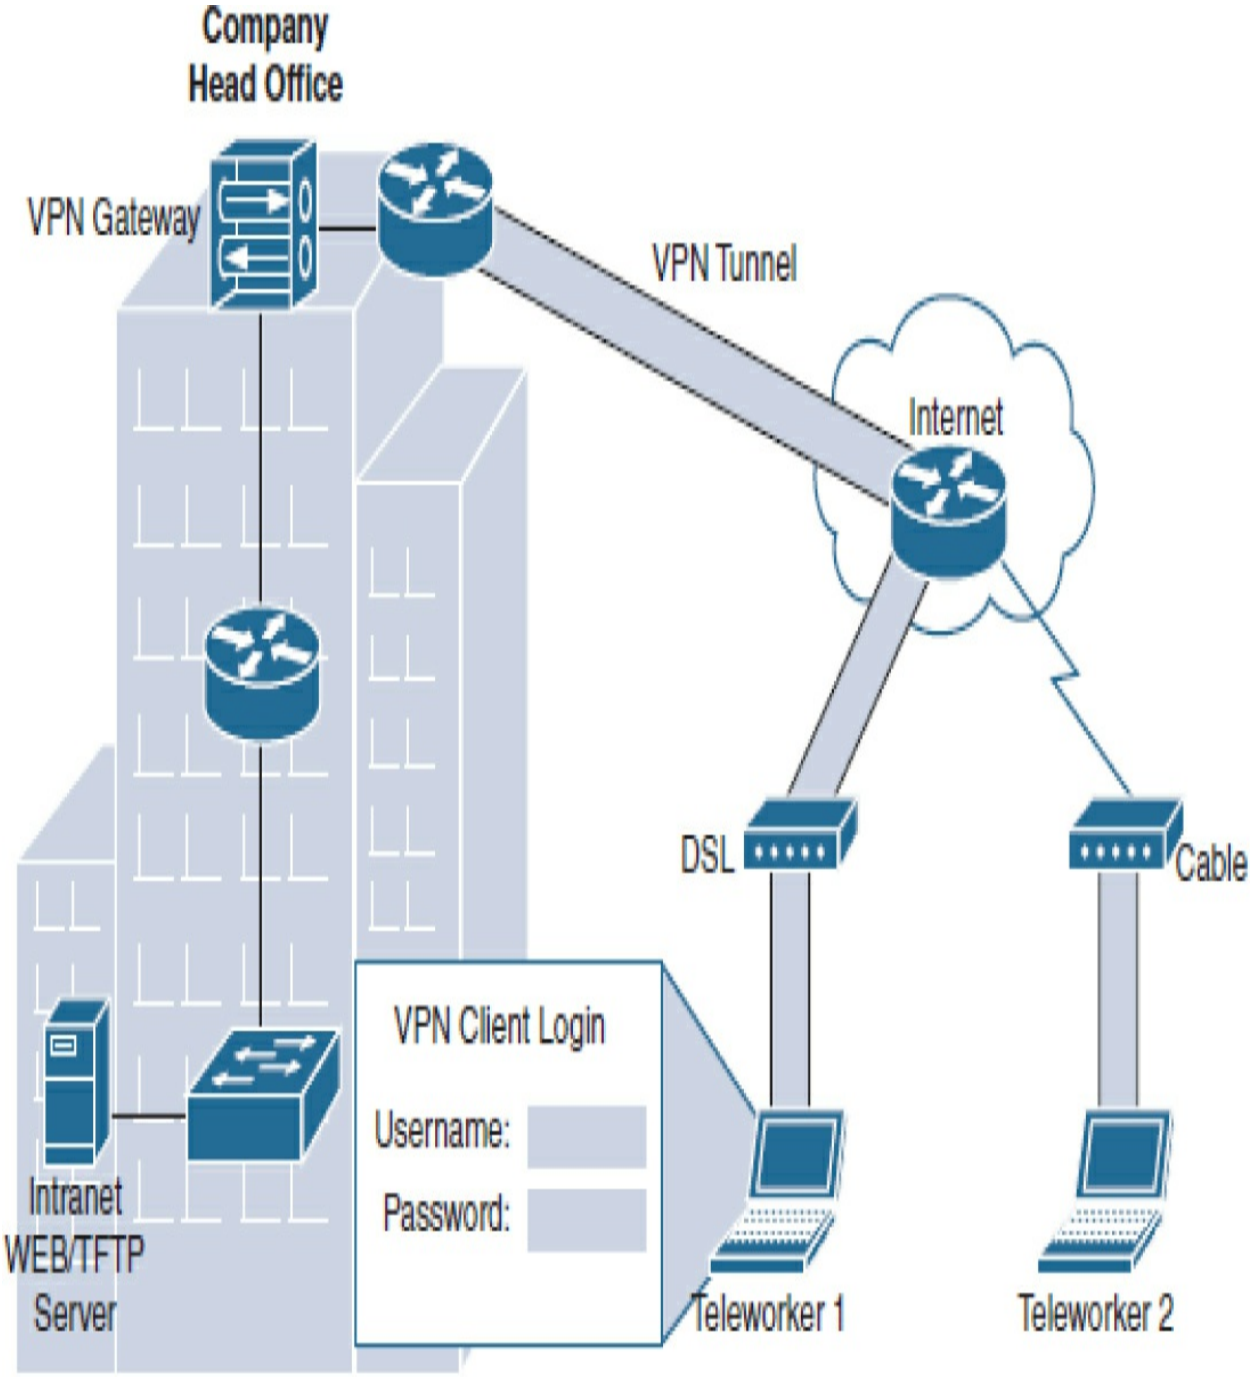

Figure 7-9 *Remote-Access VPNs*

• **Generic Routing Encapsulation (GRE):** A standard IPsec VPN (non-GRE) can only create secure tunnels for unicast traffic. GRE is a nonsecure site-to-site VPN tunneling protocol that can support multicast and broadcast traffic needed for network layer protocols. However, GRE does not by default support encryption; therefore, it does not provide a secure VPN tunnel. To solve this problem, you can encapsulate routing protocol traffic by using a GRE packet and then encapsulate the GRE packet into an IPsec packet to forward it securely to the destination VPN gateway. The terms used to describe the encapsulation of GRE over IPsec tunnel are *passenger protocol* for the routing protocol, *carrier protocol* for GRE, and *transport protocol* for IPsec, as shown in Figure 7-10.

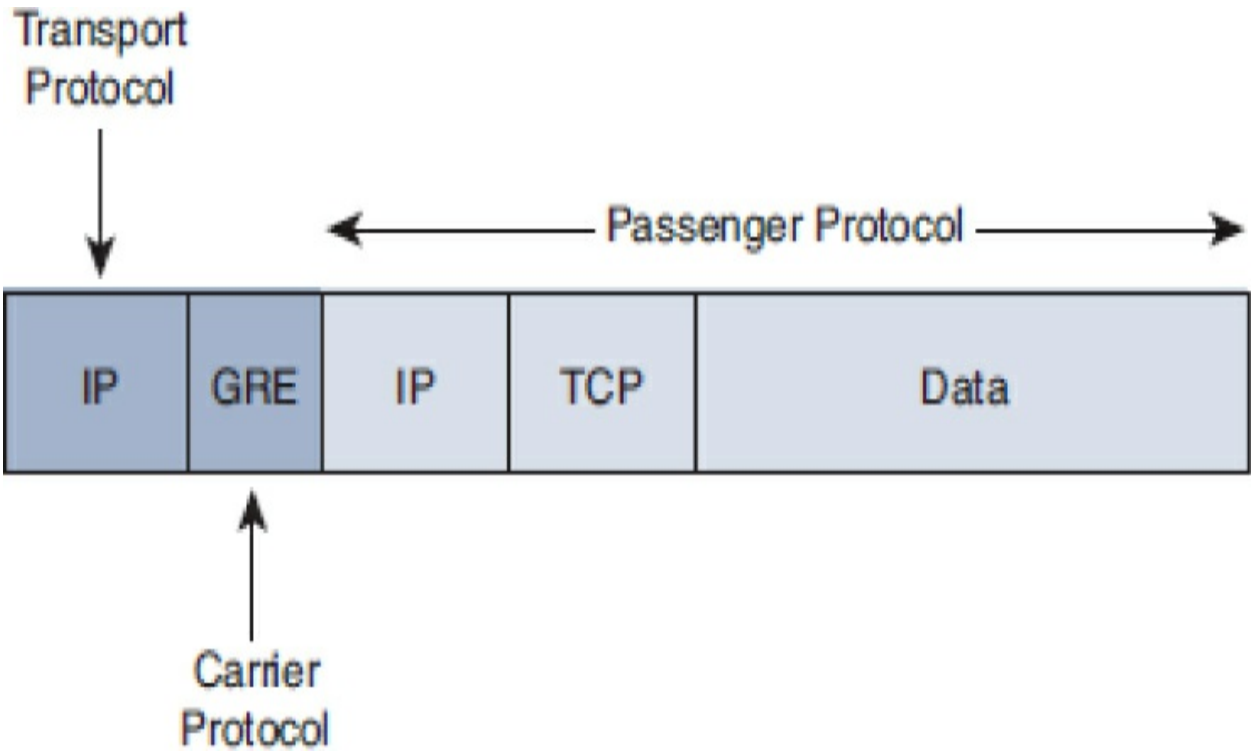

Figure 7-10 *Transport, Carrier, and Passenger Protocols*

Technet24

• **Dynamic Multipoint VPN (DMVPN):** DMVPN is a Cisco-proprietary solution for building many VPNs in an easy, dynamic, and scalable manner. DMVPN allows a network administrator to dynamically form hub-to-spoke tunnels and spoke-to-spoke tunnels, as in Figure 7-11. DMVPN simplifies the VPN tunnel configuration and provides a flexible option for connecting a central site with branch sites. It uses a hub-and-spoke configuration to establish a full mesh topology. Spoke sites establish secure VPN tunnels with the hub site. Each site is configured using Multipoint Generic Routing Encapsulation (mGRE). The mGRE tunnel interface allows a single GRE interface to dynamically support multiple IPsec tunnels.

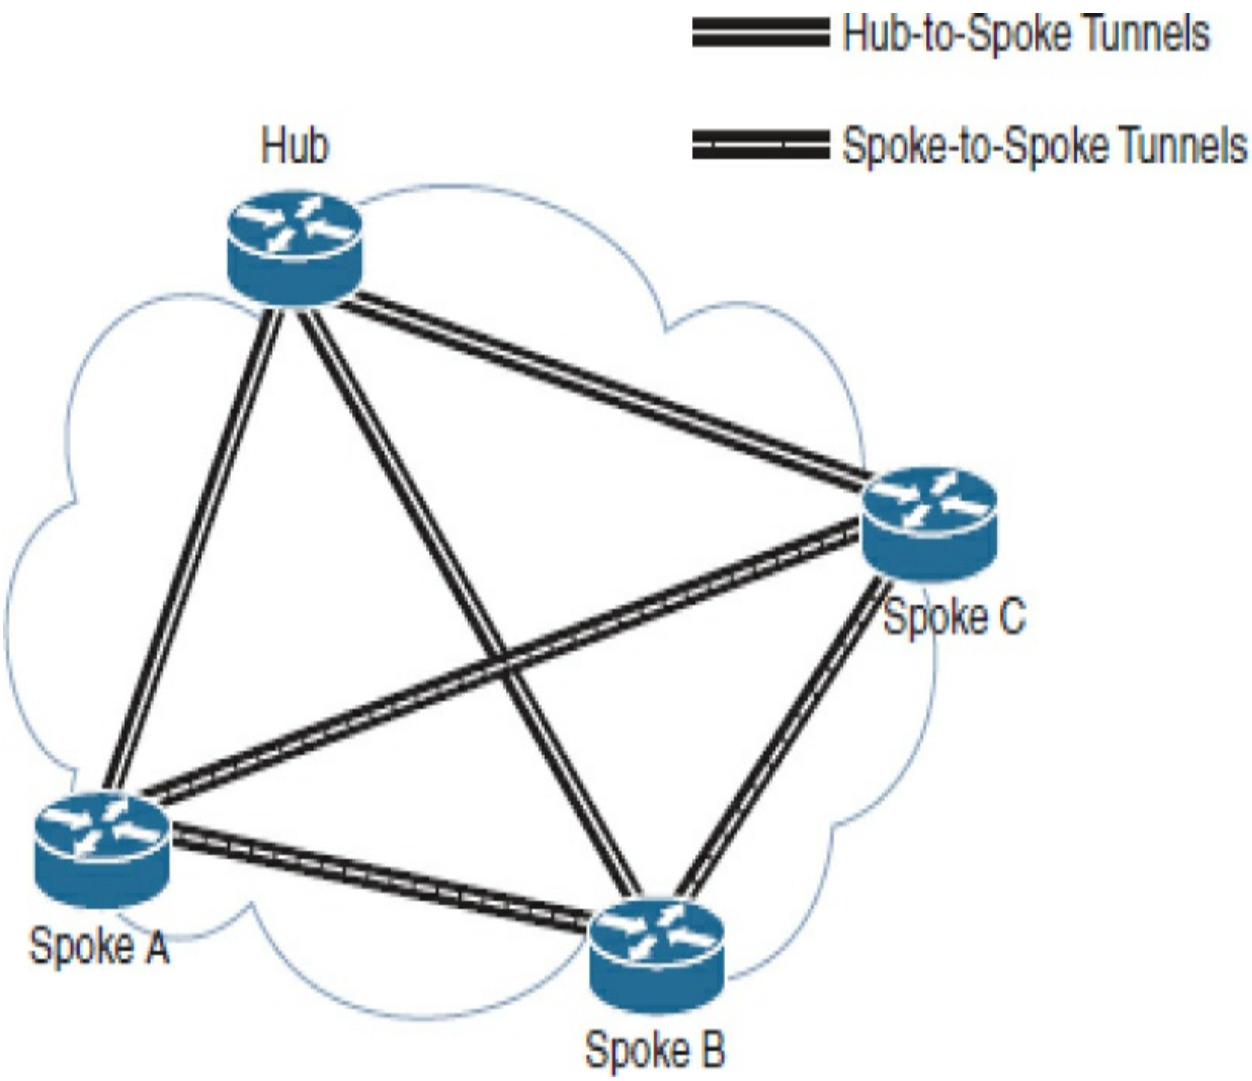

Figure 7-11 *DMVPN Sample Topology*

DMVPN uses the following technologies:

• **Next Hop Resolution Protocol (NHRP):** Maps public IP addresses for all tunnel spokes

• **IPsec encryption:** Provides the security to transport private information over public networks

Technet24

• **mGRE:** Allows a single interface to support multiple IPsec tunnels

• **IPsec Virtual Tunnel Interface (VTI):** Like DMVPN, VTI simplifies the configuration process required to support multiple sites and remote access. IPsec VTI is capable of sending and receiving both IP unicast and multicast encrypted traffic. Therefore, routing protocols are automatically supported without the need to configure GRE tunnels.

• **Service provider MPLS VPNs:** MPLS can provide clients with managed VPN solutions; therefore, securing traffic between client sites is the responsibility of the service provider. Two types of MPLS VPN solutions are supported by service providers:

• **Layer 3 MPLS VPN:** The service provider participates in customer routing, redistributing the routes through the MPLS network to the customer's remote locations.

• **Layer 2 MPLS VPN:** The service provider is not involved in the customer routing. Instead, the provider deploys Virtual Private LAN Service (VPLS) to emulate an Ethernet multiaccess LAN segment over the MPLS network. No routing is involved. The customer's routers effectively belong to the same multiaccess network.

VPN Components

Figure 7-12 illustrates a typical VPN topology. Components required to

establish this VPN include the following:

• An existing enterprise network with servers and workstations

• A connection to the Internet

• VPN gateways, such as routers, firewalls, VPN concentrators, and Adaptive Security Appliances (ASAs), that act as endpoints to establish, manage, and control VPN connections

• Appropriate software to create and manage VPN tunnels

Technet24

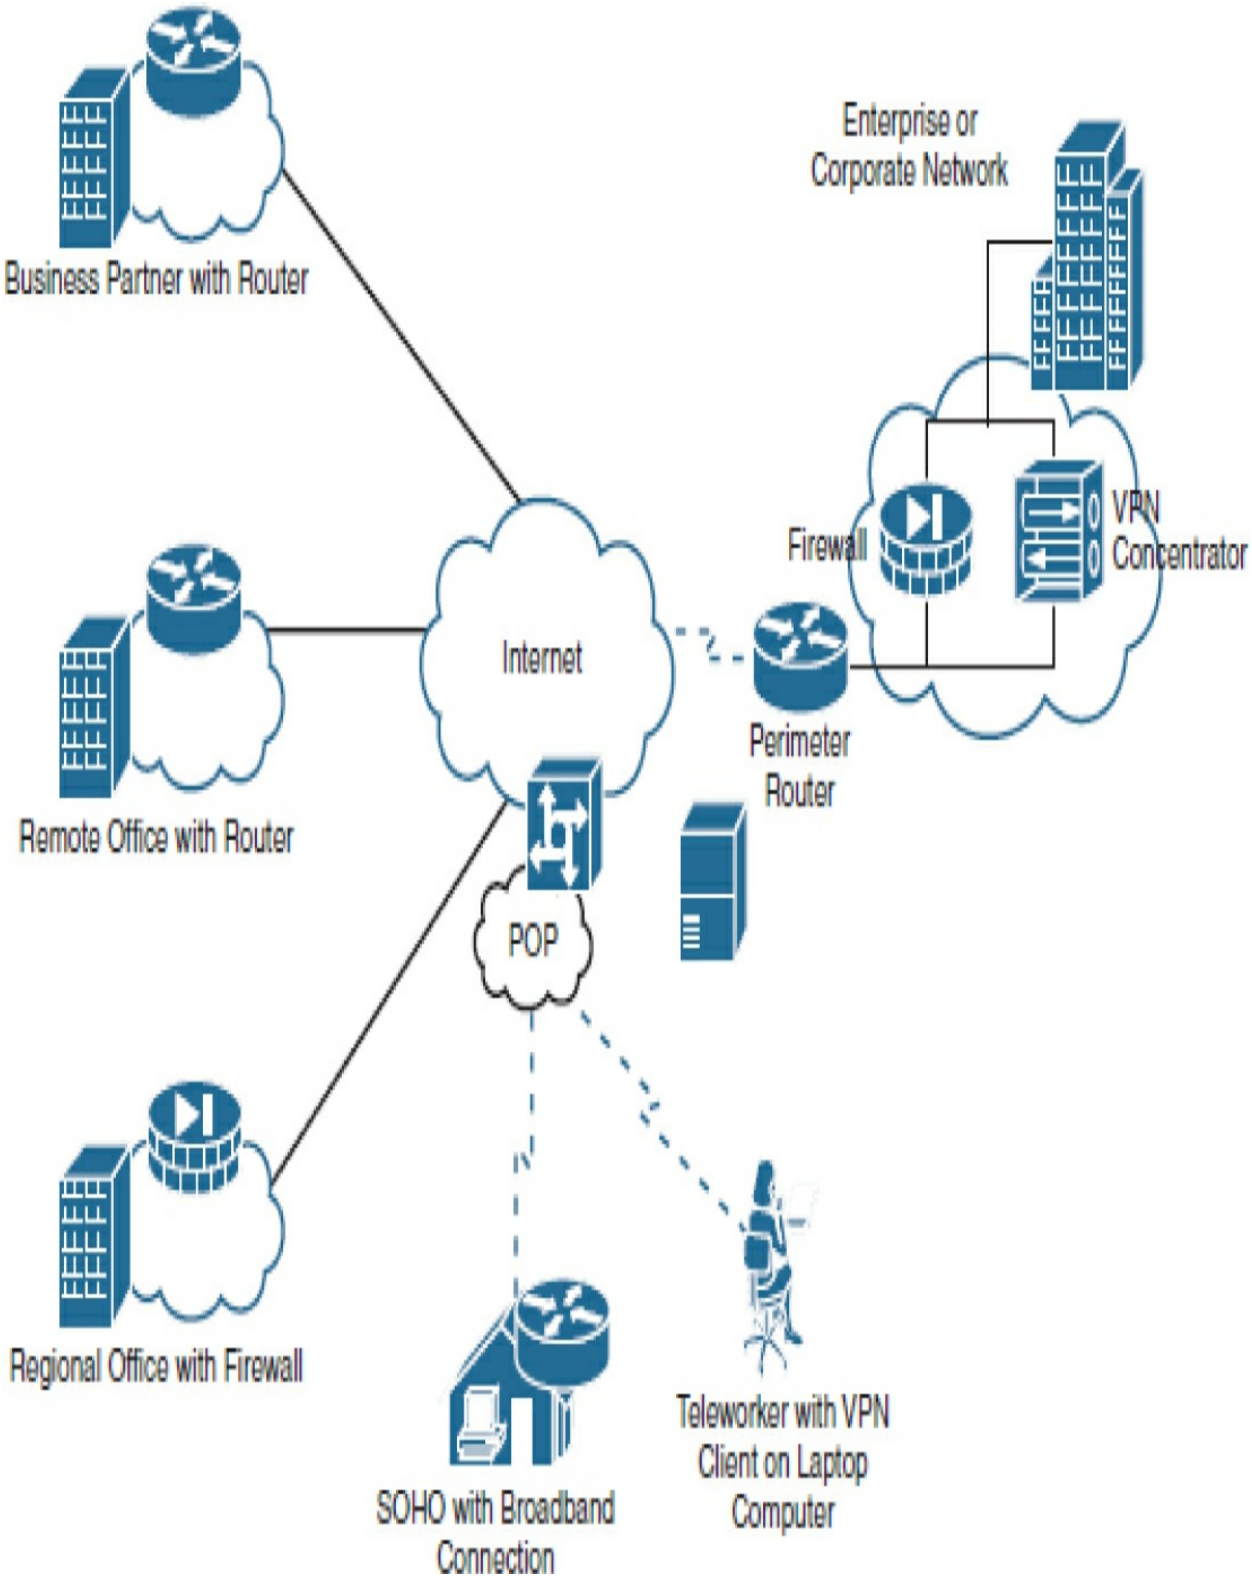

Figure 7-12 *VPN Components*

Establishing Secure VPN Connections

VPNs secure data by encapsulating and encrypting it. With regard to VPNs, encapsulation and encryption are defined as follows:

• Encapsulation is also called *tunneling* because encapsulation transmits data transparently from source network to destination network through a shared network infrastructure.

• Encryption codes data into a different format by using a secret key, which is then used on the other side of the connection for decryption.

VPN Tunneling

Tunneling uses three classes of protocols:

• **Carrier protocol:** The protocol over which information travels, such as Frame Relay, PPP, or MPLS

• **Encapsulating protocol:** The protocol that is wrapped around the original data, such as GRE, IPsec, L2F, PPTP, or L2TP

• **Passenger protocol:** The protocol over which the original data was carried, such as IPX, AppleTalk, IPv4, or IPv6

Figure 7-13 illustrates an email message traveling through the Internet over a VPN connection.

Technet24

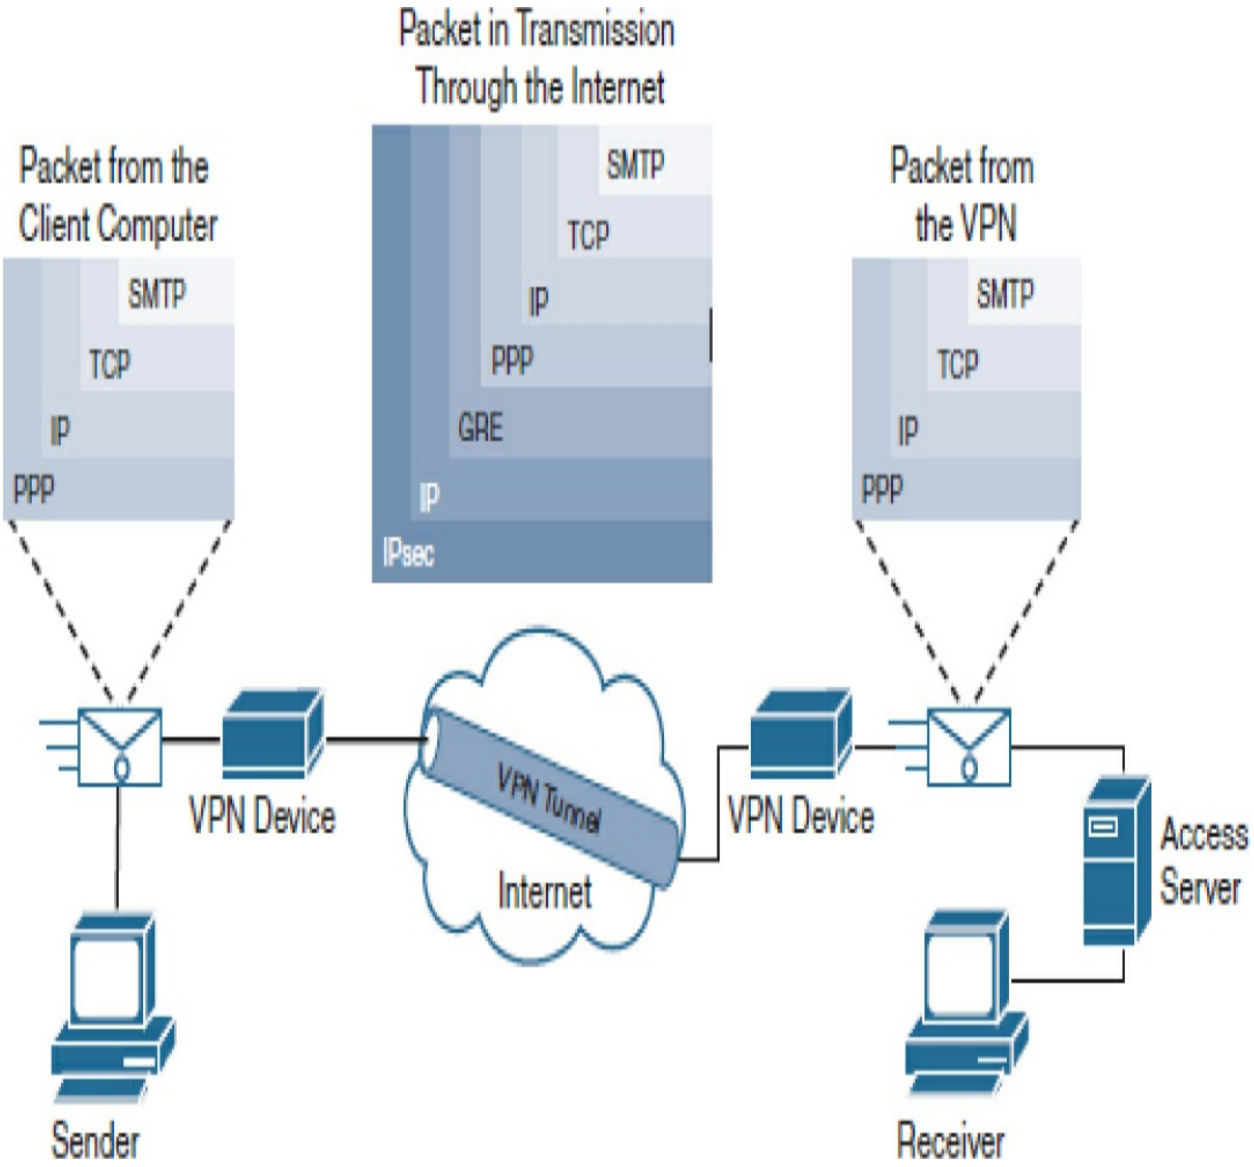

Figure 7-13 *Packet Encapsulation in a VPN Tunnel*

VPN Encryption Algorithms

The degree of security provided by any encryption algorithm depends on the key's length. Some of the most common encryption algorithms and the lengths of the keys they use are as follows:

• **Data Encryption Standard (DES) algorithm:** Uses a 56-bit key and

ensures high-performance encryption. DES is a symmetric key cryptosystem.

• **Triple DES (3DES) algorithm:** A newer variant of DES that encrypts with one key, decrypts with a different key, and then encrypts a final time with another key.

• **Advanced Encryption Standard (AES):** Provides stronger security than DES and is computationally more efficient than 3DES. AES offers three key lengths: 128-, 192-, and 256-bit keys.

• **Rivest, Shamir, and Adleman (RSA):** An asymmetric key cryptosystem. The keys use a bit length of 512, 768, 1024, or larger.

With symmetric encryption, the encryption key and decryption key are the same. With asymmetric encryption, they are different.

Hashes

VPNs use a keyed hashed message authentication code (HMAC) data-integrity algorithm to guarantee a message's integrity and authenticity without any additional mechanisms.

The cryptographic strength of the HMAC depends on the cryptographic strength of the underlying hash function, the key's size and quality, and the size of the hash output length, in bits. There are two common HMAC algorithms:

Technet24

• **Message Digest 5 (MD5):** Uses a 128-bit shared secret key

• **Secure Hash Algorithm 1 (SHA-1):** Uses a 160-bit secret key

Figure 7-14 shows an example using MD5 as the HMAC algorithm.

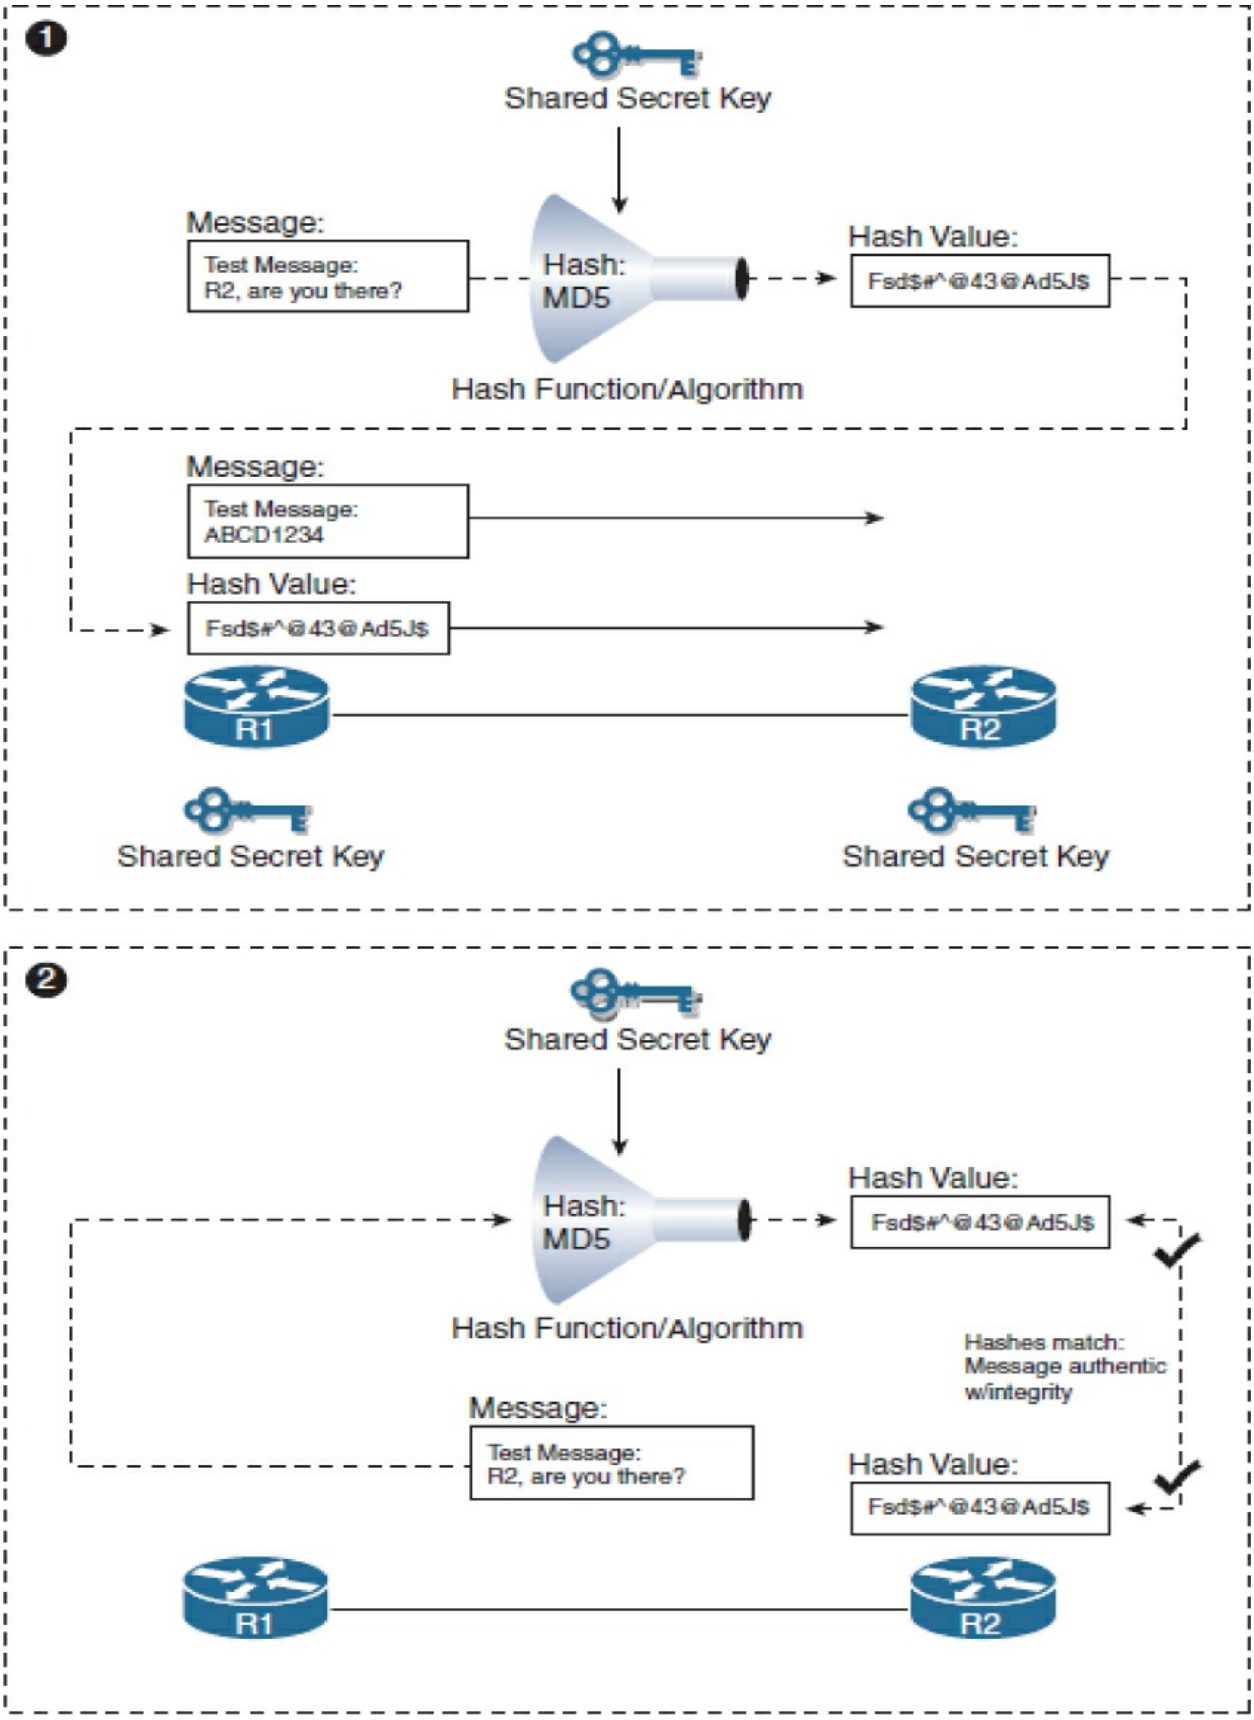
1
Shared Secret Key
Message:
Test Message:
R2, are you there?
Hash:
MD5
Hash Value:
Fsd$#^@43@Ad5J$
Hash Function/Algorithm
Message:
Test Message:
ABCD1234
Hash Value:
Fsd$#^@43@Ad5J$
R1
R2
Shared Secret Key
Shared Secret Key
2
Shared Secret Key
Hash:
MD5
Hash Value:
Fsd$#^@43@Ad5J$
Hash Function/Algorithm
Hashes match:
Message authentic
w/integrity
Message:
Test Message:
R2, are you there?
Hash Value:
Fsd$#^@43@Ad5J$
R1
R2

Technet24

Figure 7-14 *Creating and Verifying a Message Digest*

An HMAC has two parameters: a message input and a shared secret key known only to the message originator and intended recipients. In Figure 7-14, both R1 and R2 know the shared secret key. The process in Figure 7-14 uses the following steps:

R1 uses MD5 to perform the hashing function, which outputs a hash value. This hash value is then appended to the original message and sent to R2.

R2 removes the hash value from the original message, runs the same hash operation, and then compares its hash value with the hash value sent by R1. If the two hashes match, data integrity has not been compromised.

VPN Authentication

The device on the other end of the VPN tunnel must be authenticated before the communication path is considered secure. The two peer authentication methods are as follows:

• **Pre-Shared Key (PSK):** A secret key is shared between the two parties using a secure channel before it needs to be used.

• **RSA signature:** This method uses the exchange of digital certificates to authenticate the peers.

IPsec Security Protocols

Both IPsec and SSL VPN technologies offer access to virtually any network application or resource. However, when security is an issue, IPsec is the superior choice. Table 7-3 compares IPsec and SSL remote access deployments.

Table 7-3 *IPsec and SSL for Remote Access*

Technet24

Feature	IPsec	SSL
Applications supported	Extensive—All IP-based applications are supported.	Limited—Only web-based applications and file sharing are supported.
Authentication strength	Strong—Uses two-way authentication with shared keys or digital certificates.	Moderate—Using one-way or two-way authentication.
Encryption strength	Strong—Uses key lengths from 56 bits to 256 bits.	Moderate to strong—With key lengths from 40 bits to 256 bits.
Connection complexity	Medium—Requires that a VPN client be pre-installed on a host.	Low—Requires a web browser only on a host.
Connection option	Limited—Only specific devices with specific configurations can connect.	Extensive—Any device with a web browser can connect.

IPsec spells out the messaging necessary to secure VPN communications but relies on existing algorithms. The two main IPsec framework protocols are as

follows:

• **Authentication Header (AH):** Used when confidentiality is not required or permitted. AH provides data authentication and integrity for IP packets passed between two systems. It verifies the originators of any messages and that any message passed has not been modified during transit. AH does not provide data confidentiality (encryption) of packets. Used alone, the AH protocol provides weak protection. Consequently, it is used with the ESP protocol to provide data encryption and tamper-aware security features.

• **Encapsulating Security Payload (ESP):** Provides confidentiality and authentication by encrypting the IP packet. Although both encryption and authentication are optional in ESP, at a minimum, one of them must be selected.

IPsec relies on existing algorithms to implement encryption, authentication, and key exchange. Figure 7-15 shows how IPsec is structured.

IPsec provides the framework, and the administrator chooses the algorithms used to implement the security services within that framework. As Figure 7-15 illustrates, the administrator must fill the four IPsec framework squares:

• Choose an IPsec protocol.

• Choose the encryption algorithm that is appropriate for the desired level of security.

Technet24

• Choose an authentication algorithm to provide data integrity.

• The last square is the Diffie-Hellman (DH) algorithm group, which establishes the sharing of key information between peers. Choose which group to use: DH1, DH2, or DH5.

