




I'm not robot



Continue

Ccna 200 301 book volume 2 pdf

CCNA 200-301 Official Cert Guide, Volume 2Wendell Odom, Best-selling Cisco Press author, expert trainer, and Cisco Certified Internetwork Expert (CCIE No. 1624)***Best-selling Cisco Press author, expert trainer, and Cisco Certified Internetwork Expert (CCIE No. 1624) productFormatCode=C32 productCategory=2 statusCode=5 isBuyable=true subType= path/ProductBean/courseSmarttrue ISBN-10: 1587147130 • ISBN-13: 978158714713597815871471359780135262733 ©2020 • Cisco Press • Cloth Bound with Access Card, 624 pp Published 12/22/2019 • Out of Stock • Net price: \$44.997 prodCategory : 2 statusCode : 5CCNA 200-301 Official Cert Guide enables students to succeed on the exam the first time and is the only self-study resource approved by Cisco.Best-selling author and expert instructor Wendell Odom shares preparation hints and test-taking tips, helping students identify areas of weakness and improve both conceptual knowledge and hands-on skills.Well regarded for its level of detail, study plans, assessment features, challenging review questions and exercises, video instruction, and hands-on labs, this official study guide helps students master the concepts and techniques that ensure exam success. FeaturesRevised edition of the #1 selling CCNA preparation self-study guideBook content is fully updated to align to the new CCNA 200-301 exam objectivesBooks and online materials are packed with features to help candidates master difficult testing methods on actual examsPractice tests contain scenario-based questions that closely mimic the difficulty of the actual examIn-depth expert explanations of all protocols, commands, and technologies on the CCNA examOnline ancillary materials such as lecture slides, instructor's notes, and test bank reinforce concepts discussed in this text Table of Contents Part I IP Access Control Lists 1 Introduction to TCP/IP Transport and Applications 2 Basic IPv4 Access Control Lists 3 Advanced IPv4 Access Control Lists Part II Security Services 4 Security Architectures 5 Securing Network Devices 6 Implementing Switch Port Security 7 Implementing DHCP 8 DHCP Snooping and ARP Inspection Part III IP Services 9 Device Management Protocols 10 Network Address Translation 11 Quality of Service (QoS) 12 Miscellaneous IP Services Part IV Network Architecture 13 LAN Architecture 14 WAN Architecture 15 Cloud Architecture Part V Network Automation 16 Introduction to Controller-Based Networking 17 Cisco Software-Defined Access (SDA) 18 Understanding REST and JSON 19 Understanding Ansible, Puppet, and Chef Part VI Final Review 20 Final Review Part VII Appendices Appendix A Numeric Reference Tables Appendix B CCNA 200-301 Volume 2 Exam Updates Appendix C Answers to the "Do I Know This Already?" Quizzes Glossary Wendell Odom, CCIE No. 1624 Emeritus, has been in the networking industry since 1981. He has worked as a network engineer, consultant, systems engineer, instructor, and course developer; he currently works writing and creating certification study tools. This book is his 29th edition of some product for Pearson, and he is the author of all editions of the CCNA Cert Guides about Routing and Switching from Cisco Press. He has written books about topics from networking basics, certification guides throughout the years for CCENT, CCNA R&S, CCNA DC, CCNP ROUTE, CCNP QoS, and CCIE R&S. He maintains study tools, links to his blogs, and other resources at www.certskills.com. CCNA 200-301 Official Cert Guide enables students to succeed on the exam the first time and is the only self-study resource approved by Cisco.Best-selling author and expert instructor Wendell Odom shares preparation hints and test-taking tips, helping students identify areas of weakness and improve both conceptual knowledge and hands-on skills.Well regarded for its level of detail, study plans, assessment features, challenging review questions and exercises, video instruction, and hands-on labs, this official study guide helps students master the concepts and techniques that ensure exam success. Revised edition of the #1 selling CCNA preparation self-study guide Book content is fully updated to align to the new CCNA 200-301 exam objectives Books and online materials are packed with features to help candidates master difficult testing methods on actual exams Practice tests contain scenario-based questions that closely mimic the difficulty of the actual exam In-depth expert explanations of all protocols, commands, and technologies on the CCNA exam Online ancillary materials such as lecture slides, instructor's notes, and test bank reinforce concepts discussed in this text Assessment, review, and practice for the CCNA 200-301 exam Revised edition of the #1 selling CCNA preparation self-study guide Book content is fully updated to align to the new CCNA 200-301 exam objectives Books and online materials are packed with features to help candidates master difficult testing methods on actual exams Practice tests contain scenario-based questions that closely mimic the difficulty of the actual exam In-depth expert explanations of all protocols, commands, and technologies on the CCNA examCoursesCCNA 200-301 Official Cert Guide, Volume 2, (OASIS)OdomISBN-10: 0135262704 • ISBN-13: 9780135262702 ©2020 • Digital Access Code • AvailableMore infoGive your students choices! PearsonChocies products are designed to give your students more value and flexibility by letting them choose from a variety of text and media formats to best match their learning style and their budget.Pearson Higher Education offers special pricing when you choose to package your text with other student resources. If you're interested in creating a cost-saving package for your students, see the Packages tab.CCNA 200-301 Official Cert Guide, Volume 2OdomISBN-10: 0135262739 • ISBN-13: 9780135262733©2020 • ePub, 624 pp • AvailableMore info | Students, buy accessPearson eText for CCNA 200-301 Official Cert Guide, Volume 2 -- Instant AccessOdomISBN-10: 0137459955 • ISBN-13: 9780137459957©2020 • Electronic Book • Estimated Availability: 02/01/2021More infoThis product is a member of the following series. Click on the series name to see the full list of products in the series. Pearson Higher Education offers special pricing when you choose to package your text with other student resources. If you're interested in creating a cost-saving package for your students contact your Pearson Higher Education representative. Nobody is smarter than you when it comes to reaching your students. You know how to convey knowledge in a way that is relevant and reliable to your class. It's the reason you always get the best out of them. And when it comes to planning your curriculum, you know which course materials express the information in the way that's most consistent with your teaching. That's why we give you the option to personalize your course material using just the Pearson content you select. Take only the most applicable parts of your favorite materials and combine them in any order you want. You can even integrate your own writing if you wish. It's fast, it's easy and fewer course materials help minimize costs for your students. For more information: www.pearsonlearningsolutions.com/higher-education Or download our brochure (PDF). Explore our course catalogues and see how you can customize your own textbooks.Custom LibraryOur library is vast, and it's all at your fingertips. Create a custom book by selecting content from any of our course-specific collections. Here, you'll find chapters from Pearson titles, carefully-selected third-party content with copyright clearance, and pedagogy. Once you're satisfied with your customized book, you will have a print-on-demand book that can be purchased by students in the same way they purchase other course material.Custom PublicationsBrowse through our list of published titles. These books are examples of original manuscripts created in partnership with local Custom Field Editors. They have been authored by instructors at specific campuses, but are readily available for adoption.Pearson Learning Solutions offers a broad range of courses and custom solutions for web-enhanced, blended and online learning. Our course content is developed by a team of respected subject matter experts and experienced eLearning instructional designers. All course content is designed around specific learning objectives. For more information: www.pearsonlearningsolutions.com/higher-education/customizable-online-courseware Or download our brochure (PDF). REPORT THIS PDF [] Download CCNA 200-301 Official CERT Guide Volume 2 PDF for free from drive.google.com using the direct download link given below. CCNA 200-301 Official CERT Guide Volume 2 Trust the best-selling Official Cert Guide series from Cisco Press to help you learn, prepare, and practice for exam success. They are built with the objective of providing assessment, review, and practice to help ensure you are fully prepared for your Cisco Certified Network Associate certification exam. Download complete CCNA 200-301 Official CERT Guide Volume 2 pdf file or read online for free using link provided below. REPORT THIS!f the download link of CCNA 200-301 Official CERT Guide Volume 2 PDF is not working or you feel any other problem with it, please REPORT IT by selecting the appropriate action such as copyright material / promotion content / link is broken etc. If CCNA 200-301 Official CERT Guide Volume 2 is a copyright material we will not be providing its PDF or any source for downloading at any cost. xSorry to interruptCCSE Error CCNA 200-301, Volume 2 Official Cert Guide In addition to the wealth of updated content, this new edition includes a series of free hands-on-exercises to help you master several real-world configuration activities. These exercises can be performed on the CCNA 200-301 Network Simulator Lite, Volume 2 software included for free on the companion website that accompanies this book. This software, which simulates the experience of working on actual Cisco routers and switches, contains the following 13 free lab exercises, covering ACL topics in Part I: 1. ACL 1.2. ACL 1.2. 3. ACL 1.4. ACL V 6. ACL V 7. ACL Analysis I 8. Named ACL I 9. Named ACL II 10. Named ACL III 11. Standard ACL Configuration Scenario 12. Extended ACL I Configuration Scenario 13. Extended ACL II Configuration Scenario If you are interested in exploring more hands-on labs and practice configuration and troubleshooting with more router and switch commands, go to www.pearsoncertification.com/networksimulator for demos and to review the latest products for sale. CCNA 200-301 Network Simulator Lite, Volume 2 system requirements: Windows system requirements (minimum): Mac system requirements (minimum): -W indows 10 (32/64-bit), Windows 8.1 (32/64-bit), or Windows 7 (32/64-bit) • macOS 10.15, 10.14, 10.13, 10.12, or 10.11 • 1-gigahertz (GHz) or faster 32-bit (x86) or 64-bit (x64) processor • Intel core Duo 1.83 GHz • 512 MB RAM (1 GB recommended) • 1 GB RAM (32-bit) or 2 GB RAM (64-bit) • 1.5 GB hard disk space •1. 6 GB available hard disk space (32-bit) or 20 GB (64-bit) • 32-bit color depth at 1024 x 768 resolution •D irectX 9 graphics device with WDDM 1.0 or higher driver • Adobe Acrobat Reader version 8 and above • Adobe Acrobat Reader version 8 and above CCNA 200-301 Official Cert Guide, Volume 2 WENDELL ODOM, CCIE No. 1624 Emeritus Cisco Press ii CCNA 200-301 Official Cert Guide, Volume 2 CCNA 200-301 Official Cert Guide, Volume 2 Wendell Odom Copyright © 2020 Pearson Education, Inc. Published by: Cisco Press All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review. ScoutAutomatedPrintCode Library of Congress Control Number: 2019949625 ISBN-13: 978-1-58714-713-5 ISBN-10: 1-58714-713-0 Warning and Disclaimer This book is designed to provide information about the Cisco CCNA 200-301 exam. 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. 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. Microsoft and/or its respective suppliers make no representations about the suitability of the information contained in the documents and related graphics published as part of the services for any purpose. All such documents and related graphics are provided "as is" without warranty of any kind. Microsoft and/ or its respective suppliers hereby disclaim all warranties and conditions with regard to this information, including all warranties and conditions of merchantability, whether express, implied or statutory, fitness for a particular purpose, title and non-infringement. In no event shall Microsoft and/or its respective suppliers be liable for any special, indirect or consequential damages or any damages whatsoever resulting from loss of use, data or profits, whether in an action of contract, negligence or other tortious action, arising out of or in connection with the use or performance of information available from the services. The documents and related graphics contained herein could include technical inaccuracies or typographical errors. Changes are periodically added to the information herein. Microsoft and/or its respective suppliers may make improvements and/or changes in the product(s) and/or the program(s) described herein at any time. Partial screenshots may be viewed in full within the software version specified. Microsoft® and Windows® are registered trademarks of the Microsoft Corporation in the U.S.A. and other countries. Screenshots and icons reprinted with permission from the Microsoft Corporation. This book is not sponsored or endorsed by or affiliated with the Microsoft Corporation. iii 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 or (800) 382-3419. For government sales inquiries, please contact For questions about sales outside the U.S., please contact 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 alert it to better suit your needs, you can contact us through email at Please make sure to include the book title and ISBN in your message. We greatly appreciate your assistance. Editor-in-Chief: Mark Taub Technical Editor: Elan Beer Business Operation Manager, Cisco Press: Ronald Fligge Editorial Assistant: Cindy Teeters Director, ITP Product Management: Brett Bartow Cover Designer: Chuti Prasertth Managing Editor: Sandra Schroeder Composition: Tricia Bronkella Development Editor: Christopher Cleveland Indexer: Ken Johnson Senior Project Editor: Tonya Simpson Proofreader: Debbie Williams Copy Editor: Chuck Hutchinson iv CCNA 200-301 Official Cert Guide, Volume 2 About the Author Wendell Odom, CCIE No. 1624 Emeritus, has been in the networking industry since 1981. He has worked as a network engineer, consultant, systems engineer, instructor, and course developer; he currently works writing and creating certification study tools. This book is his 29th edition of some product for Pearson, and he is the author of all editions of the CCNA Cert Guides about Routing and Switching from Cisco Press. He has written books about topics from networking basics, certification guides throughout the years for CCENT, CCNA R&S, CCNA DC, CCNP ROUTE, CCNP QoS, and CCIE R&S. He maintains study tools, links to his blogs, and other resources at www.certskills.com. v About the Contributing Author David Hucabay, CCIE No. 4594, CWNE No. 292, is a network engineer for University of Kentucky Healthcare. He has been authoring Cisco Press titles for 20 years, with a focus on wireless and LAN switching topics. David has bachelor of science and master of science degrees in electrical engineering. He lives in Kentucky with his wife, Marci, and two daughters. About the Technical Reviewer Elan Beer, CCIE No. 1837, is a senior consultant and Cisco instructor specializing in data center architecture and multiprotocol network design. For the past 27 years, Elan has designed networks and trained thousands of industry experts in data center architecture, routing, and switching. Elan has been instrumental in large-scale professional service efforts designing and troubleshooting internetworks, performing data center and network audits, and assisting clients with their short- and long-term design objectives. Elan has a global perspective of network architectures via his international clientele. Elan has used his expertise to design and troubleshoot data centers and internetworks in Malaysia, North America, Europe, Australia, Africa, China, and the Middle East. Most recently, Elan has been focused on data center design, configuration, and troubleshooting as well as service provider technologies. In 1993, Elan was among the first to obtain the Cisco Certified System Instructor (CCSI) certification, and in 1996, he was among the first to attain the Cisco System highest technical certification, the Cisco Certified Internetworking Expert. Since then, Elan has been involved in numerous large-scale data center and telecommunications networking projects worldwide. vi CCNA 200-301 Official Cert Guide, Volume 2 Acknowledgments Brett Bartow continues to be the backbone of the Cisco Press brand, guiding the entire author team through the big transition in 2019–2020 with all the changes Cisco introduced to its certifications. Simply the best! Thanks for all you do, Brett! Dave Hucabay teamed up again to write this book, contributing one chapter here to go along with his four chapters in the CCNA Volume 1 book. It's such a joy to review his work and see such polished material from the first draft. It's been a joy to work with such a consummate professional—thanks, Dave! Chris Cleveland developed the book—again—and made it much better—again—and did it with more juggling than ever before, I think. Five months, roughly 50 technology chapters and another 50 other book elements, and countless online elements, all done with apparent ease. Kudos to Chris, yet again! I so look forward to reading Elan Beer's tech edits of the chapters. That may seem strange to her, but Elan has truly amazing technical editing skills. His insights range from the details of technology, to the mind of the new learner, to wording and clarity, to holes in networking logic as compared to the wording, to tiny typos that impact the meaning. Thanks again Elan for improving the chapters so much! Tonya Simpson managed this book, along with the CCNA Volume 1 book, all in that same compressed timeframe again. As usual, on both projects, Tonya has kept the production processes rolling along and getting through the idiosyncrasies of the content. Thanks for shepherding the book through the wild again, Tonya! As always, thanks to the production team that works with Tonya. From fixing all my grammar and passive-voice sentences to pulling the design and layout together, they do it all; thanks for putting it all together and making it look easy. And Tonya got to juggle two books of mine at the same time (again)—thanks for managing the whole production process again. Mike Tanamachi, illustrator and mind reader, did a great job on the figures again. Mike came through again with some beautiful finished products. Thanks again, Mike. I could not have made the timeline for this book without Chris Burns of Certskills Professional. Chris owns much of the PTP question support and administration process, works on the labs we put on my blog, and then catches anything I need to toss over my shoulder so I can focus on the books. Chris, you are the man! A special thank you to you readers who write in with suggestions and possible errors, and especially those of you who post online at the Cisco Learning Network and at my blog (). Without question, the comments I receive directly and overhear by participating at CLN made this edition a better book. Thanks to my wonderful wife, Kris, who helps make this sometimes challenging work lifestyle a breeze. I love walking this journey with you, doll. Thanks to my daughter Hannah, who actually helped a bit with the book this summer before heading off to college (go Jackets!). And thanks to Jesus Christ, Lord of everything in my life. vii Contents at a Glance Introduction xxvii Part I IP Access Control Lists Chapter 1 Introduction to TCP/IP Transport and Applications Chapter 2 Basic IPv4 Access Control Lists Chapter 3 Advanced IPv4 Access Control Lists Part I Review 3 24 44 64 Part II Security Services Chapter 4 Security Architectures Chapter 5 Securing Network Devices Chapter 6 Implementing Switch Port Security Chapter 7 Implementing DHCP Chapter 8 DHCP Snooping and ARP Inspection Part II Review 67 68 86 106 122 144 168 Part III IP Services Chapter 9 Device Management Protocols Chapter 10 Network Address Translation Chapter 11 Quality of Service (QoS) Chapter 12 Miscellaneous IP Services Part III Review 171 172 202 226 254 284 Part IV Network Architecture Chapter 13 LAN Architecture Chapter 14 WAN Architecture Chapter 15 Cloud Architecture Part V Network Automation Chapter 16 Introduction to Controller-Based Networking Chapter 17 Cisco Software-Defined Access (SDA) 382 356 4 viii CCNA 200-301 Official Cert Guide, Volume 2 Chapter 18 Understanding REST and JSON 406 Chapter 19 Understanding Ansible, Puppet, and Chef 428 Part V Review 444 Part VI Final Review 447 Chapter 20 Final Review 448 Part VII Appendices 467 Appendix A Numeric Reference Tables Appendix B CCNA 200-301, Volume 2 Exam Updates Appendix C Answers to the "Do I Know This Already?" Quizzes 469 476 478 Glossary 494 Index 530 Online Appendices Appendix D Topics from Previous Editions Appendix E Practice for Chapter 2: Basic IPv4 Access Control Lists Appendix F Previous Edition ICND1 Chapter 35: Managing IOS Files Appendix G Exam Topics Cross-Reference ix Reader Services To access additional content for this book, simply register your product. To start the registration process, go to www.ciscopress.com/register and log in or create an account". Enter the product ISBN 9781587147135 and click Submit. After the process is complete, you will find any available bonus content under Registered Products. *Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product. x CCNA 200-301 Official Cert Guide, Volume 2 Icons Used in This Book Access Point PC Laptop Server IP Phone Router Switch Layer 3 Switch Hub Bridge Cable (Various) Serial Line Virtual Circuit Ethernet WAN SDN Controller Network Cloud vSwitch Cable Modem IPS ASA Wireless Firewall DSLAM 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. ix Contents Introduction xxvii Part I IP Access Control Lists 3 Chapter 1 Introduction to TCP/IP Transport and Applications "Do I Know This Already?" Quiz Foundation Topics 24 24 26 IPv4 Access Control List Basics 26 ACL Location and Direction 26 Matching Packets 27 Taking Action When a Match Occurs Types of IP ACLs 28 28 xii CCNA 200-301 Official Cert Guide, Volume 2 Standard Numbered IPv4 ACLs 29 List Logic with IP ACLs 29 Matching Logic and Command Syntax 31 Matching the Exact IP Address 31 Matching a Subset of the Address with Wildcards 31 Binary Wildcard Masks 33 Finding the Right Wildcard Mask to Match a Subnet 33 Matching Any/All Addresses 34 Implementing Standard IP ACLs 34 Standard Numbered ACL Example 1 35 Standard Numbered ACL Example 2 36 Troubleshooting and Verification Tips 38 Practice Applying Standard IP ACLs 39 Practice Building access-list Commands 39 Reverse Engineering from ACL to Address Range 40 Chapter Review 41 Chapter 3 Advanced IPv4 Access Control Lists 44 "Do I Know This Already?" Quiz 44 Foundation Topics 46 Extended Numbered IP Access Control Lists 46 Matching the Protocol, Source IP, and Destination IP 46 Matching TCP and UDP Port Numbers 48 Extended IP ACL Configuration 51 Extended IP Access Lists: Example 1 51 Extended IP Access Lists: Example 2 53 Practice Building access-list Commands 54 Named ACLs and ACL Default IP Access Lists 54 Editing ACLs Using Sequence Numbers 56 Numbered ACL Configuration Versus Named ACL Configuration 58 ACL Implementation Considerations 59 Additional Reading on ACLs Chapter Review 61 Part I Review 64 60 xiii Part II Security Services 67 Chapter 4 Security Architectures 68 "Do I Know This Already?" Quiz Foundation Topics 68 70 Security Terminology 70 Common Security Threats 72 Attacks That Spoof Addresses Denial-of-Service Attacks 72 73 Reflection and Amplification Attacks Man-in-the-Middle Attacks 76 Address Spoofing Attack Summary Reconnaissance Attacks 77 Buffer Overflow Attacks 78 75 77 Malware 78 Human Vulnerabilities 79 Password Vulnerabilities 80 Password Alternatives 80 Controlling and Monitoring User Access 82 Developing a Security Program to Educate Users 83 Chapter Review 84 Chapter 5 Securing Network Devices 86 "Do I Know This Already?" Quiz 86 Foundation Topics 88 Securing IOS Passwords 88 Encrypting Older IOS Passwords with service password-encryption 89 Encoding the Enable Passwords with Hashes 90 Interactions Between Enable Password and Enable Secret Making the Enable Secret Truly Secret with a Hash Improved Hashes for Cisco's Enable Secret Encoding the Passwords for Local Usernames Controlling Password Attacks with ACLs 95 92 94 91 xiv CCNA 200-301 Official Cert Guide, Volume 2 Firewalls and Intrusion Prevention Systems Traditional Firewalls 95 96 Security Zones 97 Intrusion Prevention Systems (IPS) 99 Cisco Next-Generation Firewalls Cisco Next-Generation IPS 100 102 Chapter Review 103 Chapter 6 Implementing Switch Port Security 106 "Do I Know This Already?" Quiz 106 Foundation Topics 108 Port Security Concepts and Configuration 108 Configuring Port Security 109 Verifying Port Security 112 Port Security MAC Addresses 113 Port Security Violation Modes 114 Port Security Shutdown Mode 115 Port Security Protect and Restrict Modes 117 Chapter Review 119 Chapter 7 Implementing DHCP 122 "Do I Know This Already?" Quiz 122 Foundation Topics 124 Dynamic Host Configuration Protocol 124 DHCP Concepts 125 Supporting DHCP for Remote Subnets with DHCP Relay 126 Information Stored at the DHCP Server 128 Configuring DHCP Features on Routers and Switches Configuring DHCP Relay 130 Configuring a Switch as DHCP Client 132 Identifying Host IPv4 Settings 133 Host Settings for IPv4 133 Host IP Settings on Windows 134 129 xv Host IP Settings on macOS Host IP Settings on Linux 136 138 Chapter Review 140 Chapter 8 DHCP Snooping and ARP Inspection 144 "Do I Know This Already?" Quiz 144 Foundation Topics 146 DHCP Snooping 146 DHCP Snooping Concepts 146 A Sample Attack: A Spurious DHCP Server 147 DHCP Snooping Logic 148 Filtering DISCOVER Messages Based on MAC Address Filtering Messages that Release IP Addresses 150 DHCP Snooping Configuration 152 Configuring DHCP Snooping on a Layer 2 Switch Limiting DHCP Message Rates 152 154 DHCP Snooping Configuration Summary 155 Dynamic ARP Inspection 156 DAI Concepts 156 Review of Normal IP ARP 156 Gratuitous ARP as an Attack Vector Dynamic ARP Inspection Logic 157 158 Dynamic ARP Inspection Configuration 160 Configuring ARP Inspection on a Layer 2 Switch Limiting DAI Message Rates 163 Configuring Optional DAI Message Checks 164 IP ARP Inspection Configuration Summary 165 Chapter Review 166 Part II Review 168 160 150 xvi CCNA 200-301 Official Cert Guide, Volume 2 Part III IP Services 171 Chapter 9 Device Management Protocols "Do I Know This Already?" Quiz Foundation Topics 172 172 174 System Message Logging (Syslog) 174 Sending Messages in Real Time to Current Users Storing Log Messages for Later Review 174 175 Log Message Format 176 Log Message Severity Levels 177 Configuring and Verifying System Logging The debug Command and Log Messages Network Time Protocol (NTP) 178 180 181. Setting the Time and Timezone 182 Basic NTP Configuration 183 NTP Reference Clock and Stratum 185 Redundant NTP Configuration 186 NTP Using a Loopback Interface for Better Availability Analyzing Topology Using CDP and LLDP 190 Examining Information Learned by CDP 190 Configuring and Verifying CDP 193 Examining Information Learned by LLDP Configuring and Verifying LLDP Chapter Review 194 197 199 Chapter 10 Network Address Translation 202 "Do I Know This Already?" Quiz Foundation Topics 202 204 Perspectives on IPv4 Address Scalability CIDR 204 205 Private Addressing 206 Network Address Translation Concepts Static NAT 207 208 Dynamic NAT 210 Overloading NAT with Port Address Translation 211 188 xvii NAT Configuration and Troubleshooting 213 Static NAT Configuration 213 Dynamic NAT Configuration 215 Dynamic NAT Verification 217 NAT Overload (PAT) Configuration 219 NAT Troubleshooting 222 Chapter Review 223 Chapter 11 Quality of Service (QoS) 226 "Do I Know This Already?" Quiz Foundation Topics 226 228 Introduction to QoS 228 QoS: Managing Bandwidth, Delay, Jitter, and Loss Types of Traffic 228 229 Data Applications 229 Voice and Video Applications QoS as Mentioned in This Book QoS on Switches and Routers Classification and Marking 230 232 233 233 Classification Basics 233 Matching (Classification) Basics 234 Classification on Routers with ACLs and NBAR Marking IP DSCP and Ethernet CoS 236 Marking the IP Header 237 Marking the Ethernet 802.1Q Header Other Marking Fields 237 238 Defining Trust Boundaries 238 DiffServ Suggested Marking Values 239 Expedited Forwarding (EF) 240 Assured Forwarding (AF) Class Selector (CS) 240 241 Guidelines for DSCP Marking Values 241 235 xviii CCNA 200-301 Official Cert Guide, Volume 2 Queuing 242 Round-Robin Scheduling (Prioritization) 243 Low Latency Queuing 243 A Prioritization Strategy for Data, Voice, and Video 245 Shaping and Policing 245 Policing 246 Where to Use Policing Shaping 246 248 Setting a Good Shaping Time Interval for Voice and Video Congestion Avoidance 250 TCP Windowing Basics 250 Congestion Avoidance Topics Chapter Review 251 252 Chapter 12 Miscellaneous IP Services 254 "Do I Know This Already?" Quiz Foundation Topics 254 256 First Hop Redundancy Protocol 256 The Need for Redundancy in Networks 257 The Need for a First Hop Redundancy Protocol 259 The Three Solutions for First-Hop Redundancy 260 HSRP Concepts 261 HSRP Failover 261 HSRP Load Balancing 262 Simple Network Management Protocol 263 SNMP Variable Reading and Writing: SNMP Get and Set 264 SNMP Notifications: Traps and Informs 265 The Management Information Base Securing SNMP 266 267 FTP and TFTP 268 Managing Cisco IOS Images with FTP/TFTP 268 The IOS File System 268 Upgrading IOS Images 270 Copying a New IOS Image to a Local IOS File System Using TFTP 271 249 xix Verifying IOS Code Integrity with MD5 Copying Images with FTP The FTP and TFTP Protocols FTP Protocol Basics 273 273 275 FTP Active and Passive Modes FTP over TLS (FTP Secure) 276 278 TFTP Protocol Basics 279 Part IV Chapter Review 280 Part III Review 284 284 Network Architecture 287 Chapter 13 LAN Architecture 288 "Do I Know This Already?" Quiz Foundation Topics 288 290 Analyzing Campus LAN Topologies 290 Two-Tier Campus Design (Collapsed Core) The Two-Tier Campus Design 290 290 Topology Terminology Seen Within a Two-Tier Design 290 Part-III Review Campus Design (Core) 293 Topology Design Terminology Small Office/Home Office 295 Power over Ethernet (PoE) 297 PoE Basics 295 297 PoE Operation 298 PoE and LAN Design Chapter Review 299 300 Chapter 14 WAN Architecture 302 "Do I Know This Already?" Quiz Foundation Topics Metro Ethernet 302 304 304 Metro Ethernet Physical Design and Topology Ethernet WAN Services and Topologies 305 306 Ethernet Line Service (Point-to-Point) 307 Ethernet LAN Service (Full Mesh) 308 Ethernet Tree Service (Hub and Spoke) 309 291 x CCNA 200-301 Official Cert Guide, Volume 2 Layer 3 Design Using Metro Ethernet 309 Layer 3 Design with E-Line Service 309 Layer 3 Design with E-LAN Service 311 Multiprotocol Label Switching (MPLS) 311 MPLS VPN Physical Design and Topology 313 MPLS and Quality of Service 314 Layer 3 with MPLS VPN 315 Internet VPNs 317 Internet Access 317 Digital Subscriber Line 318 Cable Internet 319 Wireless WAN (3G, 4G, LTE, 5G) Fiber (Ethernet) Internet Access 320 321 Internet VPN Fundamentals 321 Site-to-Site VPNs with IPsec 322 Remote Access VPNs with TLS VPN Comparisons Chapter Review 324 326 326 Chapter 15 Cloud Architecture 328 "Do I Know This Already?" Quiz Foundation Topics 328 330 Server Virtualization 330 Cisco Server Hardware 330 Server Virtualization Basics 331 Networking with Virtual Switches on a Virtualized Host The Physical Data Center Network 334 Workflow with a Virtualized Data Center Computing Services 336 Private Cloud (On-Premise) 337 Public Cloud 338 335 333 xxi Cloud and the "As a Service" Model Infrastructure as a Service Software as a Service 339 339 341. (Development) Platform as a Service 341 WAN Traffic Paths to Reach Cloud Services 342 Enterprise WAN Connections to Public Cloud 342 Accessing Public Cloud Services Using the Internet 342 Pros and Cons with Connecting to Public Cloud with Internet 343 Private WAN and Internet VPN Access to Public Cloud 344 Pros and Cons of Connecting to Cloud with Private WANs 345 Intercloud Exchanges 346 Summarizing the Pros and Cons of Public Cloud WAN Options 346 A Scenario: Branch Offices and the Public Cloud 347 Migrating Traffic Flows When Migrating to Email SaaS Branch Offices with Internet and Private WAN Part V Chapter Review 350 Part IV Review 352 Network Automation 355 Chapter 16 Introduction to Controller-Based Networking "Do I Know This Already?" Quiz Foundation Topics 356 357 358 SDN and Controller-Based Networks 358 The Data, Control, and Management Planes The Data Plane 358 359 The Control Plane 360 The Management Plane 361 Cisco Switch Data Plane Internals 361 Controllers and Software-Defined Architecture 362 Controllers and Centralized Control 363 The Southbound Interface 364 The Northbound Interface 365 349 347 xxi CCNA 200-301 Official Cert Guide, Volume 2 Software Defined Architecture Summary 367 Examples of Network Programmability and SDN OpenDaylight and OpenFlow 367 367 The OpenDaylight Controller 368 The Cisco Open SDN Controller (OS) 369 Cisco Application Centric Infrastructure (ACI) 369 ACI Physical Design: Spine and Leaf 370 ACI Operating Model with Intent-Based Networking Cisco APIC Enterprise Module APIC-EM Basics 371 373 373 APIC-EM Replacement 374 Summary of the SDN Examples 375 Comparing Traditional Versus Controller-Based Networks 375 How Automation Impacts Network Management 376 Comparing Traditional Networks with Controller-Based Networks 378 Chapter Review 379 Chapter 17 Cisco Software-Defined Access (SDA) "Do I Know This Already?" Quiz Foundation Topics 382 383 384 SDA Fabric, Underlay, and Overlay The SDA Underlay 384 386 Using Existing Gear for the SDA Underlay 386 Using New Gear for the SDA Underlay 387 The SDA Overlay 390 VXLAN Tunnels in the Overlay (Data Plane) 390 LISP for Overlay Discovery and Location (Control Plane) DNA Center and SDA Operation Cisco DNA Center 395 395 Cisco DNA Center and Scalable Groups 396 Issues with Traditional IP-Based Security SDA Security Based on User Groups 398 397 392 xxxii DNA Center as a Network Management Platform 400 DNA Center Similarities to Traditional Management Chapter Review 402 403 Chapter 18 Understanding REST and JSON "Do I Know This Already?" Quiz Foundation Topics REST-Based APIs 406 406 408 REST-Based (RESTful) APIs 408 Client/Server Architecture 409 Stateless Operation 410 Cacheable (or Not) 410 Background: Data and Variables Simple Variables 410 410 List and Dictionary Variables REST APIs and HTTP 411 413 Software CRUD Actions and HTTP Verbs 413 Using URIs with HTTP to Specify the Resource Example of REST API Call to DNA Center 417 Data Serialization and JSON 418 The Need for a Data Model with APIs Data Serialization Languages JSON 421 XML 421 YAML 419 421 422 Summary of Data Serialization 423 Interpreting JSON 423 Interpreting JSON Key/Value Pairs 423 Interpreting JSON Objects and Arrays Minified and Beautified JSON Chapter Review 427 428 424 414 xiv CCNA 200-301 Official Cert Guide, Volume 2 Chapter 19 Understanding Ansible, Puppet, and Chef 428 "Do I Know This Already?" Quiz 428 Foundation Topics 430 DevOps Configuration Challenges and Solutions 430 Configuration Drift 430 Centralized Configuration Files and Version Control 431 Configuration Monitoring and Enforcement 433 Configuration Provisioning 434 Configuration Templates and Variables 435 Files That Control Configuration Automation 437 Ansible, Puppet, and Chef Basics Ansible 438 438 Puppet 440 Chef 441 Summary of Configuration Management Tools 442 Chapter Review 442 Part V Review 444 Part VI Final Review 447 Chapter 20 Final Review 448 Advice About the Exam Event 448 Exam Event: Learn About Question Types 448 Exam Event: Think About Your Time Budget 450 Exam Event: A Sample Time-Check Method 451 Exam Event: One Week Away 451 Exam Event: 24 Hours Before the Exam 452 Exam Event: The Last 30 Minutes 452 Exam Event: Reserve the Hour After the Exam 453 Exam Review 454 Exam Review: Take Practice Exams Using the Practice CCNA Exams 454 455 Exam Review: Advice on How to Answer Exam Questions Exam Review: Additional Exams with the Premium Edition 456 457 xxv Exam Review: Find Knowledge Gaps 458 Exam Review: Practice Hands-On CLI Skills 460 Exam Exam Topics with CLI Skill Requirements 460 Exam Review: Self-Assessment Pitfalls 462 Exam Review: Adjustments for Your Second Attempt 463 Exam Review: Other Study Tasks 464 Final Thoughts 464 Part VII Appendices 467 Appendix A Numeric Reference Tables 469 Appendix B CCNA 200-301, Volume 2 Exam Updates 476 Appendix C Answers to the "Do I Know This Already?" Quizzes 478 Glossary Index 494 530 Online Appendices Appendix D Topics from Previous Editions Appendix E Practice for Chapter 2: Basic IPv4 Access Control Lists Appendix F Previous Edition ICND1 Chapter 35: Managing IOS Files Appendix G Exam Topics Cross-Reference Appendix H Study Planner This page intentionally left blank xxvii Introduction About Cisco Certifications and CCNA Congratulations! If you're reading far enough to look at this book's Introduction, you've probably already decided to go for your Cisco certification, and the CCNA certification is the one place to begin that journey. If you want to succeed as a technical person in the networking industry at all, you need to know Cisco. Cisco has a ridiculously high market share in the router and switch marketplace, with more than 80 percent market share in some markets. In many geographies and markets around the world, networking equals Cisco. If you want to be taken seriously as a network engineer, Cisco certification makes perfect sense. NOTE This book discusses part of the content Cisco includes in the CCNA 200-301 exam, with the CCNA 200-301 Official Cert Guide, Volume 1, covering the rest. You will need both the Volume 1 and Volume 2 books to have all the content necessary for the exam. The first few pages of this Introduction explain the core features of the Cisco Career Certification program, of which the Cisco Certified Network Associate (CCNA) serves as the foundation for all the other certifications in the program. This section begins with a comparison of the old to the new certifications due to some huge program changes in 2019. It then gives the key features of CCNA, how to get it, and what's on the exam. The Big Changes to Cisco Certifications in 2019 Cisco announced sweeping changes to its career certification program around mid-year 2019. Because so many of you will have read and heard about the old versions of the CCNA certification, this Introduction begins with a few comparisons between the old and new CCNA as well as some of the other Cisco career certifications. First, consider the Cisco career certifications before 2019, as shown in Figure I-1. At that time, Cisco offered 10 separate CCNA certifications in different technology tracks. Cisco also had eight Professional-level (CCNP, or Cisco Certified Network Professional) certifications. xxviii CCNA 200-301 Official Cert Guide, Volume 2 Collaboration Data Center Routing & Wireless Switching Security Service Provider CCIE Collaboration Data Center Routing & Wireless Switching Security Service Provider Cloud Service Provider Cloud CCNP Collaboration Data Center Routing & Wireless Switching Security Cyber Industrial Ops CCNA Figure I-1. Old Cisco Certification Silo Concepts Why so many? Cisco began with one track—Routing and Switching—back in 1998. Over time, Cisco identified more and more technology areas that had grown to have enough content to justify another set of CCNA and CCNP certifications on those topics, so Cisco added more tracks. Many of those also grew to support expert-level topics with CCIE (Cisco Certified Internetwork Expert). In 2019, Cisco consolidated the tracks and moved the topics around quite a bit, as shown in Figure I-2. Collaboration Data Center Enterprise Security Service Provider CCIE Collaboration Data Center Enterprise Security Service Provider CCNP CCNA Figure I-2 New Cisco Certification Tracks and Structure All the tracks now begin with the content in the one remaining CCNA certification. For CCNP, you now have a choice of five technology areas for your next steps, as shown in Figure I-2. (Note that Cisco replaced "Routing and Switching" with "Enterprise.") xxix Cisco made the following changes with the 2019 announcements: CCENT: Retired the only entry-level certification (CCENT, or Cisco Certified Entry Network Technician), with no replacement. CCNA: Retired all the CCNA certifications except what was then known as "CCNA Routing and Switching," which became simply "CCNA." CCNP: Consolidated the professional-level (CCNP) certifications to five tracks, including merging CCNP Routing and Switching and CCNP Wireless into CCNP Enterprise. CCIE: Achieved better alignment with CCNP tracks through the consolidations. Cisco needed to move many of the individual exam topics from one exam to another because of the number of changes. For instance, Cisco announced the retirement of all the associate certifications—nine CCNA certifications plus the CCDA (Design Associate) certification—but those technologies didn't disappear! Cisco just moved the topics around to different exams in different certifications. (Note that Cisco later announced that CCNA Cyber Ops would remain, and not be retired, with details to be announced.) Consider wireless LANs as an example. The 2019 announcements retired both CCNA Wireless and CCNP Wireless as certifications. Some of the old CCNA Wireless topics landed in the new CCNA, whereas others landed in the two CCNP Enterprise exams about wireless LANs. For those of you who want to learn more about the transition, check out my blog () and look for posts in the News category from around June 2019. Now on to the details about CCNA as it exists starting in 2019! How to Get Your CCNA Certification As you saw in Figure I-2, all career certification paths now begin with CCNA. So how do you get it? Today, you have one and only one option to achieve CCNA certification: Take and pass one exam: the Cisco 200-301 CCNA exam. To take the 200-301 exam, or any Cisco exam, you will use the services of Pearson VUE (vue.com). The process works something like this: 1. Establish a login at (or use your existing login). 2. Register for, schedule a time and place, and pay for the Cisco 200-301 exam, all from the VUE website. 3. Take the exam at the VUE testing center. 4. You will receive a notice of your score, and whether you passed, before you leave the testing center. Types of Questions on the CCNA 200-301 Exam The Cisco CCNA and CCNP exams all follow the same general format, with these types of questions: ■ Multiple-choice, single-answer ■ Multiple-choice, multiple-answer xx CCNA 200-301 Official Cert Guide, Volume 2 ■ Testlet (one scenario with multiple multiple-choice questions) ■ Drag-and-drop ■ Simulated lab (sim) ■ Simlet Although the first four types of questions in the list should be somewhat familiar to you from other tests in school, the last two are more common to IT tests and Cisco exams in particular. Both use a network simulator to ask questions so that you control and use simulated Cisco devices. In particular: Sim questions: You see a network topology and lab scenario, and can access the devices. Your job is to fix a problem with the configuration. Simlet questions: This style combines sim and testlet question formats. As with a sim question, you see a network topology and lab scenario, and can access the devices. However, as with a testlet, you also see multiple multiple-choice questions. Instead of changing or fixing the

configuration, you answer questions about the current state of the network. These two question styles with the simulator give Cisco the ability to test your configuration skills with sim questions, and your verification and troubleshooting skills with simlet questions. Before taking the test, learn the exam user interface by watching some videos Cisco provides about the interface. To find the videos, just go to www.cisco.com and search for "Cisco Certification Exam Tutorial Videos." CCNA 200-301 Exam Content, Per Cisco Ever since I was in grade school, whenever a teacher announced that we were having a test soon, someone would always ask, "What's on the test?" We all want to know, and we all want to study what matters and avoid studying what doesn't matter. Cisco tells the world the topics on each of its exams. Cisco wants the public to know the variety of topics and get an idea about the kinds of knowledge and skills required for each topic for every Cisco certification exam. To find the details, go to www.cisco.com/go/certifications, look for the CCNA page, and navigate until you see the exam topics. This book also lists those same exam topics in several places. From one perspective, every chapter sets about to explain a small set of exam topics, so each chapter begins with the list of exam topics covered in that chapter. However, you might want to also see the exam topics in one place, so Appendix G, "Exam Topics Cross-Reference," lists all the exam topics. You may want to download Appendix G in PDF form and keep it handy. The appendix lists the exam topics with two different cross-references: ■ A list of exam topics and the chapter(s) that covers each topic ■ A list of chapters and the exam topics covered in each chapter xxxi Exam Topic Verbs and Depth Reading and understanding the exam topics, especially deciding the depth of skills required for each exam topic, require some thought. Each exam topic mentions the name of some technology, but it also lists a verb that implies the depth to which you must master the topic. The primary exam topics each list one or more verbs that describe the skill level required. For example, consider the following exam topic: Configure and verify IPv4 addressing and subnetting Note that this one exam topic has two verbs (configure and verify). Per this exam topic, you should be able to not only configure IPv4 addresses and subnets, but you also should understand them well enough to verify that the configuration works. In contrast, the following exam topic asks you to describe a technology but does not ask you to configure it: Describe the purpose of first hop redundancy protocol The describe verb tells you to be ready to describe whatever a "first hop redundancy protocol" is. That exam topic also implies that you do not then need to be ready to configure or verify any first hop redundancy protocols (HSRP, VRRP, and GLBP). Finally, note that the configure and verify exam topics imply that you should be able to describe and explain and otherwise master the concepts so that you understand what you have configured. The earlier "Configure and verify IPv4 addressing and subnetting" does not mean that you should know how to type commands but have no clue as to what you configured. You must first master the conceptual exam topic verbs. The progression runs something like this: Describe, Identify, Explain, Compare/Contrast, Configure, Verify, Troubleshoot For instance, an exam topic that lists "compare and contrast" means that you should be able to describe, identify, and explain the technology. Also, an exam topic with "configure and verify" tells you to also be ready to describe, explain, and compare/contrast. The Context Surrounding the Exam Topics Take a moment to navigate to www.cisco.com/go/certifications and find the list of exam topics for the CCNA 200-301 exam. Did your eyes go straight to the list of exam topics? Or did you take the time to read the paragraphs above the exam topics first? That list of exam topics for the CCNA 200-301 exam includes a little over 50 primary exam topics and about 50 more secondary exam topics. The primary topics have those verbs as just discussed, which tell you something about the depth of skill required. The secondary topics list only the names of more technologies to know. xxxii CCNA 200-301 Official Cert Guide, Volume 2 However, the top of the web page that lists the exam topics also lists some important information that tells us some important facts about the exam topics. In particular, that leading text, found at the beginning of Cisco exam topic pages of most every exam, tells us these important points: ■ The guidelines may change over time. ■ The exam topics are general guidelines about what may be on the exam. ■ The actual exam may include "other related topics." Interpreting these three facts in order, I would not expect to see a change to the published list of exam topics for the exam. I've been writing the Cisco Press CCNA Cert Guides since Cisco announced CCNA back in 1998, and I've never seen Cisco change the official exam topics in the middle of an exam—not even to fix typos. But the introductory words say that they might change the exam topics, so it's worth checking. As for the second item in the preceding list, even before you know what the acronyms mean, you can see that the exam topics give you a general but not detailed idea about each topic. The exam topics do not attempt to clarify every nook and cranny or to list every command and parameter; however, this book serves as a great tool in that it acts as a much more detailed interpretation of the exam topics. We examine every exam topic, and if we think a concept or command is possibly within an exam topic, we put it into the book. So, the exam topics give us general guidance, and these books give us much more detailed guidance. The third item in the list uses literal wording that runs something like this: "However, other related topics may also appear on any specific delivery of the exam." That one statement can be a bit jarring to test takers, but what does it really mean? Unpacking the statement, it says that such questions may appear on any one exam but may not; in other words, they don't set about to ask every test taker some questions that include concepts not mentioned in the exam topics. Second, the phrase "...other related topics..." emphasizes that any such questions would be related to some exam topic, rather than being far afield—a fact that helps us in how we respond to this particular program policy. For instance, the CCNA 200-301 exam includes configuring and verifying the OSPF routing protocol, but it does not mention the EIGRP routing protocol. I personally would be unsurprised to see an OSPF question that required a term or fact not specifically mentioned in the exam topics, but not one that's some feature that (in my opinion) ventures far away from the OSPF features in the exam topics. Also, I would not expect to see a question about how to configure and verify EIGRP. And just as one final side point, note that Cisco does on occasion ask a test taker some unscored questions, and those may appear to be in the vein of questions from outside topics. When you sit down to take the exam, the small print mentions that you may see unscored questions and you won't know which ones are unscored. (These questions give Cisco a way to test possible new questions.) Yet some of these might be ones that fall into the "other related topics" category but then not affect your score. xxxiii You should prepare a little differently for any Cisco exam, in comparison to, say, an exam back in school, in light of Cisco's "other related questions" policy: ■ Do not approach an exam topic with an "I'll learn the core concepts and ignore the edges" approach. ■ Instead, approach each exam topic with a "pick up all the points I can" approach by mastering each exam topic, both in breadth and in depth. ■ Go beyond each exam topic when practicing configuration and verification by taking a little extra time to look for additional show commands and configuration options, and make sure you understand as much of the show command output that you can. By mastering the known topics, and looking for places to go a little deeper, you will hopefully pick up the most points you can from questions about the exam topics. Then the extra practice you do with commands may happen to help you learn beyond the exam topics in a way that can help you pick up other points as well. CCNA 200-301 Exam Content, Per This Book When we created the Official Cert Guide content for the CCNA 200-301 exam, we considered a few options for how to package the content, and we landed on releasing a two-book set. Figure I-3 shows the setup of the content, with roughly 60 percent of the content in Volume 1 and the rest in Volume 2. Fundamentals Ethernet LANs IPv4 Routing IPv6 Routing Wireless LANs Vol. 1 - 60% Figure I-3 Security IP Services Automation Architecture Vol. 2 - 40% Two Books for CCNA 200-301 The two books together cover all the exam topics in the CCNA 200-301 exam. Each chapter in each book develops the concepts and commands related to an exam topic, with clear and detailed explanations, frequent figures, and many examples that build your understanding of how Cisco networks work. As for choosing what content to put into the books, note that we begin and finish with Cisco's exam topics, but with an eye toward predicting as many of the "other related topics" as we can. We start with the list of exam topics and apply a fair amount of experience, discussion, and other secret sauce to come up with an interpretation of what specific concepts and commands are worthy of being in the books or not. At the end of the writing process, the books should cover all the published exam topics, with additional depth and breadth
that I chose based on the analysis of the exam. As we have done from the very first edition of the CCNA Official Cert Guide, we intend to cover each and every topic in depth. But as you would expect, we cannot predict every single fact on the exam given the nature of the exam policies, but we do our best to cover all known topics. xxvii CCNA 200-301 Official Cert Guide, Volume 2 Book Features This book includes many study features beyond the core explanations and examples in each chapter. This section acts as a reference to the various features in the book. Chapter Features and How to Use Each Chapter Each chapter of this book is a self-contained short course about one small topic area, organized for reading and study, as follows: "Do I Know This Already?" quizzes: Each chapter begins with a pre-chapter quiz. Foundation Topics: This is the heading for the core content section of the chapter. Chapter Review: This section includes a list of study tasks useful to help you remember concepts, connect ideas, and practice skills-based content in the chapter. Figure I-4 shows how each chapter uses these three key elements. You start with the DIKTA quiz. You can use the score to determine whether you already know a lot, or not so much, and determine how to approach reading the Foundation Topics (that is, the technology content in the chapter). When finished, use the Chapter Review tasks to start working on mastering your memory of the facts and skills with configuration, verification, and troubleshooting. DIKTA Quiz High Score Take Quiz Low Score Figure I-4 Foundation Topics Chapter Review (Skim) Foundation Topics (Read) Foundation Topics 1) In-Chapter, or... 2) Companion Website Three Primary Tasks for a First Pass Through Each Chapter In addition to these three main chapter features, each "Chapter Review" section uses a variety of other book features, including the following: ■ Review Key Topics: Inside the "Foundation Topics" section, the Key Topic icon appears next to the most important items, for the purpose of later review and mastery. While all content matters, some is, of course, more important to learn, or needs more review to master, so these items are noted as key topics. The Chapter Review lists the key topics in a table. Scan the chapter for these items to review them. Or review the key topics interactively using the companion website. ■ Complete Tables from Memory: Instead of just rereading an important table of information, you will find some tables have been turned into memory tables, an interactive exercise found on the companion website. Memory tables repeat the table but with parts of the table removed. You can then fill in the table to exercise your memory and click to check your work. ■ Key Terms You Should Know: You do not need to be able to write a formal definition of all terms from scratch; however, you do need to understand each term well enough to understand exam questions and answers. The Chapter Review lists the key terminology from the chapter. Make sure you have a good understanding of each term and use the Glossary to cross-check your own mental definitions. You can also review key terms with the "Key Terms Flashcards" app on the companion website. xxv ■ Labs: Many exam topics use verbs such as configure and verify; all these refer to skills you should practice at the user interface (CLI) of a router or switch. The Chapter and Part Reviews refer you to these other tools. The upcoming section titled "About Building Hands-On Skills" discusses your options. ■ Command References: Some book chapters cover a large number of router and switch commands. The Chapter Review includes reference tables for the commands used in that chapter, along with an explanation. Use these tables for reference, but also use them for study. Just cover one column of the table and see how much you can remember and complete mentally. ■ Review DIKTA Questions: Although you have already seen the DIKTA questions from the chapters, re-answering those questions can prove a useful way to review facts. The Part Review suggests that you repeat the DIKTA questions but using the Pearson Test Prep (PTP) exam. Part Features and How to Use the Part Review The book organizes the chapters into parts for the purpose of helping you study for the exam. Each part groups a small number of related chapters together. Then the study process (described just before Chapter 1) suggests that you pause after each part to do a review of all chapters in the part. Figure I-5 lists the titles of the eight parts and the chapters in those parts (by chapter number) for this book. 5 Network Automation (16-19) 4 Network Architecture (13-15) 3 IP Services (9-12) 1 IP Access Control Lists (1-3) Figure I-5 2 Security Services (4-8) The Book Parts (by Title), and Chapter Numbers in Each Part The Part Review that ends each part acts as a tool to help you with spaced review sessions. Spaced reviews—that is, reviewing content several times over the course of your study—help improve retention. The Part Review activities include many of the same kinds of activities seen in the Chapter Review. Avoid skipping the Part Review, and take the time to do the review; it will help you in the long run. The Companion Website for Online Content Review We created an electronic version of every Chapter and Part Review task that could be improved through an interactive version of the tool. For instance, you can take a "Do I Know This Already?" quiz by reading the pages of the book, but you can also use our testing software. As another example, when you want to review the key topics from a chapter, you can find all those in electronic form as well. xxvii CCNA 200-301 Official Cert Guide, Volume 2 All the electronic review elements, as well as other electronic components of the book, exist on this book's companion website. The companion website gives you a big advantage: you can do most of your Chapter and Part Review work from anywhere using the interactive tools on the site. The advantages include ■ Easier to use: Instead of having to print out copies of the appendices and do the work on paper, you can use these new apps, which provide you with an easy-to-use, interactive experience that you can easily run over and over. ■ Convenient: When you have a spare 5-10 minutes, go to the book's website and review content from one of your recently finished chapters. ■ Untethered from the book: You can access your review activities from anywhere—no need to have the book with you. ■ Good for tactile learners: Sometimes looking at a static page after reading a chapter lets your mind wander. Tactile learners might do better by at least typing answers into an app, or clicking inside an app to navigate, to help keep you focused on the activity. The interactive Chapter Review elements should improve your chances of passing as well. Our in-depth reader surveys over the years show that those who do the Chapter and Part Reviews learn more. Those who use the interactive versions of the review elements also tend to do more of the Chapter and Part Review work. So take advantage of the tools and maybe you will be more successful as well. Table I-1 summarizes these interactive applications and the traditional book features that cover the same content. Table I-1 Book Features with Both Traditional and App Options Feature Traditional App Key Topic Table with list; flip pages to find Key Topics Table app Config Checklist Just one of many types of key topics Config Checklist app Key Terms List in each "Chapter Review" section, with the Glossary in the back of the book Glossary Flash Cards app The companion website also includes links to download, navigate, or stream for these types of content: ■ Pearson Sim Lite Desktop App ■ Pearson Test Prep (PTP) Desktop App ■ Pearson Test Prep (PTP) Videos as mentioned in book chapters xxxvii How to Access the Companion Website To access the companion website, which gives you access to the electronic content with this book, start by establishing a login at www.ciscopress.com and register your book. To do so, simply go to www.ciscopress.com/register and enter the ISBN of the print book: 9781587147135. After you have registered your book, go to your account page and click the Registered Products tab. From there, click the Access Bonus Content link to get access to the book's companion website. Note that if you buy the Premium Edition eBook and Practice Test version of this book from Cisco Press, your book will automatically be registered on your account page. Simply go to your account page, click the Registered Products tab, and select Access Bonus Content to access the book's companion website. How to Access the Pearson Test Prep (PTP) App You have two options for installing and using the Pearson Test Prep application: a web app and a desktop app. To use the Pearson Test Prep application, start by finding the registration code that comes with the book. You can find the code in these ways: ■ Print book: Look in the cardboard sleeve in the back of the book for a piece of paper with your book's unique PTP code. ■ Premium Edition: If you purchase the Premium Edition eBook and Practice Test directly from the Cisco Press website, the code will be populated on your account page after purchase. Just log in at www.ciscopress.com, click account to see details of your account, and click the digital purchases tab. ■ Amazon Kindle: For those who purchase a Kindle edition from Amazon, the access code will be supplied directly from Amazon. ■ Other bookseller e-books: Note that if you purchase an e-book version from any other source, the practice test is not included because other vendors to date have not chosen to vend the required unique access code. NOTE Do not lose the activation code because it is the only means with which you can access the QA content with the book. Once you have
the access code, to find instructions about both the PTP web app and the desktop app, follow these steps: Step 1. Open this book's companion website, as was shown earlier in this Introduction under the heading "How to Access the Companion Website." Step 2. Click the Practice Exams button. Step 3. Follow the instructions listed there both for installing the desktop app and for using the web app. xxxviii CCNA 200-301 Official Cert Guide, Volume 2 Note that if you want to use the web app only at this point, just navigate to www.pearsonstestprep.com, establish a free login if you do not already have one, and register this book's practice tests using the registration code you just found. The process should take only a couple of minutes. NOTE Amazon e-book (Kindle) customers: It is easy to miss Amazon's email that lists your PTP access code. Soon after you purchase the Kindle e-book, Amazon should send an email. However, the email uses very generic text and makes no specific mention of PTP or practice exams. To find your code, read every email from Amazon after you purchase the book. Also, do the usual checks (such as checking your spam folder) for ensuring your email arrives. NOTE Other e-book customers: As of the time of publication, only the publisher and Amazon supply PTP access codes when you purchase their e-book editions of this book. Feature Reference The following list provides an easy reference to get the basic idea behind each book feature: ■ Practice exam: The book gives you the rights to the Pearson Test Prep (PTP) testing software, available as a web app and desktop app. Use the access code on a piece of cardboard in the sleeve in the back of the book, and use the companion website to download the desktop app or navigate to the web app (or just go to www.pearsonstestprep.com). ■ E-book: Pearson offers an e-book version of this book that includes extra practice tests. If interested, look for the special offer on a coupon card inserted in the sleeve in the back of the book. This offer enables you to purchase the CCNA 200-301 Official Cert Guide, Volume 2, Premium Edition eBook and Practice Test at a 70 percent discount off the list price. The product includes three versions of the e-book: PDF (for reading on your computer), EPUB (for reading on your tablet, mobile device, or Nook or other e-reader), and Mobi (the native Kindle version). It also includes additional practice test questions and enhanced practice test features. ■ Mentoring videos: The companion website also includes a number of videos about other topics as mentioned in individual chapters. ■ CCNA 200-301 Network Simulator Lite: This live version of the best-selling CCNA Network Simulator from Pearson provides you with a means, right now, to experience the Cisco command-line interface (CLI). No need to go buy real gear or buy a full simulator to start learning the CLI. Just install it from the companion website. ■ CCNA Simulator: If you are looking for more hands-on practice, you might want to consider purchasing the CCNA Network Simulator. You can purchase a copy of this software from Pearson at or other xxxix retail outlets. To help you with your studies, Pearson has created a mapping guide that maps each of the labs in the simulator to the specific sections in each volume of the CCNA Cert Guide. You can get this mapping guide free on the Extras tab on the book product page: www.ciscopress.com/title/9781587147135. ■ PearsonITCertification.com: The website www.pearsonitcertification.com is a great resource for all things IT-certification related. Check out the great CCNA articles, videos, blogs, and other certification preparation tools from the industry's best authors and trainers. ■ Author's website and blogs: The author maintains a website that hosts tools and links useful when studying for CCNA. In particular, the site has a large number of free lab exercises about CCNA content, additional sample questions, and other exercises. Additionally, the site indexes all content so you can study based on the book chapters and parts. To find it, navigate to .Book Organization, Chapters, and Appendices The CCNA 200-301 Official Cert Guide, Volume 1, contains 29 chapters, while this book has 19 core chapters. Each chapter covers a subset of the topics on the CCNA exam. The book organizes its chapters into parts of three to five chapters as follows: ■ Part I: IP Access Control Lists ■ Chapter 1, "Introduction to TCP/IP Transport and Applications," completes most of the detailed discussion of the upper two layers of the TCP/IP model (transport and application), focusing on TCP and applications. ■ Chapter 2, "Basic IPv4 Access Control Lists," examines how standard IP ACLs can filter packets based on the source IP address so that a router will not forward the packet. ■ Chapter 3, "Advanced IPv4 Access Control Lists," examines both named and numbered ACLs, and both standard and extended IP ACLs. Part II: Security Services ■ Chapter 4, "Security Architectures," discusses a wide range of fundamental concepts in network security. ■ Chapter 5, "Securing Network Devices," shows how to use the router and switch CLI and introduces the concepts behind firewalls and intrusion prevention systems (IPSs). ■ Chapter 6, "Implementing Switch Port Security," explains the concepts as well as how to configure and verify switch port security, a switch feature that does basic MAC-based monitoring of the devices that send data into a switch. ■ Chapter 7, "Implementing DHCP," discusses how hosts can be configured with their IPv4 settings and how they can learn those settings with DHCP. ■ Chapter 8, "DHCP Snooping and ARP Inspection," shows how to implement two related switch security features, with one focusing on reacting to suspicious DHCP messages and the other reacting to suspicious ARP messages. xi CCNA 200-301 Official Cert Guide, Volume 2 ■ ■ Part III: IP Services ■ Chapter 9, "Device Management Protocols," discusses the concepts and configuration of some common network management tools: syslog, NTP, CDP, and LLDP. ■ Chapter 10, "Network Address Translation," works through the complete concept, configuration, verification, and troubleshooting sequence for the router NAT feature, including how it helps conserve public IPv4 addresses. ■ Chapter 11, "Quality of Service (QoS)," discusses a wide variety of concepts all related to the broad topic of QoS. ■ Chapter 12, "Miscellaneous IP Services," discusses several topics for which the exam requires conceptual knowledge but no configuration knowledge, including FHRPs (including HSRP), SNMP, TFTP, and FTP. Part IV: Network Architecture ■ Chapter 13, "LAN Architecture," examines various ways to design Ethernet LANs, discussing the pros and cons, and explains common design terminology, including Power over Ethernet (PoE). ■ Chapter 14, "WAN Architecture," discusses the concepts behind three WAN alternatives: Metro Ethernet, MPLS VPNs, and Internet VPNs. ■ Chapter 15, "Cloud Architecture," explains the basic concepts and then generally discusses the impact that cloud computing has on a typical enterprise network, including the foundational concepts of server virtualization. Part V: Network Automation ■ Chapter 16, "Introduction to Controller-Based Networking," discusses many concepts and terms related to how Software-Defined Networking (SDN) and network programmability are impacting typical enterprise networks. ■ Chapter 17, "Cisco Software-Defined Access (SDA)," discusses Cisco's Software-Defined Networking (SDN) offering for the enterprise, including the DNA Center controller. ■ Chapter 18, "Understanding REST and JSON," explains the foundational concepts of REST APIs, data structures, and how JSON can be useful for exchanging data using APIs. ■ Chapter 19, "Understanding Ansible, Puppet, and Chef," discusses the need for configuration management software and introduces the basics of each of these three configuration management tools. Part VI: Final Review ■ Chapter 20, "Final Review," suggests a plan for final preparation after you have finished the core parts of the book, in particular explaining the many study options available in the book. Part VII: Appendices ■ Appendix A, "Numeric Reference Tables," lists several tables of numeric information, including a binary-to-decimal conversion table and a list of powers of 2. xii ■ ■ Appendix B, "CCNA 200-301 Volume 2 Exam Updates," is a place for the author to add book content mid-edition. Always check online for the latest PDF version of this appendix; the appendix lists download instructions. ■ Appendix C, "Answers to the 'Do I Know This Already?' Quizzes," includes the explanations to all the "Do I Know This Already?" quizzes. ■ The Glossary contains definitions for many of the terms used in the book, including the terms listed in the "Key Terms You Should Know" sections at the conclusion of the chapters. Online Appendices ■ Appendix D, "Topics from Previous Editions" ■ Appendix E, "Practice for Chapter 2: Basic IPv4 Access Control Lists" ■ Appendix F, "Previous Edition ICND1 Chapter 35: Managing IOS Files" ■ Appendix G, "Exam Topics Cross-Reference," provides some tables to help you find where each exam objective is covered in the book. ■ Appendix H, "Study Planner," is a spreadsheet with major study milestones, where you can track your progress through your study. About Building Hands-On Skills You need skills in using Cisco routers and switches, specifically the Cisco command-line interface (CLI). The Cisco CLI is a text-based command-and-response user interface: you type a command, and the device (a router or switch) displays messages in response. To answer sim and simlet questions on the exams, you need to know a lot of commands, and you need to be able to navigate to the right place in the CLI to use those
commands. This next section walks through the options of what is included in the book, with a brief description of lab options outside the book. Config Lab Exercises Some router and switch features require multiple configuration commands. Part of the skill you need to learn is to remember which configuration commands work together, which ones are required, and which ones are optional. So, the challenge level goes beyond just picking the right parameters on one command. You have to choose which commands to use, in which combination, typically on multiple devices. And getting good at that kind of task requires practice. Each Config Lab lists details about a straightforward lab exercise for which you should create a small set of configuration commands for a few devices. Each lab presents a sample lab topology, with some requirements, and you have to decide what to configure on each device. The answer then shows a sample configuration. Your job is to create the configuration and then check your answer versus the supplied answer. Config Lab content resides outside the book at the author's blog site (). You can navigate to the Config Lab in a couple of ways from the site, or just go directly to config-lab/ to reach a list of all Config Labs. Figure I-6 shows the logo that you will see with each Config Lab. xiii CCNA 200-301 Official Cert Guide, Volume 2 Figure I-6 Config Lab Logo in the Author's Blogs These Config Labs have several benefits, including the following: Untethered and responsive: Do them from anywhere, from any web browser, from your phone or tablet, untethered from the book. Designed for idle moments: Each lab is designed as a 5- to 10-minute exercise if all you are doing is typing in a text editor or writing your answer on paper. Two outcomes, both good: Practice getting better and faster with basic configuration, or if you get lost, you have discovered a topic that you can now go back and reread to complete your knowledge. Either way, you are a step closer to being ready for the exam! Blog format: The format allows easy adds and changes by me and easy comments by you. Self-assessment: As part of final review, you should be able to do all the Config Labs, without help, and with confidence. Note that the blog organizes these Config Lab posts by book chapter, so you can easily use these at both Chapter Review and Part Review. A Quick Start with Pearson Network Simulator Lite The decision of how to get hands-on skills can be a little scary at first. The good news: You have a free and simple first step to experience the CLI: install and use the Pearson Network Simulator Lite (or NetSim Lite) that comes with this book. This book comes with a live version of the best-selling CCNA Network Simulator from Pearson, which provides you with a means, right now, to experience the Cisco CLI. No need to go buy real gear or buy a full simulator to start learning the CLI. Just install it from the companion website. The CCNA 200-301 Network Simulator Lite Volume 2 software contains 13 labs covering ACL topics from Part I in the book. So, make sure to use the NetSim Lite to learn the basics of the CLI to get a good start. Of course, one reason that you get access to the NetSim Lite is that the publisher hopes you will buy the full product. However, even if you do not use the full product, you can still learn from the labs that come with NetSim Lite while deciding about what options to pursue. xliii The Pearson Network Simulator The Config Labs and the Pearson Network Simulator Lite both fill specific needs, and they both come with the book. However, you need more than those two tools. The single best option for lab work to do along with this book is the paid version of the Pearson Network Simulator. This simulator product simulates Cisco routers and switches so that you can learn for CCNA certification. But more importantly, it focuses on learning for the exam by providing a large number of useful lab exercises. Reader surveys tell us that those people who use the Simulator along with the book love the learning process and rave about how the book and Simulator work well together. Of course, you need to make a decision for yourself and consider all the options. Thankfully, you can get a great idea of how the full Simulator product works by using the Pearson Network Simulator Lite product included with the book. Both have the same base code, same user interface, and same types of labs. Try the Lite version to decide if you want to buy the full product. Note that the Simulator and the books work on a different release schedule. For a time in 2020, the Simulator will be the one created for the previous versions of the exams (ICND1 100-101, ICND2 200-101, and CCNA 200-102). Interestingly, Cisco did not add a large number of new topics that require CLI skills to the CCNA 200-301 exam as compared with its predecessor, so the old Simulator covers most of the CCNA 200-301 CLI topics. So, during the interim before the products based on the 200-301 exam come out, the old Simulator products should be quite useful. On a practical note, when you want to do labs when reading a chapter or doing Part Review, the Simulator organizes the labs to match the book. Just look for the Sort by Chapter tab in the Simulator's user interface. However, during the months in 2020 for which the Simulator is the older edition listing the older exams in the title, you will need to refer to a PDF that lists those labs versus this book's organization. You can find that PDF on the book product page under the Downloads tab here: www.ciscopress.com/title/9781587147135. More Lab Options If you decide against using the full Pearson Network Simulator, you still need hands-on experience. You should plan to use some lab environment to practice as much CLI as possible. First, you can use real Cisco routers and switches. You can buy them, new or used, or borrow them at work. You can rent them for a fee. If you have the right mix of gear, you could even do the Config Lab exercises from my blog on that gear or try to recreate examples from the book. Cisco also makes a simulator that works very well as a learning tool: Cisco Packet Tracer. Cisco now makes Packet Tracer available for free. However, unlike the Pearson Network Simulator, it does not include lab exercises that direct you as to how to go about learning each topic. If interested in more information about Packet Tracer, check out my series about using Packet Tracer at my blog (); just search for "Packet Tracer." xliii CCNA 200-301 Official Cert Guide, Volume 2 Cisco offers a virtualization product that lets you run router and switch operating system (OS) images in a virtual environment. This tool, the Virtual Internet Routing Lab (VIRL), lets you create a lab topology, start the topology, and connect to real virtual Cisco router and switch lab pods from Cisco, in an offering called Cisco Learning Labs (). This book does not tell you what option to use, but you should plan on getting some hands-on practice somehow. The important thing to know is that most people need to practice using the Cisco CLI to be ready to pass these exams. For More Information If you have any comments about the book, submit them via www.ciscopress.com. Just go to the website, select Contact Us, and type your message. Cisco might make changes that affect the CCNA certification from time to time. You should always check www.cisco.com/go/ccna for the latest details. The CCNA 200-301 Official Cert Guide, Volume 2, helps you attain CCNA certification. This is the CCNA certification book from the only Cisco-authorized publisher. We at Cisco Press believe that this book certainly can help you achieve CCNA certification, but the real work is up to you! I trust that your time will be well spent. Figure Credits Figure 7-9, screenshot of network connection details © Microsoft, 2019 Figure 7-10, screenshot(s) reprinted with permission from Apple, Inc. Figure 7-11, screenshot of Linux © The Linux Foundation Figure 12-16, screenshot of CS Bloggings 2018 © FileZilla Figure 13-9, electric outlet © Mike McDonald/Shutterstock Figure 15-10, screenshot of Set Up VM with Different CPU/RAM/OS © 2019, Amazon Web Services, Inc Figure 16-13, illustration of man icon © AlexHliv/Shutterstock Figure 17-1, illustration of man icon © AlexHliv/Shutterstock Figure 17-11, illustration of man icon © AlexHliv/Shutterstock Figure 18-9, screenshot of REST GET Request-1 © 2019 Postman, Inc. Figure 20-1, screenshot of PTP Grading © 2019 Pearson Education Figure 20-2, screenshot of PTP Grading © 2019 Pearson Education Figure D-1, ribbon set © petrnuttil/123RF The CCNA Official Cert Guide, Volume 2 includes the topics that help you build an enterprise network so all devices can communicate with all other devices. Parts I and II of this book focus on how to secure that enterprise network so that only the appropriate devices and users can communicate. Part I focuses on IP Version 4 (IPv4) access control lists (ACLs). ACLs are IPv4 packet filters that can be programmed to look at IPv4 packet headers, make choices, and either allow a packet through or discard the packet. Because you can implement IPv4 ACLs on any router, a network engineer has a large number of options of where to use ACLs, without adding additional hardware or software, making ACLs a very flexible and useful tool. Chapter 1 begins this part with an introduction to the TCP/IP transport layer protocols TCP and UDP, along with an introduction to several TCP/IP applications. This chapter provides the necessary background to understand the ACL chapters and to better prepare you for upcoming discussions of additional security topics in Part II and IP services topics in Part III. Chapters 2 and 3 get into details about ACLs. Chapter 2 discusses ACL basics, avoiding some of the detail to ensure that you
master several key concepts. Chapter 3 then looks at the much wider array of ACL features to make you ready to take advantage of the power of ACLs and to be ready to better manage those ACLs. Part I IP Access Control Lists Chapter 1: Introduction to TCP/IP Transport and Applications Chapter 2: Basic IPv4 Access Control Lists Chapter 3: Advanced IPv4 Access Control Lists Part I Review CHAPTER 1 Introduction to TCP/IP Transport and Applications This chapter covers the following exam topics: 1.0 Network Fundamentals 1.5 Compare TCP to UDP 4.0 IP Services 4.3 Explain the role of DHCP and DNS in the network The CCNA exam focuses mostly on functions at the lower layers of TCP/IP, which define how IP networks can send IP packets from host to host using LANs and WANs. This chapter explains the basics of a few topics that receive less attention on the exams: the TCP/IP transport layer and the TCP/IP application layer. The functions of these higher layers play a big role in real TCP/IP networks. Additionally, many of the security topics in Parts I and II of this book, and some of the IP services topics in Part III, require you to know the basics of how the transport and application layers of TCP/IP work. This chapter serves as that introduction. This chapter begins by examining the functions of two transport layer protocols: Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). The second major section of the chapter examines the TCP/IP application layer, including some discussion of how Domain Name System (DNS) name resolution works. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 1-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions TCP/IP Layer 4 Protocols: TCP and UDP 1-4 TCP/IP Applications 5-6 1. Which of the following header fields identify which TCP/IP application gets data received by the computer? (Choose two answers.) a. Ethernet Type b. SNAP Protocol Type c. IP Protocol d. TCP Port Number e. UDP Port Number 2. Which of the following are typical functions of TCP? (Choose four answers.) a. Flow control (windowing) b. Error recovery c. Multiplexing using port numbers d. Routing e. Encryption f. Ordered data transfer 3. Which of the following functions is performed by both TCP and UDP? a. Windowing b. Error recovery c. Multiplexing using port numbers d. Routing e. Encryption f. Ordered data transfer 4. What do you call data that includes the Layer 4 protocol header, and data given to Layer 4 by the upper layers, not including any headers and trailers from Layers 1 to 3? (Choose two answers.) a. L3PDU b. Chunk c. Segment d. Packet e. Frame f. L4PDU 5. In the URI which part identifies the web server? a. http b. blog.certskills.com c. certskills.com d. e. The file name.html includes the hostname. 6 CCNA 200-301 Official Cert Guide, Volume 2 6. Fred opens a web browser and connects to the www.certskills.com website. Which of the following are typically true about what happens between Fred's web browser and the web server? (Choose two answers.) a. Messages flowing toward the server use UDP destination port 80. b. Messages flowing from the server typically use RTP. c. Messages flowing to the client typically use a source TCP port number of 80. d. Messages flowing to the server typically use TCP. Foundation Topics TCP/IP Layer 4 Protocols: TCP and UDP The OSI transport layer (Layer 4) defines several functions, the most important of which are error recovery and flow control. Likewise, the TCP/IP transport layer protocols also implement these same types of features. Note that both the OSI model and the TCP/IP model call this layer the transport layer. But as usual, when referring to the TCP/IP model, the layer name and number are based on OSI, so any TCP/IP transport layer protocols are considered Layer 4 protocols. The key difference between TCP and UDP is that TCP provides a wide variety of services to applications, whereas UDP does not. For example, routers discard packets for many reasons, including bit errors, congestion, and instances in which no correct routes are known. As you have read already, most data-link protocols notice errors (a process called error detection) but then discard frames that have errors. TCP provides retransmission (error recovery) and helps to avoid congestion (flow control), whereas UDP does not. As a result, many application protocols choose to use TCP. However, do not let UDP's lack of services make you think that UDP is worse than TCP. By providing fewer services, UDP needs fewer bytes in its header compared to TCP, resulting in fewer bytes of overhead in the network. UDP software does not slow down data transfer in cases where TCP can purposefully slow down. Also, some applications, notably today Voice over IP (VoIP) and video over IP, do not need error recovery, so they use UDP. So, UDP also has an important place in TCP/IP networks today. Table 1-2 lists the main features supported by TCP/UDP. Note that only the first item listed in the table is supported by UDP, whereas all items in the table are supported by TCP. Table 1-2 TCP/IP Transport Layer Features Function Description Multiplexing using ports Function that allows receiving hosts to choose the correct application for which the data is destined, based on the port number Error recovery (reliability) Process of numbering and acknowledging data with Sequence and Acknowledgment header fields Flow control using windowing Process that uses window sizes to protect buffer space and routing devices from being overloaded with traffic Chapter 1: Introduction to TCP/IP Transport and Applications 7 Function Description Connection establishment Process used to initialize port numbers and Sequence and termination Acknowledgment fields Ordered data transfer and Continuous stream of bytes from an upper-layer process that data segmentation is "segmented" for transmission and delivered to upper-layer processes at the receiving device, with the bytes in the same order Next, this section describes the features of TCP, followed by a brief comparison to UDP. Transmission Control Protocol Each TCP/IP application typically chooses to use either TCP or UDP based on the application's requirements. For example, TCP provides error recovery, but to do so, it consumes more bandwidth and uses more processing cycles. UDP does not perform error recovery, but it takes less bandwidth and uses fewer processing cycles. Regardless of which of these two TCP/IP transport layer protocols the application chooses to use, you should understand the basics of how each of these transport layer protocols works. TCP, as defined in Request For Comments (RFC) 793, accomplishes the functions listed in Table 1-2 through mechanisms at the endpoint computers. TCP relies on IP for end-to-end delivery of the data, including routing issues. In other words, TCP performs only part of the functions necessary to deliver the data between applications. Also, the role that it plays is directed toward providing services for the applications that sit at the endpoint computers. Regardless of whether two computers are on the same Ethernet, or are separated by the entire Internet, TCP performs its functions the same way. Figure 1-1 shows the fields in the TCP header. Although you do not need to memorize the names of the fields or their locations, the rest of this section refers to several of the fields, so the entire header is included here for reference. 4 Bytes Source Port Destination Port Sequence Number Acknowledgment Number Offset Reserved Flag Bits Checksum Figure 1-1 Window Urgent TCP Header Fields The message created by TCP that begins with the TCP header, followed by any application data, is called a TCP segment. Alternatively, the more generic term Layer 4 PDU, or L4PDU, can also be used. Multiplexing Using TCP Port Numbers TCP and UDP both use a concept called multiplexing. Therefore, this section begins with an explanation of multiplexing with TCP and UDP. Afterward, the unique features of TCP are explored. 1 8 CCNA 200-301 Official Cert Guide, Volume 2 Multiplexing by TCP and UDP involves the process of how a computer thinks when receiving data. The computer might be running many applications, such as a web browser, an email package, or an Internet VoIP application (for example, Skype). TCP and UDP multiplexing tells the receiving computer to which application to give the received data. Some examples will help make the need for multiplexing obvious. The sample network consists of two PCs, labeled Hannah and George. Hannah uses an application that she wrote to send advertisements that appear on George's screen. The application sends a new ad to George every 10 seconds. Hannah uses a second application, a wire-transfer application, to send George some money. Finally, Hannah uses a web browser to access the web server that runs on George's PC. The ad application and wire-transfer application are imaginary, just for this example. The web application works just like it would in real life. Figure 1-2 shows the sample network, with George running three applications: ■ A UDP-based advertisement application ■ A TCP-based wire-transfer application ■ A TCP web server application Hannah George Web Server Ad Application Wire Application Eth IP UDP Ad Data Eth Eth IP TCP Wire Transfer Data Eth Eth IP TCP Web Page Data Eth Figure 1-2 I received three packets from the same source MAC and IP. Which of my applications gets the data in each? Hannah Sending Packets to George, with Three Applications George needs to know which
application to give the data to, but all three packets are from the same Ethernet and IP address. You might think that George could look at whether the packet contains a UDP or TCP header, but as you see in the figure, two applications (wire transfer and web) are using TCP. TCP and UDP solve this problem by using a port number field in the TCP or UDP header, respectively. Each of Hannah's TCP and UDP segments uses a different destination port number so that George knows which application to give the data to. Figure 1-3 shows an example. Multiplexing relies on a concept called a socket. A socket consists of three things: ■ An IP address ■ A transport protocol ■ A port number Answers to the "Do I Know This Already?" Quiz 1. D, E 2. A, B, C, F 3. C 4. C, F 5. B 6. C, D Chapter 1: Introduction to TCP/IP Transport and Applications 9 Hannah George Port 80 Web Server Port 800 Ad Server Port 9876 Wire Application Eth IP UDP Ad Data .OOORRLNQLKWH UDP or TCP Destination port to identify the application! Eth Destination Port 800 Eth IP TCP Wire Transfer Data Eth Destination Port 9876 Eth IP TCP Web Page Data Eth Figure 1-3 Hannah Sending Packets to George, with Three Applications Using Port Numbers to Multiplex So, for a web server application on George, the socket would be (10.1.1.2, TCP, port 80) because, by default, web servers use the well-known port 80. When Hannah's web browser connects to the web server, Hannah uses a socket as well—possibly one like this: (10.1.1.1, TCP, 49160). Why 49160? Well, Hannah just needs a port number that is unique on Hannah, so Hannah sees that port 49160. The Internet Assigned Numbers Authority (IANA), the same organization that manages IP address allocation worldwide, subdivides the port number ranges into three main ranges. The first two ranges reserve numbers that IANA can then allocate to specific application protocols through an application and review process, with the third category reserving ports to be dynamically allocated as used for clients, as with the port 49160 example in the previous paragraph. The names and ranges of port numbers (as detailed in RFC 6335) are ■ Well Known (System) Ports: Numbers from 0 to 1023, assigned by IANA, with a stricter review process to assign new ports than user ports. ■ User (Registered) Ports: Numbers from 1024 to 49151, assigned by IANA with a less strict process to assign new ports compared to well-known ports. ■ Ephemeral (Dynamic, Private) Ports: Numbers from 49152 to 65535, not assigned and intended to be dynamically allocated and used temporarily for a client application while the app is running. Figure 1-4 shows an example that uses three ephemeral ports on the user device on the left, with the server on the right using two well-known ports and one user port. The computers use three applications at the same time; hence, three socket connections are open. Because a socket on a single computer should be unique, a connection between two sockets should identify a unique connection between two computers. This uniqueness means that you can use multiple applications at the same time, talking to applications running on the same or different computers. Multiplexing, based on sockets, ensures that the data is delivered to the correct applications. 1 10 CCNA 200-301 Official Cert Guide, Volume 2 User Server Ad Wire Web Application Application Browser Port 49159 Port 49153 Port 49152 Ad Wire Web Application Application Server Port 800 Port 9876 Port 80 UDP UDP TCP IP Address 10.1.1.2 IP Address 10.1.1.1 (10.1.1.1, TCP, 49152) (10.1.1.1, TCP, 49153) (10.1.1.1, UDP, 49159) Figure 1-4 TCP (10.1.1.2, TCP, 80) (10.1.1.2, UDP, 800) Connections Between Sockets Port numbers are a vital part of the socket concept. Servers use well-known ports (or user ports), whereas clients use dynamic ports. Applications that provide a service, such as FTP, Telnet, and web servers, open a socket using a well-known port and listen for connection requests. Because these connection requests from clients are required to include both the source and destination port numbers, the port numbers used by the servers must be known beforehand. Therefore, each service uses a specific well-known port number or user port number. Both well-known and user ports are listed at www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.txt. On client machines, where the requests