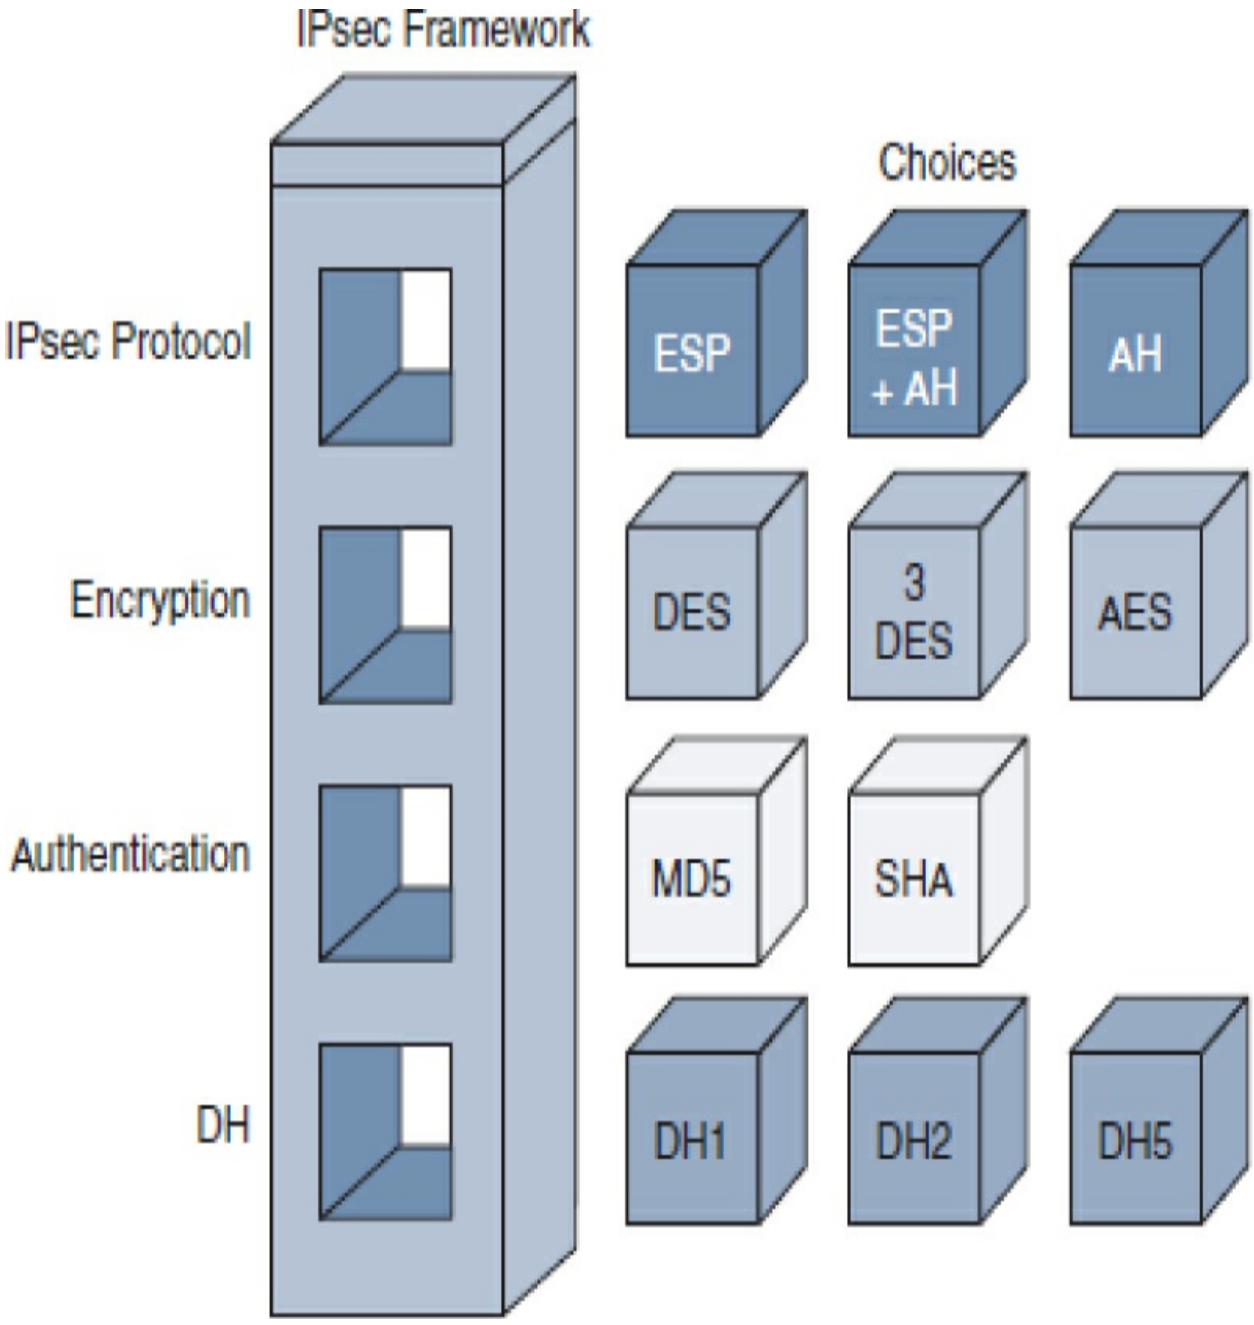

Figure 7-15 *IPsec Framework*

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	7–8
CCNA 200-301 Official Cert Guide, Volume 1, 2nd edition	3
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	19

Technet24

Day 6. QoS

CCNA 200-301 Exam Topic

• 4.7 Explain the forwarding per-hop behavior (PHB) for QoS such as classification, marking, queuing, congestion, policing, and shaping

Key Topics

Today, we review quality of service (QoS), which refers to the tools and techniques network administrators can use to prioritize traffic on a network.

QoS

Normal default operation for switches and routers is to process frames and packets in the order in which they are received. This first-in, first-out (FIFO) queueing mechanism does not discriminate between traffic types.

QoS tools are used to classify traffic types based on the following four characteristics:

• **Latency (delay):** Latency, or delay, is the amount of time it takes for data to be sent to the receiver. QoS tools can reduce the delay for time-sensitive packets, such as voice and video.

• **Jitter:** Jitter is the variance in the delay of packets. QoS tools can even out the delay of packets to improve end-user experience.

• **Loss:** Loss refers to the number of lost messages, usually as a percentage of the packets sent. QoS tools reduce packet loss, especially for time-sensitive traffic.

• **Bandwidth:** Bandwidth is a measure of the amount of data an interface can send every second. QoS tools can manage which traffic type gets to use the bandwidth next and how much of the bandwidth each type of traffic gets over time.

Figure 6-1 list the three major traffic types. For voice and video, the minimum traffic characteristic values are shown.

Technet24

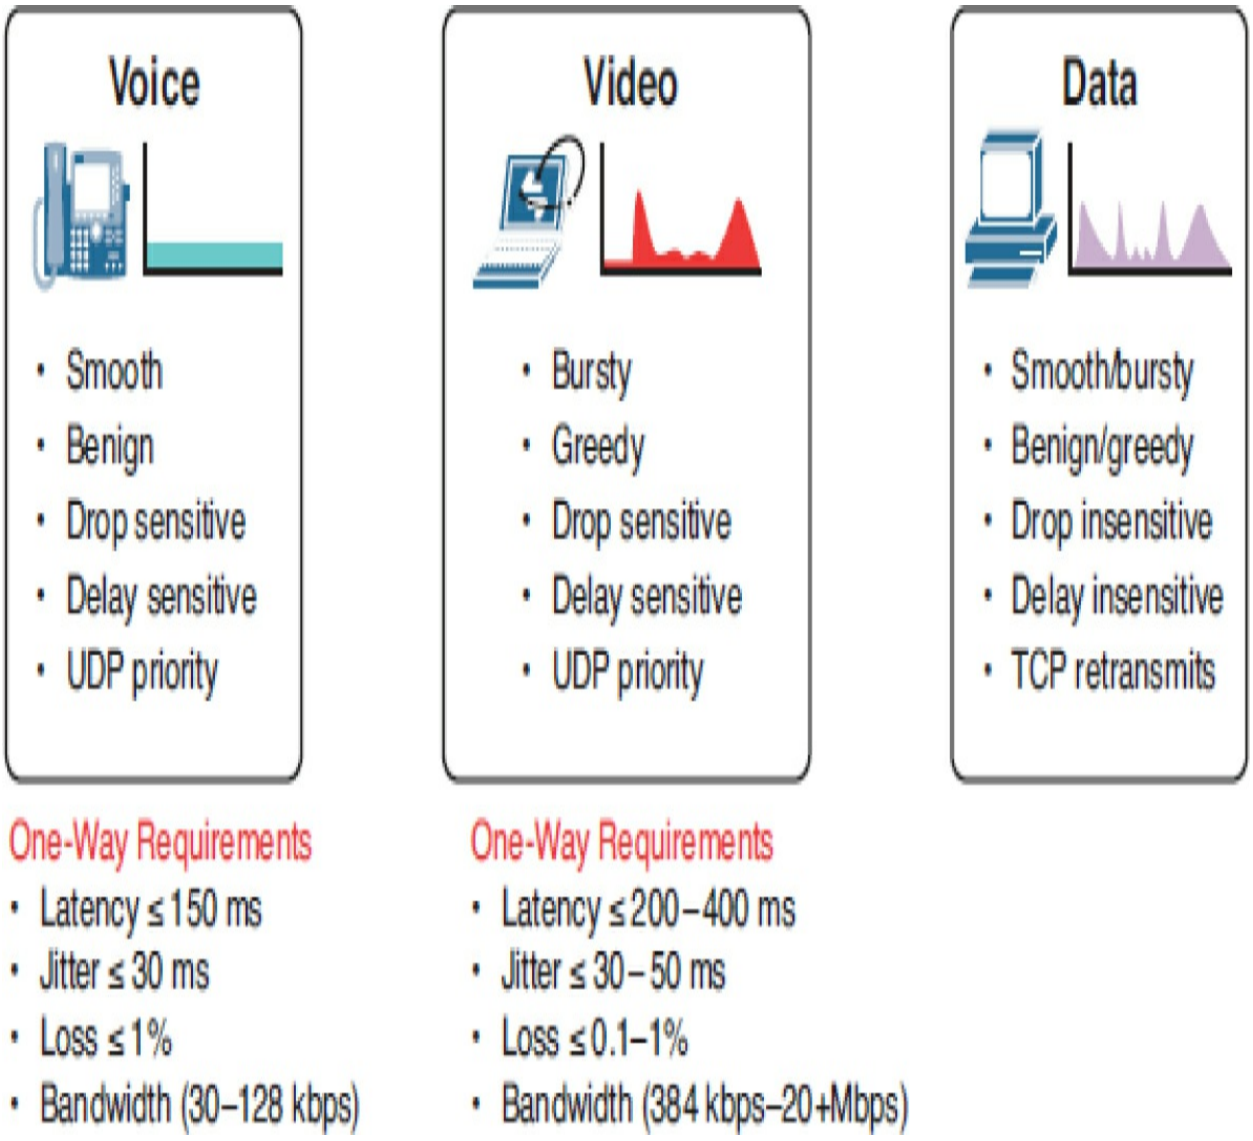

Figure 6-1 *Characteristics of Major Traffic Types*

Figure 6-2 shows the sequence of QoS events as traffic is forwarded out an interface.

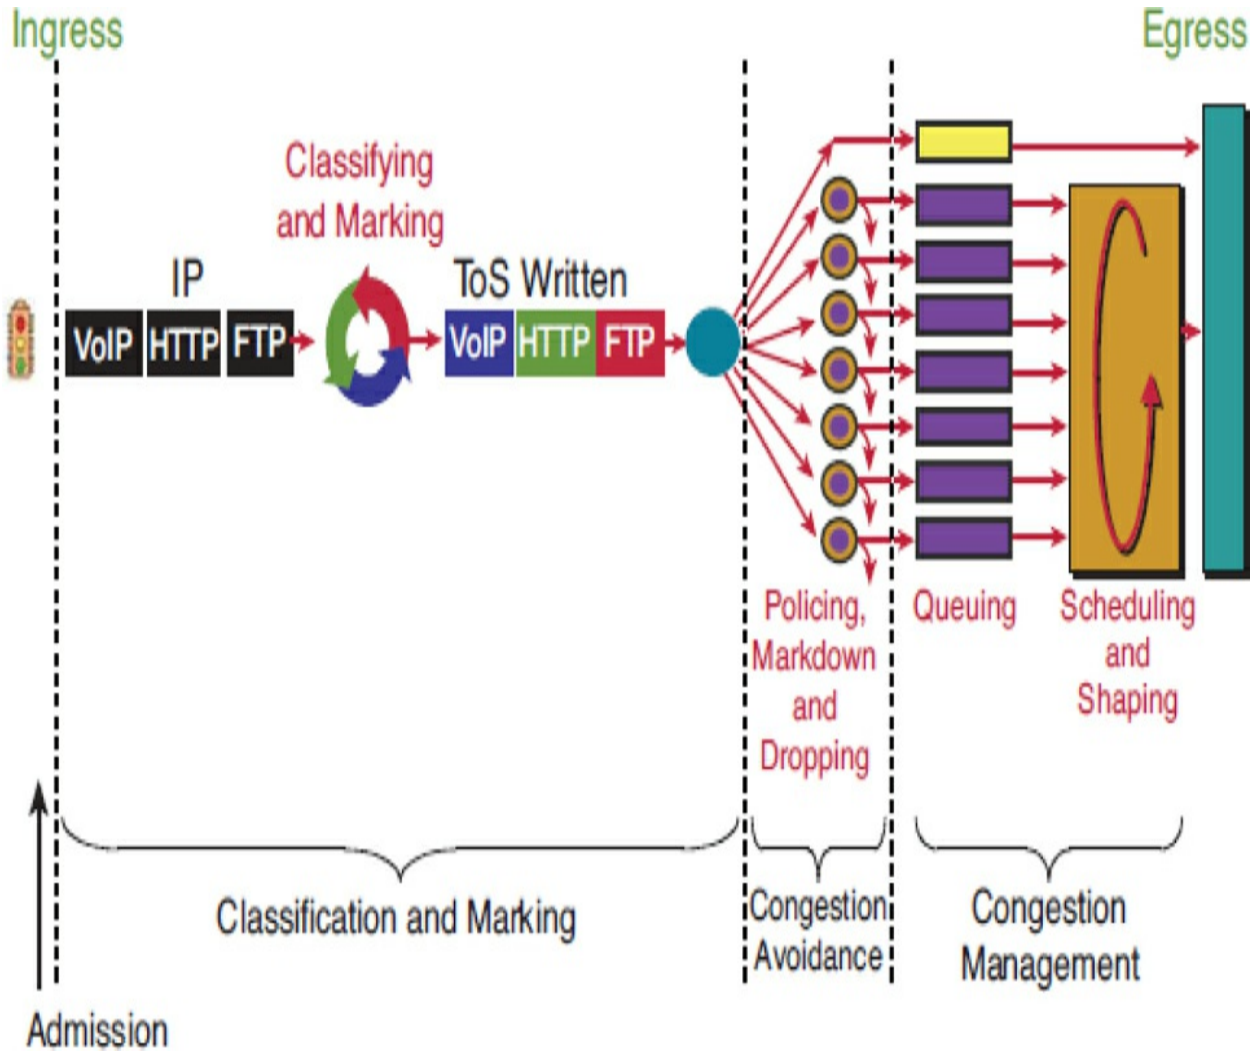

Figure 6-2 *Overview of QoS Tools*

QoS tools shown in Figure 6-2 include the following:

• **Classification and marking:** QoS tools monitor traffic flows and classify packets based on the header contents. Messages are then marked by changing bits in the header.

• **Congestion avoidance:** When traffic exceeds available network resources, some traffic might be selectively dropped, delayed, or re-marked to avoid

Technet24

congestion.

• **Congestion management:** QoS tools manage the scheduling and shaping of traffic while packets wait their turn in a queue to exit the interface.

Classification and Marking

Classification refers to the process of matching fields in the headers to take some type of QoS action on the packet. These fields can include all the normal fields filtered by ACLs, as well as the Type of Service (ToS) field in an IPv4 packet or Traffic Class field in an IPv6 packet.

Marking refers to the process of changing bit values in the ToS or Traffic Class field. The contents of these two fields are identical, as Figure 6-3 shows.

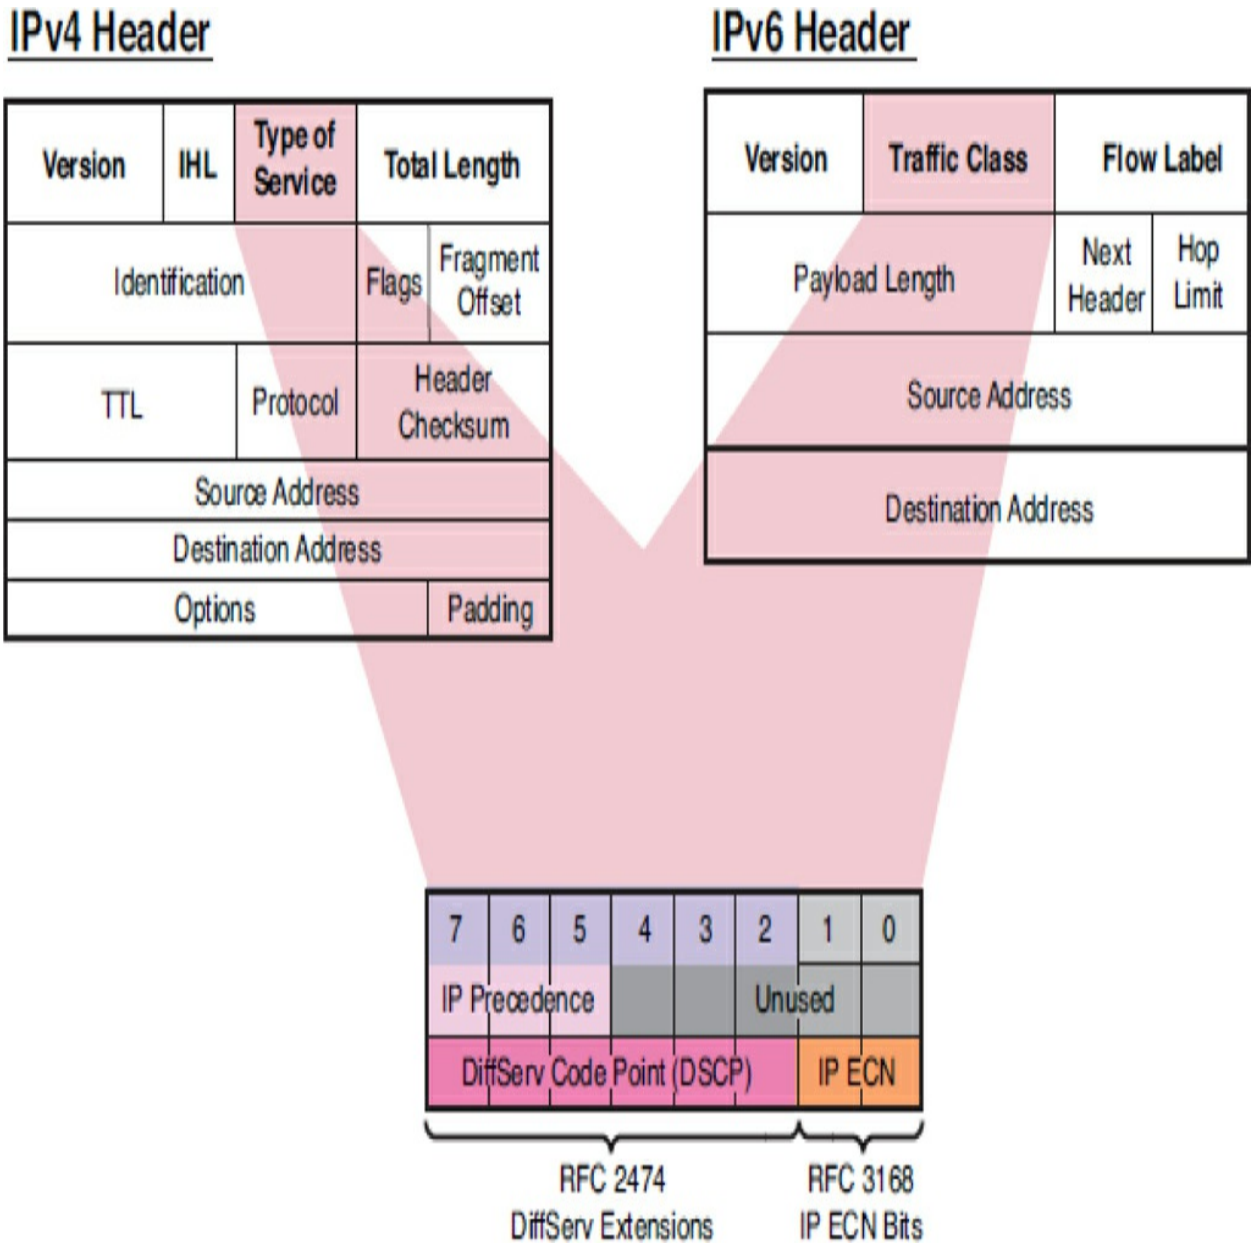

Figure 6-3 *The ToS and Traffic Class Fields in IPv4 and IPv6*

Figure 6-3 highlights the Differentiated Services Code Point (DSCP) bits, which are the core of the Differentiated Services (DiffServ) model for QoS. QoS tools can use the 2 bits allotted for IP Explicit Congestion Notification (ECN) to inform downstream routers of congestion in the traffic flow.

Technet24

DSCP and IPP

As standardized in RFC 2474, the 8 DSCP bits provide 64 different classifications that QoS can use. This is a vast improvement over the eight classifications allotted for the 3 bits in the previous IP Precedence (IPP) field (RFC 791). For backward compatibility, the DSCP bits include the Class Selector (CS) values that are designated to match the IPP bits, as in Figure 6-4.

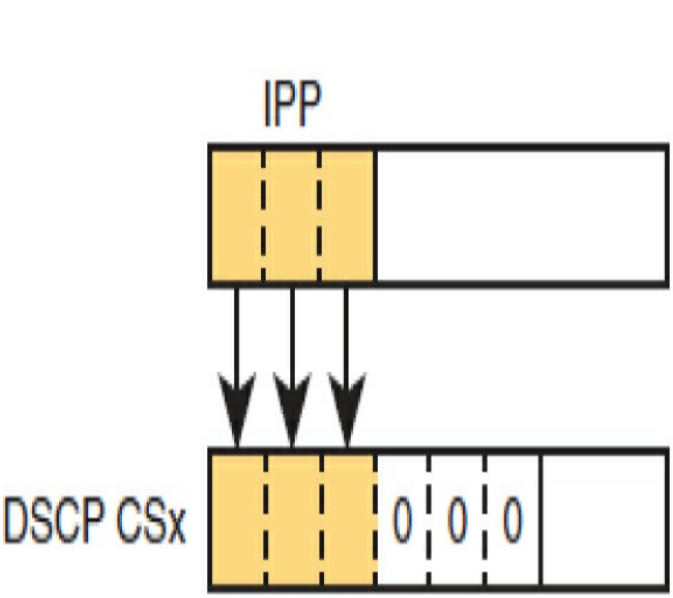

IPP	CS	Decimal DSCP
0	CS0	0
1	CS1	8
2	CS2	16
3	CS3	24
4	CS4	32
5	CS5	40
6	CS6	48
7	CS7	56

Figure 6-4 *The Class Selector Values*

For Layer 2 trunk links, the third byte of the 4-byte 802.1Q header is reserved for Class of Service (CoS), and QoS tools can use it to mark frames. However, this field exists only as long as the frame is traversing trunk links, as Figure 6-5 shows. To continue the same level of service as traffic is routed

on Layer 3, the ToS field must be marked.

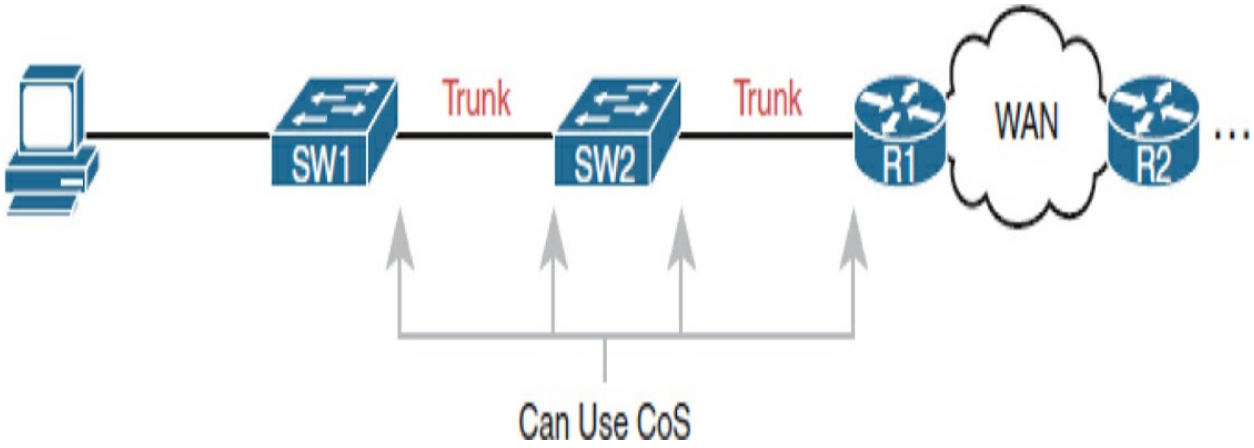

Figure 6-5 *CoS Marking Example*

Additional fields that can be marked for QoS include the Traffic Identifier (TID) field in the 802.11 frame and the EXP field in MPLS. Table 6-1 lists all the QoS fields.

Table 6-1 *QoS Marking Fields*

Technet24

Field	Name Header(s)	Length (bits)	Where Used
DSCP	IPv4, IPv6	6	End-to-end packet
IPP	IPv4, IPv6	3	End-to-end packet
CoS	802.1Q	3	Over VLAN trunk
TID	802.11	3	Over Wi-Fi
EXP	MPLS Label	3	Over MPLS WAN

Note

The MPLS EXP field was renamed the Traffic Class field in RFC 5462. However, EXP is still commonly used. The EXP name comes from the designation “experimental use.”

EF and AF

Expedited Forwarding (EF) is a single DSCP decimal value of 46 that is suggested for use with packets that require low latency, low jitter, and low loss. QoS implementations typically use EF to mark voice packets.

Assured Forwarding (AF), specified in RFC 2597, defines a set of 12 DSCP values that are arranged in a matrix, as in Figure 6-6.

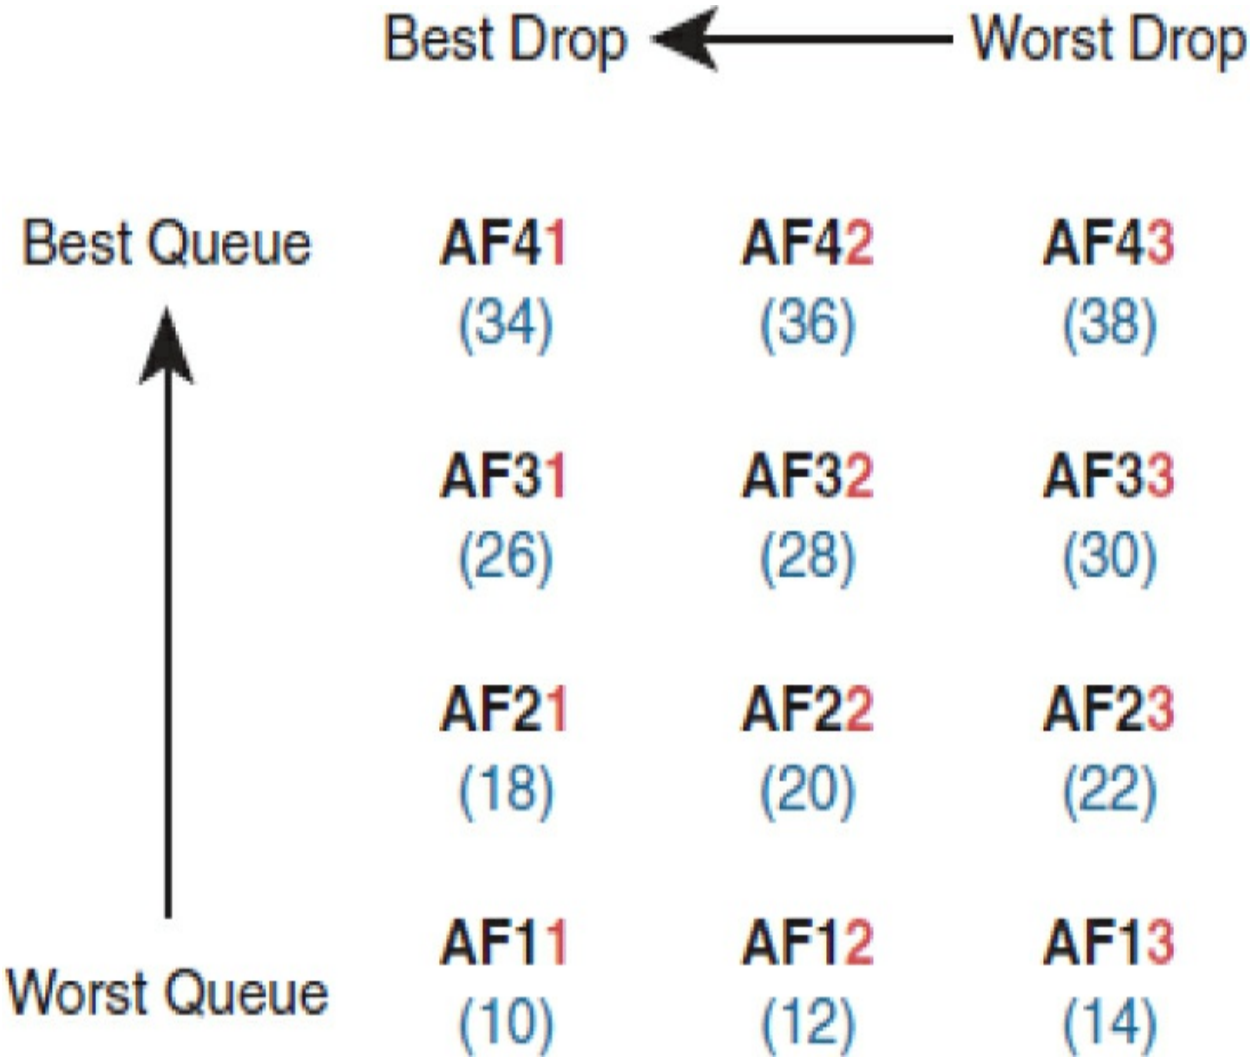

Figure 6-6 *AF DSCP Values*

The four rows in Figure 6-6 show the queue priorities. The three columns show the drop priority. The AF names follow the format AF*XY*, where *X* refers to the queue and *Y* refers to the drop priority.

Congestion Management

Technet24

Congestion management refers to the QoS tools used to manage queues as packets wait to exit an interface. Most networking devices can have a queuing system that can classify packets into multiple queues. A scheduler then decides which message to take next when the interface becomes available.

A popular tool is Class-Based Weighted Fair Queueing (CBWFQ), which assigns classes of traffic to queues and guarantees a minimum bandwidth for a queue. The scheduler then uses a round-robin algorithm to cycle through the queues in order, as in Figure 6-7.

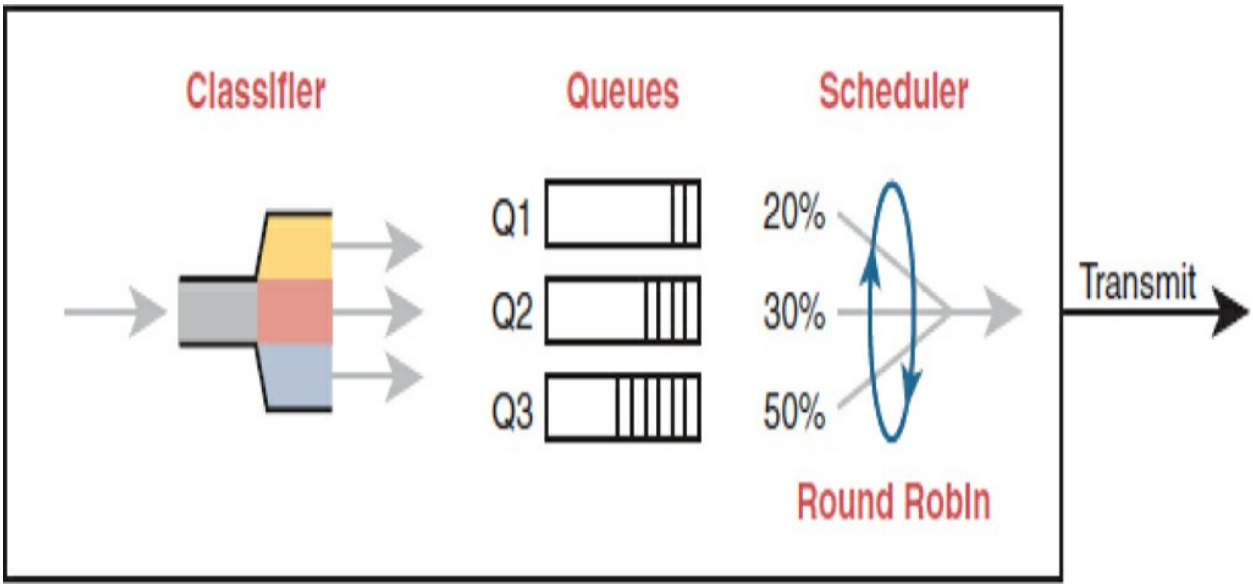

Figure 6-7 *CBWFQ Round-Robin Scheduling*

However, CBWFQ alone does not satisfy the needs of the most time-sensitive traffic type during periods of heavy bandwidth congestion. Each voice call needs between 30 and 320 kbps, maximum delay of 150 ms, maximum jitter of 30 ms, and less than 1% packet loss. The solution is to add

Low Latency Queueing (LLQ) to CBWFQ. The scheduler always takes the next voice packet from the LLQ, as Figure 6-8 shows.

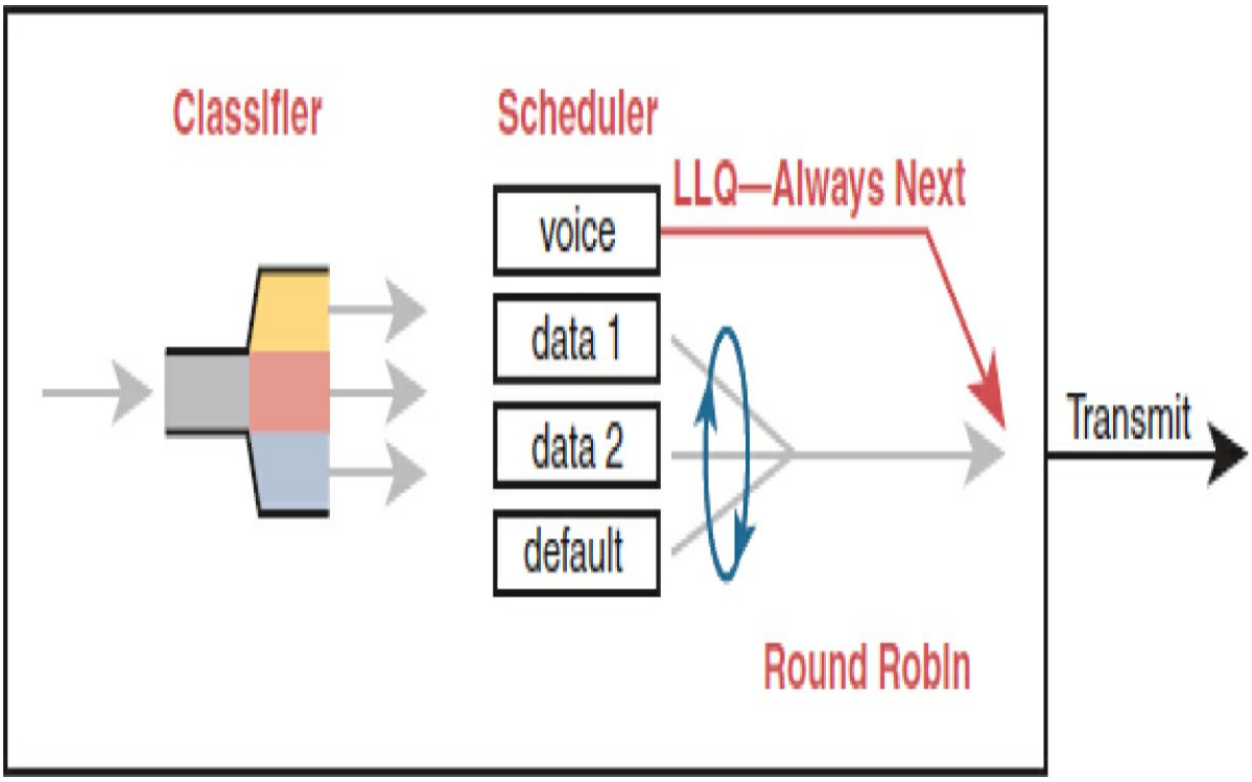

Figure 6-8 *Low Latency Queuing*

Policing, Shaping, and TCP Discards

Two tools that can help manage and avoid congestion on heavily utilized links are policing and shaping. Although these tools are not commonly used throughout the enterprise, they are particularly helpful at the WAN edge. Both tools attempt to keep the bit rate at or below a specified speed. Policers drop packets, and shapers delay packets by placing them in a queue.

Policing makes sense at the WAN edge. For example, consider a Metro Ethernet WAN link that is contracted to allow no more than 200 Mbps, as in

Technet24

Figure 6-9.

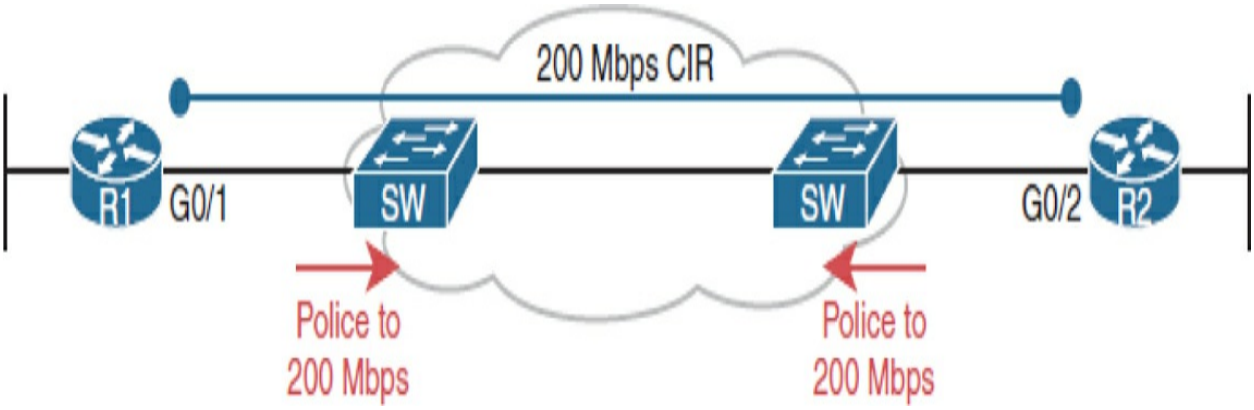

Figure 6-9 *WAN Edge with a CIR Below Link Speed*

The service provider (SP) uses policing to match the Committed Information Rate (CIR). If the customer exceeds the 200-Mbps CIR, the SP can drop the excess packets or remark the excess packets but still allow them through. Later, the excess packets can be discarded if the SP's network experiences congestion. Policing features include the following:

• Measure traffic over time and compare to a configured policing rate

• Allow for bursting traffic during slow times

• Discard excess messages or remark for discard later if congestion occurs downstream

On the customer side of the link in Figure 6-9, the network administrator can use a shaper to slow traffic to match the 200-Mbps CIR. The shaper slows traffic by queuing packets and then scheduling packets based on the shaping

rate, as Figure 6-10 shows.

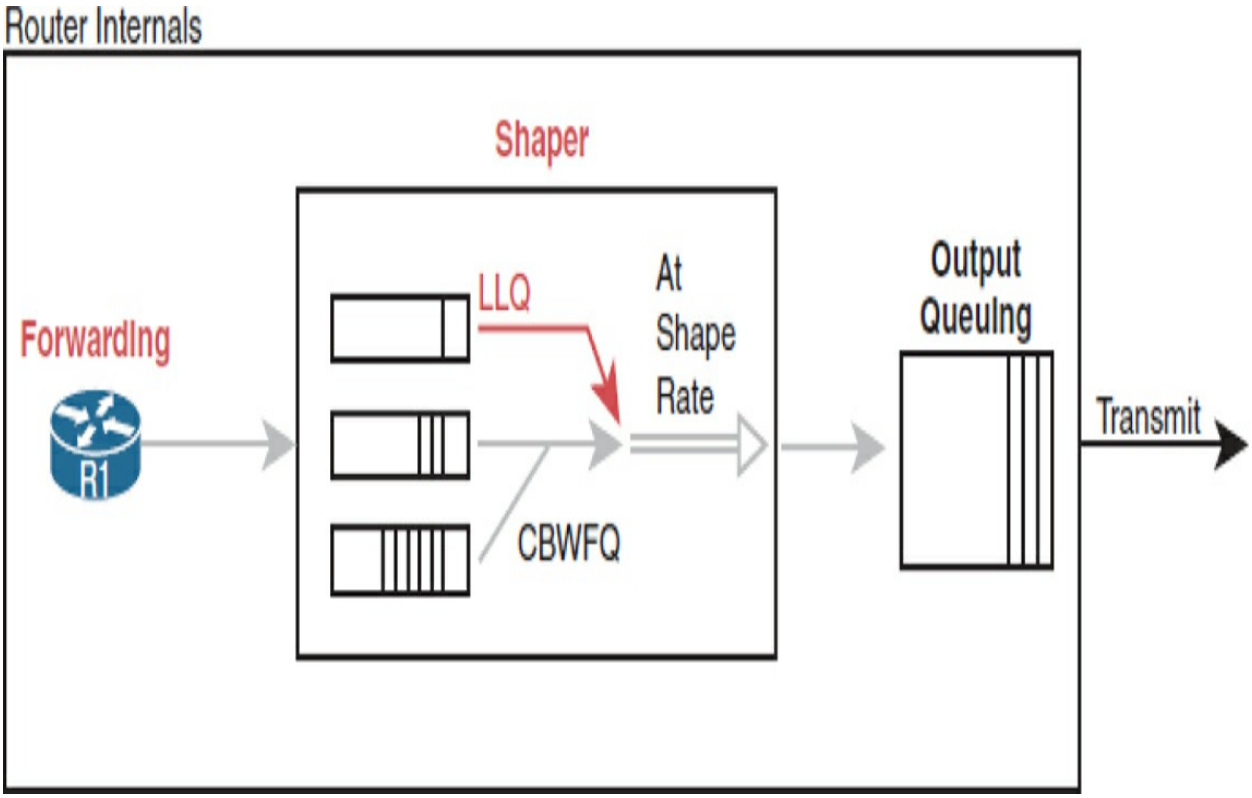

Figure 6-10 *Shaping with LLQ and CBWFQ*

Shaping cannot slow the physical speed of an interface. Instead, it sends and waits. For example, with a 200-Mpbs CIR and a 1000-Mbps interface, the shaper sends traffic at 1000 Mbps 20% of the time. The other 80% of the time, the shaper is waiting.

This send–wait tactic can adversely impact time-sensitive voice and video traffic. Therefore, it is recommended that you set the time interval to 10 ms. Then the shaper will send 1000 Mbps for 2 ms and wait for 8 ms. This ensures that a voice packet will not have to wait more than 10 ms before being sent, which is well below the 150 ms maximum delay requirement.

Technet24

The key features of shapers follow:

• Measure traffic over time and compare it to a configured shaping rate

• Allow for bursting traffic during slow times

• Slow packets by queuing them and, over time, releasing them from the queue at the shaping rate

QoS and TCP

Without congestion-avoidance tools, tail drop can occur (see Figure 6-11).