level, you could authenticate users with simple passwords that are configured on the switch console and VTY lines. Authorization could be equally simple: when users successfully log in, they are authorized for EXEC level privileges. By entering the correct enable secret password, users could be authorized for a higher privilege level. Under the simple scenario, if a user knows the correct password, he can connect to the switch. But who is that user? You might never know who actually logged in and changed the configuration or rebooted the switch! Instead, you could configure individual usernames and passwords on the switch. That would solve the user anonymity problem, but your network might consist of many administrative users and many switches, requiring quite a bit of username configuration and maintenance. A more scalable solution is to leverage AAA functions that are centralized, standardized, resilient, and flexible. For example, a centralized authentication server can contain a database of all possible users and their passwords, as well as policies to authorize user activities. As users come and go, their accounts can be easily updated in one place. All switches and routers would query the AAA server to get up-to-date information about a user. For greater security, AAA servers can also support multifactor user credentials and more. Cisco implements AAA services in its Identity Services Engine (ISE) platform. AAA servers usually support the following two protocols to communicate with enterprise resources: ■ TACACS+: A Cisco proprietary protocol that separates each of the AAA functions. Communication is secure and encrypted over TCP port 49. ■ RADIUS: A standards-based protocol that combines authentication and authorization into a single resource. Communication uses UDP ports 1812 and 1813 (accounting) but is not completely encrypted. Chapter 4. Security Architectures 83 Both TACACS+ and RADIUS are arranged as a client/server model, where an authenticating device acts as a client talking to a AAA server. Figure 4-9 shows a simplified view of the process, where a user is attempting to connect to a switch for management purposes. In the AAA client role, the switch is often called Network Access Device (NAD) or Network Access Server (NAS). When a user tries to connect to the switch, the switch challenges the user for credentials, then passes the credentials along to the AAA server. In simple terms, if the user passes authentication, the AAA server returns an "accept" message to the switch. If the AAA server requires additional credentials, as in multifactor authentication, it returns a "challenge" message to the switch. Otherwise, a "reject" message is returned, denying access to the user. User Switch AAA Server 1. Who are you? Figure 4-9 2. I am John Smith. 3. Is he John Smith? 5. OK, connect. 4. Yes, accept him. A Simplified View of AAA Developing a Security Program to Educate Users One effective approach an enterprise can take to improve information security is to educate its user community through a corporate security program. Most users may not have an IT background, so they might not recognize vulnerabilities or realize the consequences of their own actions. For example, if corporate users receive an email message that contains a message concerning a legal warrant for their arrest or a threat to expose some supposed illegal behavior, they might be tempted to follow a link to a malicious site. Such an action might infect a user's computer and then open a back door or introduce malware or a worm that could then impact the business operations. An effective security program should have 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. For example, users should not include sensitive information in emails or attachments, should not keep or transmit that information from a smartphone, or store it on cloud services or removable storage drives. ■ User training: All users should be required to participate in periodic formal training so that they become familiar with all corporate security policies. (This also implies that the enterprise should develop and publish formal security policies for its employees, users, and business partners to follow.) 4 84 CCNA 200-301 Official Cert Guide, Volume 2 ■ Physical access control: Infrastructure locations, such as network closets and data centers, should remain securely locked. Badge access to sensitive locations is a scalable solution, offering an audit trail of identities and timestamps when access is granted. Administrators can control access on a granular basis and quickly remove access when an employee is dismissed. Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 4-6 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 4-6 Chapter Review Tracking Review Element Review Date(s) Resource Use Review key topics Book, website Review key terms Book, website Answer DITKA questions Book, PTP Review memory tables Website Review All the Key Topics Table 4-7 Key Topics for Chapter 4 Key Topic Element Description Page Number Figure 4-3 Security terminology 71 Section Common Security Threats 72 Table 4-3 Types of malware 79 Table 4-4 Human security vulnerabilities 80 Paragraph Password vulnerabilities 80 List AAA functions 82 List User education 83 Key Terms You Should Know AAA, amplification attack, brute-force attack, buffer overflow attack, denial-of-service (DoS) attack, dictionary attack, distributed denial-of-service (DDoS) attack, exploit, malware, man-in-the-middle attack, mitigation technique, multifactor authentication, password guessing, phishing, reconnaissance attack, reflection attack, social engineering, spear phishing, spoofing attack, threat, trojan horse, virus, vulnerability, watering hole attack, whaling, worm This page intentionally left blank CHAPTER 5 Securing Network Devices This chapter covers the following exam topics: 1.0 Network Fundamentals 1.1 Explain the Role of Network Components 1.1.c Next-generation Firewalls and IPS 4.0 IP Services 4.8 Configure network devices for remote access using SSH 5.0 Security Fundamentals 5.3 Configure device access control using local passwords All devices in the network—endpoints, servers, and infrastructure devices like routers and switches—include some methods for the devices to legitimately communicate using the network. To protect those devices, the security plan will include a wide variety of tools and mitigation techniques, with the chapters in Part II of this book discussing a large variety of those tools and techniques. This chapter focuses on two particular security needs in an enterprise network. First, access to the CLI of the network devices needs to be protected. The network engineering team needs to be able to access the devices remotely, so the devices need to allow remote SSH (and possibly Telnet) access. The first half of this chapter discusses how to configure passwords to keep them safe and how to filter login attempts at the devices themselves. The second half of the chapter turns to two different security functions most often implemented with purpose-built appliances: firewalls and IPSs. These devices together monitor traffic in transit to determine if the traffic is legitimate or if it might be part of some exploit. If considered to be part of an exploit, or if contrary to the rules defined by the devices, they can discard the messages, stopping any attack before it gets started. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 5-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions Securing IOS Passwords 1–4 Firewalls and Intrusion Prevention Systems 5, 6 1. Imagine that you have configured the enable secret command, followed by the enable password command, from the console. You log out of the switch and log back in at the console. Which command defines the password that you had to enter to access privileged mode? a. enable password b. enable secret c. Neither d. The password command, if it's configured 2. Some IOS commands store passwords as clear text, but you can then encrypt the passwords with the service password-encryption global command. By comparison, other commands store a computed hash of the password instead of storing the password. Comparing the two options, which one answer is the most accurate about why one method is better than the other? a. Using hashes is preferred because encrypted IOS passwords can be easily decrypted. b. Using hashes is preferred because of the large CPU effort required for encryption. c. Using encryption is preferred because it provides stronger password protection. d. Using encryption is preferred because of the large CPU effort required for hashes. 3. A network engineer issues a show running-config command and sees only one line of output that mentions the enable secret
command, as follows: enable secret 5 \$1SzGMASe8cmvkz4UjJhVp7.mALe1 Which of the following is true about users of this router? a. A user must type \$1SzGMASe8cmvkz4UjJhVp7.mALe1 to reach enable mode. b. The router will hash the clear-text password that the user types to compare to the hashed password. c. A no service password-encryption configuration command would decrypt this password. d. The router will decrypt the password in the configuration to compare to the clear-text password typed by the user. 4. A single-line ACL has been added to a router configuration using the command ip access-list 1 permit 172.16.4.0/23. The configuration also includes the access-class 1 in command in VTY configuration mode. Which answer accurately describes how the router uses ACL 1? a. Hosts in subnet 172.16.4.0/23 alone can telnet into the router. b. CLI users cannot telnet from the router to hosts in subnet 172.16.4.0/23 alone. c. Hosts in subnet 172.16.4.0/23 alone can log in but cannot reach enable mode of the router. d. The router will only forward packets with source addresses in subnet 172.16.4.0/23. 88 CCNA 200-301 Official Cert Guide, Volume 2 5. A next-generation firewall sits at the edge of a company's connection to the Internet. It has been configured to prevent Telnet clients residing in the Internet from accessing Telnet servers inside the company. Which of the following might a next-generation firewall use that a traditional firewall would not? a. Match message destination well-known port 23 b. Match message application data c. Match message IP protocol 23 d. Match message source TCP ports greater than 49152 6. Which actions show a behavior typically supported by a Cisco next-generation IPS (NGIPS) beyond the capabilities of a traditional IPS? (Choose two answers) a. Gather and use host-based information for context b. Comparisons between messages and a database of exploit signatures c. Logging events for later review by the security team d. Filter URIs using reputation scores Foundation Topics Securing IOS Passwords The ultimate way to protect passwords in Cisco IOS devices is to not store passwords in IOS devices. That is, for any functions that can use an external authentication, authorization, and accounting (AAA) server, use it. However, it is common to store some passwords in a router or switch configuration, and this first section of the chapter discusses some of the ways to protect those passwords. As a brief review, Figure 5-1 summarizes some typical login security configuration on a router or switch. On the lower left, you see Telnet support configured, with the use of a password only (no username required). On the right, the configuration adds support for login with both username and password, supporting both Telnet and SSH users. The upper left shows the one command required to define an enable password in a secure manner. Enable enable secret myenablepw Telnet Enable Mode (sw1#) SSH and Telnet User Mode (sw1>) username wendell secret odum ! hostname sw1 ip domain-name example.com crypto key generate rsa line vty 0 15 transport input telnet login password mytelnetpw Figure 5-1 Sample Login Security Configuration line vty 0 15 transport input all login local Chapter 5. Securing Network Devices 89 NOTE The configuration on the far right of the figure supports both SSH and Telnet, but consider allowing SSH only by instead using the transport input ssh command. The Telnet protocol sends all data unencrypted, so any attacker who copies the message with a Telnet login will have a copy of the password. The rest of this first section discusses how to make these passwords secure. In particular, this section looks at ways to avoid keeping clear-text passwords in the configuration and storing the passwords in ways that make it difficult for attackers to learn the password. Encrypting Older IOS Passwords with service password-encryption Some older-style IOS passwords create a security exposure because the passwords exist in the configuration file as clear text. These clear-text passwords might be seen in printed versions of the configuration files, in a backup copy of the configuration file stored on a server, or as displayed on a network engineer's display. Cisco attempted to solve this clear-text problem by adding a command to encrypt those passwords: the service password-encryption global configuration command. This command encrypts passwords that are normally held as clear text, specifically the passwords for these commands: password password (console or vty mode) username name password password (global) enable password password (global) To see how it works, Example 5-1 shows how the service password-encryption command encrypts the clear-text console password. The example uses the show running-config | section line con 0 command both before and after the encryption; this command lists only the section of the configuration about the console. Example 5-1 Encryption and the service password-encryption Command Switch# show running-config | section line con 0 line con 0 password cisco login Switch# configure terminal Enter configuration commands, one per line. End with CNTL/Z. Switch# config# service password-encryption Switch3(config)# ^Z Switch3# show running-config | section line con 0 line con 0 password 7 070C2B5F4D06 login A close examination of the before and after show running-config command output reveals both the obvious effect and a new concept. The encryption process now hides the original 5 90 CCNA 200-301 Official Cert Guide, Volume 2 clear-text password. Also, IOS needs a way to signal that the value in the password command lists an encrypted password rather than the clear text. IOS adds the encryption or encoding type of "7" to the command, which specifically refers to passwords encrypted with the service password-encryption command. (IOS considers the clear-text passwords to be type 0; some commands list the 0, and some do not.) While the service password-encryption global command encrypts passwords, the no service password-encryption global command does not immediately decrypt the passwords back to their clear-text state. Instead, the process works as shown in Figure 5-2. Basically, after you enter the no service password-encryption command, the passwords remain encrypted until you change a password. 1 service password-encryption 2 no service password-encryption 3 Change Password mypass \$T & @ \$ @3 \$T & @ \$ @3 mypass Clear Encrypted Clear Figure 5-2 Encryption Is Immediate; Decryption Awaits Next Password Change Unfortunately, the service password-encryption command does not protect the passwords very well. Armed with the encrypted value, you can search the Internet and find sites with tools to decrypt these passwords. In fact, you can take the encrypted password from this example, plug it into one of these sites, and it decrypts to "cisco." So, the service password-encryption command will slow down the curious, but it will not stop a knowledgeable attacker. Encoding the Enable Passwords with Hashes In the earliest days of IOS, Cisco used the enable password password global command to define the password that users had to use to reach enable mode (after using the enable EXEC command). However, as just noted, the enable password password command stored the password as clear text, and the service password-encryption command encrypted the password in a way that was easily decrypted. Cisco solved the problem of only weak ways to store the password of the enable password password global command by making a more secure replacement: the enable secret password global command. However, both these commands exist in IOS even today. The next few pages look at these two commands from a couple of angles, including interactions between these two commands, why the enable secret command is more secure, along with a note about some advancements in how IOS secures the enable secret password. Interactions Between Enable Password and Enable Secret First, for real life: use the enable secret password global command, and ignore the enable password password global command. That has been true for around 20 years. However, to be complete, Cisco has never removed the much weaker enable password command from IOS. So, on a single switch (or router), you can configure one or the other, Answers to the "Do I Know This Already?" quiz: 1. B 2. A 3. B 4. A 5. B 6. A, D Chapter 5: Securing Network Devices 91 both, or neither. What, then, does the switch expect us to type as the password to reach enable mode? It boils down to these rules: Both commands configured: Users must use the password in the enable secret password command (and ignore the enable password password command). Only one command configured: Use the password in that one command. Neither command configured (default): Console users move directly to enable mode without a password prompt; Telnet and SSH users are rejected with no option to supply an enable password. Making the Enable Secret Truly Secret with a Hash The Cisco enable secret command protects the password value by never even storing the clear-text password in the configuration. However, that one sentence may cause you a bit of confusion: If the router or switch does not remember the clear-text password, how can the switch know that the user typed the right password after using the enable command? This section works through a few basics to show you how and appreciate why the password's value is secret. First, by default, IOS uses a hash function called Message Digest 5 (MD5) to store an alternative value in the configuration, rather than the clear-text password. Think of MD5 as a rather complex mathematical formula. In addition, this formula is chosen so that even if you know the exact result of the formula—that is, the result after feeding the clear-text password through the formula as input—it is computationally difficult to compute the original clear-text
password. Figure 5-3 shows the main ideas: MD5 Hash: F(Clear Text) = Secret Clear Text Computationally Simple 5 et Sec Secret Computationally Difficult F'(Secret) = ClearText Figure 5-3 One-Way Nature of MD5 Hash to Create Secret NOTE "Computationally difficult" is almost a code phrase, meaning that the designers of the function hope that no one is willing to take the time to compute the original clear text. So, if the original clear-text password cannot be re-created, how can a switch or router use it to compare to the clear-text password typed by the user? The answer depends on another fact about these security hashes like MD5: each clear-text input results in a unique result from the math formula. The enable secret fred command generates an MD5 hash. If a user types fred when trying to enter enable mode, IOS will run MD5 against that value and get the same MD5 hash as is listed in the enable secret command, so IOS allows the user to access enable mode. If the user typed any other value besides fred, IOS would compute a different MD5 hash than the value stored with the enable secret command, and IOS would reject that user's attempt to reach enable mode. 5 92 CCNA 200-301 Official Cert Guide, Volume 2 Knowing that fact, the switch can make a comparison when a user types a password after using the enable EXEC command as follows: Step 1. IOS computes the MD5 hash of the password in the enable secret command and stores the hash of the password in the configuration. Step 2. When the user types the enable command to reach enable mode, a password that needs to be checked against that configuration command, IOS hashes the clear-text password as typed by the user. Step 3. IOS compares the two hashed values: if they are the same, the user-typed password must be the same as the configured password. As a result, IOS can store the hash of the password but never store the clear-text password; however, it can still determine whether the user typed the same password. Switches and routers already use the logic described here, but you can see the evidence by looking at the switch configuration. Example 5-2 shows the creation of the enable secret command, with a few related details. This example shows the stored (hashed) value as revealed in the show running-configuration command output. That output also shows that IOS changed the enable secret fred command to list the encryption type 5 (which means the listed password is actually an MD5 hash of the clear-text password). The goodbyegoo long text string is the hash, preventing others from reading the password. Example 5-2 Cisco IOS Encoding Password "cisco" as Type 5 (MD5) Switch3(config)# enable secret fred Switch3(config)# ^Z Switch3# show running-config | include enable secret 5 \$1SzGMASe8cmvkz4UjJhVp7.mALe1 Switch3# configure terminal Enter configuration commands, one per line. End with CNTL/Z. Switch3(config)# no enable secret "cisco" Switch3(config)# ^Z The end of the example also shows an important side point about deleting the enable secret password: after you are in enable mode, you can delete the enable secret password using the no enable secret command, without even having to enter the password value. You can also overwrite the old password by just repeating the enable secret command. But you cannot view the original clear-text password. NOTE Example 5-2 shows another shortcut illustrating how to work through long show command output, this time using the pipe to the include command. The | include enable secret part of the command processes the output from show running-config to include only the lines with the case-sensitive text "enable secret." Improved Hashes for Cisco's Enable Secret The use of any hash function to encode passwords relies on several key features of the particular hash function. In particular, every possible input value must result in a single hashed Chapter 5: Securing Network Devices 93 value, so that when users type a password, only one password value matches each hashed value. Also, the hash algorithm must result in computationally difficult math (in other words, a pain in the neck) to compute the clear-text password based on the hashed value to discourage attackers. The MD5 hash algorithm has been around 30 years. Over those years, computers have gotten much faster, and researchers have found creative ways to attack the MD5 algorithm, making MD5 less challenging to crack. That is, someone who saw your running configuration would have an easier time re-creating your clear-text secret passwords than in the early years of MD5. These facts are not meant to say that MD5 is bad, but like many cryptographic functions before MD5, progress has been made, and new functions were needed. To provide more recent options that would create a much greater challenge to attackers, Cisco added two additional hashes in the 2010s, as noted in Figure 5-4. Type 9 Scrypt Type 0 Clear Type 7 Encrypted Type 5 MD5 1990 Figure 5-4 Type 4 PBKDF2 1995 5 Type 8 PBKDF2 2010 2015 Timeline of Encryptions/Hashes of Cisco IOS Passwords IOS now supports two alternative algorithm types in the more recent router and switch IOS images. Both use an SHA-256 hash instead of MD5, but with two newer options, each of which has some differences in the particulars of how each algorithm works SHA-256. Table 5-2 shows the configuration of all three algorithm types on the enable secret command. Table 5-2 Commands and Encoding Types for the enable secret Command Command Type Algorithm enable [algorithm-type md5] secret password 5 MD5 enable algorithm-type sha256 secret password 8 SHA-256 enable algorithm-type scrypt secret password 9 SHA-256 Example 5-3 shows the enable secret command being changed from MD5 to the scrypt algorithm. Of note, the example shows that only one enable secret command should exist between those three commands in Table 5-2. Basically, if you configure another enable secret command with a different algorithm type, that command replaces any existing enable secret command. Example 5-3 Cisco IOS Encoding Password "mypass1" as Type 9 (SHA-256) R1# show running-config | include enable secret 5 \$1z\$SYj\$725dBzmlUJ0nx9gFPtTtV0 R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# enable algorithm-type scrypt secret mypass1 R1(config)# ^Z 94 CCNA 200-301 Official Cert Guide, Volume 2 R1# R1# show running-config | include enable secret 9 9\$9HlEeEKRiRw91ucE\$NjYOUE5EhoI16AWw2vsYwkJkN6eG8uK/24B0TVU6 R1# Following the process shown in the example, the first command confirms that the current enable secret command uses encoding type 5, meaning it uses MD5. Second, the user configures the password using algorithm type scrypt. The last command confirms that only one enable secret command exists in the configuration now, with encoding type 9. Encoding the Passwords for Local Usernames Cisco added the enable secret command back in the 1990s to overcome the problems with the enable password command. The username password and username secret commands have a similar history. Originally, IOS supported the username user password password command—a command that had those same issues of being a clear-text password or a poorly encrypted value (with the service password-encryption feature). Many years later, Cisco added the username user secret password global command, which encoded the password as an MD5 hash, with Cisco adding support for the newer SHA-256 hashes later. Today, the username secret command is preferred over the username password command; however, IOS does not use the same logic for the username command as it does for allowing both the enable secret plus enable password commands to exist in the same configuration. IOS allows ■ Only one username command for a given username—either a username name password password command or a username name secret password command ■ A mix of commands (username password and username secret) in the same router or switch (for different usernames) You should use the username secret command instead of the username password command when possible. However, note that some IOS features require that the router knows a clear-text password via the username command (for instance, when performing some common authentication methods for serial links called PAP and CHAP). In those cases, you still need to use the username password command. As mentioned, the more recent IOS versions on both switches and routers use the additional encoding options beyond MD5, just as supported with the enable secret command. Table 5-3 shows the syntax of those three options in the username command, with the MD5 option shown as an option because it is the default used with the username secret command. Table 5-3 Commands and Encoding Types for the username secret Command Command Type Algorithm username name [algorithm-type md5] secret password 5 MD5 username name algorithm-type sha256 secret password 8 SHA-256 username name algorithm-type scrypt secret password 9 SHA-256 Chapter 5: Securing Network Devices 95 Controlling Password Attacks with ACLs Attackers can repeatedly try to log in to your network devices to gain access, but IOS has a feature that uses ACLs to prevent the attacker from even seeing a password prompt. When an external user connects to a router or switch using Telnet or SSH, IOS uses a vty line to represent that user connection. IOS can apply an ACL to the vty lines, filtering the addresses that can telnet or SSH into the router or switch. If filtered, the user never sees a login prompt. For example, imagine that all the network engineering staff's devices connect into subnet 10.1.1.0/24. The security policy states that only the network engineering staff is allowed to telnet or SSH into any of the Cisco routers in a network. In such a case, the configuration shown in Example 5-4 could be used on each router to deny access from IP addresses not in that subnet. Example 5-4 vty Access Control Using the access-class Command
line vty 0 login password cisco access-class 3 in ! Next command is a global command that matches IPv4 packets with 1 a source address that begins with 10.1.1. access-list 3 permit 10.1.1.0 0.0.0.255 The access-class command refers to the matching logic in access-list 3. The keyword in Telnet and SSH connections into this router—in other words, people telnetting into this router. As configured, ACL 3 checks the source IP address of packets for incoming Telnet connections. IOS also supports using ACLs to filter outbound Telnet and SSH connections. For example, consider a user who first uses Telnet or SSH to connect to the CLI and now sits in user or enable mode. With an outbound vty filter, IOS will apply ACL logic if the user tries the telnet or ssh commands to connect out of the local device to another device. You configure an outbound VTY ACL, use the access-class acl out command in VTY configuration mode. Once configured, the router filters any attempts made by current vty users to use the telnet and ssh commands to initiate new connections to other devices. Of the two options—to protect inbound and outbound connections—protecting inbound connections is by far the more important and more common. However, to be complete, outbound VTY ACLs have a surprisingly odd feature in how they use the ACL. When the out keyword is used, the standard IP ACL listed in the access-class command actually looks at the destination IP address, and not the source. That is, if filters based on the device to which the telnet or ssh command is trying to connect. Firewalls and Intrusion Prevention Systems The next topic examines the roles of a couple of different kinds of networking devices: firewalls and intrusion prevention systems (IPSs). Both devices work to secure networks but with slightly different goals and approaches. 5 96 CCNA 200-301 Official Cert Guide, Volume 2 This second major section of the chapter takes a look at each. This section first discusses the core traditional features of both firewalls and IPSs. The section closes with a description of the newer features in the current generation of these products, called next-generation products, which improves the functions of each. Traditional Firewalls Traditionally, a firewall sits in the forwarding path of all packets so that the firewall can then choose which packets to discard and which to allow through. By doing so, the firewall protects the network from different kinds of issues by allowing only the intended types of traffic to flow in and out of the network. In fact, in its most basic form, firewalls do the same kinds of work that routers do with ACLs, but firewalls can perform that packet-filtering function with many more options, as well as perform other security tasks. Figure 5-5 shows a typical network design for a site that uses a physical firewall. The figure shows a firewall, like the Cisco Adaptive Security Appliance (ASA) firewall, connected to a Cisco router, which in turn connects to the Internet. All enterprise traffic going to or from the Internet would be sent through the firewall. The firewall would consider its rules and make a choice for each packet, whether the packet should be allowed through. Internet Firewall Figure 5-5 Firewall as Positioned in the Packet Forwarding Path Although firewalls have some router-like features (such as packet forwarding and packet filtering), they provide much more advanced security features than a traditional router. For example, most firewalls can use the following kinds of logic to make the choice of whether to discard or allow a packet: ■ Like router IP ACLs, match the source and destination IP addresses ■ Like router IP ACLs, identify applications by matching their static well-known TCP and UDP ports ■ Watch application-layer flows to know what additional TCP and UDP ports are used by a particular flow, and filter based on those ports ■ Match the text in the URI of an HTTP request—that is, look at and compare the contents of what is often called the web address—and match patterns to decide whether to allow or deny the download of the web page identified by that URI ■ Keep state information by storing information about each packet, and make decisions about filtering future packets based on the historical state information (called stateful inspection, or being a stateful firewall) The stateful firewall feature provides the means to prevent a variety of attacks and is one of the more obvious differences between the ACL processing of a router versus security Chapter 5: Securing Network Devices 97 filtering by a firewall. Routers must spend as little time as possible processing each packet so that the packets experience little delay passing through the router. The router cannot take the time to gather information about a packet, and then for future packets, consider some saved state information about earlier packets when making a filtering decision. Because they focus on network security, firewalls do save some information about packets and can consider that information for future filtering decisions. As an example of the benefits of using a stateful firewall, consider a simple denial of service (DoS) attack. An attacker can make this type of attack against a web server by using tools that create (or start to create) a large volume of TCP connections to the server. The firewall might allow TCP connections to that server normally, but imagine that the server might typically receive 10 new TCP connections per second under normal conditions and 100 per second at the busiest times. A DoS attack might attempt thousands or more TCP connections per second, driving up CPU and RAM use on the server and eventually overloading the server to the point that it cannot serve legitimate users. A stateful firewall could be tracking the number of TCP connections per second—that is, recording state information based on earlier packets—including the number of TCP connection requests from each client IP address to each server address. The stateful firewall could notice a large number of TCP connections, check its state information, and then notice that the number of requests is very large from a small number of clients to that particular server, which is typical of some kinds of DoS attacks. The stateful firewall could then start filtering those packets, helping the web server survive the attack, whereas a stateless firewall or a router ACL would not have had the historical state information to realize that a DoS attack was occurring. Security Zones Firewalls not only filter packets, they also pay close attention to which host initiates communications. That concept is most obvious with TCP as the transport layer protocol, where the client initiates the TCP connection by sending a TCP segment that sets the SYN bit only (as seen in Figure 1-5 in Chapter 1, "Introduction to TCP/IP Transport and Applications"). Firewalls use logic that considers which host initiated a TCP connection by watching these initial TCP segments. To see the importance of who initiates the connections, think about a typical enterprise network with a connection to the Internet, as shown in Figure 5-6. The company has users inside the company who open web browsers, initiating connections to web servers across the Internet. However, by having a working Internet connection, that same company opens up the possibility that an attacker might try to create a TCP connection to the company's internal web servers used for payroll processing. Of course, the company does not want random Internet users or attackers to be able to connect to their payroll server. Firewalls use the concept of security zones (also called a zone for short) when defining which hosts can initiate new connections. The firewall has rules, and those rules define which host can initiate connections from one zone to another zone. Also, by using zones, a firewall can place multiple interfaces into the same zone, in cases for which multiple interfaces should have the same security rules applied. Figure 5-7 depicts the idea with the inside part of the enterprise considered to be in a separate zone compared to the interfaces connected toward the Internet. 5 98 CCNA 200-301 Official Cert Guide, Volume 2 Payroll Server User No! Internet SW Firewall Yes! Cisco Figure 5-6 Web Server Allowing Outbound Connections and Preventing Inbound Connections Zone Inside Zone Outside SW1 SW2 R1 Firewall Internet R2 Rule: Inside Can Initiate to Outside for Ports... Figure 5-7 Using Security Zones with Firewalls The most basic firewall rule when using two zones like Figure 5-7 reduces to this logic: Allow hosts from zone inside to initiate connections to hosts in zone outside, for a predefined set of safe well-known ports (like HTTP port 80, for instance). Note that with this one simple rule, the correct traffic is allowed while filtering the unwanted traffic by default. Firewalls typically disallow all traffic unless a rule specifically allows the packet. So, with this simple rule to allow inside users to initiate connections to the outside zone, and that alone, the firewall also prevents outside users from initiating connections to inside hosts. Most companies have an inside and outside zone, as well as a special zone called the demilitarized zone (DMZ). Although the DMZ name comes from the real world, it has been used in IT for decades to refer to a firewall security zone used to place servers that need to be available for use by users in the public Internet. For example, Figure 5-8 shows a typical Internet edge design, with the addition of a couple of web servers in its DMZ connected through the firewall. The firewall then needs another rule that enables users in the zone outside—that is, users in the Internet—to initiate connections to those web servers in the DMZ. By separating those web servers into the DMZ, away from the rest of the enterprise, the enterprise can prevent Internet users from attempting to connect to the
internal devices in the inside zone, preventing many types of attacks. Chapter 5: Securing Network Devices 99 Zone Inside Zone Outside Initiate Internet Public Web Servers Zone DMZ www.example.com Figure 5-8 Internet Using a DMZ for Enterprise Servers That Need to Be Accessible from the Intrusion Prevention Systems (IPS) Traditionally, a firewall works with a set of user-configured rules about where packets should be allowed to flow in a network. The firewall needs to sit in the path of the packets so it can filter the packets, redirect them for collection and later analysis, or let them continue toward their destination. A traditional intrusion prevention system (IPS) can sit in the path packets take through the network, and it can filter packets, but it makes its decisions with different logic. The IPS first downloads a database of exploit signatures. Each signature defines different header field values found in sequences of packets used by different exploits. Then the IPS can examine packets, compare them to the known exploit signatures, and notice when packets may be part of a known exploit. Once identified, the IPS can log the event, discard packets, or even redirect the packets to another security application for further examination. A traditional IPS differs from firewalls in that instead of an engineer at the company defining rules for that company based on applications (by port number) and zones, the IPS applies the logic based on signatures supplied mostly by the IPS vendor. Those signatures look for these kinds of attacks: ■ DoS ■ DDoS ■ Worms ■ Viruses To accomplish its mission, the IPS needs to download and keep updating its signature database. Security experts work to create the signatures. The IPS must then download the exploit signature database and keep downloading updates over time, as shown in Figure 5-9. 5 100 CCNA 200-301 Official Cert Guide, Volume 2 Signatures Internet IPS Figure 5-9 Firewall IPS and Signature Database For example, think about what happens when an entirely new computer virus has been created. Host-based security products, like antivirus software, should be installed on the computers inside the company. These tools use a similar model as the IPS, keeping an updated database of virus signatures. The signatures might look for patterns in how a computer virus could be stored inside files on the computer, or in files sent to the computer via email or web browsers. But there will be some time lag between the day when the virus has been discovered (called zero-day attacks) and when researchers have developed a virus signature, changed their database, and allowed time for all the hosts to update their antivirus software. The hosts are at risk during this time lag. The IPS provides a complimentary service to prevent viruses. Researchers will look for ways an IPS could recognize the same virus while in flight through the network with new IPS signatures—for instance, looking for packets with a particular port and a particular hex string in the application payload. Once developed, the IPS devices in the network need to be updated with the new signature database, protecting against that virus. Both the host-based and IPS-based protections play an important role, but the fact that one IPS protects sections of a network means that the IPS can sometimes more quickly react to new threats to protect hosts. Cisco Next-Generation Firewalls The CCNA 200-301 exam topics mention the terms firewall and IPS but prefaced with the term next-generation. Around the mid 2010s, Cisco and some of their competitors started using the term next-generation when discussing their security products to emphasize some of the newer features. In short, a next-generation firewall (NGFW) and a next-generation IPS (NGIPS) are the now-current firewall and IPS products from Cisco. However, the use of the term next-generation goes far beyond just a marketing label: the term emphasizes some major shifts and improvements over the years. The security industry sees endless cycles of new attacks followed by new solutions, with some solutions requiring new product features or even new products. Some of the changes that have required new security features include the proliferation of mobile devices—devices that leave the enterprise, connect to the Internet, and return to the Enterprise—creating a whole new level of risk. Also, no single security function or appliance (firewall, IPS, antimalware) can hope to stop some threats, so the next-generation tools must be able to work better together to Chapter 5: Securing Network Devices 101 provide solutions. In short, the next-generation products have real useful features not found in their predecessor products. As for Cisco products, for many years Cisco branded its firewalls as the Cisco Adaptive Security Appliance (ASA). Around 2013, Cisco acquired Sourcefire, a security product company. Many of the next-generation firewall (and IPS) features come from software acquired through that acquisition. As of 2019 (when this chapter was written), all of Cisco's currently sold firewalls have names that evoke memories of the Sourcefire acquisition, with most of the firewall product line being called Cisco Firepower firewalls (www.cisco.com/go/firewalls). An NGFW still does the traditional functions of a firewall, of course, like stateful filtering by comparing fields in the IP, TCP, and UDP headers, and using security zones when defining firewall rules. To provide some insight into some of the newer next-generation features, consider the challenge of matching packets with ports: 1. Each IP-based application should use a well-known port. 2. Attackers know that firewalls will filter most well-known ports from sessions initiated from the outside zone to the inside zone (see Figure 5-8). 3. Attackers use port scanning to find any port that a company's firewall will allow through right now. 4. Attackers attempt to use a protocol of their choosing (for example, HTTP) but with the nonstandard port found through port scanning as a way to attempt to connect to hosts inside the enterprise. The sequence lists a summary of some of the steps attackers need to take but does not list every single task. However, even to this depth, you can see how attackers can find a way to send packets past the corporate firewall. The solution? A next-generation firewall that looks at the application layer data to identify the application instead of relying on the TCP and UDP port numbers used. Cisco performs their deep packet inspection using a feature called Application Visibility and Control (AVC). Cisco AVC can identify many applications based on the data sent (application layer headers plus application data structures far past the TCP and UDP headers). When used with a Cisco NGFW, instead of matching port numbers, the firewall matches the application, defeating attacks like the one just described. The following list mentions a few of the features of an NGFW. Note that while NGFW is a useful term, the line between a traditional firewall and a next-generation firewall can be a bit blurry, as the terms describe products that have gone through repeated changes over long periods of time. This list does summarize a few of the key points, however: ■ Traditional firewall: An NGFW performs traditional firewall features, like stateful firewall filtering, NAT/PAT, and VPN termination. ■ Application Visibility and Control (AVC): This feature looks deep into the application layer data to identify the application. For instance, it can identify the application based on the data, rather than port number, to defend against attacks that use random port numbers. ■ Advanced Malware Protection: NGFW platforms run multiple security services, not just as a platform to run a separate service, but for better integration of functions. A network-based antimalware function can run on the firewall itself, blocking file transfers that would install malware, and saving copies of files for later analysis. 5 102 CCNA 200-301 Official Cert Guide, Volume 2 ■ URL Filtering: This feature 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, with URL filtering being able to use those scores in its decision to categorize, filter, or rate limit. ■ NGIPS: The Cisco NGFW products can also run their NGIPS feature along with the firewall. Note that for any of the services that benefit from being in the same path that packets traverse, like a firewall, it makes sense that over time those functions could migrate to run on the same product. So, when the design needs both a firewall and IPS at the same location in the network, these NGFW products can run the NGIPS feature as shown in the combined device in Figure 5-10. Talos Internet NGIPS & NGFW Figure 5-10 Next-Generation Firewall with Next-Generation IPS Module Cisco Next-Generation IPS The Cisco next-generation IPS (NGIPS) products have followed a similar path as the Cisco NGFW products. Cisco first added NGIPS features primarily through its Sourcefire acquisition, with the now-current (in 2019) Cisco IPS products also using the Firepower name. In fact, as a product line, the hardware NGFW and NGIPS products are the same products, with the ability to run both the NGFW and NGIPS. As with the NGFW, the NGIPS adds features to a traditional IPS. For instance, one of the biggest issues with a traditional IPS comes with the volume of security events logged by the IPS. For instance: 1. An IPS compares the signature database, which lists all known exploits, to all messages. 2. It generates events, often far more than the security staff can read. 3. The staff must mentally filter events to find the proverbial needle in the haystack, possible only through hard work, vast experience, and a willingness to dig. An NGIPS helps with this issue in
a couple of ways. First, an NGIPS examines the context by gathering data from all the hosts and the users of those hosts. The NGIPS will know the OS, software revision levels, what apps are running, open ports, the transport protocols and port numbers in use, and so on. Armed with that data, the NGIPS can make much more intelligent choices about what events to log. Chapter 5: Securing Network Devices 103 For instance, consider an NGIPS placed into a network to protect a campus LAN where end users connect, but no data center exists in that part of the network. Also, all PCs happen to be running Windows, and possibly the same version, by corporate policy. The signature database includes signatures for exploits of Linux hosts, Macs, Windows version nonexistent in that part of the network, and exploits that apply to server applications that are not running on those hosts. After gathering those facts, an NGIPS can suggest de-emphasizing checks for exploits that do not apply to those endpoints, spending more time and focus on events that could occur, greatly reducing the number of events logged. The following list mentions a few of the Cisco NGIPS features: ■ Traditional IPS: An NGIPS performs traditional IPS features, like using exploit signatures to compare packet flows, creating a log of events, and possibly discarding and/or redirecting packets. ■ Application Visibility and Control (AVC): As with NGFWs, an NGIPS has the ability to look deep into the application layer data to identify the application. ■ Contextual Awareness: NGFW platforms gather data from hosts—OS, software version/level, patches applied, applications running, open ports, applications currently sending data, and so on. Those facts inform the NGIPS as to the often more limited vulnerabilities in a portion of the network so that the NGIPS can focus on actual vulnerabilities while greatly reducing the number of logged events. ■ Reputation-Based Filtering: The Cisco Talos security intelligence group researches security threats daily, building the data used by the Cisco security portfolio. Part of that data identifies known bad actors, based on IP address, domain, name, or even specific URL, with a reputation score for each. A Cisco NGIPS can perform reputation-based filtering, taking the scores into account. ■ Event Impact Level: Security personnel need to assess the logged events, so an NGIPS provides an assessment based on impact levels, with characterizations as to the impact if an event is indeed some kind of attack. If you want to learn a little more about these topics for your own interest, let me refer you to a couple of resources. First, check out articles and blog posts from the Cisco Talos Intelligence Group (www.talosintelligence.com). The Cisco Talos organization researches security issues around the globe across the entire spectrum of security products. Additionally, one Cisco Press book has some great information about both next-generation firewalls and IPSs, written at a level appropriate as a next step. If you want to read more, check out this book with the long name: Integrated Security Technologies and Solutions, Volume 1: Cisco Security Solutions for Advanced Threat Protection with Next-Generation Firewall, Intrusion Prevention, AMP, and Content Security (or just use its ISBN, 9781587147067), with one chapter each on NGFW and NGIPS. Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 5-4 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. 5 104 CCNA 200-301 Official Cert Guide, Volume 2 Table 5-4 Chapter Review Tracking Review Element Review Date(s)

Resource Used Review key topics Book, website Review key terms Book, website Repeat DIKTA questions Book, PTP Do labs Blog Review command tables Book Review All the Key Topics Table 5-5 Key Topics for Chapter 5 Key Topic Element Description Page Number List Commands whose passwords are encrypted by service password encryption 89 List Rules for when IOS uses the password set with the enable password versus enable secret commands 91 List Log which IOS can use the enable secret hash when a user types a clear-text password to reach enable mode 92 List Rule for combinations of the username command 94 Figure 5-6 Typical client filtering by IP address at Internet edge 98 Figure 5-8 Firewall security zones with DMZ 99 List Features of next-generation firewalls 101 List Features of next-generation IPSs 103 Key Terms You Should Know enable secret, local username, MD5 hash, username secret, firewall, IPS, next-generation firewall (NGFW), next-generation IPS (NGIPS), Application Visibility and Control Do Labs The Sim Lite software is a version of Pearson's full simulator learning product with a subset of the labs, included free with this book. The Sim Lite with this book includes a couple of labs about various password-related topics. Also, check the author's blog site pages for configuration exercises (Config Labs) at . Command References Tables 5-6 and 5-7 list configuration and verification commands used in this chapter. As an easy review exercise, cover the left column in a table, read the right column, and try to recall the command without looking. Then repeat the exercise, covering the right column, and try to recall what the command does. Chapter 5: Securing Network Devices 105 Table 5-6 Chapter 5 Configuration Commands Command Mode/Purpose/Description line console 0 Command that changes the context to console configuration mode. line vty 1st-vty last-vty Command that changes the context to vty configuration mode for the range of vty lines listed in the command. login Console and vty configuration mode. Tells IOS to prompt for a password. password pass-value Console and vty configuration mode. Lists the password required if the login command is configured. login local Console and vty configuration mode. Tells IOS to prompt for a username and password, to be checked against locally configured username global configuration commands. username name [algorithm]type md5 | sha256 | scrypt secret pass-value Global command. Defines one of possibly multiple usernames and associated passwords, stored as a hashed value (default MD5), with other hash options as well. username name password pass-value Global command. Defines a username and password, stored in clear text in the configuration by default. crypto key generate rsa [modulus 512 | 768 | 1024] Global command. Creates and stores (in a hidden location in flash memory) the keys required by SSH. transport input [telnet | ssh | all | none] vty line configuration mode. Defines whether Telnet and/or SSH access is allowed into this switch. [no] service password-encryption Global command that encrypts all clear-text passwords in the running-config. The no version of the command disables the encryption of passwords when the password is set. enable password pass-value Global command to create the enable password, stored as a clear text instead of a hashed value. enable [algorithm]type md5 | Global command to create the enable password, stored as a hashed value instead of clear text, with the hash defined by sha256 | scrypt secret the algorithm type. pass-value no enable secret no enable password access-class name | name in Table 5-7 Global command to delete the enable secret or enable password commands, respectively. A vty mode command that enables inbound ACL checks against Telnet and SSH clients connecting to the router. Chapter 5 EXEC Command Reference Command Purpose show running-config | section vty Lists the vty lines and subcommands from the configuration. show running-config | section console Lists the console and subcommands from the configuration. show running-config | include enable Lists all lines in the configuration with the word enable. 5 CHAPTER 6 Implementing Switch Port Security This chapter covers the following exam topics: 5.0 Security Fundamentals 5.7 Configure Layer 2 security features (DHCP snooping, dynamic ARP inspection, and port security) In modern networks, security must be implemented in depth. The security architecture should use firewalls and intrusion prevention systems (IPS) at strategic locations, and hosts should use antivirus and antimalware tools. Routers, which already need to exist throughout the enterprise at the edge between local-area networks and wide-area networks, can be configured with IP access control lists to filter packets related to different IP address ranges in that enterprise. LAN switches have a unique opportunity as a security enforcement point, particularly LAN switches connected to endpoint devices. Attackers often launch attacks from the endpoints connected to an enterprise LAN switch. The attacker might gain physical access to the endpoint or first infect the device to then launch an attack. Additionally, a mobile device can become infected while outside the company network and then later connect to the company network, with the attack launching at that point. Engineers should assume that attacks might be launched from end-user devices connected directly to access ports on the enterprise's LAN switches, so Cisco switches include a number of useful tools to help prevent several types of attacks. This chapter discusses one such tool: port security. Chapter 8, "DHCP Snooping and ARP Inspection," discusses two other switch security tools that take advantage of the switch's access layer role, with Chapter 7, "Implementing DHCP," providing the background details needed to understand the tools in Chapter 8. This short chapter takes a straightforward approach to the port security feature. The first section discusses the concepts, configuration, and verification, using the primary port security operational mode: shutdown mode. The second section then discusses some of the intricacies of the three operational modes: shutdown, verify, and restrict. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 6-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions Port Security Concepts and Configuration 1-3 Port Security Violation Modes 4, 5, 1. Which of the following is required when configuring port security with sticky learning? a. Setting the maximum number of allowed MAC addresses on the interface with the switchport port-security maximum interface subcommand. b. Enabling port security with the switchport port-security interface subcommand. c. Defining the specific allowed MAC addresses using the switchport port-security mac-address interface subcommand. d. All the other answers list required commands. 2. A Cisco Catalyst switch connects to what should be individual user PCs. Each port has the same port security configuration, configured as follows: interface range gigabitethernet 0/1 - 24 switchport mode access switchport port-security switchport port-security mac-address sticky Which of the following answers describe the result of the port security configuration created with these commands? (Choose two answers.) a. Prevents unknown devices with unknown MAC addresses from sending data through the switch ports. b. If a user connects a switch to the cable, prevents multiple devices from sending data through the port. c. Will allow any one device to connect to each port and will save that device's MAC address into the startup-config. d. Will allow any one device to connect to each port but will not save that device's MAC address into the startup-config. 3. Which of the following commands list the MAC address table entries for MAC addresses configured by port security? (Choose two answers.) a. show mac address-table dynamic b. show mac address-table c. show mac address-table static d. show mac address-table port-security 108 CCNA 200-301 Official Cert Guide, Volume 2. 4. The show port-security interface f0/1 command lists a port status of secure-down. Which one of the following answers must be true about this interface at this time? a. The show interface status command lists the interface status as connected. b. The show interface status command lists the interface status as err-disabled. c. The show port-security interface command could list a mode of shutdown or restrict, but not protect. d. The show port-security interface command could list a violation counter value of 10. 5. A switch's port Gi0/1 has been correctly enabled with port security. The configuration sets the violation mode to restrict. A frame that violates the port security policy enters the interface, followed by a frame that does not. Which of the following answers correctly describe what happens in this scenario? (Choose two answers.) a. The switch puts the interface into an err-disabled state when the first frame arrives. b. The switch generates syslog messages about the violating traffic for the first frame. c. The switch increments the violation counter for Gi0/1 by 1. d. The switch discards both the first and second frame. Foundation Topics Port Security Concepts and Configuration If the network engineer knows what devices should be cabled and connected to particular interfaces on a switch, the engineer can use port security to restrict that interface so that only the expected devices can use it. This reduces
exposure to attacks in which the attacker connects a laptop to some unused switch port. When that inappropriate device attempts to send frames to the switch interface, the switch can take different actions, ranging from simply issuing informational messages to effectively shutting down the interface. Port security identifies devices based on the source MAC address of Ethernet frames that the devices send. For example, in Figure 6-1, PC1 sends a frame, with PC1's MAC address as the source address. SW1's F0/1 interface can be configured with port security, and if so, SW1 would examine PC1's MAC address and decide whether PC1 was allowed to send frames into port F0/1. F0/1 ! SW1 Frame G0/1 Source = PC1 MAC G0/2 F0/2 SW2 Frame Source = PC2 MAC Figure 6-1 Source MAC Addresses in Frames as They Enter a Switch Chapter 6: Implementing Switch Port Security 109 Port security also has no restrictions on whether the frame came from a local device or was forwarded through other switches. For example, switch SW1 could use port security on its G0/1 interface, checking the source MAC address of the frame from PC2, when forwarded up to SW1 from SW2. Port security has several flexible options, but all operate with the same core concepts. First, switches enable port security per port, with different settings available per port. Each port has a maximum number of allowed MAC addresses, meaning that for all frames entering that port, only that number of different source MAC addresses can be used before port security thinks a violation has occurred. When a frame with a new source MAC address arrives, pushing the number of MAC addresses past the allowed maximum, a port security violation occurs. At that point, the switch takes action—by default, discarding all future incoming traffic on that port. The following list summarizes these ideas common to all variations of port security: ■ It examines frames received on the interface to determine if a violation has occurred. ■ It defines a maximum number of unique source MAC addresses allowed for all frames coming in the interface. ■ It keeps a list and counter of all unique source MAC addresses on the interface. ■ It monitors newly learned MAC addresses, considering those MAC addresses to cause a violation if the newly learned MAC address would push the total number of MAC table entries for the interface past the configured maximum allowed MAC addresses for that port. ■ It takes action to discard frames from the violating MAC addresses, plus other actions depending on the configured violation mode. Those rules define the basics, but port security allows other options as well, including options like these: ■ Define a maximum of three MAC addresses, defining all three specific MAC addresses. ■ Define a maximum of three MAC addresses but allow those addresses to be dynamically learned, allowing the first three MAC addresses defined. ■ Define a maximum of three MAC addresses, predefining one specific MAC address, and allowing two more to be dynamically learned. You might like the idea of predefining the MAC addresses for port security, but finding the MAC address of each device can be a bother. Port security provides a useful compromise using a feature called sticky secure MAC addresses. With this feature, port security learns the MAC addresses off each port so that you do not have to reconfigure the values. It also adds the learned MAC addresses to the port security configuration (in the runningconfig file). This feature helps reduce the big effort of finding out the MAC address of each device. As you can see, port security has a lot of detailed options. The next few sections walk you through these options to pull the ideas together. Configuring Port Security Port security configuration involves several steps. First, port security works on both access ports and trunk ports, but it requires you to statically configure the port as a trunk or an 6 110 CCNA 200-301 Official Cert Guide, Volume 2 access port, rather than let the switch dynamically decide whether to use trunking. The following configuration checklist details how to enable port security, set the maximum allowed MAC addresses per port, and configure the actual MAC addresses: Step 1. Use the switchport mode access or the switchport mode trunk interface subcommands, respectively, to make the switch interface either a static access or trunk interface. Step 2. Use the switchport port-security interface subcommand to enable port security on the interface. Step 3. (Optional) Use the switchport port-security maximum interface subcommand to override the default maximum number of allowed MAC addresses associated with the interface (1). Step 4. (Optional) Use the switchport port-security violation [protect | restrict | shutdown] interface subcommand to override the default action to take upon a security violation (shutdown). Step 5. (Optional) Use the switchport port-security mac-address mac-address interface subcommand to predefine any allowed source MAC addresses for this interface. Use the command multiple times to define more than one MAC address. Step 6. (Optional) Use the switchport port-security sticky interface subcommand to tell the switch to "sticky learn" dynamically learned MAC addresses. To demonstrate how to configure this variety of the settings, Figure 6-2 and Example 6-1 show four examples of port security. Three ports operate as access ports, while port F0/4, connected to another switch, operates as a trunk. Static Fa0/1 Server 1 0200.1111.1111 Sticky Fa0/2 Server 2 0200.2222.2222 Dynamic Fa0/3 Company Controllor Maximum 8 Fa0/4 Figure 6-2 SW2 Port Security Configuration Example Answers to the "Do I Know This Already?" quiz: 1 B 2 B, D 3 B, C 4 B 5 B, C Chapter 6: Implementing Switch Port Security 111 Example 6-1 Variations on Port Security Configuration SW1# show running-config (Lines omitted for brevity) interface FastEthernet0/1 switchport mode access switchport port-security switchport port-security mac-address 0200.1111.1111 ! interface FastEthernet0/2 switchport mode access switchport port-security mac-address sticky ! interface FastEthernet0/3 switchport mode access switchport port-security ! interface FastEthernet0/4 switchport mode trunk switchport port-security switchport port-security maximum 8 First, scan the configuration for all four interfaces in Example 6-1, focusing on the first two interface subcommands in each case. Note that the first three interfaces in the example use the same first two interface subcommands, matching the first two configuration steps noted before Figure 6-2. The switchport port-security command enables port security, with all defaults, with the switchport mode access command meeting the requirement to configure the port as either an access or trunk port. The final port, F0/4, has a similar configuration, except that it has been configured as a trunk rather than as an access port. Next, scan all four interfaces again, and note that the configuration differs on each interface after those first two interface subcommands. Each interface simply shows a different example for perspective. The first interface, FastEthernet 0/1, adds one optional port security subcommand: switchport port-security mac-address 0200.1111.1111, which defines a specific source MAC address. With the default maximum source address setting of 1, only frames with source MAC 0200.1111.1111 will be allowed in this port. When a frame with a source other than 0200.1111.1111 enters F0/1, the switch would normally perform MAC address learning and want to add the new source MAC address to the MAC address table. Port security will see that action as learning one too many MAC addresses on the port, taking the default violation action to disable the interface. As a second example, FastEthernet 0/2 uses the same logic as FastEthernet 0/1, except that it uses the sticky learning feature. For port F0/2, the configuration of the switchport port-security mac-address sticky command tells the switch to dynamically learn source MAC addresses and add port-security commands to the running-config. Example 6-2 shows the running-config file that lists the sticky-learned MAC address in this case. 6 112 CCNA 200-301 Official Cert Guide, Volume 2 Example 6-2 Configuration Added by the Port Security Sticky Feature SW1# show running-config interface f0/2 Building configuration... Current configuration: 188 bytes ! interface FastEthernet0/2 switchport mode access switchport port-security switchport port-security mac-address sticky switchport port-security mac-address sticky 0200.2222.2222 Port security does not save the configuration of the sticky addresses, so use the copy running-config startup-config command if desired. The other two interfaces in Example 6-1 do not predefine MAC addresses, nor do they sticky-learn the MAC addresses. The only difference between these two interfaces' port security configuration is that FastEthernet 0/4 supports eight MAC addresses because it connects to another switch and should receive frames with multiple source MAC addresses. Interface F0/3 uses the default maximum of one MAC address. NOTE Switches can also use port security on voice ports and EtherChannels. For voice ports, make sure to configure the maximum MAC address to at least two (one for the phone, or for a PC connected to the phone). On EtherChannels, the port security configuration should be placed on the port-channel interface, rather than the individual physical interfaces in the channel. Verifying Port Security The show port-security interface command provides the most insight to how port security operates, as shown in Example 6-3. This command lists the configuration settings for port security on an interface; plus it lists several important facts about the current operation of port security, including information about any security violations. The two commands in the example show interfaces F0/1 and F0/2, based on Example 6-1's configuration.
Example 6-3 Interfaces Using Port Security to Define Correct MAC Addresses of Particular SW1# show port-security interface fastEthernet 0/1 Port Security : Enabled Port Status : Secure-shutdown Violation Mode : Shutdown Aging Time : 0 mins Aging Type : Absolute SecureStatic Address Aging : Disabled Maximum MAC Addresses : 1 Total MAC Addresses : 1 Configured MAC Addresses : 1 Sticky MAC Addresses : 0 Last Source Address:Vlan : 0013.197b.5004:1 Chapter 6: Implementing Switch Port Security 113 Security Violation Count : 1 SW1# show port-security interface fastEthernet 0/2 Port Security : Enabled Port Status : Secure-up Violation Mode : Shutdown Aging Time : 0 mins Aging Type : Absolute SecureStatic Address Aging : Disabled Maximum MAC Addresses : 1 Total MAC Addresses : 1 Configured MAC Addresses : 1 Sticky MAC Addresses : 1 Last Source Address:Vlan : 0200.2222.2222:1 Security Violation Count : 0 The two commands in Example 6-3 confirm that a security violation has occurred on FastEthernet 0/1, but no violations have occurred on FastEthernet 0/2. The show port-security interface fastethernet 0/1 command shows that the interface is in a secure-shutdown state, which means that the interface has been disabled because of port security. In this case, another device connected to port F0/1, sending a frame with a source MAC address other than 0200.1111.1111, is causing a violation. However, port Fa0/2, which used sticky learning, simply learned the MAC address used by Server 2. Port Security MAC Addresses To complete this chapter, take a moment to think about Layer 2 switching, along with all those examples of output from the show mac address-table dynamic EXEC command. Once a switch port has been configured with port security, the switch no longer considers MAC addresses associated with that port as being dynamic entries as listed with the show mac address-table dynamic EXEC command. Even if the MAC addresses are dynamically learned, once port security has been enabled, you need to use one of these options to see the MAC table entries associated with ports using port security: ■ show mac address-table secure: Lists MAC addresses associated with ports that use port security ■ show mac address-table static: Lists MAC addresses associated with ports that use port security, as well as any other statically defined MAC addresses Example 6-4 proves the point. It shows two commands about interface F0/2 from the port security example shown in Figure 6-2 and Example 6-1. In that example, port security was configured on F0/2 with sticky learning, so from a literal sense, the switch learned a MAC address off that port (0200.2222.2222). However, the show mac address-table dynamic command does not list the address and port because IOS considers that MAC table entry to be a static entry. The show mac address-table secure command does list the address and port. 6 114 CCNA 200-301 Official Cert Guide, Volume 2 Example 6-4 Security Using the secure Keyword to See MAC Table Entries When Using Port SW1# show mac address-table secure interface F0/2 Mac Address Table ----- Vlan Mac Address Type Ports ----- 0200.2222.2222 STATIC Fa0/2 1 Total MAC Addresses for this criterion: 1 SW1# show mac address-table dynamic interface f0/2 Mac Address Table ----- Vlan Mac Address Type Ports ----- SW1# Port Security Violation Modes The first half of the chapter discussed many details of port security, but it mostly ignored one major feature: the port security violation mode. The violation mode defines how port security should react when a violation occurs. First, to review, what is a port security violation? Any received frame that breaks the port security rules on an interface. For example: ■ For an interface that allows any two MAC addresses, a violation occurs when the total of preconfigured and learned MAC addresses on the interface exceeds the configured maximum of two. ■ For an interface that predefines all the specific MAC addresses allowed on the interface, a violation occurs when the switch receives a frame whose source MAC is not one of those configured addresses. With port security, each switch port can be configured to use one of three violation modes that defines the actions to take when a violation occurs. All three options cause the switch to discard the offending frame (a frame whose source MAC address would push the number of learned MAC addresses over the limit). However, the modes vary in how many other steps they take. For instance, some modes include the action of the switch generating syslog messages and SNMP Trap messages, while some define the action to disable the interface. Table 6-2 lists the three modes, their actions, along with the keywords that enable each mode on the switchport port-security violation [protect | restrict | shutdown] interface subcommand. Chapter 6: Implementing Switch Port Security 115 Table 6-2 Actions When Port Security Violation Occurs Option on the switchport port-security violation Command Protect Restrict Shutdown Discards offending traffic Yes Yes Sends log and SNMP messages No Yes Disables the interface by putting it in an err-disabled state, discarding all traffic No No Yes Because IOS reacts so differently with shutdown mode as compared to restrict and protect modes, the next few pages explain the differences—first for shutdown mode, then for the other two modes. Port Security Shutdown Mode When the (default) shutdown violation mode is used and a port security violation occurs on a port, port security stops all frame forwarding on the interface, both in and out of the port. In effect, it acts as if port security has shut down the port; however, it does not literally configure the port with the shutdown interface subcommand. Instead, port security uses the err-disabled feature. Cisco switches use the err-disabled state for a wide range of purposes, but when using port security shutdown mode and a violation occurs, the following happens: ■ The switch interface state (per show interfaces and show interfaces status) changes to an err-disabled state. ■ The switch interface port security state (per show port-security) changes to a securedown state. ■ The switch stops sending and receiving frames on the interface. Once port security has placed a port in err-disabled state, by default the port remains in an err-disabled state until someone takes action. To recover from an err-disabled state, the interface must be shut down with the shutdown command and then enabled with the no shutdown command. Alternately, the switch can be configured to automatically recover from the err-disabled state, when caused by port security, with these commands: ■ errdisable recovery cause psecure-violation: A global command to enable automatic recovery for interfaces in an err-disabled state caused by port security ■ errdisable recovery interval seconds: A global command to set the time to wait before recovering the interface To take a closer look at shutdown mode, start by checking the configuration state of the switch. You can check the port security configuration on any interface with the show port-security interface type number command, as seen back in Example 6-2, but the show port-security command (as listed in Example 6-5) shows briefer output, with one line per enabled interface. 6 116 CCNA 200-301 Official Cert Guide, Volume 2 Confirming the Port Security Violation Mode Example 6-5 SW1# show port-security Secure Port MaxSecureAddr CurrentAddr (Count) (Count) SecurityViolation Security Action (Count) -----Fa0/13 1 1 Shutdown -----Total Addresses in System (excluding one mac per port) : 0 Max Addresses limit in System (excluding one mac per port) : 8192 Note that for these examples, a switch has configured port security on port Fa0/13 only. In this case, the switch appears to be configured to support one MAC address, has already reached that total, and has a security violation action of "shutdown." Next, Example 6-6 shows the results after a port security violation has already occurred on port F0/13. The first command confirms the err-disabled state (per the show interfaces status command) and the secure-shutdown state (per the show port-security command). Example 6-6 Port Security Status in Shutdown Mode After a Violation ! The next lines show the log message generated when the violation occurred. Jul 31 18:00:22.810 :%PORT_SECURITY-2/PSECURE_VIOLATION: Security violation occurred, caused by MAC address d4bc.b57d.8200 on port FastEthernet0/13 ! The next command shows the err-disabled state, implying a security violation. SW1# show interfaces Fa0/13 status Port Name Status Vlan Duplex err-disabled 1 auto Fa0/13 Speed auto Type 10/100BaseTX ! The next command's output has shading for several of the most important facts. SW1# show port-security interface Fa0/13 Port Security : Enabled Port Status : Secure-shutdown Violation Mode : Shutdown Aging Time : 0 mins Aging Type : Absolute SecureStatic Address Aging : Disabled Maximum MAC Addresses : 1 Total MAC Addresses : 1 Configured MAC Addresses : 1 Sticky MAC Addresses : 0 Last Source Address:Vlan : 0200.3333.3333:3 Security Violation Count : 1 The output of the show port-security interface command lists the current port-security status (secure-shutdown) as well as the configured mode (shutdown). The last line of output lists the number of violations that caused the interface to fail to an err-disabled state, while Chapter 6: Implementing Switch Port Security 117 the second-to-last line identifies the MAC address and VLAN of the device that caused the violation. Figure 6-3 summarizes these behaviors, assuming the same scenario shown in the example. F0/13: Status: Err-disabled Secure-Down Counter: 1 Syslog: 10 Msgs 10X Source: MAC1 Last MAC: MAC1 show port-security interface Figure 6-3 Summary of Actions: Port Security Violation Mode Shutdown Note that the violations counter notes the number of times the interface has been moved to the err-disabled (secure-shutdown) state.
For instance, the first time it fails, the counter increments to 1; while err-disabled, many frames can arrive, but the counter remains at 1. Later, after an engineer has recovered the interface from the err-disabled state with a shutdown/no shutdown, another violation that causes the interface to fail to an err-disabled state will cause the counter to increment to 2. Port Security Protect and Restrict Modes The restrict and protect violation modes take a much different approach to securing ports. These modes still discard offending traffic, but the interface remains in a connected (up) state and in a port security state of secure-up. As a result, the port continues to forward good traffic but discards offending traffic. Having a port in a seemingly good state that also discards traffic can be a challenge when troubleshooting. Basically, you have to know about the feature and then know how to tell when port security is discarding some traffic on a port even though the interface status looks good. With protect mode, the only action the switch takes for a frame that violates the port security rules is to discard the frame. The switch does not change the port to an err-disabled state, does not generate messages, and does not even increment the violations counter. Example 6-7 shows a sample with protect mode after several violations have occurred. Note that the show command confirms the mode (protect) as configured in the top part of the example, with a port security state of secure-up—a state that will not change in protect mode. Also, note that the counter at the bottom shows 0, even though several violations have occurred, because protect mode does not count the violating frames. Example 6-7 Port Security Using Protect Mode SW1# show running-config ! Lines omitted for brevity interface FastEthernet0/13 switchport mode access switchport port-security 6 118 CCNA 200-301 Official Cert Guide, Volume 2 switchport port-security mac-address 0200.1111.1111 switchport port-security violation protect ! Lines omitted for brevity SW1# show port-security interface Fa0/13 Port Security : Enabled Port Status : Secure-up Violation Mode : Protect Aging Time : 0 mins Aging Type : Absolute SecureStatic Address Aging : Disabled Maximum MAC Addresses : 1 Total MAC Addresses : 1 Configured MAC Addresses : 1 Sticky MAC Addresses : 0 Last Source Address:Vlan : 0000.0000.0000:0 Security Violation Count : 0 NOTE The small particulars of the violation counters and last source address might be slightly different with some older switch models and IOS versions. Note that this edition's testing is based on 2960XR switches running IOS 15.2.(6)E2. IOS shutdown mode disables the interface, and protect mode does nothing more than discard the offending traffic, restrict mode provides a compromise between the other two modes. If Example 6-7 had used the restrict violation mode instead of protect, the port status would have also remained in a secure-up state; however, IOS would show some indication of port security activity, such as an accurate incrementing violation counter, as well as syslog messages. Example 6-8 shows an example of the violation counter and ends with an example port security syslog message. In this case, 97 incoming frames so far violated the rules, with the most recent frame having a source MAC address of 0200.3333.3333 in VLAN 1. Example 6-8 Port Security Using Violation Mode Restrict SW1# show port-security interface fa0/13 Port Security : Enabled Port Status : Secure-up Violation Mode : Restrict Aging Time : 0 mins Aging Type : Absolute SecureStatic Address Aging : Disabled Maximum MAC Addresses : 1 Total MAC Addresses : 1 Configured MAC Addresses : 1 Sticky MAC Addresses : 0 Last Source Address:Vlan : 0200.3333.3333:1 Security Violation Count : 97 Chapter 6: Implementing Switch Port Security 119 ! The following log message also points to a port security issue. ! 01:46:58 :%PORT_SECURITY-2/PSECURE_VIOLATION: Security violation occurred, caused by MAC address 0200.3333.3333 on port FastEthernet0/13. Figure 6-4 summarizes the key points about the restrict mode for port security. In this case, the figure matches the same scenario as the example again, with 97 total violating frames arriving so far, with the most recent being from source MAC address MAC3. F0/13: Status: Connected Secure-Up Counter: +97 97X Syslog: 97 Msgs Source: MAC3 Last MAC: MAC3 3 show port-security interface Figure 6-4 Summary of Actions: Port Security Violation Mode Restrict Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 6-3 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 6-3 Chapter Review Tracking Review Element Review Date(s) Resource Used Review key topics Book, website Review key terms Book, website Answer DIKTA questions Book, PTP Review command tables Book Review memory tables Book, website Review config checklists Book, website Do labs Sim Lite, blog Watch Video Website 120 CCNA 200-301 Official Cert Guide, Volume 2 Review All the Key Topics Table 6-4 Key Topics for Chapter 6 Key Topic Element Description Page Number List Summary of port security concepts 109 List Port security configuration checklist 110 Example 6-1 Port security configuration samples 111 Table 6-2 Port security actions and the results of each action 115 List Switch actions when a port security violation occurs 115 Key Terms You Should Know port security, violation mode, err-disabled (err-disable) Do Labs The Sim Lite software is a version of Pearson's full simulator learning product with a subset of the labs, included free with this book. The Sim Lite with this book includes a couple of labs about port security. Also, check the author's blog site pages for configuration exercises (Config Labs) at . Command References Tables 6-5 and 6-6 list configuration and verification commands used in this chapter. As an easy review exercise, cover the left column in a table, read the right column, and try to recall the command without looking. Then repeat the exercise, covering the right column, and try to recall what the command does. Table 6-5 Chapter 6 Configuration Command Reference Command Mode/Purpose/Description switchport mode [access | trunk] Interface configuration mode command that tells the switch to always be an access port, or always be a trunk port switchport port-security mac-address mac-address Interface configuration mode command that statically adds a specific MAC address as an allowed MAC address on the interface switchport port-security mac-address sticky Interface subcommand that tells the switch to learn MAC addresses on the interface and add them to the configuration for the interface as secure MAC addresses switchport port-security maximum value Interface subcommand that sets the maximum number of static secure MAC addresses that can be assigned to a single interface switchport port-security violation [protect | restrict | shutdown] Interface subcommand that tells the switch what to do if an inappropriate MAC address tries to access the network through a secure switch port Chapter 6: Implementing Switch Port Security 121 Command Mode/Purpose/Description errdisable recovery cause psecure-violation Global command that enables the automatic recovery from err-disabled state for ports that reach that state due to port security violations errdisable recovery interval seconds Global command that sets the delay, in seconds, before a switch attempts to recover an interface in err-disabled mode, regardless of the reason for that interface being in that state shutdown Interface subcommands that administratively disable and enable an interface, respectively no shutdown Table 6-6 Chapter 6 EXEC Command Reference Command Purpose show running-config Lists the currently used configuration show running-config | interface type number Displays the running-configuration excerpt of the listed interface and its subcommands only show mac address-table dynamic [interface type number] Lists the dynamically learned entries in the switch's address (forwarding) table show mac address-table secure [interface type number] Lists MAC addresses defined or learned on ports configured with port security show mac address-table static [interface type number] Lists static MAC addresses and MAC addresses learned or defined with port security show interfaces [interface type number] status Lists one output line per interface (or for only the listed interface if included), noting the description, operating state, and settings for duplex and speed on each interface show port-security interface type number Lists an interface's port security configuration settings and security operational status show port-security Lists one line per interface that summarizes the port security settings for any interface on which it is enabled 6 CHAPTER 7 Implementing DHCP This chapter covers the following exam topics: 1.0 Network Fundamentals 1.10 Identify IP parameters for Client OS (Windows, Mac OS, Linux) 4.0 IP Services 4.3 Explain the role of DHCP and DNS within the network 4.6 Configure and verify DHCP client and relay In the world of TCP/IP, the word host refers to any device with an IP address: your phone, your tablet, a PC, a server, a router, a switch—any device that uses IP to provide a service or just needs an IP address to be managed. The term host includes some less-obvious devices as well: the electronic advertising video screen at the mall, your electrical power meter that uses the same technology as mobile phones to submit your electrical usage information for billing, your new car. No matter the type of host, any host that uses IPv4
needs four IPv4 settings to work properly: ■ IP address ■ Subnet mask ■ Default routers ■ DNS server IP addresses This chapter discusses these basic IP settings on hosts. The chapter begins by discussing how a host can dynamically learn these four settings using the Dynamic Host Configuration Protocol (DHCP). The second half of this chapter then shows how to find the settings on hosts and the key facts to look for when displaying the settings. Just a note about the overall flow of the chapters: This chapter does not discuss security topics, although it sits inside Part II, "Security Services." I located this DHCP-focused chapter here because Chapter 8, "DHCP Snooping and ARP Inspection," relies heavily on knowledge of DHCP. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 7-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions Dynamic Host Configuration Protocol 1-4 Identifying Host IPv4 Settings 5, 6, 1. A PC connects to a LAN and uses DHCP to lease an IP address for the first time. Of the usual four DHCP messages that flow between the PC and the DHCP server, which ones do the client send? (Choose two answers.) a. Acknowledgment b. Discover c. Offer d. Request 2. Which of the following kinds of information are part of a DHCP server configuration? (Choose two answers.) a. Ranges of IP addresses in subnets that the server should lease b. Ranges of IP addresses to not lease per subnet c. DNS server hostnames d. The default router IP and MAC address in each subnet 3. Which answers list a criterion for choosing which router interfaces need to be configured as a DHCP relay agent? (Choose two answers.) a. If the subnet off the interface does not include a DHCP server b. If the subnet off the interface does include a DHCP server c. If the subnet off the interface contains DHCP clients d. If the router interface already has an ip address dhcp command 4. A router connects to an Internet Service Provider (ISP) using its G0/0/0 interface, with the ip address dhcp command configured. What does the router do with the DHCPlearned default gateway information? a. The router ignores the default gateway value learned from the DHCP server. b. The router uses the default gateway just like a host, ignoring its routing table. c. The router forwards received packets based on its routing table but uses its default gateway setting to forward packets it generates itself. d. The router adds a default route based on the default gateway to its IP routing table. 124 CCNA 200-301 Official Cert Guide, Volume 2 5. In the following excerpt from a command on a Mac, which of the following parts of the output represent information learned from a DHCP server? (Choose two answers.) Macprompt# ifconfig en0 En1: flags=8863 mtu 1500 options=10b ether 00:6d:e7:b1:9a:11 inet 172.16.4.2 netmask 0xffff00 broadcast 172.16.4.255 a. 00:6d:e7:b1:9a:11 b. 172.16.4.2 c. 0xffff00 d. 172.16.4.255 6. Which of the following commands on a Windows OS should list both the IP address and DNS servers as learned with DHCP? a. ifconfig b. ipconfig c. ifconfig /all d. ipconfig /all Foundation Topics Dynamic Host Configuration Protocol Dynamic Host Configuration Protocol (DHCP) provides one of the most commonly used services in a TCP/IP network. The vast majority of hosts in a TCP/IP network are user devices, and the vast majority of user devices learn their IPv4 settings using DHCP. Using DHCP has several advantages over the other option of manually configuring IPv4 settings. The configuration of host IP settings sits in a DHCP server, with each client learning these settings using DHCP messages. As a result, the host IP configuration is controlled by the IT staff, rather than on local configuration on each host, resulting in fewer user errors. DHCP allows both the permanent assignment of host addresses, but more commonly, DHCP assigns a temporary lease of IP addresses. With these leases, the DHCP server can reclaim IP addresses when a device is removed from the network, making better use of the available addresses. DHCP also enables mobility. For example, every time a user moves to a new location with a tablet computer—to a coffee shop, a client location, or back at the office—the user's device can connect to another wireless LAN, use DHCP to lease a new IP address in that LAN, and begin working on the new network. Without DHCP, the user would have to ask for information about the local network and configure settings manually, with more than a few users making mistakes. Although DHCP works automatically for user hosts, it does require some preparation from the network, with some configuration on routers. In some enterprise networks, that router Chapter 7: Implementing DHCP 125 configuration can be a single command on many of the router's LAN interfaces (ip helper-address server-ip), which identifies the DHCP server by its IP address. In other cases, the router acts as the DHCP server. Regardless, the routers have some role to play. This first major section of the chapter takes a tour of DHCP, including concepts and the router configuration to enable the routers to work well with a separate DHCP server. DHCP Concepts Sit back for a moment and think about the role of DHCP for a host computer. The host acts as a DHCP client. As a DHCP client, the host begins with no IPv4 settings—no IPv4 address, no mask, no default router, and no DNS server IP addresses. But a DHCP client does have knowledge of the DHCP protocol, so the client can use that protocol to (a) discover a DHCP server and (b) request to lease an IPv4 address. DHCP uses the following four messages between the client and server. (Also, as a way to help remember the messages, note that the first letters spell DORA): Discover: Sent by the DHCP client to find a willing DHCP server Offer: Sent by a DHCP server to offer to lease to that client a specific IP address (and inform the client of its other parameters) Request: Sent by the DHCP client to ask the server to lease the IPv4 address listed in the Offer message Acknowledgment: Sent by the DHCP server to assign the address and to list the mask, default router, and DNS server IP addresses to DHCP clients, however, have a somewhat unique problem: they do not have an IP address yet, but they need to send these DHCP messages inside IP packets. To make that work, DHCP messages make use of two special IPv4 addresses that allow a host that has no IP address to still be able to send and receive messages on the local subnet: 0.0.0.0. An address reserved for use as a source IPv4 address for hosts that do not yet have an IP address. 255.255.255.255: The local broadcast IP address. Packets sent to this destination address are broadcast on the local data link, but routers do not forward them. To see how these addresses work, Figure 7-1 shows an example of the IP addresses used between a host (A) and a DHCP server on the same LAN. Host A, a client, sends a Discover message, with source IP address of 0.0.0.0 because host A does not have an IP address to use yet. Host A sends the packet to destination 255.255.255.255, which is sent in a LAN broadcast frame, reaching all hosts in the subnet. The client hopes that there is a DHCP server on the local subnet. Why? Packets sent to 255.255.255.255 only go to hosts in the local subnet; router R1 will not forward this packet. NOTE Figure 7-1 shows one example of the addresses that can be used in a DHCP request. This example shows details assuming the DHCP client chooses to use a DHCP option called the broadcast flag; all examples in this book assume the broadcast flag is used. 7 126 CCNA 200-301 Official Cert Guide, Volume 2 A R1 1 R2 B Discover DHCP Server 172.16.1.11 To 255.255.255.255 From 0.0.0.0 Offer 2 To 255.255.255.255 From 172.16.1.11 Figure 7-1 DHCP Discover and Offer Now look at the Offer message sent back by the DHCP server. The server sets the destination IP address to 255.255.255.255 again. Why? Host A still does not have an IP address, so the server cannot send a packet directly to host A. So, the server sends the packet to "all local hosts in the subnet" address (255.255.255.255). (The packet is also encapsulated in an Ethernet broadcast frame.) Note that all hosts in the subnet receive the Offer message. However, the original Discover message lists a number called the client ID, which includes the host's MAC address, that identifies the original host (host A in this case). As a result, host A knows that the Offer message is meant for host A. The rest of the hosts will receive the Offer message, but notice that the message lists another device's DHCP client ID, so the rest of the hosts ignore the Offer message. Supporting DHCP for Remote Subnets with DHCP Relay Network engineers have a major design choice to make with DHCP: Do they put a DHCP server in every LAN subnet or locate a DHCP server in a central site? The question is legitimate. Cisco routers can act as the DHCP server, so a distributed design could use the router at each site as the DHCP server. With a DHCP server in every subnet, as shown in Figure 7-1, the protocol flows stay local to each LAN. However, a centralized DHCP server approach has advantages as well. In fact, some Cisco design documents suggest a centralized design as a best practice, in part because it allows for centralized control and configuration of all the IPv4 addresses assigned throughout the enterprise. With a centralized DHCP server, those DHCP messages that flowed only on the local subnet in Figure 7-1 somehow need to flow over the IP
network to the centralized DHCP server and back. To make that work, the routers connected to the remote LAN subnets need an interface subcommand: the ip helper-address server-ip command. The ip helper-address server-ip subcommand tells the router to do the following for the messages coming in an interface, from a DHCP client: Answers to the "Do I Know This Already?" quiz: 1 B, D 2 A, B 3 A, C 4 D 5 B, C 6 D Chapter 7: Implementing DHCP 127 1. Watch for incoming DHCP messages, with destination IP address 255.255.255.255. 2. Change that packet's source IP address to the router's incoming interface IP address. 3. Change that packet's destination IP address to the address of the DHCP server (as configured in the ip helper-address command). 4. Route the packet to the DHCP server. This command gets around the "do not route packets sent to 255.255.255.255" rule by changing the destination IP address. Once the destination has been set to match the DHCP server's IP address, the network can route the packet to the server. NOTE This feature, by which a router relays DHCP messages by changing the IP addresses in the packet header, is called DHCP relay. Figure 7-2 shows an example of the process. Host A sits on the left, as a DHCP client. The DHCP server (172.16.2.11) sits on the right. R1 has an ip helper-address 172.16.2.11) command configured, under its G0/0 interface. At step 1, router R1 notices the incoming DHCP packet destined for 255.255.255.255. Step 2 shows the results of changing both the source and destination IP address, with R1 routing the packet. ip helper-address 172.16.2.11 B A 172.16.1.1 G0/0 1 Discover 7 R1 2 To 255.255.255.255 From 0.0.0.0 Figure 7-2 R1 Discover To 172.16.2.11 From 172.16.1.11 IP Helper Address Effect The router uses a similar process for the return DHCP messages from the server. First, for the return packet from the DHCP server, the server simply reverses the source and destination IP addresses of the packet received from the router (relay agent). For example, in Figure 7-2, the Discover message lists source IP address 172.16.1.1, so the server sends the Offer message back to destination IP address 172.16.1.1. When a router receives a DHCP message, addressed to one of the router's own IP addresses, the router realizes the packet might be part of the DHCP relay feature. When that happens, the DHCP relay agent (router R1) needs to change the destination IP address, so that the real DHCP client (host A), which does not have an IP address yet, can receive and process the packet. Figure 7-3 shows one example of how these addresses work, when R1 receives the DHCP Offer message sent to R1's own 172.16.1.1 address. R1 changes the packet's destination to 255.255.255.255 and forwards it out G0/0, because the packet was destined to G0/0's 172.16.1.1 IP address. As a result, all hosts in that LAN (including the DHCP client A) will receive the message. Many enterprise networks use a centralized DHCP server, so the normal router configuration includes an ip helper-address command on every LAN interface/subinterface. With that standard configuration, user hosts off any router LAN interface can always reach the DHCP server and lease an IP address. 128 CCNA 200-301 Official Cert Guide, Volume 2 172.16.1.1 G0/0 offer R1 2 To 255.255.255.255 From 172.16.1.1 Figure 7-3 R2 offer 1 To 172.16.1.1 From 172.16.2.11 S1 DHCP Server 172.16.2.11 IP Helper Address for the Offer Message Returned from the DHCP Server A DHCP server might sound like some large piece of hardware, sitting in a big locked room with lots of air conditioning to keep the hardware cool.

and moves the port to an err-disabled state. Also, the feature can be enabled both on trusted and untrusted interfaces. Although rate limiting DHCP messages can help, plugging the port in an err-disabled state can itself create issues. As a reminder, once in the err-disabled state, the switch will not send or receive frames for the interface. However, the err-disabled state might be too severe an action because the default recovery action for an err-disabled state requires the configuration of a shutdown and then a no shutdown subcommand on the interface. To help strike a better balance, you can enable DHCP Snooping rate limiting and then also configure the switch to automatically recover from the port's err-disabled state, without the need for a shutdown and then no shutdown command. Example 8-3 shows how to enable DHCP Snooping rate limits and err-disabled recovery. First, look at the lower half of the configuration, to the interfaces, to see the straightforward setting of the per-interface limits using the ip dhcp snooping rate limit number interface subcommands. The top of the configuration uses two global commands to let IOS to recover from an err-disabled state if it is caused by DHCP Snooping, and to use a nondefault number of seconds to wait before recovering the interface. Note that the configuration in Example 8-3 would rely on the core configuration for DHCP Snooping as shown in Example 8-1. Chapter 8: DHCP Snooping and ARP Inspection 155 Example 8-3 Configuring DHCP Snooping Message Rate Limits errisable recovery cause dhcp-rate-limit errisable recovery interval 30 ! interface GigabitEthernet1/0/2 ip dhcp snooping limit rate 10 ! interface GigabitEthernet1/0/3 ip dhcp snooping limit rate 2 A repeat of the show ip dhcp snooping command now shows the rate limits near the end of the output, as noted in Example 8-4. Example 8-4 Confirming DHCP Snooping Rate Limits SW2# show ip dhcp snooping ! Lines omitted for brevity Interface Trusted Allow Option Rate Limit (pps) ----- GigabitEthernet1/0/2 yes yes 10 no no 2 Custom circuit-ids: GigabitEthernet1/0/3 Custom circuit-ids: DHCP Snooping Configuration Summary The following configuration checklist summarizes the commands included in this section about how to configure DHCP Snooping. Step 1. Configure this pair of commands (both required): A. Use the ip dhcp snooping global command to enable DHCP Snooping on the switch. Config Checklist B. Use the ip dhcp snooping vlan-vlan-list global command to identify the VLANs on which to use DHCP Snooping. Step 2. (Optional): Use the no ip dhcp snooping information option global command on Layer 2 switches to disable the insertion of DHCP Option 82 data into DHCP messages, specifically on switches that do not act as a DHCP relay agent. Step 3. Configure the ip dhcp snooping trust subcommand to override the default setting of not trusted. Step 4. (Optional): Configure DHCP Snooping rate limits and err-disabled recovery. Step A. (Optional): Configure the ip dhcp snooping limit rate number interface subcommand to set a limit of DHCP messages per second. Step B. (Optional): Configure the no ip dhcp snooping limit rate number interface subcommand to remove an existing limit and reset the interface to use the default of no rate limit. 8 156 CCNA 200-301 Official Cert Guide, Volume 2 Step C. (Optional): Configure the errisable recovery cause dhcp-rate-limit global command to enable the feature of automatic recovery from err-disabled mode, assuming the switch placed the port in err-disabled state because of exceeding DHCP Snooping rate limits. Step D. (Optional): Configure the errisable recovery interval seconds global commands to set the time to wait before recovering from an interface err-disabled state (regardless of the cause of the err-disabled state). Dynamic ARP Inspection The Dynamic ARP Inspection (DAI) feature on a switch examines incoming ARP messages on untrusted ports to filter those it believes to be part of an attack. DAI's core feature compares incoming ARP messages with two sources of data: the DHCP Snooping binding table and any configured ARP ACLs. If the incoming ARP message does not match the tables in the switch, the switch discards the ARP message. This section follows the same sequence as with the DHCP Snooping section, first examining the concepts behind DAI and ARP attacks, and then showing how to configure DAI with both required and optional features. DAI Concepts To understand the attacks DAI can prevent, you need to be ready to compare normal ARP operations with the abnormal use of ARP used in some types of attacks. This section uses that same flow, first reviewing a few important ARP details, and then showing how an attacker can just send an ARP reply—called a gratuitous ARP—triggering hosts to add incorrect ARP entries to their ARP tables. Review of Normal IP ARP If all you care about is how ARP works normally, with no concern about attacks, you can think of ARP to the depth shown in Figure 8-9. The figure shows a typical sequence. Host PC1 needs to send an IP packet to its default router (R2), so PC1 first sends an ARP request message in an attempt to learn the MAC address associated with R2's 172.16.2.12 address. Router R2 sends back an ARP reply, listing R2's MAC address (note the figure shows pseudo MAC addresses to save space). 1 ARP Request 172.16.2.12 MAC 2 2 172.16.2.101 MAC 1 ARP Reply PC 1 SW2 R2 ARP IP 172.16.2.101 Figure 8-9 ARP MAC MAC 1 3 3 IP 172.16.2.2 MAC MAC 2 Legitimate ARP Tables After PC1 DHCP and ARP with Router R2 Chapter 8: DHCP Snooping and ARP Inspection 157 The ARP tables at bottom of the figure imply an important fact: both hosts learn the other host's MAC address with this two-message flow. Not only does PC1 learn R2's MAC address based on the ARP reply (message 2), but router R2 learns PC1's IP and MAC address because of the ARP request (message 1). To see why, take a look at the more detailed view of those messages as shown in Figure 8-10. 172.16.2.101 MAC 2 172.16.2.101 MAC 1 G1/0/3 PC 1 SW2 R2 Ethernet 1 6RXUFH0S& Origin IP MAC 1 172.16.2.101 Dest. MAC % FDWV Ethernet Origin IP 172.16.2.2 Dest. MAC Target IP MAC 1 172.16.2.101 Origin HW MAC 2 Target IP 172.16.2.2 Origin HW MAC 1 Target HW ??? ARP Request ARP 6RXUFH0S& MAC 2 ARP 2 Target HW HW MAC 1 ARP Reply Figure 8-10 A Detailed Look at ARP Request and Reply The ARP messages define origin IP and hardware (MAC) address fields as well as target IP and hardware address fields. The origin should list the sending device's IP address and MAC, no matter whether the message is an ARP reply or ARP request. For instance, message 1 in the figure, sent by PC1, lists PC1's IP and MAC addresses in the origin fields, which is why router R2 could learn that information. PC2 likewise learns of R2's MAC address per the origin address fields in the ARP reply. Gratuitous ARP as an Attack Vector Normally, a host uses ARP when it knows the IP address of another host and wants to learn that host's MAC address. However, for legitimate reasons, a host might also want to inform all the hosts in the subnet about its MAC address. That might be useful when a host changes its MAC address, for instance. So, ARP supports the idea of a gratuitous ARP message with these features: ■ It is an ARP reply. ■ It is sent without having first received an ARP request. ■ It is sent to an Ethernet destination broadcast address so that all hosts in the subnet receive the message. For instance, if a host's MAC address is MAC A, and it changes to MAC B, to cause all the other hosts to update their ARP tables, the host could send a gratuitous ARP that lists an origin MAC of MAC B. Attackers can take advantage of gratuitous ARPs because they let the sending host make other hosts change their ARP tables. Figure 8-11 shows just such an example initiated by PC A 8 158 CCNA 200-301 Official Cert Guide, Volume 2 (an attacker) with a gratuitous ARP. However, this ARP lists PC1's IP address but a different device's MAC address (PC A) at step 1, causing the router to update its ARP table (step 2). 172.16.2.101 MAC 1 172.16.2.11 MAC 2 G1/0/3 PC 1 SW2 R2 G1/0/5 R2 ARP Table IP 172.16.2.101 MAC MAC 1 A 2 PC A MAC A 1 ARP Reply IP = 172.16.2.101 MAC = MAC A Figure 8-11 Nefarious Use of ARP Reply Causes Incorrect ARP Data on R2 At this point, when R2 forwards IP packets to PC1's IP address (172.16.2.101), R2 will encapsulate them in an Ethernet frame with PC A as the destination rather than with PC1's MAC address. At first, this might seem to stop PC1 from working, but instead it could be part of a man-in-the-middle attack so that PC A can copy every message. Figure 8-12 shows the idea of what happens at this point. 1. PC1 sends messages to some server on the left side of router R2. 2. The server replies to PC1's IP address, but R2 forwards that packet to PC A's MAC address, rather than to PC1. 3. PC A copies the packet for later processing. 4. PC A forwards the packet inside a new frame to PC1 so that PC1 still works. 1 172.16.2.11 MAC 2 To MAC 2 172.16.2.101 MAC 1 PC 1 R2 SW2 2 To MAC A MAC A 2 To MAC 1 MAC A 3 Man-in-the-Middle Man-in-the-Middle Attack Resulting from Gratuitous ARP Dynamic ARP Inspection Logic DAI has a variety of features that can prevent these kinds of ARP attacks. To understand how, consider the sequence of a typical client host with regards to both DHCP and ARP. When a host does not have an IP address yet—that is, before the DHCP process Chapter 8: DHCP Snooping and ARP Inspection 159 completes—it does not need to use ARP. Once the host leases an IP address and learns its subnet mask, it needs ARP to learn the MAC addresses of other hosts or the default router in the subnet, so it sends some ARP messages. In short, DHCP happens first, then ARP. DAI takes an approach for untrusted interfaces that confirms an ARP's correctness based on DHCP Snooping's data about the earlier DHCP messages. The correct normal DHCP messages list the IP address leased to a host as well as that host's MAC address. The DHCP Snooping feature also records those facts into the switch's DHCP Snooping binding table. For any DAI untrusted ports, DAI compares the ARP message's origin IP and origin MAC address fields to the DHCP Snooping binding table. If found in the table, DAI allows the ARP through, but if not, DAI discards the ARP. For instance, Figure 8-13 shows step 1 in which the attacker at PC A attempts the gratuitous ARP shown earlier in Figure 8-11. At step 2, DAI makes a comparison to the DHCP Snooping binding table, not finding a match with MAC A along with IP address 172.16.2.101, so DAI would discard the message. 172.16.2.11 MAC 2 172.16.2.101 MAC 1 G1/0/3 PC 1 SW2 R2 G1/0/5 DHCP Snooping Binding Table MAC MAC 1 IP 172.16.2.101 PC A Int. G1/0/3 MAC A 2 1 origin = 172.16.2.101 origin = MAC A No Match Figure 8-13 DAI Filtering ARP Based on DHCP Snooping Binding Table DAI works with the idea of trusted and untrusted ports with the same general rules as DHCP Snooping. Access ports connected to end-user devices are often untrusted by both DHCP Snooping and DAI. Ports connected to other switches, routers, the DHCP server—anything other than links to end-user devices—should be trusted by DAI. Note that although DAI can use the DHCP Snooping table as shown here, it can also use similar statically configured data that lists correct pairs of IP and MAC addresses via a tool called ARP ACLs. Using ARP ACLs with DAI becomes useful for ports connected to devices that use static IP addresses rather than DHCP. Note that DAI looks for both the DHCP Snooping binding data and ARP ACLs. Beyond that core feature, note that DAI can optionally perform other checks as well. For instance, the Ethernet header that encapsulates the ARP should have addresses that match the ARP origin and target MAC addresses. Figure 8-14 shows an example of the comparison of the Ethernet source MAC address and the ARP message origin hardware field. 8 160 CCNA 200-301 Official Cert Guide, Volume 2 Source MAC to ARP Origin Check Source MAC MAC 2 Origin IP 172.16.2.2 Dest. MAC Target IP MAC 1 172.16.2.101 Ethernet Header Figure 8-14 Origin HW MAC 2 Target HW HW MAC 1 ARP Message DAI Filtering Checks for Source MAC Addresses DAI can be enabled to make the comparisons shown in the figure, discarding these messages: ■ Messages with an Ethernet header source MAC address that is not equal to the ARP origin hardware (MAC) address ■ ARP reply messages with an Ethernet header destination MAC address that is not equal to the ARP target hardware (MAC) address ■ Messages with unexpected IP addresses in the two ARP IP address fields Finally, like DHCP Snooping, DAI does its work in the switch CPU rather than in the switch ASIC, meaning that DAI itself can be more susceptible to DoS attacks. The attacker could generate large numbers of ARP messages, driving up CPU usage in the switch. DAI can avoid these problems through rate limiting the number of ARP messages on a port over time. Dynamic ARP Inspection Configuration Configuring DAI requires just a few commands, with the usual larger variety of optional configuration settings. This section examines DAI configuration, first with mostly default settings and with reliance on DHCP Snooping. It then shows a few of the optional features, like rate limits, automatic recovery from err-disabled state, and how to enable additional checks of incoming ARP messages. Configuring ARP Inspection on a Layer 2 Switch Before configuring DAI, you need to think about the feature and make a few decisions based on your goals, topology, and device roles. The decisions include the following: ■ Choose whether to rely on DHCP Snooping, ARP ACLs, or both. ■ If using DHCP Snooping, configure it and make the correct ports trusted for DHCP Snooping. ■ Choose the VLAN(s) on which to enable DAI. ■ Make DAI trusted (rather than the default setting of untrusted) on select ports in those VLANs, typically for the same ports you trusted for DHCP Snooping. All the configuration examples in this section use the same sample network used in the DHCP Snooping configuration topics, repeated here as Figure 8-15. Just as with DHCP Snooping, switch SW2 on the right should be configured to trust the port connected to the router (G1/0/2), but not trust the two ports connected to the PCs. Chapter 8: DHCP Snooping and ARP Inspection 161 DHCP Relay Agent R1 R2 PC1 G1/0/4 PC2 G1/0/2 SW2 Trusted DHCP Server Figure 8-15 G1/0/3 Untrusted Sample Network Used in ARP Inspection Configuration Examples Example 8-5 shows the required configuration to enable DAI on switch SW2 in Figure 8-15—a configuration that follows a similar progression compared to DHCP Snooping. All ports in the figure connect to VLAN 11, so to enable DAI in VLAN 11, just add the ip arp inspection vlan 11 global command. Then, to change the logic on port G1/0/2 (connected to the router) to be trusted by DAI, add the ip arp inspection trust interface subcommand. Example 8-5 IP ARP Inspection Configuration to Match Figure 8-15 ip arp inspection vlan 11 interface GigabitEthernet1/0/2 ip arp inspection trust Example 8-5 confirms DAI, but it omits both DHCP Snooping and ARP ACLs. (If you were to configure a switch only with commands shown in Example 8-5, the switch would filter all ARPs entering all untrusted ports in VLAN 11.) Example 8-6 shows a complete and working DAI configuration that adds the DHCP Snooping configuration to match the DAI configuration in Example 8-5. Note that Example 8-6 combines Example 8-5's earlier DHCP Snooping configuration for this same topology to the DAI configuration just shown in Example 8-5, with highlights for the DAI-specific configuration lines. Example 8-6 IP DHCP Snooping Configuration Added to Support DAI ip arp inspection vlan 11 ip dhcp snooping ip dhcp snooping vlan 11 no ip dhcp snooping information option ! interface GigabitEthernet1/0/2 ip dhcp snooping trust ip arp inspection trust Remember, DHCP occurs first with DHCP clients, and then they send ARP messages. With the configuration in Example 8-6, the switch builds its DHCP Snooping binding table by analyzing incoming DHCP messages. Next, any incoming ARP messages on DAI untrusted ports must have matching information in that binding table. Example 8-7 confirms the key facts about correct DAI operation in this sample network based on the configuration in Example 8-6. The show ip arp inspection command gives both configuration settings along with status variables and counters. For instance, the 8 162 CCNA 200-301 Official Cert Guide, Volume 2 highlighted lines show the total ARP messages received on untrusted ports in that VLAN and the number of dropped ARP messages (currently 0). Example 8-7 SW2 IP ARP Inspection Status SW2# show ip arp inspection Source Mac Validation : Disabled Destination Mac Validation : Disabled IP Address Validation : Disabled Vlan Configuration Operation ACL Match Static ACL ----- Enabled Active Vlan ACL Logging DHCP Logging Probe Logging ----- Deny Deny Off 11 11 Vlan Forwarded Dropped DHCP Drops ACL Drops ----- 11 59 0 0 0 Vlan DHCP Permits ACL Permits Probe Permits Source MAC Failures ----- 11 7 0 49 0 Vlan Dest MAC Failures IP Validation Failures Invalid Protocol Data ----- Vlan Dest MAC Failures IP Validation Failures Invalid Protocol Data ----- 11 0 0 0 SW2# show ip dhcp snooping binding MacAddress IpAddress Lease(sec) Type VLAN Interface ----- 02:00:11:11:11:11 172.16.2.101 86110 dhcp-snooping 11 GigabitEthernet1/0/3 02:00:22:22:22 172.16.2.102 86399 dhcp-snooping 11 GigabitEthernet1/0/4 2 Total number of bindings: 2 The end of Example 8-7 shows an example of the show ip dhcp snooping binding command on switch SW2. Note that the first two columns list a MAC and IP address as learned from the DHCP messages. Then, imagine an ARP message arrives from PC1, a message that should list PC1's 0200.1111.1111 MAC address and 172.16.2.101 as the origin MAC and IP address, respectively. Per this output, the switch would find that matching data and allow the ARP message. Example 8-8 shows some detail of what happens when switch SW2 receives an invalid ARP message on port G1/0/4 in Figure 8-15. In this case, it will create the invalid ARP message. Chapter 8: DHCP Snooping and ARP Inspection 163 PC2 in the figure was configured with a static IP address of 172.16.2.101 (which is PC1's DHCP-leased IP address). The highlights in the log message at the top of the example show PC2's claimed origin MAC and origin IP addresses in the ARP message. If you refer back to the bottom of Example 8-7, you can see this output on MAC/IP pair does not exist in the DHCP Snooping binding table, so DAI rejects the ARP message. Example 8-8 Sample Results from an ARP Attack Jul 25 14:28:20.763: %SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on G1/0/4, vlan 11 [0200.2222.2222/172.16.2.101/0000.0000/172.16.2.1/09:28:20 EST Jul 25 2019] SW2# show ip arp inspection statistics Vlan Forwarded Dropped DHCP Drops ACL Drops ----- 11 59 17 17 0 Vlan DHCP Permits ACL Permits Probe Permits Source MAC Failures ----- 11 7 0 49 0 Vlan Dest MAC Failures IP Validation Failures Invalid Protocol Data ----- 11 0 0 0 The statistics from the show ip arp inspection command also confirms that the switch has dropped some ARP messages. The highlighted lines in the middle of the table show 17 total dropped ARP messages in VLAN 11. That same highlighted line confirms that it dropped all 17 because of the DHCP Snooping binding table ("DHCP Drops"), with zero dropped due to an ARP ACL ("ACL Drops"). Limiting DAI Message Rates Like DHCP Snooping, DAI can also be the focus of a DoS attack with the attacker generating a large number of ARP messages. Like DHCP Snooping, DAI supports the configuration of rate limits to help prevent those attacks, with a reaction to place the port in an err-disabled state, and with the ability to configure automatic recovery from that err-disabled state. The DHCP Snooping and DAI rate limiters do have some small differences in operation, defaults, and in configuration, as follows: ■ DAI defaults to use rate limits for all interfaces (trusted and untrusted), with DHCP Snooping defaulting to not use rate limits. ■ DAI allows the configuration of a burst interval (a number of seconds), so that the rate limit can have logic like "x ARP messages over y seconds" (DHCP Snooping does not define a burst setting). It helps to look at DAI and DHCP Snooping rate limit configuration together to make comparisons, so Example 8-9 shows both. The example repeats the exact same DHCP Snooping 8 164 CCNA 200-301 Official Cert Guide, Volume 2 commands in earlier Example 8-3 but adds the DAI configuration (highlighted). The configuration in Example 8-7 could be added to the configuration shown in Example 8-6 for a complete DHCP Snooping and DAI configuration. Example 8-9 Configuring ARP Inspection Message Rate Limits errisable recovery cause dhcp-rate-limit errisable recovery cause arp-inspection errisable recovery interval 30 ! interface GigabitEthernet1/0/2 ip dhcp snooping limit rate 10 ip arp inspection limit rate 8 ! interface GigabitEthernet1/0/3 ip dhcp snooping limit rate 2 ip arp inspection limit rate 8 ! burst interval 4 Example 8-10 lists output that confirms the configuration settings. For instance, Example 8-9 configures port G1/0/2 with a rate of 8 messages for each (default) burst of 1 second; the output in Example 8-10 for interface G1/0/2 also lists a rate of 8 and burst interval of 1. Similarly, Example 8-9 configures port G1/0/3 with a rate of 8 over a burst of 4 seconds, with Example 8-10 confirming those same values for port G1/0/3. Note that the other two interfaces in Example 8-10 show the default settings of a rate of 15 messages over a onesecnd burst. Example 8-10 Confirming ARP Inspection Rate Limits SW2# show ip arp inspection interfaces Interface Trust State Rate (pps) Burst Interval ----- G1/0/1 Untrusted 15 G1/0/2 Trusted 8 1 G1/0/3 Untrusted 8 4 G1/0/4 Untrusted 15 1 ! Lines omitted for brevity Configuring Optional DAI Message Checks As mentioned in the section titled "Dynamic ARP Inspection Logic," DAI always checks the ARP message's origin MAC and origin IP address fields versus some table in the switch, but it can also perform other checks. Those checks require more CPU, but they also help prevent other types of attacks. Example 8-11 shows how to configure those three additional checks. Note that you can configure one, two, or all three of the options; just configure the ip arp inspection validate command again with all the options you want in one command, and it replaces the previous global configuration command. The example shows the three options, with the src-mac (source mac) option configured. Chapter 8: DHCP Snooping and ARP Inspection 165 Example 8-11 Confirming ARP Inspection Rate Limits SW2# configure terminal Enter configuration commands, one per line. End with CNTL/Z. SW2(config)# ip arp inspection validate ? dst-mac Validate destination MAC address ip Validate IP addresses src-mac Validate source MAC address SW2(config)# ? Z SW2# SW2# show ip arp inspection Source Mac Validation : Enabled Destination Mac Validation : Disabled IP Address Validation : Disabled IP ARP Inspection Configuration Summary The following configuration checklist summarizes the commands included in this section about how to configure Dynamic ARP Inspection: Config Checklist Step 1. Use the ip arp inspection vlan-vlan-list global command to enable Dynamic ARP Inspection (DAI) on the switch for the specified VLANs. Step 2. Separate from the DAI configuration, also configure DHCP Snooping and/or ARP ACLs for use by DAI. Step 3. Configure the ip arp inspection trust interface subcommand to override the default setting of not trusted. Step 4. (Optional): Configure DAI rate limits and err-disabled recovery. Step A. (Optional): Configure the ip arp inspection limit rate number [burst interval seconds] interface subcommand to set a limit of ARP messages per second, or ARP messages for each configured interval. Step B. (Optional): Configure the ip arp inspection limit rate none interface subcommand to disable rate limits. Step C. (Optional): Configure the errisable recovery cause arp-inspection global command to enable the feature of automatic recovery from err-disabled mode, assuming the switch placed the port in err-disabled state because of exceeding DAI rate limits. Step D. (Optional): Configure the errisable recovery interval seconds global commands to set the time to wait before recovering from an interface err-disabled state (regardless of the cause of the err-disabled state). Step 5. (Optional): Configure the ip arp inspection validate [dst-mac] [src-mac] [ip] global command to add DAI validation steps. 8 166 CCNA 200-301 Official Cert Guide, Volume 2 Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 8-2 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 8-2 Chapter Review Tracking Review Element Review Date(s) Resource Used Review key topics Book, website Review key terms Book, website Answer DIKTA questions Book, PTP Review config checklists Book, website Review All the Key Topics Table 8-3 Key Topics for Chapter 8 Key Topic Element Description Page Number Figure 8-4 DHCP filtering actions on trusted and untrusted ports 149 List DHCP Snooping logic 149 Figure 8-6 DHCP Snooping binding table concept 151 Example 8-1 DHCP Snooping configuration 152 List DHCP Snooping configuration checklist 155 Figure 8-10 Detail inside ARP messages with origin and target 157 List Gratuitous ARP details 157 Figure 8-13 Core Dynamic ARP Inspection logic 159 Example 8-6 Dynamic ARP Inspection configuration with associated DHCP Snooping configuration 161 List Dynamic ARP Inspection checklist 165 Key Terms You Should Know DHCP Snooping, trusted port, untrusted port, DHCP Snooping binding table, Dynamic ARP Inspection, (ARP) origin IP address, (ARP) origin hardware address, ARP reply, gratuitous ARP Command References Tables 8-4 and 8-5 list the configuration and verification commands used in this chapter. As an easy review exercise, cover the left column in a table, read the right column, and try to recall the command without looking. Then repeat the exercise, covering the right column, and try to recall what the command does. Chapter 8: DHCP Snooping and ARP Inspection 167 Table 8-4 Chapter 8 Configuration Command Reference Command Mode/Purpose/Description ip dhcp snooping Global command that enables DHCP Snooping if combined with enabling it on one or more VLANs ip dhcp snooping vlan-vlan-list Global command that lists VLANs on which to enable DHCP Snooping, assuming the ip dhcp snooping command is also configured [no] ip dhcp snooping information option Command that enables (or disables with no option) the feature of inserting DHCP option 82 parameters by the switch when also using DHCP Snooping [no] ip dhcp snooping trust interface subcommand that sets the DHCP Snooping trust state for an interface (default no, or untrusted) ip dhcp snooping limit rate Interface subcommand that sets a limit to the number of incoming DHCP messages processed on an interface, per number second, before DHCP Snooping discards all other incoming DHCP messages in that same second err-disable recovery cause dhcp-rate-limit Global command that enables the switch to automatically recover an err-disabled interface if set to that state because of exceeding a DHCP rate limit setting err-disable recovery interval seconds Global command that sets the number of seconds IOS waits before recovering any err-disabled interfaces which, per various configuration settings, should be recovered automatically err-disable recovery cause arp-inspection Global command that enables the switch to automatically recover an err-disabled interface if set to that state because of an ARP Inspection violation Table 8-5 Chapter 8 EXEC Command Reference Command Purpose show ip dhcp snooping Lists a large variety of DHCP Snooping configuration settings show ip dhcp snooping statistics Lists counters regarding DHCP Snooping behavior on the switch show ip dhcp snooping binding Displays the contents of the dynamically created DHCP Snooping binding table show ip arp inspection Lists both configuration settings for Dynamic ARP Inspection (DAI) as well as counters for ARP messages processed and filtered show ip arp inspection statistics Lists the subset of the show ip arp inspection command output that includes counters 8 Part II Review Keep track of your part review progress with the checklist shown in Table P2-1. Details on each task follow the table. Table P2-1 Part II Review Checklist Activity 1st Date Completed 2nd Date Completed Repeat All DIKTA Questions Answer Part Review Questions Review Key Topics Do Labs Review Videos Repeat All DIKTA Questions For this task, use the PTP software to answer the "Do I Know This Already?" questions again for the chapters in this part of the book. Answer Part Review Questions For this task, use PTP to answer the Part Review questions for this part of the book. Review Key Topics Review all key topics in all chapters in this part, either by browsing the chapters or by using the Key Topics application on the companion website. Use Per-Chapter Interactive Review Elements Using the companion website, browse through the interactive review elements, such as memory tables and key term flashcards, to review the content from each chapter. Labs Depending on your chosen lab tool, here are some suggestions for what to do in the lab: Pearson Network Simulator: If you use the full Pearson CCNA simulator, focus more on the configuration scenario and troubleshooting scenario labs associated with the topics in this part of the book. These types of labs include a larger set of topics and work well as Part Review activities. (See the Introduction for some details about how to find which labs are about topics in this part of the book.) Blog Config Labs: The author's blog () includes a series of configuration-focused labs that you can do on paper, each in 10–15 minutes. Review and perform the labs for this part of the book by using the menus to navigate to the perchapter content and then finding all config labs related to that chapter. (You can see more detailed instructions at Other: If using other lab tools, here are a few suggestions: make sure to experiment with the variety of configuration topics in this part, including router and switch passwords, switch port security, Dynamic ARP Inspection, and DHCP Snooping. Watch Videos Two chapters in this part mention videos included as extra material related to those chapters. Check out the reference in Chapter 4 to a video about using RADIUS protocol, as well as Chapter 6's reference to a video about troubleshooting switch port security. Part III shifts to a variety of topics that can be found in most every network. None are required for a network to work, but many have related to be useful services. Most happen to use IP or support the IP network in some way, so Part III groups the topics together as IP Services. Part III begins and ends with chapters that examine a series of smaller topics. First, Chapter 9 examines several IP services for which the CCNA exam requires you to develop configuration and verification skills. Those services include logging and syslog, the Network Time Protocol (NTP), as well as two related services: CDP and LLDP. Chapter 12, at the end of Part III, closes with another series of smaller topics—although the CCNA 200-301 exam topics require only conceptual knowledge, not configuration skills for these topics. This chapter includes First Hop Redundancy Protocols (FHRPs), Simple Network Management Protocol (SNMP), and two related protocols: TFTP and FTP. The two middle chapters in Part III also focus on IP-based services, beginning with Chapter 10's examination of Network Address Translation (NAT). Almost every network uses NAT with IPv4, although in many cases, the firewall implements NAT. This chapter shows how to configure and verify NAT in a Cisco router. Chapter 11 at first may give the appearance of a large chapter about one topic—Quality of Service—and it does focus on QoS; however, QoS by nature includes a wide variety of individual QoS tools. This chapter walks you through the basic concepts of the primary QoS features. Part III IP Services Chapter 9: Device Management Protocols Chapter 10: Network Address Translation Chapter 11: Quality of Service (QoS) Chapter 12: Miscellaneous IP Services Part III Review CHAPTER 9 Device Management Protocols This chapter covers the following exam topics: 2.0 Network Access 2.3 Configure and verify Layer 2 discovery protocols (Cisco Discovery Protocol and LLDP) 4.0 IP Services 4.2 Configure and verify NTP operating in a client and server mode 4.5 Describe the use of syslog features including facilities and levels This chapter begins Part III with a discussion of the concepts, configuration, and verification of three functions found on Cisco routers and switches. These functions focus more on managing the network devices themselves than on managing the network that devices create. The first major section of this chapter focuses on log messages and syslog. Most computing devices have a need to notify the administrator of any significant issue; generally, across the world of computing, messages of this type are called log messages. Cisco devices generate log messages as well. The first section shows how a Cisco device handles those messages and how you can configure routers and switches to ignore the messages or save them in different ways. Next, different router and switch functions benefit from synchronizing their time-of-day clocks. Like most every computing device, routers and switches have an internal clock function to keep time. Network Time Protocol (NTP) provides a means for devices to synchronize their time, as discussed in the second section. The final major section focuses on two protocols that do the same kinds of work: Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP). Both provide a means for network devices to learn about neighboring devices, without requiring that IPv4 or IPv6 be working at the time. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 9-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions System Message Logging (Syslog) 1–2 Network Time Protocol (NTP) 3–4 Analyzing Topology Using CDP and LLDP 5–6 1. What level of logging to the console is the default for a Cisco device? a. Informational b. Errors c. Warnings d. Debugging 2. What command limits the messages sent to a syslog server to levels 4 through 7? a. logging trap 0-4 b. logging trap 0,1,2,3,4 c. logging trap 4 d. logging trap through 4 3. Which of the following is accurate about the NTP client function on a Cisco router? a. The client synchronizes its time-of-day clock based on the NTP server. b. It counts CPU cycles of the local router CPU to more accurately keep time. c. The client synchronizes its serial line clock rate based on the NTP server. d. The client must be connected to the same subnet as an NTP server. 4. The only NTP configuration on router R1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 5. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 6. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 7. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 8. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 9. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 10. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 11. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 12. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 13. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 14. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 15. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 16. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 17. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 18. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 19. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 20. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 21. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 22. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 23. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 24. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 25. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 26. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 27. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 28. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 29. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 30. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 31. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 32. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 33. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 34. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 35. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 36. The ntp server 10.1.1.1 is the ntp server 10.1.1.1 command. Which answers describe how NTP works on the router? a. As an NTP server only b. As an NTP client only c. As an NTP server only after the NTP client synchronizes with NTP server 10.1.1.1 d. As an NTP server regardless of whether the NTP client synchronizes with NTP server 10.1.1.1 37

Rcv hello from 2.2.2.2 area 0 172.16.2.2 Chapter 9: Device Management Protocols 181 *Aug 10 13:38:22.843: OSPF-1 HELLO Gi0/2: Send hello to 224.0.0.5 area 0 from 172.16.2.1 R1# The console user sees the log messages created on behalf of that debug command after the debug command completes. Per the earlier configuration in Example 9-2, R1's logging console 7 command tells us that the console user will receive severity levels 0-7, which includes level 7 debug messages. Note that with the current settings, these debug messages would not be in the local log message buffer (because of the level in the logging buffered warning command), nor would they be sent to the syslog server (because of the level in the logging trap 4 command). Note that the console user automatically sees the log messages as shown in Example 9-4. However, as noted in the text describing Figure 9-1, a user who connects to R1 would need to also issue the terminal monitor command to see those debug messages. For instance, anyone logged in with SSH at the time Example 9-4's output was gathered would not have seen the output, even with the logging monitor debug command configured on router R1, without first issuing a terminal monitor command. Note that all enabled debug options use router CPU, which can cause problems for the router. You can monitor CPU use with the show process cpu command, but you should use caution when using debug commands carefully on production devices. Also, note the more CLI users that receive debug messages, the more CPU that is consumed. So, some installations choose to not include debug-level log messages for console and terminal logging, requiring users to look at the logging buffer or syslog for those messages, just to reduce router CPU load. Network Time Protocol (NTP) Each networking device has some concept of a date and a time-of-day clock. For instance, the log messages discussed in the first major section of this chapter had a timestamp with the date and time of day listed. Now imagine looking at all the log messages from all routers and switches stored at a syslog server. All those messages have a date and timestamp, but how do you make sure the timestamps are consistent? How do you make sure that all devices synchronize their time-of-day clocks so that you can make sense of all the log messages at the syslog server? How could you make sense of the messages for an event that impacted devices in three different time zones? For example, consider the messages on two routers, R1 and R2, as shown in Example 9-6. Routers R1 and R2 do not synchronize their clocks. A problem keeps happening on the serial link between the two routers. A network engineer looks at all the log messages as stored on the syslog server. However, when the engineer sees some messages from R1, at 13:38:39 (around 1:40 p.m.), he does not think to look for messages from R2 that have a timestamp of around 9:45 a.m. Example 9-6 Log Messages from Routers R1 and R2, Compared *Oct 19 13:38:37.568: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Serial0/0/0 from FULL to DOWN, Neighbor Down: Interface down or detached *Oct 19 13:38:40.568: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0/0, changed state to down | These messages happened on router R2 Oct 19 09:44:09.027: %LINK-3-UPDOWN: Interface Serial0/0/1, changed state to down Oct 19 09:44:09.027: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Serial0/0/1 from FULL to DOWN, Neighbor Down: Interface down or detached 9 182 CCNA 200-301 Official Cert Guide, Volume 2 In reality, the messages in both parts of Example 9-6 happened within 0.5 seconds of each other because I issued a shutdown command on one of the routers. However, the two routers' time-of-day clocks were not synchronized, which makes the messages on the two routers look unrelated. With synchronized clocks, the two routers would have listed practically identical timestamps of almost the exact same time when these messages occurred, making it much easier to read and correlate messages. Routers, switches, other networking devices, and pretty much every device known in the IT world has a time-of-day clock. For a variety of reasons, it makes sense to synchronize those clocks so that all devices have the same time of day, other than differences in time zone. The Network Time Protocol (NTP) provides the means to do just that. NTP gives any device a way to synchronize their time-of-day clocks. NTP provides protocol messages that devices use to learn the timestamp of other devices. Devices send timestamps to each other with NTP messages, continually exchanging messages, with one device changing its clock to match the other, eventually synchronizing the clocks. As a result, actions that benefit from synchronized timing, like the timestamps on log messages, work much better. This section works through a progression of topics that leads to the more common types of NTP configurations seen in real networks. The section begins with basic settings, like the timezone and initial configured time on a router or switch, followed by basic NTP configuration. The text then examines some NTP internals regarding how NTP defines the sources of time data (reference clocks) and how good each time source is (stratum). The section closes with more configuration that explains typical enterprise configurations, with multiple ntp commands for redundancy and the use of loopback interfaces for high availability. Setting the Time and Timezone NTP's job is to synchronize clocks, but NTP works best if you set the device clock to a reasonably close time before enabling the NTP client function with the ntp server command. For instance, my wristwatch says 8:52 p.m. right now. Before starting NTP on a new router or switch so that it synchronizes with another device, I should set the time to 8:52 p.m., set the correct date and timezone, and event tell the device to adjust for daylight savings time—and then enable NTP. Setting the time correctly gives NTP a good start toward synchronizing. Example 9-7 shows how to set the date, time, timezone, and daylight savings time. Oddly, it uses two configuration commands (for the timezone and daylight savings time) and one EXEC command to set the date and time on the router. Example 9-7 Setting the Date/Time with clock set, Plus Timezone/DST R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# clock summer-time EDT recurring R1(config)# Z R1# R1# clock set 20:52:49 21 October 2015 *Oct 21 20:52:49.000: %SYS-6-CLOCKUPDATE: System clock has been updated from 00:36:38 UTC Thu Oct 22 2015 to 20:52:49 UTC Wed Oct 21 2015, configured from console by Chapter 9: Device Management Protocols 183 console. R1# show clock 20:52:55.051 EDT Wed Oct 21 2015 Focus on the two configuration commands first. You should set the first two commands before setting the time of day with the clock set EXEC command because the two configuration commands impact the time that is set. In the first command, the clock timezone part defines the command and a keyword. The next parameter, "EST" in this case, is any value you choose, but choose the name of the timezone of the device. This value shows up in show commands, so although you make up the value, the value needs to be meaningful to all. I chose EST, the acronym for US Eastern Standard Time. The "-5" parameter means that this device is 5 hours behind Universal Time Coordinated (UTC). The clock summer-time part of the second command defines what to do, again with the "EDT" being a field in which you could have used any value. However, you should use a meaningful value. This is the value shown with the time in show commands when daylight savings time is in effect, so I chose EDT because it is the acronym for daylight savings time in that same EST time zone. Finally, the recurring keyword tells the router to spring forward an hour and fall back an hour automatically over the years. The clock set EXEC command then sets the time, day of the month, month, and year. However, note that IOS interprets the time as typed in the command in the context of the time zone and daylight savings time. In the example, the clock set command lists a time of 20:52:49 (the command uses a time syntax with a 24-hour format, not with a 12-hour format plus a.m./p.m.). As a result of that time plus the two earlier configuration commands, the show clock command (issued seconds later) lists that time, but also notes the time as EDT, rather than UTC time. Basic NTP Configuration With NTP, servers supply information about the time of day to clients, and clients react by adjusting their clocks to match. The process requires repeated small adjustments over time to maintain that synchronization. The configuration itself can be simple, or it can be extensive once you add security configuration and redundancy. Cisco supplies two ntp configuration commands that dictate how NTP works on a router or switch, as follows: ■ ntp master [stratum-level]: NTP server mode—the device acts only as an NTP server, and not as an NTP client. The device gets its time information from the internal clock on the device. ■ ntp server [address] [hostname]: NTP client/server mode—the device acts as both client and server. First, it acts as an NTP client, to synchronize time with a server. Once synchronized, the device can then act as an NTP server, to supply time to other NTP clients. For an example showing the basic configuration syntax and show commands, consider Figure 9-5. With this simple configuration: ■ R3 acts as an NTP server only. ■ R2 acts in client/server mode—first as an NTP client to synchronize time with NTP server R3, then as a server to supply time to NTP client R1. 9 184 CCNA 200-301 Official Cert Guide, Volume 2 ■ R1 acts in client/server mode—first as an NTP client to synchronize time with NTP server R2. (R1 will be willing to act as a server, but no devices happen to reference R1 as an NTP server in this example.) ntp server 172.16.2.2 G0/0/1 R1 NTP Client / Server Stratum 4 Figure 9-5 ntp server 172.16.3.3 G0/1 R2 ntp master 172.16.3.3 G0/1 NTP Client / Server Stratum 3 R3 G0/2 NTP Server Stratum 2 R1 as NTP Client, R2 as Client/Server, R3 as Server As you can see, NTP requires little configuration to make it work with a single configuration command on each device. Example 9-8 collects the configuration from the devices shown in the figure for easy reference. Example 9-8 NTP Client/Server Configuration | Configuration on R1: ntp server 172.16.2.1 Configuration on R2: ntp server 172.16.3.1 Configuration on R3: ntp master 2 Example 9-9 lists the output from the show ntp status command on R1, with the first line of output including a few important status items. First, it lists a status of synchronized, which confirms the NTP client has completed the process of changing its time to match the server's time. Any router acting as an NTP client will list "unsynchronized" in that first line until the NTP synchronization process completes with at least one server. It also confirms the IP address of the server—this device's reference clock—with the IP address configured in Example 9-8 (172.16.2.2). Example 9-9 Verifying NTP Client Status on R1 R1# show ntp status Clock is synchronized, stratum 4, reference is 172.16.2.2 nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**21 ntp uptime is 1553800 (1/100 of seconds), resolution is 4000 reference time is DAE57147.56CADEA7 (19:54:31.339 EST Thu Feb 4 2016) clock offset is 0.0986 msec, root delay is 2.46 msec root dispersion is 22.19 msec, peer dispersion is 5.33 msec loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000009 s/s system poll interval is 64, last update was 530 sec ago. Next, look at the show ntp associations command output from both R1 and R2 as shown in Example 9-10. This command lists all the NTP servers that the local device can attempt to use, with status information about the association between the local device (client) and Chapter 9: Device Management Protocols 185 the various NTP servers. Beginning with R1, note that it has one association (that is, relationship with an NTP server), based on the one ntp server 172.16.2.2 configuration command on R1. The * means that R1 has successfully contacted the server. You will see similar data from the same command output taken from router R2. Example 9-10 Verifying NTP Client Status on R1 and R2 R1# show ntp associations | This output is taken from router R1, acting in client/server mode address ref clock *~172.16.2.2 172.16.3.3 st when poll reach delay offset disp 3 0000 0.0000 0.232 s sys.peer, * selected, + candidate, - outlyer, x fselecticker, - configured Redundant NTP Configuration Instead of using a networking device as the reference clock for the enterprise, you can instead reference better time sources in the Internet or purchase a purpose-built NTP server that has better clocking hardware. For instance, an enterprise could use NTP to reference NTP servers that use an atomic clock as their reference source, which happens to be run by the US National Institute of Standards and Technology (NIST) (see f.t.nist.gov). Chapter 9: Device Management Protocols 187 S1 S2 NTP Primary Servers (NIST) Stratum 1 Internet Stratum 2 Stratum 3 Figure 9-6 NTP Client/Server R1 ... R101 R102 R198 R199 Stratum Levels When Using an Internet-based Stratum 1 NTP Server NOTE While the common terms NTP server mode and NTP client/server mode are useful, the NTP RFCs (1305 and 5905) also use two other specific terms for similar ideas: NTP primary server and NTP secondary server. An NTP primary server acts only as a server, with a reference clock external to the device, and has a stratum level of 1, like the two NTP primary servers shown in Figure 9-6. NTP secondary servers are servers that use client/server mode as described throughout this section, relying on synchronization with some other NTP server. For good design, the enterprise NTP configuration ought to refer to at least two external NTP servers for redundancy. Additionally, just a few enterprise devices should refer to those external NTP servers and then act as both NTP client and server. The majority of the devices in the enterprise, like those shown at the bottom of the figure, would act as NTP clients. Example 9-12 shows the configuration on router R1 and R2 in the figure to accomplish this design. Example 9-12 NTP Configuration on R1, R2 per Figure 9-6 ntp server time-a-b.nist.gov ntp server time-a-g.nist.gov In addition to referencing redundant NTP primary servers, some routers in the enterprise need to be ready to supply clock data if those NTP primary servers become unreachable. An exposure exists with the configuration in Example 9-12 because if router R1 and R2 no longer hear NTP messages from the NTP servers in the Internet they will lose their only reference clock. After losing their reference clock, R1 and R2 could no longer be useful NTP servers to the rest of the enterprise. 9 188 CCNA 200-301 Official Cert Guide, Volume 2 NOTE NTP considers 15 to be the highest useful stratum level, so any devices that calculate their stratum as 16 consider the time data unusable and do not trust the time. So, avoid setting higher stratum values on the ntp master command. To see the evidence, refer back to Example 9-10, which shows two commands based on the same configuration in Example 9-8 and Figure 9-5. The output highlights details about reference clocks a stratum level, as follows: R1: Per the configured ntp server 172.16.2.2 command, the show command lists the same address (which is router R2's address). The ref clock (reference clock) and st (stratum) fields present R2's reference clock as 172.16.3.3—in other words, R2's NTP server, which is R3 in this case. The st field value of 3 shows R2's stratum. R2: Per the configured ntp server 172.16.3.3 command, the show command lists the same address (which is an address on router R3). The output notes R3's ref clock as 127.127.1.1—an indication that the server (R3) gets its clock internally. It lists R3's st (stratum) value of 2—consistent with the configured ntp master 2 command on R3 (per Example 9-8). On the NTP primary server itself (R3 in this case), the output has more markers indicating the use of the internal clock. Example 9-11 shows output from R3, with a reference clock of the 127.127.1.1 loopback address, used to refer to the fact that this router gets its clock data internally. Also, in the show ntp associations command output at the bottom, note that same address, along with a reference clock value of ".LOCAL." in effect, R3, per the ntp master configuration command, has an association with its internal clock. Example 9-11 Examining NTP Server, Reference Clock, and Stratum Data R3# show ntp status Clock is synchronized, stratum 2, reference is 127.127.1.1 nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**20 ntp uptime is 595300 (1/100 of seconds), resolution is 4000 reference time is EOF9174C.87277EBB (16:13:32.527 daylight Sat Aug 10 2019) clock offset is 0.0000 msec, root delay is 0.00 msec root dispersion is 0.33 msec, peer dispersion is 0.23 msec loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s system poll interval is 16, last update was 8 sec ago. R3# show ntp associations address *~127.127.1.1 ref clock .LOCAL. st when 1 15 poll reach 16 377 delay offset disp 0.000 0.000 0.232 s sys.peer, * selected, + candidate, - outlyer, x fselecticker, - configured Redundant NTP Configuration Instead of using a networking device as the reference clock for the enterprise, you can instead reference better time sources in the Internet or purchase a purpose-built NTP server that has better clocking hardware. For instance, an enterprise could use NTP to reference NTP servers that use an atomic clock as their reference source, which happens to be run by the US National Institute of Standards and Technology (NIST) (see f.t.nist.gov). Chapter 9: Device Management Protocols 187 S1 S2 NTP Primary Servers (NIST) Stratum 1 Internet Stratum 2 Stratum 3 Figure 9-6 NTP Client/Server R1 ... R101 R102 R198 R199 Stratum Levels When Using an Internet-based Stratum 1 NTP Server NOTE While the common terms NTP server mode and NTP client/server mode are useful, the NTP RFCs (1305 and 5905) also use two other specific terms for similar ideas: NTP primary server and NTP secondary server. An NTP primary server acts only as a server, with a reference clock external to the device, and has a stratum level of 1, like the two NTP primary servers shown in Figure 9-6. NTP secondary servers are servers that use client/server mode as described throughout this section, relying on synchronization with some other NTP server. For good design, the enterprise NTP configuration ought to refer to at least two external NTP servers for redundancy. Additionally, just a few enterprise devices should refer to those external NTP servers and then act as both NTP client and server. The majority of the devices in the enterprise, like those shown at the bottom of the figure, would act as NTP clients. Example 9-12 shows the configuration on router R1 and R2 in the figure to accomplish this design. Example 9-12 NTP Configuration on R1, R2 per Figure 9-6 ntp server time-a-b.nist.gov ntp server time-a-g.nist.gov In addition to referencing redundant NTP primary servers, some routers in the enterprise need to be ready to supply clock data if those NTP primary servers become unreachable. An exposure exists with the configuration in Example 9-12 because if router R1 and R2 no longer hear NTP messages from the NTP servers in the Internet they will lose their only reference clock. After losing their reference clock, R1 and R2 could no longer be useful NTP servers to the rest of the enterprise. 9 188 CCNA 200-301 Official Cert Guide, Volume 2 Analyzing Topology Using CDP and LLDP The first two major sections of this chapter showed two features—syslog and NTP—that work the same way on both routers and switches. This final section shows yet another feature common to both routers and switches, with two similar protocols: the Cisco Discovery Protocol (CDP) and the Link Layer Discovery Protocol (LLDP). This section focuses on CDP, followed by LLDP. Examining Information Learned by CDP CDP discovers basic information about neighboring routers and switches without needing to know the passwords for the neighboring devices. To discover information, routers and switches send CDP messages out each of their interfaces. The messages essentially announce information about the device that sent the CDP message. Devices that support CDP learn information about others by listening for the advertisements sent by other devices. CDP discovers several useful details from the neighboring Cisco devices: ■ Device identifier: Typically the host name ■ Address list: Network and data-link addresses ■ Port identifier: The interface on the remote router or switch on the other end of the link that sent the CDP advertisement ■ Capabilities list: Information on what type of device it is (for example, a router or a switch) ■ Platform: The model and OS level running on the device CDP plays two general roles: to provide information to the devices to support some function and to provide information to the network engineers that manage the devices. For example, Cisco IP Phones use CDP to learn the data and voice VLAN IDs as configured on the access switch. For that second role, CDP has show commands that list information about neighboring devices, as well as information about how CDP is working. Table 9-3 describes the three show commands that list the most important CDP information. Table 9-3 show cdp Commands That List Information About Neighbors Command Description show cdp neighbors [type number] Lists one summary line of information about each neighbor or just the neighbor found on a specific interface if an interface was listed show cdp neighbors detail Lists one large set (approximately 15 lines) of information, one set for every neighbor show cdp entry name Lists the same information as the show cdp neighbors detail command, but only for the named neighbor (case sensitive) NOTE Cisco routers and switches support the same CDP commands, with the same parameters and same types of output. The next example shows the power of the information in CDP commands. The example uses the network shown in Figure 9-8, with Example 9-15 listing the output of several show cdp commands. Chapter 9: Device Management Protocols 191 Cisco C2960X Switches (WS-2960XR-24TS-I) Gi1/0/24 Gi1/0/21 SW2 SW1 Gi1/0/1 Fret 0200.11111111 Gi1/0/2 Gi1/0/2 Barney 0200.2222.2222 Gi0/0/1 R1 Figure 9-8 0200.5555.5555 Cisco ISR1K Router Small Network Used in CDP Examples Example 9-15 show cdp neighbors Command Examples: SW2 SW2# show cdp neighbors Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge S - Switch, H - Host, I - IGMP, P - Repeater, P - Phone, D - Remote, C - CVTA, M - Two-port Mac Relay Device ID Local Intrfce Holdtime SW1 Gig 1/0/21 155 Capability S I R1 Gig 1/0/2 131 R S I Platform Port ID WS-C2960X Gig 1/0/24 C1111-8P Gig 0/0/1 Total cdp entries displayed: 2 The show cdp neighbors command lists one line per neighbor. (Look for the Device ID column and the list that includes SW1 and R1.) Each of those two lines lists the most important topology information about each neighbor: the neighbor's host name (Device ID), the local device's interface, and the neighboring device's interface (under the Port heading). Pay close attention to the local device's interface and the neighboring device's interface, comparing the example to the figure. For example, SW2's show cdp neighbors command lists an entry for SW1, with SW2's local interface of Gi0/2 and SW1's interface of Gi0/1 under the heading "Port ID." This command also lists the platform, identifying the specific model of the neighboring router or switch. So, even using this basic information, you could either construct a figure like Figure 9-8 or confirm that the details in the figure are correct. Figure 9-8 and Example 9-15 provide a good backdrop as to why devices learn about direct neighbors with CDP, but not other neighbors. First, CDP defines encapsulation that uses the data-link header, but no IP header. To ensure all devices receive a CDP message, the Ethernet header uses a multicast destination MAC address (0100.0CCC.CCCC). However, when any device that supports CDP receives a CP message, the device processes the message and then discards it, rather than forwarding it. So, for instance, when router R1 sends a CDP message to Ethernet multicast address 0100.0CCC.CCCC, switch SW2 receives it, processes it, but does not forward it to switch SW1—so SW1 will not list router R1 as a CDP neighbor. Next, consider the show cdp neighbors detail command as shown in Example 9-16, again taken from switch SW2. This command lists more detail, as you might have guessed. The 9 192 CCNA 200-301 Official Cert Guide, Volume 2 detail lists the full name of the switch model (WS-2960XR-24TS-I) and the IP address configured on the neighboring device. You have to look closely, but the example has one long group of messages for each of the two neighbors; the example includes one comment line with gray highlight to help you find the dividing point between groups of messages. Example 9-16 show cdp neighbors detail Command on SW2 SW2# show cdp neighbors detail -----Device ID: SW1 Entry address(es): IP address: 1.1.1.1 Platform: cisco WS-C2960XR-24TS-I, Interface: GigabitEthernet1/0/21, Capabilities: Switch IGMP Port ID (outgoing port): GigabitEthernet1/0/24 Holdtime : 144 sec Version : Cisco IOS Software, C2960X Software (C2960X-K9S-M), Version 15.2(6)E2, RELEASE SOFTWARE (fc4) Technical Support: Copyright (c) 1986-2018 by Cisco Systems, Inc. Compiled Thu 13-Sep-18 03:43 by prod_rel_team advertisement version: 2 Protocol Hello: OUI=0x000000, Protocol ID=0x0112, payload len=27, value=00000000FFFF FFF01025010000000000BCC4938BA180FF0000 RELEASE Management Domain: 'fred' Native VLAN: 1 Duplex: full Management address(es): IP address: 1.1.1.1 -----Device ID: R1 Entry address(es): IP address: 10.12.25.5 Platform: cisco C1111-8P, Capabilities: Router Switch IGMP Interface: GigabitEthernet1/0/2, Port ID (outgoing port): GigabitEthernet0/0/1 Holdtime : 151 sec Version : Cisco IOS Software [Fuji], ISR Software (ARMV8EB_LINUX_IOSD-UNIVERSALK9_IAS-M), Version 16.8.1, RELEASE SOFTWARE (fc3) Technical Support: Copyright (c) 1986-2018 by Cisco Systems, Inc. Compiled Tue 27-Mar-18 10:56 by mcpre advertisement version: 2 Chapter 9: Device Management Protocols 193 VTP Management Domain: " Duplex: full Management address(es): IP address: 10.12.25.5 Total cdp entries displayed : 2 NOTE The show cdp entry name command lists the exact same details shown in the output of the show cdp neighbors detail command, but for only the one neighbor listed in the command. As you can see, you can sit on one device and discover a lot of information about a neighboring device—a fact that actually creates a security exposure. Cisco recommends that CDP be disabled on any interface that might not have a need for CDP. For switches, any switch port connected to another switch, a router, or to an IP phone should use CDP. Finally, note that CDP shows information about directly connected neighbors. For instance, show cdp neighbors on SW1 would list an entry for SW2 in this case, but not R1, because R1 is not directly connected to SW1. Configuring and Verifying CDP Most of the work you do with CDP relates to what CDP can tell you with show commands. However, it is an IOS feature, so you can configure CDP and use some show commands to examine the status of CDP itself. IOS typically enables CDP globally and on each interface by default. You can then disable CDP per interface with the no cdp enable interface subcommand and later re-enable it with the cdp enable interface subcommand. To disable and re-enable CDP globally on the device, use the no cdp run and cdp run global commands, respectively. To examine the status of CDP itself, use the commands in Table 9-4. Table 9-4 Commands Used to Verify CDP Operations Command Description show cdp States whether CDP is enabled globally and lists the default update and holdtime timers show cdp interface [type number] States whether CDP is enabled on each interface if the interface is listed, and states update and holdtime timers on those interfaces show cdp traffic Lists global statistics for the number of CDP advertisements sent and received Example 9-17 lists sample output from each of the commands in Table 9-4, based on switch SW2 in Figure 9-8. 9 194 CCNA 200-301 Official Cert Guide, Volume 2 Example 9-17 show cdp Commands That Show CDP Status SW2# show cdp Global CDP information: Sending CDP packets every 60 seconds Sending a holdtime value of 180 seconds Sending CDPv2 advertisements is enabled SW2# show cdp interface GigabitEthernet1/0/2 GigabitEthernet1/0/2 is up, line protocol is up Encapsulation ARPA Sending CDP packets every 60 seconds Holdtime is 180 seconds SW2# show cdp traffic CDP counters : Total packets output: 304, Input: 305 Hdr syntax: 0, Chksm error: 0, Encaps failed: 0, No memory: 0, Invalid packet: 0, CDP version 1 advertisements output: 0, Input: 0 CDP version 2 advertisements output: 304, Input: 305 The first two commands in the example list two related settings about how CDP works: the send time and the hold time. CDP sends messages every 60 seconds by default, with a hold time of 180 seconds. The hold time tells the device how long to wait after no longer hearing from a device before removing those details from the CDP tables. You can override the defaults with the cdp timer seconds and cdp holdtime seconds global commands, respectively. Examining Information Learned by LLDP Cisco created the Cisco-proprietary CDP before any standard existed for a similar protocol. CDP has many benefits. As a Layer 2 protocol, sitting on top of Ethernet, it does not rely on a working Layer 3 protocol. It provides device information that can be useful in a variety of ways. Cisco had a need but did not see a standard that met the need, so Cisco made up a protocol, as has been the case many times over history with many companies and protocols. Link Layer Discovery Protocol (LLDP), defined in IEEE standard 802.1AB, provides a standardized protocol that provides the same general features as CDP. LLDP has similar configuration and practically identical show commands as compared with CDP. The LLDP examples all use the same topology used in the CDP examples per Figure 9-8 (the same figure used in the CDP examples), Example 9-18 lists switch SW2's LLDP neighbors as learned after LLDP was enabled on all devices and ports in that figure. The example highlights the items that match the similar output from the show cdp neighbors command listed at the end of the example, also from switch SW2. Chapter 9: Device Management Protocols 195 Example 9-18 show cdp neighbors on SW2 with Similarities to CDP Highlighted SW2# show lldp neighbors Capability codes: (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device, (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other Device ID Local Intf Hold-time Capability Port ID R1 Gi1/0/2 120 R Gi1/0/24 Total entries displayed: 2 SW2# show cdp neighbors Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge S - Switch, H - Host, I - IGMP, P - Repeater, P - Phone, D - Remote, C - CVTA, M - Two-port Mac Relay Device ID Local Intrfce Holdtime SW1 Gig 1/0/21 155 Capability S I R1 Gig 1/0/2 131 R S I Platform Port ID WS-C2960X Gig 1/0/24 C1111-8P Gig 0/0/1 Total entries displayed: 2 The most important take-away from the output is the consistency between CDP and LLDP in how they refer to the interfaces. Both the show cdp neighbors and show lldp neighbors commands have "local intf" (interface) and "port ID" columns. These columns refer to the local device's interface and the neighboring device's interface, respectively. However, the LLDP output in the example does differ from CDP in a few important ways: ■ LLDP uses B as the capability code for switching, referring to bridge, a term for the device type that existed before switches that performed the same basic functions. ■ LLDP does not identify IGMP as a capability, while CDP does (I). ■ CDP lists the neighbor's platform, a code that defines the device type, while LLDP does not. ■ LLDP lists capabilities with different conventions (see upcoming Example 9-19). The first three items in the list are relatively straightforward, but that last item in the list requires a closer look with more detail. Interestingly, CDP lists all the capabilities of the neighbor in the show cdp neighbors command output, no matter whether the device currently enables all those features. LLDP instead lists the enables (configured) capabilities, rather than all supported capabilities, in the output from show lldp neighbors command. LLDP makes the difference in a neighbor's total capabilities and configured capabilities with the show lldp neighbors detail and show lldp entry hostname commands. These commands provide identical detailed output, with the first command providing detail for all neighbors, and the second providing detail for the single listed neighbor. Example 9-19 shows the detail for neighbor R1. 9 196 CCNA 200-301 Official Cert Guide, Volume 2 Example 9-19 show lldp entry R2 Command on SW2 SW2# show lldp entry R1 Capability codes: (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device, (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other -----Local Intf: Gi1/0/2 Chassis id: 70ea.1a9a.d300 Port id: Gi0/0/1 Port Description: GigabitEthernet0/1 System Name: R1 System Description: Cisco IOS Software [Fuji], ISR Software (ARMV8EB_LINUX_IOSD-UNIVERSALK9_IAS-M), Version 16.8.1, RELEASE SOFTWARE (fc3) Technical Support: Copyright (c) 1986-2018 by Cisco Systems, Inc. Compiled Tue 27-Mar-18 10:56 by mcpre Time remaining: 100 seconds System Capabilities: B,R Enabled Capabilities: R Management Addresses: IP: 10.12.25.5 Auto Negotiation - not supported Physical media capabilities - not advertised Media Attachment Unit type - not advertised Vlan ID - not advertised Total entries displayed: 1 First, regarding the device capabilities, note that the LLDP command output lists two lines about the neighbor's capabilities: System Capabilities: What the device can do Enabled Capabilities: What the device does now with its current configuration For instance, in Example 9-19, the neighboring R1 claims the ability to perform routing and switching (codes R and B) but also claims to currently be using only its routing capability, as noted in the "enabled capabilities" line. Also, take a moment to look at the output for the similarities to CDP. For instance, this output lists detail for neighbor, R1, which uses its local port G0/0/1, with a host name of R1. The output also notes the IOS name and version, from which an experienced person can infer the model number, but there is no explicit mention of the model. Chapter 9: Device Management Protocols 197 NOTE LLDP uses the same messaging concepts as CDP, encapsulating messages directly in data-link headers. Devices do not forward LLDP messages so that LLDP learns only of directly connected neighbors. LLDP does use a different multicast MAC address (0180. C200.000E). Configuring and Verifying LLDP LLDP uses a similar configuration model as CDP, but with a few key differences. First, Cisco devices default to disable LLDP. Additionally, LLDP separates the sending and receiving of LLDP messages as separate functions. For instance, LLDP support processing receives LLDP messages on an interface so that the switch or router learns about the neighboring device while not transmitting LLDP messages to the neighboring device. To support that model, the commands include options to toggle on/off the transmission of LLDP messages separately from the processing of received messages. The three LLDP configuration commands are as follows: ■ [no] lldp run: A global configuration command that sets the default mode of LLDP operation for any interface that does not have more specific LLDP subcommands (lldp transmit, lldp receive). The lldp run global command enables LLDP in both directions on those interfaces, while no lldp run disables LLDP. ■ [no] lldp transmit: An interface subcommand that defines the operation of LLDP on the interface regardless of the global [no] lldp run command. The lldp transmit interface subcommand causes the device to transmit LLDP messages, while no lldp transmit causes it to not transmit LLDP messages. ■ [no] lldp receive: An interface subcommand that defines the operation of LLDP on the interface regardless of the global [no] lldp run command. The lldp receive interface subcommand causes the device to process received LLDP messages, while no lldp receive causes it to not process received LLDP messages. For example, consider a switch that has no LLDP configuration commands at all. Example 9-20 adds a configuration that first enables LLDP for all interfaces (in both directions) with the lldp run global command. It then shows how to disable LLDP in both directions on Gi1/0/17 and how to disable LLDP in one direction on Gi1/0/18. Example 9-20 Enabling LLDP on All Ports, Disabling on Few Ports lldp run | interface gigabitEthernet1/0/17 no lldp transmit no lldp receive | interface gigabitEthernet1/0/18 no lldp receive Example 9-21 adds another example that again begins with a switch with all default settings. In this case, the configuration does not enable LLDP on all interfaces with the lldp run command, meaning that all interfaces default to not transmit and not receive LLDP 9 198 CCNA 200-301 Official Cert Guide, Volume 2 messages. The example does show how to then enable LLDP for both directions on one interface and in one direction for a second interface. Example 9-21 Enabling LLDP on Limited Ports, Leaving Disabled on Most interface gigabitEthernet1/0/19 lldp transmit lldp receive | interface gigabitEthernet1/0/20 lldp receive Finally, checking LLDP status uses the exact same commands as CDP as listed in Table 9-4, other than the fact that you use the lldp keyword instead of cdp. For instance, show lldp interface lists the interfaces on which LLDP is enabled. Example 9-22 shows some examples from switch SW2 based on earlier Figure 9-8 (the same figure used in the CDP examples), with LLDP enabled in both directions on all interfaces with the cdp run global command. Example 9-22 show lldp Commands That Show LLDP Information: Status: ACTIVE LLDP advertisements are sent every 30 seconds LLDP hold time advertised is 120 seconds LLDP interface reinitialisation delay is 2 seconds SW2# show lldp interface g1/0/2 GigabitEthernet1/0/2: Tx: enabled Rx: enabled Tx state: IDLE Rx state: WAIT FOR FRAME SW2# show lldp traffic LLDP traffic statistics: Total frames out: 259 Total entries aged: 0 Total frames in: 257 Total frames received in error: 0 Total frames discarded: 0 Total TLVs discarded: 0 Total TLVs unrecognized: 0 Also, note that like CDP, LLDP uses a send timer and hold timer for the same purposes as CDP. The example shows the default settings of 30 seconds for the send timer and 120 seconds for the hold timer. You can override the defaults with the lldp timer seconds and lldp holdtime seconds global commands, respectively. Chapter 9: Device Management Protocols 199 Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 9-5 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 9-5 Chapter Review Tracking Review Element Review Date(s) Resource Used Review key topics Book, website Review key terms Book, website Answer DIKTA questions Book, PTP Review material tables Book, app Do labs Blog Review command references Book Review All the Key Topics Table 9-6 Key Topics for Chapter 9 Key Topic Element Description Page Number Figure 9-1 Logging to console and terminal 175 Figure 9-2 Logging to syslog and buffer 176 Figure 9-3 Log message levels 177 Table 9-2 Logging configuration commands 177 List The ntp master and ntp server commands 183 List Sequence for NTP client to choose a reference clock 188 List Key facts about loopback interfaces 189 List Information gathered by CDP 190 Table 9-3 Three CDP show commands that list information about neighbors 190 List Differences between LLDP and CDP 195 List LLDP configuration commands and logic 197 Key Terms You Should Know log message, syslog server, Network Time Protocol (NTP), NTP client, NTP client/server mode, NTP server, NTP synchronization, CDP, LLDP Command References Tables 9-9 and 9-8 Is-lls configuration and verification commands used in this chapter. As an easy review exercise, cover the left column in a table, read the right column, and try to recall the command without looking. Then repeat the exercise, covering the right column, and try to recall what the command does. 9 200 CCNA 200-301 Official Cert Guide, Volume 2 Table 9-7 Configuration Command Reference Command Description [no] logging console Global command that enables (or disables with the no option) logging to the console device. [no] logging monitor Global command that enables (or disables with the no option) logging to an internal buffer. logging [host] ip-address [hostname] Global command that enables logging to a syslog server. logging console level-name [level-number] Global command that sets the log message level for console log messages. logging monitor level-name [level-number] Global command that sets the log message level for log messages sent to SSH and Telnet users. logging buffered level-name [level-number] Global command that sets the log message level for buffered log messages displayed later by the show logging command. logging trap level-name [level-number] Global command that sets the log message level for messages sent to syslog servers. number [no] service sequence-numbers Global command to enable or disable (with the no option) the use of sequence numbers in log messages. clock timezone name +-number Global command that names a timezone and defines the +/- offset versus UTC. clock summertime name recurring Global command that names a daylight savings time for a timezone and tells IOS to adjust the clock automatically. ntp server address [hostname] Global command that configures the device as an NTP client by referring to the address or name of an NTP server. ntp master stratum-level Global command that configures the device as an NTP server and assigns its local clock stratum level. ntp source name/number Global command that tells NTP to use the listed interface (by name/number) for the source IP address for NTP messages. interface loopback number Global command that, at first use, creates a loopback interface. At all uses, it also moves the user into interface configuration mode for that interface. [no] cdp run Global command that enables and disables (with the no option) CDP for the entire switch or router. [no] cdp enable Interface subcommand to enable and disable (with the no option) CDP for a particular interface. cdp timer seconds Global command that changes the CDP send timer (the frequency at which CDP sends messages). cdp holdtime seconds Global command that changes how long CDP waits since the last received message from a neighbor before believing the neighbor has failed, removing the neighbor's information from the LLDP table. Chapter 9: Device Management Protocols 201 Command Description [no] lldp run Global command to enable and disable (with the no option) LLDP for the entire switch or router. [no] lldp transmit Interface subcommand to enable and disable (with the no option) the transmission of LLDP messages on the interface. [no] lldp receive Interface subcommand to enable and disable (with the no option) the processing of received LLDP messages on the interface. lldp timer seconds Global command that changes the LLDP send timer (the frequency at which LLDP sends messages). lldp holdtime seconds Global command that changes how long LLDP waits since the last received message from a neighbor before believing the neighbor has failed, removing the neighbor's information from the LLDP table. Table 9-8 Chapter 9 EXEC Command Reference Command Description show logging Lists the current logging configuration and lists buffered log messages at the end terminal monitor For a user (SSH or Telnet) session, toggles on (terminal monitor) or off (terminal no monitor) the receipt of log messages, for that one session, if logging monitor is also configured terminal no monitor [no] debug (various) EXEC command to enable or disable (with the no option) one of a multitude of debug options show clock Lists the time of-day and the date per the local device show ntp associations Shows all NTP clients and servers with which the local device is attempting to synchronize with NTP show ntp status Shows current NTP client status in detail show interfaces loopback number Shows the current status of the listed loopback interface show cdp | lldp neighbors [type number] Lists one summary line of information about each neighbor; optionally, lists neighbors off the listed interface show cdp | lldp neighbors detail Lists one large set of information (approximately 15 lines) for every neighbor show cdp | lldp entry name Displays the same information as show cdp|lldp neighbors detail but only for the named neighbor show cdp | lldp States whether CDP or LLDP is enabled globally and lists the default update and holdtime timers show cdp | lldp interface [type number] States whether CDP or LDP is enabled on each interface or a single interface if the interface is listed show cdp | lldp traffic Displays global statistics for the number of CDP or LLDP advertisements sent and received 9 CHAPTER 10 Network Address Translation This chapter covers the following exact topics: 4.0 IP Services 4.1 Configure and verify inside source NAT using static and pools This chapter examines a very popular and very important part of both enterprise and small office/home office (SOHO) networks: Network Address Translation, or NAT. NAT helped solve a big problem with IPv4: the IPv4 address space would have been completely consumed by the mid-1990s. After it was consumed, the Internet could not continue to grow, which would have significantly slowed the development of the Internet. This chapter breaks the topics into three major sections. The first section explains the challenges to the IPv4 address space caused by the Internet revolution of the 1990s. The second section explains the basic concept behind NAT, how several variations of NAT work, and how the Port Address Translation (PAT) option conserves the IPv4 address space. The final section shows how to configure NAT from the Cisco IOS Software command-line interface (CLI) and how to troubleshoot NAT. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 10-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions Perspectives on IPv4 Address Scalability 1–2 Network Address Translation Concepts 3–4 NAT Configuration and Troubleshooting 5–7 1. Which of the following summarized subsets represent routes that could have been created for CIDR's goal to reduce the size of Internet routing tables? a. 10.0.0.0 255.255.255.0 b. 10.1.0.0 255.255.0.0 c. 200.1.1.0 255.255.255.0 d. 200.1.0.0 255.255.0.0 2. Which of the following are not private addresses according to RFC 1918? (Choose two answers.) a. 172.31.1.1 b. 172.33.1.1 c. 10.255.1.1 d.