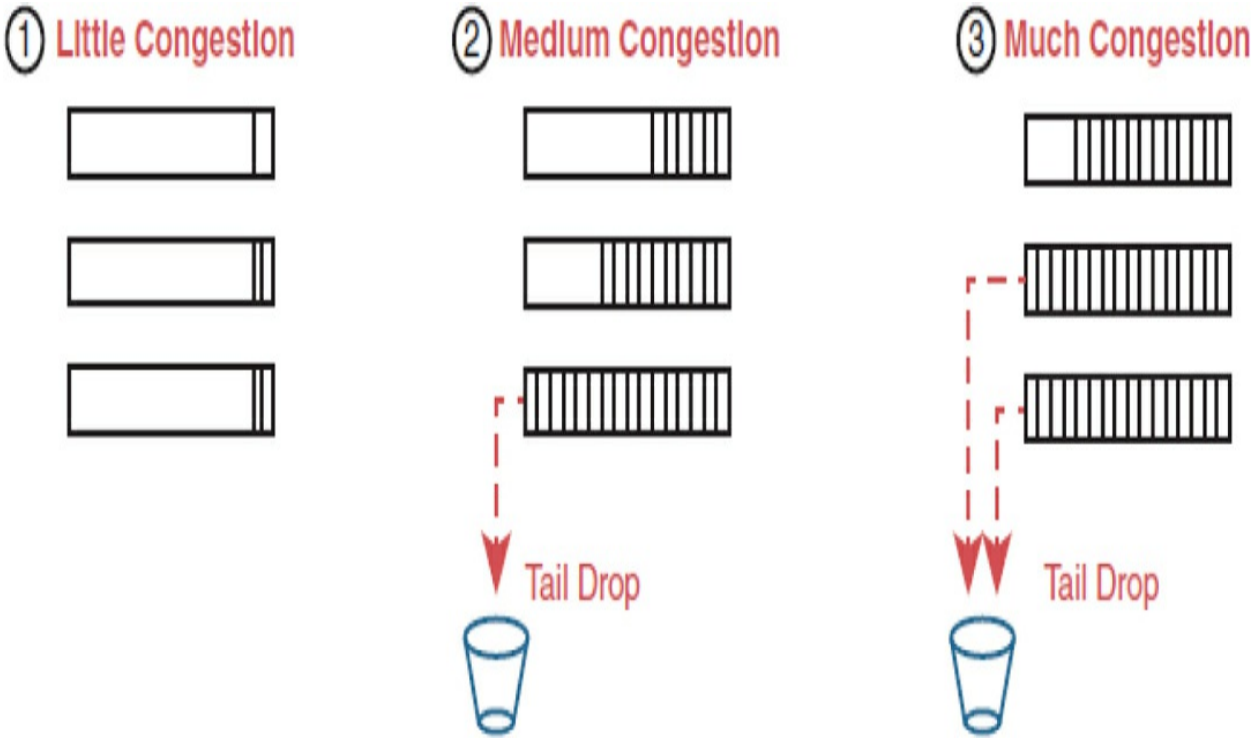

Figure 6-11 *Tail Drop Example*

As the lower queues fill up, the packets received last are dropped.

TCP’s connection-oriented services help QoS tools minimize tail drop. Recall that TCP uses a windowing process between sender and receiver to dynamically change the amount of data that is sent before an acknowledgment must be received. QoS tools can exploit this windowing feature by discarding some TCP segments before the queues fill. This forces the TCP connections to slow, reduces congestion, and avoids tail drop.

QoS tools monitor the depth of the queues over time. Configured thresholds specify what percentage of TCP packets should be dropped as the queue fills, as in Figure 6-12.

Technet24

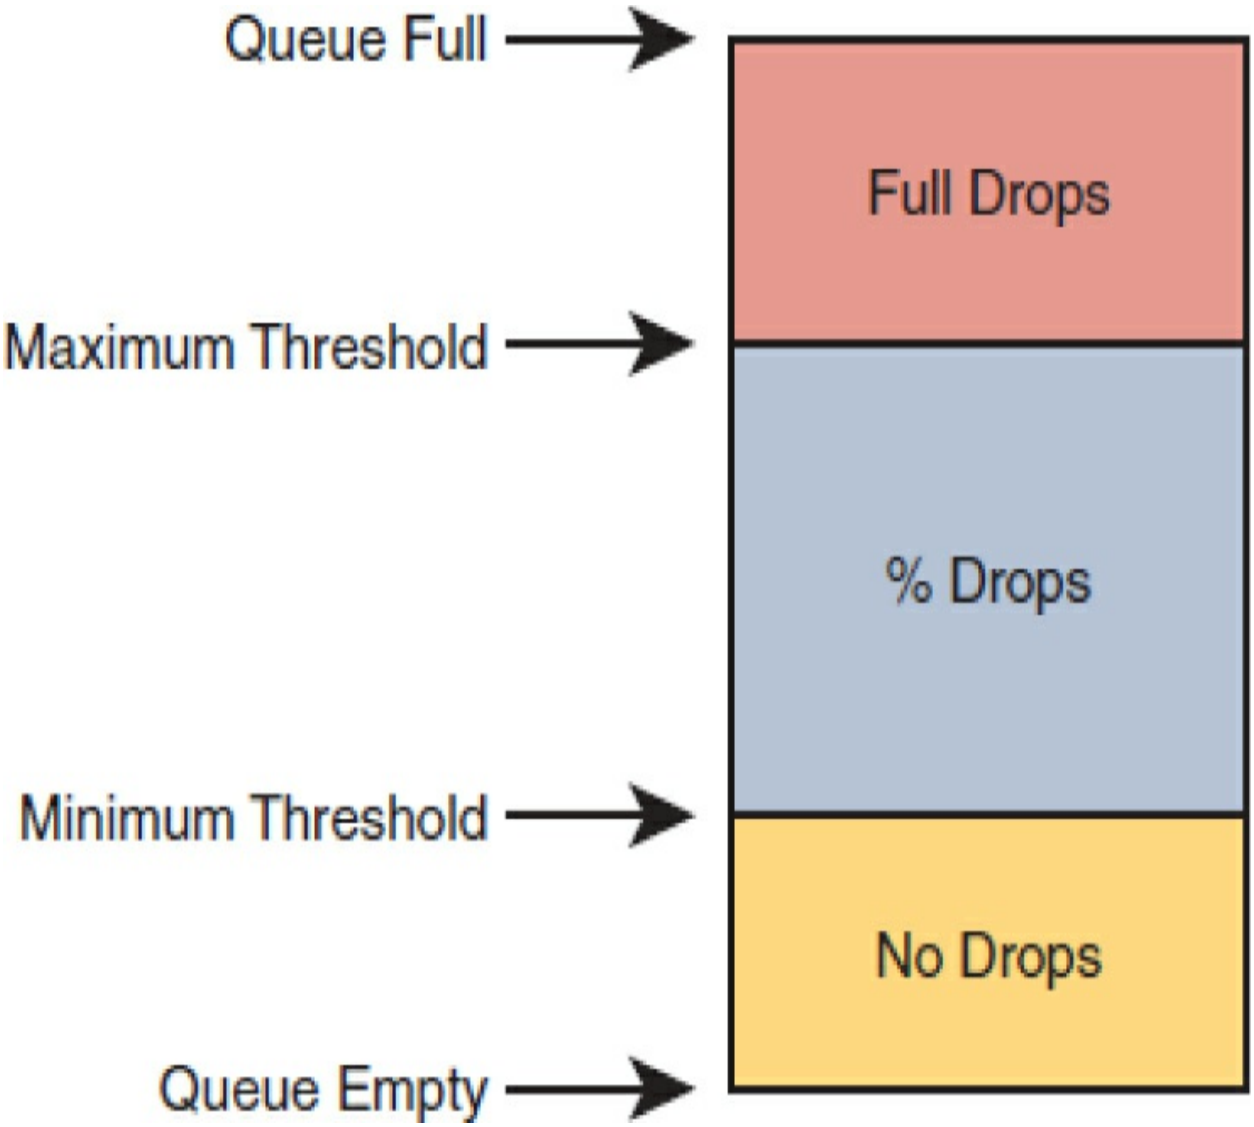

Figure 6-12 *Queue Thresholds for Discarding TCP Packets*

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	9
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	15

Technet24

Day 5. CDP and LLDP

CCNA 200-301 Exam Topic

• 2.3 Configure and verify Layer 2 discovery protocols (Cisco Discovery Protocol and LLDP)

Key Topics

Cisco Discovery Protocol (CDP) is a Cisco-proprietary Layer 2 protocol used to gather information about Cisco devices on the same data link. Cisco devices also support Link Layer Discovery Protocol (LLDP), which is a standards-based neighbor discovery protocol similar to CDP. Today we review the configuration and verification of CDP and LLDP.

CDP Overview

As Figure 5-1 shows, CDP sends advertisements to directly connected devices.

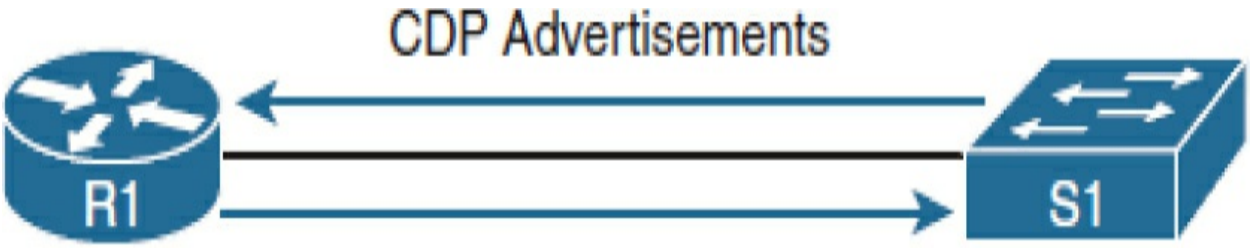

Figure 5-1 *CDP Sends Advertisements Between Directly Connected Devices*

CDP runs on all Cisco-manufactured equipment. It gathers the protocol

addresses of neighboring devices and discovers the platforms of those devices. CDP runs over the data link layer only. This means that two systems that support different Layer 3 protocols can learn about each other. Table 5-1 summarizes the CDP defaults.

Table 5-1 *CDP Defaults*

Parameter	Default
CDP	Enabled globally and on all interfaces
CDP version	Version 2
CDP timer	60 seconds
CDP holdtime	180 seconds

CDP can assist in network discovery and troubleshooting. CDP advertises the following helpful information:

- **Device ID:** The hostname of the neighboring device
- **Addresses:** The IPv4 and IPv6 addresses used by the device
- **Port ID:** The name of the local port or the remote port

Technet24

• **Capabilities:** Whether the device is a router or a switch or has other capabilities

• **Version:** The version of CDP running on the device

• **Platform:** The hardware platform of the device, such as a Cisco 1941 router or 2960 switch

CDP Configuration

CDP is enabled on a device for all interfaces, as Example 5-1 shows.

Example 5-1 *CDP Running on All Interfaces*

```
Router# show cdp interface
Embedded-Service-Engine0/0 is administratively down, l
   Encapsulation ARPA
   Sending CDP packets every 60 seconds
   Holdtime is 180 seconds
GigabitEthernet0/0 is administratively down, line prot
   Encapsulation ARPA
   Sending CDP packets every 60 seconds
   Holdtime is 180 seconds
GigabitEthernet0/1 is up, line protocol is up
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
Serial0/0/0 is administratively down, line protocol is
```

```
    Encapsulation HDLC
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
  Serial0/0/1 is administratively down, line protocol is
     Encapsulation HDLC
     Sending CDP packets every 60 seconds
     Holdtime is 180 seconds

     cdp enabled interfaces  : 5
     interfaces up           : 1
     interfaces down         : 4
  Router# show ip interface brief
  Interface IP-Address OK? Method Status Protocol
  Embedded-Service-Engine0/0 unassigned YES unset admini
  GigabitEthernet0/0 unassigned YES unset administrative
  GigabitEthernet0/1 unassigned YES unset up up
  Serial0/0/0 unassigned YES unset administratively down
  Serial0/0/1 unassigned YES unset administratively down
  Router#
```

Notice in the output in Example 5-1 that an interface does not have to be configured with a Layer 3 address to send or receive CDP advertisements. The interface only needs to be activated with the **no shutdown** command. In Example 5-2, the switch connected to the router in Example 5-1 has gathered CDP information about the router. The two devices are communicating across the Layer 2 link without any Layer 3 addressing.

Technet24

Example 5-2 *Sending Layer 2 Messages*

```
Switch# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - So
                  S - Switch, H - Host, I - IGMP, r - Re
                  D - Remote, C - CVTA, M - Two-port Mac

Device ID  Local Intrfce   Holdtme   Capability   Plat
Router     Fas 0/5         155       R B S I      CISC
Switch#
```

To disable CDP on the device, use the CDP global configuration command **no cdp run**:

```
Router(config)# no cdp run
```

Verify that the device is no longer running CDP by using the **show cdp** command:

```
Router# show cdp
% CDP is not enabled
Router#
```

After waiting for the 180-second holdtime to expire on the switch, you can verify that the switch is no longer receiving information about the router (see

Example 5-3).

Example 5-3 *Verifying That a Switch Has No CDP Neighbors*

```
Switch# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Sc
                  S - Switch, H - Host, I - IGMP, r - R
                  D - Remote, C - CVTA, M - Two-port Ma

Device ID      Local Intrfce  Holdtme  Capability  Pla
Switch#
```

You can also disable CDP on a per-interface basis. This configuration option is a security best practice for interfaces that are connected to untrusted networks. To disable CDP on an interface, use the **no cdp enable** command (see Example 5-4).

Example 5-4 *Disabling CDP on an Interface*

```
Router(config)# interface s0/0/0
Router(config-if)# no cdp enable
Router(config-if)# end
Router# show cdp interface
Embedded-Service-Engine0/0 is administratively down, l
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
```

Technet24

```
    Holdtime is 180 seconds
  GigabitEthernet0/0 is administratively down, line prot
    Encapsulation ARPA
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
  GigabitEthernet0/1 is up, line protocol is up
    Encapsulation ARPA
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
  Serial0/0/1 is administratively down, line protocol is
    Encapsulation HDLC
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds

    cdp enabled interfaces : 4
    interfaces up            : 1
    interfaces down          : 3
  Router#
```

Notice in the output of the **show cdp interface** command that the Serial 0/0/0 interface is no longer listed, as it was in Example 5-1.

To adjust the time for CDP advertisements, use the **cdp timer** global configuration command:

```
Router(config)# cdp timer seconds
```

The range is 5 to 254 seconds, and the default is 60 seconds. If you modify the CDP timer, you should also modify the holdtime with the **cdp holdtime** global configuration command:

```
Router(config)# cdp holdtime seconds
```

The range is from 10 to 255, and the default is 180 seconds.

CDP Verification

You have already seen examples of **show cdp**, **show cdp neighbors**, and **show cdp interface**. The **show cdp neighbors detail** command lists all the information CDP gathers about directly connected neighbors. In Example 5-5, switch S3 knows a variety of information about R1, including the IP address and Cisco IOS version running on the router.

Example 5-5 *CDP Detailed Information*

```
S3# show cdp neighbors detail
-------------------------
Device ID: R3.31days.com
Entry address(es):
  IP address: 192.168.1.1
Platform: Cisco CISCO1941/K9, Capabilities: Router Sou
IGMP
Interface: FastEthernet0/5, Port ID (outgoing port): G
```

Technet24

```
  Holdtime : 162 sec

  Version :
  Cisco IOS Software, C1900 Software (C1900-UNIVERSALK9-
  RELEASE SOFTWARE (fc2)
  Technical Support: http://www.cisco.com/techsupport
  Copyright (c) 1986-2015 by Cisco Systems, Inc.
  Compiled Fri 06-Feb-15 17:01 by prod_rel_team

  advertisement version: 2
  Duplex: full
  Power Available TLV:

        Power request id: 0, Power management id: 0, Pow
  management level: 0
  Management address(es):
    IP address: 192.168.1.1

  S3#
```

Note

The Device ID value shows R3.31days.com because R3 is configured with the **ip domain-name** command and SSH remote access.

When documentation is lacking or incomplete, you can use CDP to gather information about devices and discover the network topology. Example 5-6 shows how to remotely access R3 and discover that R2 is connected to R3.

Example 5-6 *Discovering More Devices*

```
S3# ssh -l admin 192.168.1.1
Password:
R3> show cdp neighbors detail
-------------------------
Device ID: S3.31days.com
Entry address(es):
    IP address: 192.168.1.2
Platform: cisco WS-C2960-24TT-L, Capabilities: Switch
Interface: GigabitEthernet0/1, Port ID (outgoing port)
Holdtime : 126 sec

Version :
Cisco IOS Software, C2960 Software (C2960-LANBASEK9-M)
RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2014 by Cisco Systems, Inc.
Compiled Thu 23-Oct-14 14:49 by prod_rel_team

advertisement version: 2
Protocol Hello: OUI=0x00000C, Protocol ID=0x0112; payl
value=00000000FFFFFFFF010221FF000000000000CD996E87400
```

Technet24

```
VTP Management Domain: ‘’
Native VLAN: 1
Duplex: full
Management address(es):
 IP address: 192.168.1.2

-------------------------
Device ID: R2.31days.com
Entry address(es):
   IP address: 192.168.10.1
Platform: Cisco CISCO1941/K9, Capabilities: Router Sou
IGMP
Interface: Serial0/0/1, Port ID (outgoing port): Seria
Holdtime : 148 sec
Version :
Cisco IOS Software, C1900 Software (C1900-UNIVERSALK9-
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2015 by Cisco Systems, Inc.
Compiled Fri 06-Feb-15 17:01 by prod_rel_team

advertisement version: 2
Management address(es):
   IP address: 192.168.10.1

Total cdp entries displayed : 2
R3>
```

You can now access router R2, with authentication, and continue discovering the network.

As Example 5-7 shows, you use the **show cdp traffic** command to verify how many CDP packets a device has sent and received.

Example 5-7 *Verifying the Number of CDP Packets*

```
R3# show cdp traffic
CDP counters :
    Total packets output: 758, Input: 724
    Hdr syntax: 0, Chksum error: 0, Encaps failed: 1
    No memory: 0, Invalid packet: 0,
    CDP version 1 advertisements output: 0, Input: 0
    CDP version 2 advertisements output: 758, Input: 7
R3#
```

LLDP Overview

In addition to supporting CDP, Cisco devices also support LLDP, which is a vendor-neutral open standard (IEEE 802.1AB). LLDP works with routers, switches, and wireless LAN access points. As with CDP, LLDP is a neighbor discovery protocol that is used for network devices to advertise information about themselves to other devices on the network. Also as with CDP, LLDP enables two systems running different network layer protocols to learn about

Technet24

each other.

Table 5-2 summarizes the CDP defaults.

Table 5-2 *LLDP Defaults*

Parameter	Default
LLDP	Disabled globally and on all interfaces
LLDP timer	30 seconds
LLDP holdtime	120 seconds
LLDP reinitialization delay	2 seconds

Note

The reinitialization delay is the number of seconds the device waits after LLDP is disabled on a port before it accepts a configuration to reenable LLDP.

LLDP Configuration

To enable LLDP globally, enter the **lldp run** command:

```
Router(config)# lldp run
```

When enabled globally, LLDP is enabled on all interfaces. To disable LLDP on an interface, use the **no lldp transmit** and **no lldp receive** commands:

```
Router(config)# interface interface-id
Router(config-if)# no lldp transmit
Router(config-if)# no lldp receive
Router(config-if)# end
Router#
```

To adjust the time for LLDP advertisements, use the **lldp timer** global configuration command:

```
Router(config)# lldp timer seconds
```

The range is 5 to 65534 seconds, and the default is 30 seconds. If you modify the CDP timer, you should also modify the holdtime with the **cdp holdtime** global configuration command:

```
Router(config)# lldp holdtime seconds
```

Technet24

The range is from 0 to 65535, and the default is 120 seconds. You can also modify the delay time for LLDP to initialize on any interface with the **lldp reinit** global configuration command:

```
Router(config)# lldp reinit seconds
```

The range is 2 to 5 seconds, and the default is 2 seconds.

For the topology in Figure 5-2, the policy is that LLDP should have the same timers as CDP. Routers should not transmit LLDP messages out LAN interfaces.

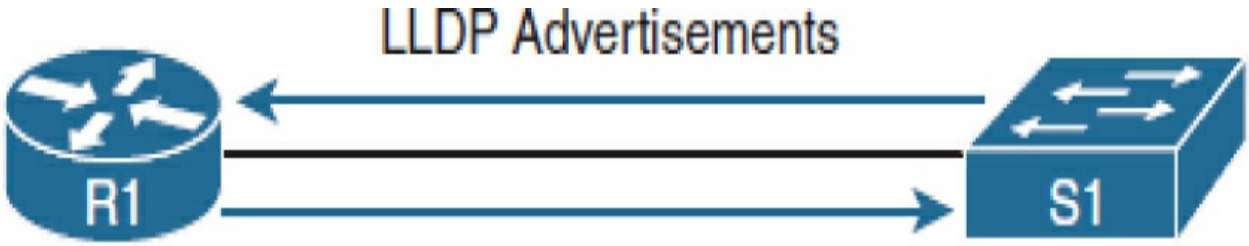

Figure 5-2 *LLDP Configuration Topology*

Example 5-8 shows the commands to implement the LLDP policy.

Example 5-8 *LLDP Configuration*

```
R1(config)# lldp run
R1(config)# lldp timer 60
R1(config)# lldp holdtime 180
R1(config)# interface g0/1
R1(config-if)# no lldp transmit
```

```
R1(config-if)# end
R1#
S1(config)# lldp run
S1(config)# lldp timer 60
S1(config)# lldp holdtime 180
S1(config)# end
S1#
```

LLDP Verification

The LLDP verification commands are similar to those in CDP. Simply replace the keyword **cdp** with **llpd**. Example 5-9 shows output from the LLDP verification commands.

Example 5-9 *LLDP Verification Commands*

```
R1# show lldp

Global LLDP Information:
    Status: ACTIVE
    LLDP advertisements are sent every 60 seconds
    LLDP hold time advertised is 180 seconds
    LLDP interface reinitialization delay is 2 seconds
R1# show lldp interface
<output omitted>
GigabitEthernet0/0:
    Tx: enabled
```

Technet24

```
    Rx: enabled
    Tx state: INIT
    Rx state: WAIT PORT OPER
GigabitEthernet0/1:
    Tx: disabled
    Rx: enabled
    Tx state: INIT
    Rx state: WAIT FOR FRAME
<output omitted>
R1# show lldp neighbors
Capability codes:
   (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS C
   (W) WLAN Access Point, (P) Repeater, (S) Station, (

Device ID Local Intf Hold-time Capability Port ID
S1         Gi0/1      180       B          Fa0/5

Total entries displayed: 1

R1# show lldp neighbors detail
------------------------------------------------
Local Intf: Gi0/1
Chassis id: 0cd9.96e8.8a00
Port id: Fa0/5
Port Description: FastEthernet0/5
System Name: S1.31days.com

System Description:
```

```
Cisco IOS Software, C2960 Software (C2960-LANBASEK9-M)
RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2014 by Cisco Systems, Inc.
Compiled Thu 23-Oct-14 14:49 by prod_rel_team

Time remaining: 127 seconds
System Capabilities: B
Enabled Capabilities: B
Management Addresses:
    IP: 172.16.1.2
Auto Negotiation - supported, enabled
Physical media capabilities:
   100base-TX(FD)
   100base-TX(HD)
   10base-T(FD)
   10base-T(HD)
Media Attachment Unit type: 16
Vlan ID: 1

Total entries displayed: 1

R1# show lldp traffic

LLDP traffic statistics:
   Total frames out: 171
   Total entries aged: 0
   Total frames in: 34
```

Technet24

```
    Total frames received in error: 0
    Total frames discarded: 0
    Total TLVs discarded: 0
    Total TLVs unrecognized: 0
R1#
```

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	10
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	13
Portable Command Guide	13

Day 4. Device Monitoring, Management, and Maintenance

CCNA 200-301 Exam Topics

• 4.2 Configure and verify NTP operating in a client and server mode

• 4.4 Explain the function of SNMP in network operations

• 4.5 Describe the use of syslog features, including facilities and severity levels

• 4.9 Describe the capabilities and function of TFTP/FTP in the network

Key Topics

The review today covers tools used to monitor, manage, and maintain routers and switches. For device monitoring, we review how to configure and verify Simple Network Management Protocol (SNMP), syslog, and Network Time Protocol (NTP). For device management, we review backing up and restoring your device configuration and Cisco IOS licensing. For device maintenance, we review file system management, Cisco IOS upgrades, and password recovery.

SNMP Operation

SNMP is an application layer protocol that provides a message format for

Technet24

communication between managers and agents.

SNMP Components

The SNMP system consists of three elements:

• SNMP manager

• SNMP agents (managed node)

• Management Information Base (MIB)

SNMP Messages

The SNMP manager is part of a network management system (NMS) and runs SNMP management software. SNMP agents are managed devices. The MIB stores SNMP variables. SNMP uses three basic messages between SNMP managers and agents: get, set, and trap messages. The SNMP manager uses get messages to poll a device for information and set messages to change a device parameter. An SNMP agent can use SNMP traps to independently notify the NMS when a problem occurs.

For example, SNMP can monitor the CPU utilization on a Cisco router. The NMS can sample this value periodically and warn the network administrator when the value deviates from the baseline. An SNMP agent can also be configured to send a trap message when CPU utilization is driving away from normal values for the network. Table 4-1 summarizes the get and set actions.

Table 4-1 *get and set SNMP Operations*

Operation	Description
get-request	Retrieves a value from a specific variable.
get-next-request	Retrieves a value from a variable within a table. The SNMP manager does not need to know the exact variable name. A sequential search is performed to find the needed variable from within a table.
get-bulk-request	Retrieves large blocks of data, such as multiple rows in a table, that would otherwise require the transmission of many small blocks of data. (This works only with SNMPv2 or later.)
get-response	Replies to a **get-request**, **get-next-request**, or **set-request** sent by an NMS.
set-request	Stores a value in a specific variable.

SNMP Versions

Several versions of SNMP exist:

Technet24

• **SNMPv1:** The original SNMP, defined in RFC 1157.

• **SNMPv2c:** Defined in RFCs 1901 to 1908. Utilizes a community string–based administrative framework.

• **SNMPv3:** Interoperable standards-based protocol originally defined in RFCs 2273 to 2275. Provides secure access to devices by authenticating and encrypting packets over the network.

SNMPv1 and SNMPv2c use community strings that control access to the MIB. Community strings are plaintext passwords. Two types of community strings exist:

• **Read-only (ro):** Provides access to the MIB variables but does not allow these variables to be changed (only read)

• **Read-write (rw):** Provides read and write access to all objects in the MIB

The Management Information Base

The MIB organizes variables hierarchically. MIB variables enable the management software to monitor and control the network device. Formally, the MIB defines each variable as an object ID (OID). OIDs uniquely identify managed objects in the MIB hierarchy. The MIB organizes the OIDs based on RFC standards into a hierarchy of OIDs, usually shown as a tree.

RFCs define some common public variables. Figure 4-1 shows portions of

the MIB structure defined by Cisco Systems.

Technet24

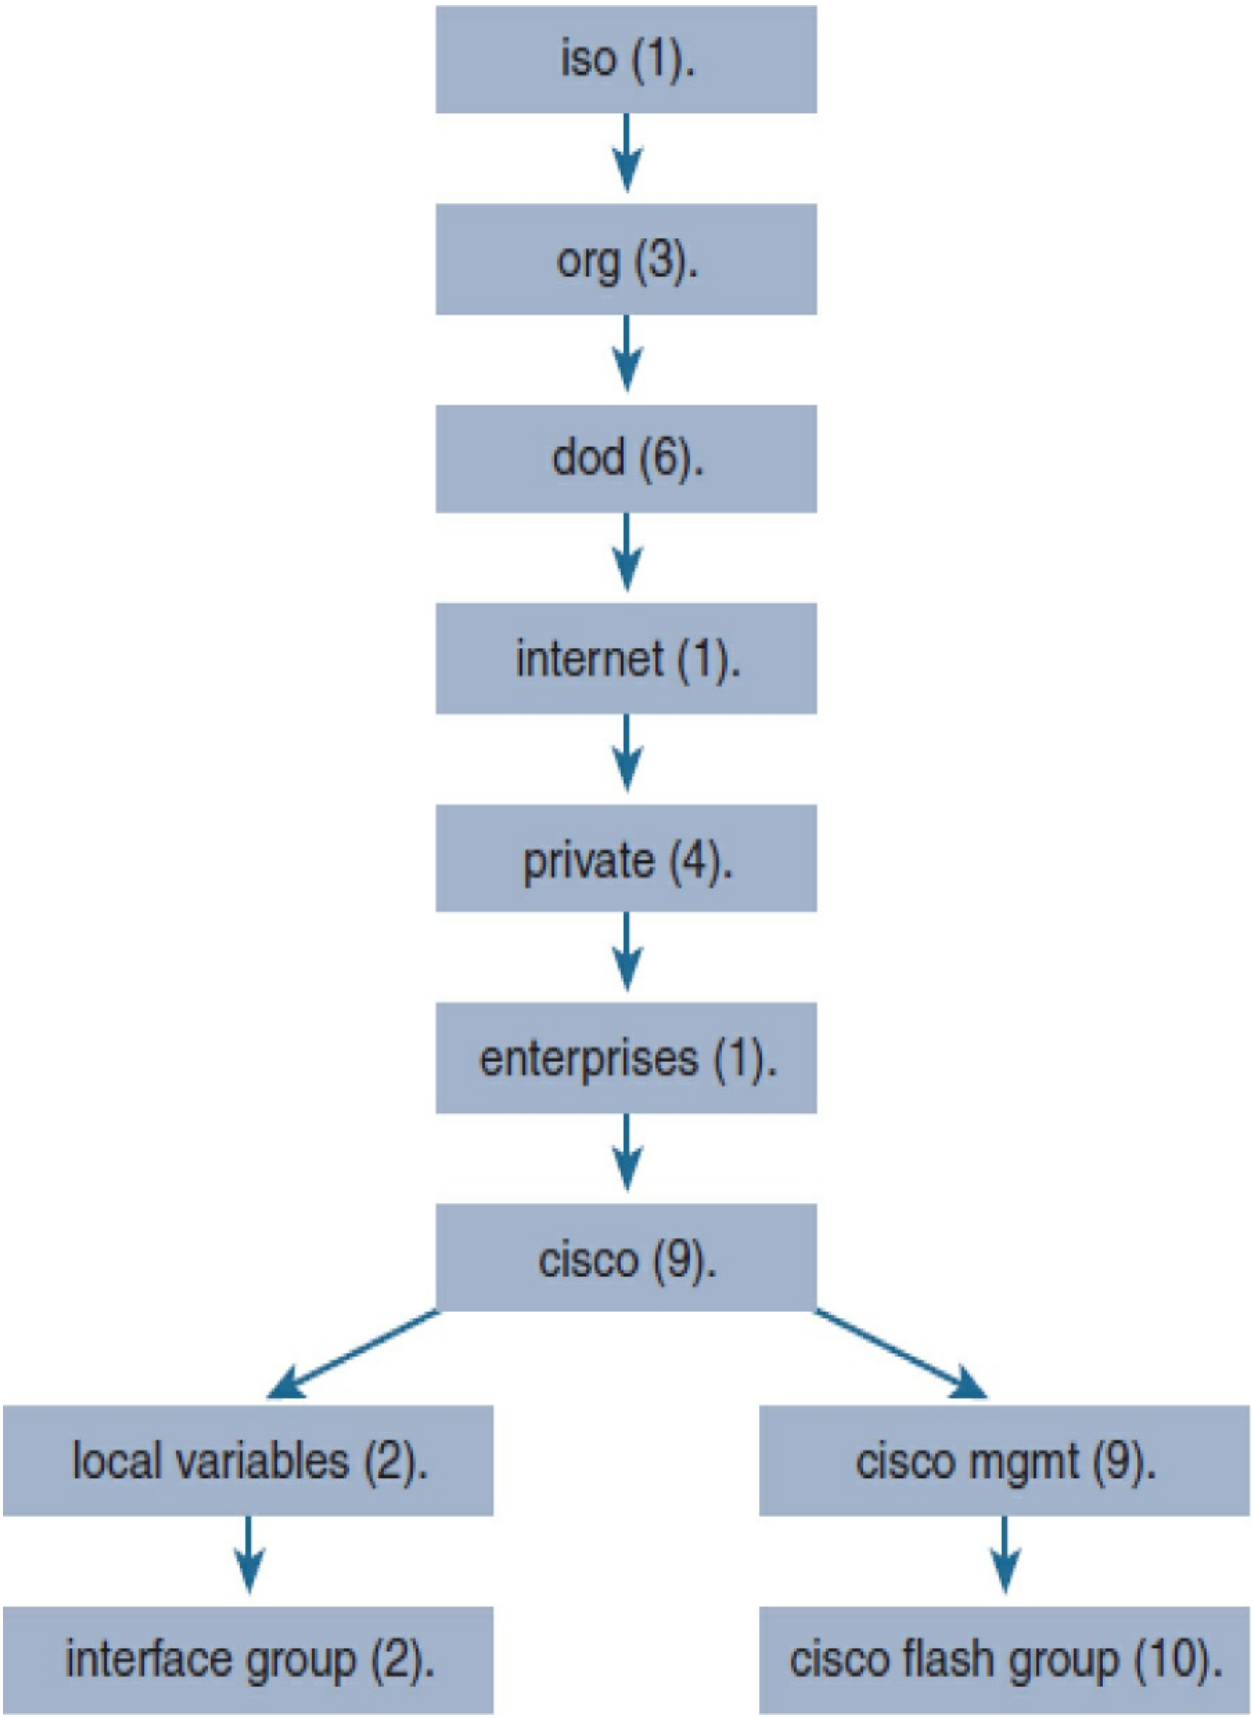
iso (1).
org (3).
dod (6).
internet (1).
private (4).
enterprises (1).
cisco (9).
local variables (2).
interface group (2).
cisco mgmt (9).
cisco flash group (10).

Figure 4-1 *Management Information Base Object IDs*

Note that the OID can be described in words or numbers to help locate a particular variable in the tree. For example, OIDs belonging to Cisco are numbered as follows: iso (1).org (3).dod (6).internet (1).private (4).enterprises (1).cisco (9). This is displayed or configured as 1.3.6.1.4.1.9.

One way to demonstrate using these OIDs is to look at how they can be implemented in the freeware SNMPGET utility. Example 4-1 shows how you might configure SNMPGET to obtain a 5-minute exponential moving average of the CPU busy percentage from a router.

Example 4-1 *Obtaining an MIB Value with SNMPGET*

```
[13:22][cisco@NMS~ ]$ snmpget -v2c -c community 10.250
1.3.6.1.4.1.9.2.1.58.0
SNMPv2-SMI::enterprises.9.2.1.58.0 = INTEGER: 11
```

The bold text shows a rather long command with several parameters highlighted:

- **-v2c:** The version of SNMP in use
- **-c community:** The SNMP password, called a *community string*
- **10.250.250.14:** The IP address of the monitored device

Technet24

• **1.3.6.1.4.1.9.2.1.58.0:** The numeric OID of the MIB variable

The last line shows the response. The output shows a shortened version of the MIB variable. It then lists the actual value in the MIB location; in this example, it means that the CPU is at 11% utilization.

Configuring SNMP

Configuring SNMPv2c on a Cisco router or switch requires only one global configuration command: **snmp-server community**. The following steps include some optional commands:

(Required) Configure the community string and access level (read-only or read-write) with the **snmp-server community** *string* {**RO**|**RW**} global command.

(Optional) Document the location of the device by using the **snmp-server location** *text-describing-location* global configuration command.

(Optional) Document the location of the device by using the **snmp-server contact** *contact-name* global configuration command.

(Optional) Restrict SNMP access to NMS hosts that are permitted by an access control list (ACL) by defining an ACL and referencing the ACL on the **snmp-server community** *string acl* global configuration command.

Example 4-2 demonstrates the use of the required and optional commands.

Example 4-2 *Configuring SNMPv2c for Read-Only Access*

```
R1(config)# ip access-list standard SNMP_ACCESS
R1(config-std-nacl)# permit host 172.16.3.110
R1(config-std-nacl)# exit
R1(config)# snmp-server community 4md!n0n1y RO SNMP_AC
R1(config)# snmp-server location Austin, TX
R1(config)# snmp-server contact Bob Smith
R1(config)# end
```

Verifying SNMP

To verify the SNMP configuration, use the **show snmp** command (see Example 4-3).

Example 4-3 *Verifying SNMP*

```
R1# show snmp
Chassis: FTX1636848Z
Contact: Bob Smith
Location: Lima, OH
0 SNMP packets input
   0 Bad SNMP version errors
   0 Unknown community name
   0 Illegal operation for community name supplied
   0 Encoding errors
```

Technet24

```
    0 Number of requested variables
    0 Number of altered variables
    0 Get-request PDUs
    0 Get-next PDUs
    0 Set-request PDUs
    0 Input queue packet drops (Maximum queue size 1000
359 SNMP packets output
    0 Too big errors (Maximum packet size 1500)
    0 No such name errors
    0 Bad values errors
    0 General errors
    0 Response PDUs
    359 Trap PDUs
SNMP Dispatcher:
   queue 0/75 (current/max), 0 dropped
SNMP Engine:
   queue 0/1000 (current/max), 0 dropped

SNMP logging: enabled
 Logging to 172.16.3.10, 0/10, 359 sent, 0 dropped.
```

The **show snmp** command output does not display information related to the SNMP community string or the associated ACL (if applicable). Example 4-4 shows the use of the **show snmp community** command to display the SNMP community string and ACL information.

Example 4-4 *Verifying SNMP Community Strings*

```
R1# show snmp community

Community name: ILMI
Community Index: cisco0
Community SecurityName: ILMI
storage-type: read-only active

Community name: 4md!n0n1y
Community Index: cisco7
Community SecurityName: 4md!n0n1y
storage-type: nonvolatile active access-list: SNMP_ACC


Community name: 4md!n0n1y
Community Index: cisco8
Community SecurityName: 4md!n0n1y
storage-type: nonvolatile active
access-list: SNMP_ACCESS
```

Syslog

Syslog is a term used to describe a standard that the IETF first documented in RFC 3164 in 2001. It is a popular protocol that many networking devices use, including routers, switches, application servers, firewalls, and other network appliances. These devices can send their messages across the network to be stored on syslog servers for later access by network administrators.

Technet24

Syslog Operation

Syslog uses UDP port 514 to send event notification messages across IP networks to event message collectors, as Figure 4-2 illustrates.

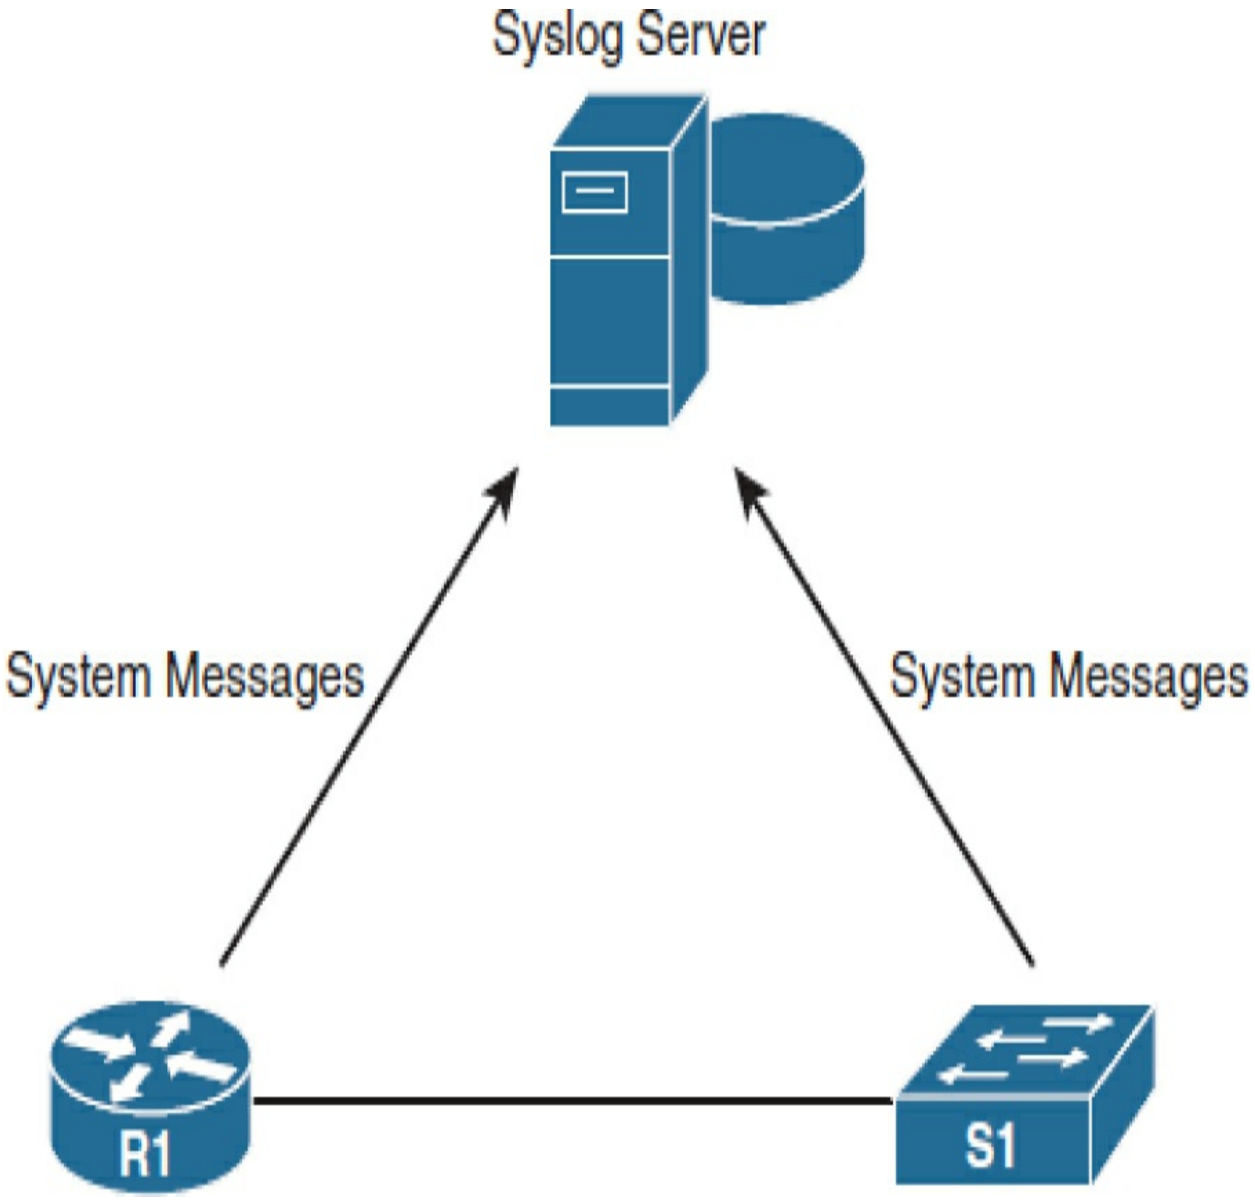

Figure 4-2 *Syslog Server Example*

The syslog logging service provides three primary capabilities:

• Gathering logging information for monitoring and troubleshooting

• Selecting the type of logging information that is captured

• Specifying the destinations of captured syslog messages

On Cisco network devices, the syslog protocol starts by sending system messages and debug output to a local logging process internal to the device. It is possible to remotely monitor system messages by viewing the logs on a syslog server or by accessing the device through Telnet, Secure Shell (SSH), or the console port.

Cisco devices produce syslog messages as a result of network events. Every syslog message contains a severity level and a facility. Table 4-2 shows the complete list of syslog severity levels.

Table 4-2 *Syslog Severity Level*

Technet24

Severity Name	Severity Level	Explanation
Emergency	Level 0	System unusable
Alert	Level 1	Immediate action needed
Critical	Level 2	Critical condition
Error	Level 3	Error condition
Warning	Level 4	Warning condition
Notification	Level 5	Normal but significant condition
Informational	Level 6	Informational message
Debugging	Level 7	Debugging message

In addition to specifying the severity, syslog messages contain information on the facility. Syslog facilities are service identifiers that identify and categorize system state data for error and event message reporting. The logging facility options that are available are specific to the networking device. Common syslog message facilities reported on Cisco IOS routers include the following:

• IP

• OSPF protocol

• SYS operating system

• IP Security (IPsec)

• Interface IP (IF)

The default format for syslog messages follows:

```
  seq no: timestamp: %facility-severity-MNEMONIC: descri
```

Table 4-3 describes the fields in a syslog message.

Table 4-3 *Syslog Message Format*

Technet24

Field	Explanation
seq no	Sequence number stamped on the log message. Appears only if the **service sequence-numbers** global configuration command is configured.
timestamp	Date and time of the message or event. Appears only if the **service timestamps** global configuration command is configured.
facility	The facility to which the message refers.
severity	Single-digit code from 0 to 7 indicating the severity of the message.
MNEMONIC	Text string that uniquely describes the message.
description	Text string containing detailed information about the event being reported.

Using the message format and Table 4-3, you can easily interpret the following message:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, cha
```

In this message, you can see that the **service sequence-numbers** command was not configured, but the **service timestamps** command was configured. The facility is LINK, the severity is 3, and MNEMONIC is UPDOWN. The rest of the message provides a description of the event.

Configuring and Verifying Syslog

By default, Cisco routers and switches send log messages for all severity levels to the console. On some Cisco IOS versions, the device also buffers log messages by default. To enable these two settings, use the **logging console** and **logging buffered** global configuration commands, respectively.

The **show logging** command displays the default logging service settings on a Cisco router, as Example 4-5 shows.

Example 4-5 *Default Logging Service Settings*

```
R1# show logging
Syslog logging: enabled (0 messages dropped, 2 message
  0 overruns, xml disabled, filtering disabled)

No Active Message Discriminator.
No Inactive Message Discriminator.

   Console logging: level debugging, 32 messages logge
             filtering disabled
```

Technet24

```
    Monitor logging: level debugging, 0 messages logged
              filtering disabled
    Buffer logging: level debugging, 32 messages logged
              filtering disabled
    Exception Logging: size (4096 bytes)
    Count and timestamp logging messages: disabled
    Persistent logging: disabled

No active filter modules.

    Trap logging: level informational, 34 message lines
       Logging Source-Interface:        VRF Name:

Log Buffer (8192 bytes):

*Jan 2 00:00:02.527: %LICENSE-6-EULA_ACCEPT_ALL: The R
  License Agreement is accepted
*Jan 2 00:00:02.631: %IOS_LICENSE_IMAGE_APPLICATION-6-
  Module name = c1900 Next reboot level = ipbasek9 and
*Jan 2 00:00:02.851: %IOS_LICENSE_IMAGE_APPLICATION-6-
  Module name = c1900 Next reboot level = securityk9 a
*Jan 2 00:01:01.619: %IFMGR-7-NO_IFINDEX_FILE: Unable
  nvram:/ifIndex- table No such file or directory
<output omitted>
```

To configure the router to send system messages to a syslog server, complete the following three steps:

Configure the IP address of the syslog server in global configuration mode:

```
R1(config)# logging 192.168.1.3
```

Control the messages that will be sent to the syslog server with the **logging trap** *level* global configuration mode command. For example, to limit the messages to levels 4 and lower (0 to 4), use one of the following two equivalent commands:

```
R1(config)# logging trap 4
```

or:

```
R1(config)# logging trap warning
```

Optionally, configure the source interface with the **logging source-interface** *interface-type interface-number* global configuration mode command. This specifies that syslog packets contain the address of a specific interface, regardless of which interface the packet uses to exit the router. For example, to set the source interface to g0/0, use the following command:

```
R1(config)# logging source-interface g0/0
```

Example 4-6 shows the output from the **show logging** command. The default

Technet24

settings have been changed, as noted by the highlights.

Example 4-6 *Verify the Logging Service After Configuration Check*

```
R1# show logging
Syslog logging: enabled (0 messages dropped, 2 message
0 overruns, xml disabled, filtering disabled)

No Active Message Discriminator.



No Inactive Message Discriminator.
    Console logging: level debugging, 41 messages logg
                     filtering disabled
    Monitor logging: level debugging, 0 messages logge
                     filtering disabled
    Buffer logging: level debugging, 41 messages logge
                     filtering disabled
    Exception Logging: size (4096 bytes)
    Count and timestamp logging messages: disabled
    Persistent logging: disabled

No active filter modules.

   Trap logging: level warnings, 43 message lines logg
     Logging to 192.168.1.3 (udp port 514, audit disab
          link up),
```

```
            4 message lines logged,
            0 message lines rate-limited,
            0 message lines dropped-by-MD,
            xml disabled, sequence number disabled
            filtering disabled
      Logging Source-Interface:    VRF Name:
      GigabitEthernet0/0
<output omitted>
```

Network Time Protocol

As you know, routers and switches issue log messages in response to different events. For example, when an interface fails, the device creates log messages. With default settings, Cisco IOS sends these messages to the console port. But Cisco IOS can be configured also to send messages to a syslog server, where they can be stored for administration review and troubleshooting. Figure 4-3 shows a topology with a syslog server.

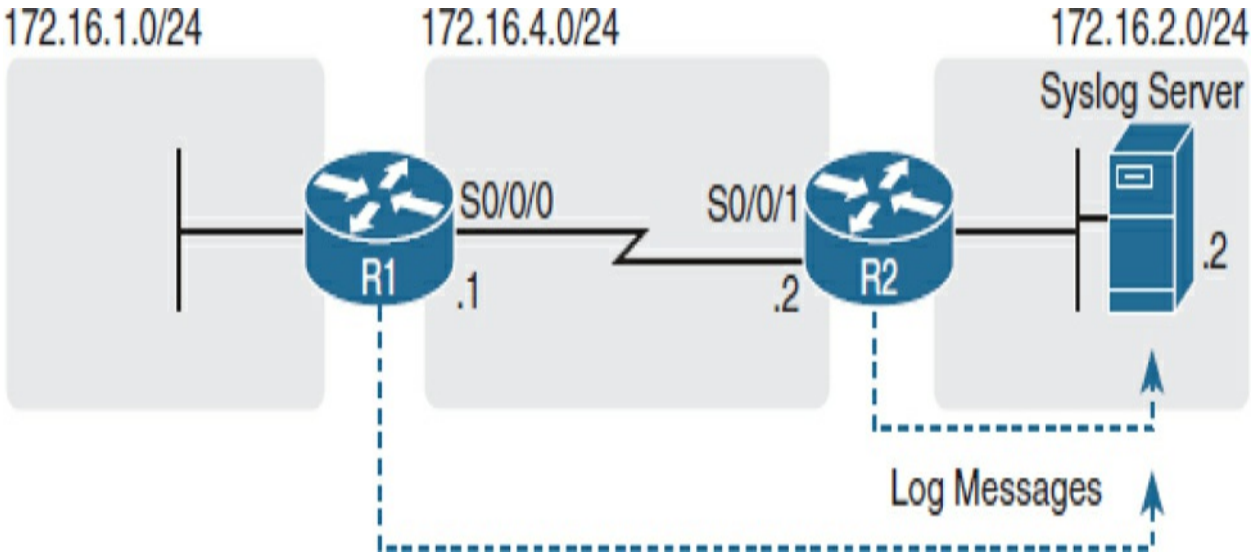

Technet24

Figure 4-3 *Sample Network with a Syslog Server*

A log message typically lists the date and time as part of the message so that a network engineer who looks back at the message knows exactly when that message occurred.

Network Time Protocol (NTP) provides a way to synchronize the time-of-day clock so that timestamps are consistent across devices, making troubleshooting easier.

To configure a router or switch to synchronize its time with an existing NTP server, use the **ntp server** command, as in Example 4-7.

Example 4-7 *Configuring and Verifying an NTP Client*

```
R1(config)# ntp server 172.16.2.2
R1(config)# ^Z
R1#
R1# show ntp status
Clock is synchronized, stratum 8, reference is 172.16.
nominal freq is 250.0000 Hz, actual freq is 250.0000 H
ntp uptime is 4700 (1/100 of seconds), resolution is 4
reference time is D42BD899.5FFCE014 (13:48:09.374 UTC
clock offset is -0.0033 msec, root delay is 1.28 msec
root dispersion is 3938.51 msec, peer dispersion is 18
loopfilter state is 'CTRL' (Normal Controlled Loop), d
system poll interval is 64, last update was 42 sec ago
```

```
R1# show ntp associations
address ref clock st when poll reach delay offset disp
*172.16.2.2 127.127.1.1 7 36 64 1 1.261 -0.001 7937.5
 * sys.peer, # selected, + candidate, - outlyer, x fal
```

The output of the **show ntp status** command gives the NTP status in the very first line; in Example 4-7, you can see that R1 is synchronized with the NTP server at 172.16.2.2. The **show ntp associations** command lists a single line of output for every other NTP device with which the router has associated.

A router or a switch can actually be the NTP server with just one command (**ntp master**) as well. In addition, NTP can use authentication so that a router or switch does not get fooled into changing its timestamp.

Cisco IOS File System and Devices

Cisco IOS devices provide a feature called the Cisco IOS Integrated File System (IFS). This system enables you to create, navigate, and manipulate directories on a Cisco device. The directories available depend on the platform.

IFS Commands

Example 4-8 shows output from the **show file systems** command.

Example 4-8 *Router Default File System Router*

Technet24

```
Router# show file systems
File Systems:
     Size(b)     Free(b)      Type   Flags    Prefixes
         -           -      opaque   rw     archive:
         -           -      opaque   rw     system:
         -           -      opaque   rw     tmpsys:
         -           -      opaque   rw     null:
         -           -      network  rw     tftp:
      262136      251594      nvram    rw     nvram:
         -           -      opaque   wo     syslog:
         -           -      opaque   rw     xmodem:
         -           -      opaque   rw     ymodem:
         -           -      network  rw     rcp:
         -           -      network  rw     http:
         -           -      network  rw     ftp:
         -           -      network  rw     scp:
         -           -      opaque   ro     tar:
         -           -      network  rw     https:
         -           -      opaque   ro     cns:
   1002143744  683163648       usbflash  rw     usbflas
*   255537152  183939072       usbflash  rw     usbflas
Router#
```

The columns show the amount of available and free memory, in bytes, and the type of file system and its permissions. Permissions include read-only (ro), write-only (wo), and read and write (rw). Although several file systems

are listed, of interest to us are the TFTP, Flash, and NVRAM file systems.

Notice that the Flash file system has an asterisk (*) preceding it, which indicates that this is the current default file system. For the 1900 router in Example 4-8, the file system is booted from USB Flash. The alias **flash:** is associated with usbflash0 so that you can use **flash:** in your commands (such as in **show flash:**).

As Example 4-9 shows, the **dir** command lists the main directory of the default file systems, whereas **show flash:** lists the entire contents of the default file system.

Example 4-9 *Default File System Directories and Files in Flash*

```
Router# dir
Directory of usbflash0:/
    1 -rw-  68831808  Jun  5  2013 18:43:02 +00:00
c1900-universalk9-mz.SPA.152-4.M3.bin
    2 -rw-      3064  Jun  5  2013 18:54:10 +00:00 cpc
    3 drw-         0  Nov  6  2013 12:40:56 +00:00 ips
   10 drw-         0  Jun  5  2013 18:54:32 +00:00 ccp
  246 -rw-      2464  Jun  5  2013 18:56:14 +00:00 hom
  247 -rw-       813  Nov  6  2013 12:44:22 +00:00 rea
  248 -rw-      2465  Nov  7  2013 11:22:20 +00:00 pre
255537152 bytes total (183939072 bytes free)
Router# show flash
-#- --length-- -----date/time------ path
```

Technet24

```
  1    68831808 Jun 5 2013 18:43:02 +00:00 c1900-univers
  2        3064 Jun 5 2013 18:54:10 +00:00 cpconfig-19xx
  3           0  Nov 6 2013 12:40:56 +00:00 ipsdir
  4         462  Nov 6 2013 13:14:52 +00:00 ipsdir/iosip
  5       14978  Nov 6 2013 12:55:56 +00:00 ipsdir/iosip
  6       38523  Nov 6 2013 12:55:58 +00:00 ipsdir/iosip
  7         304  Nov 6 2013 12:53:40 +00:00 ipsdir/iosip
  8         835  Nov 6 2013 12:53:40 +00:00 ipsdir/iosip
  9      500751  Nov 6 2013 12:56:28 +00:00 ipsdir/iosip
  10          0  Jun 5 2013 18:54:32 +00:00 ccpexp
  11          0 Jun 5 2013 18:54:32 +00:00 ccpexp/extern
  <output omitted>
  245        72 Jun 5 2013 18:56:12 +00:00 ccpexp/versio
  246      2464 Jun 5 2013 18:56:14 +00:00 home.shtml
  247       813  Nov 6 2013 12:44:22 +00:00 realm-cisco.
  248      2465  Nov 7 2013 11:22:20 +00:00 pre_autosec.

  183939072 bytes available (71598080 bytes used)

  Router#
```

Of particular interest is the first listing, which is the filename for the Cisco IOS image.

Notice that the output does not show the configuration files stored in NVRAM. To see these files, first change directories (**cd**) to the NVRAM directory (**nvram:**) and then list the contents with the **dir** command, as in

Example 4-10.

Example 4-10 *Listing Directory Contents for NVRAM*

```
Router# cd nvram:
Router# dir
Directory of nvram:/

 253  -rw-  1269       <no date> startup-config
 254  ----     5       <no date> private-config
 255  -rw-  1269       <no date> underlying-config
   1  -rw-  2945       <no date> cwmp_inventory
   4  ----     0       <no date> rf_cold_starts
   5  ----    76       <no date> persistent-data
   6  -rw-    17       <no date> ecfm_ieee_mib
   7  -rw-   559       <no date> IOS-Self-Sig#1.cer
   8  -rw-   559       <no date> IOS-Self-Sig#2.cer
   9  -rw-     0       <no date> ifIndex-table

262136 bytes total (251594 bytes free)
Router#
```

Note

You can also simply use the command **dir nvram:** to list the contents of the nvram: directory.

Technet24

The file you are most interested in as CCNA exam candidates is the startup-config file.

URL Prefixes for Specifying File Locations

File locations are specified in Cisco IFS using the URL convention shown in Figure 4-4.

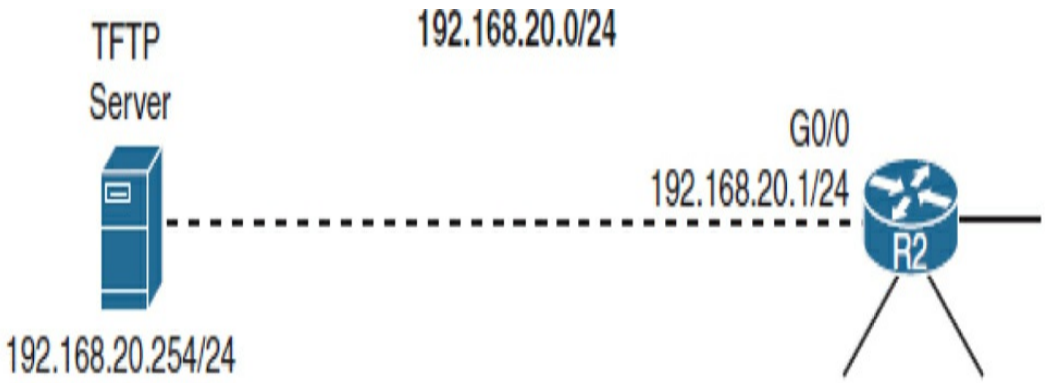

Prefix	URL Path
tftp:	[[[//location]/directory]/filename]
tftp://192.168.20.254/configs/backup-config	

Figure 4-4 *Using a URL to Specify the TFTP Location*

In Figure 4-4, the parts of the URL tftp://192.168.20.254/configs/backup-config can be dissected as follows:

• tftp: is the prefix specifying the protocol.

• Everything after the double slash (//) defines the file location.

• 192.168.20.254 is the location of the TFTP server.

• configs is the master directory on the TFTP server.

• backup-config is a sample filename.

The TFTP URL in Figure 4-4 is an example of a remote URL. Examples of URLs for accessing the local Cisco IFS include the following:

• flash:configs/backup-config

• system:running-config (which accesses RAM)

• nvram:startup-config

Commands for Managing Configuration Files

Knowing the URL structure is important because you use it when copying configuration files from one location to another. The Cisco IOS Software **copy** command enables you to move configuration files from one component or device to another, such as RAM, NVRAM, or a TFTP server. Figure 4-5 shows the command syntax.

Technet24

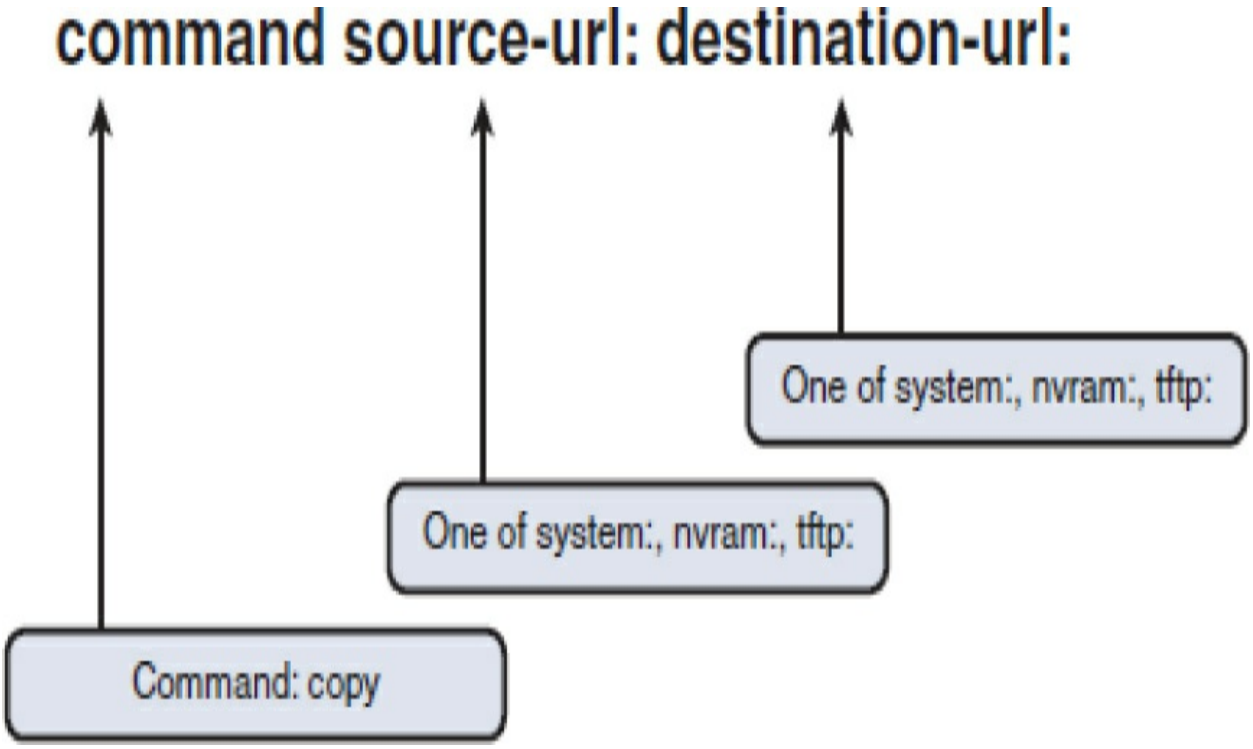

Figure 4-5 copy Command Syntax

The source URL is where you are copying from. The destination URL is where you are copying to. For example, you are already familiar with the abbreviated command **copy run start**; however, in its most verbose form, this command specifies the file locations:

```
Router# copy system:running-config nvram:startup-confi
```

The command says to "copy the running configuration from the system's RAM to NVRAM and save it with the filename startup-config."

Example 4-11 demonstrates how to copy the current running configuration to a new local directory named configs that you create.

Example 4-11 *Copying Files to a Local Directory*

```
Router# mkdir configs
Create directory filename [configs]? <Enter>
Created dir usbflash0:/configs
Router# copy system:running-config configs/backup-conf
Destination filename [/configs/backup-config]? <Enter>
1269 bytes copied in 0.648 secs (1958 bytes/sec)
```

You can also copy from RAM to TFTP:

```
Router# copy system:running-config tftp:
```

or simply:

```
Router# copy run tftp
```

You can also copy from TFTP to RAM:

```
Router# copy tftp: system:running-config
```

or simply:

```
Router# copy tftp run
```

Technet24

You can also copy from TFTP to the startup configuration file:

```
Router# copy tftp: nvram:startup-config
```

or simply:

```
Router# copy tftp nvram
```

The **copy** commands using TFTP require more configuration (covered in the next section) after you enter them to carry out the instruction.

Managing Cisco IOS Images

As a network grows, storing Cisco IOS Software images and configuration files on the central TFTP server gives you control over the number and revision level of Cisco IOS images and configuration files that must be maintained. Use the **show version** command to verify the Cisco IOS image currently running on the device. Figure 4-6 shows a sample topology with a TFTP server.

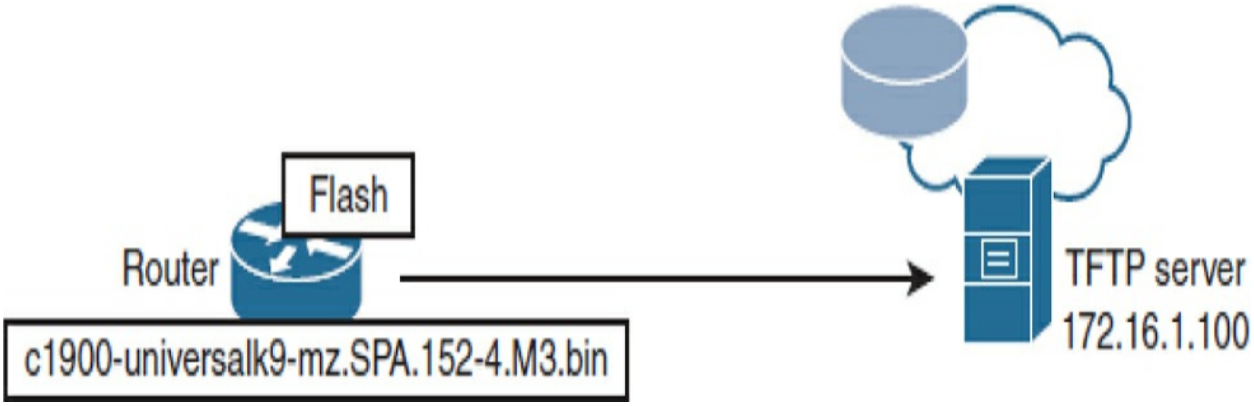

Figure 4-6 *TFTP Topology*

Backing Up a Cisco IOS Image

Make sure that a TFTP server is configured and running on the network. Then follow these steps to copy a Cisco IOS Software image from Flash memory to the network TFTP server:

Ping the TFTP server to ensure that you have access to it:

```
R1# ping 172.16.1.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.100, timeou
is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip
min/avg/max = 31/31/32 ms

R1#
```

Technet24

Copy the current system image file from the router to the network TFTP server by using the **copy flash: tftp:** command in privileged EXEC mode. You then are prompted to enter the IP address of the remote host and the name of the source and destination system image files:

```
R1# copy flash: tftp:
Source filename []?
c1900-universalk9-mz.SPA.152-4.M3.bin
Address or name of remote host []? 172.16.1.100
Destination filename
[c1900-universalk9-mz.SPA.152-4.M3.bin]? <Enter>
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
<Output omitted>
68831808 bytes copied in 113.061 secs (608802 bytes/se
R1#
```

During the copy process, exclamation points (!) indicate the progress. Each exclamation point signifies that one User Datagram Protocol (UDP) segment has successfully been transferred.

Restoring a Cisco IOS Image

You can use the **dir** command or **show flash:** command to verify that the router has sufficient disk space to accommodate the new Cisco IOS Software image, as in Example 4-12.

Example 4-12 *Output from the* **dir** *Command*

```
R1# dir
Directory of usbflash0:/

    1  -rw- 68831808 Jun    5 2013  18:43:02 +00:00 c1
  mz.SPA.152-4.M3.bin
    2  -rw-  3064    Jun    5 2013  18:54:10 +00:00 cp
    3  drw-     0    Nov    6 2013  12:40:56 +00:00 ip
   10  drw-     0    Jun    5 2013  18:54:32 +00:00 co
  246  -rw-  2464    Jun    5 2013  18:56:14 +00:00 ho
  247  -rw-   813    Nov    6 2013  12:44:22 +00:00 re
  248  -rw-  2465    Nov    7 2013  11:22:20 +00:00 pr
  249  drw-     0    Feb    8 2014  19:49:08 +00:00 co

255537152 bytes total (183926784 bytes free)
R1#
```

The **show flash:** and **dir** commands help you determine the following:

• The total amount of Flash memory on the router

• The amount of Flash memory available

• The names of all the files stored in the Flash memory and the amount of Flash occupied

Technet24

Example 4-13 shows the commands necessary to copy an image stored on the TFTP server to Flash.

Example 4-13 *Upgrading the Cisco IOS Image from a TFTP Server*

```
R1# copy tftp flash
Address or name of remote host []? 172.16.1.100
Source filename []? c1900-universalk9-mz.SPA.152-4.M3.
Destination filename [c1900-universalk9-mz.SPA.152-4.M

Loading c1900-universalk9-mz.SPA.152-4.M3.bin from 172
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!
[OK - 68831808 bytes]