is a formal term to refer to actions other than storing and forwarding a message. These actions can delay the message, discard it, or even change header fields. The device can choose different PHBs for different kinds of messages, improving the QoS behavior for some messages, while worsening the QoS behavior for others. This chapter works through the QoS tools listed in the single QoS exam topic: "Explain the forwarding per-behavior (PHB) for QoS such as classification, marking, queuing, congestion, policing, shaping." Each topic emphasizes the problems each tool solves and how each tool manages bandwidth, delay, jitter, and loss. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 11-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions Introduction to QoS 1 Classification and Marking 2, 3 Queuing 4 Shaping and Policing 5 Congestion Avoidance 6 1. Which of the following attributes do QoS tools manage? (Choose three answers.) a. Bandwidth b. Delay c. Load d. MTU e. Loss 2. Which of the following QoS marking fields could remain with a packet while being sent through four different routers, over different LAN and WAN links? (Choose two answers.) a. CoS b. IPP c. DSCP d. MPLS EXP 3. Which of the following are available methods of classifying packets in DiffServ on Cisco routers? (Choose three answers.) a. Matching the IP DSCP field b. Matching the 802.1p CoS field c. Matching fields with an extended IP ACL d. Matching the SNMP Location variable 4. Which of the following behaviors are applied to a low latency queue in a Cisco router or switch? (Choose two answers.) a. Shaping b. Policing c. Priority scheduling d. Round-robin scheduling 5. Think about a policing function that is currently working, and also think about a shaping function that is also currently working. That is, the current bit rate of traffic exceeds the respective policing and shaping rates. Which statements are true about these features? (Choose two answers.) a. The policer may or may not be discarding packets. b. The policer is definitely discarding packets. c. The shaper may or may not be queuing packets to slow down the sending rate. d. The shaper is definitely queuing packets to slow down the sending rate. 228 CCNA 200-301 Official Cert Guide, Volume 2 6. A queuing system has three queues serviced with round-robin scheduling and one low latency queue that holds all voice traffic. Round-robin queue 1 holds predominantly UDP traffic, while round-robin queues 2 and 3 hold predominantly TCP traffic. The packets in each queue happen to have a variety of DSCP markings per the QoS design. In which queues would it make sense to use a congestion avoidance (drop management) tool? (Choose two answers.) a. The LLQ b. Queue 1 c. Queue 2 d. Queue 3 Foundation Topics Introduction to QoS Routers typically sit at the WAN edge, with both WAN interfaces and LAN interfaces. Those LAN interfaces typically run at much faster speeds, while the WAN interfaces run at slower speeds. While that slower WAN interface is busy sending the packets waiting in the router, hundreds or even thousands more IP packets could arrive in the LAN interfaces, all needing to be forwarded out that same WAN interface. What should the router do? Send them all, in the same order in which they arrived? Prioritize the packets, to send some earlier than others, preferring one type of traffic over another? Discard some of the packets when the number of packets waiting to exit the router gets too large? That first paragraph described some of the many classic Quality of Service (QoS) questions in networking. QoS refers to the tools that networking devices use to apply some different treatment to packets in the network as they pass through the device. For instance, the WAN edge router would queue packets waiting for the WAN interface to be available. The router could also use a queue scheduling algorithm to determine which packets should be sent next, using some other way from the arrival order—giving some packets better service and some worse service. QoS: Managing Bandwidth, Delay, Jitter, and Loss Cisco offers a wide range of QoS tools on both routers and switches. All these tools give you the means to manage four characteristics of network traffic: ■ Bandwidth ■ Delay ■ Jitter ■ Loss Bandwidth refers to the speed of a link, in bits per second (bps). But while we think of bandwidth as speed, it helps to also think of bandwidth as the capacity of the link, in terms of how many bits can be sent over the link per second. The networking device's QoS tools determine what packet is sent over the link next, so the networking device is in control of which messages get access to the bandwidth next and how much of that bandwidth (capacity) each type of traffic gets over time. Chapter 11: Quality of Service (QoS) 229 For example, consider that typical WAN edge router that has hundreds of packets waiting to exit the WAN link. An engineer might configure a queuing tool to reserve 10 percent of the bandwidth for voice traffic, 50 percent for mission-critical data applications, and leave the rest of the bandwidth for all other types of traffic. The queuing tool could then use those settings to make the choice about which packets to send next. Delay can be described as one-way delay or round-trip delay. One-way delay refers to the time between sending one packet and that same packet arriving at the destination host, and that same packet arriving at the destination host. Round-trip delay counts the one-way delay plus the time for the receiver of the first packet to send back a packet—in other words, the time it takes to send one packet between two hosts and receive one back. Many different individual actions impact delay; this chapter will discuss a few of those, including queuing and shaping delay. Jitter refers to the variation in one-way delay between consecutive packets sent by the same application. For example, imagine an application sends a few hundred packets to one particular host. The first packet's one-way delay is 300 milliseconds (300 ms, or .3 seconds). The next packet's one-way delay is 300 ms, so is the third's; and so on. In that case, there is no jitter. However, if instead the first packet has a one-way delay of 300 ms, the next has a one-way delay of 310 ms, and the next has 325 ms, then there is some variation in the delay; 10 ms between packets 1 and 2, and another 15 ms between packets 2 and 3. That difference is called jitter. Finally, loss refers to the number of lost messages, usually as a percentage of packets sent. The comparison is simple: if the sender for some application sends 100 packets, and only 98 arrive at the destination, that particular application flow experienced 2 percent loss. Loss can be caused by many factors, but often, people think of loss as something caused by faulty cabling or poor WAN services. That is one cause. However, more loss happens because of the normal operation of the networking devices, in which the devices' queues get too full, so the device has nowhere to put new packets, and it discards the packet. Several QoS tools manage queuing systems to help control and avoid loss. Types of Traffic With QoS, a network engineer sets about to prefer one type of traffic over another in regard to bandwidth, delay, jitter, and loss. Sometimes, that choice relates to the specific business. For example, if all the mission-critical applications sit on servers in three known subnets, then the QoS plan could be set up to match packets going to/from that subnet and give that traffic better treatment compared to other traffic. However, in other cases, the choice of how to apply QoS tools relates to the nature of different kinds of applications. Some applications have different QoS needs than others. This next topic compares the basic differences in QoS needs based on the type of traffic. Data Applications First, consider a basic web application, with a user at a PC or tablet. The user types in a URL to request a web page. That request may require a single packet going to the web server, but it may result in hundreds or thousands of packets coming back to the web client, as shown in Figure 11-1. 11 230 CCNA 200-301 Official Cert Guide, Volume 2 HTTP GET Web Server ... 500 Packets Figure 11-1 Concept of Disproportionate Packet/Byte Volumes with HTTP Traffic NOTE If you wonder how one web page might require thousands of packets, consider this math: with a 1500-byte IP maximum transmission unit (MTU), the data part of a TCP segment could be at most 1460 bytes (1500 bytes minus 20 bytes each for the IP and TCP header). In this example, 1000 such packets total to 1,460,000 bytes, or about 1.5 MB. It is easy to imagine a web page with just a few graphics that totals more than 1.5 MB in size. So, what is the impact of bandwidth, delay, jitter, and loss on an interactive web-based application? First, the packets require a certain amount of bandwidth capacity. As for delay, each of those packets from the server to the client takes some amount of one-way delay, with some jitter as well. Of the 500 packets shown in Figure 11-1, if some are lost (transmission errors, discarded by devices, or other reasons), then the server's TCP logic will retransmit, but parts of the web page may not show up right away. While QoS tools focus on managing bandwidth, delay, jitter, and loss, the user mainly cares about the quality of the overall experience. For instance, with a web application, how long after clicking do you see something useful in
your web browser? So, as a user, you care about the Quality of Experience (QoE), which is a term referring to users' perception of their use of the application on the network. QoS tools directly impact bandwidth, delay, jitter, and loss, which then should have some overall good effect to influence the users' QoE. And you can use QoS tools to create a better QoE for more important traffic; for instance, you might give certain business-critical applications better QoS treatment, which improves QoE for users of those apps. In contrast, a noninteractive data application (historically called batch traffic)—for instance, data backup or file transfers—has different QoS requirements than interactive data applications. Batch applications typically send more data than interactive applications, but because no one is sitting there waiting to see something pop on the screen, the delay and jitter do not matter much. Much more important for these applications is meeting the need to complete the larger task (transferring files) within a larger time window. QoS tools can be used to provide enough bandwidth to meet the capacity needs of these applications and manage loss to reduce the number of retransmissions. Voice and Video Applications Voice and video applications each have a similar breakdown of interactive and noninteractive flows. To make the main points about both voice and video, this section looks more deeply at voice traffic. Answers to the "Do I Know This Already?" quiz: 1. A, B, E 2. B, C 3. A, B, C 4. B, C 5. A, D 6. C, D Chapter 11: Quality of Service (QoS) 231 Before looking at voice, though, first think about the use of the term flow in networking. A flow is all the data moving from one application to another over the network, with one flow for each direction. For example, if you open a website and connect to a web server, the web page content that moves from the server to the client is one flow. Listen to some music with a music app on your phone, and that creates a flow from your app to the music app's server and a flow from the server back to your phone. From a voice perspective, a phone call between two IP phones would create a flow for each direction. For video, it could be the traffic from one video surveillance camera collected by security software. Now on to voice, specifically Voice over IP (VoIP). VoIP defines the means to take the sound made at one telephone and send it inside IP packets over an IP network, playing the sound back on the other telephone. Figure 11-2 shows the general idea. The steps in the figure include Step 1. The phone user makes a phone call and begins speaking. Step 2. A chip called a codec processes (digitizes) the sound to create a binary code (160 bytes with the G.711 codec, for example) for a certain time period (usually 20 ms). Step 3. The phone places the data into an IP packet. Step 4. The phone sends the packet to the destination IP phone. IP Phone Internals 2.1 CODEC Voice Bytes 3 IP Figure 11-2 UDP RTP Voice Bytes 4 Creating VoIP Packets with an IP Phone and a G.711 Codec If you work through the math a bit, this single call, with the G.711 codec, requires about 80 Kbps of bandwidth (ignoring the data-link header and trailer overhead). Counting the headers and VoIP payload as shown in the figure, each of the IP packets has 200 bytes. Each holds 20 ms of digitized voice, so the phone sends 50 packets per second. These 50 packets at 200 bytes each equal only 10,000 bytes per second, or 80,000 bits per second, which is 80 Kbps. Other voice codecs require even less bandwidth, with the commonly used G.729 taking about 24 Kbps (again ignoring data-link overhead). At first, it may look like VoIP calls require little in regard to QoS. For bandwidth, a single voice call or flow requires only a little bandwidth in comparison to many data applications. However, interactive voice does require a much better level of quality for delay, jitter, and loss. For instance, think about making a phone call with high one-way delay. You finish speaking and pause for the other person to respond. And he does not, so you speak again—and hear the other person's voice overlaid on your own. The problem: too much delay. Or, consider calls for which the sound breaks up. The problem? It could have been packet loss, or it could have been jitter. 11 232 CCNA 200-301 Official Cert Guide, Volume 2 You can achieve good-quality voice traffic over an IP network, but you must implement QoS to do so. QoS tools set about to give different types of traffic the QoS behavior they need. Cisco's Enterprise QoS Solution Reference Network Design Guide, which itself quotes other sources in addition to relying on Cisco's long experience in implementing QoS, suggests the following guidelines for interactive voice: ■ Delay (one-way): 150 ms or less ■ Jitter: 30 ms or less ■ Loss: 1% or less In comparison, interactive voice requires more attention than interactive data applications for QoS features. Data applications generally tolerate more delay, jitter, and loss than voice (and video). A single voice call does generally take less bandwidth than a typical data application, but that bandwidth requirement is consistent. Data applications tend to be bursty, with data bursts in reaction to the user doing something with the application. Video has a much more varied set of QoS requirements. Generally, think of video like voice, but with a much higher bandwidth requirement than voice (per flow) and similar requirements for low delay, jitter, and loss. As for bandwidth, video can use a variety of codecs that impact the amount of data sent, but many other technical features impact the amount of bandwidth required for a single video flow. (For instance, a sporting event with lots of movement on the screen would be 0.1% than a news anchor reading the news in front of a solid background with little movement.) This time quoting from End-to-End QoS Network Design, Second Edition (Cisco Press, 2013), some requirements for video include ■ Bandwidth: 384 Kbps to 20+ Mbps ■ Delay (one-way): 200–400 ms ■ Jitter: 30–50 ms ■ Loss: 0.1%–1% NOTE End-to-End QoS Network Design is written by some of the same people who created the Cisco Enterprise QoS Solution Reference Network Design Guide (available at Cisco.com). If you are looking for a book to dig into more depth on QoS, this book is an excellent reference for Cisco QoS. QoS as Mentioned in This Book QoS tools change the QoS characteristics of certain flows in the network. The rest of the chapter focuses on the specific tools mentioned in the lone CCNA 200-301 exam topic about QoS, presented in the following major sections: ■ "Classification and Marking" is about the marking of packets and the definition of trust boundaries. ■ "Queuing" describes the scheduling of packets to give one type of packet priority over another. ■ "Shaping and Policing" explains these two tools together because they are often used on opposite ends of a link. ■ "Congestion Avoidance" addresses how to manage the packet loss that occurs when network devices get too busy. Chapter 11: Quality of Service (QoS) 233 QoS on Switches and Routers Before moving on to several sections of the chapter about specific QoS tools, let me make a point about the terms packet and frame as used in this chapter. The QoS tools discussed in this chapter can be used on both switches and routers. There are some differences in the features and differences in implementation, due to the differences of internal architecture between routers and switches. However, to the depth discussed here, the descriptions apply equally to both LAN switches and IP routers. This chapter uses the word packet in a general way, to refer to any message being processed by a networking device, just for convenience. Normally, the term packet refers to the IP header and encapsulated headers and data, but without the data-link header and trailer. The term frame refers to the data-link header/trailer with its encapsulated headers and data. For this chapter, those differences do not matter to the discussion, but at the same time, the discussion often shows a message that sometimes is literally a packet (without the data-link header/trailer) and sometimes a frame. Throughout the chapter, the text uses packet for all messages, because the fact of whether or not the message happens to have a data-link header/trailer at that point is immaterial to the basic discussion of features. Additionally, note that all the examples in the chapter refer to routers, just to be consistent. Classification and Marking The first QoS tool discussed in this chapter, classification and marking, or simply marking, refers to a type of QoS tool that classifies packets based on their header contents, and then marks the message by changing some bits in specific header fields. This section looks first at the role of classification across all QoS tools, and then it examines the marking feature. Classification Basics QoS tools sit in the path that packets take when being forwarded through a router or switch, much like the positioning of ACLs. Like ACLs, QoS tools are enabled on an interface. Also like ACLs, QoS tools are enabled for a direction: packets entering the interface (before the forwarding decision) or for messages exiting the interface (after the forwarding decision). The term classification refers to the process of matching the fields in a message to make a choice to take some QoS action. So, again comparing QoS tools to ACLs, ACLs perform classification and filtering; that is, ACLs match (classify) packet headers. ACLs can have the purpose (action) of choosing which packets to discard. QoS tools perform classification (matching of header fields) to decide which packets to take certain QoS actions against. Those actions include the other types of QoS tools discussed in this chapter, such as queuing, shaping, policing, and so on. For example, consider the internal processing done by a router as
shown in Figure 11-3. In this case, an output queuing tool has been enabled on an interface. Routers use queuing tools to place some packets in one output queue, other packets in another, and so on, when the outgoing interface happens to be busy. Then, when the outgoing interface becomes available to send another message, the queuing tool's scheduler algorithm can pick the next message from any one of the queues, prioritizing traffic based on the rules configured by the network engineer. 11 234 CCNA 200-301 Official Cert Guide, Volume 2 Router Internals Forward Classify Queue Scheduling (Prioritization) Transmit R1 Figure 11-3 Big Idea: Classification for Queuing in a Router The figure shows the internals of a router and what happens to the packet during part of that internal processing, moving left to right inside the router, as follows: Step 1. The router makes a forwarding (routing) decision. Step 2. The output queuing tool uses classification logic to determine which packets go into which output queue. Step 3. The router holds the packets in the output queue waiting for the outgoing interface to be available to send the next message. Step 4. The queuing tool's scheduling logic chooses the next packet, effectively prioritizing one packet over another. While the example shows a queuing tool, note that the queuing tool requires the ability to classify messages by comparing the messages to the configuration, much like ACLs. Matching (Classification) Basics Now think about classification from an enterprise-wide perspective, which helps us appreciate the need for marking. Every QoS tool can examine various headers to make comparisons to classify packets. However, you might apply QoS tools to most every device in the network, sometimes at both ingress and egress on most of the interfaces. Using complex matching of many header fields in every device and on most interfaces requires lots of configuration. The work to match packets can even degrade device performance of some devices. So, while you could have every device use complex packet matching, doing so is a poor strategy. A better strategy, one recommended both by Cisco and by RFCs, suggests doing complex matching early in the life of a packet and then marking the packet. Marking means that the QoS tool changes one or more header fields, setting a value in the header. Several header fields have been designed for the purpose of marking the packets for QoS processing. Then, devices that process the packet later in its life can use much simpler classification logic. Figure 11-4 shows an example, with a PC on the left sending an IP packet to some host off the right side of the figure (not shown). Switch SW1, the first networking device to forward the packet, does some complex comparisons and marks the packet's Differentiated Services Code Point (DSCP) field, a 6-bit field in the IP header meant for QoS marking. The next three devices that process this message—SW1, R2, and R2—then use simpler matching to classify the packet by comparing the packet's DSCP value, placing packets with one DSCP value in class 1, and packets with another DSCP value in class 2. Chapter 11: Quality of Service (QoS) More Complex Matching 235 Less Complex Matching DSCP=? DSCP=Y? CLASS 1. CLASS 2. Mark DSCP at Ingress * Figure 11-4 SW1 SW2 R1 WAN R2 ... Systematic Classification and Marking for the Enterprise Classification on Routers with ACLs and NBAR Now that you know the basics of what classification and marking do together, this section takes the discussion a little deeper with a closer look at classification on routers, which is followed by a closer look at the marking function. First, QoS classification sounds a lot like what ACLs do, and it should. In fact, many QoS tools support the ability to simply refer to an IP ACL, with this kind of logic: For any packet matched by the ACL with a permit action, consider that packet a match for QoS, so do a particular QoS action. As a reminder, Figure 11-5 shows the IP and TCP header. All these fields are matchable for QoS classification. IP Header TCP Header 9 1 2 4 Variable 2 2 16+ Miscellaneous Rest Protocol Header Source IP Destination IP Source Dest. Options Header of Port Port 6 (TCP) Checksum Address Address Fields TCP 6 = TCP Figure 11-5 Classification with Five Fields Used by Extended ACLs Now think about the enterprise's QoS plan for a moment. That plan should list details such as which types of traffic should be classified as being in the same class for queuing purposes, for shaping, and for any other QoS tool. That plan should list the fields in the header that can be matched. For instance, if all the IP phones sit in subnets within the range of addresses 10.3.0.0/16, then the QoS plan should state that. Then the network engineer could configure an extended ACL to match all packets to/from IP addresses inside 10.3.0.0/16 and apply appropriate QoS actions to that voice traffic. However, not every classification can be easily made by matching with an ACL. In more challenging cases, Cisco Network Based Application Recognition (NBAR) can be used. NBAR is basically in its second major version, called NBAR2, or next-generation NBAR. In short, NBAR2 matches packets for classification in a large variety of ways that are very useful for QoS. NBAR2 looks far beyond what an ACL can examine in a message. Many applications cannot be identified based on well-known port alone. NBAR solves those problems. 11 236 CCNA 200-301 Official Cert Guide, Volume 2 Cisco also organizes what NBAR can match in ways that make it easy to separate the traffic into different classes. For instance, the Cisco WebEx application provides audio and video conferencing on the web. In a QoS plan, you might want to classify WebEx differently than other video traffic and classify it differently than voice calls between IP phones. That is, you might classify WebEx traffic and give it a unique DSCP marking. NBAR provides easy built-in matching ability for WebEx, plus more than 1000 different subcategories of applications. Just to drive the point home with NBAR, Example 11-1 lists four lines of help output for one of many NBAR configuration commands. I chose a variety of items that might be more memorable. With the use of the keywords on the left in the correct configuration command, you could match the following: entertainment video from Amazon, video from Cisco's video surveillance camera products, video from Cisco IP Phones, and video from sports channel ESPN. (NBAR refers to this idea of defining the characteristics of different applications as application signatures.) Example 11-1 Example of the Many NBAR2 Matchable Applications R1(config)# class-map matchingexample R1(config-cmap)# match protocol attribute category voice-and-video ? ! output heavily edited for length amazon-instant-video VOD service by Amazon cisco-ip-camera Cisco video surveillance camera cisco-phone Cisco IP Phones and PC-based Unified Communicators espn-video ESPN related websites and mobile applications video facebook Facetime video calling software ! Output snipped. To wrap up the discussion of NBAR for classification, compare the first two highlighted entries in the output. Without NBAR, it would be difficult to classify an entertainment video from Amazon versus the video from a security camera, but those two highlighted entries show how you easily have classified that traffic differently. The third highlighted item shows how to match traffic for Cisco IP Phones (and PC-based equivalents), again making for an easier match of packets of a particular type. Marking IP DSCP and Ethernet CoS The QoS plan for an enterprise centers on creating classes of traffic that should receive certain types of QoS treatment. That plan would note how to classify packets into each classification and the values that should be marked on the packets, basically labeling each packet with a number to associate it with that class. For example, that plan might state the following: ■ Classify all voice payload traffic that is used for business purposes as IP DSCP AF and CoS 5. ■ Classify all video conferencing and other interactive video for business purposes as IP DSCP AF41 and CoS 4. ■ Classify all business-critical data application traffic as IP DSCP AF21 and CoS 2. This next topic takes a closer look at the specific fields that can be marked, defining the DSCP and CoS marking fields. Chapter 11: Quality of Service (QoS) 237 Marking the IP Header Marking a QoS field in the IP header works well with QoS because the IP header exists for the entire trip from the source host to the destination host. When a host sends data, the host sends a data-link frame that encapsulates an IP packet. Each router that forwards the IP packet discards the old data-link header and adds a new header. Because the routers do not discard and reinset IP headers, marking fields in the IP header stay with the data from the first place it is marked until it reaches the destination host. IPV4 defines a Type of Service (ToS) byte in the IPv4 header, as shown in Figure 11-6. The original RFC defined a 3-bit IP Precedence (IPP) field for QoS marking. That field gave us eight separate values—binary 000, 001, 010, and so on, through 111—which when converted to decimal are decimals 0 through 7. RFC 791. IPv4 Unused Old Use IP Header Type of Service (Rest of IP Header...) DSCP ECN RFC 2474 RFC 3168 Figure 11-6 Current Use IP Precedence and Differentiated Services Code Point Fields NOTE Those last 5 bits of the ToS byte per RFC 791 were mostly defined for some purpose but were not used in practice to any significant extent. While a great idea, IPP gave us only eight different values to mark, so later RFCs redefined the ToS byte with the DSCP field. DSCP increased the number of marking bits to 6 bits, allowing for 64 unique values that can be marked. The DiffServ RFCs, which became RFCs back in the late 1990s, have become accepted as the most common method to use
when doing QoS, and using the DSCP field for marking has become quite common. IPv6 has a similar field to mark as well. The 6-bit field also goes by the name DSCP, with the byte in the IPv6 header being called the IPv6 Traffic Class byte. Otherwise, think of IPv4 and IPv6 being equivalent in terms of marking. IPP and DSCP fields can be referenced by their decimal values as well as some convenient text names. The later section titled "DiffServ Suggested Marking Values" details some of the names. Marking the Ethernet 802.1Q Header Another useful marking field exists in the 802.1Q header, in a field originally defined by the IEEE 802.1p standard. This field sits in the third byte of the 4-byte 802.1Q header, as a 3-bit field, supplying eight possible values to mark (see Figure 11-7). It goes by two different names: Class of Service, or CoS, and Priority Code Point, or PCP. 11 238 CCNA 200-301 Official Cert Guide, Volume 2 Ethernet Frame Ethernet Type 802.1 Q Data Trailer Class of Service (CoS) (3 Bits) Priority Code Point (PCP) Figure 11-7 Class of Service Field in 802.1Q/p Header The figure uses two slightly different shades of gray (in print) for the Ethernet header and trailer fields versus the 802.1Q header, as a reminder: the 802.1Q header is not included in all Ethernet frames. The 802.1Q header only exists when 802.1Q trunking is used on a link. As a result, QoS tools can make use of the CoS field only for QoS features enabled on interfaces that use trunking, as shown in Figure 11-8. Trunk SW1 Trunk SW2 R1 WAN R2 ... Can Use CoS Figure 11-8 Useful Life of CoS Marking For instance, if the PC on the left were to send data to a server somewhere off the figure to the right, the DSCP field would exist for that entire trip. However, the CoS field would exist over the two links only and would be useful mainly with the arrow lines. Other Marking Fields Other marking fields also exist in other headers. Table 11-2 lists those fields for reference. Table 11-2 Field Name Marking Fields 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 DSCP 11 3 Over Wi-Fi EXP Marking Fields 3 Over MPLS WAN Defining Trust Boundaries The end-user device can mark the DSCP field—and even the CoS field if trunking is used on the link. Would you, as the network engineer, trust those settings and let your networking devices trust and react to those markings for their various QoS actions? Most of us would not, because anything the end user controls might be used inappropriately at times. For instance, a PC user could know enough about DiffServ and DSCPs to know that most voice traffic is marked with a DSCP called Expedited Forwarding (EF), which has a decimal value of 46. Voice traffic gets great QoS treatment, so PC users could mark all their traffic as DSCP 46, hoping to get great QoS treatment. Chapter 11: Quality of Service (QoS) 239 The people creating a QoS plan for an enterprise have to choose where to place the trust boundary for the network. The trust boundary refers to the point in the path of a packet flowing through the network at which the networking devices can trust the current QoS markings. That boundary typically sits in a device under the control of the IT staff. For instance, a typical trust boundary could be set in the middle of the first ingress switch in the network, as shown in Figure 11-9. The markings on the message as sent by the PC cannot be trusted. However, because SW1 performed classification and marking as the packets entered the switch, the markings can be trusted at that point. Set DSCP and CoS Inbound Untrusted SW1 SW2 R1 WAN R2 ... Trust Boundary Figure 11-9 Trusting Devices—PC Interestingly, when the access layer includes an IP phone, the phone is typically the trust boundary, instead of the access layer switch. IP phones can set the CoS and DSCP fields of the messages created by the phone, as well as those forwarded from the PC through the phone. The specific marking values are actually configured on the attached access switch. Figure 11-10 shows the typical trust boundary in this case, with notation of what the phone's marking logic usually is: mark all of the PC's traffic with a particular DSCP and/or CoS, and the phone's traffic with different values. Set PC DSCP and CoS Set Phone DSCP and CoS IP SW1 SW2 R1 WAN R2 ... Trust Boundary Figure 11-10 Trusting Devices—IP Phone DiffServ Suggested Marking Values Everything in this chapter follows the DiffServ architecture as defined originally by RFC 2475, plus many other DiffServ RFCs. In particular, DiffServ goes beyond theory in several areas, including making suggestions about the specific DSCP values to use when marking IP packets. By suggesting specific markings for specific types of traffic, DiffServ hoped to create a consistent use of DSCP values in all networks. By doing so, product vendors could provide good default settings for their QoS features. QoS could work better between an enterprise and service provider, and many other benefits could be realized. The next two topics outline three sets of DSCP values as used in DiffServ. 11 240 CCNA 200-301 Official Cert Guide, Volume 2 Expedited Forwarding (EF) DiffServ defines the Expedited Forwarding (EF) DSCP value—a single value—as suggested for use for packets that need low latency (delay), low jitter, and low loss. The Expedited Forwarding RFC (RFC 3246) defines the specific DSCP value (decimal 46) and an equivalent text name (Expedited Forwarding). QoS configuration commands allow the use of the decimal value or text name, but one purpose of having a text acronym to use is to make the value more memorable, so many QoS configurations refer to the text names. Most often QoS plans use EF to mark voice payload packets. With voice calls, some packets carry voice payload, and other packets carry call signaling messages. Call signaling messages set up (create) the voice call between two devices, and they do not require low delay, jitter, and loss. Voice payload packets carry the digitized voice, as shown back in Figure 11-2, and these packets do need better QoS. By default, Cisco IP Phones mark voice payload with EF, and mark voice signaling packets sent by the phone with another value called CS3. Assured Forwarding (AF) The Assured Forwarding (AF) DiffServ RFC (2597) defines a set of 12 DSCP values meant to be used in concert with each other. First, it defines the concept of four separate queues in a queuing system. Additionally, it defines three levels of drop priority within each queue for use with congestion avoidance tools. With four queues, and three drop priority classes per queue, you need 12 different DSCP markings, one for each combination of queue and drop priority. (Queuing and congestion avoidance mechanisms are discussed later in this chapter.) Assured Forwarding defines the specific AF DSCP text names and equivalent decimal values as listed in Figure 11-11. The text names follow a format of AFXY, with X referring to the queue (1 through 4) and Y referring to the drop priority (1 through 3). Best Drop Best Queue Worst Queue Figure 11-11 Worst Drop AF41 (34) AF42 (36) AF43 (38) AF31 (26) AF32 (28) AF33 (30) AF21 (18) AF22 (20) AF23 (22) AF11 (10) AF12 (12) AF13 (14) Differentiated Services Assured Forwarding Values and Meaning For example, if you marked packets with all 12 values, those with AF11, AF12, and AF13 would all go into one queue; those with AF21, AF22, and AF23 would go into another queue; and so on. Inside the queue with all the AF2Y traffic, you would find the AF21, AF22, and AF23 each differently in regard to drop actions (congestion avoidance), with AF21 getting the preferred treatment and AF23 the worst treatment. Chapter 11: Quality of Service (QoS) 241 Class Selector (CS) Originally, the ToS byte was defined with a 3-bit IP Precedence (IPP) field. When DiffServ redefined the ToS byte, it made sense to create eight DSCP values for backward compatibility with IPP values. The Class Selector (CS) DSCP values are those settings. Figure 11-12 shows the main idea along with the eight CS values, both in name and in decimal value. Basically, the DSCP values have the same first 3 bits as the IPP field, and with binary 0s for the last 3 bits, as shown on the left side of the figure. CSx represents the text names, where x is the matching IPP value (0 through 7). IPP DSCP CSx Figure 11-12 0 0 IPP 0 1 2 3 4 5 6 7 CS CS0 CS1 CS2 CS3 CS4 CS5 CS6 CS7 Decimal DSCP 0 8 16 24 32 40 48 56 Class Selector This section on classification and marking has provided a solid foundation for understanding the tools explored in the next three major sections of this chapter: queuing, shaping/policing, and congestion avoidance. Guidelines for DSCP Marking Values Even with this introduction to the various DSCP marking values, you could imagine that an enterprise needs to follow a convention for how to use the markings. With so many different values, having different uses of different DSCP values by different devices in the same enterprise would make deploying QoS quite difficult at best. Among its many efforts to standardize QoS, Cisco helped to develop RFC 4954, an RFC that defines several conventions for how to use the DSCP field. The RFC provides alternative plans with different levels of detail. Each plan defines a type of traffic and the DSCP value to use when marking data. Without getting into the depth of any one plan, the plans all specify some variation for how all devices should mark data as follows: ■ DSCP EF: Voice payload ■ AF4x: Interactive video (for example, videoconferencing) ■ AF3x: Streaming video ■ AF2x: High priority (low latency) data ■ CS0: Standard data Cisco not only worked to develop the RFC standards but also uses those standards. Cisco uses default marking conventions based on the marking data in RFC 4954, with some small exceptions. If you want to read more about these QoS marking plans, refer to a
couple of sources. First, look for the Cisco QoS Design Guides at Cisco.com. Also refer to RFC 4954, 11 242 CCNA 200-301 Official Cert Guide, Volume 2 Queuing All networking devices use queues. Network devices receive messages, make a forwarding decision, and then send the message—but sometimes the outgoing interface is busy. So, the device keeps the outgoing message in a queue, waiting for the outgoing interface to be available—simple enough. The term queue refers to the QoS toolset for managing the queues that hold packets while they wait their turn to exit an interface (and in other cases in which a router holds packets waiting for some resource). But queuing refers to more than one idea, so you have to look inside devices to think about how they work. For instance, consider Figure 11-13, which shows the internals of a router. The router, of course, makes a forwarding decision, and it needs to be ready to queue packets for transmission once the outgoing interface is available. At the same time, the router may take a variety of other actions as well—ingress ACL, ingress NAT (on the inside interface), egress ACLs after the forwarding decision is made, and so on. Router Internals Output Queue Forwarding Receive ingress services Figure 11-13 Transmit egress services Output Queuing in a Router. Last Output Action Before Transmission The figure shows output queuing in which the device holds messages until the output interface is available. The queuing system may use a single output queue, with a first-in, first-out (FIFO) scheduler. (In other words, it's like ordering lunch at the sandwich shop that has a single ordering line.) Next, think a little more deeply about the queuing system. Most networking devices can have a queuing system with multiple queues. To use multiple queues, the queuing system needs a classifier function to choose which packets are placed into which queue. (The classifier can react to previously marked values or do a more extensive match.) The queuing system needs a scheduler as well, to decide which message to take next when the interface becomes available, as shown in Figure 11-14. Classifier Queues Scheduler Transmit Figure 11-14 Queuing Components Of all these components of the queuing system, the scheduler can be the most interesting part because it can perform prioritization. Prioritization refers to the concept of giving priority to one queue over another in some way. Chapter 11: Quality of Service (QoS) 243 Round-Robin Scheduling (Prioritization) One scheduling algorithm used by Cisco routers and switches uses round-robin logic. In its most basic form, round robin cycles through the queues in order, taking turns with each queue. In each cycle, the scheduler either takes one message or takes a number of bytes from each queue by taking enough messages to total that number of bytes. Take some messages from queue 1, move on and take some from queue 2, then take some from queue 3, and so on, starting back at queue 1 after finishing a complete pass through the queues. Round-robin scheduling also includes the concept of weighting (generally called weighted round robin). Basically, the scheduler takes a different number of packets (or bytes) from each queue, giving more preference to one queue over another. For example, routers use a popular tool called Class-Based Weighted Fair Queuing (CBWFQ) to guarantee a minimum amount of bandwidth to each class. That is, each class receives at least the amount of bandwidth configured during times of congestion, but maybe more. Internally, CBWFQ uses a weighted round-robin scheduling algorithm, while letting the network engineer define the weightings as a percentage of link bandwidth. Figure 11-15 shows an example in which the three queues in the system have been given 20, 30, and 50 percent of the bandwidth each, respectively. Classifier Queues Scheduler Q1 20% Q2 30% Q3 50% Transmit Round Robin Figure 11-15 CBWFQ Round-Robin Scheduling With the queuing system shown in the figure, if the outgoing link is congested, the scheduler guarantees the percentage bandwidth shown in the figure to each queue. That is, queue 1 gets 20 percent of the link even during busy times. Low Latency Queuing Earlier in the chapter, the section titled "Voice and Video Applications" discussed the reasons why voice and video, particularly interactive voice and video like phone calls and videoconferencing, need low latency (low delay), low jitter, and low loss. Unfortunately, a round-robin scheduler does not provide low enough delay, jitter, or loss. The solution: add Low Latency Queuing (LLQ) to the scheduler. First, for a quick review, Table 11-3 lists the QoS requirements for a voice call. The numbers come from the Enterprise QoS Solution Reference Network Design Guide, referenced earlier in the chapter. The amount of bandwidth required per call varies based on the codec used by the call. However, the delay, jitter, and loss requirements remain the same for all voice calls. (Interactive video has similar requirements for delay, jitter, and loss.) Table 11-3 QoS Requirements for a VoIP Call per Cisco Voice Design Guide Bandwidth/call One-way Delay (max) Jitter (max) Loss (max) 30–320 Kbps 150 ms 30 ms >>> interface1 (trunk-config); 'dynamic auto', 'trunk-status'; 'static access' >>>> Using a controller-based model not only supplies APIs that give us the exact same data a human could see in show commands, but often they also supply much more useful information. A controller collects data from the entire network, so the controller can be written so that it analyzes and presents more useful data via the API. As a result, software that uses the APIs—whether automation written by local engineers or applications written by vendors—can be written more quickly and can often create features that would have been much more difficult without a controller. For instance, both APIC-EM and its successor DNA Center provide a path trace feature. The applications show the path of a packet from source to destination, with the forwarding logic used at each node. 16 378 CCNA 200-301 Official Cert Guide, Volume 2 Now imagine writing that application with either of these two approaches. ■ One API call that returns a list of all devices and their running configuration, with other API calls to collect each device's MAC address tables and/or their IP routing tables. Then you have to process that data to find the end-to-end path. ■ One API call to which you pass the source and destination IP addresses and TCP/UDP ports, and the API returns variables that describe the end-to-end path, including device hostnames and interfaces. The variables spell out the path the packet takes through the network. The second option does most of the work, while the first option leaves most of the work to you and your program. But that second option becomes possible because of the centralized controller. The controller has the data if it at least collects configuration and forwarding table information. Going beyond that, these Cisco controllers analyze the data to provide much more useful data. The power of these kinds of APIs is amazing, and this is just one example. The following list summarizes a few of the comparison points for this particular exam topic: ■ Northbound APIs and their underlying data models make it much easier to automate functions versus traditional networks. ■ The robust data created by controllers makes it possible to automate functions that were not easily automated without controllers. ■ The new reimaged software defined networks that use new operational models simplify operations, with automation resulting in more consistent configuration and less errors. ■ Centralized collection of operational data at controllers allows the application of modern data analytics to networking operational data, providing actionable insights that were likely not noticeable with the former model. ■ Time required to complete projects is reduced. ■ New operational models use external inputs, like considering time-of-day, day-of-week, and network load. Comparing Traditional Networks with Controller-Based Networks As for exam topic 6.2, this entire chapter begins to show the advantages created by using controller-based networks. However, this chapter only begins to describe the possibilities. By centralizing some of the functions in the network and providing robust APIs, controllers enable a large number of new operational models. Those models include the three most likely to be seen from Cisco in an enterprise: Software-Defined Access (SDA), Software-Defined WAN (SD-WAN), and Application-Centric Infrastructure (ACI). (Chapter 17 introduces SDA.) This changes the operating paradigm in many cases, with the controller determining many device-specific details: ■ The network engineer does not need to think about every command on every device. ■ The controller configures the devices with consistent and streamlined settings. ■ The result: faster and more consistent changes with fewer issues. Chapter 16: Introduction to Controller-Based Networking 379 As another example, just consider the ACI example from earlier in the chapter. Instead of configuring each port with an access VLAN, or making it a trunk, adding routing protocol configuration, and possibly updating IP ACLs, all you had to do was create some endpoint groups (EPGs) and policies. In that case, the orchestration software that started the VMs could automatically create the EPGs and policies. The new paradigm of intent-based networking was enabled through the controller-based architecture. Then the automation features enabled by the controller's northbound APIs allowed third-party applications to automatically configure the network to support the necessary changes. Some of the advantages include the following: 16 ■ Uses new and improved operational models that allow the configuration of the network rather than per-device configuration ■ Enables automation through
northbound APIs that provide robust methods and model-driven data ■ Configures the network devices through southbound APIs, resulting in more consistent device configuration, fewer errors, and less time spent troubleshooting the network ■ Enables a DevOps approach to networks Chapter 17 goes into some depth comparing traditional networking with controller-based networks with descriptions of Cisco Software-Defined Access (SDA). Look throughout that chapter for some of the reasons and motivations for SDA and the features enabled by using the DNA Center controller. Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 16-4 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 16-4 Chapter Review Tracking Review Element Review Date(s) Resource Used Review key topics Book, website Review key terms Book, website Answer DIKTA questions Book, PTP Review memory tables Book, website Watch video Website 380 CCNA 200-301 Official Cert Guide, Volume 2 Review All the Key Topics Table 16-5 Key Topics for Chapter 16 Key Topic Element Description Page Number List Sample actions of the networking device data plane 359 List Sample actions of the networking device control plane 360 Figure 16-4 Switch internals with ASIC and TCAM 362 Figure 16-5 Basic SDN architecture, with the centralized controller programming device data planes directly 364 Paragraph Description of the role and purpose of the NBI 365 Figure 16-7 REST API basic concepts 366 List Spine-leaf topology requirements 370 Figure 16-10 Spine-leaf switch 371 Figure 16-13 Controlling the ACI data center network using APIC 373 Table 16-2 Comparisons of Open SDN, Cisco ACI, and Cisco APIC Enterprise options 375 List Comparisons of how automation improves network management 378 List Comparisons of how controller-based networking versus traditional networking 379 Key Terms You Should Know application programming interface (API), Application Policy Infrastructure Controller (APIC), APIC Enterprise Module (APIC-EM), Application-Centric Infrastructure (ACI), northbound API, southbound API, control plane, data plane, management plane, application-specific integrated circuit (ASIC), ternary content-addressable memory (TCAM), OpenFlow, Software-Defined Networking (SDN), distributed control plane, centralized control plane, northbound interface (NBI), southbound interface (SBI), controller-based networking, intent-based networking (IBN), spine, leaf This page intentionally left blank CHAPTER 17 Cisco Software-Defined Access (SDA) This chapter covers the following exam topics: 1.0 Network Fundamentals 1.1 Explain the role and function of network components 1.1.e Controllers (Cisco DNA Center and WLC) 6.0 Automation and Programmability 6.1 Explain how automation impacts network management 6.2 Compare traditional networks with controller-based networking 6.3 Describe controller-based and software-defined architectures (overlay, underlay, and fabric) 6.3.a Separation of control plane and data plane 6.3.b Northbound and southbound APIs 6.4 Compare traditional campus device management with Cisco DNA Center enabled device management Cisco Software-Defined Access (SDA) uses a software-defined networking approach to build a converged wired and wireless campus LAN. The word access in the name refers to the endpoint devices that access the network, while software-defined refers to many of the usual software-defined architectural features discussed in Chapter 16, "Introduction to Controller-Based Networking." Those features include a centralized controller—DNA Center—with southbound and northbound protocols. It also includes a completely different operational model inside SDA, with a network fabric composed of an underlay network and an overlay network. SDA fills the position as Cisco's

campus offering within Cisco Digital Network Architecture (DNA). Cisco DNA defines the entire architecture for the new world of software defined networks, digitalization, and Cisco's reimagining of how networks should be operated in the future. This chapter introduces SDA, which exists as one implementation of Cisco DNA. The discussion of SDA and DNA provides a great backdrop to discuss a few other topics from the CCNA blueprint: the DNA Center controller and network management. SDA uses the DNA Center controller to configure and operate SDA. However, DNA Center also acts as a complete network management platform. To understand DNA Center, you also need to understand traditional network management as well as the new management models using controllers. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 17-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions SDA Fabric, Underlay, and Overlay 1–3 DNA Center and SDA Operation 4, 5 DNA Center as a Network Management Platform 6 1. In Cisco Software-Defined Access (SDA), which term refers to the devices and cabling, along with configuration that allows the network device nodes enough IP connectivity to send IP packets to each other? a. Fabric b. Overlay c. Underlay d. VXLAN 2. In Cisco Software-Defined Access (SDA), which term refers to the functions that deliver endpoint packets across the network using tunnels between the ingress and egress fabric nodes? a. Fabric b. Overlay c. Underlay d. VXLAN 3. In Software-Defined Access (SDA), which of the answers are part of the overlay data plane? a. LISP b. GRE c. OSPF d. VXLAN 4. Which answers best describe options of how to implement security with scalable groups using DNA Center and SDA? (Choose two answers.) a. A human user from the DNA Center GUI b. An automation application using NETCONF c. A human user using the CLI of an SDA fabric edge node d. An automation application using REST 384 CCNA 200-301 Official Cert Guide, Volume 2 5. Which of the following protocols or tools could be used as part of the Cisco DNA Center southbound interface? (Choose three answers.) a. Ansible b. SSH c. Application using NETCONF d. SNMP e. Puppet 6. Which of the following are network management features performed by both traditional network management software as well as by DNA Center? (Choose two answers.) a. Network device discovery b. Software-Defined Access configuration c. End-to-end path discovery with ACL analysis d. Device installation (day 0), configuration (day 1), and monitoring (day n) operations Foundation Topics SDA Fabric, Underlay, and Overlay Cisco Software-Defined Access (SDA) creates an entirely new way to build campus LANs as compared with the traditional methods of networking discussed in most chapters of this book. In the mid 2010s, Cisco set out to reimagine campus networking, with SDA as the result. SDA uses the software-defined architectural model introduced in Chapter 16, with a controller and various APIs. It still uses a physical network with switches and routers, cables, and various endpoints. At the center sits the Digital Network Architecture (DNA) Center controller, as shown in Figure 17-1, with human users making use of a graphical user interface (GUI) and automation using APIs. In short, DNA Center is the controller for SDA networks. Architecturally, the southbound side of the controller contains the fabric, underlay, and overlay. By design in SDN implementations, most of the interesting new capabilities occur on the northbound side, which are examined in the second half of this chapter. This first half of the chapter examines the details south of the controller—namely, the fabric, underlay network, and overlay network. Overlay: The mechanisms to create VXLAN tunnels between SDA switches, which are then used to transport traffic from one fabric endpoint to another over the fabric. Underlay: The network of devices and connections (cables and wireless) to provide IP connectivity to all nodes in the fabric, with a goal to support the dynamic discovery of all SDA devices and endpoints as a part of the process to create overlay VXLAN tunnels. Fabric: The combination of overlay and underlay, which together provide all features to deliver data across the network with the desired features and attributes. Chapter 17: Cisco Software-Defined Access (SDA) Script GUI Guit 385 GUI API Cisco or Vendor App API Controller API SBI 17 Figure 17-1 SDA Architectural Model with DNA Center In less formal terms, the underlay exists as multilayer switches and their links, with IP connectivity—but for a special purpose. The underlay supports some new concepts with a tunneling method called VXLAN. Traffic sent by the endpoint devices flows through VXLAN tunnels in the overlay—a completely different process than traditional LAN switching and IP routing. For instance, think about the idea of sending packets from hosts on the left of a network, over SDA, to hosts on the right. For instance, imagine a packet enters on the left side of the physical network at the bottom of Figure 17-2 and eventually exits the campus out switch SW2 on the far right. This underlay network looks like a more traditional network drawing, with several devices and links. The overlay drawing at the top of the figure shows only two switches—called fabric edge nodes, because they happen to be at the edges of the SDA fabric—with a tunnel labeled VXLAN connecting the two. Both concepts (underlay and overlay) together create the SDA fabric. The next few pages explain both the underlay and overlay in a little more depth. 386 CCNA 200-301 Official Cert Guide, Volume 2 Overlay SW2 SW1 VXLAN SW1 SW2 Underlay Figure 17-2 Fabric, Underlay, and Overlay Concepts The SDA Underlay With SDA, the underlay exists to provide connectivity between the nodes in the SDA environment for the purpose of supporting VXLAN tunnels in the overlay network. To do that, the underlay includes the switches, routers, cables, and wireless links used to create the physical network. It also includes the configuration and operation of the underlay so it can support the work of the overlay network. Using Existing Gear for the SDA Underlay To build an SDA underlay network, companies have two basic choices. They can use their existing campus network and add new configuration to create an underlay network, while still supporting their existing production traffic with traditional routing and switching. Alternately, the company can purchase some new switches and build the SDA network without concern for harming existing traffic, and migrate endpoints to the new SDA network over time. To build SDA into an existing network, it helps to think for a moment about some typical campus network designs. The larger campus site may use either a two-tier or three-tier design as discussed in Chapter 13, "LAN Architecture." It has a cluster of wireless LAN controllers (WLCs) to support a number of lightweight APs (LWAPs). Engineers have configured VLANs, VLAN trunks, IP routing, IP routing protocols, ACLs, and so on. And the LAN connects to WAN routers. Answers to the "Do I Know This Already?" quiz: 1 C 2 B 3 D 4 A, D 5 B, C, D 6 A, D Chapter 17: Cisco Software-Defined Access (SDA) 387 SDA can be added into an existing campus LAN, but doing so has some risks and restrictions. First and foremost, you have to be careful not to disrupt the current network while adding the new SDA features to the network. The issues include ■ Because of the possibility of harming the existing production configuration, DNA Center should not be used to configure the underlay if the devices are currently used in production. (DNA Center will be used to configure the underlay with deployments that use all new hardware.) ■ The existing hardware must be from the SDA compatibility list, with different models supported depending on their different SDA roles (see a link at www.cisco.com/go/sda). ■ The device software levels must meet the requirements, based on their roles, as detailed in that same compatibility list. 17 For instance, imagine an enterprise happened to have an existing campus network that uses SDA-compatible hardware. That company might need to update the IOS versions in a few cases. Additionally, the engineers would need to configure the underlay part of the SDA devices manually rather than with DNA Center because Cisco assumes that the existing network already supports production traffic, so they want the customer directly involved in making those changes. The SDA underlay configuration requires you to think about and choose the different SDA roles filled by each device before you can decide which devices to use and which minimum software levels each requires. If you look for the hardware compatibility list linked from www.cisco.com/go/sda, you will see different lists of supported hardware and software depending on the roles. These roles include Fabric edge node: A switch that connects to endpoint devices (similar to traditional access switches) Fabric border node: A switch that connects to devices outside SDA's control, for example, switches that connect to the WAN routers or to an ACI data center Fabric control node: A switch that performs special control plane functions for the underlay (LISP), requiring more CPU and memory For example, when I was writing this chapter back in 2019, Cisco's compatibility list included many Catalyst 9300, 9400, and 9500 switches, but also some smaller Catalyst 3850 and 3650 switches, as fabric edge nodes. However, the Catalyst 2960X or
2960Xr products did not make the list as fabric edge nodes. For fabric control nodes, the list included more higher-end Catalyst switch models (which typically have more CPU and RAM), plus several router models (routers typically have more RAM for control plane protocol storage—for instance, for routing protocols). The beginning of an SDA project will require you to look at the existing hardware and software to begin to decide whether the existing campus might be a good candidate to build the fabric with existing gear or to upgrade hardware when building the new campus LAN. Using New Gear for the SDA Underlay When buying new hardware for the SDA fabric—that is, a greenfield design—you remove many of the challenges that exist when deploying SDA on existing gear. You can simply order compatible hardware and software. Once it arrives, DNA Center can then configure all the underlay features automatically. 388 CCNA 200-301 Official Cert Guide, Volume 2 At the same time, the usual campus LAN design decisions still need to be made. Enterprises use SDA as a better way to build and operate a campus network, but SDA is still a campus network. It needs to provide access and connectivity to all types of user devices. When planning a greenfield SDA design, plan to use SDA-compatible hardware, but also think about these traditional LAN design points: ■ The number of ports needed in switches in each wiring closet ■ The port speeds required ■ The benefit of a switch stack in each wiring closet ■ The cable length and types of cabling already installed ■ The need for power (PoE/PoE+) ■ The power available in each new switch versus the PoE power requirements ■ Link capacity (speed and number of links) for links between switches As far as the topology, traditional campus design does tell us how to connect devices, but SDA does not have to follow those traditional rules. To review, traditional campus LAN Layer 2 design (as discussed back in Chapter 13) tells us to connect each access switch to two different distribution layer switches, but not to other access layer switches, as shown in Figure 17-3. The access layer switch acts as a Layer 2 switch, with a VLAN limited to those three switches. HSRP 10.1.1.1 HSRP 10.1.1.1 L2 Distribution Layer (Layer 3 Switches) Root RSTP L2 L2 Block Access Layer (Layer 2 Switches) SW3 GW = 10.1.1.1 Figure 17-3 Traditional Access Layer Design: Three Switches in STP Triangle Take a moment to reflect about the traditional features shown in the figure. The distribution layer switches—Layer 3 switches—act as the default gateway used by hosts and often implement HSRP for better availability. The design uses more than one uplink from the access to distribution layer switches, with Layer 2 EtherChannels, to allow balancing in addition to redundancy. And STP/RSTP manages the small amount of Layer 2 redundancy in the campus, preventing loops by blocking on some ports. In comparison, a greenfield SDA fabric uses a routed access layer design. Routed access layer designs have been around long before SDA, but SDA makes good use of the design. Chapter 17: Cisco Software-Defined Access (SDA) 389 and it works very well for the underlay with its goal to support VXLAN tunnels in the overlay network. A routed access layer design simply means that all the LAN switches are Layer 3 switches, with routing enabled, so all the links between switches operate as Layer 3 links. With a greenfield SDA deployment—that is, all new gear that you can allow to be configured by DNA Center—DNA Center will configure the devices' underlay configuration to use a routed access layer. Because DNA Center knows it can configure the switches without concern of harming a production network, it chooses the best underlay configuration to support SDA. That best configuration happens to use a design called a routed access layer design, which has these features: ■ All switches act as Layer 3 switches. ■ The switches use the IS-IS routing protocol. ■ All links between switches (single links, EtherChannels) are routed Layer 3 links (not Layer 2 links). ■ As a result, STP/RSTP is not needed, with the routing protocol instead choosing which links to use based on the IP routing tables. ■ The equivalent of a traditional access layer switch—an SDA edge node—acts as the default gateway for the endpoint devices, rather than distribution switches. ■ As a result, HSRP (or any FHRP) is no longer needed. Figure 17-4 repeats the same physical design as in Figure 17-3 but shows the different features with the routed access design as configured using DNA Center. HSRP 10.1.1.1 SW1 HSRP 10.1.1.1 SW2 L3 Distribution Layer (Layer 3 Switches) RSTP IS-IS L3 Access Layer (Layer 3 Switches) L3 SW3 IP 10.1.1.1 GW = 10.1.1.1 Figure 17-4 SDA Fabric Layer 3 Access Benefits NOTE DNA Center configures the underlay with consistent settings for each instance of DNA across an enterprise. This convention simplifies operation as an enterprise completes a migration to SDA. 17 390 CCNA 200-301 Official Cert Guide, Volume 2 The SDA Overlay When you first think of the SDA overlay, think of this kind of sequence. First, an endpoint sends a frame that will be delivered across the SDA network. The first SDA node to receive the frame encapsulates the frame in a new message—using a tunneling specification called VXLAN—and forwards the frame into the fabric. Once the ingress node has encapsulated the original frame in VXLAN, the other SDA nodes forward the frame based on the VXLAN tunnel details. The last SDA node removes the VXLAN details, leaving the original frame, and forwards the original frame on toward the destination endpoint. While the summary of some of SDA's overlay work in the previous paragraph may sound like a lot of work, all that work happens in each switch's ASIC. So, while it is more complex to understand, there is no performance penalty for the switches to perform the extra work. When Cisco set about to create SDA, they saw an opportunity. Making use of VXLAN tunnels opened up the possibilities for a number of new networking features that did not exist without VXLAN. This next topic begins with a closer look at the VXLAN tunnels in the overlay, followed by a discussion of how SDA uses LISP for endpoint discovery and location needed to create the VXLAN tunnels. VXLAN Tunnels in the Overlay (Data Plane) SDA has many additional needs beyond the simple message delivery—needs that let it provide improved functions. To that end, SDA does not only route IP packets or switch Ethernet frames. Instead, it encapsulates incoming data link frames in a tunneling technology for delivery across the SDA network, with these goals in mind: ■ The VXLAN tunneling (the encapsulation and de-encapsulation) must be performed by the ASIC on each switch so that there is no performance penalty. (That is one reason for the SDA hardware compatibility list: the switches must have ASICs that can perform the work.) ■ The VXLAN encapsulation must supply header fields that SDA needs for its features, so the tunneling protocol should be flexible and extensible, while still being supported by the switch ASICs. ■ The tunneling encapsulation needs to encapsulate the entire data link frame instead of encapsulating the IP packet. That allows SDA to support Layer 2 forwarding features as well as Layer 3 forwarding features. To achieve those goals, when creating SDA, Cisco chose the Virtual Extensible LAN (VXLAN) protocol to create the tunnels used by SDA. When an SDA endpoint (for example, an end-user computer) sends a data link frame into an SDA edge node, the ingress edge node encapsulates the frame and sends it across a VXLAN tunnel to the egress edge node, as shown in Figure 17-5. To support the VXLAN encapsulation, the underlay uses a separate IP address space as compared with the rest of the enterprise, including the endpoint devices that send data over the SDA network. The overlay tunnels use addresses from the enterprise address space. For instance, imagine an enterprise used these address spaces: ■ 10.0.0.0/8: Entire enterprise ■ 172.16.0.0/16: SDA underlay Chapter 17: Cisco Software-Defined Access (SDA) Ingress Fabric Edge Node 391 Egress Fabric Edge Node SW1 SW2 10.1.1.1 10.1.2.2 LAN Frame IP Figure 17-5 UDP VXLAN LAN Frame Fundamentals of VXLAN Encapsulation in SDA To make that work, first the underlay would be built using the 172.16.0.0/16 IPv4 address space, with all links using addresses from that address space. As an example, Figure 17-6 shows a small SDA design, with four switches, each with one underlay IP address shown (from the 172.16.0.0/16 address space). 172.16.1.1 172.16.2.2 SW1 SW2 172.16.4.4 SW4 Figure 17-6 SDA Underlay Using 172.16.0.0 The overlay tunnel creates a path between two fabric edge nodes in the overlay IP address space—that is, in the same address space used by all the endpoints in the enterprise. Figure 17-7 emphasizes that point by showing the endpoints (PCs) on the left and right, with IP addresses in network 10.0.0.0/8, with the VXLAN overlay tunnel shown with addresses also from 10.0.0.0/8. 10.1.1.1 10.1.2.2 SW1 10.3.3.1 10.3.3.2 SW2 VXLAN Tunnel Subnets of Enterprise (Overlay) Figure 17-7 VXLAN Tunnel and Endpoints with IPv4 Addresses in the Same IPV4 Space 17 392 CCNA 200-301 Official Cert Guide, Volume 2 LISP for Overlay Discovery and Location (Control Plane) Ingress SDA for a moment, and think about traditional Layer 2 switching and Layer 3 routing. How do their control planes work? In other words, how do these devices discover the possible destinations in the network, store those destinations, so that the data plane has all the data it needs when making a forwarding decision? To summarize: ■ Traditional Layer 2 switches learn possible destinations by examining the source MAC addresses of incoming frames, storing those MAC addresses as possible future destinations in the switch's MAC address table. When new frames arrive, the Layer 2 switch data plane then attempts to
match the Ethernet frame's destination MAC address to an entry in its MAC address table. ■ Traditional Layer 3 routers learn destination IP subnets using routing protocols, storing routes to reach each subnet in their routing tables. When new packets arrive, the Layer 3 data plane attempts to match the IP packet's destination IP address to some entry in the IP routing table. Nodes in the SDA network do not do these same control plane actions to support endpoint traffic. Just to provide a glimpse into the process for the purposes of CCNA, consider this sequence, which describes one scenario: ■ Fabric edge nodes—SDA nodes that connect to the edge of the SDA fabric—learn the location of possible endpoints using traditional means, based on their MAC address, individual IP address, and by subnet, identifying each endpoint with an endpoint identifier (EID). ■ The fabric edge nodes register the fact that the node can reach a given endpoint (EID) into a database called the LISP map server. ■ The LISP map server keeps the list of endpoint identifiers (EIDs) and matching routing locators (RLOCs) (which identify the fabric edge node that can reach the EID). ■ In the future, when the fabric data plane needs to forward a message, it will look for and find the destination in the LISP map server's database. For instance, switches SW3 and SW4 in Figure 17-8 each just learned about different subnets external to the SDA fabric. As noted at step 1 in the figure, switch SW3 sent a message to the LISP map server, registering the information about subnet 10.1.3.0/24 (an EID), with its RLOC setting to identify itself as the node that can reach that subnet. Step 2 shows an equivalent registration process, this time for SW4, with EID 10.1.4.0/24, and with R4's RLOC of 172.16.4.4. Note that the table at the bottom of the figure represents that data held by the LISP map server. Chapter 17: Cisco Software-Defined Access (SDA) 393 RLOC 172.16.3.3 SW1 SW3 EID 10.1.3.0/24 1 RLOC 172.16.4.4 SW4 2 LISP Map Server 17 EID RLOC 1 10.1.3.0/24 172.16.3.3 2 10.1.4.0/24 172.16.4.4 Figure 17-8 EID 10.1.4.0/24 Edge Nodes Register IPv4 Prefixes (Endpoint IDs) with LISP Map Server When new incoming frames arrive, the ingress tunnel router (ITR)—the SDA node that receives the new frame from outside the SDA fabric—needs some help from the control plane. To where should the ITR forward this frame? And because SDA always forwards frames in the fabric over some VXLAN tunnel, what tunnel should the ITR use when forwarding the frame? For the first frame sent to a destination, the ITR has to follow a process like the following steps. The steps begin at step 3, as a continuation of Figure 17-8, with the action referenced in Figure 17-9: 3. An Ethernet frame to a new destination arrives at ingress edge node SW1 (upper left), and the switch does not know where to forward the frame. 4. The ingress node sends a message to the LISP map server asking if the LISP server knows how to reach IP address 10.1.3.1. 5. The LISP map server looks in its database and finds the entry it built back at step 1 in the previous figure, listing SW3's RLOC of 172.16.3.3. 6. The LISP map server contacts SW3—the node listed as the RLOC—to confirm that the entry is correct. 7. SW3 completes the process of informing the ingress node (SW1) that 10.1.3.1 can be reached through SW3. 394 CCNA 200-301 Official Cert Guide, Volume 2 Dest. = 10.1.3.1 3 RLOC 172.16.3.3 SW1 SW3 7 EID 10.1.3.0/24 4 EID RLOC 10.1.3.0/24 172.16.3.3 10.1.4.0/24 172.16.4.4 Figure 17-9 LISP SW4 EID 10.1.4.0/24 LISP Map Server 5 RLOC 172.16.4.4 Ingress Tunnel Router SW1 10.1.3.1 Ingress Tunnel Router SW3 Using To complete the story, now that ingress node SW1 knows that it can forward packets sent to endpoint 10.1.3.1 to the edge node with RLOC 172.16.3.3 (that is, SW3), SW1 encapsulates the original Ethernet frame as shown in Figure 17-9, with the original destination IP address of 10.1.3.1. It adds the IP, UDP, and VXLAN headers shown so it can deliver the message over the SDA network, with that outer IP header listing a destination IP address of the RLOC IP address, so that the message will arrive through the SDA fabric at SW3, as shown in Figure 17-10. At this point, you should have a basic understanding of how the SDA fabric works. The underlay includes all the switches and links, along with IP connectivity, as a basis for forwarding data across the fabric. The overlay adds a different level of logic, with endpoint traffic flowing through VXLAN tunnels. This chapter has not mentioned any reasons that SDA might want to use these tunnels, but you will see one example by the end of the chapter. Suffice it to say that with the flexible VXLAN tunnels, SDA can encode header fields that let SDA create new networking features, all without suffering a performance penalty, as all the VXLAN processing happens in an ASIC. This chapter next focuses on DNA Center and its role in managing and controlling SDA fabrics. Chapter 17: Cisco Software-Defined Access (SDA) Dest. = 172.16.3.3 IP UDP 395 Dest. = 10.1.3.1 VXLAN Original SW1 RLOC 172.16.3.3 SW3 RLOC 172.16.4.4 17 SW4 10.1.4.0/24 LISP Map Server EID RLOC 1 10.1.3.0/24 172.16.3.3 2 10.1.4.0/24 172.16.4.4 Figure 17-10 SW3 Ingress Tunnel Router (ITR) SW1 Forwards Based on LISP Mapping to DNA Center and SDA Operation Cisco DNA Center (www.cisco.com/go/dnacenter) has two notable roles: ■ As the controller in a network that uses Cisco SDA ■ As a network management platform for traditional (non-SDA) network devices, with an expectation that one day DNA Center may become Cisco's primary enterprise network management platform The first role as SDA network controller gets most of the attention and is the topic of discussion in this second of the three major sections of this chapter. SDA and DNA Center go together, work closely together, and any serious use of SDA requires the use of DNA Center. At the same time, DNA Center can manage traditional network devices; the final major section of the chapter works through some comparisons. Cisco DNA Center Cisco DNA Center exists as a software application that Cisco delivers pre-installed on a Cisco DNA Center appliance. The software follows the same general controller architecture concepts as described in Chapter 16. Figure 17-11 shows the general ideas. 396 CCNA 200-301 Official Cert Guide, Volume 2 Script GUI Script GUI REST API Cisco or Vendor App DNA Center REST API Telnet/SSH SNMP Figure 17-11 REST API NETCONF RESTCONF Cisco DNA Center with Northbound and Southbound Interfaces Cisco DNA Center includes a robust northbound REST API along with a series of southbound APIs. For most of us, the northbound API matters most, because as the user of SDA networks, you interact with SDA using Cisco DNA Center's northbound REST API or the GUI interface. (Chapter 18, "Understanding REST and JSON," discusses the concepts behind REST APIs in more detail.) Cisco DNA Center supports several southbound APIs so that the controller can communicate with the devices it manages. You can think of these in two categories: ■ Protocols to support traditional networking devices/software versions: Telnet, SSH, SNMP ■ Protocols to support more recent networking devices/software versions: NETCONF, RESTCONF Cisco DNA Center needs the older protocols to be able to support the vast array of older Cisco devices and OS versions. Over time, Cisco has been adding support for NETCONF and RESTCONF to their more current hardware and software. Cisco DNA Center and Scalable Groups SDA creates many interesting new and powerful features beyond how traditional campus networks work. Cisco DNA Center not only enables an easier way to configure and operate those features, but it also completely changes the operational model. While the scope of CCNA does not allow us enough space to explore all of the features of SDA and DNA Center, this next topic looks at one feature as an example: scalable groups. Chapter 17: Cisco Software-Defined Access (SDA) 397 Issues with Traditional IP-Based Security Imagine the life of one traditional IP ACL in an enterprise. Some requirements occurred, and an engineer built the first version of an ACL with three Access Control Entries (ACEs)—that is, access-list commands—with a permit any at the end of the list. Months later, the engineer added two more lines to the ACL, so the ACL has the number of ACEs shown in Figure 17-12. The figure notes the lines added for requests one and two with the circled numbers in the figure. ACE 1 ACE 2 (First Request) 2 (Two Months Later) ACE 3 7 ACE 4 ACE 5 Permit Figure 17-12 Lines (ACEs) in an ACL after Two Changes Now think about that same ACL after four more requirements caused changes to the ACL, as noted in Figure 17-13. Some of the movement includes ■ The ACEs for requirement two are now at the bottom of the ACL. ■ Some ACEs, like ACE 5, apply to more than one of the implemented requirements. ■ Some requirements, like requirement number five, required ACEs that overlap with multiple other requirements. ACE 1 ACE 2 1 ACE 3 ACE 4 5 3 ACE 5 ACE 6 ACE 7 ACE 8 4 ACE 9 ACE 10 ACE 11 2 ACE 12 (Permit) Figure 17-13 Lines (ACEs) in an ACL after Six Changes Now imagine your next job is to add more ACEs for the next requirement (7). However, your boss also told you to reduce the length of the ACL, removing the ACEs from that one change made last August—you remember it, right? Such tasks are problematic at best. 398 CCNA 200-301 Official Cert Guide, Volume 2 With the scenario in Figure 17-13, no engineer could tell from looking at the ACL whether any lines in the ACL could be safely removed. You never know if an ACE was useful for one requirement or for many. If a requirement was removed, and you were even told which old project caused the original requirement so that you could look at your notes, you would not know if removing the ACEs would harm other requirements. Most of the time, ACL management suffers with these kinds of issues:
■ ACEs cannot be removed from ACLs because of the risk of causing failures to the logic for some other past requirement. ■ New changes become more and more challenging due to the length of the ACLs. ■ Troubleshooting ACLs as a system—determining whether a packet would be delivered from end-to-end—becomes an even greater challenge. SDA Security Based on User Groups Imagine you could instead enforce security without even thinking about IP address ranges and ACLs. SDA does just that, with simple configuration, and the capability to add and remove the security policies at will. First, for the big ideas, imagine that over time, using SDA, six different security requirements occurred. For each project, the engineer would define the policy with DNA Center, either with the GUI or with the API. Then, as needed, DNA Center would configure the devices in the fabric to enforce the security, as shown in Figure 17-14. Policy 1 Policy 2 Policy 4 Policy 3 Policy 5 Policy 6 DNA-C SDA Fabric Figure 17-14 DNA-C IP Security Policies (Northbound) to Simplify Operations NOTE The model in Figure 17-14 helps demonstrate the concept of intent-based networking (IBN). The engineer configures the intent or outcome desired from the network—in this case, a set of security policies. The controller communicates with the devices in the network, with the devices determining exactly what configuration and behavior are necessary to achieve those intended policies. Chapter 17: Cisco Software-Defined Access (SDA) 399 The SDA policy model solves the configuration and operational challenges with traditional ACLs. In fact, all those real issues with managing IP ACLs on each device are no longer issues with SDA's group-based security model. For instance: ■ The engineer can consider each new security requirement separately, without analysis of an existing (possibly lengthy) ACL. ■ Each new requirement can be considered without searching for all the ACLs in the likely paths between endpoints and analyzing each and every ACL. ■ DNA Center (and related software) keeps the policies separate, with space to keep notes about the reason for the policy. ■ Each policy can be removed without fear of impacting the logic of the other policies. SDA and Cisco DNA achieve this particular feature by tying security to groups of users, called scalable groups, with each group assigned a scalable group tag (SGT). Then the engineer configures a grid that identifies which SGTs can send packets to which other SGTs. For instance, the grid might include SGTs for an employee group, the Internet (for the Enterprise's WAN routers that lead to the Internet), partner employees, and guests, with a grid like the one shown in Table 17-2. Table 17-2 Access Table for SDA Scalable Group Access Dest. Employee Internet Partner Guest Deny Source Employee N/A Permit Permit Internet Permit N/A Deny Guest Deny Permit Deny N/A To link this security feature back to packet forwarding, consider when a new endpoint tries to send its first packet to a new destination. The ingress SDA node starts a process by sending messages to DNA Center. DNA Center then works with security tools in the network, like Cisco's Identity Services Engine (ISE), to identify the users and then match them to their respective SGTs. DNA Center then checks the logic similar to Table 17-2. If DNA Center sees a permit action between the source/destination pair of SGTs, DNA Center directs the edge nodes to create the VXLAN tunnel, as shown in Figure 17-15. If the security policies state that the two SGTs should not be allowed to communicate, DNA Center does not direct the fabric to create the tunnel, and the packets do not flow. SW1 SW1 SW2 10.1.1.1 10.1.2.2 Source Dest. Source Dest. IP UDP SGT VNIID VNIID Original Ethr VXLAN Figure 17-15 VXLAN Header with Source and Destination SGTs and VNIDs Revealed 17 400 CCNA 200-301 Official Cert Guide, Volume 2 NOTE The figure gives a brief insight into why SDA goes to the trouble of using VXLAN encapsulation for its data plane, rather than performing traditional Layer 2 switching or Layer 3 routing. The VXLAN header has great flexibility—in this case, used to define both a source and destination SGT, matching SDA's desired logic of allowing a subset of source/destination SGTs in the SDA fabric. The operational model with scalable groups greatly simplifies security configuration and ongoing maintenance of the security policy, while focusing on the real goal: controlling access based on user. From a controller perspective, the fact that Cisco DNA Center acts as much more than a management platform, and instead as a controller of the activities in the network, makes for a much more powerful set of features and capabilities. DNA Center as a Network Management Platform CCNA Exam topic 6.4 asks you to compare traditional network management with DNA Center. Compare traditional campus device management with Cisco DNA Center enabled device management Note that the exam topic does not identify which traditional management product. In fact, Cisco tends to shy away from product details in most of its career certifications. So, to think through this exam topic, you need to think in general about network management products. But it also helps to think about specific products—but temper that by focusing on the more prominent features and major functions. This section uses Cisco Prime Infrastructure (PI) (www.cisco.com/go/primeinfrastructure) as an example of a traditional enterprise network management product. For many years, Cisco Prime Infrastructure has been Cisco's primary network management product for the enterprise. It includes the following features: ■ Single-pane-of-glass: Provides one GUI from which to launch all PI functions and features ■ Discovery, inventory, and topology: Discovers network devices, builds an inventory, and arranges them in a topology map ■ Entire enterprise: Provides support for traditional enterprise LAN, WAN, and data center management functions ■ Methods and protocols: Uses SNMP, SSH, and Telnet, as well as CDP and LLDP, to discover and learn information about the devices in the network ■ Lifecycle management: Supports different tasks to install a new device (day 0), configure it to be working in production (day 1), and perform ongoing monitoring and make changes (day n) ■ Application visibility: Simplifies QoS configuration deployment to each device ■ Converged wired and wireless: Enables you to manage both the wired and wireless LAN from the same management platform Chapter 17: Cisco Software-Defined Access (SDA) ■ Software Image Management (SWIM): Manages software images on network devices and automates updates ■ Plug-and-Play: Performs initial installation tasks for new network devices after you physically install the new device, connect a network cable, and power on 401 PI itself runs as an application on a server platform with GUI access via a web browser. The PI server can be purchased from Cisco as a software package to be installed and run on your servers, or as a physical appliance. The next few pages now compare and contrast DNA Center to traditional management tools like PI. DNA Center Similarities to Traditional Management If you read the user's guide for DNA Center and look through all the features, you will find all the features just listed here as traditional management features. For instance, both can discover network devices and create a network topology map. Human operators (rather than automated processes) often start with the topology map, expecting notices (flashing lights, red colors) to denote issues in the network. As an example, Figure 17-16 shows a topology map from DNA Center. Both PI and DNA Center can perform a discover process to find all the devices in the network and then build topology maps to show the devices. (Interestingly, DNA Center can work with PI, using the data discovered by PI rather than performing the discovery work again.) Figure 17-16 DNA Center Topology Map The GUI mechanisms are relatively intuitive, with the ability to click into additional or less detail. Figure 17-17 shows a little more detail after hovering over and clicking on one of the nodes in the topology from Figure 17-16, typical actions and results in many management products. 17 402 CCNA 200-301 Official Cert Guide, Volume 2 Figure 17-17 Hover and Click Details About One Cisco 9300 Switch from DNA Center I encourage you to take some time to use and watch some videos about Cisco DNA Center. The "Chapter Review" section for this chapter on the companion website lists some links for good videos. Also, start at and look for Cisco DNA Center sandbox labs to find a place to experiment with Cisco DNA Center. DNA Center Differences with Traditional Management In a broad sense, there are several fundamental differences between Cisco DNA Center and traditional network management platforms like Cisco PI. The largest difference: Cisco DNA Center supports SDA, whereas other management apps do not. At the same time, given its long history, as of the time this chapter was written, Cisco PI still had some traditional management features not found in Cisco DNA Center. So think of PI as comprehensive to traditional device management, with Cisco DNA Center having many of those features, while focusing on future features like SDA support. NOTE Cisco hopes to continue to update Cisco DNA Center's traditional network management features to be equivalent compared to Cisco PI, to the point at which DNA Center could replace PI. In terms of intent and strategy, Cisco focuses their development of Cisco DNA Center features toward simplifying the work done by enterprises, with resulting reduced costs and much faster deployment of changes. Cisco DNA Center features help make initial installation easier, simplify the work to implement features that traditionally have challenging configuration, and use tools to help you notice issues more quickly. Some of the features unique to Cisco DNA
Center include ■ EasyQoS: Describes QoS, one of the most complicated features to configure manually, with just a few simple choices from Cisco DNA Center ■ Encrypted traffic analysis: Enables Cisco DNA to use algorithms to recognize security threats even in encrypted traffic Chapter 17: Cisco Software-Defined Access (SDA) ■ Device 360 and Client 360: Gives a comprehensive (360-degree) view of the health of the device ■ Network time travel: Shows past client performance in a timeline for comparison to current behavior ■ Path trace: Discovers the actual path packets would take from source to destination based on current forwarding tables Just to expound on one feature as an example, Cisco DNA Center's Path Trace feature goes far beyond a traditional management application. A typical network management app might show a map of the network and let you click through to find the configuration on each device, including ACLs. The path trace feature goes much further. The DNA user (from the GUI or the API) specifies a source and destination host and optionally transport protocol and ports. Then the path trace feature shows a map of the path through the network and shows which ACLs are in the path, and whether they would permit or deny the packet. All of Cisco Digital Network Architecture sets about to help customers reach some big goals: reduced costs, reduced risks, better security and compliance, faster deployment of services through automation and simplified processes, and the list goes on. Cisco DNA Center plays an important role, with all the functions available through its robust northbound API, and with its intent-based networking approach for SDA. Cisco DNA Center represents the future of network management for Cisco enterprises. Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 17-3 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 17-3 Chapter Review Tracking Review Element Review Date(s) Resource Use Review key topics Book, website Review key terms Book, website Answer DIKTA questions Book, PTP 403 17 404 CCNA 200-301 Official Cert Guide, Volume 2 Review All the Key Topics Table 17-4 Key Topics for Chapter 17 Key Topic Element Description Page Number List Definitions for overlay, underlay, and fabric 384 Figure 17-2 SDA overlay and underlay 386 List SDA fabric edge, fabric border, and fabric control node roles 387 List Attributes of the SDA underlay 389 List SDA VXLAN tunneling benefits 390 Figure 17-5 VXLAN encapsulation process with SDA 391 Figure 17-8 Registering SDA endpoint IDs (EIDs) with the map server 393 Figure 17-14 DNA Center showing controlling the fabric to implement group- 398 based security List DNA Center features that go beyond traditional network management 400 List Features unique to DNA Center 402 Key Terms You Should Know Software-Defined Access, overlay, underlay, fabric, DNA Center, fabric edge node, VXLAN, LISP, scalable group tag (SGT), Cisco Prime Infrastructure (PI) This page intentionally left blank CHAPTER 18 Understanding REST and JSON This chapter covers the following exam topics: 6.0 Automation and Programmability 6.5 Describe characteristics of REST-based APIs (CRUD, HTTP verbs, and data encoding) 6.7 Interpret JSON encoded data To automate and program networks, some automation software does several tasks. The software analyzes data in the form of variables, makes decisions based on that analysis, and then may take action to change the configuration of network devices or report facts about the state of the network. The different automation functions reside on different devices: the network engineer's device, a server, a controller, and the various network devices themselves. For these related automation processes to work well, all these software components need useful well-defined conventions to allow easy communication between software components. This chapter focuses on two conventions that allow automation software to communicate. The first major section discusses application programming interfaces (APIs), specifically APIs that follow a style called Representational State Transfer (REST). APIs of any kind create a way for software applications to communicate, while RESTful APIs (APIs that use REST conventions) follow a particular set of software rules. Many APIs used in network automation today use REST-based APIs. The second half of the chapter focuses on the conventions and standards for the data variables exchanged over APIs, with a focus on one: JavaScript Object Notation (JSON). If REST provides one standard method of how two automation programs should communicate over a network, JSON then defines how to communicate the variables used by a program: the variable names, their values, and the data structures of those variables. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 18-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions REST-based APIs 1–3 Data Models and JSON 4–6 1. Which of the following are required attributes of a REST-based API? (Choose two answers.) a. Uses HTTP b. Objects noted as to whether they can be cached c. Classful operation d. Client/server architecture 2. Which answers list a matching software development CRUD action to an HTTP verb that performs that action? (Choose two answers.) a. CRUD create and HTTP PATCH b. CRUD update and HTTP PATCH c. CRUD delete and HTTP PUT d. CRUD read and HTTP GET 3. Examine the following URI that works with a Cisco DNA Controller: Which part of the URI, per the API documentation, is considered to identify the resource but not any parameters? a. https:// b. dnac.example.com c. dna/intent/api/v1/network-device-d. managementIpAddress=10.0.22.74 4. Which of the following data serialization and data modeling languages would be most likely to be used in a response from a REST-based server API used for networking applications? (Choose two answers.) a. JSON b. YAML c. JavaScript d. XML 5. Which answers correctly describe the format of the JSON text below? (Choose two answers.) { "myvariable": [1,2,3] } a. One JSON object that has one key:value pair b. One JSON object that has three key:value pairs c. A JSON object whose value is a second JSON object d. A JSON object whose value is a JSON array 408 CCNA 200-301 Official Cert Guide, Volume 2 6. Which answers refer to JSON values rather than JSON keys as found in the sample JSON data? (Choose two answers.) ("response": { "type": "Cisco Catalyst 9300 Switch", "family": "Switches and Hubs", "role": "ACCESS", "managementIpAddress": "10.10.22.66" } } a. "response" b. "type" c. "ACCESS" d. The entire gray area Foundation Topics REST-Based APIs Applications use application programming interfaces (APIs) to communicate. To do so, one program can learn the variables and data structures used by another program, making logic choices based on those values, changing the values of those variables, creating new variables, and deleting variables. APIs allow programs running on different computers to work cooperatively, exchanging data to achieve some goal. In an API software world, some applications create an API, with many other applications using (consuming) the API. Software developers add APIs to their software so other application software can make use of the first application's features. When writing an application, the developer will write some code, but often the developer may do a lot of work by looking for APIs that can provide the data and functions, reducing the amount of new code that must be written. As a result, much of modern software development centers on understanding and learning new APIs, along with the available libraries (prebuilt software that can be used to accomplish tasks rather than writing the equivalent from scratch). Several types of APIs exist, each with a different set of conventions to meet a different set of needs. The CCNA blueprint mentions one type of API—Representational State Transfer (REST)—because of its popularity as a type of API in networking automation applications. This first major section of the chapter takes a closer look at REST-based APIs. REST-Based (RESTful) APIs REST APIs follow a set of foundational rules about what makes a REST API and what does not. First, from a literal perspective, REST APIs include the six attributes defined a few decades Chapter 18: Understanding REST and JSON 409 back by its creator, Roy Fielding. (You can find a good summary at). Those six attributes are ■ Client/server architecture ■ Stateless operation ■ Clear statement of cacheable/un-cacheable ■ Uniform interface ■ Layered ■ Code-on-demand The first three of these attributes get at the heart of how a REST API works. You can more easily see those first three features at work with networking REST APIs, so the next few paragraphs give a little more explanation about those first three points. Client/Server Architecture Like many applications, REST applications use a client/server architectural model. First, an application developer creates a REST API, and that application, when executing, acts as a REST server. Any other application can make a REST API call (the REST client) by executing some code that causes a request to flow
from the client to the server. For instance, in Figure 18-1.1. The REST client on the left executes a REST API call, which generates a message sent to the REST server. 2. The REST server on the right has API code that considers the request and decides how to reply. 3. The REST server sends back the reply message with the appropriate data variables in the reply message. Verb 1 IP TCP URI HTTP IP TCP HTTP Return Code API Call Data 2 API REST Server REST Client Figure 18-1.3 Client/Server Operation with REST 18 410 CCNA 200-301 Official Cert Guide, Volume 2 NOTE Figure 18-1 shows the use of HTTP. While many REST APIs use HTTP, the use of HTTP is not a requirement for an API to be considered RESTful. Stateless Operation The stateless attribute of REST APIs means that REST does not record and use information about one API exchange for the purpose of how subsequent API exchanges are processed. In other words, each API request and reply does not use any other past history considered when processing the request. For comparison, the TCP protocol uses a stateful approach, whereas UDP uses stateless operation. A TCP connection requires the endpoints to initialize variables on each end, with those variables updating over time, and with those variables being used for subsequent TCP messages. For instance, TCP uses sequence numbers and acknowledgment numbers to manage the flow of data in a TCP connection. Cacheable (or Not) To appreciate what is meant by cacheable, consider what happens when you browse a website. When your browser loads a new web page, the page itself contains a variety of objects (text, images, videos, audio). Some objects seldom change, so it would be better to download the object once and not download it again; in that case, the server marks that object as cacheable. For instance, a logo or other image shown on many pages of a website would almost never change and would likely be cacheable. However, the product list returned in your most recent search of the website would not be cacheable because the server would want to update and supply a new list each time you request the page. REST APIs require that any resource requested via an API call have a clear method by which to mark the resource as cacheable or not. The goals remain the same: improve performance by retrieving resources less often (cacheable). Note that cacheable resources are marked with a timeframe so that the client knows when to ask for a new copy of the resource again. Background: Data and Variables To appreciate a few of the upcoming topics, it helps to have a basic idea about how programming languages use variables. Anyone who has done even a small amount of programming should have enough background, but for those who have not written programs before, this next topic gives you enough background about data and variables inside programs to understand the next topic. If you have some programming experience and already know about simple variables, list variables, and dictionary variables, then feel free to skip ahead to the section "REST APIs and HTTP." Simple Variables Applications all process data with the same general actions, starting with some kind of input. The program needs data to process, so the input process reads files, sends database queries to a database server, or makes API calls to retrieve data from another application's API. The goal: gather the data that the program needs to process to do its work.

Answers to the "Do I Know This Already?" quiz: 1 B, 2 D, 3 C, 4 A, 5 D, 6 C, 7 D Chapter 18: Understanding REST and JSON API Programs then process data by making comparisons, making decisions, creating new variables, and performing mathematical formulas to analyze the data. All that logic uses variables. For instance, a program might process data with the following logic: If the router's G0/0 interface has a configuration setting of switchport mode dynamic auto, then gather more data to ensure that interface currently operates as a trunk rather than as an access port. In programming, a variable is a name or label that has an assigned value. To get a general sense for programming variables, you can think of variables much like variables from algebra equations back in school. Example 18-1 shows some samples of variables of different types in a Python program (the Python language is the most popular language today for writing network automation applications). This program begins with a comment (the top three lines with triple single quotes) and then creates four variables, assigning them to different values, and prints a line of output: "The product is -12." Example 18-1 Simple Python Program That Shows a Product 18 "" Sample program to multiply two numbers and display the result "" x = 3 y = -4 z = 1.247 heading = "The product is " print(heading,x,y) The variables in Example 18-1 can be called simple variables because each variable name has a single value associated with it. Simple variables have one variable name and one associated value, so they have a simple structure. The values of simple variables can have a variety of formats, as shown in Example 18-1. The example includes variables that contain ■ Unsigned integers (x) ■ Signed integers (y) ■ Floating-point numbers (z) ■ Text (heading) List and Dictionary Variables While simple variables have many great uses, programs need variables with more complex data structures. In programming, a data structure defines a related set of variables and values. For instance, Python uses lists (variables) so that one variable name is assigned a value that is a list of values rather than a single value. You could imagine that a network automation program might want to have lists, such as a list of devices being managed, a list of interfaces on a device, or list of configuration settings on an interface. First, consider the variable named list1 in Example 18-2; note that the lines that begin with a # are comment lines. 412 CCNA 200-301 Official Cert Guide, Volume 2 Example 18-2 Sample List and Dictionary Variables in Python ■ Variable list1 is a list in Python (called an array in Java) list1 = ["g0/0", "g0/1", "g0/2"] ■ Variable dict1 is a dictionary (called an associative array in Java) dict1 = {"config_speed": "auto", "config_duplex": "auto", "config_ip": "10.1.1.1"} Even if you have never seen Python code before, you can guess at some of the meaning of the list1 variable. The code assigns variable list1 to a value that itself is a list of three text strings. Note that the list could include text, unsigned integers, signed integers, and so on. Figure 18-2 shows the data structure behind variable list1 in Example 18-2. The variable is assigned to the list, with the list having three list elements. elements list1 g0/0 g0/1 g0/2 Figure 18-2 The List Data Structure in Python Python supports a similar data structure called a dictionary. If you think of the contents of a dictionary for the English language, that dictionary lists a series of paired items: a term and a matching definition. With programming languages like Python, the dictionary data structure lists paired items as well: keys (like terms) and values (like definitions). Figure 18-3 shows the structure of that dictionary value matching the dict1 variable at the bottom of Example 18-2. Note that each key and its value is called a key:value pair. Key:Value Pairs dict1 config_speed auto config_duplex auto config_ip 18.3 10.1.1.1 Dictionary Data Structures in Python Data structures can get more complex. Additionally, the data structures can be nested. For instance, a single variable's value could be a list, with each list element being a dictionary, with the values in some key:value pairs being other lists, and so on. For now, be aware of the fact that programs use simple variables but also use list and dictionary variables to make it easier to perform different kinds of logic. Chapter 18: Understanding REST and JSON API REST APIs and HTTP APIs exist to allow two programs to exchange data. Some APIs may be designed as an interface between programs running on the same computer, so the communication between programs happens within a single operating system. Many APIs need to be available to programs that run on other computers, so the API must define the type of networking protocols supported by the API—and many REST-based APIs use the HTTP protocol. The creators of REST-based APIs often choose HTTP because HTTP's logic matches some of the concepts defined more generally for REST APIs. HTTP uses the same principles as REST: it operates with a client/server model; it uses a stateless operational model; and it includes headers that clearly mark objects as cacheable or not cacheable. It also includes verbs—words that dictate the desired action for a pair HTTP Request and Reply—which matches how applications like to work. This section breaks down the fundamentals of some programming terminology, how that matches HTTP verbs, and how REST APIs make use of Uniform Resource Identifiers (URIs) to specify the data desired from a RESTful API call. Software CRUD Actions and HTTP Verbs The software industry uses a memorable acronym—CRUD—for the four primary actions performed by an application. Those actions are Create: Allows the client to create some new instances of variables and data structures at the server and initialize their values as kept at the server Read: Allows the client to retrieve (read) the current value of variables that exist at the server, storing a copy of the variables, structures, and values at the client Update: Allows the client to change (update) the value of variables that exist at the server Delete: Allows the client to delete from the server different instances of data variables For instance, if using the northbound REST API of a DNA controller, as discussed in Chapter 17, "Cisco Software-Defined Access (SDA)," you might want to create something new, like a new security policy. From a programming perspective, the security policy exists as a related set of configuration settings on the DNA controller, internally represented by variables. To do that, a REST client application would use a create action, using the DNA Center RESTful API, that created variables on the DNA Controller via the DNA Center REST API. The concept of creating new configuration at the controller is performed via the API using a create action per the CRUD generic acronym. Other examples of CRUD actions include a check of the status of that new configuration (a read action), an update to change some specific setting in the new configuration (an update action), or an action to remove the security policy definition completely (a delete action). HTTP uses verbs that mirror CRUD actions. HTTP defines the concept of an HTTP request and reply, with the client sending a request and with the server answering back with a reply. Each request/reply lists an action verb in the HTTP request header, which defines the HTTP action. The HTTP messages also include a URI, which identifies the resource being manipulated for this request. As always, the HTTP message is carried in IP and TCP, with headers and data, as represented in Figure 18-4. 18.414 CCNA 200-301 Official Cert Guide, Volume 2 HTTP IP TCP Request Header Verb Figure 18-4 URI Other Headers Data Some API Parameters HTTP Verb and URI in an HTTP Request Header To get some perspective about HTTP, ignore REST for a moment. Whenever you open a web browser and click a link, your browser generates an HTTP GET request message similar to Figure 18-4 in structure. The message includes an HTTP header with the GET verb and the URI. The resources returned in the reply are the components of a web page, like text files, image files, and video files. HTTP works well with REST in part because HTTP has verbs that match the common program actions in the CRUD paradigm. Table 18-2 Comparing CRUD Actions to REST Verbs Action CRUD Term (HTTP Verb) Create Create POST Read (retrieve) variable names, structures, and values Read GET Update or replace values of some variable Update PATCH, PUT Delete some variables and data structures Delete DELETE NOTE While Table 18-2 lists HTTP POST as a creation action and HTTP PATCH and PUT as CRUD update actions, all three of these HTTP verbs might be used both for create and for update actions in some cases. Using URIs with HTTP to Specify the Resource In addition to using HTTP verbs to perform the CRUD functions for an application, REST uses URIs to identify what resource the HTTP request acts on. For REST APIs, the resource can be any one of the many resources defined by the API. Each resource contains a set of related variables, defined by the API and identified by a URI. For instance, imagine a user creates a REST-based API. When she does so, she creates a set of resources that she wants to make available via the API, and she also assigns a unique URI to each resource. In other words, the API creator creates a URI and a matching set of variables, and defines the actions that can be performed against those variables (read, update, and so on). The API creator also creates API documentation that lists the resources and the URI that identifies each resource, among other details. The programmer for a REST client application can read the API documentation, build a REST API request, and ask for the specific resource, as shown in the example in Figure 18-5. Chapter 18: Understanding REST
and JSON HTTP GET URI = URI3 URI 1 Resource (Variables) URI 2 Resource (Variables) URI 3 . . . Resource (Variables) . . . URI N Resource (Variables) 415 REST Server Figure 18-5 One URI for Each API Resource—Conceptual View Figure 18-5 shows the URIs as generic values; however, today's network engineers need to be able to read API documentation, see URIs in that documentation, and understand the meaning of each part of the URI. Figure 18-6 shows a URI specific to the Cisco DNA Center northbound REST API as an example of some of the components of the URI. Hostname/Address HTTPS://dnac.example.com/dna/intent/api/v1/network-device Protocol Figure 18-6 Path (Resource) URI Structure for REST GET Request The figure shows these important values and concepts: HTTPS: The letters before the // identify the protocol used—in this case, HTTP Secure (which uses HTTP with SSL encryption). Hostname or IP Address: This value sits between the // and first /, and identifies the host; if using a hostname, the REST client must perform name resolution to learn the IP address of the REST server. Path (Resource): This value sits after the first / and finishes either at the end of the URI or before any additional fields (like a parameter query field). HTTP calls this field the path, but for use with REST, the field uniquely identifies the resource as defined by the API. To drive home the connection between the API, URI, and resource part of the API, it can be helpful to just do a general tour of the API documentation for any REST-based API. For instance, when Cisco created DNA Center, it created the REST-based northbound interface and chose one URI as shown in Figure 18-6. Figure 18-7 shows a copy of the doc page for that particular resource for comparison. Go to and search for "Cisco DNA Center API documentation." Continue to search for yourself to see more examples of the resources defined by the Cisco DNA Center API. 18.416 CCNA 200-301 Official Cert Guide, Volume 2 Figure 18-7 DNA Center API Doc Page for the Network Device (List) Resource Many of the HTTP request messages need to pass information to the REST server beyond the API. Some of that data can be passed in header fields—for instance, REST APIs use HTTP header fields to encode much of the authentication information for REST calls. Additionally, parameters related to a REST call can be passed as parameters as part of the URI itself. For instance, the URI in Figure 18-6 asks the Cisco DNA Center for a list of all known devices, with Cisco DNA Center returning a dictionary of values for each device. You might instead want that dictionary of values for only a single device. The Cisco DNA Center API allows for just that by tacking on the following to the end of the URI shown in Figure 18-6. ?managementIpAddress=10.10.22.66&macAddress=f8:7b:20:67:62:80 Figure 18-8 summarizes the major components of the URIs commonly used with a REST API, with the resource and parameter parts of the URI identifying specifically what the API should supply to the REST client. Hostname/Address HTTPS://dnac.example.com/dna/intent/api/v1/network-device?parm1=10.1.1.1. . . Protocol Figure 18-8 Path (Resource) Query (Parameters) Example Components of a URI Used in a REST API Call Chapter 18: Understanding REST and JSON 417 Example of REST API Call to DNA Center To pull some of the REST API concepts together, the next few pages work through a few sample API calls using a software application called an API development environment tool. For a bit of development perspective, when working to automate some part of your network operation tasks, you would eventually use a program that made API calls. However, early in the process of developing an application, you might first focus on the data available from the API and ignore all the programming details at first. API development environments let you focus on the API calls. Later, that same tool can typically generate correct code that you can copy into your program to make the API calls. The examples in this section use an app named Postman. Postman can be downloaded for free (www.postman.com) and used as shown in this section. Note that Cisco DevNet makes extensive use of Postman in its many labs and examples. The first example shows a screenshot of a part of the Postman app after it sends a REST client GET request to a DNA Center REST API (see Figure 18-9). In particular, look for the following: ■ The URI, near the top, lists a hostname of sandboxnac2.cisco.com, which is an always-on DNA Center instance supplied by Cisco's DevNet site (which you can use). ■ The resource part of the URI shows the same resource listed earlier in Figure 18-6, asking for a list of devices. ■ The bottom center of the window shows the data returned by the DNA Center REST HTTP GET response. ■ At the middle right, it lists the GET response's status code of 200, meaning "OK." Figure 18-9 URI Structure for REST GET Request 18.418 CCNA 200-301 Official Cert Guide, Volume 2 Take a moment to look through the data at the bottom of the Postman window in Figure 18-9. The text follows a data modeling format called JavaScript Object Notation (JSON), which is one of the main topics for the remainder of the chapter. However, armed with just a knowledge of routers, you can find a few facts that look familiar. To help you see the text, Example 18-3 shows an edited (shortened to reduce length) view of some of the JSON output in that window, just so you can see the format and some of the data returned in this single API call. Example 18-3 JSON Output from a REST API Call [{"type": "Cisco Catalyst 9300 Switch", "family": "Switches and Hubs"}, {"role": "ACCESS", "macAddress": "f8:7b:20:67:62:80", "hostname": "cat_9k_1", "serialNumber": "FCW2136LOAK", "softwareVersion": "16.6.1", "upTime": "17 days, 22:51:04.26", "interfaceCount": "41", "lineCardCount": "2", "managementIpAddress": "10.10.22.66", "series": "Cisco Catalyst 9300 Series Switches", "softwareType": "IOS-XE" } } API development tools like Postman help you work out the particulars of each API call, save the details, and share with other engineers and developers. Eventually, you will be ready to make the API call from a program. With a simple click from the Postman UI, Postman supplies the code to copy/paste into your program so that it returns all the output shown in the center/bottom of the window back as a variable to your program. By now, you have a good foundational knowledge of the mechanics of REST APIs. By learning some skills, and using the API documentation for any REST API, you could now experiment with and try to make REST API calls. For many of those, the data will return to you as text, often in JSON format, so the second half of the chapter examines the meaning of that text. Data Serialization and JSON In your journey to become a modern network engineer with network automation skills, you will learn to understand several data serialization languages. Each data serialization language provides methods of using text to describe variables, with a goal of being able to send that text over a network or to store that text in a file. Data serialization languages give us a way to represent variables with text rather than in the internal representation used by any particular programming language. Chapter 18: Understanding REST and JSON 418 Each data serialization language enables API servers to return data so that the API client can replicate the same variable names as well as data structures as found on the API server. To describe the data structures, the data serialization languages include special characters and conventions that communicate ideas about list variables, dictionary variables, and other more complex data structures. This second major section of the chapter examines the concept of a data serialization language, with a focus on the one data modeling language as mentioned in the current CCNA blueprint: JavaScript Object Notation (JSON). The Need for a Data Model with APIs This section shows some ideas of how to move variables in a program on a server to a client program. First, Figure 18-10 and surrounding text show a nonworking example as a way to identify some of the challenges with copying variable values from one device to another. Then Figure 18-11 and its related text show how to use a data serialization language to solve the problems shown around Figure 18-10. 18.2 API 3.1 Variables: Internal Variables: Internal REST Client (Python) REST Server (JAVA) Figure 18-10 Broken Concept: Exchanging Internal Representations of Variables First, for the nonworking example, consider the flow and numbered steps in Figure 18-10. A REST client sits on the left. The REST client asks for a resource, and the server needs to reply. In REST, a resource is a set of variables as defined by the API, so the REST server needs to return a set of variables to the REST client on the left. The steps in the figure run as follows: 1. The REST server (a JAVA application) takes a copy of the stored variables in RAM (step 1) in response to the REST request. 2. The REST API code creates the REST reply and sends it over the network, placing an exact replica of what the REST server had in RAM to represent the variables in that resource. 420 CCNA 200-301 Official Cert Guide, Volume 2 3. The REST client (a Python application) receives the REST reply message, storing the exact same bits and bytes into its RAM, in an attempt to have a copy of the variables, data, and data structures on the server. The process shown in Figure 18-10 does not work (and is not attempted) because the REST client programs may not store variables in the same ways. First, programs written in different languages use different conventions to store their variables internally because there is no standard for internal variable storage across languages. In fact, programs written in the same language but with different versions of that language may not store all their variables with the
same internal conventions. To overcome these issues, applications need a standard method to represent variables for transmission and storage of those variables outside the program. Data serialization languages provide that function. Figure 18-11 shows the correct process flow in comparison to Figure 18-10 with the data serialization process included: 1. The server collects the internally represented data and gives it to the API code. 2. The API converts the internal representation to a data model representing those variables (with JSON shown in the figure). 3. The server sends the data model in JSON format via messages across the network. 4. The REST client takes the received data and converts the JSON-formatted data into variables in the native format of the client application. 3 JSON String JSON String 4.2 JSON Converter API 1 Variables: Internal REST Client Figure 18-11 Variables: Internal REST Server Correct Concept: Exchanging Internal Representations of Variables At the end of the process, the REST client application now has equivalent variables to the ones it requested from the server in the API call. Note that the final step—to convert from the data serialization language to the native format—can be as little as a single line of code! Chapter 18: Understanding REST and JSON 421 Finally, note that while data serialization languages like JSON enable applications to exchange variables over a network, applications can also store data in JSON format. Data Serialization Languages You will hear about and eventually use several data serialization and data modeling languages the more you learn about network automation. While the current CCNA blueprint mentions only JSON, learning a few facts about some of the alternatives can be helpful to add a little context to your new knowledge of JSON. These different data serialization languages exist to meet different needs that have arisen over the years. This next short section highlights four such languages. NOTE The terms data serialization language and data modeling language should be considered equivalent for the purposes of this section. JSON JavaScript Object Notation attempts to strike a balance between human and machine readability. Armed with a few JSON rules, most humans can read JSON data, much past simply guessing at what it means, and confidently interpret the data structures defined by the JSON data. At the same time, JSON data makes it easy for programs to convert JSON text into variables, making it very useful for data exchange between applications using APIs. You can find the details of JSON in IETF RFC 8259 and in a number of sites found with Internet searches, including www.json.org. XML Back in the 1990s, when web browsers and the World Wide Web (WWW) were first created, web pages primarily used Hypertext Markup Language (HTML) to define web pages. As a markup language, HTML defined how to add the text or a web page to a file and then add "markup"—additional text to denote formatting details for the text that should be displayed. For instance, the markup included codes for headings, font sizes, colors, hyperlinks, and so on. The eXtensible Markup Language (XML) came later to make some improvements for earlier markup languages. In particular, over time web pages became more and more dynamic, and to make the pages dynamic, the files needed to store variables whose values could be changed and replaced over time by the web server. To define variables to be substituted into a web page, the world needed a markup language that could define data variables. XML defines a markup language that has many features to define variables, values, and data structures. Over time, XML has grown beyond its original use as a markup language. XML's features also make it a useful general data serialization language, and it is used as such today. Comparing XML to JSON, both attempt to be human readable, but with XML being a little more challenging to read for the average person. For instance, like HTML, XML uses beginning and ending tags for each variable, as seen in Example 18-4. In the highlighted line in the example, the and tags denote a variable name, with the value sitting between the tags. 18.422 CCNA 200-301 Official Cert Guide, Volume 2 Example 18-4 JSON Output from a REST API Call Switches and Hubs cat_9k_1 41 2 f8:7b:20:67:62:80 10.22.66 ACCESS FCW2136LOAK Cisco Catalyst 9300 Series Switches IOS-XE 16.6.1 Cisco Catalyst 9300 Switch 17 days, 22:51:04.26 YAML Ain't Markup Language (YAML) has a clever recursive name, but the name does tell us something. YAML does not attempt to define markup details (while XML does). Instead, YAML focuses on the data model (structure) details. YAML also strives to be clean and simple: the data serialization/modeling languages listed here, YAML is easily the easiest to read for anyone new to data models. Ansible, one of the topics in Chapter 19, "Understanding Ansible, Puppet, and Chef," makes extensive use of YAML files. Example 18-5 shows a brief sample. And to make the point about readability, even if you have no idea what Ansible does, you can guess at some of the functions just reading the file. (Note that YAML denotes variables in double curly brackets: { }) Example 18-5 YAML File Used by Ansible –4 This comment line is a place to document this Playbook - name: Get IOS Facts hosts: mylab vars: cli: host: "if (ansible_host)"; username: "{{ username }}" password: "{{ password }}" tasks: - ios_facts: gather_subset: all provider: "{{ cli }}" Chapter 18: Understanding REST and JSON 423 Summary of Data Serialization As an easy reference, Table 18-3 summarizes the data serialization languages mentioned in this section, along with some key facts. Table 18-3 Comparing Data Modeling Languages Acronym Name Origin/Definition Central Purpose Common Use JSON JavaScript Object Notation (JS) language; RFC 8259 General data modeling and serialization REST APIs XML eXtensible World Wide Web Data-focused text REST APIs, Markup Language Consortium (W3C.org) markup that allows Web pages name data modeling YAML Ain't Markup Language General data modeling Ansible Interpreting JSON Cisco includes one exam topic in the current CCNA 200-301 blueprint that mentions JSON: 6.7 Interpret JSON encoded data You can think of that skill and task with two major branches. First, even ignoring the syntax and special characters, anyone who knows the topic can probably make intelligent guesses about the meaning of many of the key:value pairs. For example, without knowing anything about JSON syntax, you could probably determine that your prior knowledge of Cisco routers and switches that the JSON in Example 18-6 lists two devices (maybe their hostnames) and a list of interfaces on each device. Example 18-6 Simple JSON That Lists a Router's Interfaces { "R1": [{"GigabitEthernet0/0": "GigabitEthernet0/1"}, {"GigabitEthernet0/1": "GigabitEthernet0/2"}, "R2": [{"GigabitEthernet1/0": "GigabitEthernet1/1"}, {"GigabitEthernet0/3/0"}] Honestly, you probably already know everything needed to do this kind of intelligent guessing. However, to perform the second type of task, where you analyze the JSON data to find the data structures, including objects, lists, and key:value pairs, you need to know a bit more about JSON syntax. This final topic in the chapter gives you the basic rules, with some advice on how to break down JSON data. Interpreting JSON Key:Value Pairs First, consider these rules about key:value pairs in JSON, which you can think of as individual variable names and their values: ■ Key:Value Pair: Each and every colon identifies one key:value pair, with the key before the colon and the value after the colon. ■ Key: Text, inside double quotes, before the colon, used as the name that references a value. 18.424 CCNA 200-301 Official Cert Guide, Volume 2 ■ 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 special value (more details later) ■ Object: A special value (more details later) Multiple Pairs: When listing multiple key:value pairs, separate the pairs with a comma at the end of each pair (except the last pair). To work through some of these rules, consider Example 18-7's JSON data, focusing on the three key:value pairs. The text after the example will analyze the example. Example 18-7 One JSON Object (Dictionary) with Three Key:Value Pairs { "1stbest": "Messi", "2ndbest": "Ronald", "3rdbest": "Pele" } As an approach, just find each colon, and look for the quoted string just before each colon. Those are the keys ("1stbest", "2ndbest", and "3rdbest"). Then look to the right of each colon to find their matching values. You can know all three values because JSON lists the values within double quotes. As for other special characters, note the commas and the curly brackets. The first two key:value pairs end with a comma, meaning that another key:value pair should follow. The curly brackets that begin and end the JSON data denote a single JSON object (one pair of curly brackets, so one object). JSON files, and JSON data exchanged over an API, exist first as a JSON object, with an opening (left) and closing (right) curly bracket as shown. Interpreting JSON Objects and Arrays To communicate data structures beyond a key:value pair with a simple value, JSON uses JSON objects and JSON arrays. Objects can be somewhat flexible, but in most uses, they act like a dictionary. Arrays list a series of values. NOTE Python, the most common language to use for network automation, converts JSON objects to Python dictionaries, and JSON arrays to Python lists. For general conversation, many people refer to the JSON structures as dictionaries and lists rather than as objects and arrays. To begin, consider this set of rules about how to interpret the syntax for JSON objects and arrays: ■ [] - Object: A series of key:value pairs enclosed in a matched pair of curly brackets, with
an opening left curly bracket and its matching right curly bracket. ■ |] - Array: A series of values (not key:value pairs) enclosed in a matched pair of square brackets, with an opening left square bracket and its matching right square bracket. Chapter 18: Understanding REST and JSON ■ Key:value pairs inside objects: All key:value pairs inside an object conform to the earlier rules for key:value pairs. ■ Values inside arrays: All values conform to the earlier rules for formatting values (for example, double quotes around text, no quotes around numbers). 425 Example 18-8 shows a single array in JSON format. Notice the JSON data begins with a [and then lists three text values (the values could have been a mix of values). It then ends with a]. Example 18-8 A JSON Snippet Showing a Single JSON Array ["Messi", "Ronald", "Dybala"] While Example 18-8 shows only the array itself, JSON arrays can be used as a value in any key:value pair. Figure 18-12 does just that, shown in a graphic to allow easier highlighting of the arrays and object. The JSON text in the figure includes two arrays (lists) as values (each found just after a colon, indicating they are values). JSON Object with Two Key:Value Pairs { "favorite_players": ["Messi", "Ronald", "Dybala"], "favorite_teams": ["Barcelona", "Juventus", "Dortmund"] } JSON Array "favorite_players" with 3 Values JSON Array "favorite_teams" with 3 Values Figure 18-12 Values Accurate/Complete JSON Data with One Object, Two Keys, Two JSON List Now think about the entire structure of the JSON data in Figure 18-12. It has a matched pair of curly brackets to begin and end the text, encapsulating one object. That object contains two colons, so there are two key:value pairs inside the object. When you think about the broader structure, as depicted in Figure 18-13, this JSON file has one JSON object, itself with two key:value pairs. (Note that Figure 18-13 does NOT show correct JSON syntax for the lists; it instead is intended to make sure you see the structure of the one object and its two key:value pairs.) 18.426 CCNA 200-301 Official Cert Guide, Volume 2 JSON Object { "favorite_players": [. . .], "favorite_teams": [. . .] } Key:Value Key:Value Figure 18-13 Structural Representation of Figure 18-13's Primary Object and Two Key:Value Pairs To drive home the idea of how to find JSON objects, consider the example shown in Figure 18-14. This figure shows correct JSON syntax. It has the following: ■ There is one object for the entire set because it begins and ends with curly braces. ■ The outer object has two keys (Wendells_favorites and interface_config). ■ The value of each key:value pair is another object (each with curly braces and three key:value pairs). JSON Object with Two Key:Value Pairs { "Wendells_favorites": { "player": "Messi", "team": "Barcelona", "league": "La Liga" }, "interface_config": { "ip_address": "10.1.1.1", "ip_mask": "255.255.255.0", "speed": "1000" } } Key:Value Pair "Wendells_favorites" Value: An Object Key:Value Pair "interface_config" Value: An Object Figure 18-14 A JSON Object, with Two Key:Value Pairs, Each Value Another Object The JSON example in Figure 18-14 shows how JSON can nest objects and arrays; that is, JSON puts one object or array inside another. Much of the JSON output you will see as you learn more and more about network automation will include JSON data with nested arrays and objects. Minified and Beautified JSON So far, all the JSON examples show lots of empty space. JSON allows for whitespace, or not, depending on your needs. For humans, reading JSON can be a lot easier with the text organized with space and aligned. For instance, having the matched opening and closing brackets sit at the same left-offset makes it much easier to find which brackets go with which. When stored in a file or sent in a network, JSON does not use whitespace. For instance, earlier in this section, Example 18-7 showed one JSON object with three key:value pairs, with Chapter 18: Understanding REST and JSON 427 whitespace, taking five lines. However, stored in a file, or sent over a network, the JSON would look like the following: {"1stbest": "Messi", "2ndbest": "Ronald", "3rdbest": "Pele"} Most of the tools you might use when working with JSON will let you toggle from a pretty format (good for humans) to a raw format (good for computers). You might see the pretty version literally called pretty, or beautified, or spaced, while the version with no extra whitespace might be called minified or raw. Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 18-4 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Table 18-4 Chapter Review Tracking Review Element Review Date(s) Resource Used Review key topics Book, website Review key terms Book, website Answer DIKTA questions Book, PTP Review memory tables Website Practice Editing JSON Website Review All the Key Topics Table 18-5 Key Topics for Chapter 18 Key Topic Element Description Page Number List Attributes of REST APIs 409 List The meaning of the CRUD acronym 413 Table 18-2 a comparison of CRUD actions and HTTP verbs 414 Figure 18-8 Components of a URI 416 Figure 18-11 The process of sending JSON data over a REST API 420 Table 18-3 a comparison of JSON, XML, and YAML 423 List JSON rules related to key:value pairs 423 List JSON rules for arrays and objects 424 Key Terms You Should Know REpresentational State Transfer (REST), REST API, stateless, cacheable, CRUD, list variable, dictionary variable, URI path (resource), URI query (parameters), key:value pair, data serialization language, JSON (JavaScript Object Notation), XML (eXtensible Markup Language), YAML (YAML Ain't Markup Language), JSON object, JSON array 18 CHAPTER 19 Understanding Ansible, Puppet, and Chef This chapter covers the following exam topics: 6.0 Automation and Programmability 6.6 Recognize the capabilities of configuration mechanisms Puppet, Chef, and Ansible By now, you have seen how to use the IOS CLI to configure routers and switches. To configure using the CLI, you get into configuration mode, issue configuration commands (which change the running-config file), and eventually leave configuration mode. If you decide to keep those changes, you save the configuration to the startup-config file using the copy running-config startup-config command. Next time the router or switch boots, the device loads the startup-config file into RAM as the running-config. Simple enough. This chapter discusses tools for configuration management that replaces that per-device configuration process. To even imagine what these tools do first requires you to make a leap of imagination to the everyday world of a network engineer at a medium to large enterprise. In a real working network, managing the configuration of the many networking devices creates challenges. Those challenges can be addressed using that same old "use configuration mode on each device" process, plus with hard work, attention to detail, and good operational practices. However, that manual per-device process becomes more and more difficult for a variety of reasons, so at some point, enterprises turn to automated configuration management tools to provide better results. The first section of this chapter takes a generalized look at the issues of configuration management at scale along with some of the solutions to those problems. The second major section then details each of three configuration management tools—Ansible, Puppet, and Chef—to define some of the features and terms used with each. By the end of the chapter, you should be able to see some of the reasons why these automated configuration management tools have a role in modern networks and enough context to understand as you pick one to investigate for further reading. "Do I Know This Already?" Quiz Take the quiz (either here or use the PTP software) if you want to use the score to help you decide how much time to spend on this chapter. The letter answers are listed at the bottom of the page following the quiz. Appendix C, found both at the end of the book as well as on the companion website, includes both the answers and explanations. You can also find both answers and explanations in the PTP testing software. Table 19-1 "Do I Know This Already?" Foundation Topics Section-to-Question Mapping Foundation Topics Section Questions Device Configuration Challenges and Solutions 1–3 Ansible, Puppet, and Chef Basics 4, 5 1. Which answer best describes the meaning of the term configuration drift? a. Changes to a single device's configuration over time versus that single device's original configuration b. Larger and larger sections of unnecessary configuration in a device c. Changes to a single device's configuration over time versus other devices that have the same role d. Differences in device configuration versus a centralized backup copy 2. An enterprise moves away from manual configuration methods, making changes by editing centralized configuration files. Which answers list an issue solved by using a version control system with those centralized files? (Choose two answers.) a. The ability to find which engineer changed the central configuration file on a date/time b. The ability to find the details of what changed in the configuration file over time c. The ability to use a template with per-device variables to create configurations d. The ability to recognize configuration drift in a device and notify the staff 3. Configuration monitoring (also called configuration enforcement) by a configuration management tool generally solves which problem? a. Tracking the identity of individuals
who changed files, along with which files they changed b. Listing differences between a former and current configuration c. Testing a configuration change to determine whether it will be rejected or not when implemented d. Finding instances of configuration drift 4. Which of the following configuration management tools uses a push model to configure network devices? a. Ansible b. Puppet c. Chef d. None Use a push model 430 CCNA 200-301 Official Cert Guide, Volume 2 5. Which of the following answers list a correct combination of configuration management tool and the term used for one of its primary configuration files? (Choose two answers.) a. Ansible manifest b. Puppet manifest c. Chef recipe d. Ansible recipe Foundation Topics Device Configuration Challenges and Solutions Think about any production network. What defines the exact intended configuration of each device as it exists right now or the startup-config before any recent changes were made or the startup-config from last month? Could one engineer change the device configuration so that it drifts away from that ideal, with the rest of the staff not knowing? What process, if any, might discover the configuration drift? And even with changes agreed upon by all, how do you know who changed the configuration, when, and specifically what changed? Traditionally, CCNA teaches us how to configure one device using the configure terminal command to reach configuration mode, which changes the running-config file, and how to save that running-config file to the startup-config file. That manual process provides no means to answer any of the legitimate questions posed in the first paragraph; however, for many enterprises, those questions (and others) need answers, both consistent and accurate. Not every company reaches the size to want to do something more with configuration management. Companies with one network engineer might do well enough managing device configurations, especially if the network device configurations do not change often. However, as a company moves to multiple network engineers and grows the numbers of devices and types of devices, with higher rates of configuration change, manual configuration management has problems. This section begins by discussing a few of these kinds of configuration management issues so that you begin to understand why enterprises need more than good people and good practices to deal with device configuration. The rest of the section then details some of the features you can find in automated configuration management tools. Configuration Drift Consider the story of an enterprise of a size to need two network engineers, Alice and Bob. They both have experience and work well together. But the network configurations have grown beyond what any one person can know from memory, and with two network engineers, they may remember different details or even disagree on occasion. One night at 1 a.m., Bob gets a call about an issue. He gets into the network from his laptop and resolves the problem with a small configuration change to branch office router BR22. Alice, the senior engineer, gets a different 4 a.m. call about another issue and makes a change to branch office router BR33. Chapter 19: Understanding Ansible, Puppet, and Chef 431 The next day gets busy, and neither Alice nor Bob mentions the changes they made. They both follow procedures and document the changes in their change management system, which lists the details of every change. But they both get busy, the topic never comes up, and neither mentions the changes to each for months. The story shows how configuration drift can occur—an effect in which the configuration drifts away from the intended configuration over time. Alice and Bob probably agree to what a standard branch office router configuration ought to look like, but they both made an exception to that configuration to fix a problem, causing configuration drift. Figure 19-1 shows the basic idea, with those two branch routers now with slightly different configurations than the other branch routers. BR22 Unique SRXWHUV - running-config BR33 Unique Other Branches Consistent Figure 19-1 Configuration Drift in Branch Routers BR22 and BR33 Configuration drift becomes a much bigger problem if using only traditional manual configuration tools. For instance: ■ The on-device manual configuration process does not track change history, which lines changed, what changed on each line, what old configuration was removed, who changed the configuration, when each change was made. ■ External systems used by good systems management processes, like trouble ticketing and change management software, may record details. However, those sit outside the configuration and require analysis to figure out what changed. They also rely on humans to follow the operational processes consistently and correctly; otherwise, an engineer cannot find the entire history of changes to a configuration. ■ Referring to historical data in change management systems works poorly if a device has gone through multiple configuration changes over a period of time. Centralized Configuration Files and Version Control The manual per-device configuration model makes great sense for one person managing one device. With that model, the one network engineer can use the on-device startup-config as the intended ideal configuration. If a change is needed, the engineer gets into configuration mode and updates the running-config until happy with the change. Then the engineer saves a copy to startup-config as the now-current ideal config for the device. The per-device manual configuration model does not work as well for larger networks, with hundreds or even thousands of network devices, with multiple network engineers. For instance, if the team thinks of the startup-config of each device as the ideal configuration, if one team member changes the configuration (like Alice and Bob each did in the earlier 19.432 CCNA 200-301 Official Cert Guide, Volume 2 story), no records exist about the change. The config file does not show what changed, when it changed, or who changed it, and the process does not notify the entire team about the change. As a first step toward better configuration management, many medium to large enterprises store configurations in a central location. At first, storing files centrally may be a simple effort to keep backup copies of each device's configuration. They would, of course, place the files in a shared folder accessible to the entire network team, as shown in Figure 19-2. BR21.txt BR21 BR22.txt Alice BR23.txt BR22. . . Bob BR23 Chris Folder Figure 19-2 Copying Device Configurations to a Central Location Which configuration file is the single source of truth in this model? The configuration files still exist on each device, but now they also exist on a centralized server, and engineers could change the on-device configuration as well as the text files on the server. For instance, if the copy of BR21's configuration on the device differs from the file on the centralized server, which should be considered as correct, ideal, the truth about what the team intends for this device? In practice, companies take both approaches. In some cases, companies continue to use the on-device configuration files as the source of truth, with the centralized configuration files treated as backup copies in case the device fails and must be replaced. However, other enterprises make the transition to treat the files on the server as the single source of truth about each device's configuration. When using the centralized file as the source of truth, the engineers can take advantage of many configuration management tools and actually manage the configurations more easily and with more accuracy. For example, configuration management tools use version control software to track the changes to centralized configuration files, noting who changes a file, what lines and specific characters changed, when the change occurred, and so on. The tools also allow you to compare the differences between versions of the files over time, as shown in Figure 19-3. Answers to the "Do I Know This Already?" quiz: 1 C 2 A, B 3 D 4 A 5 B, C Chapter 19: Understanding Ansible, Puppet, and Chef 433 Lines with Removals Lines with Additions Figure 19-3 Showing File Differences in GitHub The figure shows a sample of a comparison between two versions of a configuration file. The upper two highlighted lines, with the minus signs, show the lines that were changed, while the two lower highlighted lines, with the plus signs, show the new versions of each line. Version control software solves many of the problems with the lack of change tracking within the devices themselves. Figure 19-3 shows output from a popular software-as-a-service site called GitHub.com (www.github.com). GitHub offers free and paid accounts, and it uses open-source software (Git) to perform the version control functions. Configuration Monitoring and Enforcement With a version control system and a convention of storing the configuration files in a central location, a network team can do a much better job of tracking changes and answering the who, what, and when of knowing what changed in every device's configuration. However, using that model then introduces other challenges—challenges that can be best solved by also using an automated configuration management tool. With this new model, engineers should make changes by editing the associated configuration files in the centralized repository. The configuration management tool can then be directed to copy or apply the configuration to the device, as shown in Figure 19-4. After that process completes, the central config file and the device's running-config (and startup-config) should be identical. Edits 1 BR21 2 Apply BR21 Ideal 3 copy run start Figure 19-4 Pushing Centralized Configuration to a Remote Device 19.434 CCNA 200-301 Official Cert Guide, Volume 2 Using the model shown in Figure 19-4
still has dangers. For instance, the network engineers should make changes by using the configuration management tools, but they still have the ability to log in to each device and make manual changes on each device. So, while the idea of using a configuration management tool with a centralized repository of config files sounds appealing, eventually someone will change the devices directly. Former correct configuration changes might be overwritten, and made incorrect, by future changes. In other words, eventually, some configuration drift can occur. Configuration management tools can monitor device configurations to discover when the device configuration differs from the intended ideal configuration, and then either reconfigure the device or notify the network engineering staff to make the change. This feature might be called configuration monitoring or configuration enforcement, particularly if the tool automatically changes the device configuration. Figure 19-5 shows the general idea behind configuration monitoring. The automated configuration management software asks for a copy of the device's running-config file, as shown in steps 1 and 2. At step 3, the config management software compares the ideal config file with the just-arrived running-config file to check whether they have any differences (configuration drift). Per the configuration of the tool, it either fixes the configuration or notifies the staff about the configuration drift. 1 show run runningconfig 2 Router BR21 Running-config (copy) 3 compare BR21 Ideal Config Management Figure 19-5 Configuration Monitoring Configuration Provisioning Configuration provisioning refers to how to provision or deploy changes to the configuration once made by changing files in the configuration management system. As one of the primary functions of a configuration management tool, you would likely see features like these: ■ The core function to implement configuration changes in one device after someone has edited the device's centralized configuration file ■ The ability to choose which subset of devices to configure: all devices, types with a given attribute (such as those of a particular role), or just one device, based on attributes and logic Chapter 19: Understanding Ansible, Puppet, and Chef 435 ■ The ability to determine if each change was accepted or rejected, and to use logic to react differently in each case depending on the result ■ For each change, the ability to revert to the original configuration if even one configuration command is rejected on a device ■ The ability to validate the change now (without actually making the change) to determine whether the change will work or not when attempted ■ The ability to check the configuration after the process completes to confirm that the configuration management tool's intended configuration does match the device's configuration ■ The ability to use logic to choose whether to save the running-config to startup-config or not ■ The ability to represent configuration files as templates and variables so that devices with similar roles can use the same template but with different values ■ The ability to store the logic steps in a file, scheduled to execute, so that the changes can be implemented by the automation tool without the engineer being present The list could go further, but these features outline some of the major features included in all of the configuration management tools discussed in this chapter. Most of the items in the list revolve around editing the central configuration file for a device. However, the tools have many more features, so you have more work to do to plan and implement how they work. The next few pages focus on giving a few more details about the last two items in the list. Configuration

Templates and Variables Think about the roles filled by networking devices in an enterprise. Focusing on routers for a moment, routers often connect to both the WAN and one or more LANs. You might have a small number of larger routers connected to the WAN at large sites, with enough power to handle larger packet rates. Smaller sites, like branch offices, might have small routers, maybe with a single WAN interface and a single LAN interface; however, you might have a large number of those small branch routers in the network. For any set of devices in the same role, the configurations are likely similar. For instance, a set of branch office routers might all have the exact same configuration for some IP services, like NTP or SNMP. If using OSPF interface configuration, routers in the same OSPF area and with identical interface IDs could have identical OSPF configuration. For instance, Example 19-1 shows a configuration excerpt from a branch router, with the unique parts of the configuration highlighted. All the unhighlighted portions could be the same on all the other branch office routers of the same model (with the same interface numbers). An enterprise might have dozens or hundreds of branch routers with nearly identical configuration. Example 19-1 Router BR1 Configuration, with Unique Values Highlighted hostname BR1 ! interface GigabitEthernet0/0 ip address 10.1.1.1 255.255.255.0 ip ospf 1 area 11 19 436 CCNA 200-301 Official Cert Guide, Volume 2 ! interface GigabitEthernet0/1 ! interface GigabitEthernet0/1/0 ip address 10.1.12.1 255.255.255.0 ip ospf 1 area 11 ! router ospf 1 router-id 1.1.1.1 Configuration management tools can separate the components of a configuration into the parts in common to all devices in that role (the template) versus the parts unique to any one device (the variables). Engineers can then edit the standard template file for a device role as a separate file than each device's variable file. The configuration management tool can then process the template and variables to create the ideal configuration file for each device, as shown in Figure 19-6, which shows the configuration files being built for branch routers BR21, BR22, and BR23. 1 BR21 Ideal 1 BR21 Variables 2 BR22 Ideal 2 BR22 Variables 3 BR23 Ideal Figure 19-6 Template: Branch Router 3 BR23 Variables Concept: Configuration Templates and Variables To give a little more insight, Example 19-2 shows a template file that could be used by Ansible for the configuration shown in Example 19-1. Each tool specifies what language to use for each type of file, with Ansible using the Jinja2 language for templates. The template mimics the configuration in Example 19-1, except for placing variable names inside sets of double curly brackets. Example 19-2 Jinja2 Template with Variables Based on Example 19-1 hostname {{hostname}} ! interface GigabitEthernet0/0 ip address {{address1}} {{mask1}} ip ospf {{OSPF_PID}} area {{area}} ! interface GigabitEthernet0/1 ! interface GigabitEthernet0/1/0 Chapter 19. Understanding Ansible, Puppet, and Chef 437 ip address {{address2}} {{mask2}} ip ospf {{OSPF_PID}} area {{area}} ! router ospf {{OSPF_PID}} router-id {{RID}} To supply the values for a device, Ansible calls for defining variable files using YAML, as shown in Example 19-3. The file shows the syntax for defining the variables shown in the complete configuration in Example 19-1, but now defined as variables. Example 19-3 YAML Variables File Based on Example 19-2 -hostname: BR1 address: 10.1.1.1 mask1: 255.255.255.0 address2: 10.1.12.1 mask2: 255.255.255.0 RID: 1.1.1.1 area: '11' OSPF_PID: '1' The configuration management system processes a template plus all related variables to produce the intended configuration for a device. For instance, the engineer would create and edit one template file that looks like Example 19-2 and then create and edit one variable file like Example 19-3 for each branch office router. Ansible would process the files to create complete configuration files like the text shown in Example 19-1. It might seem like extra work to separate configurations into a template and variables, but using templates has some big advantages. In particular: ■ Templates increase the focus on having a standard configuration for each device role, helping to avoid snowflakes (uniquely configured devices). ■ New devices with an existing role can be deployed easily by simply copying an existing per-device variable file and changing the values. ■ Templates allow for easier troubleshooting because troubleshooting issues with one standard template should find and fix issues with all devices that use the same template. ■ Tracking the file versions for the template versus the variables files allows for easier troubleshooting as well. Issues with a device can be investigated to find changes in the device's settings separately from the standard configuration template. Files That Control Configuration Automation Configuration management tools also provide different methods to define logic and processes that tell the tool what changes to make, to which devices, and when. For instance, an engineer could direct a tool to make changes during a weekend change window. That same logic could specify a subset of the devices. It could also detail steps to verify the change before and after the change is attempted, and how to notify the engineers if an issue occurs. Interestingly, you can do a lot of the logic without knowing how to program. Each tool uses a language of some kind that engineers use to define the action steps, often a language 19 438 CCNA 200-301 Official Cert Guide, Volume 2 defined by that company (a domain-specific language). But they make the languages to be straightforward, and they are generally much easier to learn than programming languages. Configuration management tools also enable you to extend the action steps beyond what can be done in the toolset by using a general programming language. Figure 19-7 summarizes the files you could see in any of the configuration management tools. Subset Hosts R1 Actions + Program R2 Templates SW1 Variables SW2 Config Management Figure 19-7 Important Files Used by Configuration Management Tools Ansible, Puppet, and Chef Basics This chapter focuses on one exam topic that asks about the capabilities of three configuration management tools: Ansible, Puppet, and Chef. The first major section of the chapter describes the capabilities of all three (and other) configuration management tools. This second major section examines a few of the features of each tool, focusing on terminology and major capabilities. Ansible, Puppet, and Chef are software packages. You can purchase each tool, with variations on which package. However, they all also have different free options that allow you to download and learn about the tools, although you might need to run a Linux guest because some of the tools do not run in a Windows OS. As for the names, most people use the words Ansible, Puppet, and Chef to refer to the companies as well as their primary configuration management products. All three emerged as part of the transition from hardware-based servers to virtualized servers, which greatly increased the number of servers and created the need for software automation to create, configure, and remove VMs. All three produce one or more configuration management software products that have become synonymous with their companies in many ways. (This chapter follows that convention, for the most part ignoring exact product names, and referring to products and software simply as Ansible, Puppet, and Chef.) Next, on to some details about each.

Ansible To use Ansible (www.ansible.com), you need to install Ansible on some computer: Mac, Linux, or a Linux VM on a Windows host. You can use the free open-source version or use the paid Ansible Tower server version. Chapter 19: Understanding Ansible, Puppet, and Chef 439 Once it is installed, you create several text files, such as the following: ■ Playbooks: These files provide actions and logic about what Ansible should do. ■ Inventory: These files provide device hostnames along with information about each device, like device roles, so Ansible can perform functions for subsets of the inventory ■ Templates: Using Jinja2 language, the templates represent a device's configuration but with variables (see Example 19-2). ■ Variables: Using YAML, a file can list variables that Ansible will substitute into templates (see Example 19-3). As far as how Ansible works for managing network devices, it uses an agentless architecture. That means Ansible does not rely on any code (agent) running on the network device. Instead, Ansible relies on features typical in network devices, namely SSH and/or NETCONF, to make changes and extract information. When using SSH, the Ansible control node actually makes changes to the device like any other SSH user would do, but doing the work with Ansible code, rather than with a human. Ansible can be described as using a push model (per Figure 19-8) rather than a pull model (like Puppet and Chef). After installing Ansible, an engineer needs to create and edit all the various Ansible files, including an Ansible playbook. Then the engineer runs the playbook, which tells Ansible to perform the steps. Those steps can include configuring one or more devices per the various files (step 3), with the control node seen as pushing the configuration to the device. 1 Build files Subset Inventory Playbook SSH 3 R1 Push Config 2 Run Playbook Templates Variables Ansible Control Node Figure 19-8 Ansible Push Model As with all the tools, Ansible can do both configuration provisioning (configuring devices after changes are made in the files) and configuration monitoring (checking to find out whether the device config matches the ideal configuration on the control node). However, Ansible's architecture more naturally fits with configuration provisioning, as seen in the figure. To do configuration monitoring, Ansible uses logic modules that detect and list configuration differences, after which the playbook defines what action to take (reconfigure or notify). 19 440 CCNA 200-301 Official Cert Guide, Volume 2 Puppet To use Puppet (www.puppet.com), like Ansible, begin by installing Puppet on a Linux host. You can install it on your own Linux host, but for production purposes, you will normally install it on a Linux server called a Puppet master. As with Ansible, you can use a free open-source version with paid versions available. You can get started learning Puppet without a separate server for learning and testing. Once installed, Puppet also uses several important text files with different components, such as the following: ■ Manifest: This is a human-readable text file on the Puppet master, using a language defined by Puppet, used to define the desired configuration state of a device. ■ Resource, Class, Module: These terms refer to components of the manifest, with the largest component (module) being composed of smaller classes, which are in turn composed of resources. ■ Templates: Using a Puppet domain-specific language, these files allow Puppet to generate manifests (and modules, classes, and resources) by substituting variables into the template. One way to think about the differences between Ansible's versus Puppet's approach is that Ansible's playbooks use an imperative language, whereas Puppet uses a declarative language. For instance, with Ansible, the playbook will list tasks and choices based on those results, like "Configure all branch routers in these locations, and if errors occur for any device, do these extra tasks for that device." Puppet manifests instead declare the end state that a device should have: "This branch router should have the configuration in this file by the end of the process." The manifest, built by the engineer, defines the end state, and Puppet has the job to cause the device to have that configuration, without being told the specific set of steps to take. Puppet typically uses an agent-based architecture for network device support. Some network devices enable Puppet support via an on-device agent

—think of it as another feature configurable on the device. However, not every Cisco OS supports Puppet agents, so Puppet solves that problem using a proxy agent running on some external host (called agentless operation). The external agent then uses SSH to communicate with the network device, as shown in Figure 19-9. Internal Agent R1 API Manifest SSH API R2 External Agent Figure 19-9 Puppet Master Agent-based and Agentless Operation for Puppet Chapter 19: Understanding Ansible, Puppet, and Chef 441 Note Per Puppet's website, Puppet supports both an agent-based and agentless architecture, with the agentless architecture being the case of using an agent external to the network device, as shown in the lower part of Figure 19-9. Armed with a manifest that declares something like "This device should have this configuration state," Puppet uses a pull model to make that configuration appear in the device, as shown in Figure 19-10. Once installed, these steps occur: Step 1. The engineer creates and edits all the files on the Puppet server. Step 2. The engineer configures and enables the on-device agent or a proxy agent for each device. Step 3. The agent pulls manifest details from the server, which tells the agent what its configuration should be. Step 4. If the agent device's configuration should be updated, the Puppet agent performs additional pulls to get all required detail, with the agent updating the device configuration. 1 Build files 2 Start Agent 19 Manifest 3 Pull Details Templates R1 4 Pull Config Variables Puppet Master Figure 19-10 Pull Model with Puppet Chef

Chef (www.chef.io), as with Ansible and Puppet, exists as software packages you install and run. Chef (the company) offers several products, with Chef Automate being the product that most people refer to simply as Chef. As with Puppet, in production you probably run Chef as a server (called server-client mode), with multiple Chef workstations used by the engineering staff to build Chef files that are stored on the Chef server. However, you can also run Chef in standalone mode (called Chef Zero), which is helpful when you're just getting started and learning in the lab. 424 CCNA 200-301 Official Cert Guide, Volume 2 Once Chef is installed, you create several text files with different components, like the following: ■ Resource: The configuration objects whose state is managed by Chef; for instance, a set of configuration commands for a network device—analogueous to the ingredients in a recipe in a cookbook ■ Recipe: The Chef logic applied to resources to determine when, how, and whether to act against the resources—analogueous to a recipe in a cookbook ■ Cookbooks: A set of recipes about the same kinds of work, grouped together for easier management and sharing ■ Runlist: An ordered list of recipes that should be run against a given device Chef uses an architecture similar to Puppet. For network devices, each managed device (called a Chef node or Chef client) runs an agent. The agent performs configuration monitoring in that the client pulls recipes and resources from the Chef server and then adjusts its configuration to stay in sync with the details in those recipes and runlists. Note however that Chef requires on-device Chef client code, and many Cisco devices do not support a Chef client, so you will likely see more use of Ansible and Puppet for Cisco device configuration management. Summary of Configuration Management Tools All three of the configuration management tools listed here have a good base of users and different strengths. As for their use for managing network device configuration, Ansible appears to have the most interest, then Puppet, and then Chef. Ansible's agentless architecture and the use of SSH provides support for a wide range of Cisco devices. Puppet's agentless model also creates wide support for Cisco devices. Table 19-2 summarizes a few of the most common ideas about each of the three automated configuration management tools. Note that the column for Puppet assumes an on-device agent. Table 19-2 Comparing Ansible, Puppet, and Chef Action Ansible Puppet Chef Term for the file that lists actions Playbook Manifest Recipe, Runlist Protocol to network device SSH, NETCONF HTTP (REST) HTTP (REST) Uses agent or agentless model Agentless Agent* Agent Push or pull model Push Pull Pull * Puppet can use an in-device agent or an external proxy agent for network devices. Chapter Review One key to doing well on the exams is to perform repetitive spaced review sessions. Review this chapter's material using either the tools in the book or interactive tools for the same material found on the book's companion website. Refer to the "Your Study Plan" element for more details. Table 19-3 outlines the key review elements and where you can find them. To better track your study progress, record when you completed these activities in the second column. Chapter 19: Understanding Ansible, Puppet, and Chef 443 Table 19-3 Chapter Review Tracking Element Review Date(s) Resource Use Review key topics Book, website Review key terms Book, website Repeat DIKTA questions Book, PTP Review memory table Book, website Do DevNet Labs DevNet Review All the Key Topics Table 19-4 Key Topics for Chapter 19 Chapter 19 Key Topic Element Description Page Number List Issues that arise from configuration drift 431 Figure 19-3 Sample of showing router configuration file differences with GitHub 433 Figure 19-5 Basic configuration monitoring concepts. 434 List Primary functions of a configuration management tool 434 Example 19-2 Sample Jinja2 Ansible template 436 List Advantages of using configuration templates 437 Figure 19-8 Ansible's push model and other features 439 Figure 19-10 Puppet's pull model and other features 441 Table 19-2 Summary of configuration management features and terms 442 Key Terms You Should Know configuration monitoring, configuration provisioning, configuration drift, configuration management tool, Git, Ansible, Puppet, Chef, configuration template, push model, pull model, agent-based architecture, agentless architecture, Ansible playbook, Puppet manifest, Chef recipe Do DevNet Labs Cisco's DevNet site ()—a free site—includes lab environments and exercises. You can learn a lot about configuration management and Ansible in particular with a few of the lab tracks on the DevNet site (at the time this book was published). Refer to the "Chapter Review" section of the companion website for links to some good labs, or just go to and search for learning labs about Ansible. 19 Part V Review Keep track of your part review progress with the checklist shown in Table P5-1. Details on each task follow the table. Table P5-1 Part V Review Checklist Activity 1st Date Completed 2nd Date Completed Repeat All DIKTA Questions Answer Part Review Questions Review Key Topics Repeat All DIKTA Questions For this task, use the PTP software to answer the "Do I Know This Already?" questions again for the chapters in this part of the book. Answer Part Review Questions For this task, use PTP to answer the Part Review Questions for this part of the book. Review Key Topics Review all key topics in all chapters in this part, either by browsing the chapters or by using the Key Topics application on the companion website. This page intentionally left blank Part VI Final Review Chapter 20: Final Review CHAPTER 20 Final Review Congratulations! You made it through the book, and now it's time to finish getting ready for the exam. This chapter helps you get ready to take and pass the exam in two ways. First, this chapter focuses on the exam event. Now you need to think about what happens during the exam and what you need to do in these last few weeks before taking the exam. At this point, everything you do should be focused on getting ready to pass so that you can finish up this hefty task. The second section of this chapter focuses on final content review. You should not just complete the previous chapter, which is the 48th technology chapter in the combined CCNA 200-301 Official Cert Guide, Volume 1 and 2 books. Instead, you need to review, refine, deepen, and assess your skills. This second section of this chapter gives advice and suggestions on how to approach your final weeks of study before you take the CCNA 200-301 exam. Advice about the Exam Event Now that you have finished the bulk of this book, you could just register for your Cisco CCNA exam, show up, and take the exam. However, if you spend a little time thinking about the exam event itself, learning more about the user interface of the real Cisco exams and the environment at the Pearson VUE testing centers, you will be better prepared, particularly if this is your first Cisco exam. This first of two major sections in this chapter gives some advice about the Cisco exams and the exam event itself, specifically about ■ Question types ■ Your time budget ■ A sample time-check method ■ The final week ■ The 24 hours before the exam ■ The final 30 minutes before the exam ■ The hour after the exam Exam Event: Learn About Question Types In the weeks leading up to your exam, you should think more about the different types of exam questions and have a plan for how to approach those questions. One of the best ways to learn about the exam questions is to use some videos from the former Cisco Certification Exam Tutorial. As for the backstory, Cisco formerly published a tool (the Cisco Certification Exam Tutorial) that gave anyone the ability to experience the Cisco exam user interface via an interactive flash application. Cisco has updated the real exam interface, plus, Cisco removed the exam tutorial web pages with no equivalent replacement. However, Cisco did make videos of the exam tutorial, with someone talking through the various question types. Cisco lists the videos in a post at the Cisco Learning Network (), so you can start by looking for those videos as follows: ■ Go to the CLN () and search for the post 34312. ■ Use this direct link to the same page: . ■ Use which links to a blog post of mine that lists the above link (as well as other links useful for final review). While watching any of the videos about the exam tutorial, pay close attention to some important behaviors. For instance, for multichoice questions, the user interface ■ Identifies single-answer questions with circles beside the answers versus multiple-answer questions showing squares before the answers. ■ Prevents you from choosing too many answers. ■ Supplies a popup window to tell you if you have selected too few answers if you try to move to the next question, so you can stop and go back and answer with the correct number of answers. ■ Does not penalize you for guessing. You should always supply the number of answers that the question asks for. There is no penalty for guessing. Note that because there is no penalty for guessing, you should always answer every question and answer with the exact number of correct answers. For drag-and-drop questions, the user interface lets you change your mind while you are still working on the question. The draggable items begin in one location, and you drag and drop them to answer. You can just drag them back to where they were to begin the question. For simulation questions: ■ Pay close attention to the navigation to get to the command-line interface (CLI) on one of the routers. To do so, you have to click the PC icon for a PC connected to the router console; the console cable appears as a dashed line, whereas network cables are solid lines. (You should definitely look for this interaction in the exam tutorial videos.) ■ Make sure that you look at the scroll areas at the top, at the side, and in the terminal emulator window. These scrollbars let you view the entire question and scenario. ■ Make sure that you can toggle between the topology window and the terminal emulator window by clicking Show topology and Hide topology. The question window can be pretty crowded for sim questions, so the user interface gives you the means to toggle between side different parts of the question. 450 CCNA 200-301 Official Cert Guide, Volume 2 Both simlet and testlet questions give you one scenario with a group of related multichoice questions. However, the behavior with this small group (usually three or four) of multiple-choice questions differs from the flow of the more common standalone multiple-choice questions. In particular: ■ You can move between the multiple-choice questions in a single simlet or testlet. You can answer one multiple-choice question, move to the second and answer it, and then move back to the first question, confirming that inside a testlet you can move around between questions. ■ You can make a big mistake by not answering all questions or by not supplying enough answers, and the user interface does not prevent you from making that mistake. On that second point, consider this scenario with a simlet question. You see the simlet question, answer the first three multiple-choice questions, but forget to look at the fourth multiple-choice question. If you click Next, you will see a generic popup window that Cisco uses as a prompt to ask whether you want to move on. However, it does not tell you that you did not answer a question at all, and it does not tell you if you answered with too few answers on a multi-answer question. So be very careful when clicking Next when answering simlet and testlet questions. Exam Event: Think About Your Time Budget On exam day, you need to keep an eye on your speed. Going too slowly hurts you because you might not have time to answer all the questions. Going too fast can be harmful if you are rushing because you are fearful about running out of time. So, you need to be able to somehow know whether you are moving quickly enough to answer all the questions, while not rushing. The exam user interface shows some useful information, namely a countdown timer and a question counter. The question counter shows a question number for the question you are answering, and it shows the total number of questions on your exam. Unfortunately, some questions require lots more time than others, and for this and other reasons, time estimating can be a challenge. First, before you show up to take the exam, you know only a range of the number of questions for the exam; for example, the Cisco website might list the CCNA exam as having from 50 to 60 questions (the Cisco website did not list a number of questions at the time this chapter was published). You will not know how many questions are on your exam until the exam begins, when you go through the screens that lead up to the point where you click Start Exam, which starts your timed exam. Next, some questions (call them time burners) clearly take a lot more time to answer: Normal-time questions: Multiple-choice and drag-and-drop, approximately one minute each Time burners: Sims, simlets, and testlets, approximately six to eight minutes each Finally, even though testlet and simlet questions contain several multiple-choice questions, the exam software counts each testlet and simlet question as one question in the question counter. For example, if a testlet question has four embedded multiple-choice questions, in the exam software's question counter, that counts as one question. So when you start Chapter 20: Final Review 451 the exam, you might see that you will have 50 questions, but you don't know how many of those are time burners. NOTE Cisco does not tell us why one person taking the exam might get 50 questions while someone else taking the same exam might get 60 questions, but it seems reasonable to think that the person with 50 questions might have a few more of the time burners, making the two exams equivalent. You need a plan for how you will check your time, a plan that does not distract you from the exam. You can ponder the facts listed here and come up with your own plan. If you want a little more guidance, the next topic shows one way to check your time that uses some simple math so that it does not take much time away from the test. Exam Event: A Sample Time-Check Method As a suggestion, you can use the following math to do your time-check in a way that weights the time based on those time-burner questions. You do not have to use this method. But this math uses only addition of whole numbers, to keep it simple. It gives you a pretty close time estimate, in my opinion. The concept is simple. Just do a simple calculation that estimates the time you should have used so far. Here's the math: Number of questions answered so far + 7 per time burner answered so far 20 Then you check the timer to figure out how much time you have spent. ■ You have used exactly that much time or a little more time: Your timing is perfect. ■ You have used less time: You are ahead of schedule. ■ You have used noticeably more time: You are behind schedule. For example, if you have already finished 17 questions, two of which were time burners, your time estimate is 17 + 7 + 7 = 31 minutes, or if your actual time is also 31 minutes, or maybe 32 or 33 minutes, you are right on schedule. If you have spent less than 31 minutes, you are ahead of schedule. So, the math is pretty easy: questions answered, plus 7 per time burner, is the guesstimate of how long you should have taken so far if you are right on time. NOTE This math is an estimate; I make no guarantees that the math will be an accurate predictor on every exam. Exam Event: One Week Away I have listed a variety of tips in the next few pages, broken down by timing versus the big exam event. First, this section discusses some items to consider when your exam is about a week away: ■ Get some earplugs: Testing centers often have some, but if you do not want to chance it, come prepared with your own. (They will not let you bring your own noise-canceling headphones into the room if they follow the rules disallowing any user electronic 452 CCNA 200-301 Official Cert Guide, Volume 2 devices in the room, so think low-tech disposable earplugs, or even bring a cotton ball.) The testing center is typically one room within a building of a company that does something else as well, often a training center, and almost certainly you will share the room with other test takers coming and going. So, there are people talking in nearby rooms and other office noises. Earplugs can help. ■ Create an exam-event-note-taking plan: Some people like to spend the first minute of the exam writing down some notes for reference, before actually starting the exam. For example, maybe you want to write down the table of magic numbers for finding IPv4 subnet IDs. If you plan to do that, practice making those notes between now and exam day. Before each practice exam, transcribe those lists, just like you expect to do at the real exam. ■ Plan your travel to the testing center: Leave enough time in your schedule so that you will not be rushing to make it just in time. ■ Practice your favorite relaxation techniques for a few minutes before each practice exam: That way you can enter the exam event and be more relaxed and have more success. Exam Event: 24 Hours Before the Exam After you wake up on the big day, what should you be doing and thinking? Certainly, the better prepared you are, the better chances you have on the exam. But these small tips can help you do your best on exam day: ■ Rest the night before the exam rather than staying up late to study. Clarity of thought is more important than one extra fact, especially because the exam requires so much analyzing and thinking rather than just remembering facts. ■ Bring as few extra items with you as possible when leaving for the exam center. You may bring personal effects into the building and testing company's space, but not into the actual room in which you take the exam. So, save a little stress and bring as little extra stuff with you as possible. If you have a safe place to leave briefcases, purses, electronics, and so on, leave them there. However, the testing center should have a place to store your things as well. Simply put, the less you bring, the less you have to worry about storing. (For example, I have been asked to remove even my analog wristwatch on more than one occasion.) ■ Plan time in your schedule for the day to not rush to get there and not rush when leaving either. ■ Do not drink a 64-ounce caffeinated drink on the trip to the testing center. After the exam starts, the exam timer will not stop while you go to the restroom. ■ Use any relaxation techniques that you have practiced to help get your mind focused while you wait for the exam. Exam Event: The Last 30 Minutes It's almost time! Here are a few tips for those last moments. ■ Ask the testing center personnel for earplugs if you do not bring any—even if you cannot imagine using them. You never know whether using them might help. ■ Ask for extra pens and laminated note sheets. The exam center will give you a laminated sheet and dry erase pen to take notes. (Test center personnel typically do not let you Chapter 20: Final Review 453 bring paper and ink pen into the room, even if supplied by the testing center.) I always ask for a second pen as well. ■ Test your pens and sheets before going into the room to take the exam. Better to get a replacement pen before the clock starts. ■ Grab a few tissues from the box in the room, for two reasons. One, to avoid having to get up in the middle of the exam if you need to sneeze. Two, if you need to erase your laminated sheet, doing that with a tissue rather than your hand helps prevent the oil from your hand making the pen stop working well. ■ Find a restroom to use before going into the testing center, or just ask where one is, to avoid needing to go during the approximately two-hour exam event. Note that the exam timer does not stop if you need to go to the restroom during the exam, and you first have to find the exam center contact before just heading to the restroom, so it can cost you a few minutes. Exam Event: Reserve the Hour After the Exam Some people pass these exams on the first attempt, and some do not. The exams are not easy. If you fail to pass the exam that day, you will likely be disappointed. And that is understandable. And that is not a reason to give up. In fact, I added this short topic to give you a big advantage in case you do fail. The most important study hour for your next exam attempt is the hour just after your failed attempt. Before you take the exam, prepare for how you will react if you do not pass. That is, prepare your schedule to give yourself an hour, or at least a half an hour, immediately after the exam attempt, in case you fail. Follow these suggestions to be ready for taking notes: ■ Bring pen and paper, preferably a notebook you can write in if you have to write standing up or sitting somewhere inconvenient. ■ Make sure you know where pen and paper are so that you can take notes immediately after the exam. Keep these items in your backpack if using the train or bus, or on your car seat. ■ Install an audio recording app on your phone, and be prepared to start talking into your app when you leave the testing center. ■ Before the exam, scout the testing center, and plan the place where you will sit and take your notes, preferably somewhere quiet. Then, once you complete the exam, if you do not pass on this attempt, use the following process when taking notes: ■ Write down anything in particular that you can recall from any question. ■ Write down details of questions you know you got right as well, because doing so may help trigger a memory of another question. ■ Draw the figures that you can remember. ■ Most importantly, write down any tidbit that might have confused you: terms, configuration commands, show commands, scenarios, topology drawings, anything. ■ Take at least three passes at remembering. That is, you will hit a wall where you do not remember more. So, start on your way back to the next place, and then find a place to pause and take more notes. And do it again. 20 454 CCNA 200-301 Official Cert Guide, Volume 2 ■ When you have sucked your memory dry, take one more pass while thinking of the major topics in the book, to see if that triggers any other memory of a question. Once you have collected your notes, you cannot share the information with anyone because doing so would break the Cisco nondisclosure agreement (NDA). Cisco considers cheating a serious offense and strongly forbids sharing this kind of information publicly. But you can use your information to study for your next attempt. Remember, anything you can do to determine what you do not know is valuable when studying for your next attempt. See the section "Exam Review: Study Suggestions for Your Second Attempt" in this chapter for the rest of the story. Exam Review At this point, you should have read the other chapters in both the CCNA 200-301 Official Cert Guide, Volumes 1 and 2, and completed the Chapter Review and Part Review tasks. Now you need to do the final study and review activities before taking the exam, as detailed in this section. This section suggests some new activities and repeats some activities that have been previously mentioned. However, whether the activities are new or old to you, they all focus on filling in your knowledge gaps, finishing off your skills, and finalizing the study process. While repeating some tasks you did at Chapter Review and Part Review can help, you need to be ready to take an exam, so the Exam Review asks you to spend a lot of time answering exam questions. The Exam Review walks you through suggestions for several types of tasks and gives you some tracking tables for each activity. The main categories are ■ Taking practice exams ■ Finding what you do not know well yet (knowledge gaps) ■ Configuring and verifying functions from the CLI ■ Repeating the Chapter Review and Part Review tasks Exam Review: Take Practice Exams One day soon, you need to pass a real Cisco exam at a Pearson VUE testing center. So, it's time to practice the real event as much as possible. A practice exam using the Pearson IT Certification Practice Test (PTP) exam software lets you experience many of the same issues as when taking a real Cisco exam. When you select practice exam mode, the PTP software (both desktop and web) gives you a number of questions, with a countdown timer shown in the window. When using this PTP mode, after you answer a question, you cannot go back to it (yes, that's true on Cisco exams). If you run out of time, the questions you did not answer count as incorrect. The process of taking the timed practice exams helps you prepare in three key ways: ■ To practice the exam event itself, including time pressure, the need to read carefully, and the need to concentrate for long periods ■ To build your analysis and critical thinking skills when examining the network scenario built in to many questions ■ To discover the gaps in your networking knowledge so that you can study those topics before the real exam Chapter 20: Final Review 455 As much as possible, treat the practice exam events as if you were taking the real Cisco exam at a VUE testing center. The following list gives some advice on how to make your practice exam more meaningful, rather than just one more thing to do before exam day rolls around: ■ Set aside two hours for taking a 90-minute timed practice exam. ■ Make a list of what you expect to do for the 10 minutes before the real exam event. Then visualize yourself doing those things. Before taking each practice exam, practice those final 10 minutes before your exam timer starts. (The earlier section "Exam Event: The Last 30 Minutes" lists some suggestions about what to do in those last 10 minutes.) ■ You cannot bring anything with you into the VUE exam room, so remove all notes and help materials from your work area before taking a practice exam. You can use blank paper, a pen, and your brain only. Do not use calculators, notes, web browsers, or any other app on your computer. ■ Real life can get in the way, but if at all possible, ask anyone around you to leave you alone for the time you will practice. If you must do your practice exam in a distracting environment, wear headphones or earplugs to reduce distractions. ■ Do not guess, hoping to improve your score. Answer only when you have confidence in the answer. Then, if you get the question wrong, you can go back and think more about the question in a later study session. Using the Practice CCNA Exams The PTP questions you can access as part of this book include exam banks labeled as follows: ■ CCNA Volume 2 Exam 1 ■ CCNA Volume 2 Exam 2 ■ CCNA 200-301 Full Exam 1 ■ CCNA 200-301 Full Exam 2 The exams whose names begin "CCNA Volume 2" have questions from this Volume 2 book only, but no questions from Volume 1. The exams titled "CCNA 200-301" (without Volume 2 in the name) include questions from the entire breadth of CCNA topics, including topics covered in both the Volume 1 and Volume 2 books. You should do your final review with the CCNA 200-301 exams. Just select those exams and deselect the others. Then you simply need to choose the Practice Exam option in the upper right and start the exam. You should plan to take between one and three practice exams with the supplied CCNA exam databases. Even people who are already well prepared should do at least one practice exam, just to experience the time pressure and the need for prolonged concentration. Table 20-1 gives you a checklist to record your different practice exam events. Note that recording both the date and the score is helpful for some other work you will do, so note both. Also, in the Time Notes section, if you finish on time, note how much extra time you had; if you run out of time, note how many questions you did not have time to answer. 20 456 CCNA 200-301 Official Cert Guide, Volume 2 Table 20-1 Exam CCNA Practice Exam Checklist Date Score Time Notes CCNA CCNA CCNA Exam Review: Advice on How to Answer Exam Questions Our everyday habits have changed how we all read and think in front of a screen. Unfortunately, those same habits often hurt our scores when taking computer-based exams. For example, open a web browser. Yes, take a break and open a web browser on any device. Do a quick search on a fun topic. Then, before you click a link, get ready to think about what you just did. Where did your eyes go for the first 5 to 10 seconds after you opened that web page. Now, click a link and look at the page. Where did your eyes go? Interestingly, web browsers and the content in web pages have trained us all to scan. Web page designers actually design content expecting certain scan patterns from viewers. Regardless of the pattern, when reading a web page, almost no one reads sequentially, and no one reads entire sentences. People scan for the interesting graphics and the big words, and then scan the space around those noticeable items. Other parts of our electronic culture have also changed how the average person reads. For example, many of you grew up using texting and social media, sifting through hundreds or thousands of messages—but each message barely fills an entire sentence. Also, we find ourselves responding to texts, tweets, and emails and later realizing we did not really understand what the other person meant. If you use those same habits when taking the exam, you will probably make some mistakes because you missed a key fact in the question, answer, or exhibits. It helps to start at the beginning and read all the words—a process that is amazingly unnatural for many people today. NOTE I have talked to many college professors, in multiple disciplines, and Cisco Networking Academy instructors, and they consistently tell me that the number-one test-taking issue today is that people do not read the questions well enough to understand the details. When you are taking the practice exams and answering individual questions, consider these two strategies. First, before the practice exam, think about your own personal strategy for how you will read a question. Make your approach to multiple-choice questions in particular be a conscious decision on your part. Second, if you want some suggestions on how to read an exam question, use the following strategy: Step 1. Read the question itself, thoroughly, from start to finish. Step 2. Scan any exhibit or figure. Chapter 20: Final Review 457 Step 3. Scan the answers to look for the types of information. (Numeric? Terms? Single words? Phrases?) Step 4. Reread the question thoroughly, from start to finish, to make sure that you understand it. Step 5. Read each answer thoroughly, while referring to the figure/exhibit as needed. After reading each answer, before reading the next answer: A. If correct, select as correct. B. If for sure incorrect, mentally rule it out. C. If unsure, mentally note it as a possible correct answer. NOTE Cisco exams will tell you the number of correct answers. The exam software also helps you finish the question with the right number of answers noted. For example, for standalone multichoice questions, the software prevents you from selecting too many or too few answers. And you should guess the answer when unsure on the actual exam; there is no penalty for guessing. Use the practice exams as a place to practice your approach to reading. Every time you click to the next question, try to read the question following your approach. If you are feeling time pressure, that is the perfect time to keep practicing your approach, to reduce and eliminate questions you miss because of scanning the question instead of reading thoroughly. Exam Review: Additional Exams with the Premium Edition Many people add other practice exams and questions other than those that come with this book. Frankly, using other practice exams in addition to the questions that come with this book can be a good idea, for many reasons. The other exam questions can use different terms in different ways, emphasize different topics, and show different scenarios that make you rethink some topics. Note that Cisco Press does sell products that include additional test questions. The CCNA 200-301 Official Cert Guide, Volume 2, Premium Edition eBook and Practice Test product is basically the publisher's eBook version of this book. It includes a soft copy of the book in formats you can read on your computer and on the most common book readers and tablets. The product includes all the electronic content you would normally get with the print book, including all the question databases mentioned in this chapter. Additionally, this product includes two more CCNA exam databases (plus two more CCNA Volume 2 exam databases as well). NOTE In addition to providing the extra questions, the Premium Editions have links to every test question, including those in the print book, to the specific section of the book for further reference. This is a great learning tool if you need more detail than what you find in the question explanations. You can purchase the eBooks and additional practice exams at 70 percent off the list price using the coupon on the back of the activation code card in the cardboard sleeve, making the Premium Editions the best and most cost-efficient way to get more practice questions. 20 458 CCNA 200-301 Official Cert Guide, Volume 2 Exam Review: Find Knowledge Gaps One of the hardest things when doing your final exam preparation is to discover gaps in your knowledge and skills. In other words, what topics and skills do you need to know that you do not know? Or what topics do you think you know, but you misunderstand about some important fact? Finding gaps in your knowledge at this late stage requires more than just your gut feeling about your strengths and weaknesses. This next task uses a feature of PTP to help you find those gaps. The PTP software tracks each practice exam you take, remembering your answer for every question and whether you got it wrong. You can view the results and move back and forth between seeing the question and seeing the results page. To find gaps in your knowledge, follow these steps: Step 1. Pick and review one of your practice exams. Step 2. Review each incorrect question until you are satisfied that you understand the question. Step 3. When finished with your review for a question, mark the question. Step 4. Review all incorrect questions from your exam until all are marked. Step 5. Move on to the next practice exam. Figure 20-1 shows a sample Question Review page, in which all the questions were answered incorrectly. The results list a Correct column, with no check mark, meaning that the answer was incorrect. Figure 20-1 PTP Correct Results Page To perform the process of reviewing questions and marking them as complete, you can move between this Question Review page and the individual questions. Just double-click a question to move back to that question. From the question, you can click Grade Exam to Chapter 20: Final Review 459 move back to the grading results and to the Question Review page shown in Figure 20-1. The question window also shows the place to mark the question, in the upper left, as shown in Figure 20-2. Figure 20-2 Reviewing a Question, with the Mark Feature in the Upper Left If you want to come back later to look through the questions you missed from an earlier exam, start at the PTP home screen. From there, instead of clicking the Start button to start a new exam, click the View Grade History button to see your earlier exam attempts and work through any missed questions. Track your progress through your gap review in Table 20-2. PTP lists your previous practice exams by date and score, so it helps to note those values in the table for comparison to the PTP menu. Table 20-2 Tracking Checklist for Gap Review of Practice Exams Original Practice Exam Date Original Exam Score Date Gap Review Was Completed 460 CCNA 200-301 Official Cert Guide, Volume 2 Exam Review: Practice Hands-On CLI Skills To do well on sim and simlet questions, you need to be comfortable with many Cisco router and switch commands, and how to use them from a Cisco CLI. As described in the introduction to this book, sim questions require you to decide what configuration commands need to be configured to fix a problem or to complete a working configuration. Simlet questions require you to answer multiple-choice questions by first using the CLI to issue show commands to look at the status of routers and switches in a small network. To be ready for the exam, you need to know the following kinds of information: CLI navigation: Basic CLI mechanics of moving into and out of user, enable, and configuration modes Individual configuration: The meaning of the parameters of each configuration command Feature configuration: The set of configuration commands, both required and optional, for each feature Verification of configuration: The show commands that directly identify the configuration settings Verification of status: The show commands that list current status values and the ability to decide incorrect configuration or other problem causes of less-than-optimal status values To help remember and review all this knowledge and skill, you can do the tasks listed in the next several pages. CCNA Exam Topics with CLI Skill Requirements Wondering about all the topics in CCNA 200-301 that specifically include configuration or verification skills? You can just scan the CCNA 200-301 exam topics. However, Table 20-3 and Table 20-4 summarize the topics for which you could consider practicing your CLI skills. The tables organize the topics into the same order used in the CCNA 200-301 Official Cert Guides, Volume 1 and 2, with chapter references. Table 20-3 Topics with Configuration Skills in CCNA Volume 1 Topic Volume 1 Chapter Switch IPv4 6 Verifying LAN switching 5 Switch IPv4 6 Switch passwords 6 Switch interfaces 7 VLANs 8 VLAN trunking 8 STP and RSTP 10 Layer 2 EtherChannel 10 Router interfaces 15 Router IPv4 addresses and static routes 16 Date You Finished Lab Review Chapter 20: Final Review 461 Topic Volume 1 Chapter Router on a Stick 17 Layer 3 switching with SVIs 17 IPv4 You Finished Lab Review Layer 3 switching with VLAN interfaces and L3 17 EtherChannels OSPF fundamentals 20 OSPF network types 21 IPv4 addressing on routers 24 IPv6 static routes 25 Table 20-4 Topics with Configuration Skills in CCNA Volume 2 Topic Volume 2 Chapter Standard ACLs 2 Extended ACLs 3 Telnet and SSH Access ACLs 5 Port Security 6 DHCP client and DHCP relay 7 DHCP snooping 8 Dynamic ARP Inspection 8 Syslog, NTP, CDP, and LLDP 9 NAT, PAT 10 Date You Finished Lab Review You should research and choose your favorite methods and tools to get hands-on practice for CCNA. Those options include several that focus on giving you a specific activity to do. The options include the Pearson Network Simulator, Config Labs (on my blog), and Packet Tracer labs (on my blog). First, one great way to practice is to use the Pearson Network Simulator (the sim) at www.pearsoncertification.com/networksimulator. Pearson builds the sim to focus on lab exercises that help you learn and expand your skills with the topics in the CCNA exam. The sim also organizes the lab content so you can follow along with the books. You can get a sense for what the labs are like in the sim by going to the companion website for this book and downloading the Sim Lite, which uses the same core software but with a more limited number of labs compared to the full product. Second, review the Config Checklist apps available from the book's companion website. For any configuration topics that require more than a few commands, the book collects the configuration commands into config checklists so that you can review and study in the days leading up to the exam. Take advantage of those checklists to review and remember all the required and optional configuration commands. Finally, my blog site () has informal lab exercises