68831808 bytes copied in 9.656 secs (7128397 bytes/sec
R1#
```

The command asks for the IP address of the TFTP server and then the Cisco IOS image filename stored on the TFTP server that you want to copy over. When asked for the destination filename, you can change it, but doing so is not recommended.

Password Recovery

Password recovery procedures for any Cisco router or switch are readily available online. For example, search for “1941 password recovery,” and you will quickly find the procedures you need to follow to reset the password. This is why physical security is a must for all networking devices. Routers and switches should be behind locked doors.

The following steps are common to most Cisco routers. Switches have a similar process.

Use the power switch to turn off the router and then turn the router back on.

Press the break key specified by your terminal software within 60 seconds of powerup to access the ROMmon prompt. For Tera Term, for example, use the key combination **Alt+b**.

Enter **confreg 0x2142** at the ROMmon prompt. This causes the router to bypass the startup configuration, where the forgotten password is stored.

Enter **reset** at the prompt. The router reboots, but it ignores the saved configuration. However, the file still exists in NVRAM.

Press **Ctrl+C** to skip the initial setup procedure.

Enter **enable** at the Router> prompt to enter privileged EXEC mode, where you should be able to see the Router# prompt.

Technet24

Enter **copy startup-config running-config** to copy the backup NVRAM config file into memory.

Enter **configure terminal**.

Enter the **enable secret** *password* command to change the enable secret password.

Issue the **no shutdown** command on every interface that you want to activate.

From global configuration mode, enter **config-register 0x2102** to restore the original configuration registry setting.

Press **Ctrl+Z** or enter **end** to leave configuration mode.

Enter **copy running-config startup-config** to commit the changes. You can issue the **show ip interface brief** command to confirm that your interface configuration is correct. Every interface that you want to use should display up and up.

You have now completed password recovery. You can enter the **show version** command to confirm that the router will use the configured config register setting on the next reboot.

Study Resources

For today’s exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	10
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	13
	17
Portable Command Guide	19

Technet24

Day 3. Cloud, Virtualization, and SDN

CCNA 200-301 Exam Topics

• 1.2 Describe characteristics of network topology architectures

• 1.12 Explain virtualization fundamentals (server virtualization, containers, and VRFs)

• 6.2 Compare traditional networks with controller-based networking

• 6.3 Describe controller-based, software defined architecture (overlay, underlay, and fabric)

Key Topics

Cloud computing and virtualization are increasingly important services that networks offer to end users. Enterprise networks can use a number of different cloud offerings to enhance productivity and lower costs.

Network programmability through software-defined networking (SDN) is becoming integral to enterprise networks as network administrators can quickly and efficiently manage thousands of networking devices.

Today we review the technologies behind cloud computing, virtualization, and SDN.

Cloud Computing

Cloud computing involves large numbers of computers connected through a network that can be physically located anywhere. Cloud computing provides the following benefits:

• Enables access to organizational data anywhere and at any time

• Streamlines IT operations by making it possible to subscribe to only needed services

• Eliminates or reduces the need for onsite IT equipment, maintenance, and management

• Reduces costs for equipment, energy, physical plant requirements, and personnel training needs

• Enables rapid responses to increasing data volume requirements

Cloud providers rely heavily on virtualization to enable the solutions they offer to clients.

Server Virtualization

Historically, organizations bought multiple hardware servers, and the server administrator installed one or more network applications on the server, such as an email server or a file server (see Figure 3-1).

Technet24

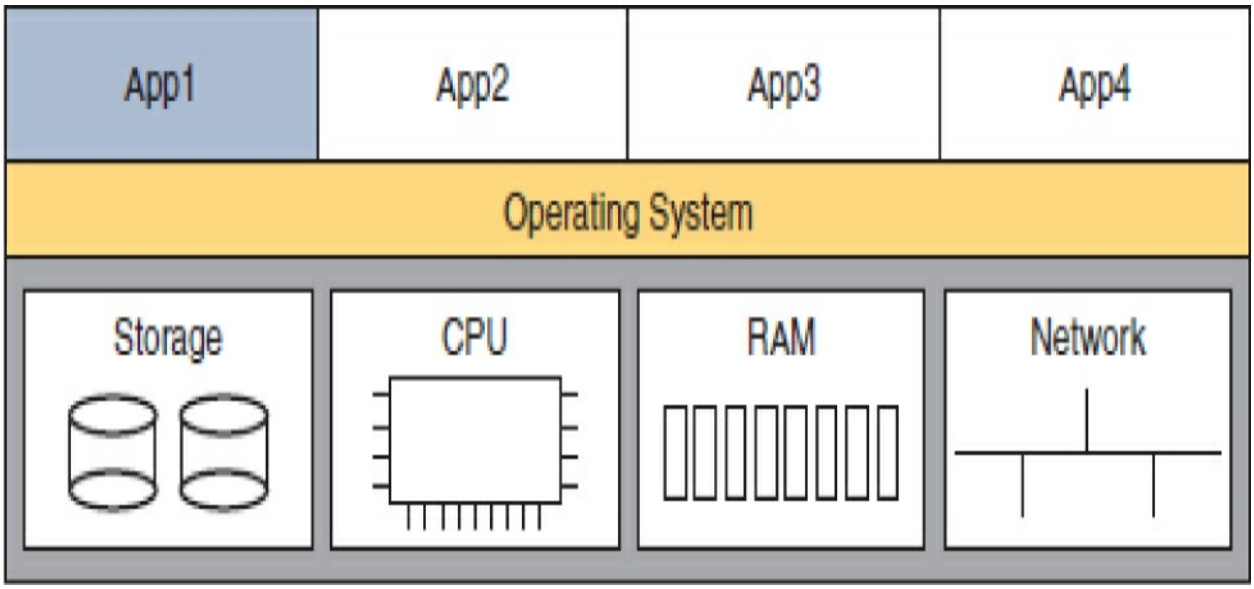

Figure 3-1 *Dedicated Server with One OS*

Each of these servers had its own CPU, memory, NIC, and disk space. However, this model faces several challenges:

• If a component fails, the service is unavailable until the component is repaired or replaced.

• Servers sometimes sit idle for long periods of time, waiting for clients to use them.

• Servers take up space and waste energy.

Server virtualization takes advantage of idle resources and consolidates the number of required servers. Virtualization separates the operating system (OS) from the hardware. This also makes it possible for multiple OSs to exist on a single hardware platform. Each instance of an OS is called a *virtual*

machine (*VM*).

A server with multiple VMs uses a hypervisor to manage access to the server's physical resources. The hypervisor sits between the VMs and the hardware, as in Figure 3-2.

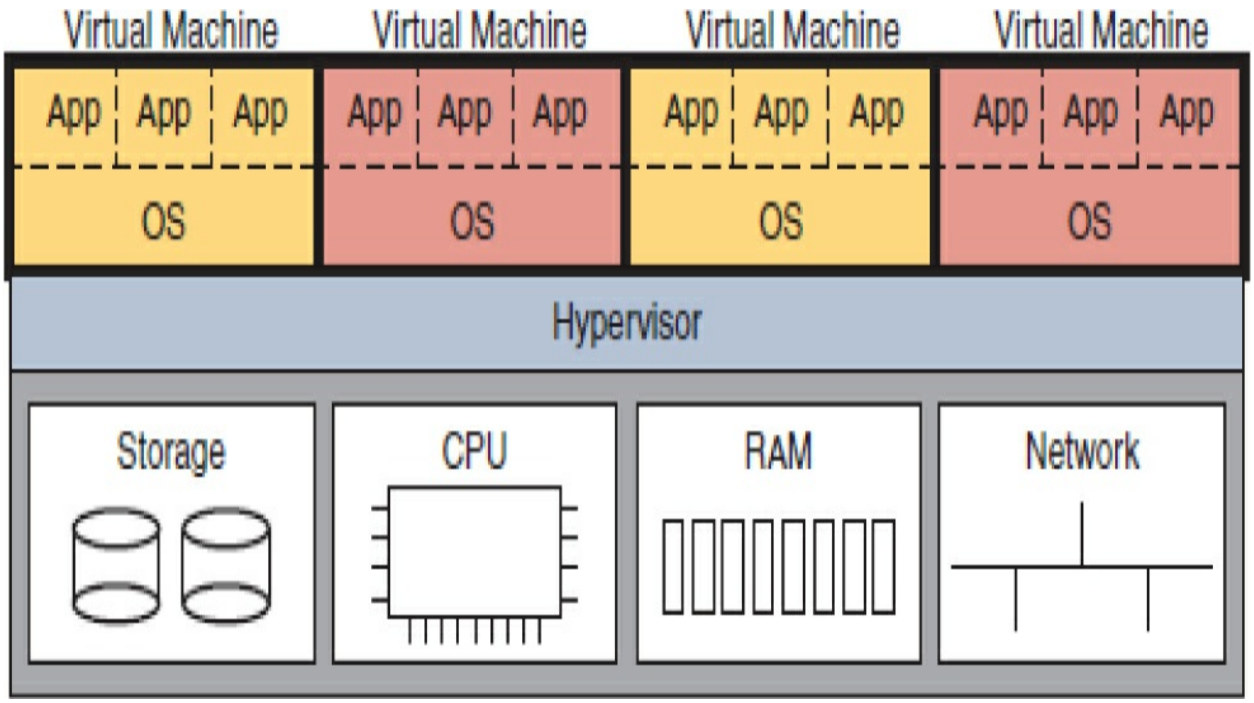

Figure 3-2 *Hypervisor Managing Four VMs*

Another method for managing a set of VMs on a server, especially in a data center environment, is to use a virtual switch that connects the VMs to physical NICs, as in Figure 3-3. An external controller (not shown) manages the server hardware.

Technet24

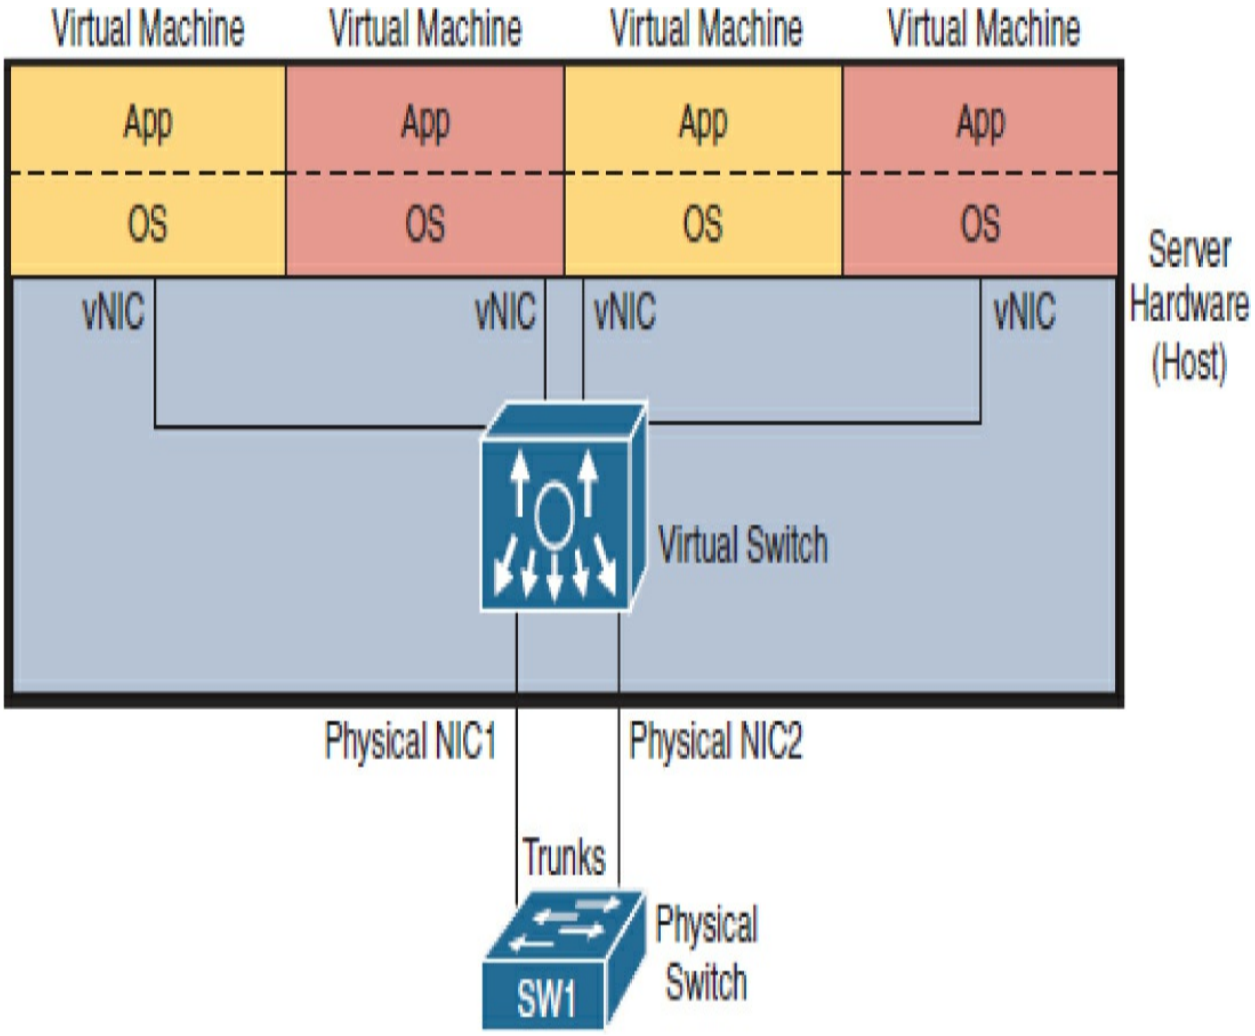

Figure 3-3 *Virtual Switch and VMs*

In a data center, multiple servers are placed in a rack. The two physical NICs in Figure 3-3 are attached to two redundant top of rack (ToR) switches. Racks are lined up in rows and managed by two redundant end of row (EoR) switches. Figure 3-4 shows this physical layout of a traditional data center.

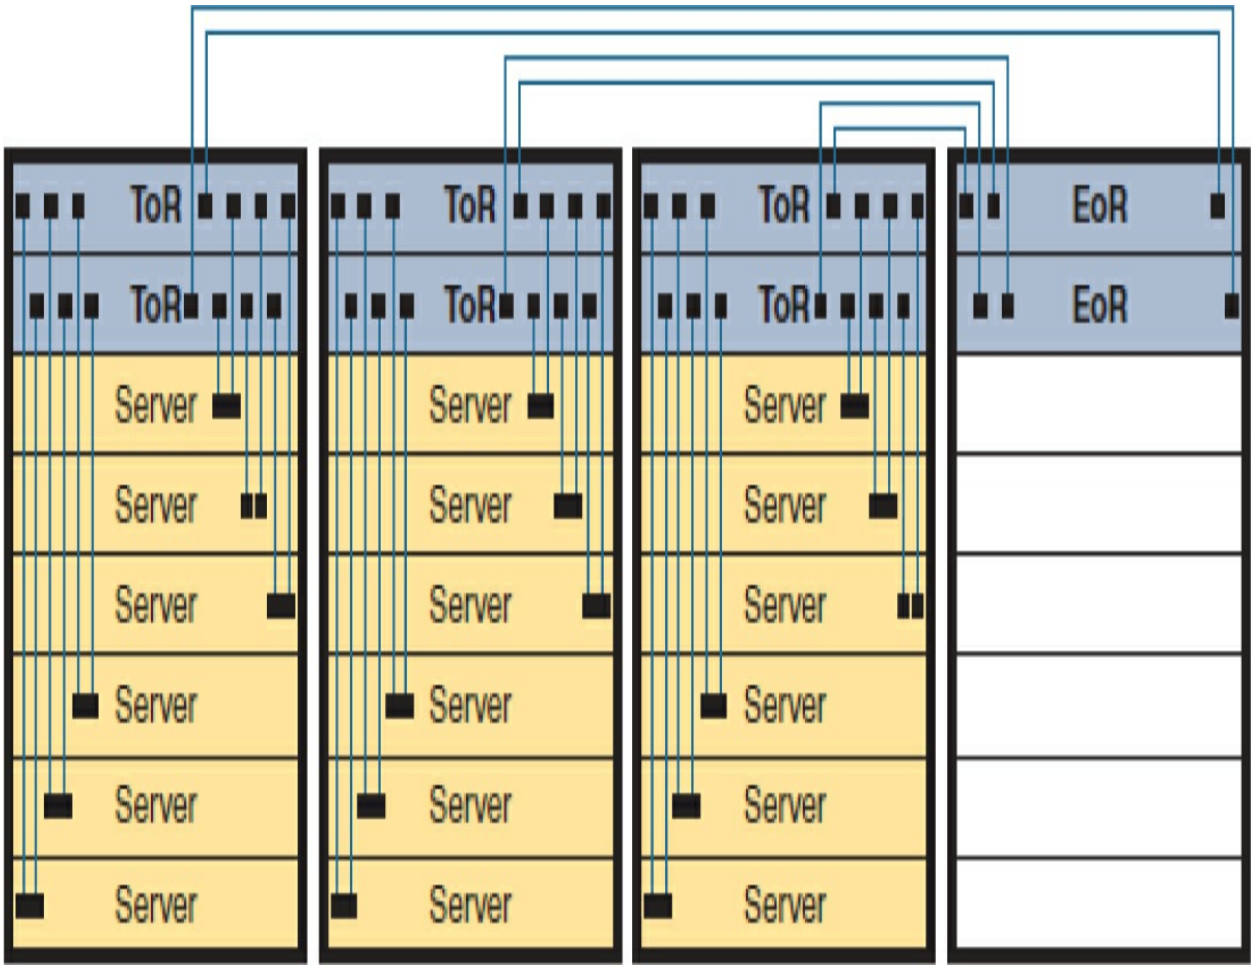

Figure 3-4 *Traditional Data Center Physical Topology*

Cloud Computing Services

To understand the value of cloud computing, consider the effort it takes to manage VMs in a traditional data center. The workflow follows:

A customer requests a VM or a new set of VMs.

The data center engineer configures virtualization software.

Virtualization software creates VMs.

Technet24

Although this process works, it does not have the characteristics of a cloud computing service, as defined by the U.S. National Institute of Standards and Technology (NIST):

• **On-demand self-service:** The user can order, modify, and end service without human interaction.

• **Broad network access:** The service can be accessed from a variety of devices across any network.

• **Resource pooling:** The provider has a pool of resources that can be dynamically allocated to users. The user typically requires no awareness of the physical location of the resources.

• **Rapid elasticity:** To the user, the resource pool appears to be unlimited; it can expand and contract as needed.

• **Measured service:** The provider can measure the usage and then report that usage to the consumer, for both transparency and billing.

Cloud providers can offer a variety of services to meet the needs of customers, including these:

• **Software as a service (SaaS):** The cloud provider is responsible for access to services that are delivered over the Internet, such as email, communication, and Office 365. Users only need to provide their data.

• **Platform as a service (PaaS):** The cloud provider is responsible for access to the development tools and services used to deliver the applications. Customers can customize the virtualized hardware.

• **Infrastructure as a service (IaaS):** The cloud provider is responsible for access to the network equipment, virtualized network services, and network infrastructure support.

Four primary cloud models exist:

• **Public clouds:** Cloud-based applications and services offered in a public cloud are made available to the general population. The public cloud uses the Internet to provide services.

• **Private clouds:** Cloud-based applications and services offered in a private cloud are intended for a specific organization or entity, such as the government. A private cloud uses the organization’s private network.

• **Hybrid clouds:** A hybrid cloud is made up of two or more clouds (for example, part private and part public). Each part remains a distinct object, and the two parts are connected int a single architecture.

• **Community clouds:** A community cloud is created for exclusive use by a specific community. The differences between public clouds and community clouds are the functional needs that have been customized for the community.

Virtual Network Infrastructure

Technet24

A virtual network infrastructure consists of a collection of virtual network functions (VNFs), including virtual switches, virtual server load balancers (SLBs), virtual routers, and virtual firewalls (see Figure 3-5).

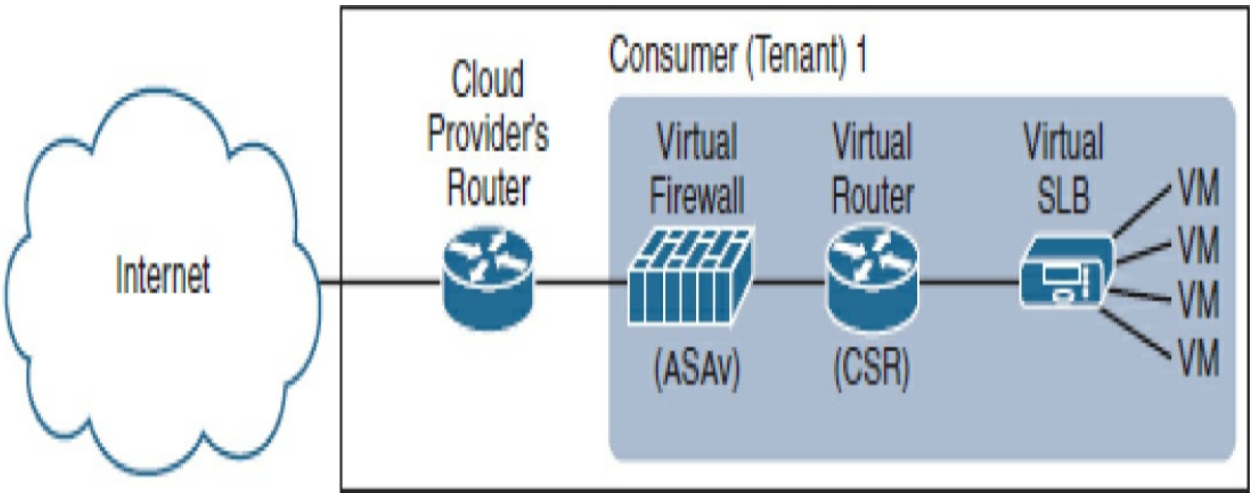

Figure 3-5 *Virtual Network Installed by a Cloud Provider's Tenant*

A cloud provider could offer to a customer (tenant) an IaaS solution that includes a virtual SLB. The customer could then install VNFs, such as a virtual version of the Cisco Adaptive Security Appliance (ASAv) and a Cisco Cloud Services Router (CSR), to manage network traffic.

Software-Defined Networking

Network programmability refers to the trend that began in the 2010s called *software-defined networking* (*SDN*). At its core, SDN decouples the data, control, and management planes from the physical device; virtualizes them; and defines the networking functions in software. This creates an architecture that can be more efficiently and effectively managed programmatically. Today, you're most likely to hear the terms *software-defined architecture* and

controller-based networking than SDN.

Data, Control, and Management Planes

A traditional networking device contains two planes. The data plane is responsible for forwarding data as quickly as possible. To do so, it relies on tables built by the control plane. Actions taken by the data plane include the following:

• Layer 2 and Layer 3 de-encapsulation/encapsulation

• Addition or removal of an 802.1Q trunking header

• MAC address table lookups

• IP routing table lookups

• Data encryption and addition of a new IP header (as in VPNs)

• Change to the source or destination IP address (with NAT)

• Message discard due to a filter (such as an ACL or port security)

The control plane does all the calculations for populating tables used by the data plane and manages control messages between other networking devices. Figure 3-6 provides an example of OSPF operating on the control plane while the data plane is responsible for forwarding packets using the best route.

Technet24

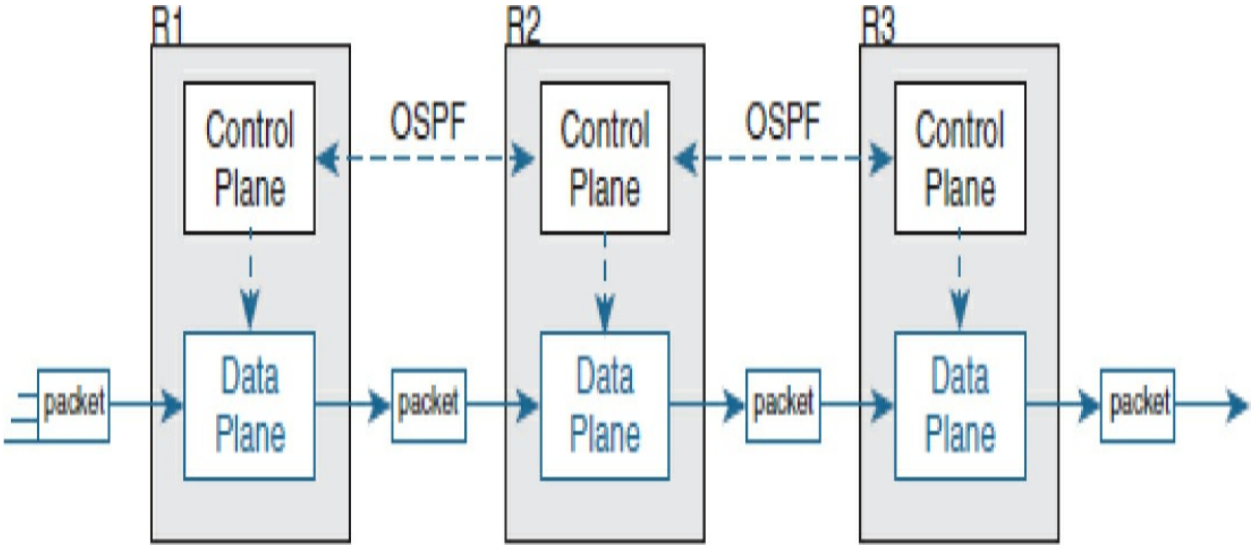

Figure 3-6 *Control and Data Plane Example*

The following are the most common control plane protocols:

• Routing protocols (OSPF, EIGRP, RIP, BGP)

• IPv4 ARP

• IPv6 NDP

• Switch MAC learning

• STP

The management plane is responsible for all functions that are not directly related to controlling the data plane. Management protocols, such as the ones in Figure 3-7, are examples of management plane functions.

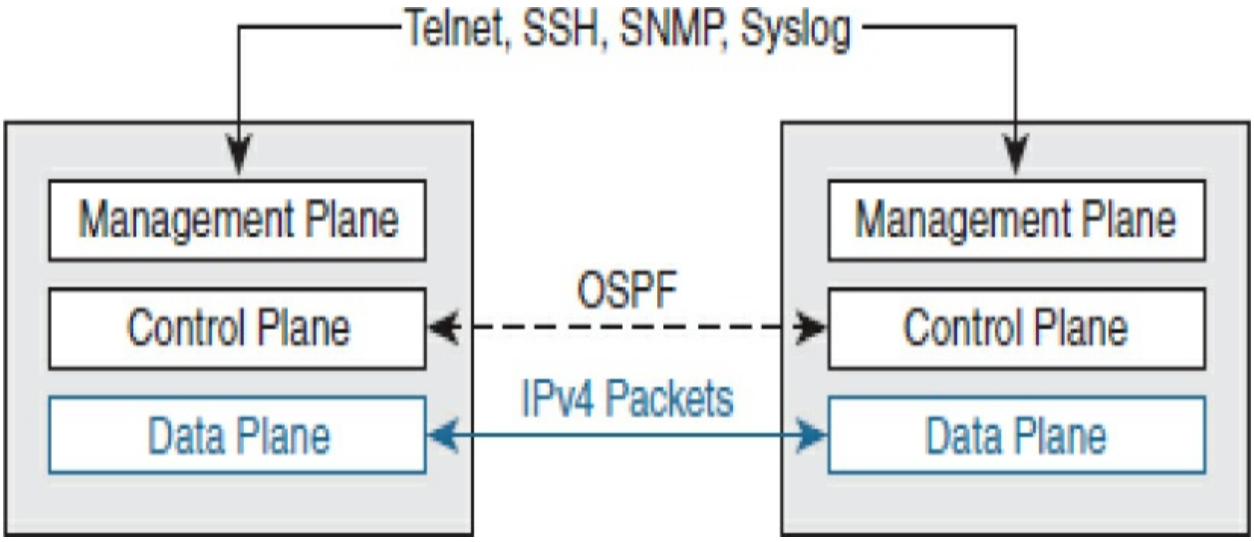

Figure 3-7 *Management Plane Example*

Controllers

Traditionally, the control plane has been part of the device OS and has been distributed across every device. That means every device must spend some resources calculating and maintaining Layer 2 and Layer 3 data structures (ARP tables, routing tables, and so on). When viewed as a whole, the network's control plane is distributed across all the networking devices.

In SDN, the functions of the control plane can be completely removed from the physical networking devices and placed in a centralized application called a *controller*. This frees up the devices to focus on data plane tasks.

The controller sits at the top of a network topology diagram, and the connections to the networking devices are called the *southbound interface* (*SBI*) (see Figure 3-8).

Technet24

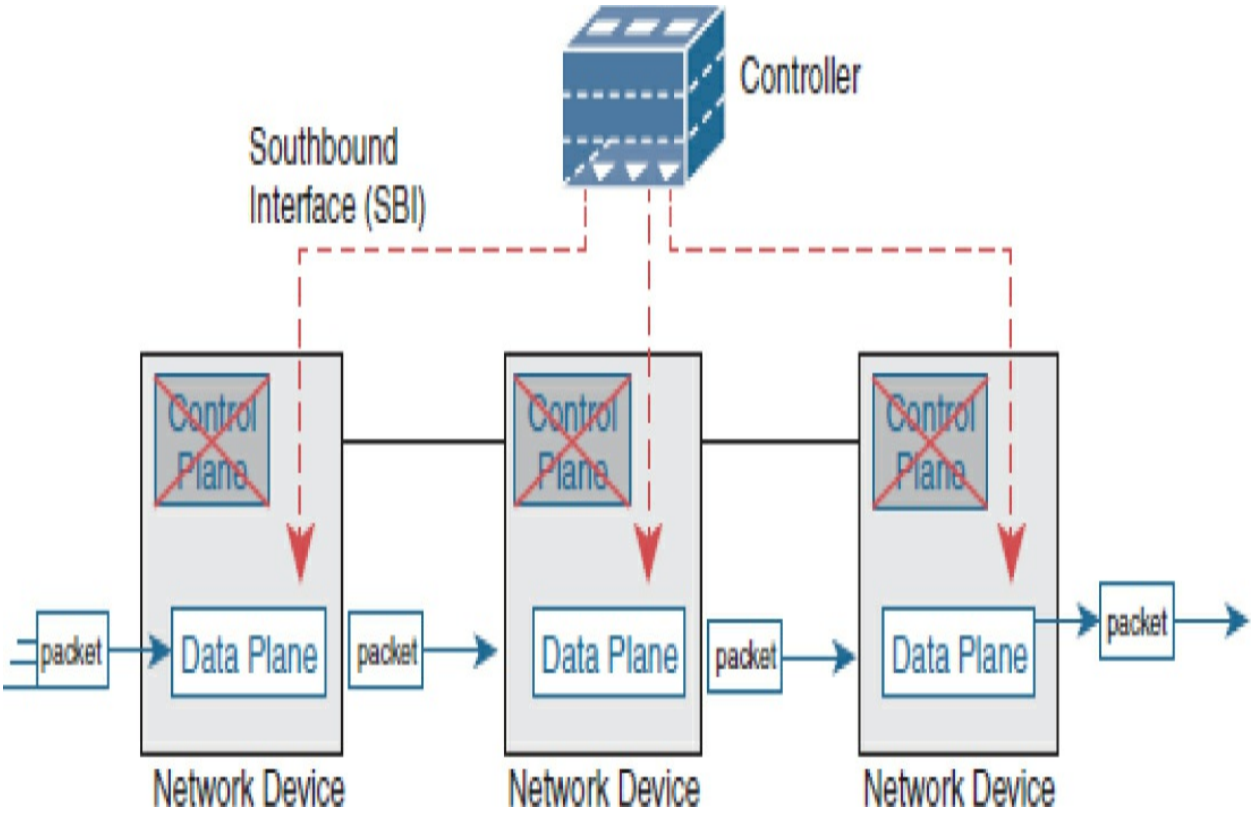

Figure 3-8 *Centralized Controller and Distributed Data Plane*

A *northbound interface* (*NBI*) also exists between the SDN controller and the applications that are installed on the controller. These applications are what enable network programmability. Applications use the NBI to pull information and use the controller's capabilities to program network flows via the network device's SBI. The controller uses APIs to enable communication between applications and itself. For example, in Figure 3-9 an application is using the controller's Java-based native API to learn about the network and program specific match/action logic into the network devices' forwarding tables.

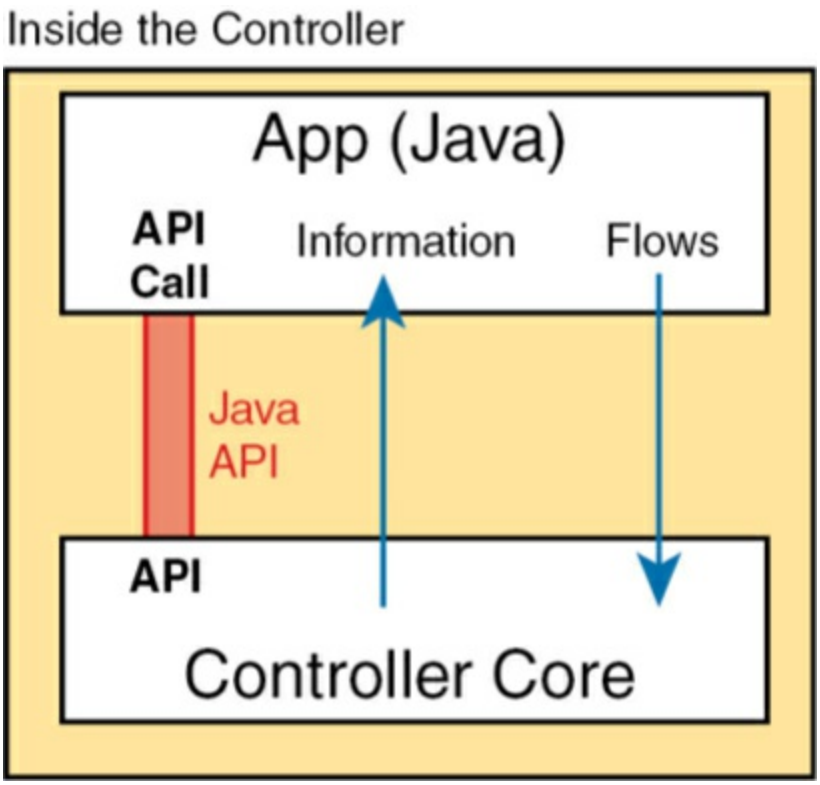

Figure 3-9 *Northbound Interface Inside the Controller*

Network Programmability and SDN Examples

Two network programmability and SDN solutions stand out due to their wide range of comparison points: OpenDaylight Controller and Cisco Application Centric Infrastructure (ACI).

OpenDaylight and OpenFlow

OpenDaylight and OpenFlow represent are what is called Open SDN. Open SDN is developed by the Open Networking Foundation (ONF), a consortium that establishes SDN standards, including protocols for SBIs and NBIs. OpenFlow defines a controller and a standardized switch abstraction, allowing switches to act as Layer 2 or Layer 3 devices with flexible

Technet24

configurations. The Open SDN model centralizes control plane functions, and the network is managed by the controller and its applications via NBIs. OpenFlow acts as the SBI protocol, necessitating OpenFlow-compatible switches.

The OpenDaylight Controller is a successful open-source SDN controller that emerged from the consolidation of various research-oriented SDN controllers. It follows open-source principles, much like Linux, with multiple vendors collaborating on a common platform. This collaboration has resulted in a versatile controller with various SBIs, such as OpenFlow, NETCONF, PCEP, BGP-LS, and OVSDB, as shown in Figure 3-10. Vendors can use OpenDaylight to create commercial products by adding proprietary features.

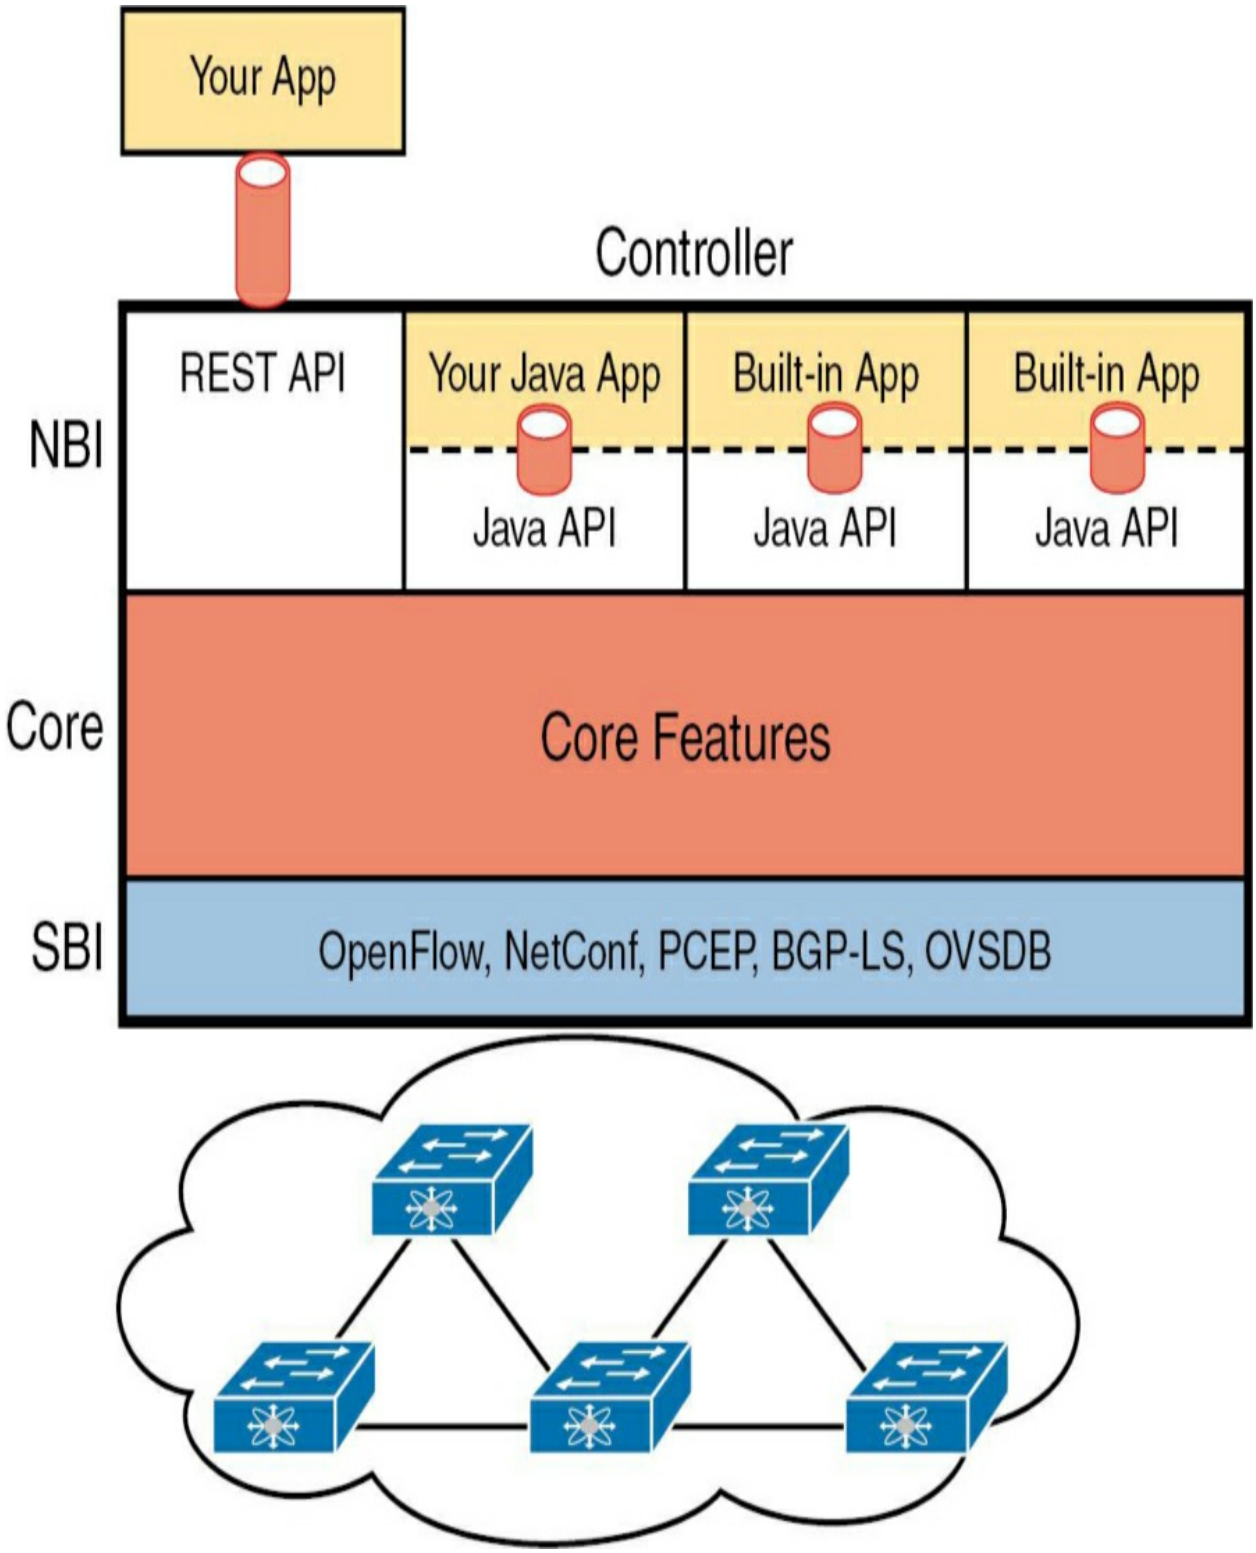

Figure 3-10 *ONF OpenDaylight Controller*

Technet24

Cisco previously offered a commercial controller that followed the OpenDaylight model: the Cisco Open SDN Controller (OSC). However, Cisco transitioned to intent-based networking (IBN), moving away from OpenFlow-based SDN.

Cisco Application Centric Infrastructure (ACI)

Cisco ACI reimagines data center networking with a focus on application needs and uses a spine-and-leaf topology. This design ensures redundancy and capacity for dynamic data center operations. ACI employs intent-based networking through Application Policy Infrastructure Controller (APIC), which manages policies and intent for endpoint communication. APIC abstracts network configuration, enabling application-centric operations without manual switch interface configuration.

Cisco ACI uses a spine-and-leaf design. As shown in Figure 3-11, each leaf switch connects to every spine switch, and endpoints connect to leaves only. This consistent design supports redundancy and scalability. Spine-and-leaf switches are connected using the following design guidelines:

• Each leaf switch must connect to every spine switch.

• Each spine switch must connect to every leaf switch.

• Leaf switches cannot connect to each other.

• Spine switches cannot connect to each other.

• Endpoints connect only to the leaf switches.

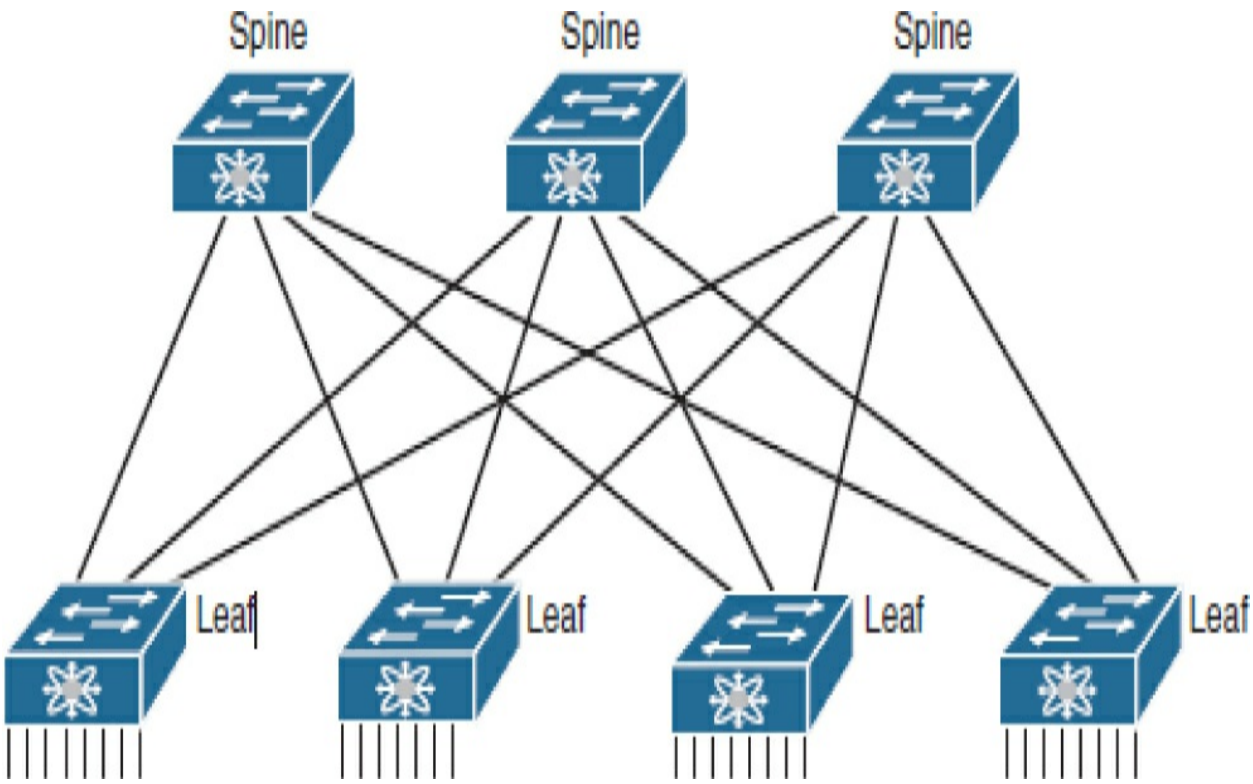

Figure 3-11 *Spine-and-Leaf Network Design*

ACI's intent-based networking model focuses on endpoints and policies managed by APIC. For example, a typical enterprise web application might consist of web servers, application servers, and database servers, each forming an endpoint group (EPG). APIC defines policies determining communication between these EPGs. This model abstracts traditional network configurations like VLANs and ACLs, shifting to an application-centric approach.

The Cisco ACI uses APIC, a centralized controller (see Figure 3-12), to create application policies for the data center infrastructure. It translates the

Technet24

intent (EPGs, policies, and so on), shifting the operational model away from configuring VLANs, trunks, EtherChannels, and ACLs.

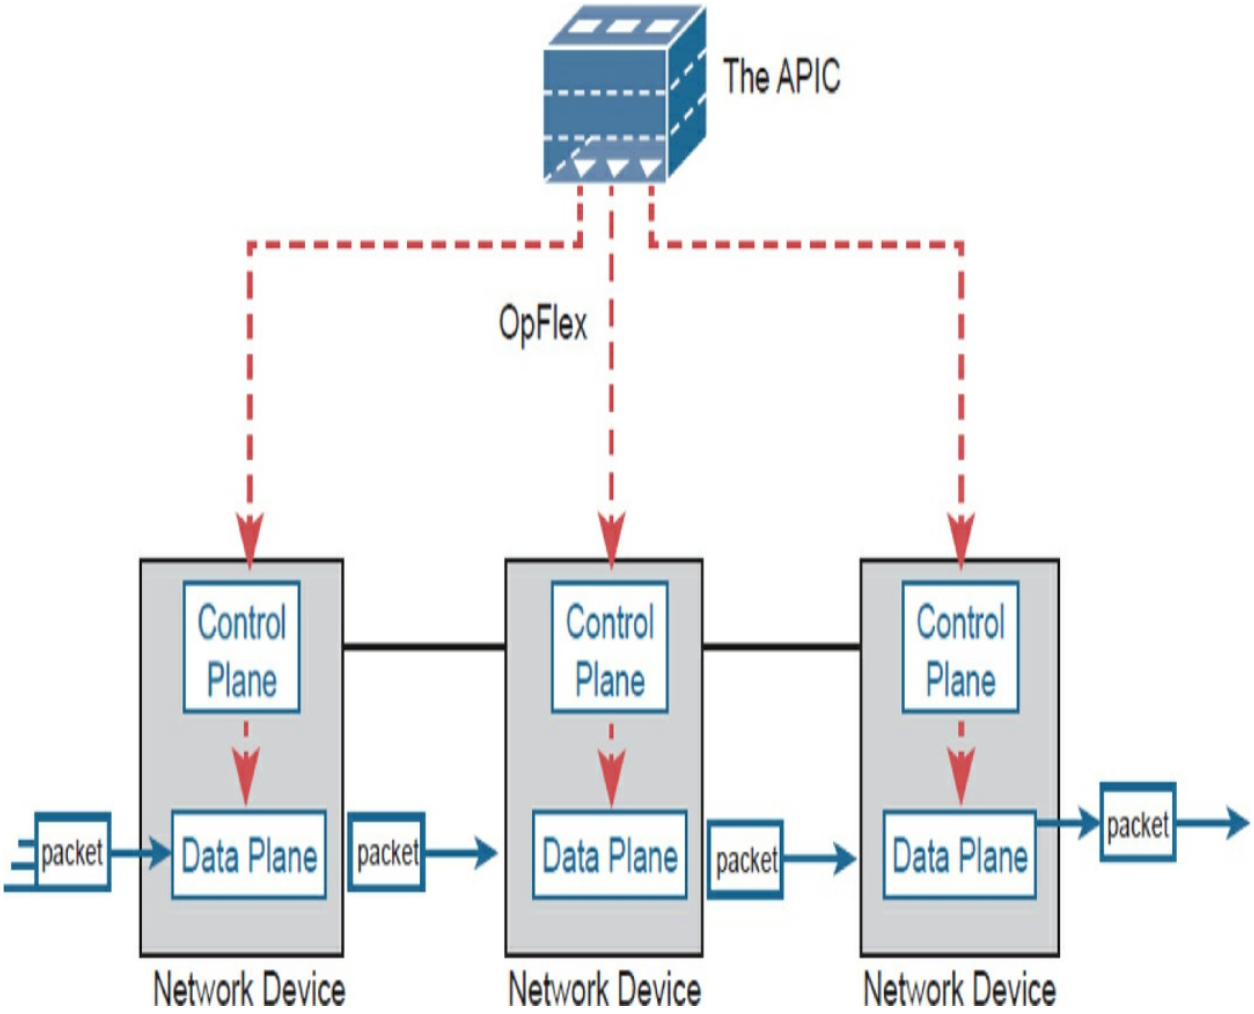

Figure 3-12 *APIC Pushes Intent to the Switch Control Planes*

APIC has a convenient GUI, but its real power lies in network programmability. Virtualization software, cloud or automation software, and scripts written by network engineers can define endpoint groups and policies for APIC. All these tools access the ACI system by interfacing with APIC, as shown in Figure 3-13. This eliminates the need for network engineers to

connect to each individual switch and configure CLI commands.

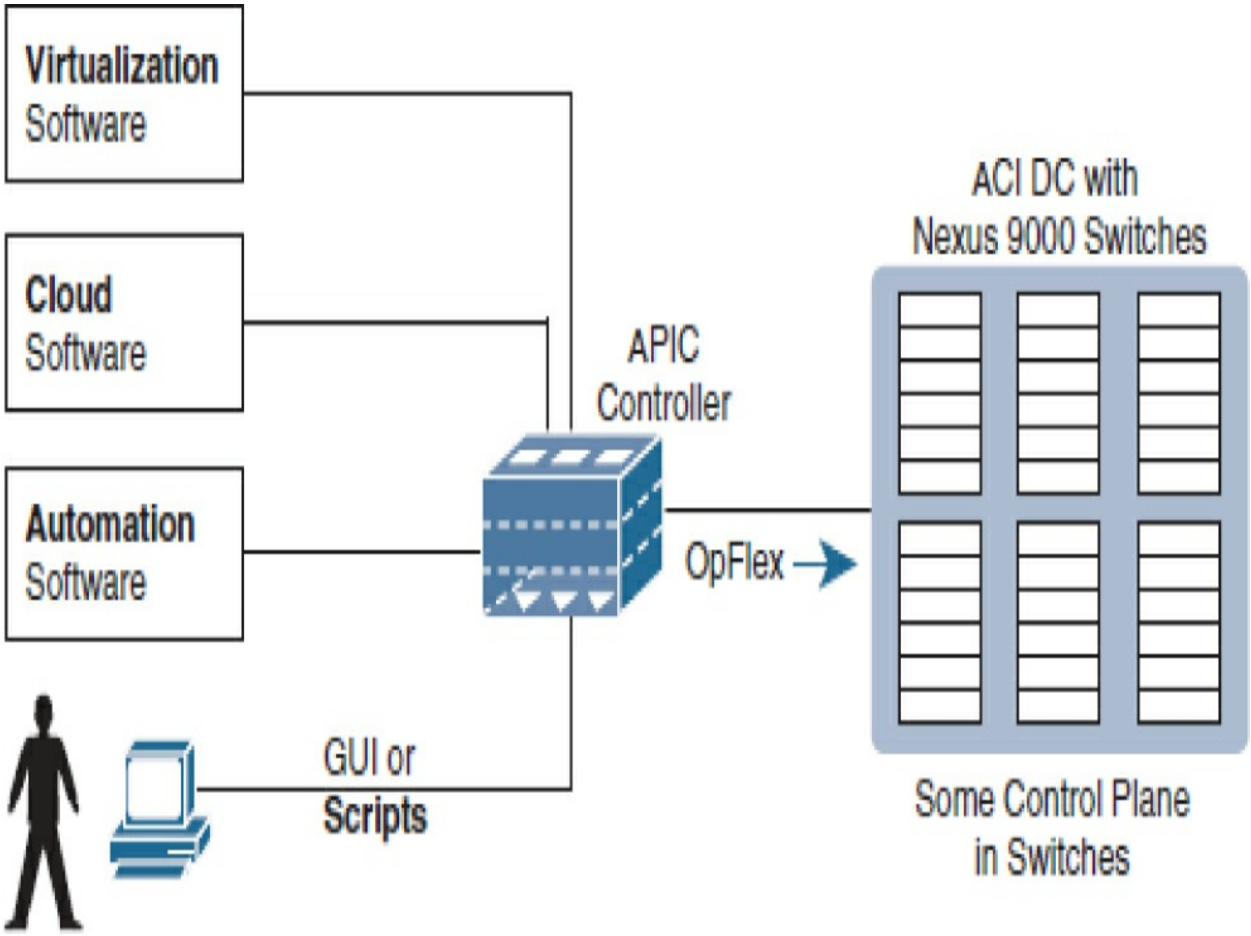

Figure 3-13 *Cisco ACI for Data Centers*

The OpenFlow and ACI SDN architectures differ in the extent to which control plane work is centralized. Table 3-1 compares them.

Table 3-1 *Comparing OpenFlow and ACI*

Technet24

Criterion	OpenFlow	ACI
Changes how the device control plane works versus traditional networking	Yes	Yes
Creates a centralized point from which humans and automation control the network	Yes	Yes
How much the architecture centralizes the control plane	Mostly	Partially
SBIs used	OpenFlow	OpFlex
Controllers mentioned in this chapter	OpenDaylight	APIC
Organization that is the primary definer/owner	ONF	Cisco

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	13
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	20—21

Technet24

Day 2. Cisco SD-Access, Ansible, and Terraform

CCNA 200-301 Exam Topics

• 6.1 Explain how automation impacts network management

• 6.3 Describe controller-based, software defined architecture (overlay, underlay, and fabric)

• 6.6 Recognize the capabilities of configuration management mechanisms such as Ansible and Terraform

Key Topics

Today we review Cisco SD-Access, including the architecture features of fabric, underlay, and overlay. Then we examine the capabilities of configuration management tools like Ansible and Terraform.

Cisco SD-Access

SDA uses a controller and application programming interfaces (APIs) to communicate via southbound interfaces (SBIs) with the network infrastructure, as shown in Figure 2-1. Cisco Catalyst Center is an example of a controller. SBIs include Telnet/SSH, SNMP, NETCONF, and RESTCONF.

Cisco SD-Access redefines campus LANs using a software-defined architectural model with Cisco Catalyst Center as the controller.

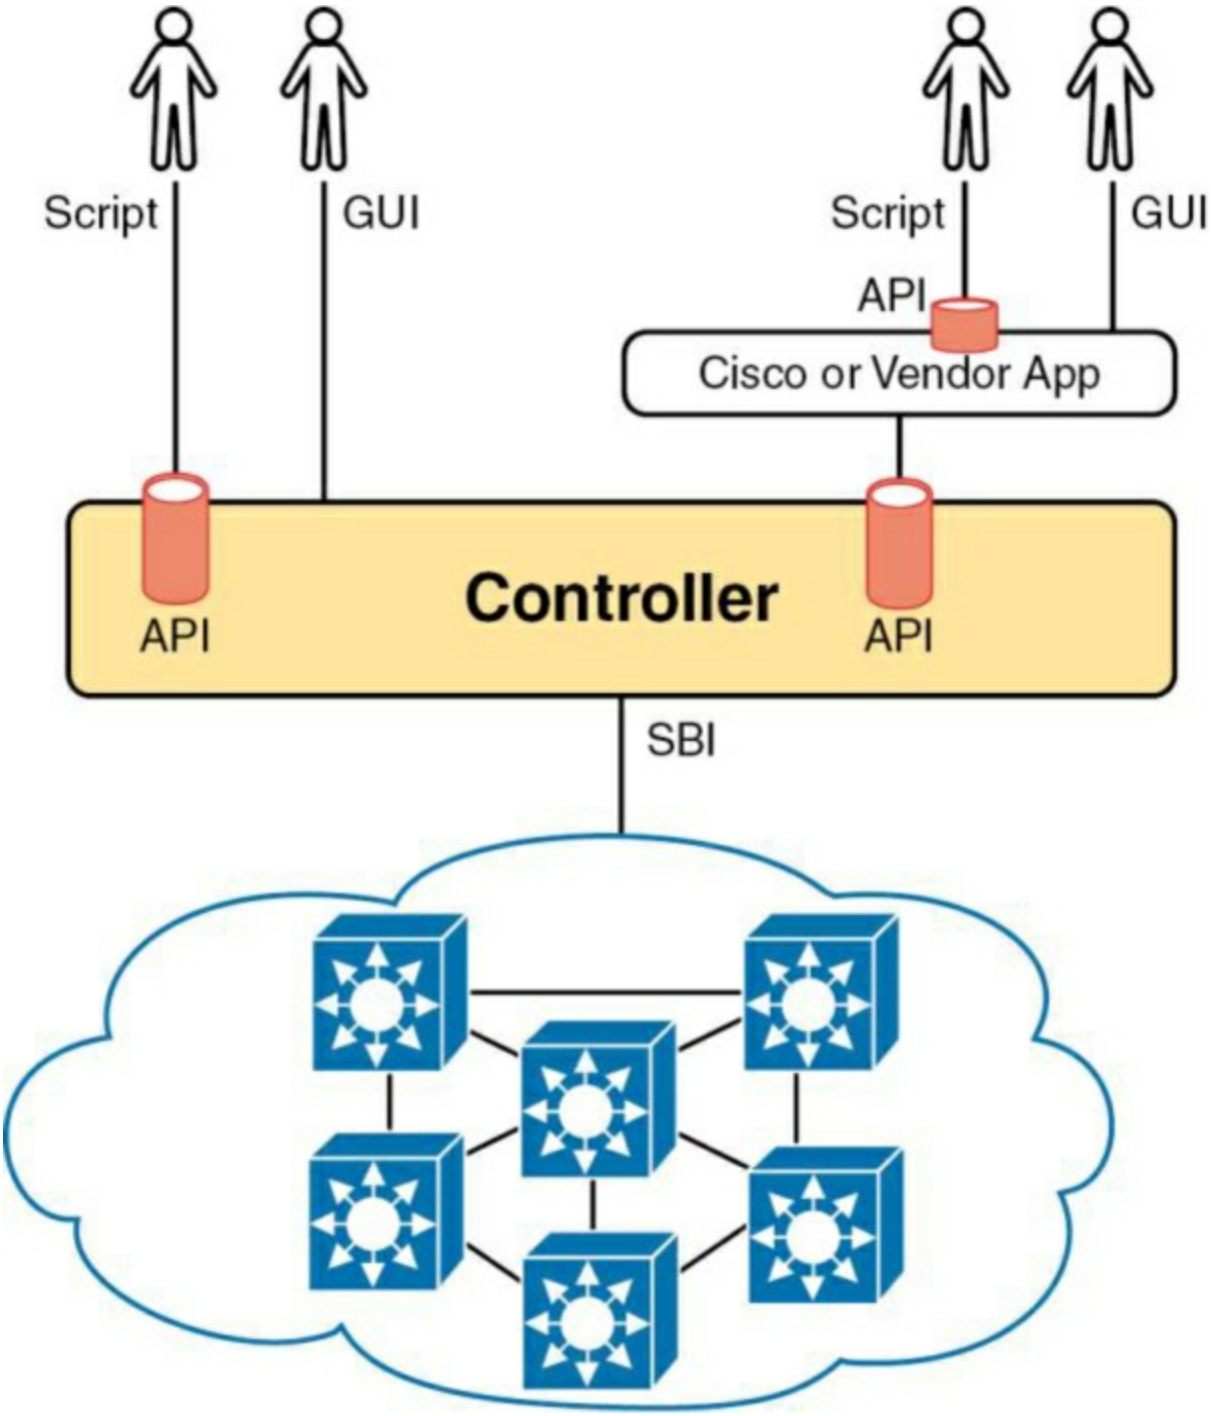

Figure 2-1 *SDA Architecture with Cisco Catalyst Center*

The architecture includes:

• **Underlay:** Physical devices and connections that provide IP connectivity