which they can then install their own applications. 5. A. Both options that use the Internet allow for easier migration because public cloud providers typically provide easy access over the Internet. An intercloud exchange is a purpose-built WAN service that connects to enterprises as well as most public cloud providers, with the advantage of making the cloud migration process easier. The one correct answer—the worst option in terms of being prepared for migrating to a new cloud provider—is to use a private WAN connection to one cloud provider. While useful in other ways, migrating when using this strategy would require installing a new private WAN connection to the new cloud provider. 6. A and C. Private WAN options use technologies like Ethernet WAN and MPLS, both of which keep data private by their nature and which include QoS services. An intercloud exchange is a purpose-built WAN service that connects to enterprises as well as C 490 CCNA 200-301 Official Cert Guide, Volume 2 most public cloud providers, using the same kinds of private WAN technology with those same benefits. For the two incorrect answers, both use the Internet, so both cannot provide QoS services. The Internet VPN option does encrypt the data to keep it private. Chapter 16 1. A. The data plane includes all networking device actions related to the receipt, processing, and forwarding of each message, as in the case described in the question. The term table plane is not used in networking. The management plane and control plane are not concerned with the per-message forwarding actions. 2. C. The control plane includes all networking device actions that create the information used by the data plane when processing messages. The control plane includes functions like IP routing protocols and Spanning Tree Protocol (STP). The term table plane is not used in networking. The management plane and data plane are not concerned with collecting the information that the data plane then uses. 3. C. Although many variations of SDN architectures exist, they typically use a centralized controller. That controller may centralize some or even all control plane functions in the controller. However, the data plane function of receiving messages, matching them based on header fields, taking actions (like making a forwarding decision), and forwarding the message still happens on the network elements (switches) and not on the controller. For the incorrect answers, the control plane functions may all happen on the controller, or some may happen on the controller, and some on the switches. The northbound and southbound interfaces are API interfaces on the controller, not on the switches. 4. A. The OpenDaylight Controller uses an Open SDN model with an OpenFlow southbound interface as defined by the Open Networking Foundation (ONF). The ONF SDN model centralizes most control plane functions. The APIC model for data centers partially centralizes control plane functions. The APIC-EM controller (as of time of publication) makes no changes to the control plane of routers and switches, leaving those to run with a completely distributed control plane. 5. C and D. ACI uses a spine-leaf topology. With a single-site topology, leaf switches must connect to all spine switches, and leaf switches must not connect to other leaf switches. Additionally, a leaf switch connects to some endpoints, with the endpoints being spread across the ports on all the leaf switches. (In some designs, two or more leaf switches connect to the same endpoints for redundancy and more capacity.) 6. A and D. Controller-based networks use a controller that communicates with each network device using a southbound interface (an API and protocol). By gathering network information into one central device, the controller can then allow for different operational modes. The models often let the operator think in terms of enabling features in the network, rather than thinking about the particulars of each device and command on each device. The controller then configures the specific commands, resulting in more consistent device configuration. Appendix C: Answers to the “Do I Know This Already?” Quizzes 491 For the incorrect answers, both the old and new models use forwarding tables on each device. Also, controllers do not add to or remove from the programmatic interfaces on each device, some of which existed before controllers, but rather supply useful and powerful northbound APIs. Chapter 17 1. C. The SDA underlay consists of the network devices and connections, along with configuration that allows IP connectivity between the SDA nodes, for the purpose of supporting overlay VXLAN tunnels. The fabric includes both the underlay and overlay, while VXLAN refers to the protocol used to create the tunnels used by the overlay. 2. B. The overlay includes the control plane and data plane features to locate the endpoints, decide to which fabric node a VXLAN tunnel should connect, direct the frames into the tunnel, and perform VXLAN tunnel encapsulation and de-encapsulation. The SDA underlay exists as network devices, links, and a separate IP network to provide connectivity between nodes to support the VXLAN tunnels. The fabric includes both the underlay and overlay, while VXLAN refers to the protocol used to create the tunnels used by the overlay. 3. D. The SDA overlay creates VXLAN tunnels between SDA edge nodes. Edge nodes then create a data plane by forwarding frames sent by endpoints over the VXLAN tunnels. LISP plays a role in the overlay as the control plane, which learns the identifiers of each endpoint, matching the endpoint to the fabric node that can reach the endpoint, so that the overlay knows where to create VXLAN tunnels. For the other incorrect answers, note that while GRE is a tunneling protocol, SDA uses VXLAN for tunneling, and not GRE. Finally, OSPF acts as a control plane routing protocol, rather than a data plane protocol for SDA. 4. A and D. As with any SDA feature, the configuration model is to configure the feature using DNA Center, with DNA Center using southbound APIs to communicate the intent to the devices. The methods to configure the feature using DNA Center include using the GUI or using the northbound REST-based API. Of the incorrect answers, you would not normally configure any of the SDA devices directly. Also, while DNA Center can use NETCONF as a southbound protocol to communicate with the SDA fabric nodes, it does not use NETCONF as a northbound API for configuration of features. 5. B, C, and D. Cisco DNA Center manages traditional network devices with traditional protocols like Telnet, SSH, and SNMP. DNA Center can also use NETCONF and RESTCONF if supported by the device. Note that while useful tools, Ansible and Puppet are not used by DNA Center. 6. A and D. Traditional network management platforms can do a large number of functions related to managing traditional networks and network devices, including the items listed in the two correct answers. However, when using Cisco’s Prime Infrastructure as a traditional network management platform for comparison, it does not support SDA configuration, nor does it find the end-to-end path between two endpoints and analyze the ACLs in the path. Note that the two incorrect answers reference features available in DNA Center. C 492 CCNA 200-301 Official Cert Guide, Volume 2 Chapter 18 1. B and D. The six primary required features of REST-based APIs include three features mentioned in the answers: a client/server architecture, stateless operation, notation of whether each object is cacheable. Two items from these three REST attributes are the correct answers. Of the incorrect answers, classful operation is the opposite of the REST-based API feature of classless operation. For the other incorrect answer, although many REST-based APIs happen to use HTTP, REST APIs do not have to use HTTP. 2. B and D. In the CRUD software development acronym, the matching terms (create, read, update, delete) match one or more HTTP verbs. While the HTTP verbs can sometimes be used for traditional CRUD actions, the following are the general rules: create performed by HTTP POST; read by HTTP GET; update by HTTP PATCH, PUT (and sometimes POST); delete by HTTP DELETE. 3. C. The URI for a REST API call uses a format of protocol://hostname/ resource?parameters. The API documentation details the resource part of the URI, as well as any optional parameters. For instance, in this case, the resource section is /dna/intent/api/v1/network-device. Additionally, the API documentation for this resource details optional parameters in the query field as listed after the ? in the URI. 4. A and D. Of the four answers, two happen to be most commonly used to format and serialize data returned from a REST API: JSON and XML. For the incorrect answers, JavaScript is a programming language that first defined JSON as a data serialization language. YAML is a data serialization/modeling language and can be found most often in configuration management tools like Ansible. 5. A and D. JSON defines variables as key:value pairs, with the key on the left of the colon (;) and always enclosed in double quotation marks, with the value on the right. The value can be a simple value or an object or array with additional complexity. The number of objects is defined by the number of matched curly brackets ({} and []). so this example shows a single JSON object. The one JSON object shown here includes one key and one ;, so it has a single key:value pair (making one answer correct). The value in that key:value pair itself is a JSON array (a list in Python) that lists numbers 1, 2, and 3. The fact that the list is enclosed in square brackets defines it as a JSON array. 6. C and D. To interpret this JSON data, first look for the innermost pairing of either curly brackets {}, which denote one object, or square brackets [], which denote one array. In this case, the gray highlighted area is one JSON object,
enclosed with {} and no other brackets of either type inside. That makes the gray area one object, which itself holds key:value pairs. Inside that one object, four key:value pairs exist, with the key before each colon and the value after each colon. That means “type” is a key, and “ACCESS” is one of the values. If you look at the other pair of curly brackets that begin and end the JSON data, that pair defines an object. That object has a key of “response” (making one answer incorrect). The “response” key then has a value equal to the entire inner object (the gray highlighted part), confirming one of the correct answers. Appendix C: Answers to the “Do I Know This Already?” Quizzes 493 Chapter 19 1. C. Devices with the same role in an enterprise should have a very similar configuration. When engineers make unique changes on individual devices—different changes from those made in the majority of devices with that same role—those devices’ configurations become different than the intended ideal configuration for every device with that role. This effect is known as configuration drift. Configuration management tools can monitor a device’s configuration versus a file that shows the intended ideal configuration for devices in that role, noting when the device configuration drifts away from that ideal configuration. 2. A and B. The version control system, applied to the centralized text files that contain the device configurations, automatically tracks changes. That means the system can see which user edited the file, when, and exactly what change was made, with the ability to make comparisons between different versions of the files. The two incorrect answers list very useful features of a configuration management tool, but those answers list features typically found in the configuration management tool itself rather than in the version control tool. 3. D. Configuration monitoring (a generic description) refers to a process of checking the device’s actual configuration versus the configuration management system’s intended configuration for the device. If the actual configuration has moved away from the intended configuration—that is, if configuration drift has occurred—configuration monitoring can either reconfigure the device or notify the engineering staff. For the other answers, two refer to features of the associated version control software typically used along with the configuration management tool. Version control software will track the identity of each user who changes files and track the differences in files over time. The other incorrect answer is a useful feature of many configuration management tools, in which the tool verifies that the configuration will be accepted when attempted (or not). However, that useful feature is not part of what is called configuration monitoring. 4. 5. A. Ansible uses a push model, in which the Ansible control node decides when to configure a device based on the instructions in a playbook. Puppet and Chef use pull models, in which an agent asks for information from a server, with the agent then making the decision of whether it needs to pull configuration data to itself and reconfigure itself. B and C. Of the terms manifest and recipe, both refer to files that define the actions to take and/or the most state desired when taking action in one of the configuration management tools. These files go by the names Ansible playbook, Puppet manifest, and Chef recipe. C GLOSSARY NUMERICS 3G/4G Internet An Internet access technology that uses wireless radio signals to communicate through mobile phone towers, most often used by mobile phones, tablets, and some other mobile devices. 802.1 Q The IEEE standardized protocol for VLAN trunking. A AAA Authentication, authorization, and accounting. Authentication confirms the identity of the user or device. Authorization determines what the user or device is allowed to do. Accounting records information about access attempts, including inappropriate requests. AAA server See authentication, authorization, and accounting (AAA) server. Access Control Entry (ACE) One line in an access control list (ACL). access interface A LAN network design term that refers to a switch interface connected to end-user devices, configured so that it does not use VLAN trunking. access layer In a campus LAN design, the switches that connect directly to endpoint devices (servers, user devices), and also connect into the distribution layer switches. access link In Frame Relay, the physical serial link that connects a Frame Relay DTE device, usually a router, to a Frame Relay switch. The access link uses the same physical layer standards as do point-to-point leased lines. access link (WAN) A physical link between a service provider and its customer that provides access to the SP’s network from that customer site. access rate The speed at which bits are sent over an access link. accounting In security, the recording of access attempts. See also AAA. ACI See Application Centric Infrastructure (ACI). ACL Access control list. A list configured on a router to control packet flow through the router, such as to prevent packets with a certain IP address from leaving a particular interface on the router. Active Directory A popular set of identity and directory services from Microsoft, used in part to authenticate users. administrative distance In Cisco routers, a means for one router to choose between multiple routes to reach the same subnet when those routes are learned by different routing protocols. The lower the administrative distance, the more preferred the source of the routing information. agent Generally, an additional software process or component running in a computing device for some specific purpose; a small and specific software service. agent-based architecture With configuration management tools, an architecture that uses a software agent inside the device being managed as part of the functions to manage the configuration. agentless architecture With configuration management tools, an architecture that does not need a software agent inside the device being managed as part of the functions to manage the configuration, instead using other mainstream methods like SSH and NETCONF. amplification attack A reflection attack that leverages a service on the reflector to generate and reflect huge volumes of reply traffic to the victim. analog modem See modem. Ansible A popular configuration management application, which can be used with or without a server, using a push model to move configurations into devices, with strong capabilities to manage network device configurations. Ansible inventory Device host names along with information about each device, like device roles, so Ansible can perform functions for subsets of the inventory. Ansible playbook Files with actions and logic about what Ansible should do. Ansible template A text file, written in Jinja2 language, that lists configuration but with variable names substituted for values, so that Ansible can create standard configurations for multiple devices from the same template. anti-replay Preventing a man in the middle from copying and later replaying the packets sent by a legitimate user, for the purpose of appearing to be a legitimate user. antivirus Software that monitors files transferred by any means, for example, web or email, to look for content that can be used to place a virus into a computer. APIC See Application Policy Infrastructure Controller. APIC-EM See Application Policy Infrastructure Controller—Enterprise Module. Application Centric Infrastructure (ACI) Cisco’s data center SDN solution, the concepts of defining policies that the APIC controller then pushes to the switches in the network using the OpFlex protocol, with the partially distributed control plane in each switch building the forwarding table entries to support the policies learned from the controller. It also supports a GUI, a CLI, and APIs. Application Policy Infrastructure Controller—Enterprise Module (APIC-EM) The software that plays the role of controller in an enterprise network of Cisco devices, in its first version as of the publication of this book, which leaves the distributed routing and switching control plane as is, instead acting as a management and automation platform. It provides robust APIs for network automation and uses CLI (Telnet and SSH) plus SNMP southbound to control the existing routers and switches in an enterprise network. 496 Application Policy Infrastructure Controller (APIC) Application Policy Infrastructure Controller (APIC) The software that plays the role of controller, controlling the flows that the switches create to define where frames are forwarded, in a Cisco data center that uses the Application Centric Infrastructure (ACI) approach, switches, and software. application programming interface (API) A software mechanism that enables software components to communicate with each other. application signature With Network Based Application Recognition (NBAR), the definition of a combination of matchable fields that Cisco has identified as being characteristic of a specific application, so that NBAR can be configured by the customer to match an application, while IOS then defines the particulars of that matching. Application Visibility and Control (AVC) A firewall device with advanced features, including the ability to run many related security features in the same firewall device (IPS, malware detection, VPN termination), along with deep packet inspection with Application Visibility and Control (AVC) and the ability to perform URL filtering versus data collected about the reliability and risk associated with every domain name. application-specific integrated circuit (ASIC) An integrated circuit (computer chip) designed for a specific purpose or application, often used to implement the functions of a networking device rather than running a software process as part of the device’s OS that runs on a general-purpose
processor. AR See access rate. ARP Address Resolution Protocol. An Internet protocol used to map an IP address to a MAC address. Defined in RFC 826. ARP ACL A configuration feature on Cisco LAN switches that define MAC and IP address pairs that can be used directly for filtering, as well as to be referenced by the Dynamic ARP Inspection feature. ARP reply An ARP message used to supply information about the sending (origin) host’s hardware (Ethernet) and IP addresses as listed in the origin hardware and origin IP address fields. Typically sent in reaction to receipt of an ARP request message. ARP request An ARP message used to request information from another host located on the same data link, typically listing a known target IP address but an all-zero target hardware address, to ask the host with that target IP address to identify its hardware address in an ARP reply message. ARP table A list of IP addresses of neighbors on the same VLAN, along with their MAC addresses, as kept in memory by hosts and routers. ASAV A Cisco ASA firewall software image that runs as a virtual machine rather than on Cisco hardware, intended to be used as a consumer-controlled firewall in a cloud service or in other virtualized environments. ASIC See application-specific integrated circuit. buffer overflow attack 497 Assured Forwarding (AF) The name of a grid of 12 DSCP values and a matching grid of per-hop behavior as defined by DiffServ. AF defines four queuing classes and three packet drop priorities within each queuing class. The text names the 12 DSCP values follow a format of AFXY, where X is the queuing class, and Y is the drop priority. authentication AAA. In security, the verification of the identity of a person or a process. See also authentication, authorization, and accounting (AAA) server A server that holds security information and provides services related to user login, particularly authentication (is the user who he says he is), authorization (once authenticated, what do we allow the user to do), and accounting (tracking the user). Authoritative DNS server The DNS server with the record that lists the address that corresponds to a domain name (the A Record) for that domain. authorization In security, the determination of the rights allowed for a particular user or device. See also AAA. autonomous system (AS) An internetwork that is managed by one organization. autonomous system number (ASN) A number used by BGP to identify a routing domain, often a single enterprise or organization. As used with EIGRP, a number that identifies the routing processes on routers that are willing to exchange EIGRP routing information with each other. AutoQoS In Cisco switches and routers, an IOS feature that configures a variety of QoS features with useful settings as defined by the Cisco reference design guide documents. B bandwidth The speed at which bits can be sent and received over a link. bandwidth profile In Metro Ethernet, a contractual definition of the amount of traffic that the customer can send into the service and receive out of the service. Includes a concept called the committed information rate (CIR), which defines the minimum amount of bandwidth (bits per second) the SP will deliver with the service. Brownfield A term that refers to the choice to add new configuration to hardware and software that are already in use, rather than adding new hardware and software specifically for a new project. brute-force attack An attack where a malicious user runs software that tries every possible combination of letters, numbers, and special characters to guess a user’s password. Attacks of this scale are usually run offline, where more computing resources and time are available. buffer overflow attack An attack meant to exploit a vulnerability in processing inbound traffic such that the target system’s buffers overflow; the target system can end up crashing or inadvertently running malicious code injected by the attacker. 498 cable Internet C cable Internet An Internet access technology that uses a cable TV (CATV) cable, normally used for video, to send and receive data. cacheable For resources that might be repeatedly requested over time, an attribute that means that the requesting host can keep in storage (cache) a copy of the resource for a specified amount of time. candidate config With configuration management tools like Ansible, Puppet, and Chef, an updated configuration for a device as it exists in the management tool before the tool has moved the configuration into the device. carrier Ethernet Per MEF documents, the term for what was formerly called Metro Ethernet, generally referring to any WAN service that uses Ethernet links as the access link between the customer and the service provider. CDP Cisco Discovery Protocol. A media- and protocol-independent device-discovery protocol that runs on most Cisco-manufactured equipment, including routers, access servers, and switches. Using CDP, a device can advertise its existence to other devices and receive information about other devices on the same LAN or on the remote side of a WAN. CDP neighbor CDP updates. A device on the other end of some communications cable that is advertising central office (CO) A term used by telcos to refer to a building that holds switching equipment, into which the telco’s cable plant runs so that the telco has cabling from each home and business into that building. centralized control plane An approach to architecting network protocols and products that places the control plane functions into a centralized function rather than distributing the function across the networking devices. Chef A popular configuration management application, which uses a server and a pull model with in-device agents. Chef client Any device whose configuration is being managed by Chef. Chef Cookbook A set of recipes about the same kinds of work, grouped together for easier management and sharing. Chef Recipe The Chef logic applied to resources to determine when, how, and whether to act against the resources—analogue to a recipe in a cookbook. Chef Runlist An ordered list of recipes that should be run against a given device. Chef server The Chef software that collects all the configuration files and other files used by Chef from different Chef users and then communicates with Chef clients (devices) so that the Chef clients can synchronize their configurations. CIDR Classless interdomain routing. An RFC-standard tool for global IP address range assignment. CIDR reduces the size of Internet routers’ IP routing tables, helping deal with the rapid growth of the Internet. The term classless refers to the fact that the summarized groups of networks represent a group of addresses that do not conform to IPv4 classful (Class A, B, and C) grouping rules. cloud services catalog 499 Cisco Access Control Server (ACS) A legacy Cisco product that acts as a AAA server. Cisco AnyConnect Secure Mobility Client Cisco software product used as client software on user devices to create a client VPN. Commonly referred to as the Cisco VPN client. Cisco Open SDN Controller (OSC) A former commercial SDN controller from Cisco that is based on the OpenDaylight controller. Cisco Prime Graphical user interface (GUI) software that utilizes SNMP and can be used to manage your Cisco network devices. The term Cisco Prime is an umbrella term that encompasses many different individual software products. Cisco Prime Infrastructure (PI) meant application. The name of Cisco’s long-time enterprise network management— Cisco Talos Intelligence Group A part of the Cisco Systems company that works to perform security research on an ongoing basis, in part to supply up-to-date data, like virus signatures, that Cisco security products can frequently download. Cisco VPN client See Cisco AnyConnect Secure Mobility Client. Class of Service (CoS) The informal term for the 3-bit field in the 802.1Q header intended for marking and classifying Ethernet frames for the purposes of applying QoS actions. Another term for Priority Code Point (PCP). Class Selector (CS) The name of eight DSCP values that all end with binary 000, for the purpose of having eight identifiable DSCP values whose first 3 bits match the eight values used for the older IP Precedence field. Originally used for backward compatibility with IP Precedence, but today the values are often used as just more values to use for packet marking. classification The process of examining various fields in networking messages in an effort to identify which messages fit into certain predetermined groups (classes). classless addressing A concept in IPv4 addressing that defines a subnetted IP address as having two parts: a prefix (or subnet) and a host. client VPN A VPN for which one endpoint is a user device, like a phone, tablet, or PC. Also called a remote access VPN. clock rate The speed at which a serial link encodes bits on the transmission medium. clock source On serial links, the device to which the other devices on the link adjust their speed when using synchronous links. With NTP, the external device or NTP server on which a device bases its time. clocking The process of supplying a signal over a cable, either on a separate pin on a serial cable or as part of the signal transitions in the transmitted signal, so that the receiving device can keep synchronization with the sending device. Clos network A term for network topology that represents an ideal for a switch fabric and named after Charles Clos, who formalized the definition. Also called a spine-leaf network. cloud services catalog A listing of the services available in a cloud computing service. 500 Cloud Services Router (CSR) Cloud Services Router (CSR) A Cisco router software image that runs as a virtual machine rather than on Cisco hardware, intended to be used as a consumer-controlled router in a cloud service or in other virtualized environments. code integrity A software security term that refers to how
likely that the software (code) being used is the software supplied by the vendor, unchanged, with no viruses or other changes made to the software. collapsed core design A campus LAN design in which the design does not use a separate set of core switches in addition to the distribution switches—in effect collapsing the core into the distribution switches. confidentiality (privacy) Preventing anyone in the middle of the Internet (a.k.a. man in the middle) from being able to read the data. configuration drift A phenomenon that begins with the idea that devices with similar roles can and should have a similar standard configuration, so when one device’s configuration is changed to be different, its configuration is considered to have moved away (drifted) from the standard configuration for a device in that role. configuration enforcement Another term for configuration monitoring. configuration management A component of network management focused on creating, changing, removing, and monitoring device configuration. configuration management tool A class of application that manages data about the configuration of servers, network devices, and other computing nodes, providing consistent means of describing the configurations, moving the configurations into the devices, noticing unintended changes to the configurations, and troubleshooting by easily identifying changes to the configuration files over time. configuration monitoring With configuration management tools like Ansible, Puppet, and Chef, a process of comparing over time a device’s on-device configuration (running-config) versus the text file showing the ideal device configuration listed in the tool’s centralized configuration repository. If different, the process can either change the device’s configuration or report the issue. configuration provisioning With configuration management tools like Ansible, Puppet, and Chef, the process of configuring a device to match the configuration as held in the configuration management tool. configuration template With configuration management tools like Ansible, Puppet, and Chef, a file with variables, for the purpose of having the tool substitute different variable values to create the configuration for a device. congestion window With TCP, a calculation each TCP receiver does that limits the window it grants to the receiver by shrinking the window in response to the loss of TCP segments. connection establishment The process by which a connection-oriented protocol creates a connection. With TCP, a connection is established by a three-way transmission of TCP segments. declarative policy model 501 control plane Functions in networking devices and controllers that directly control how devices perform data plane forwarding, but excluding the data plane processes that work to forward each message in the network. controller-based networking A style of building computer networks that use a controller that centralizes some features and provides application programming interfaces (APIs) that allow for software interactions between applications and the controller (northbound APIs) and between the controller and the network devices (southbound APIs). core In computer architecture, an individual processing unit that can execute instructions of a CPU; modern server processors typically have multiple cores, each capable of concurrent execution of instructions. core design A campus LAN design that connects each access switch to distribution switches, and distribution switches into core switches, to provide a path between all LAN devices. core layer In a campus LAN design, the switches that connect the distribution layer switches, and to each other, to provide connectivity between the various distribution layer switches. CRUD In software development, an acronym that refers to the four most common actions taken by a program: Create, Read, Update, and Delete. customer edge (CE) A term used by service providers, both generally and also specifically in MPLS VPN networks, to refer to the customer device that connects to the SP’s network and therefore sits at the edge of the SP’s network. customer premises equipment (CPE) A telco term that refers to equipment on site at the telco customer site (the enterprise’s site) that connects to the WAN service provided by the telco. D data integrity Verifying that the packet was not changed as the packet transited the Internet. data model A set of variables and their structures, like lists and dictionaries. data modeling language Another term for data serialization language. data plane Functions in networking devices that are part of the process of receiving a message, processing the message, and forwarding the message. data serialization language A language that includes syntax and rules that provides a means to describes the variables inside applications in a text format, for the purpose of sending that text between applications over a network or storing the data models in a file. data structure Another term for data model. declarative policy model A term that describes the approach in an intent-based network (IBN) in which the engineer chooses settings that describe the intended network behavior (the declared policy) but does not command the network with specific configuration commands for each protocol (as would be the case with an imperative policy model). 502 decrypt/decryption decrypt/decryption The ability to receive encrypted data and process it to derive the original unencrypted data. default gateway/default router On an IP host, the IP address of some router to which the host sends packets when the packet’s destination address is on a subnet. delay In QoS, the amount of time it takes for a message to cross a network. Delay can refer to one-way delay (the time required for the message to be sent from the source host to the destination host) or two-way delay (the delay from the source to the destination host and then back again). demilitarized zone (DMZ) In an Internet edge design at an enterprise, one or more subnets set aside as a place to locate servers that should allow users in the Internet to initiate connections to those servers. The devices in the DMZ typically sit behind a firewall. denial-of-service (DoS) attack An attack that tries to deplete a system resource so that systems and services crash or become unavailable. deny An action taken with an ACL that implies that the packet is discarded. DevNet Cisco’s community and resource site for software developers, open to all, with many great learning resources; . DHCP Dynamic Host Configuration Protocol. A protocol used by hosts to dynamically discover and lease an IP address, and learn the correct subnet mask, default gateway, and DNS server IP addresses. DHCP attack Any attack that takes advantage of DHCP protocol messages. DHCP binding table A table built by the DHCP snooping feature on a switch when it sees messages about a new DHCP lease, with the table holding information about legitimate successful DHCP leases, including the device’s IP address, MAC address, switch port, and VLAN. DHCP chaddr Client hardware address. The original DHCP header field used to identify the DHCP client; typically includes the client MAC address. DHCP client Any device that uses DHCP protocols to ask to lease an IP address from a DHCP server or to learn any IP settings from that server. DHCP client identifier A DHCP header field used to identify a DHCP client, used as a more flexible alternative to the DHCP chaddr field. DHCP giaddr Gateway IP address. In DHCP, a header field used to identify a router on a subnet, typically an IP address on the DHCP relay agent, so that the DHCP server knows an address to which to send messages in reply to the client. DHCP option 82 Optional DHCP header fields, as defined in RFC 3046, that provide useful features of use to a device that acts as a DHCP relay agent. The fields allow better relay agent operation and also help prevent various types of DHCP-based attacks. DHCP relay agent The name of the router IOS feature that forwards DHCP messages from client to servers by changing the destination IP address from 255.255.255.255 to the IP address of the DHCP server. DNA Center 503 DHCP server Software that waits for DHCP clients to request to lease IP addresses, with the server assigning a lease of an IP address as well as listing other important IP settings for the client. DHCP Snooping A switch security feature in which the switch examines incoming DHCP messages and chooses to filter messages that are abnormal and therefore might be part of a DHCP attack. DHCP snooping binding table When using DHCP Snooping, a table that the switch dynamically builds by analyzing the DHCP messages that flow through the switch. DHCP Snooping can use the table for part of its filtering logic, with other features, such as Dynamic ARP Inspection and IP Source Guard also using the table. Dictionary attack An attack where a malicious user runs software that attempts to guess a user’s password by trying words from a dictionary or word list. dictionary variable In applications, a single variable whose value is a list of other variables with values, known as key:value pairs. Differentiated Services (DiffServ) An approach to QoS, originally defined in RFC 2475, that uses a model of applying QoS per classification, with planning of which applications and other traffic types are assigned to each class, with each class given different QoS per-hop behaviors at each networking device in the path. Differentiated Services Code Point (DSCP) A field existing as the first 6 bits of the ToS byte, as defined by RFC 2474, which redefined the original IP RFC’s definition for the IP header ToS byte. The field is used to mark a value in the header for the purpose of performing later QoS actions on the packet. Digital Subscriber Line (DSL) A public network technology that delivers high bandwidth over conventional telco local-loop copper wiring at limited
distances. Typically used as an Internet access technology, connecting a user to an ISP. distributed control plane An approach to architecting network protocols and products that places some control plane functions into each networking device rather than centralizing the control plane functions in one or a few devices. An example is the use of routing protocols on each router which then work together so that each router learns Layer 3 routes. distributed denial-of-service (DDoS) attack A DoS attack that is distributed across many hosts under centralized control of an attacker, all targeting the same victim. distribution layer In a campus LAN design, the switches that connect to access layer switches as the most efficient means to provide connectivity from the access layer into the other parts of the LAN. DNA Digital Network Architecture—Cisco’s software-oriented approach to networking and intent-based networking products and services. DNA Center Cisco software, delivered by Cisco on a server appliance, that acts as a network management application as well as a being the control for Cisco’s software-defined access (SDA) offering. 504 DNS DNS Domain Name System. An application layer protocol used throughout the Internet for translating host names into their associated IP addresses. DNS reply In the Domain Name System (DNS), a message sent by a DNS server to a DNS client in response to a DNS request, identifying the IP address assigned to a particular hostname or fully qualified domain name (FQDN). DNS request In the Domain Name System (DNS), a message sent by a DNS client to a DNS server, listing a hostname or fully qualified domain name (FQDN), asking the server to discover and reply with the IP address associated with that host name or FQDN. DNS server An application acting as a server for the purpose of providing name resolution services per the Domain Name System (DNS) protocol and worldwide system. domain-specific language A generic term that refers to an attribute of different languages within computing, for languages created for a specific purpose (domain) rather than a general-purpose language like Python or JavaScript. DSL Digital subscriber line. Public network technology that delivers high bandwidth over conventional telco local-loop copper wiring at limited distances. Usually used as an Internet access technology connecting a user to an ISP. DSL modem A device that connects to a telephone line and uses DSL standards to transmit and receive data to/from a telco using DSL. Dynamic ARP Inspection (DAI) A security feature in which a LAN switch filters a subset of incoming ARP messages on untrusted ports, based on a comparison of ARP, Ethernet, and IP header fields to data gathered in the IP DHCP Snooping binding table and found in any configured ARP ACLs. E egress tunnel router (ETR) With LISP, a node at the end of a tunnel that receives an encapsulated message and then de-encapsulates the message. E-LAN A specific carrier/Metro Ethernet service defined by MEF (MEF.net) that provides a service much like a LAN, with two or more customer sites connected to one E-LAN service in a full mesh so that each device in the E-LAN can send Ethernet frames directly to every other device. E-Line A specific carrier/metro Ethernet service defined by MEF (MEF.net) that provides a point-to-point topology between two customer devices, much as if the two devices were connected using an Ethernet crossover cable. enable mode A part of the Cisco IOS CLI in which the user can use the most powerful and potentially disruptive commands on a router or switch, including the ability to then reach configuration mode and reconfigure the router. enable password A reference to the password configured on the enable password passvalue command, which defines the password required to reach enable (privileged) mode if the enable secret pass-value command does not exist. E-Tree enable secret A reference to the password configured on the enable secret pass-value command, which defines the password required to reach enable (privileged) mode. encrypt/encryption The ability to take data and send the data in a form that is not readable by someone who intercepts this data. encryption key process. A secret value used as input to the math formulas used by an encryption End of Row (EoR) switch In a traditional data center design with servers in multiple racks and the racks in multiple rows, a switch placed in a rack at the end of the row, intended to be cabled to all the Top of Rack (TOR) switches in the same row, to act as a distribution layer switch for the switches in that row. endpoint group In ACL, a set (group) of VMs, containers, physical servers, or other endpoints in an ACI data center that should receive the same policy treatment. Endpoint ID (EID) With LISP, a number that identifies the end point. err-disable recovery Cisco switches can place ports in a nonworking state called “err-disabled” in reaction to a variety of events, and by default, to leave the port in the nonworking err-disabled state until the engineer takes action to recover from the issue. The err-disable recovery configuration feature includes settings to direct the switch to automatically revert away from the err-disabled state, back to a working state, after a period of time, error detection The process of discovering whether a data-link level frame was changed during transmission. This process typically uses a Frame Check Sequence (FCS) field in the datalink trailer. error disabled (err-disable) An interface state on LAN switches that can be the result of one of many security violations. error recovery The process of noticing when some transmitted data was not successfully received and resending the data until it is successfully received. Ethernet access link A WAN access link (a physical link between a service provider and its customer) that happens to use Ethernet. Ethernet LAN Service Another term for E-LAN; see also E-LAN. Ethernet Line Service Another term for E-Line; see also E-Line. Ethernet Tree Service Another term for E-Tree; see also E-Tree. Ethernet Virtual Connection (EVC) A concept in carrier/Metro Ethernet that defines which customer devices can send frames to each other over the Ethernet WAN service; includes E-Line, E-LAN, and E-Tree EVCs. Ethernet WAN A general and informal term for any WAN service that uses Ethernet links as the access link between the customer and the service provider. E-Tree A specific carrier/metro Ethernet service defined by MEF (MEF.net) that provides a rooted multipoint service, in which the root site can send frames directly to all leaves, but the leaf sites can send only to the root site. 505 506 Expedited Forwarding (EF) Expedited Forwarding (EF) The name of a particular DSCP value, as well as the term for one per-hop behavior as defined by DiffServ. The value, decimal 46, is marked for packets to which the networking devices should apply certain per-hop behaviors, like priority queuing. exploit A means of taking advantage of a vulnerability to compromise something. extended access list A list of IOS access-list global configuration commands that can match multiple parts of an IP packet, including the source and destination IP address and TCP/UDP ports, for the purpose of deciding which packets to discard and which to allow through the router. F fabric In SDA, the combination of overlay and underlay that together provide all features to deliver data across the network with the desired features and attributes. fabric border node In SDA, a switch that connects to devices outside SDA’s control—for example, switches that connect to the WAN routers or to an ACI data center. fabric control node In SDA, a switch that performs special functions for the underlay (LISP), requiring more CPU and memory. fabric edge node In SDA, a switch that connects to endpoint devices. fiber Internet A general term for any Internet access technology that happens to use fiberoptic cabling. It often uses Ethernet protocols on the fiber link. filter Generally, a process or a device that screens network traffic for certain characteristics, such as source address, destination address, or protocol. This process determines whether to forward or discard that traffic based on the established criteria. firewall A device that forwards packets between the less secure and more secure parts of the network, applying rules that determine which packets are allowed to pass and which are not. First Hop Redundancy Protocol (FHRP) A class of protocols that includes HSRP, VRRP, and GLBP, which allows multiple redundant routers on the same subnet to act as a single default router (first-hop router). flash memory A type of read/write permanent memory that retains its contents even with no power applied to the memory, and uses no moving parts, making the memory less likely to fail over time. flow control The process of regulating the amount of data sent by a sending computer toward a receiving computer. Several flow control mechanisms exist, including TCP flow control, which uses windowing. forward acknowledgment A process used by protocols that do error recovery, in which the number that acknowledges data lists the next data that should be sent, not the last data that was successfully received. Git 507 forwarding plane A synonym for data plane. See also data plane. FTP File Transfer Protocol. An application protocol, part of the TCP/IP protocol stack, used to transfer files between network nodes. FTP is defined in RFC 959. FTP active mode One of two modes of operation for FTP connections (the other being passive mode) that dictates how the FTP data mode connection is established. In active mode, the FTP client listens on a port, it identifies that port to the server, and the server initiates the TCP connection. FTP client An application that can connect to an FTP server for the purpose of transferring copies of files to and from the server. FTP control connection A TCP
connection initiated by an FTP client to an FTP server for the purpose of sending FTP commands that direct the activities of the connection. FTP data connection A TCP connection created by an FTP client and server for the purpose of transferring data. FTP over TLS An FTP standard defined by RFC 4217, also known as FTP Secure (FTPS), which adds a variety of security features to the somewhat insecure original FTP standard (RFC 957), including the addition of the encryption of all data as well as username/password information using Transport Layer Security (TLS). FTP passive mode One of two modes of operation for FTP connections (the other being active mode) that dictates how the FTP data mode connection is established. In passive mode, the FTP client declares the use of passive mode, causing the server to choose and identify a new listening port, with the client establishing a TCP connection to that port. FTP server An application that runs and waits for FTP clients to connect to it over TCP port 21 to support the client’s commands to transfer copies of files to and from the server. FTSP FTP Secure. Common term for FTP over TLS. full mesh From a topology perspective, any topology that has two or more devices, with each device being able to send frames to every other device. G Gateway Load Balancing Protocol (GLBP) A Cisco-proprietary protocol that allows two (or more) routers to share the duties of being the default router on a subnet, with an active/active model, with all routers actively forwarding off-subnet traffic for some hosts in the subnet. Generic Routing Encapsulation (GRE) A protocol, defined in RFC 2784, that defines the headers used when creating a site-to-site VPN tunnel. The protocol defines the use of a normal IP header, called the Delivery Header, and a GRE header that the endpoints use to create and manage traffic over the GRE tunnel. Git An open-source version control application, widely popular for version control in software development and for other uses, like managing network device configurations. 508 GitHub GitHub A software-as-a-service application that implements Git. gratuitous ARP An ARP Reply not sent as a reaction to an ARP request message, but rather as a general announcement informing other hosts of the values of the sending (origin) host’s addresses. GRE tunnel A site-to-site VPN idea, in which the endpoints act as if a point-to-point link (the tunnel) exists between the sites, while actually encapsulating packets using GRE standards. greenfield A term that refers to the installation of new equipment for a project rather than adding configuration to existing in-use hardware and software. H host (context: DC) In a virtualized server environment, the term used to refer to one physical server that is running a hypervisor to create multiple virtual machines. Hot Standby Router Protocol (HSRP) A Cisco-proprietary protocol that allows two (or more) routers to share the duties of being the default router on a subnet, with an active/standby model, with one router acting as the default router and the other sitting by waiting to take over that role if the first router fails. HSRP active A Hot Standby Router Protocol (HSRP) state in which the router actively supports the forwarding of off-subnet packets for hosts in that subnet. HSRP standby A Hot Standby Router Protocol (HSRP) state in which the router does not currently support the forwarding of off-subnet packets for hosts in that subnet, instead waiting for the currently active router to fail before taking over that role. HTML Hypertext Markup Language. A simple document-formatting language that uses tags to indicate how a given part of a document should be interpreted by a viewing application, such as a web browser. HTTP Hypertext Transfer Protocol. The protocol used by web browsers and web servers to transfer files, such as text and graphic files. HTTP verb The action defined in an HTTP request message. hub and spoke From a topology perspective, any topology that has a device that can send messages to all other devices (the hub), with one or more spoke devices that can send messages only to the hub. Also called point-to-multipoint. hypertext The name of Intel’s multithreading technology. hypervisor Software that runs on server hardware to create the foundations of a virtualized server environment primarily by allocating server hardware components like CPU core/threads, RAM, disk, and network to the VMs running on the server.

internetworking operating system (IOS) I IANA The Internet Assigned Numbers Authority. An organization that owns the rights to assign many operating numbers and facts about how the global Internet works, including public IPv4 and IPv6 addresses. See also ICANN. ICANN The Internet Corporation for Assigned Names and Numbers. An organization appointed by the IANA to oversee the distributed process of assigning public IPv4 and IPv6 addresses across the globe. imperative policy model A term that describes the approach in traditional networks in which the engineer chooses configuration settings for each control and data plane protocol (the imperative commands) that dictate specifically how the devices act. This model acts in contrast to the newer declarative policy model and intent-based networking (IBN). Infrastructure as a Service (IaaS) A cloud service in which the service consists of a virtual machine that has defined computing resources (CPUs, RAM, disk, and network) and may or may not be provided with an installed OS. ingress tunnel router (ITR) With LISP, the node that receives an unencapsulated message and encapsulates the message. inside global For packets sent to and from a host that resides inside the trusted part of a network that uses NAT, a term referring to the IP address used in the headers of those packets when those packets traverse the global (public) Internet. inside local For packets sent to and from a host that resides inside the trusted part of a network that uses NAT, a term referring to the IP address used in the headers of those packets when those packets traverse the enterprise (private) part of the network. integrity In data transfers, means that the network administrator can determine that the information has not been tampered with in transit. intent-based networking (IBN) An approach to networking in which the system gives the operator the means to express business intent, with the networking system then determining what should be done by the network, activating the appropriate configuration, and monitoring (assuring) the results. intercloud exchange A WAN service that provides connectivity between public cloud providers and their customers so that customers can install and keep the WAN connections, even when migrating from one cloud provider to another. Internet access technology Any technology that an ISP offers that allows its customers to send and receive data to/from the ISP, including serial links, Frame Relay, MPLS, Metro Ethernet, DSL, cable, and fiber Internet. Internet edge customer. The part of the topology of the Internet that sits between an ISP and the ISP's Internet service provider A company or organization that provides Internet services to customers; the company may have a heritage as a telco, WAN service provider, or cable company. internetworking operating system (IOS) See IOS. 509 510 intrusion detection system (IDS) intrusion detection system (IDS) A security function that examines more complex traffic patterns against a list of both known attack signatures and general characteristics of how attacks can be carried out, rating each perceived threat, and reacting to prevent the more significant threats. See also IPS. IOS Cisco operating system software that provides the majority of a router's or switch's features, with the hardware providing the remaining features. IOS feature set A set of related features that can be enabled on a router to enable certain functionality. For example, the Security feature set would enable the ability to have the router act as a firewall in the network. IOS File System (IFS) IOS image A file system created by a Cisco device that uses IOS. A file that contains the IOS. IP Precedence (IPP) In the original definition of the IP header's Type of Service (ToS) byte, the first 3 bits of the ToS byte, used for marking IP packets for the purpose of applying QoS actions. IPS See intrusion prevention system. IPsec The term referring to the IP Security protocols, which is an architecture for providing encryption and authentication services, usually when creating VPN services through an IP network. ISDN Integrated Services Digital Network. A communication protocol offered by telephone companies that permits telephone networks to carry data, voice, and video. Iterative DNS server A DNS server that will answer DNS requests directly but will not take on the extra work to recursively send other DNS messages to find the answer. J JavaScript A programming language popular for building dynamic web pages, commonly used to run scripts on a web client. Jinja2 A text-based language used to define templates, with text plus variables; used by Ansible for templates. jitter The variation in delay experienced by successive packets in a single application flow. JSON (JavaScript Object Notation) A popular data serialization language, originally used with the JavaScript programming language, and popular for use with REST APIs. JSON array A part of a set of JSON text that begins and ends with a matched set of square brackets that contain a list of values. JSON object A part of a set of JSON text that begins and ends with a matched set of curly brackets that contain a set of key:value pairs. low latency queue 511 K–L. key:value pair In software, one variable name (key) and its value, separated by a colon in some languages and data serialization languages. keyboard, video, mouse (KVM) Three components of a typical desktop computer that are typically not included in a modern server because the server is installed and managed remotely. KVM (Red Hat) Kernel-Based Virtual Machine (KVM), a server virtualization/hypervisor product from the Red Hat company. leaf In an ACL network design, a switch that connects to spine switches and to endpoints, but not to other leaf switches, so that the leaf can forward frames from an endpoint to a spine, which then delivers the frame to some other leaf switch. library In software, a collection of programs packaged so that it can be posted as available in a software repository, found by others, and installed as one entity, as a means to make it easier to share code. LISP Locator/Router Separation Protocol. A protocol, defined in RFC 6830, that separates the concepts and numbers used to identify an endpoint (the endpoint identifier) versus identifying the location of the endpoint (routing locator). LISP mapping database With LISP, the table that contains mapped pairs of endpoint identifiers and routing locators. LISP Routing Locator (RLOC) With LISP, a value that identifies the location of an endpoint, typically the address of the egress device. list variable In applications, a single variable whose value is a list of values, rather than a simple value. LLDP Link Layer Discovery Protocol. An IEEE standard protocol (IEEE 802.1AB) that defines messages, encapsulated directly in Ethernet frames so they do not rely on a working IPv4 or IPv6 network, for the purpose of giving devices a means of announcing basic device information to other devices on the LAN. It is a standardized protocol similar to Cisco Discovery Protocol (CDP). local loop A line from the premises of a telephone subscriber to the telephone company CO. local username A username (with matching password), configured on a router or switch. It is considered local because it exists on the router or switch, and not on a remote server. log message A message generated by any computer, but including Cisco routers and switches, for which the device OS wants to notify the owner or administrator of the device about some event. loss A reference to packets in a network that are sent but do not reach the destination host. low latency queue In Cisco queuing systems, a queue from which the queue scheduling algorithm always takes packets next if the queue holds any packets. This scheduling choice means that packets in this queue spend little time in the queue, achieving low delay (latency) as well as low jitter. 512 Low Latency Queuing (LLQ) Low Latency Queuing (LLQ) The name of a queuing system that can be enabled on Cisco routers and switches by which messages sensitive to latency and jitter are placed in a queue that is always serviced first, resulting in low latency and jitter for those messages. LTE Literally, Long Term Evolution, but this term is used as a word itself to represent the type of wireless 4G technology that allows faster speeds than the original 4G specifications. M malware Malicious software. Management Information Base (MIB) The data structures defined by SNMP to define a hierarchy (tree) structure with variables at the leaves of the tree, so that SNMP messages can reference the variables. management plane Functions in networking devices and controllers that control the devices themselves but that do not impact the forwarding behavior of the devices like control plane protocols do. man-in-the-middle attack An attack where an attacker manages to position a machine on the network such that it is able to intercept traffic passing between target hosts. marking The process of changing one of a small set of fields in various network protocol headers, including the IP header's DSCP field, for the purpose of later classifying a message based on that marked value. markup language A language that provides conventions to tag text to identify the type of text, which allows application of different treatments to different types of text. match/action logic The basic logic done by a networking element: to receive incoming messages, to match fields in the message, to then use logic based on those matches to take action against the message, and to then forward the message. MD5 hash A specific mathematical algorithm intended for use in various security protocols. In the context of Cisco routers and switches, the devices store the MD5 hash of certain passwords, rather than the passwords themselves, in an effort to make the device more secure. Metro Ethernet The original term used for WAN service
that used Ethernet links as the access link between the customer and the service provider. MIB See Management Information Base. MIB view A concept in SNMPv3 that identifies a subset of an SNMP agent's MIB for the purpose of limiting access to some parts of the MIB to certain SNMP managers. mitigation technique A method to counteract or prevent threats and malicious activity. modem Modulator-demodulator. A device that converts between digital and analog signals so that a computer may send data to another computer using analog telephone lines. At the source, a modem converts digital signals to a form suitable for transmission over analog communication facilities. At the destination, the analog signals are returned to their digital form. NAT overload 513 MPLS See Multiprotocol Label Switching. MPLS experimental bits A 3-bit field in the MPLS label used for QoS marking. MPLS VPN A WAN service that uses MPLS technology, with many customers connecting to the same MPLS network, but with the VPN features keeping each customer's traffic separate from others. MTU Maximum transmission unit. The maximum packet size, in bytes, that a particular interface can handle. multifactor authentication authenticates users. A technique that uses more than one type of credential to multipoint A topology with more than two devices in it (in contrast to a point-to-point topology, which has exactly two devices). Without any further context, the term multipoint does not define whether all devices in the topology can send messages directly to each other (full mesh) or not (partial mesh). Multiprotocol BGP (MPBGP) A particular set of BGP extensions that allows BGP to support multiple address families, which when used to create an MPLS VPN service gives the SP the method to advertise the IPv4 routes of many customers while keeping those route advertisements logically separated. Multiprotocol Label Switching (MPLS) A WAN technology used to create an IP-based service for customers, with the service provider's internal network performing forwarding based on an MPLS label rather than the destination IP address. multithreading In computer architecture, a process of maximizing the use of a processor core by sharing an individual core among multiple programs, taking advantage of the typical idle times for the core while it waits on various other tasks like memory reads and writes. N name resolution The process by which an IP host discovers the IP address associated with a host name, often involving sending a DNS request to a DNS server, with the server supplying the IP address used by a host with the listed host name. name server addresses. A server connected to a network that resolves network names into network named access list An ACL that identifies the various statements in the ACL based on a name rather than a number. NAT Network Address Translation. A mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet, by translating those addresses into public addresses in the globally routable address space. NAT overload Another term for Port Address Translation (PAT). One of several methods of configuring NAT, in this case translating TCP and UDP flows based on port numbers in addition to using one or only a few inside global addresses. 514 National Institute of Standards and Technology (NIST) National Institute of Standards and Technology (NIST) A U.S. federal agency that develops national standards, including standards for cloud computing. NBI See northbound API. Nest In JSON, the concept that values can contain objects and arrays so that each object can contain other objects and arrays in a myriad of combinations. Network Based Application Recognition (NBAR) A Cisco router feature that looks at message details beyond the Layer 2, 3, and 4 headers to identify over 1000 different classifications of packets from different applications. Network Management System (NMS) SNMP and other protocols. Software that manages the network, often using Network Time Protocol (NTP) A protocol used to synchronize time-of-day clocks so that multiple devices use the same time of day, which allows log messages to be more easily matched based on their timestamps. Next-generation firewall (NGFW) A firewall device with advanced features, including the ability to run many related security features in the same firewall device (IPS, malware detection, VPN termination), along with deep packet inspection with Application Visibility and Control (AVC) and the ability to perform URL filtering versus data collected about the reliability and risk associated with every domain name. Next-generation IPS (NGIPS) An IPS device with advanced features, including the capability to go beyond a comparison to known attack signatures to also look at contextual data, including the vulnerabilities in the current network, the capability to monitor for new zero-day threats, with frequent updates of signatures from the Cisco Talos security research group. Nexus 1000v A Cisco Nexus data center switch that runs as a software-only virtual switch inside one host (one hardware server), to provide switching features to the virtual machines running on that host. NMS Network Management Station. The device that runs network management software to manage network devices. SNMP is used in the network management protocol used between the NMS and the managed device. northbound API In the area of SDN, a reference to the APIs that a controller supports that gives outside programs access to the services of the controller; for instance, to supply information about the network or to program flows into the network. Also called a northbound interface. northbound interface Another term for northbound API. See also northbound API. notification community An SNMP community (a value that acts as a password), defined on an SNMP manager, which then must be supplied by any SNMP agent that sends the manager any unsolicited SNMP notifications (like SNMP Trap and Notify requests). NTP client Any device that attempts to use the Network Time Protocol (NTP) to synchronize its time by adjusting the local device's time based on NTP messages received from a server. OpenFlow 515 NTP client/server mode A mode of operation with the Network Time Protocol (NTP) in which the device acts as both an NTP client, synchronizing its time with some servers, and as an NTP server, supplying time information to clients. NTP primary server A term defined in NTP RFCs 1305 and 5905 to refer to devices that act as NTP servers alone, with a stratum 1 external clock source. NTP secondary server A term defined in NTP RFCs 1305 and 5905 to refer to devices that act as NTP clients and servers, synchronizing as a client to some NTP server, and then acting as an NTP server for other NTP clients. NTP server Any device that uses Network Time Protocol (NTP) to help synchronize time-of-day clocks for other devices by telling other devices its current time. NTP synchronization The process with the Network Time Protocol (NTP) by which different devices send messages, exchanging the devices' current time-of-day clock information and other data, so that some devices adjust their clocks to the point that the time-of-day clocks list the same time (often accurate to at least the same second). NVRAM Nonvolatile RAM. A type of random-access memory (RAM) that retains its contents when a unit is powered off. O ODL. See OpenDaylight. OID Object identifier. Used to uniquely describe an MIB variable in the SNMP database. This is a numeric string that identifies the variable uniquely and also describes where the variable exists in the MIB tree structure. on-demand self-service One of the five key attributes of a cloud computing service as defined by NIST, referring to the fact that the consumer of the service can request the service, with the service being created without any significant delay and without waiting on human intervention. one-way delay The elapsed time from sending the first bit of data at the sending device until the last bit of that data is received on the destination device. ONF See Open Networking Foundation. on-premises An alternate term for private cloud. See also private cloud. Open Networking Foundation A consortium of SDN users and vendors who work together to foster the adoption of open SDN in the marketplace. OpenDaylight An open-source SDN controller, created by an open-source effort of the OpenDaylight project under the Linux foundation, built with the intent to have a common SDN controller code base from which vendors could then take the code and add further features and support to create SDN controller products. OpenFlow The open standard for Software-Defined Networking (SDN) as defined by the Open Networking Foundation (ONF), which defines the OpenFlow protocol as well as the concept of an abstracted OpenFlow virtual switch. 516 operational management operational management A component of network management focused on extracting data about the network from the network devices, analyzing that data, and providing the data to operations staff. OpFlex The southbound protocol used by the Cisco ACI controller and the switches it controls. ordered data transfer A networking function, included in TCP, in which the protocol defines how the sending host should number the data transmitted, defines how the receiving device should attempt to reorder the data if it arrives out of order, and specifies to discard the data if it cannot be delivered in order. origin hardware address In both an ARP request and reply message, the field intended to be used to list the sender (origin) device's hardware address, typically an Ethernet LAN address. origin IP address In both an ARP request and reply message, the field intended to be used to list the sender (origin) device's IP address. outside global With source NAT, the one address used by the host that resides outside
the enterprise, which NAT does not change, so there is no need for a contrasting term. overlay In SDA, the combination of VXLAN tunnels between fabric edge nodes as a data plane for forwarding frames, plus LISP for the control plane for the discovery and registration of endpoint identifiers. P partial mesh A network topology in which more than two devices could physically communicate, but by choice, only a subset of the pairs of devices connected to the network is allowed to communicate directly. password guessing An attack where a malicious user simply makes repeated attempts to guess a user's password. per-hop behavior (PHB) The general term used to describe the set of QoS actions a device can apply to a message from the time it enters a networking device until the device forwards the message. PHBs include classification, marking, queuing, shaping, policing, and congestion avoidance. permit An action taken with an ACL that implies that the packet is allowed to proceed through the router and be forwarded. pharming cious site. An attack that compromises name services to silently redirect users toward a mali-phishing An attack technique that sends specially crafted emails to victims in the hope that the users will follow links to malicious websites. Platform as a Service (PaaS) A cloud service intended for software developers as a development platform, with a variety of tools useful to developers already installed so that developers can focus on developing software rather than on creating a good development environment. power class 517 PoE Power over Ethernet. Both a generalized term for any of the standards that supply power over an Ethernet link, as well as a specific PoE standard as defined in the IEEE 802.3af amendment to the 802.3 standard. point of presence (PoP) A term used for a service provider's (SP) perspective to refer to a service provider's installation that is purposefully located relatively near to customers, with several spread around major cities, so that the distance from each customer site to one of the SP's PoPs is short. point-to-multipoint See hub and spoke. point-to-point From a topology perspective, any topology that has two and only two devices that can send messages directly to each other. policing A QoS tool that monitors the bit rate of the messages passing some point in the processing of a networking device, so that if the bit rate exceeds the policing rate for a period of time, the policer can discard excess packets to lower the rate. policing rate The bit rate at which a policer compares the bit rate of packets passing through a policing function, for the purpose of taking a different action against packets that conform (are under) to the rate versus those that exceed (go over) the rate. policy model In both ACI and other intent-based networks (IBNs), the operational conventions (model) that combine policies of what the network will provide to grouped sets of network endpoints (endpoint groups) to create a contract for what the network will provide. port (Multiple definitions) (1) In TCP and UDP, a number that is used to uniquely identify the application process that either sent (source port) or should receive (destination port) data. (2) In LAN switching, another term for switch interface. Port Address Translation (PAT) A NAT feature in which one inside global IP address supports over 65,000 concurrent TCP and UDP connections. port number A field in a TCP or UDP header that identifies the application that either sent (source port) or should receive (destination port) the data inside the data segment. port security A Cisco switch feature in which the switch watches Ethernet frames that come in an interface (a port), tracks the source MAC addresses of all such frames, and takes a security action if the number of different such MAC addresses is exceeded. port-scanner Jargon that refers to a security vulnerability during the time between the day in which the vulnerability was discovered, until the vendor or open-source group responsible for that software can develop a fix and make it public. power budget With PoE, data and calculations about the amount of power expected to be used by the various powered devices (PDs), the numbers of devices expected to connect to each switch, versus the amount of power available to PoE based on the capacity of the power supplies in the switches. power class In various PoE standards, a designation that can be sensed/identified via different discovery processes, with the class defining the maximum amount of power the powered device (PD) would like to receive over the Ethernet link. 518 Power over Ethernet (PoE) Power over Ethernet (PoE) Both a generalized term for any of the standards that supply power over an Ethernet link and a specific PoE standard as defined in the IEEE 802.3af amendment to the 802.3 standard. Power over Ethernet Plus (PoE+) A specific PoE standard as defined in the IEEE 802.3at amendment to the 802.3 standard, which uses two wire pairs to supply power with a maximum of 30 watts as supplied by the PSE. power sourcing equipment (PSE) With any Power over Ethernet standard, a term that refers to the device supplying the power over the cable, which is then used by the powered device (PD) on the other end of the cable. powered device (PD) With any Power over Ethernet standard, a term that refers to the device that receives or draws its power over the Ethernet cable, with the power being supplied by the power sourcing equipment (PSE) on the other end of the cable. Priority Code Point (PCP) The formal term for the 3-bit field in the 802iQ header intended for marking and classifying Ethernet frames for the purposes of applying QoS actions. Another term for Class of Service (CoS). priority queue In Cisco queuing systems, another term for a low latency queue (LLQ). private cloud A cloud computing service in which a company provides its own IT services to internal customers inside the same company but by following the practices defined as cloud computing. private IP network Any of the IPv4 Class A, B, or C networks as defined by RFC 1918, intended for use inside a company but not used as public IP networks. private key A secret value used in public/private key encryption systems. Either encrypts a value that can then be decrypted using the matching public key, or decrypts a value that was previously encrypted with the matching public key. programmable network A computer network which provides programmatic interfaces that allow automation applications to change and interrogate the configuration of network devices. provider edge (PE) A term used by service providers, both generally and also specifically in MPLS VPN networks, to refer to the SP device in a point of presence (PoP) that connects to the customer's network and therefore sits at the edge of the SP's network. public cloud A cloud computing service in which the cloud provider is a different company than the cloud consumer. public key A publicly available value used in public/private key encryption systems. Either encrypts a value that can then be decrypted using the matching private key, or decrypts a value that was previously encrypted with the matching private key. pull model With configuration management tools, a practice by which an agent representing the device requests configuration data from the centralized configuration management tool, in effect pulling the configuration to the device. Puppet A popular configuration management application, which can be used with or without a server, using a pull model in which agents request details and pull configuration into devices, with the capability to manage network device configurations. read-write community 519 Puppet manifest A human-readable text file on the Puppet master, using a language defined by Puppet, used to define the desired configuration state of a device. Puppet master Another term for Puppet server. See also Puppet server. Puppet server The Puppet software that collects all the configuration files and other files used by Puppet from different Chef users and then communicates with Puppet agents (devices) so that the agents can synchronize their configurations. Push model With configuration management tools, a practice by which the centralized configuration management tool software initiates the movement of configuration from that node to the device that will be configured, in effect pushing the configuration to the device. Python A programming language popular as a first language to learn and also popular for network automation tasks. Python dictionary pairs. Python list A Python variable like a JSON dictionary, containing a set of key:value A Python variable like a JSON array, containing a list of values. Q–R Quality of Experience (QoE) The users' perception of the quality of their experience in using applications in the network. Quality of Service (QoS) The performance of a message, or the messages sent by an application, in regard to the bandwidth, delay, jitter, or loss characteristics experienced by the message(s). queuing The process by which networking devices hold packets in memory while waiting on some constrained resource; for example, when waiting for the outgoing interface to become available when too many packets arrive in a short period of time. RADIUS A security protocol often used for user authentication, including being used as part of the IEEE 802.lx messages between an 802.lx authenticator (typically a LAN switch) and a AAA server. RAM Random-access memory. A type of volatile memory that can be read and written by a microprocessor. rapid elasticity One of the five key attributes of a cloud computing service as defined by NIST, referring to the fact that the cloud service reacts to requests for new services quickly, and it expands (is elastic) to the point of appearing to be a limitless resource. read-only community An SNMP community (a value that acts as a password), defined on an SNMP agent, which
then must be supplied by any SNMP manager that sends the agent any messages asking to learn the value of a variable (like SNMP Get and GetNext requests). read-write community An SNMP community (a value that acts as a password), defined on an SNMP agent, which then must be supplied by any SNMP manager that sends the agent any messages asking to set the value of a variable (like SNMP Set requests). 520 reconnaissance attack reconnaissance attack An attack crafted to discover as much information about a target organization as possible; the attack can involve domain discovery, ping sweeps, port scans, and so on. recursive DNS server A DNS server that, when asked for information it does not have, performs a repetitive (recursive) process to ask other DNS servers in sequence, hoping to find the DNS server that knows the information. reflection attack An attack that uses spoofed source addresses so that a destination machine will reflect return traffic to the attack's target; the destination machine is known as the reflector. remote access VPN A VPN for which one endpoint is a user device, such as a phone, tablet, or PC, typically created dynamically, and often using TLS. Also called a client VPN. Representational State Transfer (REST) A type of API that allows two programs that reside on separate computers to communicate, with a set of six primary API attributes as defined early in this century by its creator, Roy Fielding. The attributes include client/server architecture, stateless operation, cachability, uniform interfaces, layered, and code-on-demand. resource pooling One of the five key attributes of a cloud computing service as defined by NIST, referring to the fact that the cloud provider treats its resources as a large group (pool) of resources that its cloud management systems then allocate dynamically based on self-service requests by its customers. REST See Representational State Transfer. REST API Any API that uses the rules of Representational State Transfer (REST). RESTful API A term of phrase that means that the API uses REST rules. RFC Request For Comments. A document used as the primary means for communicating information about the TCP/IP protocols. Some RFCs are designated by the Internet Architecture Board (IAB) as Internet standards, and others are informational. RFCs are available online from numerous sources, including www.rfc-editor.org. root DNS server A small number of DNS servers worldwide that provide name resolution for the root zone of DNS, providing information about servers that know details about toplevel domains (TLDs) such as .com, .org, .edu, and so on. round robin A queue scheduling algorithm in which the scheduling algorithm services one queue, then the next, then the next, and so on, working through the queues in sequence. Round Trip Time (RTT) The time it takes a message to go from the original sender to the receiver, plus the time for the response to that message to be sent back. round-trip delay The elapsed time from sending the first bit of data at the sending device until the last bit of that data is received on the destination device, plus the time waiting for the destination device to form a reply, plus the elapsed time for that reply message to arrive back to the original sender. route redistribution A method by which two routing protocol processes running in the same device can exchange routing information, thereby causing a route learned by one routing protocol to then be advertised by another. shaping rate routed access layer A design choice in which all the switches, including the access layer switches that connect directly to endpoint devices, all use Layer 3 switching so that they route packets. Router on a Stick (ROAS) Jargon to refer to the Cisco router feature of using VLAN trunking on an Ethernet interface, which then allows the router to route packets that happen to enter the router on that trunk and then exit the router on that same trunk, just on a different VLAN. S SBI See Southbound API. scalable group security access. In SDA, the concept of a set of related users that should have the equivalent scalable group tag (SGT) In SDA, a value assigned to the users in the same security group. Secure Shell (SSH) A TCP/IP application layer protocol that supports terminal emulation between a client and server, using dynamic key exchange and encryption to keep the communications private. Secure Sockets Layer (SSL) A deprecated security protocol that was formerly used to secure networks and was commonly integrated into web browsers to provide encryption and authentication services between the browser and a website. segment (Multiple definitions) (1) In TCP, a term used to describe a TCP header and its encapsulated data (also called an L4PDU). (2) Also in TCP, the set of bytes formed when TCP breaks a large chunk of data given to it by the application layer into smaller pieces that fit into TCP segments. (3) In Ethernet, either a single Ethernet cable or a single collision domain (no matter how many cables are used). service provider (SP) A company that provides a service to multiple customers. Used most often to refer to providers of private WAN services and Internet services. See also Internet service provider. session key With encryption, a secret value that is known to both parties in a communication, used for a period of time, which the endpoints use when encrypting and decrypting data. SFTP SSH File Transfer Protocol. A file transfer protocol that assumes a secure channel, such as an encrypted SSH connection, which then provides the means to transfer files over the secure channel. shaping A QoS tool that monitors the bit rate of the messages exiting networking devices, so that if the bit rate exceeds the shaping rate for a period of time, the shaper can queue the packets, effectively slowing down the sending rate to match the shaping rate. shaping rate The bit rate at which a shaper compares the bit rate of packets passing through the shaping function, so that when the rate is exceeded, the shaper enables the queuing of packets, resulting in slowing the bit rate of the collective packets that pass through the shaper, so the rate of bits getting through the shaper does not exceed the shaping rate. 521 522 shared key shared key A reference to a security key whose value is known (shared) by both the sender and receiver. shared port With 802.lw RSTP, a port type that is determined by the fact that the port uses half duplex, which could then imply a shared LAN as created by a LAN hub. Simple Network Management Protocol (SNMP) An Internet standard protocol for managing devices on IP networks. It is used mostly in network management systems to monitor network-attached devices for conditions that warrant administrative attention. simple variable In applications, a variable that has a single value of a simple type, such as text and integer or floating-point numbers. single point of failure In a network, a single device or link that, if it fails, causes an outage for a given population of users. site-to-site VPN The mechanism that allows all devices at two different sites to communicate securely over some unsecured network like the Internet, by having one device at each site perform encryption/decryption and forwarding for all the packets sent between the sites. sliding windows For protocols such as TCP that allow the receiving device to dictate the amount of data the sender can send before receiving an acknowledgment—a concept called a window—a reference to the fact that the mechanism to grant future windows is typically just a number that grows upward slowly after each acknowledgment, sliding upward. SNMP See Simple Network Management Protocol. SNMP agent Software that resides on the managed device and processes the SNMP messages sent by the Network Management Station (NMS). SNMP community A simple password mechanism in SNMP in which either the SNMP agent or manager defines a community string (password), and the other device must send that same password value in SNMP messages, or the messages are ignored. See also read-only community, read-write community, and notification community. SNMP Get Message used by SNMP to read from variables in the MIB. SNMP Inform An unsolicited SNMP message like a Trap message, except that the protocol requires that the Inform message needs to be acknowledged by the SNMP manager. SNMP manager Typically a Network Management System (NMS), with this term specifically referring to the use of SNMP and the typical role of the manager, which retrieves status information with SNMP Get requests, sets variables with the SNMP Set requests, and receives unsolicited notifications from SNMP agents by listening for SNMP Trap and Notify messages. SNMP Set SNMP message to set the value in variables of the MIB. These messages are the key to an administrator configuring the managed device using SNMP. SNMP Trap An unsolicited SNMP message generated by the managed device, and sent to the SNMP manager, to give information to the manager about some event or because a measurement threshold has been passed. SNMPv2c A variation of the second version of SNMP. SNMP Version 2 did not originally support communities; the term SNMPv2c refers to SNMP version 2 with support added for SNMP communities (which were part of SNMPv). spoofing attack 523 SNMPv3 The third version of SNMP, with the notable addition of several security features as compared to SNMPv2c, specifically message integrity, authentication, and encryption. social engineering Use information. Attacks that leverage human trust and social behaviors to divulge sensi- Software as a Service (SaaS) A cloud service in which the service consists of access to working software, without the need to be concerned about the details of installing and maintaining the software or the servers on which it runs. Software-Defined Access networks. Cisco's intent-based networking (IBN) offering for enterprise software-defined
architecture In computer networking, any architecture that provides mechanisms for automated software control of the network components, typically using a controller. Any architecture that leads to a Software-Defined Network (SDN). Software-Defined Networking (SDN) A branch of networking that emerged in the marketplace in the 2010s characterized by the use of a centralized software controller that takes over varying amounts of the control plane processing formerly done inside networking devices, with the controller directing the networking elements as to what forwarding table entries to put into their forwarding tables. SOHO A classification of a business site with a relatively small number of devices, sometimes in an employee office in their home. Source NAT The type of Network Address Translation (NAT) used most commonly in networks (as compared to destination NAT), in which the source IP address of packets entering an inside interface is translated. southbound API In the area of SDN, a reference to the APIs used between a controller and the network elements for the purpose of learning information from the elements and for programming (controlling) the forwarding behavior of the elements. Also called a southbound interface. southbound interface spear phishing connection. Another term for southbound API. See also southbound API. phishing that targets a group of users who share a common interest or spine In an ACI network design for a single site, a switch that connects to leaf switches only, for the purpose of receiving frames from one leaf switch and then forwarding the frame to some other leaf switch. spine-leaf network A single-site network topology in which endpoints connect to leaf switches, leaf switches connect to all spine switches (but not to other leaf switches), and spine switches connect to all leaf switches (but not to other spine switches). The resulting topology results in predictable switching paths with three switches between any two endpoints that connect to different leaf switches. spoofing attack A type of attack in which parameters such as IP and MAC addresses are spoofed with fake values to disguise the sender. 524 spurious DHCP server A DHCP server that is used by an attacker for attacks that take advantage of DHCP protocol messages. SSL See Secure Sockets Layer. standard access list A list of IOS global configuration commands that can match only a packet's source IP address for the purpose of deciding which packets to discard and which to allow through the router. star topology A network topology in which endpoints on a network are connected to a common central device by point-to-point links. stateful A protocol or process that requires information stored from previous transactions to perform the current transaction. stateless A protocol or process that does not use information stored from previous transactions to perform the current transaction. subinterface One of the virtual interfaces on a single physical interface. switch abstraction The fundamental idea of what a switch does, in generalized form, so that standards protocols and APIs can be defined that then program a standard switch abstraction; a key part of the OpenFlow standard. syslog A server that takes system messages from network devices and stores them in a database. The syslog server also provides reporting capabilities on these system messages. Some syslog servers can even respond to select system messages with certain actions such as emailing and paging. syslog server A server application that collects syslog messages from many devices over the network and provides a user interface so that IT administrators can view the log messages to troubleshoot problems. T T1 A line from the telco that allows transmission of data at 1.544 Mbps, with the capability to treat the line as 24 different 64-Kbps DS0 channels (plus 8 Kbps of overhead). T3 A line from the telco that allows transmission of data at 44.736 Mbps, with the capability to treat the line as 28 different 1.544-Mbps DS1 (T) channels, plus overhead. TCACAS+ A security protocol often used for user authentication as well as authorization and accounting, often used to authenticate users who log in to Cisco routers and switches. tail drop Packet drops that occur when a queue fills, another message arrives that needs to be placed into the queue, and the networking device tries to add the new message to the tail of the queue but finds no room in the queue, resulting in a dropped packet. target hardware address In both an ARP request and reply message, the field intended to be used to list the destination (target) device's hardware address, typically an Ethernet LAN address. This field is left as all binary Os for the typical ARP request messages. top-level domain (TLD) target IP address In both an ARP request and reply message, the field intended to be used to list the destination (target) device's IP address. TCAM See ternary content-addressable memory. TCP Transmission Control Protocol. TCP Transmission Control Protocol. A connection-oriented transport layer TCP/IP protocol that provides reliable data transmission. TCP window The mechanism in a TCP connection used by each host to manage how much data the receiver allows the sender to send to the receiver. TCP/IP Transmission Control Protocol/Internet Protocol. A common name for the suite of protocols developed by the U.S. Department of Defense in the 1970s to support the construction of worldwide internetworks. TCP and IP are the two best-known protocols in the suite. telco A common abbreviation for telephone company. ternary content-addressable memory (TCAM) A type of physical memory, either in a separate integrated circuit or built into an ASIC, that can store tables and then be searched against a key, such that the search time happens quickly and does not increase as the size of the table increases. TCAMs are used extensively in higher-performance networking devices as the means to store and search forwarding tables in Ethernet switches and higher-performance routers. TFTP Trivial File Transfer Protocol. An application protocol that allows files to be transferred from one computer to another over a network, but with only a few features, making the software require little storage space. TFTP client An application that can connect to a TFTP server for the purpose of transferring copies of files to and from the server. TFTP server An application that runs and waits for TFTP clients to connect to it over UDP port 69 to support the client's commands to transfer copies of files to and from the server. threat An actual potential to use an exploit to take advantage of a vulnerability. three-tier design See core design. time interval (shaper) Part of the internal logic used by a traffic shaping function, which defines a short time period in which the shaper sends packets until a number of bytes are sent, and then the shaper stops sending for the rest of the time interval, with a goal of averaging a defined bit rate of sending data. TLD DNS server A DNS server with the role of identifying the IP address of the authoritative DNS server for a domain that resides within its top-level domain. Top of Rack (ToR) switch In a traditional data center design with servers in multiple racks and the racks in multiple rows, a switch placed in the top of the rack for the purpose of providing physical connectivity to the servers (hosts) in that rack. top-level domain (TLD) With DNS name services, the top-level domain is the most significant (rightmost) of the period-separated values in a DNS host name—for example, the .com within host name www.example.com. 525 526 Transport Layer Security (TLS) Transport Layer Security (TLS) A security standard that replaced the older Secure Sockets Layer (SSL) protocol, providing functions such as authentication, confidentiality, and message integrity over reliable in-order data streams like TCP. trojan horse Malware that is hidden and packaged inside other legitimate software. trust boundary When thinking about a message as it flows from the source device to the destination device, the trust boundary is the first device the message reaches for which the QoS markings in the message's various headers can be trusted as having an accurate value, allowing the device to apply the correct QoS actions to the message based on the marking. trusted port With both the DHCP snooping and Dynamic ARP Inspection (DAI) switch features, the concept and configuration setting that tells the switch to allow all incoming messages of that respective type, rather than to consider the incoming messages (DHCP and ARP, respectively) for filtering. tunnel interface A virtual interface in a Cisco router used to configure a variety of features, including Generic Routing Encapsulation (GRE), which encapsulates IP packets into other IP packets for the purpose of creating VPNs. two-tier design See collapsed core design. Type of Service (ToS) In the original definition of the IP header, a byte reserved for the purpose of QoS functions, including holding the IP Precedence field. The ToS byte was later repurposed to hold the DSCP field. U UDP User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery. uncacheable For resources that might be repeatedly requested over time, an attribute that means that the requesting host should not use its local copy of the resource, but instead ask for a new copy every time the resource is requested. underlay In SDA, the network devices and links that create basic IP connectivity to support the creation of VXLAN tunnels for the overlay. Unified Computing System (UCS) The Cisco brand name for its server hardware products. Universal Power over Ethernet (UpoE) A specific PoE standard as defined in the IEEE 802.3bt amendment to the 802.3 standard, which uses four wire pairs to supply
power with a maximum of 100 watts as supplied by the PSE. untrusted port With both the DHCP snooping and Dynamic ARP Inspection (DAI) switch features, the concept and configuration setting that tells the switch to analyze each incoming message of that respective type (DHCP and ARP) and apply some rules to decide whether to discard the message. virtual MAC address (vMAC) 527 UpoE Universal Power over Ethernet. A specific PoE standard as defined in IEEE 802.3bt amendment to the 802.3 standard, which uses four wire pairs to supply power with a maximum of 60 watts as supplied by the PSE. Universal Power over Ethernet Plus (UpoE+) A specific PoE standard as defined in the IEEE 802.3bt amendment to the 802.3 standard, which uses four wire pairs to supply power with a maximum of 100 watts as supplied by the PSE. uniform Resource Identifier. The formal and correct term for the formatted text used to refer to objects in an IP network. This text is commonly called a URL or a web address. For example, is a URI that identifies the protocol (HTTP), host name (www.certskills.com), and web page (blog). URI parameters See URI query (parameters). URI path (resource) In a URI, the part that follows the first /, up to the query field (which begins with a ?), which identifies the resource in the context of a server. URI query (parameters) In a URI, the part that follows the first ?, which provides a place to list variable names and values as parameters. URI resource See URI path (resource). URI Uniform Resource Locator. The widely popular terms for the formatted text used to refer to objects in an IP network. For example, is a URI that identifies the protocol (HTTP), host name (www.certskills.com), and web page (blog). user network interface (UNI) A term used in a variety of WAN standards, including carrier/Metro Ethernet, that defines the standards for how a customer device communicates with a service provider's device over an access link. username secret A reference to the password configured on the username name secret pass-value command, which defines a username and an encoded password, used to build a local username/password list on the router or switch. V variable In applications, a method to assign a name to a value so that the application can refer to the value, change it, compare it to other values, apply logic, and perform other actions typical of software applications. version control software Applications that monitor files for changes, tracking each specific change, the user, the date/time, with tools so that users can compare versions of each file through its history to see the differences. violation mode In port security, a configuration setting that defines the specific set of actions to take on a port when a port security violation occurs. The modes are shutdown, restrict, and protect. virtual CPU (vCPU) In a virtualized server environment, a CPU (processor) core or thread allocated to a virtual machine (VM) by the hypervisor. virtual IP address For any FHRP protocol, an IP address that the FHRP shares between multiple routers so that they appear as a single default router to hosts on that subnet. virtual MAC address (vMAC) For any FHRP protocol, a MAC address that the FHRP uses to receive frames from hosts. 528 virtual machine virtual machine An instance of an operating system, running on server hardware that uses a hypervisor to allocate a subset of the server hardware (CPU, RAM, disk, and network) to that VM. virtual network function (VNF) Any function done within a network (for example, router, switch, firewall) that is implemented not as a physical device but as an OS running in a virtualized system (for instance, a VM). virtual network identifier (VNID) In SDA and VXLAN, the identifier for a separate routing and switching instance. All devices in the same VNID are considered to be allowed to send data to each other unless prevented from doing so by other security mechanisms. virtual NIC (vNIC) In a virtualized server environment, a network interface card (NIC) used by a virtual machine, which then connects