Technet24

for nodes within the campus LAN. It supports VXLAN overlay tunnels.

• **Overlay:** Utilizes VXLAN tunnels between Cisco SD-Access devices like switches and fabric-enabled access points to transport traffic across the fabric.

• **Fabric:** Combines overlay and underlay technologies to deliver data across the network, supporting features like host mobility for seamless connectivity as users move within the campus.

For example, in Figure 2-2, the two switches, SW1 and SW2, dynamically create a VXLAN tunnel between the endpoints in the overlay. The exact path the data takes between the two switches is determined by the underlay's Layer 2 and Layer 3 process.

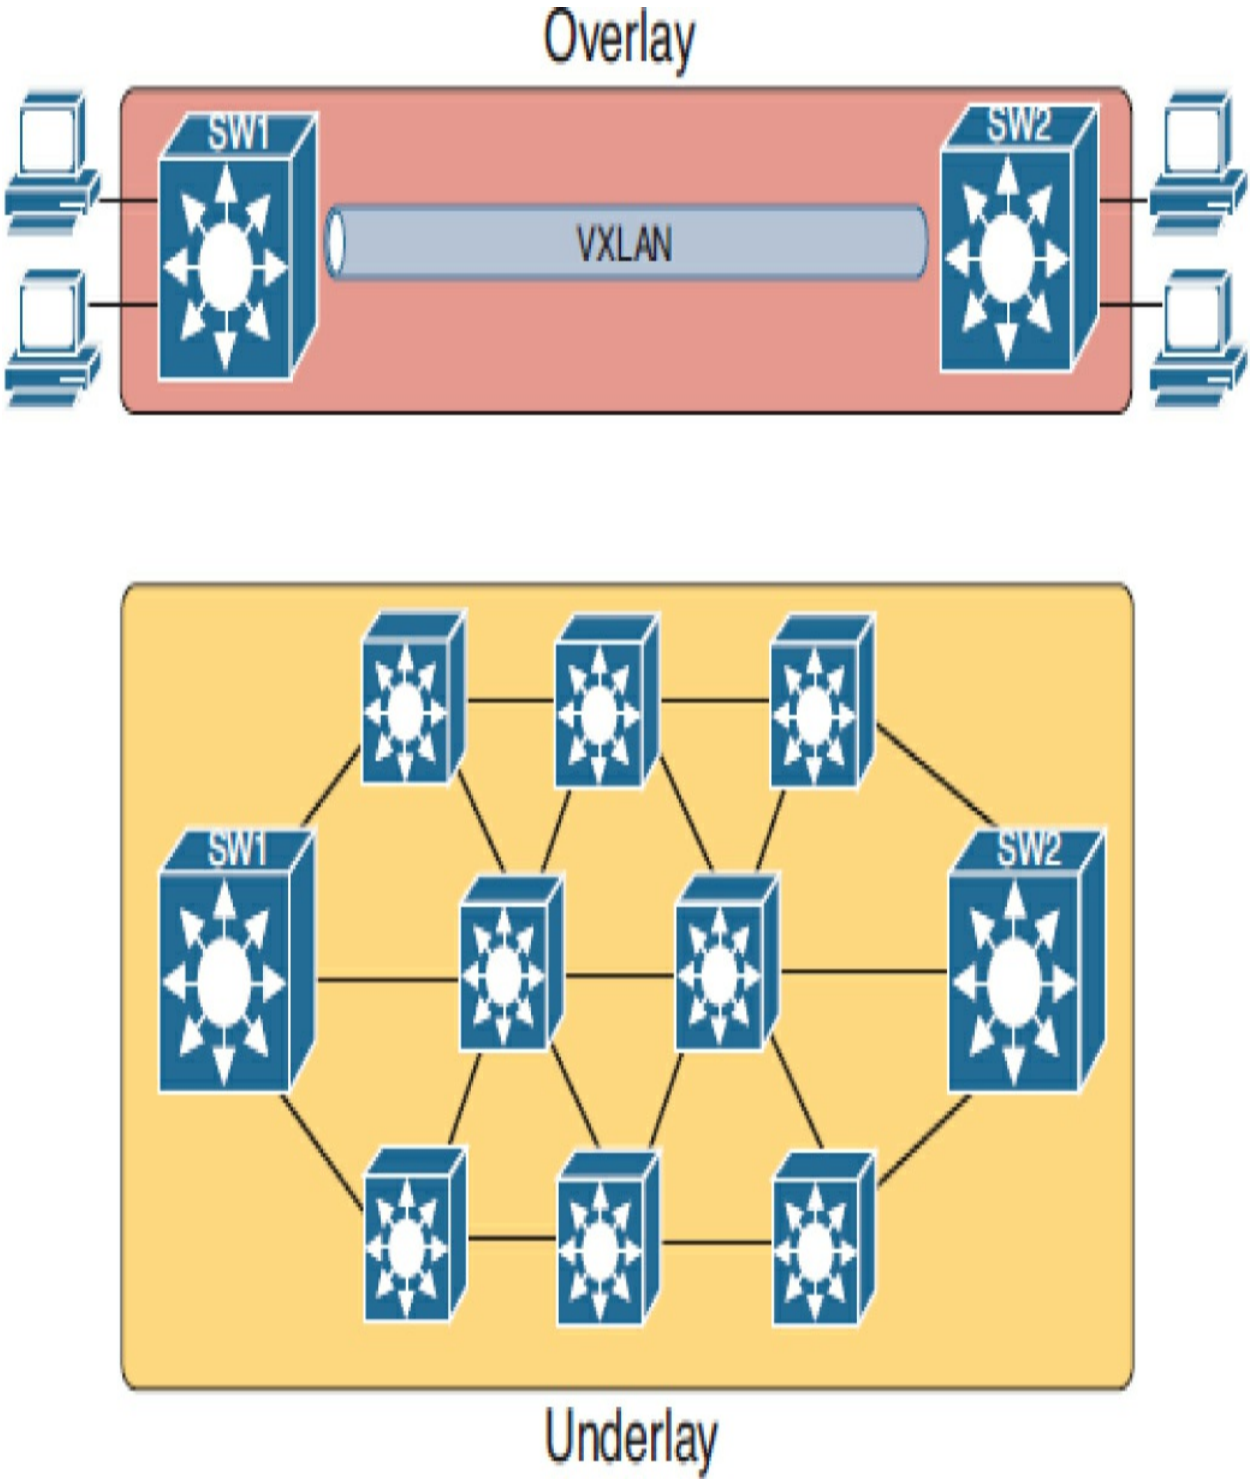

Figure 2-2 *Overlay and Underlay*

Underlay

Technet24

With Cisco SD-Access, the underlay provides connectivity between fabric edge nodes, supporting VXLAN tunnels in the overlay network. It includes all the devices that form the physical network, along with the necessary configuration and operation to support the overlay. Some existing equipment may support Cisco SD-Access. However, newer equipment enables network automation through Cisco Catalyst Center.

Existing Equipment for the Cisco SD-Access Underlay

Deploying Cisco SD-Access on existing campus networks involves using current hardware that meets compatibility requirements. The process requires careful planning to avoid disruptions to existing network operations. The key points to consider include:

• **Compatibility:** Existing hardware must be on the Cisco SD-Access compatibility list, which includes various models based on their roles within the network (fabric edge node, fabric border node, and fabric control-plane node). For example, compatible models for fabric edge nodes include Catalyst 9200, 9300, 9400, 9500, and 9600 series switches, along with some older models, like Catalyst 3850 and 3650.

• **Software requirements:** Devices must run software versions that support SD-Access functions based on their designated roles. Ensuring the correct software versions is crucial for maintaining network stability and functionality.

• **Manual configuration:** When integrating SD-Access into an existing network, the underlay configuration must be done manually. To avoid disrupting existing traffic, Cisco Catalyst Center should not be used to configure devices that are already in production.

• **Planning and roles:** Engineers need to plan the roles of each device carefully. For instance, fabric edge nodes connect to endpoint devices, fabric border nodes connect to external networks (for example, WAN routers or ACI data centers), and fabric control-plane nodes handle special control plane functions that require more CPU and memory.

• **Migration:** The migration process might involve updating IOS versions on existing devices. The goal is to integrate the new SD-Access features without negatively impacting ongoing network operations.

New Equipment for the Cisco SD-Access Underlay

For new deployments, compatible hardware can be ordered, and Cisco Catalyst Center can configure underlay features automatically. Traditional LAN design considerations, such as port requirements, cabling, and power needs, still apply. The underlay uses a routed access layer design with Layer 3 switches and IS-IS routing protocol, eliminating the need for STP/RSTP and HSRP, as shown in Figure 2-3.

Technet24

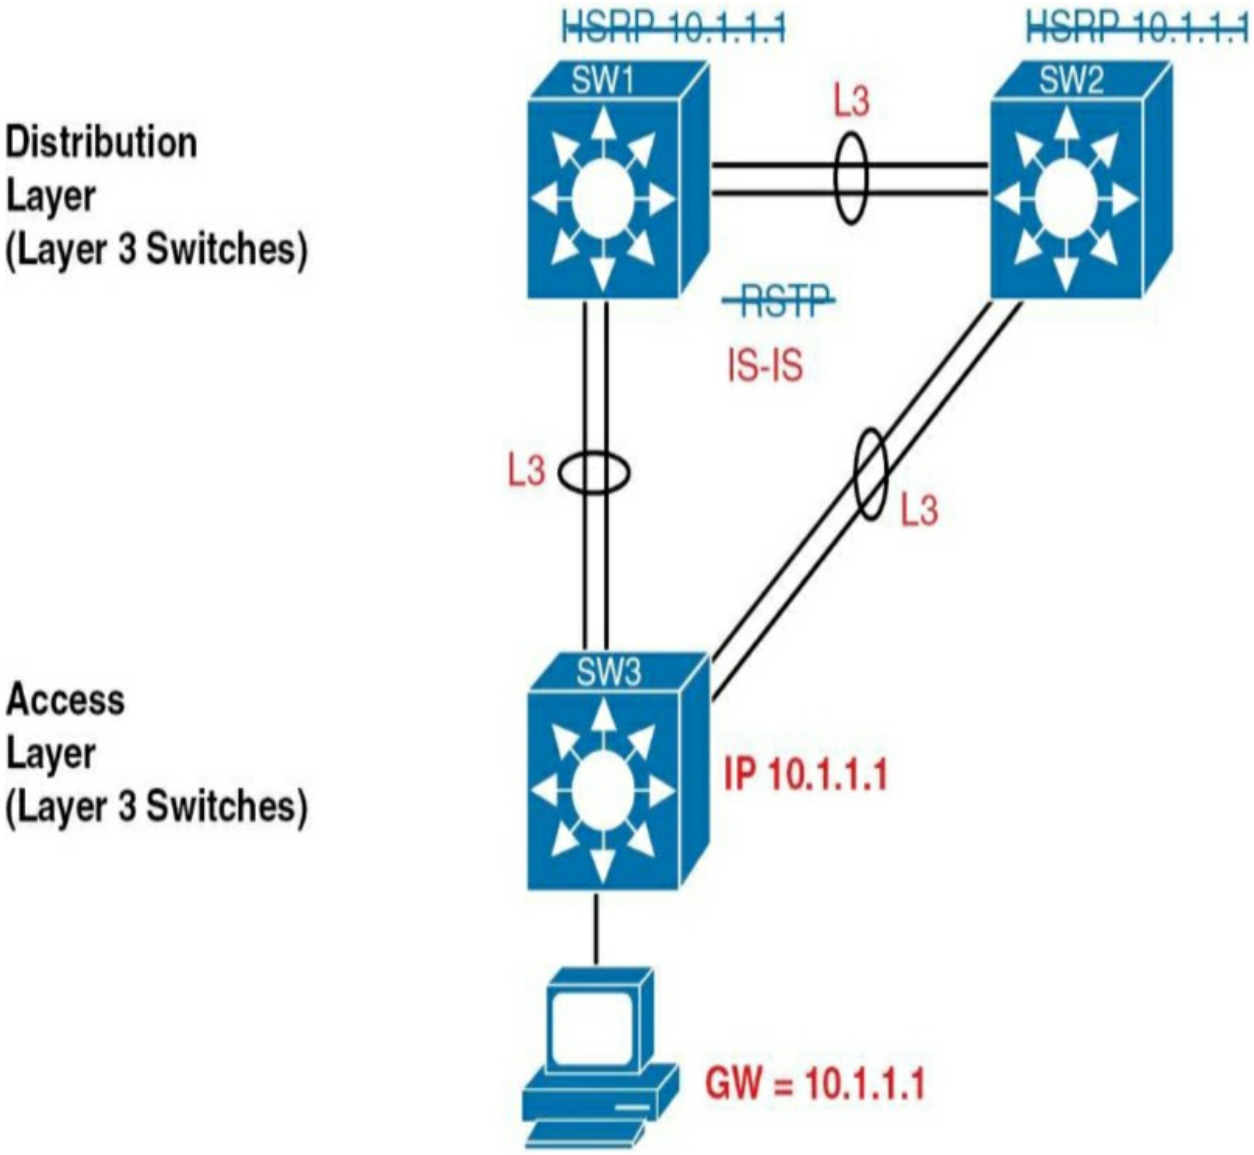

Figure 2-3 *Cisco SC-Access Layer 3 Configuration*

Deploying Cisco SD-Access with new hardware (a greenfield deployment) simplifies the process by removing concerns about disrupting existing network traffic. Key considerations include the following:

• **Ordering compatible hardware:** Organizations can order hardware and software that is fully compatible with Cisco SD-Access from the start. This

approach avoids the complexities associated with upgrading and configuring existing gear.

• **Cisco Catalyst Center configuration:** Underlay configuration automation ensures consistent and optimal setup for supporting SD-Access.

• **Traditional LAN design considerations:** While deploying a new SD-Access network, traditional LAN design factors must still be considered. These include:

• Number of ports needed in each wiring closet

• Port speeds required

• Benefits of switch stacking

• Cable types and lengths

• Power requirements (PoE/PoE+)

• Power availability versus PoE power requirements

• Link capacity (speed and number of links) between switches

• **Routed access layer design:** All LAN switches operate as Layer 3 switches with routing enabled on all interswitch links. This design maximizes available bandwidth and eliminates the need for STP/RSTP.

Technet24

• **Consistent underlay configuration:** Cisco Catalyst Center configures the underlay with consistent settings across the enterprise. Consistent configuration simplifies the operational model and ensures that the underlay network is optimized for supporting the SD-Access fabric.

Overlay

The Cisco SD-Access overlay network uses Virtual Extensible LAN (VXLAN) tunnels to encapsulate and transport data frames across the fabric, creating a virtualized network layer over the physical infrastructure. Key points include the following:

• **Encapsulation:** The overlay begins with an endpoint sending a frame that is encapsulated by the first fabric edge node into a VXLAN tunnel. This encapsulated frame is then forwarded through the fabric.

• **VXLAN tunnels:** These tunnels allow for flexible and scalable communication by extending Layer 2 networks over Layer 3 networks. VXLAN encapsulation supports both Layer 2 and Layer 3 forwarding features, making it ideal for modern data centers and campus networks.

• **Performance:** The encapsulation and forwarding process is handled by the ASICs (application-specific integrated circuits) on each switch, ensuring that there is no performance penalty for the added encapsulation complexity.

• **Endpoint mobility:** The overlay network supports host mobility, allowing

endpoints to move across the network without losing connectivity or requiring reconfiguration. The control plane keeps track of endpoint locations, ensuring seamless communication.

VXLAN Tunnels in the Overlay (Data Plane)

VXLAN tunnels are a fundamental part of the Cisco SD-Access overlay, providing a method for encapsulating data frames and transporting them across the network fabric, as shown in Figure 2-4.

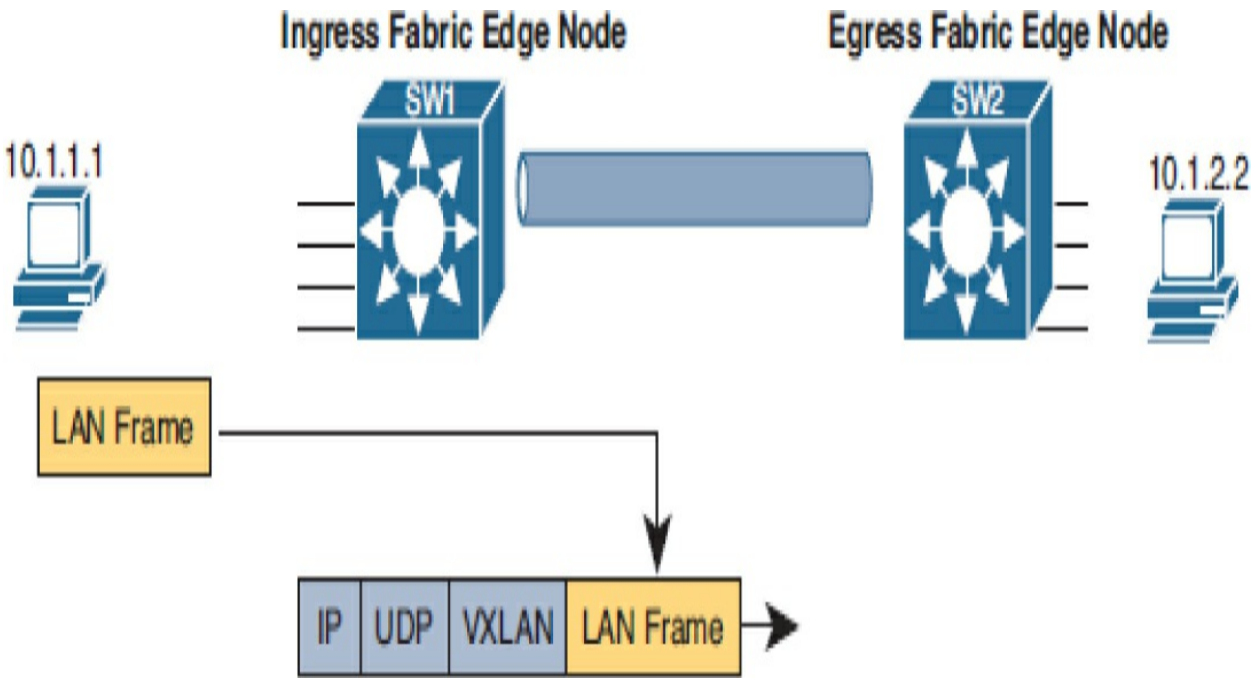

Figure 2-4 *VXLAN Tunneling Protocol Operation*

The VXLAN tunnel in the overlay works like this:

An endpoint sends a frame.

The ingress fabric edge node encapsulates the frame in the VXLAN tunneling

Technet24

specification.

The frame is forwarded to the underlay fabric.

The other nodes in the underlay forward the frame based on the VXLAN tunnel details.

The egress fabric edge node removes the VXLAN details.

The frame is forwarded to the destination endpoint.

Key details of this VXLAN encapsulation process include the following:

• **ASIC processing:** This processing ensures high performance and no impact on network speed.

• **Flexibility and extensibility:** The encapsulation process allows for additional header fields needed by Cisco SD-Access features. This flexibility supports the creation of new networking capabilities.

• **Layer 2 and Layer 3 forwarding:** By encapsulating entire data link frames rather than just IP packets, VXLAN supports both Layer 2 and Layer 3 forwarding.

• **Separate IP address spaces:** The underlay network uses a different IP address space from the overlay. For example, Figure 2-5 shows the underlay using the 172.16.0.0/16 space (SW1 and SW2), while the overlay is using

10.0.0.0/8 address space (endpoints and VXLAN tunnel). This separation ensures efficient routing and forwarding.

Figure 2-5 *Cisco SD-Access Uses Different Address Spaces for the Underlay and the Overlay*

LISP for Overlay Discovery and Location (Control Plane)

Locator/ID Separation Protocol (LISP) is used within the Cisco SD-Access overlay to manage endpoint discovery and location. This protocol enables the control plane to efficiently track and route data to endpoints.

Figure 2-6 shows how the LISP server learns about endpoints.

Technet24

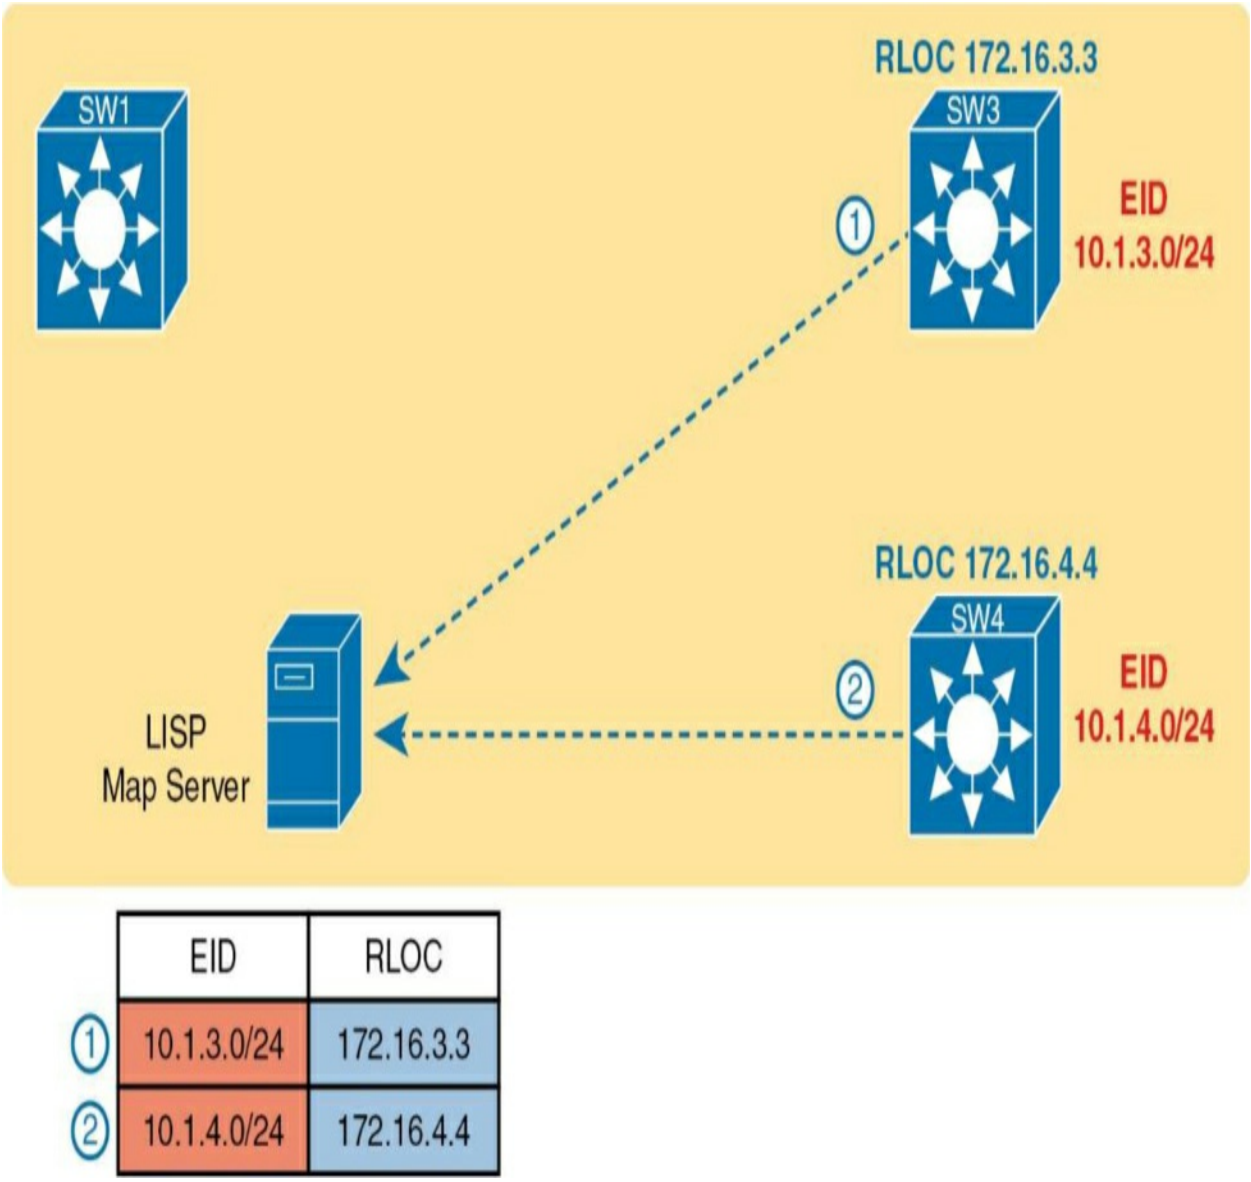

Figure 2-6 *Edge Nodes Register IPv4 Prefixes (Endpoint IDs) with the LISP Map Server*

Fabric edge nodes learn about endpoints using traditional methods (for example, MAC and IP addresses) and register these endpoints with the LISP map server (1 and 2 in Figure 2-6). The LISP map server maintains a database of endpoint identifiers (EIDs) and their corresponding routing locators (RLOCs), as shown in the table in Figure 2-6. The RLOCs identify

the fabric edge nodes that can reach specific endpoints.

In Figure 2-7, when a data frame arrives at fabric ingress node SW3 (3), the node consults the LISP map server to determine the correct VXLAN tunnel for forwarding the frame (4). The LISP map server finds the entry it built (5) and then validates with the egress node that the RLOC is still correct (6). The egress node (SW3) then informs the SW1 that 10.1.3.1 can be reached through SW3 (7).

Technet24

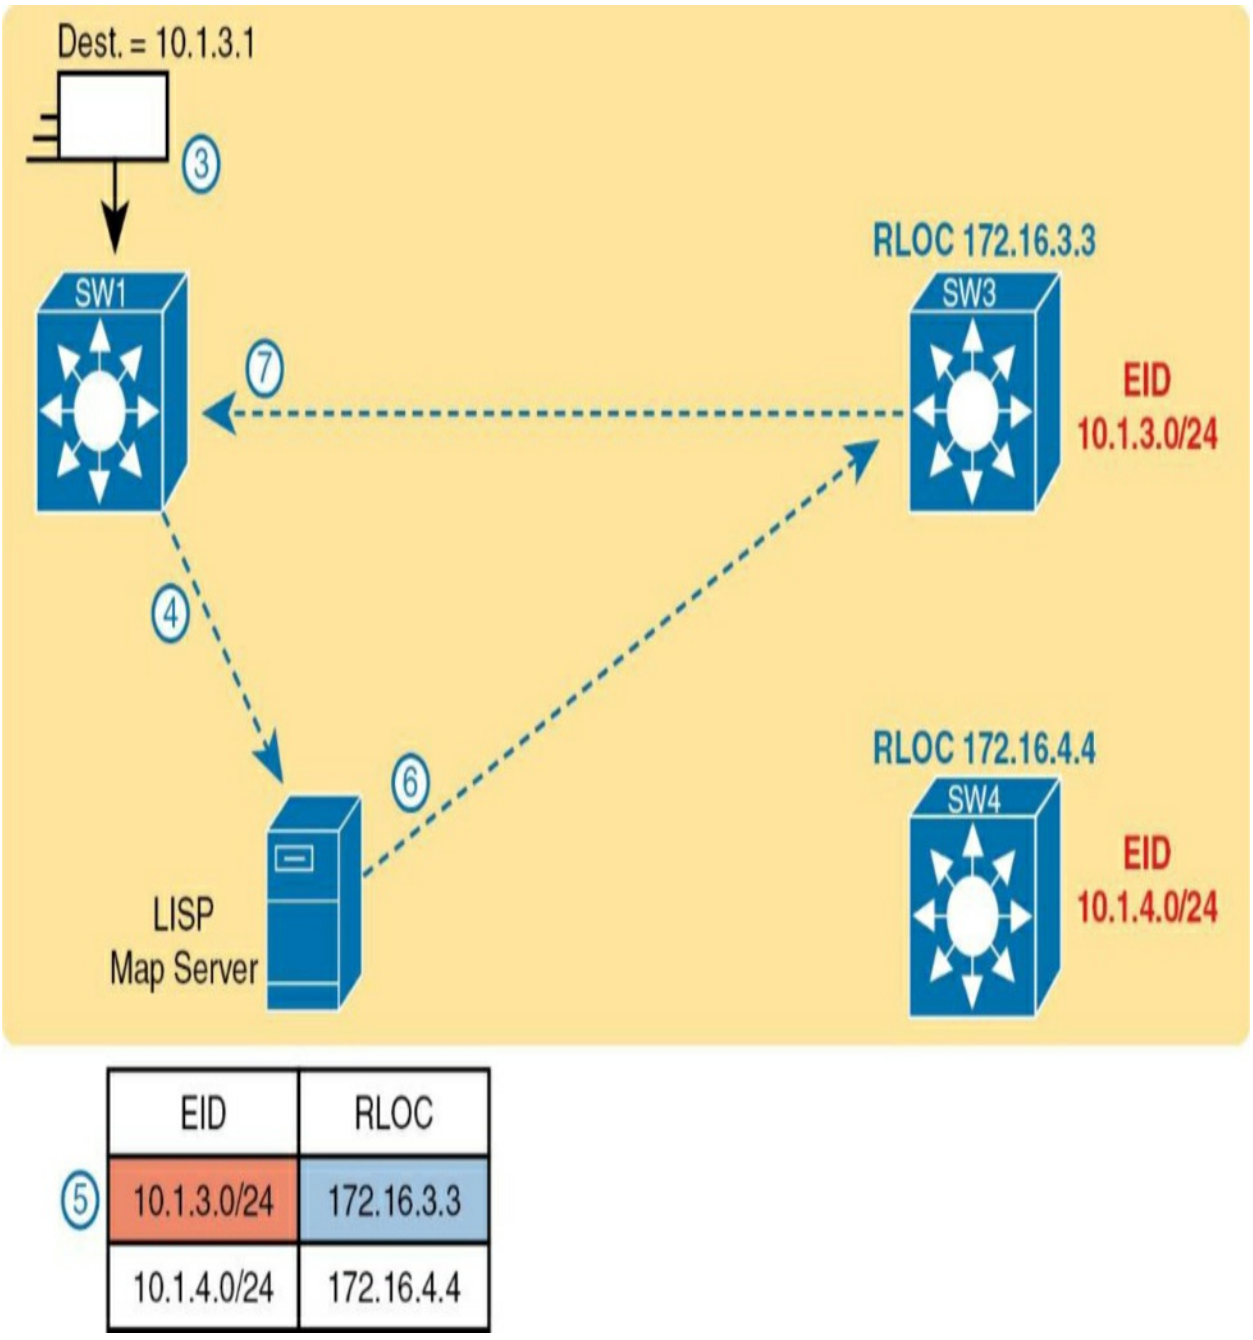

Figure 2-7 *Ingress and Egress Router Discovery*

Now SW1 and SW3 are ready to form a VXLAN tunnel, as shown in Figure 2-8. SW1 encapsulates the original Ethernet frame with the necessary IP, UDP, and VXLAN headers. The outer IP header lists the RLOC IP address as the destination, ensuring that the message is delivered through the Cisco SD-

Access fabric to SW3.

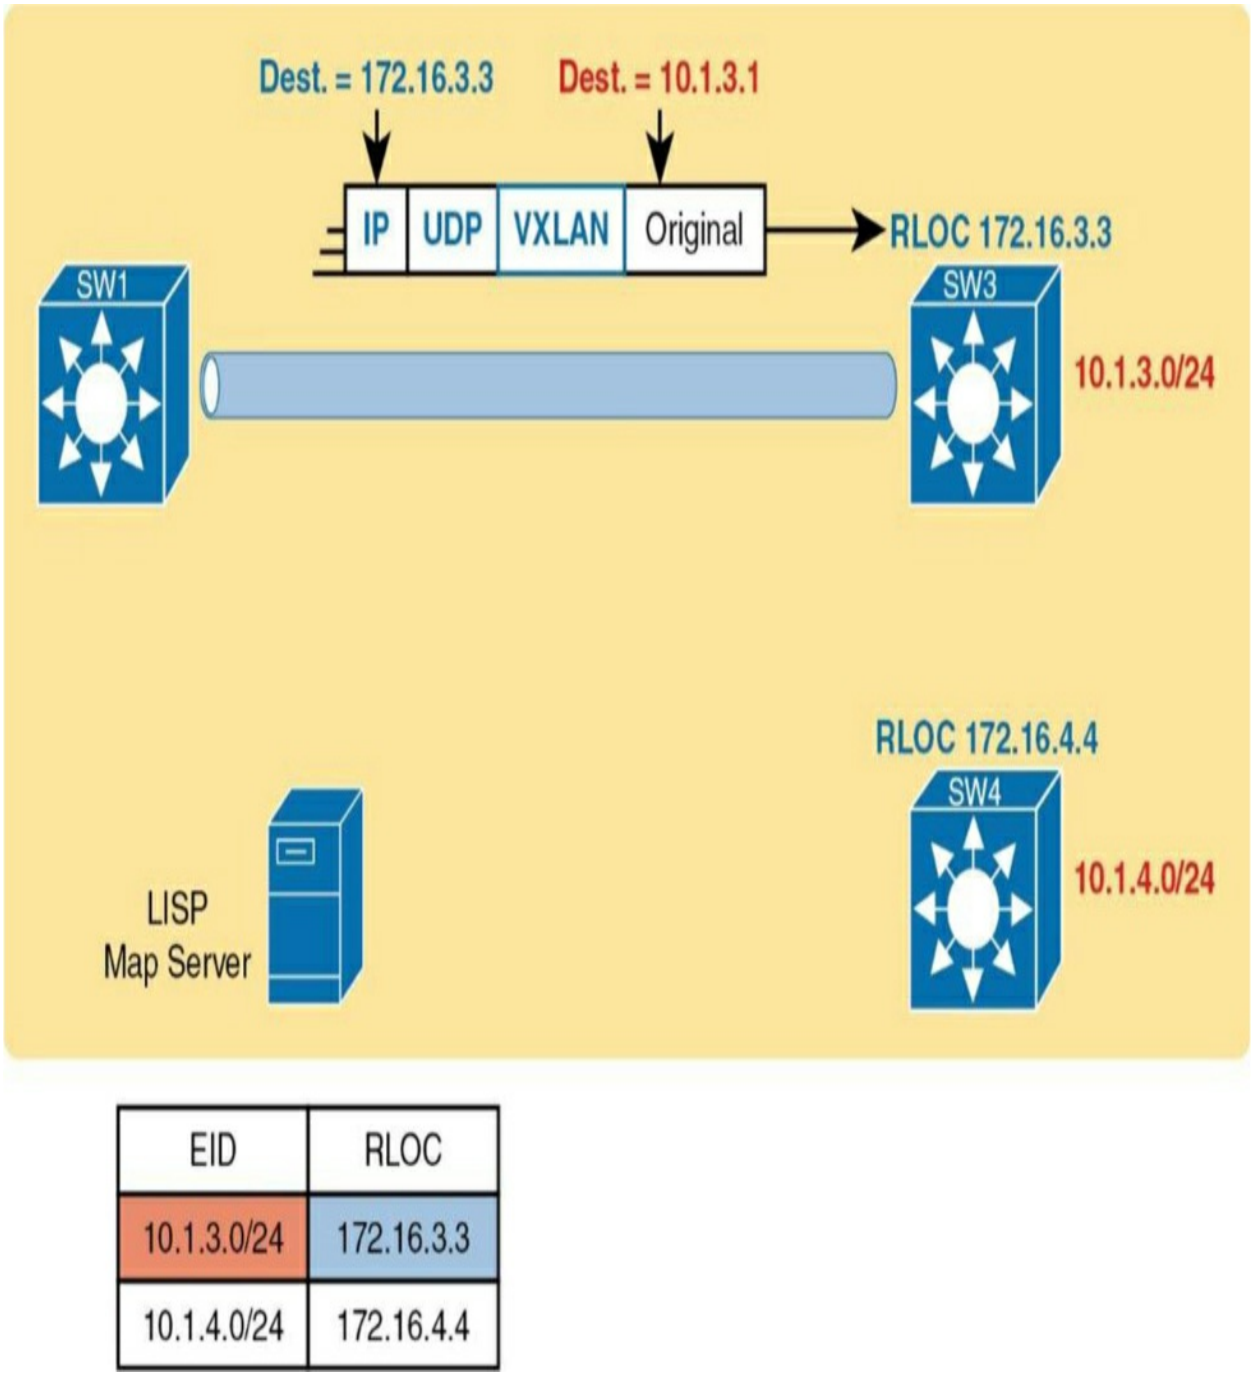

EID	RLOC
10.1.3.0/24	172.16.3.3
10.1.4.0/24	172.16.4.4

Figure 2-8 *Ingress Tunnel Router (ITR) SW1 Forwards Based on LISP Mapping to SW3*

Technet24

The LISP map server dynamically updates the location information as endpoints move within the network. This process ensures that the Cisco SD-Access fabric can dynamically and efficiently manage endpoint mobility and data forwarding, providing a robust and scalable network solution.

Cisco Catalyst Center

Cisco Catalyst Center has two roles:

• A controller in a network that uses Cisco SD-Access

• A network management platform for traditional (non-Cisco SD-Access) network devices

Cisco Catalyst Center is available as a preinstalled application on a Cisco appliance or as a virtual appliance in a hypervisor environment. It provides a GUI and northbound REST API for managing network design, configuration, and implementation. It automates operations using various APIs, supporting both SD-Access and traditional network devices. This robust architecture, shown in Figure 2-9, with northbound and southbound APIs, simplifies management, reduces manual configuration, and enhances network performance.

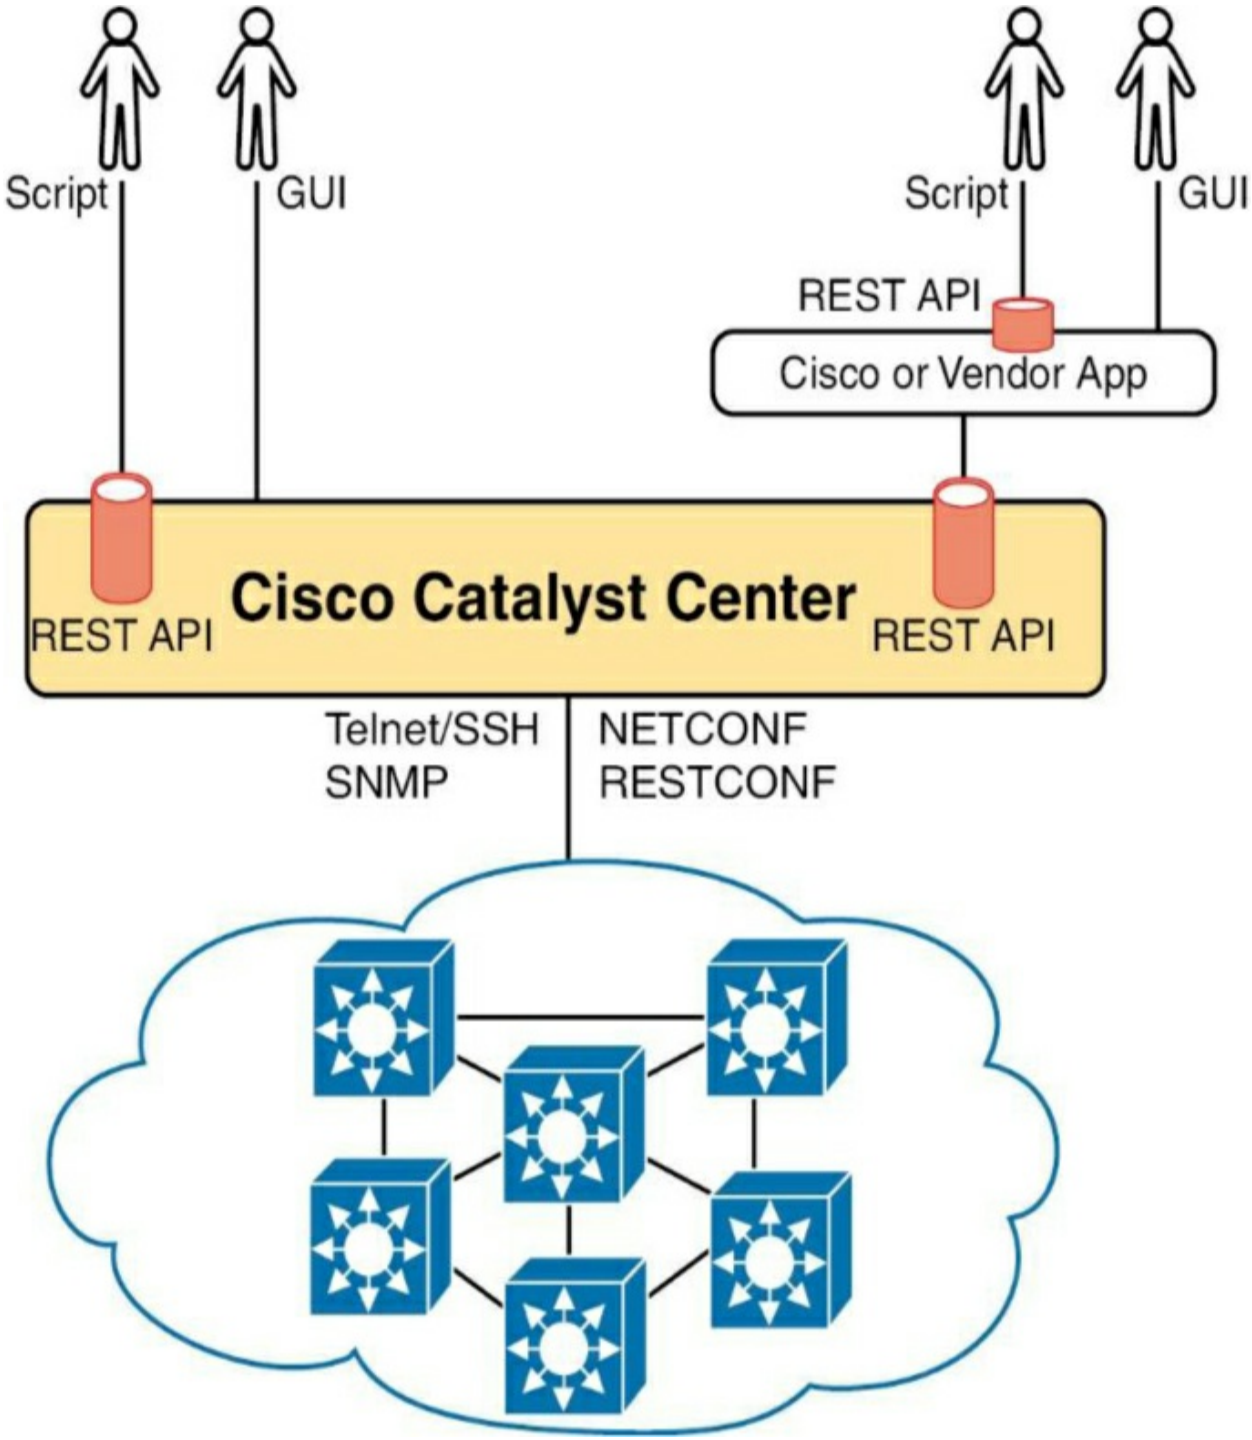

Figure 2-9 *Cisco Catalyst Center with Northbound and Southbound Interfaces*

Cisco Catalyst Center and Scalable Groups

Technet24

Cisco SD-Access introduces scalable groups to streamline security configuration. Scalable groups allow administrators to define security policies based on user groups rather than IP addresses. This approach simplifies the management of security policies, enabling dynamic and flexible control of network access.

Issues with Traditional IP-Based Security

Traditional IP-based security relies on access control lists (ACLs) to manage network access, which can become complex and unwieldy over time. As ACLs grow with new requirements, they become difficult to manage, troubleshoot, and update. The risk of inadvertently disrupting network operations increases with each change, making security management challenging and error prone.

For example, consider the ACL in Figure 2-10. Each number represents a new policy implemented over the life of the ACL.

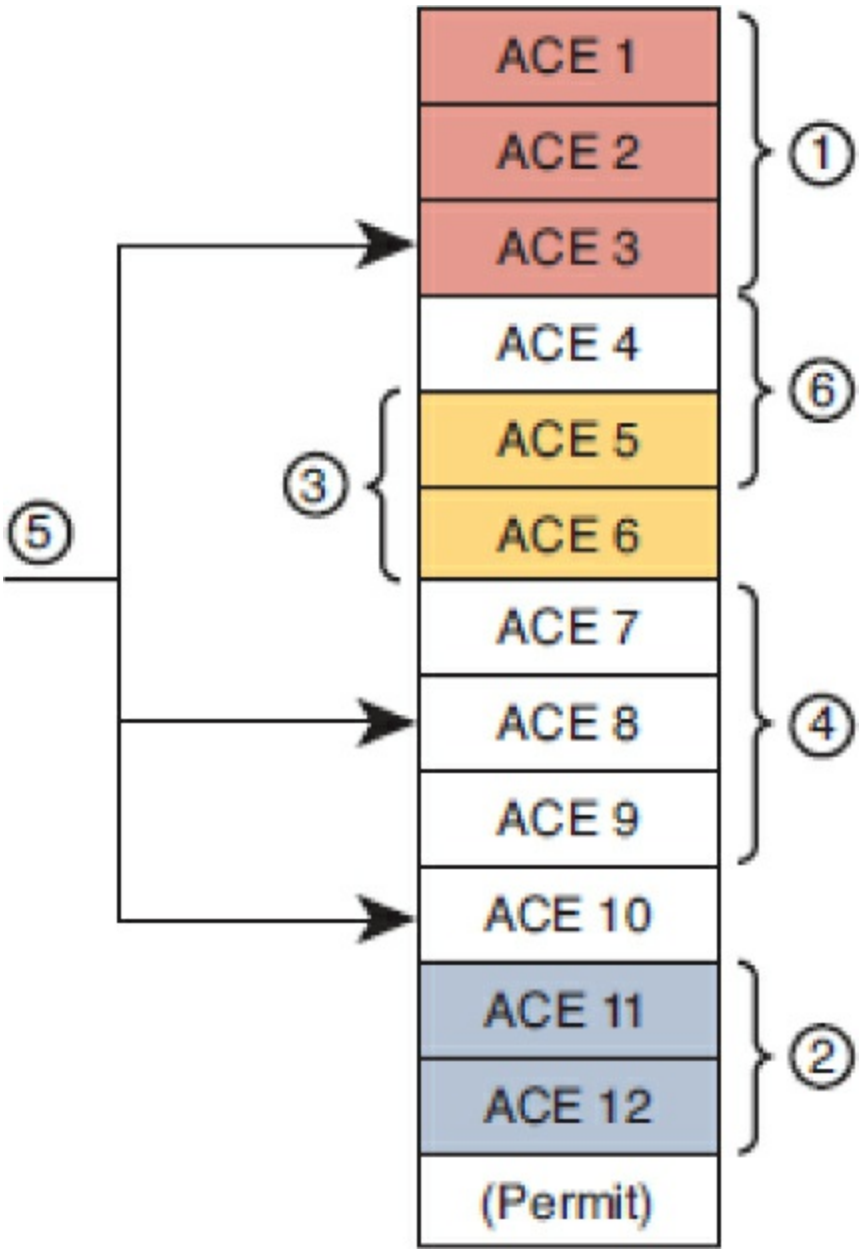

Figure 2-10 *ACEs in an ACL After Six Changes*

Determining where to place the new access control entries (ACEs) within the existing ACL can be a complex and risky process. Also, unless an ACL is fully documented, you are never quite sure what effects a new policy will have on existing policies—or even if the existing policy is still valid.

Technet24

Cisco SD-Access Security Is Based on User Groups

One of the benefits of Cisco SD-Access security groups is that you can enforce a policy without even thinking about IP address ranges and ACLs. Instead of writing new ACEs each time a policy needs to be implemented, the policy is defined in Cisco Catalyst Center. Then, as needed, Cisco Catalyst Center configures the devices in the fabric to enforce the security, as shown in Figure 2-11.

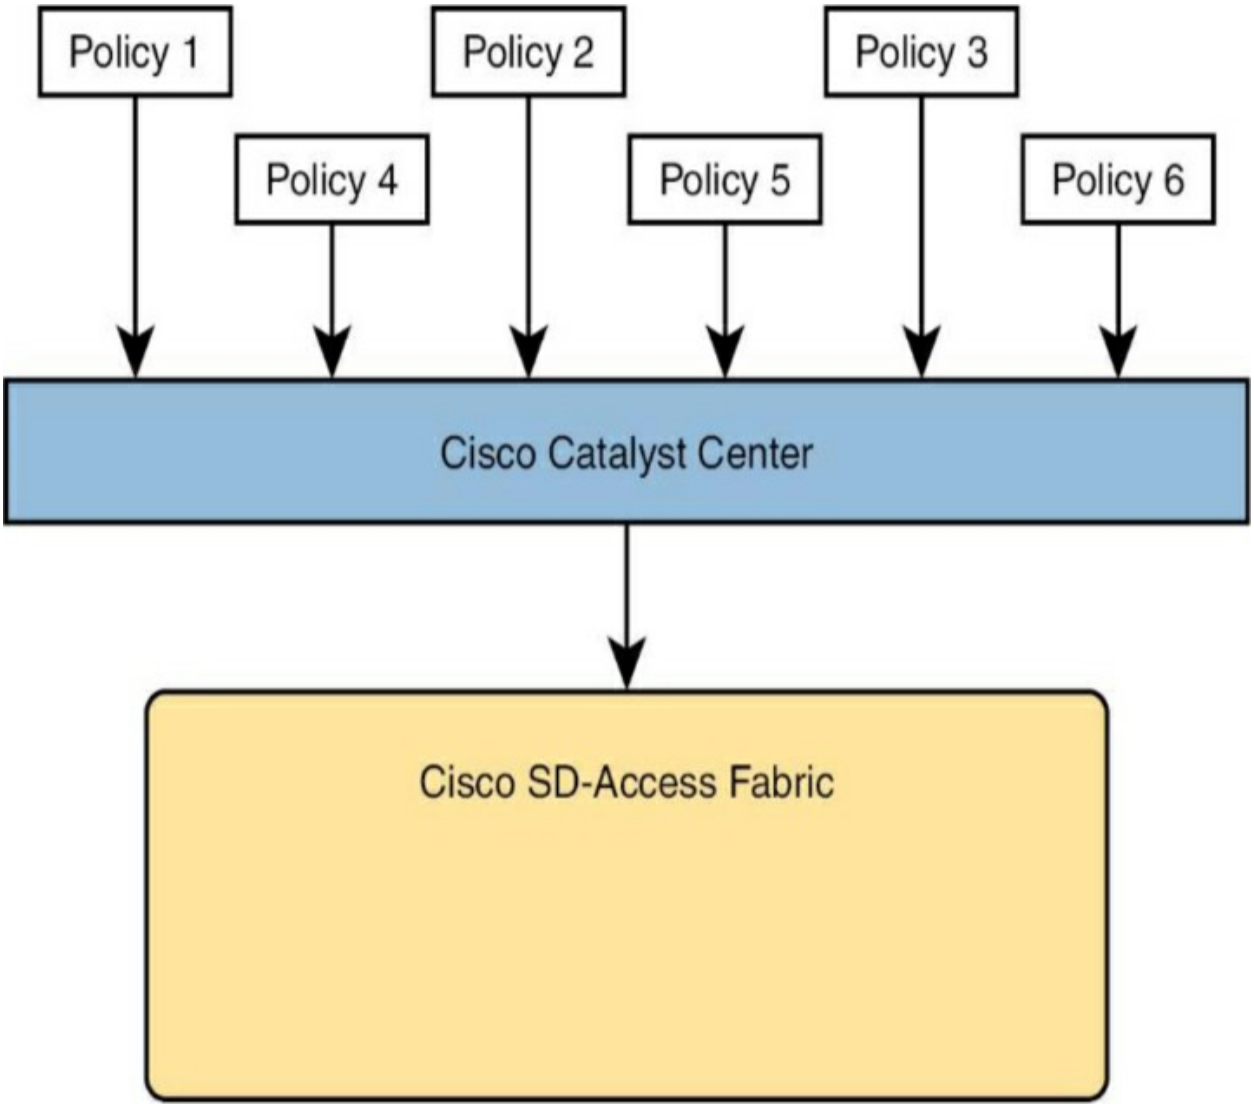

Figure 2-11 *Six Policies Implemented by Cisco Catalyst Center*

The Cisco SD-Access policy model solves the challenges with traditional ACLs:

• Each new security requirement can be considered separately, without analysis of an existing ACLs.

• Each new requirement can be considered without searching for all the ACLs in the likely paths between endpoints and analyzing each and every ACL.

• Cisco Catalyst Center keeps the policies separate.

• Each policy can be removed without fear of impacting the logic of the other policies.

To implement policies in Cisco SD-Access, you tie them to security groups, called scalable groups. A security group is identified with a scalable group tag (SGT). If Cisco Catalyst Center sees a permit action between the source/destination pair of SGTs, Cisco Catalyst Center directs the edge nodes to create the VXLAN tunnel. The SGTs for source and destination are added to the VXLAN header, along with the VXLAN IDs (VNIDs), as shown in Figure 2-12.

Technet24

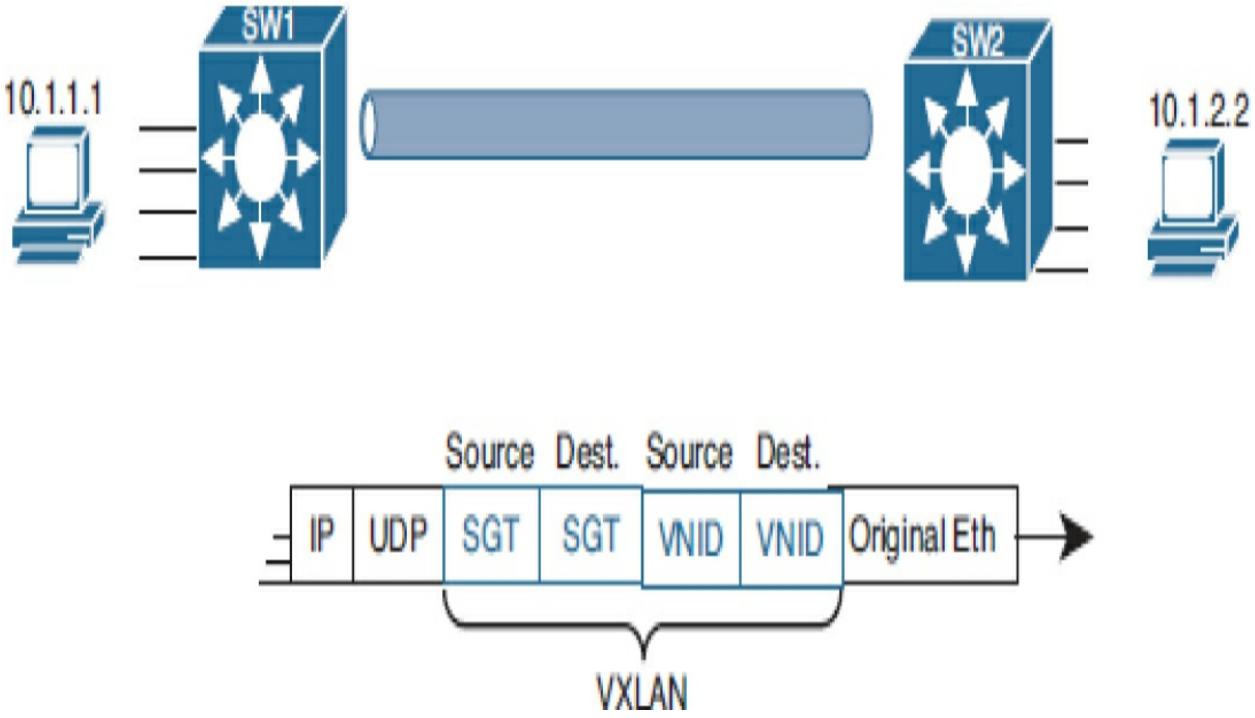

Figure 2-12 *VXLAN Header with SGTs and VNIDs*

Cisco Catalyst Center as a Network Management Platform

Cisco Catalyst Center extends beyond SD-Access, serving as a comprehensive network management platform for both SD-Access and traditional network devices. It offers features such as device discovery, inventory management, topology mapping, lifecycle management, application visibility, converged wired and wireless management, software image management, and plug-and-play capabilities. This comprehensive toolset enhances network operations, making it easier to manage and maintain complex network environments.

Cisco Catalyst Center Similarities to Traditional Management

Cisco Catalyst Center shares many features with traditional network

management systems, such as device discovery, inventory management, and topology mapping. It provides a single-pane-of-glass interface for monitoring and managing network devices, supporting protocols like SNMP, SSH, and Telnet. Event and threshold alerting, as well as detailed device information, are accessible through the intuitive GUI.

Differences Between Cisco Catalyst Center and Traditional Management

While similar to traditional network management systems, Cisco Catalyst Center offers advanced features tailored to modern network environments. It supports intent-based networking for SD-Access, enabling automated configuration and dynamic policy enforcement. Features like application policy deployment, Encrypted Traffic Analytics (ETA), Device 360 and Client 360 views, network time travel, and path trace differentiate it from traditional management platforms, providing enhanced security, reduced operational complexity, and improved network performance.

Configuration Management Tools: Ansible and Terraform

Before reviewing the features and capabilities of Ansible and Terraform, let’s first review the reasons these tools are essential in managing devices in larger enterprises that employ multiple network engineers.

Manual Configuration Management

Manually managing device configurations in a production network poses

Technet24

several challenges, including determining the intended configuration for each device, preventing configuration drift, tracking changes (version control), and maintaining consistency across multiple engineers and devices. The following list summarizes these challenges:

• **Determining the device configuration:** To address configuration management challenges, many enterprises store device configurations in a central location. This approach allows for better backup and sharing among team members. However, it raises questions about which version of the configuration should be considered the “source of truth” when engineers are allowed to change the configuration manually on the devices.

• **Preventing configuration drift:** Configuration drift happens when individual engineers make changes to resolve issues without properly documenting or communicating these changes.

• **Tracking changes with version control:** The manual configuration process lacks built-in change tracking, making it difficult to identify who made changes, when they were made, and what specific modifications occurred.

• **Maintaining consistent configurations:** Even with centralized configuration management and version control, the risk of configuration drift remains if engineers make manual changes directly on devices.

Traditional manual configuration methods that we have reviewed throughout this book do not address these issues effectively. As companies grow and

network complexity increases, the need for more sophisticated configuration management becomes apparent.

Configuration Management Tool Overview

Configuration management tools can monitor device configurations and compare them to the intended ideal configuration stored centrally. These tools can either automatically reconfigure devices to match the ideal configuration or notify engineering staff of discrepancies. This process helps maintain consistency between centralized configurations and actual device configurations.

Configuration Provisioning

Configuration provisioning refers to the deployment of changes to network devices that are made in the centralized configuration management system. Key features of configuration management tools include the ability to implement changes on specific devices or groups of devices, validate changes before implementation, revert changes if necessary, and confirm successful configuration updates.

Configuration Templates and Variables

Configuration management tools often support the use of configuration templates and variables to streamline the management of similar devices and scheduled execution of configuration changes. Using such tools allows for easier management of devices with similar roles, as engineers can edit a

Technet24

standard template file separately from each device's unique variables. In Figure 2-13, the branch router template and variables are used to pushed configurations to all the branch routers.

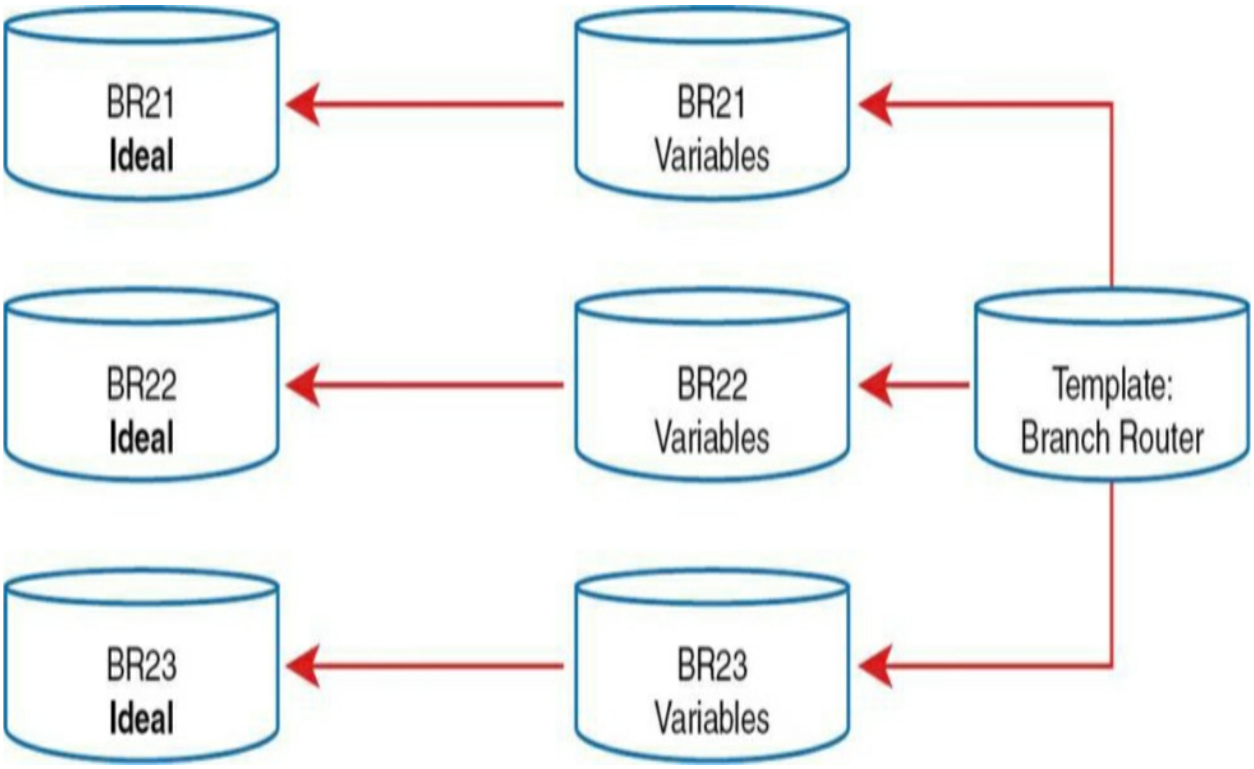

Figure 2-13 *Configuration Template and Variables*

Example 2-1 shows what a simplified template might look like for configuring the branch routers with their hostnames, IPv4 addressing, and OSPF routing. The variables are shown in the double curly brackets.

Example 2-1 *Variables for Branch Router Configuration*

```
hostname {{hostname}}
!
interface GigabitEthernet0/0
```

```
 ip address {{address2}} {{mask2}}
 ip ospf {{OSPF_PID}} area {{area}}
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/1/0
 ip address {{address2}} {{mask2}}
 ip ospf {{OSPF_PID}} area {{area}}
!
router ospf {{OSPF_PID}}
 router-id {{RID}}
```

Using templates and variables increases focus on standardization, simplifies deployment of new devices, facilitates troubleshooting, and allows for easier tracking of changes to both standard configurations and device-specific settings.

This shift away from traditional manual configuration approaches to a more automated approach using configuration management tools is known as infrastructure as code (IaC). IaC allows administrators to use these tools to create definition files, like the one shown in Example 2-1, that can set up, modify, and manage entire network infrastructures, including servers, routers, switches, and cloud resources. Both Ansible and Terraform implement IaC principles, with Ansible using YAML for its playbooks and Terraform using its own HashiCorp Configuration Language (HCL).

Technet24

Let’s briefly review Ansible and Terraform architecture, features, and configuration methods.

Ansible

Ansible operates without requiring specialized software on managed devices, instead using standard protocols like SSH. This open-source tool enables efficient, consistent, and repeatable infrastructure setup and management.

Ansible’s architecture consists of three main components: playbooks, modules, and inventories. Playbooks define the tasks to be automated, modules are the building blocks that perform specific functions, and inventories list the devices to be managed. A typical Ansible workflow, shown in Figure 2-14, involves an administrator creating playbooks and inventories, which are then processed by Ansible’s core. Ansible uses plugins and APIs to interact with external systems, such as cloud providers and the configuration management database (CMDB), and execute tasks on specified hosts according to the playbook instructions.

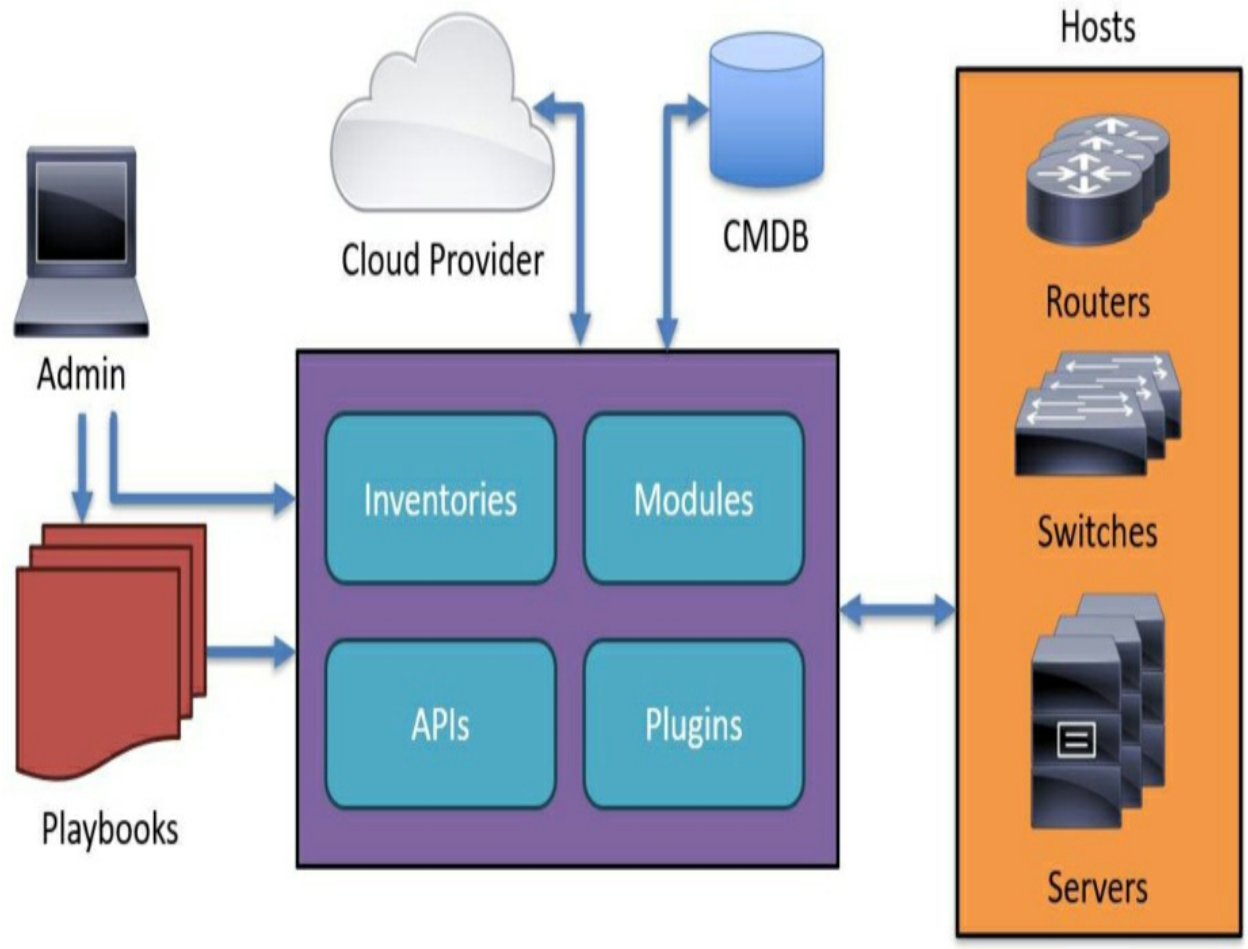

Figure 2-14 *Ansible Architecture*

Example 2-2 shows an Ansible YAML playbook that will call the Cisco routers in the inventory file named hosts and display their IP addresses using the **show ip interface brief** command.

Example 2-2 *Ansible YAML Playbook File*

```
- name: Show IP Interface Brief on Cisco IOS routers
  hosts: cisco_routers
  gather_facts: no
  tasks:
```

Technet24

```
    - name: Run 'show ip interface brief' command
      ios_command:
        commands:
          - show ip interface brief
      register: output

    - name: Display command output
      debug:
        var: output.stdout
```

Example 2-3 shows a simplified version of what the inventory host file might look like with the arguments specifying the connection details for two Cisco routers.

Example 2-3 *Ansible Inventory File*

```
[cisco_routers]
192.168.1.1 ansible_connection=network_cli ansible_net
ansible_user=your_username ansible_password=C1$c0$$hp@
ansible_become=yes ansible_become_method=enable
ansible_become_password=C1$c0pr1vp@$$w0rd

192.168.2.1 ansible_connection=network_cli ansible_net
ansible_user=your_username ansible_password=C1$c0$$hp@
ansible_become=yes ansible_become_method=enable
ansible_become_password=C1$c0pr1vp@$$w0rd
```

To execute the playbook from a Linux command line, enter the following command:

```
Linux-OS# ansible-playbook -i hosts display_ipv4_addre
```

This command runs the specified Ansible playbook (show_ip_interface_brief.yml) using the inventory file (hosts) to target the listed devices, executing the tasks defined in the playbook on those devices. Example 2-4 shows typical output after running the playbook.

Example 2-4 *Ansible Playbook Output*

```
PLAY [Show IP Interface Brief on Cisco IOS routers] **

TASK [Run 'show ip interface brief' command] ********
ok: [192.168.1.1]
ok: [192.168.2.1]

TASK [Display command output]
**********************************************
ok: [192.168.1.1] => {
    "output.stdout": [
       "Interface              IP-Address   OK? Method S
Protocol",
       "GigabitEthernet0/0   192.168.1.1  YES manual u
```

Technet24

```
        "GigabitEthernet0/1   unassigned   YES unset  a
    ]
}
ok: [192.168.2.1] => {
    "output.stdout": [
        "Interface              IP-Address   OK? Method
Protocol",
        "GigabitEthernet0/0   192.168.2.1  YES manual
        "GigabitEthernet0/1   192.168.1.2  YES manual
    ]
}

PLAY RECAP *****************************************
192.168.1.1                : ok=2    changed=0    unr
192.168.2.1                : ok=2    changed=0    unr
```

Terraform

Terraform is an open-source tool developed by HashiCorp for efficient infrastructure management. Like Ansible, it’s an IaC tool that automates the setup and management of networking resources in both on-premises and cloud environments. Terraform ensures consistent and repeatable network configurations, aligning with modern DevOps practices that emphasize automation in software and infrastructure management.

While both Ansible and Terraform are IaC tools, they have different methods

for managing infrastructure. Ansible uses SSH to directly access devices, making it ideal for configuration management on systems where SSH is the main interface. Terraform, on the other hand, uses plugins, which it calls *providers*, to interact with devices through APIs like RESTCONF or NETCONF. These providers translate Terraform configurations written in HashiCorp Configuration Language (HCL) into commands that specific technologies can understand and execute.

HCL is a configuration language designed to be both human readable and machine friendly, similar to JSON and XML. It's ideal for writing configuration files to manage infrastructure. HCL uses a structure similar to that of JSON, with braces for blocks and lists, but includes additional features for configuration management, such as comments and a more flexible syntax. Understanding HCL, along with other data structures like YAML, is crucial for effectively using tools like Terraform and Ansible in network automation and IaC.

Figure 2-15 shows a typical Terraform workflow for Cisco IOS XE devices that involves creating a configuration file using HCL, which specifies the desired settings for a device. The Terraform provider for Cisco IOS XE is then used to communicate with the device using RESTCONF or NETCONF protocols. Terraform commands such as **init**, **plan**, and **apply** are used to initialize the working directory, create an execution plan, and apply the changes to reach the desired state of configuration.

Technet24

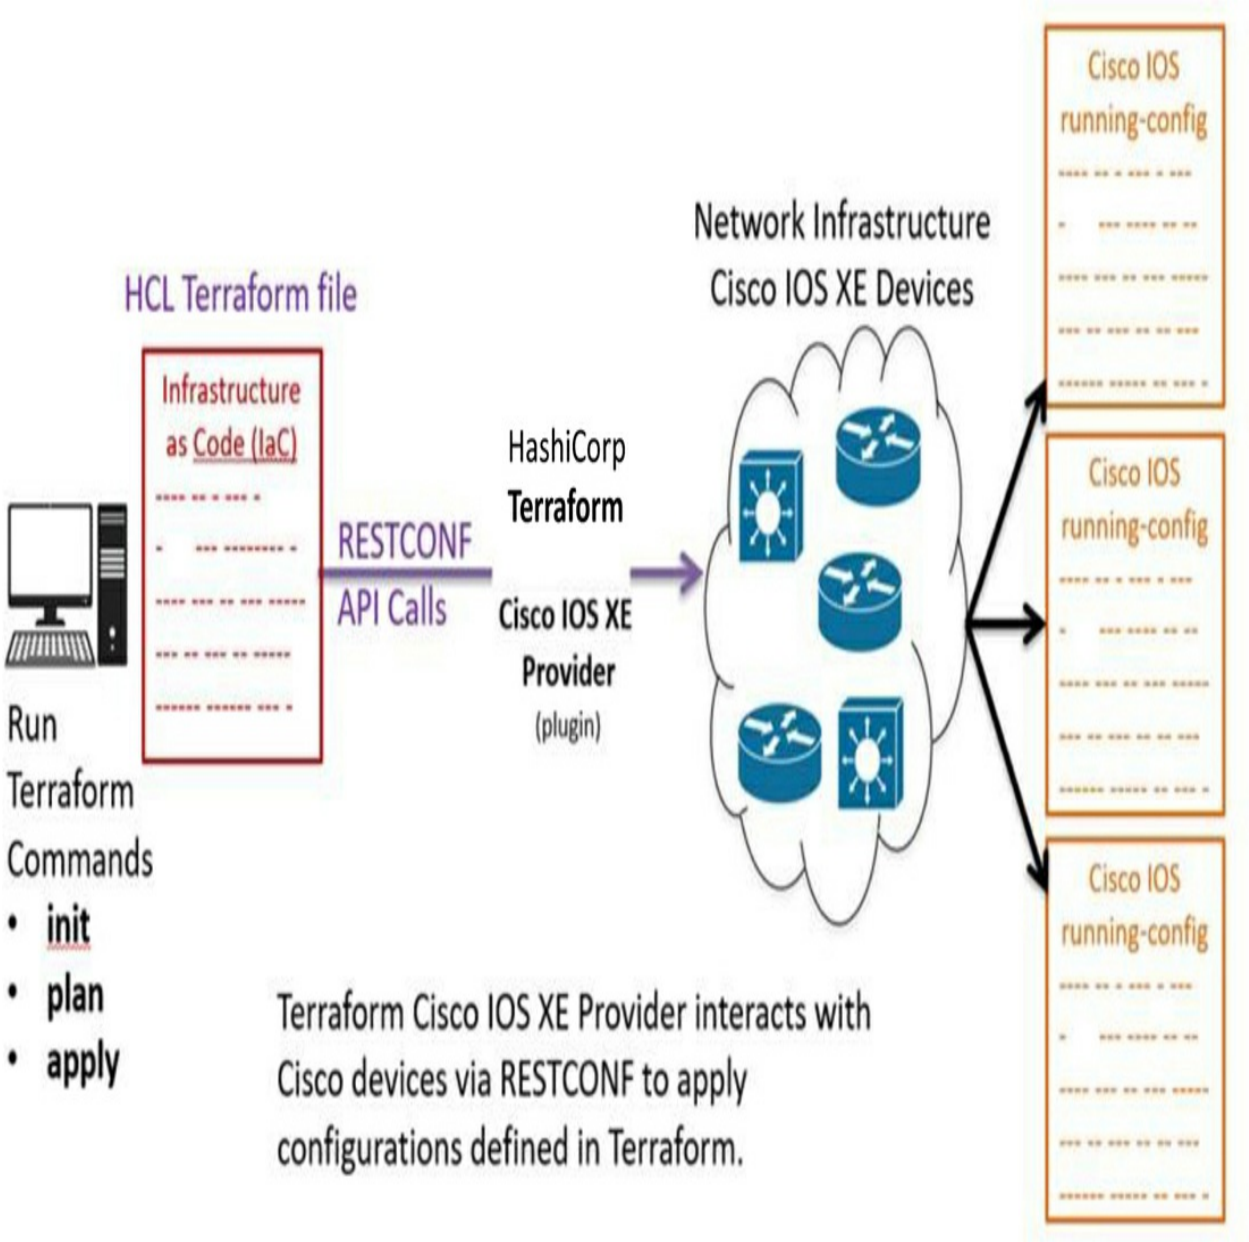

Figure 2-15 *Terraform Workflow*

Example 2-5 shows a simplified HCL configuration file stored in a Terraform directory on the administrator's PC. It first specifies Cisco IOS XE as the Terraform provider. Then, in the resource section, it defines the objects to be managed by Terraform—specifically the GigabitEthernet1 interface of the 192.168.10.1 IOS XE device. Finally, the output section requests the status of

the interface.

Example 2-5 *HCL Configuration File*

```
# Assuming there is a provider that supports Cisco IOS
provider "cisco_ios_xe" {
  host     = "192.168.10.1"
  username = "admin"
  password = "adminpass"
  secure   = true  # Assuming secure implies HTTPS/RES
}

# Resource configuration for an interface
resource "cisco_ios_xe_interface" "gigabit1" {
  name        = "GigabitEthernet1"
  description = "Uplink Interface"
  enabled     = true
}

# Output the operational status of the interface
output "interface_status" {
  value = cisco_ios_xe_interface.gigabit1.operational_
  description = "Operational status of the GigabitEthe
}
```

To use Terraform with Cisco IOS XE devices, install and configure the Cisco IOS XE Terraform provider plugin on your PC. This plugin uses

Technet24

RESTCONF or NETCONF protocols to communicate with Cisco devices. Enable these protocols on the target IOS XE device by entering the global configuration command **restconf** or **netconf**. In Example 2-6, both protocols are enabled on Router1.

Example 2-6 *Enabling RESTCONF and NETCONF on a Cisco IOS XE Router*

```
Router1> enable
Router1# configure terminal
Router1(config)# restconf
Router1(config)# netconf-yang
Router1(config)# end
Router1# write memory
```

Now you are ready to run the Terraform commands **init**, **plan**, and **apply**. Example 2-7 shows output from the **terraform init** command, which initializes the Terraform working directory and downloads the Cisco IOS XE Terraform provider plugin.

Example 2-7 *Output from the* **terraform init** *Command*

```
admin@pc:~/terraform_project$ terraform init
Initializing the backend...