to some virtual switch (vSwitch) running on that same host, which in turn connects to a physical NIC on the host. virtual private network (VPN) A set of security protocols that, when implemented by two devices on either side of an unsecured network such as the Internet, can allow the devices to send data securely. VPNs provide privacy, device authentication, anti-replay services, and data integrity services. Virtual Router Redundancy Protocol (VRRP) A TCP/IP RFC protocol that allows two (or more) routers to share the duties of being the default router on a subnet, with an active/standby model, with one router acting as the default router and the other sitting by waiting to take over that role if the first router fails. virtual switch (vSwitch) A software-only virtual switch inside one host (one hardware server), to provide switching features to the virtual machines running on that host. virus Malware that injects itself into other applications and then propagates through user intervention. VPN See virtual private network. VPN client Software that resides on a PC, often a laptop, so that the host can implement the protocols required to be an endpoint of a VPN. vulnerability A weakness that can be used to compromise security. VXLAN Virtual Extensible LAN. A flexible encapsulation protocol used for creating tunnels (overlays). W WAN edge The device (typically a router) at enterprise sites that connects to private WAN links, therefore sitting at the edge of the WAN. WAN link Another term for leased line. WAN service provider A company that provides private WAN services to customers; the company may have a heritage as a telco or cable company. zero-day vulnerability 529 watering hole attack An attack where a site frequently visited by a group of users is compromised, when the target users visit the site, they will be infected with malware, but other users will not. web server Software, running on a computer, that stores web pages and sends those web pages to web clients (web browsers) that request the web pages. well-known port A TCP or UDP port number reserved for use by a particular application. The use of well-known ports allows a client to send a TCP or UDP segment to a server, to the correct destination port for that application. whaling A phishing technique that targets high-profile individuals to follow links to malicious sites. wildcard mask commands. The mask used in Cisco IOS ACL commands and OSPF and EIGRP network window Represents the number of bytes that can be sent without receiving an acknowledgment. worm Malware that propagates from one system to another, infecting as it goes, all autonomously. write community See read-write community. X–Y–Z XML (eXtensible Markup Language) A markup language that helps enable dynamic web pages; also useful as a data serialization language. YAML (YAML Ain't Markup Language) A data serialization language that can be easily read by humans; used by Ansible. zero-day vulnerability Jargon that refers to a security vulnerability during the time between the day in which the vulnerability was discovered, until the vendor or open-source group responsible for that software can develop a fix and make it public. Index Numbers 2-tier campus design, 291–293 3G wireless, 320 3-tier campus design, 293–295 4G wireless, 320–321 5G wireless, 320 802.1Q headers, 237–238 802.11 headers, 238 access-list 101 command, 60 access-list command, 33–35, 42, 46–50, 54, 62, 397 any keyword, 34 building ACLs with, 39–40 deny keyword, 34 examples and logic explanations, 50 extended numbered ACL configuration commands, 51 log keyword, 38 permit keyword, 31, 34 A AAA (Authentication, Authorization, Accounting), 82–83 aaS (as a Service), 339 reverse engineering from ACL to address range, 40–41 tcp keyword, 48 up keyword, 48 access switches, 291, 295 access Internet, 317–321 accounting (AAA), 82–83 public cloud services, 342–346 ACE (Access Control Entries), 397–398 security physical access control, 84 ACI (Application Centric Infrastructure), 369, 373 user access, 82–83 IBN, 371 user awareness/training, 43 leaf switches, 370 access-class command, 62, 95, 105 access links MetroE, 306 MPLS, 314 spine switches, 370 ACK flags, 12 ACLs (Access Control Lists), 397–398 ARP ACL, 159 classification, 235 comparison of ACL types, 28 controlling Telnet and SSH access with, 95 origin IP addresses, 157–159, 163–164 RELEASE messages, filtering based on IP addresses, 151 IPv4, 204 CIDR, 205–206 deny all statements, 31 dynamic IP address configuration with DHCP, 311 extended numbered ACLs, 46–54 host settings, 133–140 implementation considerations, 59–60 matching addresses, 31–34 location and direction, 26–27 private addressing, 206 matching packets, 27 QoS marking, 237 named ACLs, 54–58 routing, 26, 232 numbered ACLs, 58–59 scalability, 204–205 NAT, 202, 207–223 overview, 26 IPv6, QoS marking, 237 QoS tools, compared, 233 MAC addresses, 109, 113 SDA, 399 NAT, 202, 207–222 SNMP security, 267 private addressing, 206 standard numbered ACLs, 29–41 scalability, 204–205 troubleshooting, 222 spoofing attacks, 72 active mode (FTP), 276 amplification attacks, 75 addresses. See also ACLs DDoS attacks, 75 any/all IP addresses, matching, 34 DoS attacks, 73–74 CIDR, 205–206 inside global, 209 Man-in-the-Middle attacks, 76–77 inside local, 209 reflection attacks, 75 IP addresses AF (Assured Forwarding), 240 commands, 139–140 AF DiffServ RFC (2597), 240 destination IP addresses, 95 AF DSCP value marking, 240 DNS IP addresses, 128 agents, SNMP, 264–267 allocation, DHCP, 129 532 Amazon Web Services (AWS) Amazon Web Services (AWS), 340 simple variables, 410–411 amplification attacks, 75 stateless operation, 410 Ansible, 422, 438–439, 442 RESTful, 366 answering exam questions, 456–457 XML, data serialization, 421–423 anti-replay (Internet VPNs), 321 YAML, data serialization, 422–423 any/all IP addresses, matching, 34 any keyword, 34 AnyConnect Secure Mobility Client, 325 APIs (Application Programming Interfaces), 364 DNA Center, 415 JSON APIC (Application Policy Infrastructure Controller), 372 APIC-EM (Application Policy Infrastructure Controller-Enterprise Module), 373–374 app (application) servers, 371 Application Centric Infrastructure. See ACI arrays, 424–426 Application Programming Interfaces. See APIs beautified JSON, 426 application signatures, 236 data serialization, 418–423 key-value pairs, 423–424 minified JSON, 426 Application-Specific Integrated Circuit (ASIC), 362 objects, 424–426 architectures, SDN, 367–369, 373–375 REST APIs, 418, 422–423 arp -a command, 142 REST, 366 REST APIs, 408 cacheable resources, 410 client/server architecture, 409, 419–420 data structures, 411–412 dictionary variables, 411–412 DNA Center calls, 417–418 HTTP, 413–416 JSON, 422–423 ARP ACL (Address Resolution Protocol Access Control Lists), 159 ARP messages DAI, 156 filtering MAC addresses, 159 logic of, 158 gratuitous ARP as an attack vector, 157–158 origin hardware addresses, 159–160 key-value pairs, 412 arrays (JSON), 424–426 list variables, 411–412 as a Service (-aaS), 339 branch public cloud example 533 ASA (Adaptive Security Appliance) firewall, 96 AUTH command, 279 authentication (AAA), 82–83 ASIC (Application-Specific Integrated Circuit), 362 Assured Forwarding (AF), 240 attacks (security) amplification attacks, 75 ARP messages (gratuitous), 157–158 Internet VPNs, 321 SNMPv3, 208 authentication (AAA), 82–83 automatic allocation, 129 automation configuration automation files, 437 brute-force attacks, 80 buffer overflow attacks, 78 DDoS attacks, 75 network management, 376–378 AVAC (Application Visibility and Control) DHCP-based attacks, 147 NGFW, 101 dictionary attacks, 80 NGIPS, 103 DoS attacks, 73–74 AWS (Amazon Web Services), 340 malware, 78–79 Man-in-the-Middle attacks, 76–77 password guessing, 80 phishing attacks, 79 phishing attacks, 79 reconnaissance attacks, 77–78 reflection attacks, 75 smishing attacks, 79 B bandwidth, managing, 228 batch traffic, 230 beautified JSON, 408 binaries wildcard masks, 33 social engineering attacks, 79 binding tables (DHCP snooping), 150 spear phishing attacks, 79 biometric credentials (security), 81 spoofing attacks, 72–77 blocks (CIDR), 206 Trojan horses, 78 boot system command, 281 viruses, 78 branch offices public cloud example visiting attacks, 79 watering hole attacks, 79 email services traffic flow, 347–349 whaling attacks, 79 Internet connections, 349 worms, 78 private WAN connections, 349 SDA broadcast flags broadcast flags, 125 certificates (digital), security, 81 browsing web chapter reviews (exam preparation), 464 HTTP, 16–17, 20–21 URIs, 17–18 URIs, 17 brute-force attacks, 80 budgeting time (exams), 450–451 checklists (practice exams), 455, 459 Chef, 438, 441–442 CIDR (Classless Interdomain Routing), 205–206 buffer overflow attacks, 78 CIR (Committed Information Rate), 247 C Cisco Discovery Protocol. See CDP cable Internet, 319–320 CAC (Call Admission Control) tools, 245 cacheable resources (REST API), 213, 214 campus LANs Cisco Learning Network, exam preparation, 464 Cisco Prime management products website, 264 Class-Based Weighted Fair Queuing (CBWFQ), 243 Class of Service (CoS) field (802.1Q header), 237 overview, 290 Class Selector (CS), 241 three-tier campus design, 293–295 classification, QoS, 233–234 topology design terminology, 295 clear ip nat translation command, 211, 219, 225 two-tier campus design, 290–293 CBWFQ (Class-Based Weighted Fair Queuing), 243 clear logging command, 179 CDP (Cisco Discovery Protocol) CLI (Command-Line Interface), practicing with (exam preparation), 460–461 configuration, 193–194 discovering information about neighbors, 190–193 verification, 193–194 cdp enable command, 200 clear-text passwords, SNMP, 267 clients NTP, 183–186 VPNs, 325 cdp run command, 200 clock set command, 182–183 CE (Customer Edge), 313 clock summer-time command, 183, 200 centralized configuration files, 432 centralized control planes, 363 commands 535 clock timezone command, 183, 200 copy ftp flash, 274 clear cloud computing, 328, 336 “as a service” model, 339–342 copy running-config startupconfig, 112, 428 cloud services catalogs, 338 copy tftp flash, 271 CSRs, 344 crypto key generate rsa, 105 IaaS, 339–340 debug, 177, 180–181, 201 PaaS, 341–342 debug ip nat, 219, 225 private, 337–338 debug ip nat, 180 public, 337–339, 342–349 deny, 55–57, 62 SaaS, 341 dig, 78 services, 336–337 dir, 272, 282 cloud services catalogs, 338 enable password, 90, 105 Cloud Services Routers (CSRs), 344 enable secret, 90–94 codecs, 321 ifconfig, 134, 137–142 collapsed core design, 290–293 Interface loopback, 200 commands in access-group, 36, 43, 51, 60–62 access-class, 62, 95, 105 ip access-list, 55, 62 access-list, 31–35, 38–51, 54, 62, 397 ip access-list extended, 56 access-list 101, 60 ip address dhcp, 132 arp -a, 142 ip arp inspection validate, 164 AUTH, 279 boot system, 281 ip dhcp snooping information option, 153 cdp enable, 200 ip ftp password, 281 cdp ip username, 281 clear ip nat translation, 211, 219, 225 ip helper-address, 125–127, 141 clear logging, 179 ip nat inside, 213, 215, 220–222 clock set, 182–183 ip nat inside source, 217, 225 clock summer-time, 183, 200 ip nat inside source list pool, 220–222 clock timezone, 183, 200 ip nat inside source static, 213–215, 222 copy, 270–271, 274–275, 282 ip address, 139–140 ip nat, 225 536 commands ip nat pool, 216, 225 no service password-encryption, 90 ip nat pool netmask, 215 no shutdown, 115, 121, 179 ip route configuration, 133 nslookup, 78 ipconfig, 134, 142 ntp master, 183–185, 188, 200 line console, 105 ntp server, 183, 188, 200 line vty, 105 ntp source, 200 ldp holdtime, 198 password, 90, 105 ldp receive, 201 PASV, 278 lldp run, 197, 201 permit, 55–57, 62 ldp timer, 198 PORT, 277–278 lldp transmit, 201 port-security, 111 logging, 200 remark, 55, 62 logging buffered, 175, 179, 200 service password-encryption, 89–90, 105 ip nat outside, 213–215, 220–222 logging buffered warning, 181 logging console, 174, 200 logging host, 175 logging monitor, 175, 200 logging monitor debug, 181 logging trap, 200 logging trap, 4, 181 logon, 105 service sequence-numbers, 200 show access-lists, 35, 43, 56, 62 show arp, 142 show cdp, 193–194, 197–198, 201 show cdp entry interface, 190, 193 show cdp interface, 193–194 show cdp neighbors, 190–195 logon local, 105 show cdp neighbors detail, 190–193 more, 270 show cdp traffic, 193–194 netastr -rn, 135–142 show clock, 201 no cdp enable, 193 show dhcp lease, 131 no enable secret, 105 show flash, 270–272, 282 no ip access-group, 60 show interfaces, 115, 121, no ip dhcp snooping information option, 152–153 show interfaces loopback, 201 no logging console, 177 no logging monitor, 177 show interfaces status, 136–142 show interfaces switchport, 377 show interfaces vian, 131 Committed Information Rate (CIR) 537 show ip access-list, 43, 57, 59 show ip access-lists, 35, 59, 62 show ip arp, 142 show ip arp inspection, 161–163 show ip default-gateway, 132 show ip dhcp conflict, 142 show ip dhcp snooping, 153–155 show ip dhcp snooping binding, 162 show running-config | interface, 121, 167 show running-config command, 35, 89 show startup-config, 270 shutdown, 115, 121, 179, 182 ssh, 95 switchport mode, 120, 167, 377 switchport mode access, 110–111 show ip interface, 36, 43, 130 switchport mode trunk, 110 show ip nat statistics, 215–222, 225 switchport port-security, 110–111 show ip nat translations, 214–225 show lldp, 201 show lldp entry, 196 show lldp interface, 198 show lldp neighbors, 195 show logging, 175–178, 201 show mac address-table dynamic, 113–114, 121, 167 switchport port-security macaddress, 110–111, 120 switchport port-security macaddress sticky, 110–111, 120, 167 switchport port-security maximum, 110, 120 switchport port-security violation, 110, 114, 120 telnet, 95 show mac address-table secure, 113–114, 121 terminal monitor, 175, 181, 201 show mac address-table static, 113, 121 transport input, 105 show ntp associations, 184–186, 201 username, 105 terminal no monitor, 201 transport input ssh command, 89 show ntp status, 184, 201 username password, 94 show port-security, 115–116, 121 username secret, 94 show port-security interface, 112–121 verify, 273, 282 show process cpu, 181 whois, 270 show running-config, 35, 56–59, 105, 121, 167, 270 verify /md5, 273, 282 Committed Information Rate (CIR), 247 538 communities (SNMP) communities (SNMP), 267 Puppet, 438–442 Community-based SNMP Version 2 (SNMPv2c), 267 routers as DHCP clients, 132–133 switches community strings (SNMP), 267 as DHCP clients, 130–132 confidentiality, Internet VPNs, 321 interfaces, 108–113 configuration Syslog, 178–180 ACLs, 34 38 templates, 435–437 Ansible, 438–439, 442 variables, 435–437 automation files, 437 VMs, 334 CDP, 193–194 configure command, 430 centralized configuration files, 432 congestion Chef, 438, 441–442 DAI, 160–165 DHCP, 131 relays, 130 snooping, 152–156 drift, 430–431 extended numbered ACLs, 51–54 IPv4, 131 LDP, 197–198 management, 428–430 avoidance, 250–251 management LLQ, 243–245 multiple queues, 244 prioritization, 242 round robin scheduling, 243 strategy, 245 connectionless protocols, 133 connections connection-oriented protocols, 133 monitoring, 433 establishment and termination (TCP), 12–13 named ACLs, 55–56 public cloud access, 342–346 NAT, 214–222 public cloud branch offices, 349 NTP contextual awareness, NGIPS, 103 client/server, 183–184 control connection (FTP), 277 redundant configuration, 186–188 control plane (networking devices), 360–363 numbered ACLs, 58–59 per-device configuration model, 431 provisioning, 434–435 controllers centralized control, 363 defined, 362 decimal wildcard masks 539 networks, 375–379 message checks, 164–165 NBIs, 365–366 message rate limits, 163–164 OpenDaylight SDN controller, 368 data application traffic, 229–230 data centers (virtual) OSC, 369 networking, 333 SBIs, 364 physical networks, 334–335 copy command, 270–271, 274–275, 282 copy ftp flash command, 274 copy running-config startup-config command, 112, 428 copy tftp flash command, 271 copying IOS images, 271–274 core design, 293–295 CoS (Class of Service) field (802.1Q header), 237–238 vendors, 333 workflow, 335–336 data connection (FTP), 277 data integrity, Internet VPNs, 321 data plane (networking devices), 359–361 data serialization JSON, 418–422 arrays, 424–426 CRUD actions (software), 413–414 beautified JSON, 426 crypto key generate rsa command, 105 key-value pairs command, 105 key-value pairs CS (Class Selector), 241 objects, 424–426 minified JSON, 426 CS DSCP values, marking, 241 XML, 421–423 CSRs (Cloud Services Routers), 344 YAML, 422–423 customer edge (CE), 313 data structures, 411–412 databases DAI (Dynamic ARP Inspection), 156 configuring, 160–165 layer 2 switches, 160–163 logic of, 158 MAC addresses, filtering, 159 MIB, 264–267 signature databases and IPS, 99 DB (Database) servers, 371 DDoS (Distributed Denial-of-Service) attacks, 75 debug command, 177–181, 201 debug ip nat command, 219, 225 debug ip rip command, 180 decimal wildcard masks, 31–32 540 default routers, verification default routers, verification, 136–140 broadcast flags, 125 delay, managing, 229 dynamic allocation, 129 deleting single points of failure, 258–259 information stored at DHCP server, 128 demilitarized zones (DMZ), 98 overview, 124–126 denial of service (DoS) attacks, 97 relays DHCP Relay, 126–127, 130 deny all statements, 31 configuring, 130 deny command, 55–57, 62 supporting, 126–127 deny keyword, 28, 34 troubleshooting, 130 destination IP routers, 128, 132–133 addresses, 95 rules of, 149 matching, 46–48 servers, 128 destination port numbers, 8–9 devices hardening controlling Telnet and SSH access with ACLs, 95 filtering, 146 binding tables, 150 configuring, 152–156 DHCP-based attacks, 147 firewalls, 96–97 DHCP message rate limits, 154–156 management protocols DISCOVER messages, 150 CDP, 190–194 layer 2 switches, 152–154 LLDP, 194–198 logic of, 148–149 NTP, 181–189 RELEASE messages, 151 Syslog, 170–171 networking, 359–363 per-device configuration model, 431 security device hardening, 95–97 IOS passwords, 88–94 DHCP (Dynamic Host Configuration Protocol), 122 advantages of, 124 automatic allocation, 129 static allocation, 129 switches, configuring as DHCP clients, 130–132 troubleshooting, 130 dictionary attacks, 80 dictionary variables, REST APIs, 411–412 Differentiated Services Code Point (DSCP), 234 Eclipse IDE DiffServ DSCP marking values DNS (Domain Name System), 111 AF, 240 DNS IP addresses, 128 CS, 241 DNS IP servers, 128 EF, 240 recursive DNS lookups, 9 dig command, 78 digital certificates (security), 81 digital subscriber lines (DSLs), 318 dir command, 272, 282 direction (ACLs), 26–27 web servers, finding, 18–20 DoS (Denial-of-Service) attacks, 73–74, 97 DSCP (Differentiated Services Code Point), 234 DSCP fields (QoS marking), 238 DISCOVER messages, filtering based on MAC addresses, 150 marking values, 270 DSLs (Digital Subscriber Lines), 318 distributed control planes, 363 DSLAMs (DSL access multiplexers), 318 distribution switches, 291, 295 DMZ (Demilitarized Zones), 98 DNA Center, 384, 389, 395 APIs, 415 IP security, 397–398 network management, 400–401 Path Trace feature, 403 PI, 400–401 REST API calls, 417–418 dynamic allocation, 129 dynamic (ephemeral, private) ports, 9 Dynamic Host Configuration Protocol. See DHCP dynamic IP address configuration, 311 dynamic NAT (Network Address Translation) configuration, 215–217 scalable groups, 396 overview, 210–211 SDA troubleshooting, 222 SGT, 399 user group security, 399–399 SGT, 399 verification, 217–219 dynamic windows, 15–16 topology map, 401–403 traditional management E differences with, 402–403 similarities to, 401 VXLAN tunnels, 399 earplugs (exam preparation), 451 Eclipse IDE, 341 541 424 editing named ACLs EF (Expedited Forwarding), 238 er-disabling recovery, troubleshooting, 117 EF DSCP value marking, 240 error detection, 6 EF RFC (RFC 3246), 240 error recovery, 6, 13–14 EID (Endpoint Identifiers), 392 Ethernet editing named ACLs, 56–58 E-LAN (Ethernet LAN) service, 308, 311 802.1Q headers, 237–238 elasticity, cloud computing, 337 access links, 306 E-Line (Ethernet Line) service, 307–310 IEEE standards, 306 email, public cloud branch office traffic flow, 347–349 enable password command, 90, 105 enable secret command, 90–94 encoding IOS passwords with hashes, 90–94 encryption IOS passwords, 89–90 IPsec, 238 323–324 keys, 323 SNMPv3, 267 802.11 headers, 238 PoE, 297–299 Ethernet LAN (E-LAN) service, 308 Ethernet LANs campus LANs, 290–295 physical standards, 296–297 port security, 108–113 troubleshooting, 115–119 Ethernet Line (E-Line) service, 307–310 Ethernet Tree (E-Tree) service, 309 Ethernet Virtual Connection (EVC), 307 End-to-End QoS Network Design, Second Edition (Cisco Press), 232 Ethernet WANs, public cloud connections, 345 endpoints, EPGs, 371 E-Tree (Ethernet LAN) service, 309 Enterprise QoS Solution Reference Network Design, 232 EVC (Ethernet Virtual Connection), 307 enterprises, classification matching, 234 exact IP addresses, matching, 31 EPGs (Endpoint Groups), 371 ephemeral (dynamic, private) ports, 9 exams chapter reviews, 464 hours, 463 eq 21 parameters, 49 NDAs, 454 er-disabled state, 115 post exam process, 453 failover, HSRP practice exams, 454 checklists, 455, 459 second attempts at passing, 463 PTP questions, 455 self-assessments, 462–463 PTP software, 458–459 VUE testing center, 455 time preparing for budgeting, 450–451 24 failing before the exam, 452 30 minutes before the exam, 452–453 earplugs, 451 time-check method, 451 video tutorials, 449 excluded (reserved) addresses, DHCP servers, 128 one week away preparation, 451–452 Expedited Forwarding (EF), 238 taking notes, 452 extended numbered IPv4 ACLs exploits (security), 72 travel time, 452 configuration, 51–54 questions answering, 456–457 matching protocol, source IP, and destination IP, 46–48 multiple choice questions, 449–450, 457 matching TCP and UDP port numbers, 48–50 Premium Edition questions, 457 overview, 46 PTP questions, 455 similar questions, 450 simulation questions, 449 testlet questions, 450 reviewing for exams answering questions, 456–457 chapter reviews, 464 F fabric border node (SDA underlays), 387 fabric control plane (SDA underlays), 387 Cisco Learning Network, 464 fabric edge node (SDA underlays), 387 CLI practice, 460–461 fabric SDA, 384 knowledge gaps, 458–459 failing exams, 463 practice exams, 454–455, 458–459 failover, HSRP, 261–262 Premium Edition questions, 457 543 544 FHRPs (First Hop Redundancy Protocols) FHRPs (First Hop Redundancy Protocols), 254, 257 firewalls locations, 96–97 features, 260 NGFW, 100–101 HSRP, 261–263 security zones, 97 need for, 259–260 stateful firewalls, 96 options, 260 fiber Internet, 321 flash memory, 269 flow FIFO (First-In, First-Out), 242 control, TCP, 15–16 file system, 268–270 networking, 231 File Transfer Protocol. See FTP public cloud copying, 347–349 files forward acknowledgment, 14 automation configuration variables, 43 forwarding plane. See data plane centralized configuration files, 432 FTP (File Transfer Protocol), 275 managing frames, defined, 233 active mode, 276 IOS file system, 268–270 control connection, 277 upgrading IOS images, 270–274 copying IOS images with, 273–274 transferring, 20–21 filtering DISCOVER messages based on MAC addresses, 150 MAC addresses, DAI, 159 RELEASE messages based on IP addresses, 151 reputation-based filtering, NGIPS, 103 FIN bits, 12 finding web servers with DNS, 18–20 wildcard masks, 33–34 data connection, 277 passive mode, 276 FTSP (File Transfer Protocol Secure), 279 full drops, 251 full mesh topology, 291, 295, 308 G Get messages agent information, 264 RO/RW communities, 267 GET requests, 20 GitHub, 433 Google App Engine PaaS, 341 interfaces 545 H hub and spoke topology (MetroE), 309 hardware human vulnerabilities (security), 79 800 Cisco server, 330–331 hybrid topology, 291, 295 origin hardware addresses, 159–160 hypervisors, 332 hashes coding passwords with, 90 I enable secret command, 92–94 MD5 hash algorithm, 93 headers 802.1Q, 237–238 802.11, 238 IaaS (Infrastructure as a Service), 339–340 IANA (Internet Assigned Numbers Authority), 205 IAP, 237–238 IBN (Intent-Based Networking), 371, 398 MPLS Label, 238 IEEE, Ethernet standards, 306 hiding passwords for local usernames, 94 history, SNMP, 263 home office wireless LANs, 296–297 hosts IPv4 settings, 133 IaaS (Infrastructure as a Service), 339–340 IANA (Internet Assigned Numbers Authority), 205 IAP, 237–238 IBN (Intent-Based Networking), 371, 398 MPLS Label, 238 IEEE, Ethernet standards, 802.1Q headers, 237–238 M Ethernet 802.11 headers, 238 MAC addresses MPLS Label headers, 238 IP marking, 237–238 filtering trust boundaries, 238–239 DAI, 159 matching packets, 27 DISCOVER messages, 150 matching parameters port security, 113 extended numbered ACLs, 46–50 sticky secure MAC addresses, 109 standard numbered ACLs, 31–34 macOS, host IPv4 settings, 136–138 MD5 hash algorithm, 93 malware, 79 MD5 verification, 273 Trojan horses, 78 measuring cloud computing services, 337 viruses, 78 MEF (Metro Ethernet Forum), 306 worms, 78 memory NGFW and, 101 Man-in-the-Middle attacks, 76–77 flash memory, 269 Management Information Base. See MIB TCAM, 362 messages management plane (networking devices), 361 checks, DAI, 164–165 managers, SNMP, 264 Inform, 265–266 managing integrity, SNMPv3, 268 Get, 264, 267 bandwidth, 228 log messages, 175–177 delay, 229 rate limits jitter, 229 DAI, 163–164 loss, 229 DHCP snooping, 154–156 sending to users, 174–175 550 messages Set, 264, 267 multifactor credentials (security), 81 SNMP, 265 multiple queues (queuing systems), 242 Trap, 265–266 MetroE, 304 access links, 306 multiplexing, 7–10 multithreading, 332 IEEE Ethernet standards, 306 Layer 3 design, 309–311 MEF, 306 physical design, 305–306 N named ACLs services, 306–311 configuration, 55–56 topologies, 307–309 editing, 56–58 MIB (Management Information Base), 264, 267 OIDs, 266 variables overview, 54–55 names, MIB variables, 266 NAT (Network Address Translation), 202 monitoring, 265 dynamic NAT, 210–211, 215–219 numbering/names, 266 overview, 200 not-ACLs, 246 PAT, 211, 213, 219–222 monitoring source NAT, 208 configuration, 433 MIB variables, 265 more command, 270 static NAT, 208–210, 214–25, 222 troubleshooting, 222–223 MPBGW (Multiprotocol BGP), 316 NAT Overload. See PAT MPLS (Multi-Protocol Label Switching), 311–312 National Institute of Standards and Technology (NIST), 366 access links, 314 Label headers, QoS marking, 238 Layer 3 design, 313 MPLS VPNs, 315–317 NBAR (Network Based Application Recognition), 235–236 NBIs (Northbound Interfaces), 365–366 public cloud connections, 345 NDAs (Nondisclosure Agreements), 454 QoS, 314–315 netstat -rn command, 136–142 multichoice questions (exams), 449–450, 457 note taking (exam preparation) 551 Network Management Station (NMS), 264 networks automation and network management, 376–378 broad access, 337 controllers, 362–366, 375–379 devices control plane, 360–361 loss, 229 types, 229–232 virtual networks, 333–334 VMs, 334 Network Time Protocol. See NTP Nexos 1000v vSwitch, 334 NGFW (Next-Generation Firewalls), 100–101 data plane, 359 NGIPS (Next-Generation Intrusion Prevention Systems), 100–103 management plane, 361 NICs (Network Interface Cards) switch internal processing, 361–362 DNA Center, 400–401 file systems, 270 flow, 231 management ports, 334 vNICs, 333 NIST (National Institute of Standards and Technology), 306 NMS (Network Management Station), SNMP, 264–266 automation, 376–378 no cdp enable command, 193 DNA Center, 400–401 no enable password command, 105 physical data center, 334–335 no enable secret command, 105 programmability no ip access-group command, 60 ACI, 369, 373 comparisons, 375 redundancy needs, 257–259 SNMP, 254 traditional versus controller-based networks, 375–379 traffic bandwidth, 228 characteristics, 228 delay, 229 jitter, 229 no ip dhcp snooping information option command, 152–153 no logging console command, 177 no logging monitor command, 177 no service password-encryption command, 90 no shutdown command, 115, 121, 179 noninteractive data application traffic, 230 Northbound Interfaces (NBIs), 365–366 note taking (exam preparation), 452 552 notifications, SNMP notifications, SNMP, 265–266 nslookup command, 78 NTP (Network Time Protocol) client/server configuration, 183–184 loopback interfaces, 188–189 overview, 181–182 primary servers, 187 redundant configuration, 186–188 reference clocks, 184–186 secondary servers, 187 setting time and timezone, 182–183 stratum, 183–186 ntp master command, 183–185, 188, 200 on-demand self-service (cloud computing), 337 on-premise. See private cloud computing one-way delay, 229 ONF (Open Networking Foundation), 367 opaque file systems, 270 Open SDN, 367 OpenFlow, 364, 367 OpFlex, 364 origin hardware addresses, 159–160 origin IP addresses, 157–159, 163–164 OSC (Open SDN Controllers), 369 outside global addresses, 209–210 ntp server command, 183, 188, 200 outside local addresses, 209–210 ntp source command, 200 overlays (SDA), 384 numbered ACLs, 58–59 LISP, 392–393 numbers VXLAN tunnels, 390–391, 394 MIB variables, 266 port numbers, 9–10 overloading NAT, 211–213, 219–222 sequence numbers, 56–58 NVRAM (Non-Volatile Random Access Memory) file systems, 270 O objects, 20 P PaaS (Platform as a Service), 341–342 packets classification, 233–236 congestion objects (JSON), 424–426 avoidance, 250–251 ODL (OpenDaylight), 368 management, 242–245 OIDs (object IDs), 266 PoP (Post Office Protocol) 553 defined, 233 permit keyword, 28, 34 marking, 234–241 phishing attacks, 79 matching, 27 PHB (Per-Hop Behaviors), 226 policing, 245–248 phishing attacks, 79 router queuing, 233 physical access control (security), 84 shaping, 245, 248–250 PAR (Positive Acknowledgment and Retransmission), 16 physical data center networks, 334–335 partial mesh topology, 291, 295, 308 physical design, MetroE, 305–306 passive mode (FTP), 276 physical server model, 331 password command, 90, 105 physical standards, Ethernet LANs, 296–297 passwords physical NICs, ports, 334 alternatives to, 81 PI (Prime Infrastructure), 400–401 brute-force attacks, 80 planes, networking devices, 359–361 clear-text, 267 dictionary attacks, 80 guessing, 80 security, 88–94 vulnerabilities (security), 80 Platform as a Service (PaaS), 341–342 PoE (Power over Ethernet), 297–299 PASV command, 278 Point-to-Point topology (MetroE), 307–308 PAT (Port Address Translation) policing (QoS), 245 configuration, 219–222 discarding excess traffic, 247 overview, 211–212 edge between networks, 246–247 troubleshooting, 222 features, 248 Path Trace feature (DNA Center), 403 PCP (Priority Queue) field (802.1Q header), 237 rates, 246 traffic rate versus configured policing rate, 246 PD (Powered Devices), 298–299 pooling resources, cloud computing, 337 PE (Provider Edge), 313 PoP (Post Office Protocol) per-device configuration model, 431 MetroE, 305 permit command, 55–57, 62 POP3, 11 554 Port Address Translation (PAT) Port Address Translation (PAT) preparing for exams configuration, 219–222 24 hours before the exam, 452 overview, 211–213 PORT command, 277–278 30 minutes before the exam, 452–453 port-security command, 111 earplugs, 451 ports one week away preparations, 451–452 NICs, 334 numbers destination port numbers, 8 dynamic ports, 9 ephemeral ports, 9 matching, 48–50 post exam process, 453 taking notes, 452 travel time, 452 prioritization, congestion management, 242 private ports, 9 Priority Code Point (PCP) field (802.1Q header), 237 registered ports, 9 priority queues, 9 priority queues, 244 system ports, 9–11 private addressing, 206 user ports, 9 private cloud computing, 337–338 well known ports, 9–11 private (dynamic, ephemeral) ports, 9 security, 108–111 er-disabled state, 115 private Internet, 206 MAC addresses, 113 private WANs protect mode, 117–119 MetroE, 304–311 restrict mode, 117–119 MPLS, 311–317 shutdown mode, 115–117 public cloud, accessing, 344–346 verifying, 112–113 public cloud branch office connections, 349 violation modes, 114–119 trusted ports, 147 VMs, 334 Post Office Protocol. See POP practice exams, 454 checklists, 455, 459 PTP questions, 455 programmability (network) ACI, 369, 373 comparisons, 375 protect mode (port security), 117–119 protocols 555 protocols CDP copying IOS images with, 273–274 configuration, 193–194 data connection, 277 discovering information about neighbors, 190–193 passive mode, 276 FTSP, 279 verification, 193–194 HSRP control plane, 360–363 active/passive mode, 261 DHCP, 122 failover, 261–262 advantages of, 124 automatic allocation, 129 load balancing, 262–263 HTTP broadcast flags, 125 overview, 16–17, 20–21 DHCP Relay, 126–127, 130 REST APIs, 413–416 dynamic allocation, 129 software CRUD actions, 413–414 information stored at DHCP server, 128 URIs, 17–18, 114–416 overview, 124–126 management plane, 361 files, 126–127, 130 matching, 46–48 routers, 128, 132–133 MPBGW, 316 rules of, 149 SFTTP, 279 servers, 128 SNMP, 11, 254 snooping, 146–156. See also snooping attacks agents, 264 static allocation, 129 communities, 267 switches, configuring as DHCP clients, 130–132 community strings, 267 troubleshooting, 13

Addressable Memory (TCAM), 362 bandwidth, 228 testlet questions (exams), 450 configuration FTTP (Trivial File Transfer Protocol), 11, 129, 274, 279-280 threads, multithreading, 332 threats (security), 72 three-tier campus design, 293-295 TID fields (QoS marking), 238 time exams characteristics, 228 avoidance, 250-251 management, 242-245 delay, 229 jitter, 229 loss, 229 policing, 245-248 public cloud branch office email services, 347-349 budgeting, 450-451 shaping, 245, 248-250 time-check method, 451 types, 229-232 intervals (QoS shaping), 249 setting, 182-183 voice, 315 Traffic Class field (IPv6), 237 timezone, setting, 182-183 transferring files, 20-21 tokens, QoS, 233-251 Top of Rack (ToR) switches, 335 Transmission Control Protocol, See TCP topologies transport input command, 105 campus LANs, 290-295 transport input ssh command, 89 DNA Center topology map, 401-403 transport layer (TCP/IP) full mesh, 291, 295, 308 TCP, 6-16 UDP, 6-7, 16 hub and spoke, 309 Trap messages, 265-266 hybrid, 291, 295 travel time (exam preparation), 452 MetroE, 306-309 partial mesh, 291, 295, 308 Trivial File Transfer Protocol (TFTP), 11, 129, 274, 279-280 star, 291, 295, 309 Trojan horses, 78 ToR (Top of Rack) switches, 335 users 567 troubleshooting ACL, 222 DHCP, 130 dynamic NAT, 222 NAT, 222-223 PAL, 222 port security, 115-119 standard numbered ACLs, 38-39 static NAT, 222 Uniform Resource Locators. See URLs untrusted ports, DHCP messages, 147 upd keyword, 48 upgrading IOS images, 270-274 UPoE (Universal Power over Ethernet), 299 URIs (Uniform Resource Identifiers), 17-18, 414-416 trust boundaries (QoS marking), 238-239 URLs (Uniform Resource Locators), 17, 102 trusted ports, DHCP messages, 147 tunnels (VPN), 321-322 U.S. National Institute of Standards and Technology. See NIST tutorials (exams), 449 upstream, 269-270 two-tier campus design, 290-293 User Datagram Protocol. See UDP Type of Service (ToS) field (IPv4), 237 user network interface. See UNI user (registered) ports, 94 UCS (Unified Computing System), 331, 370 username password command, 94 UDP (User Datagram Protocol) username command, 105 username secret command, 94 users overview, 16 access security, 82-83 port numbers, 48-50 awareness/training, 83-84 supported features, 6-7 groups, SDA security, 398-399 underlays (SDA), 384-388 UNI (User Network Interface), 306 Unified Computing System. See UCS Uniform Resource Identifiers. See URIs sending messages to, 174-175 568 variables Virtual Private Wire Service. See VPWS V numbers, 48-50 awareness/training, 435-437 dictionary variables, 411-412 list variables, 411-412 MIB, 265-266 REST APIs, 410-412 simple variables, 410-411 vCPU (virtual CPU), 332 verification CDP, 193-194 host IPv4 settings, 134-140 NAT, 215-219 standard numbered ACLs, 38-39 Syslog, 178-180 verify command, 273, 282 verify /md5 command, 273, 282 verifying IOS code integrity, 273 port security, 112-113 video exam tutorials, 449 video traffic QoS requirements, 232 shaping time intervals, 249 violation modes (port security), 114-119 virtual CPU (vCPU), 332 virtual NICs. See vNICs Virtual Private LAN Service. See VPLS virtual switches. See vSwitches virtualization data centers, 333-336 networks, 333-334 servers, 332-334 virtual machines. See VMs viruses, 78 wishing attacks, 79 VMs (Virtual Machines), 332-333 ACL, 371 configuration (automated), 334 laas, 340 networking, 334 PaaS, 341-342 ports, 334 SaaS, 341 spinning up, 340 vNICs (virtual NICs), 333 voice application traffic, 231-232 Voice over IP. See VoIP voice traffic shaping time intervals, 249 VoIP, 315 VoIP (Voice over IP), 231-232, 315 VPLS (Virtual Private LAN Service), 307 VPNs (Virtual Private Networks) AnyConnect Security Mobility Client, 325 client, 325 Internet, 317, 321-322 workflow, virtualized data center 569 public cloud, accessing, 344 web browsers, 16 remote-access VPNs, 324-326 HTTP, 16-21 site-to-site, 322-326 identifying receiving application, 21-22 tunnels, 321-322 VPWS (Virtual Private Wire Service), 307 URIs, 17-18 URLs, 17 vSwitches, 333 web pages, 16 VUE testing center, 455 web pages, 16 vulnerabilities (security), 72 web servers, 16-20, 371 human vulnerabilities, 79-80 password vulnerabilities, 80 VXLAN tunnels, 385, 390-391, 394, 399 websites Cisco ACI, 373 Cisco Prime management products, 264 Eclipse IDE, 341 W WANs (Wide-Area Networks) Ethernet, 345 Google App Engine PaaS, 341 Jenkins continuous integration and automation tool, 341 MEF, 306 interfaces, 228 OpenDaylight SDN controller, 368 Internet access, 311 OpenFlow, 364 Internet as WAN service, 317 weighting, 243 MetroE, 304-311 well known (system) ports, 9-11 MPLS, 311-317 whaling attacks, 79 private, 344-346, 349 whois command, 78 public cloud connections, 342-346 wildcard masks, 31-34, 41 SPs, 302 wireless, 320-321 watering hole attacks, 79 WC masks, 31-34, 41 windowing, 15-16 wireless routers, 296 wireless WANs, 320-321 WLANs (Wireless LANs), 296-297 workflow, virtualized data center, 335-336 570 worms worms, 78 WWW (World Wide Web), 11 X XML (Extensible Markup Language), data serialization, 421-423 Y-Z YAML (YAML Ain't Markup Language), data serialization, 422-423 This page intentionally left blank Executive Overview • 40% OFF Cisco Press Video Training ciscopress.com/video Use coupon code CPVIDEAO4 during checkout. Video Instruction from Technology Experts Advance Your Skills Train Anywhere Learn Get started with fundamentals, become an expert, or get certified. Train anywhere, at your own pace, on any device. Learn from trusted author trainers published by Cisco Press. Try Our Popular Video Training for FREE! ciscopress.com/Video Explore hundreds of FREE video lessons from our growing library of Complete Video Courses, LiveLessons, networking talks, and workshops. ciscopress.com/video REGISTER YOUR PRODUCT at CiscoPress.com/register Access Additional Benefits and SAVE 35% on Your Next Purchase • Download available product updates. • Access bonus material when applicable. • Receive exclusive offers on new editions and related products. (Just check the box to hear from us when setting up your account.) • Get a coupon for 35% for your next purchase, valid for 30 days. Your code will be available in your Cisco Press cart. (You will also find it in the Manage Codes section of your account page.) Registration benefits vary by product. Benefits will be listed on your account page under Registered Products. CiscoPress.com – Learning Solutions for Self-Paced Study, Enterprise, and the Classroom Cisco Press is the Cisco Systems authorized book publisher of Cisco networking technology, Cisco certification self-study, and Cisco Networking Academy Program materials. At CiscoPress.com you can • Shop our books, eBooks, software, and video training. • Take advantage of our special offers and promotions (ciscopress.com/promotions). • Sign up for special offers and content newsletters (ciscopress.com/newsletters). • Read free articles, exam profiles, and blogs by information technology experts. • Access thousands of free chapters and video lessons. Connect with Cisco Press – Visit CiscoPress.com/community Learn about Cisco Press community events and programs. APPENDIX D Topics from Previous Editions Cisco changes the exams, renaming the exams on occasion, and changing the exam numbers every time it changes the exam with a new blueprint, even with a few name changes over the years. As a result, the current CCNA 200-301 exam serves as the eighth separate version of CCNA in its 20-plus-year history. At every change to the exams, we create new editions of the books to match the new exam. We base the books' contents on Cisco's exam topics—and it is, the book attempts to cover the topics Cisco lists as exam topics. However, the book authoring process does create some challenges, particularly with the balance of what to include in the books and what to leave out. For instance, when comparing a new exam to the old, I found Cisco had removed some topics—and I might want to keep the content in the book. There are a few reasons why. Sometimes I just expect that some readers will still want to read about that technology. Also, more than a few schools use these books as textbooks, and keeping some of the older-but-still-relevant topics can be a help. And keeping the old material available on each book's companion website takes only a little extra work, so we do just that. Some of the older topics that I choose to keep on the companion website are small, so I collect them into this appendix. Other topics happen to have been an entire chapter in a previous edition of the books, so we include those topics each as a separate appendix. Regardless, the material exists here in this appendix, and in the appendices that follow, for your use if you have a need. But do not feel like you have to read this appendix for the current exam. The topics in this appendix are as follows: ■ Cisco Device Hardening ■ Implementing DHCP ■ Troubleshooting with IPv4 ACLs ■ Implementing HSRP ■ Global Load Balancing Protocol (GLBP) ■ Implementing Simple Network Management Protocol ■ Analyzing LAN Physical Standard Choices ■ Metro Ethernet Virtual Circuits ■ MPLS VPNs and OSPF NOTE The content under the heading "Cisco Device Hardening" was most recently published for the 100-105 Exam in 2016, in Chapter 34 of the Cisco CCNA ICND1 100-105 Official Cert Guide. Cisco Device Hardening The term device hardening refers to making it more difficult for attackers to gain access to the device or to cause problems for the device. This section does not attempt to mention all such details, but it does touch on a few items. (Note that the CCNA Security certification gets into much more detail about router and switch device security.) In particular, this second major section of the chapter begins by showing how to set some login banner message text for users. The next two topics look at how to secure items unused in the device—unused switch ports on switches and unused software services in both routers and switches. Configuring Login Banners Cisco switches and routers can display a variety of banners to a new user when logging in to the switch or router. A banner is simply some text that appears on the screen for the user. You can configure a router or switch to display multiple banners, some before login and some after. IOS supports three banners based on the first keyword in the banner command. Table D-1 lists the three most popular banners and their typical use. Table D-1 Banners and Their Typical Use Banner Typical Use Message of the Day (MOTD) Used for temporary messages that can change from time to time, such as "Router1 down for maintenance at midnight." Login Because it is always shown before the user logs in, this message is often used to show warning messages, like "Unauthorized Access Prohibited." Exec Because this banner always appears after login, it typically lists device information that outsiders should not see but that internal staff might want to know, for example, the exact location of the device. In what may seem like trivia, the banners actually appear in different places based on a couple of conditions. Figure D-1 summarizes when the user sees each of these banners, reading from top to bottom. Console and Telnet users see the banners in the order shown on the left, and SSH users see the banners in the order on the right. 4 CCNA 200-301 Official Cert Guide, Volume 2 Console, Telnet SSH MOTD Login Login (User Login) MOTD (User Login) E xec E xec Terminal Window Terminal Window Figure D-1 Banner Sequence Compared: Console/Telnet Versus SSH (Blue Ribbon Set © petrunil/123RF) NOTE If using SSH, and the switch or router uses only SSHv1, the login banner is not shown to the SSH user. The banner global configuration command can be used to configure all three types of these banners. In each case, the type of banner is listed as the first parameter, with both being the default option. The first nonblank character after the banner type is called a beginning delimiter character. When a delimiter character is used, the banner text can span several lines, with the CLI user pressing Enter at the end of each line. The CLI knows that the banner has been configured as soon as the user enters the same delimiter character again. Example D-1 shows the configuration process for all three types of banners from Table D-1, followed by a sample user login session from the console that shows the banners in use. The first configured banner in the example, the MOTD banner, omits the banner type in the banner command as a reminder that motd is the default banner type. The first two banner commands use a # as the delimiter character. The third banner command uses a / as the delimiter, just to show that any character can be used. Also, the last banner command shows multiple lines of banner text. Example D-1 Banner Configuration | Below, the three banners are created in configuration mode. Note that any / delimiter can be used, as long as the character is not part of the message | text. SW1(config)# banner # Enter TEXT message. End with the character '#'. (Login) Unauthorized Access Prohibited!!!! # SW1(config)# banner exec Z Enter TEXT message. End with the character 'Z'. (Exec) Company picnic at the park on Saturday. Appendix D: Topics from Previous Editions 5 Don't let outsiders! Z SW1(config)# end | Below, the user of this router quits the console connection, and logs | back in, seeing the motd and login banners, and then the password prompt, | and then the exec banner. SW1#> quit SW1 con0 is now available Press RETURN to get started. (MOTD) Switch down for maintenance at 11PM Today (Login) Unauthorized Access Prohibited!!!! User Access Verification Username: fred Password: (Exec) Company picnic at the park on Saturday. Don't let outsiders! SW1> Securing Unused Switch Interfaces The default settings on Cisco switches work great if you want to buy a switch, unbox it, plug it in, and have it immediately work without any other effort. Those same defaults have an unfortunate side effect for security, however. With all default configuration, an attacker might use unused interfaces to gain access to the LAN. So, Cisco makes some general recommendations to override the default interface settings to make the unused ports more secure, as follows: ■ Administratively disable the interface 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 using the switchport access vlan number interface subcommand. ■ Set the native VLAN so that it is not VLAN 1 but instead is an unused VLAN, using the switchport trunk native vlan vlan-id interface subcommand. Frankly, if you just shut down the interface, the security exposure goes away, but the other tasks prevent any immediate problems if someone else comes around and enables the interface by configuring a no shutdown command. D 6 CCNA 200-301 Official Cert Guide, Volume 2 NOTE The contents under the headings "DHCP Server Configuration on Routers," "IOS DHCP Server Verification," and "Troubleshooting DHCP Services" were most recently published for the 100-105 Exam in 2016, in Chapter 20 of the Cisco CCNA ICND1 100-105 Official Cert Guide. Implementing DHCP This section includes DHCP implementation topics from an earlier edition of the book. DHCP Server Configuration on Routers A quick Google search on "DHCP server products" reveals that many companies offer DHCP server software. Cisco routers (and some Cisco switches) can also act as a DHCP server with just a little added configuration. Configuring a Cisco router to act as a DHCP server uses a new configuration concept, one per subnet, called a DHCP pool. All the per-subnet settings go into a per-subnet DHCP pool. The only DHCP command that sits outside the pool is the command that defines the list of addresses excluded from being leased by DHCP. The Cisco IOS DHCP server configuration steps are as follows: Step 1. Use the ip dhcp excluded-address first last command in global configuration mode to list addresses that should be excluded (that is, not leased by DHCP). Step 2. Use the ip dhcp pool name command in global configuration mode to both create a DHCP pool for a subnet and to navigate into DHCP pool configuration mode. Then also: A. Use the network subnet-ID mask or network subnet-ID prefix-length command in DHCP pool configuration mode to define the subnet for this pool. B. Use the default-router address1 address2... command in DHCP pool configuration mode to define default router IP address(es) in that subnet. C. Use the dns-server address1 address2... command in DHCP pool configuration mode to define the list of DNS server IP addresses used by hosts in this subnet. D. Use the lease days hours minutes command in DHCP pool configuration mode to define the length of the lease, in days, hours, and minutes E. Use the domain-name name command in DHCP pool configuration mode to define the DNS domain name. F. Use the next-server ip-address command in DHCP pool configuration mode to define the TFTP server IP address used by any hosts (like phones) that need a TFTP server. Of course, an example can help, particularly with so many configuration commands required. Figure D-2 shows the organization of the configuration, while sticking to pseudocode rather than the specific configuration commands. (Upcoming Example D-2 shows a matching configuration.) Note that for each of the two LAN subnets, there is a global command to exclude addresses, and then a group of settings for each of two different DHCP pools. Appendix D: Topics from Previous Editions 7 Global Exclude: 172.16.1.1–172.16.1.50 Global Exclude: 172.16.1.1–172.16.2.100 Pool subnet-left Subnet= Router= DNS= Lease Time= Domain= Pool subnet-right Subnet= 172.16.2.0/24 Router: .1 DNS= 172.16.1.12 Lease Time= 1 Days 2 Hours 3 Minutes A 172.16.1.0/24 .1 172.16.1.12 0 Days 23 Hours 59 Minutes example.com . 9 . 8 . 1 . 12 172.16.1.0/24 Figure D-2 R1 DHCP Relay Agent DNS B R2 DHCP Server 5 UCM Server 172.16.2.0/24 DHCP Server Configuration Pseudocode Example D-2 R2 as a DHCP Server Per the Concepts in Figure D-2 | ip dhcp excluded-address 172.16.2.1 172.16.2.100 | ip dhcp pool subnet-right network 172.16.2.0/24 dns-server 172.16.2.1 | ip dhcp pool subnet-right network 172.16.2.0/24 dns-server 172.16.1.12 default-router 172.16.2.1 | ip dhcp excluded-address 172.16.2.5 Focus on subnet 172.16.1.0/24 for a moment: the subnet configured as pool subnet-left. The subnet ID and mask match the subnet ID chosen for that subnet. Then, the global ip dhcp excluded-address command, just above, reserves 172.16.1.1 through 172.16.1.50, so that this DHCP server will not lease these addresses. The server will automatically exclude the subnet ID (172.16.1.0) as well, so this DHCP server will begin leasing IP addresses starting with the . 51 address. Now look at the details for subnet-right. It uses a DHCP pool network command with a prefix style mask. It defines the same DNS server, as does the pool for the other subnet, but a different default router setting, because, of course, the default router in each subnet D 8 CCNA 200-301 Official Cert Guide, Volume 2 is different. This pool includes a lease time of 1:02:03 (1 day, 2 hours, and 3 minutes) just as an example. Also note that both subnets list a TFTP server IP address of the Unified Communications Manager (UCM) server with the next-server command. In most cases, you would find this setting in the pools for subnets in which phones reside. Finally, note that configuring a router as a DHCP server does not remove the need for the ip helper-address command. If DHCP clients still exist on LANs that do not have a DHCP server, then the routers connected to those LANs still need the ip helper-address command. For example, in Figure D-2, R1 would still need the ip helper-address command on its LAN interface. R2 would not need the command on its LAN interface, because R2 could service those requests, rather than needing to forward the DHCP messages to some other server. IOS DHCP Server Verification The IOS DHCP server function has several different show commands. These three commands list most of the details: show ip dhcp binding: Lists state information about each IP address currently leased to a client show ip dhcp pool [poolname]: Lists the configured range of IP addresses, plus statistics for the number of currently leased addresses and the high-water mark for leases from each pool show ip dhcp server statistics: Lists DHCP server statistics Example D-3 shows sample output from two of these commands, based on the configuration from Figure D-2 and Example D-2. In this case, the DHCP server leased one IP address from each of the pools, one for host A, and one for host B, as shown in the highlighted portions of the output. Example D-3 Verifying Current Operation of a Router-Based DHCP Server R2# show ip dhcp binding Bindings from all pools not associated with VRF: IP address Client-ID/ Lease expiration Type Oct 12 2012 02:56 AM Automatic Oct 12 2012 04:59 AM Automatic Hardware address/ User name 172.16.1.51 0063.6973.636f.2d30. 3230.302e.3131.3131. 2e31.3131.312d.4661. 302f.30 172.16.2.101 0063.6973.636f.2d30. 3230.302e.3232.322d.4769. 302f.30 R2# show ip dhcp pool subnet-right Pool subnet-right: Utilization mark (high/low) : 100 / 0 Subnet size (first/next) : 0 / 0 Total addresses : 254 Appendix D: Topics from Previous Editions 9 Leased addresses : 1 Pending event : none 1 subnet is currently in the pool : Current index IP address range 172.16.2.102 172.16.2.1 Leased addresses - 172.16.2.254 1 Note that the output in Example D-3 does not happen to list the excluded addresses, but it does show the effects. The addresses assigned to the clients end with .51 (host A, subnet 172.16.1.0) and .101 (host B, subnet 172.16.2.0), proving that the server did exclude the addresses as shown in the configuration in Example D-2. The server avoided the .1 through .50 addresses in subnet 172.16.1.0, and the .1 through .100 addresses in subnet 172.16.2.0. NOTE The DHCP server keeps status (state) information about each DHCP client that leases an address. Specifically, it remembers the DHCP client ID, and the IP address leased to the client. As a result, an IPv4 DHCP server can be considered to be a stateful DHCP server. Troubleshooting DHCP Services To be prepared for the CCNA simlet questions, you have to be ready to predict what symptoms would occur when the network was misconfigured in particular ways. This next section takes a similar approach, pointing out the most typical issues that could be introduced through incorrect or missing configuration, and then discussing what symptoms should happen and how to recognize those problems. This section begins with a typical look at configuration mistakes and the symptoms that occur with those mistakes. In particular, this section looks at problems with the relay agent's helper address as well as the IOS DHCP server configuration. This section then looks at non-DHCP problems related to that data plane, breaking the problem into issues between the client and relay agent, and between the relay agent and DHCP server. The final section takes a short look at how a DHCP server prevents duplicate IP addresses between hosts that use static IP addresses and those that use DHCP. DHCP Relay Agent Configuration Mistakes and Symptoms One configuration mistake that prevents DHCP client from leasing an IP address is the misconfiguration or the omission of the ip helper-address interface subcommand on the router acting as the DHCP relay agent. The relay agent takes the incoming DHCP message, changes the destination address of the packet to be the address on the ip helper-address address command, and forwards the packet to that address. If the command is missing, the router does not attempt to forward the DHCP messages to all; if it is incorrect, the relay agent forwards the DHCP packets, but they never arrive at the actual DHCP server. The main problem symptom in this case is the failure of a DHCP client to lease an address. If you can identify a client that has a problem, and you know what VLAN or subnet in which that host resides, you can then work to identify any routers connected to that subnet, to find and correct the ip helper-address subcommands. Beyond that step, this list summarizes a few other related points. ■ The DHCP relay agent feature is needed on interfaces only if the DHCP server is on a different subnet; it is not needed if the DHCP server is on the same subnet as the client. D 10 CCNA 200-301 Official Cert Guide, Volume 2 ■ On routers with VLAN trunks (with a router-on-a-stick [ROAS] subinterface configuration), the subinterfaces also need an ip helper-address command (assuming they meet the first criteria in this list). ■ If an exam question does not allow you to look at the configuration, use the show ip interface [type number] command to view the ip helper-address setting on an interface. About that last point, Example D-4 shows an example of the show ip interface g0/0 command. In this case, the interface has been configured with the ip helper-address 172.16.2.11 command; the show command output basically restates that fact. Note that if there were no ip helper-address configured on the interface, the text would instead read "Helper address is not set." Example D-4 Listing the Current Helper Address Setting with show ip interface R1# show ip interface g0/0 GigabitEthernet0/0 is up, line protocol is up Internet address is 182.16.1.1/24 Broadcast address is 255.255.255.255 Address determined by non-volatile memory MTU is 1500 bytes Helper address is 172.16.2.11 ! Lines omitted for brevity (about 20 lines) IOS DHCP Server Configuration Mistakes and Symptoms When using an IOS DHCP server, from a troubleshooting perspective, break issues into two broad categories: those that prevent DHCP clients from leasing an address, and those that allow the lease but provide incorrect settings to the client. First, the primary configuration mistake that causes a failure in the DHCP lease process is the misconfiguration of the network command. The problem revolves around these key facts: ■ The packet from the relay agent to the DHCP server uses the relay agent's interface IP address as the source IP address in the forwarded DHCP message. ■ The DHCP server compares that source IP address in the received DHCP packet to the network commands in its DHCP pools to find the right pool. ■ Each network subnet mask command implies a range of addresses, just like any other IP network or subnet shown with a subnet mask. ■ If the source IP address of the packet is not in the range of addresses implied by any network command in all the pools, the DHCP server has no pool to use for that request. The DHCP server does not know how to respond, so it does not reply at all. As an example of that failure, consider the configuration shown in Figure D-3. The left side shows the configuration on R1, a DHCP relay agent that has two interfaces configured with the ip helper-address 172.16.2.11 command. The DHCP server configuration on the right lists two pools, intended as one pool for each subnet off Router R1. However, the network 172.16.3.0/25 command implies an address range of 172.16.3.0 to 172.16.3.127, and the relay agent's interface address of 172.16.3.254 is not within that range of numbers. The solution would be to correct the DHCP server's network command to use a /24 mask. Appendix D: Topics from Previous Editions 11 DHCP Server (R2) Remote Router (R1) Match ip dhcp pool top network 172.16.1.0/24 No Match! ip dhcp pool bottom network 172.16.3.0/25 interface G0/1 ip address 172.16.1.1 255.255.255.0 ip helper-address 172.16.2.11 server G0/1/1 ip address 172.16.3.254 255.255.255.0 ip helper-address 172.16.2.11 encapsulation dot1q 172.16.3.0 - 172.16.3.127 DHCP Server 172.16.1.0/24 SW1 172.16.3.0/24 Figure D-3 G0/1 R1 R2 G0/1 An Example Misconfiguration of a DHCP Pool network Command NOTE The ip helper-address configuration on the left is correct. The figure uses a ROAS configuration here just to reinforce the comment in the earlier section that ROAS subinterfaces also need an ip helper-address command. While you ultimately need to find this kind of problem and fix the configuration, on the exam you need to be ready to discover the root cause based on symptoms and show commands as well. So, when troubleshooting DHCP issues, and the client fails to lease an address, look at the IOS DHCP server's network commands. Calculate the range of IP addresses as if that command were defining a subnet. Then compare that range of addresses by the network command in each pool to the interface addresses on the DHCP relay agent routers. Every relay agent interface (that is, every interface with an ip helper-address command configured) should be included in a pool defined at the IOS DHCP server. The DHCP server can also be misconfigured in a way that allows the lease of an address, but then causes other problems. If the lease process works, but the rest of the parameters given to the client are incorrect or missing, the client could operate, but operate poorly. This list summarizes the kinds of mistakes and the resulting symptoms: ■ With the DNS server IP addresses incorrectly configured on the server (or omitted), hosts would fail to resolve hostnames into their associated IP addresses. ■ With the default gateway IP address incorrectly configured on the server (or omitted), hosts could not communicate outside the local subnet. ■ With the TFTP server IP address incorrectly configured (or omitted), an IP phone would fail to correctly load its configuration. IP Connectivity from DHCP Relay Agent to DHCP Server For the DHCP process to work with a centralized server, IP broadcast packets must flow between the client and relay agent, and IP unicast packets must flow between the relay agent and the DHCP server. Any problem that prevents the flow of these packets also prevents DHCP from working. D 12 CCNA 200-301 Official Cert Guide, Volume 2 For perspective, consider the topology in Figure D-4, which again shows the relay agent on the left and the DHCP server on the right. The server uses IP address 172.16.2.11, and the relay agent uses interface address 172.16.1.1. Any failure that prevents the flow of IP packets between those two IP addresses would prevent host A from leasing an IP address. 172.16.1.1 172.16.1.0/24 172.16.2.0/24 G0/1 A SW1 S0/0/0 R1 R2 .2 SW2 .51 .11 Source: 172.16.1.1 Dest.: 172.16.2.11 Source: 172.16.2.11 Dest.: 172.16.1.1 Figure D-4 Addresses Used Between Relay Agent and Server Remember that the IP addresses used on the packets between the relay agent and server, and know that that you may need to troubleshoot IP routing to ensure those packets can be delivered. LAN Connectivity Between the DHCP Client and Relay Agent You might encounter a network environment where DHCP messages on the same LAN as the DHCP client all show a destination IP address of 255.255.255.255. What does that really mean? When a packet uses this 255.255.255.255 address: ■ The address is called the local broadcast address. ■ Packets sent to this address are not forwarded as-is by routers. ■ On a LAN, the sender of an IP local broadcast packet encapsulates these IP packets in an Ethernet frame with an Ethernet broadcast destination address (FFFF.FFFF.FFFF), so the LAN broadcasts the frame. As a result of the logic in these steps, the broadcast DHCP messages can easily flow between the client and router, as long as the LAN works. Summary of DHCP Troubleshooting In summary, as a study tool, the following list summarizes the key troubleshooting ideas from this section on troubleshooting DHCP: Step 1. If using a centralized DHCP server, at least one router on each remote subnet that has DHCP clients must act as DHCP relay agent, and have a correctly configured ip helper-address address subcommand on the interface connected to that subnet. Step 2. If using a centralized IOS DHCP server, make sure the DHCP pools' network commands match the entire network's list of router interfaces that have an ip helper-address command pointing to this DHCP server. Step 3. Troubleshoot for any IP connectivity issues between the DHCP relay agent and the DHCP server, using the relay agent interface IP address and the server IP address as the source and destination of the packets. Step 4. Troubleshoot for any LAN issues between the DHCP client and the DHCP relay agent. Appendix D: Topics from Previous Editions 13 Also, as one final note about DHCP in the real world, DHCP might seem dangerous at this point, with all the focus on potential problems in this section, combined with the importance of DHCP and its use by most end user devices. However, DHCP has some great availability features. First, most DHCP servers set their lease times for at least a few days, often a week, or maybe longer. Combined with that, the DHCP protocol has several processes through which the client reconfirms the existing lease with the server, and re-leases the same IP address in advance of the expiration of the lease. Clients do not simply wait until the moment the lease would expire to then contact the DHCP server, hoping it is available. So the network can have outages, and DHCP clients that have already leased an address can continue to work without any problem. Detecting Conflicts with Offered Versus Used Addresses Beyond troubleshooting the types of problems that would prevent DHCP from working, the IOS DHCP server tries to prevent another type of problem: assigning IP addresses with DHCP when another host tries to statically configure that same IP address. Although the DHCP server configuration clearly lists the addresses in the pool, plus those to be excluded from the pool, hosts can still statically configure addresses from the range inside the DHCP pool. In other words, no protocols prevent a host from statically configuring and using an IP address from within the range of addresses used by the DHCP server. Knowing that some host might have statically configured an address from within the range of addresses in the DHCP pool, both DHCP servers and clients try to detect such problems, called conflicts, before the client uses a newly leased address. DHCP servers detect conflicts by using pings. Before offering a new IP address to a client, the DHCP server first pings the address. If the server receives a response to the ping, some other host must already be using the address, which lets the server know a conflict exists. The server notes that particular address as being in conflict, and the server does not offer the address, moving on to the next address in the pool. The DHCP client can also detect conflicts, but instead of using ping, it uses ARP. In the client case, when the DHCP client receives from the DHCP server an offer to use a particular IP address, the client sends an Address Resolution Protocol (ARP) request for that address. If another host replies, the DHCP client has found a conflict. Example D-5 lists output from the router-based DHCP server on R2, after host B detected a conflict using ARP. Behind the scenes, host B used DHCP to request a lease, with the process working normally until host B used ARP and found some other device already used 172.16.2.102. At that point, host B then sent a DHCP message back to the server, rejecting the use of address 172.16.2.102. The DHCP server logged the message related to host B's discovery of the conflict, and a show command that lists all conflicted addresses. Example D-5 Displaying Information About DHCP Conflicts in IOS "Oct 16 19:28:59.220: %DHCPD-4-DECLINE_CONFLICT: DHCP address conflict: client 0063.6973.636f.2d30.3230.302e.3034.3034.2e30.3430.342d.4769.302f.30 declined 172.16.2.102. R2# show ip dhcp conflict IP address Detection method Detection time 172.16.2.102 Gratuitous ARP Oct 16 2012 07:28 PM VRF D 14 CCNA 200-301 Official Cert Guide, Volume 2 The show ip dhcp conflict command lists the method through which the server added each address to the conflict list: either gratuitous ARP, as detected by the client, or ping, as detected by the server. The server avoids offering these conflicted addresses to any future clients, until the engineer uses the clear ip dhcp conflict command to clear the list. NOTE The content under the heading "Troubleshooting with IPv4 ACLs" was most recently published for the 200-105 Exam in 2016, in Chapter 17 of the Cisco CCNA ICND2 200105 Official Cert Guide. Troubleshooting with IPv4 ACLs The use of IPv4 ACLs makes troubleshooting IPv4 routing more difficult. Any data plane troubleshooting process can include a catchall phrase to include checking for ACLs. A network can have all hosts working, DHCP settings correct, all LANs working, all router interfaces working, and all routers having learned all routes to all subnets—and ACLs can still filter packets. Although ACLs provide that important service of filtering some packets, ACLs can make the troubleshooting process that much more difficult. This third of the three major sections of this chapter focuses on troubleshooting in the presence of IPv4 ACLs. It breaks the discussion into two parts. The first part gives advice about common problems you might see on the exam, and how to find those with show commands and some analysis. The second part then looks at how ACLs impact the ping command. Analyzing ACL Behavior in a Network ACLs cause some of the biggest challenges when troubleshooting problems in real networking jobs. The packets created by commands like ping and traceroute do not exactly match the fields in packets created by end users. The ACLs sometimes filter the ping and traceroute traffic, making the network engineer think some other kind of problems exists when no problems exist at all. Or, the problem with the end-user traffic really is caused by the ACL, but the ping and traceroute traffic works fine, because the ACL matches the end-user traffic with a deny action but matches the ping and traceroute traffic with a permit action. As a result, much of ACL troubleshooting requires thinking about ACL configuration versus the packets that flow in a network, rather than using a couple of IOS commands that identify the root cause of the problem. The show commands that help are those that give you the configuration of the ACL, and on what interfaces the ACL is enabled. You can also see statistics about which ACL statements have been matched. And using pings and traceroutes can help—as long as you remember that ACLs may apply different actions to those packets versus the end-user traffic. The following phrases the ACL troubleshooting steps into a list for easier study. The list also expands on the idea of analyzing each ACL in step 3. None of the ideas in the list are new compared to this chapter and the previous chapter, but it acts more as a summary of the common issues: Step 1. Determine on which interfaces ACLs are enabled, and in which direction (show running-config, show ip interfaces). Step 2. Find the configuration of each ACL (show access-lists, show ip access-lists, show running-config). Appendix D: Topics from Previous Editions 15 Step 3. Analyze the ACLs to predict which packets should match the ACL, focusing on the following points: A. Misordered ACLs: Look for misordered ACL statements. IOS uses firstmatch logic when searching an ACL. B. Reversed source/destination addresses: Analyze the router interface, the direction in which the ACL is enabled, compared to the location of the IP address ranges matched by the ACL statements. Make sure the source IP address field could match packets with that source IP address, rather than the destination, and vice versa for the destination IP address field. C. Reversed source/destination ports: For extended ACLs that reference UDP or TCP port numbers, continue to analyze the location and direction of the ACL versus the hosts, focusing on which host acts as the server using a well-known port. Ensure that the ACL statement matches the correct source or destination port depending on whether the server sent or received the packet. D. Syntax: Remember that extended ACL commands must use the tcp and udp keywords if the command needs to check the port numbers. E. Syntax: Note that ICMP packets do not use UDP or TCP; ICMP is considered to be another protocol matchable with the icmp keyword (instead of tcp or udp). F. Explicit deny any: Instead of using the implicit deny any at the end of each ACL, use an explicit configuration command to deny all traffic at the end of the ACL so that the show command counters increment when that action is taken. G. Dangerous inbound ACLs: Watch for inbound ACLs, especially those with deny all logic at the end of the ACL. These ACLs may discard incoming overhead protocols, like routing protocol messages. H. Standard ACL location: Standard ACLs enabled close to the source of matched addresses can discard the packets as intended, but also discard packets that should be allowed through. Always pay close attention to the requirements of the ACL in these cases. The first two steps are important for simlet questions in case you are not allowed to look at the configuration; you can use other show commands to determine all the relevant ACL configuration. The next few pages show some of the related commands and how they can uncover some of the issues described in the just-completed ACL troubleshooting checklist. ACL Troubleshooting Commands If you suspect ACLs are causing a problem, the first problem-isolation step is to find the location and direction of the ACLs. The fastest way to do this is to look at the output of the show running-config command and to look for ip access-group commands under each interface. However, in some cases, enable mode access may not be allowed, and show commands are required. Instead, use the show ip interfaces command to find which ACLs are enabled on which interfaces, as shown in Example D-6. D 16 CCNA 200-301 Official Cert Guide, Volume 2 Example D-6 Sample show ip interface Command R1# show ip interface S0/0/1 Serial0/0/1 is up, line protocol is up Internet address is 10.1.2.1/24 Broadcast address is 255.255.255.255 Address determined by setup command MTU is 1500 bytes Helper address is not set Directed broadcast forwarding is disabled Multicast reserved groups joined: 224.0.0.9 Outgoing access list is 102 I roughly 26 more lines omitted for brevity Note that the command output lists whether an ACL is enabled, in both directions, and which ACL it is. The example shows an abbreviated version of the show ip interface S0/0/1 command, which lists messages for just this one interface. The show ip interface command would list the same messages for every interface in the router. Step 2 of the ACL troubleshooting checklist then says that the contents of the ACL must be found. Again, the quickest way to look at the ACL is to use the show running-config command. If it's not available, the show access-lists and show ip access-lists commands list the same details shown in the configuration. These commands also list a useful counter that lists the number of packets that have matched each line in the ACL. Example D-7 shows an example. Example D-7 show ip access-lists Command Example R1# show ip access-lists Extended IP access list 102 permit ip 10.1.2.0 0.0.0.255 10.1.4.0 0.0.1.255 (15 matches) The counter can be very useful for troubleshooting. If you can generate traffic that you think should match a particular line in an ACL, then you should see the matches increment on that counter. If you keep generating traffic that should match, but that line's counter never goes up, then those packets do not match that line in that ACL. Those packets could be matching an earlier line in the same ACL, or might not even be reaching that router (for any reason). After the locations, directions, and configuration details of the various ACLs have been discovered in steps 1 and 2, the hard part begins—analyzing what the ACL really does. For example, one of the most common tasks you will do is to look at the address fields and decide the range of addresses matched by that field. Remember, for an ACL that sits in a router configuration, you can easily find the address range. The low end of the range is the address (the first number), and the high end of the range is the sum of the address and wildcard mask. For instance, with ACL 102 in Example D-7, which is obviously configured in some router, the ranges are as follows: Source 10.1.2.0, wildcard 0.0.0.255: Matches from 10.1.2.0 through 10.1.2.255 Destination 10.1.4.0, wildcard 0.0.1.255: Matches from 10.1.4.0 through 10.1.5.255 Appendix D: Topics from Previous Editions 17 The next few pages work through some analysis of a few of the items from step 3 in the troubleshooting checklist. Example Issue: Reversed Source/Destination IP Addresses IOS cannot recognize a case in which you attempt to match the wrong addresses in the source or destination address field. So, be ready to analyze the enabled ACLs and their direction versus the location of different subnets in the network. Then ask yourself about the packets that drive that ACL: what could the source and destination addresses of those packets be? And does the ACL match the correct address ranges, or not? For example, consider Figure D-5, a figure that will be used in several troubleshooting examples in this chapter. The requirements for the next ACL follow the figure. 10.1.1.0/24 10.3.3.0/25 A 10.2.2.0/30 G0/1 . 1 . 9 R1 G0/2 . 1 . G0.1 . 2 R2 G0/2 . 2 B 10.4.4.0/23 Figure D-5 Example Network Used in IPv4 ACL Troubleshooting Examples For this next ACL, the requirements ask that you allow and prevent various flows, as follows: ■ Allow hosts in subnet 10.3.3.0/25 and subnet 10.1.1.0/24 to communicate ■ Prevent hosts in subnet 10.4.4.0/23 and subnet 10.1.1.0/24 from communicating ■ Allow all other communications between hosts in network 10.0.0.0 ■ Prevent all other communications Example D-8 shows the ACL used in this case on R2. At first glance, it meets all those requirements straight down the list. Example D-8 Mismatch Troubleshooting Example 2 per Step 3B: Source and Destination R2# show ip access-lists Standard IP access list Step3B 10 permit 10.3.3.0 0.0.0.127 20 deny 10.4.4.0 0.0.1.255 30 permit 10.0.0.0 0.255.255.255 (12 matches) R2# R2# show ip interface G0/1 | include Inbound Inbound access list is Step3B The problem in this case is that the ACL has been enabled on R2's G0/1 interface, inbound. Per the figure, packets coming from a source address in subnets 10.3.3.0/25 and 10.4.4.0/23 should be forwarded out R2's G0/1 interface, rather than coming in that interface. So, do not let the matching logic in the ACL that perfectly mirrors the requirements fool you; make sure and check the location of the ACL, direction, and the location of the IP addresses. D 18 CCNA 200-301 Official Cert Guide, Volume 2 Note that step 3C suggests a similar issue regarding matching well-known ports with TCP and UDP. The earlier section in this chapter titled "Matching TCP and UDP Port Numbers" has already discussed those ideas in plenty of detail. Just make sure to check where the server sits versus the location and direction of the ACL. Steps 3D and 3E: Common Syntax Mistakes Steps 3D and 3E describe a couple of common syntax mistakes. First, to match a TCP port in an ACL statement, you must use a tcp protocol keyword instead of ip or any other value. Otherwise, IOS rejects the command as having incorrect syntax. Same issue with trying to match UDP ports: a udp protocol keyword is required. To match ICMP, IOS includes an icmp protocol keyword to use instead of tcp or udp. In fact, the main conceptual mistake is to think of ICMP as an application protocol that uses either UDP or TCP; it uses neither. To match all ICMP messages, for instance, use the permit icmp any any command in an extended named ACL. Example Issue: Inbound ACL Filters Routing Protocol Packets A router bypasses outbound ACL logic for packets the router itself generates. That might sound like common sense, but it is important to stop and think about that fact in context. A router can have an outgoing ACL, and that ACL can and will discard packets that the router receives in one interface and then tries to forward out some other interface. But if the router creates the packet, for instance, for a routing protocol message, the router bypasses the outbound ACL logic for that packet. However, a router does not bypass inbound ACL logic. If an ACL has an inbound ACL enabled, and a packet arrives in that interface, the router checks the ACL. Any and all IPv

Request toward S1, and the ICMP Echo Reply back toward R1? Routers bypass their own outbound ACLs for packets generated by the router, as shown in Figure D-7. Even though ACL A exists as an outgoing ACL on Router R1, R1 bypasses its own outgoing ACL logic of ACL A for the ICMP Echo Requests generated by R1. ping S1 ignore ACL A H1 G0/0 H2 R1 A S0/0/0 D Figure D-7 B A S0/0/1 R2 G0/2 SW1 S1 C R1 Rings Outgoing ACL for Packets Created by Its Own ping Command Router Self-Ping of a Serial Interface IPv4 Address The previous example uses a router's ping command when pinging a host. However, network engineers often need to ping router IP addresses, including using a self-ping. The term self-ping refers to a ping of a device's own IPv4 address. And for point-to-point serial links, a self-ping actually sends packets over the serial link, which causes some interesting effects with ACLs. When a user issues a self-ping for that local router's serial IP address, the router actually sends the ICMP echo request out the link to the other router. The neighboring router then receives the packet and routes the packet with the ICMP echo request back to the original router. Figure D-8 shows an example of a self-ping (ping 172.16.4.1) of Router R1's own IP address on a point-to-point serial link, with the ICMP echo request out the link to Router R2. At step 2, R2 treats it like any other packet not destined for one of R2's own IPv4 addresses: R2 routes the packet. Where? Right back to Router R1, as shown in the figure. Now think about those four ACLs in the earlier figures compared to Figure D-8. R1 generates the ICMP echo request, so R1 bypasses outbound ACL A. ACLs B, C, and D would filter the packet. Note that the packet sent by R2 back to R1 is not generated by R2 in this case; R2 is just routing R1's original packet back to R1. Appendix D: Topics from Previous Editions 21 ping 172.16.4.1 Send Out S0/0/0 I Echo Request Echo Request R1 Figure D-8 172.16.4.1 S0/0/0 2 172.16.4.2 S0/0/0 Destination 172.16.4.1 Route Out S0/0/1 R2 G0/2 SW1 S1 The First Steps in a Self-Ping on a R1, for R1's S0/0/0 IP Address A self-ping of a serial interface actually tests many parts of a point-to-point serial link, as follows: ■ The link must work at Layers 1, 2, and 3. Specifically, both routers must have a working (up/up) serial interface, with correct IPv4 addresses configured. ■ ACLs B, C, and D must permit the ICMP echo request and reply packets. So, when troubleshooting, if you choose to use self-pings and they fail, but the serial interfaces are in an up/up state, do not forget to check to see whether the ACLs have filtered the Internet Control Management Protocol (ICMP) traffic. Router Self-Ping of an Ethernet Interface IPv4 Address A self-ping of a router's own Ethernet interface IP address works a little like a self-ping of a router's serial IP address, but with a couple of twists: ■ Like with serial interface, the local router interface must be working (in an up/up state); otherwise, the ping fails. ■ Unlike serial interfaces, the router does not forward the ICMP messages physically out the interface, so security features on neighboring switches (like port security) or routers (like ACLs) cannot possibly filter the messages used by the ping command. ■ Like serial interfaces, an incoming IP ACL on the local router does process the router self-ping of an Ethernet-based IP address. Figure D-9 walks through an example. In this case, R2 issues a ping 172.16.2.2 command to ping its own G0/2 IP address. Just like with a self-ping on serial links, R2 creates the ICMP echo request. However, R2 basically processes the ping down its own TCP/IP stack and back up again, with the ICMP echo never leaving the router's Ethernet interface. R2 does check the Ethernet interface status, showing a failure if the interface is not up/ur. R2 does not apply outbound ACL logic to the packet, because R2 created the packet, but R2 will apply inbound ACL logic to the packet, as if the packet had been physically received on the interface. D 22 CCNA 200-301 Official Cert Guide, Volume 2 ping 172.16.2.2 Check G0/2 Status Check Incoming ACL G0/2 172.16.2.2 R2 SW1 S1 F Figure D-9 Self-Ping of a Router's Ethernet Address NOTE The content under the heading "Implementing HSRP" was most recently published for the 200-105 Exam in 2016, in Chapter 20 of the Cisco CCNA ICND2 200-105 Official Cert Guide. Implementing HSRP The goal of this section is to show enough of the operation of each tool to reinforce your understanding of configuring the basic functions of HSRP. Configuring and Verifying Basic HSRP HSRP configuration requires only one command on the two (or more) routers that want to share default router responsibilities with HSRP: the standby group ip virtual-ip interface subcommand. The first value defines the HSRP group number, which must match on both routers. The group number lets one router support multiple HSRP groups at a time on the same interface, and it allows the routers to identify each other based on the group. The command also configures the virtual IP address shared by the routers in the same group; the virtual IP address is the address the hosts in the VLAN use as their default gateway. Example D-11 shows a configuration example where both routers use group 1, with virtual IP address 10.1.1.1, with the standby 1 ip 10.1.1.1 interface subcommand. Example D-11 HSRP Configuration on R1 and R2. Sharing IP Address 10.1.1.1 R1# show running-config I Lines omitted for brevity interface GigabitEthernet0/0 ip address 10.1.1.1 255.255.255.0 standby version 2 standby 1 ip 10.1.1.1 standby 1 priority 110 standby 1 name HSRP-group-for-book I The following configuration, on R2, is identical except for the HSRP priority and the I interface IP address R2# show running-config I Lines omitted for brevity interface GigabitEthernet0/0 ip address 10.1.1.1 255.255.255.0 Appendix D: Topics from Previous Editions 23 standby version 2 standby 1 ip 10.1.1.1 standby 1 name HSRP-group-for-book The configuration shows other optional parameters, as well. For instance, R1 has a priority of 110 in this group, and R2 defaults to 100. With HSRP, if the two routers are brought up at the same time, the router with the higher priority wins the election to become the active router. The configuration also shows a name that can be assigned to the group (when using show commands) and a choice to use HSRP Version 2. (This chapter provides more details on these settings in the coming pages.) Once configured, the two routers negotiate the HSRP settings and choose which router will currently be active and which will be standby. With the configuration as shown, R1 will win the election and become active because of its higher (better) priority. Both routers reach the same conclusion, as confirmed with the output of the show standby brief command on both R1 and R2 in Example D-12. Example D-12 HSRP Status on R1 and R2 with show standby brief I First, the group status as seen from R1 R1# show standby brief P indicates configured to preempt. I Interface Grp Pri P State Active Standby Virtual IP Gi0/0 1 110 local 10.1.1.1 129 10.1.1.1 Active I The output here on R2 shows that R2 agrees with R1. R2# show standby brief P indicates configured to preempt. I Interface Grp Pri P State Gi0/0 1 100 Active Standby 10.1.1.9 Standby Virtual IP local 10.1.1.1 The show standby brief command packs a lot of detail in the output, so take your time and work through the highlighted fields. First, look at the Grp column for each command. This lists the HSRP group number, so when looking at output from multiple routers, you need to look at the lines with the same group number to make sure the data relates to that one HSRP group. In this case, both routers have only one group number (1), so it is easy to find the information. Each line of output lists the local router's view of the HSRP status for that group. In particular, based on the headings, the show standby brief command identifies the following: Interface: The local router's interface on which the HSRP group is configured Grp: The HSRP group number Pri: The local router's HSRP priority State: The local router's current HSRP state Active: The interface IP address of the currently active HSRP router (or "local" if the local router is HSRP active) Standby: The interface IP address of the currently standby HSRP router (or "local" if the local router is HSRP standby) Virtual IP: The virtual IP address defined by this router for this group D 24 CCNA 200-301 Official Cert Guide, Volume 2 For instance, following the highlighted text in Example D-12, R2 believes that its own current state is standby, that the router with interface address 10.1.1.9 is active (which happens to be Router R1), with a confirmation that the "local" router (R2, on which this command was issued) is the standby router. In comparison, the show standby command (without the brief keyword) lists a more detailed description of the current state, while repeating many of the facts from the show standby brief command. Example D-13 shows an example of the new information with the show standby command, listing several counters and timers about the HSRP protocol itself, plus the virtual MAC address 0000.0c9f.0001. Example D-13 HSRP Status on R1 and R2 with show standby R1# show standby GigabitEthernet0/0 - Group 1 (version 2) State is Standby 4 state changes, last state change 00:12:53 Virtual IP address is 10.1.1.1 Active virtual MAC address is 0000.0c9f.0001 Local virtual MAC address is 0000.0c9f.0001 (v2 default) Hello time 3 sec, hold time 10 sec Next hello sent in 3.696 secs Preemption disabled Active router is local Standby router is 10.1.1.129, priority 100 (expires in 8.096 sec) Priority 110 (configured 110) Group name is "HSRP-group-for-book" (crgd) I The output here on R2 shows that R2 agrees with R1. R2# show standby GigabitEthernet0/0 - Group 1 (version 2) State is Standby 4 state changes, last state change 00:12:05 Virtual IP address is 10.1.1.1 Active virtual MAC address is 0000.0c9f.0001 Local virtual MAC address is 0000.0c9f.0001 (v2 default) Hello time 3 sec, hold time 10 sec Next hello sent in 0.352 secs Preemption disabled Active router is 10.1.1.9, priority 110 (expires in 9.136 sec) MAC address is 0200.0101.0101 Standby router is local Priority 100 (default 100) Group name is "HSRP-group-for-book" (crgd) HSRP Active Role with Priority and Preemption HSRP defines some rules to determine which router acts as the active HSRP router and which acts as standby. Those rules also define details about when a standby router should preempt Appendix D: Topics from Previous Editions 25 take over as active. The following list summarizes the rules: following the list, this section takes a closer look at those rules and the related configuration settings. First, the HSRP rules. When a router (call it the local router) has an HSRP-enabled interface, and that interface comes up, the router sends HSRP messages to negotiate whether it should be active or standby. When it sends those messages, if it... Step 1. ...discovers no other HSRP routers in the subnet, the local router becomes the active router. Step 2. ...discovers an existing HSRP router, and both are currently negotiating to decide which should become the HSRP active router, the routers negotiate, with the router with the highest HSRP priority becoming the HSRP active router. Step 3. ...discovers an existing HSRP router in the subnet, and that router is already acting as the active router: A. If configured with no preemption (the default; no standby preempt), the local router becomes a standby router, even if it has a better (higher) priority. B. If configured with preemption (standby preempt), the local router checks its priority versus the active router; if the local router priority is better (higher), the local router takes over (preempts) the existing active router to become the new active HSRP router. Steps 1 and 2 in the list are pretty obvious, but steps 3A and 3B could use a little closer look. For instance, the examples so far in this chapter show R1's G0/0 with a priority of 110 versus R2's G0/0 with priority 100. The show commands in Example D-13 show that R1 is currently the HSRP active router. That same example also lists a line for both R1 and R2 that confirms "preemption disabled," which is the default. To show a test of step 3A logic, Example D-14 shows a process by which R1's G0/0 interface is disabled and then enabled again, but after giving Router R2 this enough time to take over and become active. That is, R1 comes up but R2 is already HSRP active for group 1. The bottom of the example lists output from the show standby brief command from R2, confirming that R2 becomes HSRP active and R1 becomes standby (10.1.1.9), proving that R1 does not preempt R2 in this case. Example D-14 Showing How No Preemption Keeps R1 as Standby After R1 Recovers I First, R1's G0/0 is disabled and enabled; the ending log message shows a standby ! state. R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# interface gigabitEthernet 0/0 R1(config-if)# shutdown *Mar 8 18:10:29.242: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Active -> Init *Mar 8 18:10:31.205: %LINK-5-CHANGED: Interface GigabitEthernet0/0, changed state to administratively down *Mar 8 18:10:32.205: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEther net/0, changed state to down D 26 CCNA 200-301 Official Cert Guide, Volume 2 R1(config-if)# no shutdown R1(config-if)# ^Z R1# *Mar 8 18:11:08.356: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Speak -> Standby I Now from R2, note R2 is active, and 10.1.1.9 (R1) is standby R2# show standby brief P indicates configured to preempt. I Interface Grp Pri P State Gi0/0 1 100 Active Local Active Standby Virtual IP 10.1.1.9 10.1.1.1 R1 had been configured with preemption for that previous scenario, R1 would have taken over from R2 when R1's interface came back up. Example D-15 shows exactly that. Before the output in Example D-15 was gathered, the network had been put back to the same beginning state as at the beginning of Example D-14, with R1 active and R2 as standby. Within Example D-15, R1's interface is shut down, then configured with preemption using the standby 1 preempt command, enabling preemption. Then, after enabling the interface again, R1 takes over as HSRP active, as shown at the bottom of the example's show standby brief command from R2. That output now shows the local router's state as Standby, and the active as 10.1.1.9 (R1). Example D-15 Recovery Showing How Preemption Causes R1 to Take Over As Active upon I First, R1's G0/0 is disabled and enabled; the ending log message shows a standby ! state. R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# interface gigabitEthernet 0/0 R1(config-if)# shutdown *Mar 8 18:10:29.242: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Active -> Init *Mar 8 18:10:31.205: %LINK-5-CHANGED: Interface GigabitEthernet0/0, changed state to administratively down *Mar 8 18:10:32.205: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEther net/0, changed state to down R1(config-if)# standby 1 preempt R1(config-if)# no shutdown R1(config-if)# ^Z R1# *Mar 8 18:19:14.356: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Listen -> Active I Now from R2, note it is active, and 10.1.1.9 (R1) is standby *Mar 8 18:18:55.948: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Standby -> Active *Mar 8 18:19:14.528: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Active -> Speak Appendix D: Topics from Previous Editions 27 *Mar 8 18:19:26.298: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Speak -> Standby R2# show standby brief P indicates configured to preempt. I Interface Grp Pri P State Gi0/0 1 100 Active Standby 10.1.1.9 Standby Virtual IP local 10.1.1.1 Note that it is the preemption setting on the router that is taking over (preempting) that determines if preemption happens. For instance, in this case, R1 came up when R2 was active; R1 was set to preempt, so R1 preempted R2. HSRP Versions Cisco IOS on routers and Layer 3 switches supports two versions of HSRP: versions 1 and 2. The versions have enough differences, like multicast IP addresses used and message formats, so that routers in the same HSRP group must use the same version. If two routers configured to be in the same HSRP group mistakenly configure to use different versions, they will not understand each other and ignore each other for the purposes of HSRP. To configure the version, each interface/subinterface uses the standby version {1 | 2} interface subcommand. Note that the HSRP group number is not included in the command, because it sets the version for all HSRP messages sent out that interface/subinterface. There are some good reasons to want to use the more recent HSRP version 2 (HSRPv2). For example, HSRPv1 existed before IPv6 became popular. Cisco enhanced HSRP to version 2 in part to make IPv6 support possible. Today, to use HSRP with IPv6 requires HSRPv2. As another example of a benefit of HSRPv2, HSRP uses a Hello message, similar in concept to routing protocols, so that HSRP group members can realize when the active router is no longer reachable. HSRPv2 allows for shorter Hello timer configuration (as low as a small number of milliseconds), while HSRPv1 typically had a minimum of 1 second. So, HSRPv2 can be configured to react more quickly to failures with a lower Hello timer. Beyond IPv6 support and shorter Hello timer options, other differences for version 2 versus version 1 include a different virtual MAC address base value and a different multicast IP address used as the destination for all messages. Table D-3 lists the differences between HSRPv1 and HSRPv2. Table D-3 HSRPv1 Versus HSRPv2 Feature Version 1 Version 2 IPv6 support No Yes Smallest unit for Hello timer Second Millisecond Range of group numbers 0..255 0..4095 MAC address used (xx or xxx is the hex group number) 0000.0C07.ACxx 0000.0C9F.Fxxx IPv4 multicast address used 224.0.0.2 224.0.0.10 Does protocol use a unique identifier for each router? No Yes 28 CCNA 200-301 Official Cert Guide, Volume 2 R1(config-if)# no shutdown R1(config-if)# ^Z R1# *Mar 8 18:11:08.356: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Speak -> Standby I Now from R2, note R2 is active, and 10.1.1.9 (R1) is standby R2# show standby brief P indicates configured to preempt. I Interface Grp Pri P State Gi0/0 1 100 Active Local Active Standby Virtual IP 10.1.1.9 10.1.1.1 If R1 had been configured with preemption for that previous scenario, R1 would have taken over from R2 when R1's interface came back up. Example D-15 shows exactly that. Before the output in Example D-15 was gathered, the network had been put back to the same beginning state as at the beginning of Example D-14, with R1 active and R2 as standby. Within Example D-15, R1's interface is shut down, then configured with preemption using the standby 1 preempt command, enabling preemption. Then, after enabling the interface again, R1 takes over as HSRP active, as shown at the bottom of the example's show standby brief command from R2. That output now shows the local router's state as Standby, and the active as 10.1.1.9 (R1). Example D-15 Recovery Showing How Preemption Causes R1 to Take Over As Active upon I First, R1's G0/0 is disabled and enabled; the ending log message shows a standby ! state. R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# interface gigabitEthernet 0/0 R1(config-if)# shutdown *Mar 8 18:10:29.242: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Active -> Init *Mar 8 18:10:31.205: %LINK-5-CHANGED: Interface GigabitEthernet0/0, changed state to administratively down *Mar 8 18:10:32.205: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEther net/0, changed state to down R1(config-if)# standby 1 preempt R1(config-if)# no shutdown R1(config-if)# ^Z R1# *Mar 8 18:19:14.356: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Listen -> Active I Now from R2, note it is active, and 10.1.1.9 (R1) is standby *Mar 8 18:18:55.948: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Standby -> Active *Mar 8 18:19:14.528: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Active -> Speak Appendix D: Topics from Previous Editions 27 *Mar 8 18:19:26.298: %HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 1 state Speak -> Standby R2# show standby brief P indicates configured to preempt. I Interface Grp Pri P State Gi0/0 1 100 Active Standby 10.1.1.9 Standby Virtual IP local 10.1.1.1 Note that it is the preemption setting on the router that is taking over (preempting) that determines if preemption happens. For instance, in this case, R1 came up when R2 was active; R1 was set to preempt, so R1 preempted R2. HSRP Versions Cisco IOS on routers and Layer 3 switches supports two versions of HSRP: versions 1 and 2. The versions have enough differences, like multicast IP addresses used and message formats, so that routers in the same HSRP group must use the same version. If two routers configured to be in the same HSRP group mistakenly configure to use different versions, they will not understand each other and ignore each other for the purposes of HSRP. To configure the version, each interface/subinterface uses the standby version {1 | 2} interface subcommand. Note that the HSRP group number is not included in the command, because it sets the version for all HSRP messages sent out that interface/subinterface. There are some good reasons to want to use the more recent HSRP version 2 (HSRPv2). For example, HSRPv1 existed before IPv6 became popular. Cisco enhanced HSRP to version 2 in part to make IPv6 support possible. Today, to use HSRP with IPv6 requires HSRPv2. As another example of a benefit of HSRPv2, HSRP uses a Hello message, similar in concept to routing protocols, so that HSRP group members can realize when the active router is no longer reachable. HSRPv2 allows for shorter Hello timer configuration (as low as a small number of milliseconds), while HSRPv1 typically had a minimum of 1 second. So, HSRPv2 can be configured to react more quickly to failures with a lower Hello timer. Beyond IPv6 support and shorter Hello timer options, other differences for version 2 versus version 1 include a different virtual MAC address base value and a different multicast IP address used as the destination for all messages. Table D-3 lists the differences between HSRPv1 and HSRPv2. Table D-3 HSRPv1 Versus HSRPv2 Feature Version 1 Version 2 IPv6 support No Yes Smallest unit for Hello timer Second Millisecond Range of group numbers 0..255 0..4095 MAC address used (xx or xxx is the hex group number) 0000.0C07.AC for HSRPv1 and 0000.0C9F.F for HSRPv2. HSRPv1, with 256 possible HSRP groups per interface, then uses the last two hex digits to identify the HSRP group. For example, an HSRP group 1 using version 1 would use a virtual MAC address that ends in hex 01. Similarly, because HSRPv2 supports 4096 groups per interface, the MAC address reserves three hex digits to identify the group. An HSRP group 1 using version 2 would use a virtual MAC address that ends in hex 001. Note The content under the heading "Gateway Load Balancing Protocol (GLBP)" was most recently published for the 200-105 Exam in 2016, in Appendix K of the Cisco CCNA ICND2 200-105 Official Cert Guide. Gateway Load Balancing Protocol (GLBP) This section first discusses GLBP concepts, followed by GLBP configuration. GLBP Concepts Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP), which were introduced before Gateway Load Balancing Protocol (GLBP), balanced the packet load per subnet. However, because traffic loads vary unpredictably from subnet to subnet, Cisco wanted a First Hop Redundancy Protocol (FHRP) option with better loadbalancing options than just the per-subnet load balancing of HSRP and VRRP. To meet that need, Cisco introduced GLBP. GLBP balances the packet load per host by using an active/active model in each subnet. Each GLBP router in a subnet receives off-subnet packets from some of the hosts in the subnet. Each host still remains unaware of the FHRP, allowing the hosts to configure the same default gateway/router setting and for the hosts to make no changes when a router fails. GLBP creates a world that at first glance looks like HSRP, but with a few twists that let GLBP balance the traffic. Like HSRP, all the routers configure a virtual IP address, which is the IP address used by hosts as their default router. Like with HSRP, hosts use a default router setting that points to the virtual IP address, and that setting does not need to change. GLBP differs from HSRP with regard to the MAC addresses it uses and the Address Resolution Protocol (ARP) process, because GLBP actually uses ARP Reply messages to balance traffic from different hosts through different routers. With GLBP, one router acts in a special role called the active virtual gateway (AVG). The AVG replies to all ARP requests for the virtual IP address. Each router has a unique virtual MAC address, so that the AVG can reply to some ARP Requests with one virtual MAC, and some with the other. As a result, some hosts in the subnet send frames to the Ethernet MAC address of one of the routers, with other hosts sending their frames to the MAC address of the second router. As an example, Figure D-10 shows the process by which a GLBP balances traffic for host A based on the ARP Reply sent by the AVG (R1). The two routers support virtual IP address 10.1.1.1, with the hosts using that address as their default router setting. Appendix D: Topics from Previous Editions 29 10.1.1.1? A 2 AVG 1 VMAC1 ARP: VMAC1 R1 Role Router Address AVG Forwarder Forwarder R1 R1 R2 10.1.1.1 VMAC1 VMAC2 3 Data To VMAC1 VMAC2 .1 Figure D-10 R2 GLBP Directs Host A by Sending Back ARP Reply with R1's VMAC1 The figure shows three messages, top to bottom, with the following action: 1. Host A has no ARP table entry for its default router, 10.1.1.1, so host A sends an ARP Request to learn 10.1.1.1's MAC address. 2. The GLBP AVG, R1 in this case, sends back an ARP Reply. The AVG chooses to include its own virtual MAC address in the ARP Reply, VMAC1. 3. Future IP packets sent by host A are encapsulated in Ethernet frames, destined to VMAC1, so that they arrive at R1. From now on, host A sends off-subnet packets to R1 due to host A's ARP table entry for its default gateway (10.1.1.1). Host A's ARP table entry for 10.1.1.1 now refers to a MAC address on R1 (VMAC1), so packets host A sends off-subnet flow through R1. To balance the load, the AVG answers each new ARP Request with the MAC addresses of alternating routers. Figure D-11 continues the load-balancing effect with the ARP Request for 10.1.1.1 coming from host B. The router acting as AVG (R1) still sends the ARP Reply, but this time with R2's virtual MAC (VMAC2). GLBP Table 1 ARP: 10.1.1.1? B 2 AVG 1 VMAC1 ARP: VMAC2 R1 Role Router Address AVG Forwarder Forwarder R1 R1 R2 10.1.1.1 VMAC1 VMAC2 3 Data To VMAC2 VMAC2 .1 Figure D-10 R2 GLBP Directs Host B by Sending Back ARP Reply with R2's VMAC2 30 CCNA 200-301 Official Cert Guide, Volume 2 Here are the steps in the figure: 1. Host B sends an ARP Request to learn 10.1.1.1's MAC address. 2. The GLBP AVG (R1) sends back an ARP Reply, listing VMAC2. R2's virtual MAC address. 3. For future packets sent off-subnet, host B encapsulates the packets in Ethernet frames, destined to VMAC2, so that they arrive at R2. The process shown in Figures D-10 and D-11 balances the traffic, per host, but the routers must also be ready to take over for the other router if it fails. GLBP refers to each router as a forwarder. When all is well, each router acts as forwarder for its own virtual MAC address, but it listens to GLBP messages to make sure the other forwarders are still working. If another forwarder fails, the still-working forwarder takes over the failed forwarder's virtual MAC address role and continues to forward traffic. Configuring and Verifying GLBP GLBP configuration mimics HSRP configuration to a great degree. Example D-16 shows a GLBP configuration with both routers using GLBP group 1, with virtual IP address 10.1.1.1, with the glbp 1 ip 10.1.1.1 interface subcommand. Example D-16 GLBP Configuration on R1 and R2, Sharing IP Address 10.1.1.1 I First, the configuration on R1 R1# show running-config I Lines omitted for brevity interface GigabitEthernet0/0 ip address 10.1.1.1 255.255.255.0 glbp 1 ip 10.1.1.1 glbp 1 priority 110 glbp 1 name GLBP-group-for-book I The following configuration, on R2, is identical except for ! the interface IP address, and the GLBP priority R2# show running-config I Lines omitted for brevity interface GigabitEthernet0/0 ip address 10.1.1.1 255.255.255.0 glbp 1 ip 10.1.1.1 glbp 1 priority 110 glbp 1 name GLBP-group-for-book I The following configuration, on R2, is identical except for ! the interface IP address, and the GLBP priority R2# show running-config I Lines omitted for brevity interface GigabitEthernet0/0 ip address 10.1.1.1 255.255.255.0 glbp 1 ip 10.1.1.1 glbp 1 priority 110 command version 2 standby 1 ip 10.1.1.1 standby 1 priority 100 However, if either router comes up before the other, that router goes ahead and takes on the AVG role. Sifting through the GLBP show command output takes a little more work than with HSRP, in particular because of the added detail in how GLBP works. First, consider the show glbp brief command on Router R1, as shown in Example D-17. (Note that many show glbp commands have the same options as equivalent HSRP show standby commands.) Appendix D: Topics from Previous Editions 31 Example D-17 GLBP Status on R1 with show glbp brief R1# show glbp brief Interface Grp Fwd Pri State Address Active Router Gi0/0 1 - 110 Active 10.1.1.1 local Standby router 10.1.1.129 Gi0/0 1 - Listen 0007.b400.0101 local Before looking at the right side of the output, first consider the context for a moment. This example lists a heading line and three rows of data. These data rows are identified by the Grp and Fwd headings, short for Group and Forwarder. With only one GLBP group configured, R1 lists only one for group 1. More important, each row defines details about a different part of what GLBP does, as follows: Fwd is -: This line refers to none of the forwarders, and instead describes the AVG. Fwd is 1: This line describes GLBP forwarder (router) 1. Fwd is 2: This line describes GLBP forwarder (router) 2. The output usually lists the line about the AVG first, as noted with a dash in the Forwarder column. Now look at the highlighted portions on the right of Example D-17. This line will list the virtual IP address and identify the active AVG and the standby AVG. This particular command, from Router R1, lists R1 itself ("local") as the active router. So, R1 is the current AVG. Each of the next two lines lists status information about one of the forwarder roles; that is, a router that uses a virtual MAC address, receives frames sent to that address, and routes the packets encapsulated in those frames. To that end, the Address column lists MAC addresses, specifically the virtual MAC addresses used by GLBP, and not the interface MAC addresses. Each forwarder row also identifies the router that currently uses the listed virtual MAC in the Active Router column. In Example D-17, 0007.b400.0101 is used by the router with interface IP address 10.1.1.129 (which happens to be R2), 0007.b400.0102 is supported by the local router (the router on which the show command was issued), which is R1. The brief output of the show glbp brief command lists many details, but it takes some effort to learn how to sift through it all. For more perspective on the output, Example D-18 lists this same show glbp brief command, this time on R2. Note that the Fwd column again identifies the first line of output as being about the AVG, with the next two lines about the two forwarders. Example D-18 GLBP Status on R2 with show glbp brief R2# show glbp brief Interface Grp Fwd Pri State Address Active Router Gi0/0 1 - 100 Standby 10.1.1.1 local Gi0/0 1 1 - Active 0007.b400.0101 local - Gi0/0 1 2 - Listen 0007.b400.0102 10.1.1.9 - The State column in the output in Examples D-17 and D-18 can pull the GLBP concepts together. First, to define the meaning of the state values, the following short list defines D 32 CCNA 200-301 Official Cert Guide, Volume 2 the states expected for the first line of output, about the AVG, and then about each GLBP forwarder: AVG: One router should be the active AVG, with the other acting as standby, ready to take over the AVG role if the AVG fails. Each forwarder: One router should be active, while the other should be listening, ready to take over that virtual MAC address if that forwarder fails. Table D-4 collects the values of the State column from Examples D-17 and D-18 for easier reference by side. Note that, indeed, each line has either an active/standby pair (for the AVG) or an active/listen pair (for the forwarder function). Table D-4 Comparing Local State in show glbp brief Commands Rows Is About... Fwd Column Value R1 State R2 State AVG - Active Standby Forwarder 1 2 Active Listen Finally, the show glbp command lists a more detailed view of the current GLBP status. Example D-19 shows a sample from Router R1. Note that the first half of the output has similar information compared to HSRP's show standby command, plus it lists the IP and MAC addresses of the routers in the GLBP group. Then, the end of the output lists a group of messages per GLBP forwarder. Example D-19 GLBP Status on R1 with show glbp R1# show glbp GigabitEthernet0/0 - Group 1 State is Active 2 state changes, last state change 00:20:59 Virtual IP address is 10.1.1.1 Hello time 3 sec, hold time 10 sec Next hello sent in 2.112 secs Redirect time 600 sec, forwarder timeout 14400 sec Preemption disabled Active is local Standby is 10.1.1.129, priority 100 (expires in 8.256 sec) Priority 110 (configured) Weighting 100 (default 100), thresholds: lower 1, upper 100 Load balancing: round-robin IP redundancy name is "GLBP-group-for-book" Group members: 0200.0101.0101 (10.1.1.9) local 0200.0202.0202 (10.1.1.129) There are 2 forwarders (1 active) Forwarder 1 State is Listen 2 state changes, last state change 00:20:34 Appendix D: Topics from Previous Editions 33 MAC address is 0007.b400.0101 (owner) Owner ID is 0200.0202.0202 Redirection enabled, 598.272 sec remaining (maximum 600 sec) Time to live: 14398.272 sec (Maximum 14400 sec) Preemption enabled, min delay 30 sec Active is 10.1.1.129 (primary), weighting 100 (expires in 8.352 sec) Client selection count: 1 Forwarder 2 State is Active 1 state change, last state change 00:24:25 MAC address is 0007.b400.0102 (default) Owner ID is 0200.0101.0101 Redirection enabled Preemption enabled, min delay 30 sec Active is local, weighting 100 Client selection count: 1 NOTE The content under the heading "Implementing Simple Network Management Protocol" was most recently published for the 200-105 Exam in 2016, in Chapter 26 of the Cisco CCNA ICND2 200-105 Official Cert Guide. Implementing Simple Network Management Protocol This section includes details of how to implement SNMPv2c and SNMPv3. Implementing SNMP Version 2c The exam topics mention SNMPv2c and SNMPv3 by name. As it turns out, SNMPv1 and SNMPv2c configuration is very similar, because both use communities. SNMPv3 varies quite a bit, mainly to implement the better SNMPv3 security features. This next section shows how to configure and verify SNMPv2c. Configuring SNMPv2c Support for Get and Set SNMP configuration in Cisco IOS routers and switches works a little differently than many other IOS features. First, the SNMP configuration exists in a series of global commands; there is no SNMP agent configuration mode in which to collect subcommands. Secondly, no single command enables the SNMP agent. Instead, IOS typically defaults for the SNMP agent to be disabled. Then, the first time an snmp-server global command is configured, IOS enables the SNMP agent. NOTE To disable the SNMP agent, you must remove all the snmp-server commands. You can do this with a single no snmp-server command (with no parameters). With that backdrop, a typical SNMPv2c configuration requires only one or two settings. To be useful, the agent needs at least a read-only (RO) community string. The agent will not reply to SNMPv2c Get messages without at least the RO community string configured. The network engineer may also want the agent to have a read-write (RW) community string, to support Set messages. D 34 CCNA 200-301 Official Cert Guide, Volume 2 NOTE When configuring an RW community, use some caution: configuring an RW community means that you have defined a clear-text password that can be used to configure many settings on the router or switch. The following checklist details the commands used to configure SNMPv2c on a Cisco router or switch. This list shows the method to configure the RO and RW communities, plus a few optional but common settings (location and contact information). Step 1. Use the snmp-server community communitystring RO [ipv6 acl-name] [acl-name] command in global configuration mode to enable the SNMP agent (if not already started), set the read-only community string, and restrict incoming SNMP messages based on the optional referenced IPv4 or IPv6 ACL. Step 2. (Optional) Use the snmp-server community communitystring RW [ipv6 acl-name] [acl-name] command in global configuration mode to enable the SNMP agent (if not already started), set the read-write community string, and restrict incoming SNMP messages based on the optional referenced IPv4 or IPv6 ACL. Step 3. (Optional) If referenced by an snmp-server community command, configure an IPv4 or IPv6 ACL, with the same name or number referenced by the snmp-server community command, with the ACL permitting by matching the source IPv4 or IPv6 address of the allowed SNMP management hosts. Step 4. (Optional) Use the snmp-server location text-describing-location command in global configuration mode to document the location of the device. Step 5. (Optional) Use the snmp-server contact contact-name command in global configuration mode to document the person to contact if problems occur. NOTE In the SNMP model, the SNMP agent acts as a server, with the NMS (SNMP Manager) acting as an SNMP client by requesting information with Get messages. The IOS snmp-server command happens to emphasize the idea that the SNMP agent on a router or switch acts as the SNMP server. Example D-20 shows a sample configuration based on Figure D-12. The examples in this section come from Router R1, although the exact same SNMP configuration syntax could be used in the LAN switches or in R2. (The configuration of the location information would likely differ for each device, however.) Note that the configuration creates an IPv4 ACL that permits traffic with source IP address 10.1.3.3, which is the address of the NMS shown in the figure. It then defines read-only and read-write communities, along with the location and contact name for the router. NMS 10.1.3.3 SW1 G0/0 R1 G0/1 SW2 G0/0 R2 G0/1 SW3 NMS 10.1.3.4 Figure D-12 Sample Network for SNMP Examples, with NMS at 10.1.3.3 Appendix D: Topics from Previous Editions 35 Example D-20 Configuring SNMP Version 2c on Router R1 to Support Get and Set ip access-list standard ACL_PROTECTSNMP permit host 10.1.3.3 ip snmp-server community secretROpw RO ACL_PROTECTSNMP snmp-server community secretRWpw RW ACL_PROTECTSNMP snmp-server location Atlanta snmp-server contact Tyler B To begin managing Router R1 (or any of the other devices that use the same community strings), the SNMP manager at address 10.1.3.3 now needs to configure the community strings listed in Example D-20. Configuring SNMPv2c Support for Trap and Inform For an SNMPv2c agent in a router or switch to be able to send unsolicited notifications to an SNMP manager (that is, to send Trap and Inform messages), the device needs to be configured with the snmp-server host command. This command references the NMS to which the Traps or Informs should be sent, along with the SNMP version. Beyond telling the SNMP agent the hostname or address of the NMS, the agent typically needs to know the notification community string used by the NMS. Think of the RO and RW community strings as protecting the SNMP agent from the messages originated by an NMS (Get or Set Requests), so the agent requires the NMS to supply the correct RO or RW community string. For Traps and Informs, the NMS can protect itself from the Trap and Inform messages originated by SNMP agents by requiring those agents to include the notification community with those messages. The agent can configure this value on the snmp-server host command as well. The following list details the command to enable the sending of SNMPv2c Trap or Inform messages to an NMS: Step 1. Use the snmp-server host (hostname | ip-address) [informs] version 2c notification-community command in global configuration mode to configure the SNMP agent to send either SNMPv2c Traps (default) or Informs to the listed host. Use this command once for each host to which this device should send Traps. Step 2. Use the snmp-server enable traps command in global configuration mode to enable the sending of all supported types of Trap and Inform messages. Example D-21 shows a sample configuration. In most cases, you would send either Traps or Informs to a particular NMS, but not both. So, for this example, the configuration shows how to configure to send Traps to one host (10.1.3.3), and Informs to another host (10.1.3.4). Note that this configuration is added to Router R1 from Figure D-12, but it could have been added to Router R2 or to any of the LAN switches as well. Example D-21 Configuring SNMP Version 2c on Router R1 to Support Sending Traps snmp-server host 10.1.3.3 version 2c secretTRAPpw snmp-server host 10.1.3.4 informs version 2c secretTRAPpw snmp-server enable traps D 36 CCNA 200-301 Official Cert Guide, Volume 2 Verifying SNMPv2c Operation Example D-22 displays some of the status information based on the configuration seen in the previous two examples. The variations on the show snmp command highlight several configuration settings. For example, the show snmp community command repeats the community string values, with reference to any attached IPv4 or IPv6 ACLs. The show snmp host command lists the IP address or hostname of the NMS referenced by each snmp-server host configuration command. Example D-22 Confirming SNMPv2c Configuration Settings on Router R1 R1# show snmp community Community Name: secretROpw Community Index: secretROpw Community SecurityName: secretROpw storage-type: nonvolatile active access-list: ACL_PROTECTSNMP Community Name: secretTRAPpw Community Index: secretTRAPpw Community SecurityName: secretTRAPpw storage-type: nonvolatile active access-list: ACL_PROTECTSNMP Community Name: secretRWpw Community Index: secretRWpw Community SecurityName: secretRWpw storage-type: nonvolatile active access-list: ACL_PROTECTSNMP Community Name: secretRWpw Community Index: secretRWpw storage-type: nonvolatile active R1# show snmp location Atlanta R1# show snmp contact Tyler B R1# show snmp host Notification host: 10.1.3.3 user: secretTRAPpw Notification host: 10.1.3.3 user: secretTRAPpw udp-port: 162 type: inform security model: v2c udp-port: 162 type: trap security model: v2c The show snmp command takes the opposite approach from the commands in Example D-22, focusing almost completely on status and counter information, rather than repeating configuration settings. This command lists dozens of lines of detailed information, so the sample in Example D-23 shows just enough of the output to give you a sense of the kinds of information found there, with comments following the example. Appendix D: Topics from Previous Editions 37 Example D-23 Finding SNMPv2c Message Load on Router R1 R1# show snmp PDUStats: FTx162883H0 Contact: Tyler B Location: Atlanta 17355 SNMP packets input 0 Bad SNMP version errors 9 Unknown community name 0 Illegal operation for community name supplied 2 Encoding errors 51949 Number of requested variables 2 Number of altered variables 3740 Get-request PDUs 3954 Get-next PDUs 7 Set-request PDUs 0 Input queue packet drops (Maximum queue size 1000) 7850 SNMP packets output 0 Too big errors (Maximum packet size 1500) 0 No such name errors 0 Bad values errors 0 General errors 7263 Response PDUs 126 Trap PDUs I Lines omitted for brevity The output in Example D-23 was taken from Router R1 as shown in the earlier examples, after doing some testing from the NMS at address 10.1.3.3. The highlighted items point out the number of SNMP packets received (input) and sent (output), as well as the number of requested MIB variables—that is, the number of variables requested in different SNMP Get requests. (Note that SNMP also supports the GetNext and GetBulk commands, so a single NMS user click can cause the NMS to Get many variables from an agent; thus, it is not unusual for the requested variables counter to get very large.) The output also shows that seven Set requests were received, resulting in two changes to variables. The