Initializing provider plugins...
```

```
- Finding cisco_ios_xe versions matching ">= 1.0.0"...
- Installing cisco_ios_xe v1.0.0...
- Installed cisco_ios_xe v1.0.0 (self-signed, key ID X

Terraform has been successfully initialized!

You may now begin working with Terraform. Try running
changes that are required for your infrastructure. All
now work.
```

As shown in Example 2-8, the **terraform plan** command creates an execution plan specifying what actions Terraform will take to achieve the desired state defined in the configuration file.

Example 2-8 *Output from the* **terraform plan** *Command*

```
admin@pc:~/terraform_project$ terraform plan
Refreshing Terraform state in-memory prior to plan...
The refreshed state will be used to calculate this pla
persisted to local or remote state storage.

------------------------------------------------------

An execution plan has been generated and is shown belo
Resource actions are indicated with the following symb
  + create
```

Technet24

```
Terraform will perform the following actions:

  # cisco_ios_xe_interface.gigabit1 will be created
  + resource "cisco_ios_xe_interface" "gigabit1" {
      + description       = "Uplink Interface"
      + enabled           = true
      + id                = (known after apply)
      + name              = "GigabitEthernet1"
      + operational_status = (known after apply)
    }

Plan: 1 to add, 0 to change, 0 to destroy.

------------------------------------------------------------

Note: You didn't specify an "-out" parameter to save t
can't guarantee that exactly these actions will be per
"terraform apply" is subsequently run.
```

As shown in Example 2-9, the **terraform apply** command applies the changes required to reach the desired state of the configuration. Terraform prompts for a **yes** confirmation before making any changes. Terraform communicates with the device using RESTCONF or NETCONF through the API to handle the setup and reports back that the interface is up.

Example 2-9 *Output from the* **terraform apply** *Command*

```
admin@pc:~/terraform_project$ terraform apply
Terraform will perform the following actions:

  # cisco_ios_xe_interface.gigabit1 will be created
  + resource "cisco_ios_xe_interface" "gigabit1" {
      + description       = "Uplink Interface"
      + enabled           = true
      + id                = (known after apply)
      + name              = "GigabitEthernet1"
      + operational_status = (known after apply)
    }

Plan: 1 to add, 0 to change, 0 to destroy.

Do you want to perform these actions?
  Terraform will perform the actions described above.
  Only 'yes' will be accepted to approve.

  Enter a value: yes

cisco_ios_xe_interface.gigabit1: Creating...
cisco_ios_xe_interface.gigabit1: Creation complete aft
[id=GigabitEthernet1]

Apply complete! Resources: 1 added, 0 changed, 0 destr

Outputs:
```

Technet24

```
interface_status = "up"
```

Ansible and Terraform Comparison

Table 2-1 summarizes a few of the most common ideas about Ansible and Terraform.

Table 2-1 *Comparing Ansible and Terraform*

Attribute	Ansible	Terraform
File type for actions or config	Playbook	Configuration
Protocols to network device	SSH, NETCONF	RESTCONF, NETCONF, APIs
Execution mode	Agentless	Client/server model
Push or pull model	Push (by default) Pull with Ansible-pull program	Push
Modeling language	YAML	HashiCorp Configuration Language (HCL)

Explore DevNet Learning Labs

If you are interested in expanding your knowledge beyond the requirements

for the CCNA certification, Cisco DevNet Learning Labs offer free resources for exploring topics related to Cisco products. Users can access DevNet at developer.cisco.com. The platform provides beginner modules for learning and exploration of various networking topics, including Ansible and Terraform.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Cisco Network Academy: Enterprise Networking, Security, and Automation v7	14
Cisco Network Academy: ENSA Supplemental	1
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	22
	24

Technet24

Day 1. Data Formats, REST, and AI

CCNA 200-301 Exam Topics

• 6.4 Explain AI (generative and predictive) and machine learning in network operations

• 6.5 Describe characteristics of REST-based APIs (authentication types, CRUD, HTTP verbs, and data encoding)

• 6.7 Recognize components of JSON-encoded data

Key Topics

Today we review key concepts in modern network operations, starting with the main characteristics of the data formats JSON, XML, and YAML. We then review the characteristics of REST-based APIs, focusing on authentication, CRUD operations, HTTP verbs, and data encoding. Finally, we explore the roles of AI and machine learning in enhancing network operations.

Data Formats

Data formats provide a way to store and exchange data in a structured format. These are some common data formats used in network automation and programmability:

• JavaScript Object Notation (JSON)

• Extensible Markup Language (XML)

• YAML Ain’t Markup Language (YAML)

Table 1-1 compares the main purposes and common uses of these data formats.

Table 1-1 *Data Format Comparison*

Data Format	Origin/Definition	Central Purpose	Common Use
JSON	JavaScript (JS) language; RFC 8259	General data modeling and serialization	REST APIs
XML	World Wide Web Consortium (W3C.org)	Data-focused text markup, which makes data modeling possible	REST APIs, web pages
YAML	YAML.org	General data modeling	Ansible

Technet24

Each data format has specific characteristics:

• Syntax, which includes the types of brackets used (such as [ ], ( ), and { }), the use of whitespace, indentation, quotes, commas, and more.

• How objects such as characters, strings, lists, and arrays are represented.

• How key/value pairs are represented. The key, which is usually on the left side, identifies or describes the data. The value on the right is the data itself and can be a character, a string, a number, a list, or another type of data.

The data format that is selected depends on the format used by the application, tool, or script involved. However, the CCNA exam focuses specifically on the interpretation of JSON data.

JSON Data Format

JSON is a human-readable data format used by applications for storing, transferring, and reading data. It is easy to parse and can be used with most modern programming languages, including Python.

Example 1-1 shows partial IOS output from the **show interface GigabitEthernet0/0/0** command on a router.

Example 1-1 *IOS Router Output*

```
  GigabitEthernet0/0/0 is up, line protocol is up (conne
```

```
    Description: Wide Area Network
    Internet address is 172.16.0.2/24
```

Example 1-2 shows how this information can be represented in JSON format.

Example 1-2 *JSON Output*

```
{
    "ietf-interfaces:interface": {
        "name": "GigabitEthernet0/0/0",
        "description": "Wide Area Network",
        "enabled": true,
        "ietf-ip:ipv4": {
            "address": [
                {
                    "ip": "172.16.0.2",
                    "netmask": "255.255.255.0"
                }
            ]
         }
    }
}
```

JSON Syntax Rules

JSON data is a collection of key/value pairs that follow these rules:

Technet24

• **Key/value pair:** One key/value pair

• **Key:** Text inside double quotes and before the colon that is used as the name that references a value

• **Value:** The item after the colon that represents the value of the key, which can be:

• **Text:** Listed in double quotes

• **Numeric:** Listed without quotes

• **Array:** A list of values enclosed in square brackets ([ ])

• **Object:** One or more key/value pairs enclosed in braces ({ })

• **Multiple Pairs:** When listing multiple key/value pairs, separate the pairs with a comma at the end of each pair (except the last one)

A list of IPv4 addresses might look as shown in Example 1-3. The key here is addresses. The value is an array. Inside the array, each item is a separate object, and those objects are separated by braces ({ }). The objects are two key/value pairs: an IPv4 address (ip) and a subnet mask (netmask) separated by a comma. The objects in the array are also separated by commas, with a comma following the closing brace for each object. After all the items in the array are defined, the array is closed with the right bracket.

Example 1-3 *JSON List of IPv4 Addresses*

```
{
  "addresses": [
    {
      "ip": "172.16.0.2",
      "netmask": "255.255.255.0"
    },
    {
      "ip": "172.16.0.3",
      "netmask": "255.255.255.0"
    },
    {
      "ip": "172.16.0.4",
      "netmask": "255.255.255.0"
    }
  ]
}
```

RESTful APIs

APIs exist to allow two programs to exchange data. Some APIs are for interprogram communications within a single operating system (OS). Other APIs are available to programs that run on other computers. These APIs must define the networking protocol. Many are based on REST.

REST is an architectural style for designing web service applications. A

Technet24

REST API is an API that works on top of HTTP. It defines a set of functions developers can use to perform requests and receive responses through HTTP, such as GET and POST. An API can be considered RESTful if it has the following features:

• **Client/server:** The client handles the front end, and the server handles the back end. Either can be replaced independently of the other.

• **Stateless:** No client data is stored on the server between requests. The session state is stored on the client.

• **Cacheable:** Clients can cache responses to improve performance.

RESTful Implementation

A RESTful web service is a collection of resources with four defined aspects:

• The data format supported by the web service, which is often JSON, XML, or YAML

• The set of operations supported by the web service using HTTP methods

• The API, which must be hypertext driven

• The base uniform resource identifier (URI) for the web service, such as http://example.com/resources

RESTful APIs use common HTTP methods, including POST, GET, PUT, PATCH, and DELETE. As shown in Table 1-2, these methods correspond to the RESTful operations create, read, update, and delete (or CRUD).

Table 1-2 *HTTP Methods and RESTful Operation*

HTTP Method	RESTful Operation
POST	Create
GET	Read
PUT/PATCH	Update
DELETE	Delete

REST Authentication Types

With the convenience of RESTful services comes the critical need for securing these communications to protect sensitive data and ensure authorized access. There are several types of REST-based authentication methods. Some of the most common methods include basic authentication, token authentication, API key authentication, and OAuth authentication. Let's review each:

- **Basic authentication:**

Technet24

• **Summary:** Sends a username and password in the Authorization header of each HTTP request.

• **Use case:** Simple and quick to implement, suitable for internal applications or basic services over HTTPS.

• **Advantages:** Easy to implement and use.

• **Disadvantages:** Less secure, as credentials are sent with every request and do not expire unless manually changed.

• **Access, or bearer, token authentication:**

• **Summary:** Issues a token after an initial authentication process. The token is sent in the Authorization header of each HTTP request.

• **Use case:** Suitable for scenarios where security is a priority, such as web and mobile applications and with APIs requiring user-specific access.

• **Advantages:** Tokens can be short-lived and scoped, reducing the risk of misuse.

• **Disadvantages:** Requires a token issuance and management process.

• **API key authentication:**

• **Summary:** Includes a unique API key in the request, either in the headers,

query parameters, or body.

• **Use case:** Common for public APIs, third-party service integrations, and simple rate-limiting and tracking.

• **Advantages:** Easy to generate, distribute, and manage.

• **Disadvantages:** Less secure if keys are exposed; usually static and may not expire unless manually configured.

• **OAuth authentication:**

• **Summary:** An authorization framework that allows third-party applications to access a user's resources without exposing their credentials. Uses an authorization code and exchanges it for an access token.

• **Use case:** Ideal for applications that require user-specific permissions and third-party integrations, such as social media logins, as well as access to user data across multiple services.

• **Advantages:** Highly secure, supports fine-grained permissions, and is scalable.

• **Disadvantages:** More complex to implement and requires understanding of OAuth flows and token management.

Each method has strengths and is suitable for different scenarios based on

Technet24

security requirements, implementation complexity, and the specific needs of the application.

RESTful API Requests

A RESTful API is requested by using a URI, which is a string of characters that identifies a specific network resource. As shown in Figure 1-1, a URI has two specializations:

• **Uniform resource name (URN):** Identifies only the namespace of the resource without reference to the protocol.

• **Uniform resource locator (URL):** Defines the network location of a specific resource on the network.

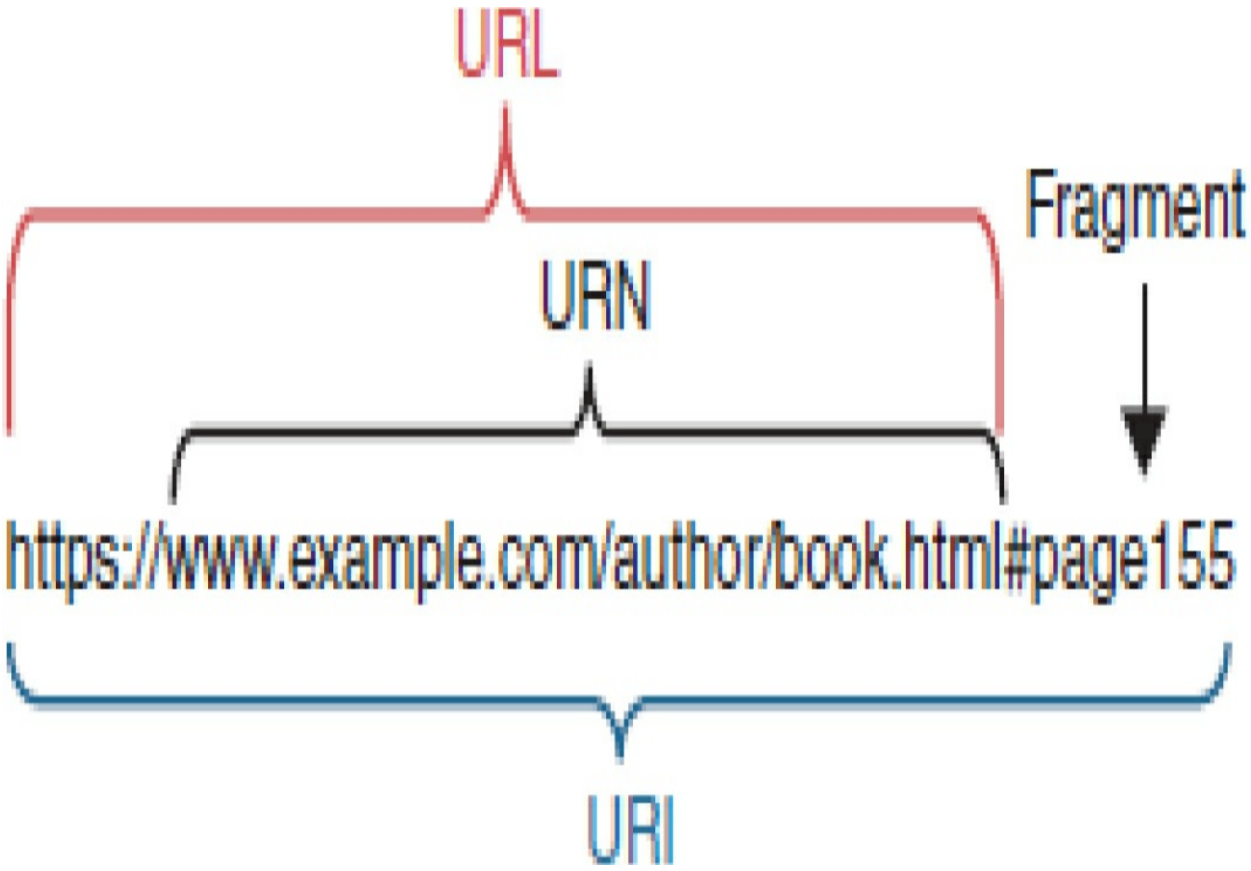

Figure 1-1 *Structure of a URI*

These are the parts of a URI, as shown in Figure 1-1:

• **Protocol/scheme:** HTTPS or another protocol, such as FTP, SFTP, mailto, or NNTP

• **Hostname:** In this case, www.example.com

• **Path and file name:** In this case, /author/book.html

• **Fragment:** In this case, #page155

A RESTful API request elicits a response from the API server. For example,

Technet24

the URI in Figure 1-2 is a properly formed GET request to the MapQuest API server for directions from San Jose to Monterey in JSON format.

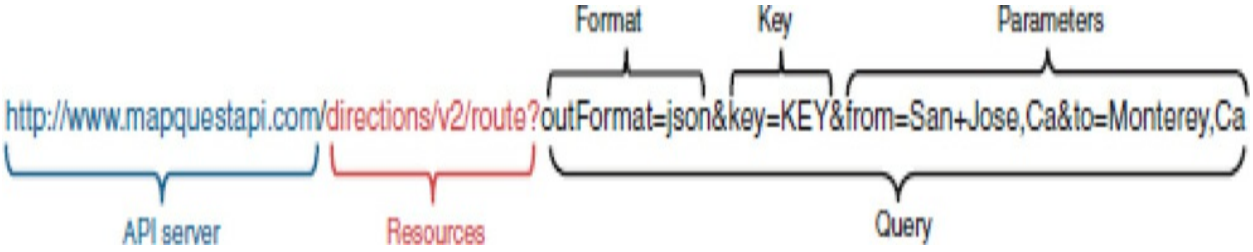

Figure 1-2 *RESTful API Request to the MapQuest API Server*

These are the different parts of the API request:

• **API server:** The URL for the server that answers REST requests.

• **Resources:** The API that is being requested.

• **Query:** The data format and information the client is requesting from the API service. Queries can include:

• **Format:** This is usually JSON but can be YAML or XML.

• **Key:** The key is for authorization, if required.

• **Parameters:** Parameters are used to send information pertaining to the request.

The beginning of the JSON payload delivered by the request in Figure 1-2 would look as shown in Example 1-4.

Example 1-4 *JSON Payload Received from an API Request*

```
{
  "route": {
    "hasTollRoad": false,
    "hasBridge": true,
    "boundingBox": {
      "lr": {
        "lng": -121.667061,
        "lat": 36.596809
      },
      "ul": {
        "lng": -121.897125,
        "lat": 37.335358
      }
    },
    "distance": 71.712,
    "hasTimedRestriction": false,
    "hasTunnel": false,
    "hasHighway": true,
    "computedWaypoints": [],
    "routeError": {
      "errorCode": -400,
      "message": ""
    },
(output omitted)
```

Technet24

AI and ML in IT Operations

Artificial intelligence for IT operations (AIOps) leverages AI and machine learning (ML) to optimize and automate various aspects of IT operations. The primary goal of AIOps is to improve the efficiency, reliability, and performance of IT systems while reducing manual interaction and human error through proven, pretested workflows or configurations.

Common AIOps use cases include the following:

• **Automation:** Reduces human intervention in routine tasks.

• **Monitoring and data analytics:** Continuously analyzes data to identify patterns and anomalies.

• **Predictive analysis:** Uses ML to predict issues by analyzing historical data.

• **Collaboration and communication:** Enhances information sharing and troubleshooting across IT operations.

• **Root cause analysis:** Identifies underlying causes of problems through deep visibility into IT environments.

Automation is a straightforward entry point for AIOps, improving workflows and procedures. Monitoring establishes a network baseline to detect abnormal behavior, aiding in identifying issues like data leaks. Predictive analytics is critical for forecasting and preventing issues during high-traffic events.

Collaborative chat tools facilitate real-time communication and quick resolution when network issues arise. Root cause analysis (RCA) ensures a documented and repeatable process for resolving outages and failures.

Generative AI for Network Configuration, Troubleshooting, and Simulation

When interacting with generative AI tools like ChatGPT, Claude, or Gemini, the phrasing of your questions significantly impacts the quality of the responses. Precise and clear language helps the models understand what you're asking and provide the specific information you need, while vague or ambiguous questions may lead to broader, less specific answers. These AI models also use context from previous interactions to tailor their responses to your preferences. For example, if you frequently ask networking questions and prefer Cisco IOS, the model will likely provide Cisco IOS-specific answers. This context-aware approach ensures that the information is relevant to your interests.

You can use one of these generative AI tools to help you with a configuration, troubleshoot output and error messages, or learn networking skills

ChatGPT Configuration Example

When prompting a generative AI tool, it's best to be as detailed as you can. For example, this prompt generates a detailed description of how to configure

Technet24

a trunk port:

Prompt:

Please give me the Cisco IOS commands to configure a trunk port. The port is gig 1/0/48, the allowed VLANs are 1, 10, 20, and 30, the native VLAN is 888, and we want to disable DTP (nonegotiate).

In my testing, ChatGPT was verbose, explaining every command. ChatGPT gave me the following response and configuration to the prompt above:

Response:

To configure a trunk port on a Cisco switch with the given requirements, you can use the following commands:

1. Enter global configuration mode.

2. Enter interface configuration mode for the desired port.

3. Set the port to trunk mode.

4. Allow the specified VLANs.

5. Set the native VLAN.

6. Disable DTP (Dynamic Trunking Protocol).

Here are the exact commands:

```
configure terminal
```

```
interface gigabitEthernet 1/0/48
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,10,20,30
```

```
switchport trunk native vlan 888
```

```
switchport nonegotiate
```

```
end
```

Here's a breakdown of each command:

• **configure terminal:** Enter global configuration mode.

• **interface gigabitEthernet 1/0/48:** Enter interface configuration mode for GigabitEthernet 1/0/48.

• **switchport mode trunk:** Set the port to trunk mode.

• **switchport trunk allowed vlan 1,10,20,30:** Allow VLANs 1, 10, 20, and 30 on the trunk.

• **switchport trunk native vlan 888:** Set VLAN 888 as the native VLAN.

Technet24

• **switchport nonegotiate:** Disable DTP to prevent the port from attempting to negotiate trunking.

After applying these commands, the port GigabitEthernet 1/0/48 will be configured as a trunk port with the specified VLAN settings and DTP disabled.

ChatGPT Troubleshooting Example

You can paste in output from **show** commands and error messages to get help from ChatGPT. The tool can analyze and decode complex error outputs, providing clear explanations and actionable steps to resolve issues. For example, this prompt is asking for help with a VLAN mismatch error message:

Prompt:

I have two switches, access-1 and access-2, that are interconnected by a trunk link that was just configured on FastEthernet 0/1 on both switches. I am receiving the following error message.

%CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch discovered on FastEthernet0/1 (1), with access-2 FastEthernet0/1 (254).

Can you tell me what the problem might be?

In testing, ChatGPT provided me with extensive detail. The response shown

here is abbreviated:

Response:

<output omitted>

The error message %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch discovered on FastEthernet0/1 (1), with access-2 FastEthernet0/1 (254) indicates that there is a native VLAN mismatch on the trunk link between FastEthernet0/1 on access-1 and FastEthernet0/1 on access-2.

Problem:

• On access-1, the native VLAN on FastEthernet0/1 is VLAN 1.

• On access-2, the native VLAN on FastEthernet0/1 is VLAN 254.

Solution:

To resolve this issue, you need to ensure that the native VLAN is the same on both ends of the trunk link.

ChatGPT as a Network Simulator

On of the most powerful features of generative AI tools for learners trying to hone their networking skills is its ability to act as a network simulator. The following prompt gives ChatGPT a topology and then asks for it to explain

Technet24

the addressing scheme:

Prompt:

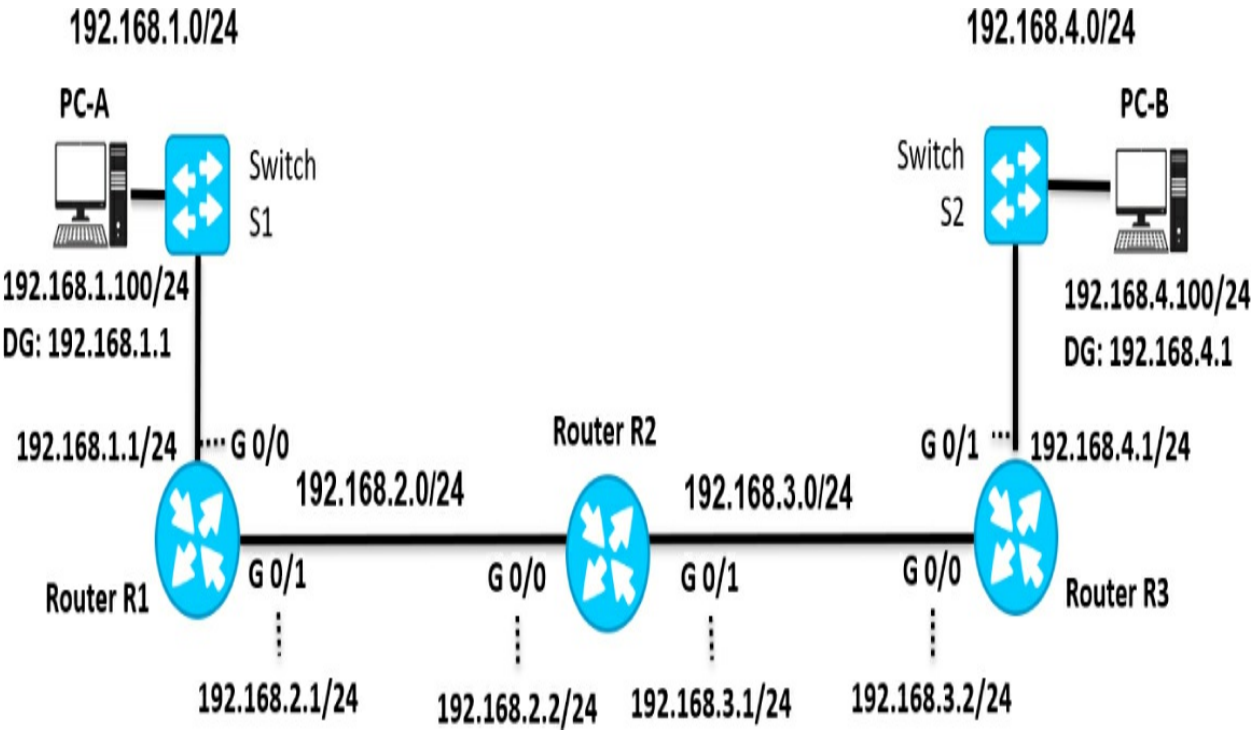

Please provide me a detailed explanation of the addressing scheme for the devices in this topology.

As you can see, you can copy and paste in a topology and then prompt ChatGPT with what you want.

Understanding the addressing scheme is important so that you can correct any errors before proceeding. ChatGPT will most likely provide a lot of detail. Your main task is to validate that the devices, interfaces, and addressing are all correct. During my testing, ChatGPT didn't always match up the correct IPv4 address to the correct interface. In this case, though, ChatGPT gave me

an accurate addressing scheme:

Response:

The network topology provided consists of three routers, two switches, and two PCs. Below is the detailed explanation of the IP addressing scheme for the devices in this topology:

Router R1:

• **Interface G0/0:**

• IP Address: 192.168.1.1

• Subnet Mask: 255.255.255.0 (or /24)

• **Interface G0/1:**

• IP Address: 192.168.2.1

• Subnet Mask: 255.255.255.0 (or /24)

Router R2:

• **Interface G0/0:**

• IP Address: 192.168.2.2

Technet24

• Subnet Mask: 255.255.255.0 (or /24)

• **Interface G0/1:**

• IP Address: 192.168.3.1

• Subnet Mask: 255.255.255.0 (or /24)

Router R3:

• **Interface G0/0:**

• IP Address: 192.168.3.2

• Subnet Mask: 255.255.255.0 (or /24)

• **Interface G0/1:**

• IP Address: 192.168.4.1

• Subnet Mask: 255.255.255.0 (or /24)

Switch S1:

• This switch connects PC-A and Router R1. It operates within the 192.168.1.0/24 network. The switch itself does not require an IP address for basic switching operations, but it may have one for management purposes which is not specified here.

Switch S2:

• This switch connects PC-B and Router R3. It operates within the 192.168.4.0/24 network. Similar to S1, it does not require an IP address for basic switching operations but may have one for management purposes which is not specified here.

PC-A:

• IP Address: 192.168.1.100

• Subnet Mask: 255.255.255.0 (or /24)

• Default Gateway: 192.168.1.1

PC-B:

• IP Address: 192.168.4.100

• Subnet Mask: 255.255.255.0 (or /24)

• Default Gateway: 192.168.4.1

<output omitted>

This address scheme matches the topology. So I didn't have to correct ChatGPT by prompting something like *"You have made some mistakes"* and then list the mistakes. Make sure you validate all the addressing and interface

Technet24

designations. If there are any errors, correct ChatGPT before proceeding.

When you are satisfied that ChatGPT understands the topology, you can proceed with prompting it to be a simulator. For example, this prompt worked really well in my testing:

Prompt:

I want you to act as a Cisco IOS terminal. I will type in commands and you will only reply with the terminal output inside one unique code block, and nothing else. Do not write explanations. Do not add commands or configurations. Do not type commands unless I instruct you to do so. When I tell you something in English, I will do so by putting text inside curly braces {like this}. I am using the topology I attached previously.

The devices are configured with the following:

• *The routers are configured with hostnames R1, R2 and R3, IPv4 addressing, and activated interfaces. No routing has been configured. The routers are Cisco ISR 4331 running IOS XE Software, Version 16.*

• *The switches are configured with hostnames S1 and S2. No other configuration has been done. The switches are Cisco 2960 running IOS Software, Version 12.*

• *The PCs are configured with IPv4 addressing and are running Windows 11.*

I will add the hostname before the command followed by a colon ":". For example, R1: show running-config or PC-A: ipconfig.

ChatGPT will respond with something like this:

Sure, I understand. Please go ahead with your commands.

Now you are ready to ask it anything about the topology or proceed through a list of configuration requirements. For example, use the following prompt to configure the routers with OSPF:

Prompt:

{Configure the following list of commands on R1}

router ospf 1

router-id 1.1.1.1

network 192.168.1.0 0.0.0.255 area 0

network 192.168.2.0 0.0.0.255 area 0”

ChatGPT should respond in a code window with only the configuration commands, acting as if it were Router1:

Response:

Technet24

`R1(config)#` **router ospf 1**

`R1(config-router)#` **router-id 1.1.1.1**

`R1(config-router)#` **network 192.168.1.0 0.0.0.255 area 0**

`R1(config-router)#` **network 192.168.2.0 0.0.0.255 area 0**

`R1(config-router)#` **end**

Prompt ChatGPT to configure the other routers. Then validate routing with **show ip route** and pings between PC-A and PC-B. Use other commands, such **as show ip ospf neighbors**, or configure IPv6 addressing and routing. However, understand that ChatGPT might begin to lose context and start to hallucinate. Regardless, this experiment shows you the emerging power of using generative AI as a Cisco IOS simulator.

Study Resources

For today's exam topics, refer to the following resources for more study.

Resource	Module or Chapter
Enterprise Networking, Security, and Automation	14
CCNA 200-301 Official Cert Guide, Volume 2, 2nd edition	22
	17
	18
	19

Technet24

Exam Day

Today is your opportunity to prove that you have what it takes to manage a small enterprise branch network. Just 120 minutes stand between you and your CCNA certification. Use the following information to focus on the process details for the day of your CCNA exam.

What You Need for the Exam

Write down the exam location, date, exam time, exam center phone number, and proctor's name. Here's what will happen when you get to the exam center:

• You must have two forms of ID that include a photo and signature, such as a driver's license, passport, or military identification. In addition, the test center admission process requires the capture of a digital photo and digital signature.

• The test proctor will take you through the agreement and set up your testing station after you have signed the agreement.

• The test proctor will give you a sheet for scratch paper or a dry erase pad. Do not take these out of the room.

• The testing center will store any personal items while you take the exam. It is best to bring only what you need.

• You will be monitored during the entire exam.

Note

You can take the exam at a remote location such as your home or office by following the setup guidelines required by the testing center prior to and during your session. The setup can be quite extensive but is an option for those who would prefer an online proctored exam. Visit https://www.cisco.com/c/en/us/training-events/training-certifications/online-exam-proctoring.html for more information or search for “Cisco online exam proctoring” if the URL is no longer available.

What You Should Receive After Completion

When you complete the exam, you will see an immediate electronic response about whether you passed or failed. The proctor will give you a certified score report with the following important information:

• The minimum passing score, your score on the exam, and a breakout displaying your percentage for each general exam topic

• Identification information required to track your certification

Do not lose your certified examination score report.

Technet24

Summary

Your state of mind is a key factor in your success on the CCNA exam. If you know the details of the exam topics and the details of the exam process, you can begin the exam with confidence and focus. Arrive early to the exam. Bring earplugs just in case a testing neighbor has a bad cough or any loud nervous habits. Do not let an extremely difficult or specific question impede your progress. You cannot return to questions on the exam that you have already answered, so answer each question confidently and move on.

Post-Exam Information

Signing up for and actually taking the CCNA exam is no small accomplishment. Many network engineers have avoided certification exams for years. The following sections discuss your options after exam day.

Receiving Your Certificate

If you passed the exam, check the CertMetrics website after a few days to access your downloadable PDF certificate and logos:

https://cp.certmetrics.com/cisco/en/credentials/status

You might want to print your certificate, frame it, and put it on a wall. You should absolutely add the CCNA logo to your profile on your social media accounts—most importantly on LinkedIn. Also update your profile to reflect the CCNA under your list of certifications and degrees.

Your CCNA is valid for 3 years. To keep your certificate valid, you must either pass the CCNA exam again or pass a more advanced certification before the end of the 3-year period. Alternatively, you can visit the Cisco Continuing Education Program website, https://ce.cisco.com, for flexible options to recertify by completing a variety of preapproved eligible activities.

Determining Career Options

As mentioned earlier, after you pass the CCNA exam, be sure to add your

Technet24

CCNA certification to your resume and social media accounts (especially LinkedIn). However, avoid placing certifications after your name, as doing so can appear presumptuous. Although achieving certifications requires a lot of effort, it's much less effort than the years of effort required for PhD and other advanced degrees that would be listed after your name. Instead, list your certifications and degrees in a section titled "Education and Certifications."

The most important factor in career development is progressing toward your career goals. You may not secure your desired title or job immediately after graduation. However, mastering relevant skills in your current position and building a network of contacts can lead you to your dream job. Career growth is a gradual process that requires patience and perseverance.

The CCNA certification is part of an overall professional skill set that must be continually enhanced to advance in an IT career. Holding a CCNA certificate demonstrates discipline and commitment to a rigorous course of study and professional goals. While having a CCNA alone may not secure a job, it can place you ahead of other candidates. To supplement the CCNA certification on your resume, highlight any networking skills related to the CCNA in the job and skills descriptions.

Examining Certification Options

Passing the CCNA exam is no easy task, but it is the starting point for more advanced Cisco certifications, such as CCNP-level exams. When you log in

to the online certification tracking tool, be sure to search for specific information about certifications you can achieve with CCNA as the base.

If You Did Not Pass the Exam

If you did not pass your first attempt at the CCNA, you must wait at least 5 calendar days after the day of the exam to retest. Stay motivated and sign up to take the exam again within 30 days of your first attempt. The score report outlines your weaknesses. Find a study group and use The Cisco Learning Network (http://learningnetwork.cisco.com) online community to get help with the topics on which you scored poorly.

If you are familiar with the general concepts tested on the CCNA exam, focus on taking practice exams and memorizing the small details that make the exam so difficult. If you are a Cisco Networking Academy alumnus, you have access to the curriculum. Packet Tracer also provides an excellent network simulator. Consider your first attempt as a formal practice exam and excellent preparation for passing on your second attempt.

Summary

Whether you display your certificate and update your resume or prepare to conquer the exam on your second attempt, remember to marvel at the innovation and creativity behind each concept you learn. The ability of our society to continually improve communication will keep you learning, discovering, and employed for a lifetime.

Technet24

CCNA Countdown Calendar

The lines after the countdown number allow you to add the actual calendar days for reference.

31 ____ Networking Models, Devices, and Components	**30** ____ Ethernet Switching	**29** ____ Switch Configuration Basics	**28** ____ IPv4 Addressing	**27** ____ IPv6 Addressing	**26** ____ VLAN and Trunking Concepts and Configuration	**25** ____ STP
24 ____ EtherChannel and HSRP	**23** ____ DHCP and DNS	**22** ____ Wireless Concepts	**21** ____ WLAN Configuration	**20** ____ LAN Security and Device Hardening	**19** ____ Basic Routing Concepts	**18** ____ Basic Router Configuration
17 ____ The Routing Table	**16** ____ Inter-VLAN Routing	**15** ____ Static and Default Route Configuration	**14** ____ OSPF Operation	**13** ____ Single-Area OSPF Implementation	**12** ____ Fine-Tuning and Troubleshooting OSPF	**11** ____ Network Security Concepts
10 ____ ACL Concepts	**9** ____ ACL Implementation	**8** ____ NAT	**7** ____ WAN, VPN, and IPsec	**6** ____ QoS	**5** ____ CDP and LLDP	**4** ____ Device Monitoring, Management, and Maintenance
3 ____ Cloud, Virtualization, and SDN	**2** ____ Cisco SD-Access, Ansible	**1** ____ Data Formats, REST, and AI	**EXAM DAY** Time _____ Location _____ _____			

Technet24

Exam Checklist

CCNA Checklist Days 31–21

Checkbox	Date	Objective
		Schedule to take the CCNA 200-301 exam at http://www.vue.com.
		Take at least one practice CCNA exam.
		Create a diagram of the layered models.
		Describe the details of sending an email from source to destination. Use a topology with several routers and switches.
		Describe the CSMA/CD process to someone who knows nothing about networking.
		Design a set of requirements to configure a basic switched network that includes SSH for remote access. Implement your design and verify the configurations.
		Show someone the MAC address on their smartphone or another connected device. Explain the purpose of the MAC address and the meaning of each part.
		Describe the structure and operation of IPv4. List and describe the uses for the various types of IPv4 addresses.
		Develop several VLSM addressing schemes with various host requirements and implement them in a lab or simulator.
		Describe the structure and operation of IPv6. List and describe the uses for the various types of IPv6 addresses.
		Design a set of requirements to configure a two-switch network with VLANs and trunking.

Technet24

		List and describe the uses for various types of VLANs.
		Describe how trunking works and the impact of Dynamic Trunking Protocol.
		Design a set of requirements to configure a three-switch network with trunking and VLANs. Implement your design and verify the configurations.
		Describe the process of STP convergence.
		Compare the varieties of STP.
		Describe the operations allowed for each PVST port state.
		Explain the difference between PVST+ and Rapid PVST+.
		Describe the benefits of EtherChannel. Compare the two EtherChannel protocols.
		Design a set of requirements to implement a two-switch topology with EtherChannel. Review implementation issues by changing the configuration parameters.
		Describe how HSPR provides default gateway redundancy.
		Explain to a friend how wireless networks operate.
		Practice configuring a wireless router. Use a simulator like Packet Tracer to practice configuring a WLC.
		Read and review Days 31–21 in this book.

Checkbox	Date	Objective
		Take at least two practice CCNA exams.
		Describe the various methods that a router can use to learn about and share knowledge of remote networks.
		Design a set of requirements to configure a three-router network with both IPv4 and IPv6 addressing using only static and default routes. Implement your design and verify the configurations.
		Describe how a router uses the routing table to determine the best path to the destination.
		Design and configure a one-router, two-switch topology for router-on-a-stick inter-VLAN routing.
		Describe the types of OSPF packets and how they are used to reach the full state of OSPF convergence.
		Describe the ways to modify OSPFv2, including redistributing a default route, modifying timers, and controlling the DR/BDR election.
		Read and review Days 20–14 in this book.

CCNA Checklist Days 13–7

Technet24

Checkbox	Date	Objective
		Take an additional CCNA practice exam.
		Design a set of requirements to configure a three-router network with IPv4 addressing and OSPFv2. Implement your design and verify the configurations.
		Design a set of requirements to configure a two-router, two-switch network with IPv4 addressing and inter-VLAN routing. Include default routing and OSPFv2. Implement your design and verify the configurations.
		Describe basic security threats and the methods used to mitigate them.
		Describe how access control lists work and the various types used by Cisco IOS Software.
		Search the Internet for various scenarios to practice designing and implementing ACLs. Most of the study resources have excellent examples.
		Design a set of requirements to implement a routed network that includes basic device security, IPv4 and IPv6 addressing, VLANs, DHCP, NAT, ACLs, and routing. Implement your design and verify the configurations.
		Design a set of requirements to implement DHCP service on a router. Dual-stack the design to include IPv4 and IPv6. Implement your design and verify the configurations.
		Design a set of requirements to implement NAT on a router. Include static, dynamic, and PAT considerations. Implement your design and verify the configurations.
		Define common WAN terminology.
		Compare various WAN connection options.
		Describe the characteristics of GRE.
		Read and review Days 13–7 in this book.

CCNA Checklist Days 6–1

Technet24

Checkbox	Date	Objective
		Describe to a friend how QoS prioritizes Netflix streaming data over web browsing data.
		Compare and contrast the implementations of CDP and LLDP.
		Describe the Cisco IOS file system and the process for backing up and restoring files.
		Explain the basic SNMP, NTP, and syslog operation.
		Design a set of requirements to implement SNMP, NTP, and syslog in a two-router, one-server topology.
		Describe to a friend the concept of cloud computing. Include a discussion of virtualization.
		Describe software-defined networking.
		Describe the fabric of the network infrastructure and the difference between overlay and underlay.
		Describe to a friend how data formats are used to store and exchange information between systems.
		Describe the JSON data format, including how brackets, braces, and commas are used to distinguish key/value pairs, arrays, and objects.
		Describe the structure of a properly formatted RESTful API request and its authentication methods.
		Compare the configuration management tools Ansible and Terraform.
		Read and review Days 6–1 in this book.
		Visit the testing center and talk with the proctor at least 2 days before the exam.
		Eat a decent meal, watch a good movie, and get a good night's rest before the exam.