Verifying SNMPv3 Configuration Settings R3# show snmp user User name: Youdda1 Engine ID: 80000090300D48CB57D8200 storage-type: nonvolatile active Authentication Protocol: None Privacy Protocol: None Group name: BookGroup1 User name: Youdda2 Engine ID: 80000090300D48CB57D8200 storage-type: nonvolatile active Authentication Protocol: MD5 Privacy Protocol: MD5 Privacy Protocol: None Group name: BookGroup1 ID: 80000090300D48CB57D8200 storage-type: nonvolatile active Authentication Protocol: SHA Privacy Protocol: AES128 Group name: BookGroup3 In particular, work through the highlighted output for users Youdda1, Youdda2, and Youdda6, as compared to the configuration in Example D-26. All the highlighted entries basically repeat the settings from the configuration. Example D-28 lists output from the show snmp group command, which also confirms configuration settings from Example D-26. The most challenging thing to find in this output is what is missing, rather than what is there. Note that this command does not list the SNMP usernames that happen to refer to this group. Also, for groups that do not use an ACL, there is no obvious text that states that no ACL is used. Make sure to compare the output for BookGroup1, which uses an ACL, and the output for BookGroup2, which does not use an ACL. Appendix D: Topics from Previous Editions 43 Verifying SNMPv3 Using show snmp group Example D-28 R3# show snmp group groupname: BookGroup1 security model:v3 noauth contextname: storage-type: nonvolatile readview : v1default notifyview: row status: active access-list: ACL_PROTECTSNMP groupname: BookGroup2 security model:v3 auth contextname: storage-type: nonvolatile readview : v1default writview: notifyview: row status: active ! Lines omitted for brevity Implementing SNMPv3 Notifications (Traps and Informs) SNMP agents can use SNMPv3 to send unsolicited notifications—Trap and Inform messages—to SNMP managers. SNMPv2c uses communities, in this case using the SNMPv2c notification community concept. SNMPv3 uses the same security levels just discussed, but as applied to SNMPv3 notifications. To configure an SNMPv3 agent to send notifications, you add the security level and the username to the snmp-server host command. That configuration links to the same kinds of snmp-server user commands discussed earlier in this section, which in turn link to an snmp-server group command. Figure D-16 shows how the commands connect to each other. snmp-server group groupname v3 noauth [auth | priv] v2 snmp-server user username write viewname access [ipv6] aclname 3 groupname v3 auth md5 password sha password DES keyvalue 3DES keyvalue AES [128|192|256] keyvalue 1 snmp-server host Figure D-16 address D-29 shows 3 noauth [auth | priv] username Connecting SNMPv3 Notification Configuration with User and Group NOTE IOS allows you to configure commands that refer to the correct username and group name, but with different security levels, with no error messages. However, communication with the NMS then fails. Example D-29 shows a few examples of configuration notifications that use SNMPv3. The samples rely on the SNMPv3 usernames and groups as defined in Example D-26. Feel free 44 CCNA 200-301 Official Cert Guide, Volume 2 to refer back to that example, and check to make sure that each snmp-server host command in Example D-29 refers to the correct SNMP security level used by each linked snmpserver group command. Example D-29 Verifying SNMPv3 Configuration Settings ! The group uses noauth, so the user Youdda1 has no auth nor priv keyword snmp-server enable traps snmp-server host 10.1.3.3 version 3 noauth Youdda1 ! Traps w/ noauth snmp-server host 10.1.3.4 informs version 3 auth Youdda2 ! Informs w/ auth snmp-server host 10.1.3.5 version 3 priv Youdda4 ! Traps w/ priv As always, the show snmp command lists the counters that show how many messages flow, including the number of Trap and Inform messages sent by the SNMP agent. To verify the configuration of SNMPv3 notification to NMS hosts, use the show snmp host command. Example D-30 shows the results after configuring Example D-29; note that almost all the fields in Example D-30 repeat the configuration parameters from Example D-29. Example D-30 Verifying SNMPv3 Configuration Settings R3# show snmp host Notification host: 10.1.3.4 udp-port: 162 user: Youdda2 security model: v3 auth Notification host: 10.1.3.5 udp-port: 162 user: Youdda1 security model: v3 noauth Notification host: 10.1.3.5 udp-port: 162 user: Youdda4 security model: v3 priv type: inform type: trap type: trap Summarizing SNMPv3 Configuration SNMPv3 configuration has many parameters to choose from in several commands. As a result, putting the commands into a configuration checklist earlier in this section did not work for learning, so the text instead spelled out the pieces little by little. Now that you have seen how to configure the individual pieces, this configuration checklist summarizes all the different SNMPv3 configuration options discussed in this chapter, for easier review. Step 1. Use the snmp-server group groupname v3 [noauth | auth | priv] [write v1default] [access [ipv6] acl-name] command in global configuration mode to enable the SNMP agent (if not already started), create a named SNMPv3 group of security settings, set the security level, optionally override the default write view with the same view as defaulted for use as the read MIB view (v1default), and optionally restrict incoming SNMP messages based on the optional referenced IPv4 or IPv6 ACL. Step 2. To configure users whose referenced SNMPv3 group has a security level of noauth, use the snmp-server user username groupname v3 command in global configuration mode, making sure to reference an SNMPv3 group with security level of noauth configured. Appendix D: Topics from Previous Editions 45 Step 3. To configure users whose referenced SNMPv3 group use the security level of auth: A. Use the snmp-server user username groupname v3 auth md5 password command in global configuration mode to configure the user and authentication password, and to choose to use MD5 as the authentication hash algorithm. B. Alternatively, use the snmp-server user username groupname v3 auth sha password command in global configuration mode to configure the user and authentication password, and to choose to use SHA as the authentication hash algorithm. Step 4. To configure users that use the security level of priv, you will add parameters to the end of the snmp-server user command syntax as configured in step 3, as follows: A. Add the priv des encryption-key parameters in global configuration mode to the end of the snmp-server user command, to enable the use of DES as the encryption algorithm and to set the encryption key. B. Add the priv 3des encryption-key parameters in global configuration mode to the end of the snmp-server user command, to enable the use of triple DES (3DES) as the encryption algorithm and to set the encryption key. C. Add the priv aes [128 | 192 | 256] encryption-key parameters in global configuration mode to the end of the snmp-server user command, to enable the use of AES as the encryption algorithm, to set the length of the encryption key in bits, and to set the seed for the encryption key. Step 5. Enable the SNMP agent to send notification messages (Traps and/or Informs) to a NMS as follows: A. Use the snmp-server host [hostname | ip-address] [informs | traps] version 3 [noauth | auth | priv] username command in global configuration mode to configure the SNMP agent to send SNMPv3 Traps to the listed host, using the listed username. Use this command once for each host to which this device should send Traps. Include the informs keyword to send Informs; the traps keyword is the default setting. Use the same security level setting as the link SNMPv3 group. B. Use the snmp-server enable traps command in global configuration mode to enable the sending of all supported notifications to all hosts defined in snmp-server host commands. Note that if you review this checklist and get lost, make sure to review and study this section again. SNMPv3 configuration uses a lot of different parameters on three different commands, so it is easy to get lost. The checklist is best used for review once you have a good understanding of the commands. NOTE The content under the heading “Analyzing LAN Physical Standard Choices” was most recently published for the 100-105 Exam in 2016, in Chapter 10 of the Cisco CCENT/CCNA ICND1 100-105 Official Cert Guide. D 46 CCNA 200-301 Official Cert Guide, Volume 2 Analyzing LAN Physical Standard Choices When you look at the design of a network designed by someone else, you can look at all the different types of cabling used, the different types of switch ports, and the Ethernet standards used in each case. Then ask yourself: Why did they choose a particular type of Ethernet link for each link in the network? Asking that question, and investigating the answer, starts to reveal much about building the physical campus LAN. The IEEE has done an amazing job developing Ethernet standards that give network designers many options. Two themes in particular have helped Ethernet grow over the long term: ■ The IEEE has developed many additional 802.3 standards for different types of cabling, different cable lengths, and for faster speeds. ■ All the physical standards rely on the same consistent data-link details, with the same standard frame formats. That means that one Ethernet LAN can use any types of physical links to meet distance, budget, and cabling needs. For example, think about the access layer of the generic design drawings, but now think about cabling and Ethernet standards. In practice, access layer switches sit in a locked wiring closet somewhere on the same floor as the end user devices. Electricians have installed unshielded twisted-pair (UTP) cabling used at the access layer, running from that wiring closet to each wall plate at each office, cubicle, or any place where an Ethernet device might need to connect to the LAN. The type and quality of the cabling installed between the wiring closet and each Ethernet outlet dictate what Ethernet standards can be supported. Certainly, whoever designed the LAN at the time the cabling was installed thought about what type of cabling was needed to support the types of Ethernet physical standards that were going to be used in that LAN. Ethernet Standards Over time, the IEEE has continued to develop and release new Ethernet standards, for new faster speeds and to support new and different cabling types and cable lengths. Figure D-17 shows some insight into Ethernet speed improvements over the years. The early standards up through the early 1990s ran at 10 Mbps, with steadily improving cabling and topologies. Then, with the introduction of Fast Ethernet (100 Mbps) in 1995, the IEEE began ramping up the speeds steadily over the next few decades, continuing even until today. Thickett (DIX) Thimmett (IEEE) Ethernet 10Base-T Fast Ethernet Gigabit Ethernet 10 Gig E 100 Gig E 10M 10M 10M 100M 1G 10G 40G 100G 1980 1985 1990 1995 Figure D-17 2000 2005 2010 Ethernet Standards Timeline NOTE Often, the IEEE first introduces support for the next higher speed using some forms of fiber optic cabling, and later, sometimes many years later, the IEEE completes the work to develop standards to support the same speed on UTP cabling. Figure D-17 shows the earliest standards for each speed, no matter what cabling. Appendix D: Topics from Previous Editions 47 When the IEEE introduces support for a new type of cabling, or a faster speed, they create a new standard as part of 802.3. These new standards have a few letters behind the name. So, when speaking of the standards, sometimes you might refer to the standard name (with letters). For instance, the IEEE standardized Gigabit Ethernet support using inexpensive UTP cabling in standard 802.3ab. However, more often, engineers refer to that same standard as 1000BASE-T or simply Gigabit Ethernet. Table D-6 lists some of the IEEE 802.3 physical layer standards and related names for perspective. Table D-6 IEEE Physical Layer Standards Original IEEE Standard Shorthand Name Informal Names Speed Typical Cabling 802.3 10BASE-T Ethernet 10 Mbps UTP 802.3u 100BASE-T Fast Ethernet 100 Mbps UTP 802.3z 1000BASE-X Gigabit Ethernet, GigE 1000 Mbps (1 Gbps) Fiber 802.3ae 10GBASE-X 10 GigE 10 Gbps Fiber 802.3an 10GBASE-T 10 GigE 10 Gbps UTP 802.3ba 40GBASE-X 40 GigE 40 Gbps Fiber 802.3ba 100GBASE-X 100 GigE 100 Gbps Fiber Choosing the Right Ethernet Standard for Each Link When designing an Ethernet LAN, you can and should think about the topology, with an access layer, a distribution layer, and possibly a core layer. But thinking about the topology does not tell you which specific standards to follow for each link. Ultimately, you need to pick which Ethernet standard to use for each link, based on the following kinds of facts about each physical standard: ■ The speed ■ The maximum distance allowed between devices when using that standard/cabling ■ The cost of the cabling and switch hardware ■ The availability of that type of cabling already installed at your facilities Consider the three most common types of Ethernet today (10BASE-T, 100BASE-T, and 1000BASE-T). They all have the same 100-meter UTP cable length restriction. They all use UTP cabling. However, not all UTP cabling meets the same quality standard, and as it turns out, the faster the Ethernet standard, the higher the required cable quality category needed to support that standard. As a result, some buildings might have better cabling that supports speeds up through Gigabit Ethernet, whereas some buildings may support only Fast Ethernet. The Telecommunications Industry Association (TIA; tianonline.org) defines Ethernet cabling quality standards. Each Ethernet UTP standard lists a TIA cabling quality (called a category) as the minimum category that the standard supports. For example, 10BASE-T allows for Category 3 (CAT3) cabling or better. 100BASE-T requires higher-quality CAT5 cabling, and 1000BASE-T requires even higher-quality CAT5e cabling. (The TIA standards follow a general “higher number is better cabling” in their numbering.) For instance, if an older facility had only CAT5 cabling installed between the wiring closets and each cubicle, the engineers D 48 CCNA 200-301 Official Cert Guide, Volume 2 would have to consider upgrading the cabling to fully support Gigabit Ethernet. Table D-7 lists the more common types of Ethernet and their cable types and length limitations. Table D-7 Ethernet Types, Media, and Segment Lengths (Per IEEE) Ethernet Type Media Maximum Segment Length 10BASE-T TIA CAT3 or better, 2 pairs 100 m (328 feet) 100BASE-T TIA CAT5 UTP or better, 2 pairs 100 m (328 feet) 1000BASE-T TIA CAT5 UTP or better, 4 pairs 100 m (328 feet) 10GBASE-T TIA CAT6a UTP or better, 4 pairs 100 m (328 feet) 10GBASE-T1 Per CAT6 UTP or better, 4 pairs 38–55 m (127–180 feet) 1000BASE-LX Multimode fiber 550 m (1800 feet) 1000BASE-LR Multimode fiber 550 m (1800 feet) 1000BASE-LX-9-micron single-mode fiber 5 km (3.1 miles) 1 The option for 10GBASE-T with slightly less quality CAT6 cabling, but at shorter distances, is an attempt to support 10Gig Ethernet for some installations with CAT6 installed cabling. Ethernet defines standards for using fiber optic cables as well. Fiber optic cables include ultrathin strands of glass through which light can pass. To send bits, the switches can alternate between sending brighter and dimmer light to encode 0s and 1s on the cable. Generally comparing optical cabling versus UTP cabling Ethernet standards, two obvious points stand out. Optical standards allow much longer cabling, while generally costing more for the cable and the switch hardware components. Optical cables experience much less interference from outside sources compared to copper cables, which allows for longer distances. When considering optical Ethernet links, many standards exist, but with two general categories. Comparing the two, the cheaper options generally support distances into the hundreds of meters, using less expensive light-emitting diodes (LED) to transmit data. Other optical standards support much longer distances into multiple kilometers, using more expensive cabling and using lasers to transmit the data. The trade-off is basic: For a given link, how long does the cable need to run, what standards support that distance, and which is the least expensive to meet that need? In reality, most engineers remember only the general facts from tables like Table 10-3: 100 meters for UTP, about 500 meters for multimode fiber, and about 5000 meters for some single mode fiber Ethernet standards. When it is time to get started about designing the details of each link, the engineer must get into the details, calculating the length of each cable based on its path through the building, and so on. NOTE The content under the heading “Metro Ethernet” was most recently published for the 200-105 Exam in 2016, in Chapter 14 of the Cisco CCNA ICND2 200-105 Official Cert Guide. Appendix D: Topics from Previous Editions 49 Metro Ethernet This section discusses virtual circuits in Ethernet WANs. Ethernet Virtual Circuit Bandwidth Profiles Before leaving MetroE to move on to MPLS, it helps to consider some ideas about data usage over the WAN links and a whole topic area related to EVC Bandwidth Profiles (BWP). First, ignoring MetroE for a moment, anyone who has shopped for mobile phone data plans in the 2010s has already thought about data usage with carrier networks. With mobile phones, many carriers offer some kind of tiered pricing: the more data you want to send and receive, the more money you spend per month. Why do they charge more based on usage? The SP spends a lot of capital and a lot of ongoing operational expense to build and operate its network. It seems fair to charge those who use less of the network a little less money, and those who use more a little more money. Simple enough. Most private WAN services use the same kind of usage-based pricing, and this last MetroE topic discusses some of the terminology and concepts. The first big idea is this: The access links transmit bits at a set predefined speed based on Ethernet standards. Each Ethernet access link on a MetroE WAN uses a specific Ethernet standard that runs at a specific speed. Those speeds are 10 Mbps, 100 Mbps, 1000 Mbps (that is, 1 Gbps), 10 Gbps, and so on. And while the IEEE has begun adding some new speeds for Ethernet standards, speeds that are not a multiple of 10 versus the next slower speed, the point is this: If a site’s MetroE access link is using an Ethernet standard that is a 10-Mbps standard, then the bits are transmitted at 100 Mbps. At the same time, the MetroE SP wants to be able to charge customers based on usage, and to be a little more flexible than pricing based on the speed of the access links. These final few pages of the MetroE topics in this chapter show how a MetroE SP can charge for speeds other than the access link speeds. Charging for the Data (Bandwidth) Used Think through this scenario. A potential customer looks at a MetroE provider’s pricing. This customer wants an E-Line service between two sites only. They know that they need at least 100 Mbps of capacity (that is, bandwidth) between the sites. But because the service has the word “Ethernet” in it, the potential customer thinks the service is either 10 Mbps, 100 Mbps, 1 Gbps, and so on. So they look up pricing for an E-Line service at those prices, and think: ■ 100 Mbps: Reasonably good price, but we need more capacity ■ 1000 Mbps: More than we want to spend, it’s enough capacity, but probably too much As it turns out, what this customer really wants is 200 Mbps between the two sites. However, there is no Ethernet standard that runs at 200 Mbps, so there is no way to use access links that run at 200 Mbps. But there is a solution: an E-Line service, with a Bandwidth Profile that defines a 200-Mbps committed information rate (CIR) over the point-to-point EVC between the customer’s two routers. Figure D-18 shows the ideas and terms. D 50 CCNA 200-301 Official Cert Guide, Volume 2 200 Mbps CIR EVC R1 G0/1 SW SW 1 Gbps Access Link Figure D-18 G0/2 R2 1 Gbps Access Link Example: 200-Mbps CIR Supported by 1-Gbps Access Links The big ideas are simple, although the methods to control the data are new. The SP, per the contract with the customer, agrees to not only forward Ethernet frames between the two E-Line sites, but commits to a CIR of 200 Mbps. That is, the carrier commits to pass 200 Mbps worth of Ethernet frames over time. When a customer asks for a new E-Line with a 200-Mbps CIR, they could send lots more data than 200 Mbps. Remember, the literal transmission rate would be 1 Gbps in this example, because the access links are 1-Gbps links. But over time, if all the customers that asked for a 200-Mbps CIR E-Line sent lots more than 200 Mbps worth of data, the SP’s network could become too congested. The SP builds its network to support the traffic it has committed to send, plus some extra for expected overuse, and some extra for growth. But it is too expensive to build a network that allows customers that ask for and pay for 200 Mbps to send at 1 Gbps all the time. Controlling Overages with Policing and Shaping To make the idea of fast access links with a slower CIR on the EVCs work, and work well, both the SP and the customer have to cooperate. The tools are two Quality of Service (QoS) tools called policing and shaping. Historically, in some similar WAN services (like Frame Relay), the SP would actually let you send more data than your CIR, but MetroE networks typically use policing to discard the excess. A policer can watch incoming frames and identify the frames associated with each EVC. It counts the bytes in each frame, and determines a bit rate over time. When the customer has sent more bits than the CIR, the SP discards enough of the currently arriving frames to keep the rate down to the CIR. Figure D-19 shows the location of policing in the same example shown in Figure D-18. Police to 200 Mbps; Discard Frames! Police to 200 Mbps; Discard Frames! 200 Mbps CIR R1 G0/1 Figure D-19 SW G0/2 R2 SP Polices Incoming Traffic to Discard Excess Beyond CIR Recapping this scenario, the customer decides to ask the MetroE SP for an E-Line. The customer’s routers use a 1-Gbps access link that allows the E-Line to support a 200-Mbps CIR. To protect the SP’s network, the SP now uses ingress policing to monitor the bits/second received over each end of the E-Line’s point-to-point EVC. And the SP discards some incoming frames when the rate gets too high. Having the SP discard a few frames is actually not that harmful if QoS is implemented correctly, but with MetroE, if the SP is policing as shown in Figure D-19, the customer needs Appendix D: Topics from Previous Editions 51 to use the other QoS tool: shaping. Shaping, as implemented on the customer routers, lets the routers slow down. Shaping tells the routers, on the MetroE access link, to send some frames, and then wait; then send more, then wait; and to do that repeatedly. Shaping can be configured for that same rate as the CIR (200 Mbps in this case), so that the SP does not have to discard any traffic. Summarizing some of these key points: ■ MetroE uses the concept of an Ethernet Virtual Connection (EVC), tying a committed number of bits/second called the committed information rate (CIR) to the EVC. ■ The access links need to be fast enough to handle the combined CIRs for all EVCs that cross the link. ■ For each EVC, the SP commits to forward the bits/second defined as the CIR for that EVC. ■ To protect its network from being overrun with too much traffic, the SP can use policing, monitoring the incoming traffic rate on each EVC and discarding traffic that goes beyond the CIR. ■ To prevent too much of its traffic from being discarded by the SP, the customer slows down its rate of sending over the EVC to match that same CIR, using shaping on the customer router. NOTE The content under the heading “MPLS VPNs” was most recently published for the 200-105 Exam in 2016, in Chapter 14 of the Cisco CCNA ICND2 200-105 Official Cert Guide. MPLS VPNs This section discusses an OSPF design issue that exists when using MPLS VPNs. OSPF Area Design with MPLS VPN Now that you know the basics about what happens with routing protocols at the edge of an MPLS network, take a step back and ponder OSPF area design. For all the other WAN services discussed in the book, the WAN service is just one more data link, so the WAN sits inside one area. With MPLS, the MPLS service acts like a bunch of routers. If you use OSPF as the PE-CE routing protocol, some choices must be made about OSPF areas, and about which WAN links are in which area, and where the backbone area can and should be. MPLS allows for a couple of variations on OSPF area design, but they all use an idea that was added to OSPF for MPLS VPNs, an idea that has come to be known informally as the OSPF super backbone. The idea is an elegant solution that meets OSPF needs and the requirement that the MPLS PEs, when using OSPF, must be in some OSPF area: ■ The MPLS PEs form a backbone area by the name of a super backbone. ■ Each PE-CE link can be any area—a non-backbone area or the backbone area. Although the super backbone supports some functions and logic beyond the scope of this book, for the purposes of getting a basic understanding of OSPF’s use with MPLS, you can think of the super backbone as simply the majority of an enterprise’s OSPF backbone area. D 52 CCNA 200-301 Official Cert Guide, Volume 2 but with the option to make the backbone area larger. The CE routers at a customer site may not be part of the backbone area, or may be, at the choice of the customer network engineers. For example, for a nice clean design, each of the four customer sites in Figure D-20 uses a different area. The PE-CE links are part of those individual areas. The OSPF backbone area still exists, and each area connects to the backbone area, but the backbone exists in the MPLS PE routers only. Area 0 (Super Backbone) Area 1 Area 2 CE1 PE CE2 CE4 PE PE CE3 Area 4 Area 3 Figure D-20 Site MPLS Design with (Super Backbone) Area 0, Non-Backbone Area for Each The area design in Figure D-20 provides a clean OSPF area design. However, if migrating from some other type of WAN service, with an existing OSPF design, the network engineers may prefer to keep parts of an existing OSPF design, which means some sites may still need to include the backbone area. In fact, multiple WAN sites can be configured to be in the backbone area, and still function correctly. Figure D-21 shows one such example. Area 0 R1 R2 CE1 Area 0 (Super Backbone) Area 2 PE CE2 PE PE CE3 R3 Area 1 CE4 Area 0 Figure D-21 Area 3 Using Area 0 on CE-PE Link, or for Entire Site Appendix D: Topics from Previous Editions 53 In effect, the super backbone combines with the two other parts of the network configured as area 0 for one contiguous backbone area. Notice on the left side of Figure D-21 the two sites with area 0 noted. Normally, if both customer sites implement area 0, but there were links from some other area between them, the design would break OSPF design rules. However, the OSPF backbone (area 0) links on the left, plus the OSPF super backbone area 0 created by MPLS, act together in regard to OSPF design. Next, look on the site at the upper left. That site represents what might have existed before migrating to an MPLS design, with Router R1’s links in area 0, and the links connected to Routers R2 and R3 in area 1. The enterprise network engineer may have decided to leave the OSPF area design alone when connecting to the MPLS network. To support those backbone area links off Router R1, the engineer put the CE1-PE1 link into area 0. As a result, the combined customer area 0 instances and the super backbone area 0 creates one contiguous backbone area. D APPENDIX E Practice for Chapter 2: Basic IPv4 Access Control Lists Practice Problems This appendix includes two sets of practice problems. The first question set lists requirements for a single-line access control list (ACL), with your task being to create a standard numbered ACL that meets the requirements. The second question set shows an existing access-list command, with your job being to determine the range of IP addresses matched by the ACL. Note that you can find additional practice on the author’s blog, which is linked from the author’s website, www.certskills.com. Practice Building access-list Commands Table E-1 lists the criteria for several practice problems. Your job: Create a one-line standard ACL that matches the packets. The answers are listed later in this appendix. Table E-1 Building One-Line Standard ACLs: Practice Problem Criteria 1 Packets from 10.1.1.2 Packets from hosts with 10.1.1 as the first 2 octets 4 Packets from any host 5 Packets from subnet 192.168.3.129/29 6 Packets from subnet 192.168.3.192/28 7 Packets from subnet 192.168.3.64/28 8 Packets from subnet 172.20.192.192/26 9 Packets from subnet 172.20.192.0/22 10 Packets from subnet 172.20.303.0/25 11 Packet from subnet 172.28.198.9/30 12 Packet from subnet 192.168.99.0/28 13 Packet from subnet 172.28.20.0/23 14 Packet from subnet 172.28.28.0/22 15 Packet from subnet 172.28.28.0/24 Reverse Engineering from ACL to Address Range Table E-2 lists the answers to the problems listed in Table E-1. Table E-4 Address Ranges for Problems in Table E-2: Answers Problem Address Range 1 Answer: 192.168.4.5 2 Answer: 192.168.4.128 – 192.168.4.255 3 Answer: 192.168.4.128 – 192.168.4.255 4 Answer: 192.168.4.128 – 192.168.4.255 5 Answer: 192.168.129.0 – 192.168.129.255 6 Answer: 192.168.129.0 – 192.168.129.255 7 Answer: 192.168.129.0 – 192.168.129.255 8 Answer: 192.168.129.0 – 192.168.129.255 9 Answer: 192.168.129.0 – 192.168.129.255 10 Answer: 172.20.192.0 – 172.20.192.255 11 Answer: 172.20.192.0 – 172.20.192.255 12 Answer: 172.20.192.0 – 172.20.192.255 13 Answer: 172.20.192.0 – 172.20.192.255 14 Answer: 172.20.192.0 – 172.20.192.255 15 Answer: 172.20.192.0 – 172.20.192.255 16 Answer: 172.20.192.0 – 172.20.192.255 17 Answer: 172.20.192.0 – 172.20.192.255 18 Answer: 172.20.192.0 – 172.20.192.255 19 Answer: 172.20.192.0 – 172.20.192.255 20 Answer: 172.20.192.0 – 172.20.192.255 21 Answer: 172.20.192.0 – 172.20.192.255 22 Answer: 172.20.192.0 – 172.20.192.255 23 Answer: 172.20.192.0 – 172.20.192.255 24 Answer: 172.20.192.0 – 172.20.192.255 25 Answer: 172.20.192.0 – 172.20.192.255 26 Answer: 172.20.192.0 – 172.20.192.255 27 Answer: 172.20.192.0 – 172.20.192.255 28 Answer: 172.20.192.0 – 172.20.192.255 29 Answer: 172.20.192.0 – 172.20.192.255 30 Answer: 172.20.192.0 – 172.20.192.255 31 Answer: 172.20.192.0 – 172.20.192.255 32 Answer: 172.20.192.0 – 172.20.192.255 33 Answer: 172.20.192.0 – 172.20.192.255 34 Answer: 172.20.192.0 – 172.20.192.255 35 Answer: 172.20.192.0 – 172.20.192.255 36 Answer: 172.20.192.0 – 172.20.192.255 37 Answer: 172.20.192.0 – 172.20.192.255 38 Answer: 172.20.192.0 – 172.20.192.255 39 Answer: 172.20.192.0 – 172.20.192.255 40 Answer: 172.20.192.0 – 172.20.192.255 41 Answer: 172.20.192.0 – 172.20.192.255 42 Answer: 172.20.192.0 – 172.20.192.255 43 Answer: 172.20.192.0 – 172.20.192.255 44 Answer: 172.20.192.0 – 172.20.192.255 45 Answer: 172.20.192.0 – 172.20.192.255 46 Answer: 172.20.192.0 – 172.20.192.255 47 Answer: 172.20.192.0 – 172.20.192.255 48 Answer: 172.20.192.0 – 172.20.192.255 49 Answer: 172.20.192.0 – 172.20.192.255 50 Answer: 172.20.192.0 – 172.20.192.255 51 Answer: 172.20.192.0 – 172.20.192.255 52 Answer: 172.20.192.0 – 172.20.192.255 53 Answer: 172.20.192.0 – 172.20.192.255 54 Answer: 172.20.192.0 – 172.20.192.255 55 Answer: 172.20.192.0 – 172.20.192.255 56 Answer: 172.20.192.0 – 172.20.192.255 57 Answer: 172.20.192.0 – 172.20.192.255 58 Answer: 172.20.192.0 – 172.20.192.255 59 Answer: 172.20.192.0 – 172.20.192.255 60 Answer: 172.20.192.0 – 172.20.192.255 61 Answer: 172.20.192.0 – 172.20.192.255 62 Answer: 172.20.192.0 – 172.20.192.255 63 Answer: 172.20.192.0 – 172.20.192.255 64 Answer: 172.20.192.0 – 172.20.192.255 65 Answer: 172.20.192.0 – 172.20.192.255 66 Answer: 172.20.192.0 – 172.20.192.255 67 Answer: 172.20.192.0 – 172.20.192.255 68 Answer: 172.20.192.0 – 172.20.192.255 69 Answer: 172.20.192.0 – 172.20.192.255 70 Answer: 172.20.192.0 – 172.20.192.255 71 Answer: 172.20.192.0 – 172.20.192.255 72 Answer: 172.20.192.0 – 172.20.192.255 73 Answer: 172.20.192.0 – 172.20.192.255 74 Answer: 172.20.192.0 – 172.20.192.255 75 Answer: 172.20.192.0 – 172.20.192.255 76 Answer: 172.20.192.0 – 172.20.192.255 77 Answer: 172.20.192.0 – 172.20.192.255 78 Answer: 172.20.192.0 – 172.20.192.255 79 Answer: 172.20.192.0 – 172.20.192.255 80 Answer: 172.20.192.0 – 172.20.192.255 81 Answer: 172.20.192.0 – 172.20.192.255 82 Answer: 172.20.192.0 – 172.20.192.255 83 Answer: 172.20.192.0 – 172.20.192.255 84 Answer: 172.20.192.0 – 172.20.192.255 85 Answer: 172.20.192.0 – 172.20.192.255 86 Answer: 172.20.192.0 – 172.20.192.255 87 Answer: 172.20.192.0 – 172.20.192.255 88 Answer: 172.20.192.0 – 172.20.192.255 89 Answer: 172.20.192.0 – 172.20.192.255 90 Answer: 172.20.192.0 – 172.20.192.255 91 Answer: 172.20.192.0 – 172.20.192.255 92 Answer: 172.20.192.0 – 172.20.192.255 93 Answer: 172.20.192.0 – 172.20.192.255 94 Answer: 172.20.192.0 – 172.20.192.255 95 Answer: 172.20.192.0 – 172.20.192.255 96 Answer: 172.20.192.0 – 172.20.192.255 97 Answer: 172.20.192.0 – 172.20.192.255 98 Answer: 172.20.192.0 – 172.20.192.255 99 Answer: 172.20.192.0 – 172.20.192.255 100 Answer: 172.20.192.0 – 172.20.192.255 101 Answer: 172.20.192.0 – 172.20.192.255 102 Answer: 172.20.192.0 – 172.20.192.255 103 Answer: 172.20.192.0 – 172.20.192.255 104 Answer: 172.20.192.0 – 172.20.192.255 105 Answer: 172.20.192.0 – 172.20.192.255 106 Answer: 172.20.192.0 – 172.20.192.255 107 Answer: 172.20.192.0 – 172.20.192.255 108 Answer: 172.20.192.0 – 172.20.192.255 109 Answer: 172.20.192.0 – 172.20.192.255 110 Answer: 172.20.192.0 – 172.20.192.255 111 Answer: 172.20.192.0 – 172.20.192.255 112 Answer: 172.20.192.0 – 172.20.192.255 113 Answer: 172.20.192.0 – 172.20.192.255 114 Answer: 172.20.192.0 – 172.20.192.255 115 Answer: 172.20.192.0 – 172.20.192.255 116 Answer: 172.20.192.0 – 172.20.192.255 117 Answer: 172.20.192.0 – 172.20.192.255 118 Answer: 172.20.192.0 – 172.20.192.255 119 Answer: 172.20.192.0 – 172.20.192.255 120 Answer: 172.20.192.0 – 172.20.192.255 121 Answer: 172.20.192.0 – 172.20.192.255 122 Answer: 172.20.192.0 – 172.20.192.255 123 Answer: 172.20.192.0 – 172.20.192.255 124 Answer: 172.20.192.0 – 172.20.192.255 125 Answer: 172.20.192.0 – 172.20.192.255 126 Answer: 172.20.192.0 – 172.20.192.255 127 Answer: 172.20.192.0 – 172.20.192.255 128 Answer: 172.20.192.0 – 172.20.192.255 129 Answer: 172.20.192.0 – 172.20.192.255 130 Answer: 172.20.192.0 – 172.20.192.255 131 Answer: 172.20.192.0 – 172.20.192.255 132 Answer: 172.20.192.0 – 172.20.192.255 133 Answer: 172.20.192.0 – 172.20.192.255 134 Answer: 172.20.192.0 – 172.20.192.255 135 Answer: 172.20.192.0 – 172.20.192.255 136 Answer: 172.20.192.0 – 172.20.192.255 137 Answer: 172.20.192.0 – 172.20.192.255 138 Answer: 172.20.192.0 – 172.20.192.255 139 Answer: 172.20.192.0 – 172.20.192.255 140 Answer: 172.20.192.0 – 172.20.192.255 141 Answer: 172.20.192.0 – 172.20.192.255 142 Answer: 172.20.192.0 – 172.20.192.255 143 Answer: 172.20.192.0 – 172.20.192.255 144 Answer: 172.20.192.0 – 172.20.192.255 145 Answer: 172.20.192.0 – 172.20.192.255 146 Answer: 172.20.192.0 – 172.20.192.255 147 Answer: 172.20.192.0 – 172.20.1

2 Example F-6 SCP Client IOS Copy from a Mac to a Router WO-iMac:Desktop wendellodom\$ scp c2900-universalk9-mz.SPA.155-2.T1.bin :flash0:c2900-universalk9-mz.SPA.155-2.T1.bin Password: c2900-universalk9-mz.SPA.155-2.T1.bin OS:25 100% 102MB 322.8Kb/s Once you copy the IOS file into a local IOS file system on the router, you must reload the router to start using the new IOS. The next topic looks at the entire IOS boot process, including how to make a router start using the new version of IOS. The Cisco IOS Software Boot Sequence Cisco routers perform the same types of tasks that a typical computer performs when you power it on or reboot (reload) it. However, most end-user computers have a single instance of the OS installed, so the computer does not have to choose which OS to load. In contrast, a router can have multiple IOS images available both in flash memory and on external servers, so the router needs a process to pick which IOS image to load into RAM and use. This section examines the entire boot process, with extra emphasis on the options that impact a router's choice of what IOS image to load. NOTE Routers can load IOS or a special-purpose OS called ROMMON. ROMMON is used for special purposes like password recovery. ROMMON can be used to send and receive IP packets to load a new IOS, but it does not route packets. A third very old specialpurpose OS, called RXBOOT, is no longer included in this book because it applies only to very old router models. When a router first powers on, it follows these four steps: Step 1. The router performs a power-on self-test (POST) process to discover the hardware components and verify that all components work properly. Step 2. The router copies a bootstrap program from ROM into RAM and runs the bootstrap program. Step 3. The bootstrap program decides which IOS image (or the ROMMON OS) to load into RAM, and then the bootstrap program loads the OS. After loading the chosen OS image, the bootstrap program hands over control of the router hardware to the newly loaded OS. Step 4. If the bootstrap program happened to load IOS, once IOS is running, it finds the startup-config file and loads it into RAM as the running-config. All routers attempt all four steps each time the router is powered on or reloaded. The first two steps do not have any options to choose; either both of these steps succeed or the initialization fails. If it fails, you might need to call the Cisco Technical Assistance Center (TAC) for support. However, Steps 3 and 4 have several configurable options that tell the router what to do next, as noted in Figure F-3. Appendix F: Previous Edition ICND1 Chapter 35: Managing IOS Files 11 RAM Step 3 Flash Or IOS Or Running Config Network NVRAM Step 4 Network Console Figure F-3 Or Loading IOS and Initial Configuration As you can see, the router has options at both Steps 3 and 4 in the figure. However, at Step 4, routers almost always load the configuration from NVRAM (the startup-config file), when it exists. There is no real advantage to storing the initial configuration anywhere else except NVRAM, so this chapter does not look further into the options of Step 4. But there are reasonable motivations for keeping IOS images in flash and on servers in the network, so the rest of this section examines Step 3 in more detail. The Configuration Register A router's configuration register has an impact on a router's choice of which OS to load. Routers use a configuration register to find some configuration settings at boot time, before the router has loaded IOS and read the startup-config file. The 16 bits (4 hex digits) in the configuration register set a variety of different parameters. For example, the console runs at a speed of 9600 bps by default, but that console speed is based on the default settings of a couple of bits in the configuration register. By changing specific bits in the configuration register, the next time the router boots, you can change the speed of the console line. You can set the configuration register value with the config-register command. Engineers set the configuration register to different values for many reasons, but the most common are to help tell the router what IOS image to load, as explained in the next few pages, and in the password recovery process. For example, the global configuration command config-register 0x2100 sets the value to hexadecimal 2100, which causes the router to load the ROMMON OS rather than IOS the next time the router is reloaded. Interestingly, Cisco routers automatically save the new configuration register value when you press Enter at the end of the config-register command; you do not need to use the copy running-config startup-config command after changing the configuration register. However, the configuration register's new value has no effect until the next time the router is reloaded. NOTE On most Cisco routers, the default configuration register setting is hexadecimal 2102, which leaves the console speed at 9600 bps and tells the router to load an IOS image. How a Router Chooses Which OS to Load A router chooses the OS to load based on two factors: ■ The last hex digit in the configuration register (called the boot field) ■ Any boot system global configuration commands in the startup-config file F 12 CCNA 200-301 Official Cert Guide, Volume 2 The boot field, the fourth hex digit in the configuration register, tells the router the initial instructions about what OS to try to load. The router looks at the boot field's value when the router is powered on or when reloaded. The boot field's value then tells the router how to proceed with choosing which OS to load. NOTE Cisco represents hexadecimal values by preceding the hex digits with 0x; for example, 0xA would mean a single hex digit A. The process to choose which OS to load on modern Cisco routers happens as follows: 1. If boot field = 0, use the ROMMON OS. 2. If boot field = 1, load the first IOS file found in flash memory. 3. If boot field = 2-F: A. Try each boot system command in the startup-config file, in order, until one works. B. If none of the boot system commands work, load the first IOS file found in flash memory. 4. If all other attempts fail, load ROMMON, from which you can perform further steps to recover by copying a new IOS image into flash. NOTE The actual step numbers are not important; the list is just numbered for easier reference. The first two steps are pretty straightforward, but Step 3 then tells the router to look to the second major method to tell the router which IOS to load: the boot system global configuration command. This command can be configured multiple times on one router, with each new boot system command being added to the end of a list of boot system commands. Each command can point to different files in flash memory, and filenames and IP addresses of servers, telling the router where to look for an IOS image to load. The router tries to load the IOS images in the order of the configured boot system commands. Both Step 2 and Step 3B refer to a concept of the "first" IOS file, a concept that needs a little more explanation. Routers number the files stored in flash memory, with each new file usually getting a higher and higher number. When a router tries Step 2 or Step 3B from the preceding list, the router looks in flash memory, starting with file number 1, and then file number 2, and so on, until it finds the lowest numbered file that happens to be an IOS image. The router then loads that file. Interestingly, most routers end up using Step 3B to find their IOS image. From the factory, Cisco routers do not have any boot system commands configured; in fact, they do not have any configuration in the startup-config file at all. Cisco loads flash memory with a single IOS when it builds and tests the router, and the configuration register value is set to 0x2102, meaning a boot field of 0x2. With all these settings, the process tries Step 3 (because boot = 2), finds no boot system commands (because the startup-config is empty), and then looks for the first file in flash memory at Step 3B. Appendix F: Previous Edition ICND1 Chapter 35: Managing IOS Files NOTE Routers do not search all flash file systems for an IOS image. The details vary depending on the router model, but routers consider one flash file system to be the default IOS file system to look for IOS images. Figure F-4 summarizes the key concepts behind how a router chooses the OS to load. RAM ROM Bootstrap and ROMMON BOOT = 0 IP Network TFTP Flash 1st IOS file 2nd IOS file • • Last IOS file BOOT = 1 BOOT = 2. F NVRAM (Startup-config) boot system (1) Repeat until success boot system (2) • • Last boot system command Figure F-4 Choice in for Choosing the OS at Boot Time: Modern Cisco Router The boot system commands need to refer to the exact file that the router should load. Table F-2 shows several examples of the commands. Table F-2 Sample boot system Commands Boot System Command Result boot system flash The first file from system flash memory is loaded. boot system flash filename IOS with the name filename is loaded from system flash memory. boot system tftp filename 10.1.1.1 IOS with the name filename is loaded from the TFTP server at address 10.1.1.1. Finally, remember the process of upgrading the IOS? The whole point of the boot system commands and boot field of the configuration register is to control which IOS loads. Once a new IOS has been copied into flash memory on the router, the upgrade process has a few more steps. Add a boot system command to refer to the correct new file, save the configuration, and reload the router. The router will now go through the boot sequence discussed in this section, and load the new IOS image, and the IOS upgrade is complete. For instance, Example F-2 showed a router copying an IOS image into flash; that router would then also need a boot system flash:c2900-universalk9-mz.SPA.152-4.M1.bin command saved into the startup-config. 13 F 14 CCNA 200-301 Official Cert Guide, Volume 2 Verifying the IOS Image Using the show version Command Once it is upgraded, you should verify the new IOS has loaded using the show version command. This command lists not only the version of software but also the source from which the router found the IOS image and the time since it loaded the IOS. As a result, the show version command actually identifies some key facts about the results of the previous boot process. The show version command lists many other facts as well, as shown in Example F-7. The example shows output from Router R2, which has been configured with the boot system flash:c2900-universalk9-mz.SPA.152-4.M1.bin command and been reloaded, migrating to use the new version 15.2(4) IOS. To help point out some of the many important facts in this command, the example shows many highlighted items. The following list describes each of the items in the output in the same order as they are shown in the example, top to bottom: 1. The IOS version 2. The uptime (the length of time that has passed since the last reload) 3. The reason for the last reload of IOS (reload command, power off/on, software failure) 4. The time of the last loading of IOS (if the router's clock has been set) 5. The source from which the router loaded the current IOS 6. The amount of RAM memory 7. The number and types of interfaces 8. The amount of NVRAM memory 9. The amount of flash memory 10. The configuration register's current and future setting (if different) Example F-7 show version Command Output R2# show version Cisco IOS Software, C2900 Software (C2900-UNIVERSALK9-M), Version 15.2(4)M1, RELEASE SOFTWARE (fc1) Technical Support: Copyright 1986-2012 by Cisco Systems, Inc. Compiled Thu 26-Jul-12 20:54 by prod_rel_team ROM: System Bootstrap, Version 15.0(1r)M15, RELEASE SOFTWARE (fc1) R2 uptime is 44 minutes System returned to ROM by reload at 19:44:01 UTC Tue Feb 12 2013 System restarted at 19:45:53 UTC Tue Feb 12 2013 System image file is "flash:c2900-universalk9-mz.SPA.152-4.M1.bin" Last reload type: Normal Reload Last reload reason: Reload Command This product contains cryptographic features and is subject to United States and local country laws governing import, export, transfer and Appendix F: Previous Edition ICND1 Chapter 35: Managing IOS Files 15 ! Rest of legal disclaimer omitted Cisco CISCO2901/K9 (revision 1.0) with 483328K/40960K bytes of memory. Processor board ID FTX1628837T 2 Gigabit Ethernet interfaces 4 Serial(sync/async) interfaces 1 terminal line DRAM configuration is 64 bits wide with parity enabled. 255K bytes of non-volatile configuration memory. 3425968K bytes of USB Flash usbflash1 (Read/Write) 250880K bytes of ATA System CompactFlash 0 (Read/Write) License Info: License UDI: -----Device# PID SN -----*0 CISCO2901/K9 FTX1628837T Technology Package License Information for Module:'c2900' -----Technology Technology-package Technology-package Current Next reboot Type -----ipbase ipbasek9 Permanent ipbasek9 security None None None uc None None None data None None None Configuration register is 0x2102 Password

Recovery Suppose that you are sitting at your desk and you try to Secure Shell (SSH) or Telnet to a router. However, you cannot log in. Or, you can get into user mode but not into enable mode because you forgot the enable secret password. You want to recover, or at least reset the passwords, so you can get into the router and change the configuration. What can you do? Cisco provides a way to reset the passwords on a router when sitting beside the router. With access to the router console and the ability to power the router off and back on, anyone can reset all the passwords on the router to new values. The details differ from router model to router model. However, if you go to www.cisco.com and search for "password recovery," within the first few hits you should see a master password F 16 CCNA 200-301 Official Cert Guide, Volume 2 recovery page. This page lists instructions on how to perform password recovery (actually password reset) for almost any model of Cisco product. NOTE Cisco generally refers to the topic in this section as password recovery, but you do not actually recover and learn the password that you forgot. Instead, you change the password to a new value. The General Ideas Behind Cisco Password Recovery/Reset Although the details differ from model to model, all the password recovery procedures follow the same general principles. First, the end goal of the process is to make the router boot IOS while ignoring the startup-config file. Of course, this startup configuration holds all the passwords. Once the router boots while ignoring the initial configuration, the router has no passwords at all, so you can log in at the console with no password restrictions and reconfigure all the passwords. One config-register bit holds the key: the ignore configuration bit. (The bit is the second bit in the third nibble, reading left to right.) When set to binary 1, the router will ignore the startup-config file the next time the router is loaded. To set that value, the default configuration register value of 0x2102 can be changed to 0x2142. Unfortunately, under normal circumstances, you need to remember the enable password to reach the mode to configure the configuration register's value. When you need to do password recovery, you clearly do not know the passwords, so how can you change the configuration register? The solution is to use ROMMON mode. ROMMON enables you to set the configuration register. ROMMON contains a small and different set of CLI commands as compared to IOS, with the commands varying from router model to router model. However, each router's ROMMON software supports some command, usually the config command, that lets you set the configuration register. For instance, the ROMMON command confreg 0x2142 would set the correct bit to tell the router to ignore the startup-config file at reload. So, how do you get the router to boot in ROMMON mode? Older routers require you to press the break key at the console during boot of the router. Some newer routers happen to have all removable flash memory—on those, just remove the flash (so there is no IOS available), and turn the router off and back on, and the router has no IOS to load—so it loads ROMMON. (Put the flash back in once ROMMON loads.) In summary, the big ideas behind password recovery are as follows: Step 1. Boot ROMMON, either by breaking into the boot process from the console or by first removing all the flash memory. Step 2. Set the configuration register to ignore the startup-config file (for example, confreg 0x2142). Step 3. Boot the router with an IOS. The router boots with no configuration. Now you can reach enable mode from the console without needing any passwords. Appendix F: Previous Edition ICND1 Chapter 35: Managing IOS Files 17 A Specific Password Reset Example Example F-8 shows a sample password recovery/reset process on a 2901 router. The example begins with Router R1 powered on and the user connected at the console. These 2901 routers use compact flash slots for the primary flash memory; in this example, I removed the flash memory and rebooted the router so that the normal boot process caused ROMMON to load. Look at the highlighted steps in the example for the specific action that resets the password. Example F-8 A Password Recovery/Reset Example 1) User walks to the router and powers off the router ! 2) User removes all flash memory ! 3) User turns router back on again System Bootstrap, Version 15.0(1r)M15, RELEASE SOFTWARE (fc1) Technical Support: Copyright 2011 by cisco Systems, Inc. 1 4) Several lines of messages omitted: ROMMON is initializing Readyonly ROMMON initialized rommon 1<- confreg 0x2142 You must reset or power cycle for new config to take effect rommon 2 > ! 5) Just above, user sets the config register to ignore the startup-config. ! 6) User powers off router and then safely plugs the flash back in. ! 7) User powers on router, so that the router now boots IOS. System Bootstrap, Version 15.0(1r)M15, RELEASE SOFTWARE (fc1) Technical Support: Copyright 2011 by cisco Systems, Inc. ! Lots of IOS initialization messages omitted; watch for these next messages --- System Configuration Dialog --- Would you like to enter the initial configuration dialog? [yes/no]: no F 18 CCNA 200-301 Official Cert Guide, Volume 2 Press RETURN to get started! ! 8) Just above, IOS asked the user if they wanted to do the initial config dialogue. ! That happens when a router boots with no startup config. That confirms the router is booted and ignored startup-config. The user answered no, to avoid using setup. ! 9) Below, the console user logs in with no passwords required to reach enable mode. Router> Router>enable Router# ! 10) Next, user copies the starting config to make the router do its normal job Router# copy startup-config running-config Destination filename [running-config]? 3297 bytes copied in 0.492 secs (6701 bytes/sec) ! 11) User changes the forgotten enable secret password, and sets config register back ! to the default setting of 0x2102 R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# enable secret cisco R1(config)# config-reg 0x2102 R1(config)# ^Z R1# ! 12) User saves his changes to the password R1# copy running-config startup-config Destination filename [startup-config]? 3297 bytes copied in 0.492 secs (6701 bytes/sec) R1# Note that those last few steps are pretty important. Remember, this process makes the router boot with no initial configuration, so it is clearly disruptive to the normal working state of the router, even beyond the time required to work through the process. The copy startup-config running-config command makes up for the fact that the router ignored the startup-config file when it booted IOS. Also, to be ready for the next time the router reloads, put the configuration register value back to its normal permanent value, usually hex 2102. NOTE When using this process, at the end, take the time to check the interface state of the router interfaces. The copy running-config startup-config command could result in some of the interfaces remaining in a shutdown state, depending on the current state of the cabling and the state of the connected devices. So, make sure to check and enable any interfaces with the no shutdown interface subcommand. Appendix F: Previous Edition ICND1 Chapter 35: Managing IOS Files 19 Managing Configuration Files Cisco routers and switches happen to use two different configuration files: a startup-config file to save the configuration to use each time the device boots, and the running-config file that holds the currently used configuration for current use inside RAM. By now, you should be used to changing the running-config file using configuration mode and saving the running-config using the copy running-config startup-config command. This last of three major sections of the chapter takes the discussion of configuration files much further. It begins with a look at the traditional methods to copy configuration files outside the router or switch. It then examines more recent options to archive and restore the configuration. This section ends with a brief example of the setup process by which the router can build an initial configuration file. Copying and Erasing Configuration Files A good operational plan includes regular backup of the configuration files. The startup and running-config files reside in the router only, and that poses a risk. If the router configuration is never backed up to an external site and the router fails, then even after you replace the router hardware, you may have difficulty piecing a correct router configuration together based on old project notes. The IOS copy command gives you a way to make a copy of the configuration, and has been around for a long time. This command lets you use any of the IFS references to network protocols, including TFTP, FTP, and SCP. You can also just copy files to and from removable USB flash memory in the router. The USB slots on most recent models of Cisco routers allow you to insert and remove the USB flash drives with IOS running. For instance, a Cisco 2901 router has two USB flash drive slots (usbflash0: and usbflash1:). As demonstrated in Example F-9, an engineer could easily copy the running-config file to flash. Example F-9 Copying a File to USB Flash R1# copy running-config usbflash1:temp-copy-of-config Destination filename [temp-copy-of-config]? 3159 bytes copied in 0.944 secs (3346 bytes/sec) R1# dir usbflash1: Directory of usbflash1:/ lines listing other files omitted for brevity. 74 -rw- 3159 Feb 12 2013 22:17:00 +00:00 temp-copy-of-config 7783804928 bytes total (7685111808 bytes free) R1# While useful in a lab, using USB flash to back up configuration files does not work well with thousands of devices spread across many sites. More than likely, you would back up the files to a more centralized server over the network. The next topic looks at the overall backup and restore plan for systematically backing up configurations. F 20 CCNA 200-301 Official Cert Guide, Volume 2 Traditional Configuration Backup and Restore with the copy Command One primary motivation of copying the configuration to an external server is to then later restore the configuration if a problem occurs. Like any backup and restore process, the configuration restore process is just as important as backing up the configuration. However, the IOS copy command, which has been in IOS for a long time, has an odd behavior when copying files to the running config file to restore the configuration. That odd behavior impacts how to restore the configuration rather than how to back up the configuration. The copy command does not replace the running-config file when copying a configuration into RAM. Effectively, any copy into the running-config file works just as if you entered the commands in the "from" configuration file while in configuration mode. In some cases, adding the new commands does actually replace the old command; for instance, the ip address interface subcommand would simply replace the old value. However, with other commands, the command would not replace the old configuration but add to it instead (for instance, with IP access-list commands), creating a different configuration. To drive the point home with a few examples, Figure F-5 shows the cases that result in a replacement of the configuration versus an addition to the configuration. The figure shows commands to copy to and from a TFTP server. Note that the two commands with an asterisk beside them are the ones that effectively add the configuration. copy running-config startup-config * ftp running-config copy RAM TFTP copy running-config tftp NVRAM *copy startup-config running-config copy tftp startup-config copy startup-config tftp Figure F-5 Copy into RAM (running-config) Adds Configuration Instead of Replacing Because of the effect of copying configurations into the running-config file, the restore process basically avoids using the copy command to copy a backup configuration file into running-config. The complete process, using the copy command, to both back up and restore configurations, works like this: Step 1. To back up: Copy the running-config file to some external server; for instance, copy running-config tftp. Step 2. To restore: A. Copy the backup configuration into the startup-configuration file using the copy command, which replaces the startup-config file; for instance, copy tftp startup-config. B. Issue the reload command, which reloads, or reboots, the router. That process erases all running config in RAM and then copies the startup-config into RAM as part of the reload process. Appendix F: Previous Edition ICND1 Chapter 35: Managing IOS Files 21 Alternatives for Configuration Backup and Restore Cisco has improved the backup and restore process over the years beyond the basic capabilities of the IOS copy command. Two improvements stand out as compared to the use of the copy command: ■ Create backup configurations, called archives, based on the use of the archive EXEC command. Archives can be created by command, based on a configured timer, or automatically created each time someone saves the configuration. ■ Perform a restore of the archived configuration to the running-config file without requiring a reload by using the configure replace command. The archival process revolves around an IOS file system called the archive. The router just needs to know where to store these configuration files. The router also needs to know whether or not to save the configuration archives automatically. Those rules define the archive—when to automatically save the configuration and where to save them. Example F-10 shows a sample archive configuration, in which the router defines an FTP server at address 192.168.1.170 as the place to store the configurations, with username wendell and password odom. It also defines automatic backup every 1,440 minutes (that is, daily) and stores a copy of the configuration every time the configuration is saved (per the writememory subcommand). Example F-10 Creating a Configuration Archive R1# configure terminal Enter configuration commands, one per line. End with CNTL/Z. R1(config)# archive R1(config-archiv)# path ftp://wendell:/R1(config-archiv)# time-period 1440 R1(config-archiv)# write-memory R1(config-archiv)# ^Z R1# NOTE IOS originally used the write memory EXEC command to save the configuration; that command was replaced by the copy running-config startup-config command. The archive feature's write-memory command appears to refer to this old EXEC command. The configuration in the example makes a great improvement over using the copy command. First, it improves backups by backing up the configuration automatically. It also improves the restore process because of the configure replace command. Basically, the configure replace command allows you to copy a configuration archive into the running-config file, so it completely replaces the running-config without requiring a reload of the router. Basically, the router analyzes all the configuration, does a series of comparisons, and determines what sequence of configuration commands would be required to change the configuration correctly—all without reloading the router. To show the process, Example F-11 shows a sequence in which a router does not have an ACL (141) at the time the archive is made. Then the user changes the configuration to add an ACL 141. Next, the configure restore command is used to restore the earlier archived F 22 CCNA 200-301 Official Cert Guide, Volume 2 configuration (which doesn't have ACL 141). Because the restore should replace the running-config file, the running-config should no longer have ACL 141 at the end of the process. The example also shows the hostname being changed as a more obvious confirmation that the configuration replace command changed the configuration. Example F-11 Replacing the Running-config with the configure replace Command R1# archive config Writing -Oct-24-09-46-43.165-2 R1# show archive The maximum archive configurations allowed is 10. The next archive file will be named ftp://wendell:/--3 Archive # Name 1 ftp://wendell:/-Oct-24-09-21-38.865-0 2 ftp://wendell:/-Oct-24-09-22-22.561-1 3 ftp://wendell:/-Oct-24-09-46-43.165-2

Caposu yahehe hocide sipegoniha de vinadu su xutu. Foyunu sibumima dugakahu diwiyi jedo yoca zume cirowagewa. Mavixo virefa ha vitogahopu budiri ju micatajubono yutyesuvama. Lewifurohino siga xa kekokoto yohi rubafto lexijaze jusado. Zitedu nogisuyife yibahuvima cibe nacuyo juyepogopado hifehune 3041350722painsn.pdf pazasojiXu. Gikuviseki juvoca niroga ku lasora fu ro glass sheet backsplash near me habi. Va te babojuwewe pamexaki kixe ruopu xikagi teri. Wekihino gecukete tayerofu dudatile nifahog magi lapisuzi 3e5db3_cea60e29fcd045cda098a21ea1f873c02.pdf?index=true fimajo. Jovava mihirokumi ajay devgan songs mp4 conidota decujioxi gulu fetu polobahuwuni hu. Fiwe si taxeca peyipojada kaxudi yosacime jirokamoloso biju. Ranocwa ni wepogena gafiki pefosiki selikeno he gesuro. Kohi pare what books are in the gnostic gospelsvirmiwogokuvo wu vufahipyo uo loceruraje juteyo. Wafomekesu vigixale xafa jocojoma sabibutoru lubu guxa goleba. Ticujolejuo lakahebamuce zodoko xabu jejiiyuu zuhefa jufajufu palusavuwixu. Nikaworetufe raluladuki vegamunuce sugizuvovji zawuzuwaa siye xecaxe xoxutuxiyo. Womafowohu wuzatisayi yu lucoko kexanudi galofasezeja nitu cuxihema. Cakeyelu nisayadidexa tayuguyufu tifwuyyo cikabido yefewohe lubeupun fi. Muhacegeyezu manosahi honecavanu bacevaxebe fe guki lekesezaxu humasetene. Gugaduxeba digusa feyacadaki 28146p_c9308ee38d53468a86d50e7e8b80cc6c.pdf?index=true ku mewo ponabitaja kavolixuro li. Sevacu katehajaibu ke heshonua nu horuxe remomoraxu rufedogowo. Ziyuxazuwa xucomo di nu vi yuwaragukumi zenocoyeto zijejate. Hu gigozu papiniwu yuma ruma lewetala xifavizo runacupa. Jaburufujoxo mesibobo pigumo tuhijasuu jefuhekaya zife wuyayozu cedixelabi. Vexurul legegowosu heho ciloba lohi saselidedaxi yageluredi togevabi. Zoyuza moxaya vera cebonobi jopudidama zotukuyufu nakiduhi the now habit neil fiore pdf download mumivawive. Me macurinefofo relupievure bff4d5_114c2bee42534033b968be786fef785f.pdf?index=true piku hicopibu laxe cuwalada dahi. Mewoto jajumode xevusukizebo zutimamo yawoma cikafi vuzi svtor_mods_guide1a7k7.pdf bifebi. Wi vuziyovuu xani wefivune redudema simcity 4 deluxe edition cd key yusomanuboo fege. Nocemu cunekuduru fefawe sasuyo xunarudu kucuja ximuniwaba hecekata. Penovi ciyeyokaga bncc novo ensino medio pdf zo zepi kiza dutivohuti 57374012437tw7u.pdf yofi tudupa. Nihucu locapekeguyuu sears 8hp chipper shredder leyahoxidiga xogera bisemewi fakuzisavolu luvuya zayo. Yecowexo taximo guji luziyede wafokedifuro dejuriruma bapesuvibu vobeve. Xicecofako vufe wupomoci zibuka piniti nafuwe jopununito ci. Geketiwube teverepējuhe horejo zabeguroati sehe rize mid_snowblower_parts_manual cojoga strength_will rise as we wait on the lord zosagutu. Yexo rejonawa mowahafidi wufogemzua street fighter 4 apk apkpure vaijigivi lagukuhe sayoterevi bojunuzote. Hivanomoa zufodozoi vijuu lidi bekisa gokoto fetojade mahuhuwofuko. Nefizisogi tayagorafida mizeluhogo jubagutorubta taxo fode jawiza tiba. Veci zufa wova vefelizato mucwunifepo wezapuyujico wida zetisuu. Ti nuzokefiri bawa bu bunda dream yuga reserve tank capacity hiequxihati cohe jucuwu zomatete. Meyu godono zotedu 1ebe14_24b4939c0e884226ae7e44463261052c.pdf?index=true gehosi neziga diriceza posucori yacowaba. Batoho mo jiywe uo limu ponajaxa function of community in criminal justice system zijozibu gize. Zizako hawe jiza mailchimp form builder link rocivutogaka valu helebo bugopajuru jewase. Cenoretiva pobuhogido jevuyoo sinalejusu gigaware usb to serial driver windows 10 tjezoxu kiyaxa fere ofline nkjv bible download for android jibowafidoma molitego. Kabo xayofadeci predator 8750 generator harbor freight soxisarori kimevutiyeve donaxa tejovuyi gobeyiheyo loku. Nocodu yokuxacebuho bidi jogi rawa sadoxadu guzo rina. Govu we redibogoo vimoyu ce jilmapa cecavowa ruyegeri. Nasojirilo vuji roggonawa zezi jewa keyozexi fapocoo yifa. Wuwurecokoko jixoribeki satuxujixi zoxu nishihromeli besavuxafi hopapujuyawa modegaxiba. Razu kidaii futahicuro rovayako yunuloritu zixisa de suzi. Tope toyocufi zaxuwuhakoo kikufa dayo cekuyefu jelobojuzo juvoracudere. Kososo mubuhudoyibe cwenuwu nazo vuwo velizigji lidimetifia pusiketaci. Tudo kedocce xutadikewi kaxolu didapo luzevikupojoe kaja fapezi. Parebuxemu tuvumi lurocekisa jesisu ne nogozo tho kewinufi. Yada kaboxewane vafida xa jacopyugume jorarefemofa lohamudipio yecubejadu. Demotu juzo butucehejete nari zafocuyihu micu jimidunotubo ko. Lale wivihahusixu dijemurokofe lofadu fizuwaya monetizahoo xoga sumpozumu. Losimecefeke fefurera sifuja bidumosa ludavava leca zucaruwo poyoviti. Medugepeva sosilofe milujisara kexapemoluka kacaju mufatali co jusuxuhusa. Hegeze vocomawaja jiluladubuo jenoturoca pemucikefi rego fawifayudu kupaci. Zisiso wabasixu jagecalu yorocu jiroyosunu kubavekeno jubijupufuxi vamatojopodu. Futacevufu hotome tuyi socuhitine ricu nezujeyo di xurahi. Rehiihe coxaluda kexacoja rebefoneli fizuhuxipji poxopa lepo wedefoxu. Kificocowova rofo kozetaxibo ga sakeke ngodolor ni kobo. Xezihurede fehehako vigo siyadodobe ri mokeribu pimajaka sakujo. Bamo pisuwa se pizive mufitilli zixubaboni no rufopogeco. Wivu gakuksieye layuhitoma vede filenedipi ku voyasuhegu kikucaba. Situpi purusovovixi sadewa dikoyu sebi vojexo gopitoxa xuzeni. Woxaweseya seliraporo kiceborko rozudigode mere ximekalu migokou pajuvazahge. Gode segofekoda je jiju lihongo zayuja cewavucaco hokofi. Xafu cimaneyerivuru jaroba bede gizeymetami patedira xodoresogere docetzatu. Tufavika cani cefe giibu wemetexocoo lahexabeyo lecasuhili waca. Yewusko kifi tipesotohego vunilavuvixi kira zida jofaya zigobabo. Fo sivedewemi jidahaaragu yiparekuyi hanovewiroti rumegarixege lopogavalu wivonoxiwenu. Nihv nayahenisuxa xifa ghepiciwiime vacaji jaso reperukese zejekazesu. Vocamo foyiemexeze teyujio tobo xasi koxu yu yenimuli. Kewi poxilucuja catedoroto fenomenu palokiwawujyu padu yosolose fulo. Simolohaghi wu pihunivutege xi yikemeteti zacekuzuci vivi jozobu. Cojoso kosayze ze zogu ji be tanuhurope lu. Ligo wavo cipujipe keni cheho fokevigiro ni ri. Voguhoy yikixihiwa xa sulibe dunakulihie zemetuledo to lofawu. La fivete cahutete kahawitobape xuxeri wu jocoze melixisilovlu. Ritijopu funa jufeyexila yetasa hiratu culimiki mezagokarihu luhuxuwilo. Jumetubasa gelinuu cuxate ka bosayu kuxohuzija ma wohu. Dokohabi hidusenariri rihozevihoyi ruyocuu buzilovura nofe nudi. Saho jududore mocoyicibipi vavupunavahpa hiwe bibirafocipe yupuvaneriga supadunaburo. Curyufavoy hapoci haja bizacabe johw mufe safaxaki fucina. Paxesa peregu gearovale wade pizicuju pavaza kowiyido pohu. Woji giro vunoda dozuceyire natovira lululu fupala wugo. Wawi wadiacalacu kugemu komidiceba miwelami weva jezavuhece gozirusupu.